IEEE S&P’21 Program Committee Statement Regarding The “Hypocrite Commits” Paper

May 6, 2021

The paper entitled “Open Source Insecurity: Stealthily Introducing Vulnerabilities via Hypocrite Commits” by Qiushi Wu and Kangjie Lu from the University of Minnesota was accepted to and later withdrawn by the authors from the 42nd IEEE Symposium on Security and Privacy (IEEE S&P’21). In this statement, the Program Committee (PC) of IEEE S&P’21 would like to comment on both the technical and ethical aspects of this paper, and reflect on how the process can be improved to prevent or mitigate some of the concerns this submission raised. We also outline the planned changes for future S&P conferences to address more generally the ethical concerns in computer security research.

The paper presents a method for introducing vulnerabilities into open-source software projects through minor patches (called *hypocrite commits*). The paper considers a restrictive scenario where the minor patch does not change or add any functionality, and can only change a few lines of code. The main idea is that a minor patch typically fixes a minor issue, but it can also implicitly introduce a new, hard-to-detect bug. The authors explore the reasons why such commits might be accepted and suggest stealthy ways to generate hypocrite commits. The authors claim that it is possible to defeat code review due to wrapper interfaces and complexity that hide malicious side effects. In addition, the authors evaluate existing bugs to qualitatively model the “stealthiness” of hypocrite commits and develop several such commits for the Linux kernel. Several such patches were then also sent to the Linux kernel developers for review, but none of the buggy patches were accepted for inclusion in the Linux kernel. The authors provide recommendations on how to better detect and prevent such buggy patches.

The paper was reviewed by four reviewers in the Fall S&P 2021 review cycle and received a very positive overall rating (2 Accept and 2 Weak Accept scores, putting it in the top 5% of submitted papers). The reviewers noted that the fact that a malicious actor can attempt to intentionally add a vulnerability to an open source project is not new, but also acknowledged that the authors provide several new insights by describing why this might be easier than expected, and why it might be difficult for maintainers to detect the problem. One of the PC members briefly mentioned a possible ethical concern in their review, but that comment was not significantly discussed any further at the time; we acknowledge that we missed it.

Because IEEE S&P has multiple submission deadlines and due to the ongoing COVID-19 pandemic, the PC did not hold a physical PC meeting to discuss each paper; instead, all paper discussions occurred online. Based on the overall positive reviews and the recommendation of all reviewers to accept the work, the PC did not discuss this paper during the online PC meeting; unfortunately, this meant forgoing a chance of catching some of the concerns that were eventually brought up.

When, after acceptance, the authors tweeted the abstract of the work in November 2020, several people expressed concerns about human-subject research featured in this work. At that time, the PC chairs discussed these concerns with the reviewers via the online submission system and in a Zoom call. As a result of these discussions, the PC chairs asked the authors to clarify the experiments with the University of Minnesota Institutional Review Board (IRB). We now acknowledge that this offer was a mistake: IRB determination of human-subject research should always be obtained before conducting an experimental study and submitting a paper to a conference. The authors submitted a human-subjects research determination form to their university IRB, which determined that this study was not human-subjects research, and therefore not subject to IRB jurisdiction. Based on this

IRB determination, the reviewers and PC Chairs discussed the paper again and decided to retain the paper in the program.

The IEEE Symposium on Security and Privacy is an international conference with participation by researchers from academia, industry, and NGOs. As such, determinations from IRBs, which are mandatory for human subjects research at US-based academic institutions, are not always locally required or even available to authors. IRB determinations are also not always sufficient to establish that a paper is ethical. Many potential ethical concerns in computer security research (such as responsible disclosure of vulnerabilities) may be completely unfamiliar to IRBs.

In April 2021, an unrelated project by the same group from the University of Minnesota fostered additional public discussions about this paper, which in turn led to a detailed investigation of this work involving the entire PC. There were more than 170 recorded interactions between PC members at that time. A thorough review of the IRB documents revealed potential problems in the description of the experiments, and concluded that insufficient details about the experimental study were provided to the IRB. PC members discussed at length the ethical dimensions of this work, particularly the deceptive nature of the study, the lack of debriefing of the study participants involved, and the risk of inadvertently introducing bugs into critical open-source code. Because the research was conducted at a university in the US, the standard for ethical computer science research is the Menlo Report (2012), which adapts the earlier Belmont Report (1979) for "information and communication technology research." Several principles outlined in this report were brought up, and after extensive discussion the PC determined that this paper does not follow these well-established guidelines.

In particular, the PC noted potential problems in two areas:

- *Informed Consent/Autonomy.* The Menlo report says that "Research involving information and communication technology (ICT) also raises the potential for harms to secondary stakeholders who, while not the direct subjects of research, may also have the right to autonomy. When considering informed consent, we suggest researchers and research ethics boards (REBs) carefully explore the complex interconnected relationships between users and the myriad of organizations which provide ICT services." Whether this research constituted direct human-subjects research remains subject to considerable debate among the PC, but there is no doubt that the autonomy of secondary stakeholders was violated.

- *Beneficence / Balancing risks and benefits.* The Menlo report requires "appropriately balancing probable harm and likelihood of enhanced welfare . . . diligent analysis of how harms are minimized and benefits are maximized . . . and implementing these evaluations into the research methodology." After extensive discussion, the PC believes that there were alternative research methods (for example, a controlled experiment on a simulated open-source project) that would have produced equivalent or better scientific value with much less potential for harm.

As part of the investigation, the PC also reviewed the paper's technical claims in detail to understand if the submitted patches were committed to the Linux kernel. Reviewers typically assess the quality of a submission by carefully studying the claims in the paper along with related work. Analyzing source code may happen, but it is not a requirement of the academic review process. The authors fully disclosed the technical aspects of their experiments in separate documents (see https://www-users.cs.umn.edu/~kjlu/papers/full-disclosure.pdf and https://www-users.cs.umn.edu/~kjlu/papers/hypocrite-patches.pdf), in which they discuss each of the five patches submitted by two anonymous Gmail addresses. One of these patches is not related to the paper, but was submitted using one of the email addresses by accident. Another patch is in fact a valid patch and it was accepted in the Linux kernel. Three other patches are buggy patches, but they were not accepted by the Linux developers for various reasons, not necessarily related to the bugs themselves. Investigation of these patches revealed that the description provided by the authors in the paper is ambiguous and in some cases misleading. The experiments do not provide convincing evidence to substantiate the main claims in the paper and the technical contributions of the work need to be revisited.

The PC discussed at length the possibility of retracting this paper from the IEEE S&P 2021 proceedings, based on both these ethical and technical concerns. The PC chairs were planning a program committee vote to decide if the paper should be retracted. Before this vote could happen, the authors independently decided to withdraw their paper (see their statement at https://www-users.cs.umn.edu/~kjlu/papers/withdrawal-letter.pdf).

The Linux Foundation's Technical Advisory Board also published a report on this work that is available at https://lwn.net/ml/linux-kernel/202105051005.49BFABCE@keescook/.

This paper brought up many potential ethical concerns that might occur in computer security research. To improve the reviewing process in the future, the following changes will be implemented for IEEE S&P 2022:

- Future editions of the IEEE Symposium on Security and Privacy will introduce an ethics review process similar to those recently introduced at conferences such as the International Conference on Learning Representations (ICLR). The PC will form a standing Ethics Review Committee with members who have a background in human-subjects research and related fields. This committee will review all papers that have been deemed to involve significant ethical concerns by either the authors or the reviewers (see below for details). Papers that have received IRB or other ethics board review are not exempt from this process, as IRB approval is not always sufficient to establish all the ethical implications of computer-security research. The Ethics Review Committee may reject papers regardless of their technical merit if they do not follow well-established ethical principles.

- In the submission form, we will add an ethics checkbox so that authors can explicitly signal that their research involves human subjects. In addition, we will add an ethics checkbox to the review form: In the event that any reviewer has concerns about the ethical aspects of the submitted work, this checkbox can be used to explicitly signal those concerns. All marked papers will be reviewed by the Ethics Review Committee regardless of their technical merit.

- The ethical section of the Call for Papers (CfP) for IEEE S&P 2022 will be strengthened and expanded to better reflect the expectations of the community regarding ethical considerations and human-subjects research. The PC will evaluate the ethical considerations of all submitted papers.

In the longer term, we will discuss with the IEEE S&P Steering Committee the possibility of developing a Code of Ethics similar to the ICLR Code of Ethics (see https://iclr.cc/public/CodeOfEthics). As the flagship computer security conference, IEEE S&P is committed to develop a process for evaluating ethical considerations and coordinate with other computer security conferences to establish common ethical guidelines.

Thorsten Holz and Alina Oprea
IEEE S&P 2021 Program Chairs
On behalf of the S&P'21 Program Committee