Evolutionary Equilibria in Computer Networks: Specialization and Niche Formation

Oliver Schulte
Petra Berenbrink

Simon Fraser University

oschulte@cs.sfu.ca

Modelling User Communities

- A system provides users with access to resources, e.g. a **network**.
- Centralized planning: gather requests, compute optimal allocation.
- "Anarchy": users individually choose resources, e.g. **routes for messages**.
- Individual choice $\rightarrow$ strategic interactions ($\approx$ traffic models).

Central Allocation

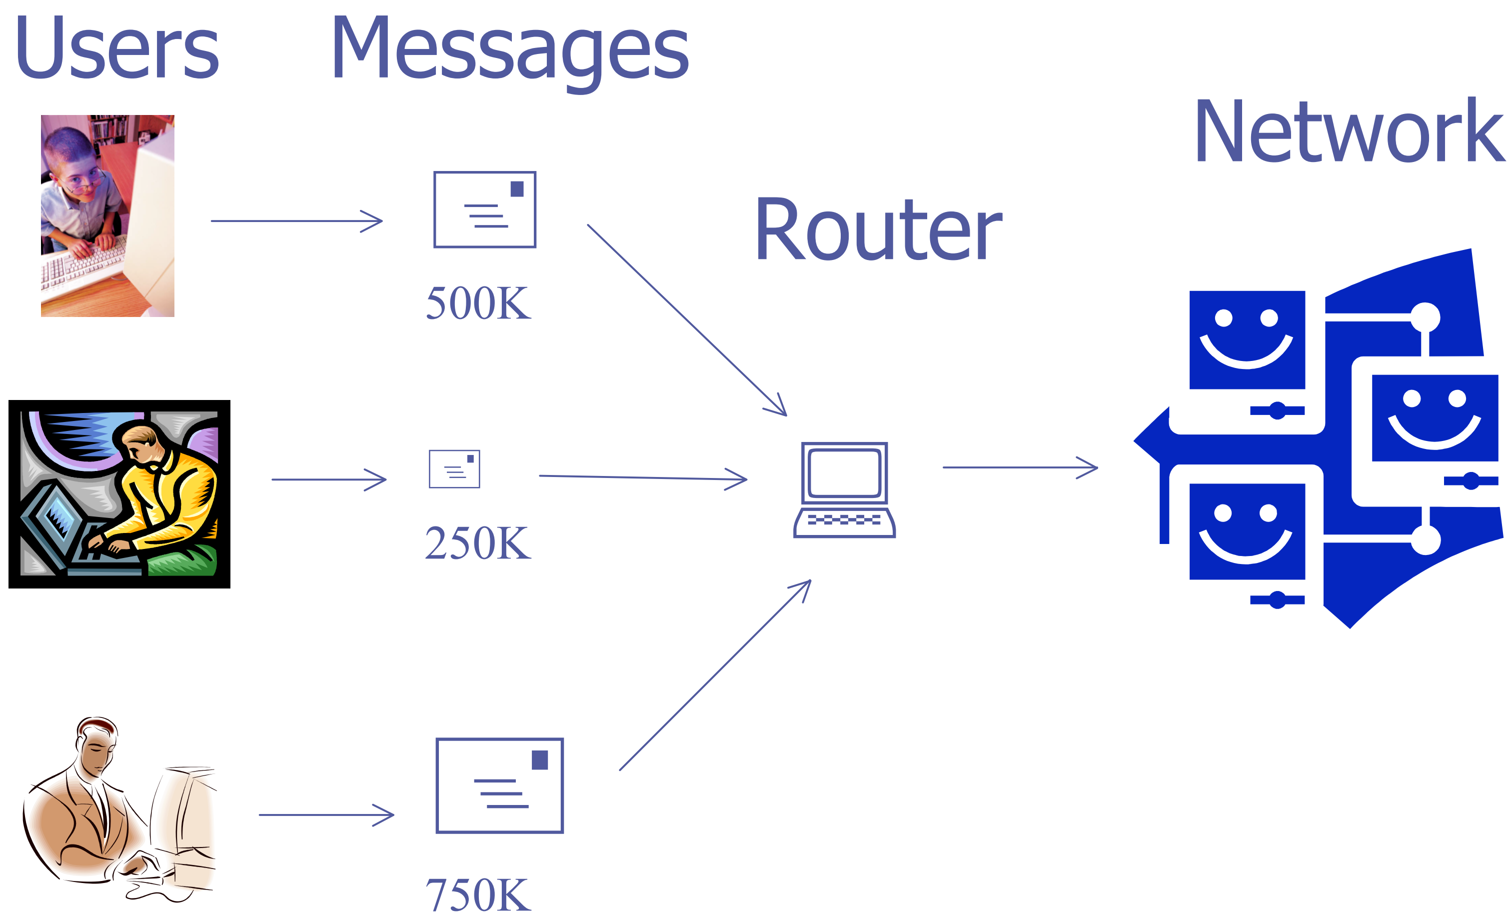

Decentralized Individual Choice

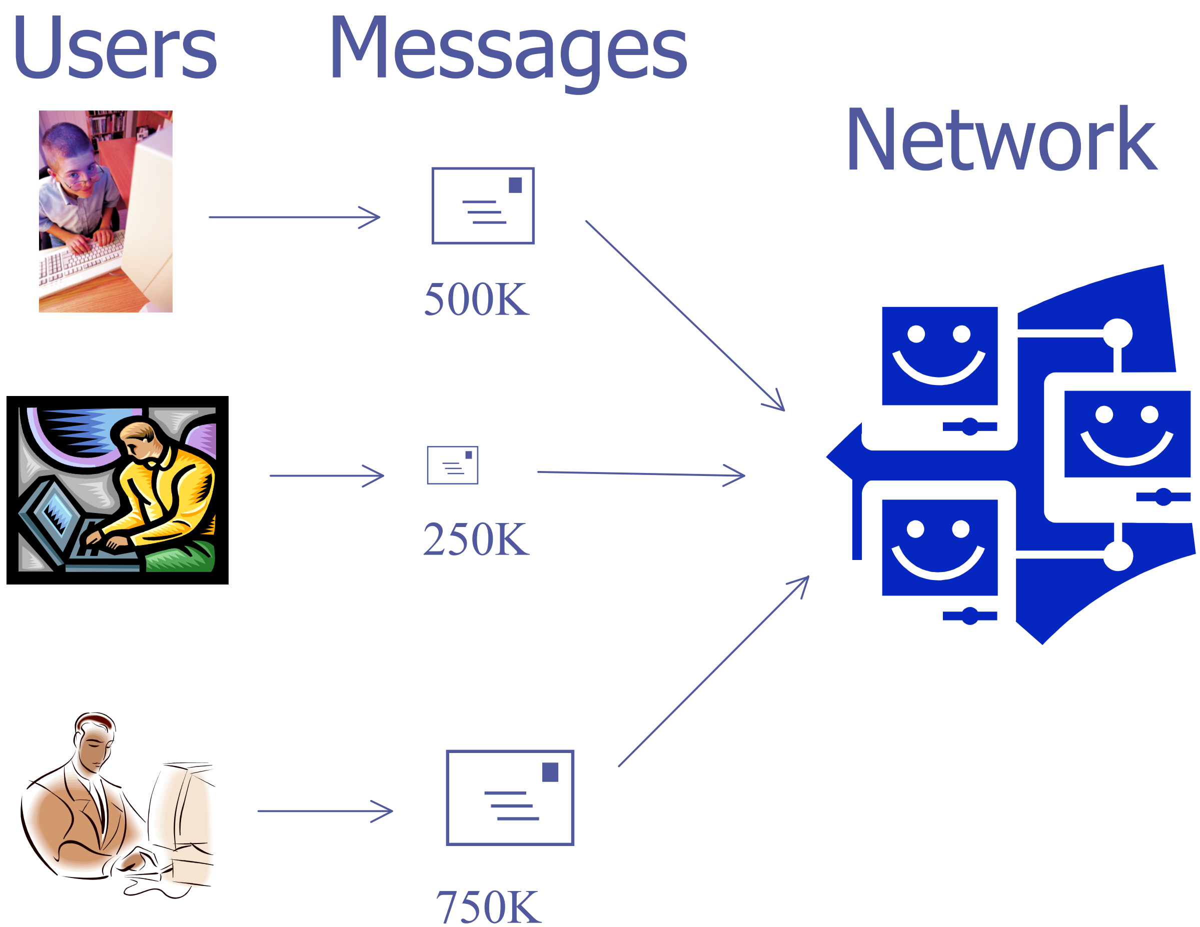

Motivation for Game-Theoretic Modelling

Use game theory to predict outcome of "selfish" user choices (Nash equilibrium)

↓

1. Assess "price of anarchy"
2. Improve network design/protocols

Outline

- Parallel Links Model
- Bayesian Parallel Links Game
- Intro to Evolutionary Stability
- ESS for Parallel Links Game
  - Characterization
  - Structural Conditions

Parallel Links Model

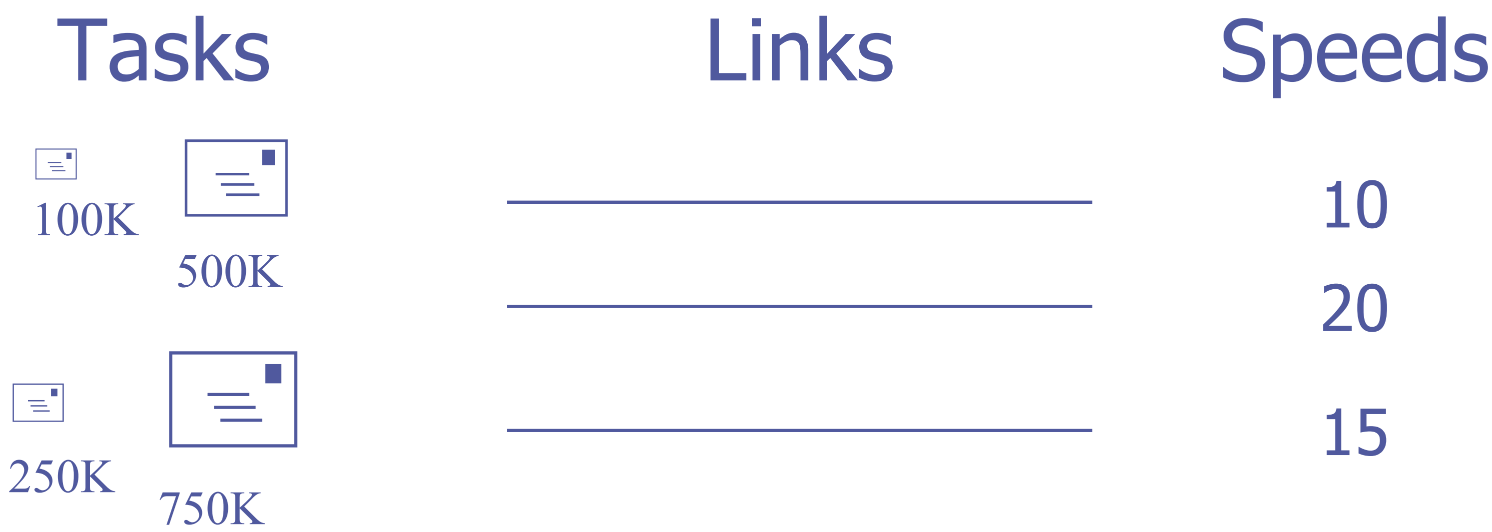

delay of task w on link l = w / (speed of l)

Parallel Links Model as a Game (Koutsoupias and Papadimitriou 1999)

1. Players 1,..,n with tasks $w_1, \ldots, w_n$
2. Pure strategy = (choice of) link
3. Fix choices $(w_1, l_1), \ldots, (w_n, l_n)$.

$\Rightarrow$ load on link $l = \Sigma_{i=1..n}\, w_i$ for $l_i = l$.

$\Rightarrow$ utility u_i for player i =

$$- \frac{\text{load on link } l_i}{\text{speed of link } l_i}$$

Bayesian Routing Game
(Gairing, Monien, Tiemann 2005)

- Agents are uncertain about tasks.
  - common dist. μ over tasks W
  - strategy ~ "program" *p* for routing tasks
  - $p(l|w)$ = probability that program *p* chooses link l when given task *w*.
- $u_i(p_1,\ldots,p_n) = \Sigma_{\text{task assignments}} <w_1,\ldots,w_n> \Pi_{j=1..n}\, \mu(w_j) \cdot u_i\, [(w_1,p_1|w_1),\ldots,(w_n,p_n|w_n)]$

Motivation for Evolutionary Analysis

1. Under "anarchy", we expect successful strategies to spread → evolutionary dynamics.
2. Highly successful predictions in biology.
3. Distinguishes stable from unstable equilibria.
4. May be useful in network design:
   see W. Sandholm's (2002) pricing scheme for traffic congestion. "evolutionary implementation in computer networks seems an important topic for future research".

Hawk vs. Dove As A Population Game

	Hawk (H)	Dove (D)
Hawk	-2,-2	6,0
Dove	0,6	3,3

- Assume a large population of agents.
- Agents are either hawks (H) or doves(D).
- We randomly draw 2 at a time to play.

Population Interpretation of Nash Equilibrium

1. Consider a population of agents with **frequency distribution** π.
   e.g. [H,H,H,H,H,H,D,D,D,D]
2. π is in equilibrium
   $\Leftrightarrow$ H does as well as D
   $\Leftrightarrow$ (π, π) is a **symmetric Nash equilibrum.**
3. (π, π) does **not** represent the choices of 2 players.
4. (π, π) says that both positions are drawn from the same population of agents with distribution π.

Stable vs. Unstable Equilibrium

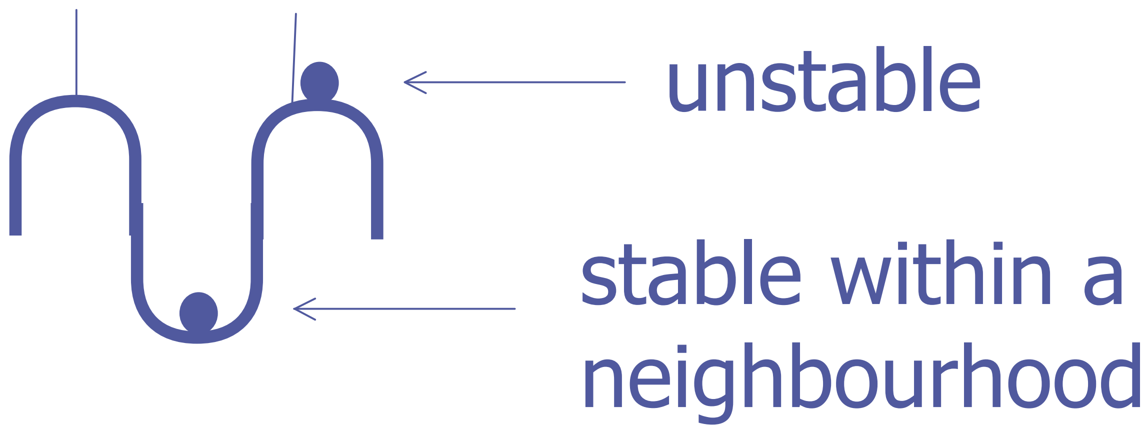

Evolutionarily Stable Strategies (ESS)

mixed population dist. = $(1-\varepsilon)\ \pi^* + \varepsilon\pi$

current dist π^*	mutant dist π
HHHHHH DDDD	H D
$10/12 = 1-\varepsilon$	$2/12 = \varepsilon$

← mutant plays mutant: $u(1/2,1/2;\ \pi)$

← incumbent plays mutant: $u(6/10,4/10;\ \pi)$

---

1. A distribution π^* is an ESS $\Leftrightarrow$ for all sufficiently small mutations π the incumbents in π^* do better in the mixed population than the mutants.

2. A distribution π^* is an ESS $\Leftrightarrow$ there is an ε^* such that for all sizes $\varepsilon < \varepsilon^*$ $u(\pi^*;\ (1-\varepsilon)\ \pi^* + \varepsilon\pi) > u(\pi;\ (1-\varepsilon)\ \pi^* + \varepsilon\pi)$ for all mutations $\pi \neq \pi^*$.

Characterization of ESS in Bayesian Routing Game *B*

Define:

- the load on link l due to strategy p:
load(p,l) = $\Sigma_{\text{tasks } w}\ \mu(w) \cdot p(l|w) \cdot w$

- the (marginal) probability of using link l:
prob(p,l) = $\Sigma_{\text{tasks } w}\ \mu(w) \cdot p(l|w)$

Theorem. A strategy p* is an ESS in *B* ⇔
for all best replies $p \neq$ p* we have
$\Sigma_{\text{links } l}$ [load(p*,l)-load(p,l)] · [prob(p*,l)-prob(p,l)] > 0

Intuition: to defeat mutation p:

- if load on link increases, use link less (- x -)
- if load decreases, use link more (+ x +)

Necessary Condition: Same Speed, Same Behaviour

Proposition. Let *B* be a Bayesian routing routing game with ESS p*. If two links l_1, l_2 have the same speed, then $p^*(l_1|w) = p^*(l_2|w)$ for all tasks *w*.

	Links	Speeds
w_1:50%, w_2:50%, w_3:0		10
w_1:50%, w_2:50%, w_3:0		10
w_1:0, w_2:0, w_3:100%		15

Necessary Condition: bigger tasks get faster links

Proposition. Let *B* be a Bayesian routing game with ESS p*. Suppose that

1. link 1 is faster than link 2
2. p* uses link 1 for task w1, link 2 for task w2.

Then w1 ≥ w2.

	Links	Speeds
w2= 10: 50%	———	10
w_1 = 20:100%, w2: 50%	———	15

Single Task: Unique ESS

Proposition. Let *B* be a Bayesian network routing game with just one task *w*.

1. *B* has a unique ESS p*.
2. If all *m* links have the same speed, $p^*(l_j|w) = 1/m$ is the unique ESS.

	Links	Speeds
w: 1/3		10
w: 1/3		10
w: 1/3		10

Strong Necessary Condition: No Double Overlap

- Fix a Bayesian network game *B*.
- Strategy p* **uses** link *l* for weight *w* ⇔ $p^*(l|w) > 0$.
- **Proposition.** Let p* be an ESS in *B*. Suppose that p* uses two distinct links $l_1 \neq l_2$ for task *w*. Then p* does not use both l_1 and l_2 for any other task w'.

	Links	Speeds
w_1:70%, w_2:30%	———	10
w_1:30%, w_2:70%	———	20
w_3:100%	———	15

>2 Tasks, Uniform Speeds: No ESS

Proposition. Let *B* be a Bayesian network game with >1 link, >1 task, all links the same speed. Then there is no ESS for *B*.

double overlap ↓

	Links	Speeds
w_1:50%,w_2:50%	———	10
w_1:50%,w_2:50%	———	10

Clusterings are typical ESS's

- Fix a Bayesian network game *B* with strategy p*.
- A link *l* is **optimal** for task *w* given p* ⇔ *l* minimizes *w*/speed(*l*) + load(*l*,p*).
- A strategy p* **clusters** ⇔ if two distinct links $l_1 \neq l_2$ are optimal for task *w*, then neither l_1 nor l_2 is optimal for any other task w' ≠ *w*.
- **Proposition.** If p* clusters, then p* is an ESS.

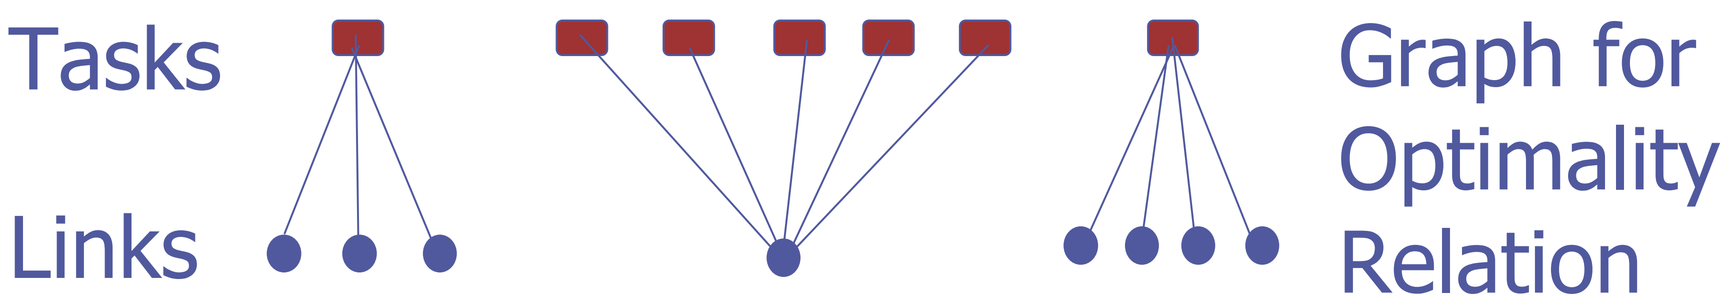

Does A Clustered Equilibrium Exist?

- Fix an assignment *A* of links to tasks.
- **Proposition.**

1. There is *at most one* clustered ESS p* whose clustering is *A*.
2. The candidate p* can be computed in polynomial time.
3. The question: is there a clustered ESS p* for a game *B*? is in NP.

Future Work

- Conjecture: if an ESS exists, it's unique.
- Conjecture: the "no double overlap" condition is sufficient as well as necessary.
- Computational Complexity and Algorithms for computing ESS's.

Conclusion

- ESS *refines* Nash equilibrium and defines *stable* equilibria.
- Analysis of evolutionary stability in Bayesian network games:
  - characterization of *successful mutations*
  - *structure* of stable *task/link allocations*.
- Finding:
  - evolutionary dynamics leads to formation of "niches" or **clusters** for task/link combinations.
  - Symmetric outcomes tend to be socially suboptimal.