Note that our proposed fast algorithms achieve 69× runtime speedup over the simple implementation without using these fast algorithms. While the problem size in this example is small (i.e., only 33 sensors are selected in Figure 1), the runtime speedup of the proposed fast algorithms would be more pronounced when applied to larger-size problems.

4.2 Measurement Experiment

In this sub-section, we further demonstrate the efficacy of the proposed entropy method by collecting the temperature sensor data from an industrial dual-core microprocessor that contains 24 temperature sensors distributed in both cores and caches. In this experiment, all 24 temperature sensors are first calibrated using infrared imaging. Next, temperature readings are recorded from these 24 sensors when running a subset of SPEC2006 benchmarks. Similar to the previous example, the measurement data are partitioned into two non-overlapped groups: the training set and the testing set. Our objective in this example is to identify the most important temperature sensors out of these 24 candidates. Once the important sensors are found, the temperature values of other sensors can be estimated and, hence, these unimportant sensors may be eliminated in future design to reduce silicon area and/or power consumption.

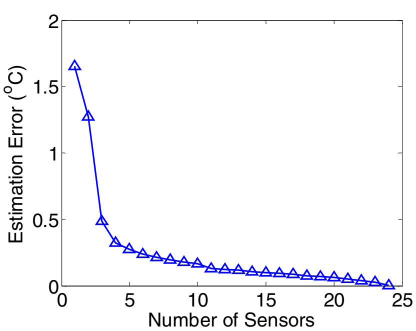

Figure 3. Thermal estimation error decreases, as the number of selected sensors increases.

Figure 3 shows the thermal estimation error as a function of the number of selected temperature sensors. Here, the estimation error is calculated by (27) using the measurement data in the testing set. Note that the estimation error is about 0.3 °C, when 5 sensors are selected. Selecting more than 5 sensors has diminishing returns in information collected. It, in turn, demonstrates that the optimal choice is to deploy 5 sensors for thermal monitoring in this example. The aforementioned analysis has potential benefits in both system and design areas. First, it reduces the number of sensors and information processing for thermal-aware OS (i.e., operating system) task scheduling [5]. Second, it demonstrates that our proposed sensor allocation technique is effective in real world and can be used for future microprocessor design.

5. CONCLUSIONS

In this paper, we propose a new information-theoretic framework to efficiently find the optimal spatial locations of on-chip temperature sensors for full-chip thermal monitoring. The key idea is to quantitatively model the uncertainty of on-chip temperature variation by differential entropy. Our experimental results demonstrate that the proposed entropy-based method achieves superior accuracy (around 1.4× error reduction) for full-chip thermal monitoring over other traditional methods. The techniques developed in this paper can be further incorporated into various thermal management schemes for hotspot-limited microprocessors.

6. ACKNOWLEDGEMENTS

The authors acknowledge the support of the C2S2 Focus Center, one of six research centers funded under the Focus Center Research Program (FCRP), a Semiconductor Research Corporation entity. This work is also supported in part by the National Science Foundation under contract CCF–0915912.

7. REFERENCES

[1] H. Hamann, A. Weger, J. Lacey, Z. Hu and P. Bose, "Hotspot-limited microprocessors: direct temperature and power distribution measurements," *IEEE JSSC*, vol. 42, no. 1, pp. 56-65, 2007.

[2] M. Pedram and S. Nazarin, "Thermal modeling, analysis, and management in VLSI circuits: principles and methods," *Proceedings of the IEEE*, vol. 94, no. 8, pp. 1487-1501, 2006.

[3] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," *IEEE HPCA*, pp. 171-182, 2001.

[4] K. Skadron, M. Stan, H. Wei, S. Velusamy, K. Sankaranarayanan and D. Tarjan, "Temperature-aware microarchitecure," *IEEE ISCA*, pp. 2-13, 2003.

[5] J. Choi, C. Cher, H. Franke, A. Weger, and P. Bose, "Thermal-aware task scheduling at the system software level" *IEEE ISLPED*, pp. 213-218, 2007.

[6] Y. Wang, K. Ma and X. Wang, "Temperature-constrained power control for chip multiprocessors with online model estimation," *IEEE ISCA*, pp. 314-324, 2009.

[7] T. Ebi, M. Faruque and J. Henkel, "TAPE: Thermal-aware agent-based power economy multi/many-core architectures," *IEEE ICCAD*, pp. 302-309, 2009.

[8] C. Cher and E. Kursun, "Exploring the effects of on-chip thermal variation on high-performance multicore architectures," *ACM TACO*, vol. 8, no. 1, pp. 2:1-2:22, 2011.

[9] S. Bokar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi and V. De, "Parameter variations and impact on circuits and microarchitecture," *IEEE DAC*, pp. 338-342, 2003.

[10] S. Memik, R. Mukherjee, M. Ni and J. Long, "Optimizing thermal sensor allocation for microprocessors," *IEEE TCAD*, vol. 27, no. 3, pp. 516-527, 2008.

[11] A. Nowroz, R. Cochran and S. Reda, "Thermal monitoring of real processors: techniques for sensor allocation and full characterization," *IEEE DAC*, pp. 56-61, 2010.

[12] Y. Zhang, B. Shi and A. Srivastara, "A statistical framework for designing on-chip thermal sensing infrastructure in nano-scale systems," *IEEE ISPD*, pp. 169-176, 2010.

[13] F. Zanini, D. Atienza, C. Jones and G. Micheli, "Temperature sensor placement in thermal management systems for MPSoCs," *IEEE ISCAS*, pp. 1065-1068, 2010.

[14] S. Sharifi and T. Rosing, "Accurate direct and indirect on-chip temperature sensing for efficient dynamic thermal management," *IEEE TCAD*, vol. 29, no. 1, pp. 1586-1599, 2010.

[15] Y. Zhang, A. Srivastava and M. Zahran, "Chip level thermal profile estimation using on-chip temperature sensors," *IEEE ICCD*, pp. 432-437, 2008.

[16] R. Cochran and S. Reda, "Spectral techniques for high-resolution thermal characterization with limited sensor data," *IEEE DAC*, pp. 478-483, 2009.

[17] Y. Zhang and A. Srivastava, "Accurate temperature estimation using noisy thermal sensors," *IEEE DAC*, pp. 472-477, 2009.

[18] W. Zhang, X. Li and R. Rutenbar, "Bayesian virtual probe: minimizing variation characterization cost for nanoscale IC technologies via Bayesian inference," *IEEE DAC*, pp. 262-267, 2010.

[19] D. Brooks, V. Tiwari and M. Martonosi, "Wattch: a framework for architecture-level power analysis and optimizations," *IEEE ISCA*, pp. 83-94, 2000.

[20] C. Bishop, *Pattern Recognition and Machine Learning*, Prentice Hall, 2007.

[21] R. Horn and C. Johnson, *Matrix Analysis*, Cambridge Press, 2007.