Performance of Passive 2-Phase Immersion Cooling of Server Hardware

Benson Chan and How Lin
Endicott Interconnect
Endicott, NY, USA
benson.chan@eitny.com, how.lin@eitny.com

Stephen Polzer, Nathan Harff, Erik Daniel and Barry Gilbert
Mayo Clinic
Rochester, MN USA
polzer.stephen@mayo.edu
harff.nathan@mayo.edu

Phillip Tuma
3M Electronic Markets Materials Division
St. Paul, MN, USA
petuma@mmm.com

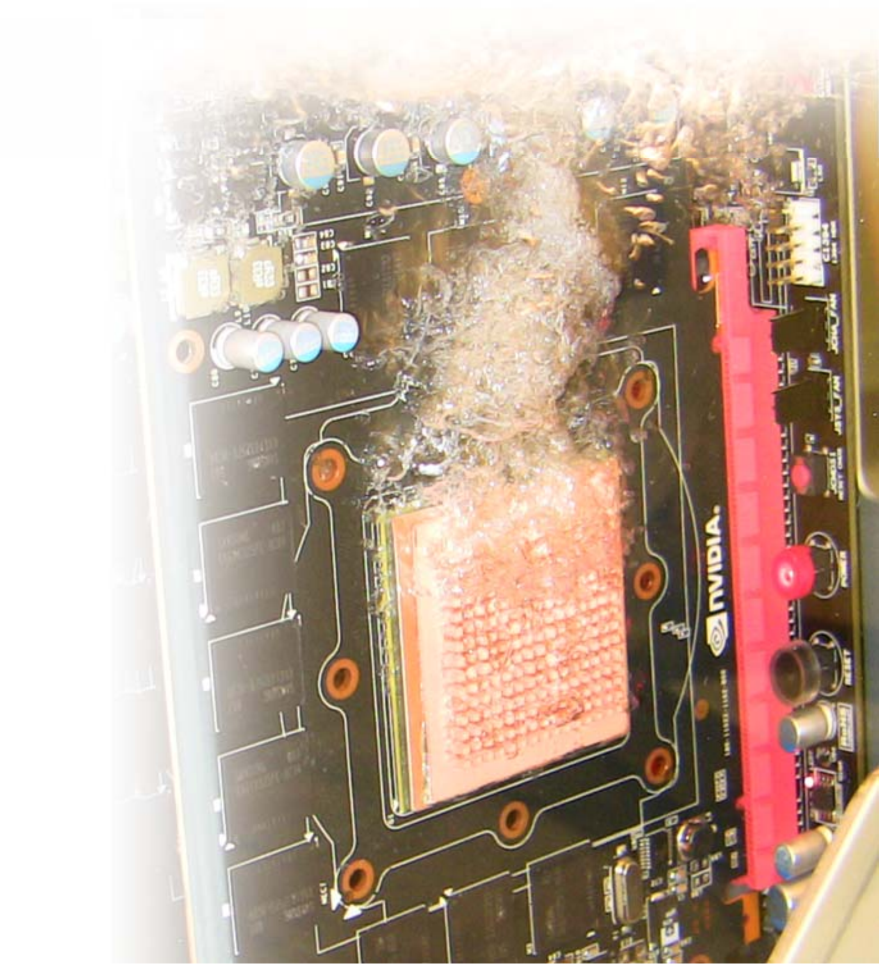

IMAPs ATW on Thermal Management, Palo Alto, Ca, USA, Sept 28-30, 2010

Outline

I. INTRODUCTION
  A. Motivation
    1. Inadequacy of Air Cooling
    2. Traditional Pumped Liquid Cooling
    3. Traditional Immersion Cooling
  B. Open Bath Immersion (OBI) Cooling Concept
    1. Inspiration
    2. Concept Overview
    3. Advantages
    4. Working Fluids

II. SIGNAL INTEGRITY DATA
  A. Motivation
  B. Procedure and Results
    1. Electrical Data
    2. Optical Data

III. 1.8kW OBI-COOLED DESKTOP
  A. Motivation
  B. Procedure and Results

IV. 350W OBI-COOLED APPLE G5

V. CONCLUSIONS

VI. FUTURE WORK

I. INTRODUCTION

A. Motivation

1. Inadequacy of Air Cooling

- Power density and efficiency are fundamentally limited by need to provide space for airflow both at the server chassis and rack or facility level.
- Consumes natural resources and real estate for fans, blowers, heatsinks, PCB, etc.
- Lacks the ability to capture and utilize waste heat because the thermodynamic availability of the heat is low

2. Traditional Pumped Liquid Cooling

- Advantages
  - Power density is greatly increased by elimination of airflow paths.
  - Energy efficiency is high because liquid and device temperatures are tightly coupled. This enables "warm" cooling techniques.
  - High temperature liquid stream permits distribution and utilization of waste heat
- Disadvantages
  - It is inherently costly and complex to build leak proof networks of pumps, valves, hoses, manifold, cold plates, fittings, etc.
  - Power density at chassis level is limited by need to place plumbing

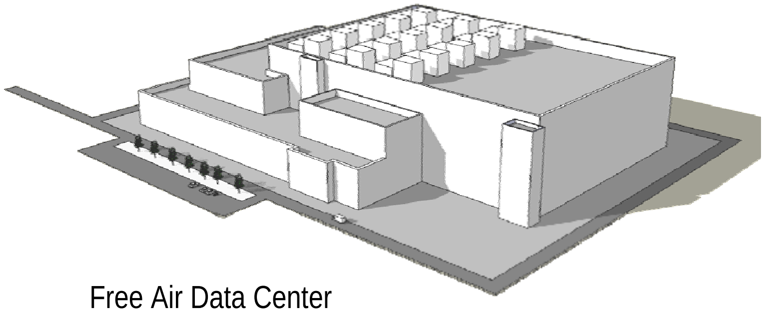

Free Air Data Center

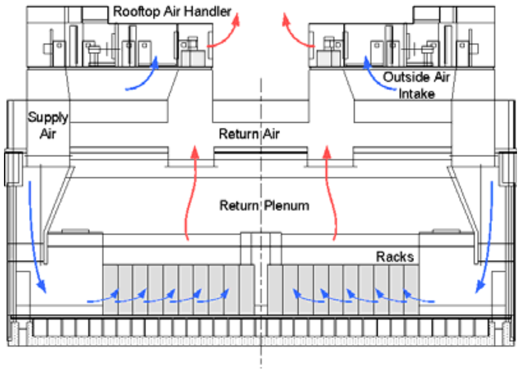

Water-Cooled IBM BladeCenter HS22

Credit: IBM Research – Zurich

A. Motivation (cont.)

3. Traditional Immersion Cooling

- Advantages
  - An elegant way to capture ALL heat generated by a complex and very dense electronics assembly
- Disadvantages
  - It is inherently costly and complex to build hermetic electronic enclosures particularly when many conductors must penetrate that enclosure.

Spray Cooled Military Enclosure
(photo courtesy of Spraycool)

B. Open Bath Immersion Cooling Concept [1]

1. Inspiration

- Vapor Degreaser
  - Ubiquitous machines used to clean parts, PCBs, etc. using a refluxing solvent
  - These use very little solvent despite the fact that they are ostensibly wide open 3 shifts per day 5 days a week with parts moving in and out constantly
  - Parts leave dry

Photo courtesy
of Forward Technology

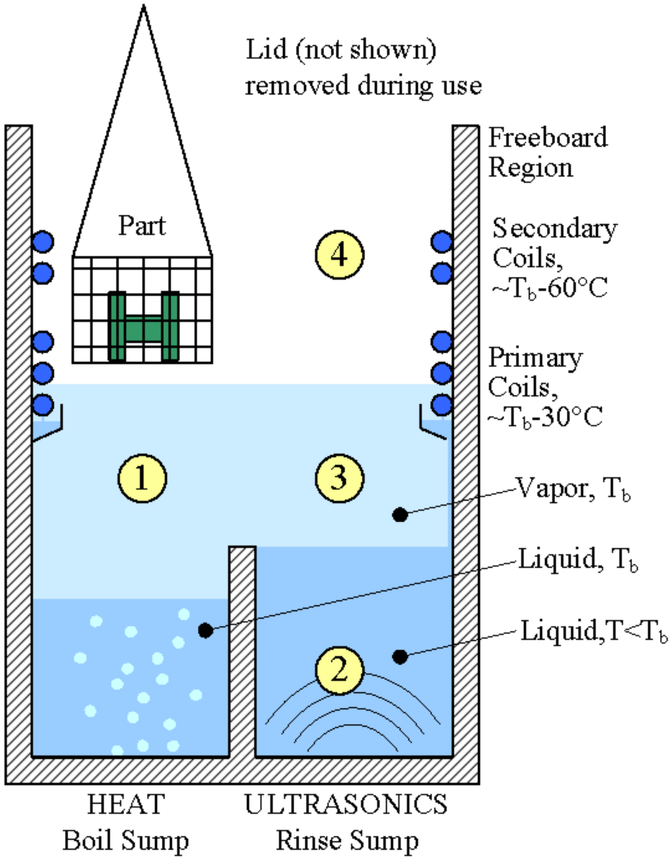

Vapor Degreaser and Cleaning Process

B. Open Bath Immersion Cooling Concept [1]

2. Concept Overview

- Servers are placed side-by-side in a semi-open bath of dielectric fluid
- Devices cause fluid to boil
- Rising vapor condenses on water-cooled coil
- Servers plug into an immersed backplane
- Power and IO enter/exit through a simple conduit that terminate below the liquid level
- Servers can be hot-swapped and leave the bath dry
- The bath is “semi-open” as it is at atmospheric pressure. Bath breathes through a trap

3. Advantages

- No fans, heatsinks, cold plates, pumps, quick disconnects (QDs), manifolds, hoses, clamshells, hermetic connectors, controls, etc.
- Fluid leakage at one easily regulated point
- No temperature glide between devices
- Eliminates advective resistance, same Tj on all devices
- Intrinsic fire protection
- Server power densities as high as 4KW/liter are possible
- **Performance, Efficiency, Density, Cost, Complexity**

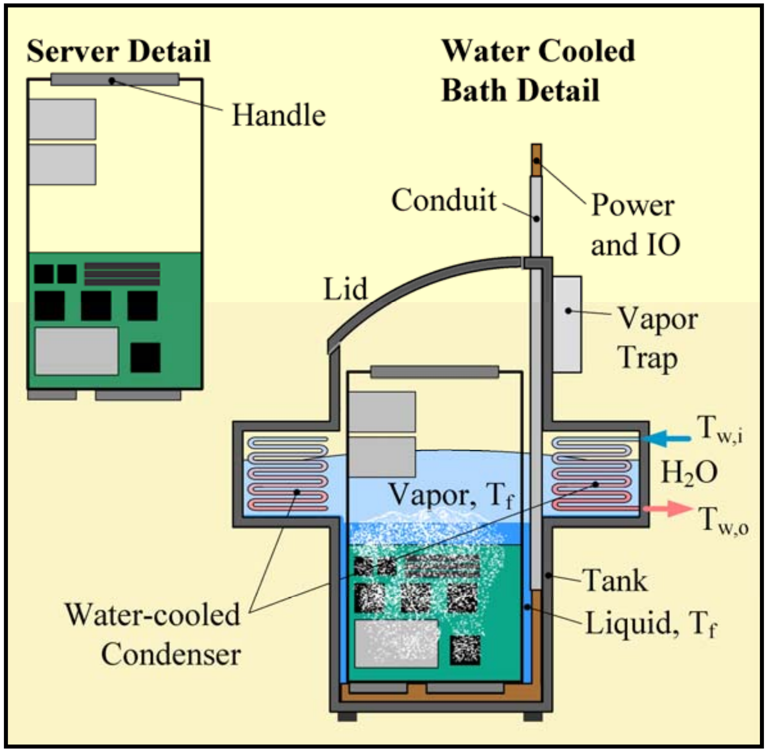

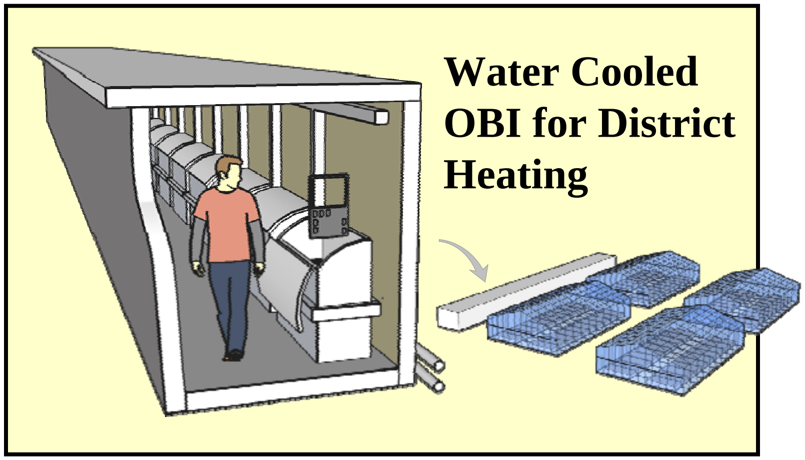

Efficiency and Performance

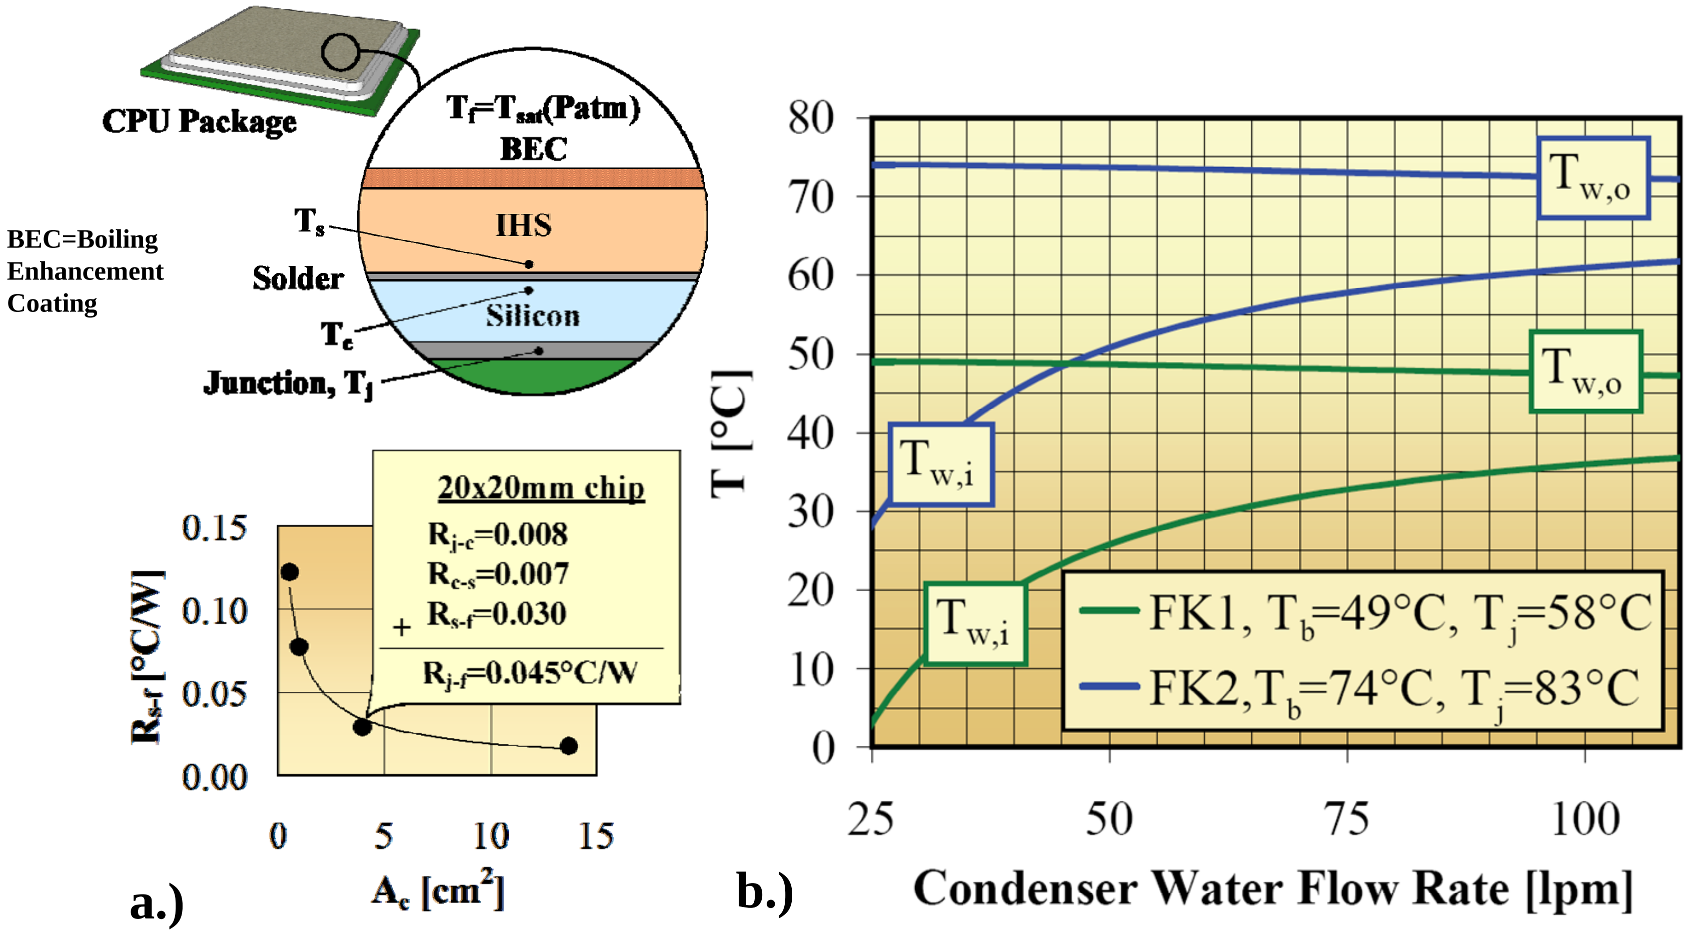

a.) Experimental sink-to-fluid performance as a function of die size and calculated junction-to-fluid performance for a 20x20mm die in a hydrofluoroether (HFE) liquid and b.) calculated condenser water temperatures as a function of flow rate for an 80kW bath for two fluoroketone bath fluids ($0.055>R_{jw,o}>0.045$°C/W)

Power Density

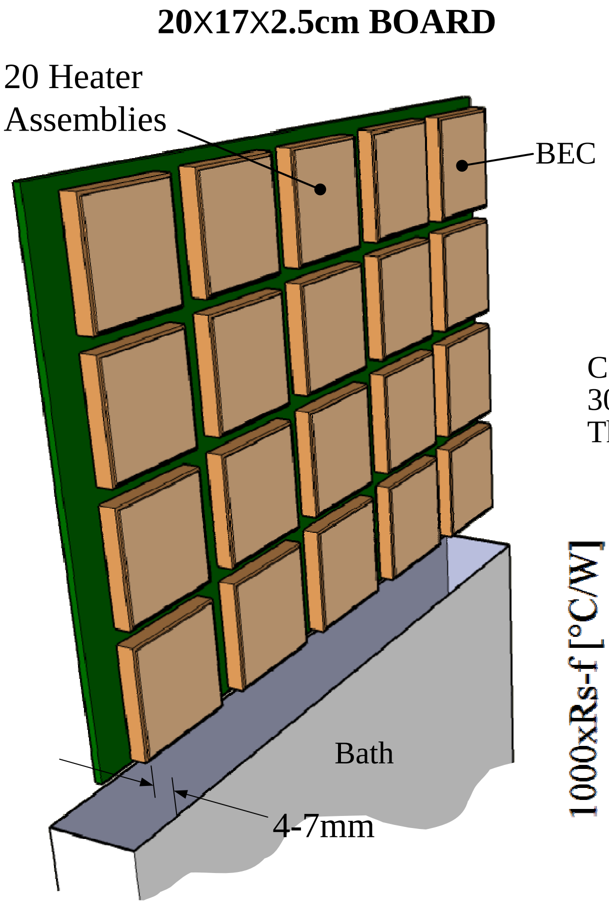

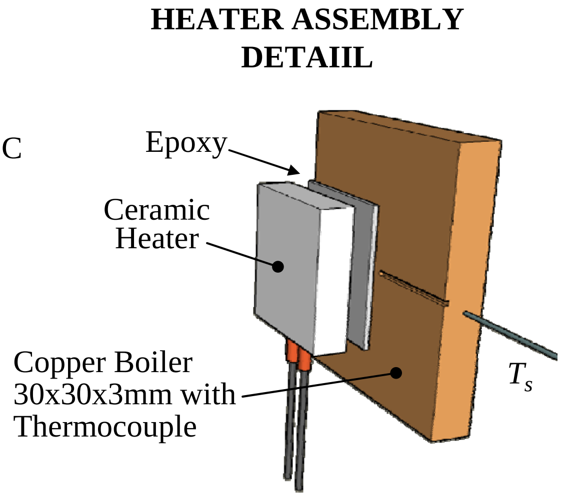

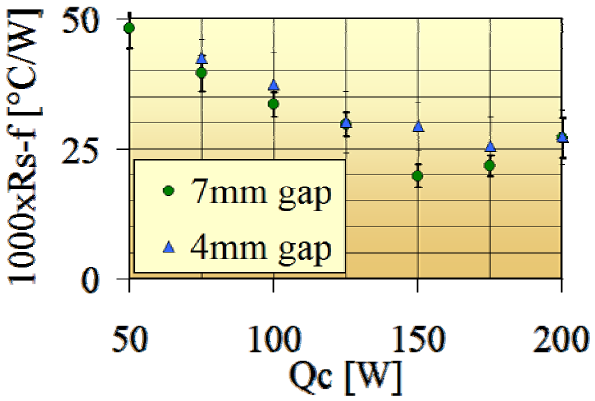

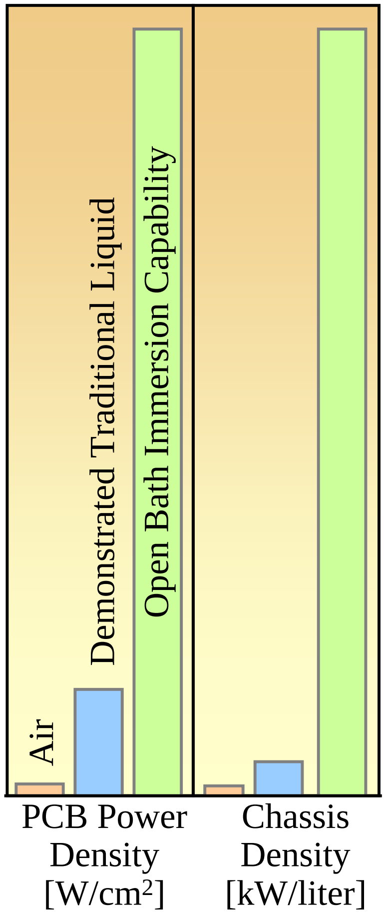

Experiment to demonstrate PCB and chassis level power density capabilities of OBI cooling. Comparison of PCB and chassis power density for OBI compared with typical air and commercial pumped liquid technologies.

4kW in 1 liter with 200cc of fluid

Comparison of projected costs associated with a server's thermal management solution as a function of life cycle

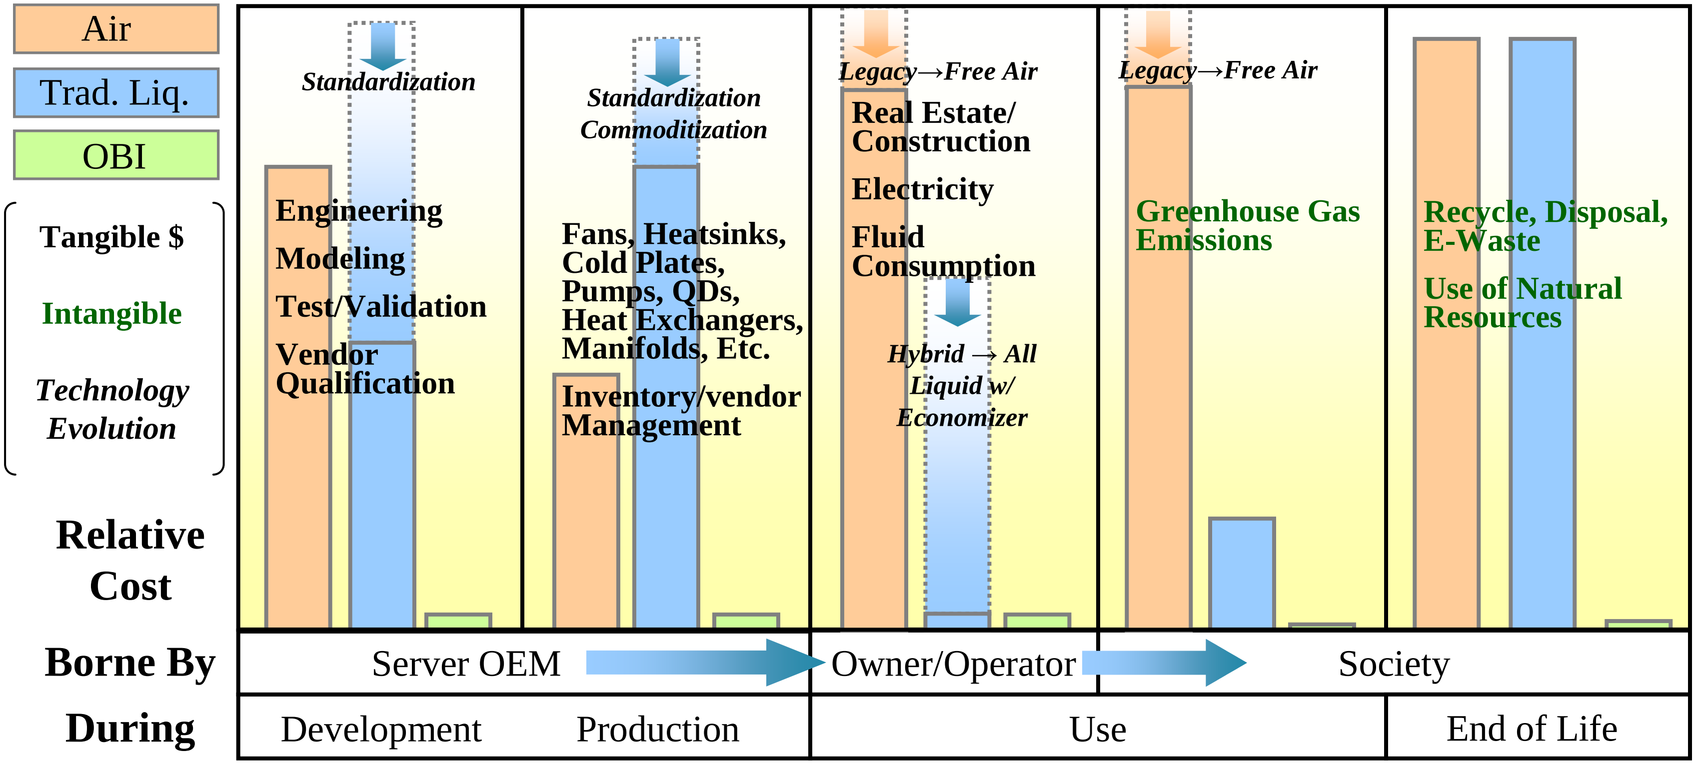

B.Open Bath Immersion Cooling Concept (cont.)

3. Working Fluids

- Immersion-cooled computer systems have historically used perfluorocarbon (PFC) fluids
  - Though Non-ozone depleting, PFCs have high Global Warming Potential (GWP)
  - Their use is inconsistent with a growing trend to use materials with low environmental impact
- Next generation Fluids
  - Segregated hydrofluoroethers (HFEs)
    - May lack requisite dielectric properties for some applications
  - Hydrofluorocarbons (HFCs)
    - Tend to have high GWPs
    - Electrical properties similar to HFEs
  - Fluoroketones (FKs)
    - C6F12O boils at 49°C and very low GWP

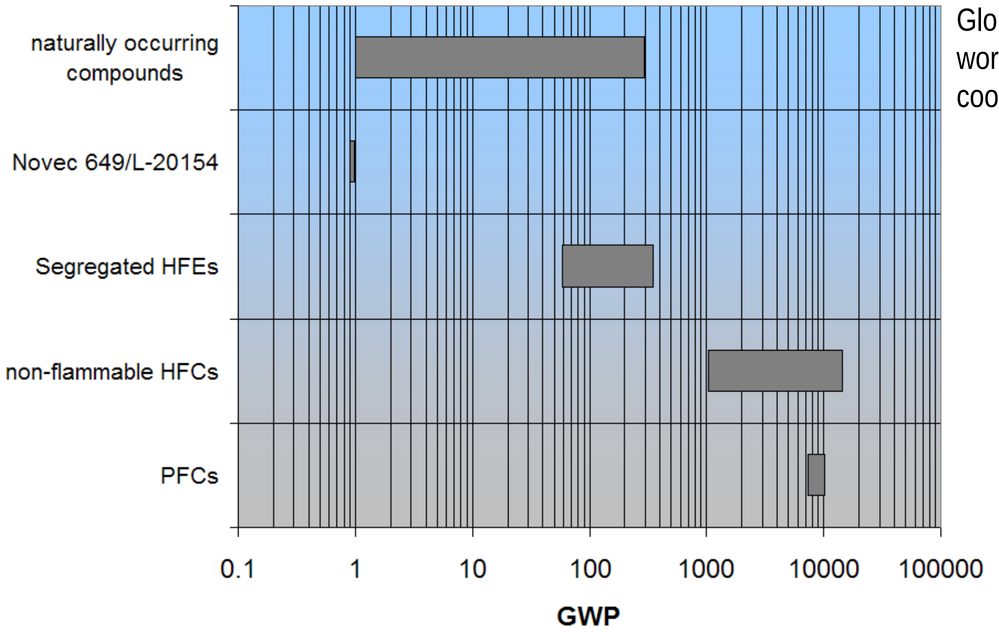

Global Warming Potential of various potential working fluids for passive 2-phase immersion cooling of datacom equipment.

II. SIGNAL INTEGRITY DATA

A. Motivation

1. The viability of OBI for High Performance Computing Applications depends upon the integrity of electrical signals moving through immersed components both electrical and optical.
    - PCBs, and electrical and optical connectors are designed to be placed in air
    - In addition to the difference in properties between air and a dielectric liquid, there are transient vapor/liquid phase discontinuities within a boiling liquid could cause similar discontinuities in electrical properties.
    - There are also differences in the properties of fluids. Hydrofluoroethers, for example, have dielectric constants much higher than those of perfluorocarbons, the traditional immersion fluids for computing applications. Fluoroketones are much more like perfluorocarbons.

B. Procedure and Results…

DIELECTRIC CONSTANT AND LOSS TANGENT MEASUREMENTS OF FK-$C_6F_{12}O$ AND HFE-$C_6F_9OH_5$ ENGINEERED FLUIDS

(Test Equipment = HP 8722ES S-Parameter Network Analyzer, HP 85070B Dielectric Probe Kit; Test Frequency Range = 50 MHz – 20 GHz)

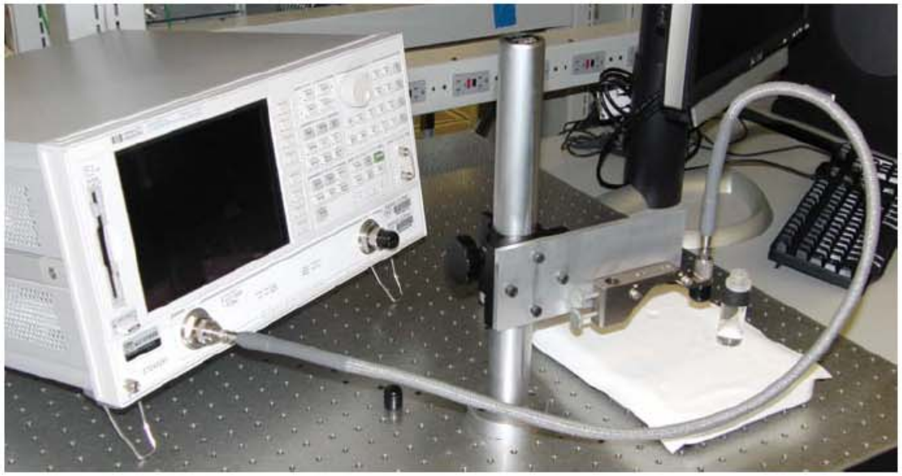

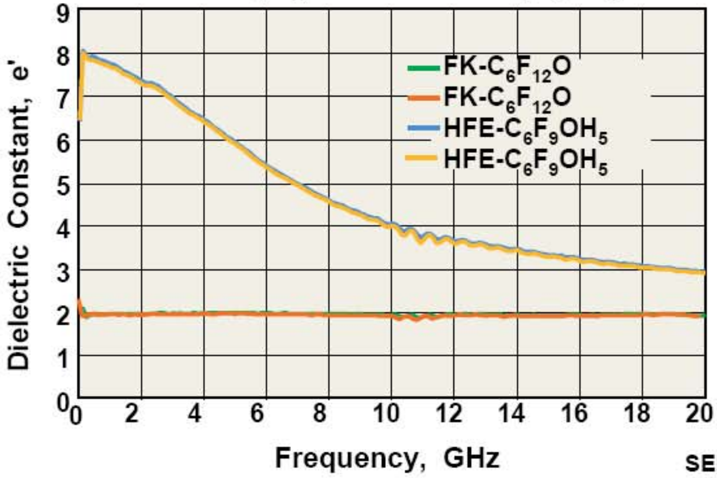

Dielectric Constant Measurement Results For FK-$C_6F_{12}O$ and HFE-$C_6F_9OH_5$

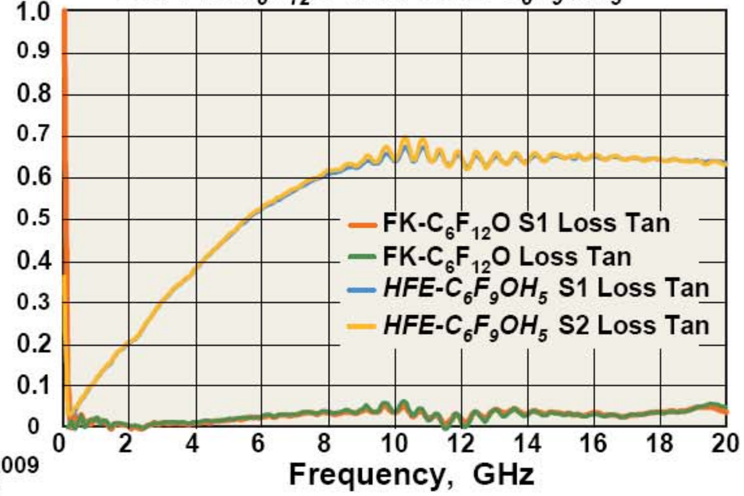

Loss Tangent Measurement Results For FK-$C_6F_{12}O$ and HFE-$C_6F_9OH_5$

SEP_02 / 2010 / SCP / 42009

MAYO CLINIC
SPPDG

TEST SETUP FOR DIELECTRIC FLUID S-PARAMETER MEASUREMENTS

Agilent Technologies E8361A 4 Port Precision Network Analyzer Test Setup

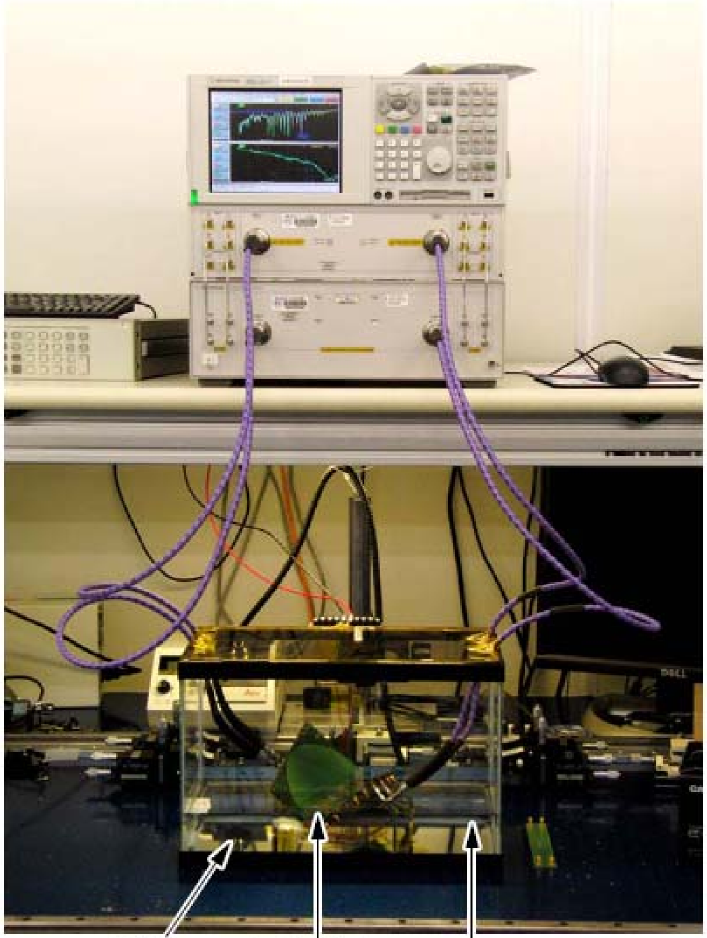

Experimental Test Vehicles

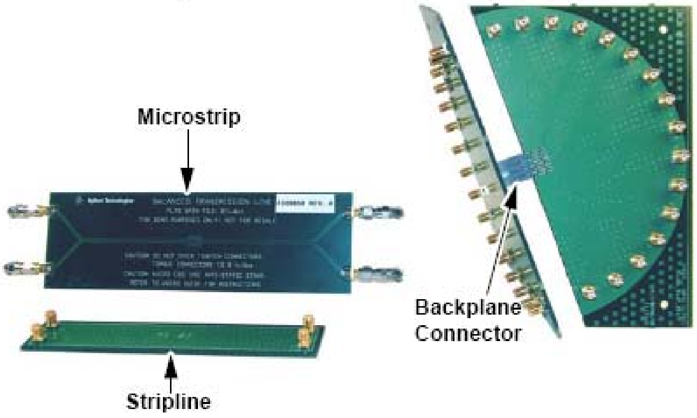

Backplane Connector Test Vehicle Immersed in Two-Phase Dielectric Fluid

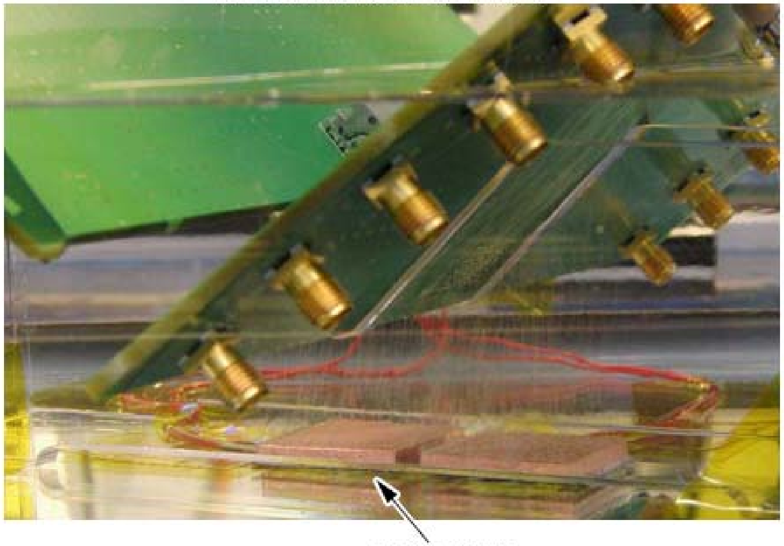

INSERTION LOSS AND RETURN LOSS MEASUREMENTS ON STRIPLINE TRANSMISSION LINE IMMERSED IN FK-$C_6F_{12}O$ AND HFE-$C_6F_9OH_5$ ENGINEERED FLUIDS

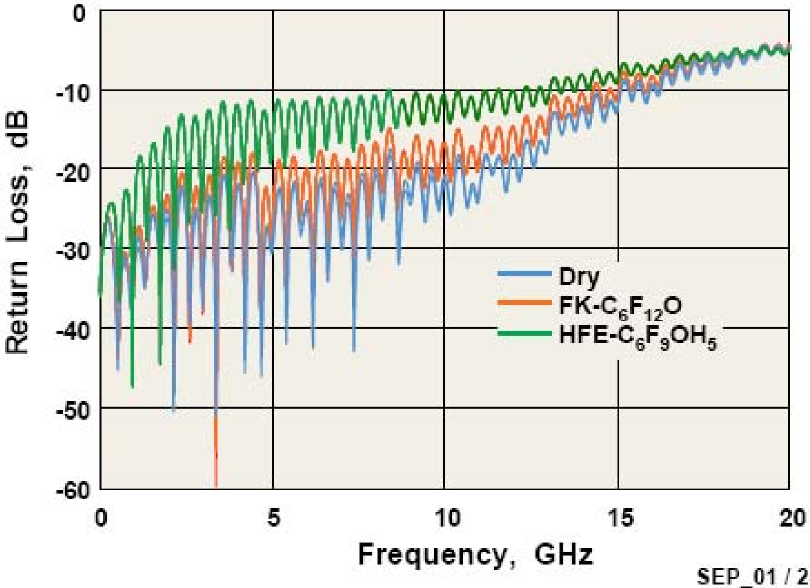

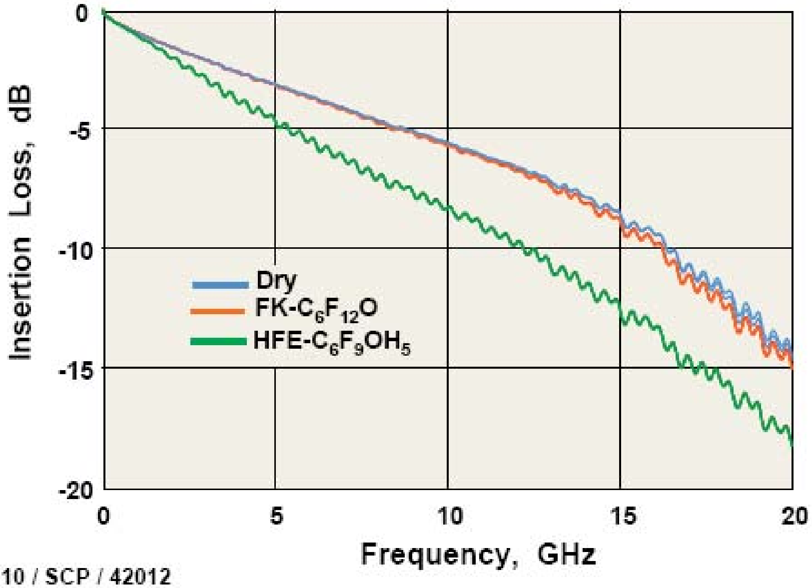

SEP_01 / 2010 / SCP / 42012

MAYO CLINIC
SPPDG

INSERTION LOSS AND RETURN LOSS MEASUREMENTS ON AGILENT TECHNOLOGIES BALANCED MICROSTRIP TRANSMISSION LINE TEST BOARD IMMERSED IN FK-$C_6F_{12}O$ AND HFE-$C_6F_9OH_5$ ENGINEERED FLUIDS

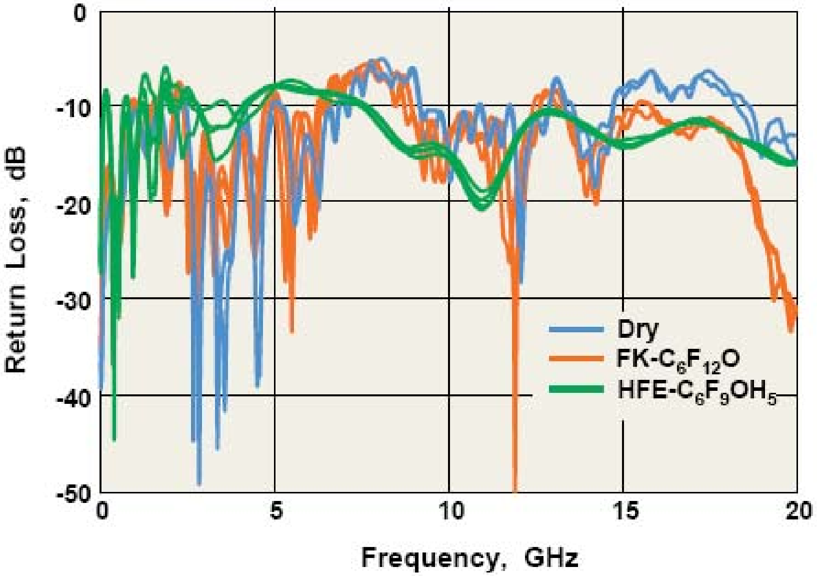

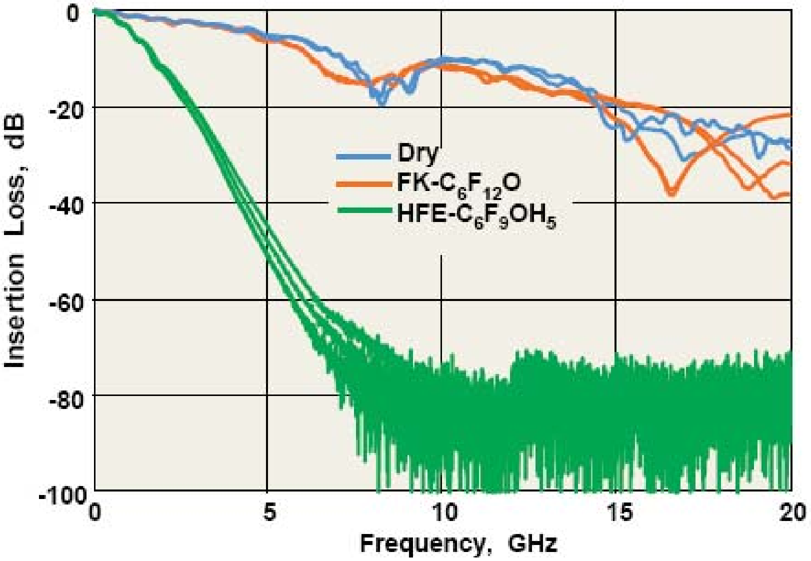

SEP_02 / 2010 / SCP / 42011

MAYO CLINIC
SPPDG

INSERTION LOSS AND RETURN LOSS MEASUREMENTS ON XCEDE BACKPLANE CONNECTOR IMMERSED IN FK-$C_6F_{12}O$ AND HFE-$C_6F_9OH_5$ ENGINEERED FLUIDS

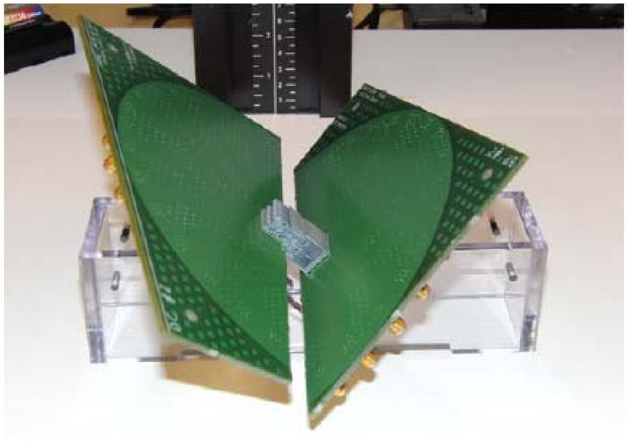

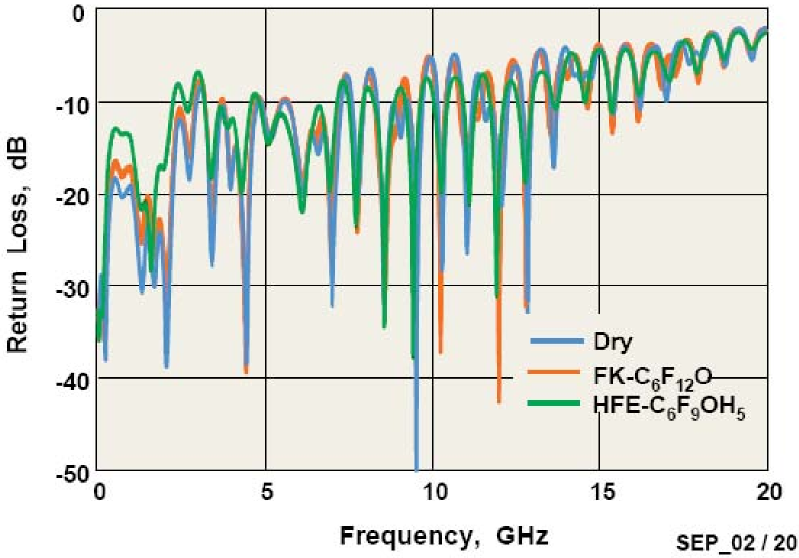

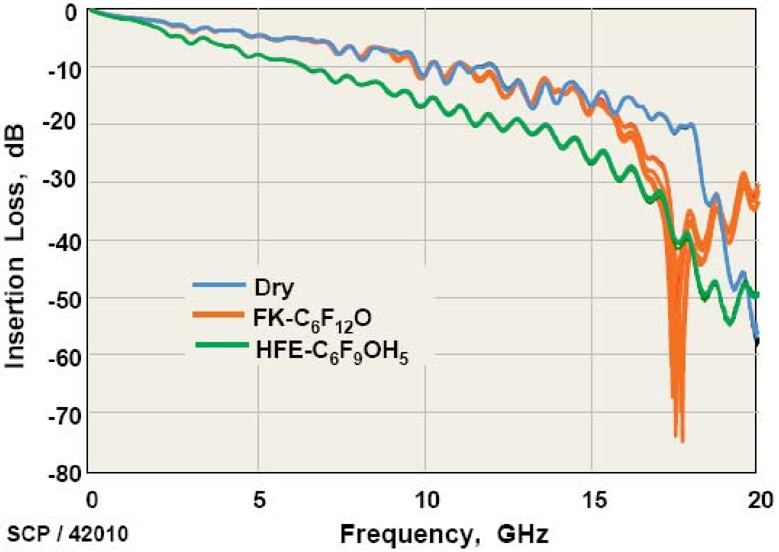

SEP_02 / 2010 / SCP / 42010

MAYO CLINIC
SPPDG

Optical Link Bit Error Rate Test Setup Using Engineered Fluids (Use Fiber Breakout Cable and FC Connector to Loop Back From Laser Transmitter to Photodetector In the Same Module)

Advantest Bit Error Rate Tester

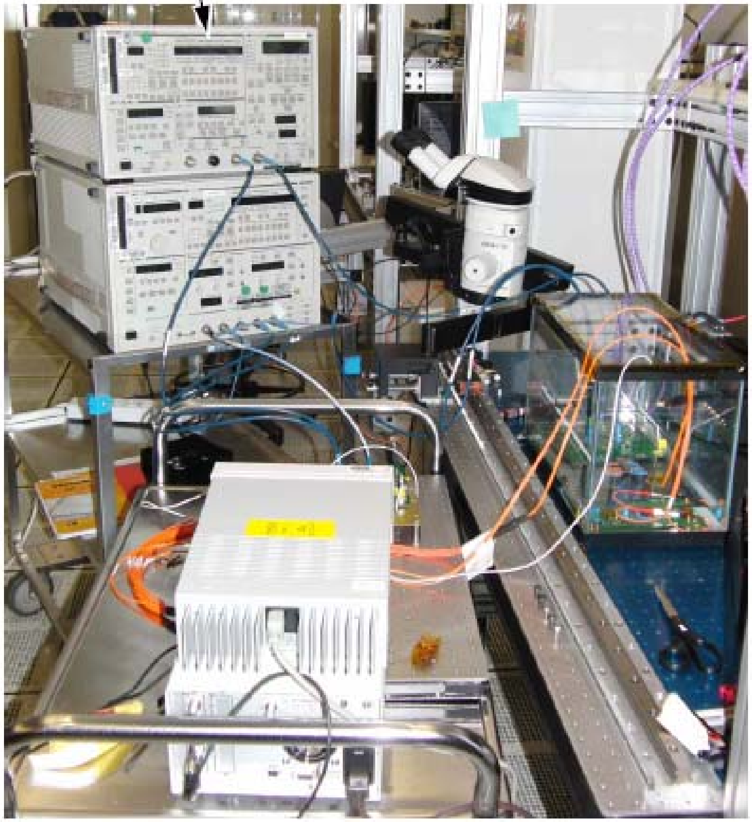

Fujitsu 4-Channel o-microGiGaCN™ 850 nm Transceiver Provides 6.25 Gbps x 4 Channel Bi-directional Data Transfer

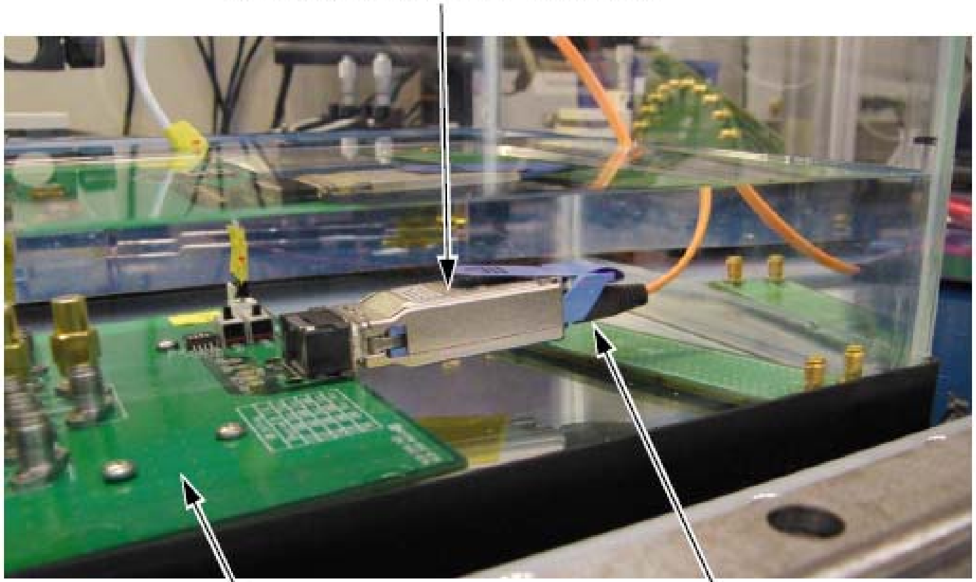

Fujitsu microGiGaCN™ Evaluation Board

12-Fiber Ribbon Cable (Only 8 Fibers Used) With MPO Connector

SEP_03 / 2010 / NEH / 42029

MAYO CLINIC
SPPDG

OPTICAL LINK TESTING
USING ENGINEERED FLUIDS

- Use Fiber Breakout Cable And FC Connector To:
    - Loop back from laser transmitter to photodetector
    - Attach fiber to optical power meter to measure laser CW output power
- FC-FC connectors through bulkhead connector
    - No change in optical power when submerged in FK-$C_6F_{12}O$ or HFE-$C_6F_9OH_5$
    - No gap between fiber cores so there is no fluid in the optical path
- Fujitsu transceiver module
    - Power
        - FK-$C_6F_{12}O$: No change in optical power when submerged
        - HFE-$C_6F_9OH_5$: Immediate drop in optical power (~5 dB) when submerged
    - Bit Error Rate Test
        - 2^7-1 pseudo-random bit sequence
        - Some difficulties with measurement due to optical power loss when fibers are bent while moving module to submerge it
        - FK-$C_6F_{12}O$: 6 Gbps operation with no errors while submerged
        - HFE-$C_6F_9OH_5$: Significant number of errors when submerged
            - Possibly due to reduced optical power

III. 1.8kW OBI COOLED DESKTOP

A. Motivation

1. Much of the junction-to-fluid data presented so far is calculated from experimental sink-to-fluid data gathered with electric heaters.
2. It is desirable to gather data with real hardware that is challenging to cool by conventional means.

B. Procedure and Results

1. An immersion test platform was built
   - EVGA 170-BL-E762-A1 LGA 1366 Intel X58 4-way SLI Class. XL ATX Motherboard
   - INTEL I7 930 CPU (die area 263mm^2)
   - (4) NVIDIA GeForce GTX 480 GF100 DX11 Video Cards (GPU die area 530mm^2)
   - Two 1,100 W power supplies, One 700W power supply
2. GPU and CPU were modified for immersion
   - The plating on the integrated heat spreaders (IHS) of the CPU and GPU were removed by lapping
   - BEC enhanced copper boilers were soldered (52In 48Sn) atop the IHS using a vapor phase process
     - CPU boiler was flat 30x30x1mm
     - GPU boiler was 40x40x1mm but covered with 1x1x4mm pin fins at 2mm pitch
3. System was immersed in the HFE C4F7OH3, T_b=34°C
4. A 12x24x6cm water-cooled condenser above the liquid condenses the vapor
   - Sized for 4kW at 10°C approach

Immersion Cooled 4-Way SLI Setup

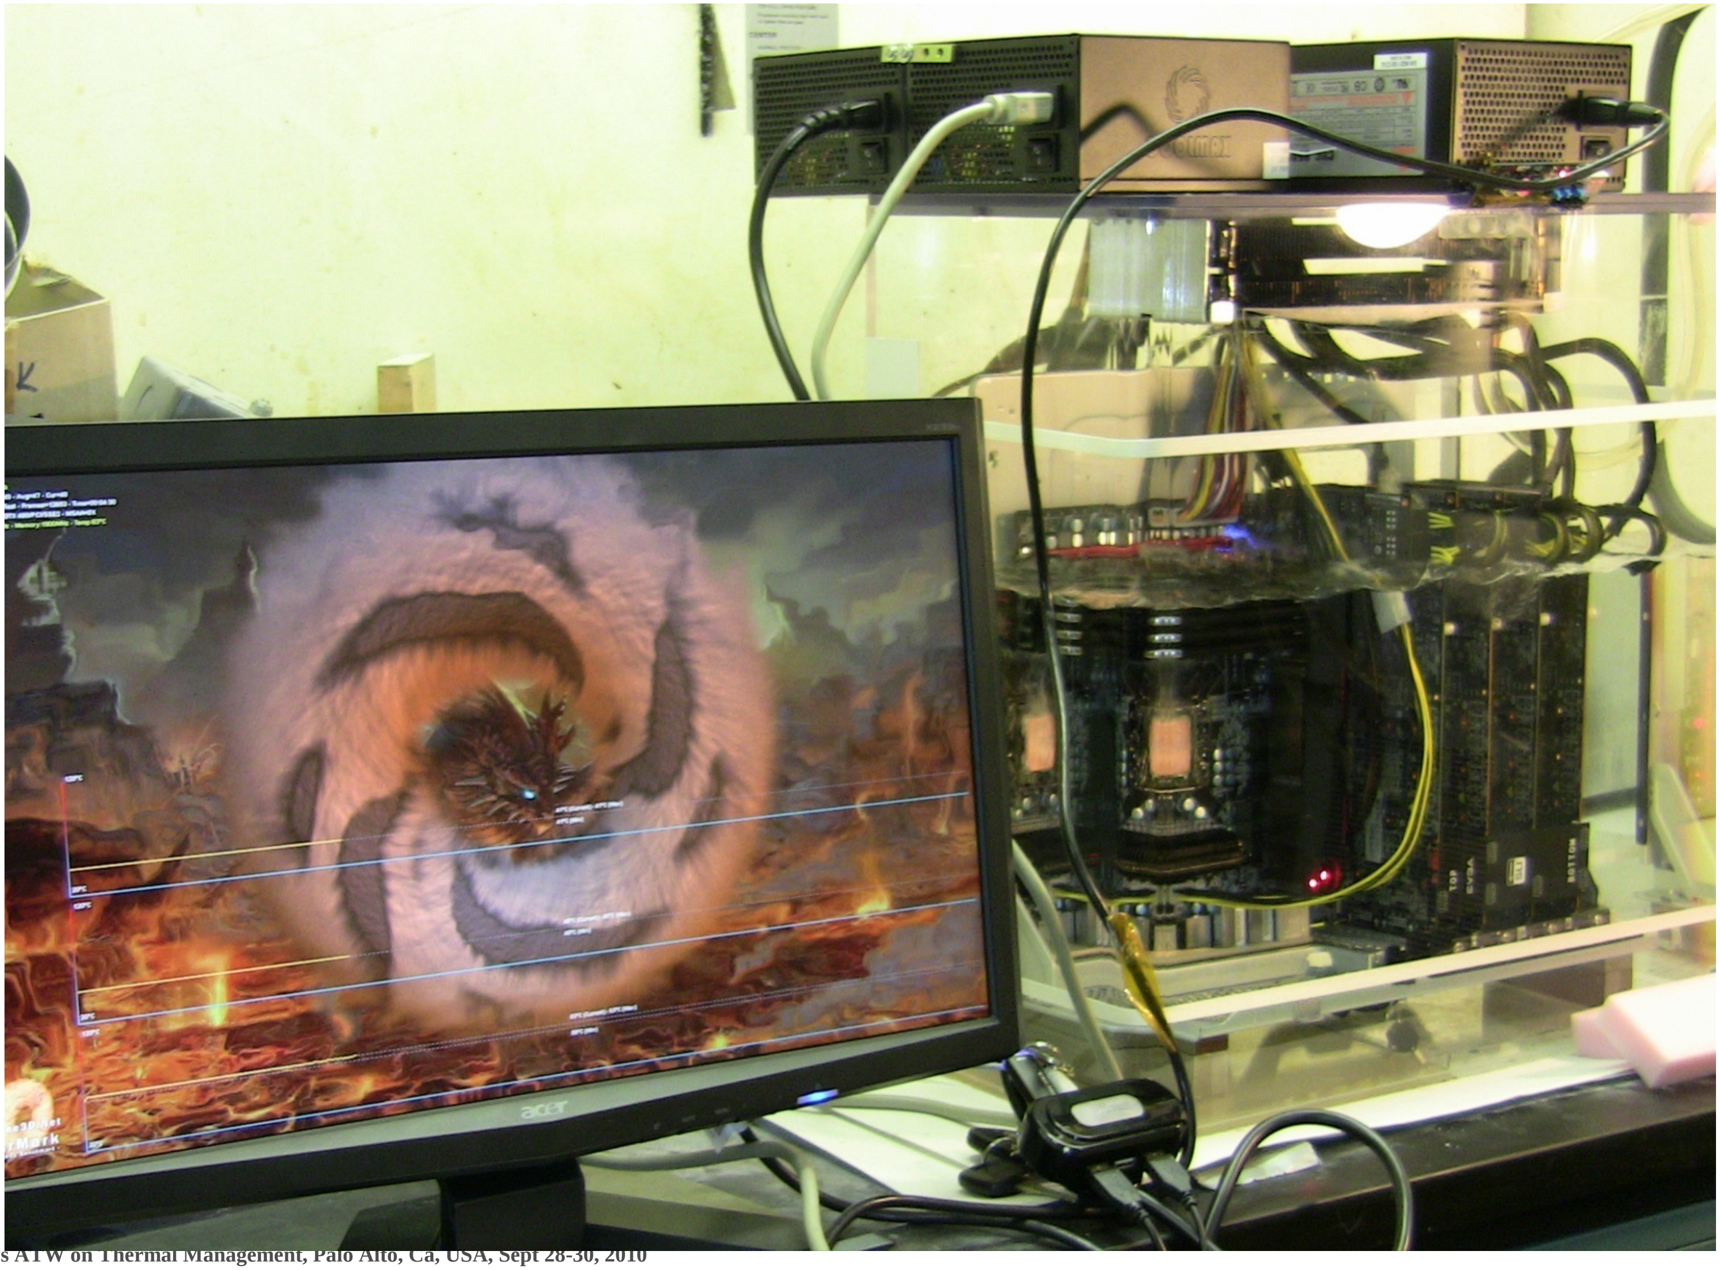

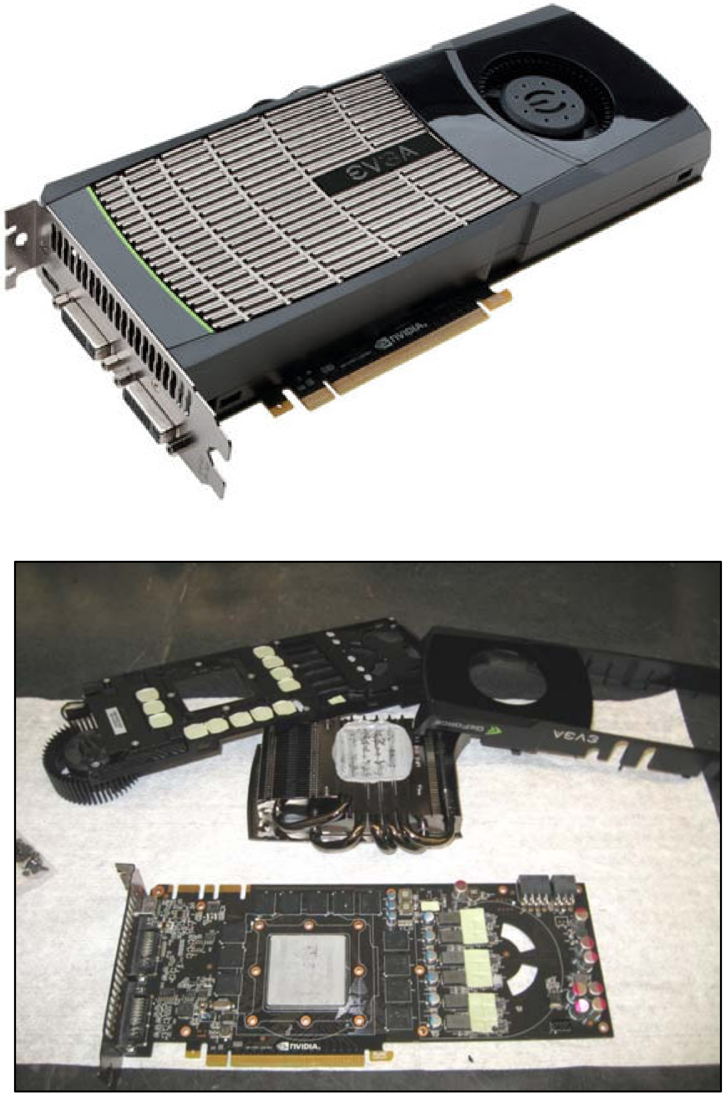

NVIDIA GeForce GTX 480 GF100 DX11 Video Card and Teardown

Graphics Processor before modification. Die is ~23x23mm and integrated heat spreader is ~42x42mm.

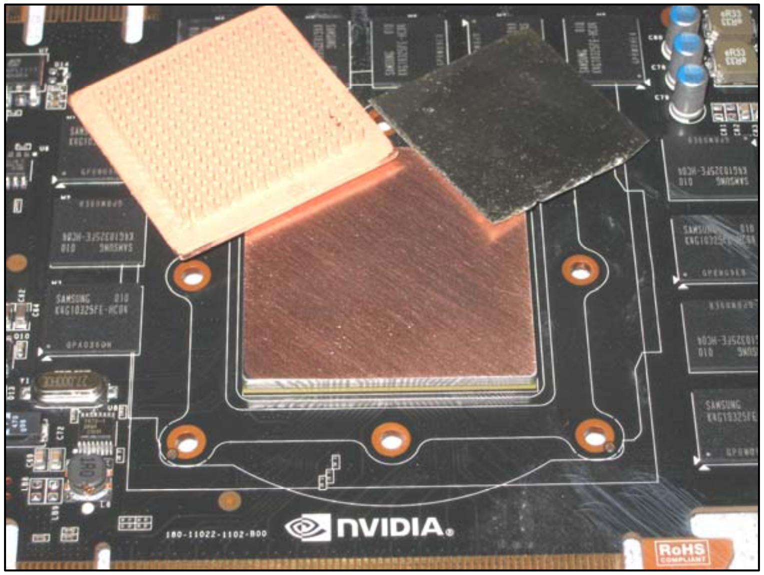

GPU with nickel plating removed from IHS. Also shown are 40x40x1mm BEC-coated pin fin boiler (1x1x4mm fins at 2mm pitch) and InBi solder pre-form.

CPU with Nickel plating removed (die is 263mm^2). BEC-Boiler, 30x30x1mm

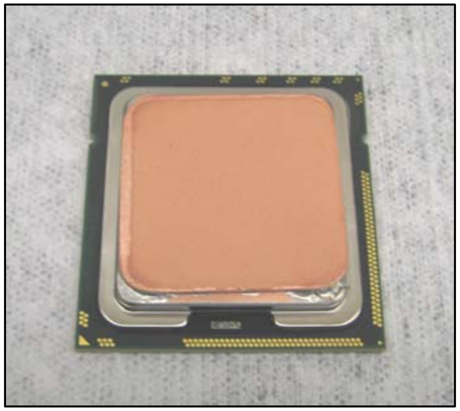

Modified CPU

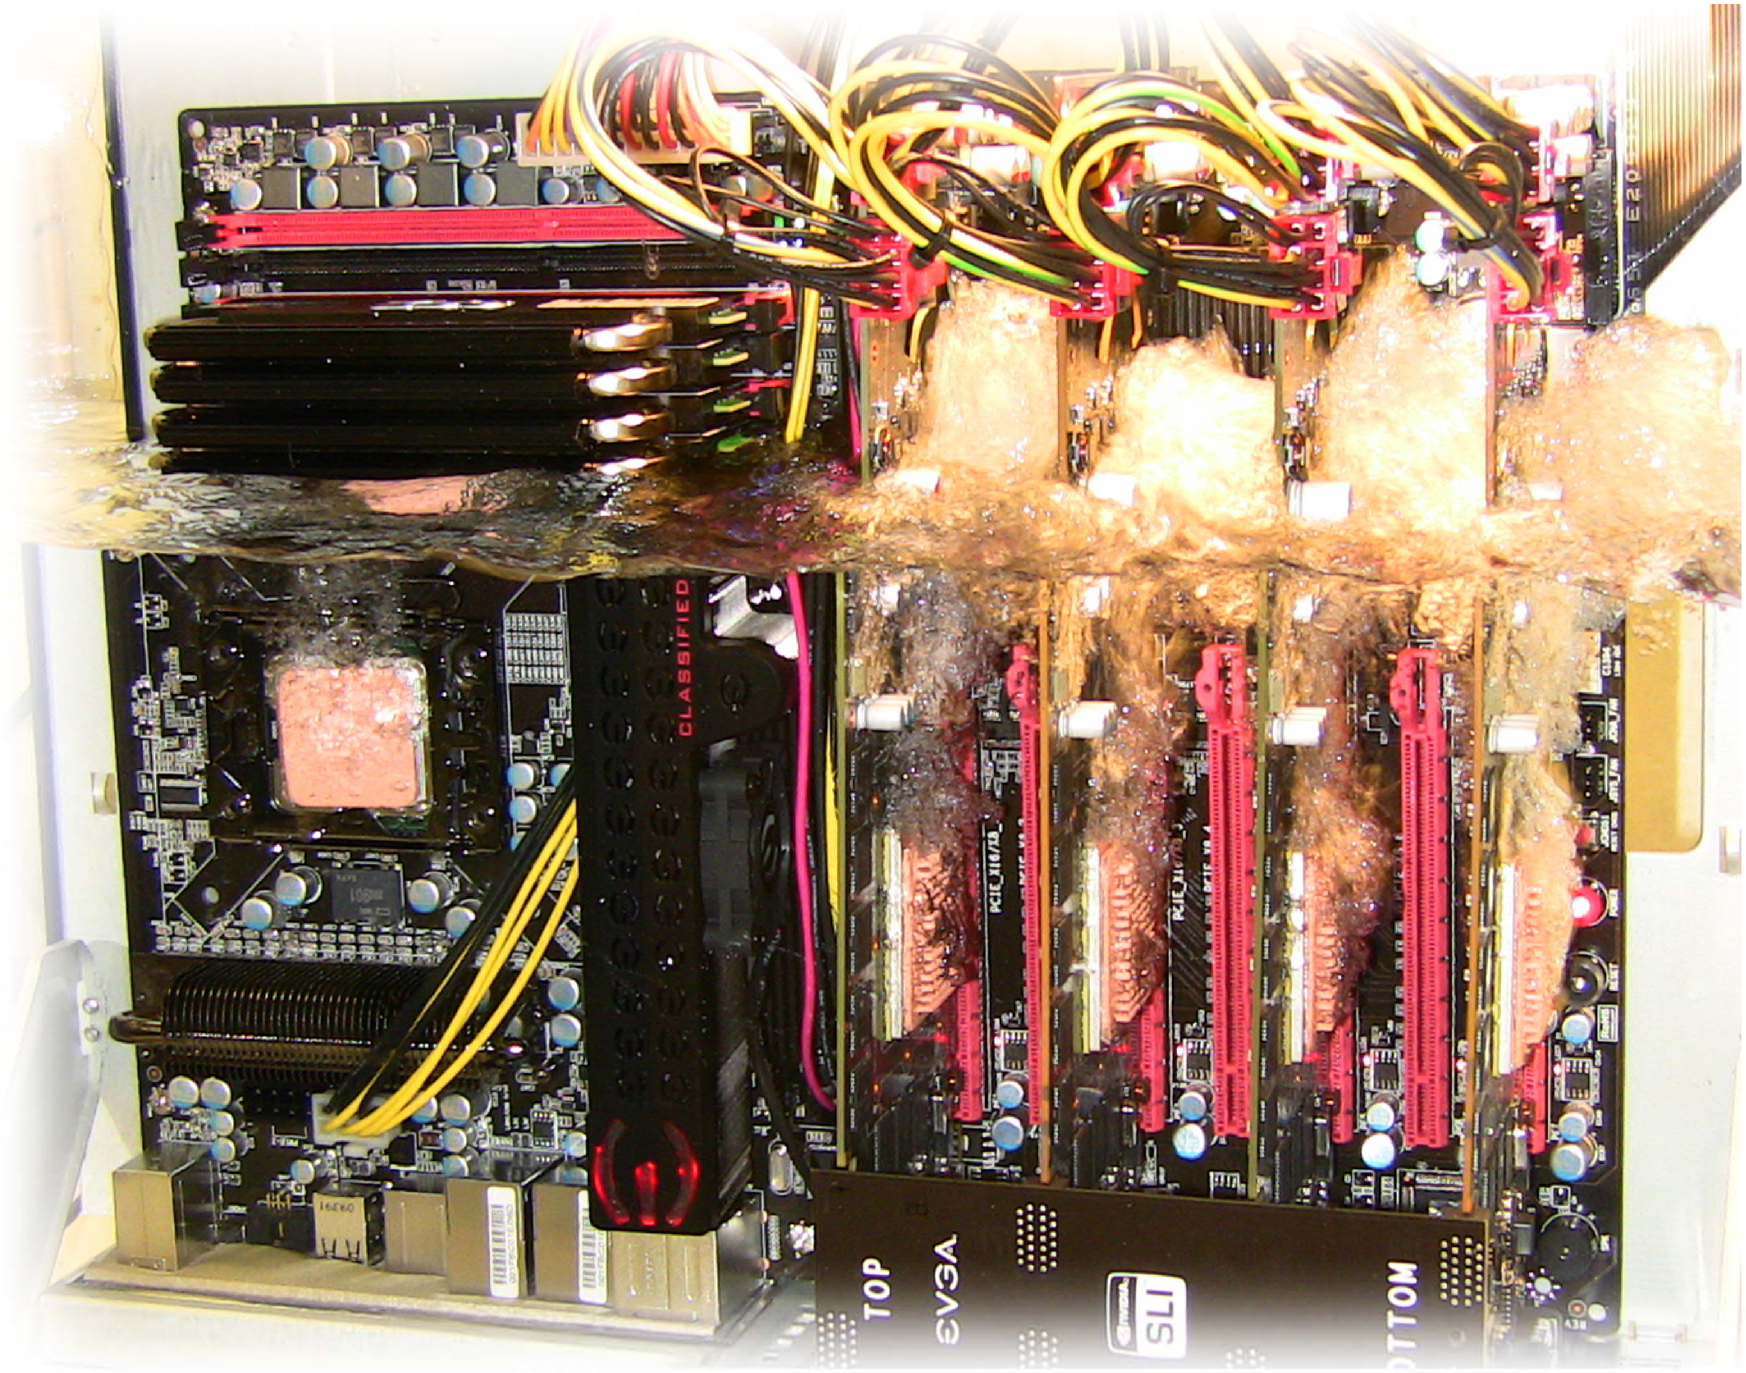
CLASSIFIED
TOP
EVGA
SLI
BOTTOM

B. Procedure and Results (cont.)

1. Baseline Air Cooled Data, T_a=23°C

- GPU average T_j=99°C, R_{ja}<0.29°C/W, at 260W
- CPU Average T_j=50°C, R_{ja}<0.27°C/W, at 100W

2. Immersion CPU Results

- T_j<40°C based upon the on-chip sensor
- With T_f=34°C, results in R_{jf}~0.06°C/W

3. Preliminary Immersion GPU Results

- Average T_j~72°C based upon the on-chip sensor
- With T_f=34°C, results in R_{jf}<0.15°C/W
- Factory TIM1 found to be lacking, possibly damaged by lapping and soldering when attaching boiling surface.

4. Final Immersion GPU Results

- IHS/boiler assembly reattached with ShinEtsu grease
- With T_f=34°C, R_{jf}~0.076°C/W

GPU Clock [MHz]	GDR Clock [MHz]	Frame Rate [fps]	Graphics Card Power [W]	GPU Power [W]	T_j [C]	R_{jf} [C/W]
725	1900	47	252	192	48	0.073
800	1900	49	272	207	51	0.082
850	2000	55	291	221	50	0.072
900	2000	58	308	234	51	0.073
924	2001	59	311	236	53	0.080

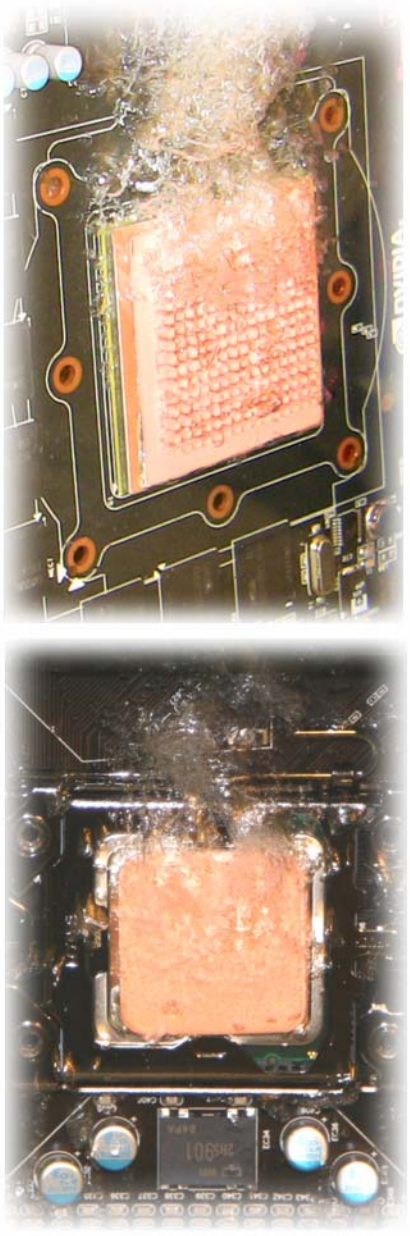

VI. OBI-COOLED APPLE G5

- Apple G5 (Dual 2.5GHz)
  - This formerly water-cooled machine has two 7x9 mm bare die CPUs that dissipate about 100W each with a core heat flux of ~200W/cm^2
    - When water cooled, typical CPU temperatures were about 70°C with water at 33°C.
    - R_{jw}~0.37°C/W
  - The immersion cooled G5 has BEC-enhanced boilers applied with Hi-Flux metal alloy TIM
    - Boilers 2mm thick, 25mm diameter
    - T_j~67°C when immersed in the HFE C4F7OH3 (T_{sat}=34°C)
    - R_{jf}~0.33°C/W

VI. FUTURE WORK

- Long Term Performance Testing
  - 3M Electronic Markets Materials Division
    - Four compute nodes with eight 8-core AMD Processors
    - Operational November 2010
    - Project trackable online via Facebook
    - Track thermal performance, efficiency, fluids loss, gas purge energy consumption

- Signal Integrity and Power Density
  - –Mayo SPPDG
    - 12”x12” custom PCB populated with 16 – 280W Heaters
    - Continue SI experiments
    - Investigate long term reliability effort

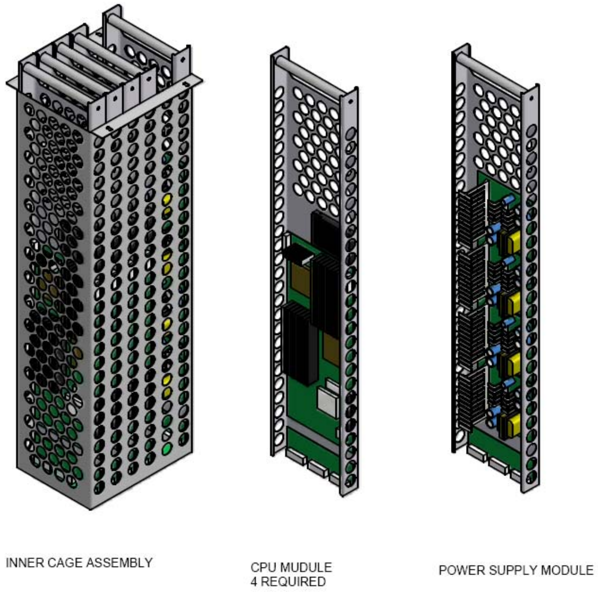

V. CONCLUSIONS

A. Preliminary investigation of electrical signal integrity though immersed components:
  - Stripline, microstrip, backplane connector components immersed in a fluoroketone (FK) fluid appear able to maintain electrical signal integrity at frequencies exceeding 15 GHz.
  - Hydrolfuoroether (HFE) fluids may be unable to maintain electrical signal integrity above a few GHz

B. Preliminary investigation of bit error rate through immersed optical connector at 850nm and 6.25Gbps
  A. Optical connectors appear able to transmit without error in FK fluids. Use of an HFE fluid resulted in 5dB optical power loss and significant bit error rates.

C. Heat Transfer Test Results based upon on-chip temperature sensors
  - Passive 2-phase immersion technology was able to remove 100W from a 23x23mm CPU with a junction-to-fluid resistance of <0.06°C/W.
  - Passive 2-phase immersion technology was able to dissipate 240W from an overclocked 23x23mm GPU with a junction-to-fluid resistance of 0.076°C/W
    - TIM1 was a silver grease.

Reference

1. Tuma, P.E., “The Merits of Open Bath Immersion Cooling of Datacom Equipment,” Proc. 26th IEEE Semi-Therm Symposium, Santa Clara, CA, Feb. 21-25, 2010.