Fifty Shades of Yay

A Multi-Shot Workflow from Design to Final

Dave Quirus
Matt Roach
Matt Wilson
Blue Sky Studios

Figure 1: Colored Explosions, Spies In Disguise. ©2020 Twentieth Century Fox Film Corporation.

ABSTRACT

At the peak of our climactic third act in “Spies in Disguise,” our heroes unleash a barrage of saturated powder grenades around the villain’s army of drones. This meant filling the screen with billions of finely detailed, multi-color voxels and particles (aptly named “Fifty Shades of Yay”) with a distinct performance and look. We needed to come up with a solution that was quick to turn around, yet highly directable and allowed for unfettered lighting and shading development. The resulting approach was an extremely tight collaboration between several departments and spearheaded by a few folks in Lighting and Effects.

CCS CONCEPTS

• **Computing methodologies** → **Procedural animation**; **Physical simulation**.

KEYWORDS

Volumetrics, Blue Sky Studios, Spies In Disguise

ACM Reference Format:

Dave Quirus, Matt Roach, and Matt Wilson. 2020. Fifty Shades of Yay: A Multi-Shot Workflow from Design to Final. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH ’20 Talks), August 17, 2020.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3388767.3407341

1 SIMULATION DEVELOPMENT

The directors’ vision called for dramatic towers of color composed of a repeating multitude of rocket launches encircling our villain. The volumetric portion of this effect could not simply billow like smoke. We needed to drive it upwards quickly like pyroclastic rocket trails and have them expand outward like chalk blowing off of disintegrating clumps. We drove the powder up with an underlying particle system that emitted density and velocity into the simulation. The expansion of the trails was handled by having the particles emit divergence into the fluid sim as well.

Aesthetically, we wanted the powder to move like a volume, but also feel more like particulate when up close. We ultimately layered in texture and scale by adding a power law distributed particle system to the mix. Pure advection by the volumetric fields wasn’t enough to maintain the proper energy and timing, so we included the original source particle attributes, along with sculpted velocities from the point of emission to add interest and spatial coherence with the volumes.

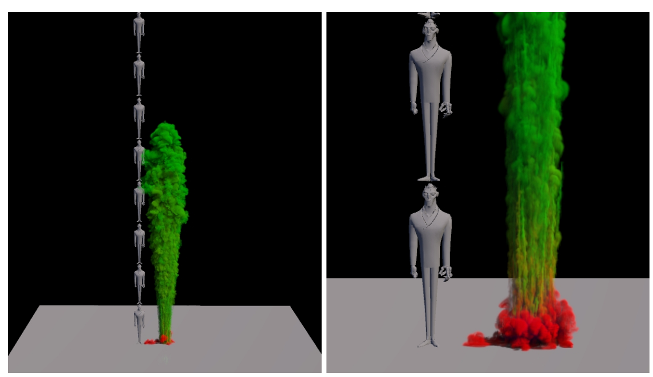

Figure 2: Pigment Advection and Integration. ©2020 Twentieth Century Fox Film Corporation.

Finally, both the volumes and the particulate needed to dynamically mix their pigment composition upward throughout each column to provide a natural range of hue from their source. Instead of running each column in each of the six art directed colors, we

SIGGRAPH ’20 Talks, August 17, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-7971-7/20/08.
https://doi.org/10.1145/3388767.3407341

ran two primary colors that would still intermingle within the simulation but could be entirely manipulated via shading networks during render.

Once simulated, our various elements needed to support enough resolution for close up and far away distances while supporting character and set interaction. Each column had controls to be time-remapped for fast-paced beats and provided render-time parameters for both lighting and effects to adjust a final look in an extremely efficient manner.

2 MULTI-SHOT WORKFLOW

To accommodate the ranging series of extremely heavy data across each of the shots, we took advantage of utilizing packed disk primitives. With visual requirements pushing us above a total of 11.5 billion sparse voxels on-screen, this allowed for real-time placement of each column and its accessory elements in addition to significant disk I/O speed ups. Memory allocation was greatly reduced during both staging and rendering due to these delayed load techniques as well. Time slipping could be handled on the fly and our particles were further packed into spatial tiles based on their resolution for even more efficiency.

Figure 3: Packed Disk Primitives Set Dressed as Spatial Tiles. ©2020 Twentieth Century Fox Film Corporation.

Material overrides were then built into the shaders for dynamic IPR updates and a near-instantaneous re-cache of our heavy data for render farm submissions. To take advantage of cryptomatte output for compositing, we structured a standardized string convention that would generate procedurally for every effects element that was dynamically added or removed for consistency and clarity across all shots.

Once each shot had timing, placement, and color approved, final touches were individually added to each shot for set and character integration. The benefit of this approach allowed us to nail down what mattered most in a very short amount of time with greater flexibility for change. It was only during this final stage that everyone in the room could make clear in-frame decisions about where these hero interactions needed to occur rather than being committed to it from the very beginning. Our final iteration cycle was reduced so that notes could be easily identified and addressed in a direct and controllable way.

3 LOOK DEV & FINISHING

The final steps were to revisit the color keys, evaluate the lighting/shading implications, and see how far we could push the look. From a computational perspective it was clear that we would be dealing with an extremely complex lighting environment. Combined voxel counts in the billions completely surrounding and overlapping our characters brought up several red flags. We started wondering if combining outputs from multiple render engines would get us what we wanted with the most control over the look. It became a game of deciding the composition of the elements on a pixel to pixel level. The solution was breaking the frame down into lighting components both on the volumetric effects themselves and their lighting/material contribution to the scenes. In the end we transferred the color parameters onto iso surface approximations to account for secondary illumination as well as shadow casting and diffuse reflections. This meant we could render the characters and the sets in the standard pipeline in Blue Sky's proprietary renderer, CGIStudio. All the primary layers and AOVs for the color clouds were generated in SideFX's Mantra, split out into their individual light paths per light/per LPE and then re-assembled in Nuke. We also generated per element Cryptomattes (one for each column, foot and particle systems). Character holdouts were built using a combination of depth renders from both render engines to account for displacement and other geometry specific to Studio++ (smoothing, hair, displacement, etc.). What could have easily ended up requiring thousands of hours of render time per frame was eventually optimized down to an overnight render with substantially more control over the nuances of the image.

4 WHAT WE GAINED

The philosophical advantages learned through this process outweighed the technical ones. In the future we know that artists can avoid the brute force method of constant in-shot simulation and step back to think holistically from a multi-shot perspective and establish a flexible workflow around it.

We also know that small teams can adapt quickly to changing requirements and they can accomplish goals set to a high bar in a tight schedule when given the opportunity to step out of an established pipeline, as long as there is a developed plan put into action. This kind of collaboration gave artists ownership over their process and ensured an extremely successful end result for everyone involved.