

Distributed Route Aggregation on the Global Network (DRAGON)

João Luís Sobrinho¹

Laurent Vanbever², Franck Le³, Jennifer Rexford²

¹*Instituto Telecomunicações, ¹Universidade de Lisboa*

²*Princeton University, ³IBM T. J. Watson Research*

Recently in the news (August 2014)

THE WALL STREET JOURNAL. ≡ | TECH

TECHNOLOGY

Echoes of Y2K: Engineers Buzz That Internet Is Outgrowing Its Gear

Routers That Send Data Online Could Become Overloaded as Number of Internet Routes Hits '512K'

By DREW FITZGERALD [CONNECT](#)

Updated Aug. 13, 2014 7:38 p.m. ET



14 August 2014 Last updated at 12:05 GMT

Browsing speeds may slow as net hardware bug bites

By Mark Ward

Technology correspondent, BBC News

+512 K IPv4 prefixes propagated to \approx 50.000 ASs

Not a scalable Internet routing system

- Most prefixes propagated (by BGP) to all ASs
 - Routing & forwarding tables growth
 - Churn & convergence time increase
 - S*BGP processing requirements escalation

Outline

- Characterizing the Internet for scalability
- DRAGON: basic ideas
- DRAGON: filtering strategy
- DRAGON: additional aspects
- DRAGON: performance
- Conclusions

Outline

- Characterizing the Internet for scalability
- DRAGON: basic ideas
- DRAGON: filtering strategy
- DRAGON: additional aspects
- DRAGON: performance
- Conclusions

Decentralization: each AS decides...

- Where to acquire address space
 - provider? provider-independent, Internet registry?
- Where to connect
 - multi-homing? peering at an exchange point?
- How to announce assigned address space
 - de-aggregate first?
- How to treat routes learned from neighbors
 - which routing policies?

Structure: opportunities to scale?

Classless Inter Domain
Routing



Hierarchy: **IP prefixes**

Structure: opportunities to scale?

Provider-customer
agreements



Hierarchy: **ASs**


Classless Inter Domain
Routing



Hierarchy: **IP prefixes**

Structure: opportunities to scale?

Provider-customer
agreements



Hierarchy: **ASs**



Classless Inter Domain
Routing




Hierarchy: **IP prefixes**

Structure: opportunities to scale?

Provider-customer
agreements



Hierarchy: **ASs**



Geography
(rough)

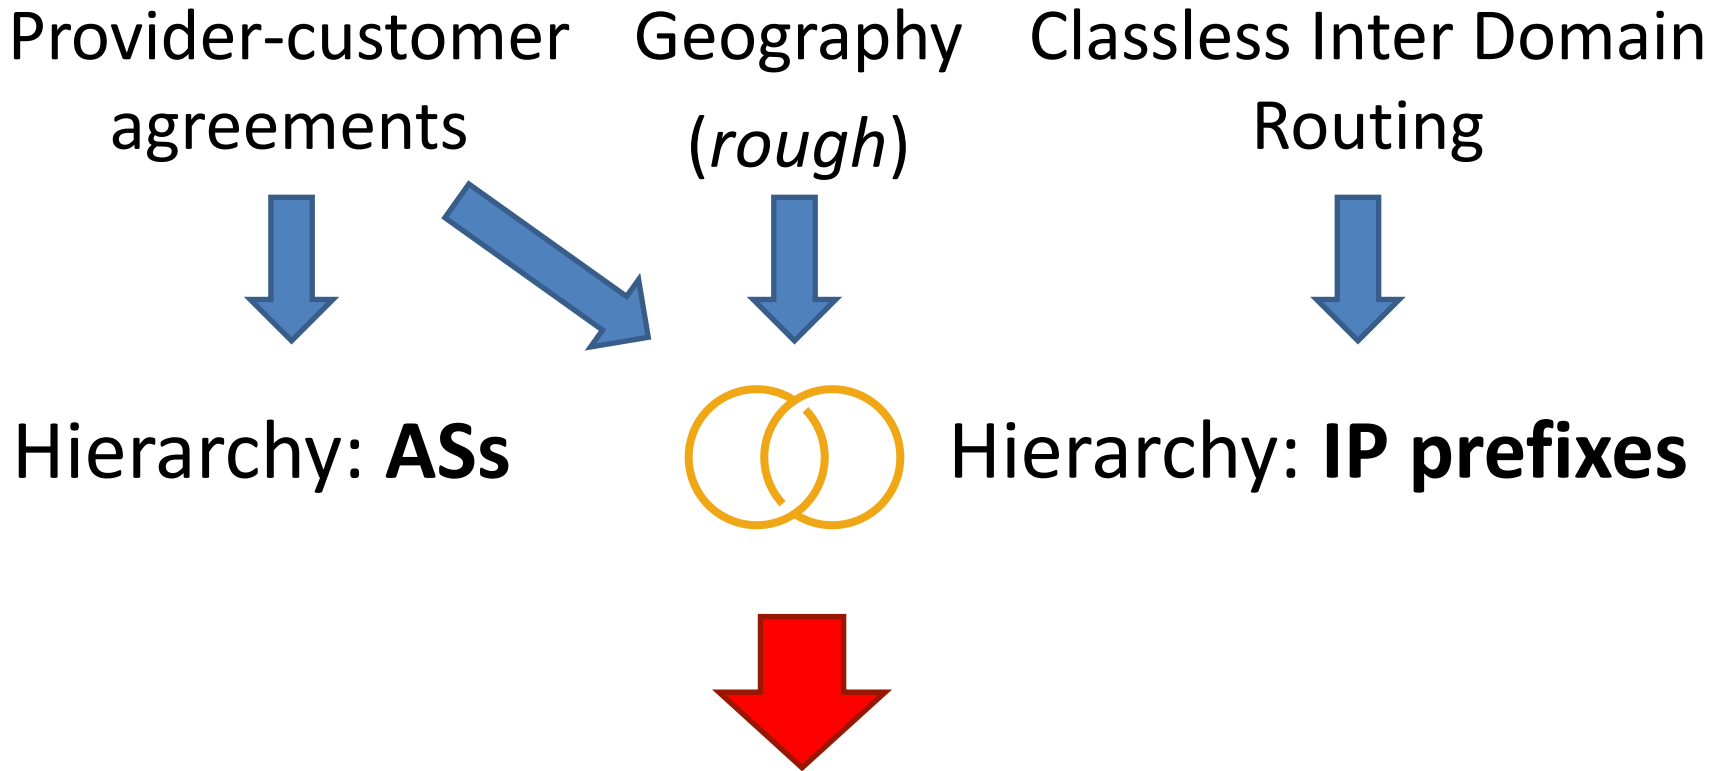


Classless Inter Domain
Routing



Hierarchy: **IP prefixes**

Structure: opportunities to scale?



Outline

- Characterizing the Internet for scalability
- **DRAGON: basic ideas**
- DRAGON: filtering strategy
- DRAGON: additional aspects
- DRAGON: performance
- Conclusions

Filtering strategy

- Filter the more specific prefixes when possible
 - no black holes
 - strive to preserve global forwarding behavior
- Use incentives to filter locally
 - save on routing and forwarding state
 - forward data-packets along best possible route
- *Make standard usage of BGP routing messages*


Generation of aggregation prefixes

- **Generate aggregation prefixes when beneficial**
 - permit filtering of provider-independent prefixes
 - new address space is not created
- **Announce as in BGP**
 - self-organization when more than one AS generates the same aggregation prefix


Outline

- Characterizing the Internet for scalability
- DRAGON: basic ideas
- **DRAGON: filtering strategy**
- DRAGON: additional aspects
- DRAGON: performance
- Conclusions

Providers, customers, and peers




Prefixes




q more specific than p

BGP: Gao-Rexford routing policies



route attributes: “learned from ...”




exportation:

all routes *from* customers


all routes *to* customers

← *q*-route (route pertaining to *q*)

BGP: Gao-Rexford routing policies



route attributes: “learned from ...”




exportation:


all routes *from* customers

all routes *to* customers

BGP: Gao-Rexford routing policies



route attributes: “learned from ...”




exportation:


all routes *from* customers

all routes *to* customers

BGP: Gao-Rexford routing policies



route attributes: "learned from ..."




exportation:


all routes *from* customers

all routes *to* customers


Final state for prefix q




route attributes: “learned from ...”




Final state for prefix p





route attributes: “learned from ...”



Combined states for q and p



route attributes: "learned from ..."



forwarding: **longest prefix match rule**


Filtering Code (FC)

Filtering Code (FC)

Other than the owner of p , in the presence of p , filter q if only if:

attribute of p -route
same or preferred to
attribute of q -route

Filtering Code (FC)



Filtering Code (FC)

Other than the owner of p , in the presence of p , filter q if only if:


attribute of p -route
same or preferred to
attribute of q -route

AS 1, AS 2, AS 3, AS 4, AS 5, AS 8

filter q on executing the FC (—)


Arbitrary AS applies the FC

AS 4 applies the FC




← withdrawal of q -route

Arbitrary AS applies the FC




AS 4 applies the FC



attribute of q -route worsens at AS 3: double incentive to apply the FC (==)

- saves on forwarding state
- restores attribute of route used to forward data-packets with destination in q

All ASs apply the FC




AS 6, AS 7, AS 9

detailed information q

AS 1, AS 2, AS 3, AS 4, AS 5, AS 8

coarse-grained information p


Global property: correctness



Correctness

no routing anomalies
(no black holes)

Global property: route consistency




Route consistency

attribute of route used to forward data-packets is preserved

Optimal route consistency


set of ASs that forgo q is maximal for route consistency

Route consistency: partial deployment




1. AS 5, AS 8 filter $q \rightarrow$ route consistency

Route consistency: partial deployment




1. AS 5, AS 8 filter q → route consistency
2. AS 1, AS 2 filter q → route consistency

Route consistency: partial deployment



1. AS 5, AS 8 filter q → route consistency
2. AS 1, AS 2 filter q → route consistency
3. AS 3 filters q → route consistency

Route consistency: partial deployment



1. AS 5, AS 8 filter q → route consistency
2. AS 1, AS 2 filter q → route consistency
3. AS 3 filters q → route consistency
4. AS 4 filters q → route consistency

Filtering strategy: general case

- Correctness
 - for all routing policies for which BGP is correct
- Route consistent states culminating in optimality
 - for *isotone* routing policies (includes Gao-Rexford)
 - otherwise, some stretch

Optimal route consistency is not synonymous with *efficiency* (think shortest paths)

Outline

- Characterizing the Internet for scalability
- DRAGON: basic ideas
- DRAGON: filtering strategy
- **DRAGON: additional aspects**
- DRAGON: performance
- Conclusions

Additional aspects of DRAGON

- Prefixes at multiple levels of specificity
 - parent prefix and child prefixes

Additional aspects of DRAGON

- Prefixes at multiple levels of specificity
 - parent prefix and child prefixes
- Generation of aggregation prefixes
 - permit filtering of provider-independent prefixes

Additional aspects of DRAGON

- Prefixes at multiple levels of specificity
 - parent prefix and child prefixes
- Generation of aggregation prefixes
 - permit filtering of provider-independent prefixes
- Network dynamics
 - adapts to link failures and additions

Outline


- Characterizing the Internet for scalability
- DRAGON: basic ideas
- DRAGON: filtering strategy
- DRAGON: additional aspects
- **DRAGON: performance**
- Conclusions

Filtering efficiency

$$\frac{\# (\text{FIB entries BGP}) - \# (\text{FIB entries DRAGON})}{\# (\text{FIB entries BGP})}$$

Current set of prefixes

50% of the prefixes
without parent

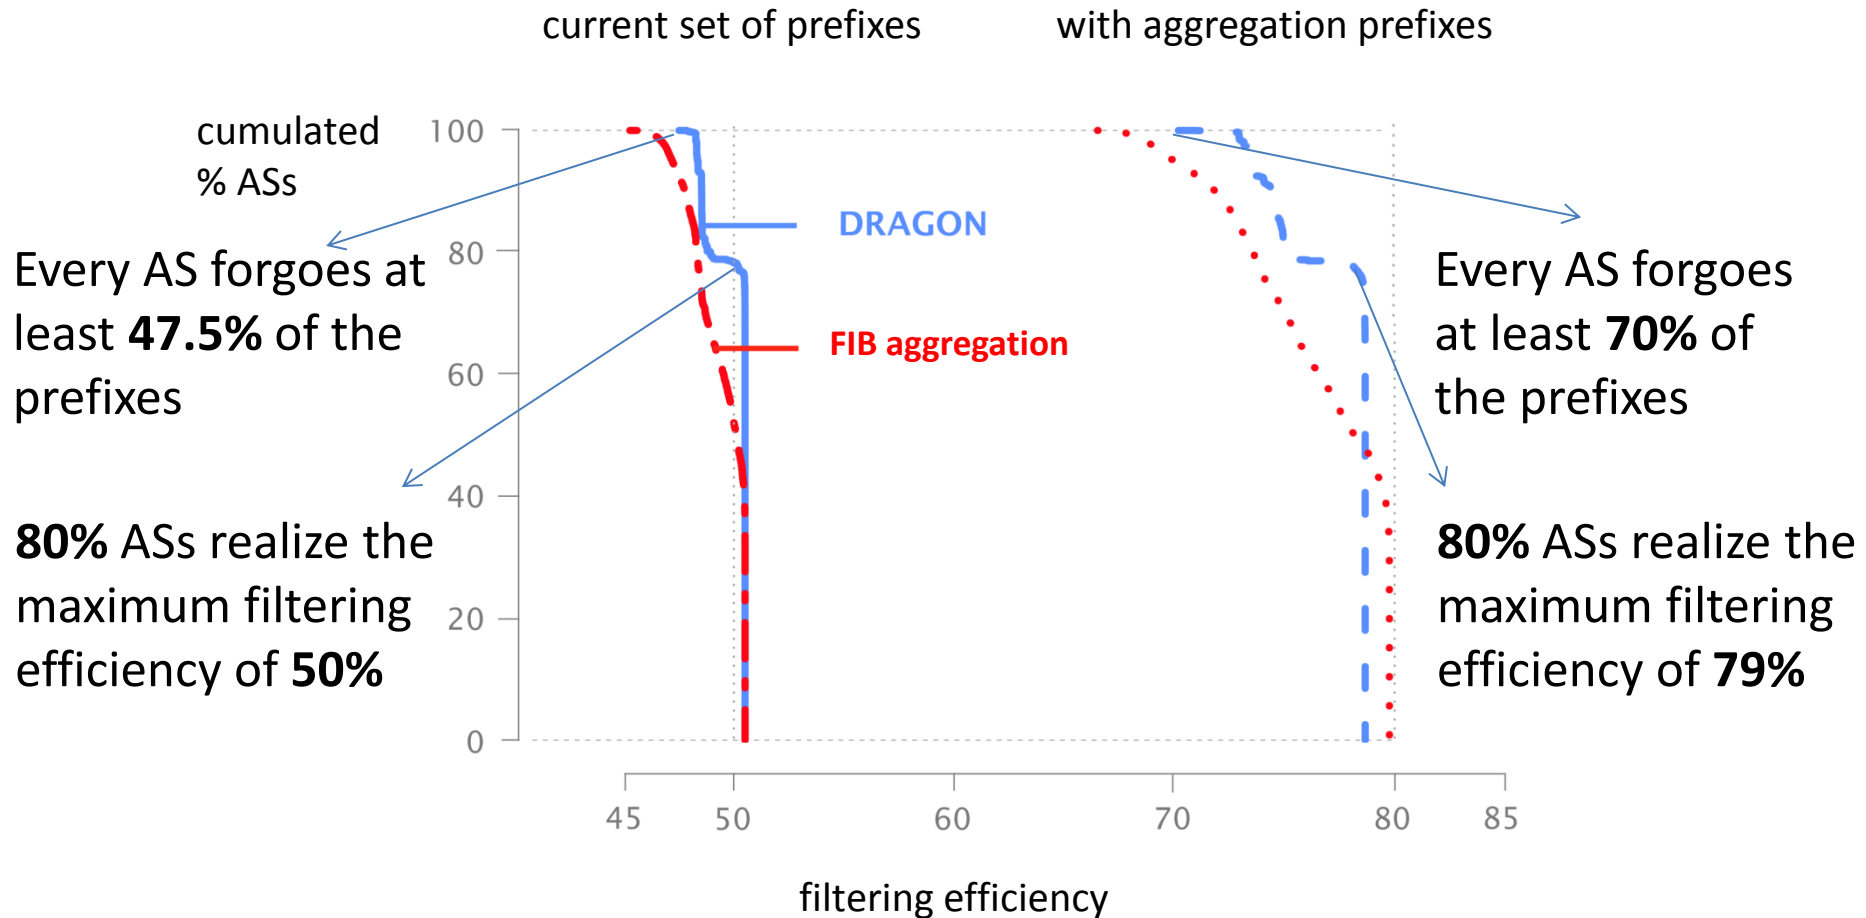


Filtering efficiency
bounded at 50%

With aggregation prefixes

Bound on filtering efficiency rises to 79%

Performance of DRAGON



Conclusions


- DRAGON is a BGP add-on to scale the Internet routing system
- DRAGON can be deployed incrementally
- DRAGON can reduce the amount of state in the Internet routing system by approximately 80%
- DRAGON is – more fundamentally – a solid framework to reason about route aggregation

Visit us at

www.route-aggregation.net

Thank you!

Correctness revisited



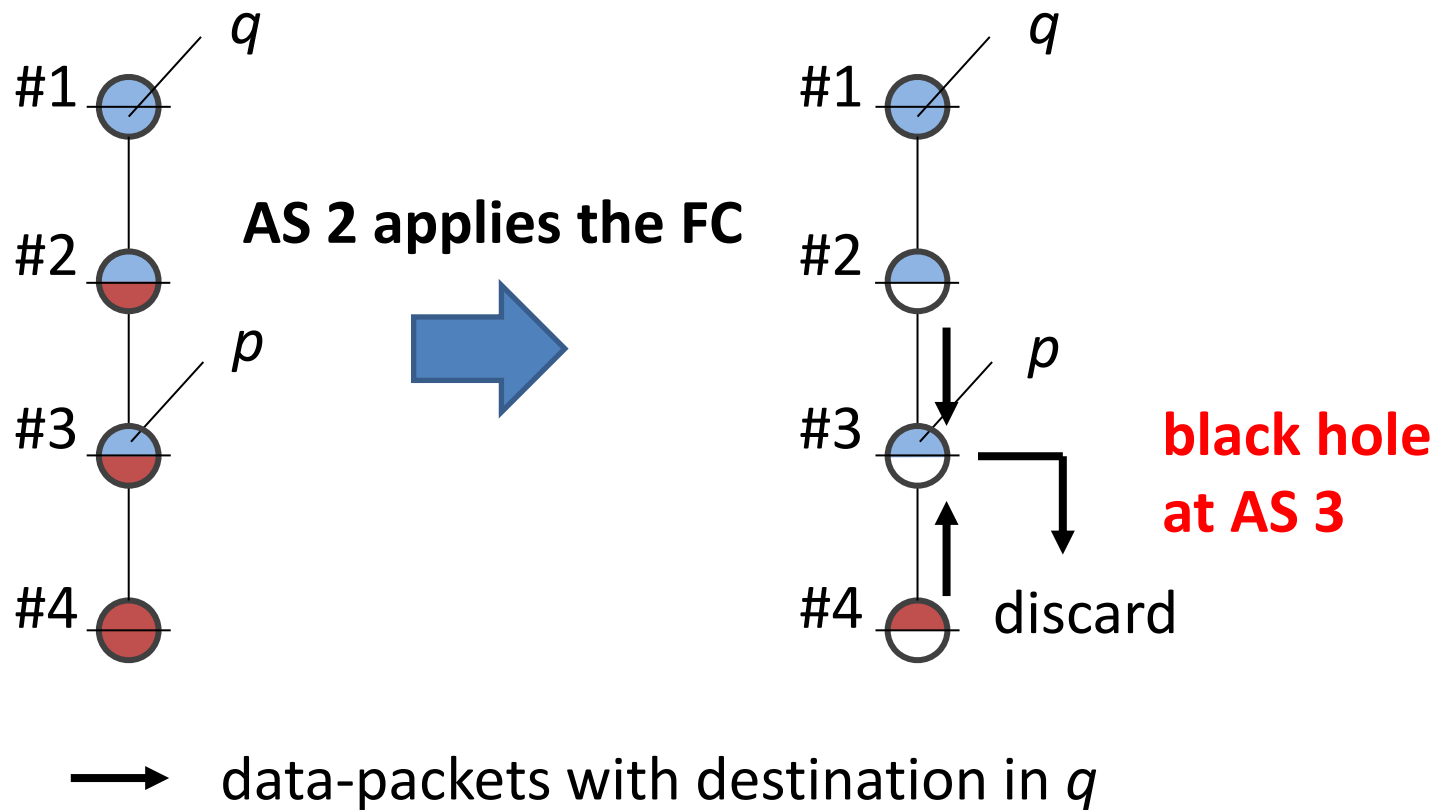
unlikely scenario:
*more specific prefix originated
higher up in hierarchy*

AS 2, AS 4

filter q on executing the FC (—)

q more specific than p

Black hole!



Origination Rule (OR)


Origination Rule (OR)

Owner of p , in the presence of q :

originate p with attribute
same or less preferred than
attribute of q -route

No black holes – guaranteed!

Origination Rule (OR)



Origination Rule (OR)

Owner of p , in the presence of q :

originate p with attribute
same or less preferred than
attribute of q -route

**AS 3 originates p with a provider route
(it announces p only to AS 4)**

No black holes – guaranteed!