

Distributed Scrum: Agile Project Management with Outsourced Development Teams

Jeff Sutherland, Ph.D.
Patientkeeper
 Newton, MA, US
jeff.sutherland@computer.org

Anton Viktorov
StarSoft Dev. Labs
 St. Petersburg, Russia
anton.viktorov@starsoftlabs.com

Jack Blount
SirsiDynix
 Provo, UT, USA
jack@dynix.com

Nikolai Puntikov
StarSoft Dev. Labs
 St. Petersburg, Russia
nick@starsoftlabs.com

Abstract

Agile project management with Scrum derives from best business practices in companies like Fuji-Xerox, Honda, Canon, and Toyota. Toyota routinely achieves four times the productivity and 12 times the quality of competitors. Can Scrum do the same for globally distributed teams? Two Agile companies, SirsiDynix and StarSoft Development Laboratories achieved comparable performance developing a Java application with over 1,000,000 lines of code. During 2005, a distributed team of 56 Scrum developers working from Provo, Utah; Waterloo, Canada; and St. Petersburg, Russia, delivered 671,688 lines of production Java code. At 15.3 function points per developer/month, this is the most productive Java project ever documented. SirsiDynix best practices are similar to those observed on distributed Scrum teams at IDX Systems, radically different than those promoted by PMBOK, and counterintuitive to practices advocated by the Scrum Alliance. This paper analyzes and recommends best practices for globally distributed Agile teams.

1. Introduction

Scrum is an Agile software development process designed to add energy, focus, clarity, and transparency to project teams developing software systems. It leverages artificial life research [1] by allowing teams to operate close to the edge of chaos to foster rapid system evolution. It capitalizes on robot subsumption architectures [2] by enforcing a simple set of rules that allows rapid self-organization of software teams to produce systems with evolving architectures. A properly implemented Scrum was designed to increase speed of development, align individual and organization objectives, create a culture driven by performance, support shareholder value creation, achieve stable and consistent communication of performance at all levels, and enhance individual development and quality of life.

Scrum for software development teams began at Easel Corporation in 1993 and was used to build the first

object-oriented design and analysis (OOAD) tool that incorporated round-trip engineering. In a Smalltalk development environment, code was auto-generated from a graphic design tool and changes to the code from the Smalltalk integrated development environment (IDE) were immediately reflected back into design.

A development process was needed to support enterprise teams where visualization of design immediately generated working code. This led to an extensive review of the computer science literature and dialogue with leaders of hundreds of software development projects. Key factors that influenced the introduction of Scrum at Easel Corporation were fundamental problems inherent in software development

- Uncertainty is inherent and inevitable in software development processes and products - Ziv's Uncertainty Principle [3]
- For a new software system the requirements will not be completely known until after the users have used it - Humphrey's Requirements Uncertainty Principle [4]
- It is not possible to completely specify an interactive system – Wegner's Lemma [5]
- Ambiguous and changing requirements, combined with evolving tools and technologies make implementation strategies unpredictable.

“All-at-Once” models of software development uniquely fit object-oriented implementation of software and help resolve these challenges. They assume the creation of software involves simultaneously work on requirements, analysis, design, coding, and testing, then delivering the entire system all at once [6].

1.1. “All-at-Once” Development Models

The simplest “All-at-Once” model is a single super-programmer creating and delivering an application from beginning to end. This can be the fastest way to deliver a product that has good internal architectural consistency and is the “hacker” model of implementation. For example, in a “skunk works” project prior to the first Scrum, a single individual surrounded by a support team spent two years

writing every line of code for the Matisse object database [7] used to drive \$10B nuclear reprocessing plants worldwide. At less than 50,000 lines of code, the nuclear engineers said it was the fastest and most reliable database ever benchmarked for nuclear plants.

IBM documented a variant of this approach called the Surgical Team and considered it the most productive approach to software development [8]. The Surgical Team concept has a fatal flaw in that there are at most one or two individuals even in a large company that can execute this model. For example, it took three years for a competent team of developers to understand the conceptual elegance of the Matisse object server well enough to maintain it. The single-programmer model does not scale well to large projects.

The next level of “All-at-Once” development is handcuffing two programmers together. Pair programming, an eXtreme Programming practice [9], is an implementation of this. Here, two developers working at the same terminal deliver a component of the system together. This has been shown to deliver better code (usability, maintainability, flexibility, extendibility) faster than two developers working individually [10]. The challenge is to achieve a similar productivity effect with more than two people.

Scrum, a scalable, team-based “All-at-Once” model, was motivated by the Japanese approach to team-based new product development combined with simple rules to enhance team self-organization (see Brooks’ subsumption architecture [2]). At Easel, the development team was already using an iterative and incremental approach to building software [11]. Features were implemented in slices where an entire piece of fully integrated functionality worked at the end of an iteration. What intrigued us was Takeuchi and Nonaka’s description of the team-building process for setting up and managing a Scrum [12]. The idea of building a self-empowered team in which a daily global view of the product caused the team to self-organize seemed like the right idea. This approach to managing the team, which had been so successful at Honda, Canon, and Fujitsu, also resonated with research on systems thinking by Professor Senge at MIT [13].

1.2. Hyperproductivity in Scrum

Scrum was designed to allow average developers to self-organize into high performance teams. The first Scrum achieved a hyperproductive state in 1993-1994 because of three primary factors. The first was the Scrum process itself, characterized by 15 minute daily meetings where each person answers three questions – what did you accomplish yesterday, what will you do today, and what impediments are getting in your way? This is now part of the standard Scrum organizational pattern [14]. Second,

the team implemented all XP engineering processes [9] including pair programming, continuous builds, and aggressive refactoring. And third, the team systematically stimulated rapid evolution of the software system.

One of the most interesting complexity phenomena observed in the first Scrum was a “punctuated equilibrium” effect [15]. This phenomenon occurs in biological evolution when a species is stable for long periods of time and then undergoes a sudden jump in capability. Danny Hillis simulated this effect on an early super-computer, the Connection Machine [16].

“The artificial organisms in Hillis’s particular world evolved not by steady progress of hill climbing but by the sudden leaps of punctuated equilibrium... With artificial organisms Hillis had the power to examine and analyze the genotype as easily as the realized phenotypes... While the population seemed to be resting during the periods of equilibrium ... the underlying genetic makeup was actively evolving. The sudden increase in fitness was no more an instant occurrence than the appearance of a newborn indicates something springing out of nothing; the population seemed to be gestating its next jump. Specifically, the gene pool of the population contained a set of epistatic genes that could not be expressed unless all were present; otherwise the alleles for these genes would be recessive.” [17]

Using Scrum with a fully integrated component design environment leads to unexpected, rapid evolution of a software system with emergent, adaptive properties resembling the process of punctuated equilibrium. Sudden leaps in functionality resulted in earlier than expected delivery of software in the first Scrum. Development tasks, originally planned to take days, could often be accomplished in hours using someone else’s code as a starting point.

This aspect of self-organization is now understood as a type of Set-Based Concurrent Engineering (SBCE) practiced at Toyota [18]. Developers consider sets of possible solutions and gradually narrow the set of possibilities to converge on a final solution. Decisions on how and where to implement a feature is delayed until the last possible moment. The most evolved component is selected “just in time” to absorb new functionality, resulting in minimal coding and a more elegant architecture. Thus emergent architecture, a core principle in all Agile processes, is not random evolution. Properly implemented, it is an SBCE technique viewed as a best business practice in some of the world’s leading corporations.

2. The SirsiDynix Distributed Scrum

Scrum was designed to achieve a hyperproductive state where productivity increases by an order of magnitude over industry averages. Many small, collocated teams have

achieved this effect. The question for this paper is whether a large, distributed, outsourced team can achieve the hyperproductive state.

U.S., European, or Japanese companies often outsource software development to Eastern Europe, Russia, or the Far East. Typically, remote teams operate independently and communication problems limit productivity. While there is a large amount of published research on project management, distributed development, and outsourcing strategies as isolated domains, there are few detailed studies of best project management practices on large systems that are both distributed and outsourced.

Current recommended Scrum practice is for local Scrum teams at all sites to synchronize once a day via a Scrum of Scrums meeting. Here we describe something rarely seen. At SirsiDynix, all Scrum teams consist of developers from multiple sites. While some Agile companies have created geographical transparency on a small scale, SirsiDynix uses fully integrated Scrum teams with over 50 developers in the U.S., Canada, and Russia. This strategy helped build a new implementation of platform and system architecture for a complex Integrated Library System (ILS). The ILS system is similar to a vertical market ERP system with a public portal interface used by more than 200 million people.

Best practices for distributed Scrum seen on this project consist of (1) daily Scrum team meetings of all developers from multiple sites, (2) daily meetings of Product Owner team (3) hourly automated builds from one central repository, (4) no distinction between developers at different sites on the same team, (5) and seamless integration of XP practices like pair programming with Scrum. While similar practices have been implemented on small distributed Scrum teams [19] this is the first documented project that demonstrates Scrum hyperproductivity for large distributed/outsourced teams building complex enterprise systems.

3. Distributed Team Models

Here we consider three distributed Scrum models commonly observed in practice.

Isolated Scrums - Teams are isolated across geographies. In most cases off-shore teams are not cross-functional and may not be using the Scrum process.

Distributed Scrum of Scrums – Scrum teams are isolated across geographies and integrated by a Scrum of Scrums that meets regularly across geographies.

Totally Integrated Scrums – Scrum teams are cross-functional with members distributed across geographies. In the SirsiDynix case, the Scrum of Scrums was localized with all ScrumMasters in Utah.

Most outsourced development efforts use a degenerative form of the Isolated Scrums model where

outsourced teams are not cross-functional and not Agile. Requirements may be created in the U.S. and developed in Dubai, or development may occur in Germany and quality assurance in India. Typically, cross-cultural communication problems are compounded by differences in work style in the primary organization vs. the outsourced group. In the worst case, outsourced teams are not using Scrum and their productivity is typical of waterfall projects further delayed by cross-continent communications lag time. Implementations of Scrum in a data rich CMMI Level 5 company simultaneously running waterfall, incremental, and iterative projects, showed productivity of Scrum teams was at least double that of waterfall teams, even with CMMI Level 5 reporting overhead [20]. Outsourced teams not using Scrum will typically achieve less than half the velocity of a primary site using Scrum.

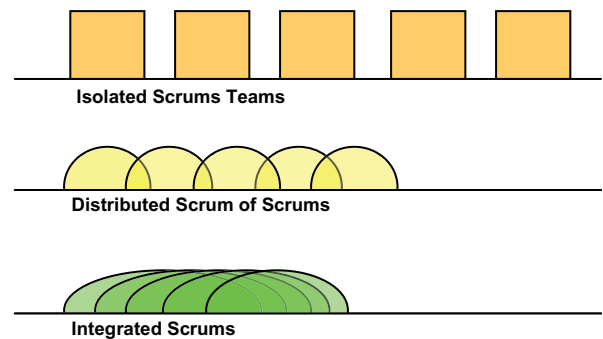


Figure 1: Strategies for distributed Scrum teams [21].

The latest thinking in the Project Management Institute Guide to the Project Management Body of Knowledge (PMBOK) models is a degenerative case of isolated non-Scrum teams [22]. This is a spiral waterfall methodology which layers the Unified Modeling Language (UML) and the Rational Unified Process (RUP) onto teams which are not cross-functional [23]. It partitions work across teams, creates teams with silos of expertise, and incorporates a phased approach laden with artifacts that violate the principles of lean development [24].

Best practice recommended by the Scrum Alliance is a Distributed Scrum of Scrums model. This model partitions work across cross-functional, isolated Scrum teams while eliminating most dependencies between teams. Scrum teams are linked by a Scrum-of-Scrums where ScrumMasters (team leaders/project managers) meet regularly across locations. This encourages communication, cooperation, and cross-fertilization and is appropriate for newcomers to Agile development.

An Integrated Scrums model has all teams fully distributed and each team has members at multiple locations. While this appears to create communication and coordination burdens, the daily Scrum meetings help to

break down cultural barriers and disparities in work styles. On large enterprise implementations, it can organize the project into a single whole with an integrated global code base. Proper implementation of this approach provides location transparency and performance characteristics similar to small co-located teams. A smaller, but similar, distributed team at IDX Systems Corporation during 1996-2000 achieved almost ten times industry average performance [19]. The SirsiDynix model approached this level of performance for distributed/outsourced Integrated Scrums. It appears to be the most productive distributed team ever documented for a large Java enterprise system with more than one million lines of code. This Integrated Scrums model is recommended for experienced Agile teams at multiple locations.

4. SirsiDynix Case Study

SirsiDynix has approximately 4,000 library and consortia clients, serving over 200 million people through over 20,000 library outlets in the Americas, Europe, Africa, the Middle East and Asia-Pacific. Jack Blount, President and CEO of Dynix and now CTO of the merged SirsiDynix company, negotiated an outsource agreement with StarSoft who staffed the project with over 20 qualified engineers in 60 days. Significant development milestones were completed in a few weeks and joint development projects are efficiently tracked and continue to be on schedule.

StarSoft Development Labs, Inc. is a software outsourcing service provider in Russia and Eastern Europe. Headquartered in Cambridge, Massachusetts, USA, StarSoft operates development centers in St. Petersburg, Russia and Dnepropetrovsk, Ukraine, employing over 450 professionals. StarSoft has experience handling development efforts varying in size and duration from just several engineers working for a few months to large-scale projects involving dozens of developers and spanning several years. StarSoft successfully uses Agile development and particularly XP engineering practices to maintain CMM Level 3 certification.

5. Hidden Costs of Outsourcing

The hidden costs of outsourcing are significant, beginning with startup costs. Barthelemy [25] surveyed 50 companies and found that 14% of outsourcing operations were failures. In the remainder, costs of transitioning to a new vendor often canceled out anticipated savings from low labor costs. The average time from evaluating outsourcing to beginning of vendor performance was 18 months for small projects. As a result, the MIT Sloan

Management Review advises readers not to outsource critical IT functions.

The German Institute for Economic Research analyzed 43,000 German manufacturing firms from 1992-2000 and found that outsourcing services led to poor corporate performance, while outsourcing production helped [26]. While this is a manufacturing study rather than software development, it suggests that outsourcing core development may provide gains not seen otherwise.

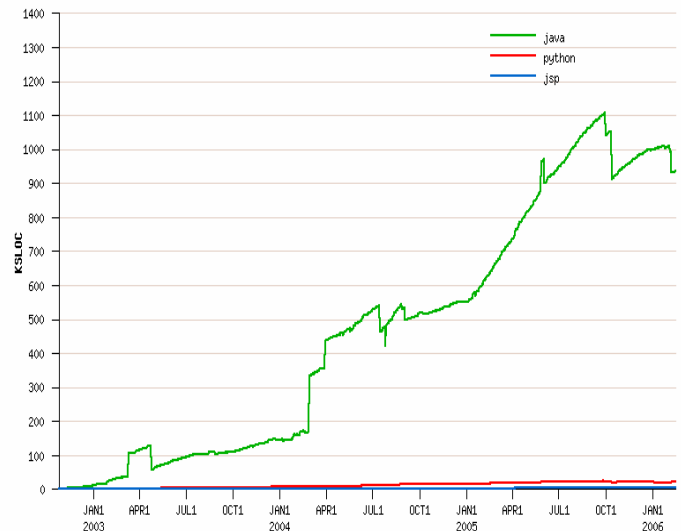


Figure 2 - SirsiDynix lines of new Java code in thousands from 2003-2006.

Large software projects are very high risk. The 2003 Standish Chaos Report show success rates of only 34%. 51% of projects are over budget or lacking critical functionality. 15% are total failures [27].

SirsiDynix sidestepped many of the hidden costs, directly outsourced primary production, and used Integrated Scrums to control risk. The goals of increasing output per team member and linearly increasing overall output by increasing team size were achieved. Production velocity more than doubled when the 30 person North American development team was augmented with 26 Russians from StarSoft in December 2005.

6. Intent of the Integrated Scrums Model

An Agile company building a large product and facing time-to-market pressure needs to quickly double or quadruple productivity within a constrained budget. The local talent pool is not sufficient to expand team size and salary costs are much higher than outsourced teams. On the other hand, outsourcing is only a solution if Agile practices are enhanced by capabilities of the outsourced teams. The primary driver is enhanced technical capability resulting in

dramatically improved throughput of new application functionality. Cost savings are a secondary driver.

7. Context

Software complexity and demands for increased functionality are exponentially increasing in all industries. When an author of this paper flew F-4 aircraft in combat in 1967, 8% of pilot functions were supported by software. In 1982, the F16 software support was 45%, and by 2000, the F22 augmented 80% of pilot capabilities with software [22]. Demands for ease of use, scalability, reliability, and maintainability increase with complexity.

SirsiDynix was confronted with the requirement to completely re-implement a legacy library system with over 12,500 installed sites. Large teams working over many years in a changing business environment faced many new requirements in the middle of the project. To complicate matters further, the library software industry was in a consolidating phase. Dynix started the project in 2002 and merged with Sirsi in 2005 to form SirsiDynix.

Fortunately, Dynix started with a scalable Agile process that could adapt to changing requirements throughout the project. Time to market demanded more than doubling of output. That could only happen by augmenting resources with Agile teams. StarSoft was selected because of their history of successful XP implementations and their experience with systems level software.

The combination of high risk, large scale, changing market requirements, merger and acquisition business factors, and the SirsiDynix experience with Scrum combined with StarSoft success with XP led them to choose an Integrated Scrums implementation. Jack Blount's past experience with Agile development projects at US Data Authority, TeleComputing and JD Edwards where he had used Isolated Scrums and Distributed Scrum of Scrums models did not meet his expectations. This was a key factor in his decision to structure the project as Integrated Scrums.

8. Forces

8.1. Complexity Drivers

The Systems and Software Consortium (SSCI) has outlined drivers, constraints, and enablers that force organizations to invest in real-time project management information systems. Scalable Scrum implementations with minimal tooling are one of the best real-time information generators in the software industry.

SSCI complexity drivers are described as [22]:

- Increasing problem complexity shifting focus from requirements to objective capabilities that must be met by larger teams and strategic partnerships.
- Increasing solution complexity which shifts attention from platform architectures to enterprise architectures and fully integrated systems.
- Increasing technical complexity from integrating stand alone systems to integrating across layers and stacks of communications and network architectures.
- Increasing compliance complexity shifting from proprietary to open standards.
- Increasing team complexity shifting from a single implementer to strategic teaming and mergers and acquisitions.

SirsiDynix faced all of these issues. Legacy products were difficult to sell to new customers. They needed a new product with complete functionality for the library enterprise based on new technologies that were highly scalable, easily expandable, and used the latest computer and library standards,

The SirsiDynix Horizon 8.0 architecture supports a wide variety of users from publication acquisition to cataloging, searching, reserving, circulating, or integrating information from local and external resources. The decision was made to use Java with J2EE, a modular design, database independency, maximum use of free platforms and tools, and wide support of MARC21, UNIMARC, Z39.50 and other ILS standards.

The project uses a three-tier architecture and Hibernate as a database abstraction layer. Oracle 10g, MS SQL, and IBM DB2 support is provided. The JBoss 4 Application server is used with a Java GUI Client with WebStart bootstrap. It is a cross-platform product supporting MS Windows 2000, XP, 2003, Red Hat Linux, and Sun Solaris. Built-in, multi-language support has on-the-fly resource editing for ease of localization. Other key technologies are JAAS, LDAP, SSL, Velocity, Xdoclet, JAXB, JUnit, and Jython.

8.2. Top Issues in Distributed Development

The SSCI has carefully researched top issues in distributed development [22], all of which had to be handled by SirsiDynix and StarSoft.

- **Strategic:** Difficult leveraging available resources, best practices are often deemed proprietary, are time consuming and difficult to maintain.
- **Project and process management:** Difficulty synchronizing work between distributed sites.
- **Communication:** Lack of effective communication mechanisms.
- **Cultural:** Conflicting behaviors, processes, and technologies.
- **Technical:** Incompatible data formats, schemas, and standards.

- **Security:** Ensuring electronic transmission confidentiality and privacy.

The unique way in which SirsiDynix and StarSoft implemented an Integrated Scrums model carefully addressed all of these issues.

9. Solution: Integrated Scrums

There are three roles in a Scrum: the Product Owner, the ScrumMaster, and the Team. SirsiDynix used these roles to solve the strategic distribution problem of building a high velocity, real-time reporting organization with an open source process that is easy to implement and low-overhead to maintain.

For large programs, a Chief ScrumMaster to run a Scrum of Scrums and a Chief Product Owner to centrally manage a single consolidated and prioritized product backlog is essential. SirsiDynix located the Scrum of Scrums and the Product Owner teams in Utah.

9.1. Team Formation

The second major challenge for large projects is process management, particularly synchronizing work between sites. This was achieved by splitting teams across sites and fine tuning daily Scrum meetings.

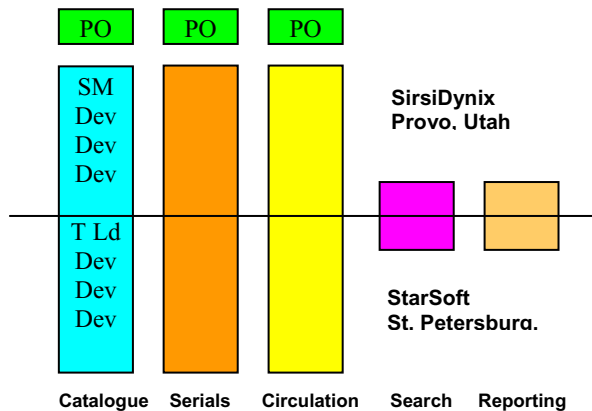


Figure 3 - Scrum teams split across sites. PO=Product Owner, SM=ScrumMaster, TLD=Technical Lead.

Teams at SirsiDynix were split across the functional areas needed for an integrated library system. Half of a Scrum team is typically in Provo, Utah, and the other half in St. Petersburg. There are usually 3-5 people on the Utah part of the team and 4 or more on the St. Petersburg portion of the team. The Search and Reporting Teams are smaller. There are smaller numbers of team members in Seattle, Denver, St. Louis, and Waterloo, Canada.

9.2. Scrum Meetings

Teams meet across geographies at 7:45am Utah time which is 17:45 St. Petersburg time. Teams found it necessary to distribute answers to the three Scrum questions in writing before the Scrum meeting. This shortens the time needed for the join meeting teleconference and helps overcome any language barriers. Each individual reports on what they did since the last meeting, what they intend to do next, and what impediments are blocking their progress.

Email exchange on the three questions before the daily Scrum teleconference was used throughout the project to enable phone meetings to proceed more smoothly and efficiently. These daily team calls helped the people in Russia and the U.S. learn to understand each other. In contrast, most outsourced development projects do not hold formal daily calls and the communication bridge is never formed.




Figure 4 – Scrum Team meetings

Local sub-teams have an additional standup meeting at the beginning of the day in St. Petersburg. Everyone uses the same process and technologies and daily meetings coordinate activities within the teams.

ScrumMasters are all in Provo, Utah or Waterloo, Canada, and meet in a Scrum of Scrums every Monday morning. Here work is coordinated across teams. Architects are directly allocated to production Scrum teams and all located in Utah. An Architecture group also meets on Monday after the Scrum of Scrums meeting and controls the direction of the project architecture through the Scrum meetings. A Product Owner resident in Utah is assigned to each Scrum team. A chief Product Owner meets regularly with all Product Owners to assure coordination of requirements.

SirsiDynix achieved strong central control of teams across geographies by centrally locating ScrumMasters, Product Owners, and Architects. This helped them get consistent performance across distributed teams.

9.3. Sprints

Sprints are two weeks long on the SirsiDynix project. There is a Sprint planning meeting similar to an XP release planning meeting in which requirements from User Stories are broken down into development tasks. Most tasks require a lot of questions from the Product Owners and some tasks take more time than initial estimates.

The lag time for Utah Product Owner response to questions on User Stories forces multitasking in St. Petersburg and this is not an ideal situation. Sometimes new tasks are discovered after querying Product Owners during the Sprint about feature details.

Code is feature complete and demoed at the end of each Sprint. Up until 2006, if it met the Product Owner's functional requirement, it was considered done, although full testing was not completed. It was not deliverable code until SirsiDynix strengthened its definition of "done" to include all testing in 2006. Allowing work in progress to cross Sprint boundaries introduces wait times and greater risk into the project. It violates the lean principle of reducing work in progress and increases rework.

9.4. Product Specifications

Requirements are in the form of User Stories used in many Scrum and XP implementations. Some of them are lengthy and detailed, others are not. A lot of questions result after receiving the document in St. Petersburg which are resolved by in daily Scrum meetings, instant messaging, or email.

Story for Simple Renewals Use Case:

Patron brings book or other item to staff to be renewed.

Patron John Smith checked out "The Da Vinci Code" the last time he was in the library. Today he is back in the library to pick up something else and brings "The Da Vinci Code" with him. He hands it to the staff user and asks for it to be renewed. The staff user simply scans the item barcode at checkout, and the system treats it as a renewal since the item is already checked out to John. This changes the loan period (extends the due date) for the length of the renewal loan. Item and patron circulation history are updated with a new row showing the renewal date and new due date. Counts display for the number of renewals used and remaining. The item is returned to Patron John Smith.

Assumptions:

- Item being renewed is currently checked out to the active patron
- No requests or reservations outstanding
- Item was not overdue

- Item does not have a problem status (lost, etc)
- No renew maximums have been reached
- No block/circulation maximums have been reached
- Patron's subscriptions are active and not within renewal period
- No renewal charges apply
- No recalls apply
- Renewal is from Check Out (not Check In)
- Staff User has renewal privileges

Verification (How to verify completion):

- Launch Check Out
- Retrieve a patron who has an item already checked out but not yet overdue
- Enter barcode for checked out item into barcode entry area (as if it is being checked out), and press <cr>.
- System calculates new due date according to circulation rules and agency parameters.
- The renewal count is incremented (Staff renewal with item)
- If user views "Circulation Item Details", the appropriate Renewals information should be updated (renewals used/remaining)
- Cursor focus returns to barcode entry area, ready to receive next scan (if previous barcode is still displayed, it should be automatically replaced by whatever is entered next)
- A check of the item and patron circulation statistics screens show a new row for the renewal with the renewal date/time and the new due date.

For this project, St. Petersburg staff likes a detailed description because the system is a comprehensive and complex system designed for specialized librarians. As a result, there is a lot of knowledge that needs to be embedded in the product specification.

The ways libraries work in St. Petersburg are very different than English libraries. Russian libraries operate largely via manual operations. While processes look similar to English libraries on the surface, the underlying details are quite different. Therefore, user stories do not have sufficient detail for Russian programmers.

9.5. Testing

Developers write unit tests. The Test team and Product Owners do manual testing. An Automation Test team in Utah creates scripts for an automated testing tool. Stress testing is as needed.

During the Sprint, the Product Owner tests features that are in the Sprint backlog. Up until 2006, testers received a stable Sprint build only after the Sprint demo. The reason for this was a lower tester/developer ratio than recommended by the Scrum Alliance.

There are 30 team members in North America and 26 team members in St. Petersburg on this project. The St. Petersburg team has one project leader, 3 technical team leaders, 18 developers, 1 test lead, and 3 testers. This low tester/developer ratio initially made it impossible to have a fully tested package of code at the end of the Sprints.

The test-first approach was initially encouraged and not mandated. Tests were written simultaneously with code most of the time. GUIs were not unit tested.

<i>Component</i>	<i>Test</i>	
	<i>Cases</i>	<i>Tested</i>
Acquisitions	529	384
Binding	802	646
Cataloging	3101	1115
Circulation	3570	1089
Common	0	0
ERM	0	0
Pac Searching	1056	167
Serials	2735	1714
Sub Total	11793	5115

Figure 5 – Test Cases Created vs. Tested

In the summer of 2006, a new CTO of SirsiDynix, Talin Bingham, took over the project and introduced Test Driven Design. Every Sprint starts with the usual Sprint Planning meeting and teams are responsible for writing functional tests before doing any coding. Once functional tests are written and reviewed, coding starts. Test-first coding is mandated. When coding is complete, developers run unit tests and manually pass all the functional tests before checking in changes to the repository.

Functional Area	Reserve Book Room
Task Description	Check that items from Item List is placed under Reserve with “Inactive” status
Condition	<ol style="list-style-type: none"> 1. User has right for placing Items under Reserve 2. At least one Item List exists in the system 3. Default Reserve Item Status in Session Defaults is set to “Inactive”
Entry Point	Launcher is opened
Test Data	No specific data
Action	<ol style="list-style-type: none"> 1. Reserve > Reserve Item 2. Select “Item Search” icon 3. Select “Item List” in the Combo box list of search options and enter appropriate Item list name 4. Press Enter

	<ol style="list-style-type: none"> 5. Select all Items which appear in the Item Search combo box and press “OK”
Expected Results	<ol style="list-style-type: none"> 1. Items that were in Item list should appear in the list in Reserve Item 2. Status of all items that has been just added should be shown as “Inactive” 3. Save button should be inactive 4. All corresponding Item should retain their original parameters

Figure 6 – Functional Test Example

Automation testing is done using the Compuware TestPartner tool, but there is still room for improvement of test coverage.

9.6. Configuration Management

SirsiDynix was using CVS as source code repository when the decision was made to engage an outsourcing firm. At that time, SirsiDynix made a decision that CVS could not be used effectively because of lack of support for distributed development, largely seen in long code synchronization times. Other tools were evaluated and Perforce was chosen as the best solution.

StarSoft had seen positive results on many projects using Perforce. It is fast, reliable and offers local proxy servers for distributed teams. Although not a cheap solution, it has been very effective for the SirsiDynix project.

Automated builds run every hour with email generated back to developers. It takes 12 minutes to do a build, 30 minutes if the database changes. StarSoft would like to see faster builds and true concurrent engineering. Right now builds are only stable every two weeks at Sprint boundaries.

9.7. Pair Programming, Refactoring, and other XP practices

StarSoft is an XP company and tries to introduce XP practices into all their projects. Pair programming is done on more complicated pieces of functionality. Refactoring was planned for future Sprints and not done in every iteration as in XP. Some radical refactoring without loss of functionality occurred as the project approached completion. Continuous integration is implemented as hourly builds. On this project, these three engineering practices were used with Scrum as the primary project management methodology.

9.8. Measuring Progress

The project uses the Jira <<http://www.atlassian.com>> issue tracking and project management software. This gives everyone on the project a real-time view into the state of Sprints. It also provides comprehensive management reporting tools. The Figure below shows the Sprint burn-down chart, a snapshot of Earned Business, and a synopsis of bug status.

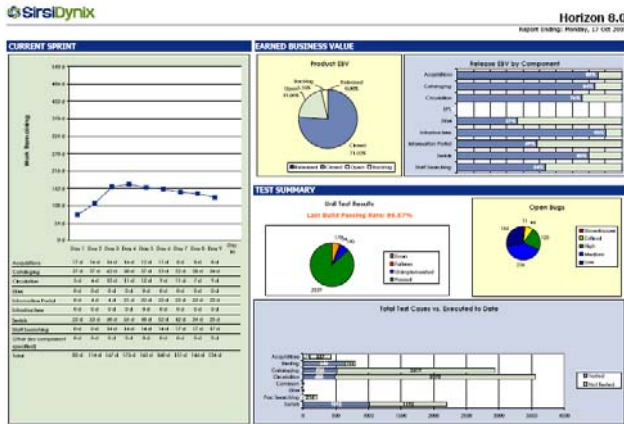


Figure 6 – SirsiDynix Horizon 8.0 Project Dashboard

Data from Jira can be downloaded into Excel to create any requested data analysis. High velocity projects need an automated tool to track status across teams and geographies. The best tools support bug tracking and status of development tasks in one system and avoid extra work on data entry by developers. Such tools should track tasks completed by developers and work remaining. They provide more detailed and useful data than time sheets, which should be avoided. Time sheets are extra overhead that do not provide useful information on the state of the project, and are de-motivating to developers.

Other companies like PatientKeeper [28] have found tools that incorporate both development tasks and defects that can be packaged into a Sprint Backlog are highly useful for complex development projects. Thousands of tasks and dozens of Sprints can be easily maintained and reviewed real-time with the right tool.

10. Integrated Scrums Model Resulting Context

Collaboration of SirsiDynix and StarSoft turned the Horizon 8.0 project into one of the most productive Scrum projects ever documented. For example, data is provide in the table below on a project that was done initially with a waterfall team and then re-implemented with a Scrum team [29]. The waterfall team took 9 months with 60 people and generated 54000 lines of code.

It was re-implemented by a Scrum team of 4.5 people in 12 months. The resulting 50,803 lines of code had more functionality and higher quality.

	SCRUM	Waterfall	SirsiDynix
Person Months	54	540	827
Java LOC	50,803	54000	671,688
Function Points	959	900	12673
FP per dev/month	17.8	2.0	15.3
FP per dev/month (industry average)	12.5	12.5	3

Figure 7 – Function Points/Developer Month for collocated vs. distributed projects.

Capers Jones of Software Productivity Research has published extensive tables on average number of function points per lines of code for all major languages [30]. Since the average lines of code per function point for Java is 53, we can estimate the number of function points in the Scrum application. The waterfall implementation is known to have fewer function points.

Distributed teams working on Horizon 8.0 generated 671,688 lines of code in 14.5 months with 56 people. During this period they radically refactored the code on two occasions and reduced the code based by 275,000. They have not been penalized for refactoring as that is rarely done in large waterfall projects in the database from which Capers derived his numbers. They have also not been rewarded for refactoring even though reducing lines of code is viewed as important as adding new code on well-run Agile projects.

Jones has also shown from his database of tens of thousands of projects that industry average productivity is 12.5 function points per developer/month for a project of 900 function points and that this drops to 3 for a project with 13000 function points [31]. Some of this is due to 4GL and other code-automation tools used on small projects, many of which are not implemented in third generation languages like Java.

The SirsiDynix project is almost as productive as the small Scrum project with a collocated team of 4.5 people. For a globally dispersed team, it is one of the most productive projects ever documented at a run rate of five times industry average.

11. Conclusion

This case study is a proof point for the argument that distributed teams and even outsourced teams can be as productive as a small collocated team. This requires excellent implementation of Scrum along with good engineering practices. The entire set of teams must function as a single team with one global build repository, one tracking and reporting tool, and daily meetings across geographies.

Outsourced teams must be highly skilled Agile teams and project implementation must enforce geographic transparency with cross-functional teams at remote sites fully integrated with cross-functional teams at the primary site. In the SirsiDynix case, the teams were all run from a central site giving strong central control.

It is highly unlikely that distributed outsourced teams using current Agile Alliance best practices of distributing work to independent Scrum teams across geographies could achieve the level of performance achieved in this case study. Therefore, SirsiDynix sets a new standard of best practices for distributed and outsourced teams with a previously demonstrated high level of Agile competence.

12. References

- [1] C. G. Langton, "Life at the Edge of Chaos," in *Artificial Life II, SFI Studies in the Sciences of Complexity*, Held Feb 1990 in Sante Fe, NM, 1992, pp. 41-91.
- [2] R. A. Brooks, "Intelligence without representation," *Artificial Intelligence*, vol. 47, pp. 139-159, 1991.
- [3] H. Ziv and D. Richardson, "The Uncertainty Principle in Software Engineering," in *submitted to Proceedings of the 19th International Conference on Software Engineering (ICSE'97)*, 1997.
- [4] W. S. Humphrey, *A Discipline for Software Engineering*: Addison-Wesley, 1995.
- [5] P. Wegner, "Why Interaction Is More Powerful Than Algorithms," *Communications of the ACM*, vol. 40, pp. 80-91, May 1997.
- [6] P. DeGrace and L. H. Stahl, *Wicked problems, righteous solutions : a catalogue of modern software engineering paradigms*. Englewood Cliffs, N.J.: Yourdon Press, 1990.
- [7] Matisse Software, "The Emergence of the Object-SQL Database," Mountain View, CA 2003.
- [8] F. P. Brooks, *The Mythical Man Month: Essays on Software Engineering*: Addison-Wesley, 1995.
- [9] K. Beck, *Extreme Programming Explained: Embrace Change*. Boston: Addison-Wesley, 1999.
- [10] W. A. Wood and W. L. Kleb, "Exploring XP for Scientific Research," *IEEE Software*, vol. 20, pp. 30-36, May/June 2003.
- [11] C. Larman, *Agile & Iterative Development: A Manager's Guide*. Boston: Addison-Wesley, 2004.
- [12] H. Takeuchi and I. Nonaka, "The New New Product Development Game," *Harvard Business Review*, 1986.
- [13] P. M. Senge, *The Fifth Discipline: the Art and Practice of the Learning Organization*. New York: Currency, 1990.
- [14] M. Beedle, M. Devos, Y. Sharon, K. Schwaber, and J. Sutherland, "Scrum: A Pattern Language for Hyperproductive Software Development," in *Pattern Languages of Program Design*, vol. 4, N. Harrison, Ed. Boston: Addison-Wesley, 1999, pp. 637-651.
- [15] S. J. Gould, *The structure of evolutionary theory*. Cambridge, Mass.: Belknap Press of Harvard University Press, 2002.
- [16] W. D. Hillis, *The Connection Machine*. Cambridge, MA: MIT Press, 1985.
- [17] S. Levy, *Artificial Life : A Report from the Frontier Where Computers Meet Biology*, 1st ed. New York: Vintage, Reprint edition, 1993.
- [18] D. K. I. Sobek, A. C. Ward, and J. K. Liker, "Toyota's Principles of Set-Based Concurrent Engineering," *Sloan Management Review*, vol. 40, pp. 67-83, 1999.
- [19] J. Sutherland, "Agile Can Scale: Inventing and Reinventing Scrum in Five Companies," *Cutter IT Journal*, vol. 14, pp. 5-11, 2001.
- [20] C. Jakobson, "The Magic Potion for Code Warriors! Maintaining CMMI Level 5 Certification With Scrum.," J. Sutherland, Ed. Aarhus, Denmark: Agile 2007 paper in preparation, 2006.
- [21] H. Takeuchi and I. Nonaka, *Hitotsubashi on Knowledge Management*. Singapore: John Wiley & Sons (Asia), 2004.
- [22] K. E. Nidiffer and D. Dolan, "Evolving Distributed Project Management," *IEEE Software*, vol. 22, pp. 63-72, Sep/Oct 2005.
- [23] R. Zanoni and J. L. N. Audy, "Projected Management Model for Physically Distributed Software Development Environment," in *HICSS'03*, Hawaii, 2003, p. 294.
- [24] M. Poppendieck, "A History of Lean: From Manufacturing to Software Development," in *JAOO Conference*, Aarhus, Denmark, 2005.
- [25] J. Barthelemy, "The Hidden Costs of Outsourcing," *MITSloan Management Review*, vol. 42, pp. 60-69, Spring 2001.
- [26] B. Gorzig and A. Stephan, "Outsourcing and Firm-level Performance," German Institute of Economic Research October 2002.
- [27] StandishGroup, "2003 Chaos Chronicles," The Standish Group International, 2003.
- [28] J. Sutherland, "Future of Scrum: Parallel Pipelining of Sprints in Complex Projects with Details on Scrum Type C Tools and Techniques," PatientKeeper, Inc., Brighton, MA May 30 2005.
- [29] M. Cohn, *User Stories Applied : For Agile Software Development*: Addison-Wesley, 2004.
- [30] C. Jones, "Programming Languages Table, Release 8.2," Software Productivity Research, Burlington, MA 1996.
- [31] C. Jones, *Software assessments, benchmarks, and best practices / Capers Jones*. Boston, Mass.: Addison Wesley, 2000.