UE4 Mobile Performance

Niklas Smedberg
Senior Engine Programmer, Epic Games

Content

- Part 1: Understanding mobile performance
  - Mobile GPU Hardware
  - Thermal limits
  - Performance guidelines
- Part 2: Adapt and conquer
  - Cross-platform profiling
  - Platform-specific profiling
  - Scaling your game based on device

Part 1: Understanding Mobile Performance

- Mobile hardware is evolving at a crazy rapid rate
- Next-generation mobile GPUs:
  - Fully featured (DirectX 11)
  - Peak performance comparable to Xbox 360 and PS3
    - 300+ GFLOPS and 26 GB/s
  - Able to run full UE4 desktop high-end rendering pipeline (e.g. NVIDIA K1)
- Phone users upgrade hardware very frequently
  - But tablet users don't
  - Also, new large low-price markets are opening up
  - Result: Extremely wide performance range

Performance Trends (FP16 GFLOPS)

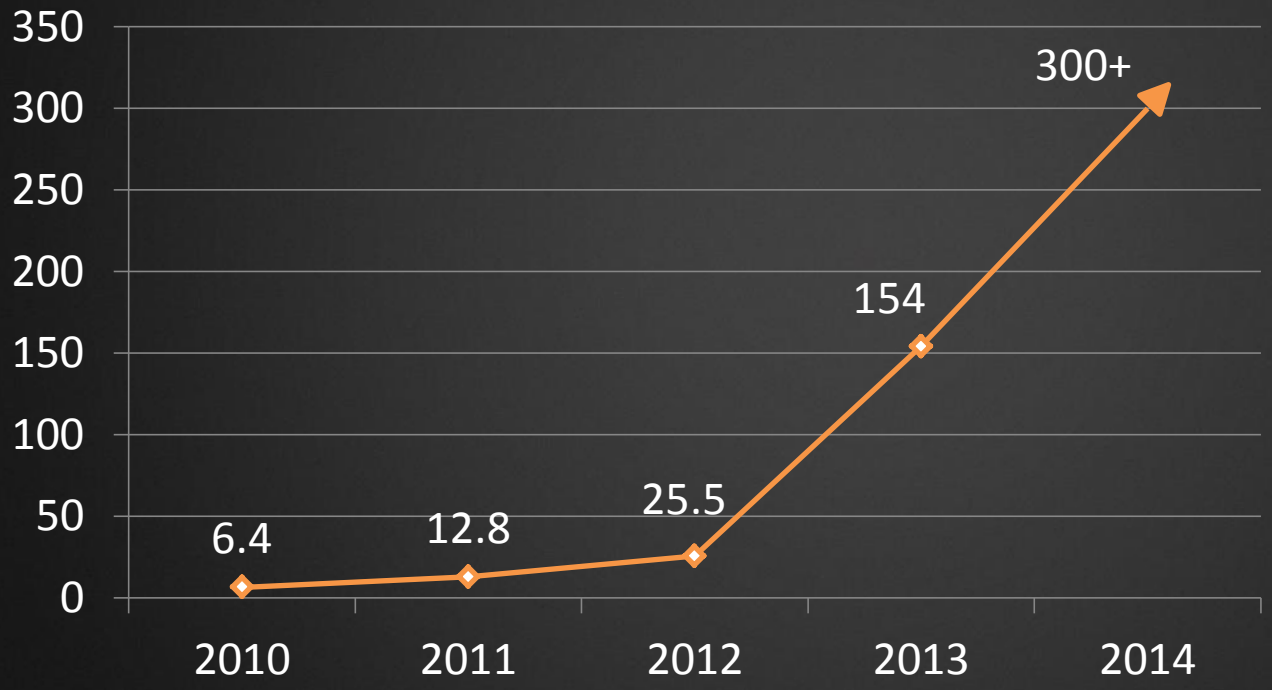

2010 SGX 535
2011 SGX 543MP2
2012 SGX 543MP3
2013 G6430
2014 Adreno, K1, GX6650

Common Mobile GPU Families

Qualcomm Snapdragon Adreno

Old: Adreno 2xx	Now: Adreno 3xx	Soon: Adreno 4xx

ARM Mali

Old: 400	Now: T604, T628	Soon: T720, T760

Imagination Technologies

Old: SGX 5xx	Now: Series 6	Soon: Series 6XT

NVIDIA Tegra

Old: Tegra 3, 4	Now: K1	Soon: ...

Tile-based Mobile GPU

- Mobile GPUs are usually tile-based (next-gen too)

  Tile-based: ImgTec, Qualcomm*, ARM

  Direct: NVIDIA, Intel, Qualcomm*, Vivante

* Qualcomm Adreno can render either tile-based or direct to frame buffer

  – Extension: GL_QCOM_binning_control

Tile-Based Mobile GPU

Summary:

- Split the screen into tiles
  - E.g. 32x32 pixels (ImgTec) or 300x300 (Qualcomm)
- The whole tile fits within GPU, on chip
- Process all drawcalls for one tile
  - Write out final tile results to RAM
- Repeat for each tile to fill the image in RAM

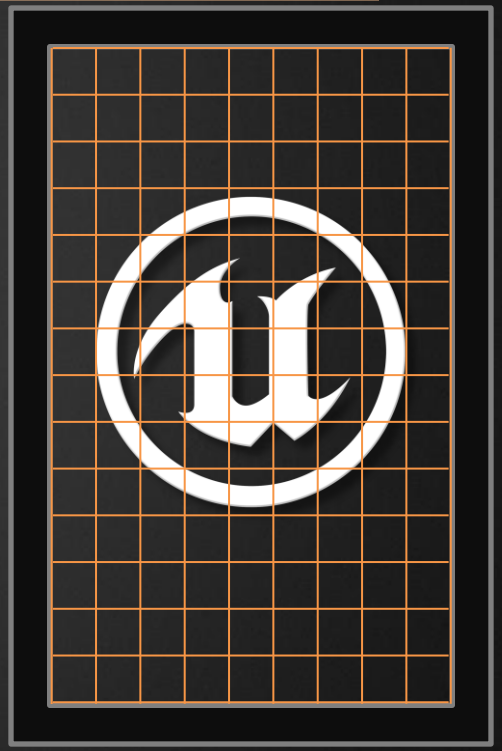

ImgTec Tile-based Rendering Process

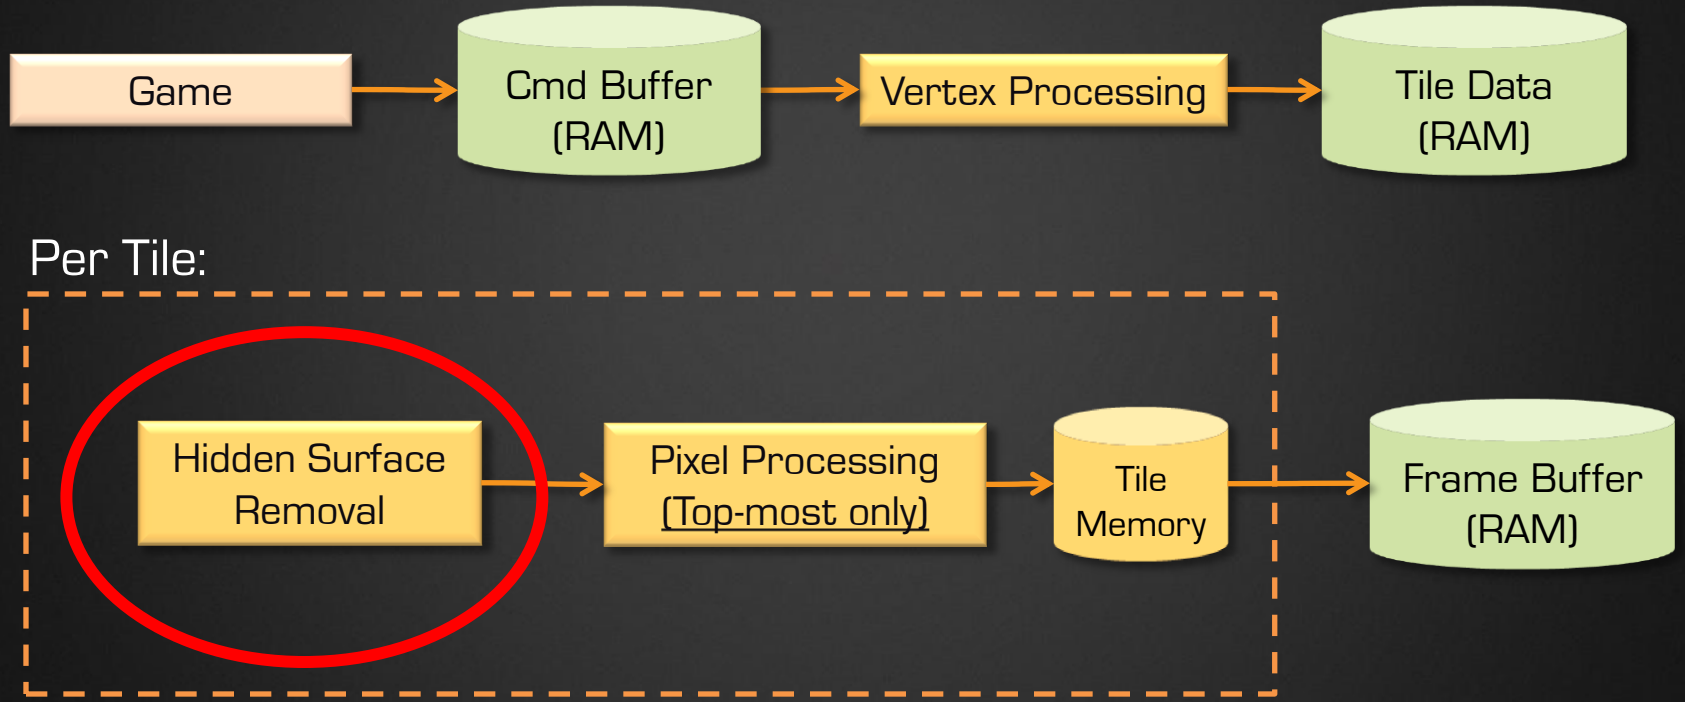

Framebuffer Resolve/Restore

- Expensive to switch Frame Buffer Object on Tile-based GPUs
  - Saves the current FBO to RAM
  - Reloads the new FBO from RAM
- Best performance:
  - A single rendertarget for the entire frame
  - No post-processing passes
- Does not apply to NVIDIA Tegra GPUs!
  - This made it simpler for us to use our full desktop rendering pipeline on K1
  - *"Rivalry"* tech demo (showing 5:00 pm today)

Thermal Limits

- Hardware CPU and GPU clock frequencies change all the time!
  - Many times per milli-second!
  - To save battery
  - To prevent overheating
- Qualcomm Trepn Profiler
  - https://developer.qualcomm.com/mobile-development/increase-app-performance/trepn-profiler

Thermal Limits

- Check your performance when device is cool
- Check again when it's hot
- CPU uses much more power and heat than the GPU
  - Also, memory bandwidth generates a lot of heat
- Avoid unnecessary CPU usage
  - Spin-loops
  - Frequently waking up threads just to put them to sleep again

Performance Guidelines

- Always make sure lighting has been built before looking at performance
- Use as little post-process effects as you can get away with
- Make sure precomputed visibility has been set up properly
- Minimize overdraw (translucent or masked materials)
- Target 100-700 draw calls per frame
- Use as few texture lookups as possible in your materials
- Documentation:
  - https://docs.unrealengine.com/latest/INT/Platforms/Mobile/Performance/index.html

Performance Tier 1 – 2

1. LDR (Low Dynamic Range)
   - Fastest mode
   - Use when you don't need lighting or post-process effects
   - Disable "Mobile HDR" in Rendering section in your Project Settings

2. Basic Lighting
   - Allows HDR lighting and some post-process effects
   - Use only static lights
   - Use only fully rough materials, not shiny (specular)
   - Disable Bloom and anti-aliasing

Performance Tier 3 – 4

3. Full HDR Lighting
   - High-quality lighting with best support for normal maps
   - Realistic specular reflections on surfaces with per-pixel roughness
   - Use only static lights
   - Bloom and anti-aliasing are recommended
   - Place reflection captures carefully for best results

4. Full HDR Lighting with per-pixel lighting from the Sun
   - Specify one directional light as stationary (the Sun)
   - All other lights are static
   - High-quality distance field shadows

Interlude: End of Part 1

Questions?

Keep going?

Coffee break?

Ready for more?

Part 2: Adapt and Conquer

- Very difficult to scale on CPU performance
  - Gameplay features can't easy be switched off
  - Also, CPUs aren't as different as GPUs are
  - Make sure you are never gamethread-bound on any device
- Scale your game purely based on GPU performance
  - Primarily resolution and post-process effects
  - Ship it!

Cross-platform Console Commands

- Common commands:
  - Stat Unit
  - Stat UnitGraph
  - Stat FPS
  - Stat SceneRendering
  - Stat Slow
  - ViewMode ShaderComplexity
- Documentation:
  - https://docs.unrealengine.com/latest/INT/Engine/Rendering/PerformanceProfiling/StatCommands/index.html

Console Command: Stat Unit

- Always the first step when checking performance

Console Command: Stat SceneRendering

- Shows Renderthread CPU performance and drawcalls

Console Commmand: ViewMode ShaderComplexity

- Visualize expensive materials in the PC ES2 previewer
- Shows approximate performance cost per material
- Green is good, red is bad. Pink or white is extremely expensive!

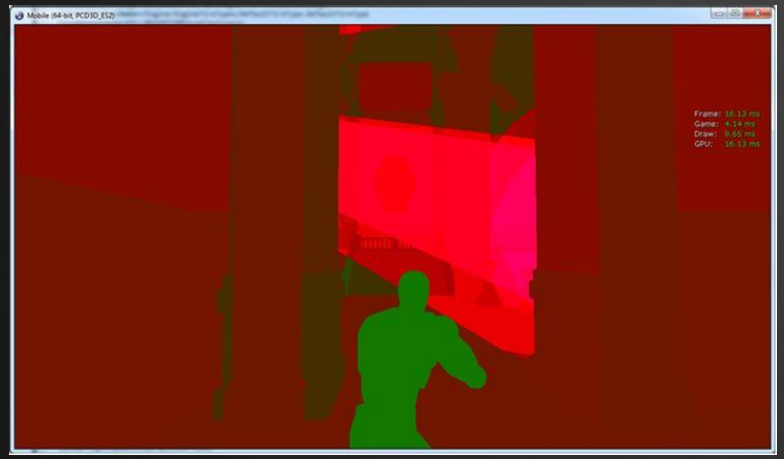

iOS Performance

- New Metal graphics API in iOS 8
  - Much faster on CPU
  - Up to 20x faster on renderthread
  - Allows for thousands of drawcalls on iOS devices with A7 processors
- Scale graphics quality based on exact device model
  - Still very few different device models, easy to target each one specifically
  - Resolution (MobileContentScaleFactor)
  - Post-process features
  - Etc...

Platform-Specific Profiling

- Each GPU family has their own profiling tools
  - Apple: Xcode GL Debugger (and Metal)
  - Qualcomm: Adreno Profiler
  - NVIDIA: Tegra Graphics Debugger
  - ImgTec: PVRTune, PVRTrace
  - ARM: Mali Graphics Debugger

- For CPU profiling
  - Apple: Instruments (Time Profiler)
  - NVIDIA: Tegra System Profiler
  - ARM: DS-5

UE4_FromPC.xcodeproj — Soul.gputrace

Soul_RunIOS › Epic_3335

Running Soul on Epic_3335

Soul
Captured OpenGL ES Frame

- Context 1 | RenderingThread 2
- Frame5405
  - InitViews
  - 19 Clear(ColorBuffer | StencilBuffer |...
  - BasePass
  - Translucency
  - FinishRendering
    - PostProcessBloomSetup
    - PostProcessBloomDown
    - PostProcessBloomDown
    - PostProcessBloomDown
    - PostProcessBloomDown
    - PostProcessBloomUp
    - PostProcessBloomUp
    - PostProcessBloomUp
      - 5438 Clear(ColorBuffer)
      - 5446 DrawElements(Trian...
    - PostProcessSunMerge
      - 5450 Clear(ColorBuffer)
      - 5459 DrawElements(Trian...
    - PostProcessSunAvg
      - 5463 Clear(ColorBuffer)
      - 5469 DrawElements(Trian...
    - PostProcessTonemap
      - 5473 Clear(ColorBuffer)
      - 5483 DrawElements(Trian...
    - PostProcessAa
      - 5487 Clear(ColorBuffer)
      - 5495 DrawElements(Trian...
  - 5513 DrawElements(Triangles,...
  - 5523 DrawElements(Triangles,...
  - 5528 DrawElements(Triangles,...
  - SlateUI

Soul.gputrace › Frame5405 › FinishRendering › PostProcessAa › 5495 DrawElements(Triangles, 3, Unsigned Short, <data>)

Color Attachment 0

P... › Fragment Shader

```
//*****************************************************************
// Edits to captured shaders are not saved to the sou
// Copy them back to your shader files when done or r
// the GPU Debug Log.
//*****************************************************************
#version 100
precision mediump float;
precision mediump int;
precision mediump sampler2D;
precision mediump samplerCube;

uniform highp vec4 pu_h[1];
uniform sampler2D ps1;
uniform sampler2D ps0;
varying highp vec2 var_TEXCOORD0;
varying highp vec2 var_TEXCOORD1;
varying highp vec2 var_TEXCOORD2;
varying highp vec2 var_TEXCOORD3;
varying highp vec2 var_TEXCOORD4;
varying highp vec2 var_TEXCOORD5;
void main()                                          20.5%
{
    mediump vec4 t0;
    vec4 t1;
    t1.xyzw = texture2D(ps0,var_TEXCOORD0);
    vec4 t2;
    t2.xyzw = texture2D(ps0,var_TEXCOORD1);
    vec4 t3;
    t3.xyzw = texture2D(ps0,var_TEXCOORD2);
    vec4 t4;
    t4.xyzw = texture2D(ps0,var_TEXCOORD3);
    vec4 t5;
    t5.xyzw = texture2D(ps0,var_TEXCOORD4);
    vec4 t6;
    t6.xyzw = texture2D(ps1,var_TEXCOORD5);
    highp float t7;
    t7 = max((clamp((min(t3.w,t6.w)*4.0),0.          28.2%
    t0.xyz = mix(t3.xyz,t6.xyz,vec                   51.3%
    gl_FragColor.xyzw = t0;
}
```

Soul › F... › F... › PostProcessAa › 5495 DrawElements(Triangles, 3, Unsigned Short, <data>)

- **Program #310 Performance** 1.73 ms (4.8%)
  - **Program Duration** = 1.73 ms (4.8%)
  - **Draw #5495 Duration** = 1.73 ms (4.8%)
  - **Fragment Shading** 1.73 ms (4.8%)
    - **Shader Duration = 1.73 ms (4.8%)**
  - **Vertex Shading** 0.00 ns (0.0%)
- **Program #310**
- **Vertex Array Object #0**
- **Texture Unit 0** 2D:971"Tonemap"
- **Texture Unit 1** 2D:969"Tonemap"

Auto

- **Current Program** = Program #310
- **Draw Framebuffer** = Framebuffer #2
- **Read Framebuffer** = Framebuffer #2
- **Renderbuffer** = Renderbuffer #1 "OnScreenColorRB"
- **Array Buffer** = 0
- **Vertex Array** = Vertex Array Object #0
- **Texture Unit 0** 2D:971"Tonemap"
- **Texture Unit 1** 2D:969"Tonemap"

Bound GL Objects

Qualcomm Adreno Profiler

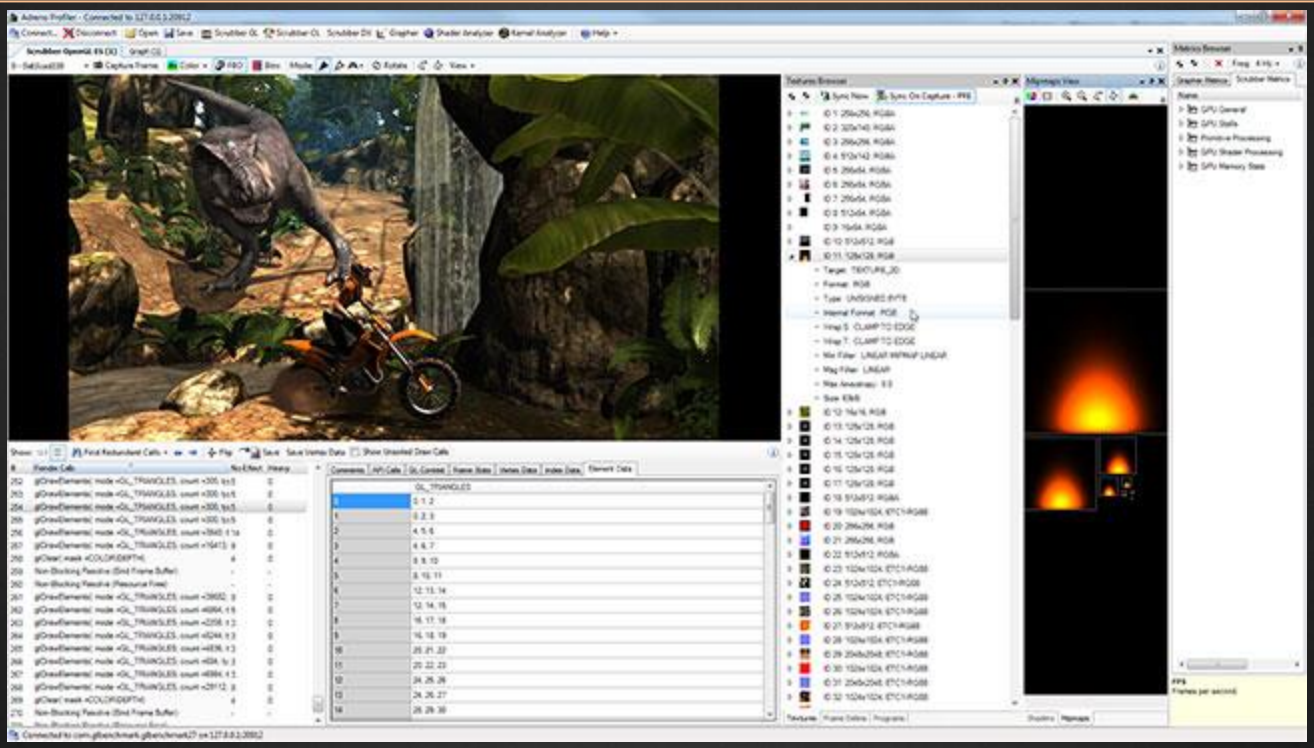

NVIDIA Tegra Graphics Debugger

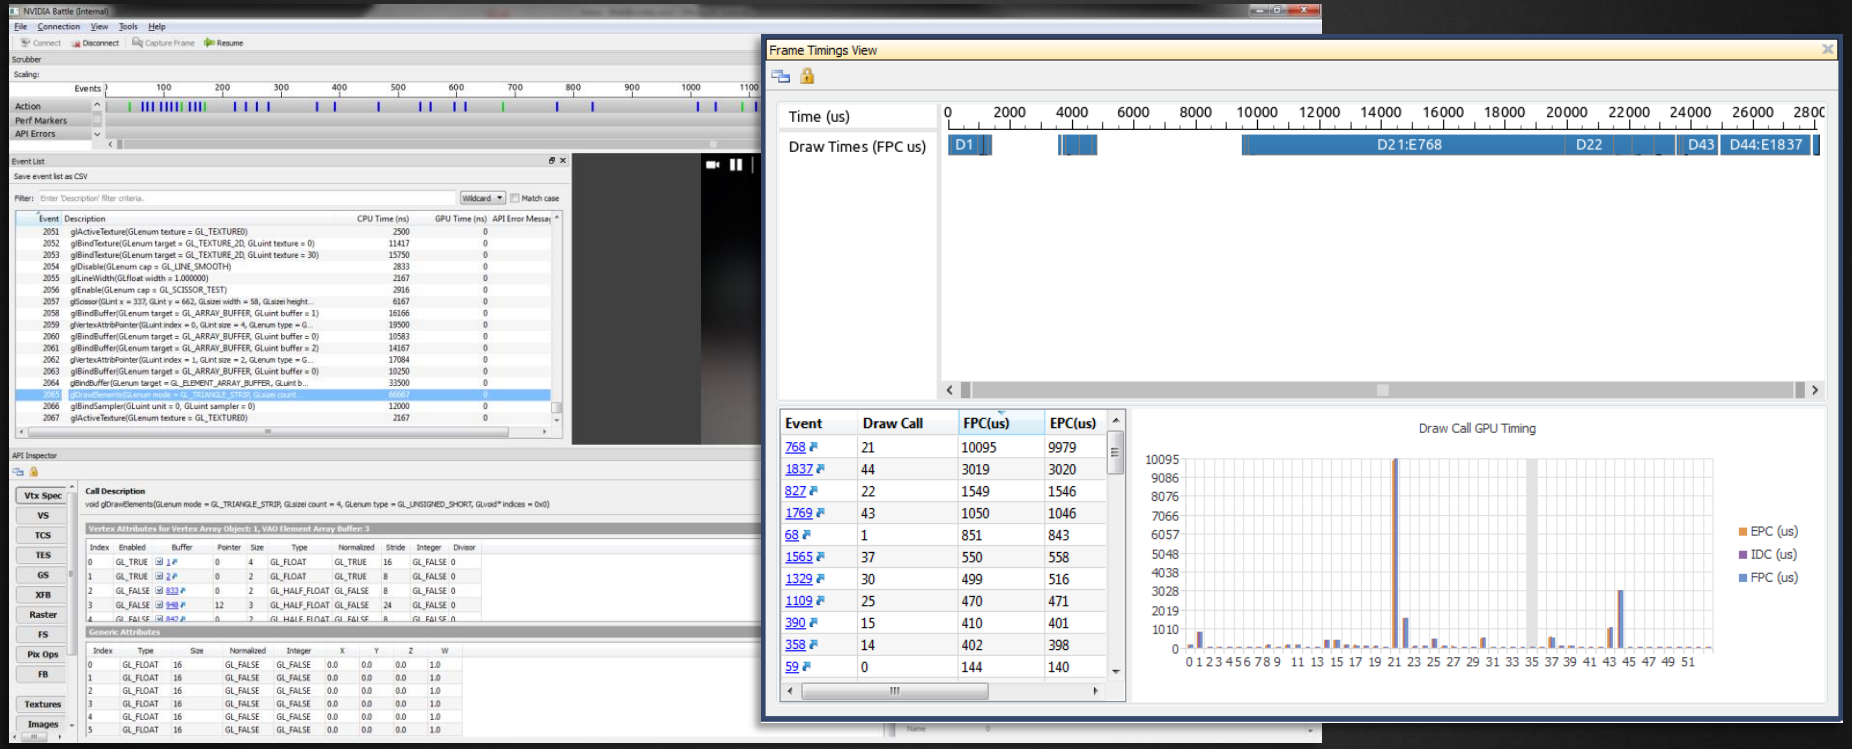

ImgTec PVRTune and PVRTrace

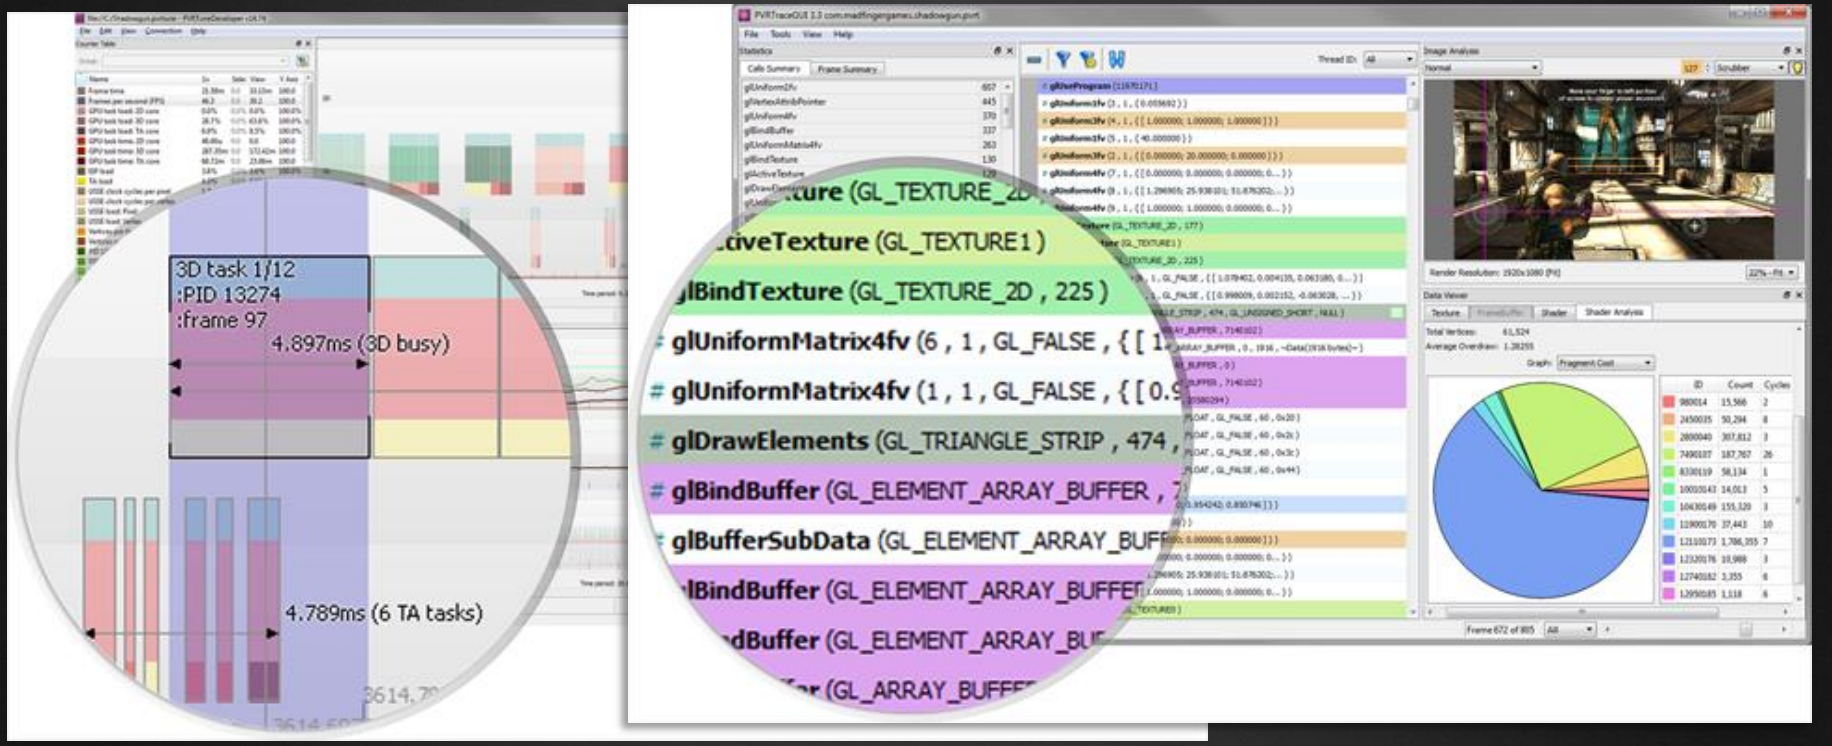

ARM Mali Graphics Debugger

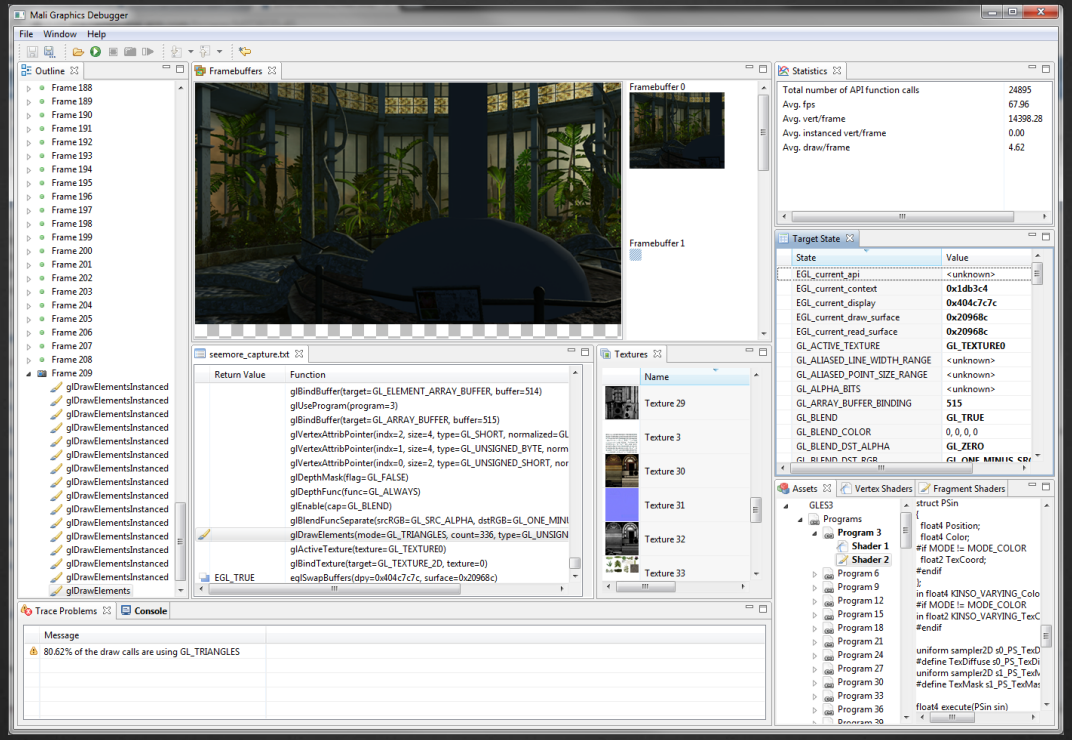

Device Profiles

- UE4 selects one device profile at startup
  - Detects device model and capabilities
- Tweak each device profile for your game
  - Config/DefaultDeviceProfiles.ini
  - Each Device Profile can customize engine features, like:
    - +CVars=r.MobileContentScaleFactor=2
    - +CVars=r.BloomQuality=1
    - +CVars=r.DepthOfFieldQuality=1
    - +CVars=r.LightShaftQuality=1
- Documentation:
  - https://docs.unrealengine.com/latest/INT/Platforms/DeviceProfiles/index.html

UE4 Mobile Performance Questions?

Documentation, Tutorials and Help at:

- AnswerHub: http://answers.unrealengine.com
- Engine Documentation: http://docs.unrealengine.com
- Official Forums: http://forums.unrealengine.com
- Community Wiki: http://wiki.unrealengine.com
- YouTube Videos: http://www.youtube.com/user/UnrealDevelopmentKit
- Community IRC: #unrealengine on FreeNode

Unreal Engine 4 Roadmap

- lmgtfy.com/?q=Unreal+engine+Trello+