OR 7310: Logistics, Warehousing, and Scheduling
Course Project:
A Telecommunications Network Design Problem

Fall 2016

Wednesday, December 14, 2016

Announcements

- Teams must be of 3-4 people. You may choose your own groups. Send me an e-mail if you need help with finding a team by Monday November 7th .
- Each group need only turn in one electronic copy of anything asked for. Please ensure that your write-up is coherent and complete (and correct!). Be sure to record the names of all team members.
- You will give a short (less than 10 minutes) presentation 1:30-4:30pm on Wednesday, December 14, 2016.

Goals

This assignment gives you an opportunity to tackle a real-world problem using the techniques you have learned in the course. It will also give you a chance to exercise your creativity and insightfulness.

Table of Contents

1. A Telecommunications Network Design Problem

1.1 Overview

A telecommunications company is building networks that require the installation of expensive equipment (the major components cost at least $1 million). There are many places where the equipment might be installed, and by selectively choosing where to put the most-expensive pieces of equipment, the company might be able to reduce its expenditures by millions of dollars. Your knack for solving network problems has earned you widespread fame, and the company hires you to find out their least costly configuration of equipment for two different networks: the Caldata network and the Barry network.

It seems simple enough, and you easily formulate a model and code it up in AMPL to solve. (We will give these to you!) You hand this formulation to your engineer sidekick and call it a day.

Unfortunately, two days later, you find the following troubling email in your inbox:

```
I tried to run the Barry network the night you handed it off to me. It was slow, so I
just let it run overnight. The next day, I found that our systems had completely run out
of memory. I'm trying to run the Caldata network now, since it's smaller. But it's
already been going for 6 hours, and it doesn't seem promising...
```

Indeed, you will face the exact same situation if you try to run this model on your computer. With a very powerful computer and a lot of time, CPLEX *may* be able to solve the Caldata network fully, but that's very unlikely. And we can promise that you will certainly not be able to solve the Barry network to optimality with your initial formulation. With less than two weeks before the company needs an answer, what do you do?

1.2 The Problem at Hand: ADM Selection and Routing

The company is building two large SONET telecommunications networks in the western United States. Each network has a small number of hubs (you can think of these as cities), labeled by letters (because the actual locations are confidential). Hubs are connected by arcs, and in the network architecture the company is using, arcs are created in sets of parallel rings.

For example, in Figure 1 (the Caldata network), there are 14 hubs (labeled A through L), and each hub has at least two rings going through it. The upper-left-hand rings, for example, connect hubs A, B, E, and C; as shown in the figure, there are four parallel rings that make up this cycle.

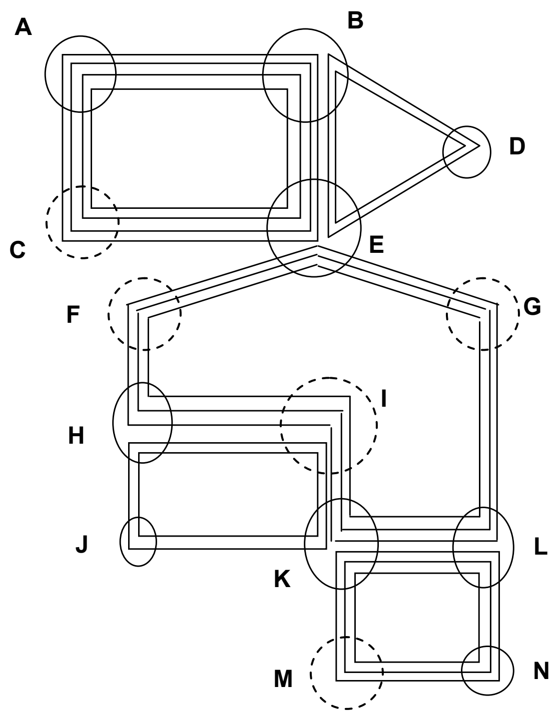

Figure 1. Plan for the company's Caldata network. Hubs with a solid circle (ABDEHJKLN) have a BBDX installed, and hubs with a dashed circle (CFGIM) do not.

As data flows from one hub to another in this network, it will often have to switch from one ring to another. For example, to go from A to D, the data will need to switch from one of the ABEC rings to one of the BED rings. Ring-switching is only possible at hubs that have a broadband digital cross-connect (BBDX) installed. In the figure above, all hubs with a solid circle (A, B, D, E, H, J, K, L, N) have a BBDX installed. Hubs with a dashed circle (C, F, G, I, M) do not have a BBDX installed. BBDX installation is based on equipment available from other networks (and each BBDX costs tens of millions of dollars), so the locations of BBDXes are fixed; they can't be changed.

Data can only enter or leave the network at hubs. Moreover, when data enters or leaves the network, it must do so on a specific ring. For example, data can enter the network at hub A on the first ring, or on the second ring, or the third ring, or the fourth ring. At hub B, there are 6 choices: data could enter on one of the four ABEC rings, or it could enter on one of the two BDE rings.

However, data can only enter or leave a ring at a hub if the ring has an add-drop multiplexer (ADM) installed at that hub. This restriction applies not only to data entering and leaving the network, but also to data switching from one ring to another (so in order for data to be able to switch rings at a hub, the hub needs to have a BBDX *and* the two rings both need to have an ADM; see Figure 2). The primary questions to be answered in this project are the locations of ADMs (at which hubs and on which rings should they be installed?).

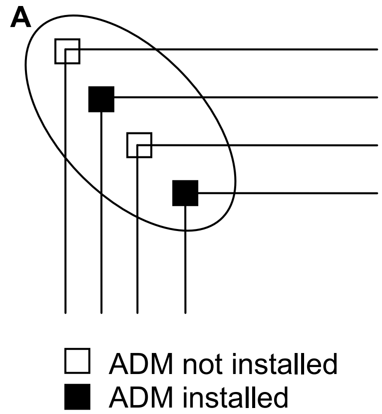

Figure 2. Hub A has four potential ADM locations, one on each ring. In this figure, ADMs are installed on the second and fourth rings (as shown by the dark squares). In this configuration, data can only enter and leave the network at Hub A at the second and fourth rings (because they have ADMs installed). Similarly, at Hub A, data can only transfer through the BBDX between the second and fourth rings. However, if another ADM was installed on the first or third ring, then data could enter, leave, and transfer on that ring as well.

In the Caldata network, each arc can carry a total of 24,000 units of data flow (sum of flow in the two directions). Because ADMs are attached to two arcs, one in each direction, they can handle a total of 48,000 units of data flow. BBDXes are large enough that they can handle as much flow as necessary (but note that because of the ring capacities, there would never be any reason to send more than 48,000 units of flow from one ADM to another at a BBDX). Figure 3 shows an example of these capacities.

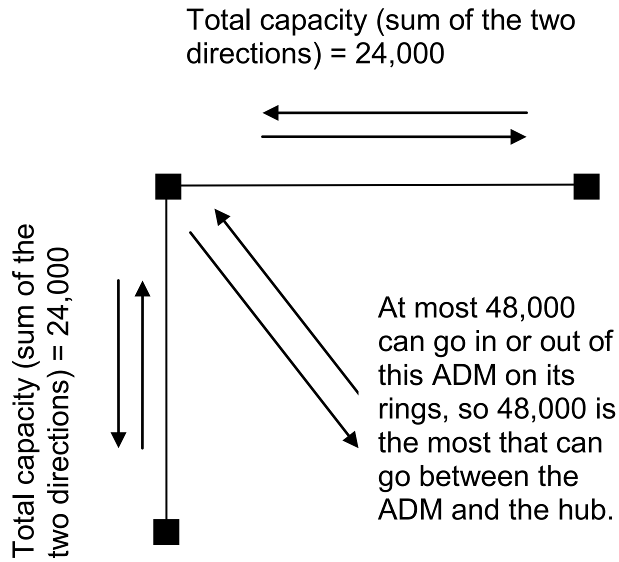

Figure 3. Example of the arc capacities and the *de facto* capacity of arcs between an ADM and its hub in the Caldata network. At most 24,000 units can flow in the left/right directions and at most 24,000 units can flow in the up/down direction, so the ADM will never need to handle more than a total of 48,000 units.

Each ADM costs $1,000,000 per year (this includes operating cost and prorated installation cost), so the company understandably does not want to install an ADM at every hub on every ring. However, they need to install enough ADMs that all of the network's demand can be met. Table 1 shows the hub-to-hub demand that the company expects its Caldata network.

Origin	Destination	Demand
A	B	38,000
A	D	7,000
B	K	13,000
B	D	19,000
B	C	1,000
B	A	42,000
B	N	31,000
B	E	7,000
B	F	6,000
B	J	9,000
B	H	2,000
K	I	1,000
K	M	2,000
K	A	11,000
K	N	32,000
K	J	14,000
K	H	6,000
K	L	6,000
D	C	1,000
D	A	12,000
D	N	1,000
D	J	2,000
M	J	1,000
C	A	1,000
A	N	27,000
A	E	5,000
A	J	1,000
N	J	20,000
N	L	11,000
E	G	3,000
F	J	2,000
J	H	4,000
J	L	2,000

Table 1. Expected Caldata network demand

In addition to considering the cost of ADMs, the company also needs to consider the cost of data flow. They estimate that because of maintenance, upkeep, etc., the annual cost per-unit of data using an ADM to enter or leave the network is $2, the annual cost per unit of data using a BBDX to change rings is a total of $8 ($2 per ADM used, plus $4 for the BBDX), and the annual cost per unit of data flow on each arc for the Caldata network is shown in Table 2. (Note that costs are symmetric, so, for example, the flow cost from K to M is the same as the flow cost from M to K.)

Hub 1	Hub 2	Cost Per Unit of Flow
K	M	1
M	N	1
N	L	1
L	K	1
J	H	1
H	I	1
I	K	1
K	J	1
L	G	2
G	E	3
E	F	3
F	H	2
B	D	1
D	E	1
E	B	1
B	A	1
A	C	1
C	E	1

Table 2. Flow costs for the Caldata network

In addition to the Caldata network (Figure 1 and Tables 1 and 2), the company is planning to build a second network, called the Barry network (see Figure 4 and Tables 3 and 4). The fibers that will be used in the Barry network are twice as large, so they can carry 48,000 units of flow (double that of Caldata).

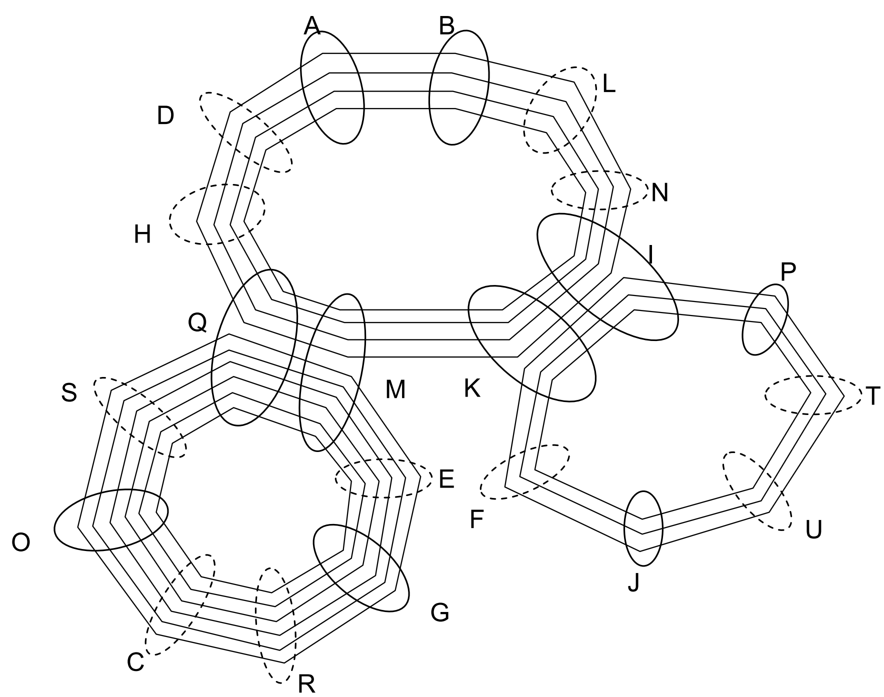

Figure 4. Plan for the company's Barry network. Hubs with solid oval have a BBDX installed, and hubs with a dashed oval do not.

Origin	Destination	Demand	Origin	Destination	Demand
A	B	66,000	F	G	6,000
A	C	1,000	F	I	19,000
A	D	34,000	F	J	61,000
A	E	2,000	F	K	10,000
A	F	2,000	F	P	2,000
A	G	23,000	F	S	1,000
A	H	10,000	G	I	13,000
A	I	16,000	G	J	23,000
A	J	13,000	G	K	4,000
A	K	6,000	G	M	17,000
A	L	9,000	G	O	118,000
A	M	23,000	G	P	6,000
A	N	17,000	G	R	177,000
A	O	57,000	G	S	15,000
A	P	9,000	H	M	1,000
A	Q	12,000	I	J	32,000
A	R	1,000	I	K	20,000
A	S	21,000	I	L	1,000
T	F	2,000	I	M	10,000
T	I	3,000	I	N	31,000
T	J	1,000	I	O	18,000
T	P	30,000	I	P	31,000
B	F	1,000	I	R	1,000
B	G	3,000	J	K	24,000
B	I	8,000	J	M	8,000
B	J	3,000	J	N	3,000
B	K	2,000	J	O	16,000
B	L	2,000	J	P	11,000
B	N	1,000	J	U	9,000
B	O	1,000	J	R	1,000
B	P	2,000	K	M	2,000
B	S	1,000	K	N	2,000
C	G	14,000	K	O	10,000
C	I	1,000	K	P	2,000
C	K	1,000	K	R	2,000
C	O	21,000	L	N	10,000
C	R	3,000	M	O	4,000
C	S	3,000	M	P	1,000
D	H	2,000	M	S	2,000
E	G	32,000	N	O	1,000
E	I	8,000	O	P	9,000
E	J	8,000	O	R	23,000
E	K	8,000	O	S	113,000
E	M	2,000	R	S	4,000
E	R	1,000			

Table 3. Expected Barry network demand

Hub 1	Hub 2	Cost Per Unit of Flow
A	B	1.08
B	L	1.08
L	N	1.04
N	I	1.07
I	K	1.05
K	M	1.07
M	Q	1.06
Q	H	1.01
H	D	1.03
D	A	1.03
G	R	1.01
R	C	1.05
C	O	1.04
O	S	1.04
S	Q	1.10
M	E	1.06
E	G	1.05
J	F	1.01
F	K	1.03
I	P	1.04
P	T	1.03
T	U	1.02
U	J	1.02

Table 4. Flow costs for the Barry network

2. An Integer Programming Formulation for the Problem

Your initial formulation of the problem is the following:

SETS

Node Sets

- Type 0 Nodes: These are the hubs/cities themselves, where demand starts and ends.
- Type 1 Nodes: These are potential ADM locations on rings, with one location per ring that goes through each hub.

Arc Sets

- Type 0 Arcs: These go from a ring to the hub/city. A commodity is allowed to flow on a Type 0 arc only if the hub is the commodity's destination, & only if an ADM is installed on the ring at this hub.
- Type 1 Arcs: These go from the hub/city to a ring. A commodity is allowed to flow on a Type 1 arc only if the hub is the commodity's origin, and only if an ADM is installed on the ring at this hub.
- Type 2 Arcs: These go from one hub to the next along a ring. Any commodity may flow on these arcs, whether or not an ADM is installed. An arc is created in each direction, so "Hub A to Hub B" and "Hub B to Hub A" are two different arcs.
- Type 3 Arcs: These go from one ring to another at a hub that has a BBDX. Any commodity may flow on these arcs, but only if an ADM is installed at the hub on both rings. ArcTable automatically includes only arcs at hubs with a BBDX.

VARIABLES

Decision Variables

- One binary variable per Type 1 node (each valid hub/ring combination), to show whether an ADM should be installed there or not
- One non-negative variable per commodity/arc pair to show how much of each commodity flows from the arc's start node to the arc's end node

MODEL

Minimize Costs

- ADM installation cost
- Flow cost on arcs and through ADMs and BBDXes

Constraints

- The total bi-directional flow on an arc (the sum of the two directions) can be no more than the bi-directional capacity.
- No commodities may flow on a Type 0 arc unless their destination is at that hub.
- No commodities may flow on a Type 1 arc unless their origin is at that hub.
- No commodities may flow on a Type 0 arc unless an ADM is installed on that ring at that hub.
- No commodities may flow on a Type 1 arc unless an ADM is installed on that ring at that hub.
- No commodities may flow on a Type 3 arc unless an ADM is installed at that hub on both rings.
- The total flow of each commodity into its destination hub (i.e., into a Type 0 node on Type 0 arcs) must equal the demand for that commodity.
- The total flow of each commodity out of its origin hub (i.e., out of a Type 0 node on Type 1 arcs) must equal the demand for that commodity.
- The total flow of each commodity into each Type 1 node must equal the total flow of that commodity out of that Type 1 node

3. AMPL Files (NOTE: you do not have to use AMPL, just document and submit what software and/or code you use)

The following files are provided to you:

`original.mod`	AMPL model file of the above basic integer formulation
`original.run`	AMPL run file that prints useful output information; modify data file line as needed
`caldata.dat`	AMPL data file that fully specifies the Caldata network
`barry.dat`	AMPL data file that fully specifies the Barry network
`caldata_toptest.run`	AMPL run file that runs the original model on the `caldata_toptest.dat` data file
`caldata_toptest.dat`	AMPL data file for the top two cycles of the Caldata network; this is provided for testing purposes and will be able to fully solve to optimality on any machine

Run `caldata_toptest.run` to see what you get!

4. Hints to Get Started

Now you've been given everything you need to solve the problem. But still, you face the problem of not being able to solve such enormous problems on your personal computer, and you cannot return to the company empty-handed. What will you do? We leave this part to you to figure out – be creative!

Here are some questions (no need to answer in the writeup – these are just for thought) to help you get started:

- If you're unable to find an optimal integer solution, can you at least provide a feasible solution?
- Would a feasible solution serve as an upper bound or a lower bound?
- How might you find an upper bound to the optimal solution?
- Will you be able to report how far off you are from the optimal? We call this the "optimality gap."
- How might you improve this optimality gap? How might you improve your lower bound? Upper bound?
- If you add constraints to the problem, will you get a better lower bound or a better upper bound?
- How might you relax the problem? Will this provide a lower bound or an upper bound?

You should aim to get at least below a 5% optimality gap. In the beginning, your gaps will be enormous. Don't worry; it takes some patience and creativity to reduce that gap

5. Deliverables

The Writeup

Please submit one writeup per team that describes all the additional ideas you added to your model or new heuristics you used. This writeup should describe the journey you took to find your best feasible solution and your best upper bound.

You must submit via e-mail to Dr. Ergun an electronic submission before **1:00 PM, Wednesday, December 14, 2016**. Gather all AMPL model and data files you have created and scripts containing the AMPL commands you used to solve the problems. Create a README file that will let us know which files refer to which problems. Put all these files in a zip archive, and email the archive to `o.ergun@neu.edu`. For the subject line, type: `submit OR7310 project teamname`, where you are free to choose any team name as long as it isn't lame. In the body of the message, include the names of your team members.

The Presentation

You will be giving a short presentation (around 10 minutes) **1:30pm-4:30pm Wednesday, December 14, 2016**. The purpose of the presentation is to show off the creative ideas you used.
We recommend limiting yourself to 3 slides in order to keep to your time limit. If using slides, be sure to email your slides to `o.ergun@neu.edu` by **1:00 PM, Wednesday, December 14, 2016**.

6. Grading

You will be graded on the following, in order of decreasing importance:

1. Creativity and thoughtfulness of gap-improving methods.
2. Correctness of methods: are your techniques implemented correctly?
3. Quality of final configuration: tightest gap.
4. Concise, clean, and polished final writeup.
5. Presentation.

Good luck and have fun!