

Best-Response Mechanisms (Application to BGP)

In this lecture we study the *Border Gateway Protocol (BGP)* which is the protocol responsible for routing packets on the Internet. There are two fundamental questions here:

1. Does this protocol stabilize? (convergence to an equilibrium)
2. Why do the Autonomous Systems implement it? (incentive compatibility)

A complicating factor will be the *asynchronous* nature of this problem (players do not necessarily move one by one).

1 Warm up

It may be useful to keep in mind the three games we saw in Lecture 1:

	-1	1
1		-1
	1	-1
-1		1

Matching Pennies

	1	0
2		0
	0	2
0		1

Battle of Sexes


	A	S
A	-2	-3
-2		0
S	0	-1
-3		-1

Prisoners' Dilemma

(Synchronous) Best-Response Dynamics: Players play their best response infinitely many times, one by one in a fixed order (round robin).

We know what happens for each of these games.

Example 1. Two nodes, 1 and 2, want to send traffic to another destination node d . Their strategy is to choose the *next hop* the traffic is sent to (one of the neighbors). The following picture shows the physical network and the preferences of each node (which path to use) near the corresponding node:



Each node prefers to reach d via the other node, but if they both send their own traffic to each other they fail (which is the least preferable option for both). ■

Question: What happens if the two nodes move (play) always simultaneously? What happens if node 1 plays “1 → 2” at each step (while the other node plays best-response)?

Best Response:

1. No convergence in *asynchronous* settings.
2. Not incentive compatible.

For which games this does not happen?

Asynchronous Best-Response Dynamics: At each step an adversary activates an arbitrary subset of players who best respond to the current profile (the adversary also chooses a starting strategy profile). The adversary must activate each player an infinite number of times.

The choice of the adversary and the “response strategies” of each player determine an infinite sequence

$$s^0 \implies s^1 \implies \dots s^t \implies \dots$$

If the game converges (after finitely many steps T we have $s^T = s^{T+1} = s^{T+2} = \dots$) then the utility of each player i is $u_i(s^T)$. If the game keeps “oscillating” then we consider an upper bound on what the player can get (the worst case for us and the best for the player) that is $\limsup_{t \rightarrow \infty} u_i(s^t)$.

Base game G	\implies Repeated game G^*
$s_i \in S_i$	response strategy $R_i(\cdot) \in S_i$
$u_i(s)$	total utility $\Gamma_i := \limsup_{t \rightarrow \infty} u_i(s^t)$

Definition 2. Best-response are **incentive compatible** for G if repeated best-responding is a Nash equilibrium for the repeated game G^* , that is, for every i

$$\Gamma_i \geq \Gamma'_i$$

where Γ_i is the total utility when all players best respond and Γ'_i is the total utility when all but i best respond (starting from the same initial profile s^0 and applying the same activation sequence).

2 “Nice” Games

Consider this game (with a unique PNE):

		Player 2	
		A	B
Player 1	A	1 2	0 0
	B	0 3	2 1

Best response works as follows

$$(A, A) \xrightarrow{\text{Player 1}} (B, A) \xrightarrow{\text{Player 2}} (B, B) \xrightarrow{\text{Player 1}} (B, B) \xrightarrow{\text{Player 2}} (B, B) \dots \implies (B, B)$$

Player 1 improves if he/she does not best response (keep playing A):

$$(A, A) \xrightarrow{\text{Player 1}} (A, A) \xrightarrow{\text{Player 2}} (A, A) \xrightarrow{\text{Player 1}} (A, A) \implies \dots \implies (A, A)$$

Convergence but no incentive compatibility

Exercise 1. For the following game

		Player 2	
		A	B
Player 1	A	1 1	1 1
	B	1 1	1 1

find best response strategies that *never converge* (keep oscillating between different profiles). Find other best response strategies for which we *do have convergence*. ■

Two intuitions/ideas:

1. Introduce tie breaking rule.
2. Eliminate “useless” strategies.

2.1 Convergence

Consider this game

	A	B	C
a	1 2	0 0	0 0
b	2 1	-1 1	1 -1
c	-2 -1	1 -1	-1 1

Exercise 2. Prove that for this game best-response dynamics converge to a unique PNE.

Note that in the previous game no strategy is dominant and no strategy is dominated. Strategy C satisfies the following (weaker) definition:

Definition 3 (never best response (NBR)). A strategy $s_i \in S_i$ is a never best response (for tie breaking rule \prec) if there is always another strategy that gives a better payoff or that gives the same payoff but is better w.r.t. to this tie breaking rule: for all s_{-i} there exists $s'_i \in S_i$ such that one of these holds

1. $u_i(s_i, s_{-i}) < u_i(s'_i, s_{-i})$ or
2. $u_i(s_i, s_{-i}) = u_i(s'_i, s_{-i})$ and $s_i \prec_i s'_i$.

The following condition is enough to guarantee convergence:

Definition 4 (NBR-solvable). A game G is NBR-solvable if iteratively eliminating NBR strategies results in a game with one strategy per player. That is, there exists a tie breaking rule \prec , sequence p_1, \dots, p_ℓ of players, and a corresponding sequence of subsets of strategies E_1, \dots, E_ℓ such that:

1. Initially $G_0 = G$ and G_{i+1} is the game obtained from G_i by removing the strategies E_i of player p_i ;
2. Strategies E_i are NBR for \prec in the game G_{i-1} .
3. The final game G_ℓ has one strategy for each player (this unique profile is thus a PNE for G).

A sequence of players and of strategies as above is called an elimination sequence for the game G .

Exercise 3. Prove that the game described at the beginning of this section is NBR-solvable. Provide also a bound on the parameter ℓ .

In the analysis we consider rounds, where a round is a sequence of consecutive time steps in which each player has been activated at least once. For example,

$$\underbrace{\{1\}, \{2\}, \dots, \{n\}}_{\text{round 1}}, \underbrace{\{1, 2\}, \{1, 2\}, \dots, \{1, 2\}}_{\text{round 2}}, \underbrace{\{1, \dots, n\}, \{1, 2, \dots, n\}, \dots}_{\text{round 3}}, \dots$$

Lemma 5 (rounds vs subgames). Let p_1, \dots, p_ℓ be the players of any elimination sequence for the game under consideration. Suppose that players p_1, \dots, p_k always best respond (according to the prescribed tie breaking rule \prec). Then, for any initial profile and for any activation sequence, every profile after the k^{th} round is a profile in the subgame G_k .

Before proving the lemma we observe that it implies convergence:

Theorem 6 (convergence). For NBR-solvable games best response (according to the prescribed tie breaking rule \prec) converge even in the asynchronous case.

Proof. Take $k = \ell$ and observe that G_ℓ contains only one profile. □

PROOF OF LEMMA 5. Denote by $round_j$ the last time step of the j^{th} round in the activation sequence. Obviously for any t we have $s^t \in G_0 = G$. Now consider $t \geq round_1$ and observe that, since player p_1 has been activated at least once the corresponding strategy satisfies ¹

$$s_{p_1}^t \notin E_1$$

which is equivalent to $s^t \in G_1$ for all $t \geq round_1$.

To prove the analogous for player p_2 we observe that, in the 2^{nd} round player p_2 is activated and, since $s^t \in G^1$ and since p_2 plays best response, for $t \geq round_2$ we have $s_{p_2}^t \notin E_2$. Since we have previously proved $s_{p_1}^t \notin E_1$, this implies $s^t \in G_2$ for $t \geq round_2$.

We can then continue and prove, by induction, that after the k^{th} round player p_k does not play any strategy in E_k and thus $s^t \in G_k$ for all $t \geq round_k$. \square

2.2 Incentive Compatible

Look (again) at this game:

		Player 2	
		A	B
Player 1	A	1 2	0 0
	B	0 3	2 1

Bad for incentive compatibility: The unique PNE does not give Player 1 the highest possible payoff he/she can get in this game.

Definition 7 (NBR-solvable with clear outcome). *A NBR-solvable game G has a clear outcome if there exists a tie breaking rule \prec such that the following holds. For every player i there exists an elimination sequence consisting of players $p_1, \dots, p_a, \dots, p_\ell$ and strategies $E_1, \dots, E_a, \dots, E_\ell$ (according to Definition 4) such that,*

1. p_a denotes the first appearance of i in the sequence, that is,

$$p_a = i \neq p_1, p_2, \dots, p_{a-1};$$

2. in the corresponding subgame

$$G_{a-1} = G \setminus (E_1 \cup E_2 \cup \dots \cup E_{a-1})$$

the PNE s^* is globally optimal for i , that is,

$$u_i(\hat{s}) \leq u_i(s^*) \quad \text{for all } \hat{s} \in G_{a-1}.$$

(Recall that s^* is the unique profile in the final subgame G_ℓ .)

¹More in detail, if the player is activated at time t' then at time $t' + 1$ his/her profile is not in E_1 ; If the player is not activated at time t' then her strategy at time $t' + 1$ remains the same.

Theorem 8 (incentive compatibility). *For NBR-solvable games with clear outcome best response (according to the prescribed tie breaking rule \prec) are also incentive compatible.*

Proof. Compare the case in which all players best respond to the case in which player i does not best respond (while the others best respond). In particular, we consider the two sequences of profiles

$$\begin{aligned} \text{All best respond: } s^0 &\implies s^1 \implies s^2 \implies \dots \implies s^* \implies s^* \dots \\ \text{All but } i \text{ best respond: } s^0 &\implies \hat{s}^1 \implies \hat{s}^2 \implies \dots \implies \hat{s}^t \implies \hat{s}^{t+1} \dots \end{aligned}$$

We want to show that starting from some finite T the utility of i in the second sequence is not better than the “final” utility in the first sequence:

$$u_i(\hat{s}^t) \leq u_i(s^*) \quad \text{for all } t \geq T \tag{1}$$

This implies $\hat{\Gamma}_i \leq \Gamma_i$ that is the incentive compatibility condition (see Definition 2). Consider the elimination sequence of definition of NBR-solvable game (Definition 7) and let $p_k = i$ be the first occurrence of i in the sequence (i.e. $i \neq p_1, \dots, i \neq p_{k-1}$):

Player:	p_1	\dots	p_{k-1}	i	p_{k+1}	\dots
NBR Strategies:	E_1	\dots	E_{k-1}	E_k	E_{k+1}	\dots
Current Game:	G_0	\dots	G_{k-2}	G_{k-1}	G_k	\dots

We know from Lemma 5 that after round $k-1$ the profile must be in the game G_{k-1} (since i does not appear in the elimination sequence before position k , all players p_1, \dots, p_{k-1} are different from i and thus they all play best response). Since the PNE s^* is globally optimal for i in this subgame, we have $u_i(s^t) \leq u_i(s^*)$ for all $t \geq \text{round}_{k-1}$. This proves Inequality (1) and thus the theorem. □

2.3 Best-Response Mechanism Framework

This is how these results can be used.


Best-Response Mechanism

Design a game (rules) that satisfy the two conditions of Definition 7.


We shall see several mechanism design problems in next lecture. Now we go back to our initial problem.

3 BGP Games

Several Autonomous Systems are connected to each other:



The Border Gateway Protocol (BGP) specifies how to forward traffic. Each node in this graph chooses neighbor (“next hop”):



3.1 BGP “in Theory”...


BGP game (static version)

1. Players = Nodes
2. Strategies = Neighbors
3. Strategy profile = Set of paths (or loops)
4. Utilities = Order over the paths connecting i to d

$$P_1 \prec_i P_2 \prec_i \dots \prec_i P_k$$


and any path \emptyset which does **not** connect i to d is strictly worse: $\emptyset \prec_i P_1$.

Consider this instance:



There is no PNE.

Dispute Wheel: every node prefers routing over the next one in the “wheel”



with preferences

$$Q_i \prec_{w_i} R_i Q_{i+1}$$

no convergence + no incentive compatible


3.2 ...BGP “in Practice”

Gao-Rexford Model \implies No Dispute Wheel \implies BPG Converges Incentive Compatible

There are two types of **commercial relationships** between ASs:




Each node i classifies paths according to its commercial relationship with the neighbor in the path (first hop): (1) **customer paths**, (2) **peer paths**, and (3) **provider paths**:



The top path is a customer path because the first hop is from i to a customer of i . Similarly, we have peer and provider paths (all neighbors of i can be grouped into these three classes). The preferences of each node i respect this classification:


Gao-Rexford model (first version):
 (GR1) provider paths \prec peer paths \prec customer paths

Dispute wheel is still possible:



Gao-Rexford model (second version):
 (GR1) provider paths \prec peer paths \prec customer paths
 (GR2) transit traffic to/from my customers only


Consider this path:




It may happen that node j does not allow **transit traffic** from node i :

- Node j chooses k as its next hop, but
- Node j does not forward the traffic coming from i to node j

There are **only two cases** where a node j allows transit traffic:




In all other cases a node j does **not** allow transit traffic:



No transit \Rightarrow zero utility


If node j does not allow transit traffic from node i then any path $P = i \rightarrow j \rightarrow \dots d$ represents a “failure” for i which we denote with the symbol \emptyset . Such “failing” paths have always the lowest utility 0.

Example 9. Reconsider our previous example with all nodes having “peer-to-peer” relationships:



Now we consider the path “12d” as a **failure for node 1** because its traffic will not be forwarded by node 2, though node 2 is forwarding its own traffic to d. Therefore the **preferences** of node 1 must be as shown in the picture. A similar argument holds for the path “21d” with respect to node 2.

Exercise 4. Show that the following dispute wheel is still possible:



that is, these preferences do not necessarily violate conditions GR 1 and GR 2 (find the commercial relationships for which this is the case). ■

Gao-Rexford model (final version):
 (GR1) $\emptyset \prec$ provider paths \prec peer paths \prec customer paths
 (GR2) transit traffic to/from my customers only
 (GR3) no customer-provider cycles

(GR3) says that no AS is indirectly a provider of itself.

(GR1) can be rewritten in terms of utilities as

$$0 = u_i(\emptyset) < u_i(\text{provider-path}) < u_i(\text{peer-path}) < u_i(\text{customer-path})$$


for any provider-path, any peer-path and any customer-path of i .

3.3 Gao-Rexford \implies No Dispute Wheel

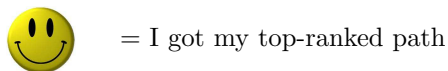
See Section A.1 and related exercises.

3.4 No Dispute Wheel \implies NBR-solvable with clear outcome


The key idea to construct an appropriate elimination sequence is to identify what we call “happy paths”:



where



Here is an example of “unhappy” path (not all players are happy):




A path $h_1 \rightarrow h_2 \rightarrow \dots \rightarrow h_l \rightarrow d$ in a subgame G_i is a **happy path** if this path gives the highest possible payoff to all of these nodes:

$$h_a \rightarrow h_{a+1} \rightarrow \dots \rightarrow h_l \rightarrow d$$

is h_a 's top ranked path among those that are available in the **subgame** G_i .

To see the idea of how happy paths give an elimination sequence:



The elimination sequence goes “**from right to left**”:

- 1) h_l eliminates all strategies other than “ $h_l \rightarrow d$ ” from the current subgame G_i and this gives us G_{i+1} . In this subgame G_{i+1} it is still true that the path is an happy path and thus h_{l-1} can eliminate all strategies other than “ $h_{l-1} \rightarrow h_l$ ”. We can continue until the first node in the happy path has eliminated all but the “ $h_1 \rightarrow h_2$ ” strategy.
- 2) In the resulting subgame we find another happy path and repeat the previous step until there are no happy paths that start with a node with at least two strategies.

Suppose at the end of this process we included all nodes:

Every node belongs to some happy path. (2)

Then the final subgame consists of a game with one strategy per player. At each step we eliminate strategies that give the node a non-optimal payoff in the current subgame. So the starting game is NBR-solvable with clear outcome.

3.4.1 No Dispute Wheel \Rightarrow Condition (2)

We show that if there is no happy path then there must be a Dispute Wheel. Given that there is no happy path, starting from a node w_0 its top ranked path is *not* an happy path:

$$TR_{w_0} = w_0 \rightarrow i_1 \rightarrow \dots \rightarrow w_1 \rightarrow i_a \rightarrow \dots \rightarrow i_l \rightarrow d$$

and w_1 is the rightmost node (closest to d) for which the subpath

$$w_1 \rightarrow i_a \rightarrow \dots \rightarrow i_l \rightarrow d$$


is not w_1 's top ranked available path which is instead

$$TR_{w_1} = w_1 \rightarrow j_1 \rightarrow \dots \rightarrow w_2 \rightarrow j_{a'} \rightarrow \dots \rightarrow j_{l'} \rightarrow d$$

where w_2 is (again) the rightmost node in this path for which the corresponding subpath is not top ranked for it (this because there is no happy path). Since there is no happy path this can go on until we get some w_k such that

$$TR_{w_k} = w_k \rightarrow n_1 \rightarrow \dots \rightarrow w_{k+1} \rightarrow n_{a''} \rightarrow \dots \rightarrow n_{l''} \rightarrow d$$

and w_{k+1} is one of the previously considered w_j 's. For instance, if $w_{k+1} = w_0$ then we get the Dispute Wheel



by setting $R_i Q_{i+1} := TR_{w_i}$. If $w_{k+1} = w_s$ then we get a smaller Dispute Wheel with nodes w_s, w_{s+1}, \dots, w_k .

BGP “in Practice” (Gao-Rexford model):
 YES convergence + YES incentive compatible

Recommended Literature

The best-response mechanism framework presented here can be found here:

- Noam Nisan, Michael Schapira, Gregory Valiant, and Aviv Zohar. Best-response mechanisms. In *Innovations in Computer Science (ICS)*, pages 155–165, 2011.
(including several applications that we present in the next lecture)


A more detailed description of the Gao-Rexford model and BGP along with a proof of convergence and incentive compatibility is in:

- Hagay Levin, Michael Schapira, and Aviv Zohar. Interdomain routing and games. *SIAM Journal on Computing*, 40(6):1892–1912, 2011.

A Omitted parts (Exercises)

A.1 Gao-Rexford \implies No Dispute Wheel

We show that the network cannot contain nodes and paths that form a dispute wheel. We prove the result only for these simpler wheels (paths P_i and Q_i consist of a single link):



Recall that \emptyset denotes any path that does not allow w_i to reach d (in particular if w_{i+1} does not allow transit traffic from w_i) and the utility is $u_{w_i}(\emptyset) = 0$. This and the preferences of the nodes

$$Q_i \prec_{w_i} R_i Q_{i+1}$$

imply that w_{i+1} must allow transit traffic from w_i . This is possible only in one of these two cases (GR2):



Exercise: show that in either case we must have a dispute wheel.

Exercises


(during the exercise class - 23.11.2021)

We shall discuss and solve together this exercise.

Exercise 5. Consider the following variant of BGP games:

1. The network consists of an undirected **weighted** graph.
2. Like in BGP games, each node is a player, and the strategies of a player consists in choosing one of the neighbors.
3. The rank of paths **respect the distances** meaning that for each path we consider its length (sum of the weights) to the destination d , and each node prefers shorter paths over longer paths.² As usual, a “fail” is the worst possible scenario for a player.³

Consider the following example:



1. Compute the cheapest (shortest) path from each node to d
2. Consider the scenario in which
 - (a) Player 2 points (chooses) 1 and
 - (b) Player 3 points (chooses) dWhat is the best response of player 1?
3. Compute a PNE for the game above. How many PNE are there?
4. Is the game NBR-solvable with clear outcome? (If yes, prove it. If no, give a formal argument.)

²Let's assume there are no parallel links, and weights on edges are strictly positive, so there are no ties.

³You can imagine that in this case the distance to the destination is “infinite”.