

# A new Origin Preference Attribute for BGP traffic engineering

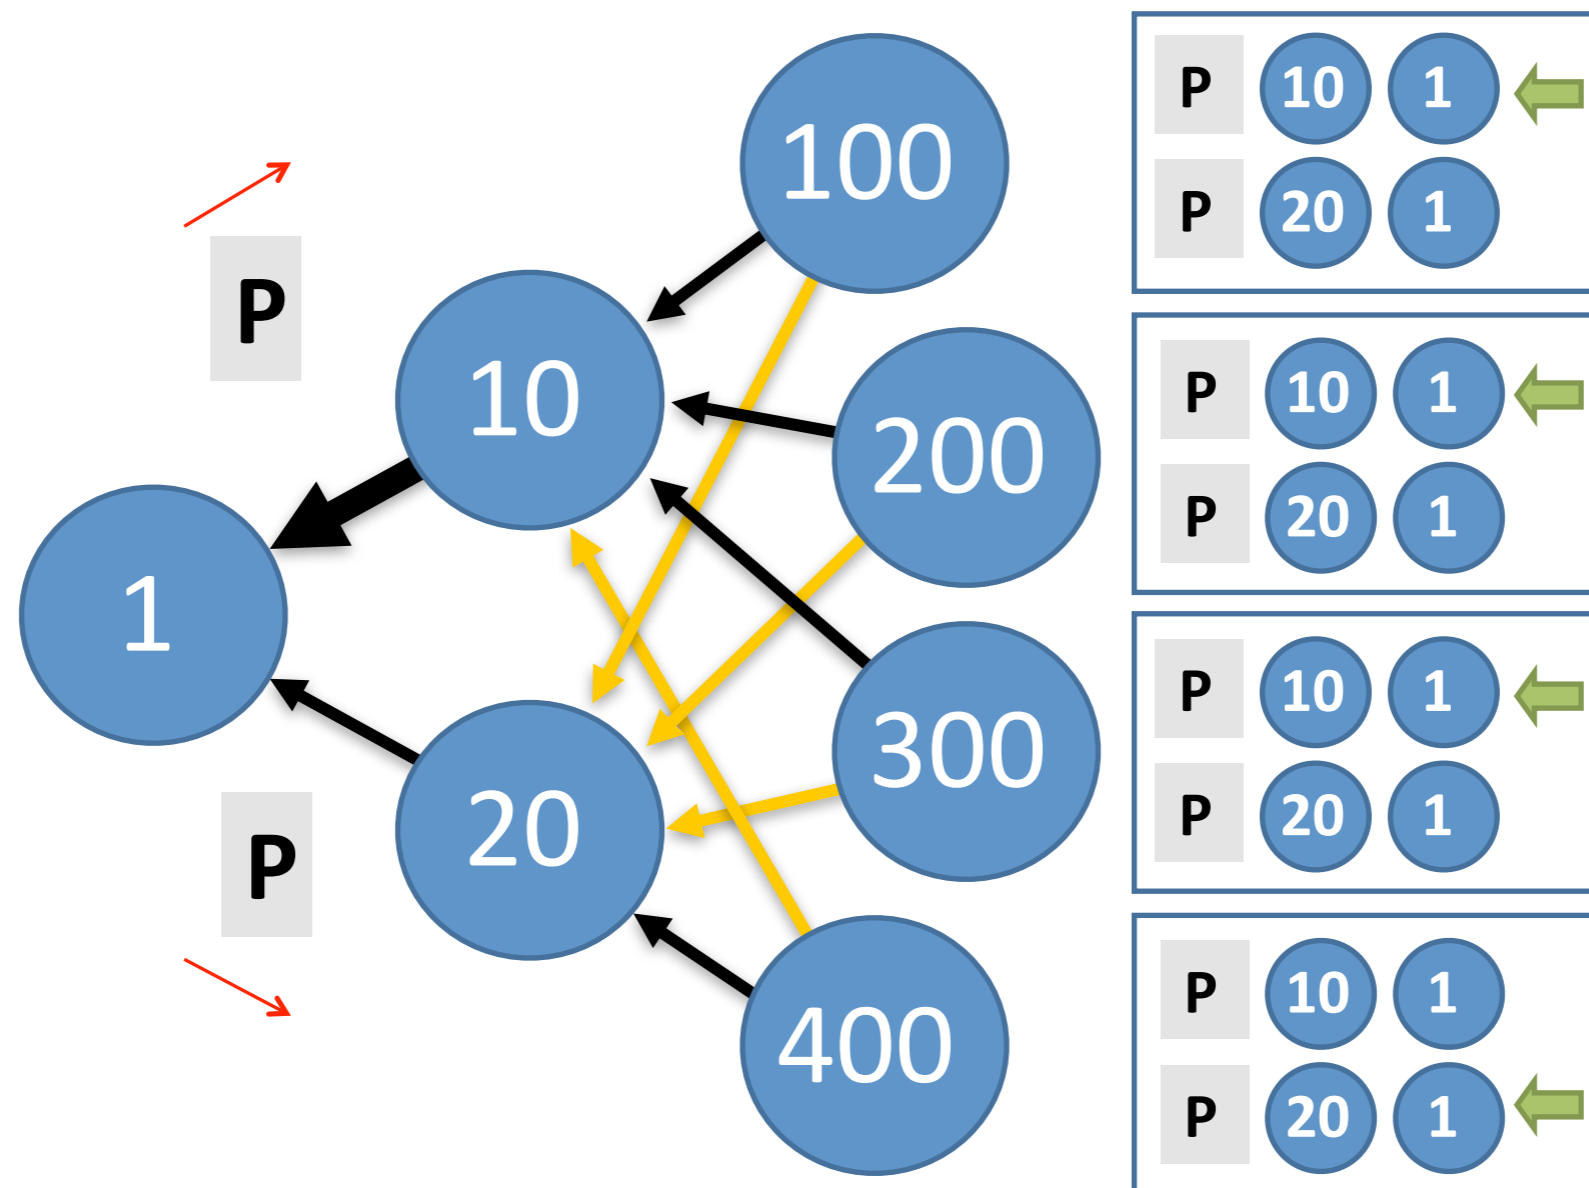
RIPE 65 in Amsterdam  
September 2012

Rolf Winter,  
Iljitsch van Beijnum

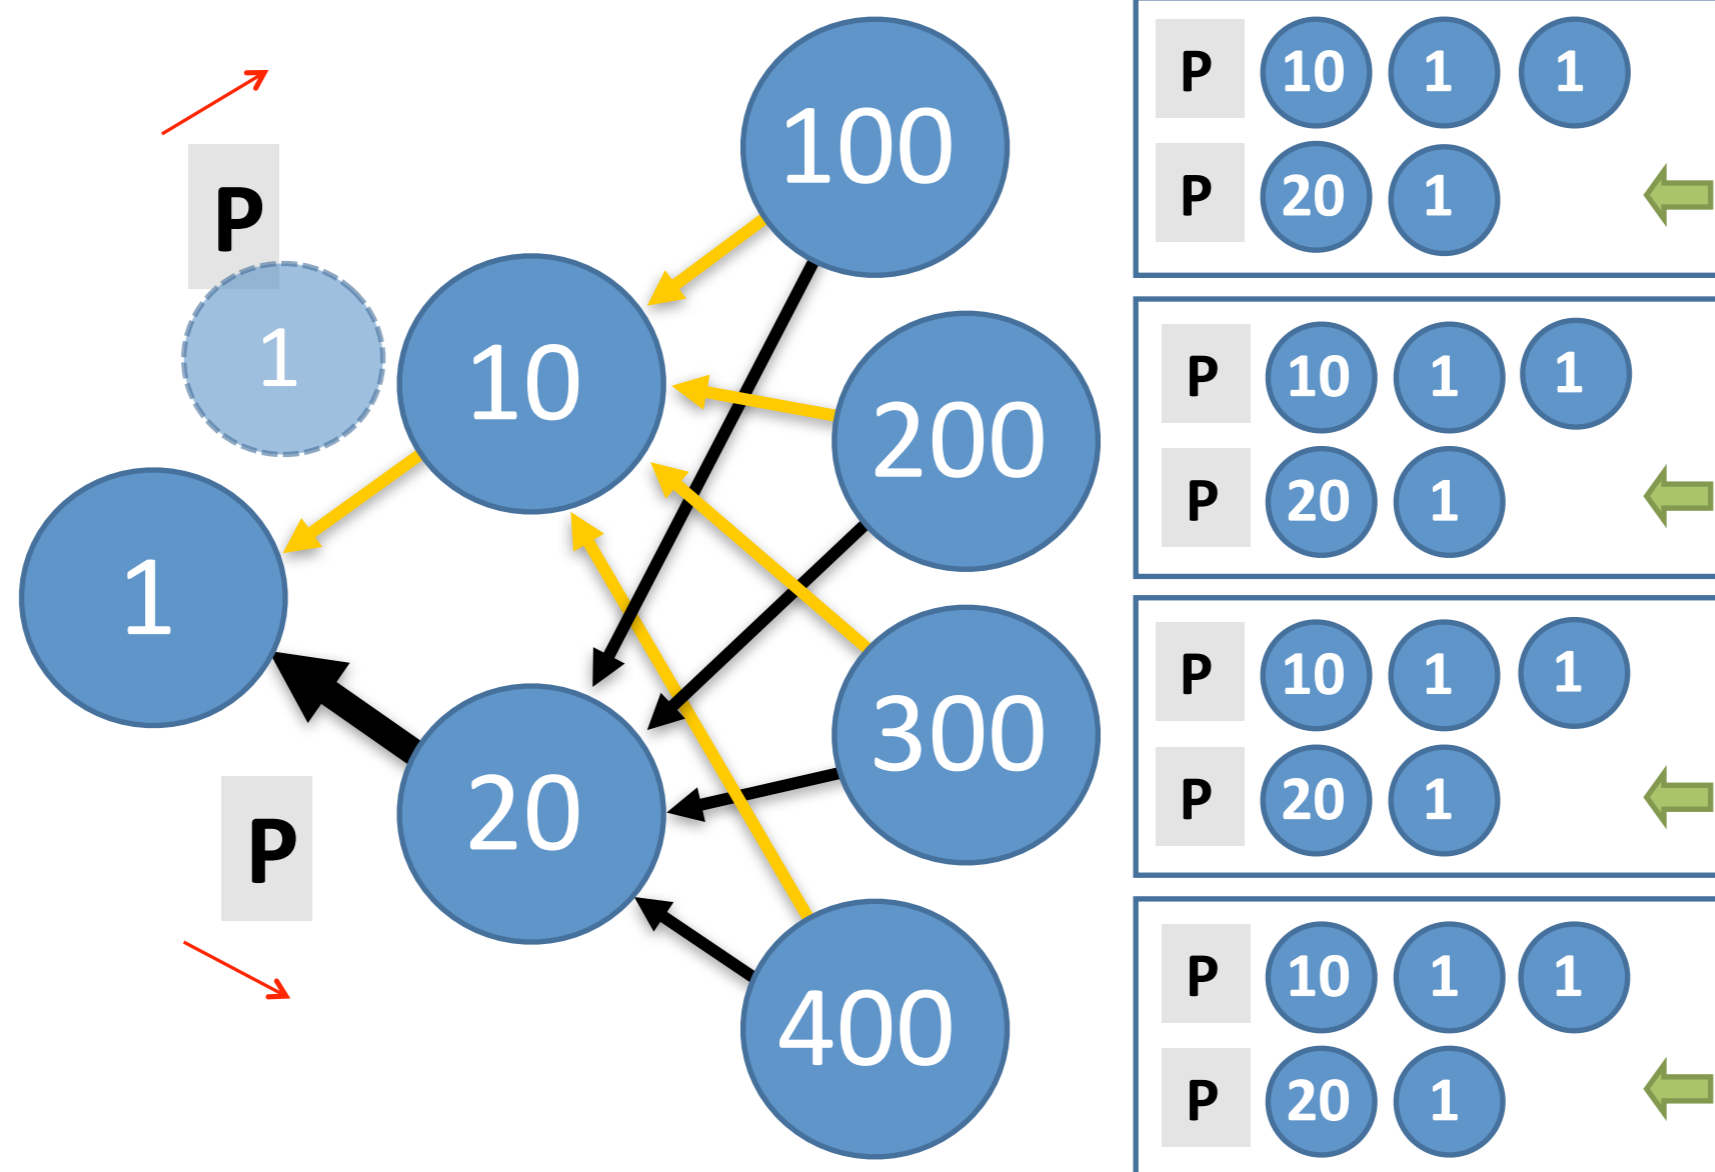
# The problem

- Traffic engineering incoming traffic is hard
  - only AS path [length] is communicated to traffic sources
  - can manipulate AS path length, but not granular enough
- So we need something new

# No inbound TE




# AS path prepending



# Prefix deaggregation

- Path hunting problem with outages
- Also not granular
- Makes the (already big) BGP tables larger




**What if we add a new  
inter-AS metric?**

# OPA

- Origin Preference Attribute
- New optional transitive path attribute containing signed 16-bit OPA value
  - (ignored but propagated if not understood)
- Set by the prefix originator
- No modifications in transit
- Used in BGP path selection after AS path

# No OPA



P	10	1	111
P	20	1	113 ←

-2

P	10	1	195
P	20	1	221 ←

-26

P	10	1	295
P	20	1	313 ←

-18

P	10	1	411 ←
P	20	1	389


22



# Randomization factor


AS100 01100100  
 AS10 00001010  
 AS1 00000001  
 01101111 = 111

**P**




AS100 01100100  
 AS20 00010100  
 AS1 00000001  
 01110001 = 113


**P**




-2



-26

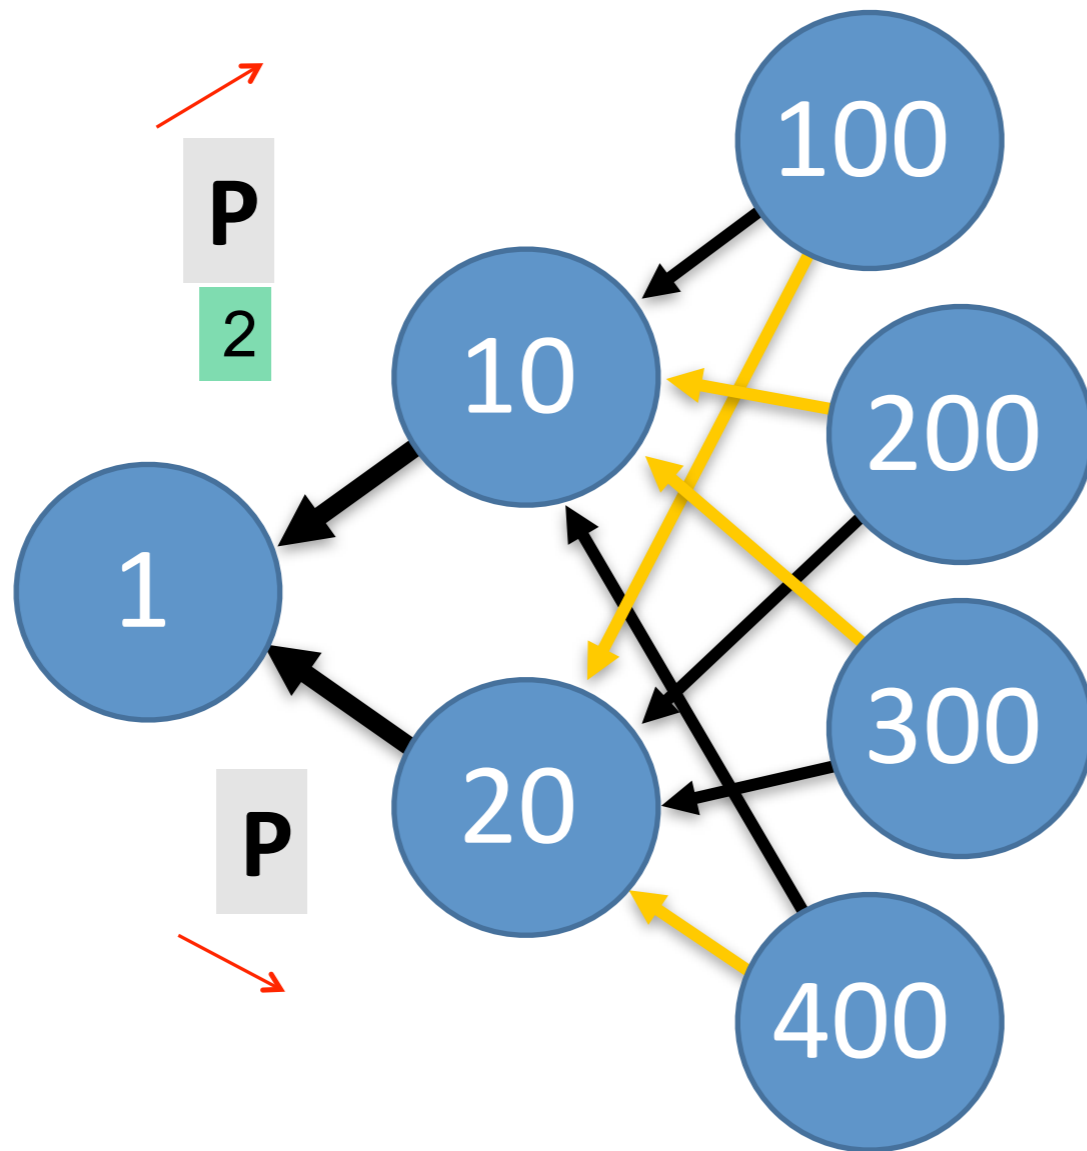


-18



22

# OPA = 2 / OPA = 0



P	10	1	115	←
P	20	1	113	

2

P	10	1	199	
P	20	1	221	←

-22

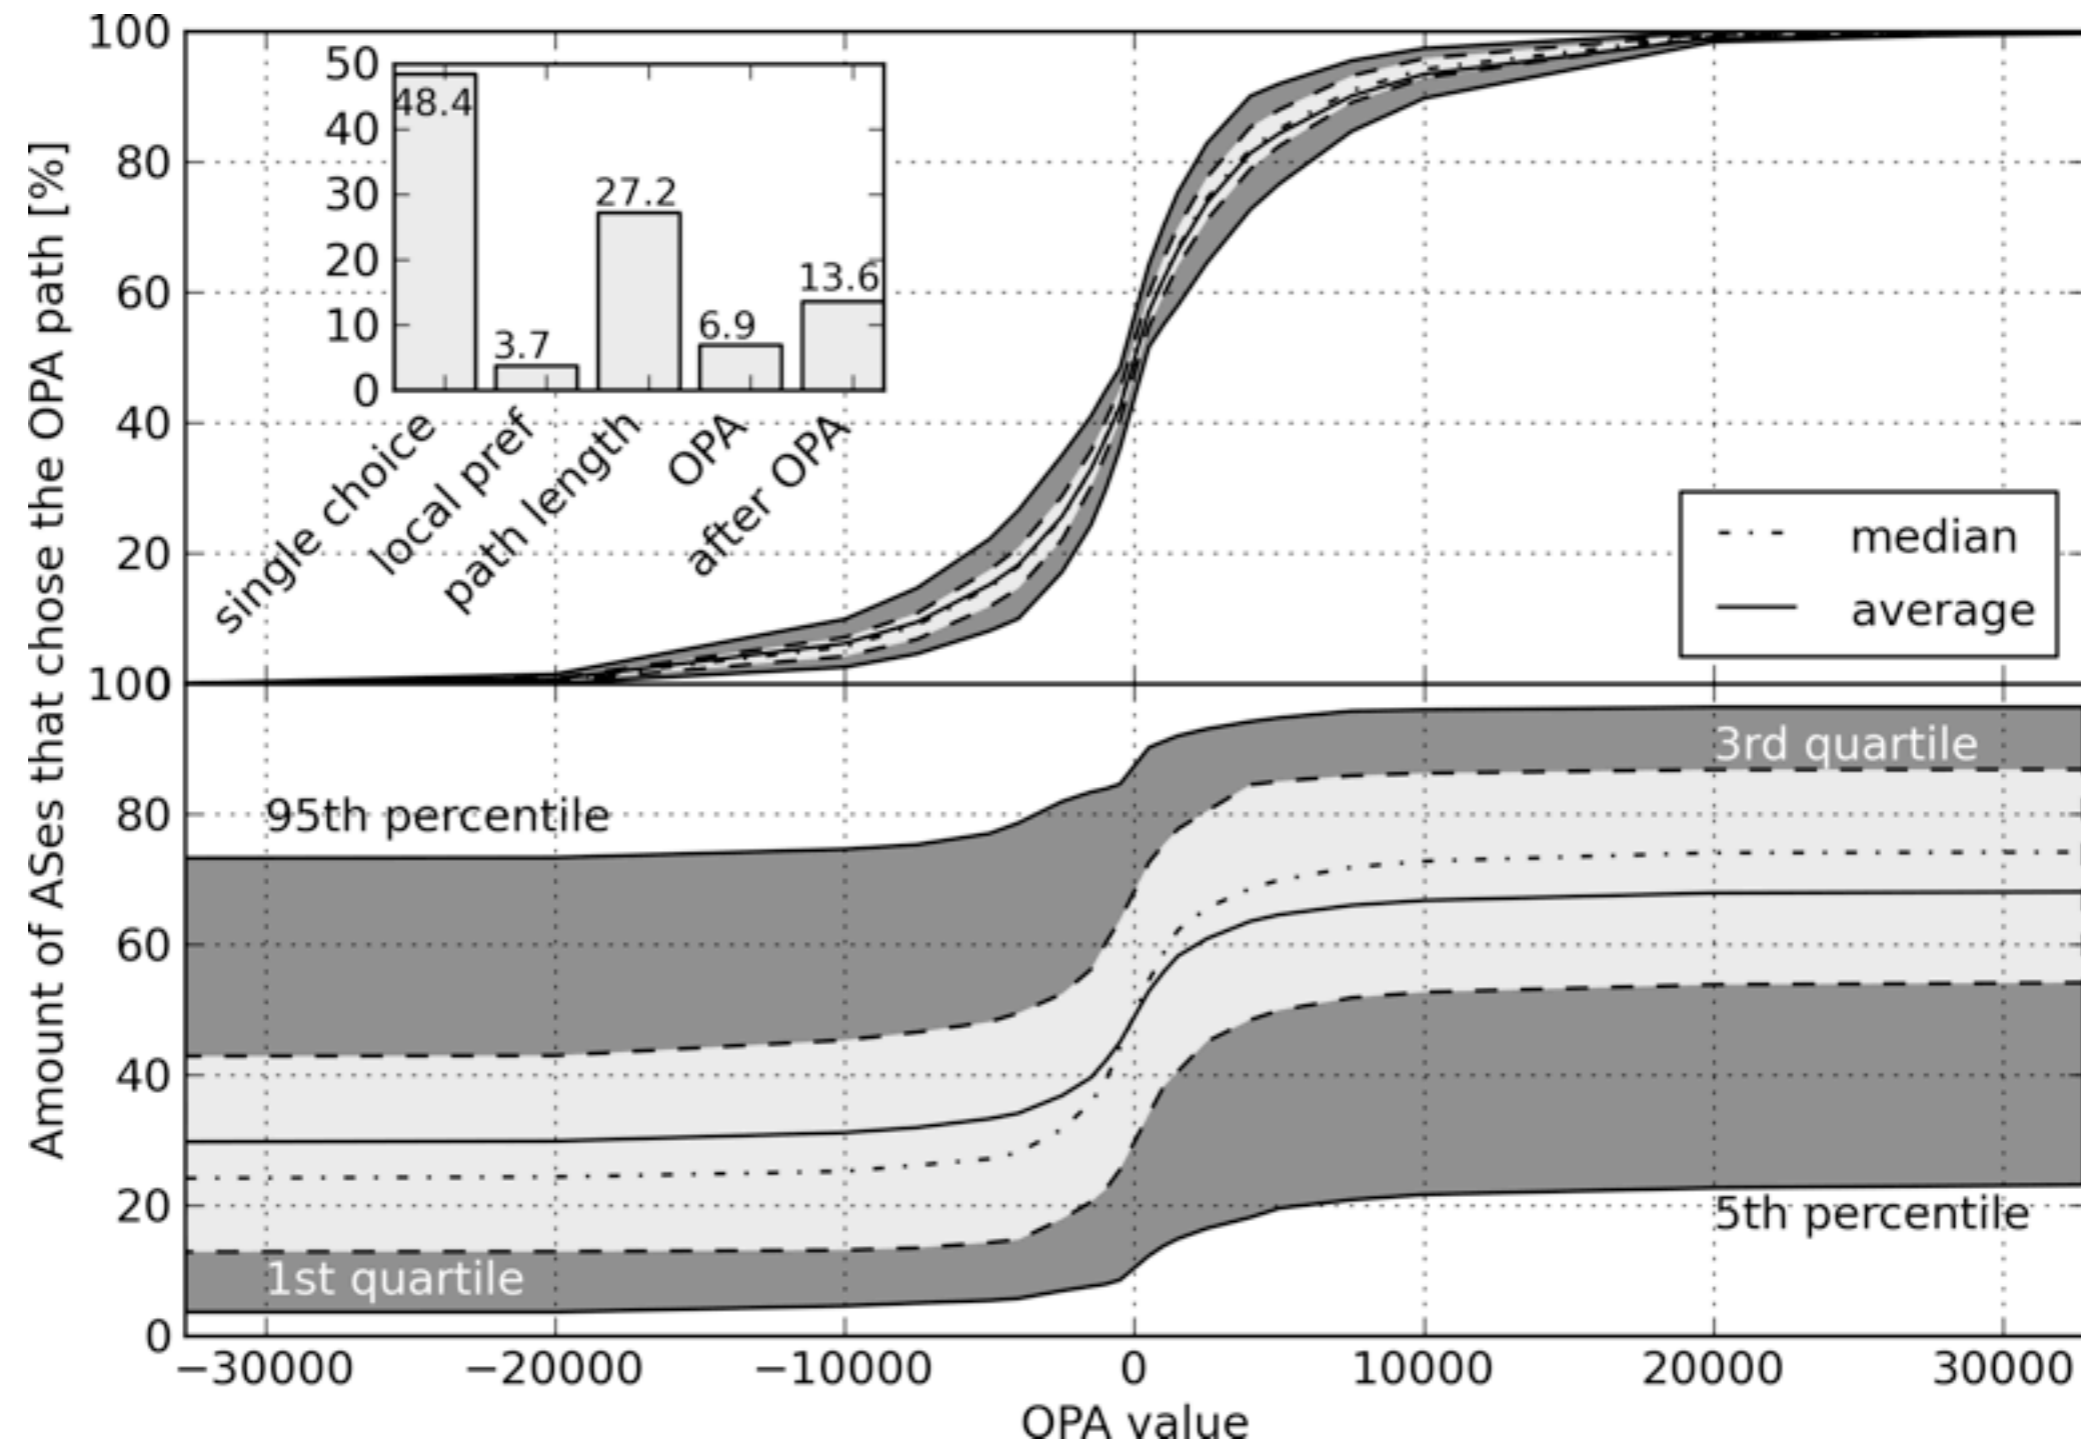
P	10	1	299	
P	20	1	313	←

-14

P	10	1	415	←
P	20	1	389	

26

# Simulations look good!



# Questions, remarks, complaints?

iljitsch@bgpexpert.com

Somewhat old:

[draft-van-beijnum-idr-iac-02](#)

For the academically inclined:

"Explicitly accommodating origin preference for inter-domain traffic engineering"  
SAC'12, Proceedings of the 27th Annual ACM Symposium on Applied Computing