May 2018

Scrum.org The Home of Scrum Whitepapers

Little's Law for Professional Scrum with Kanban

Daniel Vacanti, CEO Actionable Agile and co-creator of Professional Scrum with Kanban Course

Yuval Yeret's blog article "4 Key Flow Metrics and How to Use Them in Scrum's Events" dealt with the basic metrics of flow necessary for Professional Scrum with Kanban. These metrics are required because they give us tremendous insight into the overall health of a process and they suggest certain interventions that can be taken when a process is not performing the way we would expect it to. But there is another important reason to monitor these metrics. In what is a remarkable result, three of these metrics (Work in Progress (WIP), Cycle Time, and Throughput) are intrinsically linked by a very straightforward and very powerful relationship known as Little's Law:

Average Cycle Time = Average Work In Progress / Average Throughput

The math of Little's Law is simple. But this article is not about the math. What we do care about—and I cannot stress this point enough if we want to gain a greater appreciation of the law's applicability to our world—is looking far beyond the elegance of the equation to get a deeper understanding of what is needed to make the law work*. Much has been written about Little's Law in the Agile community; some good and some bad. My point here is not to give a detailed proof or explanation of Little's Law, but rather to highlight some of the important things to consider when applying the law to Professional Scrum with Kanban.

Little's Law

The fundamental result of Little's Law is that for a given process, in general, the more things that you work on at any given time (on average) the longer it is going to take for each of those things to finish (on average). As a case in point, managers who are ignorant of this law panic when they see that their Cycle Times are too long and perform the exact opposite intervention of what they should do: they start more work. After all, they reason, if things take so long, then they need to start new items as soon as possible so that those items finish on time—regardless of what is currently in progress. The result is that items only take longer and longer to complete. Thus, managers feel more and more pressure to start things sooner and sooner. You can see how this vicious cycle gets started and perpetuates itself. After studying Little's Law, you should realize that if Cycle Times are too long then the first thing you should consider is lowering WIP. It feels uncomfortable, but it is true. In order to get more stuff done faster, you need to work on less (again, on average).

In other words, just having a qualitative understanding that WIP is directly proportional to Cycle Time (without quantitatively understanding the mathematics) is good enough for our purposes here. This relationship between WIP and Cycle Time is upon which most of the other elements of Kanban are built.

But there are a few pitfalls to avoid when applying Little's Law to an Agile process and so I'll discuss a couple of those here. Remember these are by no means an exhaustive list, but they do represent some things to watch out for whenever you hear conversations about the law.

Forecasting

My guess is that you were expecting me to say that once you understand Little's Law all you need to do is to plug in the numbers and out will pop the forecasting result that you are looking for (à la Newton's F = ma or Einstein's $E=mc^2$). However, nothing could be further from the truth.

The first thing that you need to know about Little's Law is that it is concerned with looking backward over a time period that has completed. It is not about looking forward; that is, is not meant to be used to make *deterministic* predictions. As Dr. Little himself says about the law, "This is not all bad. It just says that we are in the measurement business, not the forecasting business".

This point requires a little more discussion as it is usually where people get hung up. The "law" part of Little's Law specifies an exact relationship between average WIP, average Cycle Time, and average Throughput, and this "law" part only applies only when you are looking back over historical data. The law is not about—and was never designed for—making deterministic forecasts about the future. For example, let's assume a team that historically has had an average WIP of 20 work items, an average Cycle Time of 5 days, and an average Throughput of 4 items per day. You cannot say that you are going to increase average WIP to 40, keep average Cycle Time constant at 5 days and magically Throughput will increase to 8 items per day—even if you add staff to the keep the WIP to staff ratio the same in the two instances. You cannot assume that Little's Law will make that prediction. It will not. All Little's Law will say is that an increase in average WIP will result in a change to one or both of average Cycle Time and average Throughput. It will further say that those changes will manifest themselves in ways such that the relationship among all three metrics will still obey that law. But what it does not say is that you can deterministically predict what those changes will be. You have to wait until the end of the time interval you are interested in and look back to apply the law.

But that restriction is not fatal. The proper application of Little's Law in our world is to understand the assumptions of the law and to develop process policies that match those assumptions. If the process we operate conforms—or mostly conforms—to all of the assumptions of the law then we get to a world where we can start to trust the data that we are collecting off of our system. It is at this point that our process is probabilistically predictable. Once there we can start to use something like Monte Carlo simulation on our historical data to make forecasts and, more importantly, we can have some confidence in the results we get by using that method.

The second, potentially more important reason to not use Little's Law for forecasting is that it is a **relationship of averages**. I mention this because even if you could use Little's Law as a forecasting tool (which you cannot), you would not want to as you would be producing a forecast based on averages. There are all kinds of reasons why you should not forecast based on averages—too many to go into

here. It turns out we can (and should) do better than average forecasting, and we have powerful tools at our disposal to do so.

Having said all that, though, there is no reason why you cannot use the law for quick, back-of-the-envelope type estimations about the future. Of course, you can do that. I would not, however, make any completion commitments, staff hiring/firing decisions, or project cost calculations based on this type of calculation alone. I would further say that it is negligent for someone to even suggest doing so. But this simple computation might be useful as a quick gut-check to decide if something like a project is worth any further exploration.

Size Does Not Matter

I have one last topic I want to cover before wrapping up. Strange as it may seem, there is nothing about Little's Law that requires all work items to be of the same size. Most people assume that an application of Little's Law specifically—and limiting WIP in general—necessitates that all work items be of the same size. That is simply not true. The precise reasons why would fill up a chapter in its own right, so I am going to limit my comments to two brief points.

First, work items size does not matter because for Little's Law because we do not necessarily care about each item individually, we care about what all items look like on average.

Second, and more importantly, the variability in work item size is probably not the variability that is killing your predictability. Your bigger predictability problems are usually too much WIP, the frequency with which you violate Little's Law's assumptions, etc. Generally, those are easier problems to fix than trying to arbitrarily make all work items the same size. Even if you were in a context where size did matter, it would be more about right-sizing your work and not same-sizing your work.

Conclusion

The bottom line is that Little's Law is not about understanding the mathematics of queuing theory. It is about understanding the assumptions that need to be in place in order for the law to work (again, see "AAMFP"). We can use those assumptions as a guide, or blueprint, or model for our own process policies. Whenever your process policies are in violation of the assumptions of Little's Law then you know that you have at least diminished—or possibly eliminated—your chance of being predictable.

For instance, as you operate your process think about the times and reasons why work flows in at a faster rate than work flows out. Think about why items age unnecessarily due to blockages or poor pull policies. Think about why work is abandoned when only partially complete (and how your team accounts for that abandonment). Think about how these occurrences are violating the assumptions Little's Law and how they are ultimately affecting your ability to be predictable. But more importantly, think about how your understanding of Little's Law should result in behavior changes for you and your team. When violations of Little's Law occur, it is usually because of something you did or chose (intentionally or not) not to do. Remember, you have much more control over your process than you think you do.

For a more detailed discussion around the mechanics of this relationship, you can read the book "Actionable Agile Metrics for Predictability" by Daniel Vacanti