



# Combining STPA and BDD for Safety Analysis and Verification in Agile Development: A Controlled Experiment

Yang Wang<sup>(✉)</sup> and Stefan Wagner

University of Stuttgart, Stuttgart, Germany  
{yang.wang, stefan.wagner}@informatik.uni-stuttgart.de

**Abstract.** *Context:* Agile development is in widespread use, even in safety-critical domains. *Motivation:* However, there is a lack of an appropriate safety analysis and verification method in agile development. *Objective:* In this paper, we investigate the use of Behavior Driven Development (BDD) instead of standard User Acceptance Testing (UAT) for safety verification with System-Theoretic Process Analysis (STPA) for safety analysis in agile development. *Method:* We evaluate the effect of this combination in a controlled experiment with 44 students in terms of productivity, test thoroughness, fault detection effectiveness and communication effectiveness. *Results:* The results show that BDD is more effective for safety verification regarding the impact on communication effectiveness than standard UAT, whereas productivity, test thoroughness and fault detection effectiveness show no statistically significant difference in our controlled experiment. *Conclusion:* The combination of BDD and STPA seems promising with an enhancement on communication, but its impact needs more research.

## 1 Introduction

Agile practices have been widely used in software industries to develop systems on time and within budget with improved software quality and customer satisfaction [1]. The success of agile development has led to a proposed expansion to include safety-critical systems (SCS) [2]. However, to develop SCS in an agile way, a significant challenge exists in the execution of safety analysis and verification [3]. The traditional safety analysis and verification techniques, such as failure mode effect analysis (FMEA) and fault tree analysis (FTA) are difficult to apply within agile development. They need a detailed and stable architecture [4].

In 2016, we proposed to use System-Theoretic Process Analysis (STPA) [6] in agile development for SCS [5]. First, STPA can be started without a detailed and stable architecture. It can guide the design. In agile development, a safety analyst starts with performing STPA on a high-level architecture and derives the relevant safety requirements for further design. Second, Leveson developed STPA based on the systems theoretic accident modeling and processes (STAMP) causality model, which considers safety problems based on system theory rather

than reliability theory. In today’s complex cyber-physical systems, accidents are rarely caused by single component or function failures but rather by component interactions, cognitively complex human decision-making errors and social, organizational, and management factors [6]. System theory can address this.

The safety requirements derived from STPA need verification. However, there is no congruent safety verification in agile development. Most agile practitioners mix unit test, integration test, field test and user acceptance testing (UAT) to verify safety requirements [2]. In 2016, we proposed using model checking combined with STPA in a Scrum development process [7]. However, using model checking, a suitable model is necessary but usually not available in agile development. In addition, the formal specification increases the difficulties of communication, which should not be neglected when developing SCS [8]. BDD, as an agile technique, is an evolution of test driven development (TDD) and acceptance test driven development (ATDD). The developers repeat coding cycles interleaved with testing. TDD starts with writing a unit test, while ATDD focuses on capturing user stories by implementing automated tests. BDD relies on testing system behavior in scenarios by implementing a template: Given[Context], When[Event], Then[Outcome] [31]. The context describes pre-conditions or system states, the event describes a trigger event, and the outcome is an expected or unexpected system behavior. It could go further into low-level BDD<sup>1</sup>. Yet, it has not been used to verify safety requirements. Leveson said [6]: “*Accidents are the result of a complex process that results in system behavior violating the safety constraints.*” Hence, in agile development, we need safety verification to: (1) be able to guide design at an early stage, (2) strengthen communication and (3) focus on verifying system behavior. Thus, we believe that BDD might be suitable for safety verification with STPA for safety analysis in agile development.

### Contributions

We propose a possible way to use BDD with STPA for safety verification in agile development. We investigate its effects regarding productivity, test thoroughness, fault detection effectiveness and communication effectiveness by conducting a controlled experiment with the limitation that we execute BDD only in a test-last way. The results show that BDD is able to verify safety requirements based on system theory, and is more effective than UAT regarding communication for safety verification.

## 2 Related Work

Modern agile development processes for developing safety-critical systems (SCS) advocate a hybrid mode through alignment with standards like IEC 61508, ISO 26262 and DO-178. There have been many considerable successes [9–11]. However, a lack of integrated safety analysis and verification to face the changing architecture through each short iteration is a challenge for using such standards.

---

<sup>1</sup> Low-level BDD is possible to define low-level specifications and intertwined with TDD [16].

In 2016, we proposed to use STPA in a Scrum development process [5]. It showed a good capability to ensure agility and safety in a student project [12]. However, we verified the safety requirements only at the end of each sprint by executing UAT together with TDD in development. A lack of integrated safety verification causes some challenges, such as poor verification and communication. The previous research regarding safety verification in agile development suggested using formal methods [13, 14]. However, they need models and make intuitive communication harder [7]. In addition, they have not considered specific safety analysis techniques.

Hence, we propose using BDD to verify safety requirements. BDD is specifically for concentrating on behavior testing [15]. It allows automated testing against multiple artifacts throughout the iterative development process [17]. Moreover, it bridges the gap between natural language-based business rules and code language [18]. Okubo et al. [19] mentioned the possibilities of using BDD for security and privacy acceptance criteria. They define the acceptance criteria by creating a threat and countermeasure graph to write attack scenarios. They verify the satisfaction of security requirements by testing the countermeasures, to see whether they can make the attack scenarios or unsecure scenarios fail. Lai et al. [20] combined BDD with iterative and incremental development specifically for security requirements evaluation. They defined the behavioral scenarios by using use case diagram and misuse case diagram. STPA encompasses determining safe or unsafe scenarios. We aim to use BDD verifying these scenarios.

To investigate the effect of using BDD for safety verification, we design a controlled experiment referring to a set of TDD experiments. Erdogmus et al. [23] conducted an experiment with undergraduate students regarding programmer productivity and external quality in an incremental development process. For safety verification in agile development, a high productivity of safety test cases promotes high safety. Madeyski [26] conducted an experiment comparing “test-first” and “test-last” programming practices with regard to test thoroughness and fault detection effectiveness of unit tests. BDD for safety verification covers also low-level tests. Thus, we decided to investigate productivity, test thoroughness and fault detection capability in this experiment. [21, 22, 28–30] provided evidence of using these three measures. In addition, George and Williams [29] focused on the understandability of TDD from the developer’s viewpoint. Using BDD for safety verification, we notice the importance of communication between developers and business analysts. We investigate understandability in the measure of communication effectiveness.

### 3 STPA Integrated BDD for Safety Analysis and Verification (STPA-BDD)

In this article, we propose STPA-BDD. We mainly focus on safety verification. As we can see in Fig. 1, we have two main parts: STPA safety analysis and

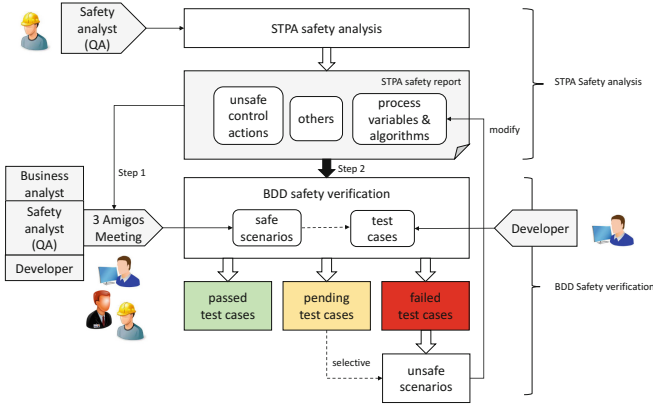


Fig. 1. STPA-BDD concept

BDD safety verification. A safety analyst<sup>2</sup> (QA) starts performing STPA safety analysis with a sufficient amount of code<sup>3</sup>. STPA is executed by firstly identifying potentially hazardous control actions, and secondly determining how unsafe control actions (UCAs) could occur. STPA derives the safety requirements, which constraint the UCAs, as well as system behaviors. Additionally, it explores the causal factors in scenarios for each UCA. The output from the safety analyst (QA) is an STPA safety report with system description, control structure, accidents, hazards, UCAs, corresponding safety requirements, process variables and algorithms.

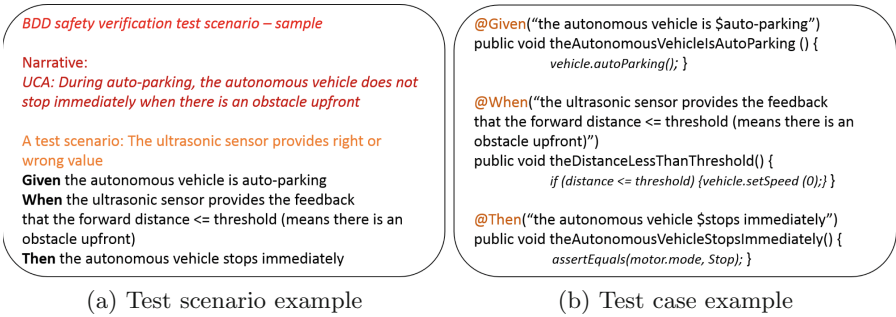
In BDD safety verification, to generate and test scenarios, the UCAs (in STPA step 1), process variables and algorithms (in STPA step 2) from the STPA safety report are needed. We write other data into “others”. BDD safety verification has two steps: In step 1, the business analyst, the safety analyst (QA) and the developer establish a “3 Amigos Meeting” to generate test scenarios. In a BDD test scenario<sup>4</sup>, we write the possible trigger event for the UCA in **When [Event]**. The other process variables and algorithms are arranged in **Given [Context]**. **Then [Outcome]** presents the expected behavior - a safe control action. In Fig. 2(a), we present an example. The safety analyst (QA) has provided a UCA as *During auto-parking, the autonomous vehicle does not stop immediately when there is an obstacle upfront*. One of the process variables with relevant algorithms detects the forward distance by using an ultrasonic sensor. The developer considers a possible trigger as the ultrasonic sensor provides the

<sup>2</sup> Since we focus on safety in our research, we assign a safety analyst as the QA role in our context.

<sup>3</sup> More descriptions of STPA for safety analysis are given in [7] concerning an example of using STPA in an airbag system and [12] concerning the use of STPA in a Scrum development process.

<sup>4</sup> We illustrate a BDD test scenario using only three basic steps “Given” “When” “Then”. More annotations, such as “And”, can also be added.

wrong feedback. Thus, a BDD test scenario should test if *the ultrasonic sensor provides the feedback that the forward distance  $\leq$  threshold (means there is an obstacle upfront)* and whether the vehicle stops. They write this after **When**. The context could be *the autonomous vehicle is auto-parking*. We write them after **Given**. **Then** constraints the safe control action as *the autonomous vehicle stops immediately*. More possible triggers are expected to be generated after **When** to test them. In step 2, after the three amigos discuss and determine the test scenarios, the developer starts generating them into test cases, as shown in Fig. 2(b). BDD test cases use annotations such as **@Given**, **@When**, and **@Then** to connect the aforementioned test scenarios with real code. The developer produces code to fulfill each annotation. We can identify unsafe scenarios when the test cases fail. We correct the trigger event to pass the test cases to satisfy the safety requirement.



**Fig. 2.** BDD safety verification example

## 4 Experiment Design (We follow the guideline by Wohlin et al. [32].)

### 4.1 Goal

Analyze BDD<sup>5</sup> and UAT<sup>6</sup> for safety verification.

**For the purpose of** comparing their effect.

**With respect to** *productivity* by measuring the number of implemented (tested) user stories per minute; *test thoroughness* by measuring line coverage; *fault detection effectiveness* by measuring a mutation score indicator; *communication effectiveness* by conducting a post-questionnaire.

**From the point of view** of the developers and business analysts.

**In the context of** B.Sc students majoring in software engineering or other related majors executing acceptance testing.

<sup>5</sup> We have a limitation in our experiment that we execute BDD only in a test-last way. More discussion of this issue can be found in Sect. 6.2.

<sup>6</sup> To execute a standard UAT, we mainly refer to [38] with fictional business analysts.

## 4.2 Context

*Participants:* The experiment ran off-line in a laboratory setting in an “Introduction to Software Engineering” course at the University of Stuttgart. Since the course includes teaching BDD and UAT technology, the students are suitable subjects for our experiment. We arrange them based on Java programming experiences (not randomly). According to a pre-questionnaire (see footnote 13), 88.6% of the students are majoring in software engineering. We conclude from Table 1 that they have attended relevant lectures and handled practical tasks relating to Java programming, acceptance testing, SCS (with a median value  $\geq 3$  on a scale from 1 to 5). The agile techniques show less competency (with a median value of 2 on a scale from 1 to 5). We provide a 1-to-1 training, which lasts 44 h overall, to reduce the weaknesses.

*Development environment:* We use a simplified Java code with mutants from a Lego Mindstorms based Autonomous Parking System (APS) and Crossroad Stop and Go System (CSGS). These two systems are comparable by lines of code and number of functional modules (see footnote 13). To ease writing test cases, we use a lejo TDD wrapper, Testable Lejos<sup>7</sup> to remove deep dependencies to the embedded environment. The BDD groups (Group A1 and Group A2) operate in an Eclipse IDE together with a JBehave plug-in (based on JUnit)<sup>8</sup>. We use Eclipse log files and JUnit test reports for calculating the number of implemented (tested) user stories. Finally, we use PIT Mutation Testing<sup>9</sup> to assess line coverage and a mutation score indicator. The UAT groups (Group B1 and Group B2) write the test cases in Microsoft Word.

**Table 1.** Medians of the student’s background

Area	Group A1	Group A2	Group B1	Group B2
Java programming	3	3	3	3
Acceptance testing	4	5	3	3
Safety-critical systems	3	4	4	4
Agile techniques	3	3	3	2

Note: The values range from “1” (little experience) to “5” (experienced). Group A1 and Group A2 use BDD, while Group B1 and Group B2 use UAT.

## 4.3 Hypotheses

We formulate the null hypotheses as:

$H_0_{PROD}$ : There is no difference in productivity between BDD and UAT.

$H_0_{THOR}$ : There is no difference in test thoroughness between BDD and UAT.

<sup>7</sup> <http://testablelejos.sourceforge.net/>.

<sup>8</sup> <http://jbehave.org/eclipse-integration.html>.

<sup>9</sup> <http://pitest.org/>.

$H_0_{FAUL}$ : There is no difference in fault detection effectiveness between BDD and UAT.

$H_0_{COME}$ : There is no difference in communication effectiveness between BDD and UAT.

The alternative hypotheses are:

$H_1_{PROD}$ : BDD is more productive than UAT when producing safety test cases.

$H_1_{THOR}$ : BDD yields better test thoroughness than UAT.

$H_1_{FAUL}$ : BDD is more effective regarding fault detection than UAT.

$H_1_{COME}$ : BDD is more effective regarding communication than UAT.

#### 4.4 Variables

The independent variables are the acceptance testing techniques. The dependent variables are: (1) productivity (PROD). It is defined as output per unit effort [23]. In our experiment, the participants test the user stories in the STPA safety report and produce safety test cases. We assess it via the number of implemented (tested) user stories<sup>10</sup> per minute (NIUS) [23]; (2) test thoroughness (THOR). Code coverage is an important measure for the thoroughness of test suites including safety test suites [27]. Considering a low complexity of our provided systems, line coverage (LC) [26] is more suitable than branch coverage (BC); (3) fault detection effectiveness (FAUL). Mutation testing [25] is powerful and effective to indicate the capability at finding faults [26]. In our experiment, we measure how well a safety test suite is able to find faults at the code level. We assess this via a Mutation Score Indicator (MSI) [26]; (4) communication effectiveness (COME). We assess this via a post-questionnaire with 11 questions for developers covering topics like understandability and 13 questions for business analysts covering topics like confidentiality according to Adzic [35]. The results are in a 5-point scale from  $-2$  (negative) to  $+2$  (positive).

#### 4.5 Pilot Study

Six master students majoring in software engineering took part in a pilot study. We arranged a four-hour training program. The first author observed the operation and concluded as follows: (1) The STPA safety report was too complicated to be used by inexperienced students. We used a comprehensive STPA report by using XSTAMPP<sup>11</sup> in the pilot study. However, a lot of unnecessary data, such as accidents, hazards and safety requirements at the system level, influenced the understanding. It costs too much time to capture the information. Thus, we simplified the STPA report with the process variables, algorithms, and UCAs. (2) We used the original Java code from a previous student project. The complex code affected the quick understanding. After the pilot study, we simplified it. (3) Training is extremely important. In the pilot study, one participant had not

<sup>10</sup> In this article, user stories are safety-related user stories.

<sup>11</sup> <http://www.xstampp.de/>.

taken part in the training program, which led to his experiment being unfinished. We provide a textual tutorial and system description for each participant as a backup. (4) We have only used an experiment report to record the measures. However, the pure numbers sometimes cannot show clear causalities. Thus, we use a screen video recording in parallel with the experiment report.

### 4.6 Experiment Operation

As we can see in Fig. 3, we divide the 44 participants into 4 groups. We provide 2 systems and evaluate 2 acceptance testing methods. Group A1 uses BDD for system 1. Group A2 uses BDD for system 2. Group B1 uses UAT for system 1. Group B2 uses UAT for system 2. We use two systems to evaluate the communication between developers and business analysts. The developers are the participants in each group, while the fictional business analysts are portrayed by the participants in the other group using various testing methods and systems.

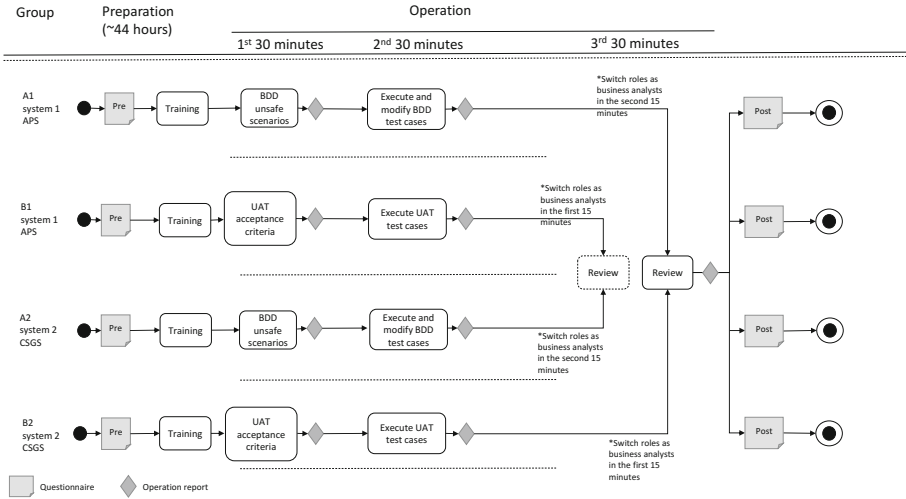


Fig. 3. Experiment operation

The experiment consists of 2 phases: *preparation* and *operation*. The *preparation* was run 2 weeks before the experiment to perform the pre-questionnaire and training. The *operation* consists of three sessions (30 min/session). In the 1<sup>st</sup> session, four groups write acceptance test cases. Group A1 (BDD) and Group A2 (BDD) write test scenarios in Eclipse with the Jbehave plug-in as a story file. Group B1 (UAT) and Group B2 (UAT) write acceptance criteria in plaintext. We provide 30 unsafe control actions (UCAs) in an STPA safety report. When the students finish all the 30 UCAs in 30 min, they record the time in minutes. After the 1<sup>st</sup> session, the participants record the NIUS and the time in the operation report. In the 2<sup>nd</sup> session, Group A1 (BDD) and Group A2 (BDD) write



each test scenario into a test case and run the test case. If it fails, they should modify the trigger (code) and pass the test case. Group B1 (UAT) and Group B2 (UAT) review Java code, execute the test cases manually and complete their acceptance test report. At the end of the 2<sup>nd</sup> session, they run PIT mutation testing. The LC and MSI are generated automatically in the PIT test report. They write down the results in the operation report. In the 3<sup>rd</sup> session, the participant portrays as a developer for 15 min and a business analyst for 15 min. The developer is expected to explain his/her testing strategy as clearly as possible, while the fictional business analyst should try to question the developer. To this end, they answer a post-questionnaire.

**Table 2.** Descriptive statistic

Measure	Treatment	Experiment	Mean	St.Dev	Min	Median	Max	95% CI lower	95% CI upper
NIUS	BDD	Group A1	0.52	0.24	0.26	0.45	1.20	0.37	0.66
		Group A2	0.69	0.19	0.42	0.65	1.00	0.58	0.80
	UAT	Group B1	0.58	0.22	0.33	0.57	1.00	0.45	0.71
		Group B2	0.67	0.29	0.27	0.60	1.20	0.50	0.84
LC	BDD	Group A1	0.02	0.01	0.01	0.02	0.05	0.02	0.03
		Group A2	0.02	0.01	0.01	0.02	0.04	0.02	0.03
	UAT	Group B1	0.02	0.01	0.01	0.01	0.03	0.01	0.02
		Group B2	0.02	0.01	0.01	0.01	0.03	0.01	0.02
MSI	BDD	Group A1	0.90	0.38	0.36	1.00	1.33	0.67	1.13
		Group A2	0.93	0.49	0.44	0.83	2.17	0.63	1.22
	UAT	Group B1	0.89	0.36	0.42	0.88	1.56	0.67	1.10
		Group B2	0.85	0.46	0.30	0.65	1.63	0.58	1.12
COME	BDD	Group A1	1.27	0.81	-2.00	<b>1.50</b>	2.00	0.79	1.75
		Group A2	1.18	0.70	-1.00	<b>1.00</b>	2.00	0.76	1.58
	UAT	Group B1	-0.05	1.20	-2.00	<b>0.00</b>	2.00	-0.75	0.66
		Group B2	0.01	1.13	-2.00	<b>0.50</b>	2.00	-0.67	0.67

Note: St. Dev means standard deviation; CI means confidence interval. NIUS means number of implemented (tested) user stories per minute. LC means line coverage. MSI means mutation score indicator. COME was assessed via questionnaire with the results in a 5-point scale from -2 (negative) to +2 (positive).

## 5 Analysis

### 5.1 Descriptive Analysis

In Table 2, we summarize the descriptive statistics of the gathered measures<sup>12</sup>. To sum up, the results from the two systems in one treatment are almost identical. BDD and UAT have only small differences regarding NIUS and MSI. However,

<sup>12</sup> Raw data is available online: <https://doi.org/10.5281/zenodo.1154350>.

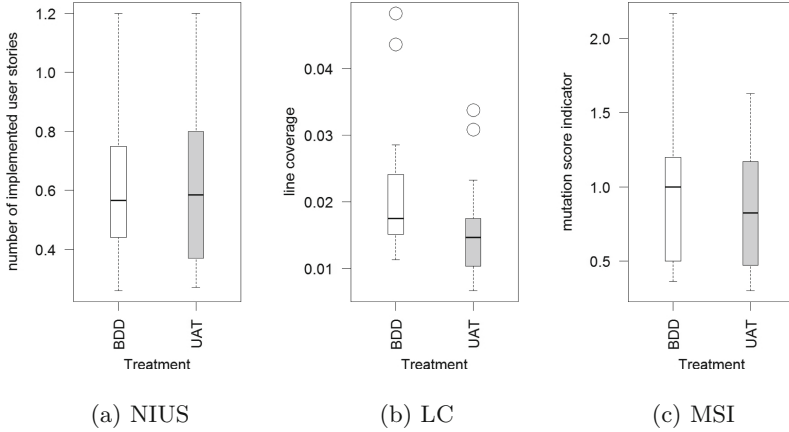


Fig. 4. Boxplot for PROD, THOR and FAUL

COME in BDD (Mean = 1.27, 1.18; Std.Dev = 0.81, 0.70) and UAT (Mean = -0.05, 0.01; Std.Dev = 1.20, 1.13) differ more strongly. LC has a small difference. In Fig. 4, we show a clear comparison and can see some outliers concerning LC. In Fig. 5, we use an alluvial diagram to show COME. We can conclude that BDD has a better communication effectiveness than UAT from the perspective of developers and business analysts respectively (depending on the length of black vertical bar on the right side of Fig. 5(a) and (b)). On the left side, we list 24 sub-aspects of assessing the communication effectiveness. The boldness of the colorful lines indicates the degree of impact. A thicker line has a larger impact on each aspect. We can see six noteworthy values from Fig. 5(a) that BDD is better than UAT: (4) Test cases have a clear documentation. (5) They could flush out the functional gaps before development. (6) They have a good understanding of business requirements.

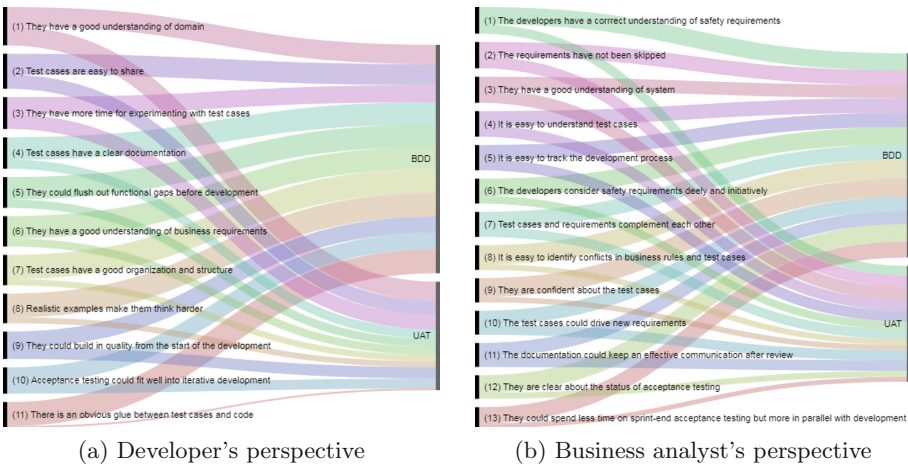


Fig. 5. Alluvial diagram for communication effectiveness

business requirements. (7) Test cases have a good organization and structure. (8) Realistic examples make them think harder. (11) There is an obvious glue between test cases and code. From Fig. 5(b), five noteworthy values show that BDD is better than UAT: (6) The developers consider safety requirements deeply and initially. (8) It is easy to identify conflicts in business rules and test cases. (9) They are confident about the test cases. (12) They are clear about the status of acceptance testing. (13) They could spend less time on sprint-end acceptance testing but more in parallel with development. In addition, the other aspects show also slightly better results when using BDD than UAT.

## 5.2 Hypothesis Testing

To start with, we evaluate the pre-questionnaire. No statistically significant differences between BDD and UAT groups are found concerning Java programming, acceptance testing, knowledge on SCS and agile techniques (t-test,  $\alpha = 0.05$ ,  $p > 0.05$  for all test parameters). Furthermore, we test the normality of the data distribution with Kolmogorov-Smirnov and Shapiro-Wilk tests at  $\alpha = 0.05$ . The results show that the data for NIUS in Group A1, for LC in Group A1, A2, B2 and for MSI in Group A1, A2 are not normally distributed. Thus, we use non-parametric tests in the analysis. In addition to the use of p-values for hypotheses testing ( $\alpha = 0.05$ , one-tailed) from the Mann-Whitney test, Wilcoxon test and ANOVA test, we include the effect size Cohen's d. Since we expect BDD to be better than UAT, we use one-tailed tests. NIUS is not significantly affected by using the BDD or the UAT approach (system 1:  $p=0.206$ ; system 2:  $p=0.359$ , non-significant). LC is not significantly affected by using BDD or UAT (system 1:  $p=0.057$ ; system 2:  $p=0.051$ , non-significant). MSI shows no statistically significant difference between using BDD or UAT (system 1:  $p=0.472$ ; system 2:  $p=0.359$ , non-significant). However, COME is significantly different (system 1:  $p < 0.00001$ ; system 2:  $p < 0.00001$ , significant). We accept the alternative hypothesis that BDD shows better communication effectiveness than UAT. Cohen's d shows the values around 0.2, which signifies small effects, around 0.5 stands for medium effects and around 0.8 for large effects. Thus, for COME, system 1 shows a large effect ( $d = 2.908$ ). For LC we have both medium effects (system 1:  $d = 0.684$ ; system 2:  $d = 0.662$ ). The rest of the effects are small.

## 6 Threats to Validity

### 6.1 Internal Validity

*First*, note that we have four groups in our experiment. To avoid a multiple group threat, we prepare a pre-questionnaire to investigate the students' background knowledge. The results of the t-tests show no statistically significant differences among the groups concerning each measure. *Second*, concerning the instrument, UAT is faster to learn than BDD regarding the use of tools. Even though we provide a training to narrow the gap, the *productivity* might have been influenced, since the students have to get familiar with the hierarchy of writing test

suites in a BDD tool. The artifacts, such as tutorials and operation report, are designed respectively with the same structure to avoid threats. In addition to the observation, we save the participants' workspaces after the experiment and video recordings for deep analysis. *Third*, the students majoring in software engineering might identify more with the developer role than the business analyst role. Thus, we design two comparable systems. The students in each pair use different systems and test approaches to reduce the influence of prior knowledge. Moreover, we provide a reference [36] on how to perform as a business analyst in an agile project. We also mention their responsibilities in the training.

## 6.2 Construct Validity

*First*, the execution of BDD is a variant. BDD should begin with writing tests before coding. However, in our experiment, we use BDD for test-last acceptance testing rather than test-driven design. Thus, we provide source code with mutants. The measures we used could be influenced. In BDD test-first, we write failing test cases first and work on passing all of them to drive coding. According to [39, 41], BDD test-first might be as effective as or even more effective than BDD test-last. *Second*, the evaluation concerning productivity, test thoroughness, fault detection effectiveness and communication effectiveness does not seem to be enough. As far as we know, our study is the first controlled experiment on BDD. We can base our measurement (PROD, THOR, FAUL) mainly on TDD controlled experiments and some limited experiments on safety verification. There might be better ways to capture how well safety is handled in testing.

## 6.3 Conclusion Validity

*First*, concerning violated assumptions of statistical tests, the Mann-Whitney U-test is robust when the sample size is approximately 20. For each treatment, we have 22 students. Moreover, we use Wilcoxon W test as well as Z to increase the robustness. Nevertheless, under certain conditions, non-parametric rank-based tests can themselves lack robustness [44]. *Second*, concerning random heterogeneity of subjects, we arranged them based on the Java programming experience. According to the pre-questionnaire, the students are from the same course and 88.6% of them are in the same major.

## 6.4 External Validity

*First*, the subjects are students. Although there are some skilled students who could perform as well as experts, most of them lack professional experience. This consideration may limit the generalization of the results. To consider this debatable issue in terms of using students as subjects, we refer to [33]. They said: conducting experiments with professionals as a first step should not be encouraged unless high sample sizes are guaranteed. In addition, a long learning

cycle and a new technology are two hesitations for using professionals. STPA was developed in 2012, so there is still a lack of experts on the industrial level. BDD has not been used for verifying safety requirements. Thus, we believe that using students as subjects is a suitable way to aggregate contributions in our research area. We also refer to a study by Cleland-Huang and Rahimi, which successfully ran an SCS project with graduate students [2]. *Second*, the simplicity of the tasks poses a threat. We expect to keep the difficulty of the tasks in accordance with the capability of students. Nevertheless, the settings are not fully representative of a real-world project.

## 7 Discussion and Conclusion

The main benefit of our research is that we propose a possible way to use BDD for safety verification with STPA for safety analysis in agile development. We validate the combination in a controlled experiment with the limitation that we used BDD only in a test-last way. The experiment shows some remarkable results. The *productivity* has no statistically significant difference between BDD and UAT. That contradicts our original expectation. We would expect BDD, as an automated testing method, to be more productive than manual UAT. Yet, as the students are not experts in our experiment, they need considerable time to get familiar with the BDD tool. The students use Jbehave to write BDD test cases in our experiment, which has strict constraints on hierarchy and naming conventions to connect test scenarios with test cases. UAT should be easier to learn. We therefore analyzed our video recordings and found that BDD developers use nearly 25% to 50% of their time to construct the hierarchy and naming. Scanniello et al. [37] also mentioned this difficulty when students apply TDD. In the future, we plan to use skilled professionals in test automation to replicate this study. This could lead to different results. The *test thoroughness* and *fault detection effectiveness* show a non-significant difference between BDD and UAT. We could imagine that our provided Java code is too simplified to show a significant difference. The mutants are easily found with a review. These aspects need further research.

The *communication effectiveness* shows better results by using BDD than UAT on 24 aspects. We highlight 11 significant aspects. The *developers* found that: **BDD has a clear documentation.** A clear documentation of acceptance test cases is important for communication [42]. The scenarios are written in plain English with no hidden test instrumentation. The given-when-then format is clear for describing test scenarios for safety verification based on system theory. **The developers using BDD could flush out functional gaps before development.** The communication concerning safety could happen at the beginning of the development. They discuss safety requirements with the business analysts and spot the detailed challenges or edge cases before functional development. UAT happens mostly at the end of the development. It makes the rework expensive and is easy to be cut in safety-critical systems. **The developers using BDD have a good understanding of the business**

**requirements.** A good understanding of safety requirements helps an effective communication. They could build a shared understanding in the “3 Amigos Meeting” to ensure that their ideas about the safety requirements are consistent with the business analysts. The developers using UAT might understand safety requirements with a possible bias. **BDD test cases have a good organization and structure.** This makes the test cases easy to understand, especially during maintenance. They include strict naming conventions and a clear hierarchy to manage test scenarios and test cases. **Realistic examples in BDD make the developers think harder.** The safety requirements are abstract with possibly cognitive diversity, which leave a lot of space for ambiguity and misunderstanding. That negatively influences effective communication. Realistic examples give us a much better way to explain how safe scenarios really work than pure safety requirements do. **There is an obvious glue between BDD test cases and code.** There is glue code in BDD safety verification, which allows an effective separation between safety requirements and implementation details. This glue code supports the understanding and even communication between business analysts and developers. In addition, it ensures the bidirectional traceability between safety requirements and test cases. The *business analysts* thought that: **The developers using BDD consider the safety requirements deeply and initiatively.** The collaboration promotes a sense of ownership of the deliverable products. That increases an initiative communication. Instead of passively reading the documents, the developers participate in the discussion about writing test scenarios and are more committed to them. **The business analysts are more confident about the BDD test cases.** Confidence promotes effective communication [43]. The business analysts could give a big picture with safety goals to the developers. Feedback from developers and their realistic unsafe scenarios give the business analysts confidence that the developers understand the safety goals correctly. **It is easy to identify conflicts in business rules and test cases when using BDD.** BDD has a set of readable test scenarios focusing on business rules (safety requirements). Each test scenario and test case are directly connected to the code. The business analysts can pull out test cases related to a particular business rule. This helps communication, especially when there is a changing request. **The business analysts are clear about the status of acceptance testing when using BDD.** It promotes a state-of-art communication. That can be attributed to the automated test suites, which might be connected with a continuous integration server and a project management tool to receive a verification report automatically. **The business analysts could spend less time on sprint-end acceptance tests but more in parallel with development.** They can verify the safety requirements periodically and therefore enhance communication throughout the project.

In conclusion, to some extent, BDD is an effective method for verifying safety requirements in agile development. As this is the first experiment investigating BDD for safety verification, further empirical research is needed to check our results. We invite replications of this experiment using our replication package<sup>13</sup>.

<sup>13</sup> <https://doi.org/10.5281/zenodo.846976>.

## References

1. Dybå, T., Dingsøy, T.: Empirical studies of agile software development: A systematic review. *Inf. Softw. Technol.* **50**(9–10), 833–859 (2008)
2. Cleland-Huang, J., Rahimi, M.: A case study: injecting safety-critical thinking into graduate software engineering projects. In: *Proceedings of the 39th International Conference on Software Engineering: Software Engineering and Education Track*. IEEE (2017)
3. Arthur, J.D., Dabney, J.B.: Applying standard independent verification and validation (IV&V) techniques within an Agile framework: is there a compatibility issue? In: *Proceedings of Systems Conference*. IEEE (2017)
4. Fleming, C.: Safety-driven early concept analysis and development. Dissertation. Massachusetts Institute of Technology (2015)
5. Wang, Y., Wagner, S.: Toward integrating a system theoretic safety analysis in an agile development process. In: *Proceedings of Software Engineering, Workshop on Continuous Software Engineering* (2016)
6. Leveson, N.: *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, Cambridge (2011)
7. Wang, Y., Wagner, S.: Towards applying a safety analysis and verification method based on STPA to agile software development. In: *IEEE/ACM International Workshop on Continuous Software Evolution and Delivery*. IEEE (2016)
8. Martins, L.E., Gorschek, T.: Requirements engineering for safety-critical systems: overview and challenges. *IEEE Softw.* **34**(4), 49–57 (2017)
9. Vuori, M.: Agile development of safety-critical software. Tampere University of Technology, Department of Software Systems (2011)
10. Stålhane, T., Myklebust, T., Hanssen, G.K.: The application of Safe Scrum to IEC 61508 certifiable software. In: *Proceedings of the 11th International Probabilistic Safety Assessment and Management Conference and the Annual European Safety and Reliability Conference* (2012)
11. Ge, X., Paige, R.F., McDermid, J.A.: An iterative approach for development of safety-critical software and safety arguments. In: *Proceedings of Agile Conference*. IEEE (2010)
12. Wang, Y., Ramadani, J., Wagner, S.: An exploratory study of applying a Scrum development process for safety-critical systems. In: *Proceedings of the 18th International Conference on Product-Focused Software Process Improvement* (2017)
13. Eleftherakis, G., Cowling, A.J.: An agile formal development methodology. In: *Proceedings of the 1st South-East European Workshop on Formal Methods* (2003)
14. Ghezzi, C., et al.: On requirements verification for model refinements. In: *Proceedings of Requirements Engineering Conference*. IEEE (2013)
15. Wynne, M., Hellesoy, A.: *The Cucumber Book: Behaviour-Driven Development for Testers and Developers*. Pragmatic Bookshelf, Dallas (2012)
16. Smart, J.F.: *BDD in Action: Behavior-Driven Development for the Whole Software Lifecycle*. Manning, New York (2015)
17. Silva, T.R., Hak, J.L., Winckler, M.: A behavior-based ontology for supporting automated assessment of interactive systems. In: *Proceedings of the 11th International Conference on Semantic Computing*. IEEE (2017)
18. Hummel, M., Rosenkranz, C., Holtz, R.: The role of communication in agile systems development. *Bus. Inf. Syst. Eng.* **5**(5), 343–355 (2013)
19. Okubo, T., et al.: Security and privacy behavior definition for behavior driven development. In: *Proceedings of the 15th International Conference on Product-Focused Software Process Improvement* (2014)

20. Lai, S.T., Leu, F.Y., Chu, W.: Combining IID with BDD to enhance the critical quality of security functional requirements. In: Proceedings of the 9th International Conference on Broadband and Wireless Computing, Communication and Applications. IEEE (2014)
21. Fucci, D., Turhan, B.: A replicated experiment on the effectiveness of test-first development. In: Proceedings of the International Symposium on Empirical Software Engineering and Measurement. IEEE (2013)
22. Fucci, D., et al.: A dissection of test-driven development: does it really matter to test-first or to test-last? *IEEE Trans. Software Eng.* **43**(7), 597–614 (2017)
23. Erdogmus, H., Morisio, M., Torchiano, M.: On the effectiveness of the test-first approach to programming. *IEEE Trans. Software Eng.* **31**(3), 226–237 (2005)
24. Kollanus, S., Isomöttönen, V.: Understanding TDD in academic environment: experiences from two experiments. In: Proceedings of the 8th International Conference on Computing Education Research. ACM (2008)
25. Hamlet, R.G.: Testing programs with the aid of a compiler. *IEEE Trans. Software Eng.* **4**, 279–290 (1977)
26. Madeyski, L.: The impact of test-first programming on branch coverage and mutation score indicator of unit tests: an experiment. *Inf. Softw. Technol.* **52**(2), 169–184 (2010)
27. Marick, B.: How to misuse code coverage. In: Proceedings of the 16th International Conference on Testing Computer Software (1999)
28. Pančur, M., Ciglarič, M.: Impact of test-driven development on productivity, code and tests: a controlled experiment. *Inf. Softw. Technol.* **53**(6), 557–573 (2011)
29. George, B., Williams, L.: A structured experiment of test-driven development. *Inf. Softw. Technol.* **46**(5), 337–342 (2004)
30. Siniaalto, M., Abrahamsson, P.: A comparative case study on the impact of test-driven development on program design and test coverage. In: Proceedings of 1st International Symposium on Empirical Software Engineering and Measurement (2007)
31. North, D.: JBehave. A framework for behaviour driven development (2012)
32. Wohlin, C., et al.: Experimentation in Software Engineering. Springer, Heidelberg (2012). <https://doi.org/10.1007/978-3-642-29044-2>
33. Falessi, D., et al.: Empirical software engineering experts on the use of students and professionals in experiments. *Empirical Softw. Eng.* **23**(1), 452–489 (2018)
34. Enoiu, E.P., et al.: A controlled experiment in testing of safety-critical embedded software. In: Proceedings of the International Conference on Software Testing, Verification and Validation. IEEE (2016)
35. Adzic, G.: Bridging the Communication Gap: Specification by Example and Agile Acceptance Testing. Neuri Limited, London (2009)
36. Gregorio, D.: How the business analyst supports and encourages collaboration on agile projects. In: Proceedings of International Systems Conference. IEEE (2012)
37. Scanniello, G., et al.: Students' and professionals' perceptions of test-driven development: a focus group study. In: Proceedings of the 31st Annual Symposium on Applied Computing. ACM (2016)
38. Crispin, L., Gregory, J.: Agile Testing: A Practical Guide for Testers and Agile Teams. Pearson Education, Boston (2009)
39. Huang, L., Holcombe, M.: Empirical investigation towards the effectiveness of Test First programming. *Inf. Softw. Technol.* **51**(1), 182–194 (2009)
40. Madeyski, L.: Impact of pair programming on thoroughness and fault detection effectiveness of unit test suites. *Softw. Process: Improv. Pract.* **13**(3), 281–295 (2008)



41. Rafique, Y., Mišić, V.B.: The effects of test-driven development on external quality and productivity: a meta-analysis. *IEEE Trans. Software Eng.* **39**(6), 835–856 (2013)
42. Haugset, B., Stålhane, T.: Automated acceptance testing as an agile requirements engineering practice. In: *Proceedings of the 45th Hawaii International Conference on System Science*. IEEE (2012)
43. Adler, R.B.: *Confidence in Communication: A Guide to Assertive and Social Skills*. Harcourt School (1977)
44. Kitchenham, B., et al.: Robust statistical methods for empirical software engineering. *Empirical Softw. Eng.* **22**(2), 579–630 (2017)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

