A Game Theoretic Analysis of Selfish Content Replication on Graphs

Valentino Pacifici and György Dán
Laboratory of Communication Networks - ACCESS Linnaeus Center
KTH Royal Institute of Technology

1 Selfish Content Replication

- **Replication group**: nodes and users located near the nodes.
  1. Node i can replicate K_i objects from the set $\mathcal{O}$
  2. Objects are accessed by local or neighboring users, the demand for object $o \in \mathcal{O}$ of the users at node $i \in N$ is $w_i^o \in \mathbb{R}_+$
  3. The cost incurred by a user accessing an object is
     - α_i, if the object is replicated locally
     - β_i, if the object is replicated in a neighboring node
     - γ_i, otherwise
  4. No central authority ⇒ No optimal solution enforced.

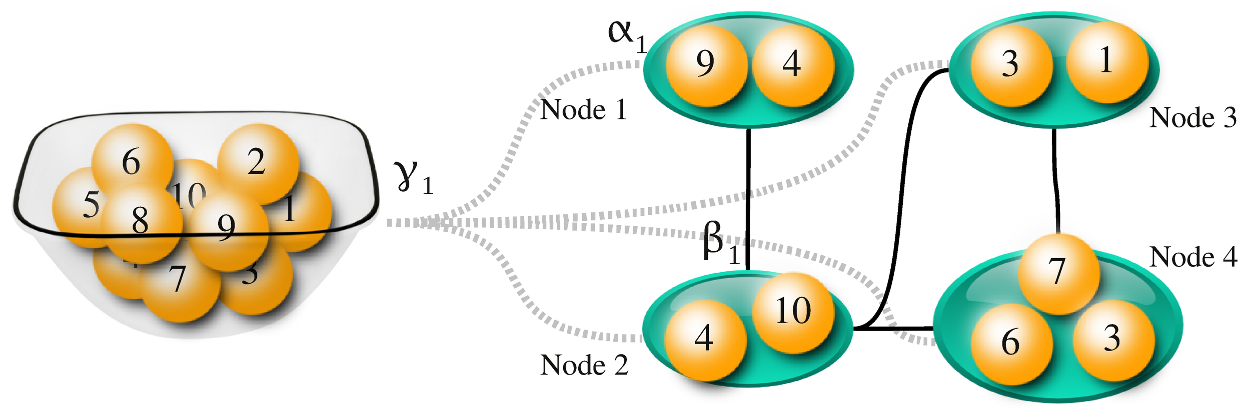

- Key questions
  1. Existence of a state where every node is satisfied with the object allocation
  2. Conditions under which the nodes reach such a state updating their decisions myopically

2 The model - Replication Game

- ***Replication game*** $< N, (\mathcal{R}_i), (U_i) >$:
  - N set of nodes,
  - $\mathcal{R}_i$ set of feasible replication configurations for player i.
  - U_i : utility function for player i.
- The nodes choose to replicate objects from the set $\mathcal{O}$
  - $r_i = (r_i^1, \ldots, r_i^{|\mathcal{O}|})$, where $r_i^o \in \{0,1\}$ is 1 if object o is replicated in node i
  - $r_i \in \mathcal{R}_i$ is feasible if $\sum_o r_i^o \leq K_i$
- Limited interactions between nodes modeled with a **social graph**
  - $\mathcal{N}(i)$ is the set of neighbors of node i.
- **Utility function**: sum of the cost savings $U_i(r_i, r_{-i}) = \sum_o U_i^o(r_i^o, r_{-i}^o)$

$$U_i^o(1, r_{-i}) = \begin{cases} w_i^o[\gamma_i - \alpha_i] & \text{if } \prod_{j \in \mathcal{N}(i)}(1 - r_j^o) = 1 \\ w_i^o[\beta_i - \alpha_i] & \text{if } \prod_{j \in \mathcal{N}(i)}(1 - r_j^o) = 0 \end{cases}$$

- *Improvement path*: a sequence of strategy profiles $r(0), r(1), ..,$ such that in every step t there is one player that strictly increases its utility by updating her strategy from $r_i(t-1)$ to $r_i(t)$.
- An update from $r_i(t-1)$ to $r_i(t)$ can be

Improvement step	Best reply step
$U_i(r_i(t+1), r_{-i}(t)) > U_i(r_i(t), r_{-i}(t))$	$U_i(r_i(t), r_{-i}(t-1)) \geq U_i(r_i(t-1), r_{-i}(t-1))$ $\forall\, r_i \in \mathcal{R}_i$

- ***Lazy improvement step*** of player i: an improvement step with minimal number of changes among all improvement steps that lead to the same utility.

3 Nash Equilibrium - Existence

- ***Nash Equilibrium***: a strategy profile r^* in which every player's strategy is a best reply to the other players' strategies

$$U_i(r_i^*, r_{-i}^*) \geq U_i(r_i, r_{-i}^*) \;\; \forall\, r_i \in \mathcal{R}_i, \;\forall\, i \in N.$$

- The following algorithm always ends in a NE:
  1. Play best replies in isolation
  2. Re-arrange the players according to the social graph
  3. Give a chance to play to every player
  4. Play in arbitrary order

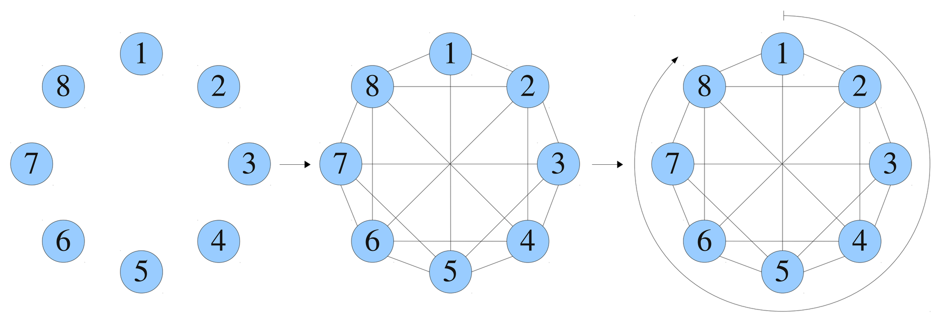

⇓

- Every graphical replication game possesses a pure strategy Nash equilibrium.

4 Nash Equilibrium - Convergence

- **Complete social graph**
  - Every best reply path in a replication game played over a complete social graph is finite.
- **Non-complete social graph**
  - The following graph topology allows a cycle in a best reply path:

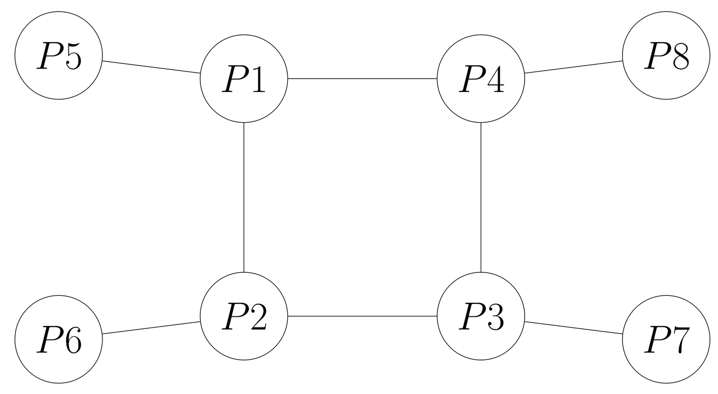

The cycle is made of the following sequence of best replies:

Player	*P1*	*P2*	*P3*	*P4*	P5	P6	P7	P8
Preferences	A<B	B<C	C<D	D<A	B	C	D	A
r(0)	A	B	D	A	B	C	D	A
r(1)	A	B	C	A	B	C	D	A
r(2)	B	B	C	A	B	C	D	A
r(3)	B	B	C	D	B	C	D	A
r(4)	B	C	C	D	B	C	D	A
r(5)	A	C	C	D	B	C	D	A
r(6)	A	C	D	D	B	C	D	A
r(7)	A	B	D	D	B	C	D	A
r(8)	A	B	D	A	B	C	D	A

- If $K_i = 1 \; \forall i \in N$

  From every strategy profile there exists a best reply path that leads to a NE in a finite number of steps.

  ⇓

  The game is weakly acyclic in best replies
- If $\beta_i = \alpha_i \; \forall i \in N$

  Every lazy improvement path is finite.

5 Fast convergence based on graph coloring

Objective: relax global synchronization requirement for convergence to NE

- If $\beta_i = \alpha_i \; \forall i \in N$

  If player i makes an improvement step at time t only if no neighboring player $j \in \mathcal{N}(i)$ makes an improvement step at time t, then every lazy improvement path is finite.

  ↓

  Plesiochronous better reply dynamic (PBRD)

Objective: maximize the convergence speed of the PBRD

1. Find a minimum vertex coloring of the social graph
2. Players with the same color update their strategy simultaneously

- Complexity: find the chromatic number of the graph ⇒ NP-hard
  - Efficient distributed algorithms exist
- Given a coloring, the number of steps required to reach the NE is significantly smaller than for ABRD

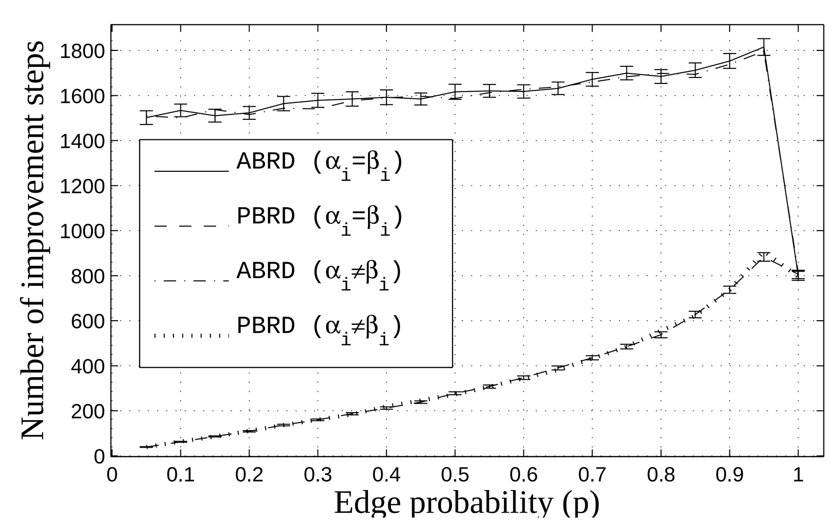

- PBRD significantly faster for sparse social graphs
- The existence of cycles when $\alpha_i \neq \beta_i$ does not affect the results
- Convergence properties different on a complete social graph than on a sparse graph, in accordance with the complexity of finding the optimal solution

References

[1] G. Dán, "Cache-to-cache: Could ISPs cooperate to decrease peer-to-peer content distribution costs?" *IEEE Trans. Parallel Distrib. Syst.*, to appear.

[2] B. Chun, K. Chaudhuri, H. Wee, M. Barreno, C. Papadimitriou, and J. Kubiatowicz, "Selfish caching in distributed systems: a game-theoretic approach," in *Proc. of ACM Symposium on Principles of Distributed Computing (PODC)*, July 2004.

[3] N. Laoutaris, O. Telelis, V. Zissimopoulos, and I. Stavrakakis, "Distributed selfish replication," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 12, pp. 1401–1413, 2006.