h

Scalable Liquid Cooling Solutions for High-Powered AI GPUs:

Evaluating Nexalus Technology on NVIDIA's Blackwell Architecture

Dr. Richard Jenkins, Dr. Anthony Robinson
11-Nov-2024

1. Introduction

With the continued rise in AI adoption and rollout, GPUs, the computing engines of AI, will continue to increase in power demand and die size. As a result, cooling solutions must also evolve to handle higher thermal loads and changing die architectures. At Nexalus, AI is a core focus, with a dedicated research and engineering team working to anticipate future needs and requirements in this evolving space.

In answer to this AI-driven GPU shift, Nexalus has developed the SoloFlux liquid-cooled cold plate heat exchanger, which boasts unprecedented thermal-hydraulic performance in a compact single-slot form factor.

With the launch of the highly anticipated NVIDIA Blackwell platform - enabling organizations to build and run real-time generative AI on trillion-parameter large language models - this white paper investigates the cooling potential of the Nexalus SoloFlux on this new GPU platform. Our estimation is based on extrapolation of measured thermal performance data obtained from experiments carried out on the NVIDIA GeForce RTX 4090, giving a realistic indicator of the performance potential of the SoloFlux for the expected power and die size of the NVIDIA Blackwell GPUs.

2. Objectives

The key objective of this report are as follows:

i. Develop a straight-forward yet accurate calculation methodology to extrapolate to larger die sizes and overall thermal power loading based on measured heat transfer performance of the SoloFlux cold plate on the NVIDIA GeForce RTX 4090,

ii. Assess the capability of the SoloFlux to meet the cooling demands of future GPUs, particularly for high-power applications such as AI training.

iii. Apply these methodologies to the NVIDIA Blackwell GPU, given the known die size of 1600 mm² and scaling the power assuming the same heat flux levels of the NVIDIA GeForce RTX 4090.

3. Methodology

Ohm's law analogy to heat transfer provides a straight-forward yet well-established means to estimate the cooling performance of heat exchangers. For the cases considered here, this method can be applied to the estimation of total thermal power threshold of the GPU, $Q_{GPU,max}$, that can achieve a specified maximum junction temperature, $T_{j,max}$;

$$Q_{GPU,max} = \frac{T_{j,max} - T_{\infty}}{R_{net}} = \frac{T_{j,max} - T_{\infty}}{\left(\frac{R''_{net}}{A_{die}}\right)} \qquad [1]$$

where T_{∞} is the average coolant temperature and R_{net} is the net source-to-sink thermal resistance, which includes the influence of all materials and interfaces in the package between the cores and the coolant (e.g. die, thermal interface material, transverse and lateral conduction, convection to coolant etc.). The term, $R''_{net} = R_{net}A_{die}$, is the thermal resistivity, and is an indicator of the overall cooling intensity at the die surface. In this first-stage calculation methodology, it is assumed that this level of cooling intensity, thus the thermal resistivity achieved with the Nexalus SoloFlux and NVIDIA GeForce RTX 4090 design, can be reasonably achieved for the NVIDIA Blackwell GPU. A further necessary simplifying assumption is that the die-level heat flux of the Blackwell GPU will be commensurate with the NVIDIA GeForce RTX 4090, such that area-based scaling of the die can be employed to estimate the GPU power thresholds;

$$Q_{GPU,max} = Q''_{die}A_{die} \qquad [2]$$

where $Q''_{die} = Q_{4090}/A_{die,4090} = 69\ W/cm^2$, is the heat flux calculated from the GPU power measurements and die area of the NVIDIA GeForce RTX 4090.

4. NVIDIA GeForce RTX 4090 Experimental Results

Experiments were performed on the NVIDIA GeForce RTX 4090 fitted with a Nexalus SoloFlux. The test facility included a pumped liquid flow loop with variable and monitored volumetric flow rate, inlet coolant temperature control, and inlet/outlet coolant temperature measurements. Tests were performed under varying flow rate conditions between 1L/min and 4.8 L/min, and the working fluid used in all tests was Nexalus FluidX, which is primarily water based, though includes necessary corrosion and microbial inhibitors. The NVIDIA GeForce RTX 4090 was stressed using FurMark and resulted in a total graphics card power draw of ~450 W, with the majority, Q_{GPU}= 415 W, associated with the GPU.

Table 1 lists the key experimental measurements along with the net core-to-coolant thermal resistance and resistivity, which were calculated as;

$$R_{net} = \frac{T_{j,max} - T_{\infty}}{Q_{GPU}} = A_{die}R''_{net} \qquad [3]$$

Here, $T_{j,max}$ is taken as the GPU hotspot temperature, which is the highest measured point on the die.

Table 1: Thermal results for NVIDIA GeForce RTX 4090 fitted with a Nexalus SoloFlux cold plate heat exchanger for varying coolant flow rate.

Flow Rate (L/min)	Water Inlet Temperature, $T_{c,i}$ (°C)	Water Outlet Temperature, $T_{c,o}$ (°C)	Mean Core Temperature, T_j (°C)	Hotspot Temperature $T_{j,max}$ (°C)	Net Thermal Resistance, R_{net} (K/W)	Net Thermal Resistivity, R_{net} (cm^2K/W)
4.8	20.0	22.1	36.0	42.3	0.051	0.31
4.0	20.0	22.8	36.8	43.0	0.052	0.31
3.0	20.0	23.7	38.0	44.6	0.055	0.33
2.0	20.0	25.1	40.1	47.4	0.060	0.36
1.0	20.0	28.5	44.9	53.8	0.071	0.43

As Figure 1 shows, the thermal resistance and resistivity decreases with coolant volumetric flow rate due to the increased convective cooling intensity. However, the influence of flow rate diminishes at higher flow rates, beyond approximately 4 L/min, showing a plateauing type of behaviour. This effect is well-known for convective cooling of electronic packages and is due to a combination of a like behaviour of the convective heat transfer coefficient, as well as the other thermal resistances in the package stack (die, thermal interface material etc.) taking up a proportionately larger contribution to the net thermal resistance as the convective-side cooling improves.

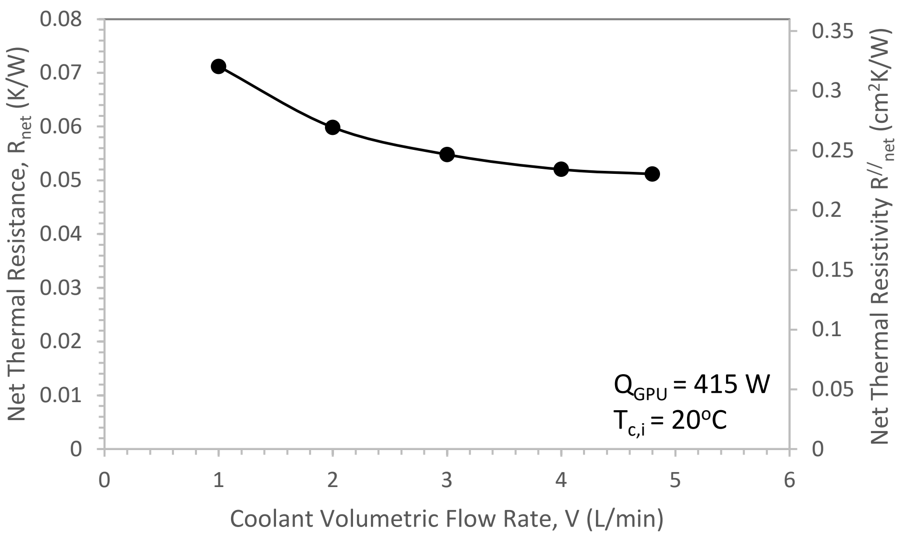

Figure 1: Net core-to-coolant thermal resistance and resistivity for NVIDIA GeForce RTX 4090 fitted with a Nexalus SoloFlux cold plate for varying coolant flow rate.

5. Maximum GPU Power Estimation for the NVIDIA GeForce RTX 4090 with Nexalus SoloFlux cold plate

The NVIDIA GeForce RTX 4090 can operate up to $T_{j,max}$ = 90°C [1]. Inserting this into Eq. 1, a first approximation of the maximum power achievable for this GPU can be calculated for varying levels of average coolant temperature, T_∞, which is subsequently plotted in Figure 2. For these estimations, the thermal resistance associated with a volumetric flow rate of 4 L/min (Table 1) are used as an illustrative example.

As the figure indicates, the Nexalus SoloFlux cold plate performance is quite exceptional, with a maximum GPU power in the region of 1.3 kW likely achievable for the NVIDIA GeForce RTX 4090 with approximately room-temperature coolant. Moreover, even at higher coolant temperatures up to 60°C, which become relevant for heat recovery and reuse applications, the Nexalus heat exchanger can manage over 40% increase in the GPU power for the same die size. Although these are first approximation calculations, and other factors such as sensible heating of the coolant at higher power levels and semiconductor thermal conductivity variations with temperature will have some degree of influence, these effects will be relatively minor and not significantly alter the main findings here.

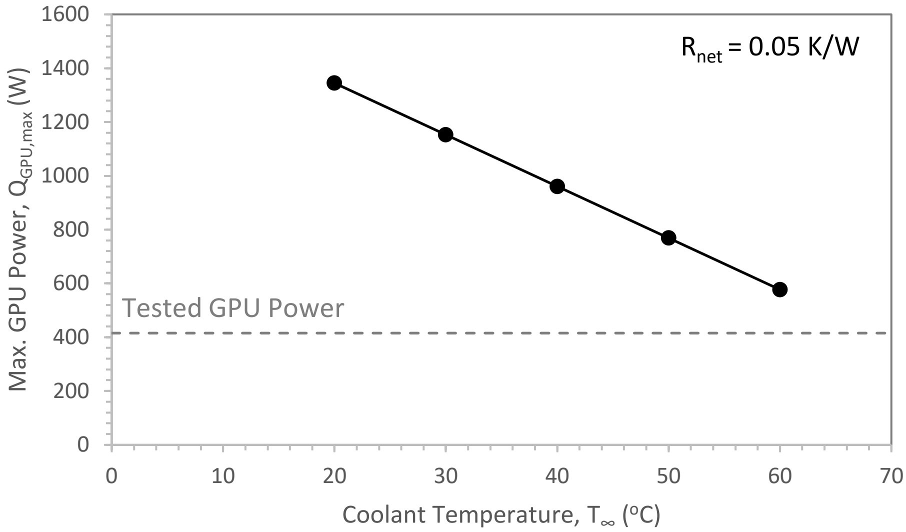

Figure 2: Approximate maximum GPU power achievable for a GPU temperature limit of 90°C for the NVIDIA GeForce RTX 4090 cooled by a Nexalus SoloFlux cold plate, for coolant flow rate of 4 L/min.

6. Blackwell GPU Specifications

The NVIDIA Blackwell GPU is specified to have a total die size of 1600 mm² (two 800mm² dies) [2], which is approximately 2.6 times larger than that of the GeForce RTX 4090 GPU (609 mm² [3]). Assuming the die-surface heat flux to be identical, the Blackwell GPU power can be estimated to be Q_{GPU}= 1100 W, which is in line with anticipated power consumption being beyond the 1 kW level. Thus, it seems a reasonable assumption that the Blackwell GPU heat flux will be of similar intensity to that of the GeForce RTX 4090 GPU. It will therefore be used here to make preliminary predictions with regard to the GPU power limit of this GPU architecture, in a similar way as had been done in the section above for the NVIDIA GeForce RTX 4090.

It will be noted here that to retain accuracy of this prediction method, it is also required that the coolant volumetric flow rate is also scaled with the GPU die surface area, in the sense that the convective cooling intensity is related to the volumetric flux of fluid. This has been taken into account in Figure 3, which shows the predicted maximum Blackwell GPU power for $T_{j,max}$ = 90°C, over a range of volumetric flow rates and coolant operating conditions.

As Figure 3 shows, the predicted maximum GPU power for the NVIDIA Blackwell cooled by a Nexalus SoloFlux cold plate generally remains significantly higher than the anticipated total power of ~ 1 kW over the range of flow rates and coolant temperatures. Choosing what can be considered a practical flow rate of 5 L/min, $Q_{GPU,max}$ = 3.1 kW, 2.2 kW, & 1.3 kW for the respective coolant temperatures of 20°C, 40°C and 60°C. Doubling the flow rate to 10 L/min further increases the estimated achievable power to 3.5 kW, 2.5 kW and 1.5 kW, respectively.

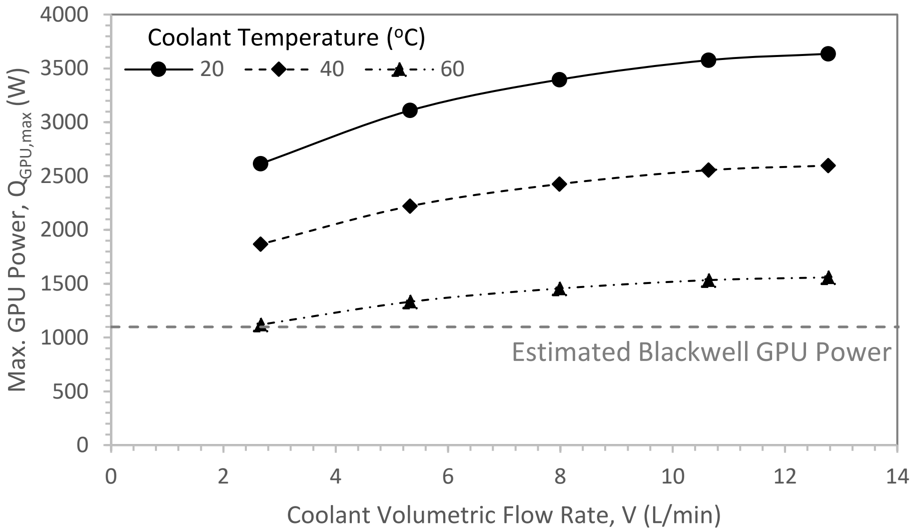

Figure 3: Approximate maximum GPU power achievable for a GPU temperature limit of 90°C for the NVIDIA Blackwell GPU cooled by a Nexalus SoloFlux cold plate.

7. Synthesis and Outlook

This report details a straight-forward yet accurate methodology to estimate the heat transfer performance of the Nexalus SoloFlux cold plate if it were to be deployed in GPU cooling scenarios that were outside of the range that it has currently been tested. Albeit an approximate method, it relies on well-established heat transfer modelling strategies, to the extent that there is confidence in the extrapolated performance results.
The prediction method was then applied to two scenarios where the GPUs are maintained at a maximum temperature of 90°C; one to predict the maximum GPU power achievable with the existing NVIDIA GeForce RTX 4090, and for the new NVIDIA Blackwell GPU series.

It is expected that the Nexalus SoloFlux provides exceptional cooling capacity with significant headroom beyond the design power of both GPUs. Specifically, for the Blackwell GPU die architecture, it is projected that over 3 kW of thermal power can be effectively managed with the SoloFlux cold plate, and done so at practical coolant operating temperatures and flow rates. The main conclusion is that the Nexalus SoloFlux cold plate design is very well positioned to manage the evolving GPU design technology space, where escalating power demands are anticipated. Importantly, the headroom is sufficiently high, such that it ensures that thermal management by Nexalus will not be the bottleneck that limits GPU performance evolution. This is of course critical for the future of generative AI.

[1] NVIDIA Corporation. (n.d.). GeForce RTX 4090 Graphics Card. NVIDIA. Retrieved from https://www.nvidia.com/en-eu/geforce/graphics-cards/40-series/rtx-4090/

[2] EEE Spectrum. (n.d.). NVIDIA Blackwell GPUs Are Revealed. Spectrum IEEE. Retrieved from https://spectrum.ieee.org/nvidia-blackwell

[3] TechPowerUp. (n.d.). GeForce RTX 4090 Ti Specifications. TechPowerUp. Retrieved from https://www.techpowerup.com/gpu-specs/geforce-rtx-4090-ti.c3917

www.nexalus.com

CONTACT US:

Nexalus Ltd,
Unit 13, South Bank,
Crosse's Green,
Cork, Ireland,
T12 XT71

Nexalus Labs,
Water St,
Loughanalla,
Castlepollard,
Co. Westmeath, Ireland,
N91 EP90

Nexalus Research,
Trinity Research & Innovation,
O'Reilly Institute, Trinity College,
Dublin 2, Ireland,

info@nexalus.com
www.nexalus.com

Registered in Ireland No. 628880,
VAT No. IE3573607SH

WITH ENDORSEMENTS AND SUPPORT FROM:

Enterprise Ireland / Taighde Éireann Research Ireland / Trinity College Dublin / CONNECT

QUALITY CERTIFICATIONS:

IATF 16949 / Automotive Standard
ISO 13485 / Medical Device Standard
ISO 50001 / Energy Management Standard
ISO 45001 / Occupational Heath & Safety Standard
ISO 14001 / Environmental Standard

NEXALUS
COOLING THE CLOUD