OAK RIDGE National Laboratory

GPU Enhanced Remote Collaborative Scientific Visualization

Benjamin Hernandez (OLCF), Tim Biedert (NVIDIA)

March 20th, 2019

U.S. DEPARTMENT OF ENERGY

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Contents

- Part I
  - GPU Enhanced Remote Collaborative Scientific Visualization
- Part II
  - Hardware-accelerated Multi-tile streaming for Realtime Remote Visualization

Oak Ridge Leadership Computing Facility (OLCF)

- Provide the computational and data resources required to solve the most **challenging problems**.

- Highly **competitive** user allocation programs (INCITE, ALCC).
  - OLCF provides **10x** to **100x** more resources than other centers

- We **collaborate** with users sharing **diversity** in **expertise** and **geographic location**

How are these collaborations ?

- Collaborations
  - Extends through the life cycle of the data from computation to analysis and visualization.
  - Are structured around data.
- Data analysis and Visualization
  - An iterative and sometimes remote process involving students, visualization experts, PIs and stakeholders

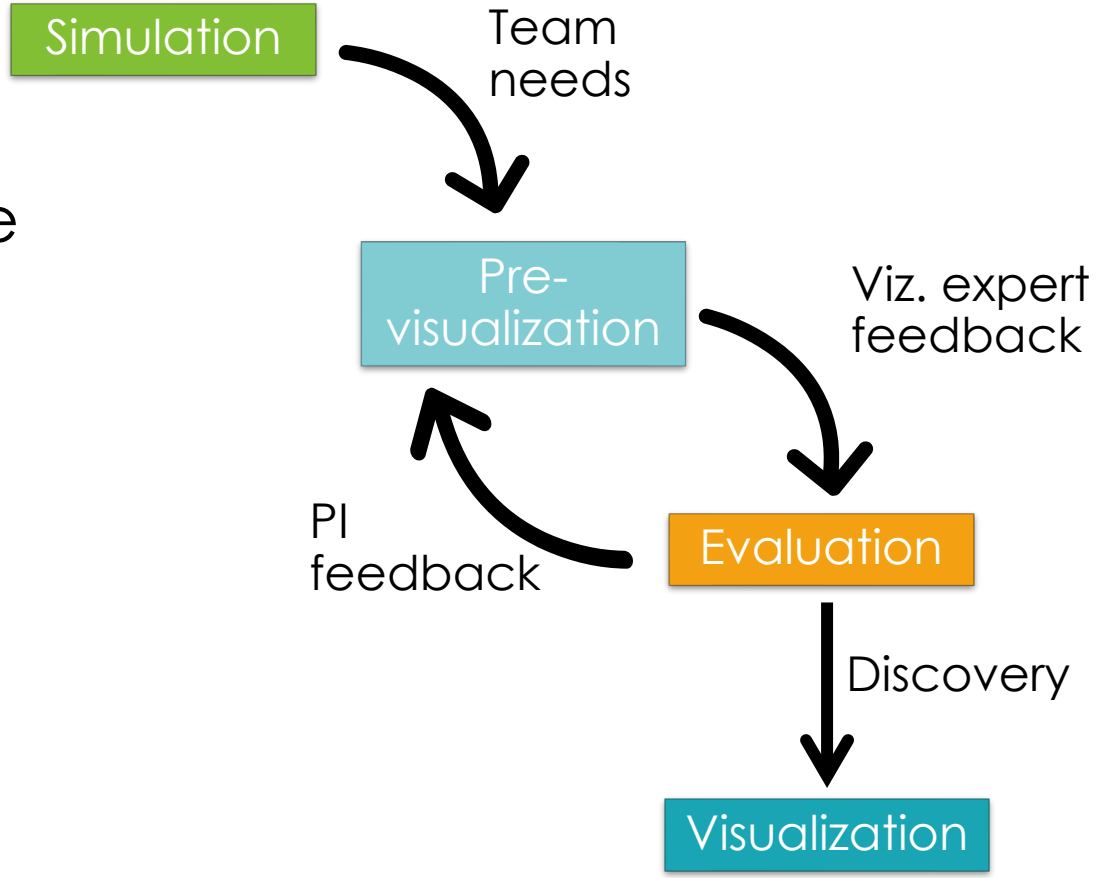

How are these collaborations ?

- The collaborative future must be characterized by[1]:

Discovery
Resources easy to find

Connectivity
No resource is an island

Portability
Resources widely and transparently usable

Centrality
Resources efficiently, and centrally supported

"Web-based immersive visualization Tools with the ease of a virtual reality game."

"Visualization ideally would be combined with user-guided as well as template-guided automated feature extraction, real-time annotation, and quantitative geometrical analysis."

"Rapid data visualization and analysis to enable understanding in near real time by a geographically dispersed team."

…

1. U.S. Department of Energy.(2011). *Scientific collaborations for extreme-scale science (Workshop Report). Retrieved from* https://indico.bnl.gov/event/403/attachments/11180/13626/ScientificCollaborationsforExtreme-ScaleScienceReportDec2011_Final.pdf

How are these collaborations ?

- INCITE "*Petascale simulations of short pulse laser interaction with metals*" PI Leonid Zhigilei, University of Virginia
  - Laser ablation in vacuum and liquid environments
  - Hundreds of millions to billion scale atomistic simulations, dozens of time steps

- INCITE "*Molecular dynamics of motor-protein networks in cellular energy metabolism*" PI(s) Abhishek Singharoy, Arizona State University
  - Hundreds of million scale atomistic simulation, hundreds of time steps

How are these collaborations ?

- The data is centralized in OLCF
- SIGHT is a custom platform for interactive data analysis and visualization.
  - Support for collaborative features are needed for
    - Low latency remote visualization streaming
    - Simultaneous and independent user views
    - Collaborative multi-display setting environment

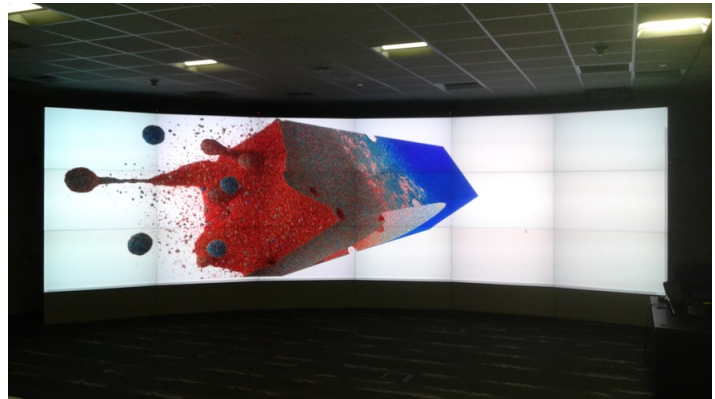

SIGHT: Exploratory Visualization of Scientific Data

- Designed around user needs
- Lightweight tool
  - Load your data
  - Perform exploratory analysis
  - Visualize/Save results
- Heterogeneous scientific visualization
  - Advanced shading to enable new insights into data exploration.
  - Multicore and manycore support.
- Remote visualization
  - Server/Client architecture to provide high end visualization in laptops, desktops, and powerwalls.
- Multi-threaded I/O
- Supports interactive/batch visualization
  - In-situ (some effort)
- Designed having OLCF infrastructure in mind.

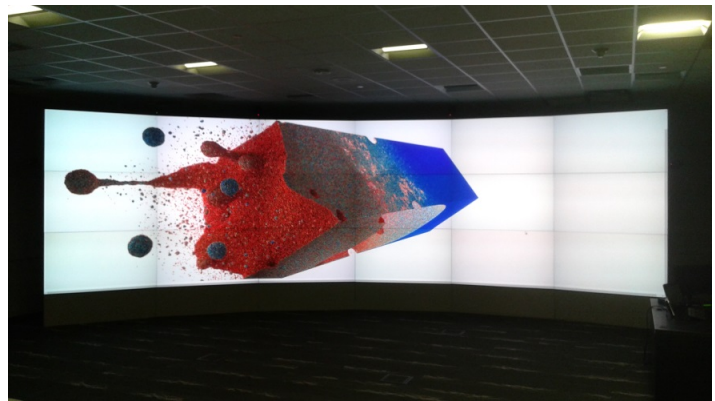

Publications

M. V. Shugaev, C. Wu, O. Armbruster, A. Naghilou, N. Brouwer, D. S. Ivanov, T. J.-Y. Derrien, N. M. Bulgakova, W. Kautek, B. Rethfeld, and L. V. Zhigilei, Fundamentals of ultrafast laser-material interaction, MRS Bull. 41 (12), 960-968, 2016.

C.-Y. Shih, M. V. Shugaev, C. Wu, and L. V. Zhigilei, Generation of subsurface voids, incubation effect, and formation of nanoparticles in short pulse laser interactions with bulk metal targets in liquid: Molecular dynamics study, J. Phys. Chem. C 121, 16549-16567, 2017.

C.-Y. Shih, R. Streubel, J. Heberle, A. Letzel, M. V. Shugaev, C. Wu, M. Schmidt, B. Gökce, S. Barcikowski, and L. V. Zhigilei, Two mechanisms of nanoparticle generation in picosecond laser ablation in liquids: the origin of the bimodal size distribution, *Nanoscale* **10**, 6900-6910, 2018.

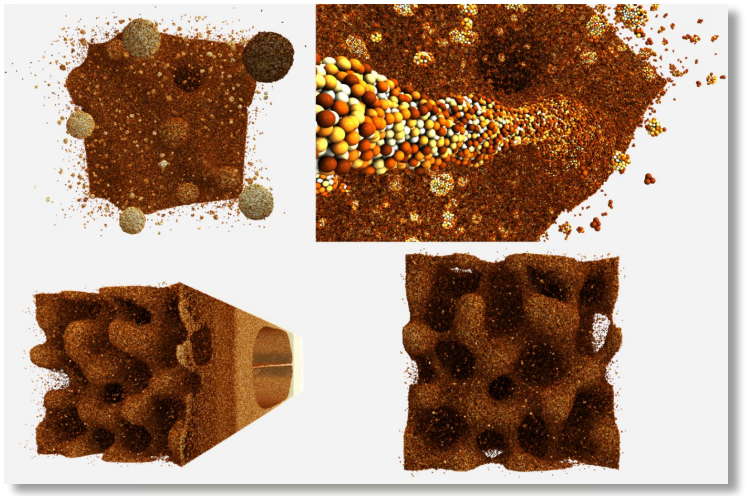

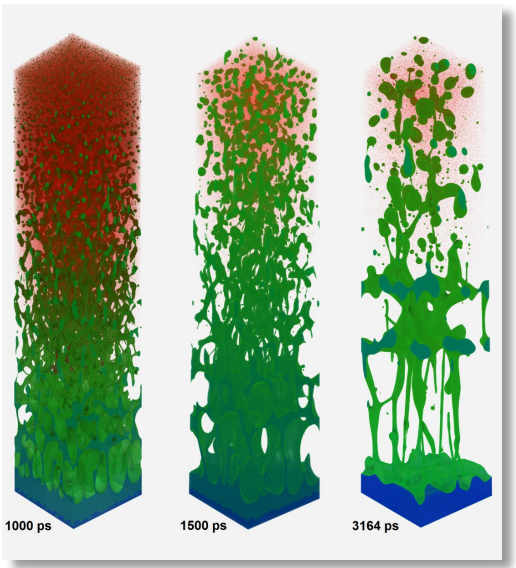

2000ps

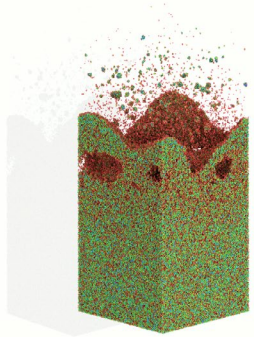

SIGHT's System Architecture

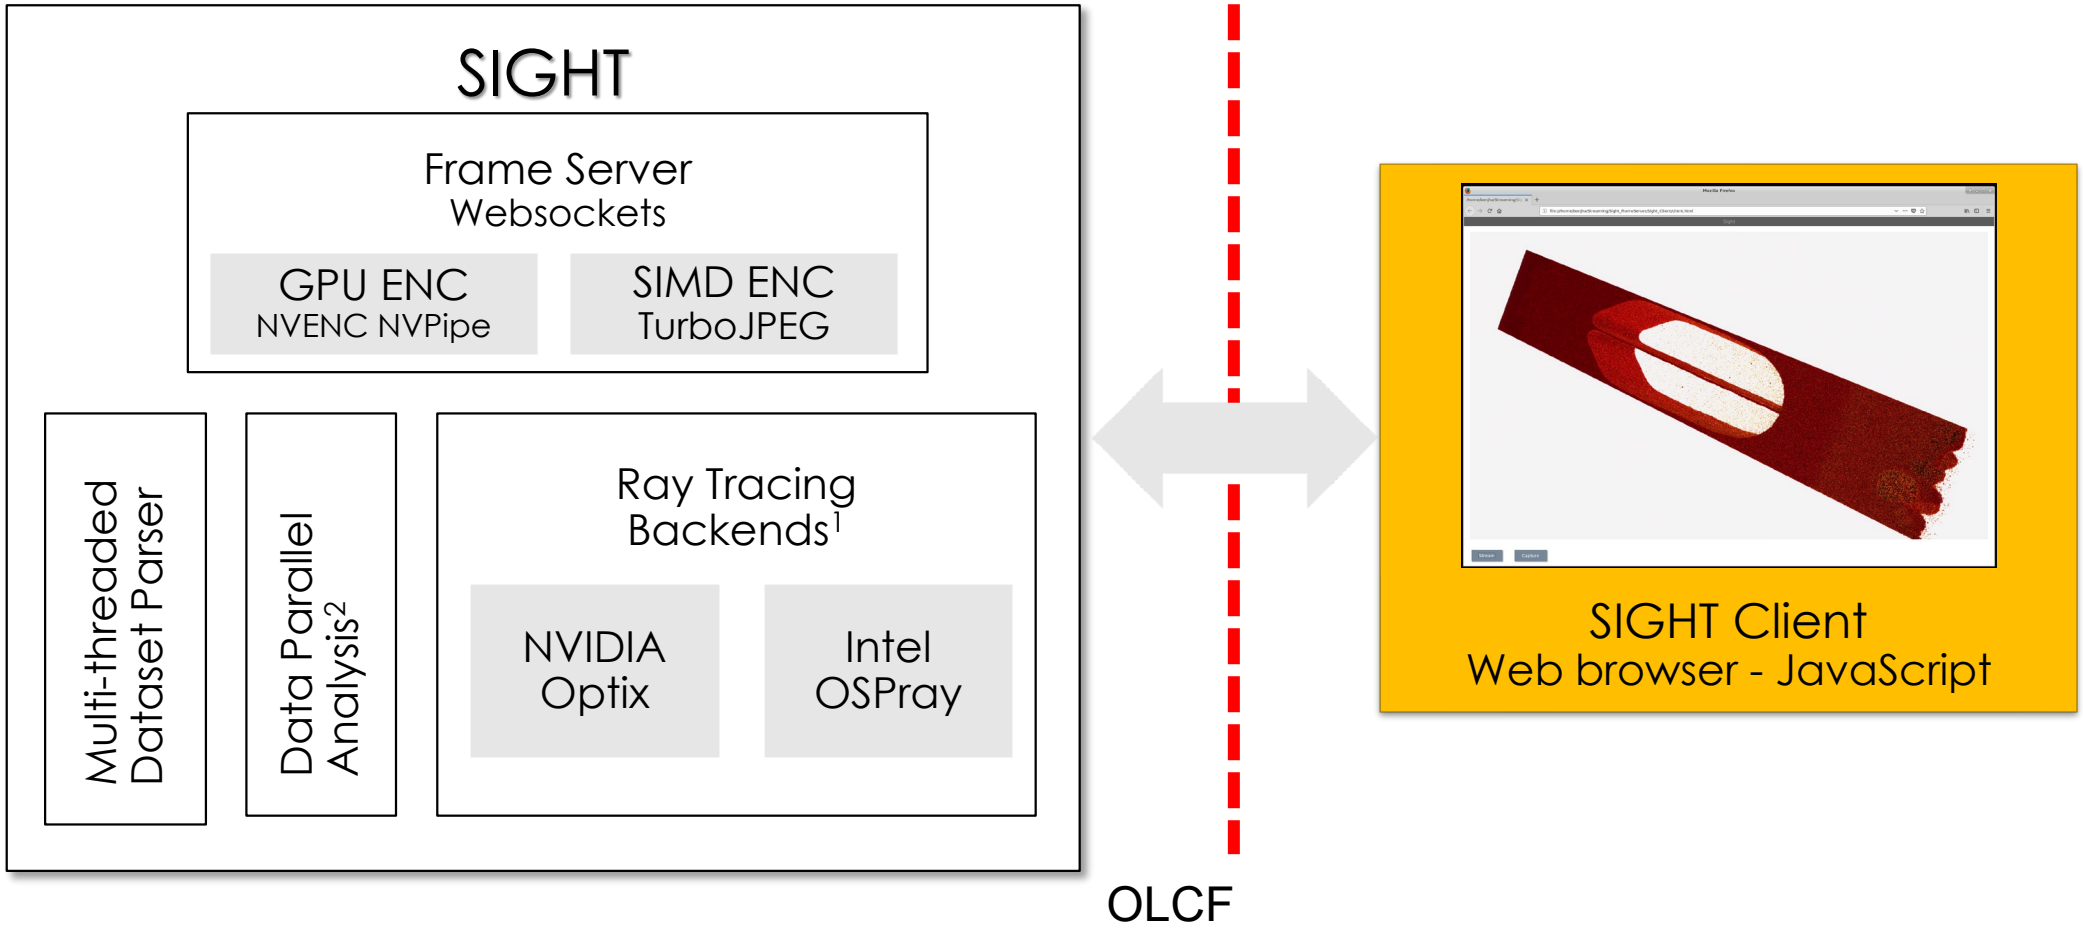

[1]S7175 *Exploratory Visualization of Petascale Particle Data in Nvidia DGX-1*
Nvidia GPU Technology Conference 2017
[2]P8220 *Heterogeneous Selection Algorithms for Interactive Analysis of Billion Scale Atomistic Datasets*
Nvidia GPU Technology Conference 2018

Design of Collaborative Infrastructure - Alternative 1

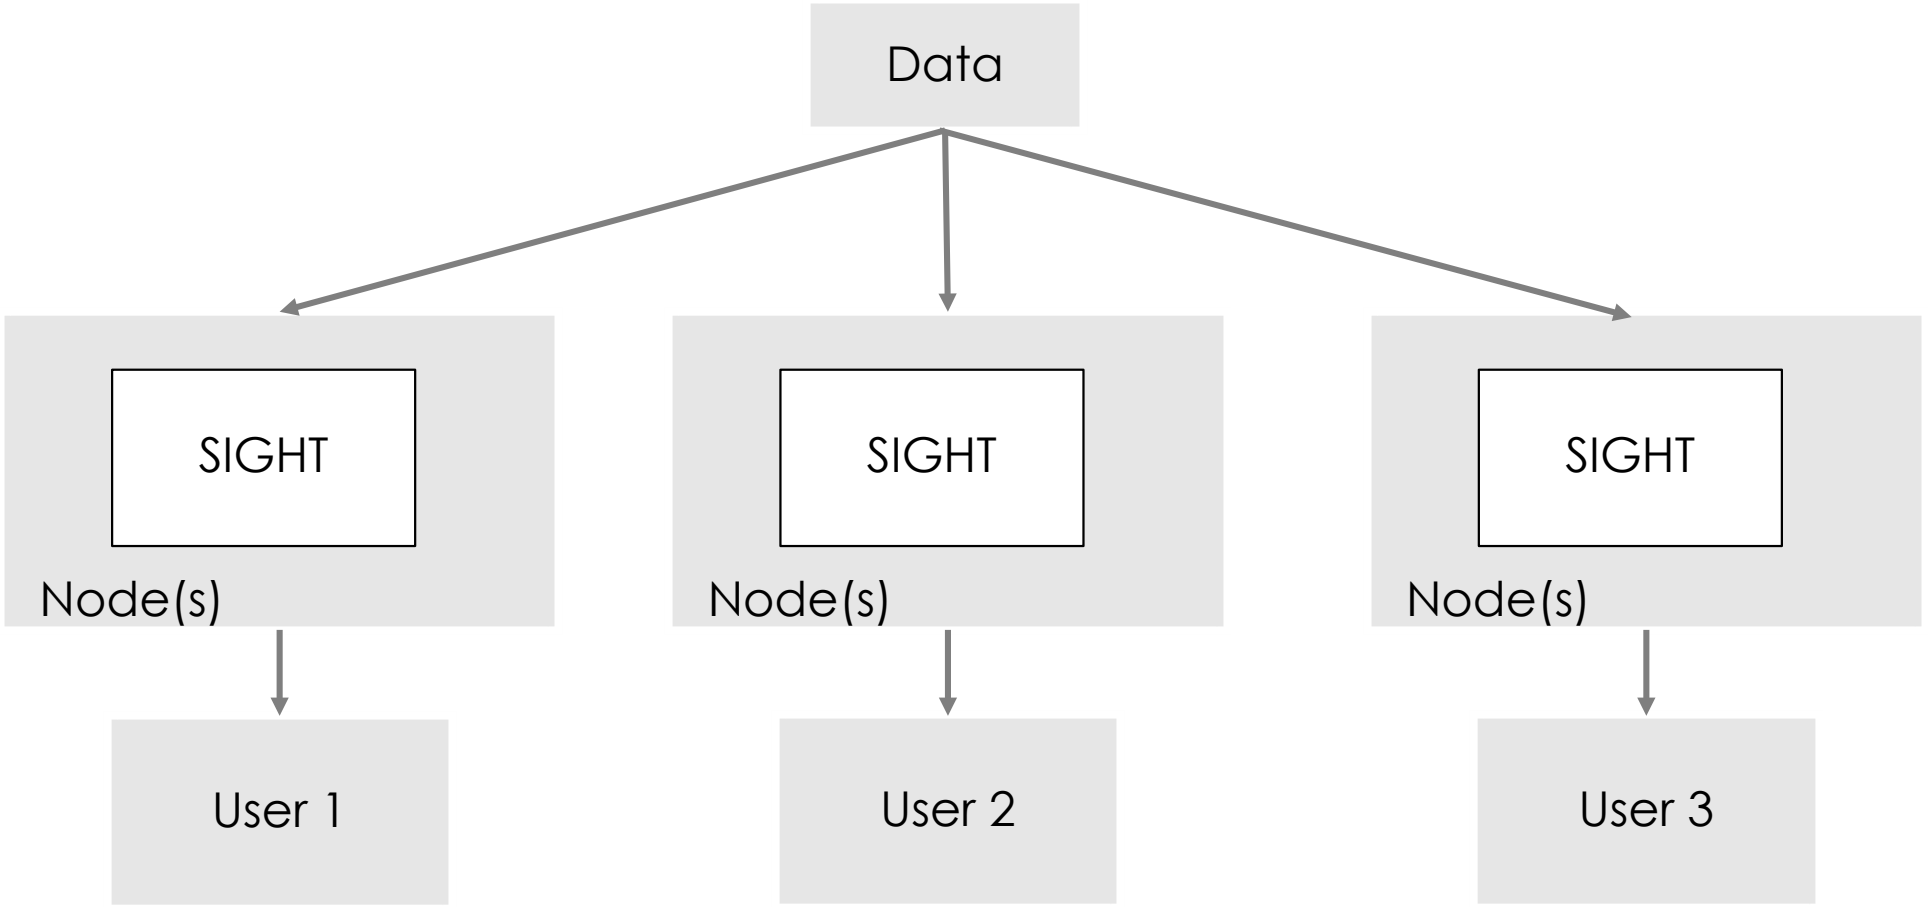

Design of Collaborative Infrastructure - Alternative 1

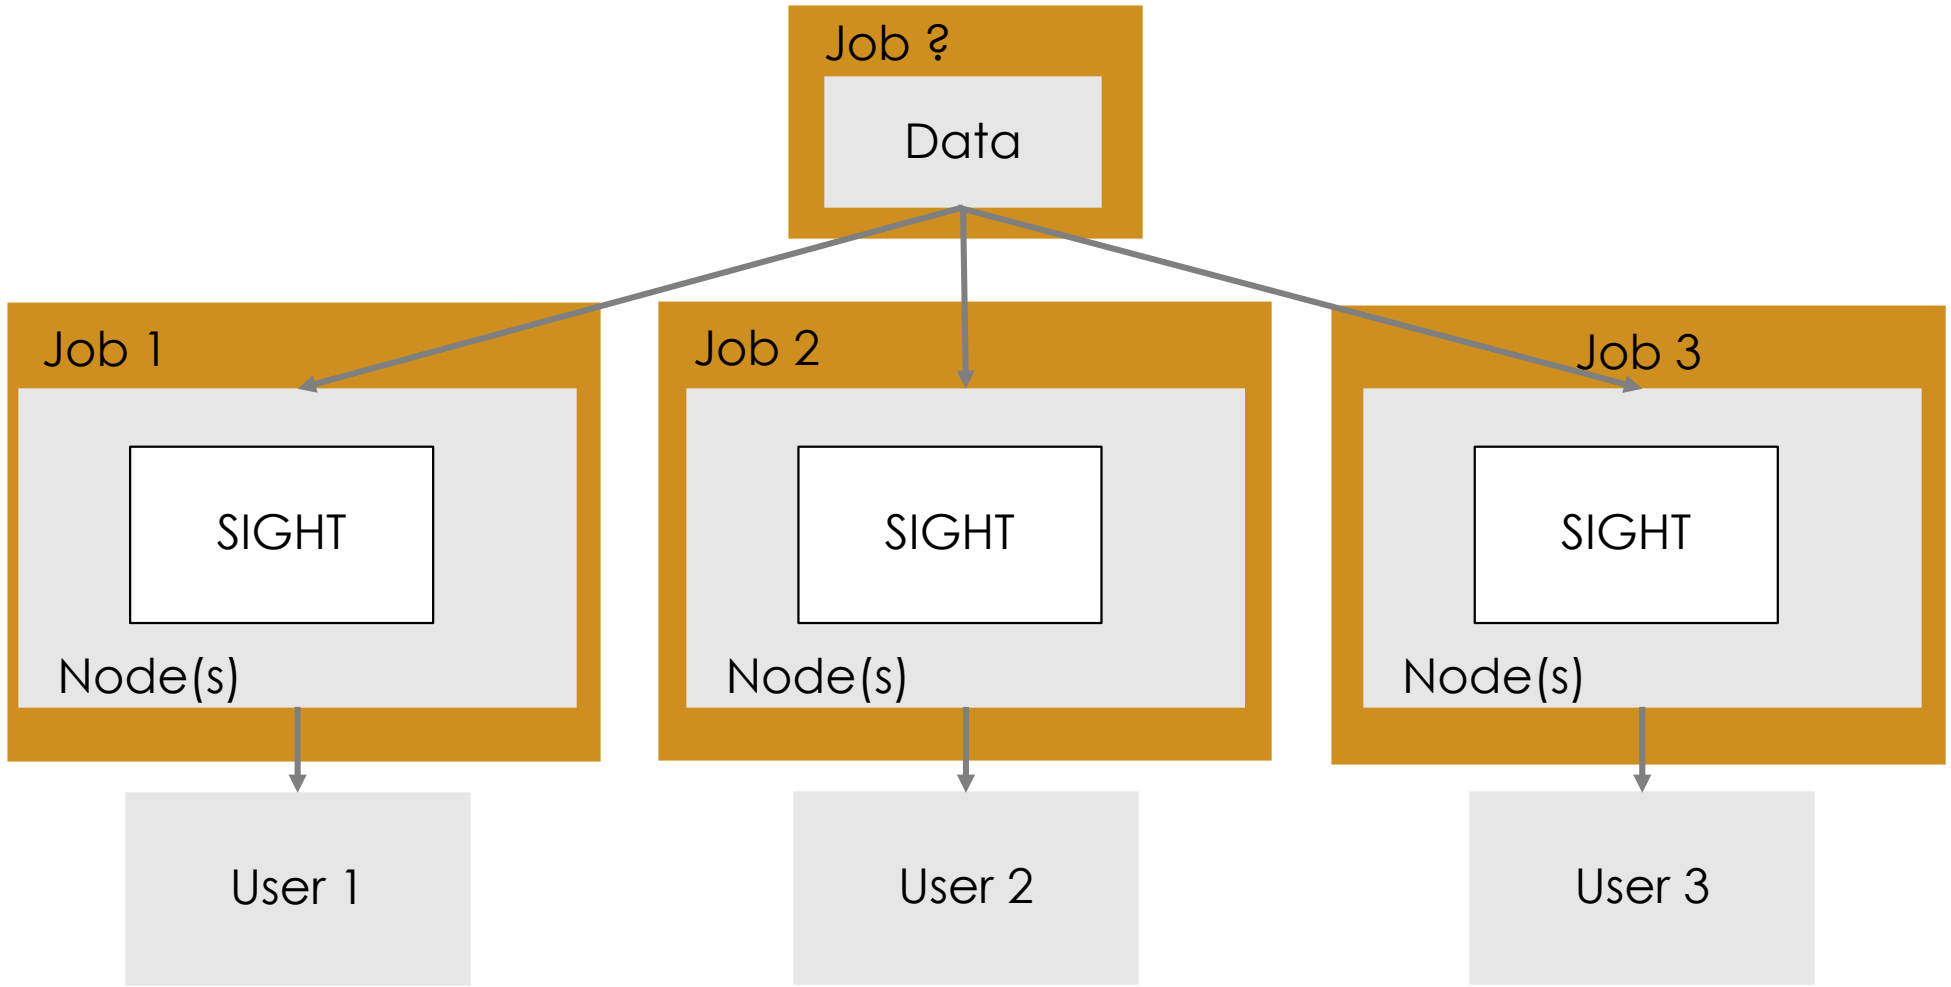

Design of Collaborative Infrastructure - Alternative 1

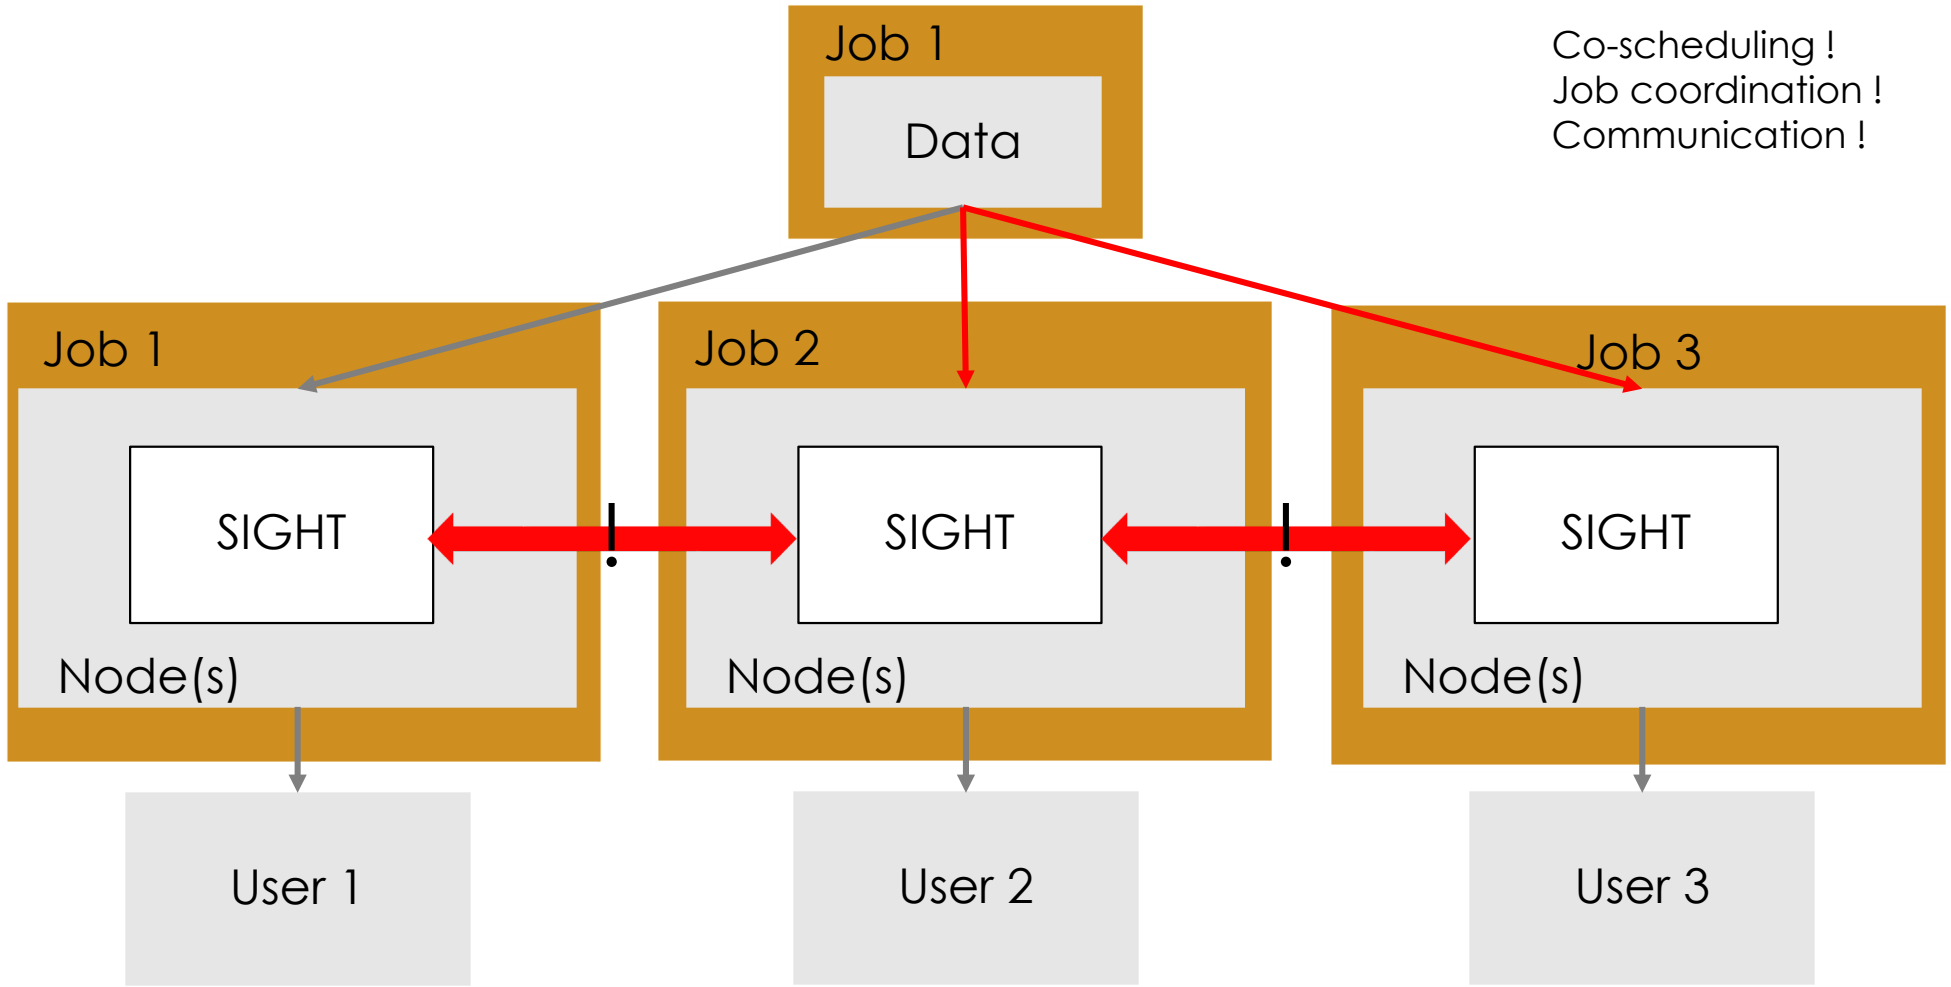

Design of Collaborative Infrastructure - Alternative 2

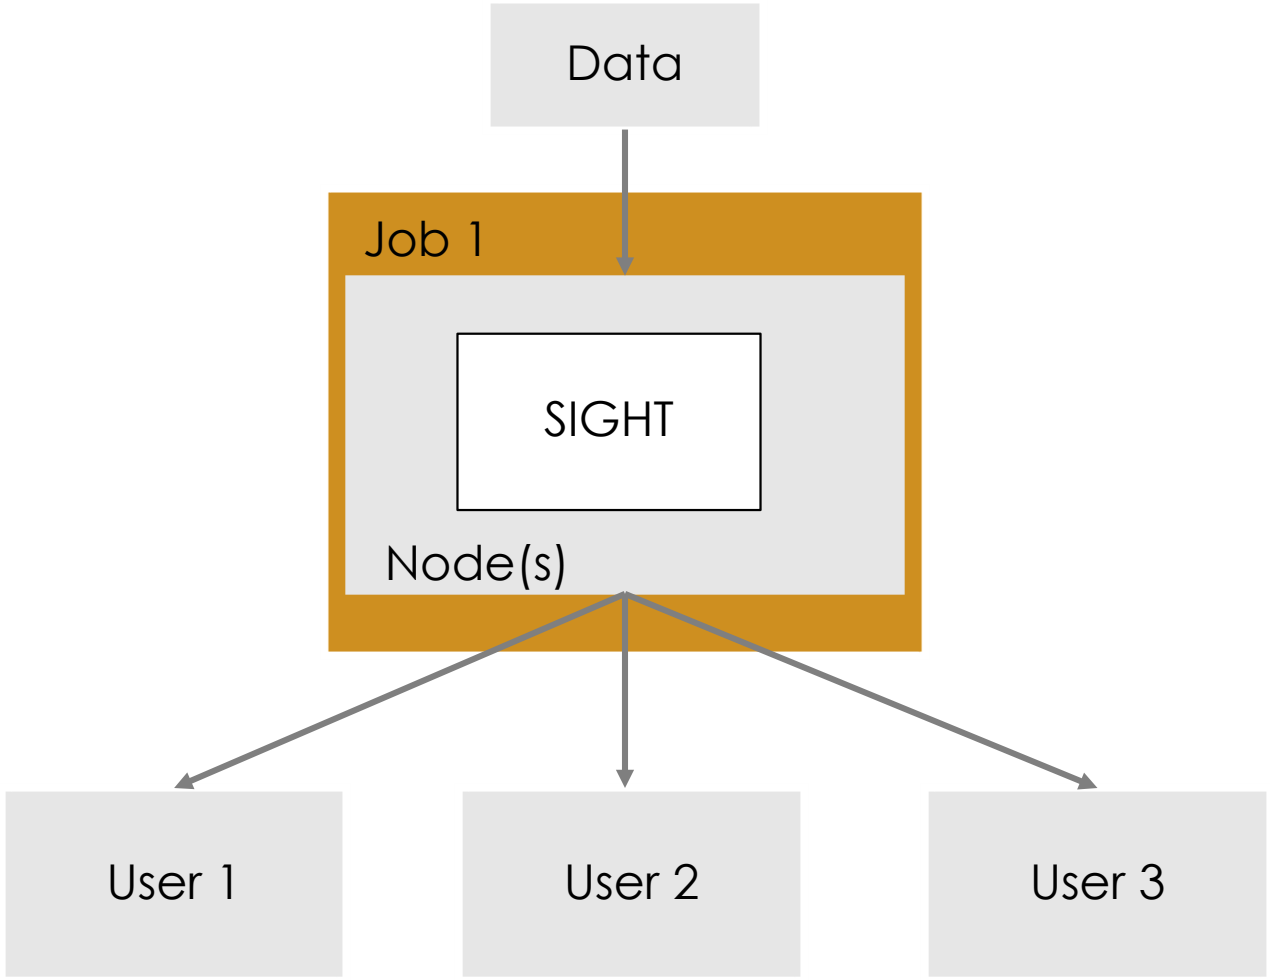

Design of Collaborative Infrastructure

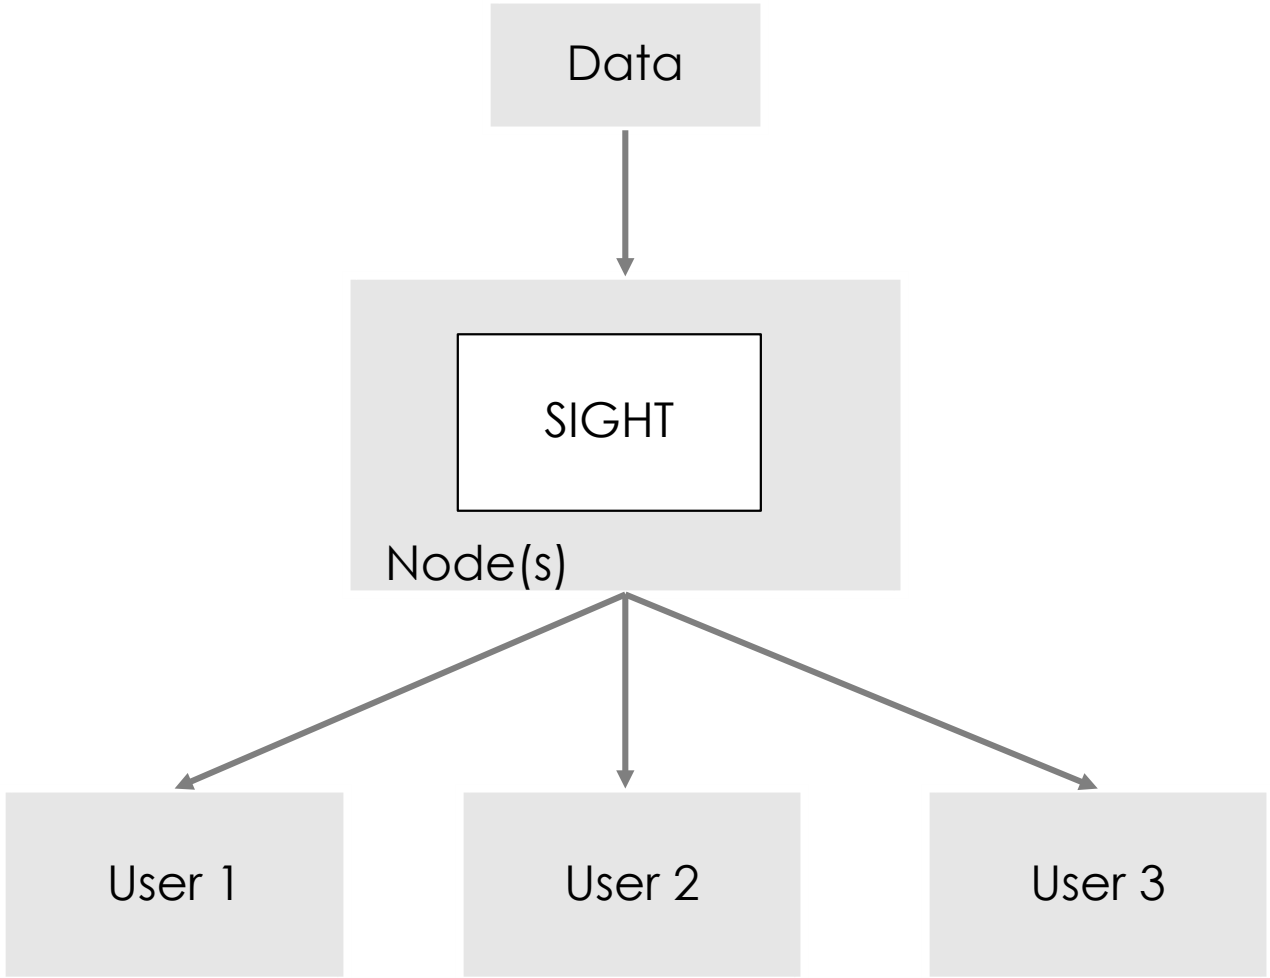

- ORNL Summit System Overview
  - 4,608 nodes
  - Dual-port Mellanox EDR InfiniBand network
  - 250 PB IBM file system transferring data at 2.5 TB/s
- Each node has
  - 2 IBM POWER9 processors
  - 6 NVIDIA Tesla V100 GPUs
  - 608 GB of fast memory (96 GB HBM2 + 512 GB DDR4)
  - 1.6 TB of NV memory

Design of Collaborative Infrastructure

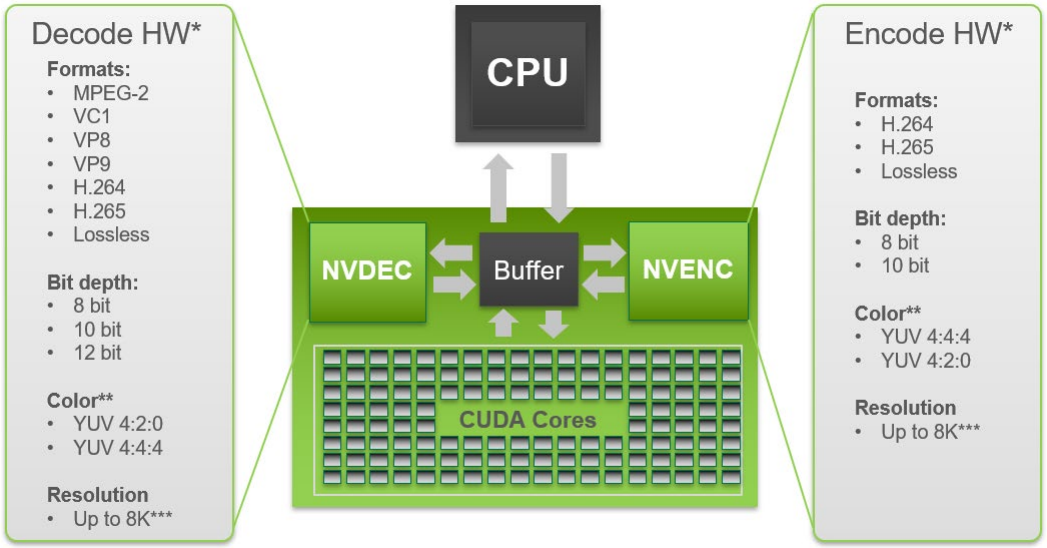

- Each NVIDIA Tesla V100 GPU
  - 3 NVENC Chips
  - Unrestricted number of concurrent sessions
- NVPipe
  - Lightweight C API library for low-latency video compression
  - Easy access to NVIDIA's hardware-accelerated H.264 and HEVC video codecs

Enhancing SIGHT Frame Server
Low latency encoding

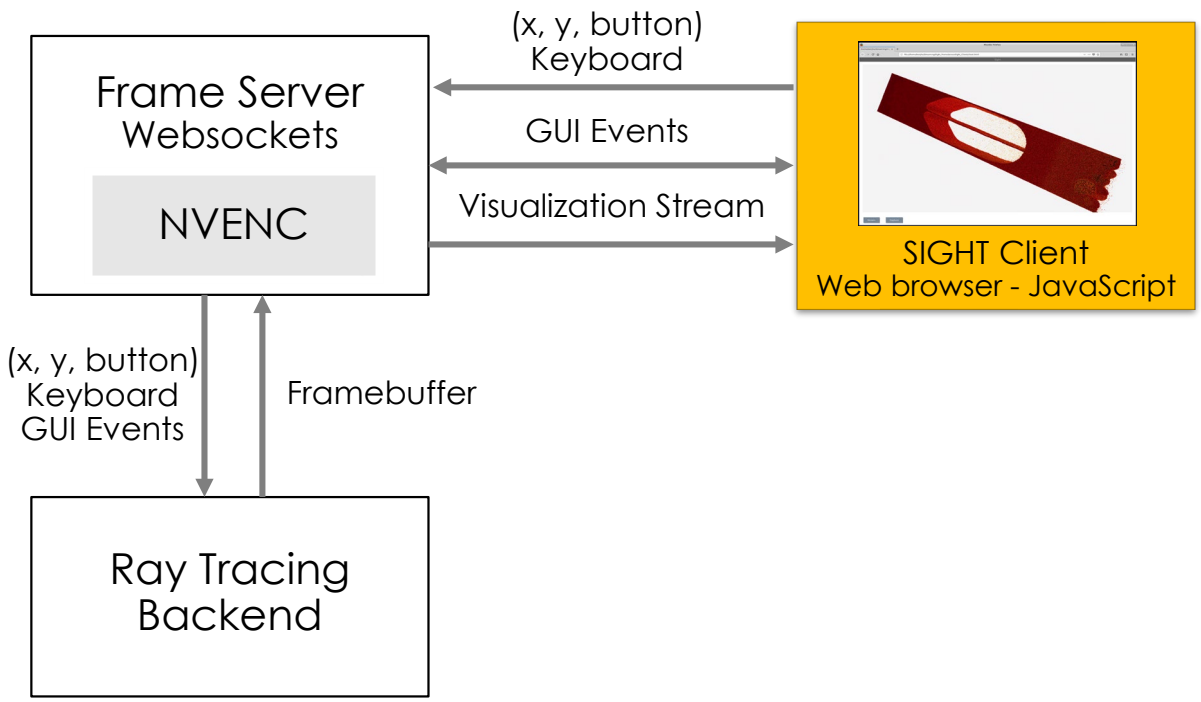

Sharing Optix buffer (OptiXpp)

```
Buffer frameBuffer = context->createBuffer(
RT_BUFFER_OUTPUT, RT_FORMAT_UNSIGNED_BYTE4, m_width,
m_height );

frameBufferPtr = buffer->getDevicePointer(optxDevice);
…

compress (frameBufferPtr);
```

Enhancing SIGHT Frame Server
Low latency encoding

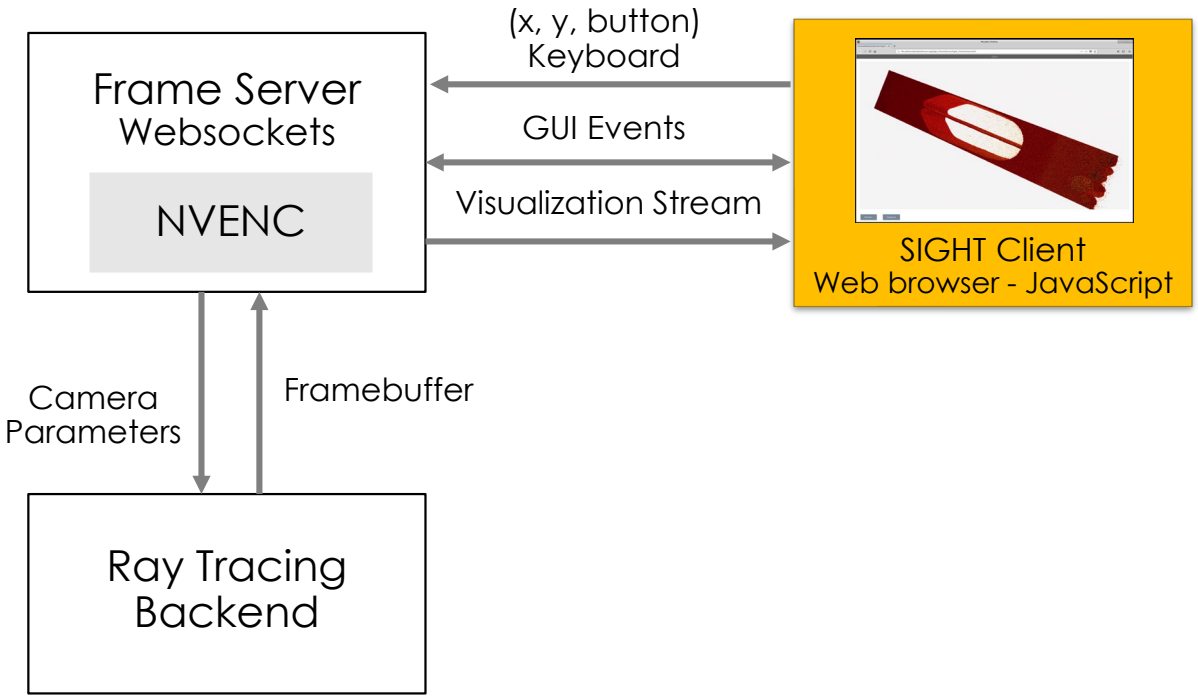

Opening encoding session:

```
myEncoder = NvPipe_CreateEncoder (NVPIPE_RGBA32,
            NVPIPE_H264, NVPIPE_LOSSY,
            bitrateMbps * 1000 * 1000, targetFps);
if (!m_encoder)
    return error;
myCompressedImg = new unsigned char[w*h*4];
```

Framebuffer compression:

```
bool compress (myFramebufferPtr)
{
...
myCompressedSize = NvPipe_Encode(myEncoder,
                   myFramebufferPtr,
                   w*4, myCompressedImg, w*h*4, w, h, true);
if (myCompressedSize == 0 )
          return error;
...
}
```

Closing encoding session:

```
NvPipe_Destroy(myEncoder);
```

Enhancing SIGHT Frame Server
Low latency encoding

- System Configuration:
  - DGX-1 Volta
  - Connection Bandwidth
    - 800Mbps (ideally 1 Gbps)
  - NVENC Encoder
    - H264 PROFILE BASELINE
    - 32 MBPS, 30 FPS
  - Turbo JPEG
    - SIMD Instructions on
    - JPEG Quality 50
  - Decoder
    - Broadway.js (FireFox 65)
    - Media Source Extensions (Chrome 72)
    - Built-in JPEG Decoder (Chrome 72)

Average

	NVENC	NVENC+MP4	TJPEG
Encoding HD (ms)	4.65	6.05	16.71
Encoding 4K (ms)	12.13	17.89	51.89
Frame Size HD (KB)	116.00	139.61	409.76
Frame Size 4K (KB)	106.32	150.65	569.04

	Broadway.cs	MSE	Built-in JPEG
Decoding HD (ms)	43.28	39.97	78.15
Decoding 4K (ms)	87.40	53.10	197.63

Enhancing SIGHT Frame Server
Low latency encoding

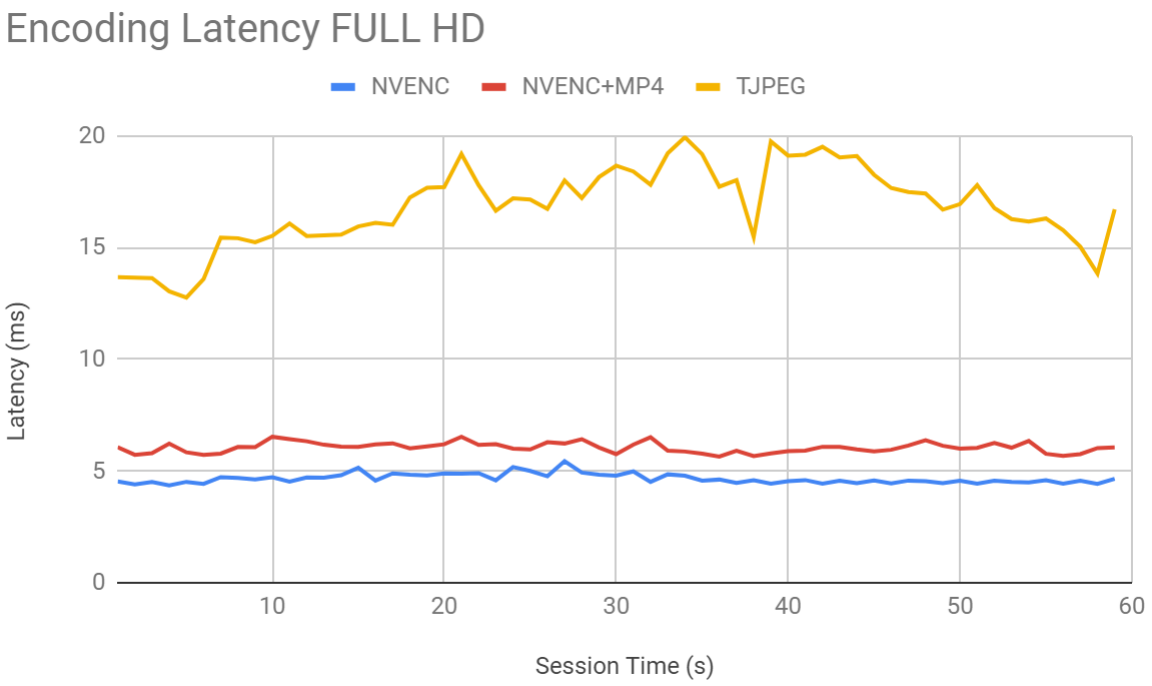

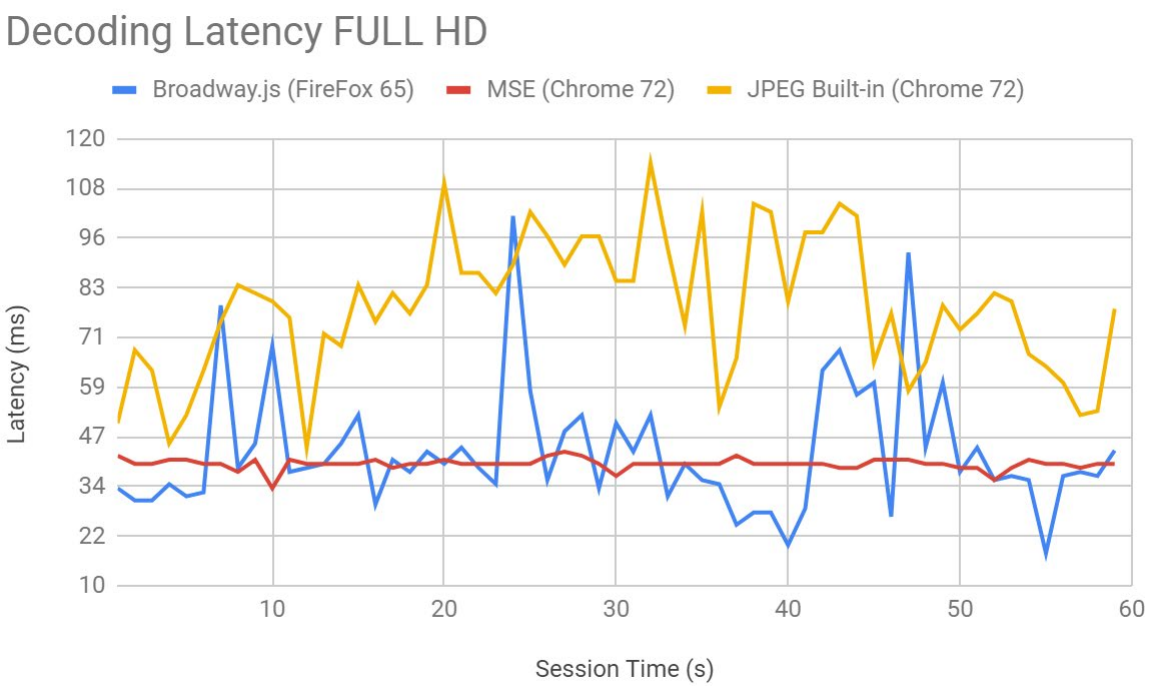

Enhancing SIGHT Frame Server
Low latency encoding

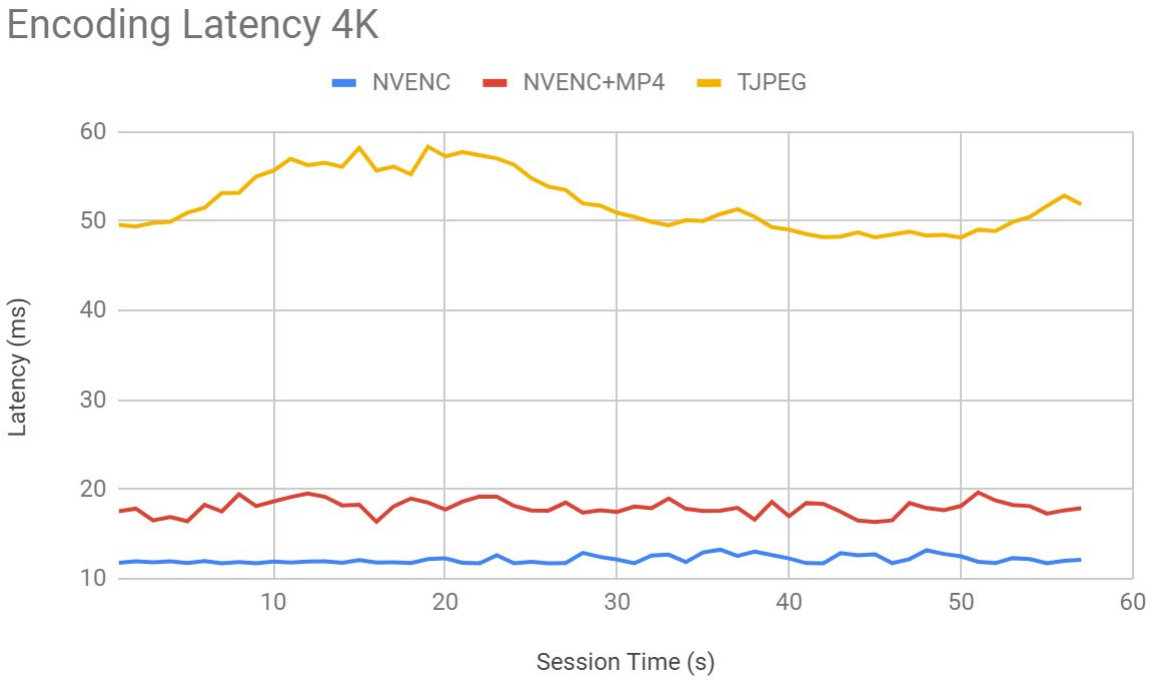

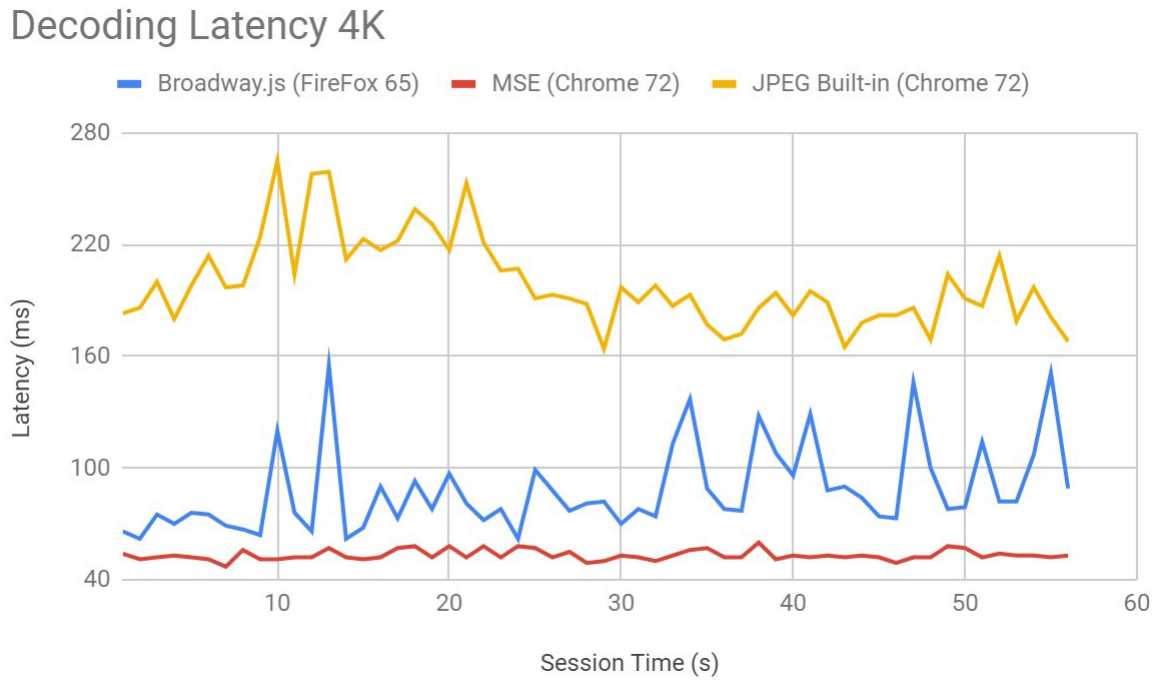

Enhancing SIGHT Frame Server
Simultaneous and independent user views

- A Summit node can produce 4K visualizations with NVIDIA Optix at interactive rates.

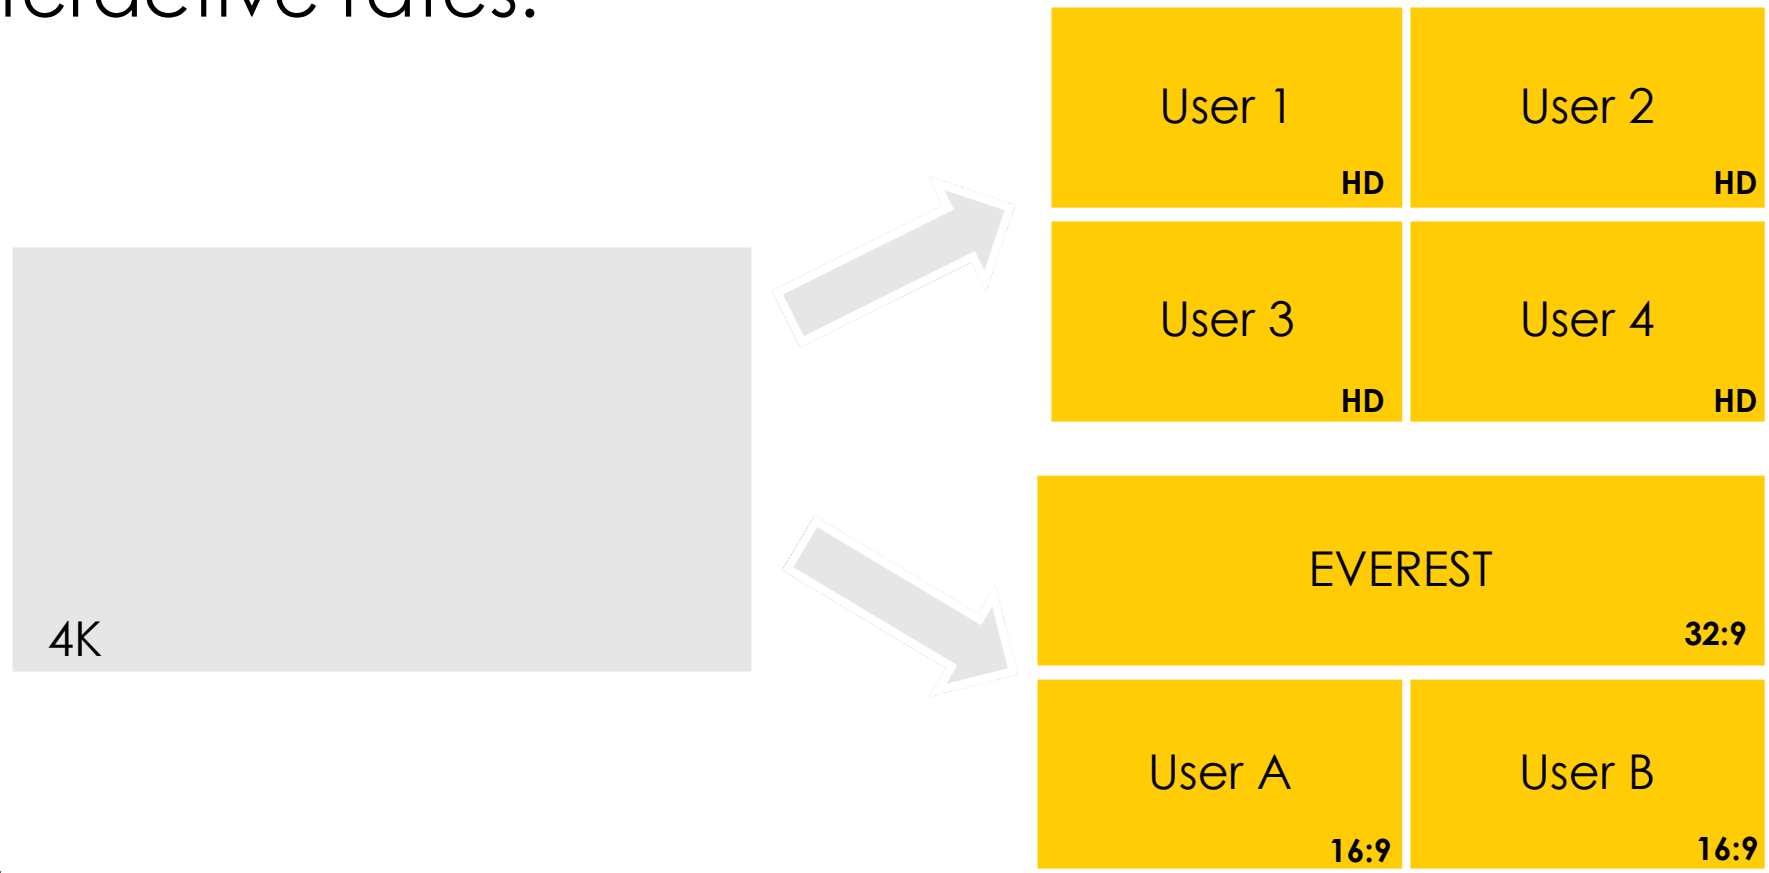

Enhancing SIGHT Frame Server
Simultaneous and independent user views

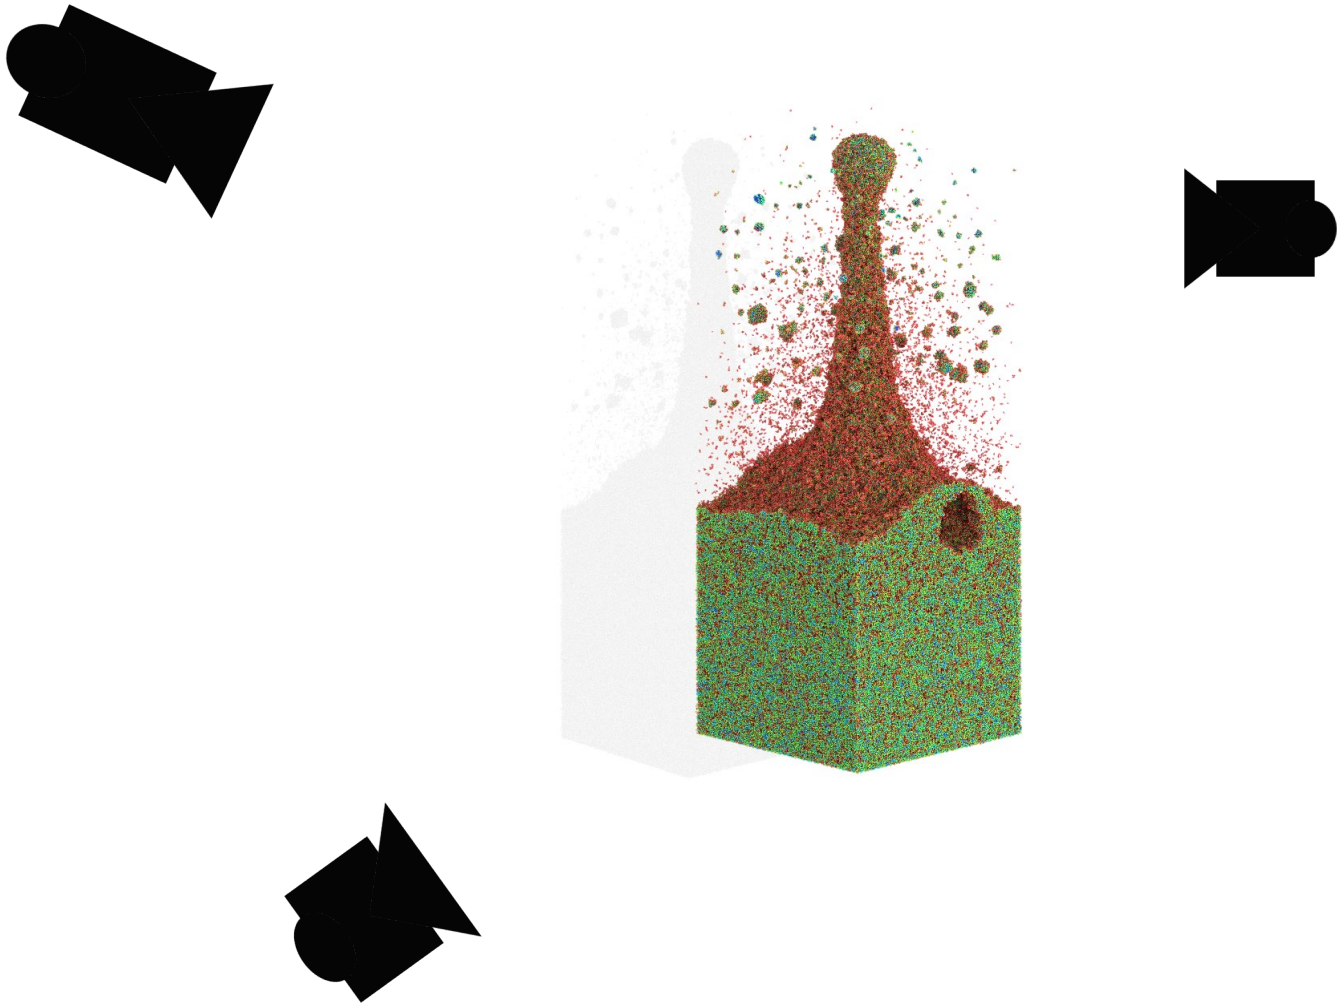

Enhancing SIGHT Frame Server

Simultaneous and independent user views

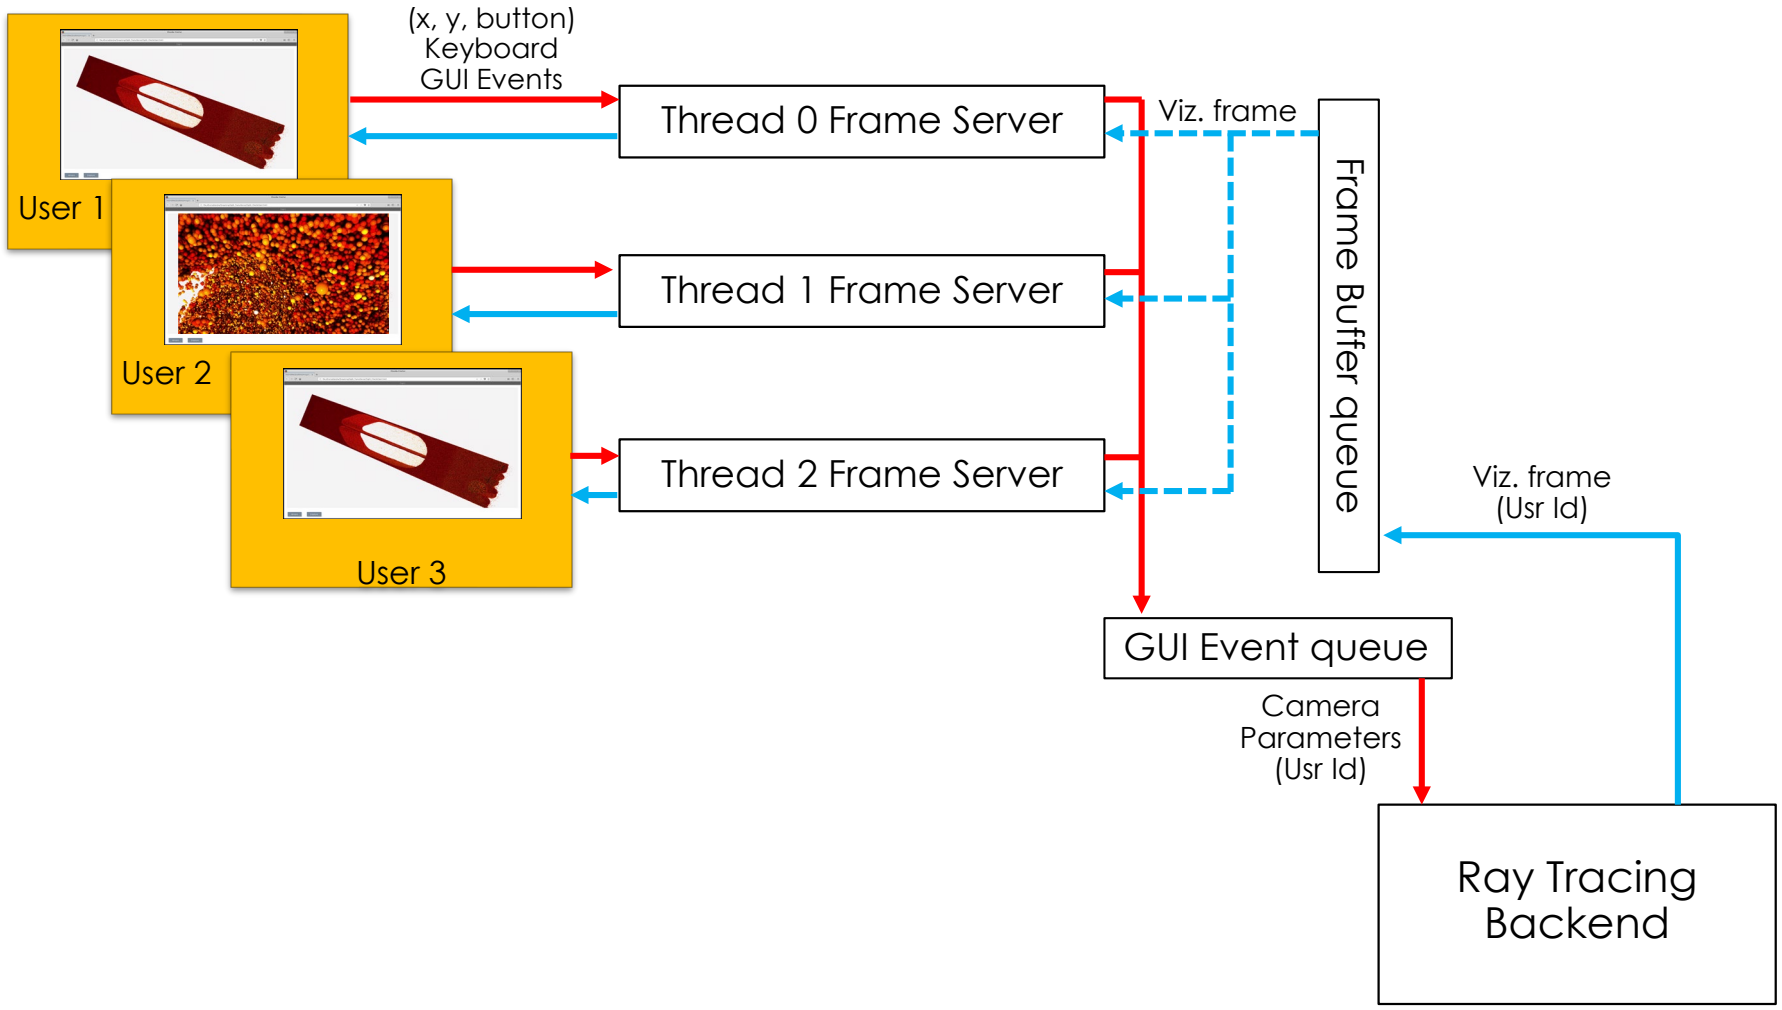

Enhancing SIGHT Frame Server
Simultaneous and independent user views

- Video

Discussion

- Further work
  - Multi-perspective/orthographic projections
  - Traditional case: stereoscopic projection
  - Annotations, sharing content between users, saving sessions
- How AI could help to improve remote visualization performance:
  - NVIDIA NGX Tech
- AI Up-Res
  - Improve compression rates
  - Different resolutions
    - Render and stream @ 720p, decoding at HD, 2K, 4K according to each user display
- AI Slow-Mo
  - Render at low framerates
- Optix Denoiser
  - Ray tracing converge faster

Thanks!

Benjamin Hernandez
Advanced Data and Workflows Group
Oak Ridge National Laboratory
hernandezarb@ornl.gov

- Acknowledgments

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Datasets provided by Cheng-Yu Shi and Leonid Zhigilei, Computational Materials Group at University of Virginia, INCITE award MAT130.

HARDWARE-ACCELERATED MULTI-TILE STREAMING FOR REALTIME REMOTE VISUALIZATION

Tim Biedert, 03/20/2019

INSTANT VISUALIZATION FOR FASTER SCIENCE

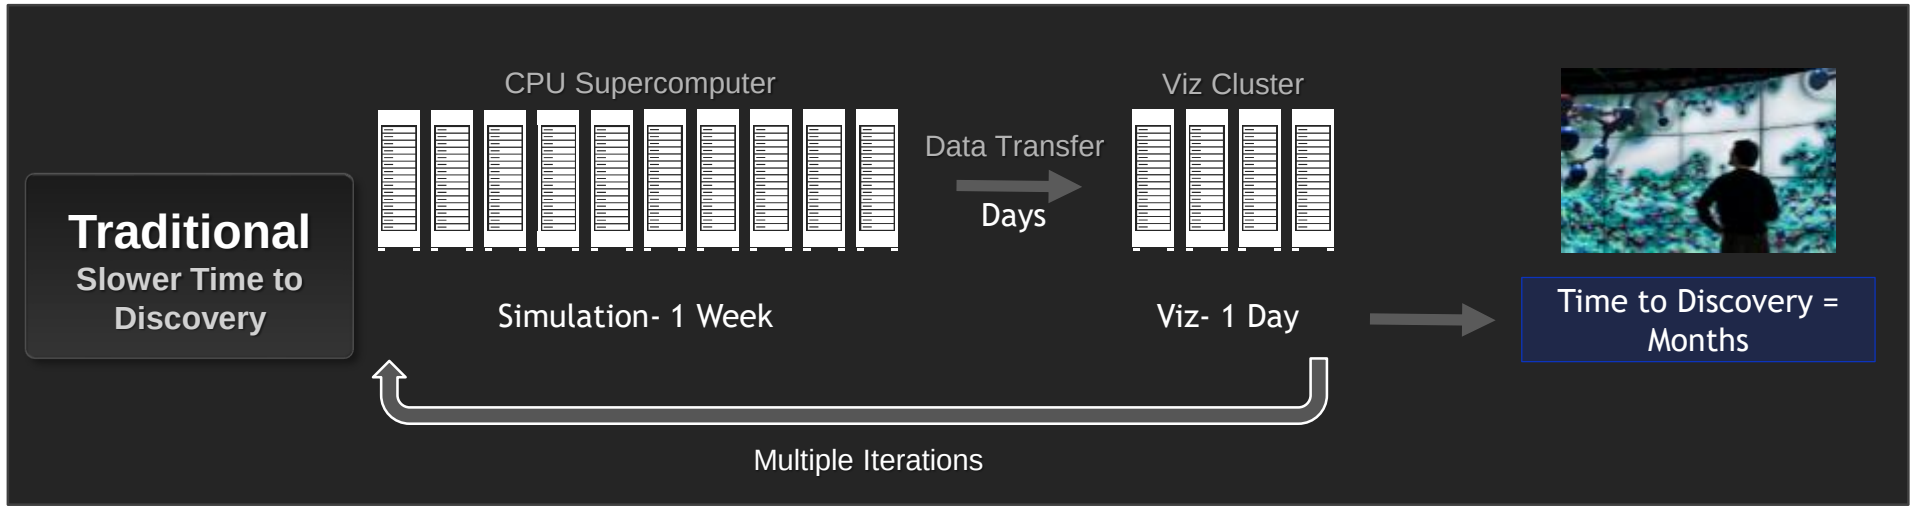

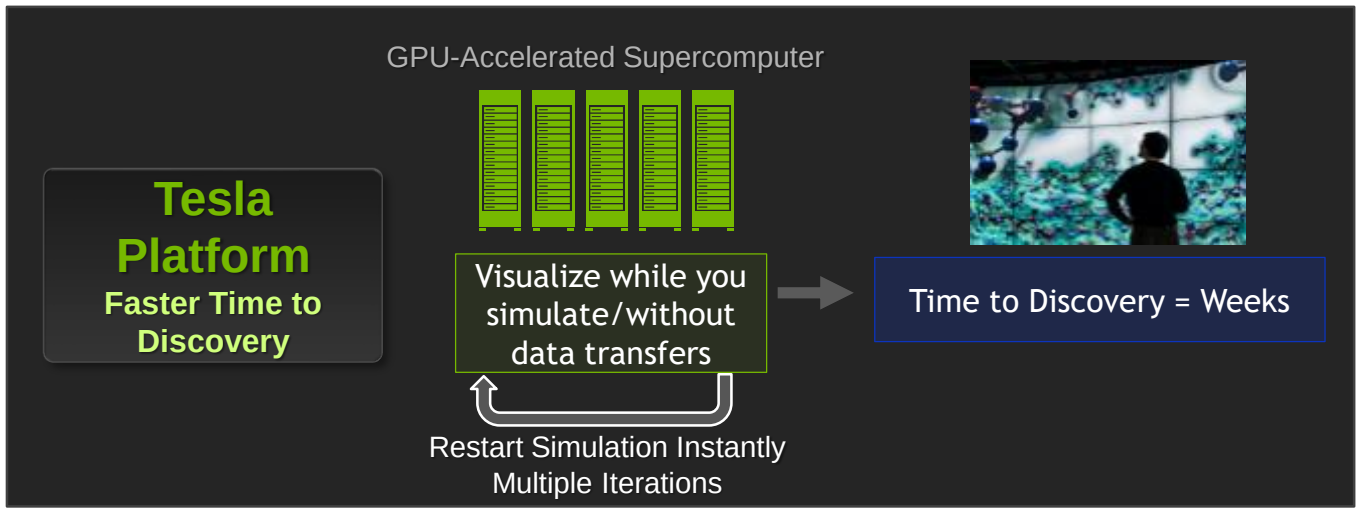

Value Proposition

- Interactivity
- Scalability
- Flexibility

SUPPORTING MULTIPLE VISUALIZATION WORKFLOWS

LEGACY WORKFLOW

Separate compute & vis system

Communication via file system

PARTITIONED SYSTEM

Different nodes for different roles

Communication via high-speed network

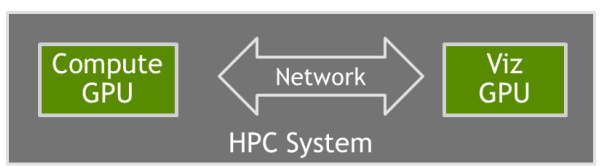

CO-PROCESSING

Compute and visualization on same GPU

Communication via host-device transfers or memcpy

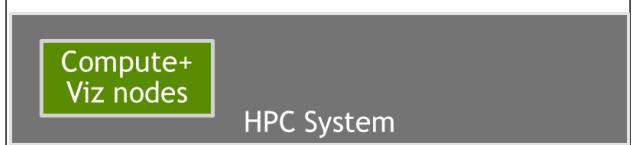

VISUALIZATION-ENABLED SUPERCOMPUTERS

CSCS Piz Daint

Galaxy formation

http://blogs.nvidia.com/blog/2014/11/19/gpu-in-situ-milky-way/

NCSA Blue Waters

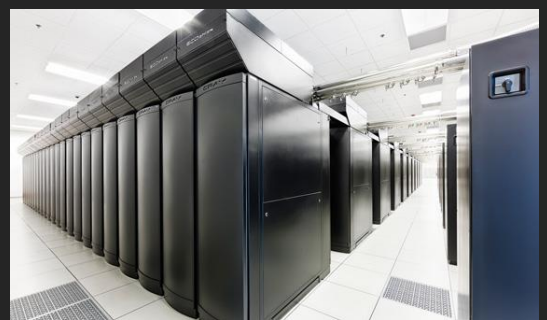

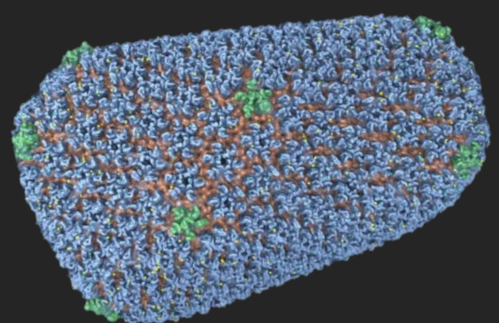

Molecular dynamics

http://devblogs.nvidia.com/parallelforall/hpc-visualization-nvidia-tesla-gpus/

ORNL Titan

Cosmology

http://www.sdav-scidac.org/29-highlights/visualization/66-accelerated-cosmology-data-anal.html

ROAD TO EXASCALE

Volta to Fuel Most Powerful US Supercomputers

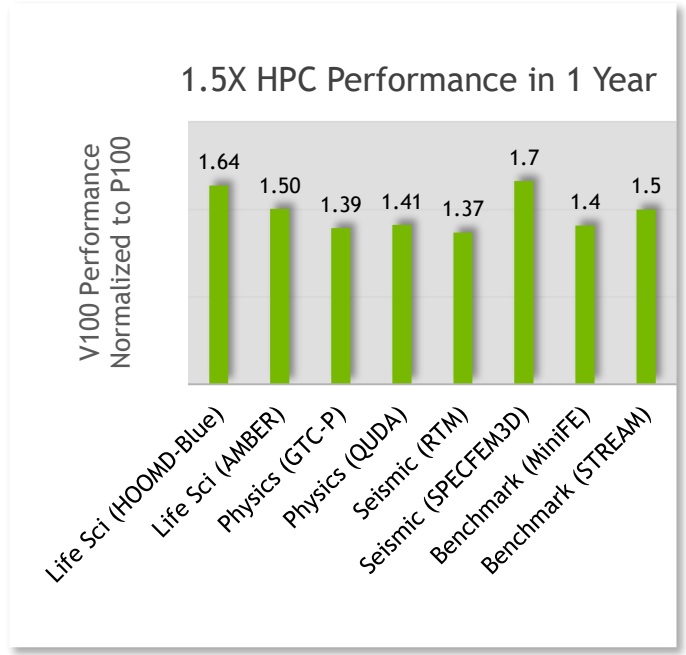

System Config Info: 2X Xeon E5-2690 v4, 2.6GHz, w/ 2X Tesla P100 or V100.

NVIDIA

VISUALIZATION TRENDS

New Approaches Required to Solve the Remoting Challenge

Increasing data set sizes

In-situ scenarios

Interactive workflows

New display technologies

Globally distributed user bases

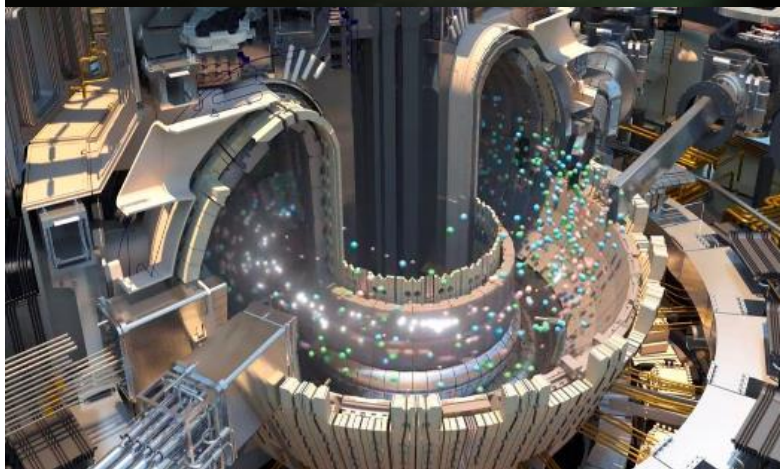

STREAMING

Benefits of Rendering on Supercomputer

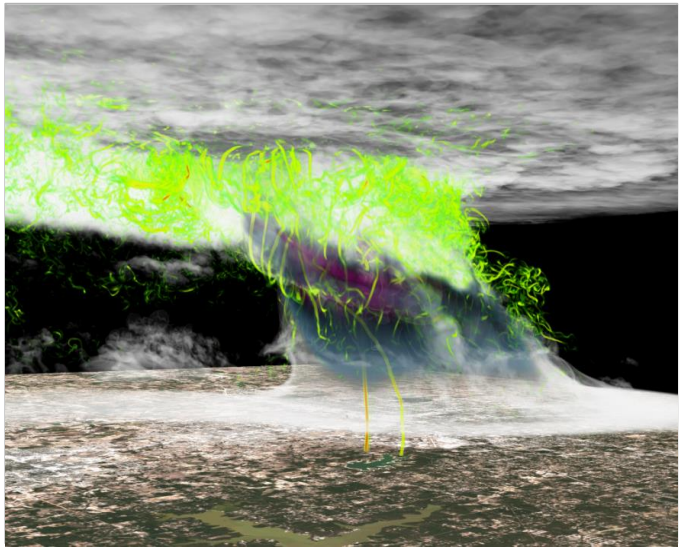

Scale with Simulation

No Need to Scale Separate Vis Cluster

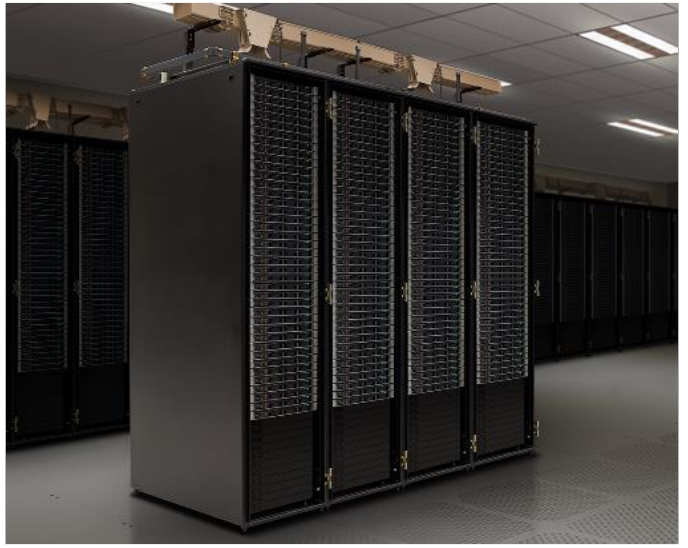

Cheaper Infrastructure

All Heavy Lifting Performed on the Server

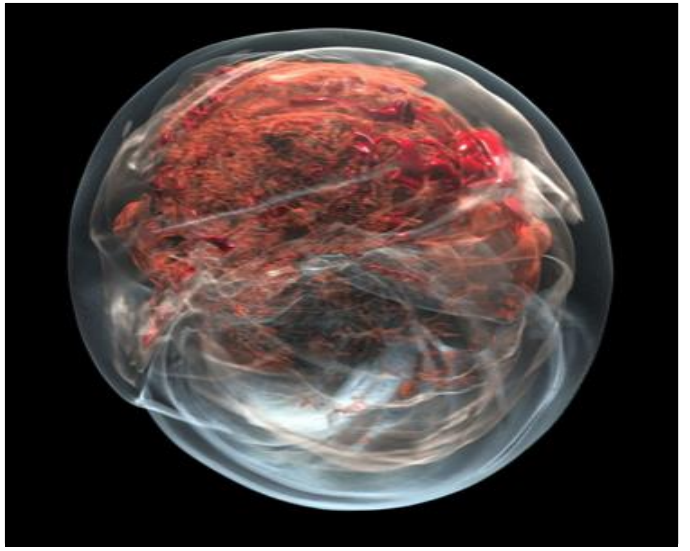

Interactive High-Fidelity Rendering

Improves Perception and Scientific Insight

FLEXIBLE GPU ACCELERATION ARCHITECTURE

Independent CUDA Cores & Video Engines

Decode HW*

Formats:
- MPEG-2
- VC1
- VP8
- VP9
- H.264
- H.265
- Lossless

Bit depth:
- 8 bit
- 10 bit
- 12 bit

Color**
- YUV 4:2:0
- YUV 4:4:4

Resolution
- Up to 8K***

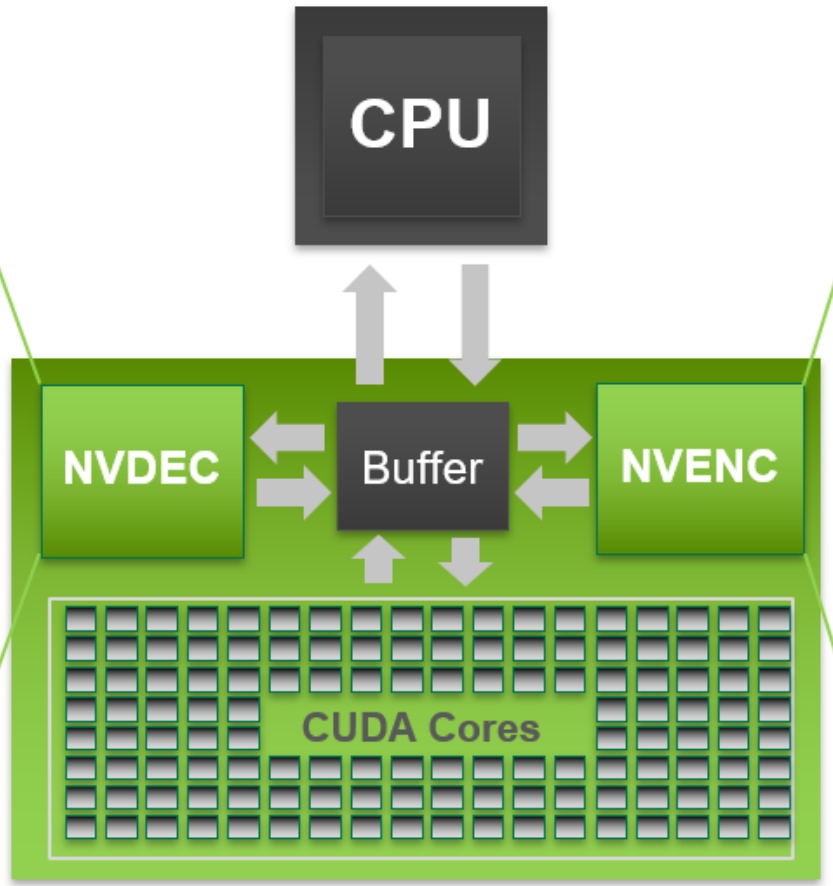

Encode HW*

Formats:
- H.264
- H.265
- Lossless

Bit depth:
- 8 bit
- 10 bit

Color**
- YUV 4:4:4
- YUV 4:2:0

Resolution
- Up to 8K***

** Diagram represents support for the NVIDIA Turing GPU family*
*** 4:2:2 is not natively supported on HW*
**** Support is codec dependent*

CASE STUDY

Streaming of Large Tile Counts

Frame rates / Latency / Bandwidth

Synchronization

Comparison against CPU-based Compressors

Strong Scaling (Direct-Send Sort-First Compositing)

Hardware-Accelerated Multi-Tile Streaming for Realtime Remote Visualization

Tim Biedert, Peter Messmer, Tom Fogal, Christoph Garth

Eurographics Symposium on Parallel Graphics and Visualization (EGPGV) 2018 (Best Paper)

DOI: **10.2312/pgv.20181093**

CONCEPTUAL OVERVIEW

Asynchronous Pipelines

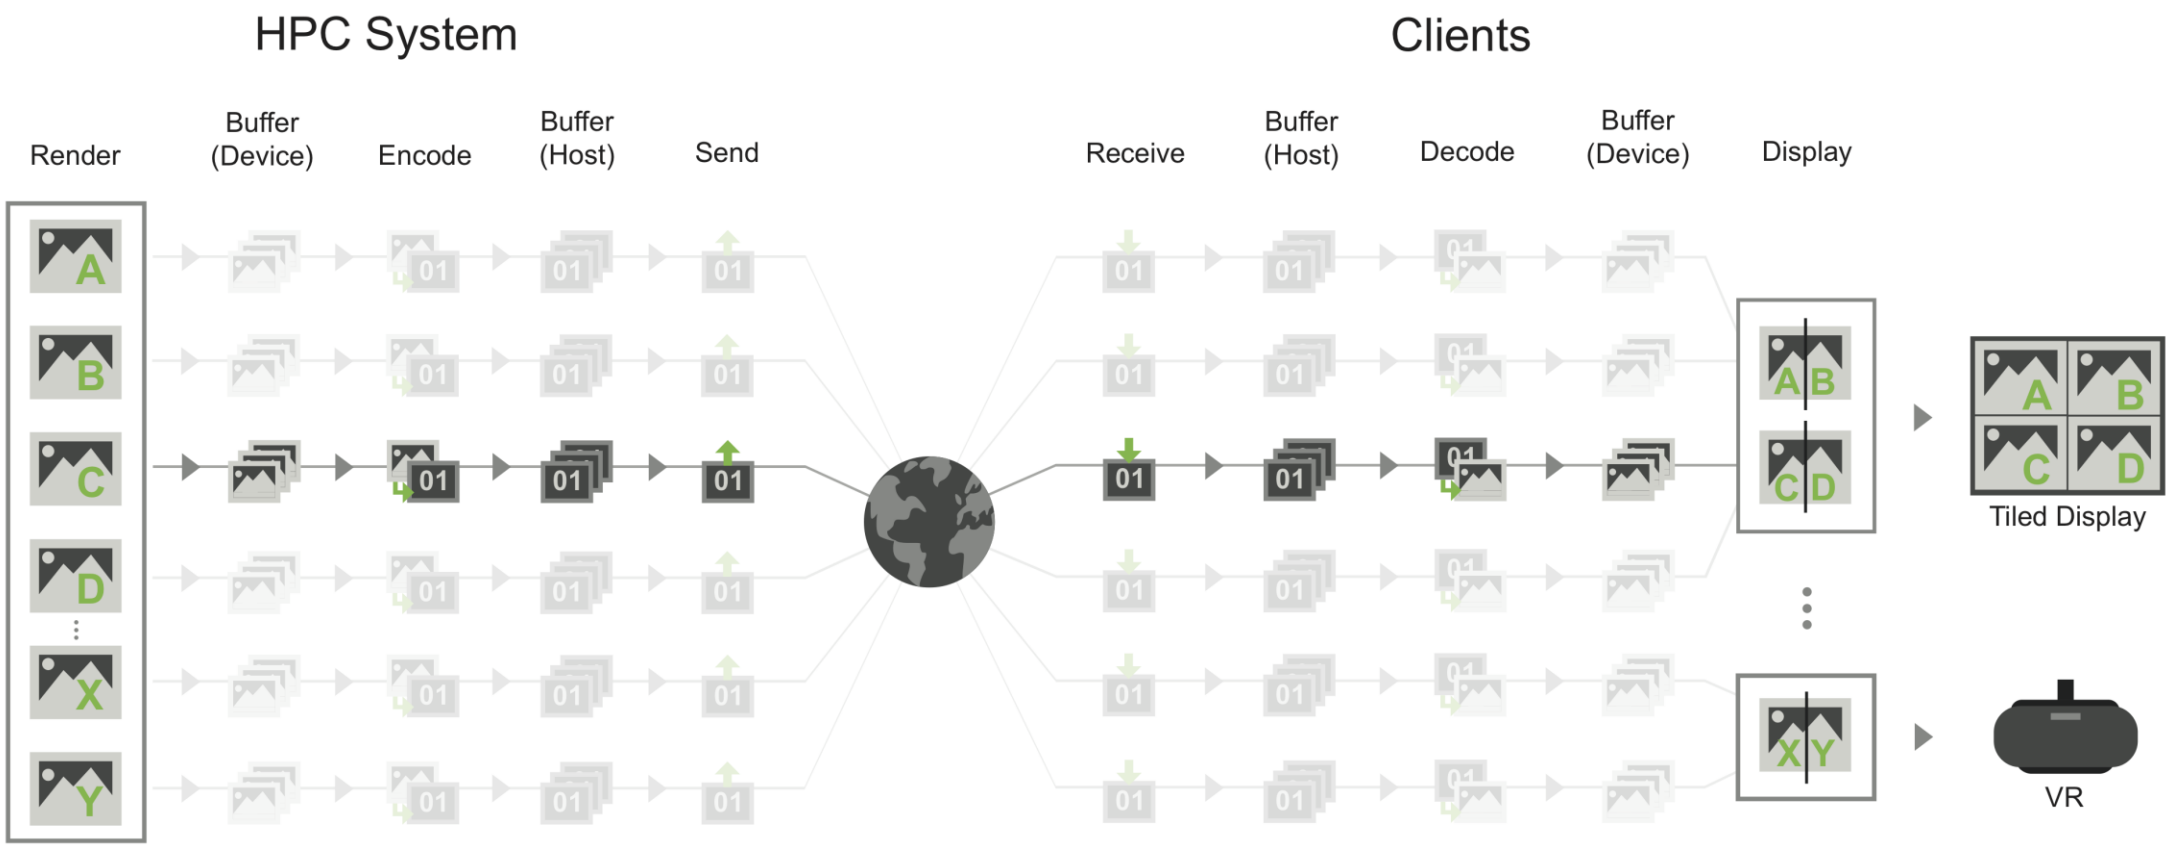

BENCHMARK SCENES

NASA Synthesis 4K

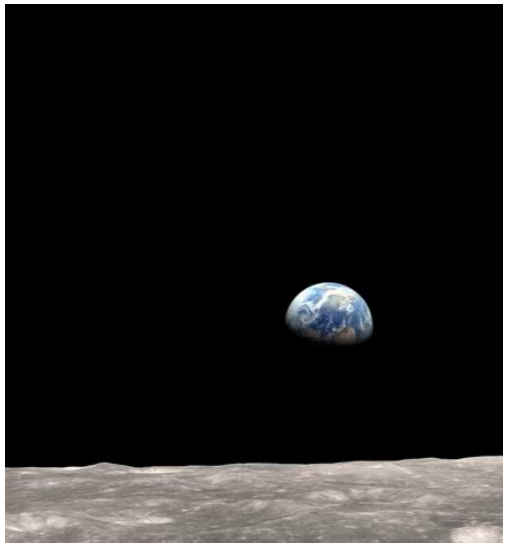

Space
Low Complexity

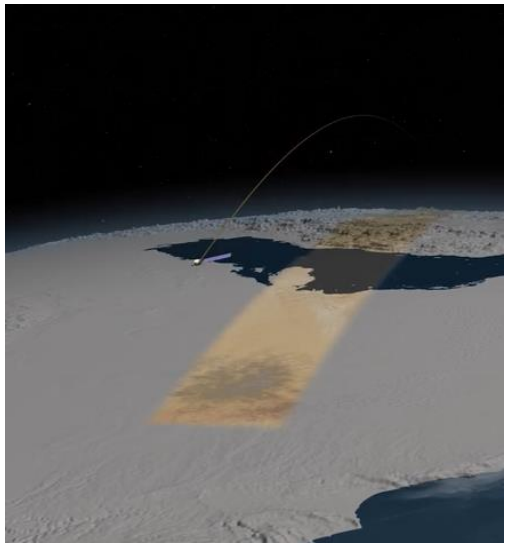

Orbit
Medium Complexity

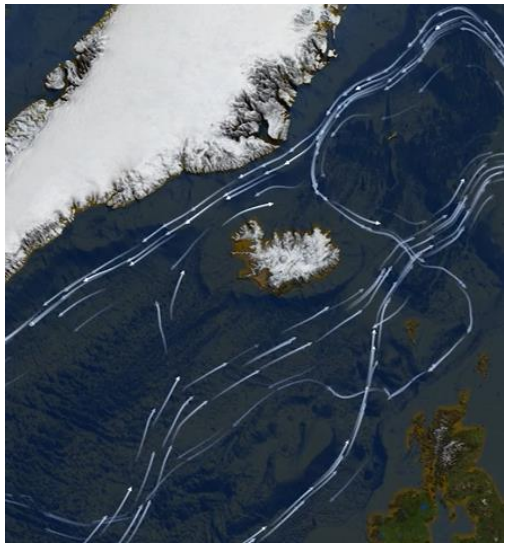

Ice
High Complexity

Streamlines
Extreme Complexity


CODEC
PERFORMANCE

BITRATE VS QUALITY

Structural Similarity Index (SSIM)

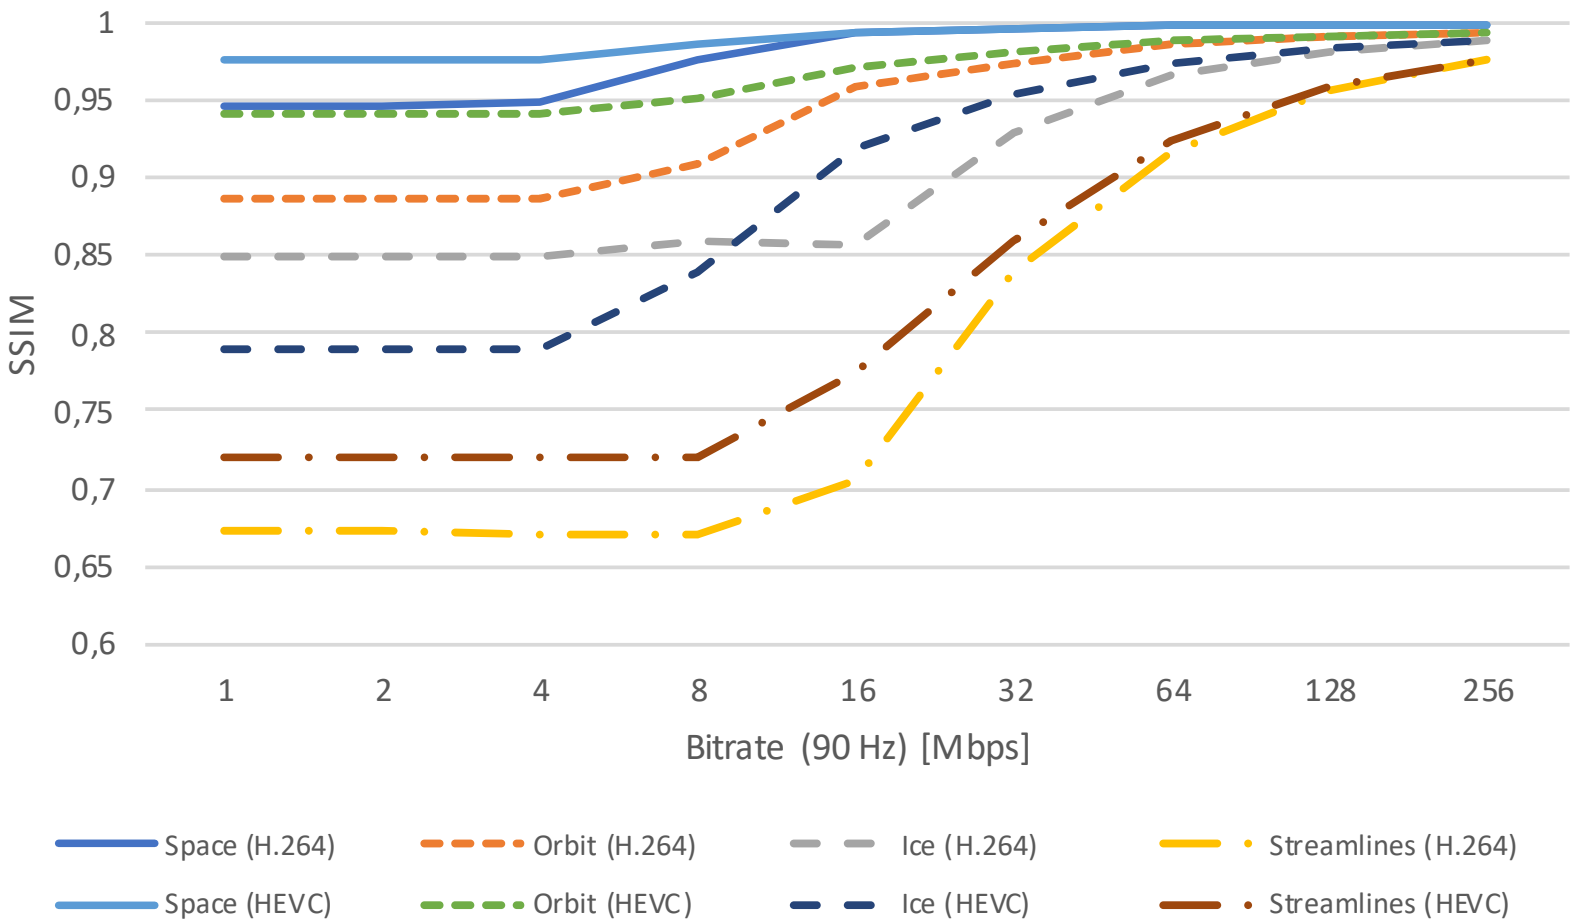

NVIDIA.

HARDWARE LATENCY

Encode HEVC - Encode H.264 - Decode HEVC - Decode H.264

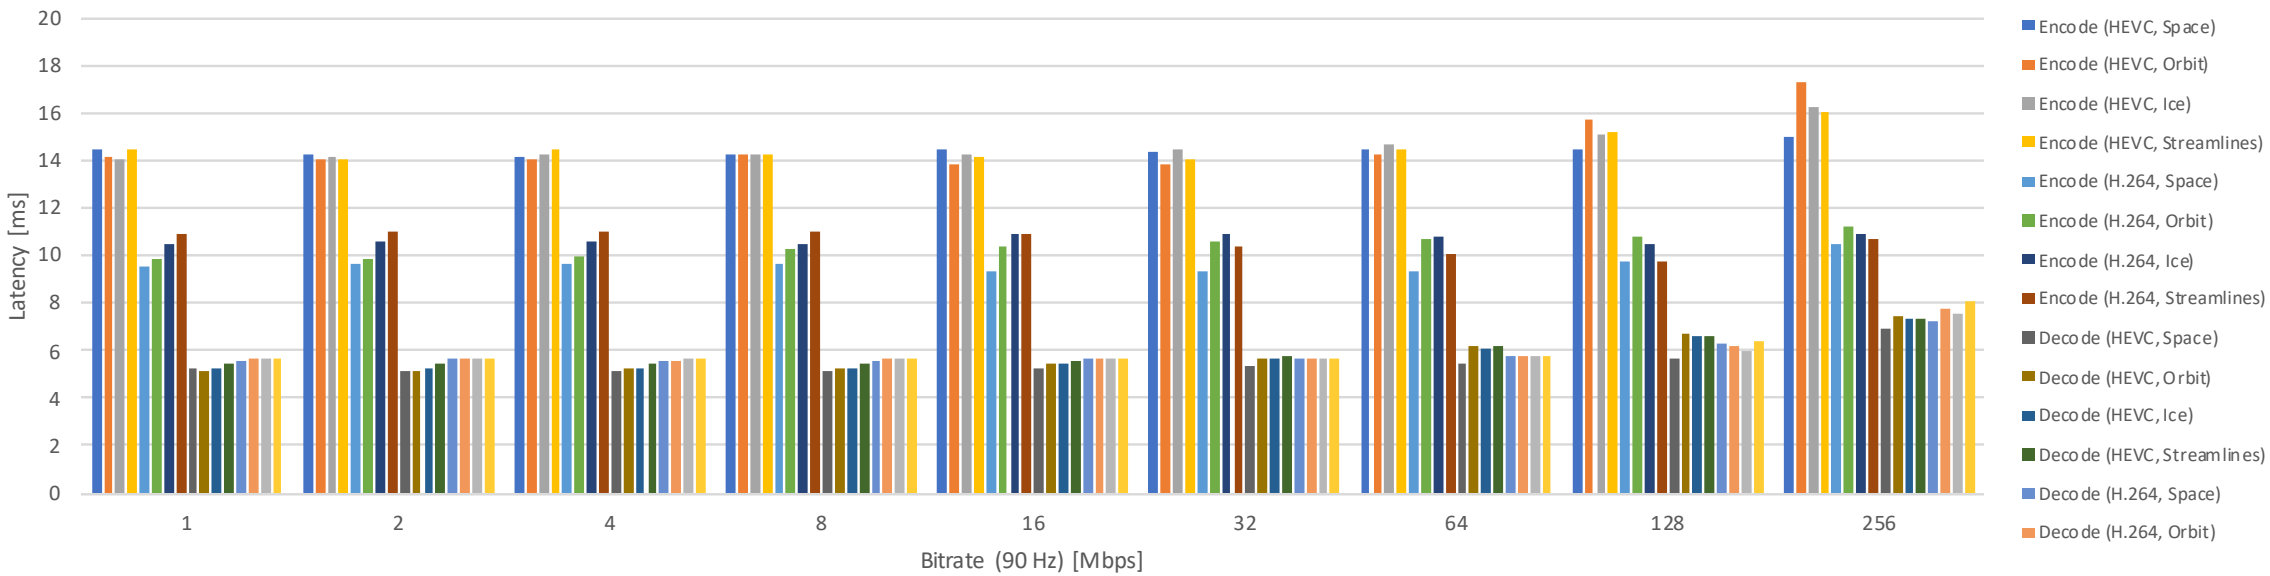

NVIDIA.

HARDWARE LATENCY

Decreases with Resolution

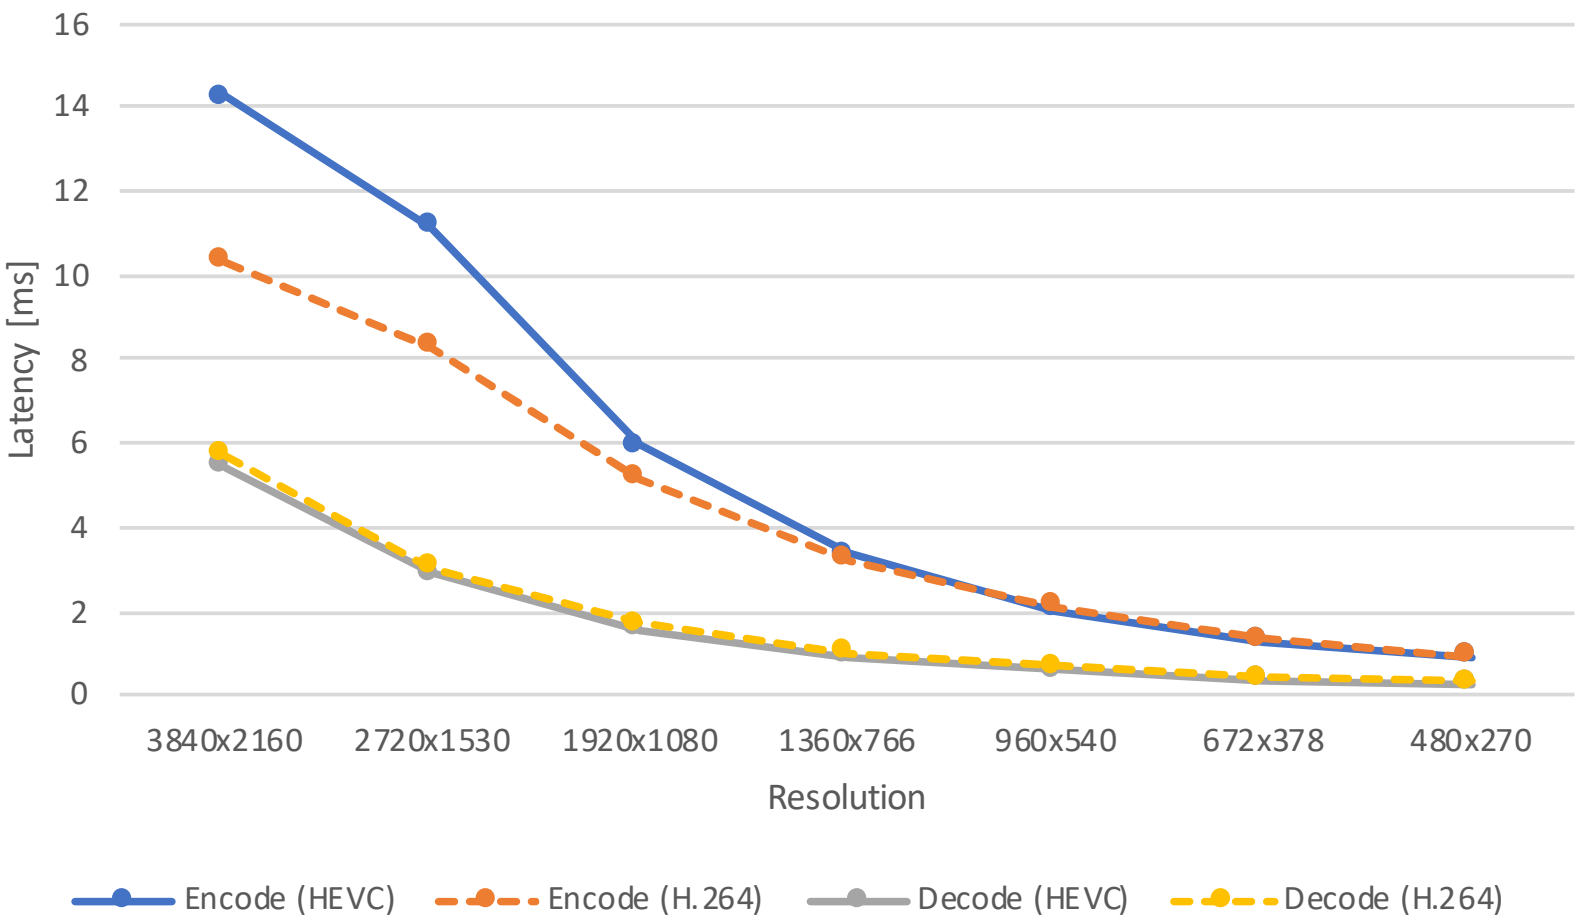

CPU-BASED COMPRESSION

Encode/Decode Latency and Bandwidth

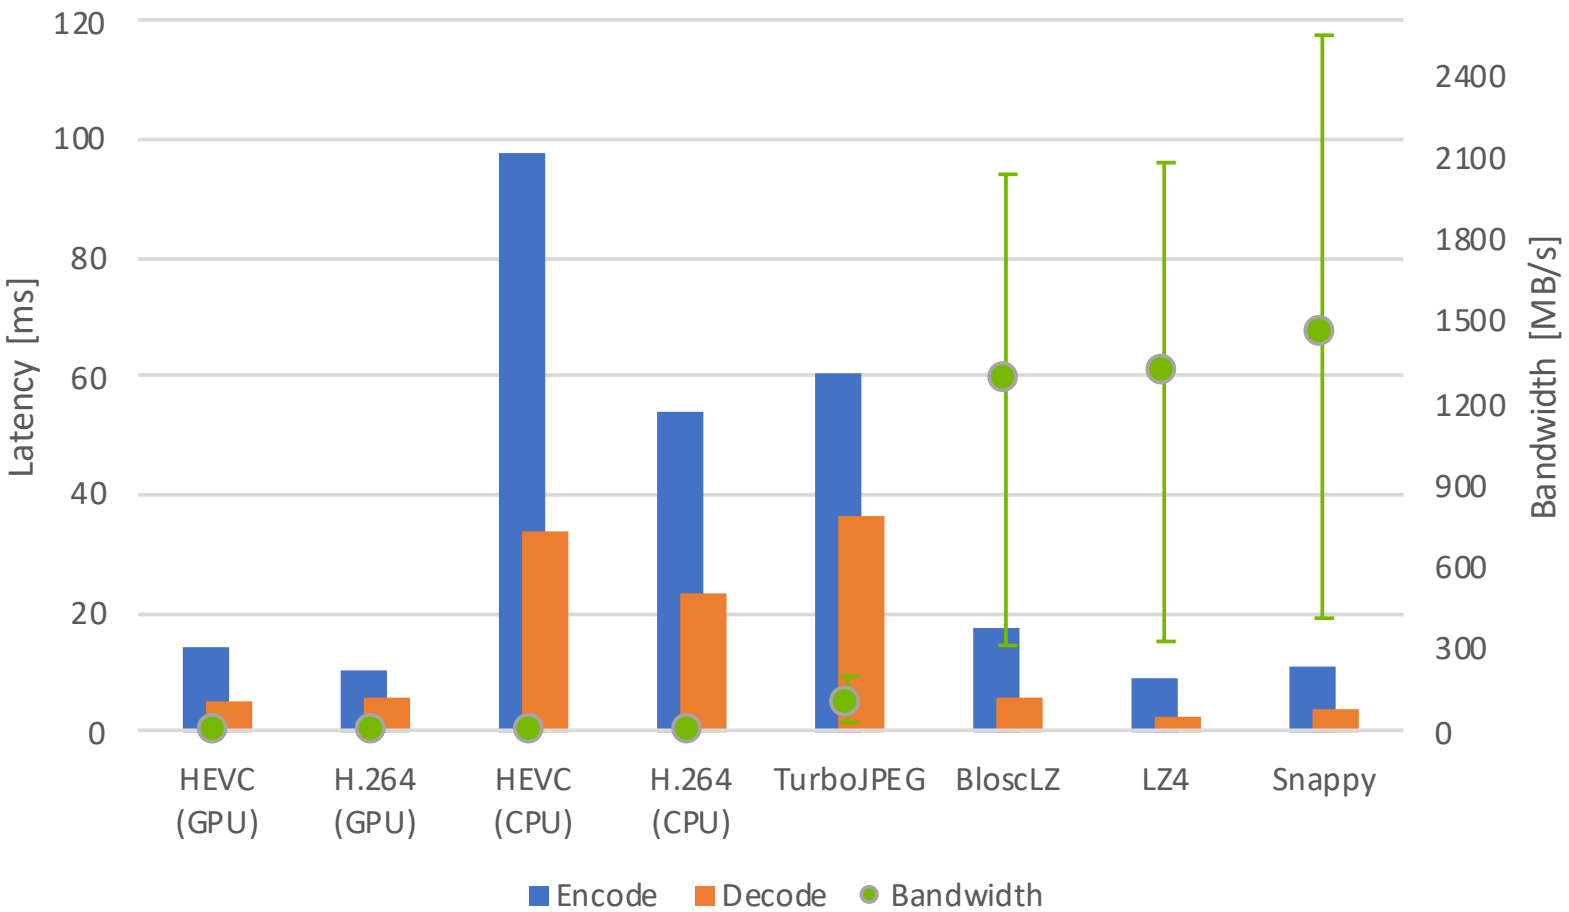


FULL TILES
STREAMING

N:N WITH SIMULATED NETWORK DELAY

Mean Frame Rates + Min/Max Ranges

Server: Piz Daint
Client: Piz Daint

Ice 4K

H.264: 32 Mbps
HEVC: 16 Mbps

Delay + 10% jitter

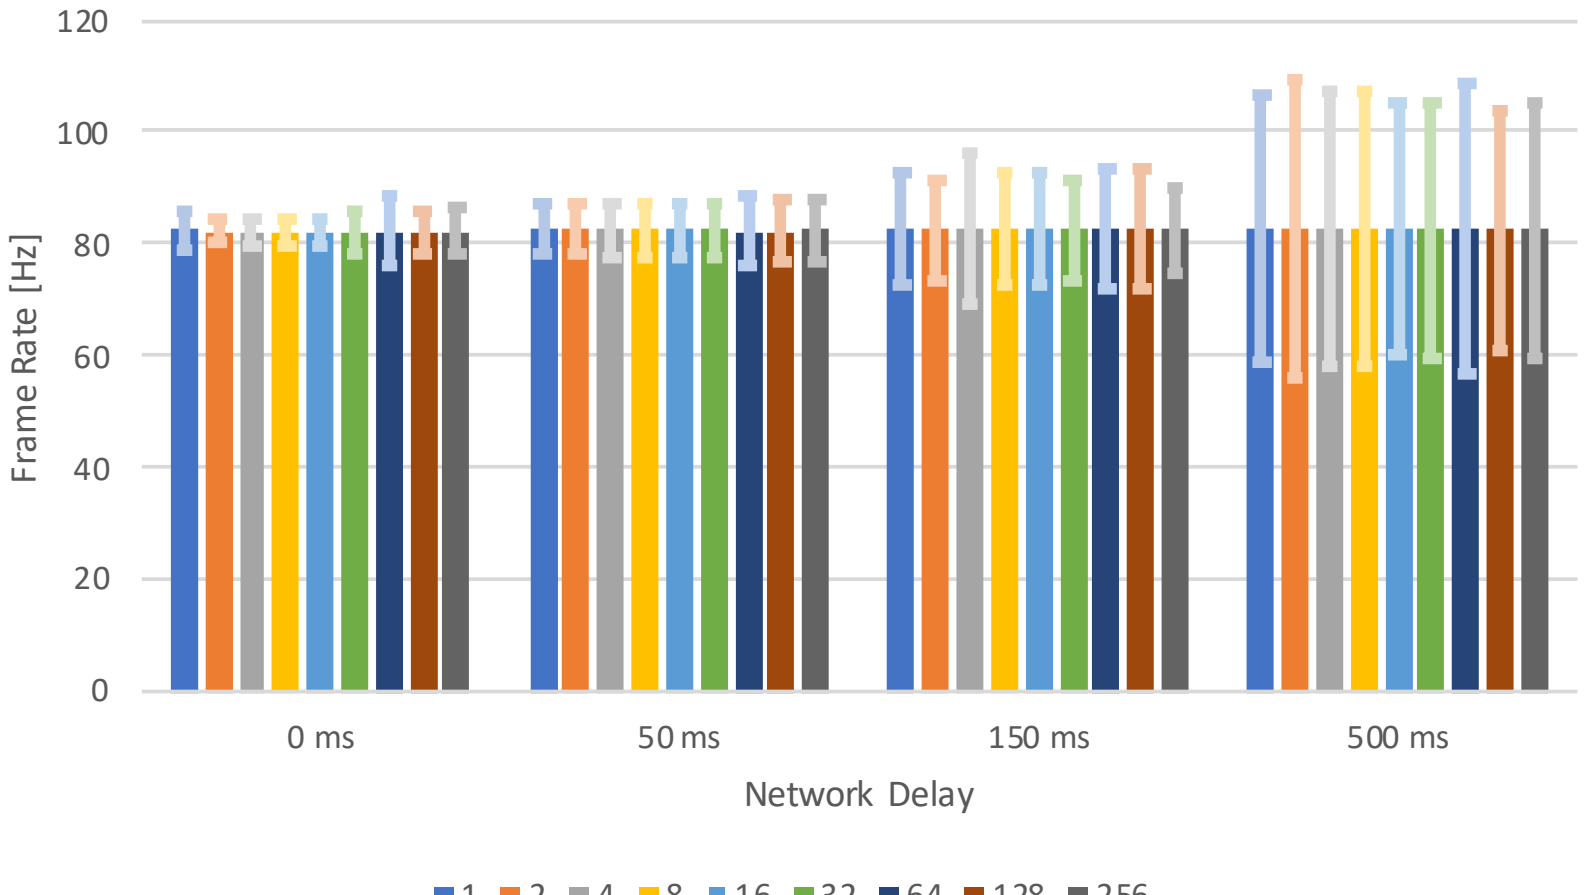

NVIDIA

N:N STREAMING

Pipeline Latencies

Server: Piz Daint
Client: Piz Daint

Ice 4K

H.264: 32 Mbps

MPI-based synchronization

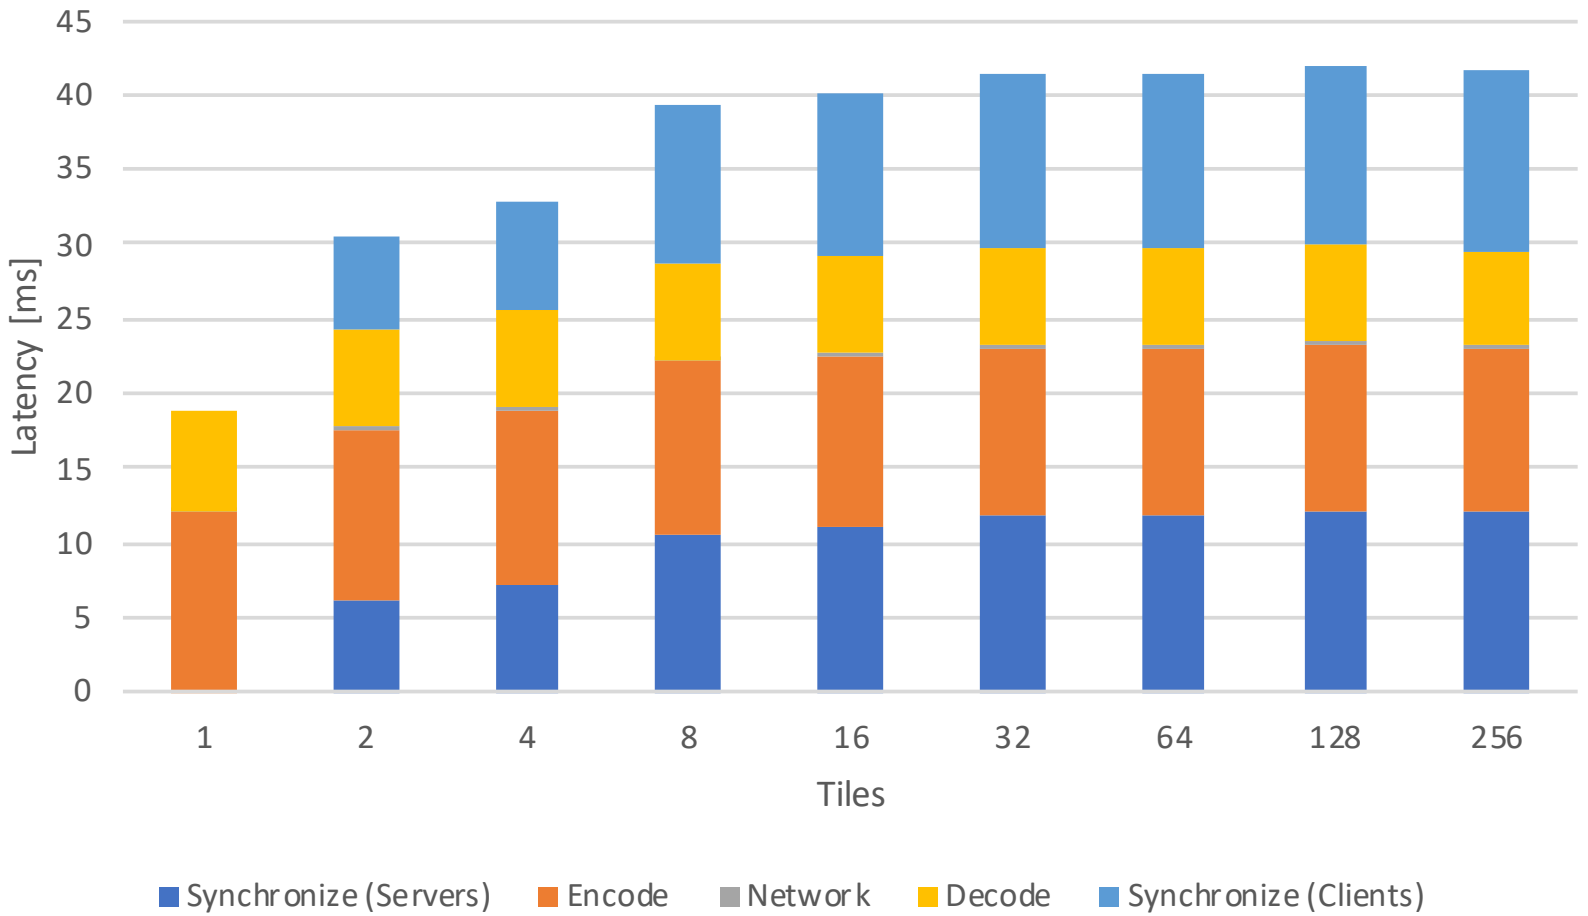

N:1 STREAMING

Client-Side Frame Rate

Server: Piz Daint
Clients: Site A (5 ms)
Site B (25 ms)

Ice 4K

H.264: 32 Mbps

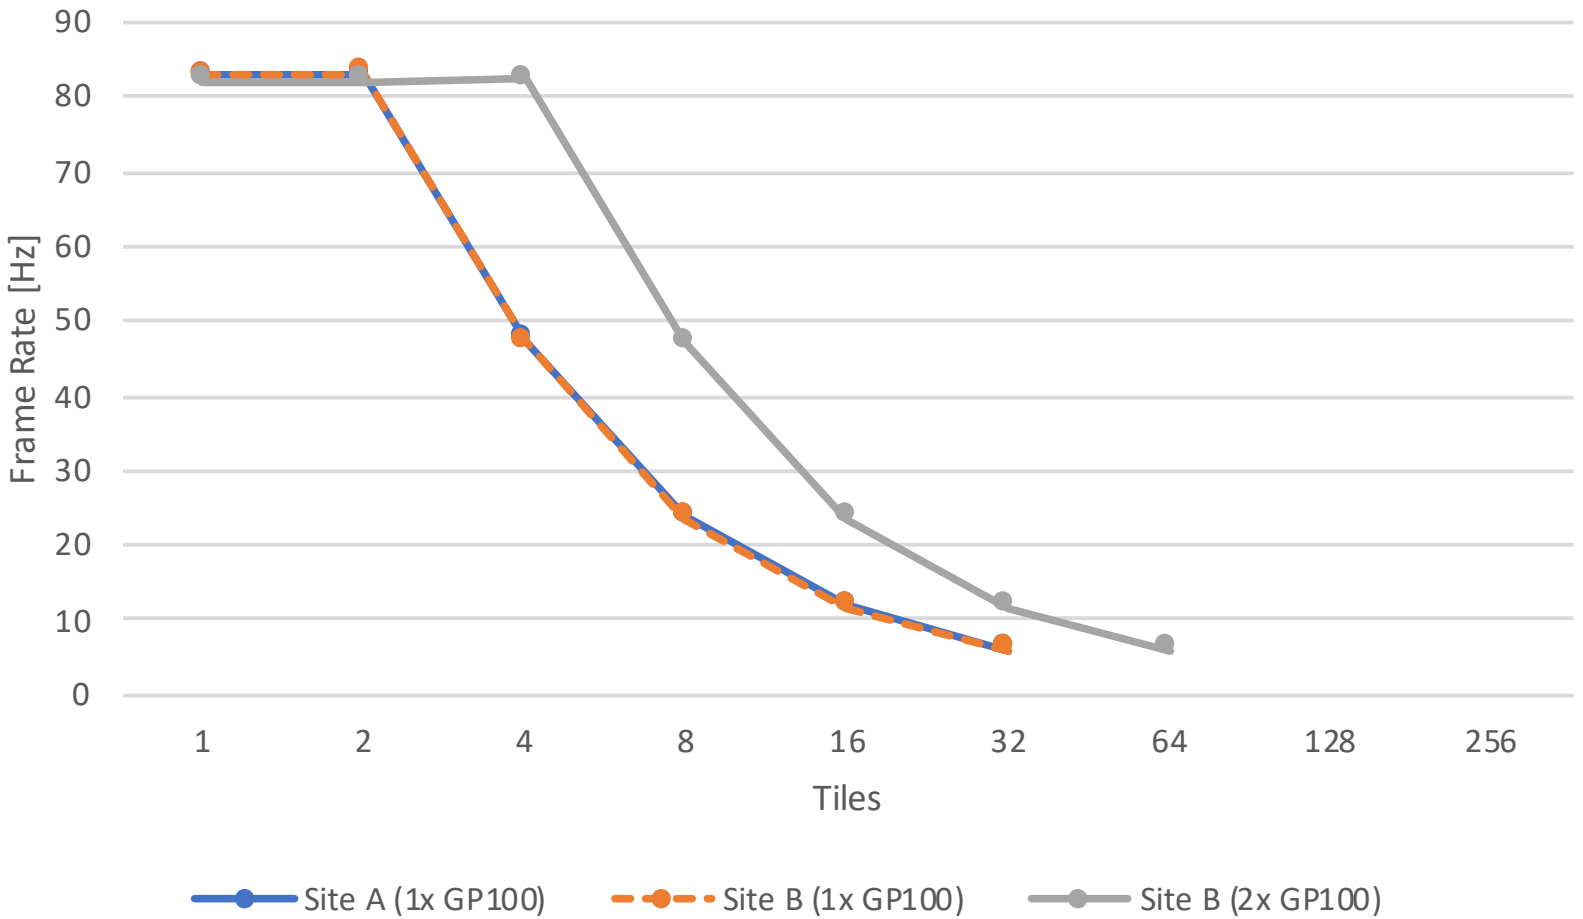

NVIDIA


STRONG SCALING

N:1 STRONG SCALING

Client-Side Frame Rate

Server: Piz Daint
Clients: Site A (5 ms)
Site B (25 ms)

Ice 4K

H.264: 32 Mbps

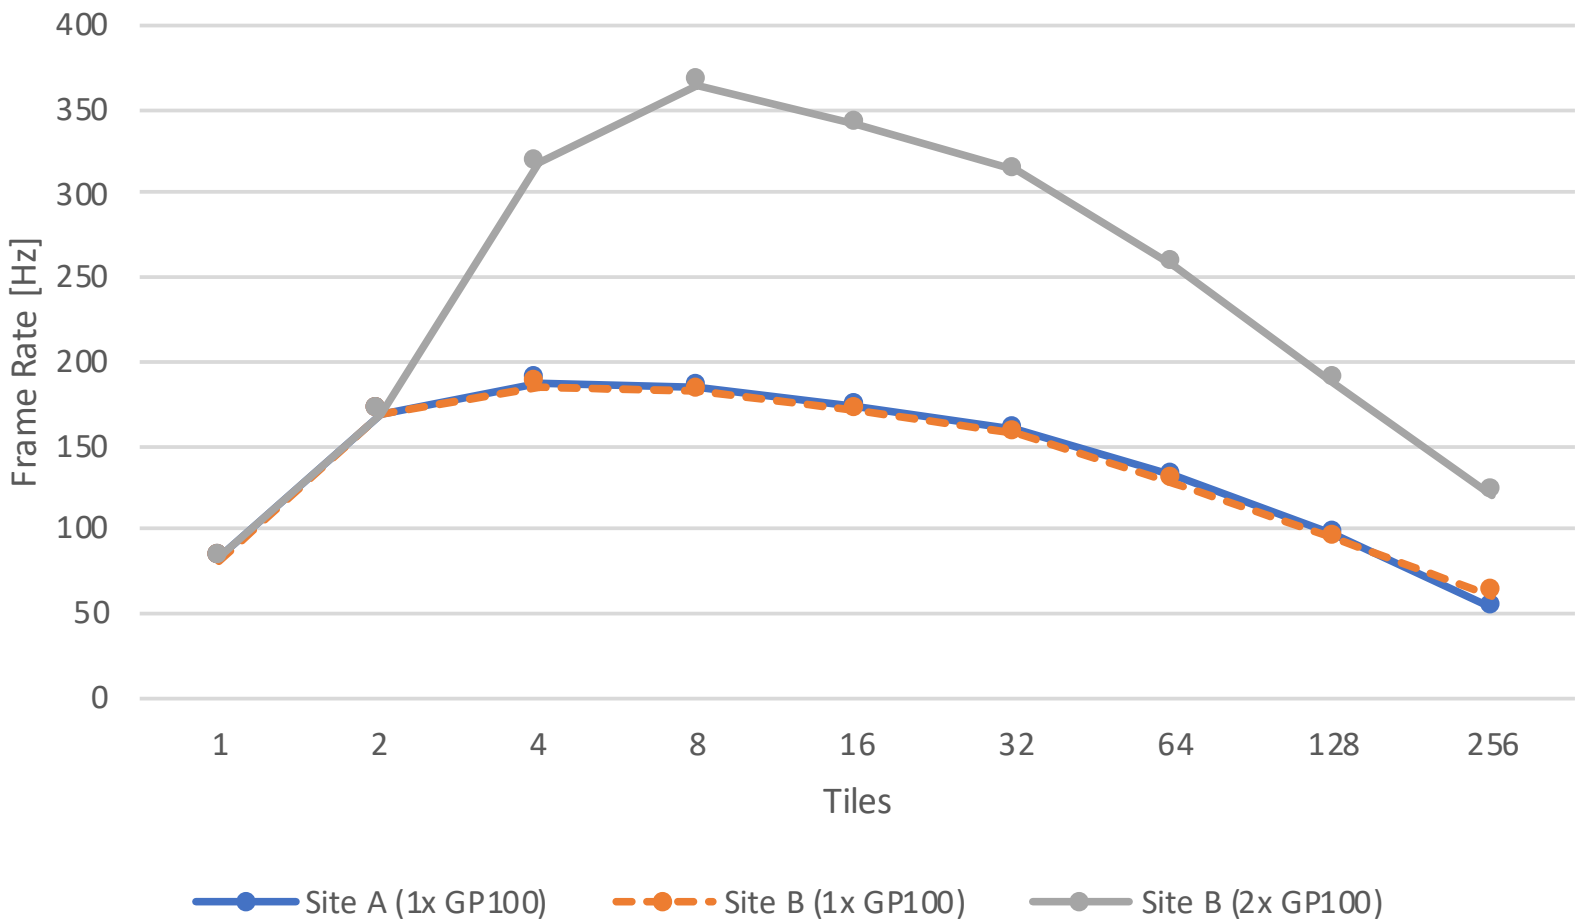

NVIDIA

N:1 STRONG SCALING

Client-Side Frame Rate For Different Bitrates

Server: Piz Daint
Client: Site A (5 ms)

Ice 4K

H.264

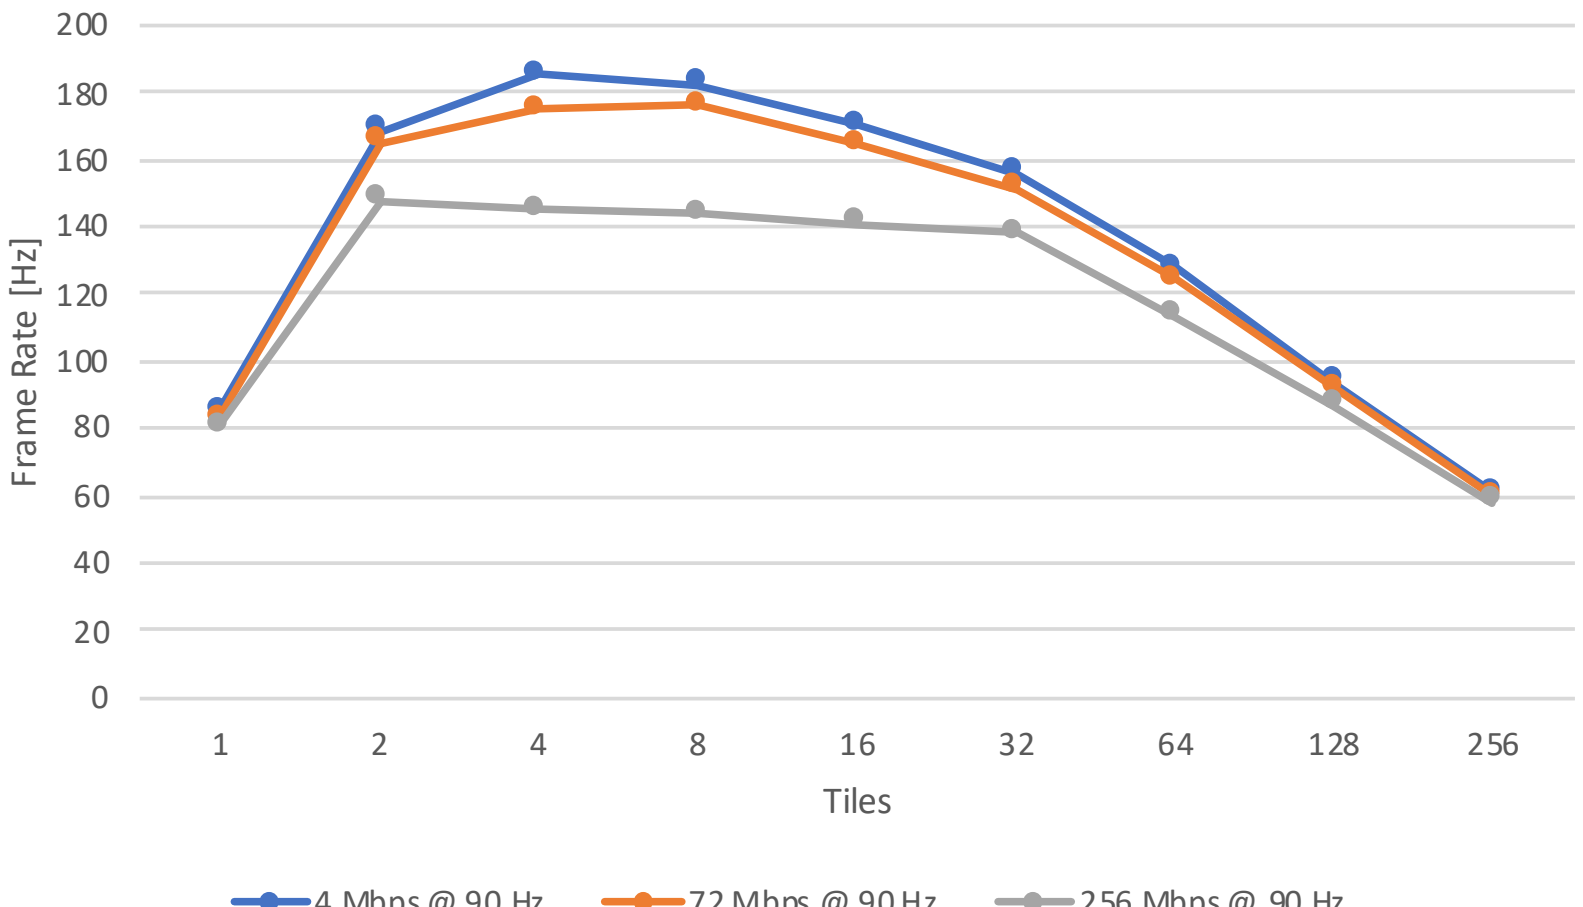

EXAMPLE: OPTIX PATHTRACER

Rendering Highly Sensitive to Tile Size

Server: Piz Daint
Client: Site A (5 ms)
H.264

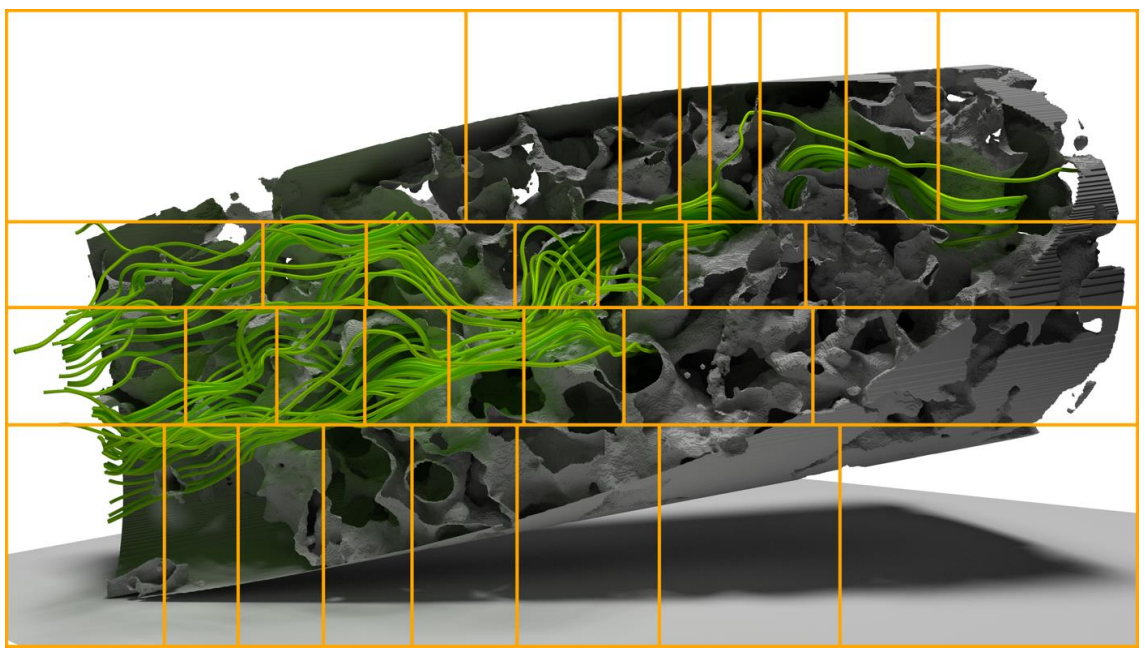

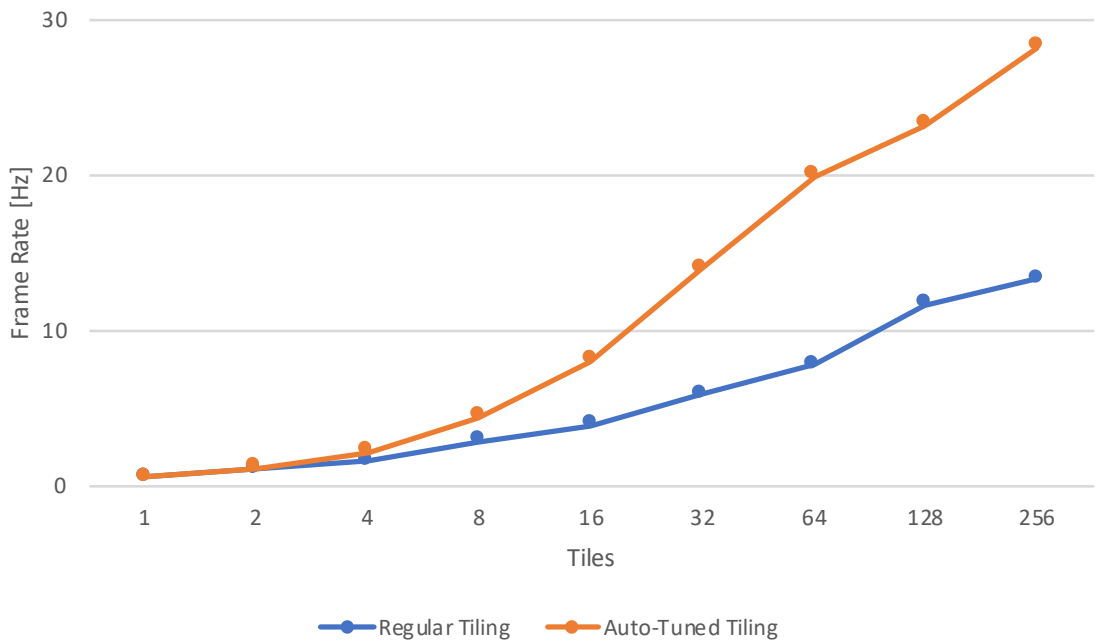


INTEROPERABILITY

STANDARD-COMPLIANT BITSTREAM

Web Browser Streaming Example

EGL-based shim GLUT

Stream unmodified simpleGL example from headless node to web browser (with interaction!)

JavaScript client

WebSocket-based bidirectional communication

On-the-fly **MP4** wrapping of **H.264**

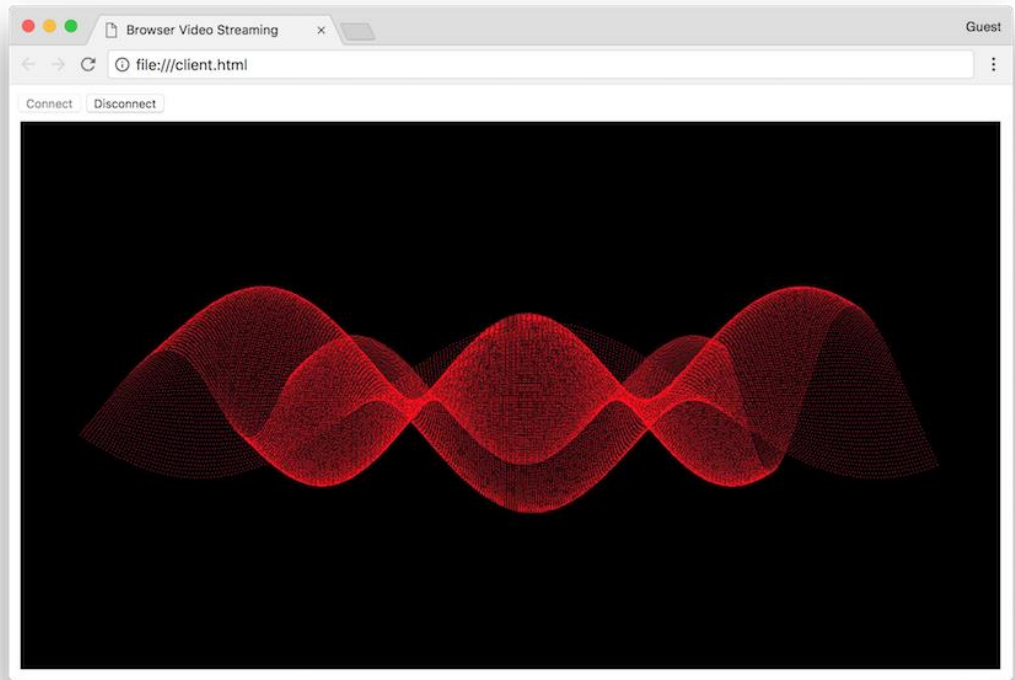

RESOURCES

VIDEO CODEC SDK

APIs For Hardware Accelerated Video Encode/Decode

What's New with Turing GPUs and Video Codec SDK 9.0

- Up to 3x decode throughput with multiple decoders on professional cards (Quadro & Tesla)
- Higher quality encoding - H.264 & H.265
- Higher encoding efficiency (15% lower bitrate than Pascal)
- HEVC B-frames support
- HEVC 4:4:4 decoding support

NVIDIA GeForce Now is made possible by leveraging **NVENC** in the datacenter and streaming the result to end clients

https://developer.nvidia.com/nvidia-video-codec-sdk

NVPIPE

A Lightweight Video Codec SDK Wrapper

Simple C API

H.264, HEVC

RGBA32, uint4, uint8, uint16

Lossy, Lossless

Host/Device memory, OpenGL textures/PBOs

https://github.com/NVIDIA/NvPipe

Issues? Suggestions? Feedback welcome!

```
#include <NvPipe.h>


// Encode
NvPipe* encoder = NvPipe_CreateEncoder(NVPIPE_RGBA32,
    NVPIPE_HEVC, NVPIPE_LOSSY, 32 * 1000 * 1000, 90);

while (...)
{
    uint64_t compressedSize = NvPipe_Encode(encoder,
        rgba, buffer, bufferSize, width, height);
    ...
}

NvPipe_Destroy(encoder);


// Decode
NvPipe* decoder = NvPipe_CreateDecoder(NVPIPE_RGBA32,
    NVPIPE_HEVC);

while (...)
{
    NvPipe_Decode(decoder, buffer, compressedSize,
        rgba, width, height);
    ...
}

NvPipe_Destroy(decoder);
```

Conclusion

GPU-accelerated video compression opens up novel and fast solutions to the large-scale remoting challenge

Video Codec SDK
https://developer.nvidia.com/nvidia-video-codec-sdk

NvPipe
https://github.com/NVIDIA/NvPipe

We want to help you solve your large-scale vis problems on NVIDIA!

Tim Biedert
tbiedert@nvidia.com


NVIDIA