

Being Extreme in the Classroom: Experiences Teaching XP

Alfredo Goldman Fabio Kon Paulo J. S. Silva
Department of Computer Science
University of São Paulo, Brazil
{gold,kon,rsilva}@ime.usp.br
<http://www.ime.usp.br/~xp>

Joseph W. Yoder
The Refactory, Inc. and
Department of Computer Science
University of Illinois at Urbana-Champaign
joeyoder@joeyoder.com
<http://www.refactory.com>

Abstract

Agile Methods propose a new way of looking at software development that questions many of the beliefs of conventional Software Engineering. Agile methods such as Extreme Programming (XP) have been very effective in producing high-quality software in real-world projects with strict time constraints.

Nevertheless, most university courses and industrial training programs are still based on old-style heavyweight methods. This article, based on our experiences teaching XP in academic and industrial environments, presents effective ways of teaching students and professionals on how to develop high-quality software following the principles of agile software development. We also discuss related work in the area, describe real-world cases, and discuss open problems not yet resolved.

1 Introduction

In the last few years, agile software development methods have become widely known and have been successfully adopted by hundreds of organizations worldwide. Agile methods such as XP [4], Scrum [27], and Crystal [10] are now used in small, medium, and large companies, universities, research institutes,

and governmental agencies. However, a large number of organizations have a long history of using old-style, heavyweight methods and many programmers and managers were educated to develop software in a bureaucratic way in which software quality is usually not the top priority.

The Manifesto for Agile Software Development [7] indicates the four most important aspects of agile methods that differentiate them from conventional software development. Agile methods value:

- **Individuals and interactions** over processes and tools;
- **Working software** over comprehensive documentation;
- **Customer collaboration** over contract negotiation;
- **Responding to change** over following a plan.

While the items on the left (in bold face) are the core principles of a successful agile software development project, most programmers and managers are educated in a culture that values the items on the right.

After participating in agile software development projects, a large number of academic and industrial software developers have come to the conclusion that agile methods are the most effective way for developing high-quality software in time-constrained projects. There are plenty of examples of successful

agile projects [2, 18, 24, 8, 15, 1, 19, 13, 17, 26]. However, there is still a lot of work to do in the field of teaching agility. According to our observations over the last ten years, most undergraduate Computer Science courses and training courses for industry are based on conventional methods that focus on tools, documentation, contracts, and following plans.

What CS education needs is a reality shock! We need to modernize our courses to show students that personal communication, working software, customer collaboration, and dynamic adaptation are, at least, as important as the traditional values that we are used to teaching.

In this article, we present our findings in agile methods education from our experiences in teaching Extreme Programming (XP) in academic and industrial environments. In Section 2, we describe a few related works in this area. In Section 3, we give a brief overview of XP and discuss the points in favor and against XP, addressing when one should and should not use this method. In Section 4, we describe the adaptations that must be made in order to carry out an XP project in an educational environment. In Section 5, we describe some of our experiences in teaching XP both in the university and in the industry and, in Section 6, we discuss problems that are still unresolved. Conclusions are presented in the last section.

2 Related Work

Many articles have recently emerged in the literature on the subject of teaching XP or some of its practices. However, special care must be taken if an instructor desires to teach or apply isolated practices. XP is clearly based on the synergy resulting from the combined use of all of 12 practices together. Applying only a subset of the practices may, in some cases, lead to disastrous results¹. Among the individual practices, two emerge as being beneficial even when detached from the others: pair programming and test-driven development.

Pair programming has been extensively studied by Williams et. al. [33, 9, 32]. Their experimental research shows that pair programming is highly efficient. Working in pairs, programmers can work twice as fast and think of more than twice as many solutions to a problem than two programmers working alone. Pair

¹Consider, for example, a team doing extensive refactoring without automated tests and pair programming to act as a safety net. It is very likely that refactoring will add a large number of bugs to the code

programming can also help attain higher defect prevention since two pairs of eyes are looking over every line of code. The results can be the faster development of higher quality production code. Experiments were performed with senior-level students in a Software Engineering course at the University of Utah and confirmed previous anecdotal evidence. Tomayko also presents an interesting experiment [30] indicating that pair programming can bring to code quality the same benefits that a method centered on formal inspections, such as the Team Software Process (TSP), would bring.

Nagappan et. al. [25] focus on the pedagogical benefits of pair programming. It stimulates cooperation, which is highly regarded in the working environment, but somehow neglected in the academic setting. Pair programming in the academic environment can lead to improved success rates, that is the rate of students that complete the course with a grade of C or better. The future performance of the paired students has shown to be about the same as the solo ones. This last observation addresses the concerns of some instructors that some paired students might pass the workload to his/her partner and not learn the course material. Another interesting benefit of pair lab sessions is that the students solve, with their partners, most of the simple questions, alleviating the lab instructors workload. Finally, the authors give some suggestions on how to use pair programming effectively in the teaching lab. For example, they suggest that the instructors must constantly emphasize the different roles in the pair (driver and navigator) and encourage role and pair changes. Another interesting advice is to use some kind of peer evaluation to avoid the “free ride” on the partner’s work.

Test-driven development [6] is another practice that fits easily in many different development methods and, hence, can be taught separately from the other practices. Actually, test-driven development should be taught before debugging techniques, as a better way to avoid and catch errors. The instructor must be prepared to face resistance from many students who feel that tests are a waste of time. Here some good anecdotal stories may be beneficial, such as McBreen’s testimony that he moved from daily debugging sections to only three or four sessions a year after adopting XP style unit tests [21].

Finally, refactoring [20] may prove beneficial to teach to advanced students. However, caution has to be taken since this practice is closely related to other XP ideas such as incremental development, not adding unneeded flexibility, tests, and continuous integration.

Astrachan, Duvall and Wallingford [3] report on teaching and using just a few practices, suggesting, for example, that teaching refactoring principles should always include information on testing. A panel discussion held recently at the ACM Technical Symposium on Computer Science Education addressed the topic of gradually introducing some of the XP practices into the CS curriculum [14].

Other works describe experiences on teaching full XP. Tomek [31] presents his experience on teaching XP in two Computer Science courses and proposes several recommendations. According to his experience, it is very important to provide a very agile environment, therefore he used VisualWorks Smalltalk and its IDE. In his first course, Tomek used two projects. The first project was used to get the feeling of XP and was followed by a more realistic one. During the second course, he focused on a single project with a real customer.

Wilson's experience [34] is similar, although he used Java instead of Smalltalk. The course project was to improve a prototype Java IDE. Finally, Lappo [16] taught an eXtreme Programming course to a group of Masters students, who spent 12 weeks working full-time to produce a Web-based resource management application with Java technologies. In all these cases, the instructors accumulated the role of mentor, coach, and occasionally, even customer.

All three experiments presented small problems such as the lack of real customers, the overload of being mentor, coach, and client, the short period that can be dedicated to the project in ordinary CS courses, and the lack of adequate space for the XP team.

At this point, a few suggestions can help: students enjoy meaningful projects that have real use, students should be relatively advanced to be able to get the most out of the course, it would be ideal to provide a dedicated room for the XP class, and it might be interesting to provide the students with an implemented core of the application. Wilson suggests that this prototype can work to give unity to the end result substituting the metaphor, which is one of the most difficult practices to teach.

The courses also revealed some problems in trying to follow all XP practices. In some cases pair programming was partially neglected, even by the instructor that refactored the code alone. Other problems were the lack of a real client, the absence of a metaphor, and some slips in the release schedule. Interestingly, it seems like test-first development was quickly assimilated by the students, who learned to appreciate its advantages.

In Sections 4 and 5, we show how to avoid many of the problems described above. In particular, we show effective ways to teach XP, putting into practice most or all of the practices in both industrial and academic environments, giving a complete XP experience to students.

3 Overview of Extreme Programming

The Extreme Programming method was formulated by Kent Beck and Ward Cunningham based on their long experience in object-oriented software development in Smalltalk. XP is composed of a collection of practices that, in isolation, have been well known and used widely for many years. The main contribution of XP is the conjunction of these practices in a cohesive method, which fosters the synergistic effects of this mixture.

3.1 The 12 Practices

When the method was first introduced in 1999 [4], it consisted of 12 practices: Planning Game, Small Releases, Metaphor, Simple Design, Testing, Refactoring, Pair Programming, Collective Code Ownership, Continuous Integration, 40-hour Week, On-site Customer, and Coding Standards. A few years later, the 40-hour Week practice was renamed to Sustainable Pace and a new rule was added: *Fix XP When It Breaks*. We will now describe briefly each of these practices, for more detailed descriptions see [4] and www.extremeprogramming.org.

Planning Game. A project starts with a short exploratory phase in which the customer expresses the requirements (through user stories written in *story cards*). The development team, together with the customer, creates a release plan specifying which story cards should be implemented for each system release. The team negotiates with the customer and dates each release based on business priorities and technical estimates. However, the most important point here is that the plan is just a plan, i.e., the team and the customer know that it is not the reality that they will face. As reality overtakes the plan, the plan must be updated. So, rather than trying to put together an unrealistic completed plan upfront, planning becomes part of an everyday activity for an XP project. A good XP team must know how to adapt dynamically to changes at any moment in the development process.

Small Releases. Rather than developing big pieces of software at a time, the team should implement a very small piece of working software first and then enhance it incrementally. Ideally, the team should deliver new releases of working software every few weeks or, in some cases, days. The time between releases should not exceed every few weeks.

In each release, the team implements a set of story cards. Each story card is assigned to a specific programmer who becomes responsible for its completion (although the programmer does receive help from his/her colleagues to complete his/her tasks). Stories that are more important to the customer receive a higher priority and are implemented in the first releases. Developers and customer may negotiate during development to move cards from one release to the other or to create, remove or modify them as the team learns new things and business requirements evolve.

A key rule of incremental development in XP is: do not code for the future and do not anticipate requirements. This spirit is usually expressed in the sentence *do the simplest thing that could possibly work*. This implies that one should not add flexibility that is not needed to complete the current task. If you think that a little more flexibility will be valuable in a couple of weeks, don't do it now; wait until it is really needed and then refactor the code to add the required flexibility.

Metaphor. A metaphor is a simple story of how the system works and it should be shared by all the stakeholders in the project. It helps all the participants to understand the basic elements and their relationships. The primary function of a metaphor is to improve communication.

Simple Design. The system should always have the simplest possible design at any moment. If extra complexity is found, it must be removed as soon as possible. And again: *do the simplest thing that could possibly work*.

Testing. Programmers write unit tests for all system components so their confidence in the correct behavior of the system becomes part of the system itself. In a more recent book [6], Beck describes *Test-Driven Development* in which the unit tests are written before the actual working code is written, which is also called *Test-First Programming*.

Customers write functional (acceptance) tests which are used to demonstrate that the required features are implemented correctly. If the customer is not

a programmer, one of the developers then pairs with the customer to write the tests.

Refactoring. Using techniques such as the ones cataloged by Fowler [12], programmers restructure the system continuously to improve it without changing its behavior. Possible improvements include simplifications, optimizations, enhancing clarity, adding flexibility, etc.

Pair Programming. Each line of production code is written with two programmers simultaneously at a single machine. As explained in Section 2, pair programming improves code quality greatly without impacting the speed of development. Communication will flow better across team members if the pairs change frequently (e.g., every day). The pairs are selected based not only on availability but also in expertise. For example, if it is necessary to build a Web interface to a database, one could select a pair in which one of the programmers is an expert in databases while the other is an expert in frameworks for building Web interfaces.

Collective Code Ownership. Any developer can change any piece of code in the system at any time without requesting permission. This introduces a high level of agility in the team. Since there are unit tests for each component, programmers are less likely to break each other's code.

Continuous Integration. The source code must be kept in a shared repository and every time a task is completed, the new code must be built, tested, and, if correct, integrated into the repository.

Sustainable Pace. The team should work in a pace that it can sustain without harming its participants, for example, 40 hours per week. A team that is physically or intellectually tired is very likely to produce low-quality software. Working overtime in a certain, special week is acceptable; however, if the team is asked to work overtime two or more weeks in a row, this is a sign that there is something very wrong with the project.

On-site Customer. A real user of the system should be included as part of the team and be available any time for answering questions. No matter what happens to the project (good or bad), there will never be a big surprise to the customer since he/she is following the development daily.

Coding Standards. In the initial phase, all the developers must agree on a common set of rules enforcing how the system must be coded. This facilitates communication and enable groups of many programmers to produce consistent code. Recent tools such as the Eclipse *Checkstyle* plug-in (see eclipse-cs.sourceforge.net) can help automate part of the process.

It is important to emphasize that the value of XP is in applying *all* the practices. Applying a subset of the practices, without careful consideration, can even be harmful. For example, applying aggressive refactoring without a good collection of unit tests may lead to disastrous results as the programmers cannot verify if their changes are breaking the code or not. Adopting the planning game, changing the plan dynamically, without a close contact with the customer may lead the team to build a system that is not the one the customer wants.

3.2 Adapting XP

Teams that are new to XP should try to follow all 12 practices as rigorously as possible. More experienced XP developers, however, will notice that this may not be possible, or even desirable, in all situations. When this happens you may need to adapt XP by applying the *Fix XP When It Breaks* rule.

We present an example to illustrate this rule. An experienced XP developer, Klaus Wuestefeld, working on a project for a cable-TV scheduling system realized that he would not be able to have an on-site customer since the company contracting their services was located in another state. The solution was to adapt XP introducing the concept of *Customer Proxy* [35]. Klaus acted as a customer by answering programmer questions immediately. He would then call the real customers on the phone or email them later with the questions verifying his response. Most of the time the proxy's guesses were correct and the development evolved quickly. The few times in which he made the wrong guess, he simply came back to the programmers and said: "I changed my mind", which is completely acceptable within the rules of XP.

Another possibility is to select analysts that have worked closely with the customers to become the customer proxy. This adaptation has been reported by Martin Fowler in projects carried out by Thoughtworks.

In another project, Klaus noticed that developers were worried too much about the story cards assigned

to them and were not always willing to help their colleagues by pair programming with them². The solution he adopted in that case was to create a new role: the *Liberero* [35]. One of the programmers, called the *Liberero*, was not assigned any story card; his task was simply to pair program with the others helping them finish their cards.

A limitation of XP is that, since it requires direct communication among all team members, it does not scale well for groups with much more than 10 developers. To overcome this limitation, practitioners have extended the method to work with larger projects of up to 100 developers. This was achieved by dividing the team into sub-groups of at most 10 people and integrating periodically the software produced by each of the groups. Ron Crocker has worked many years with large-scale agile projects for Motorola. His extension of XP is called the Grizzly method and a new book on the subject is coming out in 2004 [11].

As a last example, sometimes teams allow for more individual spike solutions to be developed. Then, these solutions are released into the main code base only after test cases are developed and a pair of eyes looks over the solution. This can be a solution when pair programming is not always possible.

Taber and Fowler [29] provide a few more examples of interesting adaptations of XP.

3.3 When not to use XP

There are some situations when using XP should be avoided. The possible pitfalls for XP adoption fall in three categories: resistance from the development team to embrace XP, resistance from the organization that houses the XP team or from the customer to accept XP corollaries, and inadequacies inherent to the software that have to be developed. McBreen has recently written an interesting book on this subject called *Questioning Extreme Programming* [21].

The resistance from the development team could have come from habits acquired during the team member's life as a programmer. At first glance, even pair programming may be odd, test-driven development a burden, and simple design seen as an excuse for lazy minds. After an adaptation phase, however, many developers learn to appreciate the practices and the agile development environment.

On the other hand, the resistance from the organization and its managers may be associated with one of XP most profound facets: XP is a subversive method,

²This is actually not very common in XP projects; usually, programmers negotiate among themselves to help each other implementing their cards.

in the sense that it requires a completely new organization of the team. In traditional water-fall methods there are distinct and well defined roles for the team members such as requirement analysts, system analysts, software architects, programmers, testers, and so on. In XP, all team members play all these roles, they are all *developers*. This creates a completely new balance of power within the team that may face great resistance [21, 28]. In order to adopt XP, the team must feel comfortable with the idea of working together as a group with the single goal of delivering high-quality software in time.

To overcome this resistance, the XP instructor or mentor should pick as members for the first XP experiment a group of people that is naturally inclined to experiment with new ideas and that are self-confident enough not to feel threatened by the new balance of power. To achieve this in industry, it is essential to have management support. This kind of support can help the mentor to identify good candidates for applying XP. In the university setting, this is better achieved by using elective courses. After the first successful XP experience, the word of mouth of the participants will spread the news and it will be much easier to introduce XP into the entire organization or to make the course mandatory for all students.

Examples of resistance from the organization and/or the customer are the lack of commitment by the customer in participating actively in the development process, a requirement for long and formal descriptions of the product to be developed before it is begun, the need to have a single person to blame should anything go wrong³, and the need for a long, detailed documentation for the maintenance phase [4, 21]. All these XP consequences must be clearly stated and understood both by managers and developers before starting the first XP project in an organization.

Finally XP is not meant for all software development projects. Certain aspects of the software to be developed may conflict with basic XP assumptions. Does the project require a very large team (e.g., more than 20 people)? Does the edit-compile-run cycle take too long to complete? Do tests need several minutes or even hours to run? Is it impossible to find an on-site customer that will faithfully represent the future users of the system? If the answer is yes to any of these questions, then XP might not be a good choice for this project or the method might need to be adapted significantly [4, 21].

It should also be noted that an organization should

³That does not go well with collective code ownership.

not start learning and experimenting with XP with a highly visible or mission critical project. Any kind of difficulty could lead people to blame the XP process. Or the temptation might be to adapt the process before giving it a chance to succeed. Therefore, it is better to learn XP in a safer and less visible project, thus gaining credibility before applying it to more visible mission critical projects.

4 Adapting XP for the Classroom

In an educational environment, not all of the aspects of a real production environment are present. Thus, when teaching XP in the university or in corporate training, some adaptations are required.

Different from most academic courses, an XP course must focus on practice rather than on theory. Students must spend most of the time programming in the lab, not attending lectures. We identified two types of courses that can produce good results: short courses and long courses.

In an academic environment, a long course would typically be a full-semester course in which the students attend, initially, a few lectures describing the method and then spend three to four months working in the lab, two to four sessions per week. Since the students work less hours in the project than a corporate employee, each release should contain only a few story cards (e.g., five to ten) and should be planned to be completed in approximately four weeks.

A short course in an academic environment can range from a full-day, six hour workshop in which the students are exposed to both theoretical and practical aspects of XP, up to a one-month Summer course in which more details can be covered.

In industrial environments, the long “course” takes the form of *mentoring*. In this case, an experienced XP consultant spends several hours per week working in a real project of interest to the company, acting as the team *coach*. The role of the coach is not to guide the development but to make sure that all XP practices are being followed and to use his/her experience to resolve conflicts and show the group how XP can help overcoming the difficulties that arise. After a few months, the role of coach can be handed over to one of the developers and the consultant becomes a *meta-coach*, gradually decreasing his/her responsibilities. It is often said that the job of an XP consultant is to put him/herself out of business in the long run by empowering the team to work by itself using XP.

Short courses in industrial environments typically take the form of immersion workshops in which devel-

opers spend 2 to 4 days working full-time in a simple project going through all the steps of an XP project, producing a few releases of working software.

Except for the mentoring case, which can mimic a production environment perfectly, the other cases may require some adaptations. The time span of the courses is very different from that of real software development projects. There might not be a real customer available. The same person (e.g., instructor or professor) may need to play the role of both coach and customer, which is probably not a good idea.

All these issues must be analyzed carefully by the instructor to enable the course participants to have an XP experience as real as possible so that they will be capable of applying the method in real life afterwards.

Our experience shows that, with proper planning, it is possible to overcome all of these difficulties and to provide a real XP experience to students. In the next section we describe some of the long and short courses we carried out in both academic and industrial environments.

5 Experiences Teaching XP

We have had many experiences from 2000 to 2004 in teaching XP at the University of Illinois at Urbana-Champaign, at the University of São Paulo⁴, and in work as consultants both in the United States and in Brazil. This wide variety of experience gives us a broad view of what it takes to teach XP principles, including teaching XP to many types of people coming from different cultures.

In this section, we describe our experiences from a full-semester course at the University of São Paulo, from mentoring for the development of a Web-based commercial product, from consulting for the Illinois Department of Public Health, and from a short-term course for a private company.

5.1 University of São Paulo

We started to disseminate the use of XP in Brazil in early 2001 with a series of lectures about different aspects of XP. These lectures included an overview of the method, refactoring, debugging, testing, and coding style⁵. In addition to these individual lectures, we

⁴The University of Illinois at Urbana-Champaign and the University of São Paulo are top 5 universities in the USA and Brazil, respectively. The students in their CS departments are in general very well prepared since there is a highly competitive student selection process.

⁵Slides, in Portuguese, available at <http://www.ime.usp.br/~xp>.

have been presenting an annual 4-month course called *Extreme Programming Laboratory*. The course is targeted to undergraduate students in the 3rd and 4th year of the Bachelors program in Computer Science. Course attendance is limited to 20 students and they are divided in groups of 6 to 10 students. This is *per se* something new for the students since they usually have not had an opportunity to work in such a large group, where all the participants work together. In fact, most courses discourage students working together. The rest of this section will describe the most important aspects that must be addressed when implementing such a course.

Workload. The students were required to be in the lab during two weekly sessions lasting two to three hours. We found that three hour sessions are much more productive. However, due to schedule restrictions, we were sometimes forced to have 2-hour sessions. A good way of keeping the students for a longer period in the lab was to provide a modest lunch. Then the students stay focused on programming for longer periods, holding their sandwiches while pair programming. Having food around a software development lab relaxes and allows developers to communicate better [4].

Besides these two mandatory sessions, it was highly suggested to the students that they come to the lab two to four additional hours per week for either pair programming or to learn about the technology used in their projects. These additional hours were not verified by the instructors.

Development site. The laboratory where the student projects were developed was set up following the guidelines to enable a high level of osmotic communication suggested by Alistair Cockburn [10]. This enabled light communication channels across teams as well as strong communication channels among the members of the same team. Cockburn has studied and experimented with many different room layouts and identified their advantages and drawbacks (see [10], Chapter 3, *Communicating, Cooperating Teams*).

The University of São Paulo lab, where the XP courses are carried out, was reorganized to follow these guidelines. As shown in Figure 1, the workstations are arranged so there is space for two people sitting in front of each computer, and the members of the team sit facing each other. This contrasts with many laboratories where the developers face a wall or in which workstations are separated by dividers or enclosed in cubicles. The two groups working in the



Figure 1: XP students in the lab

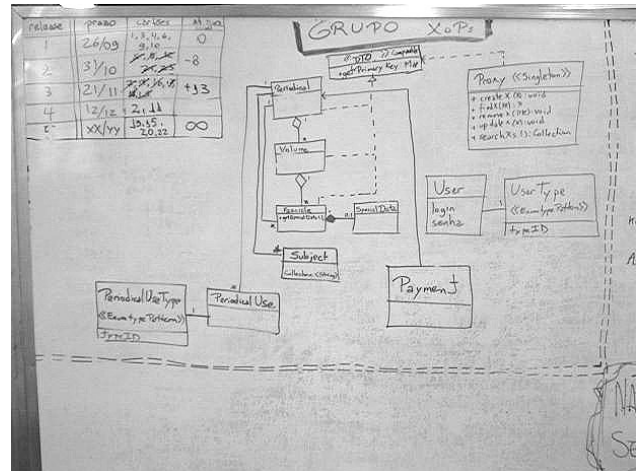
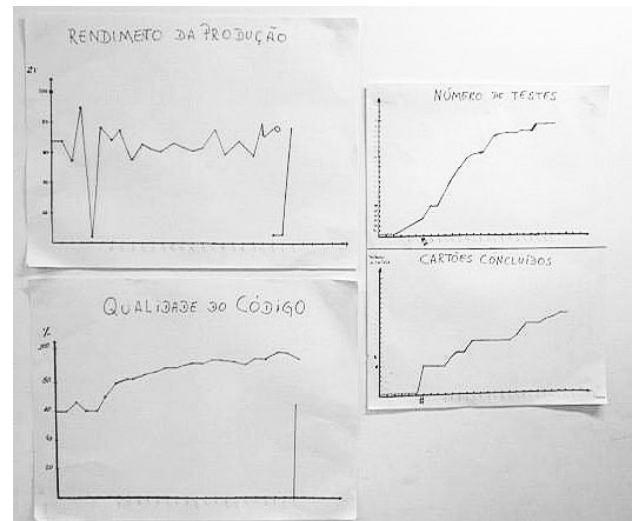


Figure 2: Whiteboard dividing the space



XOPS

Card#	Descrição	Responsável	Início	Previsão	Fim	Status
1	Atualiza de Requisitos	Stefan Dombi	08/09	18/09	26/09	OK
2	Atualiza de Requisitos	Luiz	08/09	22/09	26/09	OK
4	Atualiza de Requisitos	Luiz	08/09	22/09	26/09	OK
6	Atualiza de Requisitos	Miguel	08/09	22/09	26/09	OK
9	Atualiza de Requisitos	Miguel	11/09	25/09	26/09	OK
10	Atualiza de Requisitos	Miguel	11/09	25/09	26/09	OK
24	Atualiza de Requisitos	Stefan Dombi	08/10	13/10	15/10	OK
8	Atualiza de Requisitos	Luiz	13/10	21/10	23/10	OK
5	Atualiza de Requisitos	Miguel	08/10	09/10	09/10	OK
12	Atualiza de Requisitos	Miguel	13/10	16/10	21/10	OK
23	Atualiza de Requisitos	Stefan	06/10	09/10	09/10	OK
7	Atualiza de Requisitos	Luiz	07/10	17/10	15/10	OK
14	Atualiza de Requisitos	Stefan	30/10	11/11	19/11	OK
16	Atualiza de Requisitos	Stefan	30/10	17/11	12/11	OK
17	Atualiza de Requisitos	Miguel	17/11	17/11	19/11	OK
18	Atualiza de Requisitos	Miguel	17/11	21/11	01/12	OK
19	Atualiza de Requisitos	Stefan	17/11	21/11	01/12	OK
20	Atualiza de Requisitos	Miguel	02/12			OK

Figure 3: Information radiators maintained by the tracker

Coaching. We learned that choosing a good coach is very important for the success of an XP course. Over the years, we tried three different options for coach: a professor knowledgeable in XP, a graduate student that had attended the same course years before, and one of the students taking the course and becoming a novice in XP. We found that the best experience happened when the coach had both authority over the students and knowledge of XP. The conjunction of these two factors happened only when the professor was the coach. Nevertheless, we do believe that the other cases are also viable and, with proper care, can lead to good results; one must make sure that the two requirements are met (authority and knowledge of the method).

Customer. In the two initial years, the role of the customer was also played by CS professors. They were available during the two mandatory sessions and would be real users of the system to be built. In 2003, we developed a library management system so we invited a professor from another area and some staff members of our university library to act as customers. The experience was effective and very enlightening since the students realized that they had to use a completely different language to communicate with people that were not educated in CS.

Choosing the system. The choice of which system to build is very important: it must motivate the students, it must be interesting from a technological point of view, and it must be used often in order for us to find real future users that can act as customers. To meet all these requirements, we chose systems that the university needed to manage its resources and people.

We started with a Web-based system for managing course selection; the students could use it to express which elective courses they would like to take and the professors could express which courses they were able to teach and which ones they preferred to teach. The system then collected the results and used optimization techniques to create a course schedule for the following year. The system is now online at `mico.arca.ime.usp.br` and is used every year.

Tracking. In XP, the tracker is one of the developers who is responsible for collecting statistics about the performance of the team and who acts as its conscience, evaluating the progress of the project. Asking one student from each team to volunteer to be the tracker showed to be a very successful approach. By taking this approach we found highly motivated people for the task.

After the trackers are selected they are asked to read a few articles and book chapters on tracking, and encouraged to come up with creative ways of capturing team progress. Most of the time, the tracker chose to keep a copy of the story cards on a Web site⁶ so that team members could easily access them from any location at any time. The trackers are also responsible for maintaining the information radiators posted to the walls, as mentioned previously.

Technologies. For developing the projects, students utilize the latest real-world technology, which is very valuable for their future professional life. Most students consider this challenge motivating and work hard to learn the new tools.

The professors teaching the course do not specify any tool, language, or environment in which to develop the system. All the decisions are made by the team itself during the initial exploratory phase.

The systems were developed using modern, free software tools such as Java, Eclipse, CVS, ant, Apache, Tomcat, JSP/Struts, PostgreSQL, and Checkstyle. For unit testing, JUnit and CppUnit have been used in all our projects. User acceptance tests verifying the correct behavior of Web interfaces were carried out using HHTPUnit. Server-side Java code, such as Servlets, EJBs, and Tag Libs, has been tested using Cactus.

Student grading. Grades in Brazil are a numeric value between 0 and 10. We chose not to administer any exam during the XP course. So, the grades are calculated at the end of the semester based on four weighted criteria: attendance (30%), commitment to the XP method (35%), quality of the software produced (25%), and self-evaluation (10%).

The percentage distribution shows what is most important to us: that the students attend all the programming sessions and that the XP method be applied. Simply developing a good software system without using XP is not the objective of this course and this was made clear since the beginning.

5.2 Virtual Safe - When the Customer Wants XP

In this section we describe the mentoring performed by Alfredo Goldman and Fabio Kon at the Laboratory of Computer Networks Architecture at the University of São Paulo. A team from the lab was hired

⁶See, for example, <http://www.ime.usp.br/~xp/2003/xops/storycards>.

by a large Brazilian hardware and software company (Scopus Tecnologia S.A.) to develop a prototype of a Web-based virtual safe to store various kinds of computer files. In this case, the customer wanted to experience agile software development and chose XP as the method.

The team was composed of four programmers and it was supported by a manager who was responsible for helping the team, providing books, software, and good working conditions. Each of the two mentors visited the team twice a week to chat with the programmers, manager, and customer, to look at the source code, and to help with planning and assessment of progress. The mentors also had remote access to the source code (via CVS) and exchanged a few email messages every week. The 8 project participants subscribed to an email discussion list that was followed by all participants, although most of the posts came from the programmers.

On-site customer. The customer was an experienced systems analyst from Scopus. He agreed with the XP practices and committed to be on-site twice a week, which he did. In addition, the customer was also available by email and phone. As a former programmer, he felt comfortable participating in the project discussion list and acted as a link between the project programmers, working in the lab, and the Scopus developers and marketing department.

User Stories. After two introductory lectures on XP given by the mentors, the team started the planning phase. The customer wrote, with the help of the team, all the stories. At the end of the first planning meeting we had about 50 story cards which were divided in five releases. To provide time estimates for each release, we computed the average guessed time for each card (each of the 8 project participants made his guess and we used the arithmetic mean as the initial estimate of the group).

In addition to the story cards describing system functionalities, the customer also wrote two cards describing non-functional requirements: scalability and availability. These cards were not assigned to any particular developer, but they were hang on the wall so everyone would see them daily.

Metaphor. The system metaphor – a safe where one can store valuables – was proposed by the customer and was used during the entire development process, from the initial planning until the final marketing of the product.

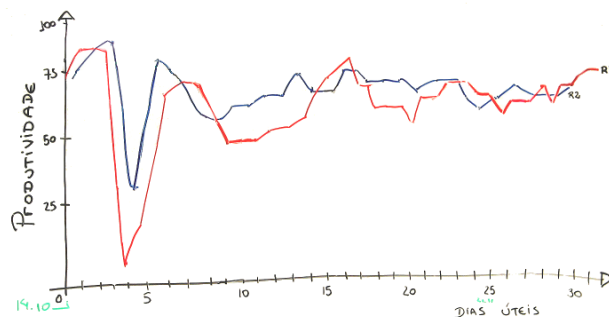


Figure 4: Subjective measure of team productivity during two releases

Releases. In the beginning of the first release, during the planning game, the programmers verified that the best technology would be J2EE. However, they had very little experience with it; they were Java programmers but did not know, for example, EJB which was considered to be a valuable technology to be used in the project. As the customer was present in this discussion, he agreed with the decision allowing an extension in the time estimate for the initial release; three additional weeks were added so that the programmers could learn and get used to EJB. The delay for the other three releases was about one week on average.

During the project, some user stories were added and others modified; some stories were also postponed to a later release. At the time of the last release (which included the least important user stories) the team agreed to cancel the release. Therefore, the 6-month project was delivered with a delay of a few days and a working version of the software. Part of its non-essential functionality was not implemented but this was not a surprise to the customer since he was following the entire process very closely.

Tracking. During the initial exploratory phase, one of the programmers volunteered to be the tracker. He collected detailed statistics about the project progress (e.g., number of tests, classes, and story cards written, number of story cards implemented, delays, etc.). In addition, one of the two available white boards was used to display tracking graphs depicting the evolution of the project. The group chose to plot, not only objective values (e.g. number of written tests), but also more subjective measures such as “team productivity” defined at the end of each day as a number between 0 and 100%. Figure 4 shows the evolution of team productivity for releases 2 and 3.

Reflection Workshop. We also used the idea of reflection workshops (also known as *retrospectives*) proposed by Alistair Cockburn [10]. After some time, all the people gathered in a short meeting to discuss how effective our method was and how we could improve it. In this meeting each programmer answered three questions:

1. What is working in the project?
2. What is not working in the project?
3. What can you suggest to improve our development methods and communication?

After one of these meetings we realized that pair programming and user stories were working fine, from the beginning. We also noticed that the programmers should pay more attention to unit tests. A suggestion was to replace the test-first idea; the team preferred to write tests after writing the functional code for a task.

Other XP practices. In this project, pair programming was used all of the time, and *move pairs around* was also applied once in a while. Tests were written with JUnit and even though the test-first practice was rarely applied, almost all the code was covered by unit tests. At the end, the programmers had written 160 unit tests. All the source code was shared in a CVS environment and, to check compliance to the group coding standards, the developers used the Eclipse *Checkstyle* plug-in. The implemented tasks were integrated to the repository continuously. At first, the team had some difficulties in refactoring, as they had only an introductory seminar on the topic. However as Eclipse provided some refactoring facilities, they got used to it, and they started to refactor the code more frequently.

5.3 Illinois Department of Public Health

In 1998 and 1999, The Refactory, Inc. provided a team to the Illinois Department of Public Health (IDPH) in order to assist with the development of medical software. Many applications at the Illinois Department of Public Health manage information about patients and people close to the patient, such as parents, children, and doctors. The programs vary in the kind of information (and the representation) they manage. However, there are core pieces of information that are common among the applications and can be shared among applications.

IDPH recruited The Refactory to assist with the development of an Enterprise framework for creating

these medical applications with the primary goals of 1) achieving reuse, 2) creating easier and quicker ways to deploy applications, and 3) to share common data across applications.

The primary development environment was Smalltalk, which was used for creating Windows-based client-server applications that interacted with a relational database running on a UNIX box. Joseph Yoder was the main software architect and led a 10-person team using XP practices (though not pure XP). This section will describe the experiences at attempting to integrate XP into IDPH and will point out some success and problems associated with incorporating XP into industry.

Open Space. IDPH used cubicles for each developer. The first thing we did was to remove the cubicles and create a shared common space. This common space was for us to pair program and to communicate more openly. We setup our workspaces in the open area to allow two people to share a single computer, primarily to facilitate pair programming. We put tables in the middle of our open area where we could gather around and share ideas. We also added a couple of whiteboards in which we could openly get together for brainstorm or shared design.

By setting up a shared area, we created an environment for good communication among developers. When new ideas were presented or new code released, everyone in our area immediately knew about it. This helped on the integration of new code. Refactoring was also easier since everyone had immediate access to all of the developers.

However, in general, most of the IDPH staff was uncomfortable with the open space idea. Our open space was almost too open. We were so open that all people within the IDPH staff could see what we were doing and hear us. Thus, we were too visible. We had some advocates to support what we were doing, but staff members not directly involved with our project were not as open to what we were doing. This sometimes stirred some inner controversy. For example, people would hear us talking amongst each other and at times our conversations would be misinterpreted.

For example, comments would be made about how we might be wasting time talking about items that were not directly related to our project. It is a common social phenomenon for people to discuss many items while dialoguing and quite often, people not involved with our project would complain to management that we were wasting time. They did not see the additional benefit that was created from the open

social environment.

Because of this, the mentor quite often had to protect or defend what the group was doing. In a sense, the mentor had to isolate or protect the team from the rest of the IDPH staff. It might have been better if we had isolated our shared space from the rest of IDPH. For example, we could have been in a separate room rather than in the middle of a large space. This would provide for an open environment that was still private from the rest of the organization. We could then prove our concept by our deeds rather than someone judging our process while in action.

However, everyone that was part of our shared space really liked what we were doing. We all felt that we were more productive in this environment based upon our previous experiences. We felt that we achieved more and had fun while we were doing it. We do not have empirical data backing up our feeling but each of us have had enough development experience to believe strongly that we produced more higher quality production code in this environment than the old “cubicle” style of developing software.

Pair Programming. Pair programming was used to develop most of the production code. Pair programming worked extremely well for us as knowledge about our frameworks was shared. Also, we all had an understanding of all of the code and we were never dependent on a single individual. We had people both leave the project and join our group. Because of pair programming we had a group understanding of the code and we were able to adapt to changes in our development team. Pair programming also provided new developers with good support for learning how to use our frameworks, thus becoming more productive in a less amount of time.

We would let individuals develop some spike solutions and what they developed would often be good enough to be incorporated into the code-base without pair programming. However, this code was released only with test cases and once the code was released into the shared repository, anyone could change the code. Therefore, there was no explicit code-ownership. Everyone “owned” the code and we worked together to make sure that we never left anything broken.

We did not rotate the pairs. Certain people tended to gravitate together and worked better together. We also had certain individuals that were very good at working with the spike solutions and integrating them into the environment. They worked well with the team but did not want to work in pairs.

Testing. Test cases and suites was an area where we went very “extreme”. We generated many test cases and suites. Our test cases were not always created first but we were very diligent about creating tests to validate our code and also to show how to use our frameworks. We always made sure to run the test cases at the end of each day.

These tests were invaluable during refactoring and integration. We all became strong believers when the test cases pointed out problems while we were integrating new functionality. Problems that would not have normally been found until late in the game were immediately found and fixed. It also made us comfortable about refactoring the code. We could apply a design pattern such as the Template Method and know immediately if we broke someone’s code.

One of the problems with building an application with our reusable frameworks was that our frameworks could be hard to understand and use. The tests provided a way to document how to use them, thus making it easier for developers to see how to use and build applications correctly with them.

Also, since we were using Smalltalk, we were able to evolve SUnit easily to make it so that we could create GUI tests. This allowed us to extend our test cases and suites to provide more extensive functional tests. We could then create complete user acceptance tests, thus ensuring the application worked according to the prescribed requirements.

Releases. We had regular internal releases and did what it took to keep a working version. This allowed us to demo the application often and get immediate feedback on what worked and what did not work. However, IDPH’s process for releasing applications did not let us release our applications to state employees on a regular basis. We could use our working version to meet with users and show them the application working. However, we did not receive the additional benefits that arise from regular feedback provided by a real customer using a current released version of the application.

Since our releases were never released to the customer until we were near completion, we did not receive the benefits of the regular feedback that XP promises. This is one of the XP principles that can be difficult in industry. Many users may feel that it is a waste of their time dealing with applications that are not completely functional.

User Stories. We did not create formal user stories. This was due to the unfortunate fact that we did

not have direct access to the users. Instead Joseph Yoder worked with the State Analysts to get the requirements and helped coordinate the team in an XP fashion. This is one of the biggest problems we had with our process.

For an application called The Refugee System, we had a customer that worked very closely with us. This helped to ensure that the system we developed was very close to what they needed. Therefore, when we were ready to go into production, the application pretty much met the needs of the end users.

However, we worked on another application called Newborn Screening (NBS), which had lots of problems. We did not have a relationship with the end user until the end of the development process. Because of the lack of a close relationship to a real customer, the system we developed was quite disparate from what the users needed. This led to many problems and complaints by the customers and management.

Upon reflection, we can see that a closer relationship with the customer was vital for success. Only relying on an analyst for the requirements was not good enough and by creating user stories for all of our applications, we might have been able to minimize some of the problems associated with NBS. Of course this is not unique to XP.

Assessment. In summary, we always kept things working, we were strongly test-driven, we did benefit from pair programming and the like. Our open space was invaluable to us though we would probably have benefited more by creating our open space in a semi-private area.

We wish we had pushed XP even further. However, it was hard to even push the principles as far as we did, given the political structure of a state organization. What we did worked well for us but our experience tells us that it could have worked even better.

For example, generating user stories and having regular releases could have helped ensure that our applications stayed on target. Rotating pairs would have helped sharing knowledge more.

There are a couple variations on XP that might have helped more such as possibly creating a proxy customer and using this proxy to generate the user stories. For example, we know of one organization that has successfully used analysts that worked very closely with the customer as a user proxy. Then user stories were created as part of the XP process [29].

5.4 Recife Short Course

In August 2003, a company called Qualiti located in Recife, Brazil presented a short XP course for industry professionals taught by Joseph Yoder. This section will outline how the course was presented along with some learning experiences.

Course Description The course was taught on site at Qualiti in Recife. We had twelve attendees which were from various areas of industry. The course duration was three, four-hour days and its description was as follows:

Evolving and adapting to changing requirements has become a crucial part of the design and programming process. Agile methods such as eXtreme Programming (XP) empowers all those that have an investment in the software being created. This ranges from the manager to the developer and end-user.

This short course will teach attendees the basic premise of Agile methods and will explore the details of the XP process. The course will consist of a mixture of lectures, reading groups, dialogs, and labs. The attendees will read some online materials, discuss the details of the techniques, and apply them in a group setting.

The three-day course was broken down by presenting, on the first day, an overview of the XP process followed by two days of hands-on experience actually working with the XP process.

Overview of XP. The first day really focused on ensuring that the students understood the main principles of XP and how the process worked. The first day overview presented: What is XP; Why XP; Principles of XP; The XP Process. This four-hour session emphasized issues such as the Customer Bill of Rights, the Programmer Bill of Rights, Rules and Practices of XP, and the overall process which included a detailed description of the iteration cycles and releases. We concluded this section with an introduction to the hands-on example that was worked on for the rest of the course.

Hands-on Example. The only way to really learn the principles of XP is by actually working with them. This is why any short course should have at least part of the course encourage students to actually try and

work through some of the principles. The students were broken down into two six-person teams working through the XP process. We would meet at regular intervals to compare notes and to learn from each other.

The primary goal of this task was to put into practice some of the principles of XP. Some of the main principles of XP that were described on the first day included items such as:

- Getting user stories from the customer
- Creating acceptance tests
- Creating spike solutions to understand the problem
- Creating a system metaphor
- Working with the customer to create a release plan
- Doing small iterations
- Doing an iteration plan
- Breaking the stories up into 1-3 day tasks
- Doing informal design such as CRC cards
- Doing test-driven development.

The example problem for practicing with XP dealt with the early design of a Conference Paper Submission System. The instructor knew this problem well and could thus work as the coach and customer; he ultimately really wanted to build such a system to use for the Pattern Languages of Programs (PLoP) conferences.

The task included creating user stories, generating acceptance tests, outlining an architectural spike to get a system metaphor, creating a release plan, and working through the start of an iteration where they broke the story up into small one to three day tasks. Then they did some initial design and outlined the unit tests for validating that the system would work properly. Rough requirements for the system were presented. This should be no surprise as this quite often happens in the real world. So, part of the task was to get better user stories from the user to make better estimates.

Analysis. Forcing the students to work through the process really emphasized how XP worked. The instructor could easily present a detailed overview of XP but many items were not understood until the students worked through the process. There was also a

huge benefit from the students interacting with one another, specifically when we came together and compared the results of the two groups.

One thing that was noted from the students was that using CRC cards for the design was difficult. Many of the students already knew UML well and they could draw class diagrams more easily than trying to learn a new way to describe their objects. The course did not dictate CRC but most students wanted to try it so that they could understand it and compare it to methods they were familiar with. XP does not dictate CRC and encourage developers to use whatever works well for them as long as they do not over design.

The main problem that the students had was trying to limit their designs. The students that attended the course were all very sharp developers from industry that had quite a lot of experience developing production systems. Thus, when they would work on an iteration, the temptation would be to go ahead and add some extra complexity or over design knowing what some of the next iterations would need. This was when refactoring and keeping it simple was emphasized. It is a difficult point to make to experienced developers. It goes against what they have learned in the past and they will probably not be convinced until they see the results by working many months on a successful XP project. From this, we can conclude that short courses are useful for introducing the concepts in industrial settings. However, this should be followed up by a long-term mentoring process where a coach works a few days a month on-site with the XP team.

6 Open Problems

Metaphor is one of the most difficult XP practice to teach. A good metaphor can be very important to improve the communication, although it usually does not receive the deserved attention. For example, in the University of São Paulo courses, the teams have not yet been able to use the Metaphor practice consistently. A possible way of introducing the use of metaphors in an organization would be to give a talk presenting examples of good and bad metaphors and emphasizing their benefits. The keynote speech given by Kent Beck at OOPSLA 2002 [5] could be used as a starting point.

In an academic environment, another problem exists related to student motivation. In the beginning, we may have an enthusiastic array of students, but this may change over the semester. In our experience, problems related to other courses (midterm and final exams, exercises, etc.) and personal problems caused

detrimental interference in the development of some XP projects. One solution would be to shorten the length of the courses. However, since similar problems also occur in real-life projects, it may be a good thing that they appear in the academic setting so that the participants learn to deal with them.

Finally, a problem often mentioned pertains to the difficulty of performing unit tests in stand-alone applications based on GUIs. In some environments it might be difficult to emulate user activities such as mouse clicks, drag and dropping, etc. However, this problem will probably not last too long since the tools for testing graphical interfaces (e.g., based on capture and replay of user movements) have been improving significantly in the last few years and we can expect the development of testing frameworks for these kinds of graphical interfaces [22, 23].

7 Conclusions

Agile software development methods, such as XP, are gradually being adopted by hundreds of organizations in the five continents. Nevertheless, the spirit of agile development is still not present in most of the organizations developing software. A new culture of agility and adaptation to change must be developed and educators must have a major role to play in this regard.

In this article, we have described our experiences in teaching XP in both academic and industrial environments and have discussed how one can be effective in teaching and implementing XP in an organization.

We have observed that, although there may be *a priori* fears of the consequences and effectiveness of XP, once developers and managers have real contact with a well-run XP project, the fears quickly dissipate. XP has proved to be a very attractive method both in academic and corporate environments due to the lack of surprises for customers and developers (thanks to the on-site customer practice) and to the high-quality of the software produced. Additionally, the environment created is optimal for developers who feel free to put all their energy in producing high-quality working code without the distractions required by bureaucratic processes that focus on tools and documents.

Within the next few years we expect that XP and agile methods will become part of the curriculum in many more universities around the world (at least as elective courses) and that industrial training and mentoring in XP will become more frequent. It is the role of educators and researchers to promote and enable this leap forward.

Acknowledgments The authors gratefully acknowledge the energy, excitement and great ideas provided by all the students and professionals that participated in our courses and mentoring projects. We thank Carlos Eduardo Ferreira, Ricardo Komatsu de Almeida and Eduardo Colli for being good customers, Klaus Wuestefeld for all the ideas shared with us and Eduardo Seiti Teruiya, César Rossi, Daniel Francis Soriano, Dairton L. Bassi Filho, and Fabio Taroda, Alexandre Freire, and Christian Asmussen for driving out fear and embracing change. We also thank Paulo Borba, Ralph Johnson, Steven Wingo, Paula Braga, Jason Frye. Finally, we acknowledge the wonderful feedback provided by the anonymous reviewers which helped us to improve the article significantly.

Alfredo Goldman, Fabio Kon and Paulo J. S. Silva were supported by CNPq under the grants 303246/2002-2, 302455/2003-5, and 304691/2002-0.

References

- [1] Johan Andersson, Geoff Bache, and Peter Sutton. XP with acceptance-test driven development: A rewrite project for a resource optimization system. In *Proceedings of the 4th International Conference on eXtreme Programming and Agile Processes in Software Engineering (XP 2003)*, volume 2675 of *Lecture Notes on Computer Science*, pages 180–188, 2003.
- [2] Gary H. Anthes. Sabre takes extreme measures. *Computer World*, see www.computerworld.com/softwaretopics/software/story/0,10801,91646,00.html, March 2004.
- [3] O. Astrachan, R. Duvall, and E. Wallingford. Bringing extreme programming to the classroom. In *Proceedings of XP Universe 2001*, Raleigh, NC, USA, 2001.
- [4] Kent Beck. *Extreme Programming Explained: Embrace Change*. Addison-Wesley, 1999.
- [5] Kent Beck. The metaphor metaphor. Keynote speech - ACM OOPSLA'02, November 2002.
- [6] Kent Beck. *Test-Driven Development: By Example*. Addison-Wesley, 2002.
- [7] Kent Beck et al. Manifesto for Agile Software Development. Home page: <http://agilemanifesto.org>, 2001.

- [8] Piergiuliano Bossi. eXtreme Programming applied: a case in the private banking domain. In *Proceedings of OOP*, Munich, 2003. Available at www.quinary.com/pagine/downloads/files/Resources/OOP2003Paper.pdf.
- [9] A. Cockburn and L. Williams. The costs and benefits of pair programming. In *Proceedings of the First International Conference on Extreme Programming and Flexible Processes in Software Engineering (XP2000)*, Cagliari, Sardinia, Italy, June 2000.
- [10] Alistair Cockburn. *Agile Software Development*. Addison-Wesley Longman, 2002.
- [11] Ron Crocker. *Large-Scale Agile Software Development*. Addison-Wesley, 2004.
- [12] Martin Fowler. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley, 1999.
- [13] Andrew M. Fuqua and John M. Hammer. Embracing change: An XP experience report. In *Proceedings of the 4th International Conference on eXtreme Programming and Agile Processes in Software Engineering (XP 2003)*, volume 2675 of *Lecture Notes on Computer Science*, pages 298–306, 2003.
- [14] Orit Hazzan, Joe Bergin, James Caristi, Yael Dubinsky, and Laurie Williams. Teaching software development methods: The case of Extreme Programming. In *Panel at the 35th Technical Symposium on Computer Science Education (SIGCSE'2004)*, Norfolk, Virginia, USA, 2004. ACM. db.grinnell.edu/sigcse/sigcse2004/viewAcceptedSession.asp?sessionType=SpecialSession&sessionNumber=28.
- [15] Donald Howard. Swimming around the waterfall: Introducing and using agile devel. In *Proceedings of the 4th International Conference on eXtreme Programming and Agile Processes in Software Engineering (XP 2003)*, volume 2675 of *Lecture Note in Computer Science*, pages 138–145, 2003.
- [16] P. Lappo. No pain, no XP: observations on teaching and mentoring extreme programming to university students. In *Proceedings of the 3rd International Conference on eXtreme Programming and Agile Processes in Software Engineering (XP2002)*, Cagliari, Italy, 2002.
- [17] Jim Little. Ats diary. C2 wiki site, see c2.com/cgi/wiki?AtsDiary, 2000.
- [18] Kim Man Lui and Keith C.C. Chan. Test driven development and software process improvement in china. In *Proceedings of the 5th International Conference on eXtreme Programming and Agile Processes in Software Engineering (XP 2004)*, volume 3092 of *Lecture Notes on Computer Science*, pages 219–222, 2004.
- [19] Katuscia Mannaro, Marco Melis, and Michele Marchesi. Empirical analysis on the satisfaction of it employees comparing xp practices with other software development methodologies. In *Proceedings of the 4th International Conference on eXtreme Programming and Agile Processes in Software Engineering (XP 2003)*, volume 2675 of *Lecture Notes on Computer Science*, pages 166–174, 2003.
- [20] K. Maruyama and K. Shima. Automatic method refactoring using weighted dependence graphs. In *Proceedings of the 21st international conference on Software engineering*, pages 236–245. IEEE Computer Society Press, 1999.
- [21] P. McBreen. *Questioning Extreme Programming*. Addison Wesley, 2003.
- [22] Atif M. Memon. *A Comprehensive Framework for Testing Graphical User Interfaces*. PhD thesis, University of Pittsburgh, July 2001.
- [23] Atif M. Memon. GUI testing: Pitfalls and process. *IEEE Computer*, 35(8):90–91, August 2002.
- [24] Roger A. Mller. Extreme Programming in a university project. In *Proceedings of the 5th International Conference on eXtreme Programming and Agile Processes in Software Engineering (XP 2004)*, volume 3092 of *Lecture Notes on Computer Science*, pages 312–315, 2004.
- [25] N. Nagappan, L. Williams, E. Wiebe, C. Miller, S. Balik, M. Ferzli, and M. Petlick. Pair learning: With an eye toward future success. In *Extreme Programming and Agile Methods - XP/Agile Universe 2003*, volume 2753 / 2003 of *Lecture Notes in Computer Science*, pages 185 – 198. Springer-Verlag Heidelberg, September 2003.
- [26] Joseph Pelrine. Modelling infection scenarios - a fixed-price eXtreme programming success story. In *ACM OOPSLA Companion Proceedings*, pages 23–24. ACM Press, 2000.
- [27] Ken Schwaber and Mike Beedle. *Agile Software Development with SCRUM*. Prentice Hall, 2001.

- [28] M. Stephens and D. Rosenberg. *Extreme Programming Refactored: The Case Against XP*. Apress, 2003.
- [29] C. Taber and M. Fowler. An iteration in the life on an XP project. *Cutter IT journal*, 13(11), November 2000. Updated electronic version: <http://www.martinfowler.com/articles/planningXpIteration.html>.
- [30] James E. Tomayko. A comparison of pair programming to inspections for software defect reduction. *Computer Science Education*, 12(3):213–222, 2002.
- [31] I. Tomek. What i learned teaching XP. In *Proceedings of the ACM OOPSLA Educators Symposium*, pages 39–46, Seattle, Washington, USA, November 2002.
- [32] L. Williams and R. Kessler. *Pair Programming Illuminated*. Addison-Wesley, 2002.
- [33] L. A. Williams and R. R. Kessler. All I really need to know about pair programming I learned in kindergarten. *Communications of the ACM*, 43(5):108–114, May 2000.
- [34] D. Wilson. Teaching XP: a case study. In *Proceedings of XP Universe 2001*, Raleigh, NC, USA, 2001.
- [35] Klaus Wuestefeld. Customer Proxy and Libero. E-mail personal communication, September 2002.