The Evolutionary Price of Anarchy

Laura Schmid

Joint work with S. Schmid and K. Chatterjee

OPODIS’19

IST AUSTRIA
Institute of Science and Technology

universität
wien

The price of anarchy

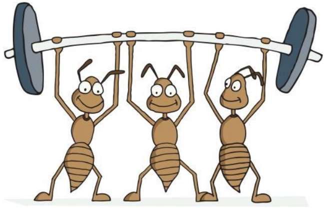

Large distributed systems often depend on cooperation, and can suffer if users behave selfishly

High costs and security issues

Price of anarchy: measure of how badly system can be affected by non-cooperative behavior

$$PoA = \frac{\text{System cost of worst Nash equilibrium}}{\text{Optimal system cost}}$$

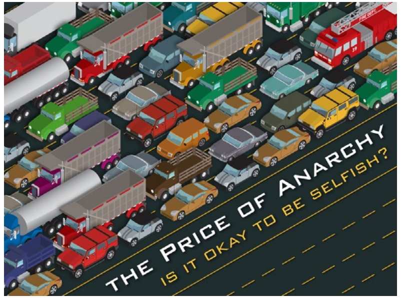

Nash: Given the other player's strategy, nobody can do better by changing their current action.

How realistic is this?

Are Nash equilibria/the PoA the right measure of efficiency?

A) The PoA typically considers one-shot interactions and fixed strategies

B) Players/nodes are rational and have global network information to play NE

C) Players enjoy unbounded resources, computation

How realistic is this?

Are Nash equilibria/the PoA the right measure of efficiency?

A) Distributed systems rely on dynamic interactions over time -> repeated games

B) Players/nodes are rational and have global network information to play NE

C) Players enjoy unbounded resources, computation

How realistic is this?

Are Nash equilibria/the PoA the right measure of efficiency?

A) Distributed systems rely on dynamic interactions over time -> repeated games

B) Nodes/players typically have local information about the network

C) Players enjoy unbounded resources, computation

How realistic is this?

Are Nash equilibria/the PoA the right measure of efficiency?

A) Distributed systems rely on dynamic interactions over time -> repeated games

B) Nodes/players typically have local information about the network

C) Players typically have limited resources and memory

How realistic is this?

Are Nash equilibria/the PoA the right measure of efficiency? No!

A) Distributed systems rely on dynamic interactions over time -> repeated games

B) Nodes/players typically have local information about the network

C) Players typically have limited resources and memory

The PoA does not account for this!

A new measure of efficiency

We would like to port the PoA to local information scenarios, where

- games are embedded in dynamical processes
- players are simple and even memoryless
- and strategies can evolve.

How?

Evolutionary games

“Classical” game theory framework: selfish individuals attempt to consciously reach best outcome for themselves

Central concept: (Nash-) equilibrium

Needs assumptions about rationality, beliefs, cognitive abilities

Evolutionary games do not!

Focus on dynamics -> equilibrium selection, off equilibrium behavior

Evolutionary games

Generic approach to evolutionary dynamics

Evolution and evolutionary dynamics

Evolutionary dynamics: mathematical principles behind evolution

Main ingredients for evolutionary change: Replication + Mutation + Selection

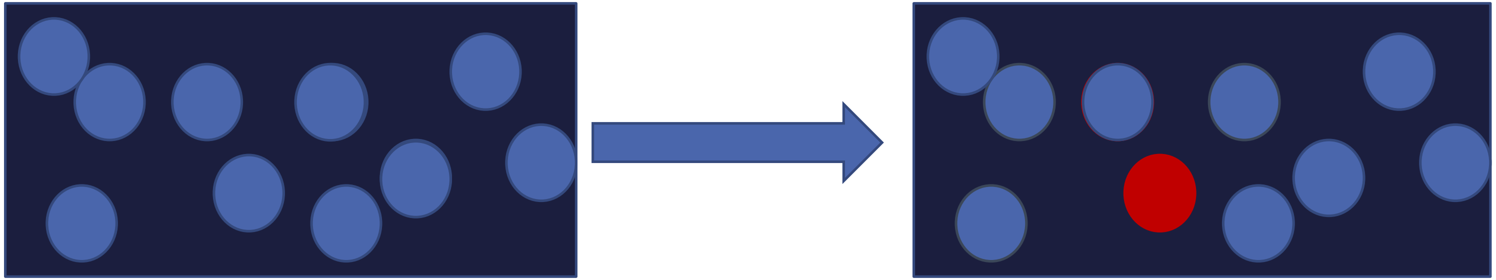

Evolution and evolutionary dynamics

Evolutionary dynamics: mathematical principles behind evolution

Main ingredients for evolutionary change: Replication + Mutation + Selection

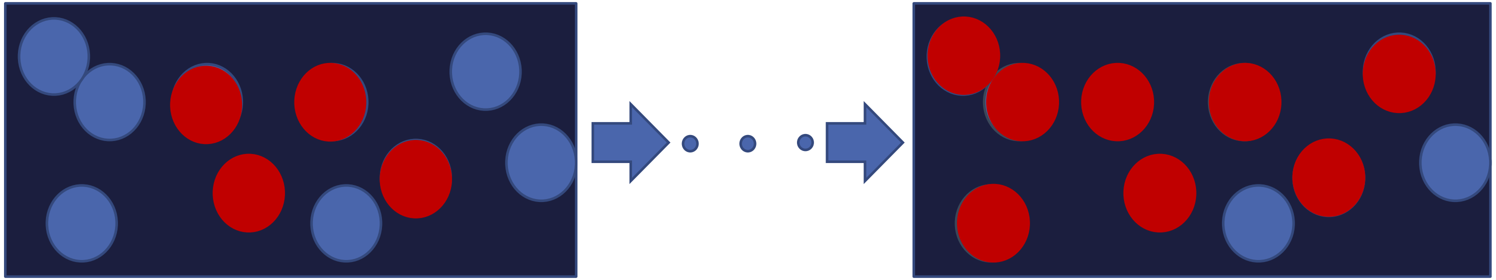

Evolutionary games

Generic approach to evolutionary dynamics

Frequency dependent fitness/selection

Population of players with individual strategies

Interactions with other players give payoffs ⇔ evolutionary fitness

Success in game -> reproductive success: good strategies reproduce faster & spread

Describe dynamics dependent on frequency of different types in population

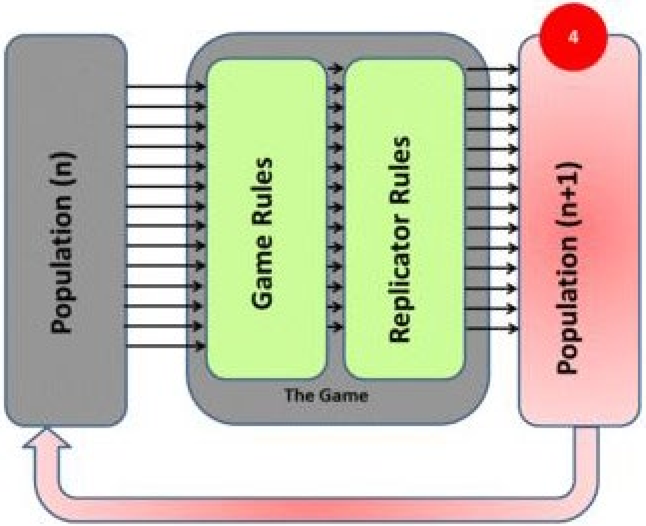

The evolutionary price of anarchy (ePoA)

The ePoA extends PoA to evolutionary games

More natural measure of efficiency than the static PoA

> We consider simple memoryless agents without perfect information, interacting repeatedly and locally
>
> Players do not even necessarily have to reach equilibrium in the game they are playing

We can study games under different evolutionary dynamics, different parameters

Exploration of equilibrium selection, long-term off-equilibrium behavior

Model game: Virus inoculation

Classic security game: N nodes in a network G must decide whether to install anti-virus software at cost V

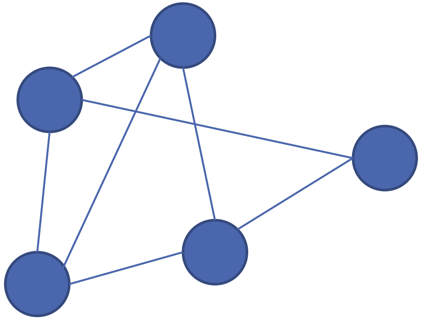

Model game: Virus inoculation

Classic security game: N nodes in a network G must decide whether to install anti-virus software at cost V

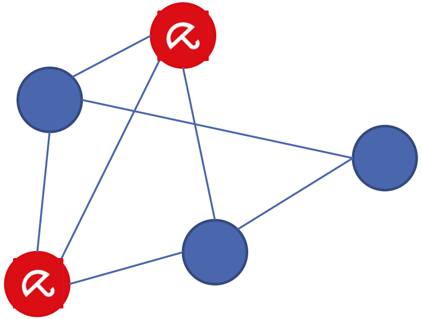

Model game: Virus inoculation

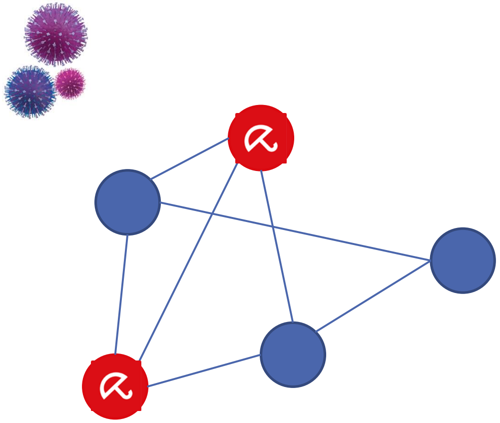

Classic security game: N nodes in a network G must decide whether to install anti-virus software at cost V

Unprotected nodes risk infection by virus spreading from a random location. Infection costs I > V.

Virus infects all unprotected nodes with direct path to an infected node. Inoculated nodes cannot be infected or transmit the virus.

Cost of strategy profile a for node i: $\mathrm{cost}_i(\vec{a}) = a_i V + (1 - a_i) I \cdot p_i(\vec{a})$

Resulting social cost: $\mathrm{cost}(\vec{a}) = \sum_{j=0}^{N-1} \mathrm{cost}_j(\vec{a})$

Well known results on Nash equilibria, optimum approximation

Evolutionary virus inoculation games

Nodes probably do not know G and each others' decision – they should only have local information!

Proposal: stochastic evolutionary process, iterated over many rounds

Three stages per round:

Evolutionary virus inoculation games

Nodes probably do not know the graph and each others' decision – they should only have local information!

Proposal: stochastic evolutionary process, iterated over many rounds

Three stages per round:

- Decision making: Two possible choices (0 or 1)

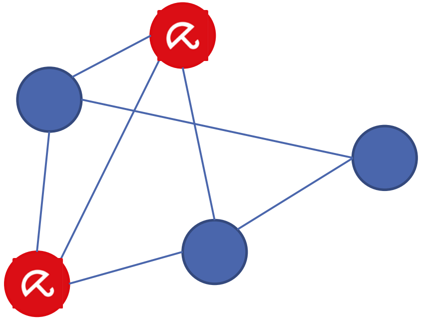

Evolutionary virus inoculation games

Nodes probably do not know G and each others' decision – they should only have local information!

Proposal: stochastic evolutionary process, iterated over many rounds

Three stages per round:

- Decision making
- Virus propagation

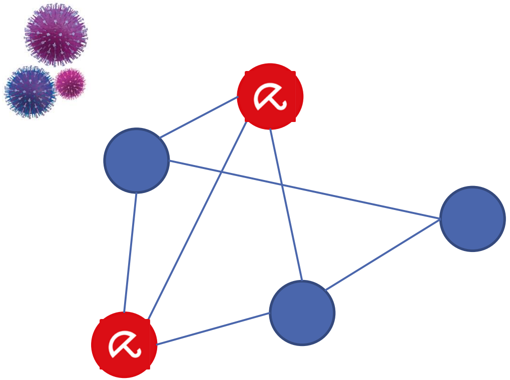

Evolutionary virus inoculation games

Nodes probably do not know G and each others' decision – they should only have local information!

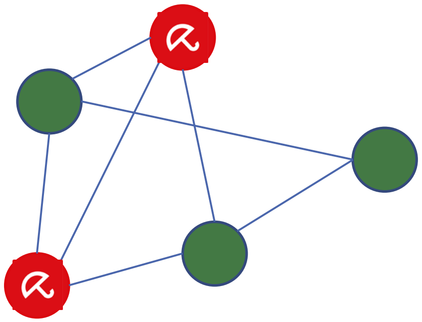

Proposal: stochastic evolutionary process, iterated over many rounds

Three stages per round:

- Decision making
- Virus propagation

Evolutionary virus inoculation games

Nodes probably do not know G and each others’ decision – they should only have local information!

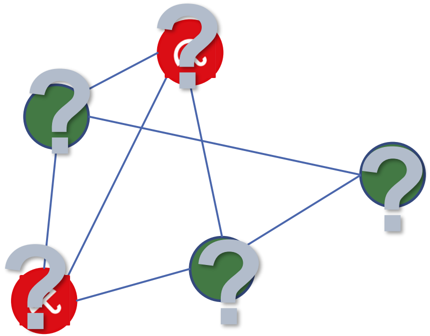

Proposal: stochastic evolutionary process, iterated over many rounds

Three stages per round:

- Decision making
- Virus propagation
- Evolution of strategies

Evolution of strategies

We take selection and mutation into account

We can compare different memoryless evolutionary dynamics

Example: Genetic evolution - The Moran Death-Birth process

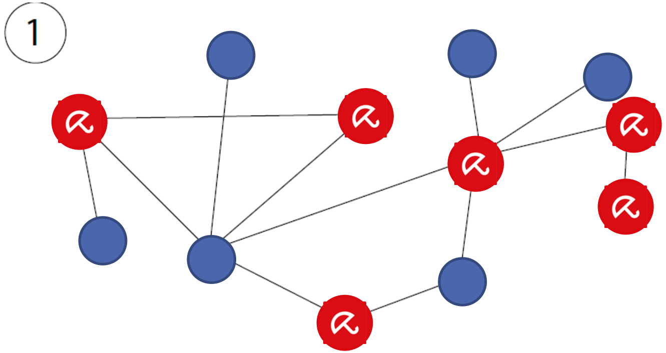

Evolution of strategies

We take selection and mutation into account

We can compare different memoryless evolutionary dynamics

Example: Genetic evolution - The Moran Death-Birth process

- Node is picked to die in each timestep

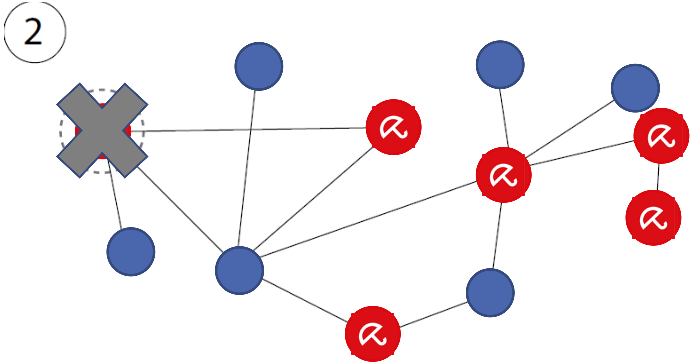

Evolution of strategies

We take selection and mutation into account

We can compare different memoryless evolutionary dynamics

Example: Genetic evolution - The Moran Death-Birth process

- Node is picked to die in each timestep
- Replaced by neighboring node proportional to the latter's payoff π

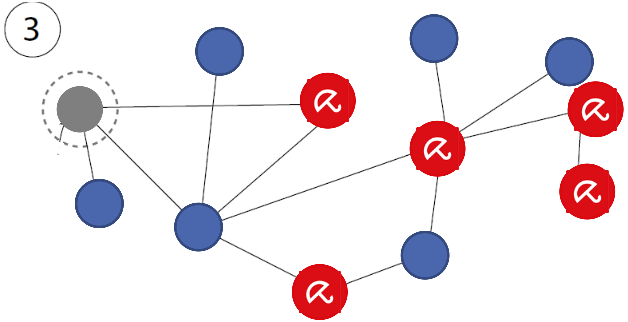

Evolution of strategies

We take selection and mutation into account

We can compare different memoryless evolutionary dynamics

Example: Genetic evolution - The Moran Death-Birth process

- Node is picked to die in each timestep
- Replaced by neighboring node proportional to the latter’s payoff

All considered dynamics can be described as ergodic Markov chains

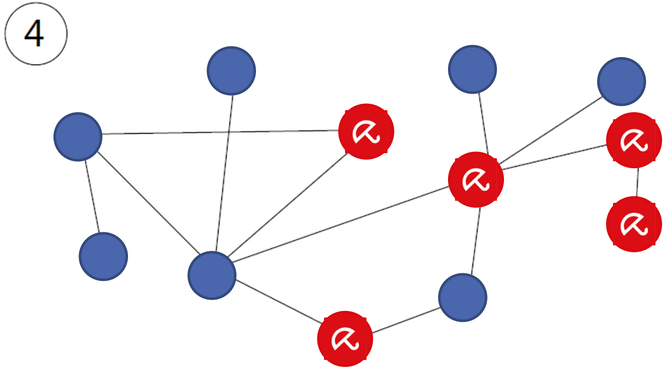

The evolutionary price of anarchy (ePoA), pt. 2

What is the stationary (limit) distribution of the underlying Markov chain = the probabilities of finding the system in different states?

$\Rightarrow$ Selection-mutation equilibrium **x** of a given evolutionary process

This allows us to find average social cost $\hat{S}$

$$\hat{S} = \mathbf{x} \cdot \mathbf{R} = \sum_i x_i R_i$$

where **R** contains the social costs of all possible system configurations

Definition: the evolutionary price of anarchy is the ratio of the average social cost of a process against the social optimum Ω.

$$ePoA = \hat{S}/\Omega$$

The virus game on a clique

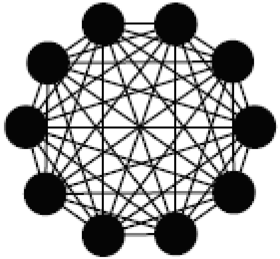

Markov chain: (N+1) states, t=0,…N (# of inoculated nodes)

The virus game on a clique

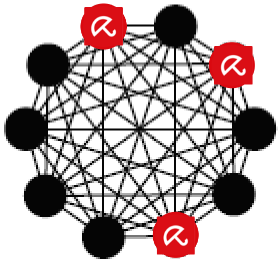

Markov chain: (N+1) states, t=0,…N (# of inoculated nodes)

Optimum at $t^* = \dfrac{N(2I-V)}{2I}$

Nash equilibria: $t = N - VN/I$ nodes inoculated => PoA

We can exactly calculate mutation-selection equilibrium **x**, for mutation rate μ and pure strategies $(a_i \in \{0,1\})$

Calculate average social cost $\hat{S}$ => ePoA

Can do this for different dynamics and compare their efficiency

Do we recover Nash equilibria in dynamics? How does ePoA compare to PoA?

Clique results

For "reasonable" parameters we recover the predicted Nash equilibria as the most abundant states of the evolutionary process

Most time is spent in states where $t = N - VN/I$ nodes are inoculated

Process is stochastic -> neighboring states also frequent, but symmetric distribution

Corollary: $\lim_{\mu \to 0} \lim_{N \to \infty} |ePoA_{Clique} - PoA_{Clique}| = 0$

Efficiency of evolutionary processes approaches static game

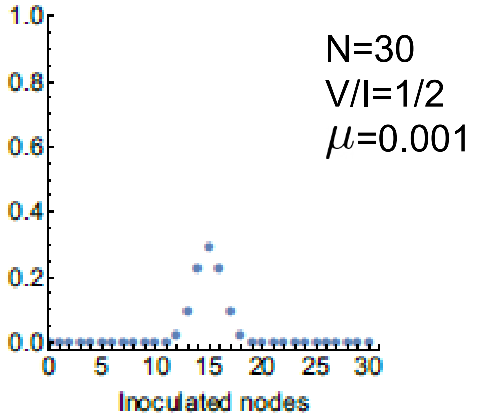

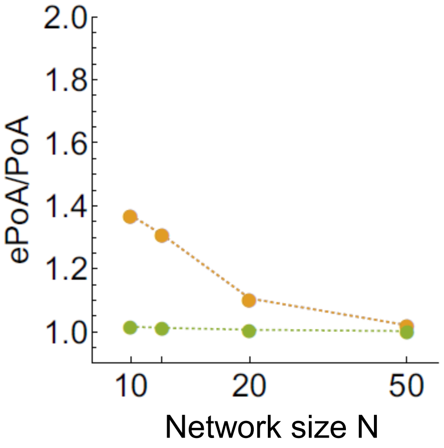

Virus game on star graphs

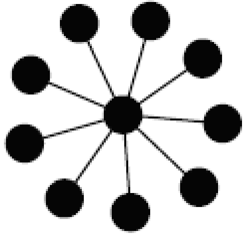

2N states, (c,l), c…inoculation state of center, l…number of inoculated leaf nodes

Virus game on star graphs

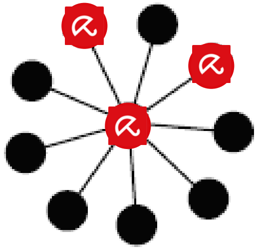

(c,l)=(1,2)

2N states, (*c*,*l*), *c*...inoculation state of center, *l*...number of inoculated leaf nodes

2 classes of Nash equilibria: $\mathcal{N}_1 = (0, N - \lfloor VN/I \rfloor)$ and $\mathcal{N}_2 = (1, 0)$

$\mathcal{N}_2$ is also the optimum!

We can again do exact calculations and get both ePoA and PoA

Now: Both classes of Nash equilibria are rare! The system exhibits strong off equilibrium behavior due to its topology

This implies that $ePoA_{Star} - PoA_{Star} \geq 0$ as long as mutation rate is small

Virus game on star graphs

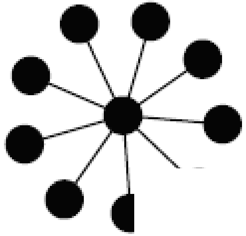

Now: Both classes of Nash equilibria are rare! The system exhibits off equilibrium behavior due to highly structured network

This implies that $ePoA_{Star} - PoA_{Star} \geq 0$ as long as mutation rate is small

Average costs are significantly higher than in traditional static model!

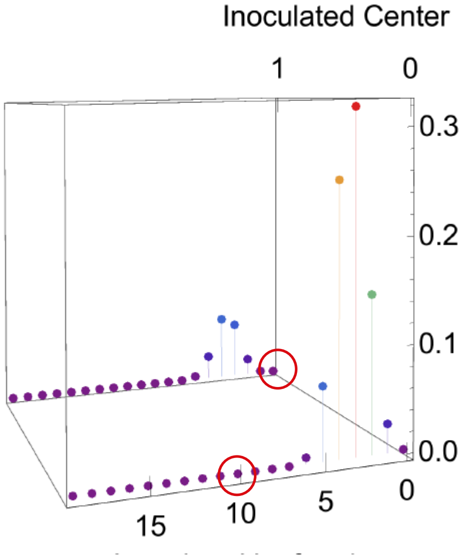

N=20
V/I=1/2
μ=0.001

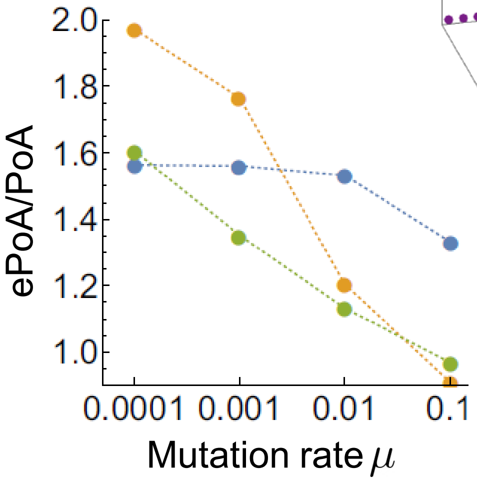

Simulations of more complex topologies

For most topologies, simulations are necessary – we cannot explicitly calculate **x**

Simulate the process, find average social welfare (or even stationary distribution approximation)

Example: 2-clique

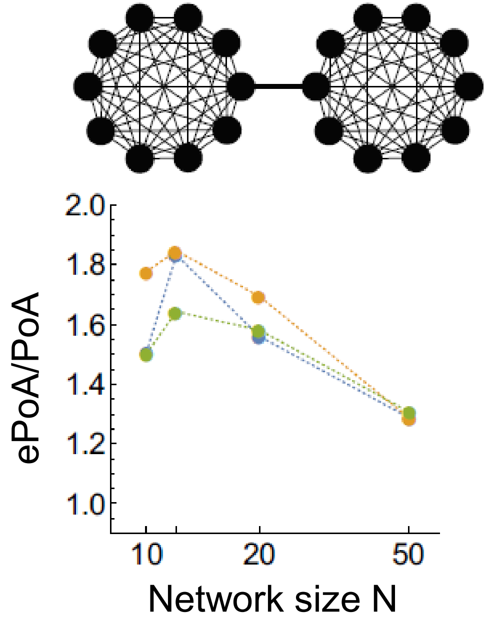

Simulations of more complex topologies

For most topologies, simulations are necessary – we cannot explicitly calculate **x**

Simulate the process, find average social welfare (or even stationary distribution approximation)

Example: 2-star

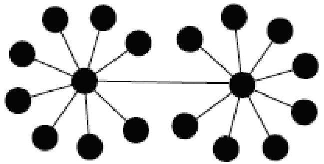

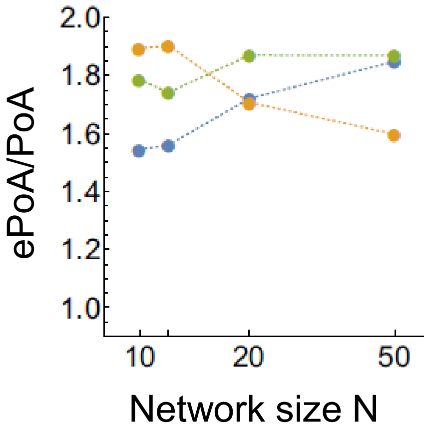

Simulations of more complex topologies

For most topologies, simulations are necessary – we cannot explicitly calculate **x**

Simulate the process, find average social welfare (or even stationary distribution approximation)

Example: cycle

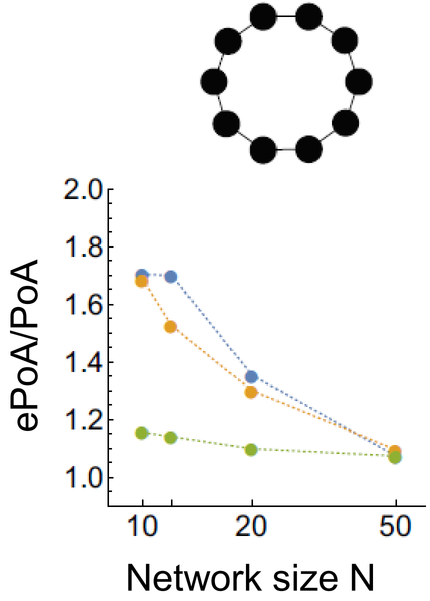

Usually no recovery of Nash equilibria for any of the considered dynamics => ePoA higher than PoA as long as μ is not too high

=> The PoA usually underestimates actual system costs for more complex topologies!

Takeaways

Static analysis of distributed systems based on the price of anarchy is falling short

We have introduced the evolutionary price of anarchy (ePoA) to study behavior of simple agents repeatedly interacting in a distributed system based on local information

Resulting stationary state can be significantly different from static equivalent/equilibria

System costs are therefore often higher than predicted by static price of anarchy

Shows impact of limited information on games in networks

Many avenues for future research

Thank you!