Solution Pluralism, Deliberation, and Metaheuristics

Extracting More Value from Optimization Models
Part 1: Motivation and Examples.

Steven O. Kimbrough

University of Pennsylvania
kimbrough@wharton.upenn.edu

MIC: Metaheuristics International Conference, Singapore, 5–8 August 2013
http://research.larc.smu.edu.sg/mic2013/

Abstract

We wish to challenge two verities in the MS/OR community as a way of promoting a conceptual shift regarding optimization and metaheuristics. The first verity, roughly, is that given a constrained optimization model the primary problem posed is to find an optimal solution to the model. We call this the *goal of optimization modeling* verity. The second, roughly, is that exactly optimal solutions (to optimization problems) are always preferred, but heuristically optimal solutions are acceptable if exactly optimal solutions are not available. We call this the *justification of heuristics* (in optimization) verity. Our challenges to these verities lie not in denying their truth so far as they go, but in denying that they have gone far enough.

Abstract (con't.)

The conceptual shift we propose may be described as *solution pluralism for deliberation with models*. In a nutshell: Given an optimization model, it is possible to define a set of *solutions of interest* (SoIs), which if well sampled would be valuable for deliberation for decision making (based on the model); further, while the problem of obtaining good samples of the SoIs is challenging, metaheuristics bid fair to be the favored approach and can be shown to be effective in many cases. Among the reasons why SoIs can usefully be defined is that, as is well known, constraints not based in logic often have somewhat arbitrarily chosen coefficient values and these constraints are amenable to adjustment if the increase in profit or decrease in cost is sufficiently large. The talk elaborates upon and illustrates these points.

Acknowledgements

Active collaborators for this talk: Hoong Chuin (HC) Lau, Frederic H. Murphy.

Also actively, Simon Caton, Christian Haas, and Magie Hall.

Thanks to the KSRI audience on August 25, 2013 for a number of useful comments.

Past collaborators: Gary Koehler, Ann Kuo, Lindawati, Ming Lu, David Harlan Wood, D.J. Wu

Outline

Outline

Motivation: Given an optimization problem...

- How can we recognize a modeling opportunity?
- How should we build the model(s)?
- How can we get a good solution?
- What should our decision be?

Our focus: this last question. The deliberation problem.

Recognize: The questions interact.

Deliberation

- Given a good solution, provide information relevant to implementing the decision and to reconsidering the model.
- Think: post-solution analysis, sensitivity analysis, post-evaluation analysis, etc.
- The candle-lighting principle: “It is better to light one candle than to curse the darkness.” [Kimbrough et al., 1993, Branley et al., 1997]
- The best response to a tight constraint may well be to loosen it.
- Generalize the point: other aspects of the model. Uncertainty, action.

Shadow prices and reduced costs

- From linear programming. Address the candle-lighting principle.
- How much would it be worth to loosen these constraints?
- How much would it cost us to tighten these constraints?
- And similar questions.
- All in support of deliberation.
- How can these questions be addressed outside of linear programming models?

Solution Pluralism & Two Verities

- **Goal of Optimization**
  $\approx$ Given a constrained optimization model, the primary problem posed is to find an optimal solution to the model.
- **Justification of Heuristics**
  $\approx$ Exactly optimal solutions (to optimization problems) are always preferred, but heuristically optimal solutions are acceptable if exactly optimal solutions are not available.

What's not to like?

Truth, But Not Enough

- Each verity misses some important aspects of its topic.
- This talk explores some of these missing aspects.
- In a nutshell regarding the Goal of Optimization verity: "Solution pluralism" captures more of the truth. Why not *all* optimal solutions? What about other solutions of interest?
- In a nutshell regarding the Justification of Heuristics verity: Heuristics have an essential role in solution pluralism; they are how we find the plurality of solutions.

Now, cases and examples to illustrate.

Structure of the Presentation

- Part 1: This deck.
  Focus on examples, arising in practice, of where finding multiple solutions to optimization problems is practicable and useful.
- Part 2: The second deck.
  Focus on the “how” of solution pluralism. The engineering and the scientific research challenges.
- Note a basic pattern: problem $\rightsquigarrow$ model $\rightsquigarrow$ defining SoIs $\rightsquigarrow$ finding SoIs $\rightsquigarrow$ deliberation and decision making aided by discovered SoIs.

Outline


HAVERHILL
AMESBURY
SALSBURY
METHUEN
N.PORT
BRADFORD
NEWBURY
ANDOVER
ROXFORD
ROWLEY
MIDDLETON
TOPSFIELD
IPSWICH
LYNNFIELD
HAMILTON
WENHAM
DANVERS
BEVERLY
MANCHESTER
CLOCESTER
SALEM
MARBLE HEAD
LYNN
CHELSEA

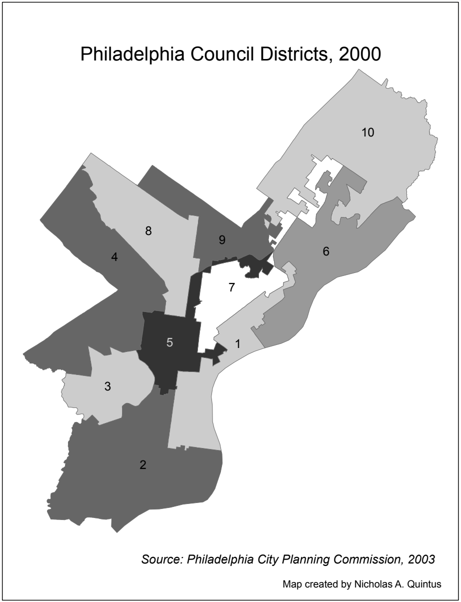
Philadelphia Council Districts, 2000
10
8
9
6
4
7
5
1
3
2
Source: Philadelphia City Planning Commission, 2003
Map created by Nicholas A. Quintus

Team Fred

- In 2011, a contest in Philadelphia, sponsored by Azavea, the *Philadelphia Daily News*, and WHYY, the local National Public Radio and television stations.
- Team Fred: Fred Murphy, Ram Gopalan, Nick Quintus, and SOK.
- What we did, what happened.
  - Team Fred Web page: [Gopalan et al., 2012b].
  - Won on most compact plan.
  - Identified 116 distinct, legally valid plans without breaking ward boundaries.
  - Testified that the most compact plan should *not* be adopted.
  - Testified that many of the 116 plans we found are quite good and can serve as starting points for deliberation.

Philly Districting and Solution Pluralism

- Wagner Prize competition
  Video and slides for presentation [Gopalan et al., 2012a].
- Resulting papers:
  [Chou et al., 2012, Chou et al., 2013, Gopalan et al., 2013]

The map City Council adopted from the 2000 census

- Terribly gerrymandered.
  Note well: Not partisan, since nearly everyone is a Democrat in Philadelphia.
- Designed with scant regard for the Latino community.
- But Maria D. Quiñones-Sánchez won anyway ...
  `http://philadelphiacitycouncil.net/council-members/councilwoman-maria-d-quinones-sanchez-7th-distr`
- And the Latino population of Philadelphia increased by about 58,000 from the 2000 census.

Philadelphia Council Districts, 2000

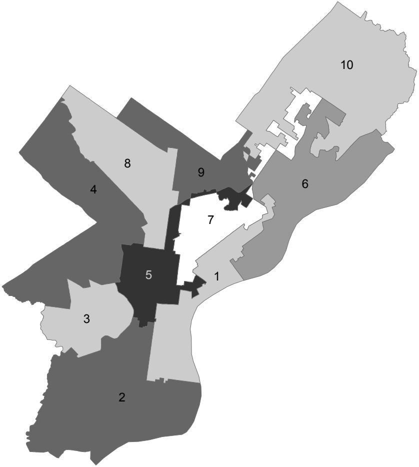

Source: Philadelphia City Planning Commission, 2003

Map created by Nicholas A. Quintus

Our favorite from the 116

- Like all of our 116, it meets all legal requirements (contiguity and population balance).
- Selected visually by “experts” from the 116.
- Does very well on population balance and on honoring existing neighborhoods.
- Does well by the Latinos, too.

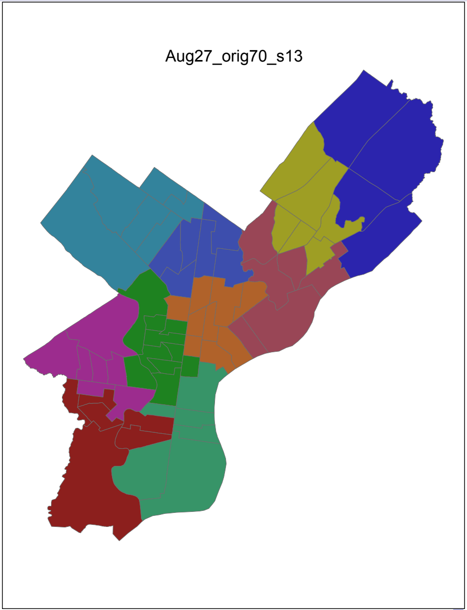
Aug27_orig70_s13

What City Council adopted for the 2010 census

- Improves on the 2000 map.
- Does better by the Latino community.
- Is still quite gerrymandered.
- Does not do very well in terms of honoring Philadelphia neighborhoods.
- Unlike our maps (all of them), this plan splits existing wards.
- Note the Darryl Clark dingle berry in Rittenhouse Square.

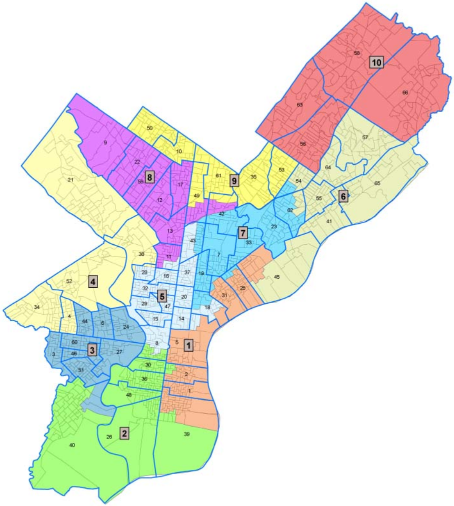
10
8
9
7
6
4
5
3
1
2

Philly districting points of emphasis: 1, Sols

- Solutions of Interest for the plurality of plans problem:
  1. Contiguous
  2. Population-balanced
  3. Comparatively good on compactness

  Note: Contiguity and population balance required for legal validity.
  They can be viewed as a convention to discipline the process, rather than as objectives.

Philly districting points of emphasis: 2, incomplete information

- Solution pluralism can help when models are known to be inaccurate.
- The presence of multiple good solutions to the model can afford reasoned deliberation with information not available to the model.
- Solution pluralism is especially important when the objective function is a proxy for larger and less-well-defined objectives (or different legitimate objectives of different people).

Philly districting points of emphasis: 2 (con't.)

In the face of model inaccuracies we have basically three choices, each of which may be appropriate depending upon circumstances:

1. Do the best we can in building the model, obtain an optimal solution for it, act on an optimal solution, and hope for the best.
2. Insist on an accurate model, then wait to decide until one is available and has a solution.
3. Use the admittedly flawed model to support discussion and deliberation, implementing a modification of an optimal solution to the model in response to considerations not fully represented in the model.

[Chou et al., 2013]

What is an accurate model in this context when the objective is ill-formed?

Philly . . . points of emphasis: 3, strategic decisions

- Districting (like many decisions) is inherently strategic. Many stakeholders and many interests.
- Solution pluralism, combined with optimization, can be used to set the rules of the game, to design the institution, so as to keep play within socially desirable bounds.
  Fred Murphy, in one of many of our conversations:

  > *We solve models as a first step to solving problems. Again, the important feature is that voter districting is a game, not an optimization. We use optimization on subjective criteria to constrain the game to make democracy more democratic.*

- Our proposal for doing districting ≈ define the SoIs *ex ante*, then let parties compete to find them and advocate particular ones. “Pick one you like and bargain from it.”

Philly districting points of emphasis: 4, OR heuristics

- We used classical OR methods to find heuristically good solutions.
- We used these good solutions as starting points for subsequent search.
  - Manually, in the case of compactness.
  - With evolutionary computation for 116.

  See forthcoming *Interfaces* paper [Gopalan et al., 2013].
- Larger pattern: Fast(er) heuristic finds initial solution(s); slower but more probing heuristics follow to improve upon and/or find a plurality of good alternatives to the initial solution(s).
  Call this: *Fast and good, then slow and better*.

Outline

A Tale from Practice

- “The [firm] that [I] dare not speak its name.”
- Rapidly-changing environment. Need to “re-plan” frequently.
- Wanted: continual updating of a pool of good solutions.
- Originally conceived under a DARPA project.
- Here: a new job arrives; quickly insert it into a number of good solutions going forward, and adjust plans accordingly.
- Also: new information arrives that changes the priority status of one or more jobs; quickly find an accommodating schedule.
- Illustration on the next slide.

```
In [8]: data1
Out[8]:
```

	0	1	2
0	453.418213	13	[-1, 23, 3, 0, 7, 21, 24, 16, 1, 20, 12, 6, 19...
1	453.650808	18	[-1, 23, 3, 0, 24, 16, 1, 21, 7, 20, 12, 6, 19...
2	453.418213	14	[-1, 23, 3, 0, 7, 21, 24, 16, 1, 20, 12, 6, 19...
3	462.963644	9	[-1, 9, 14, 10, 2, 13, 5, 22, 17, 15, 8, 11, 1...
4	463.835443	2	[-1, 23, 6, 19, 4, 18, 11, 8, 15, 17, 22, 5, 1...
5	462.107322	11	[-1, 3, 12, 20, 0, 7, 21, 24, 1, 16, 14, 9, 2,...
6	458.961756	12	[-1, 23, 3, 20, 12, 6, 19, 4, 18, 11, 8, 15, 1...
7	475.856737	16	[-1, 23, 3, 6, 19, 12, 20, 0, 7, 21, 24, 1, 16...
8	468.578015	15	[-1, 9, 14, 16, 1, 24, 21, 7, 0, 23, 3, 20, 12...
9	475.348411	4	[-1, 2, 13, 10, 5, 22, 17, 15, 8, 11, 18, 4, 1...
10	482.522823	10	[-1, 16, 1, 24, 21, 7, 0, 20, 12, 3, 23, 6, 19...
11	480.695366	1	[-1, 2, 10, 14, 9, 16, 1, 24, 21, 7, 0, 20, 12...
12	463.835443	5	[-1, 23, 6, 19, 4, 18, 11, 8, 15, 17, 22, 5, 1...
13	473.688082	6	[-1, 9, 2, 13, 4, 18, 11, 8, 15, 17, 22, 5, 10...
14	462.107322	3	[-1, 3, 12, 20, 0, 7, 21, 24, 1, 16, 14, 9, 2,...
15	489.153730	7	[-1, 9, 14, 10, 5, 22, 17, 15, 13, 2, 8, 11, 1...
16	481.202988	8	[-1, 13, 5, 22, 17, 15, 8, 11, 18, 4, 19, 6, 2...
17	485.481078	0	[-1, 2, 9, 14, 10, 5, 22, 17, 15, 13, 8, 11, 1...
18	472.392954	17	[-1, 23, 3, 6, 8, 11, 18, 4, 19, 12, 20, 0, 7,...
19	491.076392	19	[-1, 2, 9, 7, 21, 24, 1, 16, 14, 10, 13, 5, 22...

Figure : TSP with 25 cities, 20 runs of a heuristic. Raw data.

What might we do with such data?

1. New job arrival.
   Inserting the job and recalculating path lengths is fast, even for hundreds of alternate solutions. And parallelizable.
2. Priority changes.
   Consider a special interest in city 24. In the best solution, it's the 6th city visited.
   - What would it cost to move city 24 to an earlier spot? To 4th $\Rightarrow$ 0.2326.
   - City 20? From 9 to 3 (in 5 & 6), at a cost of < 9.

   Easy to imagine scores of interesting, potentially useful questions that can be answered (heuristically).

Outline

GAP4_2

- A representative GAP test problem.
- OR-Library, J.E. Beasley. `http://people.brunel.ac.uk/~mastjjb/jeb/info.html`, [Beasley, 2009].
- c530-2 644. 5 machines, 30 jobs. 644 at optimum.
- Solved with FI-2Pop GA.
- SoIs = FoIs + IoIs (feasibles of interest, infeasibles of interest)
- FoIs: top 1000 feasible solutions, ranked by objective value.
- IoIs: top 1000 infeasible solutions, ranked by distance to feasibility.

Can we find an optimal solution? Yes. Two!

Optimal Solution A										
	0	1	2	3	4	5	6	7	8	9
0	–	3	3	5	1	2	1	3	1	4
1	2	3	2	1	4	5	5	2	4	5
2	3	4	5	3	4	2	1	4	1	5
3	2	–	–	–	–	–	–	–	–	–

Optimal Solution B										
	0	1	2	3	4	5	6	7	8	9
0	–	3	3	5	1	2	1	4	1	4
1	2	3	2	1	4	4	5	2	2	5
2	3	4	5	3	5	3	1	4	1	5
3	2	–	–	–	–	–	–	–	–	–

Now let's compare them on their constraint slacks

Comparison of Slacks					
Solution	1	2	3	4	5
A=644	2	5	0	1	1
B=644	2	1	1	2	0

Decision makers may well prefer A over B or vice versa.

Shadow-price-like questions

Obj. Val.	1	2	3	4	5
A=644	2	5	0	1	1
B=644	2	1	1	2	0
643	2	5	1	4	1
643	2	1	1	2	0
643	0	4	1	11	0
643	5	4	1	3	0
643	2	1	1	3	0
642	4	0	1	0	1
642	2	1	0	1	0
642	2	5	4	1	1
642	5	4	3	1	0
642	0	4	3	9	0
641	2	1	1	4	0
641	2	5	0	3	1

- Can the additional slack resources be usefully deployed?
- On which resources do we *not* have much opportunity for redeployment? (Ans: 2 & 5.)

Shadow-price-like questions

- In both of my optimal solutions, constraint 1 has a slack of 2. Are there any good solutions available with a slack of at least 10?
- Yes. #23 has a slack on constraint 1 of 10 and an objective value of 640; #53, 11, 638; #97, 13, 636.
- What about constraint 2, which already has a slack of 5? #18, 641, 8; #68, 637, 15; #74, 637, 18.
- 3? At 0 or 1. #10, 642, 4; #39, 639, 8; #48, 638, 10.
- 4? At 1 or 2. #5, 643, 11.
- 5? At 0 or 1. #35, 639, 9.
- Plus, we can do combinations...

Slacks for top 100 feasible solutions

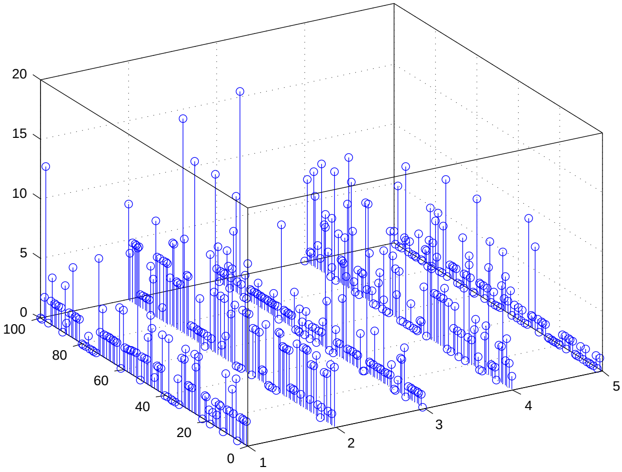

Reduced-cost-like questions

- Job 25 is assigned to machine 2 in solution A and machine 3 in solution B. What's the best solution if we assign job 25 to machine 1?
  Answer: It has an objective value of 643. The solution is

Job 25 to machine 1										
0	0	1	2	3	4	5	6	7	8	9
0	–	3	3	5	1	2	3	4	1	4
1	2	4	2	1	4	5	5	2	2	5
2	3	4	3	3	5	1	1	4	1	5
3	2	–	–	–	–	–	–	–	–	–

Its slacks:

Obj. Val.	1	2	3	4	5
643	2	1	1	3	0

Shadow-price-like questions

Obj.	1	2	3	4	5
645	0	4	1	-1	5
644	-1	5	1	3	0
645	-2	2	1	4	1
644	2	5	0	-2	0
646	-2	2	0	1	1
645	0	-2	0	3	1
644	-2	8	0	0	2
646	5	-2	0	5	0
645	-2	0	7	0	1
648	0	-2	1	0	1
644	2	-2	1	11	-2
644	2	-2	1	-2	1
646	-2	-2	1	2	0
645	-2	8	-2	4	1

- Is there opportunity to acquire more resources?
- To trade in slacks and surpluses?

Shadow-price-like questions

- We're at 644, but we need to get to 651. Show me what I need to do.
- In the top 1000 IoIs there are 75 discovered solutions with objective values at or above 651.

Shadow-price-like questions. 651?

Obj.	1	2	3	4	5
A=644	2	5	0	1	1
B=644	2	1	1	2	0
651	-6	-9	1	1	0
651	-3	-10	1	5	0
651	5	1	1	-4	-12
651	2	5	2	-1	-13
651	2	-6	0	-13	1
651	-10	-2	1	11	1
651	-6	-13	1	5	0
651	5	-7	4	-12	0
651	-5	-7	-5	-1	-12
651	2	-6	-13	6	1
651	-13	4	1	8	0
651	2	-9	0	5	-12

- Is there opportunity to acquire more resources?
- To trade in slacks and surpluses?

Shadow-price-like questions. >651?

Obj.	1	2	3	4	5
654	-9	0	-9	5	1
654	0	-6	1	-8	1
654	-6	-5	1	-4	1
653	2	-6	1	-12	1
653	-5	-11	0	6	1
653	2	-13	1	-3	0
653	2	-10	-1	-4	0
652	-2	4	-10	-5	0
652	-3	4	-9	5	0
652	-5	5	0	-8	1
652	-8	5	0	-2	1
652	2	1	1	-9	0
652	-5	1	1	-7	0
652	2	1	1	-6	1

- Is there opportunity to acquire more resources?
- To trade in slacks and surpluses?

Reduced-cost-like questions

- Job 20 is assigned to machine 3 in solutions A and B, and to 3 (and sometimes 4) in all the FoIs.
- There are 57 solutions in the IoIs in which we assign job 20 to machine 2.
- What are the most promising of these solutions?

Reduced-cost-like questions. 20 to 2?

Obj.	1	2	3	4	5
647	2	-2	0	6	-5
646	2	-2	-12	-8	1
646	2	-2	1	9	-5
645	0	-6	0	-5	-12
645	-8	-8	1	2	8
645	2	-2	-4	-9	9
645	4	-7	1	5	-5
645	4	-7	-1	6	-3
645	-5	-3	1	1	8
644	-8	-12	0	-1	-1
644	2	1	-4	2	-13
644	13	-2	0	1	-3
644	2	4	0	1	-12
644	2	-2	1	11	-2
644	0	1	-7	7	-5
644	0	1	1	9	-5
644	-2	1	7	0	-4

- Is there opportunity to acquire more resources?
- To trade in slacks and surpluses?
- Note: We do *not* find an opportunity with 20 assigned to 2 *and* $z > 647$.

GAP4_2 lols

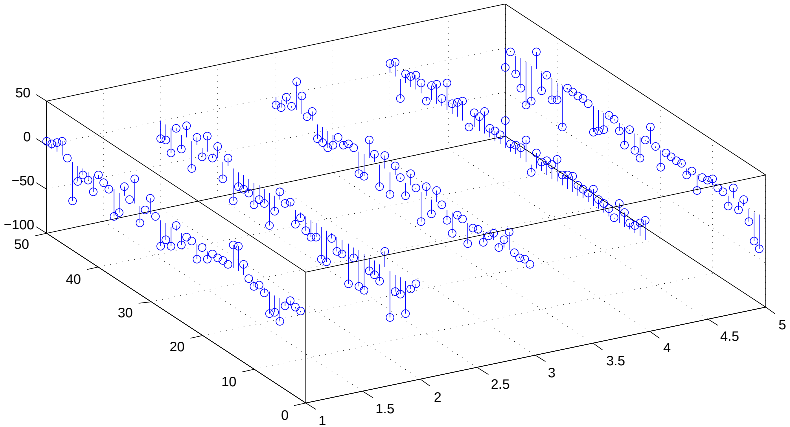

Outline

A Tale from Practice (con't.)

- "The [same firm] that [I] dare not speak its name."
- Recall: Rapidly-changing environment. Need to "re-plan" frequently.
- Wanted: continual updating of a pool of good solutions.
- Here: new jobs arrive, or traffic conditions change, or there is an unanticipated speedup or delay or . . . . Quickly insert the new information into a number of good solutions going forward, and adjust plans accordingly.
- Also: new information arrives that changes the priority status of one or more jobs; quickly find an accommodating schedule.
- Illustration on the next slide. (RJR—"rotate, jiggle, and repair"—heuristic.)

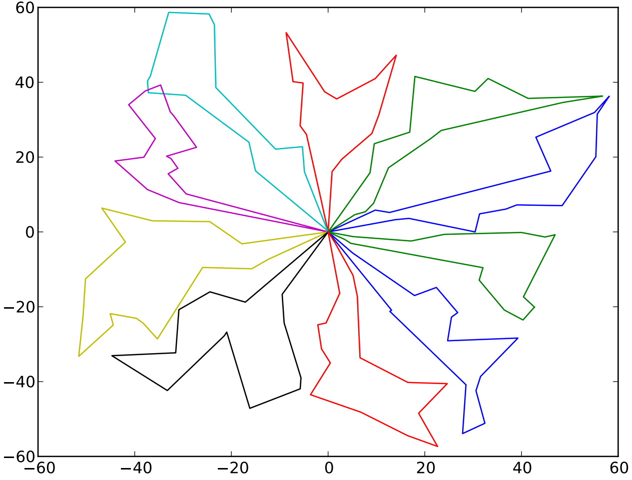

Figure : VRP with 150 jobs, 10 servers, with many constraints

What might we do with such data?

1. Note: the above RGS (reasonable general schedule) presents one of very many good quality solutions arrived at by a fast heuristic, allowing for rapid response to changing circumstances. (When the environment is stable, heuristic search ensues with slower, more exploratory algorithms.) $\Rightarrow$ It's a good solution given business constraints of balanced loads among servers.
2. New job arrival.
   Inserting jobs and recalculating path lengths is fast, even for hundreds of alternate solutions. And parallelizable.
3. Priority changes, new infeasibilities, etc. Again, can be handled with reasonable rapidity by modifying existing solutions.

Outline

Markets and two-sided matching

(Material drawn largely from [Kimbrough and Kuo, 2010].)

- Two-sided matching problems.
  Sides: X and Y. Problem: form pairs, one member from each side.
  Model for a market: buyers and sellers, etc.
- Distributed versus centralized markets.
- Market failures and the move to centralization.
  Labor markets: interns and hospitals. Admissions: students and schools.
  Widely practiced in the USA. And see [Bodin and Panken, 2003].

Markets and two-sided matching

- Agents and strategic considerations.
  Assignment versus matching.
- Stability.
- Equity.
- Social welfare.

Representative representation

- Sides X and Y. Pairs $(x, y), x \in X, y \in Y$. Match array $\mathbf{M}$ is size $|X| \times |Y|$.

$$\mathbf{M} = \begin{matrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{matrix}$$

or

	y_1	y_2	y_3	y_4
x_1	0	0	0	1
x_2	0	1	0	0
x_3	1	0	0	0
x_4	0	0	1	0

- This is the Simple Marriage Matching problem.

Simple Marriage Matching problem

- *n* men, *n* women, all of whom want to be paired up with a member of the complementary gender.
- This is very stylized. Lots of variations possible (and considered in the literature).
- More generally on matching, many elaborations possible and encountered in practice.
    - “Two-body” problem for physicians, "*n*-body" for students.
    - Admissions problem: more physicians than hospitals, more students than schools.
    - More complex preference structures: Incomplete preferences, indifference, etc.
    - Various ad hoc, special considerations.

Simple Marriage Matching problem

- Focus for the moment: the very elementary, stylized Simple Marriage Matching problem.
- Note:
  1. It is prototypical, representative in important ways.
  2. It is most favorable to the standard approach, which we are challenging, or rather seeking to augment.
     Our points are only strengthened for less stylized, more representative models.

The Gale-Shapley deferred acceptance algorithm

- The algorithm (informally):

1. Do until there are no unaccepted men:
    1. Each currently unaccepted man proposes to his most preferred woman, provided she hasn't already rejected him.
    2. Each woman with more than one proposal rejects all but her most preferred proposer, who becomes perforce currently accepted by that woman.
2. Stop. The matching is determined by the currently accepted proposals.

- Alternate version: the women propose, the men dispose.

The Gale-Shapley deferred acceptance algorithm

- Stable and unstable matches. The equilibrium concept for matching.
- Key properties of GS/DAA.
  - For special cases of matching (including the Simple Marriage Matching problem), guaranteed to terminate (quickly) and to find a stable match. (And for other problems, it isn't, doesn't.)
  - Deterministic. It can find only two distinct stable matches (one if there is only one).
  - Asymmetric. Proposers get optimal match, disposers get pessimal match.
  - Blind to considerations other than stability, e.g., fairness, social welfare.

Can we do better?

- Key property of matching problems: may be exponentially many stable matches.
- Can we find stable matches that are better on fairness and/or on social welfare?
- What about 'nearly-stable' matches?
- Must rely on heuristics.

Larger question: Can we design high-quality centralized markets that will work well in practice?

Matching with a Simple ABM

- Model of a market.
- Agents epistemically powerful; see the entire field.
- Agents randomly get the floor and make a feasible swap that is most favorable to them.

ABM Results: $n = 40$. Transaction Cost = 0 (†Mean=3, Max=124)

	1st Qu	Median	3rd Qu
Init. # unstable pairs	314	340	367
Final # unstable pairs†	0	0	0
InitialSocialWelfareSum	1570	1638	1709
Final SocialWelfareSum	476	499	529
Initial Equity	492	532	572
Final Equity	200	224	245
SwapCount	1610	5731	22389
InitialSumXScores	769	819	869
Final SumXScores	217	246	280
InitialSumYScores	772	820	869
Final SumYScores	224	253	293
Number of Runs	100		
Number of Replications	100		
TransactionCost	0		

ABM Results: $n = 40$

Points arising:

1. The initial match, which is randomly generated, is normally a very poor one. It is not stable, and it fares poorly on social welfare, equity, and individual scores.
2. In driving towards a stable solution, however myopically, the market process reliably improves the scores on social welfare and equity, as well as the individual scores. Here is an invisible hand at work with broadly sanguine effects.
3. There is, however, a large churning cost. The median number of swaps is 5731 to achieve stability in the society. Divided by 40, that's just over 573 breakups experienced on average by each agent ($5731 \times 4/40$). What a cruel system that would break so many hearts.

ABM Results: $n = 100$

	1st Qu	Median	3rd Qu
Init. # unstable pairs	1709	1808	1908
Fin. # unstable pairs	0	0	0
Init. # unstbl pairs NTC	2061	2165	2270
Fin. # unstbl pairs NTC	28	32	37
InitialSocialWelfareSum	9830	10104	10376
Final SocialWelfareSum	1984	2048	2113
Initial Equity	3173	3329	3492
Final Equity	878	934	994
SwapCount	655	1049	1901
InitialSumXScores	4859	5050	5242
Final SumXScores	941	1020	1106
InitialSumYScores	4858	5055	5249
Final SumYScores	939	1018	1103
Number of Runs	100		
Number of Replications	100		
TransactionCost	5		

ABM Results: $n = 100$

Points arising:

1. The behavior when $n = 100$, with and without transaction costs, is broadly similar to cases with $n = 40$ and smaller. That is, the market actions by myopic agents generally improve both individual and social welfare, from a random start.
2. Equilibrium, in the form of stability, is *not* achieved unless we factor in substantial transaction costs.
3. Even when equilibrium is attained, it is remarkably costly in terms of the number of matches made and abandoned.

Implemented a GA solver

- Solutions are permutations, search in permutation space.
- Every permutation is feasible.
- Used rather conventional approach (see paper for details).
- Reporting results on randomly-generated test problems. First: 25, 20$\times$20 problems.

Comparisons with GS/DAA: GA. Points arising:

1. In these 25 test problems, the GA is able to find very many stable matches that are Pareto-superior to the GS/DAA solutions. By Pareto-superior, I mean stable and strictly better than the GS/DAA solutions on both social welfare and equity.
2. The GA and the ABM do about equally well at the task of finding solutions that are Pareto-superior to those of GS/DAA.
3. These points are illustrated vividly in Figure 3, which is for test case 7.
4. We also used the GA to find and collect 'one-away' solutions, that is, matches with only one unstable pair of couples. When we relax the requirement for strict stability, many more attractive options appear. See Figure 4 as an example.

Comparisons with GS/DAA

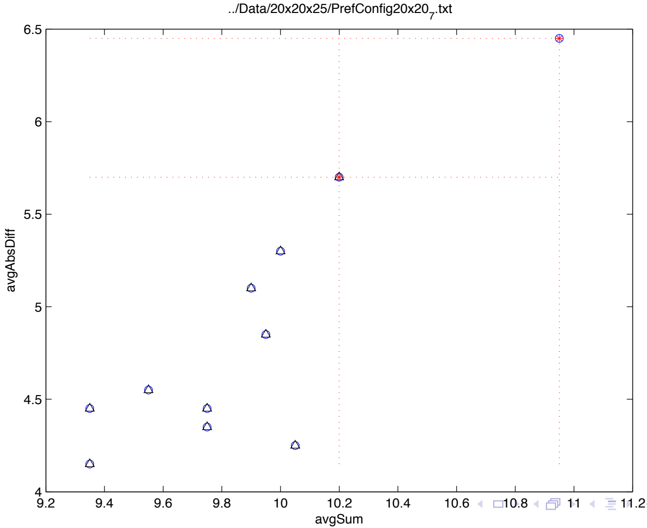

One-away solutions. 40×40.

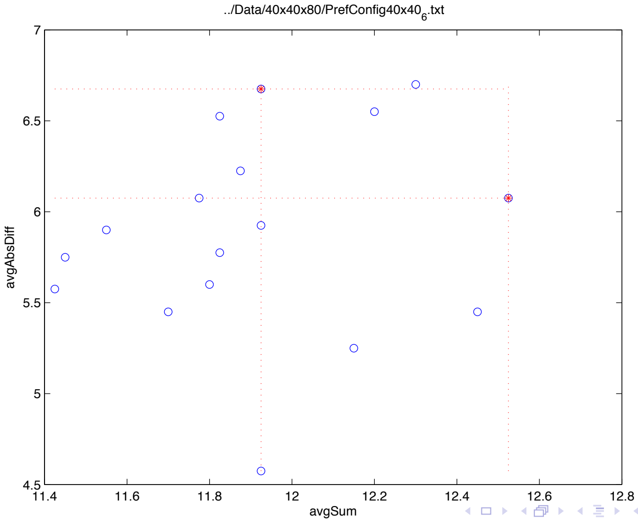

Discussion of Matching

1. Two-sided matching models apply widely.
2. The standard approach . . . Gale-Shapley deferred acceptance algorithm. . . . important virtues, including computational tractability. Operates without regard to social welfare and in contravention to equity.
3. Many stable matches in a two-sided matching problem. GS/DAA can only find two of these, which two have very special optimality and pessimality properties.
4. Metaheuristics can find stable matches that GS/DAA cannot find and that are superior with regard to equity and/or social welfare.
5. ‘one-away’ matches with attractive social welfare and equity profiles.

Again: Games and Institution Design

- Recall: Redistricting and the need to design an institution for doing it better.
- Here: Two-sided matching is for “managed markets”, for when free markets fail.
- So: The need to design an institution for doing it better.
- Finding a stable match has its attractions . . .
- But the concept has its limitations:
  1. Fairness, social welfare
  2. Stability under non-restrictive conditions

  And . . .

- On strategy-proofness, definitive results were obtained by the 1980's.
- The impossibility result by [Roth, 1982] (see [Roth and Sotomayor, 1990] for elaboration) shows that there is no mechanism which is incentive compatible for both sides and at the same time able to yield a stable outcome.
- The direct consequence of this result is that we either have to sacrifice stability to prevent the strategic misrepresentation of preferences, or if we want to always achieve stability at least one side has this type of strategic incentive.
- Very little is known to date regarding the practical consequences of this fact, although [Roth, 2008] is sanguine.

Solution Pluralism to the Rescue?

- Define SoIs, e.g., stable or “nearly-stable” solutions that dominate DA on equity and social welfare, and are not dominated by other SoIs.
- Use heuristics to find multiple solutions among the SoIs.
- Pick one of the discovered SoIs at random.

Do you think you can exercise strategic manipulation in this context?
Compare with our proposal for redistricting. These are forms of institution design predicated on solution pluralism.

Outline

Multiple reasons for seeking a plurality of solutions.

1. The problem is multi-objective.
   A known, well-established reason. Our examples: Philly districting, TSP, GAP, VRP, matching, etc. can be understood this way.
2. For doing post-solution analysis.
   Original motivation. GAP examples above (as well as others).
3. Imperfect models. “We solve models, not problems.”
   Philly districting is a prime example, but the other examples fit as well. Apply extra-model knowledge to a pool of good solutions.
4. The decision is strategic, not parametric.
   As in Philly districting, matching, etc. Use the plurality of solutions to constrain choices to what’s reasonable. Solution pluralism for institution design.
5. Speed to solution is essential, but there may be time to think.
   (i) Use an existing pool of solutions to achieve fast response to changed conditions. (ii) Use multiple solutions obtained heuristically to seed subsequent search. TSP and VRP examples here; DARPA ALP.

Solution Pluralism in Established Practice

See [Kimbrough et al., 2011].

1. *Models with stochastic outputs, e.g., simulation models, including discrete event models and agent-based models.* These models normally should be (and are) run multiple times and their outputs analyzed statistically in order to support decision making [Law, 2007]. Typically, one or more measures of performance (MoPs) are defined. Roughly, the SoIs are the MoPs as generated under a variety of specified conditions and scenarios.
2. *Variance-based sensitivity analysis of equational models.* Very many solutions are obtained with randomization of parameter values in order to estimate variances in model outputs. See for overviews [Saltelli et al., 2004, Saltelli et al., 2008]. See also [Law, 2007].
3. *Multi-Objective optimization models.* Here the principal SoIs are the solutions on the Pareto frontier and the problem is to find them. Evolutionary computation and related metaheuristics for these problems are actively being researched and used; see [Coello Coello et al., 2007, Deb, 2001] for overviews.

And in Established Practice

4. Because it is often effective in convincing clients when recommendations are surprising.
   (Thanks to Hansjoerg Fromm for this point.)

Summary

- Key for deliberation support:
  - FoIs: Feasible solutions of Interest: Feasible, non-optimal solutions with good objective values.
  - IoIs: Infeasible solutions of Interest: Infeasible solutions near feasibility with good objective values.
- These can be found by intelligent sampling of the solutions.
- This is just what metaheuristics, such as GAs, do.

Revisiting the Verities

- **Goal of Optimization**
  $\approx$ Given a constrained optimization model, the primary problem posed is to find an optimal solution to the model.
  Defining SoIs and heuristically sampling them well affords extended use of the model.
- **Justification of Heuristics**
  $\approx$ Exactly optimal solutions (to optimization problems) are always preferred, but heuristically optimal solutions are acceptable if exactly optimal solutions are not available.
  In most cases, effective sampling of the SoIs is something that must be done with heuristics.

Beasley, J. E. (Accessed July 27, 2009).
OR-Library.
World Wide Web.
http://people.brunel.ac.uk/~mastjjb/jeb/info.html.

Bodin, L. and Panken, A. (2003).
High tech for a higher authority: The place of graduating rabbis from Hebrew Union College–Jewish Institute of Religion.
Interfaces, 33(3):1–11.

Branley, B., Fradin, R., Kimbrough, S. O., and Shafer, T. (1997).
On heuristic mapping of decision surfaces for post-evaluation analysis.

In Sprague, Jr., R. H., editor, *Proceedings of the Thirtieth Annual Hawaii International Conference on System Sciences*, Los Alamitos, CA. IEEE Press.

Chou, C., Kimbrough, S. O., Murphy, F. H., Sullivan-Fedock, J., and Woodard, C. J. (2013).
On empirical validation of compactness measures for electoral redistricting and its significance for application of models in the social sciences.
Social Science Computer Review, page (forthcoming).
doi: 10.1177/0894439313484262
http://ssc.sagepub.com/content/early/2013/04/09/0894439313484262.

Chou, C., Kimbrough, S. O., Sullivan, J., Woodard, C. J., and Murphy, F. H. (2012).
Using interactive evolutionary computation (IEC) with validated surrogate fitness functions for redistricting.

In Soule, T., editor, *Proceedings of GECCO 2012 (Genetic and Evolutionary Computation Conference)*, pages 1071–1078, New York, NY, USA. ACM.
http://dx.doi.org/10.1145/2330163.2330312.

Coello Coello, C. A., Lamont, G. B., and Van Veldhuizen, D. A. (2007).
Evolutionary Algorithms for Solving Multi-Objective Problems.
Springer Science+Business Media LLC, New York, NY, second edition.

Deb, K. (2001).
Multi-Objective Optimization using Evolutionary Algorithms.

John Wiley & Sons, LTD, Chichester, UK.

Gopalan, R., Kimbrough, S. O., Murphy, F. H., and Quintus, N. (2012a).

The Philadelphia districting contest: Designing territories for city council based upon the 2010 census.
Video on the Web.
https://live.blueskybroadcast.com/bsb/client/CL_DEFAULT.asp?Client=569807&PCAT=5277&CAT=5278.

Gopalan, R., Kimbrough, S. O., Murphy, F. H., and Quintus, N. (2012b).
Team Fred home page.
World Wide Web.
http://opim.wharton.upenn.edu/~sok/phillydistricts/.

Gopalan, R., Kimbrough, S. O., Murphy, F. H., and Quintus, N. (2013).
The Philadelphia districting contest: Designing territories for City Council based upon the 2010 census.

Interfaces.
Forthcoming.

Kimbrough, S. O. and Kuo, A. (2010).
On heuristics for two-sided matching: Revisiting the stable marriage problem as a multiobjective problem.
In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2010)*. ACM.
http://dl.acm.org/citation.cfm?doid=1830483.1830712.

Kimbrough, S. O., Kuo, A., LAU, H. C., Murphy, F. H., and Wood, D. H. (2011).
Solution pluralism and metaheuristics.
In *MIC 2011: The IX Metaheuristics International Conference*, Udine, Italy.

Kimbrough, S. O., Oliver, J. R., and Pritchett, C. W. (May-June 1993).

On post-evaluation analysis: Candle-lighting and surrogate models.
Interfaces, 23(7):17–28.

Law, A. M. (2007).
Simulation Modeling and Analysis.
McGraw-Hill series in industrial engineering and management science. McGraw-Hill, New York, NY, fourth edition.

Roth, A. (1982).
The economics of matching: Stability and incentives.
Mathematics of Operations Research, 7(4):617–628.

Roth, A. E. (2008).
Deferred acceptance algorithms: History, theory, practice, and open questions.
International Journal of Game Theory, 36:537–569.

Roth, A. E. and Sotomayor, M. A. (1990).
Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis.
Cambridge University Press, Cambridge, UK.

Saltelli, A., Ratto, M., and Andres, T. (2008).
Global Sensitivity Analysis: The Primer.
John Wiley & Sons, Ltd., Chichester, England.

Saltelli, A., Tarantola, S., Campolongo, F., and Ratto, M. (2004).
Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models.
John Wiley & Sons, Chichester, UK.

$Id: deliberation-COModel-part1-master-beamer.tex 3660 2013-07-26 09:43:20Z sok $