

# UG Exam Performance Feedback

## Third Year

### 2018/2019 Semester 1

COMP37111 Advanced Computer Graphics

Toby Howard  
Steve Pettifer

- Comments**
1. Generally well answered; the question asked for two examples; some people only provided one; some gave none. A common mistake was saying that self-similarity means that all sub-parts were the same. This is true in a specific class of mathematically defined examples, but in the general case it is not "sameness", it is "similarity".
  2. I had intended this to be an easy question, and I was surprised by the variability in the answers, most of which lacked detail. The question asked for a modelling technique per element, and also a technique to closely match the synthetic element's appearance with that in the photo. Many people ignored the appearance aspect, so lost the marks.
  3. This was well answered. Most people understood that as CPU/GPU power grows, so do the demands on them. I had expected also comments about legacy, given that applications should be expected to run well on a range of systems, not just the cutting-edge. About half the class discussed this.
  4. This was an easy question and was well answered. Almost everyone chose S2 as the preferable mesh, and for the right reasons.
  5. Another straightforward question, about a fundamental aspect of the scanning process. The answer was the problem of occlusions, and most people understood this. A few people seem to have got scanning mixed up with geometry extraction using feature matching in images, and gave strange hybrid answers which alas made no sense.
  6. This question related to one of the required-reading research papers, which I had summarised in the lecture. It was clear that some people had read the paper and could talk about the technique. They scored well. Other people appeared not to have looked at the paper, and made something up which was usually expressed in very vague terms, so did not generally score well.
  7. A fairly easy question, generally well-answered. Most people were able to talk about lens distortion and one or two methods of correcting it. But, not many people actually answered the part of the question which was "describe the effects of lens distortion".
  8. This question was well-answered, and there was generally a good understanding of the use of feature detection. As with Q5, some people had conflated the ideas with those of laser scanning, which was surprising.
  9. Generally well answered; almost everyone got at least the name right (half marks for calling it the 'global illumination equation'), and most students got the general purpose of the equation right or very close to right (I was lenient in terms of use of words like 'reflection' and 'colour' for this one, but really it should have been described in terms of light energy and wavelength)
  10. Reasonably well answered. Full marks were rare, though were typically gained either by a fairly terse formal definition (irradiance, radiance etc), or a longer narrative one that covered material behaviour and angle. I was generous in the use of 'reflectance' (correct) and variations of it (e.g. reflectivity, reflection) because although the terms are used correctly in the lectures and notes, enough people got the detail of that wrong for me to conclude that I didn't make enough of a thing about the difference between the terms this year.
  11. Well answered on the whole, though marks were thrown away by just doing a braindump of everything known about ray-tracing rather than explicitly referring to the figure as the question required (For full marks I wanted to know that not only can the student recall the process, but can apply and explain it in a specific context)
  12. This question was obviously quite challenging for many, and relatively few people got the perfect answer. Quite a few punted the performance, view-dependance, 'backwardsness' or 'doesn't do matt' angles, which aren't correct (or rather, the matt one sort-of-is part of the issue, but you'd need to explain how in a fairly subtle way to get all the marks). There was a fair amount of word-soup'ing where people blurted all possible things and probably picked up a half mark or two by brute force.
  13. On the whole this was answered well, with fairly detailed descriptions of the steps involved. However, and this is my top tip for exams generally.... lots of people threw away marks by not referring to the context of the question and instead regurgitating examples from my notes or lectures. If a question has gone to the trouble of inventing a mysterious insect and encasing it in an unknown resin, then that has happened for a purpose, and not just for giggles. The difference between regurgitating my notes and being able to translate understanding from the course into a context-specific scenario matters here, so for full marks I was looking for the ability to apply knowledge, not just to recall stuff.
  14. The final question was a really interesting one to mark. When I set it, I (lazily, now I realise!) had assumed that

# UG Exam Performance Feedback

## Third Year

### 2018/2019 Semester 1

---

the obvious answer was to use an octree (and indeed, most answers were along this line, and did a decent job of explaining the how and why). Some people went for other spatial enumeration mechanisms, and as long as they also explained how you'd go from data points to objects to use these other approach that was fine too (otherwise I'd accept the approach for partial marks). What I hadn't expected was the interesting and creative alternatives that people invented, such as explicitly identifying and then throwing away all the 'resin only' cells, or using a shrinking sphere to find out which was the minimum set of cells that contained interesting stuff (i.e. both of these are very specific to the scenario I described and wouldn't work more generally, but that was great and better than the answer I had chosen. So Kudos to those that came up with these). The other thing I wasn't expecting was people to go for indirect volume rendering techniques as a solution to the problem -- the question actually just asks for an alternative data structure rather than an alternative rendering approach, but I decided that there was just enough wiggle room in the wording of the question to accept this as a valid kind of answer.

---