HDR Ecosystem for Games

Evan Hart, Principal Engineer

March 23, 2018

NVIDIA

Booth #223 - South Hall

www.nvidia.com/GDC

What is HDR and why do we care?

- HDR has existed in games since ~2005
  - Half Life 2 : Lost Coast
- High percentage of titles render in HDR
  - Necessary for proper simulation of lighting interactions
- Rendering HDR is pretty well understood
- Displaying HDR is still in it’s infancy
- HDR displays are one of the biggest advances in the past 20 years

What Makes Good HDR?

- HDR is High Dynamic **Range**
- Brighter doesn’t give you HDR
- Ultimately requires, brightness, contrast, and precision
- Good HDR will
  - Make highlights brighter
  - Maintain or improve darkness of shadows
  - Preserve more detail at both the top and bottom of the range
  - Allow more vivid, and potentially brighter mid-tones
- Ultimately, the experience is about a deeper image

Why haven’t we been doing this?

Display Technology Somewhat Stagnant for ~20 years

- CRT -> LCD
  - Lighter/thinner
  - Many visual qualities didn’t improve
  - Potentially brighter
- Overall progress has been slow
  - sRGB & Rec 709 are the standards
  - Resolution has been no better
- HDR looks to change most of this

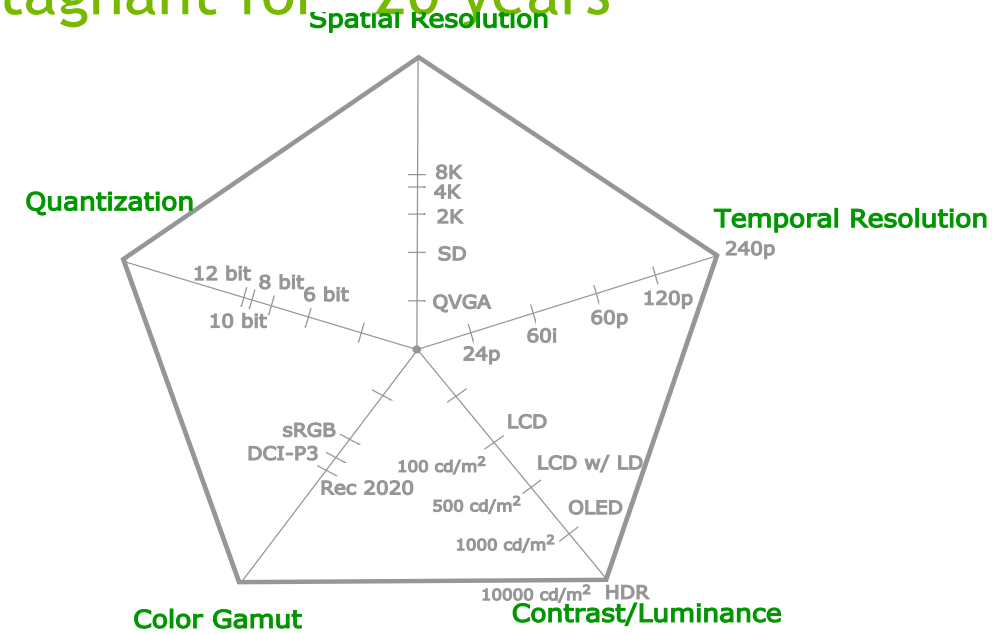

The Technology Behind HDR

HDR Display Hardware

- Understanding the strengths and limits of the devices is important
- Display devices and their interfaces have physical limitations
  - Transfer bandwidth of HDMI and DisplayPort
  - Power consumption and heat dissipation
- Knowing the HW helps you understand the results
  - Also helps you pick the right HW for your development purposes

LCD vs OLED

- Over 1000 nit peak luminance
- Local dimming matrix is best
- Wide color gamut via quantum dots
- Can suffer backlight bleed
- Well suited for bright environment
- Well suited for monitors today

- Typically limited to ~600 nits
- All pixels are independent
- Good color saturation
  - Can suffer at peak luminance – RGBW
- Excellent black levels
- Best suited for dark environment
- More suited to TVs
- Has some phosphor lifetime concerns

TV vs Monitor

- Designed for a living room/den
- Designed for dim viewing environment
- Designed for limited duty cycle
  - Couple hours a day
- Typically lots of image processing
  - Lag, distortion
- HDMI only
  - Limited refresh rate
  - Maybe only YCbCr
- TVs often obfuscate information

- Designed for an office
- Handles brighter environment
- Low latency experience
- Display Port for more bandwidth
  - Higher refresh rates
- Minimizes image processing
  - Faithful display of content

G-SYNC HDR

NVIDIA® G-SYNC™

4K 120 HZ

HDR - 1000 NITS

ULTRA-LOW LATENCY

DCI-P3 COLOR GAMUT WITH QUANTUM DOTS

NVIDIA® SHIELD™ BUILT-IN

HDR Display Standards

- Large number of overlapping standards in the space
- Many are related to or derivations of others
- Lots of organizations involved
  - VESA
  - SMPTE
  - International Telecommunications Union
  - Consumer Technology Association
  - HDMI

HDR10

- Common term for the popular HDR encoding standards tied to SMPTE/others
  - Original term referred to a ‘media profile’
  - Today, BT 2100 might be the better term for the generic usage

- For PCs it essentially means
  - SMPTE 2084 transfer function
  - BT 2020 color primaries
  - Generally 10 or 12 bits per component
  - RGB, YCbCr 4:2:2, and/or YCbCr 4:2:0 wire transfer
    - These are invisible to the developer

What does this mean?

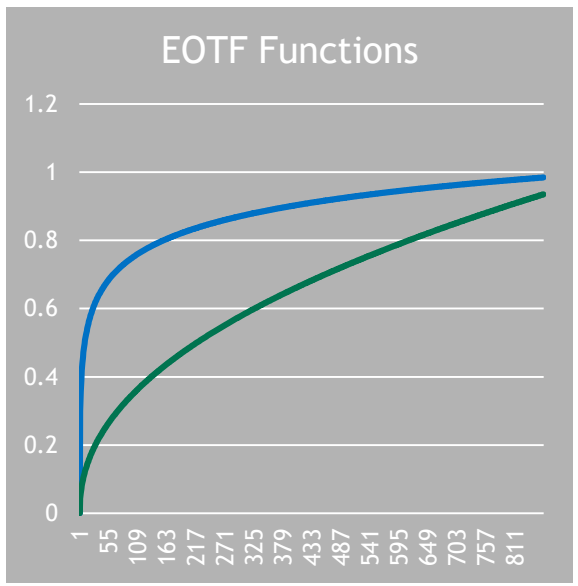

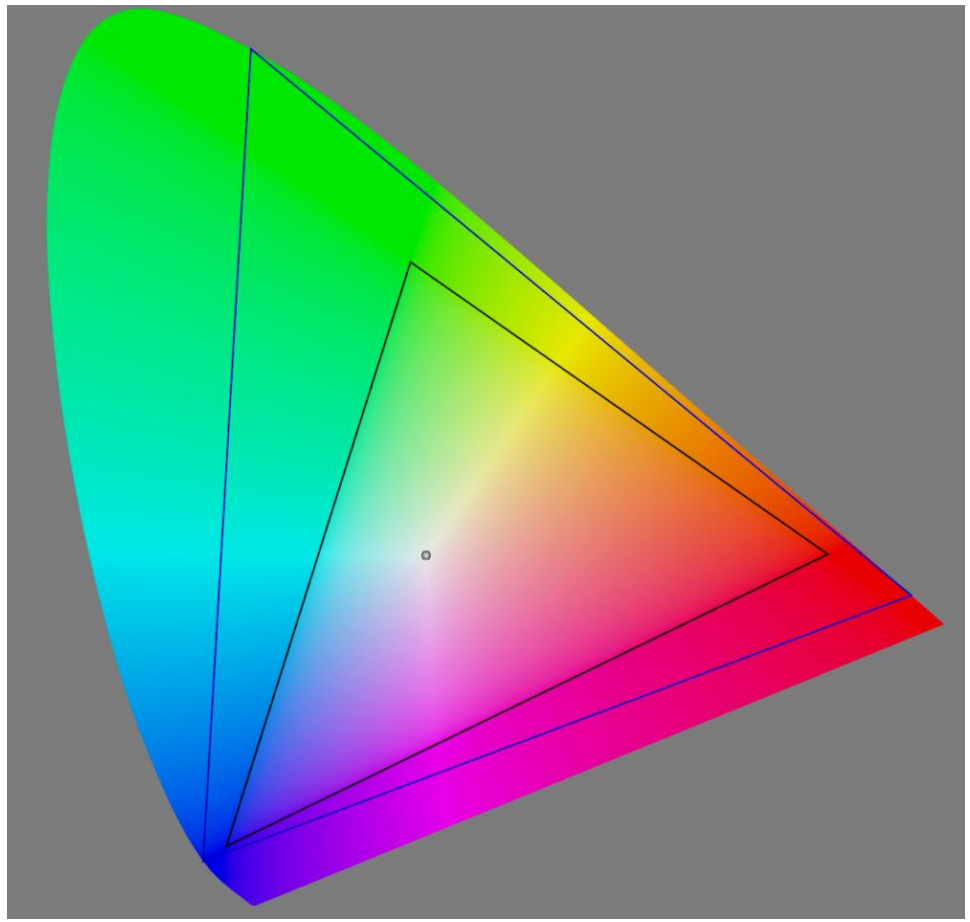

How else does it differ?

- HDR10 signal is a container format
  - Game is stating explicitly how it wants something displayed
    - Absolute nit level
    - Precise color
  - Not expected to fill the 10,000 nit + BT 2020 color volume
    - No practical display can handle this today
  - Generally, the target should be roughly P3 & ~1000 nits
- HDR10 operates best by knowing the portion of the container the game uses
  - This information is provided as metadata

Metadata

- Virtually all HDR standards use metadata to express the intent of the signal
- Two essential types
  - Static and dynamic
- Static rarely changes, just information about the stream in general
- Dynamic typically constantly changes, likely frame to frame
  - Provides more information
  - Often gets referred to as HDR10+

SMPTE 2086/CEA Static Metadata

- Reference primaries
- Reference white point
- Reference peak luminance
- Reference min luminance
- Max Content Light Level
  - Most luminous color channel for any frame
- Max Frame Average Light Level
  - Average luminance per pixel for the most luminous fame

Dynamic Metadata

- Provides more information about what is being rendered
- Allows a display to adjust how it handles things on a continuous basis
- Primary advantage is to make a lesser display produce it’s best image
  - 400 or 600 nit display attempting to display a 1000 nit signal
  - Many scenes are dim, shouldn’t adjust
  - Need to adjust bright scenes

- SMPTE 2094 defines 4(!) different flavors
  - Makes things a bit of a mess for developers today
  - Windows supports none of these

Why Games care less about dynamic

- Dynamic metadata is most useful for pre-baked content
  - Movies need to select a target max luminance when encoded
- Games are fully dynamic content
  - Post processing can use a target maximum luminance at runtime
  - Games can’t have the level of color grading video may have
    - User can do anything
- No need to hope for the display doing ‘The Right Thing’
  - Choose a level appropriate for the display and do it
  - Set the metadata to represent what you chose

Dolby Vision

- Dolby Vision is a proprietary technology/brand
- Much of the technology is strongly related to other standards
- Stream uses a special encoding
- Technology supports dynamic metadata
- Special signal encoding in the PC ecosystem is problematic

- Need to be committed to testing for the long term
- Almost every device supporting Dolby Vision supports HDR10

Programming Interface

Win 10 Creators Update and Beyond

- Introduces native HDR on the desktop
  - Enables HDR in a window
  - May need to enable it in the displays control panel
- Natively supported via DXGI
  - Extended output desc queries providing display information
  - SetColorspace1 on swap chain to control interpretation
  - Metadata set function
- Requires ‘Flip’ mode swap chain
  - Non-flip will work full-screen, but truncates if composited

Interpreting DXGI Colorspaces

- DXGI Color Space enum is overloaded for multiple purposes
  - What the display is receiving
  - What the swap chain is operating in
- Enum essentially has 5 fields
  - Color space type (RGB or YCBCR)
  - Transfer/Gamma function (G22, G10, G2084)
  - Encoding range (Full or Studio)
  - Siting or where chroma subsampling is aligned (only video relevant)
  - Color primaries (P709 or P2020)

Detecting HDR

- Obtain IDXGIOutput6
- Use GetDesc1
- HDR color space will be
  - RGB_FULL_G2084_NONE_P2020
- SDR color space will be
  - RGB_FULL_G22_NONE_P709

```
IDXGIOutput6 *output6 = nullptr;

output->QueryInterface(&output6)

DXGI_OUTPUT_DESC1 oDesc;

output6->GetDesc1(&oDesc);

supportsHDR = oDesc.ColorSpace ==
DXGI_COLOR_SPACE_RGB_FULL_G2084_NONE_P2020;
```

Additional New Display Information

- Red, green, and blue primaries
  - xy chromaticity coordinates for the display's primaries
- Minimum and maximum luminance levels
- Maximum full-frame luminance
  - Highest level that the display will produce for all pixels at once
- White point for the display

DXGI Color Spaces and Formats

DXGI Color Space	Format	Default	HDR
RGB_FULL_G22_NONE_P709	R8G8B8A8_UNORM	Yes	No
	R10G10B10A2_UNORM	Yes	No
RGB_FULL_G10_NONE_P709	R16G16B16A16_FLOAT	Yes	Yes
RGB_FULL_G2084_NONE_P2020	R8G8B8A8_UNORM	No	Yes*
	R10G10B10A2_UNORM	No	Yes*

Setting the Color Space

- IDXGISwapChain3 provides methods for color space management
- CheckColorSpaceSupport
  - Allows the app to determine if a particular color space is supported
  - Have to enumerate the list you care about and enquire
- SetColorSpace1
  - Applies the requested color space to the swap chain

Which format to use?

- Best advice is DXGI_FORMAT_R16G16B16A16_FLOAT
- Uses DXGI_COLOR_SPACE_RGB_FULL_G10_NONE_P709
- Works in HDR by default
- This color space is known as scRGB or CCCS
  - It is what the OS uses for compositing
  - Luminance is 80 nits for (1,1,1), so 1000 nits white is 12.5
  - Color primaries are Rec 709 / sRGB
  - Wide gamut supported via negative colors

Why not RGB10?

- Unfortunately doesn’t deliver on all its promises
  - Often touted as “conversion free”
    - This is not reliably true
    - May end up with a double conversion
  - Memory savings aren’t that fantastic
    - < 3% in a really aggressive case
- Sacrifices precision
  - Outputs can be 12 bits
  - May suffer requantization
- Non-default status introduces compatibility challenges

Setting the MetaData

- IDXGISwapChain4::SetHDRMetaData
  - Accepts enum, void*, and struct size
  - Today only choices are None and HDR10
- HDR10 Metadata structure has somewhat odd scaling
  - All colors are specified in xy scaled by 50,000 to produce an integer
  - Min/Max are scaled by 10,000

Earlier Versions

- NVAPI supports HDR displays back to Win7
- Restricted to full-screen exclusive mode
  - Requires driver to have full control of the display of the swap chain
  - D3D12 lacks true full-screen exclusive

- Importantly, NVAPI is still useful in latest Win10
  - Can detect HDR screens present, but not in HDR mode
  - Allows game to enable HDR and have it deactivate on exit

Making Your App

Overall Goals

- Produce a good HDR scene-referred image
  - Properly exposed image
  - Intense highlights and deep shadows where applicable
  - Not all scenes are HDR
- Preserve precision through post-processing
- Perform color grading in the expanded high-precision space
- Tone map into the monitor's capabilities
- Apply a legible UI
- Build in tools for analysis

Physically Based Rendering

- Important for generating good, plausible values
  - Will naturally generate good, detailed highlights
- Requires the artists to set things up properly
- Naturally generates the range you want

- HDR is an excellent detector for cheating in the art work
  - Emissive geometry mismatched with light values
  - Particle systems and effects can be calibrated incorrectly

Common Fault

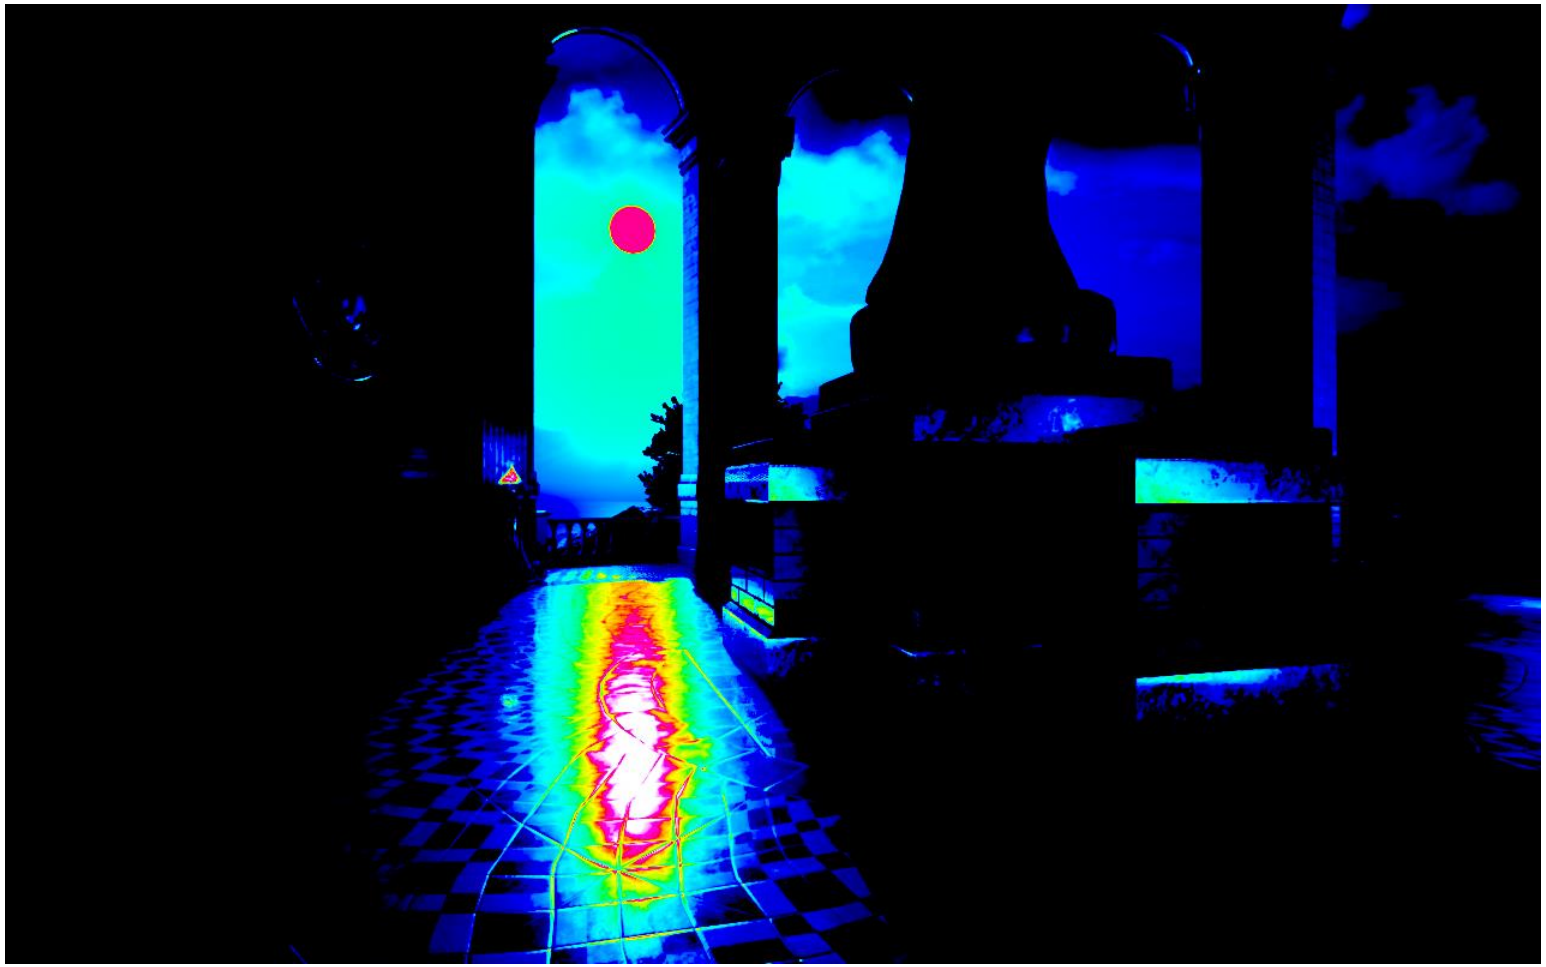

Rendering Primaries

- All color rendering is done within a specific subspace
- Rarely thought about today
  - sRGB/Rec 709 are the standard
- Altering the primaries has serious ramifications on rendering results

Best Practices on Primaries

- Use a consistent set
  - Always use the same set no matter the destination primaries
- Wider primaries would be nice
  - Allows the production of a wider gamut
  - Appears to produce more correct color interactions
  - Will generate colors outside Rec 709/sRGB
- Conservative choice: Stick to Rec 709
  - Guarantees compatibility with SDR
- Consider gamut mapping to enable better use of the wider gamut

Classic Post Processing Pipeline

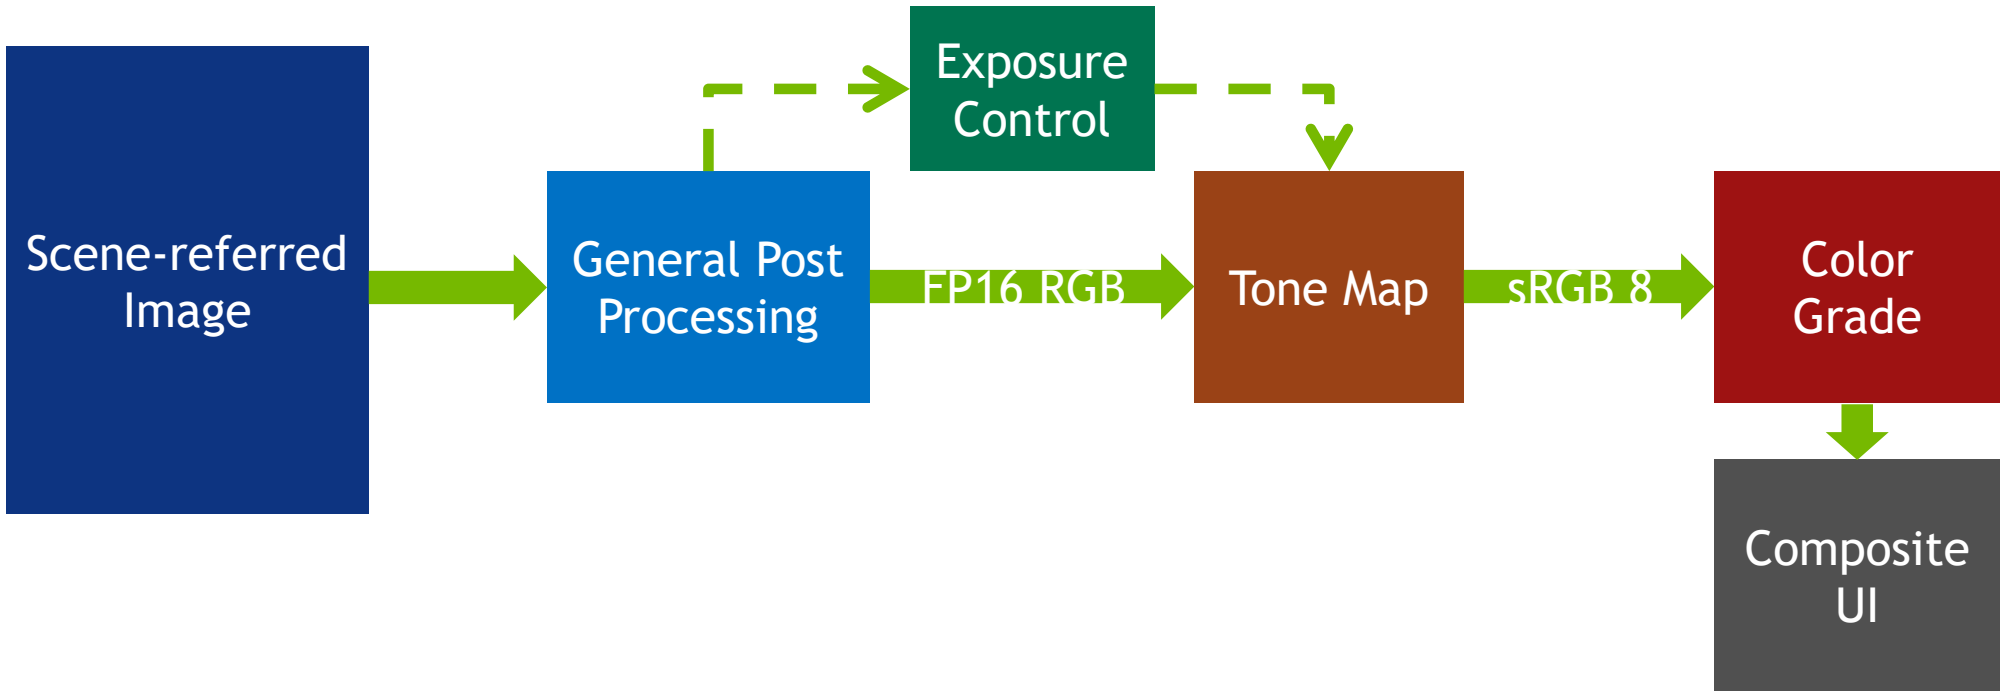

HDR-Aware Post Processing Pipeline

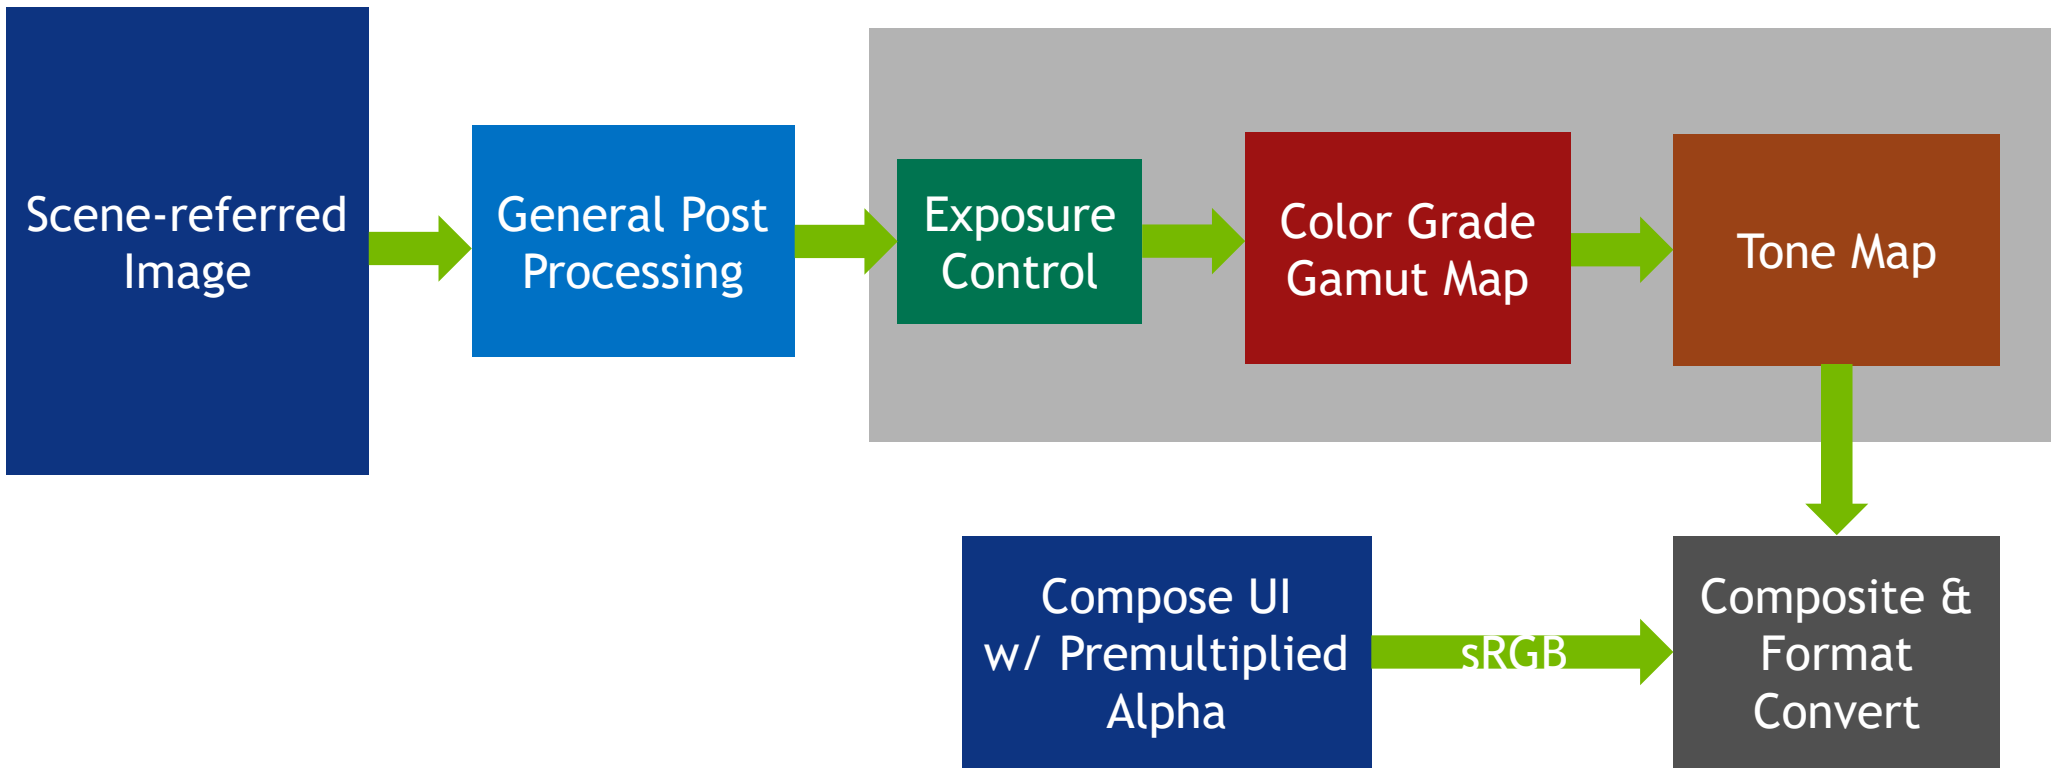

Color Grading

- One of the tougher challenges we face with HDR
- Many games today perform image referred color grading
  - Edit image in Photoshop and extract LUT
  - Often even performed in gamma space
  - This is operating in the wrong space/range for HDR
- Ideally move grading to scene-referred space before tone map
  - Now grade applies to all targets
  - Can be either pure math or LUT baked in the appropriate space

Gamut Mapping

- Gamut mapping is the process of adapting one color space to fit in another
- Allows consistent rendering primaries, with
- Simplest forms are linear stretch/squash or clip
  - Results in undesirable hue shifts
  - Colors can get too intense
  - Skin tones can get overly red
- Three solutions
  - Don’t expand, keep colors within sRGB gamut
  - Use soft expansion (desaturated colors are preserved)
  - Artist generated

Tone Mapping

- HDR displays do not remove the need for tone mapping
  - Simulations can produce 100,000 or 1,000,000 nits
  - HDR standards restrict to 10,000 nits
  - Practical displays are ~1,000 nits
- Tone map operator is still required
  - One that is adapted to higher luminance

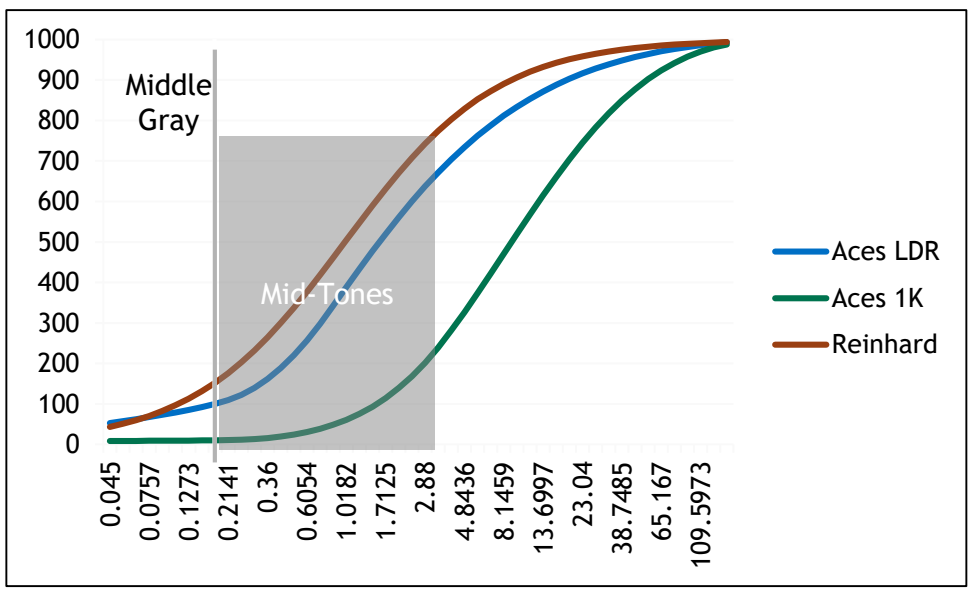

Preserving Detail

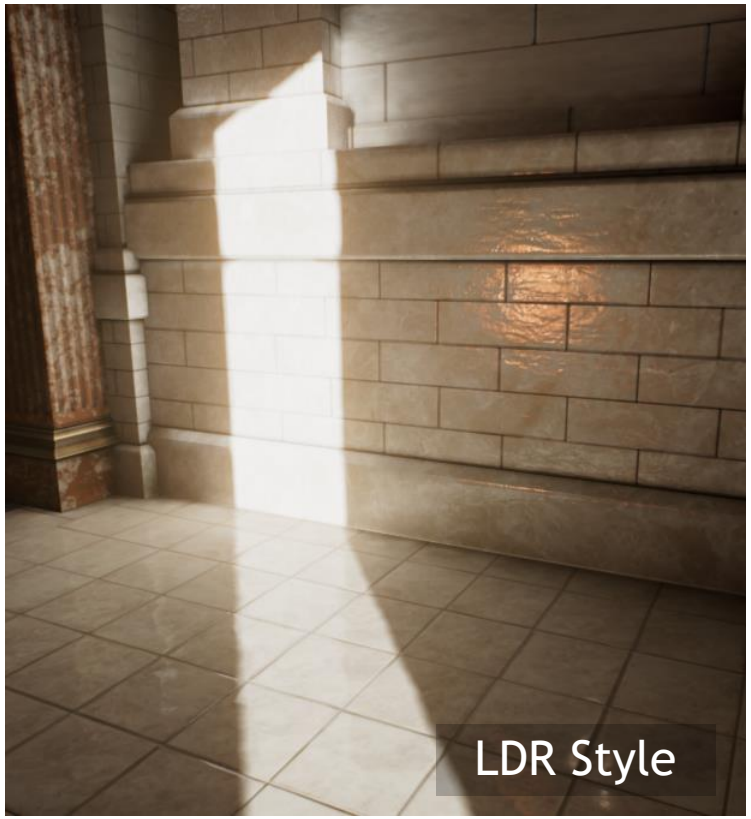

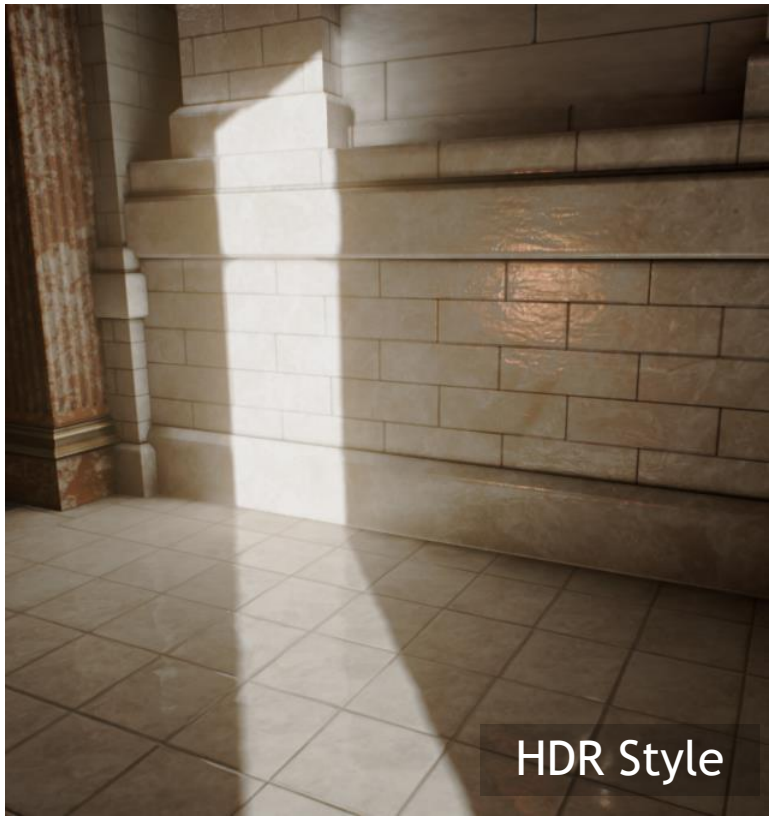

Choice of Tone Map Operator

- Ultimately, there are many choices
- Personally, local operators seem less relevant
  - Local operators are attempting to make up for what SDR lacks
- Filmic operators generally a recommended choice
  - 100+ years of science and art behind the physical basis
  - Sigmoid-like curve is similar to cone response curves
  - Natural desaturation at extremes fits well with human experience
- ACES is a good choice and has adaptable luminance ranges

ACES Pipeline

- Film-production driven pipeline for managing color
  - Includes color grade
  - Includes concept of output device transform including different display targets
  - Output transform includes tone mappers for different luminance levels
- Provides a great foundation to build from
  - A game could do much worse than using an implementation of ACES
  - Many optimizations such as baking to a LUT are possible
- Not the only solution
  - It is at least a great place to start
  - Why reinvent the wheel?

User Interface

- UI requires special care in HDR
- Traditional UI development has always been image referred
- Typically, painter's algorithm straight onto the framebuffer
- Often at least some alpha blending, and possible transparency

- For HDR
  - Need to set appropriate luminance level
  - Need to possibly adjust contrast

UI Transparency

Glow through is problem for normal transparency

LDR

HDR

A Solution for UI Composition

- Render UI offscreen as sRGB
  - Pre-multiplied alpha is probably most convenient
  - Ensure to preserve a meaningful alpha channel
- Run fullscreen pass to composite UI inside a shader
  - Detect over operations and apply simple tone map to dim HDR
    - x/(x+1) for luminance only
  - HDR should be remapped in a color space with no negative values
  - Apply scale and contrast to UI values
- Optionally, provide an additional curve on the alpha
  - Useful for fade-in alpha effects

Build in Good Tools

- Add visualizations for pre/post tone mapped range
- Allow artists to visualize things
- Makes debugging easy
- HDR helps you fix things you never quite knew were wrong
- SDR will improve too

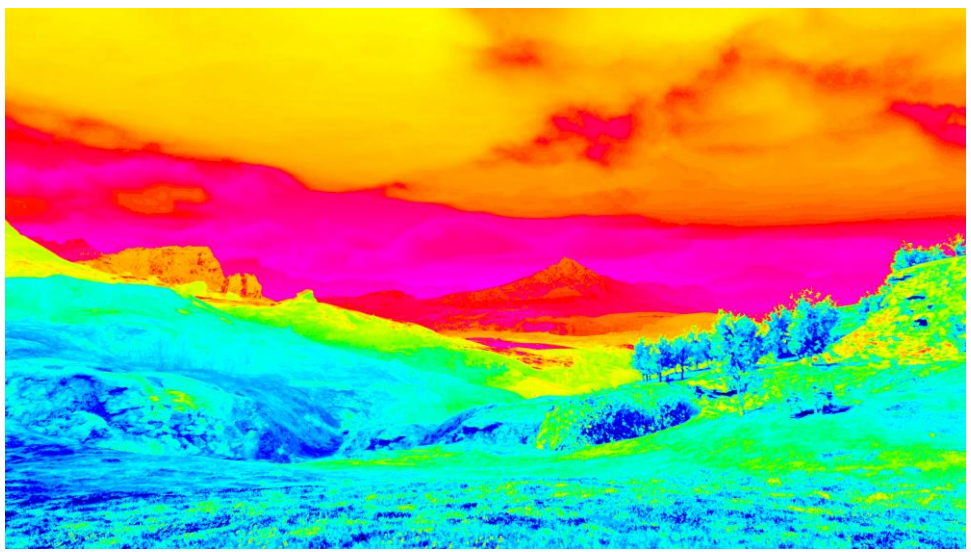

Thanks

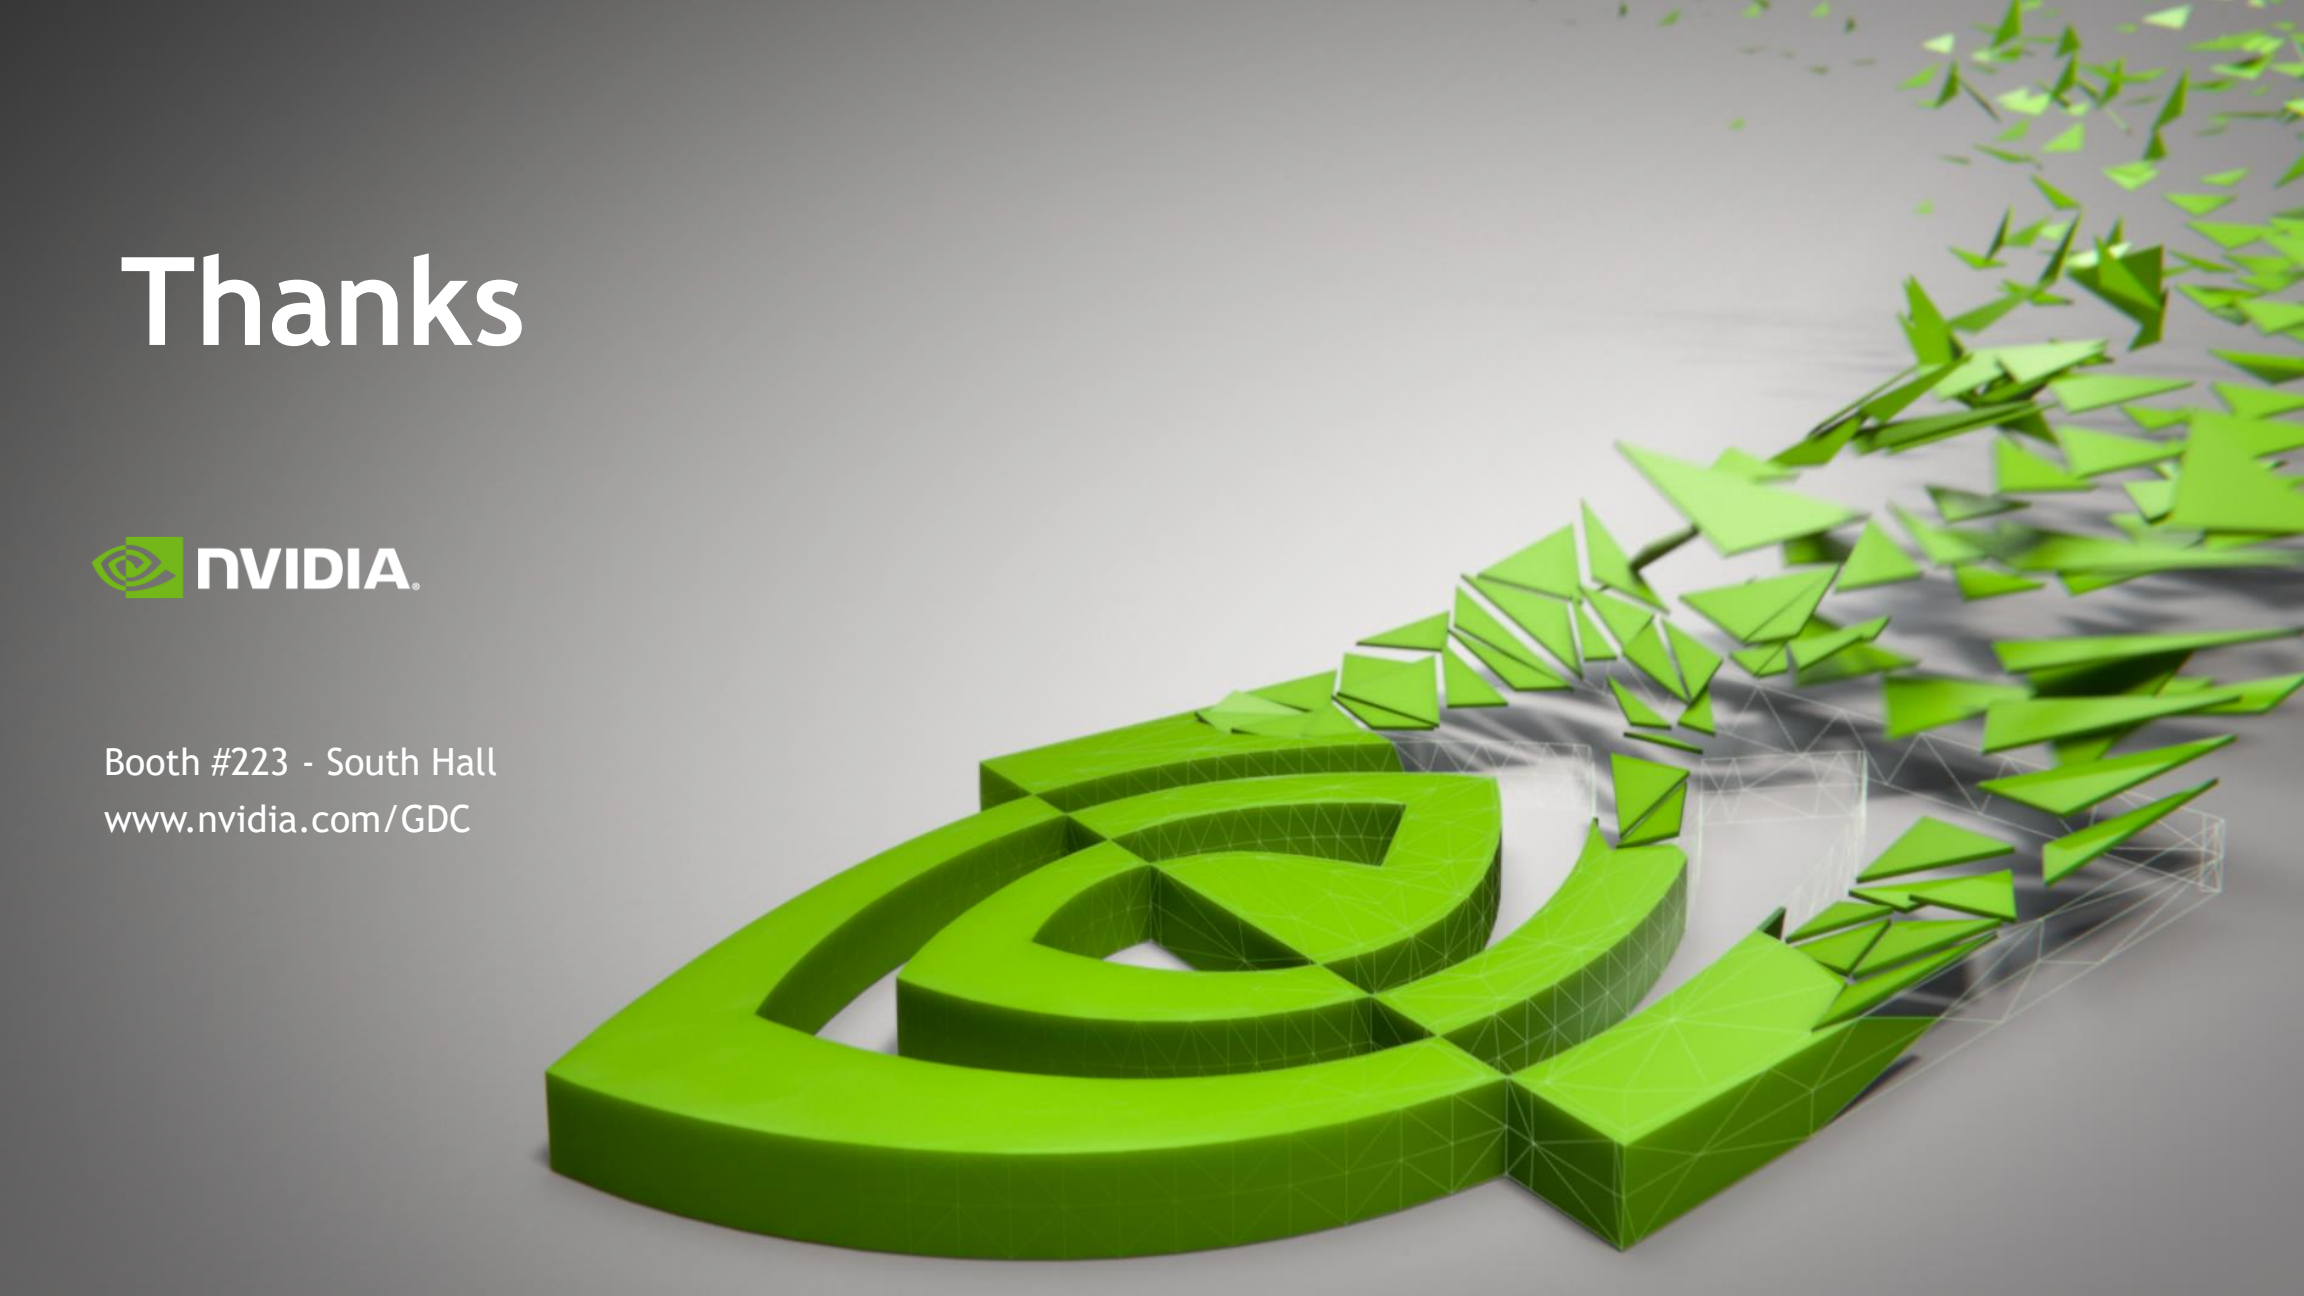

Booth #223 - South Hall
www.nvidia.com/GDC