EvilCohort: Detecting Communities of Malicious Accounts on Online Services

Gianluca Stringhini[1], Pierre Mourlanne[2], Gregoire Jacob[3], Manuel Egele[4], Christopher Kruegel[2,3], and Giovanni Vigna[2,3]

University College London[1] UC Santa Barbara[2] Lastline Inc. [3] Boston University[4]

Online services are abused by cybercriminals

- Spam
- Crawling sensitive information / documents
- Storing illegal content
- Running DoS attacks / hosting C&C servers

State of the art in malicious account detection

Current techniques leverage domain-specific elements to detect malicious activity on one type of service

Exam

- For
- Blog
- Youtube [Benevenuto2009]
- Social Networks [Mittal2009], [Grier2010], [Stringhini2010]
- Webmail accounts [Taylor2006], [Stringhini2015]

There are elements that are common to malicious activity on all online services!

Botnets accessing online accounts

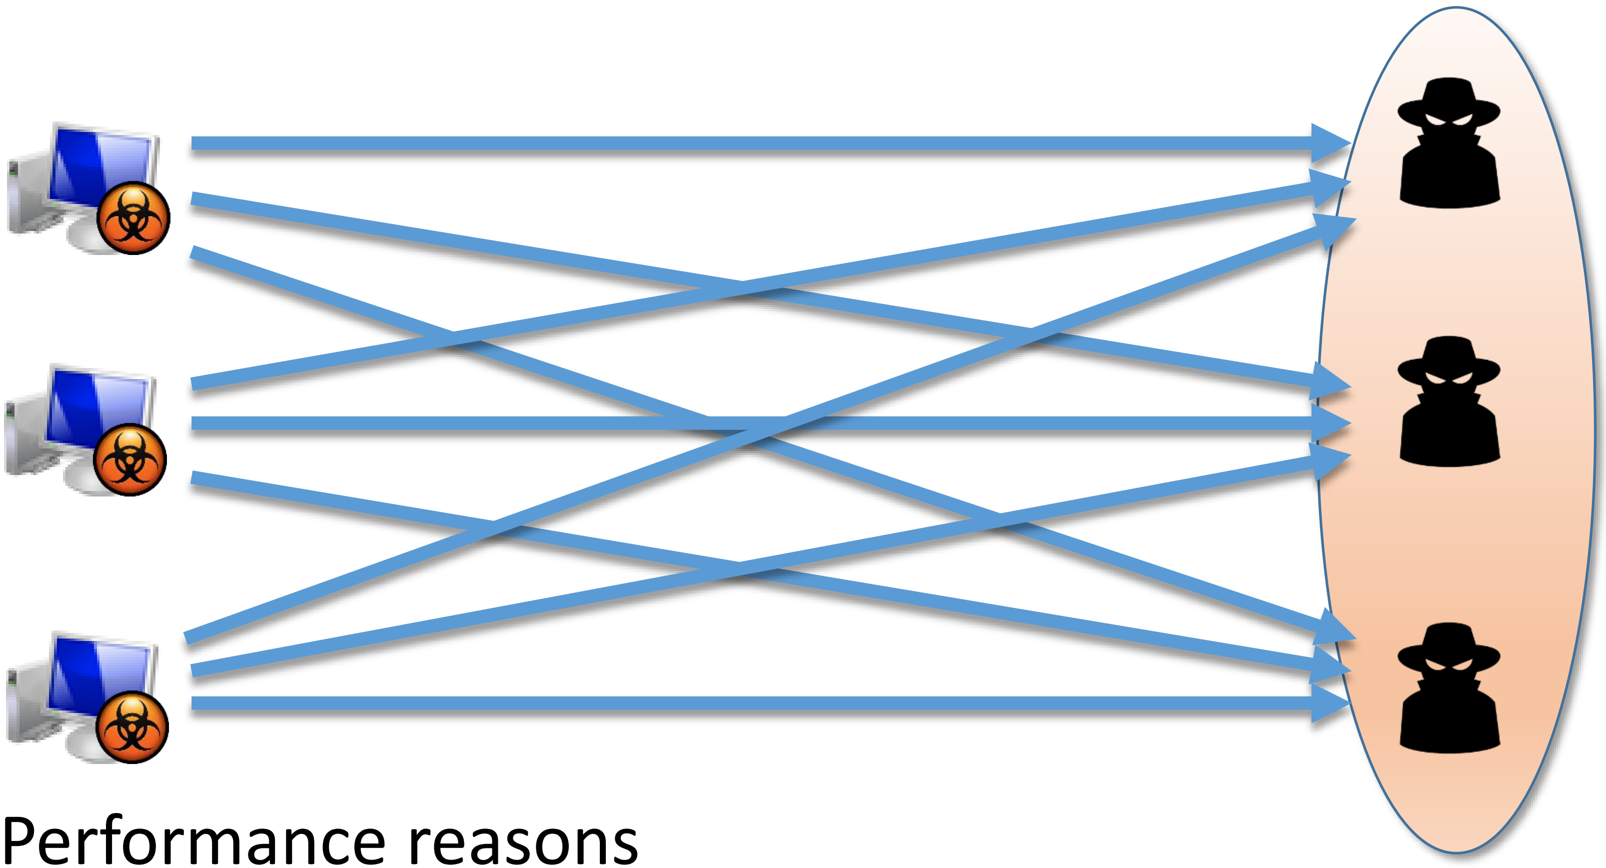

- Performance reasons
- Resilience reasons

Advantages of community detection

Service-agnostic

Can be done on any service that uses accounts

Activity-agnostic

We only look at how accounts are accessed

Different types of cybercriminal operations

- Crawl the online service
- Use the service as C&C channel
- Use the service as a “drop” service

Distributed access is prevalent

Web-based email service logs, 1 day period

72M emails sent by 21M distinct accounts

170k vetted spam accounts for ground truth

- 66k accounts accessed by a single IP address
- 104k accounts accessed by multiple IP addresses

Just looking at accounts that are accessed by many IP addresses does not work (32% FPs for accounts accessed by 10+ IPs)

Our system: EvilCohort

- Phase I: data collection
- Phase II: building the graph representation
- Phase III: finding communities
- Phase IV (optional): characterizing communities

Phase I: data collection

Timestamp_1, IP_address_1, Account_1
Timestamp_2, IP_address_2, Account_2
Timestamp_3, IP_address_3, Account_3
Timestamp_4, IP_address_4, Account_1
…

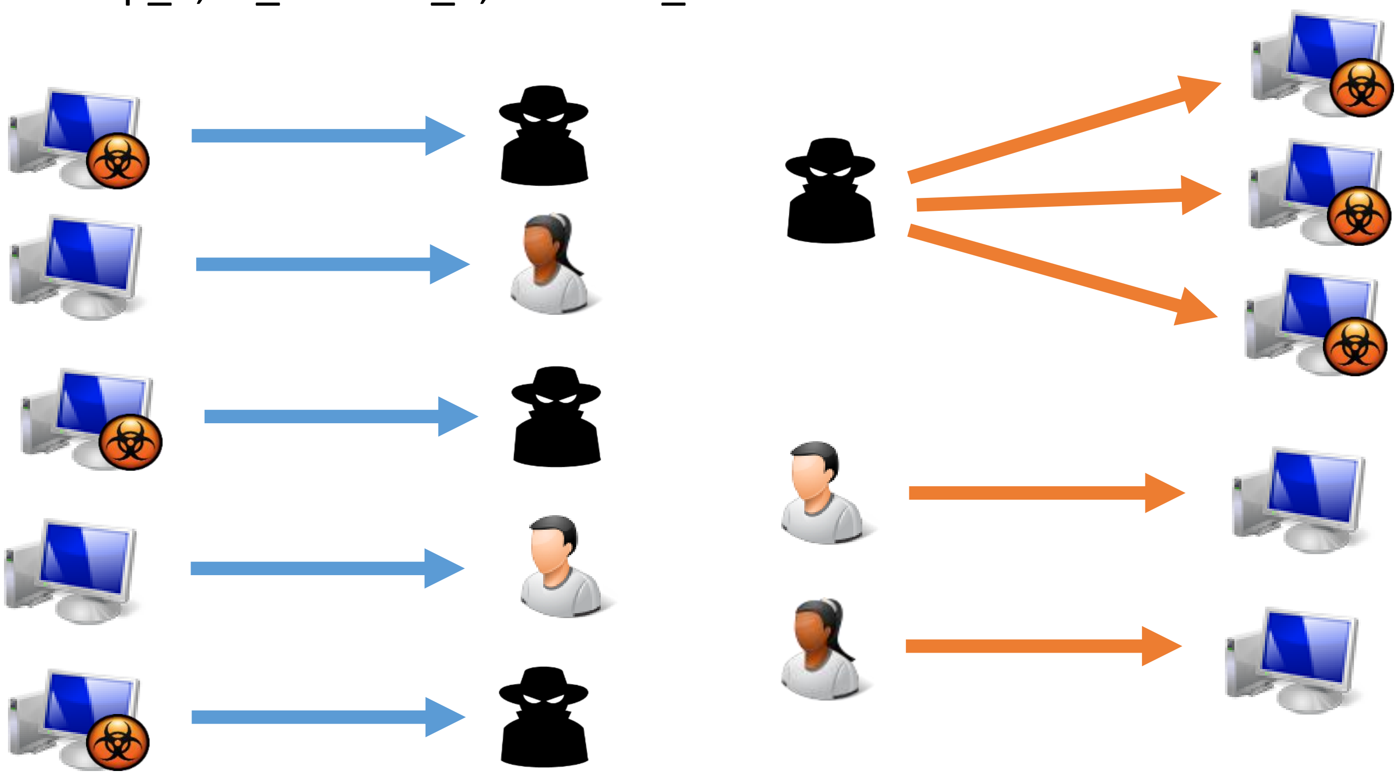

Phase II: building graph representation

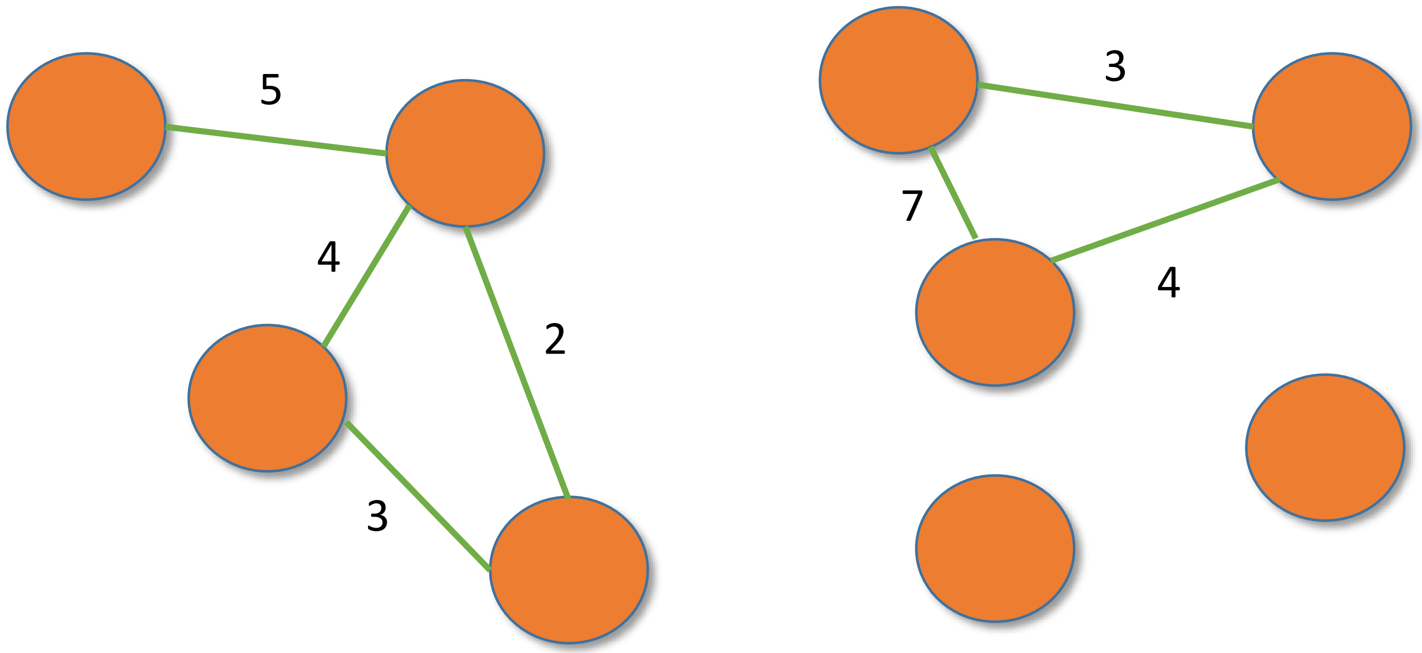

- Vertices are online accounts
- Edges' weight is number of shared IP addresses

Phase III: finding communities

We apply the ``Louvain'' method for clustering:

- Iterative method
- Based on modularity optimization

We can prune edges with low weight to improve precision (threshold *s*)

Phase IV (optional): characterizing communities

- User agent correlation
- Event-based time series
- IP address and account usage

These filters can be used to further prune false positives

Selection of *s*

Ground truth: 103k spam accounts accessed by 2+ IPs

False positive if ≤ 10% of the accounts sent spam

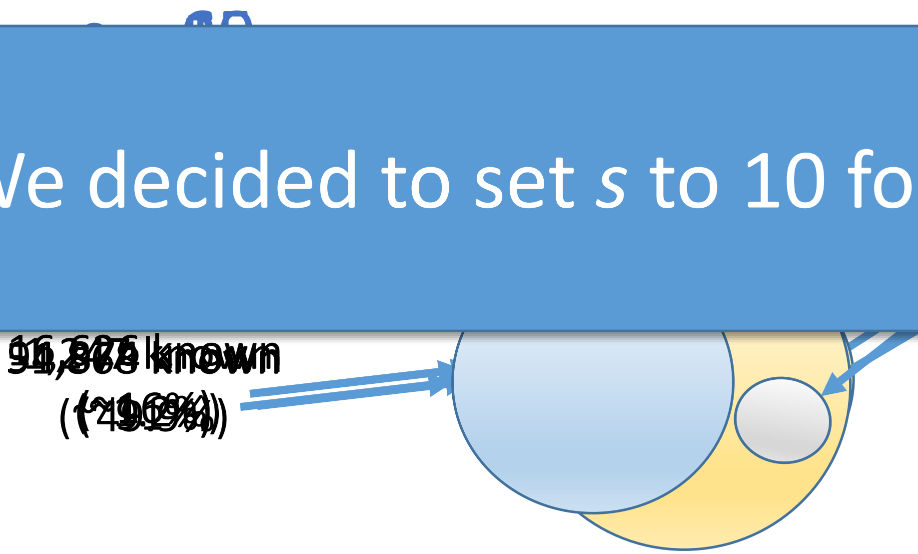

We decided to set *s* to 10 for our experiments

Grown knowledge:

Results in the wild

Webmail activity dataset: email events

5 month period, 1.2B emails

1.2M malicious accounts, 500k unknown, 23k FP (1.9%)

Online social network dataset: login events

8 day period, 14M events, 4 social networks

111k malicious accounts

Analysis of the results

111k accounts formed 83 communities

5 co
typi st
one

What is Charlie?

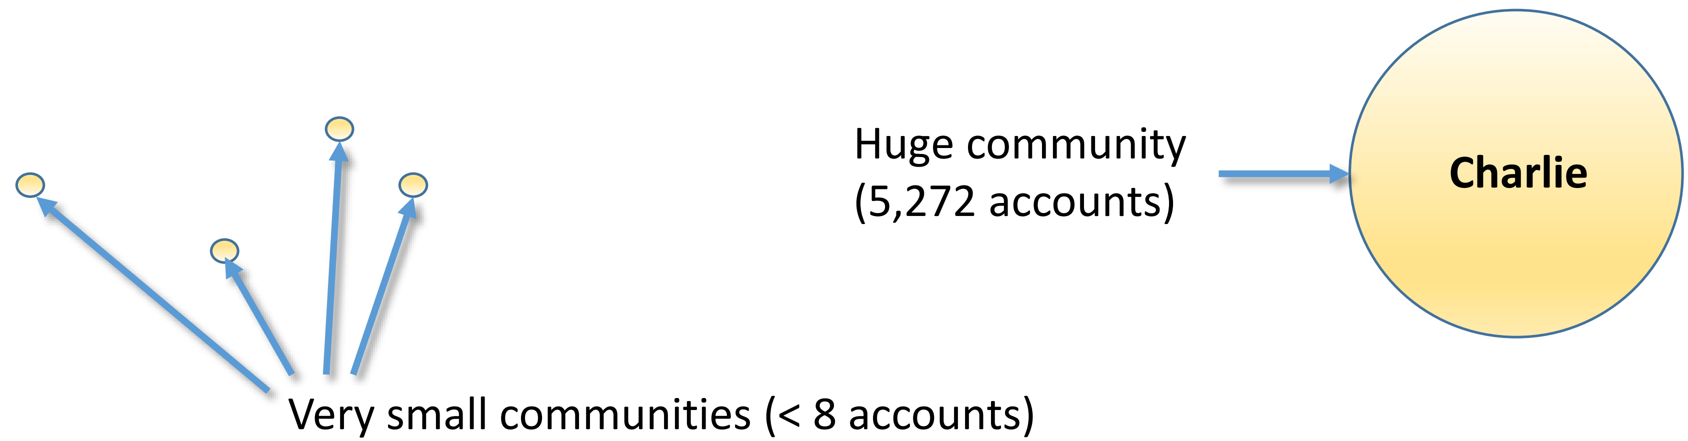

Event-based time series

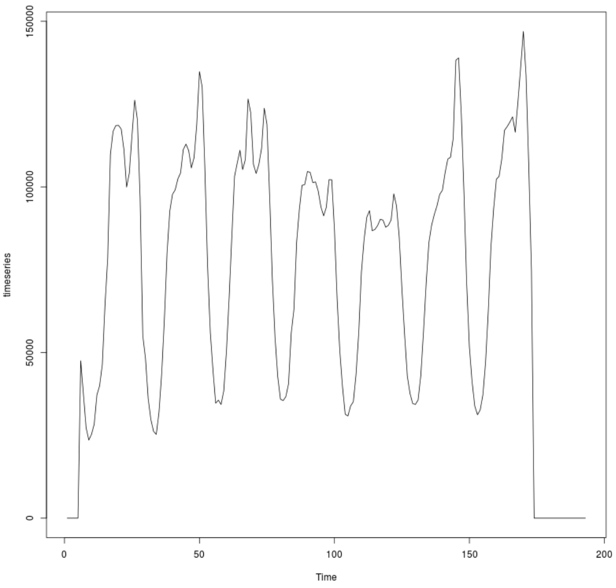

Regular accounts show diurnal patterns

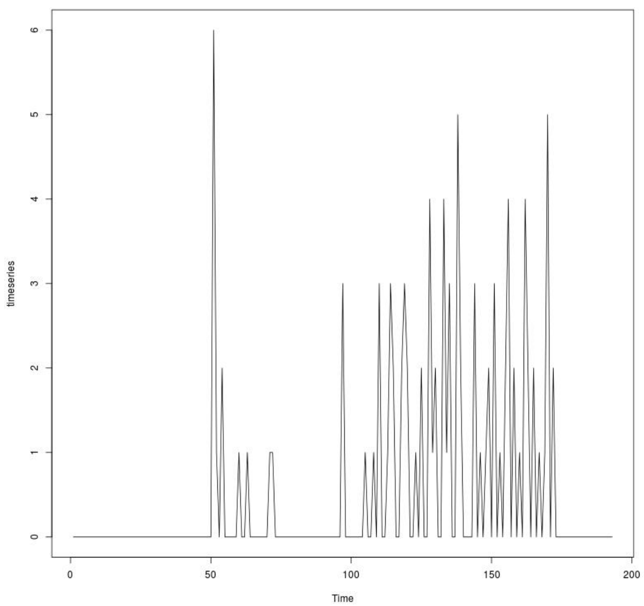

Malicious accounts show bursts in activity

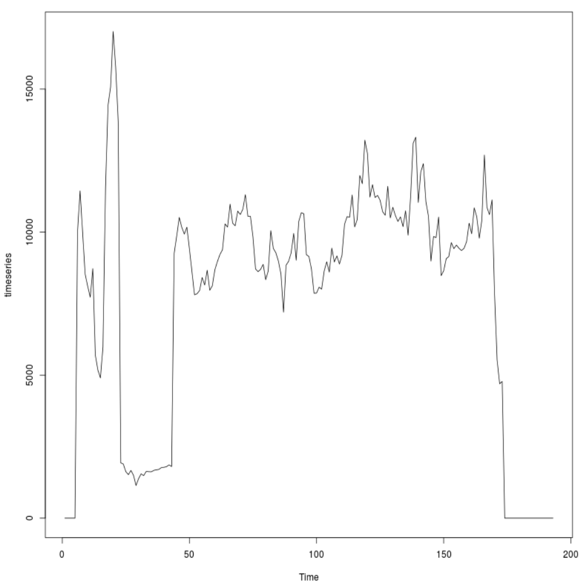

Charlie shows a weird behavior

Account usage over time

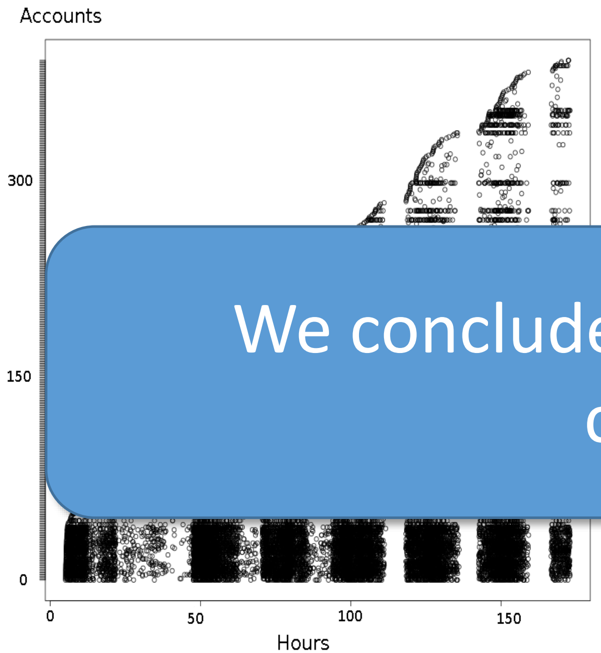

Legitimate accounts:
still diurnal patterns

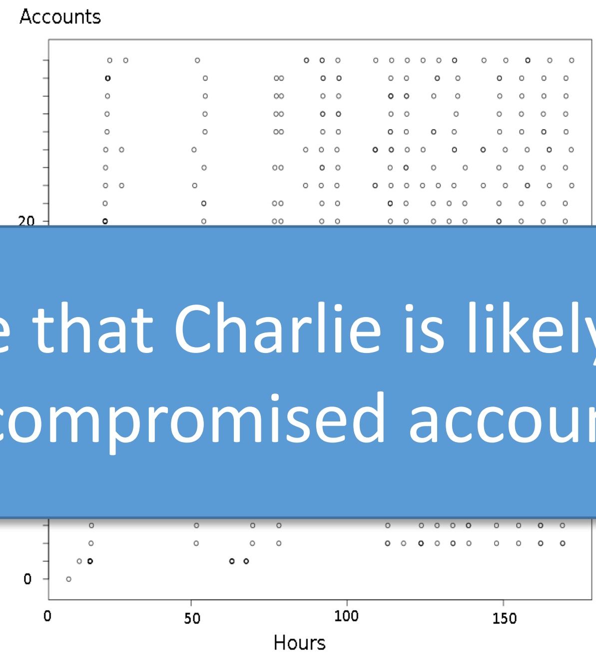

Malicious accounts:
synchronized access

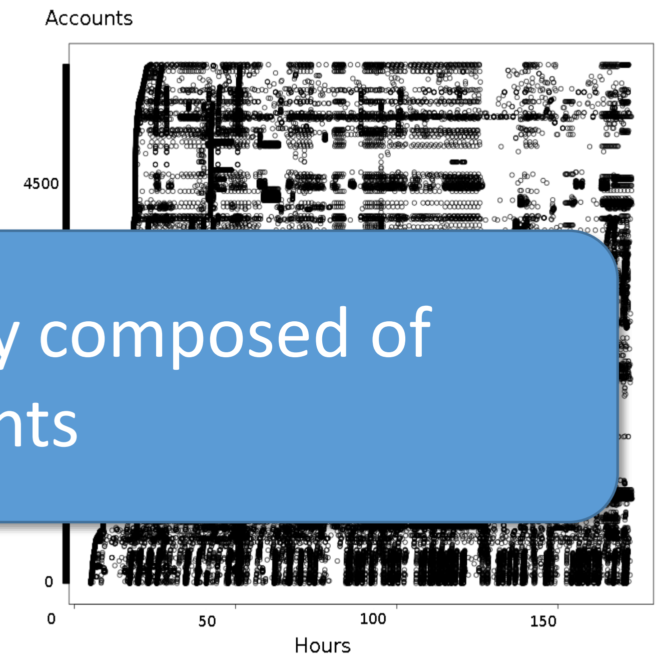

Charlie: regular +
synchronized access

We conclude that Charlie is likely composed of compromised accounts

EvilCohort: discussion

Service and activity independent

Accounts do not need to perform malicious activity to be detected

Our system detects botnet-like activity, legitimate accounts are unlikely to form communities

Limitations

- Only works on accounts accessed by multiple IP addresses
- Does not distinguish between fake and compromised accounts

Conclusions

I presented EvilCohort, a system that detects malicious accounts on online services by identifying communities of accounts that are accessed by a common set of IP addresses

We ran EvilCohort on two real-world datasets, and detected more than one million malicious accounts

Questions?

g.stringhini@ucl.ac.uk
@gianluca_string

UCL