Interesting Research in Vehicle Routing and Healthcare Analytics

Bruce L. Golden
R.H. Smith School of Business
University of Maryland

Presented at AIRO 2012 Conference, September 2012
Vietri sul Mare, Italy

Outline of Talk

- Some personal remarks
- Vehicle Routing
  - The Hierarchical Traveling Salesman Problem (HTSP)
- Healthcare Analytics
  - The Effects of Bed Utilization on Discharge and Readmission Rates
- Conclusions

Grand Hotel
Convento di Amalfi

Introduction to the HTSP

- Consider the distribution of relief aid
  - E.g., food, bottled water, blankets, or medical packs
- The goal is to satisfy demand for relief supplies at many locations
  - Try to minimize cost
  - Take the urgency of each location into account

A Simple Model for Humanitarian Relief Routing

- Suppose we have a single vehicle which has enough capacity to satisfy the needs at all demand locations from a single depot
- Each node (location) has a known demand (for a single product called an aid package) and a known priority
  - Priority indicates urgency
  - Typically, nodes with higher priorities need to be visited before lower priority nodes

Node Priorities

- Priority 1 nodes are in most urgent need of service
- To begin, we assume
  - Priority 1 nodes must be served before priority 2 nodes
  - Priority 2 nodes must be served before priority 3 nodes, and so on
  - Visits to nodes must strictly obey the node priorities

The Hierarchical Traveling Salesman Problem

- We call this model the Hierarchical Traveling Salesman Problem (HTSP)
- Despite the model's simplicity, it allows us to explore the fundamental tradeoff between efficiency (distance) and priority (or urgency) in humanitarian relief and related routing problems
- A key result emerges from comparing the HTSP and TSP in terms of worst-case behavior

Four Scenarios for Node Priorities

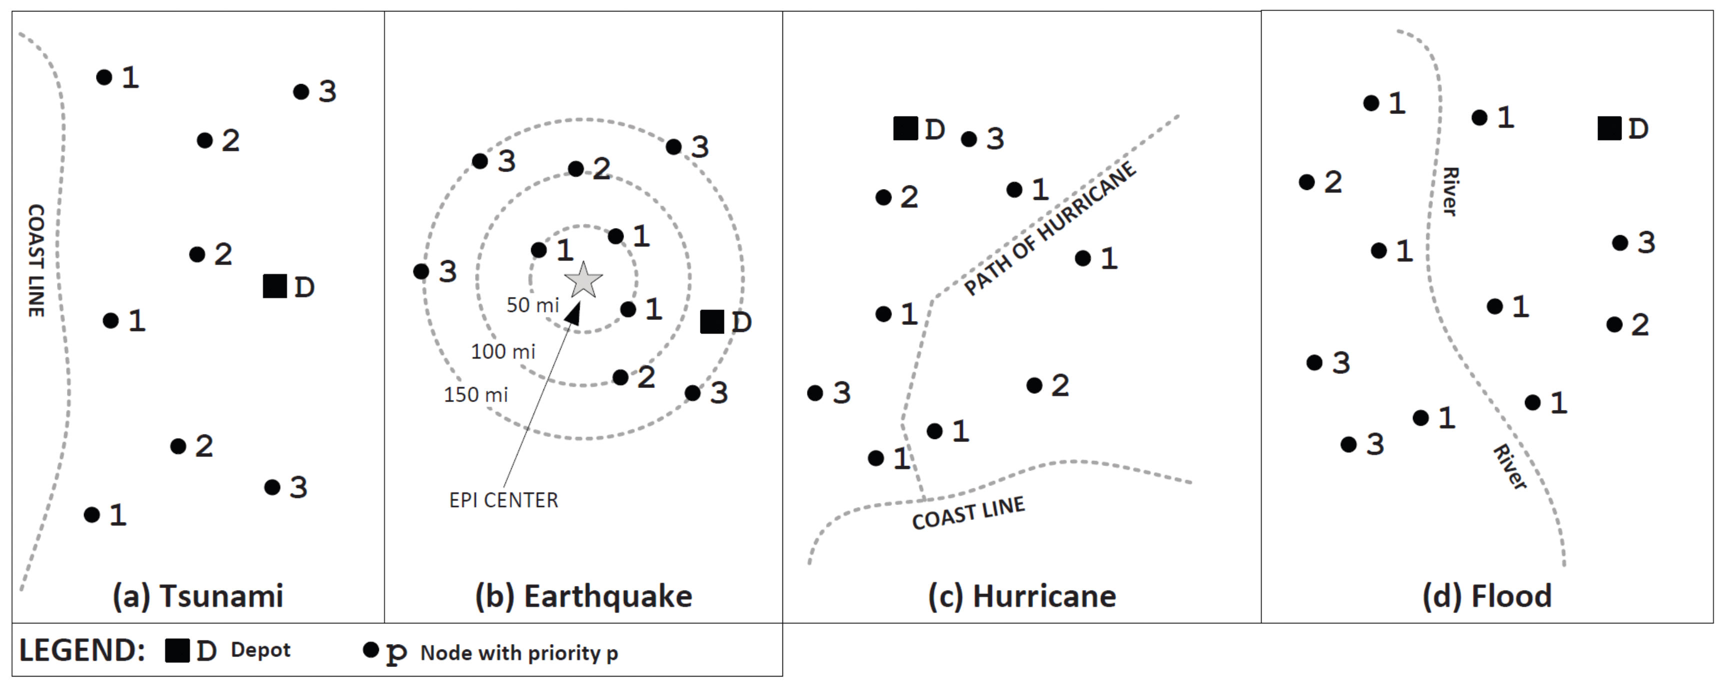

Literature Review

- Psaraftis (1980): precedence constrained TSP
- Fiala Tomlin, Pulleyblank (1992): precedence constrained helicopter routing
- Guttman-Beck et al. (2000): clustered traveling salesman problem
- Campbell et al. (2008): relief routing
- Balcik et al. (2008): last mile distribution

A Relaxed Version of the HTSP

- Definition: The d-relaxed priority rule adds operational flexibility by allowing the vehicle to visit nodes of priority $\pi + 1, \ldots, \pi + d$ (if these priorities exist in the given instance) but not priority $\pi + d + \ell$ for $\ell \geq 1$ before visiting all nodes of priority π (for $\pi = 1, 2, \ldots, P$)
- When d=0, we have the strict HTSP
- When d=P-1, we have the TSP (i.e., we can ignore node priorities)

Efficiency vs. Priority

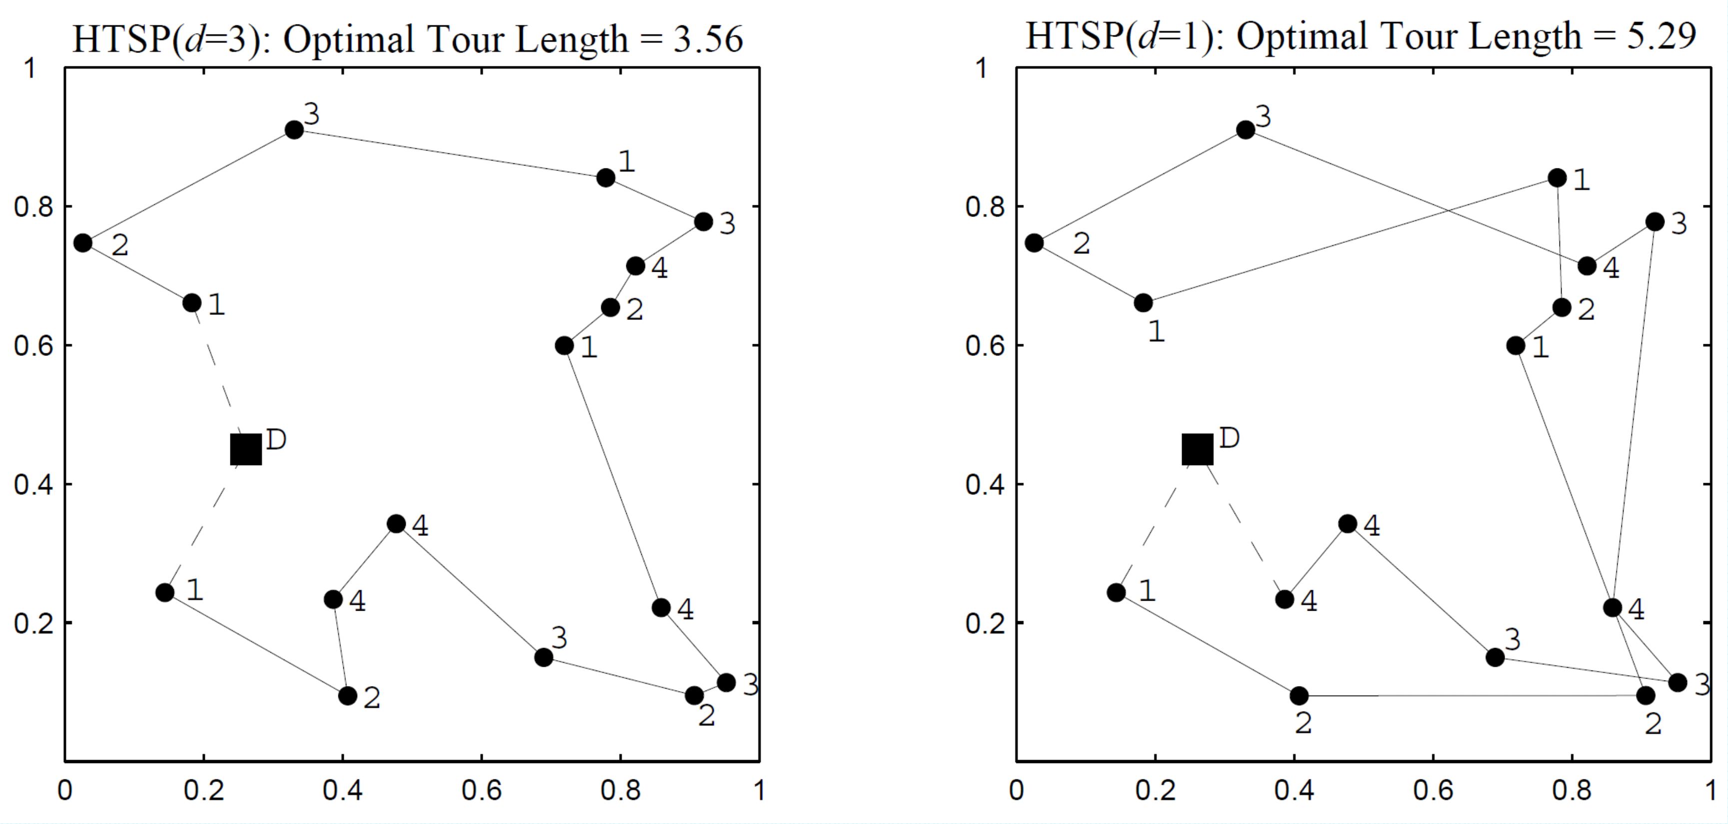

Main Results

- Let P be the number of priority classes
- Assume the triangle inequality holds
- Let $Z^*_{d,P}$ and Z^*_{TSP} be the optimal tour length (distance) for the HTSP with the d-relaxed priority rule and for the TSP (without priorities), respectively
- We obtain the following results

$$\text{(a)}\, Z^*_{0,P} \leq P Z^*_{TSP}$$

$$\text{(b)}\; Z^*_{d,P} \leq \left\lceil \frac{P}{d+1} \right\rceil Z^*_{TSP}$$

Sketch of Proof (a)

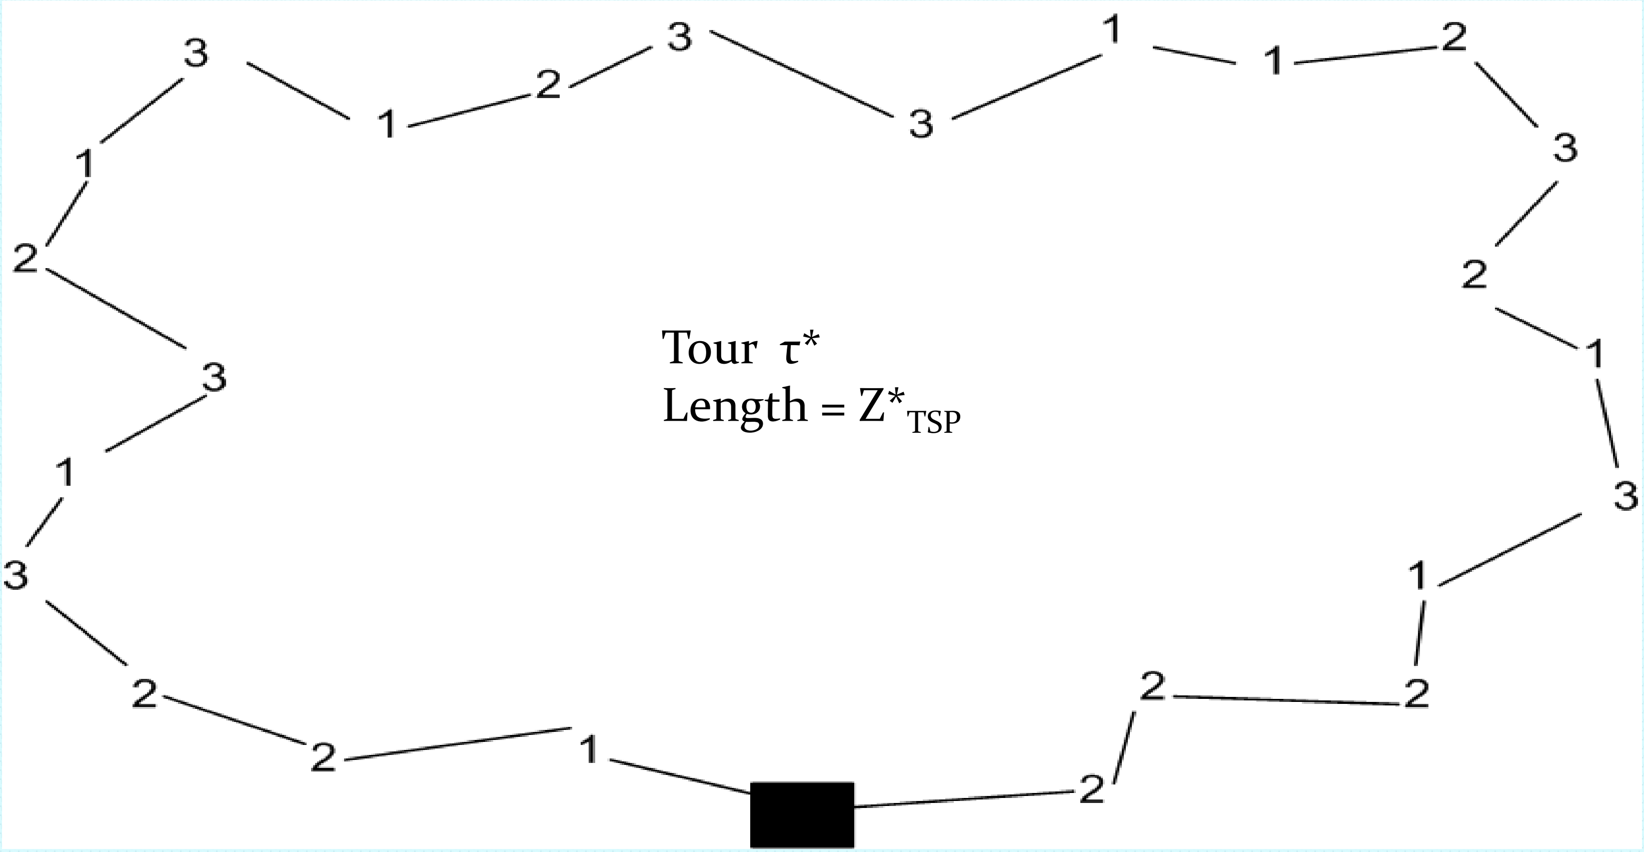

Sketch of Proof (a)

- Construct tours $\tau(1)$, $\tau(2)$, and $\tau(3)$
- Visit nodes in the same order as they appear in τ^*
- From the triangle inequality, the lengths of $\tau(1)$, $\tau(2)$, and $\tau(3)$ are each $\leq Z^*_{TSP}$
- It is easy to construct a feasible solution τ to the HTSP from $\tau(1)$, $\tau(2)$, and $\tau(3)$
- The length of $\tau \leq \sum_{i=1}^{3}$ {length of τ (i)} $\leq 3Z*_{TSP}$

Sketch of Proof of (b)

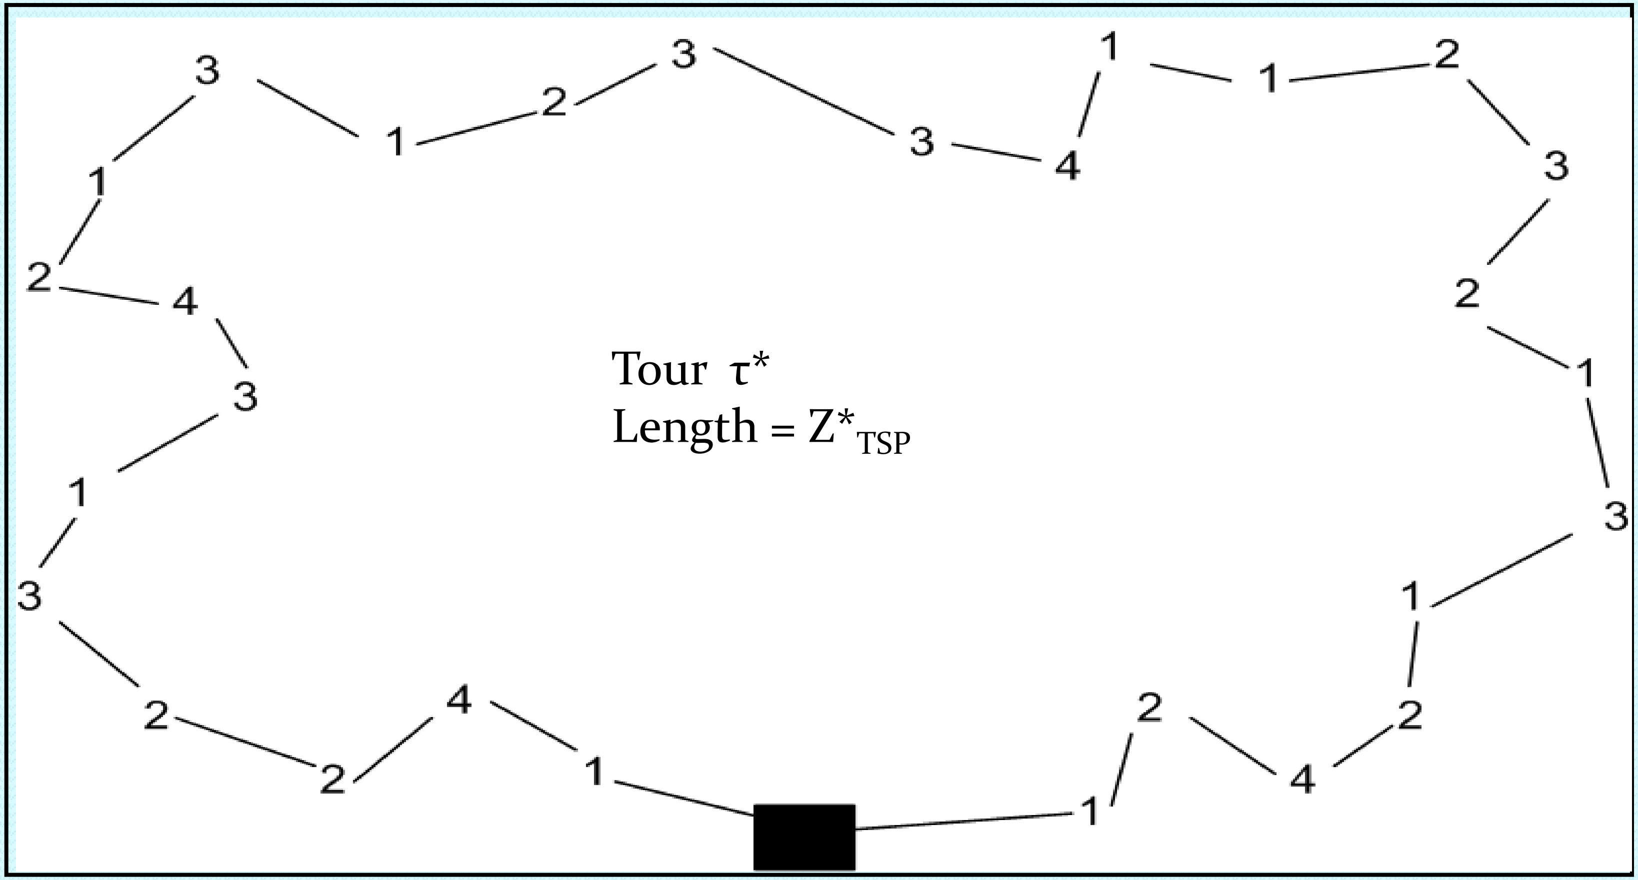

Sketch of Proof of (b)

- In our example, P=4 and d=1
- In the worst case, we can't visit a priority 3 node until we have visited all priority 1 and 2 nodes
- Visit nodes in the same order as they appear in τ^*
- $\tau(1)$ includes priority 1 and 2 nodes
- $\tau(2)$ includes priority 3 and 4 nodes
- As before, we can construct τ from $\tau(1)$ and $\tau(2)$
- The length of $\tau \leq \sum_{i=1}^{2}$ {length of τ (i)} $\leq 2Z*_{\text{TSP}}$

The General Result and Two Special Cases

- $Z^*_{d,P} \leq \left\lceil \frac{P}{d+1} \right\rceil Z^*_{TSP}$

- If d=0, we have part (a)

- If d=P-1, then $Z^*_{d,P} = Z^*_{TSP}$

Worst-case Example

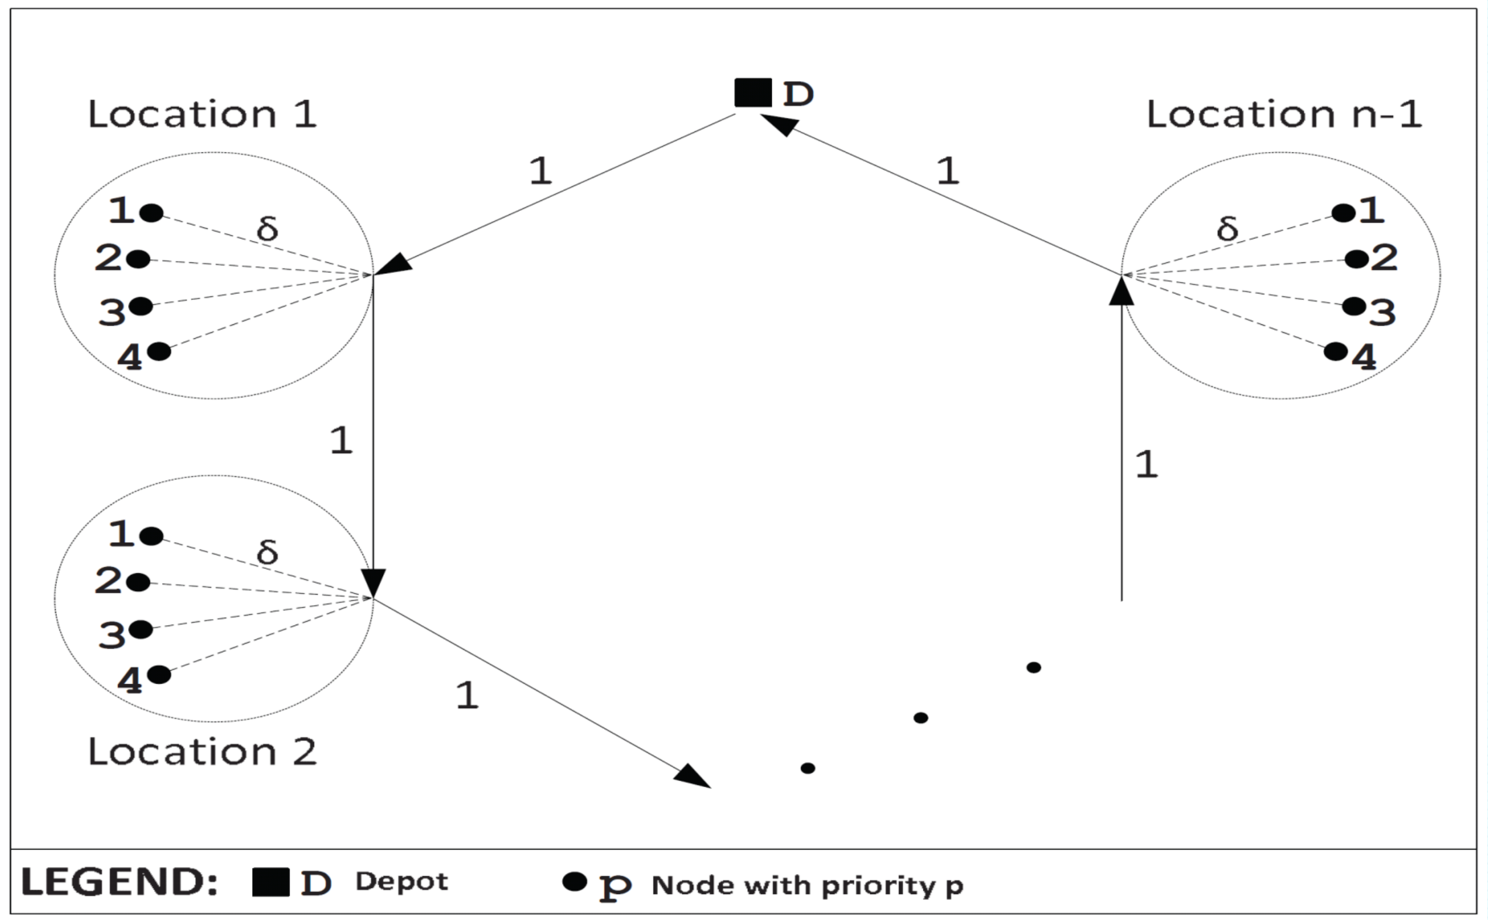

Several Observations

- Observation 1. The worst-case example shows that the bounds in (a) and (b) are tight and cannot be improved
- Observation 2. We can "solve" a TSP over the entire set of nodes using our favorite TSP heuristic and obtain a feasible tour for the HTSP by following the part (b) proof
- Observation 3. Suppose we select Christofides' heuristic and let $Z^h_{d,P}$ be the length of the resulting feasible solution to the HTSP, then we have $Z^h_{d,P} \leq \frac{3}{2} \cdot \left\lceil \frac{P}{d+1} \right\rceil Z^*_{TSP}$

Observations and Extensions

- Observation 4. The HTSP (with d=0) can be modeled and solved as an ATSP
- Observation 5. Other applications of the HTSP include routing of service technicians and routing of unmanned aerial vehicles
- We can obtain similar worst-case results (with tight bounds) for the HTSP on the line and the Hierarchical Chinese Postman Problem (HCPP)

Extensions and Future Work

- The HTSP and several generalizations have been formulated as mixed integer programs

- HTSP instances with 30 or so nodes were solved to optimality using CPLEX

- Future work
  - The Hierarchical Vehicle Routing Problem (HVRP)
  - A multi-day planning horizon
  - Uncertainty with respect to node priorities

Emergence of Healthcare Analytics within INFORMS

Number of Healthcare Talks at INFORMS Annual Meetings

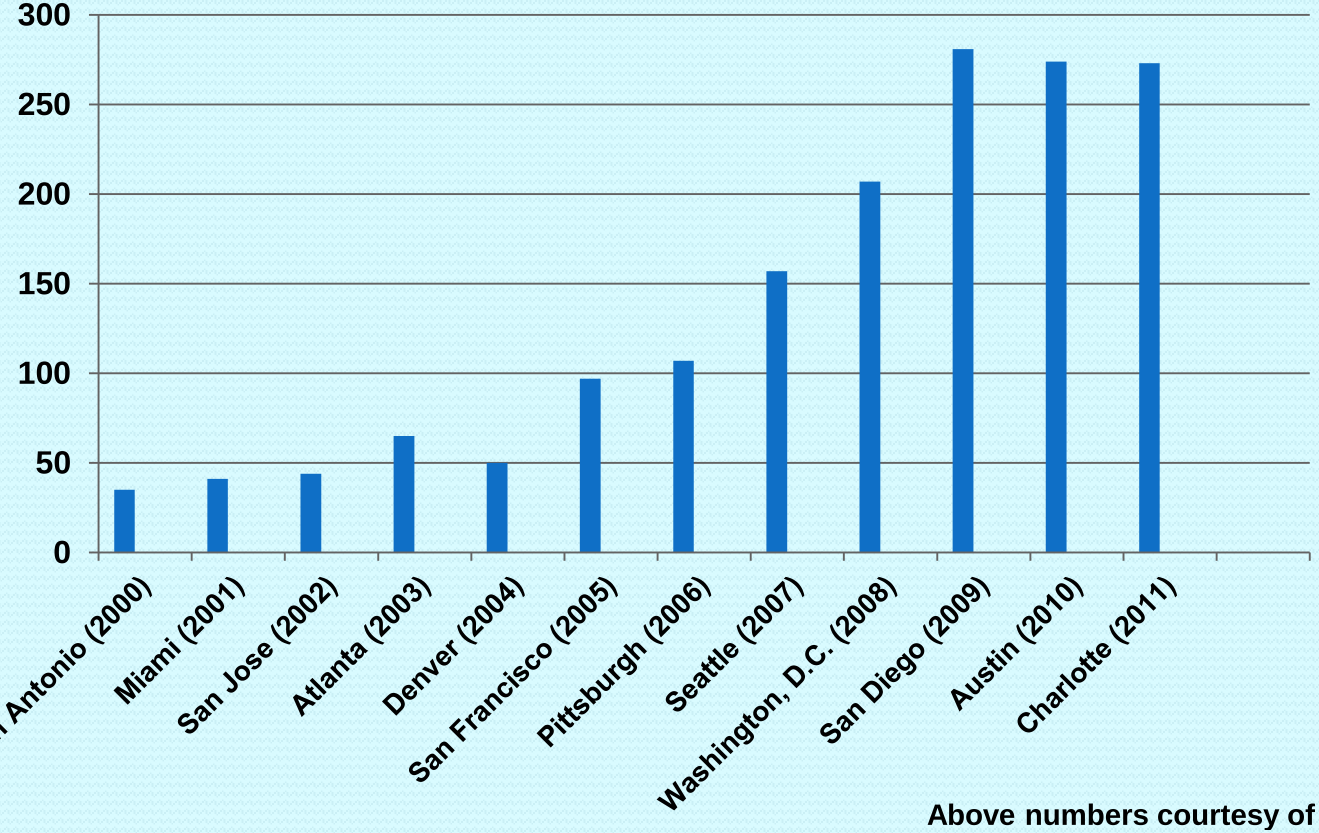

Above numbers courtesy of Brian Denton

Strength in Numbers

- There is more healthcare data available than ever before
  - Careful analysis of healthcare data can lead to smarter decisions, better quality healthcare, and cost savings
- A larger number of healthcare decision makers have MBAs than ever before
  - They understand that we can help
- A larger number of us in OR/OM are working on healthcare applications than ever before

The Effects of Bed Utilization on Discharge and Readmission Rates

- Many hospital resources are required for surgery
  - Operating rooms
  - Nurses & Physicians
  - Anesthesia team
  - Post-operative beds for recovery
- If downstream beds are unavailable, surgery might be postponed or cancelled
- Surgeons decide when patients are discharged
  - Surgeons are paid to do surgery

Research Question 1

- Does the utilization of downstream beds affect the discharge decisions of surgeons?
  - Hypothesis: There is an increased discharge rate on days when post-operative utilization is high

Data

- Data collected on every surgery performed at a large US hospital from Jan 1, 2007 to May 31, 2007
- 7808 patients, of which 6470 were admitted to the hospital and stayed for at least one night
- These patients stayed a total of 35,478 days
- Data provided on age, race, gender, surgical line, date of surgery, discharge date, and surgery type (scheduled vs. emergency)
- Utilization of post-operative beds varies widely

Discharge Rates

- Discharge rates have positive correlation with utilization

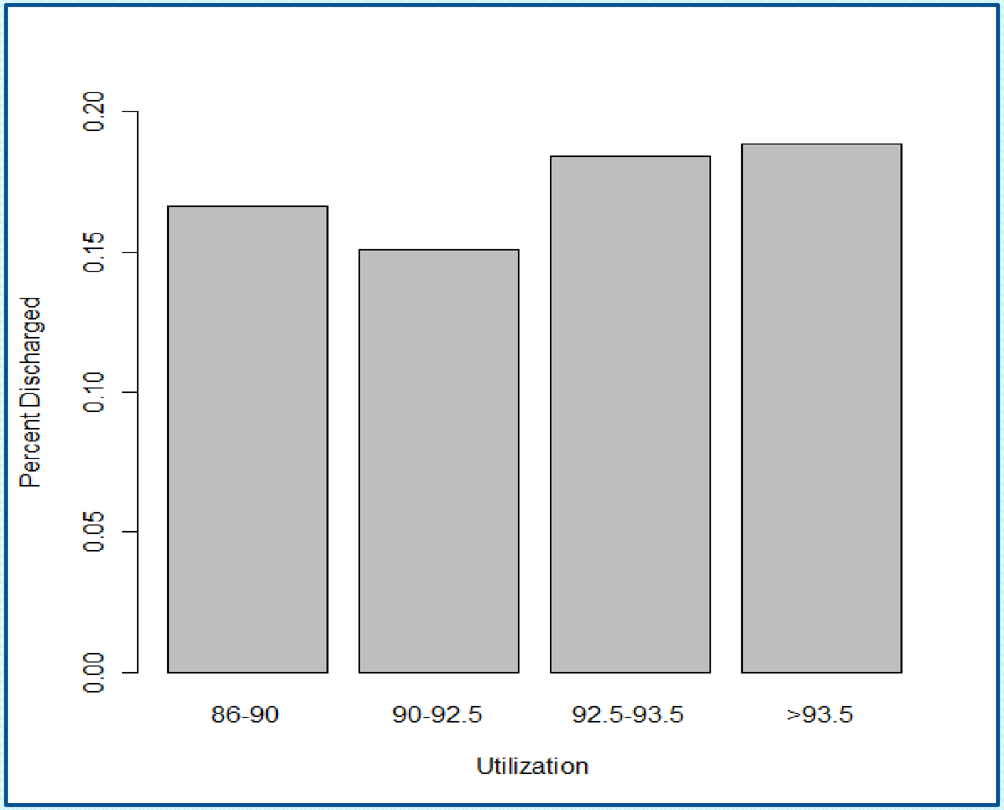

Utilization Measures

- We compute two measures of utilization
  - Discrete measure – a variable that is 1 when utilization exceeds a given threshold (e.g., 93%), and 0 otherwise
  - Continuous measure – a variable that counts the number of beds in use on each day
- Compare marginal effect of each bed in use vs. a discrete change in discharge probability when utilization exceeds a threshold

Discrete Time Survival Analysis

- Can't use logistic regression because observations are correlated -- a patient discharged on the fifth day cannot be discharged on the first four days
- Singer and Willet (1993) show how to handle discrete time survival data
- For each day, we record whether or not each patient is discharged, and use this as the outcome variable
- The outcome variable is regressed on our utilization measures and our control variables
- We control for the patient's age, race, gender, severity, and surgery type

Models and Results

- Model 1: $logit(DISCHARGE) = AGE + ELECTIVE + FULL + CARDIAC\ SURGERY$
  $+ CARDIOLOGY + \ldots + DONOR\ SERVICE + D1$
  $+ D2 + \ldots + D59 + \varepsilon$

- Model 2: $logit(DISCHARGE) = AGE + ELECTIVE + BEDS + CARDIAC\ SURGERY$
  $+ CARDIOLOGY + \ldots + DONOR\ SERVICE + D1$
  $+ D2 + \ldots + D59 + \varepsilon$

- When the utilization threshold is exceeded, the odds of discharge for any given patient increase. The estimate for Full is positive and significant for threshold above 91.5%.

- Each additional bed in use increases the odds that a patient will be discharged. The estimate for Beds is positive and significant.

Observations

- Discharge rates increase as utilization increases, regardless of how utilization is measured
- Either some patients are held too long and discharged when space is needed, or some patients are discharged too early when utilization is high
- Our results cannot distinguish between these two explanations

Research Question 2

- Are patients who are discharged when utilization is high more likely to be readmitted?
  - Hypothesis: An increase in the discharge rate will lead to some patients with shortened lengths of stay. This will cause an increase in the readmission rate for those patients.

Analysis

- Using the same data set, we apply logistic regression to study the effect that utilization has on the probability of readmission for a specific patient

- We use readmission within 72 hours as our dependent variable

- Hypothesized logistic regression model

logit(READMISSION72) = AGE + BLACK + ASIAN + HISPANIC + FULL (or BEDS)+ ELECTIVE + TRANSPLANT + TRAUMA + … + NEURO + MALE + ε

Results

- Model with Full: Controlling for race, age, gender, and the type of surgery, being discharged from a full post-operative unit increases the odds of readmission by a factor of 2.341
- Model with Beds: Controlling for race, age, gender, and the type of surgery, each bed in use at the time of discharge increases the odds of readmission by a factor of 1.008

Utilization-Readmission Relationship

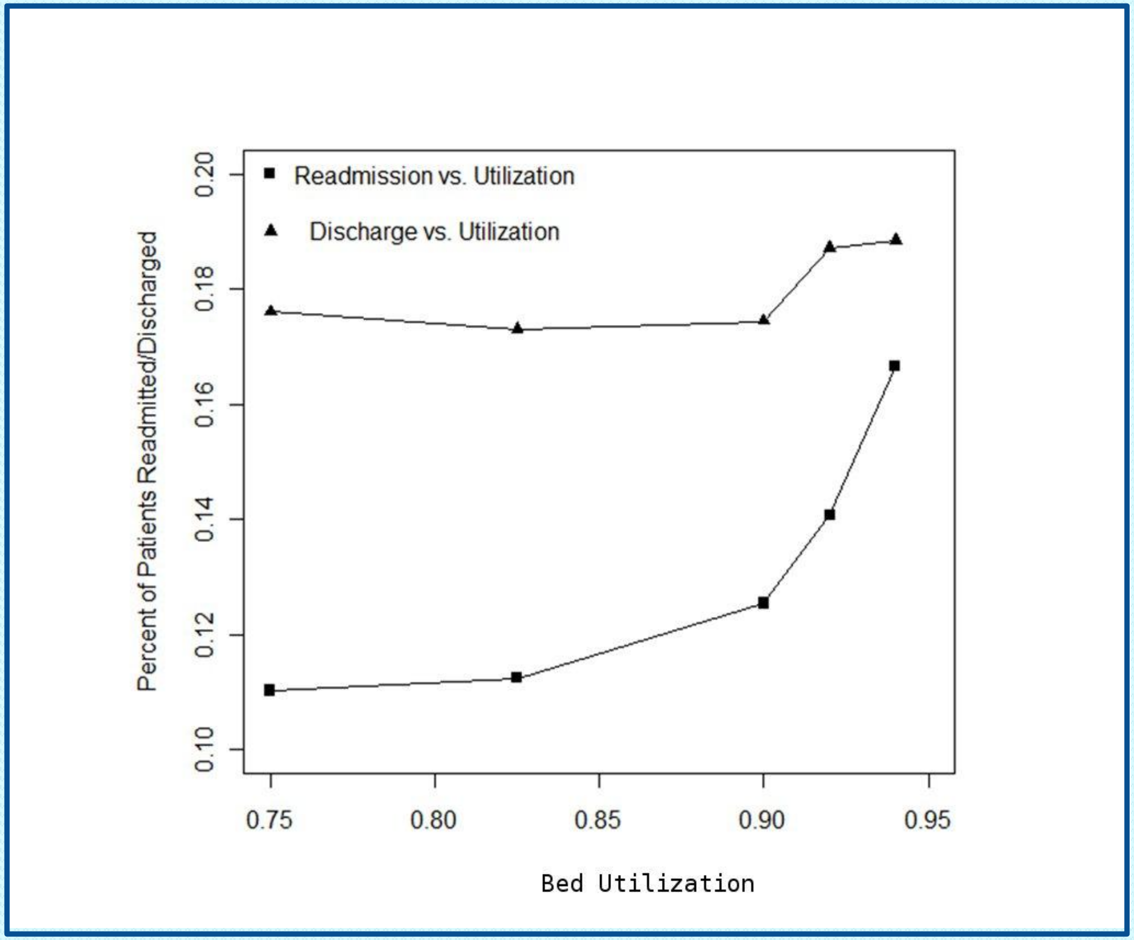

The discharge rate and readmission rate both increase as utilization increases

Survival Analysis

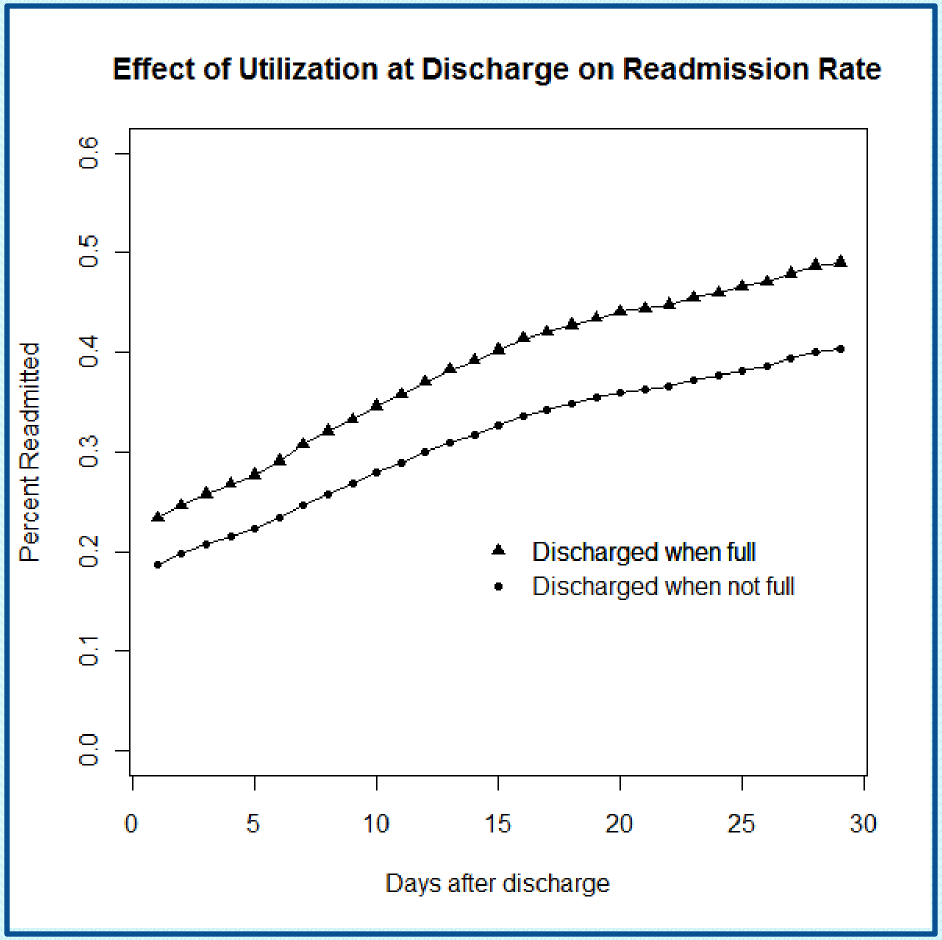

Over the course of a month, patients discharged from a full hospital are much more likely to be readmitted

Discussion

- The discharge rate rises when utilization is high
- This corresponds to an increase in the readmission rate
- We conclude that some patients are discharged too soon when utilization is high
- Surgeons have an incentive to clear space for their surgeries
- Mitigation strategy: Use a checklist before discharging a patient—force the surgeon to think about whether the discharge is for the right reason

Conclusions

- Research opportunities in vehicle routing, disaster relief, and healthcare analytics are plentiful
- The HTSP work presented here will appear in Optimization Letters
- The healthcare analytics work presented here has appeared in Health Care Management Science (2011, 2012)
- Thank you!