Strategic Inventory Placement Model (SIPModel) Assignment

The assignment is to experiment with the SIPModel software package. With this software application you can build a network model of a supply chain. The software has the capability to determine strategic safety stock levels that minimize the total inventory costs across the supply chain. The software can also determine the costs of an existing system based on its current inventory levels.

You can download the software from

http://web.mit.edu/lfmrg3/www/

And click on "SIPModel". Follow the instructions to download an education version of software. (You will actually be taken to a company web page which now hosts the trial software)

The following is an overview of how to use the application. There is also a help file with the application. The download site also has a very extensive tutorial on using the software

NOTE – there are many types of inputs and controls in the software – you only need to use the ones described in this note. Also, this is the academic version of the software – as such, it restricts your access to some of the software features.

To Build a Network Model: When you open the application you will find a tool palette on the right hand side of the second row. By using the tools in the tool palette you can build a network. By selecting the "rectangle" you can create stages, or nodes for your network, by clicking in the window. Stages represent a major processing function in the supply chain; examples of stages include final assembly, the procurement of a part, and the shipment of product from the regional distribution center to a retailer. By selecting the blue "arrow" you can then connect stages, by first clicking on the upstream stage and then clicking on the downstream stage; this creates the arcs for the network. The third tool (an "x") can be used to erase either a node or an arc in the network by clicking on it.

You can use the options in the "displayed" menu to view the data associated with each stage in the network. We often use the **icon view**, in which the circle denotes a processing function and the triangle denotes that safety stock is to be held at the stage. If there is no triangle at the stage, then the stage does not hold any safety stock.

Inputs for the Model: From the icon view, one obtains access to an input record by double-clicking with pointer (white arrow) on a stage (node). The following inputs are possible for each stage:

Procurement Time/Production Time -- for raw material stages, the time required to obtain the part from the supplier. For other stages, the time required to produce the product once

it has the raw materials to begin production. It includes wait time, processing time and any transportation involved in getting the product from the stage's supplier.

Service Level -- the probability that the stage will fill demand each period within its quoted service time (this cannot be changed in the academic version)

Cost Added – per-unit direct costs added at the stage; typical costs include direct labor and materials.

Stages that do not have any outgoing arcs correspond to demand nodes; each demand node represents one end item. These stages have additional inputs:

Mean Demand - the average per period demand at the stage

Standard Deviation of Demand - the one period standard deviation of demand at the stage

Max FGI Service time - the maximum service time which the stage can quote to its customers.

Decision Variables for the Model

Service Time - The service time for a stage is the amount of time allowed for the stage to fulfill a demand request from an adjacent downstream stage.

As an example, consider a two stage supply chain where stage 2 supplies stage 1. Let us assume that stage 2 has a production lead time of 5 days and immediate access to its raw materials. If the service time for stage 2 is 0, then stage 2 plans to meet all of stage 1's demand directly from inventory; a service time of 0 means that demand in a period will be satisfied in the same period. Therefore, stage 2 must hold enough safety stock to cover its own replenishment time.

If the service time of stage 2 is 5, stage 2 will not hold any inventory but will instead produce product only when it receives a demand signal from stage 1. Stage 2 does not need to hold any finished inventory because the stage can procure and produce the part within its quoted service time.

If the service time is between 0 and 5, this implies that stage 2 must cover some, but not all, of the demand over its replenishment time. For example, if the service time for stage 2 is 2, then stage 2 needs to cover the demand variability over the 3 days (5 - 2 = 3) that it takes to produce additional product once an order is placed on the stage.

What the Software Does

The model determines the service times that will minimize the total cost of the safety stock. To do this, the model has to calculate the required safety stock as a function of the service times. Time permitting, we will describe in class how this is done.

Selecting "Optimize" under the tools menu determines the optimal service times as described in paper by Graves and Willems.

Alternatively, the user can change the service times in the stages directly; this corresponds to using the software like a calculator, or spreadsheet, to determine the costs of your own stocking strategy. This feature is very useful when trying to find near optimal solutions.

Finally, in each stage there is a check box titled "Bound service time." This feature allows the user to optimize the rest of the system given a bound (lower and/or upper) on the stage's service time. For instance you can use this feature to fix a service time at a stage by setting the lower bound equal to the upper bound; then you can optimize the rest of the supply chain with this service time fixed. Service times could be fixed based on management discretion or because the resulting solution is easier to implement.

Exercises

To get some familiarity with the model, please do the following exercises.

1.   Build the following supply chain depicted below (note – the figures below are from an earlier version of the software, so the 'look' has changed slightly):
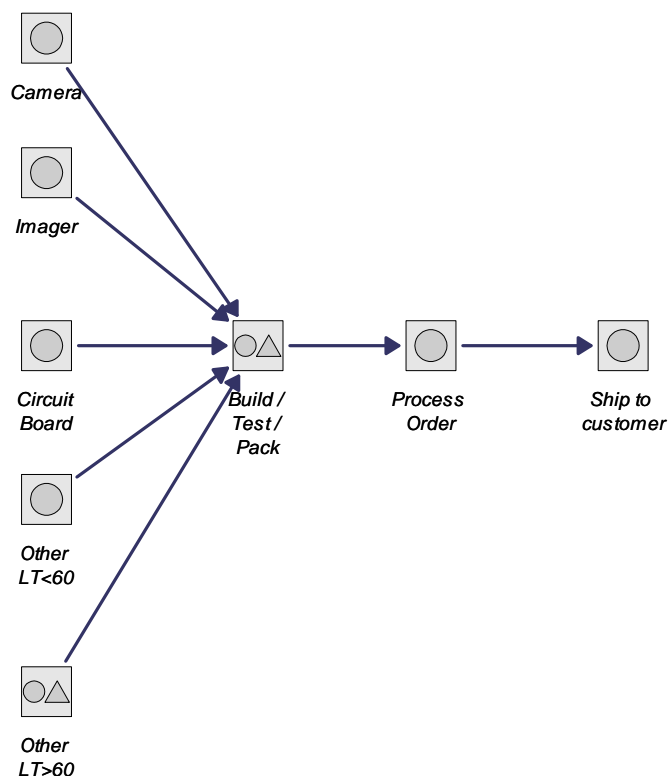
To save space (and allow you to check your work), the figure above depicts the optimal solution. The parameters of interest for the supply chain are:

Stage Name	Production time	Cost added

Camera	60	750
Imager	60	2000
Circuit Board	40	350
Other LT<60	60	150
Other LT>60	150	200
Build/test/pack	8	250
Process order	2	50
Ship to customer	3	0

The mean daily demand (entered at the Ship to customer stage) is 9 units and the standard deviation of demand is 3. The maximum service time to the customer is 5 days, meaning that the company has a maximum of 5 days for the product to arrive to the customer after the order has been placed. Assume the service level is .95 for all stages.

To gain a better understanding of the supply chain, the model depicts the production of a digital camera. The three key inputs are the camera, imager and circuit board; other parts are aggregated and included only to make the cost numbers come out correctly. "Other LT<60" is an aggregation of all components with lead times of less than 60 days; "Other LT>60" is an aggregation of all components with lead times greater than 60 days. The camera is a traditional (35 mm) camera. The circuit board and imager are produced internally; the imager digitizes the traditional camera input and the circuit board processes this picture. For the purposes of this model, the build/test/pack operations have been aggregated. At this stage, the back of the camera is taken off and the imager and circuit board are added. The digital camera is then tested to make sure there are no defects in the imager. Once this is complete, the order is processed through distribution and then the product is shipped UPS to the customer; in this model, customer refers to high end photo shops or computer retailers.

Four different departments are included in the digital-camera supply chain. The imaging group produces the imager, circuit board assembly produces the circuit board, final assembly performs the build/test/pack operations and the process order and ship stages are run by distribution. Finally, safety stock held at a stage corresponds to the functional department holding it. For example, if process order holds safety stock this means that the order can be shipped directly from distribution; i.e., if process order has safety stock then this safety stock is located at the distribution center. If build/test/pack holds safety stock, this means that fully tested digital cameras are being held at the assembly stock vault and are available as soon as distribution asks for product.

a. Optimize this system and review the outputs. Why is there no safety stock held at the process order stage?

b. After looking at the optimal solution, the consultant decides to examine a near optimal solution by setting the service time at the camera, imager and circuit board stages to zero. Why might the consultant recommend this?

c.	The company likes the near optimal solution but decides to set the service time at process order to zero. The company likes this because even though it has more inventory than the near optimal policy, it still has less inventory than their existing system; they also like the fact that distribution has inventory they can take right off the shelf. Note that in order to implement this change, you need to reduce the maximum service time to the customer from 5 days to 3 days. .

	Assuming this additional requirement that distribution must hold safety stock, would you recommend changing any other stage's stocking policy? What is the rationale for your recommendation?

2.	Now, starting from scratch, create the imager supply chain depicted below. This requires closing the current file and choosing new under the file menu.
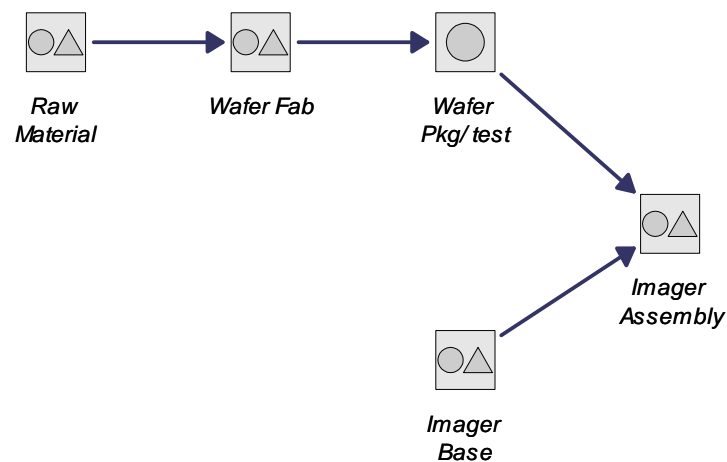
	The figure above denotes the current situation in which each stage stocks enough safety stock to decouple each stage from the rest of the supply chain. The parameters of interest for the supply chain are:

| Stage Name	Production time	Cost added
Raw Material	130	50
Wafer Fab	30	700
Wafer pkg/test	5	200
Imager base	50	250
Imager assembly	5	500

	As before, the mean demand (at imager assembly) is 9 units per day and the standard deviation is 3. The imager assembly's maximum service time is zero; i.e., they want to supply final assembly directly from inventory.

a.	Optimize the imager supply chain. Provide some intuition for the optimal solution.

b.   What happens to the optimal solution as you reduce the image base lead-time from 50 to 20?

The academic version of the software limits you to networks with at most ten nodes. For question 3, you need to make a small "fix" to the software that expands its capability to fifteen nodes. The instructions are on the last page of this assignment.

3.   Now, starting from either the model you created in problem 1 or problem 2, create the following integrated digital-camera supply chain:
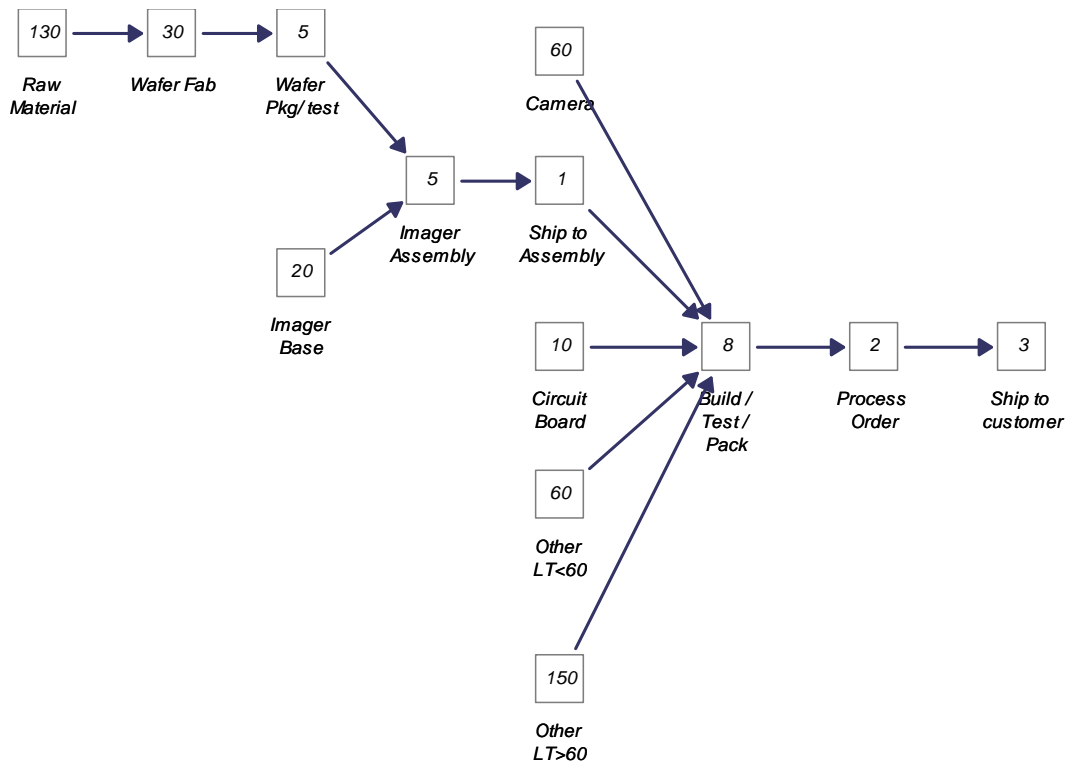
The figure above depicts the production lead-time at each stage. With the exception of the imager base, all of the inputs are still given in the tables from problems 1 and 2. For the imager base, assume the production lead-time is 20 days. As before, the maximum service time that ship to customer can quote is 5 days and the mean demand is 9 and the standard deviation is 3.

The ship to assembly stage acts to link the imager supply chain to the rest of the digital camera supply chain; holding safety stock at ship to assembly corresponds to holding safety stock at the assembly area. The shipment time takes one day and costs $25.

a.   Optimize the supply chain given the constraint that cameras and circuit boards must still quote a service time of zero. How does this solution compare to the combined solution inferred from 1b and 2b? Discuss the differences.

b.   How does the solution and its cost change as you reduce the maximum service time that is quoted by ship to customer? Consider reducing this from 5 days to three days? To one day? To Zero?

Instructions for expanding capability to 15 nodes

1. Install the regular PowerChain Inventory Academic setup first. Make a note of where you install the application.

2. After installation is complete, make sure you are not running PowerChain Inventory Academic.

3. Open the ZIP file, extract the file, and copy this file into the folder in which you installed PowerChain Inventory Academic. It should overwrite an earlier version of the same file; if not, then you have the wrong directory.

4. Launch PowerChain Inventory Academic and it should now enable up to 15 nodes (though the error message when you attempt to add a 16^{th} node will still say the limit is 10.)