Revisiting AS-ranking

Mickael MEULLE France Telecom Division R&D

June 2007

orange™

Agenda

- Interdomain reachability with BGP
- Why studying AS ranking?
- A new approach to rank transit ASs
- Classification overview

Internet routing with BGP

- The Internet network is:
  - More than 20 000 interconnected Autonomous Systems (ASs)
    - **stub** Networks (clients): ASs only send and receive traffic for themselves
    - Internet Service Providers (ISPs): ASs provide **transit** to other ASs
  - A partition of the public IP address space into ranges (**prefixes**)
    - companies, DSL users, universities, administrations…
    - web farms, "Content Delivery Networks", hosting services …
- Internet routing is handled by Border Gateway Protocol
  - BGP routers propagate reachability of prefixes
  - BGP routers maintain routing tables toward approximately 200 000 prefixes
- Internet connectivity
  - each router know a route toward any public IP address

End-to-end Internet routes

- Internet paths between hosts rely on:
  - > BGP routes selected and propagated by BGP routers
  - > Router paths inside an AS are matter of the AS's internal routing
- BGP route selection and propagation rely on:
  - > Network topology: AS-level and router-level topologies
  - > Routing Policies of ASs: prefix announcements, route filters, route preferences…

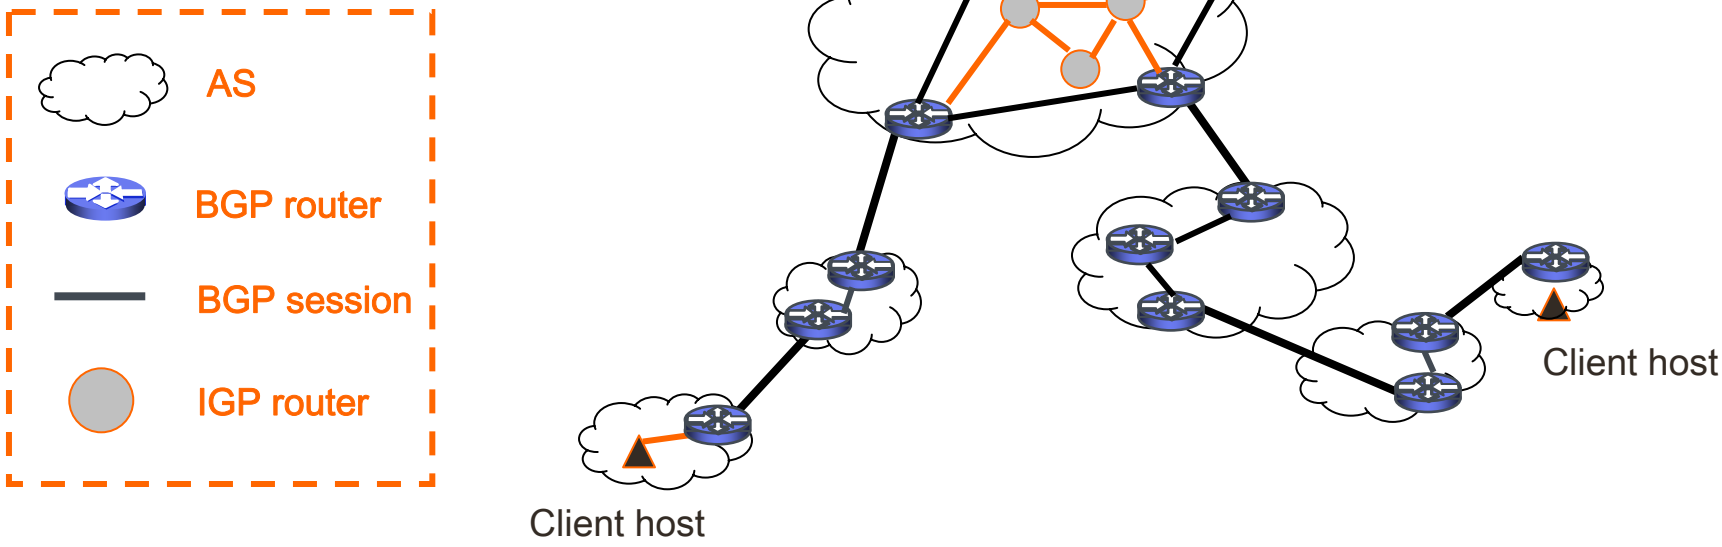

Internet is a huge business place

- business agreements are negotiated for AS links
  - > **Customer/Provider**: C2P/P2C
    - Client pays provider for incoming and outgoing traffic
    - Client routers send to provider routers own & clients BGP routes
    - Provider routers send to client routers all their BGP routes
  - > **Peering** ("Sender Keep All"): PEER

    No money, no guaranty (no SLA)
    - Transit only between clients of clients
    - on peering sessions, routers send own & clients BGP routes
  - > **Others** (closer to reality):
    - Regional / national transit & peering
    - IP prefix-based
    - Sibling SIB (same administration)

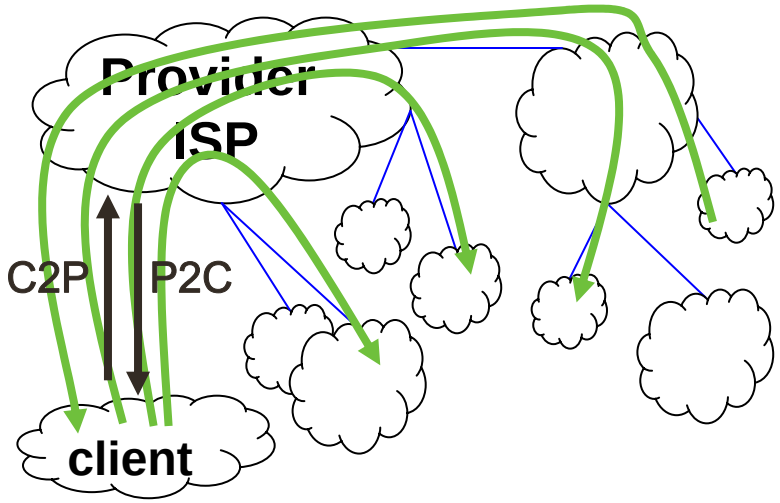

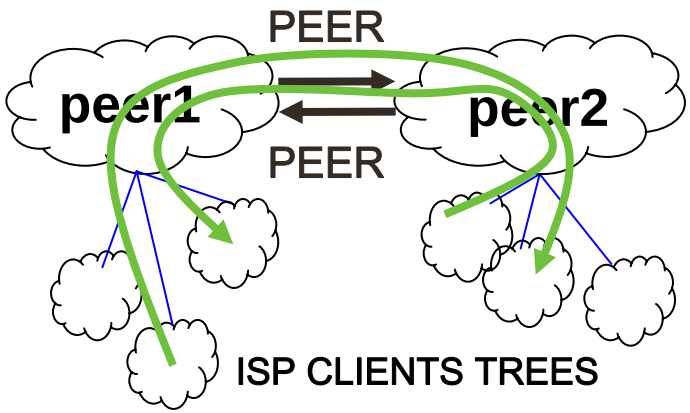

Decision support for AS interconnections

- The Reachability of an AS: BGP paths to prefixes
  - > Depends on many 'unknown' factors
  - > Business agreements, Traffic Engineering…
- A given reachability implies an underlying quality for IP packet forwarding
  - > business agreements and neighbors AS shape possible AS-level paths from a given AS
- Accurate negotiation of new inter-AS links:
  - > Needs knowledge about reachable prefixes and routes provided
    - Which prefixes? Are they cheaper for example?
    - Which hosts belonging to prefixes? Internet clients, provided contents…
    - What kind of BGP paths toward prefixes?
- prefix and paths are unknown, real impact of a new link also!

AS rankings: motivations

- AS rankings: a way to compare transit ASs
  - > Hierarchical position (economic)
  - > Topological position (connectivity)

- Reachability of an AS is strongly correlated with its position in the Internet Hierarchy
  - > Customer-to-provider & provider-to-customer agreements create hierarchy

~22 000 AS

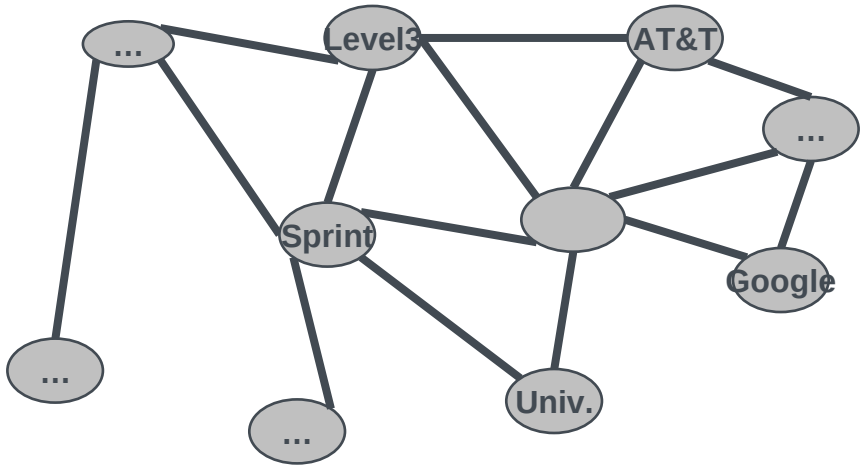

AS logical link = set of BGP sessions between routers

AS rankings overview

- Approaches are mainly based on "BGP tomography"
  - Input: BGP tables at multiple vantage points → (AS_PATH, prefix)
  - Output: estimation about AS position in Internet hierarchy
- In state of the art: no universal ranking adopted
  - > "Black boxes"
    - Netconfigs: http://www.netconfigs.com/general/ranks.htm
    - FIxedOrbit: http://www.fixedorbit.com/metrics.htm
    - Renesys: http://www.renesys.com/products_services/market_intel/rankings/
  - > Heuristics for Internet hierarchy:
    - Gao et al., "On the Hierarchical Structure of the Logical Internet Graph"
    - Subramanian et al., "Characterizing the Internet Hierarchy from Multiple Vantage Points"
  - > Caida AS rankings: http://www.caida.org/analysis/topology/rank_as/
    - Gaol: compute size of customer cones for each AS
    - Method: use AS-level graph annotated with inferred business agreements
  - > Linear algebra approach
    - Clerot et al., "*A Social Network Approach for the Ranking of the Autonomous Systems of the Internet*"
    - Wang et al., "*Inter-Domain Access Volume Model: Ranking Autonomous Systems*"

Issues in the estimation of AS ranks

- Graph-based metrics are not taking care of routing policies
  - > Topology considerations are erroneous because BGP paths are not shortest paths!
  - > Examples: betweenness centrality, triangles, clustering coefficient…
- Lacks of BGP tomography to accurately estimate *reachability* from all AS to all prefixes
  - > A small number of measurement AS
    - bias in data: prefixes seen and where, AS links seen….
  - > If all data is merged from all collectors
    - Some indicators can be over-estimated
- Uncertainty of inferred business relationship to identify customers
  - > degenerated problem: multiple solutions exist
  - > Model of agreements is too coarse

Mining BGP tomography for Transit

- We still use BGP feeds from Route-Views, RIPE and looking glass routers: more than 100 ASs with full BGP tables

- How to mind BGP tomography for AS ranking computation?
  - > Goal is to estimate amount of customer prefixes provided by each AS
    - without relying on a model of business agreements…
    - without bias introduced by measurement points…
    - With freedom on the way to give importance to prefixes
  - > Keep in mind bias in estimation of amount of prefix observed
    - For ASs (limited view of measurement points)
    - On AS links (measurement point location)
- We search for the amount of IP space “behind” an AS
  - > End-to-end reachability is maintained:
    - Transit provider are responsible for propagation of reachability announces
    - Transit providers are observed on paths to customer prefixes

IP space transited by an AS

- IP Space transited by an AS as seen in a path at an observation point
  - For a path (X-Y-Z-T) to prefix p, we record:
    - transit (X,Y,Z) of prefix p for Y
    - transit (Y,Z,T) of prefix p for Z
- Given a set of AS paths, we can compute:
  - Set of prefixes transited by each AS X
    - Prefixes on sub AS paths of length 3 (*,X,*)
- Rank of an AS is the percentage of IP space transited
  - This rank can be a weighted sum of each prefix importance
- We compute a rank for each AS, from each measurement AS (set of paths is split)
  - We use only measurement AS with full routing tables
- AS ranking: average rank = average IP space transited

Average IP space transited

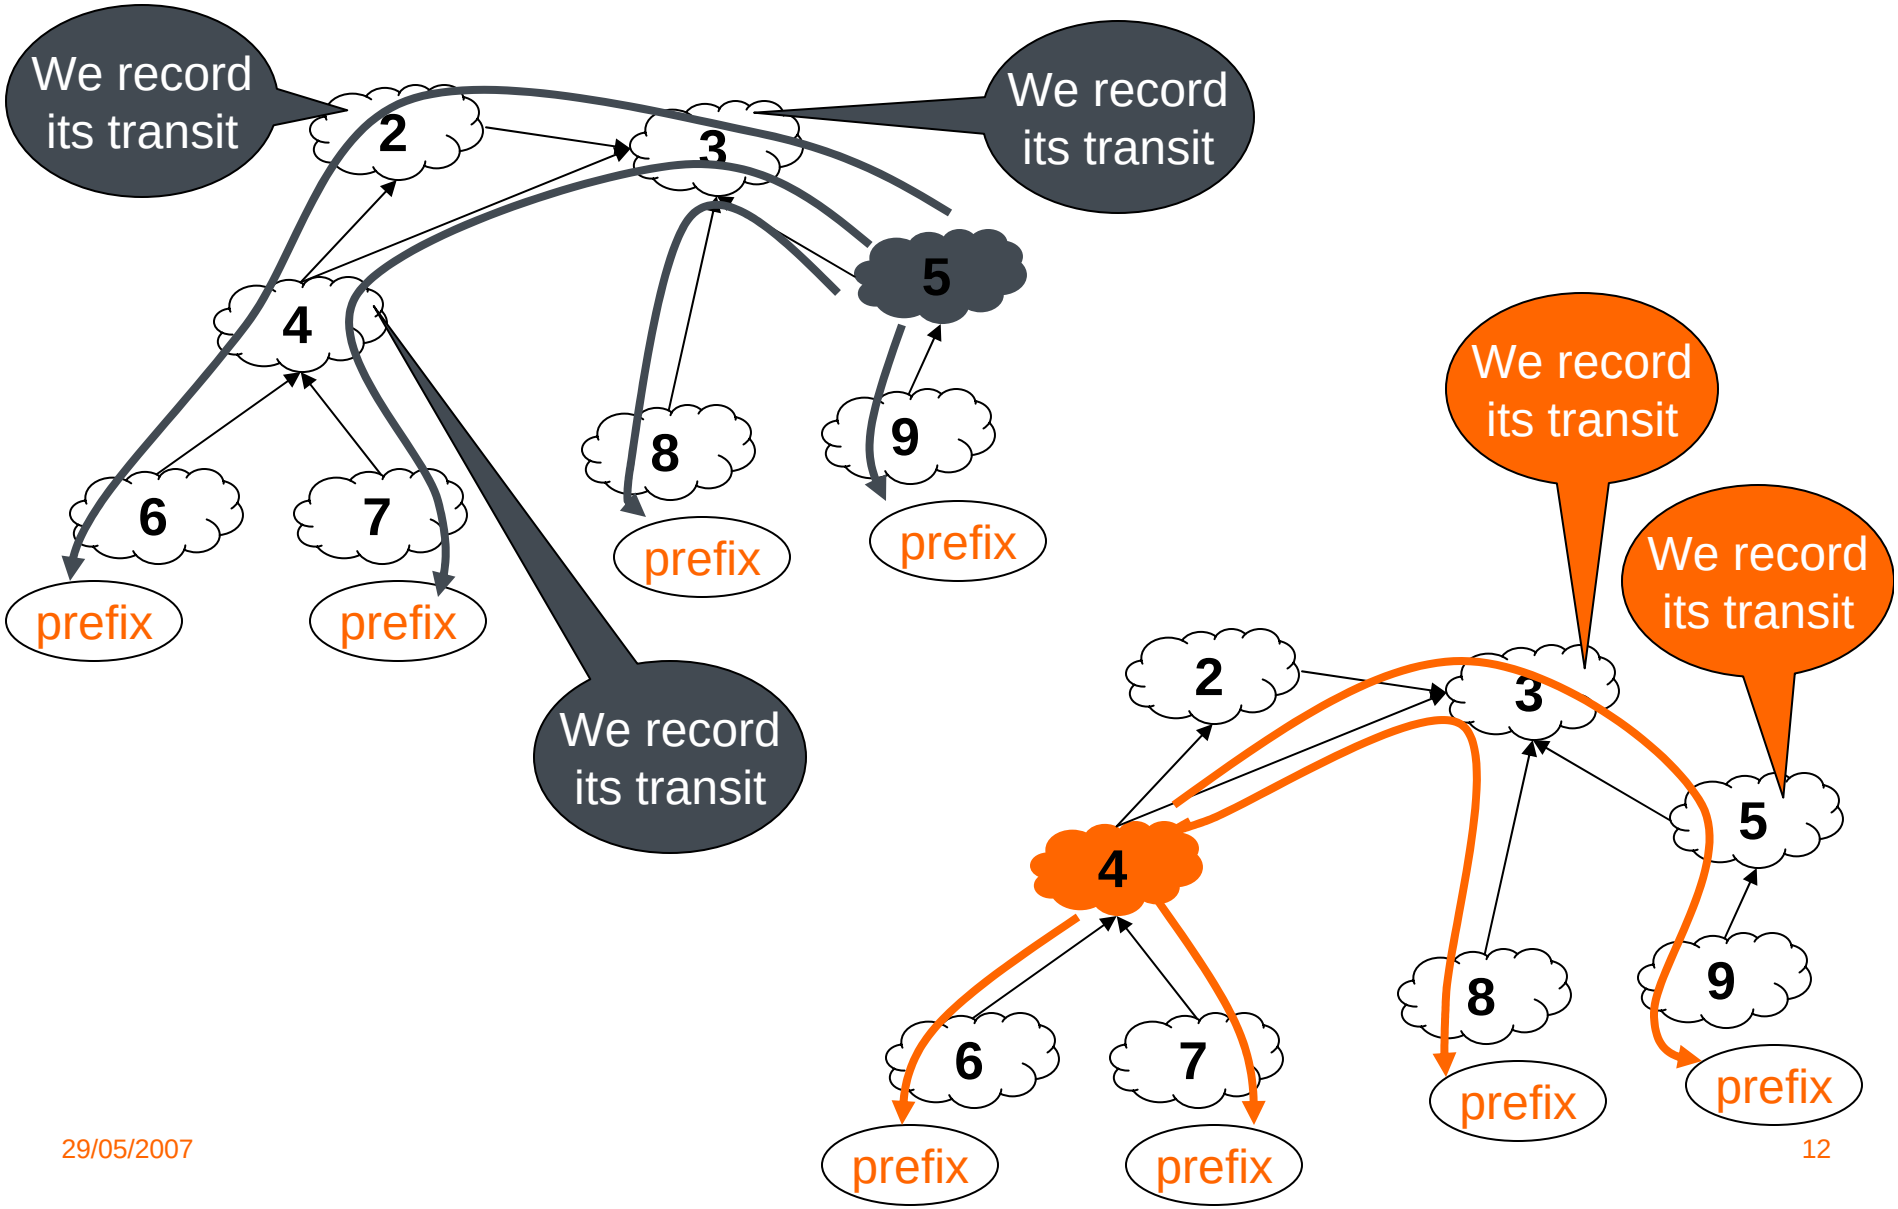

Advantages of this ranking

- Intuitively
  - > Any AS « far » from a measurement AS is seen to transit its customer prefixes
  - > Any AS « close » to a measurement AS has an over-estimated number of prefixes transited
- narrowed bias due to placement of measurement ASs
  - > If some ASs are providers of one or several measurement AS, it will not change so far the results
- Prefix granularity limits errors due to pure graph-based estimation of connectivity
  - > example: transit of some prefixes between two Tiers1 AS is not mis-understood!
- Compute the average IP space transited by an AS as seen from many measurement
  - > The sample of hundreds measurement AS (Route-Views and RIPE neighbors) becomes representative for the quantity measured

Results overview

- Sample taken in March 15th, 2007: 109 measurement ASs

1.	[3356]	LEVEL3 (14.388)
2.	[1239]	SPRINTLINK (11.460)
3.	[1299]	TELIANET (10.480)
4.	[3549]	GBLX (9.480)
5.	[174]	PSINET (8.023)
6.	[2914]	VERIO (7.767)
7.	[209]	ASN-QWEST (7.522)
8.	[7018]	ATT-INTERNET4 (5.610)
9.	[701]	ALTERNET-AS (5.489)
10.	[3561]	CWUSA (3.064)
11.	[3257]	TISCALI-BACKBONE (2.954)
12.	[6461]	ABOVENET (2.876)
13.	[5511]	OPENTRANSIT (2.839)

Results overview

- 14. [2516] KDDI (2.767)
- 15. [6453] TELEGLOBE-AS (2.701)
- 16. [4637] REACH (2.457)
- 17. [1273] ECRC (2.405)
- 18. [3320] UNSPECIFIED (2.364)
- 19. [3491] CAIS-ASN (2.061)
- 20. [8928] INTEROUTE (1.701)
- 21. [2497] IIJ (1.666)
- 22. [721] DLA-ASNBLOCK-AS (1.557)
- 23. [5400] CIPCORE (1.518)
- 24. [20965] GEANT (1.471)
- 25. [702] AS702 (1.319)
- 26. [286] UNSPECIFIED (1.280)

Conclusion

- We introduce a new AS ranking algorithm
  - Use BGP tomography with public BGP feeds
    - Route-Views & RIPE projects + looking glass
  - It records « foreign » transit of prefixes by ASs
  - The value computed can be directly interpreted

- Advantages compared to state-of-art methods
  - It is not a « black box »
  - Take care of measurement bias
    - Over-estimated reachability of some ASs is averaged
  - Take advantages of multiple routing tables used
    - The more tables as input, the more precise is result
  - One can give more importance to some prefixes or restrict the set of prefixes for which ranking is computed

Thanks!