Graph Drawing Aesthetics in User-Sketched Graph Layouts

Helen C. Purchase
Department of Computing Science
University of Glasgow
Glasgow, G12 8QQ, UK
hcp@dcs.gla.ac.uk

Beryl Plimmer, Rosemary Baker, Christopher Pilcher
Department of Computer Science
University of Auckland
Auckland 1142, New Zealand, UK
beryl@cs.auckland.ac.nz
{rbak026,cpil016}@aucklanduni.ac.nz

Abstract

Empirical work on appropriate layout aesthetics for graph drawing algorithms has concentrated on the interpretation of existing graph drawings. A more recent experiment has considered layout aesthetics from the point of view of users moving nodes in an existing graph drawing so as to create a desirable layout. The project reported here extends this research further, by asking participants to use sketching software to draw graphs based on adjacency lists, and to then lay them out – removing any bias caused by an initial configuration. We find, in common with many other studies, that removing edge crossings is the most significant aesthetic, but also discover that aligning nodes and edges to an underlying grid is important, especially to male participants who have Computer Science experience. We observe that the aesthetics favoured by participants during creation of a graph drawing are often not evident in the final product.

Keywords: Graph drawing creation, graph aesthetics, sketching.

1 Introduction

Typically the aesthetic criteria for graph layout algorithms have been based on the intuition of the algorithm designers. Some empirical work has been done in an attempt to verify the usefulness of these criteria in reading and understanding graphs (as measured by graph tasks, often shortest-path tasks) (Purchase, 1997). More recently, work has been done on the process of graph comprehension, using eye-tracking data (Huang, 2007). These empirical studies have all related to the reading and understanding of graphs, rather than their creation. A recent study by van Ham and Rogowitz (2008) considered the creation of graph layouts, looking at how people prefer to layout graphs when they have the opportunity to move nodes and edges around. This is a very different empirical task to that of reading and interpreting a graph, and Van Ham and Rogowitz investigated, in particular, the depiction of clusters, the presence of edge crossings, edge length distribution and orientation.

The work presented here builds on the work of van Ham and Rogowitz (2008), but differs in two important ways. First, the participants drew the graph from scratch, rather than simply moving nodes in a pre-drawn graph – this removed any layout bias present in the original drawing. Second, they used a sketching tool with a stylus – this allowed for the physical action of creating the graph to be done as easily as if it were on paper, reducing any cognitive distance between the participant's desired drawing and what is presented. Using the interface of a formal graph drawing tool is a less natural way of drawing a graph than the free-form hand movements made possible by a sketching tool.

This paper describes previous graph drawing empirical work, discusses the work of van Ham and Rogowitz and its conclusions, outlines the graph sketching experiment, and presents its results. Three aspects of the human graph drawing process are discussed: the product, the process and the preferences.

2 Related Work

The many graph layout algorithms that have been devised over several decades (Battista et. al., 1998) have typically been designed in accordance with the intuitions of the algorithm designers. Over the years, a set of assumed 'graph drawing aesthetics' has emerged, defining the criteria by which the 'goodness' of the graph drawing produced by a layout algorithm can be assessed (Coleman and Stott Parker, 1996, Purchase et. al., 1995). Such aesthetics include, for example, a minimum number of edge crossings, as few edge bends as possible, a display of symmetric sub-structures, and large angles between edges incident at a node. Graph layout algorithms therefore tend to be valued for the extent to which their output graph drawings conform to these aesthetic criteria.

Graph drawings should also be assessed according to the extent to which they assist human comprehension of the relational information represented in the graph. Some empirical work has been done to this end, investigating whether the aesthetic criteria used by algorithm designers do indeed assist with comprehension. Findings include the overwhelming evidence for the reduction of edge crossings (Huang et. al., 2006, Ware et. al., 2002, Purchase, 1997), some evidence for the reduction of bends and depiction of symmetry (Purchase, 1997), placement of important nodes at the top of the graph (Huang et. al., 2007) and large angles between incident edges (Huang 2007). All these studies have been conducted by asking participants to answer graph-based questions on a variety of presented graph drawings, each carefully controlled for the aesthetic criteria.

A recent publication by van Ham and Rogowitz (2008) has taken a different empirical approach to determining the best graph layout for human use. They asked participants to manually adjust the layout of existing graph drawings: "rearrange the nodes in the network in a way that you think best reflects their interconnections." They used four graphs of 16 nodes, each with two clusters: these clusters were separated by one, two, three and four edges respectively. These were both presented in a circular and a spring layout (Gansner et. al., 2005), giving a total of eight starting diagrams, which were presented in random order. Users of the Many Eyes visualisation service (Viegas et. al., 2007) were invited to take part in the experiment, by adapting the layout of the drawings. They collected 73 unique drawings, which they visually analysed according to the number of edge crossings and evidence of clustering, as well as other computational measures such as edge length distribution and cluster distance. They found that most participants separated the two node clusters, the human drawings contained 60% fewer edge crossings than the automatically produced drawings, and that humans did not value uniform edge length as much as the spring algorithm did.

Van Ham and Rogowitz (2008) acknowledge the limitations of their work. In particular, using a web-based experiment means that they have no information on their participants, apart from the fact that they had 'some sophistication in data visualisation.' In addition, the fact that the graphs were presented with an initial layout may have biased the resultant drawings, and the names used as node labels may have had an effect (for example, when participants attempt to avoid label overlaps).

The experiment reported here improves on van Ham and Rogowitz' (2008) methodology in several important ways:

- the experiment was conducted face-to-face, so demographic information about the participants is available and we know that all participants did all drawings;
- the participants drew the graphs from scratch, so were not biased by any initial layout;
- a sketching tool was used, so the physical drawing process was unhindered by a clumsy editing process and participants could draw curved or bent lines if they wished;
- we collected video data, so were able to analyse both the process and product of creation;
- we discussed layout preferences with the participants in a post-experiment interview, so were able to find out more about their thoughts on the process;
- our node labels were simple letters, enclosed within the node boundary.

Our more comprehensive and face-to-face methodology and our choice of equipment resulted in a smaller version of the experiment with 17 participants. Including the task of graph creation as well as layout meant that we used four graphs (two practise graphs and two experimental graphs), so as to make the duration of the experiment acceptable to participants. The screen size of the tablet PC sketching tool limited the graphs to 10 nodes each, with the clusters joined by one and two edges respectively.

Our experiment has produced extensive and rich data, in terms of product, process and preferences. The results can inform the design of automatic graph drawing algorithms by highlighting those features that users consider important when they are unconstrained in their own drawing of graphs and are not subject to any layout bias.

3 Graph Sketching Experiment

3.1 Equipment

A graph-drawing sketch tool, SketchNode (Reid et. al., 2007, Figure 1) was used on a tablet PC. This tool allows nodes and edges to be drawn with a stylus on the tablet screen, laid flat, thus allowing the same hand-movements as pen-and-paper, giving a more natural interaction than using an editing tool. Unlike pen-and-paper, however, the SketchNode interface allows nodes (or groups of nodes) to be selected and relocated (with corresponding movement of attached edges), and nodes and edges to be erased. It thus has the advantages of pen-and-paper, as well as the advantages of a graph drawing editing tool.

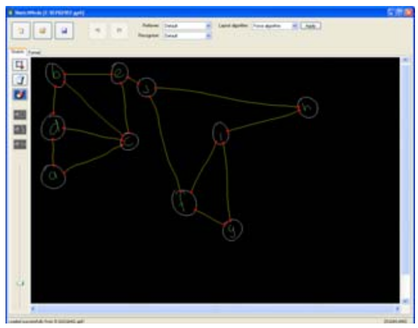

Figure 1: The SketchNode system

3.2 Task

Participants were given an adjacency list of edges (Figure 2) and asked to draw the graph in SketchNode using the stylus, with the instruction to *Please draw this graph as best as you can so to make it "easy to understand"*. They were deliberately not given any further instruction as to what "*easy to understand*" means. In particular, they were not primed with any information about common graph layout aesthetics, for example, minimising edge crossings, use of straight lines etc. They were given as long as they liked to draw and adjust the layout of the graphs.

Graph A	(A,D) (A,C) (B,D) (C,D) (B,C) (B.E) (C,E) (E,J) (F,G) (J,F) (F,I) (G,I) (J,H) (I,H)
Graph B	(J,F) (J,I) (G,I) (H,I) (G,H) (F,H) (G,J) (F,A) (F,G) (G,E) (A,E) (D,E) (D,C) (D,B) (C,B) (A,B) (A,C) (B,E)

Figure 2: The adjacency lists for the two experimental graphs

3.3 Graphs

We designed two experimental graphs: graph A had 10 nodes and 14 edges; graph B had 10 nodes and 18 edges. We were unable to use graphs as big as those used by Van Ham and Rogowitz (2008), as we were required to keep the duration of the experiment to a reasonable time, and the screen size of the tablet PC was limited. However, these graphs were still designed with the aims of Van Ham and Rogowitz (2008) in mind, as they both had identifiable clusters: graph A had two clusters separated by one cut edge; graph B had two clusters separated by two cut edges: Figure 3 shows these two graphs as drawn by participant 4, clearly showing the two clusters.

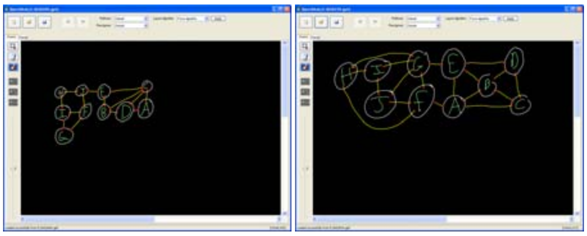

Figure 3: 4A and 4B: The drawings of graphs A and B by participant 4, using the SketchNode system

3.4 Experimental Process

After reading the information sheet and signing the consent form, participants answered a pre-experiment questionnaire which asked for some demographic information, including information about the extent of their experience with maths, theoretical computer science, graphs, and pen-based technology.

The participants were then given a demonstration of the SketchNode system, including all the interface features: node and edge creation using the stylus, selecting and moving nodes and edges, selecting and moving sub-graphs, labelling nodes, erasing, undoing and redoing actions, zooming and scrolling. They were also shown how to represent a simple four-element adjacency list as a four-node and four edge graph drawing. Participants were given ample chance to ask questions about the SketchNode system and the process of representing an adjacency list as a graph drawing at this stage.

Besides the two experimental graphs, A and B, two practise graphs were defined, P1 (n=5,e=6) and P2 (n=8,e=8), and these were presented first. Participants were not aware that these were practise graphs – their use ensured that the participants were comfortable with the task and with the system before they drew the two experimental graphs that we were interested in. Exactly the same instructions were given to the participants for the practise graphs as for the two experimental graphs which followed: *Please draw this graph as best as you can so to make it "easy to understand"*.

The two experimental graphs A and B, were then presented to the participants, with the edges listed in a different random order for each participant. Each experiment was conducted individually, with only the experimenter and participant present. So as to control for any possible ordering effects, seven participants were given A before B, while other ten were given B before A.[1]

3.5 Participants

The data from 17 participants (numbered 3-19) was collected for analysis. Participants were friends, family and classmates of the student experimenters, and were students and non-students, of both genders. Only some of the student participants were studying Computer Science (Table 1).

Gender	7F, 10M
Current occupation	10 students, 7 non-students. None of the non-students were younger than 21.
Age	max 49, min 20, mode 22, median 24, mean 27.65
Education	13 have some university-level education, 4 have not
Experience of graphs	8 yes, 9 no
Experience of university maths or computer science	8 yes, 9 no (matching the 'experience of graphs' data)
Pen-based technology experience	4 none, 9 minimal, 3 moderate, 1 extensive.

Table 1: Demographic information about the participants

3.6 Data Collection

All interactions with the tablet were recorded using Morae software (2009), producing screen casts of all the participants' interactions with the SketchNode tool, as well as a corresponding audio track. The time taken for the drawing of graphs A and B was recorded.

At the end of the experiment, the participants were asked "Why did you arrange the graphs in the way you did?" in a recorded interview.

4 Results

4.1 The product: what do the graph drawings look like?

Of the 34 drawings produced (Figure 4), three were incorrect. 15A had an additional edge (G,H), 19A represented the node H twice, and 16B was missing the (D,C) edge: this had been drawn by the participant at the start, but had been lost in the later editing process. As the focus of the experiment was on how participants represented graphs (and not on whether they drew the

[1] These numbers (7 and 10) are not half of the total number of the participants, as the data from two participants (participants number 1 and 2) who were given A before B has been removed from the analysis - in both these cases the experiment was affected by unexpected interruption or data collection errors.

graphs correctly or not), these graphs were not removed from the analysis.

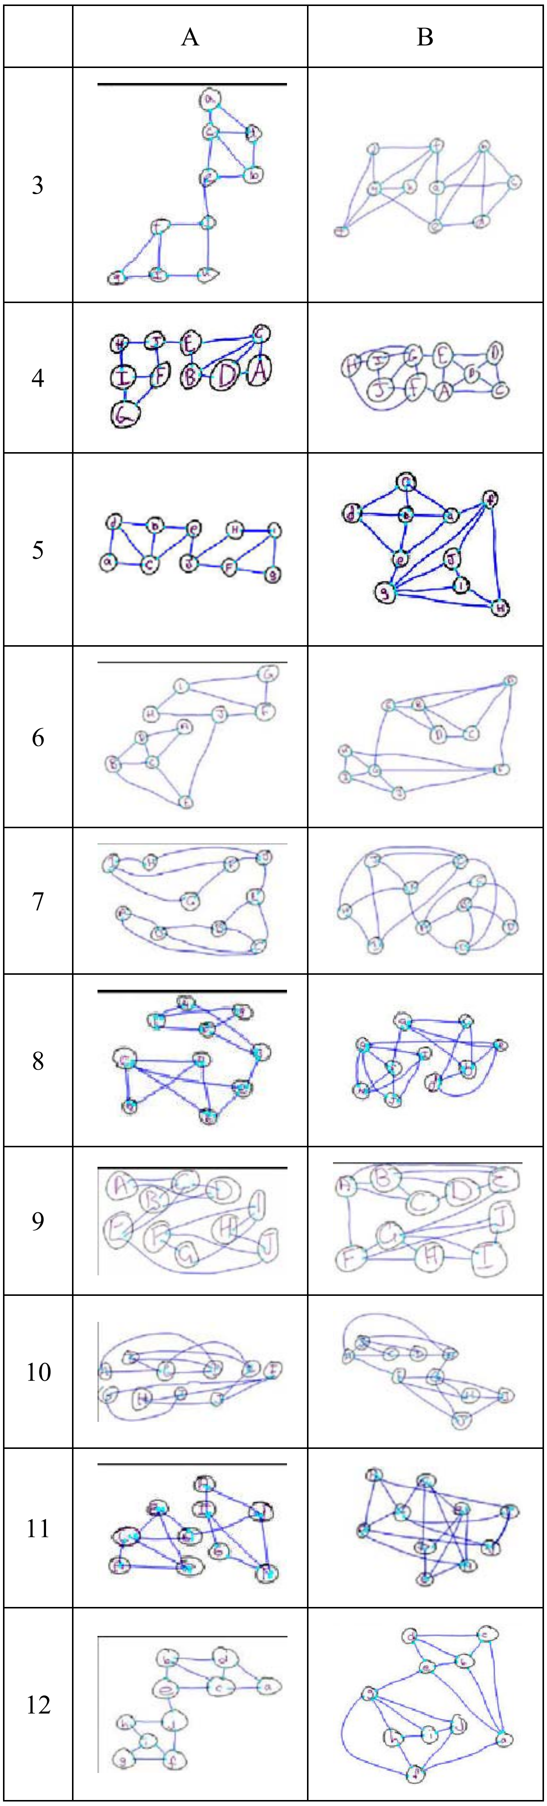

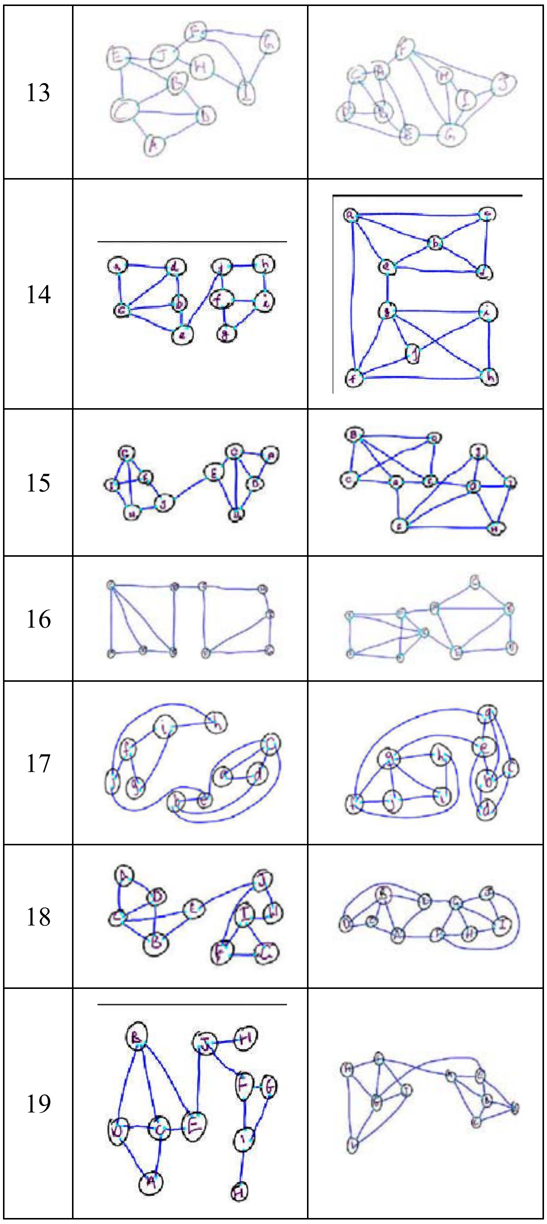

Figure 4: The graph drawings

The graph drawings were visually analysed for the following layout features (Table 2).[2]

- *Number of edge crossings*. A crossing is defined as any point outside of the node boundaries where one or more edges cross.
- *Representation of clusters*. This feature is true if a straight line can be drawn through the graph drawing so that the two pre-defined clusters in the graph can be visually separated.
- *Number of hulls*. A hull is defined as a cluster that is bounded by edges. This feature is either 1 or 2, and is not applicable if the 'Representation of clusters' feature is false.
- *Number of straight lines*. In a sketching system, the lines are unlikely to be geometrically straight, so a visual assessment of straightness was made.

[2] Our notational convention is that version of graph A drawn by participant 3 is called 3A, and the version of graph B drawn by participant 8 is referred to as 8B.

- *Number of vertical or horizontal edges*. In both cases, a visual assessment was made as to whether the edge was intended to be horizontal or visual.

All analysis was done by visual observation, rather than computationally. A detailed computational analysis of these drawings (such as that done by van Ham and Rogowitz (2008)) is beyond the scope of this paper.

The drawings revealed that even though some graphs were drawn with several edge crossings (for example, 11B and 8B), most drawings had been drawn so as to minimise crossings: 16 drawings had zero crossings, while 4 had one crossing and 4 had two. It is clear that this was the most important layout aesthetic used by the participants, and this result concurs with other empirical studies (e.g. Purchase, 1997) and with van Ham and Rogwitz' (2008) findings.

Most participants recognised the presence of the two clusters, and separated them appropriately. Van Ham and Rogowitz (2008) explicitly asked their participants to layout the graphs so that they best *...reflect[ed] their interconnections;* in our experiment we were not so explicit in our instructions regarding the connectedness of the graph – despite this, our participants still performed well when it came to grouping clustered nodes together. In addition, most of them surrounded their clusters with edges, forming clear hulls.

		GA (n=10,e=14)	GB (n=10,e=18)
Edge crossings	Mean	1.35	3
	Median	0	2
	Max	7	22
	Number of graphs with zero crossings	11	6
Clusters and Hulls	Number of graph drawings with two clusters	10	11
	Percentage of visible clusters represented as hulls	70%	100%
Vertical/ horizontal edges	Mean percentage edges horizontal or vertical	38%	29%
	Number of graphs clearly drawn with a grid arrangement in mind	6	4
Straight lines	Mean percentage of edges straight	89%	87%
	Number of graphs drawings with all lines straight	12	6
	Number of graph drawings with no edges straight [3]	1	1

Table 2: Features of the 17 drawings of graph A and 17 drawings of Graph B

[3] The two drawings with no straight edges were drawn by the same participant.

Participants were reluctant to produce final drawings with curved lines: although only 18 of the 34 graphs comprised only straight lines, the overall percentage of straight lines is high, and curved lines are typically used at the edge of the drawing to avoid edges crossing nodes or other edges. SketchNode straightens associated curved edges when nodes are moved, but this affected the nature of the final drawings for only two participants (4 and 14). In both these cases the participants drew the whole graph first before moving any nodes, indicating that they were unconcerned with the initial shape of the graph (including the curved edges) as they knew that they were going to subsequently adapt the whole drawing to make the layout more acceptable.

Approximately one third of all the edges were aligned along the horizontal or vertical axes, and ten of the 34 drawings were clearly drawn with a grid-like formation in mind. Van Ham and Rogowitz (2008) do not analyse their drawings with respect to horizontal or vertical edges, and their published examples show little evidence of this feature. It is possible that their starting configurations of circular and spring layouts, neither of which favour the presentation of edges along grid lines, may have meant that their participants were biased away from this feature.

4.2 The process: how were the graphs drawn?

Analysing the screen cast videos proved very revealing, in that it showed that the final product seldom represented the layout strategy used by the participant when drawing the graph, and that aesthetic criteria emphasised by participants early in the creation process were often compromised as the graph grew in size.

The videos of the creation process were analysed for the following features (Table 3):

- The order and timing of drawing the nodes.
- If and when nodes or sub-graphs were relocated.
- Use of straight or curved lines.
- The variation in the length of edges.
- Alignment to a horizontal/vertical grid.
- Evidence of participants analysing the adjacency list and planning ahead before drawing.

As expected, Graph B took significantly longer to draw than Graph A, as it had more edges. However the time taken to draw the graphs varied considerably between participants (Figure 5). Analysis of the screen casts revealed that there was seldom a clear break between the process of creating the drawing (i.e., representing all the information in the adjacency list) and the process of laying the graph drawing out so as to make it "easy to understand." In most cases, node positioning decisions were made during the creation process. This means that no timing data could explicitly be associated with the process of graph layout. This is unlike the research of Van Ham and Rogowitz, (2008) who, because graph creation was not part of the task that they set their participants, have clear data on the time taken for the graph layout process and are therefore able to make layout time comparisons between different graphs and different initial layouts.

	Strategy	Number of participants (out of 17)	Notes
Drawing nodes	Draw all nodes first	6	Two participants changed their strategy: for their first graph, they drew nodes where convenient, for their second graphs, they drew all the nodes first.
	Draw nodes where convenient for the placement of the next edge	11	One participant used a mix of the strategies
Moving nodes	Move nodes during creation of the graph	10	One participant changed strategy: in the first graph, he moved nodes and sub-graphs during creation; for the second graph, he didn't.[4]
	Move sub-graphs during creation of the graph	9	One participant moved nodes during creation, but not sub-graphs
	Move nodes after creation	8	
Lines	Mainly straight lines	14	This process feature is derived independently of the number of straight lines in the final product, as SketchNode straightens curved edges when nodes are moved.
	Only use curved lines to avoid crossings	10	These participants favoured straight lines, and only introduced curved lines when necessary
	Use similar length lines during creation	11	
Alignment	Favour horizontal and vertical edges during creation	10	
Analyse	Plan ahead	5	This was evidenced by excessive pauses, or by self-reporting

Table 3: Features of the graph drawing process, showing how many participants adopted the different strategies

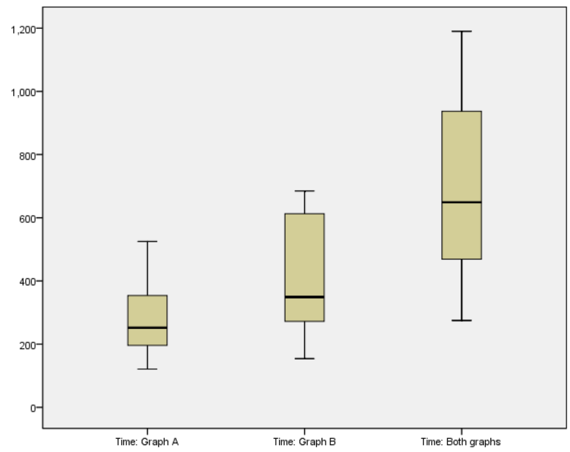

Figure 5: Box plots showing the median and range (in seconds) of the time taken for the participants to draw the graphs.

Figure 6 shows snapshots of the creation of three of the drawings, each demonstrating a different strategy. 12A has all the nodes drawn first, and then moved around as necessary, so as to allow for easy insertion of edges; 13B has the nodes drawn wherever convenient, avoiding crossings at the start, introducing crossings as the graph becomes more complex, and then relocating nodes at the end so as to remove crossings; 17A has nodes placed where convenient, and never relocated, with curved edges used so as to avoid edge crossings.

There were clearly two different strategies in the creation of the graph drawing: drawing all the nodes first and then moving them as necessary when creating edeges, or drawing nodes where convenient. The latter was favoured. Participants tended to place the next node they came across in the adjacency list near to the node it needed to be connected to so as to avoid edge crossings.

There were also two different strategies when it came to producing the layout of the graph: movement of nodes during creation of the graph, and movement after the whole graph was complete. The former was favoured. There were even a few participants who did not move the nodes at all once they had been placed conveniently next to their first adjoining node (e.g. participants 7, 8 and 10), even if this meant introducing edge crossing or curved edges.

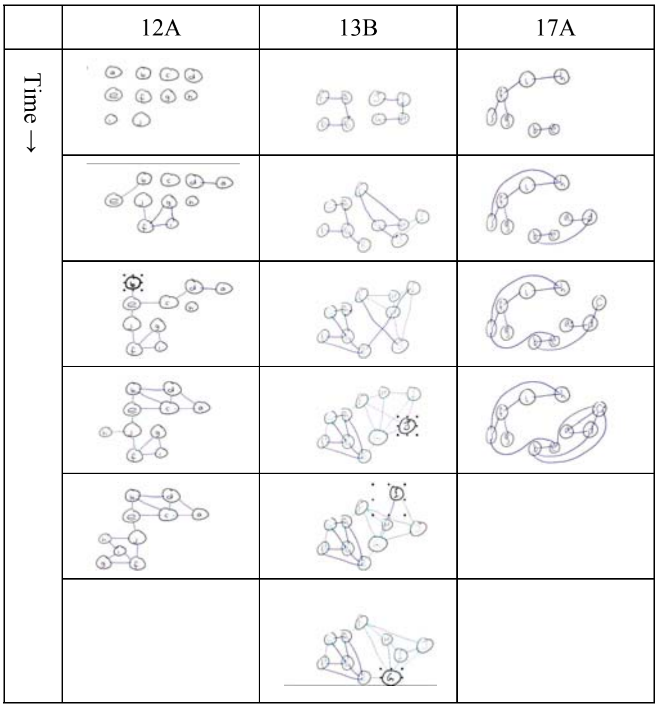

Figure 6: Snapshots of the graph creation process for three graphs, showing three different strategies.

The use of these different strategies, and the fact that only 8 participants moved nodes after the graph was created, suggest a reluctance on the part of some participants to move nodes once placed, and a possible misunderstanding of the task (which was to arrange the final graph so as to make it easy to understand). It appeared that some participants may have focussed on ease of process (i.e. ease of creating the graph) rather than on the form of the final product. While our experiment removed the layout bias of van Dam and Rogowitz' (2008) work by asking participants to draw the graph

[4] This participant is one of the two who changed their node drawing strategy.

from scratch, asking them to effectively do two tasks (creation and layout) makes it more difficult (and in some cases impossible) to focus our analysis on the process of layout.

Analysing the participants' drawing and layout processes allowed us to see where layout aesthetics had been favoured during the creation process, but abandoned later when the graph became more complex. These aesthetics are therefore not evident in the final product, even though there were considered important. Most noticeable of these was the tendency for participants to place their nodes as if on the intersections of an underlying unit grid, and to draw edges vertically and horizontally. So while only 10 of the resultant graph drawings had evidence of a grid-like formation (representing 6 participants), a further four participants attempted to use horizontal and vertical lines, but abandoned this feature as more edges were added.

Similarly, most participants attempted to use only straight edge, and the curved edges tended to be the very last edges to be drawn. Most participants also attempted to use edges of similar length: a feature that is evident in very few of the resultant drawings.

Five students performed some pre-drawing analysis and planned ahead. Planning ahead is, of course, unnecessary, as any nodes could be moved to a more appropriate position at any time. This again suggests a reluctance of the participants to move nodes and change the overall layout of the drawing.

The appendix shows snapshots of the creation of two versions of graph A (3A and 3B) and two versions of graph B (15B and 16B) both of which have similar final products.

4.3 The preferences: what did the participants think?

In the post-experiment interview, the participants were encouraged to articulate their strategy in drawing the graphs, both in terms of the product and the process. No specific features of graphs were suggested to the participants in this interview, and they were encouraged to describe the features in their own words.

All but four of the participants emphasised the need to avoid edge crossings; three of these said that avoiding crosses was not important, one of whom said that edge crossings are acceptable if the cross is at right angles. Two of those who said that edge crossing were not important specifically said that avoiding edges crossing nodes was more important.

Four participants mentioned a preference for horizontal or vertical alignment of edges, with one also favouring 45 and 60 degree angles. Five participants said that straight lines were preferred, while one specifically said that they were not important. Other features mentioned were maintaining a similar distance between nodes (2), spreading the nodes out (2), putting the nodes in groups (clusters) (2) and symmetry (1).

With respect to the graph creation process, four participants specifically mentioned their use of the node with highest degree: two participants placed it centrally, one participant placed it at the top (while mentioning a desire to draw the graphs in a tree-structure), and another participant just made sure that the node with the highest degree was drawn first.

Two participants said that they preferred to place the nodes in alphabetical order at the beginning of the drawing process.

When asked why they drew the graphs the way that they did, most participants used phrases like 'most logical' (3), 'easier' (4), and 'neater' (4). Two participants specifically said that they looked ahead, making room for future edges that would be added.

4.4 Demographic effects

There are insufficient data points for a complete statistical analysis, but we can make some informal observations on product, process and preferences when considered with respect to the important demographic properties of Computer Science/graph experience, and gender, both of which have appropriate splits to allow for reasonable analysis (Table 4).

The data suggest that the male participants favoured straight lines, horizontal and vertical lines and a grid formation (both during the process and in the final product) more than the female participants. There was no obvious gender difference in strategy, edge crossings or clusters.

	Gender		CS/graph experience	
	F (n=7, 14 graphs)	M (n=10, 20 graphs)	Yes (n=8,16 graphs)	No (n=9, 18 graphs)
Product				
Number of crosses (mean)	2.21	2.15	0.75	3.44
Percentage lines straight	76%	96%	98%	79%
Percentage lines horizontal/vertical	23%	41%	45%	23%
Number of drawings aligned to a grid	2	8	9	1
Number of clusters (mean) (max 2)	1.1	1.3	1.5	1.0
Process (number of participants, out of 17)				
Drawing all nodes first	2	4	5	1
Applying layout after creation	3	5	6	2
Aligning to a grid during creation	2	8	6	4
Planning in advance	3	2	1	4
Mean time for both graphs	11m36s	11m35s	13m7s	10m32s
Preferences (number of participants stating preferences, out of 17)				
No crosses	5	8	8	5
Grid configuration	1	3	2	2
Straight lines	2	3	1	4
Special use of node with highest degree	2	2	4	0

Table 4: The graph drawing product, process and preferences with respect to the demographics

Computer Science and graph experience seems to have made a difference in several factors: reduction in the number of edge crossings, increased percentage of straight edges and number of clusters, as well as a higher instance of horizontal and vertical edge and grid-like graph drawings. However, it is interesting to note that three of the participants who did not have Computer

Science experience aligned nodes and edges to an underlying grid while creating the drawings, even if their final graph drawings do not show evidence of this feature being favoured. Participants with Computer Science and graph experience tended to explicitly choose a strategy (drawing nodes first or laying out all nodes at the end), while the others were more likely to pause to plan their actions during the creation process. Only participants with Computer Science and graph experience recognised the importance of the node with highest degree. Non-Computer Science participants tended to perform the graph drawing tasks more quickly.

5 Discussion

5.1 Layout aesthetics and the graph drawing process

With respect to the layout aesthetics of the resultant graph drawing, it is not surprising that the minimisation of edge crossings has again been revealed as most important. The data presented here gives more weight to the principle of fixing edges and nodes to an underlying unit grid than has previously been shown to be the case in empirical studies of graph drawing comprehension (Purchase, 1997). This latter point is surprising, suggesting that users would like to see their graphs fixed to a grid, even if doing so does not necessarily assist in improving their performance in graph reading tasks.

Only one of the published graph drawings from the research of van Ham and Rogowitz (2008) shows any indication that the participants were interested in horizontal or vertical edges: we suggest that this is due to the layout bias (circular or spring) in the initial drawing presented to the participants.

When considering the creation strategies, two particular methodology issues stood out. First, a final graph drawing often does not reveal the aesthetic principles that the participants have tried to adhere to during the drawing process, and much can be missed by concentrating on the product rather than the process. Second, researchers in the area of Graph Drawing tend to make a clear distinction between the process of creating a graph and the algorithm for laying it out. This does not appear to be the case for our participants: the action of presenting relational information in as best a way as possible using a graph drawing encompasses both the initial visual representation of the nodes as well as their relocation. The fact that we cannot separate the process of creation from the process of layout in our data is a confounding factor when comparing our timing data of with that of van Ham and Rogowitz (2008), as in their case the participants' task was purely one of layout.

Informal demographic analysis suggests a tendency for male participants with some Computer Science experience to particularly favour a grid layout.

5.2 Limitations and Future work

While this experiment has successfully addressed some of the limitations of the work of van Ham and Rogowitz (2008) (as described in section 2 above) it has itself been subject to some methodological constraints. We only have 17 participants and 34 graphs, which makes any extensive demographic analysis and generalisation of results difficult; the strategy employed by most of our participants was to interleave the creation and layout processes, making it difficult to analyse the process of layout separately, and both our graphs were small. Future work would, of course, entail the drawing of larger graphs (possibly using a digital whiteboard or similar technology), and would attempt to clarify to the participants the usefulness of manipulating the nodes in the graph after they have been created.

No computational analysis has been performed on this data; such analysis could, for example, determine (as in van Ham and Rogowitz (2008)) the edge-length distribution, the orientation and the extent of clustering as derived from inter-node differences.

6 Conclusion

The van Ham and Rogowitz (2008) experiment considered how users would manipulate graph drawings so as to improve their layout. The research reported here extends this work significantly with the use of sketch-based graph drawing software, by removing any initial layout bias, and by including consideration of key participant demographics. While some results (e.g. the importance of minimising crosses) clearly follow several previous findings, we have also established the importance of a grid-based layout, and shown that the layout of a graph drawing should not simply be judged by the product, but should also be considered in the light of the process that created it.

7 Acknowledgements

We are grateful to the experimental participants, and to Hong Yul Yang, who implemented SketchNode. Ethical clearance for this study was granted by the University of Auckland, 2008.

8 References

Battista, G.D., Eades, P., Tamassia, R. and Tollis, I.G. (1998): Graph Drawing: *Algorithms for the Visualization of Graphs.* Prentice Hall.

Coleman, M.K. and Stott Parker, D. (1996): Aesthetics-based graph layout for human consumption. *Software—Practice and Experience,* 26(12):1415-1438.

Gansner, E.R., Koren, Y. and North, S. (2005) Graph Drawing by Stress Majorization. *Proceedings of 12th International Symposium on Graph Drawing,* 239-250. Springer Verlag.

Huang, W., Hong, S. & Eades, P. (2006): Layout effects on sociogram perception. *Proceedings of 13th International Symposium on Graph Drawing,* 262-273. Springer Verlag.

Huang, W. (2007): Using Eye Tracking to Investigate Graph Layout Effects. In Proc. *Asia Pacific Symposium on Visualization* (APVIS2007), IEEE, 97-100.

Huang, W., Hong, S.-H. and Eades, P. (2007): Effects of Sociogram Drawing Conventions and Edge Crossings in Social Network Visualization. *Journal of Graph Algorithms and Applications* . 11 (2): 397-429.

Morae Usability Software. www.TechSmith.com/Morae. Accessed 11th August 2009.

Purchase, H.C., Cohen, R.F. and James, M.I. (1995): Validating Graph Drawing Aesthetics. *Proceedings of 3rd International Symposium on Graph Drawing,* 239-250. Springer Verlag.

Purchase, H.C. (1997): Which Aesthetic has the Greatest Effect on Human Understanding? *Proceedings of the 5th International Symposium on Graph Drawing.* 248-261. Springer Verlag.

Reid, P., Hallett-Hook, F., Plimmer, B and Purchase, H.C. (2007): Applying Layout Algorithms to Hand-drawn Graphs, *OzCHI,* 203-206.

van Ham, F. & Rogowitz, B.E. (2008) :Perceptual Organization in User-Generated Graph Layouts. *IEEE Transactions on Visualization and Computer Graphics.* 14(6): 1333-1339.

Viegas, F,B., Wattenberg, M., van Ham, F., Kriss, J., and McKeon, M. (2007): ManyEyes: a Site for Visualization at Internet Scale. *IEEE Transactions on Visualization and Computer Graphics.* 13(6):1121-28.

Ware, C., Purchase, H.C., Colpoys, L. and McGill, M. (2002): Cognitive measurements of graph aesthetics. *Information Visualization.* 1(2):103-110.

9 Appendix

Graph drawings 3A and 3B are similar, as are 15B and 16B. However, the process of drawing in each case is different, as shown by the snapshots shown below.

	3A	4A
Time →	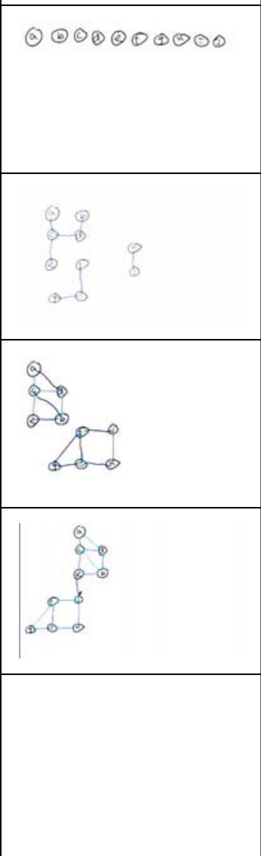	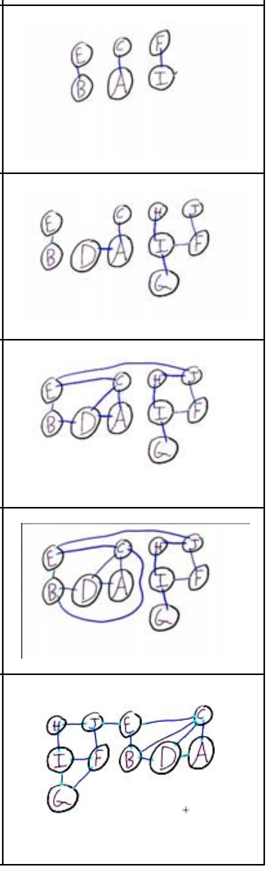

	15B	16B
Time →