arpa·e
CHANGING WHAT'S POSSIBLE

HOLISTIC CO-DESIGN OF NOVEL HYBRID COOLING TECHNOLOGY FOR THE DATA CENTER OF THE FUTURE

COOLERCHIPS™

Dereje Agonafer, University of Texas Arlington
Team Members: University of Maryland, University of Illinois Urbana-Champagne, Illinois Institute of Technology, Consultant.

Project Vision

- We are offering a transformative and extendable cooling technology for data centers of the future utilizing a hybrid cooling approach for servers that combines novel evaporative cooling technique for the primary high-powered components (CPUs & GPUs) and heatsink design for the secondary high-power component (DIMMs) that will allow for rack powers much higher than COOLERCHIPS targets.
- This hybrid cooling approach will use less cooling power than existing solutions and can be deployed with minimal alteration to existing infrastructure.

Total Project Cost:	$2.8M
Length	36 mo.

COOLERCHIPS Kickoff Meeting
October 18 & 19, 2023

Brief COOLERCHIPS Project Overview

Fed. funding:	$2.8M
Length	36 mo.

Team member	Location	Role in project, core competencies
University of Texas Arlington	Arlington, Texas	Lead, Optimized heatsink design for secondary component cluster, rack level manifolds, DCEC Testing – both at the server / Rack Level with existing TIMS and Novel UIUC TIMS.
University of Maryland	College Park, Maryland	DCEC – manufacturing, server level manifolds, and testing
University of Illinois Urbana-Champagne	Urbana-Champagne, Illinois	TIMs –Performance characterization, Integration with DCEC and optimized heatsinks.
Illinois Institute of Technology	Chicago, Illinois	ML models for next level CFD and active flow control .
Consultant	Poughkeepsie, New York	System level cooling and Technology-to-Market Lead.

Context/history of the project

- Original team had expertise in evaporative cooling, TIM material, system / component level direct to chip cooling and additively manufactured heatsinks with a potential to be scalable. The team was expanded to include expert from IIT in “Physics-Based Modeling and Neural Network Model Development.” Also, Dr. Schmidt, IBM Fellow Emeritus and NAE Member (consultant) will help us in rack level liquid cooling, system availability goals and Technology-to-Market Approach.
- The confidence level high - favorable responses received from Dr. Saket Karajgikar, Data Center R&D Engineer – Technical Strategy, META, who will provide an L11 Rack and Dr. Ashish Gupta, Director of Thermal-Mechanical, Data Platforms Group, Intel Corporation, Portland, OR, has agreed to provide guidance from an industry perspective.

Concept Detail:

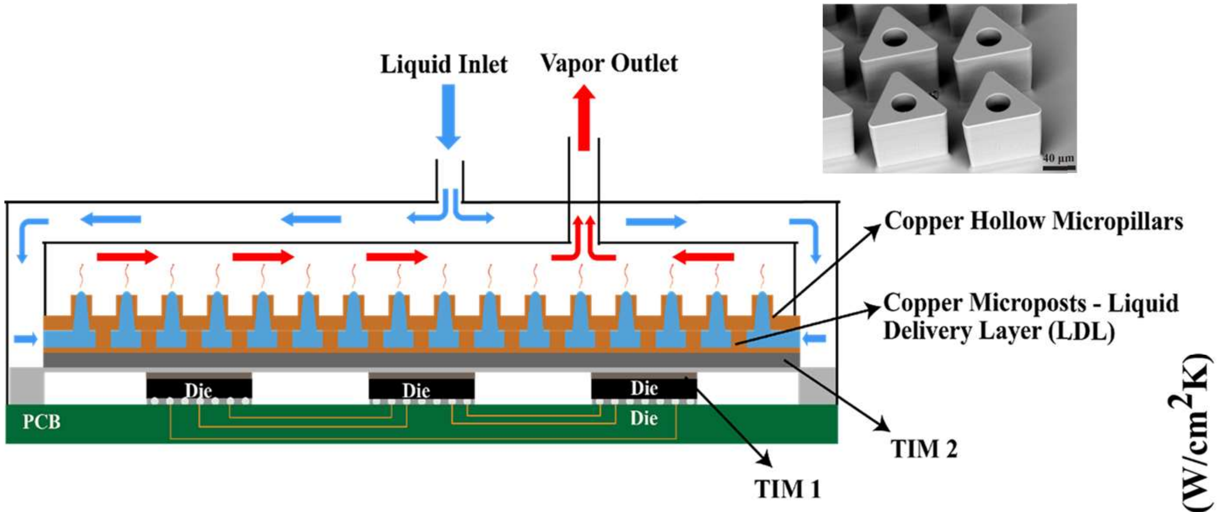

- Hollow micropillars with non-axisymmetric shapes (triangle etc.) are employed to form non-spherical microdroplets for enhanced evaporative cooling.
- Evaporator consists an array of microposts for delivering coolant from manifold to hollow micropillars.

$$R_{case\text{-}to\text{-}coolant} = R_{TIM} + R_{sub} + R_{LDL} + R_{evaporator} \leq 0.01\ ^{o}C/W$$

Heat Transfer Mode	Source	Thermal Resistance (cm² K/W)
Air Cooling	Joo et al.	11.67
Single Phase (refrigerant)	Joshi et al.	0.384
Single Phase (water)	Brunchwiller et al.	0.09
Flow Boiling	Joshi et al.	0.167
Evaporative Cooling (refrigerant)	Damena Agonafer et al.	0.04
Evaporating Cooling (Water)	Damena Agonafer et al.	0.012

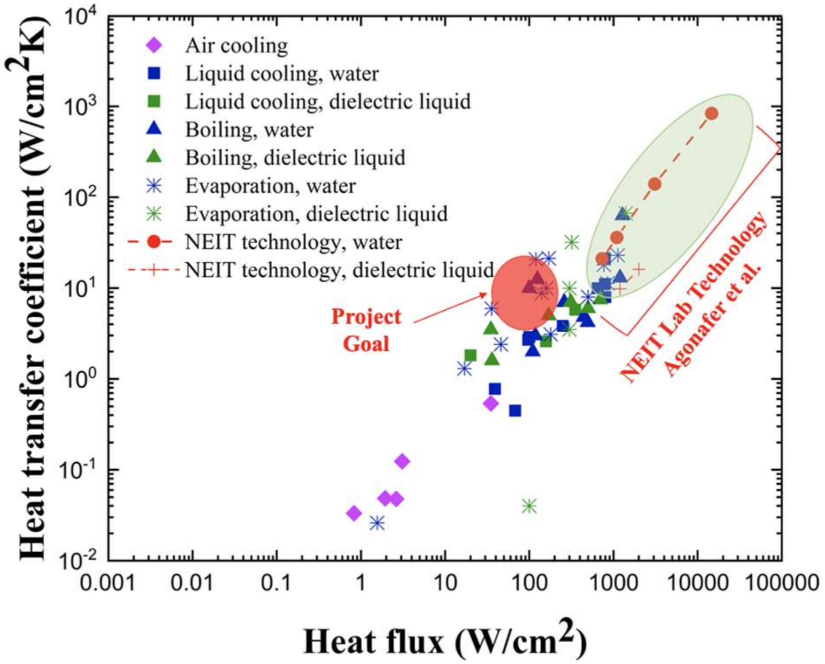

Concept Detail:

- Direct-to-chip evaporative cooler (DCEC) integrated on top of an integrated heat spreader (case).
- Coolant is actively pumped to the DCEC and evaporates resulting in an exit vapor quality of 1.
- The flow line to the DCEC is fitted with filters and flow valves to prevent clogging in the DCEC and to achieve stable thin-film evaporation.
- **Targeted $R_{case\text{-}to\text{-}ambient} \le 0.01$ °C/W**

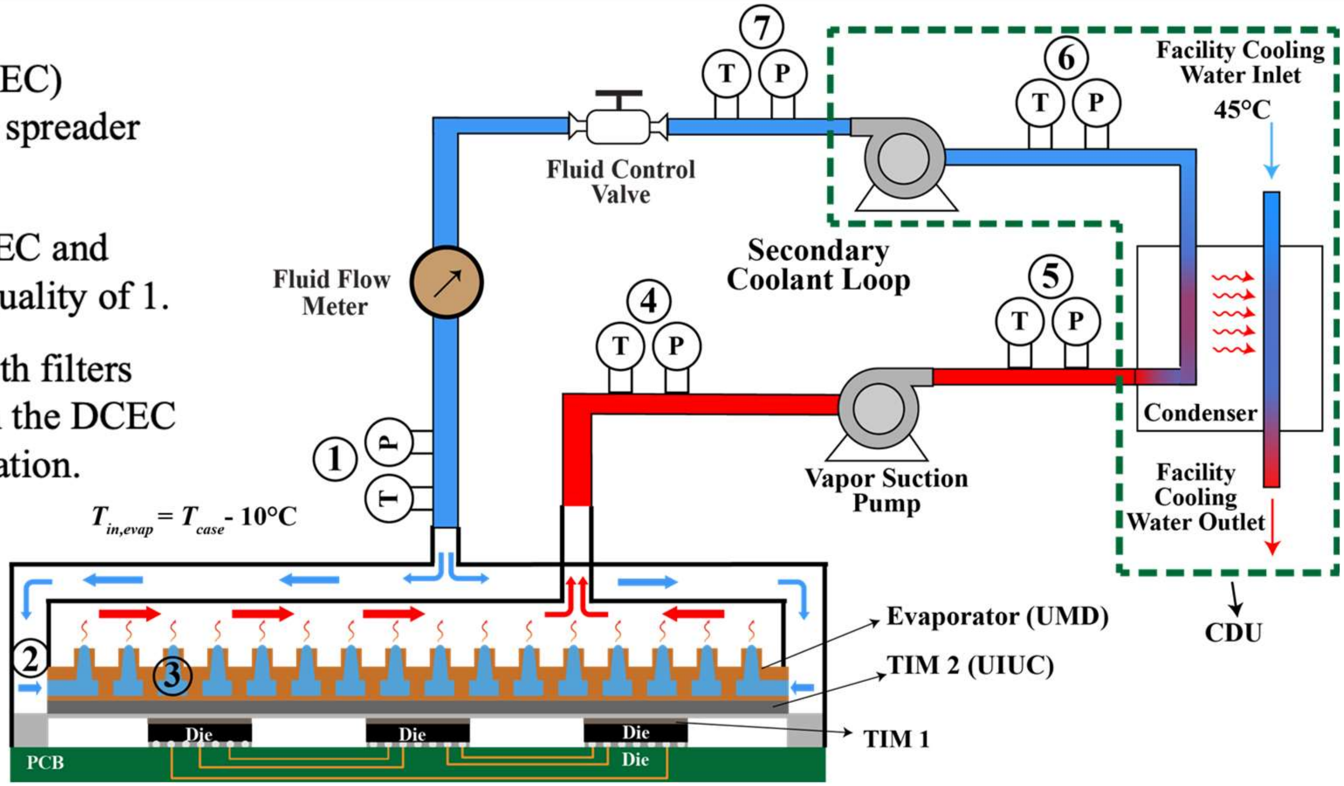

CHANGING WHAT'S POSSIBLE

Concept Detail:

- *Reliable Thermal Interface Materials*

1)

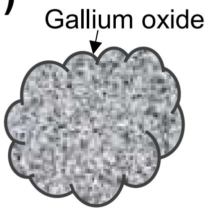

Strategy 1
: suppressing
oxidation

2)

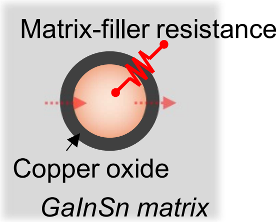

Strategy 2
: reducing
matrix-filler
thermal resistance

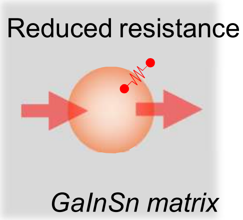

3)

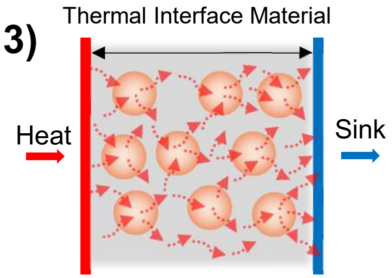

Strategy 3
: forming additional
heat-paths
within the matrix

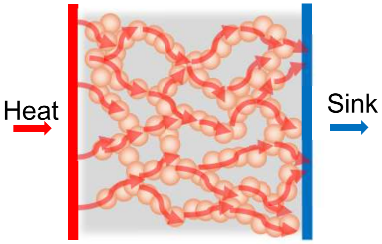

International Journal of Heat and Mass Transfer **170**, 121012, 2023

- High thermal conductivity > 70 W/(m-K)
- High reliability
- High wetting characteristics
- **Goal: < 0.004 K/W**

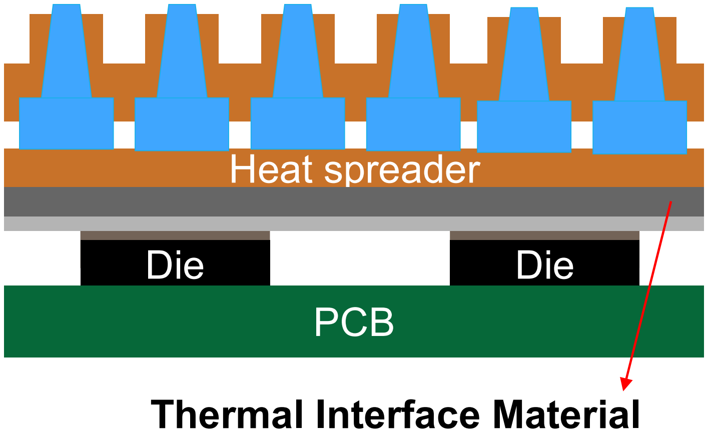

FOA Metrics	Units
Chip-to-Coolant Resistance Target	**< 0.01 K/W**
Chip Power	320 W/cm^2 for 6.25 cm^2
Rack power density	109.78 kW/m^3 (Based on 210kW rack power, ORV3 Rack (1.41 m^3), and RDHx including fans (0.503 m^3))
Rack size in demonstration	42U ORV3 with L11 rack integration
Server power per U in demonstration	Up to 5 kW
Rack power in demonstration	Up to 210 kW

Concept Detail:

- *Reliable Thermal Interface Materials*

- Galium-based TIM incorporting Cu NPs

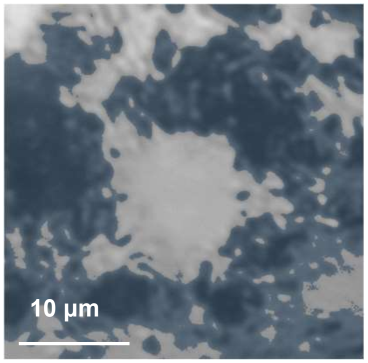

Nanoclusters in liquid metal

- Si-Cu interface filled with Pb/Sn solder

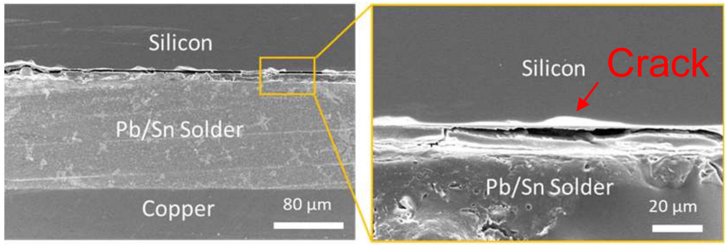

- Si-Cu interface filled with GaInSn/Cu NPs

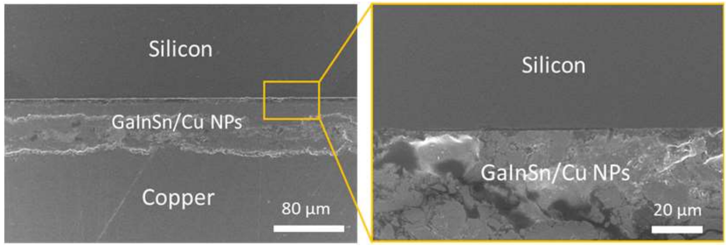

International Journal of Heat and Mass Transfer **170**, 121012, 2023

- High thermal conductivity > 70 W/(m-K)
- High reliability
- High wetting characteristics
- **Goal: < 0.004 K/W**

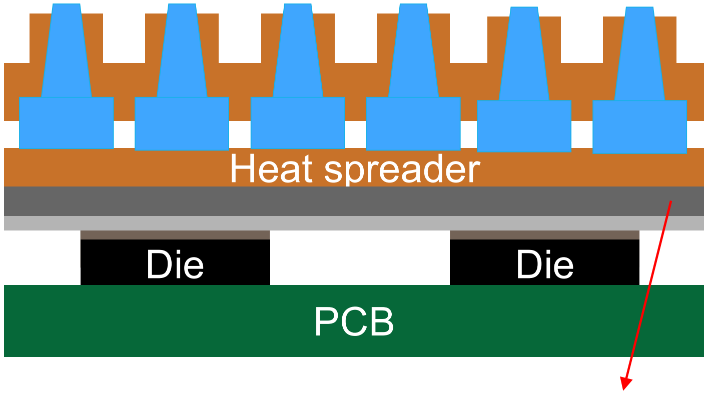

Thermal Interface Material

FOA Metrics	Units
Chip-to-Coolant Resistance Target	**< 0.01 K/W**
Chip Power	320 W/cm^2 for 6.25 cm^2
Rack power density	109.78 kW/m^3 (Based on 210kW rack power, ORV3 Rack (1.41 m^3), and RDHx including fans (0.503 m^3))
Rack size in demonstration	42U ORV3 with L11 rack integration
Server power per U in demonstration	Up to 5 kW
Rack power in demonstration	Up to 210 kW

Concept Detail:

- *Reliable Thermal Interface Materials*

- Thermal cycling

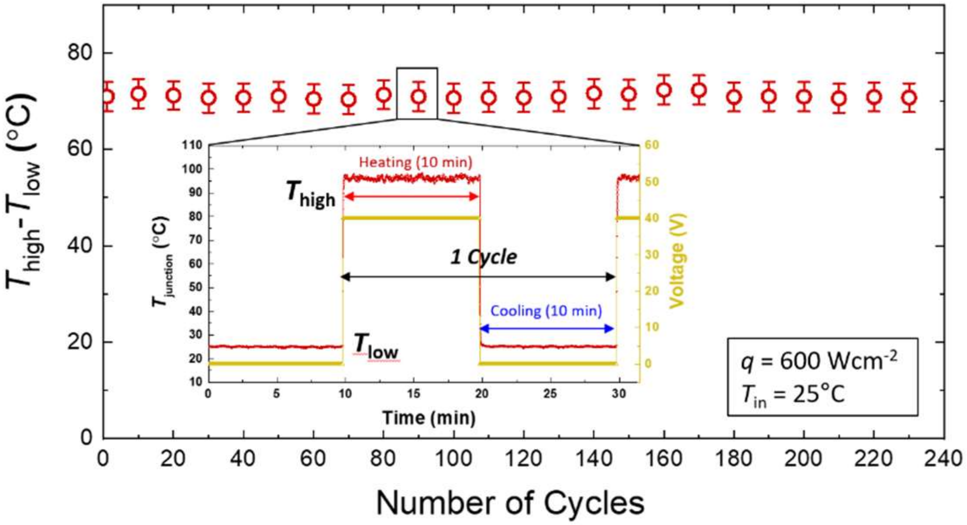

International Journal of Heat and Mass Transfer **170**, 121012, 2023

- Microcrack issues due to mismatch of thermal expansion eliminated due to liquid-phase TIM

- Comparision with comercially available TIMs

Type	Supplier	Material	Thermal conductivity (W/(m-K))	Melting point
Thermal grease	Laird	Silicone-based	0.7	·
	Laird	Zinc Oxide-based	0.92	·
	Laird	Zinc Oxide-based	5	·
Sheet	Smart High Tech	Graphene-based	50-90	·
	Carbice	Al + CNT	12	·
Liquid-metals	Indium Corporation	GaInSn (62.5/21.5/16)	16.5	17°C
		GaInSn (66.5/20.5/13)	16.5	15.7°C
		GaIn (78.6/21.4)	21	15.7°C
Low-melting point metals	RotoMetals	Field's metal (32.5Bi/16.5Sn/51In)	~19	62.0°C
		Rose's metal (50Bi/25Pb/51In)	16.3	98.0°C
		Pure Indium	83	157°C

Concept Detail

Optimized Heatsink Design for High-Power Secondary Component Clusters (DIMMs)

- Source of power usage in next gen servers will be DIMM clusters.
- DDR5 DIMMs TDP expectation – >20 W each
- Assuming 25 W/DIMM makes DIMM cooling non-trivial.
    - 16 DIMMs system power usage expectation – 400 W / server.
    - 32 DIMMs system power usage expectation – 800 W / server.

- Multi - Design Variable and Objective Function optimization approach to create a heatsink for DIMM cluster (4 -8 DIMM form factor).
- Design conceptualized with and without heat-pipes.
- Complements novel evaporative cooling technique and TIM to reduce IT load to < 4%.

2U DIMM Heat spreader with heat pipe enhancement

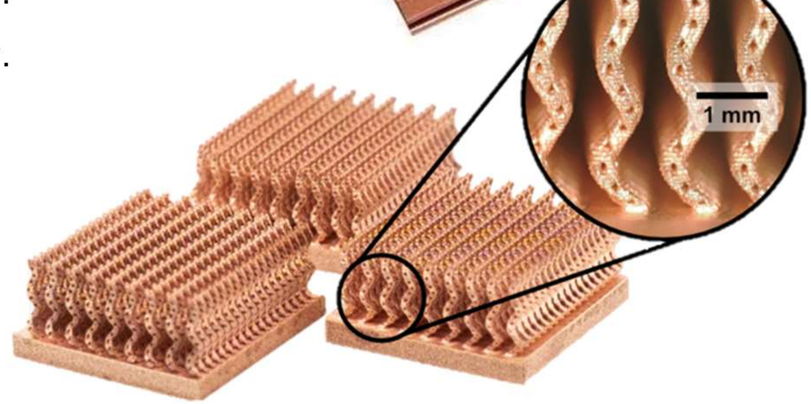

Samples – Fabric-8 additively manufactured heat sink

DIMM heat spreaders are designed to attach vertically on the chip whereas, heatsinks for associated electronic chips in the motherboard/server remain horizontal. This opens the potential for designing 3D additively manufactured heat sink/enhanced heat spreaders

Hierarchical Modeling-Concept Details

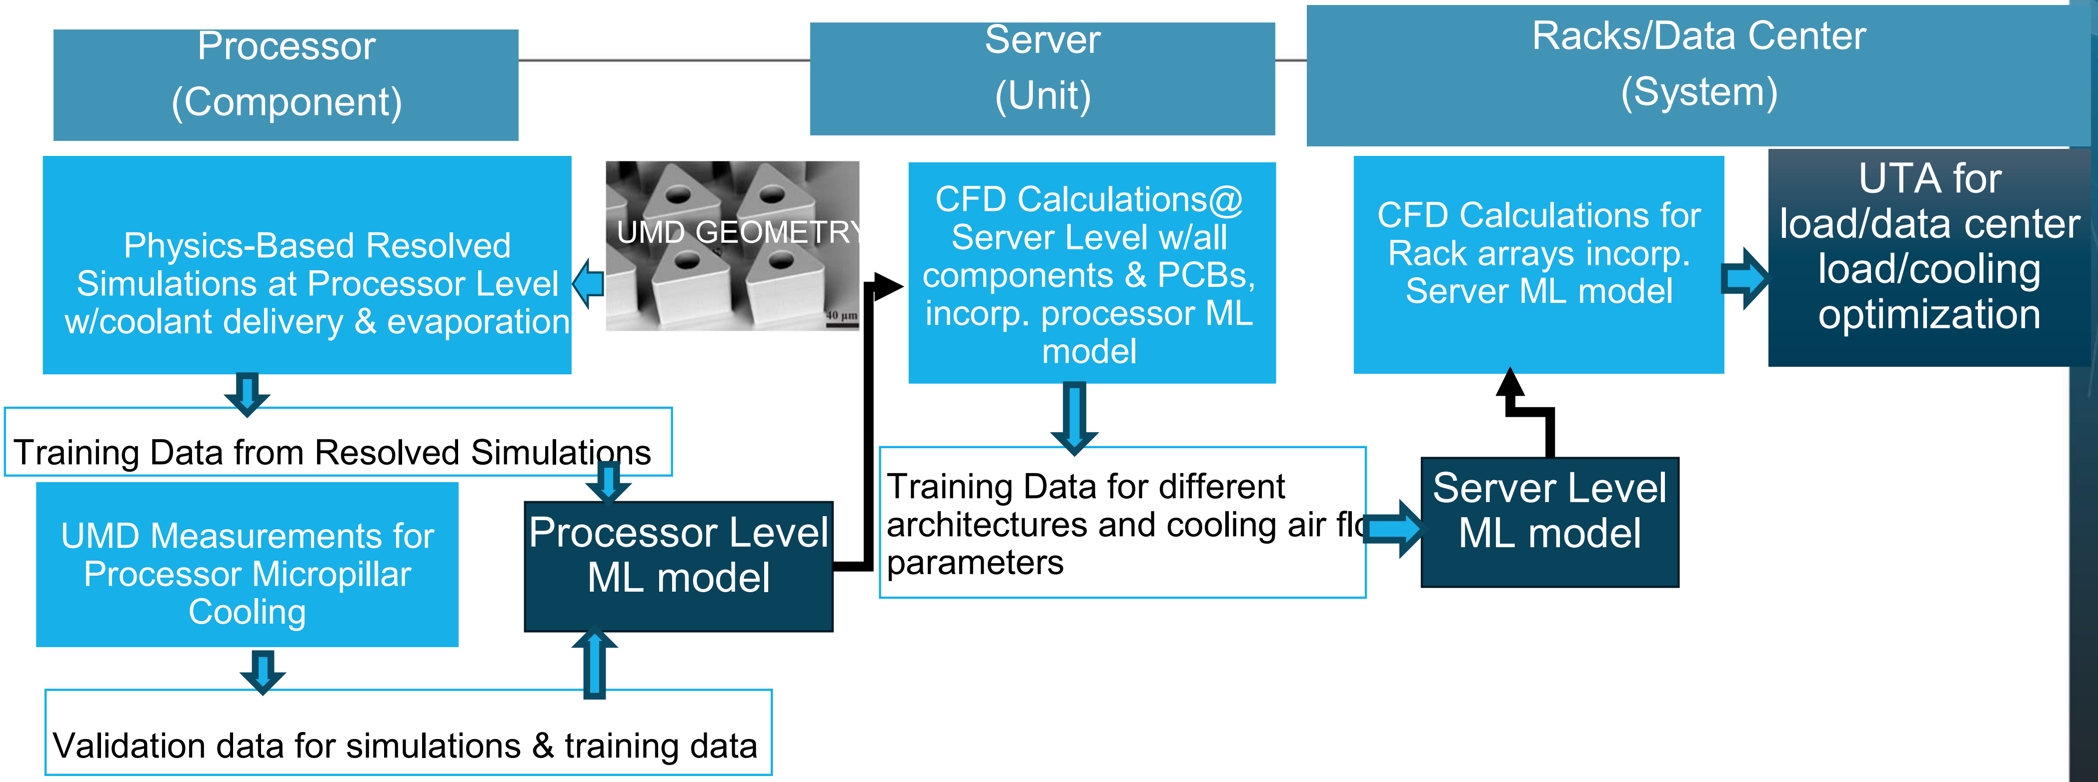

Processor (Component)

Server (Unit)

Racks/Data Center (System)

Physics-Based Resolved Simulations at Processor Level w/coolant delivery & evaporation

UMD GEOMETRY

CFD Calculations@ Server Level w/all components & PCBs, incorp. processor ML model

CFD Calculations for Rack arrays incorp. Server ML model

UTA for load/data center load/cooling optimization

Training Data from Resolved Simulations

UMD Measurements for Processor Micropillar Cooling

Processor Level ML model

Training Data for different architectures and cooling air flo parameters

Server Level ML model

Validation data for simulations & training data

- Hierarchical physics-based modeling, simulation and measurement training data sets, to develop ML models that are used as inputs to enable CFD calculations at the next scale
- Close collaborations between IIT, UMD & UTA on simulations, validation and optimization

Concept Detail

Rack Level Testing and Optimization of Hybrid Cooling Approach

- Showcase scalability of hybrid cooling approach by retrofitting an 42U ORV3 with L11 rack integration.

- Showcase passive and active flow control techniques.
- Active flow control achieved by training Machine learning algorithm.

- Showcase reliability of Hybrid cooling approach – at least 2000 hours of testing and reliability studies.

- Reliability studies - focus on processor transient powers from 0 to 1000 watts with multiple processors all at different power levels; showcase flow and temperature control features for each processor within a rack using one liquid pump and one evaporator pump per rack; maximize cooling reliability by minimizing no. of flow control components.
- Specific reliability tests to include -
  - Thermal cycling/shock
  - Vibration testing
  - Burn-In Testing
  - Thermal Performance Testing

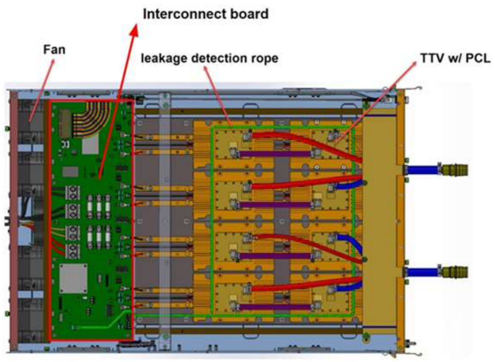

L11 rack and associated servers acquired from META for COOLERCHIPS project period.

Task Outline & Technical Objectives

Tasks	Member	Starting Quarter	Ending Quarter
Task 1: Copper heat spreader development and integration	UIUC	Q1	Q10
Task 2: Copper wick development and integration into evaporator	UIUC/UMD	Q1	Q9
Task 3: Design and fabrication direct-to-chip evaporator	UMD	Q1	Q9
Task 4: System level design and testing of evaporator	UMD	Q1	Q9
Task 5: Physics-Based Modeling at the Chip and Server Levels and Neural Network Model Development	IIT/UTA	Q1	Q12
Task 6: Optimized air cooled heat sinks	UTA	Q1	Q5
Task 7: Hybrid cooling at server/rack level	UTA	Q6	Q12
Task 8:Technology to market strategy	All Members + Consultant*	Q1	Q12

* Roger Schmidt, PhD, IBM Fellow Emeritus, Member NAE

Task Outline & Technical Objectives

- *Technical Objectives*

1. Thermal Resistance, chip surface-to- facility coolant, of ≤ 0.01 K/W
2. Total cooling power of secondary loop including all ancillary equipment (CDUs, pumps, heat exchangers, etc.) ≤ 2 % IT load, doing better than the secondary loop total cooling energy target of ≤ 3% IT load.
3. Total cooling power to reject all heat to ambient ≤ 4 % IT load, doing better than the system-wide total cooling energy target of ≤ 5% IT load. (Estimated by simulation) (at power density ≥ 20kW/m^3 and ambient conditions specified in Table 1).
4. Support for >210 kW of rack power and enabling ASHRAE-A3 class servers (up to 40°C inlet).
5. System Availability > 99.982% for economic life of the system.

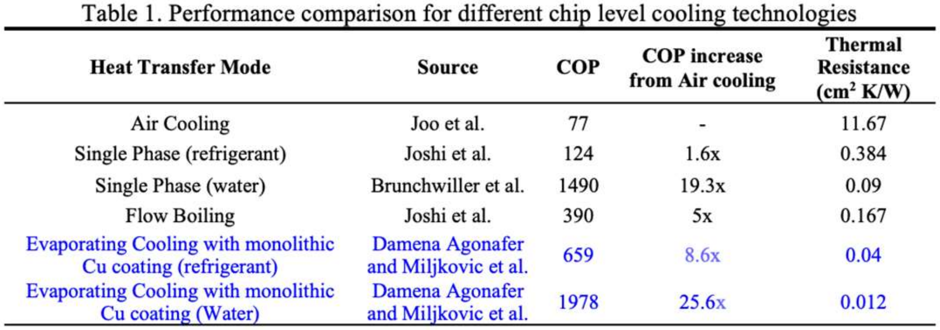

Table 1. Performance comparison for different chip level cooling technologies

Heat Transfer Mode	Source	COP	COP increase from Air cooling	Thermal Resistance (cm^2 K/W)
Air Cooling	Joo et al.	77	-	11.67
Single Phase (refrigerant)	Joshi et al.	124	1.6x	0.384
Single Phase (water)	Brunchwiller et al.	1490	19.3x	0.09
Flow Boiling	Joshi et al.	390	5x	0.167
Evaporating Cooling with monolithic Cu coating (refrigerant)	Damena Agonafer and Miljkovic et al.	659	8.6x	0.04
Evaporating Cooling with monolithic Cu coating (Water)	Damena Agonafer and Miljkovic et al.	1978	25.6x	0.012

Challenges and Risks

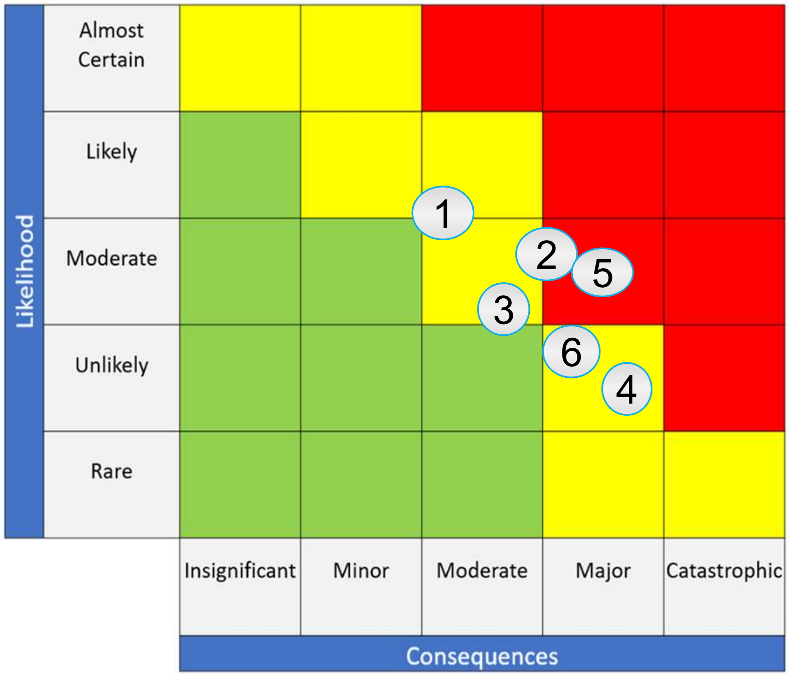

Risk Status

Risk	#
Technology does not meet the FOA target for thermal resistance due to insufficient thermal performance from the TIM prototype design	1
Photolithography templated electrochemical deposition cannot manufacture the array of hollow micropillars – including liquid delivery layer	2
Failure to address cooling energy target of secondary components (PSU, DIMMs, SSD.) with Optimized 3D printed heatsinks at the server level	3
Failure to meet the cooling energy target at the rack level.	4
Technology does not meet the FOA target for thermal resistance due to device flooding, clogging, and manufacturing difficulties.	5
System level- TIM and module testing and refrigerant leakage.	6

Technology-to-Market Approach

- Step one – publish the results of the design and testing that verifies our deliverables. And the main item will be to show that the design is very reliable and proven after many hours of testing.
- Step two – once the industry has shown an interest then possible suppliers can be sought after or they will come forward with an interest to purchase the patent technologies; specific manufacturers and manufacturing techniques will need to be researched; these activities will ultimately drive the business model
- Step 3 – based on feedback from industry and their needs follow on investments may need to be made based on the current needs, unique manufacturing processes required, etc.
- Step 4 – since this is design it focused on high power chip cooling the market will most likely focus on those markets which will utilize these high-power chips – AI, scientific computing, and analytics
- Step 5 – cost and performance targets will need to be developed as the design progresses.

T2M Lead – Dr. Roger Schmidt, IBM Fellow Emeritus & NAE Member.

Needs and Potential Partnerships

Potential partnerships, they have agreed to help as needed

- Intel
- META
- Fabric8 – Additive manufacturing of heatsinks for air cooled components.
- Chemours – Supply appropriate coolant as needed.

Q & A

U.S. DEPARTMENT OF ENERGY

https://arpa-e.energy.gov