



<b><i>Publication Year</i></b>	2023
<b><i>Acceptance in OA</i></b>	2023-03-03T10:32:40Z
<b><i>Title</i></b>	GPU water cooling investigations for the digital correlator of the MeerKAT+
<b><i>Authors</i></b>	PODDIGHE, Antonio, MELIS, Andrea, MAHGOUB, Omer, VAN BALLA, Tyrone, CONCU, Raimondo, ISAACSON, Adam
<b><i>Handle</i></b>	<a href="http://hdl.handle.net/20.500.12386/33967">http://hdl.handle.net/20.500.12386/33967</a>

# GPU water cooling investigations for the digital correlator of the MeerKAT+

A. Poddighe<sup>1</sup>, A. Melis<sup>1</sup>, O. Mahgoub<sup>2</sup>, T. VanBalla<sup>2</sup>,  
R. Concu<sup>1</sup>, A. Isaacson<sup>2</sup>

<sup>1</sup>INAF - Italian National Institute for Astrophysics

<sup>2</sup>SARAO – South African Radio Astronomy Observatory

January 18, 2023

## Abstract

GPUs represent a valuable solution in the framework of digital signal processing in radio astronomy. Currently, GPUs have been selected as the primary option for building the digital correlator of the MeerKAT+ array, for both the F-Engine and the X/B-Engine. Thermal management is one of the major issues affecting GPUs in comparison to FPGAs and in this paper we outline how water cooling solutions may assist.

## 1 Introduction

MeerKAT is a SKA-mid precursor telescope array composed of 64 antennas of 13.5m diameter each, the facility is operational and remarkable scientific results have already been achieved. In view of the upcoming construction of SKA-mid, an improvement of MeerKAT was financed; MeerKAT+ will extend MeerKAT by a further 16 dishes, bringing the total to 80. The current SKARAB-based correlator, which was designed to be suitable for 64 elements or multiples thereof, is therefore unfit for use with MeerKAT+. As a consequence, a new digital correlator is being designed and the first version is going to be entirely GPU-based. There are currently a wide variety of GPUs available, both in terms of cost and performance, and ideally one wishes to use a model which provides the best trade-off. In this regard, the GPUs used for video games need to be offered at a reasonable price and, contextually, provide superior performance to the ones required for "normal" processing.

Regardless of factors such as the silicon shortage, crypto-mining and COVID-19 which all resulted in prices increasing tremendously, the RTX 30XX series appears to be one of the best solutions. We therefore purchased a few of the RTX 3080TI and

RTX 3090 GPUs to evaluate their capabilities. Since the building where the correlator will be located is situated in the desert of the Karoo region, it is very important to provide a reliable system in which human interventions are limited as much as possible. In order to help achieve that, the temperature of the GPUs needs to be kept low. In this paper, we provide a comparison of the air cooling and water cooling solutions for the RTX 3080TI and the RTX 3090 and highlight the improvement achieved with water cooling techniques.

## **2 Air cooling vs water cooling**

In this section, we will show all of the tests that were carried out on the GPUs with both air cooling and water cooling. The tests were conducted in a computer room with controlled temperatures, varying from 21 to 23 °C.

The chosen benchmark tool was the well-known OctaneBench, which perform roughly 10 minutes long sessions; all test were carried out two days after the power on. The original GPUs air cooling stuff were replaced with waterblocks; proper thermal pads were employed to improve the heat exchange.

### **2.1 Water cooling kits**

The water cooling solution is comprised of the following major components:

- 1x 4U rackmount case (waterstation)
- 1x SANSO PDH-E 054 IT3 12V pump (see fig. 1)
- 1x Alphacool UT60 360mm full copper radiator cooled by 6 Alphacool ES 120mm 4000rpm fans (see fig. 2)
- 1x Alphacool ES Reservoir 4U (see fig. 3)
- 1x Alphacool Flow Indicator Aurora High-Flow G1/4 - D-RGB (see fig. 4)
- 2x EK-Quantum Vector ZOTAC® Trinity RTX 3080/3090 D-RGB - Nickel + Acetal (waterblock) (see fig. 5)
- 2x EK-Quantum Vector ZOTAC® Trinity RTX 3080/3090 Backplate - Black (waterblock backplate) (see fig. 5)
- Alphacool Tubo AlphaTube TPV 12,7/7,6
- Aquacomputer additive Double Protect Ultra Clear

Fig. 6 illustrates the block diagram showing all of the components the watestation is composed of, as well as the water flow direction; fig 7 shows the assembled water cooling station employed for our tests, whilst fig. 8 lists the details of the components which make up the water cooling kits that have been used.



Figure 1: SANSO Pump




Figure 2: Alphacool Radiator



Figure 3: Alphacool Reservoir



Figure 4: Alphacool Flow Indicator


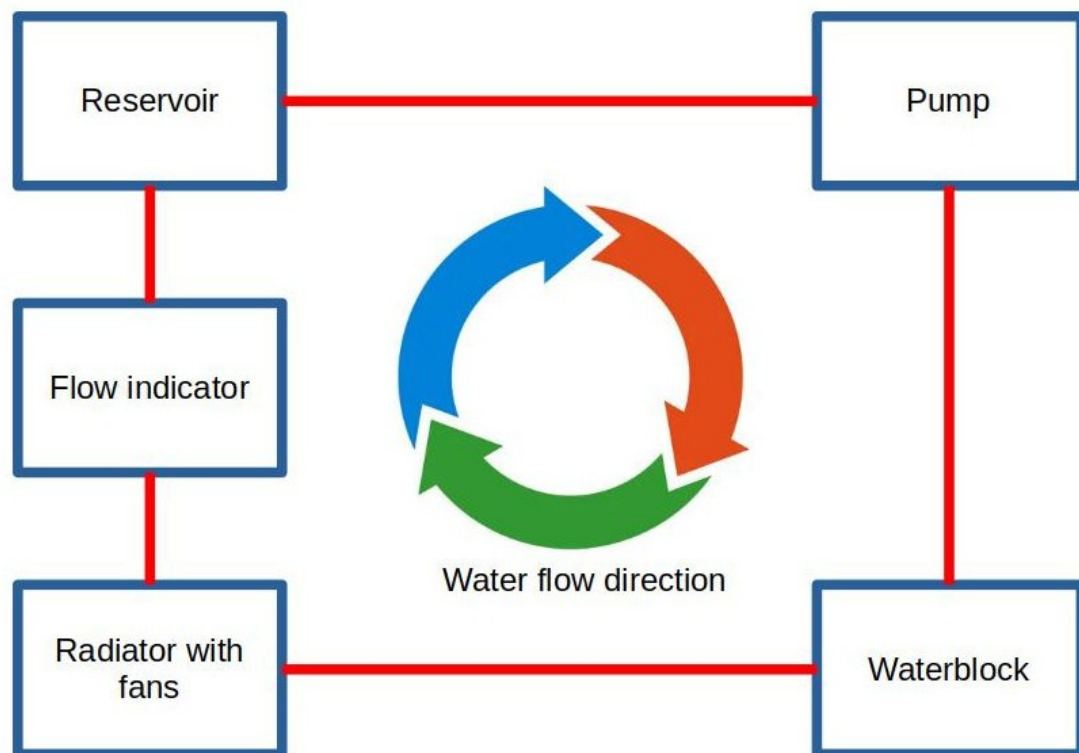


Figure 5: Alphacool Waterblock and Waterblock Backplate



*Figure 6: Water flow direction*




*Figure 7: Waterstation*



Cod. art.	Descrizione	Um	Quantità ordinata	Quantità consegnata	Quantità da consegnare
<b>Rif. Offerta 22611/2020 del 25/11/2020:</b>					
1017607	Alphacool ES 4U - 19" - ServerRack - Watercooling	nr	2		2
01260	SANSO PDH-E 054 IT3 12V Attacchi filettati	nr	4		4
1019834	Alphacool ES Reservoir 4U - Single (Bulk)	pz	2		2
1011037	Alphacool HF screw-in seal plug G1/4 - Chrome	nr	2		2
1017333	Alphacool Tappo V2 G1/4 - Deep Black	pz	2		2
EK46933	EK-AF FillPort G1/4 - Black	pz	2		2
EK46902	EK-AF Y-Splitter 3F G1/4 - Black	nr	4		4
32234	Alphacool Valvola di non ritorno F/F G1/4 - Deep black	pz	4		4
35280	Alphacool Radiatore UT60 360mm full copper	nr	2		2
1019158	Alphacool ES 120mm 4000rpm Fan ( 120x120x25mm	nr	12		12
1011350	Alphacool NexXos Viti per radiatore M3 x 30 mm - 4	nr	6		6
1013592	Alphacool Passaparete HF G1/4 Femmina/Femmina -	nr	4		4
32058	Slot Cover PCI con passaparete 1/4G F	pz	2		2
1017404	Alphacool ES Quick Release Connector Kit TPV	pz	20		20
1016781	Alphacool Raccordo HF TPV 12.7/7.6 - - Dritto - Black -	cfz	17		17
1016779	Alphacool Tubo AlphaTube TPV 12.7/7.6 - black -	cfz	8		8
EK13492	EK-Pro Collettore 2CPU 2GPU - Acetal	pz	2		2
EK02311	EK-Quantum Velocity sTR4 D-RGB - Full Nickel	nr	4		4
EK33070	EK-Quantum Vector ZOTAC® Trinity RTX 3080/3090 D-	pz	4		4
EK33087	EK-Quantum Vector ZOTAC® Trinity RTX 3080/3090	pz	4		4
30263	Aquacomputer additivo Double Protect Ultra Clear	pz	1		1
1017207	Alphacool Flussmetro Aurora High-Flow G1/4 - D-RGB	nr	4		4
81033	Phobya extension 3Pin Molex extra long 60cm -BLACK	nr	12		12
81130	Phobya Y-adaptor 4Pin Molex to 2x 4Pin PWM and 3Pin	nr	10		10
87273	Phobya 4Pin Molex power extension 60 cm - black	nr	10		10

Figure 8: List of the items that compose the water cooling kits





*Figure 9: NVidia 3080TIs with their water blocks in a daisy chain configuration*



*Figure 10: Four GPUs test*





*Figure 11: Four GPUs test*





Figure 12: Four GPUs test

## 2.2 Test sessions

All of the tests were conducted using the popular OctaneBench 2020.1 (<https://render.otoy.com/octanebench/>) benchmarking utility with the GPU's temperature read at regular 5 second intervals. The various hardware configurations tested are listed below:

- A single Supermicro server equipped with two NVidia GeForce 3080TI (air cooling)
- A single Supermicro server equipped with two NVidia GeForce 3090 (air cooling)
- A single Supermicro server equipped with two NVidia GeForce 3080TI (water cooling)
- A single Supermicro server equipped with two NVidia GeForce 3090 (water cooling)
- Two Supermicro servers equipped with two RTX3080TI and two 3090 respectively (water cooling)

The results of the air cooling tests are shown in Figures 13 and 14. As displayed, at full load both 3080TIs and 3090s have reached a temperature of about 70°C, which is approximately 50°C above ambient. Figures 15 and 16 refer to the water cooling tests, in both cases with daisy-chained water blocks (see fig. 9). These plots show that the improvement is quite significant; 25°C above ambient, which means a 50% improvement compared to air cooling. In order to figure out the order of deterioration if more than two GPUs are water cooled with the same water station, we fed the aforementioned two daisy-chained water blocks in parallel, accommodating the four GPUs in two different servers, see fig. 10, 11, 12; these results are shown in fig. 17. Despite the degradation in performance, the improvement (20°C) with respect to the air cooling techniques remains remarkable. Finally, we'd like to report some as yet unexplained and unexpected behavior which occurred when the cables of the daisy-chained water block were swapped around. As shown in fig. 18 and 19, in the swapped setup can be noticed a five degrees discrepancy between the second plot and the initial one. It can also be noticed that, in all of the plots, a slight temperature difference between the GPUs appears systematically; this is likely due to how fans and GPUs are placed in the case.




Figure 13: 3080 TI air cooled




Figure 14: 3090 air cooled






Figure 15: 3080 TI liquid cooled




Figure 16: 3090 liquid cooled




Figure 17: 3080 TI and 3090 liquid cooled




Figure 18: 3080 TI reverse liquid flow




Figure 19: 3080 TI initial liquid flow



### **3 Conclusion and remarks**

In this paper, we have described how water cooling solutions proved to dramatically improve the thermal management of the NVidia RTX 3080TIs and 3090s. Specifically, we got an improvement of at least 25°C, which represents an impressive result. Both these GPUs have an attractive performance to price ratio when compared with higher-end models, which are essentially the same but with decreased performance. On the other hand, beyond the increased cost and the additional space required, water cooling solutions come with their own logistical issues and risks due to water flowing within servers. The question as to whether or not to use water cooling solutions for the digital correlator of the MeerKAT+ radio telescope is still pending. The lifetime of a water cooled GPU is at least twice that of an air cooled one, and if done correctly, the probability of any damage is substantially reduced; moreover, the overall cost would be considerable but surely lower than using GPUs such as the NVidia A100 or similar. However, it cannot be ignored that air cooling solutions would greatly simplify the entire design, with the drawback of having to worry more about maintenance and procuring a larger number of spare parts.

Test are going to be performed again in the middle of 2023 in the desert of Karoo, results will be compared with the ones we achieved in this paper, the final decision is going to be taken soon after.