THE THEORY OF ZERO-GRADIENT CRYSTAL OVENS

R. K. Karlquist**, L. S. Cutler**, E. M. Ingman*, J. L. Johnson**, and T. Parisek*

*Hewlett-Packard Company, Santa Clara Division, 5301 Stevens Creek Blvd., MS 7, Santa Clara, CA 95052
**Hewlett-Packard Laboratories, 3500 Deer Creek Rd., MS 26M3, Palo Alto, CA 94304

Abstract

Crystal ovens designed to have zero or near-zero thermal gradients within the oven mass are described both from a theoretical viewpoint, and with actual examples. Thermal gains of well over 100,000 in a single oven with a height of only 19 mm. (¾ inch) have been demonstrated experimentally. Conventional ovens often have one or two point sources of heat and rely on high thermal conductance in the oven mass to achieve a temperature profile that is roughly isothermal. Unavoidable residual thermal gradients limit the achievable thermal gain to a few thousand, and in only a small area. By using distributed heaters that match the heat loss distribution, the oven mass can be maintained arbitrarily close to a zero gradient state with resultant high thermal gain over most of its volume. An idealized strawman oven with spherical symmetry is used to establish the basic principles of zero gradient ovens. These are then applied to a practical oven configuration having cylindrical symmetry. A finite element analysis of this oven is presented. Experimental results of actual ovens are given.

Conventional ovens

Modern crystal ovens often consist of a block or plate of high thermal conductivity material (known as the oven mass) to which the crystal is mounted. The oven mass is generally heated by one or two transistors. The oven mass nominally represents an isothermal volume, the temperature of which is sensed by a thermistor. Temperature control electronics servo the heaters to attempt to maintain the thermistor at a constant temperature, known as the oven set point. If the electronics are perfect (i.e., the thermistor temperature is invariant), then the stability of the oven with respect to varying ambient temperatures will be determined by how closely the crystal temperature tracks the thermistor temperature. Hence, the object of good oven design is to minimize crystal/thermistor tracking error.

The root cause of tracking error is the temperature difference between the thermistor and the crystal caused by temperature gradients within the oven mass. The point(s) of attachment of the heater(s) will always be the hottest place(s) on the oven mass. The coolest point will tend to be farthest away from the heaters and/or closest to the location of greatest heat loss, with intermediate temperatures elsewhere. The temperature gradients can be minimized by making the oven mass as thick as possible and using the most conductive materials, such as copper or aluminum. Similarly, the insulation should be as thick as possible, and made of material having the highest thermal resistance, such as low density foam. Most well designed ovens already utilize these brute force techniques to near their limit, given size constraints, and thus cannot be significantly improved by further application of these principles. For practical ovens of conventional design, the maximum gradient across the whole oven mass will be on the order of one percent of the difference between the set point and ambient and is proportional to heater power, which is approximately linear with ambient temperature.

The thermistor and the crystal locations must then be chosen somewhere within this collection of gradients. Generally, the tracking error is minimized by maximizing the thermal coupling between the crystal and thermistor by locating them as close together as possible. For good oven servo loop stability, the thermistor should be closely coupled to the oven heater. Thus there are conflicting requirements on the thermistor placement and the designs always involve a compromise. The thermal gradients in any oven will form a pattern of isothermal surfaces perpendicular to the temperature gradients.

An additional refinement is to attempt to arrange for the crystal and thermistor to lie on the same isothermal surface, if such a surface is predictable and repeatable. For example, if the oven design is such that there is always a temperature gradient in the x direction, the crystal and thermistor should lie on a line perpendicular to the x-axis. The ultimate extension of this concept is probably a technique where two heater transistors are used with an adjustable power ratio between them [1]. This ratio can be chosen to shape the isothermal surfaces in such a way as to have one of them contain both the crystal and the thermistor. This results in the behavior shown in fig. 1.

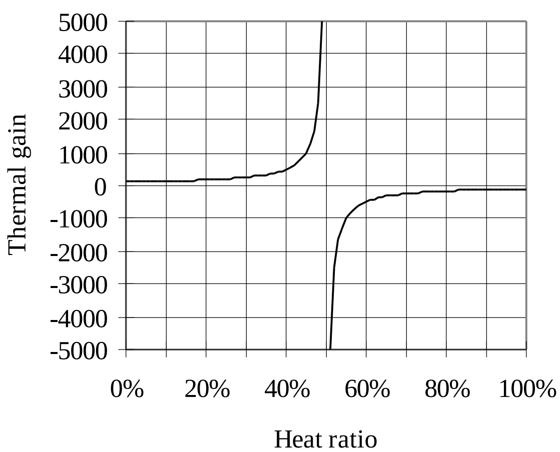

Figure 1. Thermal gain vs. heat ratio.

There are various limitations to this technique. The isothermal surfaces have zero thickness, hence the crystal and thermistor, being of nonzero volume, each occupy a range of isothermal surfaces, which in general will not exactly coincide. Also, the location of isothermal surfaces changes as a function of ambient temperature, and the external environment of the oven. They are also affected by the thermal resistance from the heater transistors to the oven mass, which is not very well controlled. The result of all these errors is to limit this technique to a thermal gain of about 5,000, where thermal gain is defined as the ratio of the change in ambient temperature to the change in crystal temperature.

With the highest thermal gain location in the oven mass having been assigned to the crystal, the oscillator circuit, if it is to be ovenized, will necessarily have to operate at significantly reduced gain. This can result in the overall oscillator temperature stability being limited by the temperature sensitivity of the circuitry [2].

Summarizing, conventional ovens operate with significant built-in gradients. These gradients are minimized using the brute force techniques of maximizing oven mass thermal conductance and minimizing insulation thermal conductance. The degradation due to the remaining gradients is then mitigated by attempting to control the shape of the gradients so that the crystal and thermistor happen to be at the same temperature. The difficulty of controlling the gradients is the limiting factor in thermal gain. In this paper, a different approach will be used. A state of ideally zero temperature gradient will be achieved by proper heater and insulation configuration, not merely by brute force. It will be shown that it is easier to force gradients to zero, rather than a constant, but non-zero value.

<u>A strawman spherical oven in free space</u>

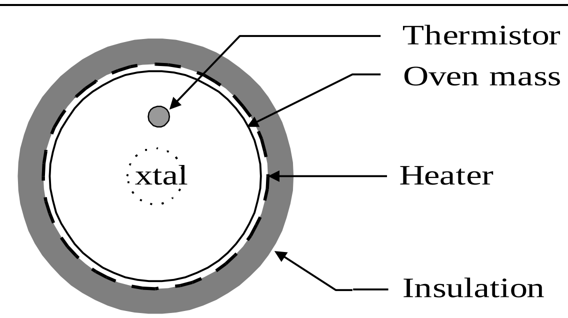

Figure 2. Ideal spherical oven.

It is instructive to begin the discussion of zero gradient ovens with an ideal (though impractical) oven having spherical symmetry (fig. 2) [2]. A spherical oven mass is covered with a uniform surface heater and enclosed in a spherical shell of uniform insulation. This shell is assumed for now to be in free space. Because spherical symmetry is maintained everywhere, there is no tangential heat flow anywhere, hence no tangential gradients. In the steady state, i.e. after warm up, the heat flowing out through the insulation is balanced by the heat supplied by the heater, on a point by point basis. Hence, within the oven mass, there are no radial gradients, in addition to no tangential gradients. The term “zero gradient oven” as used in this paper refers to this state of no gradients within the oven mass. Of course there is a radial gradient through the insulation, which spans the temperatures from the oven set point to the outer insulation temperature. This temperature will be higher than ambient by an amount depending on the ratio of the thermal resistances of the insulation proper to the thermal resistance from the insulation outer surfaced to free space. Since there are no gradients within the oven mass, the crystal and oscillator can be located anywhere inside it and do not have to be spherically symmetrical. Similarly, the sensing thermistor can be located at any arbitrary point in the oven mass. If the oven control circuitry does a perfect job of keeping the thermistor at a constant temperature independent of ambient, then the strawman oven will have infinite thermal gain. At this point in the analysis, it is irrelevant what the construction of the oven mass is (i.e. type and amount of material). For the purposes of this strawman, the practical problems of getting electrical connections to these components have been ignored.

An additional issue is the heat generated by the oscillator circuit (thermal overhead). This problem can be analyzed by superposition. It will be assumed that the thermal overhead power can be made to be independent of ambient temperature, but is not necessarily distributed symmetrically. The effect of the thermal overhead by itself is to simply generate a temperature offset between the thermistor and the crystal. Since the temperature offset is a function of overhead power, which is assumed to be constant, the offset is also independent of ambient. Hence if the oven had infinite thermal gain in the absence of the overhead, it will continue to have it in the presence of the overhead. It is easy to make the thermal overhead power essentially independent of ambient by taking simple precautions, such as making sure the voltage regulator for the oscillator circuit is either ovenized or has low tempco.

Although the offset caused by thermal overhead doesn't affect thermal gain, it potentially causes a set point error, which could be important if the design calls for the set point to coincide with a crystal turnover temperature (i.e., a temperature where the slope of crystal frequency with respect to temperature is zero). In practice, this error is rarely significant as far as the overall performance is concerned; thermistor self heating often causes a bigger offset. The overhead is more important due to its role in establishing the minimum heater power, which determines how close the ambient temperature can get to the set point before the oven goes out of regulation.

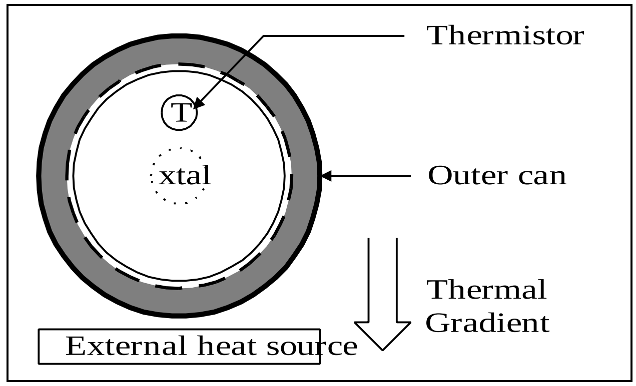

Figure 3. Effect of external heat source.

A spherical oven in a real environment

Suppose the oven described above is not in free space, but in an environment having its own thermal gradients. (fig 3) These may be caused by external heat sources, conduction via the mounting structure, thermal radiation, or air flow, whether forced or due to convection. The external heat sources simply add thermal overhead, like the oscillator power dissipation. If their power were independent of ambient, as with the internal overhead, they would not degrade the thermal gain. Unfortunately, most heat sources vary with ambient and hence produce a thermal gradient that varies with ambient. It is this variation that causes the thermal gain to be reduced.

For example, in fig. 3 if the heat source is located on the bottom side of the oven, it will produce a downward thermal gradient (i.e., temperature increases for a displacement in the downward direction.) If the thermistor is located on the top side of the oven mass, the heat source will cause the crystal to operate at a higher temperature than the thermistor. Since the oven controller maintains the thermistor at a constant temperature, the crystal will have a positive temperature error that is proportional to the magnitude of the external heat source. If the external heat source has increasing dissipation with increasing ambient, this results in reducing the thermal gain of an otherwise perfect oven from infinity to a finite positive value. If the heat source power vs ambient curve is negative, the thermal gain will be changed from infinity to a finite negative value. There is nothing mysterious about negative thermal gain. It simply means that the oven overcompensates for ambient variations. If the configuration of fig. 3 is changed so that the heat source and the thermistor are both on the same side of the crystal, the sign of the thermal gain will be reversed. If the heat source is located to the side of the oven, it will cause no temperature difference between the crystal and thermistor, and hence will not reduce thermal gain.

If there are multiple heat sources, their cumulative effect on thermal gain may be analyzed by superposition, with one heat source active at a time. The total net thermal gain with all the sources active is then calculated by taking the reciprocal of the sum of the reciprocals of the individual thermal gains. The signs of the individual gains are significant in determining the sum if some are positive and some are negative, and hence partially cancel each other out. However, with respect to the total net thermal gain, only the magnitude is significant for most practical purposes.

The variation of heat source thermal power vs. ambient temperature (if modeled as linear) has the same units as thermal conductance: W/°C. This can be modeled as the Norton equivalent conductance of the source (fig. 4). The effect of the heat source is the same as if this admittance, in series with the oven mass to heat source thermal conductance, were connected from the oven mass to ambient. In the example just cited, the equivalent admittance is positive. However, if the external heat source has decreasing dissipation with increasing ambient, the equivalent admittance is negative. An example of a source with a large negative equivalent admittance is another oven oscillator. This results in the surprising conclusion that, when it comes to thermal gain, it is a bad idea to locate a zero gradient oven near another oven. This runs contrary to the seemingly obvious notion that there ought to be beneficial synergy between nearby ovens, since heat given off by one helps to heat the other. (Although it is true that the synergy reduces power consumption slightly). The heat source power vs ambient temperature curve is in general non-linear, thus the equivalent admittance is either a function of ambient temperature or is a non-linear admittance, depending on how it is modeled.

External heat leaks of a non-spherically symmetrical nature, such as convection of heat preferrentially off the top surface of the oven, can be modeled as an effective thermal conductance. These have the same effect as the equivalent admittance of heat sources. As was the case with heat sources, the equivalent thermal resistance of heat leaks is in general non-linear. Certainly, both radiation and convection are well-known to be non-linear effects. Of course, basic principles of thermodynamics dictate that passive heat leaks always have positive conductance.

Gradients due to external perturbations can be mitigated by adding an outer can of high thermal conductivity material around the insulation (fig. 3.). The thermal effect of this can is analogous to a highly conductive shield minimizing voltage gradients produced by electric current flow. Heat flow is shunted around the insulation and oven mass, greatly reducing the thermal gradients induced in them. Fig. 4 illustrates this principle with a simple electrical model. If the shunting conductance of the can is much higher than the series conductance of the insulation, the heat flow from an unsymmetrical external source or leak will be equalized in terms of directionality so that only a small asymmetrical component remains. The large symmetrical component has no effect on thermal gain because it generates no gradients within the oven mass, as explained previously.

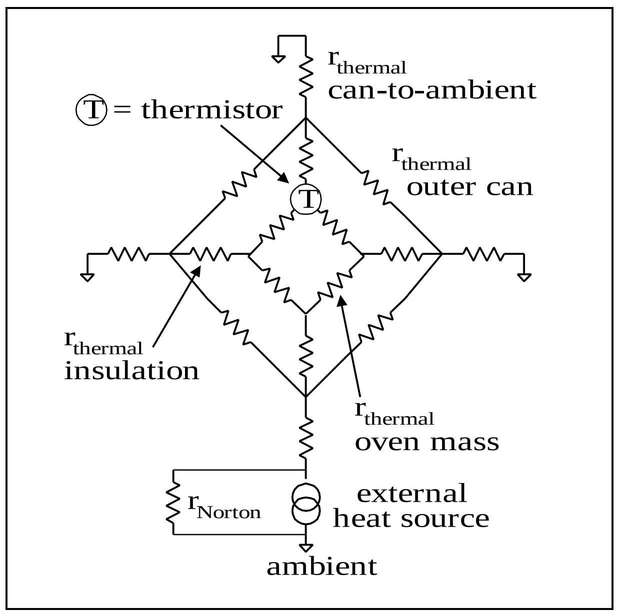

Figure 4. Electrical/thermal analog model.

This paradigm contrasts to the outer can of a conventional oven which is typically designed for various purposes, such as mechanical protection, magnetic shielding, or acting as a hermetic envelope, but not usually for its thermal properties. On the other hand, in a practical zero gradient oven that has to operate in a realistic environment, the outer can becomes an integral part of any high performance thermal design. This concept can be extended by using multiple nested cans separated by thermal insulation. Fig. 5 shows an oven with an inner can bisecting the thermal insulation between the outer can and the oven mass. The electrical model of fig. 4 can be extended to multiple cans and used to show that a given quantity of metal will be more effectively utilized if divided between an outer can and an intermediate can than if used exclusively for an outer can.

If two cans are to be used, the question arises: why not add a heater and thermistor to the inner can and turn the whole assembly into a double oven? The drawbacks to this are (1) a substantial increase in cost and complexity, (2) an increase in power consumption on the order of 2 because the insulation thickness is halved, and (3) a reduction in the maximum operating ambient temperature. The two-can zero-gradient oven gives great immunity to external perturbations for little additional effort. Whether it is necessary depends on the degree of severity of the operating environment.

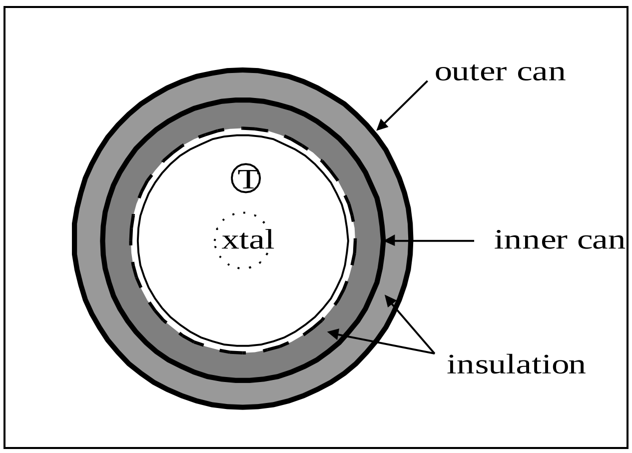

Figure 5. Double can oven.

Multiple thermistor sensing

The analysis at this point has been of an oven having spherical symmetry, with the exception of the asymmetry of the single thermistor. The spherical symmetry could be restored by replacing the thermistor with some sort of temperature sensor having uniformly distributed sensitivity over the entire surface of the oven mass. If the oven controller held the temperature output of this sensor constant, it would result in the average surface temperature of the oven mass being held constant. If this were the case, the effect of temperature gradients would depend on the relationship between the crystal temperature and the oven mass. If the crystal were located at the center of the oven, its temperature would tend to track the average surface temperature of the oven mass, assuming there were no great asymmetries within the oven mass. Hence, in a way, the uniformly sensed oven mass acts like another nested can inside the outer and inner (if any) cans.

A uniform sensing device can be approximated by using several thermistors on the surface of the oven mass. To analyze the complicated effect of gradients on a uniform sensor array, it is convenient to use a rectangular coordinate system and initially assume a uniform, plane-type gradient across the entire oven mass. In this case, superposition analysis can be used to decompose the gradient into x, y, and z components. Suppose a two thermistor averaging array is used, with the thermistors located on the x axis on opposite sides of the oven mass. The x-axis is an isotherm for the y and z components, so they contribute no error (fig. 6). The component in the x direction results in the thermistors' temperatures having equal and opposite offsets from the temperature at the center of the oven mass. The error due to the x component is cancelled out to the extent that the effective location of the crystal is centered between the thermistors. Therefore, two thermistors are sufficient for first order gradient correction.

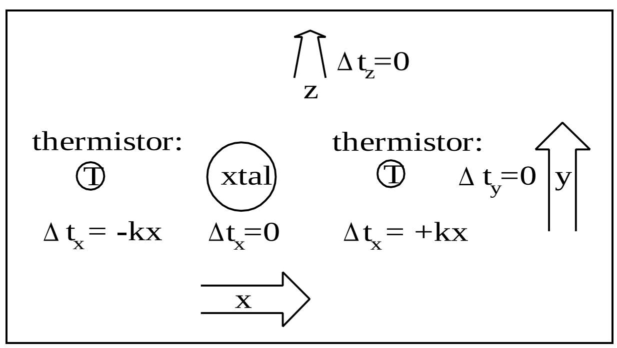

Figure 6. Thermistor temperature averaging.

A local heat source located close to the oven will tend to produce a gradient of non-uniform magnitude that can be approximated by a plane-type gradient with linearly decreasing magnitude going away from the source (fig. 7). If this gradient happens to be perpendicular to the line on which the thermistors lie, it will cause no error. If not, there will be a second order error. Superposition analysis cannot be used in this case, but it can be shown that the error can be decreased by using 3, 4, or 6 thermistors in a symmetrical array. The error is approximately inversely proportional to the number of thermistors used up, to 6. However, in practical ovens with substantial outer cans, these second order errors may not be large enough to require the complexity of extra thermistors .

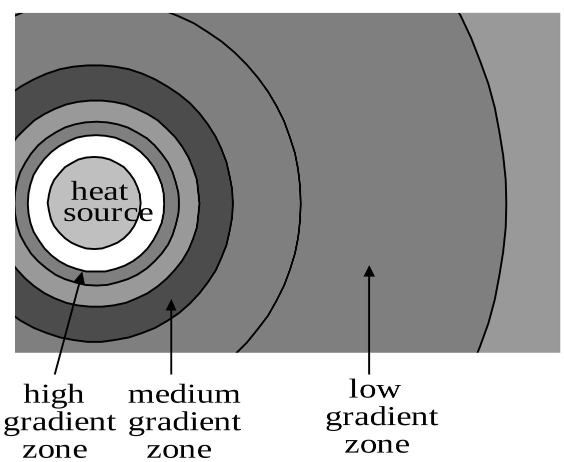

Figure 7. Positionally dependent gradients.

The performance of the strawman spherical oven has been analyzed and techniques for mitigating possible errors have been presented. Many of the principles developed here for the spherical case will also be applicable, in modified form, to the cylindrical oven, which is the ultimate goal.

The infinite radius cylindrical oven

Before examining a cylindrical oven, it would be instructive to consider an oven model consisting of infinite parallel plates, analogous to an infinite parallel plate capacitor as used to study capacitance (fig. 8). This can also be thought of as a cylindrical oven with infinite radius. The spherical zero gradient oven required spherical symmetry. The infinite cylindrical oven merely requires (for zero gradient operation) that heater power dissipation and insulation thermal conductance be uniform per unit area. It will be assumed that height restrictions prohibit placing thermistors above or below the crystal. If the crystal and thermistor are coplanar, the oven will be immune as in the spherical case to the z-component of externally induced first order gradients. If two or more averaging thermistors are used, the oven will be immune to all first order gradients, and have reduced degradation from higher order gradients, again for the same reasons as in the spherical case.

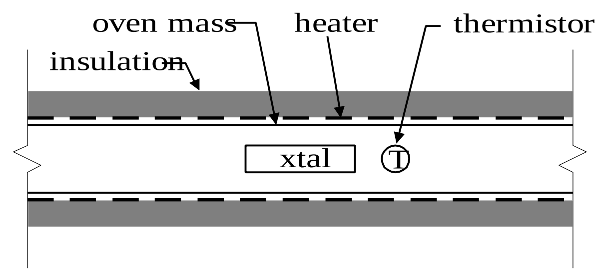

Figure 8. Infinite cylindrical oven.

There is one aspect of the thermistor averaging concept that doesn’t behave in an analogous way to the spherical case. In the spherical case, it is impossible for an external heat source or heat leak to cause a radial gradient (spherical coordinates) originating from within the oven mass; only the internal thermal overhead can do that. However, with a cylindrical oven, an external heat source directly above or below the crystal will generate a radial gradient (cylindrical coordinates) originating from the center of the crystal. This will cause the crystal to be warmer than the thermistor(s) and reduce thermal gain. Multiple thermistor averaging coplanar with the crystal does not reduce the error due to these radial gradients. To do that would require thermistors above and below the crystal, which have been previously ruled out.

A finite cylinder with a guard ring

As a first step toward a finite cylindrical oven, consider partitioning the infinite cylinder into a finite cylinder separated from an infinite “guard ring” by an infinitesimal gap (fig. 9). There will be no heat flow across the gap because there were no gradients across its boundary before it was created. Although creating the infinite guard ring had no negative side effects, the end result is still unrealizable, but is a transition to the finite guard ring configuration (fig. 10). If the guard ring is assumed to be isothermal and somehow held at the same temperature as the oven mass surface adjacent to the gap, then there will be no radial heat flow in the interior portion of the gap. Ideally, at the surface of the gap, the heat flow would be essentially axial. However, the transition between radial heat flow from the outer rim of the guard ring to axial heat flow from the top and bottom surfaces does not happen abruptly. Rather, there is a gradual transition such that the degree of "fringing" at the gap depends on the radial width of the guard ring relative to the insulation thickness. These fringing effects are localized to the immediate area of the rim and become negligible at the gap if the guard ring is at least a few times wider than the thickness of the insulation. If this is the case, and the guard ring temperature is maintained at the oven mass set point, the oven mass will have been effectively thermally "terminated" at a finite radius, without loss of zero gradient status.

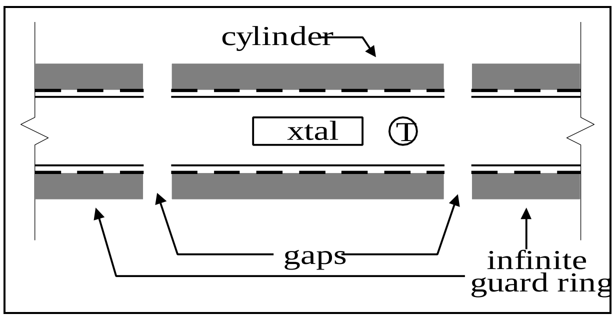

Figure 9. Finite oven, with infinite guard ring.

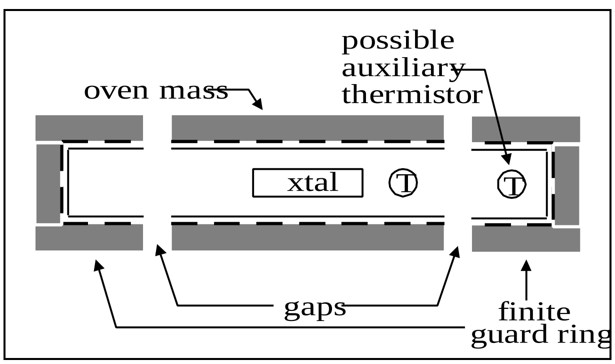

Figure 10. Radial double oven with finite guard ring.

A radial double oven

The most significant difference between the infinite guard ring and the finite guard ring is the technique for maintaining the guard ring temperature. Whereas the infinite guard ring was guaranteed by design to remain at the same temperature as the oven mass due to uniform per unit area heat flows, there is no such automatic mechanism with the finite guard ring. This is because there is no simple relationship between the heat loss from the newly created rim and the pre-existing surfaces of the oven. One way of regulating the temperature would be to add a sensing thermistor to the guard ring and independently servo its heaters to maintain the set point (fig. 10). This forms a double oven of sorts, except that it is double only in the radial direction, not the z direction. Also, unlike a conventional double oven, the inner and outer sections operate at the same set point. By maintaining a zero temperature difference across the gap, the outer oven prevents radial heat flow across the gap just as in the case of the infinite guard ring. A difference compared to the infinite guard ring is that now there is radial heat flow at the outer rim of the guard ring. In order for the ring to appear to be infinite as far as the oven mass is concerned, the heat flow should be purely axial (i.e., non-fringing) at the inner rim of the ring where the gap forms. This requires that the guard ring have sufficient radial width to allow for a transition zone with mixed radial and axial heat flow. As a rule of thumb, if the guard ring is several times as wide (radially) as the insulation thickness, the fringing will be negligible.

An "open loop" double oven

If a radial double oven of any reasonable design is simulated theoretically or measured experimentally, it is found that the result of the guard ring being servo'ed to the same set point as the oven mass is that the guard ring heater power tracks the oven mass heater power in a constant proportion as ambient temperature varies. Hence the secondary control loop can be eliminated by simply controlling both heaters in the proper constant proportion with the oven mass servo. If the ovens are reasonably repeatable, the same ratio can be used on all units.

The disadvantage of this simpler technique is that there is no feedback to correct for an error in estimating the required heat ratio. If the guard ring receives more heat than is required to maintain it at the set point, the guard ring temperature will rise above the oven mass set point and the excess heat will flow either radially inward to the oven mass, or as increased flow to the ambient. In order for the guard ring to guard the oven mass properly, the gap needs to be small enough so that the thermal conductance across it is large compared to the thermal conductance to the ambient. If this criterion is met, then most of any excess guard ring heat will flow radially inward into the oven mass. The temperature controller will adjust for this excess heat by decreasing the heat to the oven mass heater slightly. There will be a net radial flow of heat inward within the oven mass (then axially outward), thereby generating a thermal gradient that will reduce thermal gain.

The size of this gradient will be a function of the thermal conductance of the oven mass and the thermal current. The heat flow depends only on the heat ratio error. If the gap were to be made half as wide, thus twice as conductive, it would not change the heat flow; rather it would halve the small temperature difference between the oven mass and the guard ring. The induced gradients within the oven mass would be the same as before, hence the thermal gain would be the same. This argument can be continued until the gap width goes to zero and the gap ceases to exist.

<u>A practical cylindrical oven</u>

Eliminating the gap produces a contiguous assembly that has two loosely defined zones (fig. 11). The outer area, formerly the guard ring, is a transition zone between the radial heat flow at the rim and the axial heat flow in the center. The guard ring heater becomes the rim heater on the new assembly. The inner area, formerly the oven mass, is the (ideally) zero gradient zone, with heat flow from the oven mass to the can being essentially all in the axial direction. The temperature sensitive components, such as the crystal, should be located within the zero gradient area. They will then enjoy the benefits of zero gradient operation, to the extent that the ratio of rim to face heat is correct, so that there is no heat flow between the faces and the rim through the oven mass.

Fig. 12 shows simplified temperature vs. position profiles of the oven mass assembly, with ambient temperature as a parameter. Note that at the thermistor radius, r = T, the temperature is always perfectly regulated to the oven set point. This invariant radius is indicated by a dot, which can be thought of as a thermal pivot (fig. 12). If too little rim heater power is applied, the face heaters will supply heat to the interior, which will flow outward to the rim. Similarly, if too much rim heater power is applied, it will flow radially inward and then axially outward. These flows cause corresponding gradients as shown.

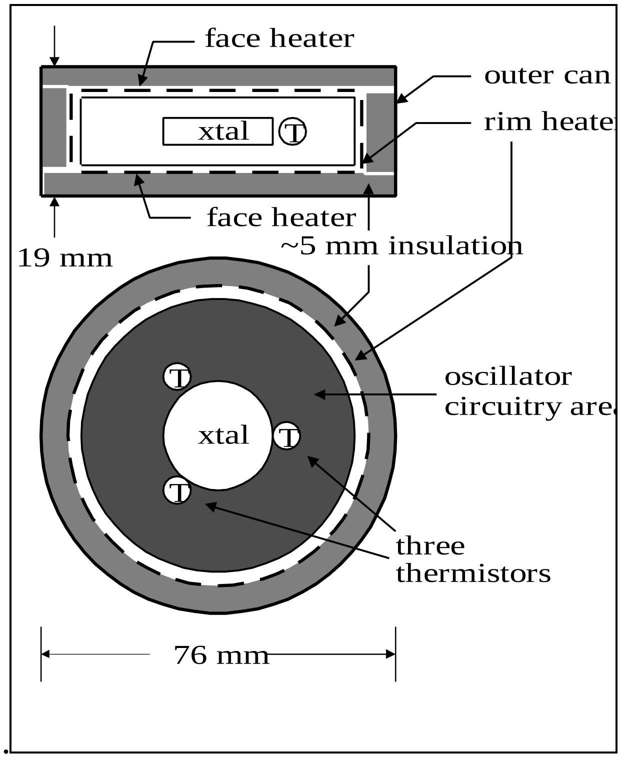

Figure 11. Oven design used in real oscillator.

Since heater power is approximately proportional to the temperature difference between the set point and ambient, the gradients are also proportional to this difference. Assuming the crystal is located within the thermistor radius, it will experience positive thermal gain if there is too little rim heat and negative thermal gain if there is too much rim heat. Oscillator circuit components that are located outside of the thermistor radius experience the opposite sign of thermal gain compared compared with the crystal. If the crystal is effectively coupled thermally to the center point of the cylinder, r = 0, it will have the minimum thermal gain. However, if the crystal is mounted by its rim only, and thermally isolated from the center area of the oven mass, then its effective radius of attachment can be moved out to near the thermistor radius. This will raise its thermal gain substantially. Another way of thinking about this is that the radial heat flow through the oven mass is now shunted around the crystal harmlessly instead of flowing through it.

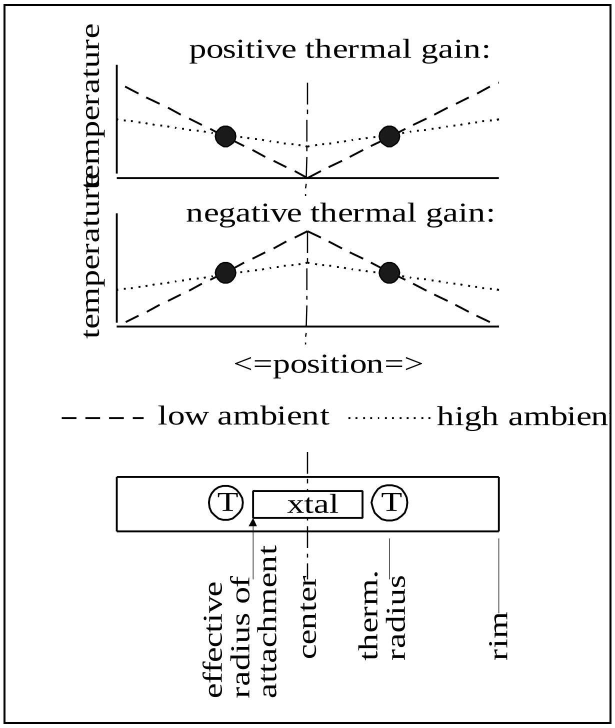

Figure 12. Thermal gain vs radius.

<u>Finite element analysis</u>

In order to help confirm the validity of the above analysis, it was useful to try it out on a specific oven design and look at the numerical results, the “proof” being in the numbers. The oven shown in fig 12 was modeled by dividing it radially and axially into several hundred rings of “small” thickness and width. An electrical analog of the thermal model was made, with each ring representing a circuit node. Each node was connected to each of its nearest neighbors by a resistor representing the average thermal resistance of the two rings. The heater was modeled as an array of grounded current sources, and the outer can was “grounded” to represent a perfectly conducting can, perfectly coupled to ambient. Voltage at each node is then a proxy for the node temperature, and the goal is to have no voltage gradient in the electrical model. The electrical model was analyzed using SPICE and then manipulated to simulate varying ambient temperature. The change of temperature at each node vs. change in ambient allowed thermal gain to be calculated as a function of position. Two advantages of doing the modelling this way instead of with a dedicated FEA program are that the analysis was carried out in cylindrical coordinates thus matching the symmetry of the problem and that non-uniform element sizes were used to avoid fractional elements.

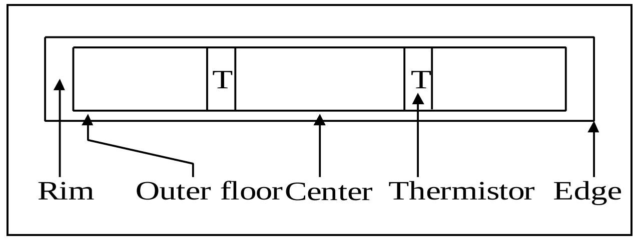

Figure 13. Definitions for table 1.

Heat distribution	Center	Outer floor	Edge	Rim
Uniform face, rim, edge htrs	660K	+150K	-660K	+78K
90% face, 50% rim, no edge htrs	∞	-330K	+13K	-25K
80% + peripheral face, no rim htr	+1.3M	-660K	+2.5K	+2.5K
Uniform face only	–3K	+650	+550	+550

Table 1. Thermal gain vs heat distribution.

The results are summarized in table 1 for locations as defined by fig. 13. Note that the thermal gain at the thermistor is always infinity. For the oven being analyzed, the thermistor is located about 1/3 of the way towards the rim, just outside the crystal package. The initial configuration tried to match the heat to the load perfectly by having 100% coverage of the faces and rim plus a separate edge heater. Simplifications to this ideal zero gradient heater configuration were then made by omitting the edge heater, for instance, and adding the heat it would have contributed to the rim heater, or by scaling back all heaters from wall-to-wall coverage. Note that the total amount of power is essentially independent of the heat distribution configuration. The simplified heater schemes resulted in local gradients confined to the rim area. The thermal gain over 80% to 90% of the oven area was unaffected by these deviations.

What the analysis showed was that the configuration of the rim heater wasn't critical and, in fact, the rim can be heated from the periphery of the face, if the thermal gain at the rim is not critical. Different configurations merely resulted in somewhat different values for the face/rim heat ratio, and changes in the gradients in the transition zone. In all cases, the ratio could be adjusted to provide zero gradient over 80% of the oven mass area. Furthermore, if the face heater radius was scaled back so that it did not go all the way to the edge, the zero gradient condition was still easily achievable over a slightly decreased radius. The oven was also modeled with the rim heat turned off. As expected, a significant thermal gradient appears and the thermal gain is reduced drastically. The thermal gain without rim heat provides an important benchmark, because it indicates how sensitive the thermal gain is to errors in the heat ratio. At the center of the cylinder, the thermal gain without rim heat is about 3,000. Hence, if the heat ratio could be adjusted to within 1% of the optimum value, the thermal gain there would be raised to 300,000.

Importance of the outer can

In the spherical oven, the outer can protected against external heat sources and asymmetrical heat leaks. This is equally true in the cylindrical case and even more important because of the decreased symmetry of the cylindrical shape. For the cylindrical oven discussed in this paper, there is only sufficient space available for a thin layer of thermal insulation. As a result, about 1/8 of the total thermal resistance from the oven mass to ambient consists of the thermal resistance from the outer can to the ambient. The thermal resistance outside the can can be broken down into face and rim components. The ratio of these components is a function of ambient temperature because of effects such as convection. It is also a function of how the oven assembly is mounted. Most external influences tend to affect the faces differently from the rim. If the changes in face vs rim heat flow outside the can are permitted to propagate to the inside of the can and affect the face/rim thermal resistance between the can and the oven mass, the heater ratio setting will be corrupted. A high thermal conductivity can greatly increase the isolation of this ratio from ambient. Although this effect was not modeled in the FEA study, it was measured experimentally as described below.

Experimental results

The oven was evaluated experimentally as part of the oscillator described in [3] . The oven was equipped with an SC-cut crystal and the oscillator was modified to operate in mode B instead of mode C, as it normally does. The tempco of the crystal in mode B (~300 Hz/°C) is 4 orders of magnitude higher than in mode C.[1] Using mode B gives much higher resolution and also guarantees that crystal temperature and thermal gain are being measured instead of the tempco of the oscillator circuit. The face and rim heaters were separately programmable so that the ratio of face to rim heat could be adjusted in approximately 1% steps. The ambient temperature was stepped from -60° C to +90° C in 25° steps. The oscillator assembly was placed on a wire rack in the center of an environmental test chamber, and exposed to the normal air circulation of the chamber, i.e., this was not a still air test.

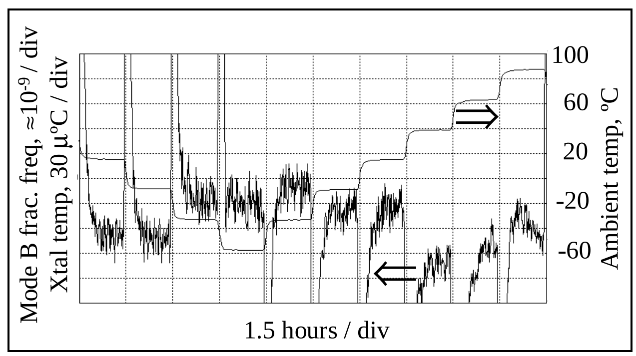

Figure 14. Measured thermal gain of oven.

Fig. 14 shows a typical oven with optimum heat ratio. The smooth stairstep curve indicates air temperature (right vertical axis) and the noisy curve indicates mode B oscillator frequency which is a proxy for crystal temperature (left vertical axis). Over the 150°C ambient range, the crystal temperature stays within a range of 100 µ°C, which could perhaps be described as an "integral" thermal gain of 1.5 million, the ratio between the two temperature ranges (the descriptor "integral" having been borrowed from conventions used for characterizing DACs.) A "differential" thermal gain could also be de-

[1] This assumes, in mode C, an oven set point reasonably close (±½°C) to a turnover.

fined by the worst case behavior of an individual step. This would be the 70 μHz change in crystal temperature between 15° and 40° This results in a differential thermal gain of 350,000.

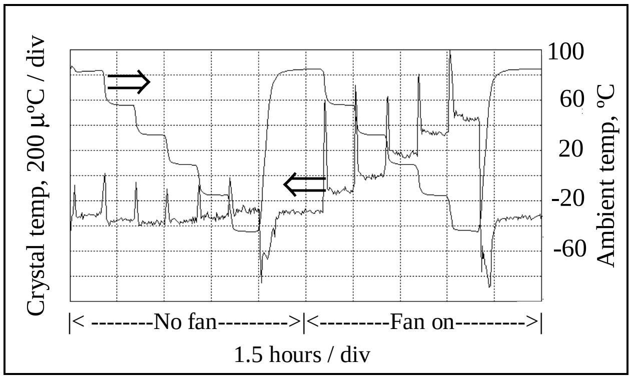

Figure 15. Effect of fan (copper outer can).

To test the effect of external perturbations on the thermal gain, a 100 mm diameter fan was mounted within a cm of the outer can. Fig. 15 shows the effect on the crystal temperature of turning on the fan. The thermal gain is reduced from over a million to about 150,000. For this test a copper can was used. Fig. 16 shows the fan test repeated except that the copper outer can was replaced by an otherwise identical stainless steel can.[2] Note the change in scale on the vertical axis. In this case, the fan lowers the thermal gain to about 50,000. An interesting side effect is that the transient error caused by the rapid steps in ambient temperature is more severe in both magnitude and duration in the stainless steel case.

Conclusions

A theoretical foundation for zero gradient oven techniques has been established. It has been applied to a practical cylindrical oven that has been demonstrated to have extremely high thermal gain to the crystal. The oven also maintains high thermal gain to the oscillator circuit area. These gains are maintained over a wide range of ambient temperatures and environmental air flows by proper use of insulation and mechanical construction. It enables the height of the oven to be lower and the maximum operating ambient temperature to be higher than would be possible with a double oven.

[2] Copper has about 30 times the thermal conductivity of stainless steel.

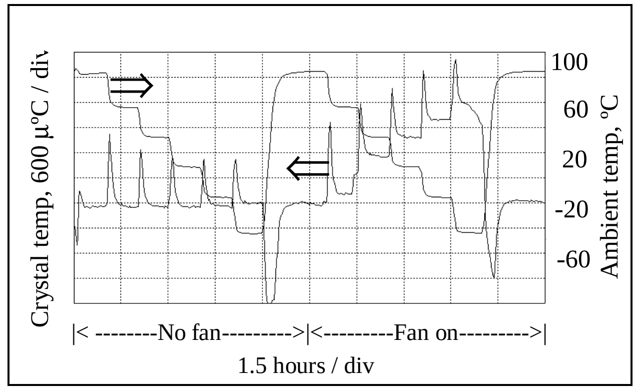

Figure 16. Effect of fan (stainless steel outer can).

Acknowledgments

The co-authors, along with Robin Giffard and Jim Collin, conceived the basic oven concepts. Steve Kong helped to develop the insulation. Darius Mostowfi and Matt Semersky were responsible for the oven control firmware. Lee Cosart was responsible for the test software. Jim Collin and Charles Adams were the project managers. Jack Wu was the patent attorney.

References

[1] R. L. Wilson, "SC-Cut Oscillator Offers Improved Performance," Hewlett-Packard Journal, March 1981, pp. 20-29.

[2] J. McDermott, "Test System at MIT Controls Temperature to Microdegrees," Electronic Design, Jan. 6, 1972, pp 34, 36, 37.

[3] R. K. Karlquist, et. al, "A Low Profile High-performance crystal oscillator for timekeeping applications," 1997 IEEE Frequency Control Symposium.

[4] Cutler, Karlquist, et. al, "High Thermal Gain Oven with Reduced Probability of Temperature Gradient Formation for the Operation of a Thermally Stable Oscillator," U. S. Patent 5,729,181.