Algorithmic Game Theory

Reading Group

Date:

What Is AGT?

What Is AGT?

- It is Not about developing computer games :(

What Is AGT?

- It is Not about developing computer games :(
- It is the formal study of interactions of different entities in strategic settings.

What Is AGT?

- It is Not about developing computer games :(
- It is the formal study of interactions of different entities in strategic settings.

People, Companies, Countries, Computers...

What Is AGT?

- It is Not about developing computer games :(
- It is the formal study of interactions of different entities in strategic settings.

People, Companies, Countries, Computers...

Players

What Is AGT?

- It is Not about developing computer games :(
- It is the formal study of interactions of different entities in strategic settings.

People, Companies, Countries, Computers...

Players

What Is AGT?

- It is Not about developing computer games :(
- It is the formal study of interactions of different entities in strategic settings.

People, Companies, Countries, Computers...

Players

- Success of a player may not only depend on their own action, but also the action of other players.

Why do (theoryCS) we care about Game Theory?

Why do (theoryCS) we care about Game Theory?

Two way interaction:

Why do (theoryCS) we care about Game Theory?

Two way interaction:

1. To use techniques from game theory to model and study certain aspects of computing

Why do (theoryCS) we care about Game Theory?

Two way interaction:

1. To use techniques from game theory to model and study certain aspects of computing
2. Contribute to game theory by offering an algorithmic perspective.

Course Goal 1: Designing systems with strategic participants.

Example: London Olympics

Example: London Olympics

A

B

C

D

Example: London Olympics

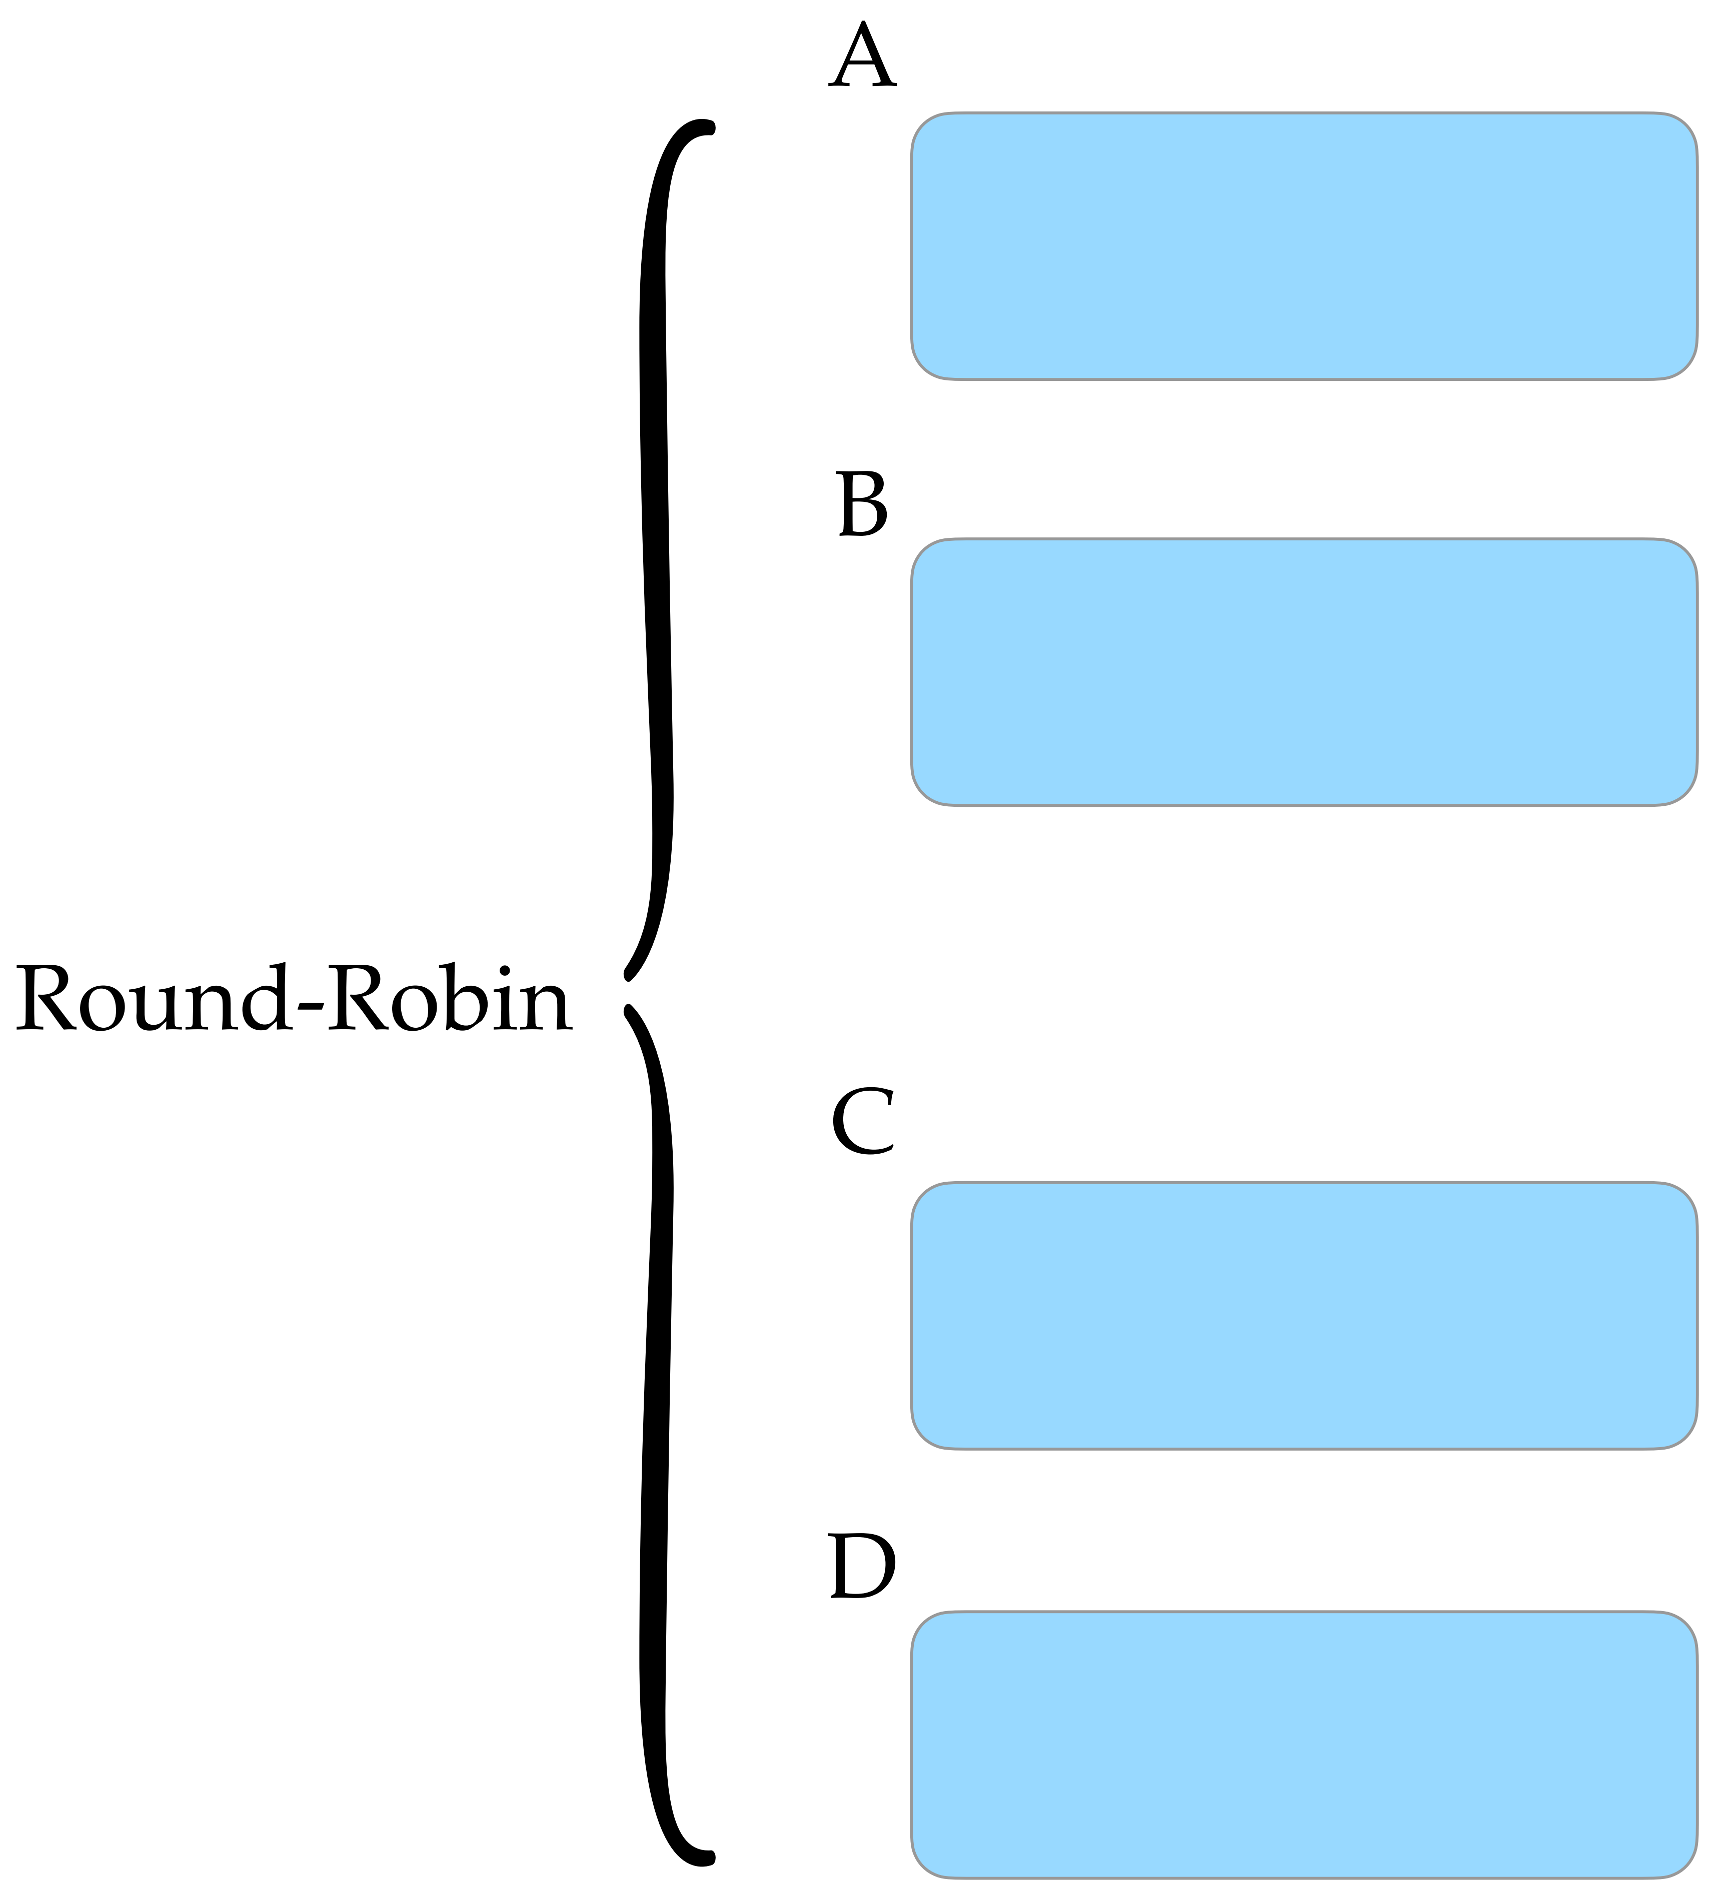

Example: London Olympics

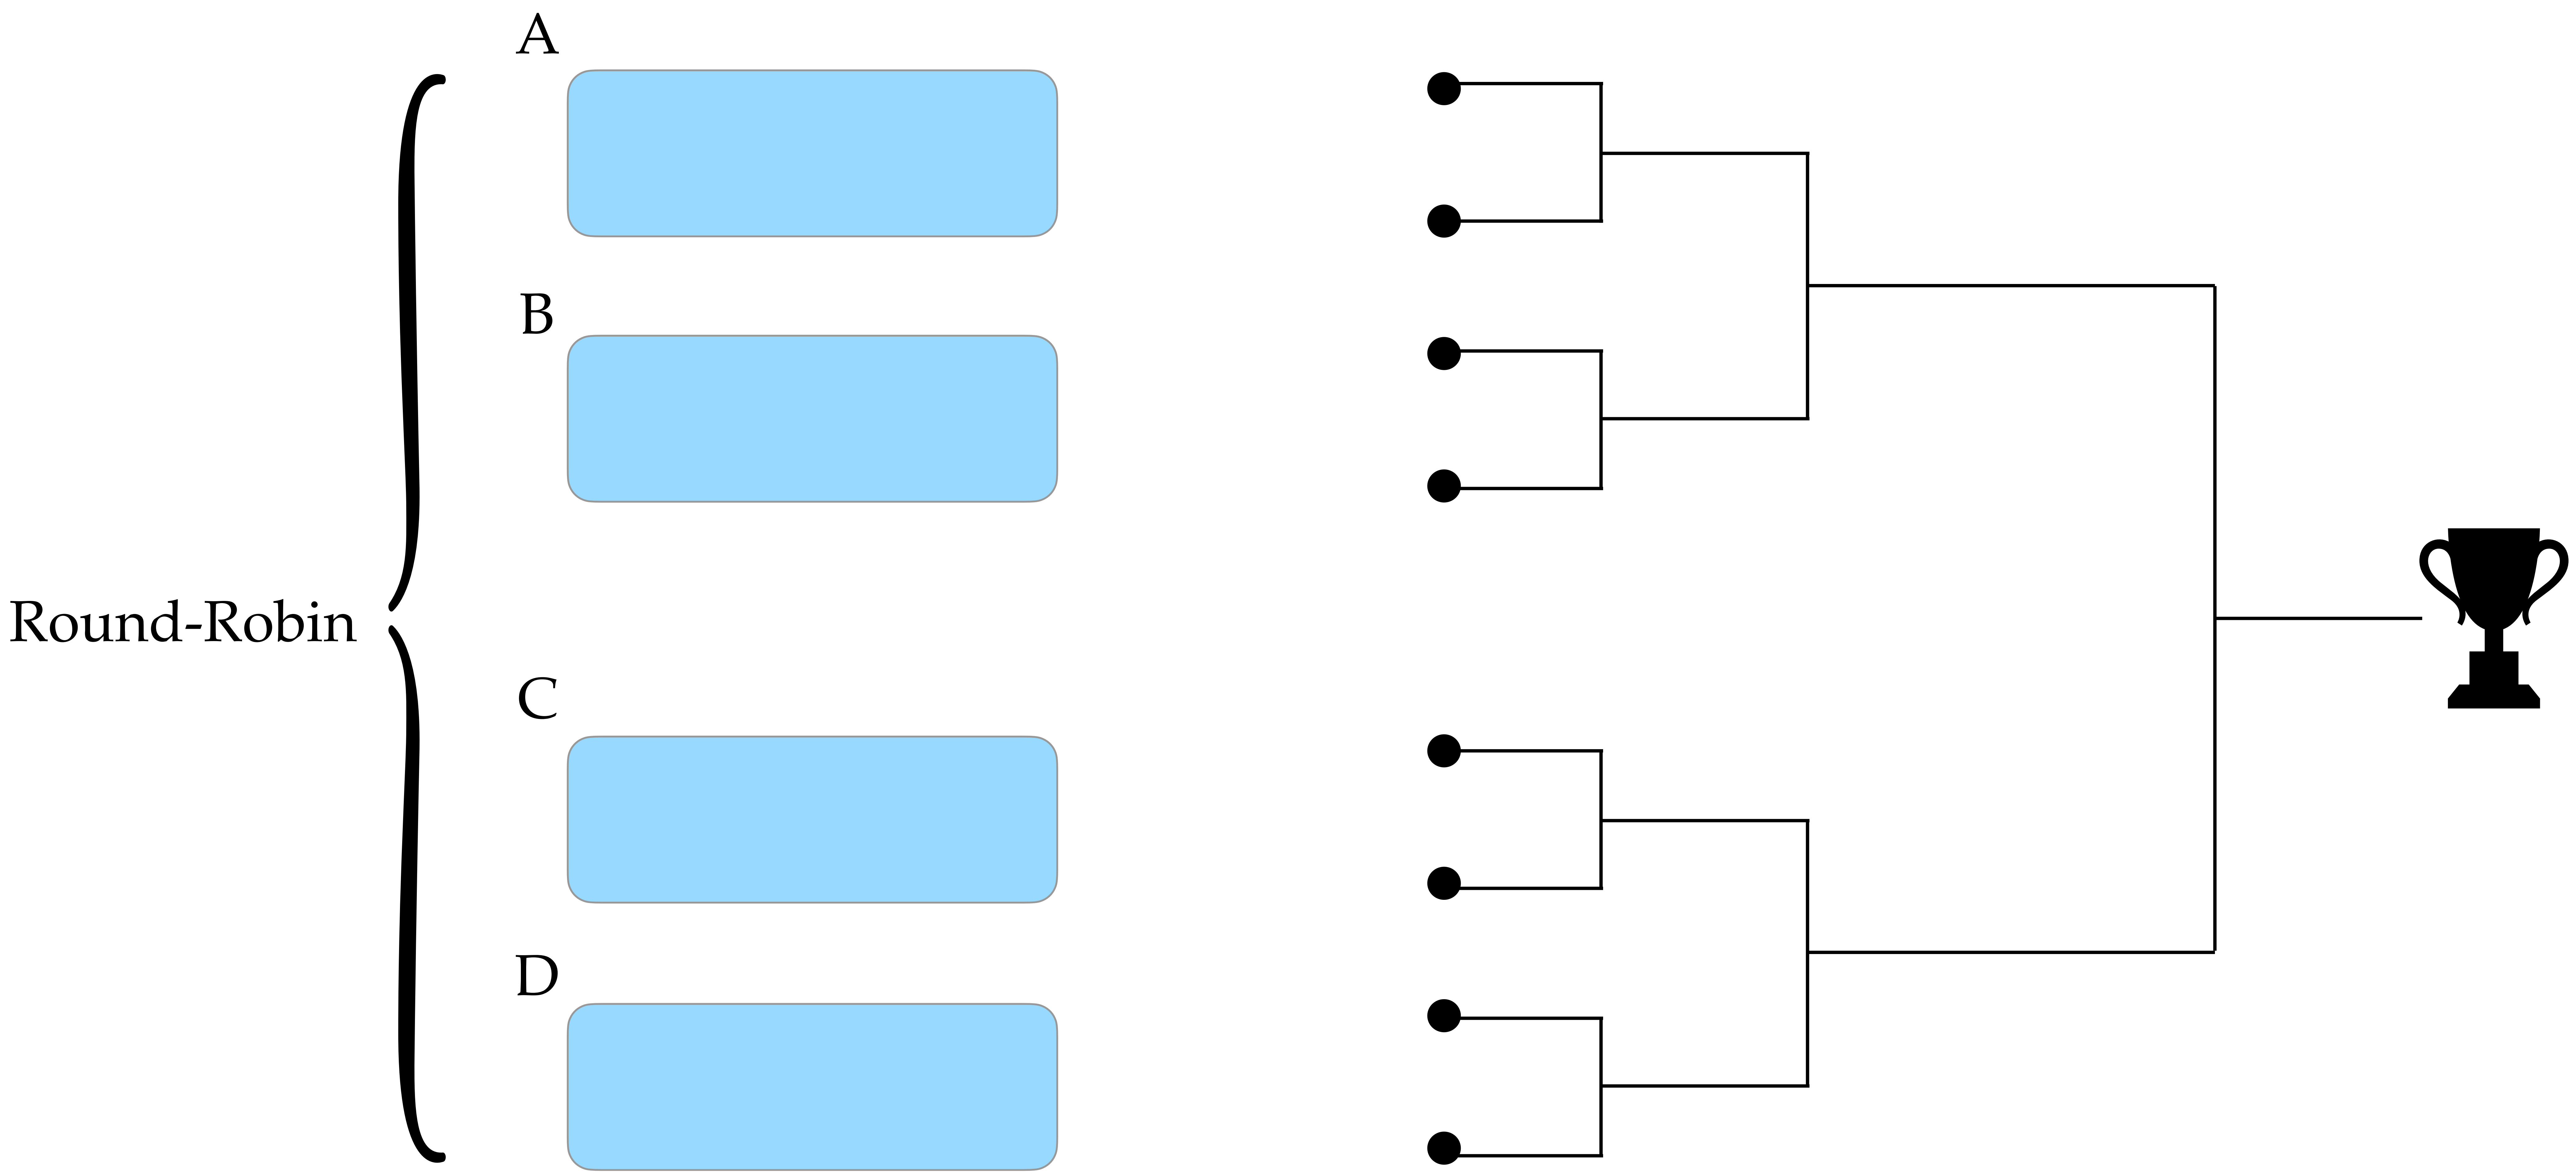

Example: London Olympics

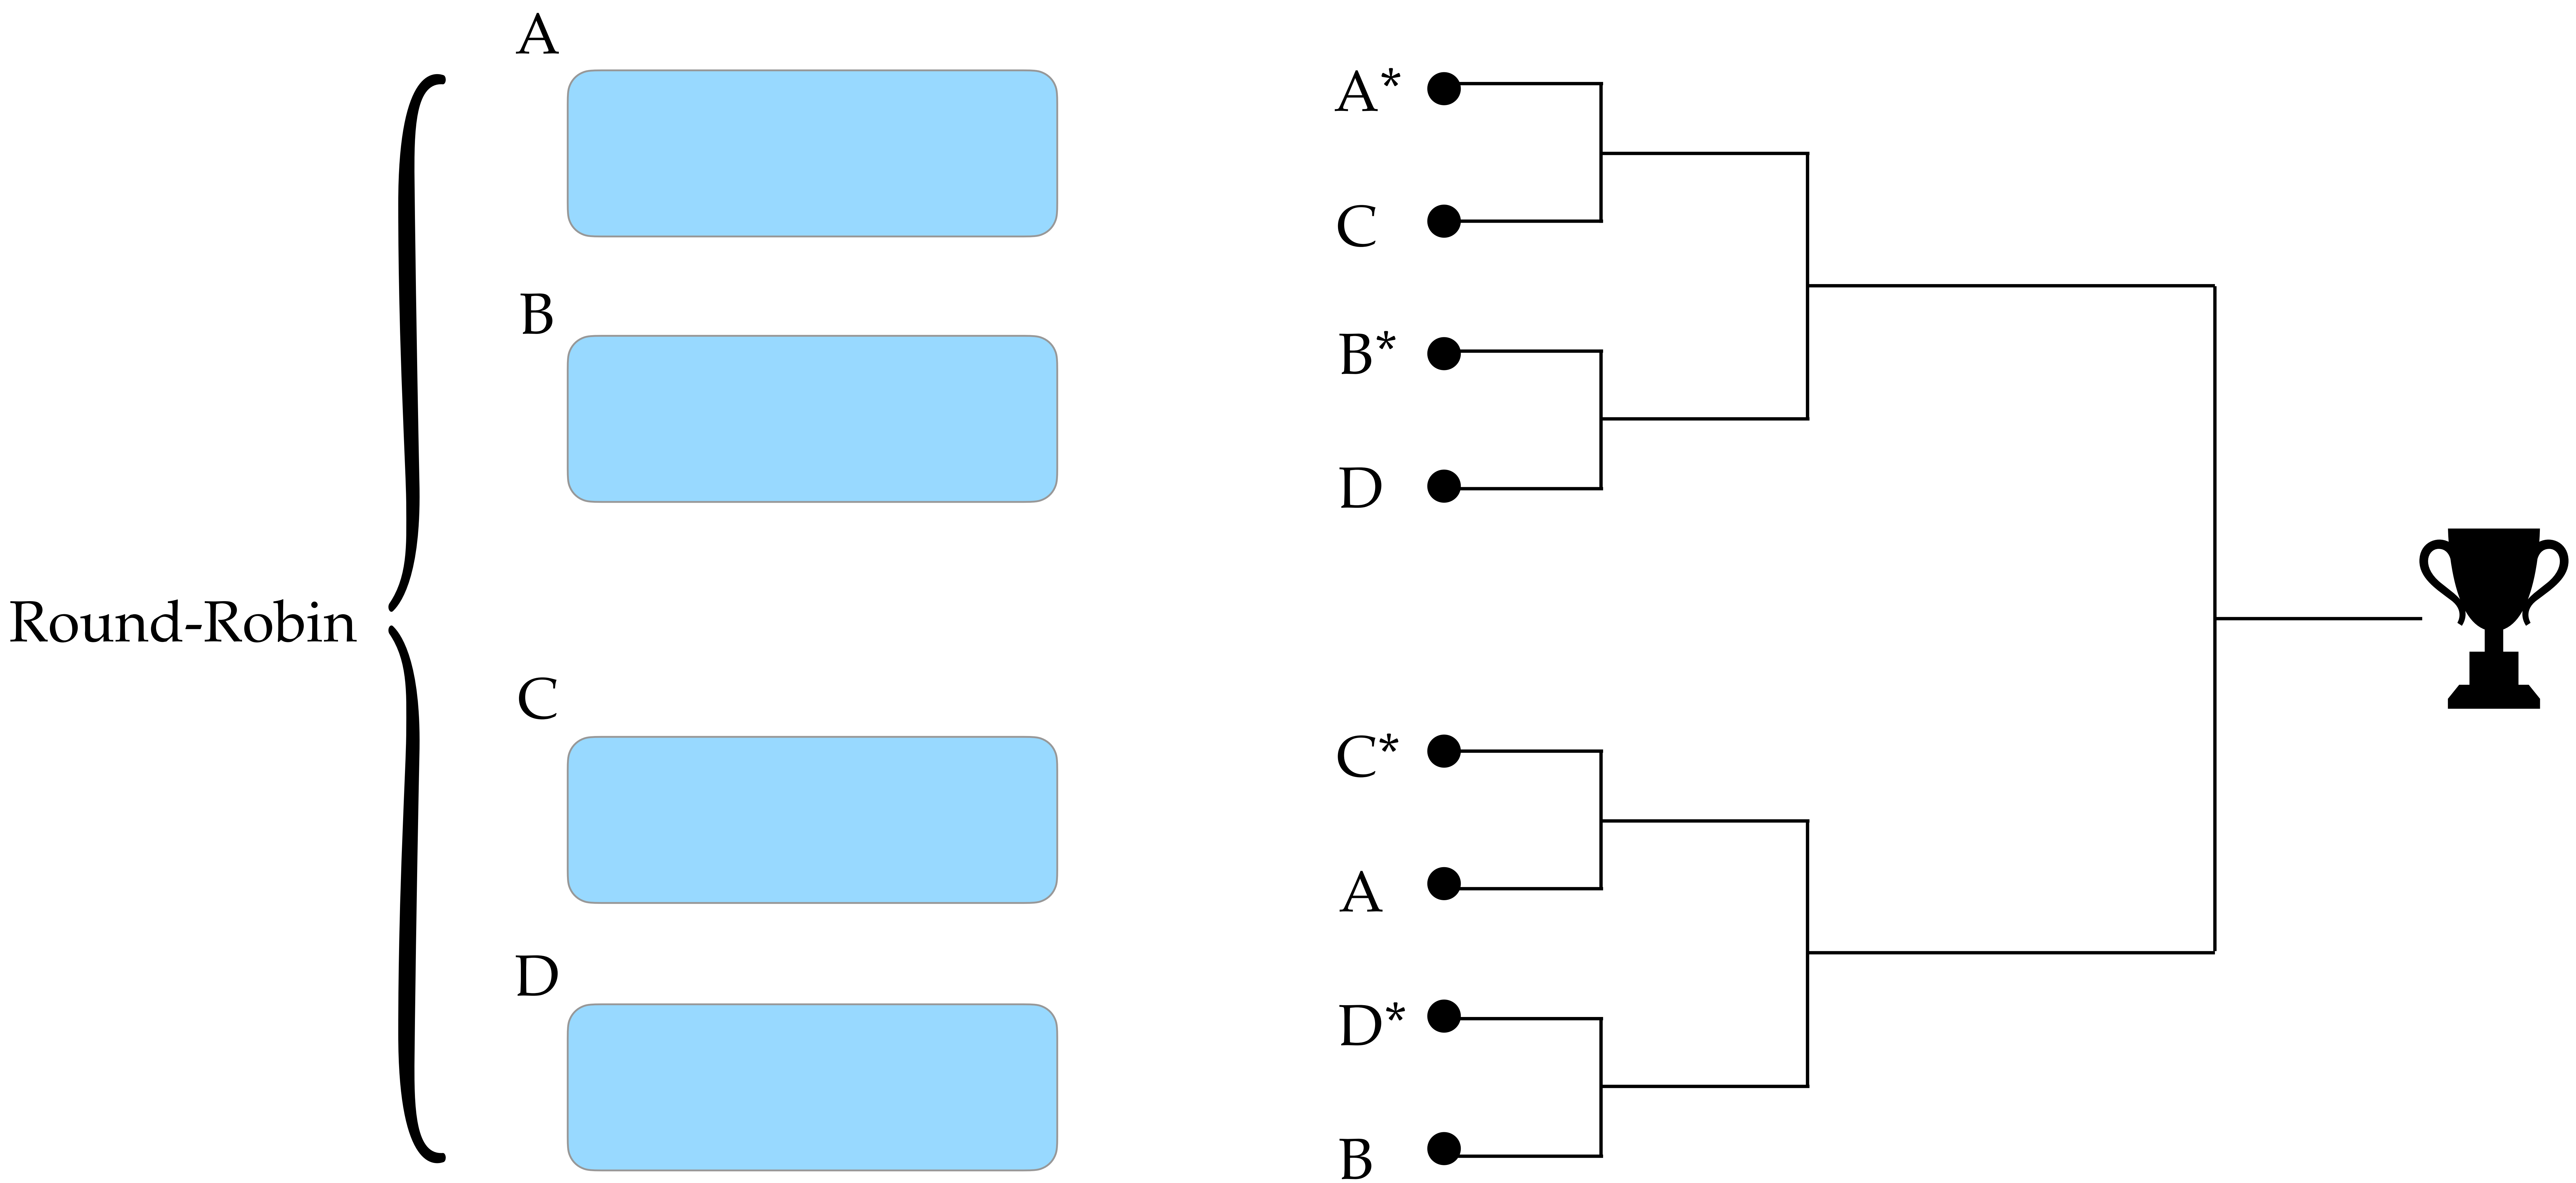

Example: London Olympics

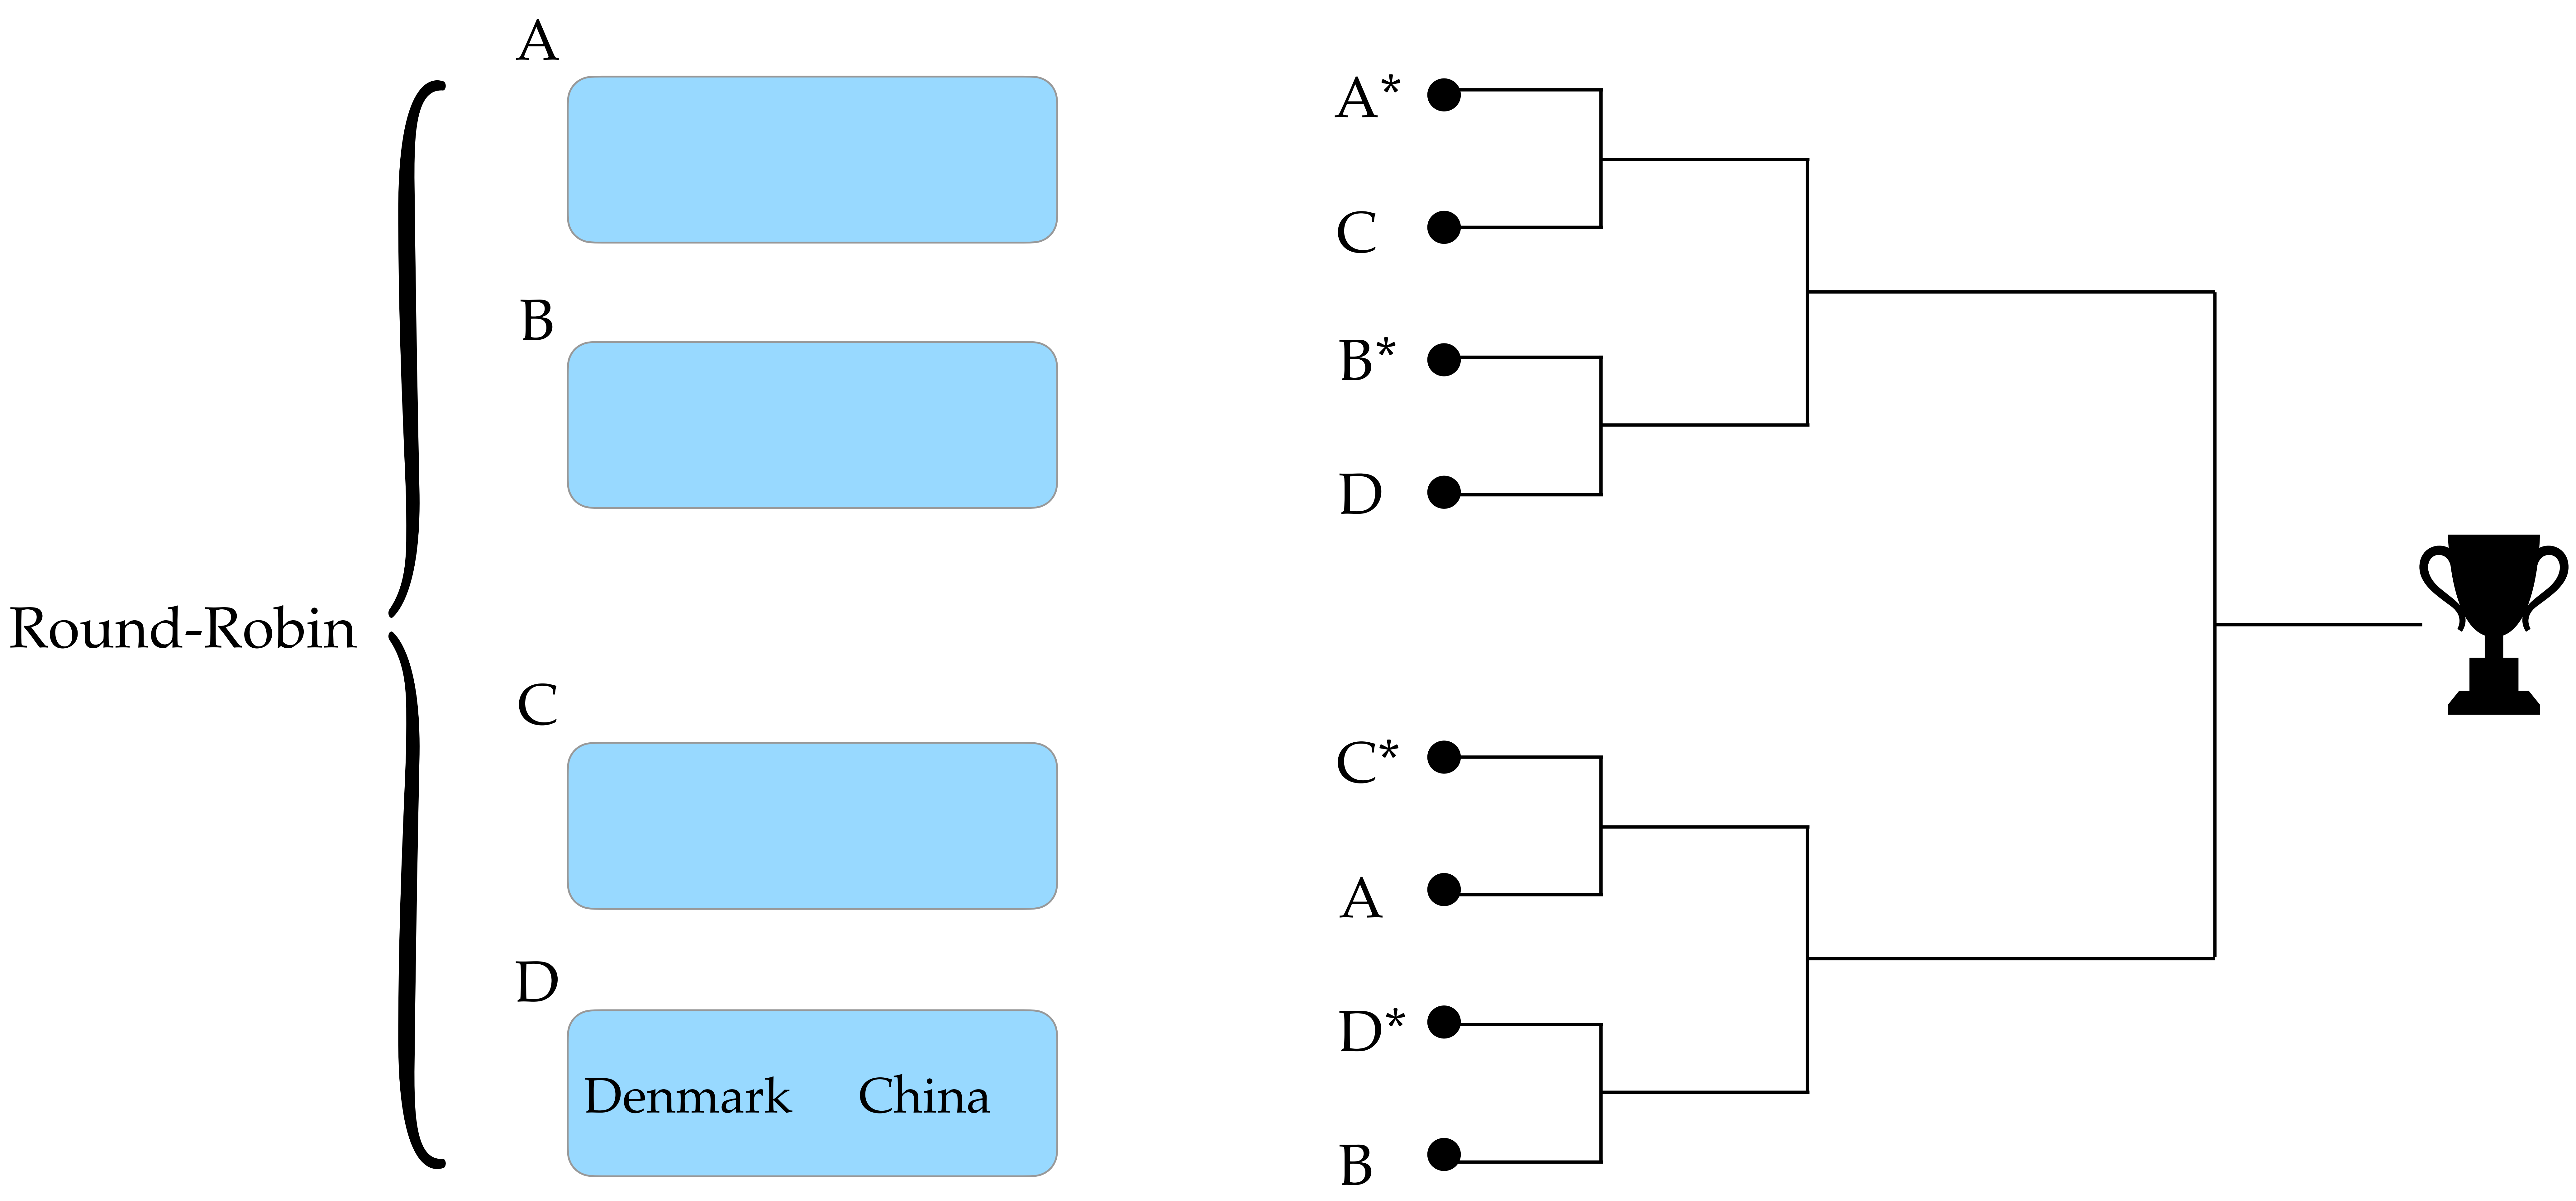

Example: London Olympics

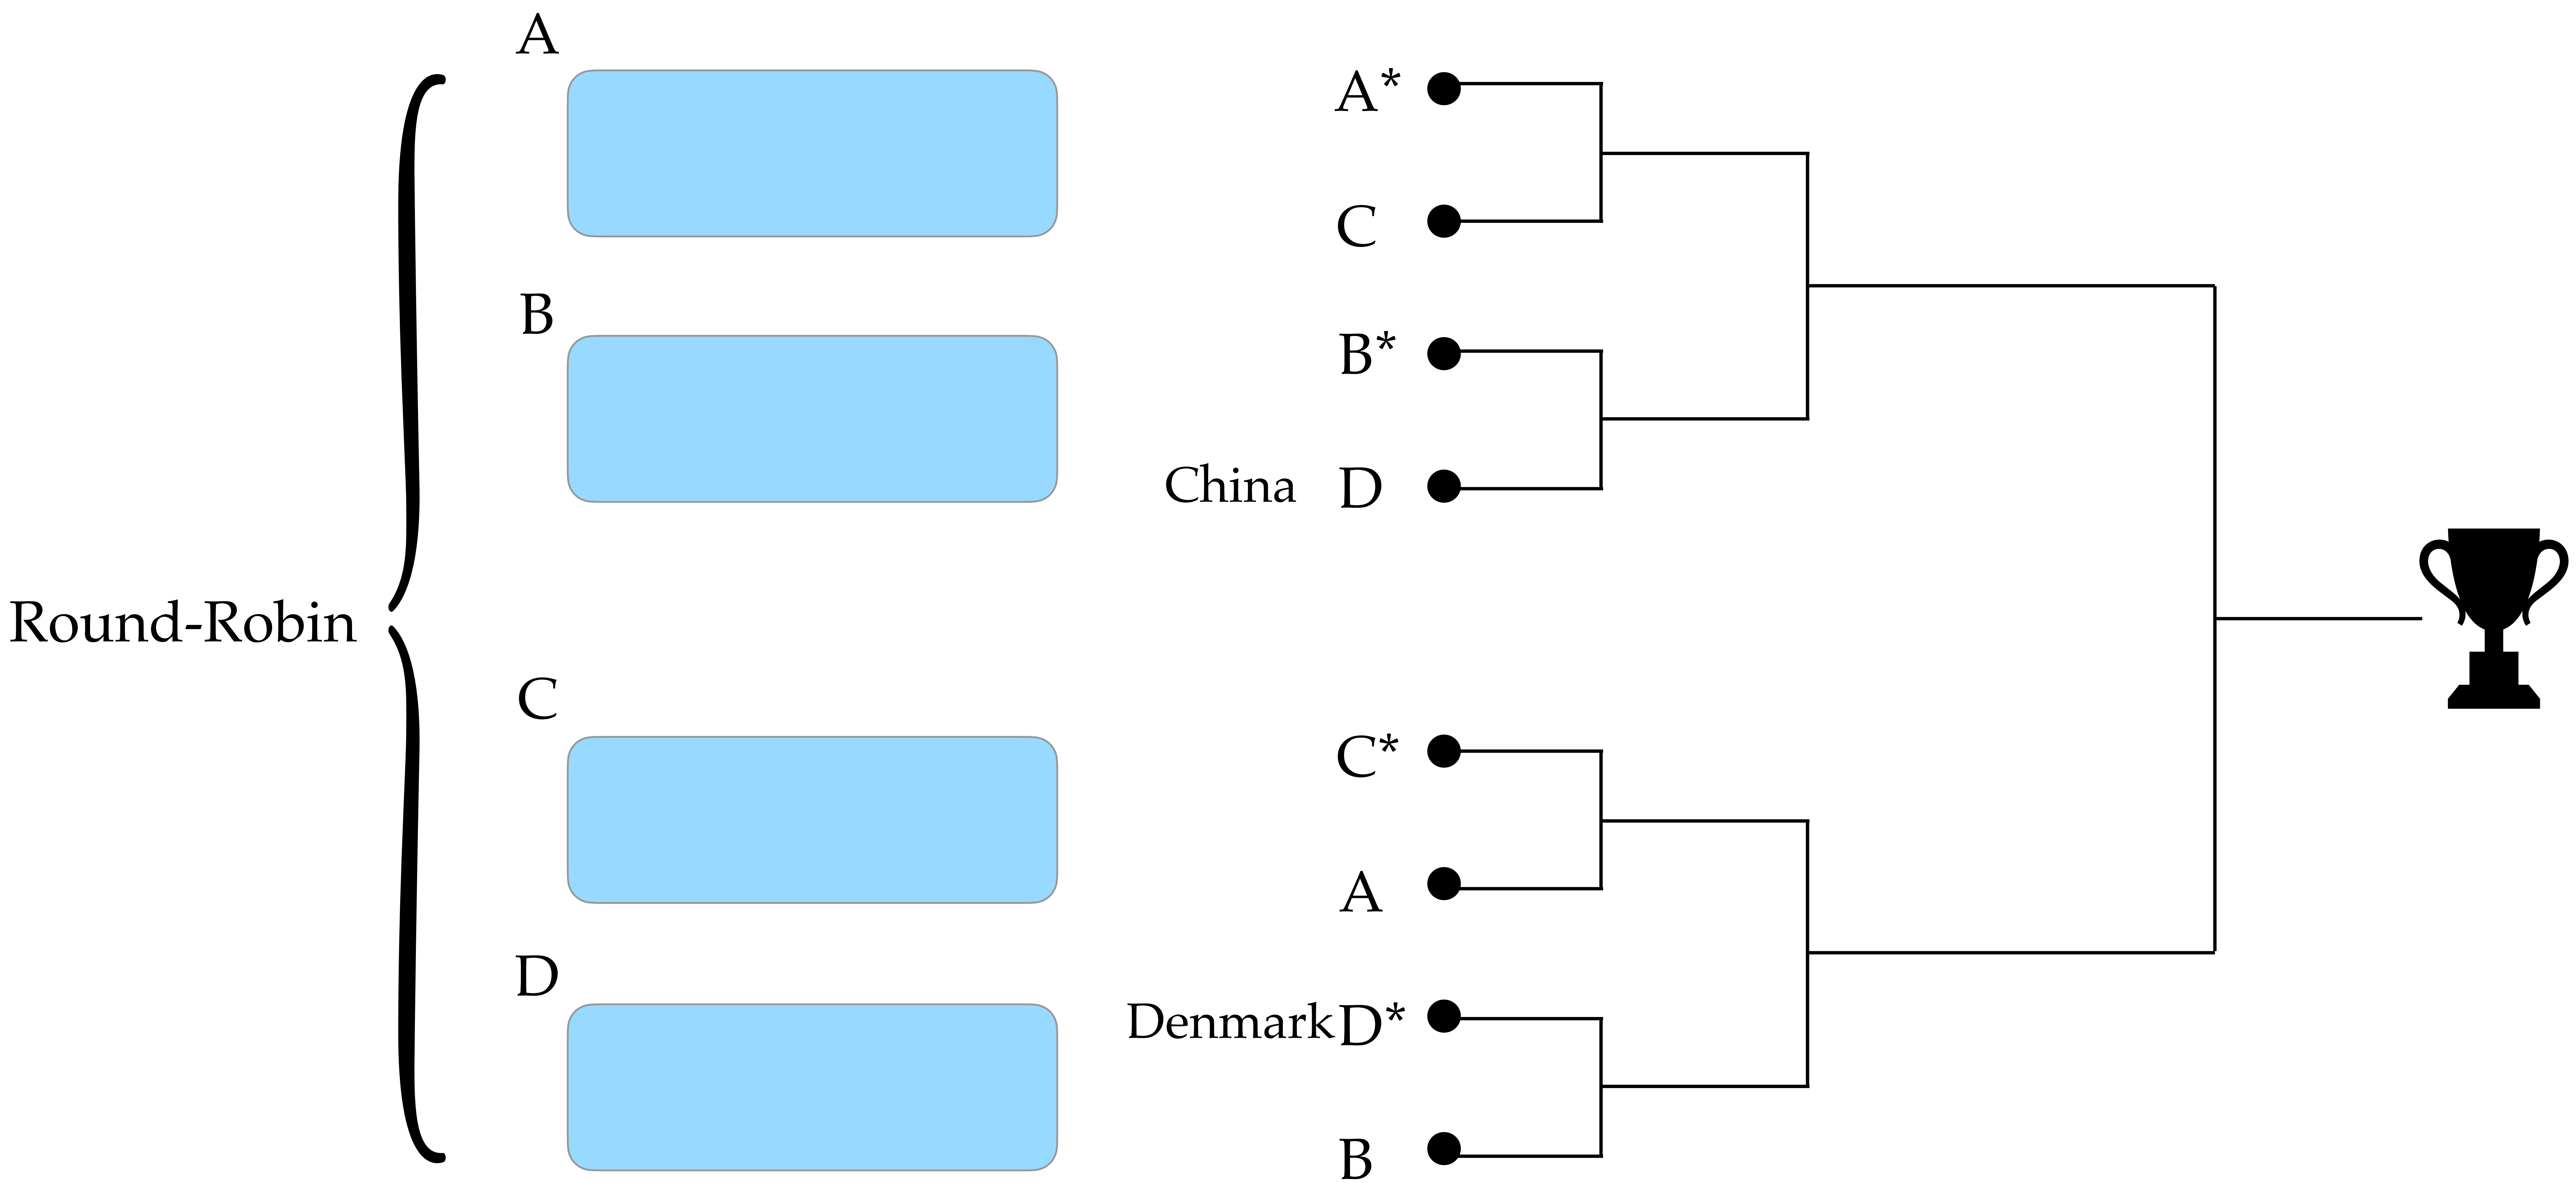

Example: London Olympics

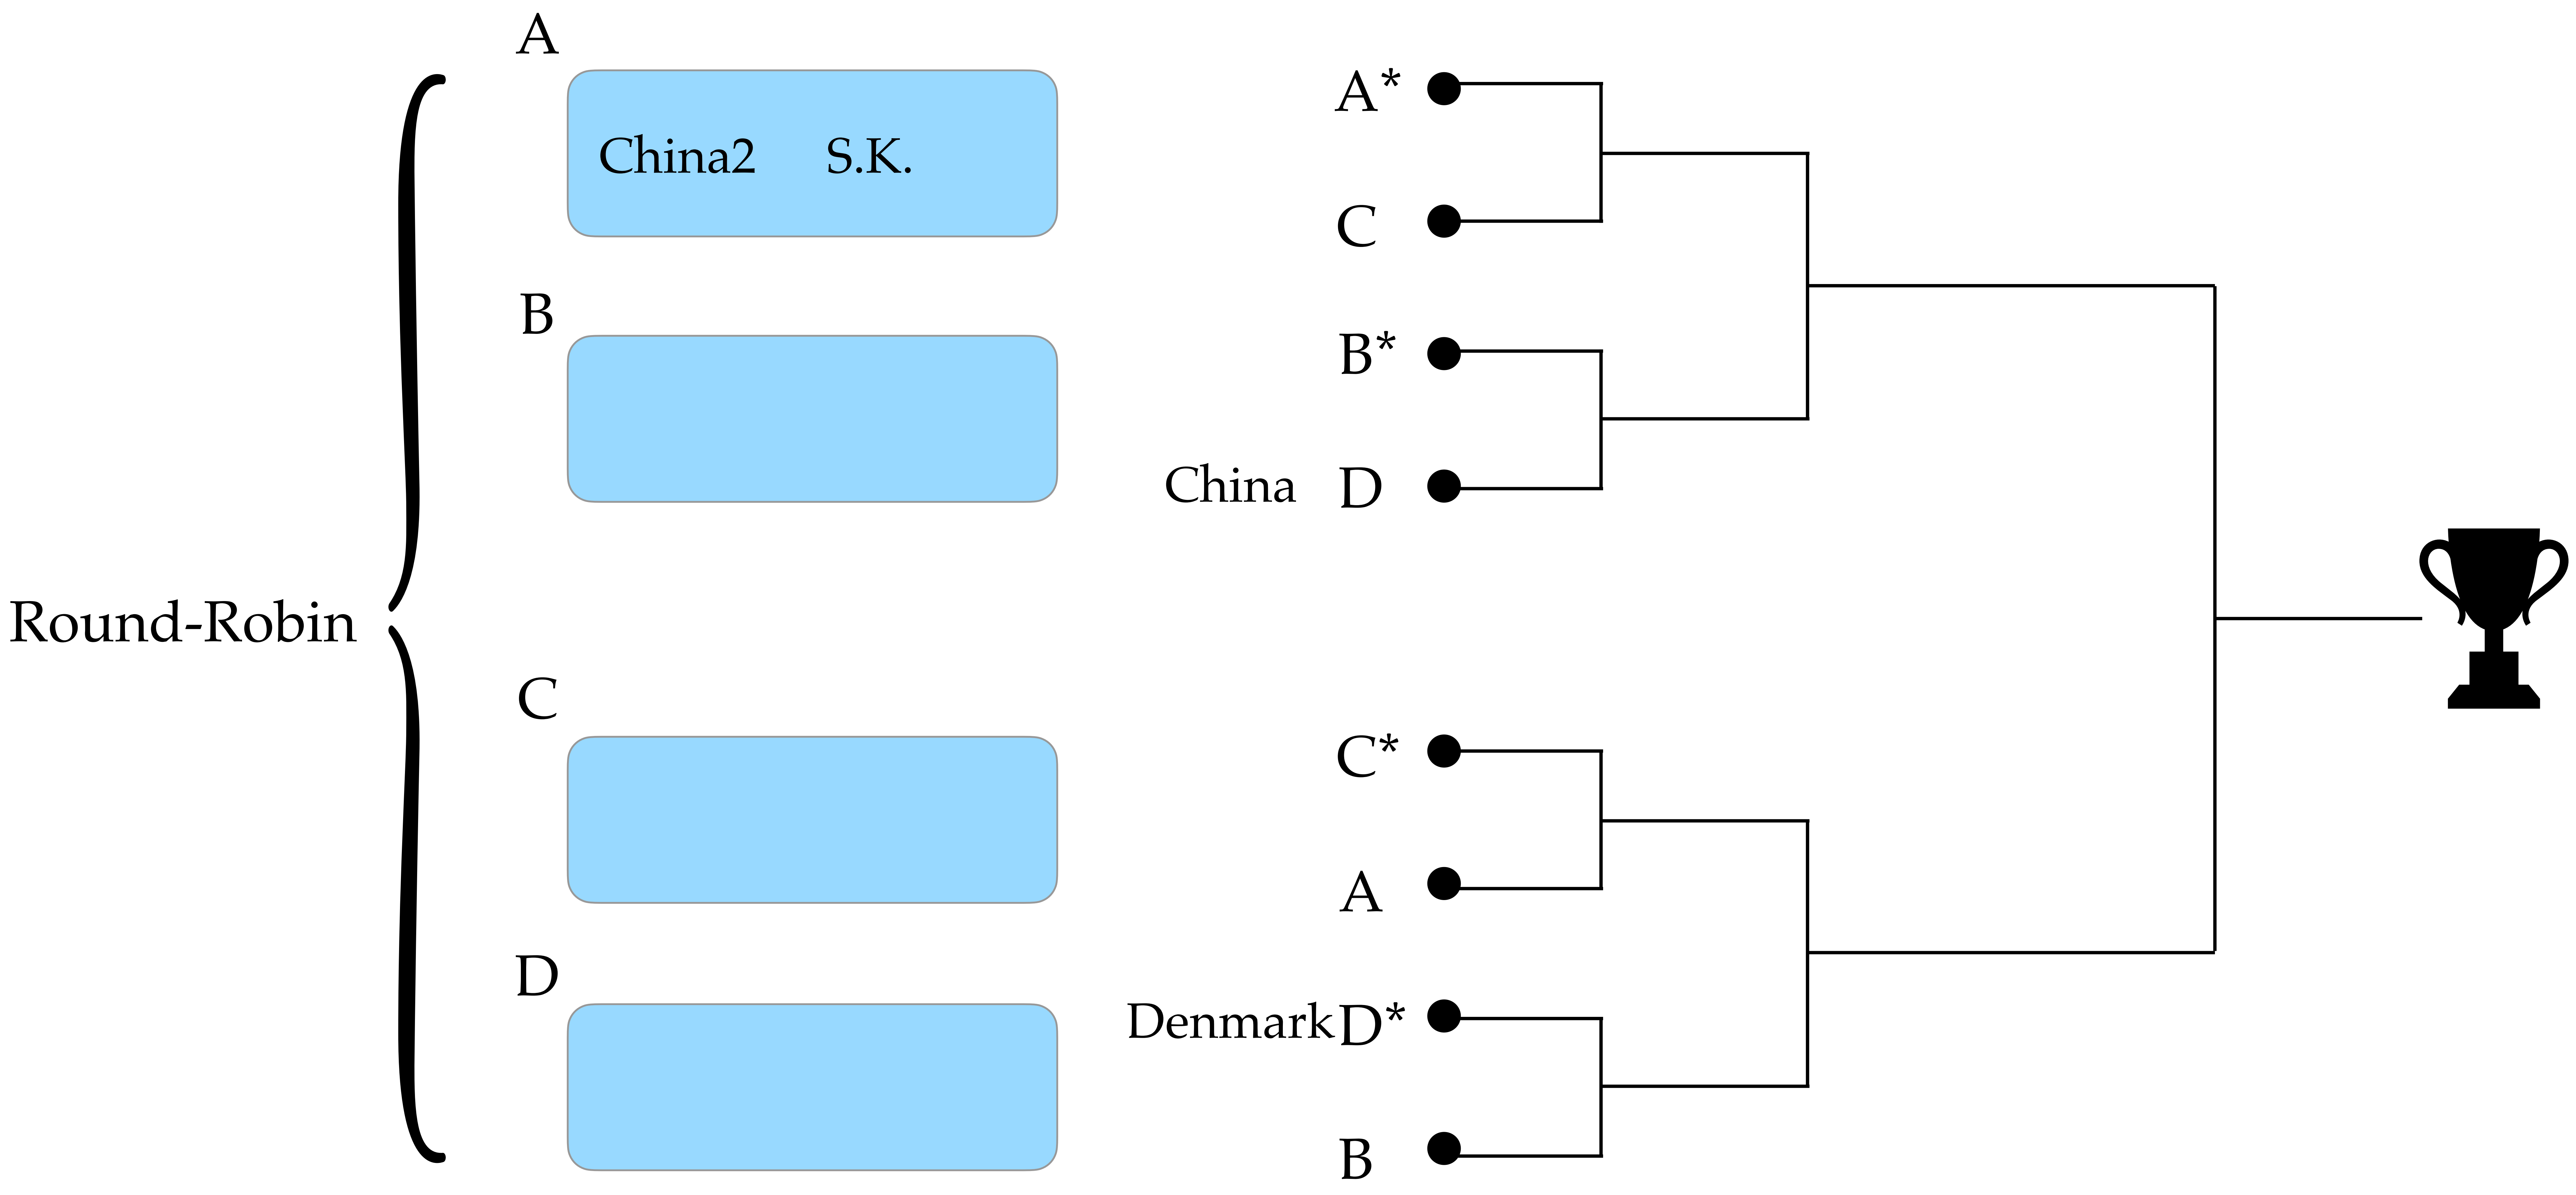

Example: London Olympics

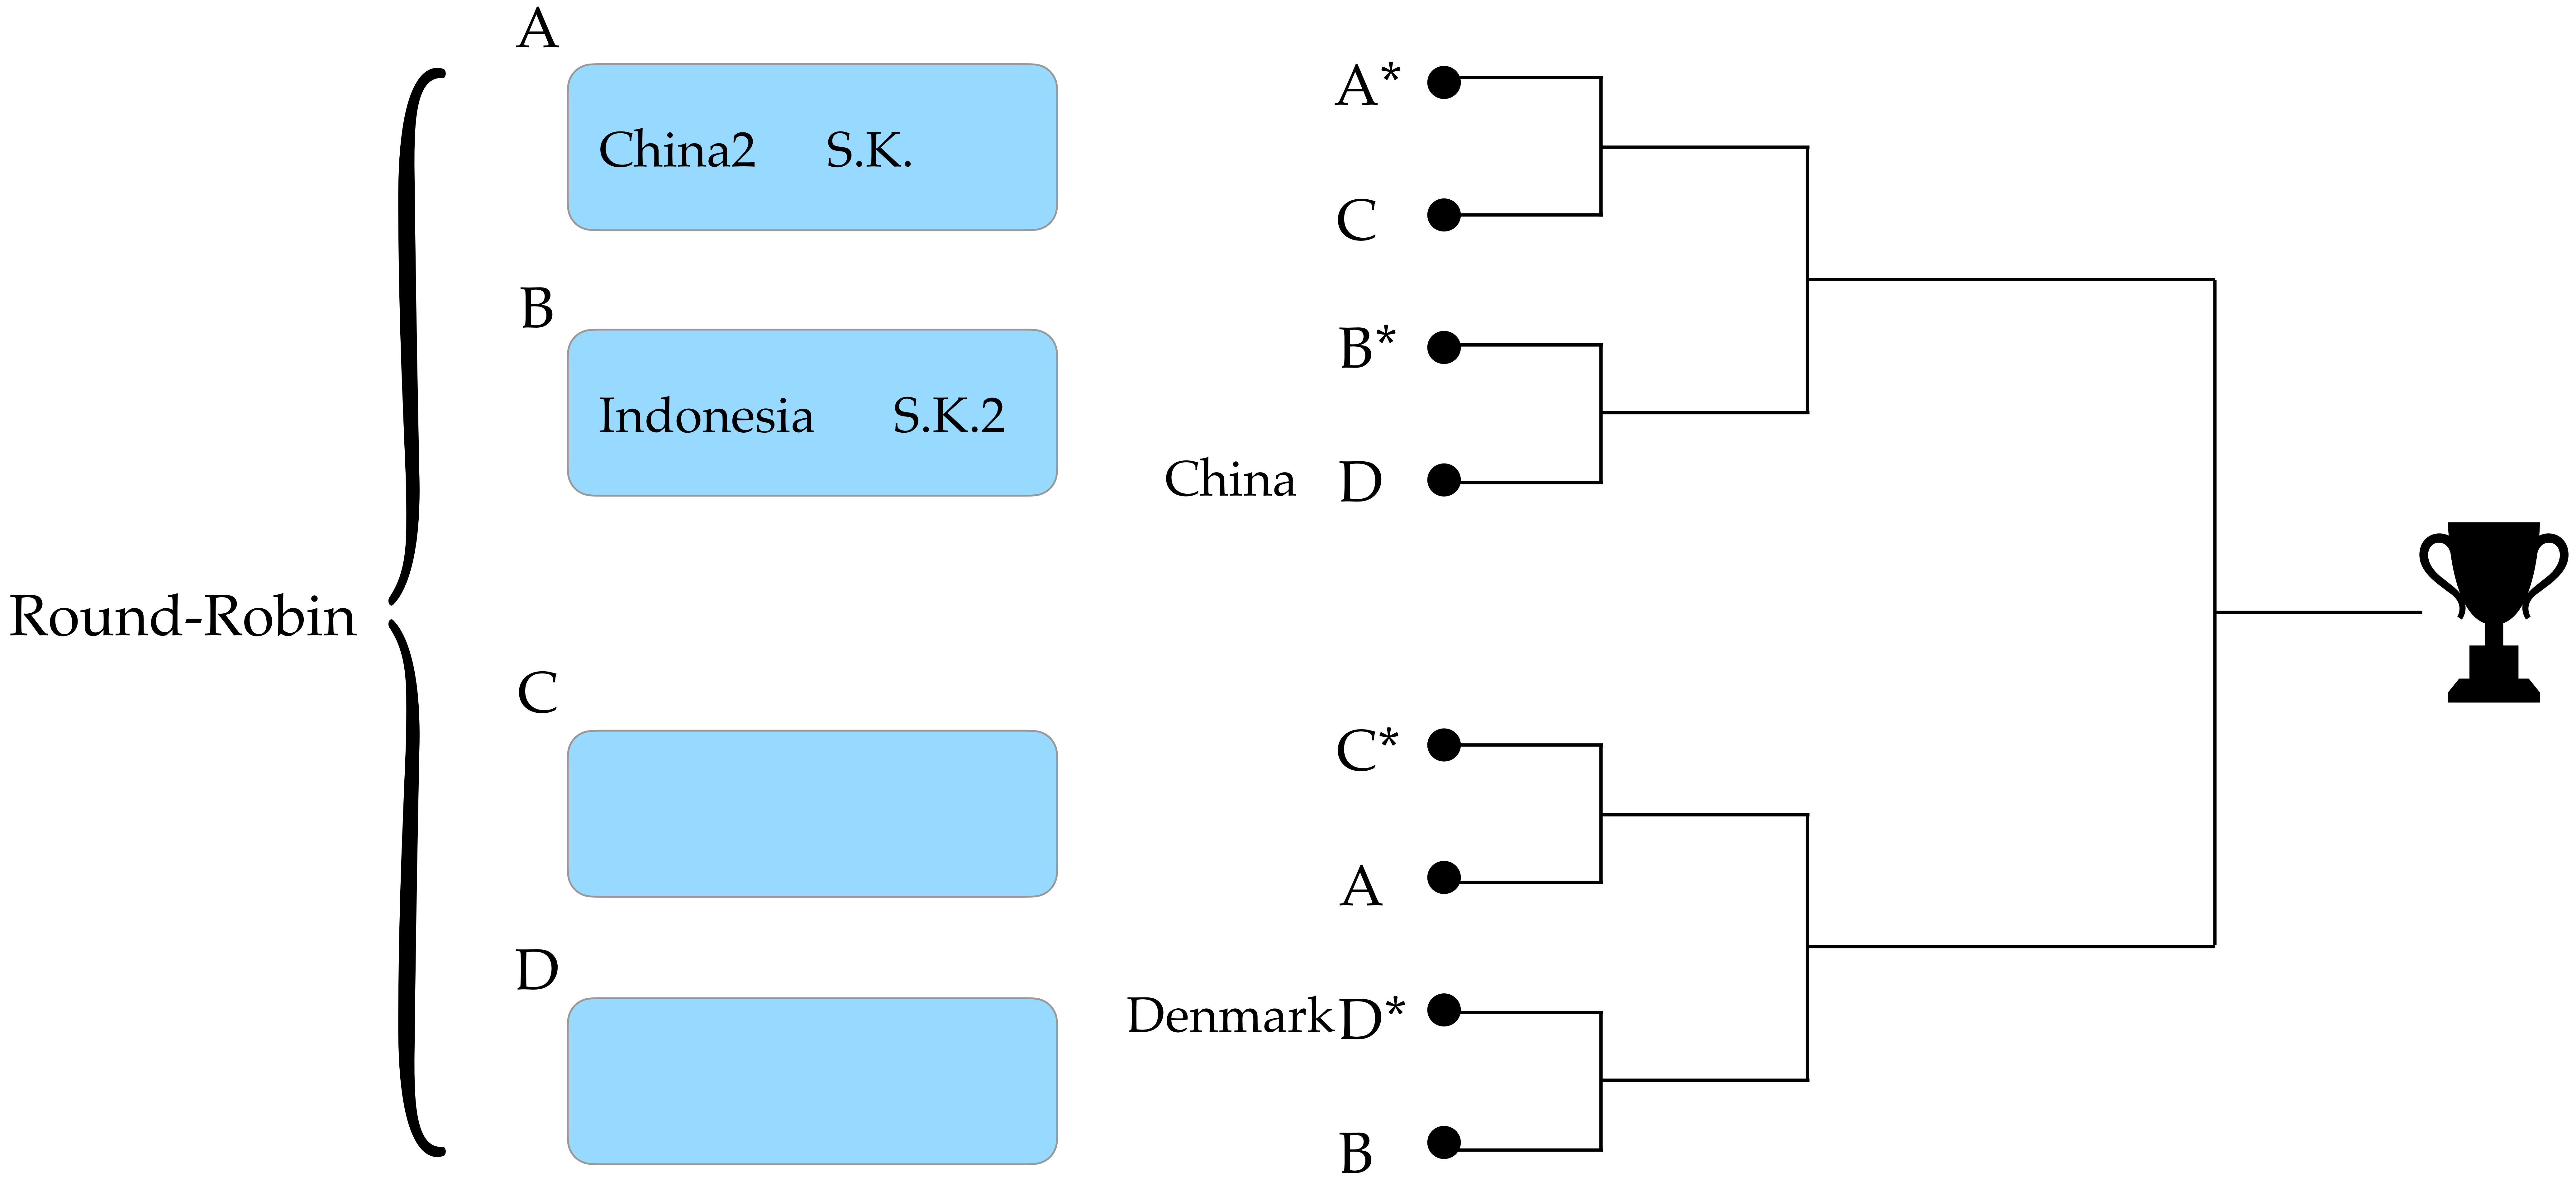

Example: London Olympics

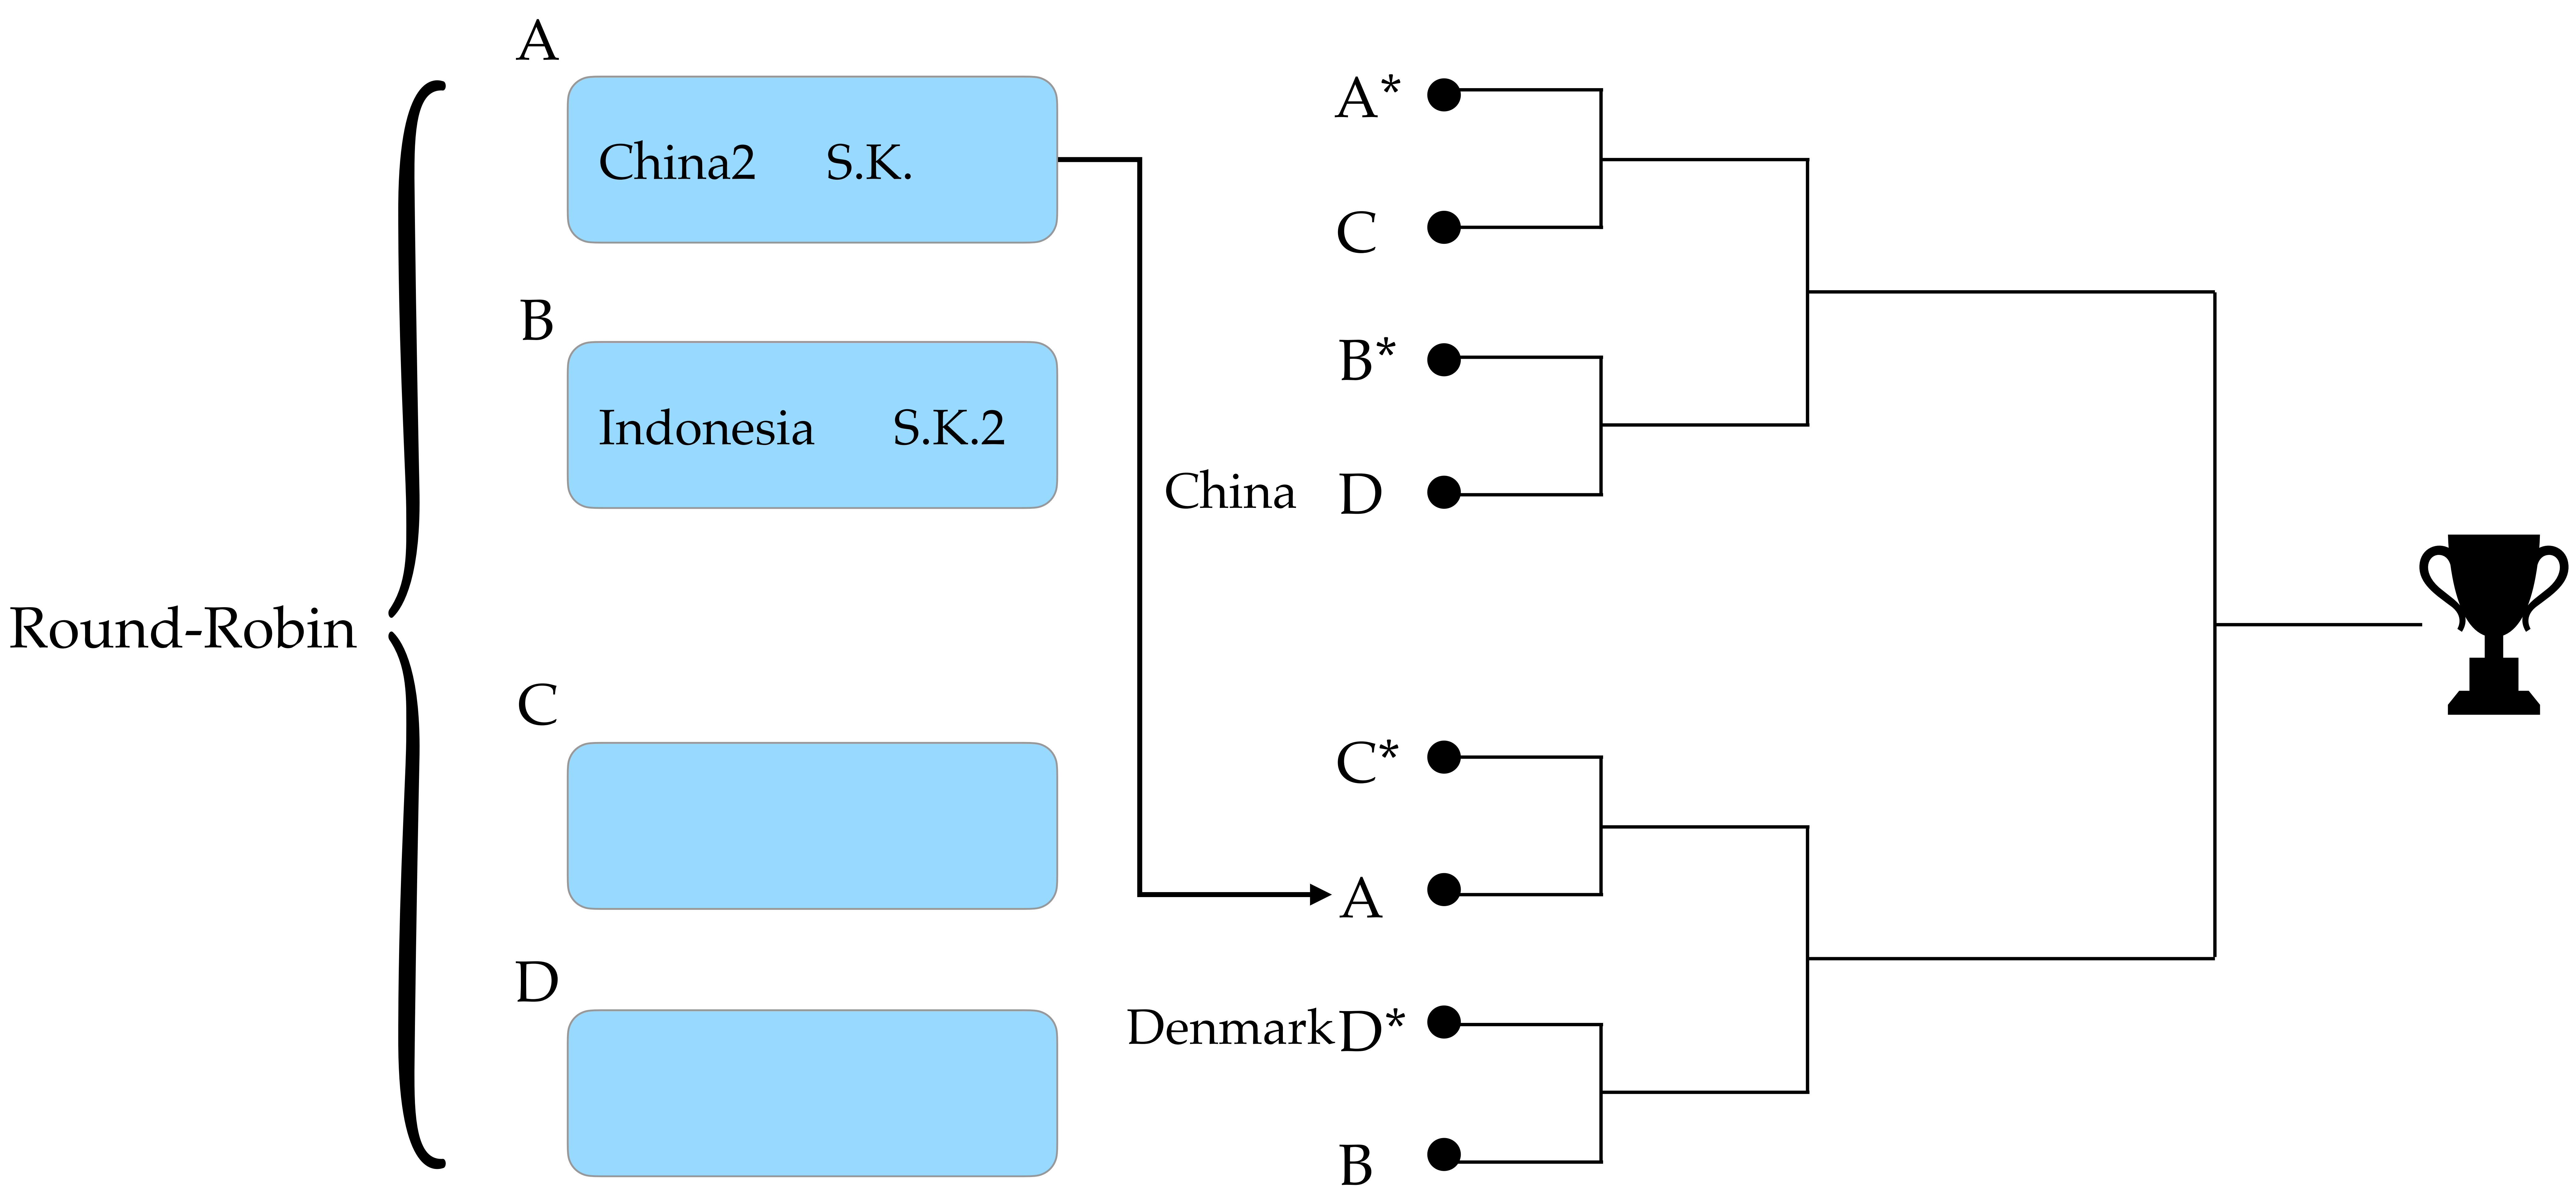

Example: London Olympics

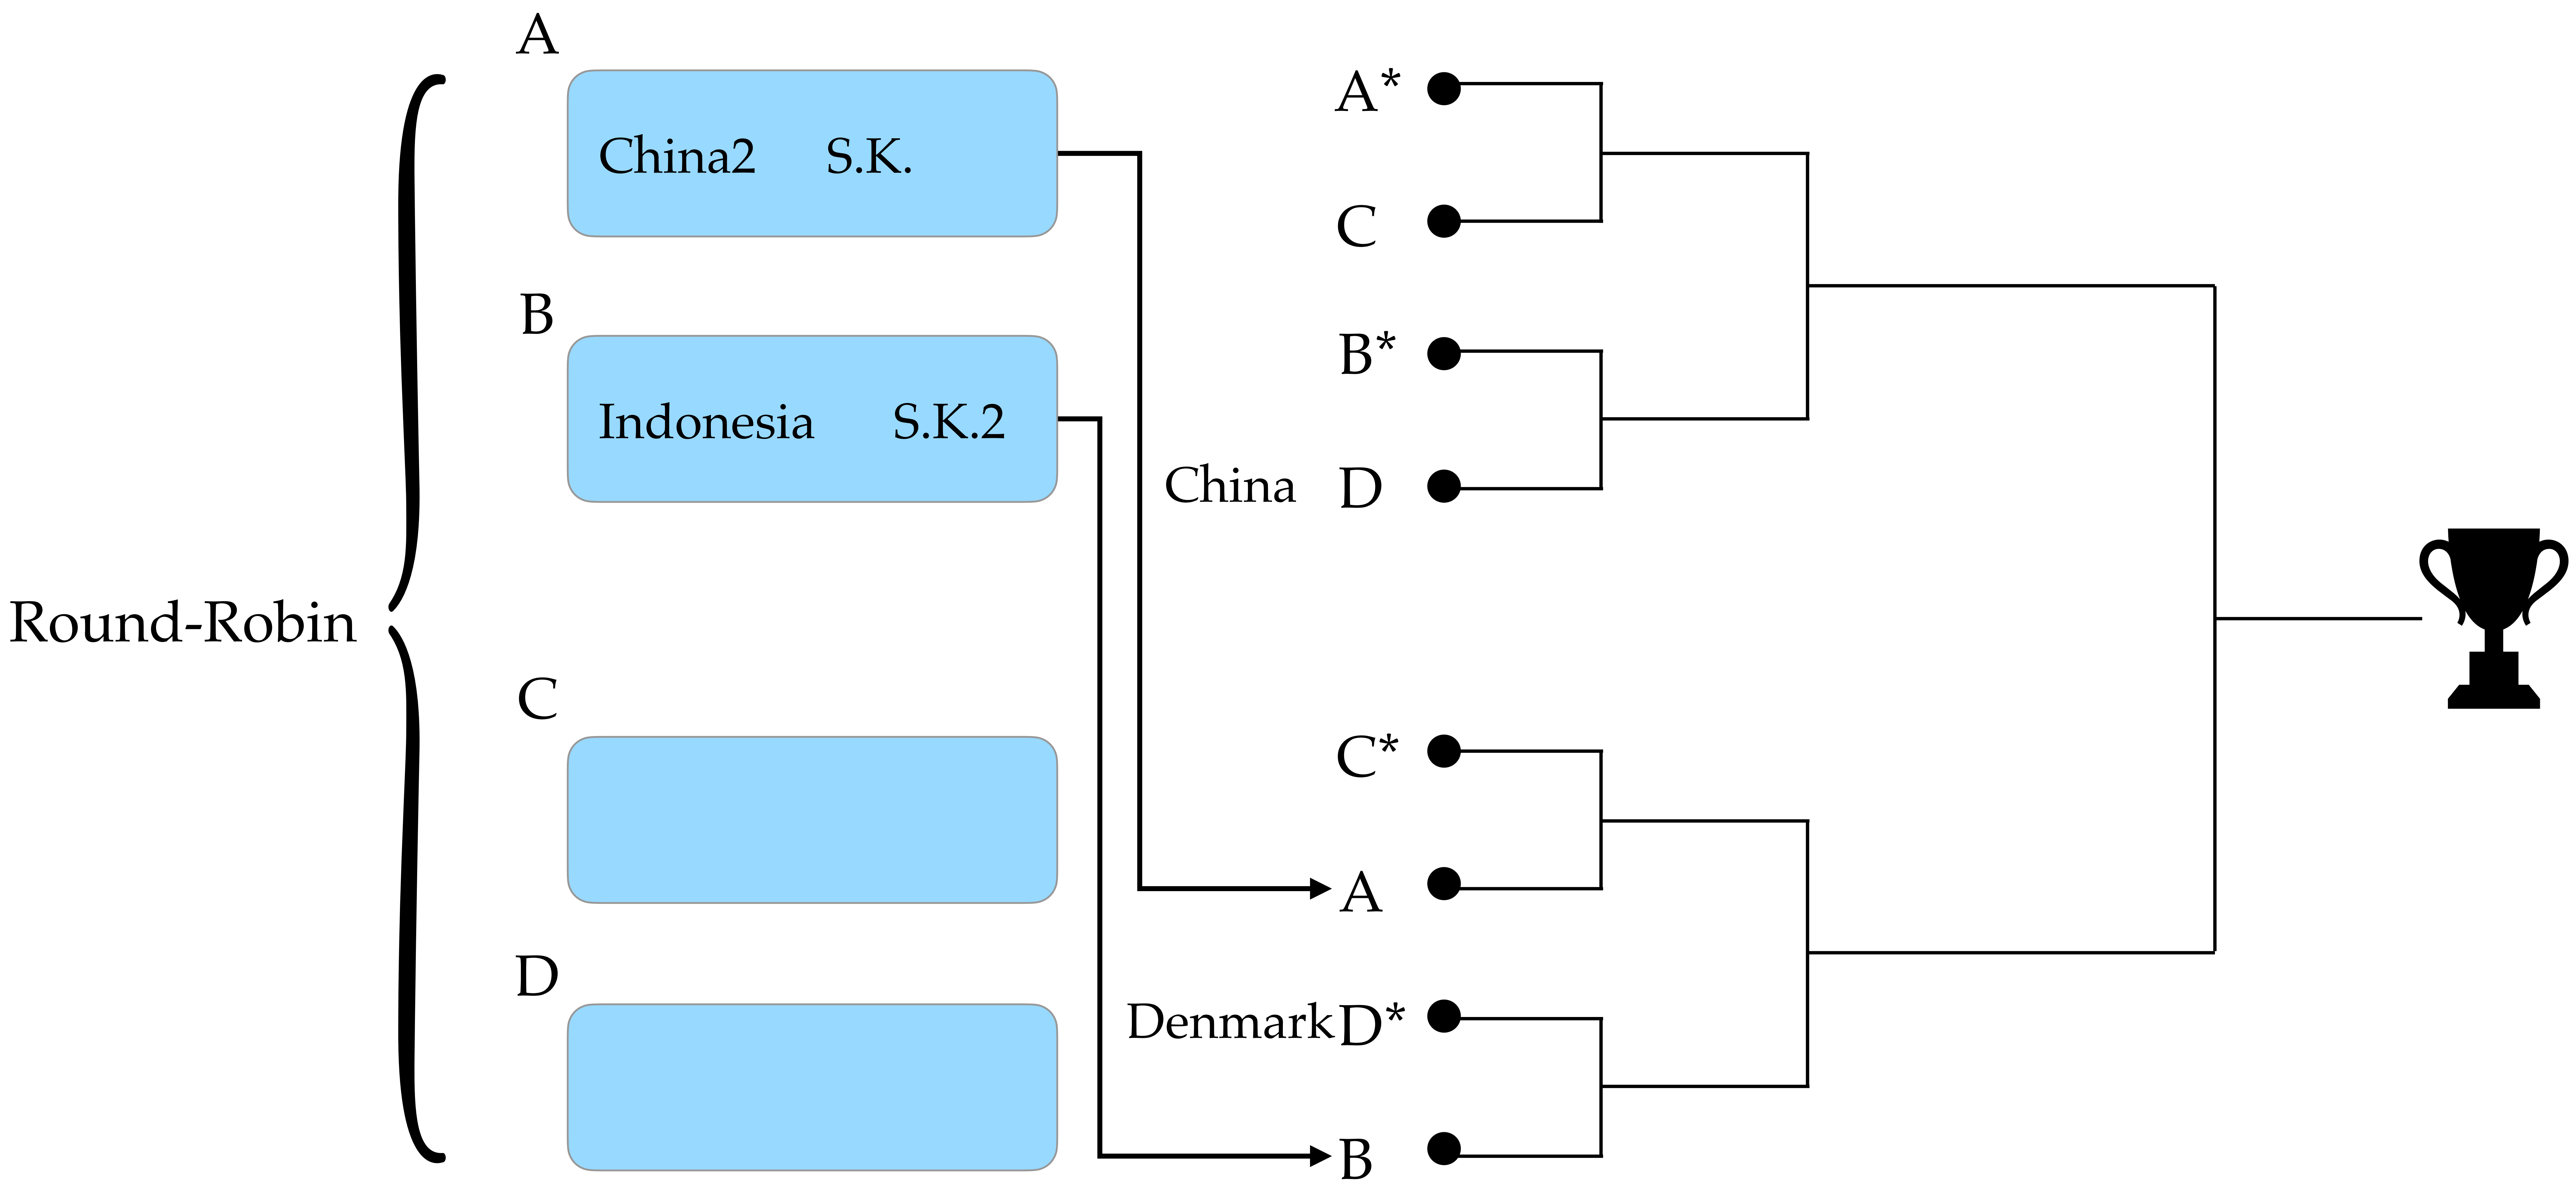

Example: London Olympics

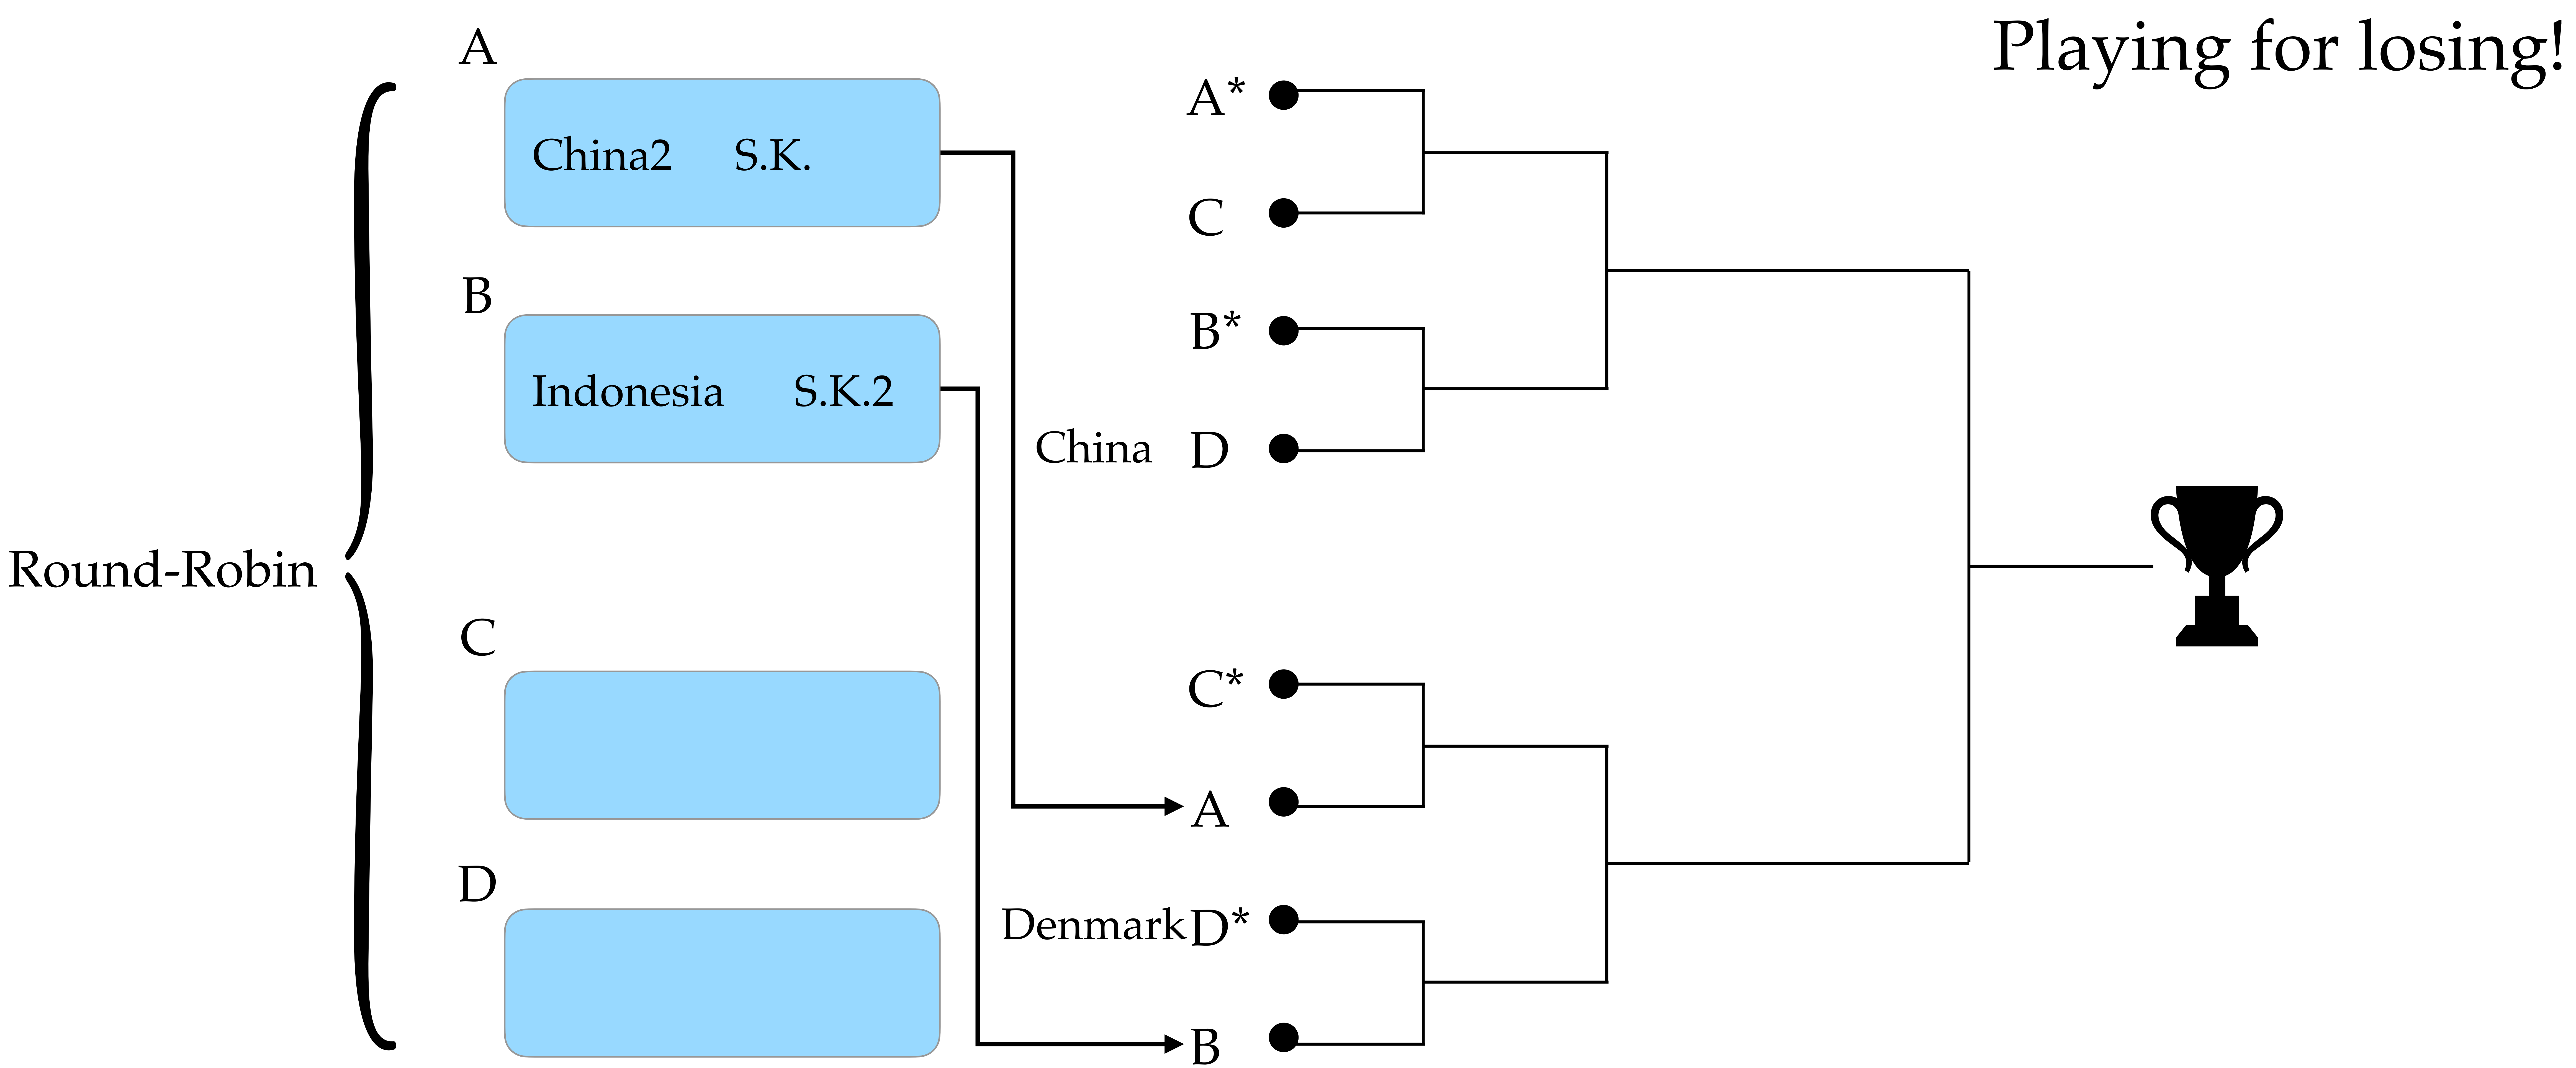

What went wrong?

What went wrong?

- In systems with strategic participants, the rules matter.

What went wrong?

- In systems with strategic participants, the rules matter.

- Poorly designed systems suffer from unexpected and undesirable results.

What went wrong?

- In systems with strategic participants, the rules matter.
- Poorly designed systems suffer from unexpected and undesirable results.
- The burden lies on the system designer to anticipate strategic behaviour, not on the participants to behave against their own interests.

What went wrong?

- In systems with strategic participants, the rules matter.
- Poorly designed systems suffer from unexpected and undesirable results.
- The burden lies on the system designer to anticipate strategic behaviour, not on the participants to behave against their own interests.

What's wrong with the people?

What went wrong?

- In systems with strategic participants, the rules matter.
- Poorly designed systems suffer from unexpected and undesirable results.
- The burden lies on the system designer to anticipate strategic behaviour, not on the participants to behave against their own interests.

~~What's wrong with the people?~~

What went wrong?

- In systems with strategic participants, the rules matter.
- Poorly designed systems suffer from unexpected and undesirable results.
- The burden lies on the system designer to anticipate strategic behaviour, not on the participants to behave against their own interests.

~~What's wrong with the people?~~

What's wrong with the system?

The Science of Rule Making

The Science of Rule Making

Mechanism Design: Design rules so that strategic behaviour by the participants leads to desirable outcomes.

The Science of Rule Making

Mechanism Design: Design rules so that strategic behaviour by the participants leads to desirable outcomes.

Applications:

The Science of Rule Making

Mechanism Design: Design rules so that strategic behaviour by the participants leads to desirable outcomes.

Applications:

- Internet search auctions

The Science of Rule Making

Mechanism Design: Design rules so that strategic behaviour by the participants leads to desirable outcomes.

Applications:

- Internet search auctions
- Wireless spectrum auctions

The Science of Rule Making

Mechanism Design: Design rules so that strategic behaviour by the participants leads to desirable outcomes.

Applications:

- Internet search auctions
- Wireless spectrum auctions
- Matching medical residents to hospitals

The Science of Rule Making

Mechanism Design: Design rules so that strategic behaviour by the participants leads to desirable outcomes.

Applications:

- Internet search auctions
- Wireless spectrum auctions
- Matching medical residents to hospitals
- Matching children to schools

The Science of Rule Making

Mechanism Design: Design rules so that strategic behaviour by the participants leads to desirable outcomes.

Applications:

- Internet search auctions
- Wireless spectrum auctions
- Matching medical residents to hospitals
- Matching children to schools
- Kidney exchange markets

Course Goal 2: When is selfish behaviour essentially benign?

The Price of Anarchy

When Is Selfish Behaviour Near-Optimal?

The Price of Anarchy

When Is Selfish Behaviour Near-Optimal?

- The game is given - the rules are made already

The Price of Anarchy
When Is Selfish Behaviour Near-Optimal?

- The game is given - the rules are made already
    - Internet, road network etc... (that are already built)

The Price of Anarchy

When Is Selfish Behaviour Near-Optimal?

- The game is given - the rules are made already
  - Internet, road network etc… (that are already built)
- What is the cost of selfish behaviour in this game?

The Price of Anarchy
When Is Selfish Behaviour Near-Optimal?

- The game is given - the rules are made already
    - Internet, road network etc… (that are already built)
- What is the cost of selfish behaviour in this game?
    - (Compared to optimal)

Braess‘s Paradox

Braess's Paradox

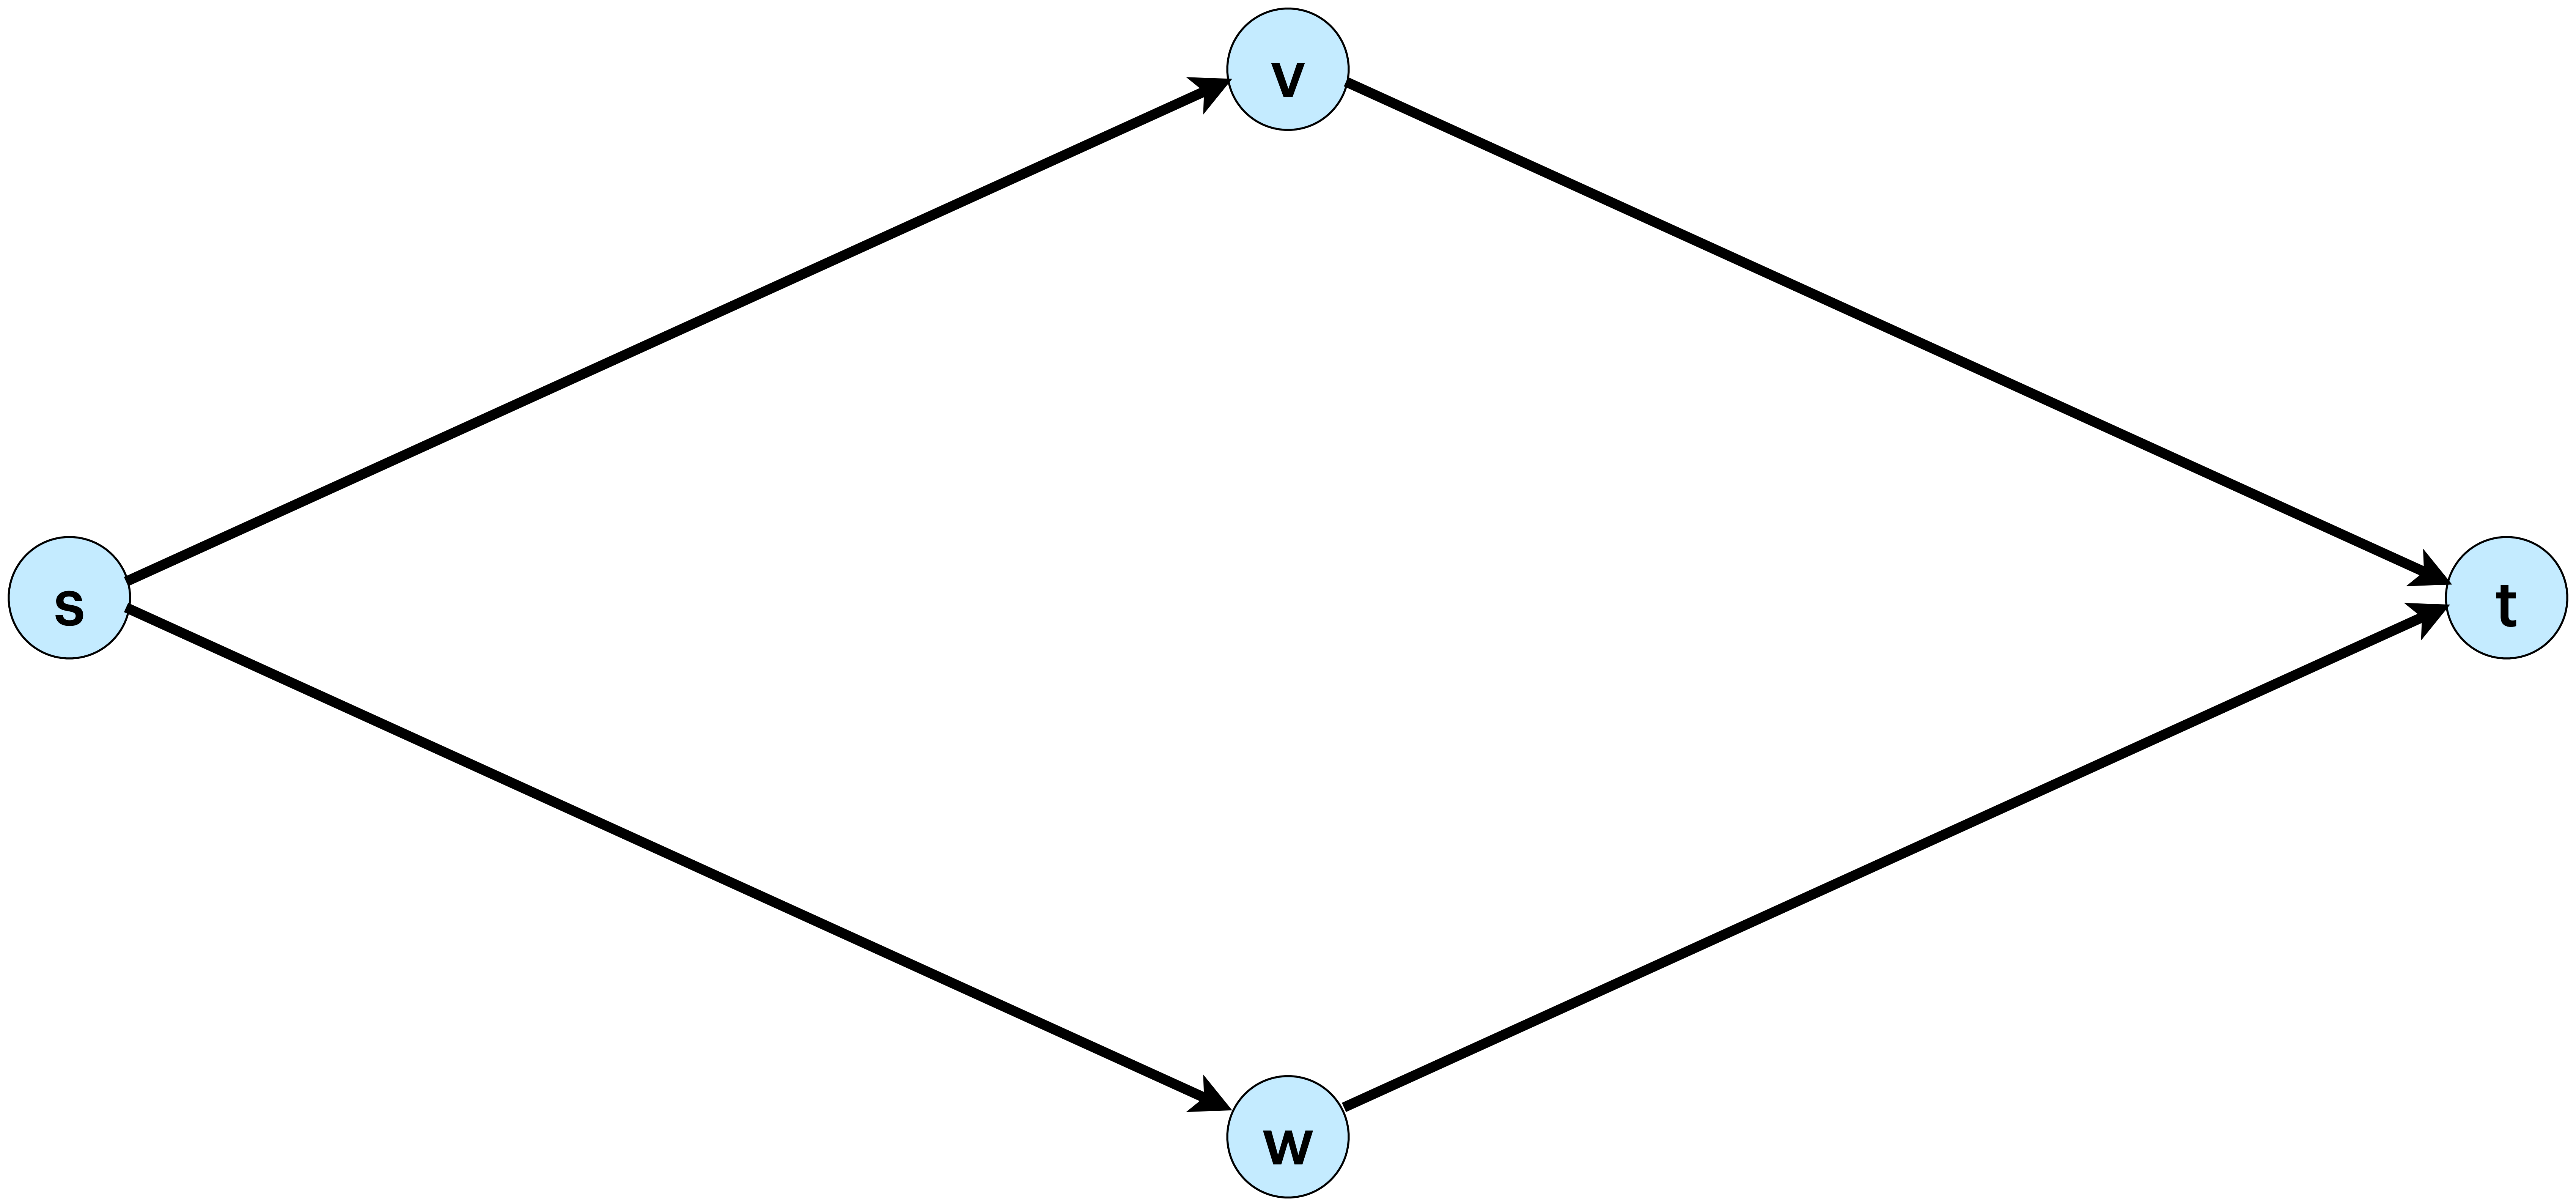

Braess's Paradox

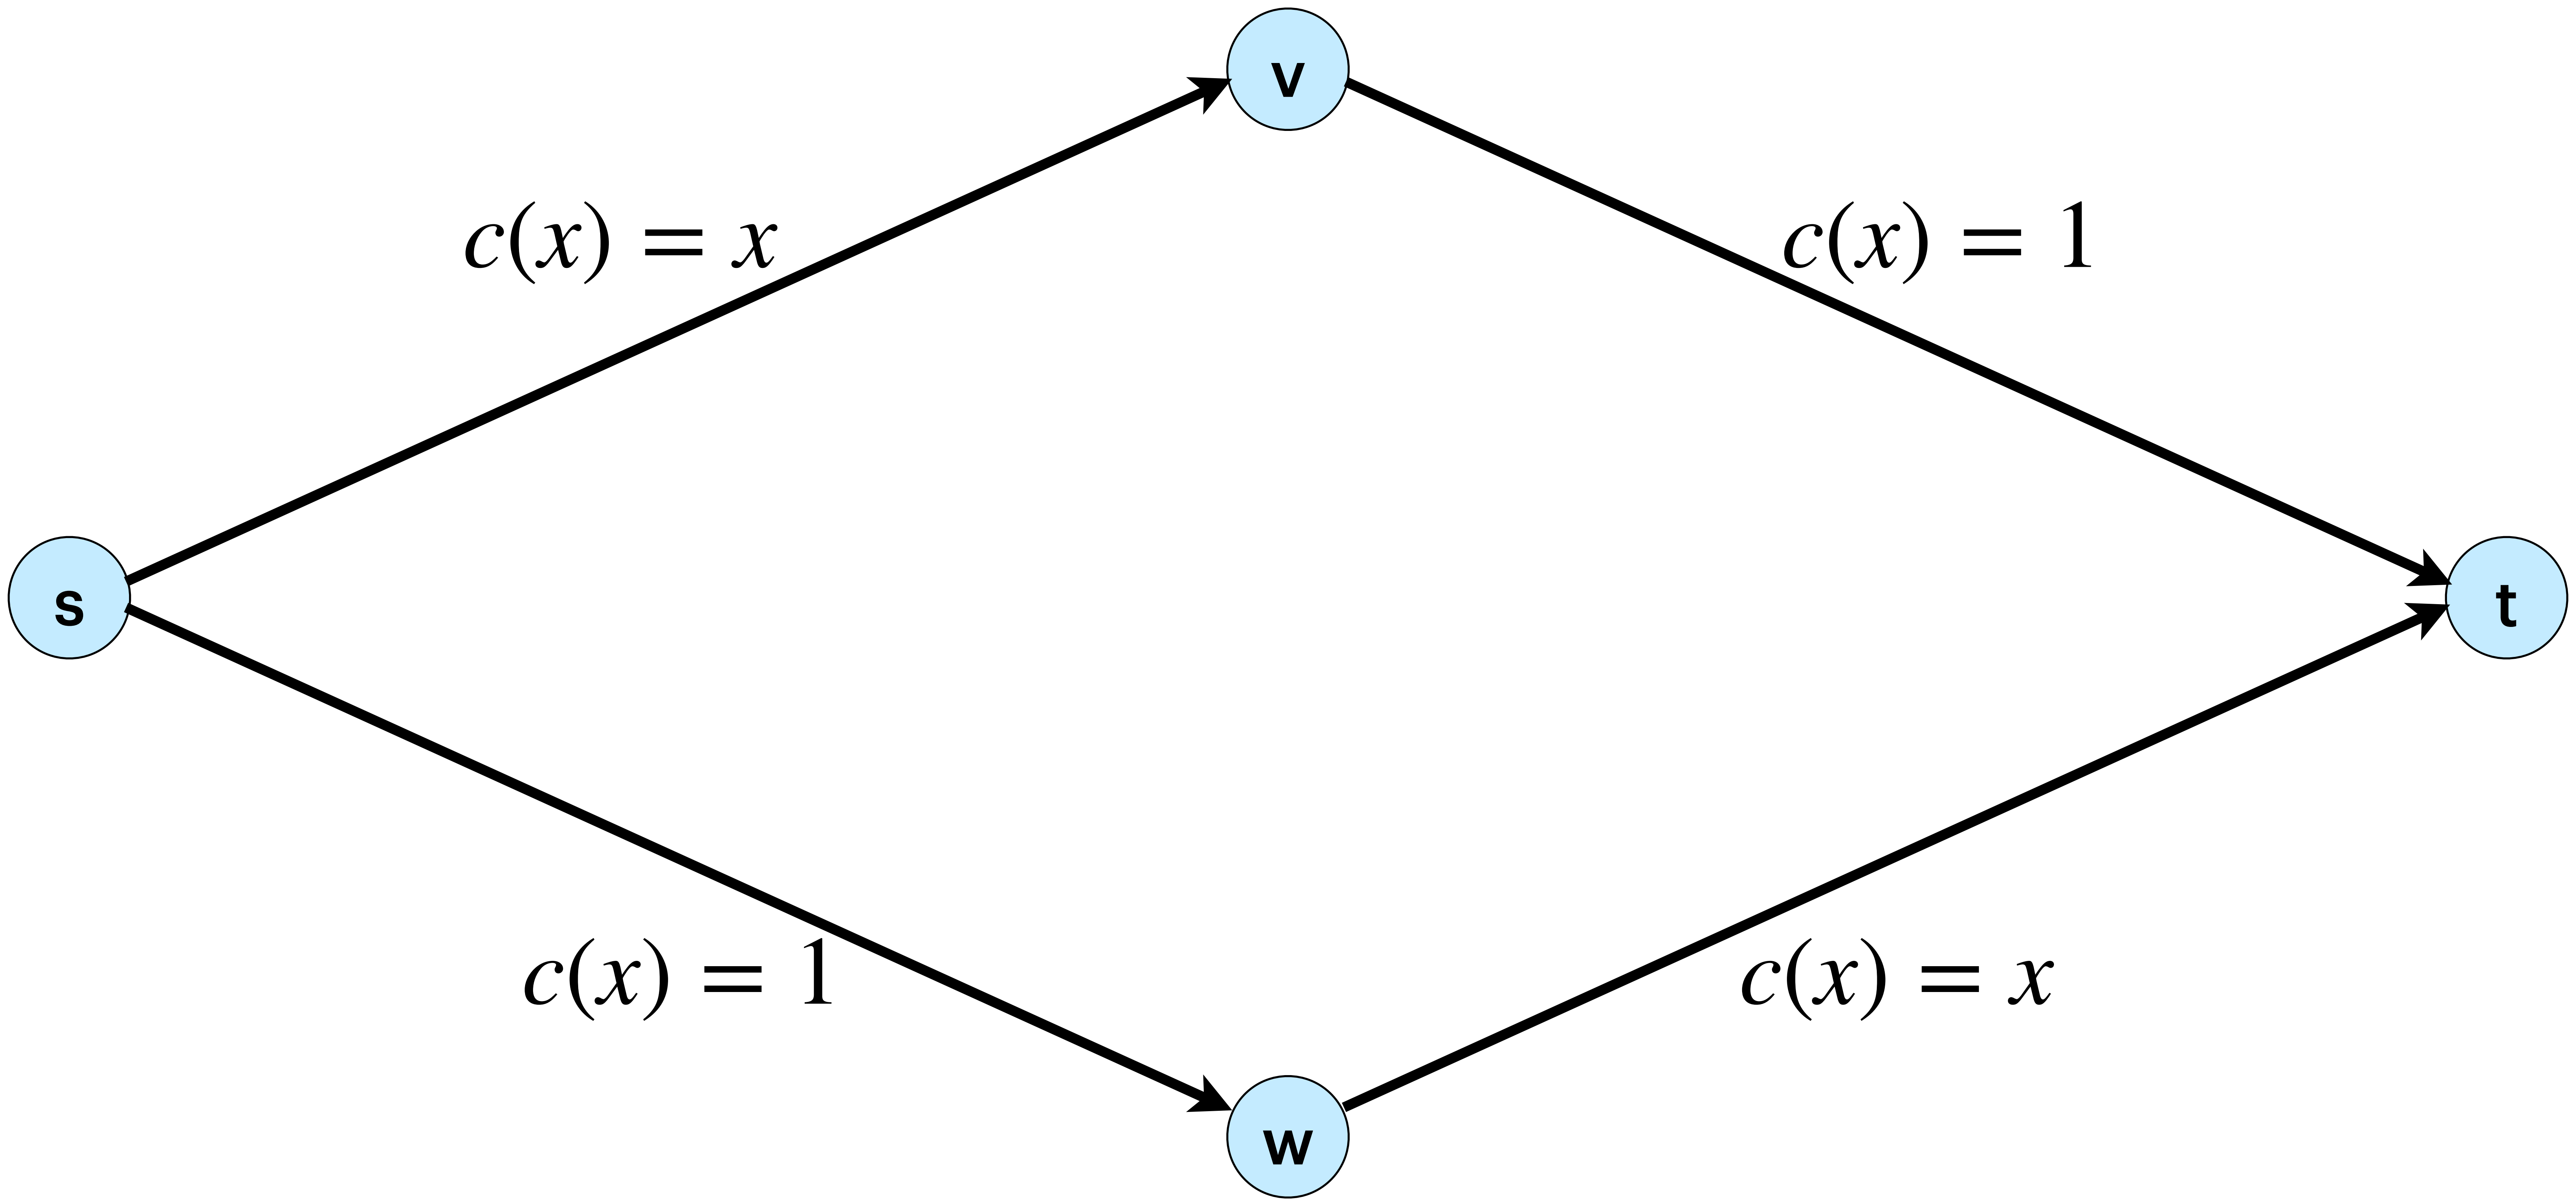

Braess's Paradox

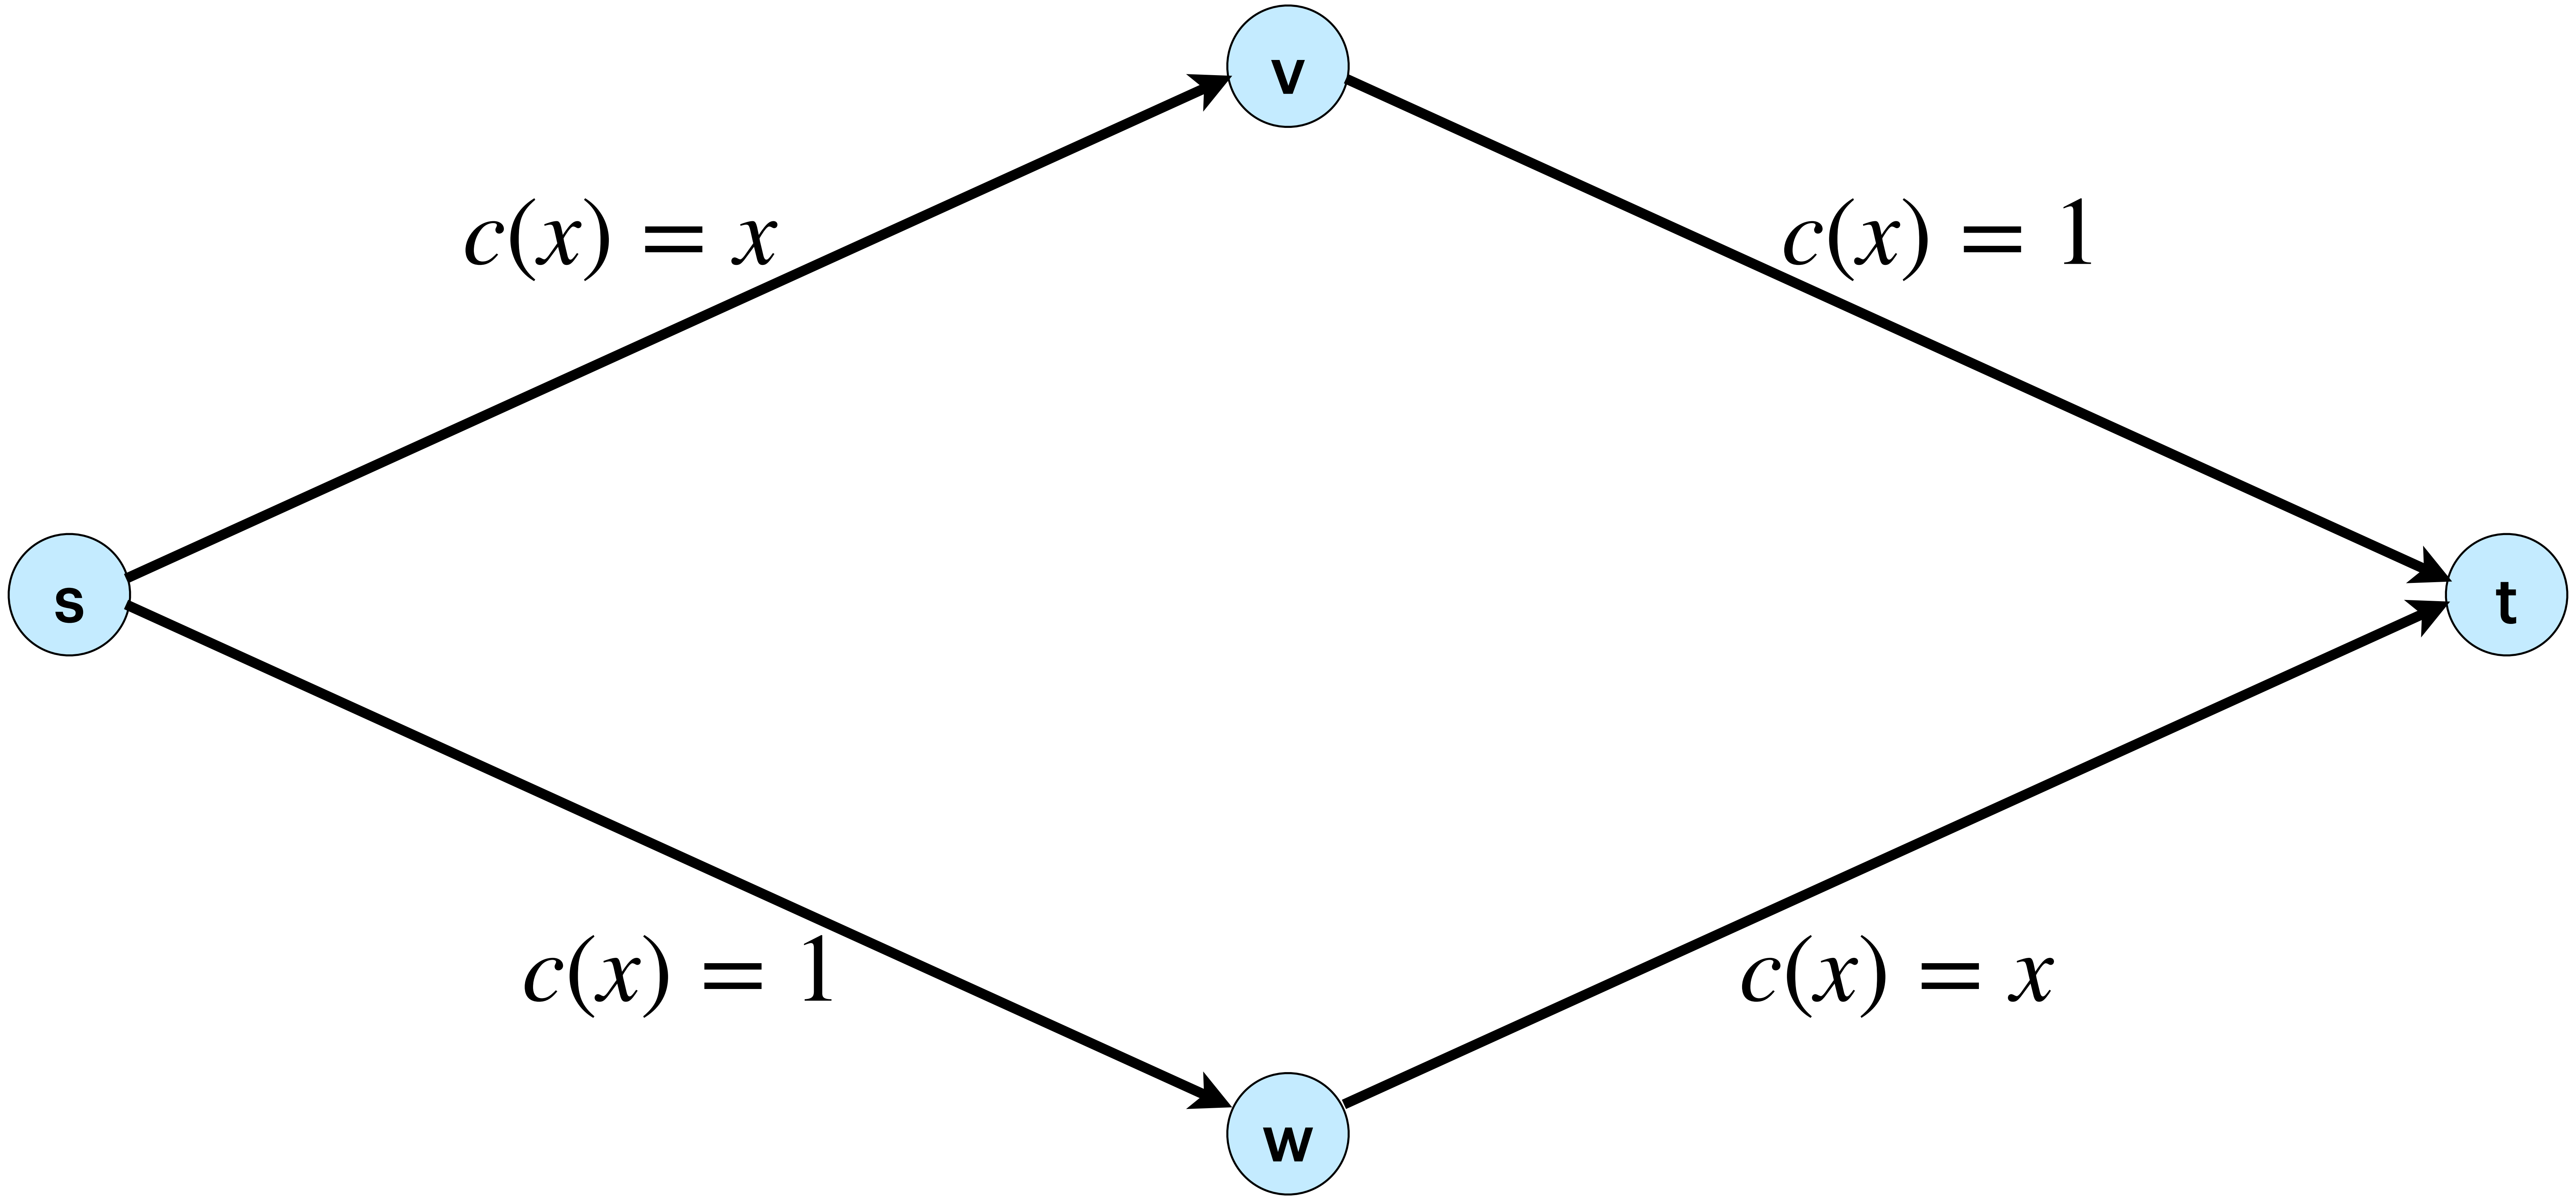

Identical paths, combined travel time: $1 + x$ each

Braess's Paradox

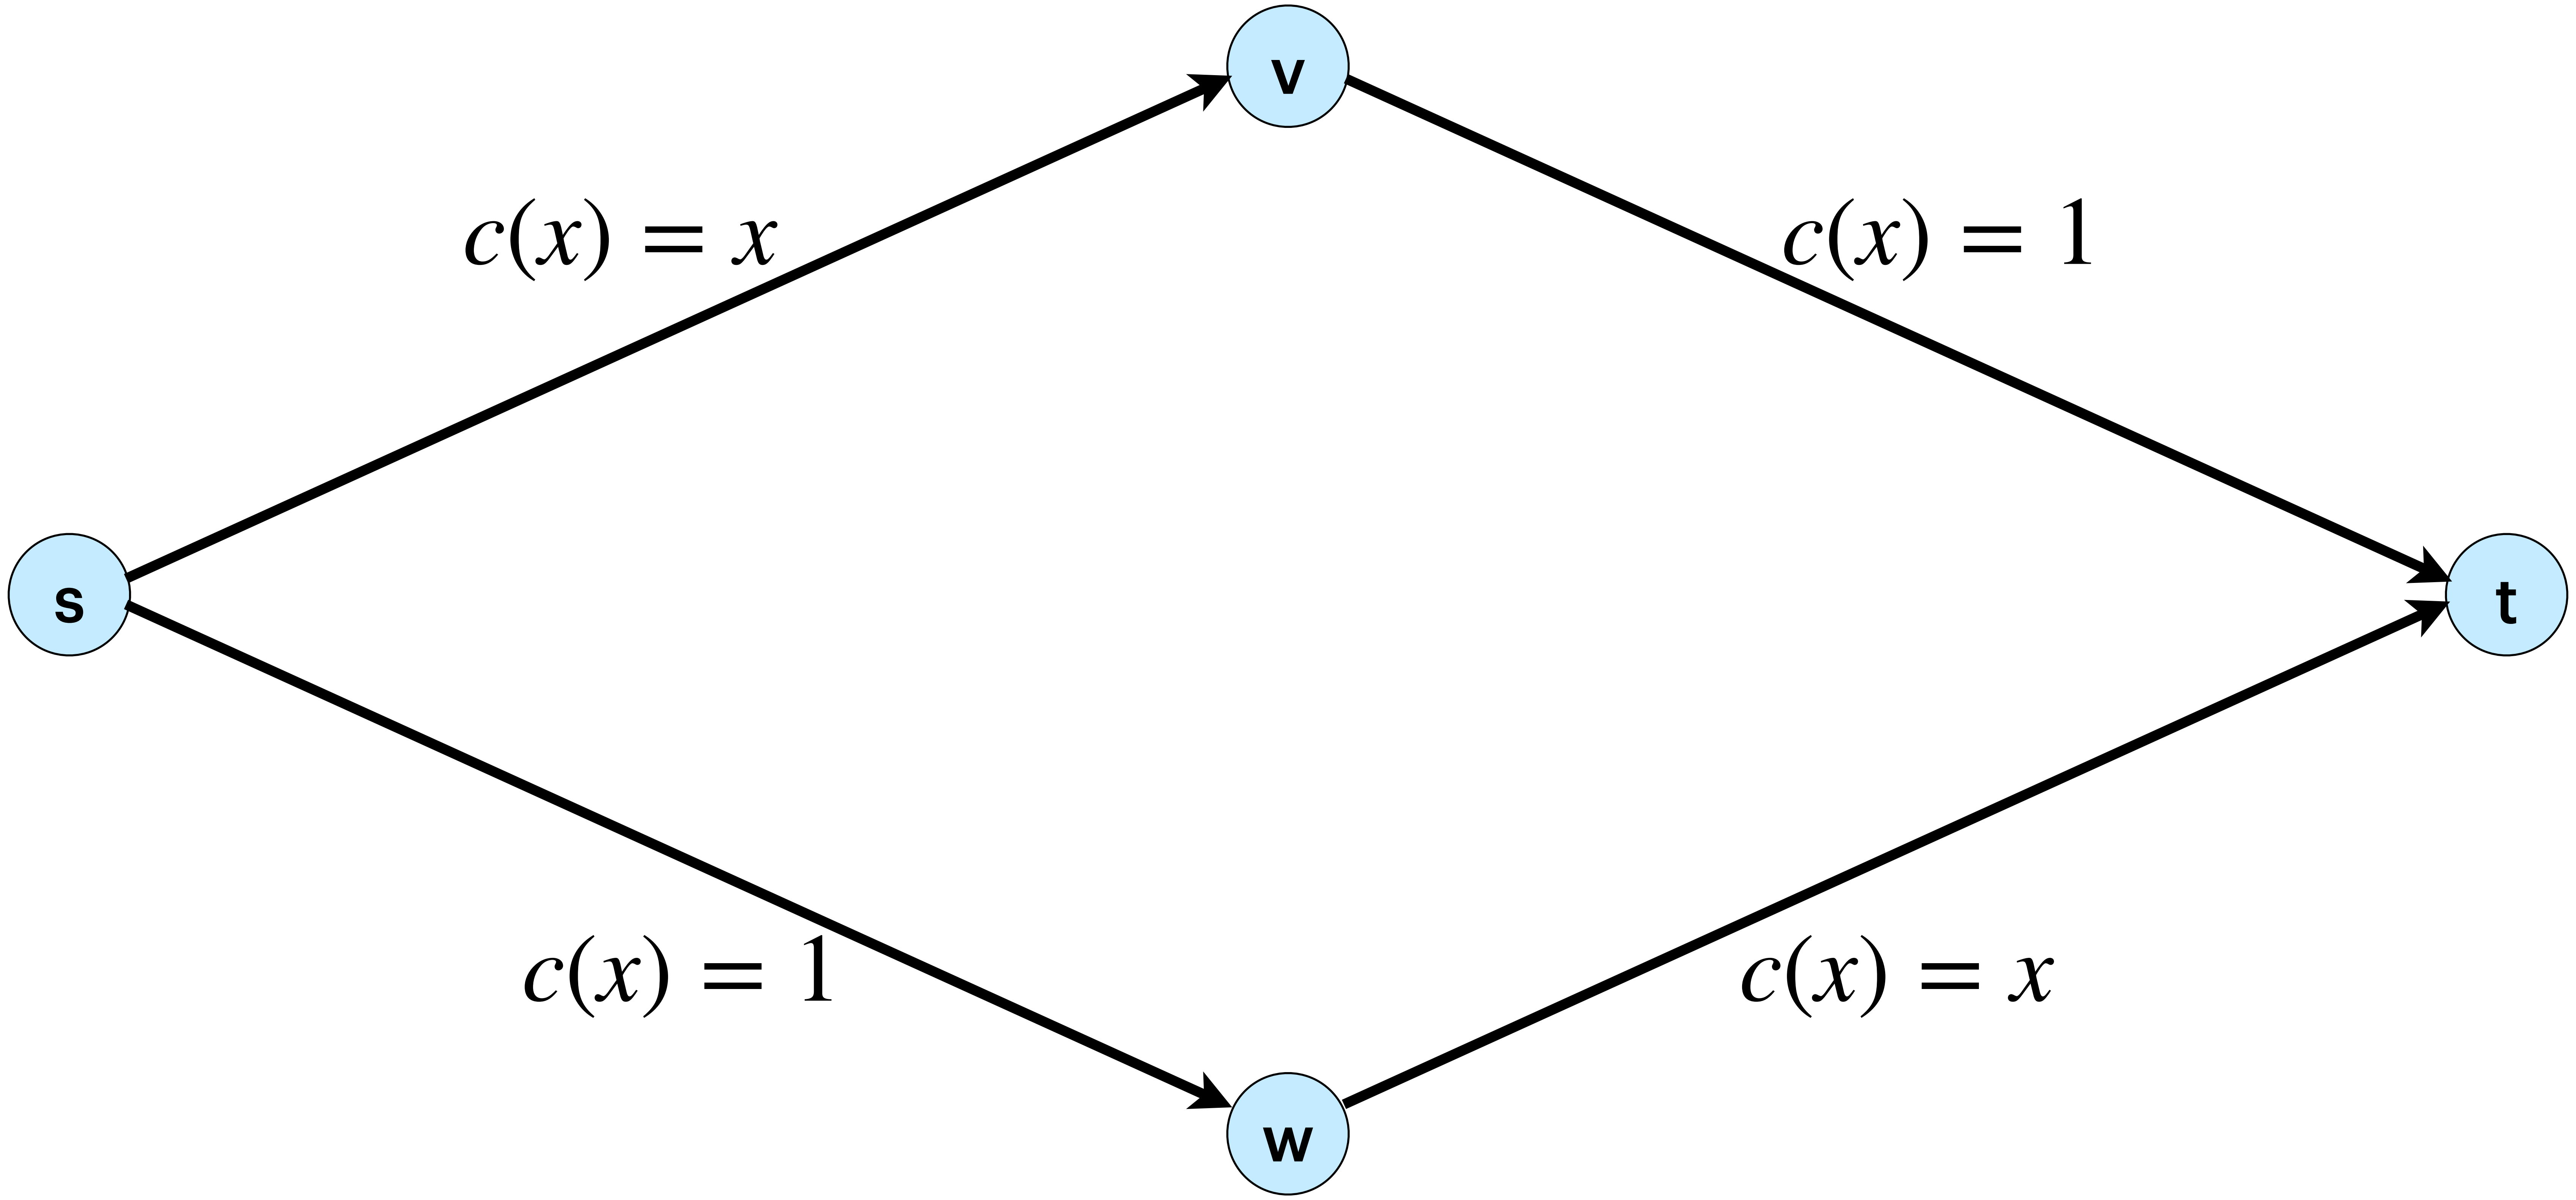

Identical paths, combined travel time: $1 + x$ each

Therefore, time taken: 90 mins

Braess's Paradox

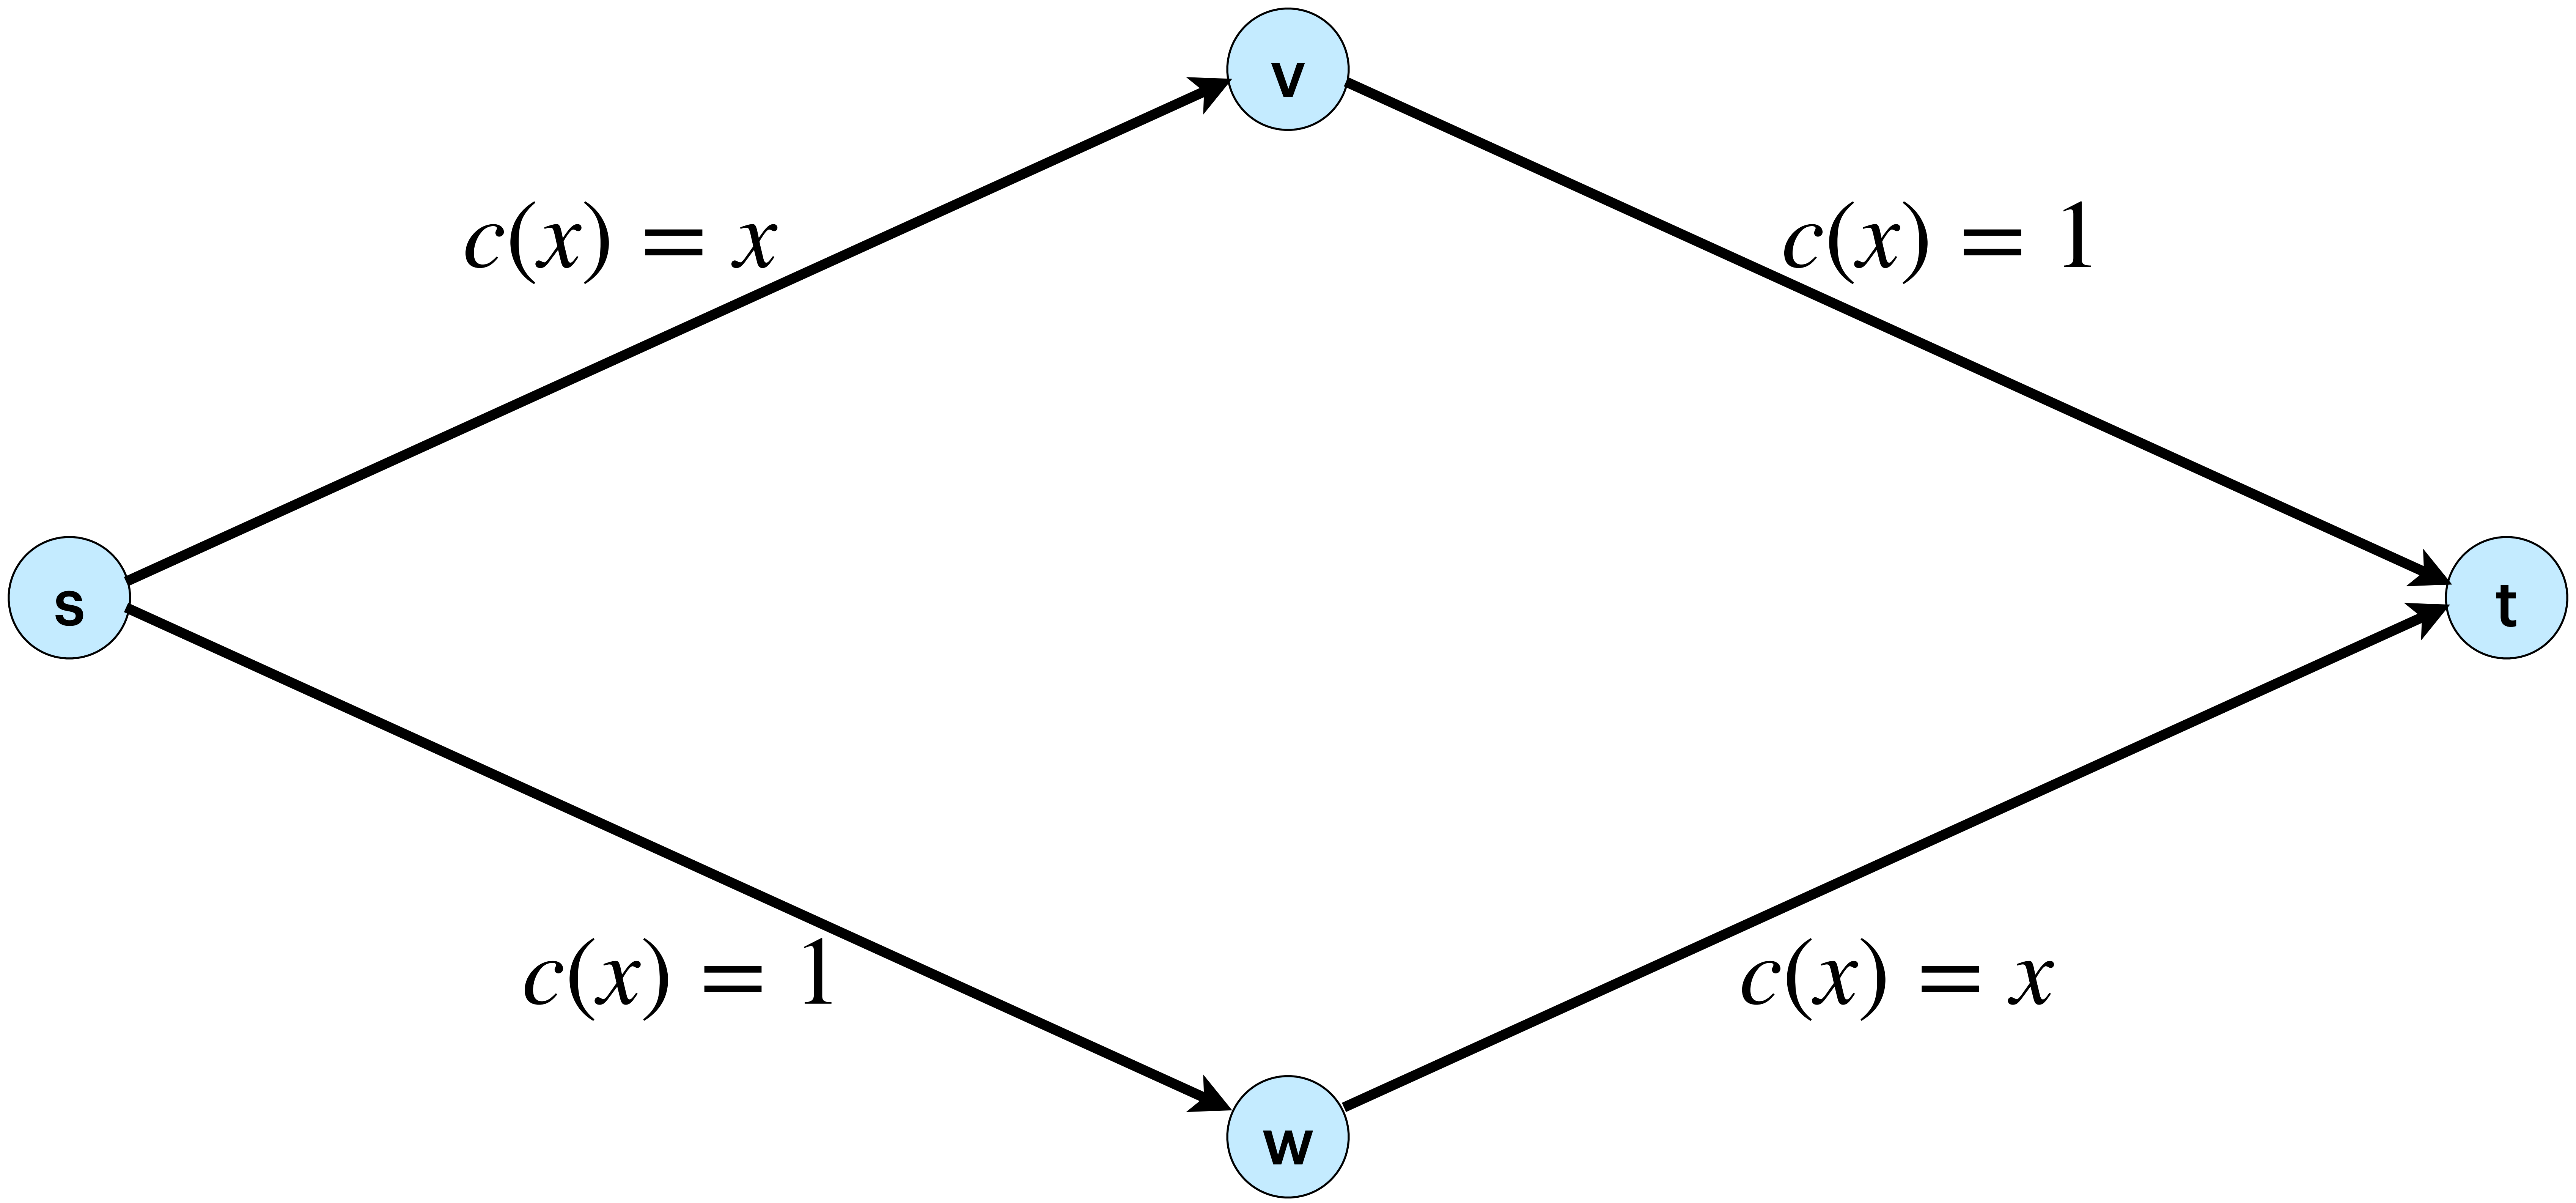

Braess's Paradox

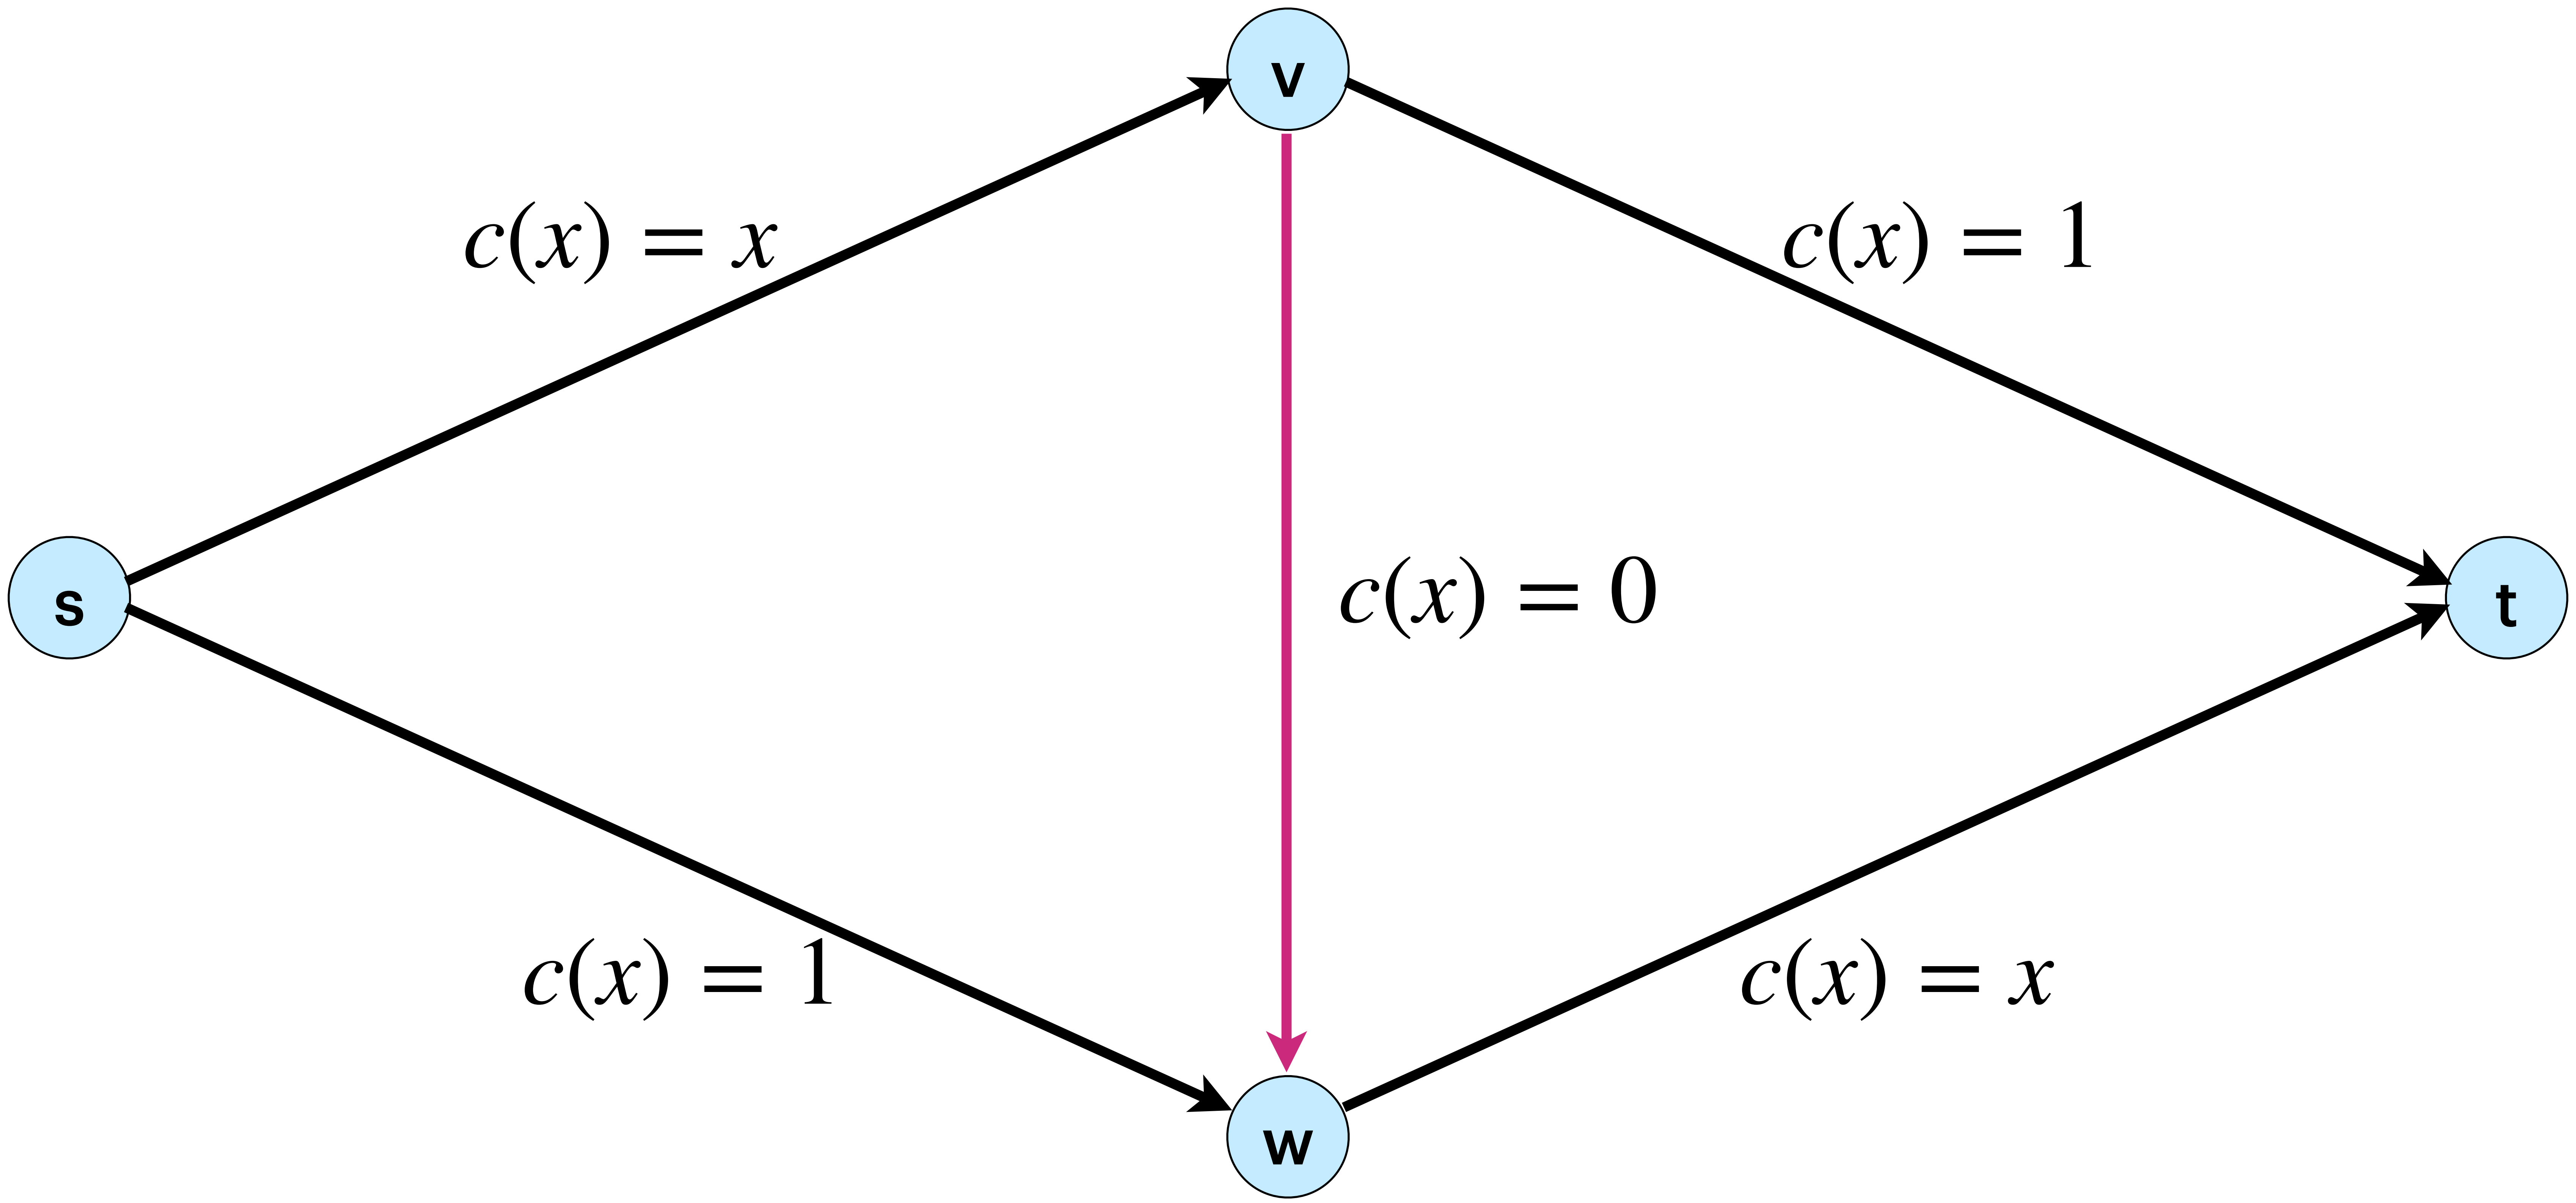

Braess's Paradox

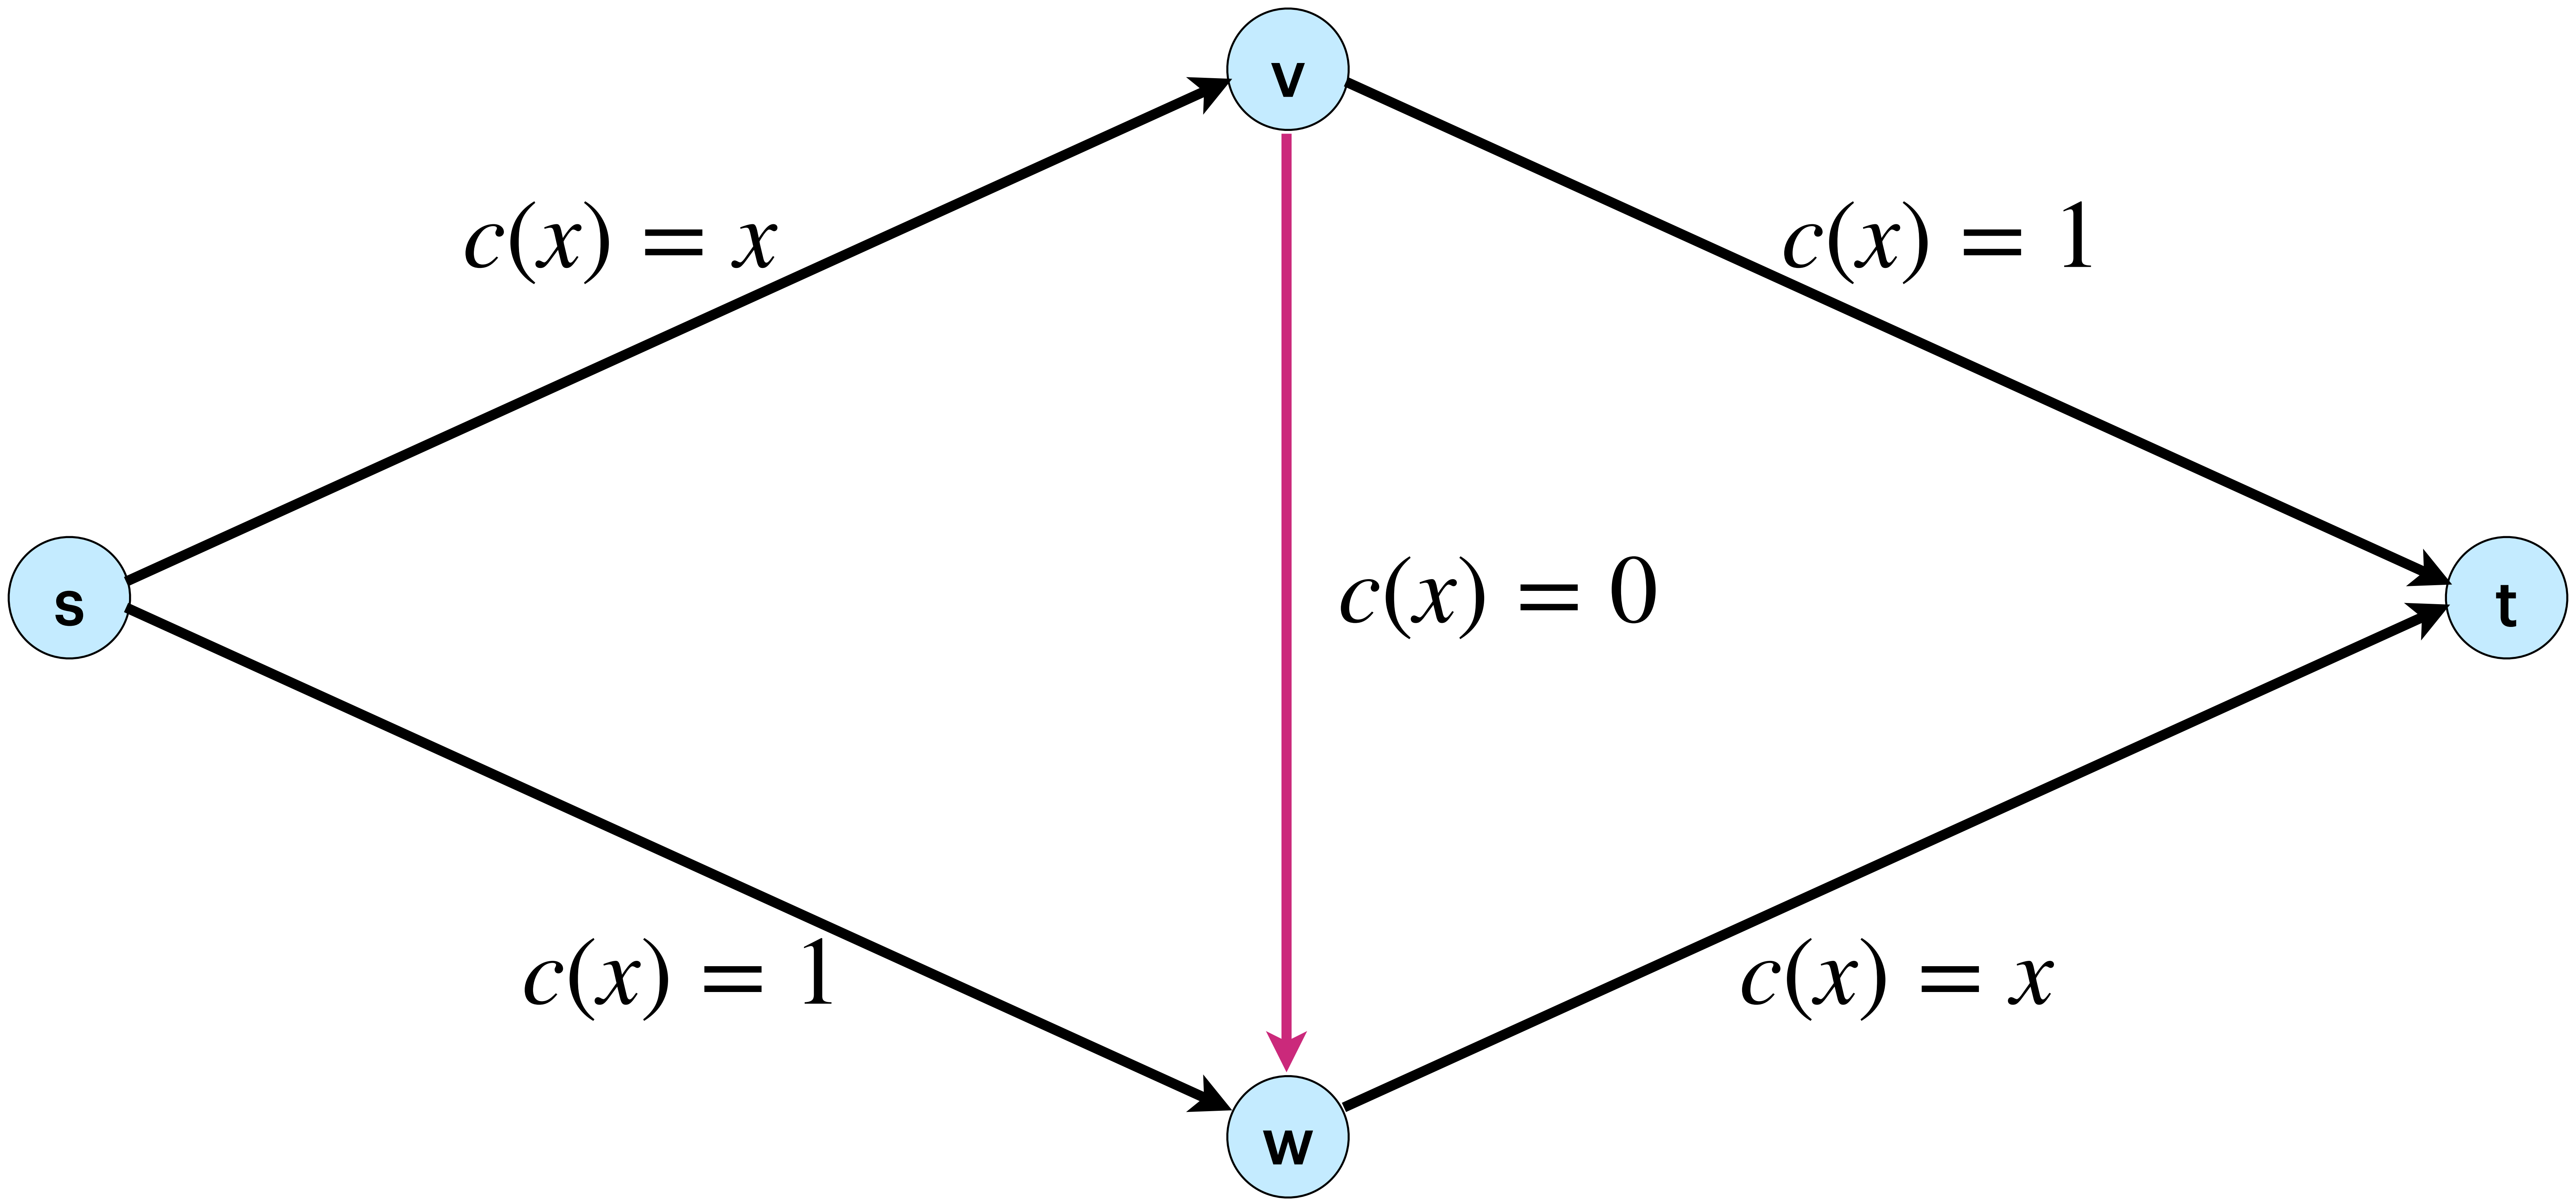

Since $2x < 1 + x$, the route $s \to v \to w \to t$ is always better.

Braess's Paradox

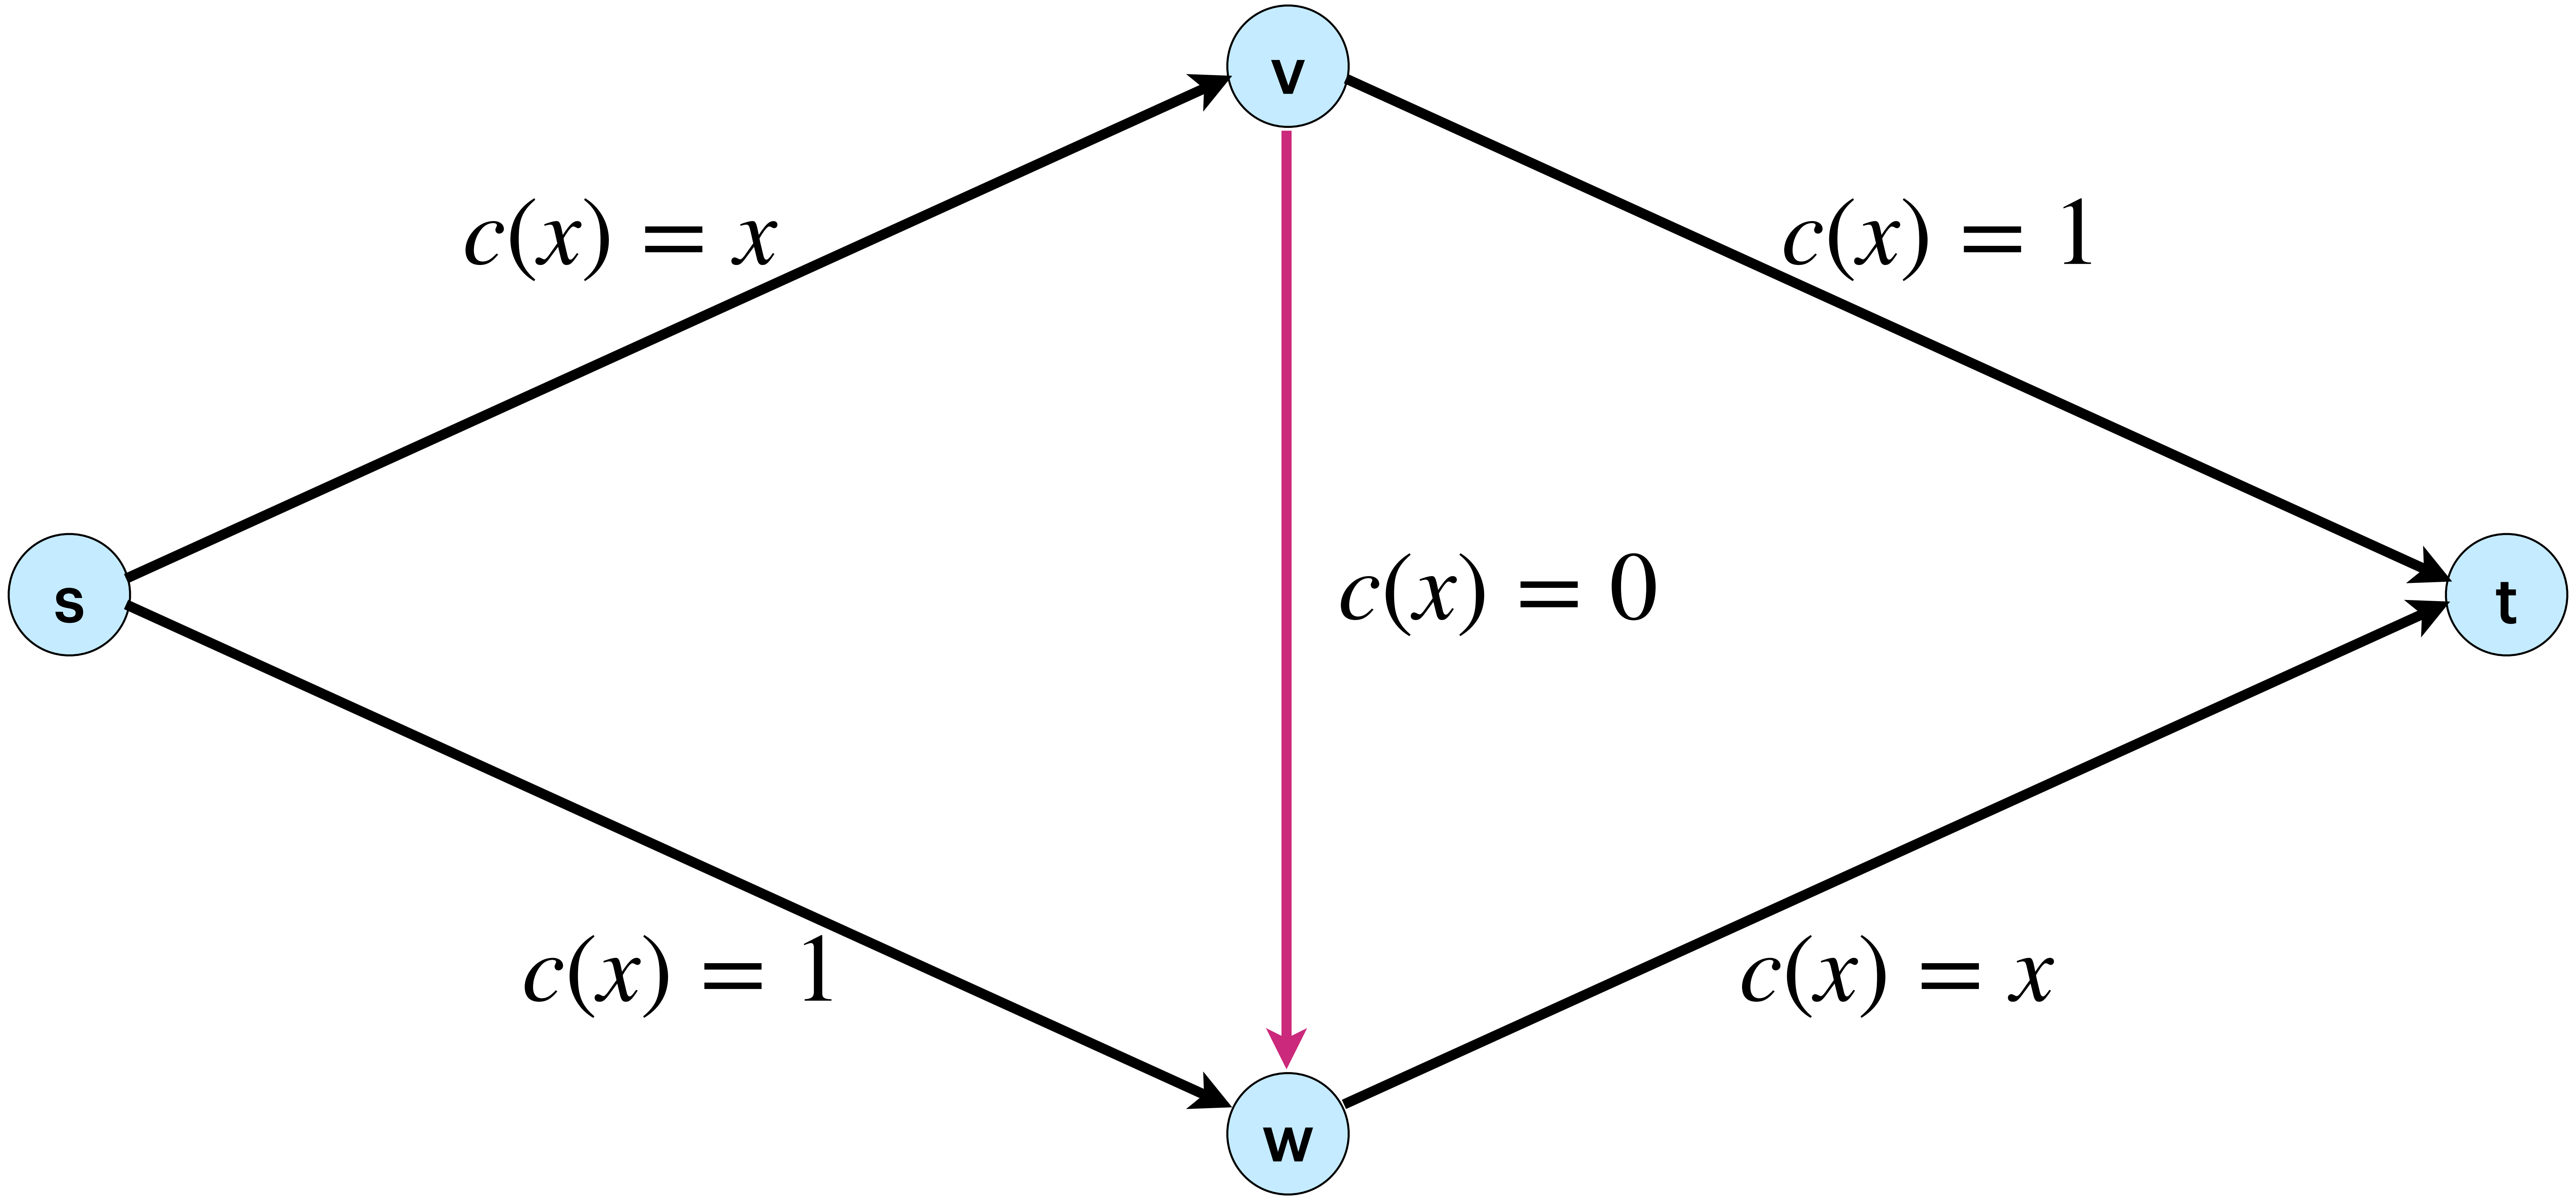

Since $2x < 1 + x$, the route $s \to v \to w \to t$ is always better.

Therefore, time taken: 120 mins

Braess's Paradox

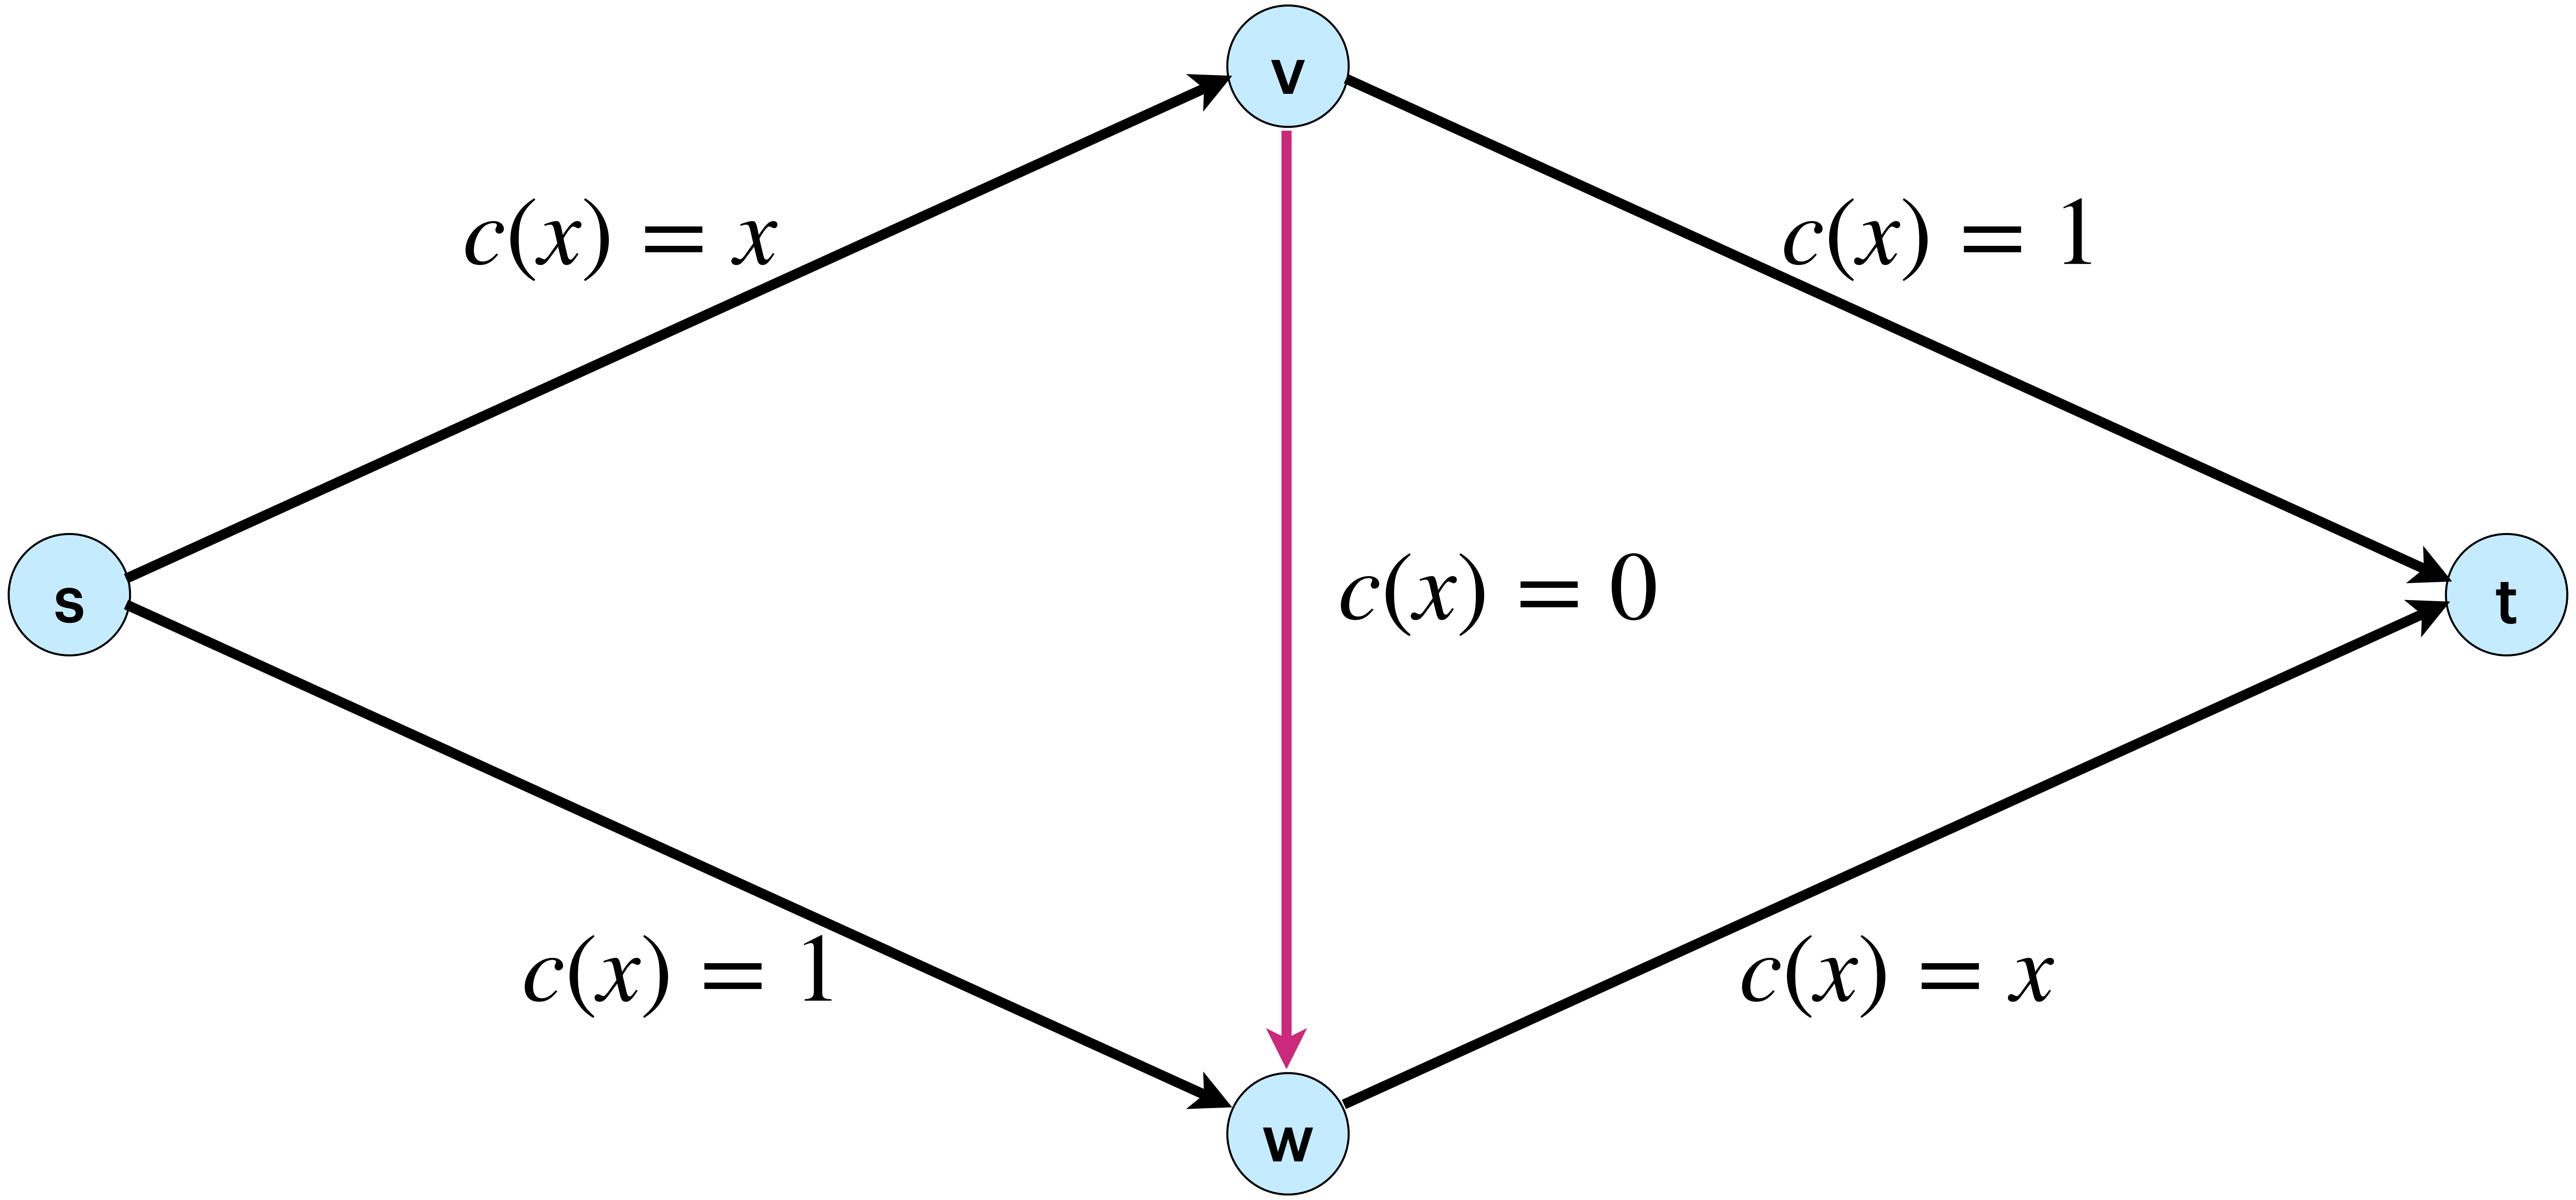

Braess's Paradox

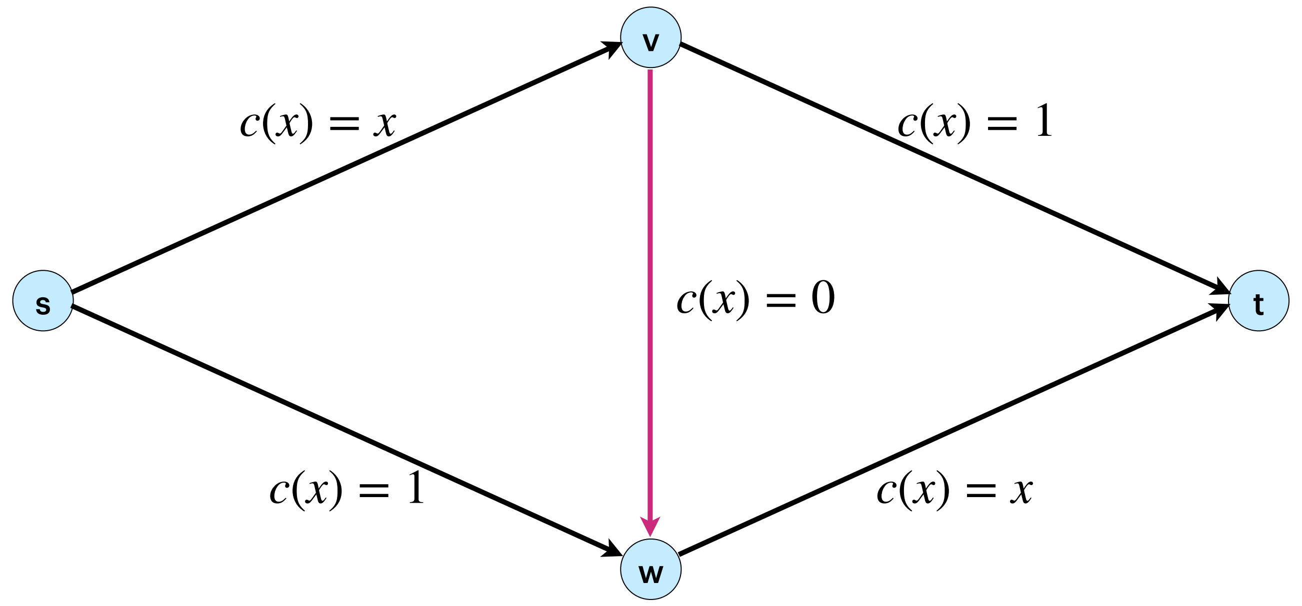

Braess's Paradox

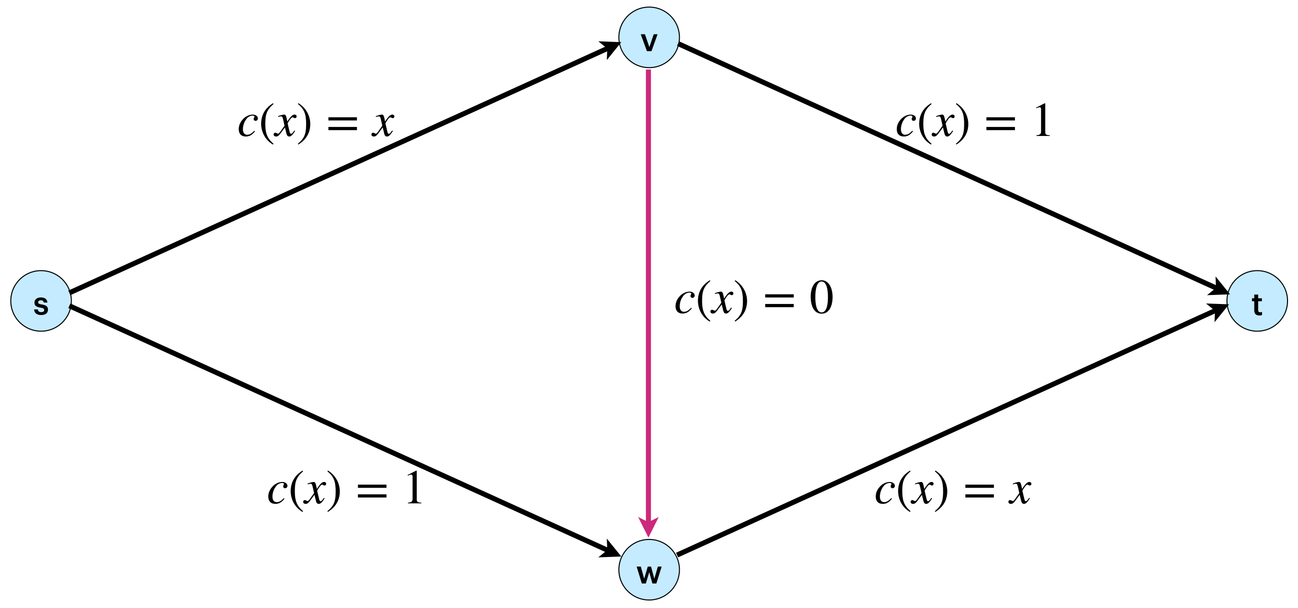

- If there was a benevolent dictator, she could ask everyone to take the previous routs, ignoring the teleportation machine. Then travel time: 90 mins

Braess's Paradox

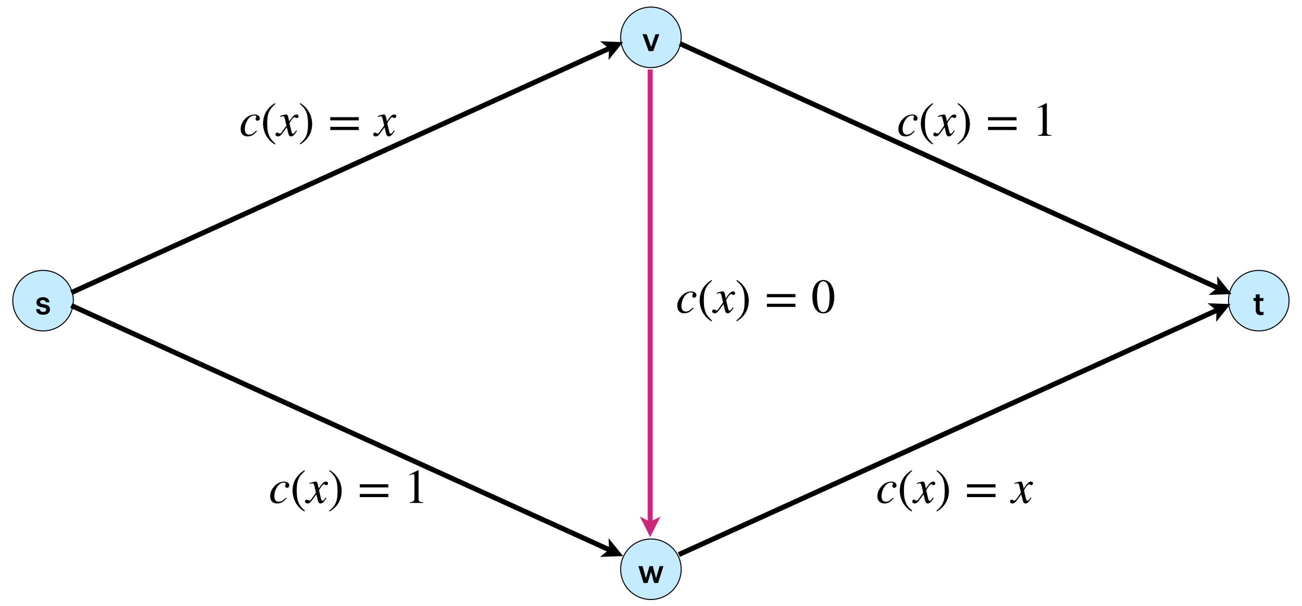

- If there was a benevolent dictator, she could ask everyone to take the previous routs, ignoring the teleportation machine. Then travel time: 90 mins

Price of Anarchy (POA): The ratio between the system performance with strategic players and the best possible system performance

Braess's Paradox

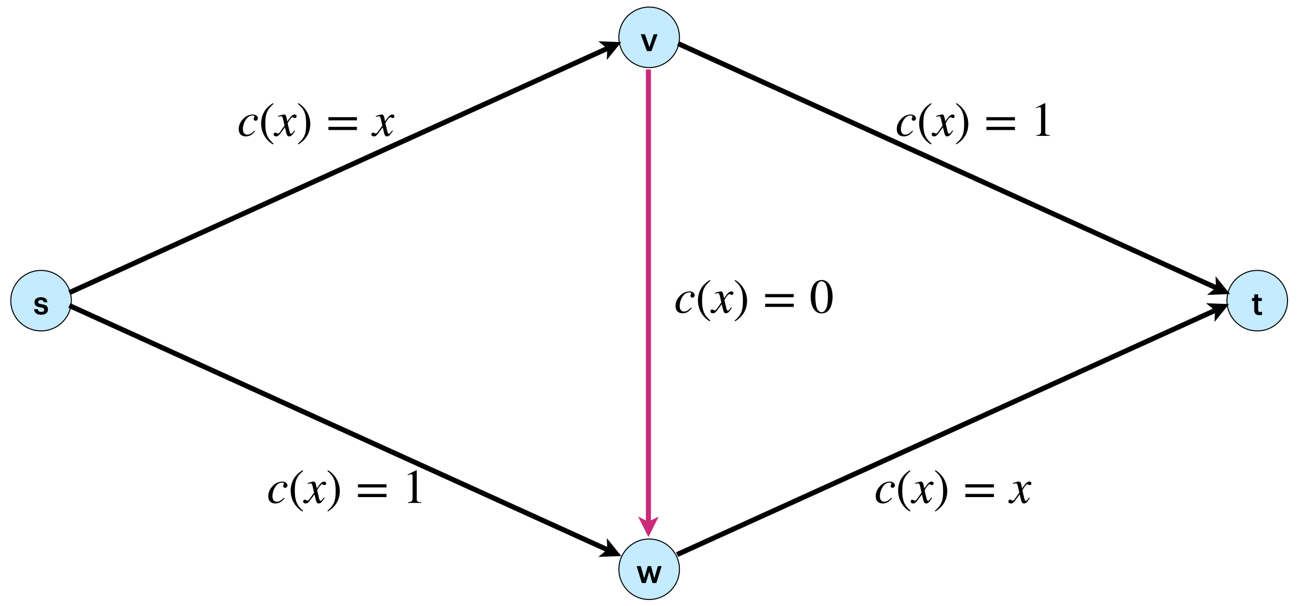

- If there was a benevolent dictator, she could ask everyone to take the previous routs, ignoring the teleportation machine. Then travel time: 90 mins

> Price of Anarchy (POA): The ratio between the system performance with strategic players and the best possible system performance

$$\text{POA} = \frac{120}{90} = \frac{4}{3}$$

No Netflix Please...

Suppose n people on a hotel in a remote location share a single internet connection.

Problem: Capacity of this communication channel is limited.

Sending email is okay. Streaming movies breaks the internet for everyone.

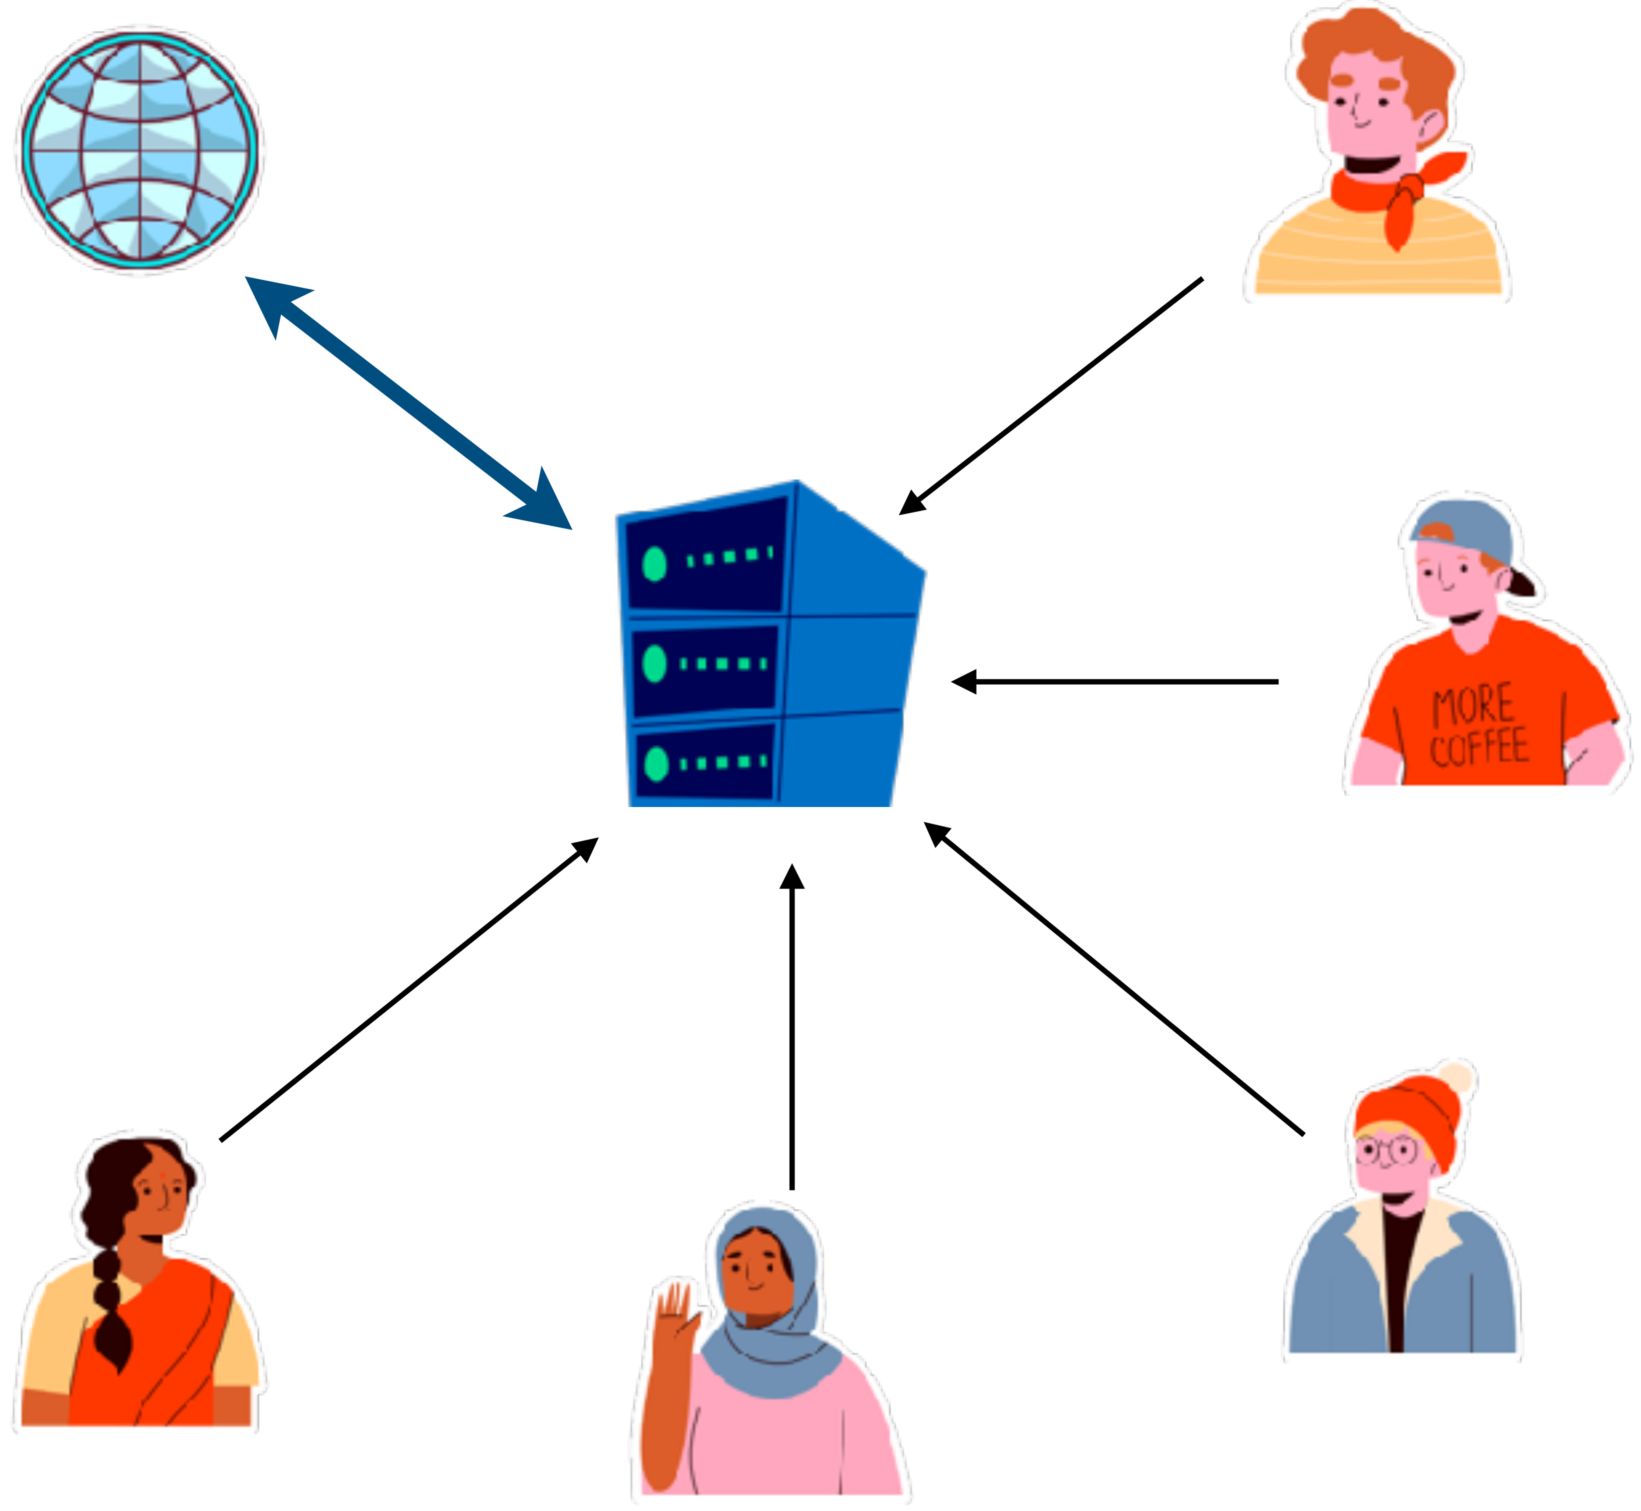

Model

Model

- Total number of agents sharing the channel : n

Model

- Total number of agents sharing the channel : n
- Quality of the channel deteriorates linearly with increasing amount of traffic.

Model

- Total number of agents sharing the channel : n
- Quality of the channel deteriorates linearly with increasing amount of traffic.
    - If total traffic is $t \leq 1$, the quality is $1 - t$

Model

- Total number of agents sharing the channel : n
- Quality of the channel deteriorates linearly with increasing amount of traffic.
    - If total traffic is $t \leq 1$, the quality is $1 - t$
    - If total traffic is $t > 1$, the quality is 0

Model

- Total number of agents sharing the channel : n
- Quality of the channel deteriorates linearly with increasing amount of traffic.
    - If total traffic is $t \leq 1$, the quality is $1 - t$
    - If total traffic is $t > 1$, the quality is 0
- Each agent $i \in [n]$ can choose $x_i \in [0,1]$, and she gets $x_i \cdot$ (Quality of channel)

Model

- Total number of agents sharing the channel : n
- Quality of the channel deteriorates linearly with increasing amount of traffic.
    - If total traffic is $t \leq 1$, the quality is $1 - t$
    - If total traffic is $t > 1$, the quality is 0
- Each agent $i \in [n]$ can choose $x_i \in [0,1]$, and she gets $x_i \cdot \overbrace{\text{(Quality of channel)}}$

Model

- Total number of agents sharing the channel : n
- Quality of the channel deteriorates linearly with increasing amount of traffic.
  - If total traffic is $t \leq 1$, the quality is $1 - t$
  - If total traffic is $t > 1$, the quality is 0
- Each agent $i \in [n]$ can choose $x_i \in [0,1]$, and she gets $x_i \cdot \overbrace{\text{(Quality of channel)}}^{\text{Payoff}}$

Model

- Total number of agents sharing the channel : n
- Quality of the channel deteriorates linearly with increasing amount of traffic.
  - If total traffic is $t \leq 1$, the quality is $1 - t$
  - If total traffic is $t > 1$, the quality is 0
- Each agent $i \in [n]$ can choose $x_i \in [0,1]$, and she gets $\overbrace{x_i \cdot (\text{Quality of channel})}^{\text{Payoff}}$
  - Payoff $p_i = x_i \cdot \left(1 - \sum_{j \in [n]} x_j\right)$

Agents are *rational* and *selfish*. Each agent wants to maximize their payoff

What will happen?

Analysis

Let $t_i = \sum_{j \neq i} x_j$

If $t_i \geq 1$, agent i gets payoff 0

If $t_i < 1$, agent i will get a payoff of $x_i \cdot \left(1 - \sum_{j \in [n]} x_j\right) = x_i \cdot (1 - t_i - x_i)$

Therefore, payoff is maximized at $x_i = \dfrac{1 - t_i}{2}$

Stability/Equilibrium

Stability/Equilibrium

Question: Can each agent pick a number in such a way that no agent would want to revise their choice?

Stability/Equilibrium

Question: Can each agent pick a number in such a way that no agent would want to revise their choice?

Answer: Yes, there exists a unique stable solution $X^* = \langle x_1^*, x_2^*, \ldots, x_n^* \rangle$ such that $x_i^* = \dfrac{1}{n+1} \forall i \in [n]$

Proof

In an equilibrium, each agent gets maximum payoff. Therefore,

$$\forall i \in [n], \quad x_i = \frac{1 - t_i}{2} \quad (1)$$

$$\implies t_i = 2x_i - 1$$

$$\implies t_i - t_j = 2x_i - 1 - 2x_j + 1 = 2(x_i - x_j) \quad (2)$$

We know that $t_i = \sum_{j \neq i} x_j$, therefore, $t_i - t_j = x_j - x_i \quad (3)$

Substitute (3) in (2), we get $\forall i \in [n], \quad x_i = x_j \quad (4)$

Now, solving for x_i in (1) with $t_i = (n-1)x_i$, gives $x_i = \dfrac{1}{n+1}$

Why is it a Tragedy?

Why is it a Tragedy?

In the stable solution,

Why is it a Tragedy?

In the stable solution,

the payoff of each player i is $p_i = \frac{1}{n+1} \cdot \left(1 - \frac{n}{n+1}\right) = \frac{1}{(n+1)^2}$

Why is it a Tragedy?

In the stable solution,

the payoff of each player i is $p_i = \frac{1}{n+1} \cdot \left(1 - \frac{n}{n+1}\right) = \frac{1}{(n+1)^2}$

However, had they chosen $\forall i \in [n], \quad x_i = \frac{1}{2n}$, then $p_i = \frac{1}{4n}$

Why is it a Tragedy?

In the stable solution,

the payoff of each player i is $p_i = \frac{1}{n+1} \cdot \left(1 - \frac{n}{n+1}\right) = \frac{1}{(n+1)^2}$

However, had they chosen $\forall i \in [n]$, $x_i = \frac{1}{2n}$, then $p_i = \frac{1}{4n}$ ← **Optimal**

Why is it a Tragedy?

In the stable solution,

the payoff of each player i is $p_i = \frac{1}{n+1} \cdot \left(1 - \frac{n}{n+1}\right) = \frac{1}{(n+1)^2}$

However, had they chosen $\forall i \in [n]$, $x_i = \frac{1}{2n}$, then $p_i = \frac{1}{4n}$ ← **Optimal**

Price of Anarchy (POA) $= \frac{4n}{(n+1)^2}$

Tragedy of the Commons: *Why does this happen?*

Tragedy of the Commons: *Why does this happen?*

- The action of a player brings large benefits to the player itself. But also, brings a cost to everybody.

Tragedy of the Commons: *Why does this happen?*

- The action of a player brings large benefits to the player itself. But also, brings a cost to everybody.
- The benefit only goes to the acting player. But the cost is shared by many players.

Tragedy of the Commons: *Why does this happen?*

- The action of a player brings large benefits to the player itself. But also, brings a cost to everybody.
- The benefit only goes to the acting player. But the cost is shared by many players.
- The large benefit to the player outweighs the share of the additional cost that the player incurs.

Tragedy of the Commons: *Why does this happen?*

- The action of a player brings large benefits to the player itself. But also, brings a cost to everybody.
- The benefit only goes to the acting player. But the cost is shared by many players.
- The large benefit to the player outweighs the share of the additional cost that the player incurs.
- But the total cost introduced to the system outweighs the total benefit

Examples

Examples

Pollution

Examples

Pollution

Over Fishing

Examples

Pollution

Over Fishing

Antibiotic Overuse

Examples

Pollution

Over Fishing

Antibiotic Overuse

Fast Fashion

Examples

Pollution

Over Fishing

Antibiotic Overuse

Fast Fashion

Evolutionary Extinction

Examples

Pollution

Over Fishing

Antibiotic Overuse

Fast Fashion

Evolutionary Extinction

Over Population

Course Goal 3: Complexity of Equilibria

Equilibria

- Informally, an equilibrium is a "steady state" of a system where each participant, assuming everything else stays the same, want to remain as-is.
- Computing Equilibria:
    - Easy in some games - Eg: using the teleport
    - Not so easy always (As we will see later)

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

No "deterministic equilibrium"

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

- Idea (by von Neumann) is to allow randomized (a.k.a. mixed) strategies.

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

- Idea (by von Neumann) is to allow randomized (a.k.a. mixed) strategies.
- If both players randomize uniformly in Rock-Paper-Scissors, then neither player can increase their expected payoff via a unilateral deviation. every unilateral deviation yields zero expected payoff to the deviator.

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

- Idea (by von Neumann) is to allow randomized (a.k.a. mixed) strategies.
- If both players randomize uniformly in Rock-Paper-Scissors, then neither player can increase their expected payoff via a unilateral deviation. every unilateral deviation yields zero expected payoff to the deviator.
- A pair of probability distributions (mixed-strategy) with this property is a Nash equilibrium.

Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	1,-1	1,-1	0,0

- Idea (by von Neumann) is to allow randomized (a.k.a. mixed) strategies.
- If both players randomize uniformly in Rock-Paper-Scissors, then neither player can increase their expected payoff via a unilateral deviation. every unilateral deviation yields zero expected payoff to the deviator.
- A pair of probability distributions (mixed-strategy) with this property is a Nash equilibrium.

Nash's Theorem (1951): Every bimatrix game has a Pure Nash Equilibrium

Rock-Paper-Scissors

Nash's Theorem (1951): Every bimatrix game has a Pure Nash Equilibrium

Rock-Paper-Scissors

Nash's Theorem (1951): Every bimatrix game has a Pure Nash Equilibrium

Rock-Paper-Scissors

Nash's Theorem (1951): Every bimatrix game has a Pure Nash Equilibrium

- Nash's theorem holds more generally in games with any finite number of players.

Rock-Paper-Scissors

Nash's Theorem (1951): Every bimatrix game has a Pure Nash Equilibrium

- Nash's theorem holds more generally in games with any finite number of players.
- If a bimatrix game is zero-sum (like in Rock-Paper-Scissors) then a Nash equilibrium can be computed in polynomial time.

Rock-Paper-Scissors

Nash's Theorem (1951): Every bimatrix game has a Pure Nash Equilibrium

- Nash's theorem holds more generally in games with any finite number of players.
- If a bimatrix game is zero-sum (like in Rock-Paper-Scissors) then a Nash equilibrium can be computed in polynomial time.
  - Using linear programming

Rock-Paper-Scissors

Nash's Theorem (1951): Every bimatrix game has a Pure Nash Equilibrium

- Nash's theorem holds more generally in games with any finite number of players.
- If a bimatrix game is zero-sum (like in Rock-Paper-Scissors) then a Nash equilibrium can be computed in polynomial time.
  - Using linear programming
- General games - Computing Nash Equilibrium is PPAD-Hard (not NP -hard (unless NP = coNP))

Rock-Paper-Scissors

Nash's Theorem (1951): Every bimatrix game has a Pure Nash Equilibrium

- Nash's theorem holds more generally in games with any finite number of players.
- If a bimatrix game is zero-sum (like in Rock-Paper-Scissors) then a Nash equilibrium can be computed in polynomial time.
    - Using linear programming
- General games - Computing Nash Equilibrium is PPAD-Hard (not NP -hard (unless NP = coNP))
    - 2NASH is NP-Complete

Computing Equilibrium

- Nash equilibria is a fundamental computational problem of “intermediate” difficulty.
  - Like Graph Isomorphism, Factoring...
- If all parties are boundedly rational, then an equilibrium can be interpreted as a credible prediction only if it can be computed with reasonable effort.
  - "Easier" equilibrium concepts: correlated equilibria and coarse correlated equilibria

Exercise 1

- Prove that RPS has <u>unique</u> Nash Equilibria.