

Data Placement, Replication, Distribution and Streaming Services: State-Of-The-Art Approaches and Open Challenges

Despoina Antonakaki


What is Twitter?

● Fast growing OSN

- **Registered users:** >500 million (On 5/7/2013: 554,750,000 users)
- **Daily tweets:** 6th birthday last year ~340 million....now ~400 million
- **Queries:** ~18,000 queries per second, >1.6 billion per day

"the SMS of the Internet.": microblogging service

Twitter Growth 2006-2011



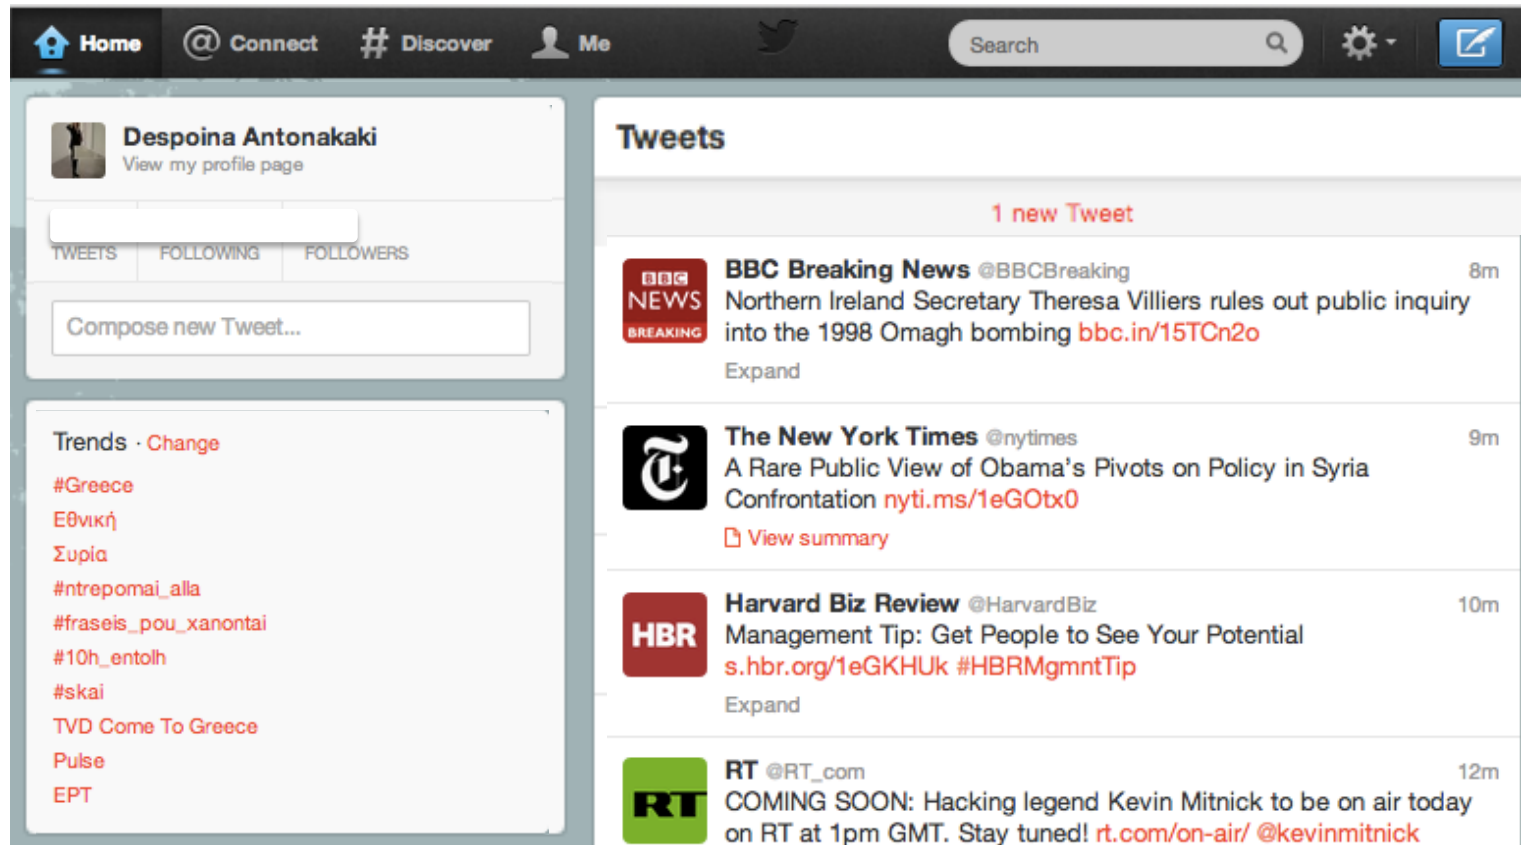
What is Twitter?

Components:

- **Tweets:** 140 character messages
- **Followers:** follower is someone who subscribes to receive your updates
 - Your messages in their timeline
- **Friends:** Users who both follow each other

Relationships **not bidirectional**,

- Users can receive tweets without revealing their tweets



The screenshot displays a Twitter interface. At the top, navigation tabs include Home, Connect, Discover, and Me. A search bar and settings icon are also visible. The profile section for Despoina Antonakaki shows her name, a profile picture, and a 'View my profile page' link. Below this are tabs for TWEETS, FOLLOWING, and FOLLOWERS, and a 'Compose new Tweet...' input field. The 'Trends' section lists several hashtags, including #Greece, Εθνική, Συρία, #ntrepomai_alla, #fraseis_pou_xanontai, #10h_entolh, #skai, TVD Come To Greece, Pulse, and EPT. The 'Tweets' section shows three tweets: 1) BBC Breaking News (@BBCBreaking) with 8m retweets, about Theresa Villiers ruling out a public inquiry. 2) The New York Times (@nytimes) with 9m retweets, about Obama's policy in Syria. 3) Harvard Biz Review (@HarvardBiz) with 10m retweets, about a management tip. A fourth tweet is partially visible, starting with 'RT @RT_com' and 12m retweets.

Four 4-tiers of relationships:

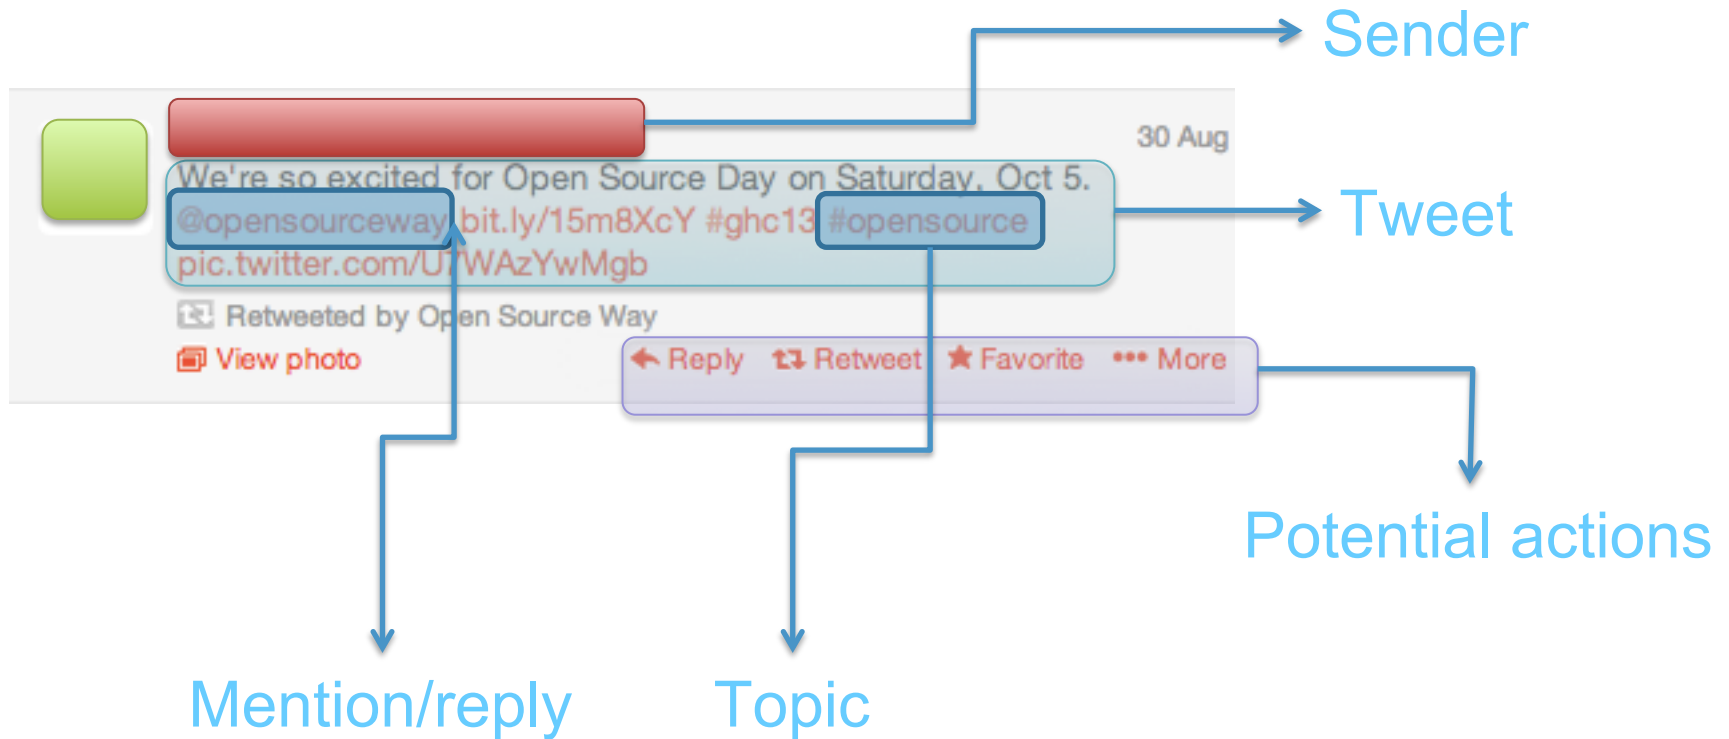
- People who follow you, but you don't follow back
- People who don't follow you, but you follow them
- You both follow each other (Friends!)
- Neither of you follow each other

What is Twitter?

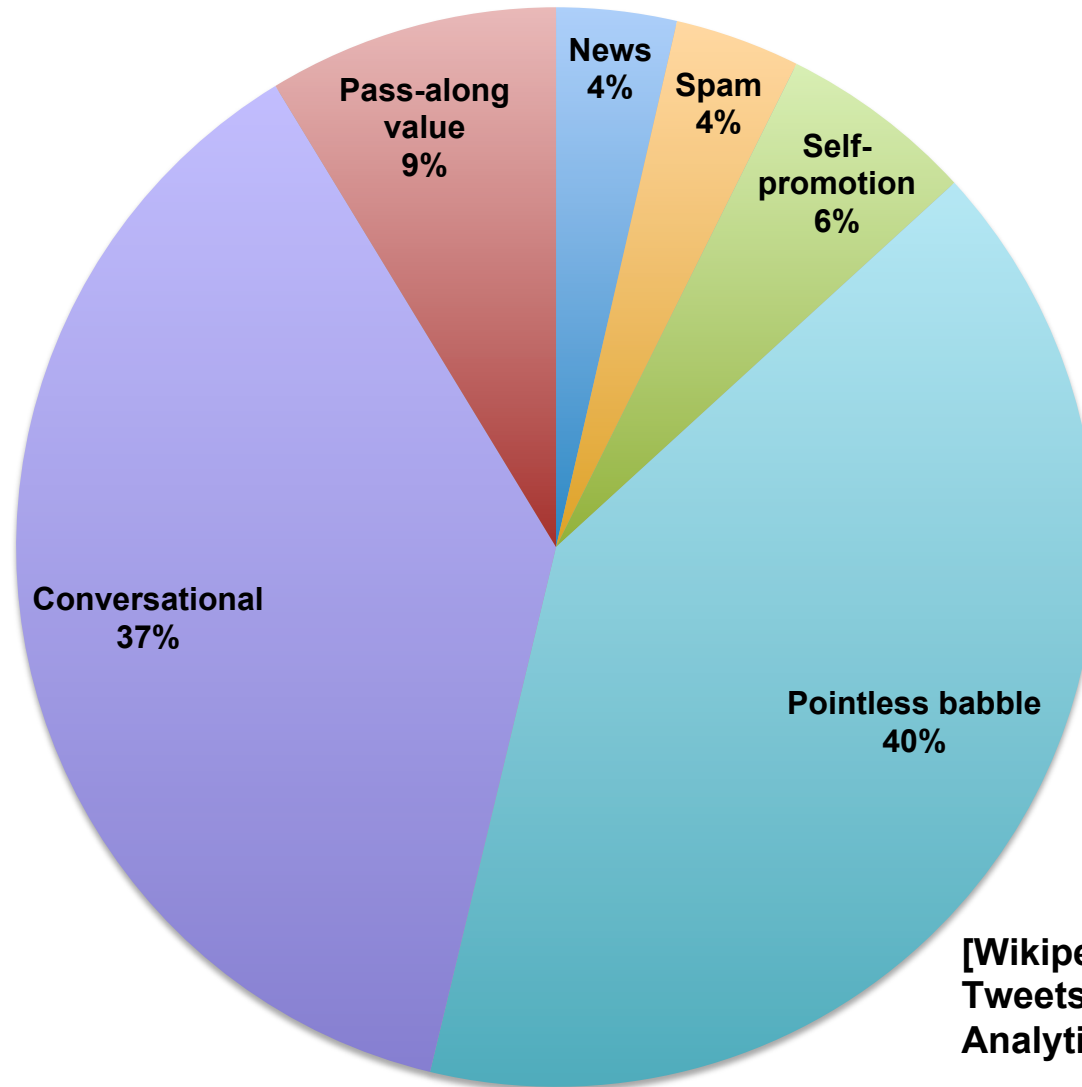
Features:

- **#hashtag**: grouping of messages in topics
 - Get free followers #FF #Follow Justin Bieber
- **@user - mention or reply to user**: a Twitter update (a tweet) directed to another user in reply to their update.
 - @justinbieber PLEASE FOLOOWW ME!!! <333
- **Retweet** message: Repost a tweet from other user
 - RT @JBieberCrewz: RT this if u <3 justin bieber

What is Twitter?



Tweet types



[Wikipedia: Content of Tweets according to Pear Analytics]

The problem: Spam in Twitter

Twitter is **ideal** place for spammers:

- **Content can be public**
- Network of followers created **without approval** of users
- **Missing filtering mechanism** for spam prevention
 - Except malware, blocked using Google's Safebrowsing API
 - Only heuristics to quantify spamming activity:
 - According to Thomas SIGCOMM'11 77% are suspended within their first day on twitter

Twitter Features used by spammers:

- **Callouts.** Mentions used to personalize
 - Example: Win an iTouch AND a \$150 Apple gift card @victim!
<http://spam.com>

[Chris Grier, CCS2010]

Features used by spammers

- **Tweet hijacking.** Hijack users tweets & re-tweet, prepending the tweet with spam URLs
 - Example: <http://spam.com> RT @barackobama A great battle is ahead of us
- **Retweets.** According to Chris Grier et al. 1.8-11.4% are re-tweets of blacklisted URLs
 - Re-tweets purchased by spammers from respected Twitter members,
 - Spam accounts re-tweeting other spam,
 - Hijacked re-tweets
 - Example: RT @scammer: check out the Ipads there having a giveaway <http://spam.com>

Features used by spammers

- **Trend setting.** Spammers attempting create trending topic:
 - 70% of phishing and malware spam containing hashtags
 - generating over 52,000 tweets containing a single tag
 - Example: Buy more followers! <http://spam.com> #fwlr
- **Trend hijacking.** Append currently trending topics to their own spam & ride on the success of other topics
 - 86% of trends used by spammers also appear in benign tweets, with popular trends
 - According to [Thomas SIGCOMM '11], 17% of spam account rely on hijacking trends
 - Example: Help donate to #haiti relief: <http://spam.com>

Spam behavior

- Aggressive followers
- Not all users follow back
- Number of URLs in tweets
- Automated structure of messages
 - Or Twitter API to generate messages
- Profile names from templates

[Gianluca Stringhini, ACSAC '10]

- Abuse compromised accounts rather than accounts generated solely for spamming

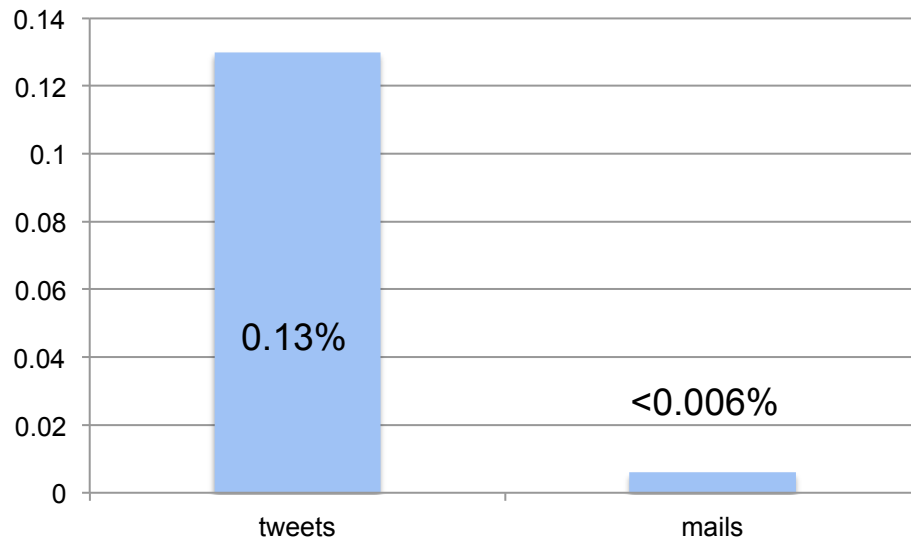
[Chris Grier, CCS2010]

Click-through rate in twitter

- **Measures success** of online advertising campaign
- Is the **number of clicks** on the advertisement **divided** by the times an advertisement was served

Why? Naivety & lack of information:

- Users judge if spam or not only by 140 characters
- Implicit trust for accounts befriend




[Chris Grier, CCS2010]


Feature frequency in spam by blacklists

- According to **Chris Grier et al.** 8% of 25 million URLs posted **point to phishing, malware, and scams** listed on popular blacklists.
- Note **blacklists lag behind Twitter**: most spam messages **appear 4–20 days before** the URLs embedded in messages **become flagged**.

According to Google



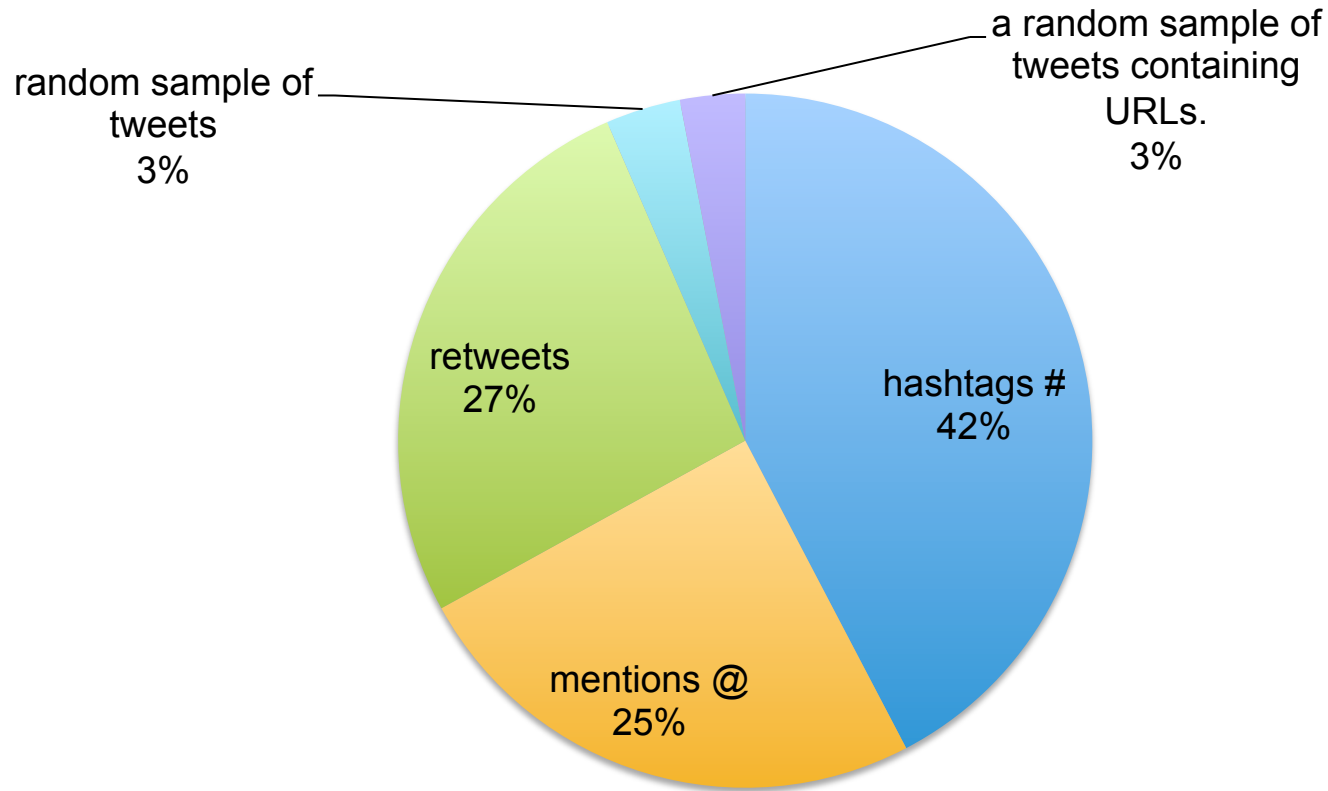
According to Joewein



[Chris Grier, CCS2010]

Feature frequency in spam by blacklists

According to URIBL



[Chris Grier, CCS2010]

Twitter: A marketplace for spammers

Twitter account sellers:

- **Assetize accounts:** > 815 accounts, composing & sending tweets for each account. Owner would be paid.
- **Sponsored Tweets:** registration & ads sent. Currently 905 accounts.

Ad-based URL shorteners

- eca.sh and vur.me: public URL shortening services
 - embed the destination page in an Iframe

[Thomas SIGCOMM '11]

Twitter: A marketplace for spammers

Spam affiliate programs (for underground market diversification)

Automated marketing programs where Web advertisers/merchant recruits webmasters to place merchant's banner ads/buttons on their Web site.

Webmasters get fee/commission on user click (purchase/ fill out a contact form)

- clickbank.com: middleman for vendors - affiliates
- affiliates paid by clickthrough - vendors charged a fee

[Thomas SIGCOMM '11]

State of the art

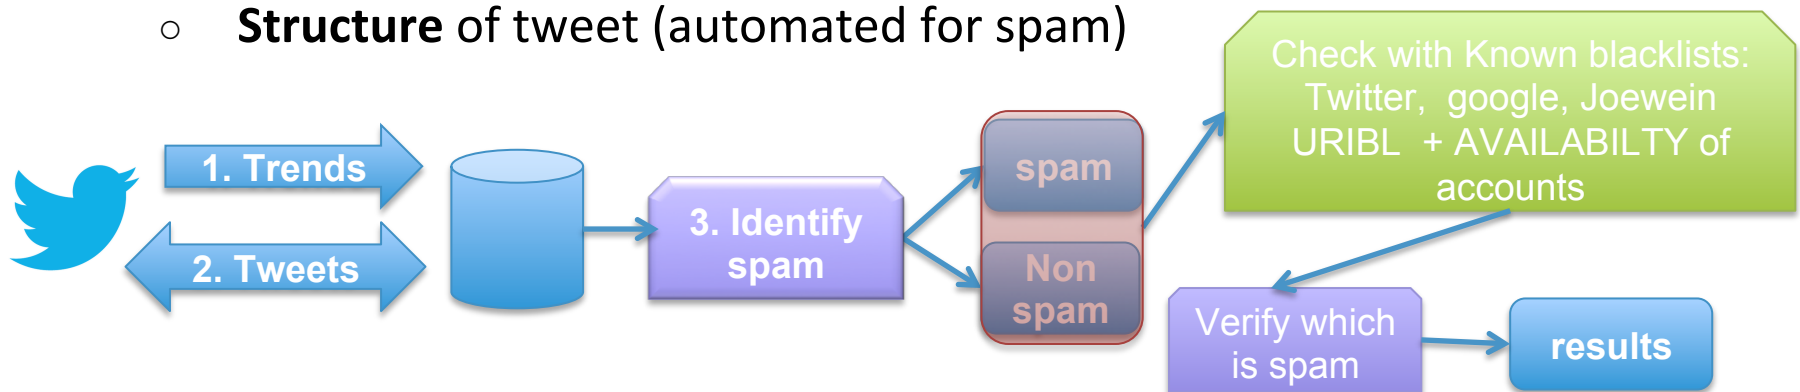
- Redirection chains used by spammers using search engine optimization techniques to increase traffic [Wang WWW '07]
- Measure the success of Twitter spam by analyzing click through statistics for spam URLs [Chris Grier, CCS'10]
 - Count of frequent terms & classify to spam categories

State of the art

- Detection of tweets containing spam instead of detecting spam accounts [Juan Martinez, Expert Systems with Applications 40 (2013)]
- Machine learning techniques used:
 - Classify spam tweets [K. Lee SIGIR '10]
 - Apply ML methods to automatically distinguish spam accounts from normal ones [Alex Hai Wang, SECRIPT '10]

Our proposal

- **Retrieve** today's fresh **trending hashtags** from Twitter
- **Collect corresponding tweets**
- **Identification of spam:**
 - Find **tweets** with **common text** but **different trending hashtags**
 - **Count #trends** used for each account or found in message in different time periods
 - Which trending topics are **more likely to be exploited?**
 - **Structure** of tweet (automated for spam)



Our proposal

- **Identification** of spam:
 - **Metadata** of users:
 - Investigate **spam patterns** for:
 - Language,
 - Geographic location,
 - Favorites
 - #Re-tweets, #URLs

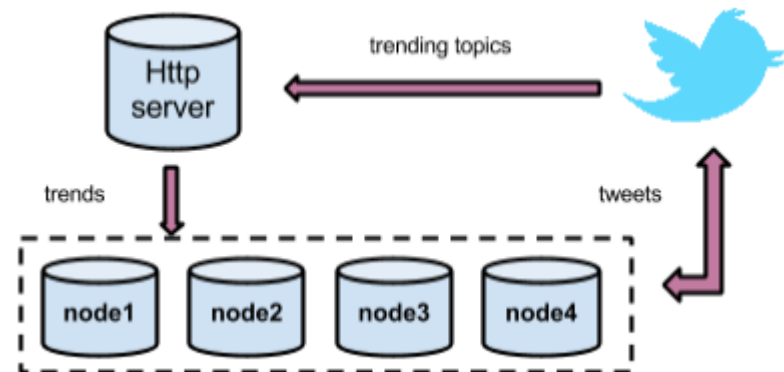
Our proposal

- **Identification of spam:**
 - Number of friends (huge #following, small #followers)
 - Density of the friends network (non spammers more strong connected- friend of friend probably connected)
 - Repetition of tweets
 - Follow/Unfollow technique
- **Verification of spam:**
 - Check URLs in tweets if exists in known blacklists
 - Check availability of accounts, spammers are suspended
 - Potential correlation with other tools

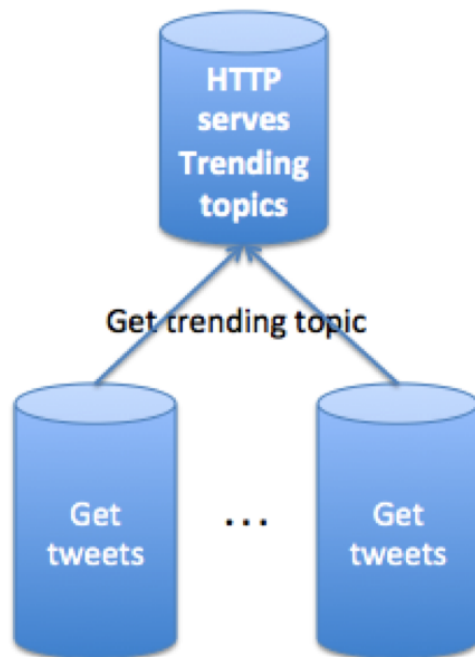
A decentralized trending tweets collector

- Components:
 - Simple http server
 - for each POST request returns a trend

```
#> curl -X POST 10.10.10.10:8080  
JustinBieber
```
 - SpamTweet script
 - i. n nodes connect to server & receive trends
 - ii. Downloads the relevant tweets
 - Identification of spam
 - Verification of technique



A decentralized trending tweets collector



- Runs on planet-lab: <http://www.planet-lab.org/>
- Can scale up easily
- Extensive use of python generators

- Data collected (so far):
- #Trends > 20.000
- #Days ≈ 100
- #Tweets > 150.000.000

The spam identification process is an idea under development ...

References

- Chris Grier et al. @spam: the underground on 140 characters or less. Proc 17th ACM conf C&C security. p 27-37
<http://www.icir.org/vern/papers/ccs2010-twitter-spam.pdf>
- Amleshwaram et al. CATS: Characterizing automation of Twitter spammers. Communication Systems and Networks (COMSNETS), 2013 Fifth International Conference.
<http://cesg.tamu.edu/wp-content/uploads/2012/02/comsnets2013-2.pdf>
- Thomas et al. Suspended Accounts in Retrospect: An Analysis of Twitter Spam. Proc. 2011 ACM SIGCOMM conf. on Internet measurement. p. 243-258
<http://www.icir.org/vern/papers/twitter-susp-accounts.imc2011.pdf>
- Juan Martinez-Romo et al. Detecting malicious tweets in trending topics using a statistical analysis of language. Expert Systems with Applications 40 (2013) 2992–3000.
<http://nlp.uned.es/~juaner/papers/jmromo-eswa-2013.pdf>

References

- Gianluca Stringhini, Christopher Kruegel, Giovanni Vigna, Detecting spammers on social networks. ACSAC '10 Proceedings of the 26th Annual Computer Security Applications Conference <http://dl.acm.org/citation.cfm?id=1920263>
- Y. Wang, M. Ma, Y. Niu, and H. Chen. Spam double-funnel: Connecting web spammers with advertisers. In Proceedings of the International World Wide Web Conference, pages 291–300, 2007.
- Alex Hai Wang, DON'T FOLLOW ME Spam Detection in Twitter College of Information Sciences and Technology, The Pennsylvania State University, PA 18512, Dunmore, U.S.A. hwang@psu.edu
- K. Lee, J. Caverlee, and S. Webb. Uncovering social spammers: Social honeypots + machine learning. In Proceeding of the SIGIR conference on Research and Development in Information Retrieval, pages 435–442, 2010.