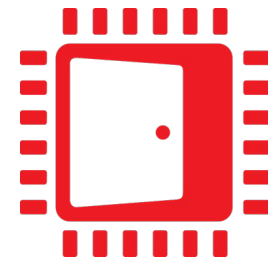




together we advance_

MESH SHADERS AND WORK GRAPHS - PERFECT PAIR

DR. ŁUKASZ IZDEBSKI
DEVELOPER TECHNOLOGY ENGINEER



AMD 
GPUOpen

INTRODUCTION

- Mesh Shaders are getting more popular at Digital Dragons 2024
 - GPU-driven Rendering with Mesh Shaders in Alan Wake 2 by Erik Jansson (Remedy Entertainment) – Monday 20.03.2024 at 3:00 PM
 - Mesh Shaders – The Future of Rendering - Radosław Paszkowski (PixelAnt Games) – Monday 21.03.2024 1:00 PM

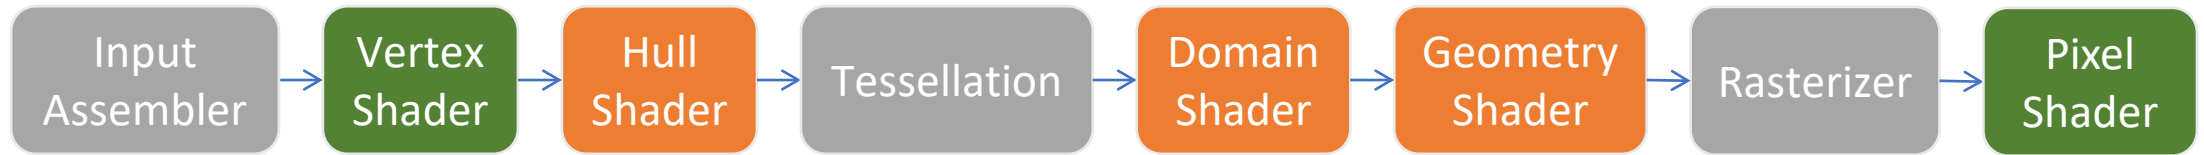
AGENDA

- Quick Mesh Shader technology introduction.
- Mesh Shader Example 1
- Mesh Shader Example 2
- Mesh Shader Example 3
- Graphics API new feature
- Mesh Shader Example 4
- Mesh Shader Example 5

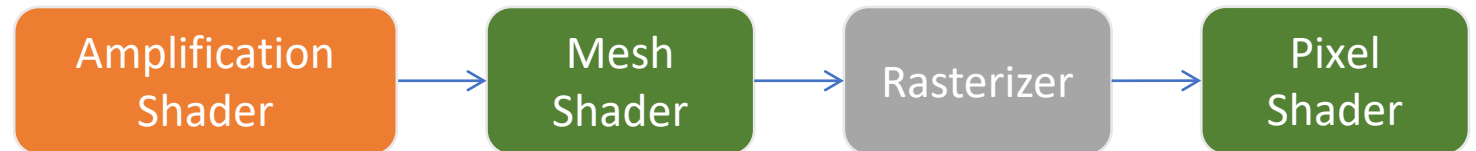
SPOILER ALERT !!!
NO SHADER CODE INCLUDED

MESH SHADERS TECHNOLOGY 101

- Traditional Graphics Pipeline

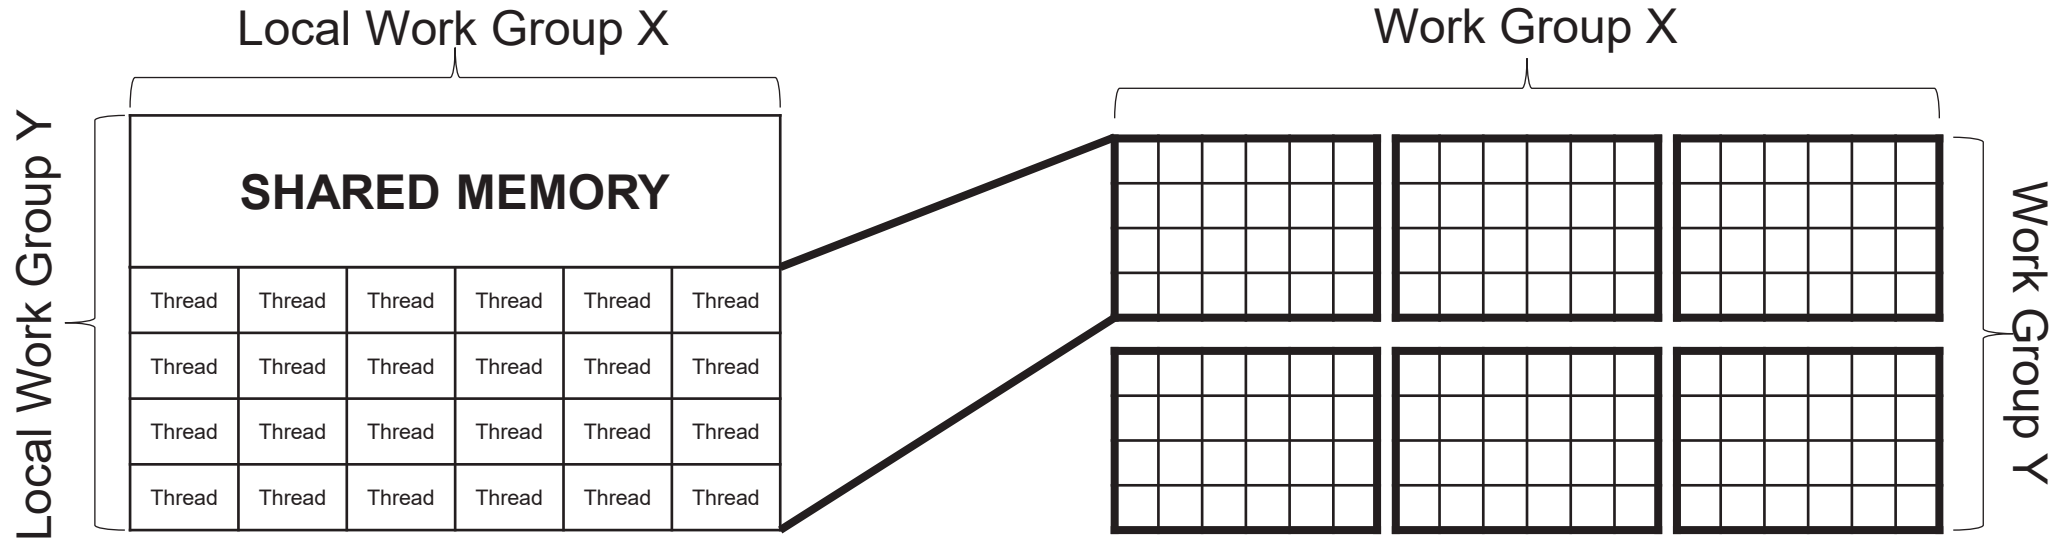


- Mesh Shader Pipeline

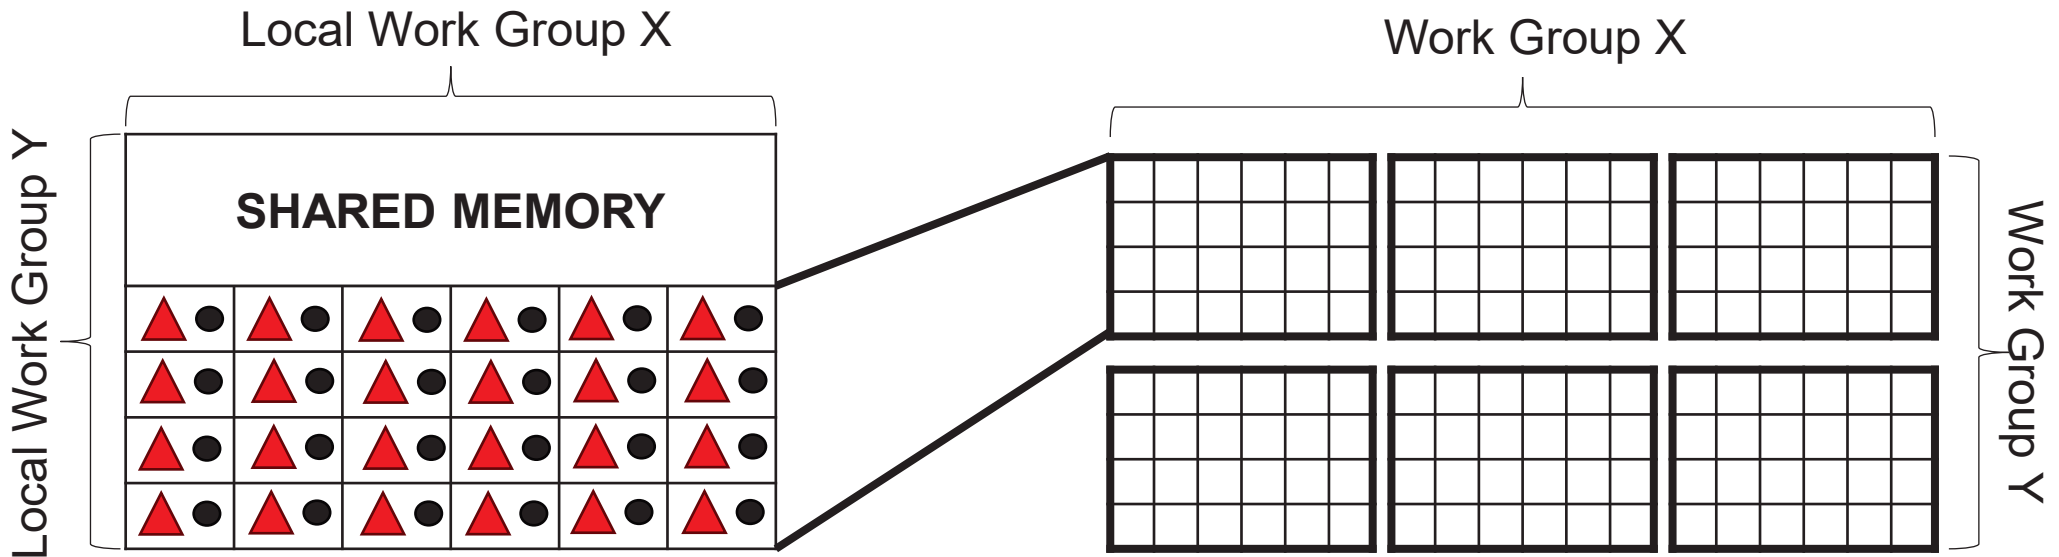


MESH SHADER VS COMPUTE SHADER

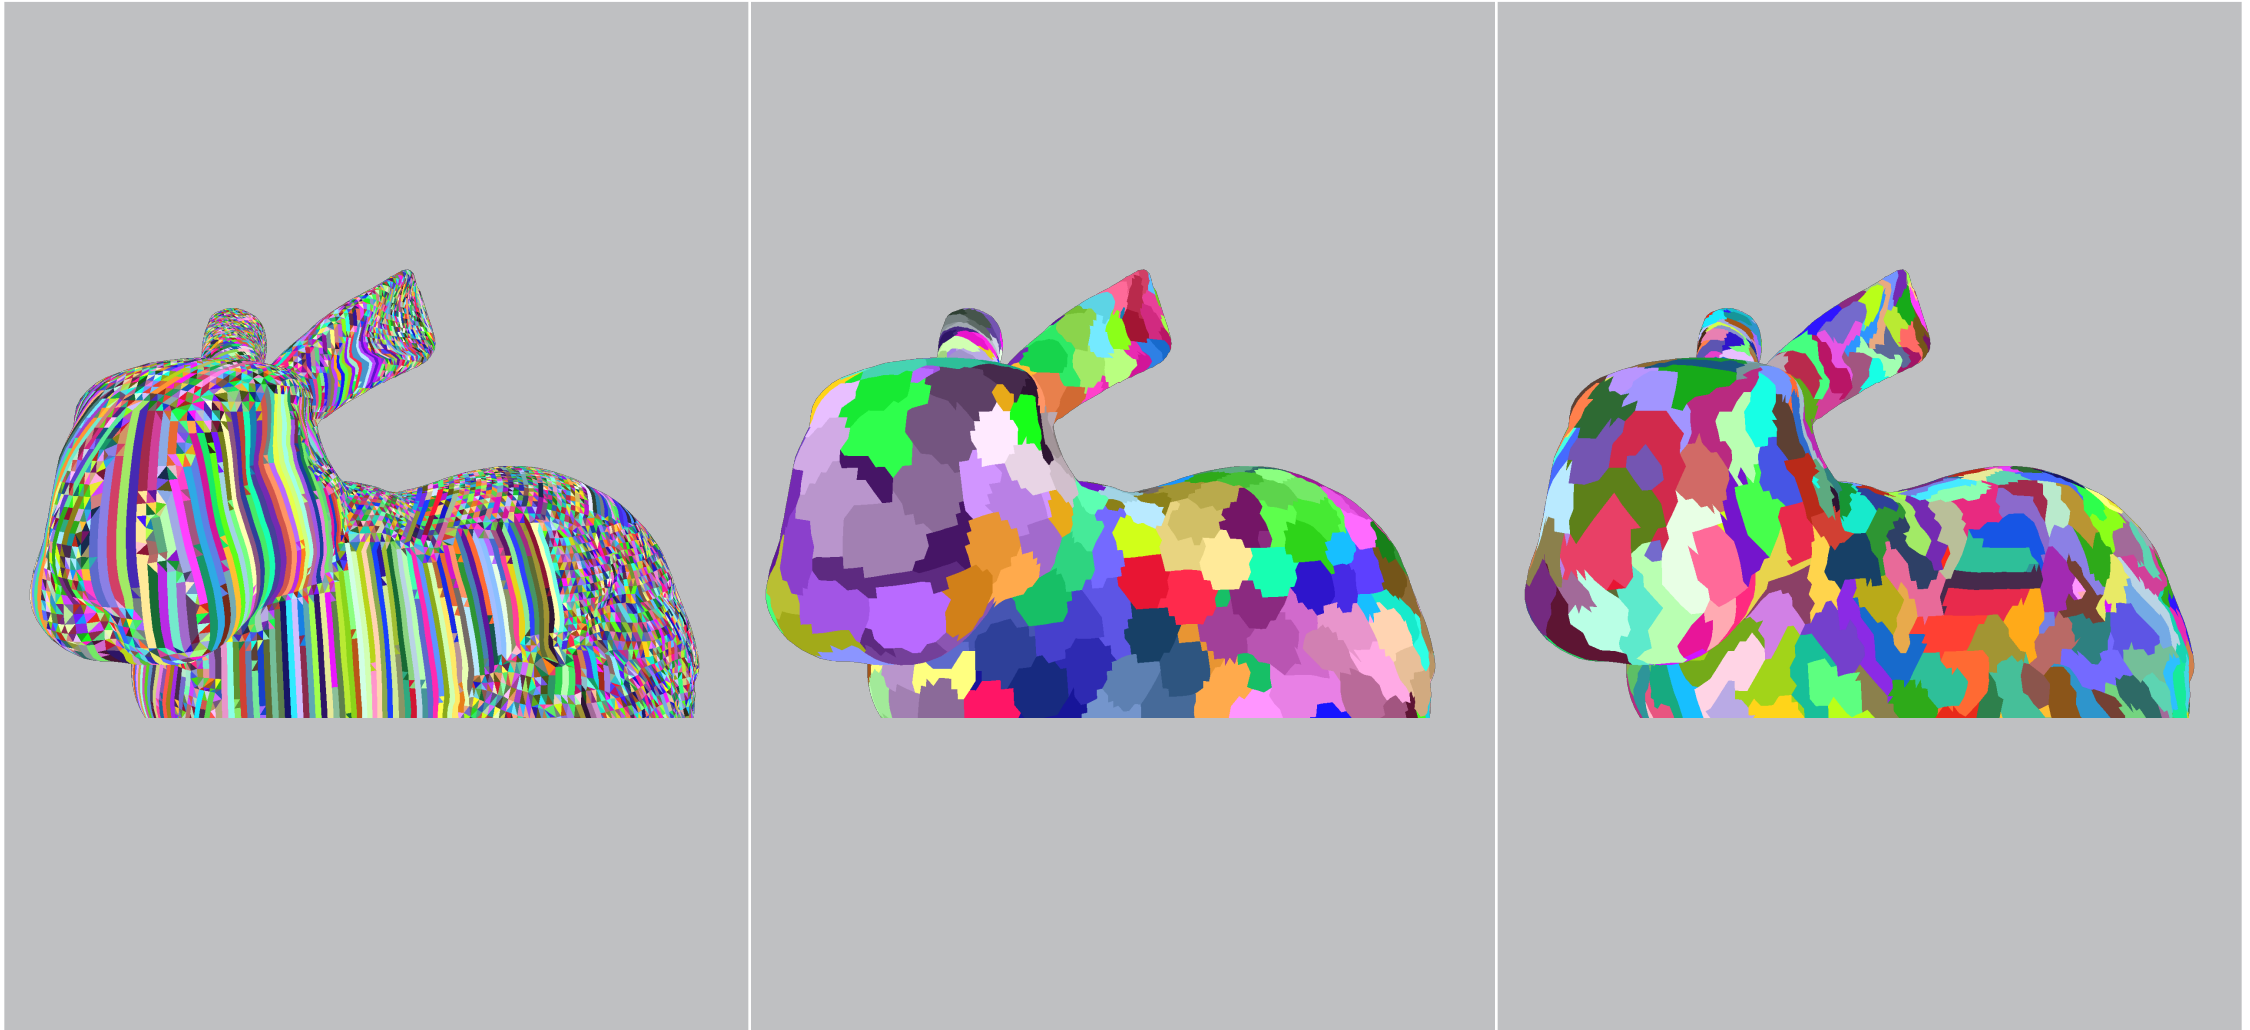
COMPUTE SHADER



MESH SHADER



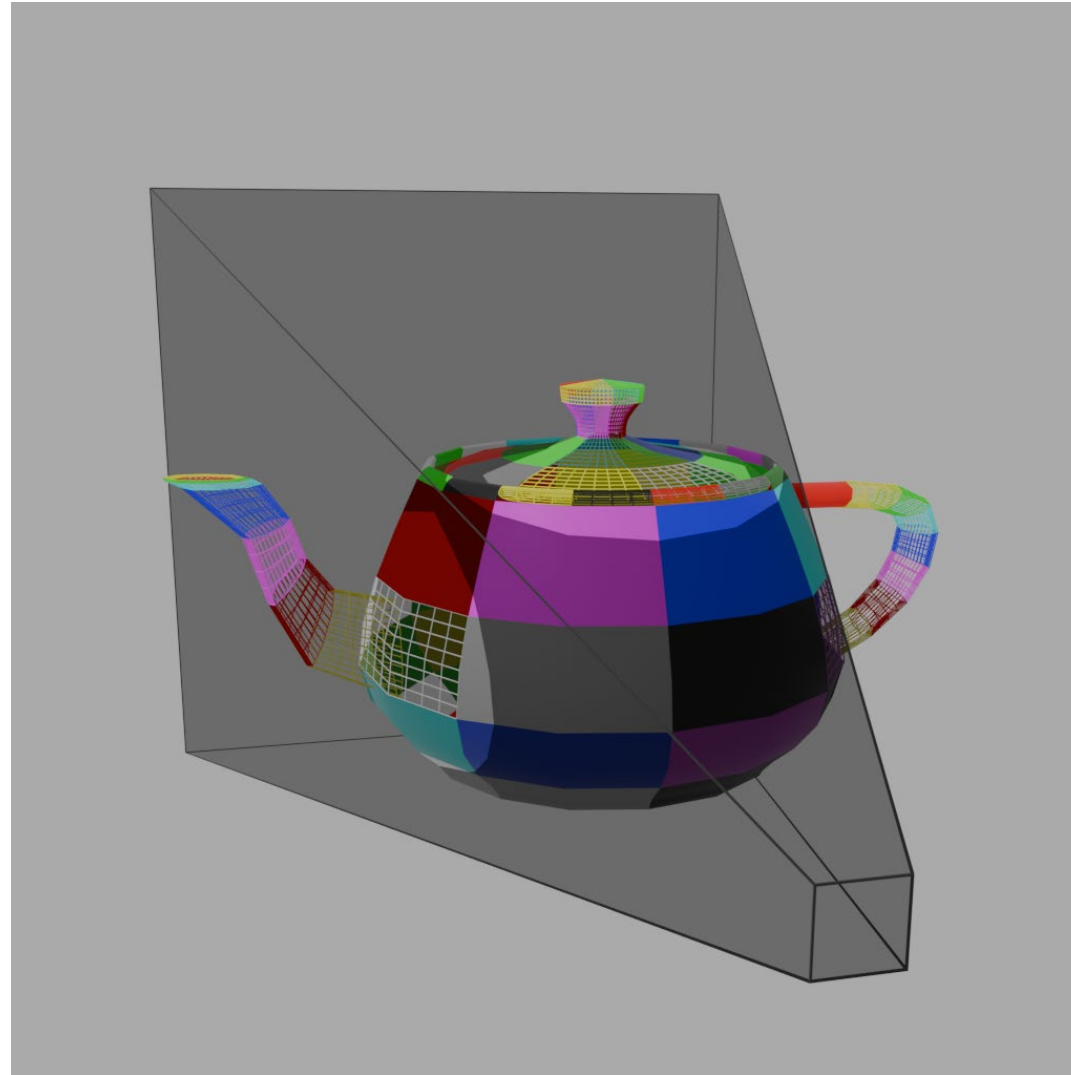
MESHLETS ZOO



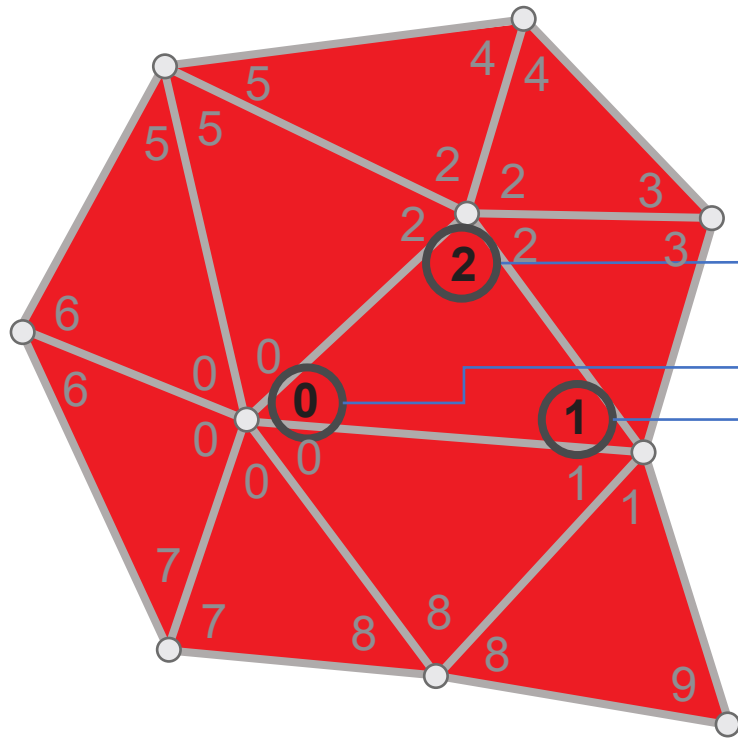
MESHLET GENERATION LIBRARIES

- More information about the generation of Meshlets:
 - https://gpuopen.com/learn/mesh_shaders/mesh_shaders-optimization_and_best_practices
 - <https://meshoptimizer.org>
 - <https://github.com/microsoft/DirectXMesh>

MESH SHADERS CULLING AND LODS



MESH SHADER COMPRESSION

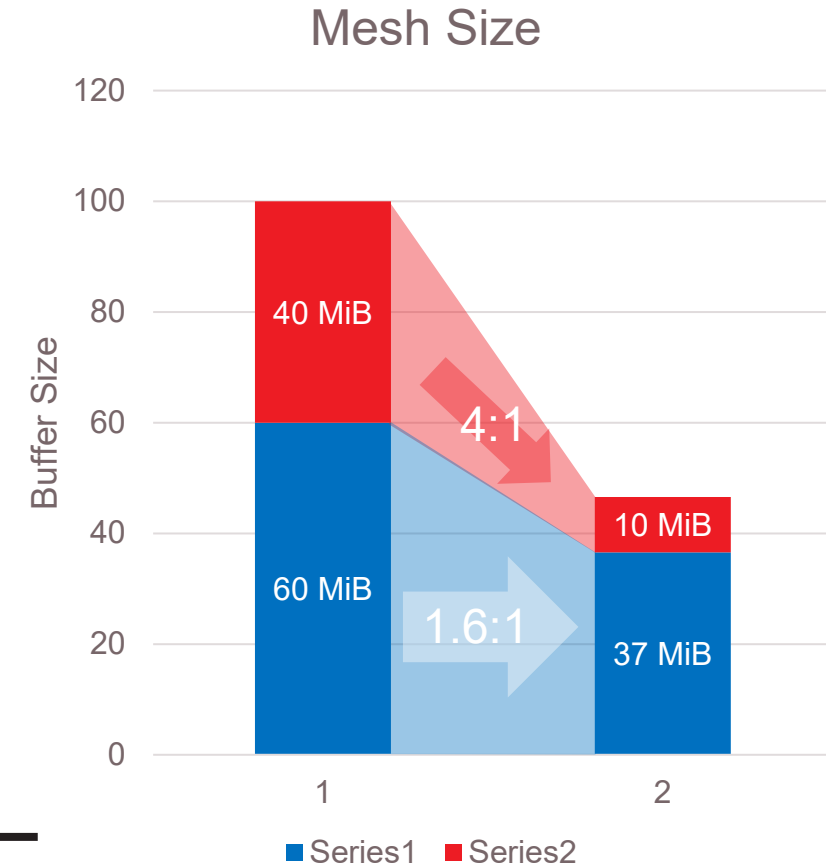


8 Bit Index

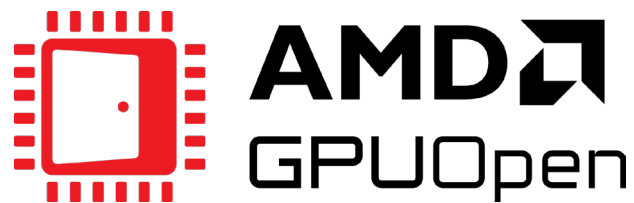
8 Bit Index

8 Bit Index

24 Bits per Triangle



MESH SHADER COMPRESSION



Keep an eye on
GPUOpen.com



Follow us on X

<https://twitter.com/GPUOpen>



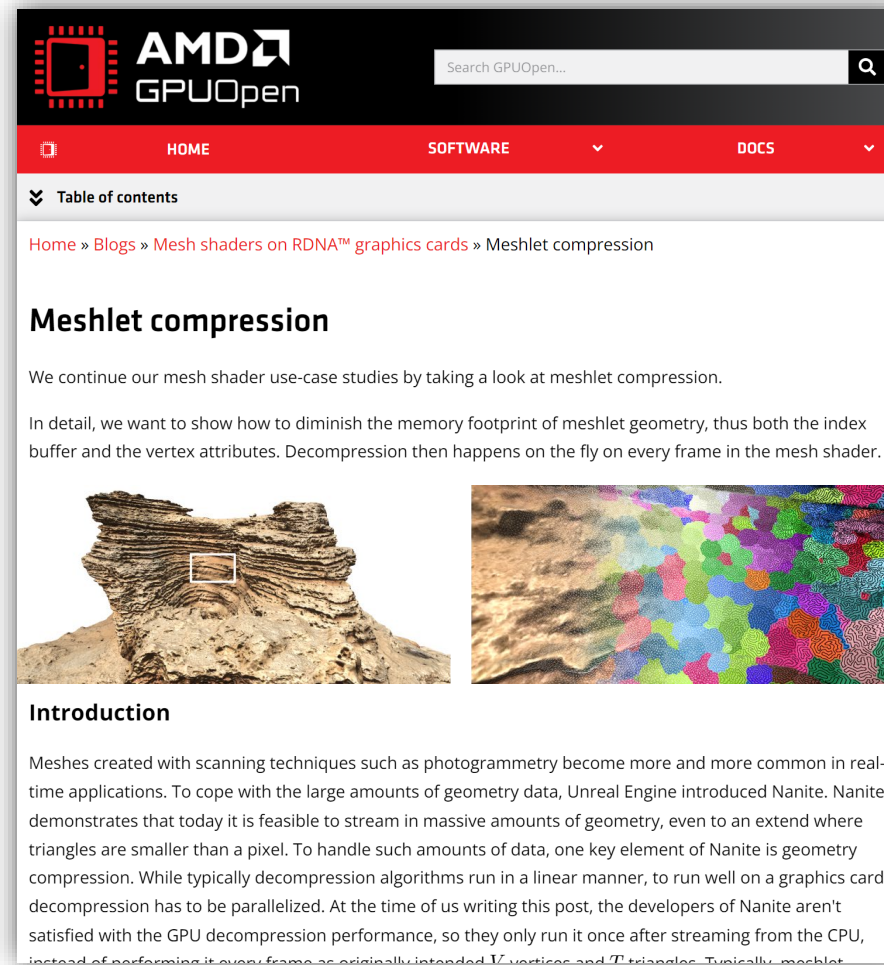
Follow us on Mastodon

<https://mastodon.gamedev.place/@gpuopen>

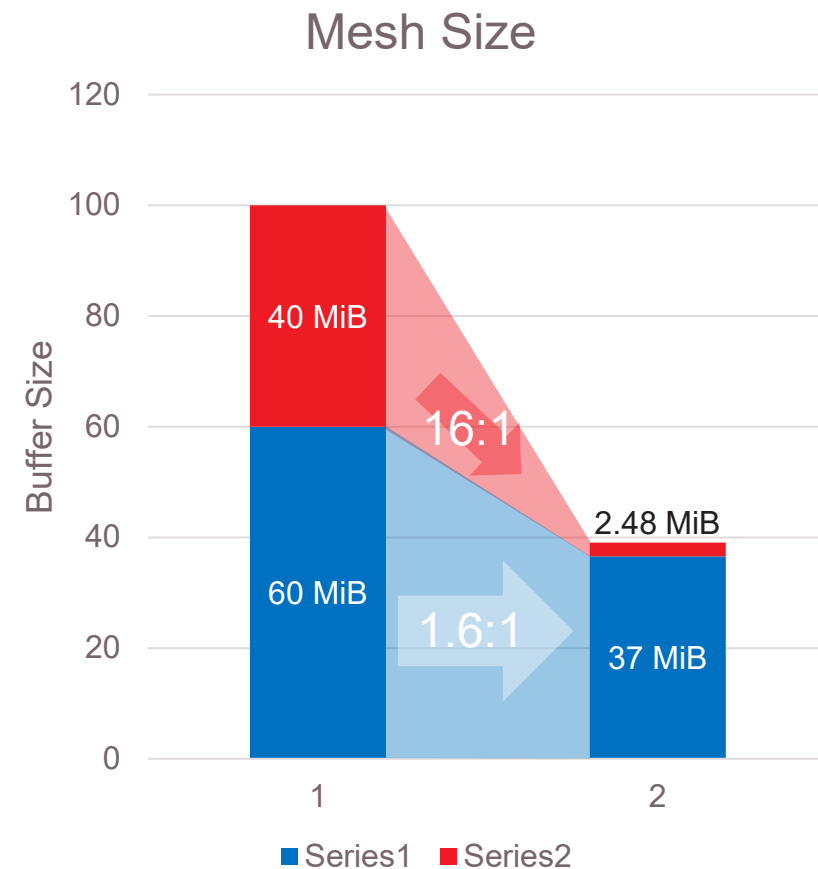
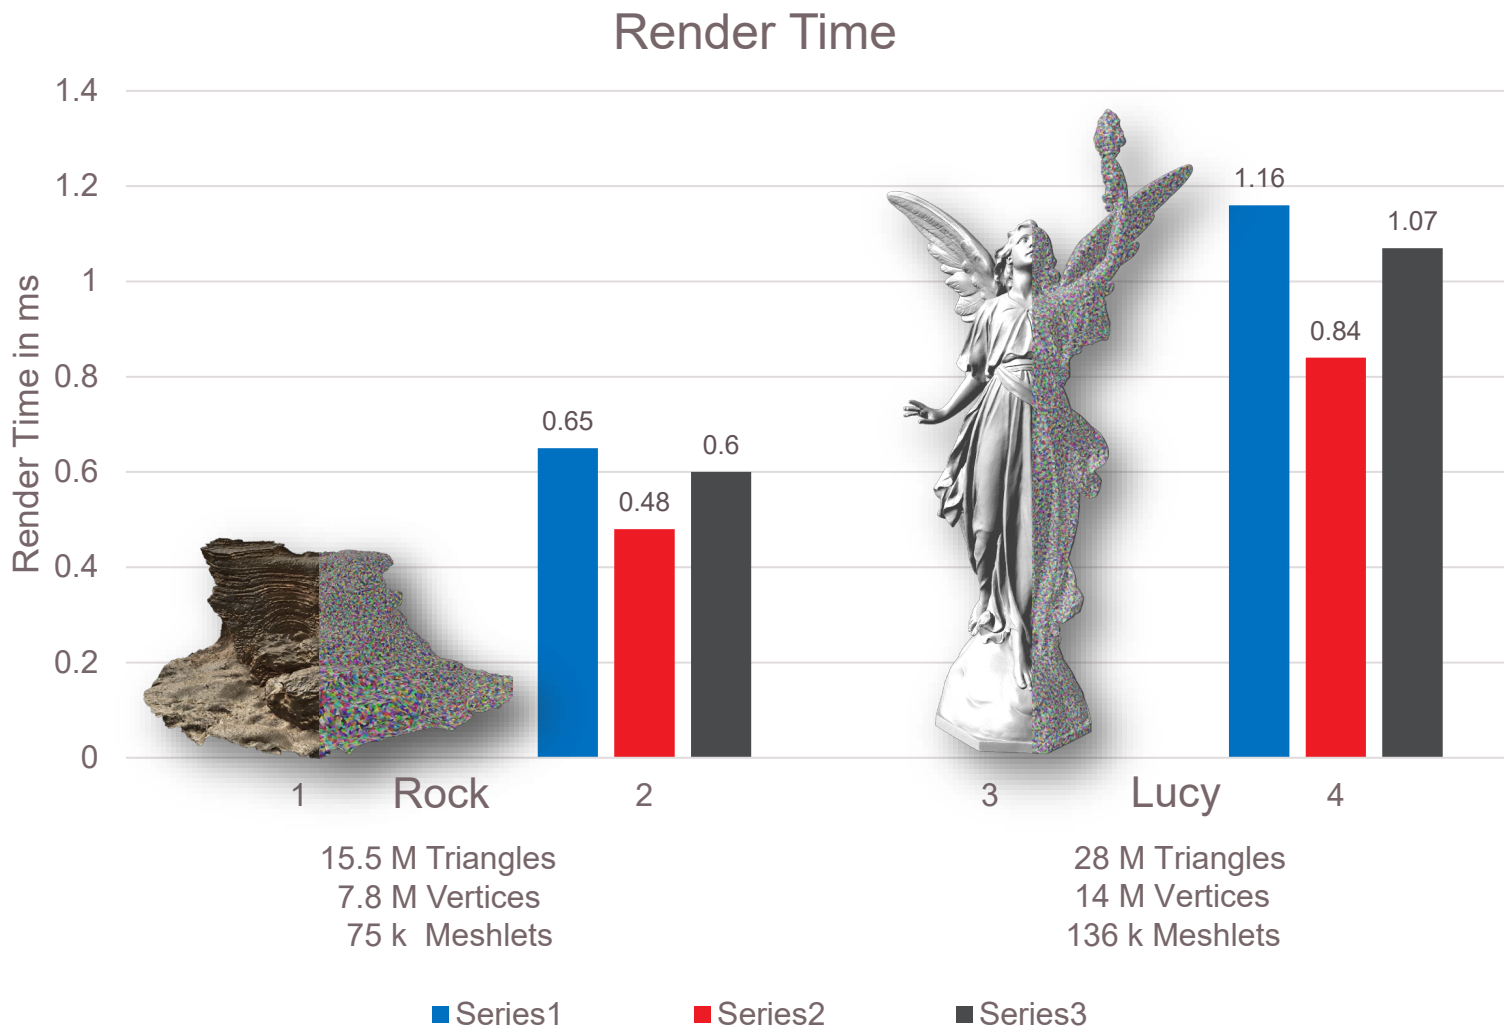


Follow us on Zhihu

<https://www.zhihu.com/org/gpuopen-7>



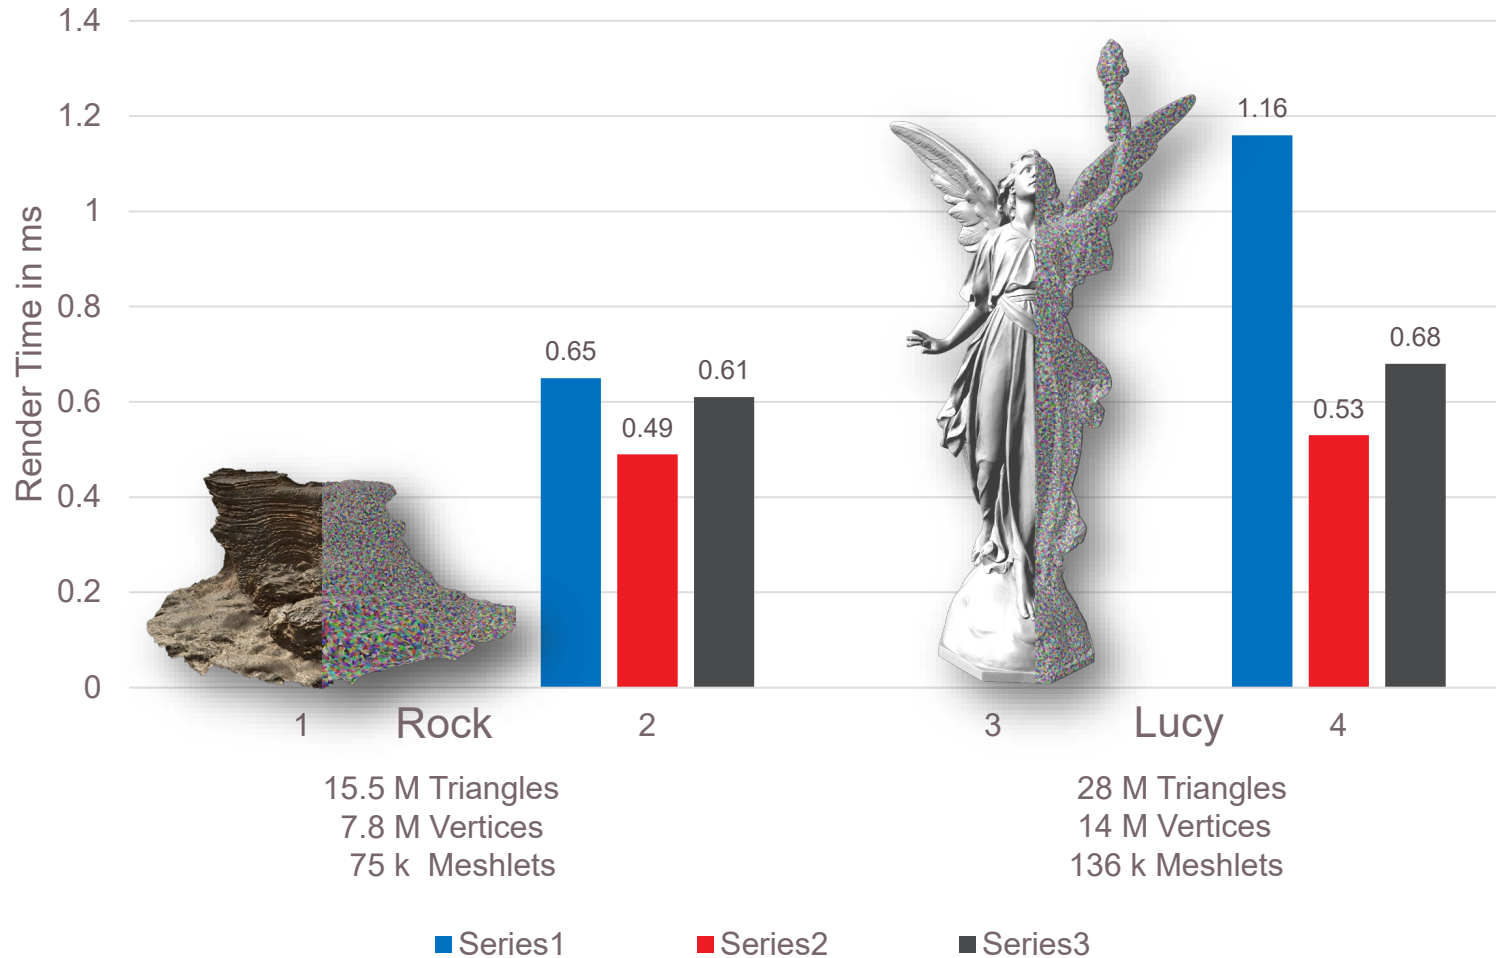
MESH SHADER COMPRESSION PERFORMANCE



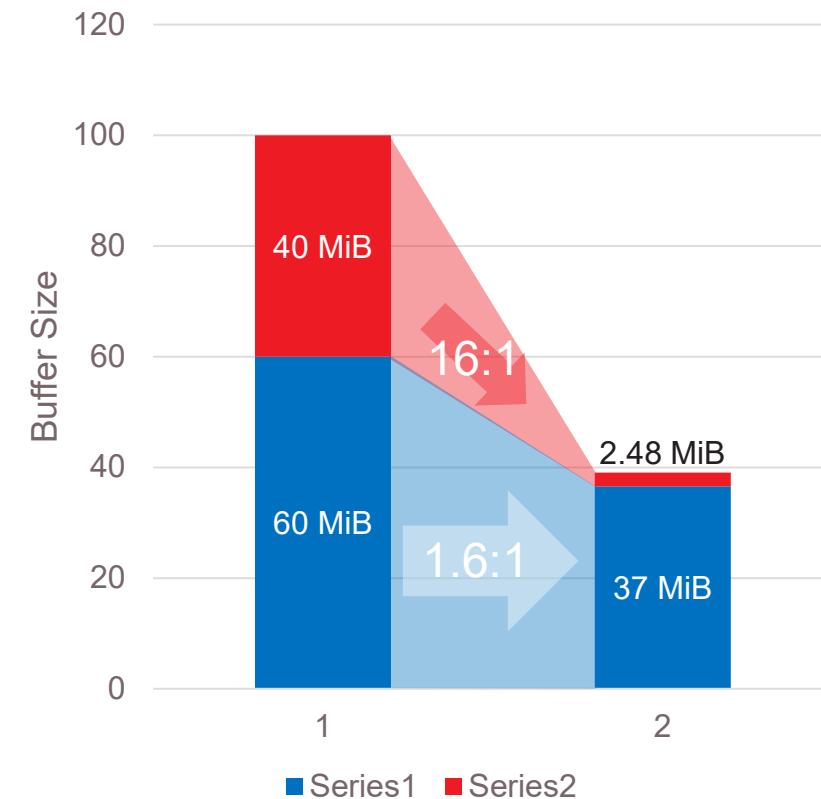
GPU: AMD Radeon™ RX 7900 XTX; driver version: 24.1.1; Meshes: Layered Rock by Aixterior, Lucy by Stanford Computer Graphics Laboratory

MESH SHADER COMPRESSION PERFORMANCE

Render Time with Culling



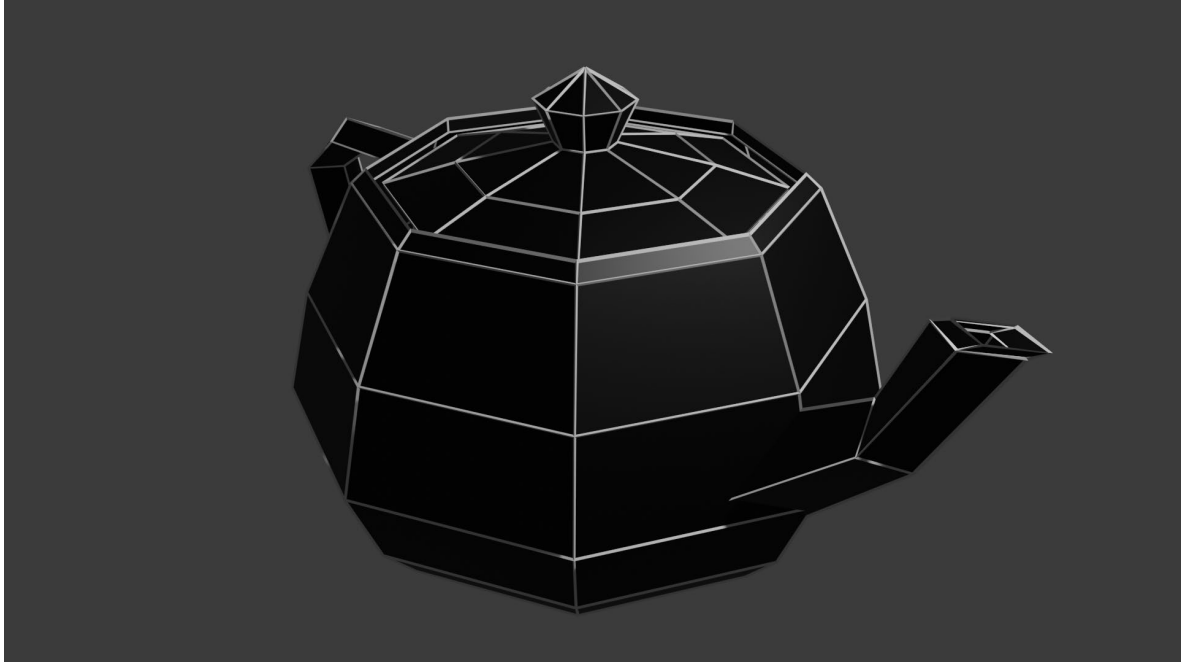
Mesh Size



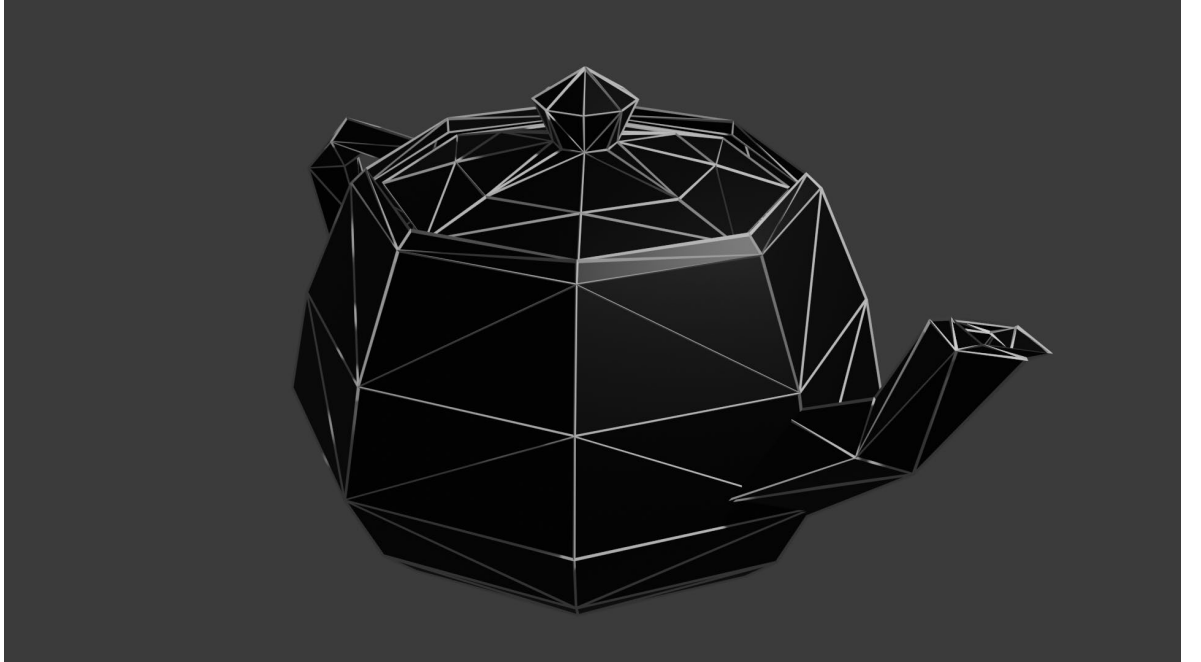
GPU: AMD Radeon™ RX 7900 XTX; driver version: 24.1.1; Meshes: Layered Rock by Aixterior, Lucy by Stanford Computer Graphics Laboratory

QUAD BASED AND TRIANGLE BASED MESH

QUAD BASED MESH



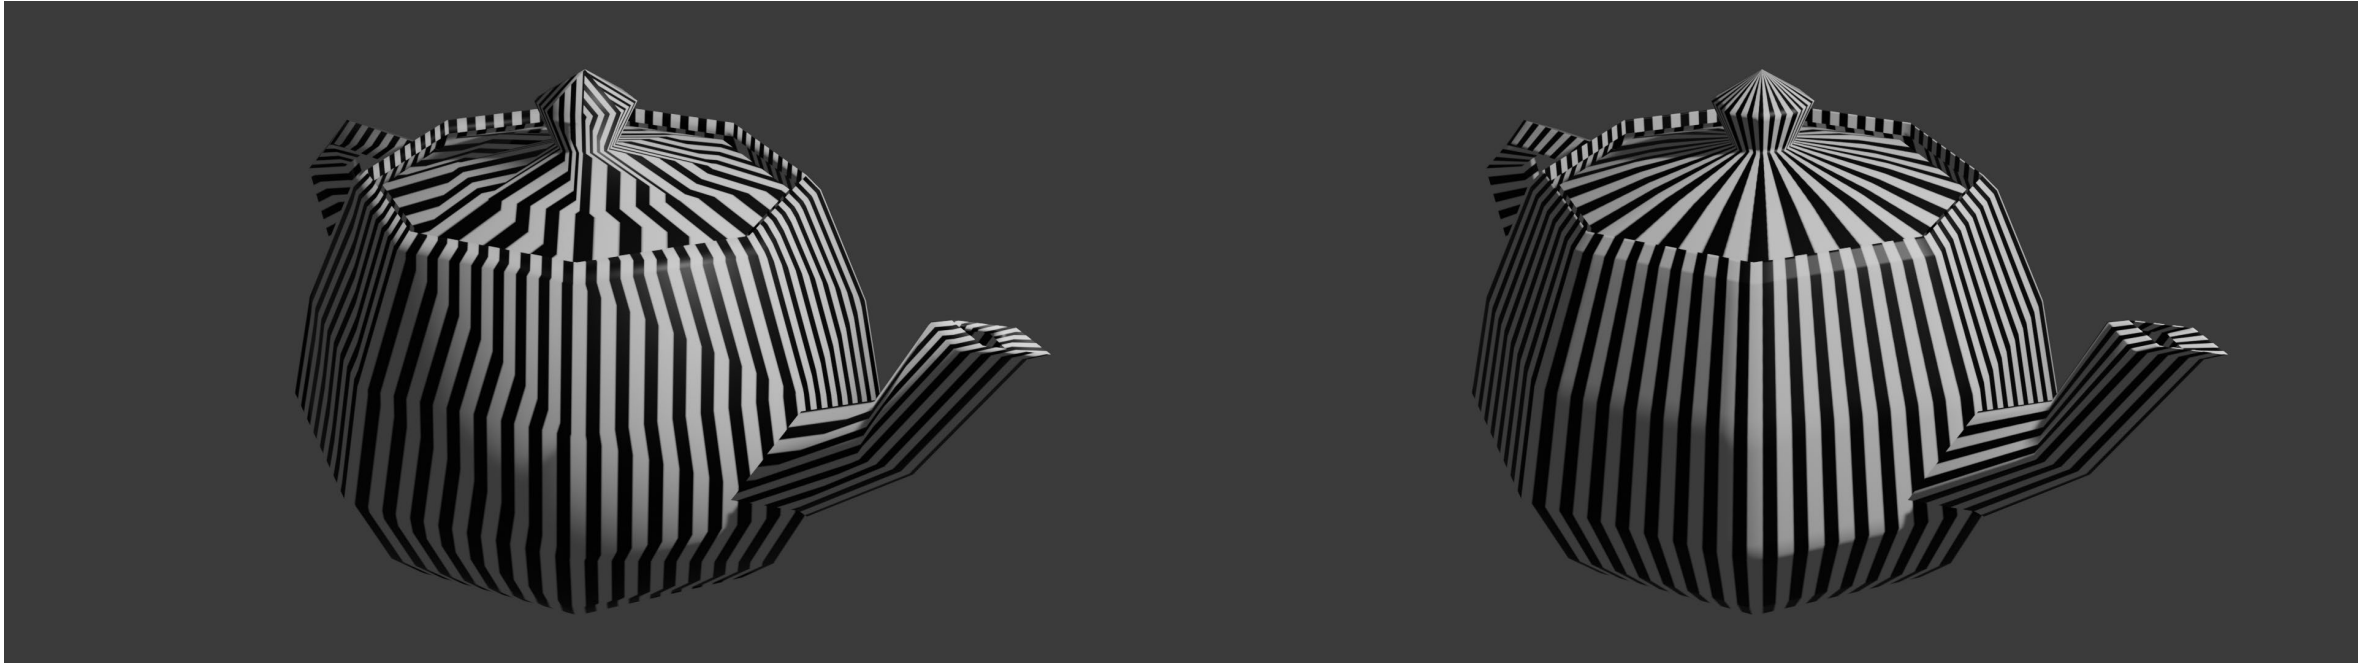
TRIANGLES BASED MESH



VERTEX ATTRIBUTES INTERPOLATION

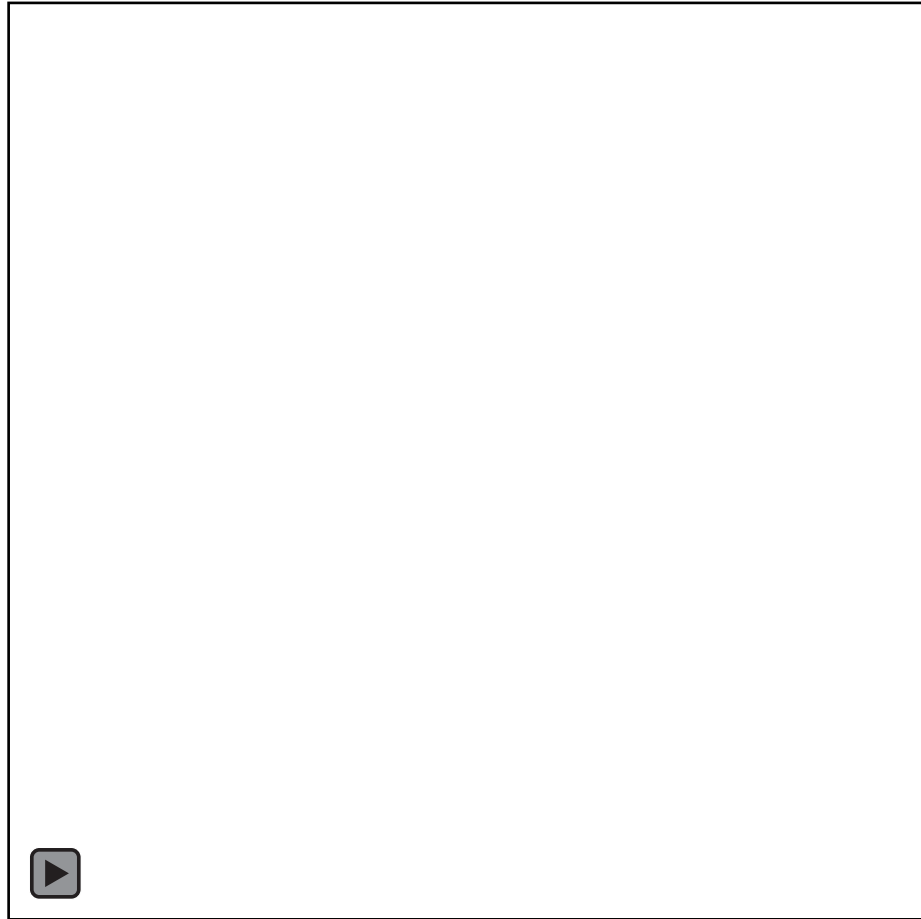
BARYCENTRIC COORDINATES

BILINEAR INTERPOLATION

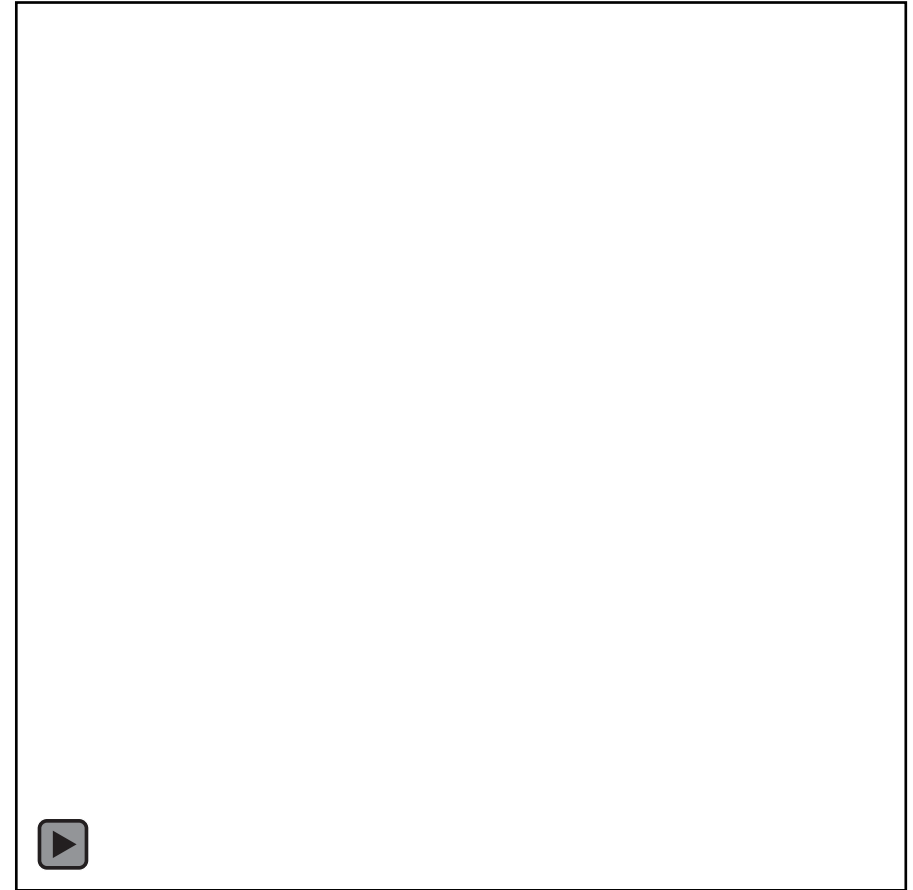


INTERPOLATION AND ANIMATION

BARYCENTRIC INTERPOLATION



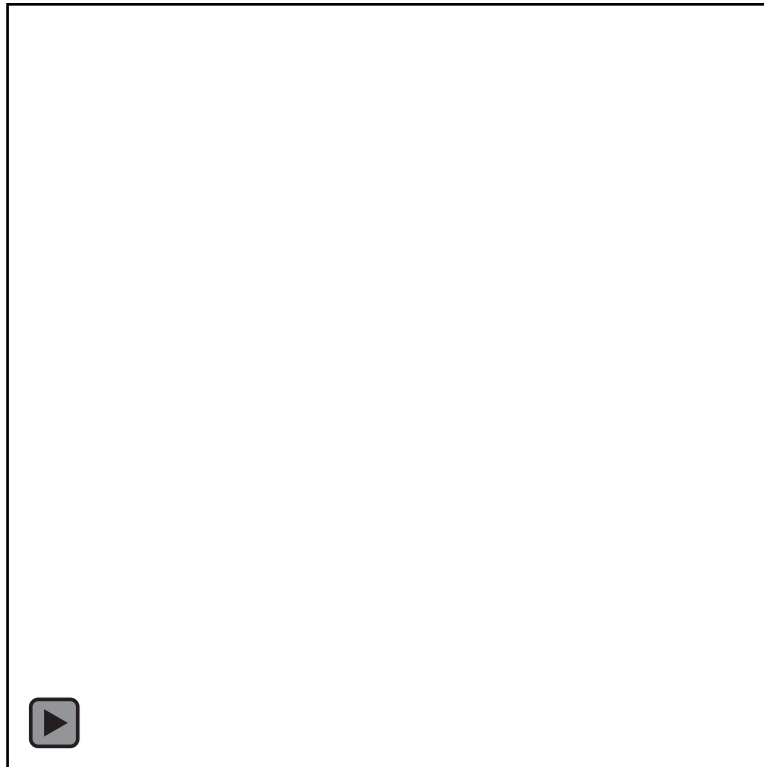
BILINEAR INTERPOLATION



QUADRILATERAL PRIMITIVE RASTERIZATION

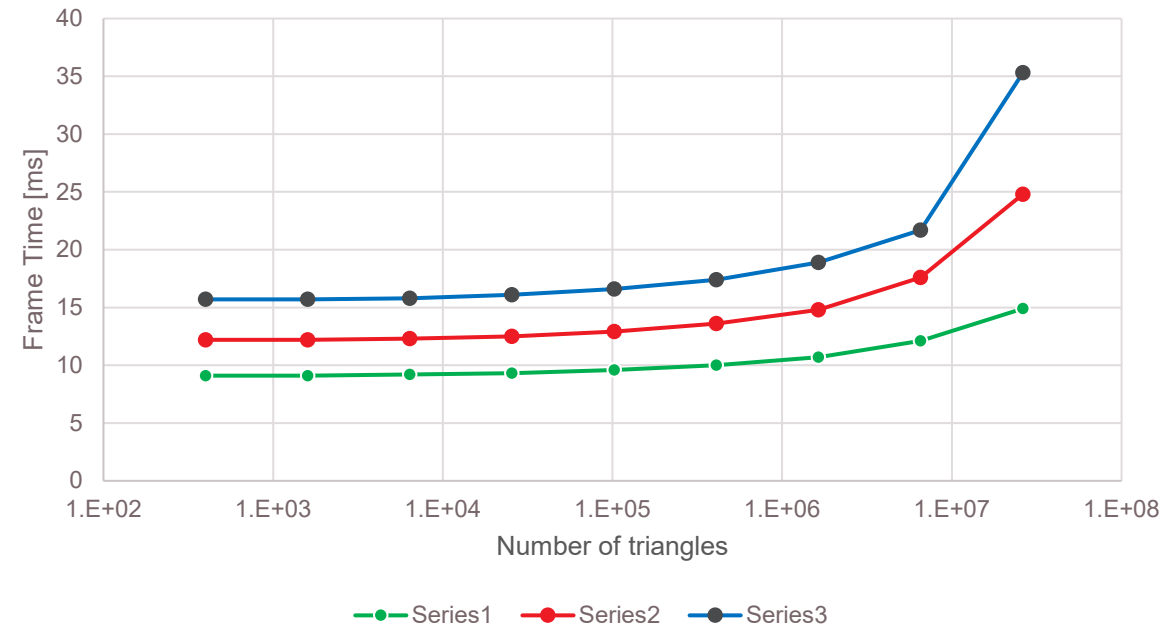
See endnotes for full test system configuration.

Bilinear Interpolation



Faster than Tessellation/Geometry pipeline!

Performance of three implementations of bilinear interpolation



GPU: AMD Radeon™ RX 7900 XTX; driver version: 24.1.1

MESH SHADERS SUBDIVISION SURFACES

High-Performance Graphics 2023
J. Bikker and C. Grimm (Eds.)

COMPUTER GRAPHICS Forum

Edge-Friend: Fast and Deterministic Catmull-Clark Subdivision Surfaces

Hans-Joachim Käfer¹, Max Oberberger², Matthias Czuprin², Quinn Meyer¹
¹Coburg University of Applied Sciences and Arts, Germany
²AMD, Germany

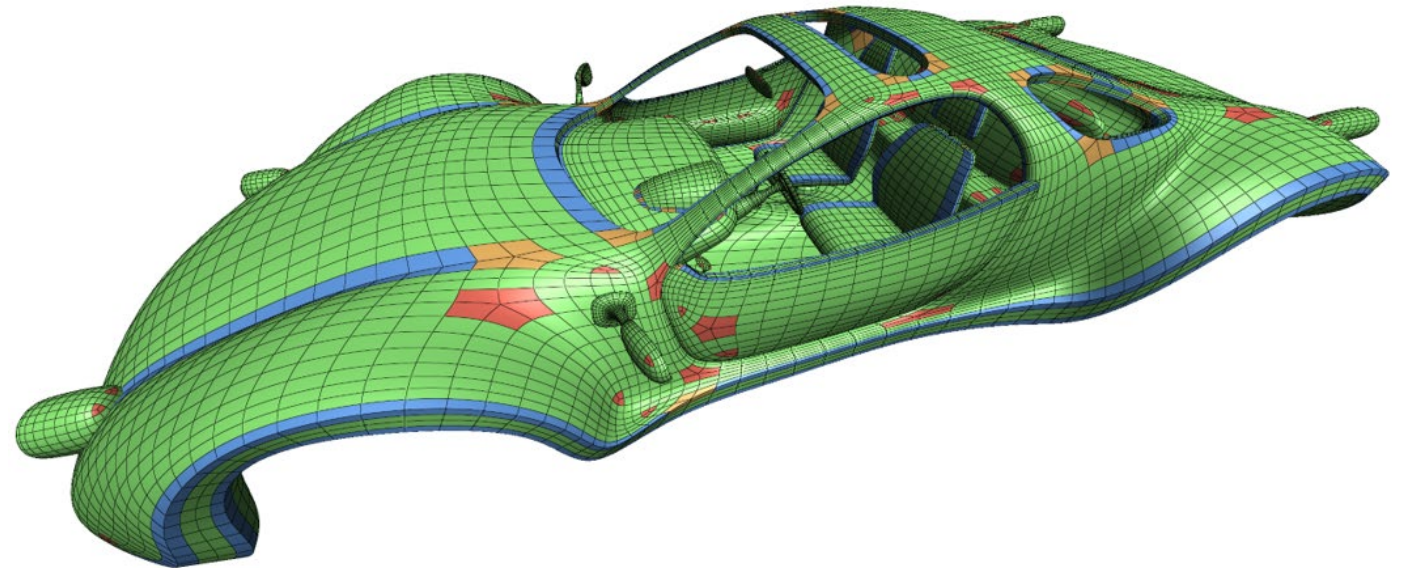
Figure 1: (a) Control mesh after a pre-processing subdivision iteration is quad-only. (b) One edge-friendly data structure implicitly stores two opposing edges (red) in each quad. Each quad stores only two edges in neighboring quads (blue). (c) We refine the edge-friendly structure through four levels of subdivision. (d) Refining and rendering the eleven model sides under alias on an AMD Radeon RX 7900 XTX GPU.

Abstract
 We present edge-friendly, a data structure for quad meshes with access to neighborhood information required for Catmull-Clark subdivision surface refinement. Edge-friendly maintains edge sets and does subdivision neighbor rendering. In particular, the rendering algorithm is deterministic, does not require hardware support for atomic float-point arithmetic, and is optimized for efficient rendering on GPUs. Edge-friendly explicitly stores two edges per quad, and edges can be uniquely and implicitly assigned to each quad. Additionally, edge-friendly is a compact data structure, adding little overhead. Our algorithm is simple to implement in a single compute shader kernel, and requires minimal synchronization which makes it particularly suited for asynchronous execution. We extend our kernel to support relevant Catmull-Clark subdivision surface features, including some mesh-to-mesh conversion. We extend our kernel to support relevant Catmull-Clark subdivision surface features, including some mesh-to-mesh conversion. We extend our kernel to support relevant Catmull-Clark subdivision surface features, including some mesh-to-mesh conversion. In case of topology changes, our data structure requires little pre-processing, making it amenable for a variety of applications, including real-time editing and animations. Our method can process and render billions of triangles per second on modern GPUs. For a single mesh, our algorithm processes and renders 2.9 million triangles in 0.25ms on an AMD Radeon RX 7900 XTX GPU.

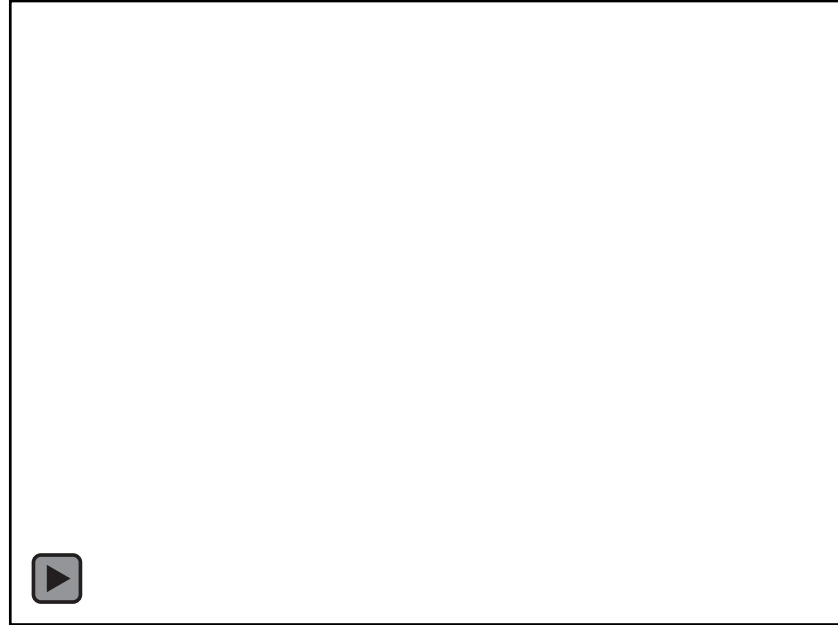
CCS Concepts
 • Computing methodologies → Rendering; Parametric curve and surface models; Massively parallel algorithms;

© 2023 The Author. Content not for redistribution. The European Conference on Computer Graphics and the Wiley Online Library on behalf of the European Association for Computer Graphics, and the ACM on behalf of the ACM SIGGRAPH Association. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

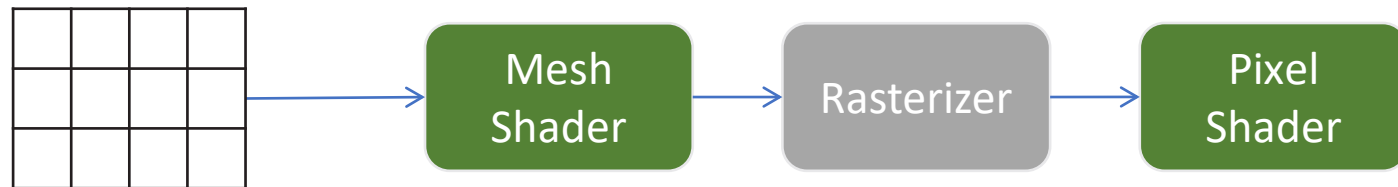
X2311-014883



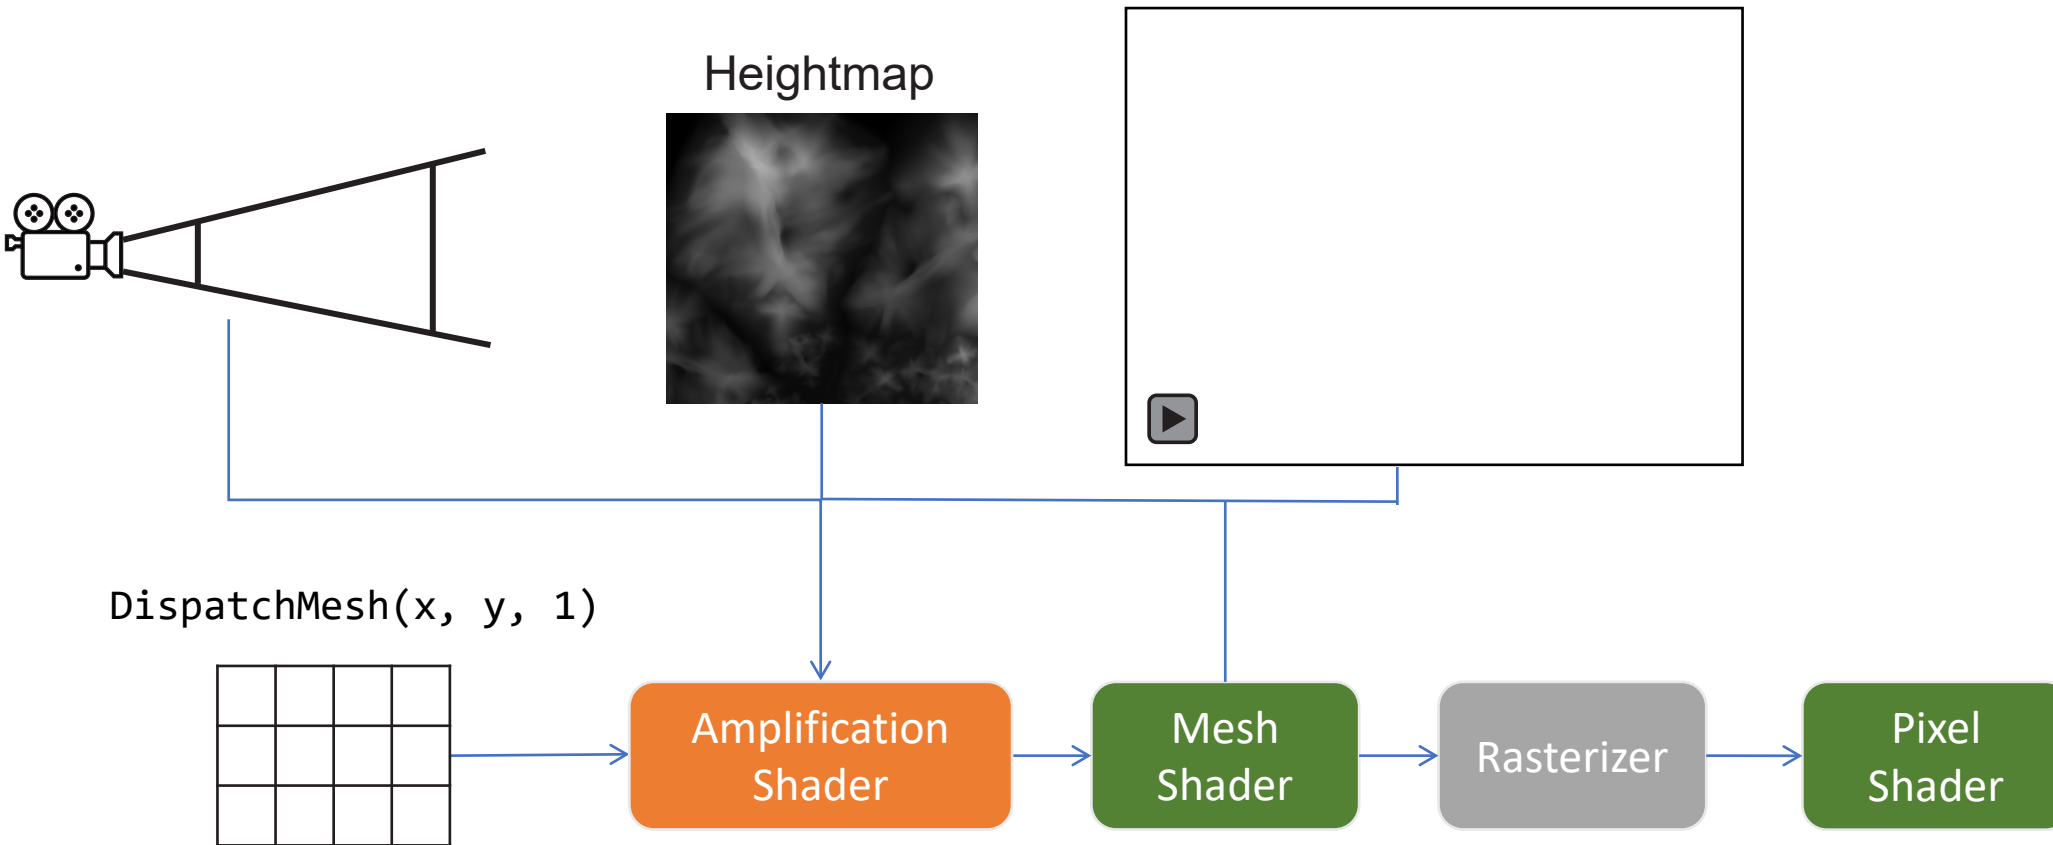
PROCEDURAL GEOMETRY



`DispatchMesh(1, 1, 1)`

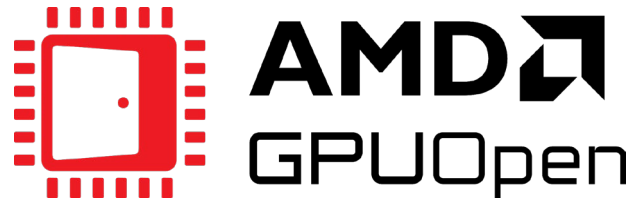


PROCEDURAL GEOMETRY





PROCEDURAL GEOMETRY



Learn more on
[GPUOpen.com](https://gpuopen.com)



Follow us on X

<https://twitter.com/GPUOpen>



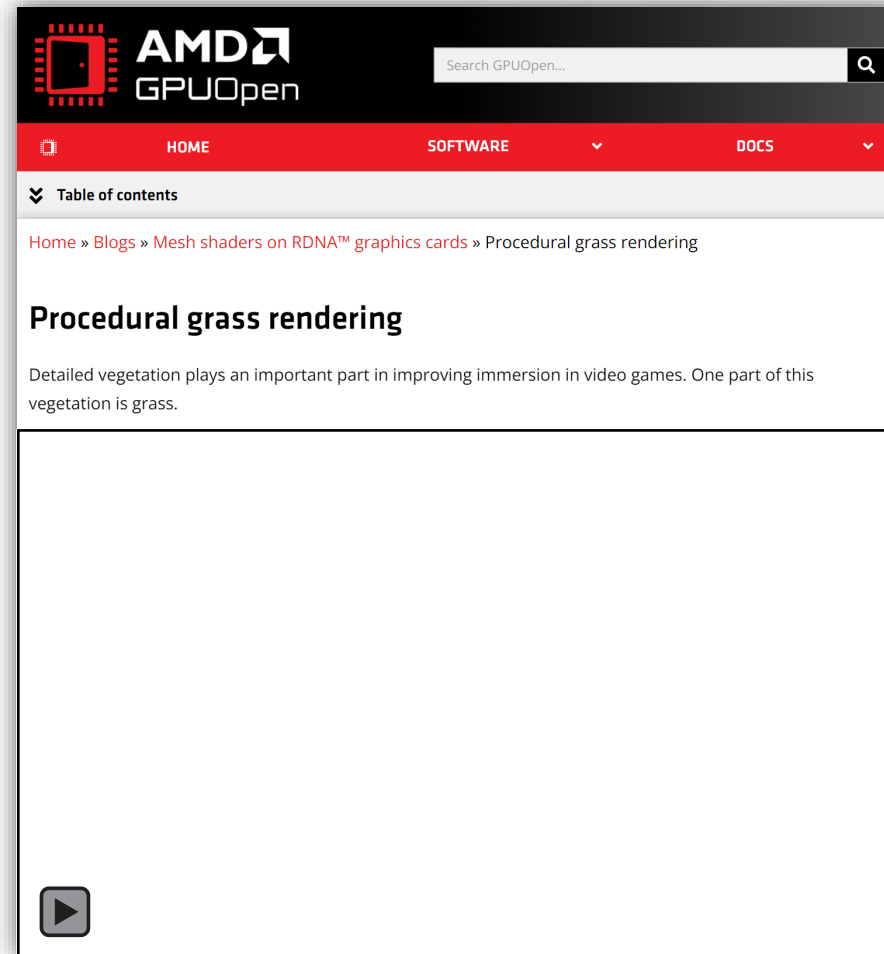
Follow us on Mastodon

<https://mastodon.gamedev.place/@gpuopen>



Follow us on Zhihu

<https://www.zhihu.com/org/gpuopen-7>

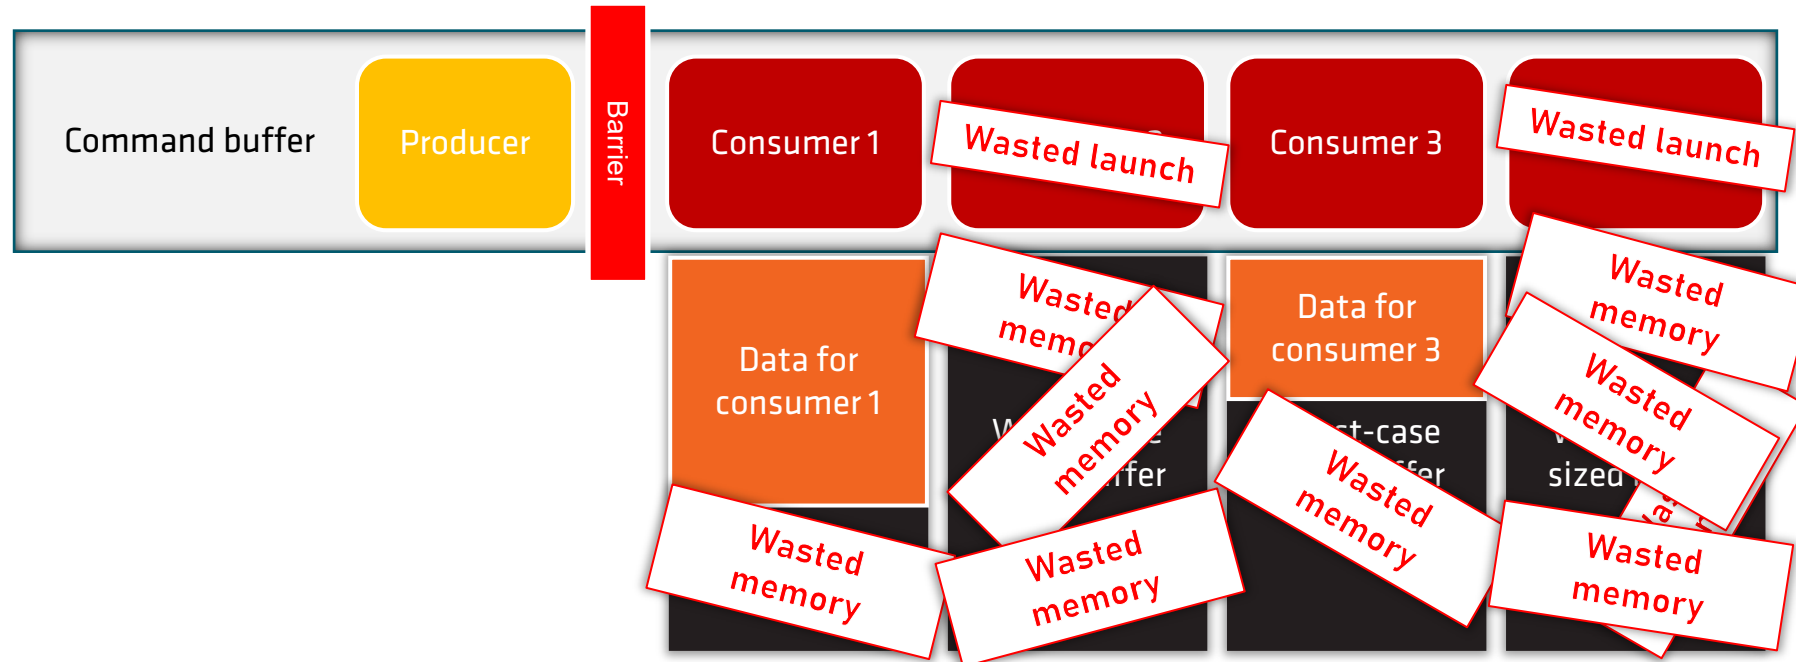


WORK GRAPHS 1.0

- Multi-year collaboration between Microsoft and industry partners
- Developer preview released June 2023
- Out of preview from March 2024
- <https://devblogs.microsoft.com/directx/d3d12-work-graphs/>

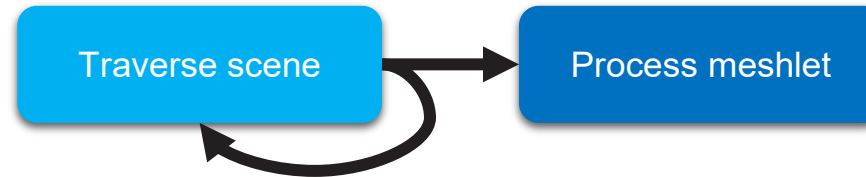
EXECUTE INDIRECT OVERHEAD

- Using execute indirect can introduce memory and performance overhead.

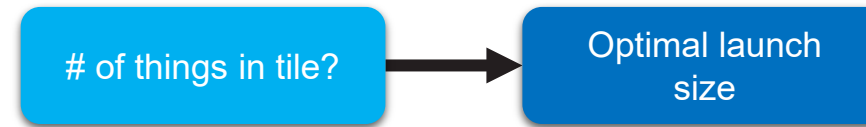


WORK GRAPHS ANSWER FOR EXECUTE INDIRECT LIMITATIONS

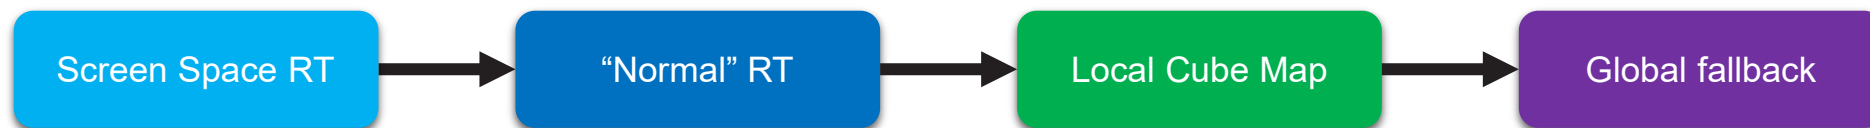
- Recursive algorithms: Scene traversal, ...



- Adaptive algorithms (launch more/less work): Physics, ...

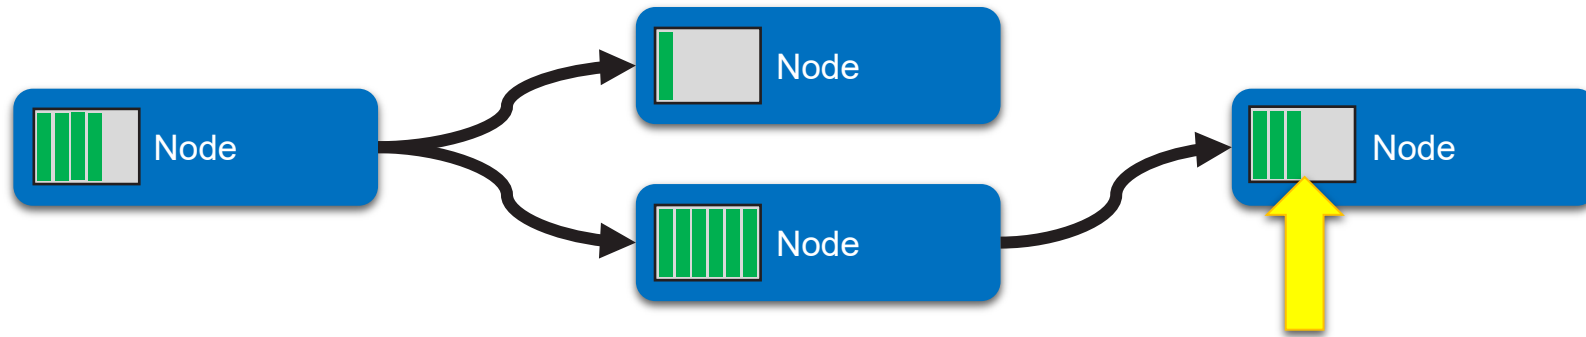


- Long execution chains: Lighting algorithms, ...



WORK GRAPHS IN A NUTSHELL

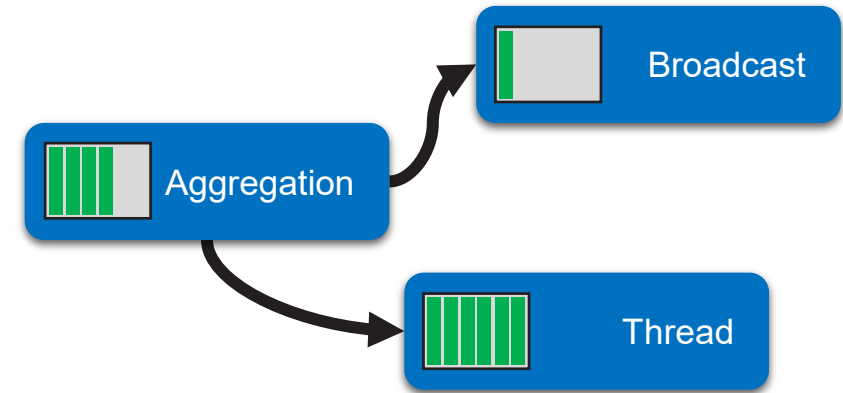
- **Nodes** connected with **edges**
- Each node has a virtual **queue**
- Nodes launch as soon as “enough” work waits for them
 - Enough depends on the GPU, driver, ...
 - Runtime can merge/fuse nodes, reorder outputs, sort, etc.



Scheduler launches dispatch to consume the data

GPU WORK GRAPHS

- a data flow model
- Work moves from node to node in the form of small “work items” (think: a struct)
- Work items get “queued up”
- Once enough work is pending, the GPU launches a dispatch

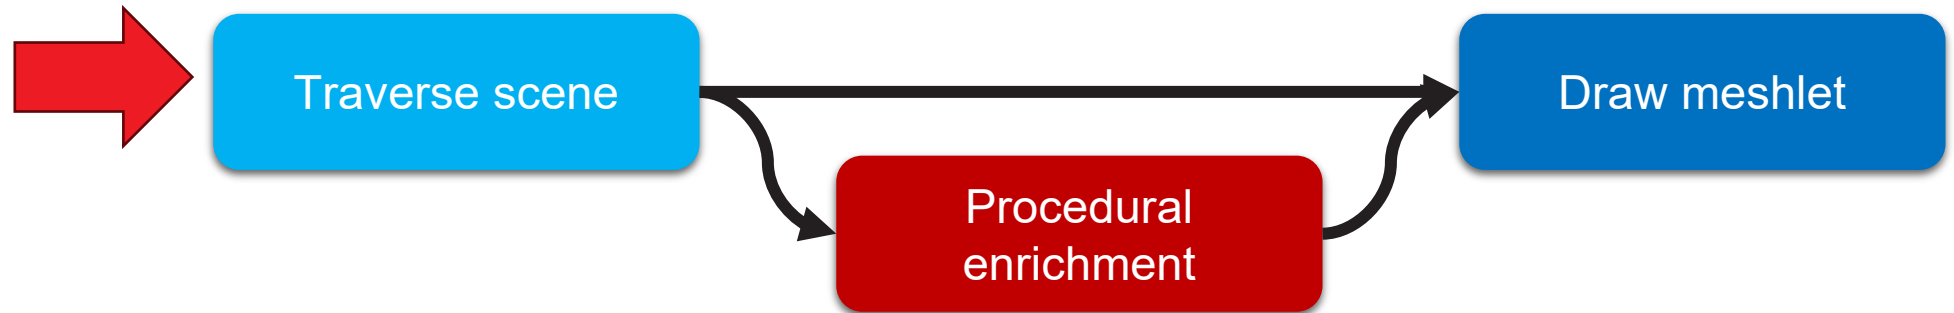


WORK GRAPHS – DRAW NODES

! PREVIEW !

- **Preview feature** announcement: Draw nodes
- Draw “inside” the work graph using “draw nodes”
- Enables fully compute-driven scene traversal (**with PSO switching**)
- Draw nodes: Feed into a **mesh shader pipeline**
- Work graph acts like an amplification shader on steroids

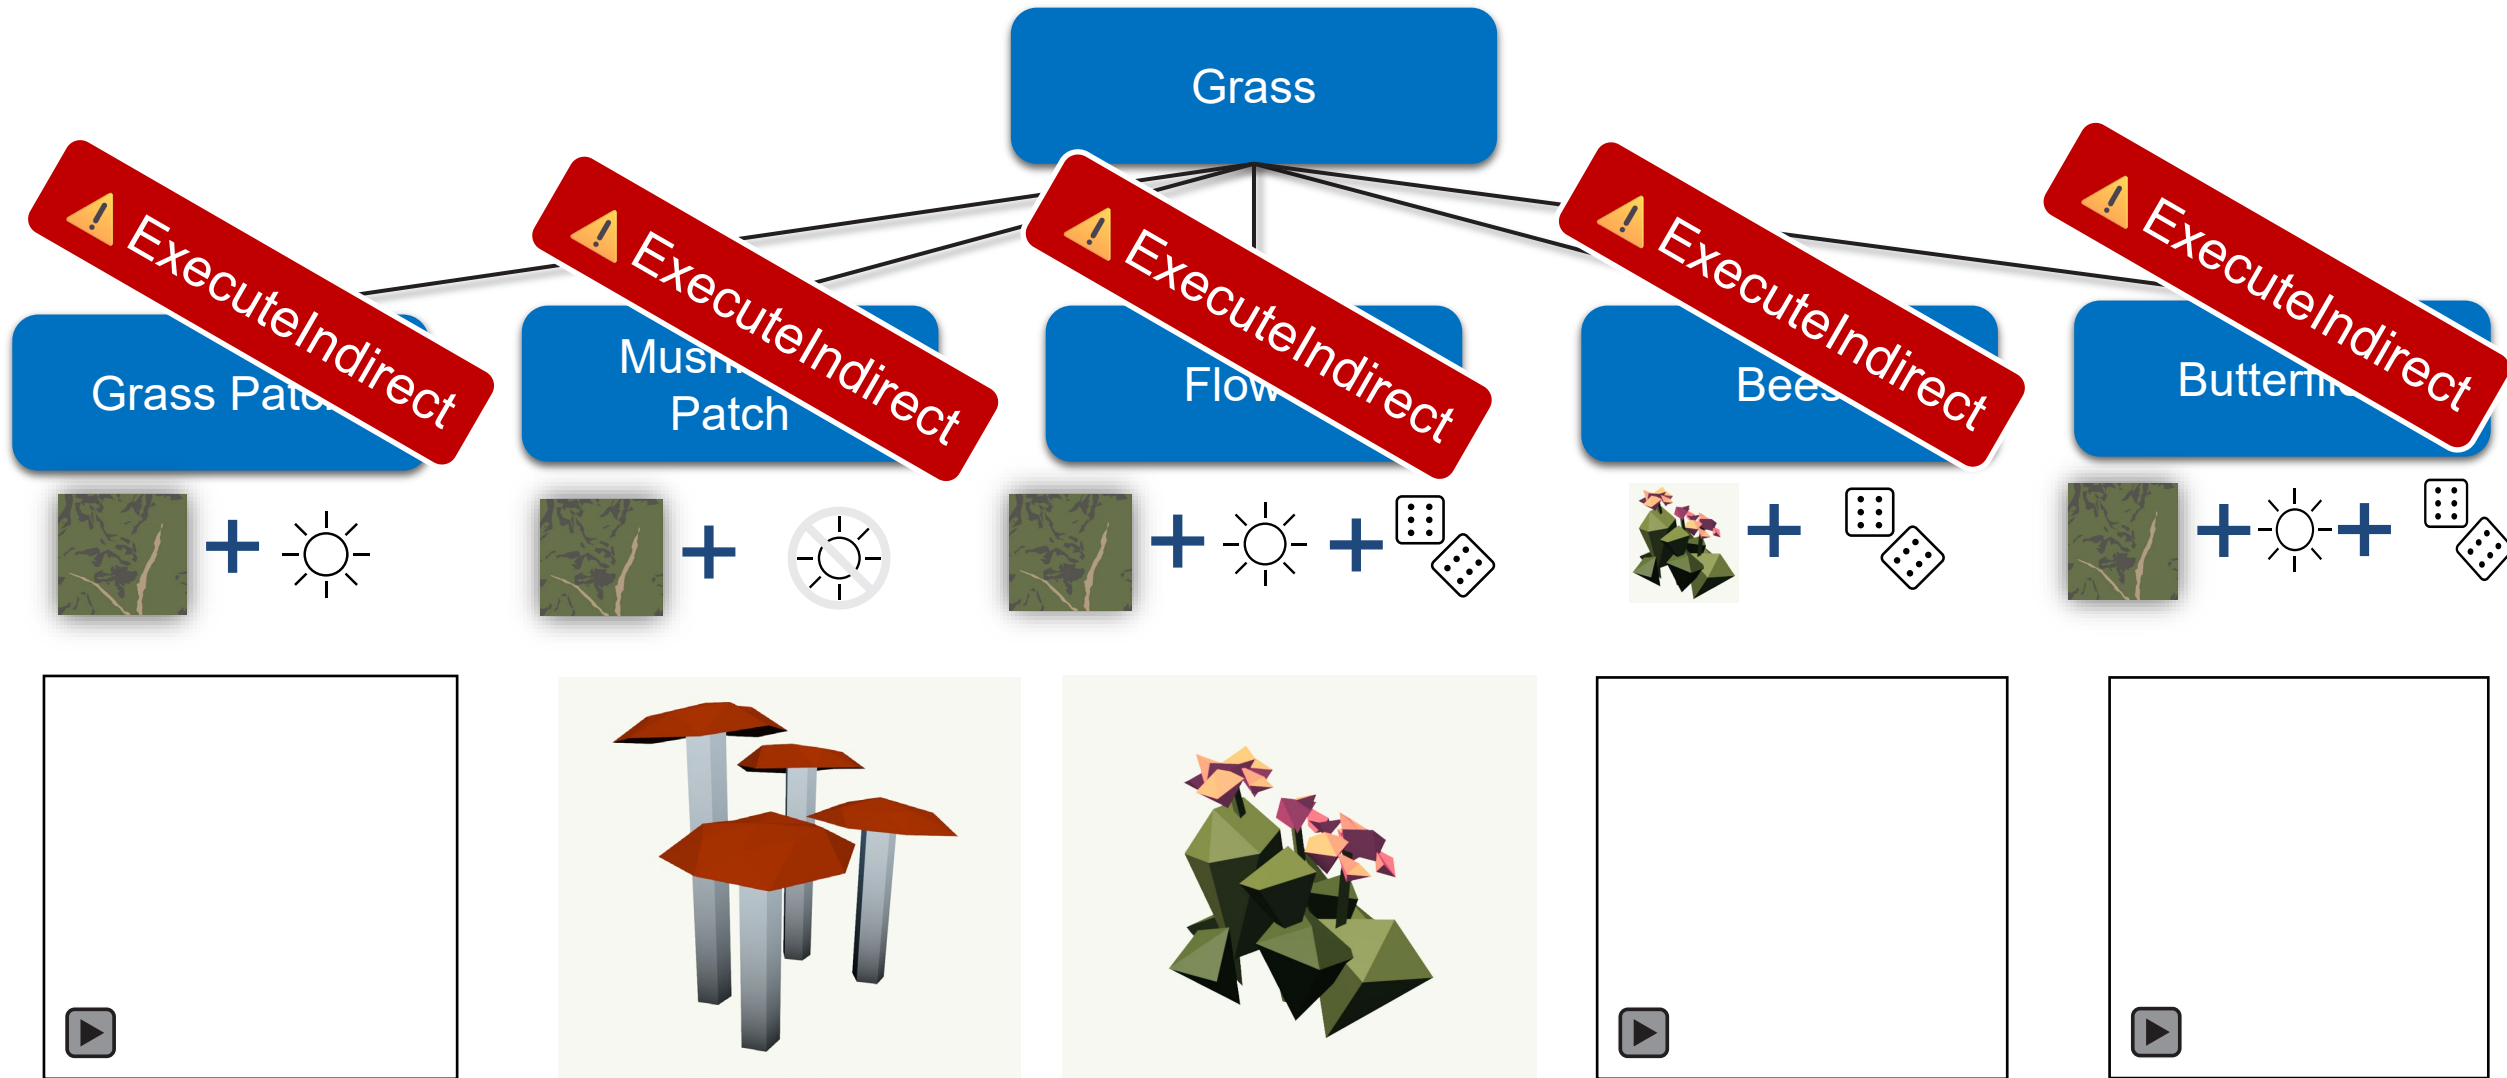
All in **one**
graph



“THE MEADOW” EXAMPLE



GRASS GENERATION WITHOUT WORK GRAPHS



GRASS GENERATION WITH MESH SHADERS

“⊘ No ExecuteIndirect”
one

Grass

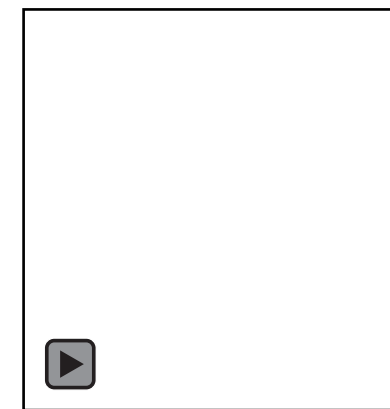
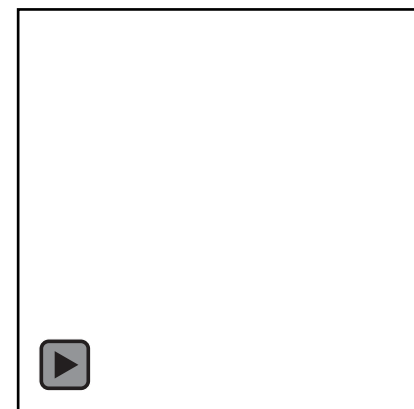
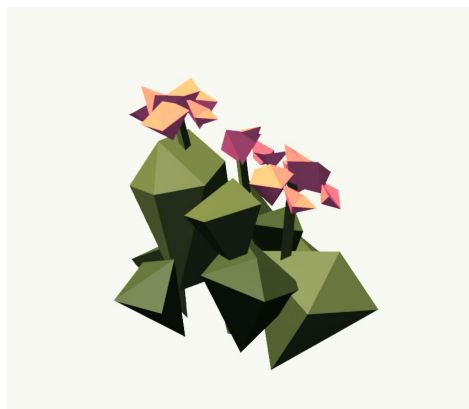
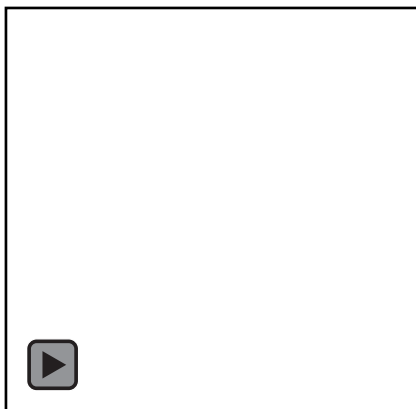
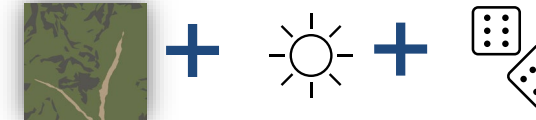
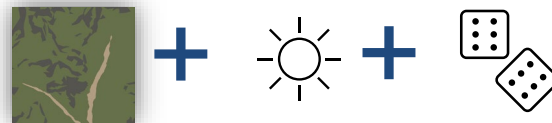
Grass Patch

Mushroom
Patch

Flower

Bees

Butterflies

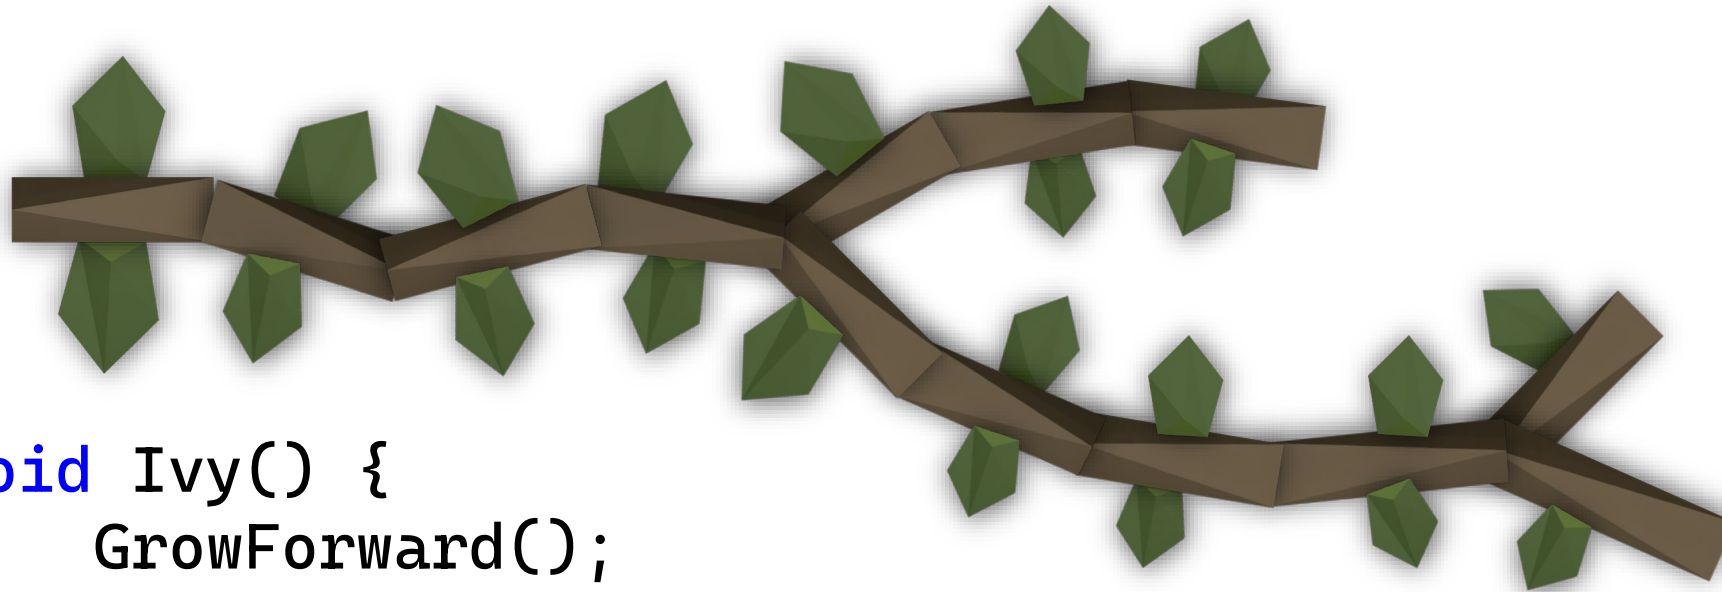


PREVIEW

“THE BRIDGE” EXAMPLE



RECURSIVE ALGORITHM EXAMPLE



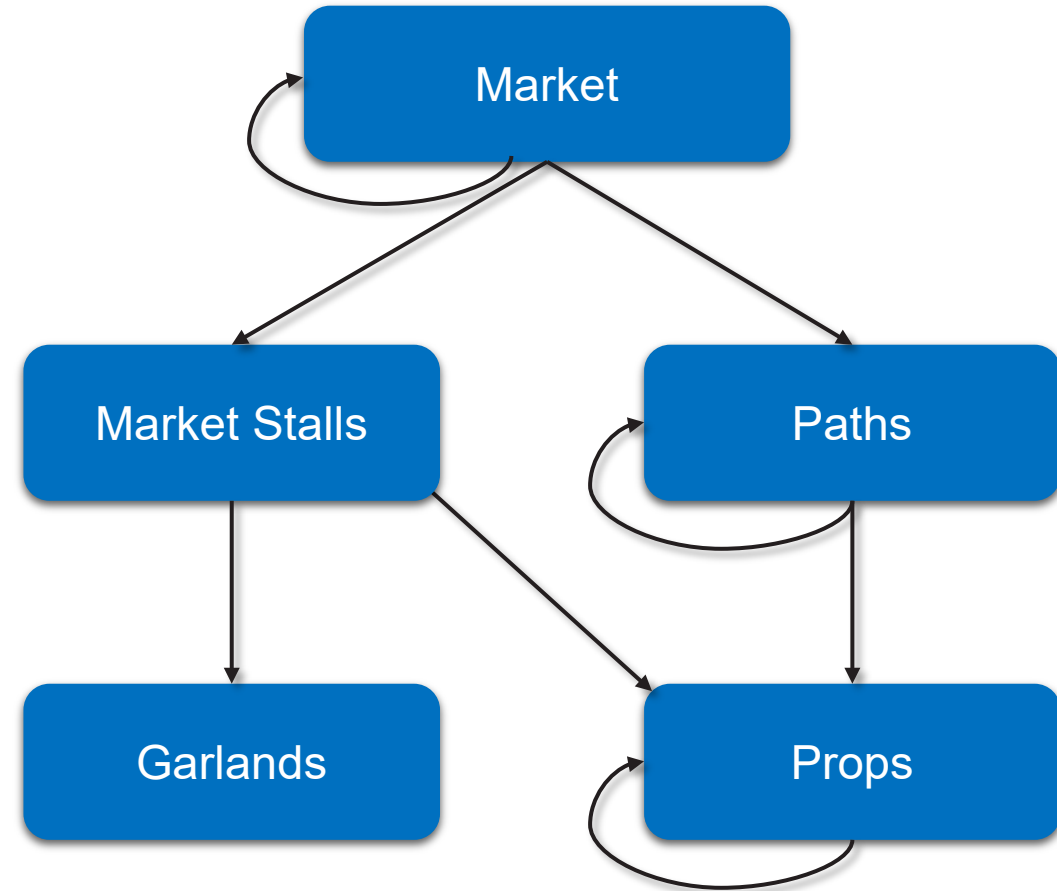
```
void Ivy() {  
    GrowForward();  
    EmitNextRecord();  
  
    if (forked) {  
        EmitNextRecord();  
    }  
}
```



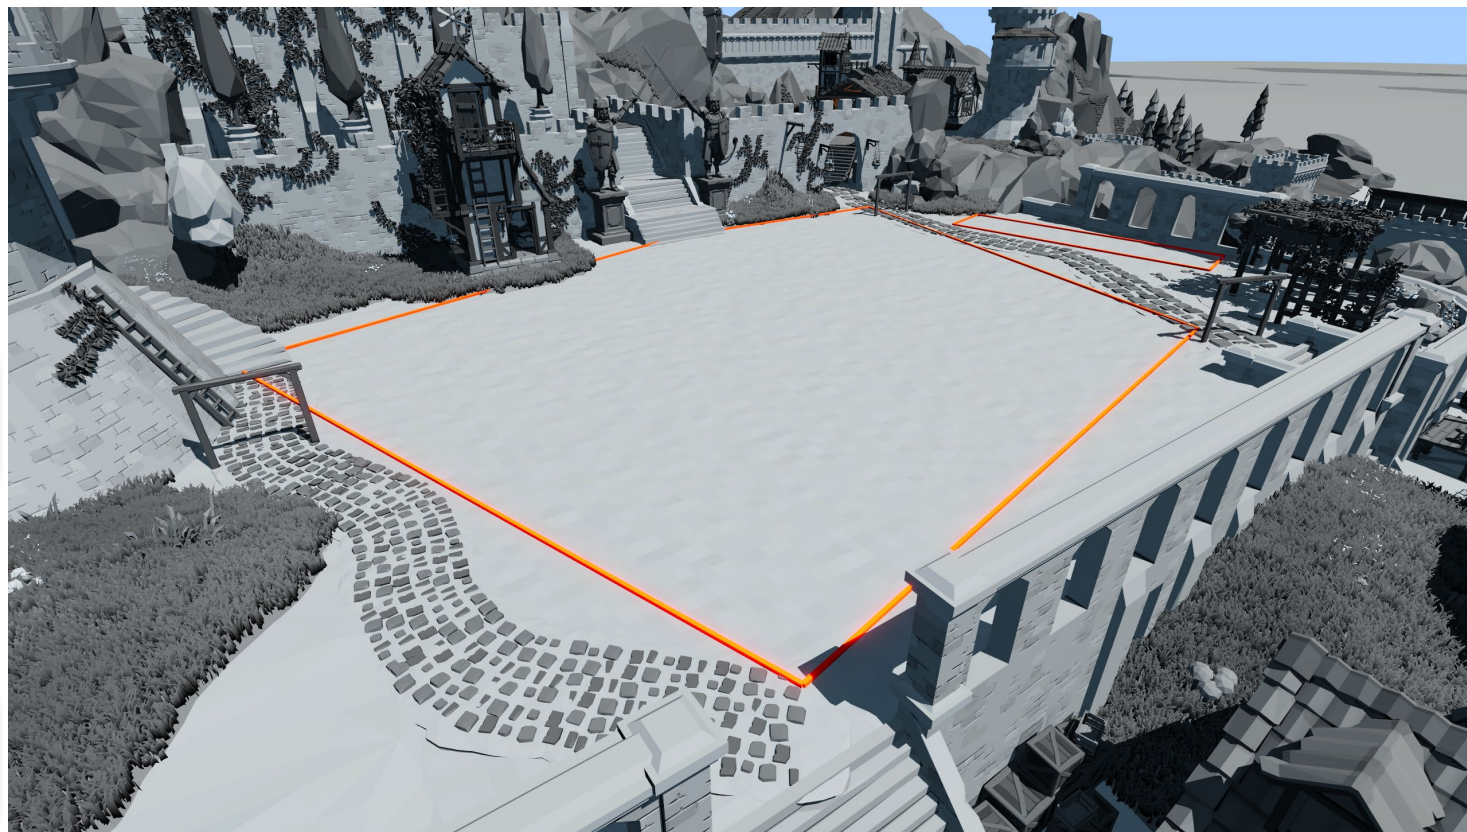
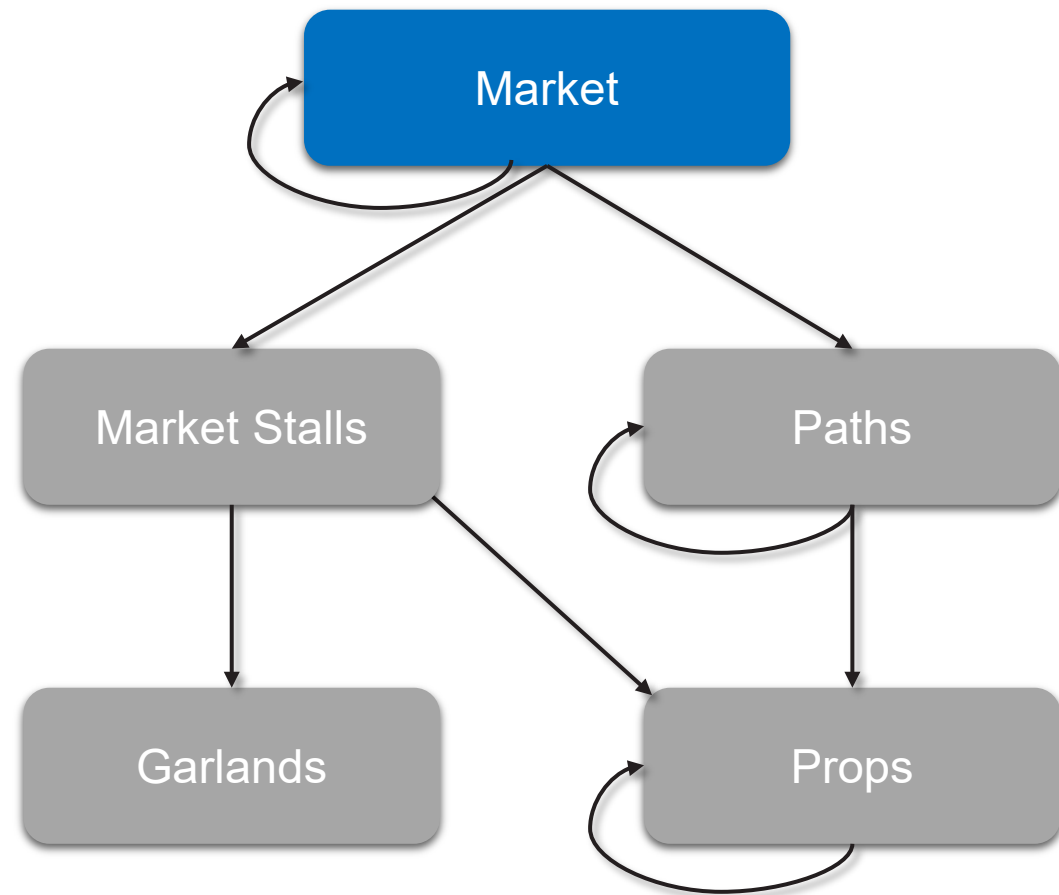
“THE MARKET” EXAMPLE



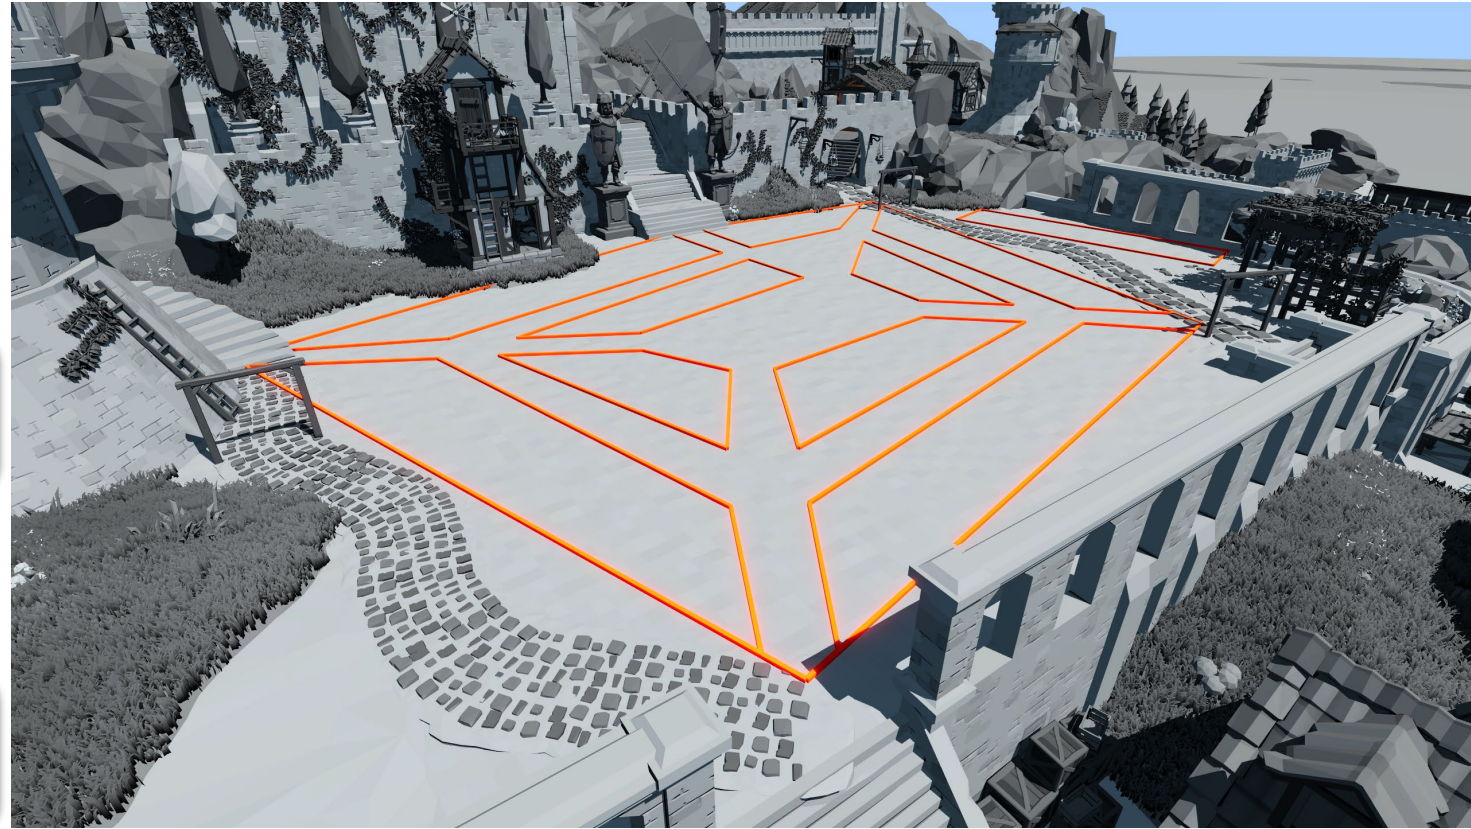
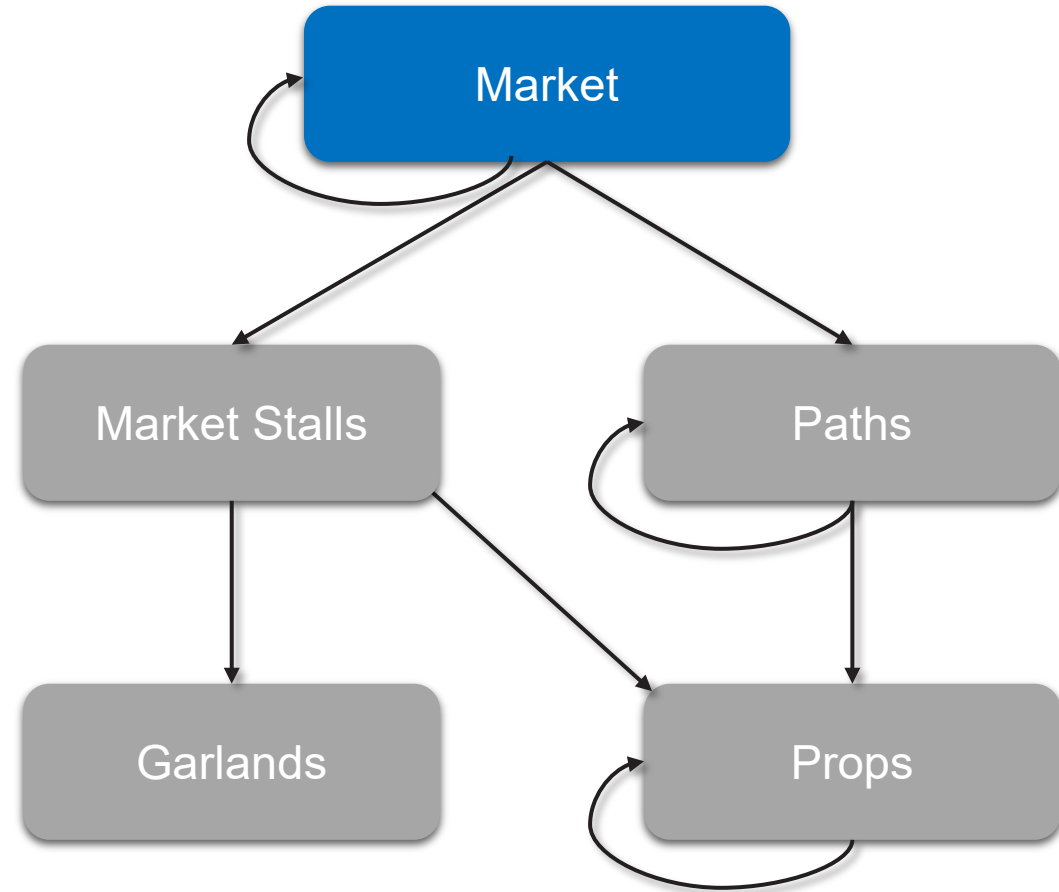
PROCEDURAL GENERATION EXAMPLE



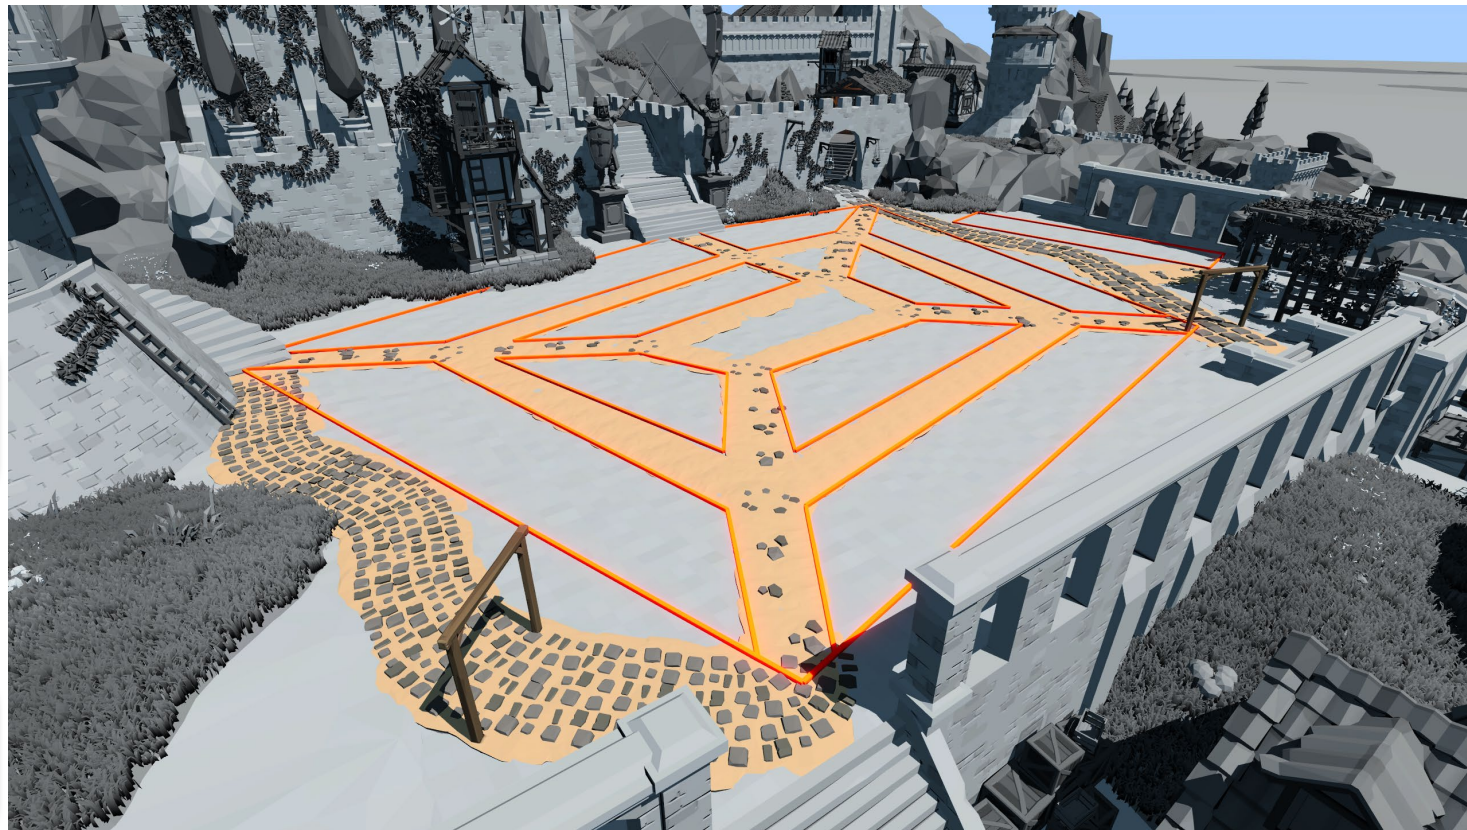
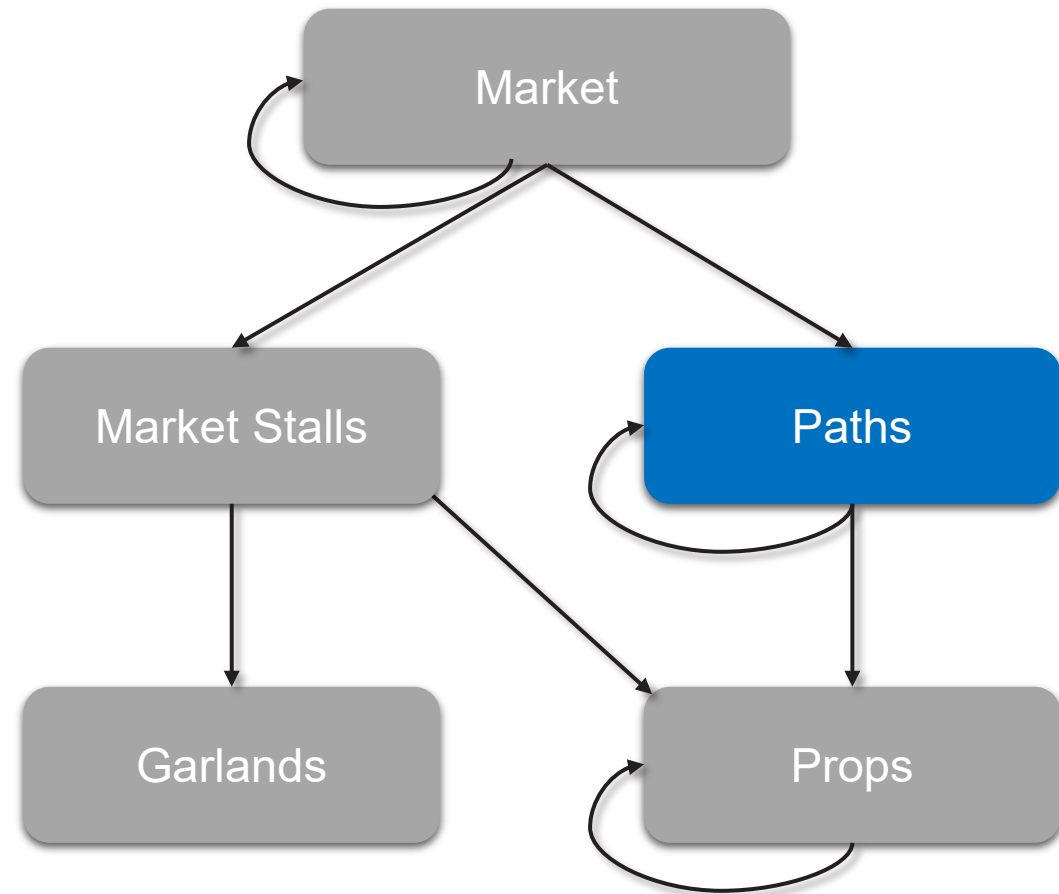
PROCEDURAL GENERATION EXAMPLE



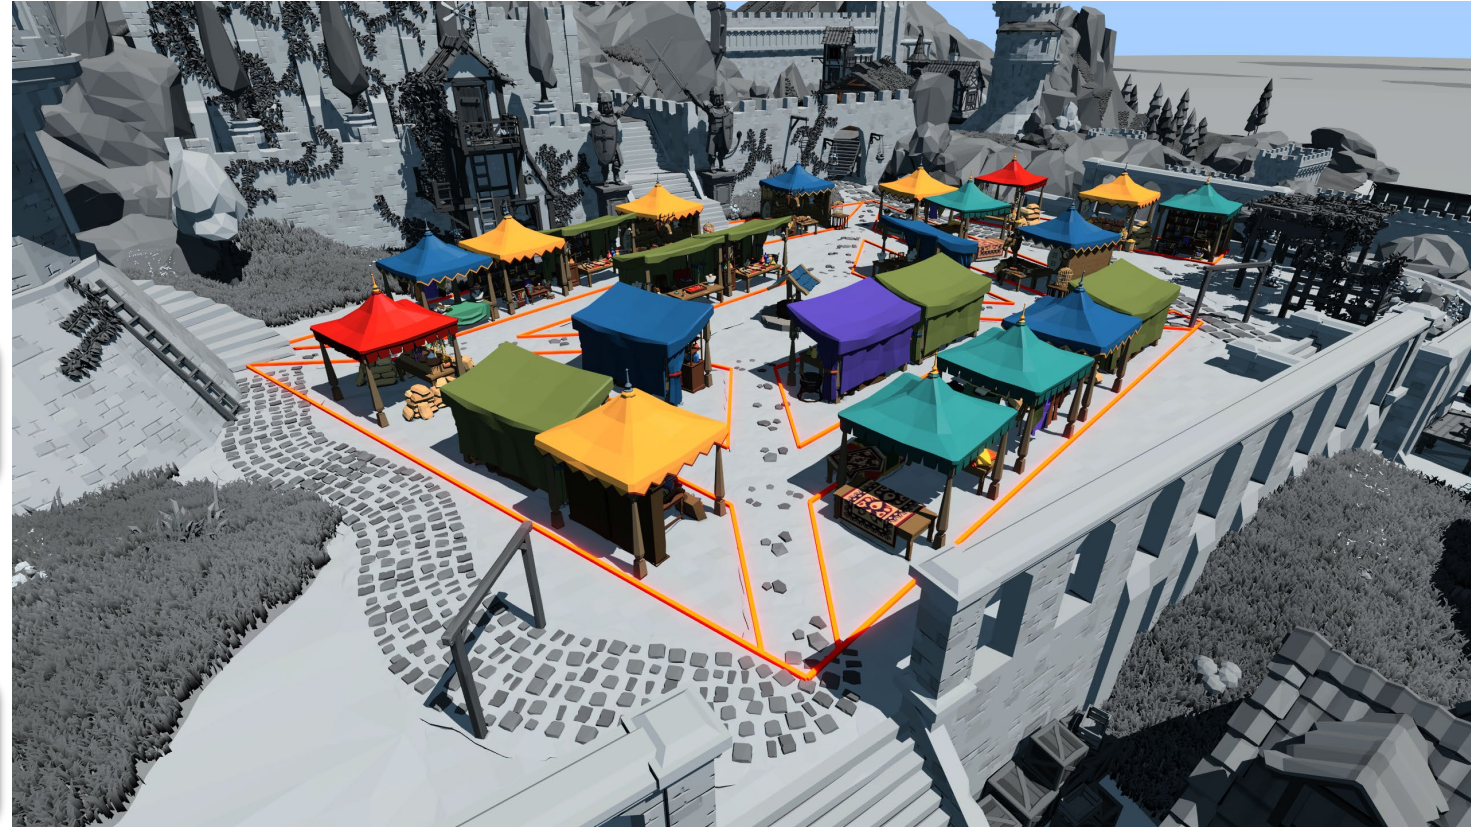
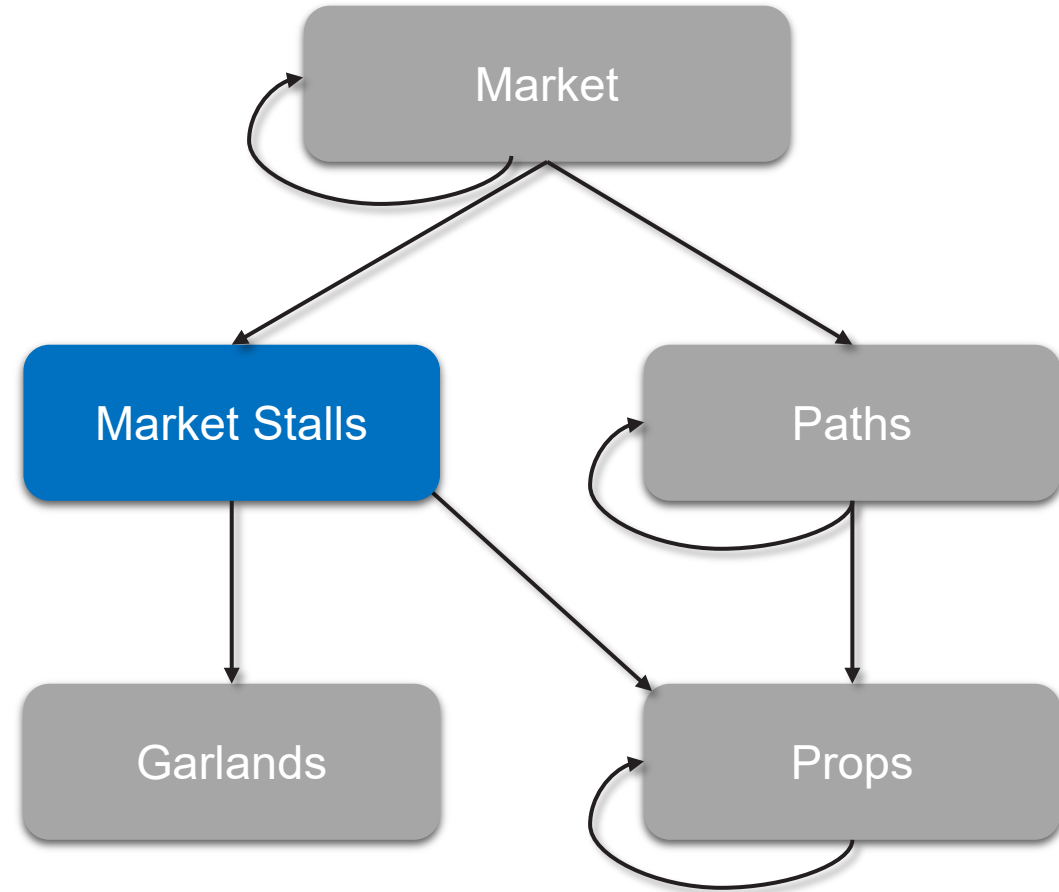
PROCEDURAL GENERATION EXAMPLE



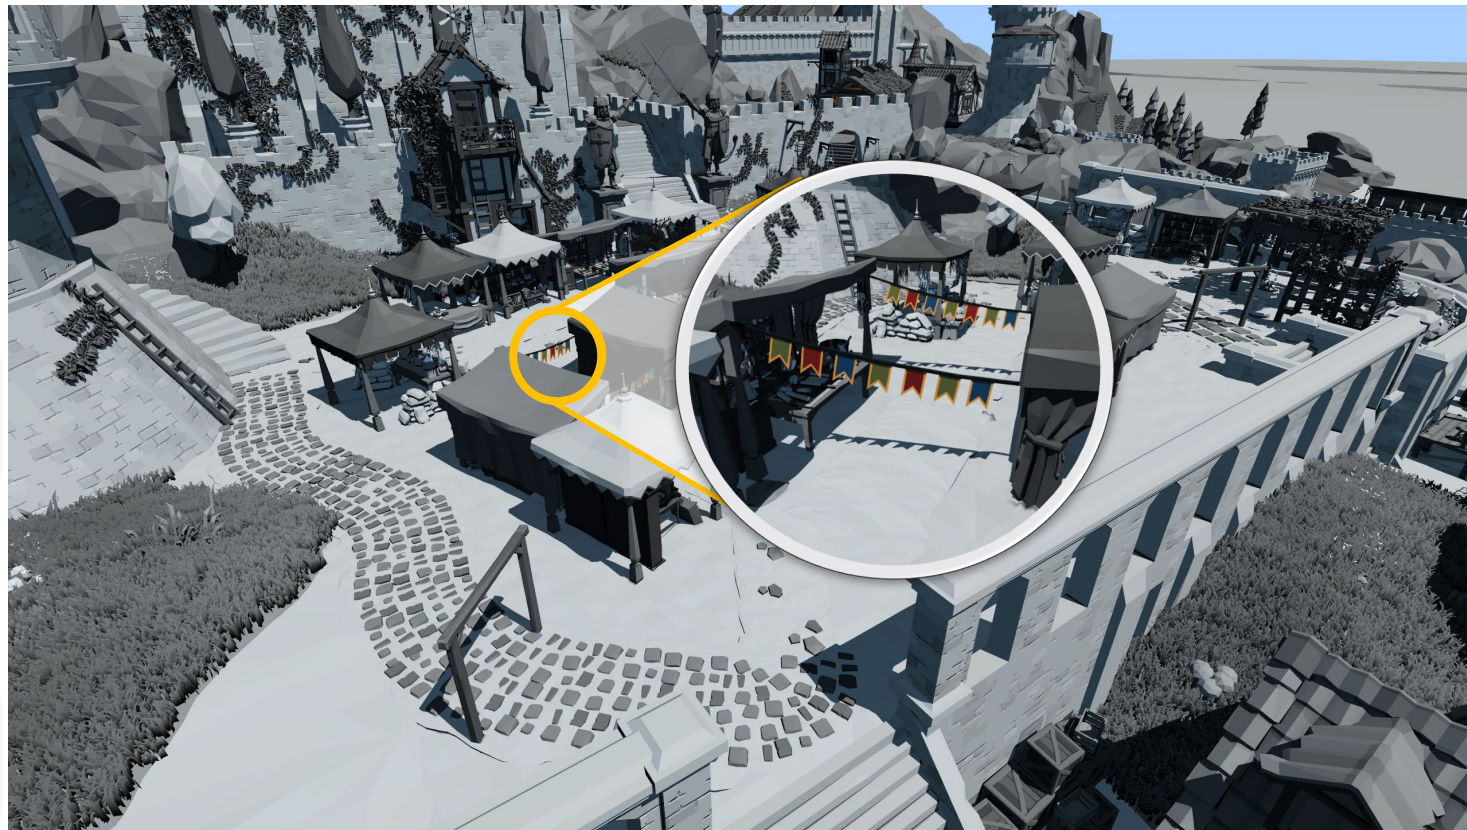
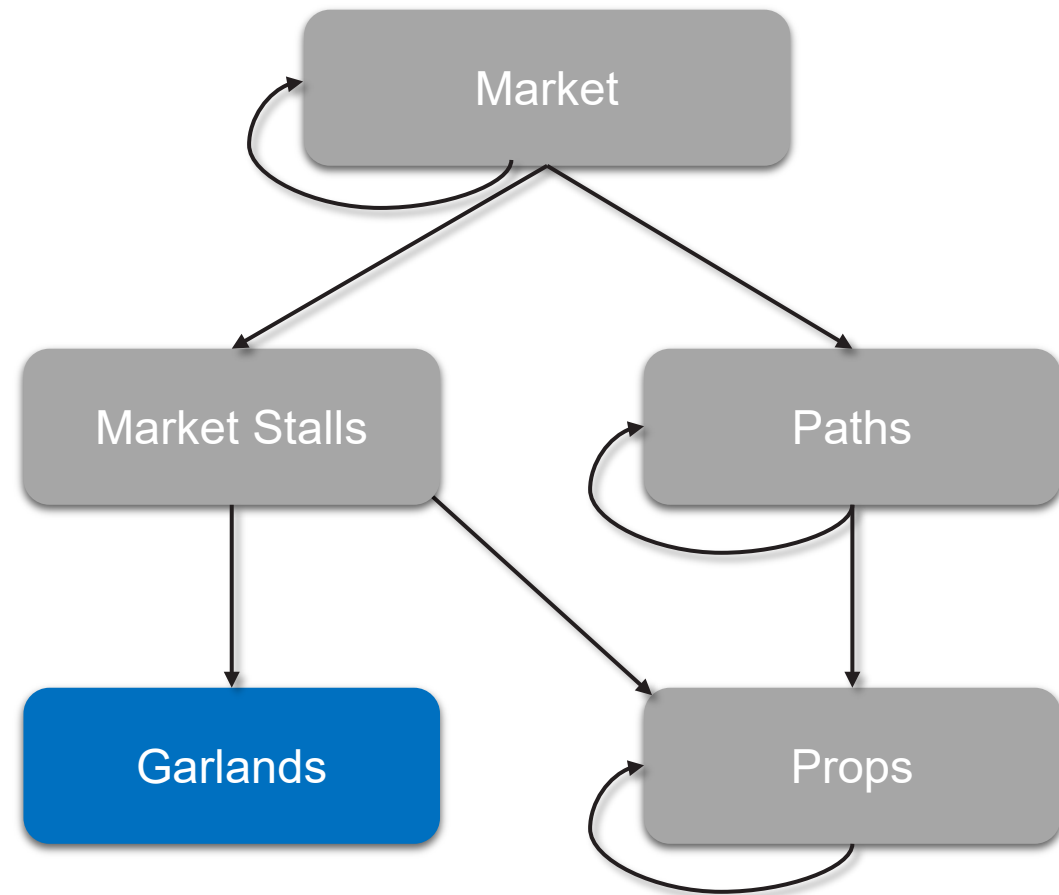
PROCEDURAL GENERATION EXAMPLE



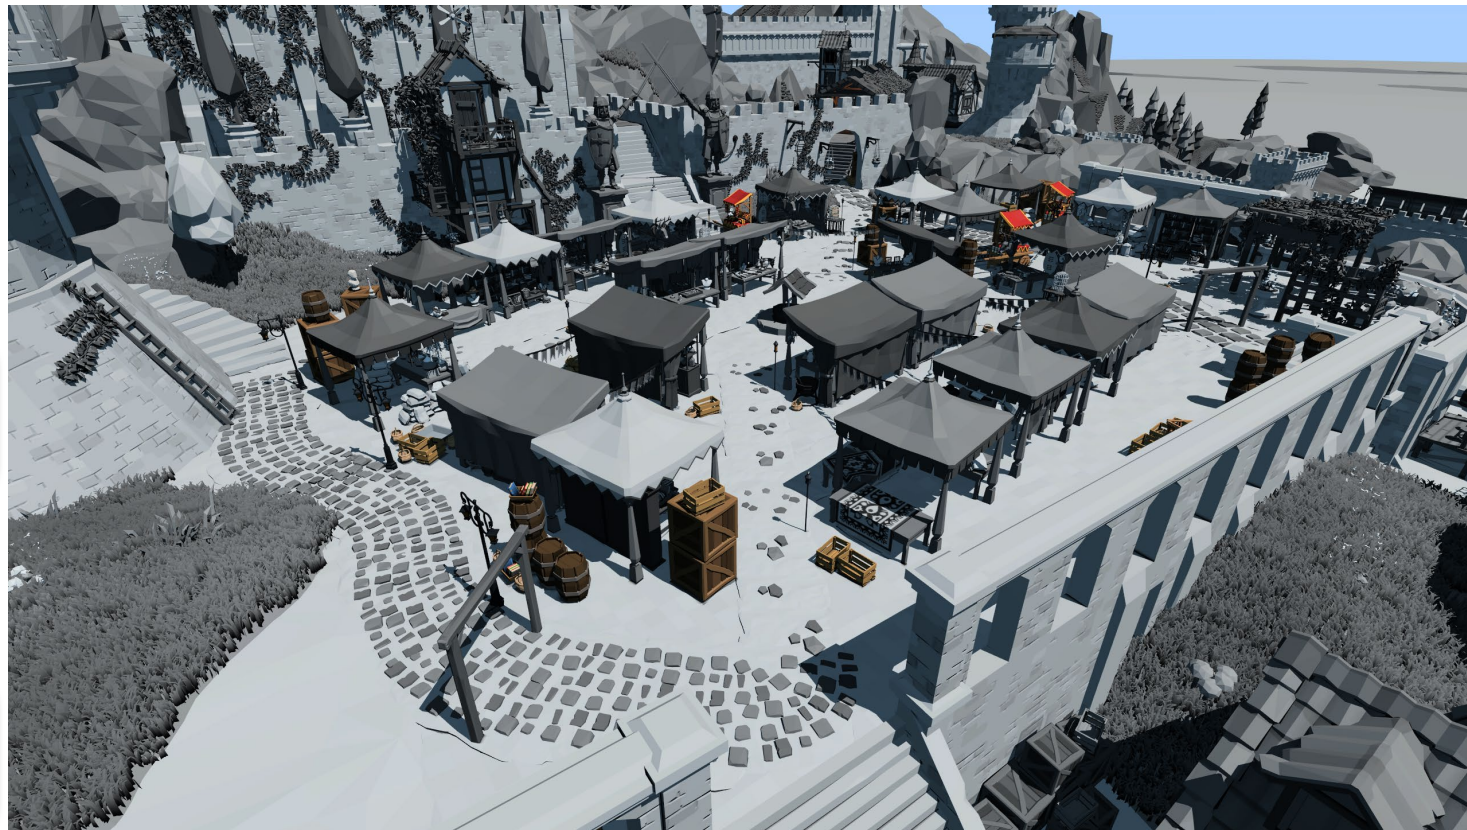
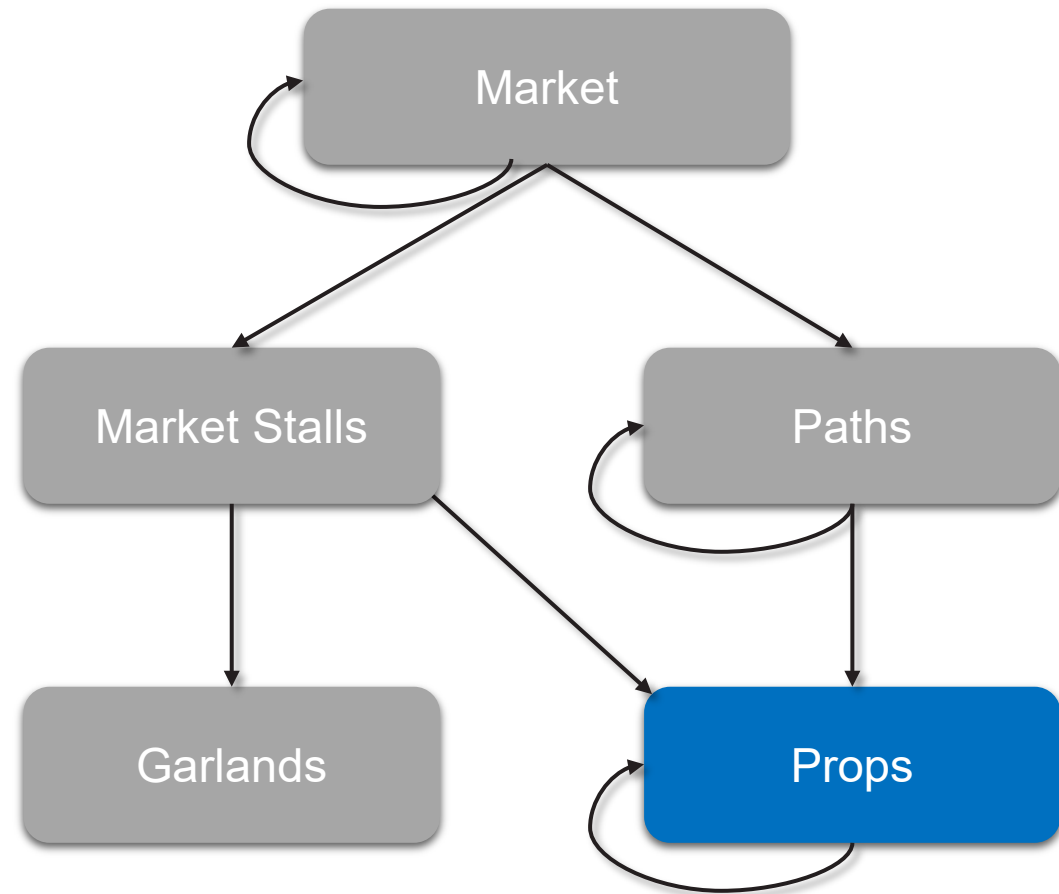
PROCEDURAL GENERATION EXAMPLE



PROCEDURAL GENERATION EXAMPLE



PROCEDURAL GENERATION EXAMPLE



SPECIAL THANKS

- Lou Kramer
- Max Oberberger
- Matthäus G. Chajdas

DISCLAIMER

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Performance testing (slides 23, 30, and 31) done on following system: AMD Ryzen™ 9 7950X, 128GB DDR5-6000 memory, Asus ROG CROSSHAIR X670E HERO Motherboard, XFX RX 7900 XTX, 4TB M.2 NVME SSD, Windows® 11 Pro 22H2, AMD Software: Adrenalin Edition 23.40.14.01

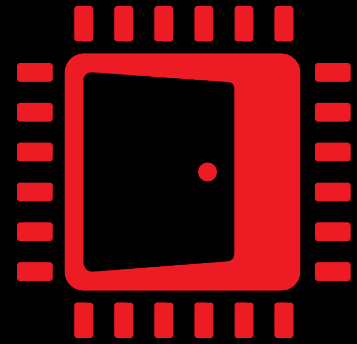
Performance testing (slide 39) done on following system: AMD Ryzen™ 7 5800X, AMD Radeon™ RX 7900 XTX, AsusTeK TUF Gaming X570-Plus, 32GB DDR4-3600, Windows 10 Home 22H2, AMD Software: Adrenalin Edition 24.2.1

© 2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Radeon, RDNA, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Vulkan and the Vulkan logo are registered trademarks of the Khronos Group Inc.

DirectX is a registered trademark of Microsoft Corporation.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners.



AMD 
GPUOpen

AMD 
together we advance_

AMD 
EPYC

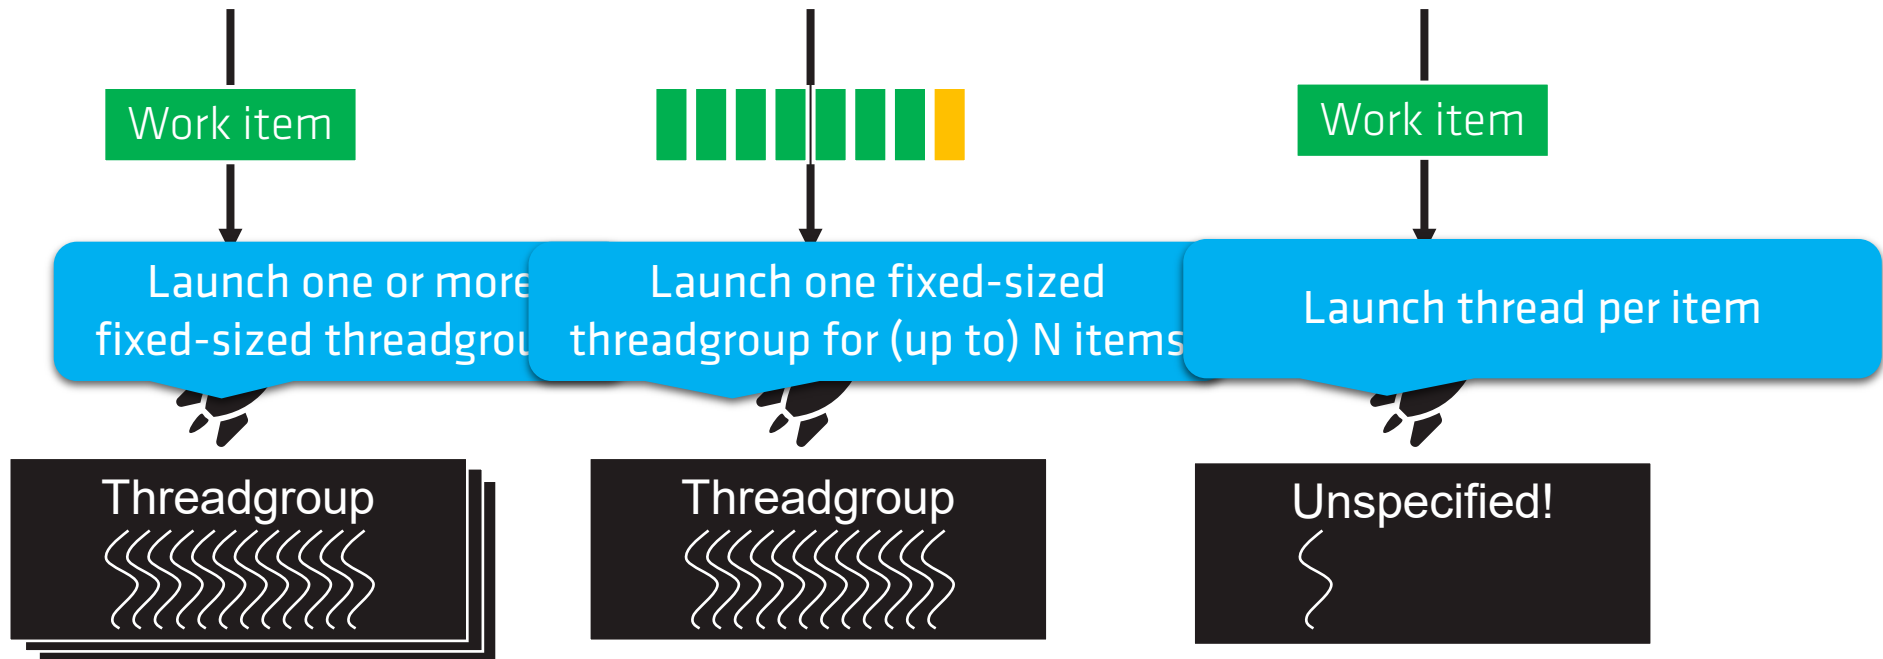
AMD 
RYZEN

AMD 
RADEON

AMD 

together we advance_

- You can select how things launch. Work items can ...
 - trigger a dispatch (“broadcast”)
 - be aggregated (“coalescing”)
 - be treated as independent launches (“thread”)



WORK GRAPHS – DRAW NODES

- Draw nodes: Feed into a mesh shader pipeline
- Work graph acts like an amplification shader on steroids
- Runtime ensures PSO switching isn't too expensive
- Will buffer up draw calls per state
- Will optimize state changes
- The more similar the states are, the better – the cheapest state change is swapping out shaders only

