DARDEN BUSINESS PUBLISHING
UNIVERSITY of VIRGINIA

UV7180
Aug. 31, 2016

Operations Management: An Introduction to Process Analysis

Firms exist to create and deliver value to customers, shareholders, employees, and society. Operations Management is focused on the means—the underlying processes—by which firms create and deliver value. A firm's processes transform inputs such as labor, capital, materials, and information into outputs of greater value, in the form of products and services. For example, in a bank loan–approval process, bank staff and computer systems process information about a loan applicant and return either an approval or a rejection of the loan application. In an airport security check-in process, security personnel use equipment to inspect passengers and their luggage and either deem them fit for boarding or detain them for further inspection. In a car assembly process, workers and equipment convert car components into assembled vehicles.

In this technical note, we illustrate several process-analysis fundamentals. To do this, we utilize a simple example of a Hawaiian-shirt production process. The note is divided into the following six sections:

1. Mapping a Process
2. Calculating Capacity and Identifying Bottlenecks
3. Load Balancing
4. Setup Times and Batches
5. Little's Law
6. Summary

Appendix 1: Process-Map Example

Appendix 2: Practice Problem

Appendix 3: Glossary

1. Mapping a Process

To manage or improve a process, it is important to first understand it. Almost always, the first step of any process analysis should be to map or draw out the existing process. Process mapping or process-flow diagramming is a visual tool that helps achieve this. Consider the following example of a Hawaiian-shirt production process. This process includes the following four steps:

Step 1 (cutting): A roll of fabric (with the desired shirt pattern) is gathered and laid out on the cutting table. The necessary pieces for a shirt (one torso, two sleeves, one collar, and one shirt-pocket) are then cut using a semiautomatic fabric cutter. The cut pieces for the shirt are then placed into a tote and staged before step 2.

Step 2 (sewing–base): When ready, the sewing machine operator at step 2 takes the next available tote and stitches the two sleeves and the pocket to the torso. The base shirt is then placed back into the tote (along with the remaining raw materials) and staged before step 3.

Step 3 (sewing–extras): The sewing machine operator at step 3 takes the next available tote and stitches the collar, two company logo patches (one for the back of the collar and one for the left shirt sleeve), and six buttons to the base shirt. The completed shirt is then placed back into the tote and staged before the pack-and-ship station.

Step 4 (pack and ship): The pack-and-ship worker takes the next available tote and places the completed shirt in a plastic bag. The bag is then packaged in a standard-size shipping box along with a customer receipt and return slip, both of which are printed out at the pack-and-ship station. The box is then sealed with packaging tape, and a shipping label is attached to the outside. Completed boxes are staged at the shipping door for pickup by a parcel carrier.

For simplicity, assume that there is one worker at each step of the process. **Figure 1** illustrates a high-level process map of the four-step process. If desired, additional details can be shown within each step of the process (see **Appendix 1** for a more detailed version of the production process).

Figure 1. Hawaiian-shirt production process.

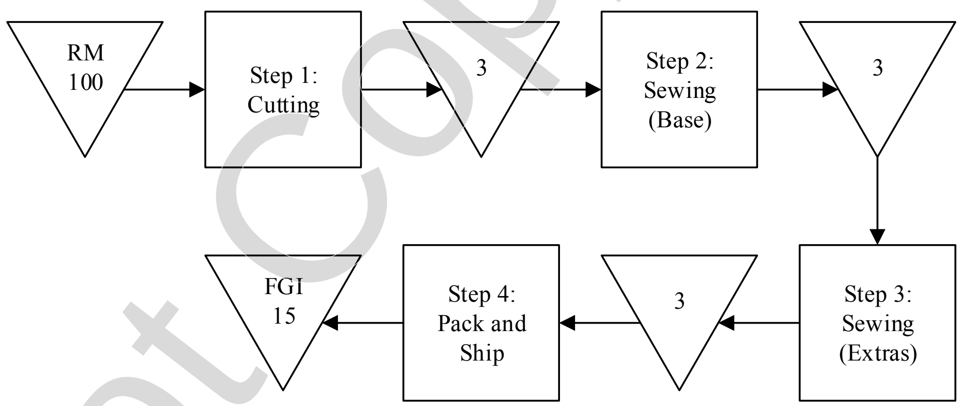

Source: All figures created by author.

In a process map, boxes are typically used to represent steps where work is done. Solid arrows show the flow of the product through the process. Inverted triangles are used to depict inventory between steps with the number illustrating how many units of inventory are in that location on average. This note will refer to three types of inventory:

Raw materials (RM): Materials that have not yet started production are referred to as raw materials. For example, the rolls of fabric, logo patches, and shipping boxes in the Hawaiian-shirt example.

Work in process (WIP): Any items that are currently in process, either at a process step or between steps, are referred to as work-in-process inventory or WIP. Often, WIP will either naturally build up between steps or be strategically placed there in order to decouple steps in a process. The area where WIP inventory is stored is often referred to as the buffer.

Finished goods inventory (FGI): Items that have been completely processed are referred to as finished goods inventory.

In **Figure 1**, the 100 shirts in inventory at the raw materials buffer (RM) represent 100 shirts' worth of raw materials inventory at this stage (these are not completed shirts). Notice that **Figure 1** does not include an inventory triangle for raw materials such as buttons or company logos. Although these can easily be incorporated, when analyzing a process it is often easier to focus on only the core components of what will eventually become the finished product.

A process can be drawn in more or less detail depending on the goal of the exercise. Additional components can be added where needed. For example, dotted lines can be used to depict the flow of information. In the Hawaiian-shirt example, a purchase order starts the production process at step 1. We could show a purchase order (i.e., information) being delivered to step 1 of the process. If a resource is shared across steps in a process, this action can be depicted by enclosing the relevant steps in a larger dotted box. For example, a third sewer and sewing machine could be added to the process, with this station sharing time between steps 2 and 3.

Finally, when analyzing a process, it is important to clearly define the process boundaries. For example, additional steps can be added to **Figure 1** either for collecting and taking customer orders to production (before step 1) or for the parcel-loading process for the shipping carrier (after step 4). To decide what activities to include in a process map, we should take into consideration what the managerial goals are for the analysis.

2. Calculating Capacity and Identifying Bottlenecks

Once the existing process has been mapped out, we can proceed to analyze it. Often, the first question we will want to address is "How much can the existing process produce?" or "How many customers can the existing process service?"

For our analysis, we will assume a steady state regarding the workload of the process. Hence there is not a ramp-up period in which work must fill the system. Instead, at the beginning of the workday, the assumption is that a worker can arrive at his or her station and immediately continue working on whatever unit he or she was working on at the end of the previous day. This is a valid assumption for many production settings (e.g., unfinished cars on a production line remain on the line overnight to be completed the next day).

Suppose that time studies were performed on the Hawaiian-shirt process with the following minutes-per-shirt values found for each step (**Table 1**).

Table 1. Run times per step.

Step	Description	Run Time Minutes/Shirt
1	Cutting	12.5
2	Sewing (base)	10
3	Sewing (extras)	15
4	Pack/ship	5

Source: All tables created by author.

The values listed in the third column of **Table 1** represent the average run times for each step.

Run time: The time it takes to process an item at a process step.

Notice that run time is measured in time per unit, not just time. When analyzing a process, keeping track of the units of measure is valuable for verifying that calculations are performed correctly. The run time of a step is independent of any setup times that may be involved in the process. It can also be computed for a batch of items (e.g., the run time for a batch of Hawaiian shirts being processed all at once). In section 4, we discuss setup times and batches in more detail.

Given the run times in **Table 1**, the capacity for each step can then be calculated. We first examine each step in isolation. Although the steps of the process are dependent on one another, as will be shown, analyzing each step in isolation first offers a simplified way to determine the capacity of the process without having to get into the details regarding the scheduling of work.

Capacity: The maximum number of items that can be processed by a resource in a given time period. Workers, production equipment, testing equipment, and computers are all examples of resources.

Why is capacity important? Capacity provides us with an upper bound on the amount of units we can produce or customers we can serve per unit of time. Knowing the capacity of a process helps to determine the extent to which demand can be met. Observe that we use the term "resource" and not worker or labor. The capacity of a process is not strictly constrained by the number of personnel available. Oftentimes, equipment is the constrained resource in a process. Take, for example, a restaurant kitchen where both the cooks and the equipment (e.g., the ovens and grills) have capacity limitations.

Based on the definition of capacity, we can determine the capacity for each step in the process. The run time for cutting (step 1) is 12.5 min./shirt. Assume once again that there is 1 worker per step and that the operation runs 8 hours per day, 5 days per week. Given this, the capacity for step 1 is as follows:

$$Capacity = \frac{5\,\frac{days}{week} \times 8\,\frac{hours}{day} \times 60\,\frac{mins}{hour}}{12.5\,\frac{mins}{shirt}} = \frac{2{,}400\,\frac{mins}{week}}{12.5\,\frac{mins}{shirt}} = 192\,\frac{shirts}{week}$$

Applying a similar method, we calculate the capacity for each independent step as follows in **Table 2**:

Table 2. Capacity per step.

Step	Description	Run Time Minutes/Shirt	Capacity Shirts/Week
1	Cutting	12.5	192
2	Sewing (base)	10	240
3	Sewing (extras)	15	160
4	Pack/ship	5	480

It is important to remember that the capacity calculations above are *average* values that are meant to be used as a guide for decision making. For any given week and any given step, a resource may be capable of producing more or less than its average capacity. Some steps may exhibit considerable variability in their run times. Although working with averages ignores this variability, it provides us with a straightforward way in which to begin analyzing a process.

To calculate the capacity of a step, we divide the total time available (i.e., hours worked in a week) by the run time. If we were to write out a general equation for capacity, it would be as follows:

$$Capacity = \frac{Total\ Available\ Time}{Time\ Required\ to\ Process\ a\ Unit\ (or\ Batch)} \quad (1)$$

Equation 1 provides us with a good starting point for examining ways to increase the capacity of a process. Just by simple examination of the formula, two ways to increase capacity become apparent. The first is to increase the numerator (i.e., Total Available Time). How can we increase the time available for a process? Common methods include increasing the number of hours the process operates per day or per week (e.g., adding a second shift), hiring more workers, or adding more machines. Notice that while these are often the solutions that immediately come to mind when looking for ways to increase capacity, they can also be very costly.

The alternative way to increase capacity is to shrink the denominator in **Equation 1**. This is often a more difficult but less expensive task. To reduce the time required to process a unit, we can apply basic Lean principles to look for and eliminate waste in the process.[1] For example, redesigning a workplace layout can help to eliminate unnecessary movement or transport in a production process.

Once we have calculated the capacity for each step in the process, we can then determine the capacity for the overall process. To do this, we must first find the bottleneck of the process.

Bottleneck: The resource that limits the production or service delivery of a process.

The capacity of a process is the capacity of its limiting resource. To determine capacity, one needs to find the limiting resource or slowest step, known as the bottleneck. The bottleneck can be found by calculating the capacity of every step or resource and finding the step with the lowest capacity. For example, in the Hawaiian-shirt process, step 3 (sewing extras) has the lowest capacity. This means that step 3 is the bottleneck for the process. Because step 3 can produce at most only 160 shirts per week, the overall process capacity is also limited to producing 160 shirts per week. The process as a whole cannot operate at a faster rate than that of step 3, even though every other step in the process has the capacity to operate at a faster rate. We emphasize that this does not necessarily mean that the process will produce 160 shirts per week; it simply means that it *can* produce that many shirts per week. The number of actual shirts produced per week will be determined by demand and customer orders.

We have identified the bottleneck as step 3; however, we have not stated the actual resource that is constrained. In the current process both the sewer and the sewing machine at step 3 would be considered the bottleneck since there is a one-to-one relationship between the two, but this is not always the case. Resources are often shared in a process. For example, imagine a very extreme case in which a single sewing machine was shared between steps 2 and 3. Although the sewers would still have capacities of 240 shirts per week (step 2) and 160 shirts per week (step 3), the single shared sewing machine would have a much lower capacity than either of these steps. A single sewing machine, available for 5 days a week, 8 hours a day, can perform steps 2 and 3 in 25 min. per shirt resulting in a capacity of:

[1] Rebecca Goldberg and Elliot N. Weiss, *The Lean Anthology: A Practical Primer in Continual Improvement* (New York: Productivity Press, 2014).

$$Capacity = \frac{5\,\frac{days}{week} \times 8\,\frac{hours}{day} \times 60\,\frac{mins}{hour}}{10\,\frac{mins}{shirt} + 15\,\frac{mins}{shirt}} = \frac{2400\,\frac{mins}{week}}{25\,\frac{mins}{shirt}} = 96\,\frac{shirts}{week}$$

In this case, the shared resource (i.e., the single sewing machine) would be the bottleneck of the process, and the new process capacity would then be 96 shirts per week.

Given the capacity of the process, we can then measure the extent to which the assets and resources are utilized.

Capacity utilization: The ratio of the amount of a resource used to the amount available in that time period.

Capacity utilization provides us with a measure of how much of the capacity of an operation is being used. Suppose that the Hawaiian-shirt process is operating at capacity (i.e., producing 160 shirts per week). Recall, however, that step 1 can operate at a capacity of 192 shirts per week. This implies that the (average) capacity utilization at step 1 is as follows:

$$Capacity\ Utilization = \frac{160\,\frac{shirts}{week}}{192\,\frac{shirts}{week}} = 83.3\%$$

For step 2, the capacity utilization is 66.6%, step 3 is 100% (i.e., the bottleneck step), and step 4 is 33.3%. Calculating the utilization of each step in a process provides an illustrative way to examine where there may be excess slack in the system.

Capacity is a rate and is always measured in units per time. In a practical sense, when managing a process it may help to measure the process in time per unit instead.

Cycle time: The average amount of time that elapses between the completion of successive items at a stage, assuming that the process is operating at capacity. The cycle time at a process step is defined as the inverse of the capacity at that step.

For example, consider step 1 again. The cycle time can be calculated as follows:

$$Cycle\ Time = \frac{1}{Capacity} = \frac{1}{192\,\frac{shirts}{week}} = \frac{1}{192}\frac{weeks}{shirt}\quad OR\quad 12.5\,\frac{mins}{shirt}$$

If we were to calculate the cycle times for steps 2 through 4, not surprisingly, we would find that the cycle times for each step of the Hawaiian-shirt process exactly match the original run times. This is because we have not included any setup times at any of the individual steps. Section 4 demonstrates how incorporating a setup time alters the cycle time. Given the cycle times for each step in the process, the cycle time of the *overall process* is then 15 min. per shirt (i.e., the cycle time of the bottleneck step).

Cycle time provides us with a different way to think about capacity and a different way to visualize the process. For example, at a car assembly line that operates nonstop on an 8-hour shift and produces 480 cars per shift, the cycle time is 1/480 shifts per car, or 1 minute per car (an 8-hour shift lasts 8 hr./day × 60 min./hr. = 480 min./day, thus 480 min./day to produce 480 cars/day gives a cycle time of 1 min. per car). This means that if we were to stand at the end of the production line, we should see, on average, a fully assembled car come off the line every minute.

Because cycle time is the inverse of capacity, if the cycle time of a process is known, then the capacity can be calculated (and vice versa). Having said this, one *does not need* to calculate cycle time in order to calculate process capacity. Cycle time simply provides an *alternative* way to think about capacity.

Finally, revisiting the process design, the question may arise of why steps 2 and 3 in **Figure 1** need to be performed sequentially, with two sewers splitting the tasks for sewing a shirt. Realistically, we could combine the two steps, where each sewer sews a complete shirt. To determine the capacity for this alternate process configuration, we assume that a single sewer can sew a shirt in 25 min./shirt (adding the run times for steps 2 and 3). We then calculate the capacity for the combined steps 2 and 3 (with two sewers) as follows:

$$Capacity = \frac{2\ sewers \times \left(5\,\frac{days}{week\ per\ sewer} \times 8\,\frac{hours}{day} \times 60\,\frac{mins}{hour}\right)}{25\,\frac{mins}{shirt}} = \frac{4800\,\frac{mins}{week}}{25\,\frac{mins}{shirt}} = 192\,\frac{shirts}{week}$$

The numerator now captures the fact that the two sewers are available for a total of 2 × 2,400 min./week to sew complete shirts. Observe that the capacity of 192 shirts/week for the combined steps 2 and 3 is higher than our previously calculated capacity of 160 shirts/week for step 3 (the bottleneck step). By combining steps 2 and 3 we eliminate idle time that the sewer at step 2 incurs in the original process configuration (since he or she can only process as fast as the rate of step 3). The overall capacity for the new process configuration is 192 shirts/week, which is equal to the capacity of both step 1 and the new combined steps 2 and 3. Given this, an important question to ask is why do we split the tasks at steps 2 and 3 between sewers (as shown in **Figure 1**)? Why not always have workers in a process perform all steps—for example, sewing a complete shirt or making an entire sandwich at a sub-sandwich counter? What are the advantages of separating tasks between workers?

In sections 1 and 2 we have mapped out the existing process and determined how many shirts per week it can produce. Next, in section 3, we will examine one way to increase the capacity of the process. Going forward, we will continue to assume that steps 2 and 3 are performed sequentially by two different sewers (as shown in **Figure 1**).

3. Load Balancing

One way to potentially increase the capacity of a process is to shift tasks away from the bottleneck. This section illustrates how load balancing and operator-loading charts can be used to increase the capacity of the Hawaiian-shirt production process.

To load balance the steps of the process, we first need to determine a limit to constrain how much time per shirt we can spend at each step in the process. Given a time frame and a target number of units to produce (per unit of time), we can divide the former by the latter and get the *maximum time available* to make a single unit and still meet the target production quantity. We refer to this value as *takt* time, from the German term for "pace" or "rhythm."

$$Takt\ time = \frac{Amount\ of\ time\ available}{Units\ demanded\ per\ unit\ time} \tag{2}$$

It's important to emphasize the difference between takt time and cycle time. Takt time refers to the target time per unit as derived from customer demand. Cycle time refers to the actual time per unit that can be achieved on the production floor. Takt time provides a metric for evaluating and improving cycle times. If the cycle time for any step of a process exceeds the takt time, then the process cannot fully meet demand and that cycle time must be reduced. Otherwise, the step will not be able to keep pace with demand.

Assume that the demand for shirts per week is 172 shirts per week. The Hawaiian-shirt production process runs 8 hours a day, 5 days a week, which means there is a total of 2,400 min. per week to meet demand. Applying **Equation 2**, the takt time can be calculated as follows:

$$\text{Takt time} = \frac{2{,}400\ \frac{mins}{week}}{172\ \frac{shirts}{week}} = 13.96\ \frac{mins}{shirt}$$

For the process to meet demand, no step can have a cycle time greater than 13.96 min. per shirt. To think about this another way, if we were to stand at the end of the production line, we would need to see a shirt come off the line every 13.96 minutes (or less), on average, to ensure the process is fully meeting demand.

To create an operator-loading chart, in **Figure 2** we graph the cycle times for each step in the process along with the takt time.

Figure 2. Operator-loading chart (Hawaiian-shirt process).

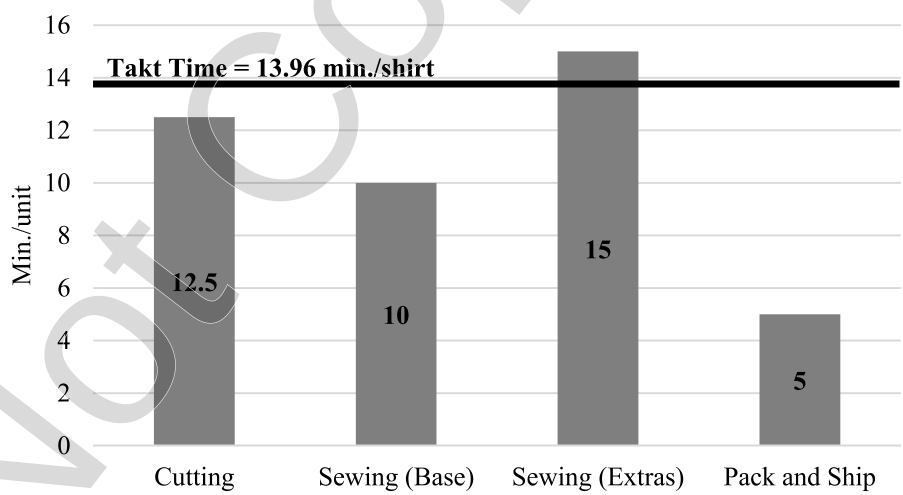

As shown in **Figure 2**, step 3, sewing (extras), is not only the bottleneck, but its cycle time also exceeds the takt time. Therefore, meeting a demand of 172 shirts per week is not possible with the current process configuration. To increase the capacity of the process, we could increase the time available at the sewing (extras) step by either adding a second sewing machine and operator or by increasing the number of hours the current sewer works. A second option is to reduce the time required to process a unit at this step by either (i) eliminating unnecessary tasks, (ii) improving the time for specific tasks, or (iii) rebalancing tasks between steps. We demonstrate (iii) below.

Realistically, given the difference in process, tasks probably cannot be shared between sewing (extras) and cutting (step 1) or pack and ship (step 4). Instead, we look for ways to share tasks between step 3 and step 2, sewing (base). **Figure 3** replicates **Figure 2** but also includes additional detail on the average cycle times (in minutes per task) for the tasks involved in steps 2 and 3.

Figure 3. Detailed operator-loading chart.

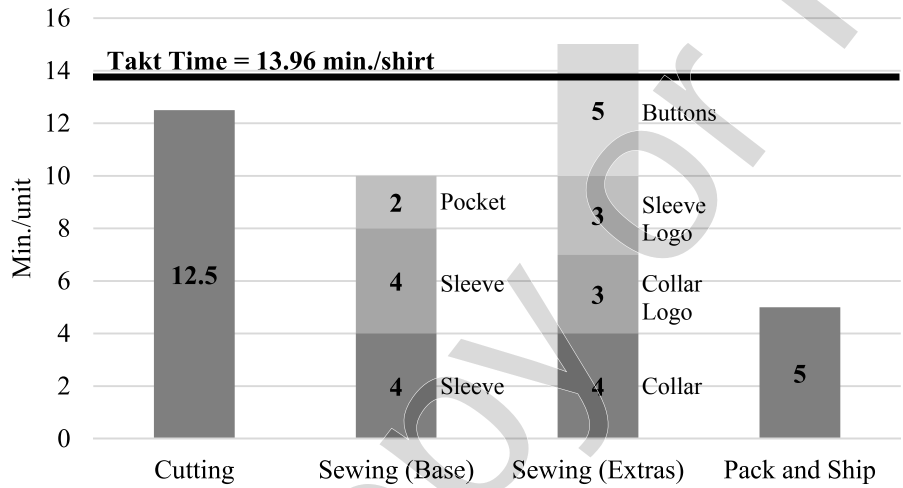

The goal is to reduce the cycle time of step 3 below the takt time of 13.96 min. per shirt by balancing the workload between steps 2 and 3. There are in fact multiple ways to rearrange tasks between steps 2 and 3. As shown in **Figure 3**, sewing the pocket onto the torso takes 2 min. per shirt (step 2), whereas sewing the collar onto the torso takes 4 min. per shirt (step 3). Exchanging these two tasks reduces the cycle time for the bottleneck step, sewing (extras), to 13 min. per shirt. **Figure 4** demonstrates a balanced operator-loading chart, with step 2 taking a total of 12 min. per shirt and step 3 taking 13 min. per shirt. Because step 3's cycle time is below the takt time of 13.96 min. per shirt, the process can now fully meet demand.

Figure 4. Balanced operator-loading chart.

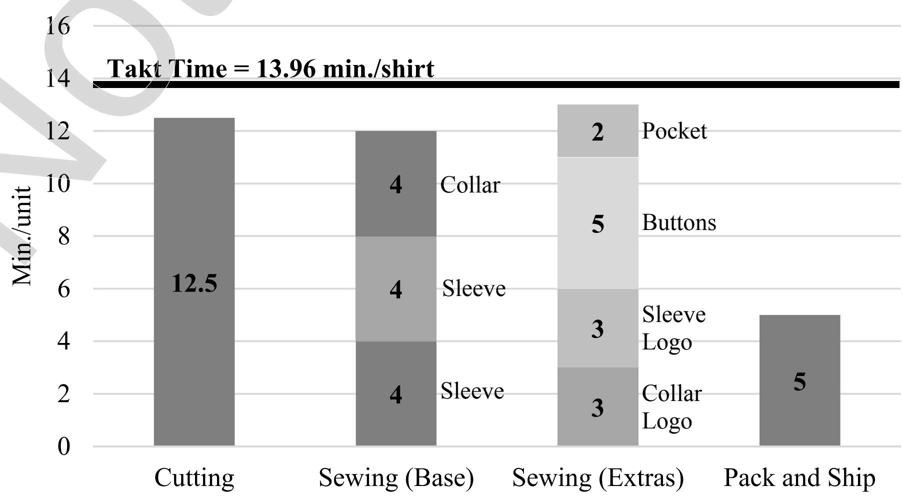

We can further verify that the updated process can now meet demand by recalculating the capacity. The capacity at step 2 reduces to 200 shirts per week (from 240 shirts per week), while the capacity at step 3 increases to 184.6 or approximately 185 shirts per week (from 160 shirts per week). Although the capacity at step 2 has decreased, realize that this is secondary since (i) step 2 is not the bottleneck, and (ii) the overall process capacity has increased from 160 shirts per week to 185 shirts per week. Also observe that there are alternate ways to rebalance the tasks at steps 2 and 3. For example, switching the task of sewing the sleeve logo (3 min. per shirt) to step 2 shifts the bottleneck to step 2 but results in the same overall process capacity. Realistically, moving this task to step 2 might make sense, since the sewer at step 2 is already attaching the sleeves to the torso.

Operator-loading charts can be used to not only examine ways to rebalance tasks between steps, but also to look for ways to reduce the number of resources required in a process. For example, consider a sub-sandwich restaurant where the tasks involved in making a sandwich are shared between multiple workers. An operator-loading chart can be used to determine the minimum number of workers required on the sandwich-making line in order to meet customer demand.

Operator-loading charts allow us to visualize opportunities for improvement and easily identify bottlenecks. Based on the improvements derived from an operator-loading chart, the design of a process can be revisited, and resources and tasks can be reallocated when necessary.

4. Setup Times and Batches

So far we have assumed that Hawaiian shirts are being processed at each step one at a time. Also, we have ignored the potential setup time required for the steps in the process in **Figure 1**. Next, we introduce the concepts of batch processing and setup times.

Batch: A group of items or orders that is processed all at one time. The batch size is then the number of items or orders that are processed as a batch.

In this section, we will assume that Hawaiian shirts are processed in batches of five shirts per batch. This means that a batch of five shirts traverse the process together in a tote, with each shirt being processed at a step before the tote (and all five shirts) moves on to the next step. Often in a production process successive batches will be different from one another. For example, in the Hawaiian-shirt process, batches may be differentiated by fabric type. Note that we will assume shirts are processed as batches throughout the process and that the run times remain the same. But sometimes items may be batched together only in certain stages of a process. Also, cases can occur in which batching makes the task for successive items in a batch easier due to repetition.

Setup time: The amount of time it takes to set up or prepare before processing an item or batch of items at a process step.

Examples of setup time include tasks such as calibrating the settings on a sewing machine, setting up a large-scale printer for a new production run, and, in general, preparing the workplace for a different product or service. Since setup time is often incurred in order to change from working on one type of product or service to another, it is also known as changeover time. We consider setup time non-value-add time because products are not being produced or customers are not being serviced while the setup occurs. Thus, during setups, we are depleting the time available while not producing anything to show for it. This is often why we look for ways to either reduce setup times or perform setups concurrently with production (if possible). Typically, the time it takes to set up does not change based on batch size.

In the Hawaiian-shirt example, we have ignored setup times. Yet from a practical standpoint, some of the steps in the process likely require preparation time before shirts can be worked on. For example, to cut the necessary pieces for sewing in step 1 requires time to retrieve and lay out the necessary fabric and calibrate the tool being used for cutting (if, for example, different-size shirts were being cut). It is important to incorporate this non-value-add time into the analysis since it can drastically impact the capacity of the process.

In the Hawaiian-shirt example, assume that there is a setup time of 22 min. per batch for step 1. This setup time represents the time to select and lay out the fabric as well as calibrate the necessary tools for cutting.[2] Based on the new batch size and setup time, we recalculate the capacity for step 1 as follows:

$$Capacity = \frac{5\,\frac{days}{week} \times 8\,\frac{hours}{day} \times 60\,\frac{mins}{hour}}{22\,\frac{mins}{batch} + \left(5\,\frac{shirts}{batch} \times 12.5\,\frac{mins}{shirt}\right)} = \frac{2400\,\frac{mins}{week}}{84.5\,\frac{mins}{batch}} = 28.4\,\frac{batches}{week}$$

The cycle time for a batch of five shirts at step 1 is 84.5 min. per batch, which includes 62.5 min. per batch of run time and 22 min. per batch of setup time. Converting batches per week to shirts per week, we find our updated capacity for step 1 is:

$$28.4\,\frac{batches}{week} \times 5\,\frac{shirts}{batch} = 142\,\frac{shirts}{week}$$

This represents a decrease from the original step 1 capacity of 192 shirts per week (see section 2). Why did capacity decrease? It decreased because there is now additional time during the production process that must be dedicated to setting up the cutting step for each batch. This setup time takes away from the time available and thus decreases the capacity for step 1.

The updated cycle time in minutes per shirt for step 1 can then be found as follows:

$$Cycle\ Time = \frac{1}{Capacity} = \frac{1}{142\,\frac{shirts}{week}} \times 2{,}400\,\frac{mins}{week} = 16.9\,\frac{mins}{shirt}$$

The updated capacity and cycle time for each step are listed in **Table 3** (note that for steps 2 and 3, we use the new run times and capacities calculated in section 3):

Table 3. Updated capacity and cycle times.

Step	Description	Run Time Min./Shirt	Setup Time Min./Batch	Capacity Shirts/Week	Cycle Time Min./Shirt
1	Cutting	12.5	22	142	16.9
2	Sewing (base)	12	0	200	12
3	Sewing (extras)	13	0	185	13
4	Pack/Ship	5	0	480	5

The bottleneck of the process has now shifted from step 3 to step 1. Furthermore, notice that the run time and cycle time for step 1 are no longer the same. Recall that the cycle time represents the inverse of capacity.

[2] Realistically, the sewing steps may also require setup times; to simplify the analysis we will continue to assume that there is no setup time for steps 2 and 3.

Therefore, since the capacity calculation now incorporates a setup time for each batch, the cycle time has increased for step 1. This new cycle time allocates the setup time across each unit of the batch. This is why, as previously stated in section 2, run time does not necessarily always equal cycle time.

For steps 2 through 4, the capacity and cycle times did not change. Although shirts are being processed in batches instead of one at a time, the run time for processing five shirts one at a time is the same as the run time for processing a batch of five shirts.

The choice of batch size can have a surprisingly large impact on the capacity of a process. An interesting question to consider is how does changing the batch size impact the Hawaiian-shirt process? For example, in the Hawaiian-shirt process, if the batch size were increased, would this increase or decrease capacity? What are the potential benefits and risks of having a larger or smaller batch size?

In section 5, to simplify the calculations, we will not use batches or setup times but, instead, revert to the assumptions from sections 2 and 3 (no setup times and processing shirts one at a time at each step).

5. Little's Law

Up to this point, we have focused on the question: "How much can the process produce?" However, we have not addressed how responsive the process can be. Identifying and reducing how long it takes an item or a customer to pass through a process leads to better customer service and thus directly impacts revenues by either increasing market share or allowing a firm to charge a premium for the better service provided. This section addresses the question: "How long does it take the process to produce an item?" Below we define two important terms for answering this question: throughput rate and throughput time.

Throughput rate: The rate at which units flow through a process. It represents the *actual* rate at which a process or system generates its products or serves its customers. The unit of measure for throughput rate is units per time (e.g., shirts per week or cars per month).

Notice that there is a natural relationship between capacity and throughput rate. In the Hawaiian-shirt example, the capacity was calculated as 185 shirts per week (assuming there are no batches or setup times and that steps 2 and 3 are balanced). This represents the maximum number of shirts we *can* produce per week; however, this is not necessarily how much we will produce. We may not (and likely will not) operate at capacity all the time. Instead, depending on demand, we may produce fewer than 185 shirts per week. The *actual* rate of production is the throughput rate. In this regard, capacity represents a ceiling or upper limit for throughput rate.

Our goal is to align the throughput rate with demand to ensure that customer needs are met and to prevent us from overproducing. For example, if we produced at capacity each week (i.e., 185 shirts per week) but then only received a demand of 120 shirts per week, eventually we would run out of storage space and working capital due to all the excess shirts we would have on hand.

Throughput time: The time it takes for a specific item, job, or order to go through the entire process. In manufacturing settings, throughput time is also known as manufacturing lead time. Throughput time is measured in units of time (e.g., seconds, minutes, and so forth).

To visualize throughput time, imagine we were at step 1 in the Hawaiian-shirt process and we started a stopwatch when the next shirt began processing. If we followed that shirt throughout the process and stopped the stopwatch when it completed step 4, we would find that the time it took the shirt to traverse the entire process represents the throughput time. Observe that throughput time depends on where the boundaries of the process are drawn. We could just as easily start the stopwatch once the raw materials inventory for the shirt entered the raw materials buffer and stop it when the finished shirt left the finished goods inventory. The choice of boundary depends entirely on what we are interested in measuring or analyzing.

Throughput time provides a measure of how much time on average it takes to produce a specific item. In this respect, throughput time is important because it provides a sense of how responsive a process is to customer demand. A very common mistake is to assume that throughput time can be determined by simply adding up the cycle times or run times of each process step. But this is incorrect as it does not capture the waiting time that often occurs in a system due to inventory buildup or the overall throughput rate of the process. Instead, to determine throughput time, we must apply Little's Law.

For a system in equilibrium, Little's Law[3] governs the relationship between the average rate of flow through a system (i.e., average throughput rate), the average amount of inventory in the system, and the average amount of time that a unit spends in the system (i.e., average throughput time). For a system with defined boundaries that is in equilibrium,[4] let:

I = Average amount of inventory (i.e., the average number of units) within the system

R = Average throughput rate through the system

T = Average throughput time for a unit to pass through the system.

Little's Law states that the relationship between I, R, and T is given by **Equation 3**:

$$I = R \times T \tag{3}$$

The inventory is defined as the average number of units or customers that are within the boundaries of the process at any point in time. Like throughput time, inventory depends on how the process boundaries are defined. The inventory in a process will fluctuate over time; I can be thought of as a snapshot of the inventory on the production floor.

To visualize Little's Law, it's often easiest to think of customers in a line. Take, for example, a single line at a grocery store checkout. Customers are processed at an average throughput rate of R = 32 customers per hr. On average, there are I = 4 customers in line (including in service). Therefore, the average throughput time is $T = I/R$ = 4 customers/32 customers per hr. = 0.125 hr. or 7.5 min. This means that a customer arriving in the line should expect to wait, on average, 7.5 minutes to be processed and out the door.

Notice that the units for throughput time in the grocery store example are hours whereas the units for throughput rate are customer per hour. A common mistake that students often make is to not recognize the difference in units. Consider the following example: Let's say we were taking a trip and wanted to determine how far we could drive over the next 3 hours. If we were to drive 60 miles per hour for 3 hours, we would travel 180 miles. From this simple calculation, we can recognize the three elements of Little's Law. The inventory is 180 miles; the throughput *rate* is 60 *miles per hour*; and the throughput *time* is 3 *hours*. Hence, applying

[3] J. D. C. Little, "A Proof of the Queuing Formula L = λ W," *Operations Research* 9 (1961): 383–87.

[4] For a system to be in equilibrium (i.e., a state of balance), the average rate at which units enter into the system must equal the average rate at which units leave the system. Thus, in the Hawaiian-shirt example, for the process to be in equilibrium, units must flow into and out of the system at a throughput rate less than or equal to 185 shirts per week (i.e., the calculated process capacity from section 3).

$I = R \times T$, we find 180 miles = 60 miles per hr. × 3 hr. The units of measure for throughput time and throughput rate are then hours and miles per hour—two very distinct and different units of measure.

In the Hawaiian-shirt example, assume that the average throughput rate is $R = 120$ shirts per week and that there is one unit of inventory in process, on average, at each step of the process and three units of inventory between each of the successive processes. To determine the average throughput time, T, from the start of step 1 to the end of step 4, we first need to add the inventory at each step of the process and in each buffer between steps 1 through 4. Revisiting **Figure 1**, we find that this gives us $I = 1 + 3 + 1 + 3 + 1 + 3 + 1 = 13$ shirts of WIP inventory. Applying Little's Law, we find:

$$T = \frac{I}{R} = \frac{13\ shirts}{120\,\frac{shirts}{week}} = 0.11\ weeks\ \ OR\ \ 4.3\ hours$$

This means that the next shirt to start production should be completed in 4.3 hours (on average). Notice that 4.3 hours is much longer than if we were to add just the cycle times of each step in the process. This is because using Little's Law to calculate T takes into account the amount of inventory in the system and the rate at which items are being processed through the system. A shirt ready to start production must wait for all the inventory ahead of it in line to be processed before it can be processed. In addition, although a step such as pack and ship *can* complete 480 shirts per week, it will not because the overall system is constrained by the rate of demand (i.e., the throughput rate).

The boundaries of the process can be redefined to determine the throughput time for a shirt to traverse through raw materials and finished goods. If we included the raw materials and finished goods inventories, the new $I = 128$ shirts. Recalculating the throughput time, we find that $T = I/R = 128$ shirts/120 shirts per week = 1.1 weeks or 42.7 hr. That is, it takes, on average, 42.7 hours for a shirt to go through the entire process from entering the raw materials inventory to leaving the finished goods inventory. It is important to note that this calculation assumes that the inventory enters and leaves raw materials and finished goods on a first-in, first-out basis (which may or may not be the case depending on how the raw materials and finished goods inventories are managed).

Little's Law can also be used to determine the average amount of inventory in the system. For example, assume that the inventory values shown in **Figure 1** are not known. Instead, sample data has been collected and the average throughput time, T, for the process from the start of step 1 to the end of step 4 was found to be 20 hours or 0.5 weeks. Based on this, we can calculate $I = R \times T = 120$ shirts per week × 0.5 weeks = 60 shirts. Thus, on average, we would expect to find 60 shirts in process at any one time (and between the start of step 1 and the end of step 4).

6. Summary

In this note, we have introduced some basic concepts used to analyze a process. These include process mapping, capacity analysis, operator-loading charts, and Little's Law. By better understanding the current and potential capabilities of a process, we can then incorporate this information into our decision making. It is important to recognize the types of managerial questions that we should be able to address after having read through this technical note. They include but are not limited to the following:

- How much can the process produce? How many customers can the process serve? How well are resources being utilized?

- What levers are available for increasing capacity? Should we increase the time available for production (e.g., adding a second shift or purchasing new equipment) or improve the existing process (e.g., eliminating unnecessary process steps or load balancing steps in the process)?
- How long does it take the process to produce an item? How long does it take the process to serve a customer?
- What levers are available for reducing throughput time?

We have made some simplifying assumptions in this note that deserve further discussion. First, we ignore rework. If, for example, a worker at the pack-and-ship station found a blemish or poorly attached logo on a shirt, then this shirt would be returned to the sewing step. This rework would affect the throughput time for that item and also impact the overall capacity. If a large number of items required rework, then the capacity of the process would decrease. We also ignore machine/operator downtimes, which would have similar effects on capacity.

Another consideration is that the calculations assume a steady stream of work throughout each day. In a manufacturing or production environment, the stream of incoming work can often be controlled. But in a service environment, a process is often subject to random arrivals (e.g., customers arriving at a restaurant). In addition, since orders may be customer specific, processing times may be highly variable. Although we can still perform calculations using averages, if arrivals and processing times are highly variable, then applying basic queueing-analysis techniques may be a better solution than assuming averages.[5]

[5] See Elliot N. Weiss, *Managing Queues*, (Boston: Harvard Business Publishing, 2013).

Appendix 1

Operations Management: An Introduction to Process Analysis

Hawaiian-Shirt Production Process (detailed example)

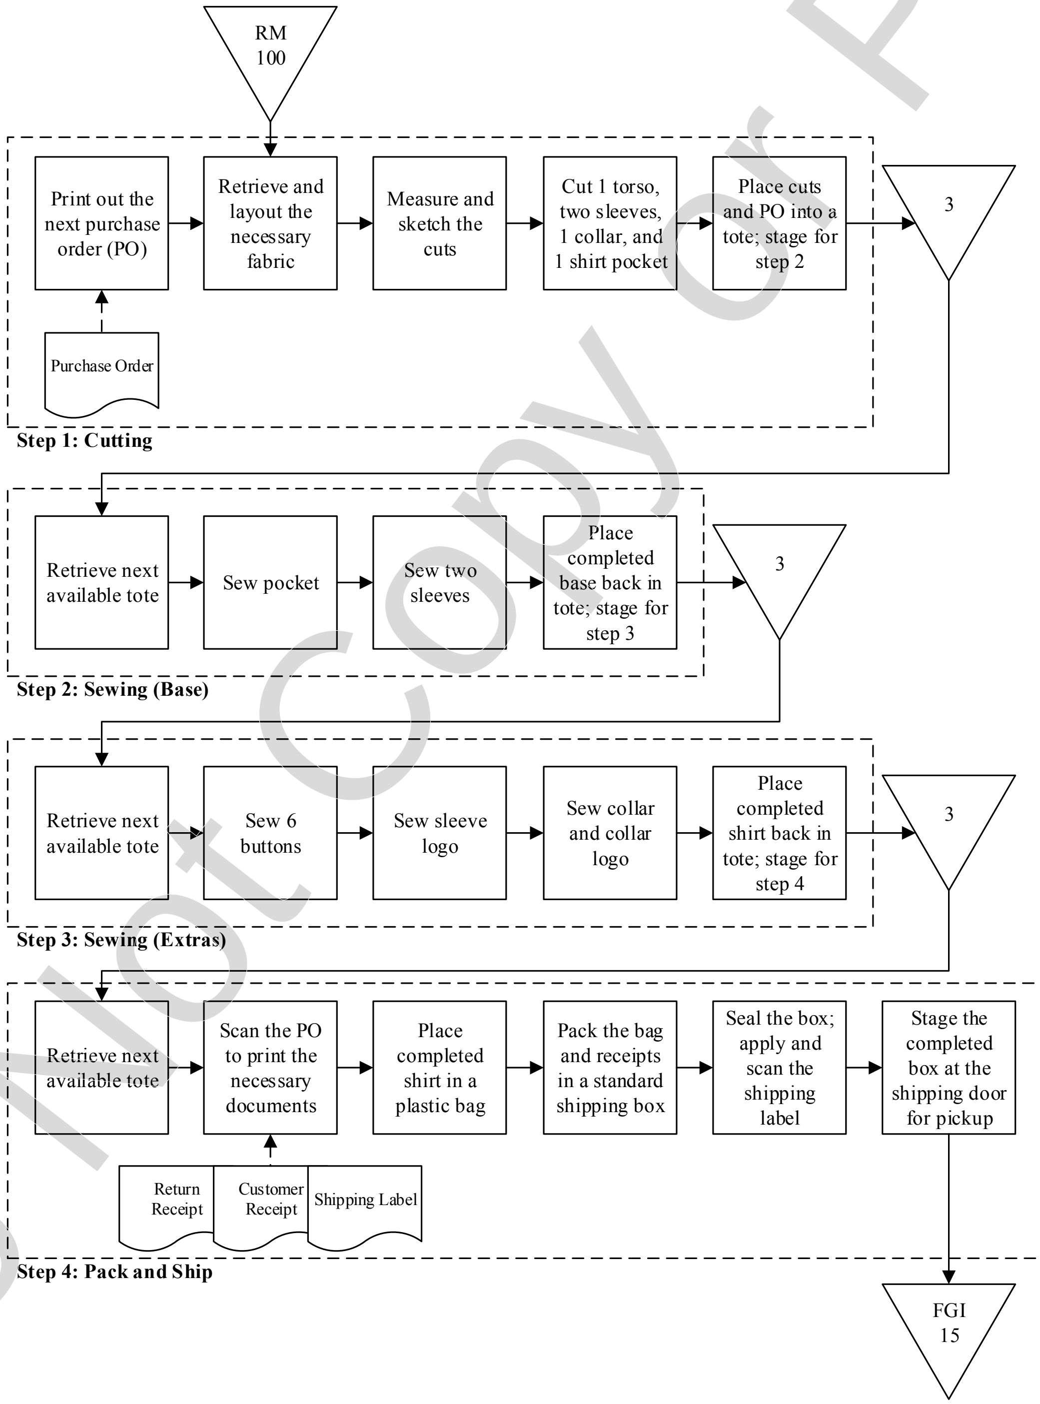

Appendix 2

Operations Management: An Introduction to Process Analysis

Practice Problem

Consider the following process to manufacture bunk beds. Assume that the beds are made in batches of 4, and that there are 40 hours available per week for production. Each bed goes through the following nine steps.

Step	Setup Time Minutes per Setup	Run Time Minutes per Bunk Bed
1. Cut	15	37.5
2. Plane	10	38.4
3. Router	10	48
4. Drill	10	28
5. Sand	25	24
6. Stain	15	12
7. Polyurethane	45	90
8. Kit assembly	0	40
9. Cleating	0	15

1. Assume that there is only 1 worker in total for building bunk beds (i.e., 1 worker performs each step). What is the capacity in beds per week for the process?

For questions 2 through 7, assume that there is 1 worker at each step in the process and that, on average, there is 1 batch of inventory before each step and 1 batch in process at each step.

2. What is the capacity in beds per week for each step in the process? Which step is the bottleneck? What is the overall capacity of the process?
3. If the demand is 20 bunk beds per week, what is the capacity utilization for each step in the process?
4. How would capacity change if we increased the batch size to 8 bunk beds?
5. List three recommendations for increasing the capacity of the process. What are the potential costs/difficulties of each recommendation?
6. If the demand is 20 bunk beds per week, how long does it take for a bunk bed to complete the entire process, from start to finish?
7. List three recommendations for decreasing the throughput time of the process. What are the potential costs/difficulties of each recommendation?

For question 8, assume that the batch size is 4 bunk beds and that the process steps must be performed in the order shown in the table.

8. What is the maximum demand the process can meet (in beds per week) if there are only 5 workers with which to staff the line? What if there are only 4 workers? How should the process steps be allocated between the workers?

Appendix 3

Operations Management: An Introduction to Process Analysis

Glossary

Batch: A group of items or orders that is processed all at one time. The batch size is then the number of items or orders that are processed as a batch.

Bottleneck: The resource that limits the production or service delivery of a process.

Buffer: Area where work-in-process inventory can be stored.

Capacity: The maximum number of items that can be processed by a resource in a given time period. Workers, production equipment, testing equipment, and computers are all examples of resources.

Capacity Utilization: The ratio of the amount of a resource used to the amount available in that time period.

Cycle Time: The average amount of time that elapses between the completion of successive items at a stage, assuming that the process is operating at capacity. The cycle time at a process step is defined as the inverse of the capacity at that step.

Finished Goods Inventory: Items that have been completely processed are referred to as finished goods inventory.

Raw Materials: Materials that have not yet started production are referred to as raw materials.

Run Time: The time it takes to process an item at a process step.

Setup Time: The amount of time it takes to set up or prepare before processing an item or batch of items at a process step.

Throughput Rate: The rate at which units flow through a process. It represents the actual rate at which a process or system generates its products or serves its customers.

Throughput Time: The time it takes for a specific item, job, or order to go through the entire process. In manufacturing settings, throughput time is also known as manufacturing lead time.

Work in Process: Any items that are currently in process, either at a process step or between steps, are referred to as work-in-process inventory or WIP.