Social and Information Networks

University of Toronto CSC303
Winter/Spring 2023

Week 11: March 27-31

This week's agenda

- Congestion Networks
  - Traffic networks (Ch 8.1)
  - Nash Equilibrium
  - Braess' Paradox (Ch 8.2)
  - Social cost
    - ⋆ Tragedy of the commons
    - ⋆ Price of anarchy
- Kidney Exchange
- Recap

New topic: A congestion game

- We will now be considering a game (as in game theory) that models a highway network system (Ch 8)
  - Surprisingly, we will see that building more roads can be harmful!

- Our model is as follows:
  - We have many agents
    - ⋆ i.e., drivers commuting at the same time and in the simple model we study they are all going from some point A to some point B
  - They are using a highway network of roads (i.e., edges in the network)
  - Travel time per road depends on the number of drivers using that road
- This is a game; drivers have a self interest in arriving as soon as possible
  - The social objective of say the government (in this model) is to minimize the average (over all drivers) driving time
- Note that we are saying "roads" here but edges could be links in a commuter rail or subway network
  - ... however, our model uses linear congestion, therefore roads make more sense

A simple but interesting example

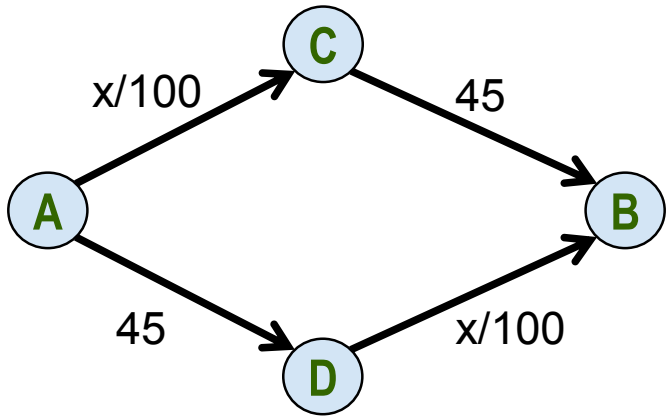

- The edge label "45" means that it takes 45 minutes no matter how many people are using that edge
- The edge label "$x/100$" means that the time on that edge takes $x/100$ time units (e.g., minutes) if there are x people using that road
- In this network, drivers (commuters) have two possible paths to go from A to B
- Question What route should they decide to take?

The traffic network example continued

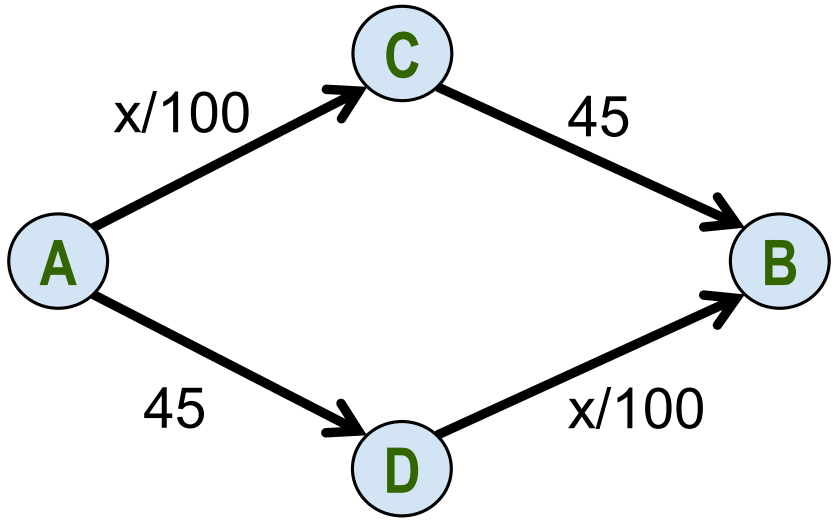

Suppose we have 4000 commuters, each deciding to travel via C or D.

- Formally, there are 2^{4000} possible outcomes
- However, we view commuters as equivalent and thus there are 4001 outcomes
- All outcomes with x people using the path via C (and $4000 - x$ using the path via D) are all equivalent and we will just view them as one outcome

What is a Nash Equilibrium for this traffic network game?

- We are interested in a Nash Equilibrium (NE)
  - an "outcome x" (i.e., with x using the path via C) such that no individual will want to change routes in order to save time (under the assumption that no one else changes their route)

Claim: The solution $x = 2000$ is the unique NE.

Proof of Claim: In the outcome with $x = 2000$ commuters using the path via C (and hence also 2000 commuters using the path via D), if any individual changes their route, then their commute time increases from $t = 45 + 2000/100 = 65$ to $t' = 45 + 2001/100 > 65$.

While this would unlikely be noticed by a single individual, what happens when more and more decide to switch?

The NE optimizes social welfare

- The outcome $x = 2000$ is a unique NE, and **also** the unique optimal outcome in terms of the social welfare (i.e., the average or total commute time)
- Consider the outcome when 2001 go via C and 1999 via D:
  - 2001 commuters will increase their commute time by .01 minutes while only 1999 will save .01 minutes
  - Therefore the total of the commute times increased by .02 minutes
  - A similar observation applies for the outcome when 1999 go via C and 2001 go via D
- While any individual commuter is unlikely to notice this, larger deviations are apparent
  - Suppose now that 3000 go via C, then the total commute time will now increase by 20,000 minutes $\approx$ 2 weeks worth of time
  - If everyone takes the same route, then the total commute time will increase by 80,000 minutes $\approx$ 2 months of time

What happens in "practice"

What would happen if everyone started using the same route? Would it be likely that they would *all* switch to the other route?

I think the NE outcome is something that we would likely see (approximately) as the result of individuals gradually adapting to traffic.

Of course, real traffic networks are more complicated and individuals do not know what others will do, but still, it is plausible to believe that individuals will converge to something resembling an equilibrium. How would you imagine this happening?

Essentially we would expect random uncoordinated decisions will gradually lead individuals to work towards solutions that come close to an equilibrium. The study of the Braess paradox comes, of course, before the use of GPS systems. With GPS people can change routes dynamically based on real-time information.

Braess' Paradox

Suppose the premier decides to build a new superhighway (or super fast rail line) and add this to the existing traffic network.

Lets even imagine that the time to traverse this new additional link is negligible (and hence approximated by 0 time). It seems that this can only improve the life of commuters. So lets add a directed link from C to D in our example traffic network.

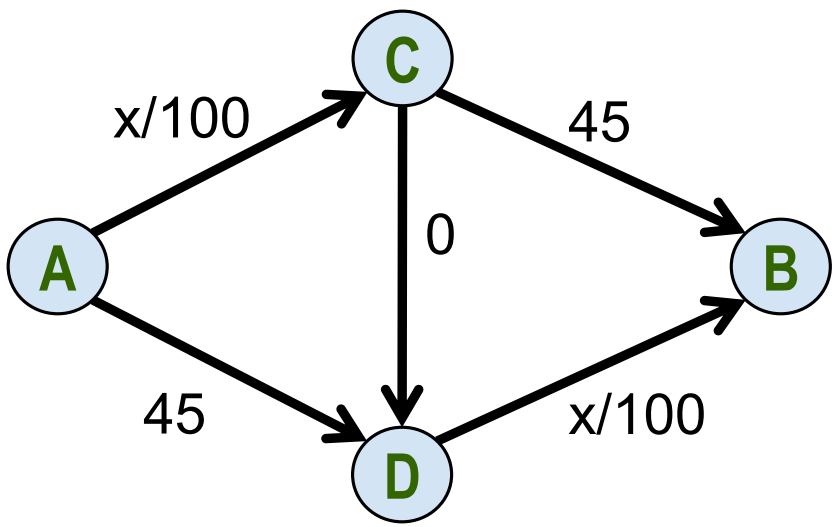

Braess' paradox continued

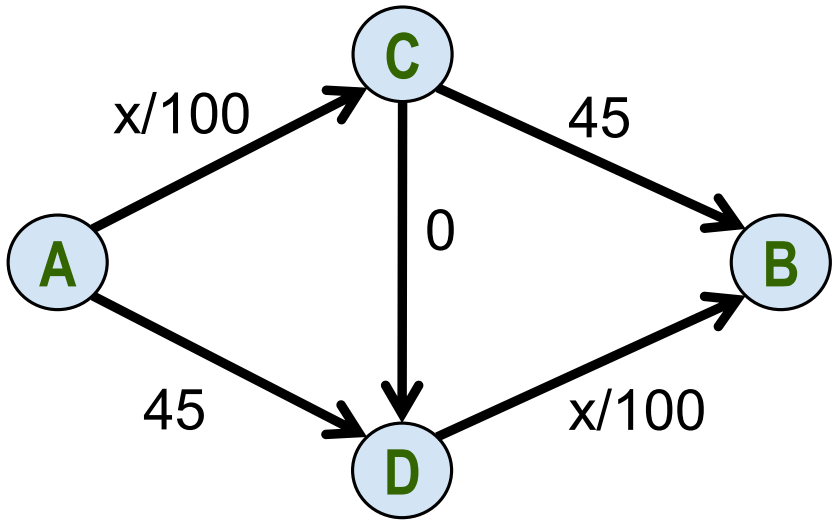

Claim: There is a new unique NE. Everyone now will want to take the route $A \to C \to D \to B$. And the individual commute time of this NE is 80 minutes! That is, by building the new superhighway (rail link) everyone has an additional 15 minutes of commuting.

Proof of claim for Braess' paradox

Claim: All 4000 commuters taking $A \to C \to D \to B$ is an NE

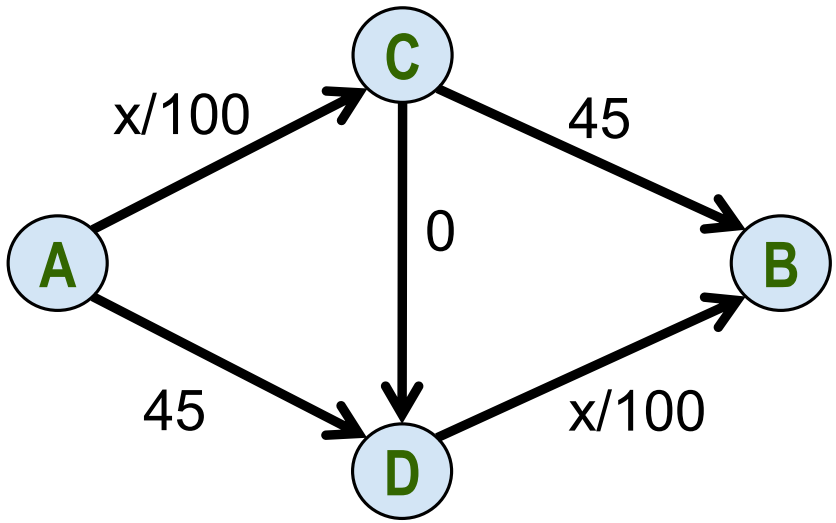

- Consider any individual wanting to deviate:
  - Deviating by taking the direct (A, D) edge is worse (for the one person deviating) than taking the indirect path to D via C
  - Therefore the potential deviating commuter will want to first go to C
  - From C, it is better to take the indirect path (via D) to B than taking the direct (C, B) link

Braess' paradox continued

- An equivalent way to state Braess' paradox:
  - In some traffic networks, closing a road or rail link might speed up the commute time! ... assuming that individuals will find their way to an equilibrium
  - This has been observed in some cases

The new link and social welfare

Is there any sense in which this new link can be beneficial?

Consider the social welfare that is now possible with the new link

- We now have three paths amongst which to distribute the load
- Furthermore, the old NE is still a possible solution, therefore the new optimal social welfare solution cannot be worse

Claim: The following is a socially optimal solution:

- 1750 take $A \to C \to B$ route
- 500 take $A \to C \to D \to B$ route
- 1750 take $A \to D \to B$ route

Society wins but many people lose

We can compare the solution welfare in this new "improved highway" network compared to the social welfare in the original network.

- 500 commuters taking the $A \to C \to D \to B$ route will each have travel time 45 minutes saving 20 minutes each in comparison to the 65 minute commute time without the new 0 cost link.
- On the other hand, the $1750 + 1750 = 3500$ commuters taking the more direct $A \to C \to B$ or the $A \to D \to B$ routes will each have travel time 67.50 minutes incurring an additional 2.5 minutes of commute time.

So the *total time* saved is $(500 \times 20 - 3500 \times 2.5) = 1250$ minutes each way, each day. **On average** (over the 4000 commuters), it is a saving of $1250/4000 = .3125$ minutes per commuter. If this doesn't sound sufficiently impressive, suppose time was being measure in hours; that is, we can scale the edge costs by any fixed factor.
And beyond time lost, a social optimum reduces pollution.

Understanding the partition into 3 routes

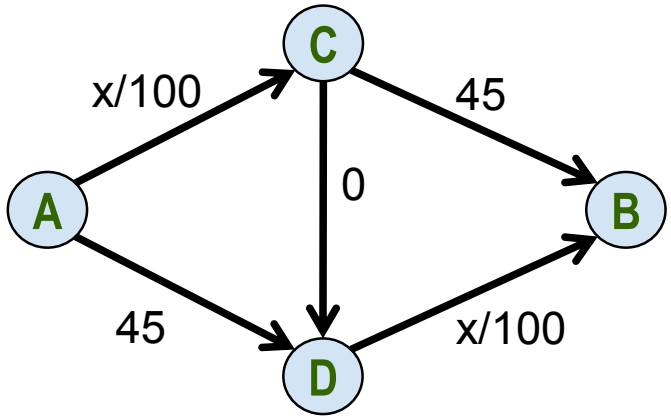

How do we argue the previous solution is a social optimum and how do we find this partition of routes?
We can lower bound the optimal solution by solving a simpler case. By introducing the edge (D, C) all previous traffic patterns are valid.

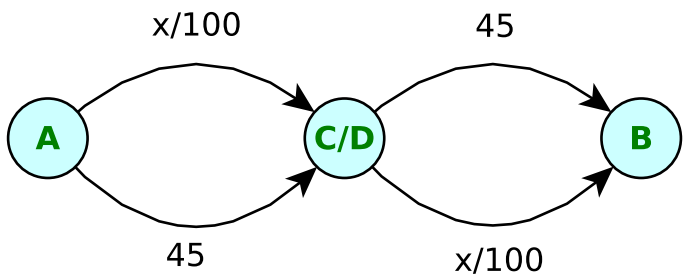

Understanding the partition into 3 routes

We can prove it is optimal by solving a quadratic equation to determine the x commuters who will directly go to C (i.e, take the $x/100$ route) and the $4000 - x$ that will directly go to D (i.e., take the 45 route).

Total time is: $T(x) = x \cdot \frac{x}{100} + (4000 - x) \cdot 45 = .01x^2 - 45x + 180000$.

With this, we're practically done!

- Taking the derivative and setting it to 0, we get: $T'(x) = .02x - 45 = 0$
  - Solving, we get $x = 2250$
    - ★ 2250 will take the $A - C$ route
    - ★ 1750 will take the $A - D$ route
- We can now pretend that all vehicles start at node C (since we merged nodes C and D), and repeat the above process
  - 2250 will take the $D - B$ route
  - 1750 will take the $C - B$ route

Understanding the partition into 3 routes

This is an optimal solution to a relaxation where we merged C and D; as the solution is valid in the original, optimality is preserved.

Therefore in the optimal solution we have 2250 going to C (with 1750 going on directly to B and 500 taking the $C - D$ road) and 1750 going to B via the $A - D - B$ route.

How could the government obtain the socially optimum solution?

- If the government selects some number (say 500) of commuters (e.g. those involved in essential services) then we can achieve the social optimum (e.g. HOV lanes). Or it can allow commuters to buy a special license for the road (e.g. HOT lanes) and hopefully let self interest lead to the social optimum
- Another implicit way to hopefully influence drivers to converge towards the socially better equilibrium is to place a toll on the new link. By adjusting the pricing on the new link, the idea would be that commuters who have the money and value their time more would start taking the new route
- They could alternatively limit the number of commuters taking the $C - D$ road by telling commuters (by say signs at the entrance to the highway system) when the road is open or closed for the commute

The Tragedy of the Commons and the Price of Anarchy

- If we believe commuters will converge to a NE, then allowing commuters to act in their own interest has a "price" (with respect to social optimality)
  - In this network road example, the price is the additional total time (1250 minutes) to commute

- This price of self interest in this or any setting where self interest is a factor is often referred to as the Tragedy of the Commons.

- The algorithmic game theory literature defines a quantitative measure of the price we pay for self-interest with respect to social optimality:

- The Price of Anarchy (POA) for any such specific "game" (where the social objective is a cost function) is a worst case ratio measuring the cost of stability; namely, taking the worst case over all NE solutions S, it is defined as : $\frac{cost(S)}{cost(OPT)}$ where OPT is an optimum solution.
- Note: In general, there can be many pure and mixed NE

The Price of Anarchy continued

The Price of Anarchy was introduced by Papadimitriou.

For a more optimistic perspective there is also a Price of Stability defined as: $\frac{cost(S)}{cost(OPT)}$ where now S is a NE solution having the least cost.

Returning to the specific setting of network congestion, the following two results (due to Roughgarden and Tardos) are early seminal results in algorithmic game theory. For *all congestion networks with linear cost functions*:

1. The POA is no more than $\frac{4}{3}$
2. This result is tight in the sense that if we change the fixed cost in the simple 4 node network from 45 to 40, the POA would be $\frac{4}{3}$.

Recap

- Congestion Networks
  - Traffic networks (Ch 8.1)
  - Nash Equilibrium
  - Braess' Paradox (Ch 8.2)
  - Social cost
    - ★ Tragedy of the commons
    - ★ Price of anarchy

New topic: Kidney exchanges

Although this is not a topic I was planning for the final exam, the topic of kidney exchanges is technically interesting and, of course, critically important for many people.
Some (admittedly outdated) statistics:

- In the US, each year there are 50,000 new cases of potentially lethal kidney disease
- There are two possible treatments: dialysis or transplant
- Transplants can come from live donations or from transplants for someone who has just died (e.g., in car accident). All else being equal, live donations are much more successful
- Each year there are $\approx 10,000$ transplants from someone deceased and $\approx 6,500$ from live donations
- The waiting list for a transplant in the US is $\approx 75,000$ people who usually wait between 2 and 5 years. During this waiting time, ≈ 4000 people die each year

More facts concerning kidney exchanges

Live donations are possible since everyone has two kidneys and only one is needed. Moreover, when people incur kidney diseases, usually both kidneys are affected so the "additional kidney" is rarely needed.

However, people are reluctant to donate kidneys and live donations usually come from close relatives and friends.

There are many biological compatibility requirements in order to do a transplant so there is often no one available and willing to do a donation.

- Blood compatibility
- Tissue compatibility

Even if possible, some donor-recipient transplants are better than others.

Pairing up transplants

So if a willing donor for a recipient is not compatible (or if the match is not that great), there may be another recipient-donor pair that are having the same issue and are willing to do a 'swap". Consider the following possibility for a pair swapping:

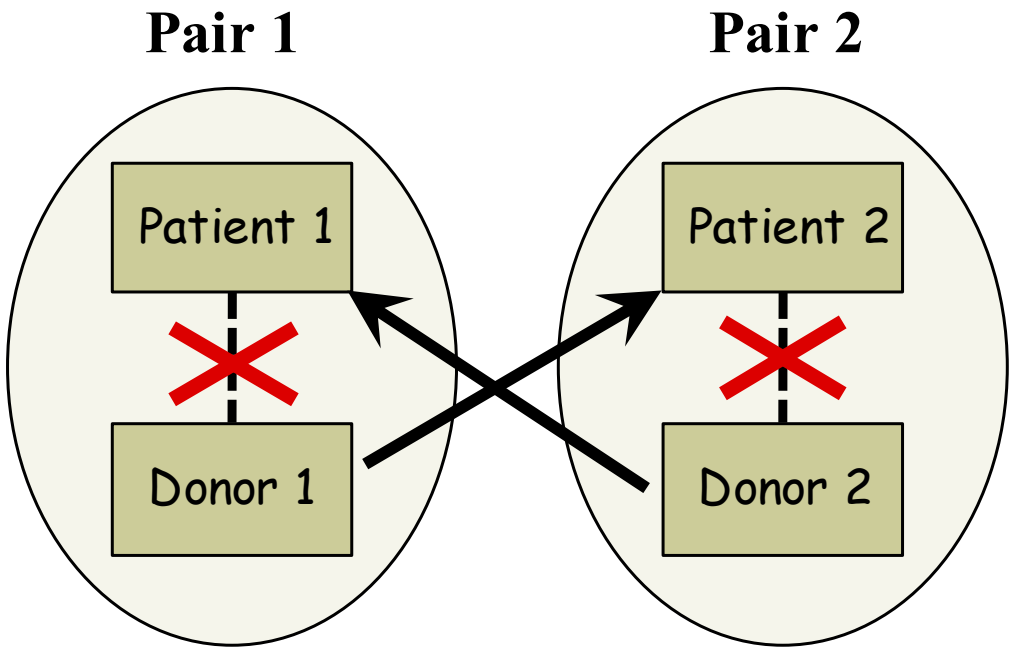

Here an edge means that the Patient (i.e. the recipient) and Donor are compatible. Edges can be weighted to reflect some objective as to how good is the match. The weight could also reflect geographic distance.

Extending to bigger cycles

The idea of pairs swapping as just illustrated was first proposed in 1986 and only realized in 2003.

This idea has been extended to bigger cycles as in the next illustration:

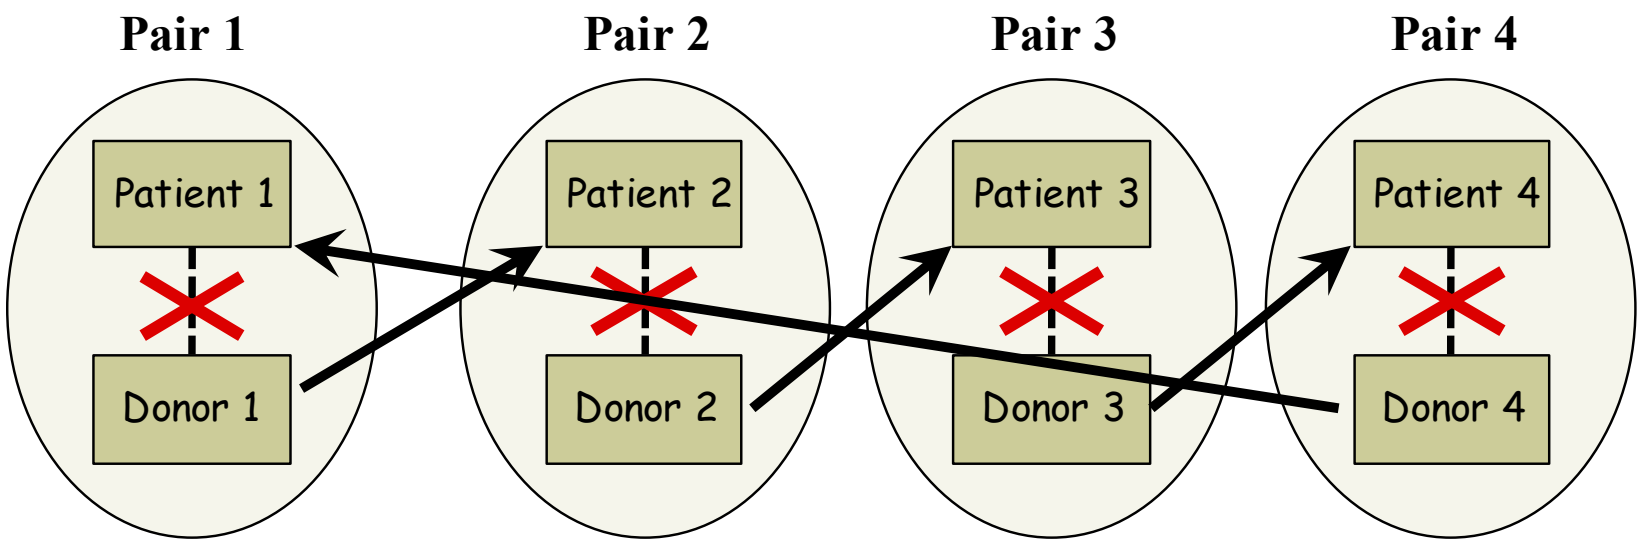

How practical are such swaps and cycles?

The are "logistical" issues that impact the practicality of such swaps and cycles, and the bigger the cycle the more problematic logistically.

What if a potential donor, say Donor i reneges (or dies, or gets ill) once his/her paired recipient Patient i has already received their (from Donor $i-1$) kidney from the person with whom they are compatible? Now Patient $i+1$ has lost a valuable resource his/her (i.e., the intended Donor they brought to the exchange) if Donor $i+1$ has already given their kidney to Patient $i+2$.

This requires that the donation and transplant must all basically be done *simultaneously*. For cycles of length k, this requires $2k$ simultaneous operations, where each transplantation requires both a donation and transplant operation.

Furthermore, live kidneys from donors travel best inside the donor, so need these operations to be geographically close (i.e. same or nearby hospital). Note: Some hospitals will not accept organs transplanted by air.

The net effect is that this severely limits the length of cycles in practice.

Altruistic donors

What happens if we have one altruistic donor who is willing to donate a kidney without having someone with whom he/she wishes to be paired? Once there is such an altruistic donor, we can eliminate the need for simultaneity, and create an altruistically initiated donor chain.

After we have an altruistic donor, we can proceed in what potentially can be an arbitrarily long chain as below. Here each Patient must still be willing to bring a willing Donor to the exchange. But now if some donor reneges, etc, the next recipient has not lost their paired donor.

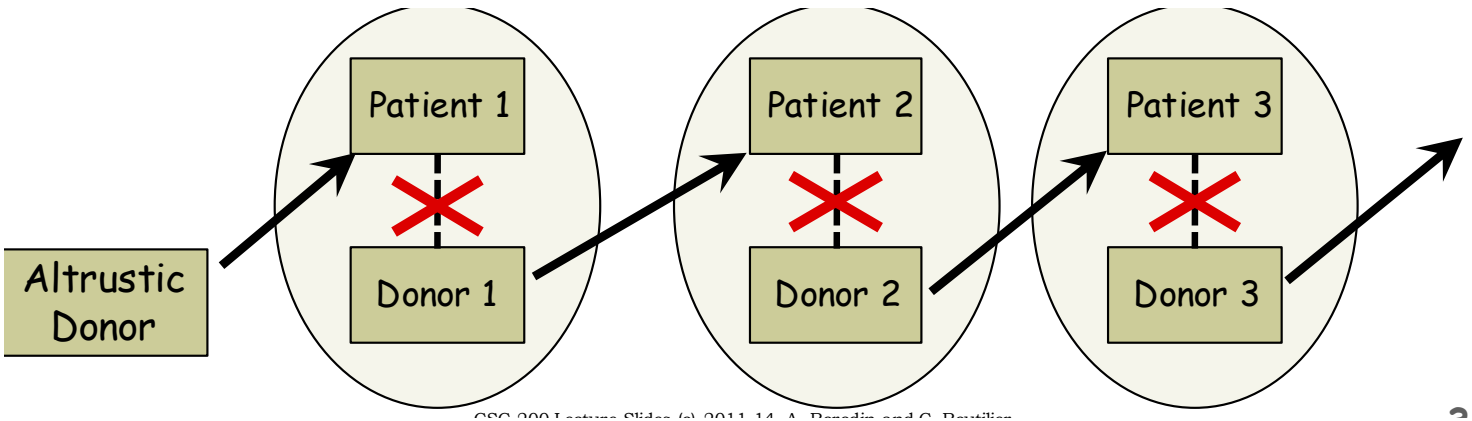

There has been at least one chain of length 30 (ending in February 2012) and some chains may be still be ongoing.

Some final comments

- Due to biological, logistical, and incentive issues, kidney exchange requires efficient algorithmic solutions
  - Fairly large scale networks; i.e., say tens of thousands of nodes when considered nationwide.
- When restricted to pairs, this is a (possibly weighted) matching problem in a non-bipartite graph
- Introducing cycles and chains makes the problem becomes much harder
  - Need to find "practically feasible" cycles and chains
- Additionally, the market is not static; there are arrivals and departures:
  - Is it better to use a current match, or wait for new donors and recipients to arrive?
  - When an altruistic donor arrives, do you use up that valuable resource now or wait for a better match that might lead to a longer chain?
  - Are there incentive issues for say hospitals to want to do more of the transplants by themselves than join in a broader exchange?

A recap of the course

- I would say that the central theme of the course is the attempt to more precisely model sociological phenomena
  - I'd argue this includes "information networks" (e.g., the web) as it is humans that create this network
  - How we link and rank documents, and "navigate" within this network of documents fits into social networks

- The main mathematical framework (and hence the course name) centres around networks
  - Modeling social networks presents significant challenges
  - Often, we have only initial insights and we are far from realistic models and analysis of social phenomena

Recap continued

- As current social networks are often extremely large, it is necessary to be able to "think algorithmically"
- It's also necessary to appreciate the fundamental insights and studies that have evolved and continue to evolve from sociology, economics, biology, physics, and other fields
  - While we can often abstract problems to the same framework, understanding the area, and the constraints and details lost in abstraction, is crucial for correctness!
- Reason about stochastic models is also often necessary
- As the text often emphasizes, in what may be called algorithmic social networks, the approach taken follows what we see in other sciences
  - Our initial efforts are informed by real world networks and phenomena
  - We then formulate precise models, draw some insights and possibly some preliminary conclusions, and then calibrate the model and insights against real world or synthetic data
  - Based on the experimental results, we are then able to iterate the process; that is, modify the model and continue to draw insights and again evaluate by experiments
  - These insights may also be directed towards concrete applications!

Recap continued

The text properly cautions that these models are just that, *only models* of real world network behaviour and that we are often far from having confidence in any preliminary conclusions

In other words, all together now... It's reduced to *lines* and *circles*!

We can get shockingly good results from such a crude abstraction, but we should never forget the simplifying assumptions we've made, and the information we've left behind

In some cases, it is surprising how much information one can obtain just from basic network models and assumptions!

- Identification of romantic ties in the Backstrom and Kleinberg paper
- Labelling of strong and weak ties in the Sintos and Tsaparas paper

Needless to say however, the more we know about the content relating to the nodes and edges in a network, the more we should be able to make informative findings

Some of the major topics in the text and the course

Here are some of the major topics in course:

- The concept of strong and weak ties and their relative role in obtaining "social capital"
- Different types of closing of triangles: triadic closure, focal closure, membership closure
- Homophily and influence
  - Are our friendships derived from similar interests and behaviour, or are our friendships influencing our interests and behaviour?
  - A central issue in social relations and one where any findings can be controversial!
    - ★ e.g., whether or not "obesity is contagious"
- A number of topics relate to different equilibrium concepts
  - Schelling's segregation model, structural balance in friend/enemy networks, balanced outcomes in bargaining networks, stable matchings, and Nash equilibria in a congestion network, debatably page rank, etc...

Some major topics continued

- A number of topics relate to navigation in a social network and in particular to the small world phenomena based on geographic or social distance
  - Also was related to power law distributions in social and information networks
- Influence spread in social networks and disease spread in contact networks
- Am I missing any major themes that we discussed?

Thank you!

You've been a wonderful class, and it's been a pleasure teaching you. I hope to see you in the last lecture, but if not, then best of luck on the final and have a great summer!