

CSE 202

Basic Information: Fall, 2020

Instructor: Russell Impagliazzo

email: russell@cs.ucsd.edu

webpage: www-cse.ucsd.edu/classes/fa20/cse202-a

Piazza:

Class:

TAs: Shaobo Cui, Sam McGuire, Yash Pande

Russell's CSE 202 Office Hours: Tu, Th: 1-2, zoom.

TA office hours: TBA

Useful websites: Gradescope (for turning in assignments), piazza (discussion, announcements, video links,), course web-site (assignments, schedule, lecture notes, study guides)

Prerequisites: We assume some undergraduate exposure to discrete mathematics, and to algorithms and their analysis, and the ability to read, recognize and write a valid proof. For example, CSE 20, CSE 21, and CSE 101 cover the prerequisite material. We will cover many of the same topics from an undergraduate algorithms courses. However, after a quick review of the basics, we will move on to related advanced material for each topic. If your background in these areas is weak, you will need to do a significant amount of extra work to catch up. Please talk to me about this.

Text Book: Kleinberg and Tardos, Algorithm Design

Optional Supplementary Texts: Neapolitan and Naimipour, Foundations of Algorithms; Jeff Edmonds, How to Think About Algorithms A helpful text for proofs is Solow: How to Read and Do Proofs.

Assignments There will be six homework assignments, a project, and a take-home final exam. Each homework assignment will have exercises, four theoretical problems and one implementation problem. Exercises will not be graded, and are optional but recommended. They might be warm-ups or hints for the problems. You should budget 10-20 hours for each homework assignment, including time spent in office hours, and 20-40 hours for the project. Homework and the project can be done in groups up to five. Most people stay with the same group for the quarter, but you

are free to switch groups (or kick people out of your group if they aren't pulling their weight.) You can use the piazza site to help find a study group.

The first homework assignment is on pre-requisite material and should be started immediately. It will be given two grades, a grade for "sincere effort" in which every problem attempted is given full credit (which will be the grade recorded), and a standard grade with the same grading criteria as all future assignments (which is only for your use.)

Evaluation: Homework will account for 30 % of the grade, the project, 20%, and the final will account for the remaining 50 % of the grade. A B- or better on the final can be used as passing the MS Comprehensive Exam for the course.

Standards for evaluation: Most problems will be algorithm design problems, where you are given a computational problem and asked to design and analyze an efficient algorithm for the problem. Any solution to an algorithm design question MUST contain the following:

Problem statement A clear unambiguous statement of the problem to be solved.

Algorithm description A clear, unambiguous description of the algorithm. This can be in (well-documented, clear) pseudocode (with explanations in English) or in (precise, mathematical, well-defined) English. There should be no room for interpretation in the steps carried out by your algorithm. Any mathematically and computer literate individual should be able to follow the steps presented, or to implement the algorithm as a computer program. Such implementations by different people should still be essentially identical.

Correctness proof: A convincing mathematical argument that the algorithm described solves the computational problem described. Sometimes this can be brief, but it must always be given. Sometimes this can be brief, but it must always be given.

Time analysis: A time analysis of the algorithm, up to order, in terms of all relevant parameters. You must prove this analysis is correct. Frequently, this will be brief, but occasionally the time analysis will be the heart of the question, and can be quite challenging.

What is a proof? A proof is a compelling argument that forces the reader to believe the result, not just notes from which a proof could be reconstructed. Note that I will mark off for any unsubstantiated and non-trivial claims in a proof, even correct claims.

Don't assume knowledge or sophistication on the part of the reader. A good rule of thumb is to pretend you are teaching an undergraduate class

in algorithms, and write in a way that explains what is going on to the students in the class. Don't think of your reader as a trained algorithms professor who already knows the answer. We have to evaluate your answer based on WHAT YOU WRITE not WHAT WE KNOW or WHAT WE THINK YOU MEANT.

Your answer will be graded on the following criteria:

1. Your algorithm must be clearly and unambiguously described.
2. You need to prove your algorithm correctly solves the problem. My rule is: an algorithm is incorrect until PROVED correct. I will use this rule in grading even if I KNOW your algorithm is correct. If I think you don't KNOW why your algorithm works, I will grade it as if your algorithm does NOT work.
3. Your time analysis must be proved correct. A time analysis is usually an upper bound on worst-case time. You get credit for what you claim and prove about the running time of your algorithm, even if the algorithm is actually faster. Logic is as important as calculations in a time-analysis. At a minimum, a time analysis requires an explanation of where the calculations come from. If the analysis is "easy" (e.g., with a simple nested loop algorithm), these explanations can be brief (e.g., "The outside loop goes from 1 to n , and each iteration, the inside loop iterates m times, so the overall time is $O(nm)$."). Other times, the time analysis is a tricky, mathematical proof. If you give just calculations or just a short explanation, and I think the time bound is NOT easy and clear from what you wrote, you will lose points even if you give the correct time.
4. Your algorithm must be efficient. Although we sometimes use "polynomial-time" as a benchmark for "efficient", this isn't a hard and fast rule. An algorithm is efficient for a problem if there is no competing algorithm that is much faster. To be efficient might require being sub-linear time for applications where input sizes are huge, or almost linear time when the problem is trivially poly-time, or an improved exponential time algorithm for an NP-hard problem.

If there is a faster algorithm, then you may not get full credit. Nothing you said is false, but I'll still deduct some points because you COULD have been cleverer. How can you avoid this? There's no guaranteed way, but if the correct algorithm seems easy, you should still try to think of improvements or other approaches. Is this fair? Not really, but it is the facts of life. Your algorithm won't be used if there is a faster equivalent algorithm.

A good algorithm designer considers a variety of approaches and picks the best one, rather than prematurely converging to the first adequate algorithm.

There will also be some implementation problems on the homework. The goal for these problems is to conduct a meaningful algorithmic experiment and present data and conclusions in a clear format. Implementing the algorithms in question is necessary for these problems, but isn't the goal in itself.

For these problems, you should describe clearly what the algorithms you implemented were, what their asymptotic time analyses are, your results from the challenge instances described in the problems, and timing information about various problem instances. Then you need to make a reasonable inference and reach a conclusion from this data about the question asked. (Since this is not well-defined, there might be multiple reasonable inferences possible, and possibly the best conclusion is "There is no statistically significant conclusion possible from this experiment".) **WE WILL NOT READ ACTUAL CODE SO DON'T BOTHER HANDING IT IN.** If the actual times seemed different from the asymptotic analysis, give a short discussion of possible reasons. Give the programming language used and your computer's speed rating so that I can normalize. Your grade will be based on your answer's completeness, and your success and time for challenge problems.

One handy tool is to graph performance on a log-log scale, the log of the problem size on one axis and the log of the time taken on the other. This allows you to compare running times on a variety of scales, and to show the results in a comprehensible way. More importantly, it converts polynomials to lines, with the slope of the line revealing the exponent of the polynomial.

Project The project for the class is meant to develop the ability to model applications as algorithmic problems. In order to avoid looking at applications that are already the focus of intense research, where many formulations are already known, I want you to pick a somewhat frivolous application based on a hobby, puzzle or game, and to avoid those addressed by large amounts of previous research.

Your group should hand in reports on the following schedule:

Proposal By Monday, October 19, hand in a two to three page description in English of the hobby or game your project will focus on, and what computational questions, in terms of the hobby, you want to address.

Formulation By Monday, Nov. 9, , give a mathematically clear formulation of computational problems related to the above. Give a brief justification of how your formal description captures the aspects you described in the first part.

Algorithmic Solutions: By Monday, Nov. 30, hand in descriptions of algorithms you have come up with to solve the problems in your formulation, or other results about the difficulty of these problems (e.g., NP-completeness). Your algorithms might be exact algorithms or approximation algorithms, or heuristics if all else fails, but you should definitely give time analyses and correctness proofs as applicable.

Experiments: By Monday, Dec.7, implement your algorithm and perform some experiments using the implementation. Submit a summary of the results, and a conclusions about the applicability of the algorithm to the game or hobby.

Academic Honesty Golden rule of Academic Honesty The point of all research is to get people to share ideas. In order to properly reward people for sharing ideas, it is necessary to give them credit when they do so. So the golden rule is: Whenever you use someone's ideas, you MUST give them credit.

(The exceptions are those being remunerated for providing you ideas for this course: myself, the TA's, the textbook, class materials on the website or piazza site for this quarter's section. These do not need to be credited. Any other person, site, or source does.)

Working in groups Students will be allowed to solve and write up all homework assignments and the project in groups of size from 2 to 5. This needs to be a *collective* effort on the part of the group, not a *division of labor*. Students are responsible for the correctness and the honesty of all problems handed in with their names attached. If a group member did not participate in discussion for one of the problems handed in by the group, the group must write a note on the front page of the solution set to that effect. However, verifying someone else's solution counts as a contribution.

The point of working in groups is:

Practicing communication skills: A major part of algorithm design and analysis is clearly communicating what your algorithm is and why it works. You need to practice with your group explaining algorithms to others and verifying claims about algorithms.

Brainstorming: Your group needs to come up with a variety of approaches to problems. Working by yourself, it is too easy to get stuck on one approach. Other people are often better at catching our mistakes, too.

Preparing for later careers: Whether in academia, research lab, government or industry, almost all work is collective, not individual. Learning to work as part of a team is essential.

Reducing our workloads: This is a big class. The TA's and I would not be able to properly grade individual assignments in a timely way. So hand in one assignment per group.

Other sources: Students should not look for answers to homework problems in other texts or on the internet. Yes, it is interesting to find the best known solutions. **However, do the literature search after you hand in your own solution.** Students may use other texts as a general study tool, and may accidentally see solutions to homework problems. In this case, the student should write up algorithms to others and verifying claims about algorithms. By text, I mean any

Brainstorming: Your group needs to come up with a variety of approaches to problems. Working by yourself, it is too easy to get stuck on one approach. Other people are often better at catching our mistakes, too.

Preparing for later careers: Whether in academia, research lab, government or industry, almost all work is collective, not individual. Learning to work as part of a team is essential.

Reducing our workloads: This is a big class. Fjola and I would not be able to properly grade individual assignments in a timely way. So hand in one assignment per group.

Be sure to follow the following guidelines:

1. Do not discuss problems with people outside your group whether students or not (except the TA and me, of course).
2. Do not share written solutions or partial solutions with other groups.
3. Prepare your final written solution without consulting any written material except class notes and the class text.
4. Acknowledge all supplementary texts or other sources that had solutions to homework problems. Acknowledge anyone who helped with assignments, except the TA and myself.
5. Do not discuss the final exam with anyone except the instructor and TA.

Remember the Golden Rule: if you use someone's ideas (whether that is an anonymous Wikipedia editor or your spouse), you **MUST** give them credit. The exceptions are for standard ideas that everyone knows and the people whose job it is to give you ideas (TA's, myself, class website, the textbooks).

Lateness Policy Late homework will be accepted until I give out an answer key and no later. So you have to be no later than me.

Reading Schedule When we cover a topic in class, I will not always use the textbook examples, or follow the textbook descriptions. The point is not that you shouldn't read the textbook, but to give you MORE examples and viewpoints (the ones in class AND the textbook's). You are expected to be reading the relevant sections of the textbook as or before we cover them in class. Much of the prerequisite background will never be explicitly covered in class, but if you aren't already familiar with it, you are expected to teach yourself from the textbook (or undergraduate textbooks in Algorithms).

To help you plan, I will give a tentative schedule of general topics to be covered in class, and the corresponding sections of the text to be read: Our mileage may vary as we actually cover the topics in class.

Background: Chapters 1-3. Also, recurrence relations from 5.1, 5.2. The first homework is on background, so if you are having trouble with understanding that, you need to spend more time on background.

Introduction, formulating problems 1 lecture (Oct. 2).

Graph search Chapter 3. Breadth-first and depth-first search, applications. Dijkstra's algorithm. Incorporating data structures and preprocessing into algorithms. Reductions and how to use them. Analogies and why they need to be reworked. 4 lectures (October 5-12).

Greedy algorithms Chapter 4, some of chapters 11 and 12. Some examples: Interval scheduling (4.1), Minimum Spanning Tree(4.5); (4.4), greedy approximation algorithms (11.1-11.3). Methods for proving greedy algorithms optimal. Analogs for proving approximation ratios. Continuing discussion of optimizing data structures for algorithms, amortized analysis. 6 lectures. (Oct. 14-26)

Divide-and-Conquer : Chapter 5. Mergesort (5.1, in passing), Integer Multiplication (5.5) and Fast Fourier Transformation, Closest Pair of points (5.4). Median Finding and selection (13.5). Recurrences, the Master Theorem, and what to do when the Master Theorem doesn't apply. Multiple parameters. Parallel computation. (see paper on website) 5 lectures. (Oct 28-Nov. 6)

Search and Optimization problems, and the class NP See section 8.3 (1/2 lecture)

Back-tracking, Depth-first Search, Breadth-first Search, Branch-and-Bound

Not in text, but see Chapter 10. Example problems: Maximum Independent set (p. 16), Graph Coloring (8.7), Hamiltonian cycle(8.5) addition chains (not in text). (1 1/2 lectures: Nov 9-13).

Dynamic Programming: Chapter 6. Weighted interval scheduling (6.1), edit distance (6.6), All-pairs shortest paths (6.8); Subset sum (6.5)

and DP-based approximation algorithms (11.8). 5 lectures (Nov. 16-25).

Network Flows, Hill-climbing, linear programming: Network flows and the Ford-Fulkerson algorithm. (Chapter 7.1-7.3). Hill-climbing as an algorithmic technique (not in text, but see Chapter 1, stable marriage, and Chapter 12, local search) . Metropolis and simulated annealing. (2 lectures, Nov. 30-Dec 2).

Linear programming: Linear programming, duality, simplex algorithm, ellipsoid algorithms, approximation algorithms using LP. (2 lecture: Dec. 4-7)

Reductions and NP-completeness (Chapter 8.2, 8.4-8.7). (Nov. 9 lecture)

Coping with intractibility: Heuristics, average-case analysis, and approximations (Chapter 11, 12). 1 lecture (Nov. 11).

Assignment Schedule Oct. 9 Homework 1 due

Oct. 19 Project proposal due

Oct. 23 Homework 2 (path-finding, reductions, using data structures) due

Oct. 30 Homework 3 (greedy algorithms, approximation algorithms) due

Nov. 9 Project problem specification due

Nov. 13 Homework 4 (divide and conquer) due

Nov. 23 Project algorithms and analysis due

Dec. 4 Homework 5 (BT, DP) due

Dec. 7 Project implementation and testing due

Dec. 10 Take home final available

Dec. 11 Homework 6 (NF, LP, NPC) due

Dec. 17 Take home final due