Introduction & Motivation

EE599: Social Network Systems

Keith M. Chugg

Fall 2014

USC University of Southern California

Overview

- My motivation and background
- Big picture overview in this context
- Class format
- Discussion

My Background

My Firsts:

- x-domain email (1990)
- Personal computer (1990)
- Non-tech email (~1993)
- Web browser (~1994)
- Web purchase (~2004)

1967

First program

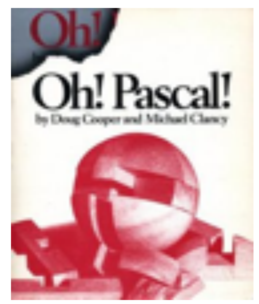

1985-1989

HARVEY MUDD COLLEGE

college

1989-1995

grad school

Google

2000

TrellisWare TECHNOLOGIES

2007 iPhone

2009 Facebook

Biggest Changes I've Lived

- Availability
  - Connected 24/7
  - Instant information on everything
  - Technology changes who we are
- Bombardment - attention/ competition
  - Advertising
  - Propaganda

This is becoming a science

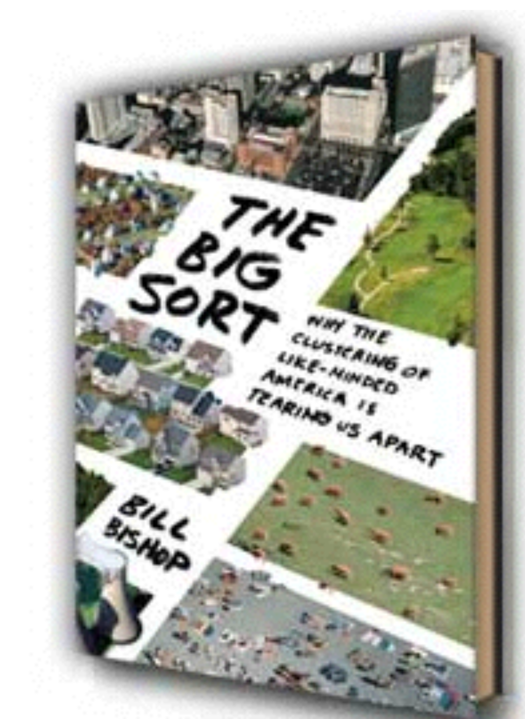

Arbitrary Example

"1985 Coke Commercial"

https://www.youtube.com/watch?v=yJoocpy7UBc

"2014 Coke Commercial"

https://www.youtube.com/watch?v=XTfVsAs4Ybs

BS General Engineering

- Electrical, Mechanical, Chemical

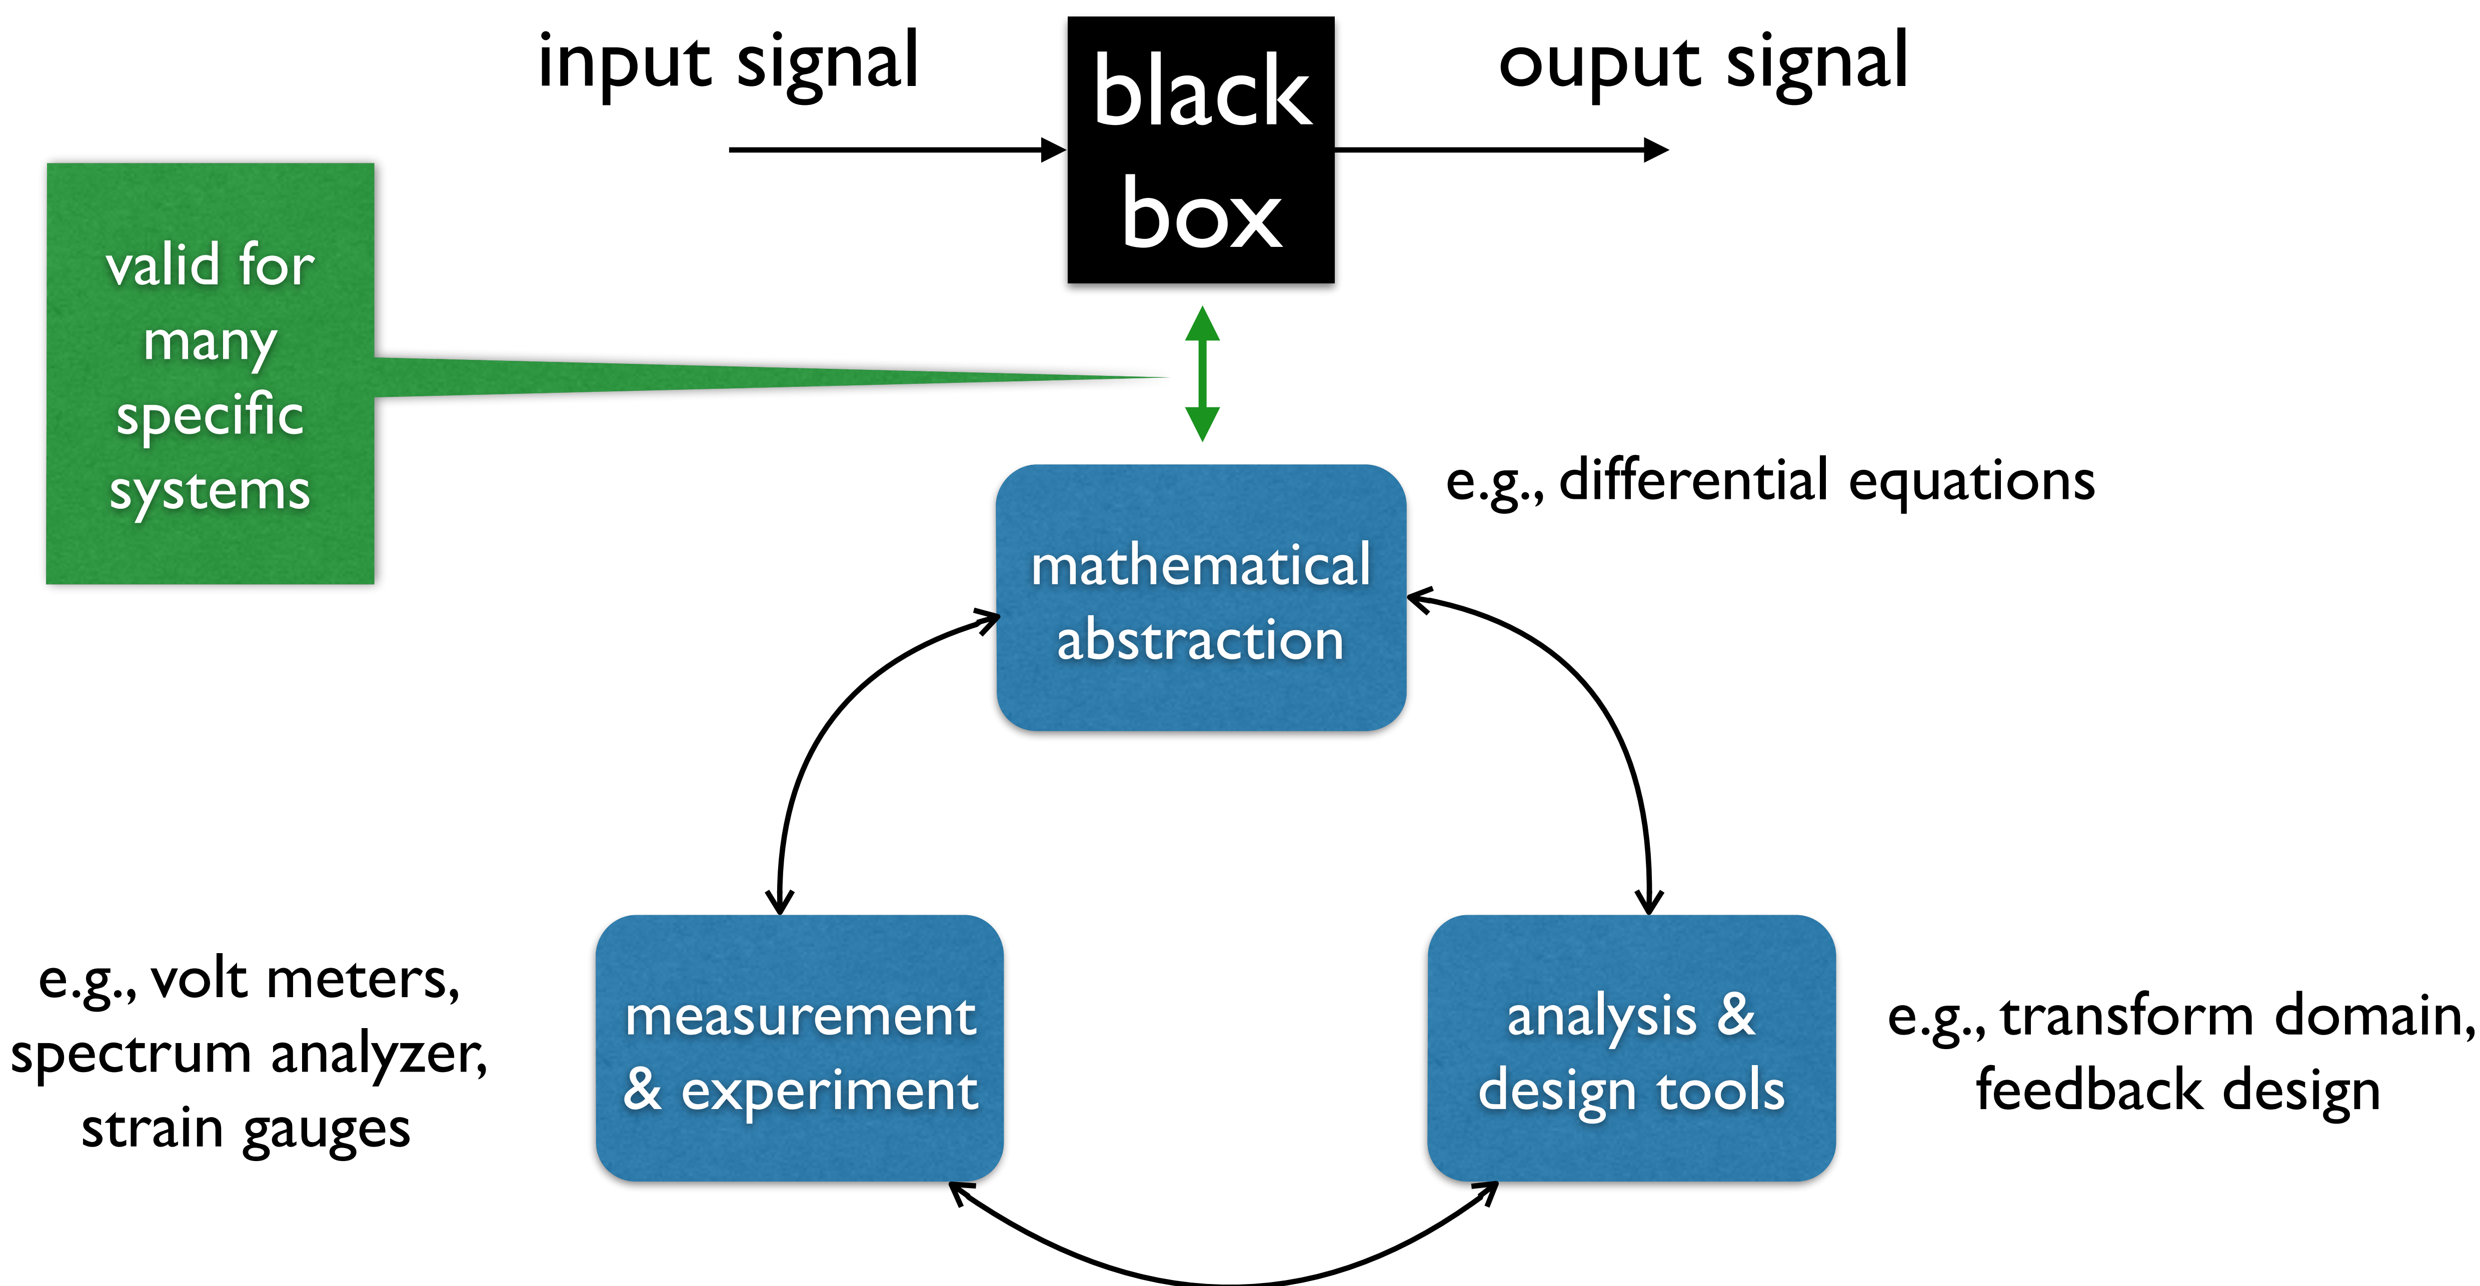

My Background

- Core Expertise (Digital Communications)
  - Probability, decision theory, estimation, graphical models, algorithms, implementation architectures
- Weak Areas (wrt EE599)
  - Programming tools
  - Computer Architecture

Example: 3 Kayak Trips & 6 People

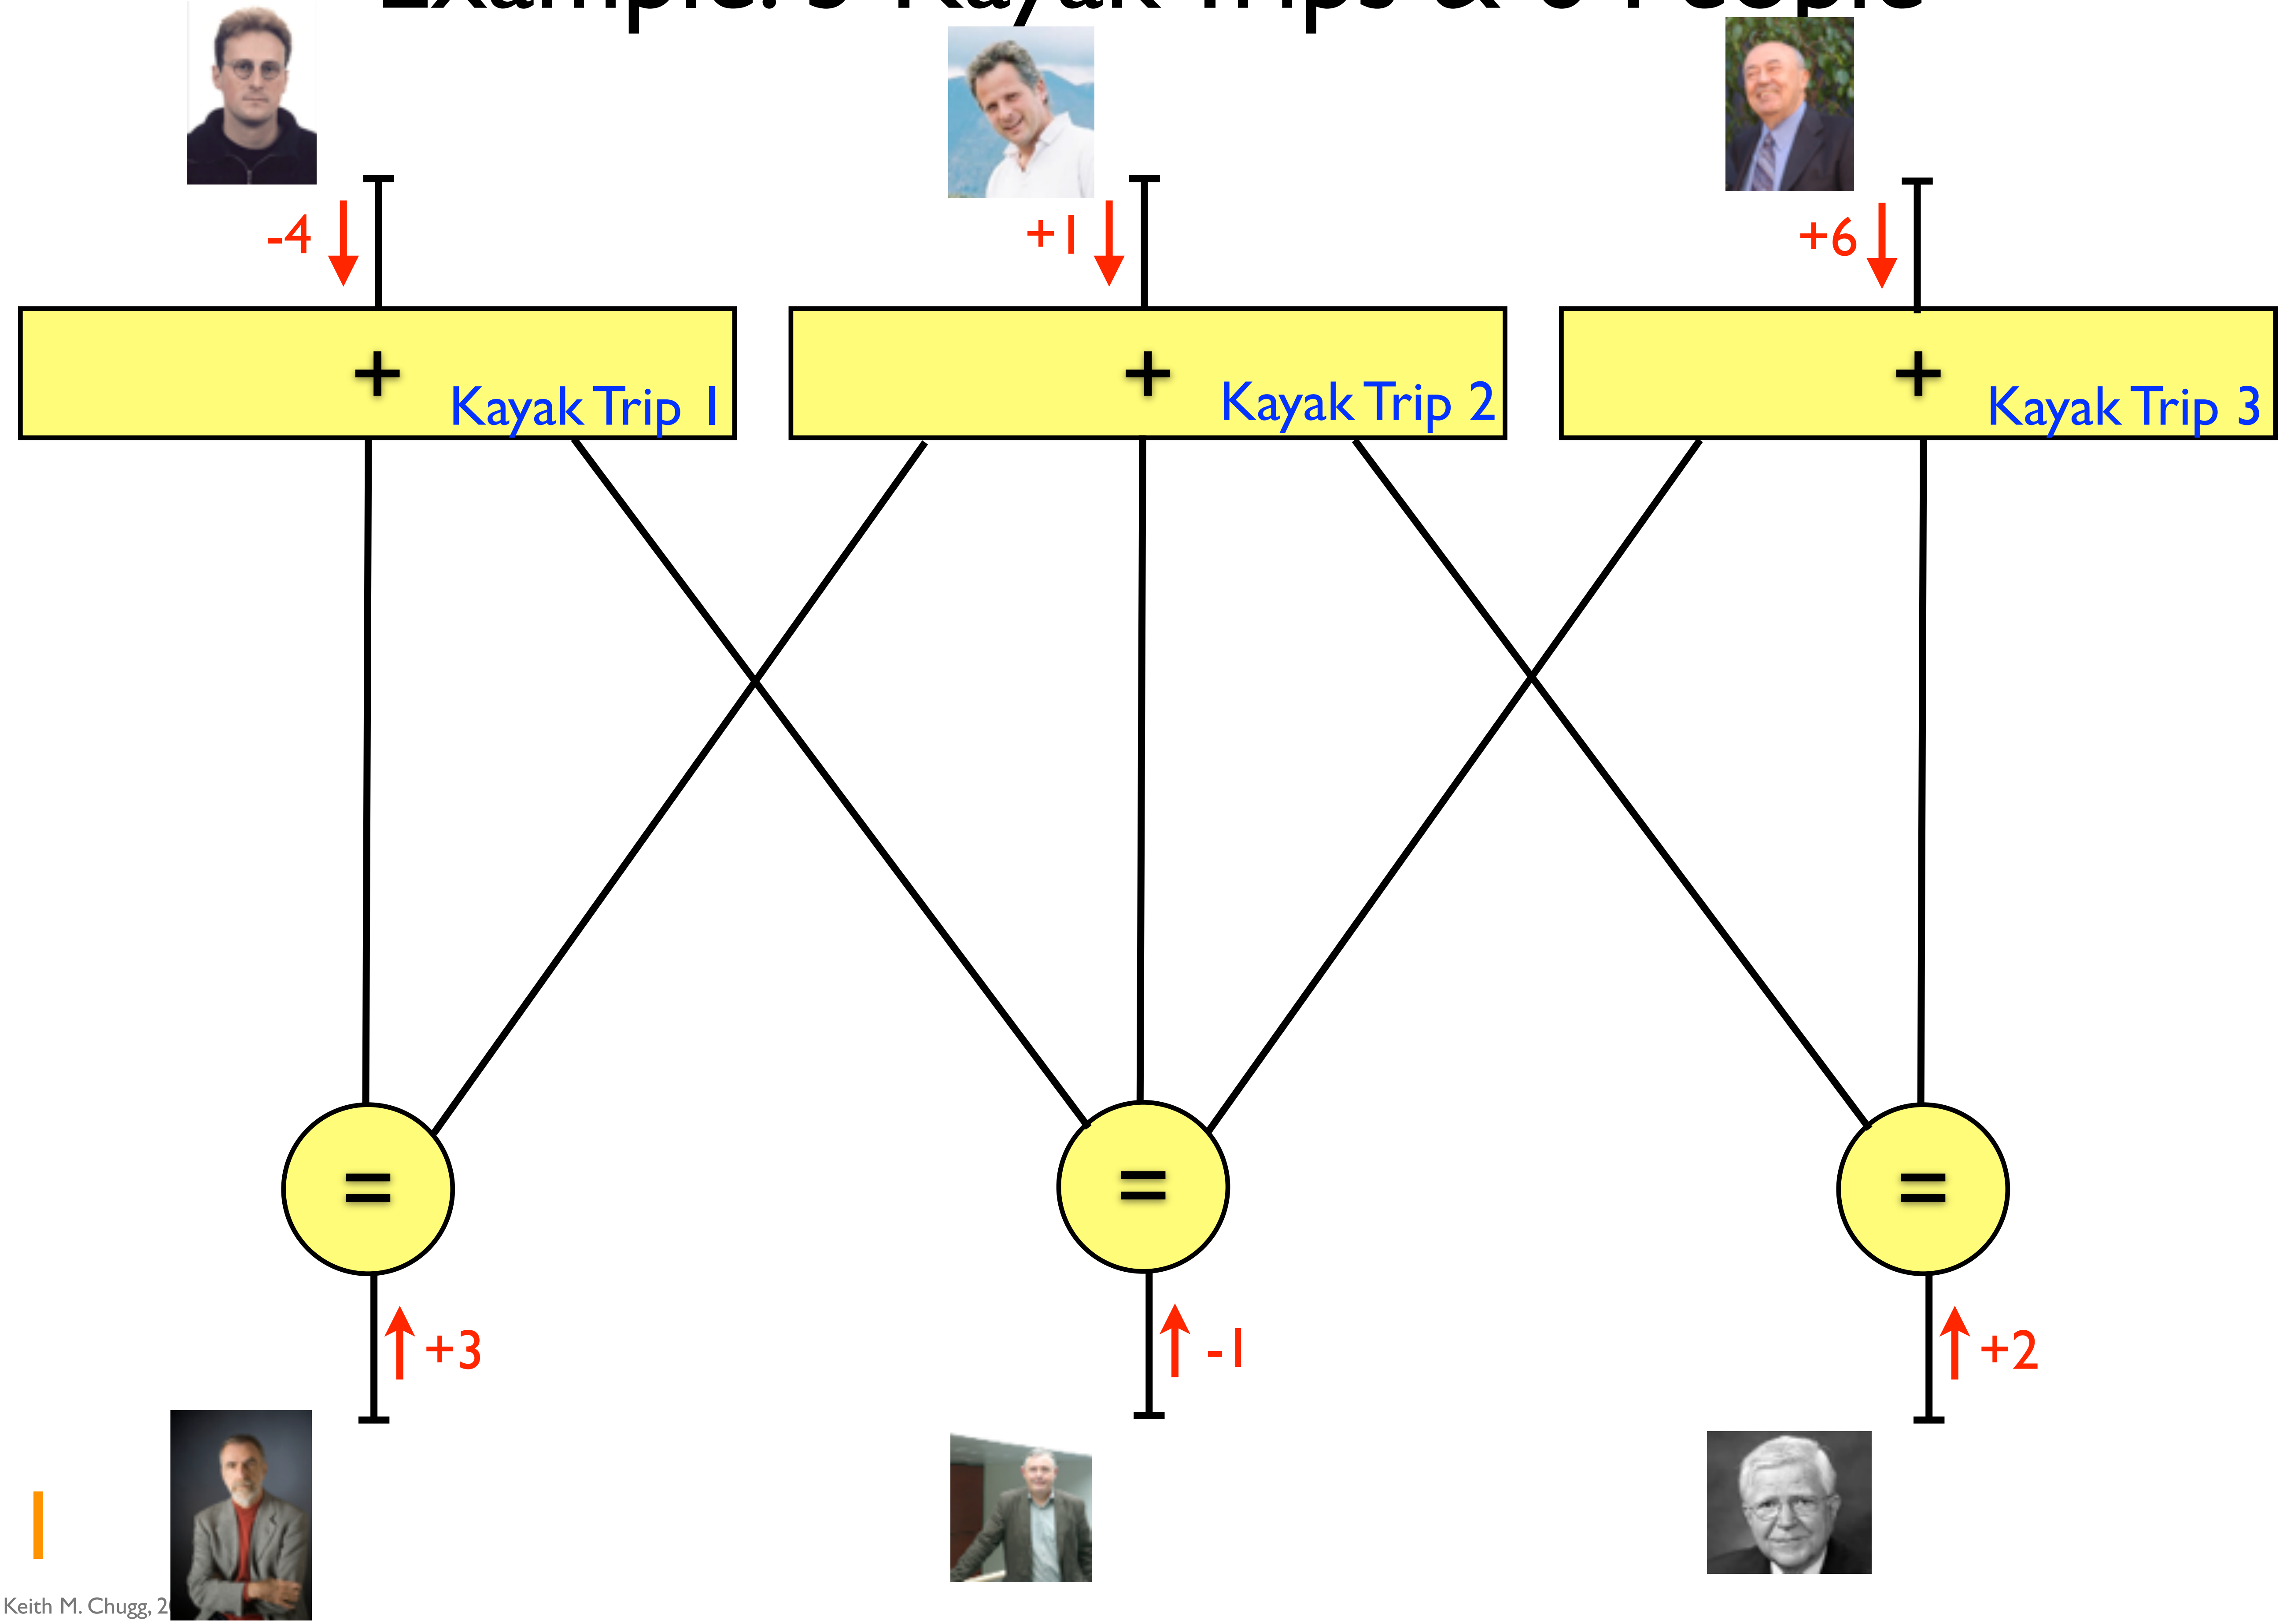

Overview

- My motivation and background
- Big picture overview in this context
- Class format
- Discussion

Emerging SysEng View

- Complex systems, networks, data science

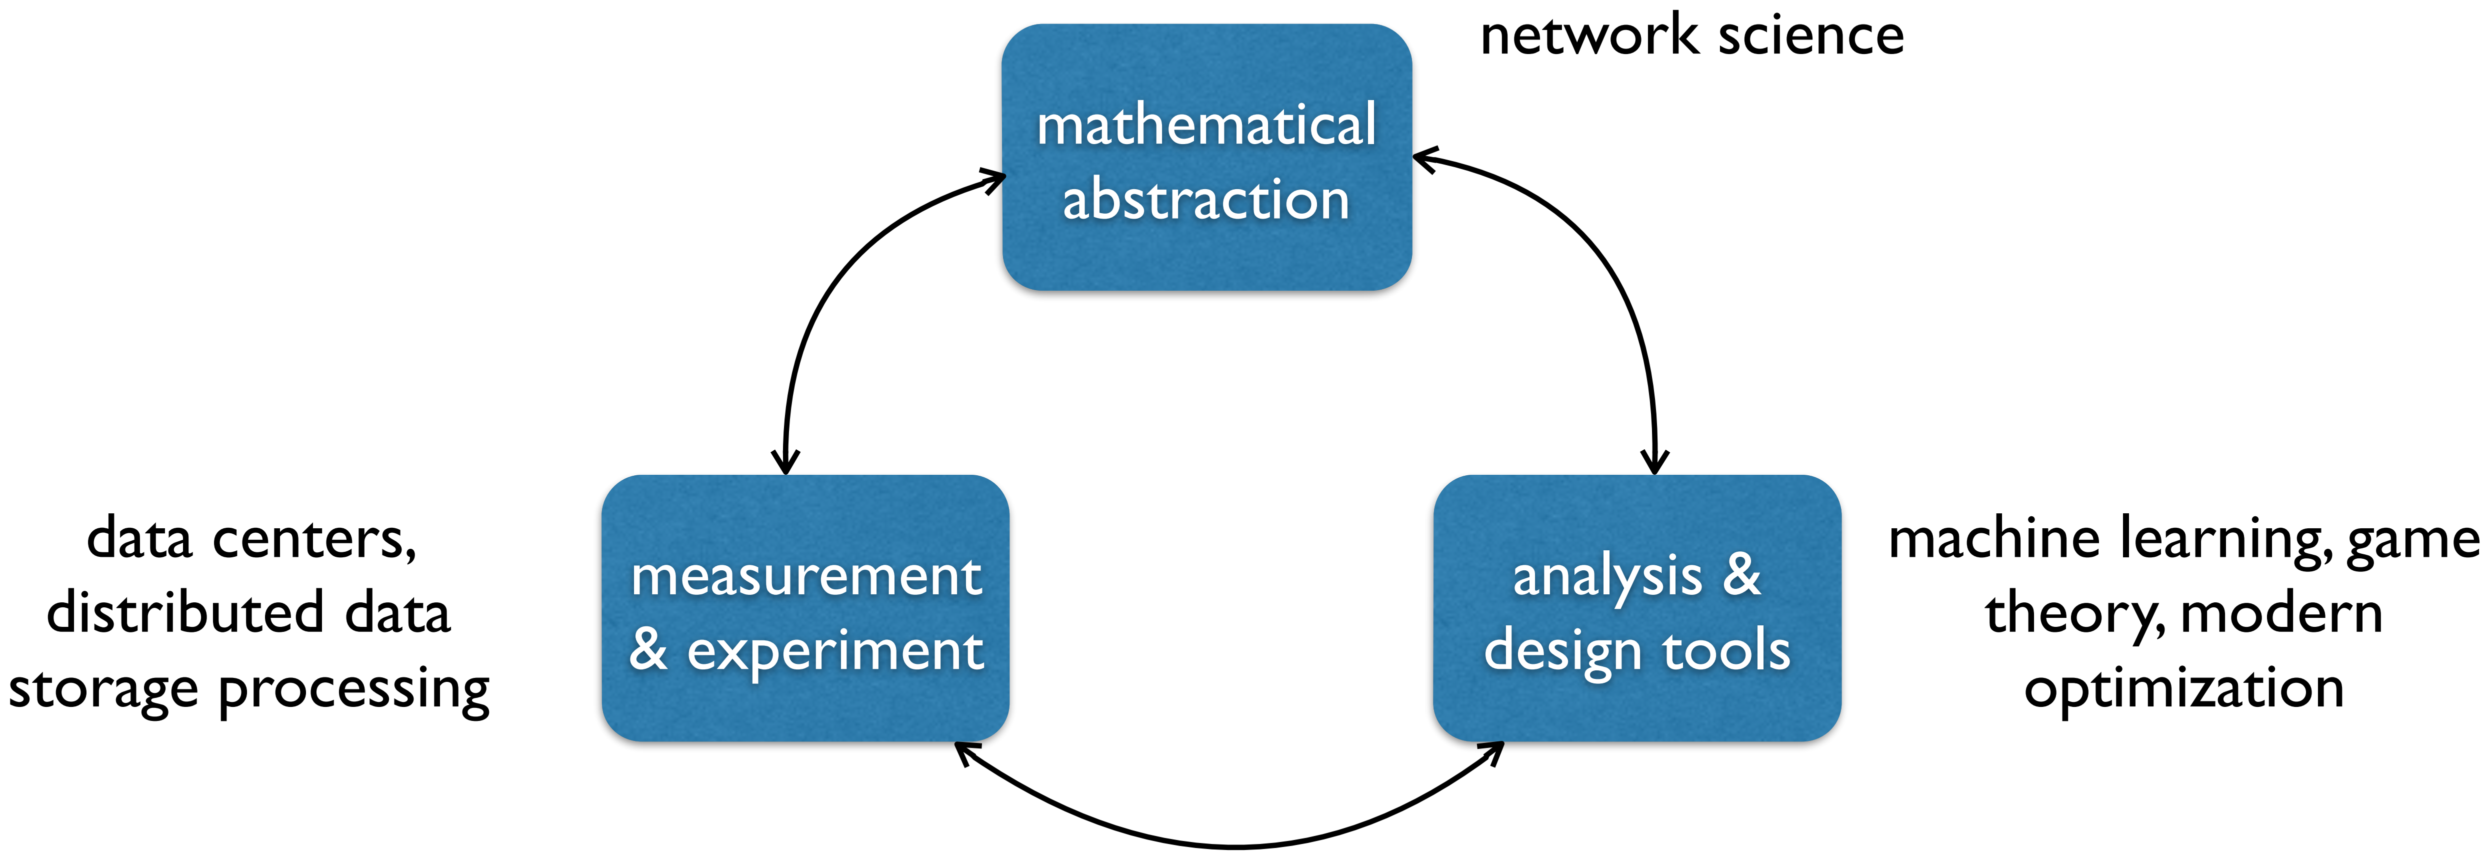

Example: 6-degrees of Separation

Image 3.13

Stanley Milgram (1933–1984)

American social psychologist known for his experiments on obedience and authority. He designed and carried out the small world experiment in 1967 as part of his Harvard dissertation.

Baraba'si, Ch3

The first experimental study of small world phenomena took place four decades after Karinthy, in 1967, when Stanley Milgram turned the idea into an experiment probing the structure of social networks [23]. Milgram chose a stock broker in Boston and a divinity student in Sharon, Massachusetts as "targets". Randomly selected residents of Wichita, Kansas and Omaha, Nebraska received a letter containing a short summary of the study's purpose, a photograph, the name, address and information about the target person. They were asked to forward the letter to a friend, relative or acquaintance, who is more likely to know the target person. Milgram wrote in 1969: *"I asked a person of intelligence how many steps he thought it would take, and he said that it would require 100 intermediate persons, or more, to move from Nebraska to Sharon."* Yet, within a few days the first letter arrived, passing through only two links. Eventually 42 of the 160 letters made it back, some requiring close to a dozen intermediates. These completed chains allowed Milgram to determine the number of individuals required to get the letter to the target. He found that the median number of intermediates was 5.5, a relatively small number and remarkably close to Karinthy's 1929 insight.

Baraba'si, Ch3

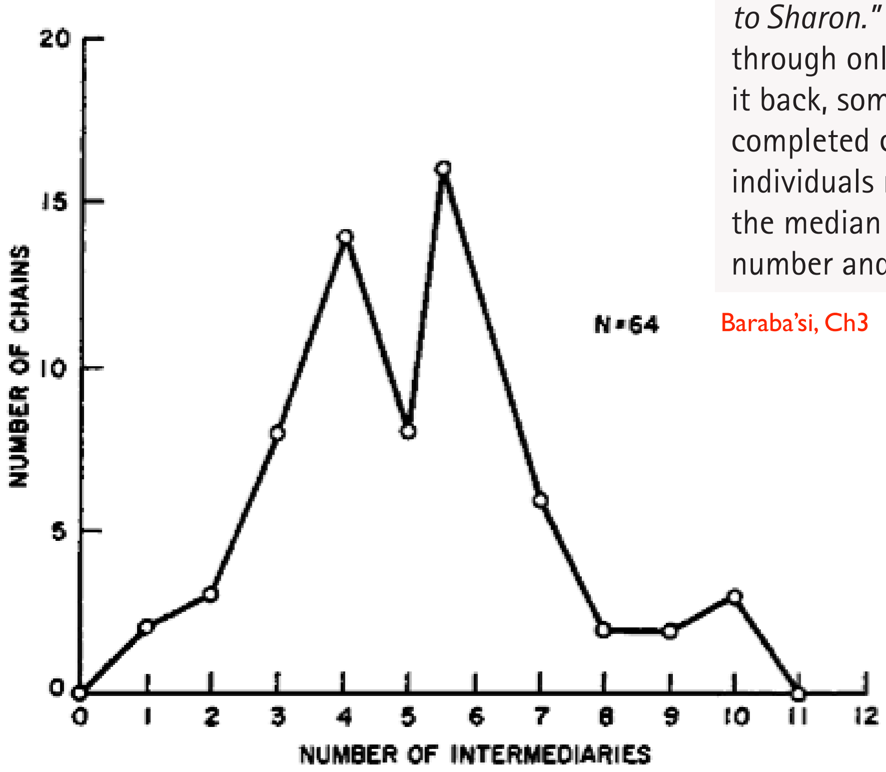

Figure 2.10: A histogram from Travers and Milgram's paper on their small-world experiment [391]. For each possible length (labeled "number of intermediaries" on the x-axis), the plot shows the number of successfully completed chains of that length. In total, 64 chains reached the target person, with a median length of six.

Easley & Kleinberg -see slightly different historical description, pg. 35

Example: 6-degrees of Separation

Milgram's experiment could not detect the true distance between his study's participants, as he lacked an accurate map of the full social network. Today Facebook has the most extensive social network map ever assembled. Using Facebook's social graph of May 2011, consisting of 721 million active users and 68 billion symmetric friendship links, the average distance between the users was 4.74. The figure shows the distance distribution, p_d, for all pairs of Facebook users worldwide (full dataset) and within the US only. Therefore, instead of 'six degrees' researchers detected only 'four degrees of separation' [4], closer to the prediction of Eq. (20) than to Milgram's six degrees [23]. Using Facebook's N and L Eq. (19) predicts the average degree to be approximately 3.90, not far from the reported four degrees.

Baraba'si, Ch3

small world phenomena is common in social networks

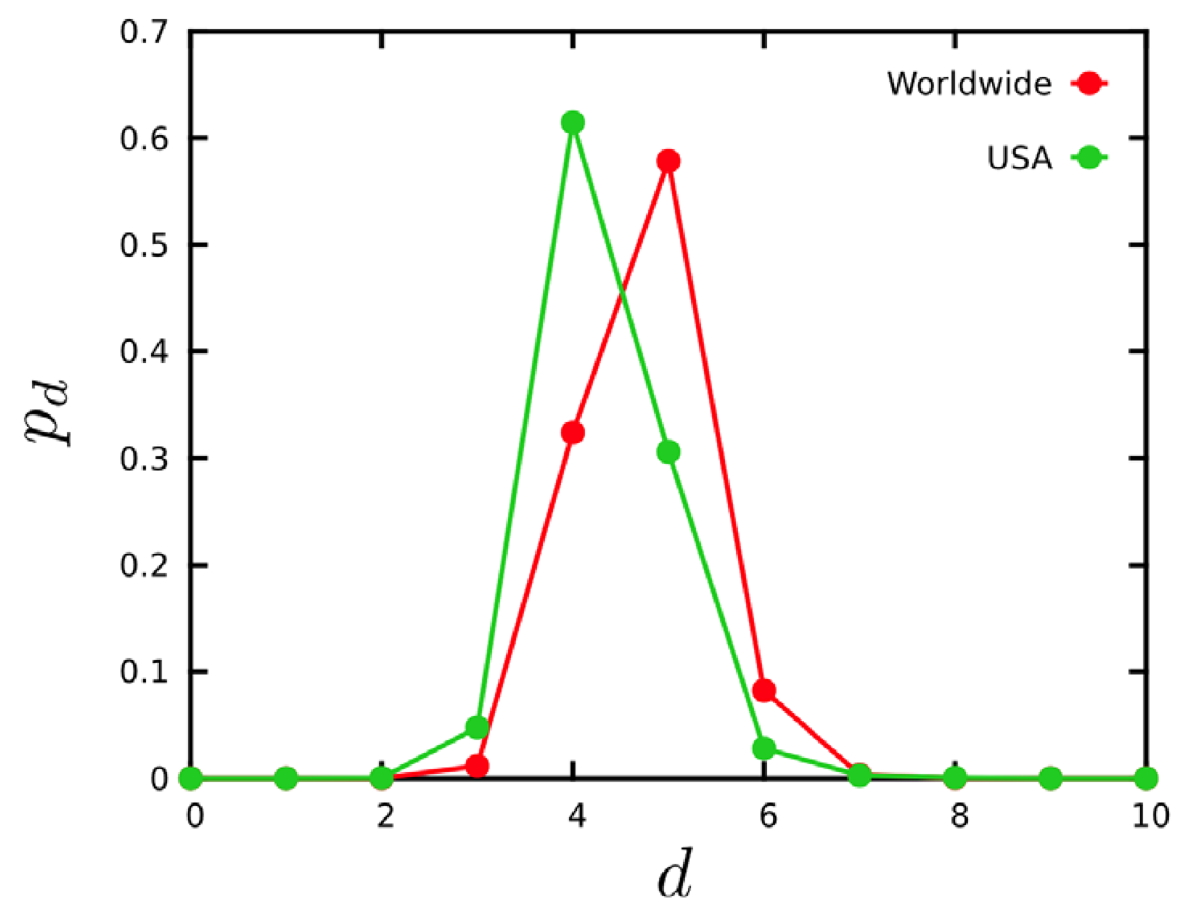

Baraba'si, Ch3

Example: Giant Component

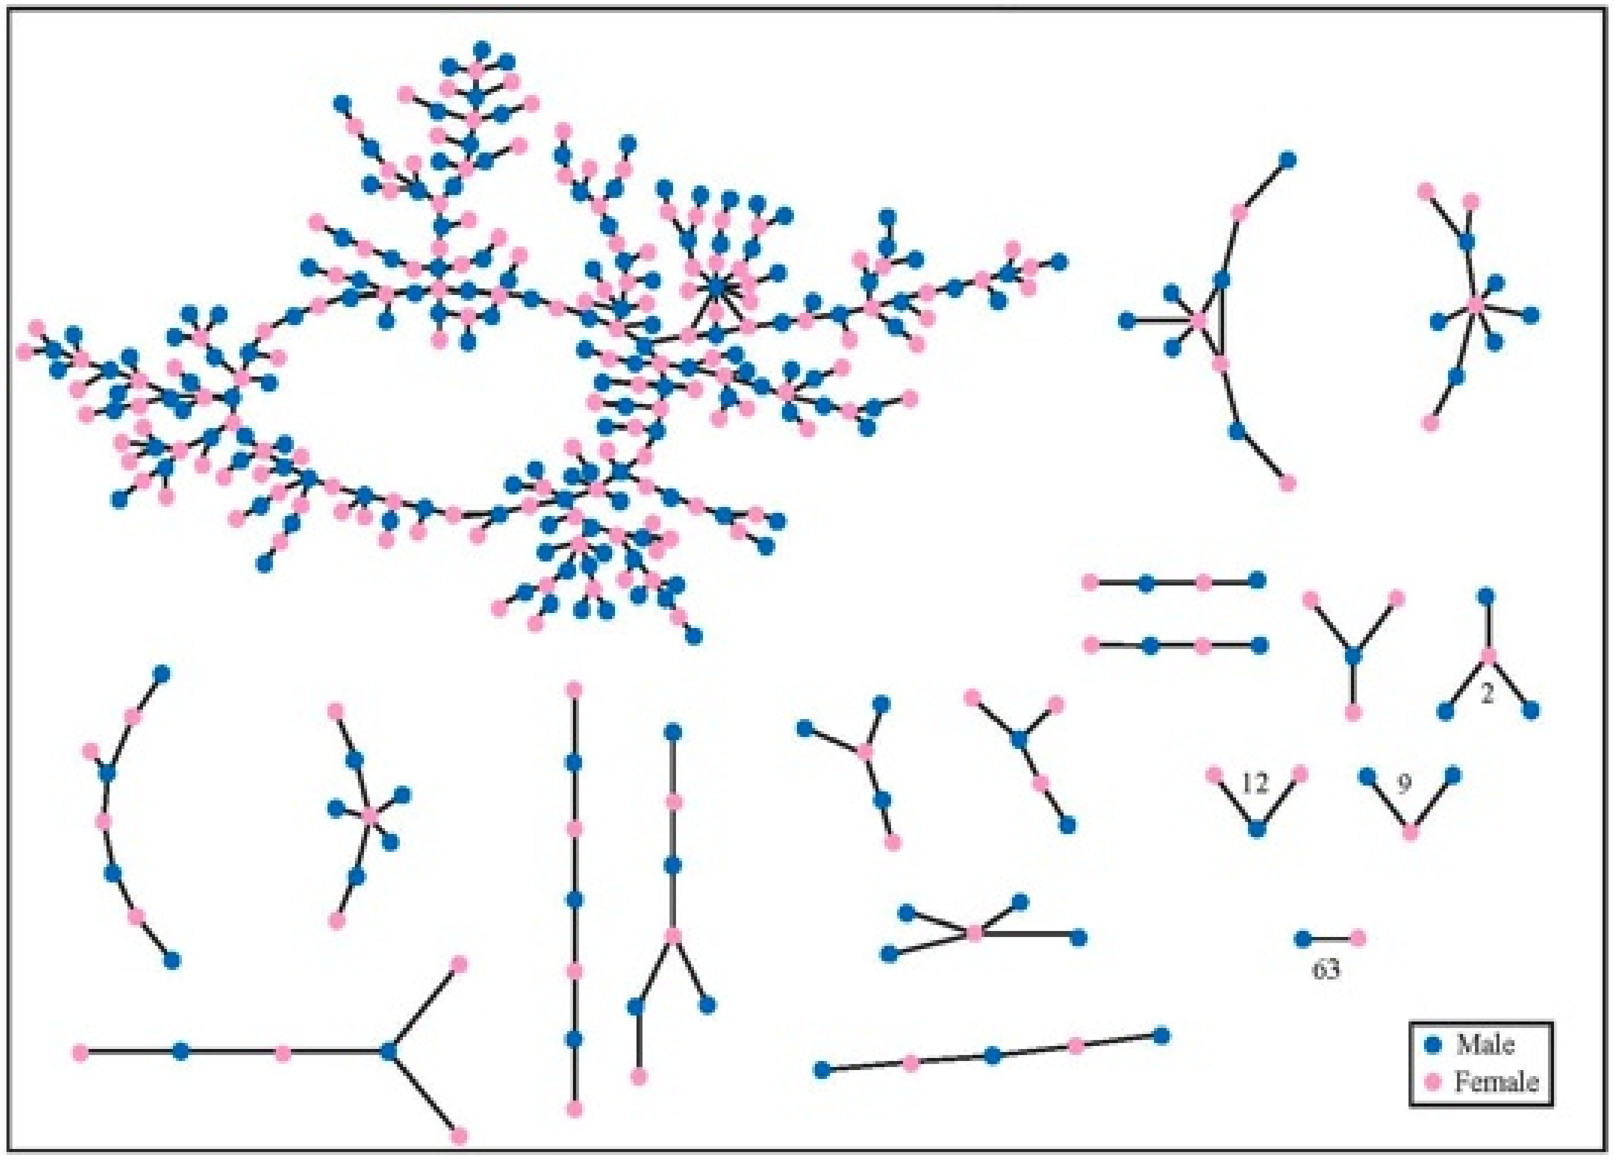

Figure 2.7: A network in which the nodes are students in a large American high school, and an edge joins two who had a romantic relationship at some point during the 18-month period in which the study was conducted [49].

Easley & Kleinberg

Example Informal Org-Chart

Image 1.10c

Understanding the inner workings of an organization.

The position of the leadership within the company's informal network is illustrated on this map, where we colored the nodes based on their company rank within the company. None of the company directors, including the CEO, shown in red, are hubs. Nor are the top managers, shown in blue. The hubs are managers, group leaders and associates, or workers. The biggest hub, hence the most influential individual, is an associate, shown as a gray node in the center.

Baraba'si, Ch1

Example: Political Polarization

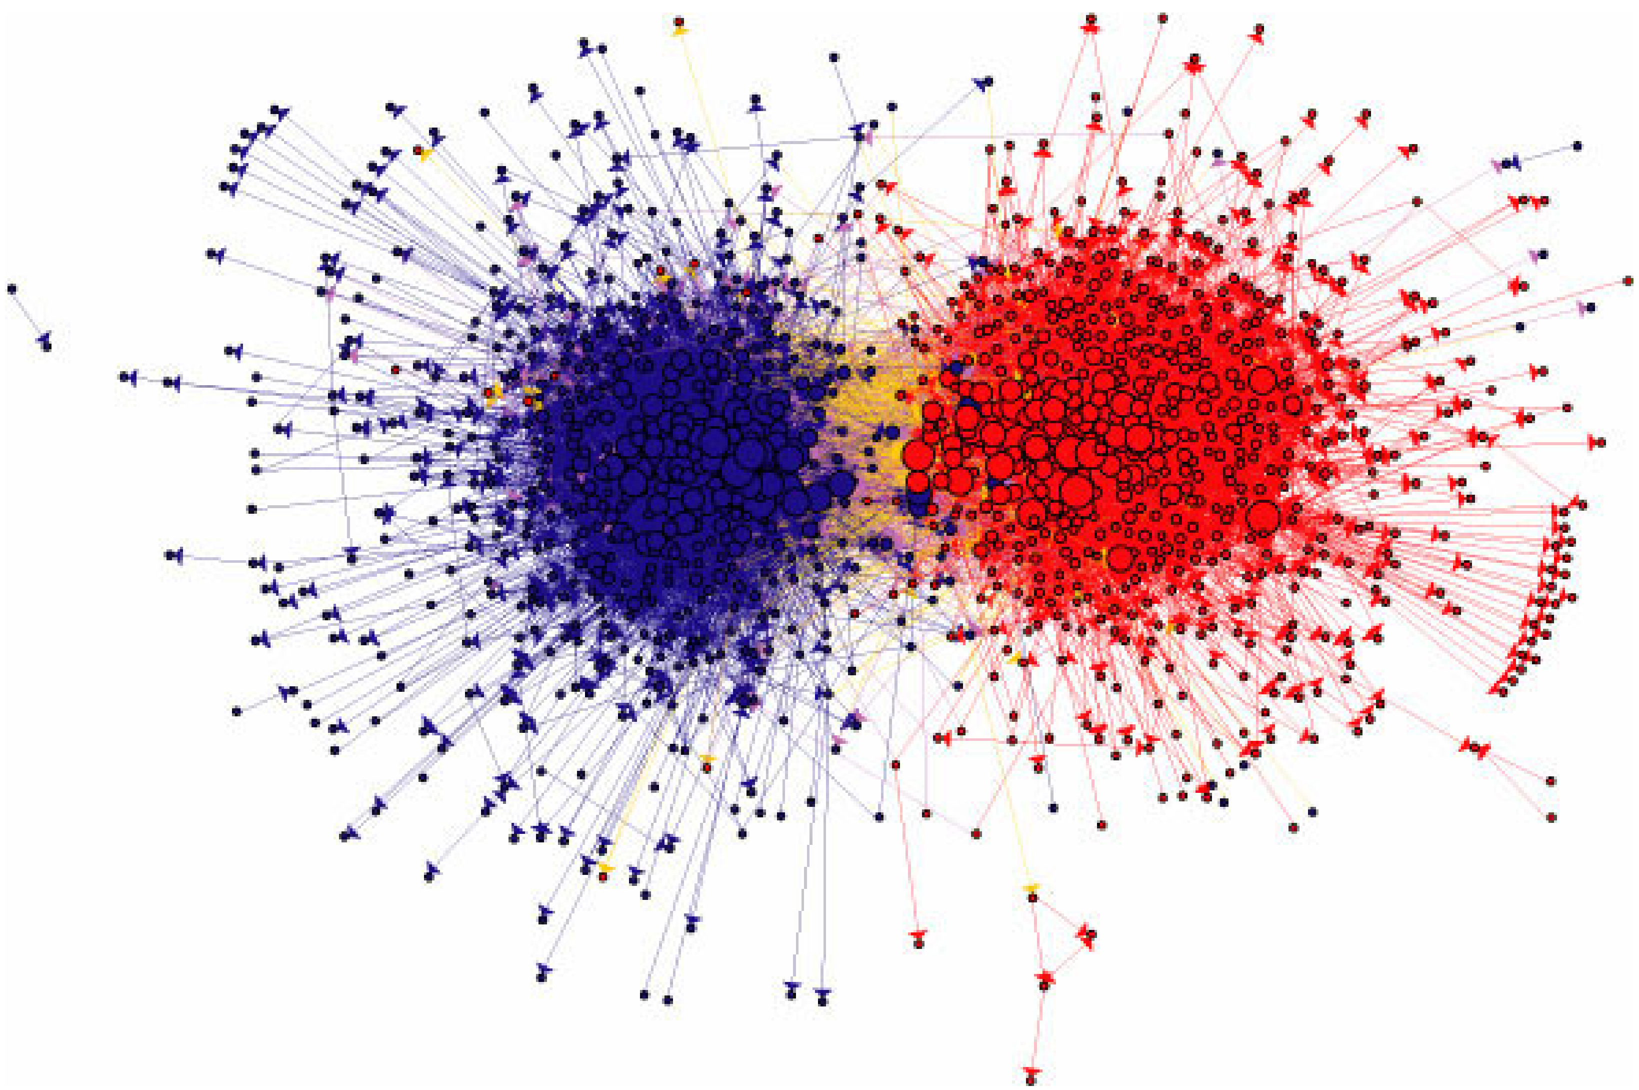

Figure 1.4: The links among Web pages can reveal densely-knit communities and prominent sites. In this case, the network structure of political blogs prior to the 2004 U.S. Presidential election reveals two natural and well-separated clusters [5]. (Image from http://www-personal.umich.edu/ ladamic/img/politicalblogs.jpg)

Easley & Kleinberg

Computational Social Science

- Social scientists research into social network phenomena has a long impressive history
- Field was disrupted in early 2000s when massive online data sets became possible
  - This is ongoing - multidisciplinary
  - Many early models have been verified, refined, extended

The Measurement Machinery

- Data centers
  - massively expensive
- Methods for distributed, regenerative storage
- Methods for distributed computation

Data Centers

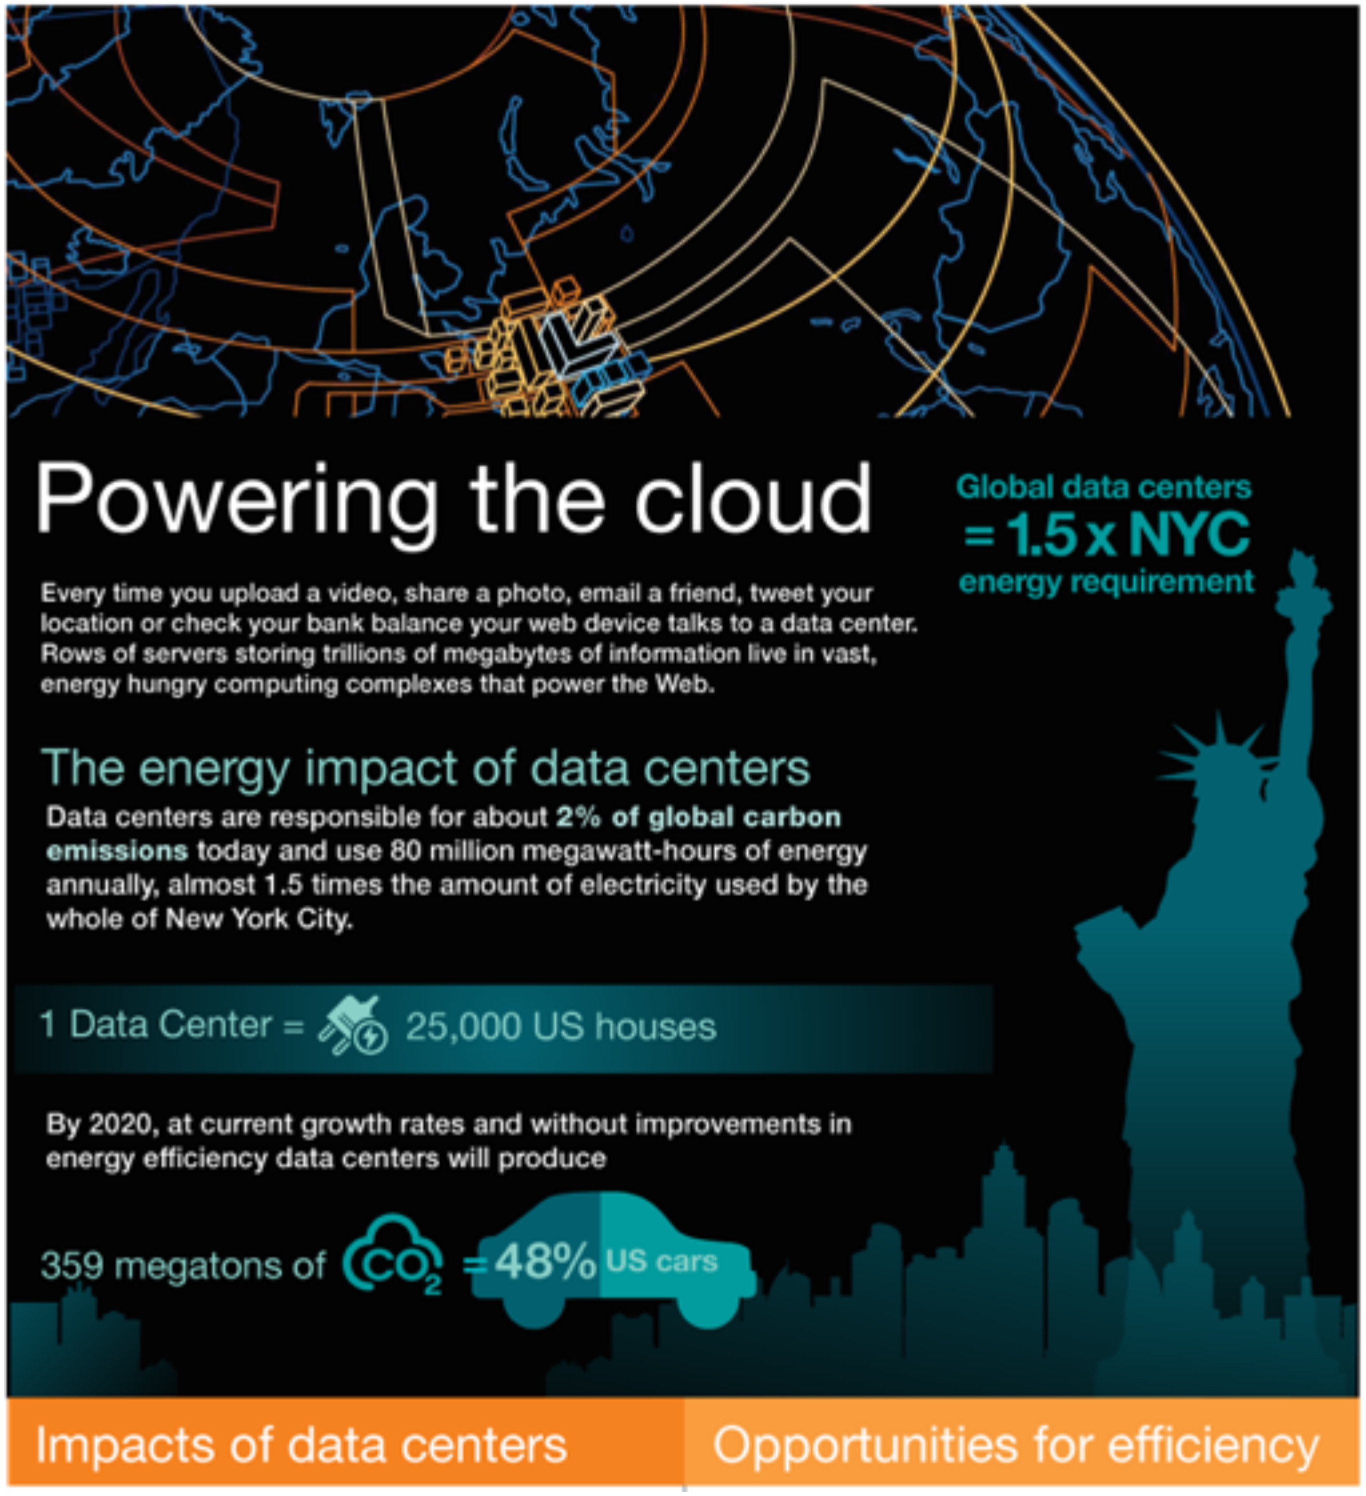

http://www.abb.com/product/ap/db0003db004052/e950c90f13518ffbc125788f0030bda0.aspx

Data Centers

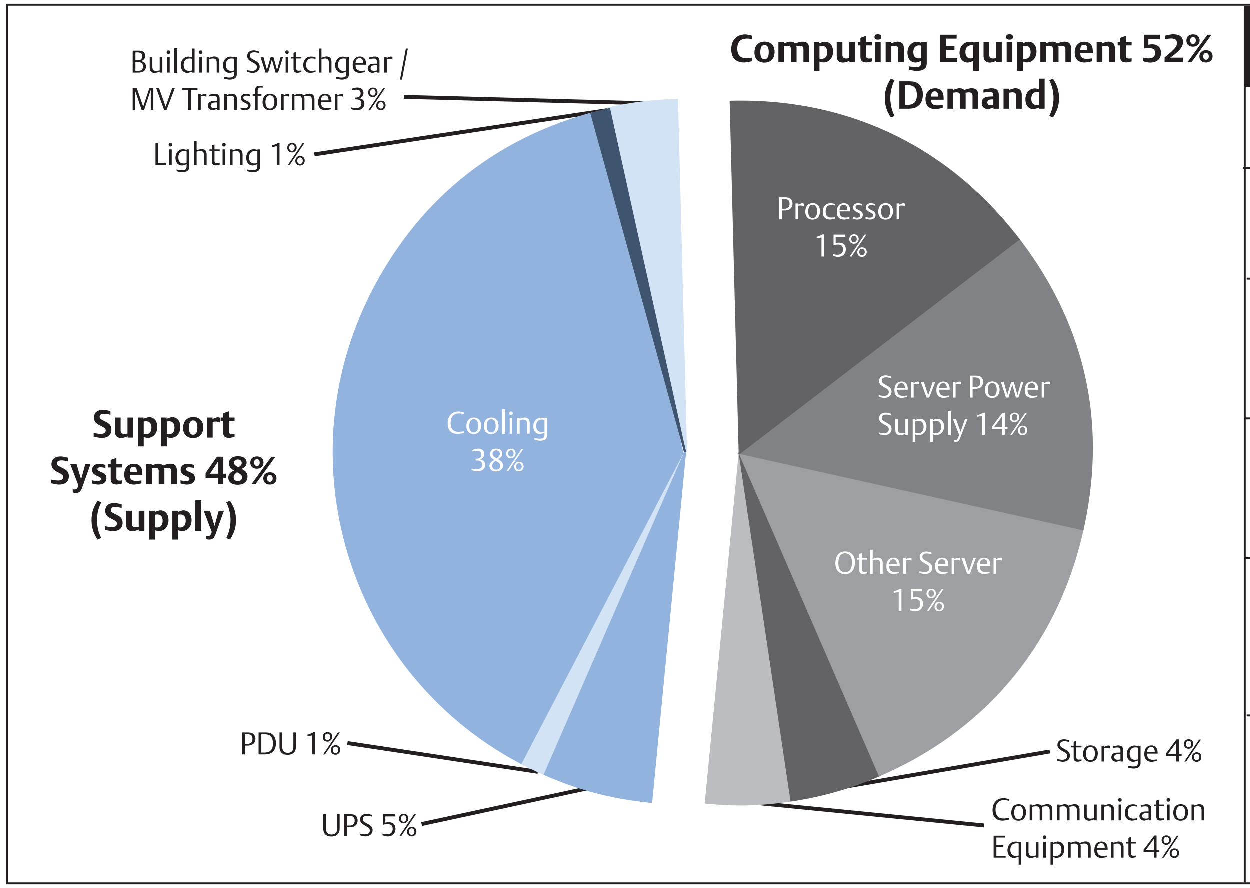

Category	Power Draw*
Computing	588 kW
Lighting	10 kW
UPS and distribution losses	72 kW
Cooling power draw for computing and UPS losses	429 kW
Building switchgear/MV transformer/other losses	28 kW
TOTAL	1,127 kW

Figure 1. Analysis of a typical 5,000-square-foot data center shows that demand-side computing equipment accounts for 52 percent of energy usage and supply-side systems account for 48 percent.

http://www.emersonnetworkpower.com/documentation/en-us/latest-thinking/edc/documents/white%20paper/energylogicreducingdatacenterenergyconsumption.pdf

Trends in Data Centers

- Becoming specialized to services, core competitive advantage
  - Architectures, chipsets, interconnects, etc.
- Optimization goals
  - Max processor utilization, min energy consumption, min interconnect delay

Return: Emerging SysEng View

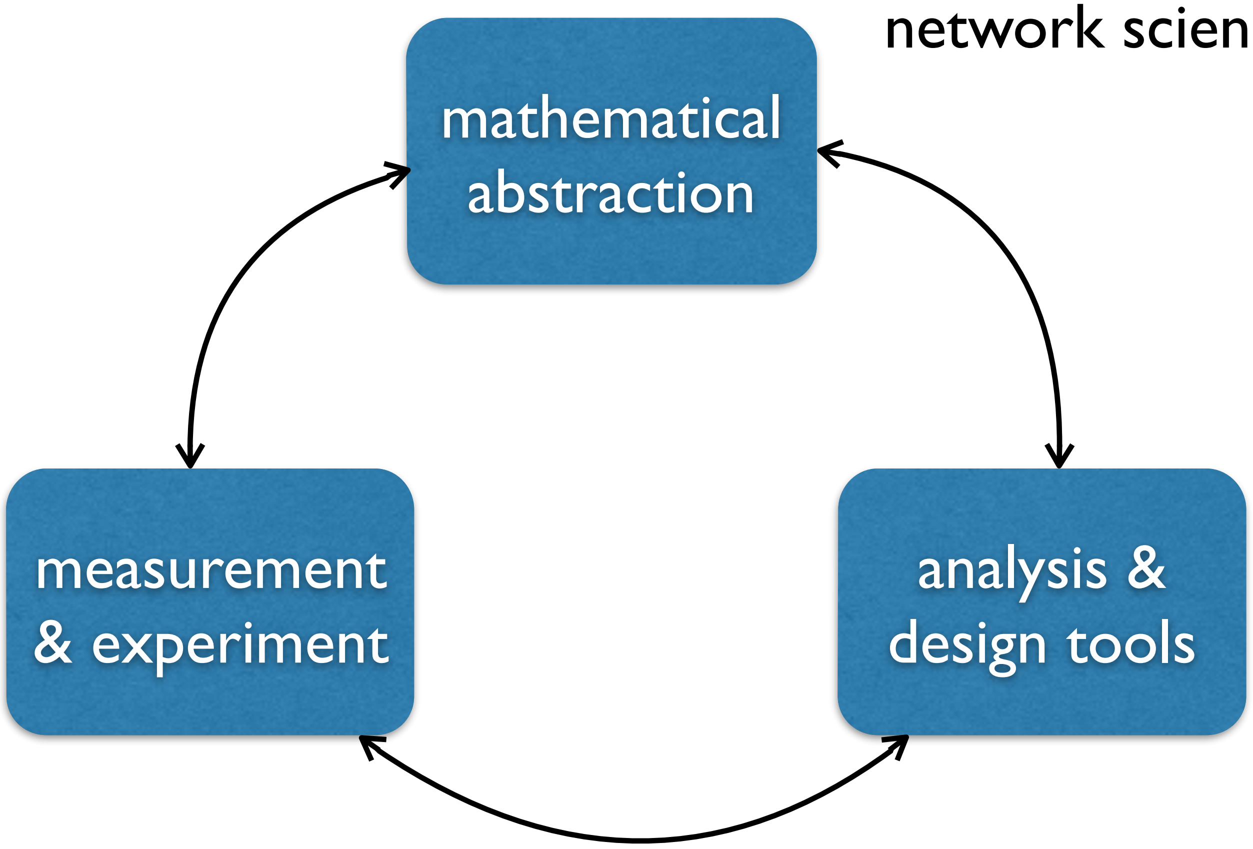

network science

data centers, distributed data storage processing

machine learning, game theory, modern optimization

What's Our Future?

- Data, data, data
- More tuning, stabilization, destabilization, optimization of social networks
    - Closing the feedback loop — more than targeted advertising

Who's Interested

- Organizational leadership (management consulting)
  - Employee engagement, flexible org-charts, team formation, leadership decision aids
- Advertisers
  - Sell before you want
- Political organizations
  - Polarize you, anger you, inspire you…
- Government Intelligence/Military
  - Anticipate crises, manipulate, detect manipulation

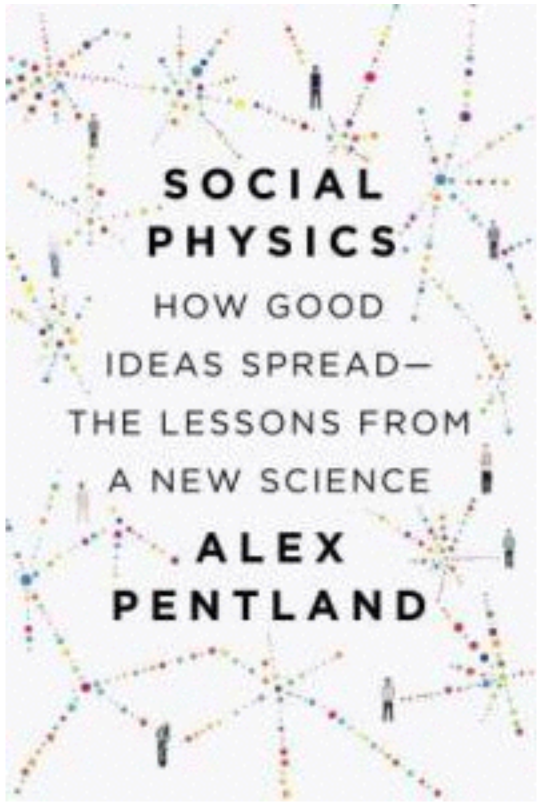

Return: Emerging SysEng View

Data trend from eToro Financial trading social net

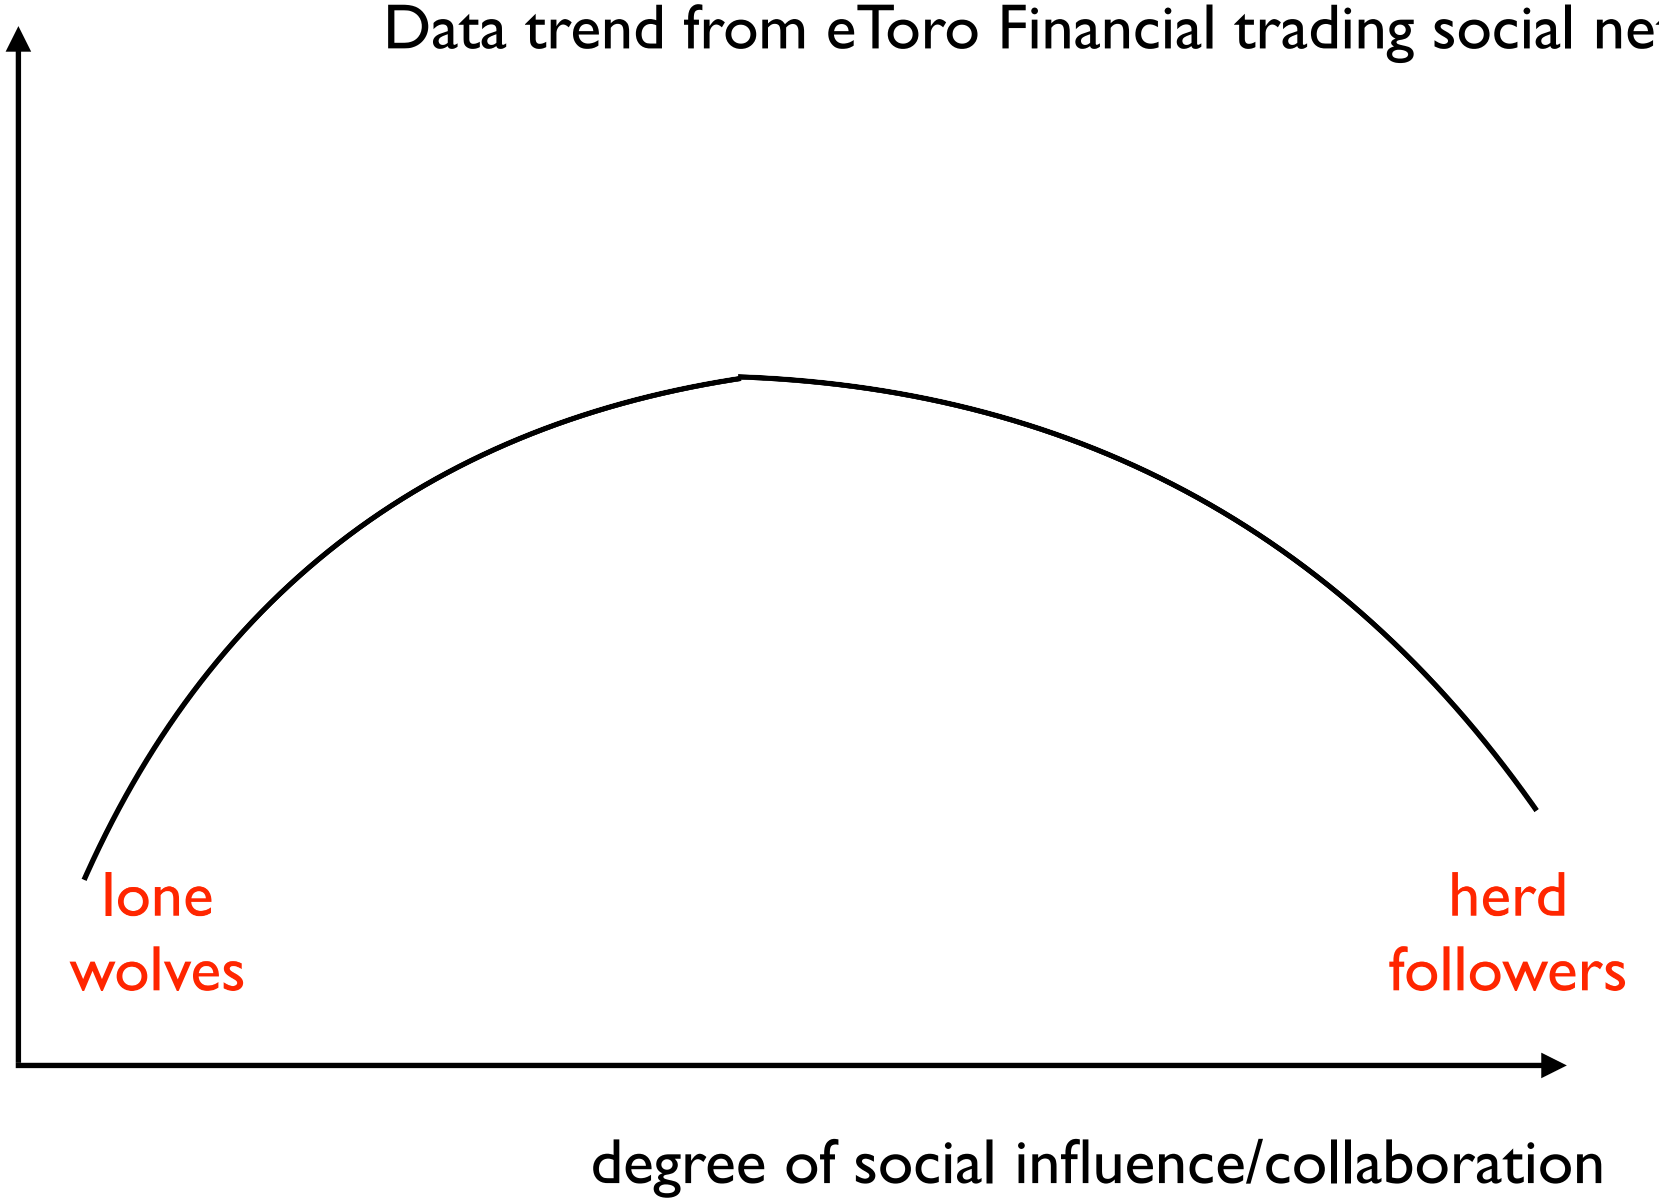

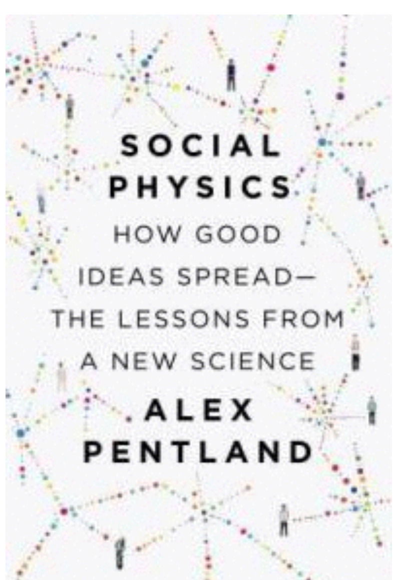

Figure 3

Return: Emerging SysEng View

Researchers Draw Romantic Insights From Maps of Facebook Networks

By STEVE LOHR

October 28, 2013 8:00 am

It's not in the stars after all. Instead, it seems, the shape of a person's social network is a powerful signal that can identify one's spouse or romantic partner — and even if a relationship is likely to break up.

So says a new research paper written by Jon Kleinberg, a computer scientist at Cornell University, and Lars Backstrom, a senior engineer at Facebook. The paper, posted online on Sunday, will be presented at a conference on social computing in February.

how do friend suggestion algorithms affect you individually and society more generally?

Why This Class Experiment?

- Emerging area that is here to stay
  - Expect it to be taught at BS, MS levels
  - Expect many opportunities in industry and academic research
- Is it advantageous to consider a systems engineering view?
  - Future offerings?

Overview

- My motivation and background
- Big picture overview in this context
- Class format
- Discussion