CSC287/487: Algorithms and Elections Information Handout/Syllabus Fall 2012 Version 1.03 2012/9/6 (last revised 2012/10/15/9:54AM) (the most recent revision of this document can always be found on the course web site)

Location/Time: The course meets in Room CSB 632, M/W 4:50PM-6:05PM.

Instructor: Lane A. Hemaspaandra, CSB 618, 275-1203, lane "at" cs.rochester.edu. Office Hours by appointment (but please always feel free to stop by any time without an appointment or, easiest of all, grab me right after any of our class sessions).

Prerequisites: Although the course has formal prerequisites (see registrar-land), I'll try to make it possible for people to take the course even if they have different background. Basically, though, the course wants you to come in with knowledge of NP, NP-completeness, and what an algorithm (and its running time) is.

Course Goal/Description/Mechanics: This course's topic will be algorithms and elections, especially how to use algorithms to attack elections and how to use complexity to protect elections from manipulative attacks. We will cover probably about twenty-six research papers in this area—generally, one per class. Among the key computationally interesting problems regarding election systems are the following four central problems: the winner problem; bribery; control; and manipulation. So we will see quite a bit about those problems. But we probably also will touch upon other issues, which potentially might include some of the following: heuristics; range restrictions such as single-peakedness; approximations; possible and necessary winners; and various issues regarding preferences, preference representation, and preference elicitation. The reason I say "potentially" is because the students in the course will in a very important way help shape what material is covered, as will be made clear in the following paragraph.

The course will work as follows. I will myself present the first four class sessions. The first class will be to go over the goals and logistics of the course. The next three sessions after that will be lectures by me, presenting the material of three papers. But after that, the sessions will be presentations, by the class's members, of literature papers. These will be organized as follows. A "unit" will be a series of three lectures (in rare cases, if needed given the number of slots in the term, a group might be given a 2 or 4-lecture span for a unit; my current estimate is there will be one 2-lecture unit and all the rest will be 3-lecture units), with as its focus a towering researcher on algorithms and elections. The group that is assigned to present the unit will give three (or whatever number of lectures the unit is assigned to be) consecutive lectures on papers by that person. (Note: Most natural will be to speak on three papers, one per lecture; but if you wish and I don't ask you to cover three specific papers, you can speak on two papers, distributed over your three class sessions.

But in general, the best approach is to cover one paper per lecture.) In some cases, I'll when assigning your focus person also specify some or all of the papers for you to cover (for example, for the "Sandholm" unit, there are (at least) three very important papers to cover, and so that one will probably be fully specified by me), but in many cases, you yourselves will choose some (or perhaps even all) of which papers by that focus person to present. The papers of course should be on the topic of algorithms and elections (note: we use the term algorithms broadly, to fully include both p-time algorithms and complexity results such as NP-hardness that suggest the nonexistence of p-time algorithms; as to elections, we'll not interpret that tremendously broadly, but rather it really means things that are directly about elections, except also issues of preferences/preference representations/preference elicitation we'll consider fair game since those issues are so important to elections; additional note: you should not speak on a paper that some earlier group has already spoken on, even if your focus person is an author of it). However, though your papers will all be by the focus researcher, you are free (give or take whichever papers by him or her I ask you to make sure to cover) either to choose three related papers by that researcher or to cover quite differing papers by that researcher; in fact, most fun might be to do the latter—these researchers are quite broad, and you can thus present three of their very best papers (but, of course, always just those within the focus on algorithms and elections).

The group members do not have to split each lecture equally between them, but each group member should do a substantial amount of lecturing during the group's unit. For example, a 2-person group could have one person give two lectures and one person give one lecture.

I'll assign the groups and the focus person for each group. The size of the groups (which might even vary from one unit to the next, and if the class is extremely small might even be one person per group for some or all units) and the number of groups each person will be in during the term will depend on the number of people in the class. At the first class session, when we see how many people are in the class, I'll be able to estimate those numbers. (My off-hand guess is that the typical group size might be two, or maybe sometimes or always three—it depends on how many people take the course—and that each person might be in around 2–4 groups. I'll try to make sure that no two people in the class differ by more than 1 regarding the number of groups they are in, and will lean toward students-taking-it-as-487-rather-than-287 or volunteers as to who will end up doing the one-more if some do one-more; and it is quite possible that things will work out so that everyone will be in the exact same number of groups.)

Logistically, class sessions 2, 3, and 4, presented by me, will give you a sense of what is expected in a unit/presentation (but my three lectures will differ from the other units in a few ways, for example, yours will be about towering researchers' work and mine will just be about some of my work; you should make your slides yourselves and for my talks in some cases I may use slides prepared by my coauthors; nonetheless, my talks should give you a sense of what a unit can look like). In particular, your presentation should be done with nice slides (namely, PowerPoint or pdf) made by your group, projected using the 3-color projector from your laptop. You during your group's 3 lectures will present the (typically 3) papers by the focus researcher. Immediately after each lecture (to be specific, this *must* occur after class but also by at latest 8PM of the day of your lecture, or your group's grade for the unit will be lowered), you will mail to the entire class an email that gives (1) a bibliographic reference to the paper that you covered (e.g., "Why P equals NP," by Dr. Seuss, Journal of Irreproducible Results, Volume 121, number 3, pages 4-143, January 2011; if the paper itself is available on the web, you might want to also provide a URL for where the paper is, though journal papers can almost always be gotten by any one of us by using our the UR library system's "eJournals" service) (2) a URL pointing to a file that is the slides that your group presented (namely, as .ppt/.pptx/etc. if they are PowerPoint slides and as .pdf if they are pdf slides). Do this by sending a very short such email to the class mailing list address, which is the username "csc487f12" in the domain "cs.rochester.edu" (email sent there will go right to everyone who is subscribed to that "mailman" list; you should each *immediately* add yourselves to that list, unless I myself already put you on it based on your preregistration). Please be aware that the setting of the mailing list will itself probably bounce/block emails that are too long, and various internal or external mail servers could impose even tighter constraints; thus, to avoid the headaches and confusion of some people not getting large emails (and slides often create huge files, e.g., due to pictures or scans), note that your email should *not* send the slides, but rather should send a URL that is the slides at a web location where you have put the slides. (Note: You should NOT change the file at that location after you have sent out your email!) If a given class is for example about the same paper as the previous class and if the current class used slides that already were distributed by your group associated with the previous class session, you nonetheless should definitely send an email by 8PM, mentioning again the paper and the URL, but also mentioning that one or both are the same as you sent around associated with the previous class. (Simply put, if you speak on day N, you should send an email out after class ends but before 8PM.)

Also, the department might in time try to create an offering that is more permanent and lecture-like than this course (e.g., perhaps at first a 19X course and then a permanentnumbered course mostly for undergraduates), and if it does, your slides might be very helpful if you'd allow them to be used for that. So by default, turning in your slides is implicit permission to use them (of course making clear the fact that you created them; by the way, the names of all your group members should be mentioned always on the very first slide of your presentations, probably write below the title, authors, and bibliographic information on the paper the presentation is on; the natural way to do this is by a line of the form "Slides/Presentation by [your names]") in future course offerings; if you wish to not have a particular set of your slides open to being used that way that is totally fine, and please to signify that in that same transfer email just include the line "please do not use these slides in any future course offerings" (including that line in no way impacts Note that at the end of each unit, I will assign a letter grade to the group for the unit (I may do this in person immediately after class on your 3rd lecture day, meeting with you right after you dismiss the rest of the class). All members of the group get the same grade for their joint unit; you are working as a team, and so should each help each other make every aspect of the unit as well-done as possible, as you all are responsible jointly for all of it, and will be graded jointly.

Each presentation should be timed so that it *and whatever questioning happens during or after the talk* will fill (but fit within) the 75-minute class. Keep in mind that each student not in the group giving the talk *must* ask a question at each talk; so do leave time for that. For most groups, it will be natural to time your talks so that without any questions or interruptions the talks take 50-60 minutes (certainly not more than 60 minutes); that will allow 15-25 minutes for answering questions. Each group at each lecture should keep track of who has asked questions so far during that lecture, and at the end of the talk should during the question is called on to ask a question before any other student asks a question (warning: I myself might stick in questions); it is part of your task as talk-givers to help ensure that everyone asks a question (but, if a person does not ask a question, it is that person who will mostly reap the bad consequences, though the group also will a bit). A good way to handle this is to while giving a talk simply keep a piece of paper and check off people as they ask questions. (See below regarding the nature of questions—what is "substantial" enough to count.)

Preparing/giving your presentations is the largest part of your coursework (and also, see below, is most of your course grade). Your group should do that preparation completely yourselves—this is a central part of the course experience (and is not an easy thing, but it will get easier as the course goes on and you get more experienced at it); you should not seek help from any live person (no asking people for help, not even me, and no posting questions to the internet or to the class mailing list, etc.) and you should not lift slides from the web site of authors of the paper (or anyone else) even if they have up slides. However, you certainly can look at such slides sets if they exist on the web (they usually don't but sometimes do), but your slides should be your own take on the work; also, it of course is fine to read lots of papers/etc. while preparing your talk—the library and the internet are great resources. After a talk is prepared, you will want to do what experienced people do: Give it, in private, to an empty room or to your group, to fine-tune it. For example, at the Dagstuhl workshop I was at earlier this year, the night before my talk, I was in the lecture

hall, giving in full, to an empty room, my talk.

By the way, the flavor of each talk is largely up to the group. Among the things each group may cover are such things as positioning of the paper and its relation to the literature, the models of the paper, the proofs/proof techniques of the paper, the results of the paper, and the open issues presented in or created by the paper. Most talks will of course cover various of those things, with stresses set to the group's taste. My only caution to you is this: Slides filled with technical proof details almost never are followed by audiences, and so to spend a talk all purely in technical proof presentations is almost impossible to pull off well; usually it is far better to just try to convey the core idea fueling the paper's proofs, or to give examples showing whatever is creative in the proof technique, or to give proofs but to keep them clear and accessible and not to let them take up so much of your presentation that people lose site of what is being proved or why it is important. For some papers, you may well choose not to give any proofs at all, but rather just to mention the flavor or idea behind the proofs and refer interested people to the source paper for the details. Let me also mention that it is particularly important for you to, as you present papers (my unit will be a bit of an exception as I'll be presenting my own work, and to do this there wouldn't make as much sense as one is always biased about one's own work), say what you found particularly strong in the paper and what you found particularly weak, and, relatedly, so very important as to be something you should make explicit at the start and end of each talk is to crystallize into a sentence or two the core contribution of a paper. (For example, "This paper's core contribution is to show that whether there are sparse sets in NP-P is completely determined by whether NE and E differ, and it shows this by a stunningly creative new encoding technique that spreads information about a set into logsized strings that allow the set to be understood and recreated. The only possible weakness I see in the paper is that this is building on an earlier and analogous result regarding tally sets; but the advance from tally to sparse seems to me quite powerful and nontrivial." As another example, this one perhaps historically unsound, "In this paper, Alan Turing introduces a computing model that is very creative and shows that it in a certain sense is 'universal,' yet I'm just not sure this notion is well connected to actual real-world computing machines in terms of running times tightly enough to ever catch on or be further studied.")

Your presentations should definitely not assume that people have read the paper before seeing your presentation. (Of course, interested audience members might well be inspired by your talk to, after you send your 8PM email, get and look at the paper! By the way, let me mention that UR gives you online access to the content of a huge range of journal (and other) articles, and almost all of our focus authors have their articles easily available directly or indirectly from their web sites.) What we're loosely modeling here is the experience people have at conferences and workshops and department seminars (and to some extent, in a sort of strange way, undergraduate and graduate courses at most lectures!); at conferences and workshops, in fact, one within a few days has to process dozens of talks on papers one has never seen before, and through focusing on the content as it unfolds and through asking questions the audience must help ensure that the speaker doesn't lose them and that they get the most possible out of the talk. One important goal of this course is for you to improve your skill at focused, active, interactive listening/learning/participation during this type of talk and the question session at its end. (By the way, unless you are a member of the group giving the talk and are using your laptop to help control the slides/talk, you should not have your laptop or cellphone open or in use during talks. If you have a documented disability relevant to this issue, please speak to me early in the course about this, as I of course will not apply this to you in that case.)

Grading (and More Mechanics, esp. Regarding Class Participation, Attendance, and the Project): The overall course grades (A, A-, etc.) will not be any fixed curve, but rather will reflect the instructor's judgment about what grade is appropriate based on performance quality, i.e., based on the grades you get on the various components, weighted as described below. However, you will certainly not be given a course grade lower than your grade average as computed as described below and rounded to the nearest grade that exists for course grades for you for the course; (e.g., if your grade average is a 3.85 then you'll get at least an A, and if your grade average is 3.84 then you'll get at least a B+); be aware that the grad school and the Engineering school disallow certain grades. I realize that that might let some of you realize well before the end that you're on track to an A, but I certainly expect you to even if so keep giving the course your best (slacking at the end may result in a low second-half class participation grade, but I know you'll work hard not because of that grade effect but rather because you're dedicated to learning—which is why you are taking this!).

The average of your unit presentation grades will count for 60% of your grade. Your class participation/attendance grade will count for 20% of your grade. Your project grade will count for 20% of your grade. See below for more details on these grades. Grades are by letter, and for doing the weighted averages, we'll use the following values for the grades: A=4.0; A-=3.7; B+=3.3; B=3; ...; E=0.0. (UR doesn't have an A+ grade, but if I give an A+ on something I'll average it as if it is a 4.3.)

Your group's unit presentation grade for a given unit will be given to you by me at or shortly after the end of the unit; see above.

Your class participation/attendance grade will be the average of three grades: your class participation grade for the first half of the term, your class participation grade for the second half of the term, and your attendance grade for the term. The first of those will be given to you around the middle of the term and the second will be given to you during or shortly after the last one or two class sessions. You are *required* to pay close attention to each of the lectures not by your groups (and also those by your unit-mates), and to, at each lecture not by your groups, *ask at least one question*. The question you ask might be at the end or might be during the talk; each group will try to ensure that everyone gets a

chance to ask a question at each talk, but it is also up to each of you to do the asking. Now, at first I was thinking of trying to define or require the questions to be substantial to count. However, one person's substantial question might be someone else's dumb question, yet might be valuable to the person asking it and to others. So, what each group will try to enforce is that everyone gets a chance to ask at least one question. And please don't hesitate to ask "dumb" questions during a talk; if you are confused about something, probably a few other people are too but are too shy to ask. However, if your question/comment is for example a joke that doesn't have any real substance, please make very sure to ask another question/comment during that same class, that does have more interest/content to it. You shouldn't ask questions just for the sake of asking a question (though sometimes you have to, but if so, that is a bad case); rather, you should be alert during the talks, and ask questions about things in the model that lack clarity, or things in the talk that you don't follow, or about open directions, or etc.--ideally, the questions should be things you care about and are interested in hearing more about or things you would like clarified. By the way, I mostly wrote "question" above, but actually I always mean "question/comment," e.g., if you make a comment such as "aha... if what you say is true then you have proven that P=NP" and you explain why, well, that is more interesting than most questions. (That sounds crazy, but that is exactly what one of the titans of the AI world said a few days ago at a conference at a talk on one of our papers, and he explained his proof of why our theorem was implying P=NP. My coauthor explained to him how the "obvious" fact he was using to support his claim had the problem of not being true, and indeed of having counterexamples known.) I hope everyone will get A's for their participation grades, but that isn't automatic, even if you ask questions at each talk; you need to seriously contribute to the questioning. For example, if at each talk your only question is, "What time is it?", you'd get an "E" (fail) on participation.

As to attendance, I expect you to attend every class session, unless of course you are ill. Logistically, though, that grade is completely mechanical: Your attendance grade starts as an A, and if you miss 0 or 1 or 2 classes remains an A, and for every class beyond two that you miss drops by a notch through the UG/G grades (whichever applies to the number you are taking the course under) at UR (so if you miss 3 classes during the term, your grade drops to an A- and if you miss 5 it drops to a B). Those 2 allowed misses are a bow to sicknesses, broken cars, and other compelling excuses, but are the only misses allowed; the only exceptions to this drop-2 rule will be in cases where you bring in an appropriate letter from a convincing Dean who in writing certifies that he or she has read and understands the course's policy that no quizzes beyond two are dropped and feels that your case is so extraordinary that there is a compelling case to not follow this fixed course policy in your case. Of course, if you are healthy all term, please miss not two classes but rather zero classes.

As to what we'll call your "project," that is not a group project, but is an individual project. And here is the project. You will write a research evaluation of one of the focus

people covered by one of your groups (you can choose which focus person to write on, but it is cheating to discuss your paper with anyone before the due date, and so if someone else is writing on the same person, it would be cheating to discuss it with that other person; note: if you're in one of the last groups, you might in light of the project's due date want to choose a from-an-earlier-one-of-your-groups focus person, of course). So, basically, if the person is an assistant professor you'll be writing a tenure-case-like research evaluation letter, if the person is an associate professor you'll be writing a promotionto-full-case-like research evaluation letter, and if the person is a full professor, you'll be writing a should-this-person-be-given-a-University-Professorship-like research evaluation letter (a University Professorship is a special distinguished level that some schools have; UR doesn't, though there are some named chairs; even if the person you're writing about is already a University Professor at his or her school, ignore that fact and write a letter as if the person currently were not one). Your letter should be a clear setting out of the strengths and weaknesses of the person's research. You'll probably want to mostly focus on the about-3 papers your group covered, but you'll also probably want to look at some other papers by the person to get a fuller picture. Your evaluation letter *must* meet all the following format/size constraints: It must be 6-10 pages in LaTeX if you are taking the course as CSC287 (and must be exactly 10 pages in LaTeX if you are taking the course as CSC487), using 11-point Computer Modern, using linespread factor 1.3 (that is, right after the documentclass line, include "\linespread{1.3}"; this will increase the interline spacing), with margins of 1.4–1.6 inches on all 4 sides (the page number itself doesn't count, that is, the space from the bottom of the page to the last text line of the page should be 1.4–1.6 inches, except perhaps on the final page; also, the top-of-page gap on the first page might be different if you use a "maketitle" or other header and that is fine). That might sound short, but evaluation letters need to convey their points succinctly enough that they are read by administrators and chairs and so on (and no, no, please don't make the same point about course information documents—anyway, the very fact that you are seeing this parenthetical proves that you did get to it, by hook or crook or perhaps because I'm reading through it out loud in class on the first day!). In writing your letter, do keep in mind that the hypothetical eyes that will see your evaluation letter are not specialists in algorithms and elections, but rather are a bunch of faculty members from all areas of the person's department, and also the department chair, the Dean, the Provost, and an outside "ad hoc" committee of faculty members every one of whom is *not* from the person's home department. So your letter must be readable by people who don't even know what NP or Big-Oh are. And so in your letter, you will have to express the person's strengths and weaknesses, and the person's advances and their importance and value and what those advances even mean at all, in terms that are clear to people who are smart but don't have any relevant background. And you'll have to do it in 6-10 pages. This is a truly demanding task, that is actually more difficult than writing for technical experts. (I know this all too well. As a faculty member, I often have to write external evaluation letters; it is way-hard!) Your project should be turned in by email, as a pdf document, no later than 11:59PM, Wednesday, November 28th, 2012 (be careful—this is just a few days after the end of Thanksgiving break, so please do plan ahead and start this early so you don't get caught in a last-minute rush); email your project not to our course mailing list, but rather directly to me, namely, to "lane" in the domain "cs.rochester.edu." (Just to be explicit about how lateness on this will affect your grade, turning it in by 11:59:59.999 on 11/28 means there is no reduction due to lateness. And on the hour, each hour, starting at midnight sharp there is a 0.25 point deduction, on a 4-point scale. So for example if your paper is turned in at 1:00:00AM and your grade is a 4.0 (i.e., an A), the lateness will reduce your grade to a 3.5 (i.e., halfway between A and B). Submitting it right at midnight would lower a 4.0 to a 3.75. And if you submit it at 8:00:00AM on 11/29, a 4.0 paper will get a 1.75, and if you submit it at 3:00:00PM on 11/29, a 4.0 paper will get a 0.00 grade. In short, since improving a paper by 0.25 points in an hour is not likely: You have every incentive to submit it on time, and submitting it late will drastically drop your grade. So please make sure to submit it safely on time. Finally, let me be clear about the issue of times-send time versus receive time and so on. I'll use (unless the entire header trace makes it clear that that one machine itself has a crazy clock, e.g., if the machines after it show much earlier times, in which case I'll use my own judgment after looking at the whole header trace) as the time the first time in the email hand-off trail, *excluding* any time value generated by your own machines. So if your laptop's clock or the clocks on other machines inside your controlled-by-you machine network/world is/are running slow, that won't give you extra time. And, after sending your email, don't put your laptop to sleep so quickly that the email doesn't get sent out at that time (on my Apple laptop, if I hit send on an email and shut the cover quickly, which puts the laptop to sleep, the email often doesn't go out until the next time I bring that laptop back to the world of the living, so this really can happen)-doing so might turn your paper from an A to a fail.)

There will be no quizzes or exams.

This course is available only for credit, not as an audit.

As to grades (and what is made public and what is not), each student/group is given his/her/its grade on each item, and I of course assume you'll keep track of those grades (and I too will, of course), and at the end of the course he or she will receive his or her course grade either from the registrar or from me. However, I in general do not give out grade/distribution info beyond that, though of course if a student wishes to share with some classmate his or her grades on some items or the course, that is up to him or her.

Class Attendance: Required.

Academic Honesty: On all aspects of the course, it is very important to follow appropriate norms of academic honesty and professional conduct, most importantly including those regarding plagiarism; do not plagiarize. Please see the link from the course's main web page to the school's site on academic honesty.

Textbook/Library Reserves: There are no required texts. I have put on reserve at Carlson over a dozen books, mostly books or handbooks on algorithms, complexity, LaTeX, and voting. These might be useful to you, for example, if you want to look up some complexity class or something about algorithms or something about voting. Most are on 1-day reserve except the book edited by Ravi and Shulka, which is on 2-hour reserve. (Note: As of when I'm writing this, the library hasn't yet processed my reserve request, but I trust they will very soon.)

Newsgroup/BlackBoard: We will not use a course newsgroup or BlackBoard (except the library makes the reserved books list available through BlackBoard).

Fun with Numbers: Although the course is both CSC287 and CSC487, depending on which you are registered under, I'll often just refer to it by one or the other of the numbers (or sometimes as CSC287/487). Note also that although students taking 287 can submit a project as short as 6 pages, those taking the course as 487 are required to submit a 10-page project; also, those taking it as 487 will be leaned toward as to who has to do an extra unit if the numbers don't come out so that the load is uniform in terms of number of units.

Email: So, please note this carefully: mail sent to lane "at" cs.rochester.edu goes just to me, and mail sent to csc487f12 "at" cs.rochester.edu (note that it is "csc" not just "cs" and that there is an "f" before the year) goes immediately to the entire class (including me). Please do not send emails to the class list except for your "by 8PM the day of your talks" emails.

Web: The course web page is www.cs.rochester.edu/courses/287/fall2012/, and you can find the syllabus and some helpful links here.

A Final Comment: In this course, you'll be yourself mastering and presenting, and hearing others present, a large number of very important, cutting-edge research papers. You may find some things difficult to follow, and if a talk is hard to follow, raise your hand right away and ask the speaker whatever question is worrying you or about whatever you are not following; as I mentioned before, if you are confused, so too are probably many other more shy people. And don't feel bad if occasionally a talk is unusually hard to follow even after you ask questions; even professors sometimes-more often than they would like to admit—at workshops attend talks, even in their own area, and find the talks hard to fully follow. (Here is a cute, relevant anecdote from Moshe Vardi's introduction to the September 2011 issue of CACM, where he mentions asking, at a theoretical computer science conference, how many people understood at least N% of at least N% of the papers—to get 50% of the people to say yes, he had to drop N down to 50. Understanding talks is a nontrivial challenge.) Nonetheless, I suspect for most of the talks, it will be (thanks to your strong presentations!) easy to follow the core idea of what the paper did; so pay attention and participate, but above all, enjoy—almost all the papers you'll present did truly lovely things or opened/expanded exciting directions.