



Network Capacity



by
Aaron D. Jaggard
James M. Kupetz

Network Capacity

Synopsis

The module develops the skills and perspectives used in modeling real-world situations as networks. Using the demand on networks familiar to students as a motivating point (text messaging, cell phone, Internet), the module attempts to model some simple networks (school hallways) and build upon that to more complicated networks. It concentrates more on allowing students to experiment with and think about networks and network capacity, rather than describing the algorithms used to determine optimal network settings/behavior. As the simulations increase in complexity, the idea that computers are essential for solving the more complex problems is introduced.

Note to teachers: Teacher notes appear in red in the module, allowing faculty to pull these notes off the teacher version to create a student version of the module.

Suggested Uses

Grades 10-14 mathematics survey classes or introductory computer science/computer information systems classes.

Prerequisites

Ability to effectively use Internet search engines.

Comfort with basic probability and statistics concepts, including die roll probabilities, mean, and standard deviation.

Minimal calculations will be required.

Ability to read spreadsheets and copy, paste, and edit formulas.

Key Ideas

Network, capacity, traffic, overload, repeated simulation

Preview activity: Wikispaces brainstorming on networks and the limitations they may place on the students or other users. Open with the questions: What are things you can think of that use the term “network?” What are examples of networks? What do you think defines a network? What activities are made possible (or better) by networks? What are limitations imposed by networks? As the module progresses, students can add more examples of networks and modify their definition of network.

Lesson 1 – Creating Network Models

This video clip http://www.youtube.com/watch?v=JkMZH_VZ3WE of SMS text messaging in Amsterdam around New Year's Eve 2007 shows the sudden spike that occurred near midnight. (Note the timestamp, using European formatting, in upper left.) What are the implications of this sudden and intense spike? **We want students to think about the possibility that the network might be overloaded.** What happens if you build a network that can handle 5 times as many messages as are sent on a regular day? **New Year's Eve traffic was described as multiple orders of magnitude greater than usual.**

What is a network?

There are many different examples of networks, some of which we will examine in this unit in detail, some we will mention briefly, and others will come up in class discussions. As you progress through the unit, you should look at the world around you and see what other collections of objects or ideas can be represented as networks.

We will start with a commercial-flight network to help us illustrate the essential components of networks. All networks have a collection of **nodes** as well as **connections** between the nodes. We may visualize these as dots or circles (the nodes) with some of the dots connected by lines (the connections). In our example to illustrate these components, the nodes would correspond to airports, and, if two nodes have a connection, then there is a direct flight between the two corresponding airports. (If we simply care about whether or not there is a flight between two airports, we will never have more than one connection between the corresponding nodes. In some settings, though, we might allow more than one connection between a pair of nodes, indicating multiple flights, for example.)



Figure 1- A map of US Airways air routes from their Pittsburgh hub as it appeared in 2011 and its representation as a graph

<source: http://www.airlineroutemaps.com/USA/US_Airways_PIT.shtml>

Networks can also contain some additional information. An important concept in modeling problems as networks is that there may be a **capacity** associated with each connection. In our commercial-flight example, this would tell us how big the airplanes are that fly between two cities (perhaps by giving the total number of seats on flights between the cities each day). There is usually also a measure of the **traffic** that makes use of each connection. In our example, this could either be a number representing how many seats are used, or it could be a number that indicates the fraction of the seats that are used.

In some problems, like this example, we may want to separately consider the traffic moving in one direction (the people flying from New York to Denver) and the traffic moving in the other direction (the people flying from Denver to New York). In those cases, we may use *arrows* instead of lines for the connections (which are then called **directed connections**). Each arrow has its own capacity. In the flight example, where the same plane may be used to fly both directions, the capacities will often be the same, but that is not necessarily the case in all circumstances.

When we encounter a new problem that we want to model as a network, we need to figure out how to set up the model. In particular, we need to decide what the nodes and connections will be (and whether the connections will be directed). We also need to decide whether to include capacities and traffic; if we do include these, we need to decide what their values are.

As part of this process, we are deciding which aspects of the real-world problem to model and which to ignore. This choice often depends on what we are trying to do with the network model. In our network with airports and flights, we have set up the network in a way that might be used to model whether a traveler could get from one city to another. In this setting, if an airport has 30 gates, it does not make a difference whether the airport has two terminals with 15 gates each or three terminals with 10 gates each; the layout is something we can safely ignore. However, in other settings (perhaps modeling movement within a single airport), such a difference might significantly affect the solution to a problem. We thus need to carefully consider the problem in deciding which information to keep. We also need to remember these decisions so that we do not try to use one network model to solve a problem that relies on information that we ignored when constructing the network.

In the airport/flight network, if we only care whether some passenger can get from one airport to another by taking one or more flights, we could ignore the capacities and traffic, although we would want to give directions to the connections. (For example, there would be a directed connection from LaGuardia airport in New York City to San Francisco airport if there is at least one flight from LaGuardia to San Francisco.) If we care about possible delays arising from overbooked flights (when an airline sells more tickets for a flight than there are seats on the plane), then it would probably be important to include capacities and traffic levels.

Let's apply these ideas to a new example. Consider the following diagram, which shows the layout of classrooms (the shaded regions) in a school building. The arrows indicate doors through which students can enter and exit the classrooms; the white areas are hallways.

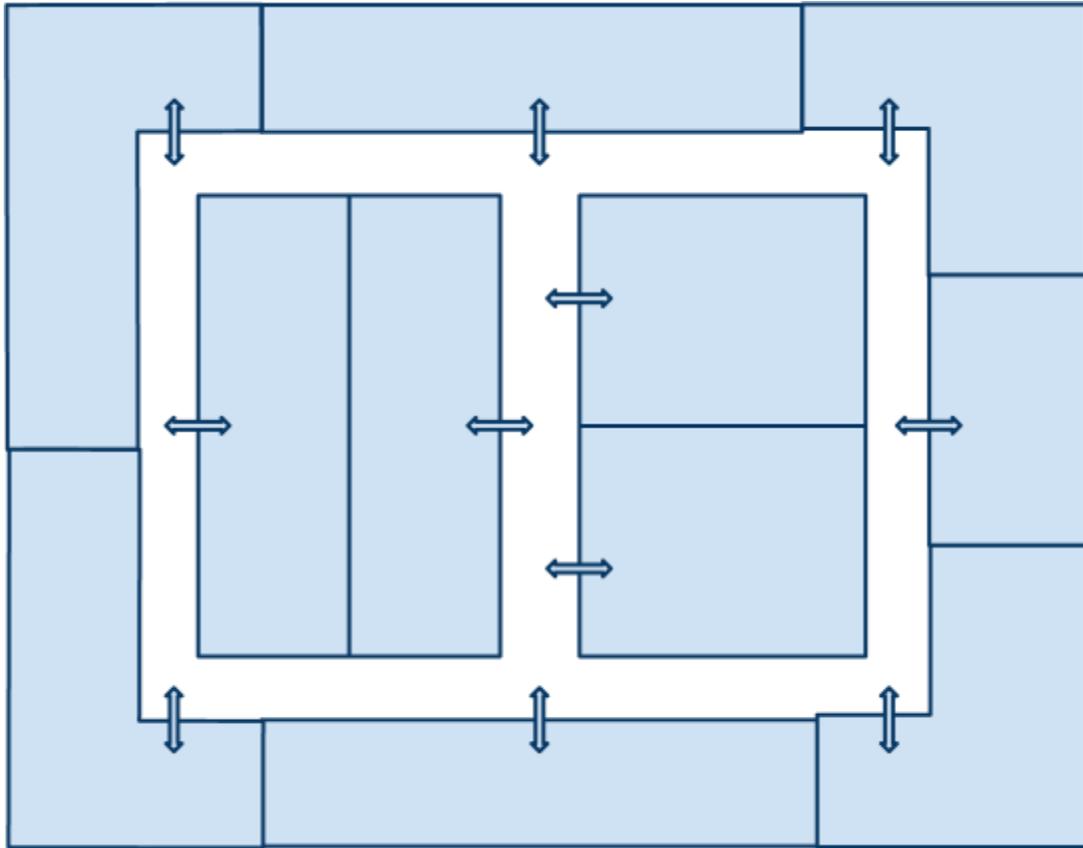


Figure 2: Classrooms (shaded regions) and hallways (un-shaded region); arrows indicate classroom doorways.

How might we model this as a network? First, we should decide which types of questions we might want to use the network to answer. When we find a problem that we might want to model using a network, we often ask the following questions:

1. What are the nodes?
2. What are the connections? Are they directed or undirected?
3. Are there measures of capacity and traffic? What do these values (in the model) correspond to in the real-world setting?
4. If there are measures of capacity and traffic, what indicates that the capacity has been exceeded? What are the consequences of exceeding a connection's capacity?

Let's look at these in the context of the example above with classrooms and hallways. One problem we might try to solve is whether students can move between classrooms without being late for their next classes. In that case, we would care about whether the hallways are full. Asking whether there are "too many" students in a hallway is the same type of question as "Does the traffic on this connection exceed the capacity of the connection?" This means that we will

probably want the hallways to be the connections (with capacities); it does not yet answer the questions of what the nodes should represent and whether the connections should be directed.

If the network traffic corresponds to students in the hallways, then that traffic is moving between classrooms. This suggests that the nodes---which the traffic moves between---should be classrooms, or, more specifically, the places in the hallway where there are classroom doorways.

We still need to decide whether the connections are directed. Because we have said we are interested in hallway overcrowding, one approach to answering this is to decide whether students moving in one direction affect how much space (capacity) is left for students moving in the opposite direction.

Question: Explain why knowing that we are interested in hallway overcrowding would help us answer the question of whether to make the connections directed.

It typically is the case that students moving in one direction take up space that could be used for students moving in the other direction. For now, we will choose to have undirected connections, since whichever direction they are moving, they are taking space away from people who could be moving in either direction.

Question: If we were concerned with whether passengers could get to San Francisco airport from New York's LaGuardia airport, would we be more likely to choose directed or undirected connections?

In the case of airline flights, the number of passengers on New York to California flights does not change the availability of seats on California to New York flights. In the airline example, it would make more sense to have directed connections. (This would apply both if we simply care whether such travel is possible or if we care whether there are available seats. If we were asking questions about available seats, we would use directed connections with capacities.)

Using the answers to the questions about how to set up the school hallway network, we may model the classrooms and hallways from Figure 2 using the following network:

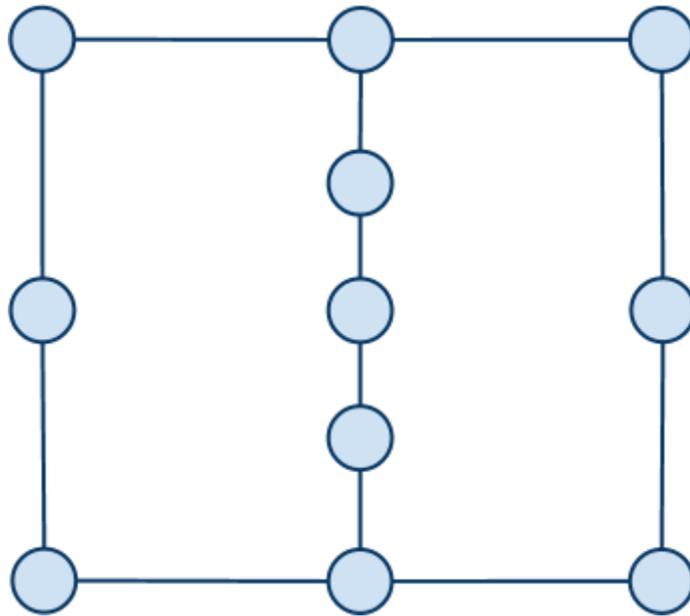


Figure 3: Network corresponding to the diagram of classrooms (nodes) and hallways (connections) shown in Figure 2.

We should make a couple of comments on this drawing of the network. First, we could have drawn it in many other ways. For example:

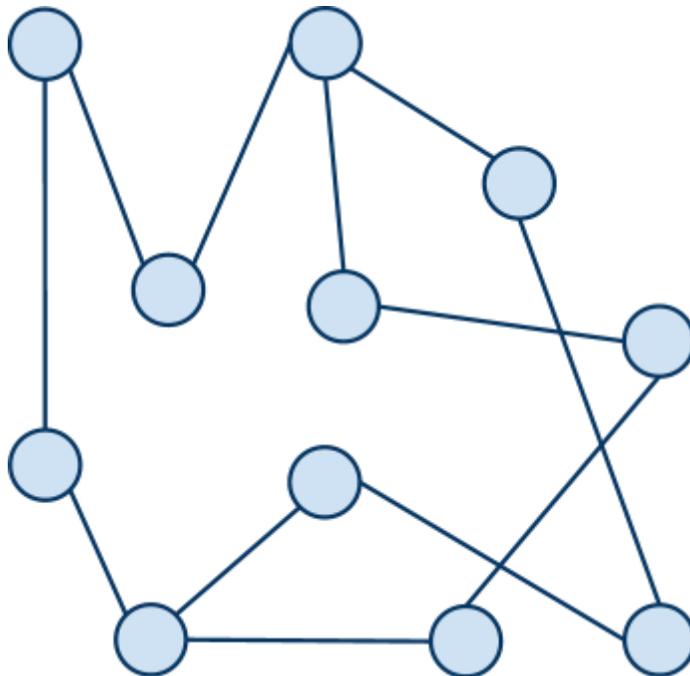


Figure 4: Another drawing of the network in Figure 3.

This drawing of the network was obtained by moving the nodes around. The connections are still the same, but it is much less clear that this network corresponds to the classroom/hallway diagram in Figure 2.

These are two drawings of the same network; in creating the network itself, we implicitly decided that one thing we would ignore was the spatial relationship of the classrooms; we only care whether their doors are connected by a hallway.

Question: Can you find a way to label the nodes in Figures 3 and 4 so that it is clear that two classrooms are connected in Figure 3 if and only if they are connected in Figure 4?

Answer: Students should be able to label the nodes so that connected nodes in Figure 3 are also connected in Figure 4. If they are having trouble, you might suggest that they look at how many other nodes each node is connected to. There are two nodes that are connected to three other nodes; all other nodes are connected to two other nodes. Note that there are multiple ways to label the drawings in the requested way.

Classroom discussion questions: What if there are two classrooms whose doors are on opposing sides of a single hallway? What can we do to model multi-way intersections?

These questions are discussed in the material below, but if students have not yet been assigned this as reading, you may want to have them suggest their own ideas as part of class discussion.

Sometimes, trying to answer the questions listed above will reveal that the problem setting cannot be modeled by a network in the way that we were hoping. Thinking more about the problem might show us a way that would work, but not always. Also, even if the problem setting can be modeled as a network, it may be that the questions we are asking cannot be described in terms of that particular network.

For example, if the nodes in the network correspond to people on Facebook, and if two nodes are connected if the people are Facebook friends, the network model does not help us answer the question of who could run the fastest in a 1500-meter race. In that case, we would need to either model the problem with a different network or with a different approach. For example, for the 1500-meter race, we could represent people as nodes, and we could draw a *directed* edge (an arrow) from a node representing a slower person to a node representing a faster person. The nodes with no arrows pointing away from them are the possible candidates for the fastest person. However, this ignores the fact that, in any particular race, someone might run faster than expected (they may have been training for the race) or slower than expected (they may be injured). Even this different way of representing people wouldn't be able to predict, with certainty, which person would win a particular race.

Before looking at different examples, let's look at a couple of variations on the example above. First, what if three or more hallways meet at a single place that doesn't have a door (so that we don't include a node at the intersection)? For example, we could have classroom layouts that look like either of the two shown in Figure 5.

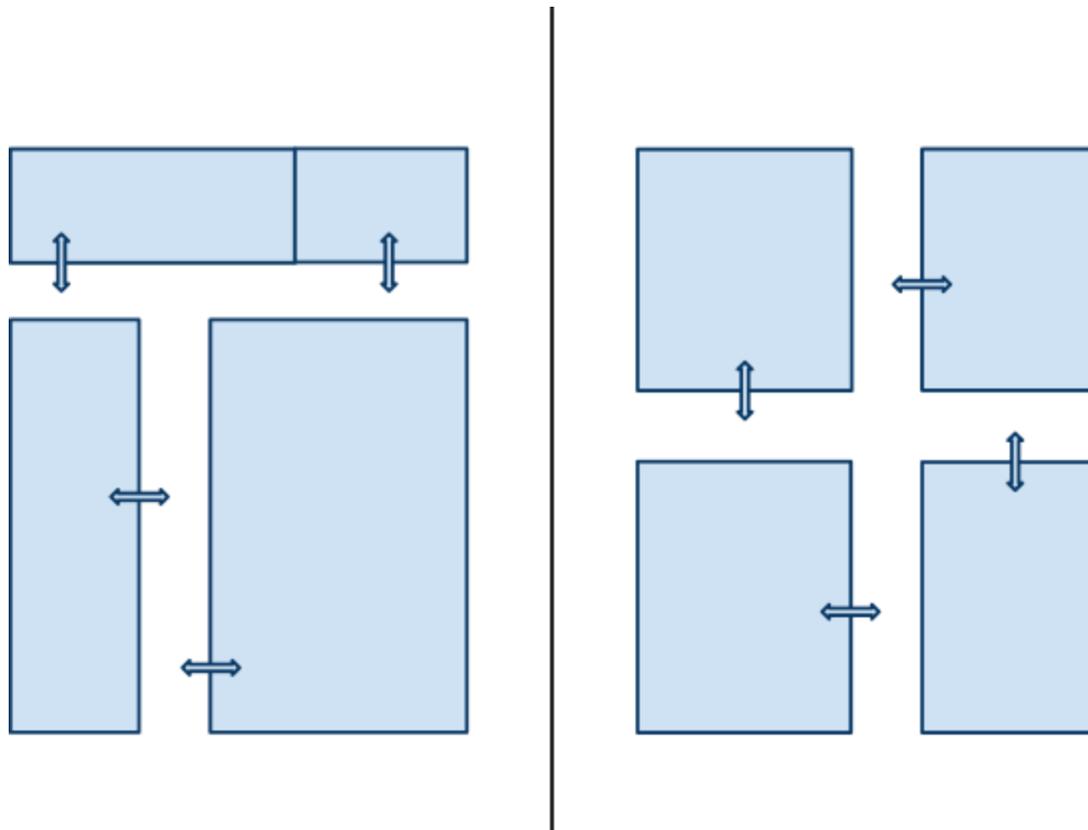


Figure 5: Parts of two other classroom layouts in which multiple hallways meet at a place without a classroom doorway to serve as a network node (three hallways meeting at a “T” on the left, four hallways meeting in one place on the right). These might be part of larger classroom layouts that are not shown here.

On the left in Figure 5, there is a three-way intersection. On the right side of Figure 5, there is a four-way intersection. Previously, classroom doorways were modeled as nodes and connections connected one node to one other node. In these examples, even if the hallways are not overcrowded, the intersections between them could be overcrowded. For example, think about what happens when the bottom hallway is completely full of students walking towards the top hallway, and the left hallway is completely full of students walking towards the right hallway. When the students get to the intersection, they want to keep going straight (up or right, depending on which hallway they're in). The way the intersection is laid out in the right part of Figure 5 suggests that the real-life intersection would be overcrowded if all the students tried to move through it at the same time. How can the situations in this figure be modeled in order to capture this?

One way that we could model these networks would be to add a node at the intersection of the hallways. This would produce networks like those shown in Figure 6. There, the node that is smaller than the others represents the intersection of hallways, and the other nodes represent doorways as before.

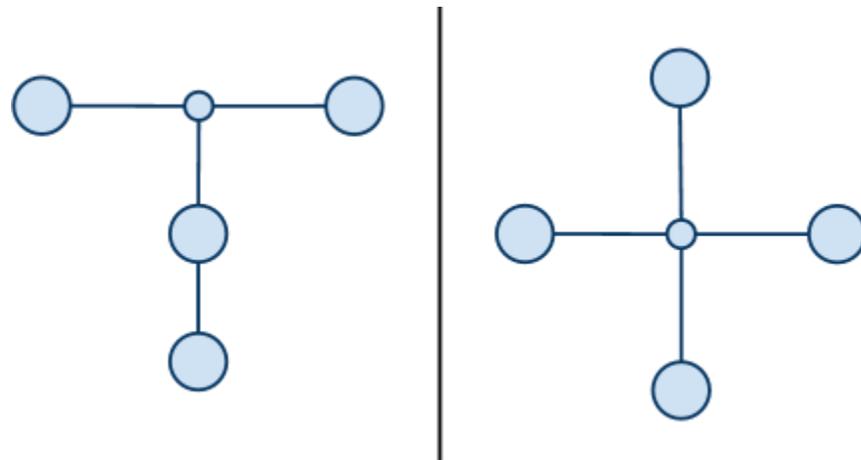


Figure 6: Parts of networks corresponding to the partial classroom/hallway layouts shown in Figure 5. These are obtained using the approach of replacing the hallway intersections with a new node (shown as smaller than the nodes that correspond to classroom doorways).

There are other workarounds to this problem. If you are trying to solve a real-world problem in which the intersections of hallways might be filled, the most natural thing to do is to also assign a capacity to the *nodes* and to add a node for each intersection, even if there isn't a classroom doorway there. The meaning of a node's capacity might depend on the particular problem that you're trying to model, and we won't discuss that approach in detail here. There are many additional ways to model these intersections, but whenever considering adding complexity to a model, you must balance that complexity with the increased ability of the model to answer questions about the situation of interest.

Throughout this module, we will try to avoid adding additional complexity by leaving our models with capacities on connections only. We will, however, allow the option of adding nodes to represent places where traffic is merging, crossing, or changing in some way.

Teacher's note: An optional exercise, listed at the end of the module, develops the idea of node capacities. If you want to use that, this would be a good place to do so.

Another consideration is, what if there were two classrooms whose doors are exactly opposite each other? If we take the approach we used initially, in which nodes are the points in the hallways where there are classroom doorways, then these two doorways would be combined into a single node. The corresponding part of the network might look like what is shown in Figure 7.

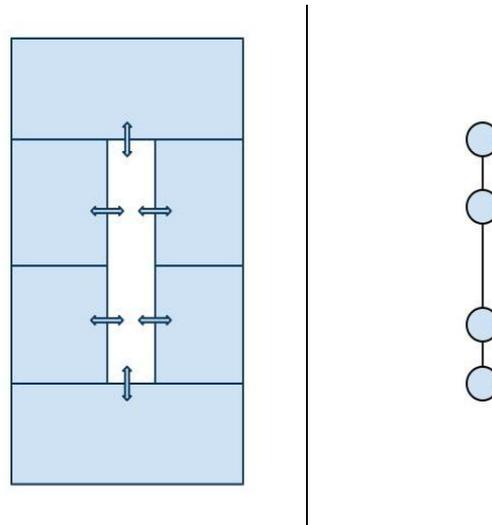


Figure 7

Students may need to get from one of these classrooms to the other one. In fact, it could be that all of the students in the classroom on the left in Figure 7 need to move to the classroom on the right, and all of the students in the classroom on the right need to move to the classroom on the left. If we model this using the fragment of the network shown in the right half of Figure 7, then all of these students moving between the two classrooms would have *no effect* on the amount of traffic in the hallway (and on how that traffic compares to the hallway's capacity).

Question: In real life, do you think that students moving between classrooms whose doorways are arranged as in Figure 7 would contribute to how full the hallway is?

Answer: Yes.

With this approach in mind, start to think about how you might be able to answer the following questions by using networks. At this point, do not worry about trying to solve the problem. Just think about how you might approach the problems keeping in mind the list of questions above.

1. Can you drive from your school to Honolulu, Hawaii on public roads?
2. Pick five cities connected by interstate highways. Could you start in one and visit all five without visiting the same city twice?
3. Knowing the capacity of a bridge, and the amount of traffic on the roads leading to the bridge, can you determine whether or not there will be a delay on the bridge?
4. There are multiple power plants that supply power to two neighboring cities. On a hot summer day, will all of the people in the two cities be able to run their air conditioners?

Lesson 2 – Simulation: Will There Be a Delay?

Now that the terminology is in place, we can create a simple **simulation**.

Simulating a simple road network using Excel

Consider the following scenario:

Alice and Bob both work at Acme Corp.; Catherine and David both work at World Wide Widgets, Inc. Alice and David live in the same subdivision, and Bob and Catherine are neighbors in an apartment building in a different town. Both companies lie on the west side of a river, while the subdivision and the apartment building lie on the east side of the river. There is only one bridge nearby, so all four people must cross the bridge to get home from work. However, the roads from the different companies to the bridge are different (Acme is north of the bridge, while World Wide Widgets is south of the bridge). Similarly, the subdivision is north of the bridge and the apartment building is south of the bridge.

Computational Thinking Alert:
Tag: Modeling
We'll now model this scenario as a network. This will omit some parts of the problem that aren't relevant to the problem.

One way to represent the roads in this scenario as a network is the following:

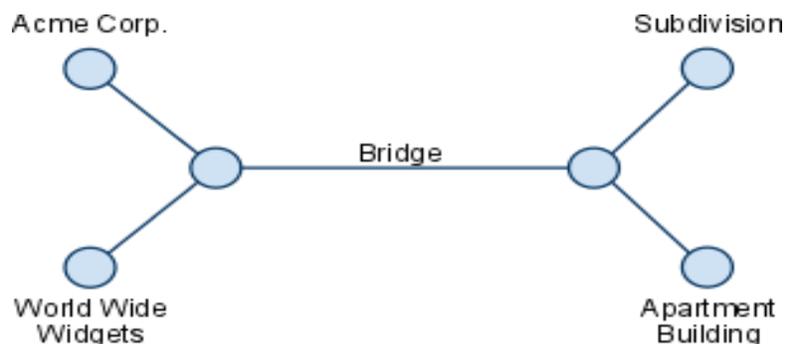


Figure 8

The local department of transportation is interested in determining whether crossing the bridge is causing delays to commuters (like Alice, Bob, Carol and David) working in the business district. Being civic minded, we offer to help answer their question using a network model. To start, we'll assume that all the roads have a capacity of 2. If there are more than two people on a road at one time, there will be a delay. (For now, we're just interested in whether or not there's a delay.) This means that there will never be a delay between an office and the bridge, and there will never be a delay between the bridge and a residence. However, because all four people are trying to drive over the bridge, there could be a delay on the bridge.

Imagine that everyone leaves work at one of two times; early or late (these might be 5:30 and 6:00). Because of different schedules at the two companies, the workers at Acme Corp. (Alice and Bob) are more likely to leave at the earlier time, while the workers at World Wide Widgets (Catherine and David) are more likely to leave at the later time. Workers who leave early drive over the bridge before workers who leave late. This means that if Alice and Bob both leave early, and if Catherine and David both leave late, then there won't be a delay. However, it's *possible* that Alice or Bob might leave late, and it's *possible* that Catherine or David might leave early. This means that there might be a delay on the bridge. We want to know how likely it is that there will be a delay.

Let's also imagine that the chance (probability) Alice leaves early is $5/6$; the chance that Bob leaves early is $2/3 = 4/6$; the chance that Catherine leaves early is $1/6$; and the chance that David leaves early is $1/3 = 2/6$. In order to get a better understanding of whether there might be a delay on the bridge, we'll run a **simulation**. This will produce four random numbers that will tell us who leaves early and who leaves late on a particular day; if three (or more) people leave early, or if three (or more) people leave late, then there will be a delay. Because we chose our probabilities as fractions with a denominator of 6, we can use regular dice to generate these random numbers.

Question: If we roll one die, what is something that happens with probability $5/6$? (**Answer:** Rolling something less than 6. Also rolling something greater than 1, or rolling something other than 2, etc.)

In our simulation of one day's traffic, if we roll one die, we'll say that Alice leaves early if we roll less than or equal to 5. We'll say that she leaves late if we roll a 6. Similarly, we'll roll another die for each of the other workers. We'll say that Bob, Catherine, and David leave early if their dice show less than or equal to 4, 1, and 2, respectively.

In-class activity: Distribute dice to students working in pairs (one die per pair). Each pair should have one student roll a die 4 times, once each for Alice, Bob, Carol, and David, in turn. The other student should record whether the person rolled for leaves early or late. Each pair should then determine whether or not there is a delay based on the number of people trying to cross the bridge at one time.

Question: Based on the values that were just rolled, was there a delay?

Repeat the die rolling activity a few times and tabulate the results. How many times was there a delay and how many times was there no delay? Express this as a probability of a delay.

Example: After rolling for 12 different days of traffic, you might have recorded 7 days with delays and 5 days without delays. In that case, the observed probability of a delay is $\frac{7}{12} \times 100\% = 58.3\%$

Question: Based on the number of trials rolled by hand, do you think you have an accurate measure of the probability of a delay?

Answer: It's tough to tell with only a few trials. We could keep doing this by hand, or we could ask the computer to help us. We could also pool the results from all students to get a better idea.

Later, we'll think about examples that would be much harder to do by hand—the computer won't just be helpful, but essential.

At this point, it should be clear that we can create a model of a simple situation by rolling the dice by hand. This is time consuming, however, to do this for more complex situations or for large numbers of repetitions. In those cases we are going to have to use a computer to handle the repetitive and time consuming task of rolling dice, checking the values against our probabilities, and recording and totaling the results. To do that, we will start by getting the computer to replicate what we just did in class with dice. After we are familiar with the process on the computer, we're going to get the computer to repeat the process much faster than we could ever do it; this will let us simulate situations that we wouldn't want to do, or that we couldn't do, by hand. This not only highlights the benefits of using computation, but it also prompts us to think about framing the problem in a way that is suitable for using computation.

Excel workbook 1: The first sheet (the sheet named “Basic (one trial)”) in the Excel workbook (network-capacity-wkbk-1.xlsx) simulates one trial rolling 4 dice. Die 1 corresponds to Alice, etc. It displays how many people leave early and how many leave late. The spreadsheet uses this to compute and display whether or not there will be a delay on the bridge. You can simulate another trial by having Excel recalculate the values in the spreadsheet. It will produce new random numbers, like rolling the dice again.

Computational Thinking Alert:

Tag: Simulation

It may be possible to construct and run a simulation by hand, if the circumstances are simple enough to run the simulation in a reasonable time. If there are many different possible behaviors of the simulation (such as when it depends on randomness or when there are a large number of options), running the simulation on a computer may help us understand the system better than we could realistically hope to do by hand

TN: See Excel's help for how to manually force recalculation. In Excel for Mac 2011, Command+= will force recalculation of the entire workbook, and Command-Shift+= will force recalculation of the single sheet. This can also be done under the Formulas toolbar at the top of

the workbook. In Microsoft Excel, F9 forces recalculation of the entire workbook, while Shift-F9 will force recalculation of the single sheet. If you do not use Excel, check your spreadsheet's help function to find how to force a recalculation.

So far, this is just like one trial of the simulation in class. We can do lots of trials very quickly, though. This will help us get a better answer. The second sheet (named "Basic (many trials)") in the workbook shows 10,000 trials instead of just 1. It now uses 1 and 0 instead of "Yes" and "No" for the delay. This means that we can sum the values in the delay column to get the total number of delays (shown in column M). We can then compute the probability of delay by dividing this by the number of trials. This is shown in column N.

Question: What happens to the probability of a delay, according to the simulation if we use 10 or 25 trials? What happens if we use 100? 1,000? 10,000?

Answer: To determine the answer to these questions, update the SUM in cell M2 and the computation in cell N2; if more trials are needed, fill down. For example, to do 10 trials, we would change the sum in cell M2 to be SUM(L2:L11); to do 100 trials it would be SUM(L2:L101).

Question: The calculated probability of a delay, with the assumptions above, is about 53.7%. Does this calculated probability come close to your observed probability from the simulation?

Answer: Will likely be "yes."

Question: Do you think we would be able to get this answer quickly and accurately by doing simulations rolling dice?

Answer: No. It would be very time consuming.

The third sheet (named "Basic (cutoffs)") lets us change the probabilities that the different workers will leave early. You guarantee that someone will leave early by setting their first value (in the "Early if <=" column) to 6 and guarantee they will leave late by setting their first value to 0. (Each person's second value is computed from their first value, so do not change those by hand.) What happens as you change these probabilities? Look at both the simulated probability in N2 and the computed probability in Q7.

Question: What happens if everyone has an equal chance of leaving early or late---is the chance of a delay 50%?

Answer: No. The chance of delay is 62.5%. If the students seem surprised by this, ask them to think about whether they think the chance of delay should be higher or lower than the previous instance in which Alice and Bob were more likely to leave early and Carol and David were more likely to leave late. Also, it might help to point out that the *only* instance in which there is no

delay is when *exactly* two decide to leave early. This means that if 0, 1, 3 or 4 decide to leave early, there will be a delay.

Finally, the fourth sheet (named “Basic (cutoffs + capacity)”) lets us vary the bridge capacity. (The computation for the theoretical probability of a delay changes with the capacity, so that is not given on this sheet.)

Question: What happens if we set the capacity to 1?

Answer: There will always be a delay. Even if only 1 person leaves early, the only way to avoid a delay at the earlier time, there would then be 3 people leaving late, causing a delay then.

Question: If we set the bridge capacity to 3, can we eliminate all delays?

Answer: Not if it’s possible for everyone to leave early or for everyone to leave late. That would put 4 people on the bridge, though the probability it occurs is quite small.

Extension 1: Update the simulation so that it can use any probability, not just multiples of 1/6 (which was motivated by the ability to use six sided dice).

Teacher’s Note: Note that it’s harder for us to run simulations with arbitrary probability values (ask students how they might physically simulate something that has probability, e.g., 0.42; perhaps note difficulty for computers as well, but computers can approximate things better than humans can), so being able to do this on a computer is even more important with decimal probabilities than it was with multiples of 1/6.

Extension 2: Assume departure times are normally distributed (possibly with different means). How should we update the model to deal with this? You can use the NORMINV (or NORM.INV in newer versions of Excel, although the old name also works) function. NORMINV(prob, mean, std_dev) takes three arguments as input; the second and third (mean and std_dev) describe the mean and the standard deviation of the distribution. A way to use this to output numbers that are chosen according to the uniform distribution with a specified mean and standard deviation is to give a number chosen uniformly at random (for example, using Excel’s RAND() function) as the first argument, 0 as the second argument, and the desired standard deviation as the third argument, and then adding the desired mean to the result. What results do we get now?

Teacher’s Note: You might ask the students to discuss how this might affect theoretical calculations. These will become harder to carry out. Other distributions of departure times might be even more difficult to work with theoretically, which makes the computer simulations especially useful.

Teacher's Note: The Excel workbooks modify the output of RAND() so that it is between 0.005 and 0.995, which ensures that the number produced by NORMINV() is within a reasonable range for the problem.

Extension 3: Consider a more complex network or a more complex behavior to simulate. How can we design a simulation for this network? What can we learn about it?

Lesson 3 – Simulation: Average vs. Peak Demand

The last page of the Excel workbook for the previous simulation allows variations in bridge capacity. The capacity was originally set at 2. What happened to the probability of a delay when we changed the capacity to 1 or 3? What would happen if we changed it to 4?

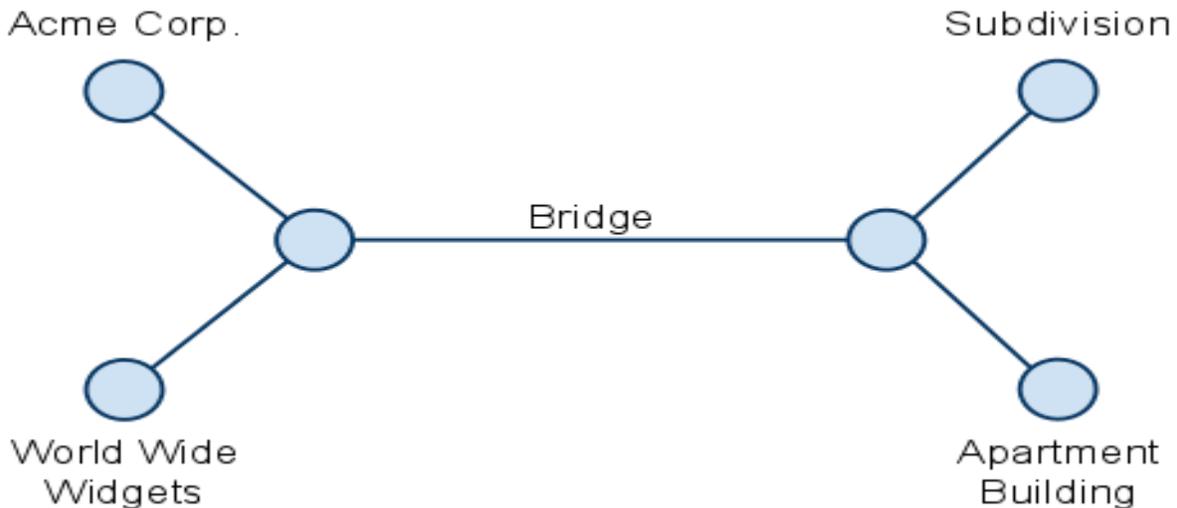


Figure 8 (again)

Exercise 1:

With capacity 1, there is 100% chance of a delay—if only 1 person crosses the bridge early or late, so that there's not a delay at that time, then 3 people try to cross the bridge at the other time, creating a delay then. With capacity 4, even if everybody leaves at the same time, either early or late, the bridge will have enough room for all of the people. With capacity 3, there will be a delay if everybody leaves early or everybody leaves late. This can happen, but the probability is much less than just having 3 people leave early or 3 people leave late. Using the same probabilities for time leaving work as before, use the simulation to find the expected probability of delay.

What probability of delay did you find when the bridge capacity is 3?

(Answer: results will vary, but should be approximately 6% of the time.)

Can we calculate this expected probability of delay? Yes, to compute the probability of everybody leaving early or everybody leaving late we start by calculating the probability of each event separately. The probability of everyone leaving early is the product of the probabilities of each person leaving early or $5/6 * 2/3 * 1/6 * 1/3 = 5/162$. Similarly, the probability of everyone leaving late is $1/6 * 1/3 * 5/6 * 2/3 = 5/162$. Because the two events can't both happen on the same day, the probability that one of these things happens is the sum of these probabilities, $10/162$, or 6.1%.

Let's look at some other problems using this same road network. Imagine that there are three possible departure times: 4:30, 5:00, and 5:30. At each of these times, some people leave both Acme Corp. and World Wide Widgets and drive across the bridge. However, now there are more people than just Alice, Bob, Catherine, and David. As before, they can leave at different times, and we're curious about whether there will be delays due to too many people leaving at one time.

Excel workbook 2: The first sheet (named "Fixed demand") in the associated Excel workbook (network-capacity-wkbk2.xlsx) considers a single day in this scenario in which exactly 5 people leave from each company at each time. The first three columns show information for 4:30: the number of people from each company and whether or not there is a delay. The second set of three columns shows this information for 5:00 departures, and the third set of three columns shows it for 5:30 departures. We can easily see that, if the bridge has capacity 10 or greater (the value shown in Q2), then there will never be a delay. However, any capacity less than 10 will produce a delay in each time period.

That's pretty straightforward. Let's look at something slightly more interesting. Over many days, we could count the number of people leaving each building on each day at each time. From those counts we could calculate the average number of people who leave at any given time on any given day. From our previous work, we have seen that we cannot just use this number to determine if there will be delays. Some days there will be more than the average leaving at any given time, and other days there will be fewer. What we want to do is generate a series of random numbers, such that the predicted average of the numbers is the average of the number of people counted leaving at each time. There is one other aspect of this average that is needed. We also need the standard deviation. The standard deviation is a measure of how spread out the numbers are. The lower the standard deviation, the closer to the average that the random numbers will be, while a high standard deviation will result in widely varying numbers. We can calculate both the average and the standard deviation from the raw data of people counts by time and then use those numbers to generate random numbers to predict how often there will be delays on the bridge.

If you are unfamiliar with the concept of standard deviation, the following example will illustrate its nature. We have two sets of numbers {4, 4, 6, 6} and {2, 2, 8, 8}. Both sets have an average (or mean) of 5, but looking at the numbers, we can see that the two sets are not the same. Standard deviation is a way of describing that difference. The standard deviation of the first set is 1, while the standard deviation of the second set is 3. To use the Excel worksheets, you do not have to know how to calculate the standard deviation, but you should understand that the smaller the standard deviation, the closer the values are (relatively) to the mean.

The second page (named “Normal”) in the workbook shows a single day in this scenario, but now the number of people who leave each company at each time is randomly generated by the spreadsheet. (These random numbers are produced using a “normal” distribution, hence the name of the Excel sheet.) At each of the three times, a random number of people leave both Acme and World Wide Widgets. This random number is generated so that, on average, it equals 5 - the same number that always left each company at each time in the example on the previous page in the workbook. The random numbers that are produced so that their standard deviations are 2 (the values given in Q4 and Q6).

This page initially uses a capacity of 10 for the bridge. Did this produce any delays? Have the spreadsheet recalculate a few times so that it generates new random numbers for the number of people leaving each business at each time. Do you often see delays? Is this surprising (given that the *average* number of people leaving at any one time---10---won’t cause a delay)? What happens if you increase the capacity to 11, 12, or more?

Exercise 2:

Before moving on, test different capacities and guess, based on a few recalculations for each capacity value, what capacity gives a delay about $\frac{1}{3}$ of the time.

We are going to increase the capacity to one more than you just guessed.

A. What is your prediction for how often this will result in a delay?

B. What is your prediction for what will happen if we once again increase the capacity by one?

Teacher Note: In the next worksheet, we’ll run many trials to help us get better estimates for these values, but it will likely be difficult to make accurate guesses based on just a few recalculations. It may be interesting to note the range of values guessed by the class. Is there a consensus? Later, it may be interesting to look back at how this compares---and how confident students were about these guesses---with the estimates produced by hundreds or thousands of simulations.

The third worksheet (named “Constant (normal)---short”) replicates the computations from the second worksheet 10 times. It now shows the number of times (and the corresponding percentages) that there is a delay for people departing at the earliest time (4:30), for people departing at the middle time (5:00), and for people departing at the latest time (5:30). If you have Excel recalculate these numbers, they may change a bit with each recalculation.

Exercise 3:

To get a better estimate change the worksheet to run 1,000 trials. How much did the numbers seem to change each time you recalculated the spreadsheet with 10 trials? What about with 1,000 trials? (You might also try 100 and 10,000 trials to see their effects on the answers.)

(Remark: This is done already in the following worksheet, named “Constant (normal).” This also shows: the total number of people used in each trial; the average number in

each trial who leave per time window; and the average, over all trials, of the number of people who leave per time window. We'll now assume that we're working with that worksheet or the equivalent one that students have modified.)

With the capacity (as given in the worksheet) set at 11, we see that there's just under a $\frac{1}{3}$ chance that there will be a delay.

Question: How did this compare with your guess earlier?

Answer: Students should be able to recognize whether or not they were close in their estimate.

Question: How big does the capacity need to be so that the probability of a delay is about 5%?

Answer: 14

With the Excel worksheets, we've looked at scenarios in which the average number of people leaving each company was 5 in each time period. (In the first scenario, this was because the number was always equal to 5.) That meant that, on average, 15 people left each company per day. Some days, there might be workers on vacation, and so fewer people leaving, while on other days there might be visitors to the company, and so more people leaving at the end of the day. The next worksheet, named "Variable (normal)," keeps 15 people per day, *on average*, for both companies, but it changes the number of people, on average, who leave at each time. Initially, this worksheet is set up so that the average number of people leaving Acme and World Wide Widgets are given in Table 1. Each time period has a standard deviation of 2 for each time interval. These values are given in columns R and S. The capacity is constant; the other values are given separately for each of the three time periods.

Company	Leaving at 4:30	Leaving at 5:00	Leaving at 5:30
Acme	4	4	7
World Wide Widgets	4	5	6

Table 1 – The average number of people leaving the three companies at three different time periods.

This means that, on average, there are still 15 people leaving each company each day, but the average number leaving at a specific time changes with the time. The model now has a rush hour built into it, so it should probably seem more realistic.

Given a capacity of 14, what do you think will happen to the delay probability as a result of these changes? Before you answer, you might think about whether or not there would be a delay at each time if the number of people leaving Acme is always 4 at 4:30, 4 at 5:00, and 7 at 5:30, and if the number of people leaving World Wide Widgets is always 4 at 4:30, 5 at 5:00, and 6 at 5:30.

Note that, even though the average number of people leaving each company is unchanged, the probability of delay has changed a great deal. Keep in mind that our previous example with a capacity of 14 gave a delay probability of around 5%.

Exercise 4:

Manipulate the capacity until you have about a 5% chance of delay for people leaving at 5:30.

A. What capacity was required to accomplish this?

(Answer: 17)

B. What does this do to the probability of a delay at the earlier times?

(Answer: The calculated delays for both earlier time periods are zero.)

Discussion Question: If you were in charge of planning the road network for the city in this example, what capacity do you think you would use? What percentage of a delay probability would you think acceptable? If city planners accept 5%, do you think it is sufficient to construct the bridge with capacity of 14, or should the capacity be 17, or something else?

This example illustrates the difference between **average demand** and **peak demand**. In all the cases discussed here, the **average demand** is 10 people per half-hour (when averaged across all three departure times). The **peak demands** are very different, though. When the demand is the same at each time (and is always 10), a capacity of 10 would guarantee that there would be no delays. When the demand is the same at each time (with an average of 10), then having enough capacity to guarantee a low probability of delay at 4:30 means that there is also a low probability of a delay at 5:00 and 5:30. However, when the demand is not the same at each time, then we need to look at the demand at different times. It is not sufficient to say that, because 30 people per day (on average) want to drive across the bridge, we can ensure a low delay probability by having a capacity of 14. Instead, we need to look at the time of peak traffic (5:30) to see if we have enough capacity to ensure that delays will be unlikely. **(Remark:** This might not be realistic because of cost and other considerations. We'll look at those separately.)

Extension: What if the bridge capacity comes at a cost, and delay probabilities come at a cost. Explore the tradeoffs.

Extension: What are other examples of networks in which the demand varies?

(Remark: Electrical networks; air-traffic networks; hallways/exits in a school; many more.) As one example, the high long distance call volume on Mother's Day and the subsequent high rate of call blocking because of insufficient network capacity was widely known throughout the telecommunication industry as the "Mother's Day Phenomenon". Though the rate of blocking is now much lower, Mother's Day is still recognized as a peak demand scenario in telecommunication.

Look up flight information for travel around a holiday, and compare it to flight information for the same day of the week a week or two earlier/later. Do you see any differences? Is the capacity the same, or does it change?

Extension: What are other, perhaps more realistic, ways of capturing traffic delays? (As a start, what if there is a delay at 4:30, and enough people leave at 5:00 to completely fill the bridge but not cause a delay. Do you think there should be a delay at 5:00?) Update the spreadsheet to model one or more of these other ways of determining delay. **(Remark: The questions just posed could be addressed simply by saying that there is a delay if there was a delay at the previous time and the bridge would be filled up by the people leaving at the next time. Another---better and more general---delay model might say that there is a delay at time t if, for any number of time slots k , in the k slots up to and including the slot at time t , there have been more than c times k people trying to cross a bridge that has capacity c . For example, if the capacity is 10 and if 14, 9, and 8 people leave at 4:30, 5:00, and 5:30, respectively, then there will be a delay at each time, even though the people leaving at 5:00 and 5:30 would not cause a delay by themselves. The intuition is that we need to wait until the backup has been cleared before we can avoid delays. Of course, this doesn't perfectly capture the real world---people will leave at times other than 4:30, 5:00, and 5:30---but it may do better than the initial model.)**

Lesson 4: How Much of a Delay?

In the simulations we studied earlier, we simply asked whether or not there would be a delay in a traffic network. We did not ask **how much** of a delay there would be, but it is easy to imagine asking this question. In the previous simulation, we discussed allowing a delay in one time period to cause a delay in a later time period if a very large number of cars had tried to use the bridge all at once. That captures some aspect of the extent of a delay, but we'll focus more explicitly on that here.

We can now introduce a new type of consequence for increasing traffic. Rather than limiting traffic to a fixed capacity, increasing traffic will increase the time to traverse the edge. As an example, the capacity of an edge can be listed as “the time it takes to traverse a hallway is $1 + S/10$, where S is the number of students trying to traverse the hallway.” This can be framed as a question like “How can we calculate the delay if it depends on the number of students in the hallway?”

Do you think modeling delay based on the amount of traffic is more realistic? Why or why not?

Does this increase the complexity of the school hallway simulation?

Ask students to provide some examples of different types of networks with different types of capacity (as either delay or hard limit). Do students have a feel for which capacity type would be better to simulate the situation?

Cell phones can be a hard limit with which they might be familiar. Either the call can be placed, or the circuits are busy and the call does not go through.

Other examples of networks with delays are road networks, airlines and trains.

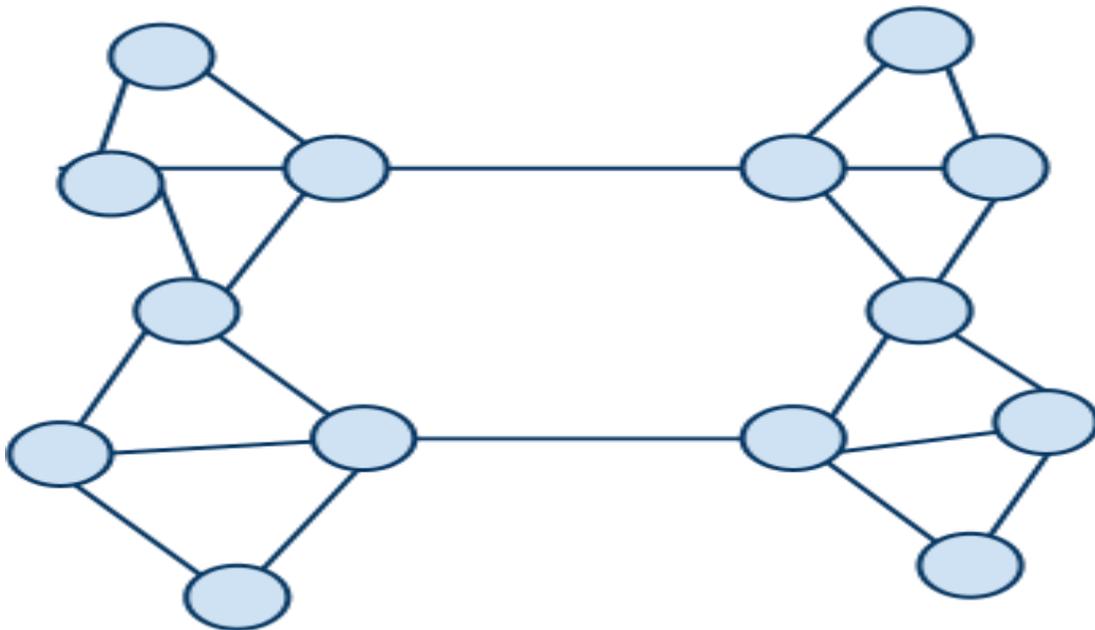


Figure 9 - A more complicated road network.

In Figure 9, the network is too complicated to simulate by hand, so a computer will have to be used to do the many calculations which will be necessary. This network, however, introduces an additional problem. How does the traffic decide which way to go? There are multiple routes between any two nodes and even if we are only concerned with the traffic ultimately traveling over the two bridges in the middle of the diagram, how do commuters decide which bridge to go to?

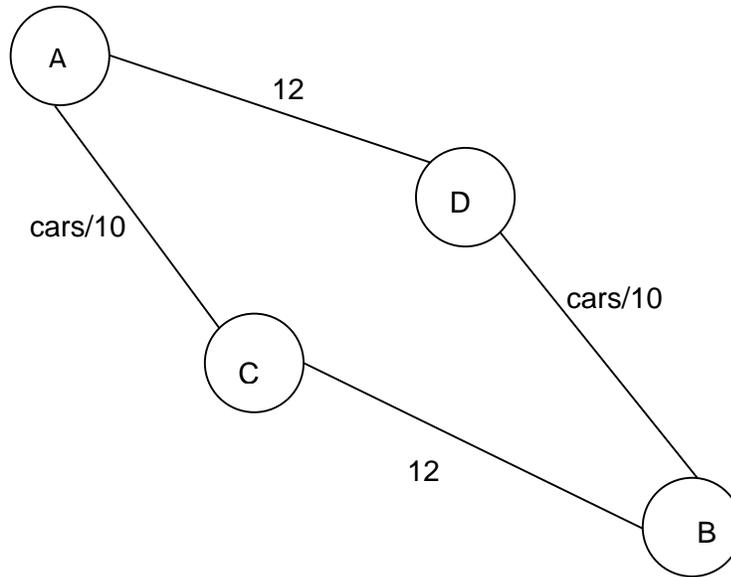


Figure 10: Braess' paradox, figure 1

Consider the preceding network diagram, (Braess' paradox, figure 1) where traffic is moving from point A to point B and the drivers have a choice of routes to take. They can either take the high road with a fixed cost of 12 minutes plus a variable cost of number of cars divided by ten minutes or the low road with a variable cost of number of cars divided by ten plus a fixed cost of 12 minutes. It should be clear that if 100 cars are driving from A to B, the best solution for any individual driver is to find the route with the least number of cars and travel that route. This will cause both routes to end up with 50 cars on it and a total traversal time of 17 minutes.

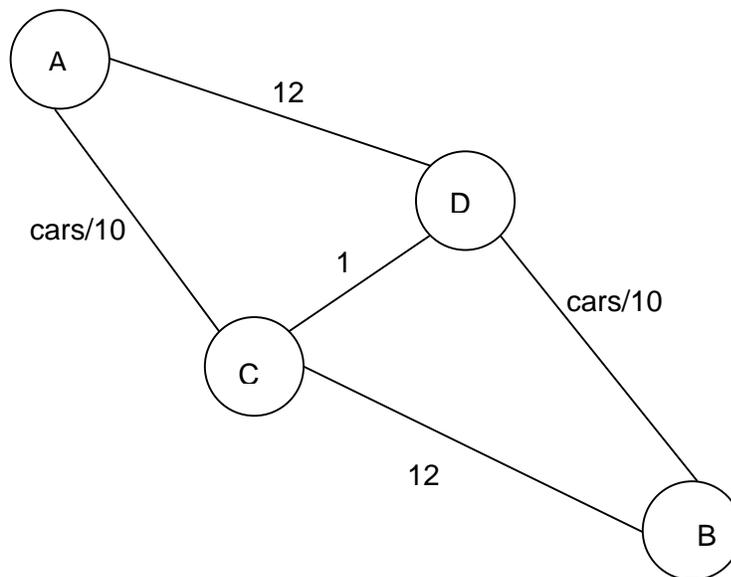


Figure 11: Braess' paradox, figure 2

To alleviate congestion, an additional connection is going to be added to connect the two roads between points C and D, as shown in Braess' paradox, figure 2. The travel time of this connection will be very low at 1 minute. Now the 100 drivers will take the low road to point C, the new connection to point D and the high road to point B. Since each of the 100 drivers will be on this route, the travel time will be 21minutes. That is longer than the old route! However, no individual driver can do better by switching to an alternate route, so they will all continue on the described route. This is known as Braess' paradox. A **paradox** is a statement or situation that seems to contradict what is expected.

Students can look up examples on the Internet of when Braess' paradox was recognized in real-world road networks when one of the roads was closed.

Lesson 5: The Cost of Capacity

We've looked at examples in which there would likely be a delay in network traffic unless the capacity of some part of the network was increased. Why do you think you don't see 20-lane bridges/roads, and why don't airlines always fly jumbo jets instead of smaller planes? **(Remark: These things would be very costly, and the costs would not be outweighed by the benefits that would be expected. For example, a 20-lane bridge might be useful---if there are also very big roads leading to/from it---for a little bit of time each day. However, the cost---both in terms of money and other issues, like the effects on neighborhoods or nature---to build all of these roads in addition to the bridge would be too high just to cut rush-hour delays by a little bit.)**

There is almost always some cost associated with adding capacity to a network. The cost may vary, but it needs to be evaluated and compared with the benefits gained by having the increased capacity. Let's return to our previous family of examples and do this.

- Imagine that the cost of a delay per person is X
- Evaluate the expected cost for a week or a year
- How much is extra capacity worth if the bridge lasts for Y years?
- Consider multiple capacity levels and associated costs (expressed in terms of some other cost/lane). How much capacity would you build if the cost/lane were Z (for different values of Z)?
- See http://www.dot.ca.gov/hq/esc/estimates/COMP_BR_COSTS_2010-eng.pdf for some sample bridge costs.

Projects and extensions for peak demands and capacity:

Have students research, on the Internet, airline departures and arrivals from a local airport on four or five different days of the year. Have students define the nodes, edges, traffic, capacity and consequences. Is there a difference in seat capacity flying from one city to another city on each of these days? Find ticket prices on those days. What is the implication of the price differences? The number of seats changes the capacity and the price is intended to change the traffic.

Contrast this to the electrical grid or hallways which have a fixed capacity. Do electric utilities also change the rates dependent upon grid usage? Does your school schedule classes in an attempt to reduce traffic in certain hallways? Teachers can have a student from the class ask the principal or guidance counselors if this is taken into account. Have students discuss where bottlenecks occur within their building and if they can suggest scheduling changes that may alleviate them. A potential project is to have students model traffic in the hallways.

Can the students think of other changes to network capacities or attempts to manipulate traffic? As students come up with examples, have them define the nodes, the edges and capacity. Examples include cell phone plans, floodgates and dams, variable traffic lanes on bridges/commuter routes/ toll booths, bus and train routes, blood vessel dilation and constriction, etc. The eventual point that students should reach is that the issues of dynamic capacity arise when

traffic is an issue. Perhaps ask students: What are examples of networks where there doesn't seem to be an attempt to manipulate traffic? Why do you think that's the case? An example of this is that there are few country roads where lanes are dynamically allocated.

Homework/projects:

1. End-of-week: Take a network that has been identified in the wiki and determine how it can be modeled as a network. Design a simulation that could be used to answer relevant questions about this network.

Extension for simulations: Use netlogo (<http://ccl.northwestern.edu/netlogo/docs/>) to simulate the networks especially if the module is being done in a programming class. This has the advantage of allowing the students to perform some decision making on the network while the simulation is running.

Optional Exercise to Develop Node Capacity

Lesson 1 mentioned that adding a concept of capacity for nodes might be the right way to model some real-world problems. Work through the following examples to see how this might be done.

First, consider the hallway diagram in the left of the following figure:

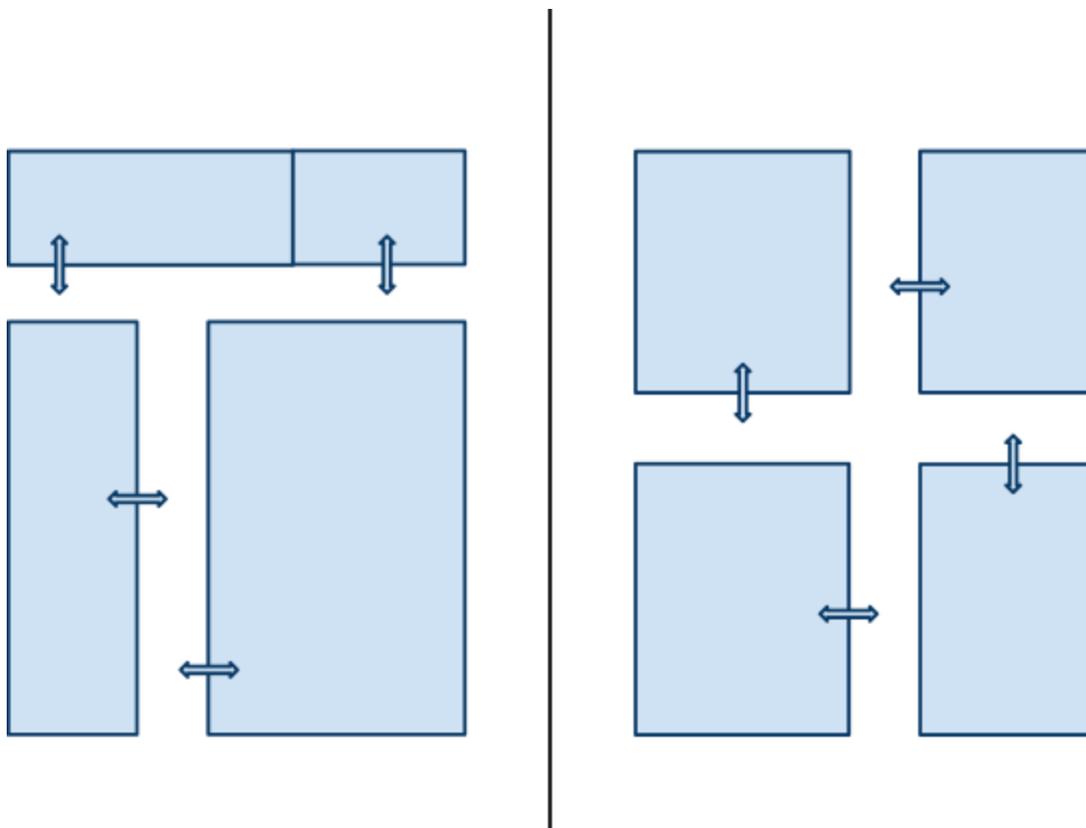
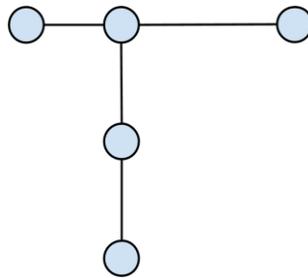


Figure 5 (again)

(This is the same as Figure 5.) Imagine that the hallways have certain capacities and that the same number of people who can fit in a hallway can also fit in the intersection. (For example, if the hallway is full, and all the people in the hallway walk through the intersection, then the intersection will also be full.) Represent this hallway diagram as a network with nodes and connections, but now add a node to represent the intersection of the hallways.

Answer: One drawing is shown below. Here, the node representing the intersection is the same size as the other nodes. The position of the nodes is intended to evoke the location of the doors and intersection in the figure above, but that is not an essential aspect of the answer.



Answer figure: Network corresponding to hallway layout in left part of Fig. 5

We now want to assign a capacity to the node that represents the intersection of the hallways. What do you think this should be?

Answer: The most natural capacity is probably the same capacity as the hallways have because we said that, if enough people to fill the hallway walked down to the intersection, then the intersection would then be filled up.

How this capacity is used is a different question. If people are just moving through the intersection from one hallway to another, then the intersection's capacity is very much like the capacity of a hallway. However, if people from hallways are entering the intersection and staying there (for example, if they are reading a bulletin board posted in the intersection), then the intersection could fill up. To make this more concrete, imagine that the three hallways and the intersection all have capacity 6. If 2 people walk down each hallway to the intersection and arrive there at the same time, the intersection will be full (but everyone can fit). However, if those people stay there and then another two people walk down each hallway to the intersection, the intersection will be full. Importantly, the capacity of the intersection captures the number of people who can be there at any one time. For hallways in which people are moving, the capacity

(in terms of how many people can be in the hallway at any one time) is similar to the number of people who can move through the hallway in a fixed amount of time (for example, 6 people every 15 seconds). However, for an intersection where people might linger for a while, there may be a mix of people moving through the intersection and people standing there.

Now, consider the example shown in Figure 12.

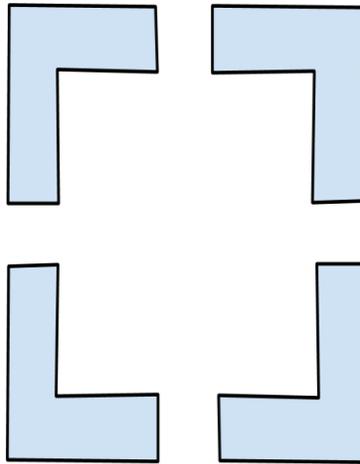


Figure 12: Four hallways leading into a large lobby

This shows four hallways leading into a large lobby. The lobby acts like an intersection of the hallways because people can move from one hallway to another through the lobby, but it is different than a simple intersection because there may be extra people in the lobby who are not moving from one hallway to another. Because the lobby is so big, even if enough people to fill up a hallway move into the lobby, there will still be additional room in the lobby. First, represent this as a network with nodes and connections.

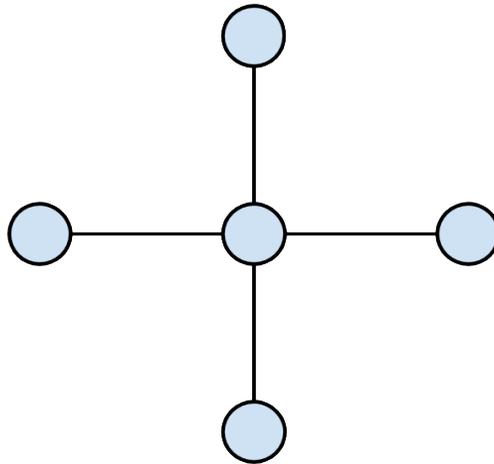


Figure 13: Network corresponding to Figure 12, with a node for the lobby

Now, think about how you might assign a capacity value to the node corresponding to the lobby. Should this be same value as the capacities of the hallways? If not, what do you think its value should be?

Transfer Assessment Problem

This exercise involves a collection of linked computers. This is often called a computer network, but that term uses “network” in a slightly different sense than the abstract way that we have been using that word. However, this should suggest to you that it will be possible to model the collection of computers as a network as we have done for other settings. Figure 14 shows how the computers are connected.

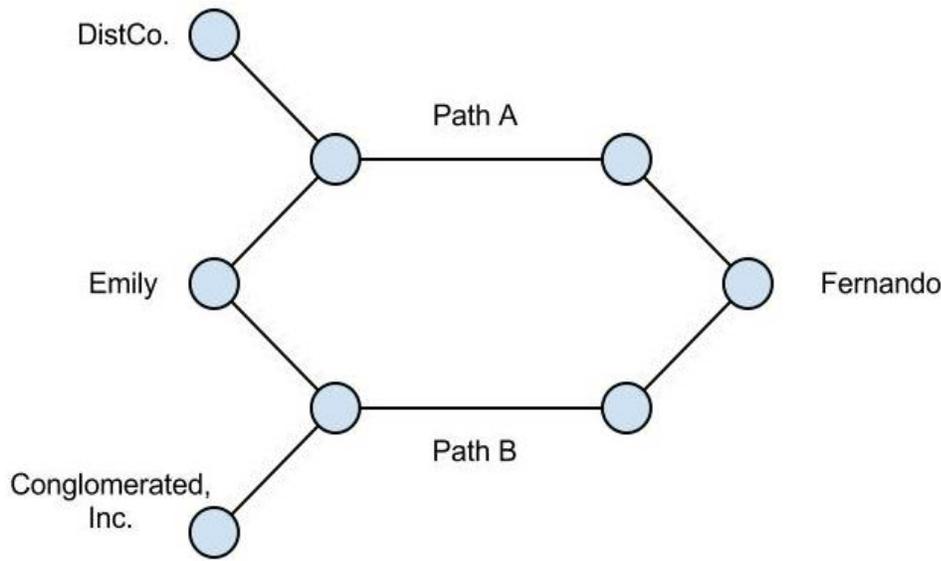


Figure 14: Emily’s and Fernando’s computer connections through Paths A and B showing other companies connected to those paths.

Two different schools need to exchange information regularly. Emily manages the computer systems at one school and can determine which of two paths (labeled as Path A and Path B) will be taken by the information that her school sends to the other school. Sometimes this is done automatically, but we will not worry about the details of how the path is chosen. Each of these paths contains multiple links between different routers—the data is sent from a computer in Emily’s school to a router, which forwards it to the next router on the path, which forwards it to the next router on the path, and so on, until the data reaches a computer at the other school. A local business on each of these paths also exchanges information using the paths. If a router attempts to send too much data across a link to the next router on the path, some of that data will be lost and will need to be resent, adding to the traffic on the connection and slowing down the process.

Emily is thinking about how to choose paths at two different times during the day. The amount of traffic that her school sends to the other school is different in the morning than in the afternoon, and the amounts of traffic that the local businesses that share the links on the paths to the other school send also differ between the morning and afternoon. Emily wants to decide how to divide the data that her school sends between the two paths at each time. Using the information below, model this problem as a network and answer the questions below.

Scenario 1:

Data transmissions are randomly distributed in a uniform manner (also known as “uniformly at random,” or “UAR”) Saying that the amount of data sent by a computer is chosen uniformly at

random from the set {4,5,6,7,8} means that each of these values has an equal chance of being chosen for the amount of data to be sent. This contrasts with the normal distribution we used in previous scenarios where some values had a greater probability than others.

The following table shows possible values for how much data each of the three sources might send at each of the two times of day (morning and afternoon). For example, DistCo is listed as possibly sending “3-7” units of data in the afternoon. We use this to mean that DistCo might send 3, 4, 5, 6, or 7 units of data in an afternoon. We assume that all of these values are chosen UAR from the ranges given.

Time	Emily	DistCo	Conglom
A.M.	5-13	4-8	2-6
P.M.	7-15	3-7	3-7

Table 2: Ranges for how much data is sent in network from Figure 14

As an example, Emily sends between 5 and 13 units of data in the morning; there are 9 (= 13 – 5 + 1) possible values. If we want to construct a corresponding simulation, we would choose a random number between 5 and 13. You could do this by drawing slips of paper from a pile. If you have a 20-sided die with the numbers 1-20 on its faces, roll the die; if the number shown is one of the values for Emily, use it, otherwise roll again. (Note that it doesn’t work to roll two six-sided dice and look at their sum, because there are more ways to get some sums than others.) We can determine the amount of data for the companies similarly.

Assume that path A has a capacity of 12 and path B has a capacity of 8. Create two spreadsheets. On one calculate the amount of traffic lost if Emily has to assign a fraction of her traffic to each path. To do this, multiply the traffic she generates by the fraction she assigns and add it to the amount generated by the other company on that path. If the amount is greater than the capacity of that path, calculate the amount of traffic lost (the difference between the traffic and the capacity.) Change the fraction devoted to each path to minimize the total of lost traffic over the number of random trials. Dividing this total by the number of trials will give you the average amount lost per trial.

On the second spreadsheet assume that Emily must assign a constant amount of traffic to path A, and that anything over that amount will be assigned to circuit B. Again try to determine what value should be used to minimize the amount of traffic lost in both the morning and evening. (Allow her to choose different values in the morning and afternoon.)

Scenario 2:

Use a normal distribution (instead of choosing the values UAR) for the data being sent by Emily and each of the companies. The average data and standard distributions are listed in the table below.

Time	Emily		DistCo		Conglom	
	Average	Standard Deviation	Average	Standard Deviation	Average	Standard Deviation
A.M.	9	2	6	1	4	1
P.M.	11	2	5	1	5	1

Table 3: Parameters for normal distributions to use with Figure 14

Set up two more spreadsheets. One with a fractional distribution of Emily's traffic and one with a constant distribution to path A with overflow going to path B. Again, try to determine which values will minimize the amount of traffic lost.

Note to the teacher. It might be worthwhile to have students broken into groups and assign one method of dividing traffic from each scenario to each group, rather than having all of the students or groups trying to create all of the spreadsheets.

Glossary – Definitions of certain words as they are used in this module.

Capacity –The maximum amount of traffic that a connection can handle either at a given time, or over a period of time, depending on the situation being modeled.

Connection – A link between two nodes which represents a relationship between the nodes allowing traffic of some sort to flow between the nodes.

Directed connection – A connection between two nodes with traffic restricted to flowing in one direction only.

Network – A collection of nodes and connections used to model a real world problem.

Node – An abstraction of objects in a problem. Individual nodes will either have or not have connections to other nodes.

Simulation –A model of a system that captures details of how the system works, often for use in analyzing the system.

Traffic – The amount of usage a connection is currently undergoing.