Graph Mining: Laws, Generators and Tools

Christos Faloutsos
CMU

Thank you!

- Sharad Mehrotra

Outline

- **Problem definition / Motivation**
- Static & dynamic laws; generators
- Tools: CenterPiece graphs; Tensors
- Other projects (Virus propagation, e-bay fraud detection)
- Conclusions

Motivation

Data mining: ~ find patterns (rules, outliers)

- Problem#1: How do real graphs look like?
- Problem#2: How do they evolve?
- Problem#3: How to generate realistic graphs

TOOLS

- Problem#4: Who is the 'master-mind'?
- Problem#5: Track communities over time

Problem#1: Joint work with

Dr. Deepayan Chakrabarti
(CMU/Yahoo R.L.)

Graphs - why should we care?

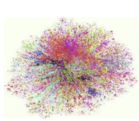

Internet Map
[lumeta.com]

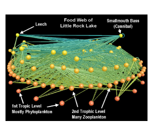

Food Web
[Martinez '91]

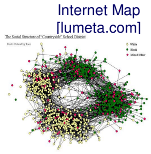

Friendship Network
[Moody '01]

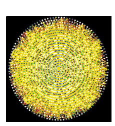

Protein Interactions
[genomebiology.com]

CMU SCS

Graphs - why should we care?

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

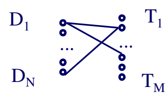

- web: hyper-text graph
- ... and more:

CMU SCS

Graphs - why should we care?

- network of companies & board-of-directors members
- ‘viral’ marketing
- web-log (‘blog’) news propagation
- computer network security: email/IP traffic and anomaly detection
- ....

CMU SCS

Problem #1 - network and graph mining

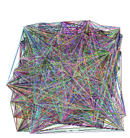

- How does the Internet look like?
- How does the web look like?
- What is ‘normal’/‘abnormal’?
- which patterns/laws hold?

CMU SCS

Graph mining

- Are real graphs random?

UCI 2007 C. Faloutsos 10

CMU SCS

Laws and patterns

- Are real graphs random?
- A: NO!!
  - Diameter
  - in- and out- degree distributions
  - other (surprising) patterns

UCI 2007 C. Faloutsos 11

CMU SCS

Solution#1

- Power law in the degree distribution [SIGCOMM99]

internet domains

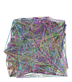

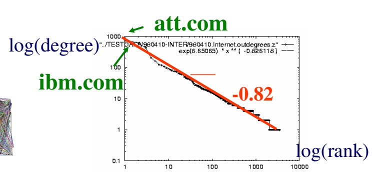

UCI 2007 C. Faloutsos 12

CMU SCS

Solution#1': Eigen Exponent *E*

Eigenvalue

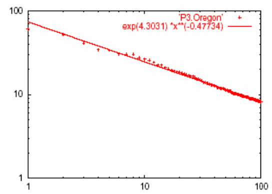

Rank of decreasing eigenvalue

Exponent = slope

E = -0.48

May 2001

- A2: power law in the eigenvalues of the adjacency matrix

CMU SCS

But:

How about graphs from other domains?

CMU SCS

The Peer-to-Peer Topology

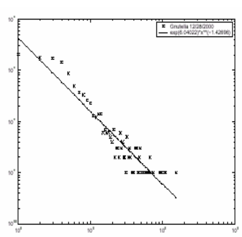

(a) Gnutella snapshot from Dec. 28, 2000 (|r|=0.94)

[Jovanovic+]

- Frequency versus degree
- Number of adjacent peers follows a power-law

CMU SCS

More power laws:

citation counts: (*citeseer.nj.nec.com* 6/2001)

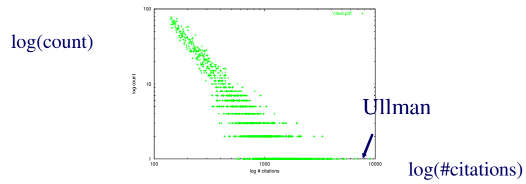

CMU SCS

More power laws:

- web hit counts [w/ A. Montgomery]

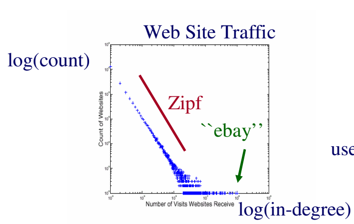

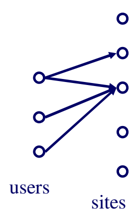

CMU SCS

epinions.com

- who-trusts-whom [Richardson + Domingos, KDD 2001]

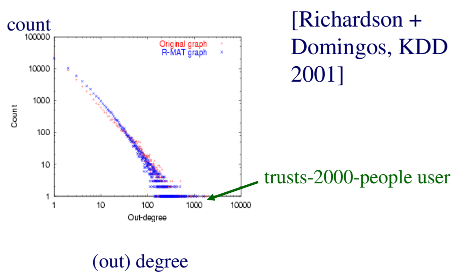

Outline

- Problem definition / Motivation
- ➡ Static & **dynamic** laws; generators
- Tools: CenterPiece graphs; Tensors
- Other projects (Virus propagation, e-bay fraud detection)
- Conclusions

Motivation

Data mining: ~ find patterns (rules, outliers)

- ✓ Problem#1: How do real graphs look like?
- Problem#2: How do they evolve?
- Problem#3: How to generate realistic graphs

TOOLS

- Problem#4: Who is the 'master-mind'?
- Problem#5: Track communities over time

Problem#2: Time evolution

- with Jure Leskovec (CMU/MLD)

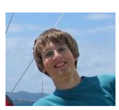

- and Jon Kleinberg (Cornell – sabb. @ CMU)

Evolution of the Diameter

- Prior work on Power Law graphs hints at **slowly growing diameter**:
  - diameter ~ O(log N)
  - diameter ~ O(log log N)
- What is happening in real data?

Evolution of the Diameter

- Prior work on Power Law graphs hints at **slowly growing diameter**:
  - diameter ~ O(log N)
  - diameter ~ O(log log N)
- What is happening in real data?
- Diameter **shrinks** over time

Diameter – ArXiv citation graph

- Citations among physics papers
- 1992 –2003
- One graph per year

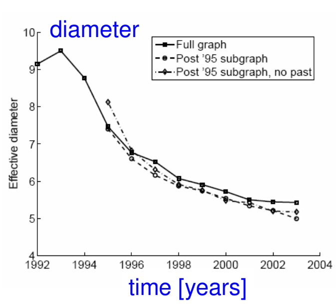

CMU SCS

Diameter – "Autonomous Systems"

- Graph of Internet
- One graph per day
- 1997 – 2000

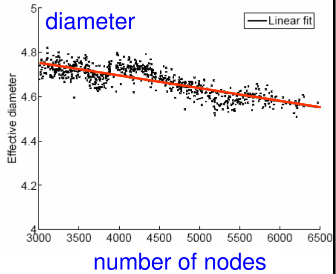

CMU SCS

Diameter – "Affiliation Network"

- Graph of collaborations in physics – authors linked to papers
- 10 years of data

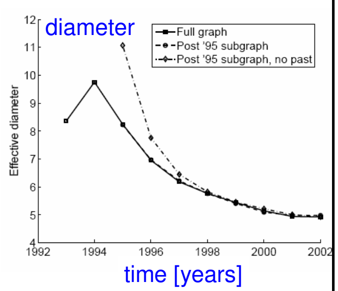

CMU SCS

Diameter – "Patents"

- Patent citation network
- 25 years of data

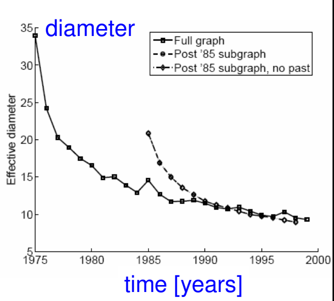

Temporal Evolution of the Graphs

- N(t) … nodes at time t
- E(t) … edges at time t
- Suppose that
  - N(t+1) = 2 * N(t)
- Q: what is your guess for
  - E(t+1) =? 2 * E(t)

Temporal Evolution of the Graphs

- N(t) … nodes at time t
- E(t) … edges at time t
- Suppose that
  - N(t+1) = 2 * N(t)
- Q: what is your guess for
  - E(t+1) =? 2 * E(t)
- A: over-doubled!
  - But obeying the ``Densification Power Law''

Densification – Physics Citations

- Citations among physics papers
- 2003:
  - 29,555 papers, 352,807 citations

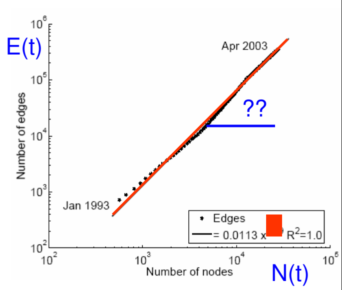

Densification – Physics Citations

- Citations among physics papers
- 2003:
  - 29,555 papers, 352,807 citations

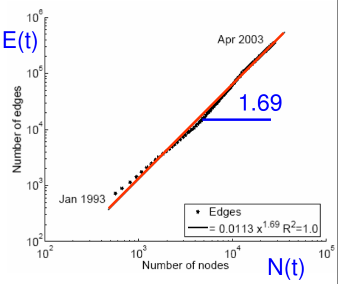

Densification – Physics Citations

- Citations among physics papers
- 2003:
  - 29,555 papers, 352,807 citations

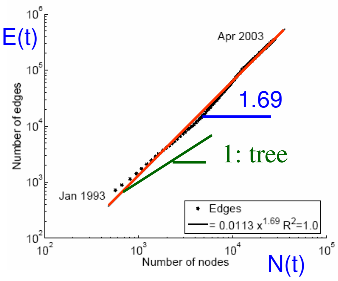

Densification – Physics Citations

- Citations among physics papers
- 2003:
  - 29,555 papers, 352,807 citations

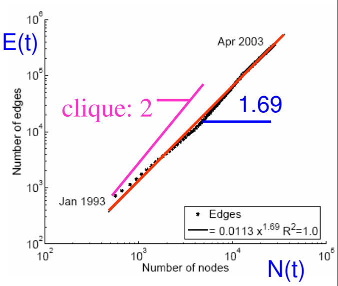

Densification – Patent Citations

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

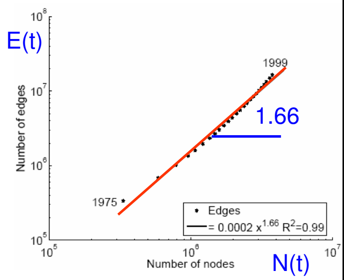

Densification – Autonomous Systems

- Graph of Internet
- 2000
  - 6,000 nodes
  - 26,000 edges
- One graph per day

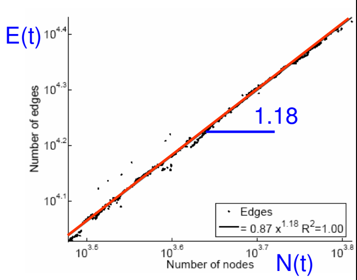

Densification – Affiliation Network

- Authors linked to their publications
- 2002
  - 60,000 nodes
    - 20,000 authors
    - 38,000 papers
  - 133,000 edges

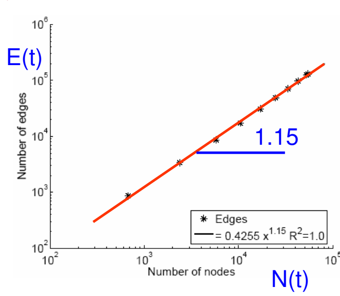

CMU SCS

Outline

- Problem definition / Motivation
- ➡ Static & dynamic laws; **generators**
- Tools: CenterPiece graphs; Tensors
- Other projects (Virus propagation, e-bay fraud detection)
- Conclusions

CMU SCS

Motivation

Data mining: ~ find patterns (rules, outliers)

- ✓ Problem#1: How do real graphs look like?
- ✓ Problem#2: How do they evolve?
- Problem#3: How to generate realistic graphs

TOOLS

- Problem#4: Who is the 'master-mind'?
- Problem#5: Track communities over time

CMU SCS

Problem#3: Generation

- Given a growing graph with count of nodes N_1, N_2, …
- Generate a realistic sequence of graphs that will obey all the patterns

CMU SCS

Problem Definition

- Given a growing graph with count of nodes N_1, N_2, …
- Generate a realistic sequence of graphs that will obey all the patterns
  - Static Patterns
    - Power Law Degree Distribution
    - Power Law eigenvalue and eigenvector distribution
    - Small Diameter
  - Dynamic Patterns
    - Growth Power Law
    - Shrinking/Stabilizing Diameters

CMU SCS

Problem Definition

- Given a growing graph with count of nodes N_1, N_2, …
- Generate a realistic sequence of graphs that will obey all the patterns

- **Idea: Self-similarity**
  - Leads to power laws
  - Communities within communities
  - …

CMU SCS

Kronecker Product – a Graph

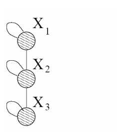

1	1	0
1	1	1
0	1	1

G_1

Adjacency matrix

Kronecker Product – a Graph

- Continuing multiplying with G_1 we obtain G_4 and so on …

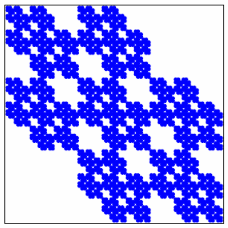

G_4 adjacency matrix

Kronecker Product – a Graph

- Continuing multiplying with G_1 we obtain G_4 and so on …

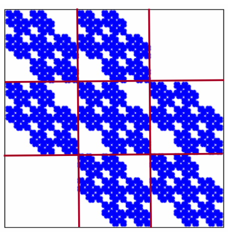

G_4 adjacency matrix

Kronecker Product – a Graph

- Continuing multiplying with G_1 we obtain G_4 and so on …

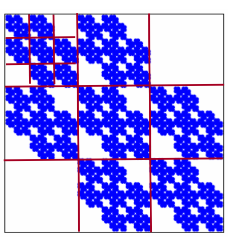

G_4 adjacency matrix

Properties:

- We can PROVE that
  - Degree distribution is multinomial ~ power law
  - Diameter: constant
  - Eigenvalue distribution: multinomial
  - First eigenvector: multinomial
- See [Leskovec+, PKDD'05] for proofs

Problem Definition

- Given a growing graph with nodes N_1, N_2, …
- Generate a realistic sequence of graphs that will obey all the patterns
  - Static Patterns
    - ✔ Power Law Degree Distribution
    - ✔ Power Law eigenvalue and eigenvector distribution
    - ✔ Small Diameter
  - Dynamic Patterns
    - ✔ Growth Power Law
    - ✔ Shrinking/Stabilizing Diameters
- First and only generator for which we can **prove** all these properties

skip

Stochastic Kronecker Graphs

- Create $N_1 \times N_1$ probability matrix P_1
- Compute the k^{th} Kronecker power P_k
- For each entry p_{uv} of P_k include an edge (u,v) with probability p_{uv}

0.4	0.2
0.1	0.3

P_1

Kronecker multiplication →

0.16	0.08	0.08	0.04
0.04	0.12	0.02	0.06
0.04	0.02	0.12	0.06
0.01	0.03	0.03	0.09

P_k

→ Instance Matrix G_2

flip biased coins

CMU SCS

Experiments

- How well can we match real graphs?
  - Arxiv: physics citations:
    - 30,000 papers, 350,000 citations
    - 10 years of data
  - U.S. Patent citation network
    - 4 million patents, 16 million citations
    - 37 years of data
  - Autonomous systems – graph of internet
    - Single snapshot from January 2002
    - 6,400 nodes, 26,000 edges
- We show both static and temporal patterns

CMU SCS

Arxiv – Degree Distribution

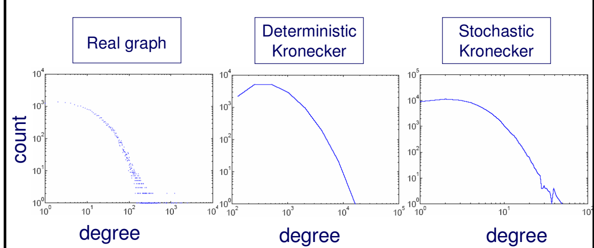

CMU SCS

Arxiv – Scree Plot

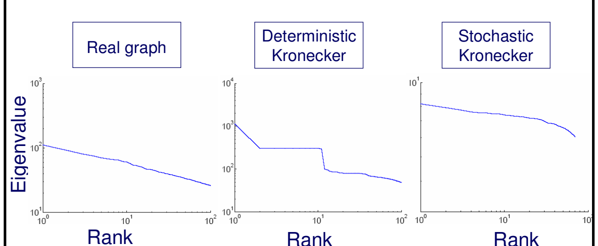

CMU SCS

Arxiv – Densification

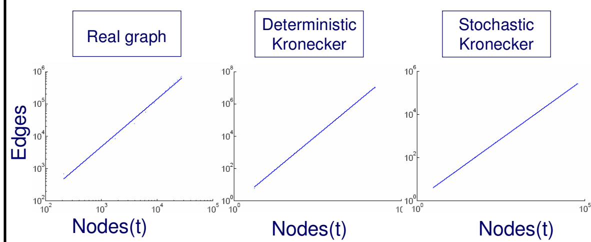

CMU SCS

Arxiv – Effective Diameter

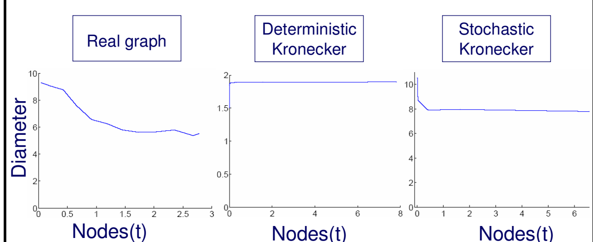

CMU SCS

(Q: how to fit the parm's?)

A:

- Stochastic version of Kronecker graphs +
- Max likelihood +
- Metropolis sampling
- [Leskovec+, ICML'07]

Experiments on real AS graph

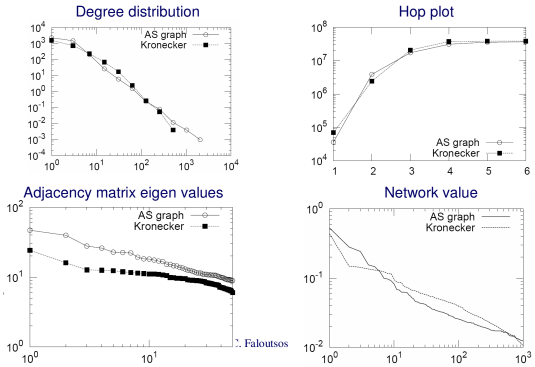

Conclusions

- Kronecker graphs have:
  - All the static properties
    - ✓Heavy tailed degree distributions
    - ✓Small diameter
    - ✓Multinomial eigenvalues and eigenvectors
  - All the temporal properties
    - ✓Densification Power Law
    - ✓Shrinking/Stabilizing Diameters
  - We can formally prove these results

Outline

- Problem definition / Motivation
- Static & dynamic laws; generators
- ➡ Tools: **CenterPiece graphs**; Tensors
- Other projects (Virus propagation, e-bay fraud detection)
- Conclusions

CMU SCS

Motivation

Data mining: ~ find patterns (rules, outliers)

✓ Problem#1: How do real graphs look like?

✓ Problem#2: How do they evolve?

✓ Problem#3: How to generate realistic graphs

TOOLS

- Problem#4: Who is the 'master-mind'?
- Problem#5: Track communities over time

UCI 2007 C. Faloutsos 58

CMU SCS

Problem#4: MasterMind – 'CePS'

- w/ Hanghang Tong, KDD 2006
- htong <at> cs.cmu.edu

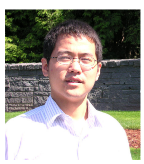

UCI 2007 C. Faloutsos 59

CMU SCS

Center-Piece Subgraph(Ceps)

- **Given** Q query nodes
- **Find** Center-piece ($\leq b$)
- App.
  - Social Networks
  - Law Inforcement, …
- Idea:
  - Proximity -> random walk with restarts

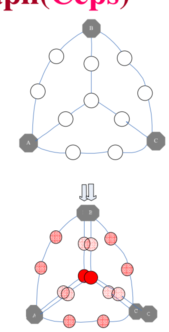

UCI 2007 C. Faloutsos 60

CMU SCS

Case Study: AND query

CMU SCS

Case Study: AND query

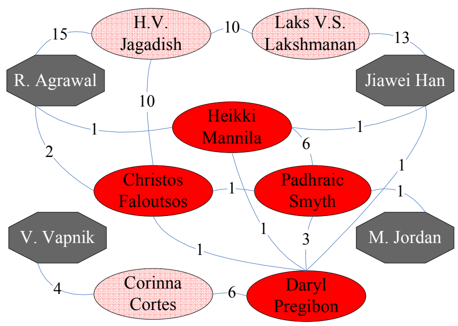

CMU SCS

Case Study: AND query

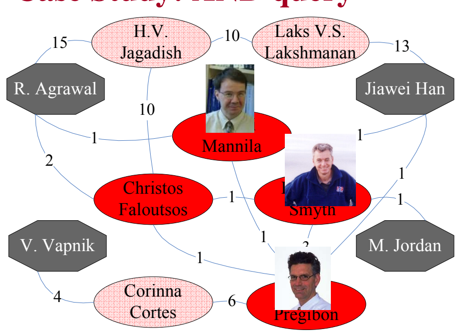

CMU SCS

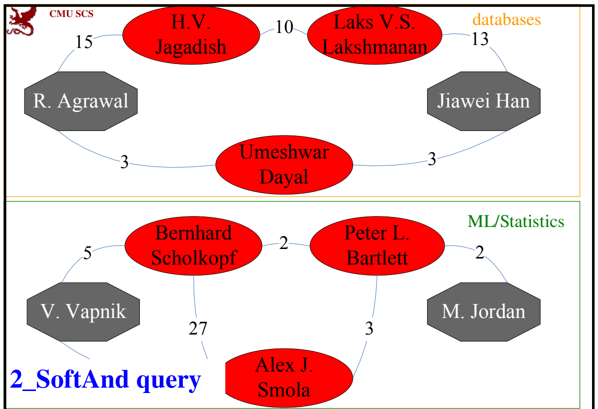

CMU SCS

Conclusions

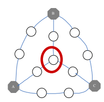

- Q1:How to measure the importance?
- A1: RWR+K_SoftAnd
- Q2: How to find connection subgraph?
- A2:"Extract" Alg.
- Q3:How to do it efficiently?
- A3:Graph Partition (Fast CePS)
  - ~90% quality
  - 6:1 speedup; 150x speedup (ICDM'06, b.p. award)

CMU SCS

Outline

- Problem definition / Motivation
- Static & dynamic laws; generators
- ➡ Tools: CenterPiece graphs; **Tensors**
- Other projects (Virus propagation, e-bay fraud detection)
- Conclusions

CMU SCS

Motivation

Data mining: ~ find patterns (rules, outliers)

✓ Problem#1: How do real graphs look like?

✓ Problem#2: How do they evolve?

✓ Problem#3: How to generate realistic graphs

TOOLS

✓ Problem#4: Who is the ‘master-mind’?

- Problem#5: Track communities over time

CMU SCS

Tensors for time evolving graphs

- [Jimeng Sun+ KDD’06]
- [“ , SDM’07]
- [CF, Kolda, Sun, SDM’07 tutorial]

CMU SCS

Social network analysis

- Static: find community structures

1990

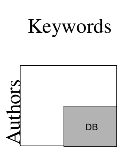

CMU SCS

Social network analysis

- Static: find community structures
- Dynamic: monitor community structure evolution; spot abnormal individuals; abnormal time-stamps

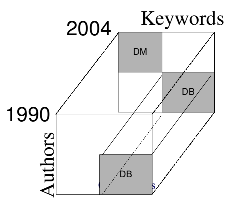

CMU SCS

Application 1: Multiway latent semantic indexing (LSI)

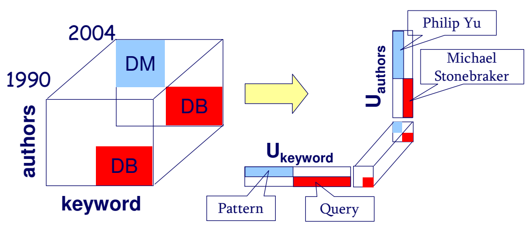

- Projection matrices specify the clusters
- Core tensors give cluster activation level

CMU SCS

Bibliographic data (DBLP)

- Papers from VLDB and KDD conferences
- Construct 2nd order tensors with yearly windows with <author, keywords>
- Each tensor: 4584×3741
- 11 timestamps (years)

CMU SCS

Multiway LSI

Authors	Keywords	Year
michael carey, michael stonebraker, h. jagadish, hector garcia-molina	queri,parallel,optimization,concurr, objectorient	1995
surajit chaudhuri,mitch cherniack,michael stonebraker,ugur etintemel	distribut,systems,view,storage,servic,process,cache	2004
jiawei han,jian pei,philip s. yu, jianyong wang,charu c. aggarw	streams,pattern,support, cluster, gener,queri	2004

DB

DM

- Two groups are correctly identified: Databases and Data mining
- People and concepts are drifting over time

UCI 2007 C. Faloutsos 73

CMU SCS

Conclusions

Tensor-based methods (WTA/DTA/STA):

- spot patterns and anomalies on time evolving graphs, and
- on streams (monitoring)

UCI 2007 C. Faloutsos 74

CMU SCS

Outline

- Problem definition / Motivation
- Static & dynamic laws; generators
- Tools: CenterPiece graphs; Tensors
- ➡ Other projects (**Virus propagation**, e-bay fraud detection)
- Conclusions

UCI 2007 C. Faloutsos 75

Virus propagation

- How do viruses/rumors propagate?
- Will a flu-like virus linger, or will it become extinct soon?

The model: SIS

- 'Flu' like: Susceptible-Infected-Susceptible
- Virus 'strength' s= β/δ

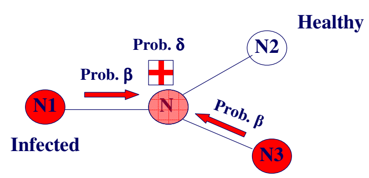

Epidemic threshold τ

of a graph: the value of τ, such that

if strength $s = \beta / \delta < \tau$

an epidemic can not happen

Thus,

- given a graph
- compute its epidemic threshold

Epidemic threshold τ

What should τ depend on?

- avg. degree? and/or highest degree?
- and/or variance of degree?
- and/or third moment of degree?
- and/or diameter?

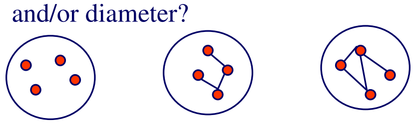

Epidemic threshold

- [Theorem] We have no epidemic, if

$$\boxed{\beta/\delta < \tau = 1/\lambda_{1,A}}$$

Epidemic threshold

- [Theorem] We have no epidemic, if

recovery prob.

epidemic threshold

$$\boxed{\beta/\delta < \tau = 1/\lambda_{1,A}}$$

attack prob.

largest eigenvalue of adj. matrix ***A***

Proof: [Wang+03]

CMU SCS

Experiments (Oregon)

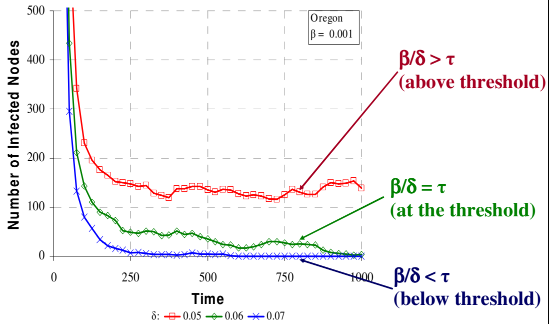

CMU SCS

Outline

- Problem definition / Motivation
- Static & dynamic laws; generators
- Tools: CenterPiece graphs; Tensors
- Other projects (Virus propagation, **e-bay fraud detection**)
- Conclusions

CMU SCS

E-bay Fraud detection

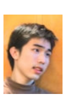

w/ Polo Chau &
Shashank Pandit, CMU

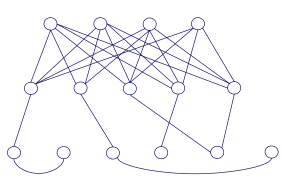

E-bay Fraud detection - NetProbe

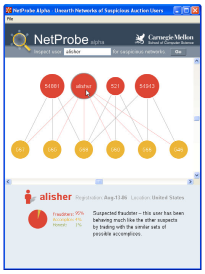

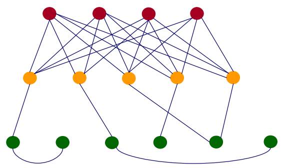

OVERALL CONCLUSIONS

- Graphs pose a wealth of fascinating problems
- self-similarity and power laws work, when textbook methods fail!
- New patterns (shrinking diameter!)
- New generator: Kronecker

‘Philosophical’ observation

Graph mining brings together:

- ML/AI / IR
- Stat, Num. analysis,
- Systems (DB (Gb/Tb), Networks)
- sociology, epidemiology
- physics, ++…

CMU SCS

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

CMU SCS

References

- Jure Leskovec, Jon Kleinberg and Christos Faloutsos *Graphs over Time: Densification Laws, Shrinking Diameters and Possible Explanations* KDD 2005, Chicago, IL. ("Best Research Paper" award).
- Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos *Realistic, Mathematically Tractable Graph Generation and Evolution, Using Kronecker Multiplication* (ECML/PKDD 2005), Porto, Portugal, 2005.

CMU SCS

References

- Jure Leskovec and Christos Faloutsos, *Scalable Modeling of Real Graphs using Kronecker Multiplication*, ICML 2007, Corvallis, OR, USA
- Jimeng Sun, Dacheng Tao, Christos Faloutsos *Beyond Streams and Graphs: Dynamic Tensor Analysis,* KDD 2006, Philadelphia, PA

CMU SCS

References

- Jimeng Sun, Yinglian Xie, Hui Zhang, Christos Faloutsos. *Less is More: Compact Matrix Decomposition for Large Sparse Graphs*, SDM, Minneapolis, Minnesota, Apr 2007. [pdf]

CMU SCS

THANK YOU!

Contact info:

www. cs.cmu.edu /~christos

(w/ papers, datasets, code, etc)