

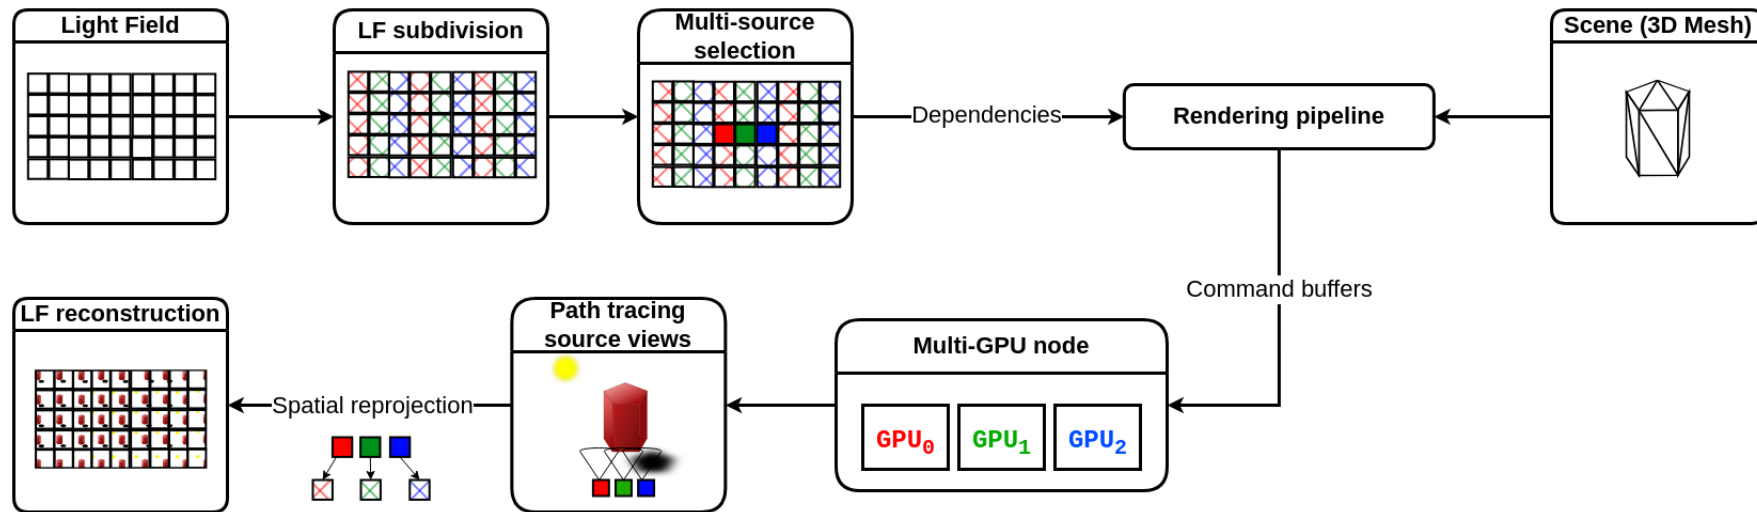
Interactive Multi-GPU Light Field Path Tracing Using Multi-Source Spatial Reprojection

Erwan Leria¹, Markku Mäkitalo¹, Pekka Jääskeläinen¹, Mårten Sjöström², Tingting Zhang²

¹Tampere University, Finland

²Mid Sweden University, Sweden

(erwan.leria@tuni.fi)



Introduction

Light field displays




*Boomerang video of retinal image at different accommodation depth passing through a numerical **near-eye** light field display*



*Picture of the San Miguel scene in a **multiscope** display (Looking Glass 32")*

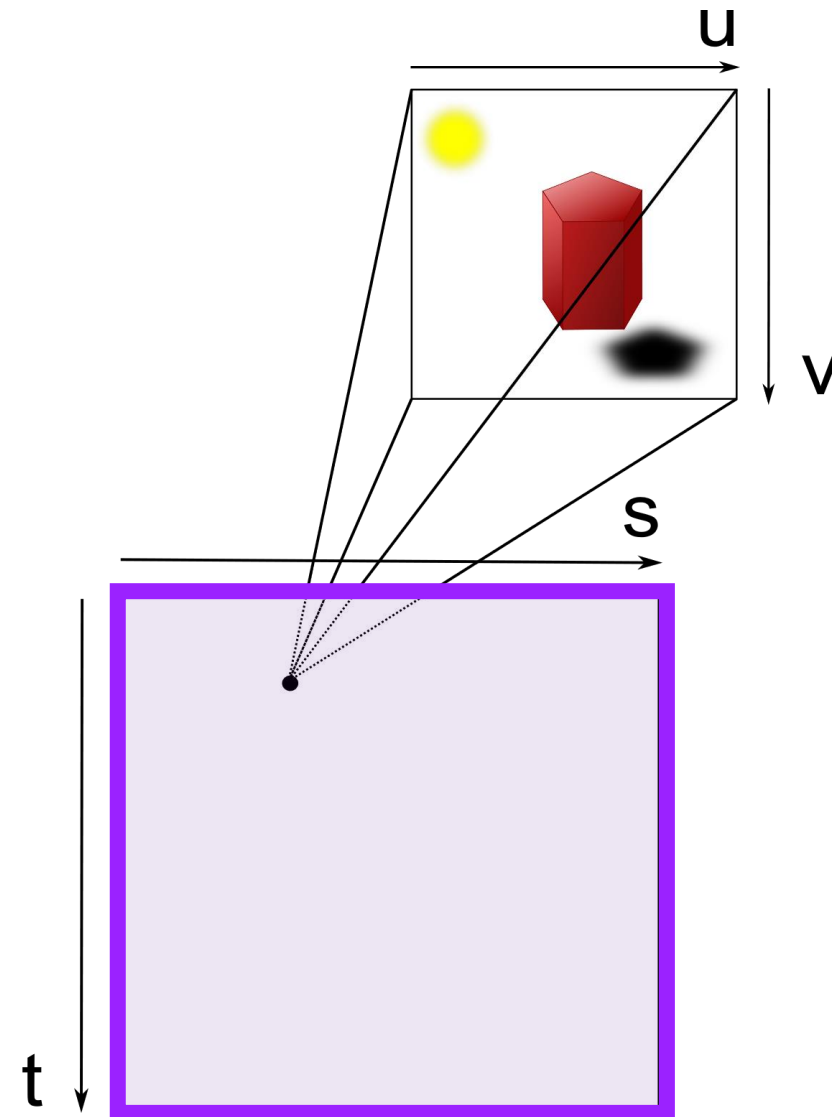
Light fields




Light field 2-plane parameterization $L(s, t, u, v)$

Light fields

- Collection of multiple views (pinhole cameras)




Light field 2-plane parameterization $L(s, t, u, v)$




Reducing computational cost

- Spatial reprojection of redundant pixels between views (Adelson, 1993)



Reducing computational cost

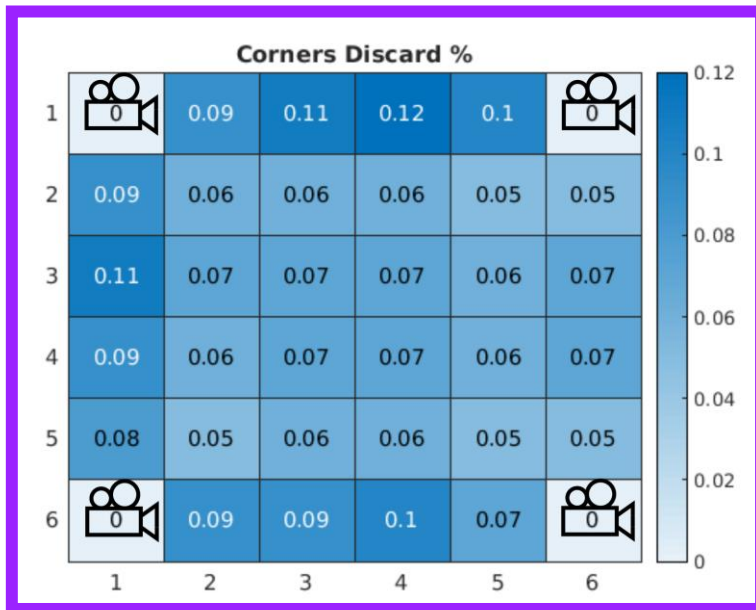
- **Spatial reprojection of redundant pixels between views** (Adelson, 1993)
- **Disocclusions** → Hidden parts of objects (in the source) cannot be reprojected
 - Disocclusion pixels are **discarded** in the target



Discarded pixels

- Depends on:

- Number of sources

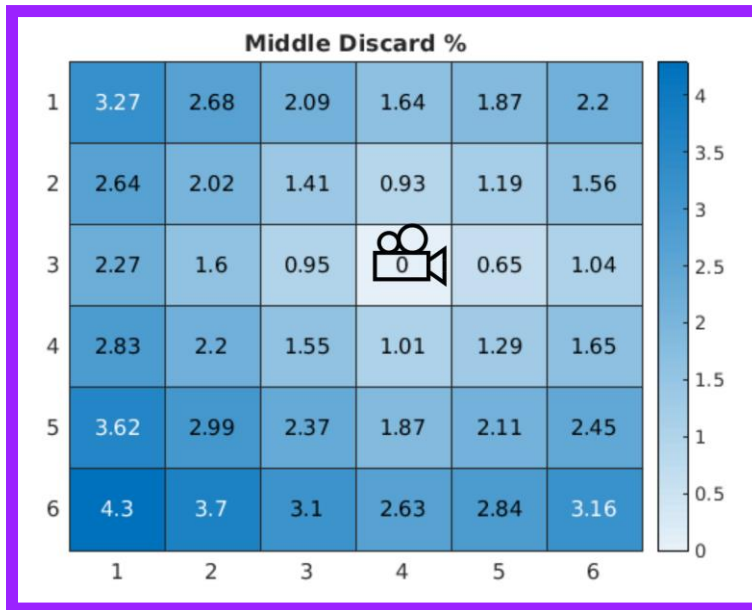
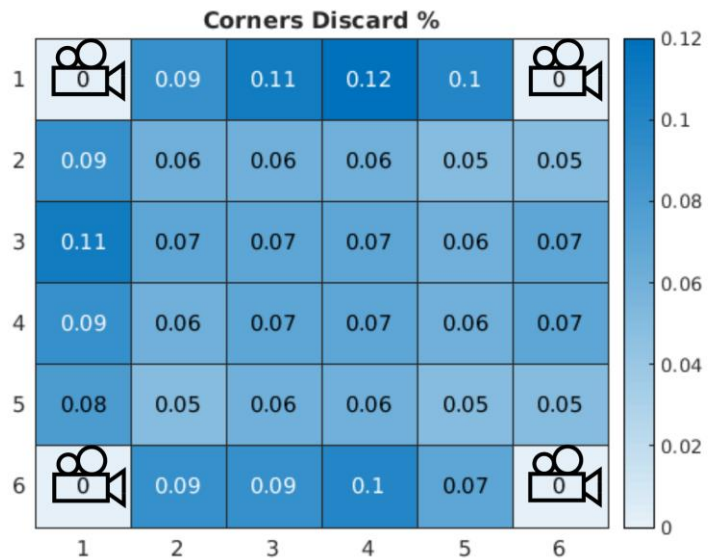


*Heatmap the ratio of discarded pixels in a 6x6 light fields at some position in the Sponza scene
(Mäkitalo 2023)*

Discarded pixels

- Depends on:


- Number of sources





Heatmap the ratio of discarded pixels in a 6x6 light fields at some position in the Sponza scene
(Mäkitalo 2023)

Discarded pixels

- Depends on:


- Number of sources 
- **Location of the sources**



Heatmap the ratio of discarded pixels in a 6x6 light fields in the Sponza scene
(Mäkitalo 2023)

Problem

- Photorealistic light field rendering:
 - Large number of pixels to process

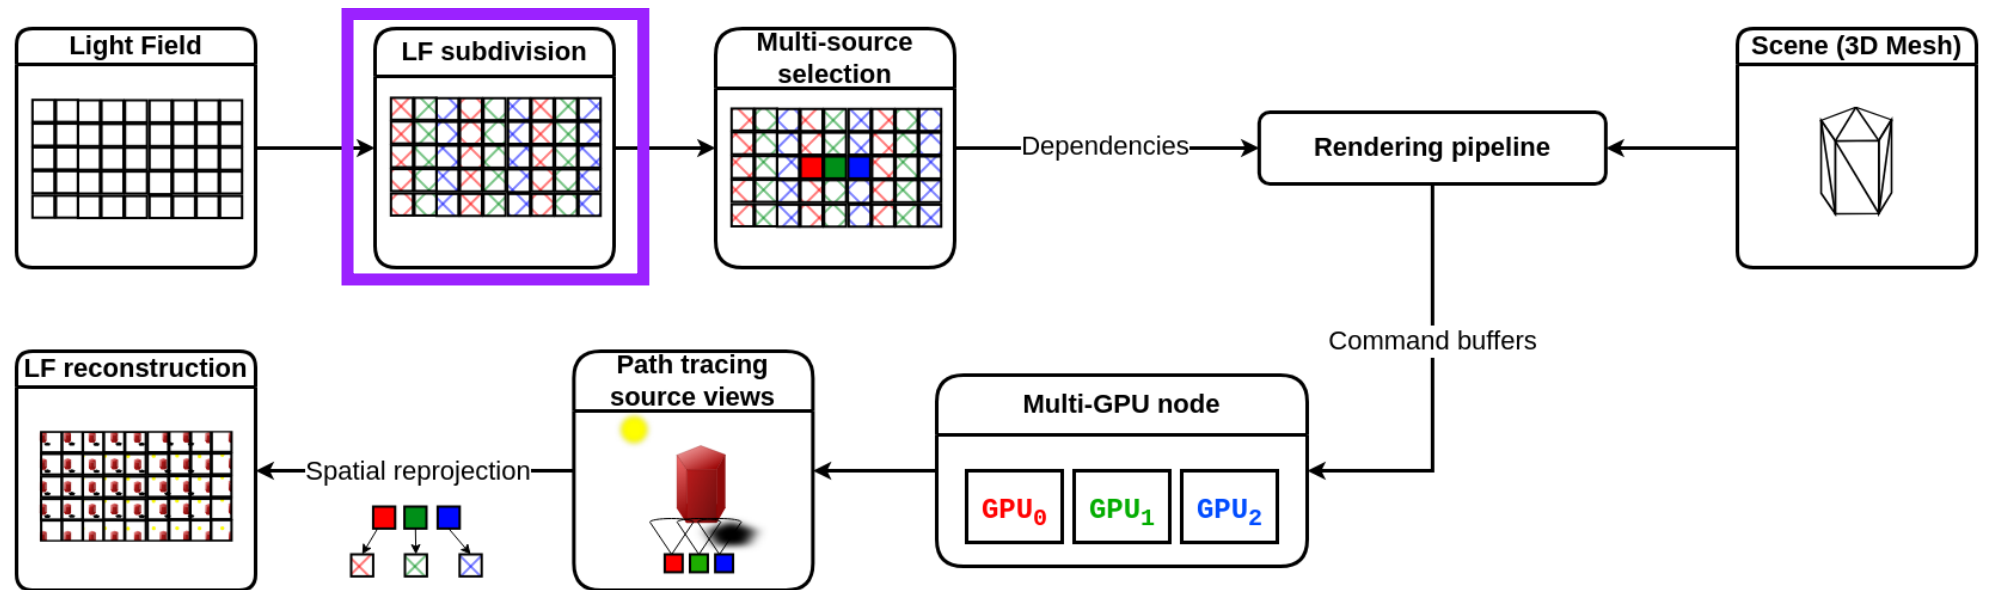


- Light field displays:
 - 50—90 Hz refresh rate = Number of light fields needed per second
- Discarded pixels must be rendered
- **Need for high computing power: Multiple GPUs**

Multi-source spatial reprojection pipeline

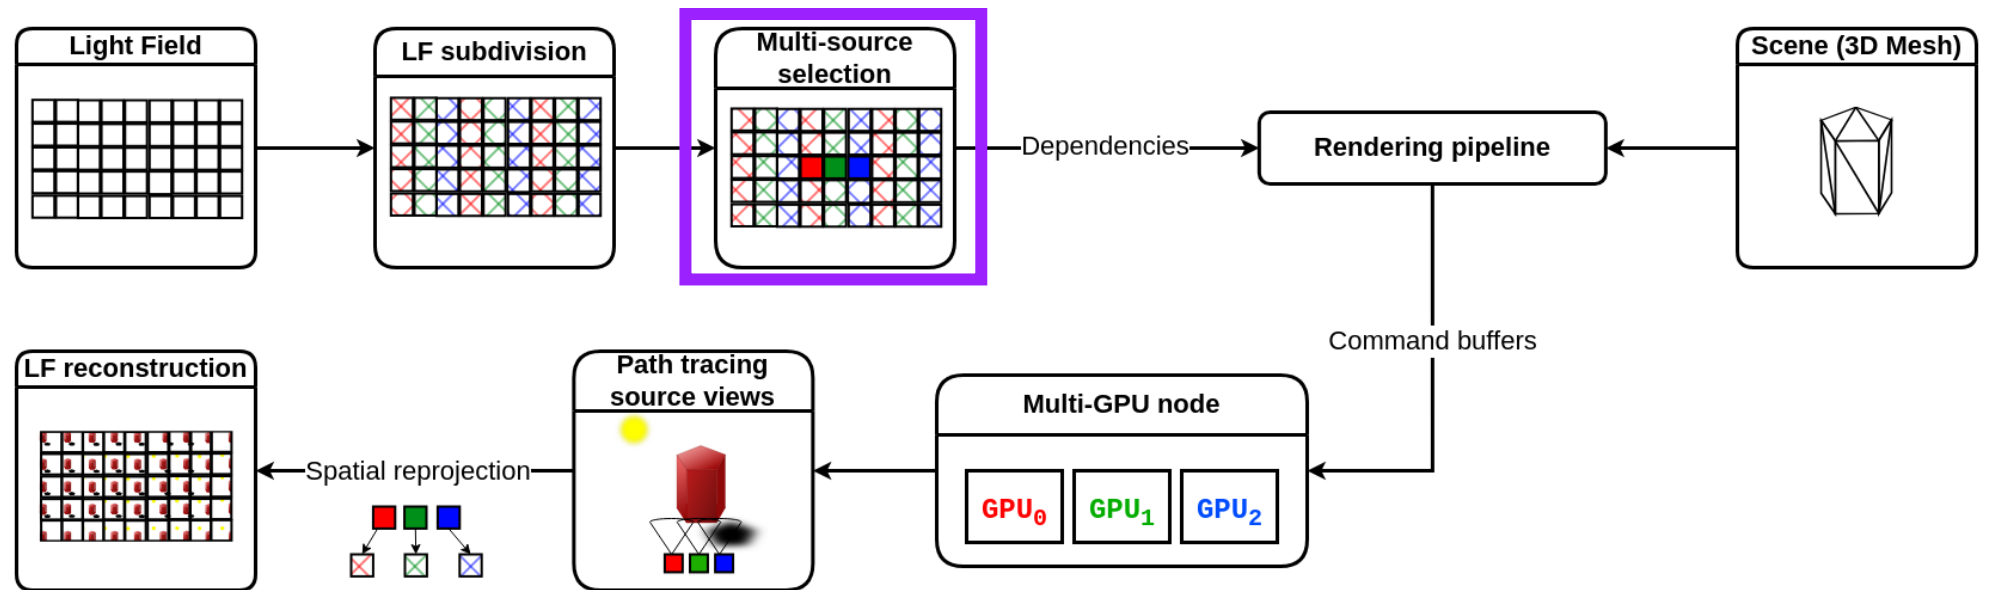
Overview of the proposed method

1. Form sets of views from a light field



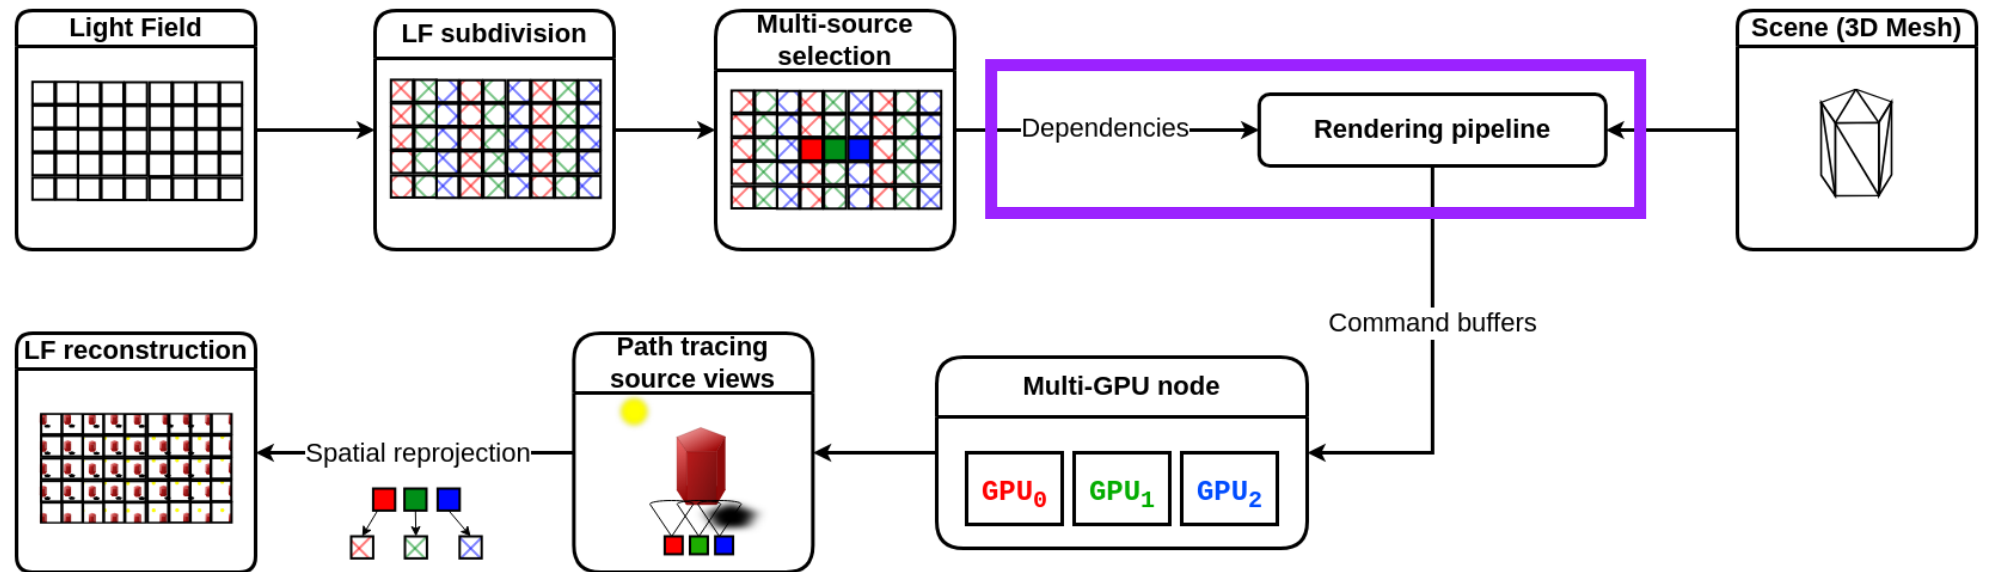
Overview of the proposed method

1. Form sets of views from a light field (1 set per GPU)
2. Find the central view in each set



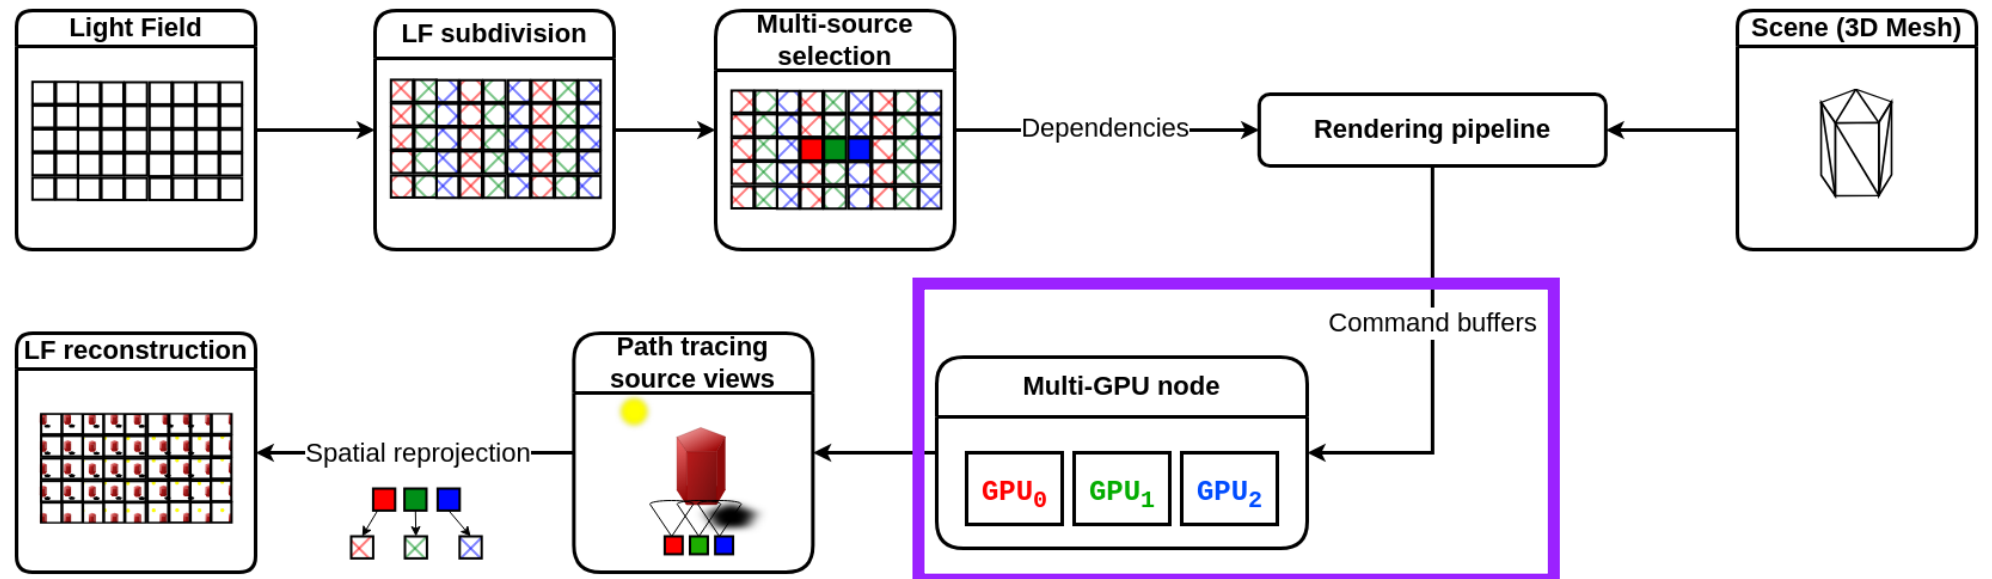
Overview of the proposed method

1. Form sets of views from a light field
2. Find a medoid (central view) in each set
3. **Send dependencies to the pipeline**



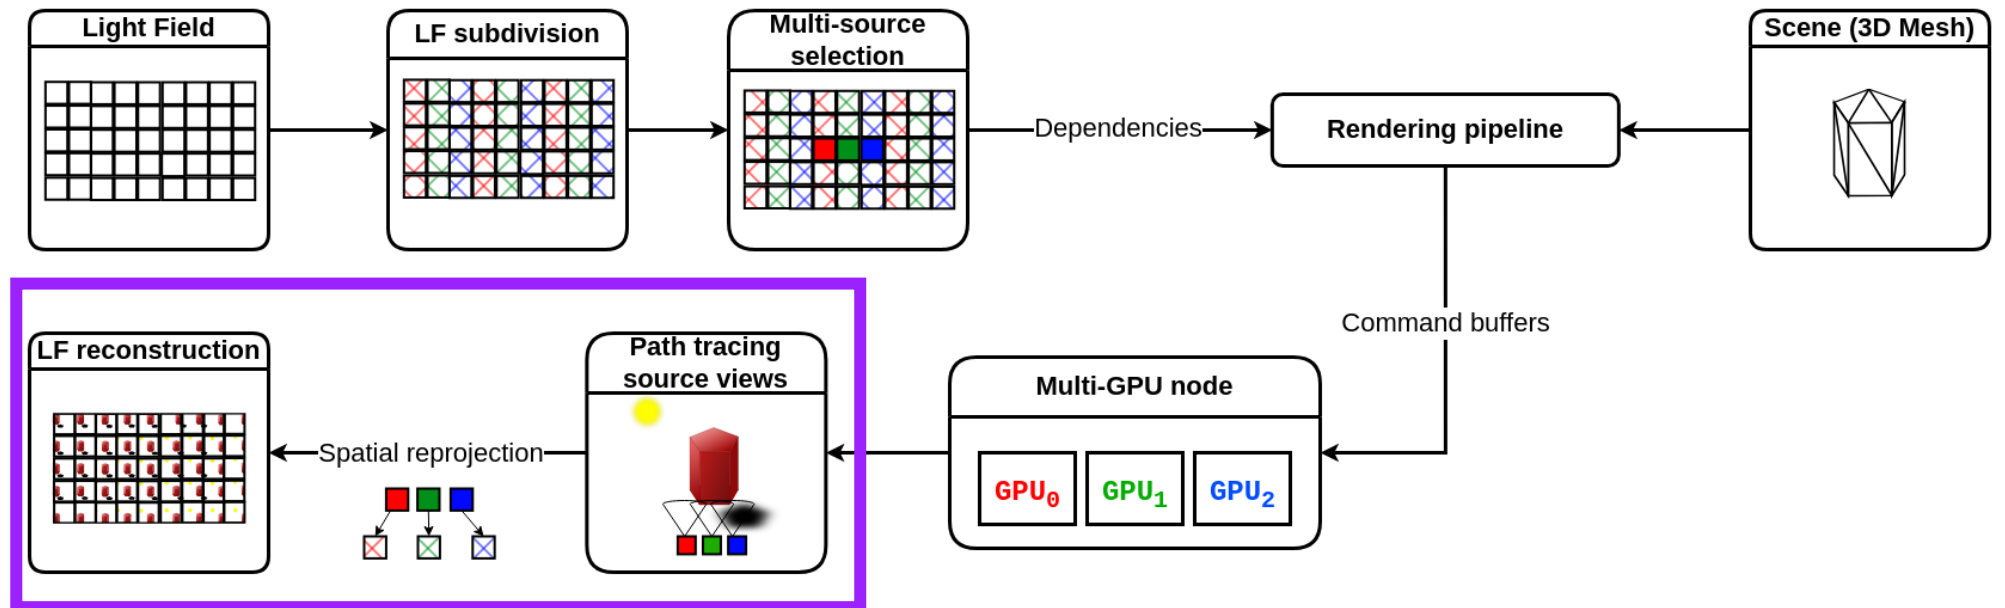
Overview of the proposed method

1. Form sets of views from a light field
2. Find a medoid (central view) in each set
3. Send dependencies to the pipeline
4. **Create and schedule work to GPUs**




Overview of the proposed method

1. Form sets of views from a light field
2. Find a medoid (central view) in each set
3. Send dependencies to the pipeline
4. Create and schedule work to GPUs
5. **Render**



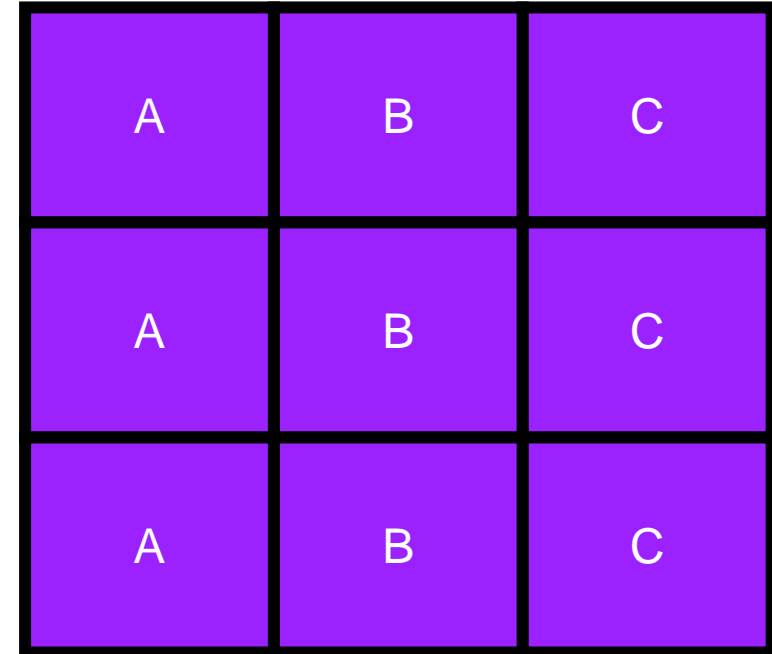

Less dependencies, more efficiency

- Workload distribution in **prior works**
- Example: a 3x3 light field and 3 GPUS (*A, B and C*)




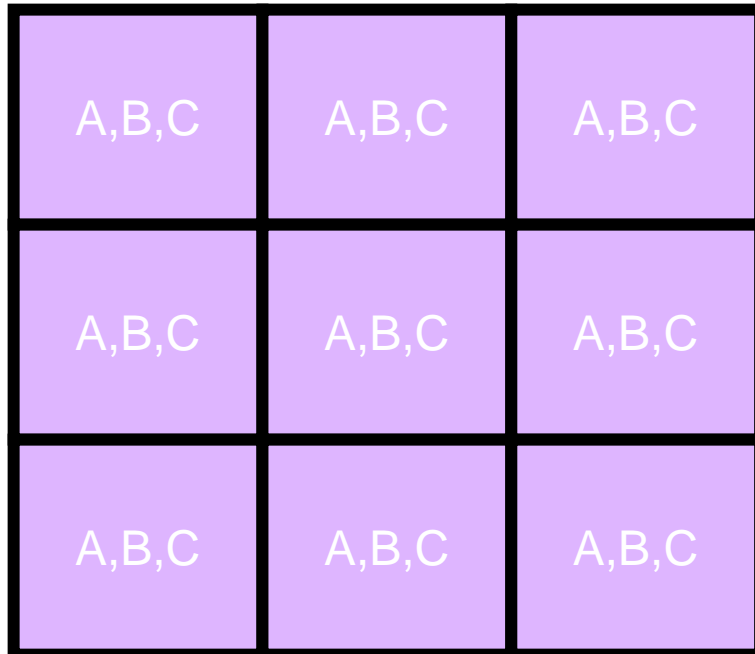
Less dependencies, more efficiency

- Workload distribution in **our approach**
- Example: a 3x3 light field and 3 GPUS (A, B and C)



Less dependencies, more efficiency

- The **central view** in each set is selected as the source
- 1 source per set

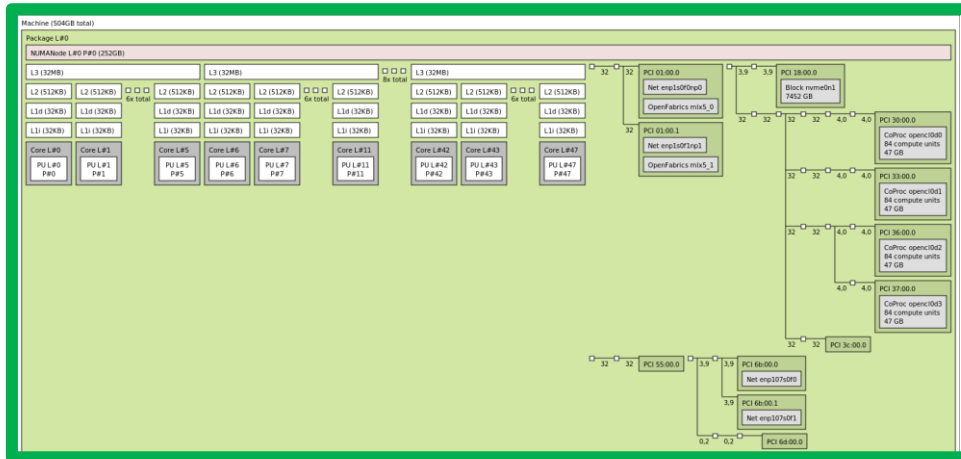


Experimental setup

Experimental setup

- **Single socket multi-GPU node**
 - **4 RTX A6000 GPUs**

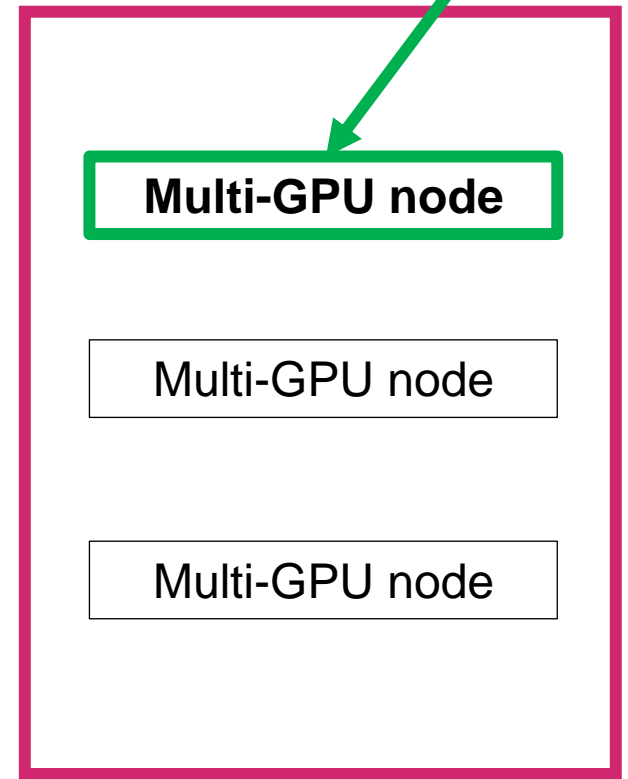
(Topological map of the node)



Rack

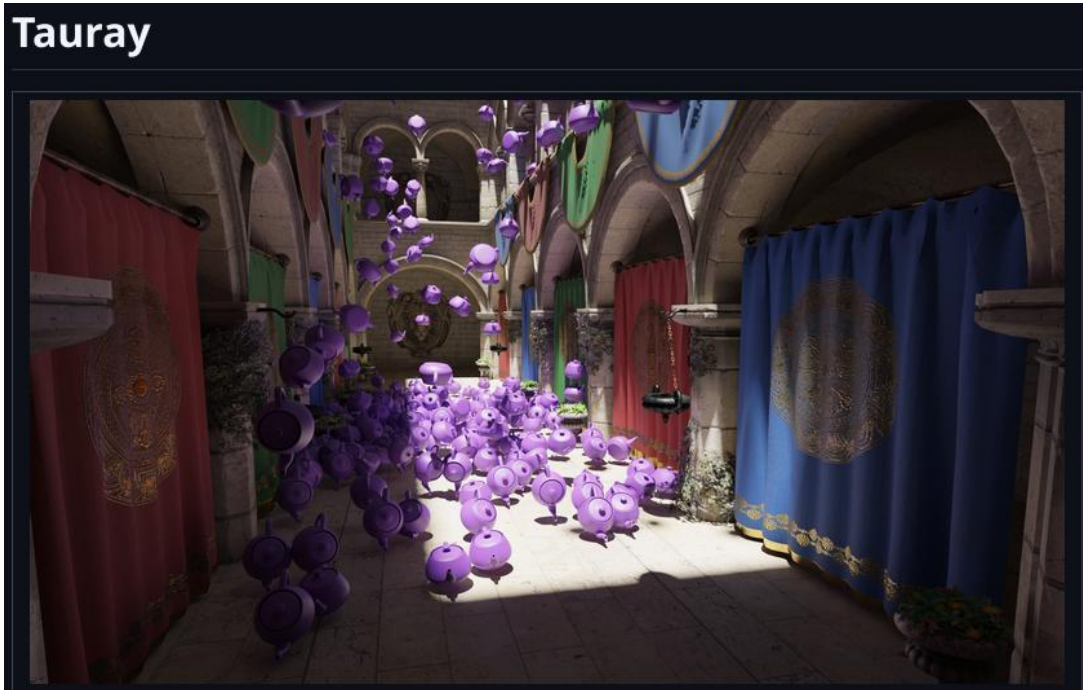


Rack



x4 RTX A6000 GPUs

Experimental setup



- Tauray: **open-source multi-GPU stereo and light field path tracer**
 - C++ / Vulkan
 - <https://github.com/vga-group/tauray>

Experimental setup




4 scenes:


- 1) Bistro Exterior (reference image)
- 2) Sponza (reference image)
- 3) San Miguel (Left: Numerical near-eye light field display model, Right: Looking Glass display 32")
- 4) Eternal Valley FPS (Left: Missing pixels -pink- in our pipeline, Right: in prior work –green, red, blue and pink)

Experimental setup

- Two light fields configuration
 1. Multiscopic 100x1 (1 dimension – horizontal): **100 views**
 2. NELF 12x6 (2 dimensions): **72 views (36 views per eye)**



1

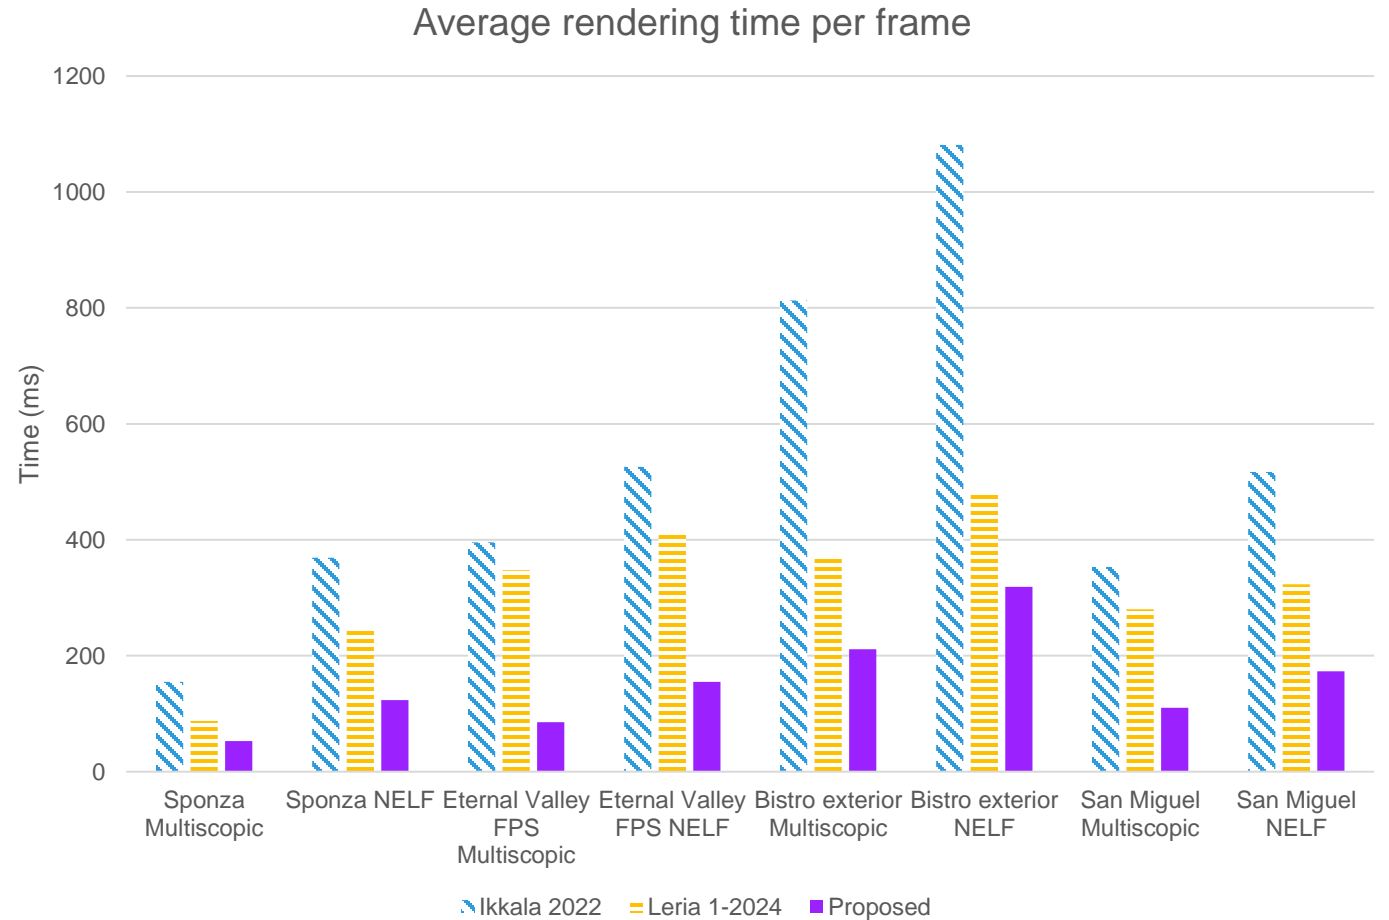


2

Results

Execution time performance per scene


- Speedup factor:
 - **x1.51** up to **x4.63**
- **Outperforms prior pipelines**
- Light field rendering frequency: **3 up to 19Hz**



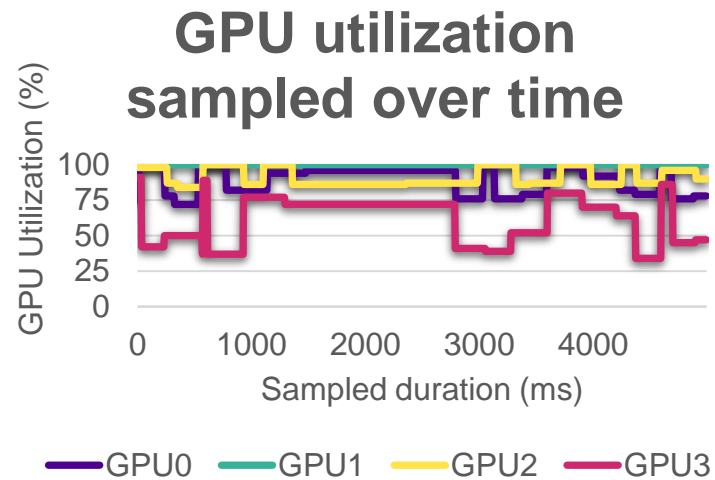
(Lower is better)

GPU utilization


- **Proposed pipeline:** High GPU utilization (99-100%)



(Ikkala 2022)





(Leria 1-2024)



Proposed

Scaling

- Proposed pipeline: 1 source view per GPU
 - **Near perfect linear** scaling when increasing the number of GPUs



Conclusions

- Importance of dependency to scalability tradeoff
- Automatic selection of source views for multi-GPU platforms
- Interactive framerates (less than 30Hz) but not real-time (50 – 90 Hz)

Interactive Multi-GPU Light Field Path Tracing Using Multi-Source Spatial Reprojection

Thanks for listening!

Acknowledgements

- U. Gudelek, K. Akbar, J. Käpylä, J. Solanti, U. Akpınar
- Centre for Immersive Visual Technologies (Tampere University)
- Anonymous reviewers

Authors

Virtual reality and Graphics Architectures (tuni.fi/vga/),
Tampere University (Finland)
Email: first.last@tuni.fi



Erwan Leria
(Presenter)



Markku Mäkitalo



Pekka Jääskeläinen

Realistic 3D,
Mid Sweden University (Sweden)
Email: first.last@miun.se



Mårten Sjöström



Tingting Zhang



This work was supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 956770 (Plenoptima)