Getting a Grip on the Development Cycle: How to Gain Speed and Predictability

Many drug companies think they have done everything possible to improve speed and predictability in their product development process, but many conventional "time-saving" tactics—including multitasking, increasing workloads, and thinking departmentally instead of looking at the company as a whole—may actually cause delays. Shifting strategies is necessary to uncover areas where progress can be made.

Lisa Scheinkopf
Life Sciences Industry Practice Leader
Chesapeake Consulting, Inc.
Severna Park, Maryland

When it comes to product development, pharmaceutical companies dream the impossible dream: to create predictable timelines and complete them on schedule and fast enough to win the first-to-market race. In a development process measured in years, delays of weeks or even days may have devastating consequences. A wrong guess about how long developing a product will take could mean, at worst, finishing second or later and forfeiting millions in sales. At best, it means postponing the day when spending is converted to revenue.

Many pharmaceutical companies have nearly written off the dream of improving speed and predictability: There are simply too many unknowns and too much risk. However, it is possible, often without additional capital expenditure or headcount, to achieve significant improvements in speed, make reliable forecasts, and complete development projects on time. Strategic shifts in belief and behavior can make a profitable difference.

From Busyness to Effectiveness

Established practice suggests that keeping everyone at full capacity—working 80-hour weeks—will speed things up. Research shows, though, that most projects take longer than planned and that keeping everyone working longer does not improve speed. Instead, it is much more important to know what factors limit the ability to work efficiently, and to focus resources there.

By definition, every system has a constrained resource, or bottleneck. In some systems, where capacity is well balanced, the bottleneck shifts unpredictably; one week it might be in analytical chemistry, and the next it is in manufacturing or clinical trials. Not knowing where the constraint will occur is, obviously, undesirable.

In other systems, the same logjam recurs. Imagine a fast-food restaurant with only one cashier, even at the busiest times of day. There is plenty of food, ample cooking staff, and a long line of customers. The constraint is clearly in order taking, and improvements elsewhere will not help.

The same principles apply to large, complex systems. For example, a pharmaceutical manufacturer is attempting to implement literally dozens of improvement projects; at the same time, it experiences delay after delay in getting products through research and development (R&D). Closer examination shows that many of the improvement projects are aimed at areas other than the company's main constraint, a manufacturing area involved with R&D. By focusing only on improvements that have a positive effect on the constrained operation, the company is able to eliminate nearly 100 "improvement" projects and be well on its way to doubling R&D throughput. Clearly, tying up resources in other areas would not help R&D. Speed is improved by identifying the bottleneck, and maximizing the use of resources there.

In pharmaceutical companies, the constraint is often a particular skill set. If, for example, there is a formulations backlog, keeping all employees busy will not increase workflow through that stage; in fact, doing so may strain it. The simplest fix may be to hire more persons with the valued skill. If getting the product to market faster depends heavily on that skill set, and millions of dollars in sales are at stake, it hardly makes sense to try to save the cost of a few salaries. If the skill set is in short supply, offloading less-critical tasks and ensuring that the work arrives fully prepared can help protect this resource.

Productivity can suffer from exhaustion, too. It may be prudent to ensure that high-skill employees do not work 80-hour weeks. Ulcers, nervous breakdowns, or overloaded talent fleeing to a competitor will not speed up the project.

Continued on page 34

Continued from page 32

Therefore, targeting inefficiencies is never as simple as declaring, "The formulators are to blame." Instead, it is a matter of looking at the whole system and understanding where the constraint is. Once this is achieved, one can apply the appropriate resources to move work through the bottleneck.

Converting Waste to Wisdom

Classical wisdom says that idle system capacity is wasteful and should be eliminated. In reality, not having some extra capacity is a sure way to slow down the process and miss key deadlines. This is easy to understand at the individual level.

If a drive to the airport takes 20 minutes under the best conditions, no one would consider leaving home just 20 minutes before departure. Everyone knows that an accident or security check might consume precious minutes. Safety margins are built in to allow adequate time.

In Toronto, Canada, a multilane tollway is seldom a problem to drive, even during rush hour. The multiple lanes represent extra capacity. Driving on this toll road costs extra, but if getting to the airport as fast as possible is sufficiently important, then the extra cost is justified. Likewise, in development projects where the goal is "as fast as possible," the investment in protective capacity is worthwhile.

Too often, timeline forecasts overlook the unpredictable. Mishaps occur, and projects fall behind. With no cushion, the system cannot handle even minor variances, let alone major ones. A certain amount of extra capacity is not excessive, but protective. It allows the system to cope with the unexpected yet inevitable problems that occur in any human endeavor.

On the flip side, excessive cushions are costly. Often, project duration is forecast by summing task-time estimates. This results in conservative, "cushioned" estimates for every task. Without an incentive to finish early, the invisible cushion is always consumed.

A more useful approach is to build visible time buffers into the project, instead of adding a little time for each task. Visible buffers can be shifted to address real variance, like a late shipment or an emergency facility shutdown. When tasks are completed ahead of schedule, the extra time can be applied elsewhere.

In product development, where variances inevitably affect timelines, flexible capacity can actually reduce overall project cost and time. Consider the pharmaceutical company involved in a technology transfer—shifting manufacturing from one facility to another. The technology transfer will enable the company to enhance product quality, better meet FDA requirements, and reduce product costs. A number of departments are involved, each accustomed to developing its own schedules. Using this approach, total project time estimates would be both inflated and unpredictable. Despite adequate total time, the buffers would not be strategically positioned to address problems as they occurred.

Instead, each department helps create a single schedule, with visible buffers placed at the project level. Projects are managed relative to the buffer consumption rate, rather than by task. Workers are encouraged to work as well and as fast as they can, and are measured accordingly, not by due dates of individual tasks or projects for individual departments. The emphasis on whole-project speed comes to the forefront, and the result is a project with a predictable, achievable end-date. The buffers provide the predictability. They are strategically applied to protect against the total unpredictability; the sum of aggressive task estimates plus buffers equals the total project time.

Workers are encouraged to work as well and as fast as they can, and are measured accordingly, not by due dates of individual tasks or projects for individual departments.

As a result of implementing visible buffers, the company meets the transfer schedule for more than a year. The manufacturing process in the new facility is enhanced, reducing costs and making the product more profitable. Predictability swiftly becomes profitability. As a bonus, the FDA in turn is happy with predictable, reliable schedules.

Reordering Chaos

The popular view is that multitasking—working on many things at once—is a desirable and productive behavior. In fact, it is one of the best ways to lengthen cycle time. Suppose, for example, that a Saturday routine includes buying groceries, doing laundry, mowing the lawn, and working on a volunteer presentation project. In multitasking, one might start a grocery list, separate laundry, add more items to the grocery list, change into gardening clothes and mow half the yard, change clothes again and drive to the grocery store, drive home, review the presentation notes, start a load of laundry, realize there is no detergent, and return to the store. Suddenly, the day is half over and not a single task has been completed. A bit of each has been done, but much time has been wasted in between.

The alternative is to "get it, work it, move it"—working on one task until it is finished or no more can be done. Using the household example, that might mean making a grocery list, driving to the store and buying everything at once, starting the laundry, changing clothes and mowing the lawn, putting the laundry in the dryer and starting another load, and

working on the presentation. By lunchtime, the shopping and the lawnmowing have been completed, and the laundry and presentation are underway.

In business organizations, the results of this strategy can be dramatic. One pharmaceutical company doubled the output of its analytical chemistry lab within one week of making the decision to go from multitasking to single tasking. That translates to doubling the speed with which work moved through the lab. It took only a little training, moving all work not designated for immediate consumption to a "holding area," and a prioritization procedure for supervisors to assign work to chemists as they finished tasks. It is almost always possible to speed up work by 10% to 30% just by single tasking. The truth is that many people cannot do more than one task at a time; they can only switch among tasks and lose time in between.

It is almost always possible to speed up work by 10% to 30% just by single tasking.

Opening the Window to a Global View

The intuitive assumption, and one which underlies many improvement methodologies, is that any improvement will contribute to improved project outcome—that local excellence always benefits the whole. Unfortunately, it is not true. Organizations are living systems, and their subsystems are interdependent. It is not always obvious how a change in one part of the system will affect others, or whether it will help or hurt. Various departments and divisions attend to their own performance goals, not necessarily thinking of the good of the whole. Becoming faster as a company simply cannot be achieved by optimizing the silos.

Sustainable improvement, such as the ability to predictably complete lengthy, complex development projects on schedule, requires looking at the whole system and then deciding how the system—not its individual components—will be measured. Buffered projects shift the balance of power. Obviously, people will resist and subvert change that hurts them in the ego or the pocketbook, so making such a shift requires a change in what is tracked and emphasized during implementation.

For example, instead of focusing on departmental milestones, a management system might evaluate the effect of buffer depletion on project-completion dates. Such a shift might place value on celebrating when people have quickly and effectively depleted their queues, rather than on keeping busy. If people "act" busy, and boast about how much work is waiting for them, it is possible that busyness is valued more than effectiveness.

The technology-transfer example also has relevance here. In historic practice, each functional area planned its own tasks and measurements, adding cushions and fighting with other departments over handoff dates. Developing an effective transfer plan required a fundamental shift from local to global thinking, and an accompanying shift in expectations. Departments were no longer held accountable for just their parts, but for the entire project. Each department participated in weekly buffer-management meetings to keep everyone rowing in the same direction. As a result, not only is the company getting the transferred products submitted and approved faster, it is also shipping products to customers and making money.

Stating with confidence how long it will take to complete a development project depends on having an operation in which people act for the good of the whole. They should be rewarded accordingly. This demonstrates a different mindset than the traditional one that pits marketing against manufacturing, or glorifies one group of employees at the expense of another.

"Paradigm shift" is an overused and often trivialized phrase. However, improving speed and predictability requires no less than a fundamental change in understanding what makes the work flow smoothly. It is not about keeping everyone busy all the time, eliminating all waste, juggling tasks, or focusing on local improvements. The way to speed and predictability is through stronger alignment—of behavior and belief, of practices and purpose—and sharper focus on both where the constraints are and where the opportunities lie.

Lisa Scheinkopf, Life Sciences Industry Practice Leader, Chesapeake Consulting, 877 B&A Boulevard, Severna Park, Maryland 21146
Phone: (410) 647-6880; Fax: (410) 647-9738
E-mail: lscheinkopf@chesapeak.com