

Benefits and Pitfalls of Using Capture the Flag Games in University Courses

Jan Vykopal*
Masaryk University
Czech Republic
vykopal@ics.muni.cz

Valdemar Švábenský
Masaryk University
Czech Republic
svabensky@ics.muni.cz

Ee-Chien Chang
National University of Singapore
Singapore
changeec@comp.nus.edu.sg

ABSTRACT

The concept of Capture the Flag (CTF) games for practicing cybersecurity skills is widespread in informal educational settings and leisure-time competitions. However, it is not much used in university courses. This paper summarizes our experience from using jeopardy CTF games as homework assignments in an introductory undergraduate course. Our analysis of data describing students' in-game actions and course performance revealed four aspects that should be addressed in the design of CTF tasks: scoring, scaffolding, plagiarism, and learning analytics capabilities of the used CTF platform. The paper addresses these aspects by sharing our recommendations. We believe that these recommendations are useful for cybersecurity instructors who consider using CTF games for assessment in university courses and developers of CTF game frameworks.

CCS CONCEPTS

• Security and privacy; • Social and professional topics → Student assessment;

KEYWORDS

cybersecurity, capture the flag, summative assessment, homework assignment, plagiarism, learning analytics, educational data mining

ACM Reference Format:

Jan Vykopal, Valdemar Švábenský, and Ee-Chien Chang. 2020. Benefits and Pitfalls of Using Capture the Flag Games in University Courses. In *The 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20)*, March 11–14, 2020, Portland, OR, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3328778.3366893>

1 INTRODUCTION

Capture the Flag (CTF) games are widely used for cybersecurity competitions and awareness events [8]. Teams of players solve several problems of varying complexity in a limited time ranging from hours to days. Assignments usually contain only a file or the IP address of a system that has to be analyzed and very little or unclear instructions on what to do. The only clear goal is to find the flag (a string). Figuring out a feasible approach and clues provided

*A part of this work was done during a stay at National University of Singapore as a research fellow in National Cybersecurity R&D Laboratory.

SIGCSE '20, March 11–14, 2020, Portland, OR, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *The 51st ACM Technical Symposium on Computer Science Education (SIGCSE '20)*, March 11–14, 2020, Portland, OR, USA, <https://doi.org/10.1145/3328778.3366893>.

by the file or system is considered a part of the game. Players can submit (un)limited number of attempts without any penalty and immediately see whether their submission is correct as well as the score of their competitors. Some games also provide hints, which may cost some points the players already earned in the game.

On the contrary, summative assessment in university courses is usually less interactive and dynamic. Final exams, midterm tests, and homework assignments are traditionally completed individually rather than in a group. They provide more information in the assignment, and students must submit more evidence than just one string in a time range of several hours (exams and tests) or weeks (homework). In addition, only one submission is accepted, and students have to wait several days or weeks for its marking. Hints might be provided only for homework.

From the comparison above, CTF games seem to be a better assessment method of skills acquired during the semester, especially for large classes. Gamification features should bring students a more enjoyable learning experience, including not only technical skills but also teamwork. Instructors should benefit from automatic scoring of students' submissions and spend time consumed by the manual marking of students' submissions more efficiently.

However, there is not much research literature investigating the suitability of the use of CTF games in university courses. This paper fills this gap by discussing the results of a case study of using jeopardy CTF games as homework assignments in an introductory computer security course taught at a public university. We are interested in how students apply taught skills and knowledge in CTF games and what are the advantages and drawbacks of using such gamification in the context of tertiary education. We share our experience and provide recommendations for instructors and developers of CTF platforms. We also contribute to the field by developing two open-source software plugins for a popular jeopardy CTF platform [6].

2 RELATED WORK

Although it is generally accepted that CTF games are an engaging and popular education tool, only a few works in the context of tertiary education can be found in the literature.

2.1 CTF in competitions

The original purpose of CTF was competitive [6], and most CTFs remain highly focused on competition [19]. Similarly to programming contests [11], their goal is to showcase and evaluate the performance of already skilled participants [3]. Competitive CTFs cover many cybersecurity topics and offer recruitment opportunities and reputation building [18] to the participants. Next, a competitive setting can motivate and engage students, especially those who

are attracted to cybersecurity, have extensive prior experience, or possess skills required by the competition [20]. By solving the competition tasks, participants deepen their understanding of cybersecurity [23] and practice creative approaches to both known problems and those outside the traditional curriculum [4]. Moreover, competitions offer considerable learning benefits also before and after the event. Preparing for a CTF involves developing new tools, studying vulnerabilities, and discussing strategies [22], which exposes participants to new skills [9]. After a CTF, the competitors or organizers publish *wroteups*: walkthroughs that report solutions and explain the vulnerabilities involved in the game. Both writing and reading these is beneficial [21].

However, the competitive setting of CTF games might discourage or even alienate some students [1, 19], especially beginners [20], for three main reasons. First, the tasks are usually too difficult for less-experienced participants [24]. Second, some of the tasks are also intentionally ambiguous, require a lot of guessing or include artificial obstacles to make them harder to solve [7]. Third, the participants receive limited individual feedback about their progress. They are often unsure if they are on the right track and usually receive only information about whether the submitted flag was correct or wrong [7]. Although the unguided progress inherent for competitions suits advanced learners and can lead to creative solutions [10], it is highly ineffective for beginners [12]. Without guidance, novice students miss essential learning goals and take longer to learn a concept [23].

2.2 CTF in university courses

Class CTFs [14] are small-scoped competitions that challenge teams of students against each other in realistic attack-defense scenarios. They have been used in an undergraduate course Introduction to Security at University of Southern California in 2013 and 2014. Students were graded based on their contributions to the team. However, no further details about the assessment are provided.

Chothia and Novakovic [5] studied whether CTFs are effective as an assessment tool in academic cybersecurity courses. They showed that the ability of students to acquire flags in CTF-style challenges is highly correlated with their marks from the written submissions for the *same* challenges, and that flag-only marking may lead to more widespread plagiarism.

Leune and Petrilli [13] conducted a study in which ten undergraduate students taking a cybersecurity class were surveyed before and after a two-week-long CTF. The students developed stronger practical skills by participating in the CTF and overwhelmingly enjoyed it. Nonetheless, no relation of performance in CTF and other assessments in the class is reported.

The automatic generation of CTF problems for preventing flag plagiarism is discussed in [2]. Each player receives a different version of a task, which leads to a unique flag that must be submitted. Still, 14% of teams participating in a case study submitted at least one shared flag. Since the game lasted 12 days, limiting online flag disclosure was challenging. Similarly to the automatic problem generation, a scaffolded, metamorphic CTF for reverse engineering proposed by Feng [3] provides each student with a unique binary for analysis.

3 CTF GAMES AND TEACHING CONTEXT

This paper presents experience learned from two jeopardy CTF games that were a part of a computer security course. The course was taught in English at a public university in Singapore in the first semester of the academic year 2018/2019. The games were used as homework assignments in an introductory course on information and system security. We collected game events generated by students using the CTF portal, answers from two surveys, and analyzed students' marks from other forms of summative assessment of the course. The study was approved by the Institutional review board of the university.

Procedure. At the beginning of the semester, all students of the course were asked to participate in the study by the first author, who was a guest instructor. Those who agreed to participate did not receive any incentives or reimbursement for taking part in the study. First, they filled in an introductory survey about basic demographic information, including their work and study experience. Then, they completed individual homework assignments set as CTF games. Participants' interactions with the CTF portal (e.g., login, submissions of correct and incorrect solutions, or displaying hints) were automatically logged. Finally, the participants were surveyed about their learning experience with CTFs in the course. The marks from other forms of assessment in the course were obtained directly from the course gradebook in the university learning management system. All collected data were analyzed to find common patterns and anomalies describing the performance and engagement of various students in the CTF games.

Participants. Out of 120 students enrolled in the course, 37 students agreed to participate in the study. The introductory survey was completed by 25 students. The median age of the participants was 23 ($\sigma = 1.49$). Only two participants had played any CTF game before. No participant was a member of any CTF team. Six were employed in a part-time IT-related job.

3.1 Features of the selected CTF games

The CTF games have common features determined by using CTFd version 1.2.0, a popular open-source CTF framework [6]:

- *Challenge value* – Each problem (question) has a set point value, which is known to students. This value usually reflects the difficulty of the problem.
- *Immediate response* – After a flag (answer) is submitted to the CTF portal, players immediately see whether it is correct or not. They can also submit an unlimited number of incorrect flags with no penalties.
- *Scoreboard* – The total current score of all students is available to all students playing the game. Anonymized names of the players were displayed to protect their privacy.
- *Scaffolding* – Some challenges offer one or more hints. Each hint has its set cost: penalty points that will be deducted from the current player's score once the hint is displayed to the player. Hints can also be provided for free (cost 0 points). Hints can be released together with challenge assignment or later during the game, based on current players' progress.
- *Challenge chains* – By default, there is no hierarchy or dependency of challenges in the CTF portal, so the players can

display and solve any challenge of their choice. Although this design choice is suitable for competitions, we feel that some guidance on which challenge should be solved first can be helpful in educational settings. Our team, therefore, developed a plugin for CTFd, which provides a feature of linear unlocking of challenges [16]. Instructors can then group some challenges to chains with a defined order of challenges that are unlocked to players once they solve the previous challenge in the chain. For instance, challenges to practice SQL injection, which are of increasing difficulty, can be locked in a chain to guide the student to start with the easiest challenge. This approach can be viewed as a means of achieving game balance [17].

3.2 CTF content and parameters

The course contains two homework assignments (A1 and A2), which were run as CTF games for individuals. A1 consisted of 8 challenges covering topics taught in the first part of the semester: substitution ciphers, hashing, symmetric and asymmetric cryptography, RSA, and cryptanalysis. A2 consisted of 15 challenges on topics of the second part of the semester: network traffic analysis, port knocking, access control, buffer overflow, command injection, format string attack, and SQL injection.

The difficulty of challenges varied largely – from a simple execution of one command to a multi-step solution involving binary debugging and writing a helper exploit script. The difficulty of each challenge was indicated by assigning to a respective category (basic, medium, advanced) and its point value (from 5 to 25, median 15).

A1 contributed by 10% to the final grade and was due in 26 days, and A2 by 15% was due in 24 days. Both assignments included optional bonus challenges for those who were interested in exploring the topics in more depth.

The majority of challenges contained hints which cost 0 penalty points. Next, there was one chain of 3 challenges in A1 and two chains of 3 challenges each in A2.

The first assignment also contained optional challenges for familiarization of students with CTF portal and conventions: flag format, hint displaying, and unlocking challenge in the chain.

4 EXPERIENCE REPORT

4.1 Students’ performance

We hypothesized that students who are struggling with CTFs games exercising topics taught at lectures and tutorials would struggle in other forms of continuous assessment and at the exam. Therefore, we sought for dependencies between variables collected in the study. The most indicative variables seemed to be the total score from both CTF games and the number of wrong flags submitted by each student. We used Spearman’s rank correlation because our dataset is not normally distributed (e.g. one participant submitted extreme numbers of wrong flags in both A1 and A2). To support our hypothesis, we then focused only on students who finally achieved lower than average total marks from all types of course assessment, excluding homework assignments (i.e., exam, midterm quiz, group presentation, and online quiz).

	CTFs	Bonuses	Midterm	Exam	All	Wrong flags
CTFs	1.00	0.63	0.35	0.31	0.31	0.11
Bonuses		1.00	0.50	0.49	0.50	-0.15
Midterm			1.00	0.89	0.92	-0.47
Exam				1.00	0.99	–
All but CTFs					1.00	–
Wrong flags						1.00

Table 1: Statistically significant ($p \leq 0.05$) correlation coefficients of selected variables captured in both CTF games

4.1.1 Observations about all students. Table 1 shows statistically significant ($p \leq 0.05$) Spearman’s rank correlation coefficients of variables captured in A1 and A2 games. The strongest non-obvious positive correlation among variables capturing performance in CTF and other types of assessment are marked red. These correlations were present between 1) the total CTF score of A1 and A2 including bonuses (*Bonuses*) and marks from all other types of assessment in the course (*All*): exam, midterm quiz, group presentation, online quiz ($r = .50, p \leq 0.001$), 2) the total CTF score including bonuses (*Bonuses*) and marks from the midterm (*Midterm*) ($r = .50, p \leq 0.001$), and 3) the total CTF score (*Bonuses*) and the exam ($r = .49, p \leq 0.001$). The strongest non-obvious negative correlation (marked blue) is between the total number of wrong flags in A1 and A2 (*Wrong_flags*) and marks from the midterm ($r = -.47, p \leq 0.05$).

4.1.2 Observations about low-performing students. Besides obvious strong correlations, we found interesting strong positive correlations between the score from A2 and the A2 game session duration ($r = .61, p \leq 0.005$), and between the score from A2 and time difference between solving the first and last challenge in A2 ($r = .59, p \leq 0.01$).

4.1.3 Discussion. Although we observed some statistically significant correlations between variables mined from CTF games and other forms of the assessment, the correlation coefficients range only from -0.5 to 0.61. We believe three factors were affecting these results. First, the total score of A1 and A2 has a skewed distribution. Their medians were the maximum possible scores without bonus challenges (i.e. 100 and 150, respectively) and the standard deviations 8.4 and 31.8, respectively. Second, all the hints were for free. Their usage is not reflected in the score, which remains the same even for students who needed different levels of help. Third, the used CTF portal did not log important game events. For instance, an event of displaying the challenge by a student can be used to analyze how much time the student spent on solving the challenge. However, this event is currently not logged by the CTF portal and cannot be easily and reliably reconstructed from other sources such as webserver access logs. This negatively affects the feasibility or accuracy of some variables relying on the missing event. The manual reconstruction of these events from the logs revealed the strongest non-obvious correlation of any two variables of our dataset.

4.1.4 Recommendation for instructors.

Examine what kind of game events are logged by the CTF portal. In our case, both assignments were open for almost a month, which might eventually enable the vast majority of students to complete all challenges with the full score regardless of how much time they needed. In addition, the instructors did not want to discourage




Figure 1: Time differences between a hint display and a correct flag submission in A1. Green triangles denote means.

students from displaying hints by deducting penalty points. Therefore, displaying these “free hints” did not affect the variance of the total score. We recommend focusing on logging features of a CTF platform, in particular, more detailed information about the game progress, such as the duration of students’ interactions with the platform. This will deliver more accurate metrics than the points awarded for the submission of correct flags.

4.2 Usefulness of hints

The majority of challenges was equipped with one or more hints. We analyzed how they helped students with solving the challenges.

4.2.1 Observations. Figure 1 and Figure 2 show times between displaying hints and solving the respective challenge. Only the hints used more than 10 times are plotted. In both games, there were challenges that the students did not solve in a reasonable time despite viewing hints. High median times of 1 day or more can be observed in both figures. In contrast, only 7 challenges out of 19 were solved with the median time less than 1 hour after displaying the hint.

The answers in the after-game survey suggest that some hints did not reach their aim. Only one student assessed hints as *Very much useful*, two as *Much useful*, 8 as *Moderately useful*, 4 as *Slightly useful* and 1 as *Not useful at all*. In addition, 6 students out of 20 mentioned hints in their answer to the question “If you could choose one element of your CTF experience to improve, what would it be?”

4.2.2 Discussion. The major differences among median times of individual hints indicate some hints were more helpful than others. This can be illustrated by hint 1 of challenge 2 in A1. This challenge was, in fact, straightforward instruction on how to display and use hints. The median time was the lowest (only 1 minute). The instant feedback provided by students on hints right after solving the challenge in the CTF portal supports that some hints were more useful than others. One example is challenge 2 in A2: all 7 students who provided feedback assessed hints as *Rather Useful*, *Somewhat Useful* or *Useful*. The challenge was easy, and hints were clear and guiding, e.g., *You need to view the _image file_ to get the flag*.

4.2.3 Recommendations for instructors.

Indicate what a hint is about. The information about hint cost (if any) is not enough for students to decide whether they will benefit from displaying the hint. We recommend adding a short description of what they can expect, such as “what tool to use”, particularly

in challenges offering two or more hints. Otherwise, students may display a hint which tells them what they already figured out.

Test challenge assignments and hints before the game. Ask teaching assistants or peer instructors to test the challenge descriptions and hints to balance what should be placed in the challenge assignments and what can be left for one or more hints. While the challenge assignments can be a bit fuzzy, hints should be clear and straightforward.

Prepare backup hints. Although hints have been tested, students may still struggle with some (advanced) challenges. Monitor the ongoing game (submissions, wrong flags, and hint usage) and be ready to add a new hint if needed.

4.2.4 Recommendations for developers of CTF frameworks.

Actively offer hints. In our experience, some students tend to beat the challenge without displaying any hints even though the hint may speed up their progress. Consider adding a feature which will offer a hint to a student after some time of the challenge solving.

Support adaptive hints. Any step toward dynamically generated hints would be beneficial. If the CTF framework provides logs capturing game events and even players’ behavior in a virtual environment, these data can be used as input to a module serving the appropriate hint at the right time.

4.3 Flag sharing in individual CTF

We investigated four unusual patterns indicative of plagiarism.

4.3.1 Observations.

Submissions of the same flag in a time vicinity. We supposed that if students are solving challenges individually in a time frame of almost a month, it is unlikely that they will submit the correct flag (almost) at the same time. Figure 3 presents how many pairs of students submitted the correct flag for the same challenge shortly after each other. Two challenges were solved by 8 and 7 pairs of students, respectively, and the rest of the challenges by 4 and fewer pairs.

Correct flag submitted as an incorrect flag to another challenge. We observed 8 cases in total. The time between the submission of any valid flag of other challenges as the wrong flag varied from 14 seconds to almost 18 hours. Notably, there were two cases in A2 where a flag from a still locked challenge was submitted as an incorrect flag to the preceding challenge.

Challenges solved without downloading the file to analysis. Some challenges require students to visit online services running at other machines than the CTF portal. Others contain attachments that have to be downloaded first. For this analysis, we inspected logs of the web server hosting the CTF portal. In total, seven challenges contain 12 attached files required for solving the challenges. We discovered 11 cases of solving any such challenge by 6 different students without prior download of a required file.

Quick solves of consecutive challenges. Finally, we focused on linearly locked challenges in both assignments. We estimated the minimal possible solve time by an expert player who immediately




Figure 2: Time differences between a hint display and a correct flag submission in A2. Green triangles denote means.




Figure 3: Cumulative number of pairs of students who both submitted correct flag for the same challenge within 10 minutes.




Figure 4: Time between submission of correct flags of two consecutive (locked) challenges. Blue lines indicate the minimal possible solve time of the first challenge.

figures out all steps required for solving a challenge, performs these steps without any mistakes, and submits the correct flag to the portal. The minimal possible solve times range from 20 seconds to 2 minutes. Figure 4 depicts the time differences of the submissions of the correct flag of a locked challenge and a preceding challenge in the same chain. A1 contains one chain with 3 challenges (two pairs of consecutive challenges) and A2 two chains with 3 challenges each (four pairs of consecutive challenges). We discovered 7 very quick solves ranging from 9 to 53 seconds in the fourth pair of challenges in A2 with the minimal possible solve time of 1 minute and 15 seconds.

4.3.2 Discussion. These different views on game events revealed some students might have obtained flags unexpectedly quickly, without downloading the necessary file(s), had submitted flags almost at the same time, or had submitted the correct flag to another challenge. While the time vicinity of submissions may report only

weak indications of plagiarism, others can be considered as more serious pieces of evidence (such as quick solves of consecutive locked challenges). In A2, instructors questioned such students. Three of them eventually confessed they used flags shared by their peers. Some students argued this was only a coincidence since they consulted their approach in a group and then solved the challenge and submitted the flag each on their own.

4.3.3 Recommendations for instructors.

Set rules for students' collaboration during the game in advance. Decide what will constitute plagiarism in your class. Is any discussion about challenges among students forbidden, or do you allow non-detailed discussions about challenge principles or techniques that can be used? Communicate these rules clearly and explicitly to students.

Inform students about how you will check suspicious submissions in advance. Describe a procedure that will be applied if instructors spot suspicious behavior. For instance, the instructor may (randomly) select several students for in-person demonstration of how they solved particular challenges.

Structure related problems to challenge chains. Challenge chains help not only with revealing plagiarism but also explicitly guide students what must be solved first.

4.3.4 Recommendations for developers of CTF frameworks.

Support challenge chains or dependencies. If your platform does not have this feature, add it.

Provide built-in analyses for revealing flag sharing. All analyses we performed were done outside the CTF platform using ad-hoc scripts and third-party tools. However, these analyses work only with generic game events such as time and date of flag submissions so that they can be run automatically for any CTF game. If the results of these analyses were easily accessible in the CTF portal, instructors would definitely benefit from them.

4.4 Students' perceptions of the CTF games

One week after the games were over, we asked students for their voluntary feedback in the online survey. Table 2 lists the questions.

4.4.1 Observations. The crucial question was the first one. The vast majority of students (13 out of 16) would prefer CTF games, only

No.	Question
Q1	Would you rather complete the Capture the Flag games or normal homework assignments in your future security courses?
Q2	How useful did you find the hints?
Q3	How useful was the instant feedback on your submissions (a response of the CTF server whether your flag is correct or not)?
Q4	How useful was seeing your overall progress in assignments?
Q5	How motivating was the scoreboard?
Q6	How frequently have you talked about your score with your peers during the semester?
Q7	How much did you like the unlocking feature?

Table 2: Wording of questions in the after-game survey.




Figure 5: After-game students’ assessment of CTF features.

two students normal assignments and one student was not sure. Students who would prefer the games liked that the games were fun, hands-on, more interactive, objective, allow to learn to work with security tools, and allow lots of trial and error in exploration. The two students who would prefer the normal assignments considered the games difficult and very time-consuming. One of them felt he had not learned much from them. The student who was not sure mentioned that “the normal assignment is little harder to simply copy than CTF games.”

The answers to the other questions (Q2–Q7) are depicted in Figure 5, which shows students’ assessment of various game features on a 5-point Likert scale. Students could elaborate on their answers to questions 2, 5, and 7. Student comments’ of answers to questions 2 was already discussed in Section 4.2. Regarding the scoreboard (Q5), some students would be more motivated if the scoreboard listed real (non-anonymized) names of their peers. Finally, the challenge unlocking (Q7) was used appropriately since it allows to “solve simpler challenges first and used the thought process and methods to help solve the harder challenges.”

4.4.2 Limitations. The final after-game survey was answered by 16 out of 37 students. This could introduce some bias since the rest of the students who did not provide any input may not have liked their CTF experience.

5 CONCLUSIONS

Capture the Flag games can be an enjoyable way of assessment in university computer security courses. We studied the use of jeopardy CTF games in the summative assessment. Based on this experience, we provided recommendations for both instructors and developers of CTF platforms.

5.1 Summary of our experience

We conclude that replacing traditional homework assignments by CTF games is generally favorable for both instructors and students. The instructors can save time spent on marking the students’ submissions and enable students learning practical skills in an interactive and enjoyable way. The vast majority of students who answered the after-game survey in our study would rather complete the CTF games than regular homework assignments in their future security courses.

We highlighted several pitfalls of using CTFs in summative assessment. Instructors should carefully consider the game format, scoring (distribution of points of game challenges and hint costs), a CTF platform for running the game, and game duration.

First, although CTFs lasting several weeks provide more opportunities and less stressful environment for exploring and learning the topic than the hours-long CTFs, they are much more vulnerable to flag sharing between students than invigilated intensive CTFs.

Next, the interactive nature of the games, particularly instant feedback about the correctness of a solution, can be utilized for identifying students at risk only if the CTF platform provides advanced analyses of students’ progress. The scoreboard is not sufficient in weeks-long games since there is a high chance the vast majority of students will finish almost all tasks, as we witnessed in our study.

Finally, the challenges presented in CTFs usually require investigating several dead ends before they are solved. Current CTF platforms, however, provide only static scaffolding by offering hints which do not consider the present performance and experience of a player. This results in the low usefulness of hints offered for medium and hard challenges. Serving dynamically generated hints thus presents an interesting direction for further research.

5.2 Practical contributions

This paper reports not only the results of the study but also presents several replicable methods of analysis of game events. These methods apply to any CTF game with common basic features regardless of the content of challenges. For example, an excessively high number of submitted wrong flags may indicate some students at risk, or unusually quickly solved challenge may reveal a flaw in challenge design or illicit flag sharing. The only prerequisite for these analyses is the availability and sufficient level of detail of game events logged by CTF platforms.

Finally, we also contributed to the current practice of running CTF games by developing two open-source software plugins. These extend the popular CTFd platform [6] by features of linear unlocking of challenges and collecting players’ feedback right after solving the challenge. The first plugin [16] is useful for structuring the challenges and revealing shared flags. The second [15] serves both CTF organizers and researchers interested in players’ thoughts about the solved challenge. In our future work, we plan to develop more plugins implementing learning analytics methods and other features that will enhance students’ learning experience and help instructors to design and run CTFs more efficiently.

ACKNOWLEDGMENTS

This research was supported by ERDF project *CyberSecurity, Cyber-Crime and Critical Information Infrastructures Center of Excellence*

(No. CZ.02.1.01/0.0/0.0/16_019/0000822). We also thank NUS Greyhats for their effort spent on the homework assignments and Heng Le Ong for developing the two CTFd plugins.

REFERENCES

- [1] John Ayccock, Andrew Groeneveldt, Hayden Kroepfl, and Tara Copplestone. 2018. Exercises for Teaching Reverse Engineering. In *Proceedings of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE 2018)*. ACM, New York, NY, USA, 188–193. <https://doi.org/10.1145/3197091.3197111>
- [2] Jonathan Burket, Peter Chapman, Tim Becker, Christopher Ganas, and David Brumley. 2015. Automatic Problem Generation for Capture-the-Flag Competitions. In *2015 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 15)*. USENIX Association, Washington, D.C., 8. <https://www.usenix.org/conference/3gse15/summit-program/presentation/burket>
- [3] Wu chang Feng. 2015. A Scaffolded, Metamorphic CTF for Reverse Engineering. In *2015 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 15)*. USENIX Association, Washington, D.C., 8. <https://www.usenix.org/conference/3gse15/summit-program/presentation/feng>
- [4] Peter Chapman, Jonathan Burket, and David Brumley. 2014. PicoCTF: A Game-Based Computer Security Competition for High School Students. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*. USENIX Association, San Diego, CA, 10. <https://www.usenix.org/conference/3gse14/summit-program/presentation/chapman>
- [5] Tom Chothia and Chris Novakovic. 2015. An Offline Capture The Flag-Style Virtual Machine and an Assessment of Its Value for Cybersecurity Education. In *2015 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 15)*. USENIX Association, Washington, D.C., 8. <https://www.usenix.org/conference/3gse15/summit-program/presentation/chothia>
- [6] Kevin Chung. 2017. Live Lesson: Lowering the Barriers to Capture The Flag Administration and Participation. In *2017 USENIX Workshop on Advances in Security Education (ASE 17)*. USENIX Association, Vancouver, BC, 6. <https://www.usenix.org/conference/ase17/workshop-program/presentation/chung>
- [7] Kevin Chung and Julian Cohen. 2014. Learning Obstacles in the Capture The Flag Model. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*. USENIX Association, San Diego, CA, 7. <https://www.usenix.org/conference/3gse14/summit-program/presentation/chung>
- [8] CTFtime. 2019. CTFtime.org / All about CTF (Capture The Flag) – Past events. Retrieved October 10, 2019 from <https://ctftime.org/event/list/past>
- [9] Chris Eagle. 2013. Computer security competitions: Expanding educational outcomes. *IEEE Security & Privacy* 11, 4 (2013), 69–71.
- [10] Manu Kapur and Katerine Bielaczyc. 2012. Designing for productive failure. *Journal of the Learning Sciences* 21, 1 (2012), 45–83.
- [11] Keith Kirkpatrick. 2016. Coding As Sport. *Commun. ACM* 59, 5 (April 2016), 32–33. <https://doi.org/10.1145/2898967>
- [12] Paul A Kirschner, John Sweller, and Richard E Clark. 2006. Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational psychologist* 41, 2 (2006), 75–86.
- [13] Kees Leune and Salvatore J. Petrilli, Jr. 2017. Using Capture-the-Flag to Enhance the Effectiveness of Cybersecurity Education. In *Proceedings of the 18th Annual Conference on Information Technology Education (SIGITE '17)*. ACM, New York, NY, USA, 47–52. <https://doi.org/10.1145/3125659.3125686>
- [14] Jelena Mirkovic and Peter A. H. Peterson. 2014. Class Capture-the-Flag Exercises. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*. USENIX Association, San Diego, CA, 8. <https://www.usenix.org/conference/3gse14/summit-program/presentation/mirkovic>
- [15] Heng Le Ong. 2018. Challenge feedback plugin for CTFd. Retrieved 2019-11-25 from <https://github.com/janvykopal/ctfd-challenge-feedback>
- [16] Heng Le Ong. 2018. Linear unlocking challenges plugin for CTFd. Retrieved 2019-11-25 from <https://github.com/janvykopal/ctfd-linear-unlocking>
- [17] Portia Pusey, Sr. David Tobey, and Ralph Soule. 2014. An Argument for Game Balance: Improving Student Engagement by Matching Difficulty Level with Learner Readiness. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*. USENIX Association, San Diego, CA, 5. <https://www.usenix.org/node/183458>
- [18] Portia Pusey, Mark Gondree, and Zachary Peterson. 2016. The outcomes of cybersecurity competitions and implications for underrepresented populations. *IEEE Security & Privacy* 14, 6 (2016), 90–95.
- [19] Clark Taylor, Pablo Arias, Jim Klopchic, Celeste Matarazzo, and Evi Dube. 2017. CTF: State-of-the-Art and Building the Next Generation. In *2017 USENIX Workshop on Advances in Security Education (ASE 17)*. USENIX Association, Vancouver, BC, 11. <https://www.usenix.org/conference/ase17/workshop-program/presentation/taylor>
- [20] David H. Tobey, Portia Pusey, and Diana L. Burley. 2014. Engaging Learners in Cybersecurity Careers: Lessons from the Launch of the National Cyber League. *ACM Inroads* 5, 1 (2014), 53–56.
- [21] Erik Trickle, Francesco Disperati, Eric Gustafson, Faezeh Kalantari, Mike Mabey, Naveen Tiwari, Yeganeh Safaei, Adam Doupé, and Giovanni Vigna. 2017. Shell We Play A Game? CTF-as-a-service for Security Education. In *2017 USENIX Workshop on Advances in Security Education (ASE 17)*. USENIX Association, Vancouver, BC, 10. <https://www.usenix.org/conference/ase17/workshop-program/presentation/trickle>
- [22] Giovanni Vigna, Kevin Borgolte, Jacopo Corbetta, Adam Doupé, Yanick Fratantonio, Luca Invernizzi, Dhilung Kirat, and Yan Shoshitaishvili. 2014. Ten Years of iCTF: The Good, The Bad, and The Ugly. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*. USENIX Association, San Diego, CA, 7. <https://www.usenix.org/conference/3gse14/summit-program/presentation/vigna>
- [23] Richard Weiss, Frankly Turbak, Jens Mache, Erik Nilsen, and Michael E. Locasto. 2016. Finding the Balance Between Guidance and Independence in Cybersecurity Exercises. In *2016 USENIX Workshop on Advances in Security Education (ASE 16)*. USENIX Association, Austin, TX, 8. <https://www.usenix.org/conference/ase16/workshop-program/presentation/weiss>
- [24] Joseph Werther, Michael Zhivich, Tim Leek, and Nickolai Zeldovich. 2011. Experiences in Cyber Security Education: The MIT Lincoln Laboratory Capture-the-flag Exercise. In *Proceedings of the 4th Conference on Cyber Security Experimentation and Test (CSET'11)*. USENIX Association.