Perspectives on mechanism design

Juho Hirvonen, Aalto University and HIIT
HIIT Foundations Friday, 27 October 2023

The setting

- We want to implement a ***social choice*:** for example, assign an item to some agent, assign students to school seats, or decide what will be the price of electricity tomorrow
    - The assignment should satisfy some optimisation criteria such as *utility maximisation* or *fairness*
- Participants are strategic agents that have ***private information hidden from the mechanism***
- Algorithm must compute the assignment in a way that ***incentivises agents to participate*** (i.e. reveal enough information to compute the correct outcome)

Today's talk

- What ***is a mechanism?*** I will overview some of the most important mechanisms
- Mechanism design focuses on ***truthful mechanisms***: this follows from the ***revelation principle***
- ***How does truthfulness affect computation?*** How does this affect the usability of a mechanism?
- General design idea for mechanisms: pay for ***maximising other agents' utility***

What is a mechanism?

A. Mechanism is a ***game of incomplete information*** that ***implements a social choice*** function (the real definition)

B. Mechanism is an ***optimisation algorithm*** with additional ***game-theoretic guarantees*** (a computer scientist's naive definition)

Ex: Deferred acceptance

- ***Deferred acceptance algorithm*** (or *Gale-Shapley algorithm*) is one of the most important mechanisms in economics (2012 Nobel Memorial Prize in Economic Sciences to Lloyd Shapley and Alvin E. Roth)
- Computes a ***stable matching*** in a bipartite network where agents have a private linear (preference) order over their neighbours
  - ***stable*** = *no pair of neighbors prefer each other over their assigned match*
- The proposing side is incentivised to propose in the order of their preferences

Deferred acceptance

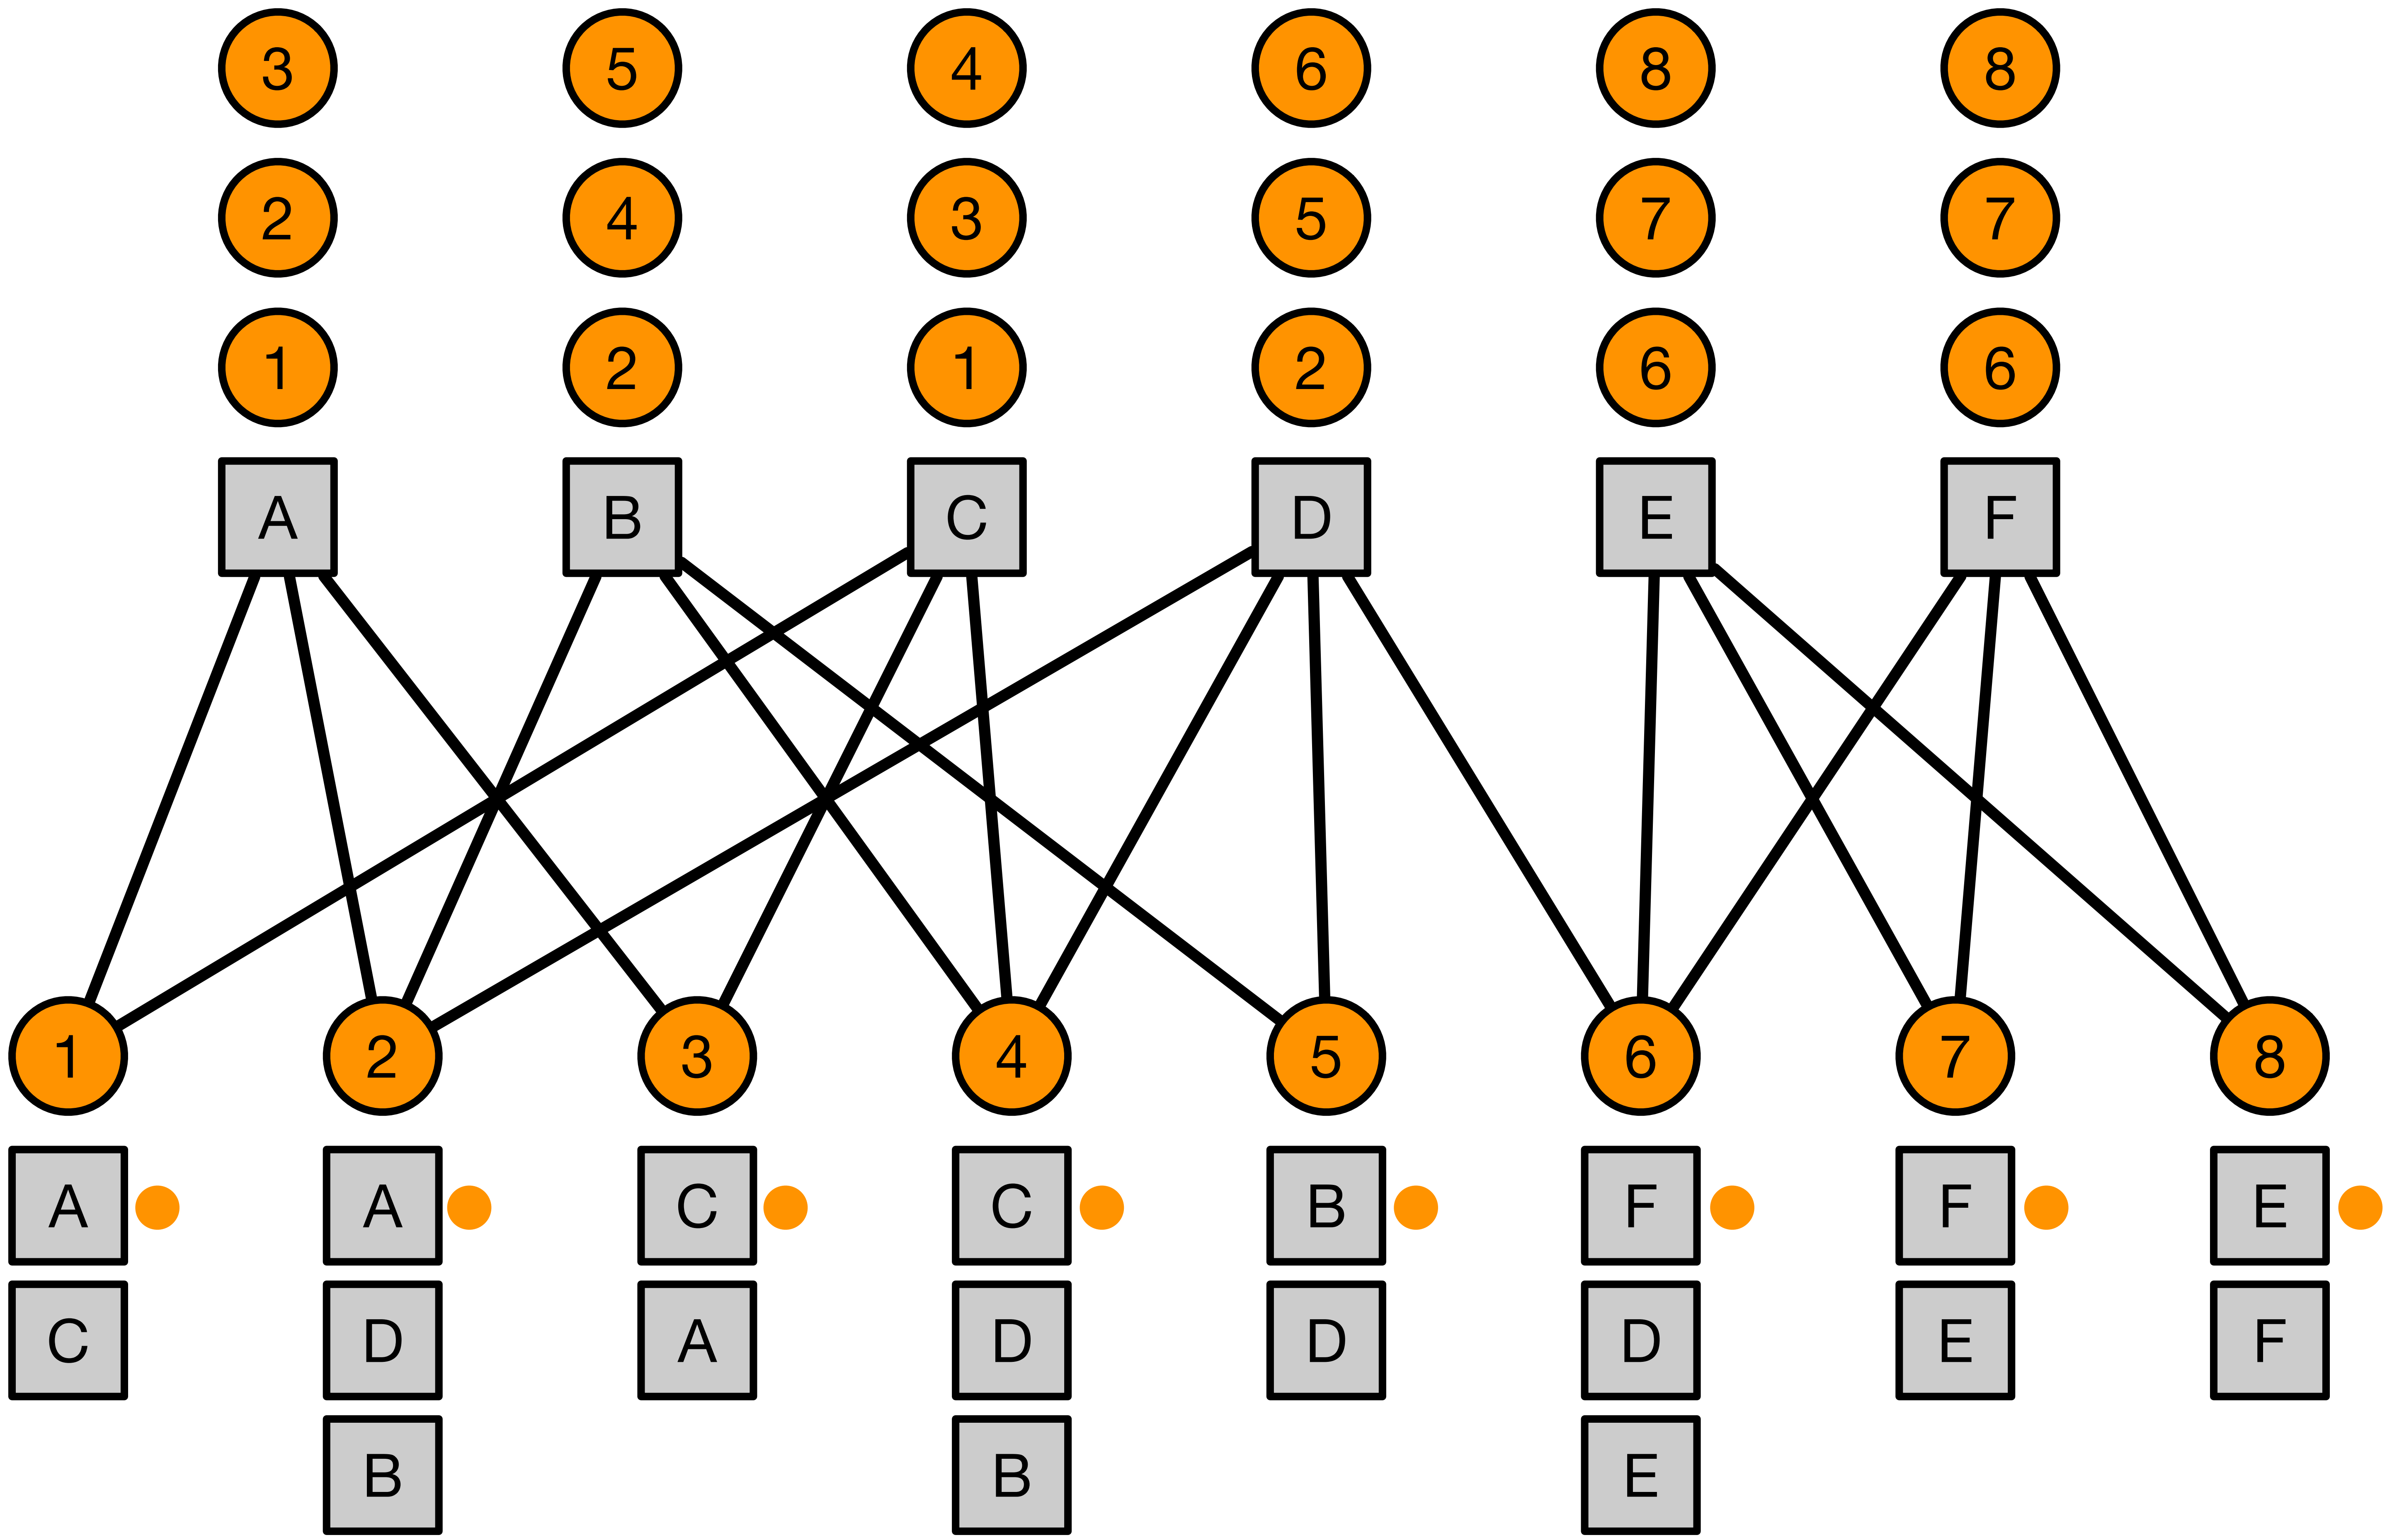

Deferred acceptance

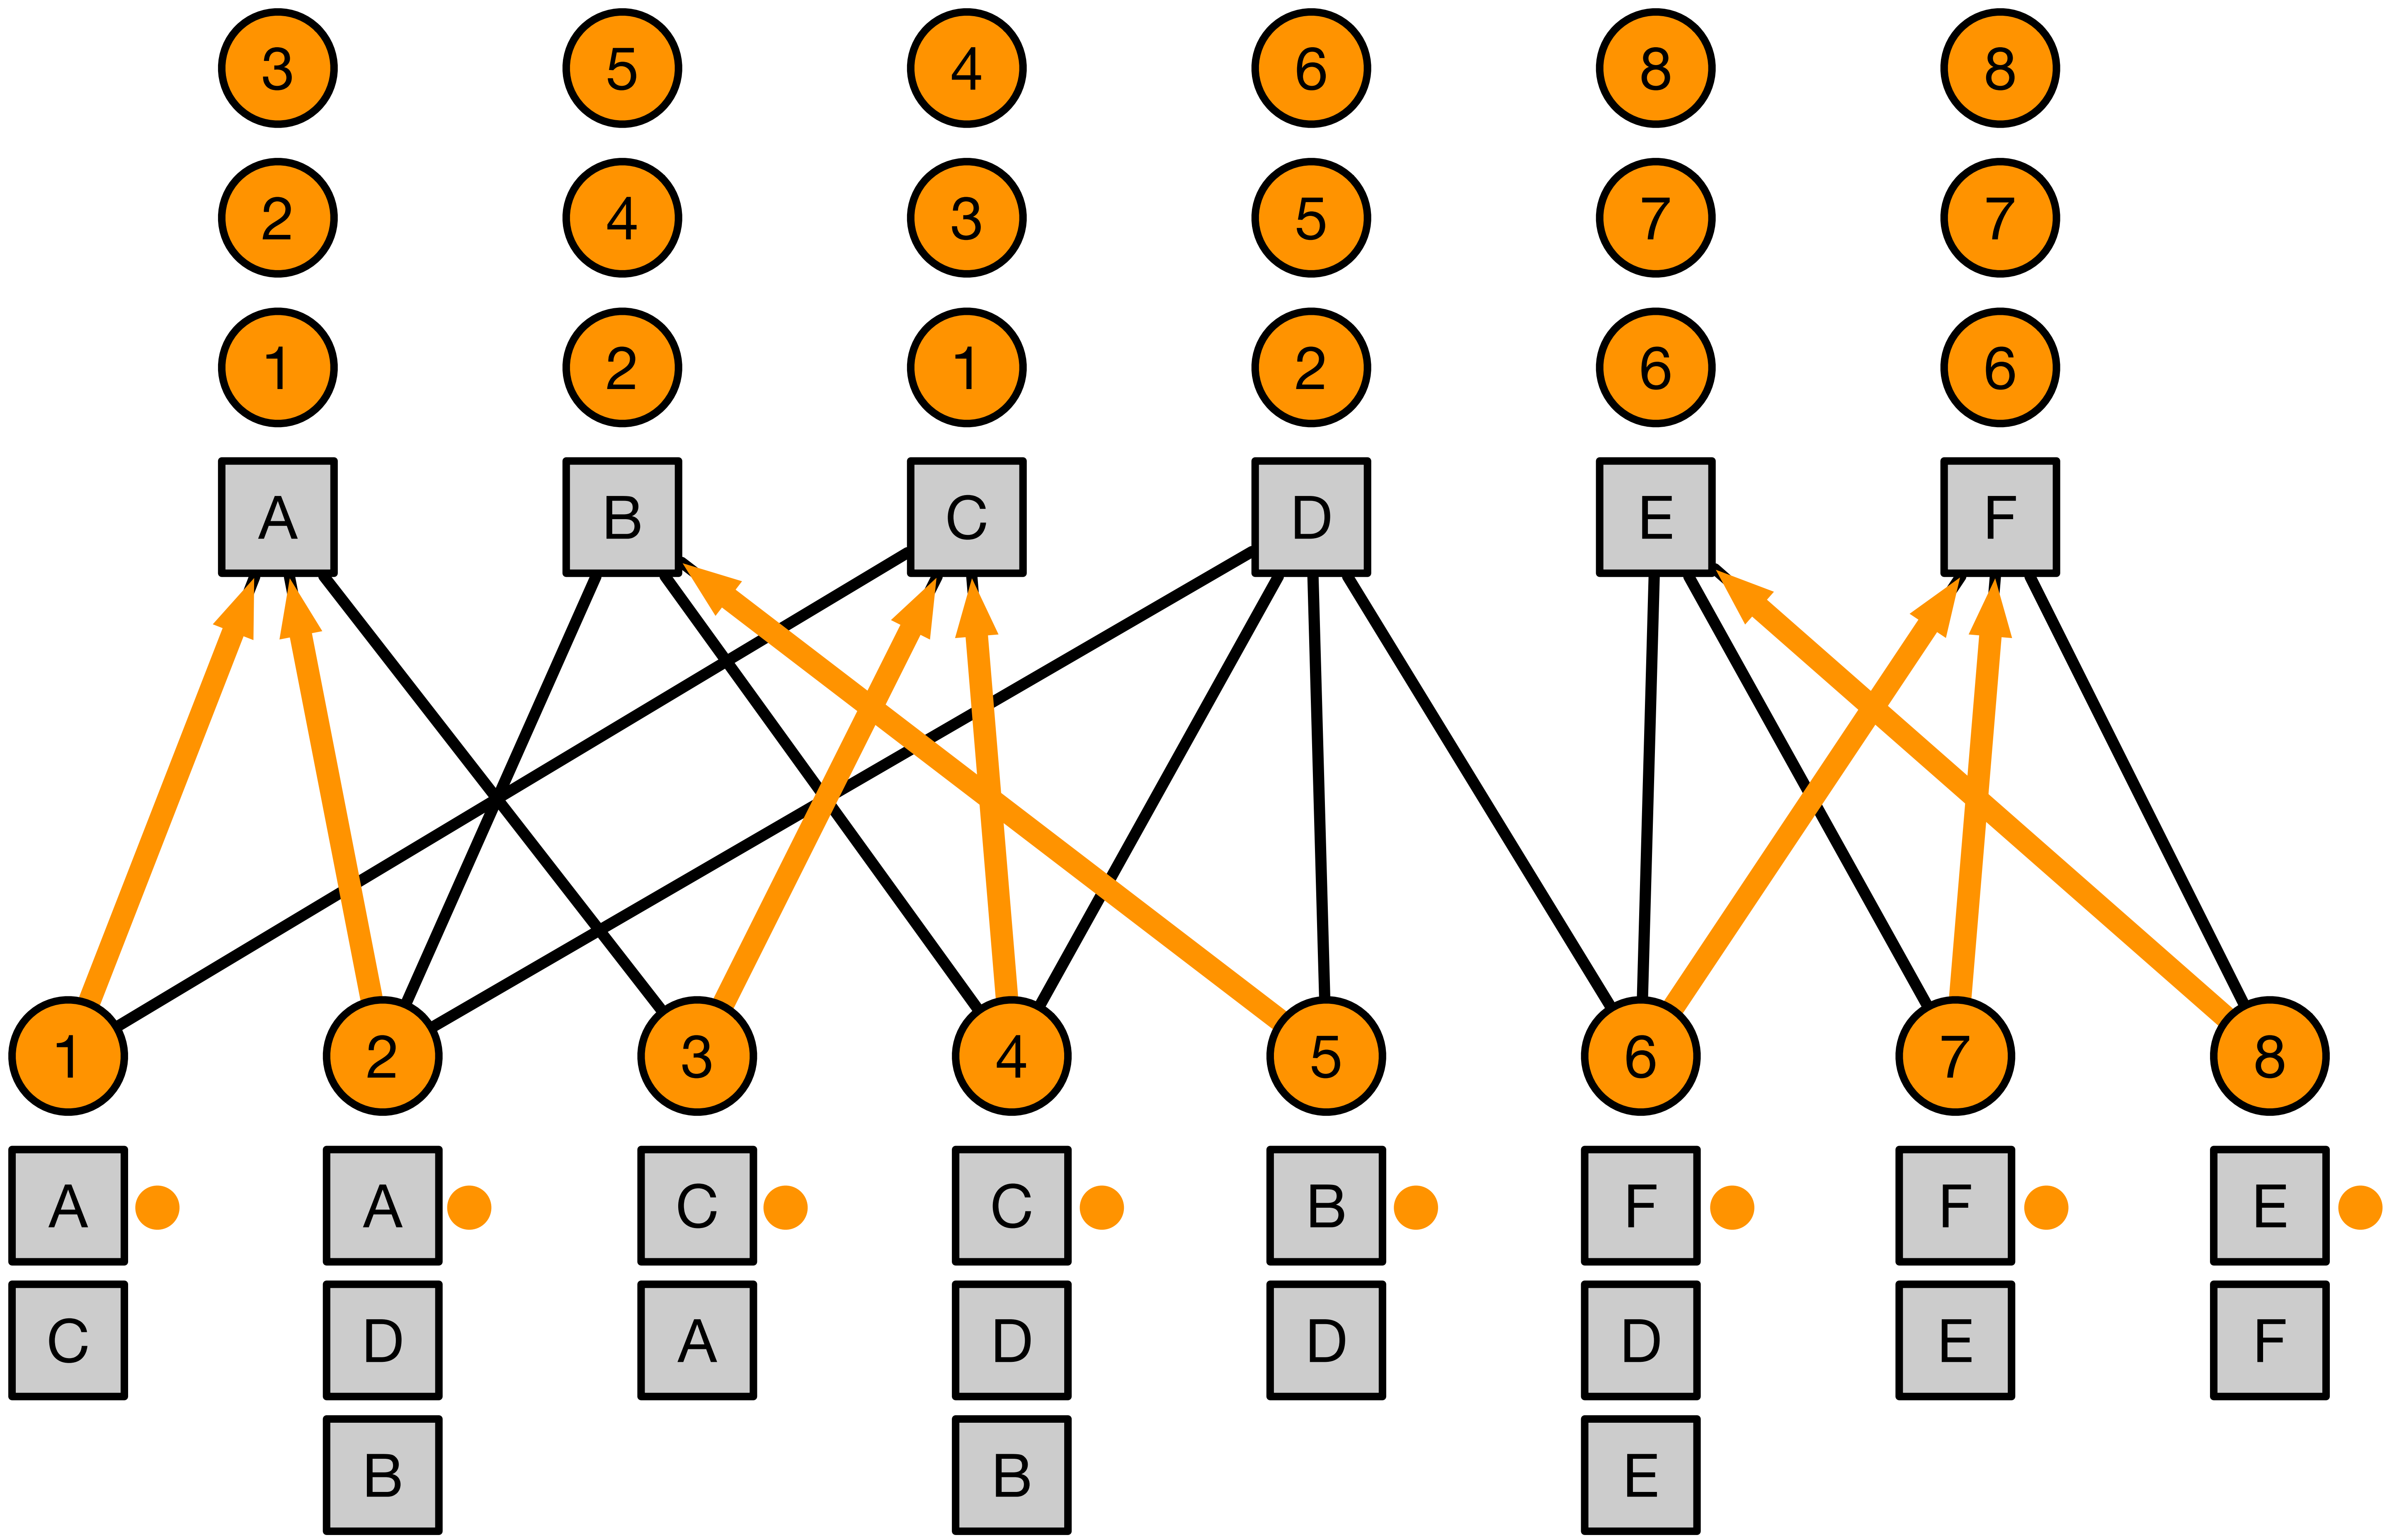

Deferred acceptance

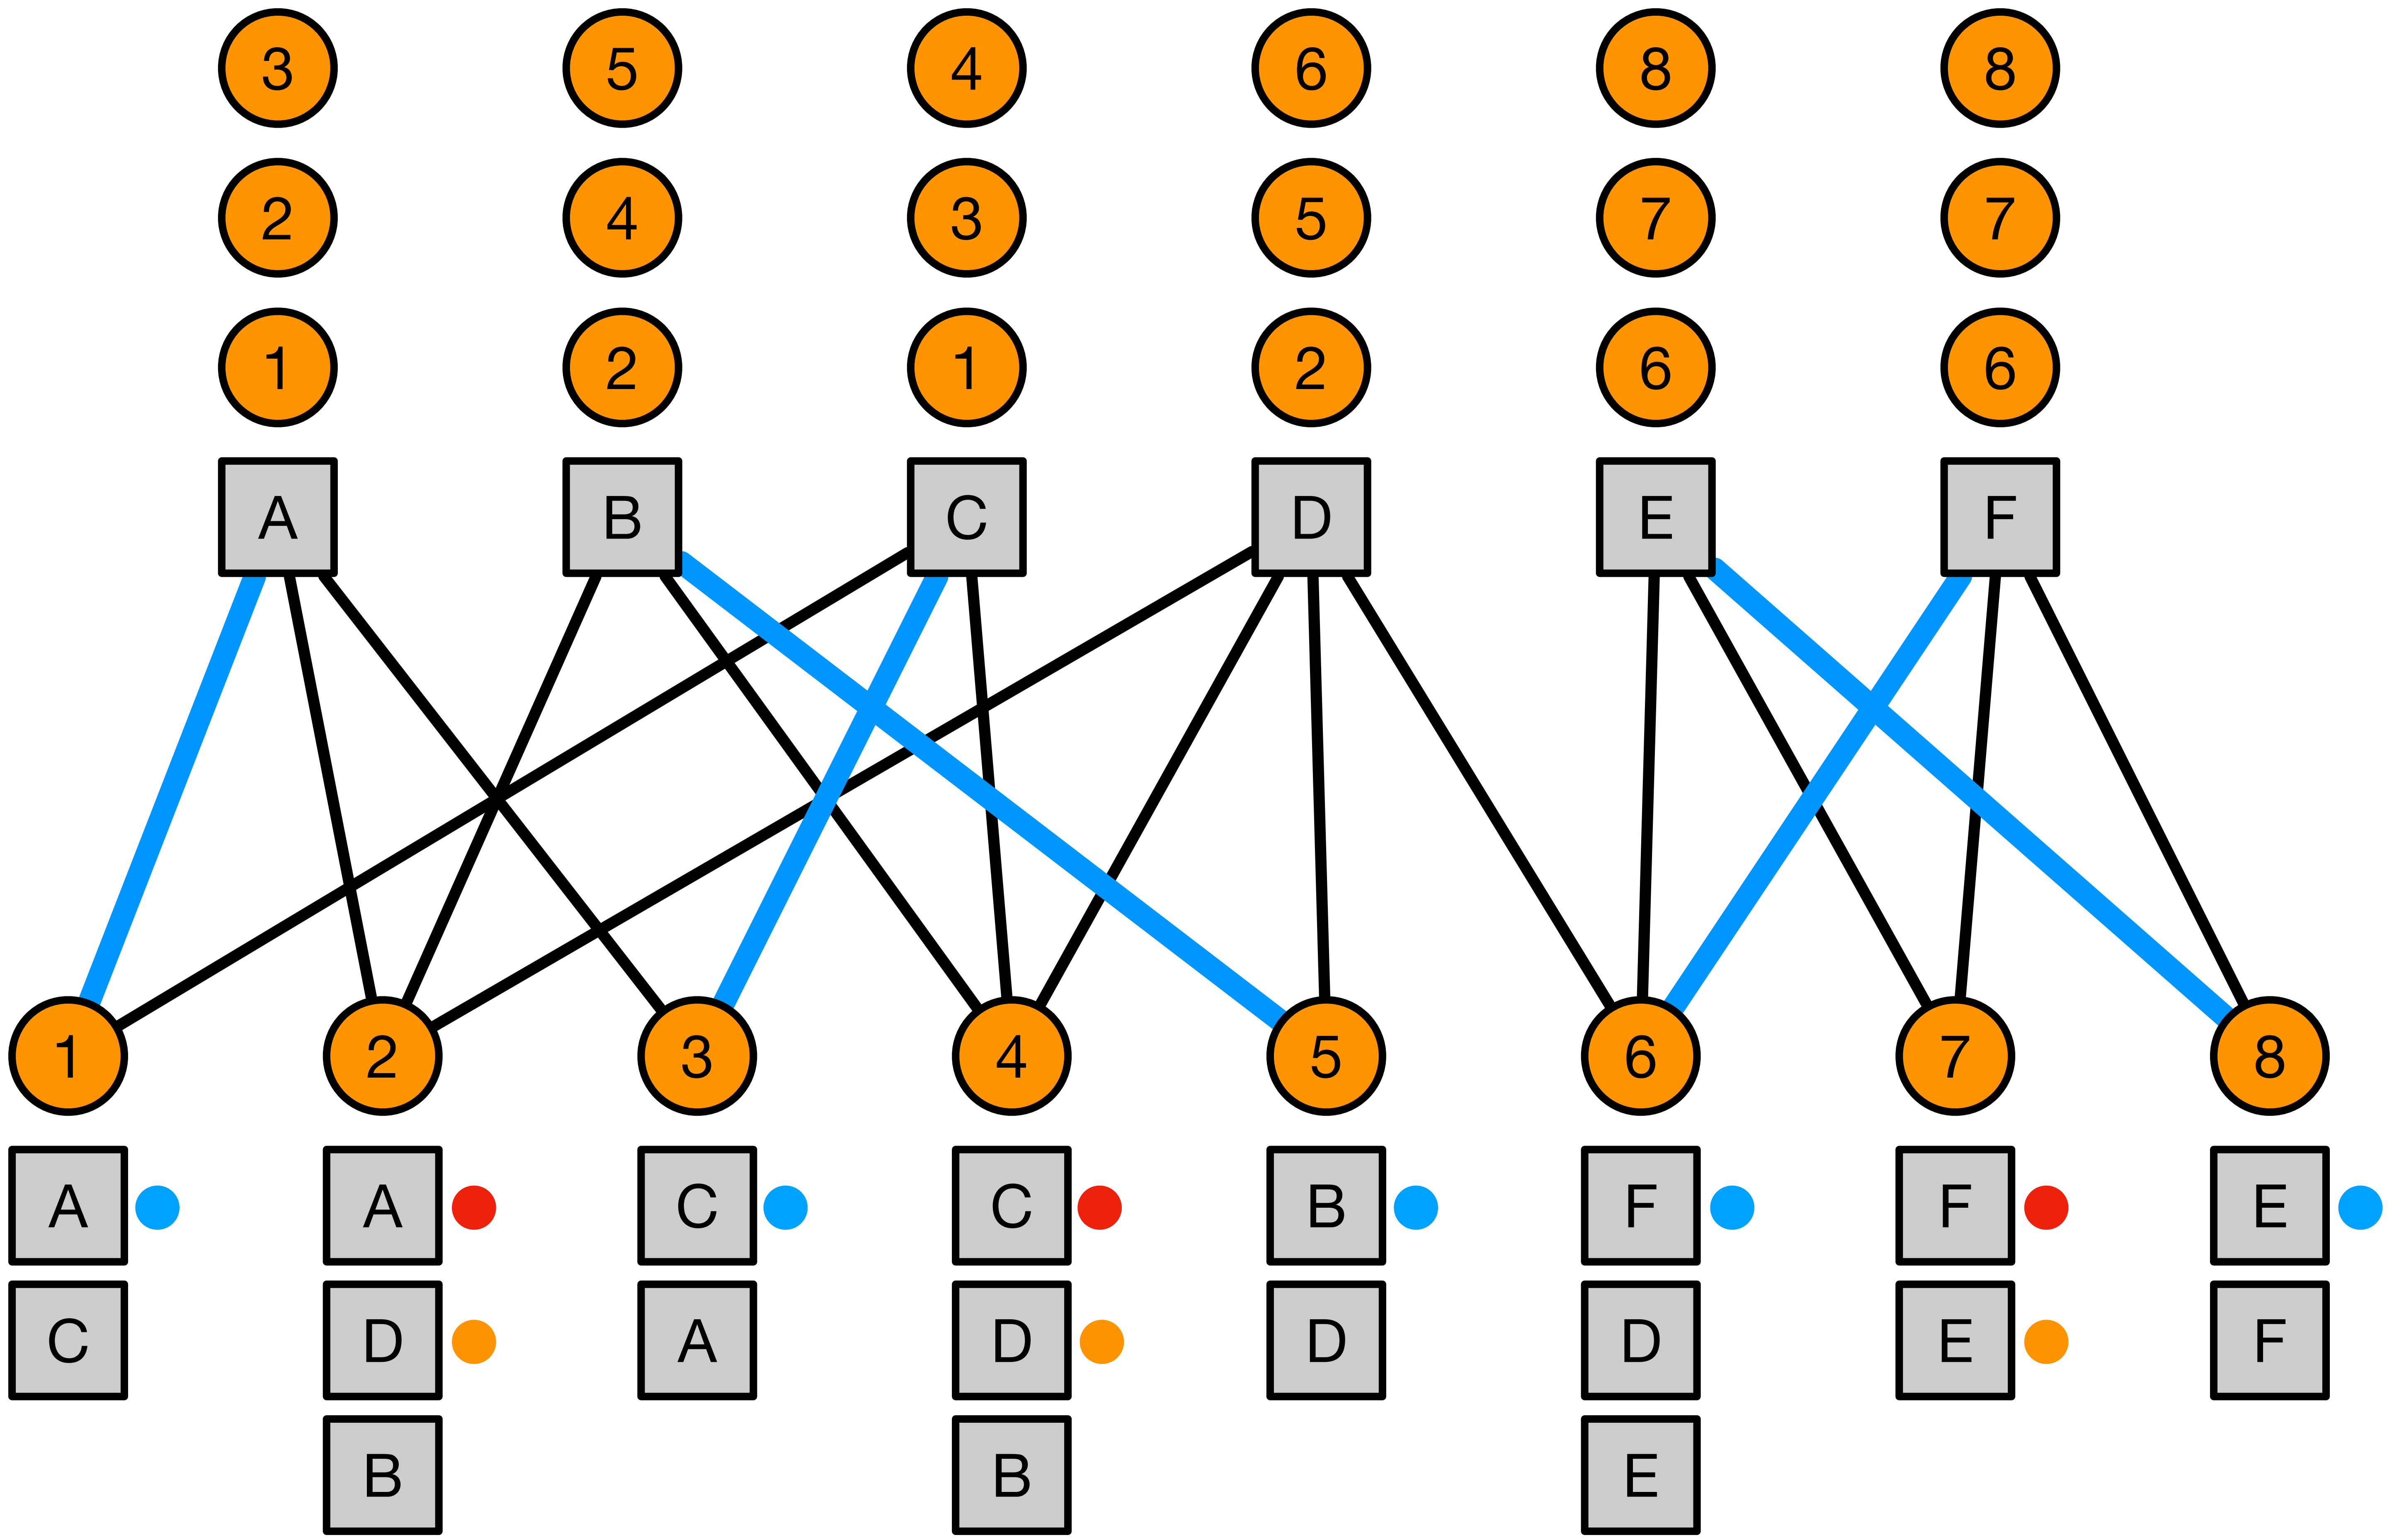

Deferred acceptance

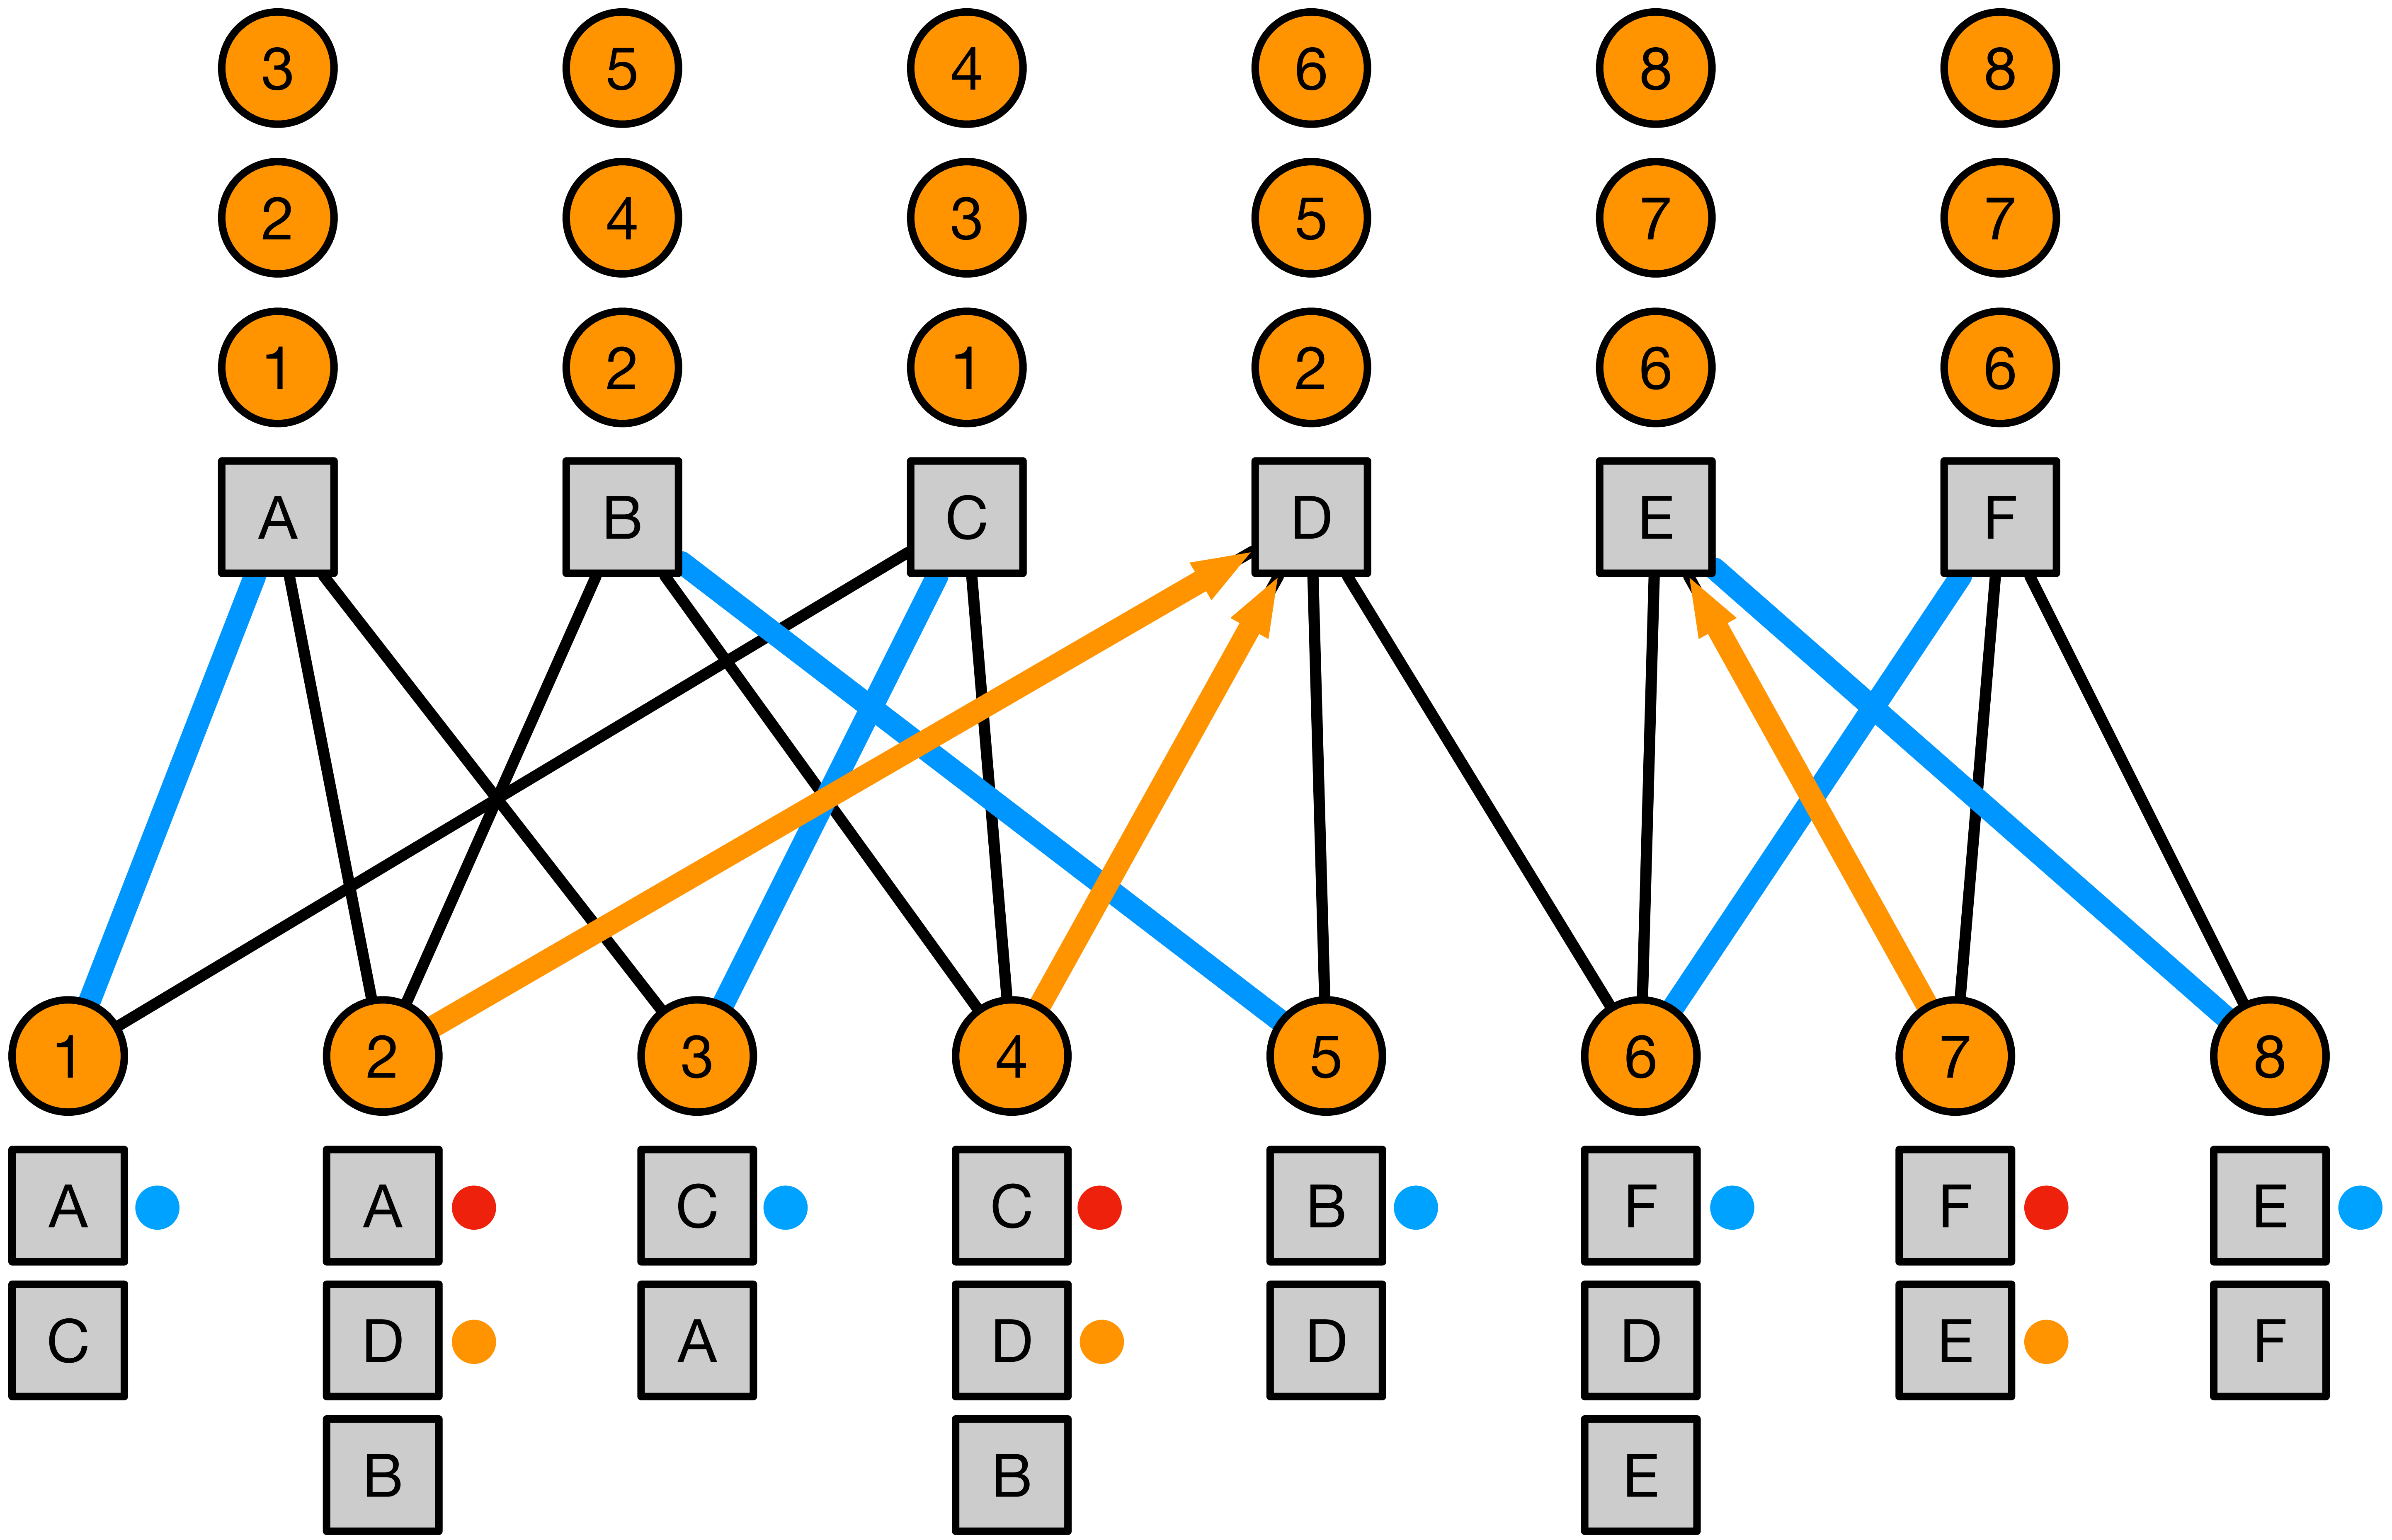

Deferred acceptance

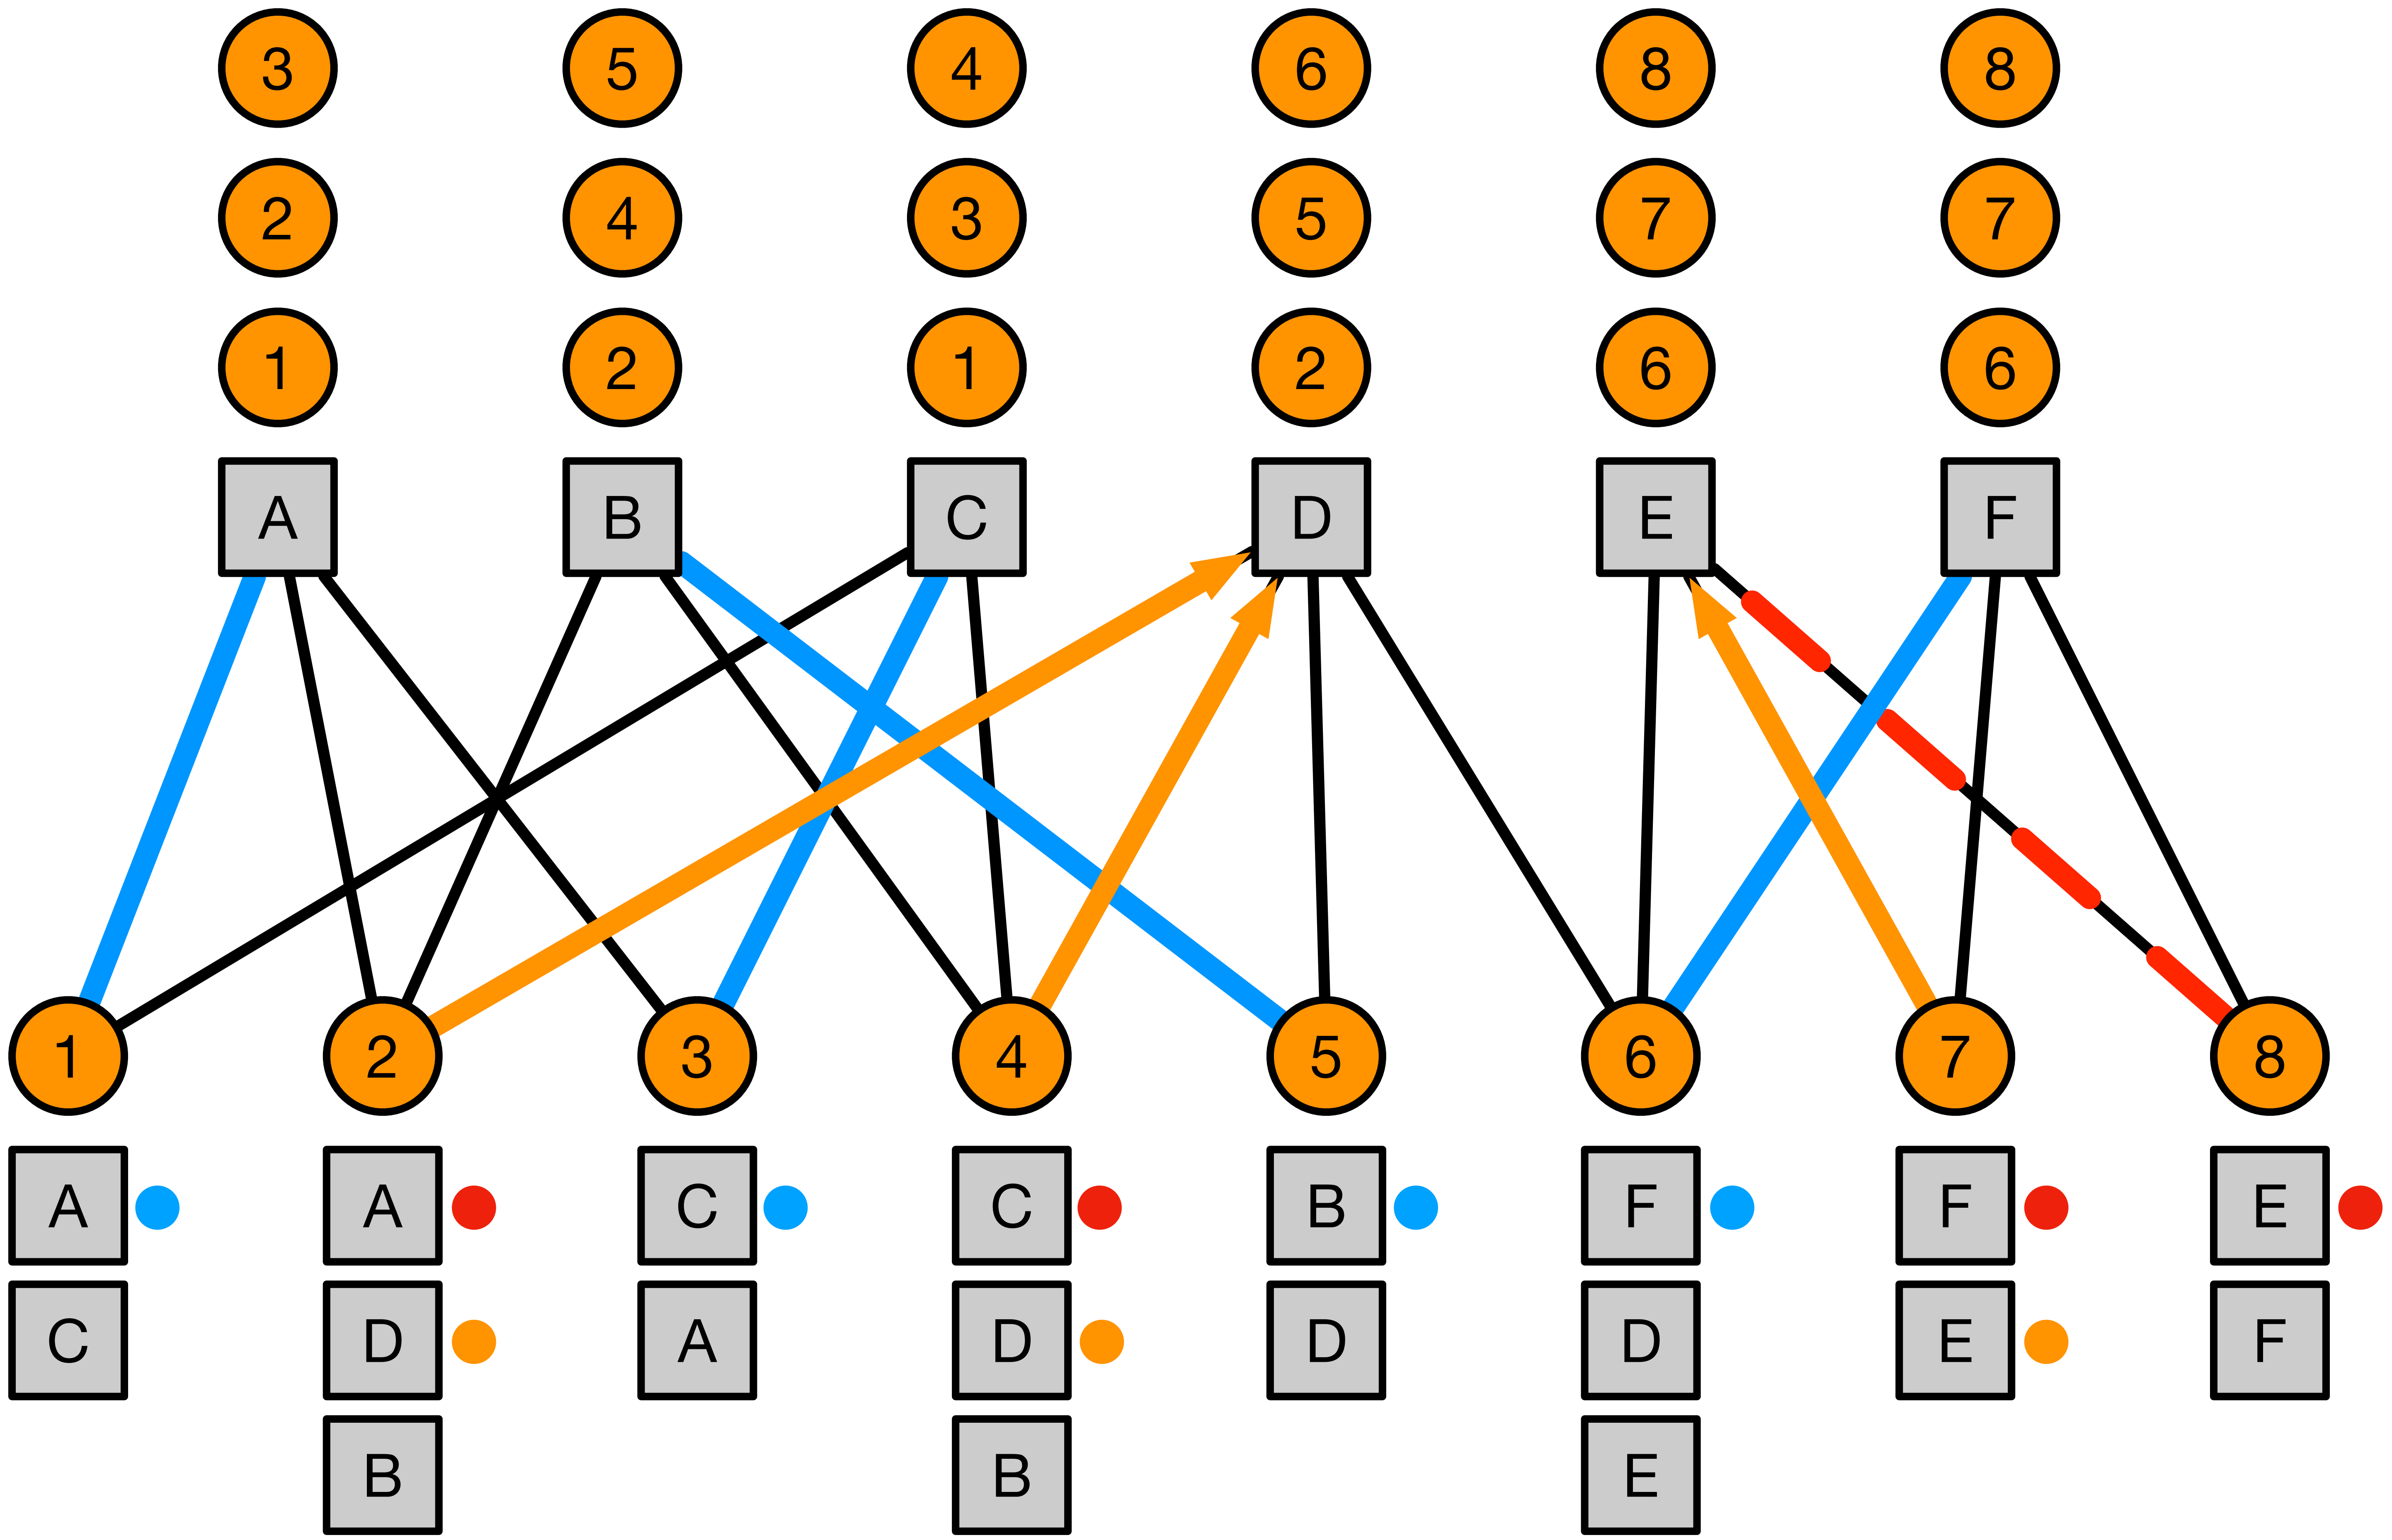

Deferred acceptance

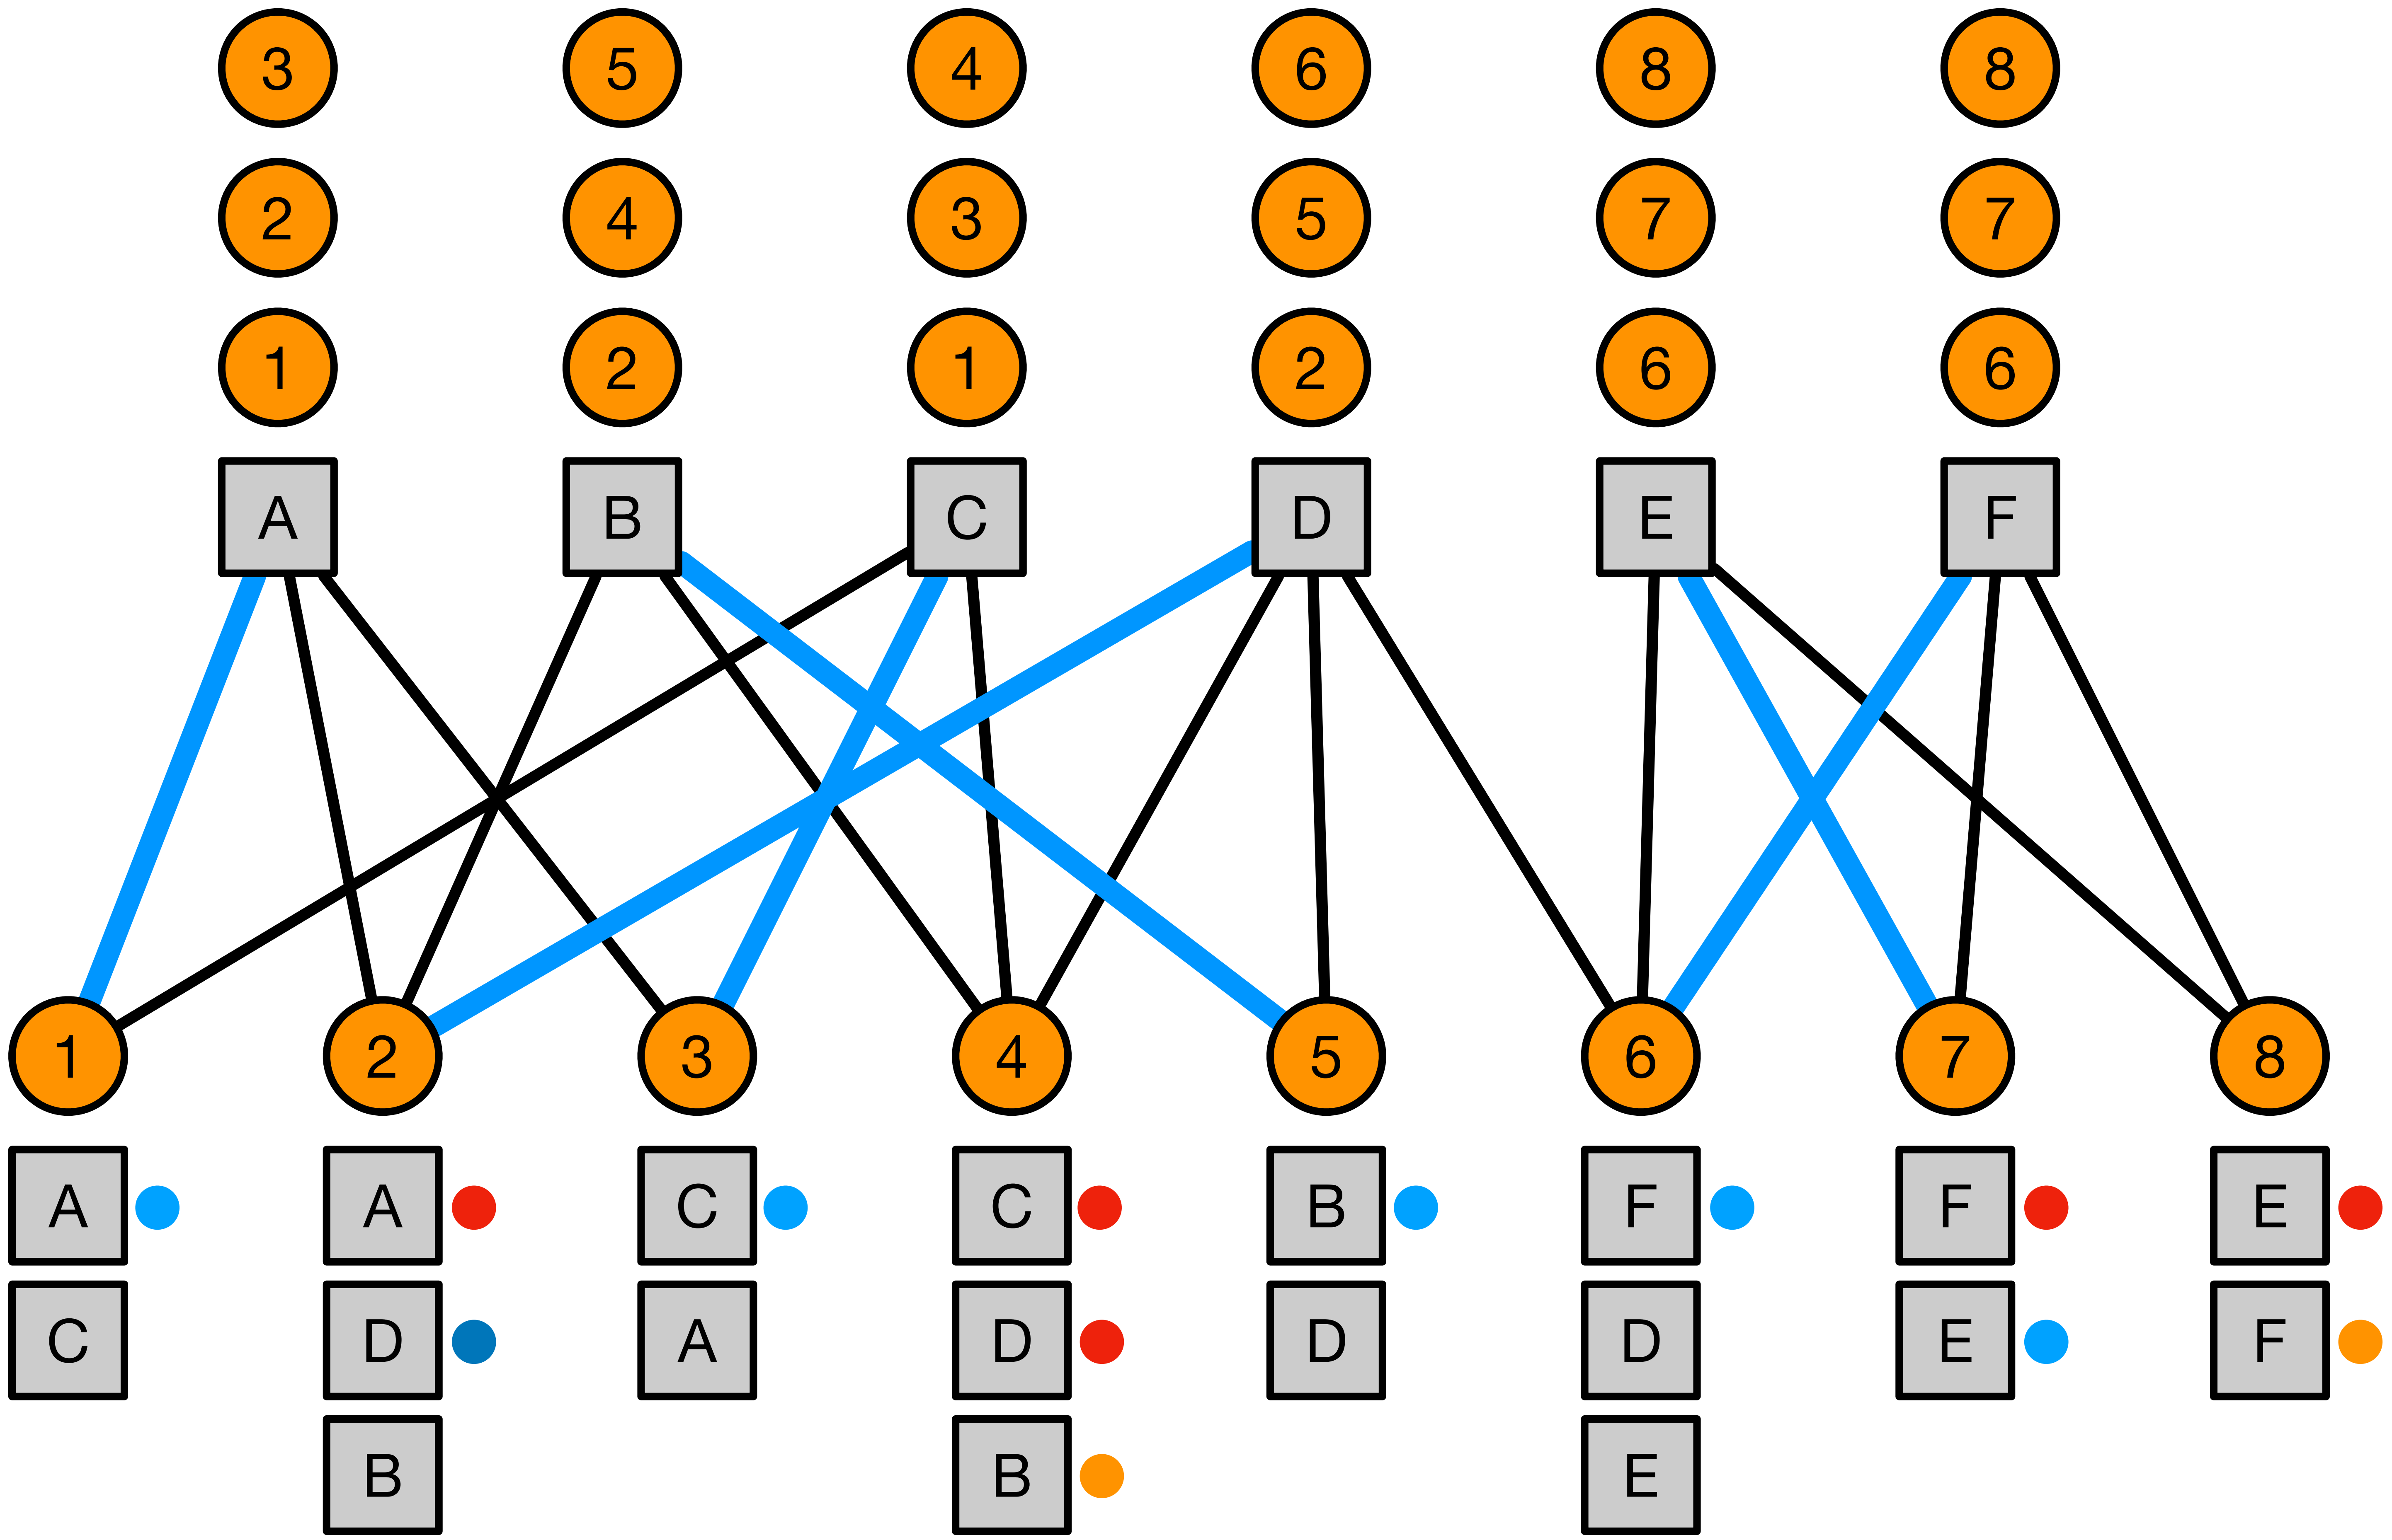

Deferred acceptance

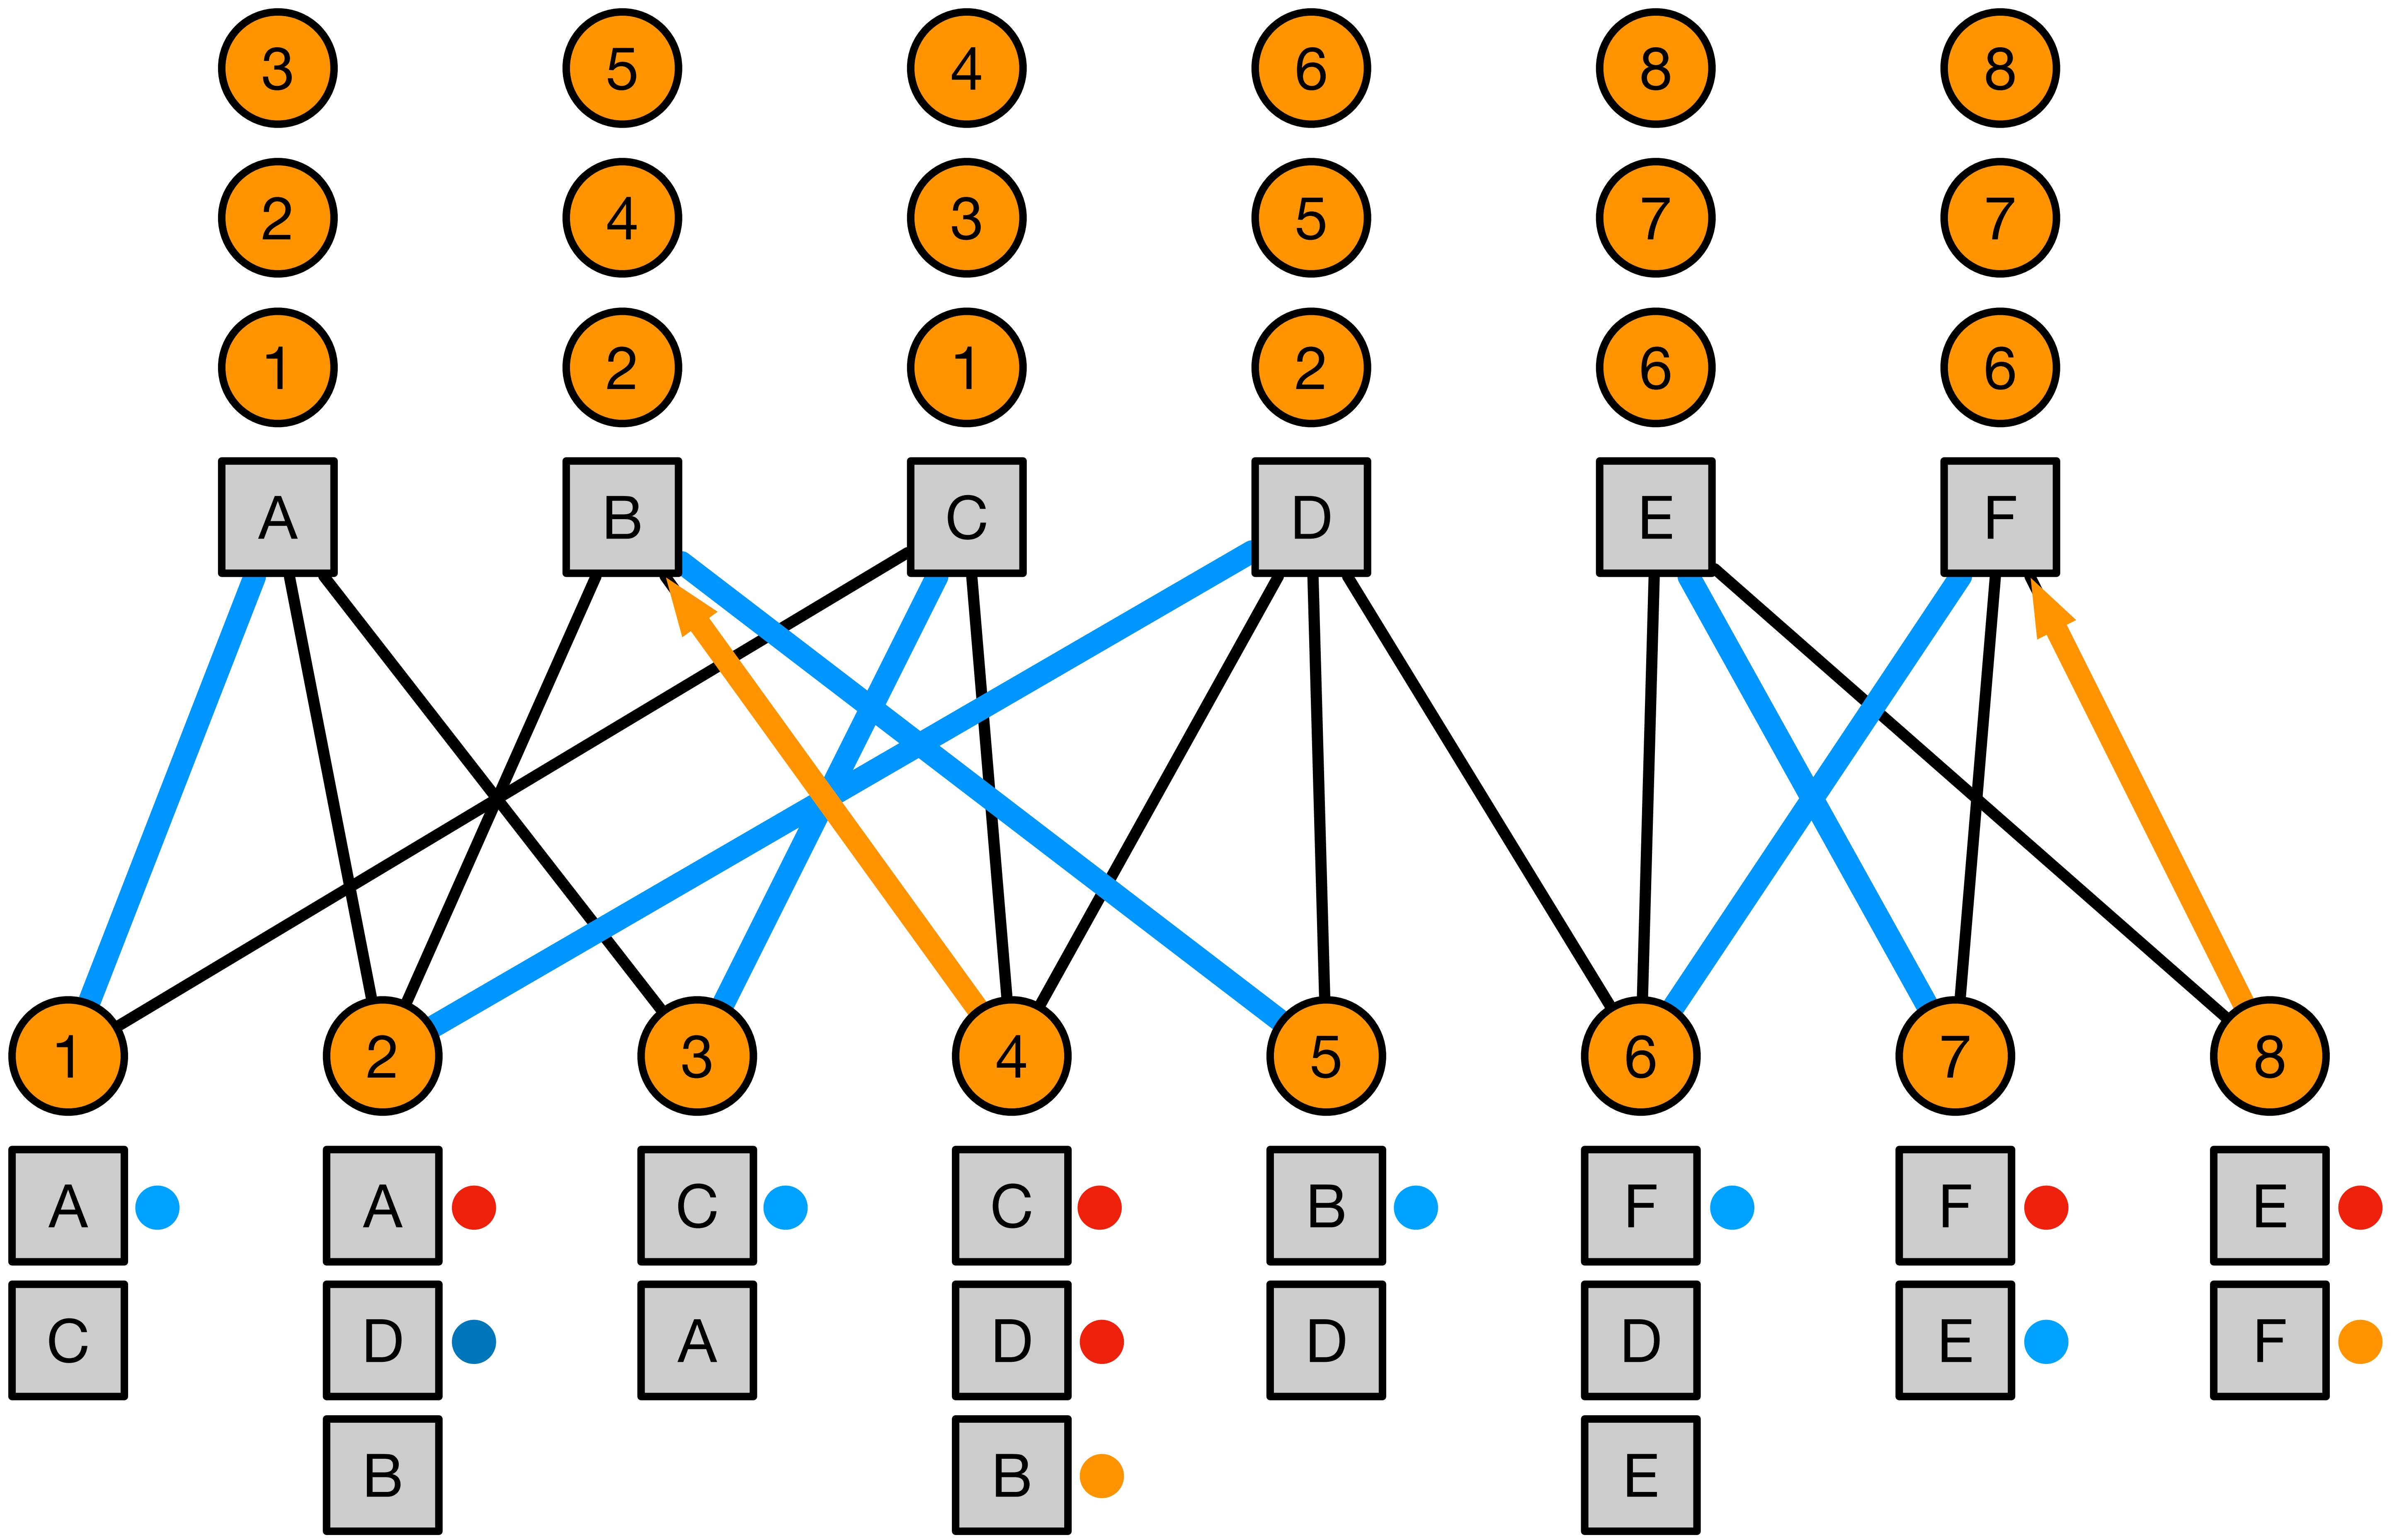

Deferred acceptance

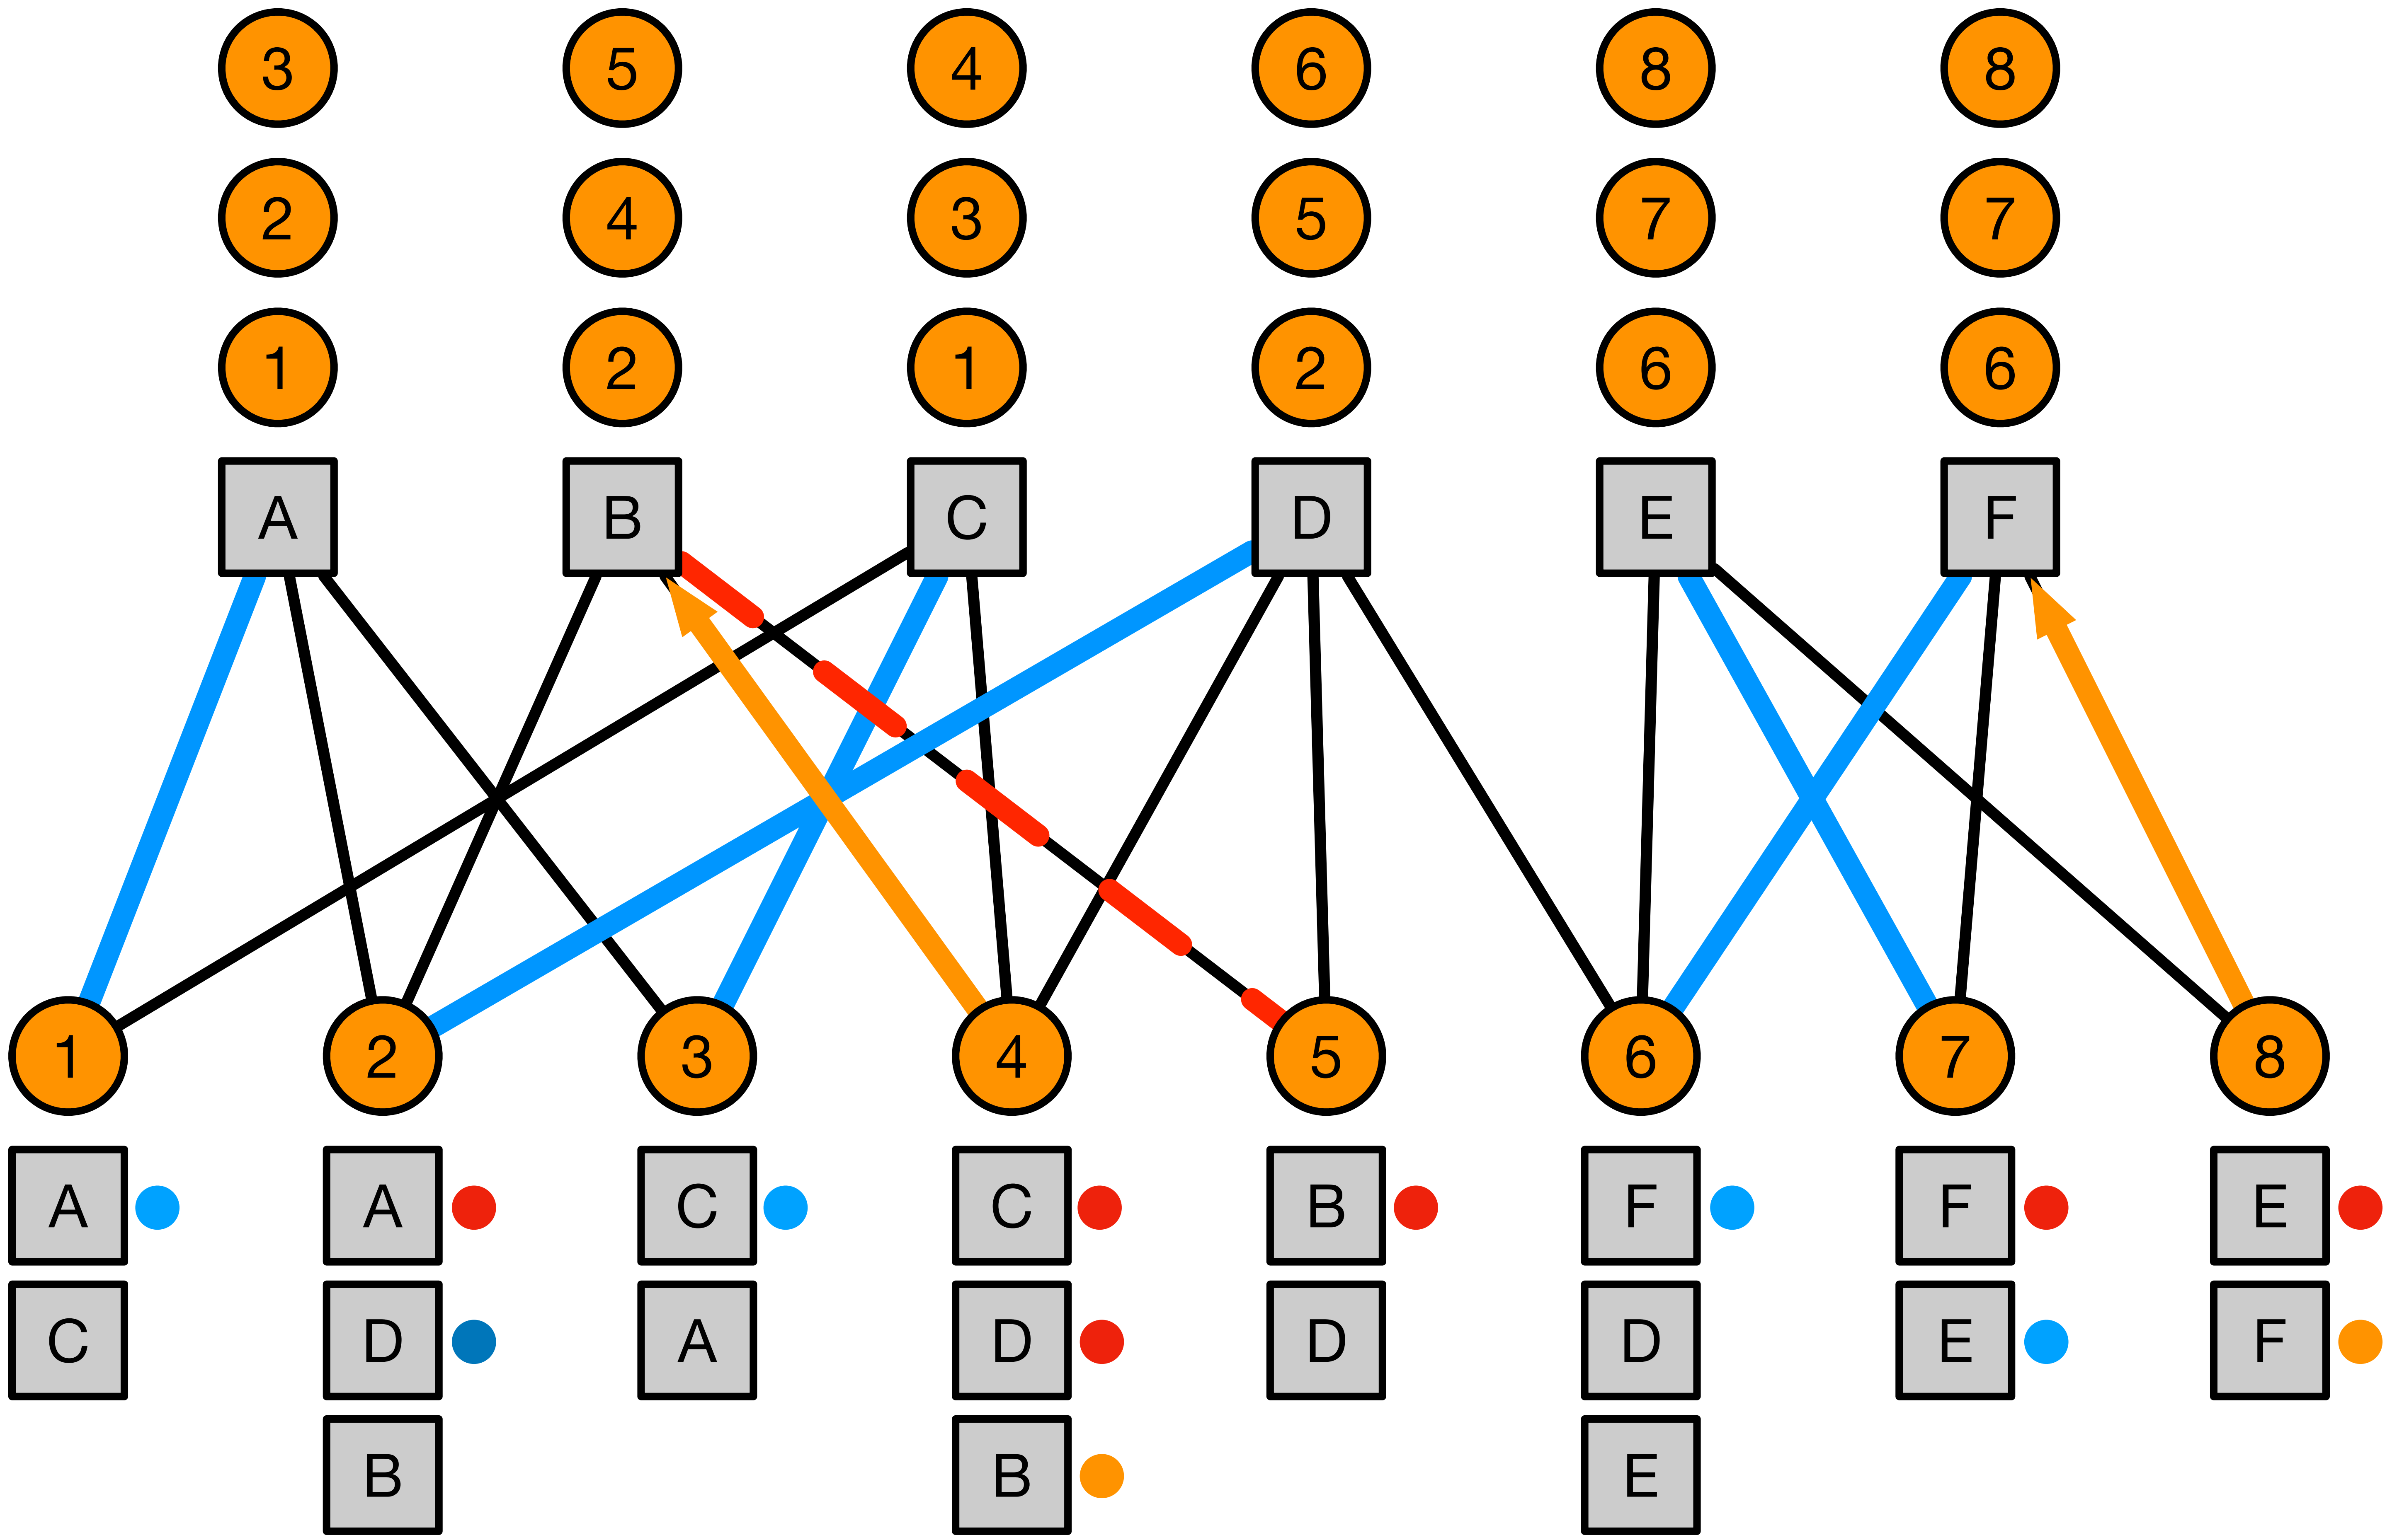

Deferred acceptance

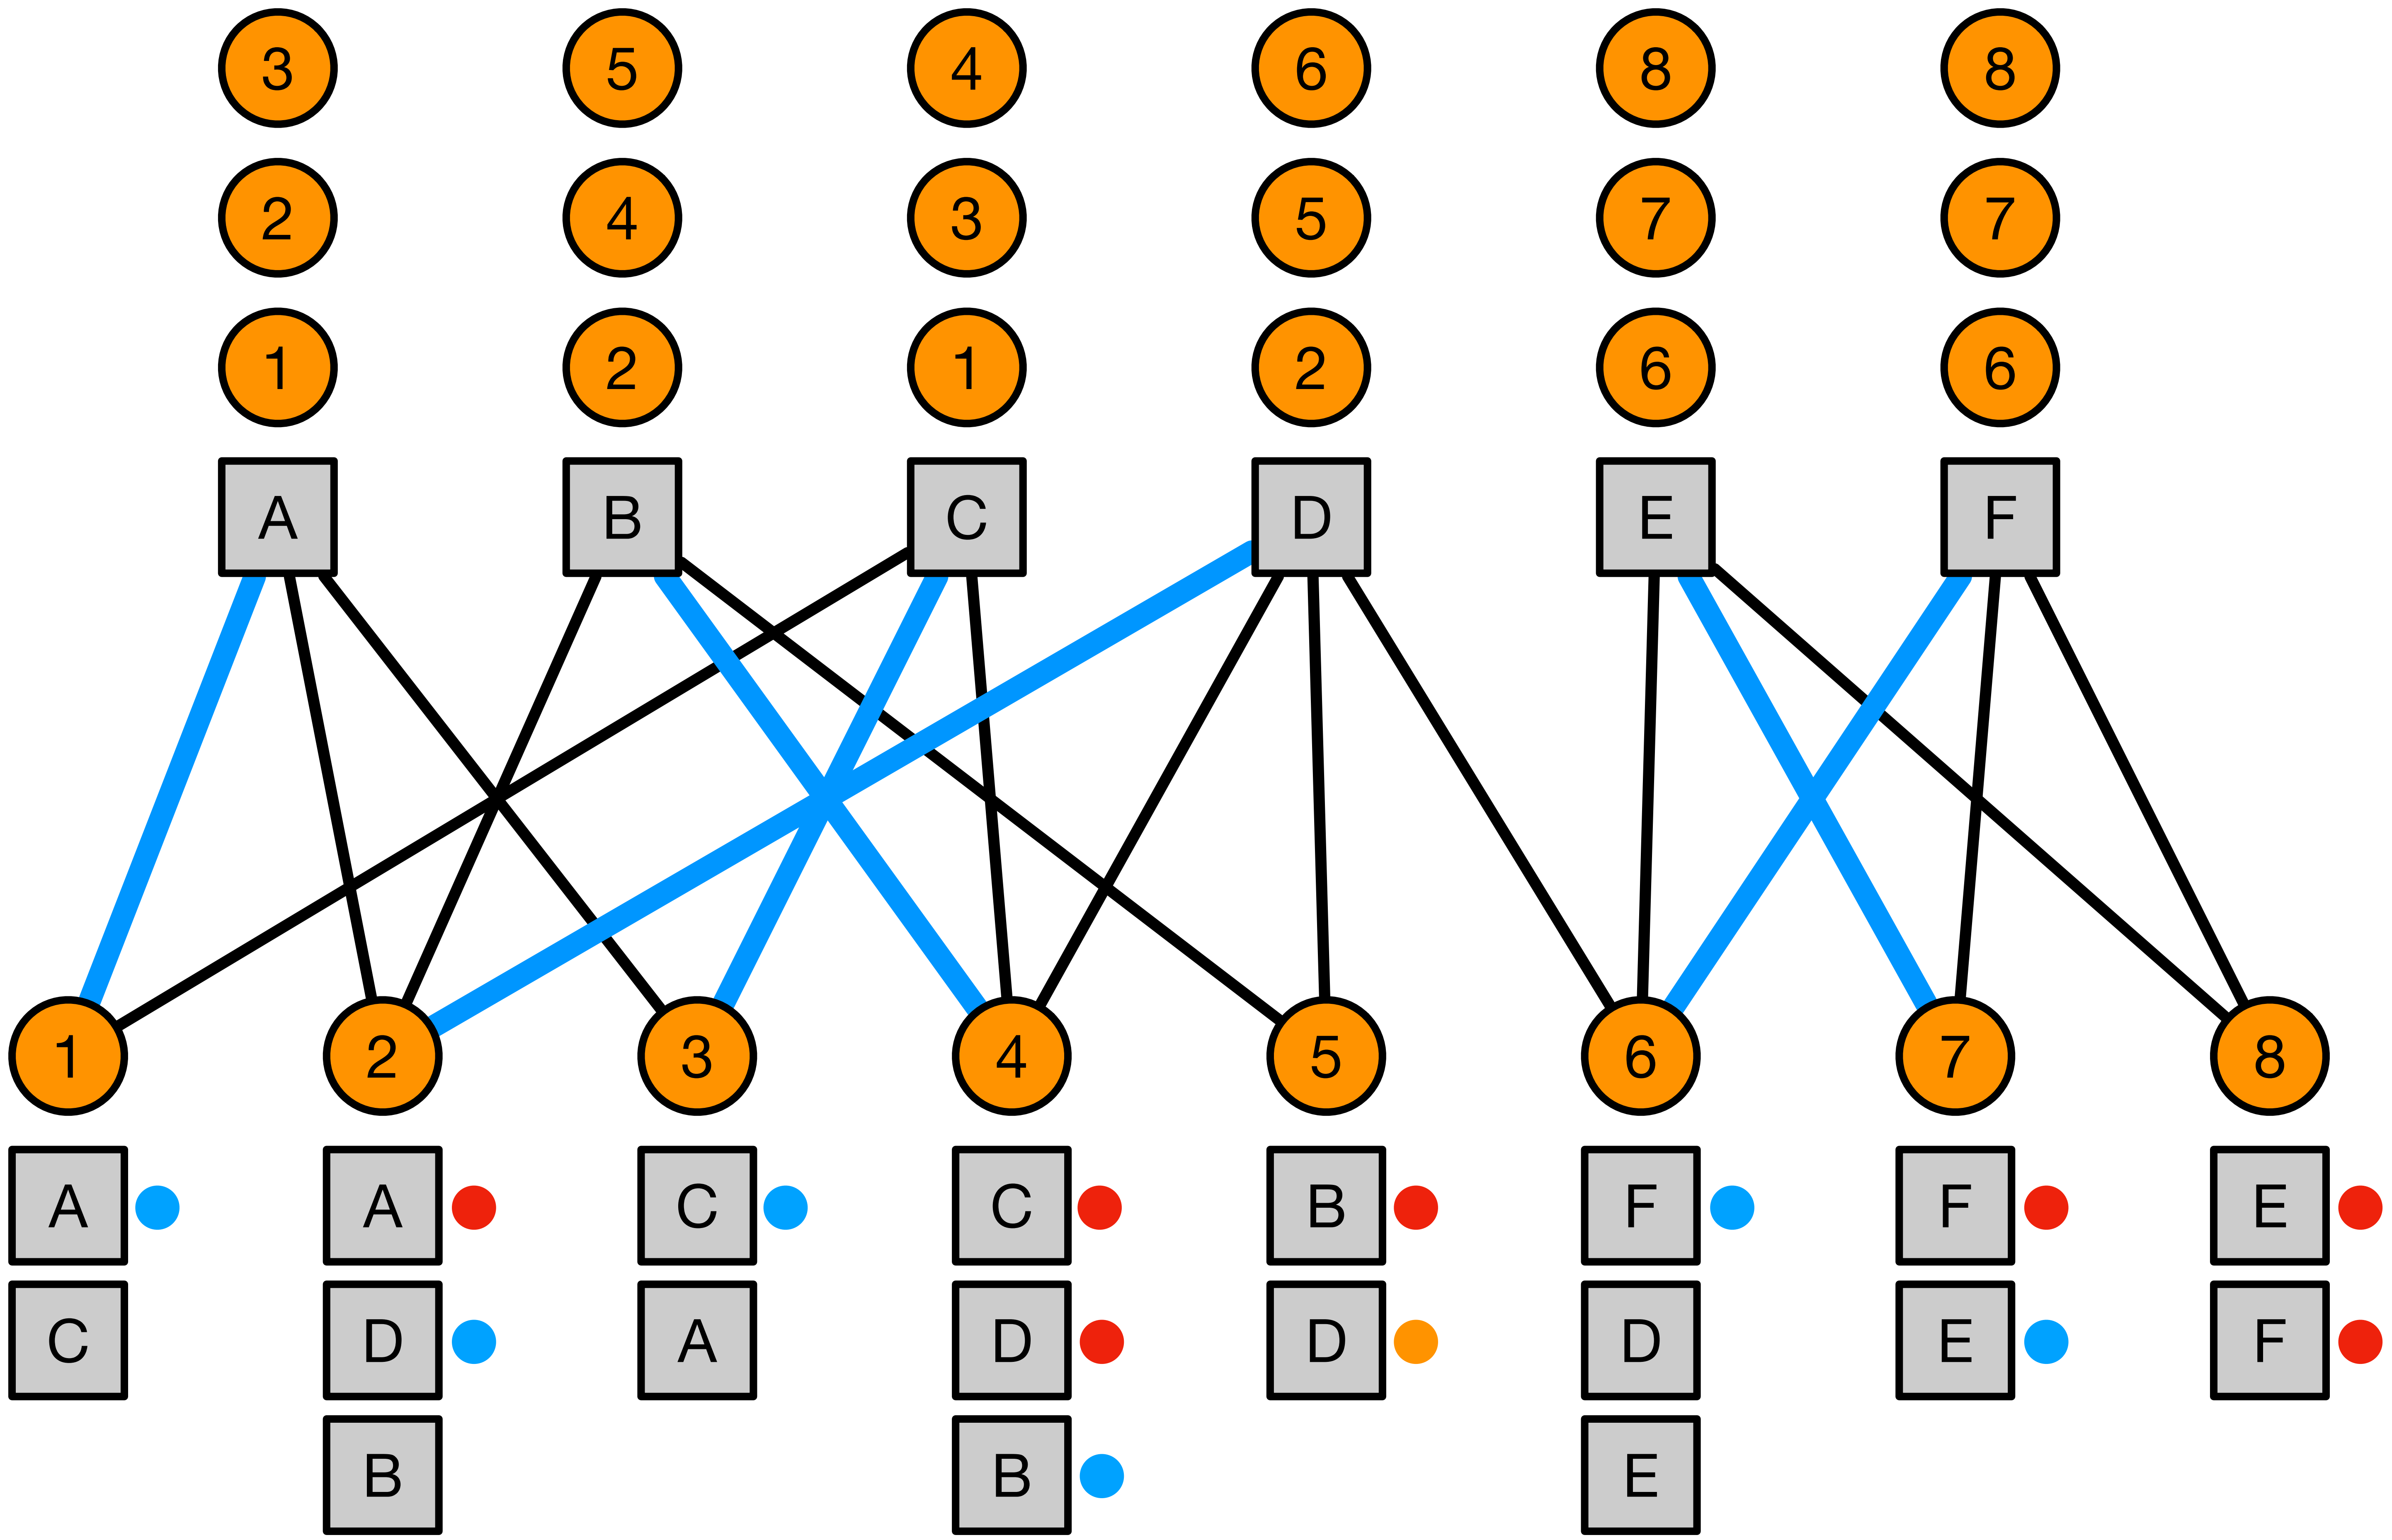

Deferred acceptance

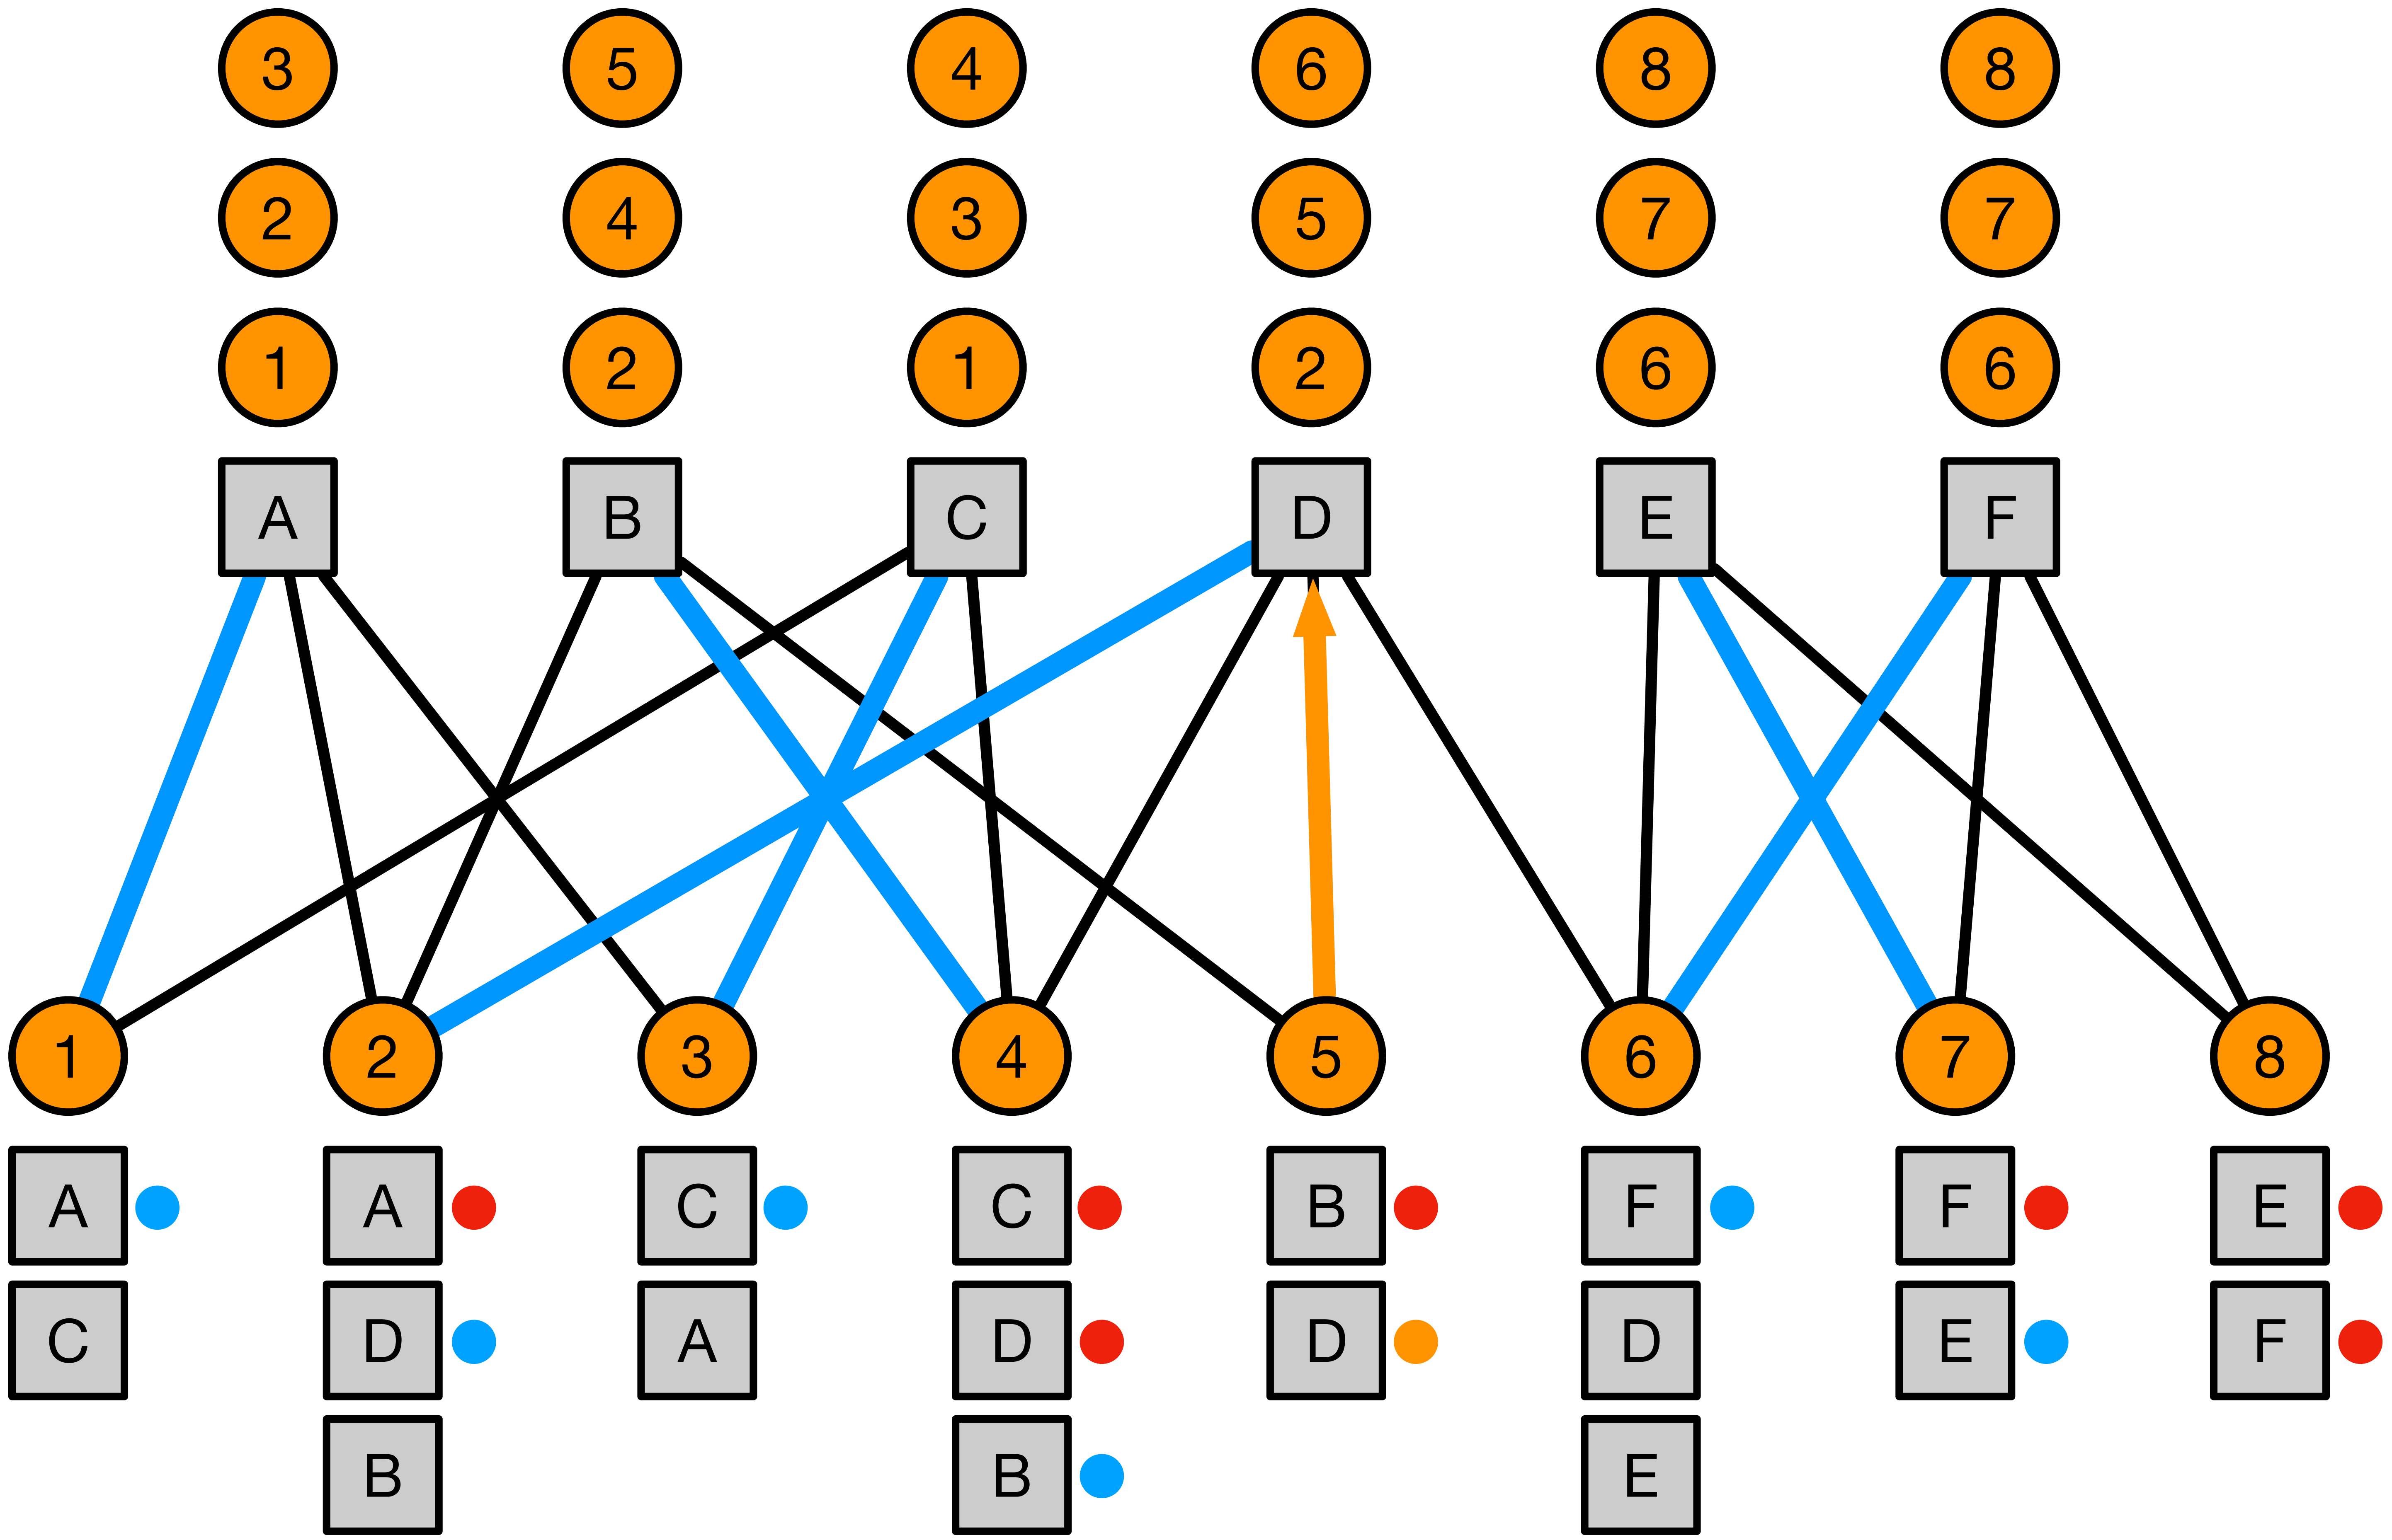

Deferred acceptance

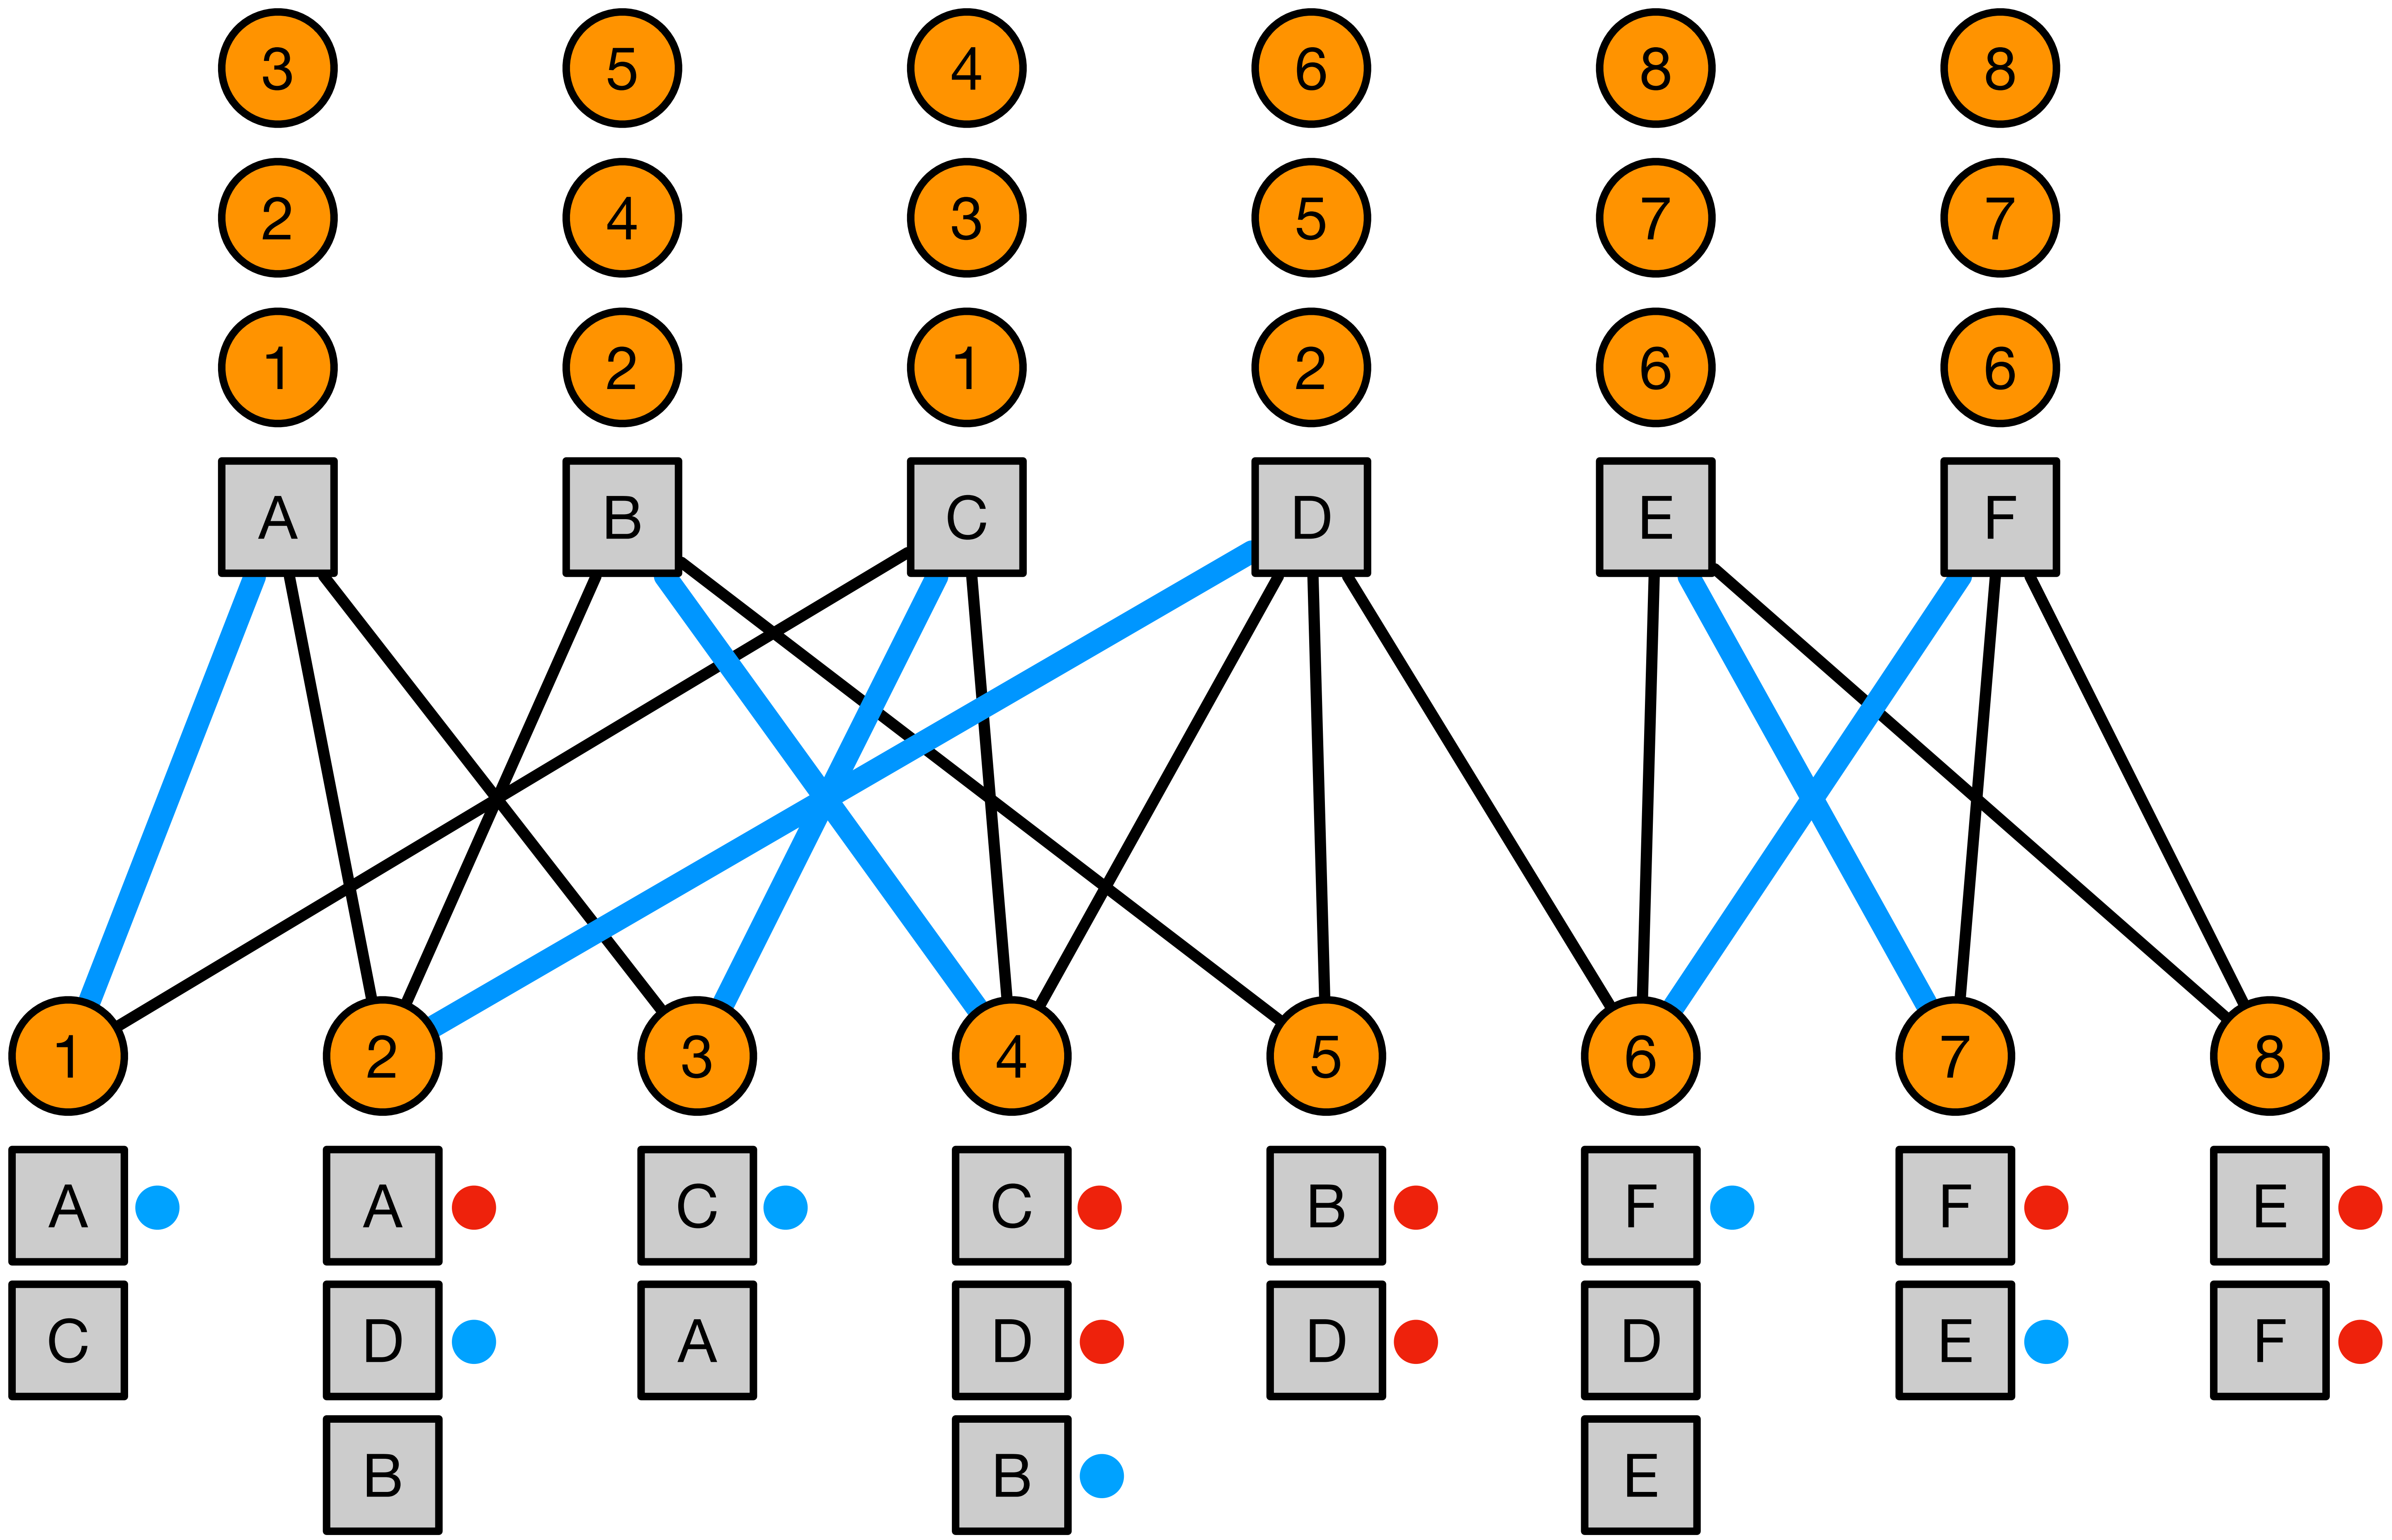

What is a mechanism?

- Each ***agent v*** has a hidden ***type t_v*** (*private information*) from a known set of ***types T*** and an ***action space A*** *(different algorithms)*

- The ***strategy*** of an agent *v* is a function ***$s_v : T \to A$***

- There is a set of ***outcomes O*** and each agent has a ***utility function $u_v : T \times O \to \mathbb{R}$***

- ***Mechanism M*** maps actions (a_1, a_2, …, a_n) to an *outcome*; a mechanism is a game of incomplete information

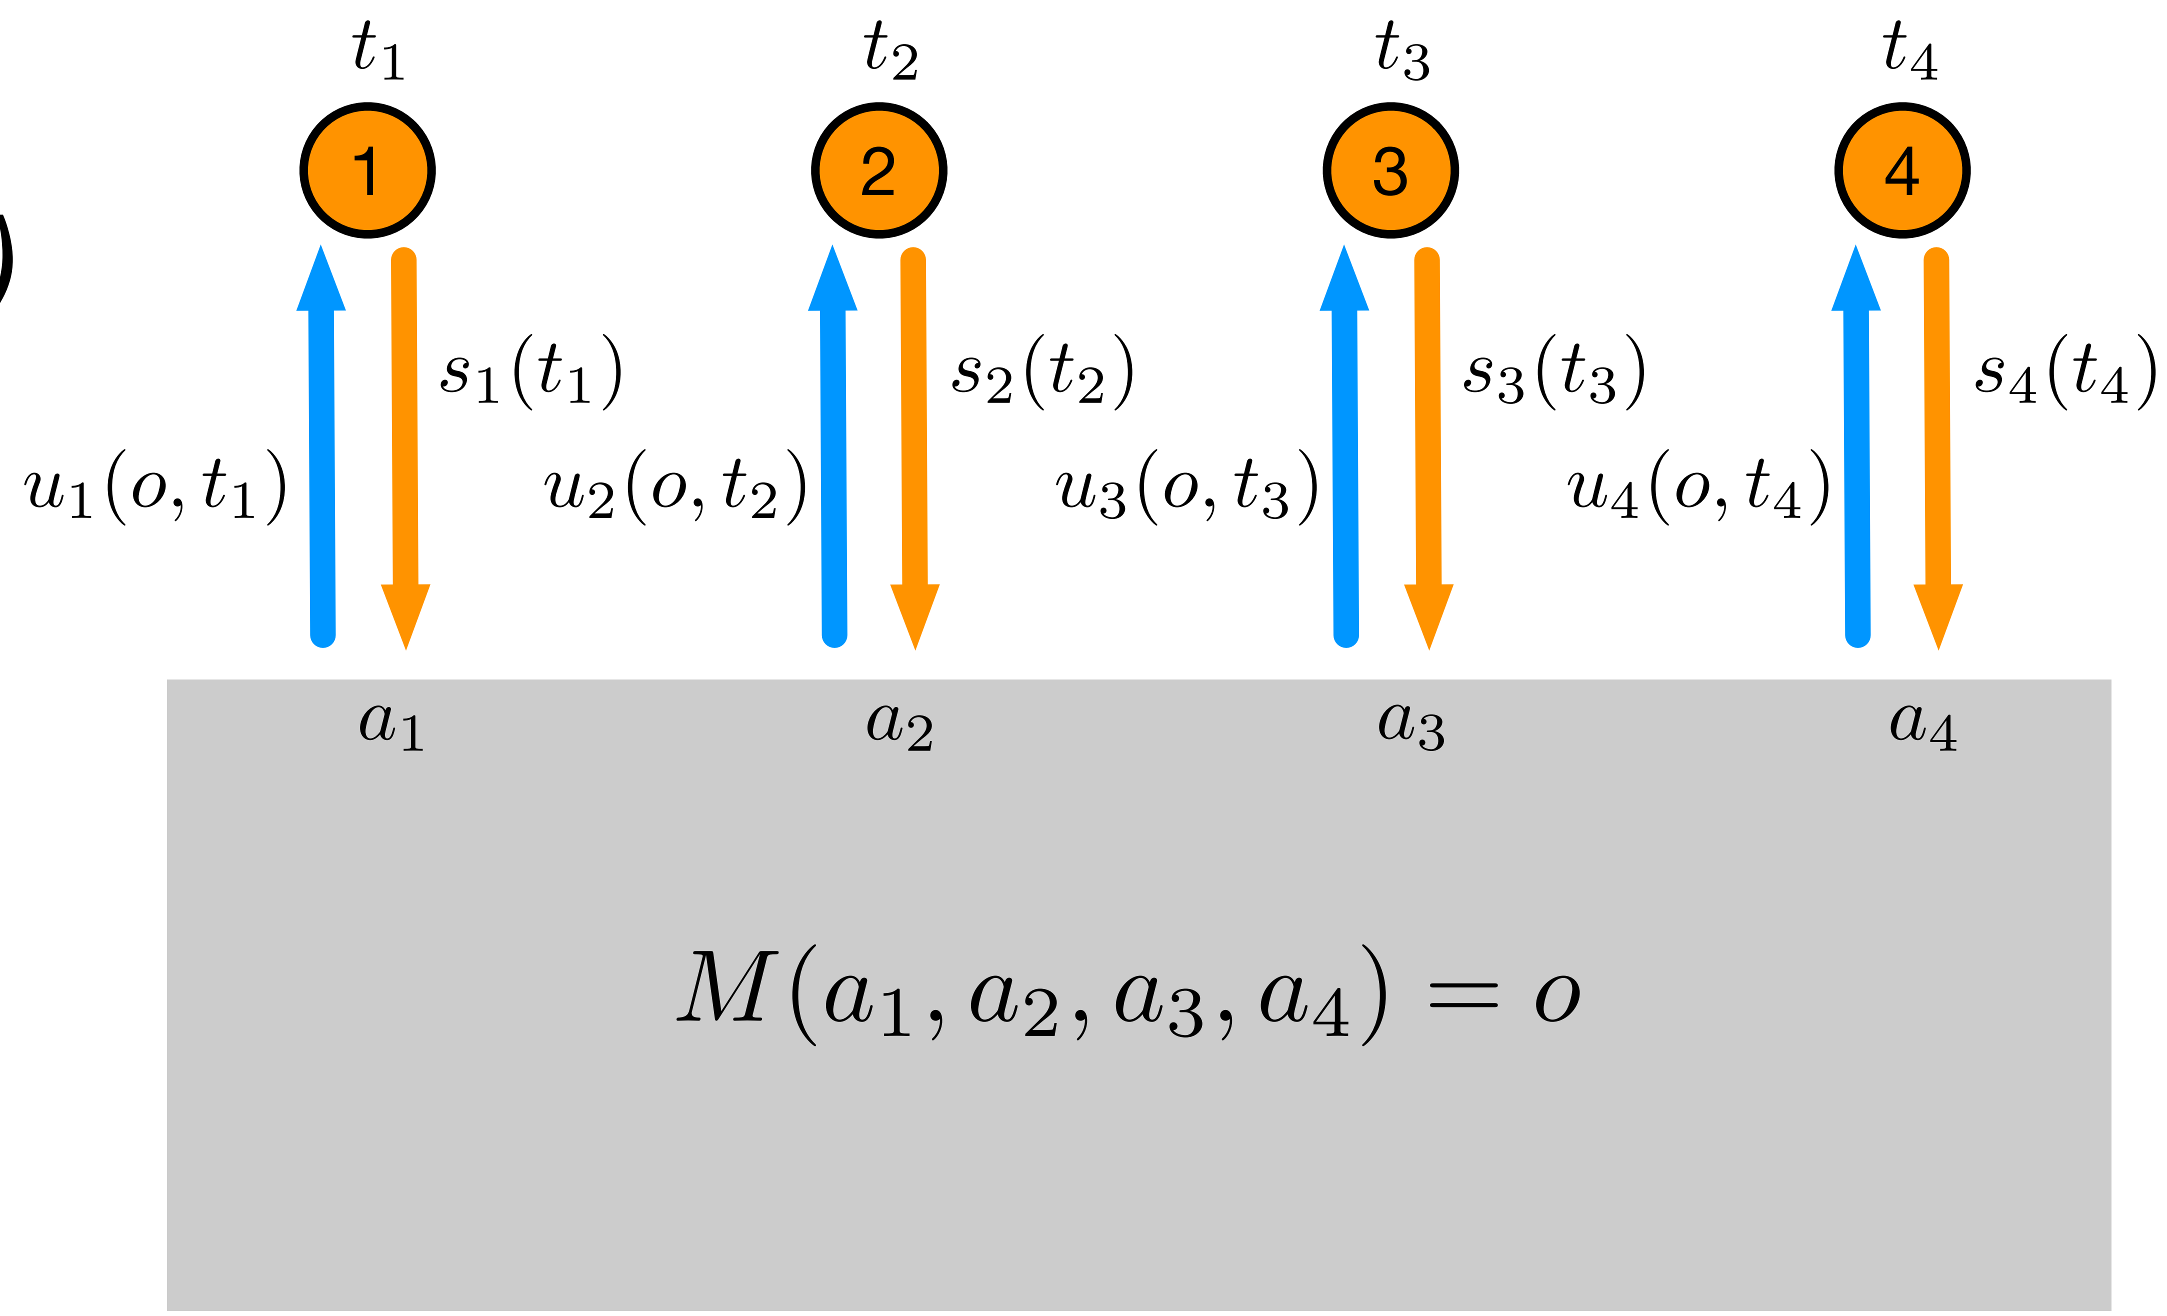

What is a mechanism?

preference order

- Each ***agent v*** has a hidden ***type t_v*** (*private information*) from a known set of ***types T*** and an ***action space A*** *(different algorithms)*

proposal orders

- The ***strategy*** of an agent *v* is a function ***$s_v : T \to A$***

matchings

- There is a set of ***outcomes O*** and each agent has a ***utility function $u_v : T \times O \to \mathbb{R}$***

- ***Mechanism M*** maps actions (a_1, a_2, …, a_n) to an *outcome*; a mechanism is a game of incomplete information

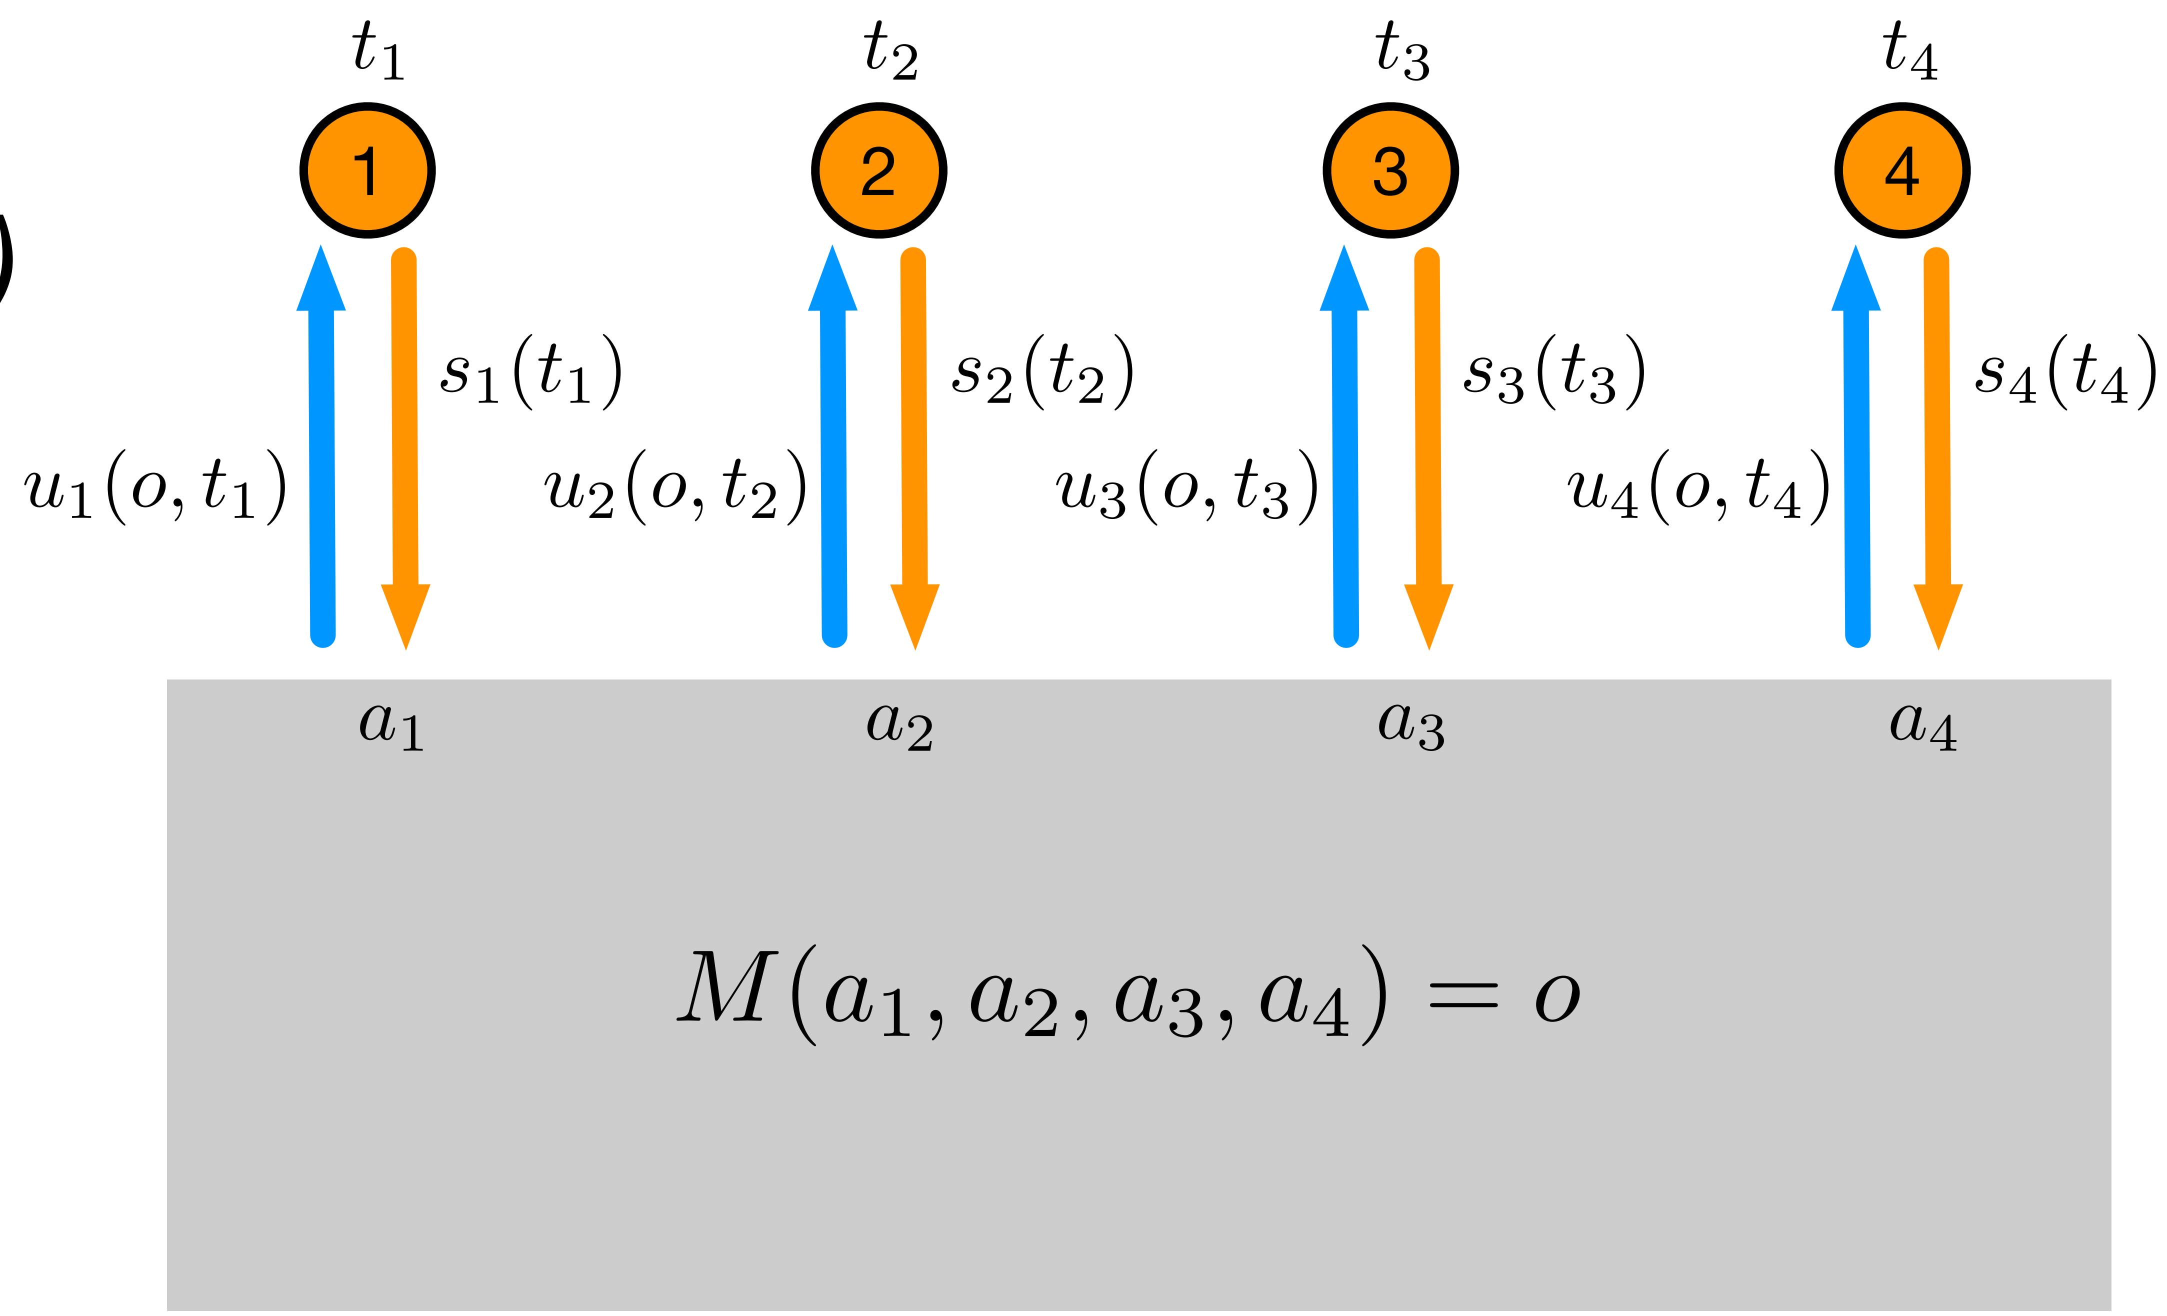

Computation in a mechanism

- Typically mechanism *implements* some social choice function; a ***social choice function F*** is a mapping from the (private) types to the outcomes
- ”Give the item to the agent that prefers it the most” or ”Find an assignment that maximises total welfare”
- A mechanism ***M implements*** *(in dominant strategies)* ***F*** if there ***exists*** a Nash equilibrium strategy $\underline{s}$ such that for all $(t_1, t_2, \ldots, t_n)$ we have that $M(\underline{s}_1(t_1), \underline{s}_2(t_2), \ldots, \underline{s}_n(t_n)) = F(t_1, t_2, \ldots, t_n)$
- ***Nash equilibrium*** = no agent has incentive to unilaterally deviate from $\underline{s}$

Revelation principle

- Limits the family of relevant mechanisms to ***direct*** and ***truthful*** (incentive compatible) mechanisms
- Mechanism is ***direct*** if the actions of the agents are the types (i.e. ”reveal your type”)
- Mechanism is ***truthful*** or ***incentive-compatible*** if reporting the true type is a dominant strategy (truth gives an outcome at least as good as any lie)
- ***Revelation principle*****:** If there exists a mechanism that implements a social choice function *F*, then there exists a direct truthful mechanism that implements *F* (Meyerson, 1981)

Revelation principle

- ***Theorem***: If there exists a mechanism *M* that implements a social choice function *F*, then there exists a ***direct truthful mechanism that implements F***

- ***Proof***: Direct mechanism *M'*: Let agents report their types t_v (can lie). Take the NE $\underline{\mathbf{s}}$ for *M*, for each *v* compute $\underline{s}_v(t_v) = t'_v$ and simulate *M*.

- Can focus on very specific types of mechanisms, ***truthfulness comes for "free"***

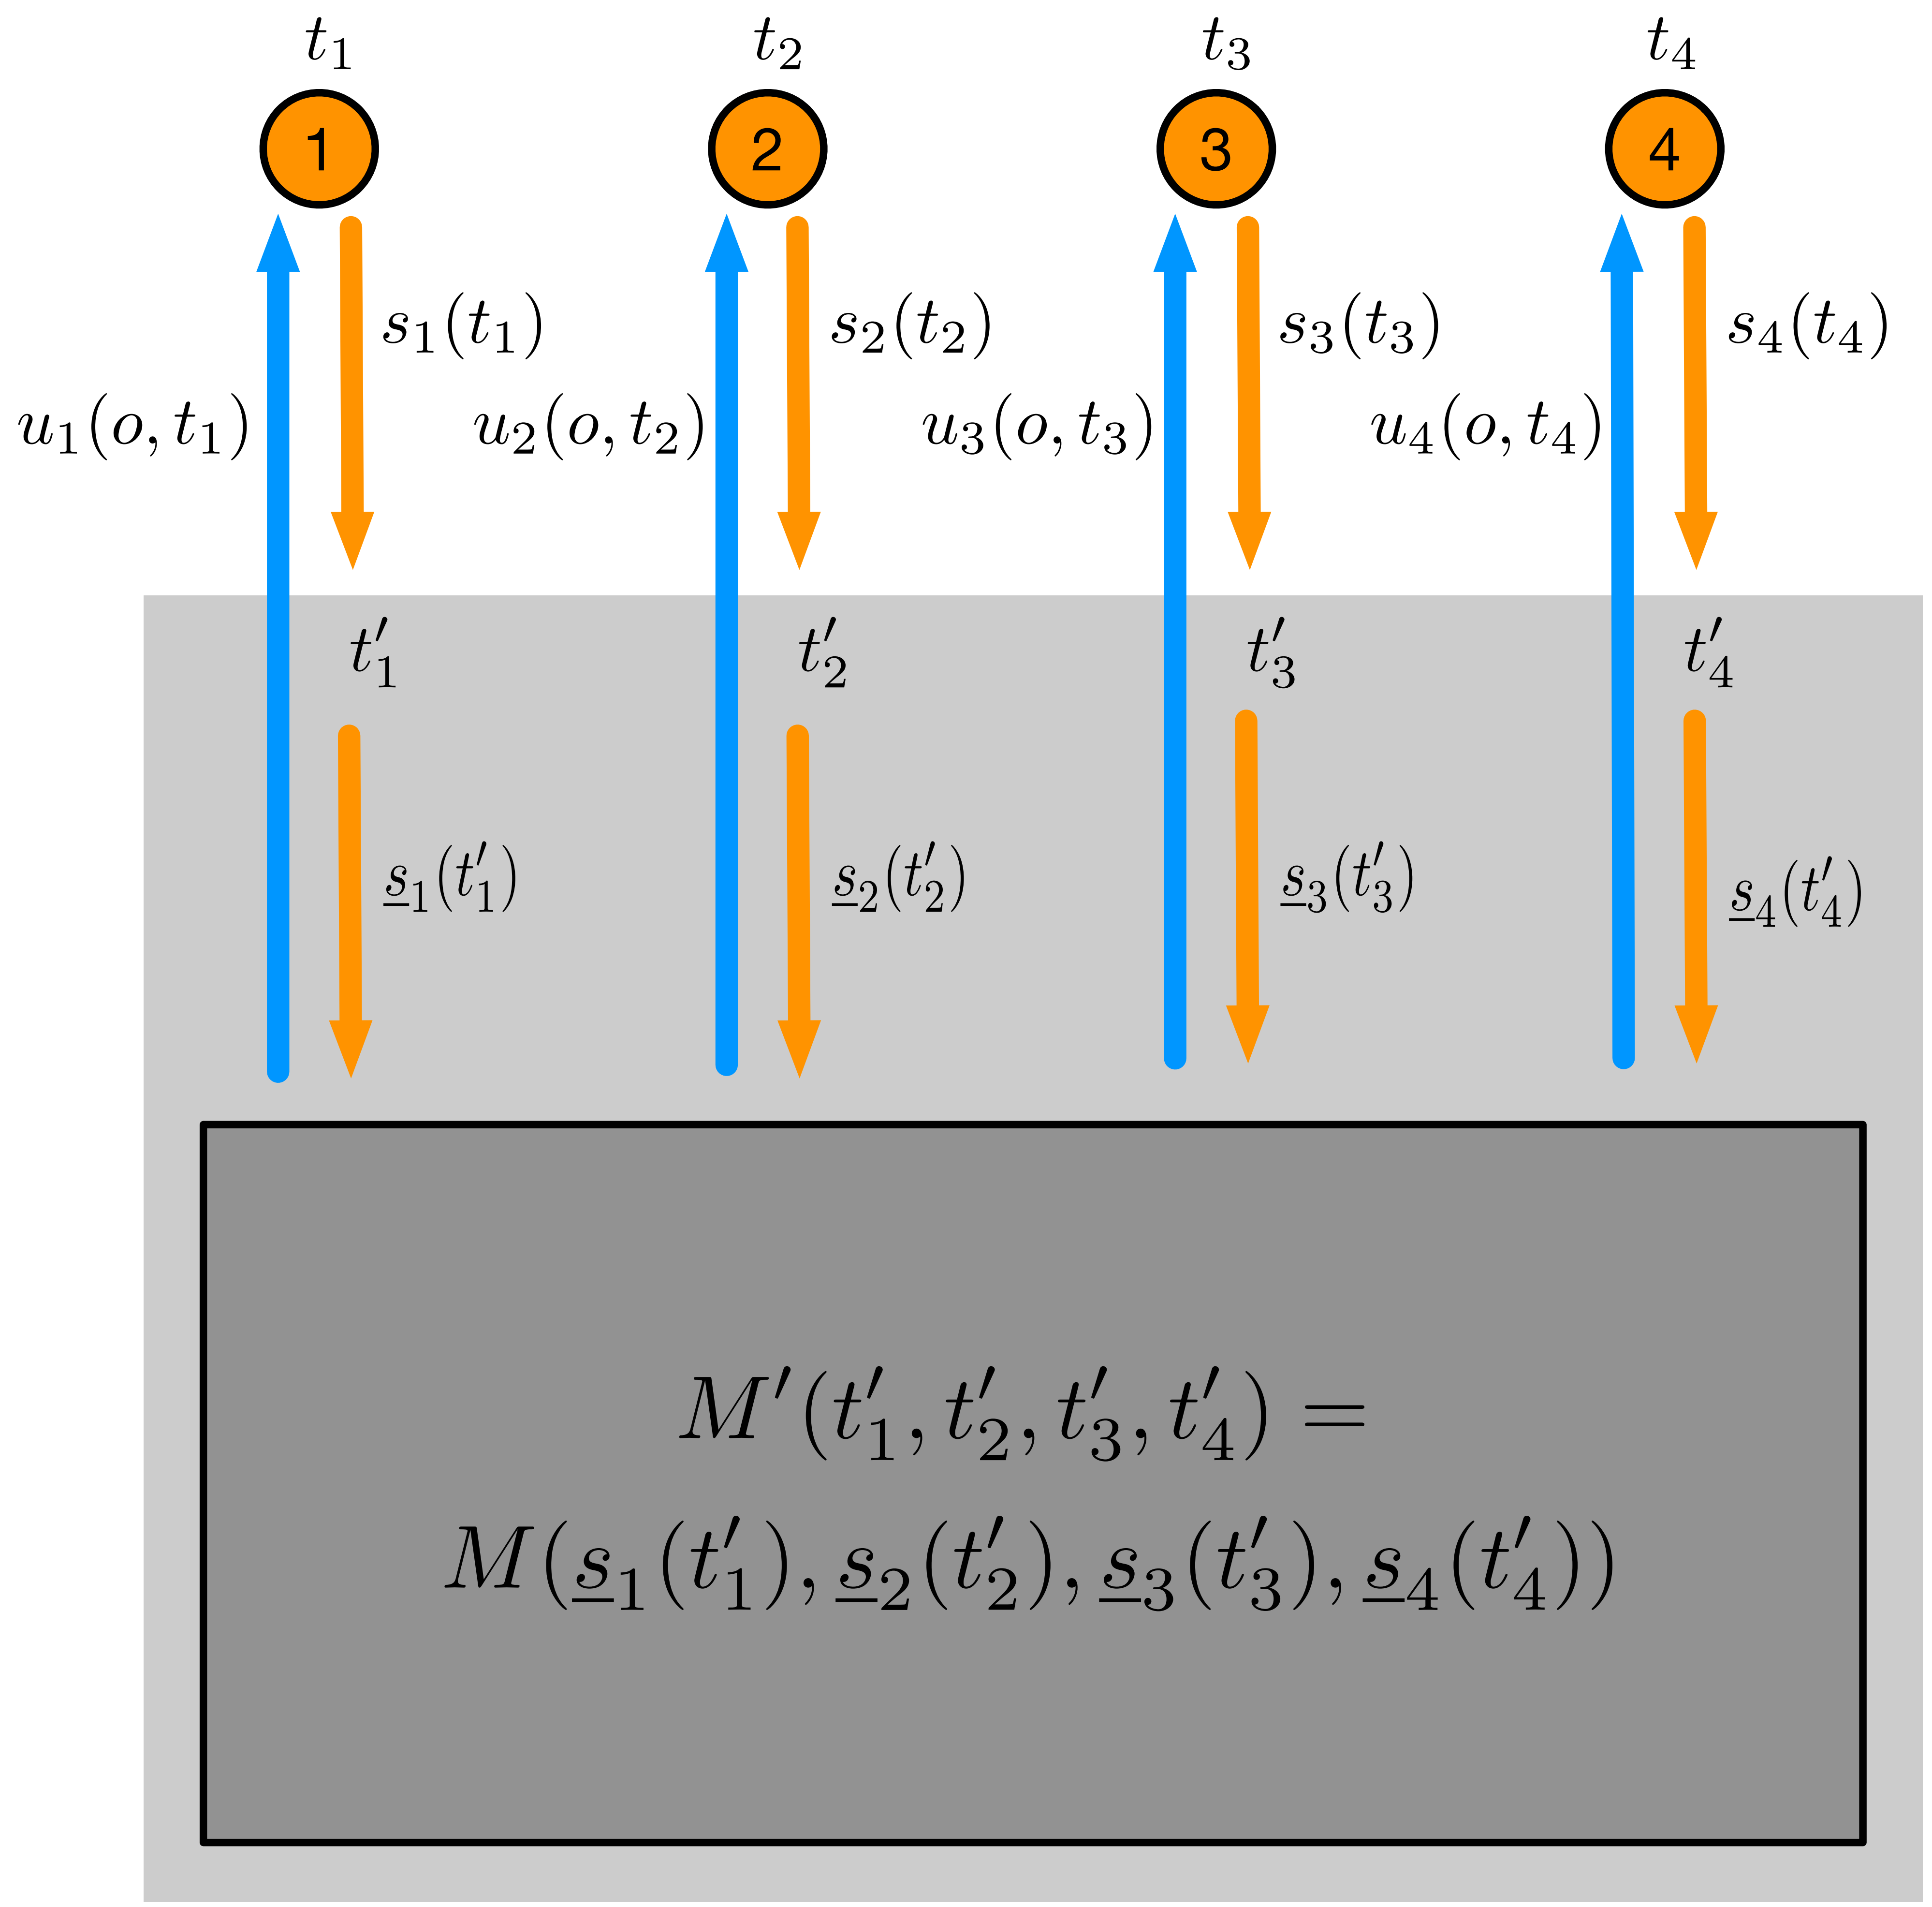

Truthfulness

- In a sense, revelation principle ***transfers the maximal amount of computation from the agents to the mechanism*** function M: agents only need to know (compute) their own type
- More generally, truthfulness can be seen as a ***computational guarantee for the participating agents***: mechanism is ”as easy as possible”
- Unfortunately this is not always as simple: type can be cumbersome to list, knowing type also ***requires computation***
- Indirect mechanisms can have ***exponentially smaller communication complexity*** (Conitzer and Sandholm, 2004)

Truthfulness and other goals

- Given revelation principle, truthfulness has become a standard assumption in mechanism design

- There are also other goals, such as ***individual rationality*** (utility > 0), ***budget balance*** (sum of payments ≥ 0), ***effectiveness*** (optimal solution), and ***revenue maximisation*** (maximise sum of payments)

- While optimisation is ”just” one of the goals, as we will see next, it often also is necessary

Existence of mechanisms

- Deferred acceptance is an example of a ***mechanism without money***
- Revelation principle with the Gibbard-Satterwaite Theorem implies that for general preferences, only dictatorial rules can be implemented in dominant strategies → there is ***no general framework for mechanisms without money***
- One solution is to include ***payments*** (money) in the mechanism

”The” truthful mechanism: Vickrey-Clarke-Groves (VCG)

- Family of direct, truthful mechanisms with payments: in addition to an outcome, ***mechanism computes a payment p*** for each agent (can be positive or negative)
- Works in any *utilitarian setting*: goal is to ***maximise the total welfare of the agents***
- Generally ***computationally infeasible to implement***: requires computing the optimum assignment

VCG mechanism

- Each agent submits their type, mechanism computes the assignment that ***maximises total utility***
- Each agent gets their utility + a ***payment function p*** consisting of two parts:
  + *The total utility of all other agents* (the incentive)
  - *The total utility of all other agents* in the assignment that maximises this sum (normalisation, known as *Clarke pivot*)
  - The difference is the *"externality"* of the agent (how much loss does it cause to the rest of the system)
- In total the mechanism has to compute the optimal assignment and the optimal assignment of a "subproblem" corresponding to each agent

Example: second price auction

- Let bids be b_1, b_2, …, b_n in decreasing order
- p = $-b_2$ for the agent with the highest bid, other agents get 0 utility
- p = 0 for agents that don't get the item
- This is ***truthful*** (each agent should bid true valuation)

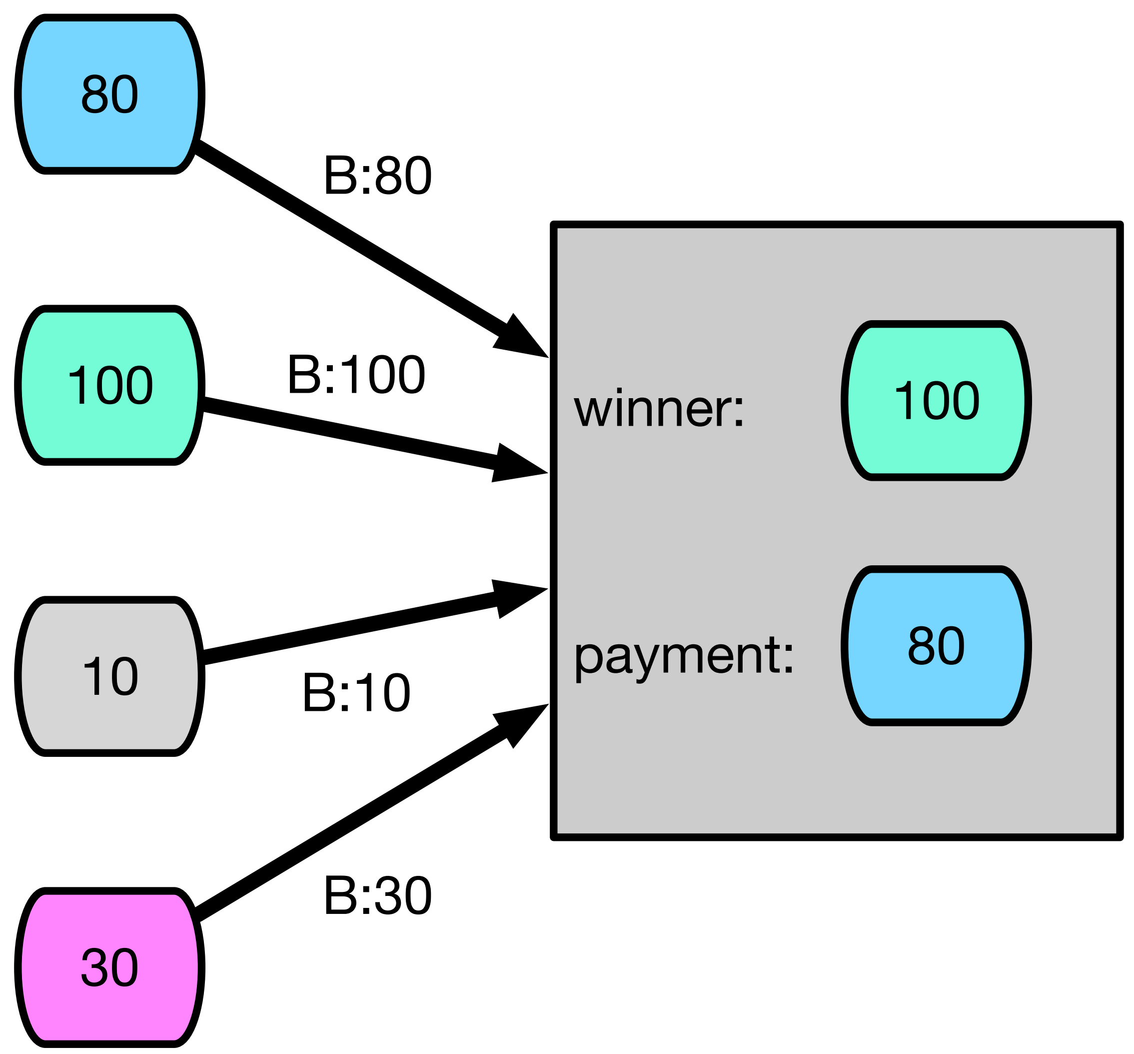

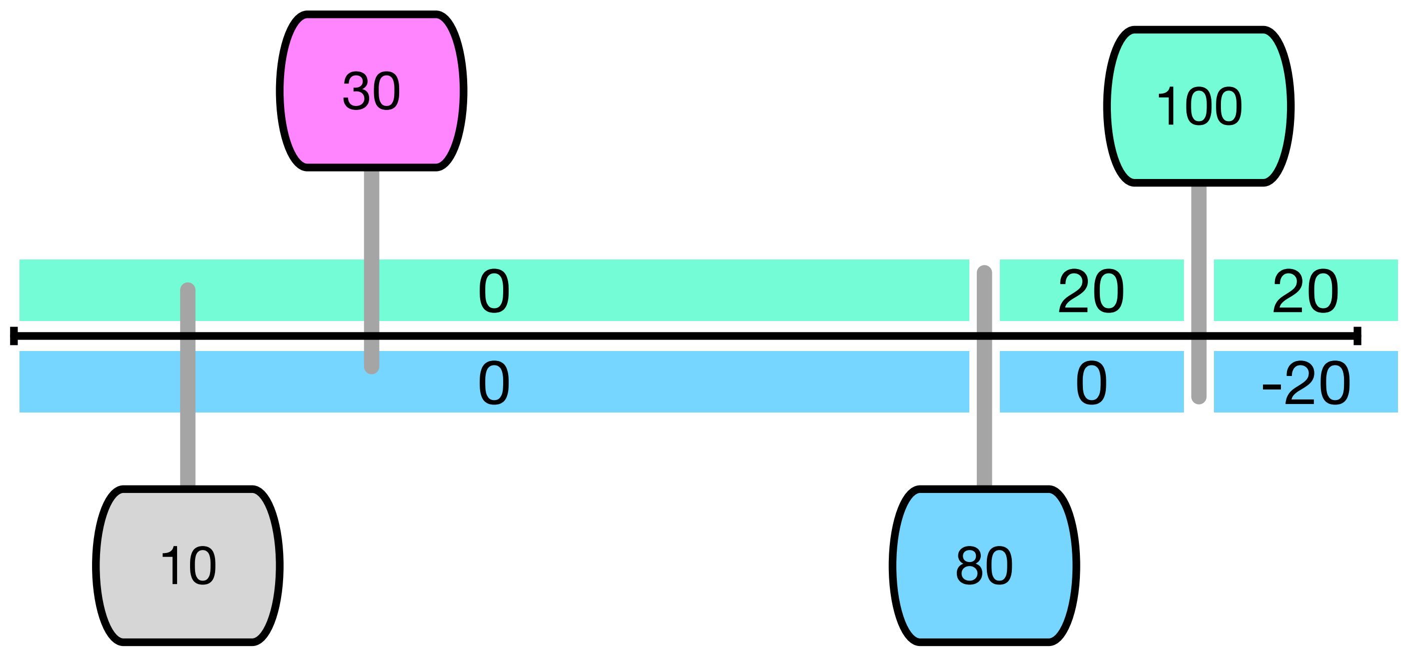

The payment function

- Optimal to the system, but not the mechanism designer: ***does not maximise payments***
- Payment function ensures that all agents are incentivised to maximise the total utility of the system: ***payment + own utility = total utility***
- Intuitively makes sense that the ***payment should not depend on agent's own type*****:** could incentivise to lie about type to increase payment
- ***Roberts' Theorem*** (1979): for many settings, only truthful mechanisms are affine maximisers (weighted generalisations of the VCG)

Computability

- While the VCG-mechanism in general is not computationally feasible, ***there are special cases that are***

- One significant use case are ***keyword advertisements on Google***: for each search, a multi-item auction is held to sell the advertisement spots on top of the search*

- A variant of the second-price auction that is computationally efficient

***As far as I understand, Google does not actually use the VCG-mechanism, but something similar that is computationally more efficient but lacks the theoretical guarantees**

Approximations?

- Unfortunately ***VCG does not generalise to approximation!***
- ***Problem*:** assume mechanism does not guarantee optimum assignment; a lie could ”accidentally” cause the algorithm to perform better, improving utility of the other agents → improves agent’s own payment
- Asymmetry of computation: typically cannot prove agents are computationally incapable of doing this
- There exist individual approximate and truthful mechanisms, but *no general theory*
- My work (in progress): ***distributed truthful mechanisms*** based on local guarantees

Applying in practice

- Mechanisms are actually deployed and tested in practice
- ***Examples:*** school choice, electricity markets, spectrum auctions
- Practical considerations often cause slight modifications to the mechanisms, this can cause serious problems!
- I will look at ***two use cases from Finland:*** school choice and spectrum auctions

Example: School choice in Finland

- Secondary school ***joint application system*** runs the deferred acceptance allocation algorithm (allocation to vocational / academic track or ammattikoulu / lukio)
- In 2019 students could apply to at most 5 tracks (has been changed to 7)
- Limit causes the system to lose truthfulness: applicants must estimate what are realistic schools to apply to with their GPA and what are the best safe options
- Causes ***computational burden on applicants*** and gives an advantage to sophisticated applicants

Otaniemi 2019

- Espoo has 11 upper secondary schools and Helsinki 14
- In 2018 the most competitive school in Espoo was Etelä-Tapiolan lukio with GPA requirement of 9.17/10
- In 2019 a new upper secondary school started in Otaniemi
- In 2019 joint application ***students with GPA 7.0 were admitted*** to Etelä-Tapiolan lukio (7.0 is minimum GPA required to apply to USS in Espoo); ***35 seats were left unfilled***

Espoo | HS Espoo

Otaniemestä kasvoi hetkessä yli tuhannen lukiolaisen keskittymä, kun metroaseman vieressä aloitti Espoon suurin lukio

Kun lukuvuosi alkoi uudessa Otaniemen lukiossa, luokista puuttui pöytiä ja tuoleja. Nyt Espoon suurimmassa lukiossa arki alkaa löytää muotoaan.

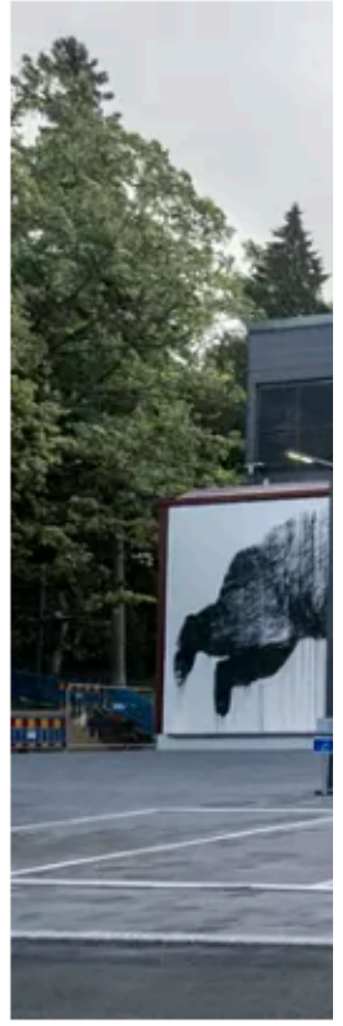

Elokuussa aloittanut
Seinää koristaa kuva

Johanna Juupaluor
25.9.2019 10:13 | Päi

Espoo | HS Espoo

Ennennäkemätön romahdus yhdessä Suomen parhaista lukioista: 35 opiskelupaikkaa täyttämättä, sisään jopa 7,0:n keskiarvolla

Espoon Etelä-Tapiolan lukiossa jäi 35 opiskelupaikkaa täyttämättä. Viime vuonna Etelä-Tapiolan lukioon oli korkein keskiarvoraja, nyt sinne oli helpointa päästä sisään.

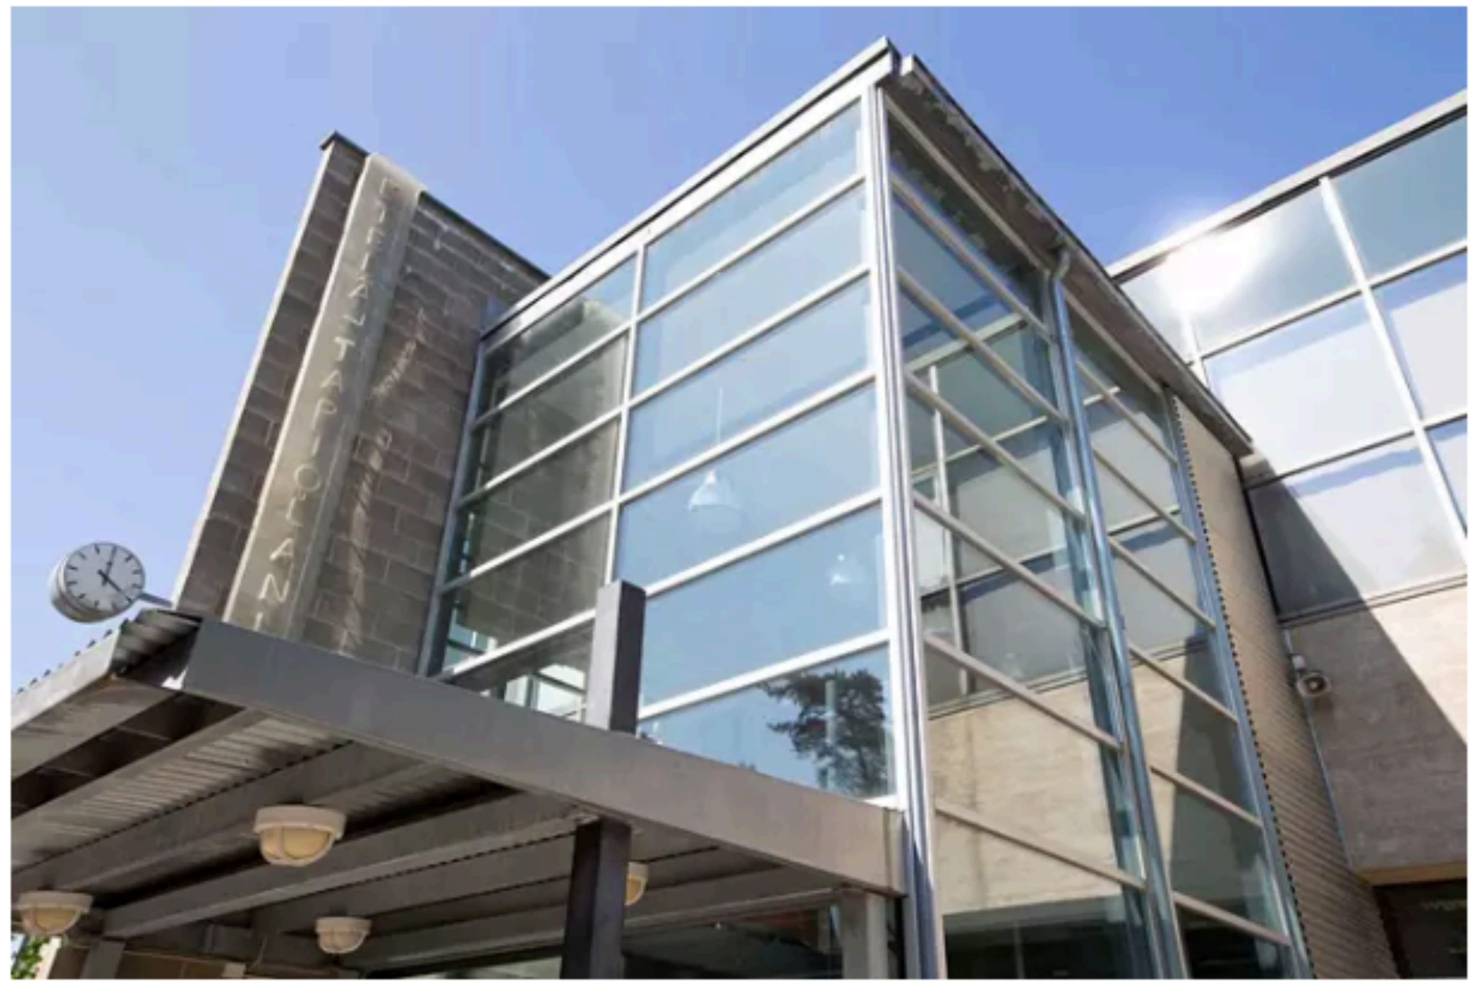

Etelä-Tapiolan lukio on ollut perinteisesti Espoon vaikein lukio päästä sisään ja siellä on kirjoitettu Suomen parhaimpia ylioppilastuloksia useampana vuotena. Tänä vuonna Etelä-Tapiolan lukioon oli helpointa päästä sisään. KUVA: MARKO SORSA

Johanna Juupaluoma HS
13.6.2019 13:33 | Päivitetty 13.6.2019 14:27

Example: combinatorial auctions

- Auctions where multiple items are sold and buyers have ***different valuations for different subsets of items***
  - e.g. buyer needs at least two items, some buyers view all items as substitutes while others require a certain item, …
- Spectrum auction: government sells ***the right to transmit signals*** on specific bands of the electromagnetic spectrum
- A teleoperator requires a sufficient volume and spread of bands, geographic coverage

Spectrum auctions

- VCG is generally not practical due to having to bid for every possible combination of items
- In the US, FCC has been using competitive auctions designed by prominent economists since 1994 (*simultaneous multiple-round ascending auction*)
  - Variants have been adopted e.g. in Germany and Finland
- Does not have game-theoretic guarantees, but is considered to be a reasonably good auction system in practice
- Relatively complicated to participate in: Milgrom et al. (2017) report a case where a ***US company saved $1 bn compared to competitors with sophisticated auction strategy***

2020 Finnish Spectrum auction

- In 2020 the operating ***licences for 25.1-27.5 GHz frequency band*** were auctioned
- Auction system was a ***simultaneous multiple-round ascending auction***, band was split into three items, and the minimum bid for each sub-band was set at 7 000 000€
- Finland has ***three established teleoperators***, Elisa, Telia, and DNA

Frequency auction ended

The auction ended on 8th of June 2020 at 11.40 am. Results of the auction:

Frequency bands	Winner	Winning bid
25.1 - 25.9 GHz (A)	Elisa Corporation	7 000 000 €
25.9 - 26.7 GHz (B)	Telia Finland Oyj	7 000 000 €
26.7 - 27.5 GHz (C)	DNA Plc	7 000 000 €

https://traficom.fi/en/communications/communications-networks/spectrum-auction-26-ghz-frequency-band

Even good mechanisms can fail…

- Systemic issue compounded by multiple factors:
  - When selling multiple items, VCG-type auctions often ***vulnerable to collusion***
  - ***Hard for new operators to enter the market*** without a portfolio of spectrum licences; bands for sale not sufficient and existing licences
  - Frequency band split into three intervals benefiting existing operators

Concluding

- ***Revelation principle:*** truthful mechanisms come for ”free”
- General truthful mechanisms without money do not exist as social choice functions ***do not exist***
- ***Money does help:*** VCG-mechanisms are a general framework for truthful utilitarian optimisation
  - Payment function ***incentivises maximising the common good***
- Truthfulness helps with the ***computational load of the participants***, important in practice!

Concluding

- ***Revelation principle:*** truthful mechanisms come for ”free”
- General truthful mechanisms without money do not exist as social choice functions ***do not exist***
- ***Money does help:*** VCG-mechanisms are a general framework for truthful utilitarian optimisation
  - Payment function ***incentivises maximising the common good***
- Truthfulness helps with the ***computational load of the participants***, important in practice!

Thank you for your attention!