Liquid Cooling for Optical Networking Equipment

Behzad Mohajer[1], Peter Ajersch[2], Michael Bishop[3], Simon Shearman[4], Peter Saturley[5], Marko Nicolici[6]
Packet Optical Devision, Ciena, 383 Terry Fox Drive, Ottawa, ON, K2K 0L1, Canada
bmohajer@ciena.com[1], pajersch@ciena.com[2], mbishop@ciena.com[3], sshearma@ciena.com[4], psaturle@ciena.com[5], mnicolic@ciena.com[6]

Abstract: This article provides insights into a successful upgrade of an air-cooled coherent metro router into a Hybrid Liquid/Air-cooled system. Additionally, an innovative solution is presented for integrating liquid-cooling into the body of pluggable optical modules.

1. Introduction

While liquid cooling technologies have long been utilized in high-performance computing systems, their application in optical networking equipment has remained relatively limited. The performance and longevity of optical modules depend on the temperature they operate at. Although most FPGAs, ASICs, and electronic components need to be maintained below 100°C, most optical components including pluggable optics must be maintained at a temperature typically below 70°C to ensure reliable transmission of data. The ever-increasing transceiver throughput and power draw coupled with the low temperature ratings associated with optical modules makes the thermal management of optical networking equipment particularly challenging. This article explores the integration of liquid-cooling technology into optical networking equipment.

At Ciena, we upgraded our all-air-cooled optical transport platform (Waveserver 5) and router platform (WaveRouter) into Hybrid Liquid/Air-Cooled systems. The hybrid systems leverage liquid cooling for high-power ASICs and coherent pluggable optical modules, while employing air-cooling at low fan speeds for the remaining low-power devices. This article provides insights into the upgrade of the router platform, which resulted in 60-80% reduction in fan power consumption, 12-15dBA reduction in fan noise level, and a significant reduction in facility cooling power consumption.

Additionally, a novel cooling solution is introduced to address the challenging thermal demands associated with present and future small form factor pluggable optical modules. This innovative solution integrates liquid cooling directly into the body of pluggable optical modules. To ensure seamless integration, dripless, blind-mate liquid connectors are utilized, allowing the pluggable modules to be connected to the coolant lines automatically upon insertion. This solution unlocks the potential for pushing the boundaries of power and functionality within extremely limited physical space, leading to advanced and compact optical networking systems.

2. Liquid Cooling Architecture

A liquid cooling system consists of several components (Fig. 1-d): coolant distribution unit (CDU), primary loop, secondary loop, and coldplates. The coolant distribution unit (CDU) includes a coolant reservoir, redundant pumps, and a heat exchanger. The CDU is responsible for chilling the coolant to a set temperature and circulating it through the liquid-cooled systems. The primary loop transfers the building's chilled water to the CDU, where it removes heat from the heated coolant. The secondary loop navigates the coolant through the IT equipment, absorbing the heat generated by the electronics and optics. The heated coolant is then transported back to the CDU for re-chilling. Coldplates are liquid-cooled heatsinks that are attached to heat-generating devices.

3. Hybrid Liquid/Air-Cooled Router Platform

WaveRouter platform is currently offered as an air-cooled product. The platform comprises a 21 RU chassis that houses 15 removable router modules. At the rear of the chassis, 8 fan trays are installed to draw air through the router modules (Fig. 1-a). Each router module hosts several thermally demanding components, including a high-power multi-die ASIC, and 15 coherent pluggable optical modules (QSFP-DD) each capable of transmitting at 400Gb/s. Notably, upcoming generations of these devices are expected to operate at significantly higher power levels. The growing transport capacity and the resulting surge in power consumption present significant challenges for thermal management of modern optical networking systems. To tackle these issues, we undertook an innovative approach upgrading the existing air-cooled router platform into a hybrid liquid/air-cooled configuration with no changes to the PCB assembly.

To develop the hybrid system, we substituted two of the existing fan trays (Fig. 1-b) with two coolant distribution manifolds equipped with blind-mate quick disconnects (Fig. 1-c). These manifolds, connected to a 4 RU 80kW Coolant Distribution Unit (CDU), facilitate the distribution of coolant to the router modules. When a router module is inserted into the chassis, both electrical and liquid connections are mated in a blind-mate fashion at the back of the chassis. The coolant then flows from the chassis-level manifold into the module-level manifold, from where it is further distributed to the primary and secondary sides of the router module.

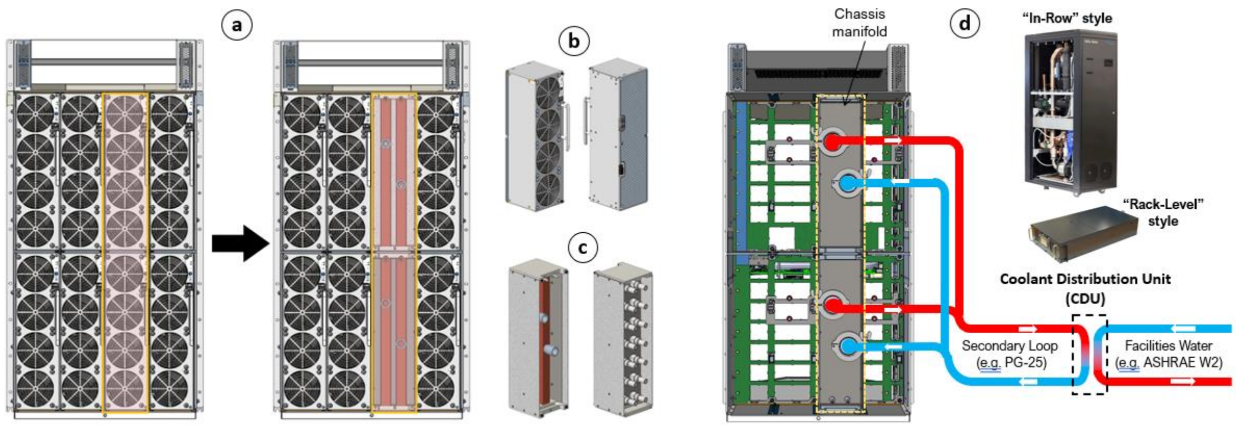

Figure 1: a) Conversion from Air-Cooled to Hybrid Liquid/Air-Cooled, b) rear and front views of a fan tray, c) rear and front views of a manifold, d) liquid-cooling infrastructure.

On the primary side, two parallel coolant lines are routed: one line connects to the coldplate attached to the high-power ASIC (Fig. 2-b), and the other line goes to a low-profile manifold (Fig. 2-c) responsible for distributing the coolant among 8 QSFP-DD low-profile riding coldplates (Fig. 2-d). The riding coldplates are spring loaded to the QSFP-DD cages. When a pluggable optical module is inserted into the cage, the riding coldplates move vertically, allowing for proper thermal contact. On the secondary side of the board, the coolant flows as a third parallel coolant line into another low-profile manifold and is distributed among 7 QSFP-DD riding coldplates.

The main coldplate attached to the ASIC has been successfully tested at power levels up to 1000W in an ASHRAE W4 environment (T_{WATER} = 45°C). In an ASHRAE W1 environment (T_{WATER} = 17°C), this coldplate is capable of providing cooling up to 2300W. The low-profile QSFP-DD riding coldplates are capable of effectively cooling future pluggable optical modules up to 35W in an ASHRAE W4 environment, and 55W in an ASHRAE W1 environment.

The incorporation of liquid-cooling into the WaveRouter system eliminates the need for airflow over selected high-power devices, allowing the system fans to run at a lower speed to manage the remainder of the system power. This reduction in fan speed results in a significant decrease in acoustic noise levels of 12-15 dBA.

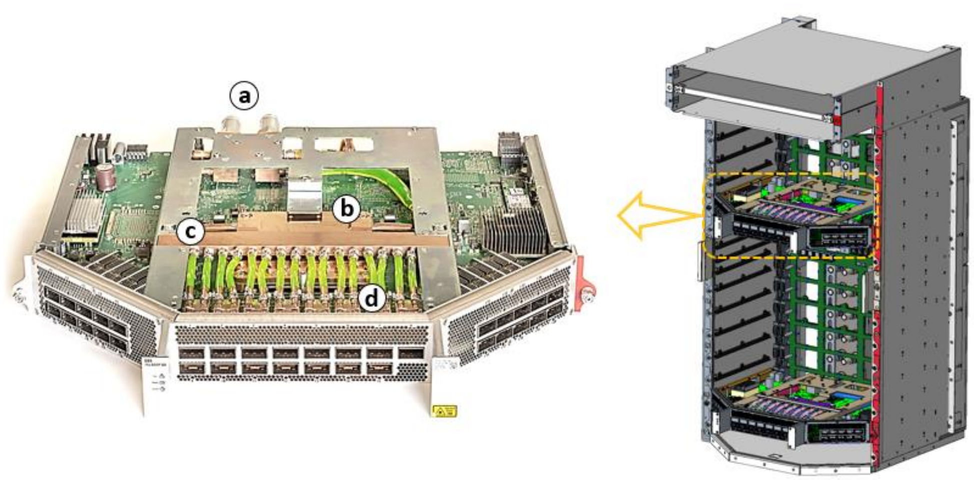

Figure 2: Hybrid liquid/air-cooled router module: a) Quick disconnects, b) ASIC coldplate, c) QSFP-DD manifold, d) QSFP-DD riding coldplate.

4. Liquid-Cooled Pluggable Optical Modules

The evolution of datacenter architectures toward highly interconnected designs requires high-capacity, highly compact optical transport interconnects. This has led to the addition of further functionalities, such as advanced digital signal processing and optical capabilities, into the same compact form factor of pluggable optical modules, which has resulted in a rapid increase in their power consumption.

Conventionally, pluggable optical modules are cooled using air-cooled riding heatsinks, where the pluggable module interfaces with the heatsink in a dry/sliding nature. Replacing the air-cooled riding heatsink with a liquid-cooled riding coldplate can considerably enhance the cooling performance. This approach has been implemented in the hybrid liquid/air-cooled WaveRouter system mentioned earlier, improving the cooling efficiency of pluggable optical modules. However, the dry/sliding interface between the plug and the riding coldplate introduces a substantial thermal penalty, which considerably degrades the cooling performance.

We have developed a solution to integrate a liquid coldplate into the body of a pluggable optical module. This integration involves replacing the module's case-top with a coldplate equipped with small-size, dripless, blind-mate liquid connectors, eliminating the need for an external heatsink or coldplate, i.e. eliminating the dry interface. The primary challenge however is the limited space available for establishing reliable liquid connections/disconnections between the pluggable optical module and the coolant lines on the host board.

This technology has been successfully demonstrated on an OSFP module, which conventionally features an integrated heatsink (Fig. 3-a). By replacing the integrated heatsink with an integrated coldplate equipped with liquid connectors, the standard height of an OSFP module is also maintained (Fig. 3-b). To achieve dripless, blind-mate operation, specifically engineered liquid connectors were developed, which have undergone rigorous testing at Ciena and proven to be a robust and reliable solution. The integrated liquid-cooled plug is capable of cooling 55W plugs in an ASHRAE W4 environment, and 85W plugs in an ASHRAE W1 environment. The addition of micro-fins to the integrated coldplate can further improve the cooling capability of this technology.

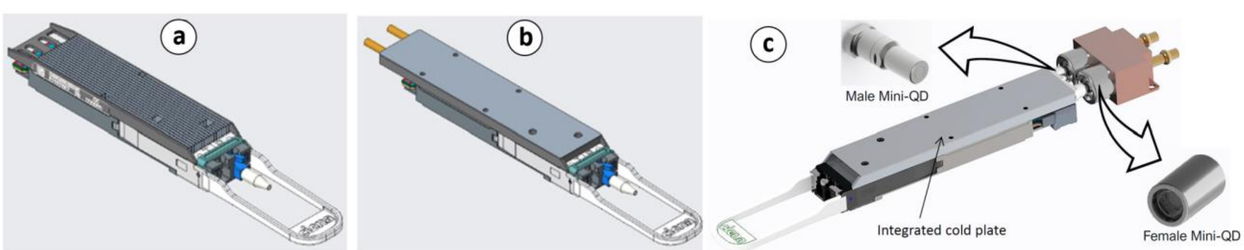

Figure 3: a) Air-cooled OSFP, b) Liquid-cooled OSFP, c) Liquid-cooled OSFP liquid connector assembly

5. Conclusion

The adoption of liquid cooling technology, particularly in the realm of optical networking, not only optimizes heat dissipation and temperature control but also brings significant energy savings in the context of the facility's HVAC system. Through this paper, we demonstrated the feasibility of upgrading an all-air-cooled system into a hybrid liquid/air-cooled configuration where key components were targeted for liquid-cooling and low-power components were air-cooled at low fan speeds. This upgrade provides the ability to use higher-power ASICs and higher-power pluggable optical modules, a significant reduction in fan power consumption and acoustic noise, as well as a meaningful reduction in facility power consumption. Furthermore, we introduced a method to bring liquid cooling directly into the body of pluggable optical modules, effectively addressing the thermal constraints of such small form factor pluggable optics.