

Final Technical Report
Date of Report: December 13, 2012

Project Title: Recovery Act: Advanced Refrigerant-based Cooling Technologies for Information and Communications infrastructure (ARCTIC)

DOE Award Number: DE-EE0002895

Project Period: January 31, 2010 to August 30, 2012

Principal Investigator: Todd Salamon, Ph.D., Room 1E-436, 600 Mountain Avenue, Murray Hill, NJ 07974 (o) (908) 582-8766, (f) (908) 582-4868
todd.salamon@alcatel-lucent.com

Recipient Organization: Bell Laboratories, Alcatel-Lucent Technologies, 600 Mountain Ave, Murray Hill, NJ 07974

Partners: Modine (cost-sharing partner), US Hose (cost-sharing partner)

Business Contact: Darrell Kidd, 5440 Millstream Road, McLeansville, NC 27301 (o) (336) 697-3224, (f) (336) 697-3259
darrell.kidd@alcatel-lucent.com

DOE Project Officer: Dibyajyoti ('Debo') Aichbhaumik, 720-356-1423,
debo.aichbhaumik@go.doe.gov

DOE Project Monitor: Chad T. Sapp, 720-356-1302, chad.sapp@go.doe.gov

DOE HQ Contact: Bob Gemmer, 202-586-5885, bob.gemmer@ee.doe.gov

DOE Contract Specialist: Christina Kouch, 720-356-1674,
christina.kouch@go.doe.gov

Acknowledgment: "This report is based upon work supported by the U. S. Department of Energy under Award No. DE-EE0002895-".

Disclaimer:

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. Any findings, opinions, and conclusions or recommendations expressed in this report are those of the authors and do not necessarily reflect those of the United States Government or any agency thereof.

Document Availability: Reports are available free via the U.S. Department of Energy (DOE) Information Bridge Website: <http://www.osti.gov/bridge>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange (ETDE) representatives, and Informational Nuclear Information System (INIS) representatives from the following source:

Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831
Tel: (865) 576-8401
FAX: (865) 576-5728
E-mail: reports@osti.gov
Website: <http://www.osti.gov/contract.html>

Table of Contents

List of Figures	5
List of Tables	11
Chapter 1: Executive Summary	12
Chapter 2: Introduction	15
Chapter 3: Background	17
Chapter 4: Results and Discussion.....	28
Section 4.1: Heat Exchanger Optimization for Enhanced Shelf-Level Cooling (Task 1)	28
Section 4.2: Component-level air-cooled heat sink development (Task 2)	49
Section 4.3: Thermal backplane development (Task 3).....	65
Section 4.4: Device-level liquid cooling development (Task 4).....	75
Section 4.5: Rack-level test apparatus build and prove-in (Task 5)	90
Section 4.7: Develop and research new connector technologies (Task 7).....	93
Section 4.8: System-level testing (Task 8)	102
Section 4.10: Training (Task 10)	113
Section 4.11: Project Management and Reporting (Task 11).....	114
Chapter 5: Benefits Assessment.....	115
Section 5.1: Economic Benefits and Capital Cost Advantages	115
Section 5.2: Estimated energy savings across U.S. ICT facilities	116
Section 5.3: Estimated reductions in greenhouse gas emissions and other environmental emissions.....	117
Chapter 6: Commercialization	118
Section 6.1: Commercialization Strategy	118
Section 6.2: Applications of the technology: Customer Success Stories.....	118
Section 6.3: Industrialization	120

Chapter 7: Accomplishments	121
Chapter 8: Conclusions	125
Chapter 9: References	128

List of Figures

Figure 1: Advance refrigerant-based cooling for rack mounted equipment.	13
Figure 2: a) Conventional ICT facility cooled by a chilled-water CRAC unit; and b) Shelf-level implementation of the ALU modular liquid cooling system. Note that in contrast to Figure 1a, cool air is exhausted into the aisle, eliminating hot spots and the mixing of hot and cold air.	15
Figure 3: Photographs of 10-rack modular cooling system showing: a) overhead refrigerant piping and equipment racks outfitted with shelf-level heat exchanger / fan assemblies (left); and b) retractability of heat exchanger / fan assembly for equipment access (right).	24
Figure 4: Production “Overhang” Fin (Top) versus “Flush Face” fin design (Bottom).....	28
Figure 5: Pressure-enthalpy (PH) diagram for R134a.	31
Figure 6: HX Build 3 (CF, 2-row HX).	34
Figure 7: Finless HX build.....	34
Figure 8: Modine test setup schematic.....	35
Figure 9: Pumping Station and air duct behind calorimeter room (a), test HX and manifold piping (b), and XDV modules (c).	36
Figure 10: Chiller shown to the side of the pumping station.....	37
Figure 11: Thermal image of flush face HX during testing.....	37
Figure 12: Calibration of MPower model for baseline HX test data (heat transfer and airside pressure drop).....	40
Figure 13: Calibrated model results of test data and prediction for "Baseline" HX at 10°C incoming saturated refrigerant temperature.	41
Figure 14: Summary of calibrated performance at 40°C ambient and 10°C R134a.....	43
Figure 15: Summary of calibrated performance at 50°C ambient and 10°C R134a.....	44
Figure 16: Summary of calibrated airside pressure drop.....	45
Figure 17: Salt spray test.....	46
Figure 18: Option "1A" EPR25309 at 0 (left) and 96 hours (right).....	47
Figure 19: Option "1A" EPR25309 Collection from DI rinse. No metal flakes at 96 hours. Examination at 35x magnification.	47

Figure 20: Three-dimensional heat sink structures studied by Krishnan et al. (2012): a) Schwartz structure; b) finned foam; and c) slotted honeycomb. 49

Figure 21: Schematic of layering approach to making a slotted honeycomb heat sink, where a pattern of N layers of the orange stamping (WAFER-F) followed by M layers of the blue stamping (WAFER-B) are repeated a total of K times. Cylindrical rods passing through holes patterned in the heat sink base are used to register stampings..... 50

Figure 22: (bottom) Slotted honeycomb heat (corresponding to Design C in Figure 23) comprised of repeat patterns "F" and "B" consisting of multiple wafers within each pattern. Heat sink overall dimensions are 125 mm in width x 72 mm in length by 5.7 mm in height. Repeat pattern length is 6 mm for each "F" and "B" section and consists of nominally 23 or 24 wafers of thickness 0.010" each; (top) sections of wafers "F" and "B" making up each repeat pattern. 51

Figure 23: (left) Sample slotted honeycomb heat sink designs A, C and D depicting the two wafer layers "F" and "B" for each design; and (right) expanded view of designs..... 52

Figure 24: Linear fin heat sink 125 mm in width by 72 mm in length by 5.7 mm in height with a base thickness of 0.8 mm and having 50 fins, each with a height of 4.9 mm, width of 0.5 mm and a pitch of 2.5 mm that is CNC machined from aluminum with a conductivity of 160 W/mK..... 53

Figure 25: CAD drawings showing key components of the wind tunnel test setup. 54

Figure 26: (bottom) Wind tunnel setup for evaluation of air-cooled heat sink performance. Note that the Duocel metal foam homogenizer, depicted schematically in the top figure, is incorporated within the aluminum box located on the right-hand-side of the bottom figure; and (top) expanded view showing mass flow meter and direction of air flow..... 55

Figure 27: Computational geometry used for 3DHS Design C. Note the unit cell shown in the top and bottom figures, with the unit cell in the bottom figure being bounded by the lateral symmetry planes. 57

Figure 28: Comparison of laminar and transitional turbulent numerical model with experimental results for 3DHS Design C and LFHS. (top) junction-to-ambient thermal resistance; and (bottom) pressure drop..... 58

Figure 29: Experimentally measured junction-to-ambient thermal resistance and pressure drop as a function of average inlet flow velocity for 3DHS designs A, C and D and LFHS with 1.5 mm top flow bypass. Note that all 3DHS designs were made in aluminum, with design A also being made in copper (3DHS-A-Cu). 59

Figure 30: Comparison of 3DHS designs A, C and D relative to LFHS for 1.5 mm top flow bypass: (left) junction-to-ambient thermal resistance reduction as a function of pumping power;

and (right) relative reduction of pumping power as a function of junction-to-ambient thermal resistance..... 60

Figure 31: Experimentally measured junction-to-ambient thermal resistance and pressure drop as a function of average inlet flow velocity for 3DHS design C and LFHS with no flow bypass.... 61

Figure 32: Comparison of 3DHS design C relative to LFHS for no top flow bypass: (left) junction-to-ambient thermal resistance reduction as a function of pumping power; and (right) relative reduction of pumping power as a function of junction-to-ambient thermal resistance. .. 62

Figure 33: (a) Photograph of an open ALU shelf-level modular cooling unit seen from the rear of the rack. The surfaces of the 2-phase heat exchanger serve as cold plates in the proposed thermal backplane. (b) Schematic of a hot device on a circuit board in the rack shelf being cooled by the thermal backplane. A thermal module transfers the heat from the device to the thermal backplane. 65

Figure 34: Photograph of a modified finless heat exchanger, intended to serve as a thermal backplane. The finless heat exchanger was made by Modine. The 12 multiport heat exchangers run horizontally in this view. Thermal modules will be clamped to these surfaces for cooling of on-board devices. 66

Figure 35: Engineering drawings of a 250-mm prototype thermal module fabricated by Vendor 3. The module consists of three 8-mm-diameter heat pipes and copper heat-spreader blocks..... 67

Figure 36: Drawing of the thermal module test apparatus. The evaporator (left) end of the thermal module is clamped to the mock IC by a spring/load-cell assembly. The condenser end is clamped to a water-cooled cold plate..... 68

Figure 37: (a) Photograph of the clamping structure of Fig. 3.4 applied to the test of a Vendor 7 air-cooled heat sink module. (b) Photograph of tests of a 200-mm Vendor 3 heat-pipe thermal module. On the right, the condenser- section heat spreader is clamped to a rectangular water channel for cooling. 69

Figure 38: The mock-IC junction temperature is plotted as a function of the total length of the thermal module. 69

Figure 39: The mock-IC junction temperature is plotted as a function of heater power for the 200-mm Vendor 3 heatpipe module (blue points and line) and for the Vendor 7 air-cooled heat sink (red). The thermal resistance of the heatpipe-based thermal module is about 30% less than that of the air-cooled heat sink..... 70

Figure 40: Photographs of three T-shaped prototype thermal modules tested using the refrigerant-cooled thermal backplane. 71

Figure 41: Sketches of fixturing used to simultaneously test three thermal modules in the refrigerant-cooled thermal backplane.	71
Figure 42: Photograph of the thermal modules under test, using the finless heat exchanger of Figure 34 as a thermal backplane.	72
Figure 43: Photograph of two methods used to clamp the condenser sections of the thermal modules to the thermal backplane, illustrated using a heat-pipe module from Vendor 3 as shown in Figure 35. Left: the condenser and heat exchanger are sandwiched between two aluminum plates with screw holes and are held together by screws. Right: the left-hand clamping plate is replaced by a U-shaped fixture which is held in place by commercially-available split pins and springs.	73
Figure 44: The mock-IC junction temperatures obtained with all the thermal modules tested are plotted as functions of heater power.	74
Figure 45: Two prototype embedded-heatpipe heat-transfer modules fabricated by AVC. The gray material is aluminum, while the brown stripes are embedded copper heat pipes. Left: horizontal configuration; heat is transferred to the top edge. Right: vertical configuration; heat is transferred to the right edge.	76
Figure 46: Junction temperatures are plotted vs. total heater power, with equal voltages on all 5 mock ICs. Left: horizontal configuration. Right: vertical configuration.	77
Figure 47: Two views of a Vendor 5 gull-wing cold plate.	77
Figure 48: Left: a photo of the Vendor 6 micro-fin cold plate, with attached fluid connections and ports for refrigerant pressure measurements. Right: micrograph of the internal cooling fins.	78
Figure 49: The junction temperature of a single mock IC is plotted as a function of dissipated power for several flow rates in the high-flow-rate regime, for both the Vendor 5 and Vendor 6 cold plates.	79
Figure 50: The junction temperature of a single mock IC cooled by a Vendor 5 gull-wing cold plate is plotted as a function of dissipated power for two low flow rates.	80
Figure 51: The refrigerant temperature increase is plotted as a function of dissipated power for two low flow rates, for the experiments of the previous figure.	80
Figure 52: The junction temperature of a single mock IC cooled by a Vendor 6 micro-fin cold plate is plotted as a function of dissipated power for three low flow rates.	81
Figure 53: The refrigerant temperature increase is plotted as a function of dissipated power for three low flow rates, for the experiments of the previous figure.	81

Figure 54: The refrigerant pressure drop is plotted as a function of dissipated power for three low flow rates, for the experiments of the previous two figures.	82
Figure 55: Engineering drawing of a cooling module constructed with two aluminum Vendor 6 cold plates. The four T fittings are pressure taps.	82
Figure 56: The junction temperature of device 1 is plotted as a function of the power it dissipates, with device 2 dissipating 500 W.	83
Figure 57: The refrigerant temperature increase in CP1 is plotted as a function of the power dissipated at CP2, with CP1 power held at 500 W.	84
Figure 58: Stability diagram of the two-cold-plate system. The system is stable between the solid and dashed curves. Black points and curves: $f = 0.406(2)$ L/min. Blue points and curves: $f = 0.302(1)$ L/min.	85
Figure 59: Photos of the dual-loop cooling system. Left: front view of equipment rack showing blue coolant hoses. The coolant pump and reservoir are mounted on the back of the rack. Right: close-up of a Vendor 5 gull-wing cold plate mounted in one circuit pack.	86
Figure 60: Left: junction temperatures are plotted as functions of individual heater powers for a run in which the applied voltages were set to keep those junction temperatures approximately equal. Right: Blue junction temperature and input/output coolant temperatures are plotted vs. heater power for a run in which this device alone was cooled and powered. Full pump speed was used for both runs.	87
Figure 61: Laboratory setup at Alcatel-Lucent's Murray Hill, NJ facility.	90
Figure 62: (left) chilled water pumping station (gray box) connected to 18 kW heat load (black box); and (right) portable chilled water plant.	91
Figure 63: Still frame from a video of an underwater helium leak test used for evaluating permeability characteristics of polymer-lined hose assemblies.	94
Figure 64: Photograph of wind tunnel setup for evaluating performance of shelf-level heat exchangers and hose assemblies at ALU facility in Murray Hill, NJ: (left) inlet and side of test setup; and (right) rear of test setup.	102
Figure 65: (left) Wind tunnel setup used for evaluating performance of hose assemblies; and (right) test setup for evaluating ability of hose assembly to maintain a vacuum.	103
Figure 66: Test fixtures developed for evaluating the mechanical integrity of the non-conductive hose assembly: (left) random vibration test fixture; and (right) drop test fixture.	104
Figure 67: 17-rack 100 kW server farm (Server Farm 1) located in Plano, TX.	105

Figure 68: ALU facilities in Naperville, IL used for evaluation of layout of underfloor refrigerant piping: (left) Prototype Data Center; and (right) Adjacent Data Center. Note shelf-level modular cooling units attached to rectifier cabinets..... 106

Figure 69: Photographs of facilities used for system-level testing of shelf-level heat exchangers and hose assemblies, in ALU's Plano, TX facility: (left) Server Farm 1; and (right) Planned Lab Expansion..... 108

Figure 70: Alcatel-Lucent stage-gate criteria for commercial product development. 109

Figure 71: Business development activities in support of commercialization of liquid cooling technology..... 110

Figure 72: Business development activities performed in support of training. 113

Figure 73: Screenshot of start-up window for the Modular Cooling Solution Quote Tool..... 113

Figure 74: Comparison of cooling power requirements for a 1.5 MW data center cooled by chilled water CRAC units (left) and ALU modular liquid cooling system (right). 115

Figure 75: Alcatel-Lucent's Modular Cooling Solution installed in ALU's 17-rack, 100 kW Network Systems Integration and Test IPTV Laboratory in Plano, TX..... 119

List of Tables

Table 1: ALU requested modeling and testing inputs.	28
Table 2: Summary of HX models evaluated.....	29
Table 3: MPower original predicted performance.....	32
Table 4: Summary of options built and delivered to Alcatel-Lucent.....	33
Table 5: Actual test results with saturated refrigerant temperature deviation.	39
Table 6: MPower predictions at 10°C incoming saturated refrigerant temperature using calibrated model.....	42
Table 7: Heat exchangers subject to salt spray test.....	46
Table 8: Summary of results of thermal backplane experiments.....	74
Table 9: Summary of results of liquid cooling experiments.....	88
Table 10: Non-conductive polymer tubing material evaluated for the polymer-lined hose assembly.....	94
Table 11: Polymer braid materials evaluated for polymer-lined hose assembly.....	95
Table 12: List of evaluated potential replacement fittings and their associated features and cost.	100
Table 13: Tests performed on baseline technology in 100 kW prototype data center in ALU's Naperville, IL facility.....	105
Table 14: Tests performed for evaluation of system performance with underfloor refrigerant piping.	107
Table 15: Estimated annual U.S. emission reductions (column 3) for modular liquid cooling based on 9.6 billion kWh savings. Column 2 data scaling emissions with energy usage are calculated from the “ICT_Benefits_Spreadsheets.xls” originally located at http://sites.energetics.com/ICT_benefits09 that was provided by DOE during the proposal submission phase.	117

Chapter 1: Executive Summary

Faster, more powerful and dense computing hardware generates significant heat and imposes considerable data center cooling requirements. Traditional computer room air conditioning (CRAC) cooling methods are proving increasingly cost-ineffective and inefficient. Studies show that using the volume of room air as a heat exchange medium is wasteful and allows for substantial mixing of hot and cold air. Further, it limits cabinet/frame/rack density because it cannot effectively cool high heat density equipment that is spaced closely together.

A more cost-effective, efficient solution for maximizing heat transfer and enabling higher heat density equipment frames can be accomplished by utilizing properly positioned “phase change” or “two-phase” pumped refrigerant cooling methods. Pumping low pressure, oil-free phase changing refrigerant through microchannel heat exchangers can provide up to 90% less energy consumption for the primary cooling loop within the room. The primary benefits of such a solution include reduced energy requirements, optimized utilization of data center space, and lower OPEX and CAPEX.

Alcatel-Lucent recently developed a modular cooling technology based on a pumped two-phase refrigerant that removes heat directly at the shelf level of equipment racks. Figure 1 is a schematic illustrating the key elements that comprise the modular cooling technology. A pump delivers liquid refrigerant to finned microchannel heat exchangers mounted on the back of equipment racks. Fans drive air through the equipment shelf, where the air gains heat dissipated by the electronic components therein. Prior to exiting the rack, the heated air passes through the heat exchangers, where it is cooled back down to the temperature level of the air entering the frame by vaporization of the refrigerant, which is subsequently returned to a condenser where it is liquefied and recirculated by the pump. All the cooling air enters and leaves the shelves/racks at nominally the same temperature.

Results of a 100 kW prototype data center installation of the refrigerant-based modular cooling technology were dramatic in terms of energy efficiency and the ability to cool high-heat-density equipment. The prototype data center installation consisted of 10 racks each loaded with 10 kW of high-heat-density IT equipment with the racks arranged in a standard hot-aisle/cold-aisle configuration with standard cabinet spacing. A typical chilled-water CRAC unit would require approximately 16 kW to cool such a heat load. In contrast, the refrigerant-based modular cooling technology required only 2.3 kW of power for the refrigerant pump and shelf-level fans, a reduction of 85 percent. Differences in hot-aisle and cold-aisle temperature were also substantially reduced, mitigating many issues that arise in purely air-based cooling systems, such as mixing of hot and cold air streams, or from placing high-heat-density equipment in close proximity.

The technology is also such that it is able to retro-fit live equipment without service interruption, which is particularly important to the large installed ICT customer base, thereby providing a means of mitigating reliability and performance concerns during the installation, training and




Figure 1: Advance refrigerant-based cooling for rack mounted equipment.

validation phases of product integration. Moreover, the refrigerant used in our approach, R134a, is a widely-used, non-toxic dielectric liquid which, unlike water, is non-conducting and non-corrosive and will not damage electronics in the case of a leak—a triple-play win over alternative water-based liquid coolant technologies. Finally, through use of a pumped refrigerant, pressures are modest (~60 psi), and toxic lubricants and oils are not required, in contrast to compressorized refrigerant systems—another environmental win.

Project Activities

The ARCTIC project goal was to further develop and dramatically accelerate the commercialization of this game-changing, refrigerant-based, liquid-cooling technology and achieve a revolutionary increase in energy efficiency and carbon footprint reduction for our nation’s Information and Communications Technology (ICT) infrastructure. The specific objectives of the ARCTIC project focused in the following three areas:

- i) advanced research innovations that dramatically enhance the ability to deal with ever-increasing device heat densities and footprint reduction by bringing the liquid cooling much closer to the actual heat sources;
- ii) manufacturing optimization of key components; and
- iii) ensuring rapid market acceptance by reducing cost, thoroughly understanding system-level performance, and developing viable commercialization strategies.

The project involved participants with expertise in all aspects of commercialization, including research & development, manufacturing, sales & marketing and end users. The team was lead by Alcatel-Lucent, and included subcontractors Modine and USHose.

The Alcatel-Lucent effort included contributions from a number of internal organizations. ALU Services Business Division was responsible for overseeing the sales and marketing efforts and bringing the technology to the end user via a service product offering focused around retro-fits and consolidations of existing data centers. ALU Bell Labs Chief Technology Office (CTO) and ALU Network System Integration and Test (NSIT) Lab focused on system-level testing in

prototype and live data center environments to understand how new component designs impact typical cooling conditions, as well as how the system responds to cooling failures. ALU Bell Labs efforts focused on research and development of component-level cooling technologies for addressing continued increases in device heat densities, with emphasis on both air- and liquid-based technologies.

Modine Manufacturing Company and USHose played key roles in the research and development and manufacturing activities. Modine, headquartered in Racine, WI, is a worldwide leader in thermal management, and designs, engineers, tests, and manufactures heat transfer products for a wide range of applications and markets. Modine's efforts focused on the refrigerant-to-air heat exchangers that cool air leaving an equipment rack, in particular optimizing heat exchanger performance to accommodate much higher cooling capacities while minimizing required pumping station and fan power. U.S. Hose Corporation, headquartered in Romeoville, IL, manufactures a wide range of hose technologies including metallic flexible hose and braid, braided and rubber covered fluoropolymer hose and petro-chemical composite transfer hose and hose assemblies. U. S. Hose Corporation's efforts focused on optimizing the hoses and connectors in the system with respect to functionality, reliability and cost.

Commercialization

Commercialization activities focused on the shelf-level pumped refrigerant technology successfully demonstrated the market need for the technology, resulting in 6 customer installations in the U.S. The primary commercial benefits related both to reduced energy usage and increase equipment heat densities, which was a significant issue for several installations where the operator felt that a purely air-based cooling approach would be insufficient to meet requisite cooling requirements

During the project it was determined that Alcatel-Lucent would find a licensee to bring the modular cooling solution to market. Emerald Technologies, Inc. signed a License Agreement that was executed in December 2010. Emerald Technologies is a leading provider of energy efficient data center cooling solutions. The company has made commercial sales of the modular cooling solution to leading universities in the United States market. Information on the company and their product can be found at their website: <http://www.emeraldtechnologies.com/>.

Two-phase pumped refrigerant technology offers tremendous potential for improving the energy efficiency of cooling in data centers. The technology also allows better utilization of data center space, reduces acoustic noise levels, is able to retro-fit live equipment without service interruption and utilizes a widely-used, non-toxic dielectric liquid which, unlike alternative cooling approaches that use water as the working fluid, is non-conducting and non-corrosive and will not damage electronics in the case of a leak. Recent commercial implementations clearly demonstrate these benefits, and validate that there is a definite and compelling market need for this cooling technology.

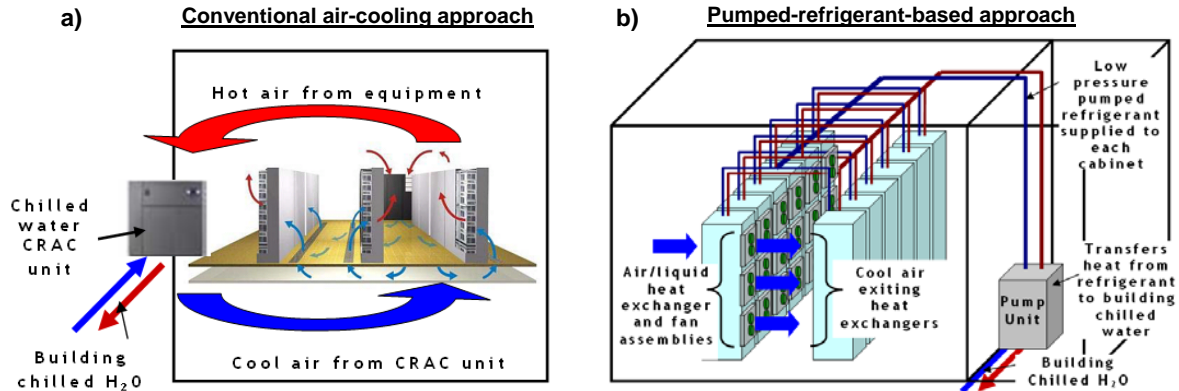


Figure 2: a) Conventional ICT facility cooled by a chilled-water CRAC unit; and b) Shelf-level implementation of the ALU modular liquid cooling system. Note that in contrast to Figure 2a, cool air is exhausted into the aisle, eliminating hot spots and the mixing of hot and cold air.

Chapter 2: Introduction

At present, ICT facilities are typically cooled by computer-room air-conditioning (CRAC) units, which provide cold air that is drawn into racks of heat-dissipating equipment by rack- and component-level fans, and with the heated exhaust air recirculated back to the CRAC units via room-level blowers and air movers (see Figure 2a). This approach has several drawbacks: i) the transport and cooling of large quantities of air is energy inefficient; ii) mixing of hot- and cold-air streams results in localized hot spots and regions of poor cooling and often necessitates additional sub-cooling of air to overcompensate for inefficiencies; and iii) the heat-carrying capacity of air is limited, so that increasing levels of heat dissipation are stretching air-cooling technologies to their limits, with acoustic noise restrictions and fan reliability concerns often responsible for placing practical limits on the achievable level of air cooling.

The technology innovation underpinning ARCTIC is the use of a pumped refrigerant to remove heat at the rack, shelf and component levels. The use of a pumped refrigerant has a number of advantages over traditional air- and liquid-based cooling. First, the heat-carrying capacity relative to pumping power is substantially greater for the refrigerant than air due to the refrigerant's large heat of vaporization. This results in projected energy savings in excess of 90% for the pumped refrigerant-based technology relative to conventional CRAC-based air cooling. Second, the pumped refrigerant allows significant increases in device densities at the shelf, circuit-pack and component levels not readily achievable via air-based approaches. Third, the R134a refrigerant is a dielectric liquid which, unlike water, will not damage electronics in the case of a leak. R134a is also non-toxic, and, since the system is based on a pumped refrigerant, it does not require toxic lubricants and oils required by refrigerant systems. Fourth, system pressures are relatively modest (~60 psi) in comparison to a compressorized refrigerant system. Finally, the technology developments in ARCTIC address important issues of reliability and cost that are critical for market acceptance.

A schematic of the ALU modular liquid cooling system is shown in Figure 2b. Liquid refrigerant R134a is pumped to finned microchannel heat exchangers placed at the air exhaust in

shelves of individual equipment racks. The heat exchangers are similar in form to automobile radiators and consist of alternating stacks of flat, microchannel tubes which carry refrigerant, and strips of fins through which the air passes and is cooled. Fan trays push air through the equipment shelf, where it gains heat dissipated by the electronic components therein. Prior to exiting the rack, the heated air passes through the heat exchangers, where it is cooled back down to the temperature level of the air entering the frame. The vaporized refrigerant is then returned to a condenser where it is liquefied and recirculated by the pump. All the cooling air enters and leaves the shelves/racks at nominally the same temperature.

The objectives of the ARCTIC proposal are three-fold: i) advanced research innovations (novel heat-sink structures, thermal backplane, device-level liquid cooling) that dramatically enhance the ability to deal with ever-increasing device heat densities and footprint reduction by bringing the liquid cooling much closer to the actual heat sources; ii) manufacturing optimization of key components (heat exchangers, pumps, connectors, hoses); and iii) ensuring rapid market acceptance by reducing cost, thoroughly understanding system-level performance and developing viable next-gen commercialization strategies. ARCTIC will deliver substantial benefits in three key areas important for promoting market acceptance:

i) reduced cooling costs: ARCTIC's projected energy savings-- hence cooling cost savings -- are in excess of 90% relative to conventional CRAC-based air cooling, since the heat-carrying capacity relative to pumping power is substantially greater for the refrigerant than for air.

ii) reduced real-estate footprint: As compared to air-cooled approaches, our pumped-refrigerant approach can handle significant increases in rack, shelf, circuit-pack and component level heat loads, allowing the much higher equipment densities required by burgeoning ICT facility needs.

iii) enhanced reliability: ARCTIC's ability to retro-fit live equipment without service interruption is particularly important to our nation's large installed ICT customer base, providing a means of mitigating reliability and performance concerns during the installation, training and validation phases of product integration. Moreover, the refrigerant used in our approach, R134a, is a widely-used, non-toxic dielectric liquid which, unlike water, is non-conducting and non-corrosive and will not damage electronics in the case of a leak—a triple-play win over alternative water-based liquid coolant technologies. Finally, through use of a pumped refrigerant, pressures are modest (~60 psi), and toxic lubricants and oils are not required, in contrast to compressorized refrigerant systems—another environmental win.

Meeting the proposal objectives and achieving full market penetration will allow an estimated reduction in U.S. greenhouse gas emissions of 1.46 MMTCE/year and will reduce data center and ICT facility energy costs by 9.6 billion kWh/year. Higher equipment densities will result in real estate utilization improvements valued at \$1.6 billion annually.

Chapter 3: Background

Section 3.1: Current state-of-the-art in liquid-cooling technology

In the following we include details on the current state-of-the-art in liquid cooling of electronics. This includes information that was current at the time of proposal submission for the project in July of 2009, as well as recent developments.

Several companies produce liquid-cooling-based technologies to address issues associated with air cooling. Liebert, APC, HP, IBM and Coolcentric/Vette have independently introduced systems wherein chilled water is pumped directly to heat exchangers near or within the equipment racks. This approach can lower energy costs for cooling and support increased device densities. However, water-based cooling systems have a tremendous disadvantage: a leak can damage or destroy expensive electronic equipment. The downtime associated with spill remediation may also result in service interruption. ARCTIC, by contrast, uses a non-conducting, non-toxic and non-corroding refrigerant.

Thermal Form and Function (TFF)¹, in partnership with Parker-Hannifin's² Precision Cooling Business Unit, proposes refrigerant-based solutions for electronics cooling. TFF offers refrigerant pumps with capacities ranging from 2.5 kW up to 100 kW along with so-called meso-scale heat exchangers for cooling of high-heat-density computer chips. Although TFF's approach has similarities to ours, it is not clear that they adequately address critical issues of system-level performance, reliability and ease-of-use. For example, it is unclear how TFF's cold-plate solution accomplishes thermal attachment (low-thermal-resistance interfaces to circuit-pack components) and fluidic coupling (low-leakage, high-reliability connectors) in ways that allow hot-swapping of circuit packs. ARCTIC proposes several innovations to address these key issues. TFF's transport of vaporized refrigerant for condensation may also be extremely inefficient if the transport distance is large – company literature mentions distances up to hundreds of meters – which may result in poor pump utilization, as well as undesirably high evaporation temperatures, which will decrease cooling efficiency. Finally, it is also unclear how TFF addresses humidity control and condensation, which can result in decreased reliability and component failure.

Liebert Corp.³ offers pumped-refrigerant-based cooling technology in their XD equipment series. Note that APC⁴ has a pumped-refrigerant based product portfolio which has similar overall features and capabilities to the Liebert XD equipment series, so that the following comments also apply to APC's technology. Liebert's pumped-refrigerant technology is an improvement over CRAC-based systems as it brings the cooling coils closer to the heat-generating equipment. Liebert's technology thus helps to minimize localized hot spots and to decrease the energy required to move large volumes of air over long distances. However, the energy, cost and device-density benefits of Liebert's technology are substantially less than those proposed in ARCTIC, for the following reasons: i) the energy efficiency of ARCTIC's cooling system is substantially greater than Liebert's due to novel heat-exchanger design and shelf integration; ii) Liebert's XD-series cooling modules are energy-inefficient, as they require significant fan power and high room air temperatures to achieve maximum efficiency; and iii) Liebert uses air as the primary medium for intermediate and long-distance heat transfer. Inefficiencies in air-side heat transfer and fan performance thus greatly limit the range of applicability of Liebert's technology for high-heat-density systems.

Liebert recently announced the addition of the XD Coolframe and the XDR rear-door cooling units to their XD equipment series. The Liebert XD Coolframe consists of two refrigerant-cooled heat exchangers that mount to the back of Egenera's Bladeframe rack. Although this unit has a number of similarities to ARCTIC's modular shelf-level heat exchanger, the XD Coolframe requires that data center operators and/or equipment vendors adopt the Egenera Bladeframe stateless server and high-speed fabric technology, which may be too restrictive from a functionality and use perspective. This requirement also appears to preclude application of the XD Coolframe technology to the large installed base of data center equipment racks, which is not an issue for the ARCTIC modular shelf-level heat exchanger. The Liebert XDR rear-door unit appears to have more flexibility in accommodating existing equipment racks. However, its lack of supplemental fans to assist in air flow will require that server fans be run harder to overcome the additional pressure drop associated with the rear-door heat exchanger. This is a concern, as the server fan size is small, and may require significant increases in fan RPM, which will increase fan power usage and acoustic noise. The supplemental cooling fans used in the ARCTIC approach are much larger in diameter, and hence have energy efficiency and acoustic noise benefits relative to smaller diameter server fans.

Clustered Systems⁵ has developed a liquid-cooled server rack that removes heat directly from servers via direct attachment to cold plates incorporated adjacent to the server blades within the rack. The technology is offered by Liebert within their XD family of cooling solutions as the Liebert XDS server cooling rack. Fans and heat sinks are replaced by a set of custom designed "heat risers", which transport heat to the server lid, which includes a super compliant thermal interface material between the server lid and cold plate. Although the configuration offers tremendous potential for energy savings as it allows the elimination of server fans, it is unclear what implications removal of equipment fans has on other components that do not have a direct thermal attachment to the cold plate via the heat risers; the ARCTIC project proposes a hybrid air- and liquid-based cooling solution, whereby the highest heat density components are cooled by liquid and low heat density components are cooled by air. Furthermore, it is unclear how efficient the heat risers are at transporting heat from the component case to the cold plate. This is in addition to the multiple thermal interfaces that are used in the approach and which can degrade thermal performance.

Green Revolution Cooling (GRC)⁶ has developed a system based on immersion cooling, in which the electronics are submerged in a dielectric liquid, such as white mineral oil. This technique is similar to what is used in cooling high-voltage transformers. The benefits of this approach include substantially reduced energy costs due to the much greater heat capacity of the liquid relative to air as well as the removal of server and power supply fans and the ability to cool much larger rack-level heat densities (GRC claims up to 100 kW in a 42U rack). However, there are still a number of special considerations and concerns regarding the approach, including: i) company literature says the cooling fluid dissolves thermal grease, and a metal-based (Indium) foil is needed for the thermal interface material between components and heat sinks. Solubility of the circuit board and associated components would also need to be fairly well understood, as the fluid might leach other chemical moieties; ii) hard drives need to be encapsulated to properly function; iii) it is not clear how the liquid could affect the performance and reliability of optical connections, which are increasingly being used for short to medium range data transport within and between data center racks; iv) the cooling liquid dramatically increases the weight of the equipment rack, with GRC's custom 42U rack weighing 200 kg without fluid and 1200 kg when

filled with liquid, a difference of 1000 kg or 1.1 tons. This is a significant amount of weight, which could have structural implications, particularly for applications in telecom, which needs to pass NEBS Earthquake and Office Vibration testing; v) the cooling fluid (white mineral oil), has a price of approximately \$1 to \$2 per lb, which can add significant cost to the overall system, especially if the fluid needs to be replaced or replenished; and vi) the cooling fluid has a high auto-ignition temperature, but one would also want to know the fire point temperature, e.g., the temperature at which the fluid can sustain a flame - 1000 kg of a potentially flammable fluid in a rack is a lot of fuel and could create a massive fire. There appears to be a push in the transformer space to increase the fire point of the fluids above 300°C (from 160°C) because of the flammability concern.

Google⁷ appears to have one of the best records for power efficiency amongst all data centers operators, reporting average PUE numbers in the range of 1.13 across all of their data centers and in all seasons. One reason Google achieves such a low PUE is through well-optimized thermal management, including: i) raising the cold-aisle air temperature to 80°F; ii) utilize thermal modeling to plan the data center layout to avoid hot spots and strategically place CRAC units in such a way as to minimize hot spots and thereby reduce the overall time that CRAC units operate; iii) manage air flow by segregating the hot aisle and cold aisle air streams, thereby reducing inefficiencies due to mixing of hot and cold air. This is accomplished by using either ducting or permanent enclosures for their large data centers, and similar measures for smaller data centers; iv) use of fresh air, evaporation and sea water for cooling; and v) the use of so-called “Hot Huts”, which trap the hot air leaving server racks and prevent it from mixing with cold air within the data center. Fans located on the top of each “Hot Hut” pull the hot air through water-cooled coils. The cooled air is then returned to the ambient air within the data center, where it is available to be pulled into servers.

The Google “Hot Hut” has a number of similarities to the water-cooled solutions described by Liebert, APC, etc. above, in that hot air exiting a server rack is cooled down prior to re-entering the data center ambient environment. The approach has the added benefit that, by sealing the hot air within the “Hot Hut”, significant inefficiencies due to hot and cold air mixing, are eliminated, which is not true of the aforementioned approaches. However, the approach has a number of drawbacks, including: i) it relies on placing water in close proximity to the electronics, and a water leak could damage or destroy expensive electronic equipment, as well as result in data loss; ii) it is not clear that the use of a purely air-based approach at the server level allows scalability to higher server heat densities; and iii) the PUE metric that Google uses includes the server fan power as part of the IT equipment power, and hence will not accurately reflect the overhead of server fan energy usage, which will be expected to increase substantially with increased server heat density.

Section 3.2: Project Goals and Objectives

The goals and objectives of this project are to dramatically accelerate the commercialization of a game-changing, refrigerant-based, liquid-cooling technology and achieve a revolutionary increase in energy efficiency and carbon footprint reduction for our nation’s Information and Communications Technology (ICT) infrastructure. The specific objectives of the ARCTIC project are three-fold: i) advanced research innovations that dramatically enhance the ability to deal with ever-increasing device heat densities and footprint reduction by bringing the liquid cooling much closer to the actual heat sources; ii) manufacturing optimization of key

components; and iii) ensuring rapid market acceptance by reducing cost, thoroughly understanding system-level performance, and developing viable commercialization strategies. The ARCTIC project will deliver substantial benefits and capabilities in three key areas important for promoting market acceptance, namely, reduced cooling costs, reduced equipment footprint, and increased reliability. These activities, broken down into specific tasks along with details on the task owner, goals, objective and planned approach, are presented below.

Task 1: Heat exchanger optimization for enhanced shelf-level cooling

Task owner: Modine Manufacturing Company

This task enhances the performance of the liquid cooling system by optimizing the shelf-level heat exchangers. The performance of the heat exchanger will be optimized with the goal of accommodating much higher cooling capacities while minimizing required pumping-station and fan power. A concurrent goal is to minimize the flow resistance of the heat exchangers for the current heat removal capacity so that integrated fans are not required.

This objective will be addressed by using computational fluid dynamics (CFD) software to devise up to four new heat-exchanger designs. Ten samples of each will be produced for testing, and will be required to match numerical predictions to within 10% for heat exchanger coefficient of performance (COP), defined as $COP = (\text{cooling capacity} / \text{power for cooling})$ and a COP that is at least 15% larger than current designs. Up to 150 heat exchangers will then be produced for large-scale system testing in Task 8.0.

Milestone: Build of up to 150 heat exchangers for system-level testing having COP (= cooling capacity / power for cooling) that is at least 15% larger than current designs.

Task 2: Component-level air-cooled heat sink development

Task owner: Alcatel-Lucent Bell Labs

The goal of this task is to improve component-level air cooling using novel three-dimensional heat sinks (3DHS). Analytical and numerical methods will be used to investigate and optimize new 3DHS designs. The anticipated result will be a computational toolkit for 3DHS design and a number of application-specific prototypes. We expect to be able to design, fabricate, and verify that at least one 3DHS design can achieve a thermal resistance at least 20% lower than an optimized finned heat sink of the same volume. Low-cost manufacturing techniques will also be investigated.

Task 3: Thermal backplane development

Task owner: Alcatel-Lucent Bell Labs

Develop refrigerant-cooled surfaces to provide a high-conductance direct thermal connection between the powered devices and the rack-level heat exchangers. The primary objective is to design and test a conductive structure that couples a powered device in an exemplary circuit pack to a mating structure in a shelf-level heat exchanger.

The goal of this task is thus to develop a structure that achieves a 30% reduction in thermal resistance and/or a 30% reduction in form factor relative to an optimized air-cooled heat sink.

Incorporation of lower cost technologies and the development of a mechanically-reliable, re-attachable direct thermal connection are concurrent goals. Rules which allow the designer to choose the heat-transfer structure appropriate for a given circuit pack will also be established.

Milestone: Show ability to achieve a 30% reduction in thermal resistance and/or a 30% reduction in form factor of the thermal backplane relative to an optimized air-cooled heat sink.

Task 4: Device-level liquid cooling development

Task owner: Alcatel-Lucent Bell Labs

For devices which are too hot to be cooled even by the thermal backplane approach, direct cooling of individual components using device-level liquid-cooled heat sinks will be investigated. The task will begin by assessing available component-level, liquid-cooling options and evaluating the trade-offs between performance, capacity, cost, reliability, and ease of installation. Next, the focus will be on designing and building a board-level test apparatus to evaluate the performance of component-level heat exchangers. Finally, the test apparatus will be integrated into an equipment frame fitted with the rack-level liquid cooling solution to evaluate system-level implications (see Task 5.0). The end goal is to understand the requirements for implementation of device-level liquid cooling and to evaluate the commercial feasibility of this approach.

Task 5: Rack-level test apparatus build and prove-in

Task owner: Alcatel-Lucent Bell Labs

An experimental apparatus will be built to enable testing and evaluation of technologies developed for improved cooling at the device, board, shelf, and rack-level. Such a test apparatus is critically important for evaluating the ability of these cooling technologies to be readily incorporated into commercial equipment, as well as to understand potential impacts on overall system-level performance.

Task 6: Optimization of refrigerant-circulation system

Task owner: Modine Manufacturing Company

The pumping station used for refrigerant circulation accounts for a significant fraction of the cooling system's overall cost and energy usage. The goal of this task is to design and build the hardware and control infrastructure to provide a more energy-efficient, reliable, and cost-effective refrigerant-circulation system. Guidelines will also be developed for supporting a range of power levels, so as to match cooling capacity with end user heat load, and interfacing with a compressorized secondary cooling loop, to accommodate sites without building chilled water.

Milestone: Manufacture of up to 2 pumping stations for system-level testing based on developed design guidelines for performance (energy efficiency, cost, reliability) and features (cooling capacity matched to end-user heat load, accommodation of a compressorized secondary cooling loop) that are as good as or better than the existing commercial pumping station.

Task 7: Develop and research new connector technologies

Task owner: USHose Corporation

The hoses and connectors used in our prototype system have a high cost, and will be optimized to balance functionality, reliability, and cost. Specific activities include:

Subtask 7.1: Optimization of Hose Assemblies

Optimize the cost and performance of the hose assemblies using existing materials and manufacturing processes.

Subtask 7.2: Research and Develop New Hose Materials

Research new polymer-lined hose assemblies as a replacement to the current metal hose assemblies. The polymer hose manufacturing process is less labor intensive than the metal hose assembly process due to its use of a crimped, as opposed to welded, joint design, and hence offers opportunities for cost reduction. An important design criterion will be making polymer hose assemblies that are impermeable to refrigerant.

Subtask 7.3: Source New or Similar Wrench-Less Fitting for Metal Hose Assembly

Reduce cost and improve fitting availability and lead time by developing multiple suppliers.

Subtask 7.4: Develop and Source New End Fittings for Polymer-Lined Hose Assembly

Design, validate and source new end fittings for the polymer-lined hose assembly, including a wrench-type design that attaches to the manifold portion of the system, and a wrenchless-type design that attaches to the shelf-level heat exchanger.

Milestone: Optimization of performance and cost of hose assembly and end fittings using new and existing materials and manufacturing processes.

Task 8.0: System-level testing

Task owner: ALU CTO Lab, ALU NSIT Lab, ALU Bell Labs and ALU Services BD

System testing will be used to understand how new component designs impact two key areas: 1) system performance under typical cooling operation; and 2) how the system responds to cooling failures. Data collected will include: temperature measurements internal to servers, after the cooling coils, and at various room locations; equipment heat loads; cooling system energy usage (pumping station and fans); room humidity level; building chilled water characteristics; refrigerant supply and return temperature and quality; and acoustic noise characteristics. Failure modes to be studied will include: loss of refrigerant cooling loops at different equipment scales (shelf, rack, row of racks, every other shelf, every other row of racks, etc.); pump failure; power failure; building chilled water flow reduction; and room humidity level changes.

Testing will be performed in the following facilities:

- A Prototype Data Center (PDC) consisting of 10 equipment frames hosting a range of servers and 100 kW total power. The PDC will also enable testing of several important configurations with respect to distributed heat loads and refrigeration distribution system

layout by providing cooling to equipment frames having 30 kW of power in an Adjacent Data Center (ADC).

- A server farm (Server Farm 1) consisting of 17 frames with approximately 400 servers of various vendor makes and models and 100 kW total power.
- A Planned Lab Expansion (PLE) consisting of between 60 and 70 equipment frames.
- An Environmental Chamber/Test Rack (ECTR) facility that will allow rigorous testing of new pumping stations under precisely controlled temperature and humidity conditions not readily achievable in a data center or lab environment.

Note that since Server Farm 1 and the Planned Lab Expansion are running live traffic for Tier 1 and 2 carriers, some of the more disruptive failure tests will be performed in the Prototype Data Center and Environmental Chamber/Test Rack facility.

Subtask 8.1: Testing of new component designs prior to full system-level evaluation

Perform evaluation of new heat exchangers, connectors and refrigerant pumping stations in the Prototype Data Center and the Environmental Chamber/Test Rack facility. Consultation and feedback on new component design choices will also be provided.

Subtask 8.2: Baseline system-level testing of current cooling technology

Perform baseline testing of the current generation cooling technology that is already installed in Server Farm 1 and the Prototype Data Center.

Subtask 8.3: System-level testing with new refrigerant distribution system layout

Evaluate the current generation cooling technology (pumping station, heat exchangers, connectors) with the new refrigerant distribution system in the Prototype Data Center and Adjacent Data Center and perform system-level testing so as to identify and remediate problems caused by the change in layout. Data will be compared to performance of the current baseline system layout in the Prototype Data Center.

Subtask 8.4: System-level testing of new components

Perform system level evaluation of new designs for heat exchangers and connectors. Testing will be performed in Server Farm 1 and/or the Planned Lab Expansion and compared to current cooling technology.

Milestone: Qualification and system-level testing of current and new cooling technology in the Prototype Data Center, Adjacent Data Center, Environmental Chamber/Test Rack facility, Server Farm 1 and the Planned Lab Expansion.

Task 9.0: Business development

Task owner: ALU Services Business Division

The goal of this task is to develop a successful product offering for the liquid cooling system. This will include incorporating the multiple and synergistic enhancements of component, shelf-

and system-level cooling being developed as part of this project into a viable commercialization strategy so as to accelerate introduction into next generation products and service offerings. Efforts will include: i) product cost analysis; ii) identification and establishment of vendors; iii) development of customer business cases; iv) establishment of pricing; v) establishment of total addressable market (TAM) and market share; vi) development of a sales plan and determination of cost; vii) development of Alcatel-Lucent business case; and viii) management review.

Milestone: Develop a promising business case based on the following metrics: direct margin, contribution margin, payback period (internal), market share, and overall projected revenue.

Task 10.0: Training

Task owner: ALU Services Business Division

Develop training procedures for installation, operation and maintenance of liquid-cooling technology. This will include creating documentation and user manuals for installation, operation and maintenance of equipment, as well as training of individuals.

Task 11.0: Project Management and Reporting

Task owner: ALU Bell Labs Business Development

Reports and other deliverables will be provided in accordance with the Federal Assistance Reporting Checklist (DOE F 4600.2) following the instructions included therein. These deliverables address Federal and American Recovery Reinvestment Act (ARRA) requirements, as well as ITP programmatic requirements including but not limited to program status reports and financial status reports.

Section 3.3: Past Experience

Alcatel-Lucent implemented a prototype of the ALU modular liquid cooling system in a test data center (see Figure 3) housing 10 racks, each equipped with 10 kW of IT equipment, for a total power of 100 kW. The racks are arranged in a standard hot-aisle/cold-aisle configuration with standard cabinet spacing. Modine heat exchangers are installed on the hot-air exhaust side of each shelf of individual racks, with a fan tray placed on the heat exchanger exterior to pull hot air through (see Figure 3a, b). A 160-kW-capacity Liebert XDP pump feeds refrigerant to overhead copper plenum piping, while vapor and liquid are returned to the chilled-water condenser by copper return plenum piping. Flexible hoses connect the heat exchangers to the supply and return lines. The lab is completely isolated from the building HVAC system, so that all



Figure 3: Photographs of 10-rack modular cooling system showing: a) overhead refrigerant piping and equipment racks outfitted with shelf-level heat exchanger / fan assemblies (left); and b) retractability of heat exchanger / fan assembly for equipment access (right).

generated heat is removed exclusively by the ALU modular liquid cooling system.

The performance of the prototype clearly demonstrates the feasibility of this approach. The power required to operate the system at full capacity consists of the pump (800 W) and shelf-level fans (1500 W = 2 fans per shelf x 30 shelves x 25 W/fan) for a total of 2.3 kW. For comparison, a comparably-sized chilled-water CRAC unit would require 16.0 kW. This results in coefficients of performance (COP), defined as $COP = (\text{cooling capacity} / \text{power for cooling})$, of **43.5** for the liquid cooling prototype and **6.3** for the chilled-water CRAC system. A second test installation, a “hot” retro-fit on live equipment in a 17-rack, 100-kW ALU IPTV center, shows similar energy efficiency improvements.

Section 3.4: Qualifications

Alcatel-Lucent / Bell Laboratories ALU is the world’s pre-eminent communications solutions company. Bell Laboratories plays a number of different, yet equally important roles for ALU including: serving as part of Alcatel-Lucent’s innovation foundation; contributing expertise to current and future products and services; impacting national technical priorities; and exerting global influence on science and technology. Bell Labs' innovations cover physical, computer, software and mathematical sciences; wireless and wireline networking; and network security, standards, and planning. Bell Labs researchers have received six Nobel Prizes in Physics, nine U.S. National Medals of Science and seven U.S. National Medals of Technology®. Our scientists and engineers have played a pivotal role in inventing or perfecting most of the key communications technologies in use today. With more than 25,000 patents, Bell Labs' research centers have experience in all aspects of network design and support and play a central role in responding to requests for information and requests for proposals (RFI/RFP’s) from the world’s largest service providers. Bell Labs is recognized as a global leader in devising, testing, developing and deploying state-of-the-art hardware and software systems for carrier-class communication networks. This requires expertise ranging from physical technology for high-capacity networks to scientific computing for simulating system performance to combinatorial optimization for high-level design.

Our Murray Hill, NJ laboratory is equipped with a 30-node Linux cluster consisting of 2- and 4-processor CPUs for intensive computing. We have licenses for a number of commercial software packages for CAD/CAE, including Fluent, CFX, ANSYS and Icepak. We possess a laboratory equipped for making thermal measurements and a large, well-equipped machine shop for prototype production. We also have access to heated and unheated wind tunnels for characterizing 3D heat sinks, 3D printers and an established fabrication process for making prototype 3D heat sinks, as well as a state-of-the-art apparatus for characterizing thermal interface materials at our Blanchardstown, Ireland laboratory.

ALU Services / Services Business Division With more than 18,000 network experts supporting the top 30 service providers in over 130 countries, ALU Services is the industry's most experienced and knowledgeable services partner, providing a broad and comprehensive set of services that encompass the entire network lifecycle, including consultation & design, integration & deployment, and maintenance & operation. ALU Services designs, builds, implements and manages facilities to help meet the demands of business and residential customers, saving costs and reducing time to revenue.

ALU Bell Labs CTO Lab The ALU Bell Labs Chief Technology Office (CTO) facility includes a 10-rack version of the modular liquid-cooling system hosting a range of operational servers with 100-kW total power and with a 160-kW capacity pumping facility. Extensive testing of the builds of optimized microchannel heat exchangers, pumping stations and connector and hose assemblies will be carried out at this facility.

ALU NSIT Lab The ALU Network Systems and Integration Test (NSIT) lab supports the testing of the ALU Internet Protocol Television (IPTV) solution, as well as new network and IPTV hardware and software releases for all current and future ALU customers. The lab maintains a test environment which features a combination of under-floor and ceiling drop-down cooling. This area and its configuration of server cabinets closely emulate, and in several cases exceed, the cabinet capacity of deployed IPTV server solutions for ALU customers in the U.S. and Canada. The ALU NSIT lab has facilities for rigorous system-level testing of enhancements to the modular liquid cooling system. The baseline first generation shelf-level modular cooling technology is currently installed in an existing NSIT server farm which houses 17 frames equipped with ~400 servers of various vendor makes and models and total power dissipation of ~100 kW. A lab expansion incorporating the modular liquid cooling system is ongoing and will provide another representative system for quantifying performance improvements due to the modular liquid cooling system.

Modine Manufacturing Co. Modine specializes in thermal management, bringing heating and cooling technology to diversified markets. Modine's products are used in light-, medium-, and heavy-duty vehicles, HVAC equipment, industrial equipment, refrigeration systems, fuel cells and electronics. In FY 2008, Modine's sales reached \$1.8 billion dollars. Modine is known throughout the automotive and non-automotive sectors as an innovative company with more than 1,000 patents awarded worldwide since its founding in 1916. Key technologies invented by Modine that are now used worldwide in heat exchanger design include the serpentine louvered fin and the microchannel heat exchanger, both helping the US to become more energy efficient and globally competitive.

Modine's facilities include a thermal testing laboratory that has extensive capabilities for evaluating the microchannel heat exchangers: single-phase calorimetry (up to 300kW of cooling), two-phase refrigeration calorimetry (up to 12kW of cooling) and isothermal pressure drop testing (flow rates up to 3240 m³/hr). Our structural laboratory provides full mechanical test capabilities: multi-, single- and high-frequency vibration testing, drop testing, burst testing (up to 55,200 kPa), and pressure (7-17,240 kPa) and thermal cycling (24-110°C). We also have an extensive suite of virtual design tools that will be applied to the proposed research, including CAD and thermal design software, statistical analysis software, and multiple tools for computational fluid dynamics simulation and finite-element analysis. We have also developed two proprietary thermal prediction codes.

US Hose Corp. US Hose has over 100 years of experience in the manufacture of corrugated metal, fluoropolymer, and composite hose and hose assemblies for a wide array of markets. US Hose has approximately 100,000 sq. ft. of fabrication, warehouse, and office space. The total number of employees is approximately 95. US Hose has three product lines: corrugated metal hose, PTFE hose, and composite hose in diameters ranging for 1/8" to 12" and beyond.

US Hose has extensive support for hose fabrication and welding: TIG/MIG welding, orbital welding, 5 welders certified to ASME Section IX BPVC. Test facilities include helium and nitrogen leak testing to 5,000 psig, mass-spectrometer helium leak testing to 1.0×10^{-10} sccs, water-free hydrogen leak testing to 400 psig, hydrostatic testing to 20,000 psig, and partnerships with NDT sources for additional testing requirements. Other capabilities include CGA G4.1 approved oxygen delivery cleaning and CNC machining. US Hose is ISO 9000 compliant and has earned a wide array of industrial approvals, including nuclear, aerospace, industrial, and defense specifications.

Chapter 4: Results and Discussion

Section 4.1: Heat Exchanger Optimization for Enhanced Shelf-Level Cooling (Task 1)

The purpose of Task 1 was to enhance the performance of the existing liquid cooling system by optimizing the shelf-level heat exchangers. Optimizing the shelf level heat exchanger (HX) leads to a more energy efficient data center cooling solution. The goal of Task 1 was to reduce the input power requirements by reducing the airside resistance of the HX while maintaining or increasing the heat transfer performance. Additionally, sample heat exchanger designs (without fins) were made and supplied to support other approaches to providing liquid cooling of components being investigated by Alcatel-Lucent as part of the project.

Subsection 4.1.1: HX modeling and Design

The depth of the fin on current parallel-flow (PF) production HX's for Alcatel-Lucent overhangs the tube on each side, as shown at the top of Figure 4. It is known to add excess airside pressure drop and material with minimal performance benefit. This project investigated the performance of a “flush face” fin that is the same depth as the tube. In conjunction with the fin improvements, a change in tube design was also investigated. This new tube has enhanced internal geometry that increases the internal surface area and increases the free flow area of the tubes. The combination of the increased surface area and free flow area yield improved heat transfer with lower refrigerant side pressure drop.




Figure 4: Production “Overhang” Fin (Top) versus “Flush Face” fin design (Bottom).

Modeling conditions were defined by Alcatel-Lucent as shown in Table 1. Inlet saturated refrigerant temperature was set to 10° C for all conditions with air inlet dry bulb temperatures at 40 and 50° C and air velocities between 1.1 and 3.0 m/s.

Alcatel-Lucent Model Conditions (Inlet Saturated Refrigerant Temp = 10°C)									
AIR VELOCITY	m/s	1.1		2.0		2.5		3.0	
INLET AIR TEMPERATURE	°C	40	50	40	50	40	50	40	50
REF MASS FLOW	FLOATS TO MATCH INLET TEMPS AND AIRFLOW								
<i>Heat exchanger will have nearly identical overall dimensions and mechanical connections</i>									

Table 1: ALU requested modeling and testing inputs.

In an attempt to further reduce airside pressure drop and material costs, a range of fin densities and material thicknesses were evaluated. A two-row counter-flow (CF) heat exchanger was also explored. A summary of all of the designs that were modeled for this award are shown in Table 2.

DESIGN	OPTION	HX PN	Description	Tube	FPI	Fin Thk
Baseline	-	1E006447	OEM Production, Overhang Fin	Current	20	.0045"
1	A	EPR25309	Flush Face	New	20	
	B	EPR25310	Flush Face	New	19	
	C	EPR25311	Flush Face	New	18	
2	A	EPR25312	Flush Face	New	20	.0037"
	B	EPR25313	Flush Face	New	19	
	C	EPR25314	Flush Face	New	18	
3	C	EPR25468D1	CF - 2 row, Flush Face	New	13	.0045"

Table 2: Summary of HX models evaluated.

Modeling of the HX's was performed using Modine's internally developed software, MPower. It is Modine owned software that was developed in conjunction with the University of Valencia, Spain. MPower is a user friendly, robust software package that includes comprehensive parametric study capabilities and provides the opportunity to design round tube plate fin coils and microchannel heat exchangers in a general or detailed system model. The software was utilized to predict heat exchanger thermal and pressure drop performance and was later used to correlate test data to the model to provide more accurate modeling capabilities.

Figure 5 is a Pressure-Enthalpy (PH) diagram for R-134a and is referenced for the following description of the method used to model the heat exchanger performance. In this liquid pumped refrigerant system, the refrigerant enters the HX in a slightly subcooled condition and exits the HX in a high quality two-phase liquid / vapor mixture. This is shown in the figure with a horizontal dotted line at 10°C connecting the two triangle icons that represent the inlet and outlet conditions for a delta enthalpy (dh) labeled "dh1". However, the modeling software is designed for direct expansion, two-phase refrigerant evaporators where the inlet condition is a low quality two-phase mixture and the outlet is superheated vapor. In order to model a pumped liquid HX using MPower, the refrigerant had to enter the HX at a slightly two-phase condition (quality $x \approx .02$) and leave slightly superheated (superheat $\leq 0.5^\circ \text{C}$). This modeled condition is shown as a solid red horizontal line connecting the blue dots, where the blue dots represent the inlet and outlet for a delta enthalpy labeled as "dh2".

This method of modeling the heat exchanger performance is believed to be a reasonable prediction of performance for the following reasons. It is assumed that the enthalpy difference between the actual and modeled conditions, "dh1" and "dh2" are nearly equal. The latent heat of vaporization results in a constant surface temperature of the heat exchanger surface, which is driven by the saturated temperature, or pressure, of the refrigerant. Furthermore, the actual volume of the heat exchanger that contains sub-cooled refrigerant is very small and probably does not extend past the inlet header as the liquid is quickly converted to a two-phase mixture upon entering. A detailed example of the effects of subcooling on performance proves this point. If 5° C of subcooled R134a were present at the inlet connection of the HX, where the saturation temperature is 10°C and the outlet refrigerant condition is saturated vapor, the

theoretical improvement in capacity is 3.4%. This is derived from the dh without subcooling having a value of 190.7 kJ/kg as compared to a dh of 197.5 kJ/kg with the 5°C of subcooling. It should be noted that this increase in capacity is probably not realized due to the phase change of the refrigerant occurring in the inlet header, before it reaches the portion of the heat exchanger that encounters the entering air.

Reliability of the predictions is further improved once the model has been correlated to the test data by adjustments made to the “enhancement factors”. Any improved accuracy that may have been obtained by creating a new model to work with actual subcooled inlet and saturated outlet conditions would not justify the additional cost and time that would have been required.

The experimental setup was controlled by monitoring the inlet and outlet pressure and temperature. Outlet superheat is calculated real time and manual adjustments to a bypass valve were made to maintain the desired saturated outlet condition.

It should be noted that the saturated refrigerant temperature of the heat exchanger will be slightly lower at the refrigerant outlet due to the pressure drop through the heat exchanger. This pressure drop is unavoidable and increases with refrigerant flow rate. The effects of refrigerant pressure drop are represented in Figure 5 with the dotted yellow line extending to the green dot at the exiting condition. MPower accounts for this pressure drop in the performance predictions.

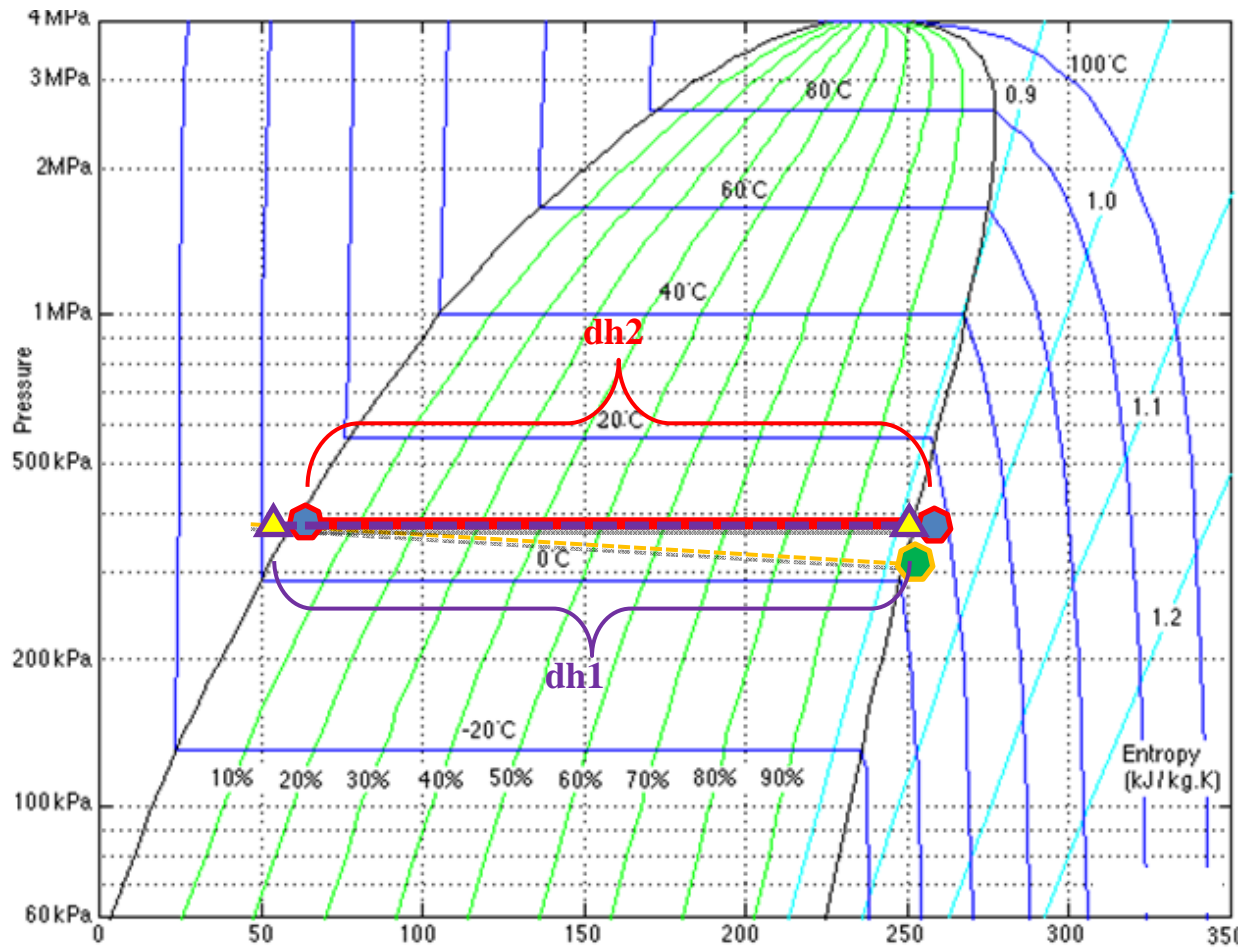


Figure 5: Pressure-enthalpy (PH) diagram for R134a.

Table 3 gives the predicted performance of the Baseline and Task 1 options at a 10° C saturated refrigerant inlet temperature and airside conditions from Table 1. Airside velocities were calculated at the inlet face of the heat exchanger. The flush face fin designs at both material thicknesses were predicted to significantly decrease the airside pressure drop with a minimal decrease in capacity.

The MPower software allows for calibration of the model if test data is available for the heat exchanger. “Enhancement factors” for refrigerant and airside performance and pressure drop can be adjusted to match the test data. The test data from this project was used to calibrate the model and allow for more accurate predictions going forward. Detail on the calibration methods is provided later in the report.

Table 3: MPower original predicted performance.

	PREDICTED PERFORMANCE [At 10°C Saturated Refrigerant Temp]						
	Airside						R-134a
	T air in °C	Velocity air in m/s	DPa Pa	Delta DPa %	Capacity kW	Delta Q %	DPr kPa
Baseline 20 FPI Overhung Fins .0045 fin	40.0	1.1	15.8		3.8		4.3
		2.0	38.4		6.1		10.0
		2.5	54.3		7.2		13.8
		3.0	71.8		8.2		17.9
	50.0	1.1	15.8		5.1		7.3
		2.0	38.2		8.1		17.5
		2.5	53.9		9.5		24.1
		3.0	71.3		10.8		31.0
Option 1A 20FPI / .0045" Flush Face EPR25309	40.0	1.1	13.8	-12.5%	3.7	-3.3%	2.5
		2.0	33.6	-12.5%	5.8	-4.2%	5.6
		2.5	47.4	-12.6%	6.8	-4.8%	7.5
		3.0	62.8	-12.5%	7.8	-5.2%	9.6
	50.0	1.1	13.7	-12.7%	5.0	-3.0%	4.2
		2.0	33.6	-12.0%	7.8	-3.5%	9.7
		2.5	47.1	-12.5%	9.3	-2.4%	13.3
		3.0	62.5	-12.5%	10.6	-1.9%	17.3
Option 1B 19FPI / .0045" Flush Face EPR25310	40.0	1.1	13.0	-17.7%	3.6	-5.2%	2.4
		2.0	31.6	-17.7%	5.7	-6.6%	5.3
		2.5	44.6	-17.8%	6.6	-7.3%	7.2
		3.0	59.1	-17.7%	7.6	-8.0%	9.1
	50.0	1.1	12.9	-18.0%	4.9	-4.7%	4.0
		2.0	31.4	-17.7%	7.6	-5.9%	9.3
		2.5	44.3	-17.7%	9.0	-5.1%	12.7
		3.0	58.7	-17.7%	10.3	-4.7%	16.3
Option 2A 20FPI / .0037" Flush Face EPR25312	40.0	1.1	13.0	-17.5%	3.7	-3.3%	2.5
		2.0	31.7	-17.5%	5.8	-4.2%	5.6
		2.5	44.7	-17.6%	6.8	-4.8%	7.5
		3.0	59.2	-17.5%	7.8	-5.2%	9.6
	50.0	1.1	13.0	-17.8%	5.0	-3.0%	4.2
		2.0	31.7	-17.0%	7.8	-3.5%	9.7
		2.5	44.4	-17.5%	9.3	-2.4%	13.3
		3.0	58.9	-17.5%	10.6	-1.9%	17.3
Option 3C 13FPI / .0045" Flush Face 2 row	40.0	1.1	11.1	-29.8%	4.5	17.4%	7.1
		2.0	27.2	-29.3%	7.7	27.1%	20.2
		2.5	38.5	-29.1%	9.4	31.5%	30.7
		3.0	51.2	-28.7%	11.1	35.5%	43.4
	50.0	1.1	10.8	-31.5%	5.6	9.9%	11.0
		2.0	26.5	-30.5%	9.3	14.4%	30.8
		2.5	37.6	-30.2%	11.4	20.5%	49.8
		3.0	51.9	-27.3%	13.4	24.7%	75.3

Subsection 4.1.2: HX Build

Throughout the course of the project, Alcatel-Lucent placed orders for the HX options and quantities listed in Table 4 and detailed below. Decisions were influenced by the initial modeling, test results, input from Modine personnel, and consideration given to the amount of development time required of each option to be production ready. Each of these builds was delivered on time and on budget.

Table 4: Summary of options built and delivered to Alcatel-Lucent.

Task	Build	Design	Option	Description	Tube	Fin Thk	FPI	Qty	
I	1	1	A	PF, Flush Face	New	.0045"	20	5	
			C	PF, Flush Face	New		18	5	
	2	2	A	PF, Flush Face	New	.0037"	20	5	
			B	PF, Flush Face	New		19	5	
			C	PF, Flush Face	New		18	5	
	3	3	C	CF - 2 row, Flush Face	New	.0045"	13	6	
	Finless	-	-		PF, Finless, (14) tubes	New	-	-	3
					PF, Finless, (14) tubes	27mm			3
					PF, Finless, (10) tubes	27mm			3
	Final	1	A	PF, Flush Face, Star Fin Roll	New	.0045"	20	51	
B			PF, Flush Face, Star Fin Roll	New	19		9		
III	1	-	Baseline	PF, Overhang Fin, Production	Current	.0045"	20	14	

HX build 1 (Options 1A & 1C)

Build 1 was a conversion to a flush face fin with the standard .0045 inch fin thickness material. Fin densities of 18 and 20 fpi were built using fins manufactured from standard round roll tooling. A total of ten heat exchangers were shipped where five had a fin density of 18 fpi and five had 20 fpi. This build was completed on time and the heat exchangers were used for evaluation purposes.

HX build 2 (Options 2A, 2B, & 2C)

Build 2 was also a flush face conversion fin; however it utilized a lighter gage fin material of .0037 inches. Heat exchangers with a range of fin densities from 18, 19 and 20 fpi were also manufactured from standard round roll tooling. The lighter gage fin material had to be special ordered, but it arrived in time to conduct the build in parallel with the HX build 1 and also to ensure that the project remained on time and budget. Heat exchangers from build 1 and 2 were tested at Modine and at Alcatel-Lucent to determine which ones would be used for the large quantity “final build” and for salt spray tests at Modine. Five heat exchangers of each fin density were shipped for a total of 15 pieces. All of the Task 1 heat exchangers described in Table 4 (excluding the 27mm “Finless” HX) were built with the new 19mm microchannel tube, which increased free flow area by 21% and inside perimeter by 15% over the “Baseline” production tube. Please note the heat exchangers built in support of Task 3 used the current tube.

HX build 3 (Option 3C – “CF build”)

A third design was investigated and built, with the knowledge that this was not an option for the final build. This was a 2-row, counter-flow (CF) HX with a lower fin density of 13 fpi made from .0045 inch thick aluminum, as shown in Figure 6. The CF coil build also used the new tube. Six of these were shipped to Alcatel-Lucent.




Figure 6: HX Build 3 (CF, 2-row HX).

HX Finless Build

Three finless heat exchanger designs were specified by Alcatel-Lucent with varying design parameters such as tube pitch and tube depth in support of DOE ARCTIC SOPO Task 3. Modine manufactured a quantity of three of each of these for a total of (9) nine; however, neither testing nor modeling of these heat exchangers was performed by Modine. Two of the designs are shown in Figure 7.



Figure 7: Finless HX build.

HX Final Build (Options 1A & 1B)

Based on the previous sample builds, sample testing and production timing guidance from Modine, Alcatel-Lucent selected designs 1A and 1B for the large sample build. Both option 1A and 1B are made with the standard fin thickness of 0.0045 inches, which makes them a better candidate for near term commercial production. To maintain good control of the fin geometry in the large sample build, fins were run with dedicated perishable sample star roll fin tooling. A total of (60) heat exchangers were shipped with (51) of them at a 20 fpi density and the remaining (9) had a fin density of 19 fpi.

Subsection 4.1.3: Testing

Setup and System

Extensive testing was performed to compare the current production (baseline) performance to that of the options listed in Table 4. Tests were performed in Modine's Refrigeration Lab with the test HX located in one of the calorimeter test cells. This test cell was capable of producing the ambient temperature, humidity and airflow required to meet the test conditions, as defined in Table 1.

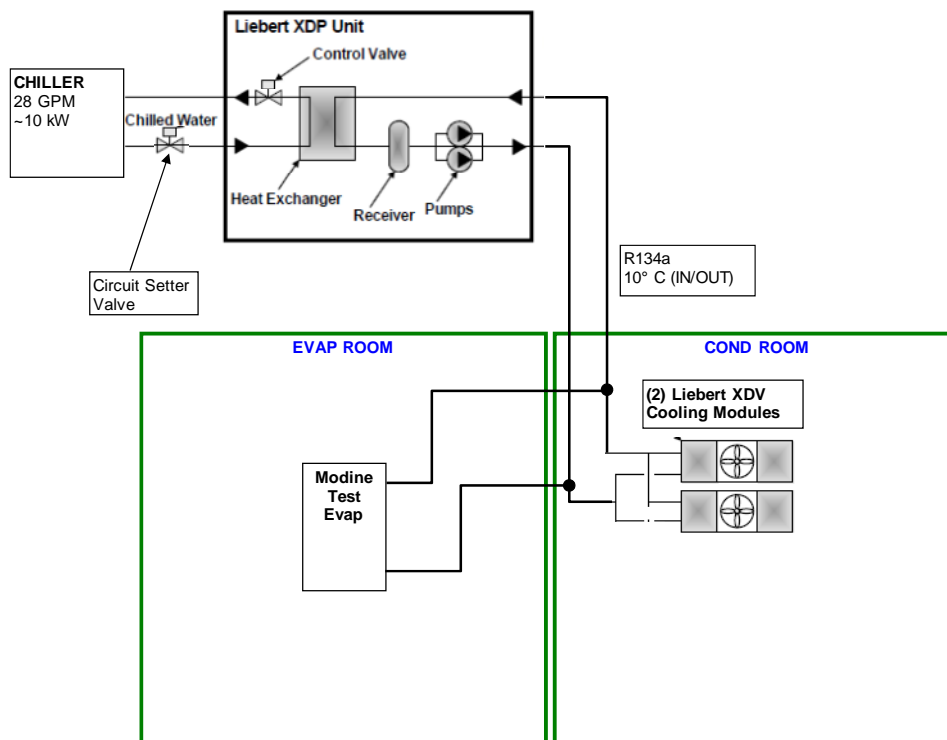


Figure 8: Modine test setup schematic.

The system, as shown in Figure 8 and Figure 9, consisted of a Liebert XDP pumping station and XDV cooling modules, a water chiller loop, and the test HX. The pumping station was located

outside of the calorimeter and connected to the test HX with pipe runs through the wall. The additional XDV cooling modules were installed on the main piping run in a calorimeter room that was adjacent to the first; however these modules remained idle throughout the testing. The test HX was mounted to an airflow measurement tunnel in the calorimeter room. Part of the Modine air tunnel ducting can be seen in Figure 9 to the left of the pumping station. Liebert and Alcatel-Lucent piping and installation bulletins were referenced during the installation of the test loop piping and components. The test HX was mounted with the tubes running vertically and the inlet connection at the bottom.

System Limitations

System pressure is the variable that controls the saturated temperature of the refrigerant. In order to maintain 10°C refrigerant, a system pressure of about 415 kPa was required. Operating at this pressure



Figure 9: Pumping Station and air duct behind calorimeter room (a), test HX and manifold piping (b), and XDV modules (c).

proved to be a challenge and was never actually achieved. Reasons for this were attributed to the short piping runs, too few bypass options for the refrigerant, and limitations of the water chiller capacity. As the bypass valves were used to control the refrigerant mass flow and exiting conditions of the test HX, a resulting elevated system pressure was unavoidable. Another limitation was the capacity and flowrate of the water cooled chiller that was used to reject the heat from the test loop. The Modine owned chiller was only capable of rejecting about 6 to 10 kW and pump water at about 106 liters per minute. The additional XDV modules were not operated due to the inability of the chiller to remove the additional heat load from these modules. Due to the low water flow rates available, the pumping station water valve actuator was disabled to allow maximum water flow to the pumping station at all times. This resulted in the loss of ability of the pumping station to control the refrigerant temperature as a means of preventing condensation on the HX; however, the chiller did not remove enough heat load for this to be an issue. A higher capacity chiller and longer piping runs may have resulted in a lower system refrigerant temperature and pressure, but these modifications would have required new equipment that was not covered in the ARCTIC project.

Performance Testing

Tests were conducted at 40° and 50°C ambient temperatures with the calorimeter room humidity maintained at a dewpoint below the refrigerant operating temperature. This was required in order to prevent condensing moisture on the HX.

The heat load from the specified ambient conditions was absorbed by the refrigerant flowing through the HX and rejected to the chiller (Figure 10) through the pumping station's refrigerant to liquid plate heat exchanger. As described in the modeling section, liquid R-134a enters the test HX at a slightly sub-cooled condition, absorbs heat from the test room and leaves the HX at a saturated vapor condition. The mass flow of the refrigerant was controlled through refrigerant bypass and shutoff valves to ensure this saturated exiting condition was maintained. The use of the chiller enabled the refrigerant entering the HX to be subcooled between 7 and 12 deg C for the range of tests. A thermal image of a flush face HX during testing is shown in Figure 11. It can be seen that the majority of the HX surface is at a uniform temperature. The small portion in the upper left hand corner shows some local superheat due to system instability.




Figure 10: Chiller shown to the side of the pumping station.




Figure 11: Thermal image of flush face HX during testing.

Due to the higher system pressures previously described, the saturated refrigerant temperatures at each condition were higher than the targeted 10°C. These actual refrigerant temperatures are shown in the test results listed in Table 5. The table also lists the actual HX capacities, airflows and pressure drops for each condition. Airside velocities were calculated at the inlet face of the HX from the airflow measurements. It is shown that a 5%-11% decrease in airside pressure drop over the baseline is possible with the flush face design of option 1A while keeping the same fin density. Additional reduction in airside pressure drop is possible with the thinner fin material and / or a reduction in fin density. The improvements made to the airside pressure drop on options 1A and 2A were achieved with only about a 0.7%-4.1% impact on HX capacity, as

compared to the baseline HX. Nominal capacity values in the table are lower than those predicted in Table 3; however this is due to the elevated refrigerant temperature that the system is operating at. These results were further evaluated in MPower, modeled at the desired 10°C saturated refrigerant temperature after the test data was used to calibrate the model. Alcatel-Lucent referenced the original predictions and the test results in selecting options 1A and 1B for the larger quantity final build. Option 1B was determined to have performance that is similar to option 2A; however, option 2A was not selected due to the thinner fin material having a longer development time to production requirement than the standard .0045 inch material used for 1B.

The two-row option, 3C, was tested under the same conditions; however the results did not influence the selection of HX for the final build due to this not being a near term production option. Test results showed a much higher airside pressure drop than originally predicted. This may be due to the geometry of the two row heat exchanger and the way it was physically modeled in MPower. Modine is currently enhancing MPower to improve the modeling accuracy for two row heat exchangers. For this project, the test data was used to calibrate MPower to provide a more accurate prediction tool. The measured capacity of the 2-row HX appeared much lower than the prediction, but this was largely due to the elevated saturated refrigerant temperature.

Table 5: Actual test results with saturated refrigerant temperature deviation.

	TEST DATA [At ACTUAL Saturated Refrigerant Temp]						
	Airside						R-134a
	T air in °C	Velocity air in m/s	DPa Pa	Delta DPa %	Capacity kW	Delta Q %	Tsat In °C
Baseline	40.1	1.1	13.5		3.5		13.3
20 FPI	40.0	2.1	32.6		4.9		15.8
Overhung Fins	39.9	2.6	46.0		5.3		16.7
.0045 fin	40.0	3.1	59.5		5.8		18.1
	49.9	1.2	11.4		4.4		15.7
	50.0	2.1	30.9		6.0		20.6
Option 1A	40.1	1.1	12.3	-9.0%	3.5	-1.5%	13.4
20FPI / .0045"	40.1	2.0	28.9	-11.5%	4.8	-1.6%	16.0
Flush Face	39.9	2.6	41.8	-9.2%	5.2	-1.9%	16.1
EPR25309	40.1	3.1	56.5	-5.1%	5.5	-4.1%	18.4
	50.0	1.1	10.2	-11.0%	4.4	-0.7%	15.4
	50.1	2.0	27.9	-9.9%	5.9	-1.1%	20.4
Option 1B							
19FPI / .0045"							
Flush Face							
EPR25310							
	Not Sampled / Test Data Not Available						
Option 2A	39.9	1.2	11.9	-12.0%	3.5	-1.6%	13.2
20FPI / .0037"	40.0	2.0	27.7	-15.3%	4.8	-1.5%	15.9
Flush Face	40.0	2.6	39.6	-14.0%	5.2	-1.7%	16.6
EPR25312	39.9	3.1	52.8	-11.2%	5.5	-5.0%	17.9
	49.9	1.2	9.9	-13.1%	4.4	0.6%	15.4
	50.0	2.1	26.7	-13.7%	5.9	-1.9%	20.4
Option 3C	40.6	1.2	16.6	22.4%	3.2	-7.8%	18.4
13FPI / .0045"	40.0	2.0	41.4	26.7%	5.2	7.5%	18.1
Flush Face	40.0	2.5	58.5	27.0%	5.9	12.0%	19.2
2 row	40.0	3.0	78.1	31.2%	6.4	10.5%	19.8

Subsection 4.1.4: MPower Correlation to Test Data

Test data is used to modify the air and refrigerant heat transfer coefficients within MPower, such that the model is fit to better predict the heat exchanger performance. Likewise, the same adjustments are made to the refrigerant and airside pressure drop coefficients. Once the model is calibrated at a certain point, it accurately predicts performance at similar conditions (i.e., Changing the inlet saturation temperature from 19C to 10C).

Figure 12 shows the actual heat transfer and airside pressure drop test results for the “Baseline” HX at 40°C ambient. This actual data is represented by the solid lines. The dotted lines represent the MPower model predictions at the same input conditions. For the Baseline HX, the default refrigerant and airside capacity coefficients of “1.0” were sufficient, as the modeled capacity was within 5% of the test capacity at each of the airflow conditions. For the other HX’s

evaluated, only minor adjustments to the heat transfer coefficients were required. However, modeling of the airside pressure drop was more involved due to the HX units operating in a “dry” condition.




Figure 12: Calibration of MPower model for baseline HX test data (heat transfer and airside pressure drop)

Two methods of modeling the airside pressure drop are shown in the figure with blue dashed or dotted lines. In one case the HX is modeled as an evaporator and in the other as a condenser. The dark blue dotted line with an airside coefficient factor of 0.47 was used when the baseline HX was modeled as an evaporator. The factor of 0.47 was selected to correlate the predicted pressure drop to the test data within 5% at each airflow condition. The 0.47 value significantly deviates from a default value of 1.0 because the evaporator model assumes a “wet” HX, and therefore higher airside pressure drop. The other way to model the airside pressure drop of the HX is to run it as a condenser (i.e., “dry”) at the same airflow conditions. Since this data center application requires a dry HX and the HX’s were tested at “dry” conditions, the pressure drop coefficient required to correlate the test data to a condenser model would be much closer to 1.0. The light blue dashed line represents the condenser model prediction with this non-optimized 1.0

factor. This factor could be optimized until the model agrees with the test data. For the purposes of this project, the evaporator model and airside pressure drop coefficient of 0.47 was used.

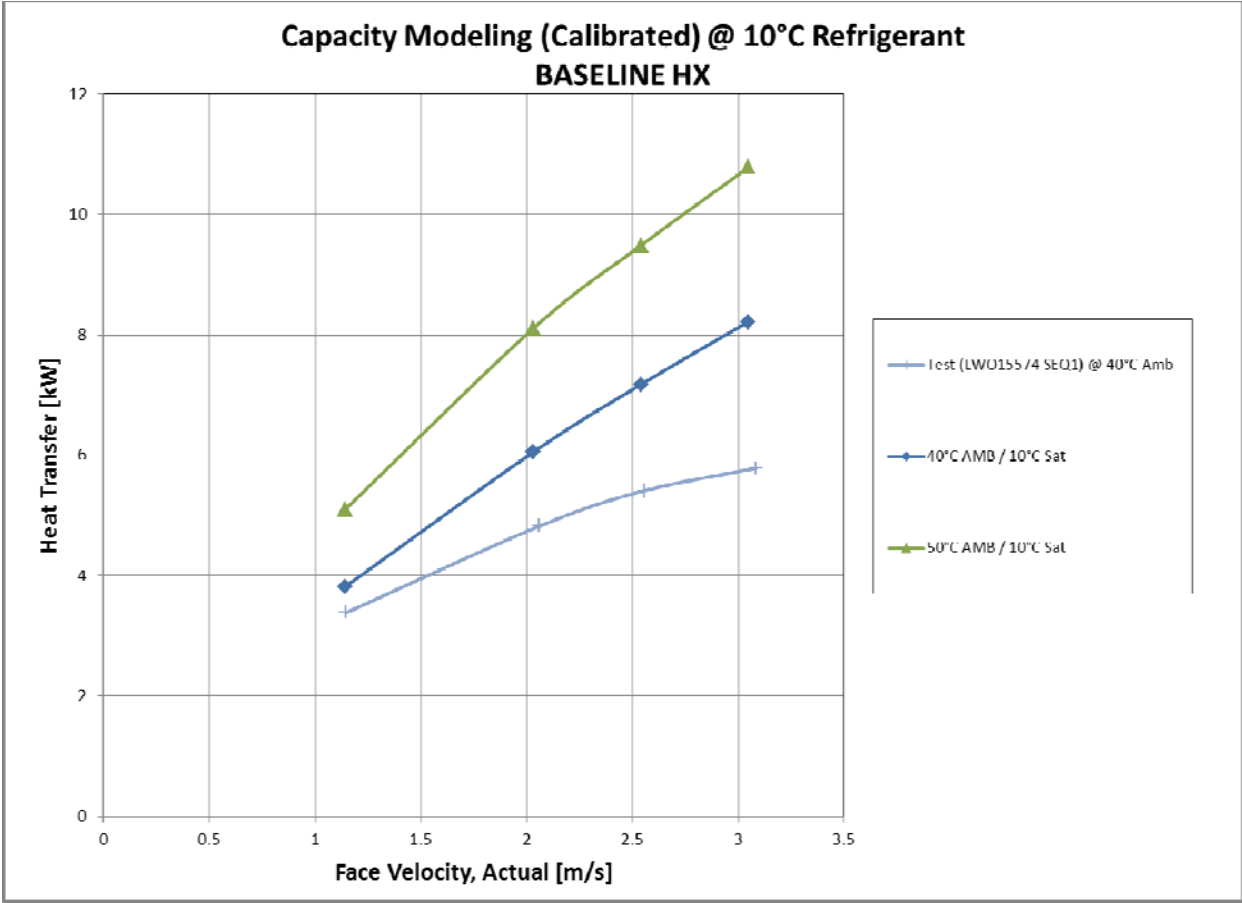


Figure 13: Calibrated model results of test data and prediction for "Baseline" HX at 10°C incoming saturated refrigerant temperature.

Table 6: MPower predictions at 10°C incoming saturated refrigerant temperature using calibrated model

	CALIBRATED PERFORMANCE [At 10°C Saturated Refrigerant Temp]						
	Airside						R-134a
	T air in °C	Velocity air in m/s	D _{Pa} Pa	Delta D _{Pa} %	Capacity kW	Delta Q %	D _{Pr} kPa
Baseline 20 FPI Overhung Fins .0045 fin	40.0	1.1	13.0		3.8		4.3
		2.0	31.9		6.1		10.0
		2.5	45.2		7.2		13.8
		3.0	60.1		8.2		17.9
	50.0	1.1	12.6		5.1		7.3
		2.0	31.0		8.1		17.5
		2.5	44.0		9.5		24.1
		3.0	58.6		10.8		31.0
Option 1A 20FPI / .0045" Flush Face EPR25309	40.0	1.1	11.8	-9.3%	3.8	-1.1%	2.6
		2.0	28.9	-9.2%	6.0	-1.4%	5.9
		2.5	41.0	-9.2%	7.1	-1.6%	8.0
		3.0	54.6	-9.2%	8.1	-1.9%	10.3
	50.0	1.1	11.4	-9.3%	5.1	-1.0%	4.3
		2.0	28.1	-9.3%	8.0	-0.8%	10.2
		2.5	39.9	-9.4%	9.5	0.6%	14.1
		3.0	53.2	-9.3%	10.8	0.3%	18.2
Option 1B 19FPI / .0045" Flush Face EPR25310	40.0	1.1	11.0	-15.7%	3.7	-2.9%	2.5
		2.0	26.9	-15.6%	5.8	-3.7%	5.6
		2.5	38.2	-15.6%	6.9	-4.2%	7.6
		3.0	50.8	-15.5%	7.8	-4.6%	9.7
	50.0	1.1	10.6	-15.7%	5.0	-2.6%	4.2
		2.0	26.2	-15.6%	7.9	-3.1%	9.8
		2.5	37.1	-15.7%	9.3	-2.0%	13.5
		3.0	49.5	-15.6%	10.5	-2.5%	17.2
Option 2A 20FPI / .0037" Flush Face EPR25312	40.0	1.1	11.3	-13.1%	3.8	-1.1%	2.6
		2.0	27.7	-13.1%	6.0	-1.4%	5.9
		2.5	39.3	-13.1%	7.1	-1.6%	8.0
		3.0	52.3	-13.0%	8.1	-1.9%	10.3
	50.0	1.1	11.0	-13.2%	5.1	-1.0%	4.3
		2.0	26.9	-13.1%	8.0	-0.8%	10.2
		2.5	38.2	-13.2%	9.5	0.6%	14.1
		3.0	50.9	-13.2%	10.8	0.3%	18.2
Option 3C 13FPI / .0045" Flush Face 2 row	40.0	1.1	15.9	22.2%	4.2	10.0%	7.0
		2.0	39.0	22.3%	6.9	14.4%	18.6
		2.5	55.3	22.3%	8.4	16.4%	27.2
		3.0	73.6	22.4%	9.7	18.2%	37.1
	50.0	1.1	15.5	22.7%	5.2	2.7%	10.4
		2.0	38.1	22.9%	8.5	4.3%	28.2
		2.5	54.1	22.9%	10.1	7.0%	42.4
		3.0	72.0	22.9%	11.7	8.7%	58.8

After the model coefficients were calibrated to the test data, the saturated refrigerant values were adjusted to 10°C to obtain accurate predictions of the HX at the originally requested conditions (i.e., 40°C and 50°C ambient). Figure 13 graphically shows the predicted performance of the baseline HX at these conditions. Performance of the baseline HX and other Task 1 options at these conditions are further detailed in Table 6. In the table, it can be seen that the capacity of flush face single row options 1A and 2A were within 2% of the baseline capacity with a 9% and 13% reduction in airside pressure drop, respectively. Option 1B improves further on the airside performance with a pressure drop reduction of 15% compared to the baseline, but with a negative capacity impact of 3 to 4%. These calibrated capacity predictions of the flush face HX units are even closer to the baseline HX capacity than what was originally predicted in Table 3.

Figure 14 and Figure 15 summarize the calibrated capacity predictions of the Baseline and options 1A, 1B, 2A and 3C at both 40°C and 50°C ambient conditions, respectively. The calibrated predictions of the airside pressure drop for the same HX models are summarized in Figure 16.

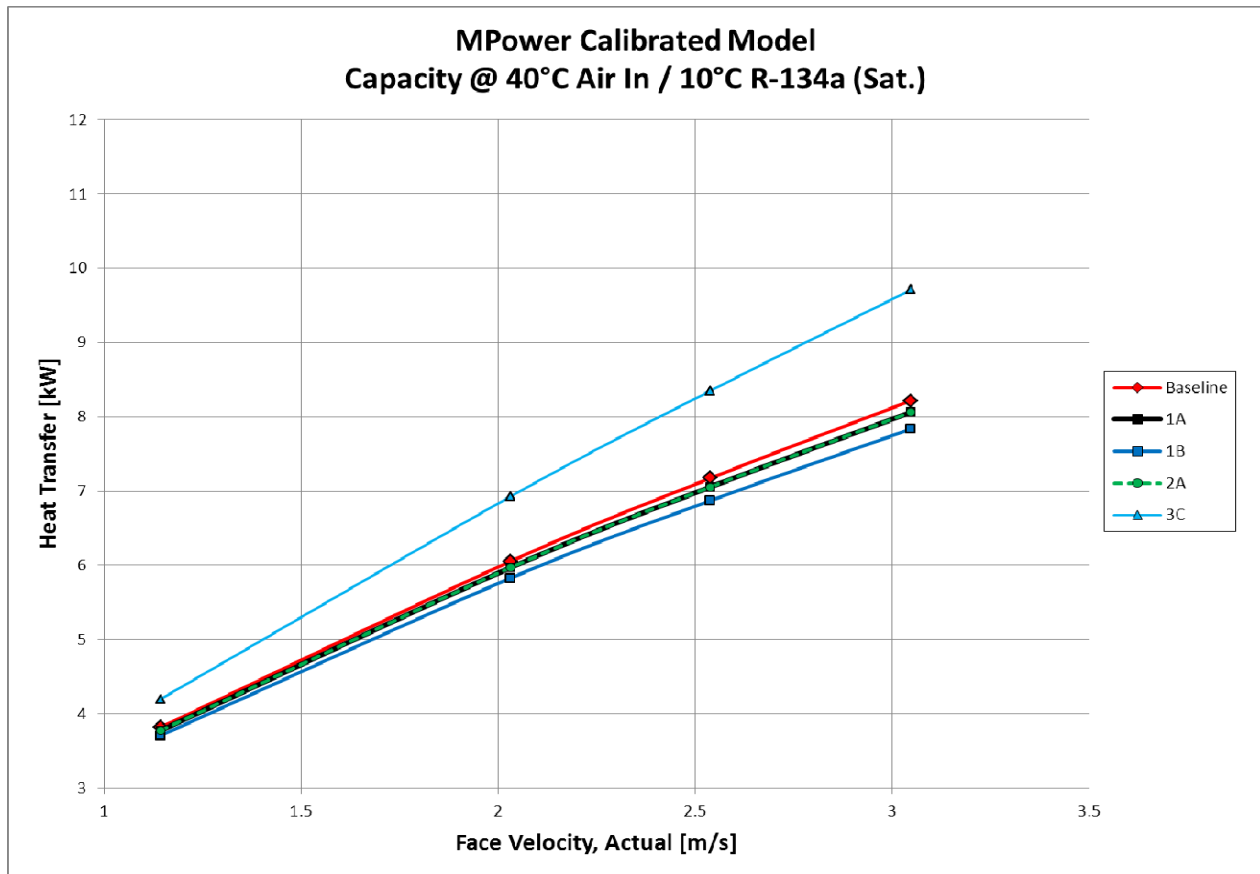


Figure 14: Summary of calibrated performance at 40°C ambient and 10°C R134a.

The MPower results clearly demonstrate that the flush face options are meeting the baseline performance at reduced airside pressure drop. It also allows a direct comparison of option 1B to 2A, and confirms that the 1B option has an even greater pressure drop reduction than 2A with a slight impact on capacity. The two row 3C option consistently provides the most performance of any option, but the airside pressure drop is also significantly higher than the baseline coil.

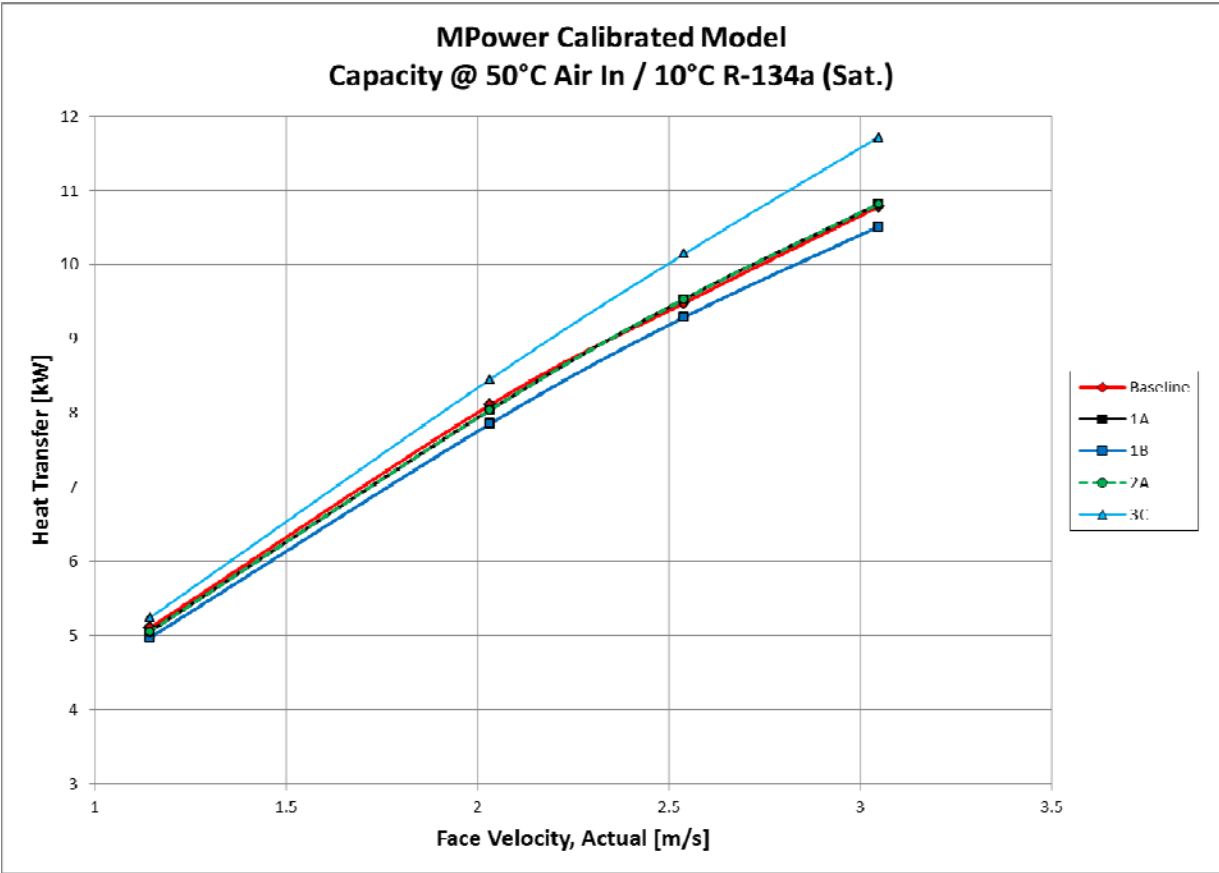


Figure 15: Summary of calibrated performance at 50°C ambient and 10°C R134a.

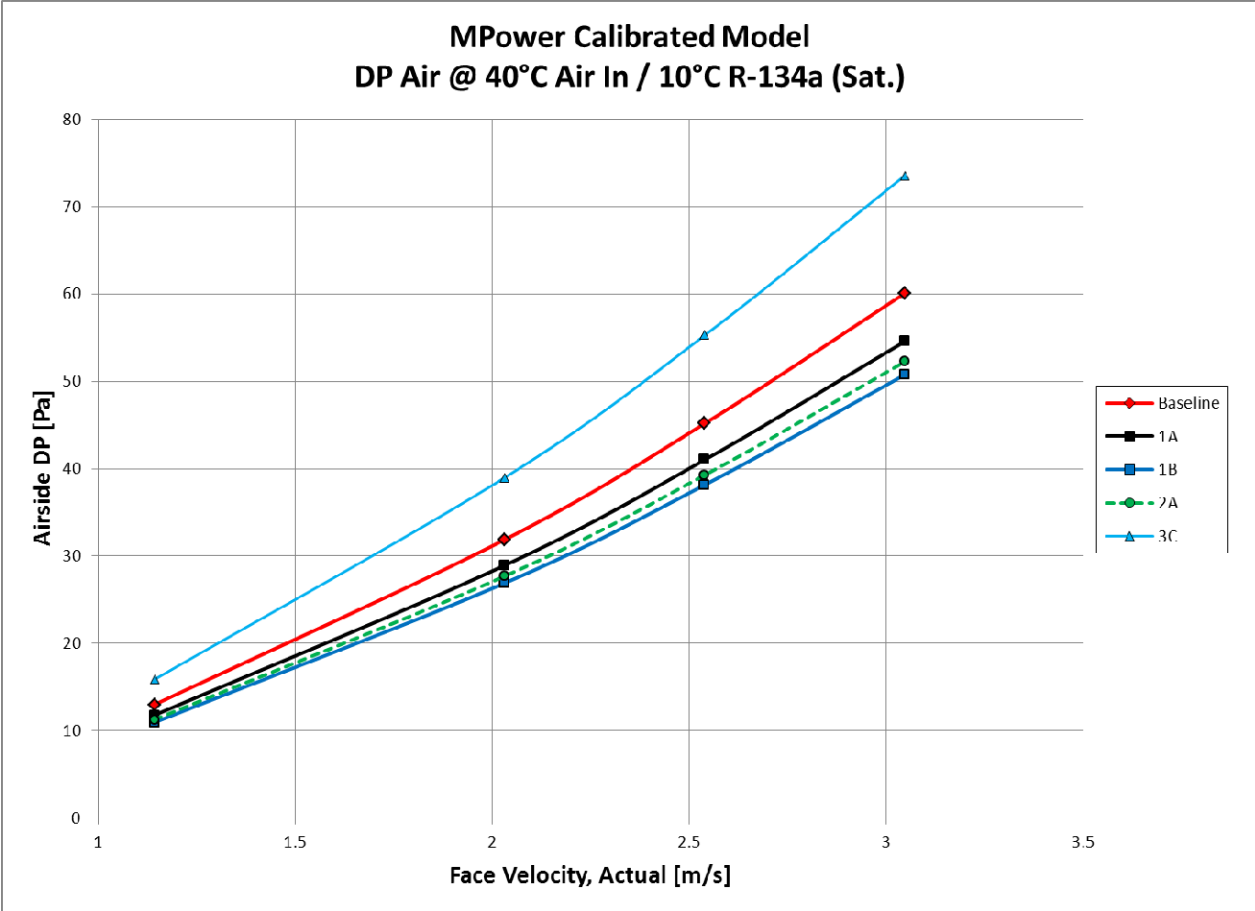


Figure 16: Summary of calibrated airside pressure drop.

Subsection 4.1.5: Salt Spray Tests

Three HX's were submitted for comparative neutral salt spray testing. They were manufactured in Racine between November 2010 and January 2011. The current production HX PN 1E006447 was used as the baseline and was compared to two of the flush face designs as presented in Table 7.

Table 7: Heat exchangers subject to salt spray test.

HX			Tube			Fin		
Design	Option	PN	PN	Spec. #	Tube Design	PN	Spec. #	Thickness
Baseline	-	1E006447	8E16801811	GM2028	Current	8E002054	GM2027	.0045"
1	A	EPR25309	8E32161811	GM2028	New	8E003216	GM2027	.0045", flush face
2	A	EPR25312	8E32161811	GM2028	New	EPR25327	GM2027	.0037", flush face

The heat exchangers were subjected to 96 hour ASTM B117 neutral salt spray test with the tubes horizontal and the face 15° from vertical as shown in Figure 17. They were removed from test, inspected and photographed every 24 hours. After the 96 hour inspection, they were rinsed with DI water. The rinse solution was collected and filtered through a Whatman 41 filter paper (22.5 micron nominal porosity). The filter paper was dried for examination.

No visual corrosion or flaking was seen at any of the time intervals (Figure 18). The dried filter paper was examined at 35x for any evidence of material flaking off of the heat exchangers. The filter papers showed only corrosion product (salt) with no metal flaking observed (Figure 19).




Figure 17: Salt spray test.



Figure 18: Option "1A" EPR25309 at 0 (left) and 96 hours (right).




Figure 19: Option "1A" EPR25309 Collection from DI rinse. No metal flakes at 96 hours. Examination at 35x magnification.

Subsection 4.1.6: Heat exchanger build in support of system-level testing

Modine supplied a total of (14) heat exchangers that were of the same design (PN: 1E006447) currently in use by Alcatel-Lucent in their Modular Cooling Solution. These samples were built in the Racine, WI sample shop and were intended to be used for an “under the floor” installation in an Alcatel-Lucent lab. This build was completed and shipped ahead of schedule.

Summary of Task 1

Task 1 efforts provided a significant amount of information regarding several optimized heat exchanger designs for a more energy efficient data center cooling solution. Performance test

data led to an improved prediction model for this application that is suitable for making performance comparisons to the current design with the ability to model at other conditions. Corrosion test data validated the changes to the new designs. In summary the Task 1 objective of providing heat exchangers that meet the required performance at reduced airside pressure drop has been achieved.

Section 4.2: Component-level air-cooled heat sink development (Task 2)

Air-based cooling technologies are widely used for thermal management of electronics in a broad range of applications. However, the continued increase in device heat densities and overall power dissipation is beginning to push the limits of air-based cooling technologies. For example, in many applications heat sink form factor is constrained, and the only approach to increasing heat sink surface area is to utilize tighter fin spacing or pitch. This in turn requires the use of more powerful fans and blowers in order to maintain large volumetric air flow rates through the heat sink. This approach, however, comes at the expense of two important factors, namely increased acoustic noise and greater energy requirements associated with the more powerful fans and blowers. There is thus a strong need for more energy-efficient and quieter fans and blowers as well as improved heat sink designs.




Figure 20: Three-dimensional heat sink structures studied by Krishnan et al. (2012)⁸: a) Schwartz structure; b) finned foam; and c) slotted honeycomb.

The state of the art in air-cooled heat sink technology is a parallel/longitudinal fin heat sink (LFHS), typically made from copper or aluminum. The heat transfer in these heat sinks is well understood and is dominated by two physical effects: i) thermal diffusion across the thermal boundary layer that develops adjacent to the fin surface; and ii) the sensible heat rise of the air flowing through the heat sink. Fin efficiency and heat spreading within the base can also be important factors in certain configurations, though their effects can be mitigated by increased fin width or a more conductive fin material, and the incorporation of heat pipes or a vapor chamber in the heat sink base, respectively. Assuming fixed fin thickness and a known pressure drop (or equivalently pumping power) across the heat sink, the optimal fin spacing typically corresponds to that pitch whereby the thermal boundary layers associated with adjacent fins merge at the heat sink exit^{9,10}.

A number of researchers have investigated novel heat sink structures such as pins, slots, bumps, foams^{11, 12}, etc. to name a few, that attempt to either increase surface area and/or the local heat transfer coefficient but without the undesirable requirement of increased fan or blower power that is observed in conventional parallel fin heat sinks. For example, several authors have studied heat sink designs incorporating unit cells having periodicity in all three spatial dimensions^{13,14,15,16,17,18}. Of particular relevance are the Schwartz, finned-foam and slotted honeycomb structures studied by Krishnan et al (2012); see Figure 20⁸. Krishnan et al. (2012)

attempted to take advantage of the following effects to enhance heat transfer in these structures: i) streaming flow in polygonal ducts; ii) introduction of slots to restart boundary layers and introduce mixing; and iii) introduce hydrodynamic instabilities to induce flow unsteadiness. Experimental results showed that the slotted honeycomb and finned-foam structures outperformed the LFHS at higher pumping powers. However, the process used to fabricate these prototypes, namely three-dimensional printing of a plastic part followed by investment casting, is an inherently expensive process and not amenable to large-scale manufacturing. The work presented in this task expands on the results of Krishnan et al. (2012) by focusing on slotted honeycomb heat sinks, as these structures, in contrast to finned-foam heat sinks, appear to have potential for being manufactured via a stamping process.

This section is thus organized as follows: i) Subsection 4.2.1 presents details on the stamping fabrication approach, with the particular application being an ultra-low-profile form factor heat sink of dimensions 72 mm in length by 125 mm in width by 5.7 mm in height; ii) test apparatus development and experimental characterization of 3DHS prototype performance are presented in Subsection 4.2.2; iii) details on the numerical models are presented in Subsection 4.2.3; and iv) experimental results are given in Subsection 4.2.4.




Figure 21: Schematic of layering approach to making a slotted honeycomb heat sink, where a pattern of N layers of the orange stamping (WAFER-F) followed by M layers of the blue stamping (WAFER-B) are repeated a total of K times. Cylindrical rods passing through holes patterned in the heat sink base are used to register stampings.




Figure 22: (bottom) Slotted honeycomb heat (corresponding to Design C in **Figure 23**) comprised of repeat patterns "F" and "B" consisting of multiple wafers within each pattern. Heat sink overall dimensions are 125 mm in width x 72 mm in length by 5.7 mm in height. Repeat pattern length is 6 mm for each "F" and "B" section and consists of nominally 23 or 24 wafers of thickness 0.010" each; (top) sections of wafers "F" and "B" making up each repeat pattern.

Subsection 4.2.1: Prototype fabrication

One of the challenges in the realization of 3DHS technology is the development of cost-effective and reliable manufacturing processes. Previous researchers have used fabrication methods based on three-dimensional printing, either directly creating the part in metal, or creating an intermediate part of the desired shape in plastic, and then using investment casting via the lost-wax process to "transfer" the shape into metal^{8,19,20}. Fabrication methods based on three-dimensional printing have the advantage of being able to create complex shapes that are monolithic. However, the approaches are, as yet, not cost-effective as they typically take several hours of printing time on an expensive three-dimensional printing machine, in addition to any post-printing operations such as investment casting. There is also concern regarding part warpage due to the creation of thermal stresses in large parts made via direct metal printing. One may also have a limited range of metals and alloys to choose from, particularly with direct metal printing.

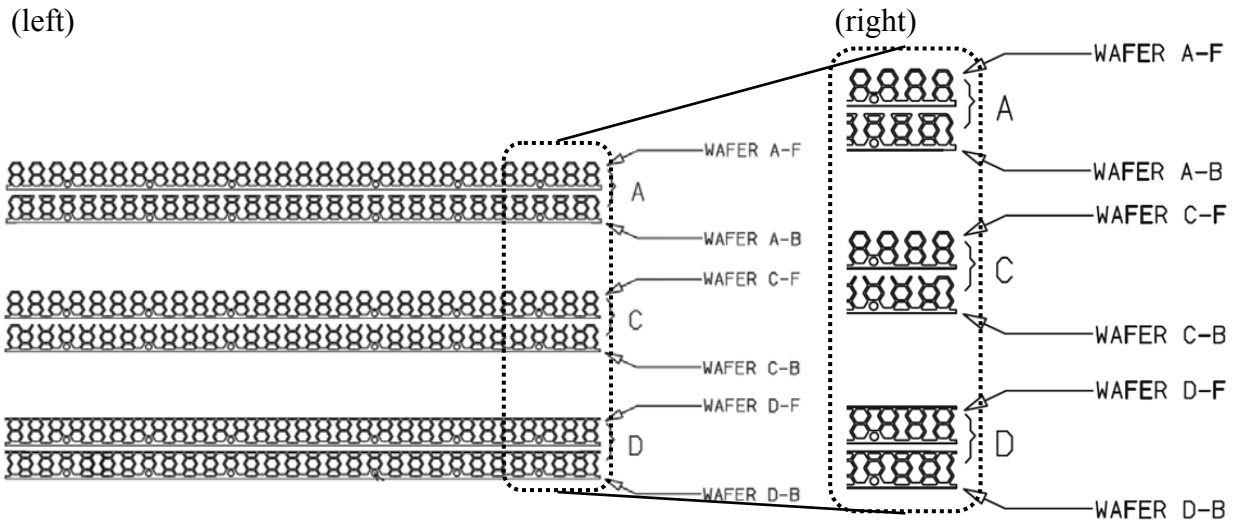


Figure 23: (left) Sample slotted honeycomb heat sink designs A, C and D depicting the two wafer layers "F" and "B" for each design; and (right) expanded view of designs.

To address the above concerns regarding cost and manufacture-ability, two approaches were used, both based on creating so-called “layers” or “wafers” that are subsequently sandwiched together to create the final assembled part. Here the slotted honeycomb heat sink can be thought of as comprising two different sets of repeating sections, e.g., an F section and a B section, whose geometry is chosen to reflect a particular slotted honeycomb heat sink design. Each section (F or B) is formed by placing multiple wafers of the same shape adjacent to each other. Figure 21 is a schematic that depicts how the layering process is accomplished, showing single wafers, denoted as “WAFER-F” and “WAFER-B”, which comprise each of the sections. Note that cylindrical rods that pass through holes patterned in the heat sink base are used to register wafers so they align precisely with each other.

Figure 22 shows a slotted honeycomb heat sink made by the above approach. The dimensions of this heat sink are 72 mm in length by 125 mm in width by 5.7 mm in height with each heat sink section F or B being 6 mm in length. Note that this is considered an ultra-low-profile heat sink because of its small height, and is characteristic of many applications in the telecommunications and data center space where there is a continued push to package greater device densities into a given form factor.. This heat sink is made by creating individual wafers, “WAFER-F” and “WAFER-B”, shown in the upper frame of Figure 22, that are subsequently stacked adjacent to each other to create the F and B sections shown in the lower frame of the figure. The wafers were made using a two-step process, whereby a mask of the wafer pattern is first created, and then the mask pattern is transferred into the wafer using a chemical etching process applied to thin sheets of metal. For the heat sinks studied in this work, several metals were used, including aluminum and copper, as well as an aluminum alloy with a thin layer of brazing material on the top and bottom surfaces of the sheet.



Figure 24: Linear fin heat sink 125 mm in width by 72 mm in length by 5.7 mm in height with a base thickness of 0.8 mm and having 50 fins, each with a height of 4.9 mm, width of 0.5 mm and a pitch of 2.5 mm that is CNC machined from aluminum with a conductivity of 160 W/mK.

The wafers made from aluminum or copper metal sheet were assembled into the desired heat sink design by registering the wafers onto 6 metal rods that aligned with holes patterned directly into the base of the heat sink. Approximately 24 wafers of thickness 0.010” make up each F or B section. The wafers are held in place by a crimp fixture located at the end of each metal rod as well as through the use of laser welding along a number of edges of the heat sink.

Figure 23 shows end-on views of the three slotted honeycomb wafer designs studied in this work. Note that the designs differ primarily with respect to their openness, or permeability, i.e., Design C has the most open structure for air flow; Design D has the least open structure, and is in fact, enclosed all along its top surface; and Design A is intermediate in openness between Designs C and D. All three designs have a base with thickness of 0.8 mm, and have a web thickness of 0.3 mm. Designs A, C and D were all manufactured in aluminum, while Design C was also fabricated in copper.


The baseline heat sink used for comparison purposes is a linear fin heat sink (LFHS) 125 mm in width by 72 mm in length by 5.7 mm in height with a base thickness of 0.8 mm and having 50 fins, each with a height of 4.9 mm, width of 0.5 mm and a pitch of 2.5 mm; the fin thickness and pitch are characteristic of a high-performance air-cooled heat sink design for use in an Alcatel-Lucent product. The LFHS is made of aluminum with a thermal conductivity of 160 W/m·K. shows a picture of the LFHS, prototype samples of which were made via CNC machining (see Figure 24).

Subsection 4.2.2: Test apparatus development


A. Design of Wind Tunnel Parts

For testing in a well-controlled laboratory environment that isolates the influence of external parameters on the measurements, a wind tunnel setup was built and instrumented. It consisted of a flow straightener, converging section, plastic screen/metal foam, test section, diverging section and a fan or a mass flow meter. The flow straightener was a honeycomb structure (see Figure 25) to obtain a near-plug profile for the air velocity field before it enters the test section. It was made, along with the converging and diverging sections (see Figure 25), by hard-printing technology


using a 3D Systems® Invision HR 3-D Modeler. A plastic screen or metal foam was employed to further homogenize the flow depending on whether a constant inlet velocity or flow rate experiments respectively were performed. The test section was machined from Plexiglas and a slot was provided in the test section for inserting the heat sink. The width of the slot was designed to allow no side flow bypass i.e., has the same width as the heat sinks. The height of the slot was designed to accommodate 1.5 mm of top flow bypass, or to incorporate a 1.5 mm thick Plexiglas plate along the base of the wind tunnel test section for the case of no top flow bypass.




Flow straightener



Converging section



Diverging section (small)



Diverging section (large)

Figure 25: CAD drawings showing key components of the wind tunnel test setup.

Note that during testing it was determined that the most appropriate experimental setup was based on constant mass flow rate. The reasons for this included comparison with numerical simulations as well as accuracy of measurement and experimental reproducibility (see Section B below). As such, we only provide details on the experimental setup for constant mass flow rate in the following.




Figure 26: (bottom) Wind tunnel setup for evaluation of air-cooled heat sink performance. Note that the Duocel metal foam homogenizer, depicted schematically in the top figure, is incorporated within the aluminum box located on the right-hand-side of the bottom figure; and (top) expanded view showing mass flow meter and direction of air flow.

B. Setup for Constant Inlet Mass Flow Rate

Even though a point-wise velocity measurement at the inlet is typically measured in the industry to gauge the thermal performance, it does not provide a good basis when performance between heat sinks with different flow/channel areas is compared or for comparing with numerical models. The more accurate parameter to be measured for a comparative thermal performance assessment is the mass flow rate. Another issue with measuring spatial velocity is that if the flow has profiles in the vertical or lateral flow directions, the value of the measured velocity might change depending on the sensor location and it can be a challenging task to ensure the sensor is reproducibly located at exactly at the same position for all the experiments with different heat sinks. Hence, a mass flow meter (Sierra Smart-Trak 2® flow meter with an accuracy of $\pm 1\%$ of full-scale of 19 g/s), which also acted as a flow controller, was used at the inlet of the test section to meter a specified mass flow rate of air into the test section. An air filter was inserted before the flow meter to prevent the chances of dust particles entering the test section. For flow homogenization before it enters the wind tunnel setup, an enclosed block of porous aluminum

foam (from Duocel, Inc.) was used. The test apparatus with the flow meter is shown in Figure 26.

Two static pressure ports were provided along the duct section located 35 mm from the inlet and the outlet of the heat sink across which the pressure drop was measured using a differential MEMS piezoresistive based amplified low pressure transducer (accuracy of ± 5 Pa over a range of 0-250 Pa) from All Sensors Corporation. The pressure ports were provided at an offset from the heat sink extremes to minimize the flow entrance and exit effects near the heat sink on the readings. A ball valve mounted in series with the wind tunnel setup (before the flow meter) was used to allow air flow into the setup. A wire screen was employed just before the flow meter to filter unwanted dust particles from entering the setup. Kapton® thin film heaters (0.2 mm thick, 350 Ω resistance) were used to provide heat to the heat sinks and, owing to their different sizes, a copper block was designed and used as an interface for effective and uniform heat spreading to the heat sink base. The heat sinks were attached to the copper block using thermal grease (OmegaTherm® 201) with a controlled (using spacers of <1% of the total heat spreading area) bond layer thickness of ~ 360 μm . The base temperatures of the heat sinks tested were obtained by measuring the temperature within the copper block and then subtracting the temperature drop due to the copper block and the thermal interface material. A power of 20 W (from a DC power supply) was supplied to the heaters in all the experiments.

Subsection 4.2.3: Numerical Model

Geometry and Boundary Conditions

The geometry of the 3DHS chosen to develop the numerical model is shown in Figure 27. The following assumptions and boundary conditions were implemented: steady, incompressible flow; uniform velocity at the inlet; convective flux/outflow at the outlet; symmetry on the left and right periodic boundaries; constant heat flux of 2222 W/m² (corresponding to a 20 W heat source) at the base. Air was assumed as the fluid and aluminum as the heat sink material.

To reduce the mesh size and the computation time, lateral periodicity of the model (in the width direction) was employed. For accurately capturing the flow-field that will be present in the experiment through the model, the inlet and outlet sections (150 mm each) of a duct (test section in the wind tunnel in the experiments) were also included. This allowed for the velocity profile to develop realistically before entering the heat sink without imposing an artificial plug velocity at the heat sink inlet. Heat source of the same size as the heat sink was considered in order to isolate the effect of spreading resistance in the heat sink base. A thermal grease layer with measured bond layer thickness was also incorporated into the model. The geometry of a conventional LFHS (design was based on the heat sinks being used in the Alcatel-Lucent products) was modeled using lateral periodicity and channel half-symmetry (same form factor as 3DHS, with channels 2 mm wide and fins 0.5 mm thickness). Aluminum was assumed as the material.

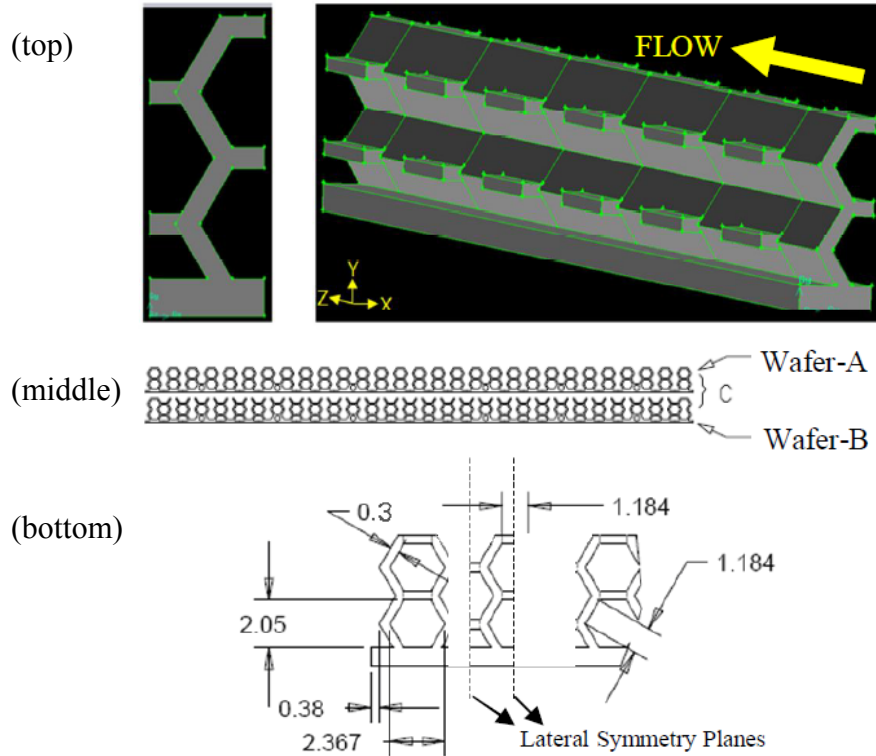


Figure 27: Computational geometry used for 3DHS Design C. Note the unit cell shown in the top and bottom figures, with the unit cell in the bottom figure being bounded by the lateral symmetry planes.

By considering the flow Re based on the air velocities at the inlet of the test section, it was found that the flow is laminar. But with slots in the 3DHS, it is likely that the flow could undergo an early transition. Hence, both laminar and low Re transitional models were employed for identifying the most appropriate model for this problem. A Shear Stress Transport (SST) correction was applied to the low Re transitional model to capture the flow physics in the near-wall regions. Commercial software, FLUENT (ANSYS, Inc.)²¹, was used to solve the set of governing equations in mass, momentum and energy implicitly using the specified boundary conditions. The pressure-velocity coupling was achieved using the SIMPLE algorithm. A second-order discretization scheme was used for momentum and energy equations, while a standard scheme was employed for the pressure field.

Changes in the pressure drop across the heat sink and the mean wall heat transfer coefficient were monitored with iterations for obtaining a true convergence (instead of monitoring default residuals for mass, momentum and energy that show premature convergence especially with separating and periodic flows). A tolerance of 10^{-3} Pa for pressure drop and 10^{-3} W/m²·K for heat transfer coefficient residuals, 10^{-6} for continuity and momentum residuals, and 10^{-10} for the energy residual were found to satisfactorily help in arriving at a truly converged solution. Default relaxation factors provided in the software were used for most of the simulations, but

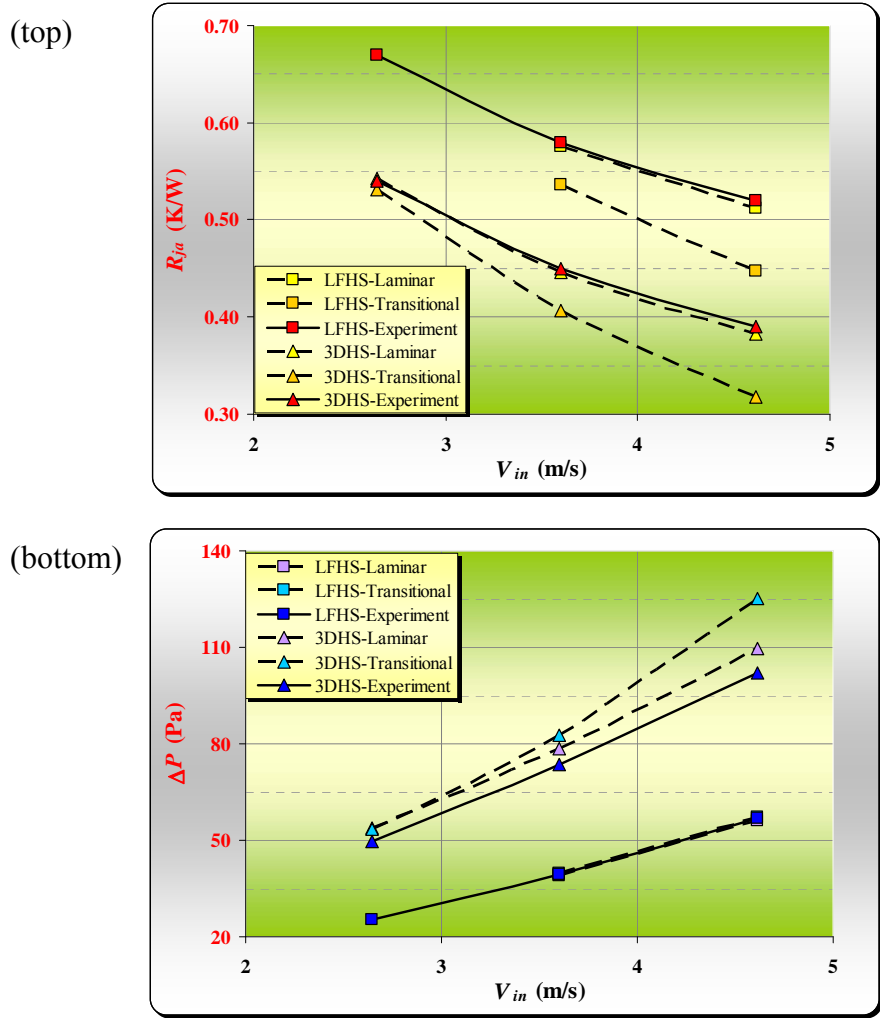


Figure 28: Comparison of laminar and transitional turbulent numerical model with experimental results for 3DHS Design C and LFHS. (top) junction-to-ambient thermal resistance; and (bottom) pressure drop.

minor under-relaxation was employed for high velocity flows. The solver was run in the double-precision mode.

Meshing and Grid Independence

The commercial geometry and grid generation code, GAMBIT (ANSYS, Inc.), was used to create and mesh the geometry. To enable faster convergence of the solution, a hexahedral meshing scheme was employed for the entire domain (including the solid and fluid portions). Grid sizes with 0.5, 1.3 and 2.7 million elements were used to check the solution for grid independence. The difference between the junction-to-ambient thermal resistance and the pressure drop across the heat sink were found to vary by less than 0.3% and 0.8% respectively for grids between 1.3 and 2.7 million elements. Hence, a grid with 1.3 million elements (with element sizes ranging from 85 μm to 250 μm) was employed for the simulations. The mesh was

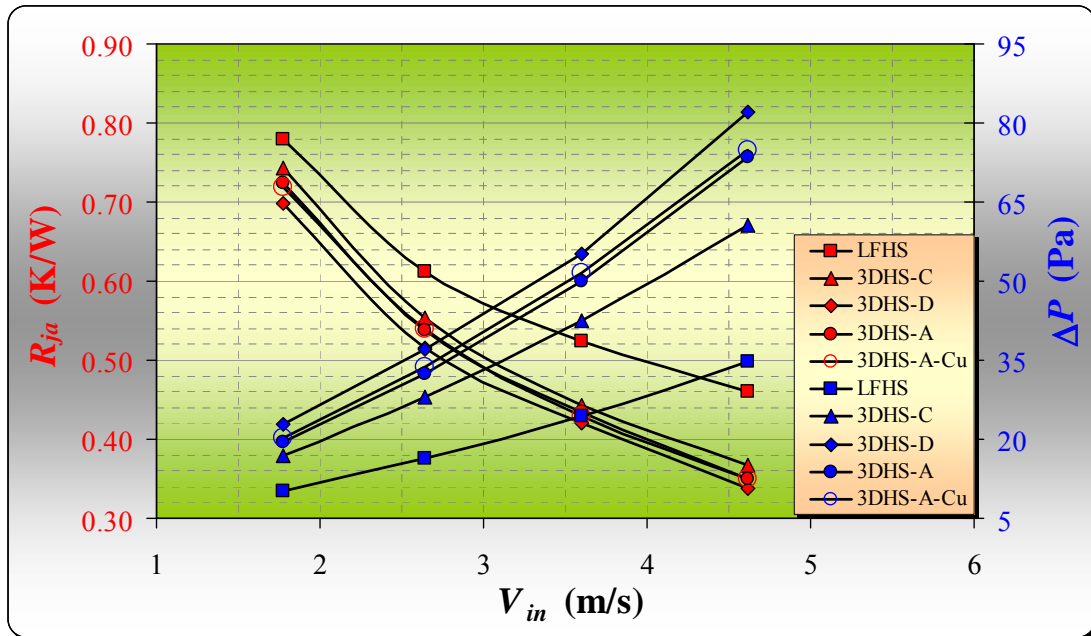


Figure 29: Experimentally measured junction-to-ambient thermal resistance and pressure drop as a function of average inlet flow velocity for 3DHS designs A, C and D and LFHS with 1.5 mm top flow bypass. Note that all 3DHS designs were made in aluminum, with design A also being made in copper (3DHS-A-Cu).

created such that it is fine in the areas close to the heat sink ligaments/walls/fins (where flow recirculation currents were expected to be present) and grows coarser in the regions farther from the ligaments (where the flow is more uniform). This technique helped in accurately capturing the flow field in the near-wall regions without having the need to resort to very large mesh sizes.

Subsection 4.2.4: Results

Verification of Numerical Model

The accuracy of the numerical model was verified with experiments. Figure 28 shows the comparison between experiments and the two numerical models employed for 3DHS Design C and the LFHS. For 3DHS, since the slots were assumed to cause a flow transition in the laminar regime, a low Re transitional flow model was also used to solve the problem. It is evident from the figure that the transitional model performs less satisfactorily compared to the laminar model but still compares reasonably well with the experiments. This implies that the flow is in the laminar regime and is predicted better by the laminar model as almost all the transitional models employ corrections to fully turbulent models at the walls and therefore perform better for moderate to high Re transitional flows (i.e., as the flow starts to develop more turbulent characteristics). The maximum differences between the results of the laminar model and the experiments were found to be 7.3% and 1.8% for pressure drop and junction-to-ambient thermal resistance respectively, and were found to occur at the velocity of 4.6 m/s. As expected, a better match between the simulation and test results was observed at low velocities.

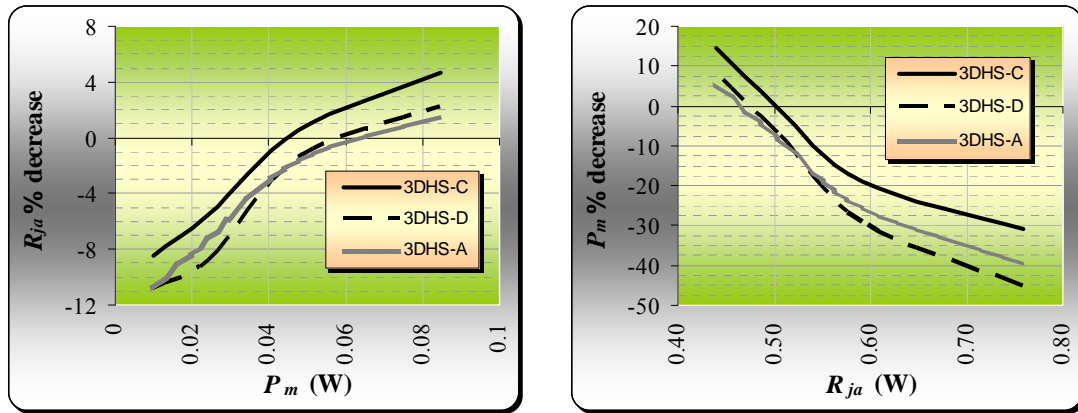


Figure 30: Comparison of 3DHS designs A, C and D relative to LFHS for 1.5 mm top flow bypass: (left) junction-to-ambient thermal resistance reduction as a function of pumping power; and (right) relative reduction of pumping power as a function of junction-to-ambient thermal resistance.

1.5 mm Top Flow Bypass

Results of the LFHS and 3DHS experiments are plotted in Figure 29. As expected, the pressure drop (ΔP) and the junction-to-ambient thermal resistance (R_{ja}) showed increasing and decreasing trends respectively with the duct inlet velocities (V_{in}) which are plotted on the x-axis. As observed from the figure, there will be an optimum heat sink design for which the improvement in thermal resistance of 3DHS over LFHS overtakes the increase in pressure drop. This condition would yield a positive performance benefit and will depend on the flow conditions and the heat sink geometry. Therefore, it is crucial to optimize 3DHS designs for specific applications to reap maximum performance benefits.

Under the considered ideal flow conditions in the wind tunnel with a 1.5 mm top flow bypass, for the condition of a constant pumping/fan power, it was found that 3DHS-C showed a maximum of 4.7% reduction in junction-to-ambient thermal resistance as shown in Figure 30. Similarly, using 3DHS-C will result in 15.1% energy savings per component compared to a LFHS for the same thermal performance (see Figure 30). In general, the overall thermal performance of 3DHS is better compared to LFHS at higher velocities under the considered test conditions.

A key finding from the results was that 3DHS with open features perform better than ones that are more closed due to low pressure drops in the former configurations (see Figure 29). This can be observed from Figure 30 wherein the performance curves for 3DHS-C consistently lie on the top of those for 3DHS-A and 3DHS-D in the considered range. This crucial observation helped

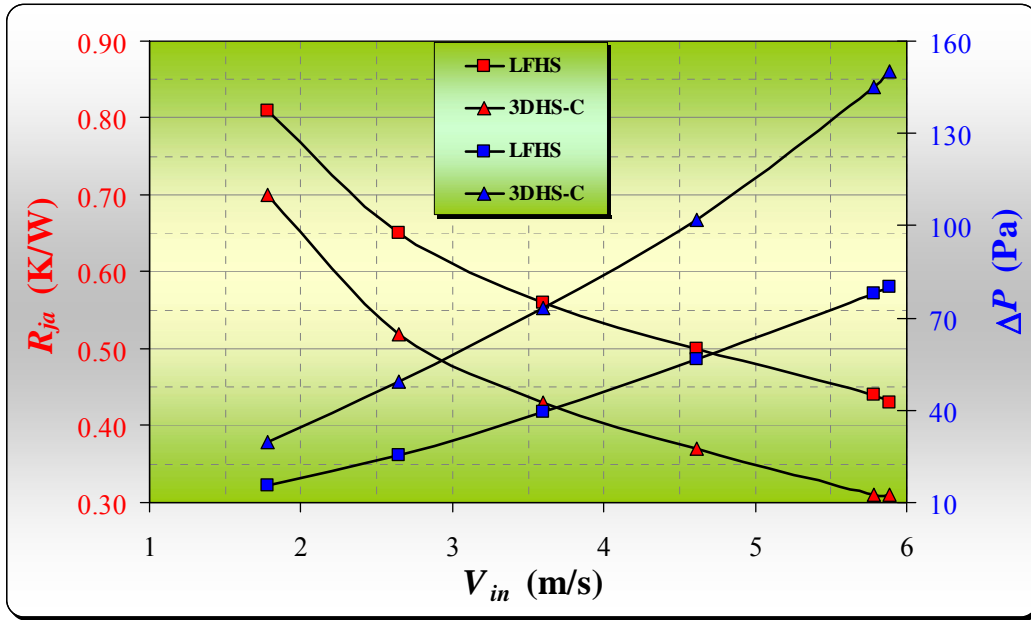


Figure 31: Experimentally measured junction-to-ambient thermal resistance and pressure drop as a function of average inlet flow velocity for 3DHS design C and LFHS with no flow bypass.

in narrowing down the focus of future designs based on 3DHS-C configuration which is also being used as a basic design for numerical optimization.

Another important observation from the experiments is that copper and aluminum 3DHS performed identically (sparing experimental errors) as could be seen in Figure 29 by comparing the curves for 3DHS-A and 3DHS-A-Cu. This implies air side thermal resistance is the dominating one and using copper for 3DHS does not help in improved performance but rather negatively aids in increasing the weight.

No Top Flow Bypass

Results comparing the performance of the LFHS and 3DHS under the condition of zero top flow bypass are plotted in Figure 31. Since it was found that 3DHS-C performed the best among other types in prototype set 1 for the case of 1.5 mm top flow bypass, it was the only design directly tested for these experiments. It is evident from the figure that with zero bypass, the pressure drop in 3DHS-C rapidly shoots up and the flow rate and the pressure drop ratio between 3DHS-C and LFHS is more compared to the same ratio for experiments with top flow bypass. But on the positive side, the thermal resistance ratio between 3DHS and LFHS is lower compared to the data from top flow bypass experiments.

Under the considered ideal flow conditions in the wind tunnel, for the condition of a constant pumping/fan power and acoustic noise, it was found that 3DHS-C showed a maximum of 19% reduction in junction-to-ambient thermal resistance, as shown in Figure 32. Similarly, using 3DHS-C will result in 57% energy savings per component compared to a LFHS for the same

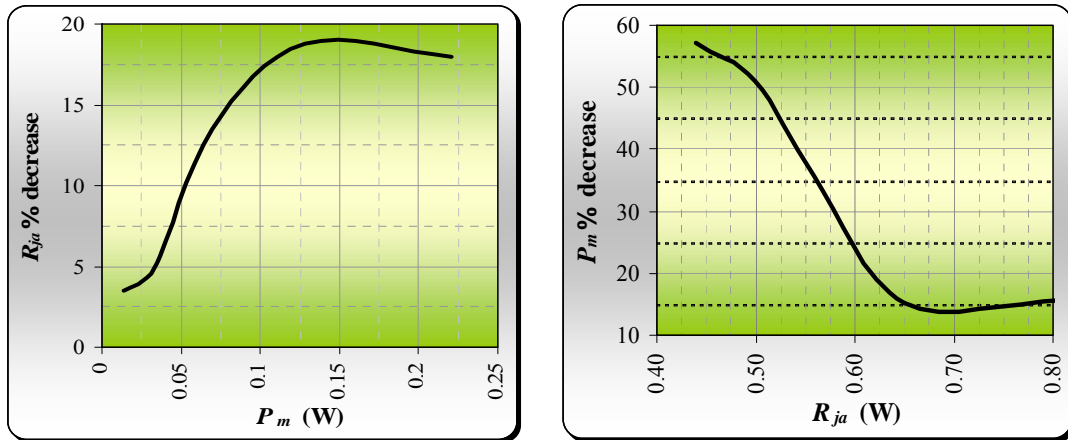


Figure 32: Comparison of 3DHS design C relative to LFHS for no top flow bypass: (left) junction-to-ambient thermal resistance reduction as a function of pumping power; and (right) relative reduction of pumping power as a function of junction-to-ambient thermal resistance.

thermal performance (see Figure 32). These improved results show that the ultra low profile 3DHS configurations outperform LFHS of the same form factor significantly only under the condition of zero flow bypass i.e., when all the air entering is forced through the heat sinks. With considerable flow bypass, air ceases to flow through the 3DHS (owing to its high pressure drop values), which will reduce the heat transfer performance, which in turn will lower the overall thermal performance improvement.

From a practical perspective, even though the application-specific optimized performance numbers might vary moderately, using 3DHS will result in notable energy savings in a data center considering the benefit will be achieved passively and without implementing complicated liquid cooling setups. The figure also shows that there exists an optimum inlet velocity for which the overall thermal performance is a maximum for any specific application. This elucidates the need to optimize 3DHS design based on the application dependent conditions (air flow rates, bypass and available volume).

Additional Experiments

Based on the above results, a numerical design of experiments (D.O.E.) was conducted to improve the performance of the 3DHS under the conditions of 1.5 mm top flow bypass. The particular flow condition that the designs were optimized for was $V_{in} = 4.617$ m/s. The numerical D.O.E. considered 3DHS Design C as the base design, which was observed to have the best performance relative to Designs A and D. Three variables, namely the width of the hexagons w , the number of repeat patterns of section F present in the structure n , and the ratio b of the length of repeat pattern F relative to the length of repeat pattern B were considered.

The D.O.E. was based on a central composite design with full quadratic response function. This described the performance of the heat sink in terms of a response function $RF(b, n, w)$ given by the following polynomial

$$RF(b, n, w) = a_0 + a_1b + a_2n + a_3w + a_{11}b^2 + a_{22}n^2 + a_{33}w^2 + a_{12}bn + a_{13}bw + a_{23}nw$$

Simulations were run to calculate the values of the parameters a_i and a_{ij} , which showed that the performance was dominated primarily by the width of the honeycomb w , with a weaker dependence on the number n of repeat patterns for section F, and least dependent on the ratio b .

The D.O.E. was used to create an additional set of prototypes. These included 3DHS Design C as well as additional designs where the width of the hexagons was doubled and tripled, and with small variations in the parameter b . This approach utilized aluminum sheet with a thin clad layer of brazing alloy on each side of the sheet. Masks were made for each design, and the particular shapes were transferred into the aluminum sheet via chemical etching. The resultant wafers were then stacked according to the desired overall geometry for the particular heat sink design, placed within a fixturing device, and heated within a salt pot, which caused the brazing alloy to melt and the stacked wafers to solidify into a single monolithic structure.

Unfortunately, due to the different rates at which the base aluminum sheet and the brazing alloy clad material chemically etched at, the resultant wafers contained very jagged exterior edges. This resulted in occlusion of the honeycomb channels, which had a significant and deleterious effect on the pressure drop, so that, although the overall thermal performance of several of the structures appeared to have improved, the pressure drop penalty negated any thermal performance benefit. Details of the experimental results are not shown here for brevity purposes.

Summary of Task 2

This task has advanced the knowledge of so-called three-dimensional heat sinks (3DHS) through modeling, fabrication and testing of prototypes. The particular focus for the task was on slotted honeycomb heat sink designs, as they appeared to offer advantages with respect to performance and manufacturability. Several manufacturing approaches based on stacking of individual wafers were investigated, including one based on registering wafers on cylindrical rods followed by laser welding along the exterior, and another based on brazing together of individual wafers containing a thin layer of cladding material so as to form a monolithic structure.

For the approach based on registering wafers on cylindrical rods followed by laser welding, 3DHS Design C was observed to have improved performance relative to a LFHS, particularly at higher pumping powers, for no top flow bypass. The performance enhancement relative to LFHS was diminished at higher pumping powers for 1.5 mm top bypass, and worse at lower pumping powers. No performance difference between copper or aluminum was observed for 3DHS Designs A, suggesting that fin efficiencies within this ultra-low-profile heat sink are fairly close

to unity. Detailed computational models based on steady laminar fluid flow and heat transfer showed excellent agreement for junction-to-ambient thermal resistance and pressure drop, whereas computational models based on steady transitional turbulent flow under-predicted junction-to-ambient thermal resistance and over-predicted pressure drop. A numerical design of experiments (D.O.E.) suggested that further performance improvement could be achieved, particular for the case of top flow bypass, by increasing the width of the honeycombs.

For the approach based on brazing together of individual wafers containing a thin layer of cladding material so as to form a monolithic structure, the heat sinks were observed to have a significantly greater pressure drop. This was attributed to unequal chemical etch rates for the aluminum sheet and the overlaid brazing alloy resulting in rough wafer edges which occluded the flow channels. We note that an alternative approach to making such a structure would involve creating a stamping die to stamp the wafer shapes out of the overlaid aluminum sheet. This process would not be expected to create rough edges, and hence may have yielded heat sinks with the desired pressure drop characteristics.

Section 4.3: Thermal backplane development (Task 3)

The presence of exposed, refrigerant-cooled surfaces in the ALU MCS modules invites their use as a “thermal backplane” for cooling of devices inside the rack shelves via direct or indirect contact instead of just extracting heat from the exhaust of cooling fans. As shown in Figure 33 below, this goal could be accomplished if we could devise thermal connections with sufficiently low thermal resistance and adequately high power capacity to make this contact.

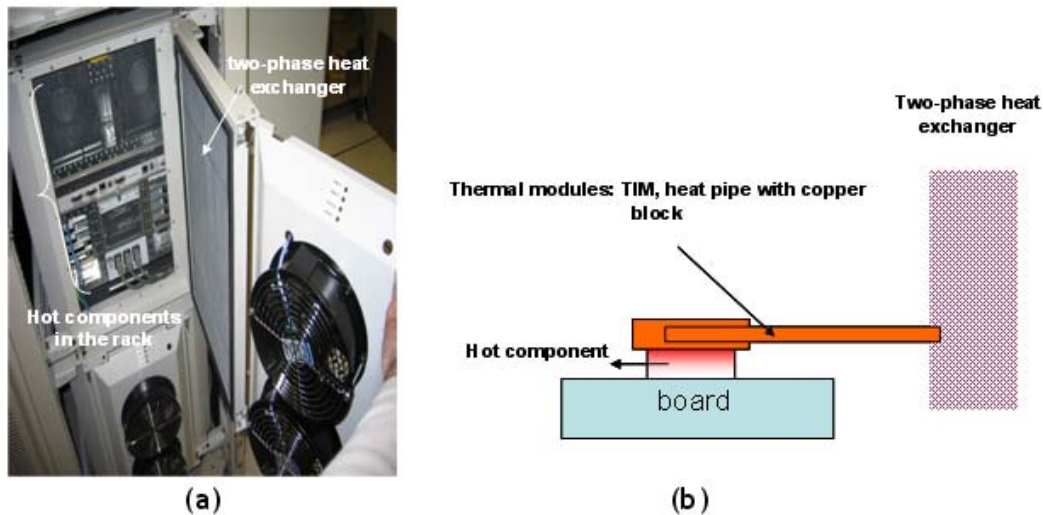


Figure 33: (a) Photograph of an open ALU shelf-level modular cooling unit seen from the rear of the rack. The surfaces of the 2-phase heat exchanger serve as cold plates in the proposed thermal backplane. (b) Schematic of a hot device on a circuit board in the rack shelf being cooled by the thermal backplane. A thermal module transfers the heat from the device to the thermal backplane.

As an estimate of the required performance of the thermal backplane and the connections to heated devices in the rack, let us assume that a 200-Watt device inside the rack shelf has a maximum allowable case temperature of 85 °C. If the thermal backplane has a refrigerant temperature of 25 °C, then the entire thermal path needs to have a thermal resistance R_{th} less than $60\text{ °C}/200\text{ W} = 0.3\text{ K/W}$. A passive conductor is unlikely to meet this specification: a copper bar of length 200 mm and cross-section 10 mm x 40 mm has a thermal resistance in excess of 1 K/W. Therefore, the thermal connection has to be made by something like a heat pipe or vapor chamber²², and we refer to such an object as a “thermal module”. The results of our investigations of several prototype thermal modules, summarized in Table 8, indicate that some embodiments of the thermal backplane can actually satisfy the stringent thermal-resistance specification of 0.3 K/W under heat loads in excess of 200 W.

The thermal backplane

To produce the thermal backplane itself, we arranged for the fabrication by our subcontractor Modine Co. of MCU heat exchangers without the air-side fins. Figure 34 below shows a photograph of this assembly. The multiport heat exchangers, which are horizontal in this

photograph, are aligned vertically when the unit is installed in a rack using the normal MCS heat-exchanger mounting frame. Thermal modules attached to devices to be mounted inside the rack shelf simply slide horizontally into alignment with the heat exchangers, to which they are clamped. All metal-to-metal thermal interfaces in these experiments were coated with Omegatherm 101 thermal grease to minimize contact resistance.



Figure 34: Photograph of a modified finless heat exchanger, intended to serve as a thermal backplane. The finless heat exchanger was made by Modine. The 12 multiport heat exchangers run horizontally in this view. Thermal modules will be clamped to these surfaces for cooling of on-board devices.

Thermal modules and mock ICs

Our test devices are copper blocks of dimensions 40 mm x 40 mm x 12.7 mm. We refer to these devices as “mock ICs”. One of the square faces is polished with 800 grit polishing paper to provide a smooth surface for contact with cooling modules. Cartridge heaters are installed in holes drilled through the blocks at mid-height, parallel to the polished surface, and a small, parallel hole is drilled close to that surface for a type T thermocouple, which serves to measure the temperature – this is referred to as the junction temperature, and its accuracy is $\pm 1^\circ\text{C}$. The cartridge heaters are powered by 117-Volt AC variacs, allowing the dissipated power to be continuously varied up to approximately 200 Watts. All metal-to-metal contacts between the mock ICs and cooling modules are made using Omegatherm 101 thermal grease. Refrigerant flow rates are measured using an ultrasonic flowmeter (accuracy 0.5 L/min), and inlet and outlet refrigerant pressures are measured using gauge pressure transducers (typical accuracy $\pm 0.5\%$ of reading). The temperature rise of the refrigerant is measured using Pt resistance thermometers attached to the inlet and outlet refrigerant manifolds; these are calibrated with an accuracy of better than $\pm 0.03^\circ\text{C}$.

Figure 35 shows the prototype design of a thermal module based on heat pipes manufactured by Vendor 3. Cooling of 150 Watts requires a design which incorporates three heat pipes. The condenser and evaporator ends of the heat pipes are embedded in copper heat spreaders with dimensions 40x50x15mm to match the size of our 40-mm-square mock ICs. The insertion lengths of the heat pipes in both the condenser and evaporator blocks is 45 mm. All three heat pipes have diameter 8mm. We have investigated the performance of modules of lengths 200 to 400 mm.




Figure 35: Engineering drawings of a 250-mm prototype thermal module fabricated by Vendor 3. The module consists of three 8-mm-diameter heat pipes and copper heat-spreader blocks.

Bench-level test apparatus details

Figure 36 below shows a drawing of the bench-level test apparatus we constructed to evaluate the performance of thermal modules. Visible on the left side of the drawing is the mock IC, held in place by alignment plates machined out of ULTEM thermoplastic. These hold the mock IC in contact with the evaporator end of the thermal module. Clamping pressure is applied by a “C” clamp which presses down the top of a large spring. The clamping force is measured by a load cell. The condenser side of the thermal module is pressed against a cold surface which consists of a rectangular copper tube with flowing water that is cooled by a refrigerating circulator. In typical experiments, the cooling water was held at 10 °C. Our results were unchanged when this temperature was increased as high as 30 °C.




Figure 36: Drawing of the thermal module test apparatus. The evaporator (left) end of the thermal module is clamped to the mock IC by a spring/load-cell assembly. The condenser end is clamped to a water-cooled cold plate.

Figure 37 shows two photographs of this test fixture. On the left, the spring-loaded clamping structure is being used to cool the mock IC with an air-cooled heat sink as a baseline for our thermal-module tests. The heat-sink module is manufactured by Vendor 7 for cooling Pentium-IV microprocessors. Its volume is substantially larger than that of the evaporator-section heat spreader in the thermal module. On the right, the full test fixture of Figure 36 is shown, with a focus on the water-cooled condenser end. In this picture, the clamp that holds the condenser assembly together has been removed for clarity.

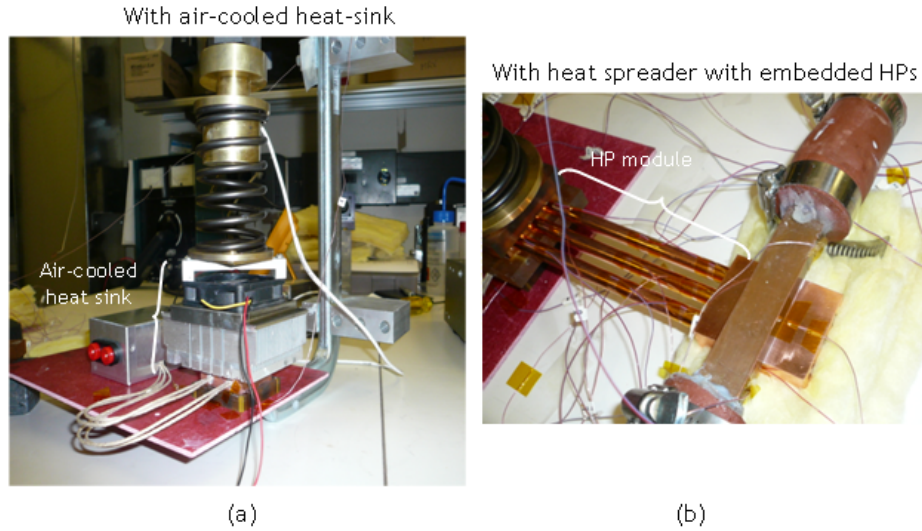


Figure 37: (a) Photograph of the clamping structure of Fig. 3.4 applied to the test of a Vendor 7 air-cooled heat sink module. (b) Photograph of tests of a 200-mm Vendor 3 heat-pipe thermal module. On the right, the condenser-section heat spreader is clamped to a rectangular water channel for cooling.

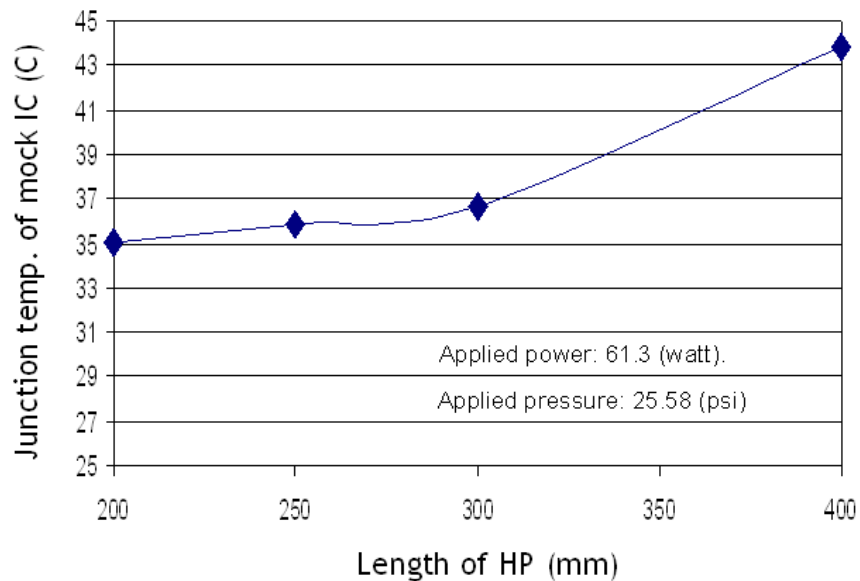


Figure 38: The mock-IC junction temperature is plotted as a function of the total length of the thermal module.

Figure 38 illustrates how the performance of the heat-pipe-based thermal module depends on its total length. Increasing the length from 200 mm to 300 mm makes only a small change, but the performance of the 400-mm module is substantially worse, with an increase in junction temperature of more than 8 °C.

Figure 39 compares the performance of a 200-mm-long thermal module with that of the reference air-cooled heat sink (Vendor 7), under identical conditions. The vertical offset of the two fit lines is related only to the difference in ambient temperature: 10 °C for the thermal module vs. room temperature for the air-cooled heat sink. More relevant to our task are the observations that (a) the thermal module is functional up to a power level of at least 170 W, and (b) the thermal resistance of the thermal module, which is equal to the slope of the straight line fit to the data points, is approximately $R_{th} = 0.29$ K/W, comparable with our estimated requirement. The thermal resistance and form factor are approximately 10% and 50% lower, respectively, than that of the air-cooled heat sink, thus meeting our milestone for this task.




Figure 39: The mock-IC junction temperature is plotted as a function of heater power for the 200-mm Vendor 3 heatpipe module (blue points and line) and for the Vendor 7 air-cooled heat sink (red). The thermal resistance of the heatpipe-based thermal module is about 30% less than that of the air-cooled heat sink.

Rack-level testing of the thermal backplane

With the basic parameters of thermal modules verified in bench-top tests, we went on to test the functionality of the complete thermal backplane, in which devices on a mock circuit board in a rack shelf are cooled by a thermal connection with the refrigerant-cooled backplane assembly shown in Figure 34. In addition to testing the heat-pipe modules from Vendor 3 (Figure 35) under these conditions, we also provided specifications to three vendors, Vendor 1, Vendor 2, and Vendor 4, to make three prototype T-shaped thermal modules of conduction length 250 mm. Photographs of these three modules are shown in Figure 40 below. The left end of each thermal module is the evaporator section, intended to be clamped permanently to the device being cooled. Thus, as the circuit board is slid horizontally into the rack, the vertical segments of the T come into alignment with the heat exchangers of the thermal backplane, to which they are clamped.

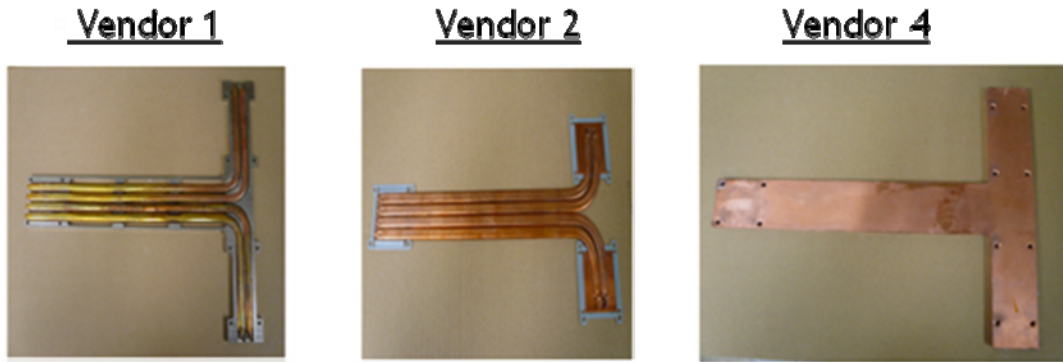


Figure 40: Photographs of three T-shaped prototype thermal modules tested using the refrigerant-cooled thermal backplane.

Figure 41 shows the mounting and clamping arrangement. For each device, the mock IC is mounted on a fiberglass mock circuit board (colored pink in the figure), which is held in the vertical plane by an aluminum base block. The thermal module (yellow) is mounted to the mock IC, and this assembly is slid into the rack just as a server blade or circuit pack would be. Once the condenser section of the thermal module comes into alignment with the finless heat exchanger (blue), it is clamped to it.




Figure 41: Sketches of fixturing used to simultaneously test three thermal modules in the refrigerant-cooled thermal backplane.

Figure 42 shows a photograph of the actual test, as viewed from the inside of the rack. We have tested individual thermal modules one at a time and have verified the results in simultaneous testing of three different thermal modules.




Figure 42: Photograph of the thermal modules under test, using the finless heat exchanger of **Figure 34** as a thermal backplane.

We have investigated three kinds of techniques to clamp the condenser sections of the thermal modules to the surfaces of the thermal backplane.

1. The baseline method is to sandwich the condenser section of the thermal module and the heat exchanger between flat clamping plates with mating screw holes, and simply screw them together, as shown in the left photograph in Figure 43. This is a reliable technique but requires time, tools, and dexterity
2. The second technique, shown in the right-hand photograph of Figure 43, utilizes split pins and springs, which are widely available for use in clamping heat sinks to ICs on circuit board. The clamping pressure is set by the geometry and the spring constants. Because commercially-available split pins come in a limited range of lengths, we designed a U-shaped clamping fixture that can use short split pins. This fixture can be assembled without tools, but disassembly requires pinching the ends of the split pins with pliers to release them.




Figure 43: Photograph of two methods used to clamp the condenser sections of the thermal modules to the thermal backplane, illustrated using a heat-pipe module from Vendor 3 as shown in **Figure 35**. Left: the condenser and heat exchanger are sandwiched between two aluminum plates with screw holes and are held together by screws. Right: the left-hand clamping plate is replaced by a U-shaped fixture which is held in place by commercially-available split pins and springs.

3. We also designed a spring-loaded latching fixture that allows even easier, tool-less assembly and disassembly simply by engaging or disengaging the latch (not shown).

Our quantitative measurements of thermal resistances did not reveal any thermal differences between these three clamping techniques.

Summary of results

Figure 44 gathers together measurements of junction temperature vs. heater power made for all the prototype thermal modules, as well as the air-cooled heat-sink baseline. The quantitative measures derived from these results are discussed with respect to Table 8 below. Here, we confine ourselves to some qualitative comments:

1. The two T-shaped modules fabricated by Vendor 1 are not capable of transporting more than 20 to 30 Watts before wick dryout, which is indicated by the sharp increase in slope of the associated curves. The 8-mm heat-pipe module could be used at powers of up to about 45 Watts, but operation in this regime is unstable.




Figure 44: The mock-IC junction temperatures obtained with all the thermal modules tested are plotted as functions of heater power.

- The Vendor 4 vapor chamber shows the best performance of all, going into dryout only at power levels above 170 Watts. The low-power slope of the associated curve – i.e., the thermal resistance, is the lowest of all. However, this module, being fabricated of full-length copper plates, is the heaviest of all.
- The Vendor 3 heat-pipe module meets the program milestone of exhibiting a thermal resistance and form factor which are at least 10% and 50% lower, respectively, than that of an air-cooled heat sink.

HT module	height above IC, including clamp	coolant	Rth	Q [W], mock IC at 85C	notes
Baseline					
air cooled HS	55.9 mm	air	0.32K/W	~215W	extrapolated from lower power
Thermal backplane configuration					
Vendor 4 VC	15.5 mm	water/R134a	0.13K/W	200W	very heavy module
Vendor 3 round 8mm HP	21.6 mm	water/R134a	0.25K/W	~260W	extrapolated from lower power
Vendor 3 flat 8mm HP	11.2 mm	water/R134a	0.86K/W	105W	hybrid-wick HP should do better
Vendor 1 flat 8mm HP	12.7 mm	water/R134a	2.7K/W	50W	

Table 8: Summary of results of thermal backplane experiments.

Section 4.4: Device-level liquid cooling development (Task 4)

Board-level test apparatus details

As in Task 3, our test devices are copper blocks of dimensions 40 mm x 40 mm x 12.7 mm. We refer to these devices as “mock ICs”. One of the square faces is polished with 800 grit polishing paper to provide a smooth surface for contact with cooling modules. Cartridge heaters are installed in holes drilled through the blocks at mid-height, parallel to the polished surface, and a small, parallel hole is drilled close to that surface for a type T thermocouple, which serves to measure the temperature – this is referred to as the junction temperature, and its accuracy is $\pm 1^\circ\text{C}$. The cartridge heaters are powered by 208-Volt AC variacs, allowing the dissipated power to be continuously varied up to approximately 880 Watts. All metal-to-metal contacts between the mock ICs and cooling modules are made using Omegatherm 101 thermal grease. Refrigerant flow rates are measured using a piston flowmeter (accuracy $\pm 0.2\%$ of reading), and inlet coolant pressures and inlet-outlet pressure drops are measured using gauge and differential pressure transducers (typical accuracy $\pm 0.5\%$ of reading). The temperature rise of the coolant is measured using Pt resistance thermometers attached to the inlet and outlet pipes; these are calibrated with an accuracy of better than $\pm 0.03^\circ\text{C}$. The entire assembly is wrapped in high-temperature foam insulation.

Indirect cooling of multiple devices on circuit boards

For this series of measurements, a test fixture was constructed using 5 mock ICs attached to a fiberglass board. The power levels and device dimensions and locations were set to match those on an ALU optical circuit board. The nominal power levels of the five mock ICs were 75W (green mock IC in the figure below), 100W (yellow), and 3 x 30W (white A, B, and C). In the experiments described herein, the junction temperatures were measured as a function of input voltage, with equal voltages on all 5 devices.

To cool these assemblies, we provided specifications to Vendor 1 to develop prototype embedded-heatpipe heat-transfer modules in two configurations, as shown in Figure 45.


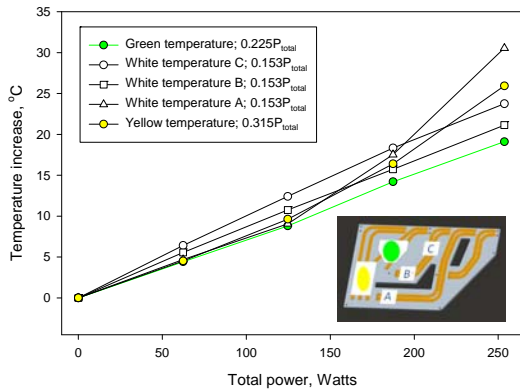


Figure 45: Two prototype embedded-heatpipe heat-transfer modules fabricated by AVC. The gray material is aluminum, while the brown stripes are embedded copper heat pipes. Left: horizontal configuration; heat is transferred to the top edge. Right: vertical configuration; heat is transferred to the right edge.

In the “horizontal” configuration, heat pipes transfer the heat to a strip along the top edge of the module. To this area, we clamped an extruded multi-port heat exchanger fitted with quick-connect fluid connectors, allowing refrigerant to be flowed through it. In the “vertical” configuration, the heat pipes conduct the heat to the right edge of the module. This allows the entire assembly to be slid into a rack shelf and contacted with a vertical heat exchanger in a finless refrigeration coil; thus allowing the circulating refrigerant in the existing ALU modular cooling solution to carry off the dissipated heat.

Figure 46 shows the junction temperatures as functions of total heater power, measured with equal voltages applied to all five devices. In the “horizontal” configuration, the refrigerant flow rate through the extruded multi-port heat exchanger was sufficiently high that only a small fraction of the refrigerant was vaporized, even at the highest power levels. In the “vertical” configuration, it was not possible to measure the flow rate, but our estimate is that this condition was also obtained in these experiments. Thus, the fact that the average slopes of the curves in the left-hand graphs is 2 to 3 times lower than those in the right-hand graph implies that the horizontal configuration is much more effective in cooling the mock ICs. We attribute this difference to the fact that the heat pipes in this module are shorter and closer to vertically-oriented, with the heat source at the bottom. In this configuration, buoyancy aids in heat transfer.

Junction temperature vs total power with equal voltages on all these devices. Vendor 1 horizontal multichip cold plate.



Junction temperature vs total power with equal voltages on all these devices. Vendor 1 vertical multichip cold plate.




Figure 46: Junction temperatures are plotted vs. total heater power, with equal voltages on all 5 mock ICs. Left: horizontal configuration. Right: vertical configuration.

Direct refrigerant cooling of individual devices

To evaluate the potential for direct liquid cooling of ICs, we utilized two refrigerant-cooled cold plates. The first is a low-profile cross-fin cold plate manufactured by Vendor 5. As shown in Figure 47, it is only 3 mm thick and has a “gull-wing” shape. The lower surface is polished using 800-grit paper for contacting to a mock IC. This was clamped against the polished surface of a mock IC using screws and a backing plate for some experiments. For others, we constructed a spring-loaded clamping device for potential use in a thermal-backplane system, allowing the heated device to be slid into place and clamped against the cold plate and later removed, both without tools.

The second cooling module we tested was a micro-fin cold plate manufactured by Vendor 6. This device is thicker than the Vendor 5 cold plate (15 mm thickness vs. 3 mm), but it promises to have much better performance due to the explicit optimization of the fin design for cooling




Figure 47: Two views of a Vendor 5 gull-wing cold plate.




Figure 48: Left: a photo of the Vendor 6 micro-fin cold plate, with attached fluid connections and ports for refrigerant pressure measurements. Right: micrograph of the internal cooling fins.

with R134a. A photo of the assembly and a micrograph of the internal fins are given in Figure 48.

Measurements were made under two complementary modes of refrigerant flow: (1) At high refrigerant flow rates ($f > 1$ L/min), the highest dissipated heater power was sufficient only to vaporize a small fraction of the flowing refrigerant. Under such conditions, the refrigerant temperature remains approximately constant, and the junction temperature increase is expected to be proportional to the dissipated power. (2) At low flow rates ($f < 0.3$ L/min), it is possible to increase the dissipated power enough to vaporize all of the refrigerant. Under such conditions, we expect to observe refrigerant superheat and a sharp increase of junction temperature.

Results – high flow rate

Figure 49 below shows the results obtained in the high-flow-rate limit for both cold plates. As expected, the junction temperature increase is proportional to the dissipated power level and is only weakly dependent on refrigerant flow rate. These results can be characterized by a thermal resistance R_{th} , which is equal to the average slope of the plotted curves for each cold plate. The values obtained in this series of experiments are $R_{th} = 0.08$ K/W for the Vendor 5 cold plate and $R_{th} = 0.05$ K/W for the Vendor 6 cold plate. Another useful index of the performance is $Q(85C)$, the power level at which a junction temperature of 85 °C is reached $Q(85C) = 925$ W and 1625 W for the Vendor 5 and Vendor 6 cold plates, respectively.




Figure 49: The junction temperature of a single mock IC is plotted as a function of dissipated power for several flow rates in the high-flow-rate regime, for both the Vendor 5 and Vendor 6 cold plates.

Results – low flow rate, Vendor 5 gull-wing cold plate

The Vendor 5 cold plate was also examined at several flow rates in the low-flow-rate regime. Figure 50 below shows the junction temperature plotted as a function of heater power for two flow rates. At the higher flow rate, the low-power slope (effective thermal resistance R_{th}) is 0.096 K/W, and the junction temperature starts to diverge at the very high end of the power range – at about 700 W. Decreasing the flow rate pushes this behavior to lower power levels, so that the junction temperature starts to take off at a heater power of around 300 W. The low-power slope is not determined quantitatively, but it is obviously higher for the dashed curve than for the solid curve. The power levels at which the junction temperatures begin to diverge are consistent with complete vaporization of the refrigerant, as determined from the known latent heat of the phase transition and the measured flow rates. The complete vaporization of the refrigerant leads to both lower heat capacity and a reduced heat-transfer coefficient, and these effects are responsible for the increase in the junction temperature at high power.




Figure 50: The junction temperature of a single mock IC cooled by a Vendor 5 gull-wing cold plate is plotted as a function of dissipated power for two low flow rates.

These effects are also reflected in Figure 51, which shows the refrigerant temperature increase observed in these two experiments. As seen in the open symbols, the refrigerant starts to superheat as the power level is increased above about 300 W, the same power at which the junction temperature starts to diverge. At the higher flow rate, the junction temperature begins to diverge at a proportionally higher power level – about 700 W. No refrigerant superheat is observed, but this is expected at higher power levels.




Figure 51: The refrigerant temperature increase is plotted as a function of dissipated power for two low flow rates, for the experiments of the previous figure.




Figure 52: The junction temperature of a single mock IC cooled by a Vendor 6 micro-fin cold plate is plotted as a function of dissipated power for three low flow rates.

Results - low flow rate, Vendor 6 micro-fin cold plate

Figure 52, Figure 53 and Figure 54 show the results of corresponding experiments performed using a Vendor 6 cold plate. The behavior is qualitatively the same as seen with the Vendor 5 cold plate, with the following differences: (1) there is less difference in the shapes of the junction-temperature curves at different flow rates, and (2) the low-power thermal resistance is lower by a factor of about 2: $R_{th} = 0.05 \text{ K/W}$.

The basic conclusion to be drawn from these experiments is there is no anomalous behavior that would prevent the use of these cold plates for direct device cooling in telecommunications equipment. At the highest power levels and power densities that are typically encountered in such equipment – a few hundred Watts per device and per cm^2 - these cold plates simply cool the devices without instability, as long as the refrigerant flow rate is sufficiently high that complete

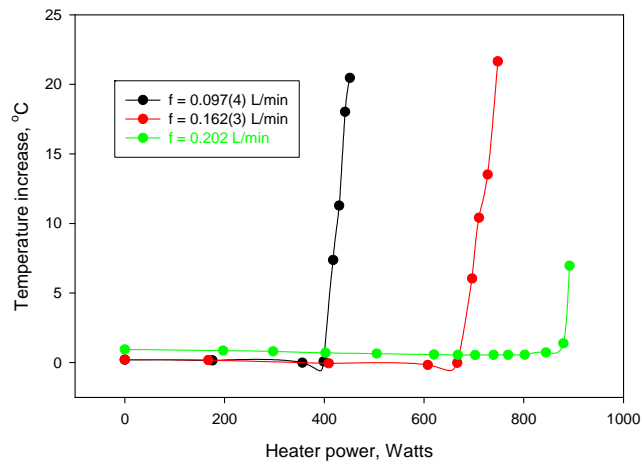


Figure 53: The refrigerant temperature increase is plotted as a function of dissipated power for three low flow rates, for the experiments of the previous figure.




Figure 54: The refrigerant pressure drop is plotted as a function of dissipated power for three low flow rates, for the experiments of the previous two figures.

vaporization does not take place. If the flow rate drops below this limit – or if the time-dependent device power exceeds the limit that corresponds to the existing flow rate - then refrigerant superheat or device over-temperature is readily detected. This implies that a control system which acts to keep the flow rate at or above this lower limit – with a safety margin – is sufficient to maintain stability.




Figure 55: Engineering drawing of a cooling module constructed with two aluminum Vendor 6 cold plates. The four T fittings are pressure taps.

Dual Vendor 6 aluminum cold plates

In order to begin an investigation of how to control a pumped-refrigerant system which provides direct cooling of many devices in parallel, we constructed a test fixture for cooling two mock ICs simultaneously. Figure 55 is a sketch of the module, which consists of two aluminum Vendor 6 micro-fin cold plates.

Each cold plate (referred to as CP1 and CP2) has a copper mock IC clamped to it, as in the previous single-cold-plate experiments. Pressure transducers connected to the T fittings measure




Figure 56: The junction temperature of device 1 is plotted as a function of the power it dissipates, with device 2 dissipating 500 W.

the input refrigerant pressure for each cold plate. A differential pressure transducer is used to measure the pressure drop ΔP_1 across cold plate 1. ΔP_2 is not measured.

The refrigerant flow rate was held constant at $f = 0.405 \pm 0.002$ L/min. At this flow rate, it takes approximately 1.8 kW to vaporize all the refrigerant – that power level is just above what we achieve with both heater blocks running at 100% power.

Results

Figure 56 shows the junction temperature of CP1 as a function of its device power, with the device on CP2 dissipating a constant 500 W. The dependence is linear, with a thermal resistance of approximately $R_{th} = 0.04$ K/W.

While this behavior is similar to what is seen with a single cold plate, it is not stable with respect to changes in the power dissipated by the other device. To demonstrate this, we fixed the power of CP1 at 500 W and varied the power dissipated on CP2. As shown in Figure 57 below, the refrigerant superheat in CP1 starts to increase dramatically if the power on CP2 is reduced below about 140 W:




Figure 57: The refrigerant temperature increase in CP1 is plotted as a function of the power dissipated at CP2, with CP1 power held at 500 W.

Our interpretation of these results is that, as the power dissipated on CP1 is reduced, a smaller fraction of the refrigerant fed to CP1 vaporizes, and its flow impedance therefore drops. It therefore begins to rob CP2 of refrigerant, leading to superheat in CP2.

Because of this behavior, the stability of the uncontrolled system depends on the two device powers. Figure 58 below shows this dependence, for two different refrigerant flow rates. The diagonal lines represent the latent-heat limit of the refrigerant: to the left of these lines, the total power dissipated by both devices is insufficient to cause complete vaporization. The data points are measurements of the actual stability limits. They were measured as follows: Starting in a condition in which the two device powers are equal, the power on CP1 is reduced until CP2 starts to exhibit superheat. The full curves are spline fits to the data points, and the dashed curves are the reflections of those curve about the diagonal.

The results of the previous measurements on cooling single devices indicate that active control can be applied to conserve refrigerant – i.e., to keep the flow rate to each device at or just above the level corresponding to total vaporization of the refrigerant. The results of the measurements on cooling two devices imply that a further layer of control will be required to compensate the fact that low-power devices tend to rob refrigerant from high-power devices, leading to overheat. How to deal with these effects in a system of many cooled devices will be a subject of future research.




Figure 58: Stability diagram of the two-cold-plate system. The system is stable between the solid and dashed curves. Black points and curves: $f = 0.406(2)$ L/min. Blue points and curves: $f = 0.302(1)$ L/min.

Dual-loop rack-level cooling of devices in circuit packs

As an optional adjunct to the measurements on direct refrigerant cooling of mock ICs, we have also conducted preliminary exploration of dual-loop cooling of devices inside a commercial equipment rack. The concept is to design and construct a primary closed-loop, single-phase cooling system for each shelf in a rack, with heat exhaust to the outside world accomplished by a commercial liquid-liquid heat exchanger in a secondary loop, which is cooled by either house chilled water or pumped R134a refrigerant. The design goal was to remove 3300 W per shelf while minimizing the volume occupied by the coolant circulation system.

Figure 59 below show the implementation of this concept. The equipment chosen was from an ALU product line. The rack has 16 full-height slots for circuit packs, mounted with 22.5-mm spacing. As shown in the right-hand photo, this spacing is just sufficient to allow mounting and plumbing of a Vendor 5 gull-wing cold plate on a blank circuit pack with enough extra space for a 12.7-mm-high mock IC and clamping plate. The front panel of this circuit pack has NPT ports which are fitted with dripless quick-connect fittings for attachment of coolant hoses. These hoses are seen in the left-hand photo, in which two circuit packs have a total of three liquid-cooled devices installed. The hoses are connected to the circulation system via overhead input and output manifolds, and these in turn are connected to the coolant pump, which is mounted on the back plate of the rack.

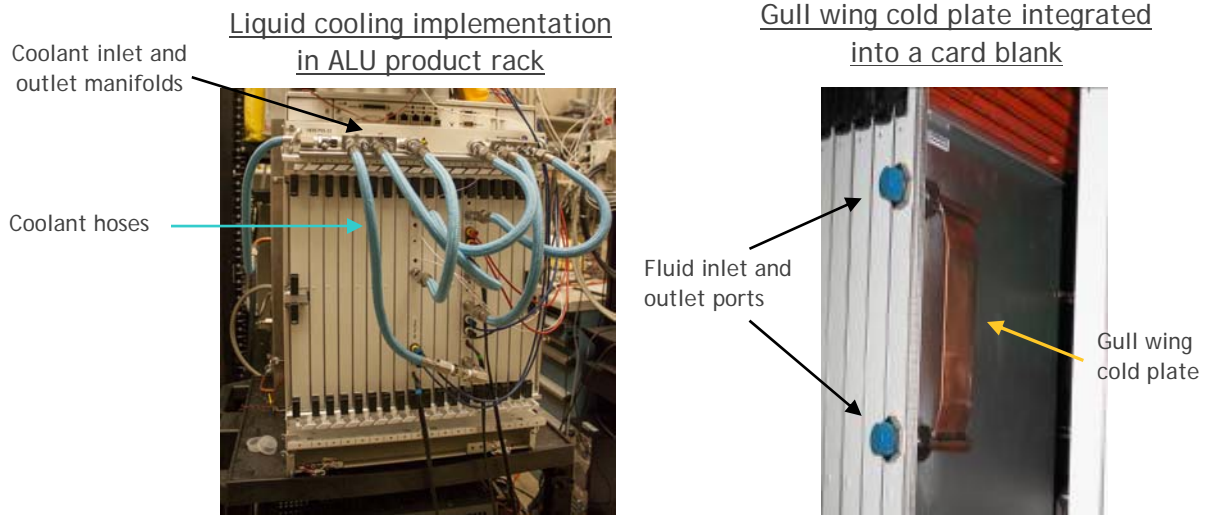


Figure 59: Photos of the dual-loop cooling system. Left: front view of equipment rack showing blue coolant hoses. The coolant pump and reservoir are mounted on the back of the rack. Right: close-up of a Vendor 5 gull-wing cold plate mounted in one circuit pack.

The test system includes one circuit pack with a single Vendor 5 gull-wing cold plate (blue curves in data plots) and a second pack with one Vendor 5 cold plate (red curves) and one Vendor 6 cold plate (green). The mock IC/cold-plate assemblies are wrapped in heavy insulation – this allows us to infer the coolant flow rate through the blue cold plate from the known coolant specific heat and coolant temperature-rise measurements.

The coolant pump is a Micropump model CA centrifugal pump, with a zero-head flow rate of 21 L/min of water at full speed. The total volume occupied by the pump and the coolant reservoirs is approximately 3500 cm³.

The heat exchanger chosen for this system is an Alfa-Laval type CB27-24H brazed-plate heat exchanger, with a design capacity of 40kW using either water or R134a as a secondary coolant. The heat-exchanger volume is approximately 5400 cm³.

The coolant chosen for this application was 3M Novec 7500. Its viscosity and specific heat at 25 °C are 0.77 cSt and 1100 J/kg-K, respectively.

Figure 60 below shows some representative results obtained with this test system using house chilled water as a secondary coolant. In the left-hand graph, the three mock ICs were operated simultaneously, with their voltages set to keep the measured junction temperatures approximately equal. The two Vendor 5 cold plates respond slightly differently, presumably because their flow impedances are slightly different. As before, the Vendor 6 cold plate exhibits a thermal resistance that is about half as much as that of the Vendor 5 unit. Both impedances are somewhat larger than measured with direct refrigerant cooling (see Table 9 below for

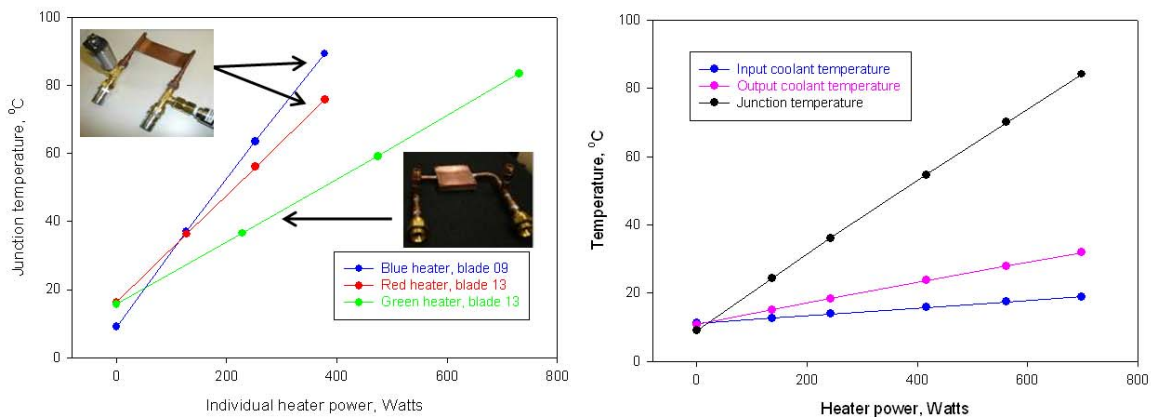


Figure 60: Left: junction temperatures are plotted as functions of individual heater powers for a run in which the applied voltages were set to keep those junction temperatures approximately equal. Right: Blue junction temperature and input/output coolant temperatures are plotted vs. heater power for a run in which this device alone was cooled and powered. Full pump speed was used for both runs.

specifics) – this is due to incomplete cooling by the liquid-liquid heat exchanger, which results in an increase in coolant temperature at the input of the cold plates as their power is increased. In this experiment, the total power level that brings all three junction temperatures simultaneously to 85 °C is approximately 1540 W, well below the design value of 3300 W.

In order to better probe the ultimate performance of this system, we increased the bore of the connectors on the blue cold plate to minimize their flow impedance, and we disconnected the other two heaters. In the right-hand graph of the figure above, we plot the blue junction temperature (black curve) and coolant input ((blue curve) and output (red curve) temperatures. The thermal resistance of this cold plate has been reduced to 0.13 K/W, close to the value of 0.10 K/W seen using direct refrigerant cooling at low flow rates. When this experiment was repeated using R134a as a secondary coolant, the thermal resistance was further reduced to 0.11 K/W. This remaining discrepancy is explained by the coolant heating, as reflected in the red and blue curves.

Summary of results of liquid cooling experiments

Table 9 below gives a numerical summary of the results of our experiments on liquid cooling, including the thermal backplane measurements in Task 3. The first point of comparison between different cooling solutions is the volume required for the hardware. The cooling modules we have tested all require some kind of plumbing, which occupies volume to the sides of the device being cooled as well as above it. The baseline solution, and air-cooled heat sink with integral fan, does not physically occupy space other than its own volume, but it casts a “thermal shadow” in the surrounding airspace that affects cooling of downstream devices. In addition, external fans occupy space that is not included in the accounting in the table. Thus, we think it is reasonable to use the total height of all the competing cooling solutions as measures of their volumes. In the

HT module	height above IC, including clamp	coolant	R _{th}	Q [W], mock IC at 85C	notes
Baseline					
air cooled HS	55.9 mm	air	0.32K/W	~215W	extrapolated from lower power
Thermal backplane configuration					
Vendor 4 VC	15.5 mm	water/R134a	0.13K/W	200W	very heavy module
Vendor 3 round 8mm HP	21.6 mm	water/R134a	0.25K/W	~260W	extrapolated from lower power
Vendor 3 flat 8mm HP	11.2 mm	water/R134a	0.86K/W	105W	hybrid-wick HP should do better
Vendor 1 flat 8mm HP	12.7 mm	water/R134a	2.7K/W	50W	
Direct cooling configuration					
Vendor 5 gull wing CP	9.4 mm	R134a	0.10 K/W	560W	flow rate 0.265 L/min; Q(85) limited by flow rate
Vendor 6 micro-fin CP	19.0 mm	R134a	~0.05K/W	885W	flow rate 0.210 L/min; Q(85) limited by flow rate
Vendor 5 gull wing CP	9.4 mm	R134a	0.08 K/W	903W	flow rate >1.1 L/min, no limit
Vendor 6 micro-fin CP	19.0 mm	R134a	0.05K/W	1625W	flow rate >1.1 L/min, no limit
Dual-loop rack configuration					
Vendor 5 gull wing C	9.4 mm	Novec 7500/water	0.16 to 0.22K/W	350 to 435W	at max flow rate, with 2 other heaters connected
Vendor 5 gull wing CP	9.4 mm	Novec 7500/water	0.13K/W	718W	at max flow rate, low-impedance plumbing; no other heaters connected
Vendor 5 gull wing CP	9.4 mm	Novec 7500/R134a	0.11K/W	730W	at max flow rate, low-impedance plumbing; no other heaters connected
Vendor 6 micro-fin CP	19.0 mm	Novec 7500/water	0.09K/W	755W	at max flow rate, with 2 other heaters connected
Vendor 6 micro-fin CP	19.0 mm	Novec 7500/water	0.07K/W	1038W	at max flow rate, no other heaters connected

Table 9: Summary of results of liquid cooling experiments.

case of the liquid-cooling solutions, we have to include clamping plates in these heights, and these have thicknesses in the range 7 – 9 mm. All of our liquid cooling solutions occupy heights that are much smaller than the fan/heat/sink combination we used as a baseline. This volume comparison is very imprecise, but our impression is that the liquid cooling solutions we have investigated use volume with comparable or superior efficiency to that required for air cooling.

The coolant listed for the thermal-backplane thermal modules is listed as water/R134a. This is meant to indicate that the condenser ends were cooled by pumped R134a refrigerant, while the working fluid inside the modules is water. The best performance among these devices was given by the Vendor 4 vapor chamber. It exhibits a very low thermal resistance (0.13 K/W) but is very heavy. The highest device powers that can be mediated by these solutions lie in the range of 200 to 260 W – comparable but a bit superior to what can be achieved using air cooling. These power limits are fundamental in the case of heat-pipe-based thermal solutions – they are imposed by wick dryout and cannot be overcome without using larger units.

Direct refrigerant cooling using cold plates such as the Vendor 6 and Vendor 5 modules is extremely effective. In both cases, we showed that the refrigerant flow rate can be safely adjusted so that the device power just evaporates all the refrigerant. With this control implemented, the maximum power that can be safely dissipated is limited only by the thermal resistance of the cold plate. Even for the Vendor 5 cold plate, whose resistance is < 0.10 K/W, this power level is conservatively estimated at $Q_{\max} = \Delta T/R_{\text{th}} = 60 \text{ }^\circ\text{C}/(0.1\text{K/W}) = 600 \text{ W}$, assuming a refrigerant temperature of $25 \text{ }^\circ\text{C}$ and a maximum junction temperature of $85 \text{ }^\circ\text{C}$. This is more than enough headroom for devices in telecommunications equipment.

The dual-loop cooling system does not yet live up to its design specifications. As shown by the Q values in the table, which correspond to the case where the mock IC junction temperature is $85 \text{ }^\circ\text{C}$, there is more than enough cooling capacity in this system for cooling individual devices, whose power levels don't exceed a few hundred Watts. However, the total capacity of the system as tested needs to be improved. This should be feasible with a higher-capacity pump and liquid-liquid heat exchanger.

Section 4.5: Rack-level test apparatus build and prove-in (Task 5)

The goal of this task was to build an experimental apparatus in Alcatel-Lucent's Murray Hill, NJ location to enable testing and evaluation of technologies developed for improved cooling at the device, board, shelf, and rack-level.

Task activities included the following: i) installation of a refrigerant pump and associated supply and return manifolding; ii) installation of an equipment rack and instrumentation to support testing of heat exchangers (Sections 4.1 and 4.8), thermal backplane (Section 4.3), device-level liquid cooling (Section 4.4) and hoses and connectors (Section 4.7). This included a data acquisition system to monitor air and refrigerant temperatures, refrigerant pressures, refrigerant and chilled water flow rates and air velocities, along with the setup of a dedicated wind tunnel within one shelf of the equipment rack for heat exchanger evaluation; and iii) installation of a second equipment rack with an additional 18 kW (3 x 6 kW) of heat load for more realistic system-level load conditions. Figure 61 shows a photograph of the laboratory setup at Alcatel-Lucent's Murray Hill, NJ location.

The task activity was completed ahead of schedule in 4Q10.




Figure 61: Laboratory setup at Alcatel-Lucent's Murray Hill, NJ facility.

Section 4.6: Optimization of refrigerant-circulation system (Task 6)

The goal of this task was to design and build the hardware and control infrastructure to provide a more energy-efficient, reliable, and cost-effective refrigerant-circulation system. Guidelines were also to be developed for supporting a range of power levels, so as to match cooling capacity with end user heat load, and interfacing with a compressorized secondary cooling loop, to accommodate sites without building chilled water.



Figure 62: (left) chilled water pumping station (gray box) connected to 18 kW heat load (black box); and (right) portable chilled water plant.

During the early stages of the project Alcatel-Lucent downscoped Modine's responsibility for this task activity as ALU became aware of a prototype pumping station that had already been developed by another manufacturer. This included moving testing of the pumping station from Modine to Alcatel-Lucent. The immediate availability of the prototype pumping station reduced program risk by realizing delivery of the pump one year ahead of schedule. Alcatel-Lucent apprised DOE of this development promptly and in a timely fashion, and requested a budget modification in early 4Q10 to reflect the downscoping of Modine's activities, acquiring of the prototype pump and movement of pump testing activities to Alcatel-Lucent. Note that this rebudgeting did not change planned task activities, though it required, and was granted, approval from the DOE Contracting Officer.

The remainder of this chapter details task activities.

Subsection 4.6.1: Chilled-water pumping station

Testing and evaluation of the initial chilled-water pump prototype was performed in Alcatel-Lucent's Naperville, IL location. Figure 62 shows a photograph of the pump attached to a bank of three 6 KW heaters to emulate heat load on the pumping station; the portable chilled water station that the pump used for heat rejection is also shown in the figure. Testing of the pump included, among other things, measurement of refrigerant and air temperatures as well as refrigerant pressures at key locations within the test setup (pump, manifolds, coils) for heat loads ranging from 0 kW to 18 kW.

Feedback based on the testing of the initial chilled-water pump prototype was provided to the pump vendor through a revised set of commercial pump design requirements; this effort was largely lead by Alcatel-Lucent's Services Business Division as part of Commercialization activities (see Section 4.9) with extensive interaction with the pump vendor. Details on the specifics findings of the pump testing as well as the revised set of commercial pump design requirements are not provided here due to the proprietary nature of the material. The vendor made subsequent modifications to the pump reflecting the new commercial pump requirements. These included a completely redesigned pump controller.

The redesigned pump prototype was subsequently tested by Alcatel-Lucent. This lead to the creation of a final set of commercial pump design requirements, again, with this activity being lead by ALU's Services BD (see Section 4.9) with extensive interaction with the pump vendor and resulting in a total of approximately 90 necessary features and 20 desirable features being specified in the design requirements. Completion of the final set of commercial pump design requirements completed activity related to this subtask.

Subsection 4.6.2: Compressorized / Condenserized Pump Design

Many ICT facilities do not have the infrastructure and facilities to support a chilled-water plant. Therefore, there is a strong commercial need for a pump design that is capable of accommodating such sites. A pumping unit that interfaces with a compressorized / condenserized refrigeration loop is thus desired.

The focus of this subtask was on the development of the commercial design requirements for such a compressorized / condenserized pump. In particular, since many compressorized / condenserized loops would be a separate system, the focus was on specifying the design requirements associated with the interface between the pump and this external heat rejection loop. This task was lead by Alcatel-Lucent's Services Business Division, with input from the pump vendor as well as Alcatel-Lucent's Bell Labs CTO organization. The design requirements for the compressorized / condenserized pump were completed in 3Q11. Note that specific details on the design requirements are not provided here due to the proprietary nature of the material. Completion of the final set of commercial pump design requirements completed activity related to this subtask.

Section 4.7: Develop and research new connector technologies (Task 7)

Subsection 4.7.1: Optimization of Hose Assemblies

The goal of this task was to design and optimize the cost and performance of the hose assemblies, both for the stainless steel hose assembly and the non-conductive polymer-lined hose assembly, using existing materials and manufacturing processes. A major component of the cost of the hose assemblies is associated with the stainless steel tubing and fittings that attach to the quick-connect end fittings on the heat exchanger and manifold sides of the hose assembly.

Brass fittings and copper tubes were investigated as possible replacement materials for stainless steel for the non-conductive polymer-lined hose assembly (see Subsection 4.7.2) due to their potential for lower material and manufacturing cost. It was found that brass fittings would not hold pressure after thermal cycling for the non-conductive hose assembly. Testing was also performed that showed that copper tubes would be too soft for the application, with a high probability of developing leaks during the assembly process and being damaged during installation and use.

With regards to cost reduction of the stainless steel hose assembly, although there was potential for the brass fitting to be used and still maintain pressure after thermal cycling, since the brass fittings didn't work with the polymer-lined hose assembly there would not have been an overall cost savings in using the brass fitting, as this would have required carrying two different part numbers as well as maintaining two different manufacturing processes to build the hose assemblies. It was thus decided that the stainless steel material for the fittings and tubing was the optimal material with respect to performance, reliability and cost.

In the course of the work four additional potential suppliers were identified for the stainless steel fittings and tubes.

Subsection 4.7.2: Research and Develop New Hose Materials

The goal of this task was to create a new polymer-lined hose assembly as a replacement to the current metal hose assembly. The polymer-lined hose and non-conductive braid material would make the hose completely non-conductive, which was a requirement for a customer application that was not met by the metal host assembly. Also the polymer-lined hose manufacturing process is less labor intensive than the metal hose assembly process due to its use of a crimped, as opposed to welded, joint design, and hence offers opportunities for cost reduction.

The polymer-lined hose assembly consists of the following components: i) an interior polymer tube that confines the refrigerant. The polymer tube is desired to have low permeability to refrigerant, good pressure rating, able to hold a partial vacuum and be readily available from suppliers; ii) an exterior polymer braid material that provides mechanical protection to the polymer tube and additional pressure margin. The polymer braid material is also desired to be non-flammable, low-cost and readily available; and iii) metal tubing and metal fittings that connect to the ends of the braided polymer tube and allow attachment of quick-connect fittings.

Tube material	Comments
PTFE	The effusion rate of refrigerant through the material was too high. Testing confirmed that refrigerant readily permeated the tube.
PVDF	This material had the lowest effusion rate compare to PTFE and ETFE. However, there were only limited suppliers that would make the tubing and it was long lead time item and hence not readily available.
ETFE	This material had a much lower effusion rate than PTFE. It was also a material US Hose had significant experience with and would be readily able to find suppliers for.

Table 10: Non-conductive polymer tubing material evaluated for the polymer-lined hose assembly.

The non-conductive hose assembly effort was completed in 3Q11; details are presented in Subsection 4.7.2.1. However, in 4Q11 it was determined in field application that the non-conductive hose assembly would build up static electricity that would result in electro-static discharge (ESD) and create pin hole leaks in the hose. This development required re-opening task activities, including requesting and obtaining a no-cost time extension from DOE to complete work on this task. Details on the development of the static-dissipative hose assembly that resolved this ESD issue are presented in Subsection 4.7.2.2.

Subsection 4.7.2.1: Non-conductive hose assembly

Evaluation of different polymer hose tubing materials

Table 10 lists the particular non-conductive polymer tubing that was evaluated. Of particular interest were materials with a good pressure rating, low permeability to refrigerant, able to hold a partial vacuum and readily-available commercially. Effusion rates were estimated either by




Figure 63: Still frame from a video of an underwater helium leak test used for evaluating permeability characteristics of polymer-lined hose assemblies.

Material and manufacturer	Comments
Flex-Guard PE (PET monofilaments) from Alta Technologies	The braid did not meet the flammability requirement.
Flex-Guard FPE (flame-retardant polyester monofilaments) from Alta Technologies	The braid did not meet the pressure requirement.
Expandable polyester monofilament sleeving from RDS Wire and Cable	The supplier would not directly braid to the tubing to meet US Hose's manufacturing process requirements.
EX920FR Expandable Monofilament Sleeving from Jenisco	The supplier would not directly braid to the tubing to meet US Hose's manufacturing process requirements.
Polyester fabric from Unigasket	The braid meets pressure rating and flammability requirements.

Table 11: Polymer braid materials evaluated for polymer-lined hose assembly.

performing underwater helium leak testing, underwater nitrogen leak testing, pressure testing or from manufacturer's data. For example, Figure 63 shows a still frame from a video of an underwater helium leak test of a polymer lined hose samples. Arrows point to helium bubbles emanating from the hose sample, indicating that the hose sample has undesirable effusion with respect to helium, which is a very stringent test due to the small molecular size of the helium molecule. ETFE tubing with an OD (outside diameter) of 10 mm and wall thickness of 1 mm was chosen as the preferred material due to its very low permeability to refrigerant, good pressure rating, ability to hold a partial vacuum and ready availability.

Evaluation of different polymer braid materials

Table 11 lists the different polymer braid materials that were evaluated. Of particular interest were materials that met requirements for pressure rating and flammability, were able to be braided on the ETFE tubing and were readily-available commercially. The polyester fabric from Unigasket was chosen as the preferred material as it was the only material to meet the aforementioned requirements.

Crimp fittings and tube materials

The polymer-lined hose assembly utilizes a crimp insert to connect the braided polymer tube to either the end fitting that connects to the manifold or the S- or U-shaped tubes that connect to the heat exchanger. Two types of materials were evaluated, namely brass and stainless steel, with the brass offering potential for cost reduction. It was found that the brass crimp fittings would not hold pressure after thermal cycling. It was thus decided to use the stainless steel crimp inserts.

Copper tubes were also tested as a possible replacement material for the stainless steel S- and U-shaped tubes. A strength test of the copper tube showed that the material was much softer relative to the stainless steel tubes, and hence could be prone to developing leaks during the

assembly process and being damaged during installation and use. It was determined that to proceed further with this material would require re-certification of the material, which would incur high cost and additional schedule delay. US Hose and Alcatel-Lucent thus jointly decided to use stainless steel because, although it had incrementally higher cost, it would have much better reliability characteristics.

Non-conductive polymer-lined hose assembly

The above efforts resulted in a completely non-conductive polymer-lined hose assembly being developed for this task. The hose subsequently passed NEBS and IEC compliance testing, certifying it for use in telecommunications applications; compliance testing included a number of NEBS (GR-63-Core) tests for temperature and humidity exposure, fire resistance, earthquake and office vibration and an IEC test for di-electric strength (voltage breakdown).

However, in field application it was determined that the hose would develop a significant build-up of static electricity that manifested as electro static discharge (ESD) which created pin hole leaks in the hose. It was postulated that the static build-up occurred on the interior surface of the polymer-lined hose due to tribologic charging from the flow of the dielectric refrigerant vapor passed the electrically insulating polymer material.

The ESD issue for the non-conductive hose came to Alcatel-Lucent's attention in late 4Q11. ALU immediately apprised DOE of this issue. It became readily apparent that, due to potential issues with material availability and lead times, coupled with the need for detailed functional and compliance testing, addressing the unanticipated hose performance issue would continue past March 30, 2012, the original end date of the project. As such, a no-cost time extension for a new project end date of Aug. 30, 2012 was formally requested from DOE on Feb. 17, 2012 and was granted by DOE on Feb. 29, 2012. Note that, due to the pressing need for realizing a hose assembly meeting commercial requirements, as well as to meet the new project end date, a number of design, testing and evaluation activities were moved from USHose to Alcatel-Lucent.

Subsection 4.7.2.2 below details the project activities undertaken to resolve the ESD issue with the non-conductive polymer-lined hose assembly.

Subsection 4.7.2.2: Static-dissipative hose assembly

ALU identified multiple mitigation strategies to address this hose performance issue. More specifically, we identified and initiated activity on eight potential solutions simultaneously (some with multiple variations). The reasons for the aggressive parallel resolution approach were the project time constraints and relatively low costs. Our mitigation planning determined that the lead times required to prove the different strategies were a significant risk if a serial approach to solving this problem was used. Specifically, lead times for some materials exceeded a month, so waiting to order material would have added excessive schedule risk. These eight efforts can be divided into two groups: those that utilize technologies already developed under this grant and

those that utilize newly identified alternative materials. Five of the eight approaches were with the current subcontractor, USHose, who actively supported the initiatives. The remaining three approaches utilized alternate hose vendors.

Testing of the different hose options was done in three phases, namely first-pass testing, functional testing and compliance testing. First-pass testing included performing leak testing (under pressure and under vacuum) and voltage breakdown testing of the initial round of samples. Functional testing included leak testing (under pressure and under vacuum), static build-up and refrigerant flow rate testing of samples that had passed first-pass testing. Compliance testing included a number of NEBS (GR-63-Core) tests for temperature and humidity exposure, fire resistance, earthquake and office vibration and an IEC test for di-electric strength (voltage breakdown) for samples passing functional testing. Note that for some hose options listed below only a subset of the first-pass or functional testing was performed, which allowed a determination to be made regarding the viability of that particular hose assembly.

1. Option 1 (USHose and ALU): Insert a static-dissipative PTFE spiral wrap within the existing non-conductive ETFE hose. The spiral wrap material exists but tools needed to be developed to pull the liner through the hose. This approach would have taken advantage of the current hose material, which has excellent permeability characteristics and already passes stringent permeability tests. Preliminary functional testing showed this approach is viable and does not affect refrigerant flow and warranted more detailed functional testing.

Recommendation: This option was not down-selected for compliance testing for the following reasons: i) more detailed functional testing showed that the liner did have an adverse effect on refrigerant flow; ii) there was concern that the PTFE static-dissipative filler material (carbon black) could leach into refrigerant which could damage or clog heat exchangers, fittings and the refrigerant pump; iii) concern of the repeatability of the manufacturing process, which could lead to the insert not making electric contact with the metal hose ends and result in static build up; and iv) galvanic corrosion concerns due to the static-dissipative nature of the material.

2. Option 2 (USHose and ALU): Insert a static-dissipative PTFE tube liner within the existing non-conductive ETFE hose. The liner material exists but tools needed to be developed to pull the liner through the hose. This approach had similarities to Option 1 above in that it takes advantage of the current hose material and its excellent permeability characteristics. However, it differs in that the liner material would have been crimped within the hose end fittings, which makes for a more secure mechanical connection for the static-dissipative path than the spiral wrap in Option 1. Preliminary functional testing showed that the liner does not affect refrigerant flow and warranted more detailed functional testing.

Recommendation - This option was not down-selected for compliance testing due to the following concerns: i) manufacturing tolerances and development of the tool required to pull the liner through the hose assembly; ii) the PTFE liner has different permeability characteristics with respect to refrigerant than the ETFE hose. Risks associated with how these different permeabilities could affect performance, for example, if refrigerant were to become trapped between the liner and hose tube, could not be evaluated within the program scope and timeline; and iii) there was concern that the PTFE static-dissipative filler material (carbon black) could leach into refrigerant which could damage or clog heat exchangers, fittings and the refrigerant pump.

3. Option 3 (USHose and ALU): Use of a plastic (PEEK) bushing located at the manifold end of the stainless steel hose assembly. The stainless steel hose exists, so the focus of the effort was on the development of the plastic bushing. Initial experiments performed in March showed a reduced and unacceptable refrigerant flow with the original PEEK bushing design. A larger internal diameter bushing was created to mitigate this issue, with more functional testing in May of 2012 showing a small but acceptable reduction in refrigerant flow rate.

Recommendation: This hose was subsequently down-selected for compliance testing as it offered the best options with respect to reliability, performance and cost. The hose passed NEBS and IEC compliance testing, certifying it for use in telecommunications applications. System-level testing of the hose also showed no static electricity build-up.

4. Option 4 (USHose and ALU): Develop a static-dissipative version of the ETFE tube material used in the existing hose assembly. This looked to be the best option with respect to static-dissipative properties, manufacturing costs and permeability to refrigerant. However, it was determined that the ETFE tube supplier would need to develop a new manufacturing process to incorporate static-dissipative material (carbon black) either within the ETFE matrix, or in a thin layer adjacent to the inner surface of the ETFE tube.

Recommendation: This option was not selected for further study for the following reasons: i) the time frame for development of a new manufacturing process was unknown as well as its associated effect on cost; ii) it was not clear what effect incorporation of carbon black would have on the permeability of the ETFE material to refrigerant; and iii) there was concern that the static-dissipative filler material (carbon black) could leach into refrigerant, which could damage or clog heat exchangers, fittings and the refrigerant pump.

5. Option 5 (USHose and ALU): Use of a static dissipative wrap applied to the exterior of the existing ETFE braided hose. This option looked optimal with respect to maintaining the excellent refrigerant flow properties of the ETFE tube liner and re-using existing non-conductive hose inventory. Preliminary testing completed in April showed that the static dissipative wrap on the exterior of the hose could be effective at dissipating static build-

up and warranted more detailed functional testing. However, it was discovered that the quality of wrap, namely the tightness, could lead to potential ESD build-up. Additionally, the static-dissipative wrap was not flame retardant, and an electrically-insulating electrical tape overwrap needed to be applied to make the assembly pass flammability testing.

Recommendation: This option was not selected for further study as the electrically-insulating overwrap would develop static build-up that would still arc to ground, albeit at a much lower frequency than the original hose.

6. Option 6 (Vendor A and ALU): Use of a polyester-braided corrugated PTFE hose advertised to be static dissipative from a new hose supplier. Hoses were ordered and basic testing with respect to hose performance were performed.

Recommendation: This hose was not selected for further testing as: i) there was concern that the static dissipative filler material (carbon black) could leach into refrigerant, which could damage or clog heat exchangers, fittings and the refrigerant pump; ii) there was concern regarding the permeability of the PTFE material to refrigerant, in particular, the manufacturer would not warranty that the PTFE had desired permeability characteristics; and iii) basic functional testing showed that the corrugated tubing used in the hose assembly resulted in significant reduction in refrigerant flow rate.

7. Option 7 (Vendor B and ALU): Use of a fiberglass-braided PTFE hose advertised to be static dissipative from a new hose supplier. Hoses were ordered in February for performing basic first-pass and functional testing, however, there was a significant lead time of 8 to 10 weeks for shipment.

Recommendation: This hose was not selected for further testing as: i) there was concern that the static dissipative filler material (carbon black) could leach into refrigerant, which could damage or clog heat exchangers, fittings and the refrigerant pump; ii) there was concern regarding the permeability of the PTFE material to refrigerant, in particular, the manufacturer would not warranty that the PTFE had desired permeability characteristics. As this was determined prior to order completion, the hose order with Vendor B was subsequently cancelled.

8. Option 8 (Vendor C and ALU): Use of a thick-walled, static-dissipative PFA tube. This option looked attractive as the thick-walled tube had excellent pressure characteristics.

Recommendation: Ruled out as a viable option because of concerns regarding the permeability of the PFA material to refrigerant, in particular, the manufacturer would not warranty that the PFA had desired permeability characteristics.

In summary, Option 3, which used a plastic (PEEK) bushing located at the manifold end of the stainless steel hose assembly, was eventually down-selected for, and successfully passed NEBS

Fitting #	Vendor #	Features	Cost
1	F1	The fitting is a wrench on fitting and will not work for this tight space application.	NA
2	F2	This fitting has a flush-face and is non-spill.	Too costly
3	F2	The fitting is hand-tightening and the required wing nut is too big to fit in the tight space. The fitting is also minimum spill, not the preferred non-spill type.	NA
4	F3	Smooth push-to-connect is ideal for hand connection. Light-weight and dry-break with minimal air inclusion are also desirable features.	Too costly
5	F3	Smooth push-to-connect is ideal for hand connection. Light-weight and dry-break with minimal air inclusion are also desirable features.	Too costly
6	F3	Smooth push-to-connect is ideal for hand connection. Flush-face and dry-break with minimal air inclusion are also desirable features. Available in different metals with different coatings and seal materials.	Too costly
7	F3	Smooth push-to-connect is ideal for hand connection. Flush-face and dry-break with minimal air inclusion are also desirable features.	Too costly
8	F4	Clean-break, quick-release	Too costly
9	F5	This is a twist on fitting, but it takes too much force to connect.	NA
10	F6	This is the closest possible fitting for the application. The fitting was a push-to-connect design that is non-spill and has no air inclusion on connection. However, the force required for connection was too large and would have required a potential redesign.	Lower than the current fitting, but not low enough to justify replacement.

Table 12: List of evaluated potential replacement fittings and their associated features and cost.

compliance testing, certifying it for use in telecommunications applications. System-level testing of the hose also showed no static electricity build-up.

Section 4.7.3: Source New or Similar Wrench-Less Fitting for Metal Hose Assembly

The goal of this task was to source a cheaper, more reliable, and easier-to-connect fitting than the current fitting. A number of potential designs from 6 different fitting manufacturers were considered, with push-to-connect and twist-on styles being the most desirable for meeting ease-

of-installation requirements. Equally important were the ability of the fitting to not leak refrigerant or entrain air on connection and disconnection and the cost of the fitting. During the course of the project it was also determined that backward-compatibility was a highly desirable feature to have from a maintenance, inventory and supply chain perspective, which conflicts with having multiple potential fitting designs.

Task efforts included sourcing sample fittings from the manufacturers and determining volume pricing for promising fitting designs. Table 12 details the different fittings that were investigated. Note that particular details on the fitting manufacturer are not supplied for proprietary and confidentiality reasons.

The result of the task was that, of the 10 fittings that were considered, only Fitting #10 had promise as a replacement to the current fitting due to its features and lower cost. However, this fitting was not chosen as a replacement for the following reasons: i) the force to connect was deemed too large and would have required a redesign, which could have adversely affected the non-spill and no air inclusion characteristics of the fitting; ii) the cost was not substantially lower than the current fitting so as to justify the overhead of maintaining two different fitting part numbers (due to the lack of backward-compatibility of any new fitting design with the current fitting) within the commercial product offering, e.g., the cost of maintaining two fitting part numbers exceeded any savings associated with the lower fitting cost; iii) cost reduction efforts reduced the cost of the current fitting; and iv) a less expensive backward-compatible version of the current fitting (wrench-type connection) for the manifold side of the hose assemblies was identified.

In summary, the recommendation of the task was that the current fitting design is the best option relative to other available fittings with respect to features, performance and cost.

Subsection 4.7.4: Develop and Source New End Fittings for Polymer-Lined Hose Assembly

The goal of this task was to design, validate and source new end fittings for the polymer-lined hose assembly, including a wrench-type design that attaches to the manifold portion of the system, and a wrench-less type design that attaches to the shelf-level heat exchanger.

This task activity was strongly dependent on the results of the tasks detailed in Sections 4.7.1, 4.7.2 and 4.7.3, and was overcome by the events associated with these activities. In particular, it was determined that a critical design criterion is “backward compatibility”, and that any new end fittings would have to be substantially lower in cost to justify the overhead of maintaining two different fitting part numbers (due to the lack of backward-compatibility of any new fitting design with the current fitting).

In summary, the results of tasks detailed in Sections 4.7.1, 4.7.2 and 4.7.3 obviated the need to complete this task activity, as the current fitting design is optimal with respect to features, performance and cost.




Figure 64: Photograph of wind tunnel setup for evaluating performance of shelf-level heat exchangers and hose assemblies at ALU facility in Murray Hill, NJ: (left) inlet and side of test setup; and (right) rear of test test setup.

Section 4.8: System-level testing (Task 8)

The goal of this task was to perform system-level testing will be used to understand how new component designs impact two key areas: 1) system performance under typical cooling operation; and 2) how the system responds to cooling failures.

Subsection 4.8.1: Testing of new component designs prior to full system-level evaluation

The focus of this task was to perform testing of shelf-level heat exchangers, hose assemblies and pumping stations prior to evaluation on operational equipment running live traffic. This allowed evaluation of the performance of different designs and their relative benefits, and provided valuable feedback to component suppliers to improve and enhance the performance of different designs.

Shelf-level heat exchanger testing

Figure 64 shows a photograph of a wind tunnel setup for evaluating the performance of shelf-level heat exchangers supplied by Modine. A fan located at the inlet of the wind tunnel drives air through the test apparatus, with the heat exchanger under test located at the exit of the wind tunnel. An Omega resistance heater capable of providing up to 6 kW of heat load was located within the wind tunnel duct and used to heat the incoming air. Sensors were used to measure the speed and temperature of the air at various locations entering and exiting the coil, as well as the temperature and pressure of the refrigerant entering and exiting the evaporator coil of the heat exchanger. Variation of the fan speed allowed control of the temperature of the incoming air and the face velocity of the air impacting the coil, although not independently. The setup also incorporated an ultrasonic flow meter to measure the flow rate of refrigerant to the evaporator coil.




Figure 65: (left) Wind tunnel setup used for evaluating performance of hose assemblies; and (right) test setup for evaluating ability of hose assembly to maintain a vacuum.

The experimental results validated a number of key performance factors for the coils, including: i) refrigerant pressure drop across the next-generation coils was lower than the baseline coil. This was expected as the next-generation coils have a larger free-flow area relative to the baseline coil; and ii) expected trends for coil overall thermal performance appeared to qualitatively follow the predictions of Modine’s modeling and experimental results. However, due to a number of factors, including unstable chilled water supply temperature; non-uniform air temperature into the coil; non-uniform air velocity into the heater and into the coil; a refrigerant saturation temperature of 18 to 20°C, which is significantly larger than the desired 10°C condition, and a variable refrigerant flow rate, it was not possible to provide a quantitative comparison with Modine’s modeling and experimental results.

These results, coupled with Modine’s modeling and experimental efforts, allowed for the recommendation of two preferred shelf-level heat exchanger designs, namely 1A and 1B, to be chosen for system-level testing on operational equipment. The recommendation was based on performance and near-term manufacture-ability.

Hose assembly testing

This effort was done in conjunction with USHose’s efforts to develop a non-conductive hose assembly. Note that this effort progressed in two phases (see Subsection 4.7.2). In the first phase an ETFE-based hose assembly was developed and commercialized. However, in late 4Q11 ALU became aware of an electro-static discharge (ESD) issue with the hose where static electricity would build up and subsequently discharge from the hose, creating pin-hole leaks in the hose.



Figure 66: Test fixtures developed for evaluating the mechanical integrity of the non-conductive hose assembly: (left) random vibration test fixture; and (right) drop test fixture.

This required a second phase to develop a replacement non-conductive hose that did not suffer from the ESD issue.

In both phases, the hose assembly was required to meet a number of stringent performance and reliability requirements. Efforts that focused on developing a polymer-lined hose assembly required that it was impermeable to refrigerant and able to hold a vacuum. USHose performed underwater helium leak testing and underwater nitrogen testing of prospective hose materials (see Subsection 4.7.2); helium and nitrogen are smaller molecules relative to the refrigerant R134a, so that hose materials that are impermeable to these molecules would be expected to also be impermeable to R134a. Alcatel-Lucent tested prospective hose materials to ensure they would be able to hold a vacuum of 150 microns for one hour, not leak refrigerant and would not impede the flow of refrigerant, which could deleteriously impact cooling performance. Figure 65 shows photographs of test apparatuses for evaluating whether a new hose assembly impacted cooling performance and for testing the ability of a hose material to maintain a vacuum; note that the wind tunnel shown in Figure 64 was also used for evaluating whether a hose assembly impacted cooling performance. Compliance testing included a number of NEBS (GR-63-Core) tests for temperature and humidity exposure, fire resistance, earthquake and office vibration and an IEC test for di-electric strength (voltage breakdown).

Note that, in the second hose development phase, the build-up of static electricity was a new concern. An ESD meter that measured voltages up to 25 kV was used to evaluate whether any of the different approaches detailed in Subsection 4.7.2 were adequately addressing this issue. Additionally, the preferred hose assembly that was eventually down-selected for full compliance testing was based on the current stainless steel hose assembly and incorporating a plastic bushing on the manifold end of the hose assembly for electrical isolation. Since there was concern regarding the mechanical integrity of the plastic bushing, two custom tests were developed to evaluate the bushing's performance with respect to random vibrations and being subject to large



Figure 67: 17-rack 100 kW server farm (Server Farm 1) located in Plano, TX.

forces associated with sudden acceleration; Figure 66 shows test apparatuses developed for evaluating each of these effects.

Refrigerant pump testing

See Section 4.6.

Test #	Test details
1	First pass testing exploring the initial capabilities of early designs.
2	Investigation of effects of changes in chilled water flow rate and changes in humidity.
3	Testing of distribution manifolds.
4	Experimentation exploring failure modes and how the system responds.
5	Examination of effects of low chilled water pressure differential.
6	Experimental study of extended length flex hose performance
7	Fan-less experiments comparing current active solution with passive designs

Table 13: Tests performed on baseline technology in 100 kW prototype data center in ALU's Naperville, IL facility.

Subsection 4.8.2: Baseline system-level testing of current cooling technology

Baseline testing of the current shelf-level pumped refrigerant technology was tested in the 10-rack 100 kW Prototype Data Center in Naperville, IL and the 17-rack 100 kW server farm with approximately 400 servers of various makes and models in Plano, TX (see Figure 67). Testing performed at the Naperville, IL prototype data center is listed in Table 13. The testing at the



Figure 68: ALU facilities in Naperville, IL used for evaluation of layout of underfloor refrigerant piping: (left) Prototype Data Center; and (right) Adjacent Data Center. Note shelf-level modular cooling units attached to rectifier cabinets.

Plano, TX facility included a comprehensive “before” and “after” energy audit, as well as temperature and acoustic noise measurements within the server farm. Note that these tests were done in conjunction with and immediately prior to the start of the ARCTIC project. Details of the tests are not provided for proprietary and confidentiality reasons.

Subsection 4.8.3: System-level testing with new refrigerant distribution system layout

The purpose of this subtask was to investigate the performance of the system with an underfloor layout of the refrigerant piping, as opposed to the current approach using overhead piping. This is to potentially support customer installations that would prefer the refrigerant piping recessed, both for functional/operational and aesthetic reasons.

All testing was performed in the Prototype Data Center and Adjacent Data Center in ALU’s facility in Naperville, IL. The testing had originally been planned to be performed in ALU’s Plano, TX facility, however, the proposed room location had a raised floor whose height was insufficient to accommodate the underfloor piping. Figure 68 shows photographs of the facilities used for evaluation of the system. The refrigerant pump and 10 racks of equipment are located in the Prototype Data Center. The Adjacent Data Center was located approximately 80 feet from the refrigerant pump, so as to evaluate the ability for refrigerant to be adequately supplied to (and returned from) heat exchangers located sufficiently remote from the refrigerant pump. Heat sources in the Adjacent Data Center included an equipment rack with 18 kW of heat load disbursed equally across three shelves, along with additional heat load from four rectifiers bays located in the Adjacent Data Center environment. All refrigerant piping supplying the Prototype Data Center and Adjacent Data Center was routed underfloor.

Refrigerant temperatures were measured at the pump supply and return as well as at the inlet and outlet for several heat exchangers located in the Prototype Data Center and the Adjacent Data

Test #	Details
1	Baseline performance of the underfloor system
2	Disturbance of system performance by adding heat load
3	Temporary power failure to refrigerant pump
4	Start-up with large thermal load
5	Start-up with small thermal load
6	Direct wind-tunnel comparison of under-floor and above floor performance

Table 14: Tests performed for evaluation of system performance with underfloor refrigerant piping.

Center. Refrigerant pressures were measured at the pump supply and return and at the inlet and outlet of two of the heat exchangers, one in the Prototype Data Center and one in the Adjacent Data Center. Air temperatures and air velocities were measured at the inlet and outlet of several heat exchangers so as to determine the cooling capacity of selected coils. Chilled water supply and return temperatures were measured. Total power (heat load) supplied to the overall system was also estimated. Table 14 lists the particular tests that were performed on the system.

Details of the test results are considered proprietary and confidential, and only the conclusions drawn from the testing are presented below:

- Use of under floor refrigerant piping does have an impact on coil cooling capacity. It is thus not recommended that under floor refrigerant piping be used. However, if under the floor piping is the only option, then proper care and engineering analysis should be used to deploy the system.
- When the thermal loads are low the refrigerant does not return to the pump unit properly, causing the pump to cavitate, which could lead to premature failure
- Significant amounts of condensation can occur on the quick-connect couplings if the under floor air is not conditioned and the humidity levels are higher than those of the data room where the pump and IT equipment are generally located.




Figure 69: Photographs of facilities used for system-level testing of shelf-level heat exchangers and hose assemblies, in ALU's Plano, TX facility: (left) Server Farm 1; and (right) Planned Lab Expansion.

Subsection 4.8.4: System-level testing of new components

System-level testing of preferred heat exchanger designs and hose assemblies was performed in Server Farm 1 and the Planned Lab Expansion (see Figure 69) in ALU's Plano, TX facility. The purpose of the testing was to evaluate any potential impact on system-level performance of the new components.

Heat exchangers

Heat exchanger testing performed by Modine in Task 1 (see Section 4.1) and ALU in Subtask 8.1 (See Subsection 4.8.1) resulted in the recommendation of coil option 1B for system-level testing. This coil had a redesigned microchannel tube with larger cross-sectional area and larger wetted perimeter as well as improved air-side heat transfer due to its use of a flush fin design with a lower fin density than the current coil design. Modine made 51 option 1B coils which were installed on 17 equipment racks in ALU's Server Farm 1. The system-level heat exchanger testing results were consistent with prior component-level testing performed by Modine and ALU in Task 1 and Subtask 8.1, and indicate that the coils provided enhanced heat transfer and improved airflow characteristics. This constituted a successful passing of system-level testing for the next-generation heat exchangers.

Non-conductive hose assemblies

A set of 3 pairs of the down-selected non-conductive hose assemblies based on a stainless steel hose assembly with a plastic bushing near the manifold side of the hose assembly were manufactured by USHose for system-level testing in the Planned Lab Expansion in ALU's Plano, TX facility (see Figure 69). The hose assemblies were installed on an equipment rack and the performance of the equipment and the hose monitored over the course of several days. No issues were observed with respect to either cooling performance or the build-up of static electricity. This constituted successfully passing system-level testing for the non-conductive hose assemblies.

1. Product cost analysis.
2. Identification and establishment of vendors.
3. Development of customer business cases.
4. Establishment of pricing.
5. Establishment of total addressable market (TAM) and market share.
6. Development of a sales plan and determination of cost.
7. Development of Alcatel-Lucent business case.
8. Re-issuing of Product Ordering Guides, Configurators, Brochures, Data Sheets, Installation Guides and Sales collateral.
9. Management review.

Figure 70: Alcatel-Lucent stage-gate criteria for commercial product development.

Section 4.9: Business Development (Task 9)

The focus of this task was to support the commercialization of the liquid cooling technology. This included commercializing of the current shelf-level technology as well as next-generation enhancements to the technology, each of which is discussed below. Also discussed are efforts to license the technology.

Commercialization of current shelf-level pumped refrigerant technology

Alcatel-Lucent devoted significant effort to commercialization of the current shelf-level pumped refrigerant technology. The stage-gate criteria used for product development are listed in Figure 70. These commercialization activities for the current shelf-level pumped refrigerant technology proceeded prior to and during the course of the project. Efforts included: i) development of prototypes demonstrating the technology, including the 10-rack 100 kW prototype data center in ALU's Naperville, IL facility; ii) retrofitting of an internal customer, namely ALU's IPTV Test Lab in Plano, TX, with the technology, including a study quantifying energy benefits by performing a detailed energy analysis before- and after implementation of the technology; iii) performing a thorough and detailed business analysis and associated development, as illustrated in Figure 70, to validate the commercial viability of the technology; and iv) performing detailed compliance testing to certify the technology for end use. The effort was successful, resulting in installation of the technology in 6 locations in the U.S. for 3 customers. Details on the particular installations are provided in Chapter 6.

1. Establish requirements and commercialization timeline
2. Create part level documentation
 - a. Create, Process, and Archive Manufacturing Drawings / Bill Of Material Data
 - b. Establish Orderable Items
3. Create System Level Documentation
 - a. Update Technical Documentation
 - i. Applications, Planning, and Ordering Guide (APOG)
 - ii. Installation, Operation, and Maintenance Manual (IO&M)
 - iii. Engineering Rules Document (ERD)
 - iv. Individual Piece Part Drawings
4. Support additional testing required for commercialization
 - a. Agency Approval
 - b. Compliancy Standards (NEBS)
5. Update configurator tools as needed

Figure 71: Business development activities in support of commercialization of liquid cooling technology.

Commercialization of next-generation enhancements

Commercialization of next-generation enhancements to the shelf-level cooling technology focused on activities listed in Figure 71; note that this is a smaller subset of activities than those listed in Figure 70 as the overall successful business case for the product offering had already been developed. Activities related to the pump unit development, new connector technologies and next-generation heat exchangers are discussed in the following.

Pump unit development

The pump unit development is highlighted in detail in Section 4.6. ALU Services Business Division’s activities focused on the specification of the requirements for the commercial pumping stations, which corresponds to bulleted item 1 in Figure 71. This included an extensive set of commercial requirements for a chilled-water refrigerant pumping station as well as commercial requirements for a refrigerant pumping station that interfaced with a compressorized / condenserized secondary loop (for those locations without facilities supporting building chilled water). This effort involved extensive interaction with the pump vendor with a total of

approximately 90 necessary features and 20 desirable features being specified in the design requirements.

ALU Services Business Division decided not to commercialize the pumping station, as, during the course of the project: i) two other pump vendors entered the market with commercially-available pumping stations; and ii) the existing commercial pump vendor announced enhancements and features to their current pumping station that addressed a number of concerns with regard to features and scale-ability.

Develop and research new connector technologies

ALU Services Business Division was actively involved in all aspects of commercializing new connector technology described in Section 4.7.

The non-conductive hose assembly development work described in Subsection 4.7.2 involved a number of the activities listed in Figure 71, both for the initial design based on a polymer-lined hose assembly as well as the final design based on a stainless steel hose with a plastic bushing placed on the manifold end of the hose assembly. Particular emphasis was placed on specification of appropriate tests to satisfy NEBS (Network Equipment Building System) compliance requirements. There was also significant effort devoted to pursuing hose assembly designs that held promise to meet these stringent compliance specifications, were manufacturable and met cost and performance targets.

The hose optimization effort described in Subsection 4.7.1 was not commercialized as it was determined that switching to new tube materials would have resulted in significant potential reliability issues for a marginal improvement in product cost. The new connector technologies effort described in Subsections 4.7.3 and 4.7.4 also was not commercialized as it was determined that, even though a new fitting could have been offered at a somewhat lower price to the end customer, there was a negative effect on Alcatel-Lucent's total cost from having to maintain two different sets of product lines associated with the different fittings.

Next-generation heat exchangers

The preferred heat exchanger design (Design option 1B) chosen for system testing offered an approximately 15% coefficient of performance (COP) enhancement, where $COP = (\text{cooling capacity})/(\text{power for cooling})$, relative to the existing commercial heat exchanger. This design had an improved tube side geometry with larger free flow area and greater wetted perimeter, and incorporated a flush face and reduced fin density on the air side. This design was manufacture-able by Modine in the near term.

The particular design which appeared to have the most promise for energy efficiency enhancement, namely the two-row counterflow design (Design option 3C), had unexpectedly large air-side pressure drop relative to modeling predictions; initial modeling predictions

suggested that the design could have allowed the elimination of supplemental cooling fans, which would have greatly reduced the energy requirements of the shelf-level pumped refrigerant system. Additionally, the two-row counterflow design was also not readily manufacture-able in the near term without significant up-front engineering costs to develop the tooling to support commercialization as well as additional costs to cover Modine's inventory of material supporting the current commercial heat exchanger design.

The current commercialization status is such that Design Option 1B is manufacture-able and available from Modine. However, Alcatel-Lucent decided not to fully commercialize this design for two reasons: i) there was not a compelling customer demand; and ii) efforts and resources were focused on commercialization of the shelf-level pumped refrigerant system as well as licensing of the technology (see below).

Licensing (see write-up in Section 6.3)

During the project it was determined that Alcatel-Lucent would find a licensee to bring the modular cooling solution to market. Emerald Technologies, Inc. a leading provider of energy efficient data center cooling solutions, signed a License Agreement that was executed in December 2010. The company has made commercial sales of the modular cooling solution to leading universities in the United States market. Information on the company and their product can be found at their website: <http://www.emeraldtechnologies.com/>.

1. Create System Level Documentation
 - a. Update Technical Documentation as new features are approved
 - i. Applications, Planning, and Ordering Guide (APOG)
 - ii. Installation, Operation, and Maintenance Manual (IO&M)
 - iii. Engineering Rules Document (ERD)
2. Update configurator tools as needed (see Figure 73 for a screenshot of the Modular

Figure 72: Business development activities performed in support of training.

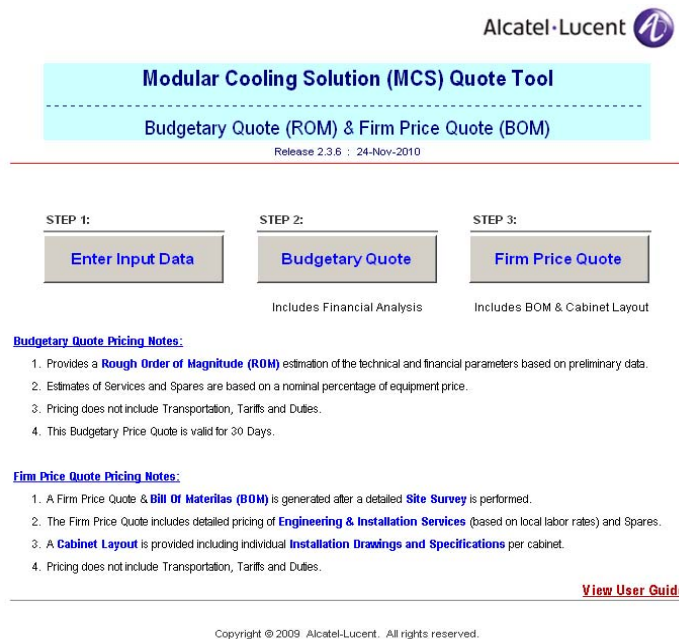


Figure 73: Screenshot of start-up window for the Modular Cooling Solution Quote Tool.

Section 4.10: Training (Task 10)

The goal of this task was to develop training procedures for installation, operation and maintenance of liquid-cooling technology. This included creating sales collateral, brochures, data sheets, configurators, documentation and user manuals for installation, operation and maintenance of equipment, as well as training of individuals.

During the course of the project a number of activities were undertaken in support of training. Some of these activities, listed in Figure 72, were directly related to project tasks, while others were indirectly related through support of the commercial product offering. The particular details for each of these activities are not included here for proprietary and confidentiality reasons. We do include a screenshot of the start-up window (see Figure 73) for the Quote Tool used for creation of Budgetary Quotes (ROMs) and Firm Price Quotes (BOMs) for the shelf-level cooling technology.

Section 4.11: Project Management and Reporting (Task 11)

This task focused on two areas, namely: i) management of the project to ensure that task activities were completed on time and within budget; and ii) meeting government reporting requirements, which included submitting formal reports and documents as well as providing informal feedback to DOE via e-mail exchanges and phone calls. Each of these areas is discussed in detail in the following.

Project management activities included holding regular meetings, typically on a weekly basis, with all project participants, both internally and externally. A detailed work-breakdown-structure (WBS) was also created to track progress on a task-by-task basis, with owners assigned to individual task activities. The WBS was reviewed and updated on a regular basis. This approach helped facilitate identification of any potential risks to the project, whether technical, resource or schedule related, as well as allowed the development of appropriate mitigation strategies for any identified risks. Project financial details were also tracked in detail internally, and invoicing activities were performed on a regular basis.

Government reporting included: Quarterly Progress Reports, ARRA reporting, revised SF424 and PMC123 forms, semi-annual Davis-Bacon reporting, presentation and attendance at Annual Review meetings, and hosting site visits. Regular phone calls were also held, as needed, with the DOE Project Officer and Project Monitor at different stages during the project to apprise them of important project developments. For example, the downscoping of Modine's activities due to the identification of an already-developed prototype refrigerant pump at another vendor (see Task 6 details in Section 4.6) along with the hose performance issue identified in late 4Q11 that subsequently required requesting a no-cost time extension to address (see Task 7 details in Subsection 4.7.2), required frequent and detailed communication and interaction.

<u>Chilled water CRAC units</u>	<u>Rack-level liquid cooling</u>
<ul style="list-style-type: none"> • 500 server cabinets at 3 kW per cabinet • 22,000 square foot service area • CRAC units cool 47 kW; consume 13 kW • 32.1 chilled water CRAC units required to cool total equipment load (Supported by actual conditions: 32 of 40 units are running in practice) • Energy required: 32 x 13 kW x 8760 hours = 3,640,000 kWh 	<ul style="list-style-type: none"> • 150 server cabinets at 10 kW per cabinet <ul style="list-style-type: none"> – Space freed up for 350 cabinets • 3 heat exchangers + 6 fans total per cabinet (2 fans per heat exchanger) • Pump unit uses 800 W • Each fan uses 25 W (150 W per cabinet) • Cooling power: 30.5 kW = (10 pumps x 0.8 kW/pump + 150 cabinets x 0.150 kW / cabinet) • Energy required: 30.5 kW x 8760 hours = 267,000 kWh

Figure 74: Comparison of cooling power requirements for a 1.5 MW data center cooled by chilled water CRAC units (left) and ALU modular liquid cooling system (right).

Chapter 5: Benefits Assessment

Section 5.1: Economic Benefits and Capital Cost Advantages

To illustrate the operating-expense and capital-cost advantages of ARCTIC, we consider the retro-fit of a data center with the ALU modular liquid cooling system. The data center that we consider is a scaled-up version of the prototype system developed internally within ALU; details of the configurations and energy usage for the cases of chilled-water CRAC units and the modular liquid cooling system are presented in Figure 74. Note that the approach using chilled-water CRAC units requires 3,640,000 kWh of energy annually, whereas the modular liquid cooling system requires only 267,000 kWh, a **93% reduction** in energy usage. At a cost of \$0.10 per kWh, this represents an annual operating-expense reduction of \$337k. We estimate the cost for this retrofit at \$1.24M, which would thus be paid back in 3.7 years. To retrofit an air-cooled system using compressor CRAC units, which consume 30 kW instead of the 13 kW required for chilled-water units, the payback period is only 1.5 years.

A significant benefit of the ALU modular liquid cooling system is improved real-estate utilization. The three-fold increase in power density frees up 350 cabinets in the above example. Assuming 12 sq ft per cabinet (6 sq ft each of cabinet and aisle space) and a construction cost of \$200 / sq ft, this results in an additional \$840k in space utilization savings, bringing the payback period to 1 year for chilled-water CRAC units and 0.75 years for compressor CRAC units.

This factor-of-3 increase in equipment density will have a profound effect on the long-term financial and energy costs of ICT cooling as well as on one-time retrofit expense. At growth rates of 5 to 7 % per year, it will take the ICT industry ~20 years to expand its equipment base by this factor of 3. Long before then, if our technology penetrates the market widely, new equipment installations will start to fit in existing real estate without the need for more building.

Thus, we estimate that the cost of building new ICT facilities will be reduced to near zero by this technology. We estimate that the present total built area occupied by the U.S. ICT industry is 0.137×10^9 sq ft (based on 120 billion kWh annually and an energy density of 100 W/sq ft). Thus, the total yearly cost savings due to better real estate utilization is

$$6\% * 0.137 \times 10^9 \text{ sq ft} * \$200 / \text{sq ft} = \$1.6 \text{ billion.}$$

Section 5.2: Estimated energy savings across U.S. ICT facilities

The estimates in this section are based on the DOE presentation entitled “Routing Telecom and Data Centers Toward Energy Efficient Use” dated May 13, 2009²³, namely:

- Total US energy expenditure for ICT facilities of 120 billion kWh in 2006, with 25% for cooling:
 - total annual US energy expenditure of 60 billion kWh for data centers
 - total annual US energy usage of 20 billion kWh for large communications centers and network trunk lines
 - local equipment, private exchanges and mobile phone towers of 40 billion kWh

Based on these estimates, the total annual U.S. ICT expenditure for cooling in 2006 was

$$25\% * 120.0 \text{ billion kWh} = 30.0 \text{ billion kWh.}$$

We estimate that 25% of this cooling expenditure is used for running CRAC units in existing ICT facilities. Based on our prototyping and typical data center and central office configurations, we estimate a reduction of this portion of the energy expenditure by a factor of 11.8. If this technology is universally adopted, total U.S. ICT cooling power would be reduced to

$$[75\% + (25\%/11.8)] * 30.0 \text{ billion kWh} = 23.1 \text{ billion kWh.}$$

Savings in the energy costs of real-estate expansion are a significant additional component of energy-use reduction enabled by our technology. Estimates for energy savings due to better use of space follow similar reasoning to that discussed in Section 5.1 for building cost savings. We estimate that 5 to 7% of the present annual energy cost associated with building new ICT facilities will be saved by deployment of our technology. Assuming that the present total built area occupied by the U.S. ICT industry is 0.137×10^9 sq ft, the total yearly energy cost saved by improved real estate utilization²⁴ is

$$6\% * 0.137 \times 10^9 \text{ sq ft} * 332 \text{ kWh/ sq ft} = 2.7 \text{ billion kWh,}$$

where the value of 332 kWh/sq ft is taken from [24]. This is added to the present U.S. ICT cooling energy usage of 30.0 billion kWh to yield a total present yearly energy cost of 32.7 billion kWh. Thus, we estimate that the total yearly energy savings in the U.S will ultimately amount to

$$32.7 \text{ billion kWh} - 23.1 \text{ billion kWh} = 9.6 \text{ billion kWh.}$$

Environmental Metric (Note: MT ≡ Metric tons)	Emissions per billion kWh (from ICT_Benefits_Spreadsheets.xls)	Annual Emission Reductions for modular liquid cooling
CO displaced (MT)	164.4	1580
CO2 displaced (MM TCE)	0.154	1.48
SO2 displaced (MT)	2070	19870
NOx displaced (MT)	1359	13050
Particulates displaced (MT)	44.1	423
VOCs displaced (MT)	17.7	170

Table 15: Estimated annual U.S. emission reductions (column 3) for modular liquid cooling based on 9.6 billion kWh savings. Column 2 data scaling emissions with energy usage are calculated from the “ICT_Benefits_Spreadsheets.xls” originally located at http://sites.energetics.com/ICT_benefits09 that was provided by DOE during the proposal submission phase.

Section 5.3: Estimated reductions in greenhouse gas emissions and other environmental emissions

In Table 15 estimates for reductions in U.S. greenhouse gas and other environmental emissions are provided based on an annual energy reduction of 9.6 billion kWh. Data for the emissions per billion kWh (middle column of Table 15) were obtained by using values for each metric as provided in the Results sheet in the ICT_Benefits_Spreadsheets.xls spreadsheet provided by the DOE for the proposal²⁵.

Chapter 6: Commercialization

Section 6.1: Commercialization Strategy

The proposed ARCTIC commercialization strategy addresses two main markets: retrofits and new installations. The retrofit market appears particularly viable, as the cooling system is not vendor dependent, and there is a large installed equipment base, with estimates of ~10 million servers in 10,000 data centers in North America. New installations include sales into new ICT facilities, as well as design and integration of shelf- and component-level cooling technology into next-generation equipment that will be particularly challenging to cool. For example, ALU is currently considering this technology for application in select product platforms.

A number of studies estimate that between 25 and 50 percent of total power is being used for cooling^{23,26,27}, and these values may be overly optimistic as they are reflective of facilities that are using industry-accepted best practices. The large percentage of total power spent on cooling illustrates that there is a pressing need for the energy-efficient modular liquid cooling technology developed in this project and speaks strongly to its economic viability. Furthermore, ALU envisions that the modular liquid cooling technology will readily meet any potential regulatory and environmental requirements. ALU has extensive experience in ICT equipment testing and expects the modular liquid cooling system to achieve NEBS, UL, CL and CSA compliance. For example, this system has passed a number of stringent NEBS requirements, such as Zone 4 Earthquake tests, typically not required for data center environments. Finally, we note the ALU modular liquid cooling system offers substantial advantages in meeting energy reduction requirements that may be mandated in the future; this benefit has been validated in internal applications of the technology (see Figure 75).

ALU Services Business Division's commercialization strategy has thus been two-fold: i) initial focus on data center retro-fit with the modular liquid cooling system. Because of the potentially improved use of space, a secondary focus will be consolidation of small data centers into a single location. New installations will also be investigated, but will not be the primary focus; and ii) research and development on current and next-generation liquid-cooling technology, to address future market requirements.

Section 6.2: Applications of the technology: Customer Success Stories

ALU Services Business Division has made a number of commercial sales of the modular liquid cooling technology. The commercial implementations include both retro-fits as well as new installations. Details on several of the implementations are provided below:

1. US Wireless Carrier: The customer had a serious heat issue at their Mobile Switch Center (MSC) site in Texas, with aisle temperatures approaching 90° F. Additionally, their present heat densities were requiring that they spread out heat load leading to poor utilization of floor space. Finally they wanted to reduce the energy required for cooling.



Figure 75: Alcatel-Lucent’s Modular Cooling Solution installed in ALU’s 17-rack, 100 kW Network Systems Integration and Test IPTV Laboratory in Plano, TX.

Solution: Refrigerant cooling systems were deployed in two separate locations in the office, addressing heating concerns in 57 cabinets (19 in a “data room”, 38 in the “telco room”).

Results: Aisle temperatures were reduced from a high of 90°F to between 67 and 72°F. Two CRAC Units were turned off, and a third avoided. Based on the difference between the power consumed by the condenser model CRAC units and the condenser model refrigerant pumps added, this resulted in an estimated power reduction of 54% using the pumped refrigerant technology, which translates into a power savings of 75.7 KW and a yearly energy savings of about 663,400 KW hours, or \$66,340 dollars based on 0.10 \$/KW hr.

2. A US Satellite TV service provider: The customer planned for very high per cabinet heat loads, and determined that standard CRAC-based cooling would not be sufficient.

Solution: The customer deployed two chilled water refrigerant systems in a redundant solution. 19 Heat exchangers were mounted in 10 cabinets.

Result: The refrigerant cooling system provided adequate cooling for the dense heat loads in the office, and the redundant design provided for a robust cooling system. Although the customer did not consider a CRAC based system, an equivalent redundant, chilled water CRAC system would require 3 CRACS. This results in an estimated power reduction of 97% using the pumped refrigerant technology, which translates into an estimated power savings of 48.8KW and a yearly energy savings of about 427,400 KW hours, or \$42,740 dollars based on 0.10 \$/KW hr.

3. A US Wireless Carrier: Desired to design their new MSC sites with reduced floor space and increased heat per cabinet. The overall design could eventually grow to 150KW of heat per cabinet row (15 cabinets per row with 10KW each) with a potential of 12 rows of equipment per site. If successful, they planned to use this design in several (up to 8) locations in the US. Their analysis determined that traditional CRAC-based cooling would not be sufficient for the densities planned or the floor space that they wanted to dedicate to the new MSC sites. Additionally, they wanted to avoid a raised floor.

Solution: The wireless carrier chose to use a chilled water solution in their new MSC design. Based on the heat load and densities planned, each row of 15 cabinets required a single chilled water refrigerant pump.

Result: To date, four sites have been deployed with 4 pumps per site. Estimates for full site power reduction compared to a purely CRAC-based system (requiring 32 units) would be 97% using the pumped refrigerant technology. This translates into a power savings of 515KW and a yearly energy savings of about 4,500,000 KW hours, or \$450,000 dollars based on 0.10 \$/KW hr for a fully-populated site.

Section 6.3: Industrialization

During the project it was determined that Alcatel-Lucent would find a licensee to bring the modular cooling solution to market. Emerald Technologies, Inc. signed a License Agreement that was executed in December 2010. Emerald Technologies is a leading provider of energy efficient data center cooling solutions. The company has made commercial sales of the modular cooling solution to leading universities in the United States market. Information on the company and their product can be found at their website: <http://www.emeraldtechnologies.com/>.

Chapter 7: Accomplishments

In the following we detail the accomplishments with respect to project objectives on a task-by-task basis. Highlights include:

- Successful commercialization of the shelf-level pumped refrigerant technology, resulting in installation at 6 customer locations in the U.S at the time of the writing of this report.
- Licensing of shelf-level pumped refrigerant cooling technology to Emerald Technologies, who will be taking over commercialization activities.
- Development and commercialization of several enhancements to the shelf-level pumped refrigerant technology.
- Filing of 5 patent applications with the USPTO, with a 6th application in process to be filed with the USPTO.
- Advanced research and development activities focusing on enhanced component-level cooling using liquid- and air-based technologies.

Task specific accomplishments are listed below:

Task 1

- Several next-generation heat exchanger designs were modeled, prototyped and tested, demonstrating an increased coefficient-of-performance (COP), defined as $COP = (\text{cooling capacity}) / (\text{energy for cooling})$, in excess of 15%.
- A large-scale build of heat exchangers meeting the above COP enhancement was completed for system-level testing in Task 8, thereby meeting a project milestone.

Task 2

- A novel approach for manufacturing air-cooled heat sinks with three-dimensional features was demonstrated through numerical modeling, prototype fabrication and experimental evaluation. The effort focused on slotted honeycomb heat sinks, which are a design that represents a good compromise between manufacture-ability and performance enhancement relative to current linear fin heat sink technology. The approach was based on stacking of different “wafers” that are subsequently sandwiched together to create the completed part.
- Two patent applications based on the novel manufacturing approach were filed with the USPTO.

- A paper describing the task activities was presented at ITHERM 2012 and included in the conference proceedings.

Task 3

- Bench-scale and rack-level test apparatuses were developed for evaluating the thermal backplane approach.
- A heat-pipe-based thermal backplane prototype was developed that demonstrated a 10% reduction in thermal resistance and 50% decrease in form factor relative to a state-of-the-art fan and heat sink assembly, thereby meeting a project milestone.
- Two patent applications based on the thermal backplane approach were filed with the USPTO.
- Several make/break thermal/mechanical connections were developed and tested.
- Key factors influencing the commercial applicability of the thermal backplane approach were determined, including: i) identification of potential supply chain partners and their capabilities, as verified through extensive prototyping and testing; and ii) bounds for heat-density enhancement due to limitations on heat pipe and vapor chamber performance, including trade-offs between form factor, orientation, operating temperature and heat transport distance.

Task 4

- Bench-scale and rack-level test apparatuses developed for evaluating direct component-level liquid cooling.
- Utilized test apparatuses to evaluate performance of several cold plate designs, allowing determination of upper bounds on heat density enhancement for single-phase and two-phase heat transfer.
- One patent application on a novel microchannel heat sink design filed with the USPTO. A second patent application related to control of two-phase liquid cooling systems is in process to be filed with the USPTO.
- Identified key supply chain partners with unique capabilities for commercial implementation of direct component-level liquid cooling technology.

Task 5

- Developed experimental capabilities enabling testing and evaluation of technologies developed for improved cooling at the device, board, shelf, and rack-level. This activity supported testing for Tasks 1, 3, 4, 7 and 8.

Task 6

- Experimentally evaluated a prototype chilled-water refrigerant pump with enhanced features relative to the existing commercial pumping unit, in particular, a more sophisticated controller and the ability to handle a broader range of power levels.
- Developed a complete set of commercial pump design requirements, including incorporating enhancements learned from the prototype pump evaluation, for a refrigerant pump that interfaces to chilled water for heat rejection.
- Developed a complete set of commercial pump design requirements for a refrigerant pump that interfaces to a compressorized / condenserized secondary loop for heat rejection.

Task 7

- Developed and commercialized a NEBS-compliant non-conductive hose assembly for application in data center and telecommunications central office environments utilizing different grounding schemes for equipment frames and the refrigerant pump.
- Investigated several approaches to cost reduce stainless steel and polymer-lined hose assemblies and connectors.

Task 8

- Performed system-level testing supporting the development of new heat exchangers (Task 1), refrigerant pumps (Task 6) and hoses and connectors (Task 7).
- Evaluated the performance of a refrigerant distribution network where the refrigerant plumbing is placed underfloor, instead of overhead.

Task 9

- Developed a successful commercial product offering based on the shelf-level pumped refrigerant technology, including installation of the technology in 6 customer locations within the U.S.
- Licensed the shelf-level pumped refrigerant technology to Emerald Technologies, who will be taking over commercialization activities.
- Commercialized the non-conductive hose assemblies (Task 7) for application in a customer site.
- Lead the development of final commercial refrigerant pump design requirements for sites with building chilled water and compressorized / condenserized secondary loops.

Task 10

- Created system-level documentation for new product features, including: Applications, Planning, and Ordering Guide (APOG); Installation, Operation, and Maintenance Manual (IO&M) and Engineering Rules Document (ERD).
- Updated configurator tools supporting creation of Budgetary Quotes (ROMs) and Firm Price Quotes (BOMs).

Task 11

- Met project requirements with respect to Quarterly Progress Reports, ARRA Reporting, Annual Reviews, Site Visits, status calls with DOE, etc.

Chapter 8: Conclusions

The pumped-refrigerant-based cooling technology developed in this project is a reliable and energy-efficient way to cool data centers and central offices and has a number of benefits over traditional air- and liquid-based cooling approaches, including: i) projected energy savings in excess of 90% for the pumped refrigerant-based technology relative to conventional CRAC-based air cooling due to the substantially greater heat-carrying capacity relative to pumping power for the refrigerant than air (due to the refrigerant's large heat of vaporization); ii) dramatic increases in device densities at the shelf, circuit-pack and component levels not readily achievable via purely air-based approaches; iii) the use of a working fluid, namely R134a refrigerant, which is a dielectric liquid that, unlike water, will not damage electronics in the case of a leak. R134a is also non-toxic, and, since the system is based on a pumped refrigerant, it does not require toxic lubricants and oils required by refrigerant systems; and iv) relatively modest system pressures (~60 psi) in comparison to a compressorized refrigerant system.

The project's technical activities focused on two primary areas, namely: i) enhancements to the shelf-level pumped refrigerant technology, including refrigerant pumps, heat exchangers, hoses and connectors, and refrigerant piping layout; and ii) next-generation technologies that enhance cooling capability and efficiency at the device level, including three-dimensional heat sinks for improved air cooling, and thermal backplane and direct component-level liquid cooling. Conclusions and observations for each of these activities is discussed in the following, along with important aspects related to commercialization of the technology.

The shelf-level heat exchanger enhancements developed by Modine showed promise to further decrease the cooling overhead by reducing required fan power. However, the particular design expected to achieve the best air-side performance, namely the two-row, counter-flow heat exchanger, was experimentally observed to have much larger air-side pressure drop than expected. This requires additional investigation to understand the source of the discrepancy between modeling and experiment, and may require a redesign to achieve desired performance, which, if realized, could result in the elimination of fans currently used for supplemental airflow. This would yield additional benefits for overall cooling efficiency as well as reliability.

The non-conductive hose assembly developed in the project, namely incorporation of a plastic bushing in a modified version of the current stainless steel hose assembly, was the best available option to meet a pressing commercial need while maintaining reliability, cost and performance targets. A polymer-based non-conductive hose assembly could still have added benefit from a cost, weight and performance perspective, though additional R&D, which was beyond the scope and timeline of the project, is needed to develop materials and the associated manufacturing processes that adequately address the issue of static dissipation as well as guarantee low refrigerant permeability, each of which impacts reliability.

The hose and connector development activity also showed that introducing a new component that is not fully-compatible with the current system may not make financial sense, particularly if

the new component is only slightly lower in cost and has marginally beneficial feature enhancements. This is because there are often hidden costs associated with non-backward compatibility that require maintaining two sets of products, including part numbers, inventory and the associated manufacturing processes. This effect was observed in the cost reduction and optimization of the hose assemblies, where it was determined that there was not a compelling financial reason to replace the existing quick-disconnect fitting, as overall cost would have increased with minimal performance enhancement.

The pump development activity was particularly relevant and timely, addressing important needs for a pump architecture that scaled appropriately with heat load, had an improved controller relative to the existing commercially-available pumping station, and accommodated facilities with and without building chilled water. The pump development activity was not pursued through commercialization as, during the course of the work: i) two other companies introduced commercial pumps with desired features; and ii) a number of enhancements were added to the existing commercially-available pumping station that addressed deficiencies in the original design. This resulted in the availability of three different pumping stations that met commercial design requirements.

System-level testing with the underfloor refrigerant distribution network showed that care must be taken in the design and implementation of such a system, otherwise there is a potential risk of undersupplying refrigerant to the shelf-level heat exchangers.

Research and development of the thermal backplane and device-level liquid cooling approaches showed significant promise for facilitating continued increases in device heat density while maintaining energy efficiency. A number of concepts were developed that investigated key requirements deemed essential for commercial applicability, including cooling capacity (heat load) and performance (pressure drop), form factor, reliability, cost and maintainability (e.g., plug-and-play capability for replacement of circuit packs). The activity also allowed the identification of key supply chain partners for further development and commercialization of the approaches.

An important conclusion drawn from the component-level cooling studies is that a stand-alone liquid cooling system at the shelf or rack level is likely the right level of granularity for implementing component-level liquid cooling. This approach effectively separates the operation of the rack-level cooling system from the room-level liquid-cooling system. This separation has a number of reliability benefits, as it allows the room-level liquid-cooling system to be isolated from any cooling failures that might occur within the equipment rack. It also allows much greater autonomy from a control perspective, as it may be particularly challenging, for example, to design a single system capable of efficiently implementing cooling across multiple scales, e.g., tens of kilowatts at the rack level to hundreds of watts at the component level. This modular approach would also give considerable flexibility to data center designers and IT equipment manufacturers in terms of how they architect their respective cooling infrastructure – they only

need to agree upon the details of the thermal interface that exchanges heat from the rack to the room-level cooling solution.

Two-phase flow systems exhibit positive feedback, which can have potential deleterious operational effects. The impact of this type of feedback, investigated via benchtop experiments with refrigerant distributed to multiple heat sources in parallel, clearly showed that, without any type of active control, there are stable and unstable regions of operation. Research into how to best design and/or control such a system for robust and predictable operation, while also maintaining adequate cost, performance and reliability, is needed. Modeling tools and design rules for predicting and/or estimating performance of two-phase flow systems would also be particularly beneficial in this regard.

The air-cooled heat sink activity yielded particularly promising results for the development of a cost-effective manufacturing process for making complex, three-dimensional heat sink structures with enhanced performance relative to the current state-of-the-art linear fin heat sink. The excellent agreement between numerical modeling and experimental results showed that simulations can be a powerful tool in the design and optimization of these structures. Additional work exploring refinements to the manufacturing process is needed, and will likely be explored in future collaboration with the supply chain partner involved in the work, who is particularly interested in the commercial application of the technology.

Commercialization activities focused on the shelf-level pumped refrigerant technology successfully demonstrated the market need for the technology, resulting in 6 customer installations in the U.S. The primary commercial benefits related both to reduced energy usage and increase equipment heat densities, which was a significant issue for several installations where the operator felt that a purely air-based cooling approach would be insufficient to meet requisite cooling requirements.

During the project it was also determined that Alcatel-Lucent would find a licensee to bring the modular cooling solution to market. Emerald Technologies, Inc. signed a License Agreement that was executed in December 2010. Emerald Technologies is a leading provider of energy efficient data center cooling solutions. The company has made commercial sales of the modular cooling solution to leading universities in the United States market. Information on the company and their product can be found at their website: <http://www.emeraldtechnologies.com/>. The licensing of the technology and resulting deployments further points to the benefits of the pumped refrigerant technology.

Chapter 9: References

- ¹ <http://www.thermalformandfunction.com/>
- ² <http://www.parker.com>
- ³ <http://www.emersonnetworkpower.com/en-US/Brands/Liebert/Pages/default.aspx>
- ⁴ <http://www.apc.com>
- ⁵ <http://www.clusteredsystems.com>
- ⁶ <http://www.grcooling.com>
- ⁷ <http://www.google.com/about/datacenters/>
- ⁸ S. Krishnan, D. Hernon, M. Hodes, J. Mullins and A. M. Lyons, "Design of Monolithic Heat Sinks of Arbitrary Structure for Enhanced Air Cooling", *IEEE Trans. Components, Packaging and Manufacturing Technology*, 2012, 2, 266-277.
- ⁹ A. Bejan and S. Lorente, "Constructal theory of generation of configuration in nature and engineering," *Journal of Applied Physics*, vol. 100, pp. 041301-27, 2006.
- ¹⁰ D. Liu and S. V. Garimella, "Analysis and Optimization of the Thermal Performance of Microchannel Heat Sinks," *International Journal for Numerical Methods in Heat and Fluid Flow*, vol. 15, pp. 7-26, 2005.
- ¹¹ S. Krishnan, J. Y. Murthy and S. V. Garimella, "Direct Simulation of Transport in Open-Cell Metal Foams," *ASME Journal of Heat Transfer*, vol. 128, pp. 793 – 799, 2006.
- ¹² S. Krishnan, S. V. Garimella and J. Y. Murthy, "Simulation of Thermal Transport in Open-Cell Metal Foams: Effect of Periodic Unit Cell Structure," *ASME Journal of Heat Transfer*, vol. 130, pp. 024503, 2008.
- ¹³ T. J. Lu, "Heat transfer efficiency of metal honeycombs," *International Journal of Heat and Mass Transfer*, vol. 42, pp. 2031-2040, 1998.
- ¹⁴ T. J. Lu, L. Valdevit and A. G. Evans, "Active Cooling By Metallic Sandwich Structures With Periodic Cores", *Progress in Material Science*, vol.50, pp. 789-815, 2005.
- ¹⁵ L. Valdevit, A. Pantano, H.A. Stone and A.G. Evans, "Optimal Active Cooling Performance of Metallic Sandwich Panels With Prismatic Cores", *International Journal of Heat and Mass Transfer*, vol. 49, pp. 3819-3830, 2006.
- ¹⁶ S. Gu, T. J. Lu and A. G. Evans, "On the Design of Two-Dimensional Cellular Metals for Combined Heat Dissipation and Structural Load Capacity", *International Journal of Heat and Mass Transfer*, vol. 44, pp. 2163-2175, 2001.
- ¹⁷ T. Wen, J. Tian, T. J. Lu, D.T. Queheillalt and H.N.G. Wadley, "Forced Convection in Metallic Honeycomb Structures", *International Journal of Heat and Mass Transfer*, vol. 49, pp. 3313-3324, 2006.
- ¹⁸ T. Wen, F. Xu and T. J. Lu, "Structural Optimization of Two-Dimensional Cellular Metals Cooled by Forced Convection", *International Journal of Heat and Mass Transfer*, vol. 50, pp. 2590-2604, 2007.
- ¹⁹ D. Hernon, M. Hodes, A. Lyons, A. O'Loughlin and S. Krishnan. "Monolithic Structurally Complex Heat Sink Designs." US Patent Application 20090321045, 2009.
- ²⁰ A. Lyons, S. Krishnan, J. Mullins, M. Hodes, and D. Hernon, "Advanced Heat Sinks Enabled by Three-Dimensional Printing", *20th International Solid Freeform Fabrication Symposium*, Austin, Texas, 2009.
- ²¹ <http://www.ansys.com>

²² A. Faghri, Heat Pipe Science and Technology. Taylor & Francis Inc., ISBN 978-1-56032-383-9 (1995).

²³ http://sites.energetics.com/ICT_roadmap09/pdfs/ICT_Vision_and_Roadmap_051309_draft.pdf

²⁴ http://architecture2030.org/current_situation/building_sector.html

²⁵ http://sites.energetics.com/ICT_benefits09

²⁶ “Energy Efficient Cooling Solutions for Data Centers”, Emerson Network Power white paper, 2007.

²⁷ R. Schmidt and M. Iyengar, “Thermo-dynamics of information technology data centers”, *IBM J. Res. & Dev.* (2009).