Feedback About 2018's Coursework in
MATH20902: Discrete Mathematics

I made the notes below while marking the coursework. I also made remarks on individual papers and would be happy to discuss these. The marked work is available at the Reception desk near the entrance to the Alan Turing Building.

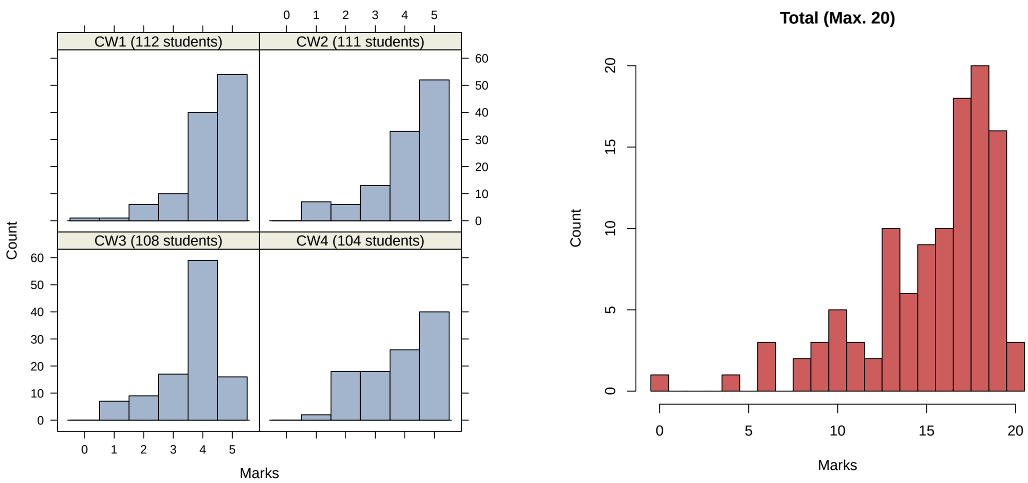

Figure 1: *Histograms for the individual questions and for the total score (right). There were 112 pieces of coursework and the average mark was* ≈ 15.3 *out of 20, or just over 76%.*

Overall Remarks

- Generally speaking, people did well: I saw a lot of really good arguments and examples while marking this work.
- I marked these papers partly with an eye to good mathematical style (see the question-by-question notes below for details). To get full marks you needed to make an argument that was technically correct, clear and concise. The last bit is especially important: long answers are not necessarily better ones. Most of the top-scoring papers used only three or four sides of A4.

Remarks about individual problems

(1) (Testing for stability: 5 marks).
As I'd hoped, this proved to be the easiest question: the average was 4.2 out of 5 marks.

(a) Almost everyone got at least one of the two available marks here, but I often deducted the second mark if the answer mentioned the Gale-Shapley algorithm or some process of choosing partners. The point is that the number of possible matchings follows from the fact that a matching is a bijection between two sets of size n and so there is one matching for each element of the permutation group S_n.

(b) This part proved harder to mark than I'd expected as quite a few people—perhaps 10%—generated all 24 possible matchings in some idiosyncratic way that I found hard to follow. As the list of possible matchings was rather long, I tended not to deduct marks for the odd mis-classified result. A few students did as I had, and wrote computer programs to list all matchings and, for those that proved unstable, all blocking pairs: see Table 1.

Arti	Betty	Carol	Deb	Stable?	Blocking Pairs
Rob	Seb	Tom	Ulf	Yes	NA
Rob	Seb	Ulf	Tom	Yes	NA
Rob	Tom	Seb	Ulf	No	(Betty, Rob), (Betty, Ulf), (Carol, Ulf), (Carol, Rob)
Rob	Tom	Ulf	Seb	No	(Betty, Rob), (Betty, Ulf), (Deb, Tom)
Rob	Ulf	Tom	Seb	No	(Betty, Rob), (Deb, Tom)
Rob	Ulf	Seb	Tom	No	(Betty, Rob), (Carol, Rob)
Seb	Rob	Tom	Ulf	Yes	NA
Seb	Rob	Ulf	Tom	Yes	NA
Seb	Tom	Rob	Ulf	No	(Betty, Ulf), (Carol, Ulf)
Seb	Tom	Ulf	Rob	No	(Betty, Ulf), (Deb, Tom)
Seb	Ulf	Tom	Rob	No	(Deb, Tom)
Seb	Ulf	Rob	Tom	Yes	NA
Tom	Seb	Rob	Ulf	No	(Arti, Seb), (Carol, Ulf)
Tom	Seb	Ulf	Rob	No	(Arti, Seb), (Deb, Tom), (Deb, Seb)
Tom	Rob	Seb	Ulf	No	(Arti, Seb), (Carol, Ulf), (Carol, Rob)
Tom	Rob	Ulf	Seb	No	(Arti, Seb), (Deb, Tom)
Tom	Ulf	Rob	Seb	No	(Arti, Seb), (Deb, Tom)
Tom	Ulf	Seb	Rob	No	(Arti, Seb), (Carol, Rob), (Deb, Tom)
Ulf	Seb	Tom	Rob	No	(Arti, Seb), (Arti, Tom), (Deb, Tom), (Deb, Seb)
Ulf	Seb	Rob	Tom	No	(Arti, Seb)
Ulf	Tom	Seb	Rob	No	(Arti, Seb), (Carol, Rob), (Deb, Tom)
Ulf	Tom	Rob	Seb	No	(Arti, Seb), (Deb, Tom)
Ulf	Rob	Tom	Seb	No	(Arti, Seb), (Arti, Tom), (Deb, Tom)
Ulf	Rob	Seb	Tom	No	(Arti, Seb), (Carol, Rob)

Table 1: All possible matchings and blocking pairs for Problem (1), presented as a list of partners for the women. Each row corresponds to a possible matching and, for those that aren't stable, includes a list of all blocking pairs.

When students lost marks in here it was typically because of confusion about what a blocking pair is: a couple (m, w) is a blocking pair if they are *not* currently matched to each other, but would prefer to be. Thus any answer that lists an actually-existing couple as a blocking pair has to be wrong.

(2) (**Number of proposals: 5 marks**).
This question also went well for most students: the average was just under 4.1 out of 5.0.

(a) Most people got both of the marks available here, but a few wrote down preference lists that lead to three *rounds* of proposals rather than three proposals in total and so lost marks.

(b) Most people used Proposition 5 from the assignment to say that $N \leq n^2 - n + 1$, then either worked directly from the definition of $f(n) = O(g(n))$ or used a result from the problem sets to say that any polynomial of degree d is $O(n^d)$. But a number of people presented a lovely alternative argument saying that, as $|\mathcal{M}| = |\mathcal{W}| = n$ and no one ever proposes to the same partner twice, we have automatically that $N \leq n \times n$.

I marked this rather section rather strictly in that, if a student wrote

$$N < c_1 n^2 \text{ for all sufficiently large } n,$$

I wanted to see an explicit value for c_1 and a number n_0 such that "sufficiently large n" meant $n \geq n_0$: here $c_1 = n_0 = 1$ works. If I didn't see these things, I deducted a mark.

(3) (Impossible partners: 5 marks).

This question was the hardest, though the average was still a respectable 3.6 out of 5.0.

(a) Many students lost one of the two marks available for this part because they did not reduce the preference lists *as far as possible.* Although straightforward application of Corollaries 3 and 4 allows one to use matching

(Rob, Carol), (Seb, Arti), (Tom, Betty) and (Ulf, Deb)

produced when the men propose in the Gale-Shapley algorithm to reduce the preference lists as follows

Men's Preferences

Rob	Seb	Tom	Ulf
Carol	Arti	~~Deb~~	Deb
Arti	Carol	Betty	Carol
Betty	Deb	Carol	Arti
Deb	Betty	Arti	Betty

Women's Preferences

Arti	Betty	Carol	Deb
Tom	Seb	Ulf	Ulf
Seb	Ulf	Rob	~~Seb~~
~~Rob~~	Tom	~~Tom~~	~~Tom~~
~~Ulf~~	~~Rob~~	~~Seb~~	~~Rob~~

it's possible to eliminate many more potential partners. The first thing to note is that impossibility is a symmetric relationship: if, as the tables above say, Rob is impossible for Arti (there are no stable matchings in which they are partners), then Arti is impossible for Rob as well. This observation allows one to reduce the preference lists to:

Men's Preferences

Rob	Seb	Tom	Ulf
Carol	Arti	~~Deb~~	Deb
~~Arti~~	~~Carol~~	Betty	Carol
~~Betty~~	~~Deb~~	~~Carol~~	~~Arti~~
~~Deb~~	Betty	Arti	Betty

Women's Preferences

Arti	Betty	Carol	Deb
Tom	Seb	Ulf	Ulf
Seb	Ulf	Rob	~~Seb~~
~~Rob~~	Tom	~~Tom~~	~~Tom~~
~~Ulf~~	~~Rob~~	~~Seb~~	~~Rob~~

Finally, if one uses the Gale-Shapley algorithm with the women as proposers, one finds the following matching

(Rob, Carol), (Seb, Betty), (Tom, Arti) and (Ulf, Deb).

Applying Corollaries 3 and 4 then yields the following, fully-reduced preference lists

Women's Reduced Preferences

Arti	Betty	Carol	Deb
Tom	Seb	Rob	Ulf
Seb	Tom		

Men's Reduced Preferences

Rob	Seb	Tom	Ulf
Carol	Arti	Betty	Deb
	Betty	Arti	

where I have removed impossible partners completely.

(b) The majority of students proved Corollaries 3 and 4 by contradiction, which is the simplest approach. Less successful strategies included

- Trying to prove Theorems 1 and 2 by repeating or reinterpreting the proofs in either Gale and Shapley's original paper or the book by Gusfield and Irving. This was much harder than what the question actually asked.
- Several students seemed to think that only two stable matchings existed—one produced when the men are proposers and the other when the women propose. There can be many, many than this, as the example from Problem 1 suggests.

(4) (Matching in labour markets: 5 marks).
The average on this question was 3.8 out of 5.0. A surprisingly large number of students seemed to think that in the labour market described in the problem, the employers are the proposers. This is clearly not the case because, at least for entry-level positions, people are seldom head-hunted: a student will only get a job offer from a firm to which she has applied. Perhaps the root of this confusion is the sense of powerlessness one can feel as a job applicant: it's hard to believe that the labour market is going to deliver an applicant-optimal stable matching.

The two most successful strategies were:

- Assume for contradiction that there is a stable matching in which an excellent graduate gets assigned a job outside Manchester: this fairly quickly leads to a blocking pair, contradicting the assumed stability of the matching. Students who took this approach often got full marks, but I sometimes deducted a mark or two if there was either (a) extraneous stuff not relevant to the argument or (b) assumed that all stable matchings must come from the Gale-Shapley algorithm.

- Think carefully about the Gale-Shapley algorithm and argue that, whoever proposes, it will yield a matching in which all the Manchester jobs go to the excellent graduates. Then use Corollaries 3 and 4 to establish that there can be no stable matchings, whether generated by Gale-Shapley or not, in which Manchester jobs go to ordinary (that is, non-excellent) graduates.

  This approach proved harder in practice. Typical mistakes included:

  - Writing preference lists such as the one below

Businesses			*Graduates*	
E	O		M	N
M	M		E	E
N	N		O	O

    Here E represents the excellent graduates, O the ordinary ones, M the Manchester jobs and N those outside the city. The problems with this table is that it lacks detail. It's not clear whether the person who wrote it appreciates that although the Manchester employers prefer the excellent graduates, their preferences need not all be the same: each Manchester business can rank the members of E differently, provided only that they place all of them ahead of anyone from O.

  - Assuming that all the employers and/or all the students have identical preferences.

  - Changing the problem, for example:

    Assume without loss of generality that job offers are firm (that is, the engagement can't be broken).