

Social and Information Networks

University of Toronto CSC303
Winter/Spring 2019

Week 11: March 25,27 (2019)

Announcements and agenda

Announcements

- All the questions for the final assignment have been posted. Questions 1-6 are due this Friday, March 29 by 3PM so that Tyrone can go over any questions. Since I am only doing stable marriage today, question 7 is optional. However, there could likely be a question concerning stable marriage on the final exam so you should try to do that question at some time.
- Later this week, I will also post an optional problem concerning congestion networks as that also could be a question on the final exam.
- If you do not submit your assignment by 3PM on Friday, Markus will not accept the assignment. Do not send the assignment to me.

Today's agenda.

- 1 For the remainder of the course, we will discuss the Stable Marriage problem and then Congestion Networks and Braess' paradox.

New topic: The stable marriage problem

Note: This material is not in the text. I am not sure if this can be viewed as part of social choice theory, but I know it has been covered in CSC304.

However, I do think it fits in nicely with the focus of CSC303. Namely, as in our last topic we will be concerned with graph matching but now restricted to bipartite graphs. And we will also be led to another important example of a “coalition equilibrium”.

The stable marriage problem and the Gale Shapley algorithm, is interesting for a number of reasons.

- Mainly because it has practical application, and it is still actively considered due to variants arising from applications.
- The algorithm is elegant and the analysis is interesting.

End of Week 10

We ended at slide 40, just beginning to introduce the stable marriage problem and the Gale Shapley Deferred Acceptance algorithm. I am leaving in the remainder of my slides. I will soon augment them and the post slides for Week 11 in advance.

Preferences vs utilities

In game theory and mechanism design, individual valuations are numeric utilities (e.g., money). In contrast in social choice theory (e.g., forming consensus as in voting) and in the stable marriage problem, individuals have preferences (that do not necessarily get translated in numeric values).

A *preference* over a set A of alternatives (e.g., candidates) is a total or partial order (also called an ordering or ranking) of the alternatives.

In many cases, we may have a hard time placing values on alternatives but we may surely know that we like alternative a_1 relative to alternative a_2 .

Suppose $A = \{a_1, a_1, \dots, a_n\}$. Consider an individual (say k). We will use \succ_k (or \prec_k) to denote k 's preference between alternatives when k has such a preference. That is, $a_i \succ_k a_j$ (alternatively $a_j \prec_k a_i$) if k definitely prefers a_i to a_j .

Total orders vs partial orders

Of course, sometimes we are not so sure about our preferences. We can use $a_i \succeq_k a_j$ to indicate that k likes a_i at least as much as a_j . And it is often the case that there are two alternatives for which we have no relative opinion.

A *total order* \succ on a set of alternatives $A = \{a_1, a_2, \dots, a_n\}$ satisfies the following:

- \succ is transitive; that is, $a_i \succ a_j$ and $a_j \succ a_\ell$ implies $a_i \succ a_\ell$.
- There is a permutation π such that $a_{\pi(1)} \succ_k a_{\pi(2)} \dots \succ_k a_{\pi(n)}$.

A *partial order* \succ satisfies the following:

- \succ is transitive
- There is a way to extend the order (i.e., to all a_i, a_j such that neither $a_i \succ a_j$ nor $a_j \succ a_i$ is given) so as to make \succ into a total order.

Two-sided matching markets

In a two-sided matching market, we are interested in a matching in a graph/network where :

- There are two sets of agents X and Y .
Note: X and Y can be the same set in some applications. This was the situation in the study of network exchanges under the 1-exchange rule assumption. It is also the situation in a kidney exchange market.
- The goal is to match agents in X to agents in Y to satisfy some objective.
- Agents have the ability to leave unfavourable matches so as to obtain a more favourable match.

Note: As we remarked in our discussion of network exchanges, we are generally interested in b matchings in many applications where say agents (and in the bipartite case, maybe only agents on one side of the graph) can be involved in up to b edges. But for now, and in keeping with the terminology of a marriage, let us restrict our attention to the standard definition of a matching.

The bipartite case and the stable marriage problem

In the stable marriage problem, we are interested in matchings in a bipartite graph $G = (V, E)$ where $V = X \cup Y$. Furthermore, we assume that every agent X has a total preference order over Y and every Y has a total preference order over X . This total order assumption, and the restriction to matchings and not b -matchings, can be eliminated (say for the basic Gale-Shapley stable marriage algorithm) but they can present issues in some applications.

Applications:

- Matching employess to specific positions (or tasks). In particular, match medical school graduates to specific residence positions.
- Matching Men and Women in marriages. This is the classical terminology used and we will stay with that terminology which at least motivates the assumption of a matching rather than a b -matching.

Aside: Arguably the most important application of the Gale-Shapley algorithm for the stable marriage problem (and variants of that problem and algorithm) is in matching doctors to residency positions at hospitals.

Stable marriages

First some notation:

Let the set of men be M (with $m \in M$) and let W be the set of women (with $w \in W$). For simplicity, we will assume $|M| = |W|$.

Let μ denote a matching; that is, $\mu(m)$ is the woman matched to m and $\mu^{-1}(w)$ is the man matched to w . Abusing notation, we will just use $\mu : M \rightarrow W$ as a 1-1 mapping between men and women.

Similar to the issue of stability in the network exchange process, the most basic objective is to find a maximum (in this case perfect since we assume $|M| = |W|$) matching between M and W that is *stable* in the following sense:

A stable matching in the stable matching problem

. A matching μ is *unstable* if there exists an unstable (also called *blocking*) pair (m, w) such that m prefers w to his current match $\mu(m)$ and w prefers m to her current match $\mu(w)$. In this case, m and w will leave their current matches to be with each other. A match is *stable* if it contains no unstable (blocking) pairs.

Some examples of stable and unstable matchings

We have to check for the presence or absence of a blocking pair; that is, a pair (m, w) such that $w \succ_m \mu(m)$ and $m \succ_w \mu(w)$.

Here are a set of preferences for the men and women :

Man	1st	2nd	3rd
x	a	b	c
y	b	a	c
z	a	b	c

Woman	1st	2nd	3rd
a	y	x	z
b	x	y	z
c	x	y	z

Which of the following matchings are stable/unstable?

- Matching 1: $a - x, b - y, c - z$ **Stable?**
- Matching 2: $a - y, b - x, c - z$ **Stable?**
- Matching 3: $a - z, b - y, c - x$ **Stable?**

Some examples of stable and unstable matchings

We have to check for the presence or absence of a blocking pair; that is, a pair (m, w) such that $w \succ_m \mu(m)$ and $m \succ_w \mu(w)$.

Here are a set of preferences for the men and women :

Man	1st	2nd	3rd
x	a	b	c
y	b	a	c
z	a	b	c

Woman	1st	2nd	3rd
a	y	x	z
b	x	y	z
c	x	y	z

Which of the following matchings are stable/unstable?

- Matching 1: $a - x, b - y, c - z$ **Stable?**
- Matching 2: $a - y, b - x, c - z$ **Stable?**
- Matching 3: $a - z, b - y, c - x$ **Stable?**

In Matching 3, we can see that (b, x) is a blocking pair. **What other blocking pairs exist?**

Stability as an equilibrium

Stability is an equilibrium concept. But like stability in the network exchange setting, and unlike Nash equilibrium, it takes two people to conspire to deviate. In the network exchange setting that was built into the experiments.

This is a form of *coalitional stability*

In some versions of the stable matching problem, we allow individuals to remain “unmarried”. This can be incorporated into the problem formulation by letting each man m (respectively, each woman) to put himself (respectively, herself) into his (resp, her) preference ordering \succ_m (resp. \succ_w).

For example, if we have $m_1 \succ_w m_2 \succ_w w \succ_w m_3 \dots \succ_w m_n$ then w would rather be by herself than with anyone other than m_1 and m_2 .

Do stable matchings always exist and, if so, how do we find them?

Aside: When there are n men and women, there are $n!$ possible matchings so we certainly cannot exhaustively check all matchings. And even if we could for a given instance of the problem (i.e., a set of preferences for each man and woman) that would not determine if there is always a stable matching.

Fortunately, we have the Gale Shapley algorithm which constructively and efficiently shows how to compute a stable matching for any instance.

There are two standard analogous varieties of the Gale Shapley algorithm:

- 1 Man proposes, woman disposes. Also called Male Proposing Deferred Acceptance (MPDA)
- 2 Female proposes, man disposes. Also called Female Proposing Deferred Acceptance (FPDA)

The FPDA and MPDA are completely analogous But in general, they will produce different matchings.

The FPDA algorithm

- The algorithm will proceed in rounds, at the end of each round, each woman will have a set P_w of people to whom they have previously proposed. There will also be a set C of current engagements. Both sets are initially empty.
- In each round t , every unengaged woman w proposes to the man $m \notin P_w$ that is highest in her preference ranking \succ_w . If every woman is engaged at the start of a round, the algorithm terminates.
- After a round of female proposals, every man m will consider his set $P_{m,t}$ of current proposals (if any).

We consider what each man m does in this round.

- 1 $P_{m,t} = \emptyset$, then m does not do anything in this round.

So now consider the case that $P_{m,t} \neq \emptyset$, and let w^* be the most preferred woman in $P_{m,t}$. That is, $w^* \succ_m w'$ for every $w' \neq w^* \in P_{m,t}$.

- 2 If m is not currently engaged, he will become engaged to w^* and C is updated accordingly.
- 3 If m is currently engaged to w (i.e., $(m, w) \in C$), then he will break this engagement if and only if $w^* \succ_m w$ and will then become engaged to w^* . In this case, $C := C \setminus \{(m, w)\} \cup \{(m, w^*)\}$

A running example for the FPDA algorithm

Wmn	1st	2nd	3rd	4th
a	x	y	z	w
b	y	x	w	x
c	x	y	z	w
d	y	w	x	z

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

Round 1

Proposals: New Engagements:

a: x

w: -

b: y

x: a

c: x

y: b

d: y

z: -

Example: Round 2

Wmn	1st	2nd	3rd	4th
a	x*	y	z	w
b	y*	x	w	x
c	x*	y	z	w
d	y*	w	x	z

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

A * indicates that the man has already been proposed to by this woman.

Round 2

Current: Proposals: New Engagements:

w: -

x: a

y: b

z: -

a: -

b: -

c: y

d: w

w: d

x: a

y: ~~b~~ c

z: -

b is "jilted"

Example: Round 3

Wmn	1st	2nd	3rd	4th
a	x*	y	z	w
b	y*	x	w	x
c	x*	y*	z	w
d	y*	w*	x	z

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

A * indicates that the man has already been proposed to by this woman.

Round 3

Current: Proposals: New Engagements:

w: d

x: a

y: ~~b~~ c

z: -

a: -

b: x

c: -

d: -

w: d

x: ~~a~~ b

y: ~~b~~ c

z: -

a is "jilted"

Example: Round 4

Wmn	1st	2nd	3rd	4th
a	x*	y	z	w
b	y*	x*	w	x
c	x*	y*	z	w
d	y*	w*	x	z

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

A * indicates that the man has already been proposed to by this woman.

Round 4

Current: Proposals: New Engagements:

w: d
x: ~~a~~ b
y: ~~b~~ c
z: -

a: y
b: -
c: -
d: -

w: d
x: ~~a~~ b
y: ~~b~~ c
z: -

**a's proposal
not accepted by y**
(no change)

Example: Round 5

Wmn	1st	2nd	3rd	4th
a	x*	y*	z	w
b	y*	x*	w	x
c	x*	y*	z	w
d	y*	w*	x	z

Man	1st	2nd	3rd	4th
w	d	b	a	c
x	b	a	d	c
y	c	b	a	d
z	d	b	c	a

A * indicates that the man has already been proposed to by this woman.

Round 5

Current:

w: d

x: ~~a~~ b

y: ~~b~~ c

z: -

Proposals:

a: z

b: -

c: -

d: -

New Engagements:

w: d

x: ~~a~~ b

y: ~~b~~ c

z: a

Stable:

a:z

b:x

c:y

d:w

FPDA results in a stable matching

Here are the ingredients of a proof that FPDA (and similarly MPDA) will always produce a stable matching.

For simplicity, we are assuming the same number, say n , of woman and men, we have to show that the algorithm always terminates and when it terminates, it results in a stable perfect matching. So first, why does the algorithm terminate.

- At the end of each round, we have a partial matching in terms of the current engagement.
- In each round, before the men get their chance to approve or disapprove, every woman has a current engagement or has proposed.
- In each round, if every man is satisfied, the algorithm terminates in what must be a perfect 1-1 perfect matching. Otherwise, at least one man has improved upon his last match.
- Since there are n woman and n men, **FPDA must terminate in at most n^2 rounds** for every (n, n) input.

Why is this matching stable?

Why is the FPDA matching stable?

Let μ be the matching produced by the FPDA. Assume (m, w) is a blocking pair for some man m and woman w .

This means that w must prefer m to $\mu(w)$ and hence must have proposed to m before proposing to $\mu(w)$.

By the assumption that (m, w) is a blocking pair, m prefers w to $\mu(m)$. This means that either

- 1 He would have rejected $\mu(m)$ if w proposed to m *after* $\mu(m)$ OR
- 2 He would not have accepted the proposal from $\mu(m)$ if w proposed *before* $\mu(m)$, since he would already be engaged to either w or someone even more preferred than w .

It follows that μ is stable since there cannot be a blocking pair.

Different stable matchings

In algorithm design (without any self interest by agents), we would be interested in algorithms that produce a maximum matching or in the edge or vertex weighted cases, a maximum weight matching.

Do we have a sense of how good is a given stable matching?

(Note: There can be exponentially many stable matchings for some instances. And some instances have a unique stable matching.)

Since we do not have numeric values (only preferences) for any individual or edge, it may not be clear at first why we would prefer one stable matching to another.

There are ways that we can define the *social welfare* of a stable matching. It is always possible (e.g., use the Border scoring rule) to transform a preference ranking to a utility for the agents based on the match they receive. We will not be interested in this issue.

However, for basic applications of the FPDA (or MPDA) we will not be interested in a numeric objective function.

Properties of the FPDA (MPDA) algorithm

From the analysis of the FPDA stability, we know that FPDA always terminates within n^2 rounds.

And we know that there are (n, n) instances on which FPDA will use $\Omega(n^2)$ rounds. **Can you construct such an instance?**

Although, we will not be concerned with social welfare, we can ask how satisfied will either the men or women be in a stable matching produced by the FPDA and MPDA algorithms.

One more important property I failed to mention in the lecture. Clearly, the order in which women propose in a given round does not depend on the order in which they propose. Since the same woman cannot propose to more than one man in a round, it also doesn't matter in what order the men accept or refuse new proposals. That is, the same woman w^* cannot be the reason for cancelling more than one engagement. Thus the matching of FPDA is completely determined no matter what order the woman propose or the order that the men make or break engagements.

Female-optimal and male-optimal stable matchings

Define $OPT(w)$ (resp. $Pess(w)$) to be the most (resp. least) preferred man she could be matched with *in a stable matching*. This is a well defined concept since there can only be a finite number of stable matchings.

A matching is *female-optimal* if every woman w is match to $OPT(w)$. **Is such a matching possible?**

Female-optimal and male-optimal stable matchings

Define $OPT(w)$ (resp. $Pess(w)$) to be the most (resp. least) preferred man she could be matched with *in a stable matching*. This is a well defined concept since there can only be a finite number of stable matchings.

A matching is *female-optimal* if every woman w is match to $OPT(w)$. **Is such a matching possible?**

Good news for women

Theorem: FPDA results in a female-optimal stable matching for all input instances.

We can also define a male-optimal stable matching in the same way.

Good news for men

Theorem: MPDA results in a male-optimal stable matching for all input instances.

Bad news for society?

FPDA (resp. MPDA) results in a *male-peessimal* (resp. female-peessimal) stable matching for all instances.

End of Monday, March 25 lecture

We ended just flashing slides for the proof of FPDA stability. We will quickly do that again.

Today (Wednesday, March 27) we will finish our discussion of the FPDA and some extensions of the basic stable matching problem we have been considering thus far. I also want to mention another property of the FPDA algorithm that I failed to mention last lecture. (See slide 22.)

Sketch of proof that FPDA is female-optimal

Suppose that FPDA is not female-optimal. Then in the FPDA, some woman w has wound up with some $m \neq OPT(w)$.

Lets suppose $m_1 \succ_w m_2 \dots \succ_w m_n$ and let $m = m_i = OPT(w)$ for some i . For each $j \leq i$, m_j must have rejected w in the FPDA or else FPDA would have produced a stable matching with w matched to m_j . (For $j < i$ this would contradict the assumption that m_i is the best she can do in a stable matching.) So, in particular, w has been rejected by $m = OPT(w)$

Let t be the earliest round in which some woman w has been rejected by $m = OPT(w)$. If there is more than one such w at this round, let this be the first such rejection in this round. Then $OPT(w)$ must be engaged to some $w^* \succ_w OPT(w)$ at round t and $w^* \succ_m w$.

We know that w^* was not rejected by $OPT(w^*)$ by round t since we defined t to be the first round where this happens.

Sketch of proof that FPDA is female-optimal continued

Since w^* has also been proposing in order of her preferences, we must have $OPT(w) = OPT(w^*)$ or $OPT(w) \succ_{w^*} OPT(w^*)$ since we know that $m = OPT(w)$ was engaged to w^* at round t . Abusing notation, denote this by $OPT(w) \succeq_{w^*} OPT(w^*)$.

Now consider a stable matching μ in which w and m are matched and say that $m^* = \mu(w^*)$. Then $OPT(w^*) \succeq_{w^*} m^*$ by definition of $OPT(w^*)$.

We then have $m = OPT(w) \succeq_{w^*} OPT(w^*) \succeq_{w^*} m^*$. At least one of these is a strict preference \succ_{w^*} since we know that $w^* \succ_m w$.

But now (m^*, w^*) is a blocking pair in the matching μ .

Should you be truthful about your preferences?

It does seem reasonable for women to propose in order of their preferences and men to accept their best offer. So why should anyone manipulate and not be truthful about their preferences?

Should you be truthful about your preferences?

It does seem reasonable for women to propose in order of their preferences and men to accept their best offer. So why should anyone manipulate and not be truthful about their preferences?

However, the Gale-Shapley algorithm can be manipulated. That is, there are instances where someone can wind up better off by not stating their true preferences. Here is an example:

First, consider the truthful set of preferences:

$m_2 \succ_{w_1} m_1 \succ_{w_1} m_3 \succ_{w_1} m_4$

$m_4 \succ_{w_2} m_1 \succ_{w_2} m_2 \succ_{w_2} m_3$

$m_1 \succ_{w_3} m_3 \succ_{w_3} m_2 \succ_{w_3} m_4$

$m_4 \succ_{w_4} m_3 \succ_{w_4} m_2 \succ_{w_4} m_1$

$w_1 \succ_{m_1} w_2 \succ_{m_1} w_3 \succ_{m_1} w_4$

$w_2 \succ_{m_2} w_1 \succ_{m_2} w_3 \succ_{m_2} w_4$

$w_3 \succ_{m_3} w_1 \succ_{m_3} w_2 \succ_{m_3} w_4$

$w_4 \succ_{m_4} w_2 \succ_{m_4} w_3 \succ_{m_4} w_1$

FPDA will compute the following stable matching:

$(w_1, m_2), (w_2, m_1), (w_3, m_3), (w_4, m_4)$

You should check this by running FPDA.

But what if m_2 is not always truthful?

Suppose that m_2 lies in round 2 and rejects the proposal from w_2 (staying engaged to w_1 even though $w_2 \succ_{m_2} w_1$).

This will result in the following matching:

$(w_1, m_2), (w_2, m_1), (w_3, m_3), (w_4, m_4)$

where now m_2 is matched to w_2 , an improvement for him.

You should check this by running FPDA with m_2 deviating as indicated.

NOTE: It is not easy to prove but in FPDA, women can never benefit by being untruthful. That is, women should always propose in the order of their preferences when using the FPDA.

Of course, it is just the opposite when using MPDA: Men cannot benefit from lying but women can sometimes gain by an untruthful rejection.

Lots of extensions of deferred acceptance (DA) and other considerations

- Many applications are *many-to-one*) and not 1-1 as in the basic formulation. For example, a University accepts many students. This extension is not difficult to handle.

One way would be to replicate a University K times if it had a quota of K students. **Is this a good solution?**

Lots of extensions of deferred acceptance (DA) and other considerations

- Many applications are *many-to-one*) and not 1-1 as in the basic formulation. For example, a University accepts many students. This extension is not difficult to handle.

One way would be to replicate a University K times if it had a quota of K students. **Is this a good solution?**

This is inefficient (especially if K is big and it imposes an artificial ranking amongst the copies).

Instead, we can extend Gale-Shapley by having each University have a quota and while that quota is not filled, they keep admitting students. When the quota is filled and they get another request, they can reject it or take it and remove the least desirable student. (Of course, they don't announce any decisions until the end of the admission process and hopefully have a reliable way to rank students.) Now Universities (the men in FPDA) can also manipulate by misreporting their quota.

Other important considerations in stable matching

- Partial preferences. In general, our preference relation is probably transitive (but not always) and usually incomplete. That is, we may not have a preference between various choices. Now there can be different ways to define stability.
 - 1 Weak stability: (m, w) is a blocking pair iff both m and w are strictly better.
 - 2 Strong stability: (m, w) is a blocking pair iff at least one of m and w is strictly better.
 - 3 Super strong stability: (m, w) is a block pair if neither m nor w is worse off.

Gale-Shapley is easily extended to handle weak stability (i.e., break ties arbitrarily), but strong and super strong stability require modifications.

- Partial preferences raises the issue as to how to possibly resolve some preferences by say interviews. But that can be costly.
- As mentioned, the number of couples graduating medical school has been increasing. (In 2015, 6% of resident applications were coupled.) Couples rank pairs of residency positions. NP-complete problem to determine if there is a stable matching. Various ways of approaching

Concluding stable matching

Very important and still active topic as stable matching is used in a number of applications. In the kidney exchange application (stable matching in a non-bipartite graph whose), it can literally be a matter of life and death.

As another indication of the importance of stable matching, the 2012 Nobel Prize in Economics was awarded to Lloyd Shapley and Alvin Roth for their work in the theory (Shapley) and application (Roth) of stable matching algorithms.

New topic: A congestion game


We will now be considering a game (as in game theory) that models a highway network system. This is the topic of Chapter 8 in the text. We will see a very surprising phenomena. Namely, building more roads in some situations could be harmful.

Here is the model: 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 and the travel time on different roads (i.e., the edges in the network) depend on the number of drivers using that road.

As we said, this is a game, and here the 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. We are saying “roads” here but edges could be links in a commuter rail or subway network. Probably better to just commuters.


A simple but interesting example

Consider the following example.



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

The traffic network example continued



Suppose we have 4000 commuters each making an individual decision whether to travel via C or via D .

Formally speaking, there are 2^{4000} possible outcomes depending on which route each individual takes. But many outcomes are equivalent since we are viewing all commuters as equivalent. So 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); that is, 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.)

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 not only a unique NE, it is also the unique outcome that is optimum 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 . Now the total of the commute times increases since 2001 commuters will increase their commute time by .01 minutes while only 1999 will save .01 minutes so that


It is unlikely that any individual commuter will notice this, but suppose now that 3000 go via C . The total commute time will now increase by 20,000 minutes \approx 2 weeks worth of time. And, if everyone takes the same route, the total commute time will increase by 80,000 minutes \approx 2 months of time. What would happen if everyone started using the same route? Would it be likely that they would *all* switch to the other route?

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?**


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 $A \rightarrow C \rightarrow D \rightarrow B$. 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

- Everyone taking $A \rightarrow C \rightarrow D \rightarrow B$ is an NE. This can be seen by considering 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 . So the potential deviating commuter will want to first go to C and then from C , it is better to take the indirect path (via D) to B than taking the direct (C, B) link.

Another equivalent way to state this paradox is that in some traffic networks, closing a road or rail link might speed up the commute time! And this has been observed in some cases. Of course, all this assumes that individuals will find their way to an equilibrium.

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. Note that we now have three paths amongst which to distribute the load.

Claim: The following is a socially optimal solution:

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

Society wins but some people lose

What is the social welfare of this solution? We have

- 500 commuters taking the $A \rightarrow C \rightarrow D \rightarrow B$ 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 $A \rightarrow C \rightarrow B$ or the $A \rightarrow D \rightarrow B$ 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. **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 measured in hours; that is, we can scale the edge costs by any fixed factor.

So do we build the new road or railway link?

Even if the cost of the new link is not a factor, do we build the new link?

Probably not. **Why?**

So do we build the new road or railway link?

Even if the cost of the new link is not a factor, do we build the new link? Probably not. **Why?**

Most of the commuters now have incurred some additional travel time and will explore other routes. We view this as an unspecified random process where different individuals exploring new routes from time to time resulting eventually in the commuters more likely to return to the solution without the new link where everyone's commute time was 65 minutes. Note that the partition into three routes is not an equilibrium but it is a social optimum.

So in order to achieve the saving in travel time, the government would have to somehow dictate the socially optimum solution. No one would voluntarily want to take the (A, D) or (C, D) links. One 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 value their time more would start taking the new route.

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.

There is something very symmetrical about the network that the new link can now exploit. Note that we can equalize the total time used between going from A to C and from going to A to D by having 2250 going to C and 1750 going to B . This can be determined by solving a quadratic equation to determine x such that x go to C and $(4000 - x)$ go to D .

Looking back from B , the new link allows us to equalize the total time going between going from C to B and from going from D to B . This is what was more easily done in the network without the link by sending equal numbers to C and D .