Adventures in Scheduling: Some Trends in Operations Research

IJCAI 2011

Michael Trick

Tepper School of Business, Carnegie Mellon University

Adventures in Scheduling: Some Trends in Operations Research

And Artificial Intelligence!

IJCAI 2011

Michael Trick

Tepper School of Business, Carnegie Mellon University


PNC PARK
BUD LIGHT

The History Starts in 1995

History

- Call from Doug Bureman (former Senior VP of the Pirates): Would you like to work at scheduling Major League Baseball?

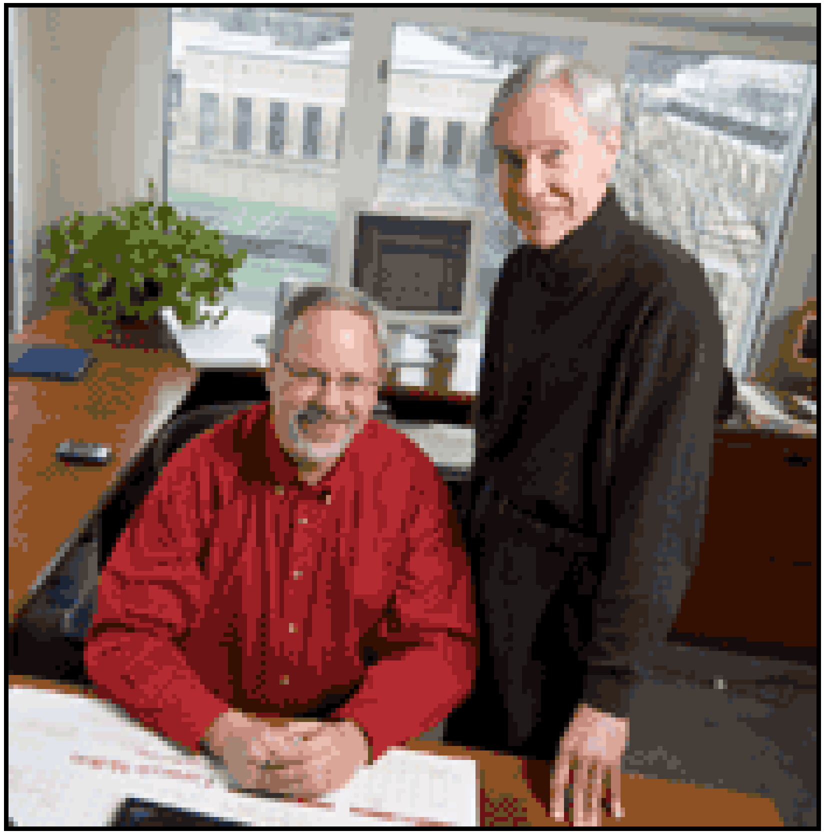

The Competition

The Incumbents:

Henry and Holly Stephenson of Long Island, New York

Schedulers since 1981

Why is this interesting?

- Huge schedule: 30 teams, 162 games each over 182 day season

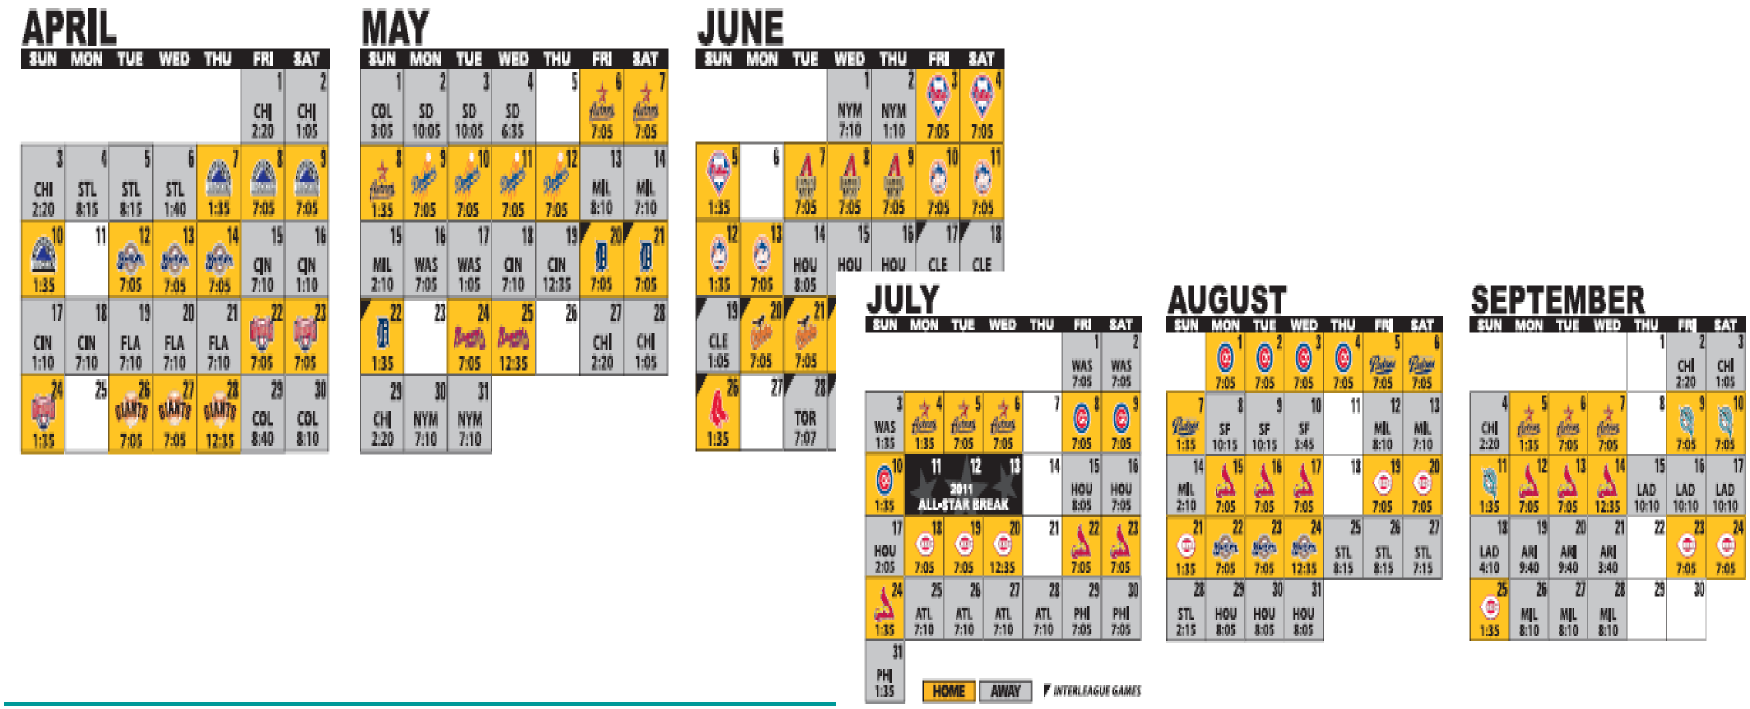

APRIL

SUN	MON	TUE	WED	THU	FRI	SAT
					1 CHI 2:20	2 CHI 1:05
3 CHI 2:20	4 STL 8:15	5 STL 8:15	6 STL 1:40	7 1:35	8 7:05	9 7:05
10 1:35	11	12 7:05	13 7:05	14 7:05	15 CIN 7:10	16 CIN 1:10
17 CIN 1:10	18 CIN 7:10	19 FLA 7:10	20 FLA 7:10	21 FLA 7:10	22 7:05	23 7:05
24 1:35	25	26 7:05	27 7:05	28 12:35	29 COL 8:40	30 COL 8:10

MAY

SUN	MON	TUE	WED	THU	FRI	SAT
1 COL 3:05	2 SD 10:05	3 SD 10:05	4 SD 6:35	5	6 7:05	7 7:05
8 1:35	9 7:05	10 7:05	11 7:05	12 7:05	13 MIL 8:10	14 MIL 7:10
15 MIL 2:10	16 WAS 7:05	17 WAS 1:05	18 CIN 7:10	19 CIN 12:35	20 7:05	21 7:05
22 1:35	23	24 7:05	25 12:35	26	27 CHI 2:20	28 CHI 1:05
29 CHI 2:20	30 NYM 7:10	31 NYM 7:10				

JUNE

SUN	MON	TUE	WED	THU	FRI	SAT
			1 NYM 7:10	2 NYM 1:10	3 7:05	4 7:05
5 1:35	6	7 7:05	8 7:05	9 7:05	10 7:05	11 7:05
12 1:35	13 7:05	14 HOU 8:05	15 HOU 8:05	16 HOU 2:05	17 CLE 7:05	18 CLE 7:05
19 CLE 1:05	20 7:05	21 7:05	22 12:35	23	24 7:05	25 7:05
26 1:35	27	28 TOR 7:07	29 TOR 7:07	30 TOR 7:07		

Interesting fomulation

- Lots of constraints to handle
  - Home half the time
  - Home half weekends
  - Team requests
  - Etc., etc.
- Interesting Objective(s)
  - Travel distance (since coast to coast in US)
  - Evenness of Flow (Home/Away sequencing)

Important problem

- The least valuable team in MLB (the Pittsburgh Pirates) has a valuation of $304 million; the most (the New York Yankees) is at $1.7 billion
- Total attendance: 100 million/year; many more watch on TV/internet
- Schedule does affect attendance and league fairness

Computationally Difficult Problem

- MLB “instance” is one/year consisting of pages and pages of requirements/requests/wishes
- Traveling Tournament Problem abstracts out two key issues of travel and length of home stands/road trips

Traveling Tournament Problem

Given an *n* by *n* distance matrix *D= [d(i,j)]* and an integer *k* find a double round robin (every team plays at every other team) schedule such that:

- The total distance traveled by the teams is minimized (teams are assumed to start at home and must return home at the end of the tournament), and
- No team is away more than *k* consecutive games, or home more than *k* consecutive games.
- (no repeat) If i is at j in slot t then j is not at i in slot t+1

Sample Instance

NL6: Six teams from the National League of (American) Major League Baseball. Distances:

0	745	665	929	605	521
745	0	80	337	1090	315
665	80	0	380	1020	257
929	337	380	0	1380	408
605	1090	1020	1380	0	1010
521	315	257	408	1010	0

k is 3

Sample Solution

Distance: 23916 (Easton May 7, 1999)

Slot	ATL	NYM	PHI	MON	FLA	PIT
0	FLA	@PIT	@MON	PHI	@ATL	NYM
1	NYM	@ATL	FLA	@PIT	@PHI	MON
2	PIT	@FLA	MON	@PHI	NYM	@ATL
3	@PHI	MON	ATL	@NYM	PIT	@FLA
4	@MON	FLA	@PIT	ATL	@NYM	PHI
5	@PIT	@PHI	NYM	FLA	@MON	ATL
6	PHI	@MON	@ATL	NYM	@PIT	FLA
7	MON	PIT	@FLA	@ATL	PHI	@NYM
8	@NYM	ATL	PIT	@FLA	MON	@PHI
9	@FLA	PHI	@NYM	PIT	ATL	@MON

Simple Problem, yes?

NL12. 12 teams

Feasible Solution: 143655 (Rottembourg and Laburthe May 2001), 138850 (Larichi, Lapierre, and Laporte July 8 2002), 125803 (Cardemil, July 2 2002), 119990 (Dorrepaal July 16, 2002), 119012 (Zhang, August 19 2002), 118955 (Cardemil, November 1 2002), 114153 (Anagnostopoulos, Michel, Van Hentenryck and Vergados January 14, 2003), 113090 (Anagnostopoulos, Michel, Van Hentenryck and Vergados February 26, 2003), 112800 (Anagnostopoulos, Michel, Van Hentenryck and Vergados June 26, 2003), 112684 (Langford February 16, 2004), 112549 (Langford February 27, 2004), 112298 (Langford March 12, 2004), 111248 (Anagnostopoulos, Michel, Van Hentenryck and Vergados May 13, 2004), 110729 (Van Hentenryck and Vergados, May 30 2007).

Lower Bound: 107483 (Waalewign August 2001), 107494 (Melo, Ribeiro, and Urrutia July 15 2006), 107548 (Mitchell, Trick and Waterer July 31 2008), 108244 (Uthus, Riddle, and Guesgen, Feb 11 2009), 108629 (Uthus, Riddle, and Guesgen January 6, 2010)

10 years of progress on NL12

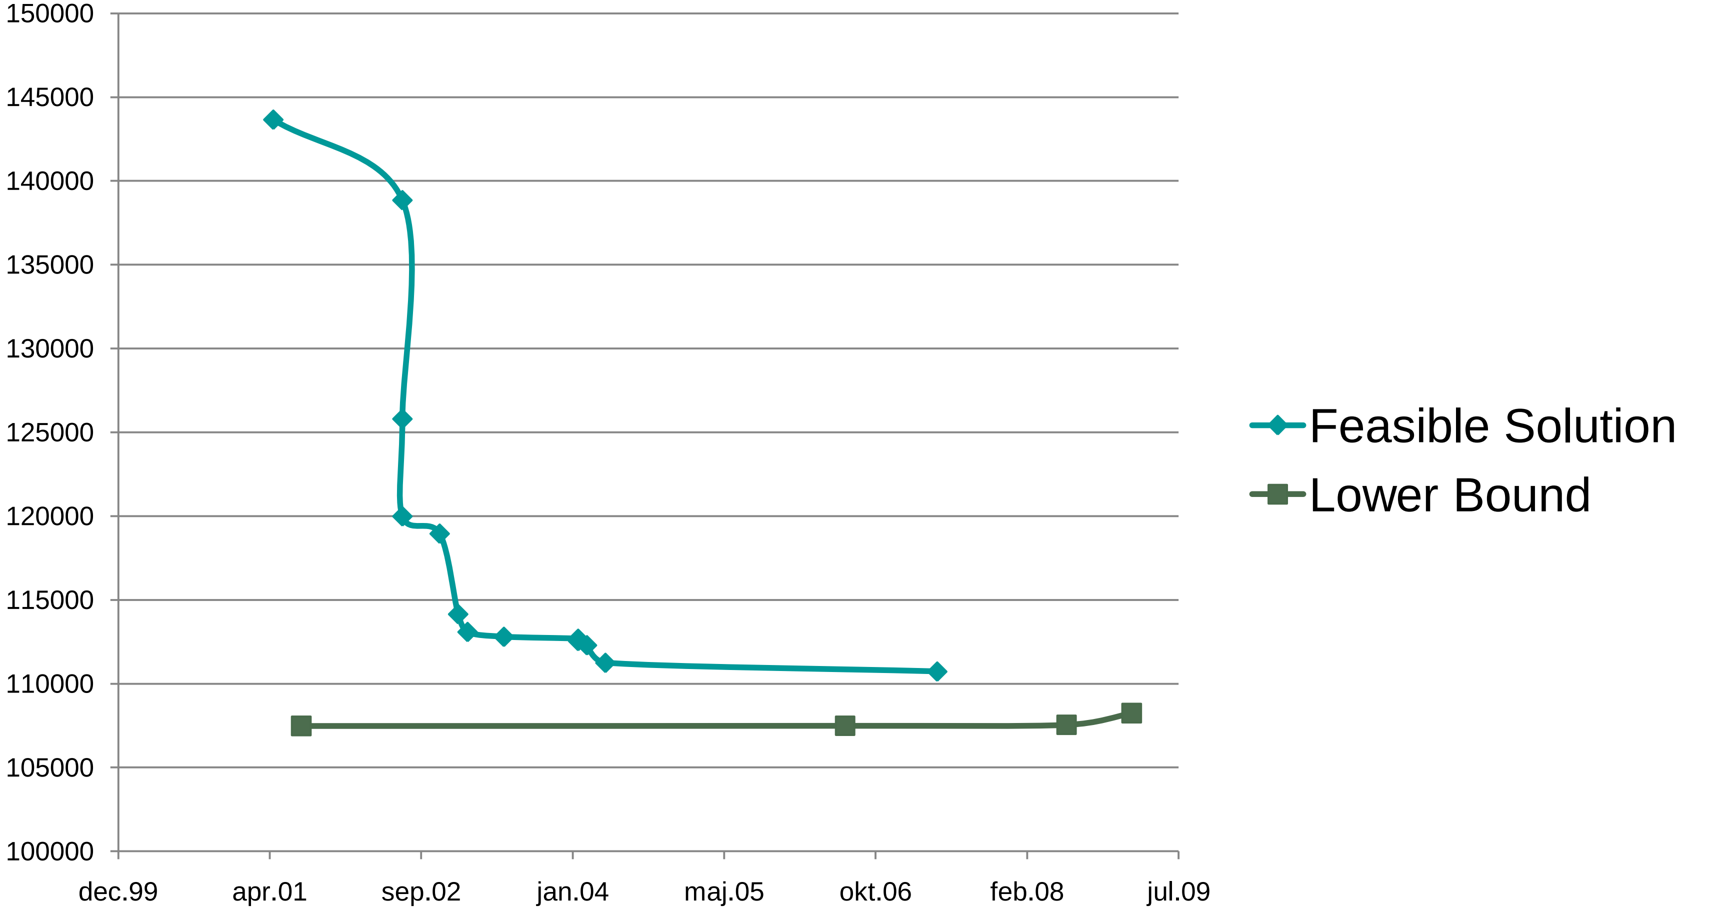

2005 SCHEDULE

SI EXCLUSIVE

EXCLUSIVE ONLINE CONTENT FOR SPORTS ILLUSTRATED SUBSCRIBERS

Matchmakers vs. the Machine

The husband and wife team that made big league scheduling an art is finally being replaced by a computer

FREE E-MAIL ALERTS   E-MAIL THIS   PRINT THIS   SAVE THIS   MOST POPULAR

By Bruce Stoff

For Holly and Henry Stephenson, this year's baseball playoffs brought bad news: They lost their spots on the major league roster.

Their release was one of the quietest milestones in the game's history. For the last 23 seasons the Stephensons had scheduled the national pastime. Working from a cluttered office in their modest home on Martha's Vineyard in Massachusetts, the couple set the date and location of every big league game. Improbable as it sounds, a $4 billion empire depended on the Stephensons' ability to juggle the schedules of two leagues, 30 teams and 1,200 players.

The baseball schedule is the most complex in sports: 2,430 games in 26 weeks. "Few people understand how hard the job is," says Gene Orza, COO of the players' association, "and the Stephensons are the masters of it." Or were. Baseball is set to approve a 2005 schedule produced by a team from Carnegie Mellon, Georgia Tech and the Butler, Pa.-based Sports Scheduling Group (SSG). It is the sport's first fully computer-generated schedule.

Front-loading the Cardinals-Cubs season series was a blemish on the Stephensons' 2004 schedule.
Elsa/Getty Images

SI Exclusive

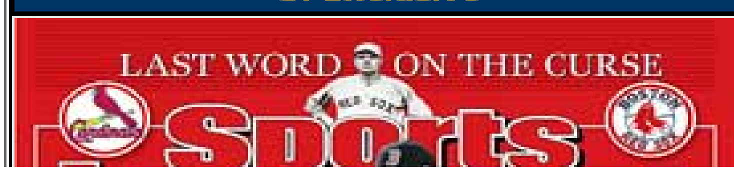

2008 SCHEDULE

Story to be continued

But first

Why were we successful in 2005 (and 2011) but not in 1995?

What changed?

The world of operations research, particularly integer programming, is much different now than in 1995.

Outline

- General Trends
- OR Specific Trends
- Challenge Trends

General Trends

- More data and more spending on data puts pressure on companies to make better decisions from that data (which is a pretty good definition of operations research)

Increased Data

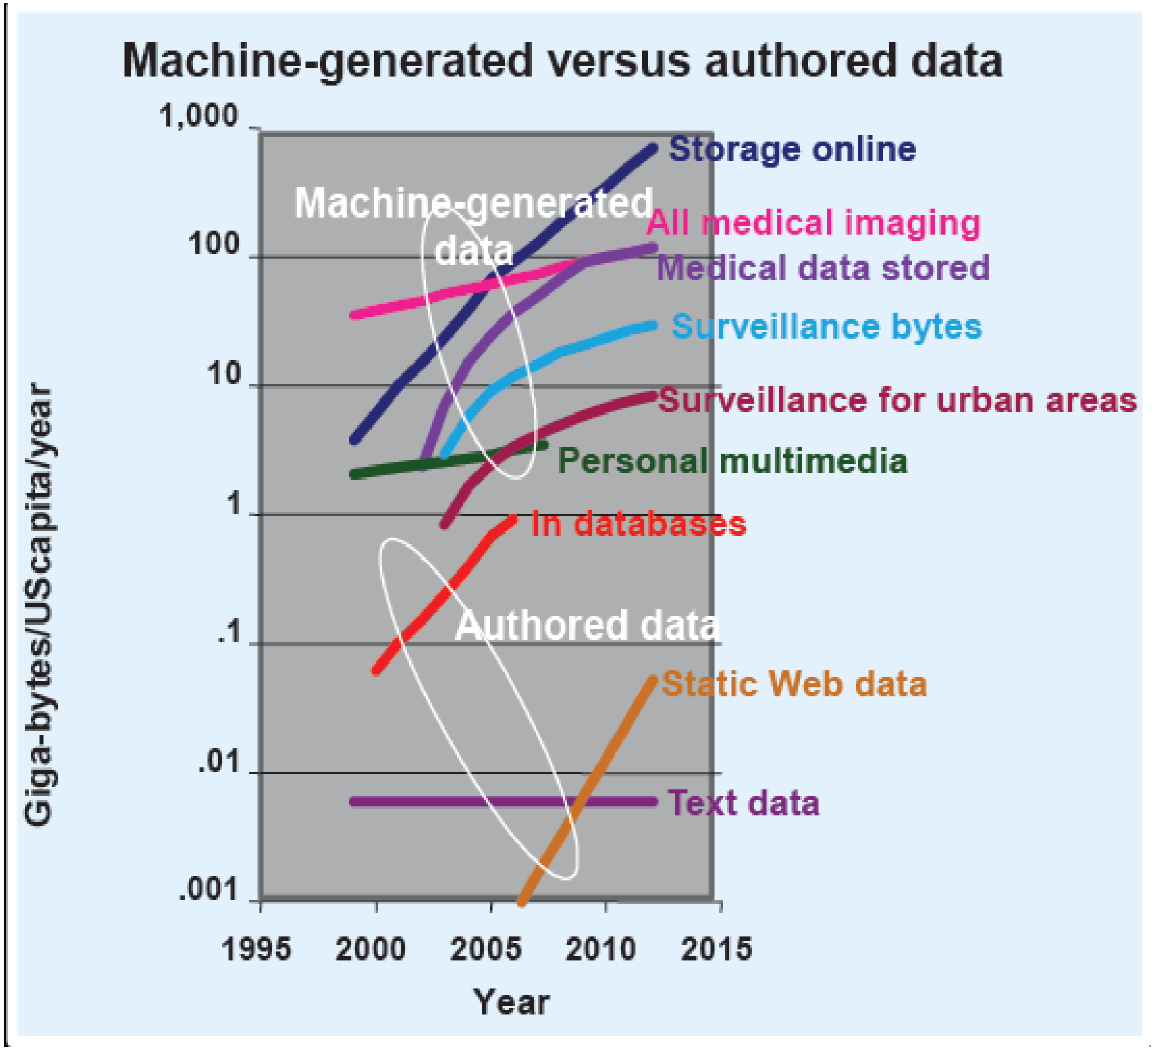

Big Spending

- Yearly global spending on information technology: $1.7 trillion.

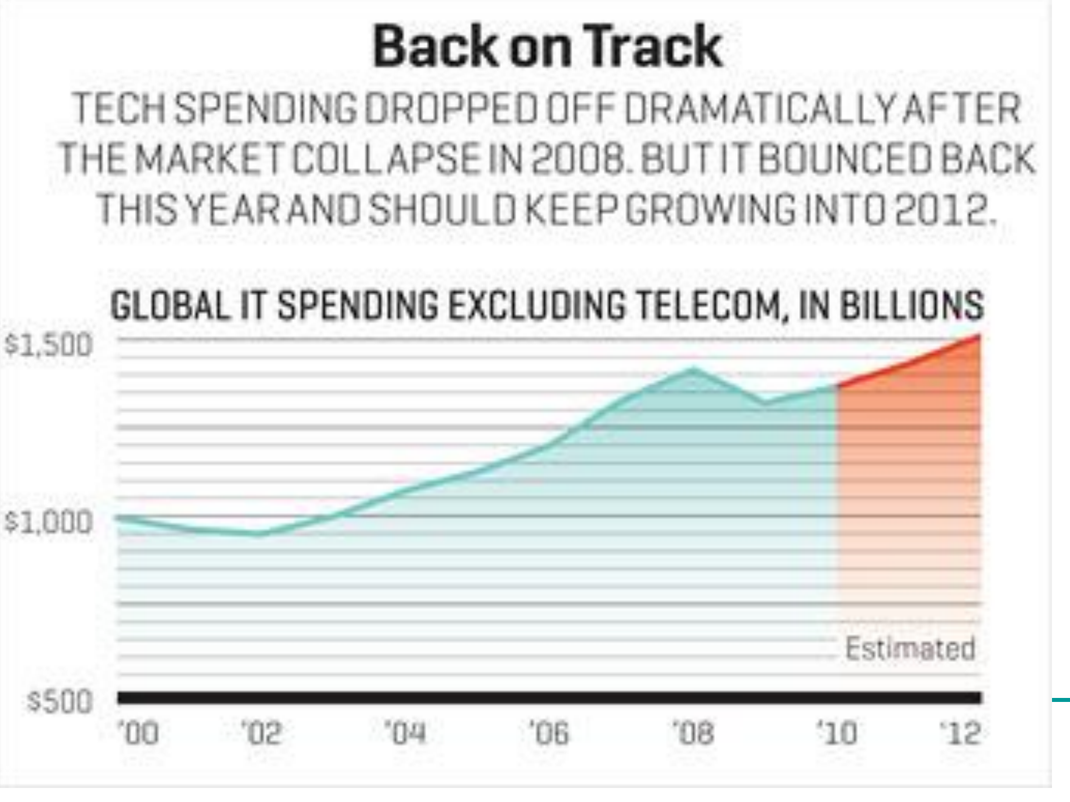

OR Role

- Data is not information; information is not improved decision making
- Operations research allows companies to make better decisions based on the data they are collecting

Faster Computers: Supercomputer and otherwise

340 Gflops
Now costs
US$8,000

www.top500.org

Faster computers increases the relevance and applicability of OR

- Example. Sports scheduling example
    - Time (1996): 18 days to solve (impractical)
    - Time (2010): 30 minutes to solve (practical)

Algorithms are getting better also!

- Linear Programming

Table 5. PDS models—solution times.

Instance	CPLEX 1.0	CPLEX 5.0 Dual	CPLEX 7.1 Primal	CPLEX 7.1 Dual
pds100	—	[illegible]	[illegible]	[illegible]
pds90	—	[illegible]	[illegible]	[illegible]
pds80	—	[illegible]	[illegible]	[illegible]
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

My sports scheduling example now solves in about 3 seconds!

Bixby, Solving Real-World Linear Programs 2002

Speed

- The combination of improved algorithms and faster computers mean that many problems can now be solved 5 million times faster than they could 15 years ago.

- Huge increase in applicability of OR methods (and lots of interest in *business analytics*: which combines predictive analytics with prescriptive analytics)

Illustration: TSP with 2392 nodes

1987: Padberg and Rinaldi

when using OSL as the LP solver: While the CYBER 205 runs of the two largest problems TK1002 and TK2392 of our study required about 7 hrs., 18 mins. and 27 hrs., 20 mins., respectively, the IBM 3090/600 runs required about 3 hrs., 10 mins.

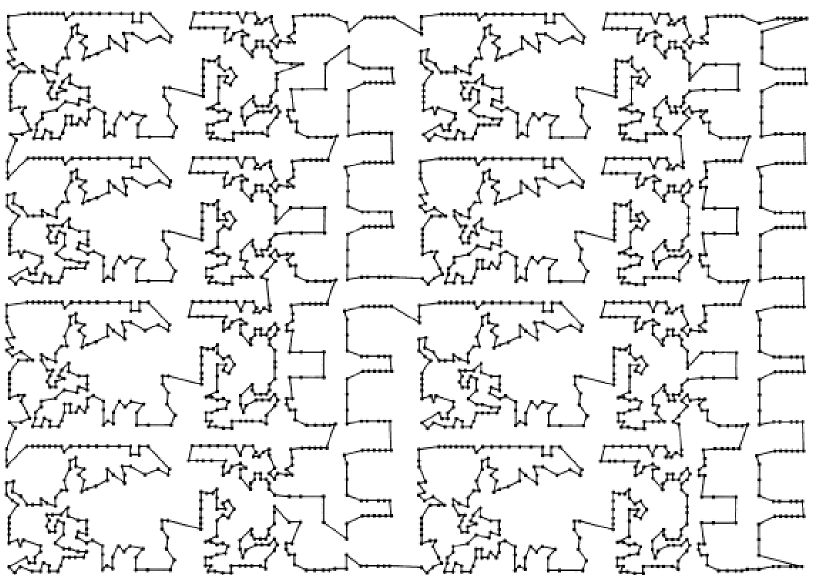

2011: Bill Cook and Concorde with iPhone

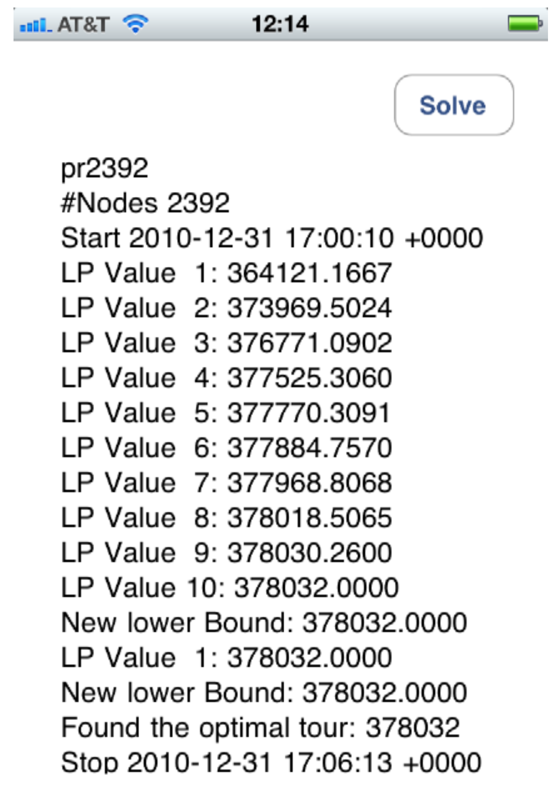

Current Trends in Operations Research (with an IP focus)

- General IP improvements, not problem specific improvements
- Putting more structure into variables
- Linking models
- Integrating metaheuristics and optimization through large neighborhood search

Current Trend 1: General IP Improvements, not Problem Specific Ones

- 1995

  Formulate problem as IP

  Try to solve with CPLEX 2 (or so)

  If fail, then look for problem specific cuts (additional constraint valid just for this problem

- 2011

  Formulate problem as IP

  Try to solve with CPLEX 12 (or Gurobi, or XPRESS)

  If fail, figure out how to help solver do better (then watch as next version does automatically)

Example

Work with United States Postal Service

Mail generated at a mail processing facility at given times; must be transported to other facilities to meet schedules.

Choose shipping contracts at minimum cost to meet processing requirements

2005

- Formulate as an integer program: x[i] = 1 if contract i is chosen; y[i,j]= 1 if mail grouping j put on truck i

- Easy integer program:

  minimize costs;

  every contract chosen at most once;

  every shipment shipped.

- Works OK under CPLEX 7.5, but not great

2005 solution

- Add redundant constraint

Capacity of contracts chosen >= amount to be shipped

Goes from solving 100 contract problem (in 5 minutes) to 5000 contract problem (in 10 seconds)!

But constraint is simply a linear combination? Why?

2005 solution

- CPLEX (and other solvers) do great with “knapsack” constraints:

  $3x + 4y + 5z \le 8$

  Causes code to generate new constraint

  $y + z \le 1$

  which is stronger and lets the code work better

  Redundant constraint is knapsack and the constraints generated by CPLEX then solve the problem

2009 solution

- Present at a few conferences, finally get around to writing up
- Implement on CPLEX 11. No improvement!
- CPLEX 11 solves 5000 contract problem out of the box
- Reason: code now automatically identifies hidden knapsacks

Subtle but important change

- Look carefully at how commercial (or open source: see COIN-OR) is solving and aid the solver.

Current Trend 2: More complicated variables

- 1995
  - Generally straightforward formulations (with some exceptions dating back to 1961)
  - Style was every more complicated, “tricky” (and ultimately weak and unsolvable) constraints
- Current (since 2000)
  - Embed more structure in the variables, even at the expense of a huge (exponential) number of variables
  - Generate as needed (“Branch and Price”)

Formulations

- Simple formulation: x[i,j,t] = 1 if i plays at j in time t
- Formulation is a mess (looks very impressive when all put together)
- Results for NL6: initial bound 2186; After 30 minutes lower bound is 5434, feasible solution of 25,650 (Optimal is 23,916)

Better formulation

- Travel issue can be embedded in “road trip” variables: the path a team takes at a given time.
  - 0-1 variable for “Pittsburgh plays at NY in time 2, then PHI in time 3, then ATL in time 4” (roadtrip)
    - Many more variables, but fewer constraints
    - Much tighter relaxation
    - Able to have many more idiosyncratic costs

Variables

- Sample Variables:

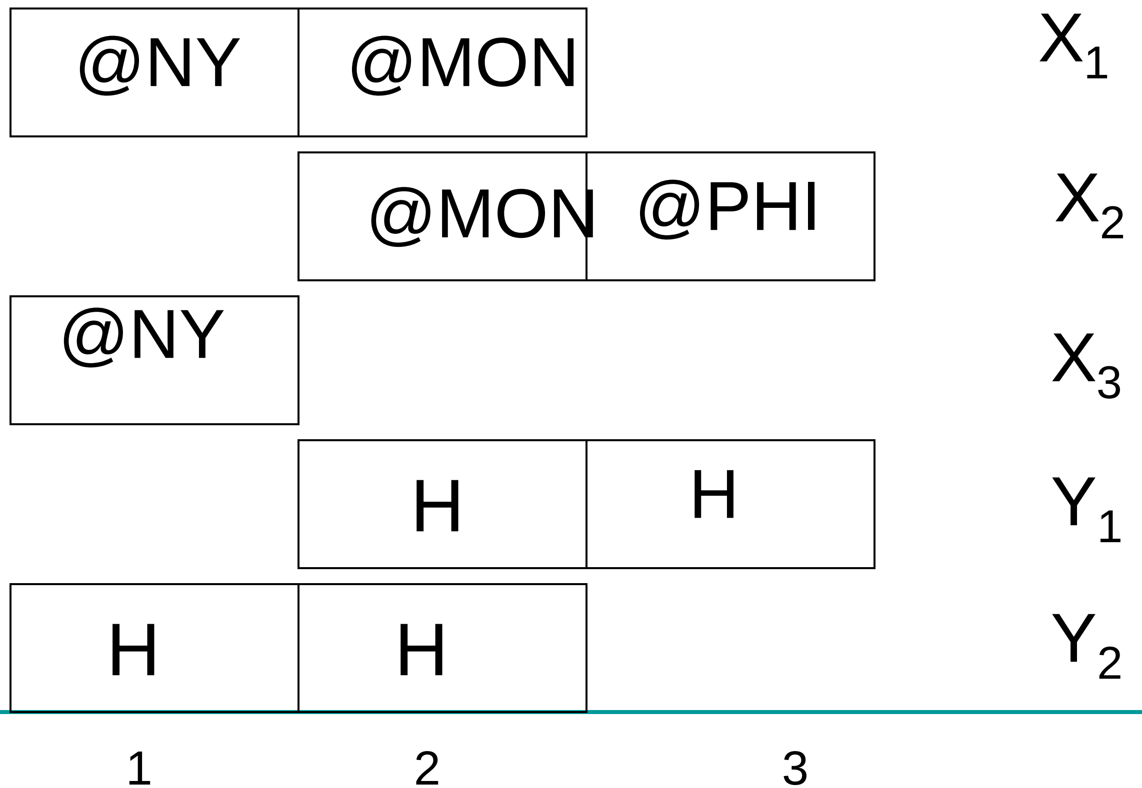

Constraints

- One thing per time: $X_1+X_2+Y_1+Y_2 \leq 1$

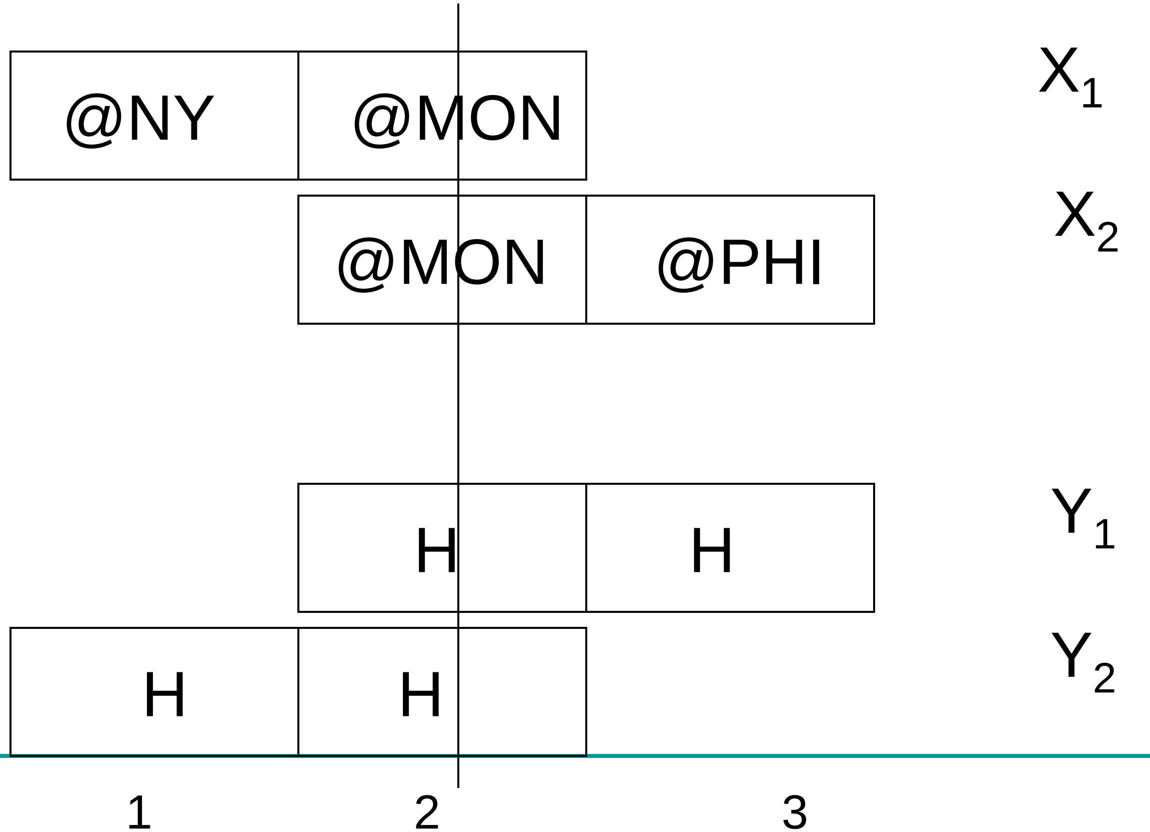

Constraints

- No Away followed by Away $X_1+X_3 \leq 1$

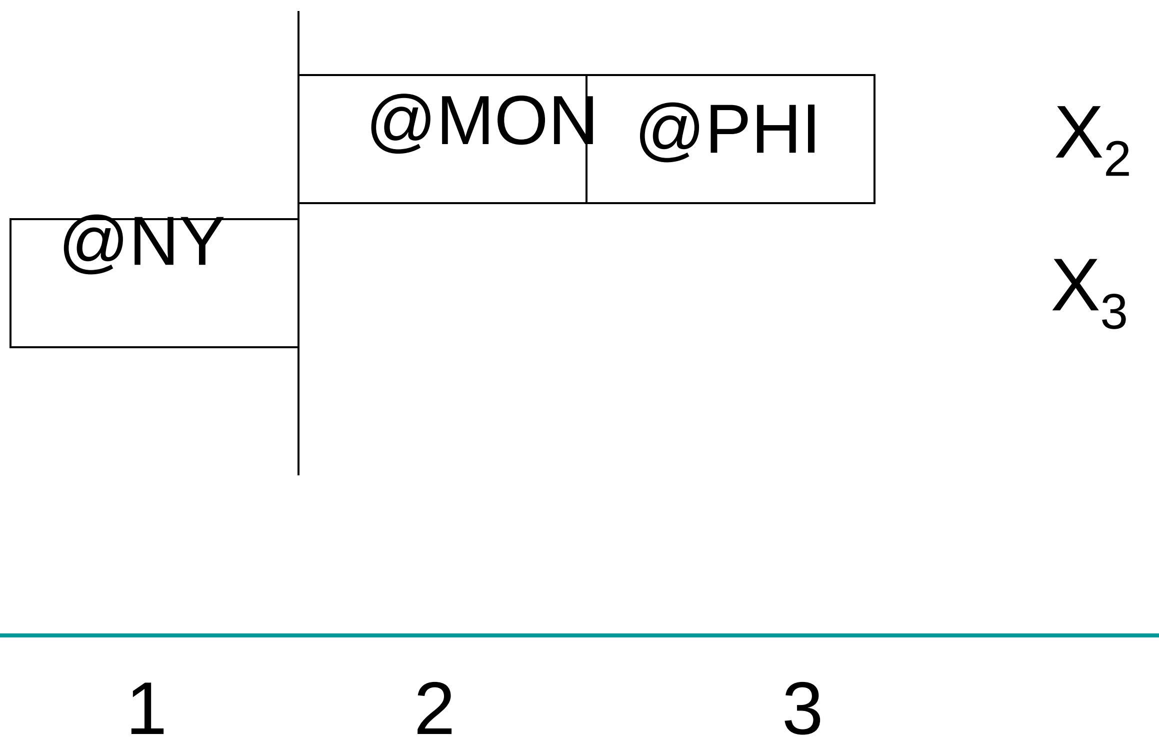

Linking Constraints

- Constraints from different teams linked by “If *a* at *b* then *b* at home” constraints:

$$X_1+X_3 - Y_{NY1}-Y_{NY2} \leq 0$$

Results

- NL 6. Initial bound is 21624.7 (an order of magnitude improvement)
  - Optimality proved within an hour

More complicated variables

- Now a routine approach to formulating integer programs
- Issues in branching and generating variables
- Often extremely effective

Current Trend 3: Linking Models

- Often the case where model A fixes some variables (Open warehouse?) that are then used by model B (Assign customer to warehouse?).

- What if result from model A is wrong for model B?

- Linkage back is provide by technique called Benders Decomposition

Back to 1962

Numerische Mathematik 4, 238—252 (1962)

Partitioning procedures for solving mixed-variables programming problems*

By

J. F. BENDERS**

I. Introduction

In this paper two slightly different procedures are presented for solving mixed-variables programming problems of the type

$$\max\{c^T x + f(y) \mid A x + F(y) \leq b,\ x \in R_p,\ y \in S\}, \quad (1.1)$$

where $x \in R_p$ (the p-dimensional Euclidean space), $y \in R_q$, and S is an arbitrary subset of R_q. Furthermore, A is an (m, p) matrix, $f(y)$ is a scalar function and $F(y)$ an m-component vector function both defined on S, and b and c are fixed vectors in R_m and R_p, respectively.

An example is the mixed-integer programming problem in which certain variables may assume any value on a given interval, whereas others are restricted to integral values only. In this case S is a set of vectors in R_q with integral-valued components. Various methods for solving this problem have been proposed by BEALE [1], GOMORY [8] and LAND and DOIG [11]. The use of integer variables, in particular for incorporating in the programming problem a choice from a set of alternative discrete decisions, has been discussed by DANTZIG [4].

Other examples are those in which certain variables occur in a linear and others in a non-linear fashion in the formulation of the problem (see e.g. GRIFFITH and STEWART [7]). In such cases $f(y)$ or some of the components of $F(y)$ are non-linear functions defined on a suitable subset S of R_q.

Obviously, after an arbitrary partitioning of the variables into two mutually exclusive subsets, any linear programming problem can be considered as being of type (1.1). This may be advantageous if the structure of the problem indicates a natural partitioning of the variables. This happens, for instance, if the problem is actually a combination of a general linear programming and a transportation problem. Or, if the matrix shows a block structure, the blocks being linked only by some columns, to which also many other block structures can easily be reduced. A method of solution for linear programming problems efficiently utilizing such block structures, has been designed by DANTZIG and WOLFE [5].

The basic idea behind the procedures to be described in this report is a partitioning of the given problem (1.1) into two sub problems; a programming

* Paper presented to the 8th International Meeting of the Institute of Management Sciences, Brussels, August 23—26, 1961.

** Koninklijke/Shell-Laboratorium, Amsterdam. (Shell Internationale Research Maatschappij N.V.).

General idea

- Have model B return a constraint that says “If you want a better solution to B, the solution to A must satisfy *this.*”
- Definition of *this* depends on structure
  - If B is a linear program, then duals provide constraint
  - Other ideas (some from the CP/AI world) work for other problems
    - No goods
    - Logical constraints

Example

- Sports scheduling
  - Problem A: Assign home/away patterns
  - Problem B: Assign games consistent with those patterns and teams to patterns
- Objective based on things like
  - Number of consecutive home games
  - …

Problem A: Find HAPs

- Find Home/Away pattern, one sequence per team

```
1: HAHAH
2: AHAHA
3: HHAAH
4: HAHHA
5: AAHHA
6: AHAAH
```

Problem B. Assign Games

- Assign games consistent with HAP (+ denotes home; - is away)

```
1: +2 -3 +6 -4 +5
2: -1 +4 -5 +6 -3
3: +6 +1 -4 -5 +2
4: +5 -2 +3 +1 -6
5: -4 -6 +2 +3 -1
6: -3 +5 -1 -2 +4
```

Iterating

- Easy to formulate and solve each problem individually
- But too many solutions to A to generate all of them.
- Can use solution to B to add constraints to A

Assigning teams to patterns

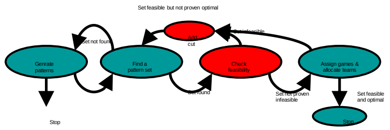

- The allocation corresponds to a matching in a bipartite graph

Use the Hungarian Method to:

- Find a set of teams which cannot be assigned to the pattern set or

-Find a feasible matching

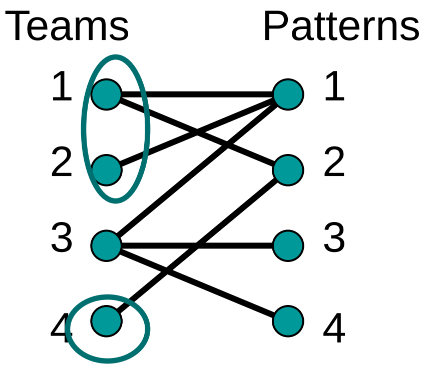

In this case, Benders constraint says “Find 3 patterns that 1, 2, and 4 can do”

Results

- Problems that don’t solve in 10 minutes can be solved in <.5 seconds
- More examples in scheduling, resource assignment, transportation, etc.

Benders

- Underused, but becoming more common
- Great for automatic linking of multi-set approaches

Current Trend 4: Large Scale Neighborhood search

For MLB size problems, integer programming finds good but not optimal solution

1995 Approach

Local search: pick a time slot and pair of games and switch opponents.

(Doesn't work well: hard to get back to feasibility)

2011 Approach

- Take a schedule; relax ¼ of it; resolve; repeat for days on end

Large Scale Neighborhood Search

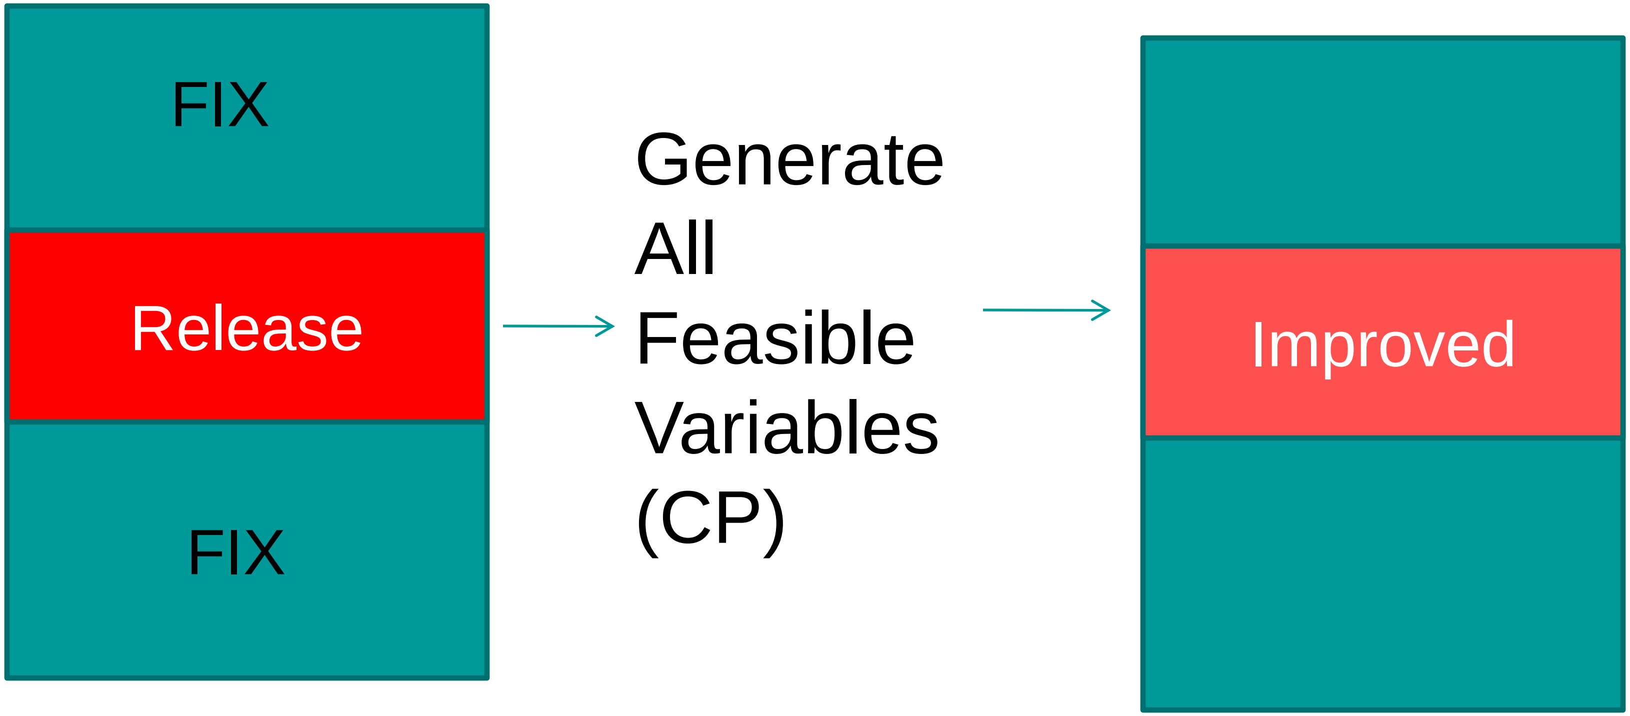

Lots to play around with

- Release set of slots
  - Consecutive
  - Nonconsecutive
- Release set of teams
- Release combination

- Often used to get rid of “bad” part of schedule (keep releasing bad part along with other parts)

Large Neighborhood Search

- *The* key idea for a lot of practical integer programming-based OR work
- Used in airline crew scheduling, machine scheduling, etc. etc.

Result for MLB

- System is a mix of
  - Complicated variables
  - Benders constraints to give information between phases
  - Large neighborhood local search to improve solutions
  - Heavy dependence on underlying optimization to identify formulation improvements automatically
- A Pretty High Barrier to Entry!

Challenge Trends

- Linkage between predictive analytics (data mining) and prescriptive analytics (optimization)
- Robust solutions
- Parallelism
- Adversaries

Challenge Trend 1: Prescriptive and Predive Analytics Example

- Large automobile parts supplier
- Receives estimates of order quantities six months in advance (which are then regularly updated)

- Uses estimates to do capacity planning, shift scheduling, etc.

Capacity Planning

- Well solved, classical operations research problem
- Issues of level of detail, scheduling, etc. etc.
- But what about data…?

Problem

- Estimates were (sometimes) terrible:
- January: “I’ll need 1000 in June”
- February: “I’ll need 2500 in June”
- March: “I won’t need any in June”
- April: “I’ll need 10,000 in June”
- May: “I’ll need 500 in June”
- June: “Where are the 5,000 I need?”

Business Analytics Approach to Capacity planning

- Capacity planning needs to adapt to the messy data environment.

One step: cluster customers based on accuracy of their forecasts

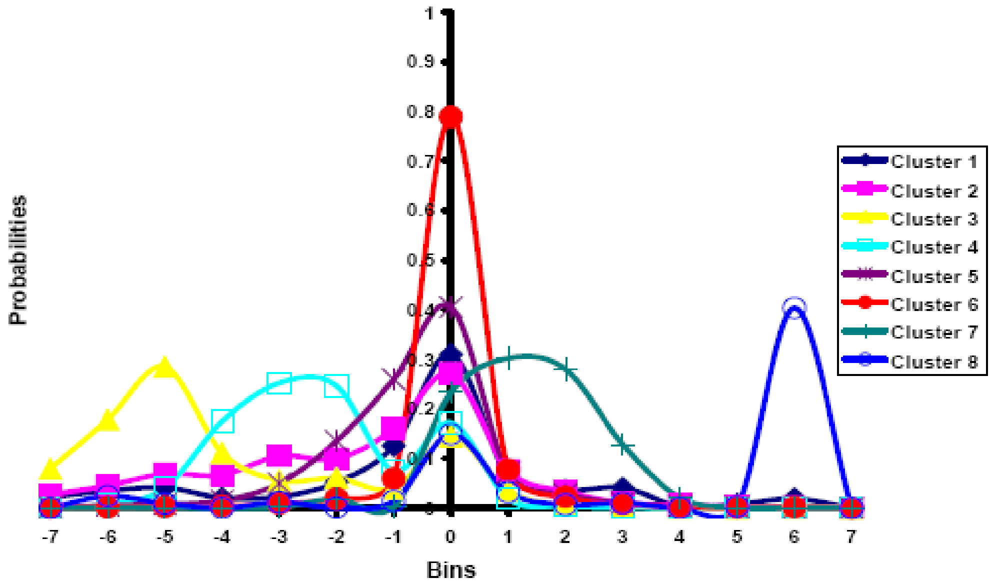

Other Examples

- Revenue management and prediction of “no-shows”
- Text mining to estimate stock variability and optimal portfolios

Challenge Trend 2: Handling Uncertainty and Robustness

- Many approaches for handling uncertainty in models
  - Replacing with mean, or some %tile (USPS capacity planning)
  - Scenario optimization
  - Stochastic programming
  - “Robust” optimization (assumptions on ranges, number of changes of data)

True measures of robustness?

- What does a robust solution really mean?
- David Ryan and others on robust airline crew schedules

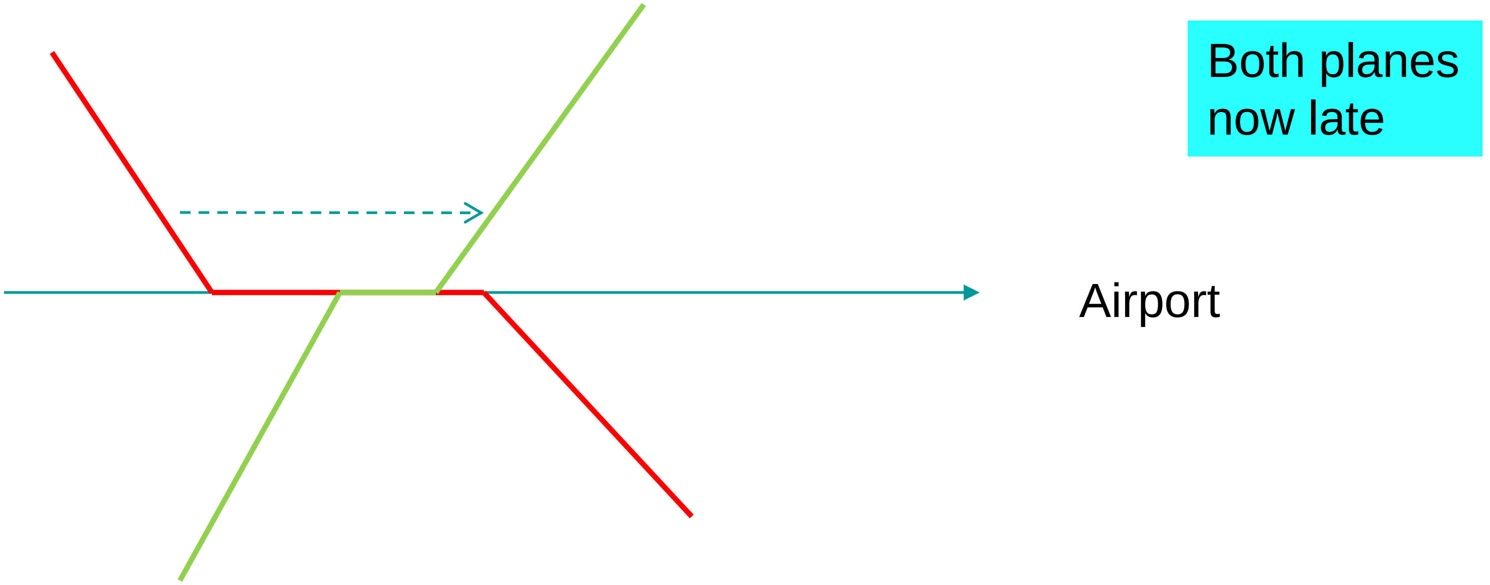

Measure of Robustness

- Unclear where uncertainty is. Model storms? Politics? Baggage handling?
- Use “crew plane changes” as surrogate. Schedules with few plane changes are naturally more robust.
- Get robustness at little cost

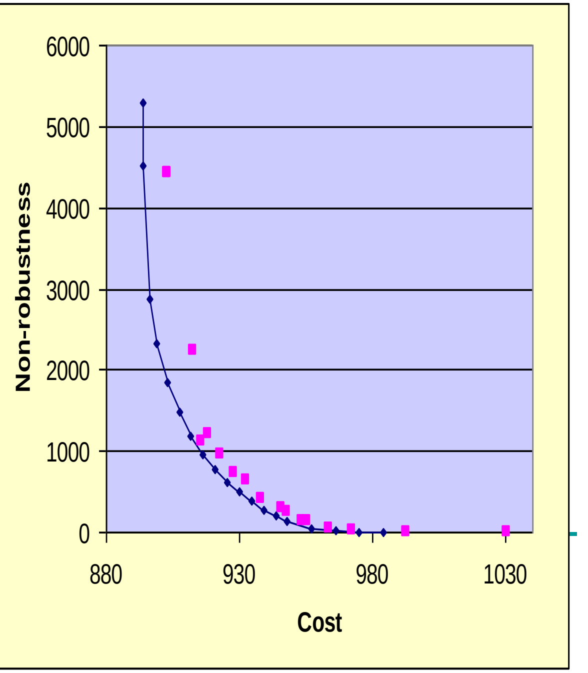

Data-free Robustness

- Bartholdi and Eisenstein: scheduling shuttle buses. Don’t! At checkpoint, use rule “Wait 60% of the time until the bus behind will arrive at the checkpoint”
- Robust to changes in traffic congestion, passenger load times, and even number of buses!

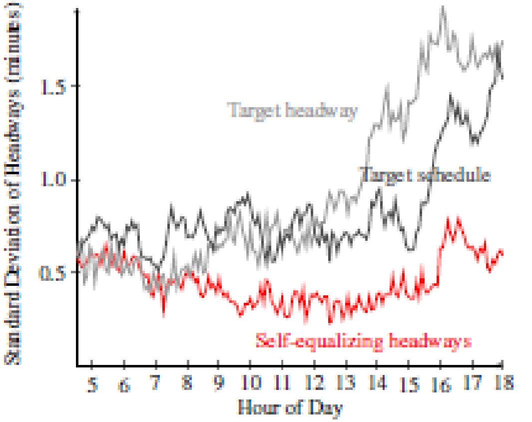

Other applications?

- Badly need some measure of robustness for financial markets: historical variance doesn't seem to work well.
- Scheduling: include simple robustness in Major League Baseball scheduling

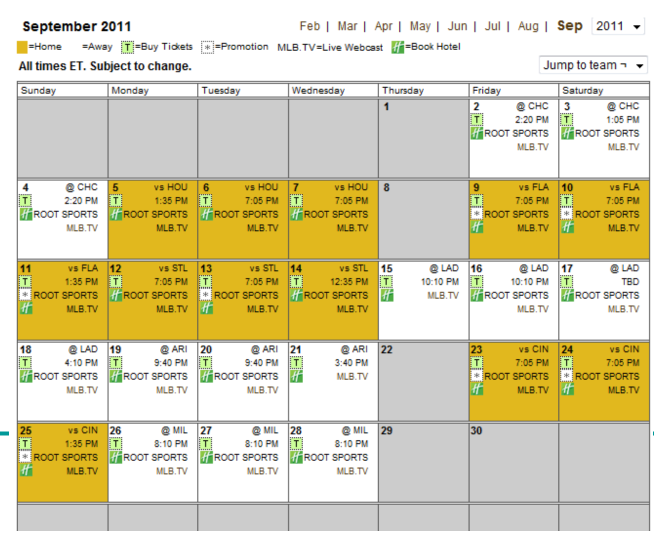

Challenge Trend 3: Parallelism

- Back to the “top 500”:

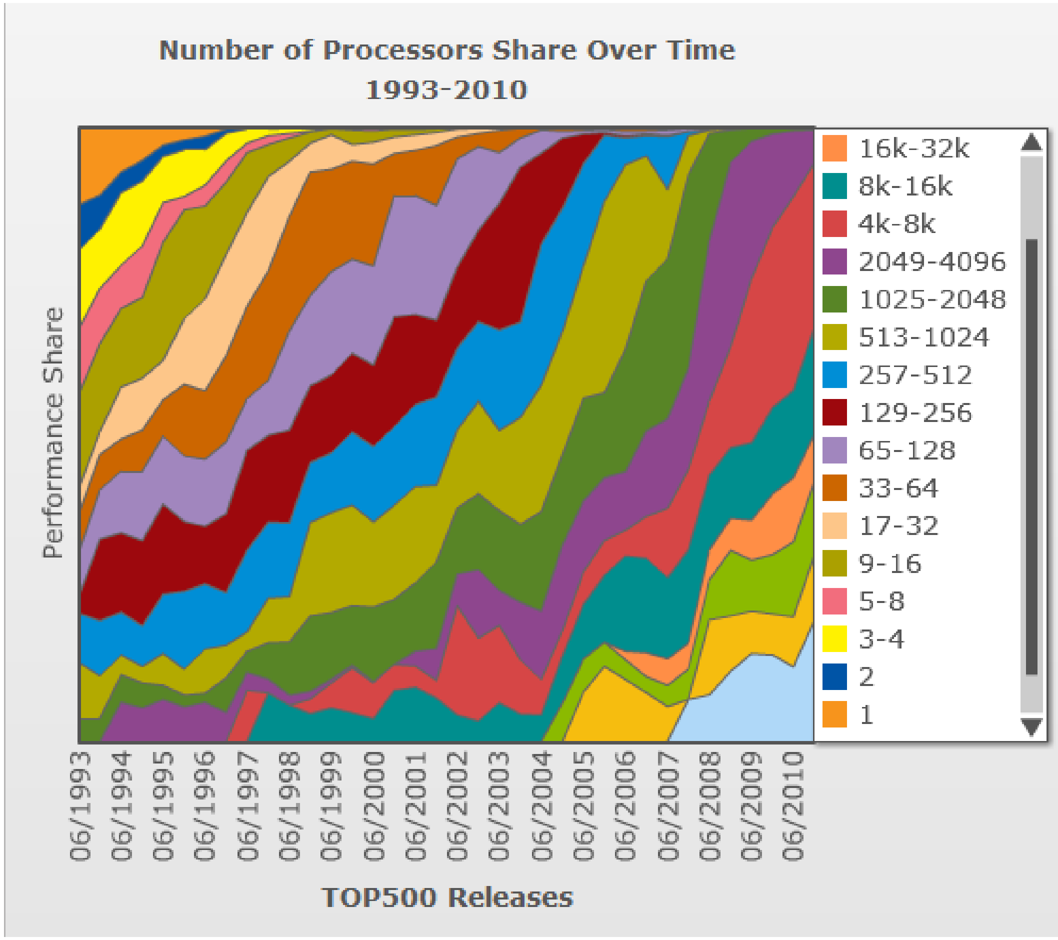

Risk for some types of operations research

- Some of what we do seems hard to parallelize:
  - Simplex-based linear programming
  - Many network algorithms
  - Dynamic Programming

Effect on Mixed Integer Programming

- Hans Mittelmann's benchmarks:

s	problem	CPLEX6	GUROBI6	CBC6	CPLEX12	GUROBI12	CBC12
1	air04	8	7	31	9	7	31
	cap6000	1	1	1	1	1	2
	mas74	103	56	166	100	47	95
	mod011	15	12	13	15	12	14
	mzzv11	25	12	97	24	12	100
	mzzv42z	13	3	67	13	3	78
	pk1	12	6	13	8	5	9
	qiu	6	7	55	4	5	35
2	ran12x21	6	10	22	6	8	12
	ran13x13	3	5	12	3	3	9
3	binkar10_1	4	2	95	6	1	65
	lrn	24	65	475	25	24	355
	prod2	18	8	17	14	8	17
4	bc1	25	30	228	14	30	178
	bienst2	13	7	390	9	5	162
	dano3_5	140	81	169	173	79	162
	mark._4_0	20	11	42	10	10	34
	mark._5_0	1053	1837	f	575	691	f
	qap10	3	56	156	3	56	157
	seymour1	37	32	693	28	32	210
	swath2	7	4	11	7	3	9
	swath3	27	9	181	27	13	60
	30_05_100	9	9	153	9	9	141

Not a lot of halving of time between the 6 and the 12 columns!

Challenges

- In order to keep up with computer speed improvements, we need to improve parallel performance
- Alternative algorithms? Perhaps formerly "failed" algorithms now seem more appealing?

Challenge Trend 4: Adversaries

- As OR combines with predictive analytics, more likely to face adversaries working against model assumptions
  - Best Buy searching for “bad customers”
  - Google in never-ending fight against spammers and those who game rankings
- But we are pretty good at that.

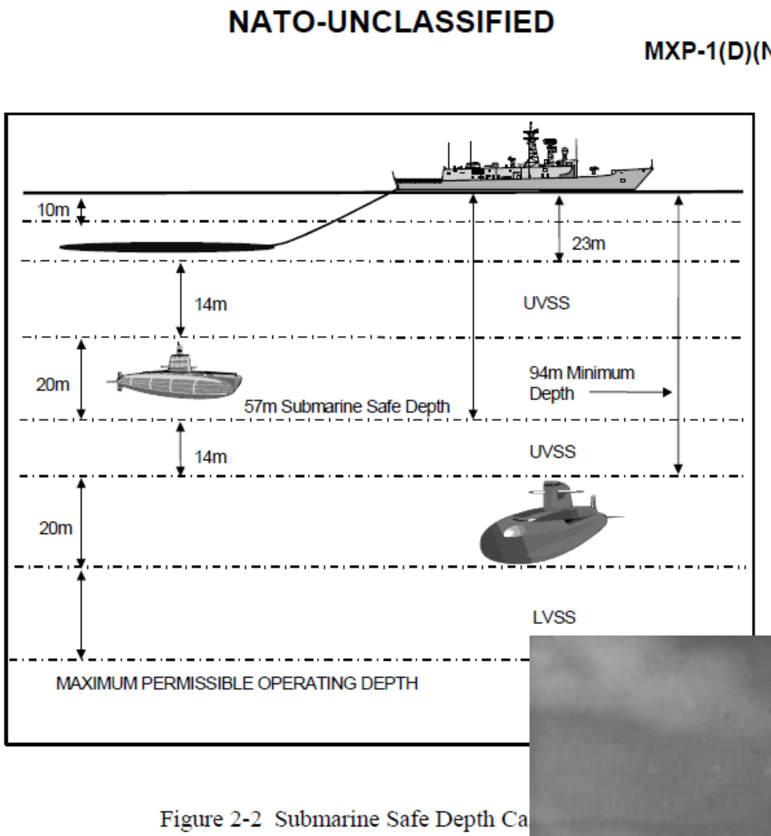

Which is where Operations Research started

Summary: Trends

- General Trends
  - More Data
  - Faster computers
  - Better algorithms
  - Lower fixed Costs
- OR (IP) Specific Trends
  - More general, less specific
  - Complicated variables
  - Linked models
  - Large scale neighborhoods
- Future Trends
  - Predictive Analytics
  - Robustness
  - Parallelism
  - Adversaries

More at

http://mat.tepper.cmu.edu/blog

Back to the story

- August 2007 rolls around: still no schedule
  - Normally due June 1
- Phone call: "Could you help us out?"

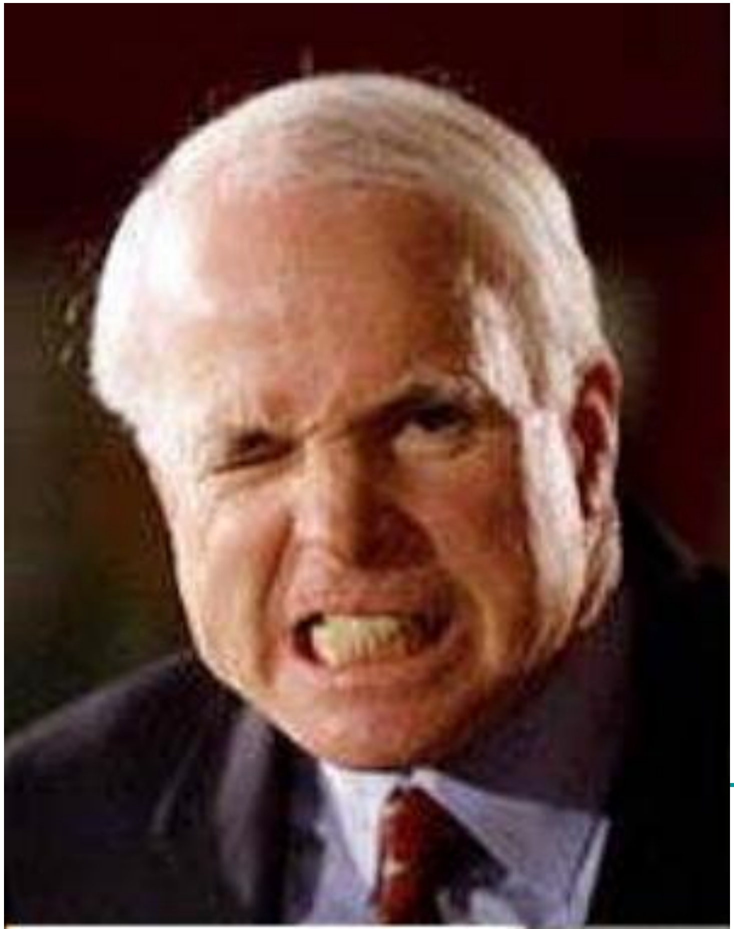

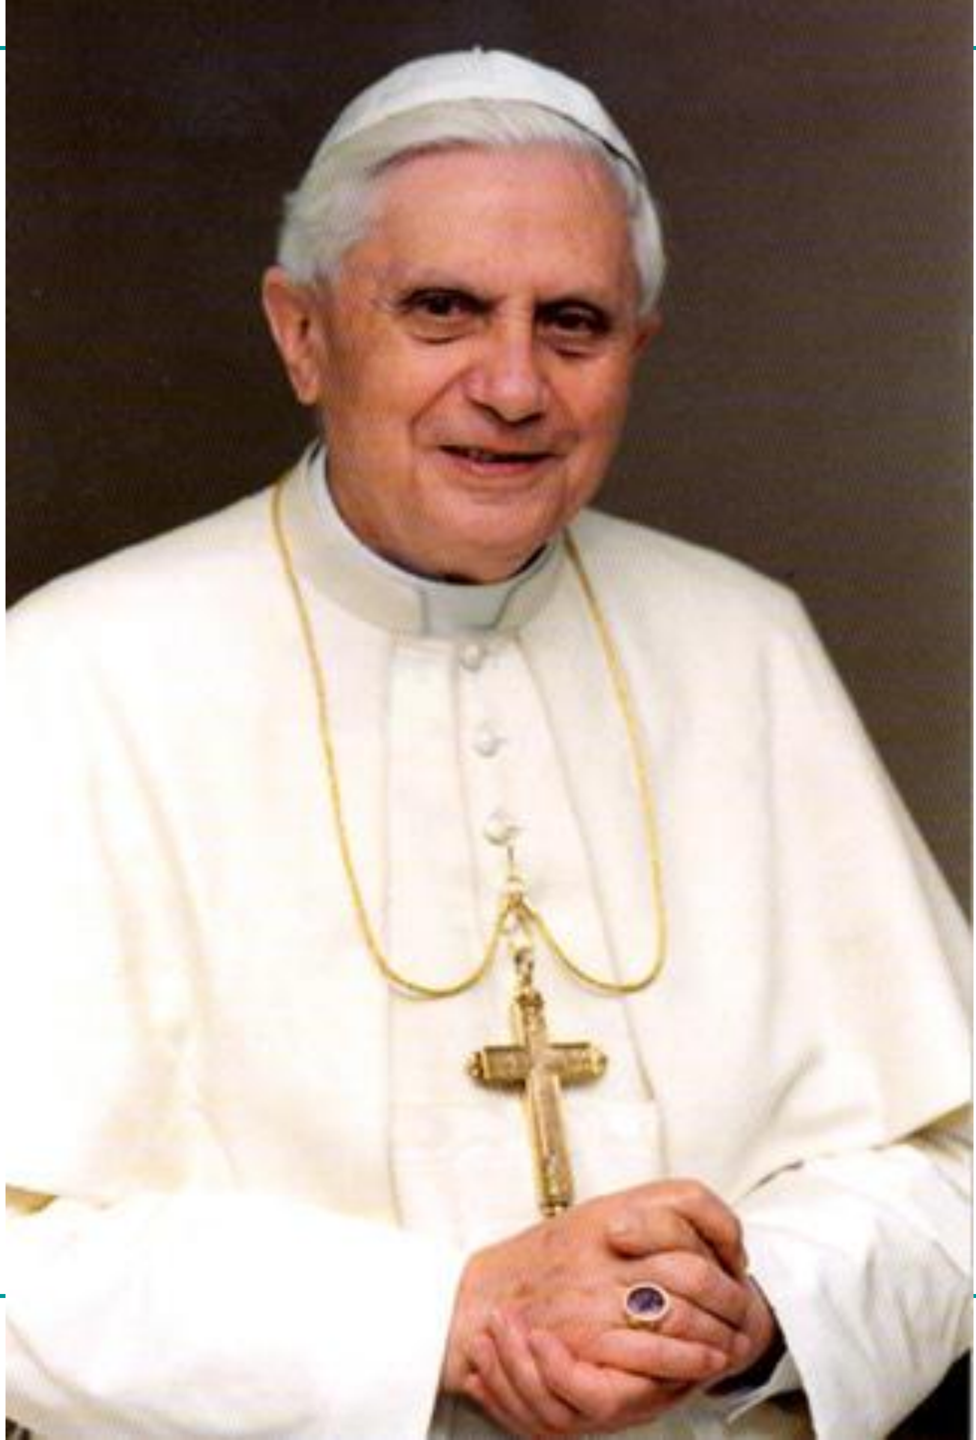

The New York Times
Friday, March 21, 2008

N.Y./Region

City Room

Back to front page »

November 12, 2007, 12:41 pm

Pope to Visit Ground Zero and Celebrate Mass at Yankee Stadium

By SEWELL CHAN

Updated, 2:12 p.m. | Pope Benedict XVI will make his first visit to the United States as pontiff over six days in April, and he plans to visit ground zero, address the United Nations and celebrate Masses at National Stadium in Washington and St. Patrick's Cathedral and Yankee Stadium in New York, officials at the Vatican and the Archdiocese of New York announced today.

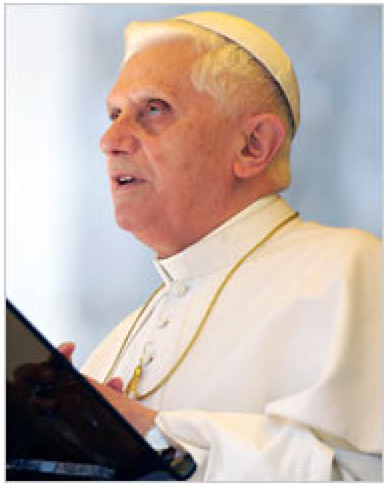
Benedict XVI's visit to the United States in April will be his first as pope. (Photo: Alessia Giuliani/AFP/Getty Images)

The papal visit will be only the fourth in New York City's history. Pope Paul VI visited in October 1965, during the first-ever papal visit to the United States. Pope John Paul II visited New York in October 1979 and October 1995.

SERIES SCHEDULE (NL): Official Schedule MLB 2008

	WAS	NYM	PHI	ATL	FLA	PIT	CHC	CIN	STL	HOU	MIL	ARZ	LAD	SF	SD	COL
3/31	@PHI	@FLA	WAS	PIT	NYM	@ATL	MIL	ARZ	COL	@ SD	@CHC	@CIN	SF	@LAD	HOU	@STL
4/4	@STL	@ATL	@CIN	NYM	PIT	@FLA	HOU	PHI	WAS	@CHC	SF	@COL	@ SD	@MIL	LAD	ARZ
4/7	FLA	PHI	@NYM	@COL	@WAS	CHC	@PIT	@MIL	@HOU	STL	CIN	LAD	@ARZ	SD	@ SF	ATL
4/11	ATL	MIL	CHC	@WAS	@HOU	CIN	@PHI	@PIT	@ SF	FLA	@NYM	COL	SD	STL	@LAD	@ARZ
4/14	NYM	@WAS	HOU	@FLA	ATL	@LAD	CIN	@CHC	MIL	@PHI	@STL	@ SF	PIT	ARZ	COL	@ SD
4/18	@FLA	@PHI	NYM	LAD	WAS	@CHC	PIT	MIL	SF	COL	@CIN	SD	@ATL	@STL	@ARZ	@HOU
4/21	@ATL	@CHC	@COL	WAS	@PIT	FLA	NYM	LAD	@MIL	SD	STL	SF	@CIN	@ARZ	@HOU	PHI
4/23	@NYM	WAS	@MIL	FLA	@ATL	STL	@COL	HOU	@PIT	@CIN	PHI	@LAD	ARZ	@ SD	SF	CHC
4/25	CHC	ATL	@PIT	@NYM	@MIL	PHI	@WAS	@ SF	HOU	@STL	FLA	@ SD	COL	CIN	ARZ	@LAD
4/28	ATL	PIT	SD	@WAS	LAD	@NYM	MIL	@STL	CIN	@ARZ	@CHC	HOU	@FLA	COL	@PHI	@ SF
5/2	PIT	@ARZ	SF	CIN	SD	@WAS	@STL	@ATL	CHC	MIL	@HOU	NYM	@COL	@PHI	@FLA	LAD
5/5	@HOU	@LAD	@ARZ	SD	MIL	SF	@CIN	CHC	@COL	WAS	@FLA	PHI	NYM	@PIT	@ATL	STL
5/9	FLA	CIN	@ SF	@PIT	@WAS	ATL	ARZ	@NYM	@MIL	@LAD	STL	@CHC	HOU	PHI	COL	@ SD
5/12	@NYM	WAS	ATL	@PHI	@CIN	@STL	SD	FLA	PIT	@ SF	LAD	COL	@MIL	HOU	@CHC	@ARZ
5/16	@BAL	@NYY	TOR	OAK	KC	@CHC	PIT	CLE	TB	@TEX	@BOS	DET	@LAA	CWS	@SEA	MIN
5/19	PHI	@ATL	@WAS	NYM	ARZ	MIL	@HOU	@LAD	@ SD	CHC	@PIT	@FLA	CIN	@COL	STL	SF
5/23	MIL	@COL	@HOU	ARZ	SF	CHC	@PIT	@ SD	@LAD	PHI	@WAS	@ATL	STL	@FLA	CIN	NYM
5/26	@ SD	FLA	COL	@MIL	@NYM	@CIN	LAD	PIT	HOU	@STL	ATL	SF	@CHC	@ARZ	WAS	@PHI
5/30	@ARZ	LAD	FLA	@CIN	@PHI	@STL	COL	ATL	PIT	@MIL	HOU	WAS	@NYM	SD	@ SF	@CHC

SERIES SCHEDULE (NL): Official Schedule MLB 2008

	WAS	NYM	PHI	ATL	FLA	PIT	CHC	CIN	STL	HOU	MIL	ARZ	LAD	SF	SD	COL
3/31	@PHI	@FLA	WAS	PIT	NYM	@ATL	MIL	ARZ	COL	@ SD	@CHC	@CIN	SF	@LAD	HOU	@STL
4/4	@STL	@ATL	@CIN	NYM	PIT	@FLA	HOU	PHI	WAS	@CHC	SF	@COL	@ SD	@MIL	LAD	ARZ
4/7	FLA	PHI	@NYM	@COL	@WAS	CHC	@PIT	@MIL	@HOU	STL	CIN	LAD	@ARZ	SD	@ SF	ATL
4/11	ATL	MIL	CHC	@WAS	@HOU	CIN	@PHI	@PIT	@ SF	FLA	@NYM	COL	SD	STL	@LAD	@ARZ
4/14	NYM	@WAS	HOU	@FLA	ATL	@LAD	CIN	@CHC	MIL	@PHI	@STL	@ SF	PIT	ARZ	COL	@ SD
4/18	@FLA	@PHI	NYM	LAD	WAS	@CHC	PIT	MIL	SF	COL	@CIN	SD	@ATL	@STL	@ARZ	@HOU
4/21	@ATL	@CHC	@COL	WAS	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	SF	@CIN	@ARZ	@HOU	PHI
4/23	@NYM	WAS	@MIL	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	@LAD	ARZ	@ SD	SF	CHC
4/25	CHC	ATL	@PIT	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	@ SD	COL	CIN	ARZ	@LAD
4/28	ATL	PIT	SD	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	HOU	@FLA	COL	@PHI	@ SF
5/2	PIT	@ARZ	SF	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	NYM	@COL	@PHI	@FLA	LAD
5/5	@HOU	@LAD	@ARZ	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	PHI	NYM	@PIT	@ATL	STL
5/9	FLA	CIN	@ SF	@PIT	@WAS	ATL	ARZ	@NYM	@MIL	@LAD	STL	@CHC	HOU	PHI	COL	@ SD
5/12	@NYM	WAS	ATL	@PHI	@CIN	@STL	SD	FLA	PIT	@ SF	LAD	COL	@MIL	HOU	@CHC	@ARZ
5/16	@BAL	@NYY	TOR	OAK	KC	@CHC	PIT	CLE	TB	@TEX	@BOS	DET	@LAA	CWS	@SEA	MIN
5/19	PHI	@ATL	@WAS	NYM	ARZ	MIL	@HOU	@LAD	@ SD	CHC	@PIT	@FLA	CIN	@COL	STL	SF
5/23	MIL	@COL	@HOU	ARZ	SF	CHC	@PIT	@ SD	@LAD	PHI	@WAS	@ATL	STL	@FLA	CIN	NYM
5/26	@ SD	FLA	COL	@MIL	@NYM	@CIN	LAD	PIT	HOU	@STL	ATL	SF	@CHC	@ARZ	WAS	@PHI
5/30	@ARZ	LAD	FLA	@CIN	@PHI	@STL	COL	ATL	PIT	@MIL	HOU	WAS	@NYM	SD	@ SF	@CHC

We've continued to schedule through 2011 (at least!)

Takeaways

- Traveling Tournament Problem: http://mat.tepper.cmu.edu/TOURN
- More on Benders? John Hooker's book *Integrated Methods in Optimization*, Springer 2008
- Blog on OR: http://mat.tepper.cmu.edu/blog
- Advances in IP and CP
  - More complicated variables
  - Large scale local search
  - Benders constraints

This has even led to a buzzword: Business Analytics

- Integration of
  - Descriptive Analytics (What is happening)
  - Predictive Analytics (What is going to happen)
  - Prescriptive Analytics (What to do about it)

Business Analytics is a Hot Topic

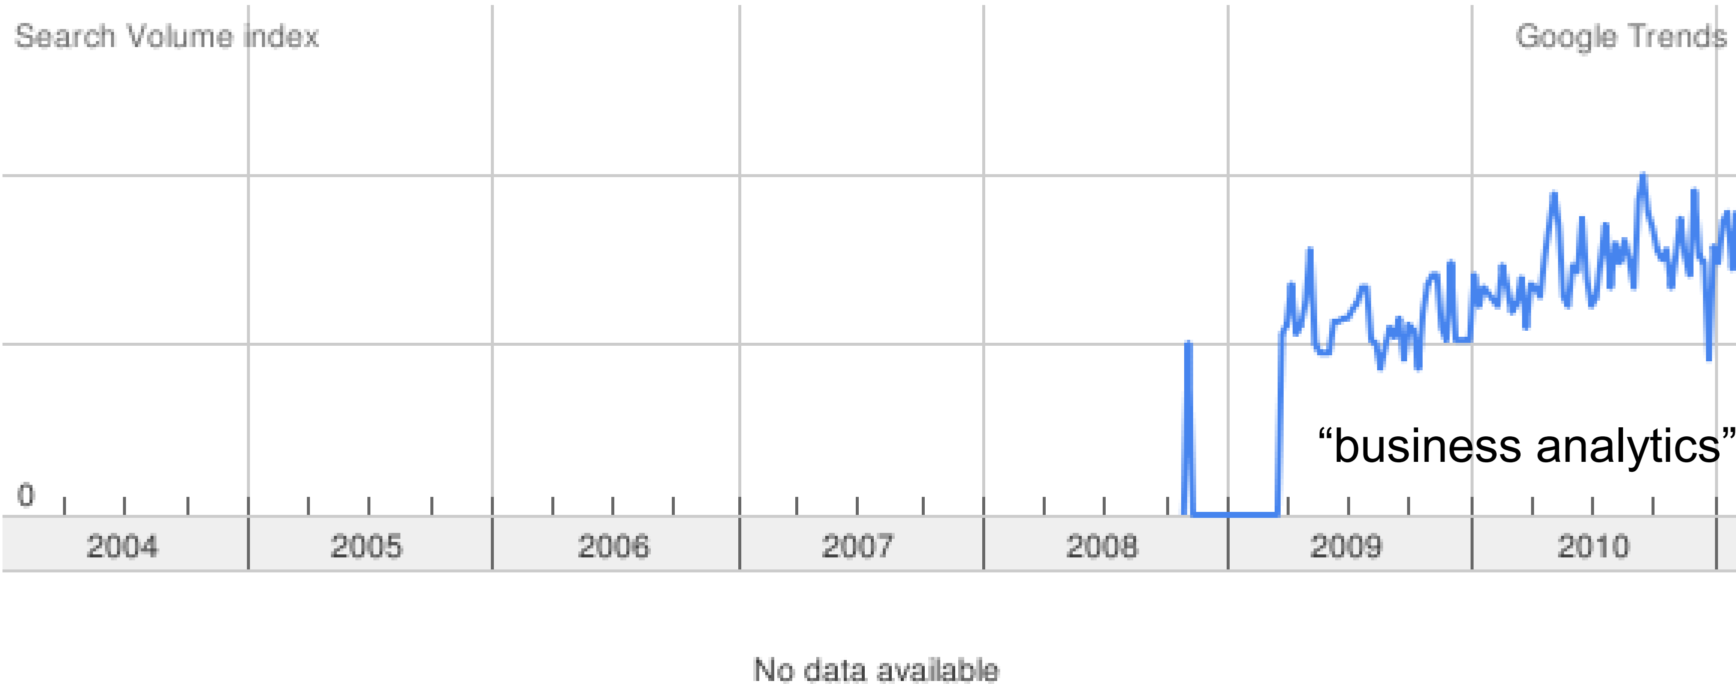

…. Maybe


Search Volume index
Google Trends
60.0
40.0
20.0
0
Operations research
Management science
Business analytics
2004
2005
2006
2007
2008
2009
2010
No data available