Team Explorer

Ruibiao Zhu (ANU)

Mentors
Maruf Ahmed (NCI)
Maggie Zhang (Nvidia)

Physics Simulation Using Graph Neural Network (GNN)

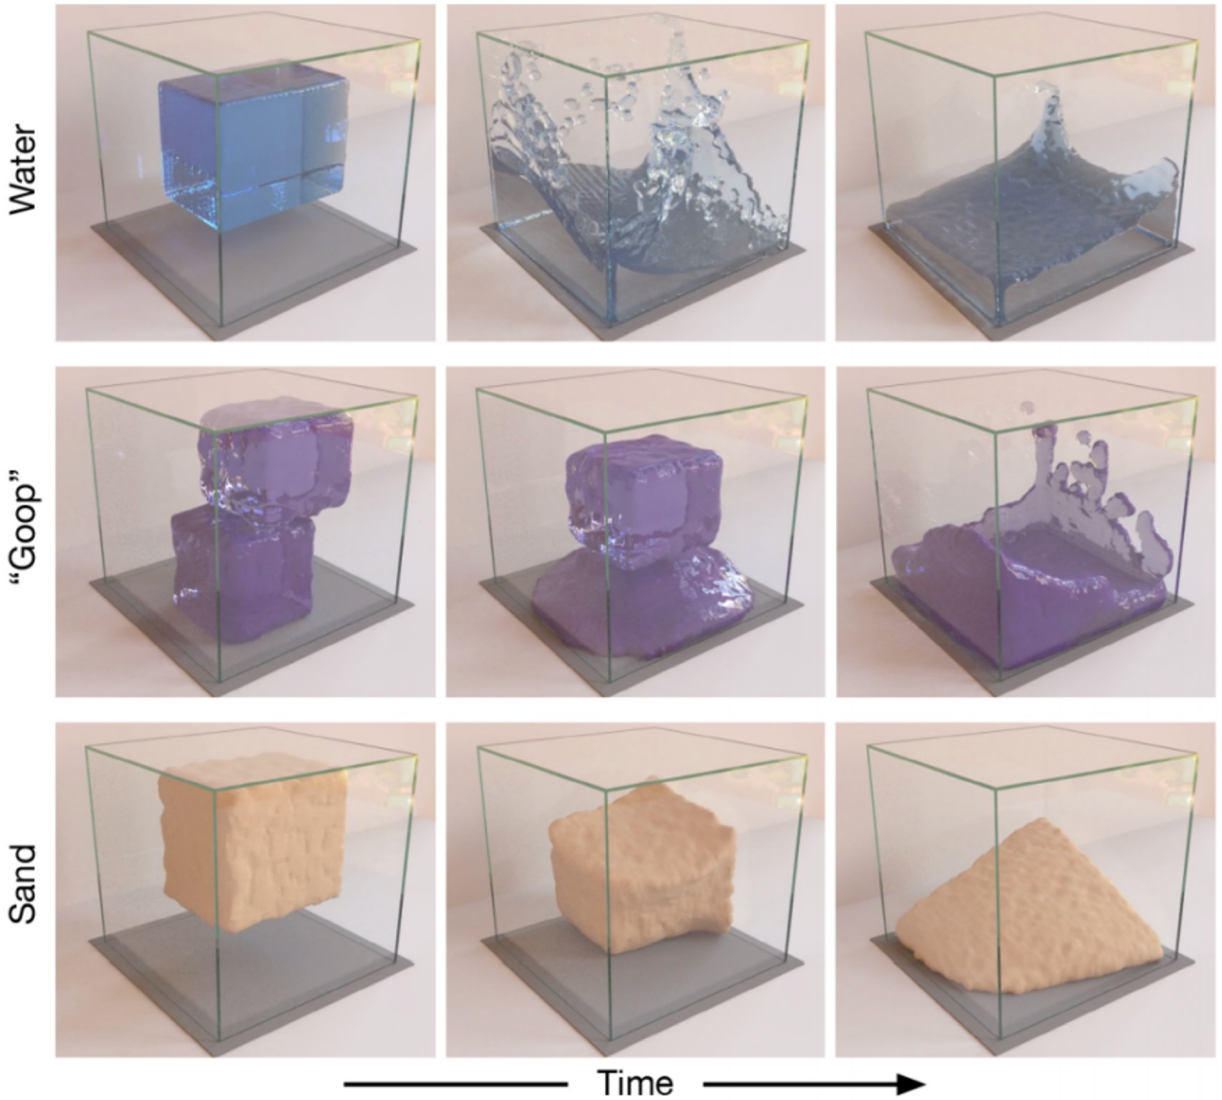

Graph Neural Network

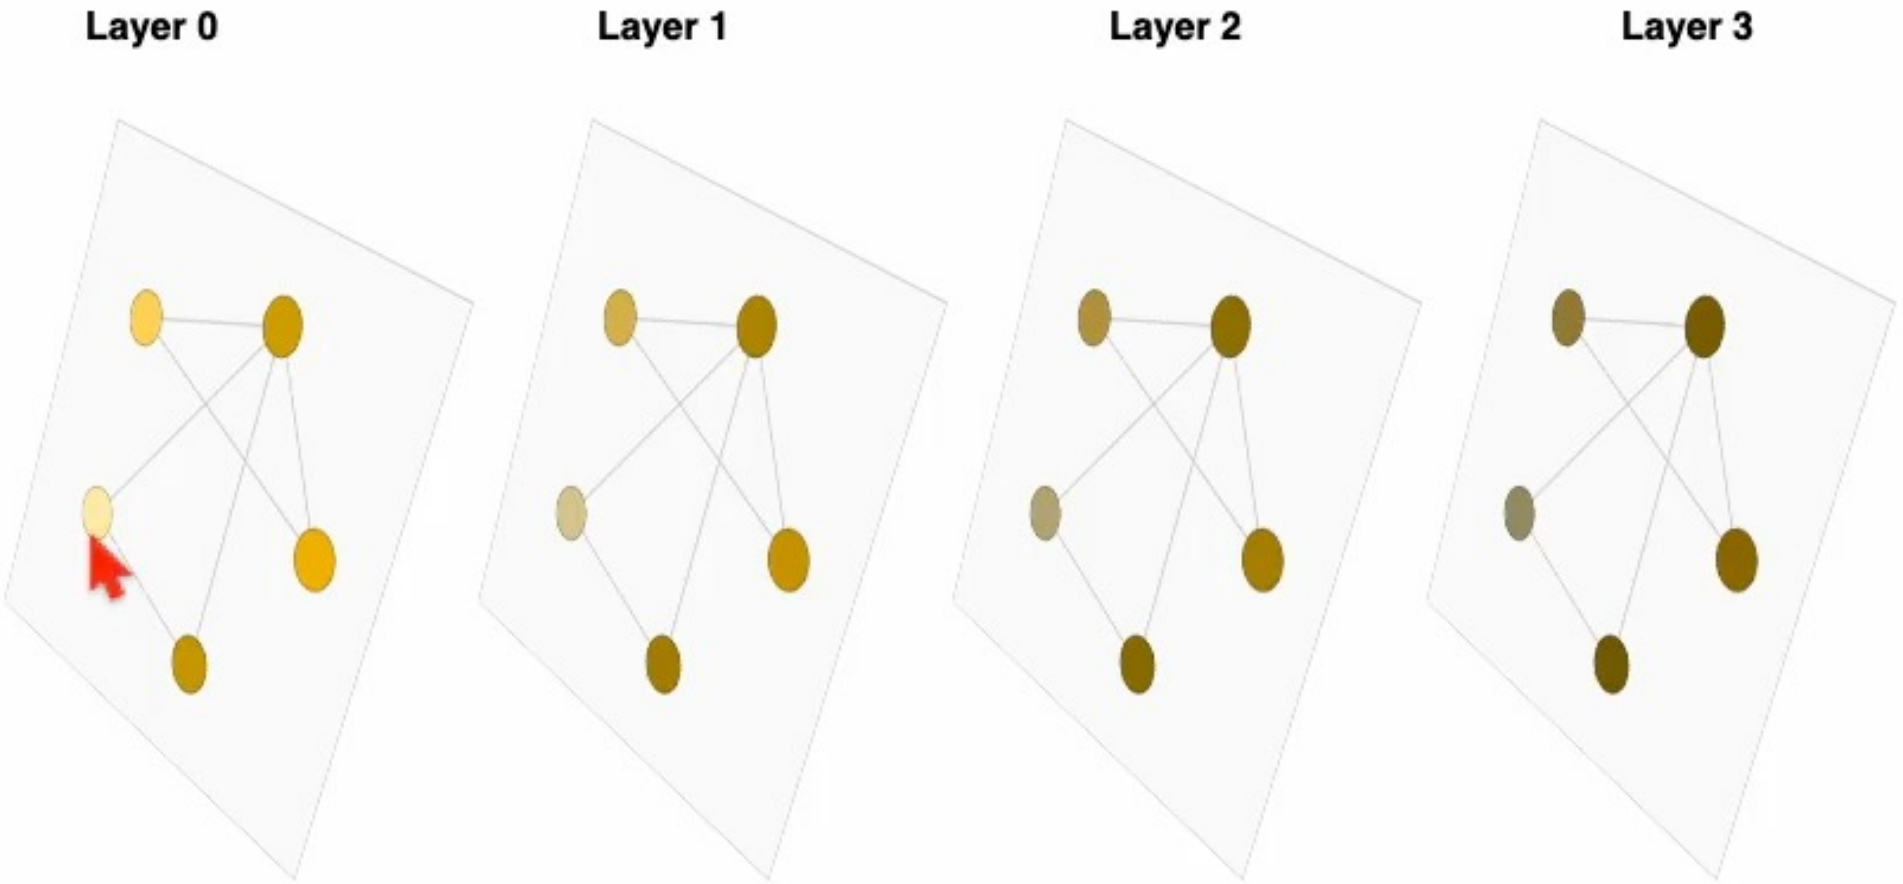

Legacy Code Ported to A100

marufahmed — mah900@gadi-gpu-v100-0109:/scratch/vp91/mah900/deepmind-research — ssh mah900@gadi.nci.org.au — 113×21

```
I0609 01:37:55.707464 22773708277248 basic_session_run_hooks.py:692] global_step/sec: 13.3664
INFO:tensorflow:loss = 0.24372683, step = 39700 (7.481 sec)
I0609 01:37:55.708581 22773708277248 basic_session_run_hooks.py:260] loss = 0.24372683, step = 39700 (7.481 sec)
INFO:tensorflow:global_step/sec: 13.2571
I0609 01:38:03.250619 22773708277248 basic_session_run_hooks.py:692] global_step/sec: 13.2571
INFO:tensorflow:loss = 0.06620553, step = 39800 (7.543 sec)
I0609 01:38:03.251495 22773708277248 basic_session_run_hooks.py:260] loss = 0.06620553, step = 39800 (7.543 sec)
INFO:tensorflow:global_step/sec: 12.7202
I0609 01:38:11.112110 22773708277248 basic_session_run_hooks.py:692] global_step/sec: 12.7202
INFO:tensorflow:loss = 0.050374176, step = 39900 (7.862 sec)
I0609 01:38:11.113271 22773708277248 basic_session_run_hooks.py:260] loss = 0.050374176, step = 39900 (7.862 sec)
INFO:tensorflow:Saving checkpoints for 40000 into /scratch/vp91/mah900/models/WaterRamps_V100/tmp/model.ckpt.
I0609 01:38:18.994716 22773708277248 basic_session_run_hooks.py:606] Saving checkpoints for 40000 into /scratch/v
p91/mah900/models/WaterRamps_V100/tmp/model.ckpt.
INFO:tensorflow:Loss for final step: 0.053839043.
I0609 01:38:20.069769 22773708277248 estimator.py:371] Loss for final step: 0.053839043.

real    54m21.031s
user    64m25.606s
sys     13m1.116s
(/scratch/fp0/mah900/env/sTF1.15) [mah900@gadi-gpu-v100-0109 deepmind-research]$
```

marufahmed — mah900@gadi-dgx-a100-0001:/scratch/vp91/mah900/deepmind-research — ssh mah900@gadi.nci.org.au — 113×20

```
INFO:tensorflow:loss = 0.12506124, step = 39700 (3.570 sec)
I0609 01:10:47.133861 22787569262912 basic_session_run_hooks.py:260] loss = 0.12506124, step = 39700 (3.570 sec)
INFO:tensorflow:global_step/sec: 27.627
I0609 01:10:50.753088 22787569262912 basic_session_run_hooks.py:692] global_step/sec: 27.627
INFO:tensorflow:loss = 0.07625614, step = 39800 (3.620 sec)
I0609 01:10:50.753701 22787569262912 basic_session_run_hooks.py:260] loss = 0.07625614, step = 39800 (3.620 sec)
INFO:tensorflow:global_step/sec: 26.6118
I0609 01:10:54.510792 22787569262912 basic_session_run_hooks.py:692] global_step/sec: 26.6118
INFO:tensorflow:loss = 0.071071, step = 39900 (3.758 sec)
I0609 01:10:54.511220 22787569262912 basic_session_run_hooks.py:260] loss = 0.071071, step = 39900 (3.758 sec)
INFO:tensorflow:Saving checkpoints for 40000 into /scratch/vp91/mah900/models/WaterRamps_A100/tmp/model.ckpt.
I0609 01:10:58.257473 22787569262912 basic_session_run_hooks.py:606] Saving checkpoints for 40000 into /scratch/v
p91/mah900/models/WaterRamps_A100/tmp/model.ckpt.
INFO:tensorflow:Loss for final step: 0.07648795.
I0609 01:10:59.554331 22787569262912 estimator.py:371] Loss for final step: 0.07648795.

real    26m41.711s
user    31m47.432s
sys     2m12.726s
(/scratch/fp0/mah900/env/sTF1.15) [mah900@gadi-dgx-a100-0001 deepmind-research]$
```

Energy Efficiency

INPUTS	
# CPU Cores	12
# GPUs (A100)	1
Application Speedup	21.2x

Node Replacement	15.9x

GPU NODE POWER SAVINGS			
	AMD Dual Rome 7742	**8x A100 80GB SXM4**	**Power Savings**
Compute Power (W)	17,498	6,500	10,998
Networking Power (W)	739	93	646
Total Power (W)	**18,237**	**6,593**	11,644

Node Power efficiency	**2.8x**

ANNUAL ENERGY SAVINGS PER GPU NODE			
	AMD Dual Rome 7742	**8x A100 80GB SXM4**	**Power Savings**
Compute Power (kWh/year)	153,285	56,940	96,345
Networking Power (kWh/year)	6,471	814	5,657
Total Power (kWh/year)	**159,756**	**57,754**	**102,002**

$/kWh	$	0.34
Annual Cost Savings	$	34,680.69
3-year Cost Savings	$	104,042.08

Metric Tons of CO2	72
Gasoline Cars Driven for 1 year	16
Seedlings Trees grown for 10 years	1,195

(source: Link)

What problems have you encountered?

- The Legacy code is written in TensorFlow 1
  - TensorFlow 1.15 is supported up to CUDA 10
  - However, modern GPUs (Ampere, Hopper, and later) Requires CUDA 11.
  - Porting legacy code and dependencies to unsupported hardware was the most challenging part.
- Run the code on Multiple GPUs

Wishlist

- Simulation is computation hungry and Full training is impossible.
- We estimate that, at least 20 million epochs are required, which would take 333.33 hours on a V100.
- We were busy porting the code to A100 and did not have the time to do the distributed training.

Was it worth it?

- Well, we made a cool animation from our work
- This 3-second prediction required half a million epochs worth of training
- Took about eight hours on an A100
- Video file:

Team Explorer-Graph Neural Network-Physics Simulation.mp4

GNN-Water motion simulation

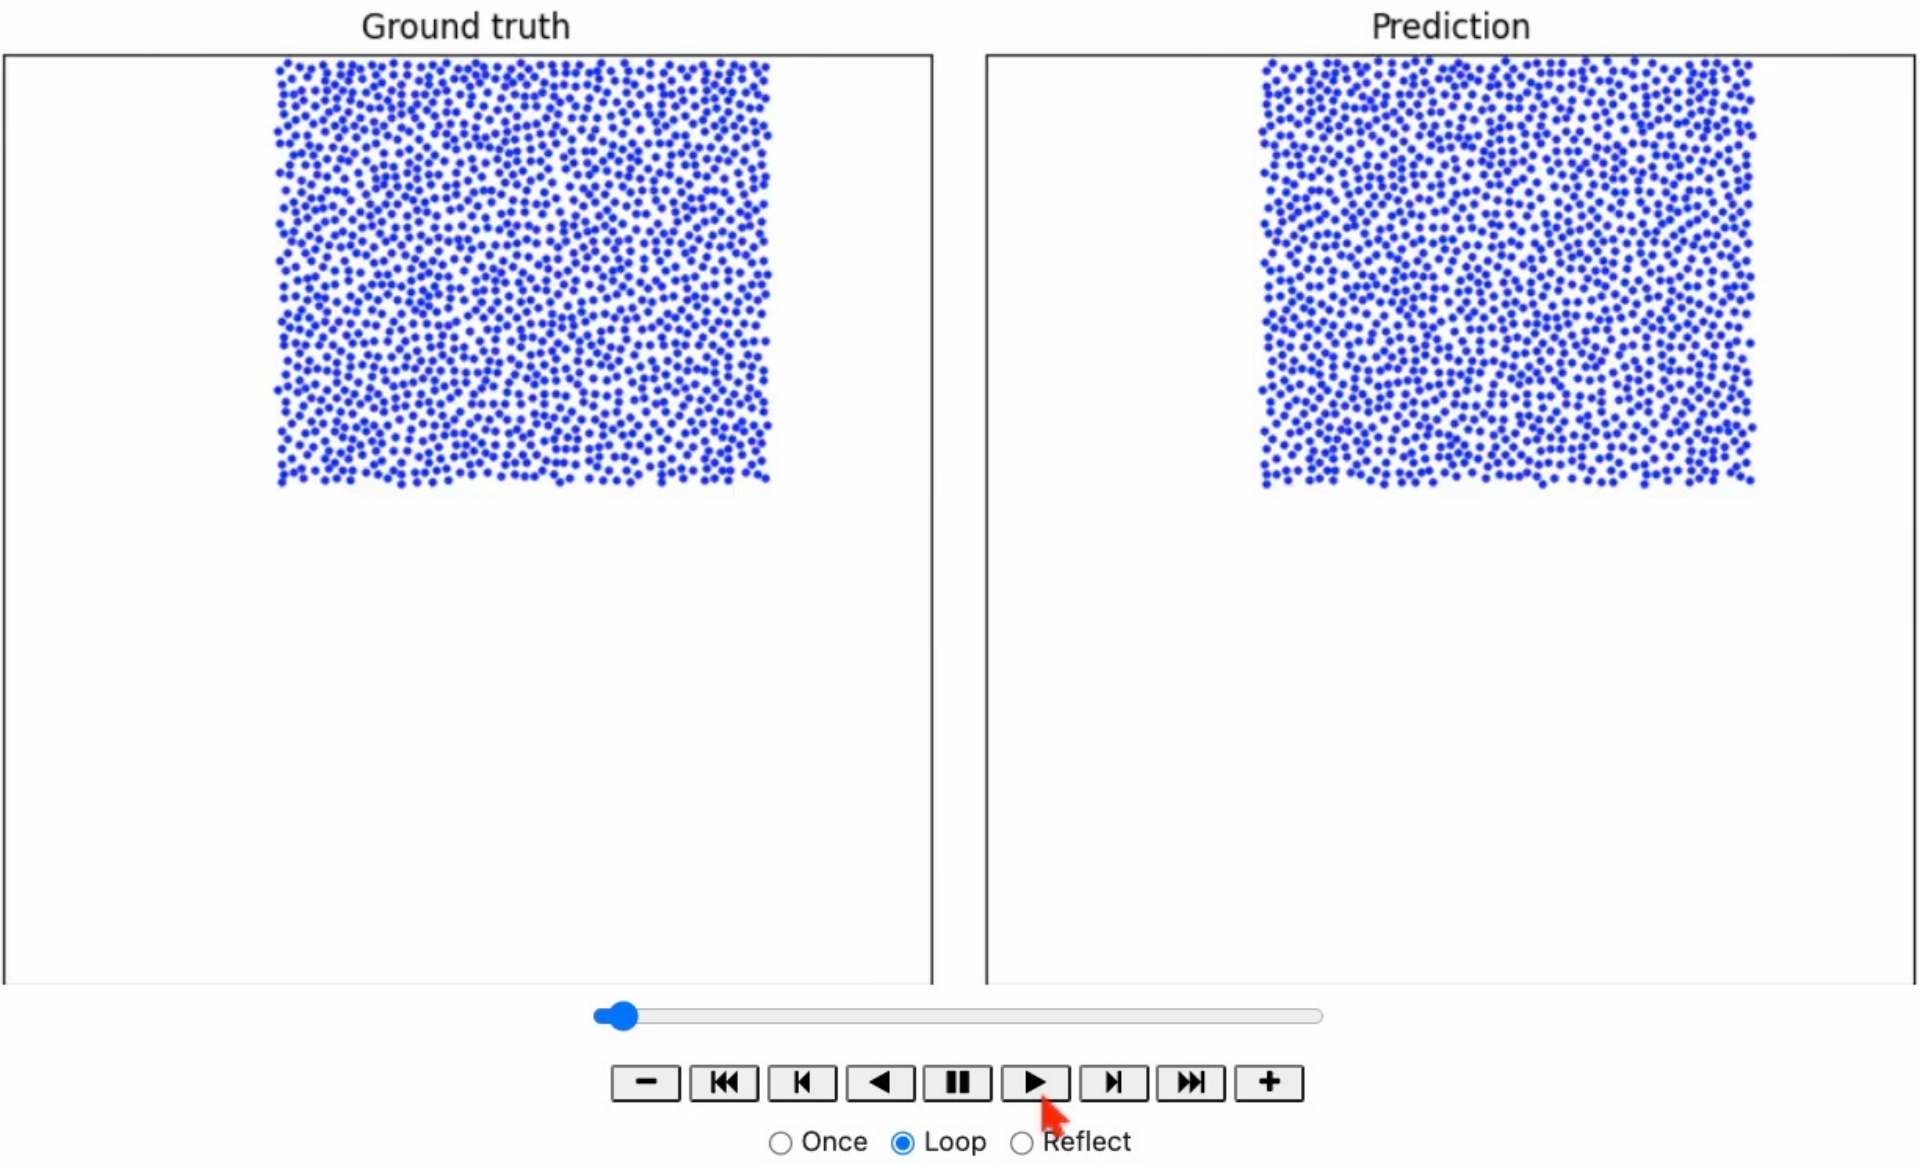

3 GPU Utilization

```
(base) [rz6525@gadi-gpu-v100-0160 ~]$ nvidia-smi
Thu Jun  8 11:53:07 2023
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 525.60.13    Driver Version: 525.60.13    CUDA Version: 12.0     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|                               |                      |               MIG M. |
|===============================+======================+======================|
|   0  Tesla V100-SXM2...  On   | 00000000:3D:00.0 Off |                    0 |
| N/A   42C    P0   142W / 300W |  31165MiB / 32768MiB |     56%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+
|   1  Tesla V100-SXM2...  On   | 00000000:3E:00.0 Off |                    0 |
| N/A   40C    P0    89W / 300W |  31165MiB / 32768MiB |     56%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+
|   2  Tesla V100-SXM2...  On   | 00000000:B2:00.0 Off |                    0 |
| N/A   40C    P0    86W / 300W |  31165MiB / 32768MiB |     47%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+
```

3 GPU simulation (Work in progress)

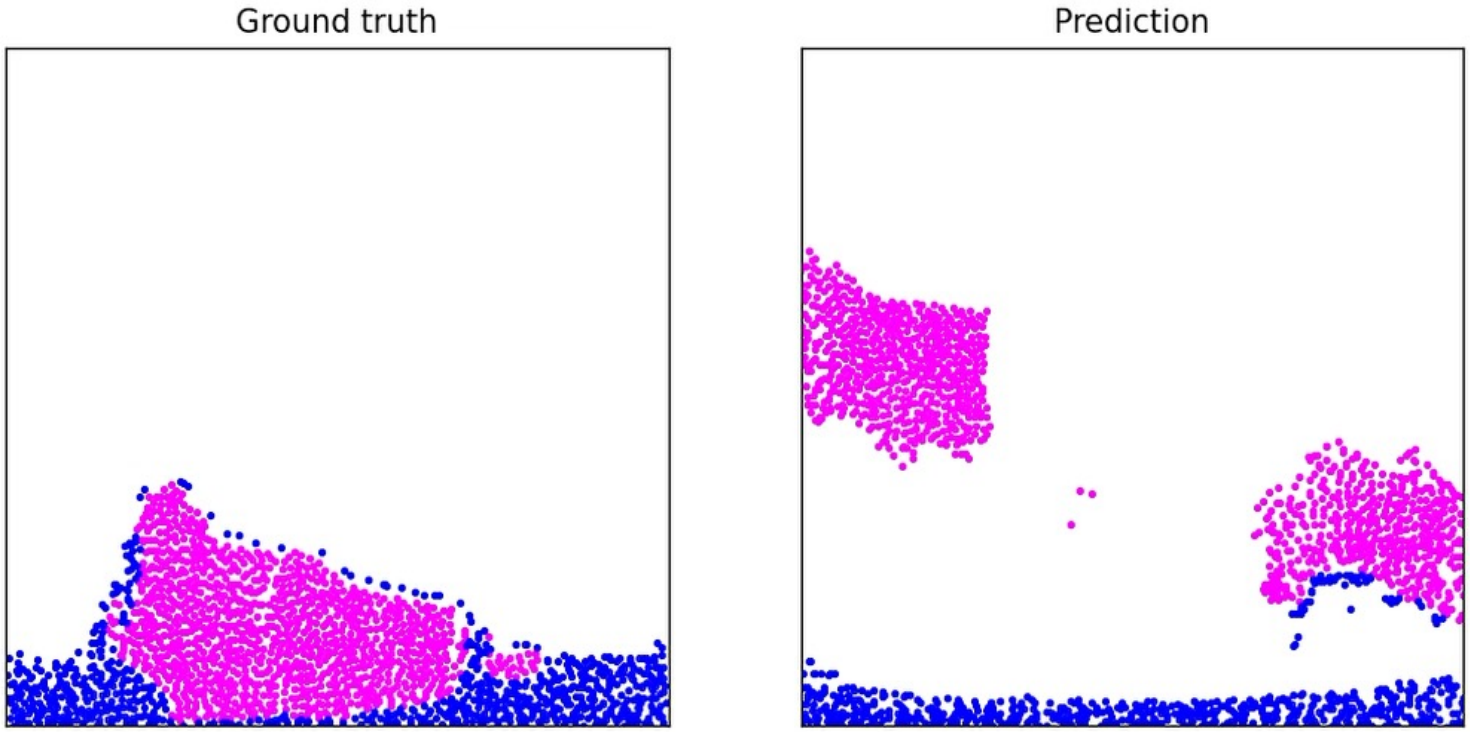

3 GPU simulation (Work in progress)

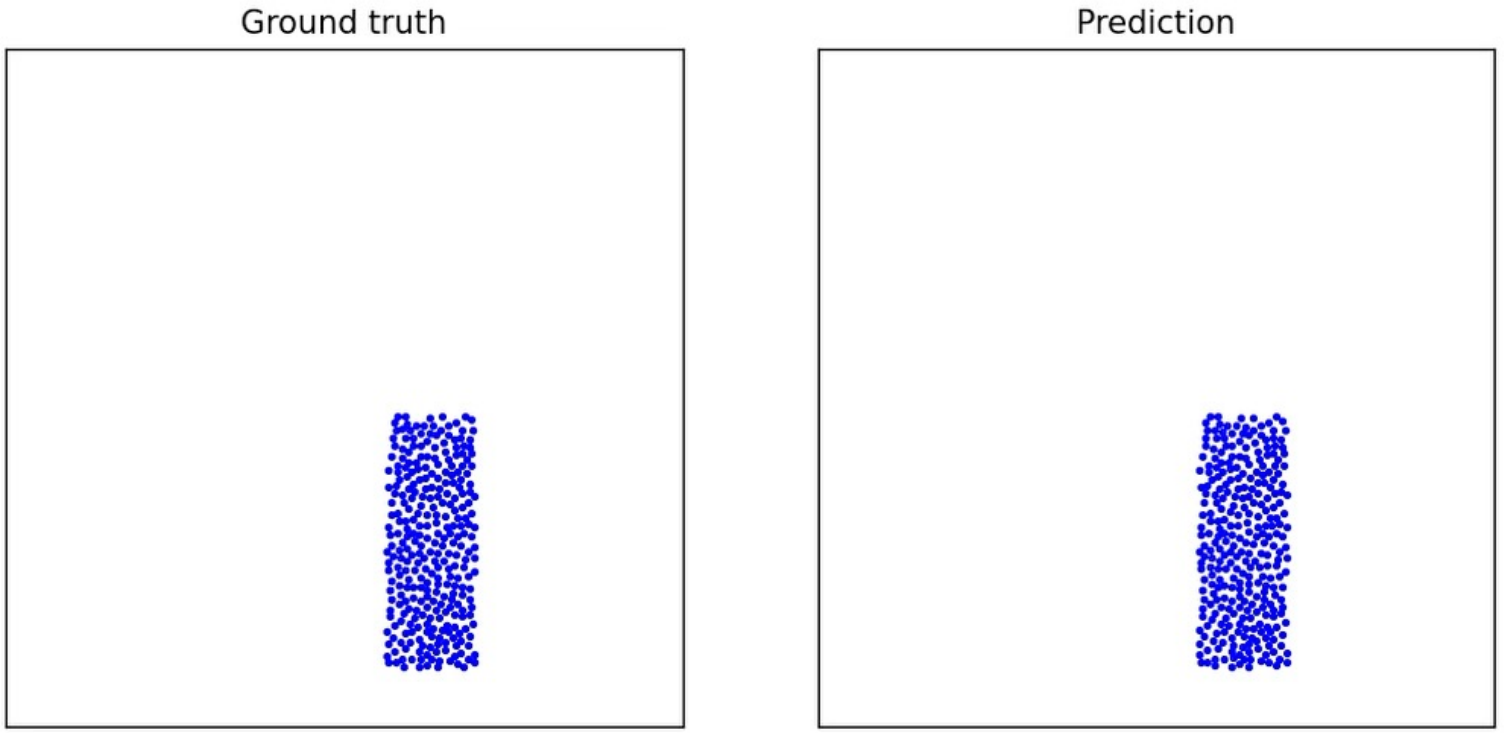

Questions?

Application Background

- High-level description of application and uses
- Light on domain-specific jargon; should be appropriate for general technical audience
- Computational motifs targeted at hackathon

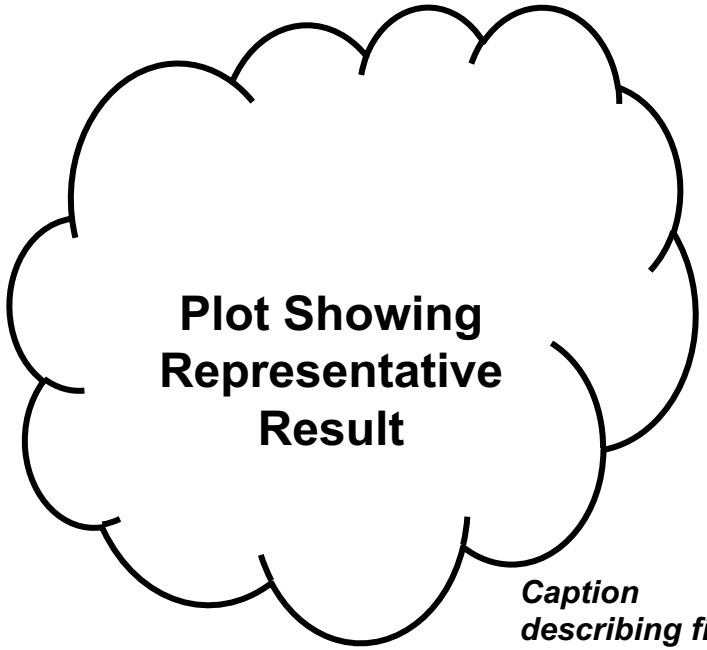

Caption describing figure in simple terms

Hackathon Objectives and Approach

- Programming models
- Profiling / hot spots
- Refactorings
- Libraries
- Performance tuning
- Other

Technical Accomplishments and Impact

- What were you able to achieve at the hackathon?
- How did you achieve it?
- Speedup
- Why does it matter / what does it enable?

Please use 100 words to summarize your team’s achievements during this Hackathon

PROMOTING YOUR WORK: AVAILABLE OPPORTUNITIES

- **Papers and Talks:** Please acknowledge the Open Hackathons program and OpenACC Organization in any planned or upcoming papers, presentations, or talks.

  "This work was completed in part at the NCI Open Hackathon, part of the Open Hackathons program. The authors would like to acknowledge OpenACC-Standard.org for their support."
- **Social Media Support:** Please feel free to promote your participation across your social media channels. Tag **@OpenACCorg** and **#OpenHackathons** and we are happy to amplify.
- **Blogs and Technical Write-ups**: Create a blog post or technical article that highlights the work being done and results achieved.
- **Quotes and Testimonials:** Highlight your quote or feedback on our channels (i.e. social, website, etc.).

***Please reach out to Izumi Barker (ibarker@nvidia.com) to discuss marketing options and opportunities.**