中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

中国科学院大学
University of Chinese Academy of Sciences

Encoding Route Origin Authorizations for Flexible and Fine-Grained Management

Yanbiao Li

CNIC, CAS / UCAS

BGP is important yet vulnerable

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

Border Gateway Protocol (BGP) is one of the key building blocks of the global Internet

a network of
Autonomous Systems (ASes)

CAIDA'S IPV4 AS CORE GRAPH
JANUARY 2020

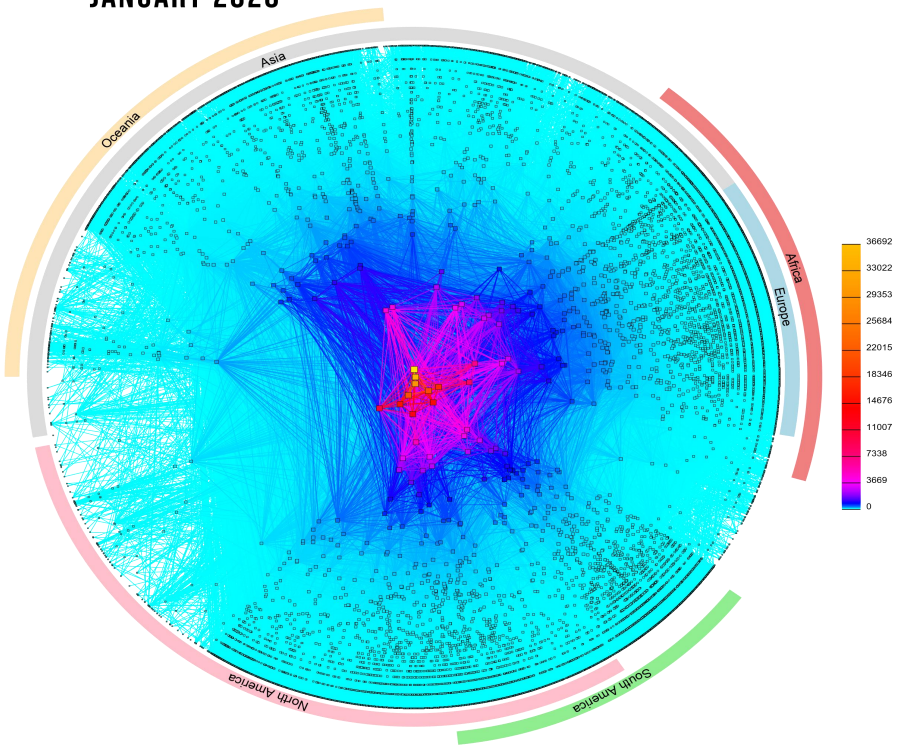

source: https://www.caida.org/projects/cartography/as-core/pics/2020/ascore-2020-ipv4-standalone.png

Mutually Agreed Norms for Routing Security (MANRS) 27 April 2018 EN ES

What Happened? The Amazon Route 53 BGP Hijack to Take Over Ethereum Cryptocurrency Wallets

By Aftab Siddiqui
Senior Manager, Internet Technology - Asia-Pacific

Doug Madory
@DougMadory

From 12:05-12:50 UTC, RU telecom RTComm (AS8342) hijacked a prefix (104.244.42.0/24) belonging to Twitter.

The hijack didn't propagate far due to a RPKI ROA which asserted AS13414 was the rightful origin.

This is the same prefix hijacked during the coup in Myanmar last year.

Prefix and sub-prefix hijacks

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

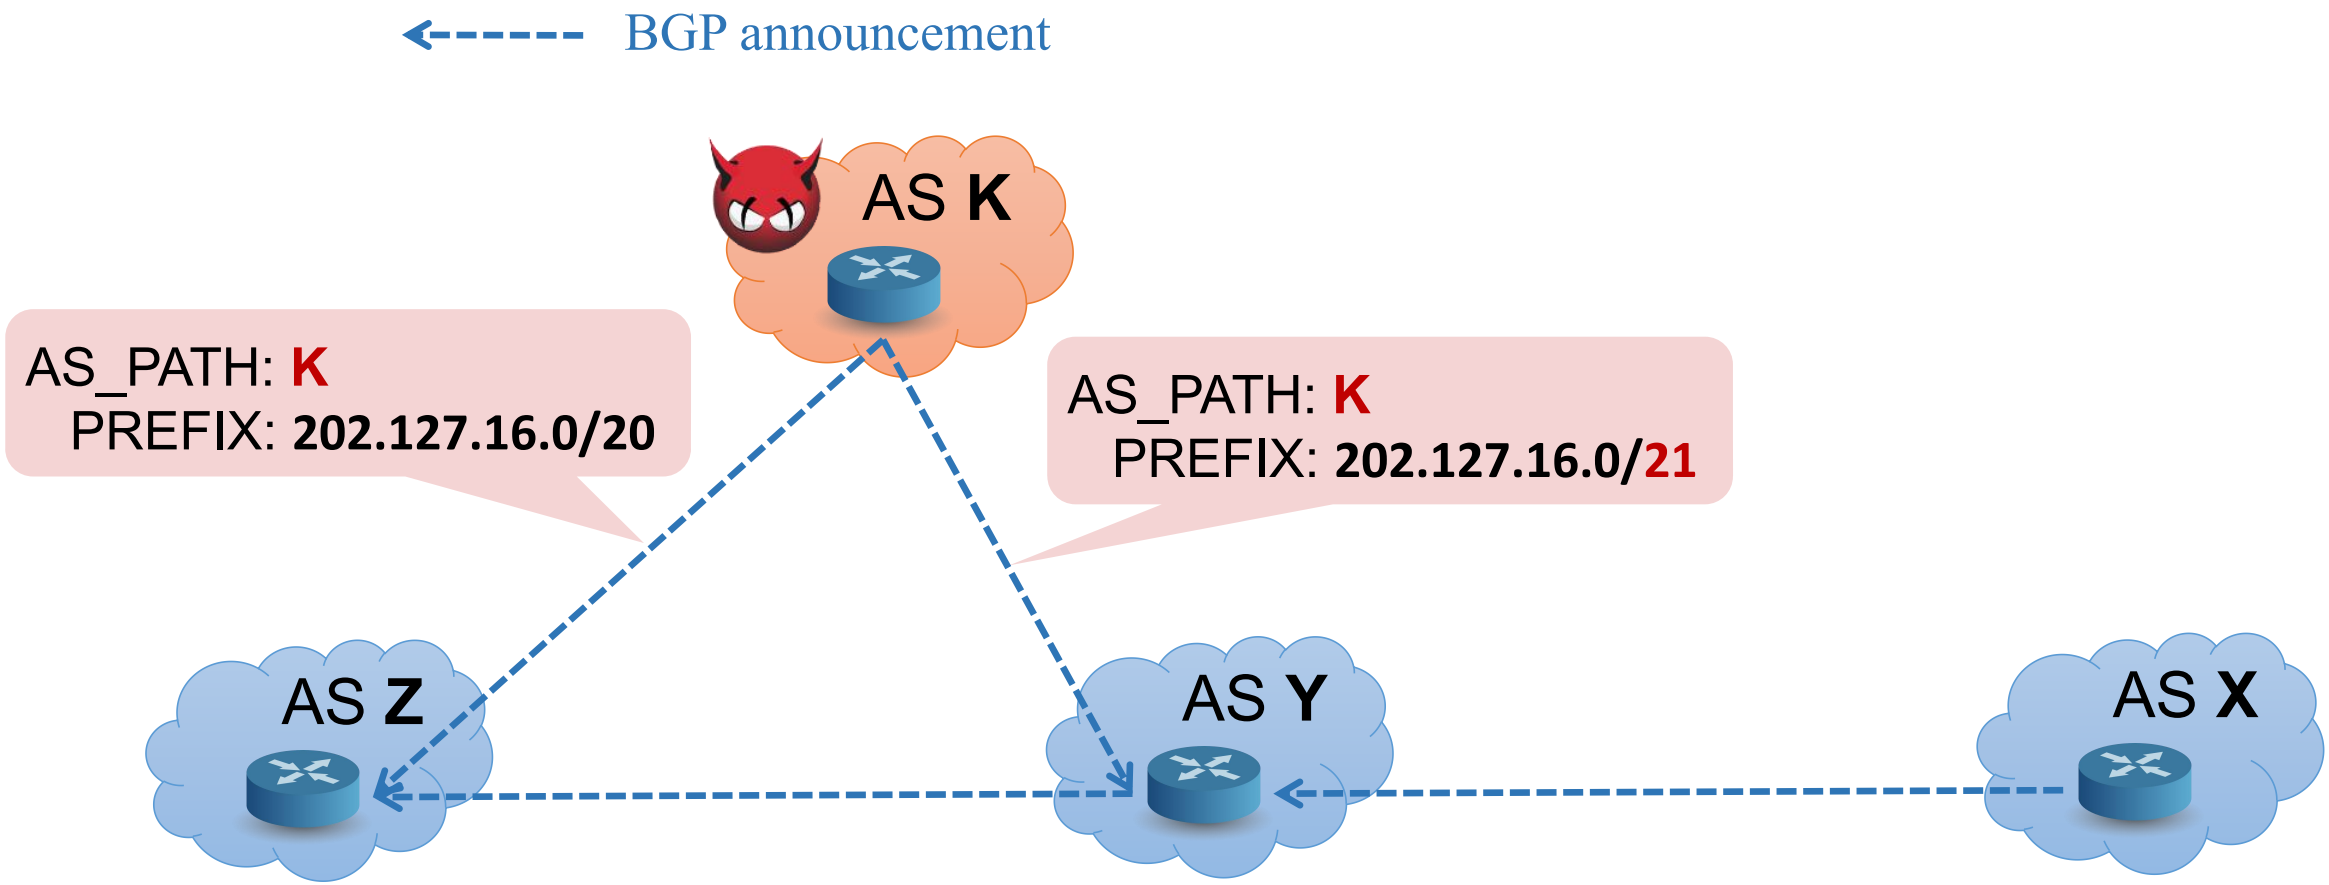

Prefix	Path
202.127.16.0/20	Y->X
202.127.16.0/20	K

shorter path

Prefix	Path
202.127.16.0/20	X
202.127.16.0/21	K

more specific

4 Stop route hijacks with RPKI

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

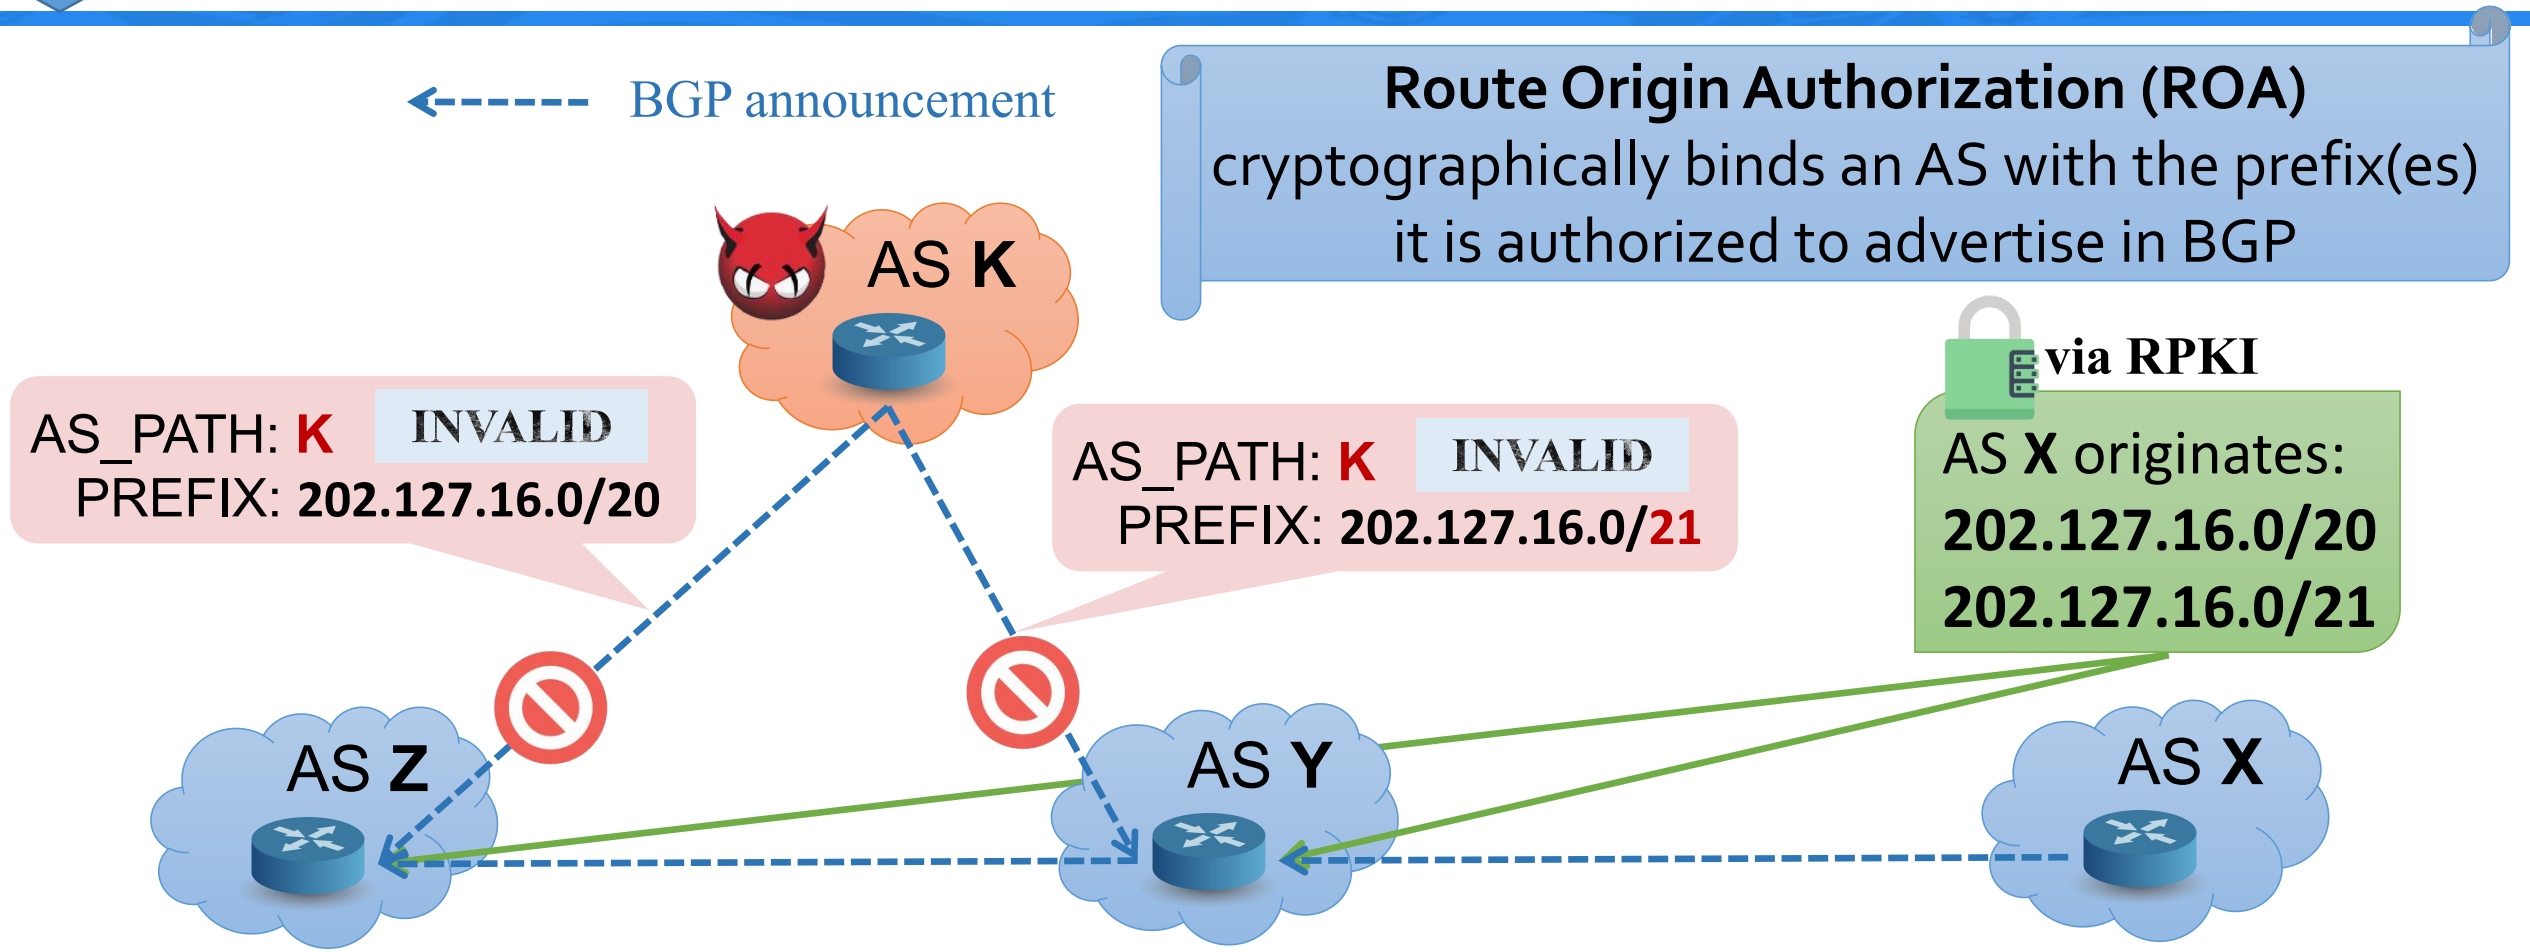

Prefix	Path
202.127.16.0/20	Y->X

Prefix	Path
202.127.16.0/20	X

ROA encoding

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

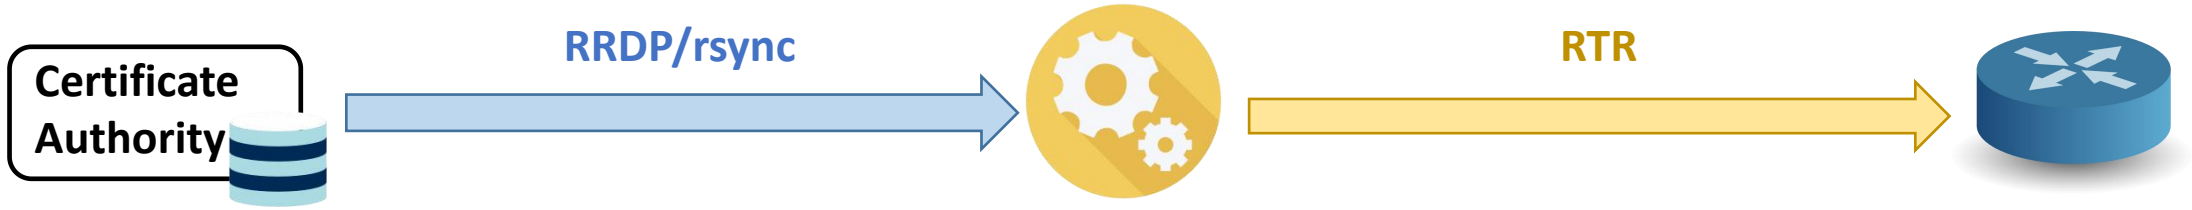

Publication Point | **Relying Party** | **BGP Speaker**

```
RouteOriginAttestation ::= SEQUENCE {
   version [0] INTEGER DEFAULT 0,
   asID  ASID,
   ipAddrBlocks SEQUENCE (SIZE(1..MAX)) OF ROAIPAddressFamily }

ASID ::= INTEGER

ROAIPAddressFamily ::= SEQUENCE {
   addressFamily OCTET STRING (SIZE (2..3)),
   addresses SEQUENCE (SIZE (1..MAX)) OF ROAIPAddress }

ROAIPAddress ::= SEQUENCE {
   address IPAddress,
   maxLength INTEGER OPTIONAL }

IPAddress ::= BIT STRING
```

ROA eContent (RFC 6482)

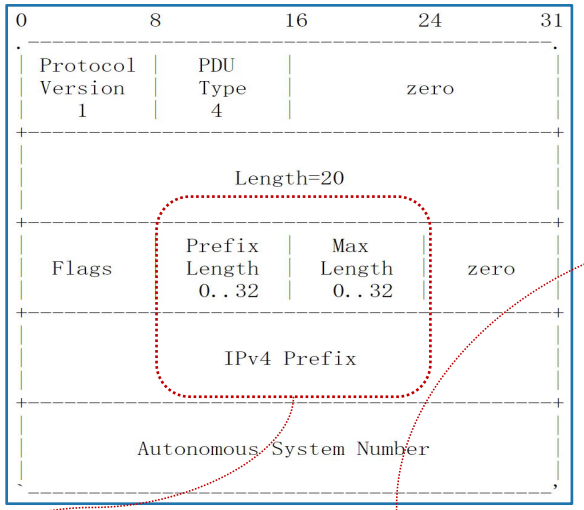

RTR Prefix (v4/v6) **PDU** (RFC 8210)

The BGP speaker loads validated objects from the cache into local storage. The objects loaded have the content (**IP address, prefix length, maximum length**, origin AS number). We refer to such a locally stored object as a "Validated ROA Payload" or "VRP".

ROA Payload (RFC 6811)

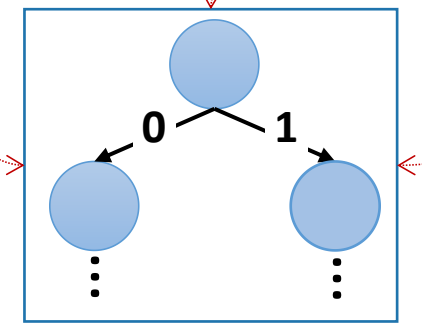

ROA prefix block: [**IP prefix** (addr / len)**, maximum length**] represents a sub-tree of the IP prefix trie

Security: 'holey' ROA

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

Operational granularity of constructing ROAs: a set of IP prefixes

202.127.16.0/20

AS X: 202.127.16.0/**20**, **21**

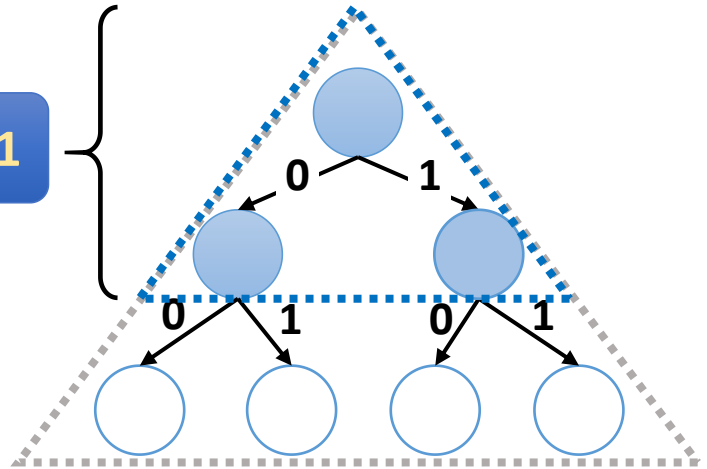

AS X: 202.127.16.0/**20**, **22**

maxLen + 1:
2^n IP prefixes are authorized

'Holey' ROA

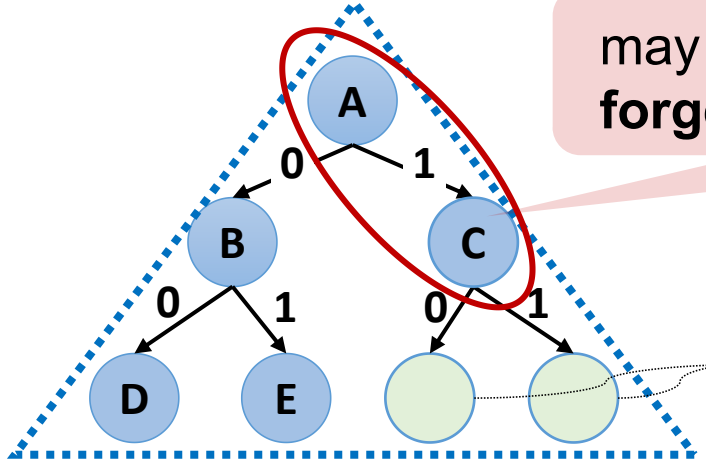

may suffer from **forged-origin sub-prefix hijacks** (RFC 9319)

AS X **does NOT annoucne** these routes in BGP; leaving two 'holes'

Statistics of 'holey' ROAs

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

500K ROAs published in 2020 are tracked
data: https://ftp.ripe.net/rpki/

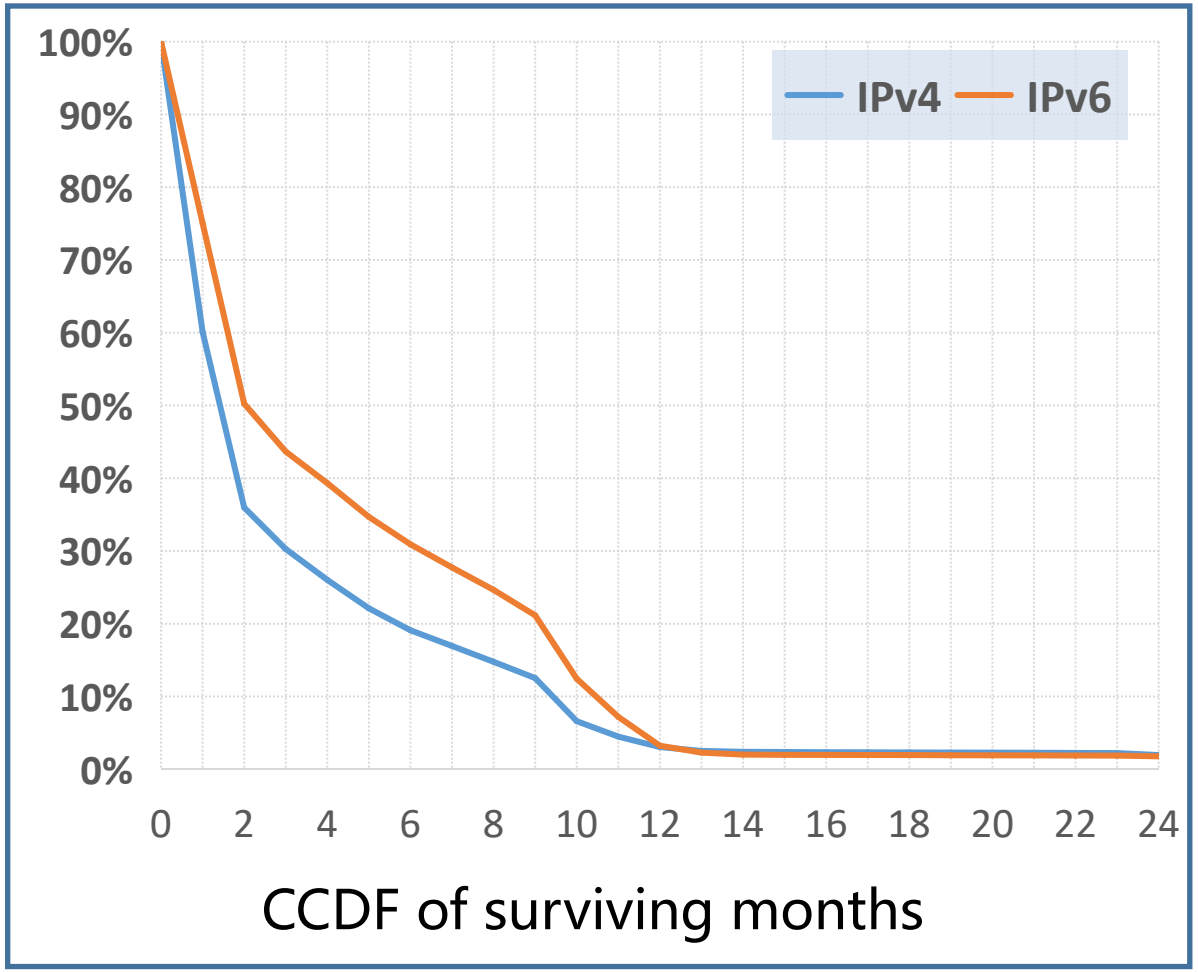

For each ROA prefix block, # of "holes" is counted as the number of its prefixes that are not in the RIB.
hole ratio = (# of holes) / (# of prefixes in the block)
data: https://www.ripe.net/analyse/raw-data-sets

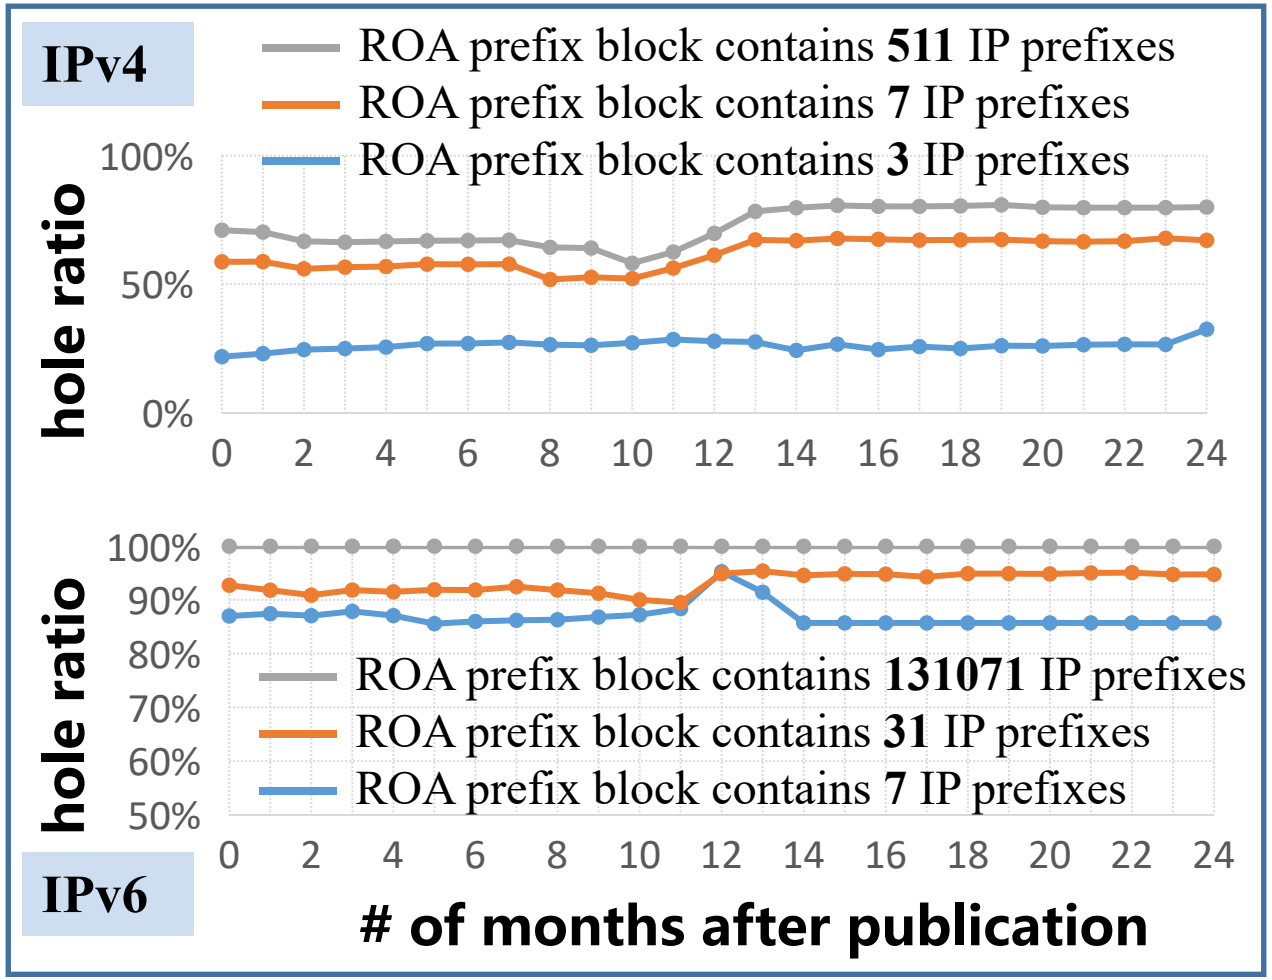

Scalability: exceesive withdrawal

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

Operational granularity of updating ROAs:
a whole ROA

AS X is authorized to originate **7 routes** in BGP

AS X: 202.127.16.0/**20, 22**

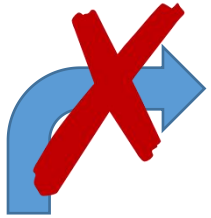

Withdraw part (3) of authorizations granted to X (202.127.24.0/**21~22**)

202.127.16.0/20

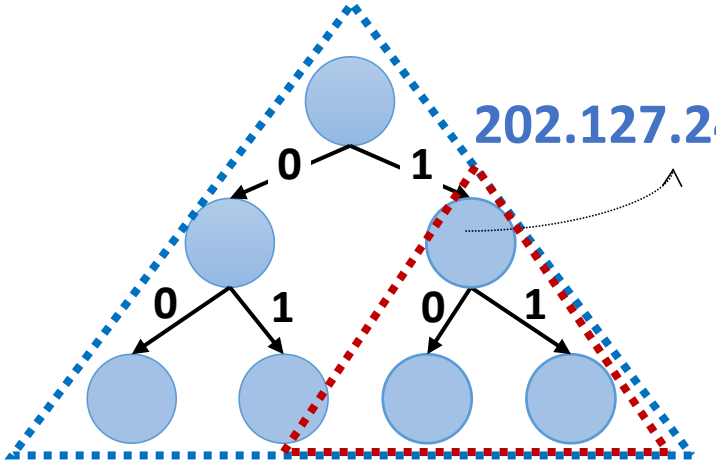

administrator

W **AS X**: 202.127.24.0/**21, 22**

trying to withdraw a non-existing ROA

```
root@ca0:/# krillc roas list --ca cert0
202.127.16.0/20-22 => 1

root@ca0:/# krillc roas update --remove '202.127.24.0/21-22 => 1' --ca cert0
Delta rejected:

Cannot remove the following unknown ROAs:
  202.127.24.0/21-22 => 1

root@ca0:/#
```

Flexibility: exceesive withdrawal

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

Operational granularity of updating ROAs:
a whole ROA

AS X is authorized to originate **7 routes** in BGP

AS X: 202.127.16.0/**20, 22**

Withdraw part (3) of authorizations granted to X (202.127.24.0/**21~22**)

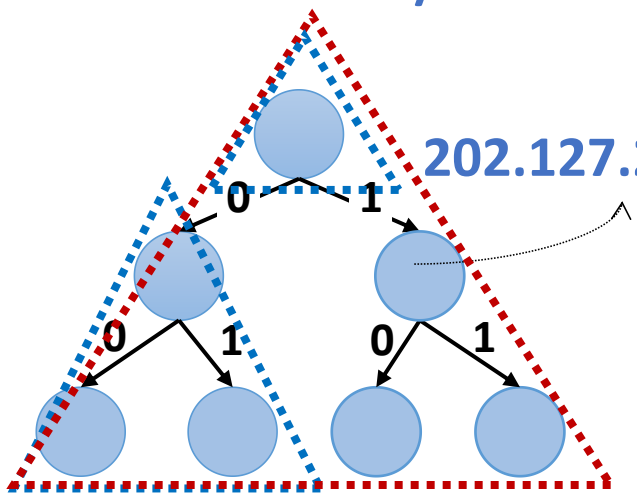

administrator

Excessive withdrawal

W **AS X**: 202.127.24.0/**21, 22**

trying to withdraw a non-existing ROA

```
root@ca0:/# krillc roas list --ca cert0
202.127.16.0/20-22 => 1

root@ca0:/# krillc roas update --remove '202.127.24.0/21-22 => 1' --ca cert0
Delta rejected:

Cannot remove the following unknown ROAs:
  202.127.24.0/21-22 => 1

root@ca0:/#
```

W **AS X**: 202.127.16.0/**20, 22**

A **AS X**: 202.127.16.0/**20, 20**

A **AS X**: 202.127.16.0/**21, 22**

Statistics of exceesive withdrawal

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

We collected **VRPs** every day from 2022/11/1 to 2022/12/31, and calculated the **difference between every two consecutive days**, where we counted total withdrawals and **# of exceesive withdrawals**.

data: https://ftp.ripe.net/rpki/

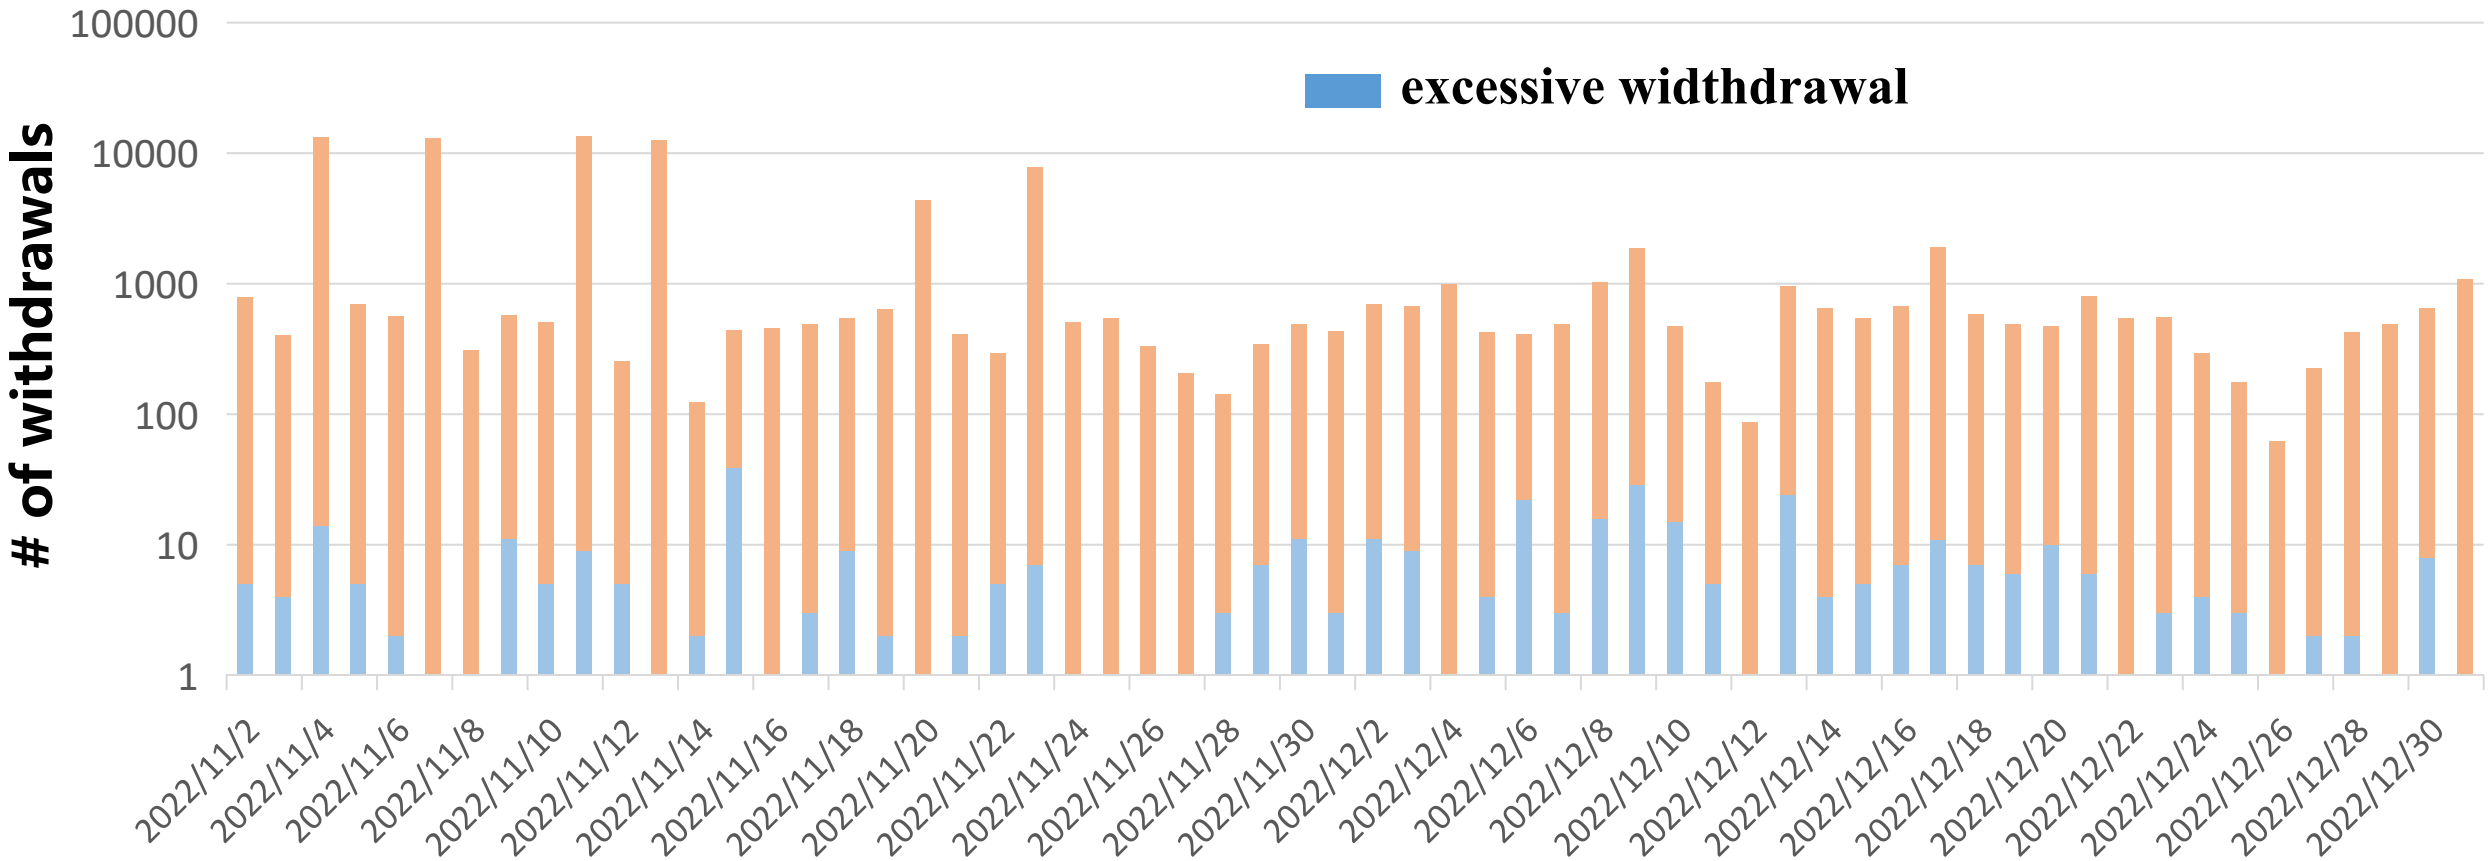

Resource allocation

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

The provider **X** allocates a sub- address block to its customer **Y**

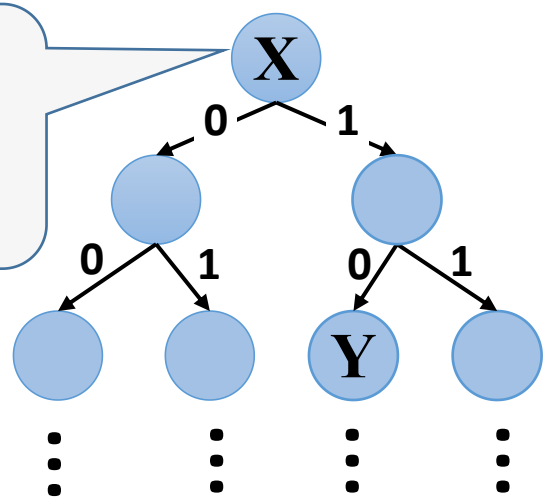

I **No ROA with X;** Y may or may not have

X's route is **unprotected**.

Once X is issued an ROA, the situation will fail into one of the **following cases**

	Y's prefix is included in X's ROA	Y's prefix is NOT included in X's ROA
No ROA with Y	**II** 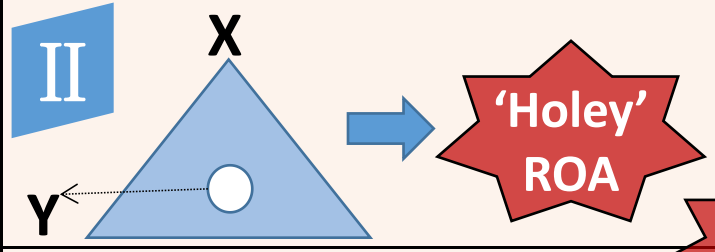	**IV** 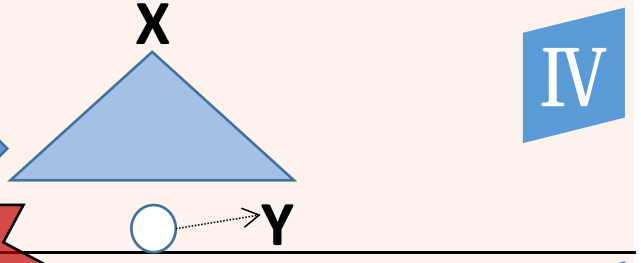
Y has an ROA	**III** 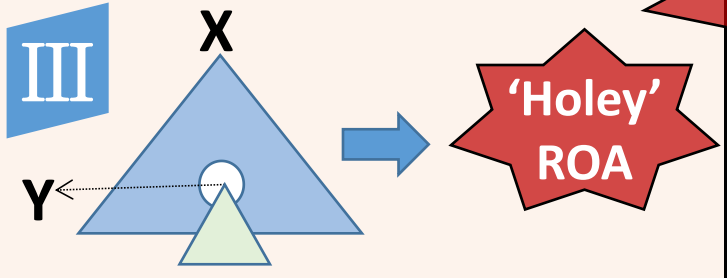	**V** 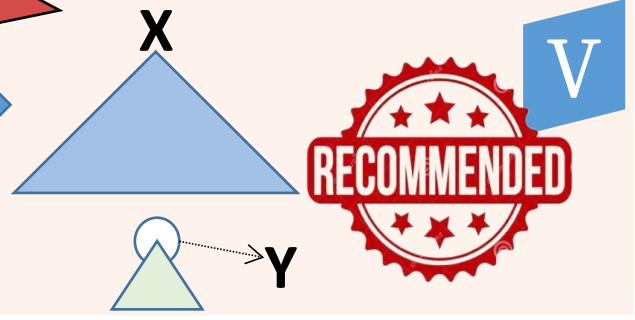

Reachibility issue

Y's routes will be identified as INVALID per RFC 6811

Statistics of overlapping routes

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

We collected 14 RIBs from 2022/1/14. For each RIB, we extracted **route-pairs** (X, Y) where X's prefix covers Y's prefix and they have **different origin ASNs**, and classified them into 5 cases with **VRPs collected on the same day**.

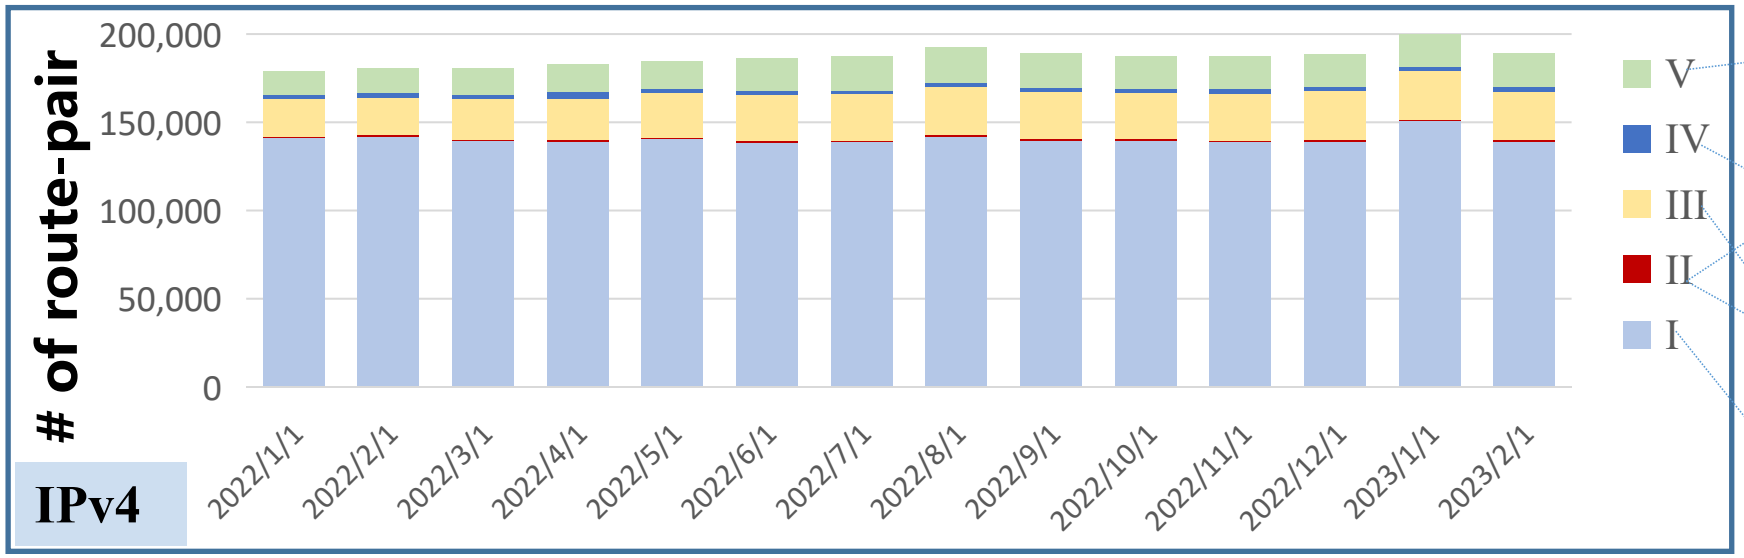

Recommended: 7.6%~10.5%

Reachibility issue: 1.6%~2.8%

Holey ROA: 12.2%~15.1%

Unprotected: 73.6%~78.8%

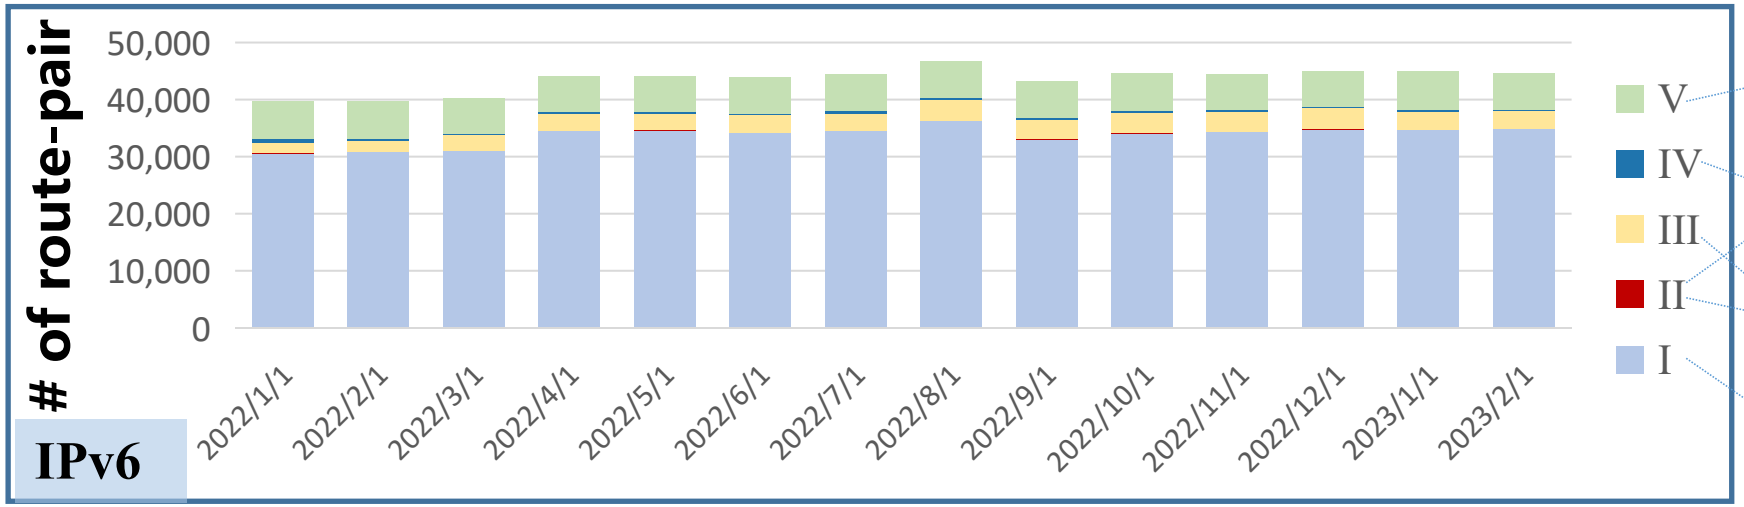

Recommended: 13.7%~16.6%

Reachibility issue: 0.7%~1.6%

Holey ROA: 4.8%~8.2%

Unprotected: 76.1%~78.6%

A brief summary

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

'Holey' ROA

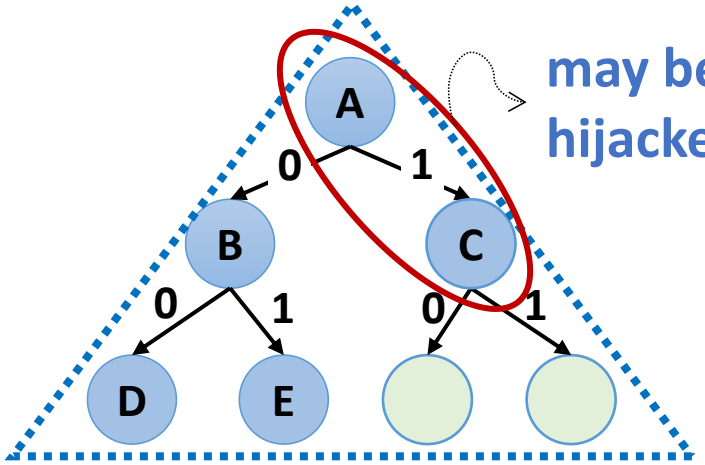

Excessive withdrawal

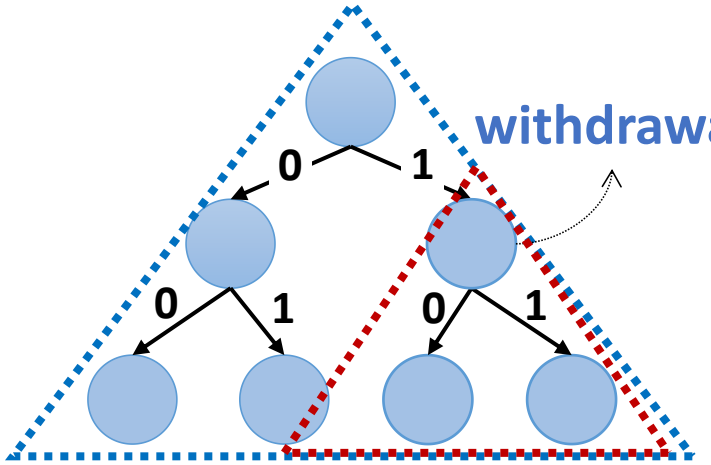

Reachibility issue

The provider **X** allocates a sub-address block to its customer **Y**

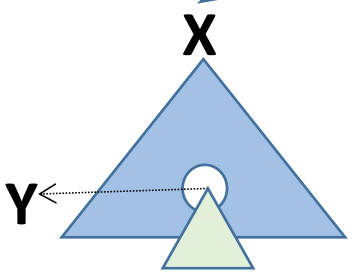

Recommendation-1: issue an ROA for Y

Recommendation-2: withdraw part of X's ROA in case it's overlapped with Y's ROA

Objective: enable ROA management at a fine graunlarity

an IP prefix

A simple solution

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

ΔL of an ROA: the difference of ***maxLen*** and IP **prefix length**
A simple solution: construct ROAs such that **ΔL of every ROA is 0**.

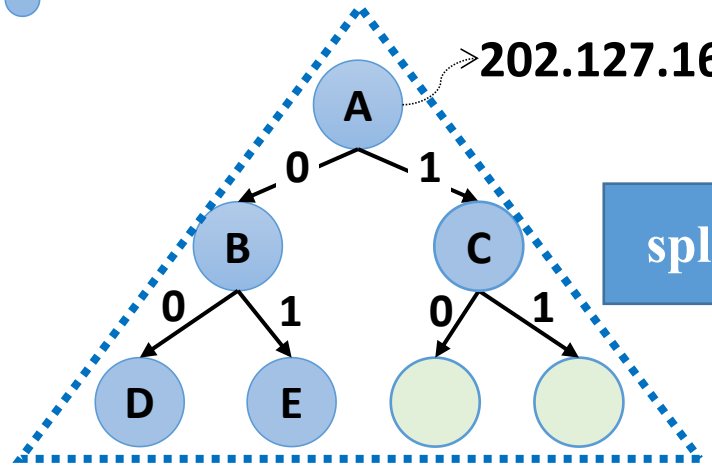

split into 5 entries

A	AS X: 202.127.16.0/**20, 20**
B	AS X: 202.127.16.0/**21, 21**
C	AS X: 202.127.24.0/**21, 21**
D	AS X: 202.127.16.0/**22, 22**
E	AS X: 202.127.20.0/**22, 22**

IPv4 ROAs of 2023.02.01

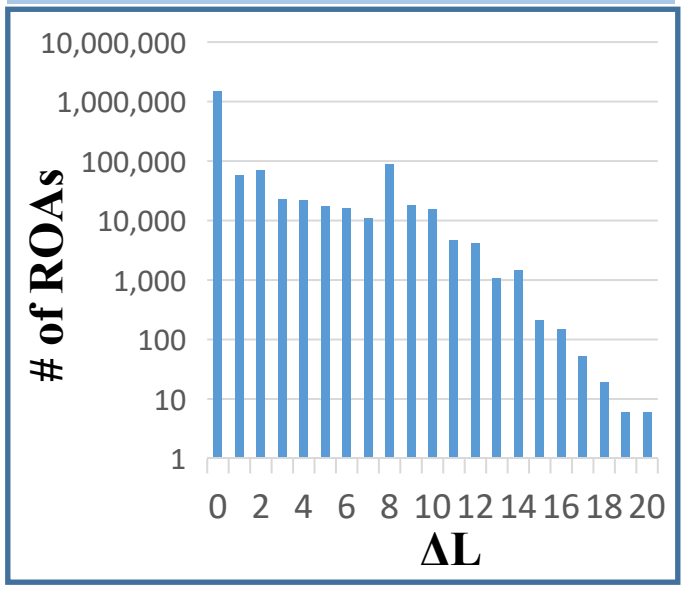

We collected ROA records from 2018.01 to 2022.12 (one set per month). data: https://ftp.ripe.net/rpki/

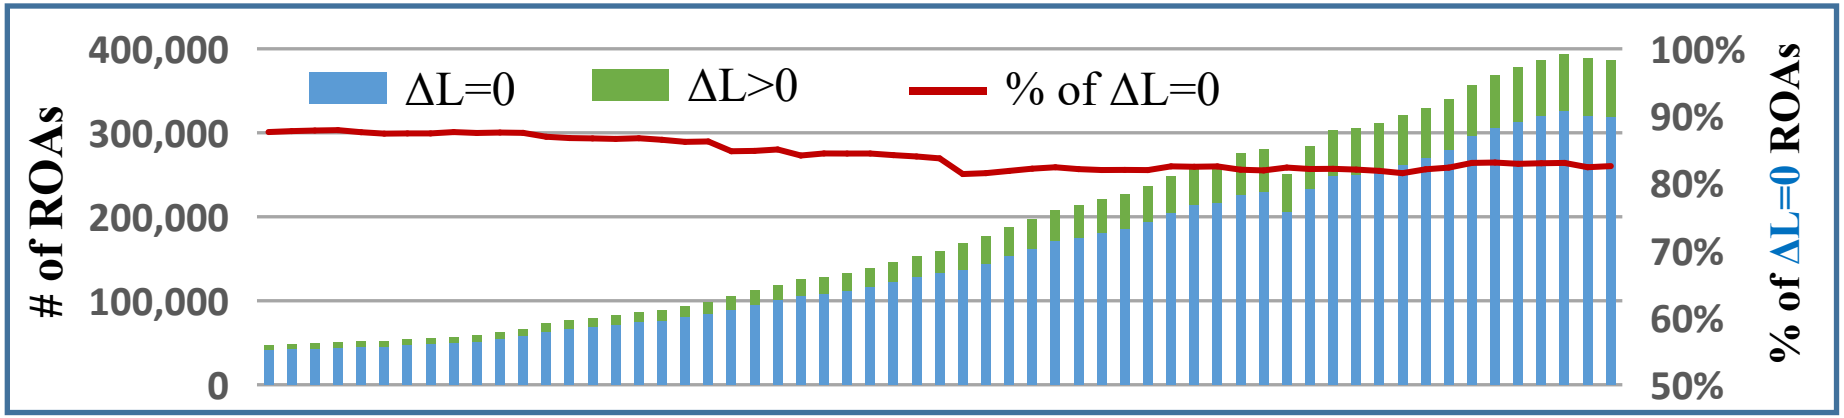

ΔL=N, may split into at most 2^N entries.

Scalability issue

Encode ROAs with bitmaps

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

Hanging ROA[1]: Enocdes an **ROA prefix block** with a **bitmap** where **a set bit** indicates the prefix corresponding to this bit is **authorized;** manage authorizations **bit by bit** (prefix by prefix)

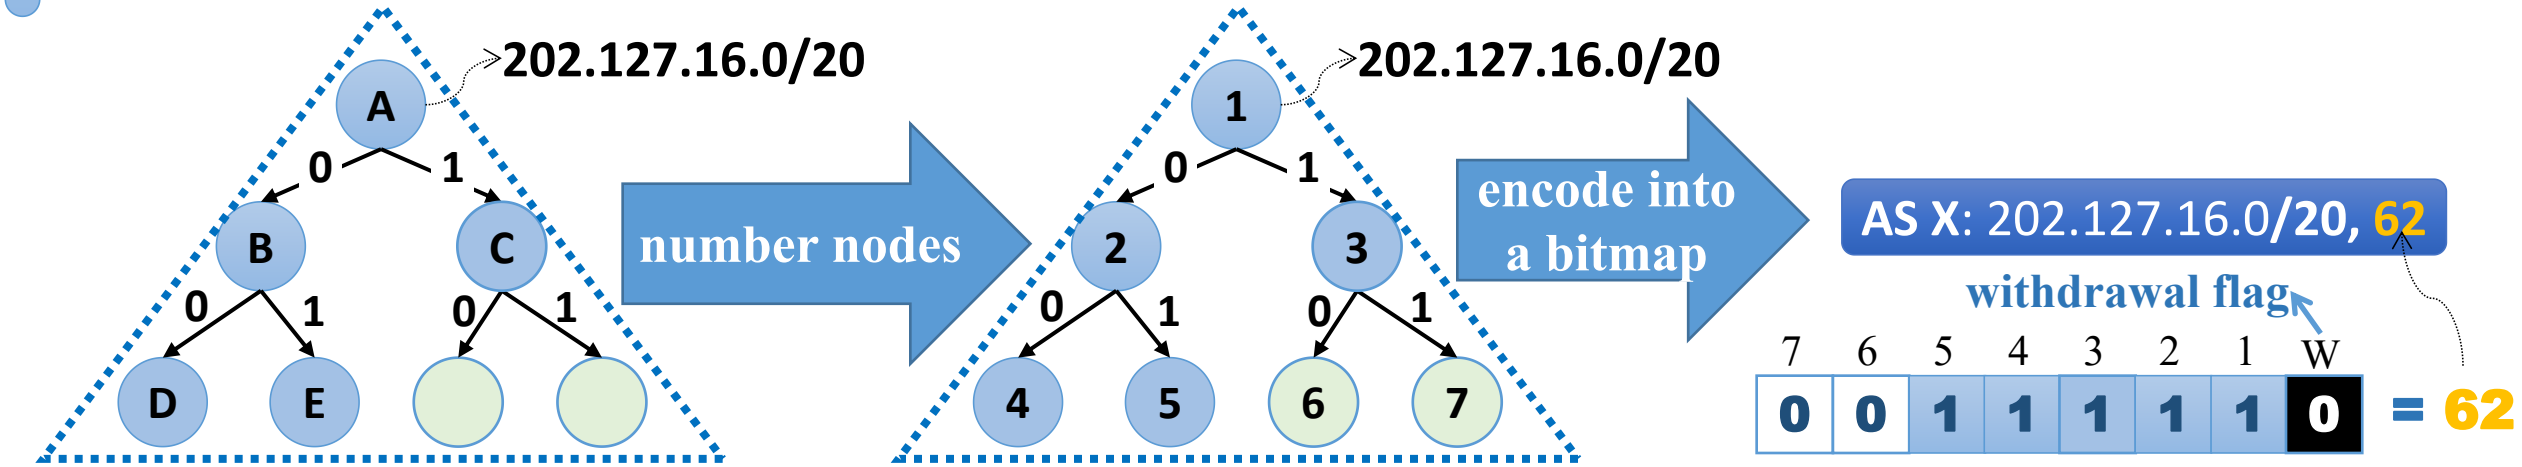

We collected VRPs and RIBs on Jan. 1st every year from 2001 to 2023.

2 data sets of authorized prefixes and 3 schemes

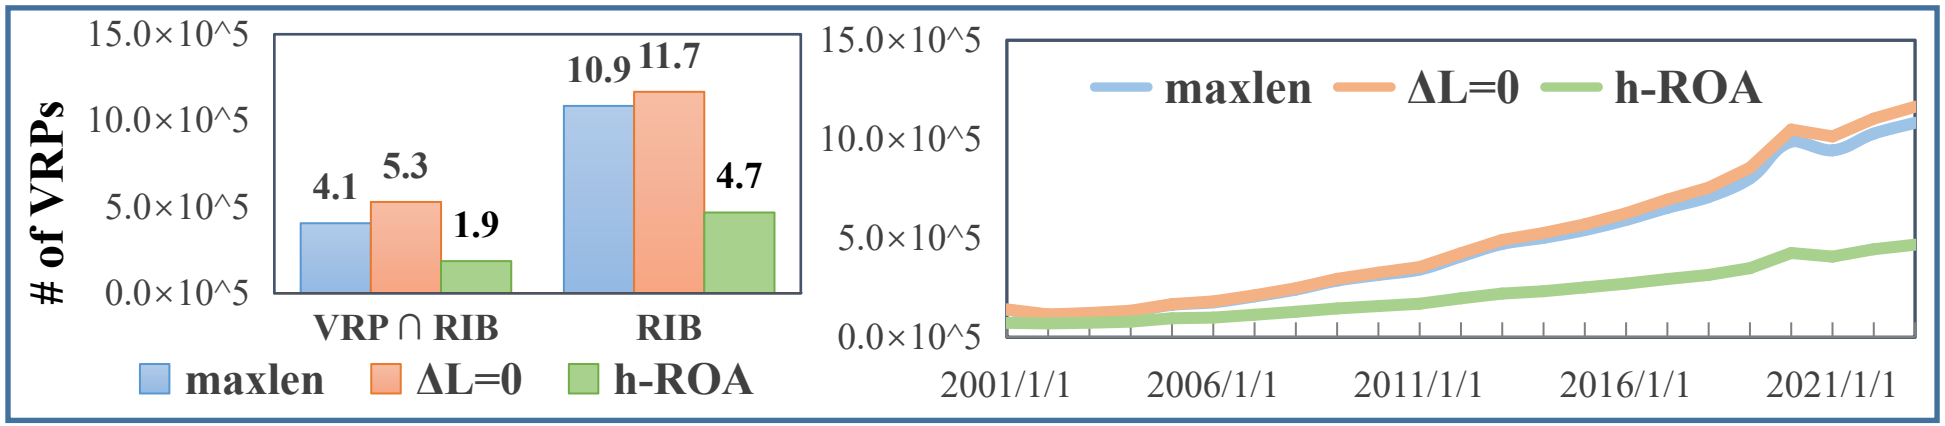

[1] The Hanging ROA: A Secure and Scalable Encoding Scheme for Route Origin Authorization[C]//IEEE INFOCOM 2022

Flexible ROA management

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

On basis of the bitmap encoding scheme, ROAs can be constructed and updated at **a prefix-level granularity**, enabling very flexible management with efficient **bitwise operations**.

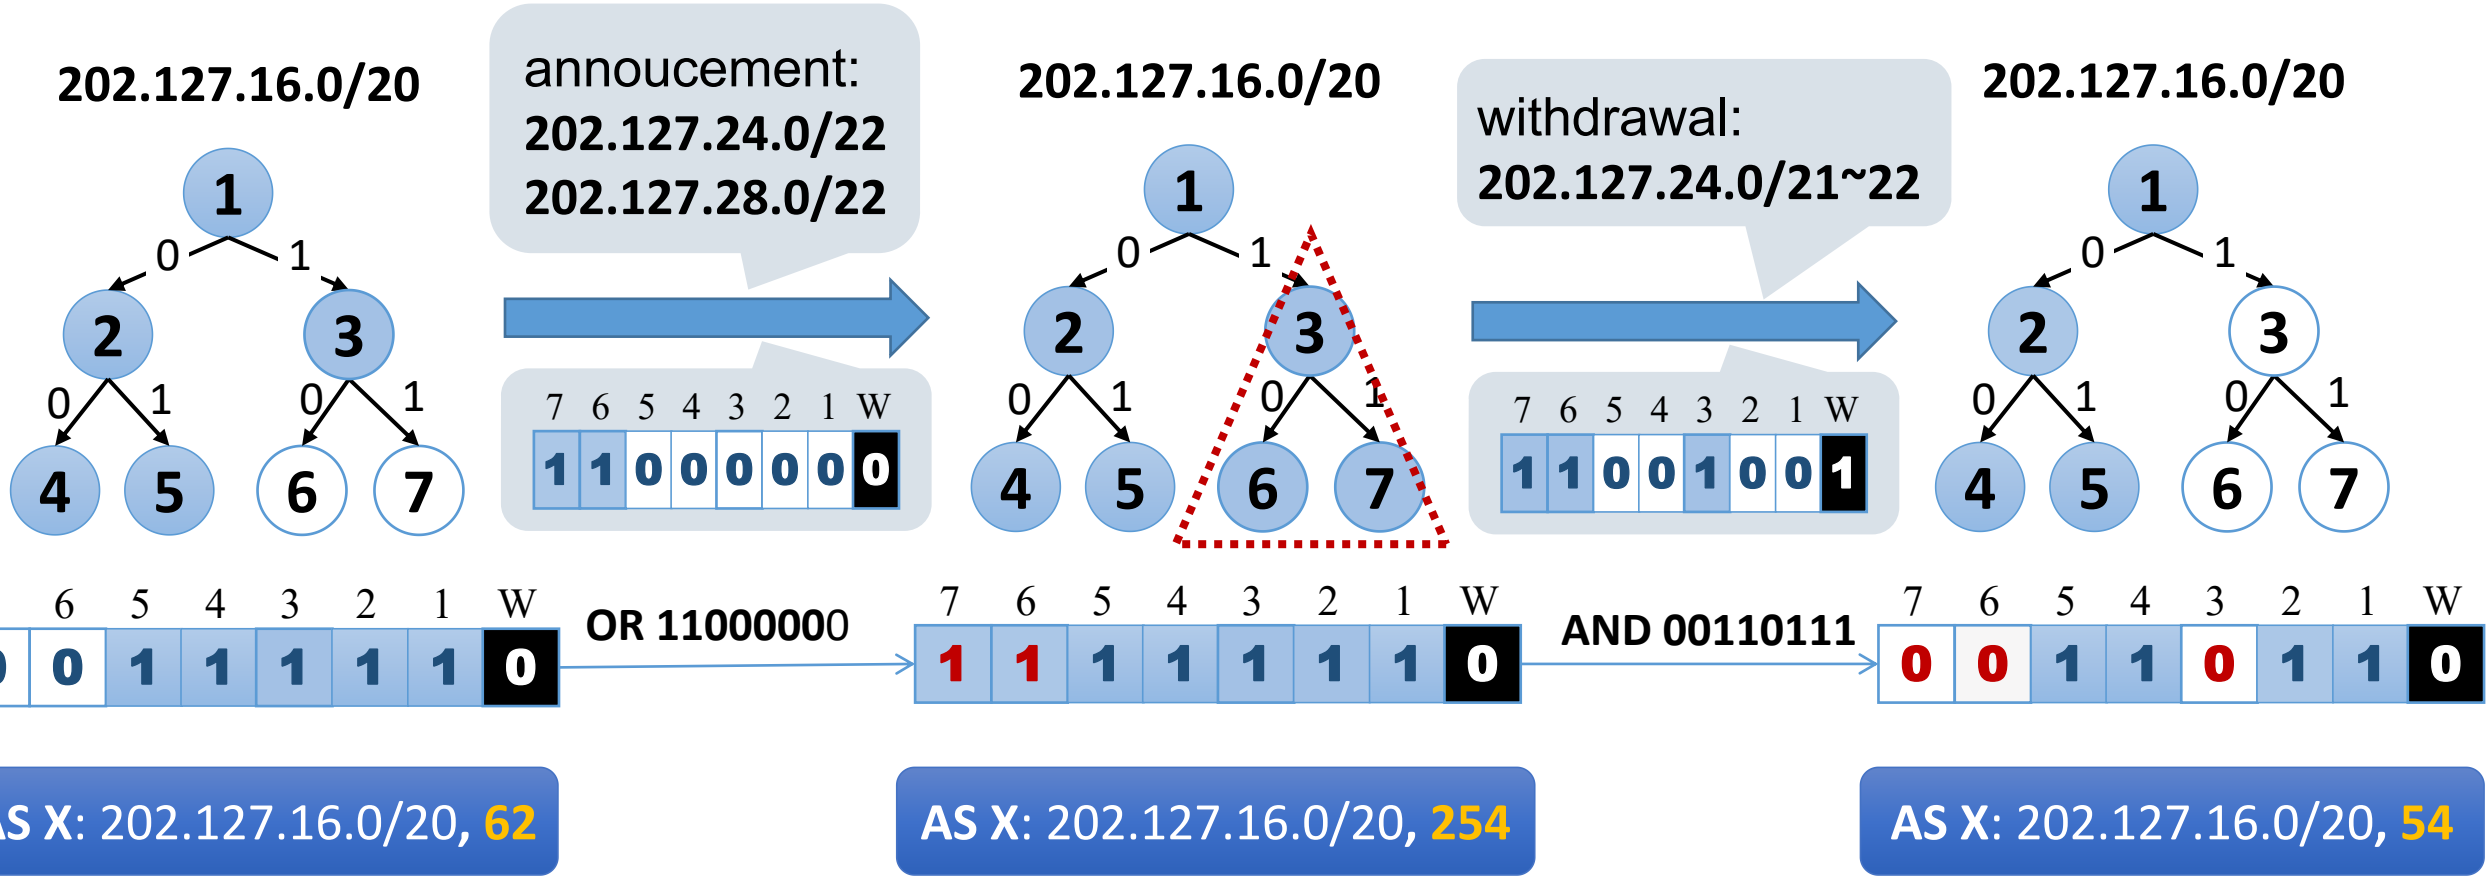

Summary and recommendations

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

Key Observations

Current **corase-grained ROA management** may lead to **security or scalability issues**, which will become more serious with the promotion of RPKI.

Technical Contributions

The **Hanging ROA** uses a **bitmap**-based encoding scheme, which enables **flexible and fine-grained ROA management** with high encoding efficiency.

Recommendations

Remember to issue ROAs for customers after IP address resource allocation, and take fine-grained operations in managing ROAs.

中国科学院计算机网络信息中心
Computer Network Information Center
Chinese Academy of Sciences

中国科学院大学
University of Chinese Academy of Sciences

Thank You For Your Attention

Any questions, please feel free to contact: lybmath@cnic.cn