(n)K RESOLUTION

THE IMMINENT NEED FOR INFINITE PIXELS

(n)K RESOLUTION: THE IMMINENT NEED FOR INFINITE PIXELS

The history of modern computing is virtually synonymous with the race to push ever more pixels onto screens. The first IBM PC had no graphics display at all; its MDA (Monochrome Display Adapter) only showed text. In quick succession, PCs adopted the 320×200-pixel CGA (Color Graphics Adapter) and 640×350-pixel EGA (Enhanced Graphics Adapter) formats before landing on 640x480-pixel VGA graphics in 1987.

In more modern times, American television's ancient NTSC standard of 525 blurry scan lines (or PAL's 625 lines, found in most of the rest of the world) gave way to High Definition (HD) in 1998. This is around the time that computer and television standards merged, giving us a veritable parade of pixel densities. 720P (1280x720 pixels) was quickly supplanted by 1080P (1920x1080 pixels). Today, 4K (3840x2160 pixels, also known as Ultra HD) reigns, but 8K (7680x4320 pixels) waits in the wings.

The moral is simple: More pixels have always been better. Modern techniques are enabling the creation of content with unlimited resolution, able to scale up to whatever display device exists. "We see a lot of industries moving from video playback to real-time 3D because it is inherently resolution independent," says Marc Petit, General Manager for Unreal Engine at Epic Games. "When you generate the frames on demand, you can get that content at 2K, 4K or 8K – you just have to tell the renderer how many pixels to render."

IMAGE RESOLUTION NEEDS ARE BOUNDLESS

We're entering an age in which the need for image resolution is boundless. From 4K to 8K to 32K to (n)K, our demand for more resolution in visual content remains insatiable.

Resolution has myriad applications. High resolution capture allows the composition and framing to be changed in post, for example, even giving the impression of multiple camera angles. It ensures the content is future-proofed, able to accommodate future display formats. And resolution isn't limited just to video capture; it's integral to graphics and visual effects as well. Textures, for example, add another dimension to the needs of creators. Especially in real-time environments like games, objects might have multiple textures that need to be loaded in real time as the viewer moves through the environment, getting close enough to reveal limitations in visual quality.

FROM 4K TO 8K TO 32K TO (n)K

"Everyone is being asked to create more and more photorealistic content," says Ross McKegney, Director of Engineering for 3D and AR at Adobe. "Historically, due to hardware limitations, we were down-sampling everything because you didn't have enough memory on the GPU, you didn't have enough compute capability. That's increasingly not the case anymore, and that expectation is driving the need for higher quality geometry and higher resolution textures everywhere."

VIRTUAL REALITY DRIVES THE NEED FOR RESOLUTION

These kinds of issues are manifesting themselves in virtual reality (VR), which is seeing significant growth in industrial applications. First-generation consumer tethered-VR systems which appeared on the market around 2016 displayed a total of 2160x1200 pixels spread across a 100-degree field of view.[1] But humans can see at least 180 degrees, and sometimes as much as 210 degrees. Lacking all that peripheral vision, there's no denying you're wearing a headset that's delivering a truncated view of the (virtual) world. For VR to be more immersive, resolution needs to increase along with field of view (FOV).

There's been some progress; the more recent HP Reverb VR headset, for example, includes a more expansive 114-degree field of view and a display resolution of 2160x2160 pixels per eye. It's a dramatic improvement compared to 1080x1200 pixels per eye found in first-generation headsets like the HTC Vive or even a more recent standalone headset like the Oculus Quest.

But some simple math illustrates why there's so much room for improvement to enhance VR realism. If you double a 100-degree FOV to 200 degrees – full human field of view – you need to quadruple the resolution from 2880x1700 to 5760x3400 pixels just to keep the same pixel density at that larger size. To drive a headset running that kind of resolution, you don't just need new display technology, but also a display adapter, processor, and memory to move all those pixels at up to 90 frames per second, at a minimum.

1 https://www.vive.com/us/product/vive-virtual-reality-system/

That's just part of VR's problem. The real world has virtually infinite resolution – everything stays razor sharp no matter how close you get to it. If VR designers are going to mimic the real world, they need to be able to simulate reality's nearly infinite resolution, and that requires the ability to scale textures in real time.

"Honestly, the goal of a real-time engine in the past wasn't to be photoreal," says McKegney. "The goal today, and especially going into the future, is to be completely photorealistic in real time, in VR and elsewhere." Petit agrees: "Even if you're creating a fantasy world it's a world that does not exist. But photorealism is the benchmark."

Because it's impossible to store the virtually unlimited resolution needed to accommodate viewing every object in a VR environment close up, the technique of only loading higher-resolution textures as they're needed is a tried and true technological shortcut borrowed from the world of video games. In other words, (n)K resolution is, at least in part, about having a library of texture resolutions stored in software to deliver the fidelity needed in any given situation.

COMPUTATIONAL PHOTOGRAPHY

There was a time when most kinds of computational photography were something of a niche – but these

The HP Reverb VR headset expands peripheral vision to a 114-degree field of view, bringing the field of vision closer to that in our everyday life.

days, computational techniques have gone mainstream. To make up for the lack of interchangeable lenses found in professional cameras, for example, smartphones increasingly rely on their powerful processors and computational techniques to take better pictures, putting this technology in the hands of millions of people. Panoramas, for example, are a rudimentary example of computational photography. As you pan a phone, it takes a series of photos and then "stitches" them together in real time, creating one long, seamless image.

Google's Pixel 3,[2] for example, captures 15 photos at a time when you press the shutter release, and then interpolates them into a super resolution photo that's sharper and larger than would be possible by taking

2 https://store.google.com/product/pixel_3

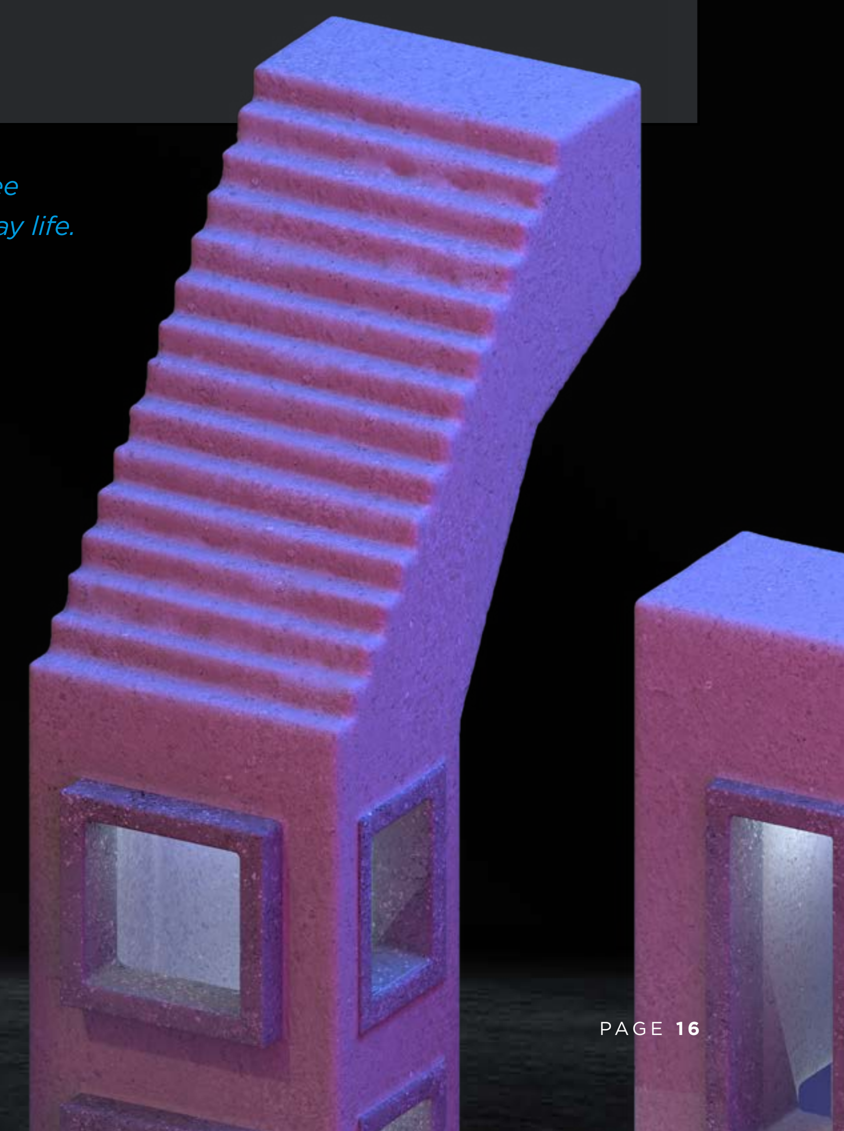

"We're at the moment in the industry where we finally can have a 'what you see is what you get' experience."

SÉBASTIEN DEGUY
Vice President of 3D and Immersive, Adobe

the photo in the traditional way. It does this, in part, by comparing similar pixel regions in each of the various photos. Ordinary camera shake from hand-holding the phone allows the software to sample each part of the image slightly differently, for an overall better image.

There's also High Dynamic Range (HDR), which uses a series of photos shot at various exposures to increase the dynamic range of photos and video. This solves an age-old problem; the human eye has significantly more dynamic range (in photography terms, about 30 stops, or a dynamic range of about 1,000,000,000:1) than film or digital cameras (both of which top out around 12 stops, or just 4,000:1). Any scene that has very bright and very dark elements will invariably require a "compromise" in setting the exposure – hence, flat and lifeless sunsets.

But (n)K resolution comes to the rescue. Like the library of textures that need to be rendered in VR, computational photography relies on an enormous library of photos to be captured and processed in the background to make modern photography possible. The numbers are sobering. When you consider the commonly cited statistic that more photos were taken this year than in the whole previous history of photography,[3] it becomes clear that we're entering a whole new epoch for processing and storage.

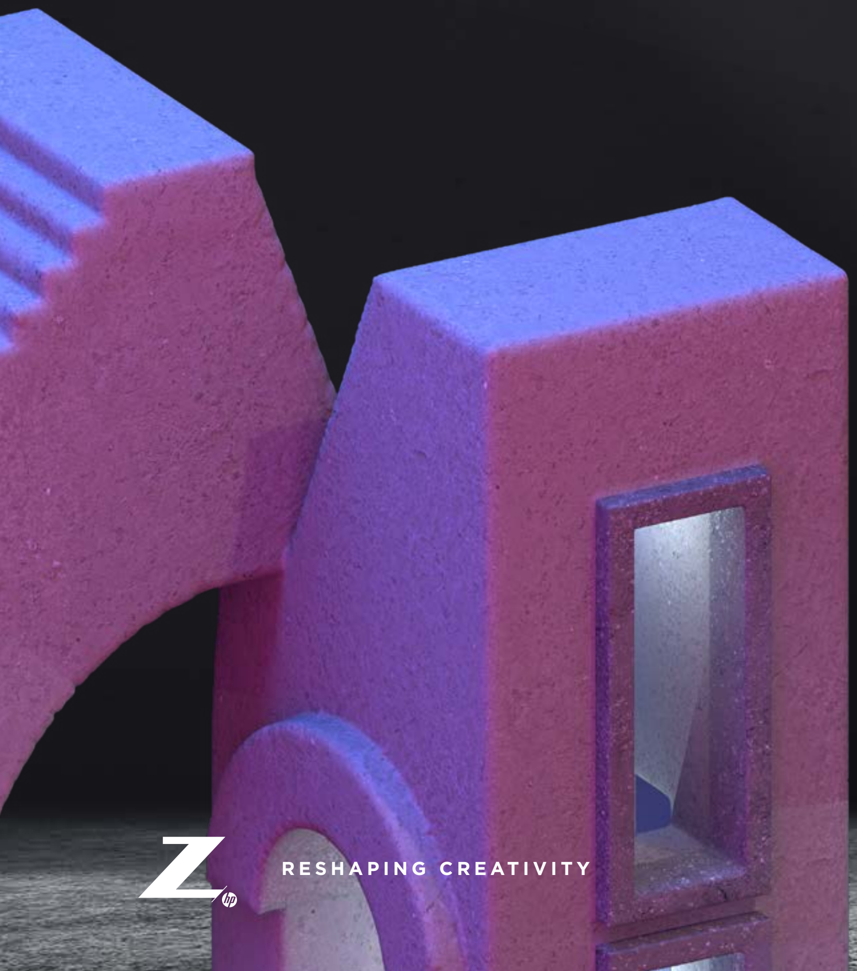

RAY TRACING

For decades, ray tracing has been an arcane side-branch of computer graphics. By tracing the path of every ray of light entering the frame, a ray-traced image can reach otherwise unapproachable levels of detail and realism, with high-fidelity reflections, water effects, lighting and shadows. Translucent materials can be accurately modeled, along with how light and colors realistically scatter through glass, crystal, water, and other tricky real-world materials. Ray-traced visuals are almost immediately distinguishable from traditionally rendered scenes, which take shortcuts to simulate what a scene would look like in its virtual environment.

Traditionally, ray tracing has required enormous amounts of computational power, and rendering even a short scene just a few years ago could take hours, days, or weeks depending upon the complexity. They were the domain of Hollywood, which had the lead time to render lengthy video sequences, and even there rarely applied.

"To get anything close to photorealism, we used to write complex shaders specific to the frames being rendered", says McKegney. But that has changed. Now, PBR materials and ray tracing are a routine part of the lighting toolkit in film, gaming, and other rendering applications. "Ray tracing is particularly geared towards managing reflections and refractions, and so it's a complementary technique to global illumination," says Petit.

Thanks to a new generation of graphics processors like the NVIDIA® Quadro RTX™ family of graphics cards that are capable of rendering ray-traced scenes in real time, the demand for real-time rendering and buffering of ray-traced imagery will soon rival that of traditionally rendered video content.

For Sébastien Deguy, Vice President for 3D and Immersive at Adobe, this is a critically enabling technology.

3 https://www.macfilos.com/2015/09/01/2015-9-1-more-photographs-taken-in-one-year-than-in-the-entire-history-of-film/

Still image from the new Unreal Engine 5 demo, "Lumen in the Land of Nanite," that highlights two of the latest generation's core technologies: Lumen, for real-time dynamic global illumination, and Nanite, for high geometric detail micropolygon geometry modeling.

"It's a total game changer. We're at the moment in the industry where we finally can have a WYSIWYG [What you see is what you get] experience in 3D. You don't have to hit the button and wait for hours to have it render. It's completely photorealistic and it's in front of you all the time during the creation process."

IMPLICATIONS

There was a time when display resolution was a major limiting factor in the pursuit of image quality. And while display resolution will continue to grow to accommodate larger displays and novel applications like VR, (n)K resolution has important implications, ensuring resolution independence for scalable and realistic texture maps, and for the enormous number of pixels being produced behind the scenes for tasks like computational photography.

Every sign points to the need for exponential growth in graphics processing, memory, and storage needs to keep pace with these files, which are likely to continue to grow in an unbounded way. "Everything is going to be photorealistic, and that means we need better hardware to create it and better hardware to experience it," says Deguy.