pulsed-heat-pump-data-center-cooling-by-infinity-turbine

Infinity Turbine LLC

12/3/2024

608-238-6001 [TEL]

greg@infinityturbine.com [Email]

Pulsed Heat Pump Data Center Cooling by Infinity Turbine

This webpage QR code

Structured Data

```
<script type= "application/ld+json">
{"@context":"http://schema.org",
"@graph":[
{
"@type" : "Organization",
"@id" : "https://infinityturbine.com/#organization",
"name" : "Infinity Turbine LLC",
"url" : "https://infinityturbine.com",
"sameAs":
["https://www.youtube.com/channel/UCsobpvy0xqc13uvhA71Cv4w",
"https://x.com/InfinityTurbine",
"https://www.instagram.com/infinityturbine/"],
"telephone" : "608-238-6001",
"email" : "greg@infinityturbine.com",
"logo" : "https://infinityturbine.com/logo.png"
},
{
"@type":"WebSite",
"@id":"https://infinityturbine.com",
"url":"https://infinityturbine.com",
"name":"Pulsed Heat Pump Data Center Cooling by Infinity Turbine",
"description":"Discover Infinity Turbine's revolutionary Pulsed Supercritical CO2 Heat Pump, a
cutting-edge cooling solution with no moving parts. Efficiently cool data center GPUs using
supercritical CO2 technology for unmatched energy savings and performance."
}
,{
"@type":"NewsArticle",
"mainEntityOfPage":{
"@type":"WebPage",
"@id":"https://infinityturbine.com/pulsed-heat-pump-data-center-cooling-by-infinity-turbine.html"},
"headline":"Pulsed Heat Pump Data Center Cooling by Infinity Turbine",
"image":"https://infinityturbine.com/images/infinity-turbine-pulsed-heat-pump-76.png",
"datePublished":"2024-12-03T08:00:00+08:00",
"dateModified":"2024-12-03T09:20:00+08:00",
"author":{
"@type":"Organization",
"name":"Infinity Turbine LLC",
"url":"https://infinityturbine.com"
},
"publisher":{
"@type":"Organization",
"name":"Infinity Turbine LLC",
"logo":{
"@type":"ImageObject",
"url":"https://infinityturbine.com/logo.png"
}}}


]}</script>
```

Discover Infinity Turbine's revolutionary Pulsed Supercritical CO2 Heat Pump, a cutting-edge cooling solution with no moving parts. Efficiently cool data center GPUs using supercritical CO2 technology for unmatched energy savings and performance.

PDF Version of the webpage (first pages)

https://infinityturbine.com/pulsed-heat-pump-data-center-cooling-by-infinity-turbine.html

12/3/2024

Revolutionizing Data Center Cooling: Using Supercritical CO2 for Efficient Cooling of Nvidia A100 GPUs

As data centers grow in size and computing power, efficient cooling systems are becoming more critical than ever. High-performance GPUs, like the Nvidia A100, generate significant amounts of waste heat, creating the need for more advanced thermal management systems. Traditional cooling methods such as air and water cooling are becoming less efficient as server densities increase.

A new concept leveraging supercritical CO2 (sCO2) for cooling Nvidia A100 GPUs offers an innovative solution by harnessing the waste heat from the GPUs and providing substantial cooling through pressure drop expansion. This article explores the potential benefits of this technology, the calculations behind it, and how it can result in significant energy savings for data centers, particularly those with high-performance workloads.

12/3/2024

The Problem with Traditional Cooling Systems

In conventional air-cooled data centers, large fans and air conditioning units are required to keep GPUs at safe operating temperatures. This method, while simple, becomes inefficient as the number of servers increases. Similarly, water cooling systems, though more efficient than air cooling, involve complex plumbing and maintenance issues and require substantial amounts of energy to pump water through the system. As data centers scale up, these cooling systems become increasingly expensive to operate and maintain.

A typical Nvidia A100 GPU consumes 300 watts (W) of power, and most of this energy is converted into heat that must be managed effectively to prevent overheating and ensure stable performance. In a data center with 50,000 Nvidia A100 GPUs, this generates 15 megawatts (MW) of waste heat that needs to be dissipated.

Harnessing Waste Heat with Supercritical CO2 Cooling

The supercritical CO2 cooling concept offers a more efficient way to manage this heat by capturing it and utilizing it for further cooling. CO2 is heated to its supercritical state (above 31°C and 73.8 bar) where it behaves like both a gas and a liquid, making it highly efficient at transferring heat.

Here’s a breakdown of how the system works:

1. Liquid CO2 enters a heat exchanger where it absorbs the waste heat from the Nvidia A100 GPUs.
2. As it heats up, the CO2 transitions into its supercritical phase, where it can absorb even more heat.
3. The CO2 is then passed through an ejector pump, creating a pressure drop that rapidly cools the GPUs. This cooling effect takes advantage of the Joule-Thomson effect, where expanding CO2 experiences a significant temperature drop.

The efficiency of this process can be illustrated with the following calculations.

Calculating the Waste Heat and Cooling Efficiency

Each Nvidia A100 GPU generates about 300 watts of heat, and we’ve already calculated that:

- Total power consumption for 50000 GPUs:
50000 GPUs times 300 W = 15000000 W (15 MW)

This equates to 15 megawatts of waste heat generated by the GPUs.

- Waste heat per GPU in BTUhour:

300 W times 3.412 BTUhour per W = 1023.6 BTUhour

For 50000 GPUs the total waste heat is:

1023.6 BTU/hour per GPU x 50000 GPUs = 51,180,000 BTU/hour

This is the amount of heat that needs to be managed.

Cooling with CO2 Expansion

Maintaining Efficiency in Data Center Cooling: How the Ejector Pump Prevents Equilibrium in the Pulsed Supercritical Heat Pump

In the rapidly evolving world of data centers, cooling high-performance GPUs like the Nvidia A100 is a critical challenge. As computational power grows, so does the need for more efficient, reliable cooling systems that can handle the intense heat generated by these powerful chips. Enter the Pulsed Supercritical Heat Pump, an innovative cooling technology that uses supercritical CO_2 to manage waste heat in data centers more effectively.

12/3/2024

How the Ejector Pump Prevents Equilibrium in the Pulsed Supercritical Heat Pump

One of the most important components of this system is the ejector pump, which ensures the cooling process remains efficient and continuous. Without this vital element, the system could potentially reach equilibrium, where temperature and pressure balance out, reducing its ability to dissipate heat effectively. This article explores how the ejector pump plays a critical role in preventing equilibrium in the Pulsed Supercritical Heat Pump, ensuring the system operates at peak efficiency.

What Is the Pulsed Supercritical Heat Pump?

The Pulsed Supercritical Heat Pump is a revolutionary cooling system designed to handle the extreme heat generated by modern data centers, particularly those using GPUs like the Nvidia A100. The system works by cycling supercritical CO2 through a series of heat exchangers, where it absorbs heat from the GPUs and provides significant cooling through pressure drop expansion.

The system relies on supercritical CO2, which has both gas and liquid properties, making it an extremely efficient medium for heat transfer. When CO2 is heated above its critical temperature of 31°C (87.9°F) and critical pressure of 73.8 bar, it enters a supercritical state. In this state, CO2 can absorb large amounts of heat, making it ideal for cooling data center GPUs.

Why Equilibrium Is a Potential Challenge

In a cascading cooling system, like the one used in the Pulsed Supercritical Heat Pump, there is always the risk of the system reaching equilibrium. Equilibrium, in this context, means that the temperature and pressure gradients driving the heat transfer and cooling process could equalize over time, causing the system to lose efficiency. If the system reaches equilibrium, the CO2 will no longer be able to absorb and transfer heat effectively, reducing the system's overall cooling capability.

The Role of the Ejector Pump in Preventing Equilibrium

To prevent equilibrium and ensure continuous, effective cooling, the Pulsed Supercritical Heat Pump relies on an ejector pump. The ejector pump is a critical component because it maintains the necessary pressure drops that drive the cooling process.

Here's how the ejector pump works to prevent equilibrium:

1. Maintaining Pressure Drops

The ejector pump creates a pressure differential by forcing the supercritical CO2 through a nozzle, creating a low-pressure zone that ensures the CO2 expands rapidly. This pressure drop is key to generating the Joule-Thomson effect, where the temperature of the CO2 decreases significantly as it expands. This cooling effect is critical for the heat exchange process, preventing the system from reaching a balanced, inefficient state.

By maintaining consistent pressure drops, the ejector pump ensures that the CO2 continues to cycle through the system effectively, providing continuous cooling for the GPUs.

2. Driving Fluid Circulation

In a closed-loop cooling system, the ejector pump plays a vital role in keeping the CO2 circulating through the system. The pump draws CO2 from the high-pressure side of the system, expands it to a lower pressure, and drives the

Nvidia A100 Chip Waste Heat

The Nvidia A100 chip, a high-performance GPU widely used in data centers for AI workloads, generates a significant amount of waste heat during operation. To estimate the recoverable waste heat from an Nvidia A100 chip, we need to consider the chip's power consumption and its thermal efficiency.

12/3/2024

Key Factors to Consider

1. Power Consumption:
- The power consumption of the Nvidia A100 chip depends on the specific model and usage scenario. However, the A100 can draw around 300 watts (W) or more under full load.

2. Heat Dissipation:
- The A100 chip, like all electronic components, converts most of the electrical power it consumes into heat. Typically, the power consumed by a GPU is converted into waste heat that needs to be dissipated to maintain safe operating temperatures.

Recoverable Waste Heat Estimate
Assuming that nearly all of the electrical power is converted into heat (as is common with GPUs), the waste heat generated by a single Nvidia A100 GPU can be estimated as follows:

- Power Consumption: 300 W
- Heat Dissipation: 300 W

Conversion to BTU/hour
To express the waste heat in BTU per hour (BTU/h), we can use the conversion factor:
1 Watt = 3.412 BTUhour

For the Nvidia A100 chip:
300 W x 3.412 BTU/hour per W = 1,023.6 BTU/hour

Therefore, the recoverable waste heat from an Nvidia A100 GPU chip is approximately 1,024 BTU per hour.

Considerations for Heat Recovery
- Cooling Method: The recoverable heat depends on the cooling system used. A water-cooled or CO2-based system (such as the supercritical CO2 concept discussed earlier) could capture a large portion of this waste heat efficiently.
- Energy Recovery Efficiency: Not all waste heat may be recoverable for reuse. The efficiency of the heat recovery system (e.g., a heat exchanger or waste heat recovery unit) will determine how much of the 1,024 BTU/hour can be used for power generation or other purposes.

In ideal conditions, a well-designed heat recovery system could capture most of this heat, especially in liquid cooling or CO2 cooling systems where the heat is transferred directly from the chips to the working fluid.

Conclusion
The Nvidia A100 chip generates approximately 1,024 BTU/hour of waste heat under full load, which can be recovered using advanced cooling and heat recovery technologies. The amount of heat recovered will depend on the efficiency of the cooling system and the design of the heat recovery infrastructure in the data center.

The temperature of the waste heat from an Nvidia A100 chip

The temperature of the waste heat from an Nvidia A100 chip depends on several factors, including the thermal characteristics of the chip, the cooling system used, and how efficiently heat is dissipated. However, we can estimate the temperature of the waste heat based on typical operating conditions for high-performance GPUs like the A100.

12/3/2024

Nvidia A100 Chip Temperature

Operating Temperature of the A100 Chip

The Nvidia A100 chip is designed to operate at safe temperatures under load. Based on general specifications for high-performance GPUs, the following temperatures are typical:

- Typical Operating Temperature: Under full load, the A100 may operate at temperatures ranging from 70°C to 85°C (158°F to 185°F).
- Thermal Throttling Limit: Most high-performance GPUs, including the A100, are designed to throttle their performance if temperatures exceed 85°C to 90°C (185°F to 194°F) to prevent damage.

Temperature of Waste Heat

The actual temperature of the waste heat will depend on how the cooling system transfers heat away from the chip. However, we can assume the following general estimates based on typical cooling systems:

1. Air Cooling:
- In air-cooled systems, the waste heat will be transferred to the surrounding air via heat sinks and fans. The exhaust temperature of the air from these systems typically ranges from 40°C to 60°C (104°F to 140°F).

2. Liquid Cooling:
- In liquid-cooled systems, the heat is transferred more efficiently to the cooling liquid, resulting in cooler exhaust temperatures compared to air cooling. The liquid temperature entering the cooling loop might be around 30°C to 40°C (86°F to 104°F), and after absorbing heat from the chip, it could reach 45°C to 60°C (113°F to 140°F) before being cooled again.

3. CO2 Cooling (Supercritical CO2):
- If the supercritical CO2 cooling system is used, the working fluid (CO2) can absorb large amounts of heat while maintaining a relatively low temperature. The CO2 temperature at the chip interface might range from 45°C to 70°C (113°F to 158°F), depending on the system's efficiency and the phase of the CO2.

Factors Affecting Waste Heat Temperature

- Cooling Efficiency: More efficient cooling systems (e.g., liquid or supercritical CO2) will result in lower waste heat temperatures at the exhaust, as they are able to capture and transfer heat more effectively.
- Workload: Under heavy loads, the GPU will produce more waste heat, potentially leading to higher exhaust temperatures.
- System Design: The design of the cooling loop, including heat exchangers and radiators, will determine the final temperature of the waste heat as it is expelled from the system.

Conclusion

- The operating temperature of the Nvidia A100 chip ranges from 70°C to 85°C (158°F to 185°F).
- The temperature of the waste heat will vary depending on the cooling method, but it can be expected to range from 40°C to 60°C (104°F to 140°F) in air-cooled systems, and 45°C to 70°C (113°F to 158°F) in liquid or supercritical CO2 cooling systems.

These temperatures provide an estimate of the heat that can be recovered and managed through advanced cooling systems in a data center.

12/3/2024

CO2 Supercritical Temperature and Pressure Drop for NVIDA A100 GPU Chipset

To calculate the resulting pressure when waste heat is transferred into CO2 and how it behaves during expansion, we need to use the properties of CO2 in both its liquid and supercritical states. We can also estimate the cooling effect when the CO2 expands and the final pressure and temperature after expansion.

12/3/2024

Temperature and Pressure Drop

Step 1: Initial Conditions of CO2

Let’s assume the CO2 starts as a liquid or subcooled liquid and is heated to a supercritical state by the waste heat from the Nvidia A100 chip.

- Waste heat: Approximately 1,024 BTU/hour (as calculated earlier).
- Temperature of the waste heat: Let’s assume an initial cooling system temperature around 40°C to 60°C (104°F to 140°F), in line with typical operating temperatures in air• or liquid-cooled systems.

In a CO2 cooling system, we are particularly interested in its behavior around the critical point of CO2:

- Critical Temperature (Tc): 31.1°C (87.9°F)
- Critical Pressure (Pc): 73.8 bar (7.38 MPa or 1,070 psi)

At temperatures above 31.1°C and pressures above 73.8 bar, CO2 enters the supercritical phase where it has properties of both a liquid and a gas, making it an excellent heat transfer medium.

Step 2: Heating CO2 to Supercritical Conditions

Let’s assume we are heating liquid CO2 from an initial condition of 25°C (77°F) up to a supercritical state of 60°C (140°F). At this temperature:

- At 25°C and 5.6 MPa (56 bar), CO2 is still in its liquid state.
- As we apply waste heat, CO2 heats up and enters its supercritical phase at temperatures above 31.1°C and pressures above 73.8 bar.

Let’s assume the waste heat raises the temperature of the CO2 to 60°C (140°F). In this supercritical state, at 60°C, the pressure would typically be in the range of 80 bar to 100 bar.

Step 3: Expansion of CO2 and Resulting Pressure

When CO2 in the supercritical phase undergoes adiabatic expansion (meaning no heat is exchanged with the surroundings during expansion), the pressure drops, and the gas cools down significantly due to the Joule-Thomson effect.

The pressure after expansion depends on how much the volume is allowed to increase. If CO2 expands through a turbine or ejector in the cooling system, the pressure drop can be significant, and the resulting cooling can be used for further thermal management.

Let’s assume the CO2 expands from 80 bar to a lower pressure of 10 bar (1 MPa or 145 psi). At this pressure, CO2 undergoes significant cooling.

Step 4: Cooling Effect and Final Temperature

The cooling effect can be approximated using enthalpy drop during the expansion. Using thermodynamic tables for CO2, we can estimate the final temperature and cooling effect:

- Before expansion: At 60°C and 80 bar, the enthalpy of supercritical CO2 is around 420 kJ/kg.
- After expansion: At 10 bar, the enthalpy drops to around 230 kJ/kg.

12/3/2024