Practical Lessons From Creating the Control-Alt-Hack Card Game and Research Challenges for Games In Education and Research

Tamara Denning, *Computer Science & Engineering, University of Washington*
Adam Shostack
Tadayoshi Kohno, *Computer Science & Engineering, University of Washington*

Abstract

We designed, produced, distributed, and evaluated *Control-Alt-Hack™*: a recreational tabletop card game intended to promote a casual awareness of high-level computer security concepts. Our experiences throughout this process gave us insights and opinions regarding the creation of games, their role in educational or outreach contexts, and opportunities and challenges for the research community. In particular, we: (1) provide a logistics-oriented reflection on our experiences, including a list of the work roles that were involved in producing the game and a timeline of the creation process; and (2) step back to consider higher-level issues for the community, including the role of games in the classroom and the challenges behind conducting and publishing evaluations of game-based learning.

1. Introduction

In November 2012 we released *Control-Alt-Hack: White Hat Hacking for Fun and Profit™* —a recreational tabletop card game where players take on the roles of white hat hackers at a security consulting company. Using strategy and luck (via dice rolls), they use their character's skills (e.g., Hardware Hacking, Network Ninja, Social Engineering) to attempt a variety of security-themed narratives.

Control-Alt-Hack was not designed to be an educational game that teaches specific knowledge or skills; instead, the game was intended to prioritize recreation and to facilitate voluntary gameplay in social settings. The narrative content of the game was designed to casually increase familiarity with high-level security concepts: for example, the creativity of attackers, the potential human impacts of different kinds of attacks, and the broad range of technology classes that are affected by computer security issues. We prioritized aspects that we felt would help the game fulfill its purpose, such as relatively high production quality (see Figures 1, 2, and 3) and availability via a mass-market online retailer.

In this paper, we reflect on our experiences designing, producing, distributing, and evaluating our game. Section 2 provides a brief overview of the game premise and pointers to topics covered in our previous publication. Section 3 focuses on providing logistical details pertaining to the game development: both to paint a picture of the process and to provide consideration points for researchers looking to produce games. Section 4 shifts gears and focuses on two major research challenges that we believe provide obstacles and opportunities to researchers looking to employ games in computer security education.

2. Game Premise and Previous Publication

The following is the premise of Control-Alt-Hack:

> *You and your fellow players work for Hackers, Inc.: a small, elite computer security company of ethical (a.k.a., white hat) hackers who perform security audits and provide consultation services. Their motto? "You Pay Us to Hack You."*
>
> *Your job is centered around Missions—tasks that require you to apply your hacker skills (and a bit of luck) in order to succeed. Use your Social Engineering and Network Ninja skills to break the region's power grid, or apply a bit of Hardware Hacking and Software Wizardry to convert your robotic vacuum cleaner into an interactive pet toy...no two jobs are the same. So pick up the dice, and get hacking!*

Figure 1 shows the game box and contents. Figures 2 and 3 show some of the game art and card contents.

Our previous publication [1] provides additional information on:

- Our target audiences for the game.
- Our goals in producing the game.
- Choosing a gaming mechanic to license.
- Our feedback process.
- Writing card content while: including technical content, mapping to the original game mechanics, offering comprehensibility (including brevity), and incorporating humor.

- Example technical security works that inspired card topics.
- An evaluation of the game in educational settings. The evaluation was performed via (primarily) open-ended surveys distributed to educators who had previously received copies of the game.

3. Logistics and the Creation Process

In this section, we provide selected logistical details from the creation process for the game.

It should be noted that, rather than design a game from scratch, we chose to license a preexisting game mechanic (e.g., card count, card attributes, basic game rules). We licensed *Ninja Burger* from Steve Jackson Games [7], then redid all textual and visual content to create an isomorphic remap. Licensing a pre-existing mechanic allowed us both to skip playtesting mechanics and to utilize the expertise of a company whose purpose is to create successful games.

3.1. Involved Parties

Below we list some of the parties who were involved in the logistics of the game's creation. This list is meant to both serve as an indication of some of the subprocesses of the game's design, as well as to indicate some of the areas of expertise on which it is useful to draw.

Lead Game Designers: The authors served as lead game designers. We created all textual content and coordinated/gave input to all subprocesses.

Mechanics Designer(s): As mentioned above, we chose to license a pre-existing mechanics. To our knowledge, game mechanics are not legally protected; however, there are good reasons for licensing mechanics. One reason is to support the original game designer(s). Another reason is to generate goodwill among game fans—or at least to avoid the ill will that might be associated with using someone else's mechanic. We found these points particularly pertinent given the fact that we were always intended to make the game available for sale (in order to reach a broader audience). Another reason to license game mechanics is that the game designer might be able to provide feedback and comments on drafts of the new game.

License Coordinator(s): We sought the aid of someone more familiar with the industry to help negotiate terms for the license agreement. Potential details to consider are: (a) whether the license payment is a flat payment or a royalty; (b) how many copies of the game may be produced under the license; (c) whether or not a certain license or royalty rate is guaranteed for future license arrangements; (d) what reporting, if any, is required on product print runs or sales; and (e) whether or not aspects of the produced game may be sublicensed.

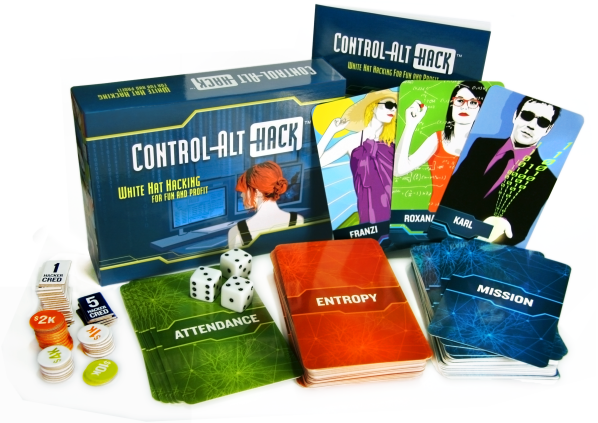

Figure 1. A photo of the game box and contents. Photo Credit: Juliet Fiss. © University of Washington

Graphic Designers: We worked with a local graphic design firm on the visual characteristics of the game. The firm had previous experience with designing for a game (Elevation of Privilege [4]).

Graphic designers are (or can be) distinct from illustrators. Illustrators create art such as the portraits on the character cards (see Figure 2); graphic designers, on the other hand, are concerned with elements of a visual layout such as colors, fonts, icons, and patterns (see Figures 1 and 3). The work that goes into a graphic design addresses issues such as aesthetics (for both genders), readability and visual comprehensibility, and colorblindness. In our case, the illustrator created the character portraits; the graphic designers did everything else.

Some details to consider when creating a contract with a graphic designer are: (a) how many different options are explored before choosing a general design direction; (b) how many iterations of different design elements are included; (c) whether or not the graphic design firm will be coordinating with the illustrator and/or the printer, and that they will be delivering files in the format the printer needs; (d) whether or not they will be sending over (or coordinating with the printer to send over) print proofs that accurately represent the colors of the end print product; and (e) whether box design, instruction booklet design, etc. will be included.

Illustrator: We located an illustrator from their online portfolio and collaborated remotely. As mentioned in [1], we chose the illustration style to further our goals for the game. Some details to consider when creating a contract with an illustrator are: (a) securing the copyright to use the produced work in all envisioned scenar-

ios (e.g., the game, merchandise, advertising); (b) securing the right to all work products that might be desired, such as layered vector files; (c) ensuring that the illustrator is positioned to offer the rights to all their work; (d) how many rough illustration directions will be explored before a style is chosen (for the game, or per illustration); and (e) how many minor illustration edits will be made.

Manufacturing Facility: There are a number of options for printing physical games, including print-on-demand services (and Kickstarter campaigns, for generating funding). We chose to use a large manufacturing facility with a print run of ~7,500 games; with such printers, large price drops might not be available unless at least 5,000 units are ordered. Roughly speaking, these options trade off: (a) minimum order size required; (b) cost per unit; (c) control over characteristics of game components; (d) quality of game components; (e) turnaround time; and (f) up-front cost required.

Before the full print run ships, the printer sends over a proof copy. One major purpose of the proof copy is to check that the printed colors appropriately match the intended colors. Ideally, the graphic designers are involved in the proofing process. Some people choose to go on site with the printer to iterate more closely on color matching; we worked with the printer remotely.

Another thing to consider is funding: many printers require a large deposit or full payment up front, while universities may have the policy of paying for work delivered.

Production Manager: Some manufacturing facilities are not available to the general public, and will only work through production managers. Additionally, someone in this position has insights into issues such as typical practices, terminology, and related details such as safety testing. With our production manager, we considered properties of the materials that we would be using including: (a) cardstock; (b) card coating; (c) box style; (d) box material; and (e) packaging.

It is useful to price with the production manager/printer early on so that the graphic designers know the dimensions for which they are designing and whether they are using a black-and-white palette, a 3-color palette, or a full-color palette.

Trademark Lawyer(s): Part of our research goal was to make the game available for sale as a regular—rather than a limited-audience or educational—game. As one result of this goal, we decided to file for a trademark on the game name. Making this decision early on in the process had cascading effects. For example, several of our first name choices were inadvisable due to potential conflicts. The game name, in turn, had direct effects on the logo design.

Distributor: Different models are available for distribution. We chose to use Amazon as the distributor (Fulfillment by Amazon) in addition to the retail platform. It would be useful to work closely with someone who has undergone the process before, as there are a lot of details and fees to keep track of that are not necessarily apparent. For example: (a) the Amazon seller platform does not automatically activate the collection of sales taxes for the seller's home state, which the seller is responsible for submitting to the state; and (b) the default process for shipping stock to Amazon fulfillment warehouses may split between multiple warehouse destinations across the country, and the number and location of the assigned warehouses that will be assigned to a shipment vary from week to week.

Retailer: We chose to coordinate the sales portion of Control-Alt-Hack through an LLC, which sublicensed Control-Alt-Hack from the University of Washington. (All conflict of interest policies were followed—see Acknowledgments.) We chose this route for a variety of reasons; however, we can offer suggestions for those who wish to coordinate sale through their institution. Some points to consider include: (a) how the university can contract with distributors (e.g., Amazon), and what the contract review process looks like; (b) who would coordinate the payment of taxes, including sales taxes, to the state; and (c) how the university is equipped to enter wholesale or similar arrangements.

3.2. Timeline

We began work on the game in late summer of 2011 and submitted our conference paper on the game in Spring 2013. The timeline below contains some highlights of events along that timeline. We present the timeline for two reasons: (1) to convey amount of time that *must* be allocated for some activities, such as printing and shipping; and (2) to give an idea of the amount of time that we took with other sections of the process. We particularly focus on the activities that depend upon interactions and iterations with other parties, since these aspects are most out of a researcher's control (in terms of timing).

September 2011: First contact with Steve Jackson Games regarding licensing the *Ninja Burger* game mechanics.

December 2011: First meeting with graphic designers. First round of playtesting with a partially completed deck. (4-5 playtesting rounds performed total. 4-5 other sessions of game "read-throughs.")

February 2012: First meeting with illustrator.

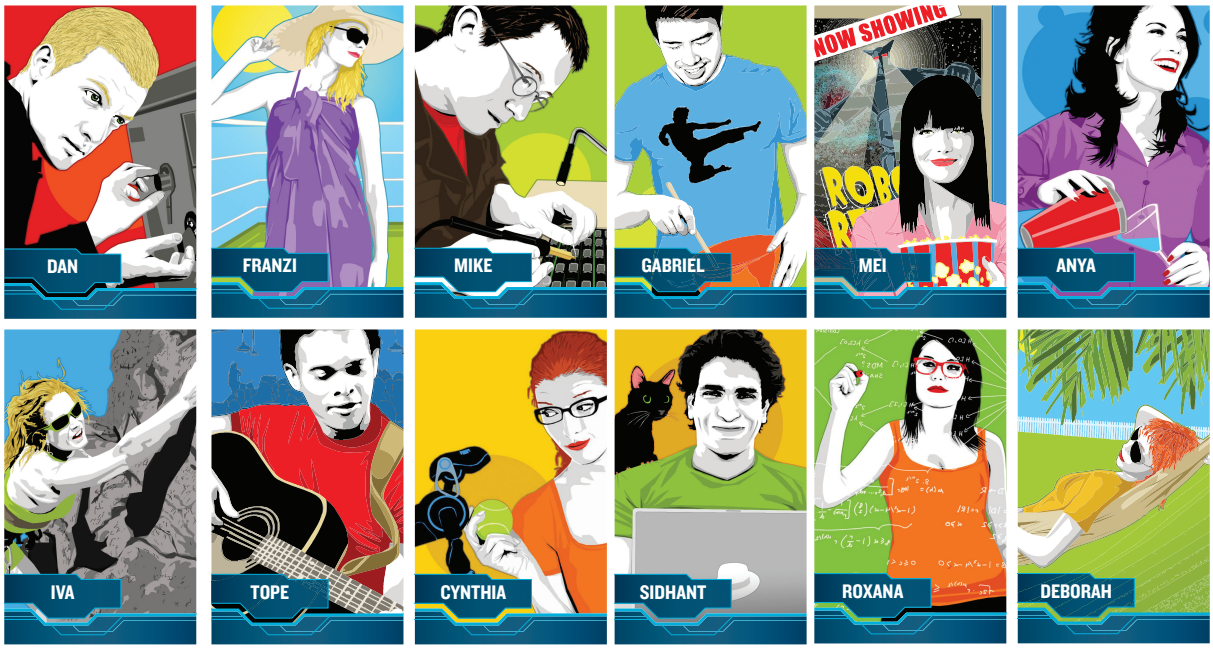

Figure 2. The character art from the portrait side of 12 of the game's 16 Hacker cards. © University of Washington

March 2012: First graphic design options reviewed. First game title (*Hackers, Inc.*) (of many) rejected due to potential trademark conflicts.

April 2012: Final game title chosen. License on mechanics signed.

June 2012: Final illustrations ready. (0-5 revisions done per illustration.)

July 2012: Print-ready files sent to printer. (Approximately 3 rounds of revision were done per card under the final design.) Talk on the game given at Black Hat USA.

August 2012: Printer prototypes (proofs) in and available for review.

September 2012: Games ship from printer.

October 2012: Trademark granted for *Control-Alt-Hack*.

November 2012: Games available for sale and start to be distributed to educators.

March 2013: Poster on *Control-Alt-Hack* at the SIGCSE conference.

April 2013: Surveys distributed to educators.

May 2013: Paper on educational evaluation submitted to CCS [1].

Two of our decisions particularly allowed us to compress our overall timeline beyond what we would have otherwise experienced. First, by licensing a game mechanic rather than developing one from scratch, we cut a great deal of time off of necessary playtesting, since we were only playtesting textual and visual choices. Second, our budget gave us the freedom to go with professional services rather than other alternatives; for example, while we could have sought to work with art and design students (for credit or for pay), that option would potentially have taken longer overall.

4. At a Higher Level: Research Challenges

The previous section focuses on logistics rather than the research process or goals. In this section, we depart from our particular experiences with producing Control-Alt-Hack; we switch contexts to reflect on a much higher level. We raise two of (we believe) the most prominent research challenges facing the usage of games in educational contexts.

4.1. Where Do Games Fit With Other Tools?

It would be helpful to educators for researchers to be able to suggest simple rules for how to best integrate non-lecture methodologies into a modern course of study. There are a great many educational approaches in use, and it is unlikely that simple rules can guide how to integrate security games into even a single pedagogic approach.

These models of non-lecture methodologies include a range of playful approaches from awareness-oriented games like *Control-Alt-Hack* to focused games such as Werewolves to multi-skill competition games like Capture-the-flag (e.g., [2][5][6]). The approach can even go as far as gamification of an entire course or degree program. Non-lecture models can also include creative writing (e.g., [3]) or blogging.

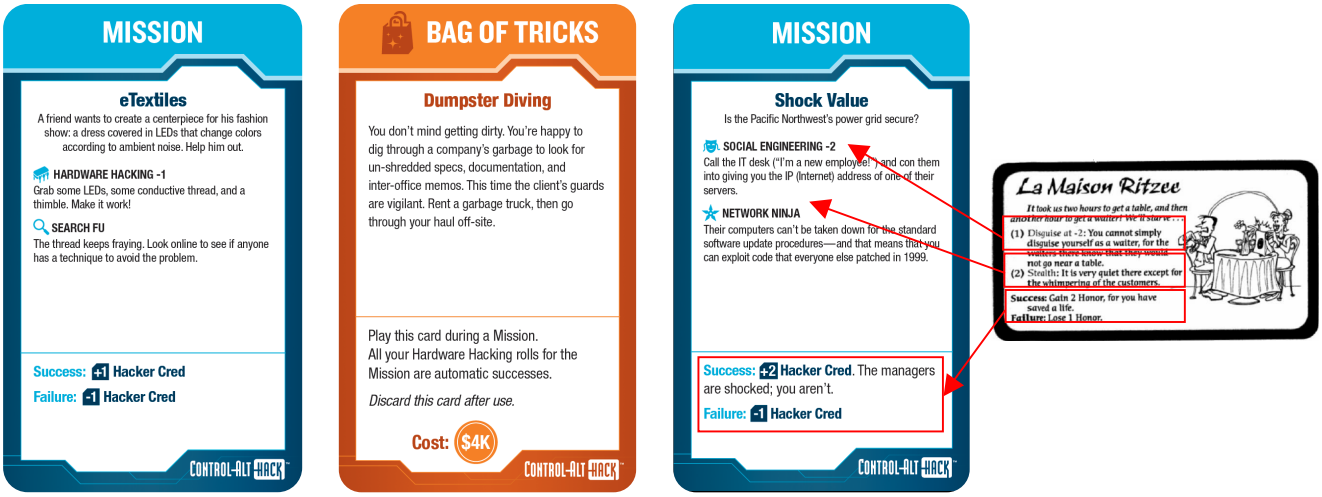

Figure 3. From left to right: (a) a Mission demonstrating the usage of technical skills for artistic purposes; (b) a Bag of Tricks card illustrating a particular attack threat; and (c) a Mission describing a social engineering attack on an SCADA system along with the mappings to the original Ninja Burger card. © University of Washington

Any use of games in an educational context is subject to risks. One risk is that the quality of games—and particularly computer games—has risen dramatically over the last 50 years. As a result, there is an exceptionally high bar for games today. This evolution was in full swing before the advent of app marketplaces and instant gratification, and the "red queen" effect is even faster since that epochal change. Many students are exposed to a wide variety of games from a very early age, and are skeptical of approaches that splash a layer of gamification on a normal lesson. Conversely, the increasing prevalence of game-building engines and a resurgence of simpler games or nostalgic, retro-style games on app stores and among indie games raises the hope that educational games do not need an entire production studio behind them to achieve success.

Another risk is that educators are supplied with (and include) too many games with similar mechanics in the same course. An increasing awareness of game mechanics, coupled with a vocabulary for their definition and discussion, means that students exposed repeatedly to the same mechanic are more likely to focus on the mechanisms, rather than the educational content.

An additional risk is that a course of study might become saturated with games, and that—instead of being a novel break from other instructional methods—the effort to learn the new games instead becomes tedious. (There is a variant of the first risk here; it is unlikely that all the educational games in a saturated course will be of the same quality.)

4.2. The Dilemma of the Evaluation, Or, The Madness In the Methods

Current (US) approaches to science and research favor uncovering cause and effect relationships, generally via reductionism and/or experiments that manipulate relevant variables while statistically controlling for all other potential influences. This approach undoubtedly has many benefits; however, over-reliance on this frame of inquiry can lead to a variety of ills, including dismissing other forms of inquiry, or steering away from the questions that most need investigating because they are much less convenient (or even infeasible) to measure via such methods.

If one were to attempt to compare the learning value of a game to a textbook for a certain set of learning goals—an obvious question to ask regarding the introduction of games into the classroom (or workplace)—the number of extraneous variables that would need to be controlled for is very large. Similarly, comparing a game such as *Elevation of Privilege* [4] to other threat modeling techniques presents complex challenges for evaluation that involves questions regarding prerequisite training, participant experience, or instructor bias.

One typical way to combat this problem is to conduct the evaluation at such a large scale that other potential causes begin to become statistically unlikely. Occasionally such a measurement can be attached to an existing mechanism—or collected from an existing data feed—in a way that provides a clear signal on a question of interest; when this happens, it is due to some combination of serendipitous circumstances and keen insight into experimental design. Frequently, however, such data can only be acquired via brute force and persistence. The volume of data that is required to determine "clean" results on a messy process as fraught with complications as human learning is somewhat prohibitive. This problem is exacerbated by the fact that the

publishing norms in computer science—the timescale and the expectations regarding the number of publications that result from a single project—tend to favor smaller, more agile projects; these norms can be contrasted with practices in fields that more traditionally deal with people, rather than technology, as their units of discovery. Moreover, it is unrealistic to expect a style of evaluation that requires such a commitment of time (and other resources) when the research mode is still so exploratory; the effort involved in the evaluation is not appropriately calibrated to expected outcomes.

Another looming problem in terms of measuring the effects of games—although one that is dependent upon the game's particular context and the goal in using it—is *how to measure success*. If the goal in using the game is to increase performance in a discrete task that is easily measured, then there is a clear avenue for measuring success. The more that goals steer away from this mode, however, the more challenging this proposition becomes.

For Control-Alt-Hack, we had the nebulous goal of increasing awareness of high-level security issues, with the idea of doing so via a recreational activity in which people might voluntarily engage. We could have attempted to create validated pre- and post-tests that measure security awareness. This would be an interesting—if daunting—approach, but one that would be susceptible to problems such as overfitting the evaluation metric to fit the characteristics of the game. Instead, we chose to take a more inductive (rather than deductive) approach. Instead of attempting to measure an answer to a pointed question, we used open-ended questions to invite educators to supply their own perspectives: on what worked and what didn't work (and why), on the things that the game is good for (versus not), and why they chose to use the game in their classroom. When examining the survey responses, it quickly became apparent that they could be characterized as either describing a positive function served by the game or critiques of the experience. The two clusters of positive functions volunteered in survey responses—Awareness and Social/Engagement—corresponded to our goals in creating the game (see [1] for more details).

5. Conclusion

In this paper, we address both the very low-level logistics of game creation and the very high-level issues of challenges facing games in education and research. We hope that sharing our experiences designing, producing, and distributing Control-Alt-Hack helps guide others who are interested in creating games for educational or research purposes. In addition, we discussed what we believe are the two greatest challenges in this area: properly situating games within the classroom curriculum, and evaluating—and publishing evaluations of—games with serious intentions. We welcome contributions and discussion in this area.

Acknowledgments

This work was supported in part by Intel Labs, an Intel PhD Fellowship, the ACM Special Interest Group on Computer Science Education, and the US National Science Foundation (NSF) under awards CNS-0846065 and CNS-1247216. Any opinions, findings, and conclusions or recommendations expressed in this article are the authors and do not necessarily reflect the views of the funding agencies.

Tamara Denning and Tadayoshi Kohno are founders and equity owners of RGB Hats, LLC, a private, for-profit company which has licensed the subject technology from the University of Washington. This research is subject to the conditions of a financial conflict of interest management plan established by the University of Washington.

We thank everyone who playtested, provided feedback, or otherwise helped improve the game. We thank those who graciously lent us their names and biographic details: Deborah Alterman, Mike Clarke, Alexei Czeskis, Karla Danson, Iva Dermendjieva, Miro Enev, Roxana Geambasu, Sidhant Gupta, Dan Halperin, Melody Kadenko, Karl Koscher, Gabe Maganis, Cynthia Matuszek, Temitope Oluwafemi, Franzi Roesner. We thank Steve Jackson for licensing his game and providing feedback, and Deborah Alterman and Mike Clarke at the University of Washington's Center for Commercialization for helping obtain the license. We also thank and acknowledge the following people: Gravity Creative (Graphic Design), Rob Kelly (Character Illustrations), Web Design / Development (Thomas Winegarden), Juliet Fiss (Product Photos), and Tina Wegner (Production Manager).

References

[1] T. Denning, A. Lerner, A. Shostack, and T. Kohno. Control-Alt-Hack: The Design and Evaluation of a Card Game for Computer Security Awareness and Education. CCS 2013.

[2] DEF CON. DEF CON Capture the Flag. https://www.defcon.org/html/links/dc-ctf.html.

[3] T. Kohno and B.D. Johnson. Science Fiction Prototyping and Security Education: Cultivating Contextual and Societal Thinking in Computer Security Education and Beyond. SIGCSE 2011.

[4] Microsoft. Elevation of Privilege. http://www.microsoft.com/security/sdl/adopt/eop.aspx.

[5] National Collegiate Cyber Defense Competition. http://www.nationalccdc.org/.

[6] PlaidCTF. http://play.plaidctf.com.

[7] Steve Jackson Games. http://www.sjgames.com.