

An Introduction to Agile Software Development

By: Victor Szalvay

What is Agile? How can Agile help improve my organization? Those are just a few questions I will answer in this introduction to the Agile school of software development.

This paper is written for IT managers and CXOs with an interest in improving development productivity. First, I introduce the two broad schools of thought common to software development: traditional sequential, a.k.a. “the waterfall method,” and iterative methods, of which Agile is one subset. My objective is to demonstrate the shortcomings of the waterfall approach, while providing a solution in iterative and, more specifically, Agile methods.

What is Agile? How can Agile help improve my organization? Those are just a few questions I will answer in this introduction to the Agile school of software development.

PART I: SHORTCOMINGS OF THE TRADITIONAL WATERFALL APPROACH

The essence of waterfall software development is that complex software systems can be built in a sequential, phase-wise manner in which all of the requirements are gathered at the beginning, all of the design is completed next, and, finally, the master design is implemented into production quality software. This approach asserts that complex systems can be built in a single pass, without revisiting requirements or design ideas in light of changing business or technology conditions. It was first introduced in an article written by Winston Royce in 1970 and was primarily intended for use in government projects¹.

Waterfall equates software development to a production line conveyor belt. “Requirements analysts” compile the system specifications until they pass the finished requirements specification document to “software designers,” who plan the software system and create diagrams documenting how the code should be written. The design diagrams are then passed to the “developers,” who implement the code from the design (Figure 1).



Figure 1: Traditional Methods: Sequential Phased Approach

Using the waterfall approach, traditional IT managers have made valiant efforts to craft and adhere to large-scale development plans. These plans are typically laid out in advance of development projects using Gantt or PERT charts to map detailed tasks and dependencies for each member of the development group months or even years down the line. However, studies of past software projects show that only nine to 16 percent are considered on-time and on-budget². In this article, I attempt to summarize current thinking among computer scientists on why waterfall fails in so many cases. I also explore a leading alternative to waterfall: “Agile” methods that focus on incremental and iterative development where requirements, design, implementation, and testing continue throughout the project lifecycle.

Imagine a scenario in which you engage a software group to build a critical software system. Do you think you could provide every last detail the developers need to know right off the bat?

Up-front Requirements Analysis

What are requirements? From the stakeholder's perspective, the requirements are the features and specifications of the system. Requirements define what developers are to build. For example, the system must have a Web site with e-commerce capability that can handle 10,000 purchases per hour or the system must be accessible 99.999 percent of the time.

One of the biggest problems with waterfall is that it assumes that all project requirements can be accurately gathered at the beginning of the project. In Figure 1, the first block represents the requirements analysis phase of a software development project. Analysts slave for weeks or months compiling everything they can glean about the proposed system into comprehensive "Software Requirements Specification" (SRS) documents. Once finished, the SRS is sent over the fence to the designers while the requirements analysts go to work on the next project.

Imagine a scenario in which you engage a software group to build a critical software system. Do you think you could provide every last detail the developers need to know right off the bat? I have yet to encounter such a customer and I am hard pressed to think I ever will. As a start, consider the areas that must be addressed: business rules and exceptions; scalability and concurrent user support; browser or OS support; user roles and restrictions; user interface standards. In fact, it is inevitable that attempts at up-front requirements specification will leave out some very important details simply because the stakeholders cannot tell developers everything about the system at the beginning of the project³. This means that the requirements typically change outside of the requirements phase in the form of "change orders," and in many waterfall projects this can be very costly. By virtue of a requirements change, the intricately planned design can be affected dramatically, which will in turn affect any implementation and test strategies. The cost of change in a waterfall project increases exponentially over time because the developer is forced to make any and all project decisions at the beginning of the project.

What if your business needs are still emerging and certain aspects of the system are rapidly changing or cannot be defined yet? Business climates and objectives often change rapidly, especially in today's age of instant information. Can you afford to lock your business into a rigid long-term project where the cost of change grows exponentially? For example, a national test preparation organization commissioned my company to build a simulator for an upcoming standardized test. Since the test itself had not been released yet, the types of questions that would appear on the test were unknown when we started development. But the system had to be done shortly after the tests were released. Markets are forcing the software development community to respond with flexible development plans that flatten the cost of change.

Bridge building relies on physical and mathematical laws, but software development has no laws or clear certainties on which to build. As a result, software is almost always flawed or sub-optimized.

People need to see and feel something before they really know what they want. The "I'll Know it When I See It" (IKIWISI) law says that software development customers can better articulate what they really want after they see and try a piece of functional software. I often use a "drawing" analogy to help explain this effect. Although I'm a terrible artist, when I draw a picture, I need to see the drawing as I progress. If I tried to close my eyes and draw the same picture, it would prove far less successful. But this is what waterfall asks customers to do: specify the entire system without having a chance to periodically evaluate the progress and make adjustments to the requirements as needed. Waterfall is an "over the fence" approach; the requirements are solicited from the user and, some time later, the finished product is presented to the user. This is entirely unnatural because customers find it difficult to specify software perfectly without observing it evolve and progress.

The problem of undefined, changing, and emerging requirements presents a very large challenge to waterfall projects because, by definition, all requirements must be captured up-front at the risk of costly changes later.

Software Development is More Like New Product Development than Manufacturing

Software development is a highly complex field with countless variables impacting the system. All software systems are imperfect because they cannot be built with mathematical or physical

certainty. Bridge building relies on physical and mathematical laws, but software development has no laws or clear certainties on which to build. As a result, software is almost always flawed or sub-optimized. Also consider that the building blocks of software projects are usually other software systems (e.g., programming languages, database platforms, etc.), and those systems that act as building blocks contain bugs and cannot be relied upon with certainty. Because the foundations of software development are inherently unstable and unreliable, organizations developing software must realize variables exist that are largely outside of management control. It is therefore fair to say that software development is more akin to new product research and development than it is to assembly line-style manufacturing. Software development is innovation, discovery, and artistry; each foray into a development project presents new and difficult challenges that cannot be overcome with one-size-fits-all, cookie-cutter solutions⁴.

The waterfall methodology assumes that up-front planning is enough to take into account all variables that could impact the development process. In fact, waterfall projects allocate ample effort to the task of detailing every possible risk, mitigation plan, and contingency. But is it possible to predict any and all variables that could possibly affect a software project? The empirical answer is “no,” considering the limited success of waterfall projects⁵.

Waterfall therefore inaccurately equates the dynamic process of software development to a static assembly line. In that idealized — and unrealistic — version of development, defined processes can be established such that, when used sequentially, they result in a successful project each time. The first step is X, the second is Y, and the result is always Z. Can research really be relegated to a series of steps that, when performed in sequence, result in a new product? If this formulaic approach were adequate, medical researchers could simply plug variables into equations to discover new medicines. On the contrary, since the late 1970s, product development companies led by Toyota, Honda, Fujitsu, 3M, HP, Canon, and NEC supplanted the sequential “Phased Program Planning” (PPP) approach to new product development with a flexible, holistic approach where the traditional phases of development overlap throughout the entire product lifecycle⁶. The results were a dramatic reduction in cost and development time to market, ultimately leading to the popular rise of “lean development” and “just-in-time manufacturing.” Following the lead of Japanese auto makers, sequential, waterfall-style approaches to new product development were effectively abandoned outside the software development industry in the 1990s⁷. But longstanding insistence from IT managers to categorize software development as a straightforward assembly line progression has kept the industry from evolving to better methods, the benefits of which other new product development industries have been reaping for decades. It’s ironic that a cutting edge technology field like software is so far behind more traditional engineering fields in terms of development methods.

Almost no software system is so simple that the development can be entirely scripted from beginning to end. The inherent uncertainty and complexity in all software projects requires an adaptive development plan to cope with uncertainty and a high number of unknown variables.

Almost no software system is so simple that the development can be entirely scripted from beginning to end. The inherent uncertainty and complexity in all software projects requires an adaptive development plan to cope with uncertainty and a high number of unknown variables.


Agile methods stress productivity and values over heavyweight process overhead and artifacts.

PART II: ITERATIVE AND AGILE METHODS

Iterative Development

The simple ability to revisit the “phases” of development dramatically improves project efficiency. This idea of revisiting phases over and over is called “iterative development.” Under iterative processes, the traditional lifecycle is supplanted by a sequence of short “iterations” (typically two to four weeks in length), each of which touches on each of the traditional “phases” of development. For example, gathering requirements is an ongoing process that is periodically revisited. As new

requirements surface and the scope of the project evolves, iterative processes continually capture the requirements iteration after iteration. Interestingly, Winston Royce (of waterfall process fame) later noted that his ideas were incorrectly interpreted and that a “single pass” framework would never work (his article actually advocates at least a second pass)⁸. Multiple “passes,” or iterations, over a project lifecycle are required to properly address complexities and risk factors.



Source: Adapted from H. Takeuchi and I. Nonaka, “The New New Product Development Game”, Harvard Business Rev., Jan. 1986, pp. 137-146.

Figure 2: Iterative Approach: Overlapping Phases of Development

This concept of iterative development hails from the “lean development” era of the 1980s. As described above, Japanese auto makers made tremendous efficiency and innovation increases simply by removing the phased, sequential approach and implementing an iterative approach, in which prototypes were developed for short-term milestones (Figure 2). Each phase was actually a layer that continued throughout the entire development lifecycle; the requirements, design, and implementation cycles were revisited for each short-term milestone. This “concurrent” development approach created an atmosphere of trial-and-error experimentation and learning that ultimately broke down the status quo and led to efficient innovation⁹. Although direct analogies between industries are never seamless, the success of lean development has influenced a broad class of “iterative” software methods including the Unified Process, Evo, Spiral, and Agile methods.

Agile Methods: Embracing Change

Agile methods stress productivity and values over heavyweight process overhead and artifacts. The Agile Manifesto¹⁰, a concise summary of Agile values, was written and signed in 2001, although Agile methods have existed since the early 90s. Agile methods promote an iterative mechanism for producing software and they further increase the iterative nature of the software lifecycle by tightening the design-code- test loop to at least once a day (if not much more frequently) as opposed to once per iteration. Agile visionary Kent Beck challenged the traditional cost of change curve pioneered by Barry Boehm¹¹ more than 20 years ago. Beck’s model asserts that the cost of change can be inexpensive, even late in the project lifecycle while maintaining or increasing system quality¹². Beck’s idealistic “flat” cost of change curve has since been revised and softened by Alistair Cockburn¹³ and Scott Ambler¹⁴ to reflect modern corporate realities. Nevertheless, Agile ideals can be applied to reduce the cost of change throughout the software lifecycle even if the cost of change is not perfectly flat.

To accomplish this “flatter” cost of change curve, Agile methods promote a number of engineering practices that enable cost-effective change. Author and speaker Martin Fowler describes testing and continuous integration as the “enabling” practices that yield Agile’s advantages, like rapid production and minimal up-front design¹⁵. “Test-driven

For a process to be considered “Agile,” individual iterations should produce some fully functional feature or enhancement. In this way, the end product is developed piece by piece, iteration by iteration.




Figure 3: Cost of Change Curves

development” is a quality-first approach in which developer tests (called “unit tests”) are written prior to the functional code itself. Rather than focusing a great deal of effort on big, up-front design analysis, small increments of functional code are produced according to immediate business need. It is the role of the automated test suite built around the rapidly evolving code to act as a harness that allows developers to make aggressive code changes without fear of undetected regression failure.

Object technology and modern integrated development environments (IDEs) boasting built-in testing and refactoring mechanisms negate Boehm’s cost of change curve and allow for the cheap change, even late in the project lifecycle.

Agile Methods: Incremental and Empirical Processes

Agile methods are also characterized and differentiated from other iterative methods (like RUP) by virtue of being “incremental.” Incremental development means that a product is built chunk by chunk throughout the product development lifecycle. For a process to be considered “Agile,” individual iterations should produce some fully functional feature or enhancement. In this way, the end product is developed piece by piece, iteration by iteration. Agile methods are therefore iterative and incremental processes.

Scrum, a popular Agile project management method, introduced the concept of empirical process control for the management of complex, changing software projects. Scrum holds that straightforward defined processes alone cannot be used to effectively manage complex and dynamic software projects. Risk factors and emerging requirements complicate software development to a point where defined processes fall short. Although it has been attempted in the past, there cannot be a single exhaustive library of defined processes to handle every situation that could possibly surface during a software project. In fact, the manufacturing industry has long known that certain chemical processes, for example, are too difficult to script and define. Instead, an empirical or adaptive management approach is employed to measure and adjust the chemical process periodically to achieve the desired outcome¹⁶. Because each iteration produces some tangible increment, in the Scrum process, project plans are continuously inspected and adapted based on the empirical reality of the development.

The frequent demonstration and release of software common in Agile approaches gives customers a chance to “try software” periodically and provide feedback. In short, Agile helps companies produce the “right product.”

Lean Thinking

Another effective way to analyze how Agile methods benefit organizations is to apply lean manufacturing principles to software development. Although cross-industry analysis can be tenuous, Agile methods have their conceptual roots in the Japanese manufacturing productivity boom of the 1980s¹⁷.

Consider, for example, the “small batch” principle: Things produced in smaller batches are of higher quality and efficiency because the feedback loop is shorter. Controls can be adjusted more frequently and resources are utilized efficiently to avoid “queuing” (see “queuing theory” and the theory of constraints). Second, Agile methods encourage delaying irreversible decisions until the last responsible moment. Many software development organizations that implement Agile software development discover the methodology delivers something they never expected: options. Rather than locking into decisions at the beginning of a project, organizations can reduce risks by leaving options open, preferring to make decisions when more accurate information is available. Third, the concept of frequent or continuous integration keeps software development teams synchronized. Teams can work independently for a while, but the code base never diverges for long periods of time, thereby reducing the risks associated with large integrations at the tail end of projects¹⁸.

Agile Requirements: A Focus on Business ROI

Agile projects avoid “up-front” requirements gathering for the reasons stated above: Customers cannot effectively define all requirements with enough detail for implementation to occur at the beginning of a project. Customers often wish to postpone making decisions about the system until

they have the information to make a well-informed decision. Agile values transparency, open communication, and direct involvement with the customer. The frequent demonstration and release of software common in Agile approaches gives customers a chance to “try software” periodically and provide feedback. In short, Agile helps companies produce the “right product.” An iterative approach also allows customers to delay decisions. Decisions can be delayed to some future iteration when better information or technology is available to optimize the choice. For example, we recently delayed selecting a database package for an application because some of the desired features were not available at that time. Thus we built the system in a database independent manner, and, a new version was released by one of the database vendors that solved our problem.

One of the biggest advantages to Agile methods is that work can begin before all of the requirements are known. Many organizations are not fully staffed with business analysts cranking out reams of requirements specs. On the contrary, in our experience, the bottleneck in the development process is often the result of the lack of availability of customer domain experts to perform detailed requirements analysis. This is especially the case with small businesses where domain experts wear many hats and often cannot commit to two or three months of straight requirements analysis. Agile methods are ideally suited, then, to take on bite-sized chunks of requirements that the customer can easily digest.

Going to production with high-priority features is better than never going to production at all.

How do Agile projects prioritize work? A study by the Standish Group shows that, in typical software systems, 45 percent of the features are never actually used by users and another 19 percent are only rarely used¹⁹. This is largely because the unused features were specified in some up-front plan before the ratio of their cost to value was considered or even knowable. Focusing on high business value features first is therefore a critical component of efficient Agile development. Because we can change direction rapidly (every iteration) and the cost of change is low, there is a valuable opportunity for the customer to re-examine business factors at the beginning of each iteration to select features for inclusion in the current iteration according to business ROI. Of course, the development team must bring technical risks to the customer, but, in the end, it is the customer that decides what the development team builds.

Convergence with Agile

One of the most commonly asked questions by those examining Agile is: “How do you know when the software will be finished if there’s no up-front plan?” Of course, that question is invariably followed by: “How can we budget for such a project?” It sounds a bit scary: Let’s start working in short iterative cycles that yield demonstrable software without actually planning everything in advance. But we already know that we cannot plan for everything in advance.

The Agile answer is to examine project progress empirically, rather than trying to guess how things might shape up a priori. Agile processes like Scrum and XP use a concept called velocity, which is the amount of estimated effort a team can complete in a time-boxed iteration. Once a team has established a velocity, a Project Burndown Chart can be utilized to estimate the eventual conclusion of an estimated backlog of work. Each point of the chart in Figure 4 represents an iteration (or Sprint in Scrum), and the Y-axis represents the total estimated effort remaining for the backlog. As iterations progress, a trend line can be established through the points to create a velocity (work the team can complete per iteration). The trend line can then be extrapolated to determine the X-intercept, which represents the empirical estimate of the completion date.

Figure 4 is a Product Burndown Chart representation of a typical project. Through the first nine iterations, the project’s “burndown” trend line indicated an X-intercept well into the future (off the charts!). By iteration 10, the product scope had been adjusted down so that the project could be completed by the budgeted 20th iteration. Notice also that the velocity (slope) changed for the better. In this case, perhaps the reduction in scope was accompanied by the removal of critical impediments to efficient progress.

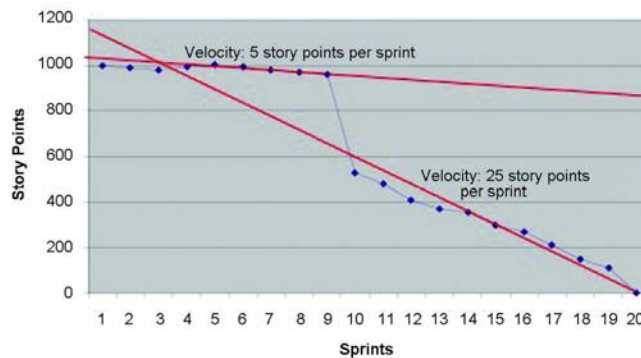


Figure 4: Product Burndown Chart with Velocity

For complex problems like project convergence, Agile methods tell us that the customer cannot specify all four of the software development variables (cost, schedule, scope, quality). To answer the questions above, if the timeline and cost variables are fixed, then the scope of the work must be variable or the definition of the scope must be at a high level so the robustness of each feature can be negotiated²⁰. That is, work will proceed on the highest priority requirements first to ensure that the most important things get done before a deadline or the money runs out. Going to production with high-priority features is better than never going to production at all, especially considering the aforementioned Standish report cited that nearly 65 percent of features are never or rarely used in reality. And quality should not be negotiable; the features built should always be high quality, adhering to strict code and testing standards. There is no guarantee all features will be built, but it is certain that the highest-priority features will go into production and that they will be built well.

CONCLUSION

But does Agile work? Of course, the proof is always in the pudding, and the 2004 Standish Group CHAOS report on the success of software projects shows a dramatic improvement in the failure rate of software projects. In 1994, Standish reported a 31 percent failure rate that has improved to 15 percent in 2004²¹. Standish Chairman Jim Johnson attributes the improvement to smaller projects using iterative processes as opposed to the waterfall method²².

Although waterfall is often referred to as “traditional,” software engineering has had a very short history relative to other engineering disciplines.

The notion that Agile is a radical deviation from the long established, tried-and-true history of waterfall software development is incorrect. Although waterfall is often referred to as “traditional,” software engineering has had a very short history relative to other engineering disciplines. Unlike bridge building, software development is not built on thousands of years of trial and error, and is therefore in a rapidly evolving infancy as an engineering discipline. Agile is simply the latest theory that is widely replacing the waterfall approach that itself will change and evolve well into the future²³.

FURTHER READING

After reading an introductory article, sometimes I wish I could ask the author for a list of good books or articles to further my knowledge on the subject. Below are a few good starting points for anyone interested in learning more about Agile software development.

Craig Larman, *Agile & Iterative Development: A Manager’s Guide*, Addison-Wesley, 2003. This book should be read by any manager interested in Agile. It has the most comprehensive empirical evidence for Agile/iterative of any book currently on the market. It also nicely summarizes and contrasts some of the major Agile/iterative approaches such as Scrum, XP, and the Unified Process.

Mary Poppendieck, Tom Poppendieck, [Lean Software Development: An Agile Toolkit](#), Addison-Wesley, 2003. This book is outstanding and each page seems to offer a valuable nugget of information. It contains the most compelling case for Agile over sequential development I have yet to uncover; much of this article leans on the Poppendiecks' work. It also delivers outstanding tools for implementing Agile.

Ken Schwaber, Mike Beedle, [Agile Software Development with Scrum](#), Prentice Hall, 2001. This book is about "Scrum," a highly effective Agile project management method. Theory and rhetoric are nice, but how do you do Agile? Scrum is a very good place to start on the management side.

Kent Beck, [Extreme Programming Explained: Embrace Change](#), Addison-Wesley, 2000. So what do your developers do differently in Agile? Extreme Programming (or XP) advocates engineering principles such as pair programming and test-driven development and Beck's book is the de facto authority.

ABOUT THE AUTHOR



Victor Szalvay serves as the Chief Technology Officer of CollabNet's Scrum business unit. With a strong background in Agile software processes and, in particular, Scrum, he is a Certified Scrum Trainer (CST) with experience implementing Agile frameworks for software development teams in both the public and private sector.

ABOUT COLLABNET

CollabNet leads the industry in Agile application lifecycle management (Agile ALM) in the Cloud. The CollabNet TeamForge ALM platform, CollabNet Subversion [software configuration management \(SCM\)](#) solution, and ScrumWorks project management software enable teams using any environment, methodology, and technology to improve productivity up to 50% and to reduce the cost of software development by up to 80%. Millions of users at more than 2,500 organizations, including Applied Biosystems, Capgemini, Deutsche Bank, Oracle, Reuters, and the U.S. Department of Defense, have transformed the way they develop software with CollabNet. For more information, visit www.collab.net.

REFERENCES

1. Winston Royce, [Managing the Development of Large Software Systems](#), Proc. Westcon, IEEE CS Press, 1970, pp. 328-339.
2. Standish Group International, Inc., "Chaos Chronicles", 1994, http://www1.standishgroup.com/sample_research/chaos_1994_1.php
3. Kent Beck, [Extreme Programming Explained: Embrace Change](#), Addison-Wesley, 2000, pp. 18-19.
4. Ken Schwaber, Mike Beedle, [Agile Software Development with Scrum](#), Prentice Hall, 2001, pp. 89-94
5. Standish Group International, Inc., "Chaos Chronicles", 1994, http://www1.standishgroup.com/sample_research/chaos_1994_1.php
6. I. Nonaka, H. Takeuchi, "The New New Product Development Game," Harvard Business Review, January 1986, pp. 137-146.
7. Mary Poppendieck, Tom Poppendieck, [Lean Software Development: An Agile Toolkit](#), Addison-Wesley, 2003
8. Craig Larman, Victor R. Basili, "Iterative and Incremental Development: A Brief History", Computer, IEEE CS Press, June 2004, p. 48.
9. I. Nonaka, H. Takeuchi, "The New New Product Development Game", Harvard Business Review, January 1986, pp. 137-146.
10. *The Agile Manifesto* is online at <http://www.agilemanifesto.org/>
11. Barry Boehm, [Software Engineering Economics](#), Prentice Hall PTR, 1981.

12. Kent Beck, Extreme Programming Explained: Embrace Change, Addison-Wesley, 2000.
13. "Reexamining the Cost of Change Curve, year 2000," by Alistair Cockburn; XP Magazine, September 2000.
14. "Examining the Cost of Change Curve," by Scott Ambler; Agile Modeling Essays excerpted from the book The Object Primer, 3rd ed.: Agile Model-Driven Development with UML2"; by Scott Ambler, Cambridge University Press, 2004.
15. Martin Fowler, "Is Design Dead," <http://martinfowler.com/articles/designDead.html> , 2004.
16. Ken Schwaber, Mike Beedle, Agile Software Development with Scrum, Prentice Hall, 2001, pp. 100-101
17. The Scrum Agile method has its roots in the Nonaka-Takeuchi article referenced above.
18. Mary Poppendieck, Tom Poppendieck, Lean Software Development: An Agile Toolkit, Addison-Wesley, 2003, p 28.
19. Jim Johnson, "ROI, It's Your Job!" Published Keynote Third International Conference on Extreme Programming, 2002.
20. Mary Poppendieck, Tom Poppendieck, Lean Software Development: An Agile Toolkit, Addison-Wesley, 2003, p 32.
21. Standish Group International, Inc., "Chaos Chronicles," 2004.
22. "Standish: Project Success Rates Improved Over 10 Years", Software Magazine and Weisner Publishing, <http://www.softwaremag.com/L.cfm?Doc=newsletter/2004-01-15/Standish>.
23. Software architect Michael James has a semi-serious theory that "test-only" development will soon be feasible with the advances in cheap, clustered supercomputing. Developers write only robust tests and supercomputers will write and compile code until all tests pass.