Games in the Interdomain Routing

imdea networks

Motivation

- Use ***game theory*** to present a different approach on well-known problems of the interdomain-routing protocol

- The design of BGP considered the AS to be ***obedient and trusted entities***
- Model the ASes as ***rational agents*** who act in a self-interested manner
- The interaction between these agents is dynamic and complex (asynchronous, repeated, with private information)

- Introduce ***incentive-compatibility*** in the interdomain-routing framework
- Study ***complexity and incentive-related issues*** in the interdomain routing, using the game theoretic models for a BGP system
- ***Provide incentives*** for ASes to adhere to the BGP prescribed behaviour

Outline

- Games and Game Theory
- Interdomain Routing – the networking approach
- Modeling Interdomain Routing as a Game
- Mechanism Design for Interdomain Routing
- Conclusions

Games and Game Theory

- ***Game Theory*** -> collection of analytical and modeling tools used to help us understand the interaction of decision-makers/players/strategic agents
  - It provides the necessary tools for predicting what might (and probably what should) happen when agents with conflicting interests interact
  - The agents are selfish entities which try to improve their outcome (behave in such a way to optimize their benefit)

- A ***game*** is…. made up of three important elements:
  - a set of **players** (the decision-makers)
  - a set of **actions** (the alternatives available to each player)
  - a set of **preferences** (the player´s evaluation of all the possible outcomes of the game)

[Robert Gibbons. *A primer in Game Theory*. Pearson Education, 1992]

Games and Game Theory

•Conventional games of strategy (poker, chess)

•Strategic negotiation (purchase a car)

•Daily group decision-making processes (where to go to lunch?)

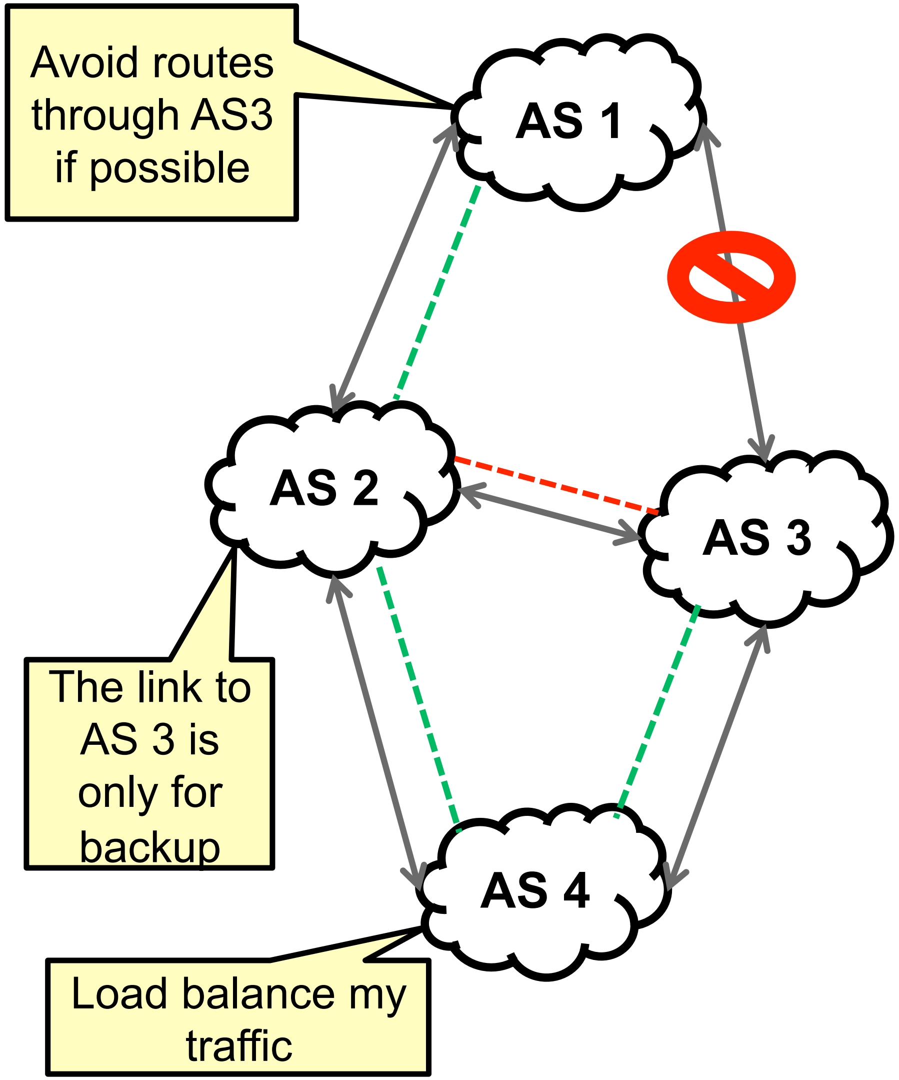

Games and Game Theory

- the ***outcome*** of the game → ***equilibrium point of the game***

  = is the set of nodes' *action choices* from which no node wishes to unilaterally deviate, because doing so would imply reducing its benefit

- Mainly , there are four **classes of games**:
  - Static games of complete information
  - Dynamic games of complete (and perfect) information
  - Static games of incomplete information
  - Dynamic games of incomplete information

- …for which we define four different **equilibrium concepts** :
  - Pure Nash Equilibrium
  - Subgame-perfect Nash Equilibrium
  - Bayesian Nash Equilibrium
  - Perfect Bayesian Equilibrium

Games and Game Theory

- ***Static games*** = single-round game in which players choose their actions simultaneously and are not aware of previous actions, after which they receive their payoff

- ***Dynamic games*** = multiple-round games (sequential games) in which players have information about the previous actions of other players

- ***Complete information*** = each player´s payoff function is common knowledge among all the players
  - *payoff function* = the function that determines the player's payoff from the combination of actions

- ***Incomplete information*** = players are uncertain about other players' payoff function
  - *perfect information* = all players know the previous actions taken by all other players

Games and strategies

- ***Strategy*** = complete plan of action for one player
  - *It specifies a feasible action for the player in every eventuality in which the player might be called upon to act*
- the combination of strategies chosen by the players defines the *"strategy profile"* of the game, which consequently determines the *outcome* of the game (the payoff for each player)

- ***The Nash equilibrium*** = the set of nodes' strategy choices from which no node wishes to unilaterally deviate and which maximizes the payoff of each player
- Formally, the strategy profile (s_1^* , s_2^* ,… s_n^*) is the ***Nash Equilibrium*** of a static game with complete information if ∀player i, s_i^* solves:

$$\max u_i(s_1^* ,\ldots, s_{i-1}^* ,s_i , s_{i+1}^* ,\ldots, s_n^*)$$

Games and Game Theory

- A game may have *more* than one equilibrium point
- The strongest equilibrium concept (that assumes the most general knowledge assumptions) is the perfect Bayesian equilibrium
- The *desired outcome* of any game is ***the perfect Bayesian equilibrium*** that is equivalent to:
  - a pure Nash equilibrium in a (induced) static game of complete information
  - a subgame-perfect Nash equilibrium in a (induced) dynamic game of complete (and perfect) information
  - A Bayesian Nash equilibrium in a (induced) static game of incomplete information
  - A perfect Bayesian equilibrium in a (induced) dynamic game of incomplete information

A normal-form game and the Nash Equilibrium

The Prisoners' Dilemma

Normal-form representation of the game*:* $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$

- ***Players****:* prisoner 1, prisoner 2
- ***Strategies*** *available to each player* => strategy spaces: {guilty, not guilty}
- *The* ***payoff*** u_i *received by each player i for each combination of strategies*

		Prisoner 2	
		Not Guilty	**Guilty**
Prisoner 1	**Not Guilty**	-2, -2	-5, -1
	Guilty	-1, -5	-3, -3

The **NE** does not guarantee the highest payoff!

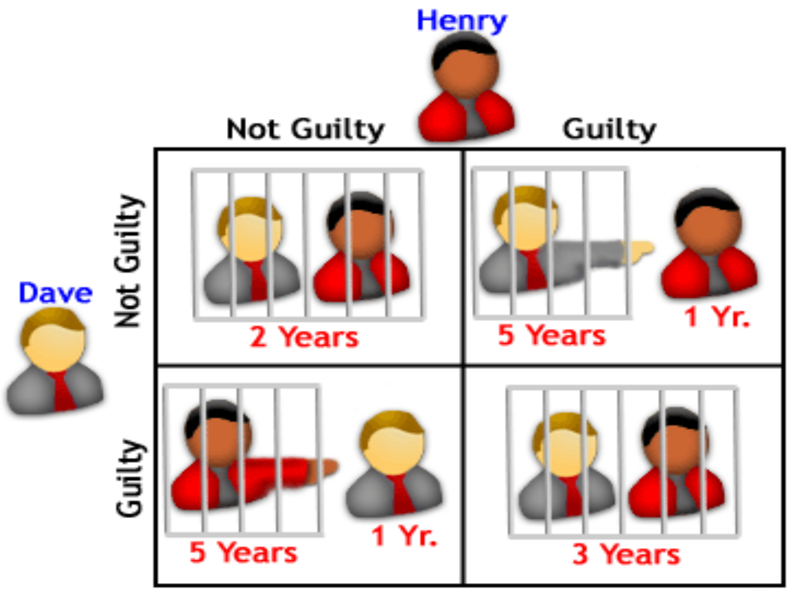

Outline

- Games and Game Theory
- Interdomain Routing – the networking approach
- Modeling Interdomain Routing as a Game
- Mechanism Design for Interdomain Routing
- Conclusions

Interdomain routing – the networking approach

- Establishing routes in the Internet between ASes is…***difficult***

- Current protocol: the **Border Gateway Protocol** (BGP)
  - Route choices ~ complex *routing policies*
  - *Routing policies* → **preferences** of the ASes on the available path choices
  - Ideally, the BGP system will use only ***stable paths*** to forward traffic, thus reaching a ***stable state*** = a state where ASes would not change their routes

  - However, routing policies are *not* coordinated
  - Routing policies may be conflicting=>

    BGP divergence =>

    persistent routing oscillations

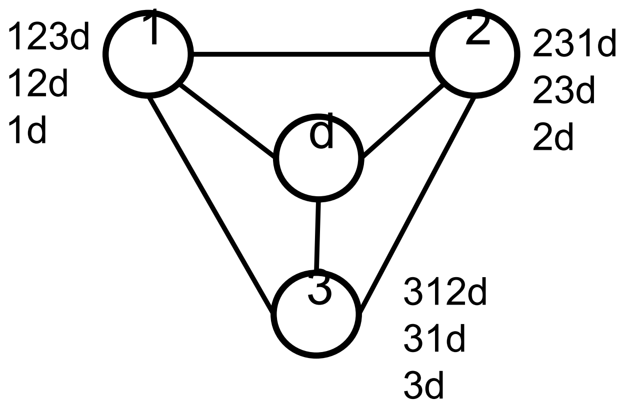

Fig 1. BAD GADGET

Interdomain routing – the networking approach

- The *route selection process*
  - distributed and asynchronous
  - triggered by advertisements and withdrawal of routes
- **State** of the network = the routes chosen by the ASes
- The BGP system ***converges*** when it arrives at a ***stable state*** after the ***activation sequence***
  - *activation sequence* = a (possibly infinite) sequence of activations
  - *activation* = a node applying the export policies of ASes on the neighbors, the import policies and the BGP route-selection
- A BGP system may not be able to *converge* to a stable state, even if one exists for that particular networking configuration

Interdomain routing – the networking approach

ABSTRACT MODEL OF BGP

- Model the network as an AS graph *G* = *<N, L>*
  - N: n source nodes (ASes) and a single destination node, d
  - L: physical links between Ases

- Every source-node *i* has a ***valuation function*** λ_i that assigns a non-negative value to each simple route from *i* to *d*
- Λ = the ranking (valuation) functions for each node *i*
- ***P*** = set of permitted paths for all the nodes in the network
- ***SPP = (G,*** Λ ***, P)*** – the Stable Path Problem

- **BGP** is a distributed manner of solving the SPP, assuming that the nodes are ***trusted and obedient parties***

Interdomain routing – the networking approach

- Important desired features of the BGP routing outcome:
  - ***Solvability***: the BGP system has at least one stable state
  - ***Safety***: the BGP system has a *stable state* and *converges* to it through any activation sequence
  - ***Robustness***: a network failure does not affect the *safety* of the routing outcome

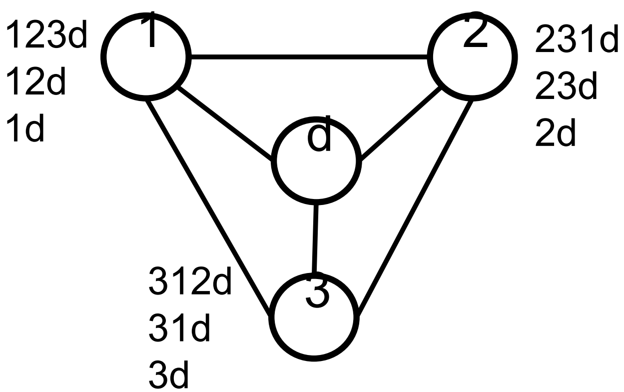

Fig 2.a. BAD GADGET

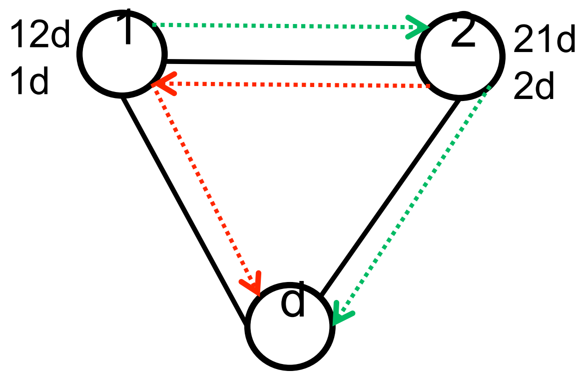

Fig 2.b. DISAGREE

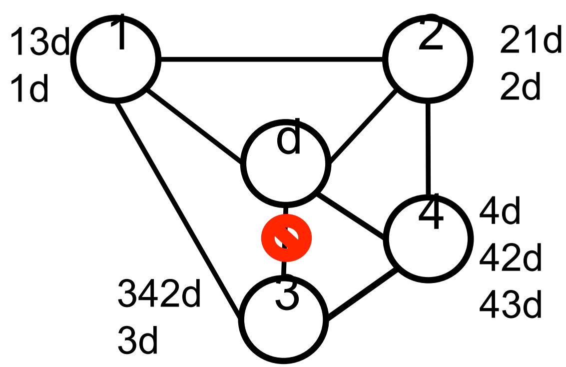

Fig 2.c. BAD BACKUP

BGP Solvability

- *Solving* the SPP ~ assign ***stable paths*** to each node in the AS graph that complies with the order of preferences
- Is ***any*** instance of SPP solvable?
- ***Offline static analysis*** => central entity analyzes the routing policies to verify that they do not contain *conflicts* that could lead to protocol divergence and compute the routing tree
- This implies that all ASes disclose their private routing policies
- The solution to the SPP is computed by a central entity that has ***complete information*** about the network

COMPLEXITY OF SPP SOLVABILITY

DETERMINING WHETHER A ***SOLUTION*** FOR THE SPP EXISTS IS NP-HARD

[Timothy G. Griffin and Gordon Wilfong. An analysis of BGP convergence properties. *SIG-COMM Computer Communication Review, 29 (4):277-288, 1999]*

BGP Safety

NO DISPUTE WHEEL

- Broadest condition known to guarantee BGP convergence to a stable solution
- $U = (u_0 \; u_1 \ldots\ldots u_{k-1})$ – k nodes (*pivot-nodes*)
- $R = (R_0, R_1\ldots\ldots R_{k-1})$
- $G = (G_0 , G_1 \ldots\ldots G_{k-1})$

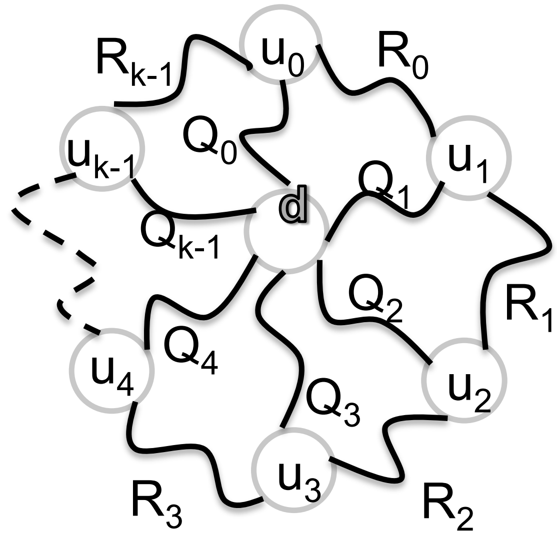

Fig. 3. A Dispute Wheel

- In a such a structure, it must hold that:
  - Each route Q_0 starts at u_i and ends at at the destination node *d.*
  - Each route R_i start at node u_i and ends at node u_{i+1} .
  - $v_i(Q_i) \leq v_i(R_i \, Q_{i+1})$ – cyclic interdependence
- The *sufficient* condition for BGP safety: ***No Dispute Wheel***

[T.G. Griffin, F.B. Sheperd, G, Wilfong. The Stable Path Problem and Interdomain Routing. *IEEE/ACM Transactions on Networking* , vol 10, no 2, April 2002]

Outline

- Games and Game Theory
- Interdomain Routing – the networking approach
- Modeling Interdomain Routing as a Game
- Mechanism Design for Interdomain Routing
- Conclusions

Modeling Interdomain Routing as a Game

- The network → G = *(N, L)*
  - *N = n source-nodes and an unique destination node, d*
  - *L = physical links in the network*
  - ∀ player *i*, it has a *routing policy* with two components:
    - *Valuation function -* v_i(R)
    - *Export policy*

- ***ONE-ROUND GAME***:
  - full-information static (non-sequentional) game
  - players = the ASes (the nodes in graph G)
  - strategy = chose outgoing edge (*i*'s choice to forward traffic)
    - Choices are simultaneous
  - The node's payoff is:
    - v_i(R), if R is a route from *i* to *d* induced by the nodes' choices
    - *zero* otherwise

Modeling Interdomain Routing as a Game

- The **SPP** – modeled as a ***game***!
- the pure Nash equilibrium = the stable path assignment (stable state)
- *ONE-ROUND GAME* = BGP network with a central unity that has complete information

Determining whether a *pure Nash equilibrium (or more)* in the One-Round game exists is **NP-hard.**

[Timothy G. Griffin and Gordon Wilfong. An analysis of BGP convergence properties.

SIG-COMM Computer Communication Review, 29(4):277-288, 1999]

BGP and THE CONVERGENCE GAME

THE CONVERGENCE GAME

- *dynamic (sequential) game* with an infinite number of rounds
- *incomplete* and *imperfect* information

- *scheduler* → model the asynchronous behavior of BGP
  - decides which players participate in each round (which nodes are activated)
  - delaying update messages
  - removing links and nodes from the AS graph
  - implements the fair activation sequence

- the nodes (the players) are ***strategic agents***
- the *type* of the player = the valuation function (private information)

BGP and THE CONVERGENCE GAME

- a node´s <u>*strategy*</u> is its choice of outgoing edge for forwarding traffic
  - executing BGP is a strategy

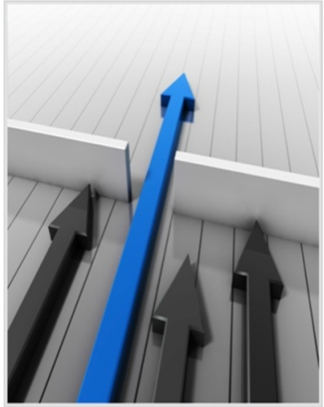

- BGP <u>*action space*</u>:
  - read update messages announcing routes
  - choose a single outgoing edge to forward traffic
  - announce simple routes to all neighbors

- A route is <u>*stable*</u> if from some route onwards every node in the route forwards traffic to the *same* next-hop on that route

- the <u>*payoff*</u> of the players is:
  - v_i(R), if R is a route from *i* to *d* induced by the nodes' choices
  - *zero* otherwise

Modeling BGP as best-reply dynamics

- *strategy profile* = best-reply dynamics (performing BGP)
- *executing* best-reply dynamics => find the Nash equilibrium of a given game

How do *best-reply dynamics* work?

- start with an arbitrary strategy profile
- in each round, some players switch their strategies to be *the best reply* to the current strategies of the other players
- if the process converges, then the pure Nash equilibrium of the game is reached

Games

Column Player

Row Player

	movie	opera
movie	2,1	0,0
opera	0,0	1,2

Pure Nash Equilibria and Best-Replies

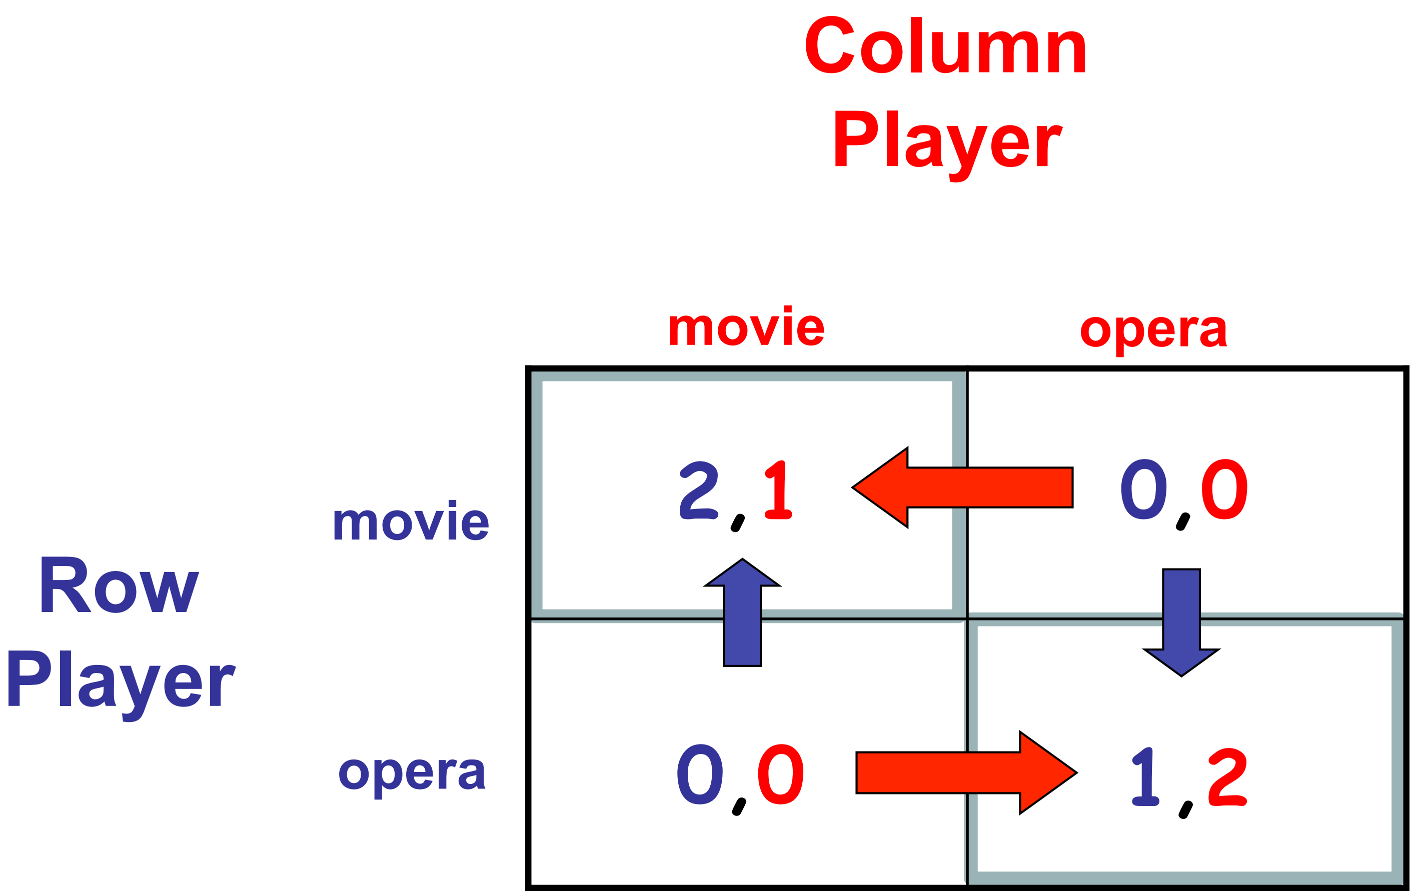

Best Reply Dynamics

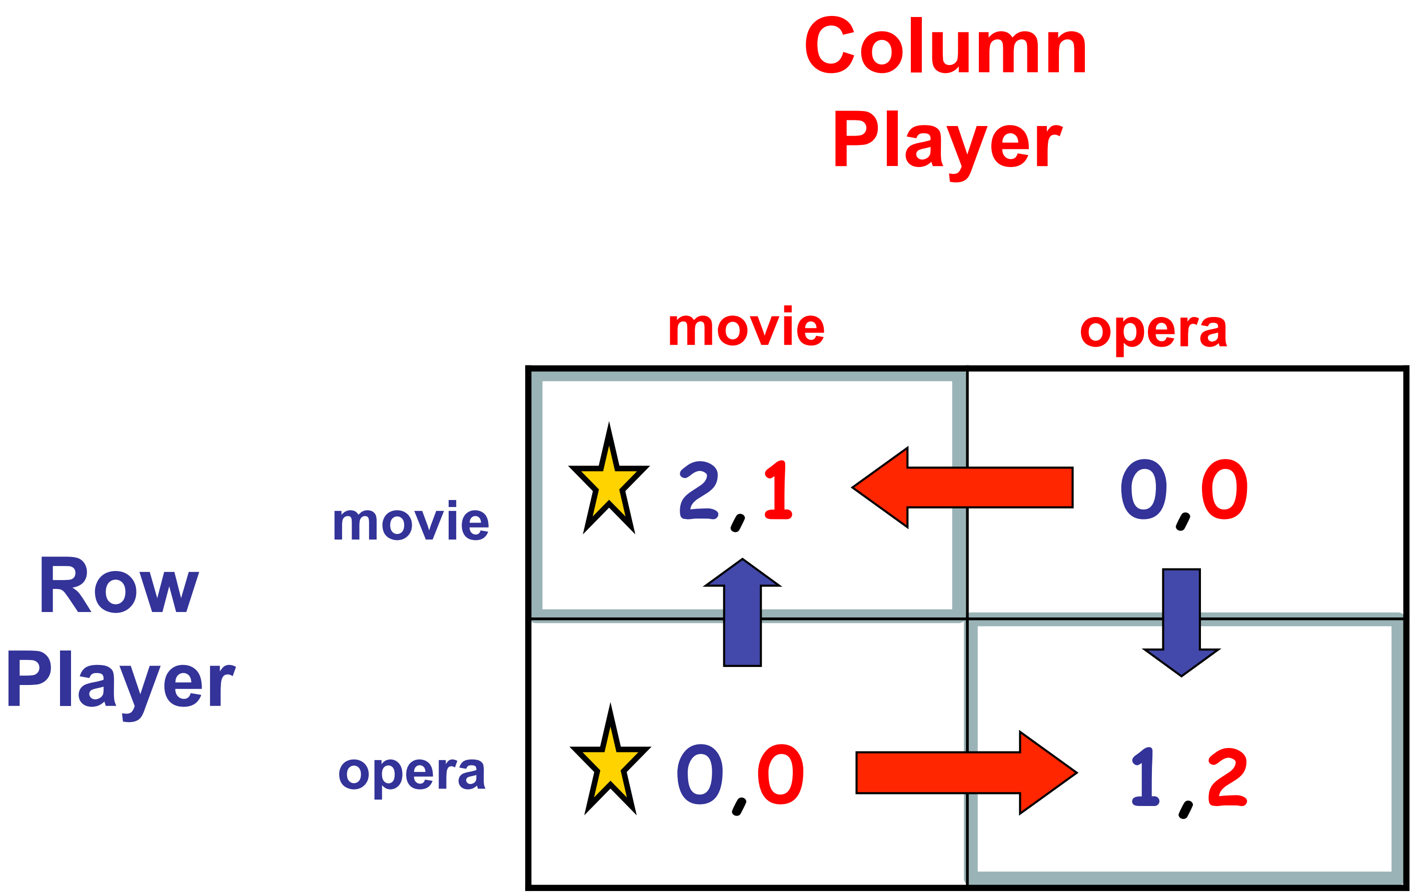

But…

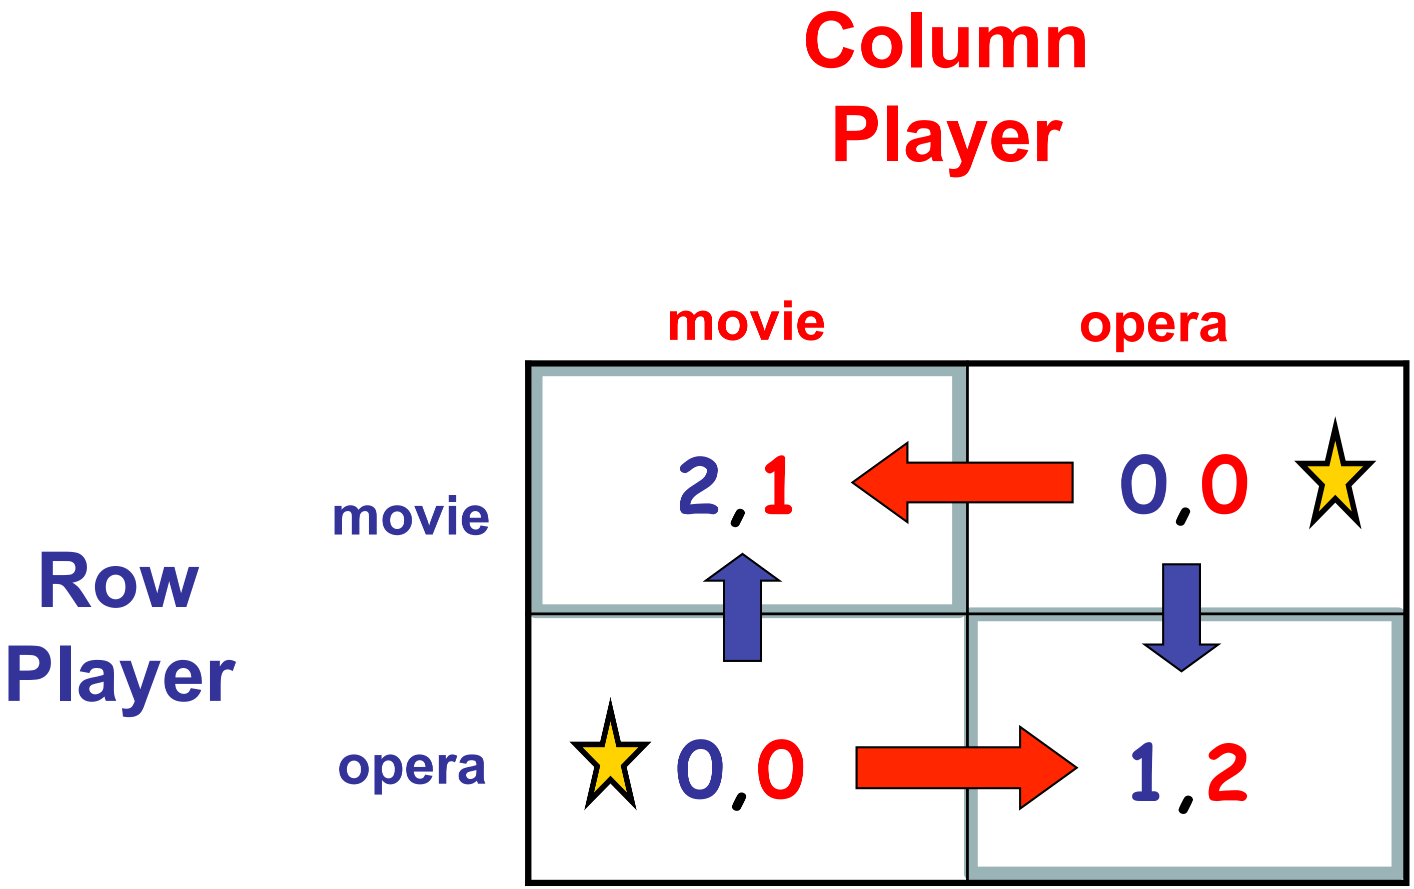

Modeling BGP as best-reply dynamics

- The outcome of *best-reply dynamics* in the CONVERGENCE GAME is the unique pure Nash equilibrium of the induced complete information ONE ROUND GAME

- IF TWO PURE NASH EQUILIBIRA EXIST IN THE ONE-ROUND GAME, THE BEST-REPLY DYNAMICS CAN POTENTIALLY **OSCILLATE** IN THE CONVERGENCE GAME.

  [Aaron D. Jaggard, Michael Shapira and Rebecca N. Wright.Towards a Unified Approach to (In)Decision: Routing, Circutits, Consensus, and Beyond. *Working paper., 2009.]*

Incentive-compatibility and BGP

- *BGP-compliant strategy* = strategy that obeys the rules of the protocol
- Until now, nodes were considered ***trusted and obedient parties,*** players that followed the prescribed strategy-profile (BGP)

- ASes are owned by *selfish economic entities* with very different (and conflicting) interests

- Do all ASes adhere to BGP?
- Incentives to *deviate* from the prescribed BGP behaviour?

- A node is said to *deviate* from BGP (*manipulate* BGP) if is does not follow BGP (does not have a BGP-compliant strategy)

Outline

- Games and Game Theory
- Interdomain Routing – the networking approach
- Modeling Interdomain Routing as a Game
- Mechanism Design for Interdomain Routing
- Conclusions

Incentive-compatibility and BGP

- Nodes = ***rational*** agents => *strategic behaviour* (nodes try to improve their payoff by deviating from the prescribed behaviour)

- ***Incentive-compatibility:*** intuitively means that any player would prefer *telling the truth* about its private information rather than any possible *lie*, since this will give him higher payoff
  - No *unilateral* deviation from BGP by any AS can strictly improve the routing outcome of that AS

- Forms of manipulation:
  - Lying about individual preferences
  - Reporting inconsistent information
    - Announcing one route, and using another
  - Announcing non-existing route
  - Denying routes
  - Etc…

Incentive-compatibility and BGP

THEOREM: Best-reply dynamics are ***not*** *incentive-compatible* in ex-post Nash even if *No Dispute Wheel* holds.

[H. Levin, M. Schapira, and A. Zohar. *Interdomain routing and games*. In ACM Symposium on Theory of Computing, May 2008]

PURE NASH EQUILIBRIUM ← nodes are familiar with the routing policies of ***all*** the other nodes in the network => strong knowledge assumptions

EX-POST NASH EQUILIBRIUM = each node obtains at least as great a utility by executing the actions in the prescribed strategy profile rather than some other strategy, regardless of the private information of all the other nodes

[J. Shneidman and D.C. Parkes. Specification Faithfulness in networks with rational nodes. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing (PODC 2004), pages 88{97. ACM, 2004*]

The Gao-Rexford constraints

The previous statement is true even if the network is *consistent* with the Gao-Rexford constraints**.**

THE GAO-REXFORD CONSTRAINTS:

- Capture the economic aspects of the Internet
- Suggest constraints and routing policies that are naturally induced by the business relationships

1. **Prefer *customer-routes* to *provider* or *peer* routes**
2. **There are no *customer-providers cycles* in the AS graph (no node is its indirect customer)**
3. **A node *a* only exports to node *b* paths through node *c* if at least one of the nodes *b* or *c* is a customer of node *a***

[Lixin Gao and Jennifer Rexford. Stable Internet routing without global coordination. *IEEE/ACM Transactions on Networking, 9(6):681{692, 2001]*

Incentive-compatibility and BGP

- **Incentive-compatibility in ex-post Nash** → *no node will deviate from the adopted strategy profile even if it would know the private information of the other players*
- **knowledge assumptions**: the rationality of the node

Fig. 3.a. Network for which the Gao-Rexfort constraints hold

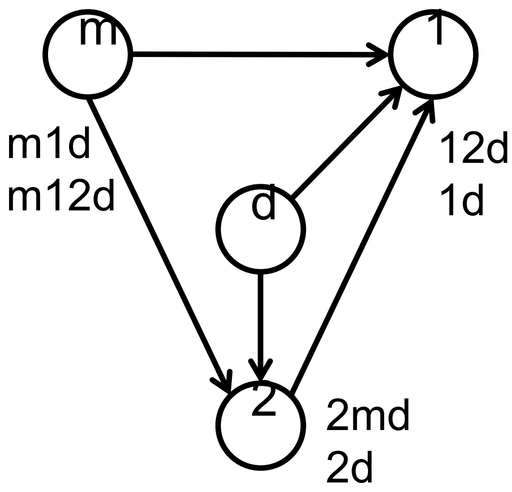

Fig. 3.b. Unique stable solution

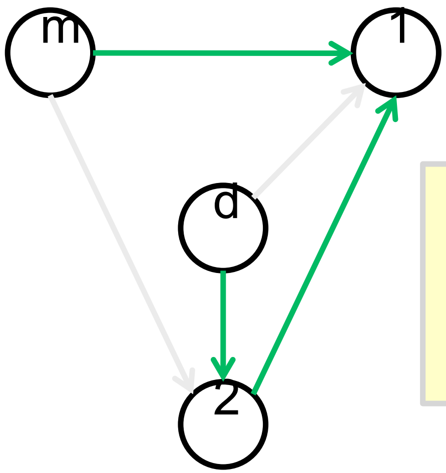

Fig. 3.c. Routing tree after manipulation

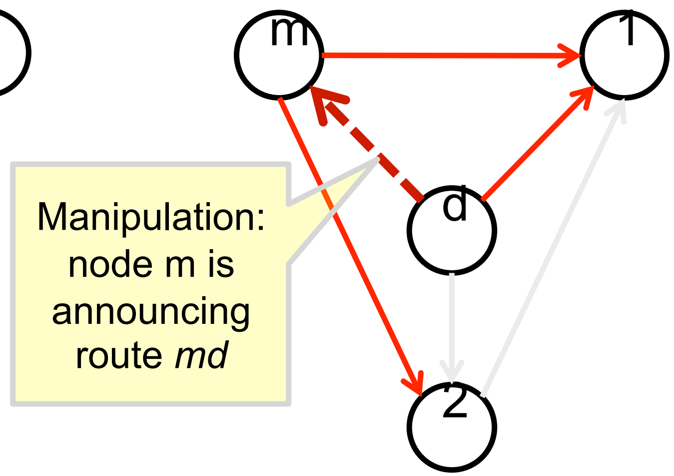

Customer → provider

Incentivizing Players in BGP

- *"The Internet is an equilibrium – we just have to identify the game"* *[S. Shenker]*

- The ***mechanism-design*** approach to interdomain routing
  - ***Mechanism design*** → creates games in which the *desired behaviour* emerges as an equilibrium of selfish participants, independently of the participants' unknown true preferences
  - *Desired outcomes* are: **Truthfulness**, individual rationality, social welfare

- **Incentivize** ASes to adhere to BGP
  - by restricting ASes' routing policies (*without money*)
  - with *monetary incentives* (VCG payments)

Collusion-proofness and BGP

THEOREM: If <u>*No Dispute Wheel*</u> and <u>*Policy Consistency*</u> hold, then BGP is ***incentive-compatible***, and even ***collusion proof***.

Policy Consistency:

if $v_k(R_1) > v_k(R_2)$,
then $v_i((i, k)\ R_1) > v_i((i, k)\ R_2)$,
$\forall\ R_1, R_2$

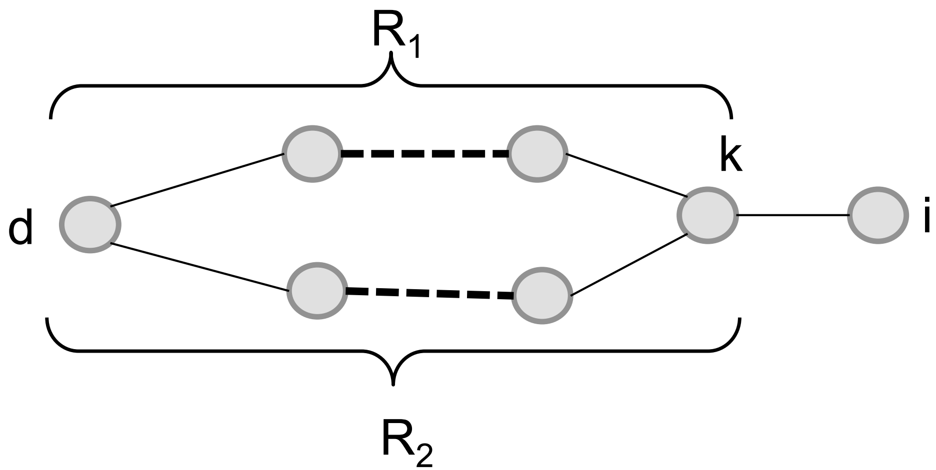

PROBLEM: *Policy consistency* is an unrealistic condition (too strong) !

[Joan Feigenbaum, Vijay Ramachandran, and Michael Schapira. Incentive-compatible interdomain routing. In *Proceedings of the 7th ACM conference on Electronic commerce, pages* 130-139, 2006]
[Joan Feigenbaum, Michael Schapira, and Scott Shenker. *Algorithmic Game Theory, chapter 14,* pages 363-383. Cambridge University Press, 2007]

Incentivizing Players in BGP

THEOREM: Best-reply dynamics are *incentive-compatible in ex-post Nash* if *No Dispute Wheel* and *Route Verification* hold.

Moreover, BGP is also *collusion-proof* in these settings.

- provides incentives *without* monetary transfers
- combine *protocol security* and *incentives*

- *Route verification* = a node can verify whether a route announced by a neighboring node is indeed available to that particular node

- *Route verification in the real BGP* => **S-BGP** (integrating security in the interdomain routing with cryptographic or other means)

BGP and the Convergence Game

UNTIL NOW WE HAVE LEARNED THAT…

- BGP is not incentive-compatible even if No Dispute Wheel holds

- A modification of BGP (Route Verification) is incentive compatible
- BGP is Pareto-optimal

BUT…

- Many other forms of misbehaving in the interdomain still possible….