

# YOUR ROUTER IS MY PROBER:

Measuring IPv6 Networks via  
ICMP Rate Limiting Side Channels

Long Pan, Jiahai Yang, Lin He, Zhiliang Wang, Leyao Nie,  
Guanglei Song, Yaozhong Liu



# I. Background

# IPv6 Measurements

## Importance and Challenges

- IPv6 becomes popular.
  - ~40% users access Google over IPv6.
  - 7.5B active IPv6 addresses observed per day.
- Measuring IPv6 networks faces challenges.
  - Lack of related resource (e.g., open datasets, vantage points)
  - More secure protocol (e.g., no IPID, huge address space, ICMP rate limiting)

<https://www.google.com/intl/en/ipv6/statistics.html>

<https://www.akamai.com/blog/trends/10-years-since-world-ipv6-launch>

# Vantage Points Never too Many!

Active Internet measurements require us to send and receive packets on vantage points.

Many measurements tasks cannot be done without appropriate vantage points.

- For instance, without a VP in a specific network, you cannot know the filtering policy of that network.

Alternatives include in-network volunteers, public probers and looking glasses....

- But you cannot have volunteers or VPs in all networks.

*How about using others' devices as our vantage points?*



## ICMP Rate Limiting

Challenge but also Opportunity

- ICMP (Internet Control Message Protocol) is one of the most common protocols on the Internet (e.g., **Ping**, **Traceroute**).
- To prevent from possible waste of resources and ICMP flood attacks, every IPv6 nodes are **required** by RFC to limit the rate of sending ICMP messages.



Usually, widespread ICMP rate limiting is considered to bring more challenges to IPv6 measurements (especially topology discovery), but the global ICMP rate limiting also opens up new side channels.

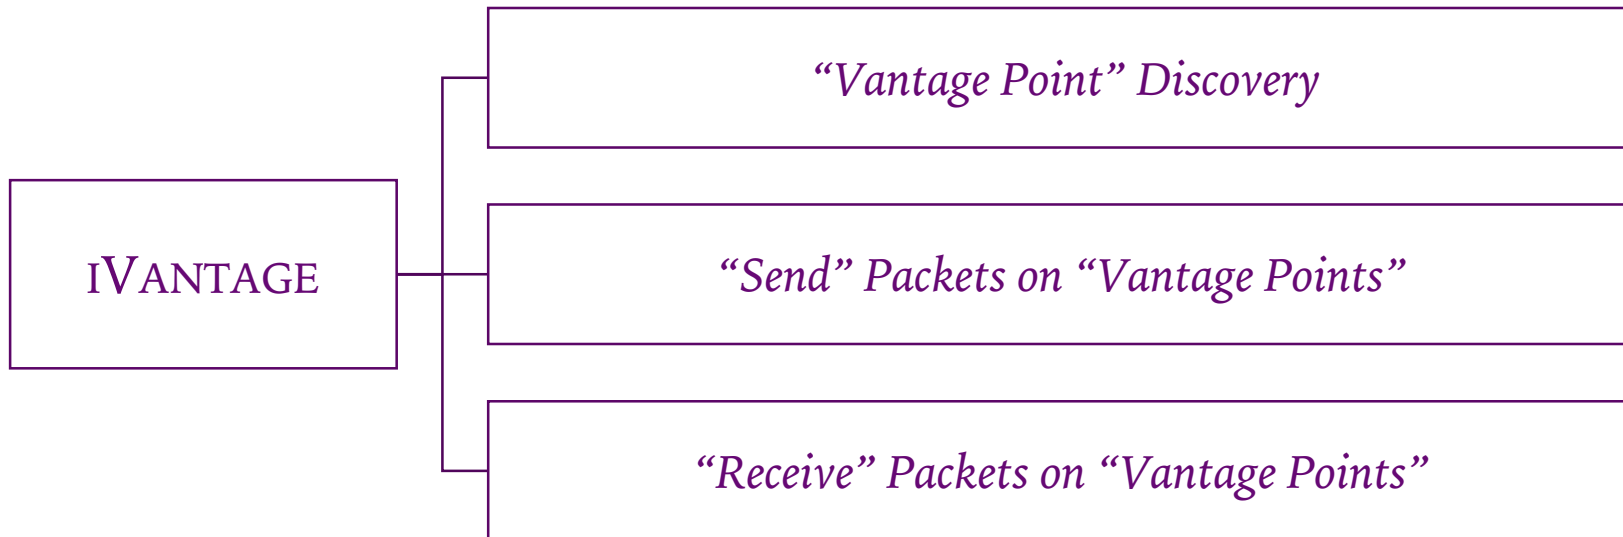
# I. Background

## *ICMP Rate Limiting Side Channel-based Measurements*

We propose a novel measurement technique based on ICMP rate limiting side channels, IVANTAGE.



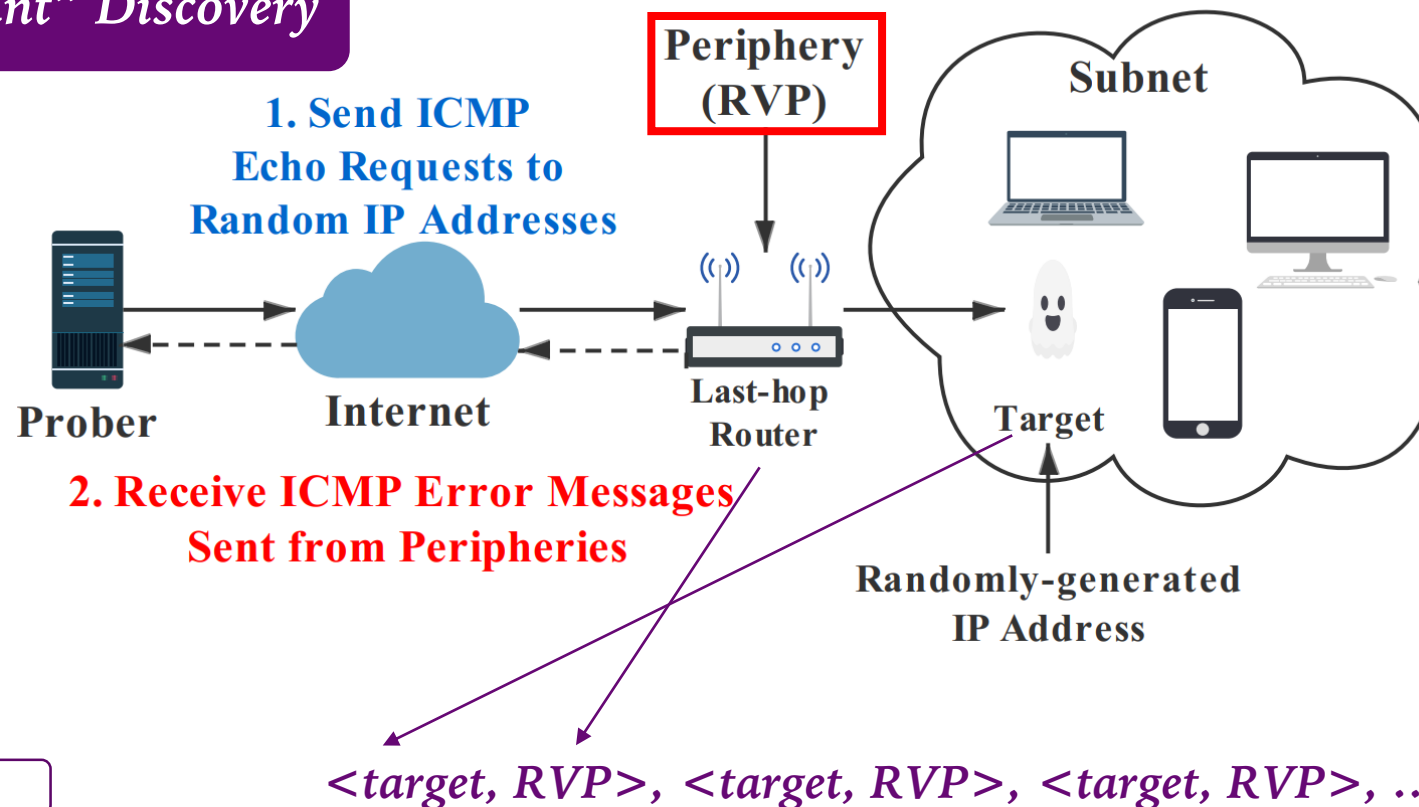
IVANTAGE can to some extent use others' device as our “vantage points”, and then “**send**” and “**receive**” packets on those “vantage points”.



## II. IVANTAGE

## II. iVANTAGE

### 1. “Vantage Point” Discovery



First, we perform active measurements to discover potential **remote “vantage points” (RVPs)**. Though, they are not “vantage points” in the true sense. They are mostly others’ routers.





## II. iVANTAGE

### 1. “Vantage Point” Discovery

*Local Prober*

Sending Probes



Receiving ICMP  
Error Messages

*2001:da8:1234::/48*

*2001:da8:1234:0001:abcd:1493:1264:e2f6:32ed*

*2001:da8:1234:0002:ed42:19c4:23d5:1d4f*

*2001:da8:1234:0003:ac2f:ed99:2443:124a*

...

*2001:da8:1234:ffff:ec4d:4429:64aa:ffae*

65536  
Probes

*Example:*

Sending probes to *2001:da8:1234:0001:abcd:1493:1264:e2f6:32ed*



*target*

Receiving ICMP Destination Unreachable from *2001:da8:1234:212::1*



*RVP*

*Data  
Pairs*

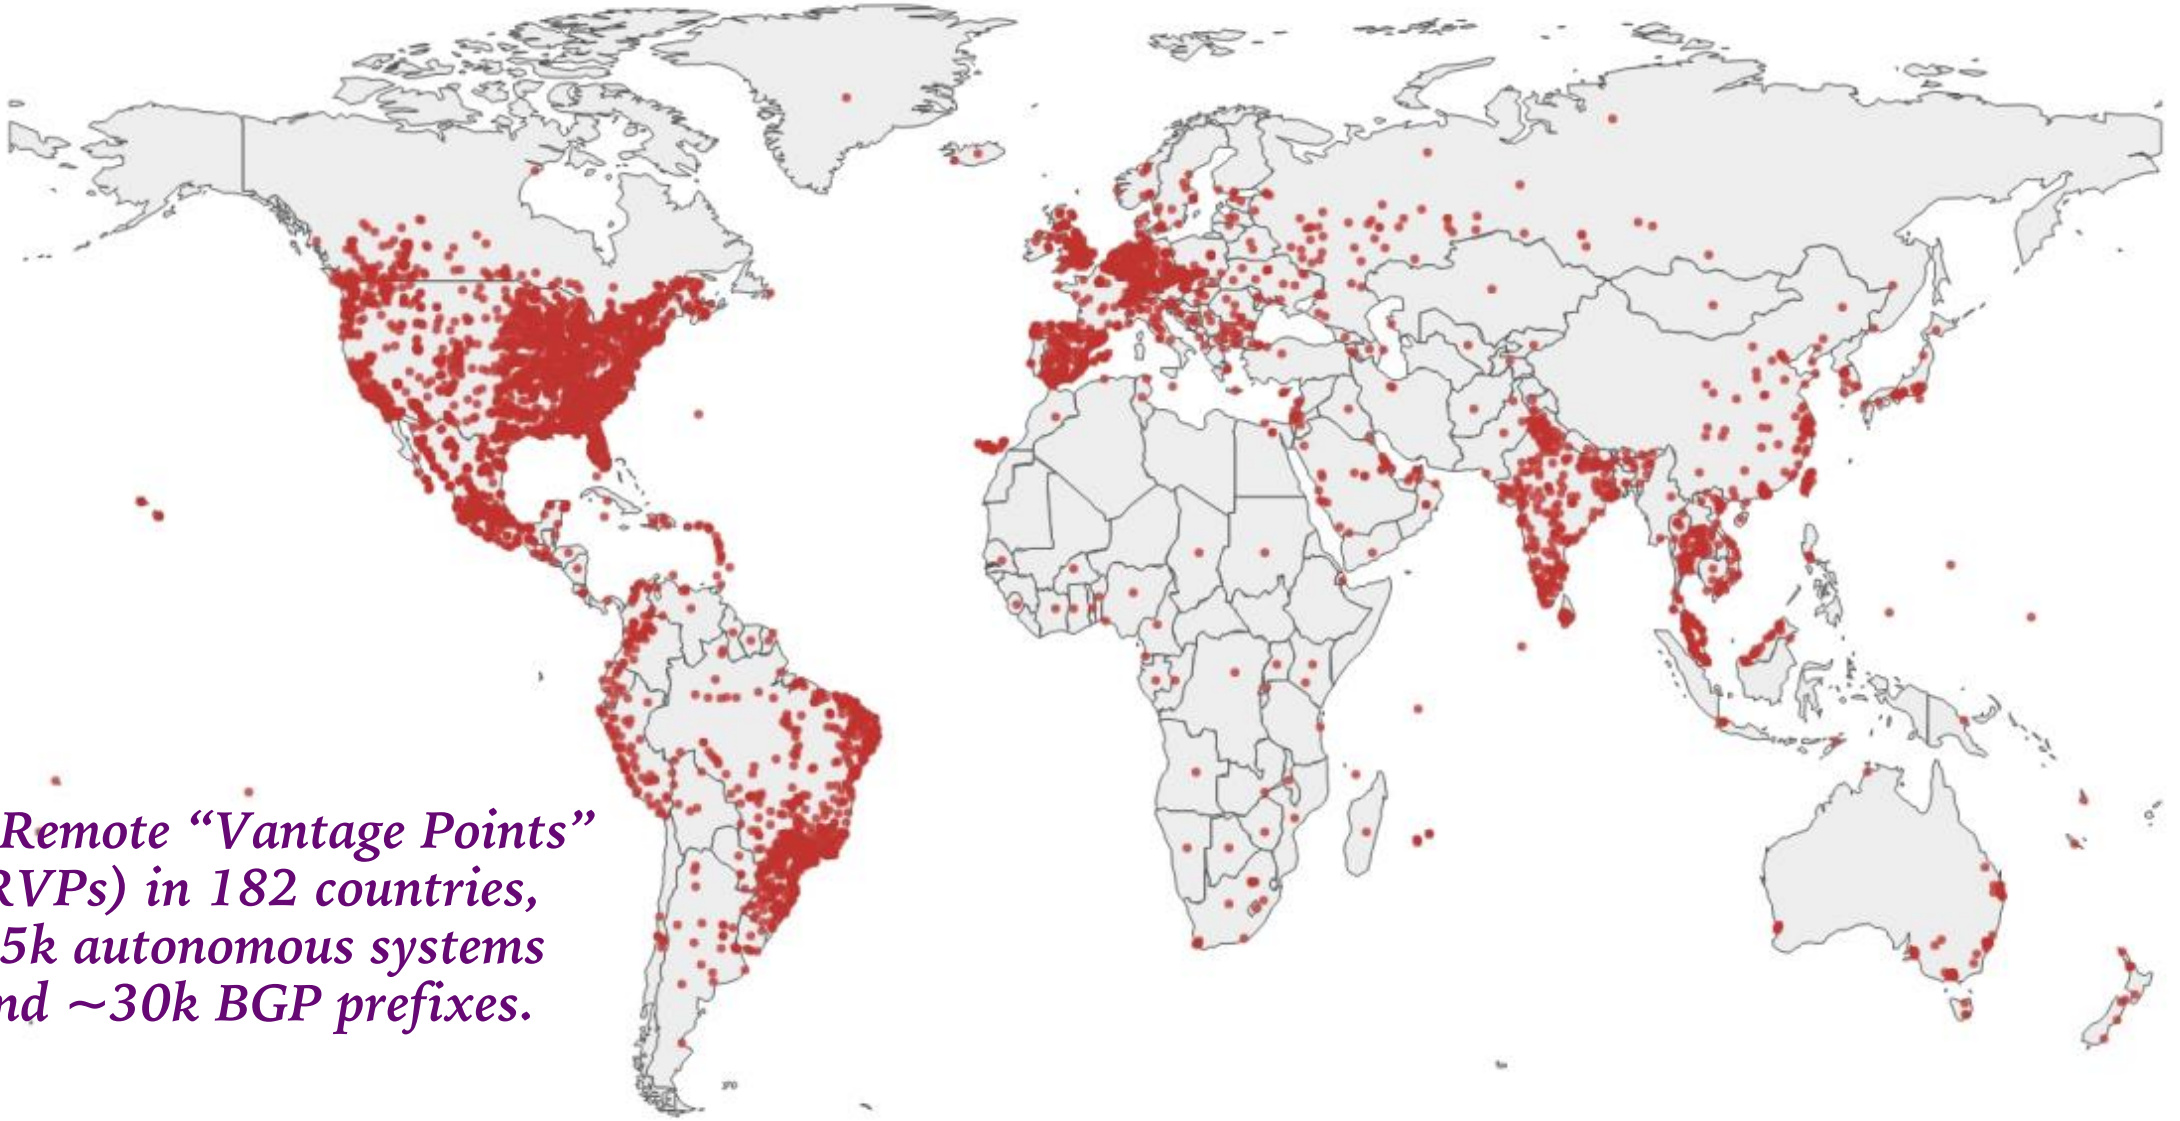
[DSN'21] Xiang Li et al. *Fast IPv6 Network Periphery Discovery and Security Implications*.  
[IMC'22] Robert Beverly et al. *Follow The Scent: Defeating IPv6 Prefix Rotation Privacy*.



清华大学  
Tsinghua University

## II. iVANTAGE

*~1M Remote “Vantage Points”  
(RVPs) in 182 countries,  
9.5k autonomous systems  
and ~30k BGP prefixes.*



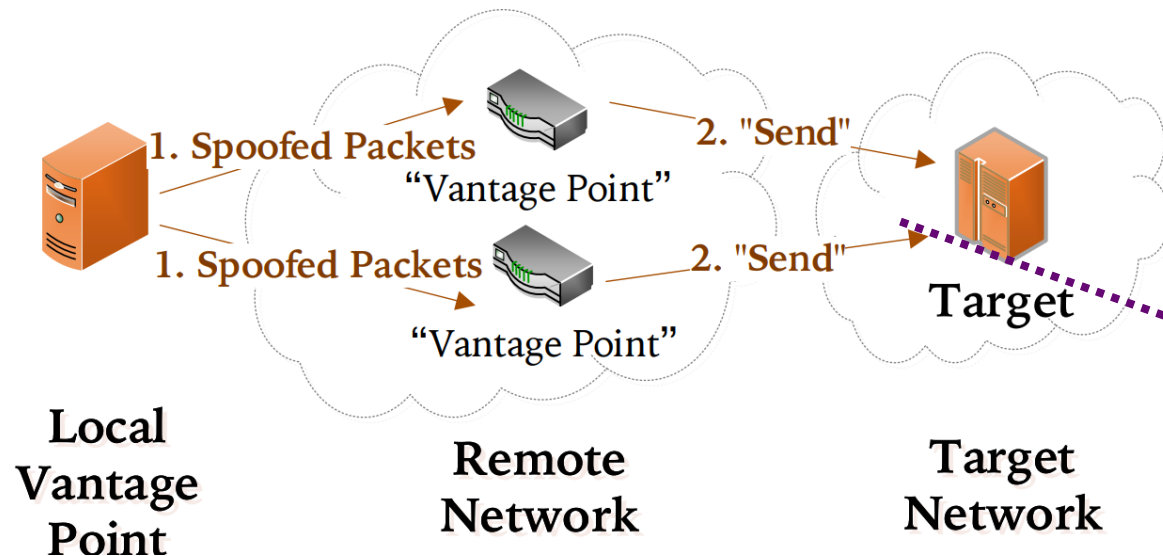
## II. iVANTAGE

### 2. "Send" Packets

We have no control over those "vantage points" (RVPs),  
how can we ask them to send packets as we wish?

An intuitive way is to send packets (that the  
receiver is required to respond) with spoofed  
source addresses to the RVPs.

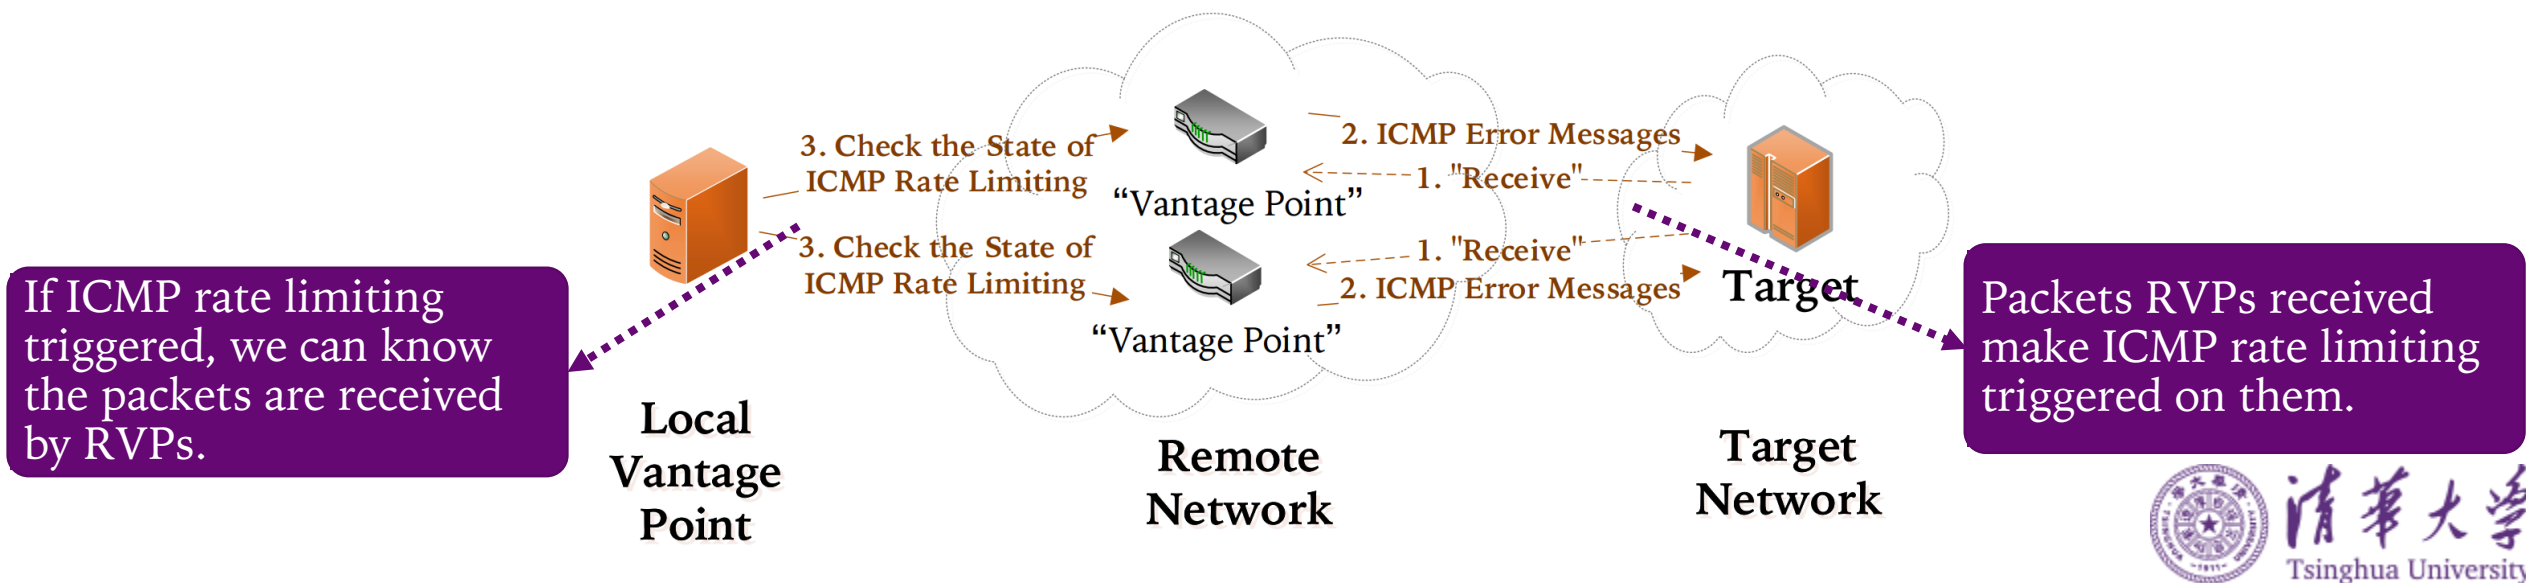
For example, ping, TCP-  
SYN, DNS Queries...



Just as if we had required the  
RVPs to send packets to the  
specified targets.

### 3. “Receive” Packets

- Receiving packets on the RVPs without controlling them seems to be even more difficult!
- However, assume that if the packets the RVPs receive will trigger ICMP rate limiting on them, then *we can know whether the packets are received by observing ICMP rate limiting on the RVPs.*



### *Measurement Tasks*

- We apply iVANTAGE into several different measurement tasks.

Measurement Task	Challenges	Our Approach
1. Deployment of Inbound Source Address Validation	Without a VP in the target network, we cannot know the spoofed packets we send are filtered or not.	“Receiving” Packets on “Vantage Points”
2. IP Reachability between two IPv6 Nodes	Without controlling either of the two nodes, we cannot know whether it can ping the other node.	“Sending” and “Receiving” Packets on “Vantage Points”
3. Discovering Hidden Machines	Without a VP in the target networks, the hidden machines cannot be discovered because it doesn't respond to probes sent from networks other than its own network.	“Sending” and “Receiving” Packets on “Vantage Points”
...	...	...

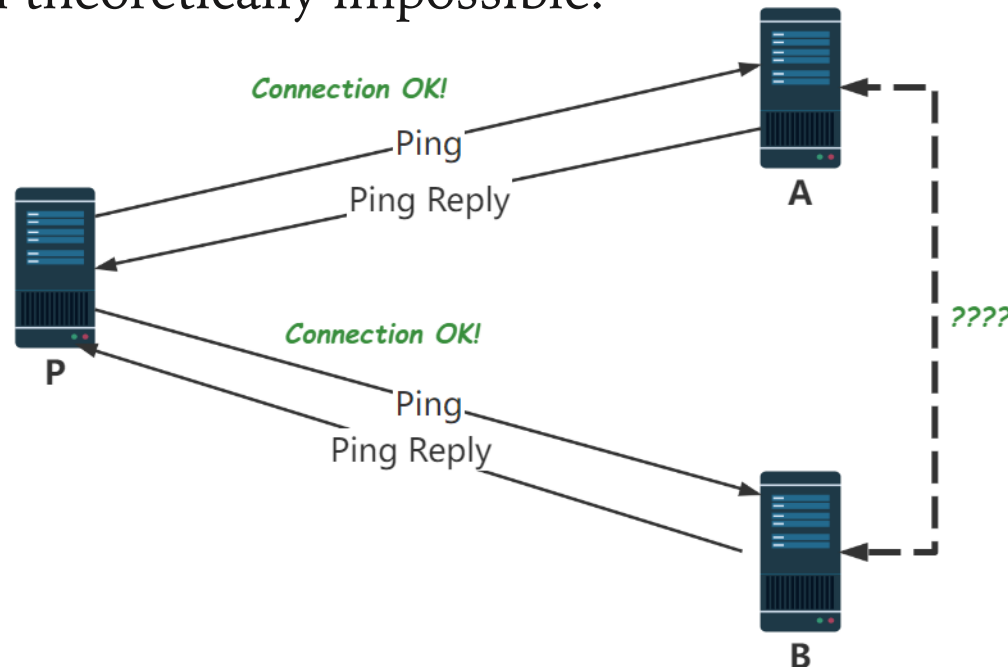
### III. Measurement Applications of IVANTAGE

*Reachability Measurements as Example*

#### Loss of Reachability

##### *Are Internet nodes always inter-connected?*

- Many causes including but not limited to link failures, routing failures and Internet censorship may lead to loss of reachability.
- However, measuring reachability between two nodes *without controlling any of them* can be really hard and even theoretically impossible.





#### *Related Work on Measuring Reachability (Censorship)*

Related Work	Based on	Limitations
[Security'17][ATC'16]	DNS	Rely on DNS. However, DNS connectivity may not reflect IP connectivity.
[NDSS'20][Security'18]	Echo Servers	Rely on echo servers. Echo servers cannot cover all networks. Few IPv6 echo servers, and it's also very difficult to discover echo servers in IPv6 address space.
[S&P'20]	VPN	Deploying or buying VPN service can cost a lot. It remains difficult to have VPNs in all networks.
[S&P'17]	IPID Side Channels	No IPID field in IPv6 fixed header.

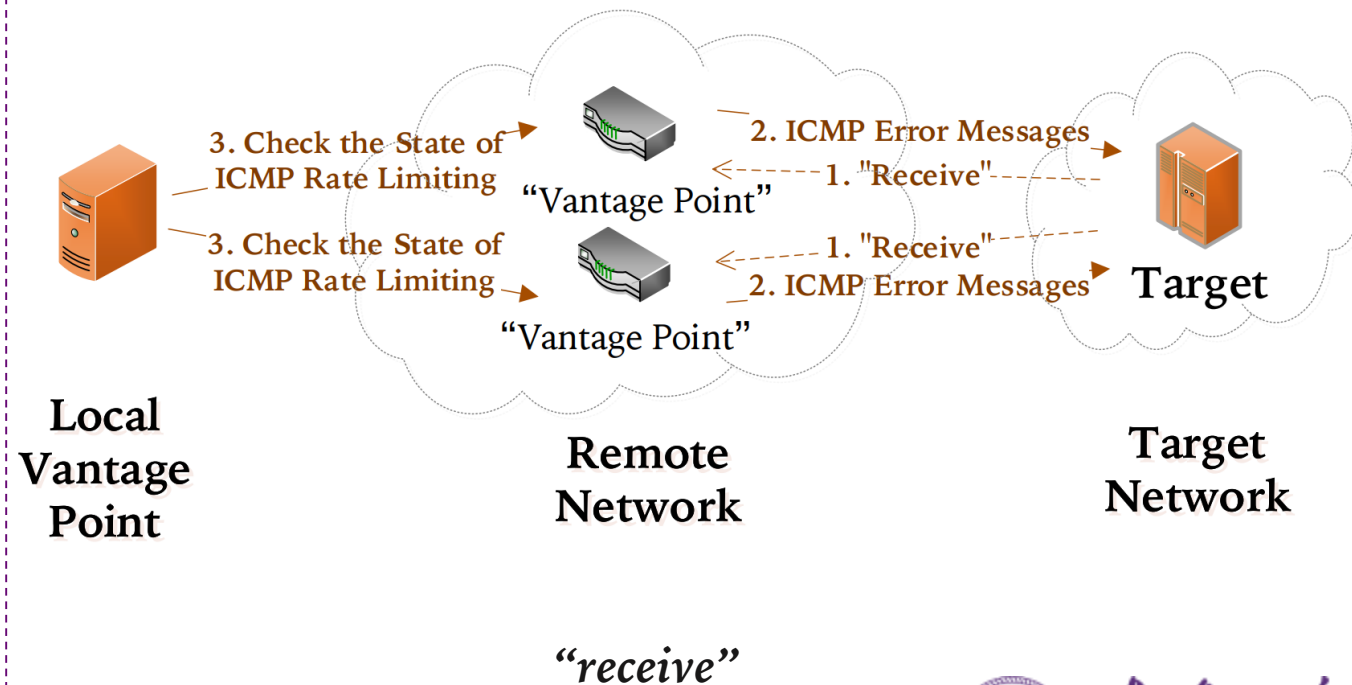
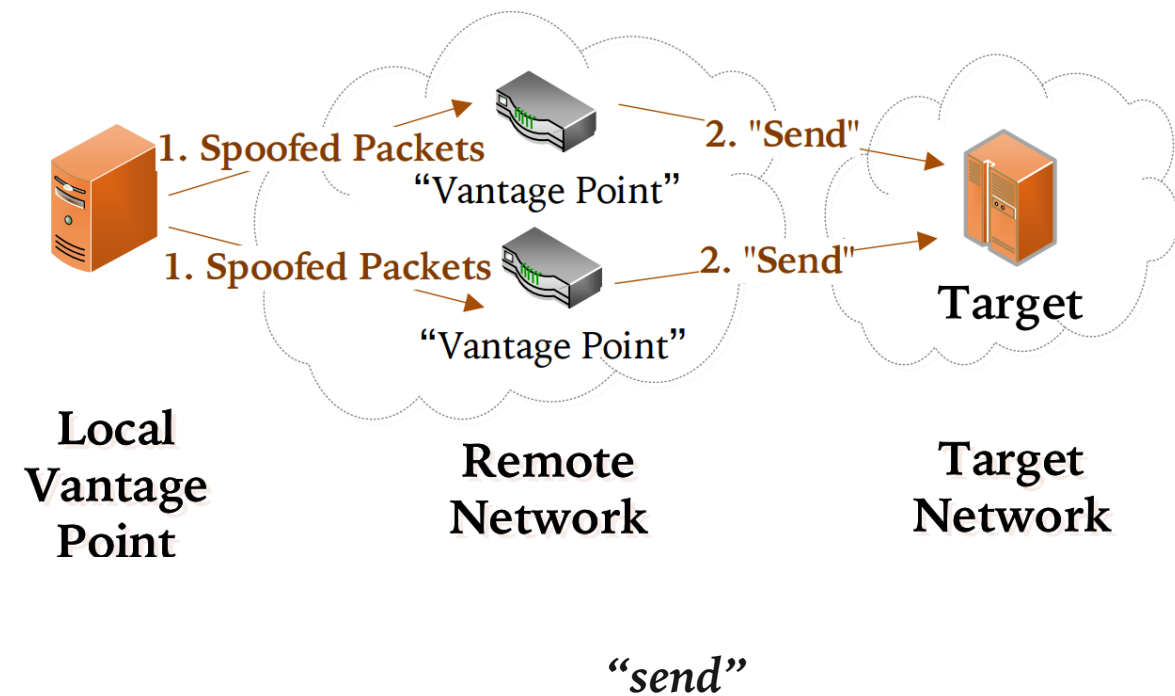
Our work focuses on network-level reachability itself instead of Internet censorship.



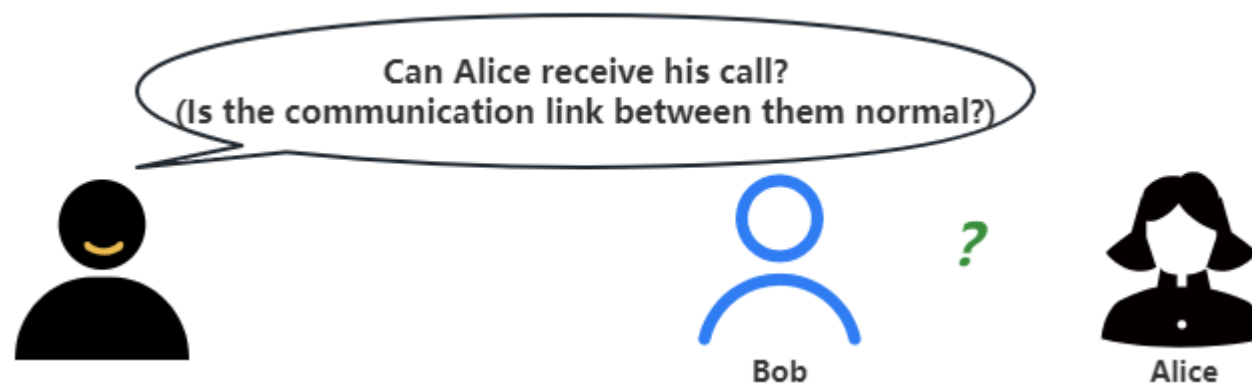
### III. Measurement Applications of IVANTAGE

#### Measuring Reachability based on IVANTAGE


- Measuring reachability between two Internet nodes without controlling any of them relies on the ability to both “send” and “receive” packets on our remote “vantage points”.



### III. Measurement Applications of IVANTAGE

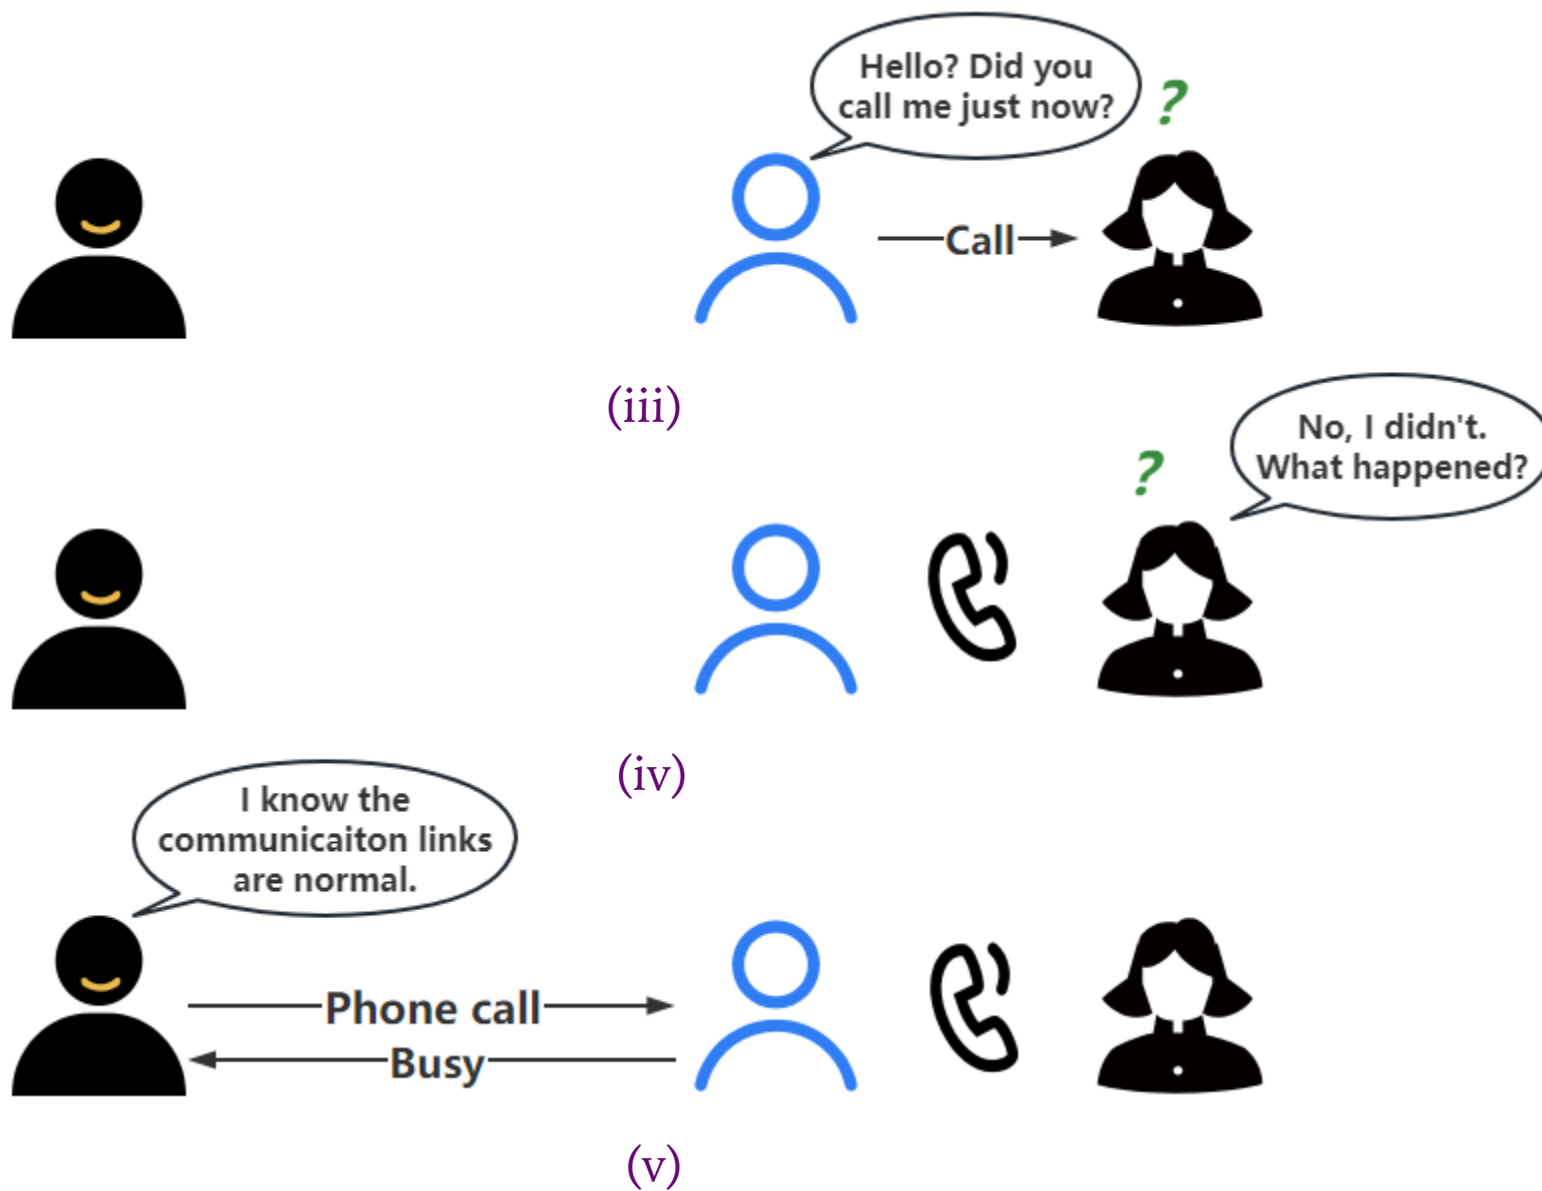


(i)

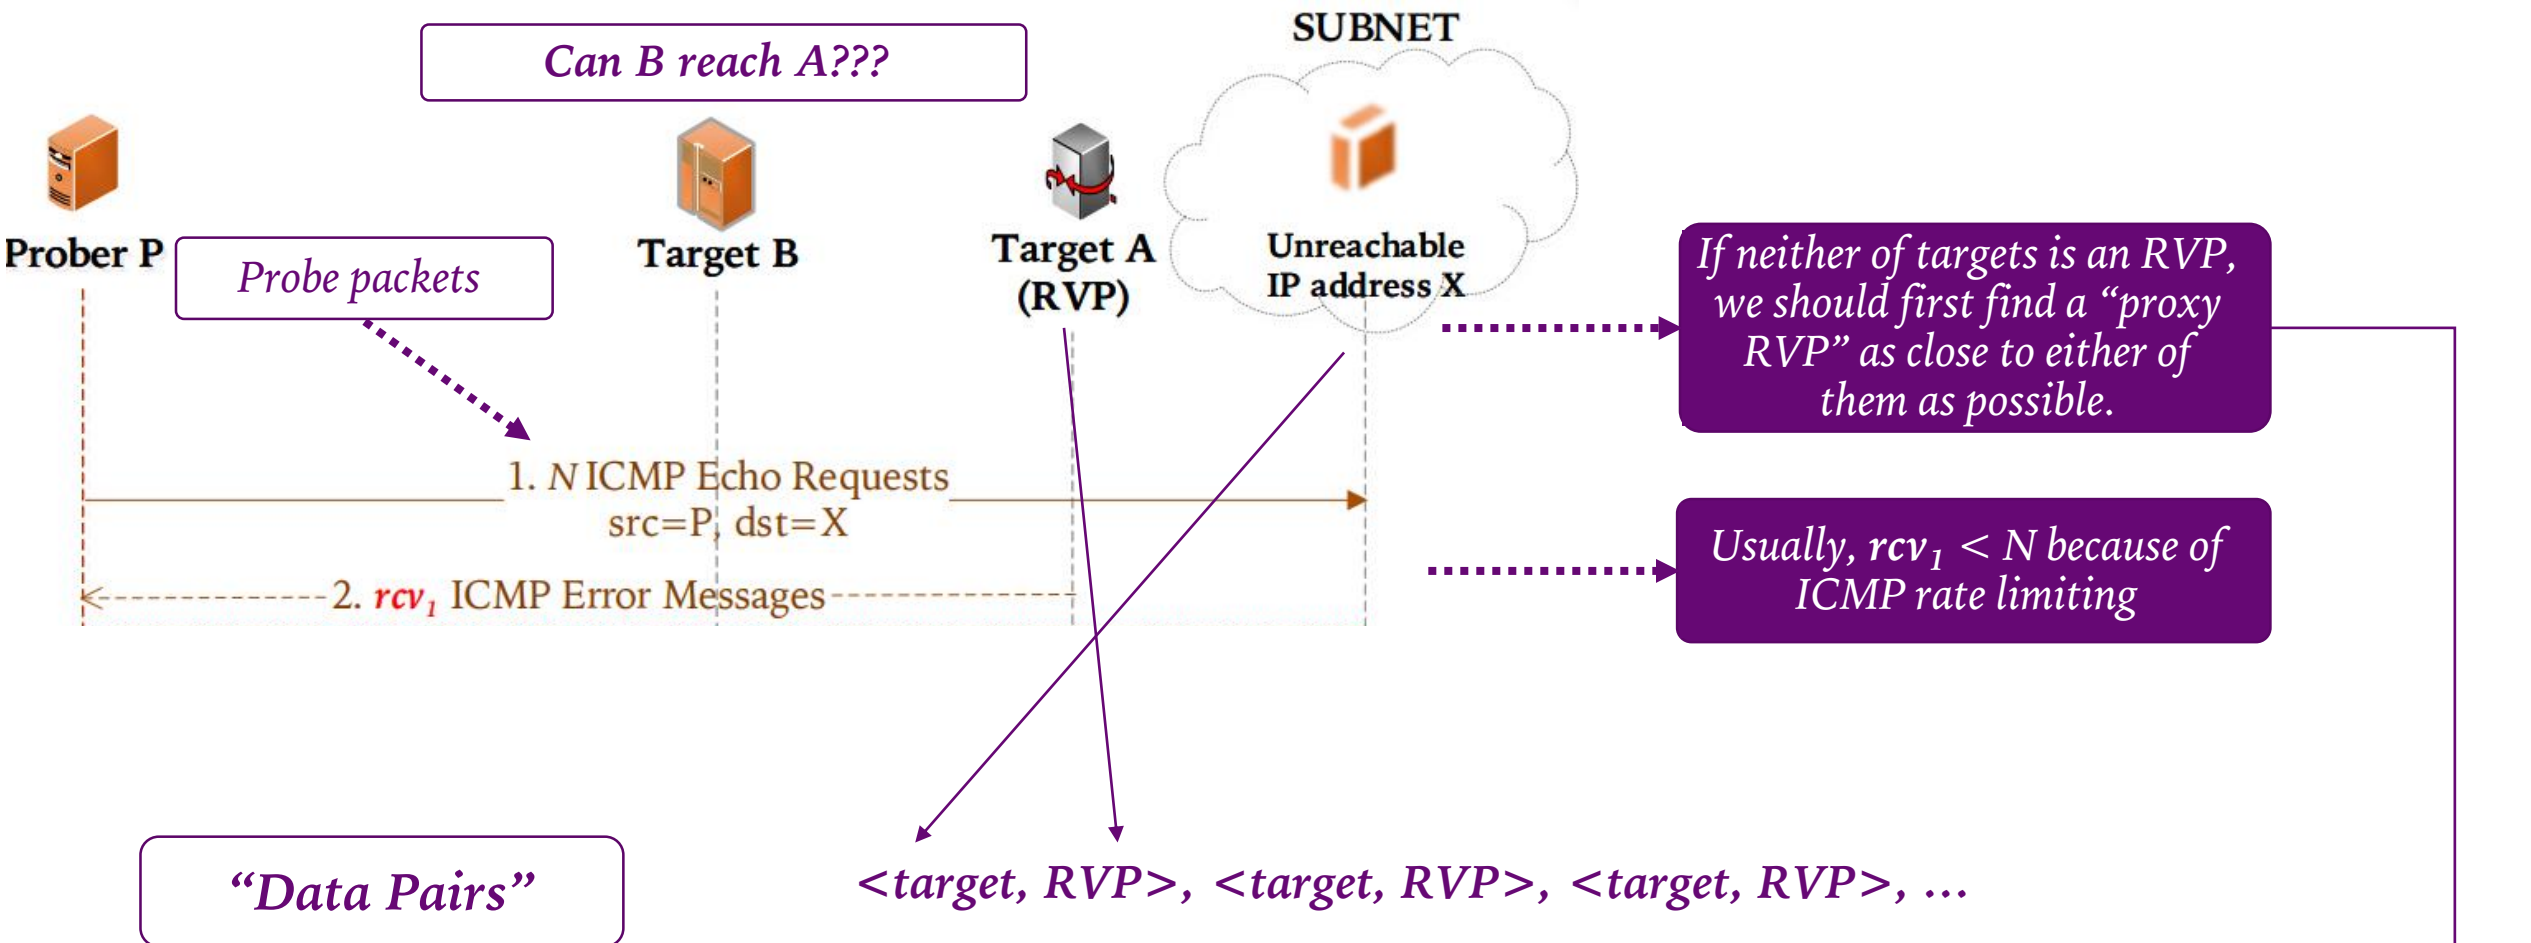


(ii)

### III. Measurement Applications of IVANTAGE

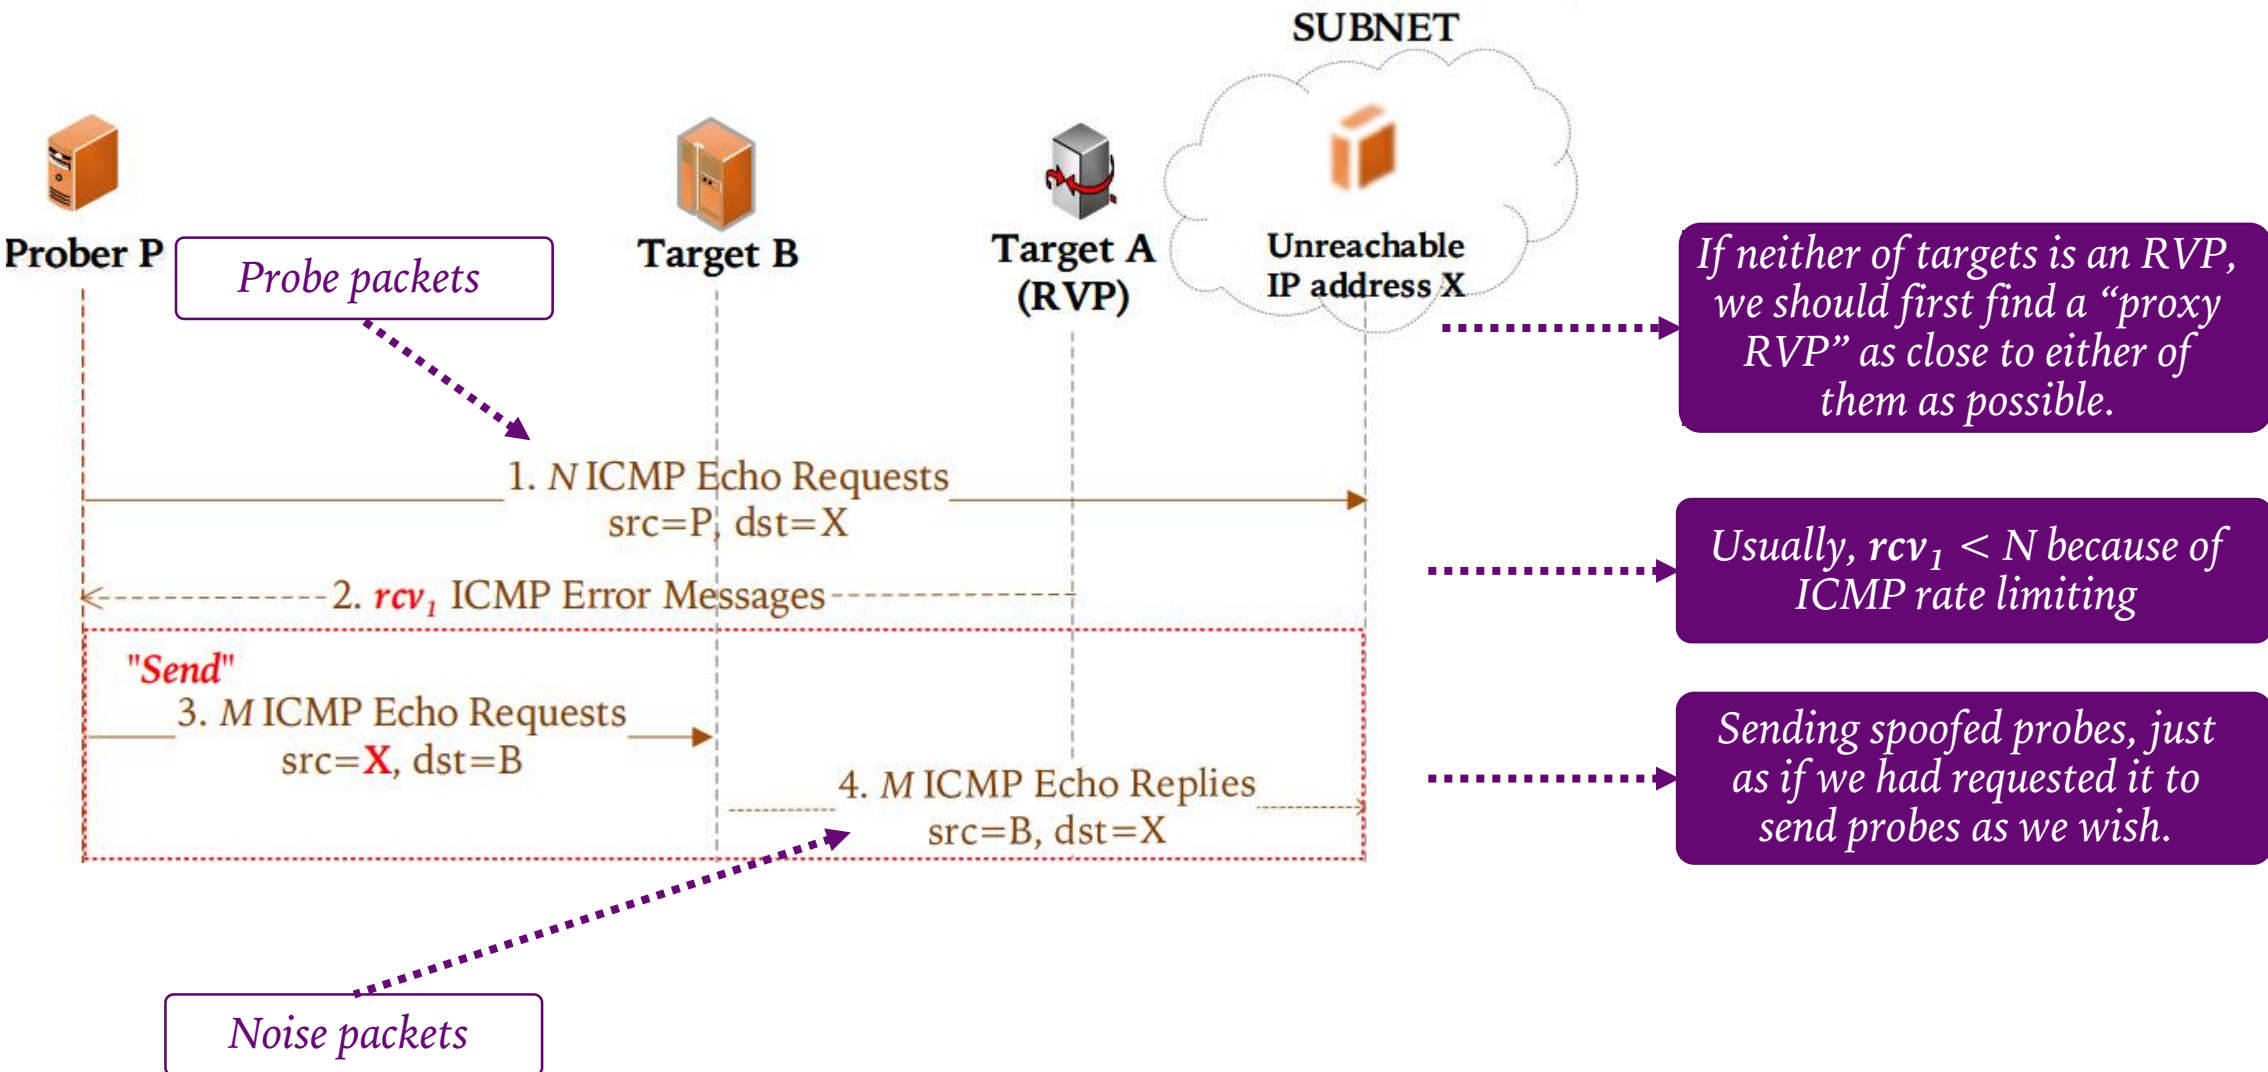


### III. Measurement Applications of IVANTAGE

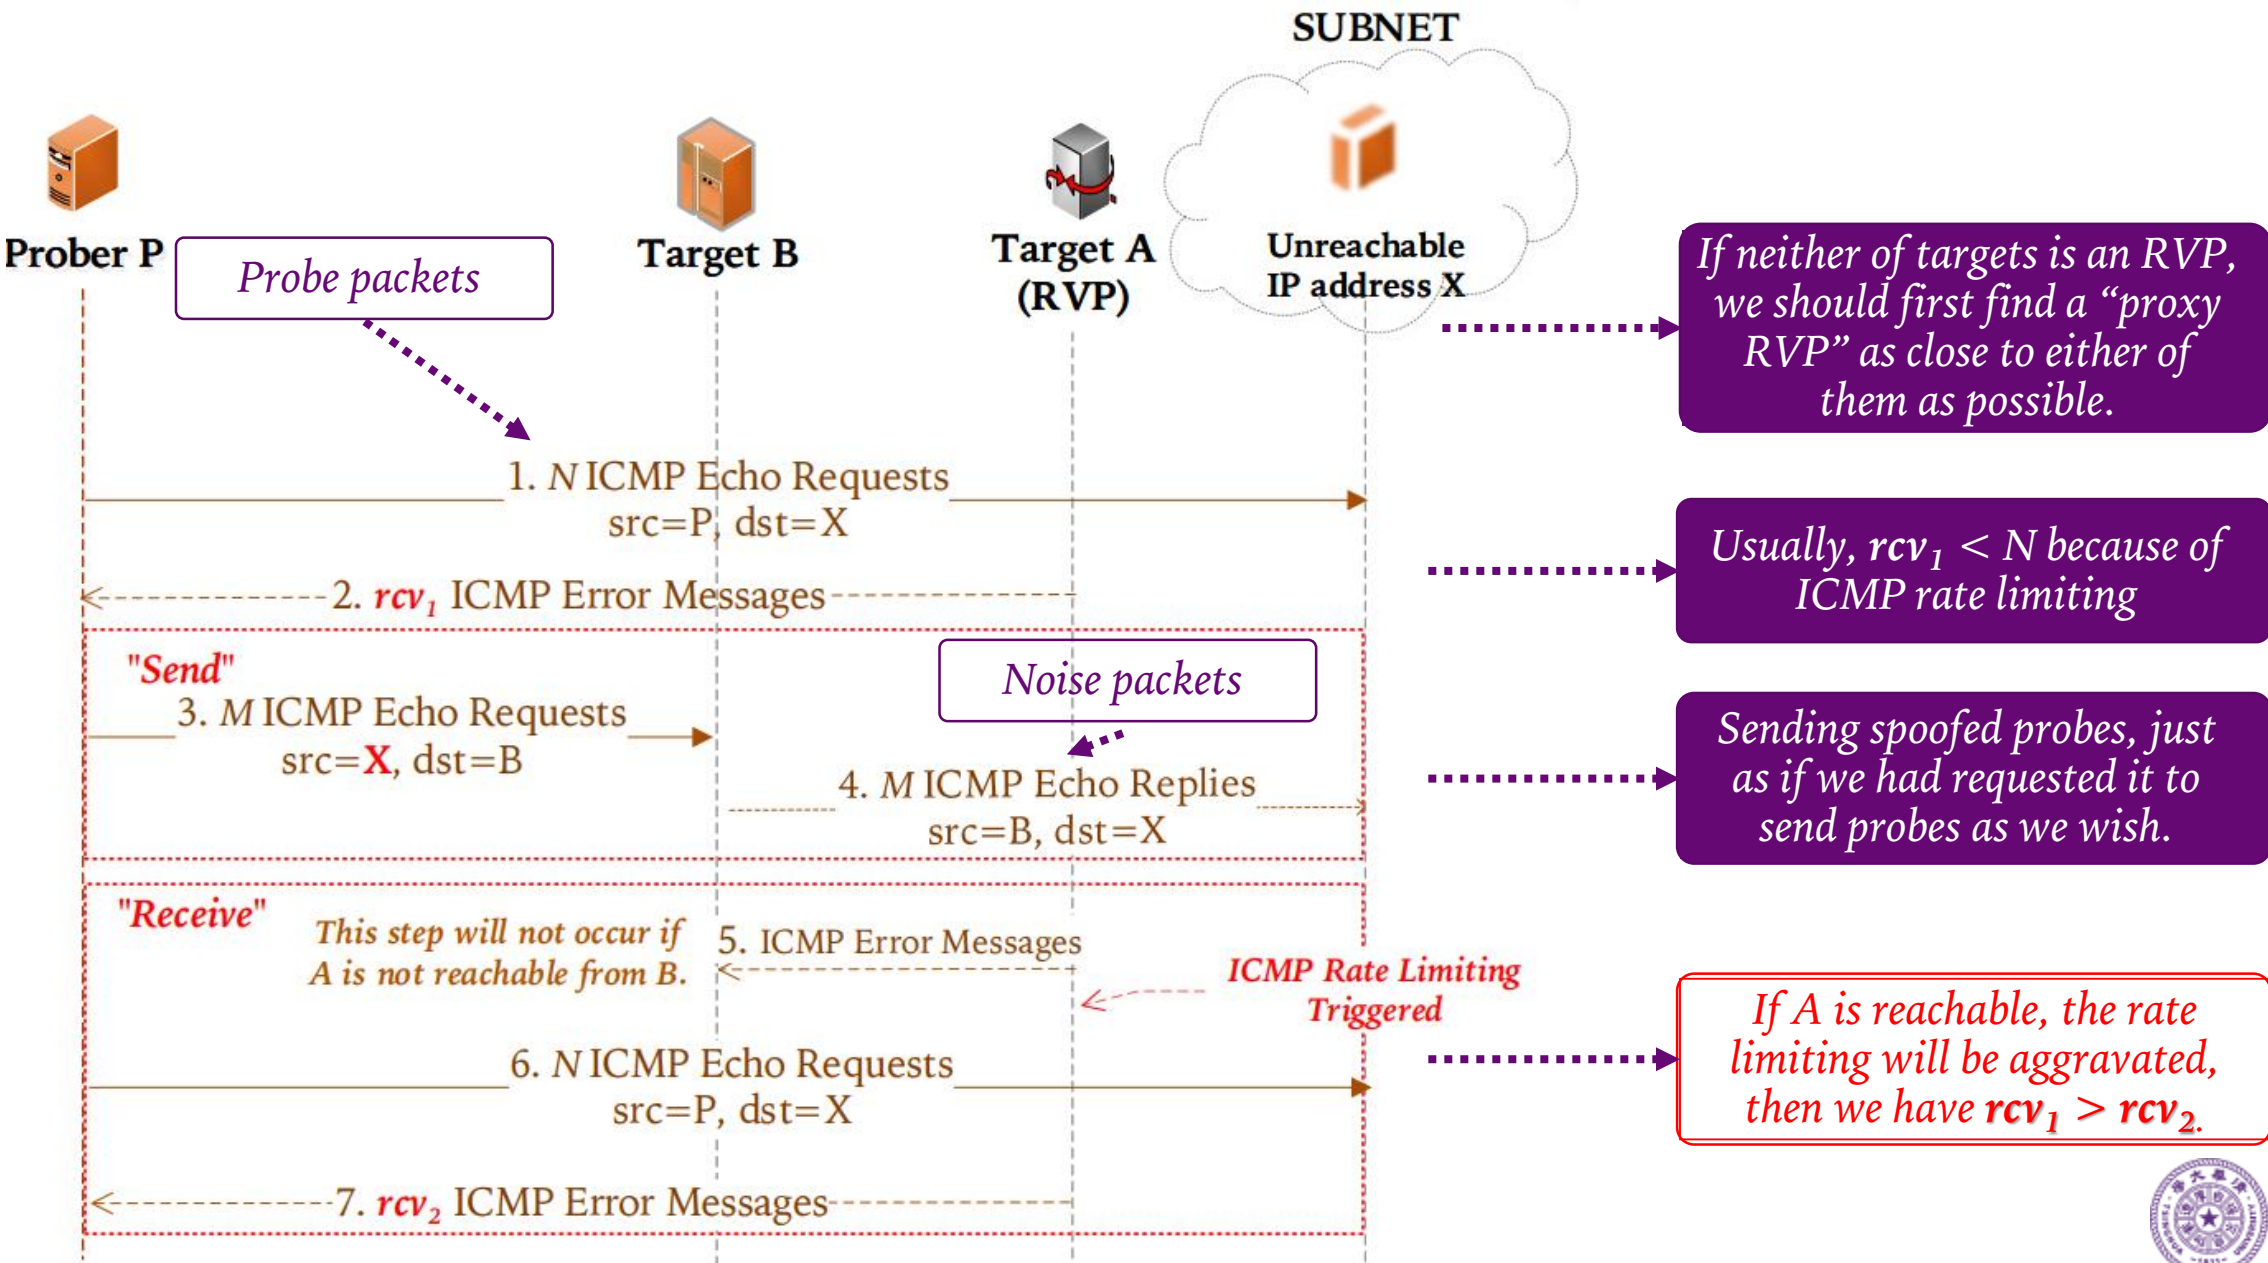


Loss of reachability is unlikely to occur between two very close Internet nodes.

### III. Measurement Applications of IVANTAGE



### III. Measurement Applications of IVANTAGE

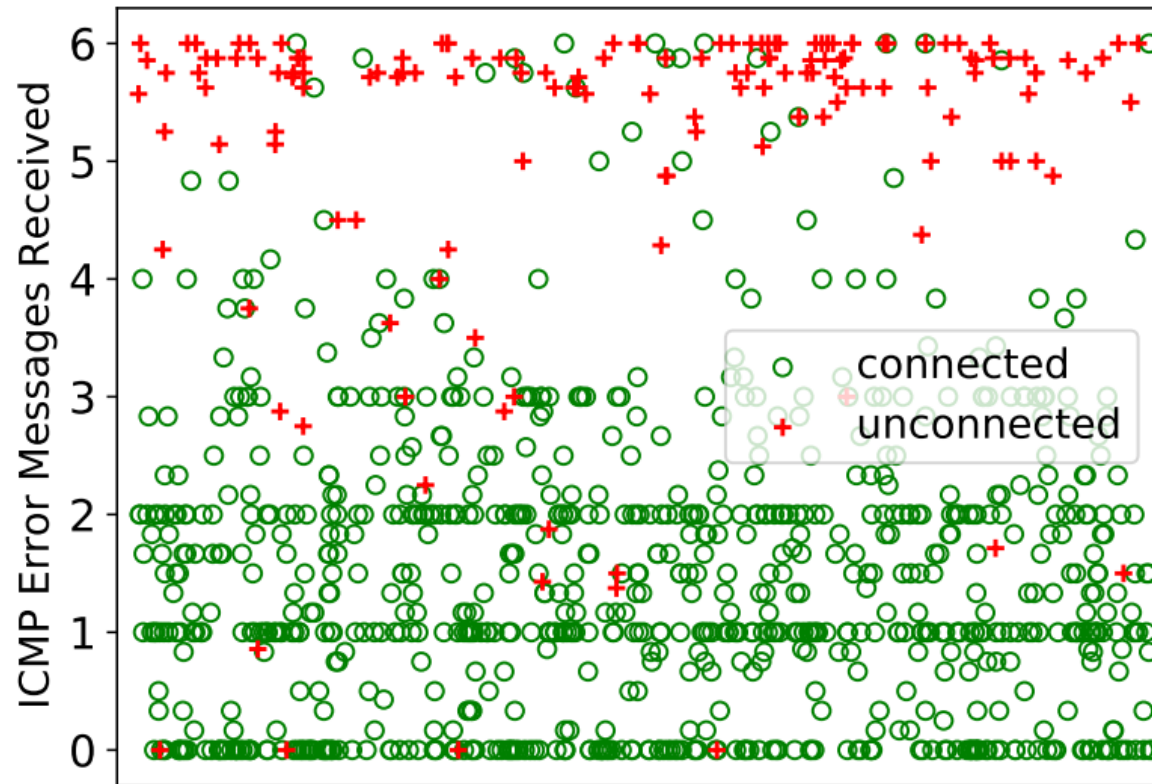





### III. Measurement Applications of IVANTAGE

#### Results

( $N=50$ ,  $M=100$ , for  $k$ -time measurements)

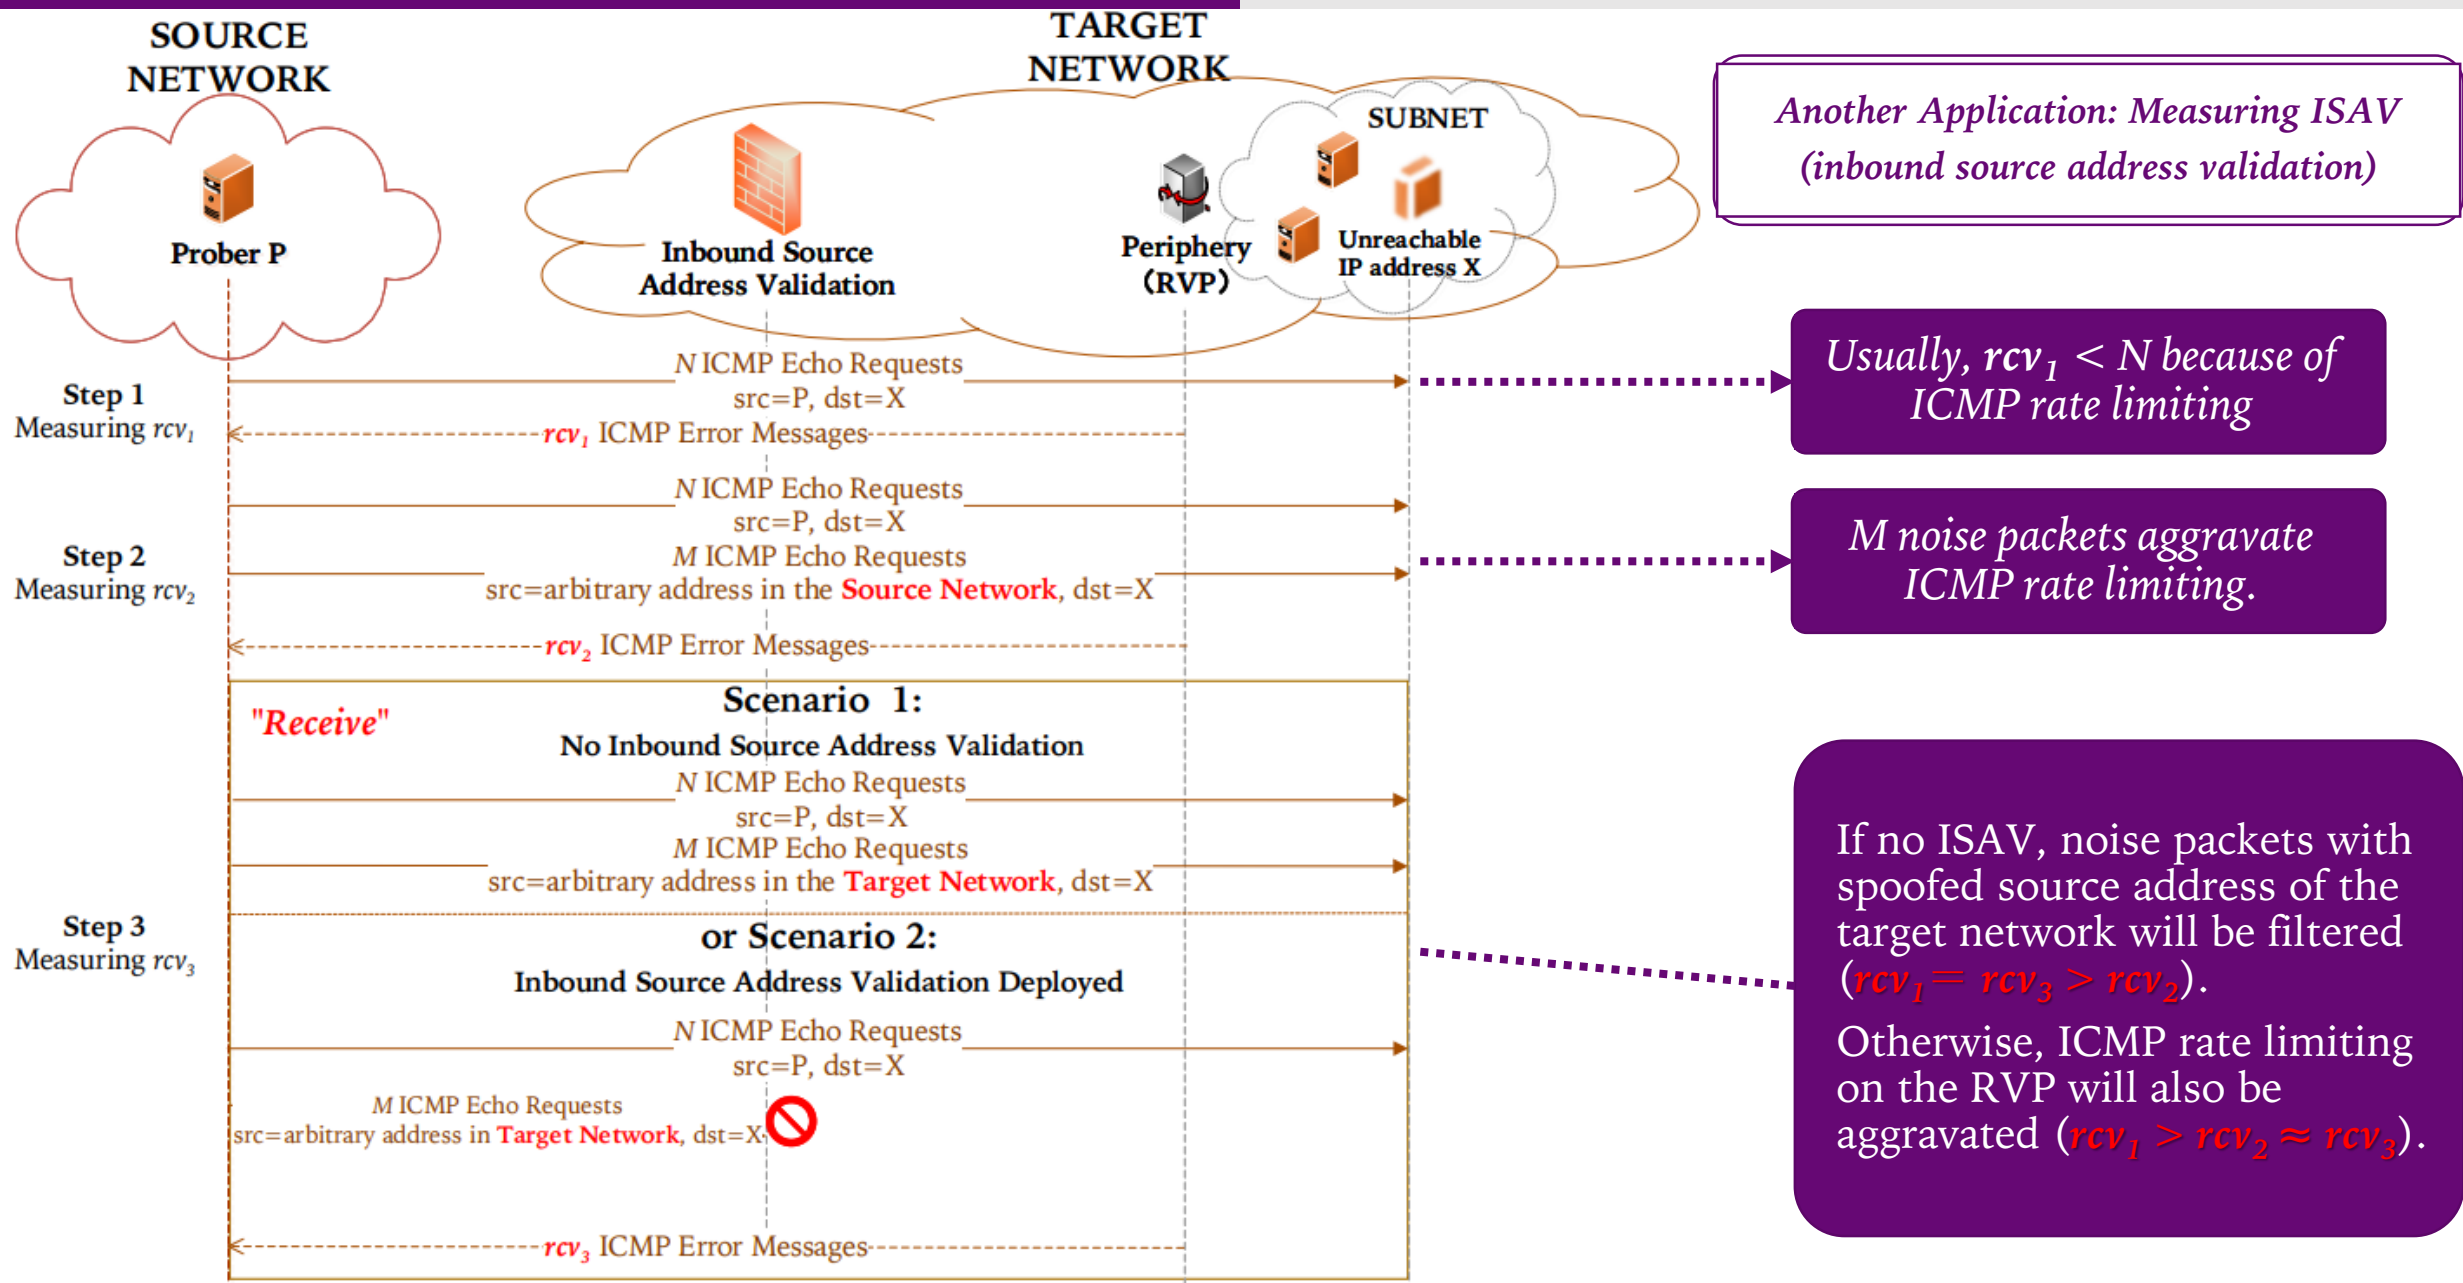


Unconnected nodes result in more error message received. (5.13 vs. 1.48 on average)



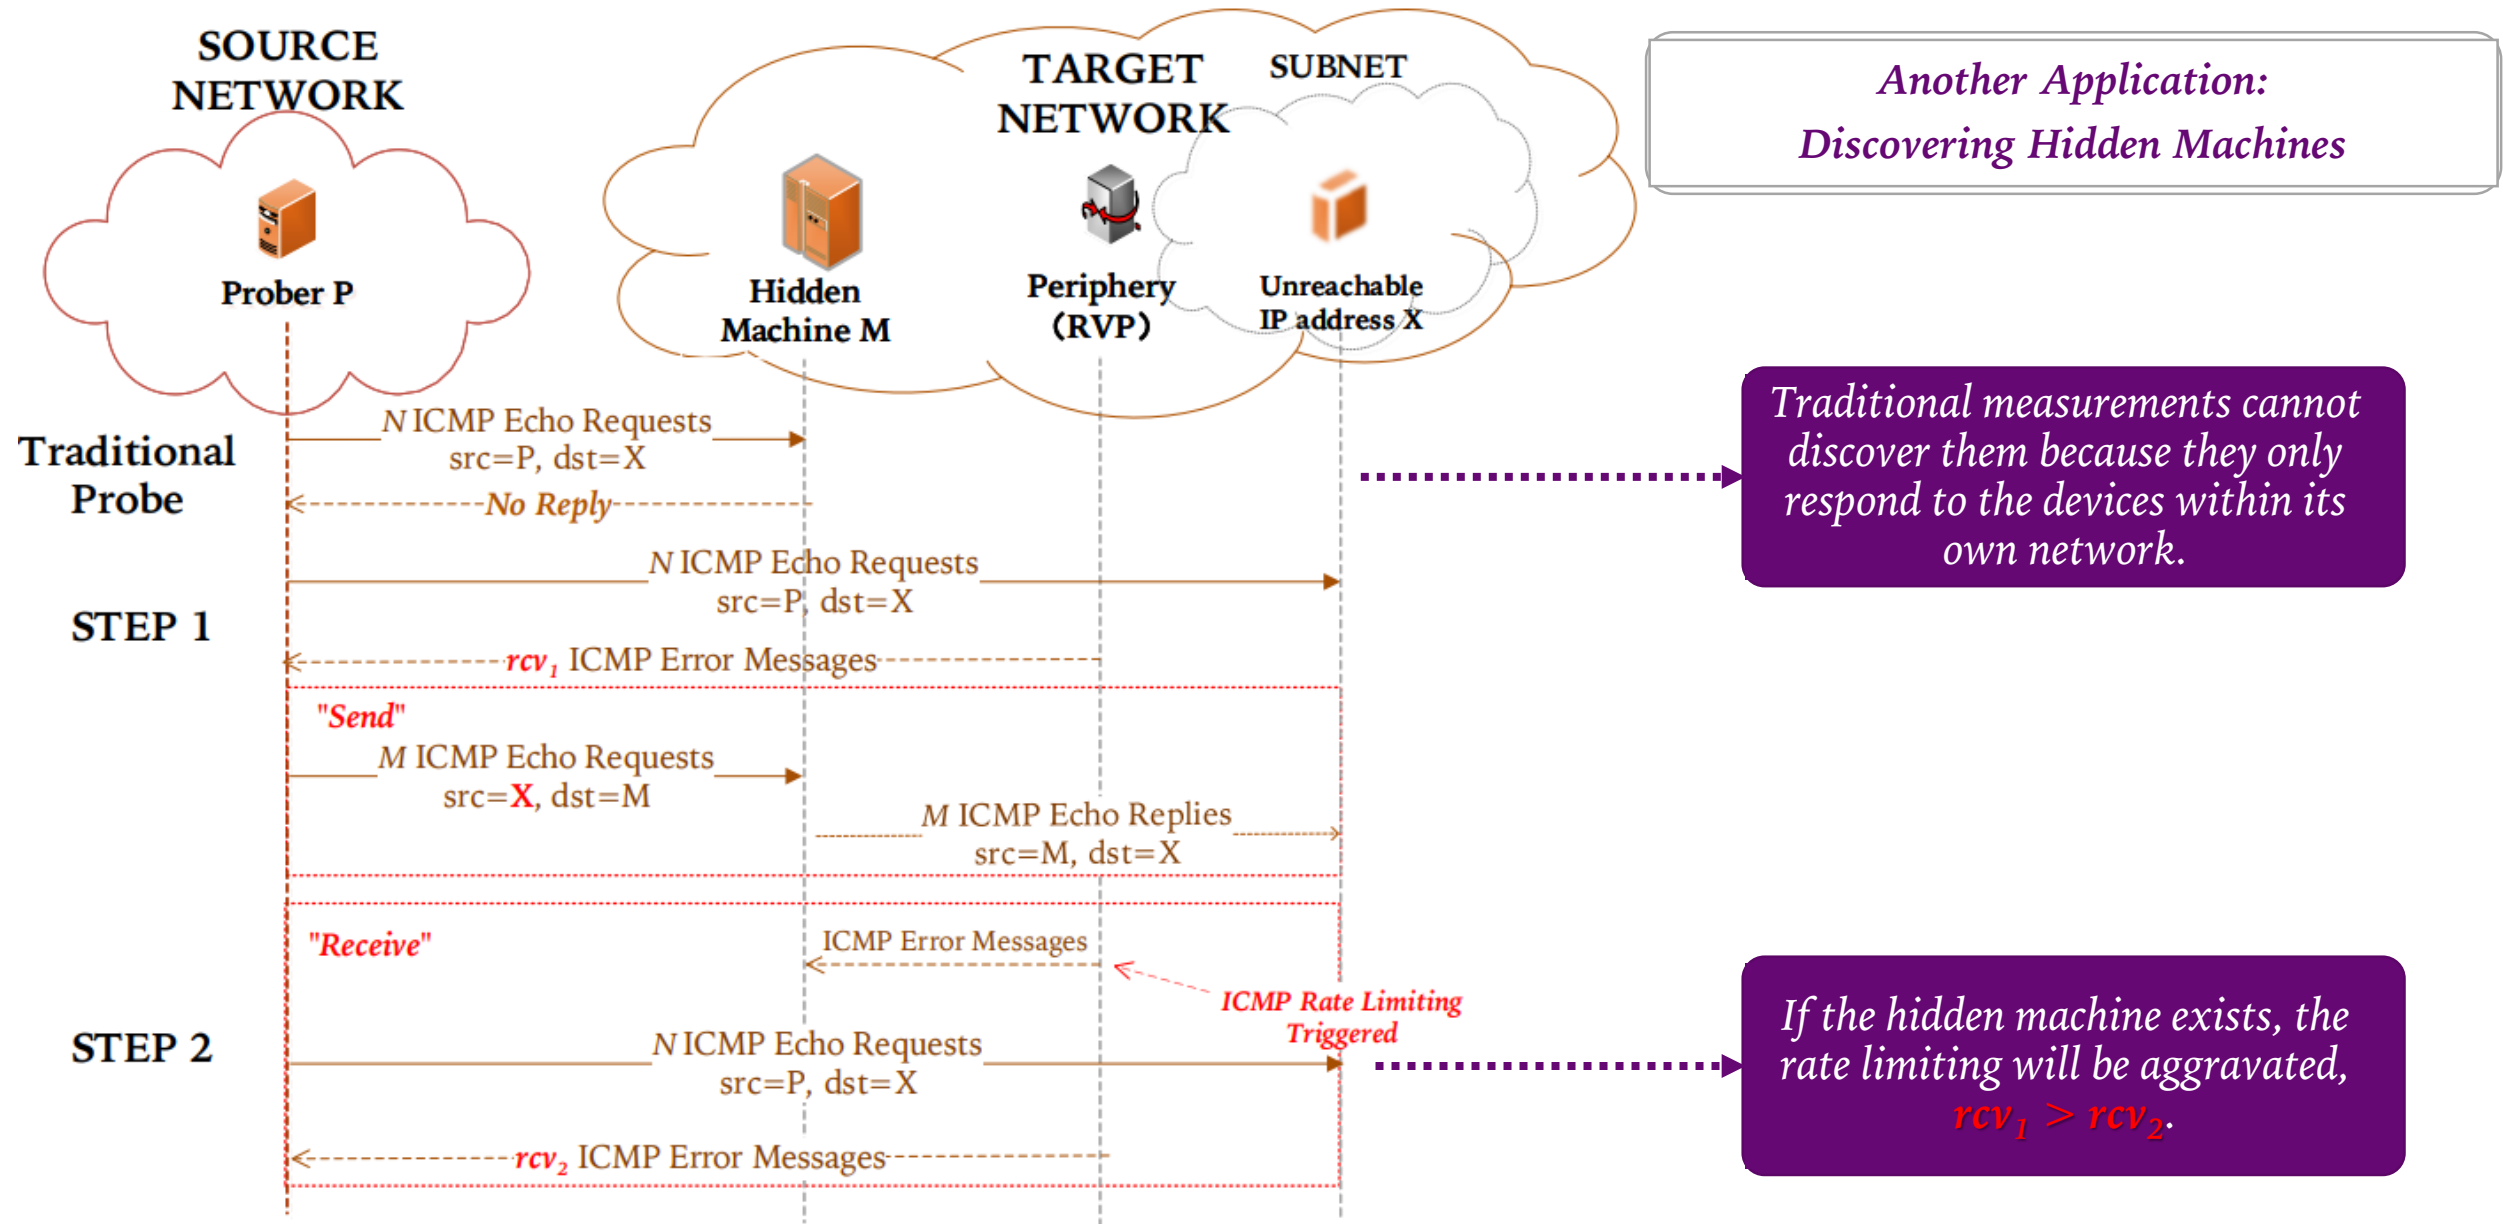
~80% precision and recall for a 5-time measurements.

### III. Measurement Applications of IVANTAGE





### III. Measurement Applications of IVANTAGE



## IV. Measuring “Rate Limiting”

## IV. Measuring “Rate Limiting”

### *How do IPv6 Nodes Implement ICMP Rate Limiting?*


There may be different kinds of ICMP rate limiting implementations:

- ***Global Rate Limiting***: Limiting the rate of originating ICMP error messages sent to all IP addresses, even if triggered by only one IP address.
- ***Strict Rate Limiting***: Limiting the rate very strictly. For example, only reply with 1 packets no matter what they received)
- ***Loose Rate Limiting***: Very loose (or even no) rate limiting. For example, 500 replies for 500 probes.

We perform Internet-scale active measurement covering ~9k IPv6 autonomous systems.

## IV. Measuring “Rate Limiting”

### *A Measurement Study on ICMP Rate Limiting Behaviors*



*$rcv_2$ : Error messages received after sending additional noise packets (with different src addr) to try to aggravate rate limiting.*

ICMP Type	Global ( $rcv_1 \gg rcv_2$ )	Strict ( $rcv \approx 1$ )	Loose ( $rcv \approx N$ )
Destination Unreachable	72.16%	15.46%	2.41%
Time Exceeded	38.84%	1.94%	21.03%
Echo Reply (※ 500 packets)	40.11%	0.88%	35.63%

## IV. Measuring “Rate Limiting”

### *Our Findings*

- ICMP rate limiting is prevalent, 65%-98% implementing significant ICMP rate limiting.
- Rate limiting of ICMP Destination Unreachable is **more strict and easy to observe** (97% not loose with  $N = 50$ , only  $\sim 18\%$  implementing strict or loose rate limiting ).
- **Global ICMP rate limiting is common**, especially ICMP Destination Unreachable.  $>70\%$  implement global rate limiting of ICMP Destination Unreachable.

iVANTAGE can make good use of ICMP Destination Unreachable without sending a large number of packets.

iVANTAGE can be widely used for different RVPs distributed across the Internet.

## IV. Measuring “Rate Limiting”

### *Risks and Mitigation Measures*

#### Global things on Internet seem harmful!

- Researchers exploited global IPID counters for alias resolution, stealthy scans, TCP hijacking. When global IPID counters become fewer, global SYN cache is used as substitutes.
- Global ICMP rate limiting, though less harmful, can also be dangerous (alias resolution, DNS poisoning and other side channel-based measurements).

#### Sending ICMP error message exposes itself!

- It is easy to find an unreachable IP address in such a large IPv6 address space, so it is also easy to induce IPv6 nodes to initiate ICMP Destination Unreachable messages. The node initiating an ICMP Destination Unreachable message exposes itself, which can then be exploited.

## IV. Measuring “Rate Limiting”

### *Risks and Mitigation Measures*

#### Strict or Non-global ICMP Rate Limiting is Recommended

- Non-global ICMP rate limiting is an intuitive solution, but may not be easy to implement and deploy (too many rate limiting counters, e.g., token buckets to maintain).
- Strict ICMP rate limiting is a more simple solution. Even though still global, strict rate limiting makes the differences much less observable. However, it cannot cope with bursty traffic.

Therefore, there may be a trade-off, and it is still difficult to find a perfect solution. Side channels of ICMP rate limiting may be exploitable for a long time to come.

#### ICMP error message should be restricted!

- Allowing IPv6 nodes to generate ICMP Destination Unreachable messages without any restrictions will be dangerous. It exposes itself.

For example, when a router receives a series of packets destined for very strange destination addresses within its subnet (especially if these packets are sent from a remote network!), it may be a safer choice to ignore them than to initiate ICMP destination unreachable messages for each packet.

## V. Limitations and Future Work



## V. Limitations and Future Work


### *More IVANTAGE-based Measurement Applications*

#### *More Measurement Applications*

- It's also possible to send other packets (e.g., TCP-SYN, UDP, DNS queries, NTP, etc.) instead of ping as probe packets. Then we can infer the reachability of them by using IVANTAGE technique.

#### *Demystifying ICMP Rate Limiting Behaviors*

- Characterizing rate limiting behaviors.
- Determining rate limiting more accurately and efficiently with fewer packets to be sent.



Different devices show different packet loss traces.

## VI. Ethics

## VI. Ethics

### Anonymity

- Ensure the anonymity of prefixes and ASes we measured to prevent those vulnerable-to-spoofing networks from being attacked.

### Relatively Harmless Probes

- Compared with other types of scans like port scans and sending DNS queries, sending ICMP Echo Requests (**ping**) is relatively harmless.

### Preventing Continuous Rate Limiting

- RVPs are used in rotation. We prevent triggering ICMP rate limiting continuously on same device.

### Laboratory and Real Internet Experiments

- ICMP rate limiting does not lead to a disruptive impact on either the data plane or the control plane of the target device.



清华大学  
Tsinghua University

Thank You!

## **YOUR ROUTER IS MY PROBER:**

Measuring IPv6 Networks via ICMP Rate Limiting Side Channels

Long Pan, Jiahai Yang, Lin He, Zhiliang Wang, Leyao Nie,  
Guanglei Song, Yaozhong Liu