Operations Research: Opportunities and Challenges

Willem-Jan van Hoeve

Associate Professor
Tepper School of Business
Carnegie Mellon University

vanhoeve@andrew.cmu.edu
@wjvanhoeve

Outline

- What is Operations Research?
  – examples
- Future opportunities
  – due to technological and economic changes
- Challenges
  – scale, uncertainty, algorithms

Acknowledgement: Thanks go to Michael Trick for an earlier version of this presentation

What is Operations Research?

Operations Research: The analysis and optimization of business decisions using mathematical models

tions, the probability that the searcher detect the object of search during the time interval dt is $f\,dt$, where the coefficient f depends only on the distance r between the two, $f = f(r)$. If the object is fixed at the point (ξ, η) while the searcher moves at constant speed v in a path $x = x(t)$, $y = y(t)$, the probability that he detect the object during the time O to T is

$$1 - e^{-\int_0^T f(r)\,dt}$$

where

$$r^2 = [x(t) - \xi]^2 + [y(t) - \eta]^2,$$

and where the constraint condition of speed v is

$$x'(t)^2 + y'(t)^2 = v^2.$$

But if the position of the object is not given except in probability; in other words, if the probability that it lie within the infinitesimal rectangle bounded by ξ, $\xi + d\xi$, η and $\eta + d\eta$ is $p(\xi, \eta)\,d\xi\,d\eta$, where the "probability density" $p(\xi, \eta)$ is known, then the probability of finding the object is given by the expression

$$\iint [1 - e^{-\int_0^T f(r)\,dt}]\,p(\xi, \eta)\,d\xi\,d\eta.$$

What is Operations Research?

Role of Operations Research

CarnegieMellon Tepper SCHOOL OF BUSINESS

- Data is not information; information is not improved decision making
- Operations research allows companies to transform data into better decision making
- New roles:
  - Competitive advantage
  - Business Opportunity
  - Unlocking the value of information

Why Now?

- Trends are moving towards more operations research
  - Increased Data
  - Faster Computers
  - Better Algorithms
  - Lower Fixed Cost for Optimization
  - Service Applications

Increased Data

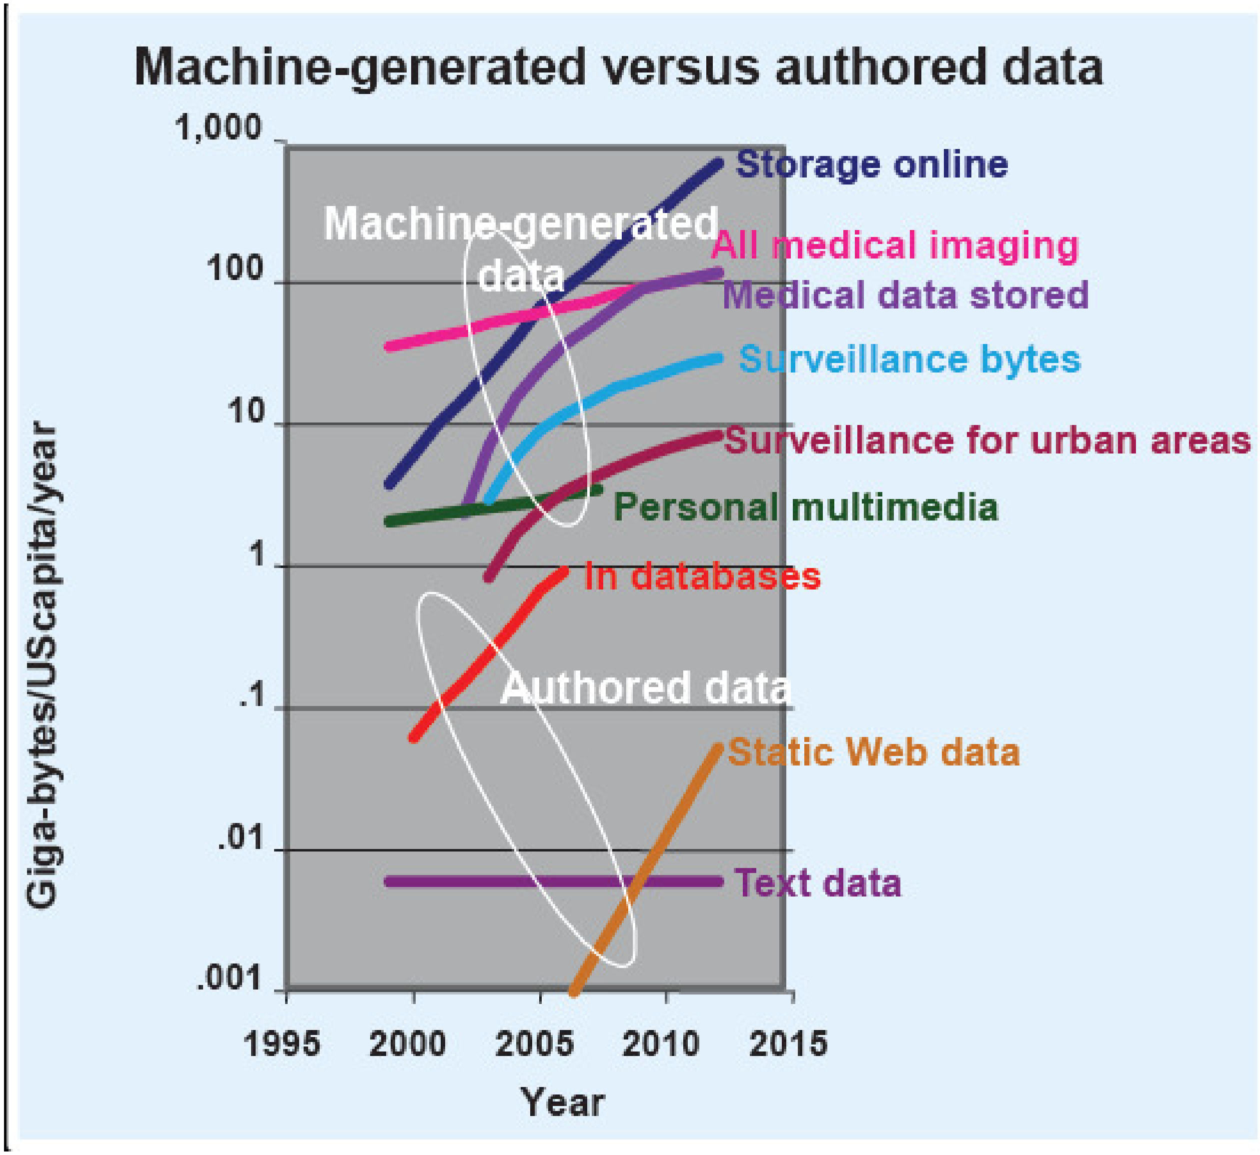

Faster Computers: Supercomputers and otherwise

16 TFlops now costs US$3,000

(AMD 7970 desktop computer)

www.top500.org

Algorithms are getting better also!

- Evolution of Linear Programming solver CPLEX

Table 5: **PDS models–Solution times**

Instance	CPLEX 1.0	CPLEX 5.0 Dual	CPLEX 7.1 Primal	CPLEX 7.1 Dual
pds100	–	50413.1	2414.8	256.3
pds90	–	59981.0	2452.2	320.3
pds80	–	42055.4	2201.5	304.4
pds70	335292.1	21120.4	1504.1	197.8
pds60	205798.3	7442.6	852.4	160.5
pds50	122195.9	8509.9	493.2	114.6
pds40	58920.3	2816.8	188.3	79.3
pds30	15891.9	1154.9	74.8	39.1
pds20	5168.8	232.6	27.9	20.9
pds10	208.9	13.0	3.7	2.6
pds06	26.4	2.4	1.4	0.9
pds02	0.4	0.1	0.1	0.1

← 546,469 variables
156,171 constraints

[Bixby, 2002]

Speed

- The combination of improved algorithms and faster computers mean that many problems can now be solved 1 million times faster than they could 10 years ago
  - solve in few seconds instead of days
- Huge increase in applicability of optimization methods
- Optimization methods can now often be applied in real-time

Lower Fixed Cost for Optimization

- Past. Big projects for big companies
  - Fighting World War II
  - Airline Crew Scheduling
  - Material Planning at Ford
- Current. Everywhere
  - Much more accessible
  - Optimization software on your computer (e.g., Solver in Excel)

New Areas of Application

- Traditional view of OR:
  - Manufacturing
  - Services limited to transportation and logistics

- New view
  - Operations Research everywhere
  - Services are a great opportunity

Rise of service economy 1800-2010

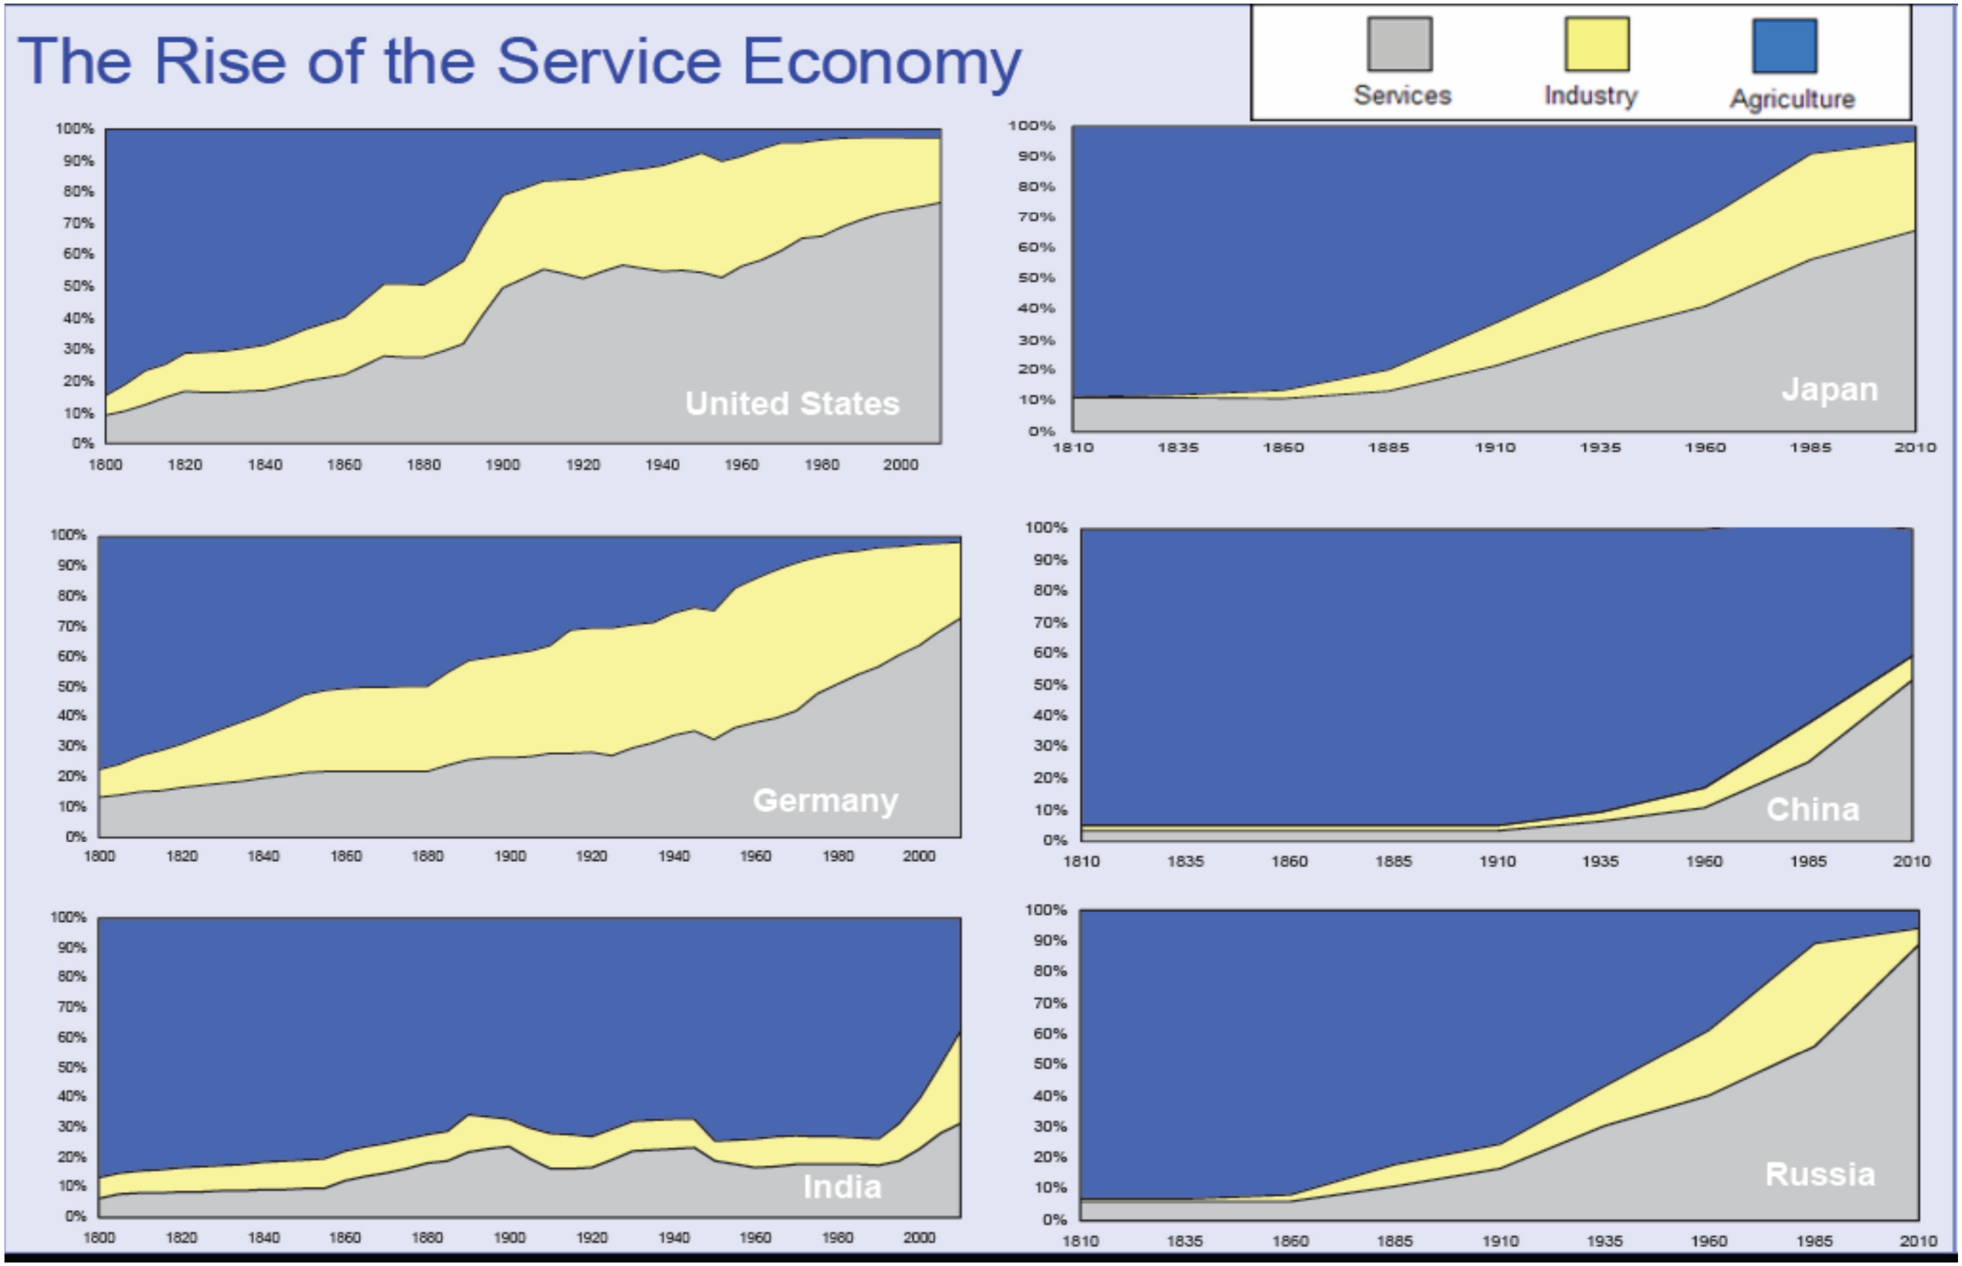

Specific OR Opportunities

- Healthcare
  - cancer treatment planning, (influenza) vaccine composition, medical staff planning, in-home health care, …
- Financial services
  - portfolio optimization (with side constraints such as transaction costs, shared budget, …), balance risk and profit, credit score estimation, …
- Energy market
  - smart grid, pricing, strategic inventory, …
- Mass media
  - next generation advertising, product marketing, political campaigning, …

Example: Ad placement

Carnegie Mellon Tepper School of Business

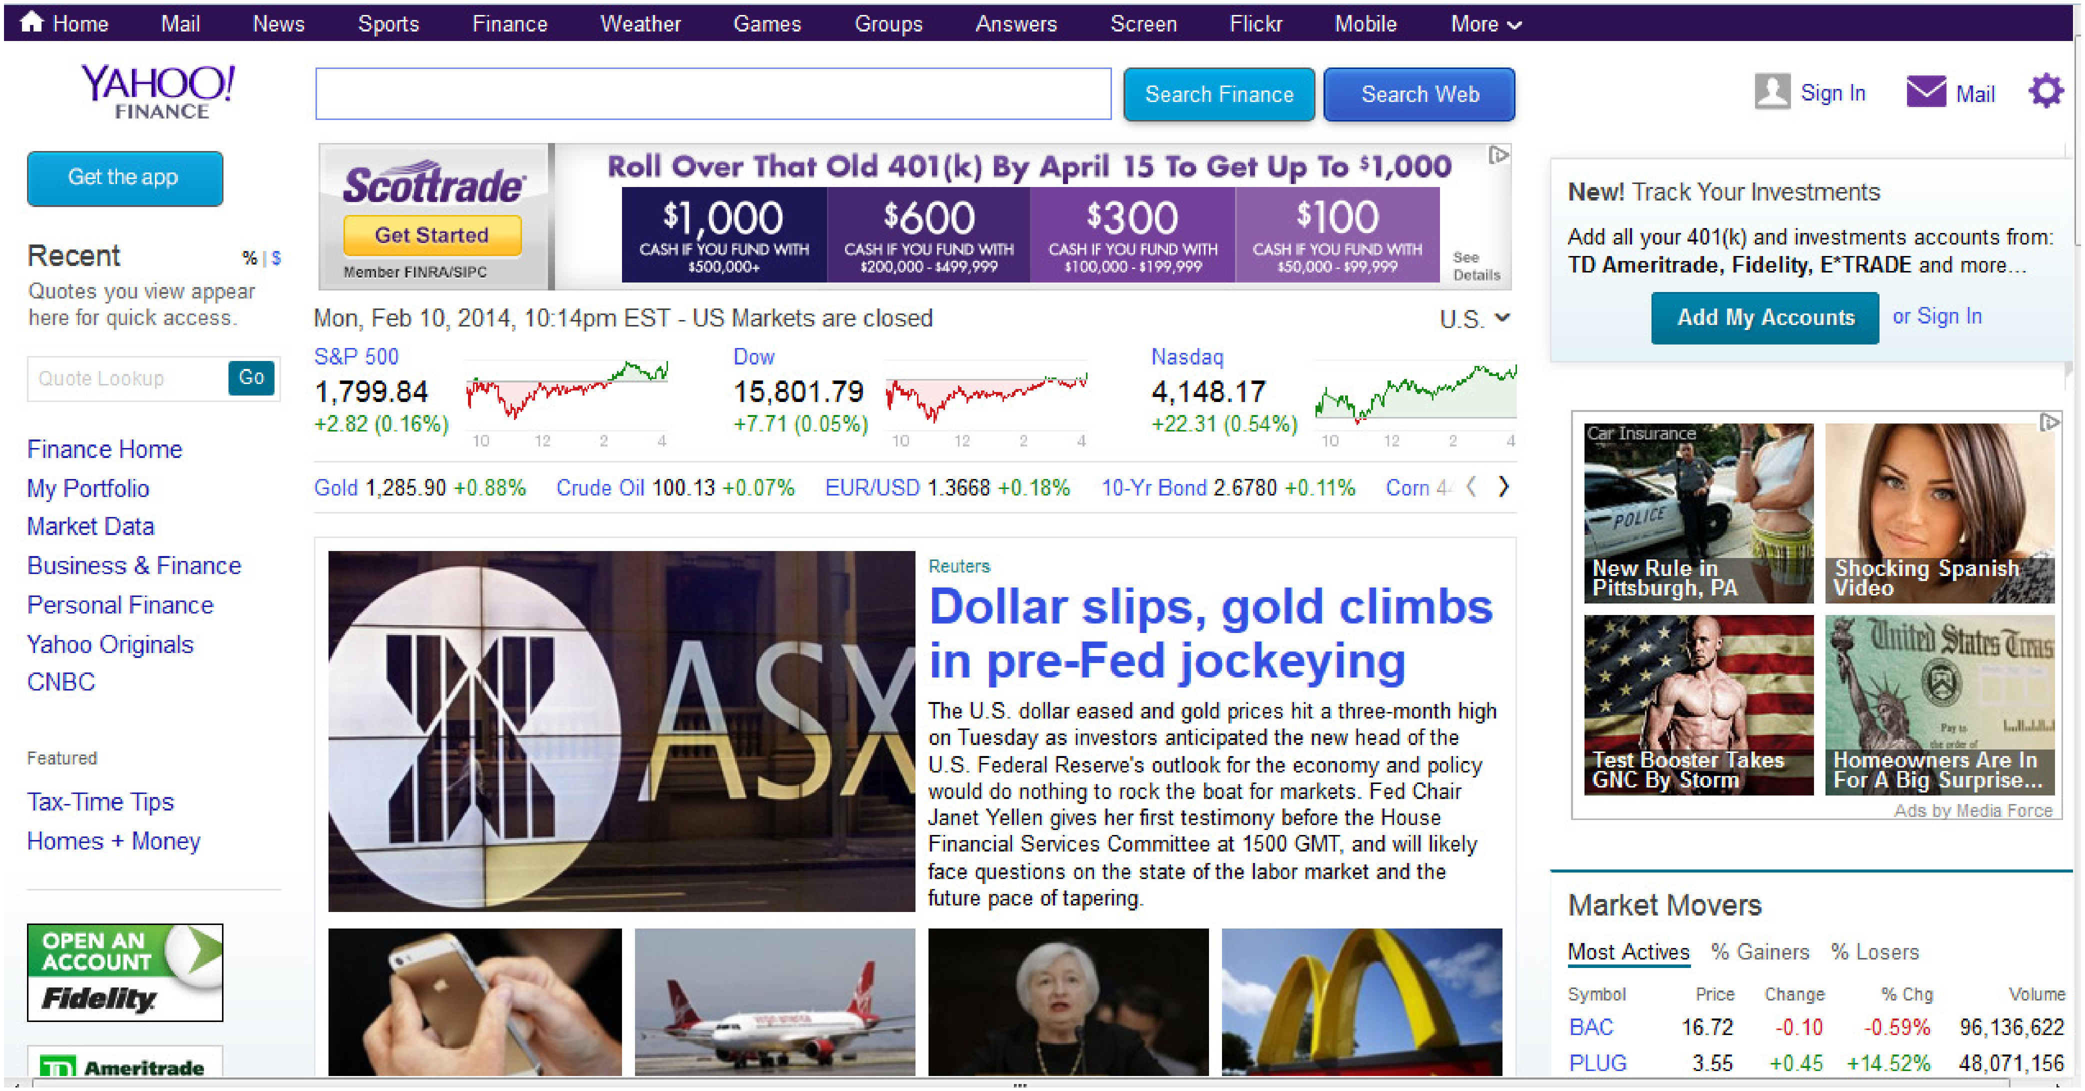

Why Now?

- Trends are moving towards more operations research
  - Increased Data
  - Faster Computers
  - Better Algorithms
  - Lower Fixed Cost for Optimization
  - Service Applications

…not present 10~15 years ago

Impact and Challenges

Carnegie Mellon Tepper SCHOOL OF BUSINESS

- Bright future for OR practitioners…
  - many job opportunities, new application areas

BEST JOBS
US News
BUSINESS
2014

#2 best business job

Best Business Jobs

Operations Research Analyst

Job Profile

Overview
Salary
Reviews & Advice
Job Listings

Overall Score ★★★★★
(7.1 out of 10)

Number of Jobs
19,500

Median Salary
$72,100

Unemployment Rate
1.3 percent

Show Jobs Near: 15217 Find Jobs

Mean Salaries Shown

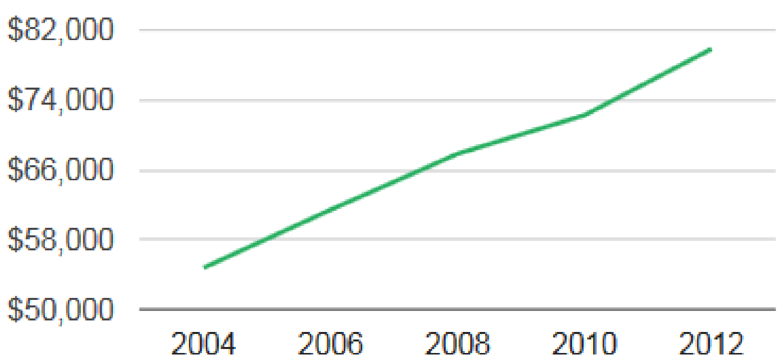

Operations research analysts are high-level problem-solvers who use advanced techniques, such as optimization, data mining, statistical analysis and mathematical modeling, to develop solutions that help businesses and organizations operate more efficiently and cost-effectively.

This Job is Ranked in

Best Business Jobs	#2
The 100 Best Jobs	#23

Impact and Challenges

- Bright future for OR practitioners…
  - – many job opportunities, new application areas
- … and for researchers!

- new application areas
- large scale problems
- real-time application
- data driven
- handle uncertainty

- different requirements
- new methodologies
- new algorithms
- combine technologies
- more funding

Challenge 1: Scale

- Graph theory is one of the corner stones of operations research (e.g., networks)
- Up to recently, many theoretic results involved the relationship between specific subgraphs
  - Specify (family of) graphs by forbidden substructures
  - Example: the *Petersen graph* is the smallest bridgeless cubic graph with no three-edge-coloring

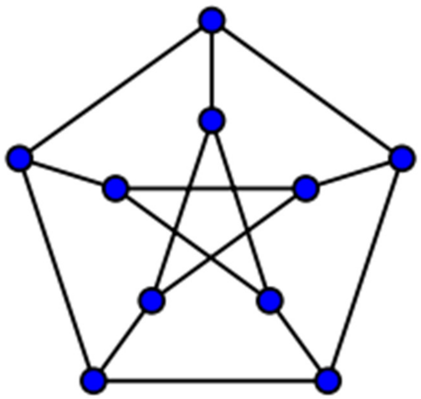

- Small graphs, can draw on piece of paper

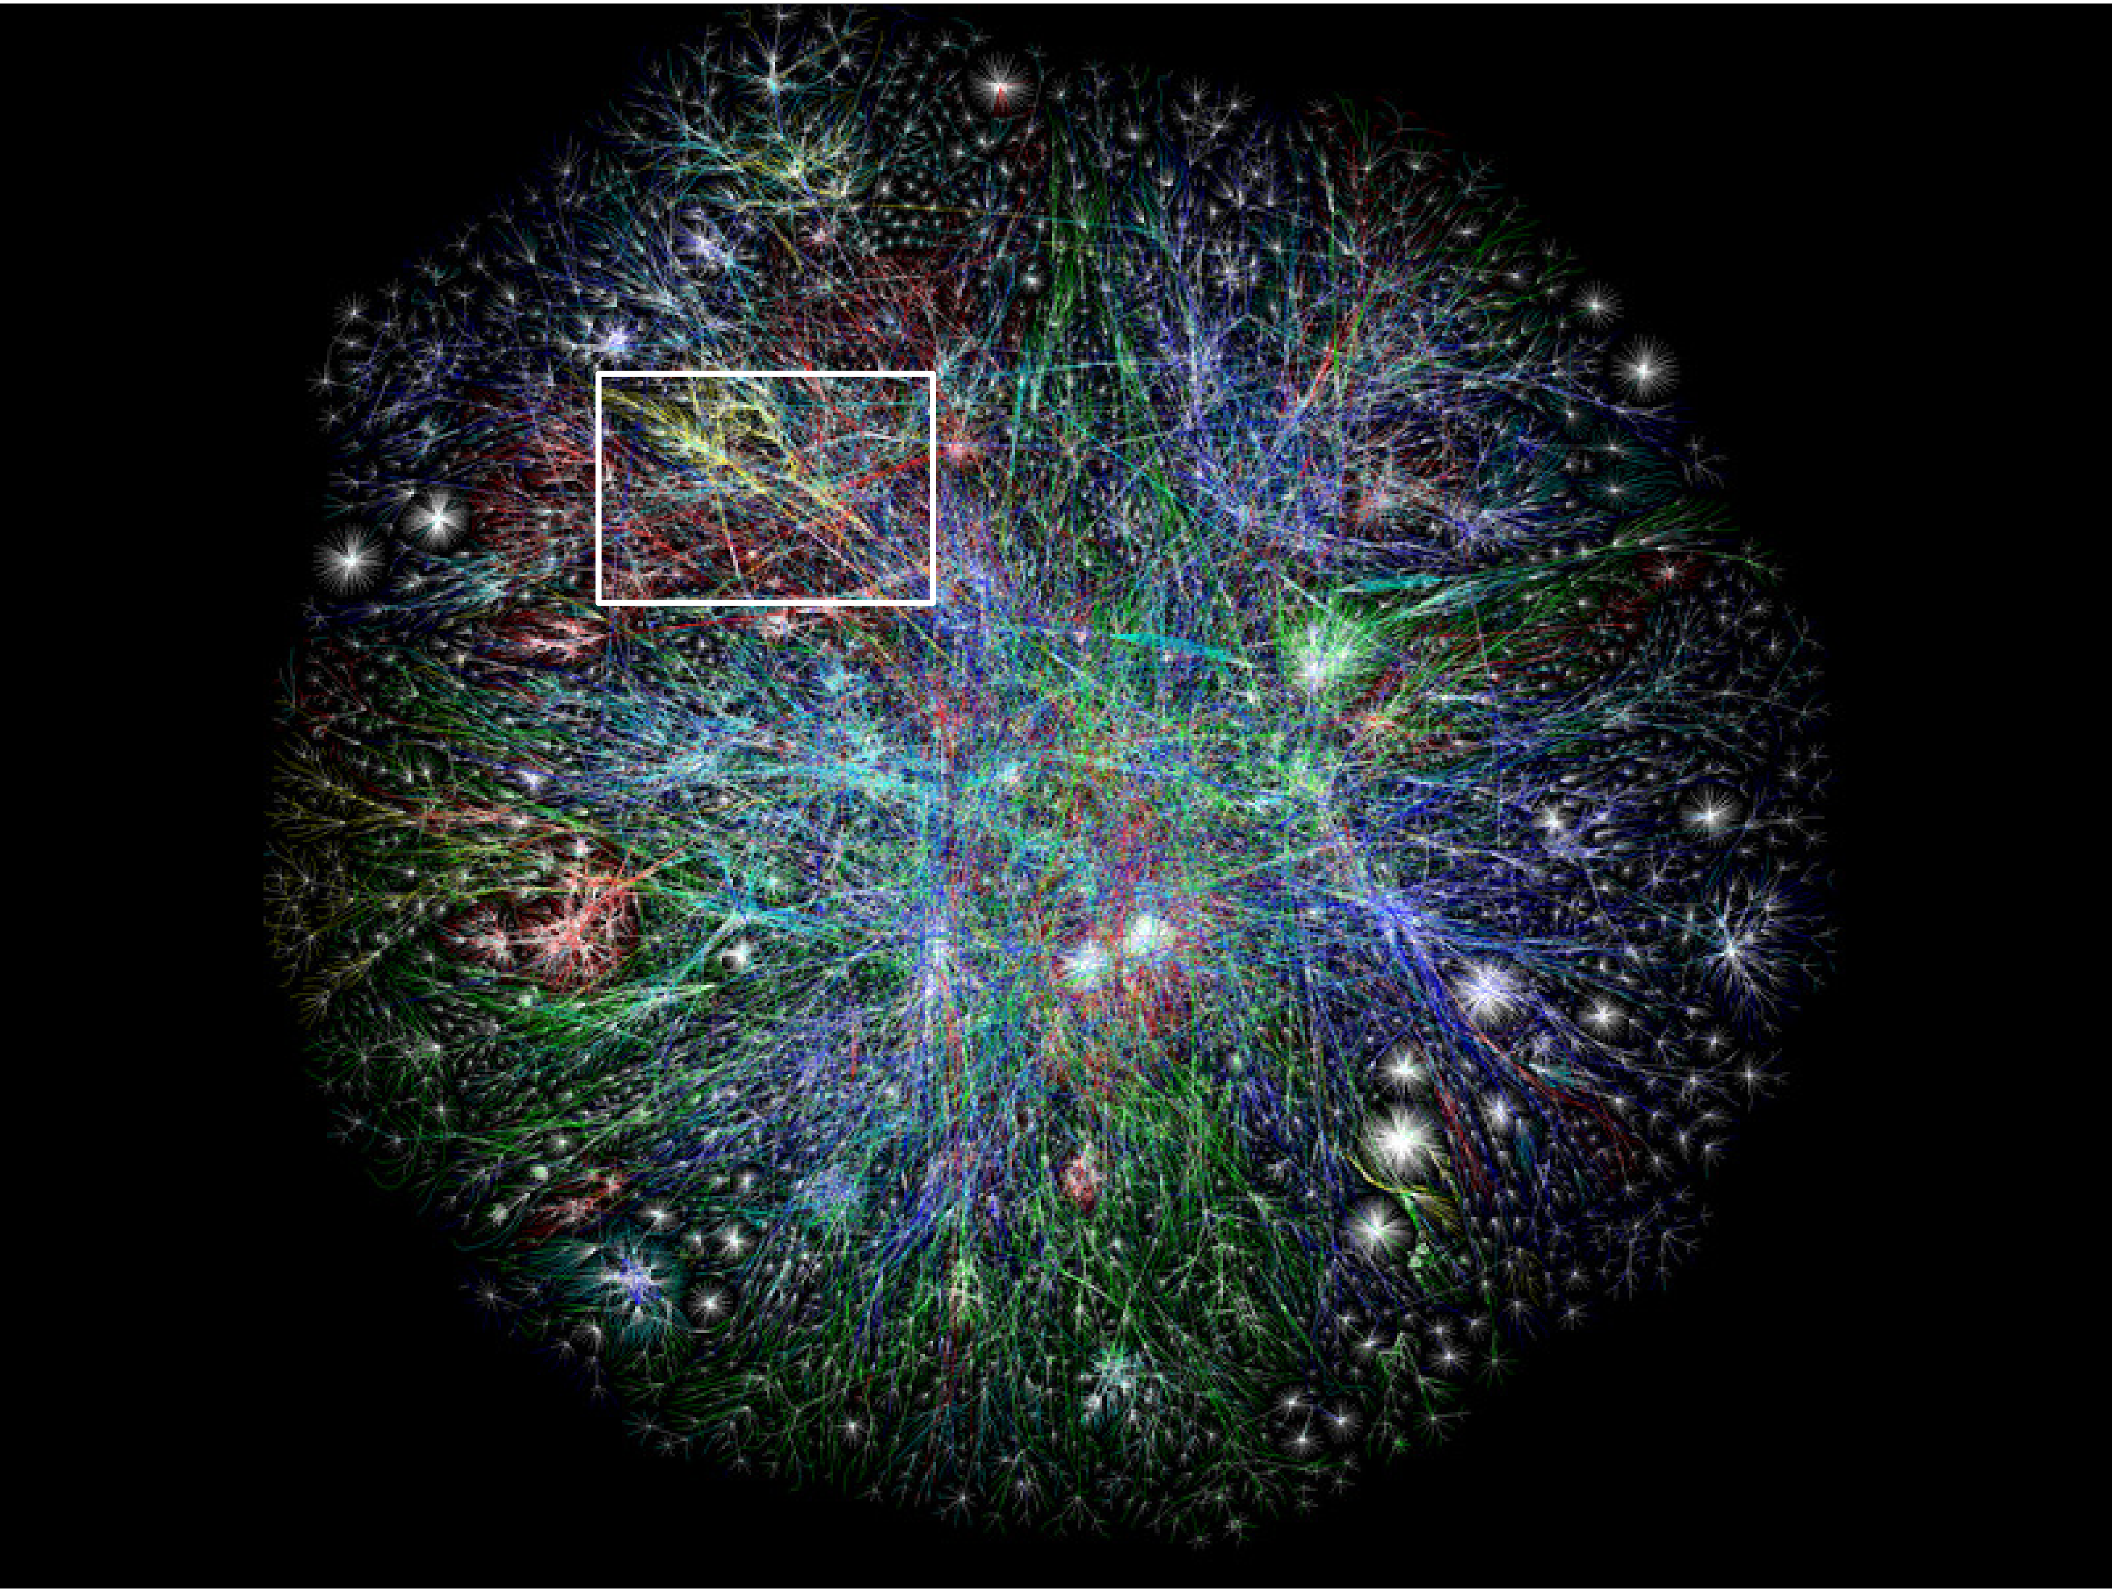

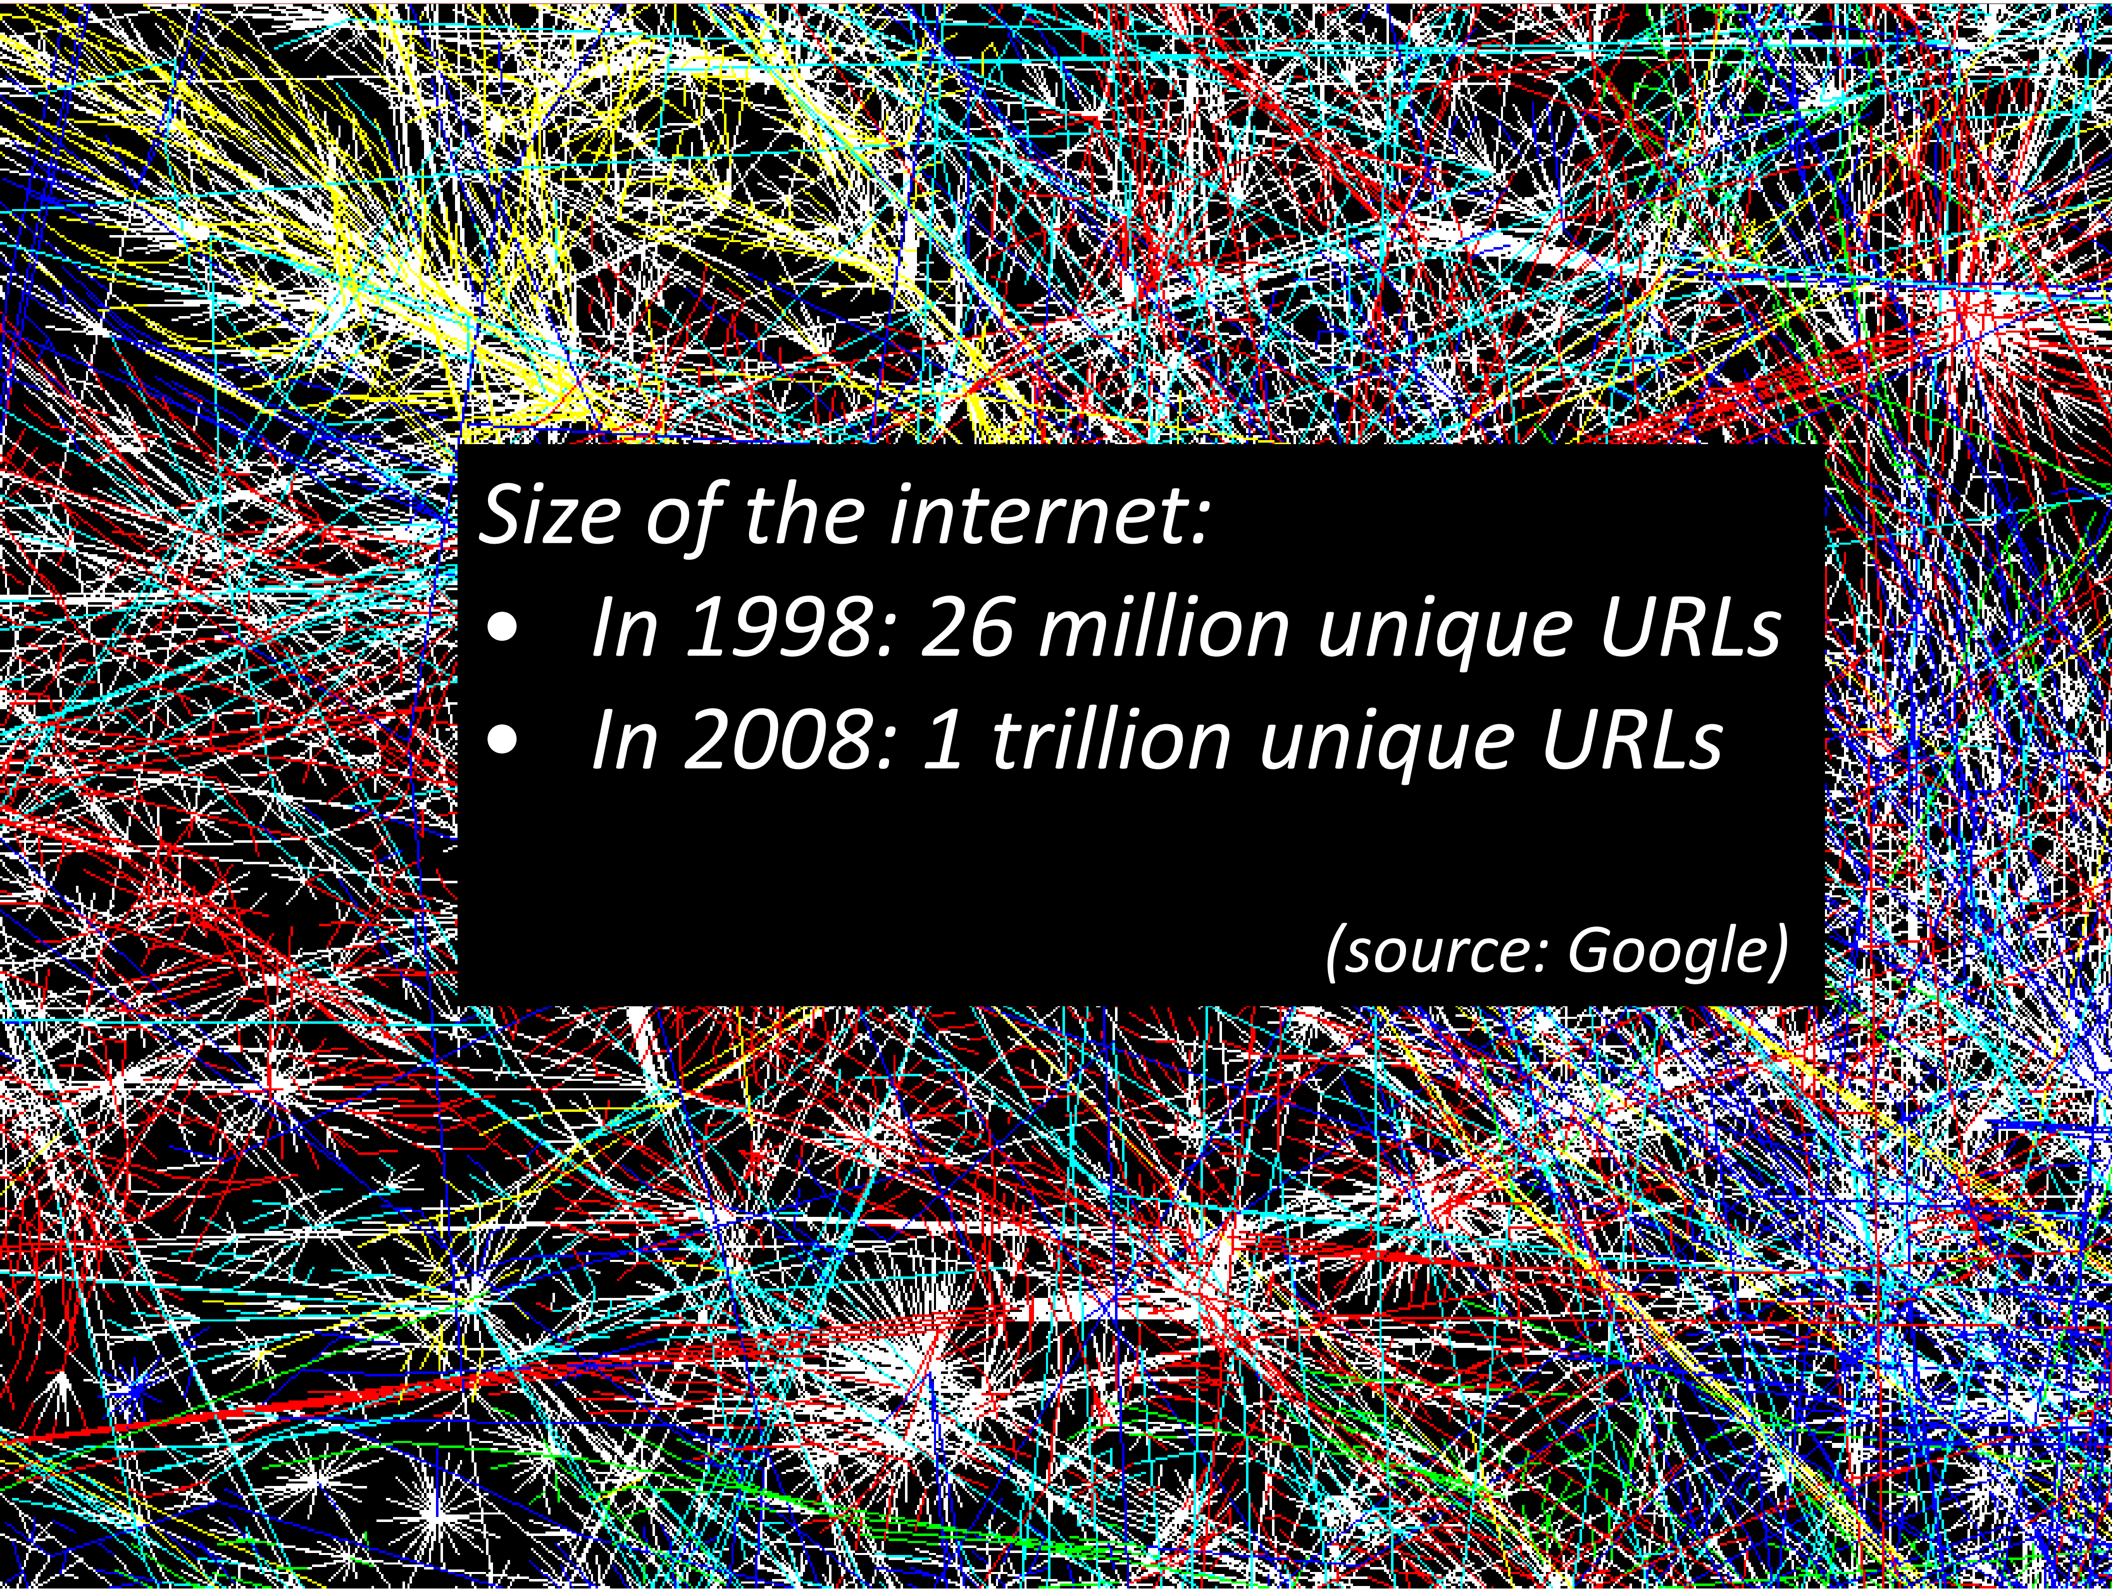

Size of the internet:

- *In 1998: 26 million unique URLs*
- *In 2008: 1 trillion unique URLs*

(source: Google)

Large scale optimization

- Huge networks
  - internet-based applications, e.g., advertising via facebook or twitter network
- Different questions can be asked
  - importance of nodes in a network
  - connectedness; small-world phenomenon
- Different graph-theoretic approaches
  - ‘social networks’ very active research area
- Other large-scale applications
  - routing applications, client management for service industry; complex supply chains, …

Challenge 2: Uncertainty

- New applications will demand better handling of uncertainty
- Example application: Vehicle routing

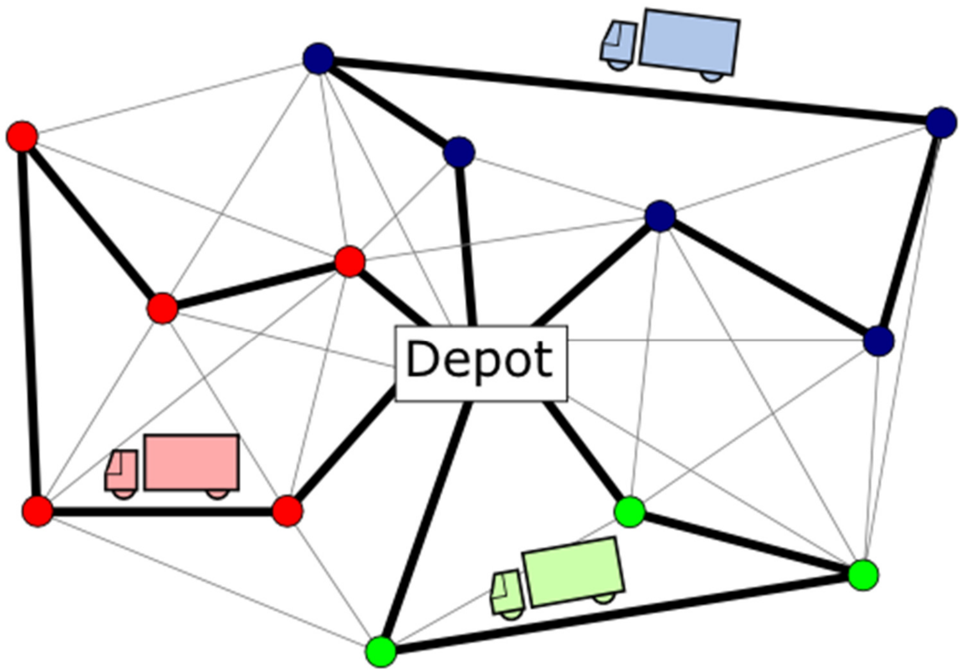

Vehicle Routing Approaches

Many approaches in the literature assume

- One commodity needs to be picked up or delivered
- Client demand is deterministic
- Vehicles are uniform and 1-dimensional
- Distances are given and fixed

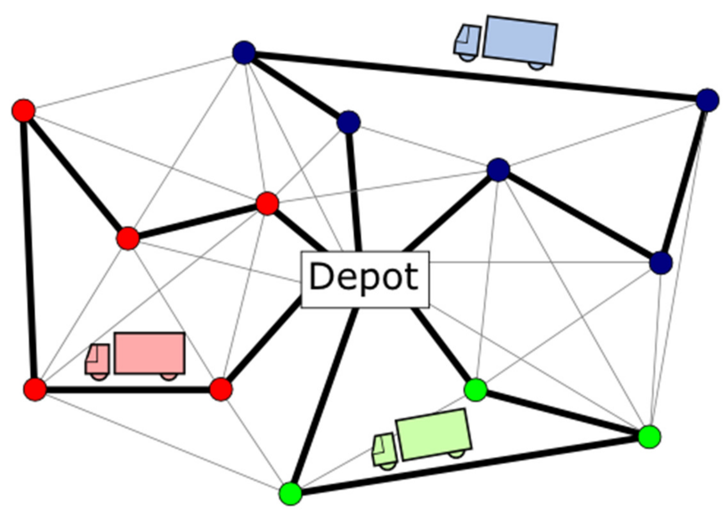

In practice, usually *none* of these assumptions applies

- Moreover, there are often side constraints (time windows, stacking conditions, …)

Next generation vehicle routing

- Need to integrate traditional optimization with data mining, statistics, forecasting, …
  - Use historical data to build demand distributions
  - Cluster clients together such that trucks do not overload with p=0.95

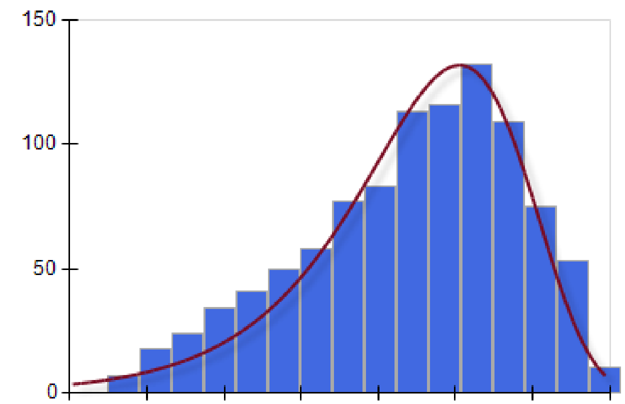

- Assume uncertain distances (again, can be based on historic traffic data)
- Try to accommodate all side constraints
- We are still far from optimal solutions…

Challenge 3: Algorithms

- Previously known 'good' algorithms may no longer be applicable

- Example: *Minimum spanning tree* for graph G = (V,E)
  - Prim: O(|E| + |V| log |V|)
  - Kruskal: O(|E| log |V|)

- Graph on 1M nodes?

- Note: often used as subroutine

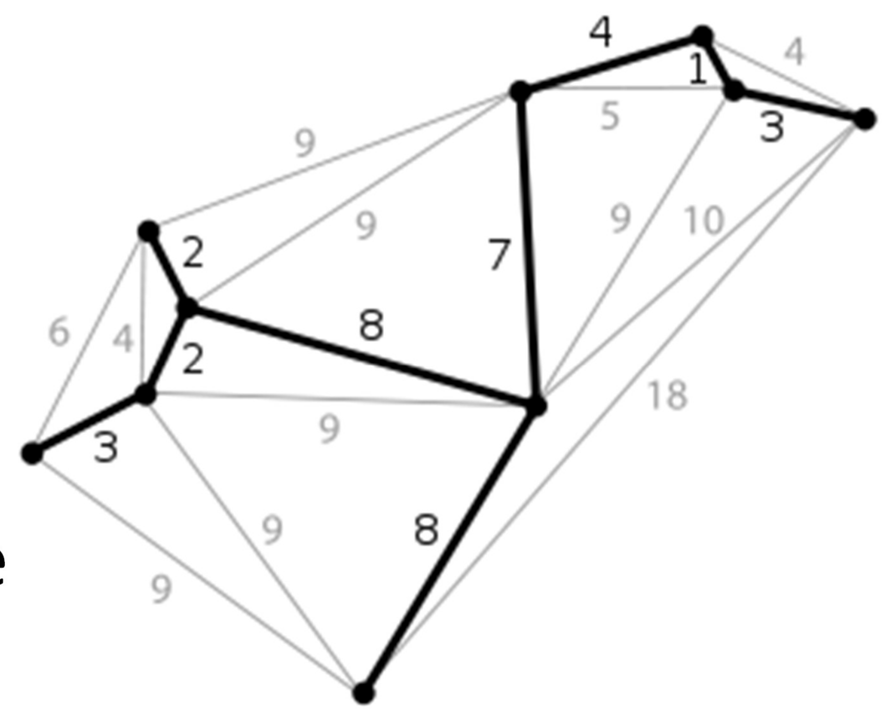

Traveling Salesman Problem

Find the shortest closed tour that visits each city exactly once

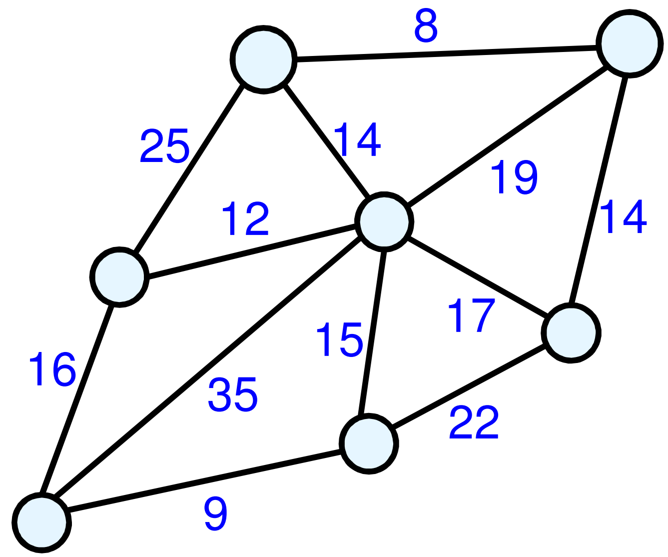

Applications:

- Truck routing
- Electronic circuit design
- Genome sequencing
- Parcel delivery services
- Robotic arm movement planning
- and many, many more

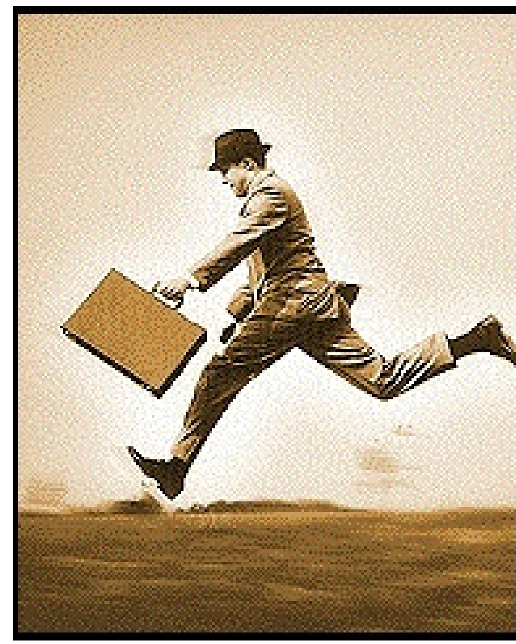

Hard to Solve Problem

For n locations, there are (n-1)! possible routes

Example:

n= 5	n! = 120
n=10	n! = 3 628 800
n= 20	n! = 2.43×10^{18}
n= 40	n! = 8.15×10^{47}
n= 60	n! = 8.32×10^{81}
	this is more than the total number of atoms in the observable universe! (estimated to be around 10^{80})

71,009 Cities in China

Carnegie Mellon Tepper SCHOOL OF BUSINESS

http://www.tsp.gatech.edu

Milestones

Year	Research Team	Size of Instance
1954	G. Dantzig, R. Fulkerson, and S. Johnson	49 cities
1971	M. Held and R.M. Karp	64 cities
1975	P.M. Camerini, L. Fratta, and F. Maffioli	67 cities
1977	M. Grötschel	120 cities
1980	H. Crowder and M.W. Padberg	318 cities
1987	M. Padberg and G. Rinaldi	532 cities
1987	M. Grötschel and O. Holland	666 cities
1987	M. Padberg and G. Rinaldi	2,392 cities
1994	D. Applegate, R. Bixby, V. Chvátal, and W. Cook	7,397 cities
1998	D. Applegate, R. Bixby, V. Chvátal, and W. Cook	13,509 cities
2001	D. Applegate, R. Bixby, V. Chvátal, and W. Cook	15,112 cities
2004	D. Applegate, R. Bixby, V. Chvátal, W. Cook, and K. Helsgaun	24,978 cities

2005 Applegate et al. 85,900 cities

chip design application for AT&T/Bell Labs, solved to optimality in 136 CPU years

(on a 250-node cluster this took around one year)

Algorithmic Advances

- Integer linear programming model
  - binary variable x_e for each edge e to represent tour
  - constraints to ensure that there are no sub-tours
  - minimize $\sum_e d_e x_e$
- Challenge:
  - huge number of edges and constraints

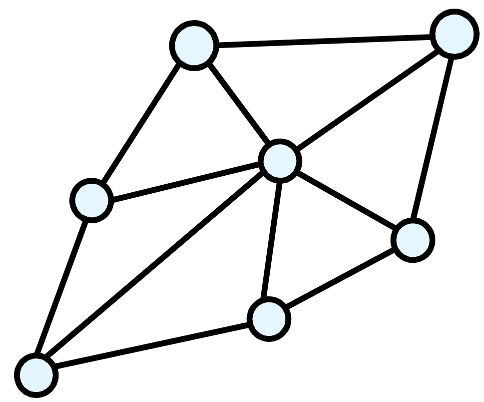

- Remedy: problem decomposition
  - variable generation based on marginal cost of edge
  - constraint generation for detected subtours, integrality, …
  - heuristic solutions to find upper bounds
  - optimality is still guaranteed, with fraction of full model

China TSP revisited

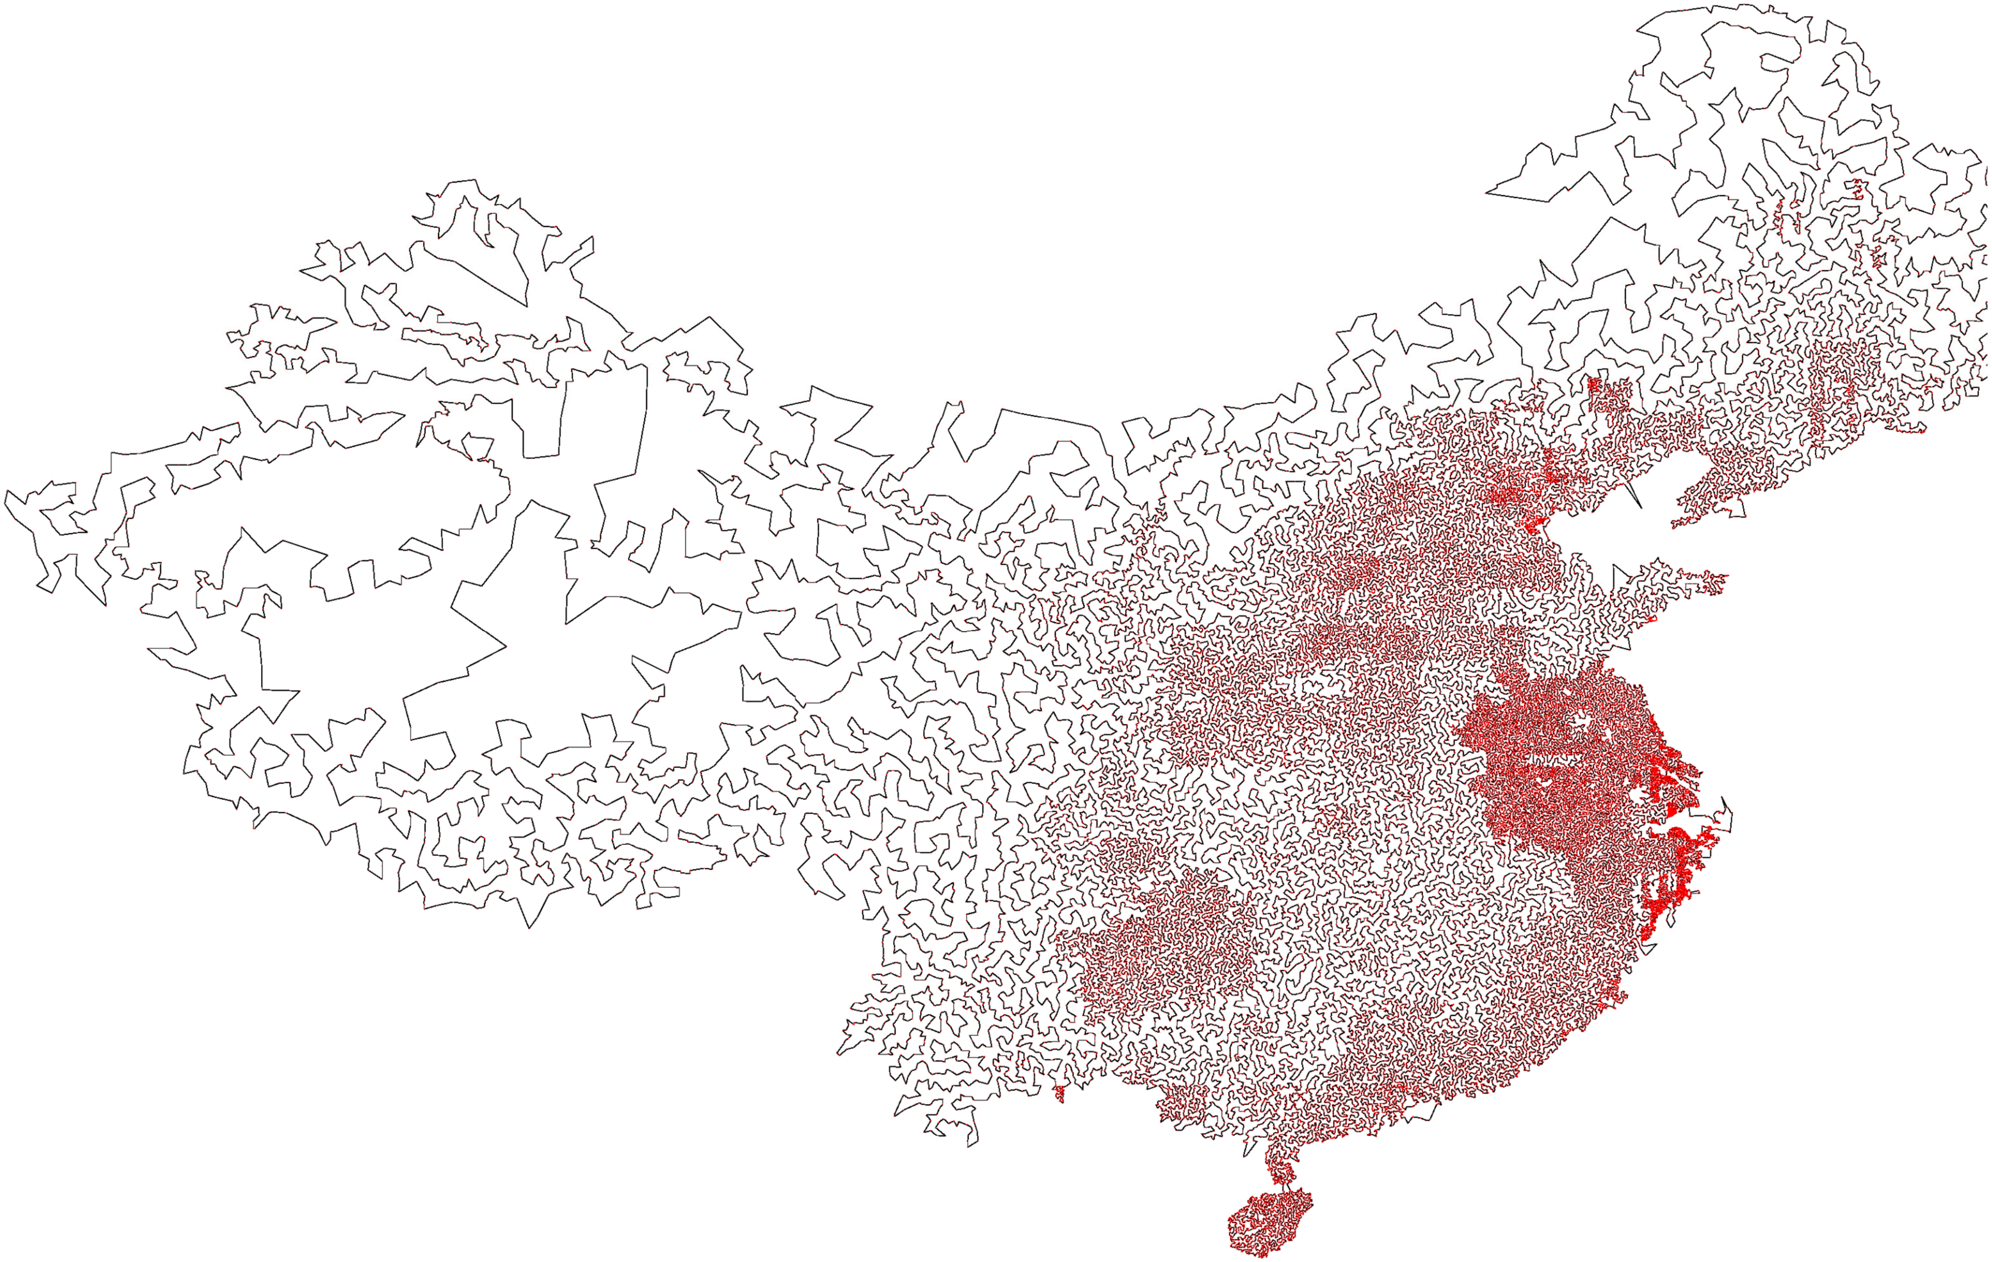

Optimal tour [Hung Dinh Nguyen]

Summary

- What is Operations Research?
  - examples
- Future opportunities
  - due to technological and economic changes
- Challenges
  - scale, uncertainty, algorithms

https://www.informs.org/