SANDIA REPORT

A Fundamentally New Approach to Air-cooled Heat Exchangers

Jeffrey P. Koplow

Sandia National Laboratories

A Fundamentally New Approach To Air-Cooled Heat Exchangers

Jeffrey P. Koplow
Remote Sensing and Energetic Materials
Sandia National Laboratories
P.O. Box 969
Livermore, California 94551-9056

Abstract

We describe breakthrough results obtained in a feasibility study of a fundamentally new architecture for air-cooled heat exchangers. A longstanding but largely unrealized opportunity in energy efficiency concerns the performance of air-cooled heat exchangers used in air conditioners, heat pumps, and refrigeration equipment. In the case of residential air conditioners, for example, the typical performance of the air cooled heat exchangers used for condensers and evaporators is at best marginal from the standpoint the of achieving maximum the possible coefficient of performance (COP). If by some means it were possible to reduce the thermal resistance of these heat exchangers to a negligible level, a typical energy savings of order 30% could be immediately realized.

It has long been known that a several-fold increase in heat exchanger size, in conjunction with the use of much higher volumetric flow rates, provides a straight-forward path to this goal but is not practical from the standpoint of real world applications. The tension in the market place between the need for energy efficiency and logistical considerations such as equipment size, cost and operating noise has resulted in a compromise that is far from ideal. This is the reason that a typical residential air conditioner exhibits significant sensitivity to reductions in fan speed and/or fouling of the heat exchanger surface. The prevailing wisdom is that little can be done to improve this situation; the "fan-plus-finned-heat-sink" heat exchanger architecture used throughout the energy sector represents an extremely mature technology for which there is little opportunity for further optimization. But the fact remains that conventional fan-plus-finned-heat-sink technology simply doesn't work that well.

Their primary physical limitation to performance (i.e. low thermal resistance) is the boundary layer of motionless air that adheres to and envelops all surfaces of the heat exchanger. Within this boundary layer region, diffusive transport is the dominant mechanism for heat transfer. The resulting thermal bottleneck largely determines the thermal resistance of the heat exchanger. No one has yet devised a practical solution to the boundary layer problem. Another longstanding problem is inevitable fouling of the heat exchanger surface over time by particulate matter and other airborne contaminants. This problem is especially important in residential air conditioner systems where often little or no preventative maintenance is practiced. The heat sink fouling problem also remains unsolved. The third major problem (alluded to earlier) concerns inadequate airflow to heat exchanger resulting from restrictions on fan noise.

The air-cooled heat exchanger described here solves all of the above three problems simultaneously. The "Air Bearing Heat Exchanger" provides a several-fold reduction in boundary layer thickness, intrinsic immunity to heat sink fouling, and drastic reductions in noise. It is also very practical from the standpoint of cost, complexity, ruggedness, etc. Successful development of this technology is also expected to have far reaching impact in the IT sector from the standpointpoint of solving the "Thermal Brick Wall" problem (which currently limits CPU clocks speeds to ~ 3 GHz), and increasing concern about the the electrical power consumption of our nation's information technology infrastructure.

Table of Contents

1. Motivation

The objective of this work was to demonstrate proof-of-concept results for a new type of air-cooled heat exchanger (Figure 1) that simultaneously addresses the four longstanding problems of conventional "fan-plus-finned-heat-sink" (FFHS) devices [Koplow, 2008]:

1) high thermal resistance (i.e. low cooling capacity),
2) performance degradation due to heat exchanger fouling (e.g. from dust, pollen, etc.),
3) high electrical power consumption (related to fan aerodynamic efficiency), and
4) heat exchanger cooling capacity limitations imposed by fan noise.

Progress in forced-air heat exchanger technology is hampered by the fundamental physical limitations of the traditional FFHS device architecture (Figure 1). For example, it is well known that boundary layer effects impose fundamental limitations on cooling performance. In qualitative terms, a "boundary layer" may be considered a stationary layer of "dead air" that clings to the surface of a structure (e.g., a finned heat sink) and acts like an insulating blanket. In conventional FFHS devices, the difference in temperature between the base of the finned heat sink and ambient air is almost entirely accounted for by the temperature drop across the boundary layer. The exception to this rule is lap top computers, where available electrical power is extremely limited. In this special case, CPU clock speeds and fan rotor speeds are reduced to conserve power, albeit at the expense of CPU performance. At these low fan speeds the residence time of air in the heat exchanger is greatly extended, resulting in much higher exhaust air temperatures.

Figure 1: The conventional "finned-heat-sink-plus-fan" device architecture for air-cooled heat exchangers.
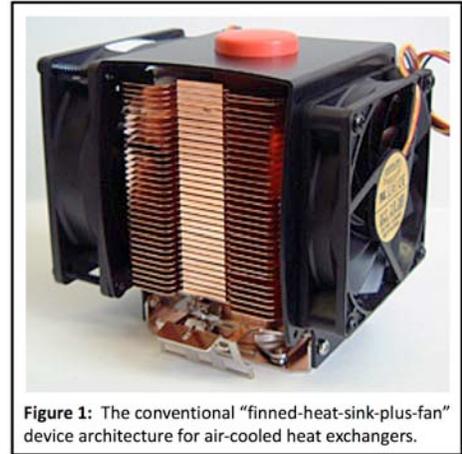
Within the boundary layer, molecular diffusion is the primary transport mechanism for conduction of heat, resulting in very poor heat transfer. Accordingly, the designs of ultra-high-performance air-cooling devices place a great deal of emphasis on boundary layer disruption. For example, air-jet-impingement cooling, in which a high-pressure pump generates a jet of compressed air that is directed at a heat sink surface, is very effective at reducing the thickness of the boundary layer. But the electrical power consumption and cost of air-jet-impingement cooling is prohibitive for the vast majority of applications. In devices such as the CPU cooler shown in Figure 1, although the fan generates a large amount of turbulence, only a modest reduction in the effective boundary layer thickness is observed relative to the case of laminar flow. This boundary-layer disruption effect can be increased to a small extent by running the fan at higher speed, but the tradeoff with respect to electrical power consumption quickly becomes very unfavorable.

Because problems such as the boundary layer effect stem from the fundamental physical limitations of the traditional FFHS device architecture, very little progress has been made to date with regard to performance improvements. For example, the current state of electronics thermal management technology was summarized by DARPA in a recent call for research proposals on new ideas for air-cooled heat exchanger technology: *"Over the past 40 years, CMOS, telecommunications, active sensing and imaging and other technologies have undergone tremendous technological innovation. Over this same historical period the technologies, designs and performance of air-cooled heat exchangers has remained unchanged. The performance data for today's state of the art heat exchangers and blowers is, in many cases, based on measurements performed in the 1960s."* DARPA has now

decided that considerable resources must be directed towards solving the air-cooling problem (www.darpa.mil/baa, DARPA Broad Agency Announcement 08-15).

As shown below, cooling of high-flux heat sources can be divided into two steps:

1) *heat extraction*, in which heat is removed from a high-thermal-density region (e.g., a CPU chip, power supply, or laser diode) and reformatted (typically to a larger area); and

2) *heat rejection*, in which the heat is transferred to the air.

The past several years have witnessed dramatic advances in the heat-extraction step, including microfluidic devices (e.g., microchannel coolers), heat pipes, heat pumps, and thermoelectric coolers. In marked contrast, technology for heat rejection to the surrounding air has not advanced significantly for the past 40 years. As a result of this technology stagnation, the specifications of commercially available devices for heat rejection are very similar from the standpoint of performance and engineering tradeoffs, despite multiple manufacturers for large and varied markets.
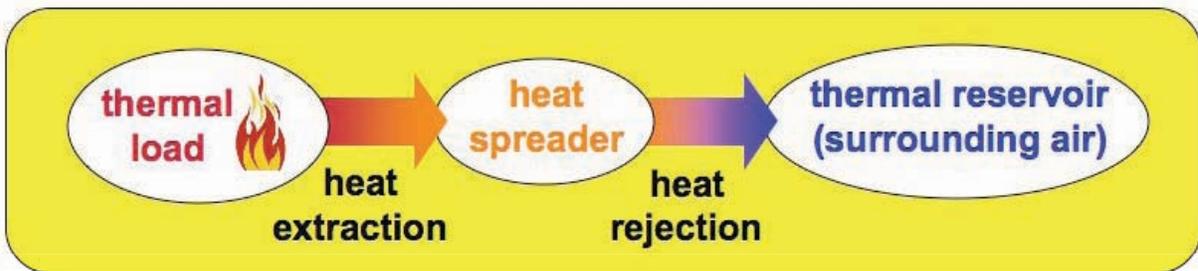
Figure 2: Schematic representation of the two steps involved in cooling a high-flux heat source. Despite significant progress in the heat-extraction step, the technology and performance of systems for heat rejection to the surrounding air have stagnated for decades.

Of course, heat disposal may also involve transfer to water or another coolant, but for the vast majority of practical applications, the goal is to transfer waste heat to the large thermal reservoir provided by the surrounding atmosphere. With the exception of heat pipes, which can be implemented in the form of a hermetically sealed metal enclosure having no moving parts, there has also been a great deal of reluctance to adopt cooling methodologies that entail any kind of liquid handling and/or containment.

In recent years, the greatly increased size, weight and power consumption of air-cooled heat exchangers used for CPU cooling have begun to reach the limits of practicality for applications such as mass produced personal computers for use in home and office environments, data centers and server farms. The high level of audible noise generated by the larger, more powerful fans used in high-
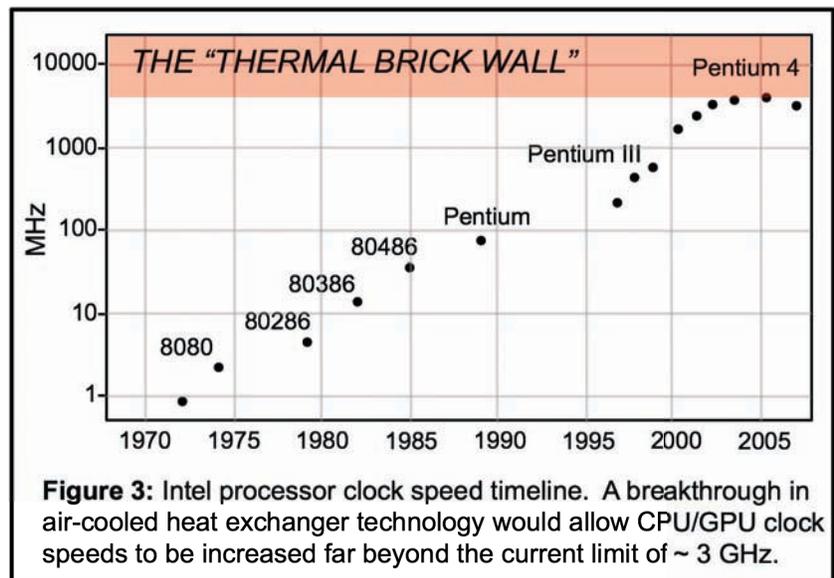
Figure 3: Intel processor clock speed timeline. A breakthrough in air-cooled heat exchanger technology would allow CPU/GPU clock speeds to be increased far beyond the current limit of ~ 3 GHz.

capacity CPU coolers has also proved a deterrent to further scaling up of such devices. Meanwhile, progress in VLSI chip fabrication technology has continued. In many real-world applications, the performance of air-cooled heat exchanger technology is now the primary limiting factor to further improvements in CPU performance. Continued progress along the growth curve of Moore's Law is no longer dictated solely by improvements in VLSI technology. Because of thermal limitations, VLSI advances such as higher transistor density and the ability to operate at higher clock speeds can no longer be readily exploited, a situation widely referred to as the "thermal brick wall" problem. Efforts must now focus on improving the three specific cooling capacity metrics for heat exchangers: cooling capacity per unit volume (W K^{-1} m^{-3}), cooling capacity per unit weight (W K^{-1} kg^{-1}), and cooling capacity per unit power consumption (K^{-1}). The essence of the thermal brick wall problem is that all practical options for increasing the specific capacity of devices such as CPU coolers appear to have already been exhausted.

The thermal brick wall problem also underscores the fact that the lack of progress in heat exchanger technology for electronics cooling in previous years has certainly not been for lack of trying. The current world market for electronics thermal management technology is ~$7B (BCC Research Report SMC024E). Despite an enormous amount of economic and technological pressure to improve upon air-cooled heat exchanger technology, very little progress has been made. In addition to the thermal brick wall problem, there is growing concern over energy usage in the information technology (IT) industry. The IT industry recently surpassed the aviation industry in annual energy consumption, and IT energy usage is projected to double in the next 3 years. Approximately 50% of this energy is used for cooling equipment.

Figure 4: Heat sink fouling in a conventional CPU cooler. The finned heat sink, which is stationary, cannot be seen because it's completely covered in dust. The fan blades, which rotate at several thousand rpm, are nearly dust-free despite operating in the same environment.

Having familiarized ourselves with one of the key applications of air-cooled heat exchanger changers, let us now consider in more detail the underlying limitations of the FFHS architecture. The efficiency of the heat exchange process (e_G) can be quantified by considering the maximum theoretical value of thermal conductance (G) for an ideal heat exchanger, in which heat transfer is limited only by the heat capacity of the flowing air stream:

$$\varepsilon_G = \frac{G}{G_{ideal}} = \frac{G}{C_p \rho \Phi} = \frac{1}{R C_p \rho \Phi}$$

where G and R are the thermal conductance and resistance, respectively, of the non-ideal heat sink, C_p is the heat capacity of air at constant pressure, ρ is the density of air, and Φ is the volumetric flow rate of air through the finned heat sink. For example, this calculation may be carried out for a commercially available CPU cooler typical of those used in many desktop computers. For example, the calculation below is for a Bitspower model NP15S CPU cooler
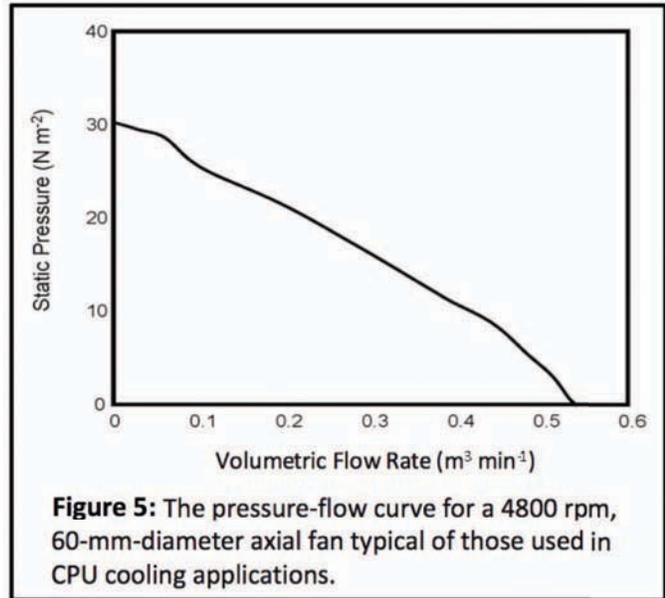
with thermal resistance rating of 0.92 K W^{-1}, fan diameter of 60 mm, and fan speed: 4800 rpm.

$$\varepsilon_G = \frac{1}{\left(0.92\,KW^{-1}\right)\left(10^3\,J\,kg^{-1}\,K^{-1}\right)\left(1.2\,kg\,m^{-3}\right)\left(9.5x10^{-3}\,m^3\,s^{-1}\right)} = 0.10$$

Accordingly, the temperature of the air discharged by such a CPU cooler is only slightly greater than the temperature of the surrounding ambient air, even if the CPU is running very hot. The above calculation is informative because it suggests a great deal of room for improvement in the efficiency of the air-heat-sink interaction. To address this issue, what is needed is a method for substantially reducing the thickness of the heat sink boundary layer without incurring substantial penalties in electrical power consumption, size, weight, cost, complexity, etc.

As shown in Figure 4, another fundamental weakness of the FFHS architecture is the inevitability of heat exchanger fouling. A conventional FFHS device comprises a high-surface-to-volume structure to which a large flow rate of ambient air is delivered on a continuous basis to affect heat transfer. But operation of such an FFHS device also delivers a steady stream of particulates and other airborne foreign matter to this same high-surface-to-volume structure. The accumulation of dust on the heat exchanger surfaces acts as an insulating blanket that interferes with heat transfer, and eventually (as shown in Figure 4) chokes off air flow. In applications where regular preventative maintenance is practiced, the solution to this problem is an intake filter that is cleaned or replaced on a regular basis. Although the pressure drop across such filters can result in a significant penalty from the standpoint of energy usage, especially towards the later portion of their service life, in most applications there is little alternative. In addition, the filter element, which must be a compromise between particle collection efficiency and impedance to flow, does not prevent fine particulates from reaching the heat exchanger surface. Despite the enormous potential payoff in real-world applications, no one has yet developed an air-cooled heat exchanger with immunity to heat sink fouling.

Figure 5: The pressure-flow curve for a 4800 rpm, 60-mm-diameter axial fan typical of those used in CPU cooling applications.

Turning our attention to the next issue, for devices based on the standard FFHS architecture, the flow rate of air is in part limited by restrictions the electrical power consumption and/or noise associated with the fan. In principle, the flow rate of air through the heat exchanger can always be increased by increasing the speed of the fan, but as before, running the fan at high rpm involves a tradeoff with respect to electrical power consumption that quickly becomes very unfavorable. To better understand the role of fan performance, it is informative to consider in absolute terms, the efficiency of the fan.

The efficiency of the fan is defined as the amount of pressure-volume (p-V) work imparted to the airflow per unit time divided by the mechanical shaft power (P_{mech}) supplied to the fan motor.

$$\varepsilon_{fan} = \frac{p_{fan}\,\Phi_{fan}}{P_{mech}} = \frac{p_{fan}\,\Phi_{fan}}{\varepsilon_{motor}\,V_{motor}\,I_{motor}}$$

where the variables ε, p Φ, P, V, and I are efficiency, pressure drop, volumetric flow rate, power, voltage, and current, respectively. Fan and CPU cooler data sheets provide data in the form of a pressure-flow (p-F) curve and specifications for voltage (V) and current (I), but rarely if ever do such data sheets provide a specification for fan efficiency. The efficiency of the small brushless motors used in such fans is typically 60%.

Figure 5 shows the pressure-flow curve for a 60-mm-diameter axial fan typical of those used in CPU cooling applications. The pressure-flow curve of such fans is typically an approximately straight line of the form:

$$p(\Phi) = p_{max}\left(1 - \frac{\Phi}{\Phi_{max}}\right)$$

The $p(\Phi)$ curve has a maximum efficiency operating point at which the product of Φ and p is maximized:

$$p_{fan} = \frac{p_{fan(max)}}{2} \qquad\qquad \Phi_{fan} = \frac{\Phi_{fan(max)}}{2}$$

where $p_{fan(max)}$ is the stagnation pressure of the fan, and $\Phi_{fan(max)}$ is the "open circuit" flow rate of the fan. An estimate of the maximum operating efficiency of the fan is therefore:

$$\varepsilon_{fan} = \frac{p_{fan(max)}\,\Phi_{fan(max)}}{4\,\varepsilon_{motor}\,V_{motor}\,I_{motor}}$$

For the 60-mm-diameter, 4800-rpm, cooling fan of Figure 5, under typical operating conditions, the mechanical efficiency of the fan is estimated to be 3.1%:

$$\varepsilon_{fan} = \frac{(30\,N\,m^{-2})(0.53\,m^3\,min^{-1})}{4\,(0.60)(12\,V)(0.30\,A)(60\,s\,min^{-1})} = 0.031$$

This is characteristic of the small-area, high-speed fans used for cooling of electronics and other small devices. Only ~3% of the mechanical power delivered to the fan is typically imparted to the airflow; the other 97% of this power is wasted on fan-blade viscous drag (at operating speed, the mechanical resistance of the fan motor bearings is negligible compared to the total drag force exerted on the fan blades). This large unproductive expenditure of mechanical energy is also the primary source of fan noise. Much larger, lower-rpm fans achieve far better efficiency, as high as 80% for large industrial fans that consume in excess of 5 kW of electrical power. In between these two extremes, the mechanical efficiency of the fan used in a typical 3-ton residential air conditioner condenser is of order 25% [Parker, 2005].

The key implication of the above calculation is that the specific cooling capacity of conventional CPU coolers falls far short of what is physically possible. An efficiency of ~3%

implies that in principle, drastic reductions in power consumption could be realized using a different approach. To a large extent the same argument also applies to fans having efficiencies of order 25% such as those used in residential air conditioners.

Figure 6: Photograph (left) and schematic diagram (right) of version 1 prototype device.
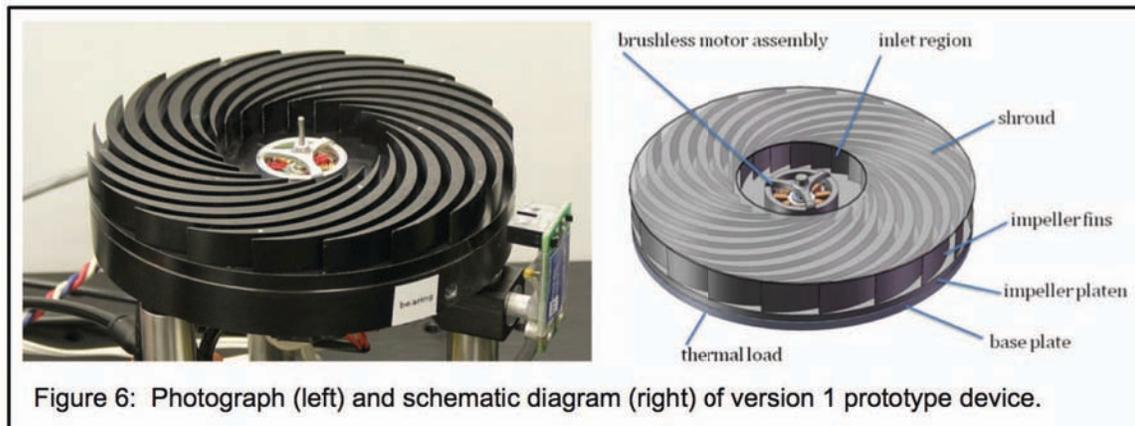
The alternative device architecture invented by Dr. Koplow at Sandia National Labs solves all four of the fundamental problems of conventional air-cooled heat exchanger technology simultaneously (the boundary layer effect, heat-sink fouling, the poor efficiency of small high-speed turbo-machinery, and noise). A prototype of this device, which we refer to as an "Air Bearing Heat Exchanger", is shown in Figure 6. This first prototype device was designed for an electronics cooling application and therefore only has a diameter of 10 cm.

As is the case in a conventional "fan-plus-heat-sink" CPU cooler, the heat load is placed in thermal contact with the bottom surface of an aluminum base plate that functions as a heat spreader. As in a conventional CPU cooler, this heat spreader plate is stationary. In a conventional CPU cooler, the top surface of the heat spreader base plate is populated with fins. In the air bearing heat exchanger, instead of having fins, the top of the heat spreader base plate is simply a flat surface.

The "heat-sink-impeller" (the finned, rotating component) consists of a disc-shaped heat spreader populated with fins on its top surface, and functions like a hybrid of a conventional finned metal heat sink and an impeller. Air is drawn in the downward direction into the central region having no fins, and expelled in the radial direction through the dense array of fins. A high efficiency brushless motor mounted directly to the base plate is used to impart rotation (several thousand rpm) to the heat-sink-impeller structure. The bottom surface of this rotating disc-shaped heat spreader is flat, such that it can mate with the top surface of the heat spreader plate described above.

During operation, these two flat surfaces are a separated by a thin (~0.03 mm) air gap, much like the bottom surface of an air hockey puck and the top surface of an air hockey table. This air gap is a hydrodynamic gas bearing, analogous to those used to support the read/write head of computer disk drive (but with many orders of magnitude looser mechanical tolerances).

Heat flows from the stationary aluminum base plate to the rotating heat-sink-impeller through this 0.03-mm-thick circular disk of air. As shown later in Figure 18, this air-filled thermal interface has very low thermal resistance and is in no way a limiting factor to device performance; its cross sectional area is large relative to its thickness, and because the air that occupies the gap region is violently sheared between the lower surface (stationary) and the upper surface (rotating at several thousand rpm). The convective mixing provided by this

shearing effect provides a several-fold increase in thermal conductivity of the air in the gap region.

One important point about the air bearing is that the ~0.03 mm air gap is not maintained by using extremely tight mechanical tolerances. Much like an air hockey puck on an air hockey table, or a hard disk read/write head, the air gap distance is self-regulating. If the air gap distance increases, the air pressure in the gap region drops, which causes the air gap distance to decrease. This built in negative feedback provides excellent mechanical stability and an extremely stiff effective spring constant (important for ruggedness). Unlike an air hockey table, which relies on gravity to counter-balance the pressure force acting on the puck, the air-bearing cooler can be mounted in an arbitrary orientation (e.g., up-side-down, sideways, etc.). And unlike a computer disk drive, incidental mechanical contact between the two air bearing surfaces does not damage either surface.

This rotating heat exchanger geometry places the thermal boundary layer in an accelerating frame of reference. Placing the boundary layer in this non-inertial frame of reference adds a new force term to the Navier-Stokes equations, whose steady state solution governs the functional form of the heat-sink-impeller flow field [Schlichting, 1979]. At a rotation speed of several thousand rpm, the magnitude of this centrifugal (in the frame of reference of the boundary layer) force term is as such that as much as a factor of ten reduction in average boundary layer thickness is predicted [Cobb, 1956]. Unlike techniques such as air jet impingement cooling, the mechanism for boundary layer thinning in the air bearing heat exchanger does not rely on a process that entails dissipation of significant amounts of energy, nor is the boundary layer thinning effect localized in a small area. Rather, the centrifugal force generated by rotation acts on all surfaces simultaneously, and all portions of the finned heat sink are subject to the resulting boundary layer thinning effect. For the limiting case of flat rotating disk, an exact solution of the Navier-Stokes equation is possible and indicates that the magnitude of the boundary-layer thinning effect is constant as a function of radial position. This high-speed rotation also provides violent shearing within the planar air gap region, providing efficient convective heat transfer between the base plate and heat-sink-impeller.

Rotation of the heat sink at several thousand rpm also provides a potent remedy to the longstanding problem of heat exchanger fouling. Consider for example the CPU cooler shown in Figure 1. The finned, metal heat sink cannot be seen because it's covered in dust. But the fan blade, which operates in the same environment, is for all intents and purposes perfectly clean. This contrast in dust accumulation is at first startling, but in hindsight entirely expected. The air bearing heat exchanger therefore provides a complete solution to the problem of performance degradation due to heat sink fouling. In specialized applications involving extremely high particle loading, a straight-radial rather than backward-swept fin design would likely be used [Bleier, 1997].

The Sandia device architecture circumvents the poor efficiency of small, high-speed fans by using the mechanical work provided by the motor to directly generate relative motion between the heat exchanger structure and the surrounding air. This highly efficient "direct drive" scheme for creating relative motion between the finned structure and the surrounding air also results in a drastic reduction in noise. In addition to the boundary layer thinning effect, immunity to heat sink fouling, and reduction in electrical power consumption, this was the other motivation for attempting to implement some form of rotating heat sink architecture. Unlike a fan or other type of conventional turbo-machinery, doing large amounts of mechanical work on the surrounding fluid is not the primary objective of the heat-sink-impeller. Rather, the impeller fin design for the most part emphasizes streamlined passage of air through the adjacent flow channels, prevention of flow separation from the
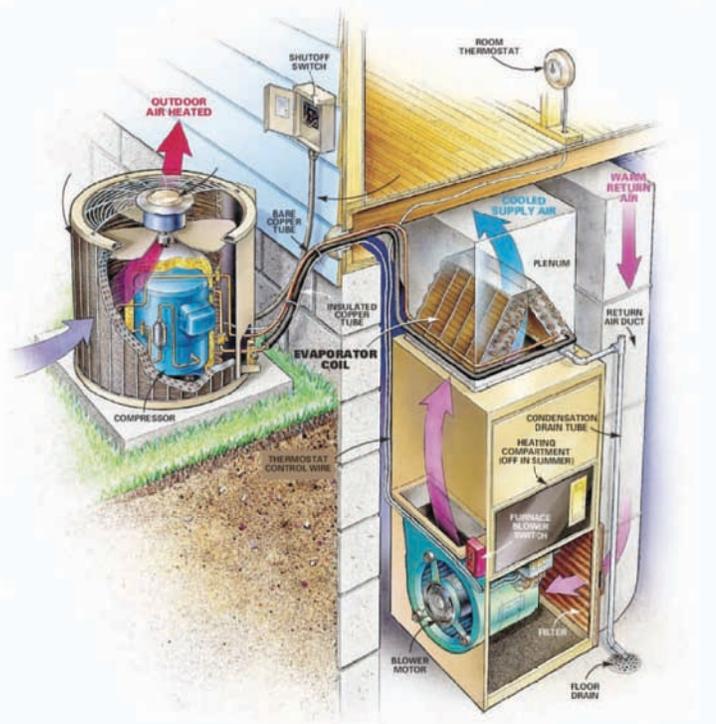
blade surface, and smooth rejoining of the adjacent air channel flow fields as the air exits the heat-sink-impeller.

One other important point concerns the simplicity of this new device architecture relative to conventional high performance CPU coolers. An integral part of the high performance CPU cooler shown in Figure 1 is an embedded heat pipe structure. Because of boundary layer effects, this device must have a very large surface area (relative to the heat load) to achieve low thermal resistance. The average thermal path length from the heat load to exterior surface of the heat exchanger is of order 80 mm. The only way to relay the heat load such a long distance without introducing a large amount of series thermal resistance is to incorporate a heat pipe structure directly into the finned copper heat sink; the additional series resistance introduced by a solid copper conduction path of the same dimensions would be prohibitively large. The device based on a rotating heat sink impeller does not require nearly as much surface area to achieve low thermal resistance because of the boundary-layer thinning effect described above. As a result, the average thermal path length is much lower (of order 10 mm). The finned heat sink therefore need not incorporate an internal heat pipe structure.

The performance of air-cooled heat exchanger technology is critically important in other portions of the energy sector as well, most notably for devices such as air conditioners, heat pumps and refrigeration equipment (i.e., any heat engine in which the source/sink temperature difference is relatively small). While the use of Sandia's new air cooled heat exchanger technology in applications such as air conditioners has many features in common with electronics cooling applications, but there are important differences as well. Consider for example the condenser unit of the split system air conditioner depicted in Figure 7. Heat exchanger thermal resistance, heat exchanger fouling, and fan noise are clearly all important considerations. Fan efficiency is of minor importance, however, because nearly all of the electrical power consumption in an air conditioner is associated with operation of the compressor.

In air conditioning applications, boundary layer thermal resistance and fan noise both end up having a bearing on heat exchanger thermal resistance, and thus air conditioner coefficient of performance (COP).

Figure 7: Typical split system residential HVAC installation.

The temperature drop across the thermal resistance of the air-cooled heat exchanger figures directly into the Carnot efficiency of the heat pump it is in thermal contact with; this is the reason that reduced airflow and heat exchanger fouling can have such a substantial effect on air conditioner electrical power consumption and cooling capacity. For example, consider a hypothetical air conditioner comprised of a heat pump sandwiched between air-cooled heat exchangers. Let us assume for the sake of discussion that the thermal resistance of the heat

exchangers on the hot and cold side of the heat pump is the same. The Carnot coefficient of refrigerator performance is given by:

$$\gamma_{Carnot} = \frac{T_{SINK} - qR}{\Delta T + 2qR}$$

where T_{sink} is internal air temperature of the residence, ΔT is difference between the outdoor and indoor temperatures, q is the amount of heat flow through the air conditioner, and R is the heat exchanger thermal resistance. If R is increased the term in numerator decreases and the term in the denominator increases. This translates directly to a reduction in γ_{Carnot}. The effect is especially pronounced under conditions of high thermal loading (high q).
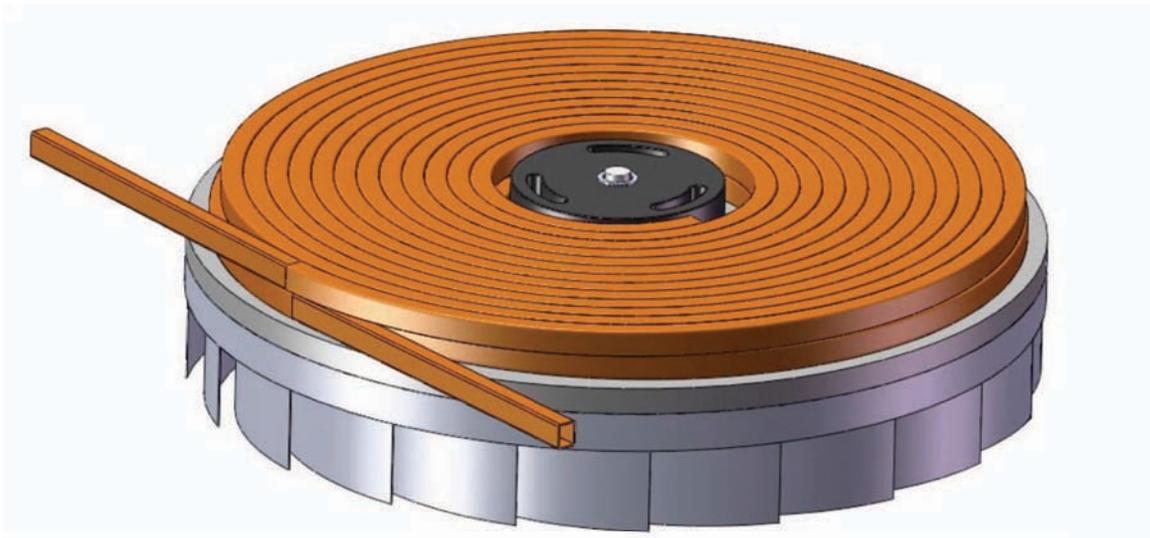
Figure 8: Schematic drawing showing application of air bearing heat exchanger concept to an air conditioner condenser. Such a unit can be made extremely compact compared to a traditional fan-plus-serpentine-finned-tube architecture, and fouling of the outer surface of the soldered copper tubing assembly has no effect on condenser performance. A second identical air bearing cooler (not shown) can be fused to the opposing face of the soldered copper tubing assembly to lower the thermal resistance of the heat exchanger by an additional factor of 2 if desired. An air conditioner evaporator unit can be implemented in an analogous structure. In the case of the air conditioner evaporator, it is also possible to enclose the heat-sink-impeller assembly in a plenum, such that it also serves to circulate air through the HVAC air duct system.

Consider for example the indirect effect that fan noise exerts on air conditioner COP. Air conditioner EER is typically specified by the manufacturer only at maximum fan speed. On the other hand, the end-user may prefer "quiet mode" operation, particularly if the manufacturer attained a high EER rating by incorporating a powerful fan. As described in a 2008 study by Riviere et al., the simplest way to reduce air conditioner noise is to reduce the speed of the heat exchanger fan, but the required reduction in fan speed may decrease air conditioner efficiency by 15 to 30% [Riviere, 2008]. This reflects the greater reality that in many applications of air-cooled heat exchangers (e.g., air conditioners, heat pumps, refrigeration equipment, personal computers), while it would clearly be advantageous to use higher air flow rates or to simply scale up the heat-exchanger-fan assembly in size to achieve a higher COP, such measures cannot be implemented because of noise restrictions. The ability to lift this noise constraint would be very valuable, both from the standpoint of energy

efficiency and consumer satisfaction. The latter consideration is relevant to the question of further enhancing market acceptance.
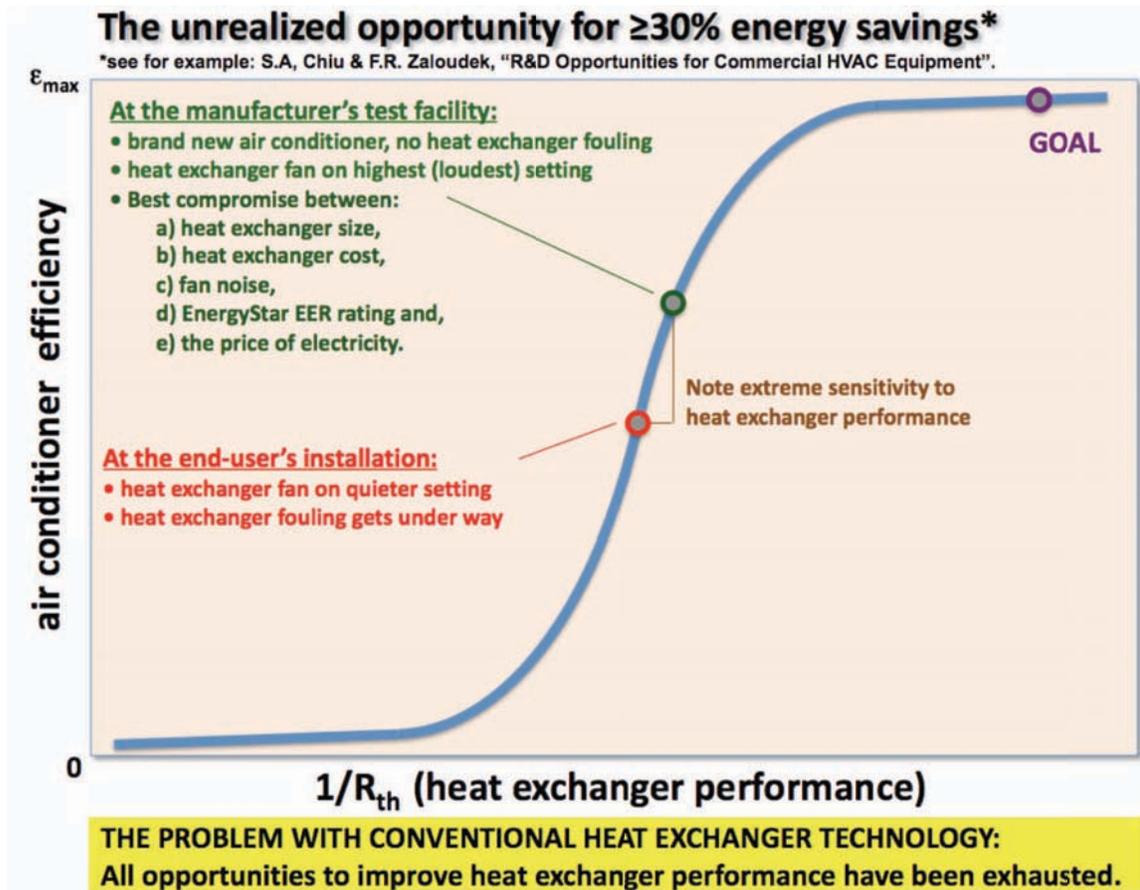
Figure 9: Better air-cooled heat exchanger technology will result in tremendous energy savings in air conditioning (above) and other applications such as heat pumps and refrigeration equipment.

From the standpoint of energy conservation, immunity to heat exchanger fouling also has significant implications. For example, as reported by Mowris in 2006, in a study conducted by Trane, the effects of heat exchanger fouling in air conditioners over a 4 to 8 year period under average contaminant loading were monitored (Mowris estimated an average air conditioner service lifetime of 15 years.). The measurements conducted by Trane indicated a 17% loss of cooling capacity and a 27% reduction in efficiency [Mowris, 2006]. Residential air conditioners in particular are often subject to extreme negligence with regard to preventative maintenance. To the extent that the consumer regards an air conditioner as a self-contained black box that provides cool air whenever it's turned on, an Energy Star unit destined to suffer a 25% increase in power consumption early in its service life provides little of its purported benefits.

Having enumerated the benefits of higher airflow rates and clean heat exchangers, the significance of reducing the thickness of the insulating boundary layer can readily be appreciated. With regard to the question of fan efficiency, in contrast to the case of a CPU cooler, reductions in heat exchanger fan electrical power consumption conferred by the use of air bearing heat exchanger technology can only be expected to provide incremental benefits in air conditioner efficiency. As mentioned earlier, air conditioner power consumption is dominated the compressor, and the problem of low fan efficiency (typically 25% in a 3-ton residential air conditioner unit) is less of a driver than in electronics cooling

where fan efficiencies are often just a few percent. While the prospect of an additional 2 to 4% increase in air conditioner efficiency that can be realized by reducing fan losses is certainly welcome [Parker, 2005], this effect will be swamped by direct reductions in heat exchanger thermal resistance (i.e., the boundary layer thinning effect), the indirect effect of drastically reduced fan noise on COP, and the elimination of heat exchanger fouling.

2. Breakthrough results

There is no single figure of merit to quantify the performance of an air-cooled heat exchanger. In most applications, thermal resistance (i.e. cooling power), electrical power size, cost, and reliability (e.g. resistance to fouling) are key considerations. These and other variables coexist in an engineering trade space that is complex and application specific. The best way to assess whether a new air-cooling technology represents a significant advance or breakthrough, is to compare its performance to state-of-the-art air-cooling technology in the context of one or more specific applications.

For example, a survey of commercial CPU coolers indicates that a conventional fan-plus-heat-sink device equal in size to the version 1 prototype device in Figure 6 has a typical thermal resistance of 0.6 to 0.8 C/W. Our version 1 prototype device on the other hand has a measured thermal resistance of 0.2 C/W. This represents a huge advance in a field that has long seen only incremental progress in cooling performance. Moreover, as discussed later, we have reason to believe that a 2^{nd} generation prototype could readily achieve a thermal resistance of ≤0.1 C/W in a device of the same size.

A second basis for comparison is provided by the performance specifications given in the 2008 DARPA MACE call for proposals on ultra-high-performance air-cooling technology. As shown in the table below, our version 1 prototype device was able to provide the same cooling performance as DARPA's state-of-the-art device while providing a factor of 4 reduction in size and more than a factor 10 reduction in electrical power consumption. Again, this clearly represents a breakthrough in performance.

Table 1: laboratory data	thermal resistance	power consumption	device volume
DARPA state-of-the-art system	0.2 C W^{-1}	100 W	1050 cm^3
Proof-of-concept device (Figure 6)	0.2 C W^{-1}	20 W	240 cm^3

Other more qualitative performance comparisons are important to consider as well. For example, state-of-the-art CPU coolers providing thermal resistances as low as 0.2 C/W are not only very large devices unsuitable for mass market applications, but in addition are prohibitively expensive. A great deal of this expense is associated with the need for an integral heat pipe to distribute heat to distant locations within the massive heat sink structure. The device developed in this study has no requirement for a heat pipe; the heat-sink-impeller is a monolithic structure that could be fabricated by die-casting, stamping, etc. Another such comparison concerns the fact that the air-cooling architecture of the present study is the only device architecture to date that provides intrinsic resistance to heat sink fouling, a problem that has long been considered unsolvable. As a final example, we note that the much lower noise levels generated by this new air cooling architecture has significant implications for scaling up air flow (and therefore cooling), as discussed in Section 1. While we have yet to record comparative laboratory data pertaining to sound levels, consistent with the discussion of fan noise in Section 1, it is clear that low noise operation will ultimately be another important advantage in real-world applications.

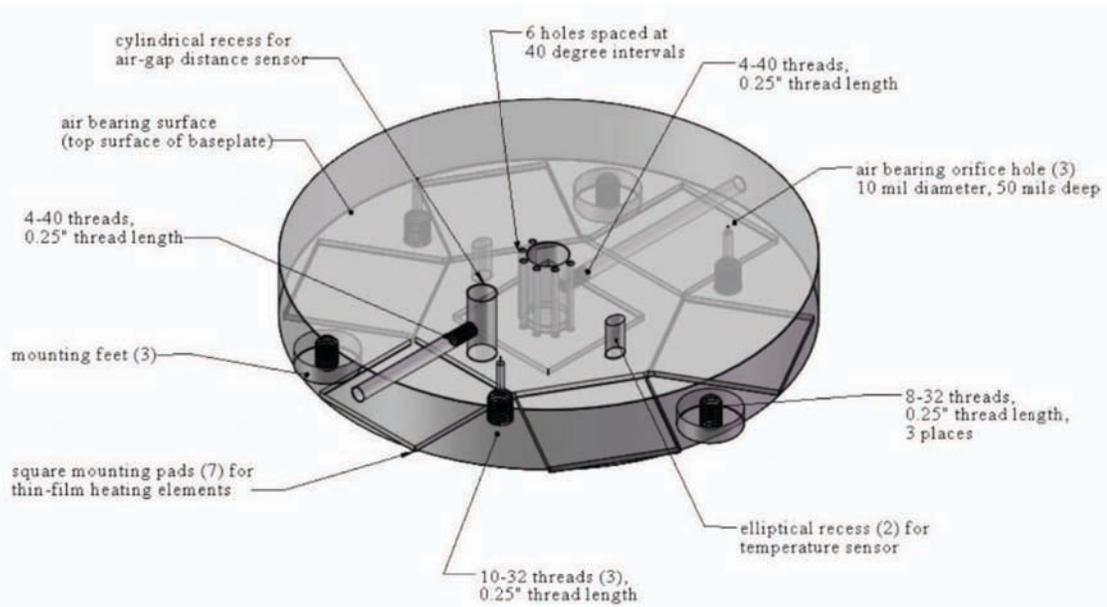## 3. Experimental apparatus and prototype device design

All laboratory measurements were conducted using the device test bed shown in Figure 10. The air bearing cooler (9) is visible at the lower right inside the protective metal mesh enclosure (7). Please refer back to Figure 6 for a close-up photograph of the prototype air bearing cooler.

In this device test bed, ambient air temperature is monitored by a thermocouple probe (3) located above the device air intake, and the air bearing cooler base plate temperature is measured by thermistor beads (not visible) potted in thermally conductive epoxy and embedded in the aluminum base plate. Two thermistor beads in different locations are used for redundancy. An array of six 1" x "1, 10-Watt, thin-film-Kapton resistive heaters are attached to the bottom of the base plate. The current and voltage readout from the dc power supply used to power the heating element array is used to calculate the input thermal loading. The calibration of the current and readouts of the power supply was in turn checked against a calibrated Fluke multimeter, and heavy gauge current delivery leads were used to ensure negligible voltage drop between the dc power supply and heating element array.

Figure 10: Experimental set up. 1) vacuum line, 2) micrometer used for secondary calibration of distance sensor, 3) thermocouple probe used to measure air intake temperature, 4) 3-way valve (vacuum/blocked/gas-delivery), 5) pressure regulator, 6) mass flow controller, 7) protective wire mesh enclosure, 8) readout for Keyence EX-305V air gap distance sensor, 9) air bearing cooler visible behind wire mesh enclosure, 10) supply manifold for externally pressurized air bearing, 11) motor phase DP3T switch (used to choose between (a) connecting the three brushless motor leads to the Motortron-9A brushless motor driver, (b) leave three brushless motor leads floating, or (c) apply a dc current to one of the three motor phases to check for current dependent z axis rotor movement (using the micrometer).

cylindrical recess for
air-gap distance sensor

6 holes spaced at
40 degree intervals

4-40 threads,
0.25" thread length

air bearing surface
(top surface of baseplate)

air bearing orifice hole (3)
10 mil diameter, 50 mils deep

4-40 threads,
0.25" thread length

mounting feet (3)

8-32 threads,
0.25" thread length,
3 places

square mounting pads (7) for
thin-film heating elements

elliptical recess (2) for
temperature sensor

10-32 threads (3),
0.25" thread length

Figure 11a: SolidWorks CAD drawing of air bearing cooler base plate. The base plate was CNC machined out of a solid block of 7075 aluminum and anodized black to facilitate the recording of infrared temperature measurements.

recess/shoulder for mounting of
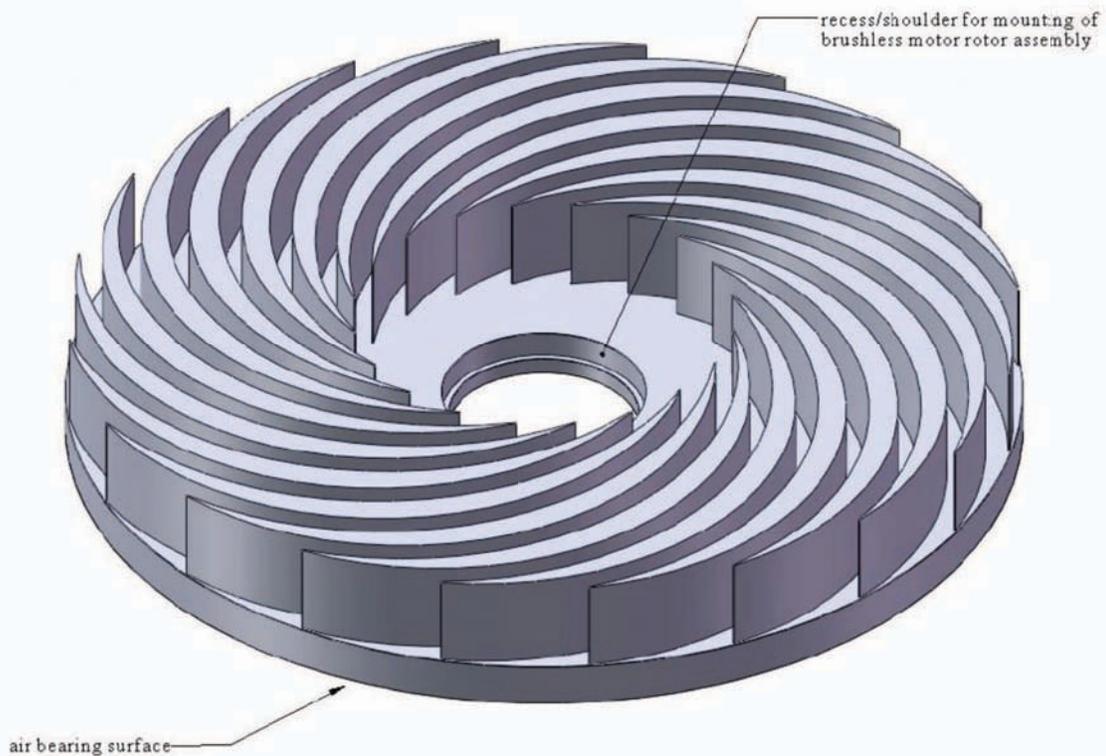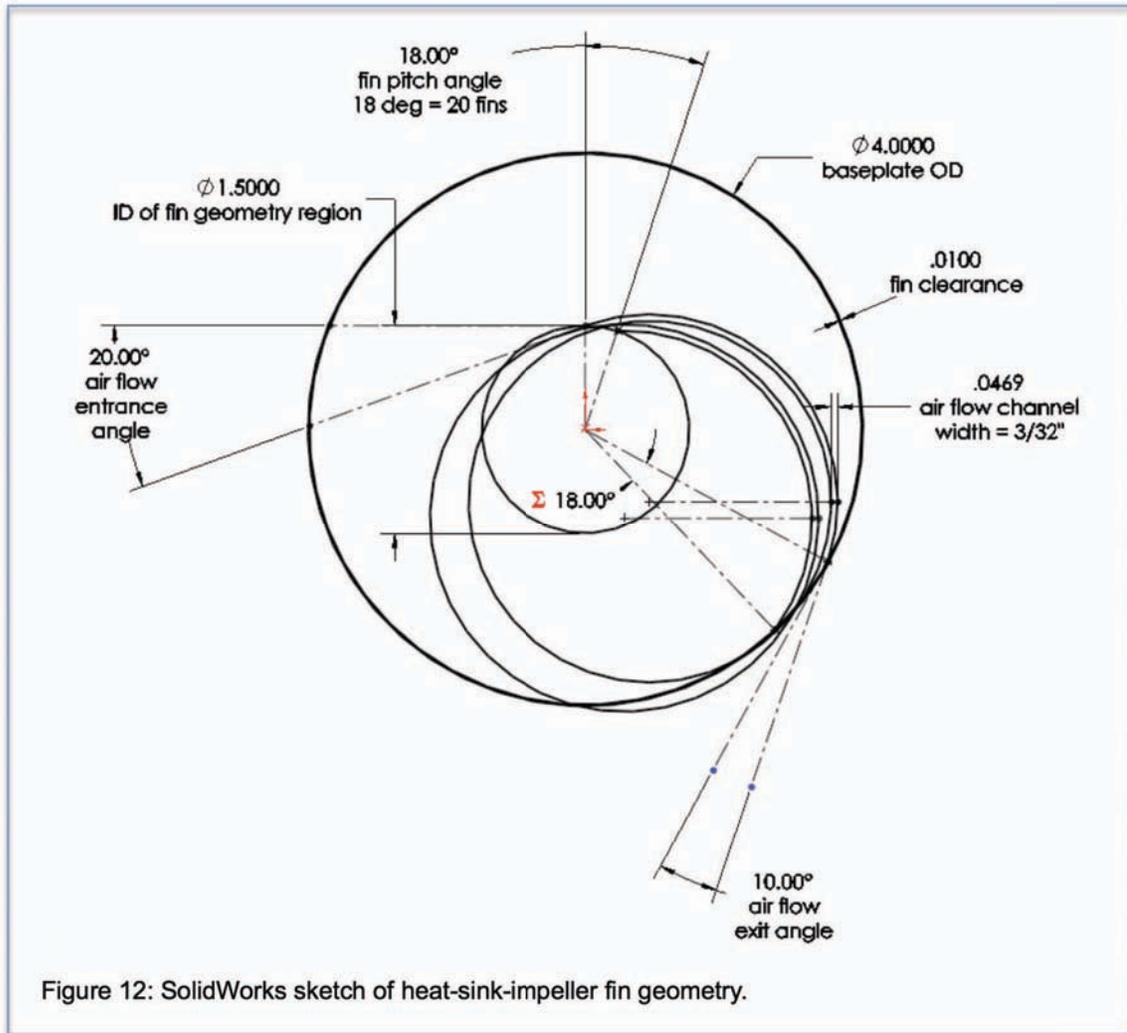brushless motor rotor assembly

air bearing surface

Figure 11b: SolidWorks CAD drawing of heat-sink-impeller. The heat-sink-impeller was CNC machined out of a solid block of 7075 aluminum and anodized black to facilitate the recording of infrared temperature measurements.

17

Figure 12: SolidWorks sketch of heat-sink-impeller fin geometry.

The prototype device is configured as a static (externally pressurized) thrust bearing. In real-world thermal management applications such an externally pressurized air bearing would be replaced by a hydrodynamic (self-pressurizing) air bearing, which uses a minute fraction of the mechanical power supplied by the brushless motor to generate the required lifting force. For experimental measurements, however, an externally pressurized air bearing is preferable because it allows the air gap distance to be varied systematically, and over a wide dynamic range. The mass flow rate of air required to float the air bearing is very small, < 10 sccm at an air gap distance of 10 μm. This externally supplied gas flow is delivered to a triangular array of orifice holes (0.010"-diameter), and each orifice is equipped with a flow restricting element in the form of a 125 μm diameter, 5-mm-long length of capillary tubing located just upstream of each orifice hole. Flow restrictors minimize the amount of dead volume added to the air gap volume by the gas delivery lines, which is desirable from the standpoint of air bearing stability (e.g., suppression of oscillations). The incorporation of such flow restrictors also ensures nominally equal division of mass flow between the three orifice holes. In addition, the inclusion of flow restrictors allows highly reproducible adjustment of the flow rate in the absence of a mass flow controller (see for example Figure 13), which in some circumstances is advantageous. The gas delivery manifold is equipped with a manual three-way value through which it can be connected to a source of compressed gas, a vacuum pump, or neither (blocked off). In turn, the gas delivery line has second three-way valve to select

either nitrogen or helium, which is necessary for determination of the air gap thermal resistance.

Figure 13: Air gap distance as a function of mass flow rate for the externally pressurized air bearing of the version 1 prototype air bearing cooler.
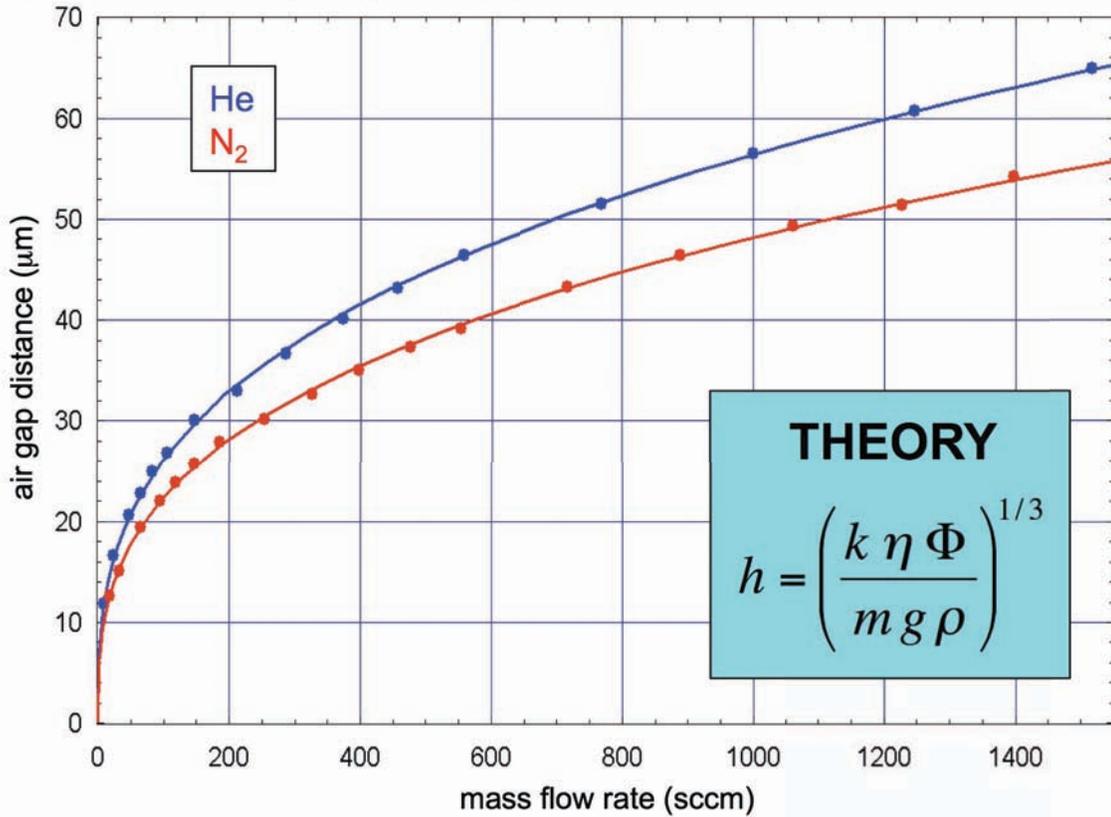
The air gap distance between the top surface of the base plate and the bottom surface of the heat-sink-impeller is monitored with an eddy current displacement sensor (Keyence, Model EX-305V, 0.4 µm resolution, 1000 µm dynamic range). The model 305V sensor has a sampling rate of 40 kHz, which ensures adequate bandwidth to record mechanical vibration waveforms in addition to average air gap distance. The 5-mm-diameter cylindrical sensor head is recessed with respect to the upper surface of the base plate by approximately 100 µm. It is mounted in a 5-mm diameter clearance hole through the base plate, held in place by a set screw. Silicone rubber sealant at the junction of the sensor head and the bottom surface of the base plate is used to prevent leakage of air between the outer diameter of the sensor head and the inner diameter of the clearance hole through which the sensor head is inserted. The gap distance sensor is zeroed by evacuating the air gap region to bring the base plate and heat-sink-impeller air bearing surfaces into positive contact. A three-point calibration protocol is then conducted using 0.5-mm and 1-mm ring-shaped precision shims. In future experiments it would be advisable to use a triangular array of three sensors rather than a single sensor, which would provide further information about heat-sink-impeller wobble and sensor measurement reproducibility.

Brushless motor current, voltage and rpm are monitored continuously using a USB-interfaced Power Analyzer Pro instrument module from Medusa Research. The Power Analyzer Pro is also used to log the ambient and base plate temperature sensor data. These data are analyzed and displayed in the Power ProView software (version 2.2.1.2). The brushless motor

controller, or "electronic speed controller (ESC)" used in to drive the brushless motor is the Model MDrive-9A ESC manufactured by Motortron. This particular ESC was chosen for its programming flexibility, although a more compact ESC will likely be used in future work. The MDrive-9A ESC is programmed, controlled, and monitored using Motortron's ESC view software (version 1.17). The user interface for the software used to run the Power Analyzer Pro instrument and MDrive-9A ESC is shown in Figure 14.
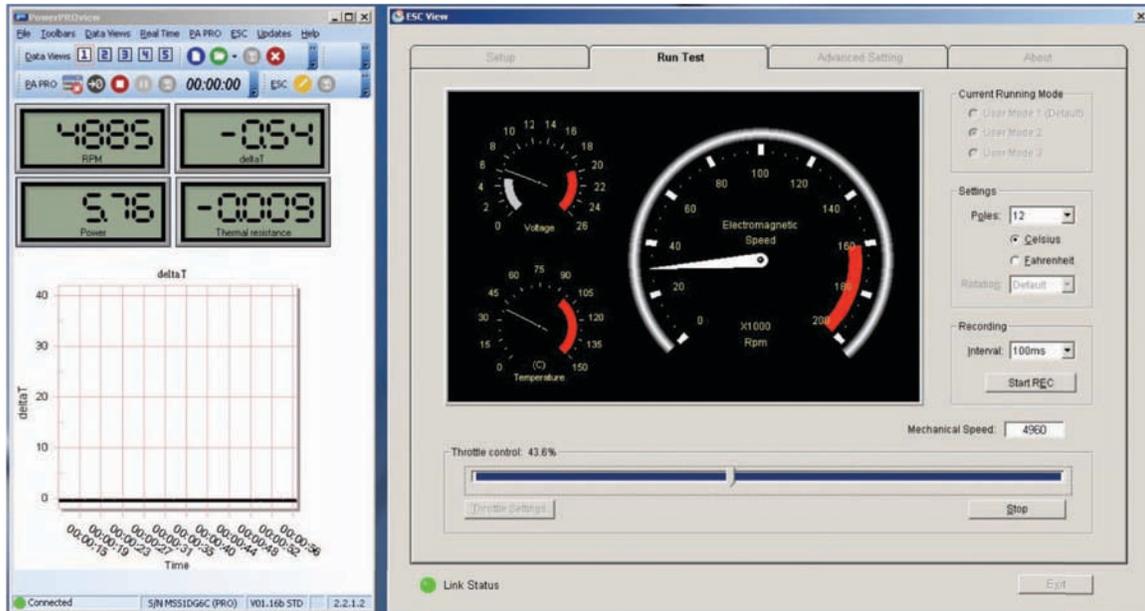

Figure 14: Software user interface used to control/monitor motor, and log/process data.

The brushless motor chosen for this project was the Model MiniGBx, which is available as a kit from GoBrushless Inc. The MiniGBx is a 3-phase, external rotor brushless motor having 9 stator poles, 12 rotor poles, and an outer diameter of 17.8 mm (Figure 15). Purchasing the motor in the form of a kit is advantageous from the standpoint of winding the stator to match the mechanical impedance of the load (e.g., the torque-speed curve of the heat-sink-impeller).

Figure 15: At left, 9-pole, 3-phase internal rotor. At right, 12-pole permanent magnet external rotor.

This is important from the standpoint of design flexibility, and from the standpoint of attaining the highest possible brushless motor efficiency. In the prototype device, each of the three phases were wound with 1.53 m of 26 AWG copper magnet wire, resulting in a total of 13 winding turns per individual stator pole, 39 windings per phase, and 117 windings total. Each of the three phases has a dc resistance of 2.05 Ω, and an inductance of 15.0 µH (which corresponds to an L/R time constant of 73 µs).

Another important distinction concerns the difference between a conventional and "sensorless" brushless motor controller. In a conventional brushless motor, Hall-effect sensors are incorporated into the motor housing to measure the position (phase) of the rotor. These timing signals are used to determine when the motor controller should "fire" each phase of the stator. Such electronic commutation avoids the numerous problems associated with brushed motors, most notably the generation electromagnetic interference and limited service life. The later parameter is critically important in the vast majority of air cooling applications. A sensorless brushless motor infers rotor phase from the back-emf waveform generated by one or more unexcited phases at any given time. Elimination of the need for hall effect sensors reduces motor size, cost, and complexity. It was for this reason that a sensorless brushless motor was chosen for the prototype device. The only "drawback" of a sensorless brushless motor is a minimum operating rpm below which the amplitude of the back-emf waveform is too low to detect with adequate signal to noise. In the current application, operation at extremely low speeds is obviously not of interest. For example, the minimum operating speed of the MiniGBx motor is in the vicinity of 1000 rpm (this number dependents somewhat on the number of windings per phase), and the typical operating speed for the air bearing cooler is of order several thousand rpm.

We now describe the design of the prototype heat-sink-impeller. The x-y geometry (i.e., footprint) chosen for the heat-sink-impeller fin array was an educated guess based on a survey of the heat transfer and turbo-machinery literature. Most likely the geometry chosen is substantially less than optimal because insufficient resources were available to conduct any formal fluid dynamic modeling. The primary design objectives were,

1) maximize fin surface area to promote heat transfer,
2) maximize the cross-sectional area of the radial air channels (to minimize flow resistance),
3) avoid constrictions in the radial air channels by maintaining constant cross-sectional area,
4) ensure fins have adequate internal thermal conductivity in the z direction,
5) smoothly separate and rejoin air flow at air channel entry/exit to minimize noise and aerodynamic drag,
6) implement a backward-swept fin geometry, which is known to provide better efficiency (i.e. requires less motor power) than a radial or forward-swept geometry,
7) avoid constricting air the flow in the central intake region of the heat-sink-impeller,
8) use a geometry that is compatible with low cost fabrication techniques such as die-casting, stamping, etc.

Obviously many of the above design objectives are conflicting. Figure 12 shows the final specifications arrived at for the xy geometry of the version 1 prototype fin array. Most notably:

1) The number of fins used was 20,
2) The width of the air flow channels was 3/32",
3) The air flow radial entrance angle was 20°,
4) The air flow tangential exit angle was 10°, and
5) The annular finned region corresponds to 80% of the platen surface area.

The other design parameter of fundamental importance is the fin height, which in the version 1 prototype was chosen to be 0.40". This choice of fin height reflected a desire to have the total device size (base plate plus heat-sink-impeller) fit within the guidelines specified in the 2008 DARPA MACE proposal for advanced air-cooling systems. In the prototype device the base plate was made far thicker than necessary (0.60") to facilitate mechanical mounting of sensors, heating elements, motor hardware and quick-connect air delivery lines. The platen of heat-sink-impeller to which the fins are mounted was also made much thicker than necessary (3/16") to guarantee the mechanical rigidity required to obtain a 6.3 micro-inch rms surface finish during fabrication (to avoid any uncertainties regarding the roll of surface finish in the air gap thermal impedance).
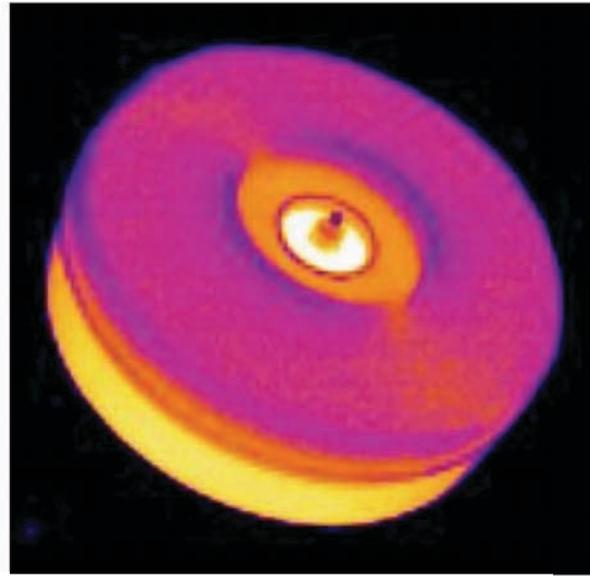

Figure 16: Infrared thermal image of air bearing cooler under typical operating conditions.

3A Cooler prototype with two watt fan
0.55°C/W

Aerocool AB7080H
0.58°C/W

Aerocool Glacier
0.64°C/W

Aerocool Hercules
0.62°C/W

Aerocool X-Facto (aluminium fan)
0.62°C/W

Aerocool X-Facto (black fan)
0.56°C/W

Alpha PAL6035
0.66°C/W

Alpha PAL8045 (variable speed fan, maximum power)
0.56°C/W

Alpha FC-PAL15
0.96°C/W

Alpha PEP66
0.7°C/W

Arctic Cooling Super Silent 2000
0.90°C/W

The resulting fin height (0.40") was later shown experimentally to be far shorter than optimum. For example, as shown in Figure 16, the temperature distribution recorded by infrared thermal imaging (FLIR Systems Thermovision model 20A camera) indicates that with the exception of the sharp leading edge of the heat-sink-impeller fins in the vicinity of the air intake region, the fin array is nearly isothermal in the z direction (here we are referring to the pink regions shown in Figure 16). This is also consistent with the very low ratio of fin internal thermal resistance (calculated to be 0.0188 C/W for a 0.40" high fin array) to overall device thermal resistance (an order of magnitude higher). Substantially better performance will therefore be realized in a version 2 prototype device by redesigning the base plate to have a thickness of order 0.20", reducing the platen thickness to of order 0.10", and extending the fin height to of order 1.00". The above information also suggests that a larger number of fins, having smaller xy cross-sectional area, should be used to increase heat exchanger surface area. Informed selection of such numerical parameters will should be informed by finite-element analysis of internal heat flow, as well as the mechanical stress field associated with operation at high rpm. In the present study, insufficient resources were available to conduct such analyses.

Such a 2[nd] generation device is expected to have approximately a factor of two lower thermal resistance (~0.10 C/W) than that achieved in the version 1 device (~0.20 C/W). Note that in CPU cooling applications, a device with a 0.2 C/W thermal resistance is considered "ultra-high performance" (and is too large, noisy and expensive for non-specialty applications), and 0.1 C/W in a device comparable in dimensions to the version 1 prototype would be truly revolutionary. The CPU coolers used in typical mass market computers (a sampling of which are shown at right) have thermal resistances in the 0.6 to 0.9

C/W range. Based on the discussion in Section 1 concerning air-cooling applications in the energy sector, presumably such a "factor-of-several" improvement in air-cooled heat exchanger performance would also resolve the longstanding efficiency penalty suffered by air conditioners, heat pumps, etc. discussed in Section 1 (please refer back to Figure 9).

A few other remarks are in order concerning the infrared image shown in Figure 16:

1) The temperature drop between the base plate and platen is readily discernible.

2) Apparent deviations from azimuthal symmetry are known to measurement artifacts associated with changes in apparent emissivity as a function of object viewing angle; the image exposure time is long compared to rotational period of the high sink impeller, and without such artifacts, no azimuthal image structure should be visible.

Figure 17: Equivalent circuit for steady state operation.

3) The color-coded temperature scale provided by the FLIR infrared camera software is monotonic in temperature but exhibits rapid changes in color over certain ranges of temperature. It is therefore difficult to make quantitative inferences about temperature differences by simply looking at Figure 16.

4. Heat transfer measurements

All laboratory measurements were conducted using the device test bed shown in Figure 10. The objectives of the heat transfer measurements were:

- sf □ □□□□□□□□□□□□□□ □□□ □□□□□ □) □□□□ x□
- uf □□□□□□□□□ □□□□□□□□□□□□□□□□□ □□□□□□□□□□□□□□□□□□□□□□□□ □□□□□□□□□□ □□□□□ □□□□□□□□□□s px□
- rf □□□□□□ □□□ □□□□□□ □□□□□□□□□□□□□□□□□□□□□□□□□□□□□ □□□x□
- cf □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ □□□□□ □□□□□□□□□) □□□ □□□□□□□□□□□□□□□□□□□□□□□□x□
- tf □□□□□□□□□□□ □□□□□) □□□□ □□ □□□□□□□□ □a□of a□□□□□□□□ □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□a□□□□□□□□□o □□□□□□i □□ □□□□□□□□□□□□□□□) a□□□□□□□□□□□) □□□ □□□□□□□□□bx□

The protocol for measuring the total thermal resistance of the device as a whole is a simple steady state measurement, in which a known amount of power is dissipated (P) into the base plate and the temperature rise of the base plate above ambient temperature is recorded:

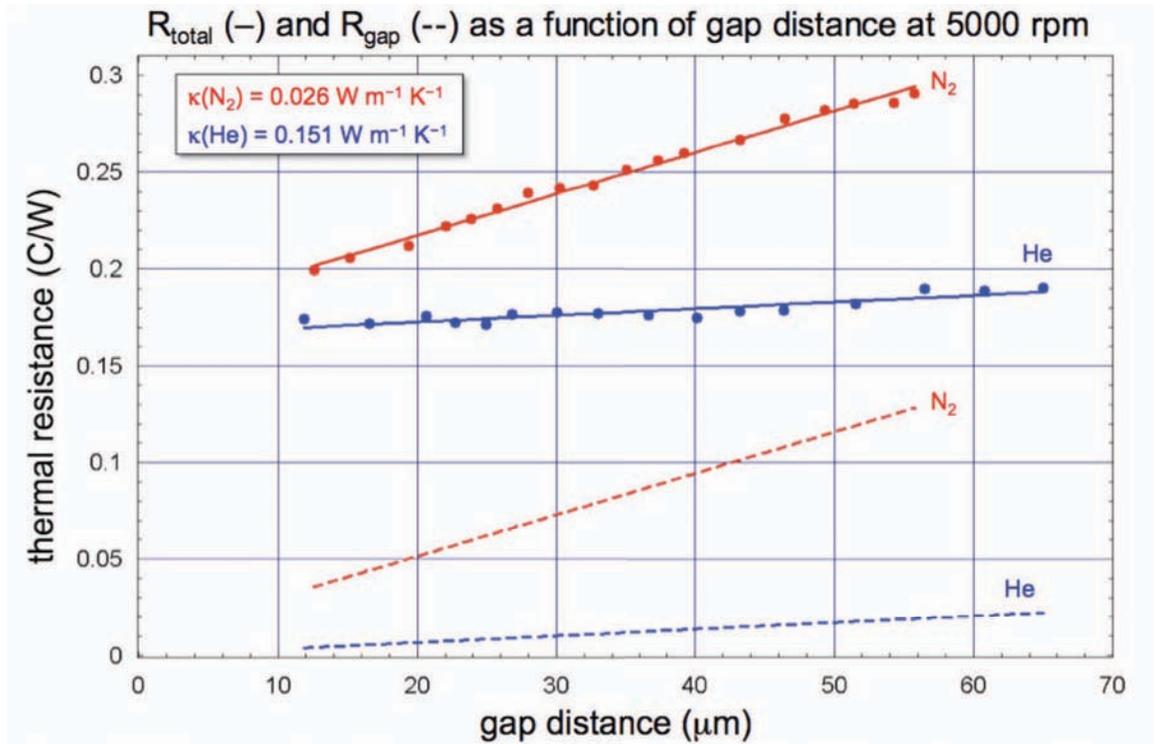$$R_{total} = \frac{T_{baseplate} - T_{ambient}}{P_{heater}}$$

Figure18: Determination of R_{total} and R_{gap} as a function of gap distance at ω = 5000 rpm.

As shown in Figure 17, R_{total} can be modeled as a chain of series thermal resistance terms for heat flow through the device, in parallel with a leakage thermal resistance that accounts for incidental leakage of heat to the surrounding air when the heat sink impeller is stationary. This effect is small but nonetheless accounted for in the analysis that follows.

Starting at the top of the resistor diagram, $R_{baseplate}$ is the thermal resistance for internal heat flow from the bottom side of the base plate (where the heaters are located) to the top side of the base plate (where the air gap region begins). It can be accurately calculated and is $R_{baseplate}$ = 10.4 mK/W. The air gap thermal resistance cannot be calculated because of the unknown effect of the violently sheared flow in the air gap region from the standpoint of heat transfer. It must be determined by an indirect experimental measurement technique and has a complicated functional dependence on h and ω. As with $R_{baseplate}$, R_{platen} and R_{fins} are both simple internal conduction calculations in which heat flow is assumed to be along the z direction. The calculated value of R_{platen} is 3.5 mK/W. The thermal resistance for heat transfer from the top to the bottom of the 0.40" height fin array is calculated to be 9.4 mK/W. As mentioned earlier, during operation, the variation in fin temperature as a function of z is negligible. Accordingly, we use the distance-averaged value of R_{fins} = 4.7 mK/W for heat transfer calculations. R_{xfer} is the thermal resistance for heat transfer from the outer surface of the heat-sink-impeller to the surrounding air as a result of forced convection. R_{xfer} depends on

ω but not on h. It is the dominant thermal resistance term and is largely driven by boundary layer thickness. We will discuss momentarily how R_{xfer} is determined experimentally. Lastly, $R_{leakage}$ is the parallel leakage resistance term that accounts for incidental leakage of heat due to conduction and natural convection (i.e. when the heat-sink-impeller is not rotating) from the base plate to the surrounding air when the heat sink impeller is stationary. $R_{leakage}$ was measured experimentally to be 2060 mK/C. While this relatively high leakage resistance doesn't represent a major perturbation to heat transfer measurements, it is worth correcting for and is included for in all subsequent analysis.

As alluded to earlier, R_{gap} is determined indirectly from a set of two measurements of R_{total} using different air gap gases of known but very different thermal conductivities (helium and nitrogen). In the following algebraic expression, α is the ratio of thermal conductivities for He/N$_2$ gas (6.0), and the symbol Y = 1/R is represents thermal admittance.

$$R_{gap(N_2)} = \frac{\alpha}{\alpha - 1}\left[\frac{1}{Y_{total(N_2)} - Y_{leakage}} - \frac{1}{Y_{total(He)} - Y_{leakage}}\right]$$

The use of dry nitrogen, rather than air, is an experimental convenience. As shown in Appendix B, the thermal conductivities air and nitrogen differ by less than 1%. The result of a typical set of experimental measurements is shown in Figure 18. A series of such measurements at different angular velocities between 2,000 and 10,000 rpm was used to determine the functional dependence R(h, ω).

isolated heat–sink–impeller thermal resistance as a function of angular velocity
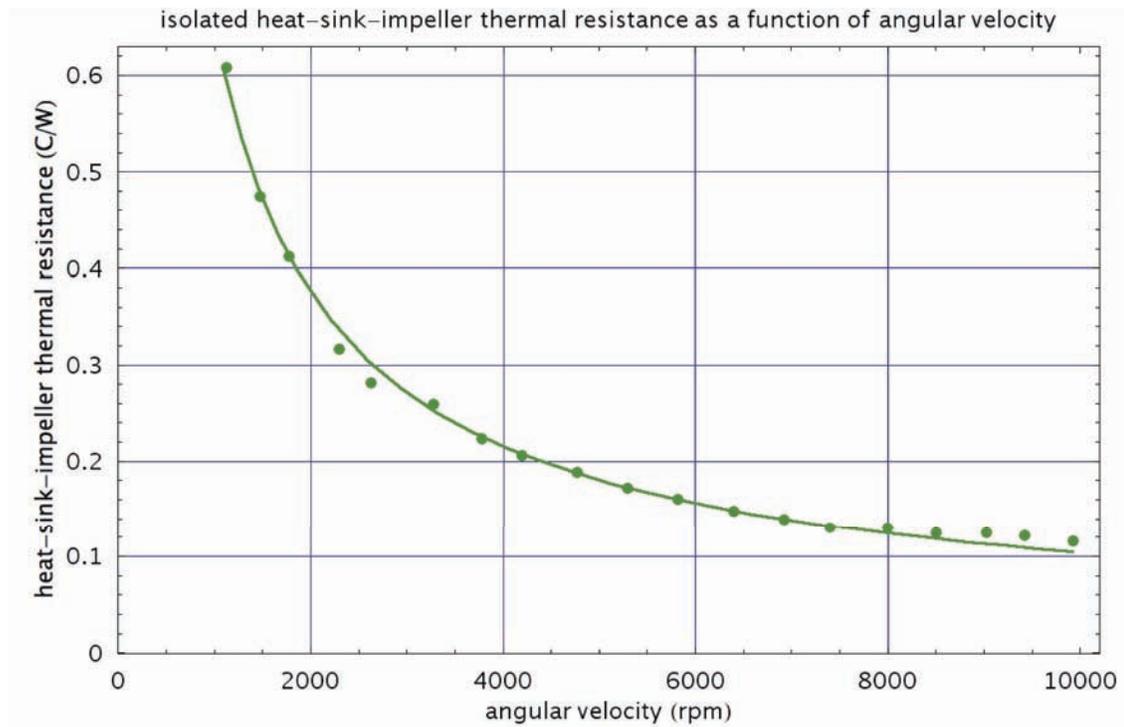
Figure 19: Experimentally measured value of R_{xfer} as a function of angular velocity.

The final thermal resistance parameter $R_{xfer}(\omega)$ can then be isolated by subtracting out the effect of all other thermal resistance terms (now known) to furnish a plot such as that shown in Figure 19. To improve R_{xfer} measurement accuracy, it is advantageous to use helium rather than nitrogen to reduce the correction for air gap thermal resistance to ~0.01 C/W. The 0.8-

order power dependence of $R_{xfer}(\omega)$ is noteworthy because it strongly suggests that the heat-sink-impeller operates in the turbulent flow regime over the range of angular velocity for which data was recorded (1000 to 10000 rpm). Such a 0.8 power dependence of heat transfer on flow velocity is observed in numerous systems having widely varying geometries. For example, Figure 20 shows that similar heat transfer behavior is observed for internal pipe flow [Kreith, 2001]. In this plot, the quantity plotted on the x axis is Reynolds number (Re, dimensionless), which is linear proportional to flow velocity, and the quantity plotted on the y axis is Nusselt number (Nu, dimensionless), which is linearly proportional to the rate of convective heat transfer per unit surface area. As shown in Figure 20, in the limit of low Re the flow is laminar and the power dependence for Nu(Re) is 0.3. In the limit of high Re the flow is turbulent and has a Nu(Re) power dependence of 0.8. The approximate range 2000 < Re < 4000 is where the flow becomes unstable, and the transition from laminar to turbulent behavior is observed. Analogous behavior is observed for the case of heat transfer from a rotating disk [Cobb, 1956].   In the case of a rotating disk, the asymptotic power dependence of Nu(Re) in the limit of high Re (turbulent regime) is also 0.8. As with the case of pipe flow, the asymptotic behavior in the limit of low Re (laminar flow) has a well defined power dependence, in this case 0.5. Note that in the case of a rotating disk Re is proportional to angular velocity rather than linear velocity (as was the case for pipe flow).
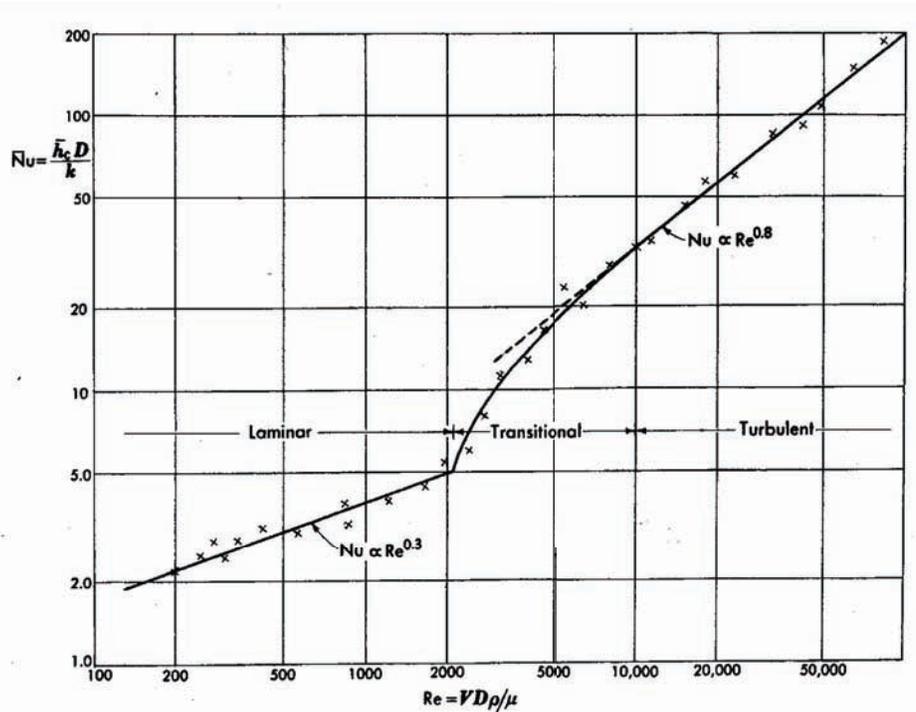

Figure 20: Experimentally measured value of Nusselt number (Nu) as a function of Reynolds number (Re) for internal flow through a round pipe.

In future work it would be very worthwhile to record data over a wide enough dynamic range in Re to map out the entire Nu(Re) dependence from laminar, to transitional, to turbulent flow. Such data would be extremely useful from the standpoint of rigorously validating computational fluid dynamic (CFD) code that will be required for device optimization. As mentioned earlier, the prototype device uses a sensorless brushless motor, which is limited to a minimum angular velocity of 1000 rpm. In future experiments the use of a more complex brushless motor having Hall-effect phase-angle sensors would allow operation down to very low rpm. Having said that, from a practical standpoint in real world applications of devices

such as the current prototype we are only interested in device behavior at several thousand rpm, where low thermal resistance is obtained.

A nonlinear least squares fit of the data points in Figure 19 (the green curve) yielded:

$$R_{xfer}(\omega) = \frac{k}{\omega^n} \qquad k = 1.6 \times 10^2 \; K \, W^{-1} \, rpm^n \qquad n = 0.80$$

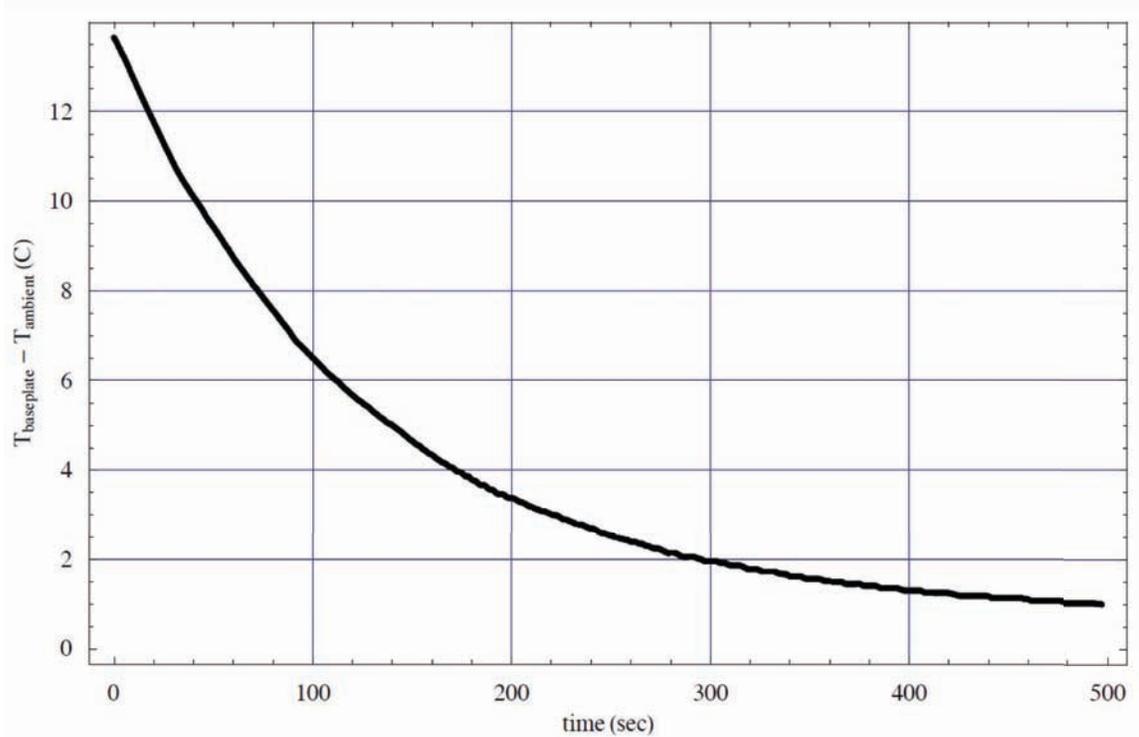We will use the above expression in subsequent system models as our mathematical model for $R_{xfer}(\omega)$.

Figure 21: Temperature decay curve data for validation of steady state measurement technique.

Because we know that R_{xfer}, which represents the thermal resistance of the heat-sink-impeller boundary layer, is the dominant source of thermal resistance, the primary focus of future work will be heat-sink-impeller design optimization. Such work would involve CFD calculations to guide heat-sink-impeller design, and experimental measurements to test performance and provide CFD model validation. Based on the experience to date concerning experimental measurement of R_{total}, R_{xfer}, etc., such an experimental effort should emphasize R_{xfer}, measurement accuracy, and the ability to quickly and easily conduct R_{xfer} measurements of prototype heat sink impeller geometries. The best way to accomplish this goal is to fabricate the CNC-machine prototype heat-sink-impellers in the form of monolithic disk in which the finned structure is present on both sides of the disk. The use of such a double-sided heat-sink-impeller eliminates the need to attempt to correct for heat transfer at the bottom (flat) surface of the platen (and the measured thermal resistance is simply be half of that for a one-sided heat-sink-impeller). Such measurements would include a hall-effect sensor brushless motor to permit operation at low rpm, and a convenient means of attaching/detaching the device under test to the motor spindle such as a keyed shaft.

The R_{xfer} measurement protocol would entail running the double-sided heat-sink-impeller at a fixed rpm, applying a heat load to the rotating structure (e.g. with a heat gun) to bring it substantially above ambient temperature (e.g. by 20 C), recording the exponential temperature to room temperature observed when the heat source is turned off using a non-contact infrared temperature sensor. The time constant of the thermal decay will be $\tau = R_{xfer}$ C_{hsi}, where C_{hsi} is the heat capacity of the heat sink impeller (calculated from its mass and specific heat). It would also be advisable to fabricate the keyed shaft from a material having relatively low thermal conductivity. If necessary, another option would be to provide for independent closed-loop control of the shaft temperature, which would then be maintained at the heat-sink-impeller temperature at all times (including during the $\sim 10^2$ seconds during which transient thermal decay is recorded). Such a thermal decay curve was recorded for the prototype heat-sink-impeller and is shown in Figure 21. This measurement was conducted to provide an independent validation of steady state measurement technique used through out this study. The thermal resistance values obtained with the steady state and transient decay measurements agreed to within a few percent.
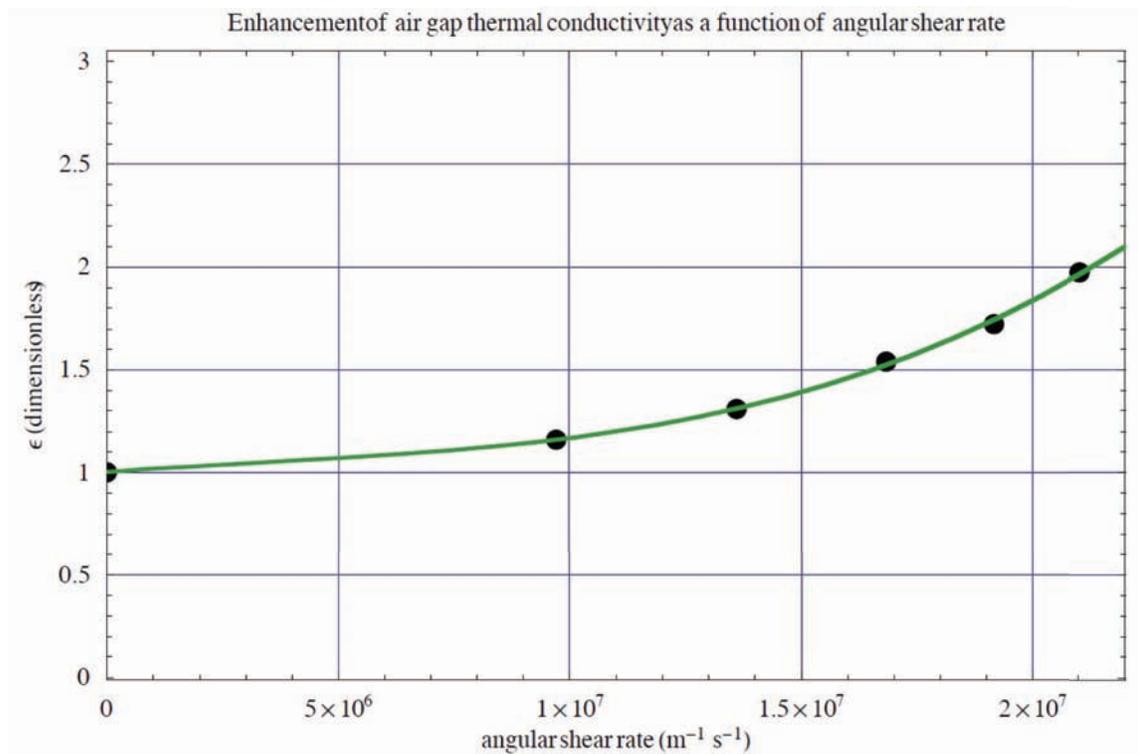

Figure 22: Measured enhancement of air gap thermal conductivity as a function of shear rate.

With regard to the experimental work done to determine the typical range of air gap thermal resistance, the most important conclusions are that R_{gap} can be made very small (a few tens of mK/W) at gap distances that are relatively large (~ 20 μm) from the standpoint of mechanical tolerance (such as flatness, surface roughness, etc.). The other important experimental measurement pertaining to air gap thermal resistance concerns the effect of sheared flow in air gap region. Figure 22 shows the experimentally observed enhancement in air gap thermal conductivity relative to the base line case of motionless air as a function of angular shear rate (which is proportional to angular velocity).

In the absence of shearing (rotation), the thermal conductivity for thermal transport across the

air gap region can be calculated simply from the geometry of the air gap and the thermal

$$R_{gap(stagnant)} = \frac{h}{\kappa A}$$

conductivity of the fill gas:

where h is the gap distance, A is the gap area, and κ is the thermal conductivity of the fill gas. During operation, shearing of the air gap region enhances the thermal conductivity by a factor of ε via convective mixing:

$$\kappa_{gap(sheared)} = \varepsilon\,\kappa_{gap(stagnant)}$$

Assuming an approximately linear velocity profile in the z direction within the gap region, the absolute shear rate (s = du/dz. where u is the tangential velocity) at a given radial position r is:

$$s = \frac{du}{dz} = \frac{\omega\,r}{h}$$

A theoretical prediction of the functional dependence of e on absolute shear rate is difficult, primarily because of questions surrounding the validity of candidate physical models. Due to time constraints, exploration of this topic is beyond the scope of this study. It is reasonable to expect however, that ε for an area element $dA = r\,d\phi\,dr$ at a given location can be expressed as a power series in s, the absolute shear rate,

$$\varepsilon = 1 + K_1\,s + K_2\,s^2 + K_3\,s^3\ \dots$$

at least up to some critical shear rate at which transition to a different flow regime occurs. ε for the entire air gap region can therefore be written as the sum of terms:

$$\varepsilon = 1 + K_1\langle s\rangle_{area} + K_2\left\langle s^2\right\rangle_{area} + K_3\left\langle s^3\right\rangle_{area}\ \dots$$

where $\langle s^n\rangle_{area}$ is the area-averaged value of s^n. The area-averaged n^{th} moment s is given by:

$$\left\langle s^n\right\rangle_{area} = \frac{\displaystyle\int_{R_{min}}^{R_{max}} r\left(\frac{\omega r}{h}\right)^n dr}{\displaystyle\int_{R_{min}}^{R_{max}} r\,dr} = \frac{2\left(R_{max}^{\,n+2} - R_{min}^{\,n+2}\right)\left(\dfrac{\omega}{h}\right)^n}{(n+2)\left(R_{max}^{\,2} - R_{min}^{\,2}\right)}$$

ε can therefore written as a power series of the independent variable (ω/h):

$$\varepsilon(\omega, h, R_{min}, R_{max}) = \sum_{n=0}^{n_{max}} \frac{2\left(R_{max}^{\,n+2} - R_{min}^{\,n+2}\right)}{(n+2)\left(R_{max}^{\,2} - R_{min}^{\,2}\right)}\left(\frac{\omega}{h}\right)^n$$

We will refer to the quantity $S = \omega/h$ as the "angular shear rate" of the heat-sink-impeller. The best-fit line that passes through data points in Figure 22 is a third-order polynomial in S.

$$\varepsilon = 1 + k_1\,S + k_2\,S^2 + k_3\,S^3\ \dots$$

$k_1 = 1.75 \times 10^{-8}\ m\ s$
$k_2 = -1.38 \times 10^{-15}\ m^2\ s^2$
$k_3 = 1.30 \times 10^{-22}\ m^3\ s^3$

Our mathematical model for air gap thermal resistance as a function of gap distance and angular velocity is therefore:

$$R_{gap}(h, \omega) = \frac{h}{A\,\kappa\left[1 + k_1\frac{\omega}{h} + k_2\left(\frac{\omega}{h}\right)^2 + k_3\left(\frac{\omega}{h}\right)^3\right]}$$

Note that the maximum shear rates likely to be encountered with the current prototype device (or a device of comparable diameter) is of order 3×10^7 m^{-1} s^{-1}, above which the frictional losses in the air gap region would likely be considered impractically large (discussed in Section 5). This corresponds to a maximum thermal conductivity enhancement of $\varepsilon \approx 4$. In the current prototype device, under typical operating conditions ($h \approx 20$ μm, $\omega \approx 5000$ rpm), $\varepsilon \approx 2$. Having said that, in small diameter high-speed devices, such as those that might be used for direct cooling of a CPU IC package, this thermal conductivity enhancement factor may be considerably larger. This effect is potentially very important because the reduction in air gap area associated with small-diameter devices will itself raise the air gap thermal resistance considerably. It should also be noted that in future devices the incorporation of a hydrodynamic bearing, which typically entails the addition of air gap surface features such as grooves, will also likely enhance convective heat transfer in the air gap region. Lastly, in some cases it may also be worthwhile to incorporate additional surface features, texture, etc. for the sole purpose of further enhancing convective heat transfer in the air gap region.
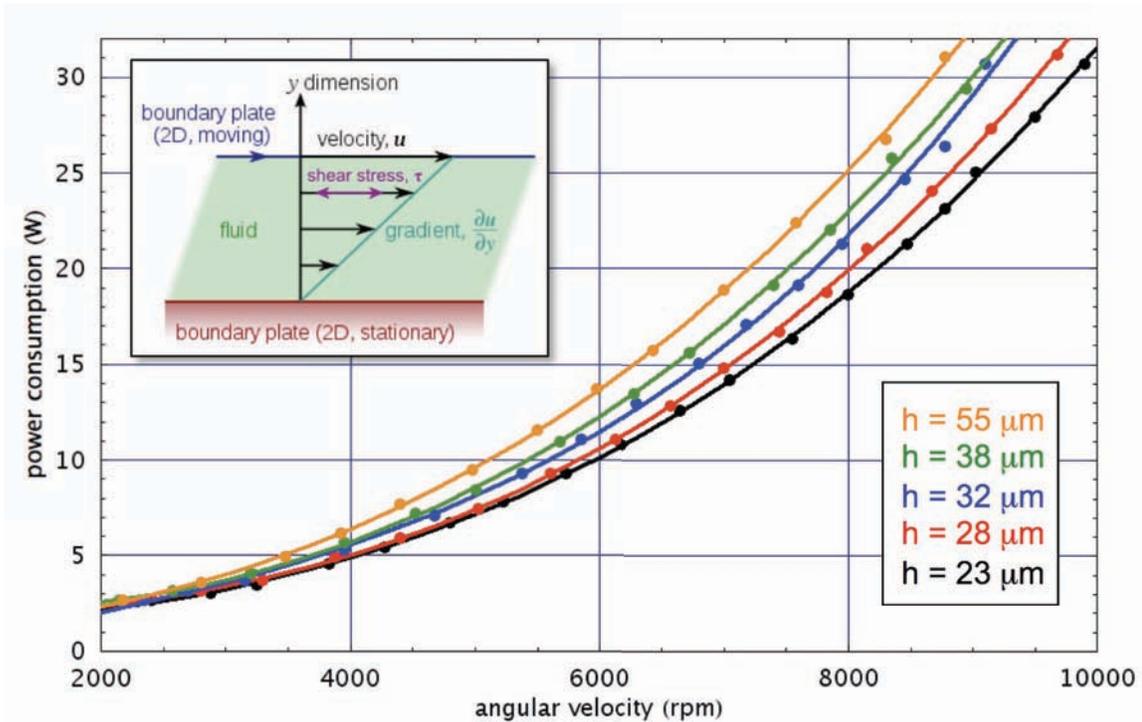
5. Electrical power consumption

Figure 23: Shearing of the of the air gap region is a dissipative process.

As discussed extensively in Section 1, the electrical power consumption of an air-cooled heat exchanger is important in most applications, and coexists in a complex engineering trade

space with other specifications such as size, cost, noise, and cooling performance. The experimental measurement of electrical power consumption is trivial. Accurate measurement of mechanical power consumption, which boils down to measurement of the shaft torque (τ) required to make the heat-sink-impeller rotate at a given angular velocity is considerably more difficult; the electrical-to-mechanical efficiency of a brushless motor in connected to a particular mechanical load is in general neither constant as a function of angular velocity nor known a priori. More difficult still is the determination of how mechanical power consumption is distributed between into its two components, that associated with rotation of the heat-sink-impeller itself (aerodynamic drag pressure-volume work performed by the impeller, brushless motor bearing friction, etc.), and that associated with shearing of the air gap region (see Figure 23 above). The latter power consumption term is in large part what determines the optimum value of air gap distance (h), which represents a compromise between the desire for low thermal resistance and low frictional shearing losses.
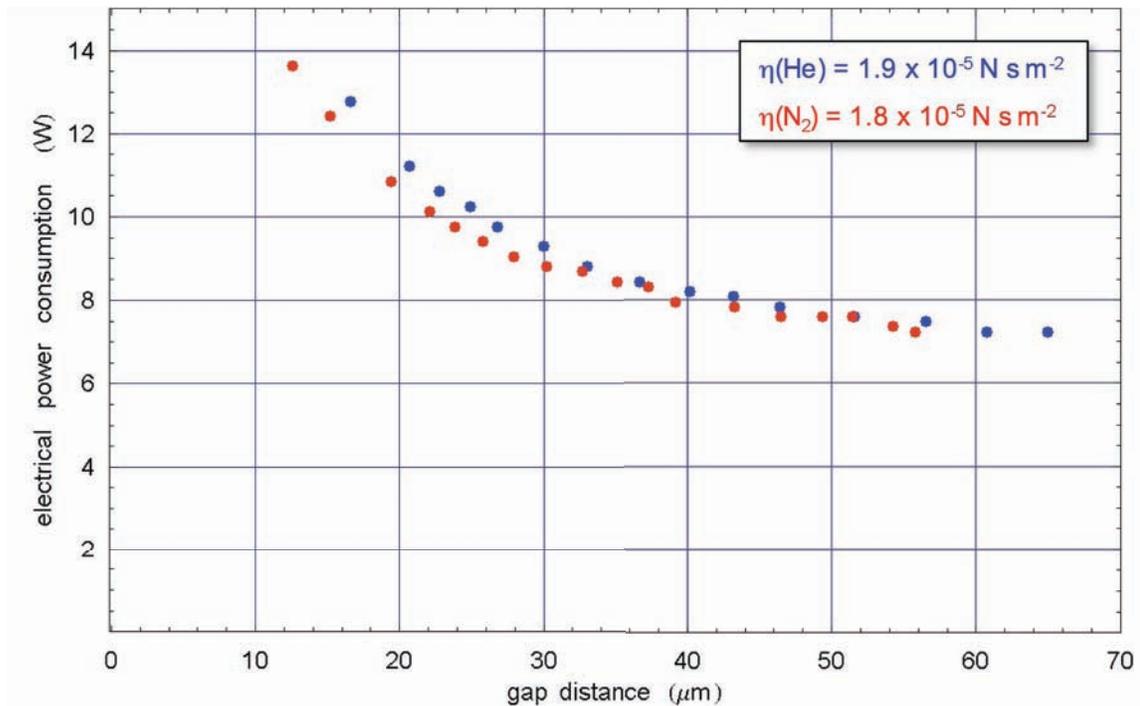

Figure 24: Brushless motor electrical power consumption as a function of air bearing gap distance for ω = 5000 rpm. The data obtained with nitrogen and helium in the air gap region are nearly identical because the absolute viscosities of He and N_2 are nearly equal.

The use of a torque sensor between the brushless motor and mechanical load is not particularly practical in the current study. Difficulties include the size, cost, and general lack of availability of good sensor for measurement of small torque values. We can however infer the value of $\tau(\omega)$ if the moment of inertia of the mechanical load (J) is known (derived from the SolidWorks solid model, see Appendix A) by recording the angular velocity decay of the mechanical load from some initial angular velocity to $\omega = 0$:

$$\tau(\omega) = J\frac{d\omega}{dt}$$

Such an angular velocity decay curve is shown in Figure 25. The best-fit line in Figure 25 is a single exponential decay. Taking the derivative of this fitting function furnishes a measurement of load torque as a function of angular velocity (Figure 26).

31

heat−sink−impeller angular velocity decay curve (Φ = 1000 sccm nitrogen)
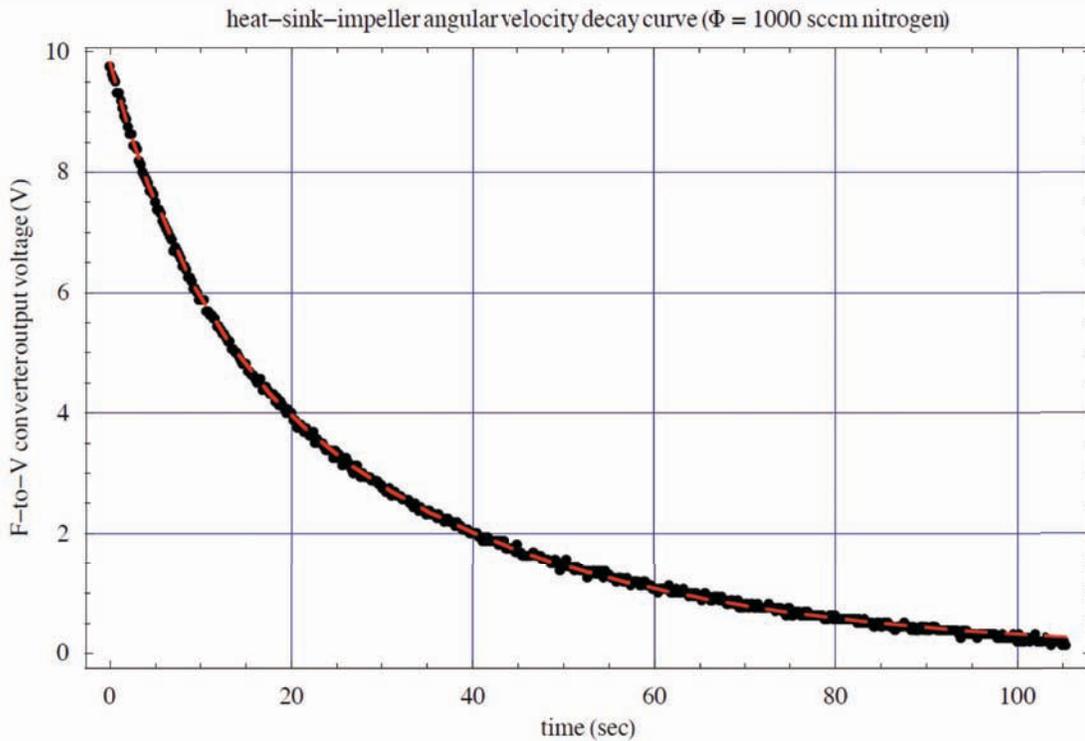

Figure 25: Heat-sink-impeller angular velocity decay curve used to determine heat-sink-impeller load torque load as a function of angular velocity. In such experiments, the externally pressurized air bearing is operated at high mass flow rate to achieve a very large value of h (gap distance). This improves measurement accuracy because it minimizes the extent to which the data must be corrected for frictional effects associated with shearing friction of the air gap region.
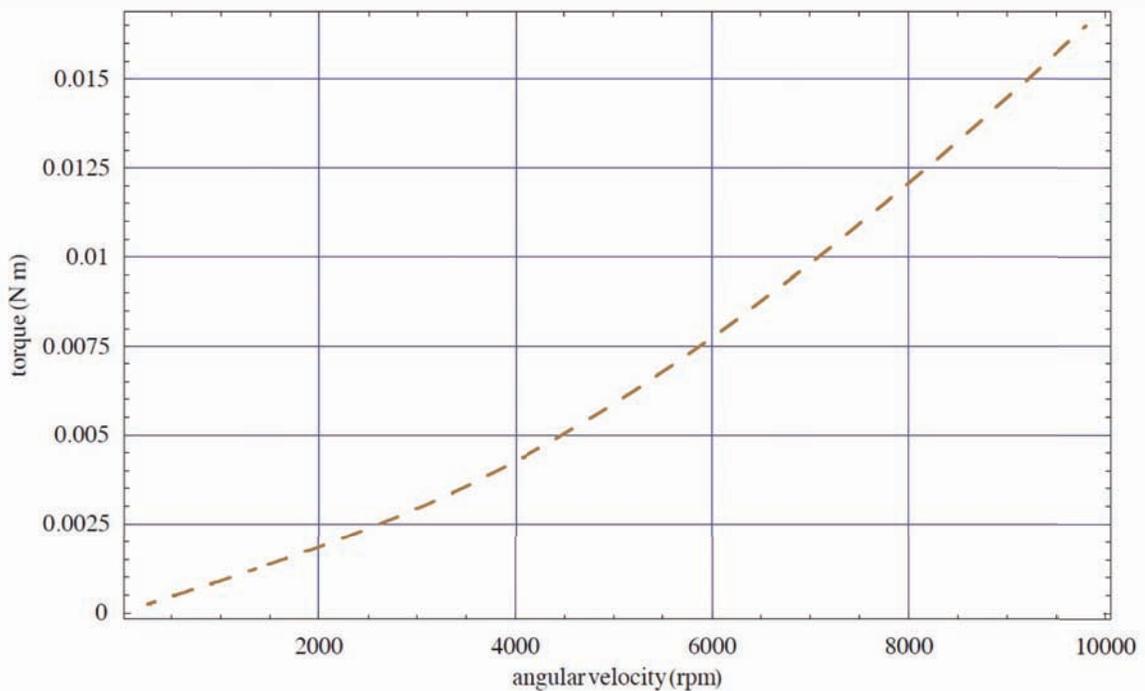

Figure 26: Torque-speed curve derived from angular velocity decay data. This shows the amount of torque that must be supplied by the brushless motor to rotate the impeller at a given rpm as a result of aerodynamic drag, shearing of the air gap region, bearing friction, etc.

Under normal operating conditions, the vast majority of the mechanical load driven by the brushless motor is associated with rotation of the heat-sink-impeller, the mechanical loading associated with other effects such as bearing friction is very small in comparison. The mechanical power consumption can in turn be broken down into two components, the mechanical power required to rotate the finned heat-sink-impeller (hsi) structure through the surrounding air (P_{hsi}), and the power associated with shearing of the gas-filled gap region (P_{sgr}):

$$P_{mechanical} = P_{hsi} + P_{sgr}$$

The viscous drag force that retards rotation of the heat-sink-impeller is in the tangential (ϕ) direction at all values of r. For an area element dA, this tangential force is:

$$dF = \mu \frac{du}{dz} dA = \frac{\mu \omega r}{h} dA$$

The corresponding torque about the axis of rotation is:
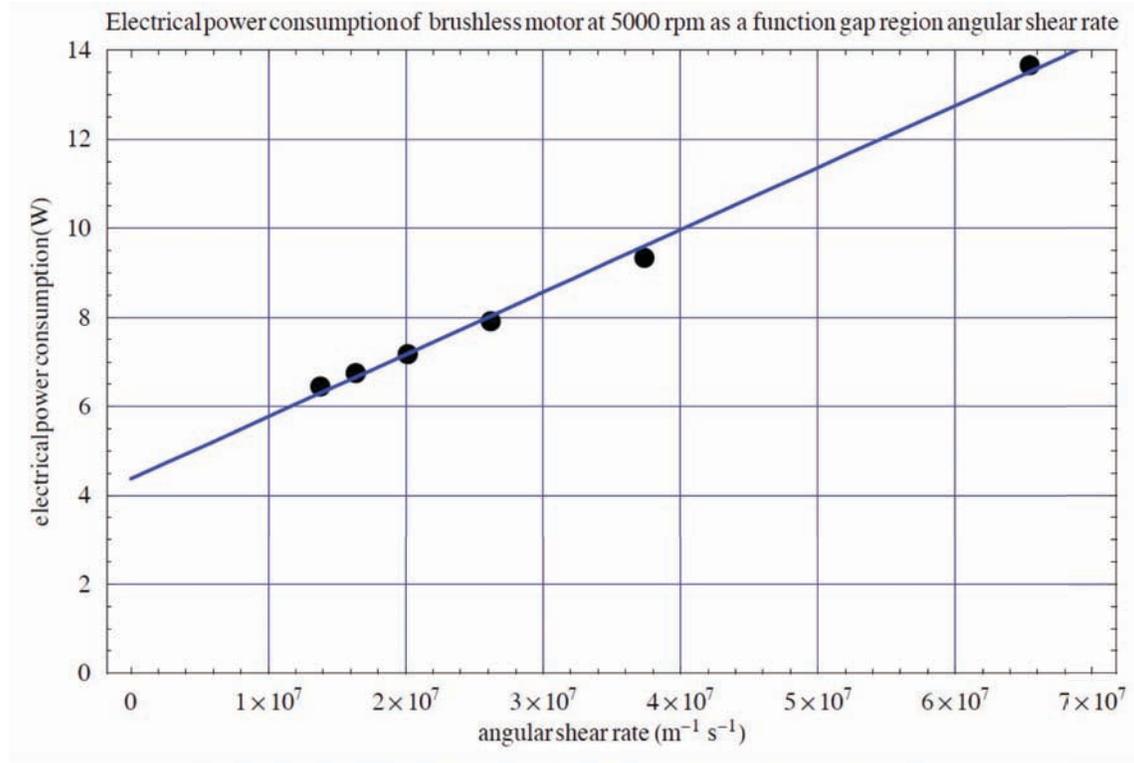
$$d\tau = \frac{\mu \omega r^2}{h} dA$$

Figure 27: Electrical power consumption as a function of shear rate at 5000 rpm.

The total torque associated with shearing of the gap region is therefore:

$$\tau = \int_{area} \frac{d\tau}{dA} dA = \int_0^{2\pi} \int_{R_{min}}^{R_{max}} \left(\frac{\mu \omega r^2}{h} \right) r \, dr \, d\phi = \frac{\pi \mu \omega}{2h} \left(R_{max}^4 - R_{min}^4 \right)$$

33

This very strong R^4 dependence stems from the fact that tangential velocity of a given area element is proportional to r, the torque generated by a tangential force,

$$\vec{\tau} = \vec{r} \times \vec{F}$$

has a first-order r dependence, and because the area of a wedge-shaped area element,

$$\int_{area} dA = \int r \, dr \, d\phi$$

has an r^2 dependence. As with the shear-rate-dependent thermal conductivity enhancement coefficient ε, the shearing torque can be expressed in terms of the independent variable S = ω/h:

$$\tau_{sgr}(S) = \frac{\pi \, \mu \left(R_{max}{}^4 - R_{min}{}^4 \right)}{2} S$$

This 4th order dependence on device diameter has significant implications. Earlier we mentioned that the reduction in air gap area associated with the construction of small-diameter devices would have the undesired effect of increasing the air gap thermal resistance, and that an important consideration is the extent to which we can compensate by operating at higher angular velocity, to take advantage of the conductivity enhancement factor, ε(h, ω). The above 4th order dependence indicates that all other things being equal, if we reduce the diameter of the device by a factor of 2, shearing losses will be reduced by a factor of 16. This gives us the freedom to increase ω and/or decrease h to compensate for the factor of 4 decrease in air gap area. In future work, prototype small diameter devices should be constructed to be compatible with very high angular velocities and reduced air gap distances.

The implication for larger-diameter devices is that limitations imposed by air gap shearing drag will necessitate reduced angular velocities. The resulting lower shear rates imply operation in the regime ε ≈ 1. Without a significant enhancement in air gap thermal conductivity due to shearing, it will be desirable to maintain relatively small gap distances to the fullest extent possible.

To validate our analytical model for the mechanical power consumption due associated with shearing of the air gap region (P_{sgr}), the data shown in Figure 27 were collected at constant angular velocity (5000 rpm) while varying the air gap distance. The resulting data set shows the electrical power consumption of the brushless motor as a function of shear rate with P_{hsi} held constant. The y-intercept of the best-fit line (4.63 W) represents the power consumption of the brushless motor associated with effects other than shearing of the air gap region, most notably, P_{hsi} at 5000 rpm, the power required to rotate the finned heat-sink-impeller structure through the surrounding air at an angular velocity of 5000 rpm (and to a much lesser extent, the power required to spin the rotor assembly of the brushless motor). The slope of the best-fit line (1.40 x 10^{-7} W m s) represents the power consumption associated with shearing of the air gap region. The theoretically calculated value of the slope,

$$\frac{dP}{dS} = \frac{\pi \, \mu \left(R_{max}{}^4 - R_{min}{}^4 \right) \omega}{2}$$

assuming a brushless motor with 100% electrical-to-mechanical efficiency, is 9.94 x 10^{-8} W

m s. For the most part the 40% disparity between these two numbers is attributable to the fact that, as will be discussed later, the brushless motor efficiency is in reality of order 60%. Because the physical basis of the analytical viscous shearing model is entirely reasonable, and because it agrees well with experimental measurements, from this point on we will therefore use the analytical expression derived above to calculate the mechanical torque (power consumption) associated with shearing of the air gap region. For the prototype system this analytical expression evaluates to:

$$P_{sgi}(h, \omega) = \frac{k_2' \, \omega^2}{h}$$

$k_2' = 1.88 \times 10^{-10}$ m W s^{-2}

This in turn allows us to determine experimentally the torque (power consumption) associated with rotation of heat-sink-impeller using the angular velocity decay technique. To improve measurement accuracy, such experiments are conducted at large values of h (~50 µm) in order to make the required correction for $P_{sgi}(h, \omega)$ as small as possible. The results of such an angular velocity decay experiment are shown in Figure 28.

$$P_{hsi} = P_{mechanical} - P_{sgi}(h, \omega)$$

This dashed curve shown in Figure 28 is our experimentally derived mathematical model for $P_{hsi}(\omega)$:
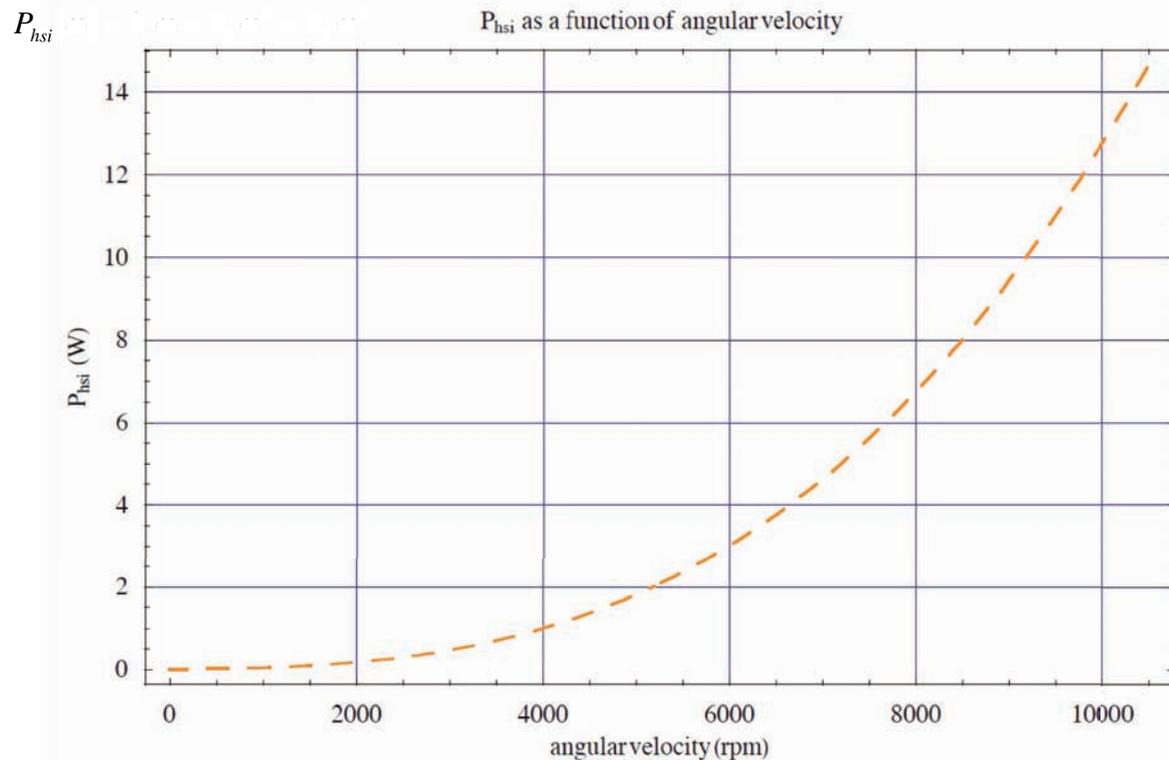

Figure 28: Electrical power consumption associated with rotation of heat-sink-impeller as a function of angular velocity (the effects of air gap shearing friction has been subtracted out).

$k_1 = 8.59 \times 10^{-5}$ W s^{-1}
$k_2 = 1.33 \times 10^{-6}$ W s^{-2}

$k_3 = 9.74 \text{ x } 10^{-9} \text{ W s}^{-3}$

Finally, these two models for air gap shearing and heat-sink-impeller mechanical power consumption are combined to obtain a system model of the form:

$$P_{mechanical}(h, \omega) = k_1 \omega + \left(k_2 + \frac{k_2'}{h} \right) \omega^2 + k_3 \omega^3$$

With regard to physical interpretation of the above model, a priori expectations are roughly that:

1) the 1st order term to be associated with mechanical (Coulombic, dry) friction, whose torque (power) is typically modeled as a 0^{th} order (1^{st} order) function of angular velocity,
2) the 2nd order term to be associated with shearing of the air gap region, whose torque (power) is modeled as a 1^{st} order (2^{nd} order) of angular velocity (as derived earlier), and
3) the 3^{rd} order term to be associated with heat-sink-impeller aerodynamic drag, whose torque (power) is modeled as a 2^{nd} order (3^{rd} order) function angular velocity because of the ω^2 dependence for stagnation pressure.
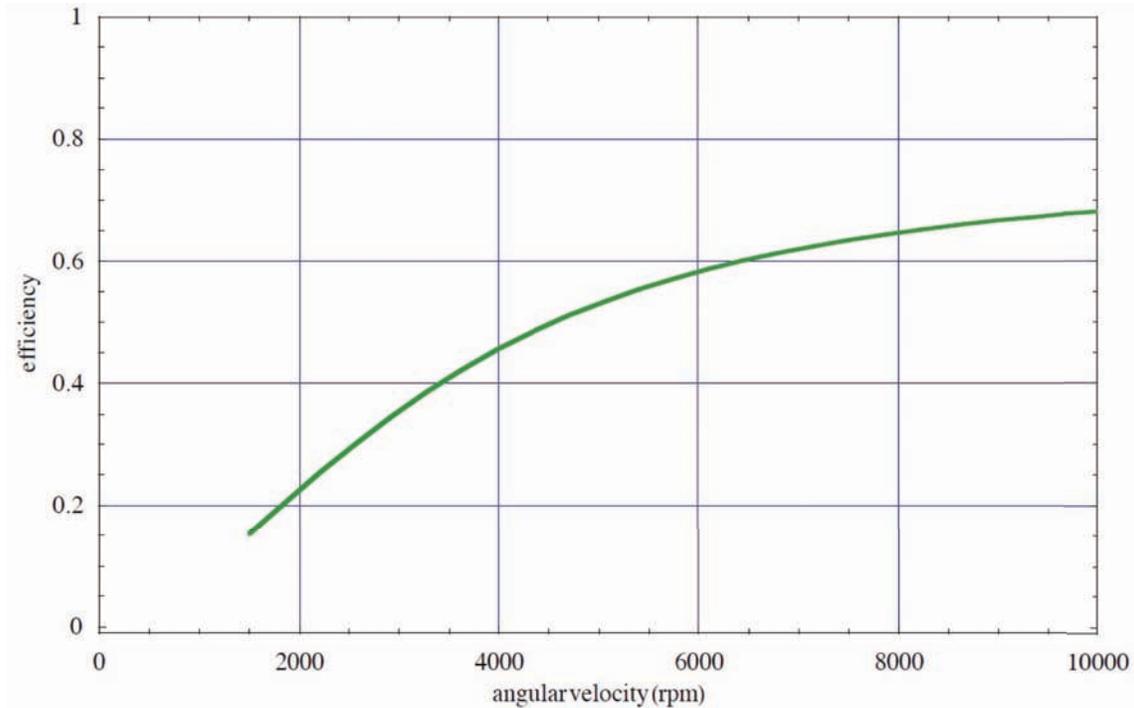

Figure 29: One of several measurements conducted to determine brushless motor efficiency. In this case the air gap distance was 40 μm and the brushless motor power supply voltage was 8.2 V.

For typical values of h (~2 x 10^{-5} m) the coefficient of the ω^2 term in the above expression for $P_{mechanical}(h, \omega)$ is dominated by k_2'/h rather than k_2. It is likely that the principal source of the ω^2 term in the equation for $P_{hsi}(\omega)$ is the effect of eddy current loss during rotational decay. Although the brushless motor leads are disconnected during recording of the angular velocity decay, from Faraday's law, relative motion between the permanent magnet rotor and the high

magnetic susceptibility stator does result in a retarding torque due to eddy current that is proportional to ω. Thus the power dissipation associated with eddy current generation has an $ω^2$ dependence.

The above system model for mechanical power consumption incorporating Coulombic friction, air gap shearing friction, eddy current losses, and aerodynamic drag allows us to determine the electrical-to-mechanical efficiency curve of the brushless motor from experimental measurements (of h, ω, and $P_{electrical}$):

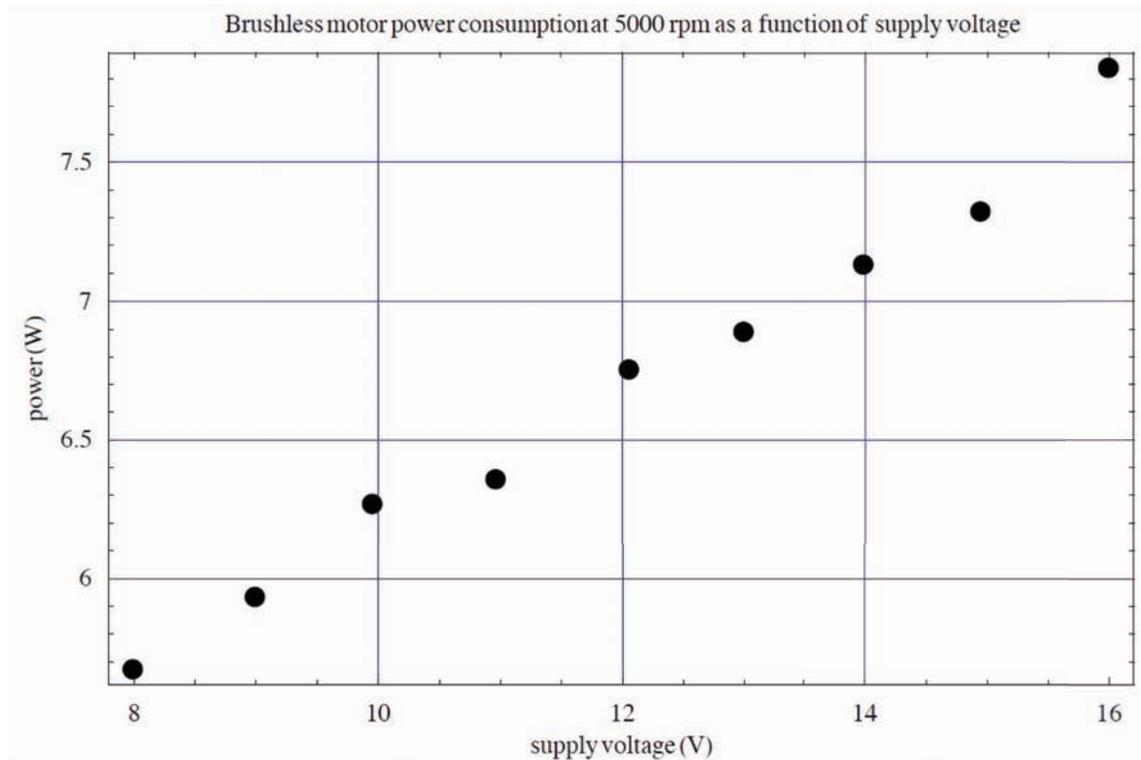$$\varepsilon(h, \omega) = \frac{P_{mechanical}(h, \omega)}{P_{electrical}(h, \omega)}$$

Figure 30: In this experiment the brushless motor power supply voltage was varied from 8 to 16 volts and the duty cycle of the brushless motor controller was adjusted to maintain an angular velocity of 5000 rpm. As predicted, a significant reduction in electrical power consumption at higher duty cycle and lower supply voltage is observed.

This is extremely important from the standpoint of brushless motor optimization (load matching). Figure 29 shows one such brushless motor efficiency curve recorded for the version 1 prototype system.

Such a brushless motor efficiency curve reflects the complex interaction of the load (torque-speed curve) and the motor (size, number of poles, supply voltage, etc.). The results shown in Figure 29 are typical: zero efficiency in the limit ω = 0, and a relatively broad peak at which the brushless motor and load are most well matched. Eventually the motor efficiency rolls over at higher speeds, primarily due to increased resistive losses associated with heating of the stator winding, heating of the high current MOSFET transistor in the output stage of the ESC, and magnetic saturation of the stator poles. For a load with a given torque speed

37

curve, the maximum brushless motor efficiency and angular velocity at which maximum efficiency is realized depends a great deal on the stator winding configuration. For a given application, once the desired operating rpm is established the stator winding configuration can be empirically, or semi-empirically optimized. Factors that affect what combination of wire gauge and number of windings will ultimately turn out to be best suited to the application include:

1) motor size (principally volume)
2) load torque-speed curve (heat-sink-impeller geometry and air gap distance)
3) power supply voltage
4) operating temperature
5) provisions for cooling stator windings

Having said that, with exception motor-load combinations in which the torque speed curve of the load has resonances or other unusual features, the peak of the brushless motor efficiency curve is broad enough to obviate exacting optimization. Assuming such optimization is carried out diligently, the maximum attainable efficiencies for modern brushless motors can >95% for relatively large motors, and are of order 70 to 80% for motors comparable in size to the GBx Mini motor used in the current prototype system. This dependence of maximum efficiency on motor volume is one example of a more general rule for magnetic machinery/circuits. In an inductor for example, the ratio of energy stored in the magnetic field to power dissipated because of winding resistance losses is:

$$\frac{\varepsilon_{magnetic}}{P_{resistive}} = \frac{\frac{1}{2} L I^2}{I^2 R} = \frac{L}{2 R}$$

For a simple geometry such as a solenoid of radius r, it is readily apparent that winding resistance scales as r while the inductance scales as r^2. The same scaling is valid for more complex geometries. A motor is a device that converts electrical power to a magnetic field through which mechanical work is transmitted. The transmission loss associated with winding resistance also obeys the same scaling law. Thus in the design of future prototype devices, the amount of volume allocated for the brushless motor should be considered carefully if low electrical power consumption is a high priority.

Another high priority is to configure the system (winding, supply voltage, or both) so that the brushless motor controller operates at 100% duty cycle. Modeling the system as an ideal dc motor in series with the winding resistance, we can see that at a given angular velocity, the voltage across the ideal motor (im) driving a mechanical load is fixed:

$$V_{im} = \frac{\omega}{K_v}$$

In the above equation, K_v is generator constant of the motor. The mechanical power delivered by the motor to the load (P_{load}) is fixed as well. The time averaged average current flowing through motor must therefore be:

$$\langle I \rangle = \frac{P_{load}}{V_{im}} = \frac{P_{load} K_v}{\omega}$$

This implies that the on-state current of the PWM waveform (I_{on}) used to drive the motor must be:

$$I_{on} = \frac{P_{load}\,K_v}{f\,\omega}$$

where f is the duty cycle of the brushless motor controller output waveform. The power lost to the parasitic resistance of the motor winding and electronic speed controller is:

$$P_{loss} = \left\langle I^2 \right\rangle (R_{wnd} + R_{esc}) = f\,I_{on}^{\ 2}\,(R_{wnd} + R_{esc}) = f\left(\frac{P_{load}\,K_v}{f\,\omega}\right)^2 (R_{wnd} + R_{esc}) = \frac{R_{wnd} + R_{esc}}{f}\left(\frac{P_{load}\,K_v}{\omega}\right)^2$$

Power loss is therefore minimized by operating the electronic speed controller at 100% duty cycle. Thus, for a given motor-load combination we should set the duty cycle of the electronic speed controller to 100% and increase the supply voltage until we obtain the desired operating rpm. This does not imply that the motor winding configuration is optimized for the load in question, but rather, that we are running the motor, such as it is, at the highest efficiency possible.

Thus, while PWM does provide a way to vary the speed of the motor, in the presence on significant parasitic resistance, this comes at a price. The efficiency of the system as a whole is:

$$\varepsilon = \frac{P_{load}}{P_{load} + \left\langle P_{loss} \right\rangle} = \frac{P_{load}}{P_{load} + \dfrac{R_{wnd} + R_{esc}}{f}\left(\dfrac{P_{load}\,K_v}{\omega}\right)^2} = \frac{1}{1 + \dfrac{(R_{wnd} + R_{esc})P_{load}\,K_v^{\ 2}}{f\,\omega^2}}$$

In situations where we have the freedom to tailor the stator winding configuration to the load, to achieve a given angular velocity, you're better off using a small value of K_v (a motor with a larger number of windings, which has a lower rpm/V rating) and as large a duty cycle as possible (ideally, 100%). A larger number of windings means a proportional increase in winding resistance (R_{wnd}) assuming the wire gauge remains unchanged. This partially cancels $K_v^{\ 2}$ effect. But typically the volume of windings is limited by the stator pole spacing. A larger number of winding turns means thinner gauge wire, which further increases the winding resistance. Eventually this effect overtakes the benefit of reducing K_v; for a given stator geometry, operating rpm, and load torque, there is an optimum wire diameter (number of windings). But for a given motor, where K_v is a constant, it is always advisable to adjust the power supply voltage to allow operation at 100% duty cycle.

## 6. Further characterization

Here we discuss other aspects of exploratory work conducted in this study that merit discussion. The first of these pertains to potential concerns about vibration and stability that arise in conjunction with rotating heat sink impeller at several thousand rpm. Typical results
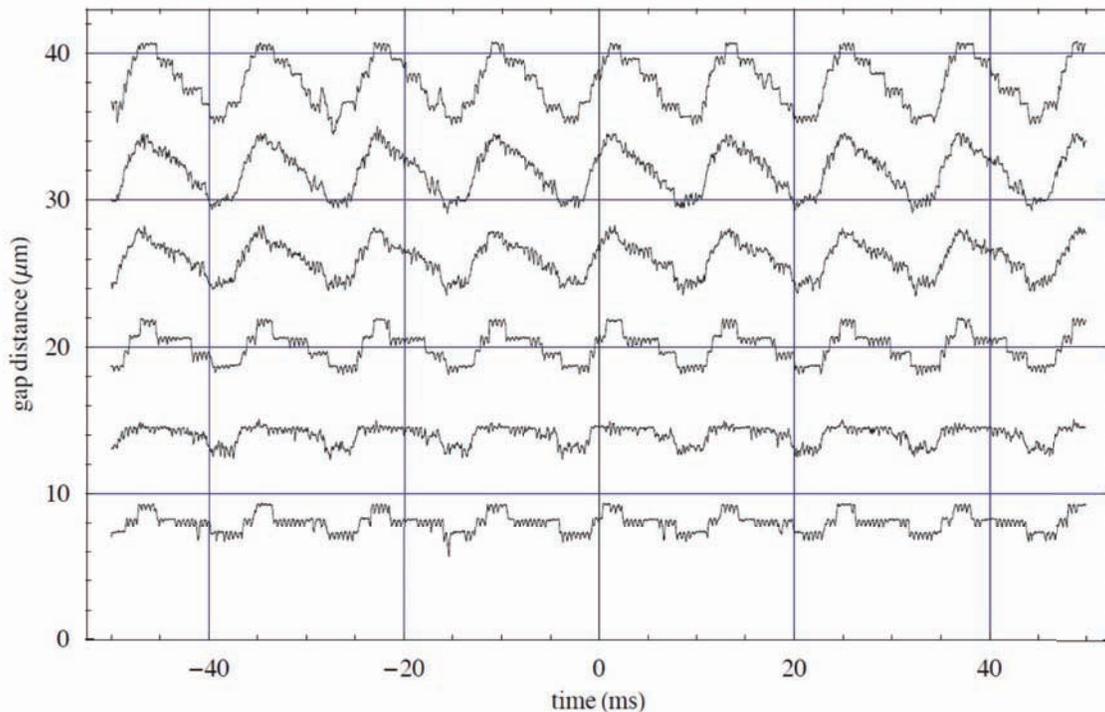
Figure 31: One of several measurements conducted to assess vibration/precession of the rotating heat-sink-impeller. This set of measurements was conducted at an angular velocity of 5000 rpm.

from experimental tests conducted on the version 1 prototype system are shown in Figure 31. The main conclusions that were drawn from these test were (1) that observed vibrations were largely driven by less than perfect concentricity between the axis of the brushless motor rotor/stator assembly and that of the heat-sink-impeller, (2) to the extent that such vibration/precession was induced by motor/impeller misalignment, no mechanical resonances or other erratic behavior was observed (testing was conducted to a maximum speed of 12.700 rpm), (3) as expected from theory, with decreasing gap distance vibration amplitude is attenuated because of the increasing stiffness of the air bearing suspension, (4) despite the relatively crude scheme used to align the heat-sink-impeller and the brushless motor external rotor housing for assembly (prior to the application of UV curing epoxy), under typical operating conditions (gap distances of order 20 μm), observed vibrations were limited to an amplitude of ~2 μm rms. Having said that, in subsequent work provisions should be made to achieve better motor/impeller concentricity.

Among the purported advantages of the air-bearing cooler with respect to conventional air-cooler heat exchangers are low noise operation and immunity to fouling. It was hypothesized that in addition to low power consumption, low noise operation would be an additional advantage of the "direct drive" scheme used to create relative motion between the finned heat sink and the surrounding air. No attempt was made to optimize the heat-sink-impeller geometry from the standpoint of noise generation other than attempting to select a fin geometry that appeared to be conducive to smooth separation and rejoining of the heat-sink-impeller-fin flow field. No quantitative acoustic measurements were conducted due to limited resources, either. Nonetheless, it was obvious during initial testing that the broadband acoustic spectrum generated by conventional small high-speed fans was largely absent in the version 1 prototype device. Similarly, despite the lack of systematic measurements to

40

Figure 32: Open impeller configuration (left) and closed impeller configuration (right).
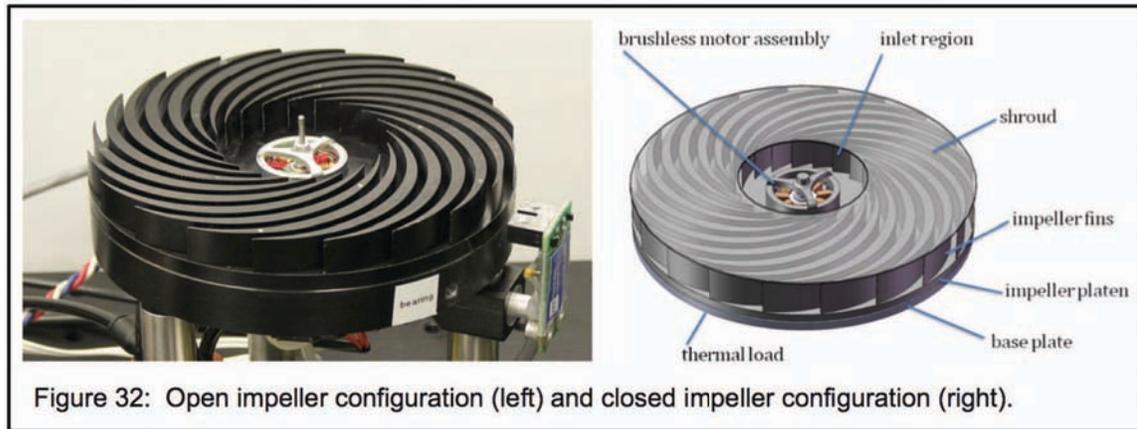
characterize particulate fouling, it became obvious over time that fouling of heat-sink-impeller heat exchange surfaces was for all intents and purposes zero. While neither of these two observations are surprising, they are nonetheless noteworthy because in most real-world applications, low-noise operation and immunity to heat sink fouling are critically important advantages over the prior art.

Although the vast majority of work conducted in this study pertains to the open-impeller configuration shown on the left hand side of Figure 32, a set of experiments was also carried out to evaluate the merit of a closed-impeller configuration such as that shown on the right-hand-side of Figure X. The main conclusions of this study were:

1) Power consumption as a function of rpm is considerably higher for the closed impeller.
2) The closed impeller generated higher noise levels.
3) Device thermal resistance (at a given rpm) was unchanged by the addition of the shroud.

At this stage, it therefore appears that the open impeller design provides better all around performance. This topic should nonetheless be revisited in future work, where we anticipate the development of powerful computational fluid dynamic analysis tools, as should theoretical considerations governing the use of open vs. closed impeller configurations in other real-world applications. Among the obvious disadvantages of a closed impeller design would be increased difficulty of fabrication and/or assembly. For example, the open impeller configuration leaves open the possibility of single-piece die casting or stamping, while the closed impeller configuration would likely require dip brazing to fuse together two separately fabricated components. A potential advantage of the closed impeller design is increased mechanical stiffness, although the relative importance of this consideration is not yet known. An important component of future work will be a basic stress-strain analysis of the rotating impeller and an understanding of the relevant scaling relationships (e.g. the effect of fin height on platen deformation).

Other relatively unremarkable but nonetheless useful results were obtained from experiments to determine the optimum firing phase angle for the brushless motor controller, and the effect of pulse-width modulation frequency on motor operation. Selection of the Motortron MDrive 9-A brushless motor controller was motivated in part by its programming flexibility, which allows the rotor phase firing angle to be adjusted in 1° increments from 0 to 25°. The

results of experiments to determine optimum phase angle, and sensitivity to phase angle, are shown in Figure 33. These data indicate that the motor-load combination of the prototype air bearing cooler exhibits little sensitivity to firing phase angle, and the that the optimum phase angle is approximately 10°. With regard to the effect of PWM frequency, the Motortron ESC has only two settings, 8 KHz and 16 kHz. Nonetheless, recording of data at both PWM frequencies was important in that it showed that operation 16 KHz provided an improvement in motor efficiency, significant reduction in audible electro-acoustic noise associated with magnetostriction of the brushless motor stator poles.
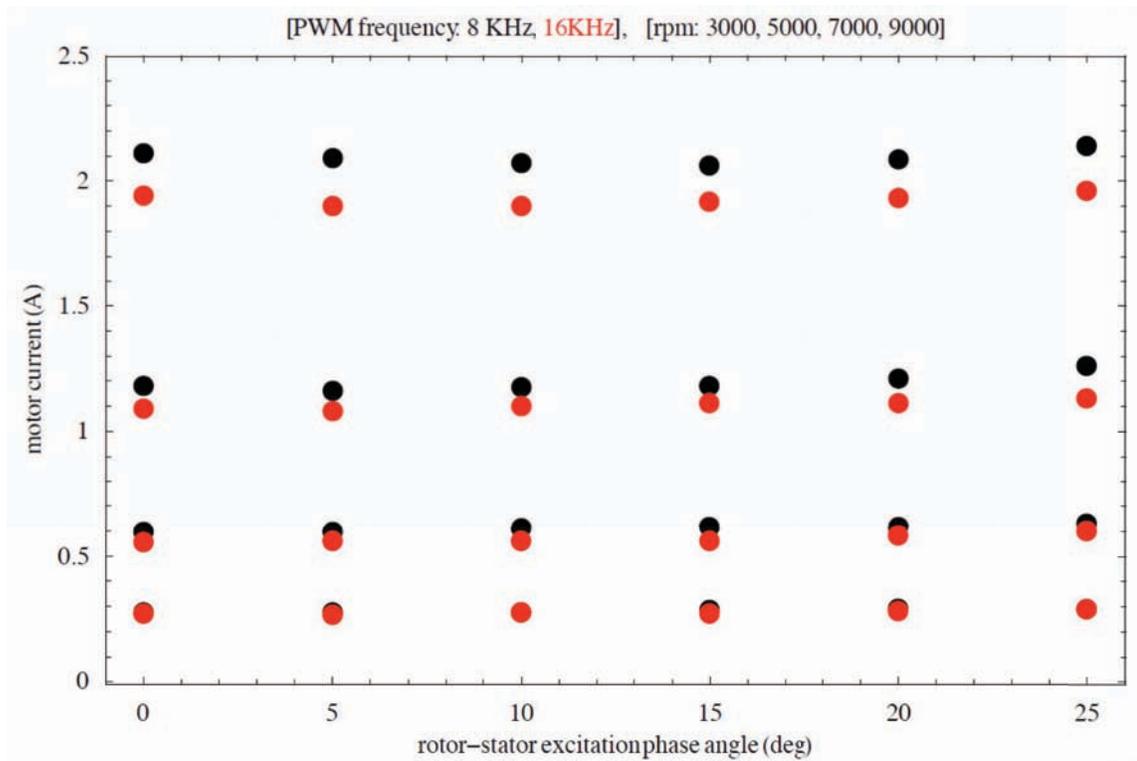

Figure 33: Motor current required to maintain angular velocities of 3000, 5000, 7000, and 9000 rpm as a function of phase angle at fixed supply voltage. The points in black (red) correspond to a pulse width modulation frequency of 8 KHz (16 KHz).

7. System model

In additional to proving the technical viability of the proposed air-cooled heat exchanger architecture, a central goal of the present study was the development of a system model that can provide quantitative insight in the following areas:

- s f ▨▨▨▨ ▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨▨▨s ▨▨▨▨▨) ▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨ ▨f ▨▨▨▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨) ωfa▨

- uf ▨▨▨▨ ▨ ▨▨▨ ▨▨▨ ▨▨▨ ▨▨ ▨▨▨▨▨▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨s ▨▨▨▨▨▨) ▨▨▨ ▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨▨ ▨▨▨▨ ▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨▨▨▨▨▨▨ ω▨▨▨▨▨▨▨▨) ▨▨▨▨▨ ▨h▨▨▨▨▨▨▨▨▨▨▨▨) ▨▨▨▨ ▨h▨▨▨▨▨▨▨▨▨▨▨▨▨b ▨▨▨▨▨▨▨▨ ▨▨▨▨▨▨▨▨▨▨▨▨▨ ▨▨fa▨

3) *scaling laws and engineering tradeoffs:* scaling laws for thermal resistance and electrical power consumption as a function of various design parameters to the fullest extent possible in the absence of a comprehensive fluid dynamic model.

The following table summarizes the various components of the system model developed in this study:

Physical process	System model	Basis
Thermal resistance for transfer of heat from rotating heat-sink-impeller to surrounding air.	$R_{xfer}(\omega) = \dfrac{k}{\omega^{0.8}}$	Empirical model based on experimental measurements and well known fluid dynamic scaling relationship.
Thermal resistance for transfer of heat across air-gap region.	$R_{gap}(h,\omega) = \dfrac{h}{A\,\kappa\left[1 + k_1\frac{\omega}{h} + k_2\left(\frac{\omega}{h}\right)^2 + k_3\left(\frac{\omega}{h}\right)^3\right]}$	Empirical model based on experimental measurements.
Thermal resistance for internal heat transfer, base plate.	$R_{baseplate} = 10.4\ mK/W$	Calculated.
Thermal resistance for internal heat transfer, heat-sink-impeller platen.	$R_{platen} = 3.5\ mK/W$	Calculated.
Thermal resistance for internal heat transfer, heat-sink-impeller fin array.	$R_{fins} = 9.4\ mK/W$	Calculated.
Thermal resistance for external heat transfer, parallel leakage.	$R_{leakage} = 2060\ mK/W$	Measured.
Power consumption (mechanical) associated with shearing of air in gap region.	$P_{sgr}(h,\omega) = \dfrac{\pi\,\mu\left(R_{max}^{\ 4} - R_{min}^{\ 4}\right)\omega}{2\,h}$	Analytical model validated by experimental measurements.
Power consumption (mechanical) associated with rotation of heat sink impeller and brushless motor rotor assembly.	$P_{hsi}(h,\omega) = k_1\,\omega + k_2\,\omega^2 + k_3\,\omega^3$	Empirical model based on experimental measurements.

Such a comprehensive system model is extremely powerful from the standpoint of understanding the complex relationships that govern device performance. Figure 34, for example, uses the above system model to illustrate how thermal resistance is broken down between the various heat transfer processes as a function of available electrical power, assuming that h and ω are globally optimized and that the brushless motor operates at an electrical-to-mechanical efficiency of 70%. The later efficiency figure is very realistic for a brushless motor that is roughly optimized to a given set of operating conditions. Such a plot is extremely informative in its ability to show realistically attainable device performance as a function of available electrical power, and how the relative importance of various heat transfer processes change (or do not change) as function of operating conditions. Note that the individual thermal resistance contributions (dotted lines) do not add up exactly to R_{total} because of the effect of the parallel leakage term $R_{leakage}$. The most obvious features of such a plot are that is the bottleneck to heat transfer is R_{xfer}, and that R_{gap} is contributes very little to R_{total}. Such a plot also indicates that there some improvement in device performance will be realized when the thickness of the base plate and platen are reduced in subsequent

prototypes. The decrease in R_{gap} as a function of available electrical power is observed because increasing angular velocity results in increasing enhancement of convective heat transfer in the air gap region. Among the obvious conclusions to be drawn from Figure 34 is that subsequent work should focus on the heat transfer properties of the rotating heat-sink-impeller (R_{xfer}). This what motivated earlier recommendations for the implementation of simplified test bed dedicated to quick and accurate measurements of R_{xfer} using "double-sided" heat-sink-impeller prototypes, and an aggressive modeling CFD modeling campaign directed towards rational design/optimization of heat-sink-impeller geometry.

Figure 34: Total thermal resistance (solid line) and series thermal resistance components (dashed lines) for air bearing heat exchanger version 1 prototype as a function of available brushless motor electrical power (assuming a brushless motor efficiency of 70%).
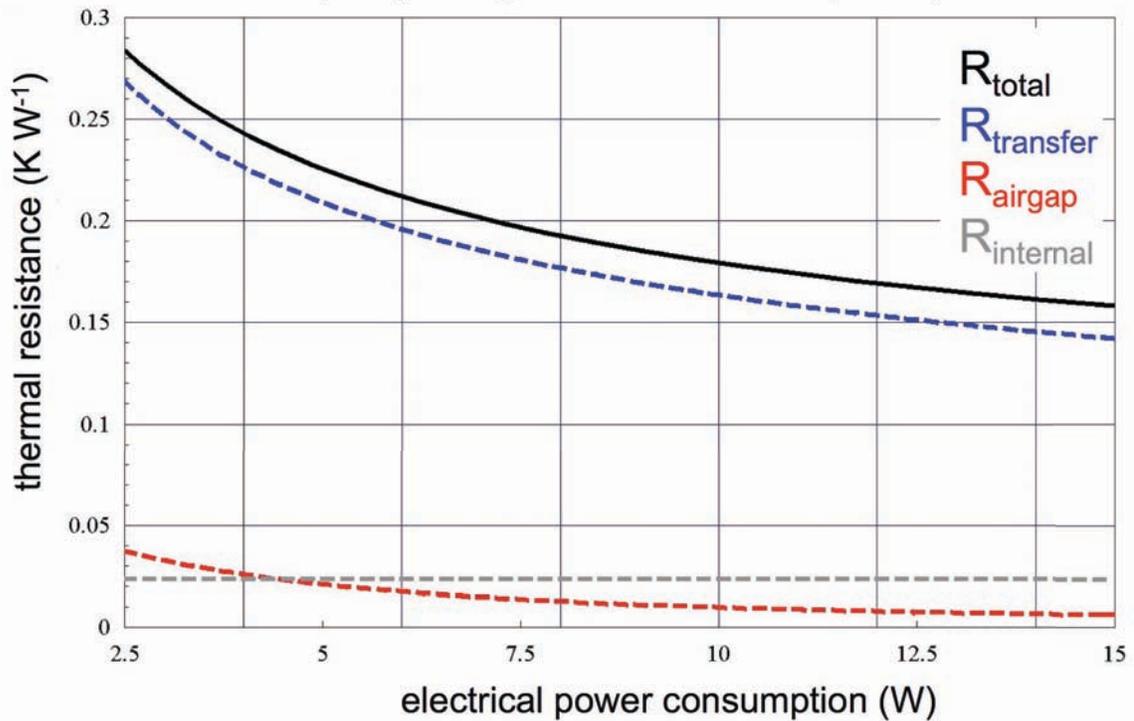
Figures 35 illustrates other informative plots constructed using the device system model that can be used to inform choice of operating conditions and better understand performance tradeoffs.

8. Conclusion

The FY09 Tier 1 LDRD and Sandia Royalty funds invested in this proof-of-concept study have resulted in a major breakthrough in air-cooling technology. The arguments put forth in the original proposal in support of the boundary layer thinning effect, negligible air gap thermal resistance, low electrical power consumption, low-noise operation, and immunity to fouling have all been shown to be correct. The performance obtained with a highly unoptimized version 1 prototype device already represents a major advance in a technology area of fundamental importance that has changed little in the past 40 years. The potential implications in the U.S. energy sector (air conditioners, heat pumps, and refrigeration
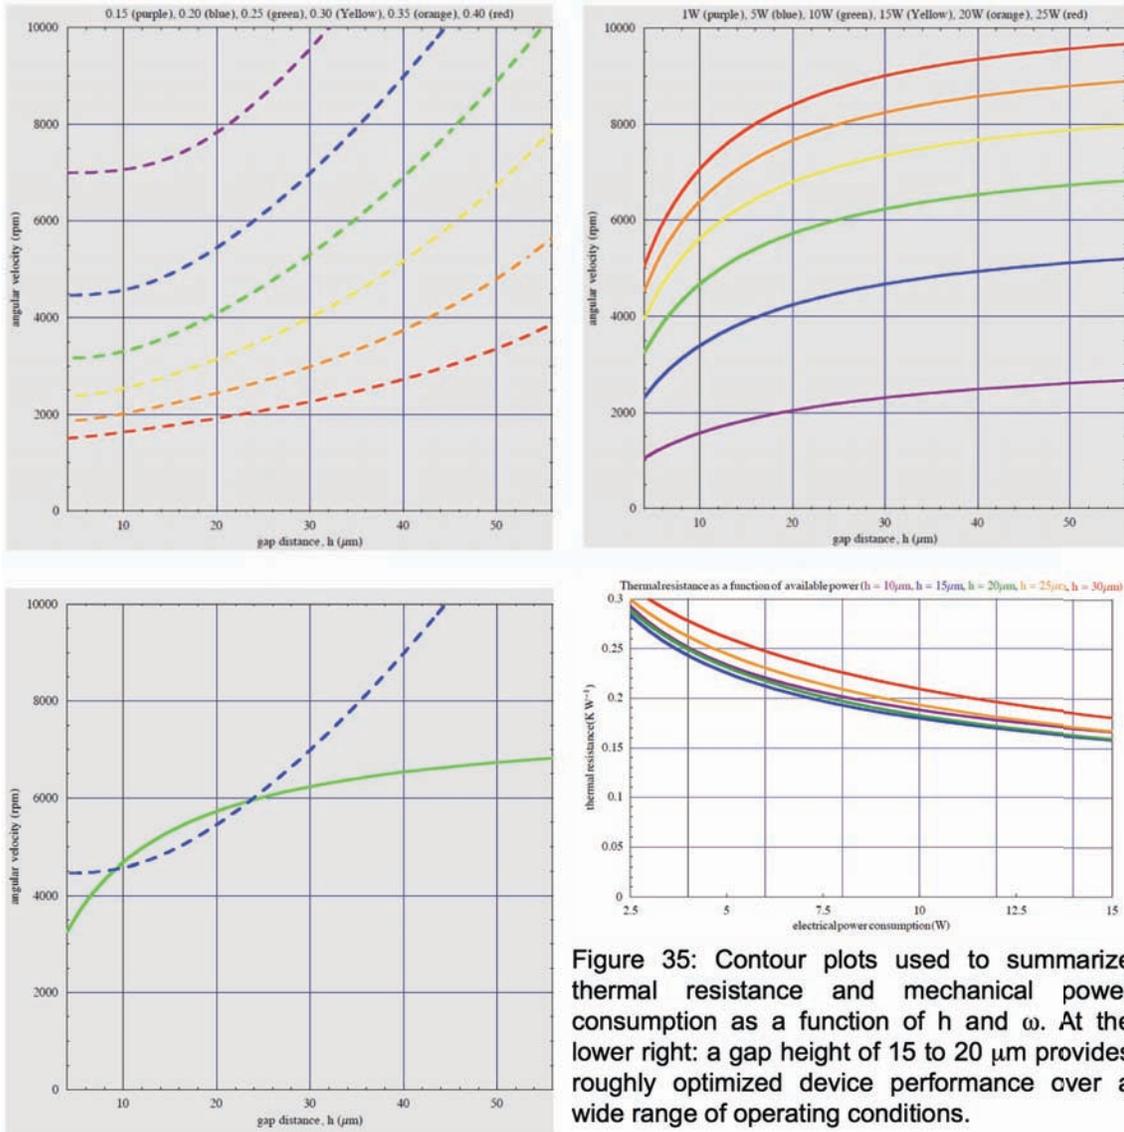
Figure 35: Contour plots used to summarize thermal resistance and mechanical power consumption as a function of h and ω. At the lower right: a gap height of 15 to 20 μm provides roughly optimized device performance over a wide range of operating conditions.

equipment) amount to a ~5% reduction in electrical power consumption, significantly increased grid operating margin, and significant reduction in heat-wave generated load spikes. The potential implications in the information technology sector (desktop computers, high-performance graphics cards, server farms, and data centers) are also very large and center on resolving the thermal brick wall problem, which has prevented CPUs from advancing beyond clock speeds of ~3 GHz, and emerging concerns about the energy consumption of data centers, half of which is associated with cooling.

The most immediate priority for future work is construction of the version 2 prototype, which is predicted to reduce thermal resistance to ~0.1 C/W. Other high priorities include:

sf ▯▯ ▯▯▯▯ ▯▯▯▯▯▯▯▯▯▯▯▯▯▯) ▯▯▯▯▯) ▯▯▯ ▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯a▯

2) ▯▯▯▯▯▯▯▯ ▯▯▯▯ ▯▯▯ ▯▯▯▯ ▯▯▯▯ ▯▯▯) ▯▯▯ ▯▯▯▯▯▯▯▯▯) ▯ ▯▯▯▯▯▯▯▯▯ ▯▯▯▯ ▯▯▯▯▯▯▯▯
    ▯▯▯▯▯e▯▯▯ ▯n▯▯▯▯▯▯▯▯▯▯▯▯ ▯▯▯▯i ▯ ▯▯▯▯▯▯▯▯▯▯▯▯a▯

3) ▯▯▯▯▯▯▯▯▯▯▯▯▯n▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯▯ ▯▯▯▯▯▯▯ ▯▯▯▯▯▯▯ ▯▯▯i ▯ ▯▯▯▯▯

thermal resistance (to facilitate rapid and accurate evaluation of candidate geometries),

4) internal heat flow analysis directed towards reducing the thermal resistance of the base plate and platen while maintaining heat spreading performance,

5) internal mechanical stress analysis directed towards reducing the thickness of the base plate and platen without adversely effecting mechanical rigidity,

6) far more extensive work on optimization of brushless motor efficiency.

It should be emphasized that the budget and time constraints imposed upon this project thus far imply that a great deal of further technical progress will likely result from additional support. The opportunity cost of not having capabilities in fluid dynamic modeling and iterative testing of multiple prototypes has been very large. Limited resources have necessitated an extremely conservative approach to device development, and at this point it has become clear that a more aggressive approach to developing this breakthrough technology is warranted.

Appendix:

A. Numerical values of device parameters

Surface area of heat-sink-impeller: 4.83×10^{-2} m^2 (all surfaces including bottom of platen)
Volume of heat-sink-impeller: 6.93×10^{-5} m^3
Outer diameter of heat-sink-impeller: 0.1016 m (4.000")
Inner diameter of fin array: 0.0381 m (1.500")
Inner diameter of rotor housing recess: 0.0203 m (0.800")
Height of heat-sink-impeller platen: 4.76×10^{-3} m (3/16")
Height of heat-sink-impeller fin array: 0.0101 m (0.40")
Height of heat-sink-impeller, total: 0.0149 m (0.5875")
Area of annular air gap region: 7.78×10^{-3} m^2 (3/16")
Width of heat-sink-impeller air channels: 2.38×10^{-3} m (3/32")
Mass of heat-sink-impeller: 0.194 kg
Mass of base plate: 0.331 kg
Mass of brushless motor rotor: 5.5×10^{-3} kg
Mass of brushless motor stator: 6.6×10^{-3} kg (includes windings)
Mass of impeller shroud: 0.035 kg
Moment of inertia of heat-sink-impeller: 2.74×10^{-4} kg m^2
Moment of inertia of impeller shroud: 0.52×10^{-4} kg m^2
Thermal resistance of base plate: 10.4 mK/W
Thermal resistance of platen: 3.54 mK/W
Thermal resistance of fin array (top-to-bottom): 18.8 mK/W
Thermal resistance of fin array (z-averaged): 9.4 mK/W
Parallel leakage thermal resistance: 2060 mK/W
Heat capacity of base plate: 317 J K^{-1}
Heat capacity of heat-sink-impeller: 186 J K^{-1}
Number of stator phases: 3
Number of stator poles: 9

Number of permanent magnet rotor poles: 12
Stator winding diameter: 26 AWG
Stator winding length: 1.53 m per phase
Number of stator windings per stator pole: 13
Number of stator windings per phase: 39
Number of stator windings per motor: 117
Stator winding resistance: 2.05 Ω per phase
Stator winding inductance: 15.0 µH per phase
Stator winding L/R time constant: 7.3 µs

B. Numerical values of physical constants

Density of alloy 7075 aluminum: 2.80×10^{-3} kg m^{-3}
Thermal conductivity of alloy 7075 aluminum: 173 W m^{-1} K^{-1}
Specific heat of alloy 7075 aluminum: 960 J kg^{-1} K^{-1}
Thermal conductivity of air: 2.62×10^{-2} W m^{-1} K^{-1}
Thermal conductivity of nitrogen: 2.60×10^{-2} W m^{-1} K^{-1}
Thermal conductivity of helium: 1.57×10^{-1} W m^{-1} K^{-1}
Dynamic/absolute viscosity of air: 1.83×10^{-5} N s m^{-2}
Dynamic/absolute viscosity of nitrogen: 1.78×10^{-5} N s m^{-2}
Dynamic/absolute viscosity of helium: 1.90×10^{-5} N s m^{-2}

C. Bibliography and References Cited

Bleier, F. P., *Fan Handbook, Selection, Application and Design*, McGraw-Hill, New York, 1997.

Chiu, S. A., Zaloudek, F. R., *R&D Opportunities for Commercial HVAC Equipment*, Pacific National Laboratory, DOE report under contract DE-AC06-76RLO 1830, March 1987.

Cobb, E. C., Saunders, O. A., *Heat Transfer from a rotating disk*, Proceedings of the Royal Society of London. Series A, Mathematical and Physical Science, 1956.

Dixon, S. L., *Fluid Mechanics and Thermodynamics of Turbomachinery*, 4th edition, Elsevier Science, Boston, MA, 1998.

Fox, R. W., McDonald, A. T., *Introduction to Fluid Dynamics*, 4th edition, John Wiley & Sons, New York, 1992.

Incropera F. P., Dewitt D. P., Bergman, T. L. and Lavine, A. S., *Fundamentals of Heat and Mass Transfer*, 6th Edition, John Wiley & Sons, New York, 2007.

Jorgensen, R., Fan Engineering, 7th edition, Buffalo Forge Company, Buffalo, NY, 1970.

Kim, N. S., Austin, T., Blaauw, D. Mudge, T., Flautner, K., Hu, J. S., Irwin, M. J., Narayanan, M. K. V., "Leakage Current: Moore's Law Meets Static Power," Computer, vol. 36, no. 12, pp. 68-75, Dec. 2003, doi:10.1109/MC.2003.1250885.

Koplow, J. P., HEAT EXCHANGER DEVICE AND METHOD FOR HEAT REMOVAL OR TRANSFER, USPTO Application #: 20090199997.

Kreith, F., *Fluid Mechanics*, CRC Press, Boca Raton, FL, 2001.

Kutz, M., *Heat Transfer Calculations*, McGraw-Hill, New York, 2005.

Masinet, E., "Energy Efficiency Potential for Data Centers", Fourth Annual California Climate Change Conference, Berkeley, CA September 11, 2007.

Mowris, R., "Strategies for improving HVAC efficiency with quality installation and service", 2006 International Energy Efficiency In Domestic Appliances & Lighting Conference, Session F4, June 22, 2006, London, U.K.

Parker, D., Sherwin, J., Hibbs, B., "Development of high efficiency air conditioner condenser fans", ASHRAE Transactions, June 2005.

Reist, P. C., *Aerosol Science and Techology*, 2nd edition, McGraw-Hill, New York, 1993.

Riviere, P., "Preparatory study on the environmental performance of residential room conditioning appliances (airco and ventilation), Technical Analysis of Best Available Technology, DGTREN service contract TREN/D1/40-2005/LOT10/S07.56606, July, 2008.

Sasakura, M., Hirao, T., Maeno, M., Oseki, S., "Compact, High-Efficiency Multi-Air-Conditioner LX", Mitsubishi Technical Review, Vol. 41, No. 2, April 2004.

Schlichting H., *Boundary Layer Theory*, McGraw-Hill, New York, 1979.

Siegel, J. A., "Particulate Fouling of HVAC Heat Exchangers", Ph. D. Thesis, University of California at Berkeley, 2002.

Steinberg, D. S., *Cooling Techniques for Electronic Equipment*, John Wiley and Sons, New York, 1991.

Tennekes, H., Lumley, J. L., *A First Course in Turbulence*, The MIT Press, Cambridge, MA, 1972.