APPENDIX F

Participants in the qualitative elegance investigations reported in chapter 9 of this thesis were invited to provide any comment on their human experience of their interactions with the early lifecycle software design interactive framework. Four of the participants provided comments which are quoted as follows:

Participant 3

In the first of three sessions I was asked by Chris Simons to view a class diagram created by his Computer Based System (CBS) for a cinema-booking system case study and, on the basis of its perceived elegance, to assign it a score from one to five where one was low and five was high. It was pointed out to me that the degree of coupling and the degree of cohesion might sensibly influence my judgement of elegance. In the depicted models, cohesion was denoted in each class by a colour: red meant high cohesion, blue meant low cohesion, and there were colours in between (yellow and grey). Coupling was denoted both by the number of connections (lines drawn) between classes and the thickness of lines drawn: high coupling was denoted by thick lines and low by thin.

I was given to understand that after I had given my score for elegance to a CBS-generated design, I could either ask the CBS to generate another design or I could quit. I was told that I could quit if I was bored or tired, or if I thought I had produced an optimal design (or at least an unimproveable design), or for any reason.

I conceived the goal of the task as one of generating the best possible design in terms of minimising coupling and maximising cohesion. After "playing" a few rounds, I had the sense that I could influence the next design generated by the CBS as follows. If the current design was "better" than the previous one, either in terms of coupling or cohesion or both, and if I assigned it a higher score than the previous one, then the next design generated was invariably better again. So I "encouraged" the CBS to move towards an optimal design by assigning scores on the basis of this assumption. When the design seemed impossible to improve on for two or three successive goes, then I would quit.

I did worry that if I had started each trial (each trial for a given scenario started with configuration using some sort of mystery number) by assigning a different score, usually a lower score, then perhaps I would have needed up with an even "better" final

design. But there didn't seem to be a way of testing this hypothesis since one couldn't repeat a trial with the same mystery number.

I was also aware that I wasn't primarily judging elegance in terms of the aesthetics of the appearance of the designs: I was more influenced by the coupling and cohesion measures, although I thought these were correlated with elegance in the sense that simpler looking redder diagrams were better.

I repeated the process described above for two other increasingly more complex class diagrams intended to model increasingly more complex scenarios. With one exception, the experience was the same as for the first scenario and set of designs. The exception occurred for the final scenario. This was so complex that I needed to write down (i.e. record using pencil and paper) both the coupling score and an assessment of the overall cohesion in terms of the number of instances of red classes, the number of yellow, and so on. This was because from one trial to the next, I could not remember precisely enough how good (or bad) the coupling or cohesion had been.

I didn't find the sessions tiring; I tended to stop to give time to other concerns. I found the sessions quite enjoyable, the enjoyment gained from working towards, and seeming to achieve, a useful goal: i.e. a good design.

Participant 5

Making a value judgement on a complex software system was interesting and maybe a little daunting. I felt sometimes that my judgement values altered during the course on a run – especially for the more complicated examples. So … the perceived lack of consistency could undermine the confidence in the value of the decisions.

The process was quite comfortable – the problems were explained well enough. The screen was readable and re-arranging the classes on screen gave time to analyse the solution a little. However it is possible that the particular layout (or even process of laying out) could influence the final decision – especially when the screen was "busy".

The use of colours was a heavy influence – reading the relative cohesion values conflicted with the colours. I felt that I had the let one aspect "lead" the decision – coupling was a strong indicator. Perhaps because the cohesion values and colours appeared misleading.

Participant 6

I believe the tool is to assist software designers to identify good quality designs. The concepts of quality in this case are specifically related to high cohesion and loose coupling. The interface of the tool is very easy to follow. The idea of being able to visualise the degree of cohesion and coupling is very good. I believe the tool helps the users to easily understand the quality of the software design.

Classes are given different colours to indicate the degree of cohesion, with red being highly cohesive and grey being low cohesion. This is especially helpful. Currently the colours are scaled based on the selected classes in the chosen design. For further improvement, I would recommend the colours being scaled based on a much large set of classes.

Participant 7

I found testing the tool to be surprisingly engaging, so much so that, at times, I think I lost sight of the aim of the task and was more focussed on looking to see how the results changed from one output run to the next. I'm not convinced that I really made any "elegance" judgements. My judgement was principally guided by the tool. I immediately viewed output runs with lots of red as potentially interesting. On reflection, this was poor judgement since "lots of red" might just have indicated that the majority of classes generated were equally (but not very) cohesive but that one had a very low cohesion value. Nonetheless, the colours had a huge impact on my decision making.

I found the cruising problem extremely difficult to concentrate on. I was very difficult to absorb the very rich information presented by the tool.

A matter of some frustration was the fact that I couldn't save any "good corners" of the designs generated. From time to time a pearl of a class was generated amongst the swine. I would like to have been able to freeze that class and save it from further evolution. In general I guess I'd have liked even more interactivity. I'd have liked to be able to use the tool to come up with some initial class diagrams that I could modify by hand and then (perhaps) present again to the tool for a bit more refinement.