



Imagination

HPG2012 - The Upwardly Mobile GPU

Simon Fenney - PowerVR Division
25 June, 2012

- **Features have followed same, albeit accelerated, progression**
 - Texturing + fixed function pipeline →
 - Programmable T&L →
 - Fragment & Compute shading
- **Always* will be a gap**
 - Due to power and relatively tiny SOC memory interfaces.
 - *But* mobile performance is growing at much higher rate than desktop
 - Top end mobile GPUs (i.e. Tablet) likely to soon equal mid-range desktop.

* Well, probably!

- **What features might we see in future mobile devices?**
 - HOS/Tessellation (but maybe not *exactly* the same as Desktop)
 - High dimensional rendering
 - Order Independent Translucency (it's been done before)
 - Ray Tracing
- **'Mobile' should imply power efficiency**
- **So any new feature 'just' has to be done efficiently**
 - But how?

- **Fully programmable HW is great but...**
- **Use flexible, fixed-function hardware for ‘common’ tasks**
 - e.g. 1080p video decode is $\approx 12\text{mW}$ with dedicated HW.
 - Hard to achieve with general purpose HW.
- **SW: Use the right tool for the job**
 - SW needs to consider what things cost. E.g. precision
 - Will LowP ($\sim 10\text{bits}$ fixed) or MediumP (16-bit float) do?
 - Don't just jump to HighP ($\sim 32\text{bit}$), (and do we *really* want 64-bit yet?)
- **Learn from Desktop – but be more aggressive.**
 - Consider Desktop as a ‘prototype’ and tweak e.g.
 - Mobile APIs: took Desktop versions and weeded out less efficient interfaces.
 - Texture compression: Desktop: 4~8bpp vs. Mobile: 2~4bpp
 - Not convinced unification a good idea *just* yet.

- **Given the rapid expansion of the mobile market...**

...might the rôles reverse?