Minimizing Thermal Gradient and Pumping Power in 3D IC Liquid Cooling Network Design

Gengjie Chen, Jian Kuang, Zhiliang Zeng, Hang Zhang, Evangeline F. Y. Young, Bei Yu

Department of Computer Science and Engineering, The Chinese University of Hong Kong

Introduction

Why 3D IC Liquid Cooling?

- **Power** is the number one problem in chip design
- **3D IC** is promising for increasing computer performance
- But 3D IC **worsens** power problem by
  - higher heat dissipation density
  - larger thermal resistance from junction to ambient
- Microchannel-based liquid cooling is proposed as a solution

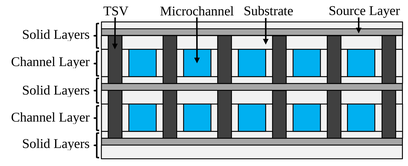

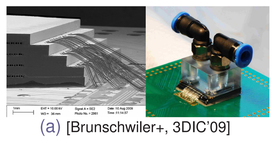

(a) [Brunschwiler+, 3DIC'09]

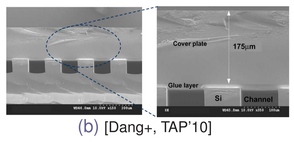

(b) [Dang+, TAP'10]

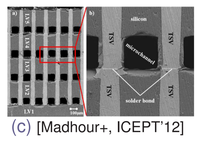

(c) [Madhour+, ICEPT'12]

Challenges for 3D IC Liquid Cooling

- Hot downstream and cool upstream $\Longrightarrow$ **large thermal gradient** $\Longrightarrow$ reliability and timing issues
- limited channel diameter $\Longrightarrow$ **high pumping requirement** $\Longrightarrow$ overhead to whole system
- Limitations of previous work
  - No considering thermal gradient
  - Assuming unidirectional straight channels
  - Assuming unrealistic constant-temperature heat source

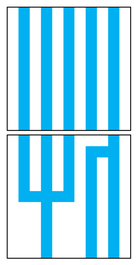

Thermal Modeling

- Most existing models assume unidirectional straight channels
- 4-register model (4RM) in 3D-ICE [Sridhar+, TOC'14]
  - Accurate
  - Has been extended for flexible topology
  - **Slow**
- We construct a fast 2-register model (2RM) for cooling network

Basics

- Divide channel layer into **basic cells** with a 2D grid
  - Either solid (white/black, black reserved for TSV) or liquid (blue)
- Solve local pressure P_i and flow rate $Q_{i,j}$ from a **linear system**
  - $Q_{i,j} = g_{fluid,i,j} \cdot (P_i - P_j)$ ($g_{fluid,i,j}$: fluid conductance)
  - $\sum_{j \in N_i} Q_{i,j} = 0$ (N_i: neighboring cells, inlet/outlet)

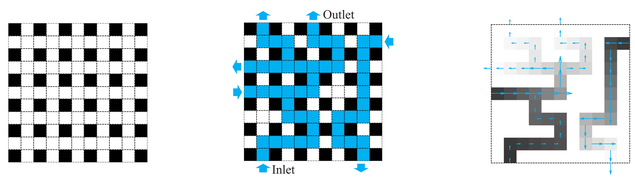

4RM Model

- **Thermal cell** = basic cell
- Solve temperature from a **linear system** considering three kinds of heat transfer
  - Solid-solid thermal conductance $g_{ss} = \frac{q_{i,j}}{T_i - T_j} = \frac{k_{solid} \cdot A_{i,j}}{l_{i,j}}$
  - Solid-liquid thermal conductance $g_{sl} = \frac{q_{i,j}}{T_i - T_j} = g^*_{sl} \parallel g^*_{ss} = \frac{g^*_{sl} \cdot g^*_{ss}}{g^*_{sl} + g^*_{ss}}$ with $g^*_{sl} = h_{conv} A_{i,j}$
  - Liquid-liquid heat transfer $q_{ll} = C_v \cdot \sum_{j \in N_i} (Q_{j,i} \cdot T^*_{j,i}) = \frac{C_v}{2} \cdot \sum_{j \in N_i} (Q_{j,i} \cdot T_j)$

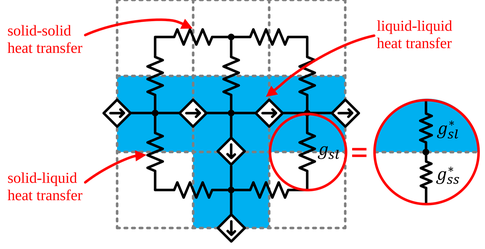

Faster 2RM Model

- No conforming channel geometry $\Longrightarrow$ **larger** and fewer thermal cells $\Longrightarrow$ **speed-up**
- Important due to frequent simulation

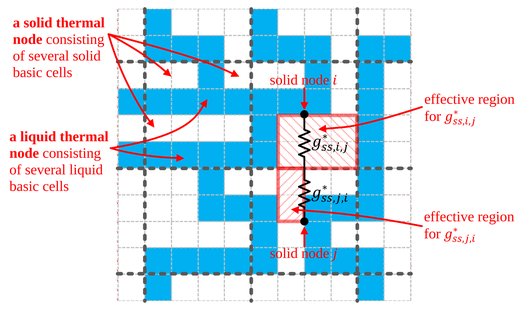

Thermal nodes

- In solid layers, $m \times m$ basic cells = a thermal node
- In channel layers, $m \times m$ basic cells = a solid thermal node + a liquid one

Heat transfer

- Solid-solid: only consider complete conducting paths
- Solid-liquid: project horizontal heat transfer to vertical direction
- Liquid-liquid: sum heat transfer over multiple channel connections

Problem Formulations

Decision variables

- **Cooling network topology** $\boldsymbol{N}$
- **System pressure drop** P_{sys}

Metrics

- **Pumping power** $W_{pump} = \frac{P_{sys} \cdot Q_{sys}}{\eta}$
  - Q_{sys}: system flow rate; η: efficiency term
- **Thermal gradient** $\Delta T = \max_i(\Delta T_i)$
  - ΔT_i: range of node temperatures in i-th source layer
- **Peak temperature** T_{max}

Design Rules

- TSV positions are at alternating basic cells in both dimensions
- Inlets and outlets can only occur at edges of channel layer
- At most one "continuous" inlet and outlet on each side

Problem 1: Pumping Power Minimization

$$\begin{aligned} &\min W_{pump}, \\ &\text{s.t. } P_{sys} \in \mathbb{R}^+,\ \boldsymbol{N} \in \mathcal{N},\ T_{max} \le T^*_{max},\ \Delta T \le \Delta T^*. \end{aligned} \tag{1}$$

($\mathcal{N}$ is the set of all legal cooling networks)

Problem 2: Thermal Gradient Minimization

$$\begin{aligned} &\min \Delta T, \\ &\text{s.t. } P_{sys} \in \mathbb{R}^+,\ \boldsymbol{N} \in \mathcal{N},\ T_{max} \le T^*_{max},\ W_{pump} \le W^*_{pump}. \end{aligned} \tag{2}$$

General considerations

ΔT is most difficult to handle among all metrics (W_{pump}, ΔT and T_{max})

- W_{pump} vs. T_{max} is a simple trade-off under a specific $\boldsymbol{N}$
- Liquid cooling alleviates T_{max} and worsens ΔT

Three inducing factors for ΔT

1. Temperature rise of coolant
2. Non-uniform power source distribution
3. Non-uniform channel distribution

Factor 3 can be used to compensate for factors 1 & 2

Pumping Power Minimization

The problem is divided into two levels:

- **Inner**: P_{sys} is varied to minimize W_{pump} for a specific $\boldsymbol{N}$, which evaluates $\boldsymbol{N}$
- **Outer**: simulated annealing (SA) searches for a good $\boldsymbol{N}$

Overall Flow of Pumping Power Minimization

Input: $\boldsymbol{N}_{init}$, ΔT^*, T^*_{max}, stack description and floorplan files.
Output: $\boldsymbol{N}$, P_{sys}.

```
N ← N_init;
while #iteration is within the limit do
    Obtain neighboring network solution N';
    W'_pump ← EvaluateNetwork(N', ΔT*, T*_max);
    N ← N' or not according to SA mechanism;
    if W'_pump converges then return N and P_sys;
end while
```

Temperature vs. Pressure

- As P_{sys} increases, T_{max} decreases and finally becomes approximately constant
- $\Delta T = f(P_{sys})$ is either uni-modal or monotonically decreasing

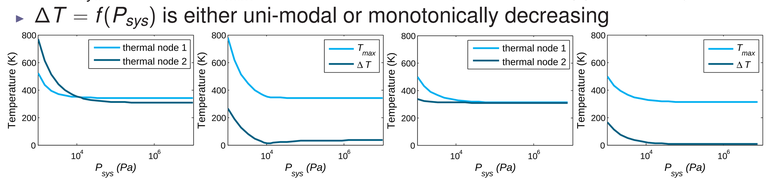

Network Evaluation

- Replace W_{pump} by P_{sys}, as W_{pump} vs. P_{sys} is monotonic for a specific $\boldsymbol{N}$
- Ignore T_{max} first, as it is easier to handle
  - Step 1: solve the problem without constraint T^*_{max}
  - Step 2: check T_{max} and find optimal solution by binary search

Network Evaluation of Pumping Power Minimization

```
function EvaluateNetwork(N, ΔT*, T*_max)
    Minimize W_pump s.t. ΔT ≤ ΔT*;
    if ΔT > ΔT* then
        return +∞;
    else if T_max > T*_max then
        Minimize W_pump s.t. T_max ≤ T*_max;
        if ΔT > ΔT* or T_max > T*_max then
            return +∞;
        else
            return W_pump;
        end if
    else
        return W_pump;
    end if
end function
```

- In step 1, by further substituting $\Delta T = f(P_{sys})$, Problem 1 becomes single-variable:

$$\begin{aligned} &\min P_{sys}, \\ &\text{s.t. } P_{sys} \in \mathbb{R}^+,\ f(P_{sys}) \le \Delta T^*. \end{aligned} \tag{3}$$

- Solve (3) by searching (with three probing points):
  - If a feasible P_{sys} exists, return optimal P_{sys}
  - Otherwise, return the P_{sys} for minimum f (show the nonexistence of feasible P_{sys})

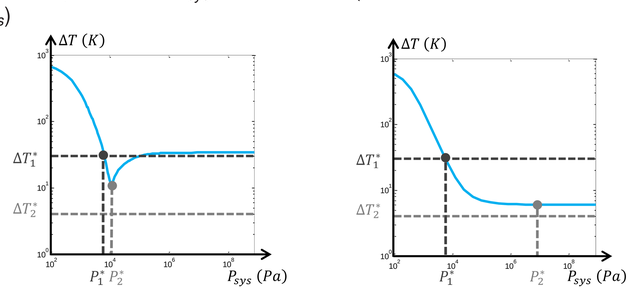

Pumping Power Minimization

Tree-like Cooling Network

Hierarchical tree-like structure is simple and can balance cooling:

- Between upstream and downstream (factor 1)
- Among different trees (factor 2)

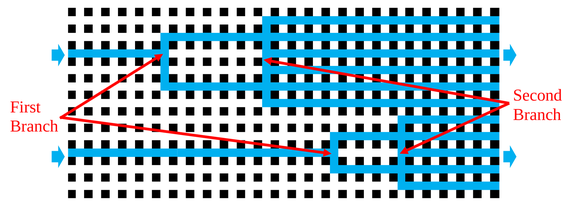

Network Topology Optimization

- In stage 1, ΔT under a **fixed** P_{sys} is used as cost function to accelerate
- In earlier stages, more **rounds** are performed to fully explore solution space
- Eight types of global flow directions are attempted

Stage #	Step Size	Objective Function	Simulator	Runtime for an Iteration
1	10	ΔT	2RM	short
2	10	W'_{pump}	2RM	medium
3	2	W'_{pump}	2RM	medium
4	2	W'_{pump}	4RM	long

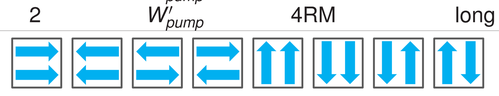

Thermal Gradient Minimization

Similar to solving pumping power minimization with some optimization

Network Evaluation

- Its simplified form becomes:

$$\begin{aligned} &\min f(P_{sys}), \\ &\text{s.t. } P_{sys} \in \mathbb{R}^+,\ P_{sys} \le P^*_{sys}. \end{aligned} \tag{4}$$

- Solving (4) is simpler:
  - If P^*_{sys} locates on falling side of f, it is optimal already
  - Otherwise, adopt golden section search

Network Topology Optimization

Minimizing W_{pump} under a fixed P_{sys} is unrelated to temperature and meaningless, but minimizing ΔT under a fixed P_{sys} is safe $\Longrightarrow$ **speed-up**

- Some iterations are evaluated by one simulation under a fixed P_{sys}
- The original stage 1 is no longer needed
- Another stage with 4RM is affordable to replace the original stage 3

Stage #	Step Size	Objective Function	Simulator	Runtime for an Iteration
1	10	$\Delta T'$	2RM	short
2	10	$\Delta T'$	4RM	medium
3	2	$\Delta T'$	4RM	medium

Experimental Results

Faster 2RM Model

- 5 benchmarks, 40 network samples, 6 thermal cell sizes and 13 pressures
- Tree-like networks, 400μm thermal cells: 0.52% errors (compared to 4RM), runtime reduced from 3.37s to 0.07s

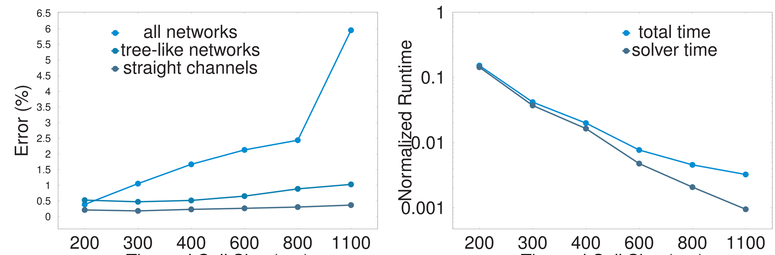

Pumping Power Minimization

- 40 min for cases 1-3 and 240 min for case 4
- **79.61%** better than baseline (unidirectional straight channels)
- **16.35%** better than 1st place in ICCAD 2015 Contest

	Case #	1	2	3	4	5
Baseline	P_{sys} (kPa)	12.98	6.23	7.85	9.71	N/A
	T_{max} (K)	322	314	321	314	N/A
	ΔT (K)	15.0	10.0	15.0	10.0	N/A
	$\boldsymbol{W_{pump}}$ **(mW)**	**10.41**	**6.91**	**8.34**	**11.65**	**N/A**
Manual (1st place in ICCAD Contest)	P_{sys} (kPa)	8.86	5.54	6.98	9.45	40.1
	T_{max} (K)	357	336	328	336	338
	ΔT (K)	15.0	10.0	15.0	10.0	10.0
	$\boldsymbol{W_{pump}}$ **(mW)**	**1.72**	**1.51**	**3.36**	**2.96**	**113.96**
Ours	P_{sys} (kPa)	8.72	5.13	5.81	8.27	40.10
	P_{system} (kPa)	358	336	337	335	338
	ΔT (K)	15.00	10.0	15.0	10.00	10.00
	$\boldsymbol{W_{pump}}$ **(mW)**	**1.66**	**1.37**	**1.90**	**2.68**	**113.96**

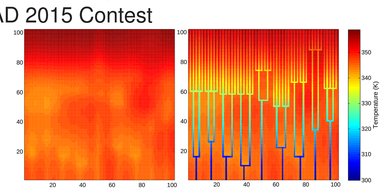

Thermal Gradient Minimization

- Constraint W^*_{pump} on W_{pump} is set to 0.1% of die power
- **37.27%** better than baseline

	Case #	1	2	3	4	5
Baseline	P_{sys} (kPa)	26.08	14.43	17.82	26.51	45.81
	T_{max} (K)	316	309	316	308	338
	W_{pump} (mW)	42.0	37.0	43.0	43.4	148.2
	$\boldsymbol{\Delta T}$ **(K)**	**8.75**	**5.42**	**11.42**	**4.76**	**26.48**
Ours	P_{sys} (kPa)	16.51	8.96	11.46	13.80	40.06
	T_{max} (K)	338	319	327	321	338
	W_{pump} (mW)	5.67	5.66	6.56	4.16	113.80
	$\boldsymbol{\Delta T}$ **(K)**	**5.54**	**3.81**	**7.12**	**3.87**	**9.64**

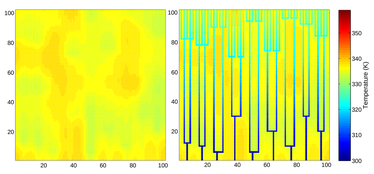

E-Mail: gjchen@cse.cuhk.edu.hk