

Measuring Service Availability Using Twitter

Marti Motoyama, Brendan Meeder, Kirill Levchenko,
Stefan Savage and Geoffrey M. Voelker

Introduction

- OSN graph properties widely studied
- More to OSNs than the “network”?
 - Large amount of information being disseminated
 - Real-time updates often reflect real events

Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors

Takeshi Sakaki
The University of Tokyo
Yayoi 2-11-16, Bunkyo-ku
Tokyo, Japan
sakaki@biz-model.t.u-
tokyo.ac.jp

Makoto Okazaki
The University of Tokyo
Yayoi 2-11-16, Bunkyo-ku
Tokyo, Japan
m_okazaki@biz-
model.t.u-tokyo.ac.jp

Yutaka Matsuo
The University of Tokyo
Yayoi 2-11-16, Bunkyo-ku
Tokyo, Japan
matsuo@biz-model.t.u-
tokyo.ac.jp

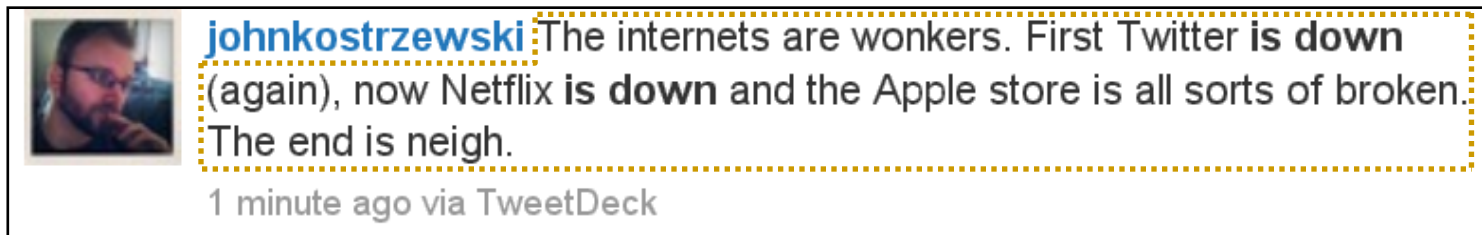
OSNs = HUMAN Sensor Networks

Introduction

The Twitter logo, consisting of the word "twitter" in a light blue, lowercase, sans-serif font, enclosed in a thin black rectangular border.

a real-time microblogging service

- Users post 140 character updates (**Tweets**)



- Twitter statistics:
 - Over **75 million** users and counting
 - Over **30 million** Tweets posted per day

Introduction

- Goal: Assess service availability using Twitter
- Motivation for looking at availability:
 - Movement towards cloud-hosted services

NETWORKWORLD

[News](#) | [Blogs & Columns](#) | [Subscriptions](#) | [Videos](#) | [Events](#)

[Security](#)

[LANs & WANs](#)

[VoIP](#)

[Infrastructure Mgmt](#)

[Wireless](#)

[Software](#)

[Data Center](#)

[Cloud Computing](#) | [Desktop](#) | [Disaster Recovery](#) | [Green IT](#) | [NAS](#) | [SAN](#) | [Server](#) | [Storage Management](#) |

Rackspace to issue as much as **\$3.5M** in customer credits **after outage**

Introduction

- OSNs offer a number of advantages:
 - Varied set of vantage points



coreyann i get to the shipping screen on AT&T and they don't have a button to click, **Apple is down** :(

3 days ago via webfrom North Canton, OH

- No need to specify services a priori
 - Observe correlated failures
- Recall: Great Gmail Outage of Sept. 1st, 2009



Server Error

The server encountered a temporary error and could not complete your request.

Please try again in 30 seconds.

Introduction

The image shows a screenshot of a TechCrunch article. The article title is "Gmail Hitting Some Turbulence" by Robin Wauters, dated Sep 1, 2009. It has 135 comments, 0 buzz, and 34 retweets. The article content shows a Gmail error message: "Temporary Error (500)". A yellow callout box is overlaid on the error message, containing the text: "I tried to log on to Gmail this morning... anyone else seeing this?". Below the error message, the text reads: "I tried to log on to Gmail this morning and this is all I've been getting for the last hour or so. Anyone else seeing this?".

HOME

TechCrunch

Xeon
Powerful.
Intelligent.

The new IBM eX5 enterprise systems
with Intel® Xeon® Processor 7500 Series

About Advertise Archives Contact Events Jobs Trends

Gmail Hitting Some Turbulence

by Robin Wauters on Sep 1, 2009


135 Comments Like Buzz 0 34 retweet

Temporary Error (500)

I tried to log on to Gmail this morning... anyone else seeing this?

I tried to log on to Gmail this morning and this is all I've been getting for the last hour or so. Anyone else seeing this?

Introduction



[dan325013](#): Not only was **Gmail** down for like a few hours earlier today, but it's still having issues running wtf Google [<http://twitpic.com/g4xy2>]

Sep 1, 2009 10:48 PM GMT · from *DestroyTwitter* · [Reply](#) · [View Tweet](#)



[newsreporter2](#): Integrated **Gmail** Transforms Your Inbox into a Website Hub
[http://](#)

Sep 1,

**Gmail goes down,
users cry to twitter**



[shmuxworld](#)

Sep 1, 2009 10:48 PM GMT · from *TwitterFox* · [Reply](#) · [View Tweet](#) · [Show Conversation](#)



[aivanet](#): NEWS: **Gmail** goes down, users cry to twitter <http://bit.ly/4GhUyg> ([expand](#))

Sep 1, 2009 10:48 PM GMT · from *twitterfeed* · [Reply](#) · [View Tweet](#)



[Jane1822](#): Make \$100+ a Day on Twitter While Sleeping. Find out here!
<http://tinyurl.com/kpfqdg> ([expand](#)) Snow Leopard DO WHILE **GMAIL**



Sep 1, 2009 10:48 PM GMT · from *web* · [Reply](#) · [View Tweet](#)

Outline

- Introduction
- Data Collection
- Detecting Outage Tweets
- Raising Alarms
- Evaluation
 - Known Events
 - Unknown Events
- Summary

Data Collection

- Methodology:



- Data Set:


- 2.8 Billion Tweets
 - Close to **800 GB** of content
- Tweets span 3 years

Detecting Outage Tweets

- Topic detection intuition:
 - Labeled 878 Tweets from 4 outages:
 - Gmail (02/24/09), Hotmail(03/12/09), PayPal (08/03/09), Bing (12/03/09)
 - Top Bi-gram:
 - **"is down"** (2.4%)
 - Top Hash Tag:
 - **"#fail"**(8.2%)

Detecting Outage Tweets

- Predicate Heuristics:
 - Check whether entity X is down:
 - IsDown(X)



ryangregg I'd really like to pre-order a new iPhone. Apparently that is impossible. AT&T is down, **Apple is down**, and AppleStore is crashing

about 5 hours ago via TweetDeck

- Fail(X)



dmcohen RT **@bayou**: It's easier to find a new iPhone 4 deserted in a bar than it is to order one through their website.
#fail #apple

9 minutes ago via web

Detecting Outage Tweets

Bing (12/3/09)		Gmail (9/1/09)		Hotmail (3/12/09)		Flickr (10/20/09)		PayPal (8/3/09)	
Phrase	Count	Phrase	Count	Phrase	Count	Phrase	Count	Phrase	Count
maps beta	702	is down	12,836	is down	61	is down	261	and its	318
bing bing	678	is back	2,774	server is	24	photo	227	payfail	190
is down	445	being down	1,630	server too	23	is	98	is down	81
the new	288	down	998	is	14	down battling	83	suffers from	80
with revamped	278	was down	994	is back	14	is back	77	shut down	66
<i>Other</i>	13,934	<i>Other</i>	171,772	<i>Other</i>	1,570	<i>Other</i>	8,235	<i>Other</i>	4,707
Total	16,325	Total	191,004	Total	1,706	Total	8,981	Total	5,442

- **IsDown(X)** provides subject detection
 - Looked at 2 words surrounding entity during 5 service outages
 - “is down” in top 5 across all outages

Raising Alarms

■ Expect noise:



[muddyxconverse](#) youtube **is down!**):

1 minute ago via web

1. No outage is actually occurring
 - Use Exponentially Weighted Moving Average (EWMA)




[katjolufan](#): so Gatorade drops **Tiger**...talk about kicking a **man when he is down**

Twitter - Feb 26, 2010

2. Subject not an internet service
 - Check for IsDown and Fail occurring in some time window

Raising Alarms

- High Level Methodology:

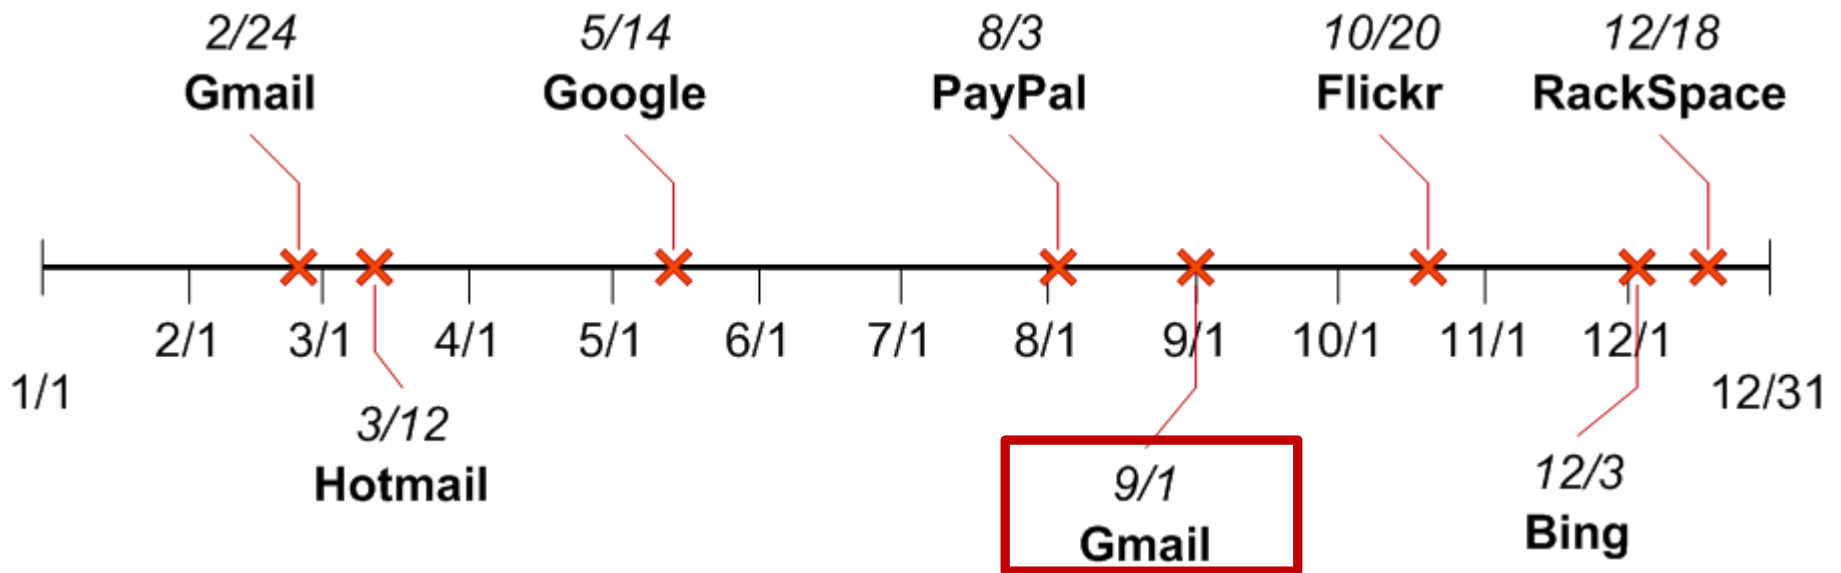


- Compute on a per entity basis:
 - EWMA on IsDown count
 - Smoothed variance using EWMA and current count
 - Threshold using EWMA and variance
- Check for consecutive threshold violations
- Optionally: check for Fail predicate

Choosing Parameters

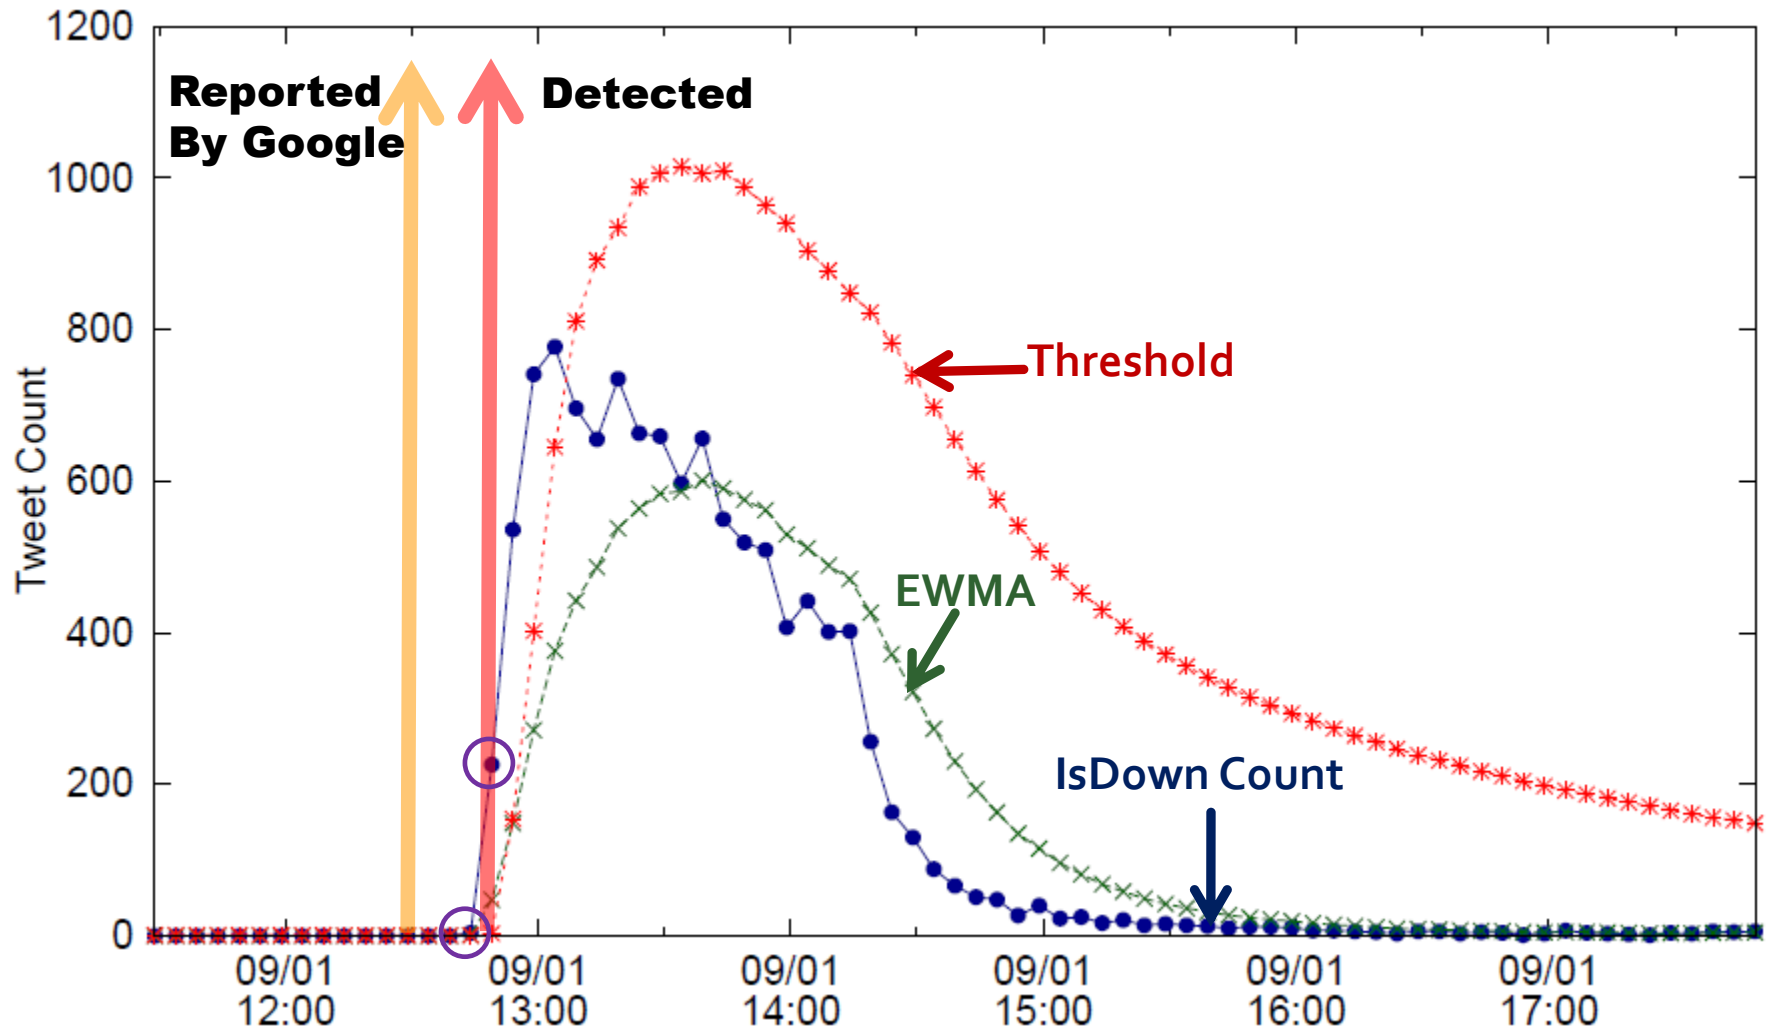
- Creating validation set:
 - Searched/checked maintenance blogs
 - Flickr, Hotmail, Ning, LiveJournal, PayPal, Tmobile
 - Found 45 outage events
- Using validation set:
 - Computed F-Scores for various parameter combinations and chose best
 - Alarm if threshold violated for 2 consecutive bins

Evaluation: Known Events



- Picked 8 well-known events
- Ran detection methodology

Gmail: Sept 1st, 2009



Known Outage Summary

- Good News:
 - Detected all 8 events
 - Also detected using Fail heuristic
- Bad News:
 - Time to detect varies (10-50 min)
 - Delay time increases using Fail heuristic
 - Possible delay causes:
 - News reports imprecise?
 - Better outage tweet detection?
 - At 12:39 pm: anybody else having problems getting on gmail?

Unknown Outages

- Ran analysis on entire corpus
 - 1+ million tweets expressing IsDown/Fail
- Without checking for Fail predicate
 - 5,358 “outages” spread over 1,556 entities
 - However, many false positive entities:

attendance

demand

pressure

tourism

usage

crime

visibility

who

spending

sun

mood

etc...

Unknown Outages

- **Solution:** Combine with Fail predicate
 - Heuristic: Fail within 30 min. of signal
 - Produces 894 outages, 245 entities



Is12styler

Hahahahahahahahahahahahaha

#Spain #fail

about 14 hours ago via HTC Peep

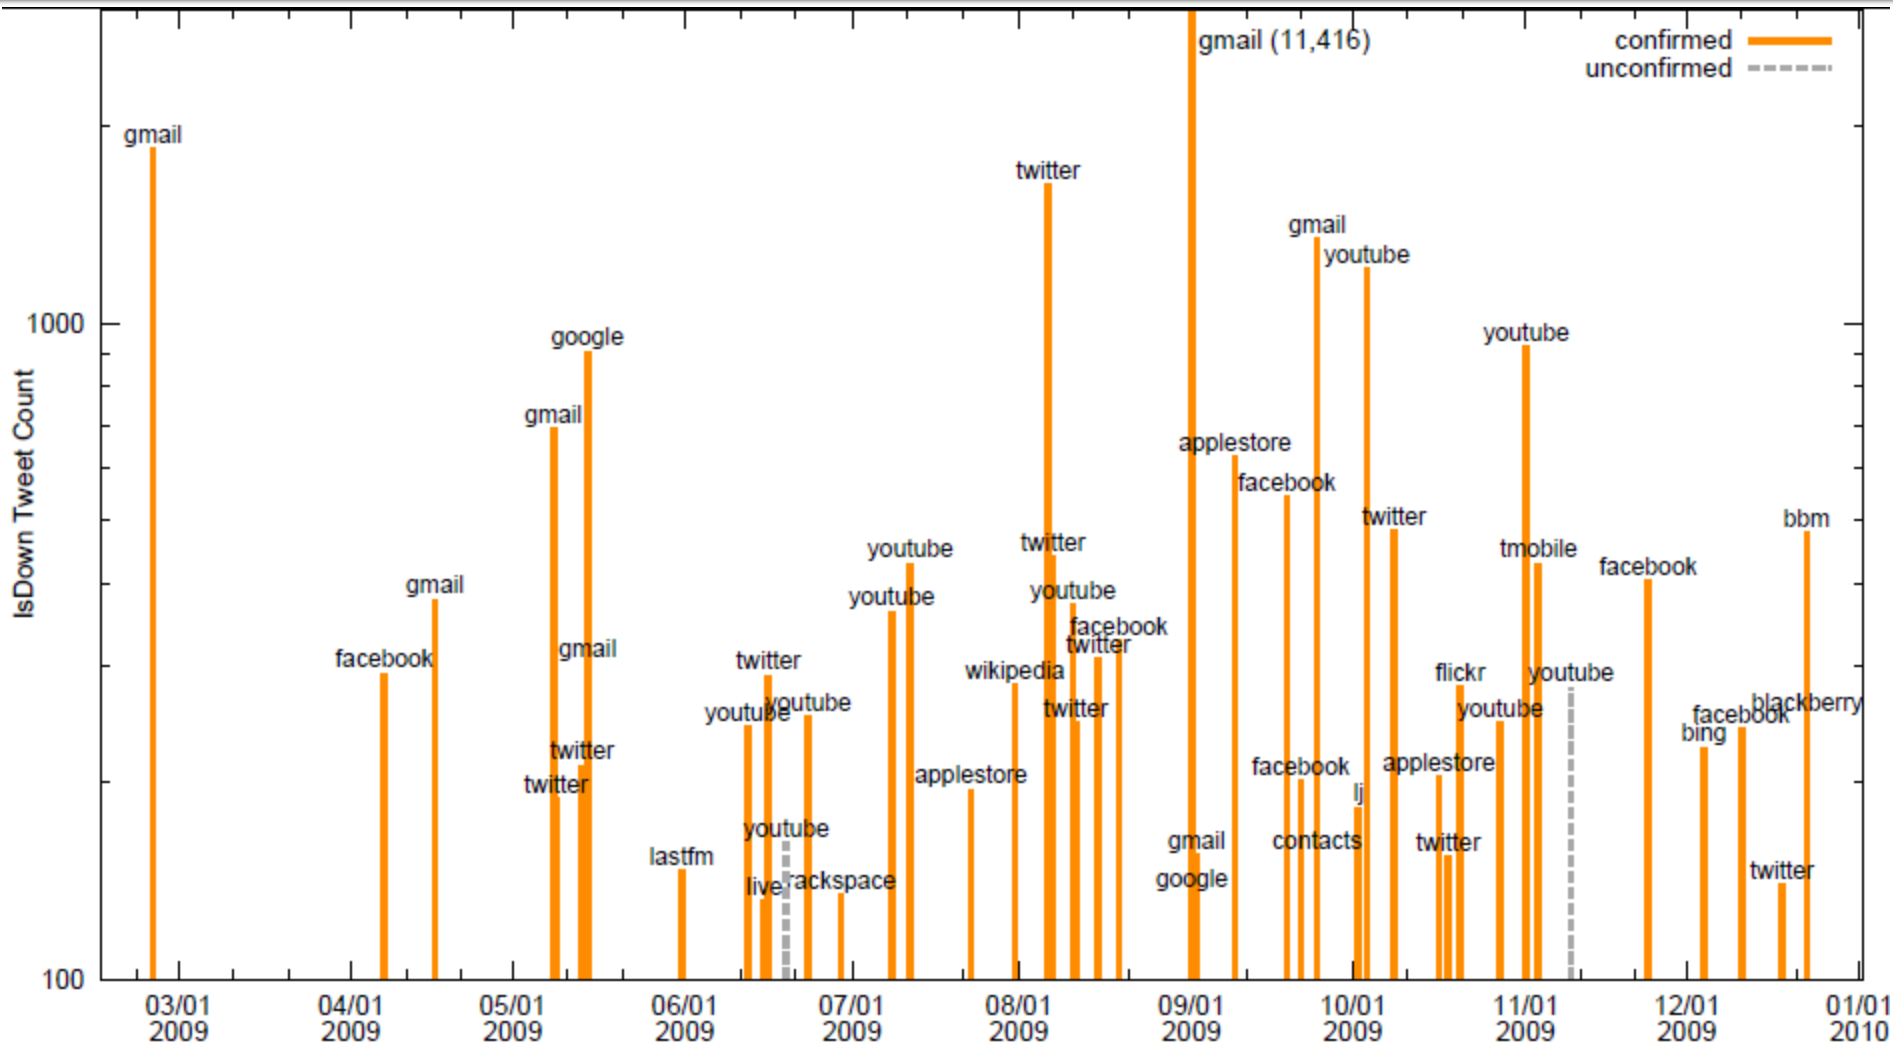


NeoHippie Spain is down 0-1. WTF?!?!?!?

about 14 hours ago via twirl

- Inspection of 245 entities reveals:
 - 59 false positive entities
 - Heuristics not robust to sporting events
 - Examples: USC, Liverpool, Federer, etc

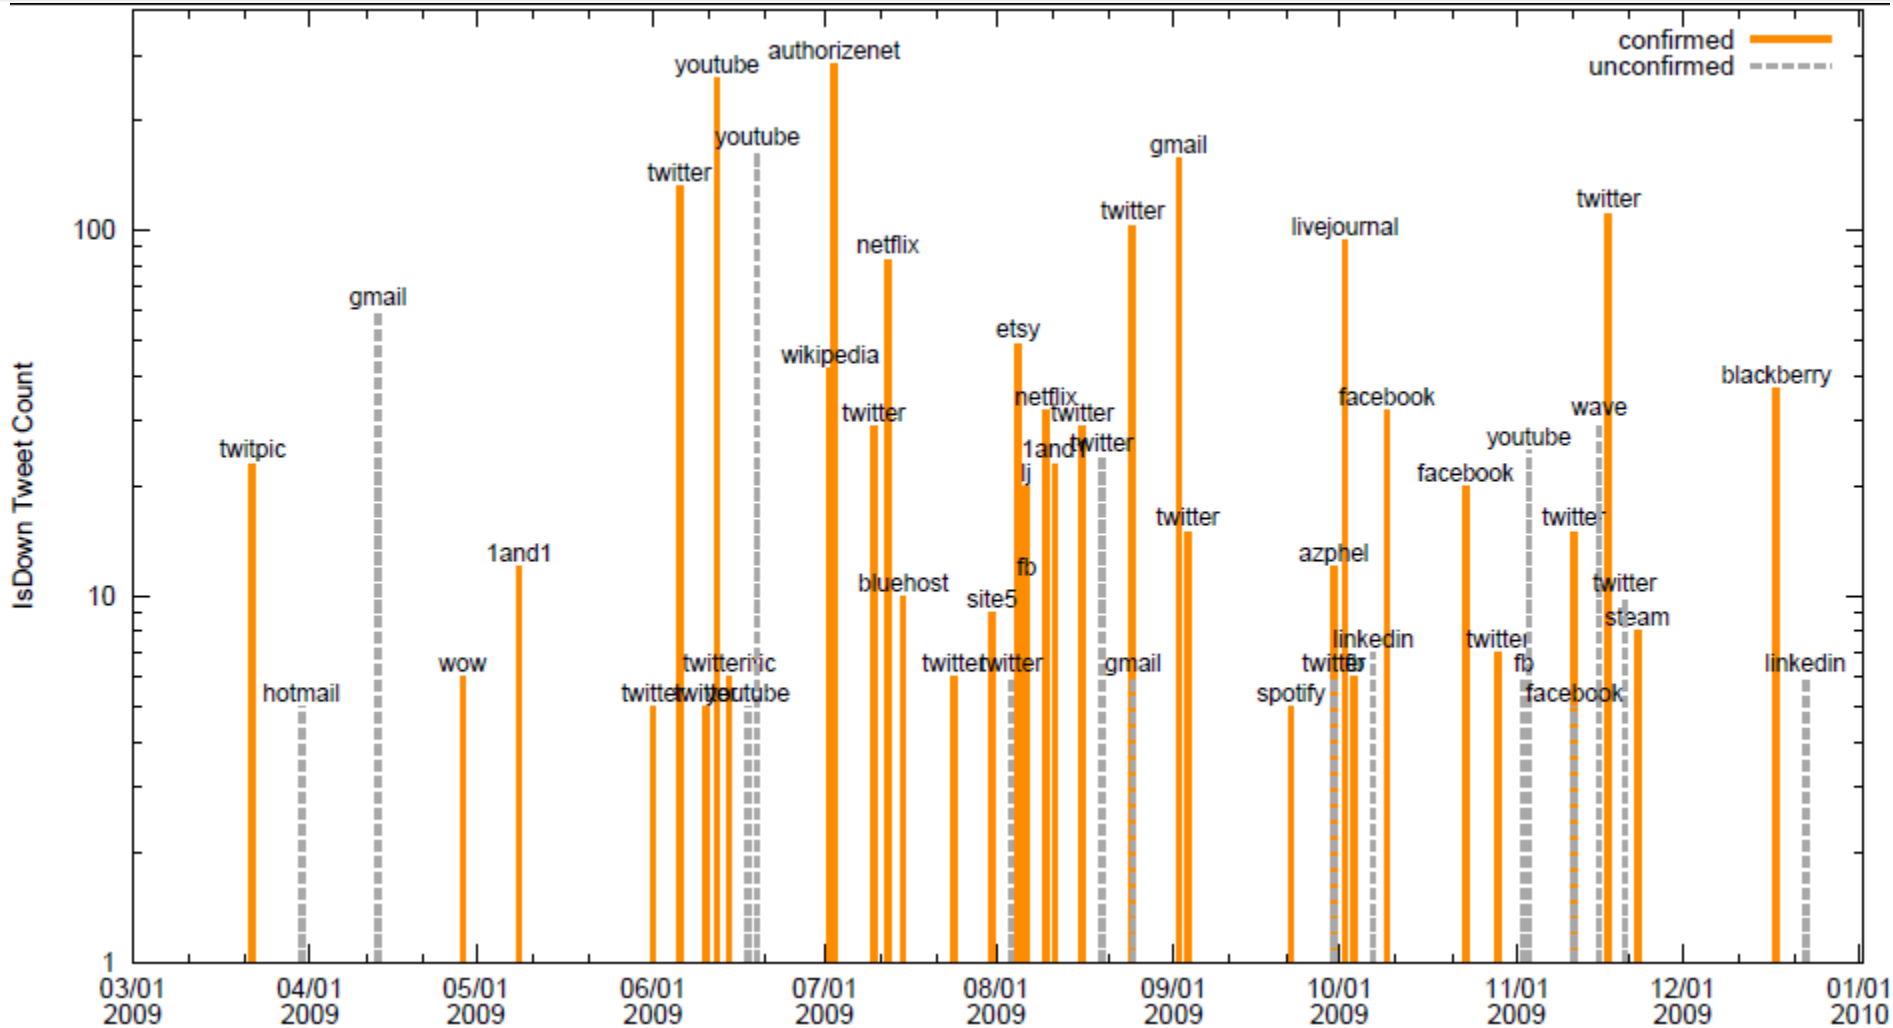
Top 50 Outages



Top 50 Outages

- 48 confirmed:
 - YouTube top with 11
 - Nine confirmed, two plausible
- Nine Twitter service disruptions?
 - Errors tend to be transient
 - Third party applications retry posts:
 - **Twitter is down once again :(#fail
#TwitterIsDown #TwitterFail**
-via TwitterFeed

Random 50 Outages



Random 50 Outages


- 35 confirmed (70%)
 - Span a variety of services
 - Azphel, WoW, Authorize.net, Netflix
- Unconfirmed:
 - At least 3 look plausible:
 - YouTube on 6/19, Gmail on 4/13, Google Wave on 11/16
 - Wave Example:
 - **wave is down, though I doubt if people noticed! RT @annkur: Twitter shows a whale .Google wave shows the entire Ocean when down :P**

Summary

- Explored application to service outages
 - Simple methods identify important events
- Future Work:
 - Improve outage tweet detection
 - Explore alternatives to EWMA
 - Monitor availability in real time
- OSNs: multipurpose sensor networks

■ Any questions?

THIS MEANS WHEN THE SEISMIC WAVES ARE ABOUT 100 KM OUT, THEY BEGIN TO BE OVERTAKEN BY THE WAVES OF POSTS ABOUT THEM.



PEOPLE OUTSIDE THIS RADIUS MAY GET WORD OF THE QUAKE VIA TWITTER, IRC, OR SMS BEFORE THE SHAKING HITS.

WHOA!
EARTHQUAKE!



SADLY, A TWITTERER'S FIRST INSTINCT IS NOT TO FIND SHELTER.

RT @ROB1163 HUGE EARTHQUAKE HERE!

