

A BGP Inter-AS Cost Attribute

draft-van-beijnum-idr-iac-02

IETF 74

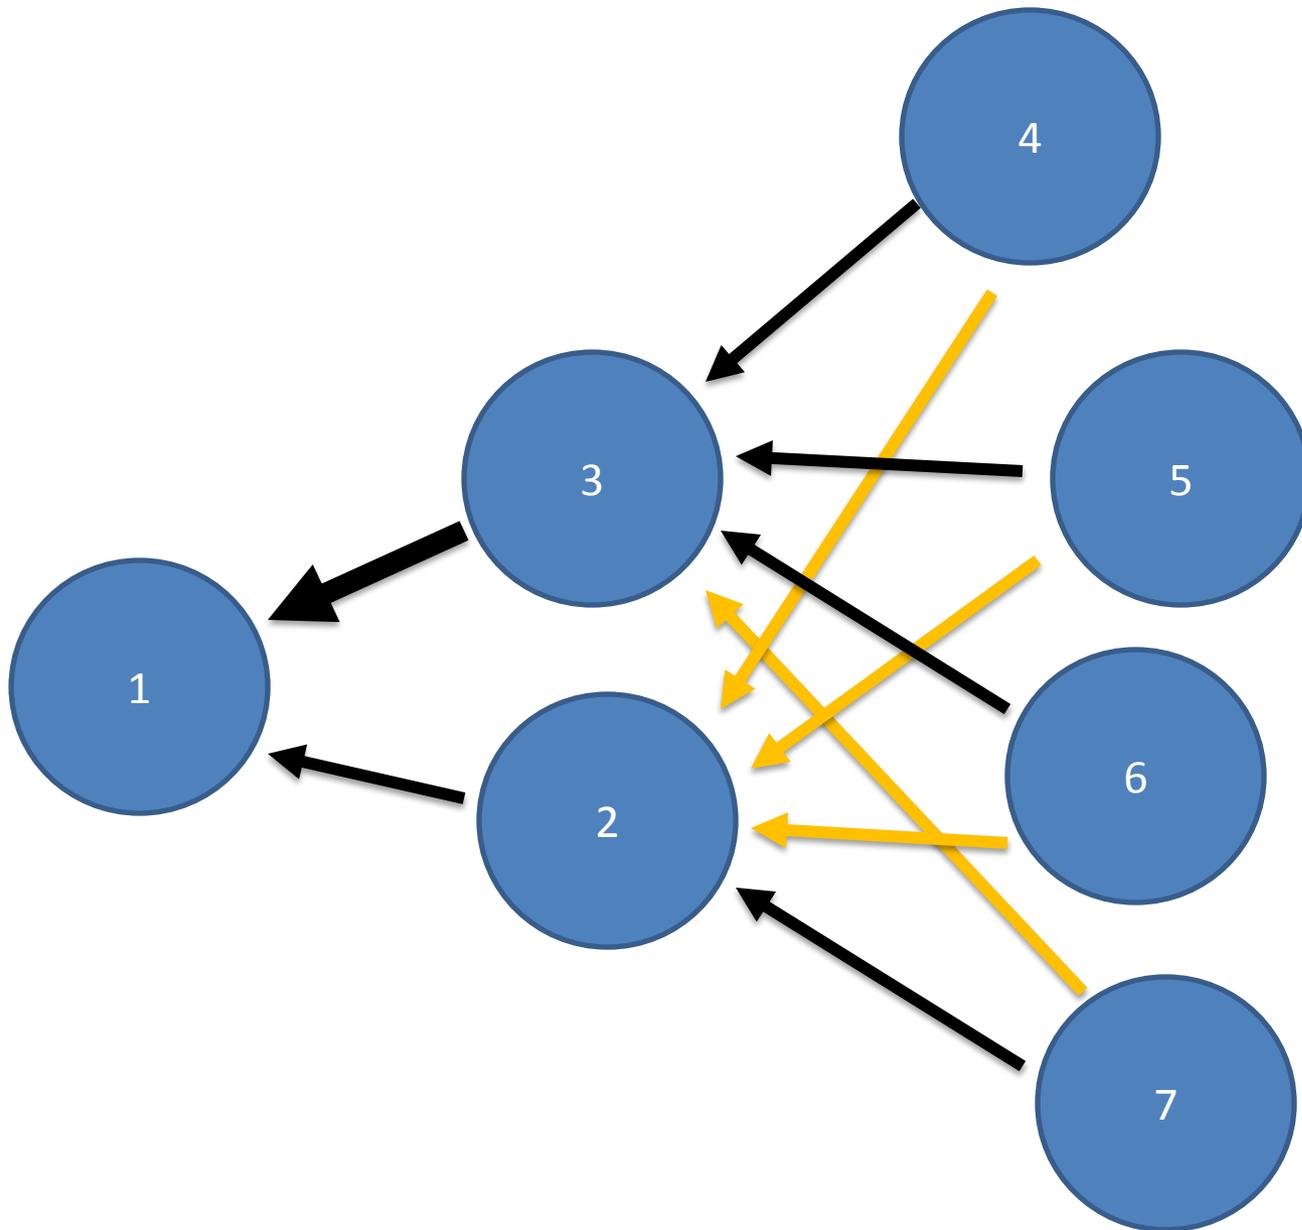
Iljitsch van Beijnum, Rolf Winter

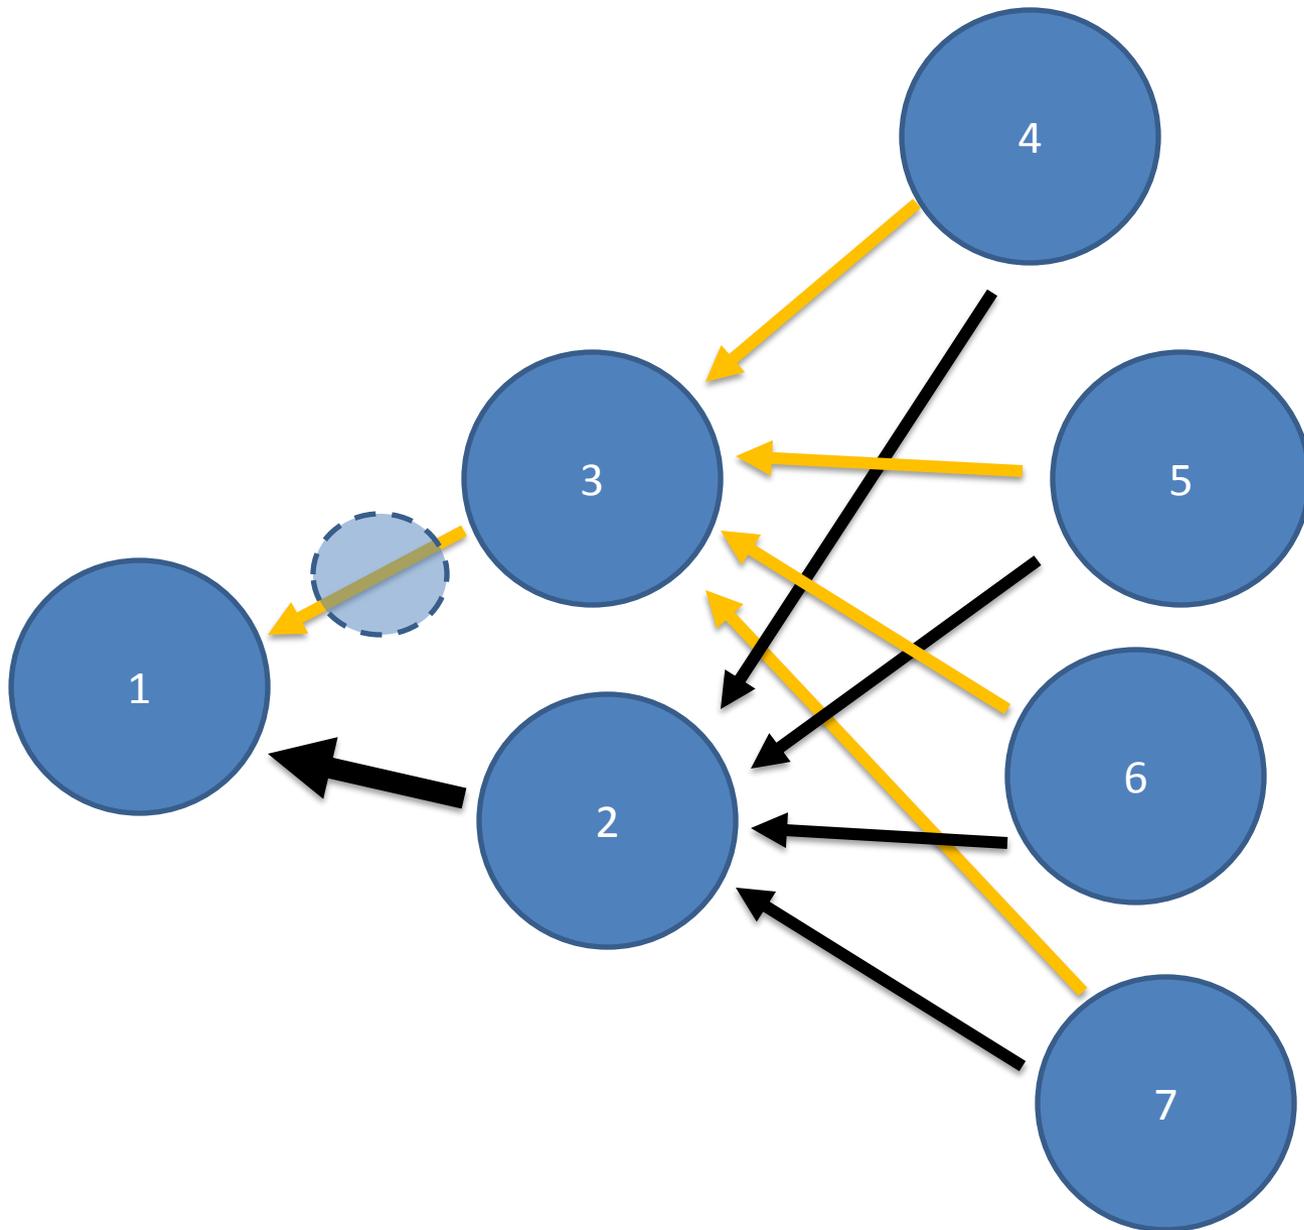
Problem

- Traffic engineering incoming traffic is either
 - Good for the AS that does the TE but suboptimal for everyone else (disaggregation)
 - Unworkable for the AS that attempts TE but has no impact on the global routing tables (AS path prepending)
- >Engineering problem to find some workable solution in-between

Changes since last version

- Everything except the motivation has changed more or less





Adding a cost metric

- Prepending too coarse, especially on the AS-level graph that we have
- Need to be able to influence a small set of Ases' decision without effecting the other ASes
- Need a place in the decision process to act on the cost metric
- Need to signal the cost metric across AS boundaries

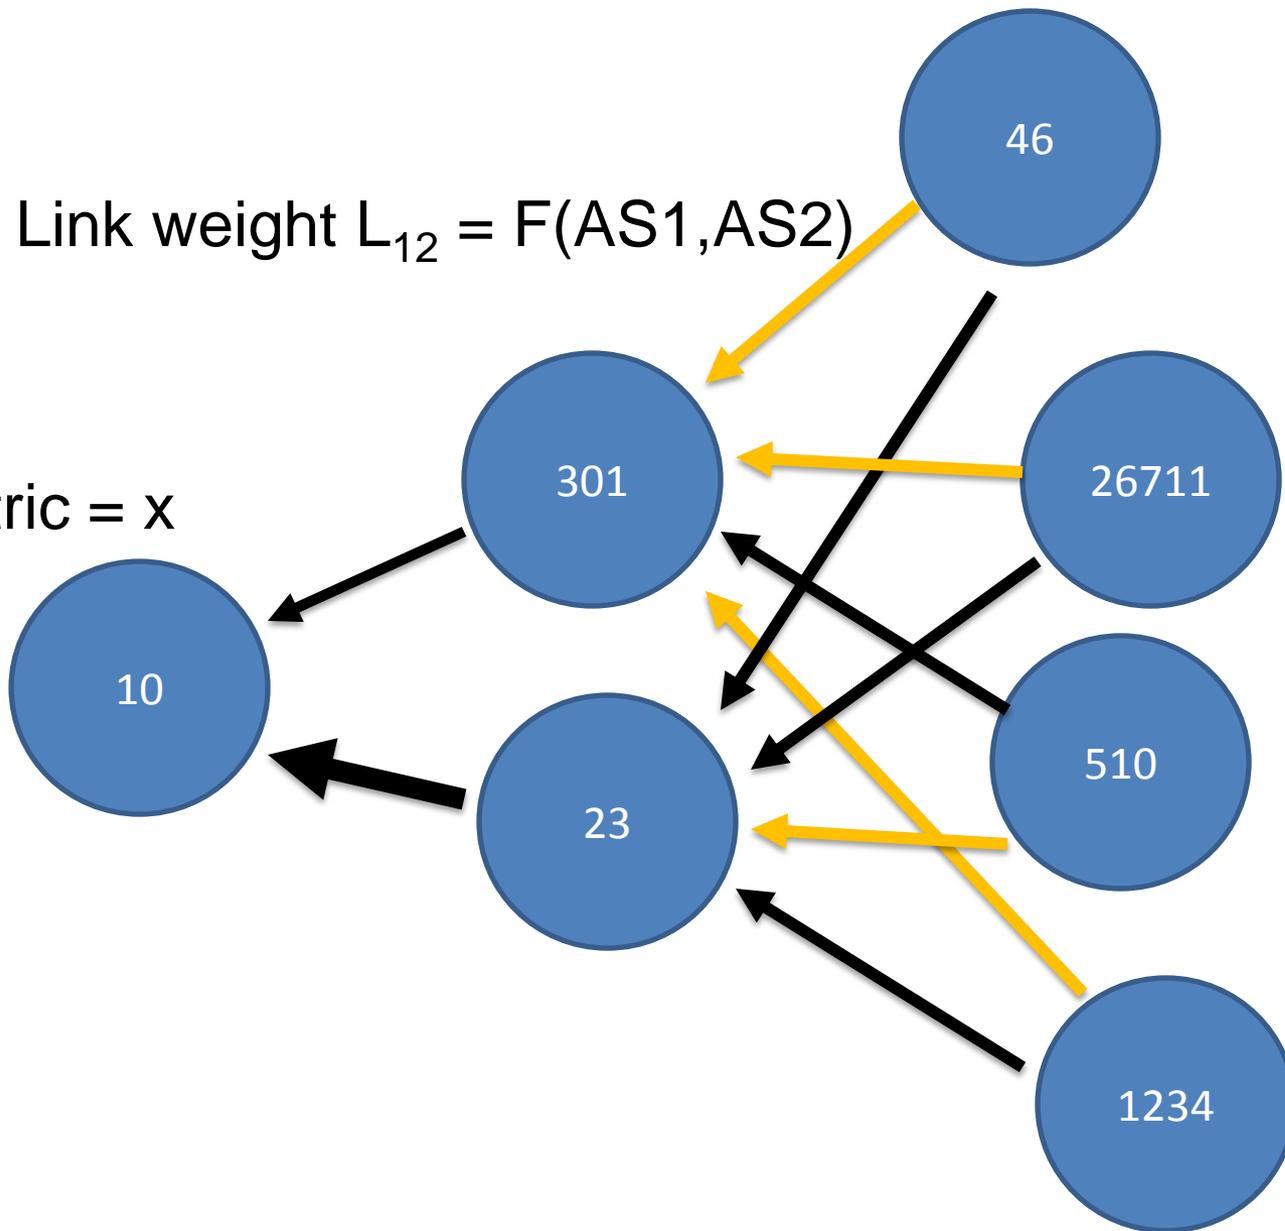
The (new) metric

- A new attribute: AS cost metric
 - Added by the origin AS, not changed after that any more
- Used in the decision process somewhere below the AS path length comparison
 - Not a replacement for it any more
- Used to influence a locally computed “link weight”

Decision process => $\min(L_i + x)$

Link weight $L_{12} = F(AS1, AS2)$

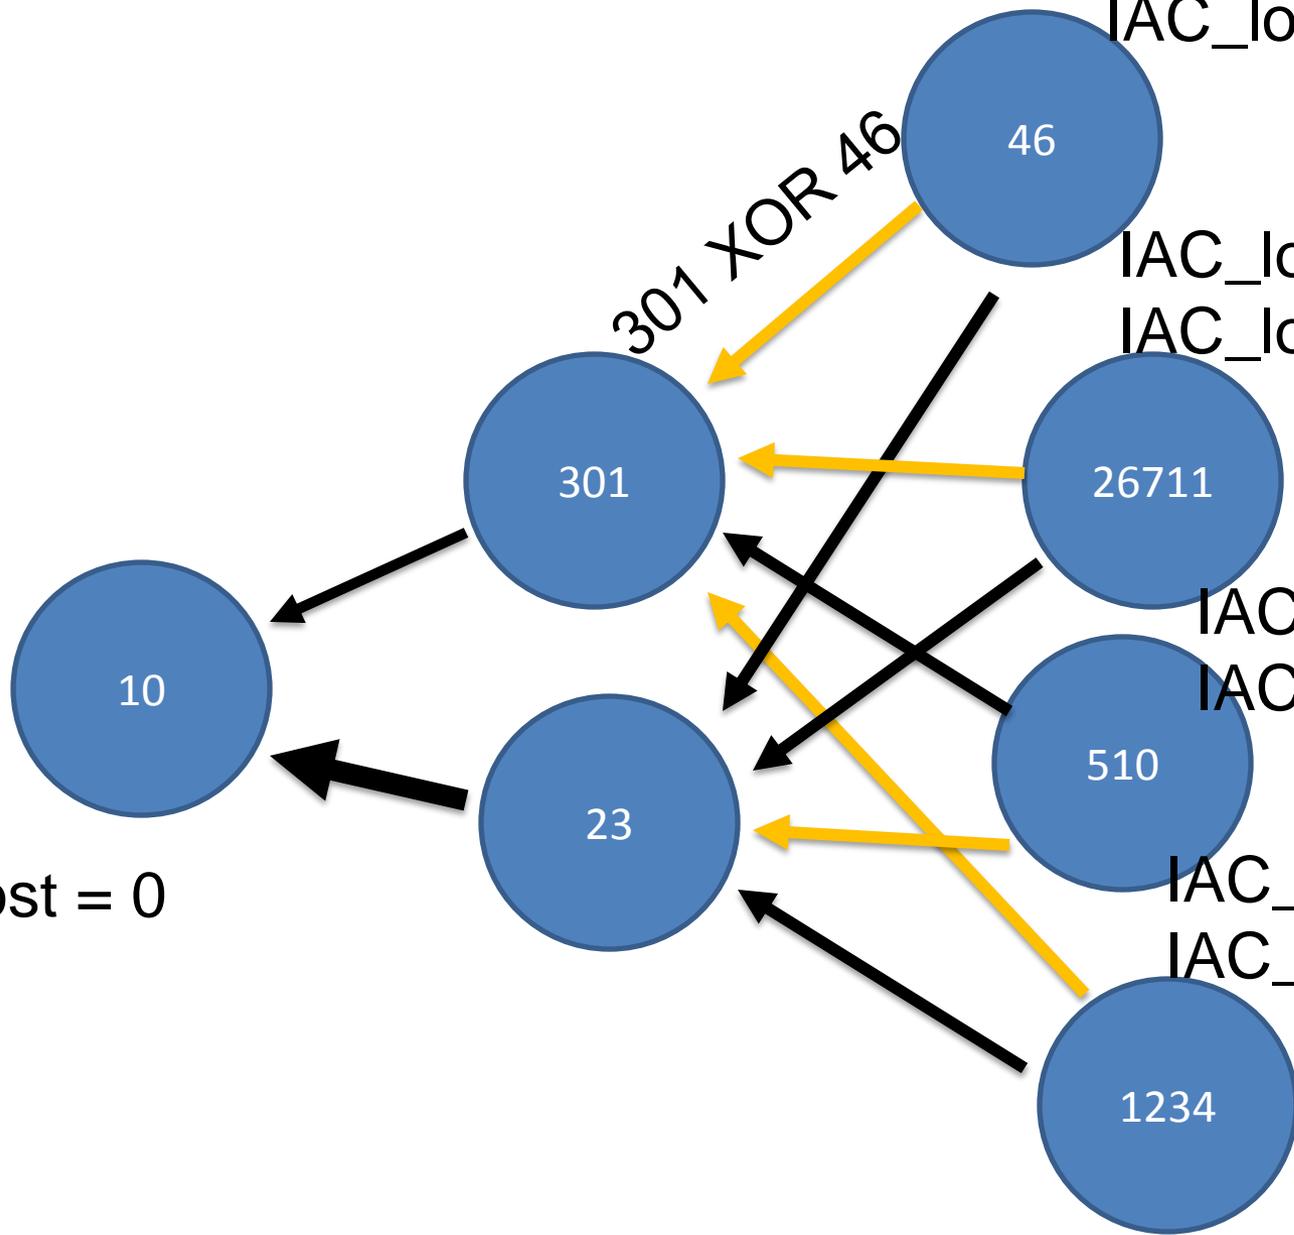
Metric = x



Link weight = IAC_local

IAC_local301 = 259
IAC_local23 = 57

AS cost = 0



IAC_local301 = 27002
IAC_local23 = 26688

IAC_local301 = 211
IAC_local23 = 489

IAC_local301 = 1535
IAC_local23 = 1221

IAC_local301 = 259

IAC_local23 = 57

202

301 XOR 46

IAC_local301 = 27002

IAC_local23 = 26688

314

IAC_local301 = 211

IAC_local23 = 489

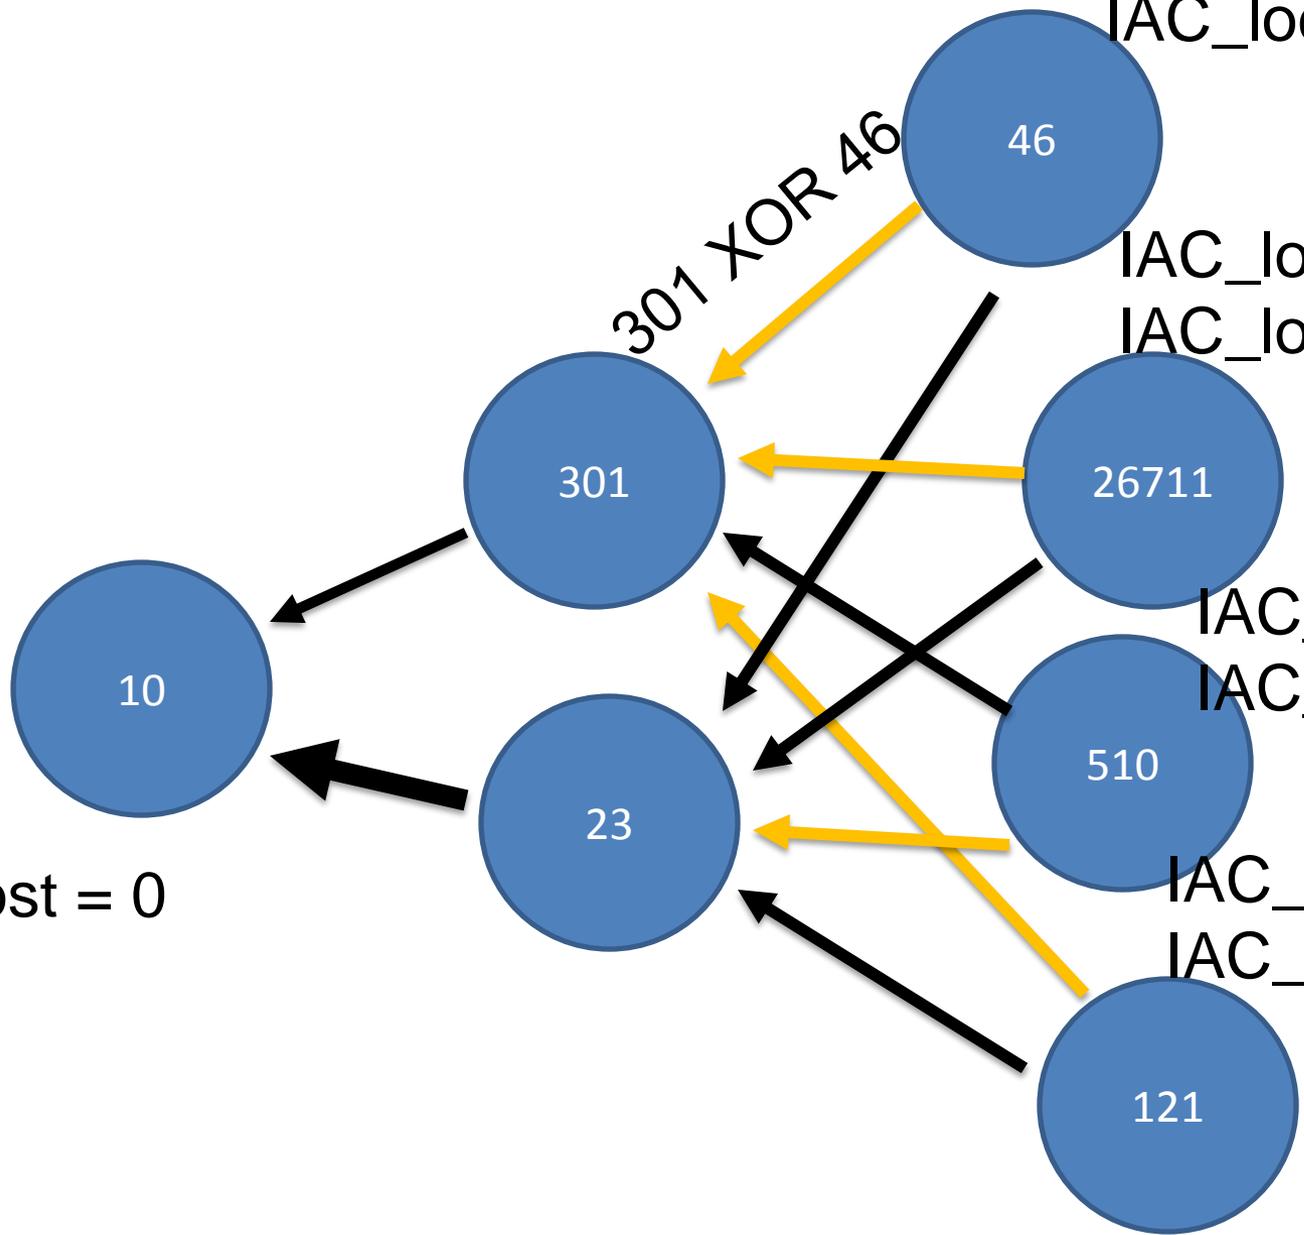
278

IAC_local301 = 340

IAC_local23 = 110

230

AS cost = 0



IAC_local301 = 259

IAC_local23 = $\frac{57 + 203}{1}$

1

301 XOR 46

IAC_local301 = 27002

IAC_local23 = $\frac{26688 + 203}{211}$

211

IAC_local301 = 211

IAC_local23 = $\frac{489 + 203}{481}$

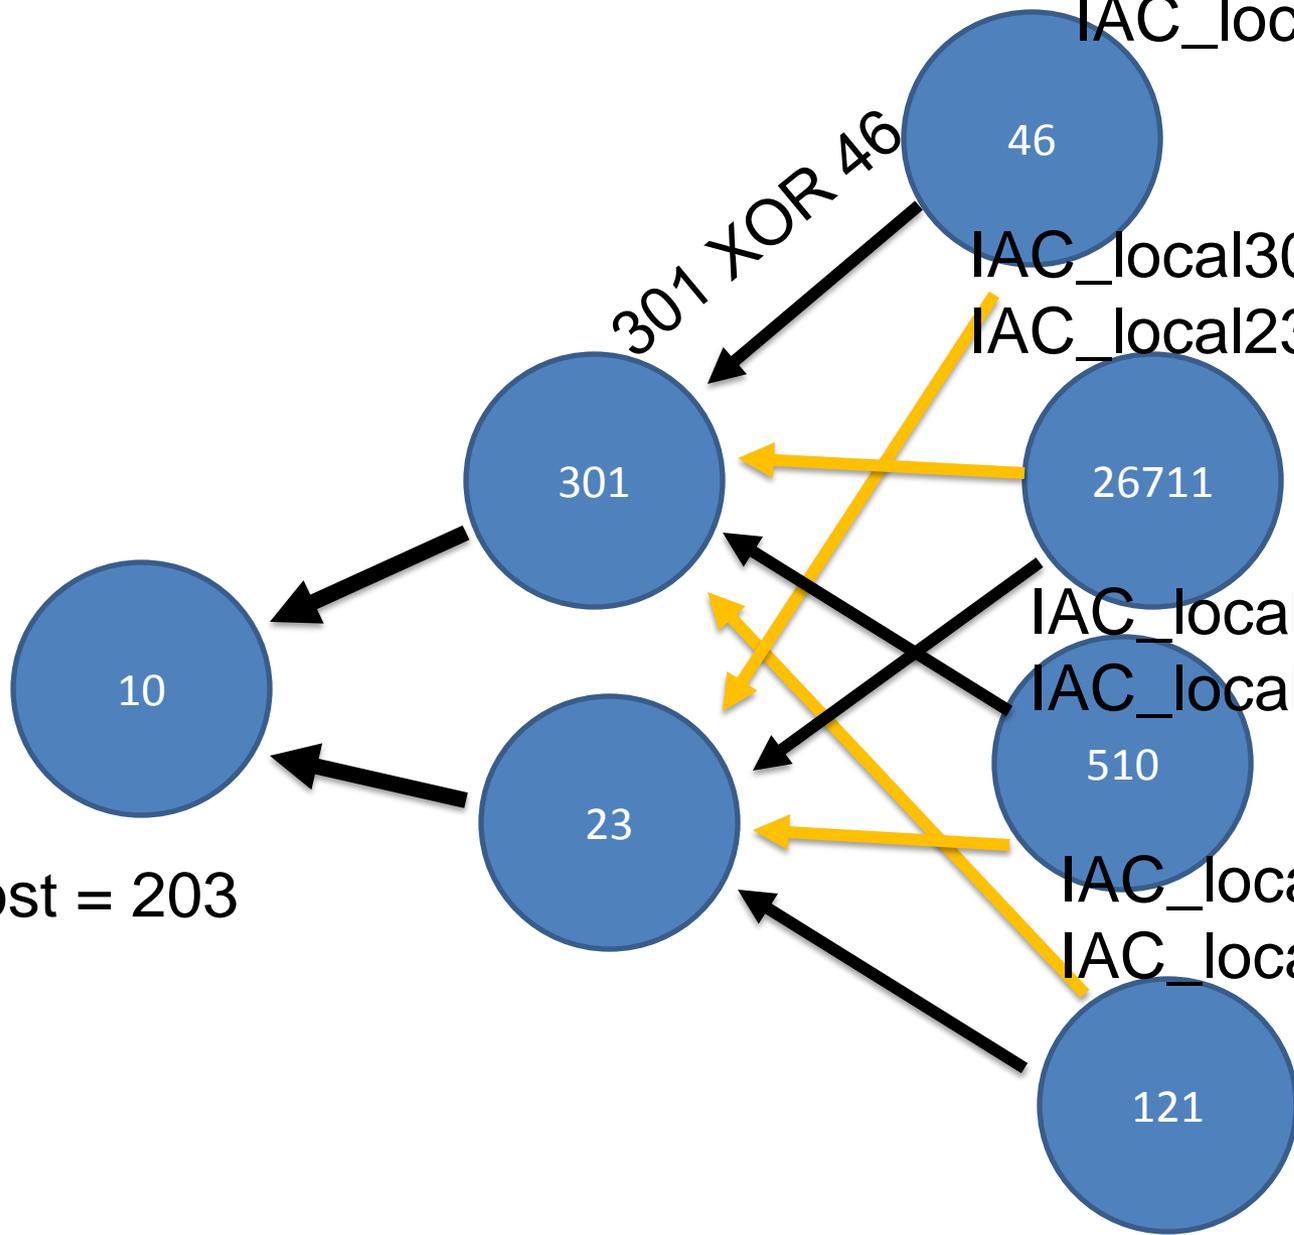
481

IAC_local301 = 340

IAC_local23 = $\frac{110 + 203}{27}$

27

AS cost = 203



IAC in the BGP decision process

Local
preference

AS path
length

IAC

- We want have reasonably short paths
- Not allow remote ASes to do things they cannot do today
- Only apply when the local operator “does not care”

Next steps

- Write it up in more detail, explain the reasoning behind some decisions we made more clearly
- Ask for feedback from the operational community
- Evaluate the idea's potential
- Go back to IDR