


# Quantifying the Impact of Blocklisting in the Age of Address Reuse

Sivaram Ramanathan , Anushah Hossain, Jelena Mirkovic, Minlan Yu and  
Sadia Afroz



# IP Blocklists

- IP Blocklists contain a list of known malicious IP addresses.
- IP Blocklists are commonly used to block attack traffic.

1. 198.38.89.61	2. 175.230.213.33	3. 182.74.165.174	4. 178.137.90.85
5. 111.40.73.83	6. 61.132.233.195	7. 193.150.72.50	8. 221.4.205.30
9. 60.172.69.66	10. 61.163.36.24	11. 60.166.48.158	12. 117.214.17.72
13. 180.121.141.117	14. 114.232.216.5	15. 183.159.83.71	16. 121.239.86.33
17. 92.73.213.217	18. 162.248.74.123	19. 183.159.95.87	20. 14.207.215.126
21. 222.191.179.90	22. 217.110.92.194	23. 156.216.145.235	24. 81.17.22.206
25. 41.251.33.175	26. 114.223.61.210	27. 114.232.193.38	28. 114.231.141.136
29. 170.51.62.241	30. 49.67.83.155	31. 180.121.141.119	32. 39.40.30.104
33. 209.54.53.185	34. 167.114.84.153	35. 223.240.208.236	36. 183.150.34.181
37. 95.37.125.239	38. 171.14.238.42	39. 1.55.199.83	40. 222.191.177.40
41. 45.234.101.139	42. 117.85.56.142	43. 123.54.107.199	44. 45.119.81.235
45. 186.47.173.213	46. 49.67.67.141	47. 95.211.149.134	48. 113.128.132.9
49. 49.67.67.140	50. 119.180.198.174	51. 103.69.46.81	52. 128.199.35.34
53. 159.255.167.131	54. 181.215.89.206	55. 192.210.201.168	56. 128.199.44.20
57. 218.72.108.217	58. 113.120.60.120	59. 111.125.140.155	60. 60.50.145.121



# IP Blocklists

- IP Blocklists contain a list of known malicious IP addresses.
- IP Blocklists are commonly used to block attack traffic.
- Blocking reused addresses can lead to unjust blocking of many more users.

1. 198.38.89.61	2. 175.230.213.33	3. 182.74.165.174	4. 178.137.90.85
5. 111.40.73.83	6. 61.132.233.195	7. 193.150.72.50	8. 221.4.205.30
9. 60.172.69.66	10. 61.163.36.24	11. 60.166.48.158	12. 117.214.17.72
13. 180.121.141.117	14. 114.232.216.5	15. 183.159.83.71	16. 121.239.86.33
17. 92.73.213.217	18. 162.248.74.123	19. 183.159.95.87	20. 14.207.215.126
21. 222.191.179.90	22. 217.110.92.194	23. 156.216.145.235	24. 81.17.22.206
25. 41.251.33.175	26. 114.223.61.210	27. 114.232.193.38	28. 114.231.141.136
29. 170.51.62.241	30. 49.67.83.155	31. 180.121.141.119	32. 39.40.30.104
33. 209.54.53.185	34. 167.114.84.153	35. 223.240.208.236	36. 183.150.34.181
37. 95.37.125.239	38. 171.14.238.42	39. 1.55.199.83	40. 222.191.177.40
41. 45.234.101.139	42. 117.85.56.142	43. 123.54.107.199	44. 45.119.81.235
45. 186.47.173.213	46. 49.67.67.141	47. 95.211.149.134	48. 113.128.132.9
49. 49.67.67.140	50. 119.180.198.174	51. 103.69.46.81	52. 128.199.35.34
53. 159.255.167.131	54. 181.215.89.206	55. 192.210.201.168	56. 128.199.44.20
57. 218.72.108.217	58. 113.120.60.120	59. 111.125.140.155	60. 60.50.145.121





# Blocklisting Reused Addresses: NAT




# Blocklisting Reused Addresses: NAT





 **CLOUDFLARE** | community WHAT WE DO BLOG SUPPORT

 **Cloudflare blocking my IP?**  
Security ■ dash-ssl-tls

 kieran.hill1796 Mar '19

I don't know who this company thinks they are but they've accessed my IP without permission, have decided that it is for some reason untrustworthy and are now blocking me from websites, servers and just making the internet unusable for me. I want them to leave me alone, delete all my information and stop blocking me from everything because they have done this with no consent given.

1 Reply  

# Blocklisting Reused Addresses: NAT

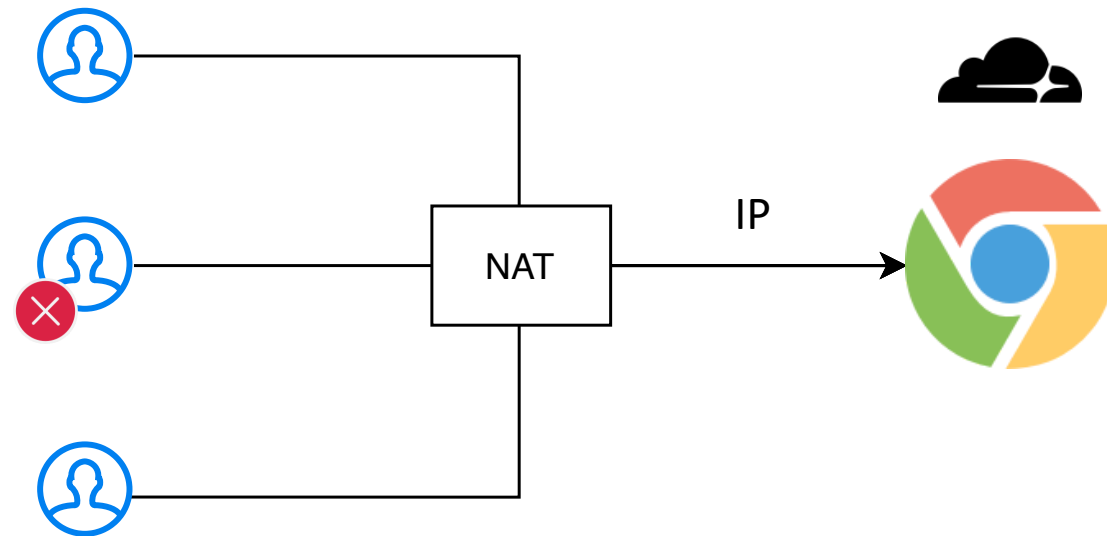


# Blocklisting Reused Addresses: NAT



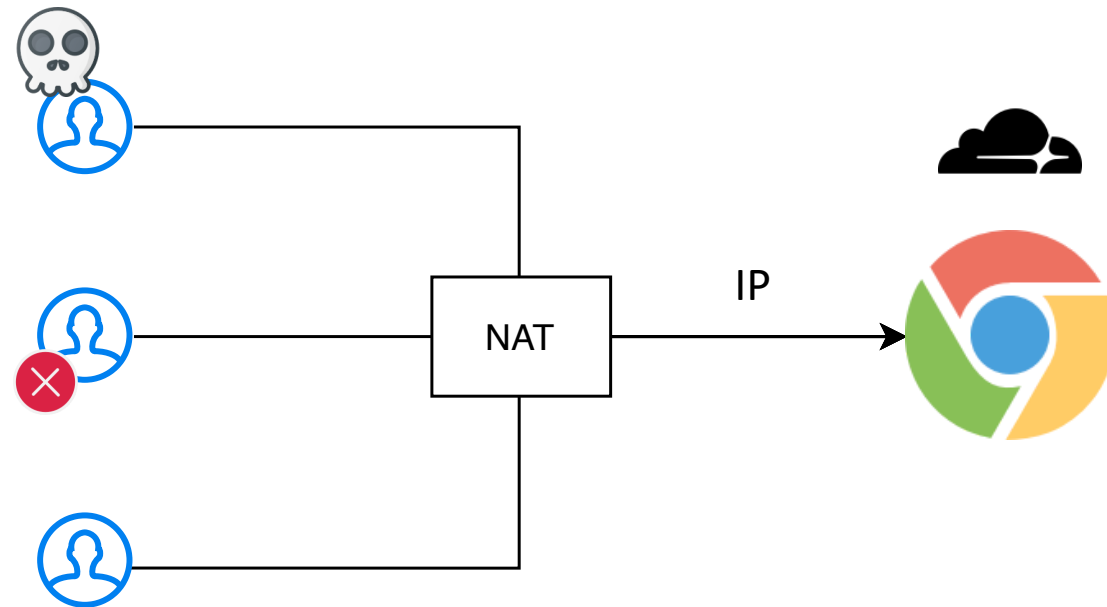
Cloudflare uses  
Dshield blocklist.

# Blocklisting Reused Addresses: NAT

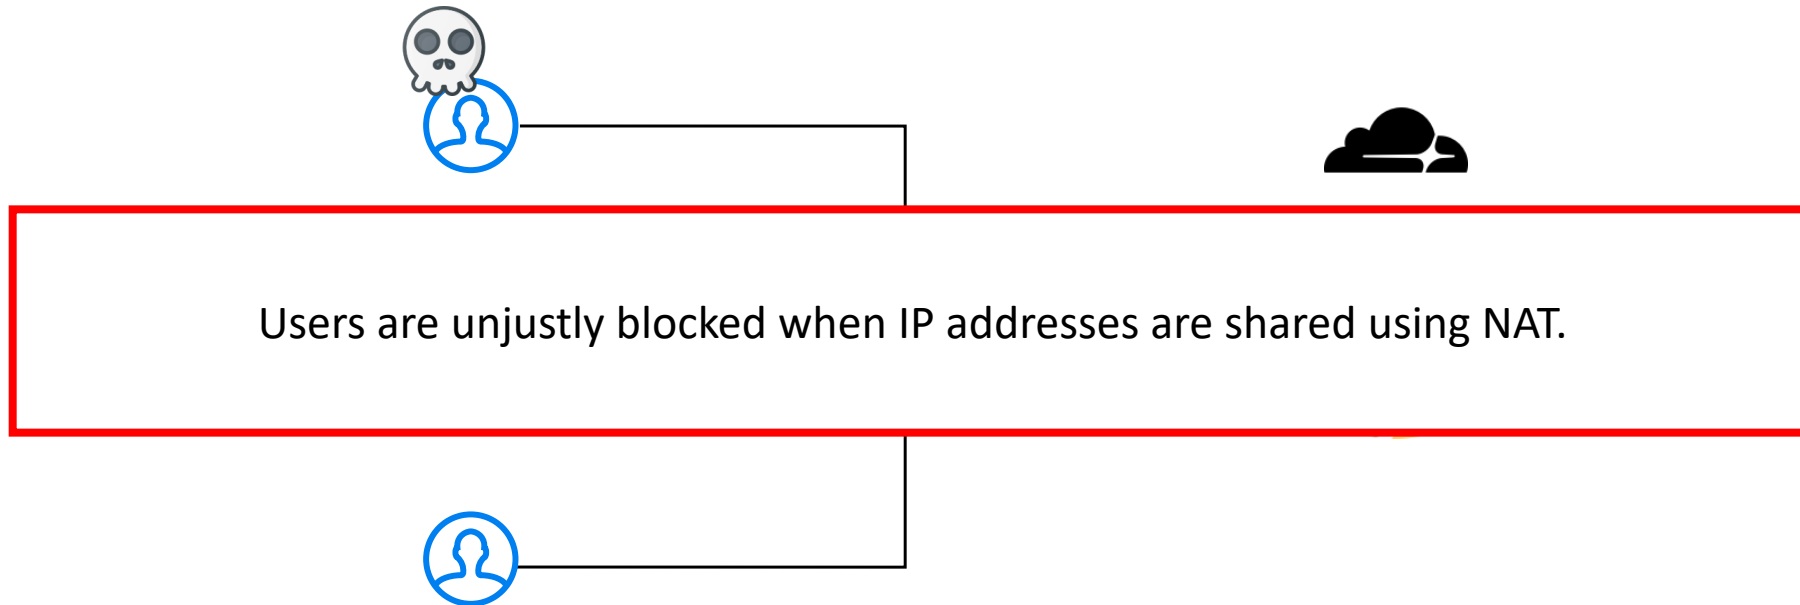





# Blocklisting Reused Addresses: NAT




# Blocklisting Reused Addresses: NAT



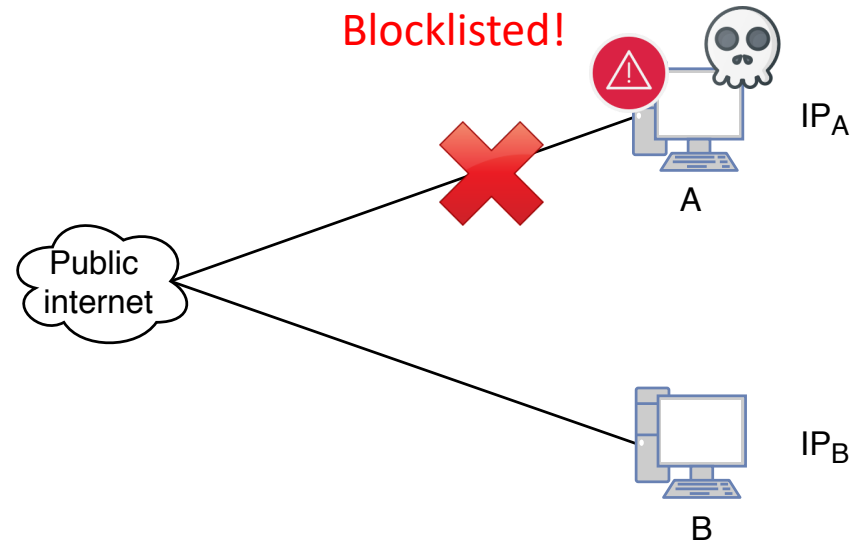
# Blocklisting Reused Addresses: Dynamic Addressing




# Blocklisting Reused Addresses: Dynamic Addressing




# Blocklisting Reused Addresses: Dynamic Addressing




# Blocklisting Reused Addresses: Dynamic Addressing



# Blocklisting Reused Addresses: Dynamic Addressing



# Blocklisting Reused Addresses: Dynamic Addressing





# Outline

- Introduction
- Usage and perception of blocklists
- Identifying reused addresses
  - Detecting NATed addresses
  - Detecting dynamic addresses
- Blocklist dataset
- Evaluation
- Summary and Conclusions

# Usage and Perception of Blocklists

- Surveyed 40 network operators to understand usage of blocklists and their anecdotal experiences on blocklisting reused addresses.
- **Blocklists are commonly used and used for active defense:**
  - 70% of operators used blocklists and 60% of them use blocklists to directly block traffic.

# Usage and Perception of Blocklists

- Surveyed 40 network operators to understand usage of blocklists and their anecdotal experiences on blocklisting reused addresses.
- **Blocklists are commonly used and used for active defense:**
  - 70% of operators used blocklists and 60% of them use blocklists to directly block traffic.
- **Blocklists can have inaccuracies due to reused addresses:**
  - About 56--76% of operators feel inaccuracies in blocklists due to reused addresses.

# Quantifying the Impact of Blocklisting

- Accurately identifying reused addresses.
- Identifying blocklists that list such reused addresses.
- Quantifying the impact of blocking reused addresses.

# Quantifying the Impact of Blocklisting

- Accurately identifying reused addresses.
  - Two techniques using a BitTorrent DHT crawler and RIPE atlas measurement logs.
- Identifying blocklists that list such reused addresses.
- Quantifying the impact of blocking reused addresses.

# Quantifying the Impact of Blocklisting

- Accurately identifying reused addresses.
  - Two techniques using a BitTorrent DHT crawler and RIPE atlas measurement logs.
- Identifying blocklists that list such reused addresses.
  - 151 publicly available blocklists used for detecting variety of malicious users.
- Quantifying the impact of blocking reused addresses.

# Quantifying the Impact of Blocklisting

- Accurately identifying reused addresses.
  - Two techniques using a BitTorrent DHT crawler and RIPE atlas measurement logs.
- Identifying blocklists that list such reused addresses.
  - 151 publicly available blocklists used for detecting variety of malicious users.
- Quantifying the impact of blocking reused addresses.
  - Impact on the number of addresses potentially affected due to blocking reused addresses.

# Outline

- Introduction
- Usage and perception of blocklists
- Identifying reused addresses
  - Detecting NATed addresses
  - Detecting dynamic addresses
- Blocklist dataset
- Evaluation
- Summary and Conclusions




# Detecting Reused NATed addresses

- We use the BitTorrent Network to identify users that are allocated the same IP address.
- The BitTorrent protocol allows two messages that helps us identify NATted users accurately.
  - *get\_nodes*: Returns a list of active neighbors to a node.
  - *bt\_ping*: Periodically pings active neighbors.
- The protocol mandates all BitTorrent users to reply to these messages.


# Detecting NATed addresses

Using *get\_node* messages.




# Detecting NATed addresses


Using *get\_node* messages.



# Detecting NATed addresses




# Detecting NATed addresses




Two active users with two different port numbers using the same IP address.

# Detecting NATed addresses




# Detecting NATed addresses



get\_nodes allows to identify IP addresses with multiple port numbers. bt\_ping verifies if users with different port numbers are active.

# Discovered NATed addresses


- 48.7M IP addresses that use BitTorrent.






# Discovered NATed addresses

- 48.7M IP addresses that use BitTorrent.
- 1.6B bt\_ping messages sent.

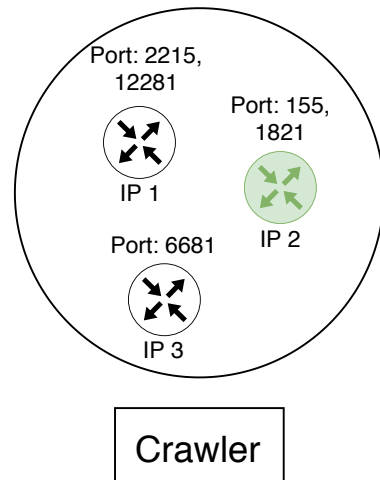


# Discovered NATed addresses



- 48.7M IP addresses that use BitTorrent.
- 1.6B bt\_ping messages sent.
- 779M responses (48.6%).

# Discovered NATed addresses




- 48.7M IP addresses that use BitTorrent.
- 1.6B bt\_ping messages sent.
- 779M responses (48.6%).
- 2M IP addresses that are NATed.

# Outline

- Introduction
- Usage and perception of blocklists
- Identifying reused addresses
  - Detecting NATed addresses
  - Detecting dynamic addresses
- Blocklist dataset
- Evaluation
- Summary and Conclusions

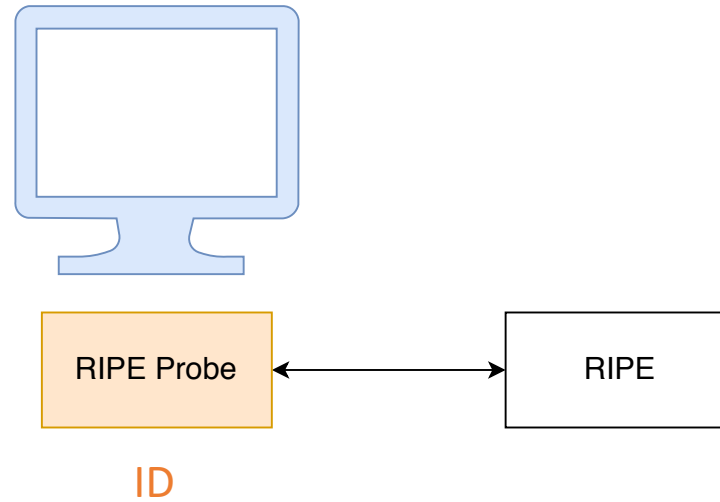
# Detecting Dynamic Addresses



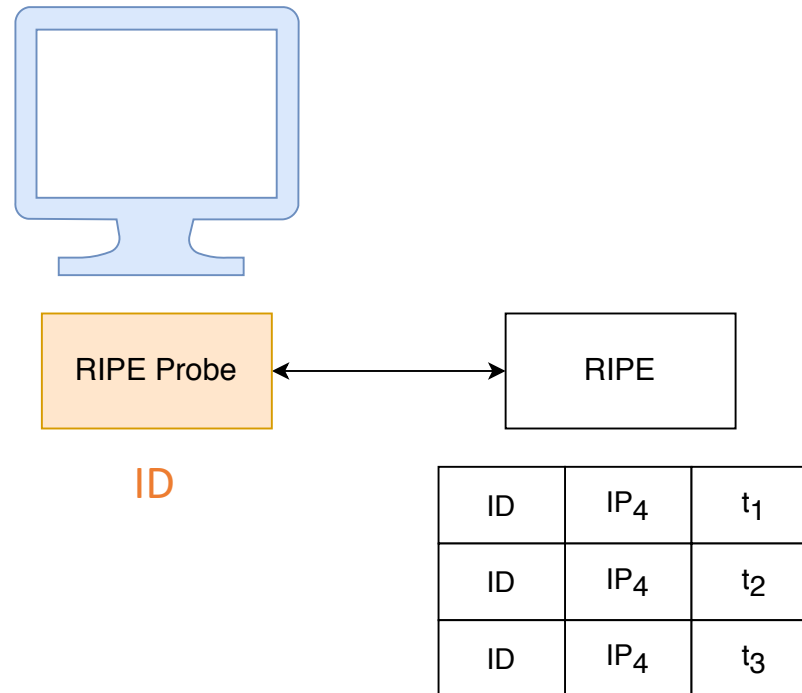
RIPE Probe

ID

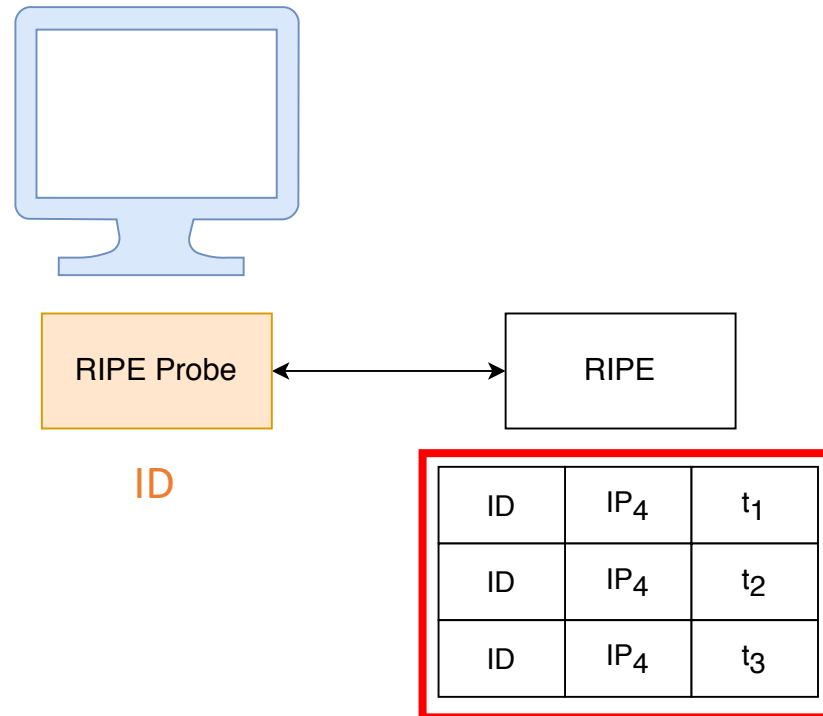
# Detecting Dynamic Addresses



# Detecting Dynamic Addresses



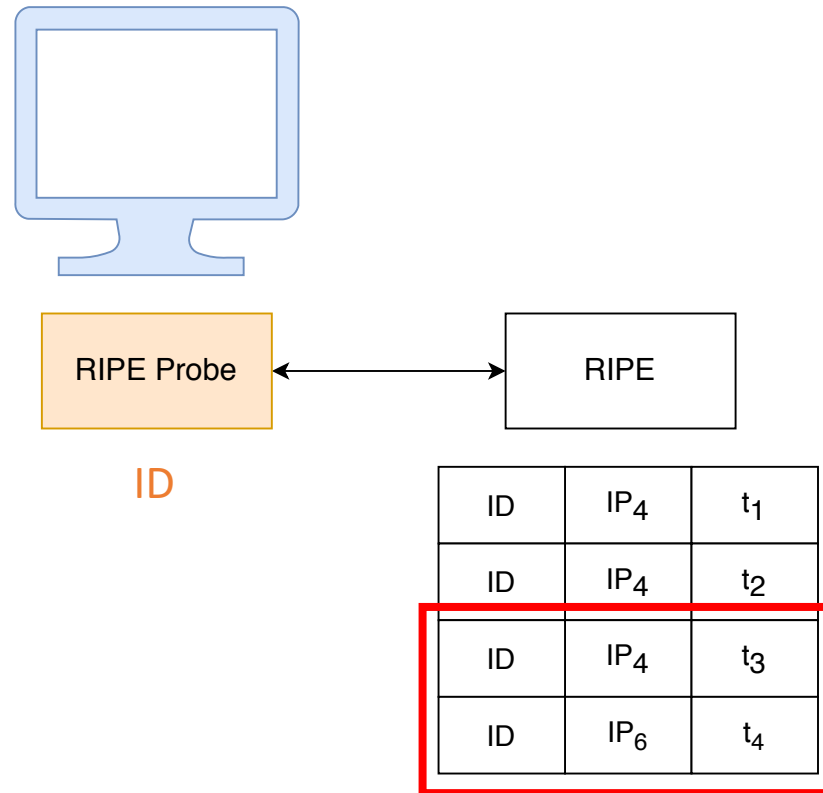
# Detecting Dynamic Addresses



Measurement logs to determine dynamically allocated addresses.



# Detecting Dynamic Addresses



IP<sub>4</sub> and IP<sub>6</sub> are potentially dynamically allocated.

# Detecting Dynamic Addresses

To prevent users  
that have changed  
ISPs.

Probes with  
addresses  
changes in the  
same AS.

**Remaining:** 13.6K RIPE probes

# Detecting Dynamic Addresses

To prevent users that have changed ISPs.

To consider probes that are potentially dynamically allocated.

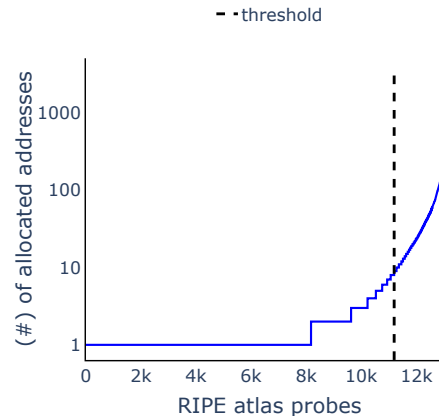


**Remaining:** 13.6K RIPE probes

2.6K RIPE probes

# Detecting Dynamic Addresses

To prevent users that have changed ISPs.




To consider probes that are potentially dynamically allocated.



Remaining: 13.6K RIPE probes

2.6K RIPE probes

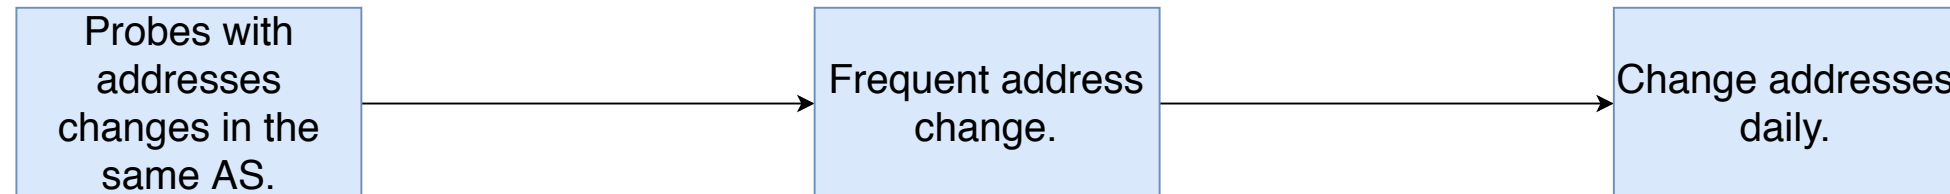
# Detecting Dynamic Addresses



To prevent users that have changed ISPs.

To consider probes that are potentially dynamically allocated.

Addresses that will have maximum impact on being blocklisted.



Remaining: 13.6K RIPE probes

2.6K RIPE probes

629 RIPE probes


# Outline

- Introduction
- Usage and perception of blocklists
- Identifying reused addresses
  - Detecting NATed addresses
  - Detecting dynamic addresses
- **Blocklist dataset**
- Evaluation
- Summary and Conclusions

# Quantifying Impact with Blocklists


- We use the BLAG dataset that actively maintains blocklisted addresses from public blocklists.
- **151 blocklists** that monitor variety of attacks including Spam, DDoS, malware hosting or reputation of IP addresses.
- Monitoring period of **83 days** over two measurement periods:
  - Aug 2019 – Sep 2019
  - Mar 2020 – May 2020
- Observed **2.2M blocklisted IP addresses**.

# Number of Reused Addresses in Blocklists







# Number of Reused Addresses in Blocklists




# Number of Reused Addresses in Blocklists




# Number of Reused Addresses in Blocklists

## NATed addresses




## Dynamic addresses




# Number of Reused Addresses in Blocklists

## NATed addresses




## Dynamic addresses




# Number of Reused Addresses in Blocklists

## NATed addresses



## Dynamic addresses




# Outline

- Introduction
- Usage and perception of blocklists
- Identifying reused addresses
  - Detecting NATed addresses
  - Detecting dynamic addresses
- Blocklist dataset
- Evaluation
- Summary and Conclusions

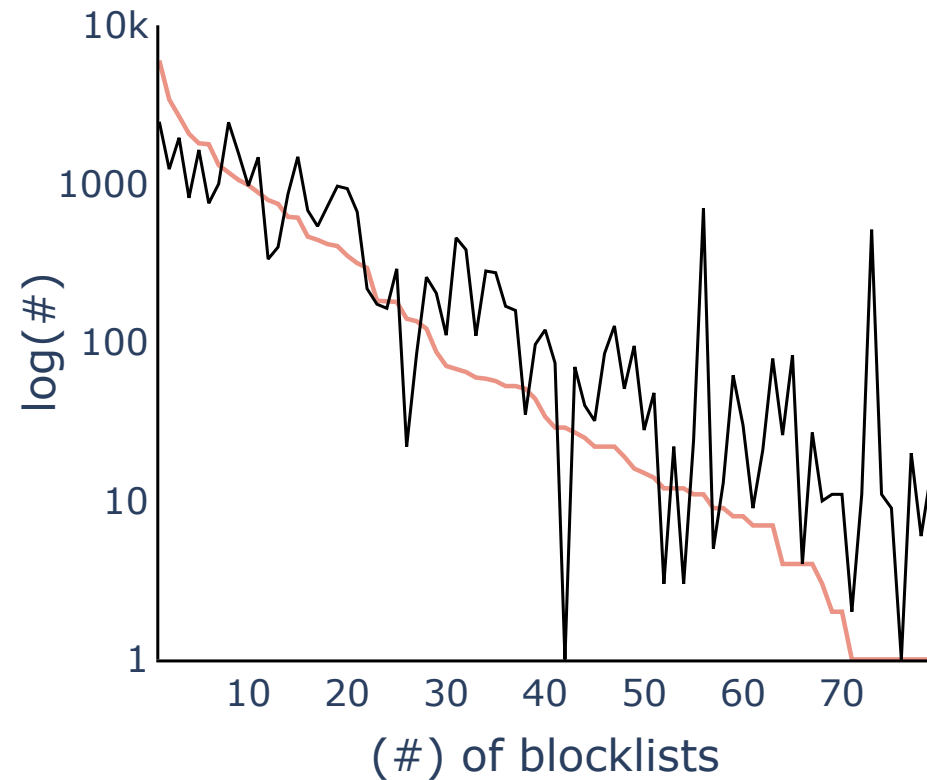
# How many Blocklists list reused addresses?

**NATed Addresses**




**Dynamic Addresses**

— RIPE — Cai et al.




# How many Blocklists list reused addresses?

## NATed Addresses




## Dynamic Addresses






# How many Blocklists list reused addresses?

## NATed Addresses

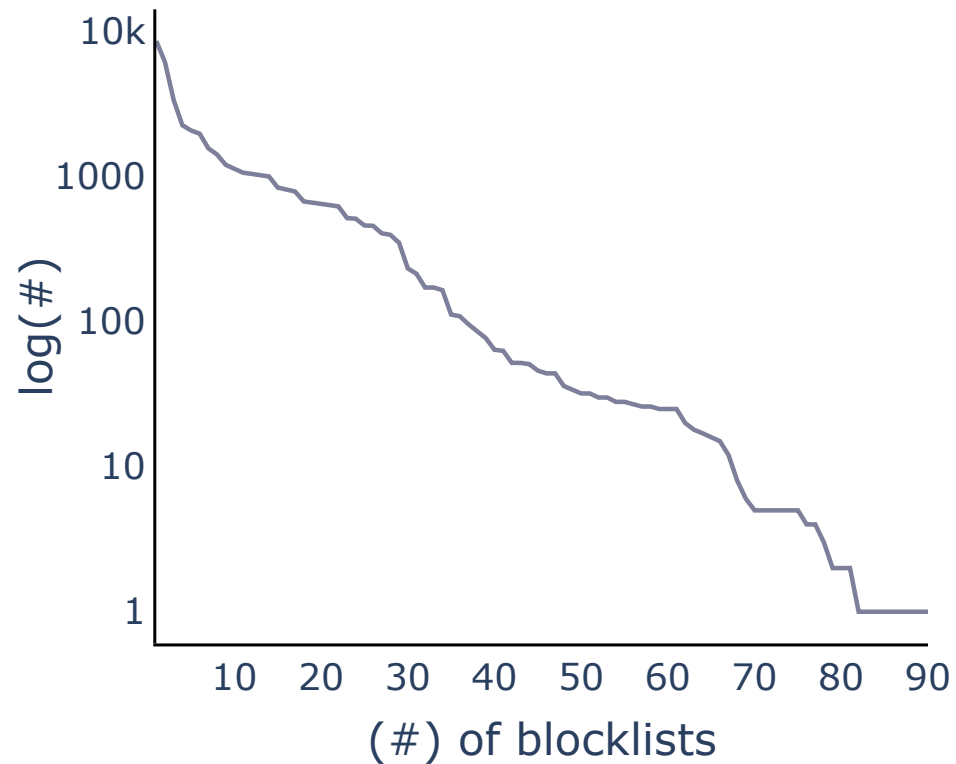


## Dynamic Addresses




# How many Blocklists list reused addresses?

**NATed Addresses**



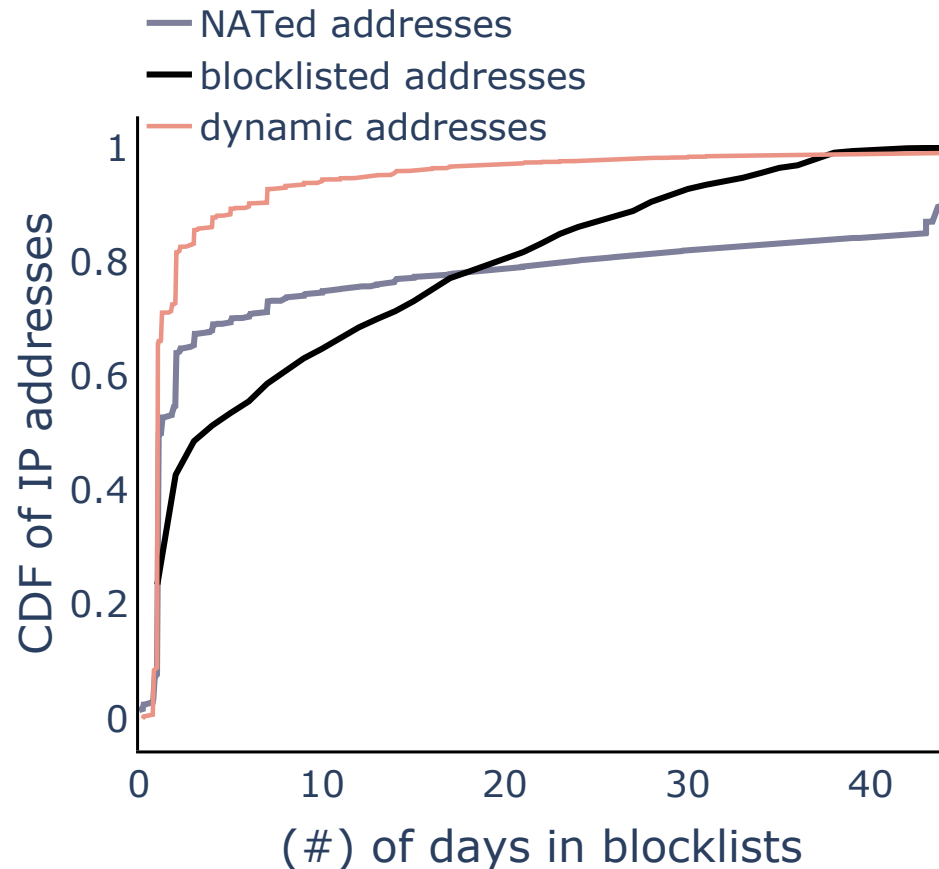
**Dynamic Addresses**

— RIPE — Cai et al.




Our technique is comparable to existing reproducible technique.

# How long are reused addresses in Blocklists?



- Reused addresses are removed faster than other addresses (3—9 days).
- Among reused addresses, dynamically allocated addresses are removed quicker.
- Within two days, 77% of dynamic addresses are removed compared to only 42% of all blocklisted addresses.

# How many users are affected?



- Some IP addresses impact many more users, affecting as many as 78 users.
- Many IP addresses have only two active users (68.5%)
- 98% of IP addresses have less than 10 active users.

# Outline

- Introduction
- Usage and perception of blocklists
- Identifying reused addresses
  - Detecting NATed addresses
  - Detecting dynamic addresses
- Blocklist dataset
- Evaluation
- Summary and Conclusions

# Making Blocklists Better

- We make our datasets public to improve blocklists.
- **Network Operators:**
  - Can use our list to treat reused addresses differently using techniques such as greylisting.
- **Blocklist Maintainers:**
  - Improve blocklist accuracies by monitoring these reused addresses.
- **Other services:**
  - Use to deploy warnings to end users that access services from reused blocked addresses.

# Limitations

- Detecting NAT

- Detects reused addresses only for BitTorrent users
- IP addresses that have more than two active users.
- Some ISPs may block BitTorrent traffic.

- Detecting dynamic addresses

- Limited only to IP prefixes that have deployed RIPE probes.
- Miss IP addresses that have been allocated to different ASes of the same ISP.
- Incorrect boundary detection of IP prefix.

# Conclusion

- We propose two techniques of identifying reused addresses in blocklists.
- We detect reused addresses in 151 publicly available blocklists and find as many as 60% of blocklists that have listed at least one reused address.
- Our NAT detection technique has found reused address in blocklists that have affected 78 users.
- Our datasets are public:
  - [https://steel.isi.edu/members/sivaram/blocklisting\\_impact/](https://steel.isi.edu/members/sivaram/blocklisting_impact/)



# Thank You! Questions?

All detected reused addresses are present in:

[https://steel.isi.edu/members/sivaram/blocklisting\\_impact/](https://steel.isi.edu/members/sivaram/blocklisting_impact/)

All monitored blocklists are available at:

<https://steel.isi.edu/Projects/BLAG/>



**USC** University of  
Southern California

