Applied Energy 138 (2015) 414–422

Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apenergy

Energy efficient hotspot-targeted embedded liquid cooling of electronics

Chander Shekhar Sharma [a], Manish K. Tiwari [a,1], Severin Zimmermann [a], Thomas Brunschwiler [b], Gerd Schlottig [b], Bruno Michel [b], Dimos Poulikakos [a,*]

[a] Department of Mechanical and Process Engineering, ETH Zurich, 8092 Zurich, Switzerland
[b] Advanced Micro Integration, IBM Research–Zurich, 8803 Rüeschlikon, Switzerland

HIGHLIGHTS

- We present a novel concept for hotspot-targeted, energy efficient ELC for electronic chips.
- Microchannel throttling zones distribute flow optimally without any external control.
- Design is optimized for highly non-uniform multicore chip heat flux maps.
- Optimized design minimizes chip temperature non-uniformity.
- This is achieved with pumping power consumption less than 1% of total chip power.

ARTICLE INFO

Article history:
Received 5 June 2014
Received in revised form 20 September 2014
Accepted 24 October 2014

Keywords:
Hotspot-targeted cooling
Microchannel cooling
Electronics cooling
Hotspots
Energy efficient computing
Multicore microprocessors

ABSTRACT

Large data centers today already account for nearly 1.31% of total electricity consumption with cooling responsible for roughly 33% of that energy consumption. This energy intensive cooling problem is exacerbated by the presence of hotspots in multicore microprocessors due to excess coolant flow requirement for thermal management. Here we present a novel liquid-cooling concept, for targeted, energy efficient cooling of hotspots through passively optimized microchannel structures etched into the backside of a chip (embedded liquid cooling or ELC architecture). We adopt an experimentally validated and computationally efficient modeling approach to predict the performance of our hotspot-targeted ELC design. The design is optimized for exemplar non-uniform chip power maps using Response Surface Methodology (RSM). For industrially acceptable limits of approximately 0.4 bar (40 kPa) on pressure drop and one percent of total chip power on pumping power, the optimized designs are computationally evaluated against a base, standard ELC design with uniform channel widths and uniform flow distribution. For an average steady-state heat flux of 150 W/cm^2 in core areas (hotspots) and 20 W/cm^2 over remaining chip area (background), the optimized design reduces the maximum chip temperature non-uniformity by 61% to 3.7 °C. For a higher average, steady-state hotspot heat flux of 300 W/cm^2, the maximum temperature non-uniformity is reduced by 54% to 8.7 °C. It is shown that the base design requires a prohibitively high level of pumping power (about 2000 fold for 150 W/cm^2 case and 600 fold for 300 W/cm^2 case) to match the thermal performance of the optimized, hotspot-targeting designs. The pumping power requirement for optimized designs is only 0.23% and 0.17% of the total chip power for 150 W/cm^2 and 300 W/cm^2 hotspot heat flux respectively. Moreover, the optimized designs distribute the coolant flow without any external flow control devices and the performance is only marginally affected by the manifold geometry used to supply the coolant to the microchannel heat transfer structure. This also attests to the robustness of the optimized embedded microchannel structures.

1. Introduction

Since the advent of first very large scale integrated circuits, electronic chip manufacturers have endeavored to keep up with Moore's law [1]. Chip performance has improved due to increasing gate density as well as fabrication of multiple cores on a single chip. However, the inability to follow the Dennard scaling for

* Corresponding author. Tel.: +41 44 632 27 38; fax: +41 44 632 11 76.
E-mail address: dpoulikakos@ethz.ch (D. Poulikakos).
[1] Present address: Department of Mechanical Engineering, University College London (UCL), Torrington Place, London WC1E 7JE, UK.

http://dx.doi.org/10.1016/j.apenergy.2014.10.068
0306-2619/

Nomenclature

A_{fs}	interfacial area density $(2/(w_c + w_w))$
h_f	specific enthalpy of fluid (J/kg)
h_c	microchannel height (m)
K	permeability (m^2)
l_{th}	length of throttling zone (m)
n,m	integral multipliers
$\dot{m}$	mass flow rate (kg/s)
$\overset{*}{m}_i$	relative mass flow rate
P	pressure (Pa)
q''	heat flux (W/m^2)
R	flow resistance (Pa s/kg)
$R_{thermal}$	thermal resistance (°C m^2/W)
T	temperature (K)
$T_{s,max}$	maximum temperature of chip (at base of silicon die)
$\Delta T_{s,max}$	chip wide temperature difference $(= T_{s,max} - T_{s,min})$
$\mathbf{U}$, $\mathbf{U}_s$	true and Superficial velocity vectors for porous domain (m/s)
$\dot{V}$	volumetric flow rate (m^3/s)
w_c, w_w	microchannel and wall widths (m)
W, L	overall chip dimensions (m)
W_m	half width of inlet manifold
$\dot{W}_{pumping}$	pumping power (W)
$\dot{W}_{chip}$	total power dissipated by the chip (W)

Greek letters

α_{fs}	interfacial heat transfer coefficient (W/m^2 K)
φ_m	fraction of total coolant flow flowing through HSBG zone
γ	volume porosity $(w_c/(w_c + w_w))$
λ	thermal conductivity (W/m K)

Subscripts

bg, hs, th	background, hotspot and throttling zones
f, s	fluid and solid
fp, sp	fluid and solid phases of porous medium
intf	interface
max	maximum
min	minimum
opt	optimum value
V	optimized value at total flow rate $\dot{V}_T$
B	base case
T	total

supply voltage has lead to high chip heat flux dissipation densities, especially during the last decade [2]. As a consequence, electronic reliability considerations are driving a shift away from air-cooling [3] towards a number of alternative approaches including liquid cooling [4,5], two-phase cooling [6,7], phase change materials (PCM) [8,9] and nanofluids [10,11]. Single-phase liquid cooling has long been identified as an effective and feasible approach for cooling high heat flux density chips. Starting with the landmark work of Tuckerman and Pease [12], liquid cooling of chips has been analyzed in great detail. This includes investigations on traditional microchannel heat sinks [13], manifold microchannel (MMC) heat sinks [5,14], spray and jet cooling [15].

Most of the research on liquid cooling of chips has focused on maximum temperature ($T_{s,max}$) reduction under uniform heat flux dissipation conditions. However, it is also necessary that the chip temperature is spatially as uniform as possible (i.e. approaching the isothermal chip condition). Large temperature gradients in the package increase thermal stresses in the chip to substrate or heat sink interface, reduce electronic reliability in regions of high temperature and create circuit imbalances in CMOS devices [4]. A few studies have focused on reducing chip temperature non-uniformity ($\Delta T_{s,max}$) under a uniform chip heat flux map. These include use of flow boiling of dielectric liquid [16], single phase liquid cooling with variable pin fin density [17], variable microchannel width in the streamwise direction [18–20] and double-layer microchannel structure [21].

During the last decade, computer architecture has witnessed a trend towards multicore microprocessors due to limitations encountered in further increasing the single-core clock frequencies [22]. In a multicore microprocessor, each core houses the execution units [23]. In modern high performance chips, the cores can dissipate on an average up to 150 W/cm^2 while the rest of the chip dissipates as low as 20 W/cm^2. Since the cores dissipate multiple times higher heat flux than the rest of the chip, we refer to the cores as the hotspots. The large difference in heat dissipation between hotspot and background regions makes the goal of isothermal junction temperatures even more challenging.

Several approaches have been proposed in the past aimed at achieving an isothermal chip condition by preferential cooling of hotspots (henceforth referred to as hotspot-targeted cooling). These include, for example, use of thermoelectric cooling [24] and electrowetting [25]. However, these approaches are limited by low device efficiencies, low heat flux pumping capacities and contact parasitic resistance for the former [24] and relatively low heat fluxes for the later [25]. Attempts using single phase liquid cooling have also been reported with varying degree of success [26–28]. Additionally, conventional backside attached MMC heat sinks are thermally joined to the chip through a thermal interface material (TIM) which, combined with the thickness of substrate and heat sink base, results in heat spreading and high thermal resistance. For effective hotspot-targeted cooling, the net thermal resistance and heat spreading can be reduced by circulating the coolant through microchannels etched into the backside of the chip (also termed as embedded or direct chip backside microchannel cooling) [29–32].

It is also critical that any cooling of microprocessors be energy efficient [33]. This becomes even more relevant for large datacenters employing thousands of multicore microprocessors. Direct electricity consumption by large data centers had almost tripled between 2000 and 2010 and reached 1.31% of the total world electricity consumption by 2010 due to increased demands for Information Technology (IT) related services. Although the growth in data center energy consumption slowed down in the period 2005–2010 [34], the still increasing data center energy bill is quickly becoming a cause of concern [35]. Cooling overhead for these datacenters contributes roughly 33% of this electricity consumption. Hence, efficient cooling can significantly reduce the IT energy bill [35,36].

In this paper, we propose a novel, hotspot-targeted and, at the same time, highly energy efficient cooling solution for typical modern day multicore microprocessors. We employ a passive, energy saving approach to alter the heat sink design by optimizing the microchannel geometry and flow rate distribution. The hotspot-targeted cooling design is described in Section 2. Sections 3 and 4 focus on the modeling, validation and optimization methodologies and the final results for two exemplar chip power maps are presented in Section 5. Section 6 presents general rules for design of hotspot-targeted microchannel heat sink. Lastly, Section 7