Exploring Innovative Cooling Solutions for IBM's Super Computing Systems: A Collaborative Trail Blazing Experience

Dr. Richard C. Chu, IBM Fellow
Member, National Academy of Engineering, USA
Academician, Academia Sinica, ROC

This talk presents a practitioner's personal recount on how collaborative innovation can be achieved to create innovative cooling solutions for high performance computing systems. It draws on my 50 years' experience, working for IBM since 1960. The emphasis is on learning from others by understanding the underpinnings of their novel ideas. It is then necessary to apply what you have learned using your imagination. During the talk, 6 milestone achievements in cooling system innovation will be presented to illustrate the process of collaborative innovation to meet the ever increasing cooling requirements for IBM super computing systems. Also included will be a summary of IBM's experience of sponsoring basic research at 12 universities exploring innovative cooling technology for future applications. Finally, I will conclude my talk by sharing my personal process for innovation by outlining 6 rules of creativity.

(A 45 minute presentation with 58 slides)

This picture was taken in 1960, the year that I joined IBM in Poughkeepsie. I just want to prove that I was young once.

"Imagination is more important than knowledge."

Albert Einstein

"Perseverance – There is no substitute for hard work."

Thomas A. Edison

"The harder you work the luckier you get."

Gary Player

"Healthy body, healthy mind – Only a healthy body can support a healthy mind."

Source Unknown

"The real battle is conquering ourself."

Clarence Thomas
U.S. Supreme Court Justice

My One Minute Story

Since joining IBM, I have been surrounded by good people who inspired me, challenged me, and most importantly guided and supported me. Occasionally, people including my own children have asked me, what is my secret to enjoy smooth sailing for so long. I honestly do not know the answer, but, I firmly believe that the key to success in professional life is an ongoing investment of hard work and perseverance. Of course luck is also a part of it. However, it has been my experience that the harder you work, the luckier you get. I also believe that you must have a balanced life. This means that you should find time to do other things as well, such as practicing the religion of your choice, regular exercise, and engaging in sports or other activities you enjoy. Personally I like winter skiing, summer sailing, and golfing whenever I find time. Finally, I have been happily married since 1963 and I will let the next few slides tell the rest of the story.

My 2004 family reunion in Southern California with my 4 children, their spouses and 6 grandchildren.

A 2005 virtual reunion of my 8 grandchildren through electronic photo editing.

My 2007 family reunion in San Diego, California with my 4 children, their spouses and 8 grandchildren.

A 2007 actual reunion of my 8 grandchildren in San Diego, California.


Holiday Greetings
From Our Family
To Yours


08/16/2009

Significant IBM Cooling Development Innovations of the Last Five Decades in Support of Continued Performance Improvements for Bipolar and CMOS Technology for High Performance Computing

- **Air/Liquid Hybrid Cooling Technology for System 360/370 Mainframe Systems (1960's – 1970's)**
- **Thermal Conduction Water-Cooling Technology for 3080, 3090 & ES9000 Systems (1980's – 1990+)**
- **Flat Plate Conduction Cooling Technology for IBM CMOS-based Server (mid-1990's – present)**
- **Modular Refrigeration Cooling Technology for High End IBM Servers (late 1990's – present)**
- **"Cool Blue" Air-Liquid Hybrid Cooling Technology for the Data Center (2005 – future)**
- **P575 Supercomputer Water Cooling System With Clustered Cold Plate Assembly and Modular Water Cooling Unit – WCU (2008 – present)**

Significant IBM Cooling Development Milestones of the Last Five Decades in Support of Continued Performance Improvements for Bipolar and CMOS Technology for High Performance Computing

- Over the past 50 years increases in computing performance were accompanied by increases in power consumed and heat produced within the computing system.
- The late 1980s was a period of dramatic increases in performance and in power. The switch from Bipolar to CMOS brought a brief respite in the early 1990's, but the upward power trend continued into the 2000's to provide increased computing performance.
- Throughout this period and continuing today IBM has been a technology leader in cooling development achieving many milestones to provide the cooling performance needed to gain dramatic improvements in computing capability.

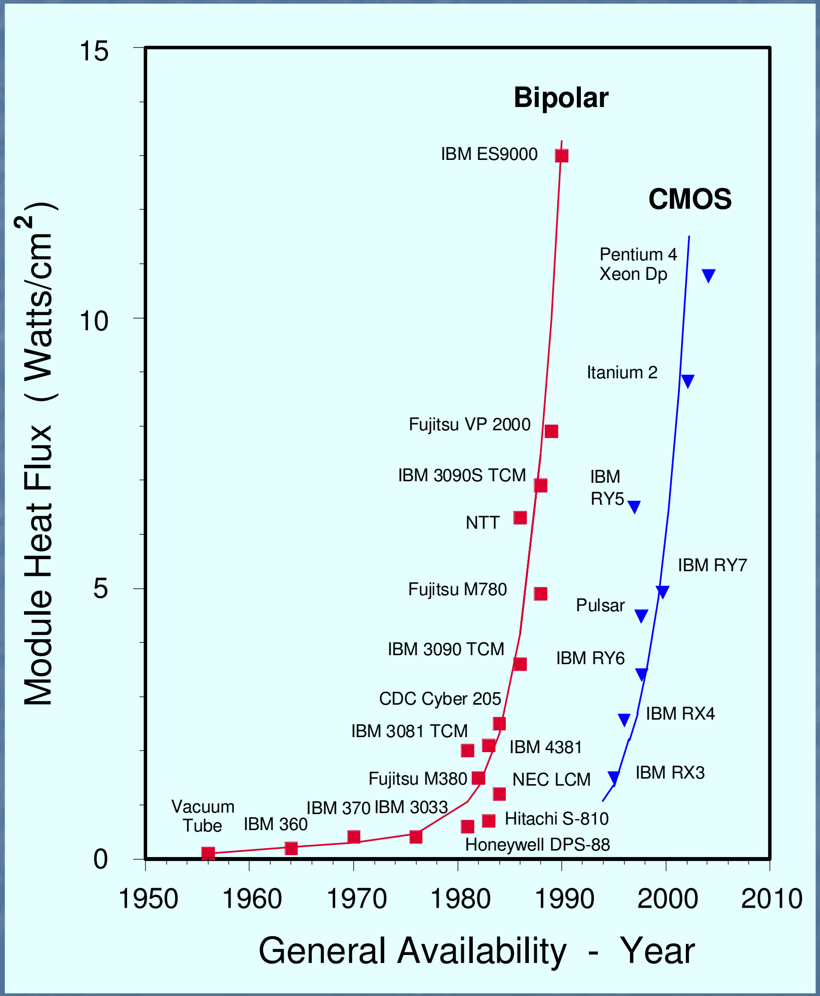

Air/Liquid Hybrid Cooling Technology for System 360/370 Mainframe Systems

- In the late 1960's through the 1970's IBM developed the hybrid air/liquid cooling concept allowing air-cooling in large mainframe racks to accommodate increased heat loads. In these systems as much as 50% of the heat dissipated was removed from the cooling air and transferred to customer facility water.

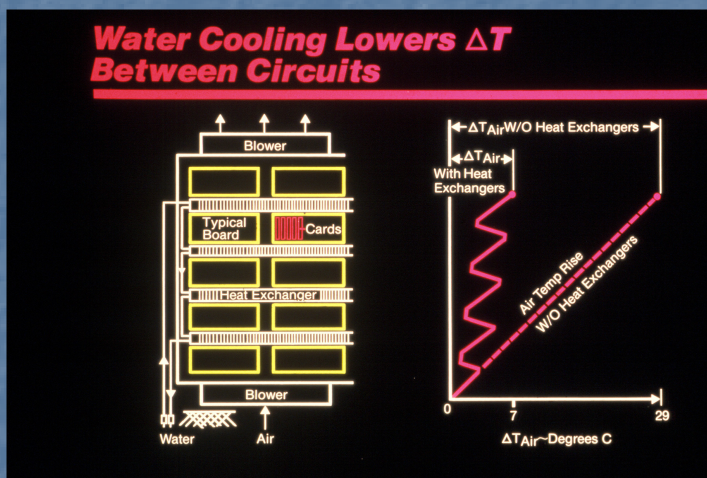

Interboard Heat Exchangers in Frames

System/360 - Card on Board Packaging

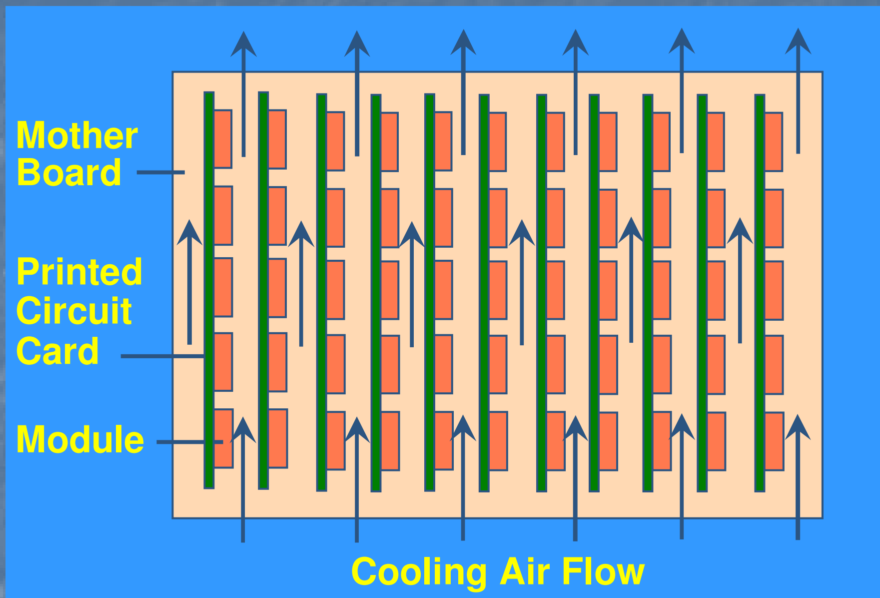

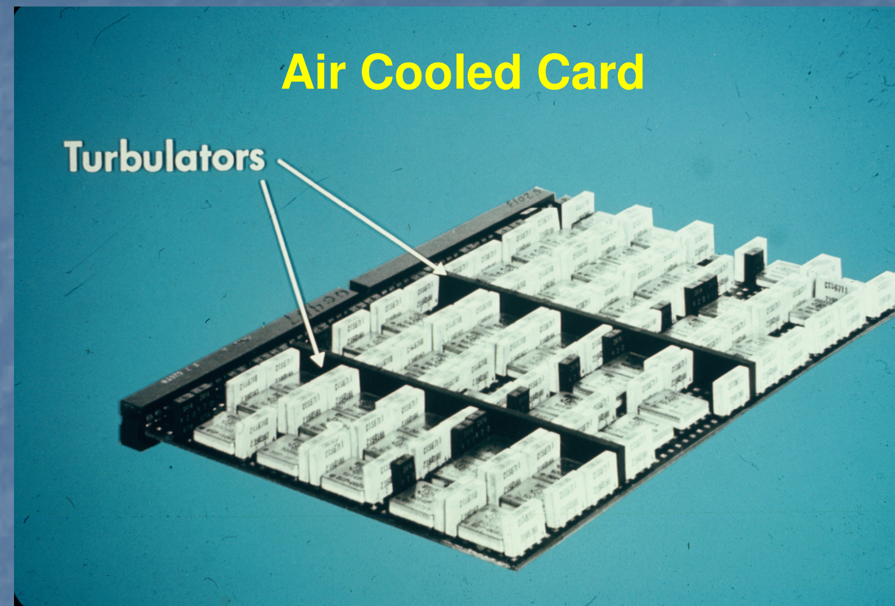

Module Circuit Packages

Hybrid Air-to-Water Cooling (ca. 1967)

- System/360 Model 91

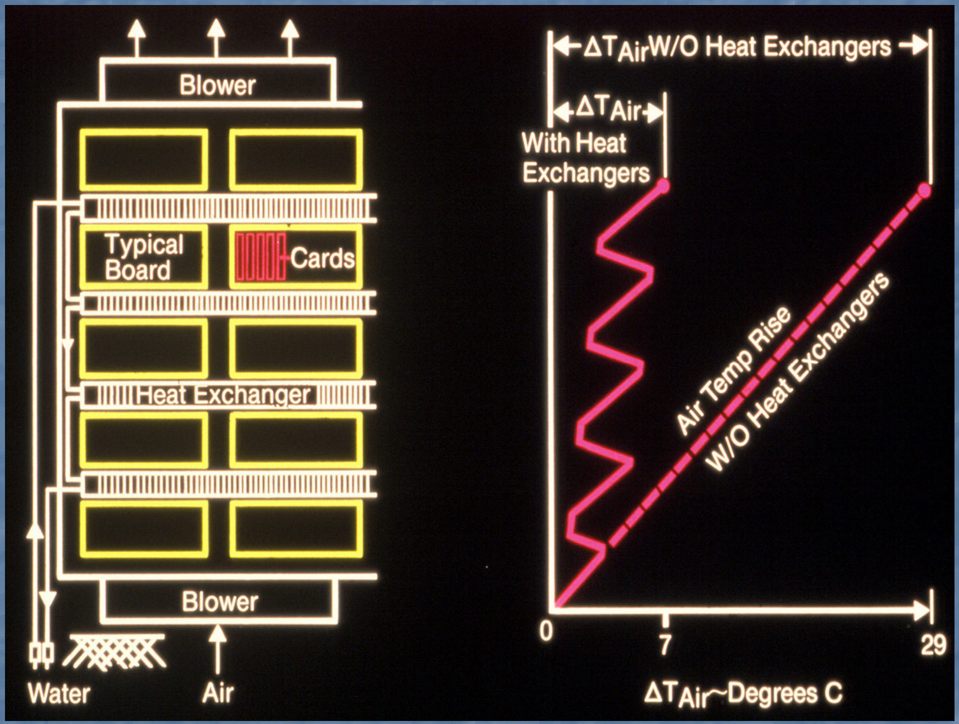

Coolant Distribution Unit (CDU)

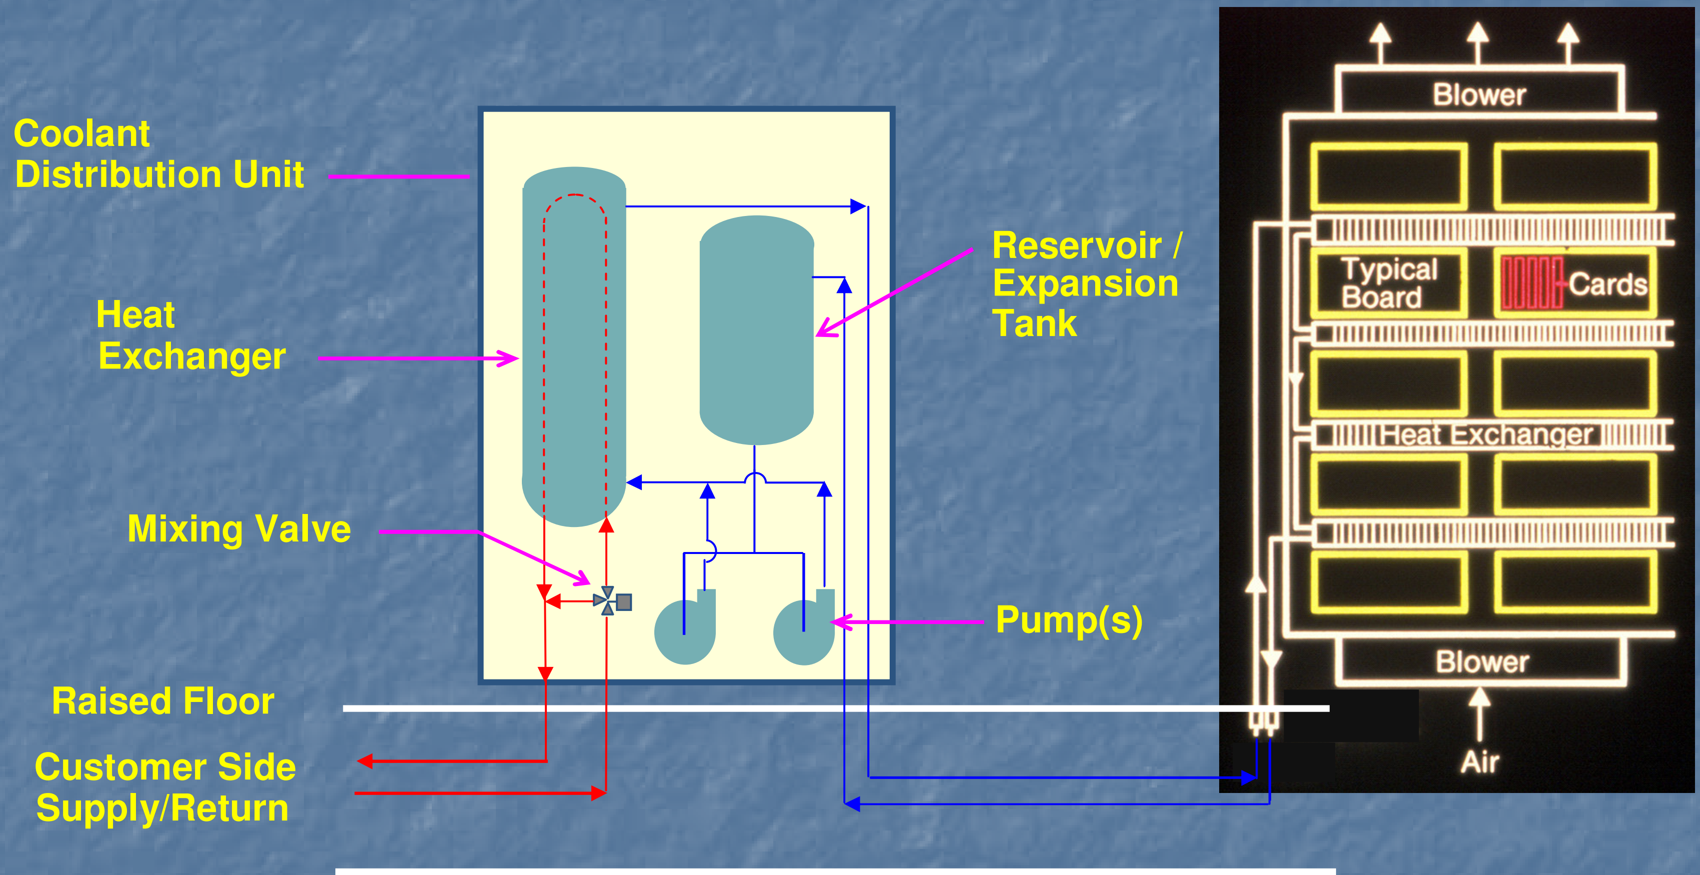

System/360 Model 91

- Heat Load to H_2O : 40 – 79 kW
- H_2O Flow Rate : 35 – 100 GPM
- P_{module} : 300 mW
- Air Flow Rate: 100 CFM/Board Column

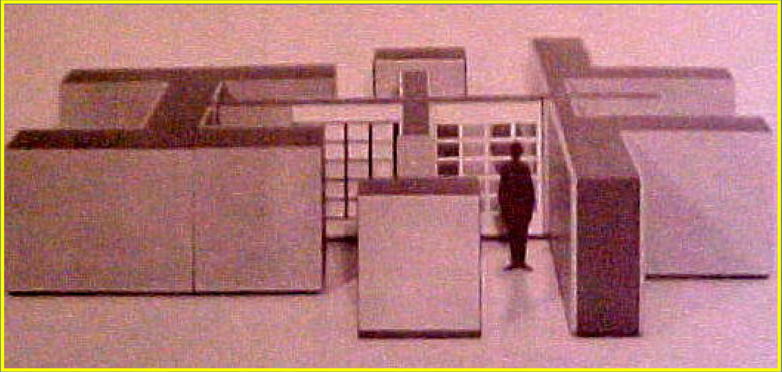

Typical System Configuration

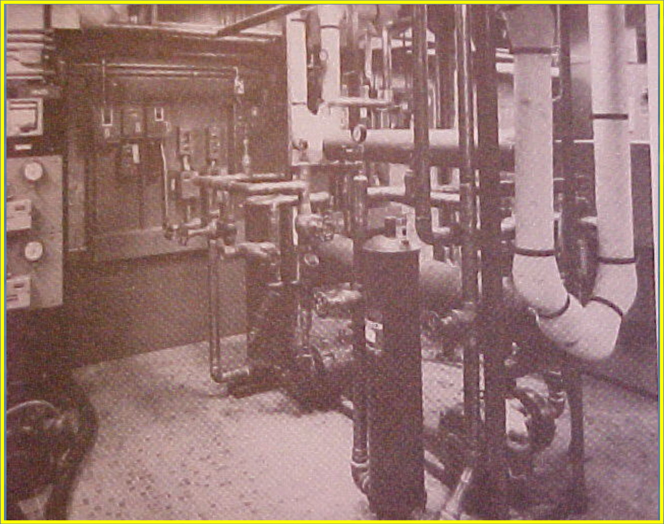

Pumping – Heat Exchange Installation

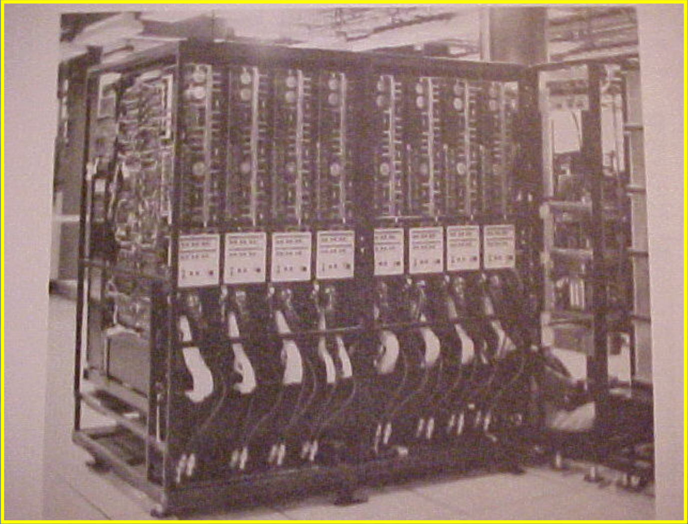

Storage Power Frame

CPU Power Frame

Thermal Conduction Cooling Technology for 3080, 3090 and ES900 Mainframe Systems

- In the early 1980's the revolutionary water-cooled Thermal Conduction Module cooling concept was introduced to manage the high heat loads produced by high performance Bipolar circuits in large multi-chip modules. The heat dissipated by the processor modules and many of the power supplies in these systems was transferred to customer facility water.

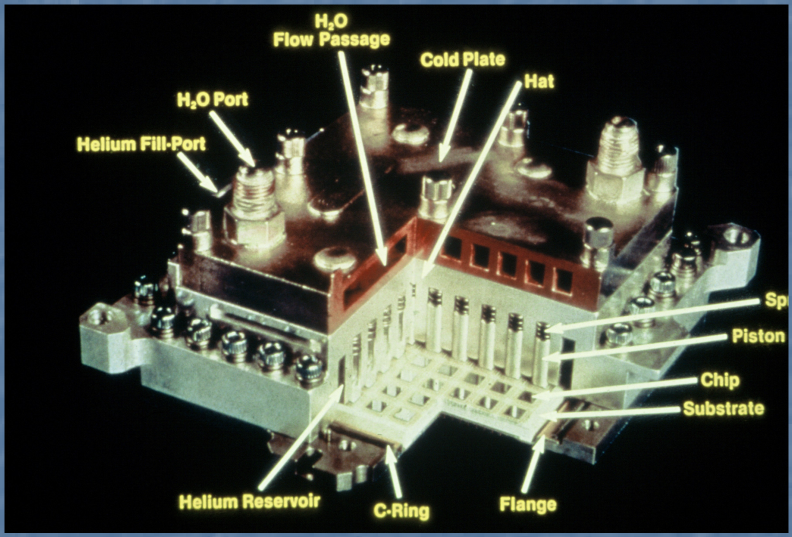

IBM Thermal Conduction Module

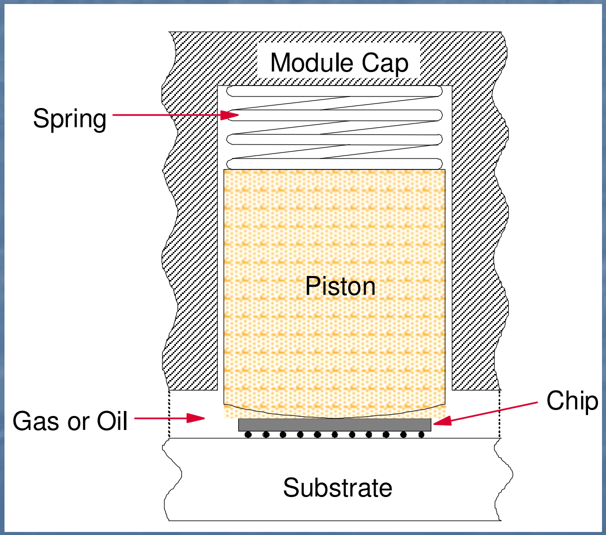

Multi-Chip Module Cooling Problem

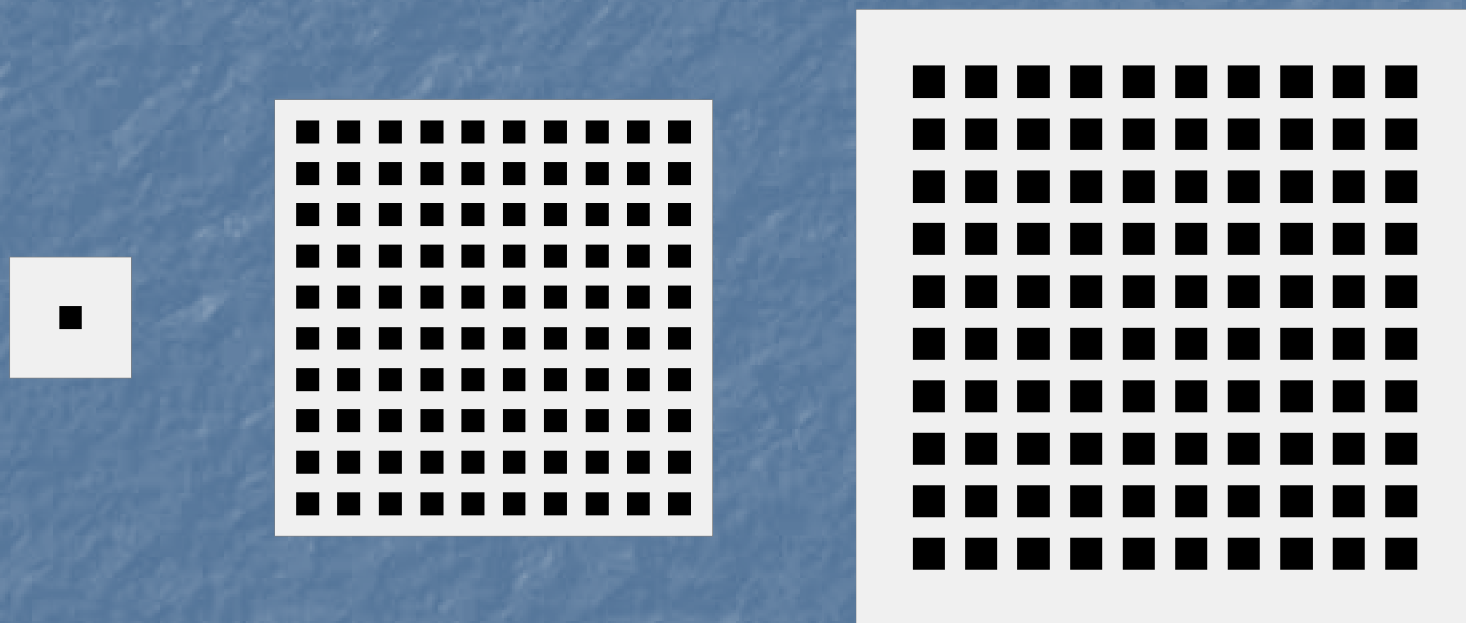

S/360
0.3 W/cm2

IBM 3081
3.7 W/cm2

ES 9000
11.8 W/cm2

Indirect Water Cooling (ca. 1980)

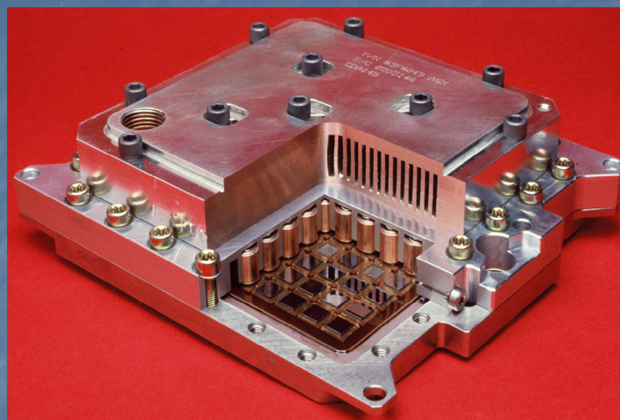

ES9000 Thermal Conduction Module (TCM)

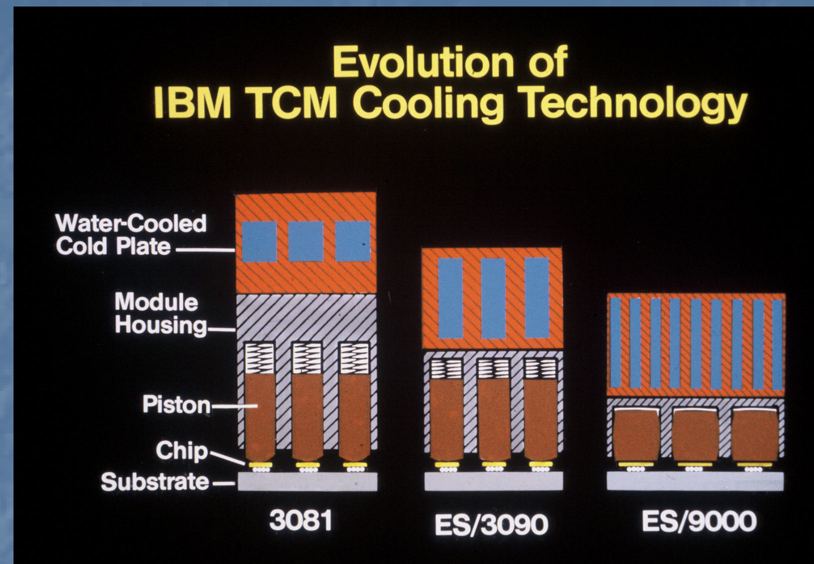

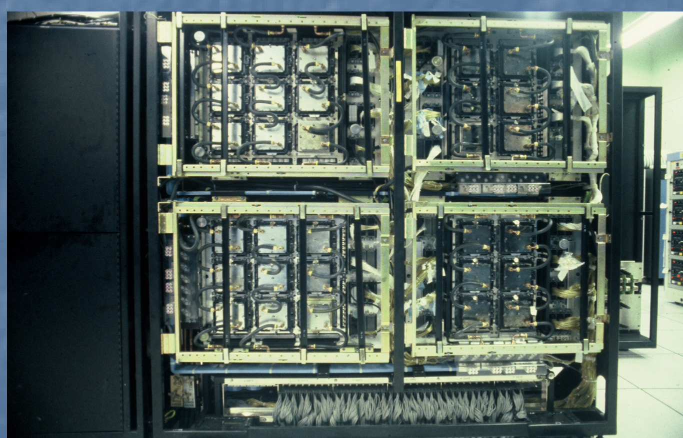

IBM System	3081	ES/3090	ES/9000
Year	1980	1985	1990
Thermal Resistance (°C/W)	11.6	7.7	1.8
Thermal Resistivity (°C/W/cm^2)	9.4	8.7	2.8
Chip Heat Flux (W/cm^2)	19.0	33.0	64.0
Module Heat Flux (W/cm^2)	3.7	5.3	11.8

S/3090 Frame w/TCMs

Thermal Conduction Module (TCM)

TCM Cutaway

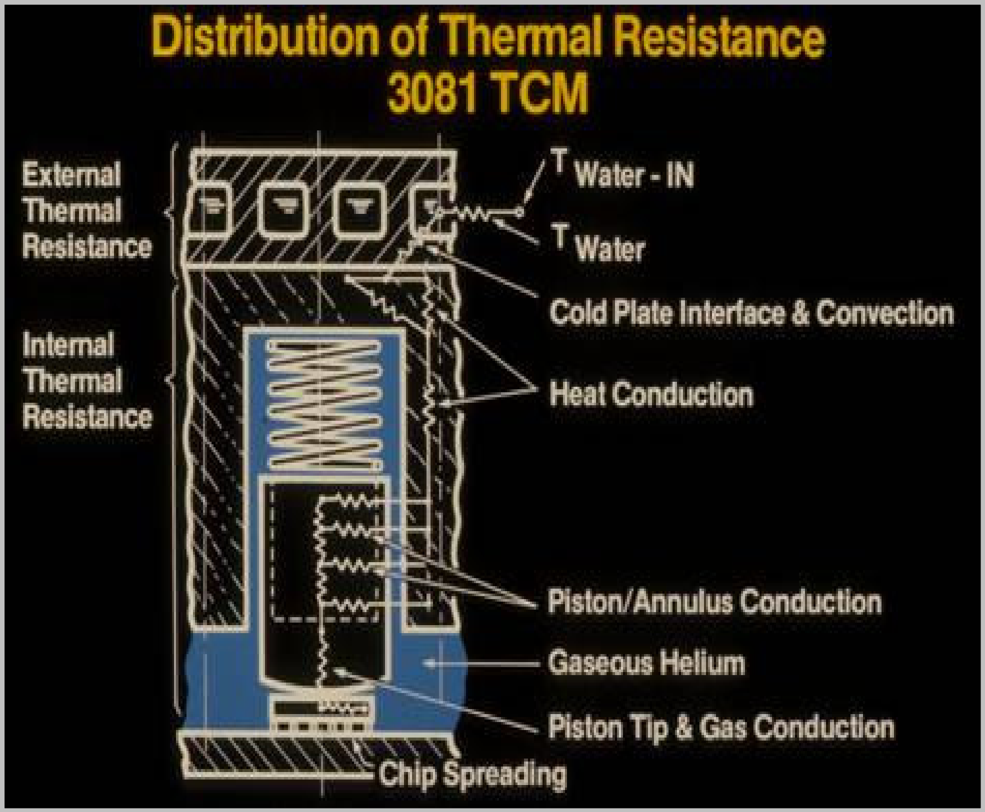

TCM Thermal Analog

Thermal Conduction Modules

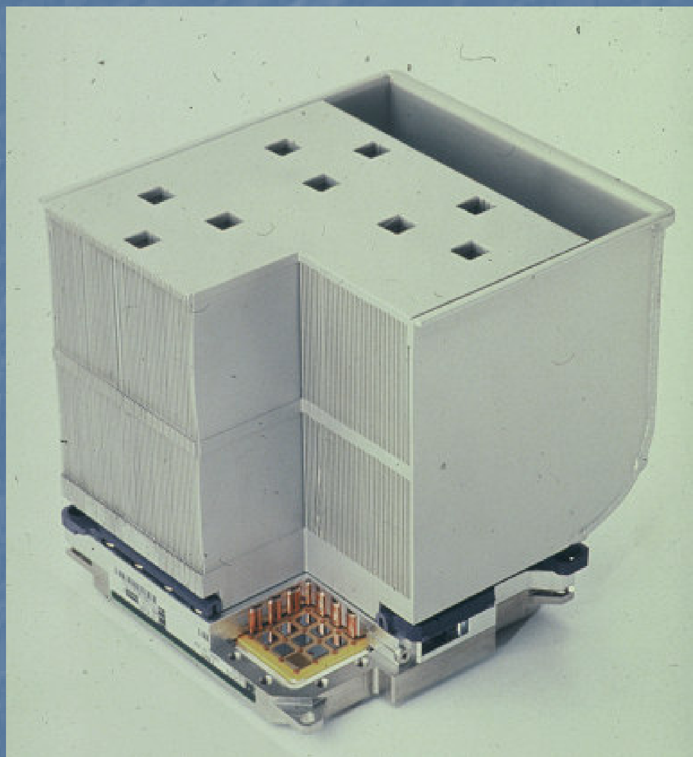

Air Cooled

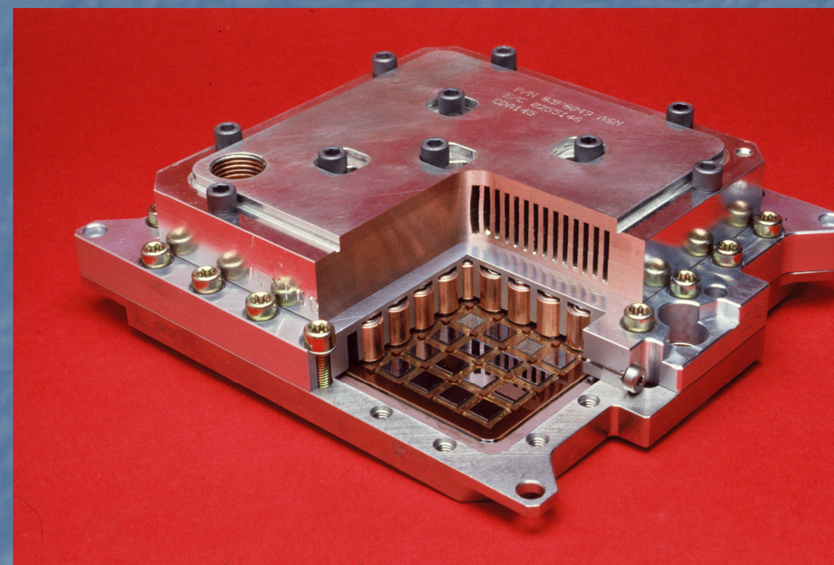

Water Cooled

Flat Plate Conduction Cooling Technology for IBM CMOS-based Servers

- In the mid-1990's IBM moved to server designs based upon CMOS circuit technology to take advantage of the lower costs and power dissipation offered by this technology. The combination of reduced power and improved thermal design within the modules using flat plate cooling allowed these servers to be totally air-cooled. All of the heat dissipated by the processors, memory and power supplies in these systems was transferred to air in the data center.

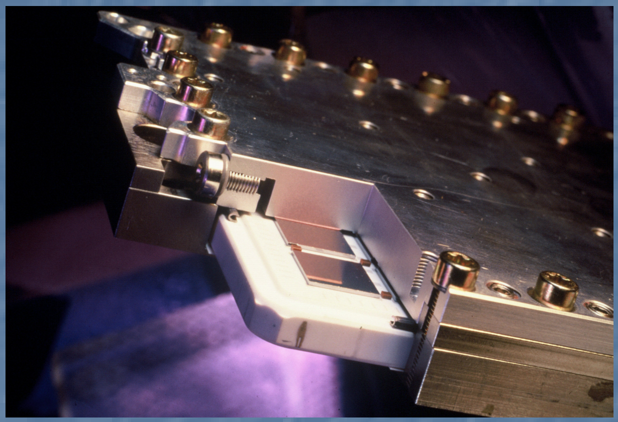

IBM Flat Plate Cooling

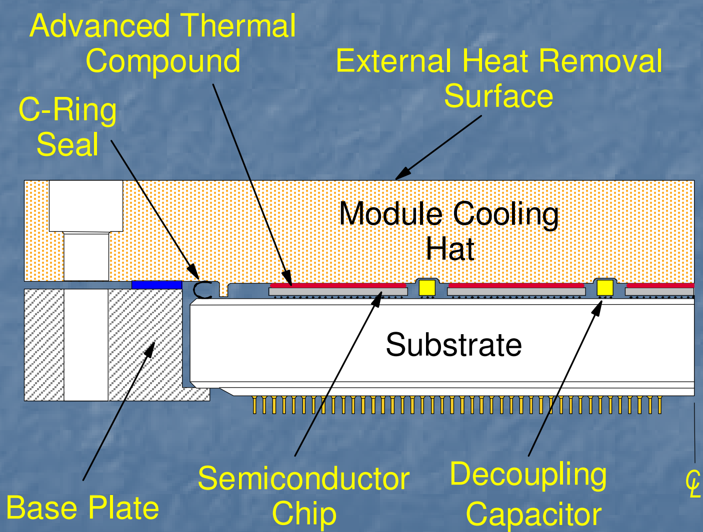

Flat Plate Cooling (ca. 1995)

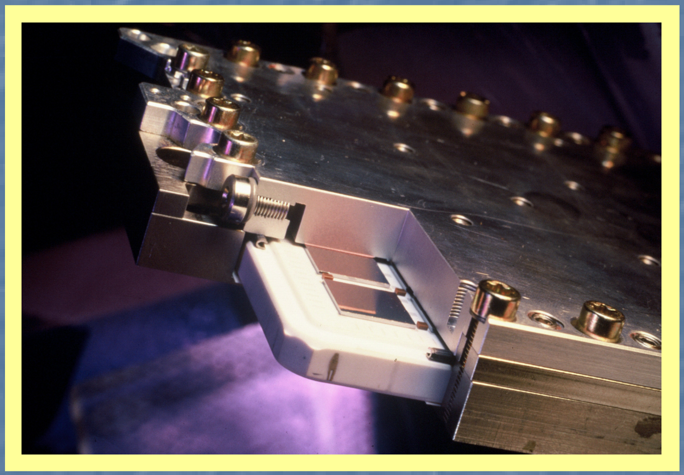

Flat Plate
Cooling Module

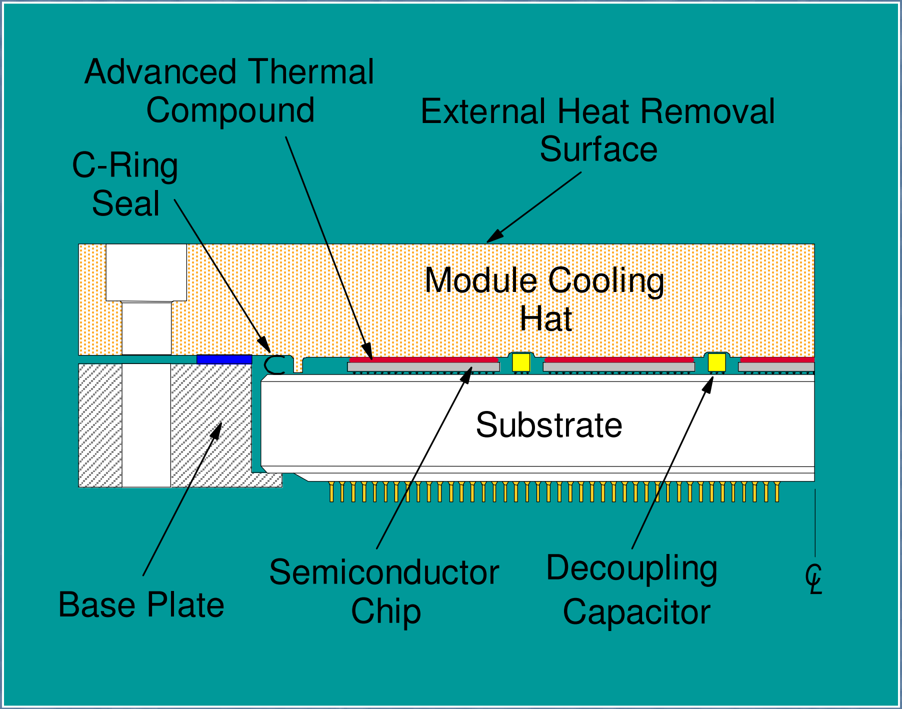

Cross-Section View

Impinging Flow Heat Sinks

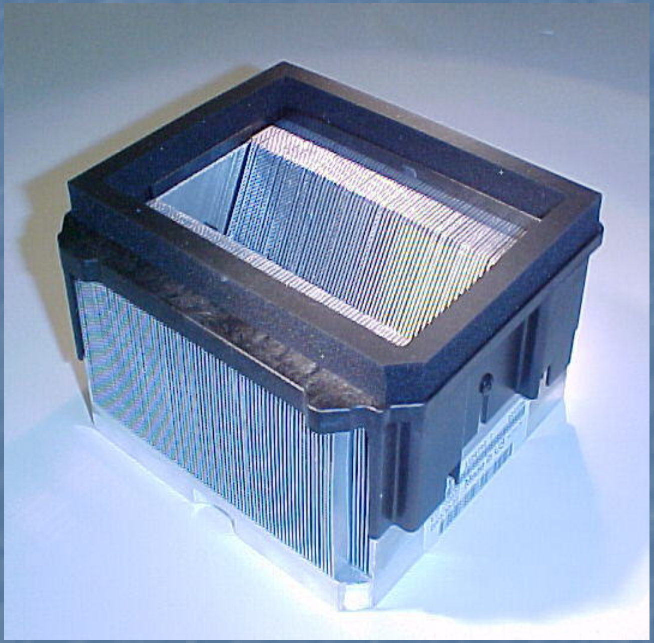

Modular Refrigeration Cooling Technology for IBM G4 → G6 to Z-series Servers

- In the late-1990's a Modular Refrigeration Unit (MRU) was introduced in high-end server designs to reduce chip operating temperatures which improved system computing performance. At the system level the server remained air-cooled rejecting the total heat load to the room air in the data center.

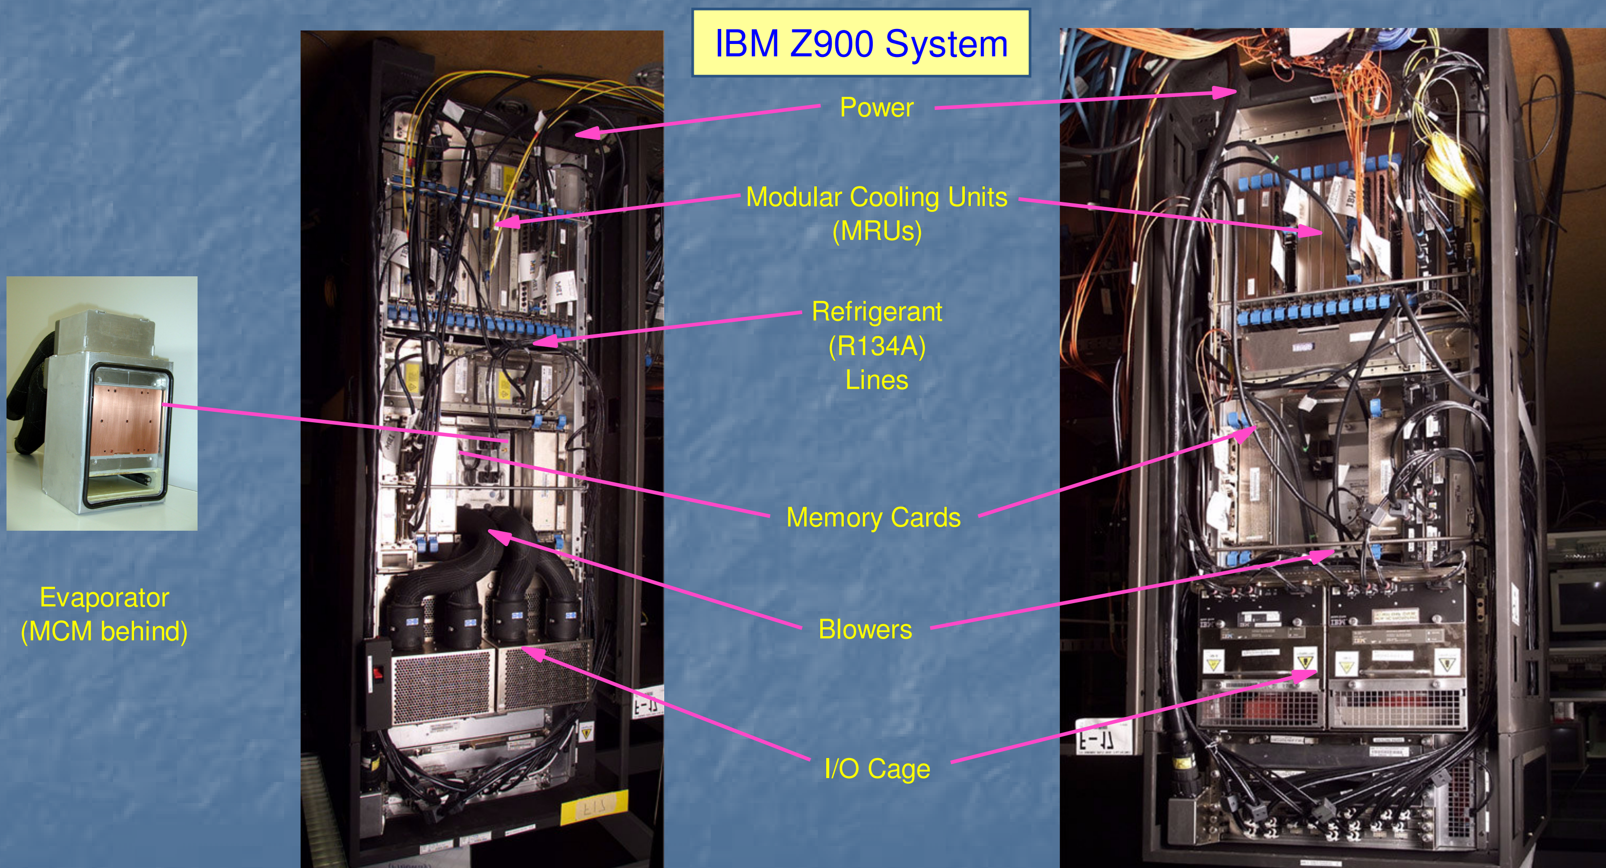

Front Side

Back Side

IBM Refrigeration Cooled S/390 G4 Computer (ca. 1997)

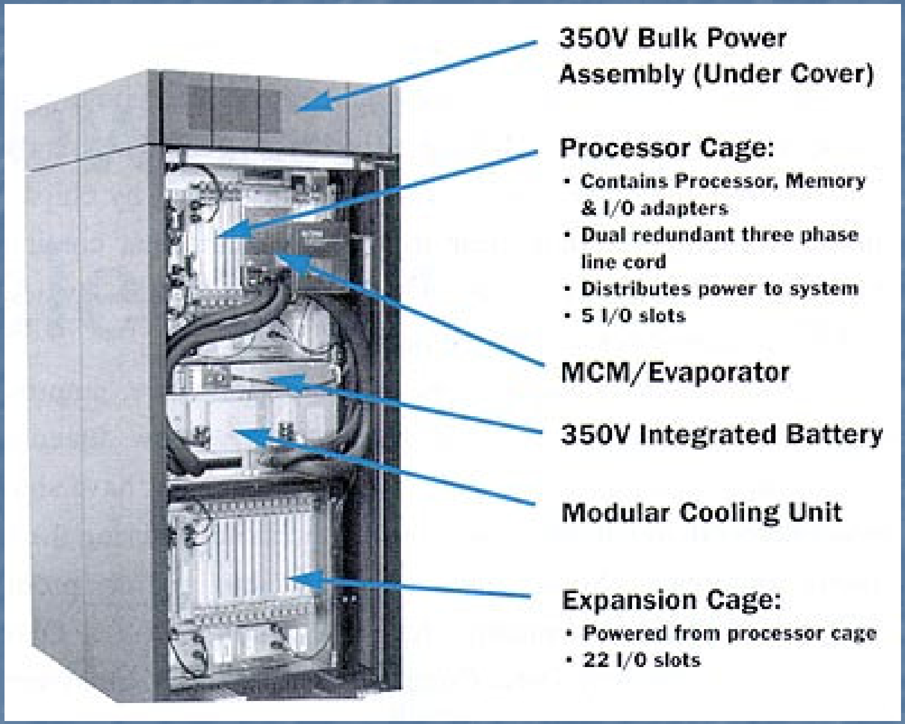

S/390 G4 Refrigerated System

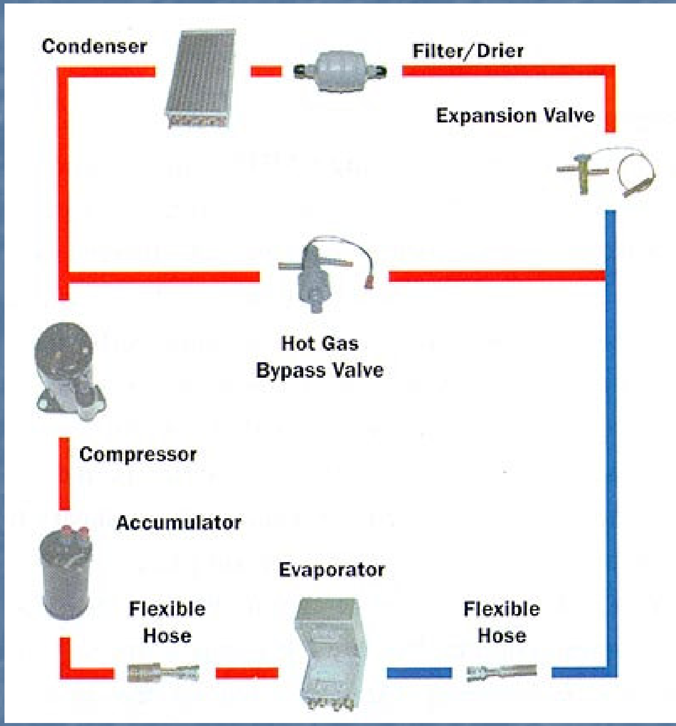

G4 Refrigeration Loop Schematic and Components

IBM zSeries 900 Server Evaporator

Front side – interface to MCM

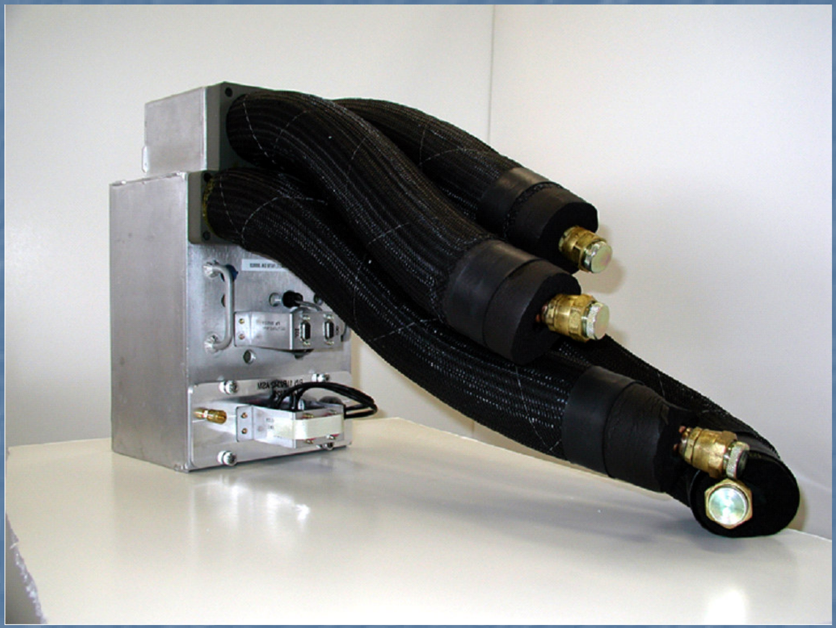

Back side – hoses to MRU

Rack Level Heat Load Trend

- From the 1990's and continuing today, the increased packaging density and heat within IBM server frames as well as others throughout the industry has resulted in an increased cooling challenge at the data center level.

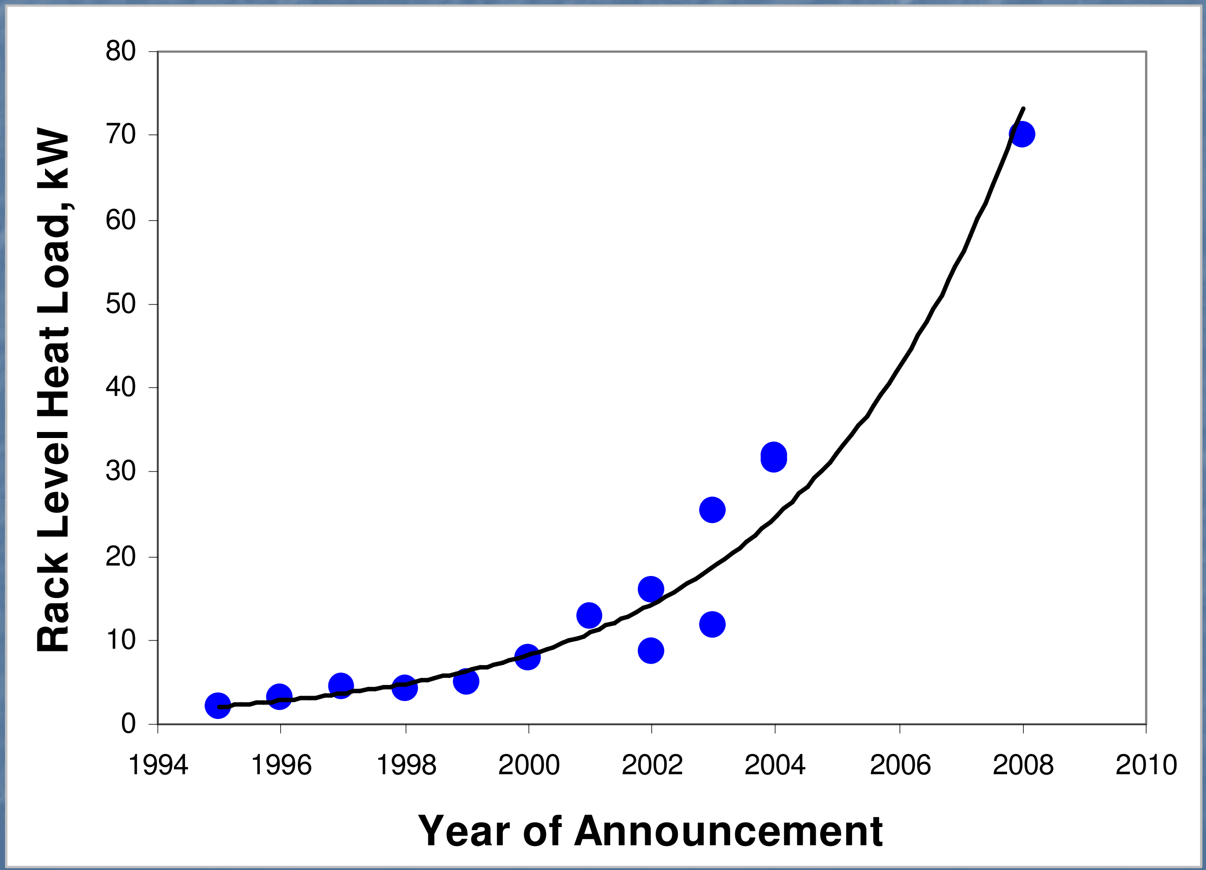

"Cool Blue" Air-Liquid Hybrid Cooling Technology For the Data Center

- With its new "Cool Blue" air-liquid hybrid cooling technology IBM is providing the means to manage increased heat loads within the rack while improving cooling energy efficiency by reducing the amount of heat rejected to data center room air. Using this technology a large portion of the heat from a rack is carried away by water.

Rear Door Heat Exchanger

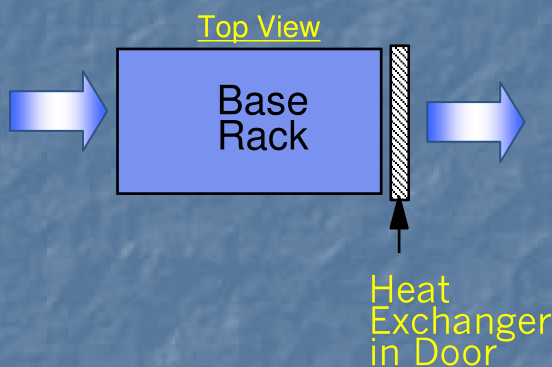

Side Car Heat Exchanger

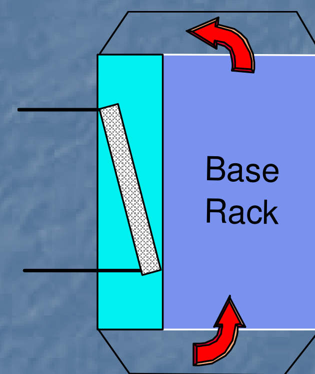

Rear Door Heat Exchanger (ca. 2005)

Air to liquid
heat exchanger

Front to
back air flow

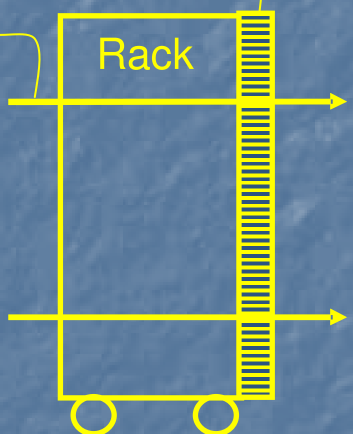

side view

inlet and exit
water carrying
hoses

Row of racks (rear) with rack
mounted heat exchangers

P575 Supercomputer Water Cooling

- Water cooling is back at the high-end in the P575 Supercomputer enabling a 34% performance increase, better reliability and with 78% of the heat load going to water, a 45% reduction in power required to transfer heat to outside ambient.

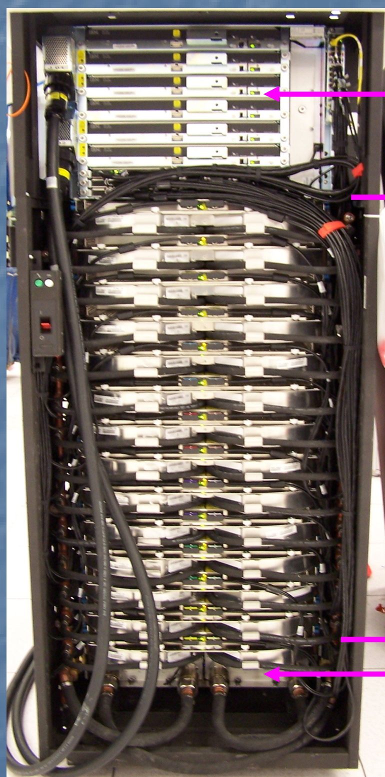

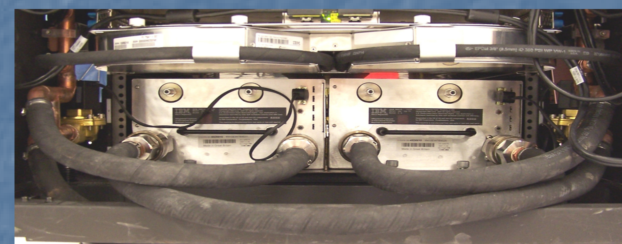

WCUs (System Side)

WCUs (Facilities Side)

P575 Water Cooling System

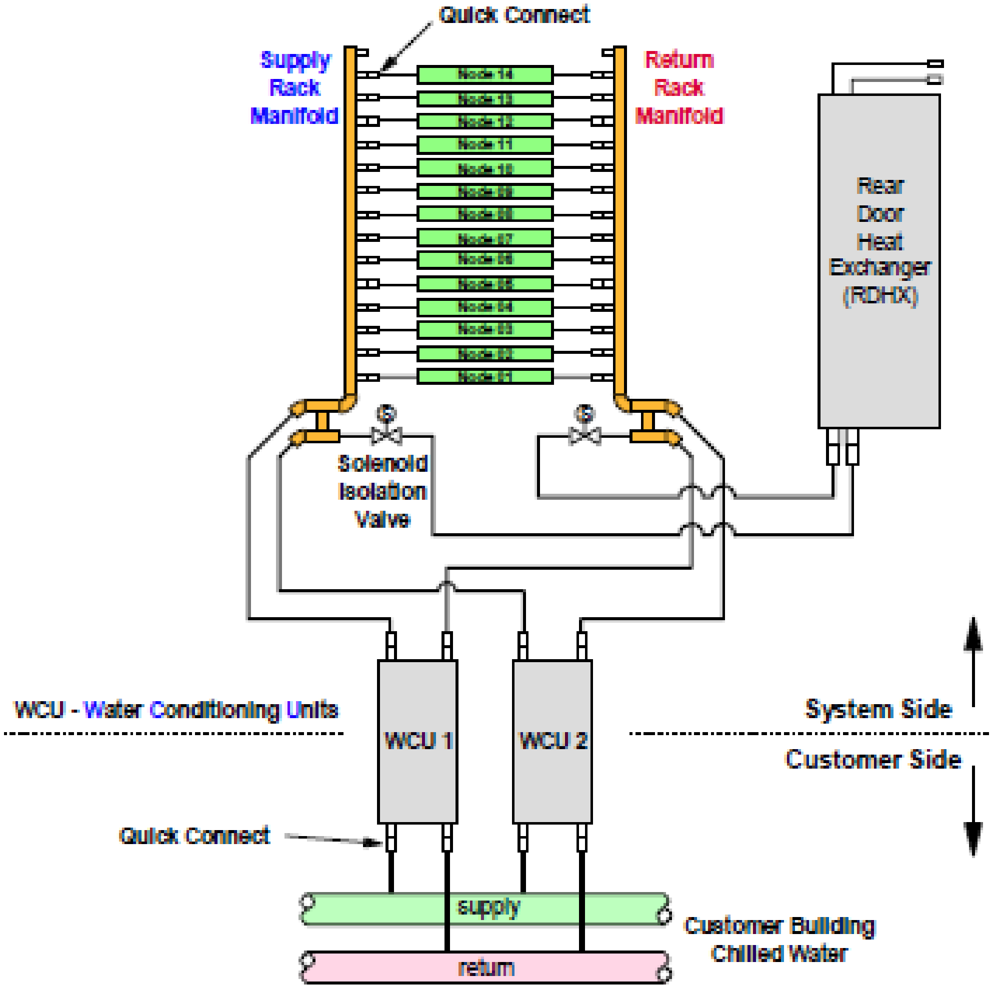

- ★ *Enables 4.7 GHz Processor*
  - *34% performance increase*
- ★ *T_{chip} < 65 °C*
  - *Better Reliability*
  - *Less Power (leakage)*
- ★ *45% reduction in power required to transfer heat to outside ambient*

System Heat Load
max config = 85 kW

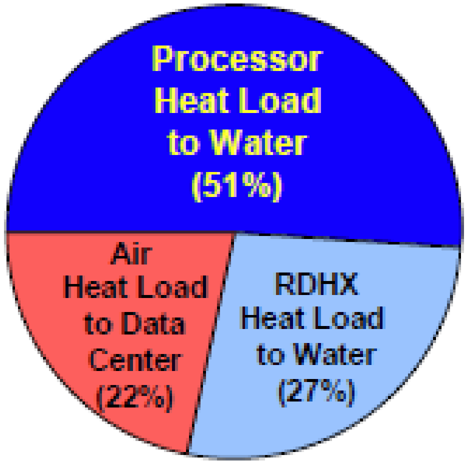

P575 Supercomputer Water Cooled Node

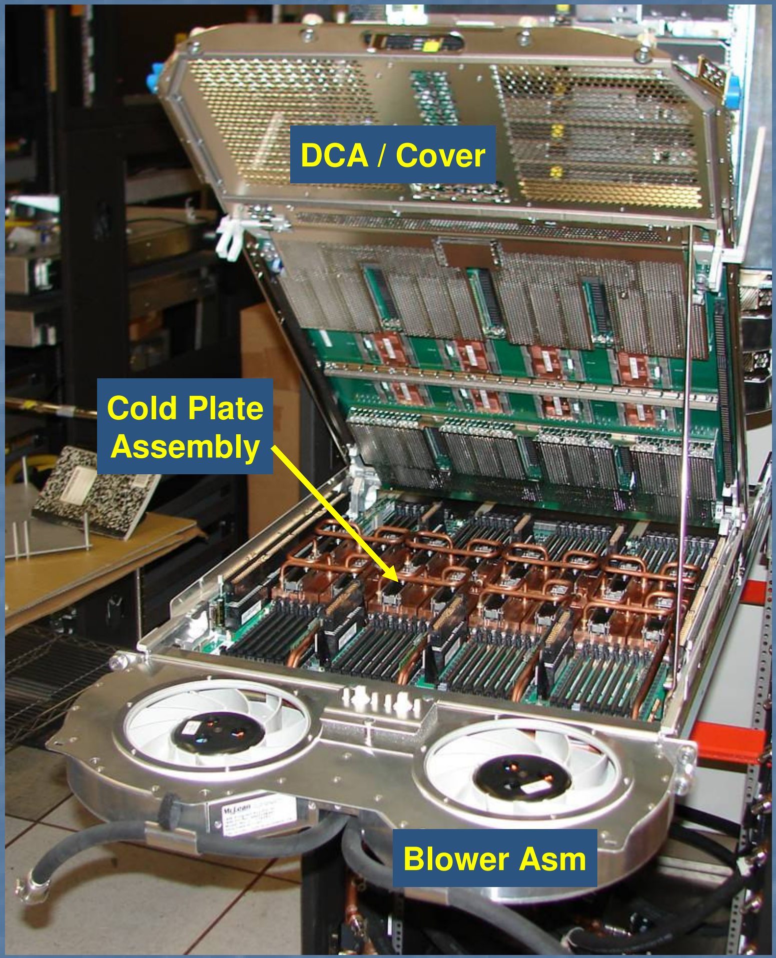

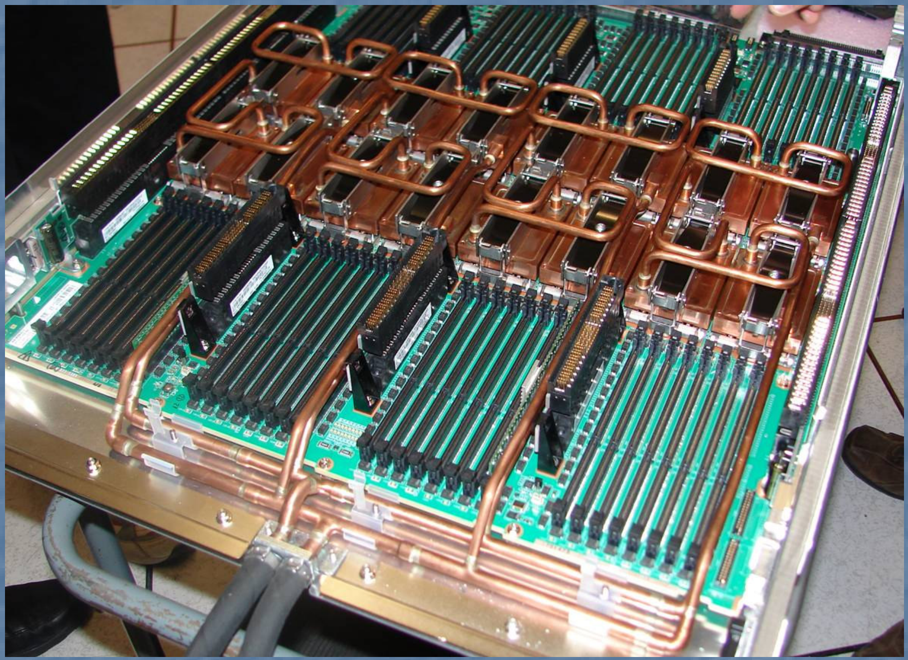

Node Cold Plate Assembly

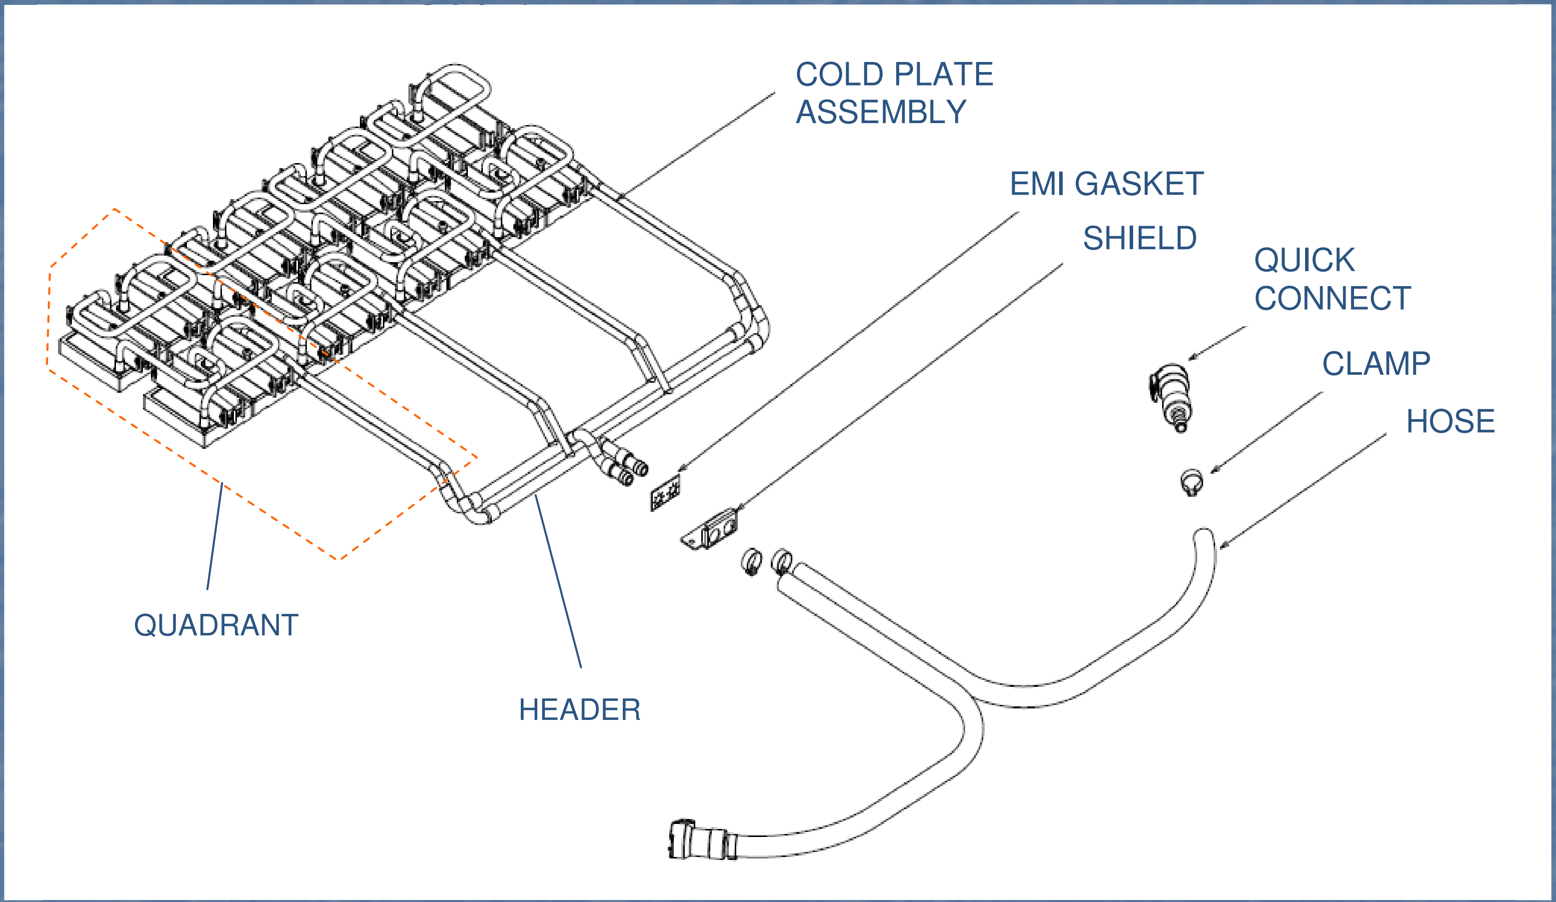

Water Conditioning Unit (WCU)

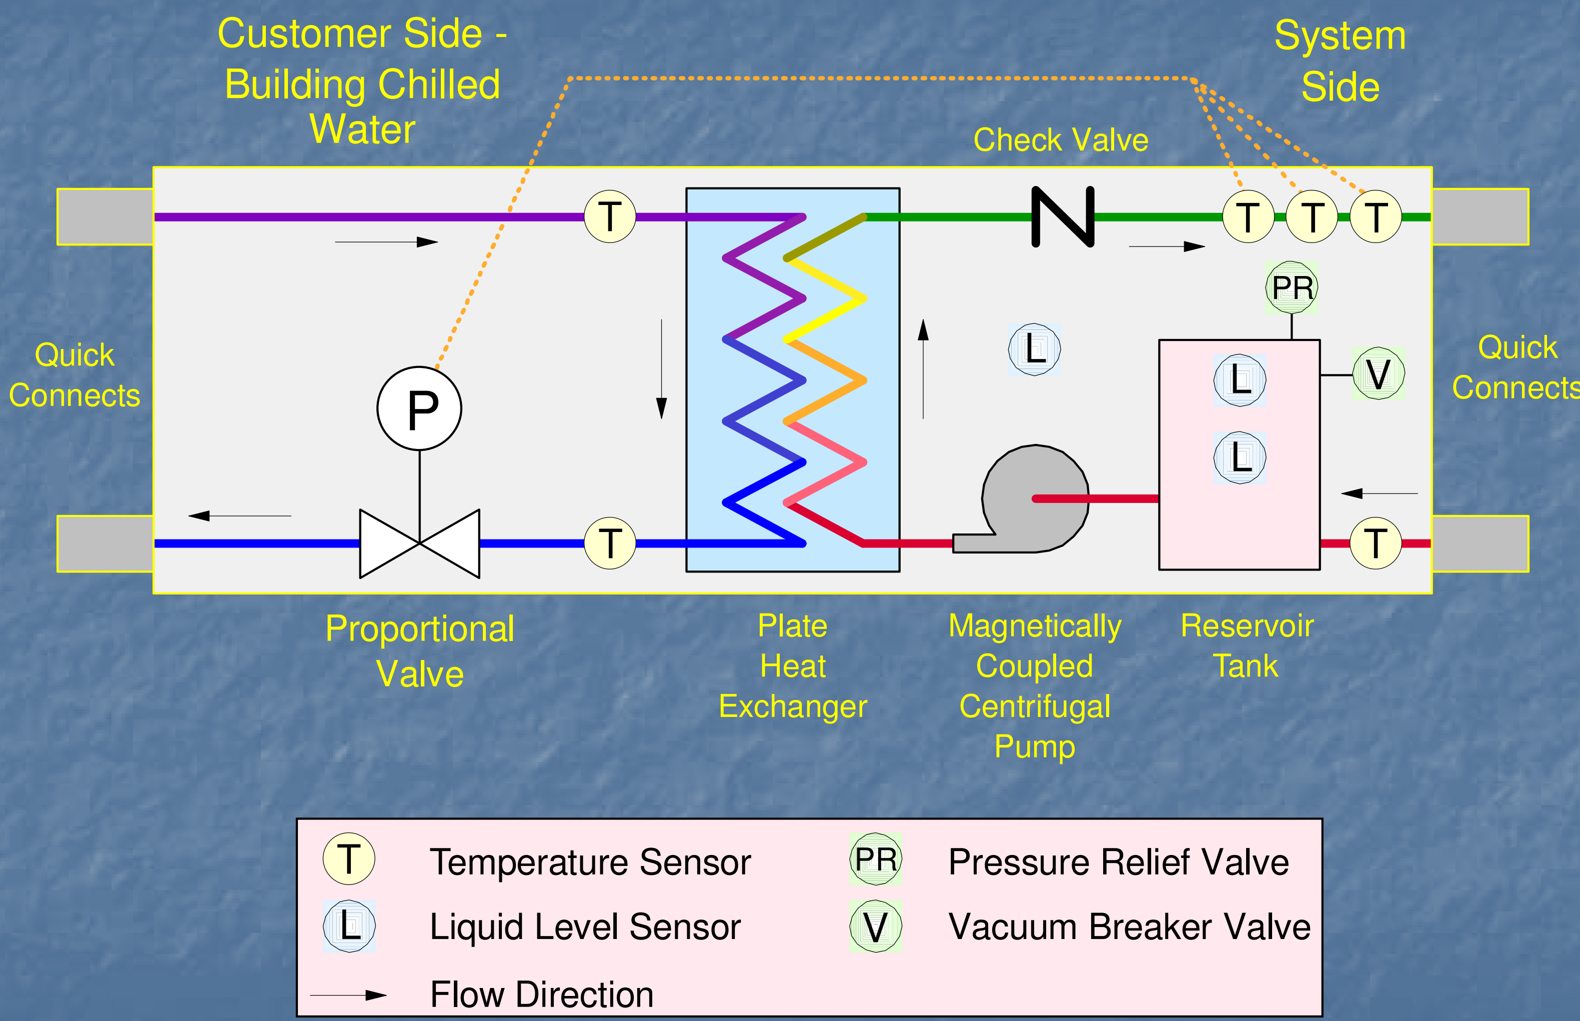

Water Conditioning Unit (WCU)

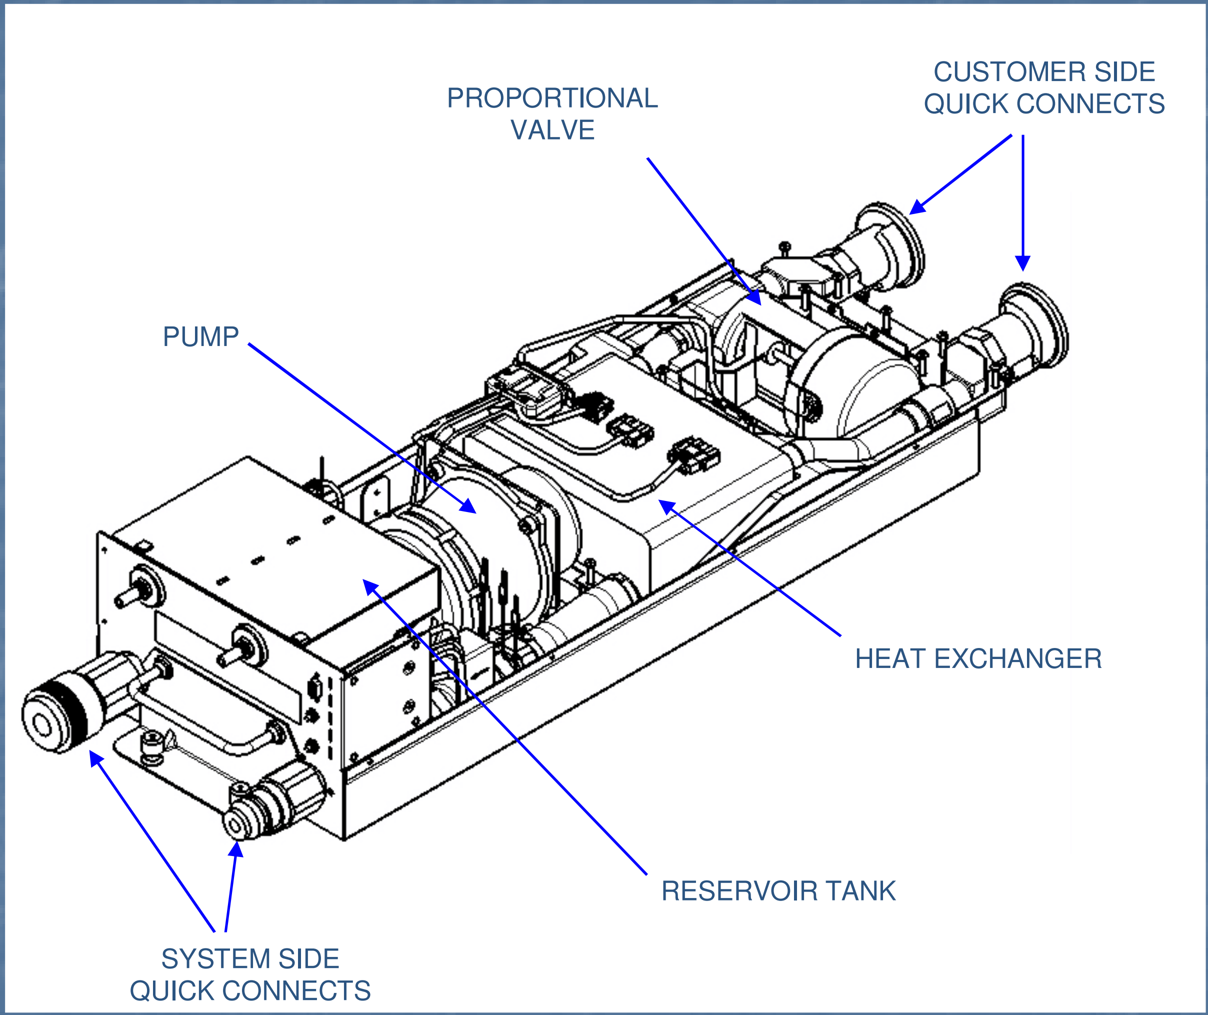

IBM Sponsored Heat Transfer Research

Investigator	University	Area of Research
E. Sparrow	University of Minnesota	Natural convection Channel flow - Effect of barriers - Effect of missing modules Impinging flow - Pin fins - Straight fins
R. Moffat	Stanford University	Channel flow - Adiabatic heat transfer coefficient
R. Wirtz	Clarkson University UNV - Reno	Channel flow - Module height effects Upstream heat transfer Downstream heat transfer
M. Faghri	University of Rhode Island	Channel flow - Entrance, periodic and fully developed
A. Ortega	University of Arizona	Channel flow - CFD code verification
A. Bejan	Duke University	Optimal PCB spacing - Natural and forced convection

Air Cooling

Investigator	University	Area of Research
A. Bergles	MIT, Iowa State University, and RPI	Natural convection Forced convection Pool and flow boiling Boiling enhancement - Ultrasonic - Surface treatment - Extended surfaces
F. Incropera and S. Ramadhyani	Purdue University	Forced convection - Single phase - Extended surfaces - Boiling Liquid jet impingement - w/o Enhancement - w/ Enhancement
I. Mudawar	Purdue University	Falling liquid film Pool boiling Flow boiling
L. Jiji and Z. Dagan	City College of City University Of New York	Liquid jet impingement - Single phase and boiling - Single and multiple jets
V. Carey and C.L. Tien	University of California at Berkeley	Flow boiling Thermosyphon cooling
T. Simon and A. Bar-Cohen	University of Minnesota	Pool boiling - Enhancement

Liquid Immersion Cooling

Range of Heat Fluxes in Direct Liquid Immersion Cooling

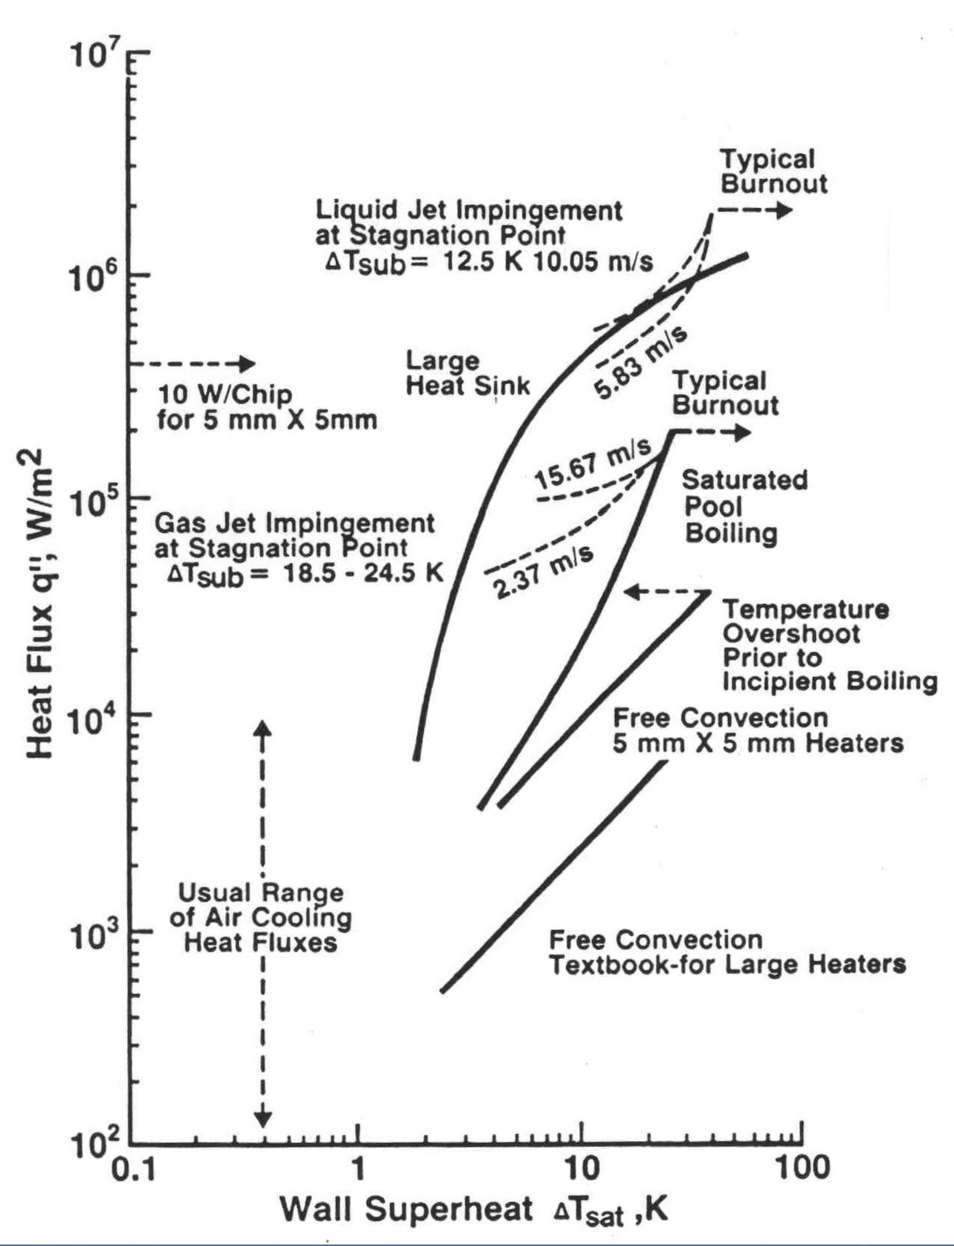

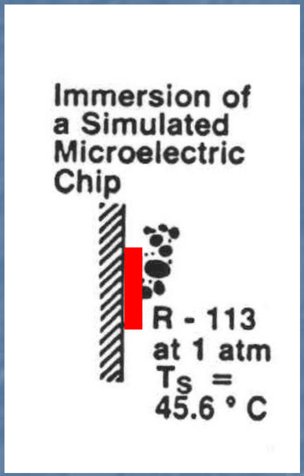

Bergles, 1988.

Purdue Pool and Flow Boiling Results

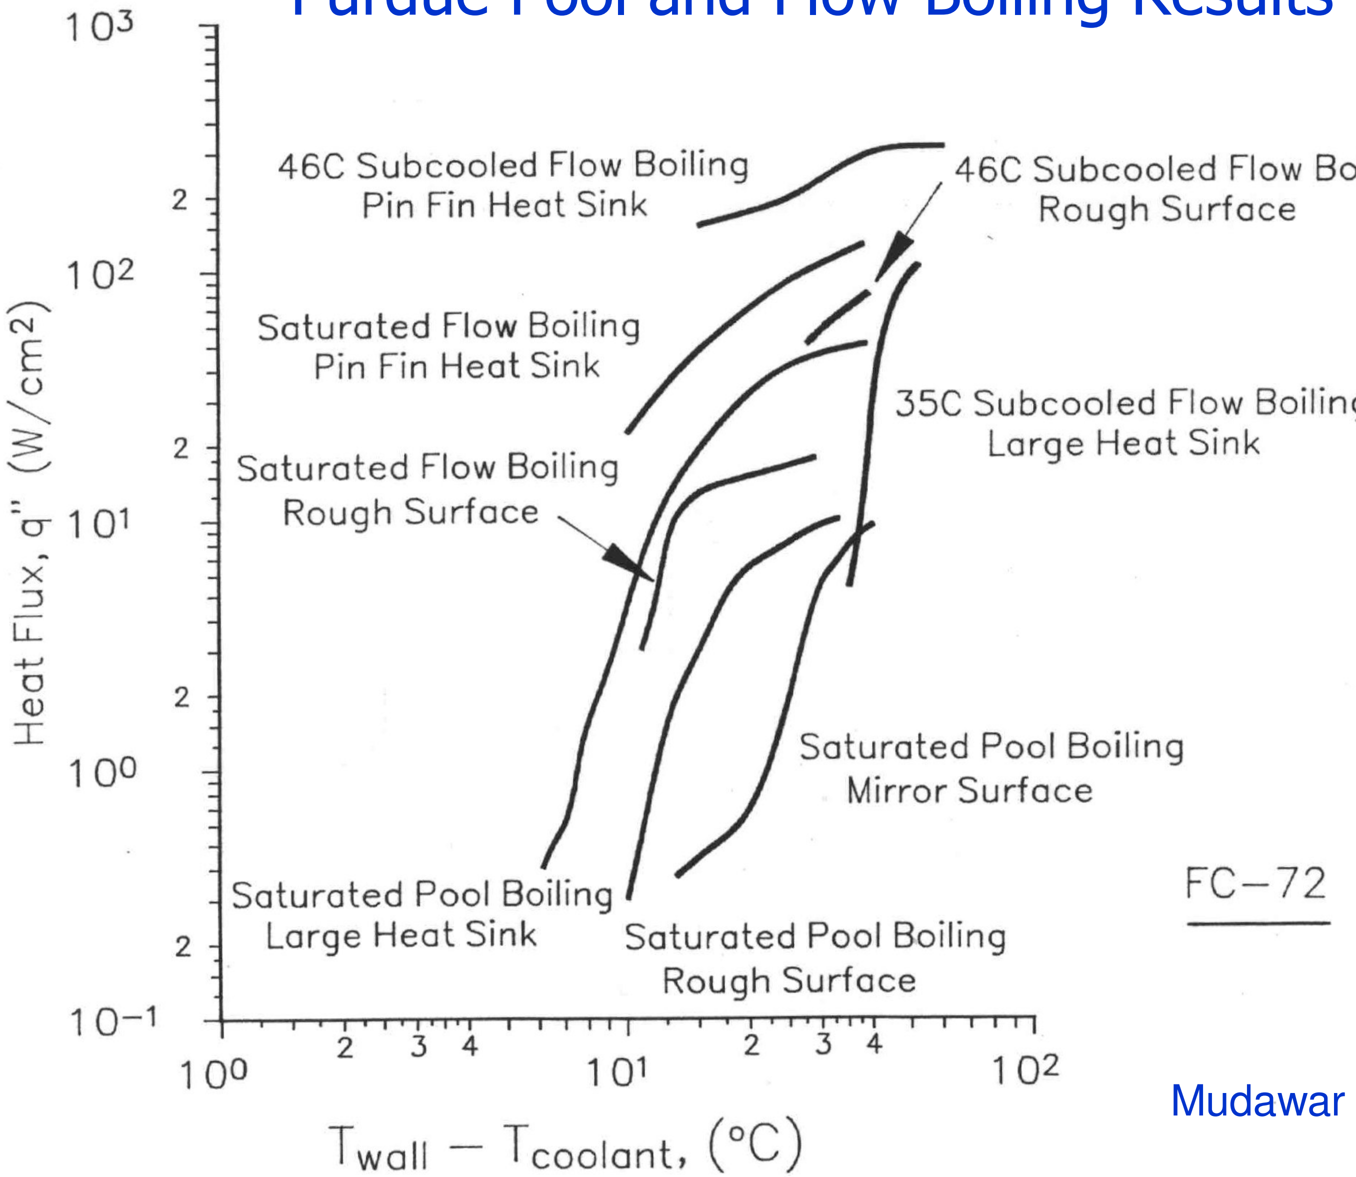

Mudawar et al.

Purdue Forced Liquid Convection Results

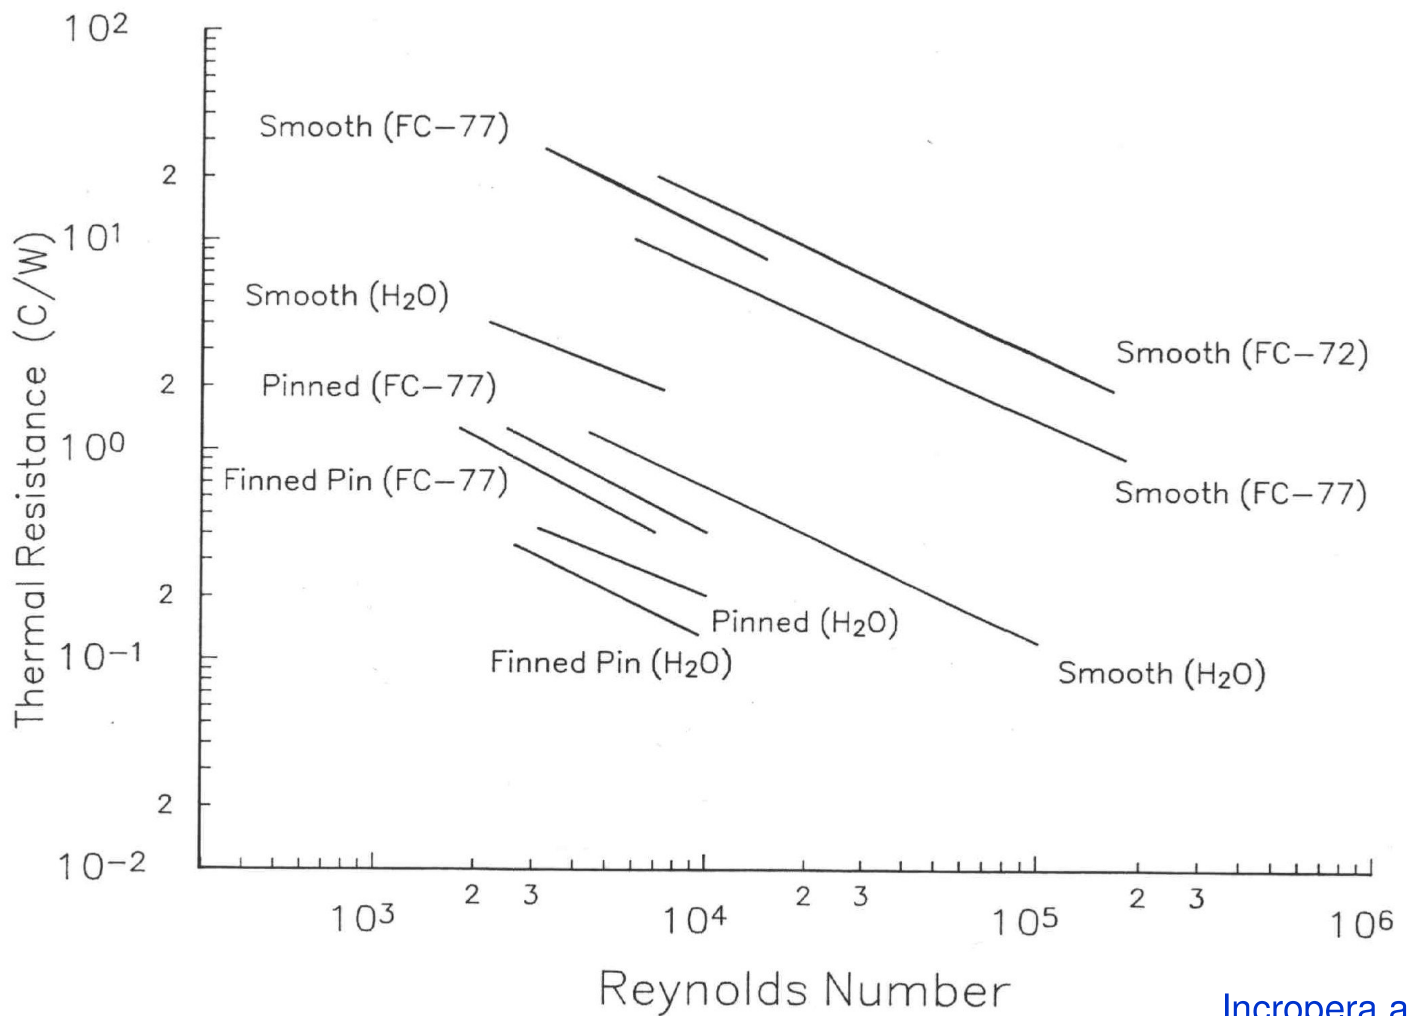

Incropera and Ramadhyani.

Boiling Curve Comparison For Jet Impingement and Channel Flows

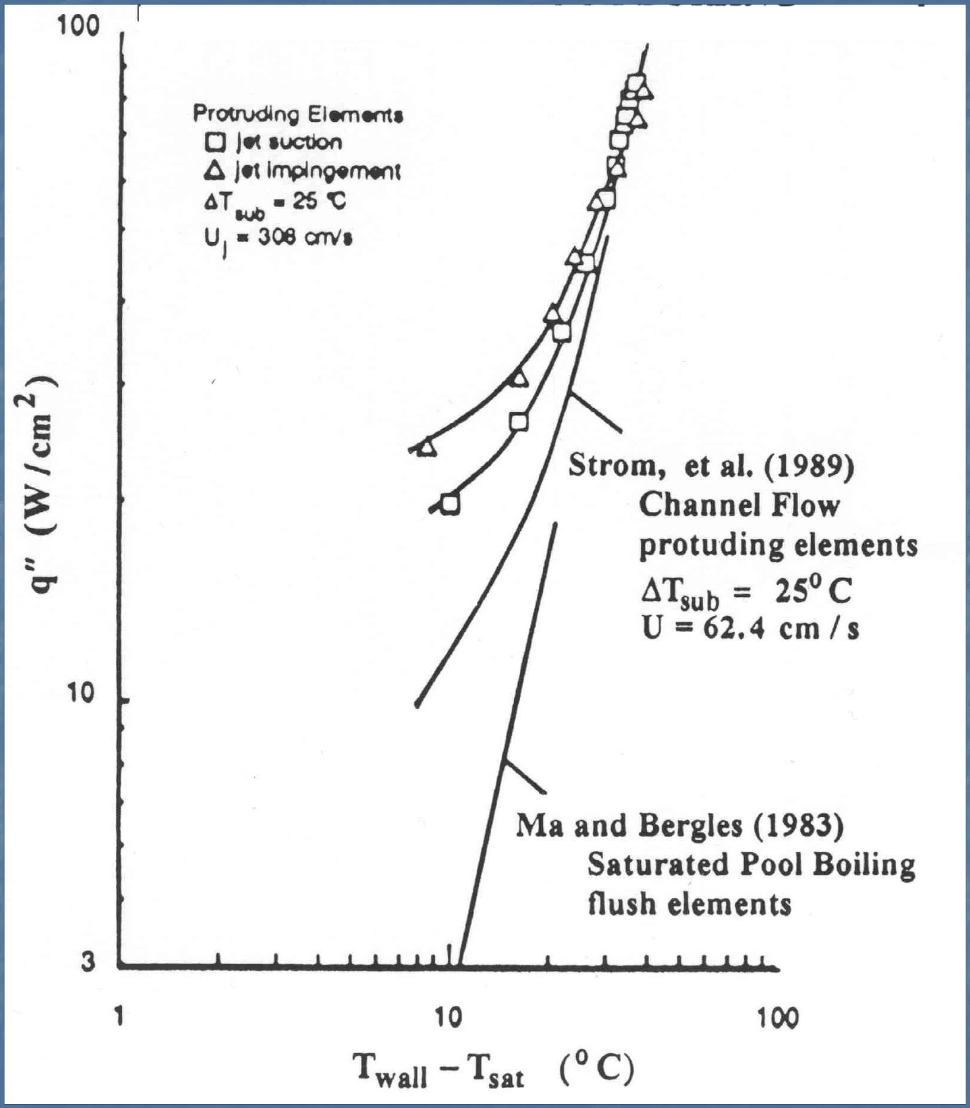

McGillis and Carey, 1990.

Early IBM Hybrid Liquid-Air Cooling System (c. 1969)

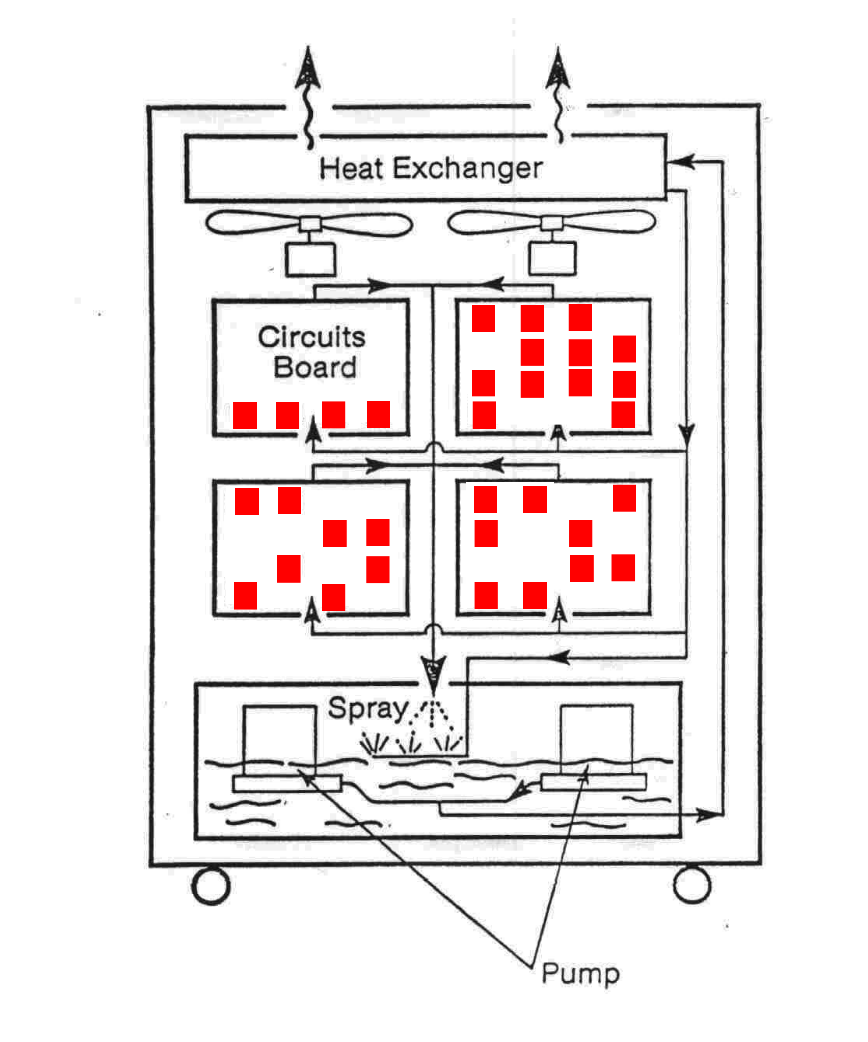

IBM Gravity-Feed Subcooled Flow Boiling System

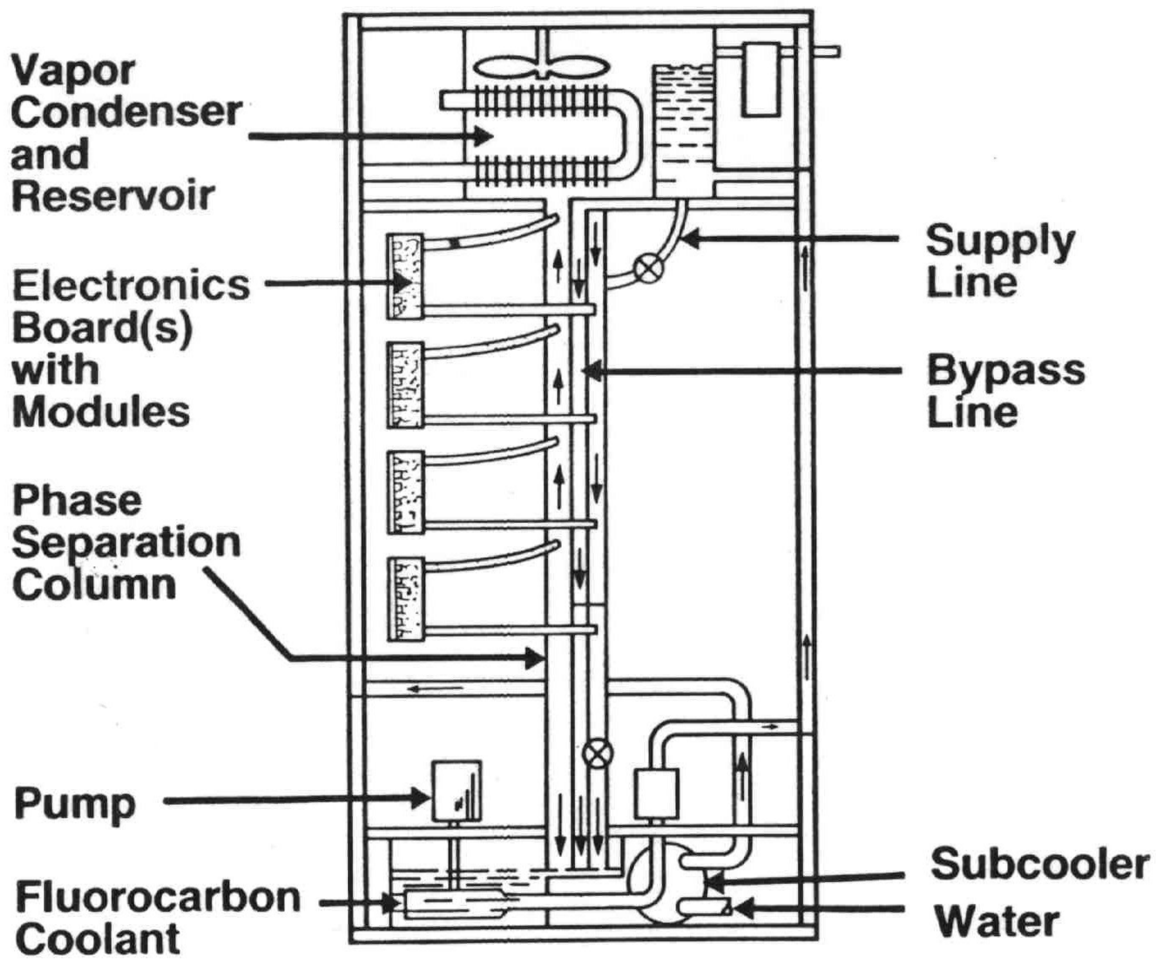

The Multi-Chip Module Cooling Problem

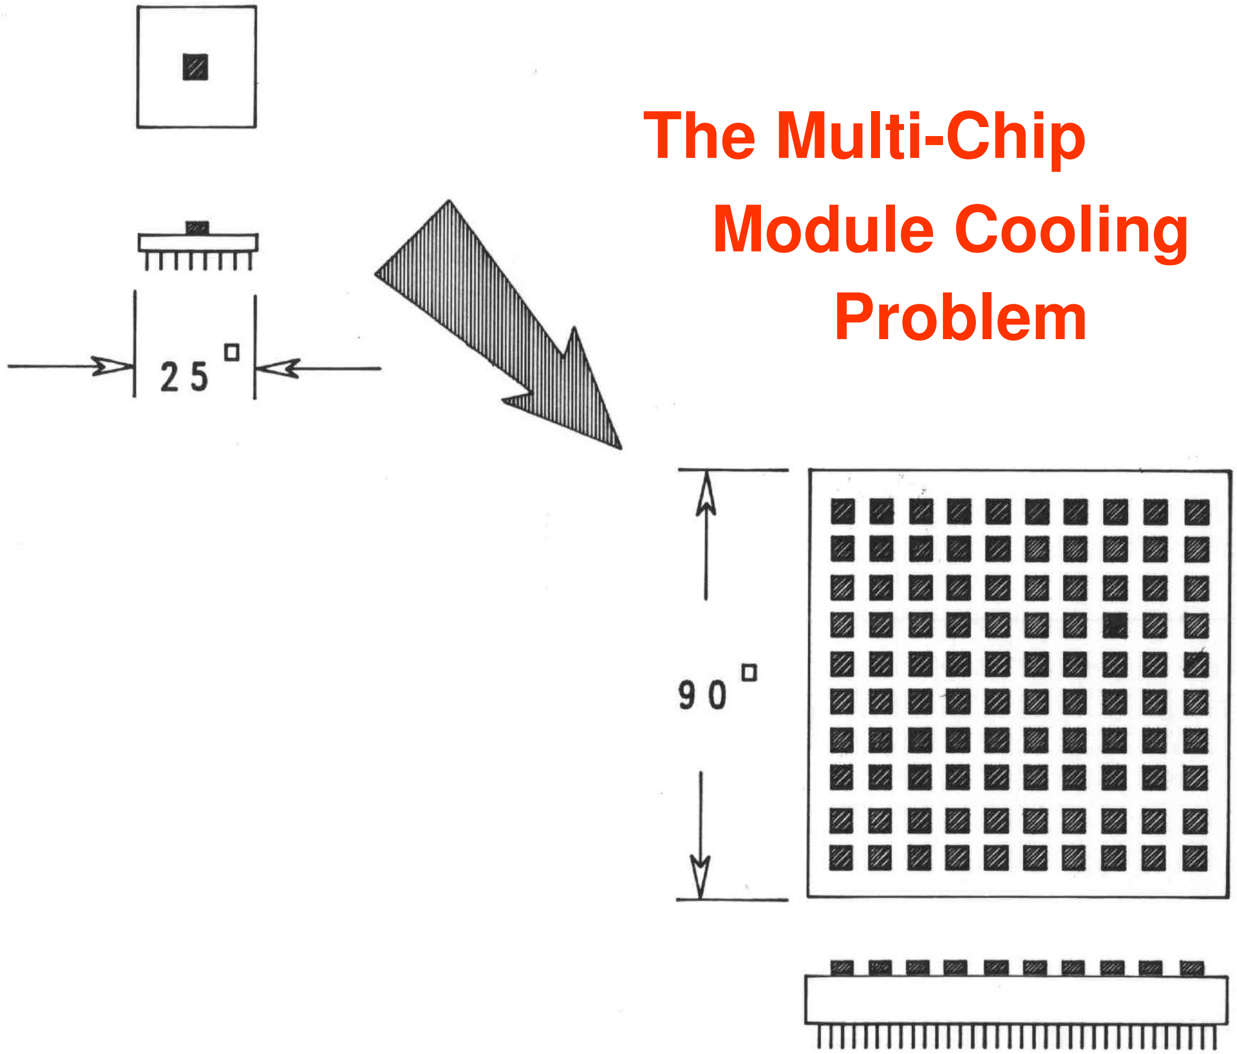

Multi-Chip Substrates Immersed in FC-72 With Boiling on Chips

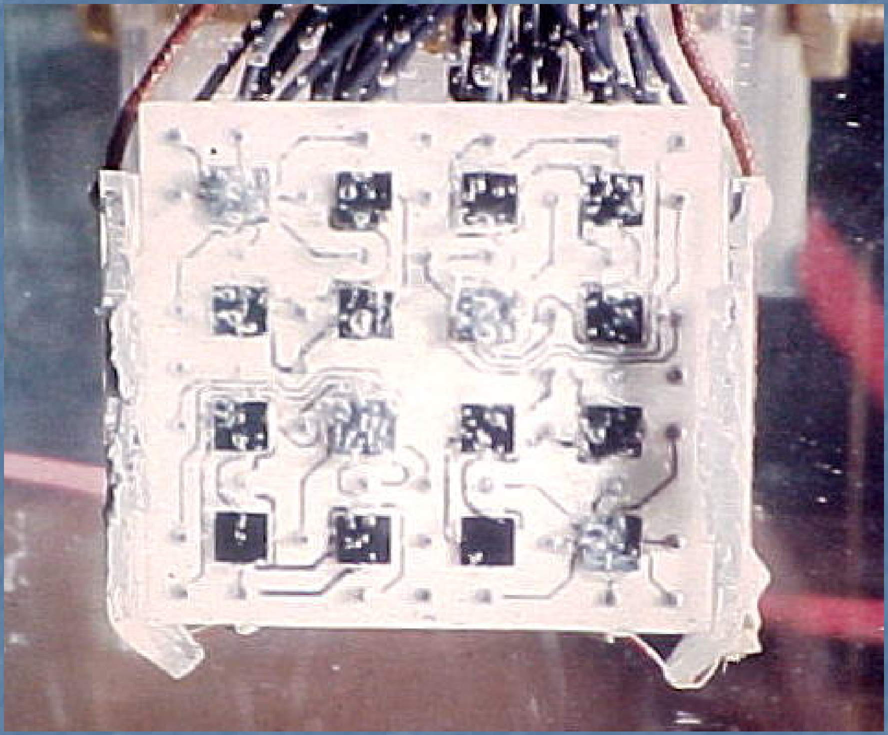

25 mm MC Substrate

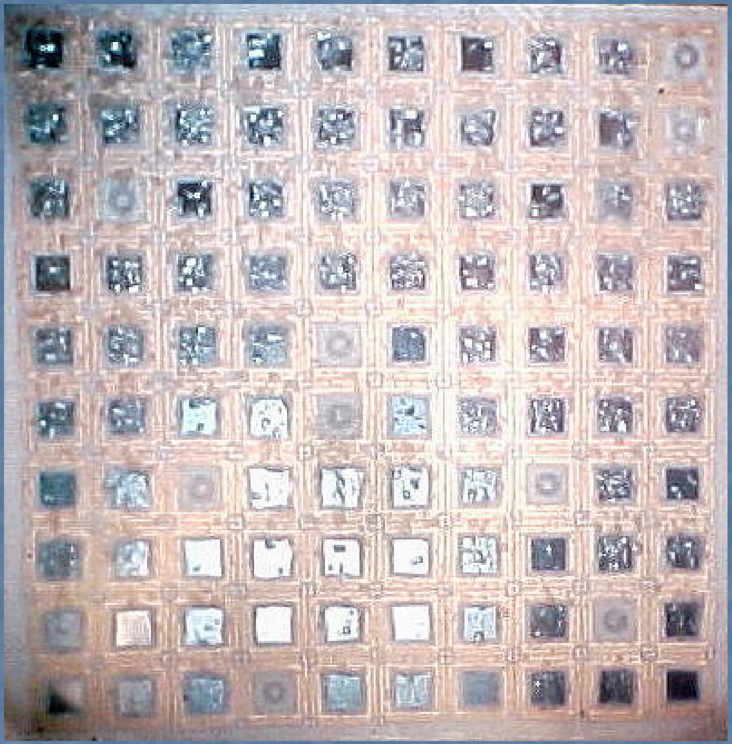

90 mm MLC Substrate

IBM Liquid Encapsulated Module (LEM)

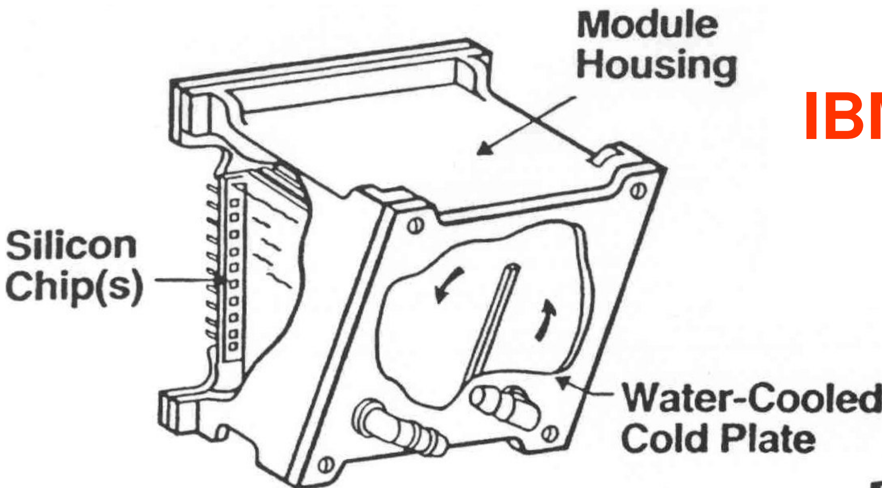

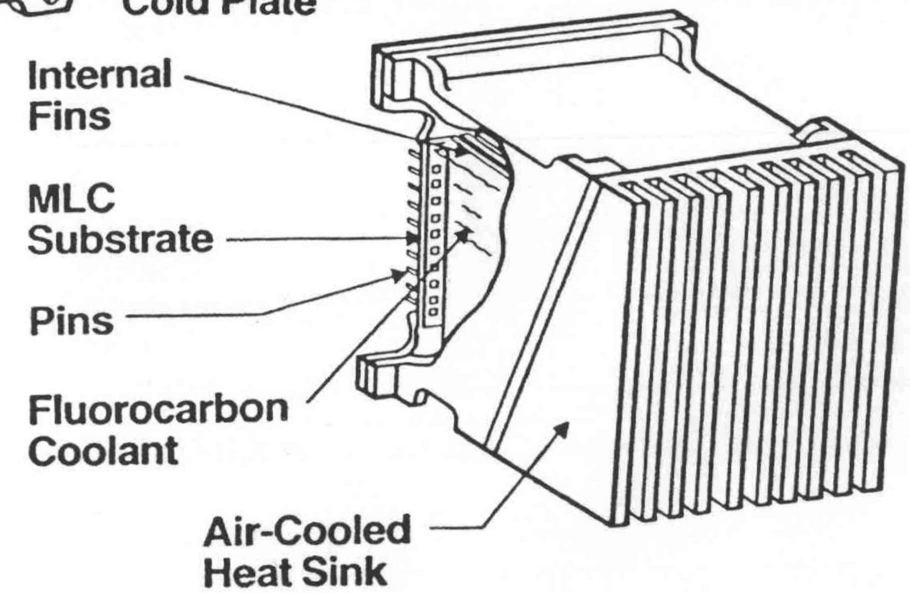

- Coolant - FC 86
- Temperature(s)
  - T_j = 85 C
  - $\Delta T_{j\text{-}c}$ = 3 C
  - $\Delta T_{j\text{-}sat}$ = 20 C
  - T_{sat} = 62 C
  - T_{liq} = 59 C
  - T_{H2O} = 24 C
- Power(s)
  - Chip = 4 W
  - Module = 200 W (air)
  - Module = 300 W (H2O)

Thermal Hysteresis on a Silicon Chip Mounted on a a Substrate Immersed in FC-86 Coolant

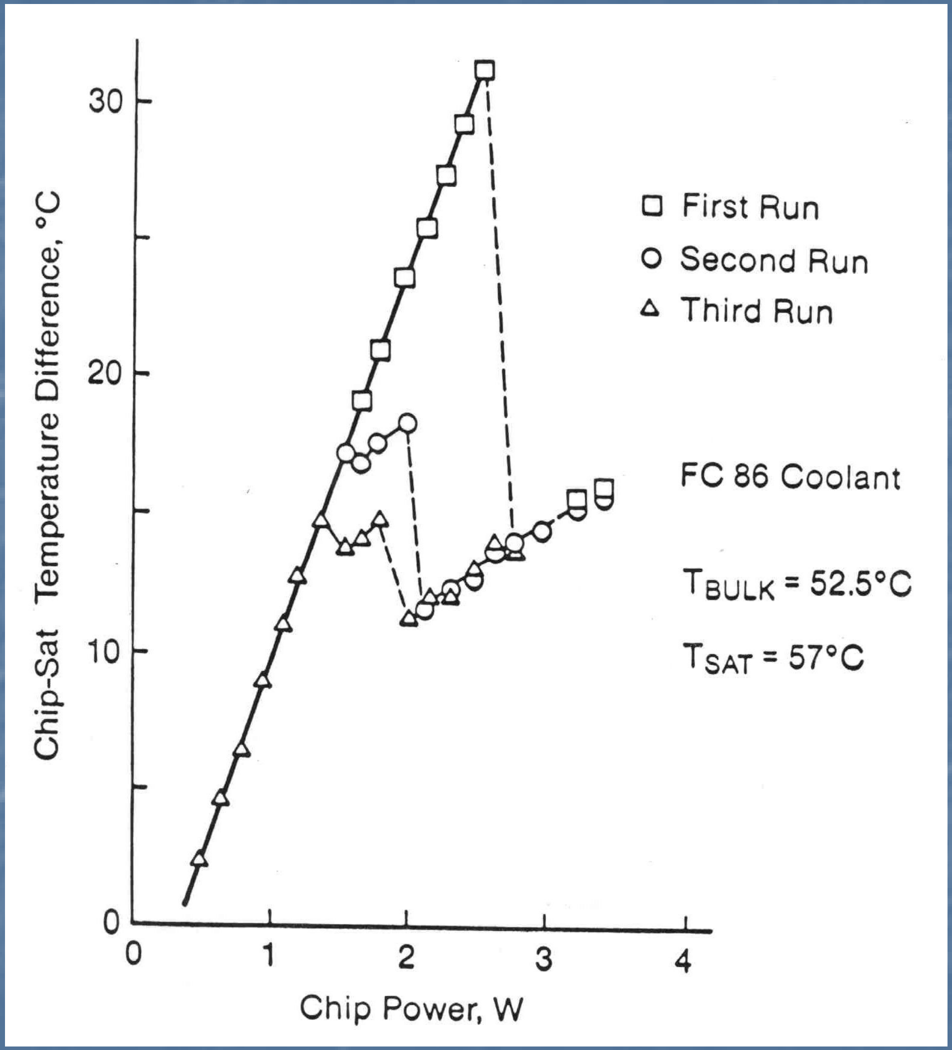

Some Techniques to Enhance Chip Pool Boiling Performance

- FC 86 Coolant

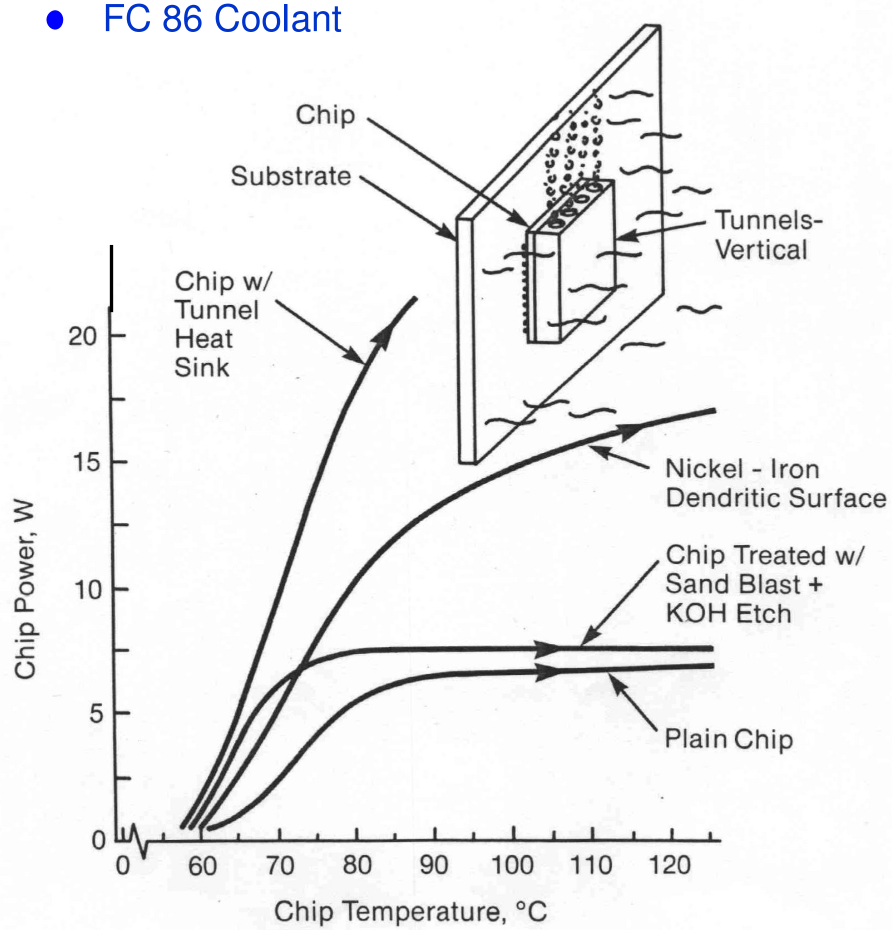

IBM Chip Subcooled Pool Boiling Data

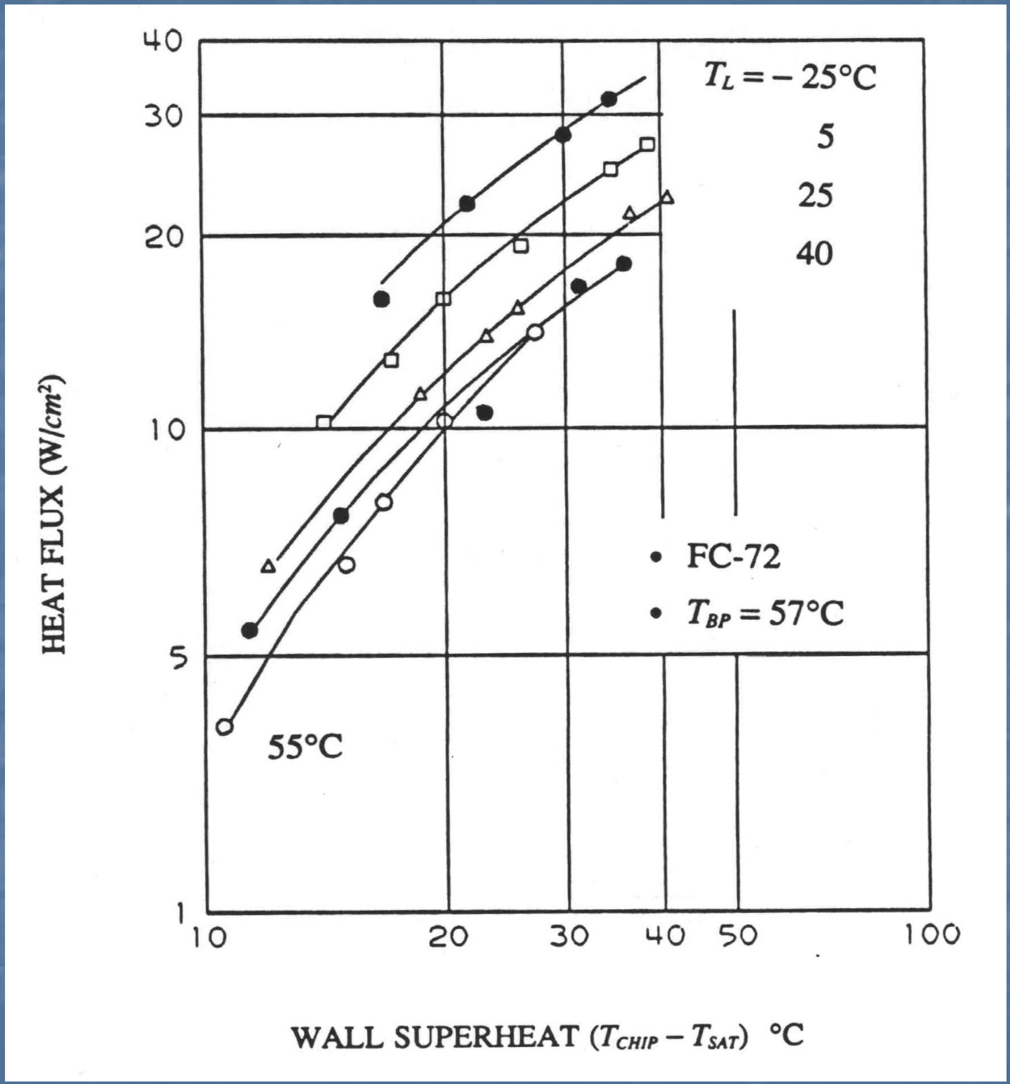

IBM Chip-on-Substrate Test in Fluorocarbon Bath

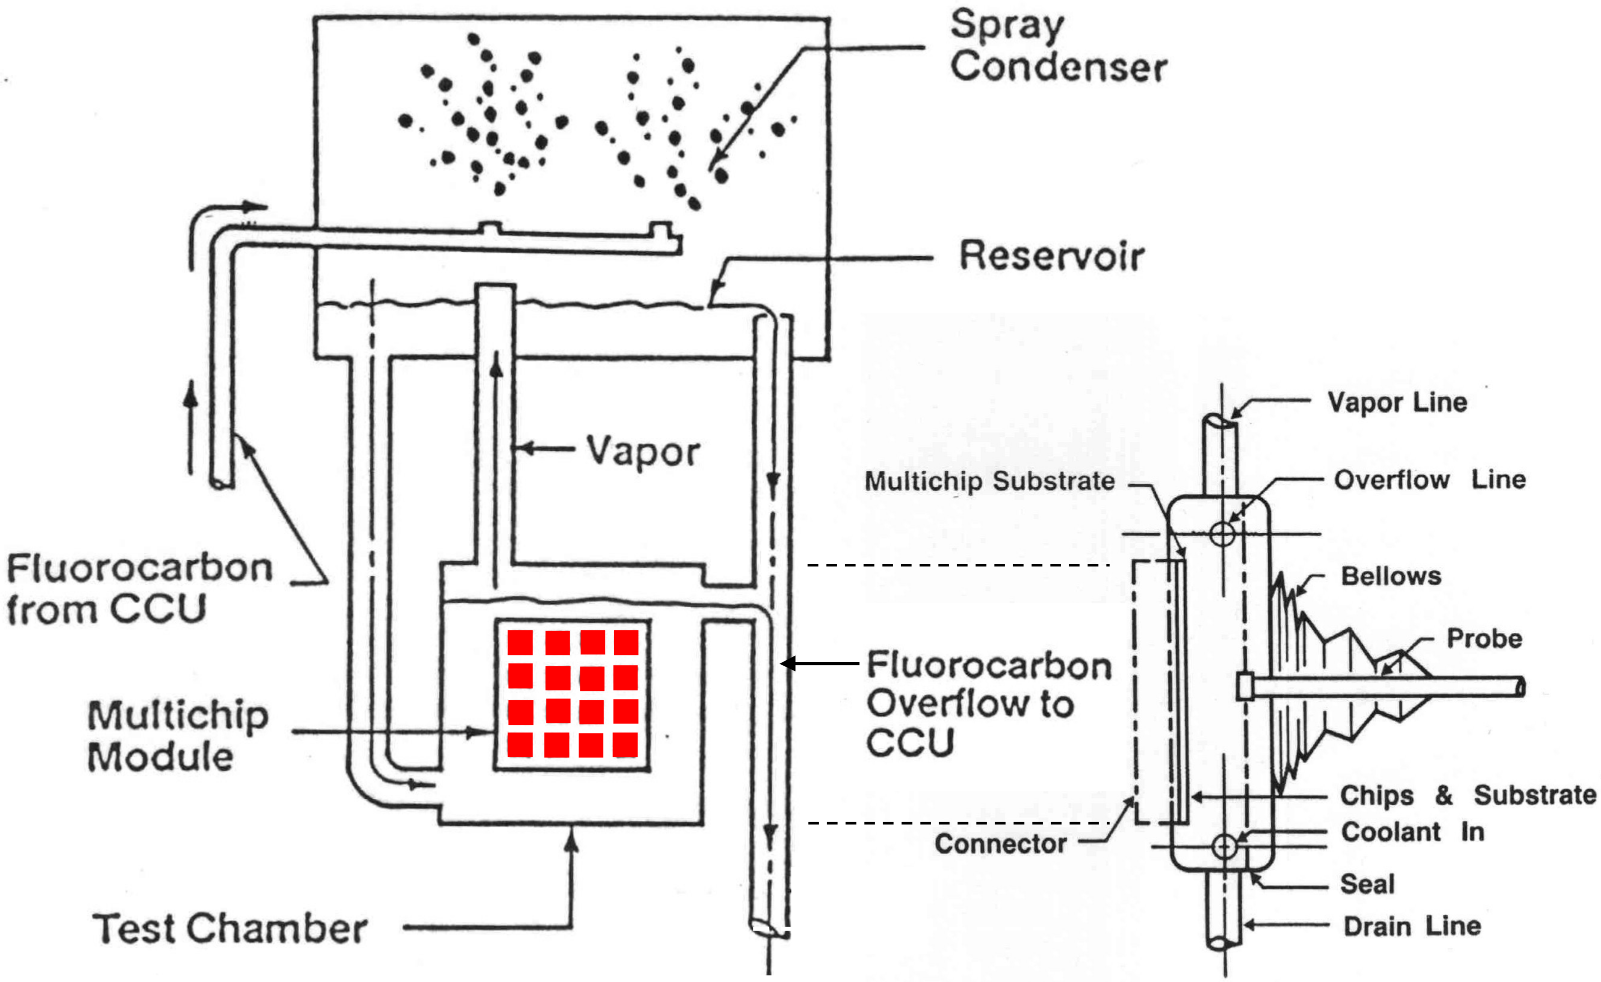

IBM Liquid Jet Impingement Data

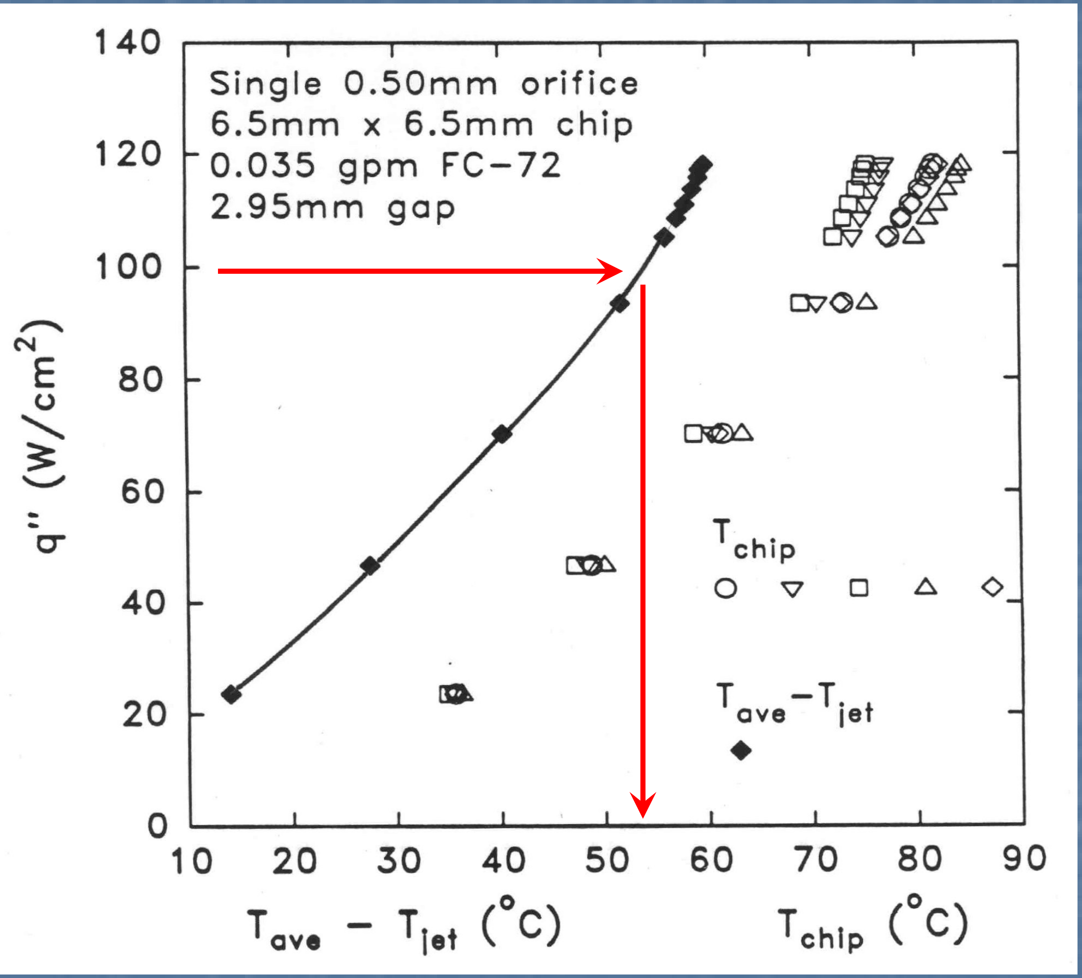

Liquid Jet Impingement Cooling for SSI Supercomputer

Cross-section view of chip-substrate package(s) and jet impingement plate(s)

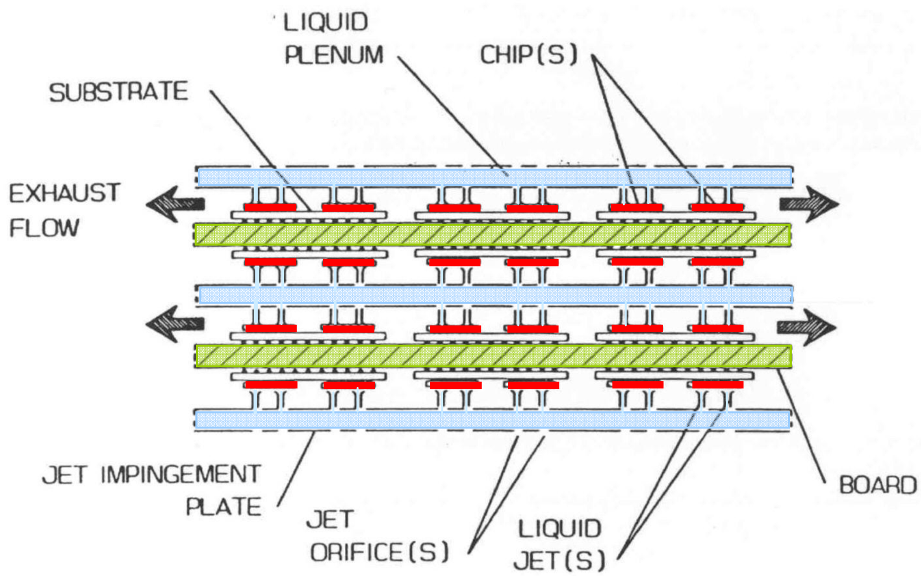

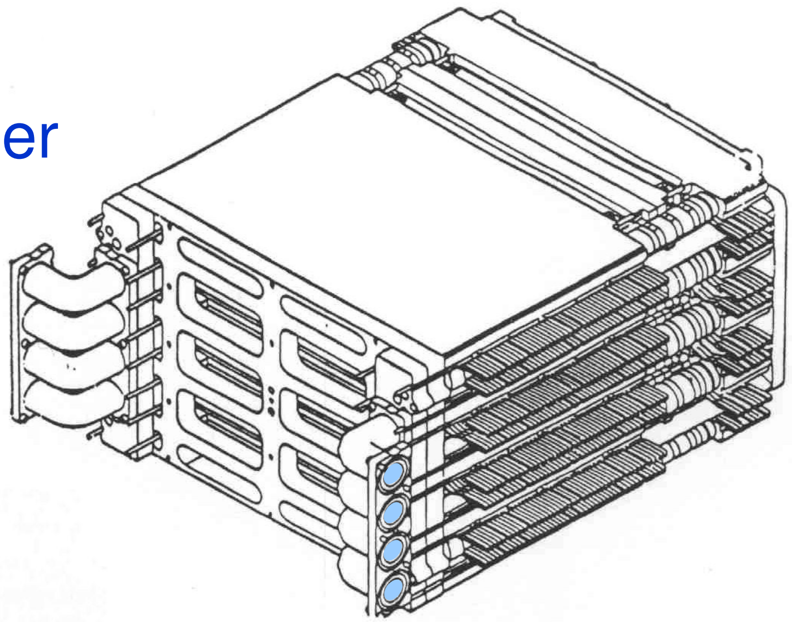

CPU LogicBrick Brick Assembly

Outlook

- Liquid cooling will become more pervasive and will migrate closer to the server electronics.
- Economizer based facility level cooling architectures have become more widespread due to the desire to eliminate refrigeration chillers.
- Future pervasive use of warmer coolant based cooling will co-exist with strategic need based localized refrigeration cooled devices.
- Servers and data centers will include sophisticated sensors and real time control algorithms to manage for better energy efficiency.
- End-to-end design will most likely result in cost and energy optimized computer systems with a life cycle perspective.
- Cooling designs will be needed for much wider range of operating conditions with requirement of cost/energy scaling with IT power.

Potential Research Areas (data centers)

- **Waste Heat Recovery**
  - use of on-site power generation with absorption chillers
  - use of waste heat for building heating
  - exploration of waste heat conversion into higher quality work
- **Weather Optimized Cooling System Operation**
  - local weather optimized data center hardware design
  - extensive control of cooling infrastructure optimized for local weather
  - exploration of pre-emptive control based on local weather forecasts
- **Energy Proportional Cooling Technology**
  - cooling energy use scales linearly with IT power from 0-100% range
  - control of variable speed fluid moving devices in the data center
  - cost effective design of liquid cooling in servers
- **System Level Modeling, Measurement, and Control**
  - generate tools/models for lowest TCO and energy efficient design
  - enable real-time prediction and control - thermal/energy parameters
  - increased focus on reliability models and data to drive design decisions

Some References

R. Schmidt, E. Cruz, and M. Iyengar, “Challenges of Data Center Thermal Management”, IBM Journal of Research and Development, Vol. 49, No. 4/5, July/September, 2005.

J.G. Kooney, “Estimating Total Power Consumption By Servers In The US and The World”, Lawrence Berkeley National Laboratory, February, 2007.

W. Tschudi, “Best Practices Identified Through Benchmarking Data Centers”, presentation at ASHRAE Summer Conference, Quebec City, Canada, June, 2006.

M. J. Ellsworth, Jr., L.A. Campbell, R.E. Simons, M.K. Iyengar, R.R. Schmidt, and R.C. Chu, "The Evolution of Water Cooling for Large IBM Large Server Systems: Back to the Future", Proceedings of the 11th IEEE ITherm Conference, Orlando, Florida, USA, May, 2008.

H. Hamann, J. Lacey, M. O'Boyle, R. Schmidt, and M. Iyengar, "Rapid Three Dimensional Thermal Characterization of Large Scale Computing Facilities", IEEE Transactions of Components and Packaging Technologies, Vol. 31, No. 2, June, 2008.

R. Schmidt and M. Iyengar, "Information Technology Industry Energy Usage and Our Planet", Proceeding of the 11th IEEE ITherm Conference, Orlando, Florida, USA, May, 2008.

R. Schmidt and M. Iyengar, “Thermodynamics of Information Technology Data Center”, IBM Journal of Research and Development, Vol. 51, No. 3, 2009.

L.A. Barroso and U. Holzle, “The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines”, Synthesis Lectures on Computer Architecture, Google Inc., Morgan & Claypool Publishers (http://www.morganclaypool.com/doi/abs/10.2200/S00193ED1V01Y200905CAC006), 2009 .

L.T. Yeh and R.C. Chu, “Thermal Management of Microelectronic Equipment", ASME Press, New York, NY, 2002.

Dick Chu's Personal Process for Creativity

Within the IBM working environment, I would define creativity as one's intellectual output in the specific forms of new ideas or concepts which could lead to better, improved products. It should be a natural by-product of one's work. In terms of creative techniques, I have found the following simple rules useful:

- Always try to trace problems to their sources and try to understand their causes as fully as possible.
- Try to understand other people's work; both that of your peers and your competitors. By doing so, you'll learn from them, particularly their mistakes.
- Always anticipate future problems and requirements.
- Make every effort to understand both the *practical* and *intrinsic* limits of the technologies in your field. It's my experience that creativity is the art and science of stretching the practical limits, bringing them closer to their intrinsic ones.
- Be critical of yourself.
- You can never be content. This is the basic driving force of creativity.

Some Tips for Enhancing I/P Productivity and Quality

- Encourage collaborative innovations by forming small teams, meeting regularly with rotational leadership.
- Establish a recognition and reward system to offer incentives for I/P accomplishment. For example, a monetary award for each patent filed plus 3 invention points towards a plateau award consisting of a certificate and money for each 12 points accumulated.
- Confer the title of “Master Inventor” on those inventors who have reached a significant. level of patents in their patent portfolio and who have consistently demonstrated leadership in mentoring young inventors.
- Host a Quarterly I/P breakfast or luncheon to honor and celebrate innovation achievements.
- Publicize the invention award system to employees.
- Educate new employees on the invention process and the business importance of I/P.

Thank You

and

Good Luck