## Large Graph Mining Patterns, Tools and Cascade analysis

Christos Faloutsos
CMU

## Thank you!

- Xuewen Chen
- Dennis Schwartz


## Roadmap

$\Rightarrow$ - Introduction - Motivation

- Why 'big data'
- Why (big) graphs?
- Problem\#1: Patterns in graphs
- Problem\#2: Tools
- Problem\#3: Scalability
- Conclusions


## Why 'big data'

- Why?
- What is the problem definition?
- What are the major research challenges?


## Main message: Big data: often $>$ experts

- 'Super Crunchers’ Why Thinking-By-Numbers is the New Way To Be Smart by Ian Ayres, 2008
- Google won the machine translation competition 2005
- http://www.itl.nist.gov/iad/mig//tests/mt/2005/doc/ mt05eval_official results release 20050801 v3.html


## Problem definition - big picture



# Tera/Peta-byte data 

Analytics
Insights, outliers

## Problem definition - big picture



# Tera/Peta-byte data 

Analytics
Insights, outliers

Main emphasis in this talk

## Problem definition - big picture


(my personal) rules of thumb: if data

- fits in memory -> R, matlab, scipy
- single disk -> RDBMS (sqlite3, mysql, postgres)
- multiple (<100-1000) disks: parallel RDBMS (Vertica, TeraData)
- multiple (>1000) disks: hadoop, pig


## (Free) Resource for graphs

Open source system for mining huge graphs:

PEGASUS project (PEta GrAph mining
System)

- www.cs.cmu.edu/~pegasus

- Apache license for $\mathrm{s} / \mathrm{w}$
- code and papers


## Research challenges

- The usual ones from data mining
- Data cleansing
- Feature engineering


## PLUS

- Scalability ( $<\mathrm{O}\left(\mathrm{N}^{* *}\right.$ 2) )
- Real data *disobey* textbook assumptions (uniformity, independence, Gaussian, Poisson) with huge performance implications


## Roadmap

- Introduction - Motivation
- Why 'big data’
- Why (big) graphs?
- Problem\#1: Patterns in graphs
- Problem\#2: Tools
- Problem\#3: Scalability
- Conclusions


## Graphs - why should we care?



Food Web
[Martinez '91]

## >\$10B revenue

>0.5B users


Internet Map [lumeta.com]

## Graphs - why should we care?

- IR: bi-partite graphs (doc-terms)


- web: hyper-text graph
- ... and more:


## Graphs - why should we care?

- 'viral' marketing
- web-log ('blog') news propagation
- computer network security: email/IP traffic and anomaly detection
- Subject-verb-object -> graph
- Many-to-many db relationship -> graph


## Outline

- Introduction - Motivation

Problem\#1: Patterns in graphs


- Static graphs
- Weighted graphs
- Time evolving graphs
- Problem\#2: Tools
- Problem\#3: Scalability
- Conclusions


## Problem \#1 - network and graph mining

- What does the Internet look like?
- What does FaceBook look like?
- What is 'normal'/‘abnormal'?
- which patterns/laws hold?


## Problem \#1 - network and graph mining

- What does the Internet look like?
- What does FaceBook look like?
- What is 'normal'/‘abnormal'?
- which patterns/laws hold?
- To spot anomalies (rarities), we have to discover patterns


## Problem \#1 - network and graph mining

- What does the Internet look like?
- What does FaceBook look like?
- What is 'normal'/‘abnormal'?
- which patterns/laws hold?
- To spot anomalies (rarities), we have to discover patterns
- Large datasets reveal patterns/anomalies that may be invisible otherwise...


## Graph mining

- Are real graphs random?


## Laws and patterns

- Are real graphs random?
- A: NO!!
- Diameter
- in- and out- degree distributions
- other (surprising) patterns
- So, let's look at the data


## Solution\# S. 1

- Power law in the degree distribution [SIGCOMM99]


## internet domains



## Solution\# S. 1

- Power law in the degree distribution [SIGCOMM99]


## internet domains



## Solution\# S.2: Eigen Exponent $E$

Eigenvalue


Exponent $=$ slope
$E=-0.48$

May 2001

Rank of decreasing eigenvalue

- A2: power law in the eigenvalues of the adjacency matrix
C. Faloutsos (CMU)


## Solution\# S.2: Eigen Exponent $E$

Eigenvalue


> Exponent = slope

$$
E=-0.48
$$

May 2001

Rank of decreasing eigenvalue

- [Mihail, Papadimitriou '02]: slope is $1 / 2$ of rank exponent
Wayne State, Feb. 2013
C. Faloutsos (CMU)


## But:

## How about graphs from other domains?

## More power laws:

- web hit counts [w/ A. Montgomery]



## epinions.com



## And numerous more

- \# of sexual contacts
- Income [Pareto] -'80-20 distribution'
- Duration of downloads [Bestavros+]
- Duration of UNIX jobs ('mice and elephants’)
- Size of files of a user
- 'Black swans’


## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Static graphs
- degree, diameter, eigen,
- triangles
- cliques
- Weighted graphs
- Time evolving graphs
- Problem\#2: Tools


## Solution\# S.3: Triangle 'Laws’



- Real social networks have a lot of triangles


## Solution\# S.3: Triangle 'Laws'



- Real social networks have a lot of triangles
- Friends of friends are friends
- Any patterns?


## Triangle Law: \#S. 3 [Tsourakakis ICDM 2008]




Wayne Sate, Feb. 2013


ASN

X-axis: \# of participating triangles
Y: count ( $\sim$ pdf)
$10^{5}$ is (CMU)

## Triangle Law: \#S. 3 [Tsourakakis ICDM 2008]



Epinions


X-axis: \# of participating triangles
Y: count ( $\sim$ pdf)

## Triangle Law: \#S. 4 [Tsourakakis ICDM 2008]




Wayne State, Feb. 2013 Degree


X-axis: degree
Y-axis: mean \# triangles
$n$ friends -> $\sim n^{1.6}$ triangles

## Triangle counting for large graphs?



Anomalous nodes in Twitter( $\sim 3$ billion edges)
[U Kang, Brendan Meeder, +, PAKDD'11]
C. Faloutsos (CMU)

38

## Triangle counting for large graphs?



Anomalous nodes in Twitter( $\sim 3$ billion edges)
[U Kang, Brendan Meeder, +, PAKDD'11]

## Triangle counting for large graphs?



Anomalous nodes in Twitter( $\sim 3$ billion edges)
[U Kang, Brendan Meeder, +, PAKDD'11]
C. Faloutsos (CMU)

## Triangle counting for large graphs?



Anomalous nodes in Twitter( $\sim 3$ billion edges)
[U Kang, Brendan Meeder, +, PAKDD'11]
C. Faloutsos (CMU)

## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Static graphs
- degree, diameter, eigen,
- triangles
- cliques
- Weighted graphs
- Time evolving graphs
- Problem\#2: Tools


## Observations on weighted graphs?

- A: yes - even more 'laws'!

M. McGlohon, L. Akoglu, and C. Faloutsos Weighted Graphs and Disconnected Components: Patterns and a Generator. SIG-KDD 2008


## Observation W.1: Fortification

Q: How do the weights of nodes relate to degree?

## Observation W.1: Fortification

## More donors, more \$ ?



Wayne State, Feb. 2013

## Observation W.1: fortification: Snapshot Power Law

- Weight: super-linear on in-degree
- exponent 'iw': $1.01<\mathrm{iw}<1.26$


## More donors, even more \$



In-weights

C. Faloutsos (CMU)

Orgs-Candidates
e.g. John Kerry, \$10M received, from 1K donors

## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Static graphs
- Weighted graphs
- Time evolving graphs
- Problem\#2: Tools


## Problem: Time evolution

- with Jure Leskovec (CMU -> Stanford)

- and Jon Kleinberg (Cornell sabb. @ CMU)



## T. 1 Evolution of the Diameter

- Prior work on Power Law graphs hints at slowly growing diameter:
- diameter ~ $\mathrm{O}(\log \mathrm{N})$
- diameter $\sim \mathrm{O}(\log \log \mathrm{N})$

- What is happening in real data?


## T. 1 Evolution of the Diameter

- Prior work on Power Law graphs hints at slowly growing diameter:
- diameter $\sim($ ( $\mathrm{L} \alpha \mathrm{I})$
- diameter $\sim \mathrm{O}($ rug $\log \mathrm{N})$

- What is happening in real data?
- Diameter shrinks over time


## T. 1 Diameter - "Patents"

- Patent citation network
- 25 years of data
-@1999
- 2.9 M nodes
- 16.5 M edges

C. Faloutsos (CMU)


## T. 2 Temporal Evolution of the Graphs

- $\mathrm{N}(\mathrm{t})$... nodes at time t
- $\mathrm{E}(\mathrm{t})$... edges at time t
- Suppose that

$$
\mathrm{N}(\mathrm{t}+1)=2 * \mathrm{~N}(\mathrm{t})
$$

- Q: what is your guess for

$$
\mathrm{E}(\mathrm{t}+1)=? 2 * \mathrm{E}(\mathrm{t})
$$

## T. 2 Temporal Evolution of the Graphs

- $\mathrm{N}(\mathrm{t})$... nodes at time t
- $\mathrm{E}(\mathrm{t})$... edges at time t
- Suppose that
$\mathrm{N}(\mathrm{t}+1)=2 * \mathrm{~N}(\mathrm{t})$
- Q: what is your guess for
$\mathrm{E}(\mathrm{t}+1)=?$ ? $\mathrm{E}(\mathrm{t})$
- A: over-doubled!
- But obeying the "Densification Power Law"


## T. 2 Densification - Patent Citations

- Citations among patents granted
- @1999
- 2.9 M nodes
- 16.5 M edges
- Each year is a datapoint



## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Static graphs
- Weighted graphs
- Time evolving graphs
- Problem\#2: Tools


## T. 3 : popularity over time

\# in links


Post popularity drops-off - exponentially?


## T. 3 : popularity over time

\# in links
(log)

days after post (log)

Post popularity drops-off - expon $e^{\dagger}$ ally? POWER LAW!
Exponent?

## T. 3 : popularity over time

\# in links
(log)

days after post (log)

Post popularity drops-off - expor ent ally? POWER LAW!
Exponent? -1.6

- close to -1.5: Barabasi's stack model
- and like the zero-crossings of a random walk

Wayne State, Feb. 2013
C. Faloutsos (CMU)

## -1.5 slope

J. G. Oliveira \& A.-L. Barabási Human Dynamics: The Correspondence Patterns of Darwin and Einstein. Nature 437, 1251 (2005) . [PDF]


Figure 1 |The correspondence patterns of Darwin and Einstein.

## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Problem\#2: Tools
- Belief Propagation
- Tensors
- Spike analysis
- Problem\#3: Scalability
- Conclusions


## E-bay Fraud detection



## w/ Polo Chau \& Shashank Pandit, CMU [www'07]



## E-bay Fraud detection



## E-bay Fraud detection



## E-bay Fraud detection - NetProbe



Wayne State, Feb. 2013
C. Faloutsos (CMU)

## E-bay Fraud detection - NetProbe

## Compatibility matrix

|  | $\mathbf{F}$ | $\mathbf{A}$ | $\mathbf{H}$ |
| :--- | :--- | :--- | :--- |
| $\mathbf{F}$ |  | $99 \%$ |  |
| $\mathbf{A}$ | $99 \%$ |  |  |
| $\mathbf{H}$ |  | $49 \%$ | $49 \%$ |$\quad$ beterophily



## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Problem\#2: Tools
- Belief Propagation
- Tensors
- Spike analysis
- Problem\#3: Scalability
- Conclusions


# GigaTensor: Scaling Tensor Analysis Up By 100 Times Algorithms and Discoveries 

U Evangelos Abhay Christos<br>Kang Papalexakis Harpale Faloutsos

## KDD'12

## Background: Tensor

- Tensors (=multi-dimensional arrays) are everywhere
- Hyperlinks \&anchor text [Kolda+,05]



## Background: Tensor

- Tensors (=multi-dimensional arrays) are everywhere
- Sensor stream (time, location, type)
- Predicates (subject, verb, object) in knowledge base


Wayne State, Feb. 2013
(48M) verbs subjects (26M)


NELL (Never Ending Language Learner) data Nonzeros $=144 \mathrm{M}$
C. Faloutsos (CMU)69

## Background: Tensor

- Tensors (=multi-dimensional arrays) are everywhere
- Sensor stream (time, location, type)
- Predicates (subject, verb, object) in knowledge base


Anomaly
Detection in
Computer
networks

IP-destination

## Problem Definition

- How to decompose a billion-scale tensor?
- Corresponds to SVD in 2D case



## Problem Definition

- Q1: Dominant concepts/topics?
$\square$ Q2: Find synonyms to a given noun phrase?
- (and how to scale up: |data|> RAM)
(48M) verbs


NELL (Never Ending
Language Learner) data
Nonzeros $=144 \mathrm{M}$

## Experiments

- GigaTensor solves $100 x$ larger problem



Number of
nonzero
= I / 50

## A1: Concept Discovery

## - Concept Discovery in Knowledge Base



Wayne State, Feb. 2013
C. Faloutsos (CMU)

75

## A1: Concept Discovery

| Noun <br> Phrase | Noun Phrase 2 | Context |
| :---: | :---: | :---: |
| Concept 1: interne file data | Web Protocol' protocol software suite | 'np1' 'stream' 'np2' 'np1' 'marketing' 'np2' <br> 'np1' 'dating' 'np2' |
| Concept 2: <br> credit <br> Credit <br> library | Credit Cards" <br> information <br> debt <br> number |  |
| Concept 3: <br> health <br> child <br> home | Health System provider providers system | 'np1' 'care' 'np2' 'np' 'insurance, 'np2' 'np1' 'service' ' $n$ n2' |

## A2: Synonym Discovery

## (Given) <br> Noun Phrase

| pollutants | dioxin, sulfur dioxide, <br> greenhouse gases, particulates, <br> nitrogen oxide, air pollutants, cholesterol |
| :--- | :--- |
| disabilities | infections, dizziness, <br> injuries, diseases, drowsiness, <br> stiffness, injuries |
| vodafone | verizon, comcast |

Christian history European history, American history, Islamic history, history
disbelief dismay, disgust, astonishment

## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Problem\#2: Tools
- Belief propagation
- Tensors
- Spike analysis
- Problem\#3: Scalability -PEGASUS
- Conclusions


## Rise and fall patterns in social media

- Meme (\# of mentions in blogs)
- short phrases Sourced from U.S. politics in 2008
"you can put lipstick on a pig"

"yes we can"

C. Faloutsósecheyrs)


## Rise and fall patterns in social media

- Can we find a unifying model, which includes these patterns?
- four classes on YouTube [Crane et al. '08]
- six classes on Meme [Yang et al. '11]



## Rise and fall patterns in social media

- Answer: YES!

- We can represent all patterns by single model


## In Matsubara+ SIGKDD 2012

## Main idea - SpikeM

- 1. Un-informed bloggers (uninformed about rumor)
- 2. External shock at time nb (e.g, breaking news)
- 3. Infection (word-of-mouth)


Time $\mathrm{n}=0$


Time $\mathrm{n}=\mathrm{n}_{\mathrm{b}}$


Time $n=n_{b}+1$

Infectiveness of a blog-post at age $n$ :
$\beta \quad$ - Strength of infection (quality of news)
$f(n)$ - Decay function

## Main idea - SpikeM



- 1. Un-informed bloggers (uninformed about rumor)
- 2. External shock at time nb (e.g, breaking news)
- 3. Infection (word-of-mouth)


Time $\mathrm{n}=0$


Time $\mathrm{n}=\mathrm{n}_{\mathrm{b}}$


Time $\mathrm{n}=\mathrm{n}_{\mathrm{b}}+1$

Infectiveness of a blog-post at age n :
$\beta \quad$ - Strength of infection (quality of news)
$f(n)$ - Decay function
$f(n)=\beta * n^{-1.5}$

## SpikeM - with periodicity

- Full equation of SpikeM

$$
\begin{array}{cc}
\Delta B(n+1)=\frac{p(n+1)}{\text { Periodicity }} \cdot\left[U(n) \cdot \sum_{t=n_{b}}^{n}(\Delta B(t)+S(t)) \cdot f(n+1-t)+\varepsilon\right] \\
\begin{array}{c}
\text { Bloggers change their } \\
\text { activity over time } \\
\text { (e.g., daily, weekly, } \\
\text { yearly) }
\end{array} & \text { noon } \\
\text { activity } \\
\text { Wayne State, Feb. 2013 }
\end{array}
$$

## Details

- Analysis - exponential rise and power-raw fall



## Details

- Analysis - exponential rise and power-raw fall




## Tail-part forecasts

- SpikeM can capture tail part




## "What-if" forecasting



## "What-if" forecasting



## SpikeM can forecast upcoming spikes

C. Faloutsos (CMU)

## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Problem\#2: Tools
- Problem\#3: Scalability -PEGASUS
- Diameter
- Connected components
- Conclusions


## Scalability



- Google: $>450,000$ processors in clusters of $\sim 2000$ processors each [Barroso, Dean, Hölzle, "Web Search for a Planet: The Google Cluster Architecture" IEEE Micro 2003]
- Yahoo: 5Pb of data [Fayyad, KDD'07]
- Problem: machine failures, on a daily basis
- How to parallelize data mining tasks, then?
- A: map/reduce - hadoop (open-source clone) http://hadoop.apache.org/


## Roadmap - Algorithms \& results

|  | Centralized | Hadoop/ <br> PEGASUS |
| :--- | :---: | :---: |
| Degree Distr. | old | old |
| Pagerank | old | old |
| Diameter/ANF | old | HERE |
| Conn. Comp | old | HERE |
| Triangles | done | HERE |
| Visualization | started |  |

## HADI for diameter estimation R.

- Radius Plots for Mining Tera-byte Scale Graphs U Kang, Charalampos Tsourakakis, Ana Paula Appel, Christos Faloutsos, Jure Leskovec, SDM'10
- Naively: diameter needs $\mathbf{O}\left(\mathbf{N}^{* *} \mathbf{2}\right)$ space and up to $\mathrm{O}\left(\mathrm{N}^{*} * 3\right)$ time - prohibitive ( $\mathrm{N} \sim 1 \mathrm{~B}$ )
- Our HADI: linear on E ( $\sim 10 \mathrm{~B}$ )
- Near-linear scalability wrt \# machines
- Several optimizations -> 5x faster


YahooWeb graph (120Gb, 1.4B nodes, 6.6 B edges)
- Largest publicly available graph ever studied.


YahooWeb graph (120Gb, 1.4B nodes, 6.6 B edges)

- Largest publicly available graph ever studied.


YahooWeb graph (120Gb, 1.4B nodes, 6.6 B edges)
-7 degrees of separation (!)
-Diameter: shrunk


YahooWeb graph (120Gb, 1.4B nodes, 6.6 B edges) Q: Shape?


YahooWeb graph (120Gb, 1.4B nodes, 6.6 B edges)

- effective diameter: surprisingly small.
- Multi-modality (?!)

Wayne State, Feb. 2013
C. Faloutsos (CMU)


Radius Plot of GCC of YahooWeb.


YahooWeb graph (120Gb, 1.4B nodes, 6.6 B edges)

- effective diameter: surprisingly small.
- Multi-modality: probably mixture of cores .

Wayne State, Feb. 2013

C. Faloutsos (CMU)


Conjecture:
EN
$\{D E$

$\sum \sum B R$

YahooWeb graph (120Gb, 1.4B nodes, 6.6 B edges)

- effective diameter: surprisingly small.
- Multi-modality: probably mixture of cores .

Wayne State, Feb. 2013

C. Faloutsos (CMU)

Conjecture:


YahooWeb graph (120Gb, 1.4B nodes, 6.6 B edges)

- effective diameter: surprisingly small.
- Multi-modality: probably mixture of cores .

Wayne State, Feb. 2013

C. Faloutsos (CMU)

## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Problem\#2: Tools
- Problem\#3: Scalability -PEGASUS
- Diameter
- Connected components
- Conclusions


# Generalized Iterated Matrix Vector Multiplication (GIMV) 

PEGASUS: A Peta-Scale Graph Mining System - Implementation and Observations. U Kang, Charalampos E. Tsourakakis, and Christos Faloutsos. (ICDM) 2009, Miami, Florida, USA. Best Application Paper (runner-up).

## Vector Multiplication (GIMV)

- PageRank
- proximity (RWR)
- Diameter
- Connected components
- (eigenvectors,
- Belief Prop.
- ...)

Matrix - vector Multiplication
(iterated)

## Example: GIM-V At Work

- Connected Components - 4 observations:



## Example: GIM-V At Work

- Connected Components



## Example: GIM-V At Work

- Connected Components



## Example: GIM-V At Work

- Connected Components



## Example: GIM-V At Work

- Connected Components



## Example: GIM-V At Work

- Connected Components



## GIM-V At Work

- Connected Components over Time
- LinkedIn: 7.5M nodes and 58M edges






## Stable tail slope after the gelling point

## Roadmap

- Introduction - Motivation
- Problem\#1: Patterns in graphs

- Problem\#2: Tools
- Problem\#3: Scalability
$\Rightarrow$ - Conclusions


## OVERALL CONCLUSIONS low level:

- Several new patterns (fortification, triangle-laws, conn. components, etc)
- New tools:
- belief propagation, gigaTensor, etc
- Scalability: PEGASUS / hadoop


## OVERALL CONCLUSIONS high level

- BIG DATA: Large datasets reveal patterns/ outliers that are invisible otherwise




## References

- Leman Akoglu, Christos Faloutsos: RTG: A Recursive Realistic Graph Generator Using Random Typing. ECML/PKDD (1) 2009: 13-28
- Deepayan Chakrabarti, Christos Faloutsos: Graph mining: Laws, generators, and algorithms. ACM Comput. Surv. 38(1): (2006)


## References

- Deepayan Chakrabarti, Yang Wang, Chenxi Wang, Jure Leskovec, Christos Faloutsos: Epidemic thresholds in real networks. ACM Trans. Inf. Syst. Secur. 10(4): (2008)
- Deepayan Chakrabarti, Jure Leskovec, Christos Faloutsos, Samuel Madden, Carlos Guestrin, Michalis Faloutsos: Information Survival Threshold in Sensor and P2P Networks. INFOCOM 2007: 1316-1324


## References

- Christos Faloutsos, Tamara G. Kolda, Jimeng Sun: Mining large graphs and streams using matrix and tensor tools. Tutorial, SIGMOD Conference 2007: 1174


## References

- T. G. Kolda and J. Sun. Scalable Tensor Decompositions for Multi-aspect Data Mining. In: ICDM 2008, pp. 363-372, December 2008.


## References

- Jure Leskovec, Jon Kleinberg and Christos Faloutsos Graphs over Time: Densification Laws, Shrinking Diameters and Possible Explanations, KDD 2005 (Best Research paper award).
- Jure Leskovec, Deepayan Chakrabarti, Jon M. Kleinberg, Christos Faloutsos: Realistic, Mathematically Tractable Graph Generation and Evolution, Using Kronecker Multiplication. PKDD 2005: 133-145


## References

- Yasuko Matsubara, Yasushi Sakurai, B. Aditya Prakash, Lei Li, Christos Faloutsos, "Rise and Fall Patterns of Information Diffusion: Model and Implications", KDD'12, pp. 6-14, Beijing, China, August 2012


## References

- Jimeng Sun, Yinglian Xie, Hui Zhang, Christos Faloutsos. Less is More: Compact Matrix Decomposition for Large Sparse Graphs, SDM, Minneapolis, Minnesota, Apr 2007.
- Jimeng Sun, Spiros Papadimitriou, Philip S. Yu, and Christos Faloutsos, GraphScope: Parameterfree Mining of Large Time-evolving Graphs ACM SIGKDD Conference, San Jose, CA, August 2007


## References

- Jimeng Sun, Dacheng Tao, Christos Faloutsos: Beyond streams and graphs: dynamic tensor analysis. KDD 2006: 374-383


## References

- Hanghang Tong, Christos Faloutsos, and Jia-Yu Pan, Fast Random Walk with Restart and Its Applications, ICDM 2006, Hong Kong.
- Hanghang Tong, Christos Faloutsos, Center-Piece Subgraphs: Problem Definition and Fast Solutions, KDD 2006, Philadelphia, PA


## References

- Hanghang Tong, Christos Faloutsos, Brian Gallagher, Tina Eliassi-Rad: Fast best-effort pattern matching in large attributed graphs. KDD 2007: 737-746


## Project info \& 'thanks'

Www.cs.cmu.edu/~pegasus


Thanks to: NSF IIS-0705359, IIS-0534205, CTA-INARC; Yahoo (M45), LLNL, IBM, SPRINT, Google, INTEL, HP, iLab

## Cast



Akoglu, Leman


Beutel, Alex


Chau, Polo


Kang, U


McGlohon, Mary


Prakash, Aditya


Papalexakis, Vagelis


Tong, Hanghang

## Take-home message



## Big data reveal insights that would be invisible otherwise (even to experts)

C. Faloutsos (CMU)

