

# Device-Generated Commands in Vulkan (VK\_EXT\_device\_generated\_commands)

Ricardo Garcia

XDC 2024 - October 10 - Montreal



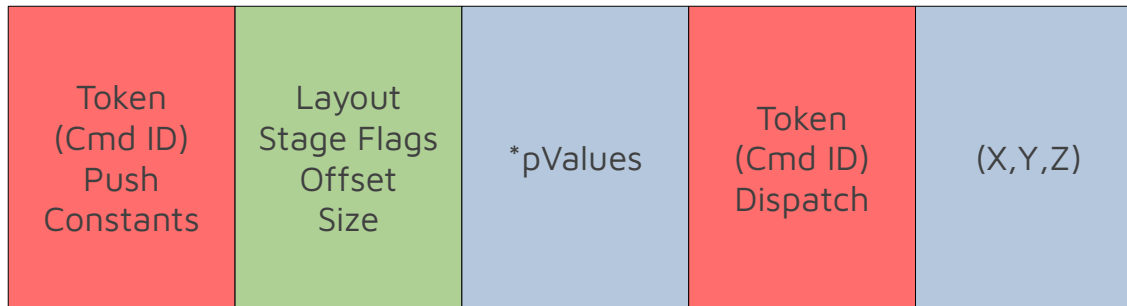
# What are Device-Generated Commands?

- One step ahead of indirect draws and dispatches
- One step behind work graphs
- Allows implementations to read command sequences from a regular buffer instead of a command buffer
- That buffer could be filled from the GPU to achieve GPU-driven rendering



# Naïve CPU-based Approach


- 1) vkCmdPushConstants(layout, stageFlags, offset, size, pValues)
- 2) vkCmdDispatch(x, y, z)




# VK\_EXT\_device\_generated\_commands

- VkIndirectCommandsLayoutEXT

- 1) vkCmdPushConstants
- 2) vkCmdDispatch



- Buffer contains a number of fixed-size sequences and each follows the layout



# Restricted Command Selection

VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_EXECUTION\_SET\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_PUSH\_CONSTANT\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_SEQUENCE\_INDEX\_EXT

VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_INDEX\_BUFFER\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_VERTEX\_BUFFER\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DRAW\_INDEXED\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DRAW\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DRAW\_INDEXED\_COUNT\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DRAW\_COUNT\_EXT

VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DISPATCH\_EXT

VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_TRACE\_RAYS2\_EXT

VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DRAW\_MESH\_TASKS\_NV\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DRAW\_MESH\_TASKS\_COUNT\_NV\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DRAW\_MESH\_TASKS\_EXT  
VK\_INDIRECT\_COMMANDS\_TOKEN\_TYPE\_DRAW\_MESH\_TASKS\_COUNT\_EXT



# Indirect Execution Sets

- A group of similar pipelines or shader objects
- All state must be identical (only shaders change)
- Each pipeline/shader has an index in the set
- The IES to use is specified beforehand and the DGC buffer contains indices into the set



# Quick How-To

- 1) Create the commands layout, and IES if needed (VkIndirectCommandsLayoutEXT, VkIndirectExecutionSetEXT)
- 2) Establish the maximum number of sequences
- 3) Query the required preprocess buffer size (vkGetGeneratedCommandsMemoryRequirementsEXT)
- 4) Allocate DGC buffer and preprocess buffer
- 5) Record commands and state almost normally (including work that fills the DGC buffer)
- 6) Dispatch work with vkCmdExecuteGeneratedCommandsEXT
- 7) If using explicit preprocessing (e.g. Proton does it to improve performance):
  - a) Use a separate command buffer for it
  - b) Pass the main command buffer in as state
  - c) Call vkCmdPreprocessGeneratedCommandsEXT and submit this work first, synchronizing with vkCmdExecuteGeneratedCommandsEXT



# Thanks for watching!

Join us!

<https://www.igalia.com/jobs>

