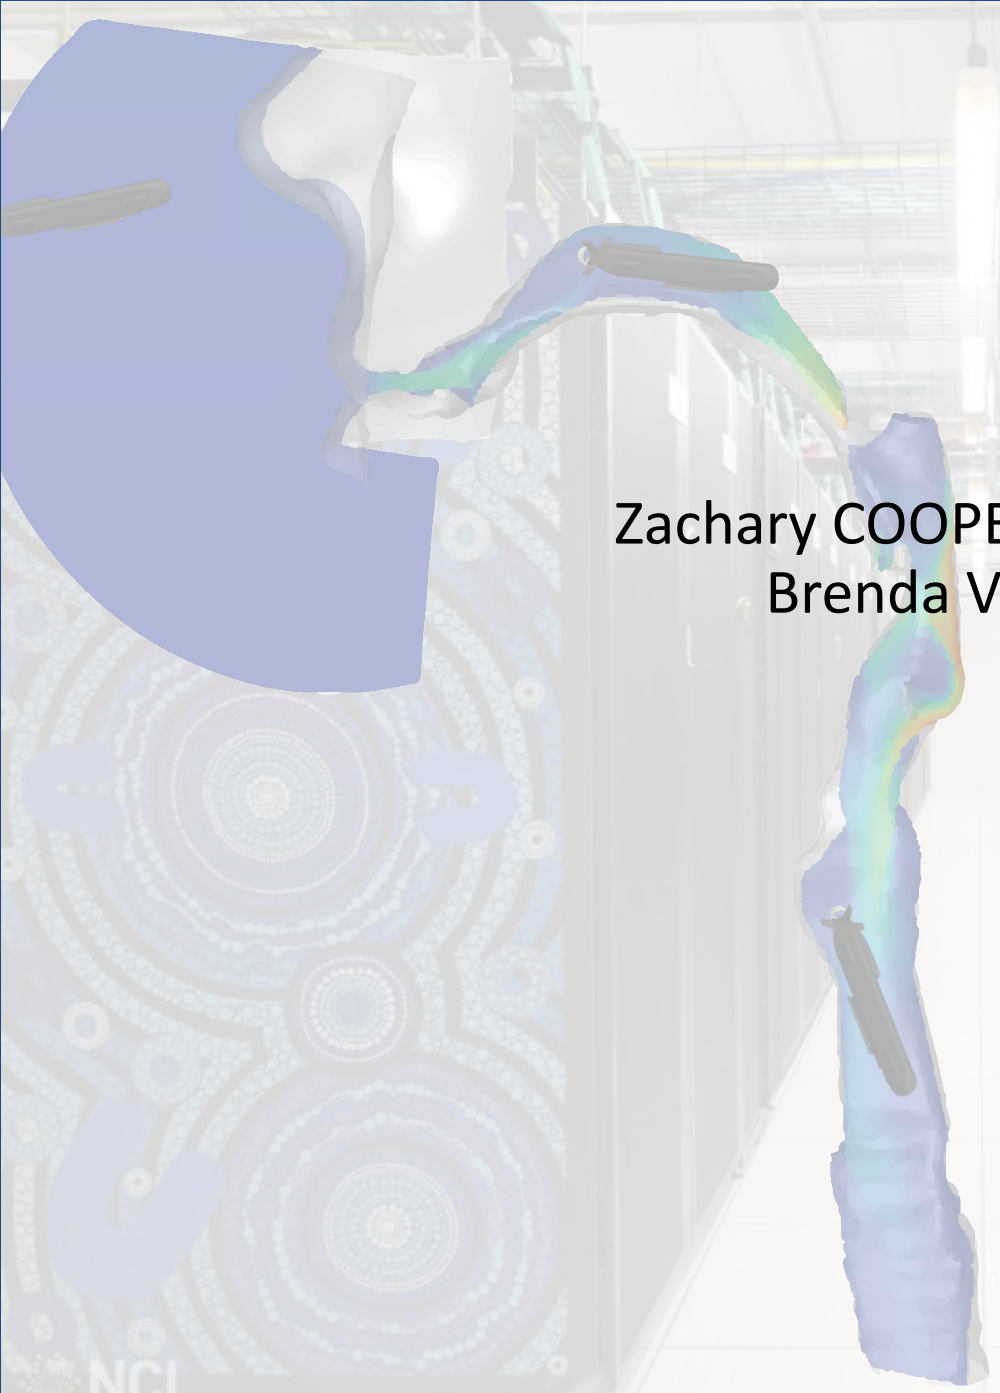

Breathing Submarines

Team Members
Zachary COOPER-BALDOCK, Flinders University
Brenda VARA ALMIRALL, RMIT University

Mentors
Frederick FUNG, ANU
Joseph JOHN, ANU)

Breathing Submarines

- **Primary problem:**
  - Run complex ANSYS fluid simulations on Gadi, accelerated by GPUs
  - Test two different models (airway tract & submarine flow)

- **Focus:**
  - Required code/flags to enable GPU use
  - Accelerate speed of computation
  - Decrease total job time
  - Decrease SU cost

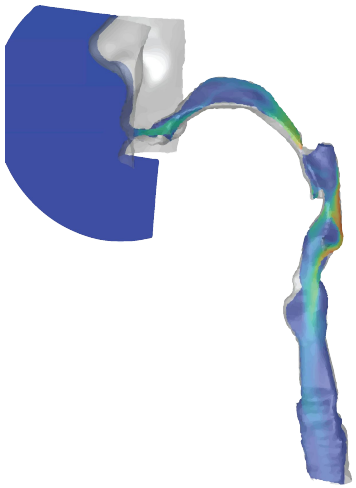

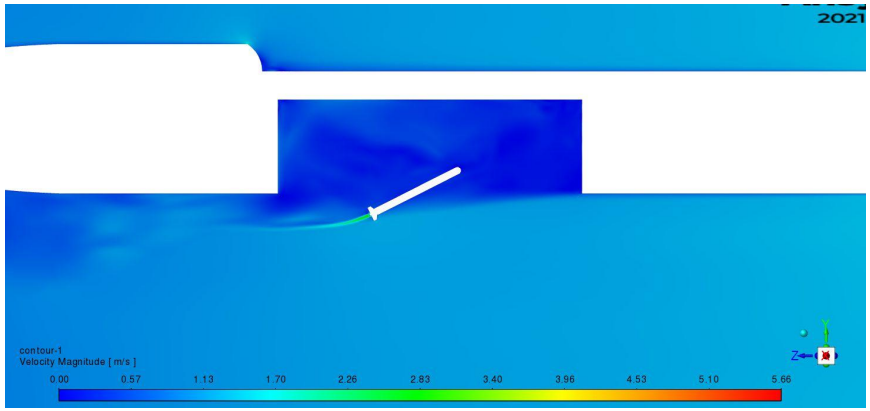

Evolution and Strategy

- **What was your goal coming here?**
  - Enable, then efficiently run, GPU simulations on Fluent
- **What was your initial strategy?**
  1. ANSYS files and documentation
  2. GADI/HPC files and documentation
  3. Optimisation of queues, CPU #, RAM #, GPU #
- **How did this strategy change?**
  - ANSYS documentation issues
  - Inconsistent flags/CLI requirements

Breathing - Energy Efficiency

The calculator will compare energy consumption of a number of CPU only nodes with dual CPUs required to perform the same amount of work as 1 GPU node with 2 CPUs and 8 GPUs.

INPUTS	
# CPU Cores	48
# GPUs (A100)	2
Application Speedup	10.0x

Node Replacement	15.0x

GPU NODE POWER SAVINGS			
	AMD Dual Rome 7742	8x A100 80GB SXM4	Power Savings
Compute Power (W)	16,500	6,500	10,000
Networking Power (W)	697	93	604
Total Power (W)	**17,197**	**6,593**	10,604

Node Power efficiency	2.6x

ANNUAL ENERGY SAVINGS PER GPU NODE			
	AMD Dual Rome 7742	8x A100 80GB SXM4	Power Savings
Compute Power (kWh/year)	144,540	56,940	87,600
Networking Power (kWh/year)	6,102	814	5,288
Total Power (kWh/year)	**150,642**	**57,754**	**92,888**

$/kWh	$	0.34
Annual Cost Savings	$	31,581.99
3-year Cost Savings	$	94,745.98

Metric Tons of CO2	66
Gasoline Cars Driven for 1 year	14
Seedlings Trees grown for 10 years	1,089

(source: Link)

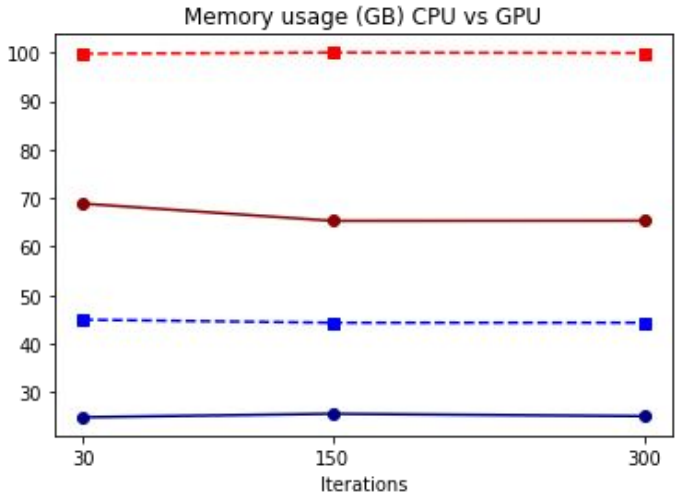

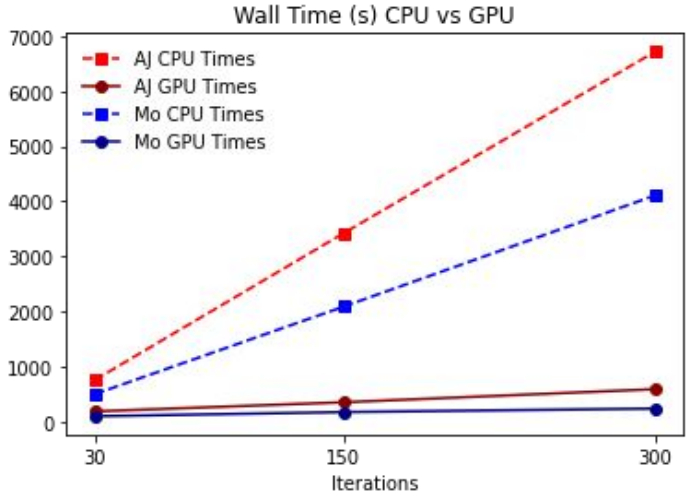

Results - Submarine

- Total calculation time of 2,664s reduced to 128s
- Single iteration time of 26.517s reduced to 1.264s
- Job cost reduced from 20.11 SU to 5.08 SU
- GPU acceleration showing speed increases* of 5x to 20x

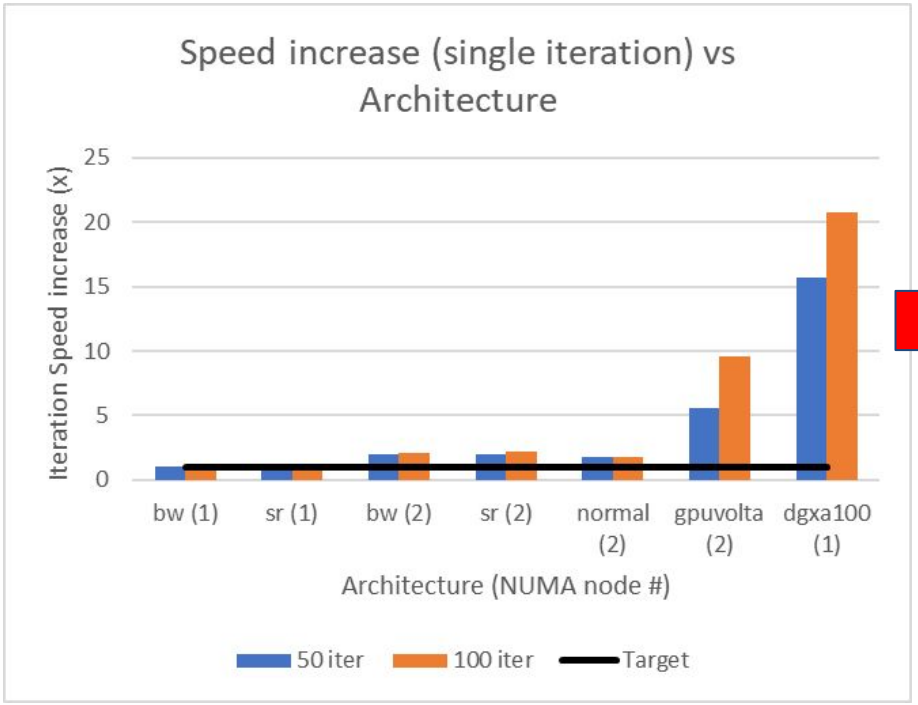

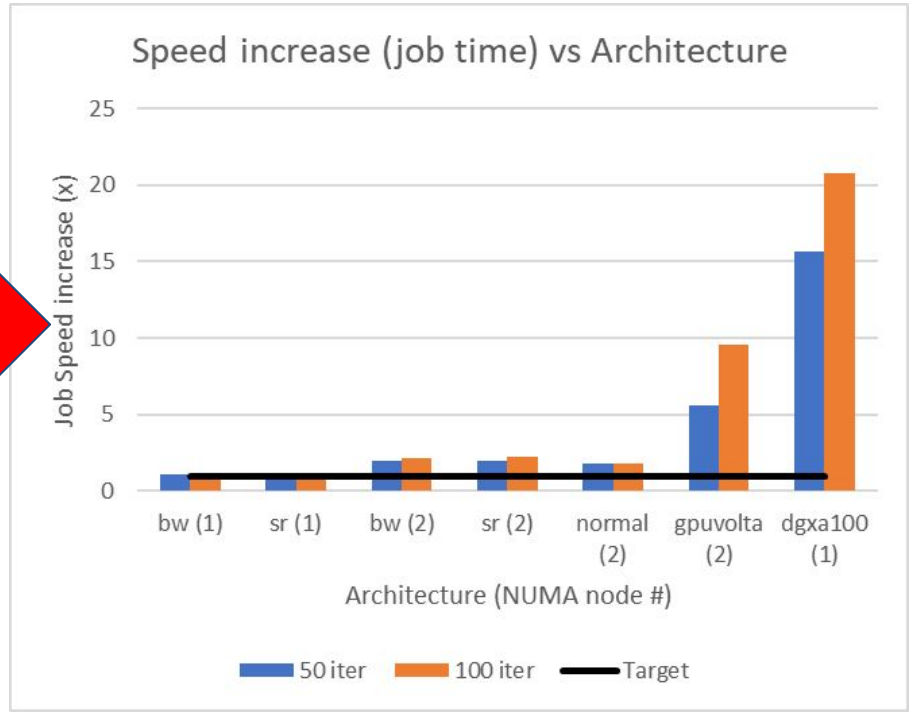

*Normalised against 1 NUMA node normalsr calculation

sr = Sapphire Rapid, bw = Broadwell, gpuvolta = V100, dgxa100 = A100

Energy Efficiency - Submarine

INPUTS	
# CPU Cores	16
# GPUs (A100)	1
Application Speedup	20.7x
Node Replacement	20.7x

GPU NODE POWER SAVINGS			
	AMD Dual Rome 7742	8x A100 80GB SXM4	Power Savings
Compute Power (W)	22,813	6,500	16,313
Networking Power (W)	963	93	870
Total Power (W)	**23,776**	**6,593**	17,183
Node Power efficiency		**3.6x**	

ANNUAL ENERGY SAVINGS PER GPU NODE			
	AMD Dual Rome 7742	8x A100 80GB SXM4	Power Savings
Compute Power (kWh/year)	199,841	56,940	142,901
Networking Power (kWh/year)	8,436	814	7,623
Total Power (kWh/year)	**208,277**	**57,754**	**150,524**

$/kWh	$ 0.34
Annual Cost Savings	$ 51,178.08
3-year Cost Savings	$ 153,534.24
Metric Tons of CO2	107
Gasoline Cars Driven for 1 year	23
Seedlings Trees grown for 10 years	1,764

(source: Link)

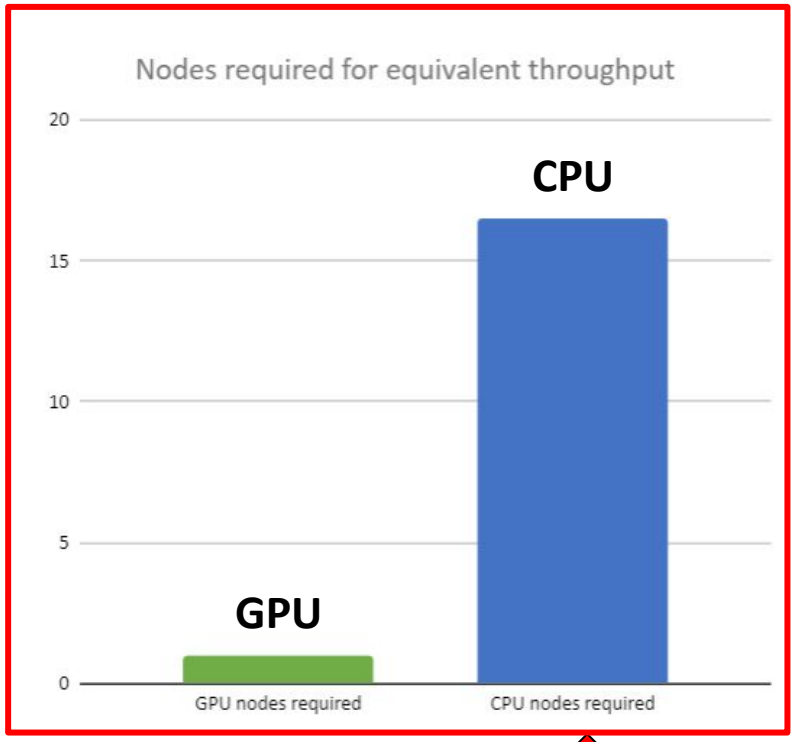

$/kWh	$ 0.34
Annual Cost Savings	$ 51,178.08
3-year Cost Savings	$ 153,534.24
Metric Tons of CO2	107
Gasoline Cars Driven for 1 year	23
Seedlings Trees grown for 10 years	1,764

What problems have you encountered?

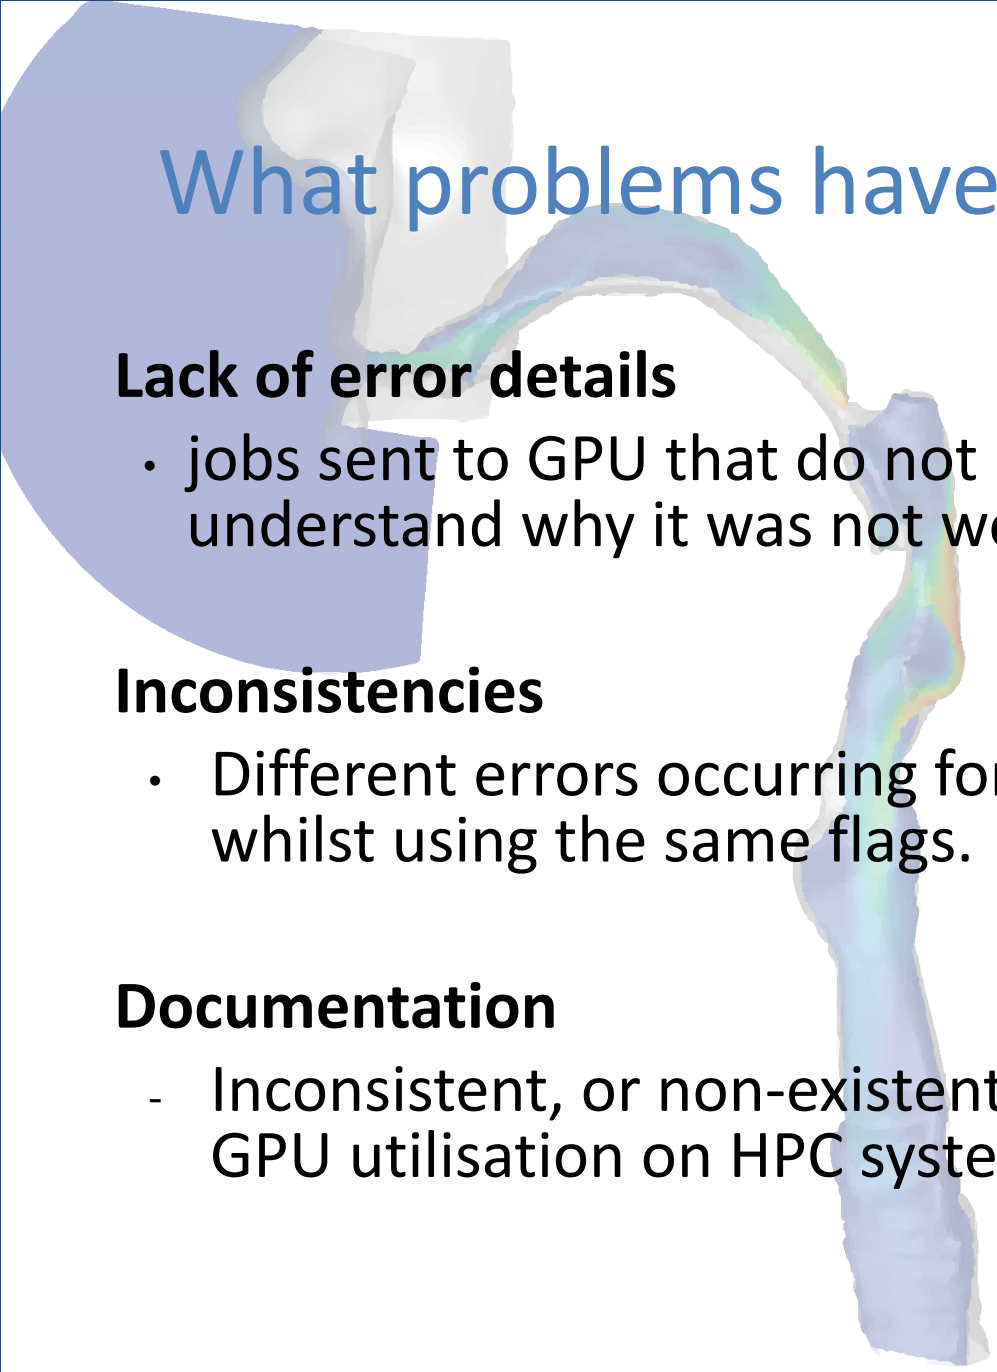

Lack of error details

- jobs sent to GPU that do not use GPU it was difficult to understand why it was not working

Inconsistencies

- Different errors occurring for different model types, whilst using the same flags.

Documentation

- Inconsistent, or non-existent, ANSYS documentation for GPU utilisation on HPC systems.

Wishlist

- **Tools** - documentation for GPU for Fluent users gpuapp flag
- **Language standards** - CLI commands, different commands in different ANSYS releases (2021/22/23)

What made this worthwhile

- We will **continue development** with simulations to run **more complex** models, specifically moving propeller and unsteady breathing conditions

- **Sustained resources/support** with team across RMIT and Flinders, NCI, ANSYS and Nvidia

100 words summary of Breathing Submarines achievements during this Hackathon

To run Ansys Fluent CFD simulations using GPU's on Gadi potentially for the first time as no one in either of our teams have been able to do this.

This means an average memory reduction of 62% and a speedup average of 10-20, that demonstrate exponential speedup increase with job size increases.