

facebook

An on-demand, optimising gLTF backend

Pär Winzell

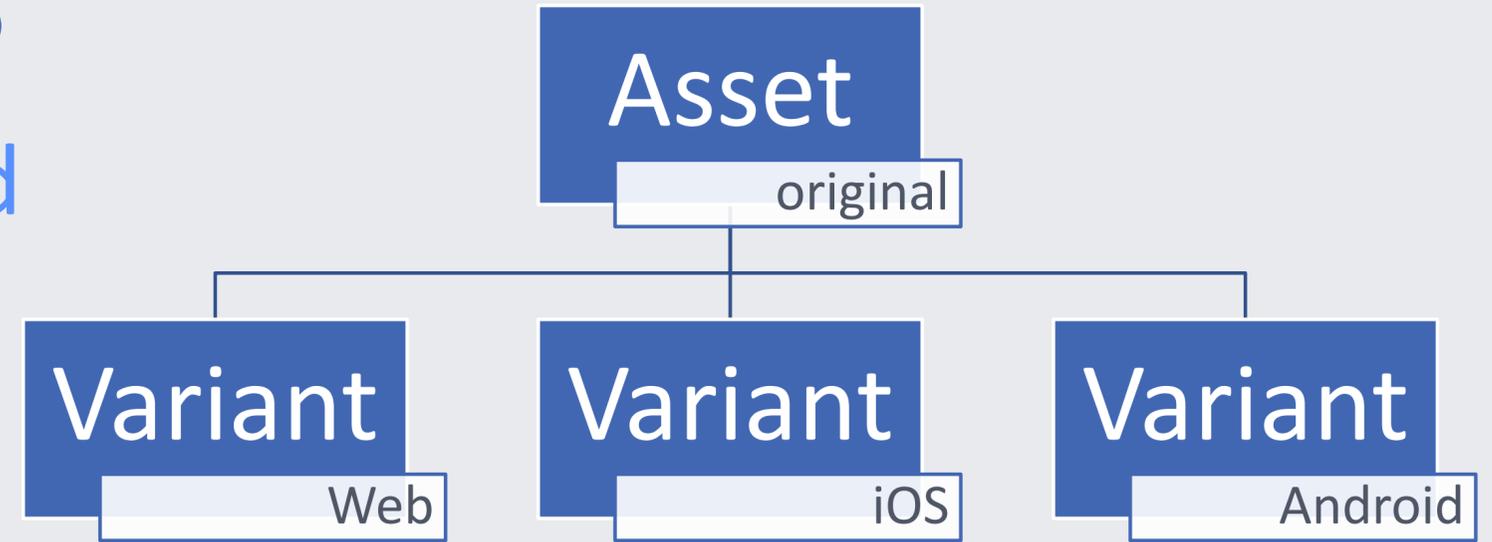
zell@fb.com

Susie Su

susiesu@fb.com

Early days: 3D Posts

Building asset variants on demand



- Tear uploaded glTF down to the level of logic & data
- Perform client-specific optimisation and customisation
- Rebuild glTF to specification, for delivery

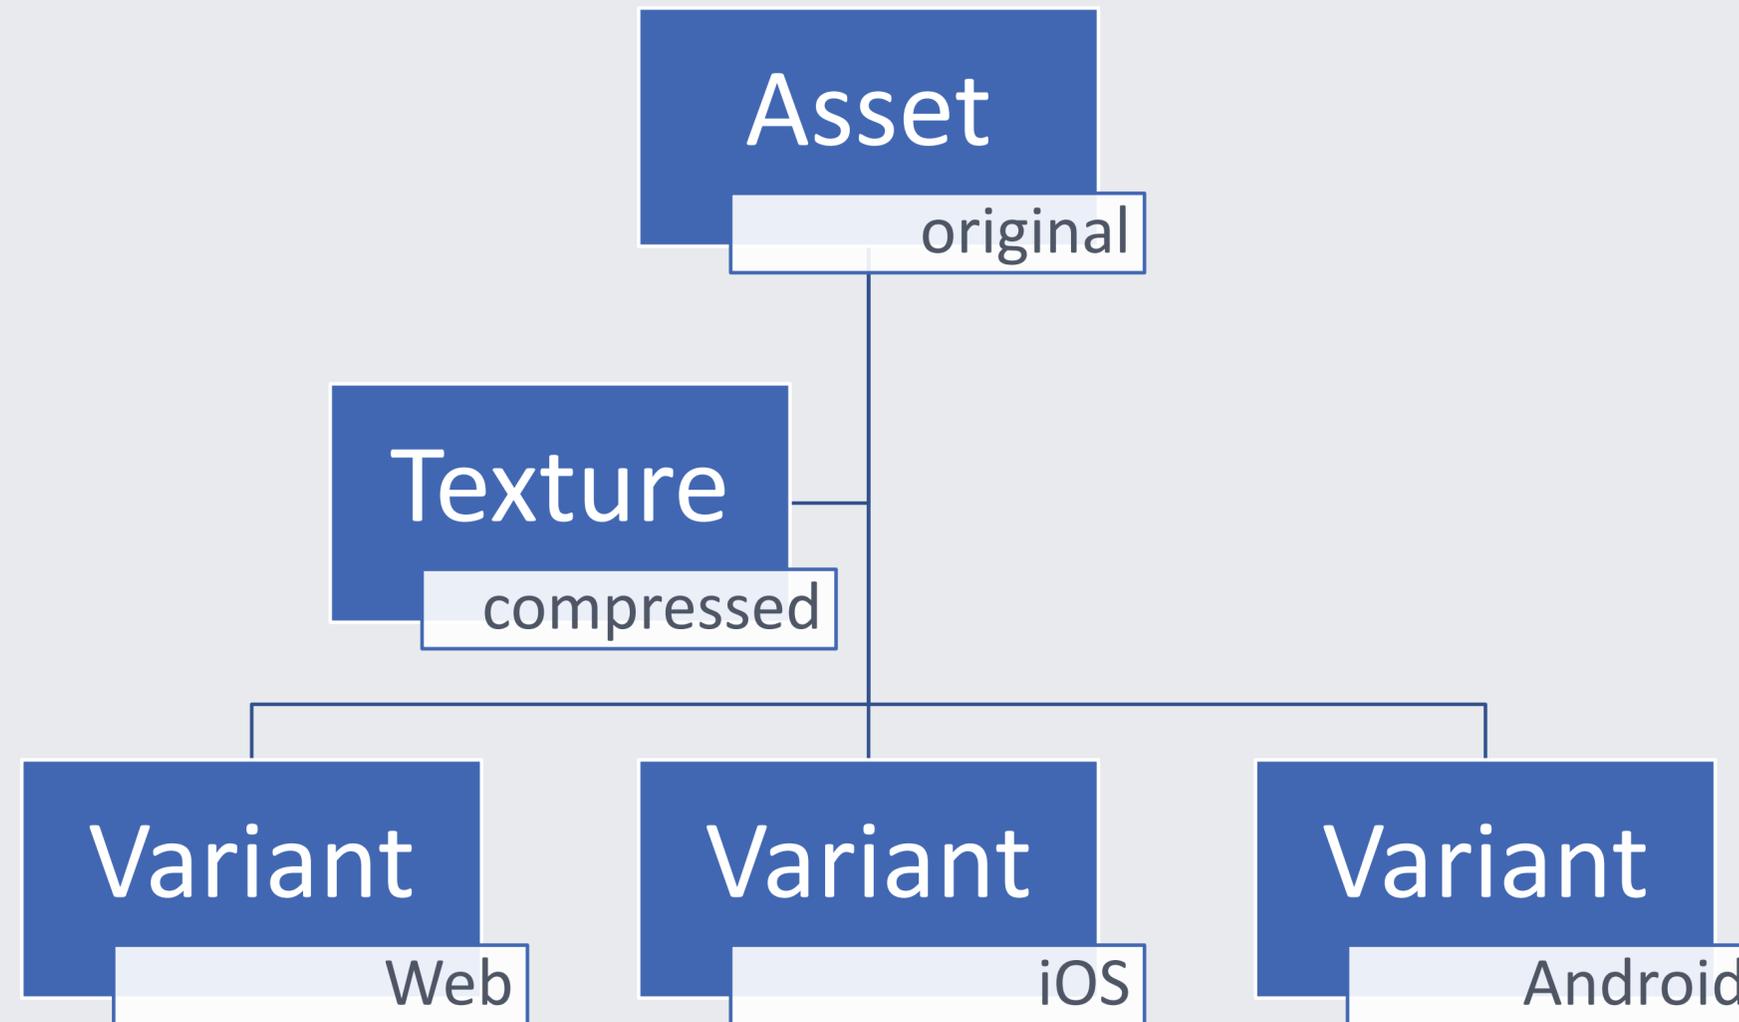
Working at scale: 3D Photos

Millions of models, millions of textures...

- Turns out older phones really need compressed textures...
 - Quickly! Build an in-house texture compression service.
- Turns out people quite like 3D Photos...
 - Quickly! Scale to compressing hundreds of textures per second.
 - Still, texture compression is at odds with on-demand generation.

Texture caching

2nd layer of caching – reuse compressed texture on variants creation



Extend system for broader adoption

Be the backend for all 3D glTF products

- Generalized the 3D infra to power up products from different apps across the company
 - Facebook 3D post and 3D Photo
 - Oculus Home

Ongoing Challenges

Uploaded models in busy online virtual worlds

- “Level of Detail” now inescapably required
 - Too computationally expensive to generate on demand
- Streaming textures and geometries
 - Very different aesthetic impact, depending on context
- Configurable asset specification
 - Single asset which holds variations of different material/colour etc.