Intelligent Modeling Essential to Get Good Results

Katta G. Murty

ISE, KFUPM, Dhahran-31261, Saudi Arabia

IOE, University of Michigan, Ann Arbor, MI-48109-2117, USA

March 2011.

Abstract: This article is based on a seminar given at KFUPM recently, and at the Edelman award talks organized by CPMS of INFORMS in 2004. It also includes a **call for submissions** of real-world case study project problems in OR for an edited book (see Section 8 for details).

1 My Introduction to Optimization

My first encounter with optimization was in 1960 in a talk given by a visiting US Professor at the Indian Statistical Institute, Kolkata, where I was a graduate student at that time. The American Professor talked about a newly evolving subject called **OR (Operations Research)** unheard of in India at that time. He talked about “algorithms” (a new term for me at that time), and a problem called the **TSP (traveling salesman problem)**, for which he said no one has been able to develop a reasonable algorithm so far. That intrigued me so much that I started thinking about the TSP, and in about a year, while on a visit supported by a Fulbright grant to Case Institute of Technology (now called the Case Western Reserve University), I developed the Branch and Bound method for it [1].

This work has helped to shift my focus from Statistics to Optimization, and in 1965 I joined the University of California, Berkeley (UCB) as a graduate student for my Ph. D. When I met my advisor there, David Gale, for the first time, he

asked me “Mr. Murty, what do you want to work on?”. I told him “Prof. Gale, I want to work on optimization”. Thinking about this today, I am so happy that I replied in this way at that time, as optimization has opened many opportunities for me. I also want to encourage young scholars planning their careers to seriously consider optimization.

I believe that the development of mathematics, and in fact of all sciences, has its roots in the human desire to optimize. So, before getting into the main topic, I will briefly survey the history of the development of optimization from ancient times.

100,000 Years Ago ...

In those days, people used to live in caves. They need water from the river, so they faced the problem of going from their cave to the river for water. Seeing the nearest point on the river bank to their cave, they discovered the fundamental result that the shortest route from the center of their cave’s entrance to that point on the river bank is the straight line joining them. This is perhaps the earliest mathematical result discovered by mankind; solution to an unconstrained minimization problem (as there are no constraints on the route under consideration).

Several Thousand Years Later...

Things changed slowly in those days. Some thousands of years later, a family of tigers moved in and occupied a region in the middle of their shortest route from their cave to the river. Now they faced the new problem of finding a shortest route from their cave to the river avoiding the zone occupied by the tigers, an instance of a constrained optimization problem.

Now...

We have seen above the two types of optimization problems encountered in real-

world decision-making. Now, OR is the branch of Science dealing with optimum decisions for situations that we face in real-world operations. The main strategy for solving these problems involves the following steps: **1.** Construct a mathematical model for the decision problem (which involves making a list of all the relevant decision variables that play a role in the problem, and whose values need to be determined optimally; identifying the bounds and constraints on them from the way the system operates; and a list of all the objective functions that need to be optimized), **2.** Solving the model using an efficient algorithm to find the optimum solution(s), **3.** Making necessary changes, and implement the final solution obtained.

By now OR Theory has developed efficient algorithms for solving several single objective decision-making models. But in many real-world applications, practitioners find that none of these well-solved models fit their application exactly; because real-world problems are usually multiobjective, and lack the nice structure of models discussed in theory. So, there is a big gap between theory and practice, and we need a bridge to cross this gap. In order to get good results in applications, it is essential to model the problems intelligently using heuristic modifications, approximations, substitute objective functions, relaxations, and hierarchical decomposition, these provide the bridge to cover the gap between theory and practice. I will illustrate this point using a project carried out at **HIT (Hong Kong International Terminals)** in Hong Kong Port, and later at PSA (Port of Singapore Authority), two of the largest container terminals in the world.

2 Brief Introduction to Container Shipping

Maritime shipping is broadly divided into two categories: one is shipping of bulk commodities like crude oil, iron ore, coal etc. in specialized vessels, and the other called **container shipping** handling all other manufactured goods like textiles, etc.

In container shipping, customers pack all the goods that they want to send to overseas destinations in steel boxes called “containers” of which there are 3 different sizes in use today: 20-foot long containers called **TEUs (twenty-foot equivalent units)** common in old times, 40-foot long containers called FEU (forty-foot equivalent units, counted as 2 TEUs) which are becoming common today, and “refers” (refrigerated containers that must be held at cold temperatures, commonly 45-foot long). The customer then takes these containers (called **outbound** or **ECs (export containers)**) to the container terminal in his own truck (referred to as **ET (external truck)**) which enter the terminal through the **TG (terminal gate)**. The gate operator tells the driver where to take it in the **SY (storage yard)**. At the SY, a **YC (yard crane)** removes the container from the truck, and stores it, where it will remain until the vessel into which it is to be loaded arrives and docks at a berth in the terminal. At that time the YC retrieves the container from its storage position and puts it on one of the terminal’s trucks (these are owned by the terminal, hence called **ITs (internal trucks)** and used by them to shuttle containers between the quay and the SY) which takes it to the vessel on the berth. This IT with that container joins a queue under a **QC (quay or shore crane)** loading/unloading containers into/out of the vessel. When this truck moves to the front of the queue, the QC removes the container from it and loads it into the vessel. The truck that carried it then returns to the SY to bring another container to be loaded.

The handling of inbound containers or **ICs (import containers)** that are sent to this port by overseas customers follows a reverse sequence to that of ECs, they arrive on vessels, unloaded on berth by QCs, and loaded onto ITs which take them to be stored in the SY until its owner sends his ET to the terminal to pick it up, then it is retrieved from storage by a YC, transferred to the ET which leaves the terminal with it through the TG.

SY is the section of the terminal used for the temporary storage of containers between land and sea transportation. Depending on the terminal’s mode of op-

eration, containers may be stored either on the trailer on which they arrived, or stacked one on top of the other. The SY is divided into rectangular regions called **storage blocks**. At HIT, the width line of each block is divided into 7 lanes, 6 of them are storage lanes and the 7th is the truck passing lane for container trucks (ET or IT) to bring containers for storage in the block, or to take away containers stored in this block. Each storage lane is divided into 26 storage spaces, each for storing one stack of containers one on top of the other up-to 6 containers high. Each block is allotted 2 or more YCs for storing arriving containers and for retrieving containers from storage.

3 Key Service Quality Metrics

There are only about 40 major shipping lines in the world, and they all take their decision on which container terminal to patronize very seriously. Once a shipping line decides to stop using a terminal's services, there is very little that the terminal can do to win them back. Regional competition among terminal operators is fierce. Every day terminal managers rack their brains to improve service level to surpass their competitors, as efficiency is a matter of business survival.

Each vessel has a lot of employees working on it to keep it sailing smoothly. When a vessel docks on its berth, the terminal begins **processing it**, i.e., unload import containers off it and load export containers onto it; once this processing is completed the vessel leaves the port on its next voyage. The amount of time taken for this processing is called the **vessel turnaround time**. While the vessel is docked on the berth, its employees have no work, but the shipping line has to pay their wages. That's why shipping lines are very anxious to keep these vessel processing times at terminals as low as possible. So the **average vessel turnaround time** at a terminal is one of the most important service quality metrics used in the industry to rate terminals, at top rated terminals this varies between 8 to 12 hours.

The average vessel turnaround time is directly influenced by another metric known as the **GCR (gross crane rate)** or **quay crane rate**, which is the average lifts achieved at the terminal per QC working hour; where a “lift” refers to either the unloading of an import container from, or the loading of an export container onto the vessel. Clearly the vessel turnaround time at a terminal is inversely proportional to its GCR, that’s why the GCR is the most commonly used productivity measure to rate container terminals. The higher the GCR the better the service quality, so it is a profit measure to be maximized. The gold standard for GCR is 40, good terminals have GCRs in the high 30s.

4 Business Environment in Hong Kong and in HIT in Mid-1990s

China’s manufacturing boom was creating a remarkable demand for terminal services in Southern China at that time. The surge of exports from the region offered terminals there a golden opportunity; whichever terminal could offer the capacity and service quality would capitalize on this expanding market. However, with land scarce in Hong Kong, HIT could not simply expand its terminals to meet the growing demand. Also it faced growing competition from several new terminals opening along China’s southern coast. With favorable labor costs and governmental subsidies in China, these ports were poised to offer HIT’s customers much cheaper prices than HIT could. So HIT realized that it had to provide premium service quality and establish itself as the industry’s benchmark in order to survive.

However, the GCR at HIT was hovering in the upper 20s at that time; consequently their average vessel turnaround time was high. The situation entered a crisis mode in 1995 when they lost a customer, and faced a dire situation of losing market share in a growing market.

The **terminal road network** refers to a directed network representation of the road system within the terminal; with blocks, berths, TG and road intersections

as nodes, and road segments joining pairs of nodes as arcs. This terminal road network is virtually a closed system with the TG as the only access point for ETs; and the only vehicles operating on it are container trucks and the occasional YC moving from one block to another.

Traffic flow in a container terminal is akin to the circulation of blood in the human body: life depends on it. With congestion on the terminal road network, ITs get stuck in traffic, consequently QCs have to wait for them, pulling the GCR down. So, the GM of HIT realized that congestion was the Gordian Knot that needs to be untied before any other improvements can benefit HIT.

The operational practice at that time at HIT (and in all other busy container terminal around the world) was to allocate 6 ITs to serve each working QC. Each IT would shuttle back and forth between the SY and its allotted QC. Realizing that QCs were losing productivity due to waiting for ITs, HIT management decided to increase the allotment of ITs/QC from 6 to 8, hoping that it would fix the problem. Six months after this change, they found that the new policy made the situation even worse than before. It was counterproductive because increasing the number of trucks operating on the finite terminal road network system naturally increases traffic congestion. HIT needed to reduce the number of working ITs, but increase the throughput of QCs through other improvements.

So, the GM of HIT decided to take the bull by the horns and reengineer HIT's processes using intensive computerization, and improving its decision making in daily operations by engaging an outside consultant. I was visiting HKUST at that time. When the GM met me, he told me "Professor Murty, to describe a problem as hard, Americans use the expression "it is like rocket science". I think that congestion on our terminal road network is harder than rocket science. If you can, help us develop a new **DSS (decision support system)** for daily operations that would help reduce it. This will form the core of our program to enhance productivity and capacity".

千里之行, 始于足下

A long journey begins with the first step

5 My Involvement With the Problem

It became clear to me that reducing congestion requires a 2-pronged approach: reduce the number of trucks operating on the terminal road network, and most importantly routing them optimally. I did not know anything about solving road congestion problems at that time, so I hesitated that by accepting to work on it, I may be setting my foot into a very dangerous territory. But I was in Hong Kong, and just heard the Chinese proverb given above, that mustered my courage to take the plunge.

The first thing that I did was to scan the literature for any publications in prestigious journals on alleviating traffic congestion in terminal road networks. Even though it is a common problem, I did not find any research publications on this problem at that time. So, we were on our own.

Reducing the Number of Vehicles Operating on the Terminal Road Network

The operational procedure at that time was having 8 ITs allotted to serving each QC, these line up in a queue to serve only that QC. Thinking about this, I was reminded about an analogous problem that we discussed in a Queueing theory course at UCB; in which we analyzed a system with many servers working simultaneously. We compared two policies: Policy 1 would maintain a separate queue in front of each server, and arriving customers can join any queue that they like at the time of arrival; Policy 2 maintains a single queue which every customer joins at the time of their arrival, and the customer at the front of the queue will be serviced by any server that becomes free next. I remembered learning that Policy 2 gives

better performance results than Policy 1.

Applying this to the container terminal situation, Policy 1 corresponds to the existing practice of having a separate batch of 8 ITs to serve each QC. Policy 2 corresponds to a new operational procedure that I called the **pooling system** in which all the ITs working in the terminal form a pool that collectively serves all the QCs. This pooling system needs a new **central dispatching unit** which will be communicating with each IT driver, and continuously monitoring the number of QCs lined up and waiting under each QC. It directs each IT returning from the SY to the quay to join under the QC with the smallest number of ITs waiting under it at that time. The queueing theory result suggests that this pooling system will have better performance than the existing system.

So, I asked the HIT Management if this pooling system can be implemented. They estimated that the required communication and monitoring equipment alone would require an initial investment of 11 million US\$, a substantial investment for this industry. But being in an urge to modernize, they decided to go ahead with it.

In a conference a couple of years later, the manager at the Port of Singapore Authority told me that they were also adopting this pooling system for ITs at their container terminals. Now every busy container terminal in the world has adopted this pooling system. This pooling system helped reduce the number of ITs operating in the terminal while simultaneously improving their service level to the QCs (see Section 6 for details).

Routing Container Trucks

While the pooling system helped improve performance, I realized that in order to achieve a significant dent in congestion, we need to route container trucks optimally. At that time HIT had 10 berths; over 200 ITs, each of them continuously on the move between the SY and the quay; 80,000 TEUs of constantly accessed storage space; and 10,000 truck passages through the TG daily. I realized that it is really impractical to develop a route that each IT driver is required to follow for each

trip between the SY and the quay. Also all the IT drivers know the terminal roads intimately, and know instantly what route to take to get to any destination point as fast as possible. So ideally we must make sure that route optimality to minimize congestion follows as a consequence of other decisions made in daily operations. I realized that the policy of allocating storage spaces in the SY to arriving containers is the key to achieve optimum routing, as that determines the routes that ITs take.

HIT organizes its work using a 4-hour interval (half-a-shift) as the **planning period**. In each period they work on planning how to make decisions during the next period. In the literature I have seen several publications on this storage space allocation problem for arriving containers. All of them modeled the problem using binary variables of the form: $x_{ijkl} = 1$ if the i -th container from the j -th QC is stored in the k -th stack of block ℓ , 0 otherwise.

The model is large with many binary variables, and took too long to solve. At that time, I attended a container terminal decision-making session in a conference. One of the speakers said that he was able to reduce the solution time to an hour. In daily operations in a terminal, to retrieve a container at the bottom of a stack of 5 stored containers, the crane operator has to first move the top 4 containers in this stack to other stacks, this is a common operation called **reshuffling**. So a storage space which is unoccupied at an instant, may not be vacant even a few minutes later due to this reshuffling. The input data for this binary variable model is the set of vacant storage positions in the SY, it keeps changing constantly in the container terminal, so I realized that this binary variable model for storage space allocation to arriving containers is not only impractical, but also totally inappropriate.

So, we decided to develop a new model for this problem on our own. We then realized that the best strategy to minimize road congestion, is to spread the container truck traffic on all the road segments evenly, i.e., to equalize the container truck flows on all the arcs of the terminal road network as much as possible. This lead us to the 2nd model for the problem which is a multicommodity flow model. The data needed for the model is the expected number of containers flowing into and

out of each node in the terminal node network. For instance, during the planning period:

- at each QC position on the quay, this data consists of how many import containers are expected to be unloaded and sent to the SY from there, and how many export containers are expected to be sent here from each block of the SY
- at each block, how many export containers are expected to be dispatched to each QC position, and how many import containers are expected to be retrieved for leaving the terminal through the TG
- at the TG how many export containers are expected to arrive for entry into the terminal
- etc.

The decision variables in this multicommodity flow model are: f_{ij} = the total number of container trucks flowing on arc (i, j) in the terminal road network during the planning period. The objective function to minimize is $\theta - \mu$, where $\theta = \text{maximum}\{f_{ij} : \text{over all arcs } (i, j)\}$, and $\mu = \text{minimum}\{f_{i,j} : \text{over all arcs } (i, j)\}$.

For each 4-hour planning period, the multicommodity flow model is a large-scale LP, but we could solve it in a few minutes of cpu-time using the best available LP software system at that time, and the output led to the routes for container trucks to minimize congestion.

We ran into very serious difficulties trying to implement the optimum solution from this model. First, truck drivers resented being told what routes to take. Their union said “We all know the terminal road network well, and can find the best route to get to any destination point by ourselves based on current traffic conditions”. Second, this model tries to optimize the handling of the total expected workload in the planning period of 4 hours, in that process its solution handles this total workload evenly over the time interval of this period. So, inherently it is assuming

చీకటిని బిడుతూ కూర్చోకు, ఒక చీరుదీపం వెలిగించు

Don't just sit there cursing the darkness, light a small candle

that all this work is available at the terminal at the beginning of the period. But in reality, in the arriving streams of ECs and ICs each container has to be handled at the instant of its arrival, and the flow rates of these streams is quite uneven over time because of the stochastic nature of vessel arrivals. So this model also, is not really appropriate for our problem.

When we realized this, I became totally frustrated, and at a loss to know what to do next. Then the Telugu proverb given above came to my mind:

I made up my mind to find that candle to shed light on an approach for solving the congestion reduction problem. I realized that I needed to have more intimate knowledge about the various operations going on inside the terminal, and the consequences of various decisions. So, I spent a good 3 weeks in the terminal just observing the operations. During this time I noticed that the : occupancy or **fill-ratio in a block** = (number of containers in storage in the block)/(number of storage positions in it) varied significantly over blocks. At HIT the number of storage positions in each block = 600. All the blocks with higher fill-ratios tended to have much more truck traffic around them compared to those with low fill-ratios.

This gave me the idea that equalizing the fill-ratios among blocks ensures equal distribution of traffic on terminal roads, minimizing congestion. Fill-ratios in blocks of course vary with time as containers are added to or retrieved from them. So, we needed to select a specific time-point in each planning period at which we try to equalize the fill-ratios in all the blocks, and we decided that this point of time will be the end of each planning period. This led to our next model for the problem, which is again an LP model but much smaller. The decision variables and data elements in this model are the following (all the data elements are either known or can be estimated with good precision for the planning period):

$x_i =$	Container quota number for block i in this period = number of arriving new containers in this period to be dispatched to block i for storage, a decision variable.
$a_i =$	Number of stored containers that will remain in block i at the end of this period if no new containers are sent there for storage during this period, a data element.
$N =$	Number of new containers expected to arrive for storage in this period, a data element.
$B, A =$	Total number of blocks in the storage yard, number of storage positions in each block, data.
$w_i(t) =$	Number of container trucks waiting in block i to be handled by the yard cranes there at time point t , data.
$x_i^R(t) =$	Remaining container quota number for block i at time point t in the period = $x_i -$ (number of new containers sent to block i for storage up to time point t), data.

5.1 Fill-Ratio Equalization Policy

This policy determines the decision variables x_i for this period to make sure that the fill-ratios in all the blocks are as nearly equal as possible at one time point in the period, the end of the period. The fill-ratio in the whole yard at the end of this period will be $F = (N + \sum_i a_i)/(A \times B)$. If the fill-ratios in all the blocks at the end of this period are all equal, they will all be equal to F . So, this policy determines x_i to guarantee that the fill-ratio in each block will be as close to F as possible by the end of this period. The linear programming (LP) formulation of this problem is

$$\begin{aligned}
 & \text{Minimize} && \sum_i (u_i^+ + u_i^-) \\
 & \text{subject to} && a_i + x_i - (u_i^+ - u_i^-) = A \times F \quad \text{for all } i
 \end{aligned}$$

$$\sum_i x_i = N$$

$$x_i, u_i^+, u_i^- \geq 0 \quad \text{for all } i$$

The number of variables in this model is B = number of blocks in the SY, which is typically of the order of 100 or so for large, busy terminals. Also, because of the special structure of this LP, it can be shown that its optimum solution is given by a simple combinatorial scheme described below:

1. Rearrange blocks so that a_i increases with i .
2. Determine x_i in increasing order of i using:

$$x_1 = \text{Minimum}\{N, A \times F - a_1\}, \quad \text{and for } i \geq 2$$

$$x_i = \text{Minimum}\{\text{Maximum}\{0, A \times F - a_i\}, N - \sum_{r=1}^{i-1} x_r\}.$$

Numerical Example

Suppose there are $B = 9$ blocks in the terminal, each with $A = 600$ storage spaces. Suppose $N = 1040$ new containers are expected to arrive for storage in the planning period. Data on a_i is given below, with a_i already increasing with i . The fill-ratio over the whole yard at the end of this period will be $F = (\sum a_i + 1040)/(9 \times 600) = (2570 + 1040)/(5400) = \approx 0.67$, and $A \times F \approx 400$.

The combinatorial scheme described above determines x_i in the order $i = 1, 2, \dots$, to bring $a_i + x_i$ to 400 until all 1040 new containers are allotted to a block for storage. The results are shown below.

i	a_i	x_i
1	100	300
2	120	280
3	150	250
4	300	100
5	325	75
6	350	35
7	375	0
8	400	0
9	450	0
Total		1040

Note that this policy only determines the container quota number for each block, not the identity of which containers will be stored in each block. That will be determined by the dispatching policy discussed below.

5.2 Arriving Container Dispatching Policy

Irrespective of how the container quota numbers x_i are determined, Congestion will be created at a block if we send a consecutive sequence of arriving container trucks to it in a short time interval. To avoid this possibility it is necessary to allow the yard crane in that block enough time to handle a tractor sent there before sending the next one.

Hence this policy dispatches each arriving truck (at the terminal gate, and each berth) at time point t in the period to a block i satisfying: $w_i(t) = \text{Min}\{w_j(t) : j \text{ satisfying } x_j^R(t) > 0\}$, i.e., a block with remaining positive quota that has the smallest number of trucks waiting in it.

The Substitute Objective Function Technique

Here instead of minimizing congestion measured directly by either the maximum

flow amount on an arc in the terminal road network, or $\theta - \mu$ defined earlier, we were able to control it by minimizing the sum of absolute deviations in the fill-ratios of the various blocks from their average. The latter objective function has a strong influence on the former, which is why the substitution produced excellent results. This is the idea behind the substitute objective functions approach developed in [Murty and Djang, 1999] for other routing problems.. By finding suitable objective functions to substitute, this approach can be used to handle complex large scale optimization problems in many different areas. The main requirements are that the substitute objective function have a strong influence on the original, and optimizing it should be simpler than the original problem.

Many other problems in industry and business deal with processing arriving streams of items, orders, etc. When each arrival has to be processed at the time of arrival, traditional approaches based on batch processing (i.e., planning for a batch of expected arrivals in a period) do not usually work well. In such cases, developing dynamic real-time strategies to handle each arrival based on the prevailing conditions at that time, as we developed here for arriving containers, may offer a good approach.

6 Performance Benefits

The policies discussed above were implemented, and they resulted in significant improvements in performance. With optimized storage space allocation, dispatching and routing, congestion has decreased as evidenced by a 16% decrease in IT cycle time. Actually we also worked on some other decision-making problems in daily operations in developing the DSS, for details see [3]. Implementing this DSS for decision-making in daily operations resulted in a 30% improvement in the GCR at HIT, and a reduction in the average vessel turnaround time from over 13 hours to 9 hours. The average number of ITs deployed/QC decreased from 8 to 4. Since this project was initiated, HIT's annual throughput has gone from about 4 million

TEUs in 1995 to about 6 million in 2002 , which they continued to handle using the same labor force resulting in significant savings.

HIT received an Edelman Finalist Award of INFORMS for this work.

7 The Lessons that I Have Learnt

As you know there is this eternal conflict between theoreticians and practioners:

Theoreticians say that practitioners do **practically nothing**

Practitioners say that theoreticians do **nothing practical**.

I consider myself belonging to both camps, and I learnt two things from my experience: one is that constructing an intelligent model for problems encountered in applications is often quite challenging. The second is that implementing the solution from the model requires listening very carefully to the viewpoints of all the stakeholders and requires being very tactful as explained in Chapter 3 of the recent book [Murty, 2010].

Summary: OR academic curricula in general emphasize the teaching of techniques and theory, and do not pay much attention to cultivating good modeling skills in their students. With significant fractions of OR graduates opting for applications-oriented careers these days, I think it is important that modeling find a prominent place in OR curricula.

8 A Call for Case Studies With Data, Suitable for Project Problems

The Railroad Applications Society (RAS) of INFORMS has set up an excellent railroad application problem (see: <http://www.informs-ras.org/problem.htm>) that

can be used as a project problem in a modeling course. By working on this problem either individually or in a group, the student gets a clear idea of the types of problems that railroad OR practitioners face in their work. I am currently editing a book of such cases from a variety of application areas, that instructors can use as project problems in teaching modeling courses in OR programs, or students can use in self-study. The aim is to provide realistic cases that will help the students in developing the modeling skill, and in deciding the approach and the algorithm to attack problems that they will face in their work.

If you are able to contribute one or more cases to this book, please contact me (murty@umich.edu , murtykatta@gmail.com) so that we can discuss the details. An example of the types of submissions that I am looking for is our paper [Nag, Murty, 2010] on the railroad problem set up by RAS, citing the website where all the data for the problem can be accessed. I hope to get enough cases to publish the book in the next three years.

9 References

1. Murty K. G., C. Karel, and J.D.C. Little, 1962, "Original Article on Branch and Bound Method for the Traveling Salesman Problem," can be seen under "Selected Publications" on my webpage at: <http://www-personal.umich.edu/~murty/>
2. Murty, K. G., and P. A. Djang, 1999, The US Army National Guard's mobile simulators location and routing problem, *Operations Research*, 47(2)175-182.
3. Murty, K. G., Y.-W. Wan, J. Liu, M. M. Tseng, E. Leung, K.-K. Lai, H. W. C. Chiu, 2005, Hongkong International Terminals gains elastic capacity using a data-intensive Decision-Support System, *Interfaces*, 35(1)61-75.
4. Murty, K. G., 2010, "Optimization for decision making: Linear and quadratic models", Springer.
5. B. Nag, and K. G. Murty, 2010, "Locomotive Fueling Problem (LFP) in railroad operations", IOE Dept., University of Michigan, Ann Arbor, MI-48109-

2117; can be accessed under “Selected Publications” on my webpage: <http://www-personal.umich.edu/~murty/>