

Enhancing Performance of Thermoelectric Coolers Through the Application of Distributed Control

R. D. Harvey, D. G. Walker, and K. D. Frampton

Abstract—The primary drawback of thermoelectric coolers (TECs) for electronics cooling applications is their thermodynamic inefficiency due to material limitations. The present work considers a control strategy to improve the overall coefficient of performance in an engineering system instead of addressing material shortcomings. Typical TECs are composed of several individual thermocouples that are powered in series and remove heat in parallel. If one of the numerous thermocouples is powered, all the thermocouples receive the same power whether or not they are needed. The fact that chips heat nonuniformly provides an opportunity for performance enhancement, by sensing and controlling the power to individual couples within the device. The current work presents evidence that applying distributed control to TEC operation can realize appreciable improvement in performance. Compared to monolithic cooling devices, a distributed control strategy can realize a factor of 2 increase in performance for the device studied. Additionally, this type of control can be used in conjunction with many of the existing material-based research initiatives to further compound the benefits.

Index Terms—Thermoelectric coolers (TECs).

I. INTRODUCTION

ALTHOUGH many research efforts to develop miniature refrigerators exist, thermoelectric coolers (TECs) are currently the only miniature cooling devices finding use in microelectronics cooling applications [1]. Unlike other miniature cooling schemes, TECs are solid-state devices and, therefore, are reliable, compact, and quiet. However, TECs generally suffer from poor device efficiencies, which limit their widespread applicability. To overcome these performance limitations, many research initiatives are currently seeking to improve the effectiveness of TECs by improving the materials used for these coolers [2]. Examples of possibilities that are being explored include skutterudites, superlattices, thin films, and other nanostructured materials [3]–[7]. Additional research related to chip cooling has also been devoted to improving fin-fan systems, minimizing contact resistances, and employing multistage TECs, which have already been implemented in chip-cooling applications [8]–[13].

Manuscript received August 1, 2005; revised October 8, 2006. This work was recommended for publication by Associate Editor T. Lee upon evaluation of the reviewers comments.

R. D. Harvey and D. G. Walker are with the Department of Mechanical Engineering, Vanderbilt University, Nashville, TN 37235 (e-mail: greg.walker@vanderbilt.edu).

K. D. Frampton is with the Institute for Sound and Vibration Research, University of Southampton, Southampton SO17 1BJ, U.K.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCAPT.2007.898376

In electronic chip-cooling applications, the packaging may consist of the chip, a TEC, a heat spreader, and finally, a fin-fan system. Thermal transport at each interface is enhanced with a thermal interface material (TIM). The heat spreader uniformly distributes the heat load produced by the chip. A TEC that is comparable in size and shape is contacted with the spreader to extract the heat from the chip. The fin-fan system must still remove the heat from the TEC and away from the computer. The purpose of the TEC is to lower the processor temperature so that the integrated circuits may operate in a desired temperature range, because the performance of modern computer chips will degrade above 85 °C [14]. However, the ultimate load on the power and cooling system increases as with any thermodynamic cycle.

The typical TEC unit that is currently marketed for chip-cooling applications consists of many individual thermocouples, called elements or cells, connected electrically in series but thermally in parallel. Each individual thermocouple is composed of two dissimilar materials, called pellets or legs, one of which is a p-type semiconductor and the other an n-type semiconductor [15]. In general, we can identify two metrics used to evaluate the performance of TECs.

At the material level, a TEC is rated by a dimensionless figure of merit, ZT . Typical values for ZT available from manufacturers are slightly less than unity. ZT s required to cool the future generations of computer chips will need to approach three and eventually even higher to be competitive to modern vapor-compression cycles [16]. This figure of merit is based solely on material properties but provides a useful comparison between the transport properties of different materials. As such, this metric has a direct effect on the performance of the system. Recent material improvements have resulted in ZT increases of between two and three fold [17]. Some sources purport theoretical increases even higher [18]. One problem with the current material improvements, however, is that they may require years to become fully implemented into production if they can be mass-produced at all.

When the TEC is placed into a system and used as a refrigerator the thermodynamic efficiency or coefficient of performance (COP) is often used to evaluate performance. This general measure of efficiency of a TEC is based on the amount of heat that it removes compared to the amount of work required to remove the heat. This value is calculated the same as any refrigeration cycle and is defined as

$$\text{COP} = \frac{Q_{\text{out}}}{W_{\text{in}}}. \quad (1)$$

In this configuration, all elements within the TEC device are powered simultaneously. If the heating in a chip is localized,

then elements within the TEC device may be needlessly powered, which will reduce the overall COP of the system.

One alternative approach is to control individual elements of the TEC independently based on the local heat load experienced by a single element [19]. This technology is independent of the improvement of ZT . Instead, this technology could improve the performance of any TEC device regardless of internal material efficiency. Because computer chips generate hot spots corresponding to the portion of the chip that is performing calculations at any given time, the heat load is not uniformly distributed. Heat spreaders improve the distribution of heat but do not transfer heat instantaneously [20]. Existing TECs can only be powered as a monolithic unit; all the elements receive the same amount of current regardless of their individual heat load. Therefore, the COP might be improved if each TEC element is only powered when it is needed, and then, powered only to the degree that is required to remove the amount of heat being produced. Therefore, we propose to enhance the system efficiency by distributed control of TEC units.

Distributed control refers to a method that allows the individual elements of a system to respond independently based on their respective state [21], [22]. Previous research has shown the potential of increased effectiveness by separately controlling individual thermocouples or groups of thermocouples in TECs [12], [23]. Another major benefit of distributed control is that it may be applied concurrently with other technological advances. Regardless of the material advances in TECs, distributed control would allow for the optimization of COP. The present work performs analysis for the basis of a strategy for using distributed control to improve the performance of TECs as a proof of concept and demonstrates preliminary effectiveness of the proposed approach.

II. THEORY

A typical TEC is a 2-D array of individual thermocouples that work together to cool larger areas. The device is in direct contact with the surface of a computer chip. For the present analysis, a one-dimensional (linear) array is considered for simplicity. Although transient effects such as periodic heating or migration of the heated spot may be important in some applications, the present analysis considers only steady-state heating.

Each section of the computer chip that is in contact with an individual couple is considered to be a separate and distinct heat-producing portion. Using the electrical analogy for heat transfer, the schematic for n number of thermocouples is given in Fig. 1. $Q_{\text{gen},i}$ represents the heat being produced by the i th portion of the chip. Q_i represents the heat being removed by each thermocouple. $Q_{\text{lat},i}$ represents the heat that is transferred by conduction between the i th and the i th+1 portions of the computer chip. R_A and R_B represent the thermal resistances between the chip and thermocouple, and between the thermocouple and the ambient environment, respectively. R_A is an equivalent resistance that accounts for chip packaging (including a heat spreader if used), TIM, and contact resistances. R_B accounts for TIM, contact resistances, and the effectiveness of the fin-fan system. Likewise, R_{lat} is an equivalent thermal resistance between portions of the computer chip and is based on the expected lateral conductivity of the chip with no heat

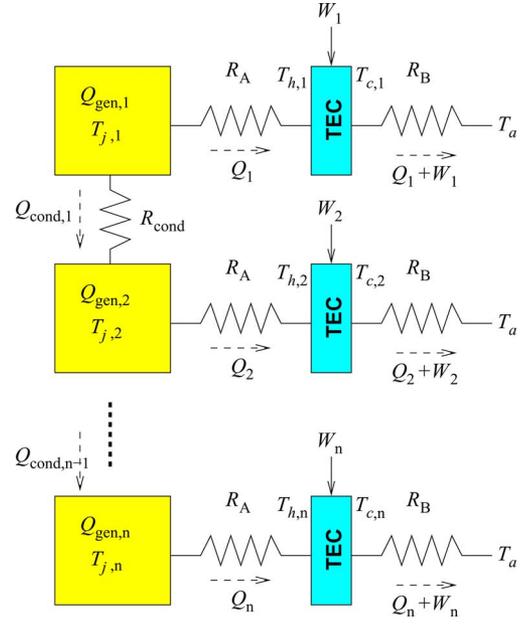


Fig. 1. Schematic diagram using the electrical analogy for heat transfer to represent n thermocouple cells removing heat from an equal number of computer chip sections.

spreader. W_i is the amount of work the thermocouple requires to remove the heat Q_i from the chip. Due to the first law of thermodynamics, however, the sum of the work and the heat, $Q_i + W_i$, must be removed from the system through R_B . $T_{j,i}$ represents the junction temperature, which is at the surface of the computer chip. $T_{c,i}$ and $T_{h,i}$ are the temperatures of the cold and hot surface of each thermocouple respectively. Finally, T_a is the ambient temperature.

The solution is found by solving a system of simultaneous algebraic equations. The heat removed by each thermocouple, Q_i , is [24]

$$Q_i = 2 \left[\alpha I_i T_{c,i} - \frac{1}{2} I_i^2 \frac{\rho}{G} - k G \Delta T_i \right] \quad (2)$$

where α is the Seebeck coefficient, I_i is the electrical current through the device, $T_{c,i}$ is the cold-side temperature, and ΔT_i is given by

$$\Delta T_i = T_{h,i} - T_{c,i}. \quad (3)$$

$T_{h,i}$ is the hot-side temperature. Referring to (2), G is a geometric factor equal to the cross-sectional area for heat removal divided by the height of a thermocouple, ρ is the electrical resistivity, and k is the thermal conductivity. The factor of two is required because each couple is composed of two pellets. The first term in brackets represents the heat being pumped by the couple due to the Peltier effect. This benefit is offset by the Joule heating in the device (the second term in brackets), and the parasitic heat flow in the couple due to the adverse temperature gradient (the last term in brackets). The parameters used for analysis are based on a specific Melcor device, Model CP5-31-06L and are given in Table I. Because the contact resistances, R_A and R_B , vary widely with temperature, surface preparation, TIM and heat removal system, they are largely unknown and regarded

TABLE I
ANALYSIS PARAMETERS

Parameter	Value
geometric factor, G	1.196cm
electrical resistivity, ρ	$1 \times 10^{-5} \Omega\text{m}$
thermal conductivity, k	1.5 W/mK
lateral conductivity, k_{lat}	10.0 W/mK
Seebeck coefficient, α	$\pm 2 \times 10^{-4} \text{V/K}$
thermal resistance, R_A and R_B	10K/W
lateral resistance, R_{lat}	20K/W
ambient temperature, T_a	25°C

as free parameters. A constant value based on other research efforts is assumed here for simplicity [25]. The work required by each thermocouple, W_i , is given by [25]

$$W_i = 2 \left[I_i^2 \frac{\rho}{G} + \alpha I_i \Delta T_i \right]. \quad (4)$$

The remaining equations used to define the system are obtained by balancing heat into and out of the chip sections and each thermocouple. That is, heat is conserved for each portion of the chip as

$$Q_{\text{gen},i} + Q_{\text{lat},i-1} = Q_{\text{lat},i} + Q_{\text{TEC},i}. \quad (5)$$

The heat transferred between two portions of the chip is given by

$$Q_{\text{lat},i} = \frac{T_{j,i} - T_{j,i+1}}{R_{\text{lat}}} \quad (6)$$

where T_j is the junction temperature. The heat removed by a thermocouple is equivalent to (2) and is given by

$$Q_{\text{TEC},i} = \frac{T_{j,i} - T_{c,i}}{R_A}. \quad (7)$$

The heat and work removed by the system is given by

$$Q_{\text{TEC},i} + W_i = \frac{T_{h,i} - T_a}{R_B}. \quad (8)$$

The algebraic system is non-linear with respect to current, so it must be solved iteratively for temperatures, heat loads, and currents. A Matlab script is used to maximize the overall COP by adjusting the applied current to each element individually for a given generation profile $Q_{\text{gen},i}$. The optimization is constrained to maintain a local junction temperature below some maximum value. Note that the effective COP of the system as a whole is different than (1), which is valid for monolithic devices. In the present work the COP is defined as

$$\text{COP} = \frac{\sum_{i=1}^n Q_{\text{gen},i}}{\sum_{i=1}^n W_i} = \frac{\sum_{i=1}^n Q_{\text{TEC},i}}{\sum_{i=1}^n W_i}. \quad (9)$$

Both forms in (9) are equivalent since all the heat generated must leave the device through the TECs during steady operation. However, because of lateral conduction, the distribution of the generated heat and the heat removed through the TECs will not necessarily be the same, even though their sums are

equal. This equation gives some insight to the inherent advantage of distributed control. If a couple is not powered, it may still provide some heat removal due to Fourier heat conduction. This advantage is termed free heat removal because the work required is zero and can result in scenarios that make analysis for the system different than typical analysis for a single TEC device. For instance, in some operating conditions the resulting COP of the present device will be infinite if all the heat can be removed by conduction without powering the device. In addition, favorable temperature gradients across the device can theoretically produce power, and COP will be negative. For performance analysis of a TEC not attached to a system, the COP is considered only when the device is powered at optimal conditions. These optimality conditions are for zero heat flux and maximum temperature drop or for zero temperature drop and maximum heat flux. While these conditions are useful for TEC design, they are not necessarily the same conditions that result in the best system efficiencies. Therefore, results obtained for the entire system may not compare to results for a cooler operating at optimal conditions and must be interpreted cautiously. The analysis in this work will concentrate on the regions where the device is powered, so COP is both positive and finite. Furthermore, we will compare the COP for a monolithic device in the system to the distributed controlled device in the same system only. Consequently, our comparison of COPs is for the system and not an isolated device.

III. RESULTS

The present work will examine an array of five thermocouples corresponding to a discretized chip with five nodes. For simplicity and ease of comparison only the first portion of the chip corresponding to the first thermocouple will produce heat; this is the worst-case scenario for a monolithic TEC that is cooling uniformly. For this example the input values of heat being generated by the first portion of the chip $Q_{\text{gen},1}$ range from 0.0 W to 5.0 W, representing a chip that is capable of generating 25 W of total energy. Because computer-chip performance degrades at 85 °C, this value is used as a maximum junction temperature $T_{j,\text{max}}$ for the optimization. At heat loads greater than 4.8 W, the five thermocouple configuration is unable to maintain $T_{j,1}$ less than 85 °C for the given resistances in the problem. These resulting computed values are not considered in the analysis because the system is operating in failure mode.

The maximum COP for the system with a TEC whose control is distributed among the elements is plotted in Fig. 2 as a function of heat load. For comparison the values of COP for the same configuration with the thermocouples connected in series (monolithic device) are also plotted. Therefore, Fig. 2 illustrates the maximum theoretical gain in COP from applying distributed control.

In order to achieve this gain, a distributed control rule must be applied, which will allow the system to respond appropriately via sensors and actuators to reproduce the optimum values. More specifically, the nodes within the system must be able to sense specified local parameters, to apply a distributed control rule in order to determine appropriate system response, and then to apply that response to the system. Ideally, the amount

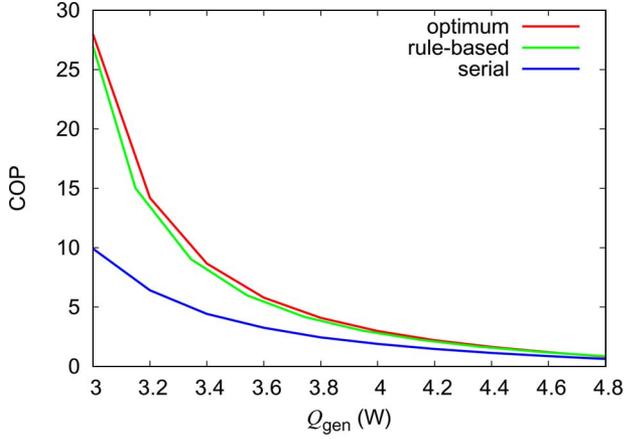


Fig. 2. Optimum COP and calculated COP from the distributed control rule as a function of the heat generated by the chip. Also shown are the comparable values of COP for a typical TEC in which the couples are powered in series.

of cooling achieved by the TEC should be directly related to the amount of heating provided by the chip. The most reasonable way to sense the amount of heat generation is by measuring local temperatures. The local temperatures that are available to be directly sensed by the TEC are the hot-side and cold-side temperatures, $T_{c,i}$ and $T_{h,i}$, of the individual couples. If the optimal current—the current that maximizes COP—may be determined by its relationship with the sensed parameters, then this relationship is the distributed control rule.

In order to find this distributed control rule, $T_{c,i}$ and $T_{h,i}$ were compared to the optimum current using curve fitting techniques. Several different methods were attempted including variations of multiple-linear and multiple-polynomial regression. Multiple-polynomial regression with squared terms provided a good fit, so higher order terms were not needed. A comparison of sensitivities of each term revealed that the $T_{c,i}$ squared term affected the current minimally. Therefore, the $T_{c,i}$ squared term was eliminated from the subsequent expression. Considering all factors, such as the number of terms and closeness of fit, the best distributed control rule found includes only one squared term and is given by

$$I_i = 1.5877 + 0.9126(T_{h,i} - 300) - 0.1491(T_{c,i} - 300) - 0.0050(T_{h,i} - 300)^2 \quad (10)$$

where $T_{c,i}$ and $T_{h,i}$ are provided in units of Kelvin and the resulting I_i is given in amperes. The resulting values of COP from applying (10) are also plotted in Fig. 2. The rule-based values are just slightly less than the optimal values throughout the operating range. Note that the rule-based values still represent a considerable improvement over the serial values. A comparison of the optimal values of current and the values of current computed with (10) are given in Fig. 3. In this figure each thermocouple has a distinct curve. The couple adjacent to the heat source has the highest temperatures and requires the most current. The other couples require less current because they are farther from the heat source. Each computed value differs by less than 1.0 A from the optimal value, and, in most cases, by less than 0.1 A. These minor variations in the computed values from the optimal values affect the COP as shown in Fig. 2. However,

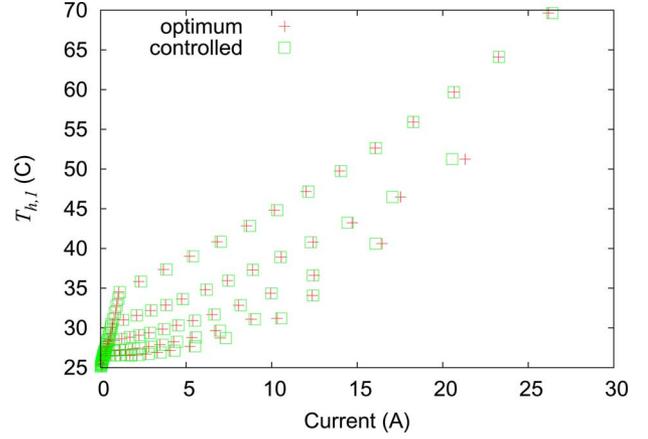


Fig. 3. Plot of the optimum COP data compared to the data calculated from the distributed control rule as calculated using multiple polynomial regression ($T_{h,i}$ squared term only).

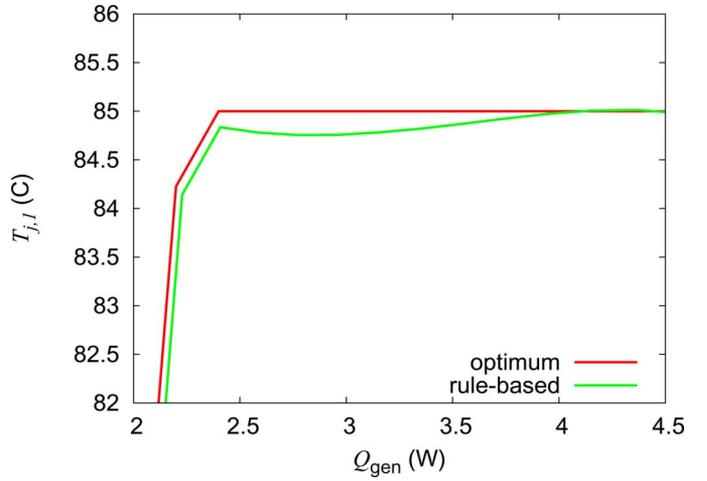


Fig. 4. Junction temperature of the heated cell for the optimal solution and the distributed control rule are shown as a function of the heat generated by the chip.

the reduction is minimal compared to the benefit gained over the serial value.

To ensure the new scheme does not enter a forbidden operating region—where the chip temperature exceeds the junction temperature—the junction temperatures $T_{j,i}$ were plotted for the optimum current values and the rule-based current values. The resulting graph for the first thermocouple is presented in Fig. 4. The resulting graph for the second thermocouple is presented in Fig. 5. For the optimum values, the junction temperature for the first thermocouple rises steadily until it reaches the maximum allowable junction temperature $T_{j,max}$. At this point, approximately 2.4 W, the couple must be powered to maintain $T_{j,max}$ throughout the operating regime. The rule-based temperature closely follows the temperatures from the optimal solution. The junction temperature $T_{j,1}$ remains just below the $T_{j,max}$ throughout most of the operating regime. Although $T_{j,1}$ does barely exceed $T_{j,max}$ between 4.0 and 4.5 W, this minor breach is considered acceptable. Fig. 4 shows that the rule-based computations can stay within system limitations and produce the desired gains in COP.

In Fig. 5, the second junction temperature $T_{j,2}$ initially rises due to the heat from conduction $Q_{lat,1}$. At approximately 2.4 W,

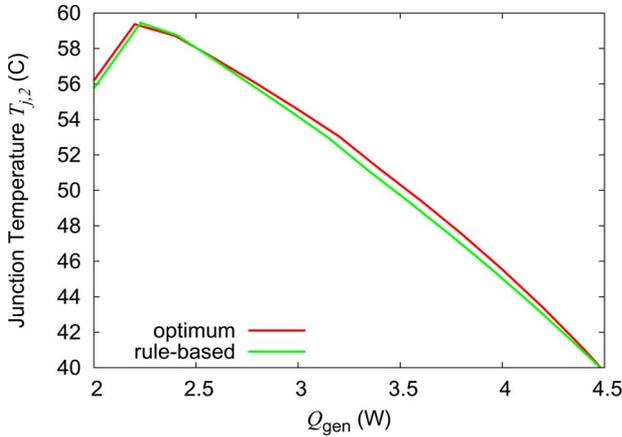


Fig. 5. Plots of the optimum and calculated junctions temperatures of the second cell versus the heat removed by the TEC.

the second thermocouple starts to consume more current and actually starts to lower $T_{j,2}$. Recall that only the first section of the chip is producing heat $Q_{gen,1}$. The first couple responds to this generation but can not satisfy the load and maintain $T_{j,1}$ less than $T_{j,max}$. This reduction in temperature of the adjacent section at first may seem counterintuitive to the goal of increasing COP, since $T_{j,2}$ is not approaching $T_{j,max}$. However, the benefit in this reduction is to increase the temperature difference between $T_{j,1}$ and $T_{j,2}$, thereby increasing $Q_{lat,1}$ as described in (5). So, the second couple, as well as the other couples, gradually assumes more of the heat load, although another portion of the chip is producing the heat. When the adjacent cell must turn on in order to reduce the load on its neighbor, the operation enters a second regime where the gradients in the chip become greater. One claimed benefit of using distributed control to power TECs is the lack of temperature gradients, yet this regime is producing larger temperature gradients. The design tradeoffs of distributed TECs, therefore, must include internal chip gradients. This drawback must be balanced with the ultimate gains in COP. The benefit of each couple assuming a portion of the heat load can be seen in the Joule heating term of (2). Each couple consumes a small amount of current to optimize the load on the primary thermocouple, thereby controlling the Joule heating. Fig. 6 illustrates the degree to which the five thermocouples assume three different heat loads. The corresponding COPs are also shown.

The distribution of heat load assumed by the other couples will be proportional to the thermal resistance between the portions of the chip (R_{cond}). As R_{cond} approaches infinity, the heat load seen by the other couples will approach zero. This limit will require each thermocouple to manage the heat load of its assigned chip portion alone. This means that the assigned chip portion can only produce an amount of heat less than or equal to the capacity of the individual thermocouple without operating in failure mode. As R_{lat} approaches zero, the heat load seen by each couple will be exactly the same regardless of where the generation occurs. Therefore, distributed control offers no benefit, and the optimum solution will be the same as a serially connected TEC. This relationship shows that a heat spreader with a

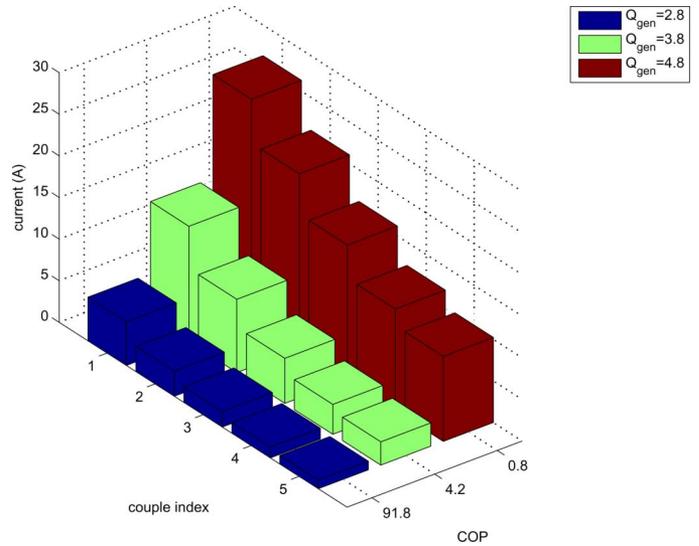


Fig. 6. Current applied to each cell for three different heat loads. The overall COP in each case is shown.

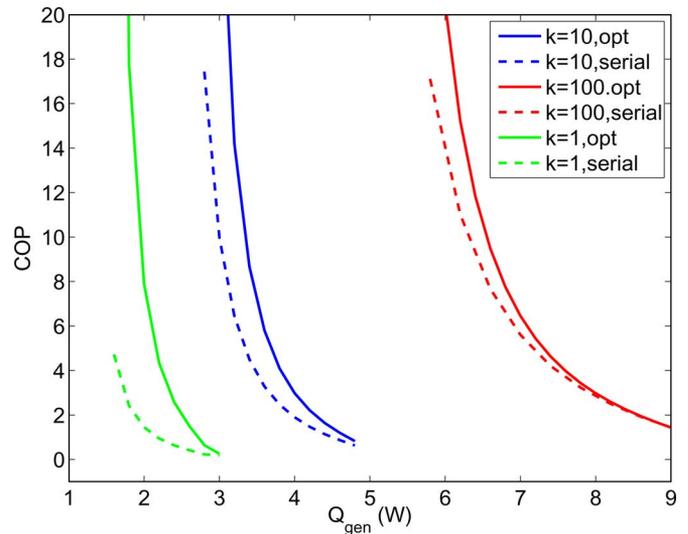


Fig. 7. Comparison of optimum COPs for the distributed control case and the COPs for the serial case for different values of lateral conductivity, k_{cond} .

high thermal conductivity may benefit a traditional TEC. However, a TEC operated with distributed control will benefit from a reduced value of thermal conductivity and should not employ a traditional heat spreader. The value of R_{lat} chosen for the current analysis approximates that anticipated from a computer chip with no heat spreader.

Fig. 7 compares the COPs for the serial and distributed control cases for three different values of lateral conductivity, k_{lat} . Lateral conductivity is inversely proportional to lateral thermal resistance, R_{lat} . The figure shows the gain from distributed control as the space between the solid and the dashed curves. As the value of k_{cond} increases, the benefit that distributed control offers decreases.

The effect of the thermal conductance of interfaces on the performance of the system is also considered. The values of R_A and R_B were selected to represent of the typical conditions of

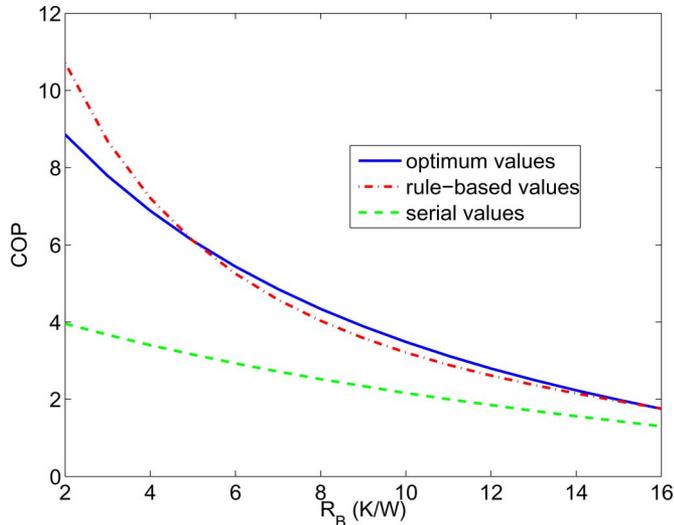


Fig. 8. Variation of COP over a range of R_B values with a constant value of $R_A = 10$ K/W and a constant heat load of 3.9 W.

existing systems. These values vary greatly based on the contact resistances, TIM, and fin-fan systems. Fig. 8 shows how the COP varies over a wide range of R_B values with R_A held constant for an intermediate heat load of 3.9 W. The optimum, rule-based, and serial solutions are shown. R_B was chosen to vary because it typically has the widest range of values. Fig. 8 shows that the optimal COPs represent a two-fold increase in serial COP values for 3.9 W regardless of the R_B value. The most significant feature in Fig. 8 is that the rule-based values exceed the optimum values. We do not expect any control rule to give better performance than the optimal operating conditions. However, the present rule was specifically designed for a resistance value of $R_B = 10$ K/W, where the rule based performance is slightly lower than the optimal performance. The control rule appears to outperform the optimal solution because the constraint used in the optimization ($T_j < 85^\circ\text{C}$) is being violated. Therefore, this result indicates that a control rule must be devised for each cooling solution and is not characteristic of a particular chip, for example. The distributed control rule is sufficient for small variances in R_B , but when R_B varies considerably from the value used in this paper a new rule must be developed.

Other considerations relevant to the success of the distributed control approach include interconnect heating and thermal fatigue. Due to the increased number of connections required by distributed control, the amount of power losses due to interconnect heating will also increase. Preliminary calculations show that due to the high current levels and the small sizes, the worst-case scenario of individual control of each couple may be prohibitive. One solution may be to control the individual couples in clusters or arrays. Such prospects should be investigated to a greater extent in a working prototype to determine feasibility.

Another immediate issue affecting feasibility of distributed control is thermal fatigue. The application of distributed control will reduce temperature gradients across the computer chip (at least in nominal operating regimes), and therefore, reduce thermal stresses internal to the chip itself. However, current TEC

technology is quite susceptible to thermal fatigue due to cycling. Traditional TECs are powered at one level constantly, and are not frequently cycled. Frequent cycling of a TEC causes the internal electrical connections to fail. Some advances in materials and manufacturing have already been considered and some have been implemented [26]. Nevertheless, the extent of the thermal stresses due to distributed control should be considered.

IV. CONCLUSION

The need for improved performance of TECs is well-documented. By managing the individual thermocouples internal to a typical TEC with distributed control, significant gains in performance can be realized. A computer model was developed that determines the optimal performance of the individual thermocouples in a simplified linear array for a non-uniform heat distribution case. This optimum performance can be closely approximated by applying a distributed control rule to the TEC. This rule accepts the sensed local temperature values and determines the current loads to individual thermocouples to maximize efficiency. The results represent improved performance throughout the operating regime. In much of the regime, this improvement is 1.5 to 2.0 times the traditional values. This improvement could, as a minimum, narrow the gap between thermal demands of modern computing equipment and cooling capability making material advances even more rewarding. Distributed control of thermoelectric devices for chip cooling applications where non-uniform heating occurs is a technology that deserves further investigation. Because the heat loads and distribution of heat loads as well as construction of chips vary dramatically between chip technologies, a general solution can not be provided here. The concepts presented in this article should be studied in the context of different chip design and load requirements, transient effects, modern interface resistance technologies, and so forth to assess the widespread applicability of this work.

REFERENCES

- [1] P. E. Phelan, V. A. Chiriac, and T.-Y. T. Lee, "Current and future miniature refrigeration cooling technologies for high power microelectronics," *IEEE Trans. Compon. Packag. Technol.*, vol. 25, no. 3, pp. 356–365, Sep. 2002.
- [2] H. J. Goldsmid and G. S. Nolas, "A review of new thermoelectric materials," in *Proc. IEEE 20th Int. Conf. Thermoelect.*, Beijing, China, Jun. 2001, pp. 1–6.
- [3] C. Uher, "Skutterudites: Promising power conversion thermoelectrics," in *Proc. 22nd Int. Conf. Thermoelect. (ICT)*, La Grande Motte, France, Aug. 2003, pp. 42–47.
- [4] G. Chen, "Phonon engineering in nanostructures for solid-state energy conversion," *Mater. Sci. Eng. A*, vol. 292, pp. 155–161, 2000.
- [5] J. W. Sharp, S. J. Poon, and H. J. Goldsmid, "Boundary scattering and the thermoelectric figure of merit," *Solid-State Phys.*, vol. 187, pp. 507–516, 2001.
- [6] L. W. da Silva and M. Kaviany, "Micro-thermoelectric cooler: Interfacial effects on thermal and electrical transport," *Int. J. Heat Mass Transf.*, vol. 47, pp. 2417–2435, 2004.
- [7] J. Nurnus, H. Böttner, C. Künzel, U. Vetter, A. Lambrecht, J. Schumann, and F. Völklein, "Thin film based thermoelectric energy conversion systems," in *Proc. IEEE 21st Int. Conf. Thermoelect.*, Long Beach, CA, Aug. 2002, pp. 523–527.
- [8] D. Astrain, J. G. Vián, and M. Domínguez, "Increase of COP in the thermoelectric refrigeration by the optimization of heat dissipation," *Appl. Therm. Eng.*, vol. 23, no. 17, pp. 2183–2200, Dec. 2003.
- [9] S. B. Riffat and X. Ma, "Improving the coefficient of performance of thermoelectric cooling systems: A review," *Int. J. Energy Res.*, vol. 28, no. 9, pp. 753–768, Jul. 2004.

- [10] X. C. Xuan, "Investigation of thermal contact effect on thermoelectric coolers," *Energy Conv. Manag.*, vol. 44, no. 3, pp. 399–410, Feb. 2003.
- [11] J. P. Gwinn and R. L. Webb, "Performance and testing of thermal interface materials," *Microelectron. J.*, vol. 34, no. 3, pp. 215–222, Mar. 2003.
- [12] J. Chen, Y. Zhou, H. Wang, and J. T. Wang, "Comparison of the optimal performance of single- and two-stage thermoelectric refrigeration systems," *Appl. Energy*, vol. 73, no. 3–4, pp. 285–298, Nov. 2002.
- [13] V. Semenyuk, "Cascade thermoelectric micro modules for spot cooling high power electronic components," in *Proc. IEEE 21st Int. Conf. Thermoelect.*, Long Beach, CA, Aug. 2002, pp. 531–534.
- [14] Semiconductor Industry Association, International Technology Roadmap for Semiconductors: Executive Summary Tech. Rep., SEMATECH, 2003.
- [15] M. Hodes, "One-dimensional analysis of thermoelectric modules," in *Proc. 9th Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst. (ITHERM)*, Jun. 2004, pp. 242–250.
- [16] A. Majumdar, "Thermoelectricity in semiconductor nanostructures," *Science*, vol. 303, no. 6, pp. 777–778, Feb. 2004.
- [17] A. Shakouri, "Nanoscale devices for solid state refrigeration and power generation," in *Proc. 20th Semi-Therm Symp.*, 2004, pp. 1–9.
- [18] R. F. Service, "Temperature rises for devices that turn heat into electricity," *Science*, vol. 306, pp. 806–807, Oct. 2004.
- [19] D. G. Walker, K. D. Frampton, and R. D. Harvey, "Distributed control of thermoelectric coolers," in *Proc. 9th Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst. (ITHERM)*, Las Vegas, NV, Jun. 2004, pp. 361–366.
- [20] R. S. Viswanath, V. S. Wakharkar, A. Watwe, and V. Lebonheur, "Thermal performance challenges from silicon to systems," *Intel Technol. J.*, vol. 4, no. 3, pp. 1–16, Aug. 2000.
- [21] F. Giuletti, L. Pollini, and M. Innocenti, "Autonomous formation flight," *IEEE Control Syst. Mag.*, pp. 34–44, Dec. 2000.
- [22] M. Kader, M. Lenczner, and Z. Mrcarica, "Distributed control based on distributed electronic circuits: Application to vibration control," *Microelectron. Rel.*, vol. 41, no. 11, pp. 1857–1866, Nov. 2001.
- [23] R. D. Harvey, D. G. Walker, and K. D. Frampton, "Distributed control to improve performance of thermoelectric coolers," in *Proc. Int. Mech. Eng. Congr. Expo.*, Anaheim, CA, Nov. 2004, [CD ROM].
- [24] M. J. Ellsworth, Jr. and R. E. Simons, *Analysis of a Thermoelectric Cooler with Non-Uniform Heating, Packaging of Electronic and Photonic Devices*, G. J. Kowalski, Ed. New York: ASME, 2000, vol. EEP-28, pp. 269–276.
- [25] R. Chein and G. Huang, "Thermoelectric cooler application in electronic cooling," *Appl. Therm. Eng.*, vol. 24, no. 14–15, pp. 2207–2217, Oct. 2004.
- [26] D. D. L. Chung and Y. M. Hoi, "Flexible graphite as a compliant thermoelectric material," *Carbon*, vol. 40, no. 7, pp. 1134–1136, Jun. 2002.

R. D. Harvey received the B.S. degree in nuclear engineering from the Rensselaer Polytechnic Institute, Troy, NY, in 1996, and the M.S. degree in mechanical engineering from Vanderbilt University, Nashville, TN, in 2005.

He is currently a Lieutenant in the United States Navy. After his career in the Navy he plans to continue research in energy conversion and energy transport.

D. G. Walker is an Assistant Professor of mechanical engineering at Vanderbilt University, Nashville, TN, with interests in energy transport at small scales, simulation of microelectronic device physics, and direct energy conversion technologies.

K. D. Frampton received the B.S. and M.S. degrees in mechanical engineering from the Virginia Polytechnic Institute and State University, Blacksburg, in 1988 and 1991, respectively, and the Ph.D. degree from Duke University, Durham, NC, in 1996.

He is currently an Assistant Professor of mechanical engineering at Vanderbilt University, Nashville, TN. He was with Wyle Laboratories, Arlington, VA, for two years as an Acoustical Engineer. He was a Research Assistant Professor at Duke University, then he joined the faculty at Vanderbilt University in 1998.