Feedback About 2019's Exam in
MATH20902: Discrete Mathematics

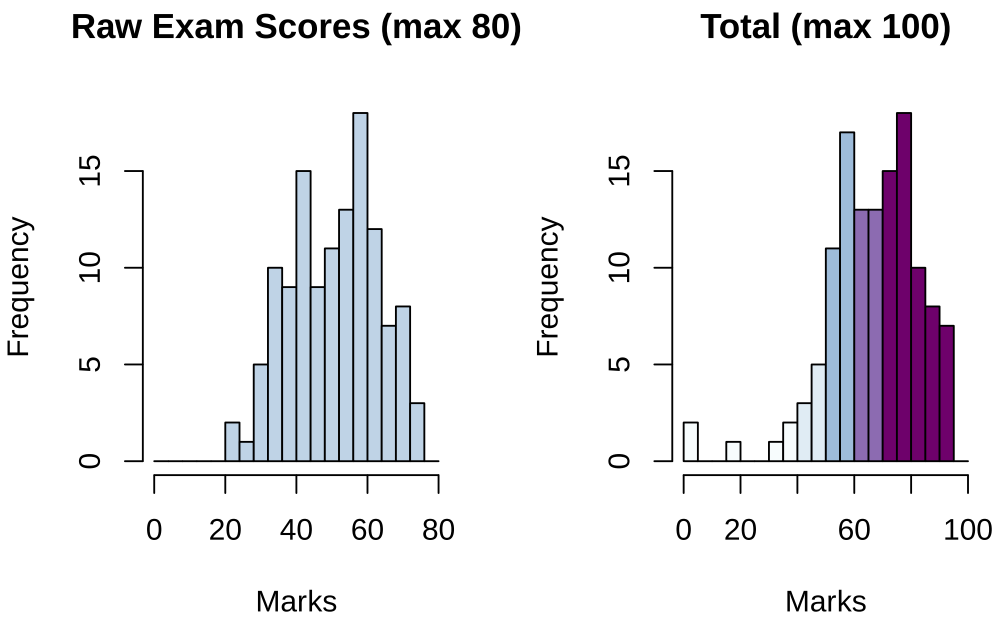

Figure 1: *Histograms of the raw exam results (left panel) and final marks (right panel). The histogram of total marks (where "total" means exam plus coursework) has been shaded to indicate degree classes. These plots include data for all 123 students who sat the exam.*

General remarks

- The final marks summarised in the panel at right above have not been adjusted to take account of Mitigating Circumstances nor have they been moderated (that is, scaled) to make them more comparable to other exams in the School.
- Many people did very well on the exam: although there were no perfect papers, 18 students had exam marks of 65 or better (where the maximum possible was 80). Here is a summary of the final marks, by degree class.

Result:	First	2(i)	2(ii)	Third	Fail
Range:	70–100	60–69	50–59	40–49	0–39
Number of students:	58	26	28	8	3
Fraction of students:	47.2%	21.1%	22.8%	6.5%	2.4%

- Around a sixth of all students attempted all three of the part B problems. For these students only the best two part B scores contributed to the total exam mark.

Part B problems attempted:	B4 & B5	B4 & B6	B5 & B6	All
Number of students:	33	61	8	21

- These notes about individual problems easier to follow if you have a copy of the exam.

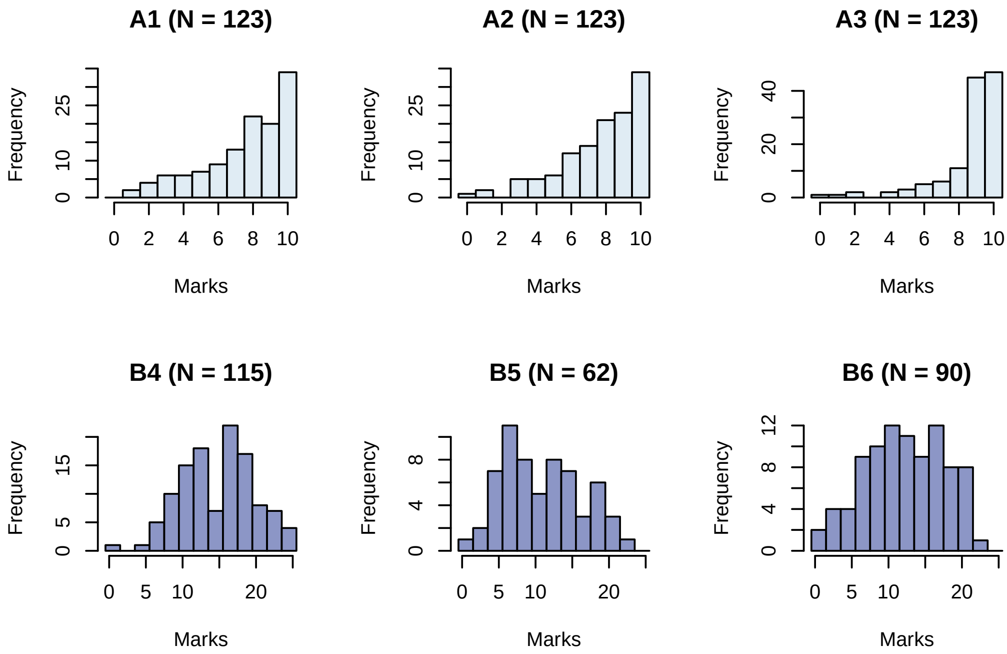

Figure 2: *Histograms for the individual questions: note that the vertical scale varies from question to question. Those for the required Part A questions (10 marks apiece) are plotted light blue, while those for the Part B questions (best two out of three, 25 marks apiece) are slightly darker. In the latter group, some of the very lowest marks come from students who attempted all three part B questions and so these scores didn't usually contribute to the total exam mark.*

Remarks about individual problems

A1 People did well on this question: the average was around 7.5, with 34 perfect scores.

(a) Almost everyone got the 3 marks associated with the definitions, though a handful of students defined the degree sequence as "the set of the degrees of the vertices …" when, in fact, it is a list of the degrees. This is an important distinction as we know, from the problem sets, that the degree sequence of a graph always includes at least one repeated entry, while a set cannot have two elements with the same value.

As seems to happen every year, a few students lost a mark in the definition of isomorphism for saying

$$\ldots (\alpha(u), \alpha(v)) \in V_2 \Rightarrow (u, v) \in V_1.$$

instead of

$$\ldots (\alpha(u), \alpha(v)) \in V_2 \text{ if and only if } (u, v) \in V_1.$$

and, finally, quite a few students included odd, extraneous remarks in the definitions of isomorphic: things such as $|V_1| = |V_2|$, which are consequences of the definition, but not part of it. I didn't deduct any marks for this, but it is bad mathematical style.

(b) Most students realised that the statement in the exam is false and offered the counterexample below

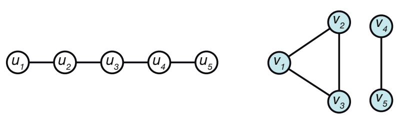

This accounted for two of the four marks available here, while the remaining two were for the proof, which most students handled by contradiction. A surprisingly large minority of students drew examples of multigraphs here, and so got little or no credit.

(c) Most students realised that $(1, 3, 4, 4, 4)$ can't be the degree sequence of any graph, and almost everyone managed to prove this by contradiction. Here again, there were occasional odd statements about multigraphs

A2 Perhaps because similar questions have appeared on many recent exams, this proved a fairly easy question: the average was 7.7/10.

(a) As in question A1, most students got most of the marks available for definitions. Those who lost marks made two main kinds of mistakes.

- Some neglected to explain key terms such as *rooted* or *spanning*.
- Some gave definitions that were too terse. So, for example, it was not enough to say

  > A *spanning arborescence* in a directed graph $G(V, E)$ is an arborescence that includes all the vertices in V.

  without explaining what an arborescence is. Similarly, it's not enough to say

  > A directed graph $T(V, E)$ is an *arborescence* if T contains a root and the graph $|T|$ that one obtains by ignoring the directedness of the edges is a tree.

  without saying what a *root* is. On the other hand, you wouldn't need to define *tree*.

  The principle here is that if a group of definitions appear in quick succession (those for *root*, *arborescence* and *spanning arborescence* are Definitions 7.6–7.8 and all appear one after the other in the printed lecture notes), then you need to spell out any and all of them that you use. So, for example, I awarded full marks to students who wrote something like

  > An *arborescence rooted at v* is a directed graph $T(V, E)$ in which all other vertices are reachable from v and which becomes a tree when one ignores the directedness of the edges.

  This definition doesn't use the term *root*, but instead explains that the vertex v must have the properties of a root. On the other hand, it does use the terms *reachable* and *tree* without further explanation, but that's okay as both were defined in earlier lectures.

(b) One can count spanning arborescences either with Tutte's Matrix Tree Theorem (and it's good to explain that this is what you're doing, though I didn't deduct any marks for correct answers that didn't include such an explanation) or by drawing all associated spregs and then counting those that are spanning arborescences: both approaches received full marks when correct. The Matrix-Tree strategy was both more common and more generally successful, though there were a lot of small mistakes in the evaluation of $\det(\hat{L}_5)$.

(c) Here I wanted students to either use the result from Proposition 9.4, which is about spregs that contain a single directed cycle, or to just use the approach from lecture to enumerate the relevant spregs, noting that the presence of the cycle (v_3, v_4, v_3) fixes the predecessors of v_3 and v_4, so that the total number of spregs with distinguished vertex v_5 that contain the given cycle is $\deg_{in}(v_1) \times \deg_{in}(v_2) = 1 \times 3 = 3$. Most successful answers used the latter approach, though some students just tried to draw the relevant spregs and, often, found only two of the three.

A3 Perhaps because we covered this topic in the very last lecture of the term, marks on this question weren't as high as they have been in recent years. The average was around 8.6/10 and 92 students got marks of either 9 or 10. The main ways people lost marks were by:

- in part (a), representing the problem with some graph other than the sort used in lecture and not explaining how this non-standard graph worked;
- also in part (a), drawing an undirected graph, even though the problem asked explicitly for a directed one (one mark deducted) or drawing a graph with no edge weights and no other indication of the times required for the various tasks;
- failing to indicate the critical paths in part (b). There were two, S-E-G-H-I-Z and S-E-F-C-D-I-Z, leading to a minimum time-to-completion of 29 days. A few students found this result by exhaustive enumeration of all paths and received full credit.
- As all vertices except A and B lie on one of the critical paths, part (c) was pretty easy: the only ways students lost marks here were by not answering the question (that is, not giving earliest and latest starts for each vertex) or by making a numerical mistake.

B4 This was the most popular of the Part B problems: 115 out of 123 students tried it. Answers varied wildly in quality, from 17 scripts with a score less than 10 on up to 14 scripts with scores in excess of 20: the average was 14.9/25.

(a) In contrast to Part A, where most students gave correct definitions and got all the associated marks, here there were more mistakes. Examples include:

- Definitions of k-colouring that didn't make any mention of k. A typical mistaken answer of this kind began

  A *k-colouring* of a graph $G(V, E)$ is a map $\phi : V \to \mathbb{N}$ ...

  Some better, but still incorrect answers began

  A *k-colouring* of a graph $G(V, E)$ is a bijection $\phi : V \to \{1, \cdots, k\}$ ...,

  but this is only true if $\chi(G) = |V|$. Finally, I saw a few mistaken answers that swapped the roles of $\{1, \cdots, k\}$ and V:

  A *k-colouring* of a graph $G(V, E)$ is a map $\phi : \{1, \cdots, k\} \to V$ ...
- A claim, in the definition of a k-colouring, that "$\phi(u) \neq \phi(v)$ if and only if $(u, v) \in E$". The right thing to say is that $(u, v) \in E \Rightarrow \phi(u) \neq \phi(v)$.
- In the definition of a *tree*, some students said only that $|E| = |V| - 1$, which is true of trees, but not sufficient to characterise them: think of a graph with two connected components, one isomorphic to K_2 and the other to C_3. This graph has $|E| = 4 = |V| - 1$, but is not connected and so not a tree.

(b) Many, perhaps most, students gave a correct degree sequence,

$$(1, 1, 1, 1, 1, 1, 1, 2, 2, 2, 3, 3, 3, 4),$$

and drew a suitable example (there are lots), but failed to prove—or offered an incorrect proof—that the degree sequence above, with 7 leaves, is the only possibility. Mistaken proofs often included incorrect remarks about evenness and oddness of degrees or false claims such as "a tree with 7 internal nodes must have 7 leaves", for which a counterexample is the path graph P_9

(c) The most successful answers began with a proof of the following lemma:

If $G_1(V_1, E_1)$ and $G_2(V_2, E_2)$ are isomorphic and G_1 is k-colourable, then so is G_2.

which establishes that $\chi(G_2) \leq \chi(G_1)$. I also saw a number of attempts at proofs by induction—typically induction on $\chi(G)$—but it was hard to see how to handle the inductive step. That is, it's hard to see how, given a pair of graphs with $\chi(G_j) = n_0 + 1$, one can find a systematic way to modify it to produce a pair G'_j with $\chi(G'_j) = n_0$.

(d) Almost everyone found the correct chromatic numbers here, though there were quite a few claims of the form

> By greedy colouring, $\chi(G) \leq 2$ ...

that didn't actually exhibit an explicit 2-colouring and so lost marks. There were also a surprising number of claims that Q contains a subgraph isomorphic to K_3, even though this isn't true. Also, quite a few students wrote "P is I_3 ..." or, sometimes "P is I_2 ...". Both these statements are false[1]. The first remark would be correct if it said "P is isomorphic to I_3" and so I tended to be generous about this mistake, but the second version, involving I_2, is just plain wrong and I deducted a mark.

(e) Here P and R are isomorphic and, for full credit, I wanted to see an *explicit* bijection between their vertex sets. Quite a few students claimed, mistakenly, that all three of P, Q and R are isomorphic because they all have the same number of edges, vertices and the same degree sequence.

B5 This was the least popular of the Part B problems, with only 62 attempts out of a possible 123. The average was 11/25.

(a) The definitions went well for the most part.

(b) I had hoped this would be easy as it is similar to problems that I suggested you study and it did, indeed, go well. The main ways that students lost marks were as follows.

- Drawing and colouring a graph without explaining the construction. The issue is that the question asks you to "Explain your approach thoroughly ..." and so an answer that does not include something like

  > One begins by constructing a graph whose vertices are the flats and which has an edge between two vertices if and only if the corresponding flats are within 150 meters of each other. The chromatic number of this graph is the number of channels required for an interference-free wireless network.

  is incomplete.
- Drawing the correct graph, but making a mistake in computing its chromatic number: it was 4.
- A handful of students drew a graph with flats as vertices, but edges between those what were *more* than 150 meters apart. None of these students gave a good explanation of why this was a useful construction, so I don't know what they were thinking, but this approach also yields the wrong numerical answer.

(c) This part asked you to prove or disprove three propositions, though many students skipped the second and/or third of them.

- *If, for some positive integer k_0, the matrix A^{k_0} has strictly positive entries, then G is strongly connected.*
  Most students who attempted this realised that the proposition is true (it follows easily from material in the lecture titled Tropical Arithmetic and Shortest Paths), but a few lost marks for confusion about the definitions of the terms *path* and *walk*. There were also a few more profoundly confused answers involving adjacency matrices for digraphs that included negative entries.
- *If, for some positive integer k_0, the matrix A^{k_0} has strictly positive entries, then A^k has strictly positive entries for all $k \geq k_0$.*
  Many students skipped this: the statement is true.
- *If G is strongly connected, then there exists a positive integer $k_0 > 0$ such that A^{k_0} has strictly positive entries.*
  Many students skipped this too, though the statement is false and a suitable counterexample, illustrated below,

[1] See the lecture notes: the vertex set of I_d is the set of strings $\{0, 1\}^d$ and those labels don't appear in the diagram in the exam)

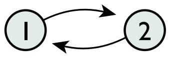

is discussed in the feedback for 2017's exam.

(d) It's possible to get the right answer, $k = 5$, by thinking carefully about walks in the graph, but it's arguably easier to just compute powers of A with the repeated-squaring algorithm. There were, however, a surprising number of elementary arithmetic mistakes and, as in the previous part, I saw a number of incorrect adjacency matrices with negative entries.

B6 This was the second most popular question in Part B, with 90/123 attempts. The average was around 12.2/25, but this was pulled down by a number of students who seemed to have run short of time, though 9 students got a score of 20/25 or better.

(a) The rather long list of definitions went well, though some students lost a mark or two in one of the following ways.

- In the definition of *girth*, saying

  If G contains more than one cycle ...

  instead of

  If G contains one or more cycles ...

- In the definition of *homeomorphic*, failing to say what a subdivision is or saying something like

  Replace an edge (a, b) with a walk ...

  instead of

  Replace an edge (a, b) with a path ...

Also, I found it odd that, as in previous years, many students wrote

A graph is *planar* if it contains a planar diagram.

Strictly speaking, this is wrong: graphs don't "contain" diagrams, they "have" them, but I was more concerned that people should give a complete definition, including one for *planar diagram*, and so ignored this minor mistake.

(b) It seems to have helped students know that the graph in question could be produced by removing 5 edges from K_{12}. In the yes/no questions that made up this part, I awarded one mark for a correct yes/no answer, then various numbers (between 2 and 4) for the supporting arguments. The main issues were:

Eulerian? I saw several incorrect claims that if the sum of the degrees of a graph's vertices is even, then $|E|$ must be even: K_3 provides a counterexample. Others claimed, incorrectly, that if the sum of degrees is odd (which is impossible: review the Handshaking Lemma), then some vertex must have odd degree. On the positive side, several students made interesting use of the Handshaking Lemma to establish the existence of one or more vertices with degree 11.

Hamiltonian? Many students (correctly) invoked one of Dirac's, Ore's or the Bondy-Chvátal Theorem, but a few just showed $\deg(v) \geq n/2 \; \forall v \in V$, which I regarded as incomplete. A handful of students tried, confusingly, to use the converses of these theorems, which are (a) false and (b) not relevant to the question.

girth? Most successful answers involved pointing out that K_{12} contains such a large number of edge-disjoint three-cycles that one can't possibly break them all by removing just 5 edges. The best answers included specific lists of three-cycles or exact formulae for the number of three cycles in K_{12} and an explicit upper bound on the number that can be broken by removing 5 edges.

planar? Essentially all students who got this far answered correctly, demonstrating that G can't planar as it has far too many edges, though, as in the question whether G

is Hamiltonian, a few students became confused and tried to use the converse of the relevant theorem. The issue is that although

$$G \text{ planar} \quad \Rightarrow \quad m \leq 3n - 6,$$

it is *not* true that

$$m \leq 3n - 6 \quad \Rightarrow \quad G \text{ planar}.$$

A counterexample to this second, erroneous statement appears in the left part of Figure 10.11, which is part of the lecture titled Planar Graphs.

contains K_5 homeomorphically? It turns out that G must contain subgraphs isomorphic (and so also homeomorphic) to both K_5 and $K_{3,3}$. Several students seemed to think that Kuratowski's Theorem says that a non-planar graph must contain an *induced* subgraph homeomorphic (or sometimes isomorphic—these students were confused) to $K_{3,3}$ or K_5 and so made a great deal of fuss about the degrees of the vertices in G, but Kuratowski's theorem concerns ordinary subgraphs, meaning that one is free to delete unwanted edges from G when hunting obstacles to planarity.

A handful of students made an interesting, novel argument saying that if we consider any one of K_{12}'s' many, many subgraphs isomorphic to K_5, the process of deleting edges to get to G can only break a few of of the subgraph's edges and that the endpoints of these broken edges would still, certainly, be connected by paths running through the rest of the graph. This means that any subgraph that's isomorphic to K_5 in K_{12} will have a companion that is homeomorphic to K_5 in G. This is true, and a great observation, so I awarded substantial partial credit, but for full credit I wanted to see a quantitative argument to establish that the necessary paths need never intersect each other.