"It was jerks on the Internet being jerks on the Internet": Understanding Zoombombing Through the Eyes of Its Victims

CHEN LING, Indiana University Bloomington, USA
GIANLUCA STRINGHINI, Boston University, USA

Zoombombing, a kind of attack in which uninvited people intrude into online meeting rooms and harass meeting participants, emerged since the lockdown of COVID-19, people rely on online meeting tools to keep functioning. To understand how Zoombombing attacks unfold and their impact on participants, we performed 15 semi-structured interviews with victims of Zoombombing, who were either hosts or participants of the targeted online meetings. Through our interviews, we find that Zoombombing attacks can be distressful for attendees and that it is difficult for hosts to effectively react to them due to multiple factors, including the difficulty in identifying the intruders in crowded public meetings and the confusion that arises in attendees when a meeting is terminated and they are asked to join a new, more secure one. We find that meeting information leaks because of the vulnerability of the password, either by posting publicly, or an easy guess. We also find that the hosts of online meetings prioritize accessibility for their attendees to encourage participation and neglect the risks of lacking security measures, which facilitates attacks. This paper provides a comprehensive overview of Zoombombing attacks and the challenges in mitigating them, offering a blueprint for the research community to further investigate this problem.

ACM Reference Format:
Chen Ling and Gianluca Stringhini. 2024. "It was jerks on the Internet being jerks on the Internet": Understanding Zoombombing Through the Eyes of Its Victims. In *Proceedings of The 2024 European Symposium on Usable Security (EuroUSEC 2024), Sep 30th – Oct 1st, 2024, Karlstad, Sweden.* ACM, New York, NY, USA, 24 pages. https://doi.org/10.1145/3688459.3688466

1 INTRODUCTION

In the post COVID-19 pandemic era, setting up online meetings remains a normal for social activities. In fact, "Zoom" is now a used as a verb. According to the report from Zoom [11], Zoom had 219,700 enterprise customers as of October 2023, four times over the number 54,600 as of January 2020. The Zoom mobile app has been downloaded 81.48 million times in the first 6 months of 2023. *Zoombombing*, first spikes with the migration of meetings to online, as a consequence of the COVID-19 pandemic [6, 15], defined as uninvited people joining online meetings to disrupt them, often by posting obscene language and content. Previous research highlighted that Zoombombing could be a coordinated attack, initiated by eligible participants of the meeting targeting online lectures, which are the bored students looking for a thrill [26].

Common security advice to prevent Zoombombing includes requiring a password to join online meetings [14], adding a waiting room to manually vet attendees before letting them in, and not sharing meeting links publicly [52]. Because

Authors' Contact Information: Chen Ling, ccling@iu.edu, Indiana University Bloomington, Bloomington, Indiana, USA; Gianluca Stringhini, gian@bu.edu, Boston University, Boston, MA, USA.

Manuscript submitted to ACM

of that, they can share the password with other attackers, and even instruct them to adopt the names of legitimate participants to the meeting to confuse the host and avoid vetting. To mitigate this, Ling et al. recommend hosts take advantage of features like requiring registration to join meetings and sending attendees a personalized link so that it cannot be shared [27]. Other scholars however point out that there is a tension between tightening the security of online meetings and the need for them to be easily accessible to as many people as possible, especially those with low technical expertise [6].

Despite the emerging threat posed by Zoombombing, the research community still lacks a good understanding of what happens during attacks, what harm is caused to participants, and what issues faced by meeting hosts when trying to remediate the damage caused by attackers. Existing research on the subject either relied on analyzing public posts on social media and anonymous message boards calling for Zoombombing attacks [26] or analyzing news coverage about Zoombombing attacks [12, 15]. These approaches either lack visibility on what happens during Zoombombing attacks and on the experience of victims or do not have the nuance that can be provided by rigorous qualitative studies.

To fill this gap, we recruit victims of Zoombombing and conduct the first interview study on the attack. Our goal is to understand the following research questions:

(1) How do attackers disrupt online meetings, and what methods do they use?
(2) What types of harm do Zoombombing attacks cause to participants?
(3) What usability challenges are associated with current security measures for online meetings?

In early 2022 we recruit 15 Zoombombing victims, including participants with different familiarity with technology, age groups, and industries. This allows us to get a multi-faceted understanding of the Zoombombing attack and of the challenges in effectively mitigating the problem. Some of the key themes that emerge from our interviews:

- Meeting hosts commonly prioritize the convenience of their attendees to adopting strict security measures, considering the likelihood of being attacked low and therefore not worth the burden.
- Meetings are usually targeted after attackers find a link to them that is posted publicly, for example on social media. This indicates that Zoombombing is largely an opportunistic attack and is in contrast to previous research that found that early in the pandemic attacks were often called by insiders who wanted others to disrupt their meetings [26].
- Once the attack starts, it is difficult for hosts to remediate. Depending on how the meeting is set up, attackers can come back, or even take over the meeting. Attackers also impersonate the host or other participants to cause confusion and make identifying and kicking them out harder. Often hosts have no choice but to start a completely new meeting, but this introduces usability challenges because many attendees get confused and are unable to join the new meeting, increasing the disruption caused by the attack.

Our study introduces novel contributions to the existing body of research in online harassment. Firstly, this paper is the first to examine Zoombombing attacks from the victims' perspective, providing unique insights into the experiences and impacts on those directly affected by Zoombombing. The study offers a victim's perspective of the process of Zoombombing, including how attackers gain access, the nature of their disruptions, and the subsequent challenges faced by meeting hosts. It highlights the tension between tightening security measures and maintaining ease of access for participants, particularly those with low technical expertise. In the meantime, this paper explore the personal and professional impacts of Zoombombing on victims, examining how these attacks affect their cybersecurity awareness and practices. Finally, based on the findings, we discuss practical recommendations for improving the security of online

meetings, balancing the need for protection with usability for all participants, and paving the way for more targeted and effective interventions.

2 RELATED WORK

Online abuse carried out on social networks has been a common topic of study by the security community, spanning from fake accounts carrying out spam campaigns [4, 9, 16, 19, 40, 47, 49], astroturfing campaigns [10, 41, 46], and online disinformation operations [3, 28, 32, 38, 39, 50].

A wealth of work has recently focused on understanding various aspects of coordinated online harassment [43]. Hine et al. [20] study calls for attacks made by anonymous users on 4chan's Politically Incorrect Board (/pol/), showing that attackers often pick YouTube videos as targets and that the synchronization patterns between comments on 4chan and YouTube can be used to identify attacks. Mariconti et al. [29] use this observation to identify a set of YouTube videos that received harassment attacks, and develop a classifier to determine the likelihood that a newly uploaded video will receive hateful attacks. Zannettou et al. [51] leverage a similar observation to study the occurrence of hateful comments in the comment section of news articles, posted after links to these articles are posted on polarized online communities. Kumar et al. [22] focus on *brigading* on Reddit, where members of a subreddit coordinate to attack another subreddit, often by posting toxic and hateful messages and disrupting its operation. Snyder et al. [37] study *doxing*, where attackers post personal information about a victim with the goal of encouraging harassment against them on several online platforms or even in real life. Aliapoulios et al. [2] develop a technique to automatically identify calls for doxing on multiple platforms and study their occurrence.

Aside from discussing how online harassment is coordinated to help us better understand Zoombombing, several works focus on its purpose and implications. Friedberg et al. [15] analyze the Zoombombing attack from a socio-technical point of view. Elmer et al. [12] study the media frames that discussed Zoombombing incidents early in the pandemic, providing an overview of how the attack emerged. Ling et al. [27] perform a mixed-method analysis of calls for Zoombombing on 4chan and Twitter, finding that attacks are often called by insiders (e.g., students in classes) who are inviting others to disrupt their meetings.

Other work focused on the lack of security measures deployed by online meeting platforms to prevent Zoombombing. Secara. [33] argues that online platforms are not taking security seriously enough, prioritizing the introduction of new functionalities instead of focusing on keeping users safe. Young et al. [48] examine privacy issues on online meeting software, arguing that these platforms should incorporate privacy by design and be transparent about their privacy practices to address the growing public interest in, and dissatisfaction with, corporate privacy practices. Siddiqui et al. [35] point out despite Zoom's advanced encryption, overlooked security features and insider threats leave users vulnerable, necessitating stricter settings and institutional punishments to reduce attacks.

Other research in this space found that similar to other coordinated harassment studies, Zoombombing attacks target race and gender when picking victims and carrying out their attacks. Lee. [24] analyzes news reports about Zoombombing incidents and observe that there is a prevalence of attacks that target people with Asian, African, and Jewish heritage. Ali. [1] also observes that racism is a major driver for Zoombombing. They argue Zoombombing as a new form of alt-right shitposting needs to be critiqued and understood as more than simply trolling. Walsh et al. [44] document how they reacted to a Zoombombing incident that targeted a meeting with 70 staff members and students, organizing a rapid response team with experts in trauma-informed psychology to address any trauma suffered by victims. Lee et al. [25] uses a grounded theory-based qualitative analysis of over 1,000 Reddit posts to explore

how Zoombombing victimization subreddits function as communities for sharing stories, perceptions, and support, highlighting the role of online forums and cyberpolicing tools in preventing and discussing victimization.

To the best of our knowledge, our paper is the first interview study to understand the Zoombombing attack. By interviewing participants across industries, and technical expertise, we build a multi-faceted view of Zoombombing incidents, which we hope can serve as a blueprint for further research conducted by the security research community in this space, especially informing better mitigations.

3 METHODS

We conducted semi-structured interviews during the first months of 2022 with victims of Zoombombing attacks to gain deeper insights into victims' views, behaviors, and attitudes towards such attack.

In this section, we first introduce how we recruited participants to our study. We then present our interview protocol and the ethical considerations that guided our work.

3.1 Participant Recruitment

The goal of this study is to gain an in-depth understanding of the Zoombombing attack from a victim's perspective. We sought self-claimed victims who met the eligibility criteria for the interview, which required them to have experienced at least one instance of Zoombombing, defined as an unwanted, disruptive intrusion into a video-conference call [1].

From January to April 2022, we recruited participants by posting a recruitment flier on social media platforms (Facebook, Twitter, Reddit, LinkedIn, Instagram.) Additionally, we reached out to individuals who publicly shared their Zoombombing experiences on social media. We recruited and interviewed participants on a rolling basis within this same timeframe.

We managed to assemble a diverse set of participants ($N = 15$) based on data saturation, i.e., when no new themes emerged from new data [11, 36]. Table 3 (in the Appendix) provides details on the socio-demographic profile of participating Zoombombing victims. Note that two of our participants experienced Zoombombing during the same online meeting (one as a host and one as an attendee). Therefore, while we have $N = 15$ participants, our meeting-wide analysis only covers $N = 14$ online meetings.

3.2 Interview protocol

We conducted virtual interviews ($N = 15$) through Zoom from January to April 2022. We sent potential interested participants a consent form before they confirm their participation. An interview session was schedule for 60 minutes. The actual interview lasted forty minutes on average. Before starting, the interviewer went through the study details and consent form with the participants which were recorded with the participants' verbal consent. Participants were given the option to disable their video if they wished. Recordings were transcribed using automated means provided by Zoom, and the transcripts were checked by the interviewer against the original audio for accuracy. we reserved time at the end of each interview to allow the participant to ask us questions about our study. We did not compensate participants for the study.

During our interviews, we asked questions about five themes, designed to cover multiple aspects of the Zoombombing experience. To minimize biases, we used a neutral tone, refrained from expressing opinions, and allowed sufficient

[1] https://en.wikipedia.org/wiki/Zoombombing

time for interviewees to remember and articulate their answers. We prepared prompts like "What else do you want to mention about?" to encourage participants to comment or ask questions at any time.

In the following section, we describe five themes of our interviews in detail, specific interview questions please see Appendix A.

1. User habits and Zoombombing experience. To understand participants' familiarity with online meeting tools, their experience with multiple Zoombombing incidents, and their roles during these attacks, we include the questions like how often they use online meeting tools, the number of Zoombombing incidents they have encountered, and their specific roles (e.g., host, participant) during these incidents.

2. Meeting settings. To explore how the meetings were organized and conducted, we involve questions identifying the purpose of the meeting (e.g., business, academic, social), the expected audience size, whether it was a recurring meeting, and any security measures the host implemented to prevent attacks.

3. Delivery of the attack. To analyze the actions and behavior of attackers during the Zoombombing incidents, we include questions examining the specific actions attackers took, whether they acted alone or as part of a group, and if the attack seemed targeted at specific participants. It also seeks to understand the nature and dynamics of the attack.

4. Mitigation of the attack. To assess the effectiveness of the host's response to the Zoombombing incident, we ask questions to look at the actions taken by the host to stop the attack, the effectiveness of these actions in resuming the meeting, and whether attackers responded to the host's mitigations. This theme aims to identify effective strategies for preventing and stopping Zoombombing.

5. Attack aftermath. To evaluate the impact of the Zoombombing experience on participants, we include questions assessing the emotional impact on participants, changes in their attitudes towards online meeting tools, and whether they have become more cautious in setting up and participating in online meetings.

3.3 Analysis of the Interview Recordings

After recording all interviews, we follow a thematic coding process [5], similar to those used in previous interview studies in top computer security venues [17].

First, the two authors of this paper, acting as independent annotators, watched the recordings to familiarize themselves with the content. One of the annotators, who also conducted the interviews, then developed a set of initial codes to characterize the responses to the questions listed in Section 3. The two annotators discussed these codes until they reached an agreement, using the threat model proposed by previous work [26] as a framework to understand the nature and impact of the attacks. They also considered user behavior vulnerabilities and different meeting settings to develop more detailed codes.

Once agreed on the codebook, the annotators independently coded the interviews, pausing periodically to discuss and refine the codes. They also noted significant statements from participants. To assess inter-rater reliability, they calculated Cohen's Kappa score [23] for each label. The results showed almost perfect agreement for 36 out of 39 labels ($\kappa > 0.8$), and substantial agreement for the remaining three labels: how the host dealt with the attack ($\kappa = 0.76$), participants' attitudes towards online meeting tools ($\kappa = 0.73$), and their preferences for meeting formats ($\kappa = 0.61$).

These findings indicate that the annotators reached a high level of consensus on most codes, while also highlighting the diverse responses of participants. These nuances are further explored in the paper through direct quotes from the participants.

3.4 Author Positionality

Given the sensitivity of the subject and the fact that previous research has highlighted that Zoombombing attacks might disproportionately target marginalized groups [26], it is important to discuss how the authors approach this research. We have been working on developing automated techniques to defend against different types of cyber threats, therefore our main goal is to better understand the *modus operandi* of attackers and the shortcomings of current mitigation tools made available by meeting platforms. Half of the authors identify as women and half identify as men, with half of them having Asian heritage and half being Caucasian. All authors identify as heterosexual. Some of the authors have experienced coordinated online harassment in the past, although this did not happen on Zoom or other meeting tools.

The diverse backgrounds and personal experiences of the authors have influence on the overall analysis. For instance, the gender diversity within the team allows for a nuanced exploration of how Zoombombing attacks might differentially impact users based on gender, particularly since women and minorities are often more targeted in online harassment [30]. Similarly, the cultural diversity provides insights into how these attacks might be perceived and experienced differently across racial and ethnic groups, ensuring that the research does not overlook the unique vulnerabilities faced by different participants.

Furthermore, the personal experiences of some authors with coordinated online harassment contribute to a deeper understanding of the psychological and emotional impacts of such attacks. This firsthand knowledge informs the development of more empathetic and user-centered defense mechanisms. By integrating these varied perspectives, the research aims to create more robust and inclusive mitigation strategies that address the specific needs of marginalized groups who are disproportionately affected by cyber threats.

3.5 Ethical Considerations

Interviewing victims of harm calls for careful ethical consideration, as bringing up potentially traumatizing memories could cause further harm. To mitigate this, when conducting interviews we emphasize to the participants that it is OK to take time off or skip any questions that may make them feel uncomfortable We allow the participants to withdraw consent at any time. Additionally, since our interviews are audio and/or video recorded, there is a risk of loss of confidentiality. To mitigate this, we let participants choose whether they want to share their cameras or not during the interviews. All recordings stay anonymous and are stored encrypted on a server at our institution and are only accessible to the authors of this study. We ensure that the use of the recordings is only on this study which is agreed upon by the participants. We delete all information after the publication of the study to avoid leaking the recordings or being used in a way that could harm the interviewees or misrepresent their views. To further protect participants, we do not disclose any personal information in this paper and paraphrase their quotes. Our study was reviewed and approved by the IRB at our institution.

4 RESULTS

In this section, we present an analysis of the results that emerge from our interviews. We organize our analysis following the five themes discussed in Section 3.2, focusing on the user experience of our participants regarding online meetings and Zoombombing incidents, the precautions that were taken when setting up these meetings, what happened during the attack, and the aftermath of Zoombombing incidents, including the harm caused to participants.

In the rest of this paper, we report paraphrased quotes attributing them to participants (P1-15, $N = 15$) if the quote is about a participant's experience or opinion. On the other hand, if the quote is about something that happened during the zoombombed meeting we attribute it to the meeting (M1-14, $N = 14$).

4.1 User habits and Zoombombing experience

To better put our participants' experience in context, we first need to understand their general experience with online meeting tools. We find that twelve of our participants use these tools at least once a day. Three of the participants attend or hold virtual meetings on a weekly basis.

When looking at the popularity of online meeting tools, Zoom is the most often used platform, with 13/15 of our participants primarily using it. One participant mostly uses Cisco Webex while another participant uses Google Meet, both in a corporate setting. These results support earlier research, which found that Zoombombing attacks typically occur on Zoom likely due to its large user base [26].

The majority of the participants (12/15) only experienced one Zoombombing incident as a result of their use of online meeting tools, while three of the participants have been through the attack twice or more.

Takeaways. Given that most participants use online meeting tools frequently, Zoombombing attacks are not prevalent but rather sporadically encountered by frequent users.

4.2 Meeting settings

Previous research on Zoombombing recommended that online meeting hosts adopt best practices to reduce the chances that their meetings will be disrupted [26, 33]. As part of this theme, we aim to understand whether hosts adopt measures to keep their meetings secure.

Table 2 (in the Appendix) reports the precautions that were set up by the hosts of the meetings that were attacked and that were witnessed by our participants ($N = 14$ meetings). Not every participant remembers the technical settings of the meeting, particularly when they are the attendees rather than the host. For six meetings, the participant could not recall if any security measures were taken. Only one participant reports that the meeting they joined required them to register in advance. Three meetings utilized a password to secure the meeting room, and one meeting was required to log in with a Zoom account to join. For one meeting, the participant reported that the host set attendees to be muted by default, with the camera turned off.

Through our discussions with participants, it became clear that hosts often prioritize ease of access for their attendees over security measures. Security is frequently seen as an afterthought, perceived as a potential barrier for participants, particularly those who are less technically savvy.

In four cases, participants believed that precautions were unnecessary because they had successfully organized similar gatherings multiple times, establishing trust and forming a community with their attendees. Most meetings were held among acquaintances, with attendees either being members of the same group or familiar with each other due to the regularity of the activities, such as a weekly research seminar.

The host of one meeting (M4) explained the rationale for minimal security precautions:

> *"Members are middle-aged people, so they were trusted - invites were only sent via email... Only members were invited to meetings, so there was no need for security." (M4)*

Similarly, the host of another meeting (M3), who gave an educational talk, noted that the meetings were open to all students and advertised only through internal channels. She mentioned that seeing unfamiliar names among attendees would not raise suspicion, as new Ph.D. students often joined these meetings:

> *"The meeting wasn't password protected because it was open to all students." (M3)*

Furthermore, the host of M4 clarified why he does not mandate registration for large gatherings of 180 people:

> *"By requiring registration, people need to deal with this ahead of the meeting. It took the speaker 10 minutes to figure out how to get in." (M4)*

We asked our participants if they had any idea how the attackers obtained the meeting details (e.g., the Zoom URL) for their meetings. In eight cases, participants said that the URL for the meeting was posted publicly, either on social media or an institutional website. Often, this link would also include the meeting password, allowing attackers to bypass this protection and enter the meeting.

> *"In the morning a tweet went out about the seminar, including the link (with the password)." (P6)*

In one case, a participant mentioned that hacking might have been involved since links to the meeting were only shared with trusted parties via email:

> *"One of the panelists thought that someone had hacked into someone's email to get the link." (P5)*

For some online meeting tools (e.g., Webex), meeting rooms can have a predictable URL, which attackers can use to join certain meetings and either disrupt or eavesdrop on conversations.

> *"With Webex, you don't need to know the link; you only need the company's domain name and the person's first and last name (which are publicly available)." (P7)*

Takeaways. Participants prioritize the ease of access to their online meetings over implementing security measures. This creates a tension between maintaining accessibility and ensuring security, suggesting that hosts may underestimate the risks or feel overly confident in their established community's trustworthiness. Our interviews reveal that attackers often gain access to meetings because the details, including passwords, are posted publicly online. Additionally, attackers exploit the predictability of meeting links for certain tools, making it easy for them to join without specific invites.

4.3 Delivery of the attack

Previous work on Zoombombing focused on public calls for attacks made by attackers [26], which did not capture what actually happens during these attacks. In our study, we investigate how attacks are carried out, identify the attackers, and examine how hosts react to these incidents. Details of the attacks are provided in Table 4 (in the Appendix).

Timing and Duration of Attacks. In all Zoombombing attacks observed by our participants, attackers launched the attack at the start of the sessions or within a few minutes. As one attendee (M3) noted:

> *"Since the beginning people were saying weird things." (M3)*

Most attacks were short-lived, with twelve attacks lasting for less than five minutes. Although the attacks only lasted a few minutes, one host mentioned that they were so distressed that it felt longer:

> *"I obviously felt like 10 minutes or whatever, but I think it was only like two minutes or so." (P6)*

Multi-modal attacks. Attackers used various methods to disrupt meetings, including:

Attack method	Description & Quotes
Sound Attacks	Yelling, shouting, or playing profane sounds were the most common, occurring in almost all (13/14) meetings. *"At the beginning, the attackers were eating chips loudly... (then) one attacker showed a video and wasn't wearing a shirt." (M13)*
Video Attacks	Attackers often turned on their cameras or played videos to escalate the situation. In three meetings, they revealed their naked bodies or played pornographic videos. M12's host described attackers playing disturbing videos, including pornography and a beheading video.
Chat Harassment	In five meetings, attackers posted offensive content in the chat to harass attendees. *"There was in the chat... a white supremacist curse, there were anti-Semitic slurs." (M4)*
Profile Picture Changes	Two meetings reported attackers changing their profile pictures to disturbing images, including a Nazi symbol (M6) and child pornography (M2).
Screen Control	In three meetings, attackers took over the screen to display obscene imagery and used the annotate function to draw hateful graffiti on each attendee's screen. *"The speaker eventually got upset and stopped screen sharing, which was a mistake because the attackers then started sharing their screen, showing indecent videos and images." (M14)*
Impersonation	Impersonation is particular harmful to the victims as they did not realize what was happening and saw racist and hateful messages being sent by someone they trusted and respected. This exemplifies how such an attack can have a far wider impact, as the impersonated victims will have to explain to the other attendees what is going on. The host of M6 confirmed that attackers adopted the screen names of legitimate attendees to avoid being kicked out. An attendee from M8 reported that these impersonation attempts can cause problems for legitimate users, who might be mistaken for attackers and get removed by the host: *"The impersonated participant was kicked out after 15 minutes and couldn't join back." (M8)*

Re-entering and attack	If attackers have the password-protected link to join the meeting and that is the only mitigation put in place by the host, they can re-enter the meeting without host approval, even after being kicked out. As a result, in seven meetings, attackers kept re-entering the room, making the host's efforts to repel their attack ineffective. Due to these issues, hosts often have to shut down their meetings or even establish another meeting with a fresh link, risking the loss of guests. *"After the meeting was restarted, the attackers entered again and had cameras on and were wearing black hoods." (M5)*

Table 1. Multi-modal attacks reported by participants ($N = 15$) during their Zoombombing experience.

There is one testimony (M7) which reports a silent eavesdrop during a corporate conference, saying that this occurs frequently when the participant utilizes Webex for work:

"It's not always somebody harassing people, they stay silent to eavesdrop." (M7)

It highlights that there are other motives that drive people to join online meetings without authorization besides disrupting them.

Attack content. When looking at the type of hateful content that attackers post to disrupt online meetings, we see profanity, targeted race group, and sexist.

Twelve meetings reported obscene sounds (M1, M5, M8, M9) or pornographic videos being played (M1, M3, M4, M5, M6, M8, M9, M10, M11, M12, M13). To make these cases even worse, several hosts (M4, M5) report that child pornography was shown as an avatar, or as video taking over the screen.

Nine meetings are disrupted by assailants shouting racial slurs or displaying white supremacist images or videos.

"It definitely felt targeted at the Jewish community... one of the attackers played footage of Hitler." (M9)

An M8 attendee also shared his experience of being impersonated to post sexist messages when he was unaware of them at the time:

"Someone just started like sending some very nasty messages using my name... Sexist messages."(M8)

Meanwhile, the attacker continued to rename themselves as attendees and send abusive sexist messages in the conversation.

"But then when I was kicked out, the person renamed themselves as another attendee. They even sent messages... like all the vile stuff you can think about."(M8)

Description of the attackers. The majority of the attacks (11/14) were carried out by multiple attackers. M1's attendee recalled:

"A guy started it, ten people participated, they had profile pictures but not sure if it was them... Generic white high school guy and generic African American high school guy picture." (M1)

Multiple of our participants mentioned young males as the attackers, such as "Naked man," "Teenagers," (M6, M12) and "From the username it appeared that they were male." (M10)

Beyond these general characterizations, no participant was able to identify the actual identity of the perpetrators.

Target of the attack. For nine meetings, participants report that the host or the speaker of the meeting was specifically targeted during the attack. We observe race and gender were used opportunistically by the attackers as an excuse to harass participants.

> *"Panelists had African and East Asian ancestry... As soon as the black panelist started speaking, the zoombombers started playing pornography sounds." (M5)*

Even worse in this case, the attendee believes the attackers knew the African panelist because they knew personal information about her. The attackers were yelling racist slurs and whispering death threats (M5).

For the remaining five meetings participants could not pinpoint a specific motivation for the attack, other than it was done to disrupt it, and likely for "fun."

Takeaways. We find that attackers use all tools at their disposal to disrupt online meetings, including audio, text, video, and the annotation tool. Attackers also impersonate legitimate participants of the call to confuse hosts and make it harder to get identified and kicked out. Obscene content and even child abuse material are shared as part of these attacks. Hosts of the meetings and people from minorities group are more often to be targeted.

4.4 Mitigation of the attack.

As part of this theme, we want to understand how meeting hosts react to Zoombombing attacks, and whether they are able to continue their meeting.

Kicking the attackers out. All of the hosts of the zoombombed meetings try kicking out the attackers. However, as discussed previously, oftentimes attackers can just re-enter, making this countermeasure ineffective. Because of this, seven hosts simply terminated the meeting. This is not necessarily an effective solution either. As the host of M12 reports, they terminated the meeting, but since there were no security measures in place attackers could just restart the meeting themselves and harass any participants that joined late and were unaware of the attack.

Start another meeting. For six meetings, the hosts started another meeting after the original one was disrupted, but encountered the challenge of communicating the new meeting location to the participants.

> *"I expelled them all and we actually closed the meeting and shared a new link via email." (M13)*

Alternative mitigations. Other mitigations that our participants tried were muting the attacker (M11), disabling the chat (M10), or allowing only the presenter to speak (M10).

Issues with mitigating attacks. A common trend that emerged during our interviews is that hosts struggle to keep up with the attackers and remediate their nefarious actions. One participant mentions that they were not familiar enough with the meeting tools to be able to perform all the actions that were needed to undo the damage that had been done by the attackers:

> *"Overwhelming experience, trying to figure out how to delete the annotations, how to mute everybody, too many attackers etc." (P6)*

The same participant also mentions that in a large meeting where not all people are known it is very difficult to identify who is legitimate and who is not:

> *"Since we had attendees from all over the world, it would have been extremely challenging to identify who was an attacker and who wasn't." (P6)*

As a potential solution, one participant suggests that meeting tools could add a "panic button" that would allow any participant to stop the meeting.

> *"Zoom should have an emergency button that any participant could hit." (P3)*

Takeaways. We find that attackers carry out Zoombombing attacks through a number of means, from posting abusive content in the chat to shouting profanities and showing obscene content. We also find that participants can be targeted because of their race and gender. From our interviews, it emerges that reacting to Zoombombing attacks is difficult, not only because attackers are actively trying to confuse hosts by impersonating other attendees, but also because if the initial meeting settings are lax then attackers will be able to re-enter or even re-start the meeting themselves.

4.5 Attack aftermath

Previous studies mention the harm suffered by the victims of Zoombombing [26], but did not speak with Zoombombing victims directly. To bridge the gap, we investigate the effect of Zoombombing on the affected victims, the harm to them, and the impact on the platforms. Table 5 (see Appendix) presents these harms as reported by participants ($N = 15$).

Effect on the meetings. We first want to understand how attacks affected their targeted meetings. Even though over half of the attacks finish within 5 minutes, much harm is caused to both the hosts and the attendees during that time.

As we said, because of lax online meeting settings it can be difficult for hosts to keep attackers out of their sessions. After being kicked out, attackers can often re-enter, resuming the attack. Shutting down the meeting without creating a new one does not help either, because if attackers have the link they can re-enter or even re-start the meeting themselves, causing further confusion and disruption.

For this reason, hosts often have no choice but to end the meeting and start a new one (with a new meeting ID and a new link). However, this can confuse the attendees, and prevent the meeting from properly resuming. The host of M4 reports that everyone was calling him because they were trying to join the old meeting and could not get in:

> *"Set up a unique link for Zoom after the attack, but participants couldn't figure out how to join - at least 20 people gave up." (P4)*

In some cases, meetings have to be ended as soon as a breach is suspected due to company policy:

> *"If an attacker is spotted, the host has to stop the meeting (it's a corporate meeting)." (P7)*

Harm on the attendees. Our participants reacted differently to the Zoombombing experience, depending on several factors, including their sensibility, their role in the attacked meetings, and their previous exposure to Internet toxicity.

Not a big deal. P1, who is a mid-career security researcher, does not make a big deal of her Zoombombing experience. She was an attendee in a virtual research seminar where attackers joined and yelled profanities. She was aware of Zoombombing before, mostly having read about it from the news, but she would have never thought this would happen to her in real life. She stated:

> *"You know it was just jerks on the Internet being jerks on the Internet. At the time I was sort of incredulous and kind of amused." (P1)*

P1 however said that if the meeting had happened in more solemn circumstances, like during a memorial, then she would have taken it much more seriously and would be mindful of the potential consequences.

Other participants expressed similar attitudes. P3, a African American young engineer who was targeted during an educational talk that he held, holds an optimistic attitude on the attack.

> *"I didn't think it was like that huge of a deal for me... You know, I have a great job, I have a great life. They really didn't affect me that much." (P3)*

Unacceptable. On the other end, victim participants who were hosting the attacked meetings often reacted more negatively to the attacks. P2, a senior engineer who organized a weekly technology hobby group over Zoom was hurt deeply by the Zoombombing attack. He was impersonated by the attackers, who posted offensive content on his behalf, and he had to apologize and explain this to its attendees.

> *"I felt humiliated and embarrassed... As the host, this kind of disruption is unacceptable." (P2)*

P4 feels the same way, as he organizes a monthly gathering for a hobby group over zoom, inviting speakers from all over the world. He and his audience both come from non-technical backgrounds. Like many others, he prioritized ease of access to security. He never set up a waiting room because he said that elderly attendees are not tech savvy and they would often leave and join back several times, this would impair their ability to attend the meetings. He mentioned that the attack was particularly disorienting and upsetting for these elderly attendees:

> *"Because I've got a lot of old people that are really clueless... That was a huge mess eventually." (P4)*

Eight of the participants report some kind of PTSD from the Zoombombing experience, not only themselves, but also the other victims of the attack.

> *"The targeted panelist didn't look shocked, although she was irritated... She was surprised at how something digital (not physical) could cause her so much harm." (P5)*

Nearly half (7/15) of participants talked with others about their experience after the Zoombombing attack. Some hosts sent emails to the speakers inviting them to resume the seminar, while others sent apologies after the meeting ended.

Support and reports. Hosts often seek support either during the meeting or after the attack ends.

Since Zoombombing occurs in real-time, any help from an organization's IT Department happens after the fact, limiting its effectiveness. For four meetings, hosts reported the incident to their organization's IT Department. In two cases, the organization started a detailed investigation. The attendees of M3, which was an educational talk held by a university, received a follow-up email about the incident:

> *"The university thinks that attackers ran a script to find public Zoom links on the Internet." (M3)*

The attendee of M7 was not satisfied with the support from the IT Department of their organization:

> *"They locked the meeting after 10 minutes so that people couldn't join - enforced by IT... Information leaks are a serious concern, and this wastes a lot of time by IT." (M7)*

However, the remaining half of the reports did not receive any follow-up from IT as the host of M6 said:

> *"Sent a message to IT, didn't hear back." (M6)*

Four of the victims sought help from online platforms or law enforcement: the host of M2 reported the attackers' accounts to Zoom. According to his description, the platform seemed to monitor the meetings continuously. When reporting the incident to Zoom, they showed him screenshots of the meeting, which surprised him. The host of M4 reported the incident to Zoom because the attackers shared child pornography, but did not receive a response from Zoom. The victims of M5 reported the racist attack they experienced to the federal police. The RCMP (Royal Canadian Mounted Police) informed them that they lack the infrastructure to deal with this type of crime and that such speech is

likely not against the law. The victim of M7 reported silent attackers appearing in their work meetings. Webex advised them not to use the personal room, which has a predictable URL, but to create new meetings with randomized IDs.

Lessons learned from the attack. Being victimized by a Zoombombing attack encouraged all participants to think about security more carefully, and adopt more precautions in the future. Examples include generating a new meeting URL for each occurrence of a recurring seminar, or sending the meeting link only through private channels:

> *"I stopped using a recurring meeting for the seminar and started creating a new link for each occurrence." (P6)*

Regarding the balancing of security and usability, most of the participants admit security is necessary to sacrifice some convenience. Five participants mentioned that security should be kept in mind by hosts when setting up meetings, although this might make it more difficult for participants to join.

> *"Improved security of Zoom meetings, but this adds hurdles to participants. It's probably worth it to keep people safe though." (P2)*

On the other hand, three participants pointed out that the onus of security should be on the platforms. They have responsibility to provide better security to their users without jeopardizing the accessibility of their meetings:

> *"The burden shouldn't be placed on users but it should be on the platforms." (P3)*

Impact on the participant perception of online platforms. Although Zoombombing experiences were harmful to participants, most of them did not consider abandoning online meeting tools or switching to other platforms. Most participants have not heard of attacks happening on platforms other than Zoom, but they also did not think that other platforms are better or more secure. Eight participants explicitly mentioned that Zoom has better features than other platforms, for example, the live stream feature or the ability to raise hands.

Some participants said that they observed improvements in the security posture of Zoom since the beginning of the pandemic:

> *"Zoom fixed issues with screen hijacking, people having to register, issues with encryption." (P1)*

At the same time, other participants were critical of how seriously Zoom takes security:

> *"Zoom engineers are not good, the security team is not capable, they messed up end-to-end encryption, etc. Jitsi has security in mind and works closely with activist groups." (P6)*

Attitudes towards online meetings. When asked about their attitude towards online meetings, most participants held a positive one, saying that they help remove physical barriers among people and reach a broader audience, as well as allow to recruit of better speakers without the need for travel:

> *"A big advantage of Zoom is that you don't have to wait for people to be local, way better people can come and speak." (P4)*

Five participants preferred the online meeting format, especially during the pandemic:

> *"A lot of people still don't want to meet in person due to health concerns, so hybrid meetings are good. Hard to maintain though, keeping track of people in physical attendance and the virtual audience too." (P2)*

Five participants would like to keep a hybrid format even for future meetings. At the same time, they noted that people are less engaged in an online meetings than in-person ones (P1).

Takeaways. We find that the attacks are often successful in disrupting meetings, making it difficult for them to resume. Victims react differently to these attacks: while some participants did not make a big deal of it, others were deeply hurt. Reporting to platforms or IT departments does not seem to lead to effective actions, and law enforcement acknowledges that they might be ill-equipped to deal with this problem as well.

5 DISCUSSION

Our interview study provided us with a multi-faceted view of the Zoombombing attack. In this section, we summarize our findings and discuss their implications for users of online meeting tools and for the platforms themselves. By examining recent privacy and security features implemented by Zoom and other video conferencing technologies, we highlight how the computer security community can develop better mitigations against this threat. We then discuss the limitations of our work.

5.1 Implications for Usable Design as Security Mitigations

A hybrid lifestyle after the pandemic forced people who were not technology-savvy to adopt online meeting tools as well. In this climate, hosts often do not adopt adequate security countermeasures because they feel that the benefit offered by them is outweighed by the burden that they put on their users. For example, previous research recommends setting up meetings that require registration and sending personalized invitation links to users [26], while our study finds this impractical, especially for large meetings.

Although our participants reported that experiencing Zoombombing made them reconsider the importance of taking preventative measures and they are now more aware of the potential risks of publicly sharing meeting links, there still efforts should be done by the platforms.

Reported by Zoom, compared to the early days of the pandemic in 2020, incidents of Zoombombing and other security breaches have decreased [53]. The introduction of mandatory passcodes, waiting rooms, and the At-Risk Meeting Notifier tool has enhanced security, reducing unauthorized access to meetings. In the meantime, Zoom and other video conferencing technologies have implemented various privacy and security features [13, 54]. These advancements focus on data encryption—Zoom introduced end-to-end encryption for meetings in 2020 and expanded it to Zoom Phone in 2022—and addressing vulnerabilities like weak anti-tampering mechanisms and susceptibility to brute force attacks.

While many advancements in online meeting tools focus on stricter default settings and advanced encryption, usable security features are less emphasized. Zoom, for example, added a feature to lock the meeting after it begins, helping to prevent Zoombombing [34]. To help meeting hosts secure their meetings, we encourage online meeting platforms to improve both usability and security. Simple adjustments, such as avoiding the reuse of the same meeting link for multiple events and sharing the link only via email, can prevent opportunistic attacks. Platforms should consider disabling the ability to copy-paste meeting links and instead encourage sharing details through secure methods like calendar invites and email.

A general trend that we observe is that Zoombombing attacks are disorientating to both hosts and participants. On one hand, hosts often do not have a clear idea of what to do to mitigate the attack, and they have to deal with attackers changing their names and impersonating them or other participants. This makes it difficult for them to identify the attackers and kick them out, forcing them to stop the meeting and start a new one. In turn, this further confuses participants, who may not be able to find the new link in time and rejoin the meeting. A potential solution to clear this confusion could be to have online platforms develop an easy way for all people who are registered for a meeting to communicate (i.e., an out-of-band chat). Developing this tool would however require properly weighing the pros and

cons, as this could become an additional vector for abuse and harassment. Alternatively, the platform could simply prohibit users from adopting the same or a similar user name as other participants or require host approval to do so.

Our interviews also shed light on the need for better mechanisms to identify and kick out attackers during a Zoombombing incident. Right now, Zoom has an indicator showing who the legitimate host of a meeting is, but this indicator only appears in the participants' window, and not in the video and chat panels, where Zoombombing attacks (and impersonation) actually take place. Small feature improvements like adding a checkmark to the host in every panel on Zoom would make impersonation more difficult. Similarly, providing a verification logo for authenticated users is a viable way to reduce participants' disorientation during Zoombombing, and help the host identify attackers and kick them from the meeting.

5.2 Need for Inclusive Online Safety

Our findings also highlight that people react very differently to Zoombombing attacks; while some users simply shrug off the experience, others are significantly harmed by it. This is particularly true for older users, who may not have encountered online toxicity before. Additionally, we find that attackers often target the race or gender of participants to cause harm, which aligns with existing online safety research [31, 42]. This research reveals the complex landscape of cyber harassment, particularly for LGBTQ+ individuals and female adolescents. These findings underscore the necessity for tailored, inclusive, and supportive online environments to mitigate the negative impacts of cyber harassment. The computer security research community has recently turned its attention to marginalized communities and their security and privacy needs [8, 17, 21, 45]. Future research in this space could investigate the experience of victims of Zoombombing, and device-tailored interventions to support them, for example by adopting trauma-informed computing [7].

5.3 Support and Legal Challenges in Addressing Zoombombing Attacks

We find that participants rarely received help from either the meeting platform or their institution's IT during or after the attack. Meeting platforms typically provide channels for incident reports after the fact, instead of real-time support. Online resources provided by the platforms generally focus on technical solutions without expressing sufficient empathy or awareness of the victims' psychological harm. We recommend that meeting platforms could provide a support channel to help victims of Zoombombing deal with the aftermath of the attack even after it took place.

We also notice that the legal area in which Zoombombing attacks operate is often unclear, and law enforcement is usually unable to act since abusive activity falls under free speech [25, 35]. However, our participants experienced actions that are clearly illegal, like the sharing of child pornography by attackers. In these cases, platforms should develop more effective channels to enable participants to report these incidents and allow law enforcement to act more swiftly. We hope that this study will raise awareness on the importance of security and privacy of online meeting tools. This is particularly important as more and more people rely on online communication tools like Zoom, to work, learn, and socialize, making it essential to understand the risks and threats associated with these platforms and establish a corresponding policy to prevent harm to users.

5.4 Limitations

Our study sheds light on the experiences of Zoombombing victims. While our sample size may seem small ($N = 15$), we followed common qualitative analysis practices to inform our decision to stop the interviews: the same themes kept repeating in the interviews as we kept going, which indicates that the study reached data saturation. We face a common

limitation of qualitative work [17, 18] in that our results might not cover the entirety of Zoombombing experiences out there.

For instance, none of our participants experienced Zoombombing attacks during large public meetings like political rallies or performances.

Although one-third of our participants work as college teachers, researchers, and graduate students, we did not manage to recruit any high school teachers for our study. Previous research [26] suggested that Zoombombing attacks are often happening in high school online lectures. We acknowledge this as a limitation of our study, although we note that since previous work focused on the attacking threads posted on polarized online communities like 4chan, the teenager demographics might have been over-represented. On the other hand, our study uses professional platforms like LinkedIn to recruit participants, showing that Zoombombing in a professional setting is a common occurrence.

6 CONCLUSION

In this paper, we conducted the first interview study of victims of Zoombombing ($N = 15$). We find that meetings usually get attacked because links to join them get posted publicly, for example on social media. We also find that in the mind of meeting hosts, security is often an after thought, and that ease of access for participants is what takes priority. Our research identifies several interesting directions that the computer security community could take to help solve the problem, from devising more frictionless mitigations that do not drive users away to developing automated techniques to detect Zoombombing attacks early on. Finally, we highlight the importance of understanding how vulnerable and marginalized demographics are targeted by Zoombombing attacks, which could inform how to best provide support to a diverse set of victims in the aftermath of an attack.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their insightful comments that helped us improve this paper. This work was supported by the National Science Foundation under Grant CNS-1942610 and by a Meta Research PhD Fellowship.

REFERENCES

[1] Kawsar Ali. 2021. Zoom-ing in on White Supremacy : Zoom-Bombing Anti-Racism Efforts. *M/C Journal* 24, 3 (June 2021). https://doi.org/10.5204/mcj.2786

[2] Max Aliapoulios, Kejsi Take, Prashanth Ramakrishna, Daniel Borkan, Beth Goldberg, Jeffrey Sorensen, Anna Turner, Rachel Greenstadt, Tobias Lauinger, and Damon McCoy. 2021. A large-scale characterization of online incitements to harassment across platforms. In *Proceedings of the 21st ACM Internet Measurement Conference* (Virtual Event) *(IMC '21)*. Association for Computing Machinery, New York, NY, USA, 621–638. https://doi.org/10.1145/3487552.3487852

[3] Adam Badawy, Emilio Ferrara, and Kristina Lerman. 2018. Analyzing the Digital Traces of Political Manipulation: The 2016 Russian Interference Twitter Campaign. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. 258–265. https://doi.org/10.1109/ASONAM.2018.8508646

[4] Fabricio Benevenuto, Gabriel Magno, Tiago Rodrigues, and Virgilio Almeida. 2010. Detecting spammers on Twitter. In *Collaboration, electronic messaging, anti-abuse and spam conference (CEAS)*. https://homepages.dcc.ufmg.br/~fabricio/download/ceas10.pdf

[5] Virginia Braun and Victoria Clarke. 2012. Thematic analysis. (2012). https://doi.org/10.1037/13620-004

[6] Barry Brown. 2020. Notes on running an online academic conference or how we got zoombombed and lived to tell the tale. *Interactions* 27, 4 (jul 2020), 16–21. https://doi.org/10.1145/3406108

[7] Janet X. Chen, Allison McDonald, Yixin Zou, Emily Tseng, Kevin A Roundy, Acar Tamersoy, Florian Schaub, Thomas Ristenpart, and Nicola Dell. 2022. Trauma-Informed Computing: Towards Safer Technology Experiences for All. In *Proceedings of the 2022*

CHI Conference on Human Factors in Computing Systems (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 544, 20 pages. https://doi.org/10.1145/3491102.3517475

[8] Kovila P.L. Coopamootoo, Maryam Mehrnezhad, and Ehsan Toreini. 2022. "I feel invaded, annoyed, anxious and I may protect myself": Individuals' Feelings about Online Tracking and their Protective Behaviour across Gender and Country. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 287–304. https://www.usenix.org/conference/usenixsecurity22/presentation/coopamootoo

[9] Clayton Allen Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. BotOrNot: A System to Evaluate Social Bots. In *Proceedings of the 25th International Conference Companion on World Wide Web* (Montréal, Québec, Canada) *(WWW '16 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 273–274. https://doi.org/10.1145/2872518.2889302

[10] Emiliano De Cristofaro, Arik Friedman, Guillaume Jourjon, Mohamed Ali Kaafar, and M. Zubair Shafiq. 2014. Paying for Likes? Understanding Facebook Like Fraud Using Honeypots. In *Proceedings of the 2014 Conference on Internet Measurement Conference* (Vancouver, BC, Canada) *(IMC '14)*. Association for Computing Machinery, New York, NY, USA, 129–136. https://doi.org/10.1145/2663716.2663729

[11] Brian Dean. 2022. Zoom User Stats: How Many People Use Zoom in 2022? https://backlinko.com/zoom-users.

[12] Greg Elmer, Stephen J. Neville, Anthony Burton, and Sabrina Ward-Kimola. 2021. Zoombombing During a Global Pandemic. *Social Media + Society* 7, 3 (2021). https://doi.org/10.1177/20563051211035356

[13] eSecurity Planet. 2023. Zoom Security Issues: A Wakeup Call for Enterprises. https://www.esecurityplanet.com/threats/zoom-security-issues-a-wakeup-call-for-enterprises/ Accessed: 2024-07-18.

[14] FBI. 2020. FBI Warns of Teleconferencing and Online Classroom Hijacking During COVID-19 Pandemic. https://www.fbi.gov/contact-us/field-offices/boston/news/press-releases/fbi-warns-of-teleconferencing-and-online-classroom-hijacking-during-covid-19-pandemic

[15] Brian Friedberg, Gabrielle Lim, and Joan Donovan. 2020. Space invaders: The networked terrain of zoom bombing. *Technology and Social Change Research Project* (2020). https://shorensteincenter.org/wp-content/uploads/2020/08/TaSC_Zoom-Bombing_August-2020.pdf

[16] Hongyu Gao, Jun Hu, Christo Wilson, Zhichun Li, Yan Chen, and Ben Y. Zhao. 2010. Detecting and characterizing social spam campaigns. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement* (Melbourne, Australia) *(IMC '10)*. Association for Computing Machinery, New York, NY, USA, 35–47. https://doi.org/10.1145/1879141.1879147

[17] Christine Geeng, Mike Harris, Elissa Redmiles, and Franziska Roesner. 2022. "Like Lesbians Walking the Perimeter": Experiences of U.S. LGBTQ+ Folks With Online Security, Safety, and Privacy Advice. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 305–322. https://www.usenix.org/conference/usenixsecurity22/presentation/geeng

[18] Nitesh Goyal, Leslie Park, and Lucy Vasserman. 2022. "You have to prove the threat is real": Understanding the needs of Female Journalists and Activists to Document and Report Online Harassment. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 242, 17 pages. https://doi.org/10.1145/3491102.3517517

[19] Chris Grier, Kurt Thomas, Vern Paxson, and Michael Zhang. 2010. @spam: the underground on 140 characters or less. In *Proceedings of the 17th ACM Conference on Computer and Communications Security* (Chicago, Illinois, USA) *(CCS '10)*. Association for Computing Machinery, New York, NY, USA, 27–37. https://doi.org/10.1145/1866307.1866311

[20] Gabriel Hine, Jeremiah Onaolapo, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Riginos Samaras, Gianluca Stringhini, and Jeremy Blackburn. 2017. Kek, cucks, and god emperor trump: A measurement study of 4chan's politically incorrect forum and its effects on the web. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 11. 92–101. https://doi.org/10.1609/icwsm.v11i1.14893

[21] Deepak Kumar, Patrick Gage Kelley, Sunny Consolvo, Joshua Mason, Elie Bursztein, Zakir Durumeric, Kurt Thomas, and Michael Bailey. 2021. Designing toxic content classification for a diversity of perspectives. In *Proceedings of the Seventeenth USENIX Conference on Usable Privacy and Security (SOUPS'21)*. USENIX Association, USA, Article 16, 19 pages. https://dl.acm.org/doi/10.5555/3632186.3632187

[22] Srijan Kumar, William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2018. Community Interaction and Conflict on the Web. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 933–943. https://doi.org/10.1145/3178876.3186141

[23] J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics* (1977), 159–174. https://pubmed.ncbi.nlm.nih.gov/843571/

[24] Claire Seungeun Lee. 2021. Analyzing Zoombombing as a new communication tool of cyberhate in the COVID-19 era. *Online Information Review* (2021). https://www.emerald.com/insight/content/doi/10.1108/OIR-05-2020-0203/full/html

[25] Claire Seungeun Lee and Ahnlee Jang. 2023. Sharing experiences and seeking informal justice online: a grounded theory analysis of Zoombombing victimization on Reddit. *Victims & Offenders* 18, 5 (2023), 988–1007. https://doi.org/10.1080/15564886.2023.2171169

[26] Chen Ling, Utkucan Balcı, Jeremy Blackburn, and Gianluca Stringhini. 2021. A first look at zoombombing. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1452–1467. https://doi.org/10.1109/SP40001.2021.00061

[27] Chen Ling, Gianluca Stringhini, Utkucan Balcı, and Jeremy Blackburn. 2022. A First Look at Zoombombing. *IEEE Security & Privacy* 20, 2 (2022), 22–30. https://doi.org/10.1109/MSEC.2021.3127392

[28] Luca Luceri, Silvia Giordano, and Emilio Ferrara. 2020. Detecting Troll Behavior via Inverse Reinforcement Learning: A Case Study of Russian Trolls in the 2016 US Election. *Proceedings of the International AAAI Conference on Web and Social Media* 14, 1 (May 2020), 417–427. https://doi.org/10.1609/icwsm.v14i1.7311

[29] Enrico Mariconti, Guillermo Suarez-Tangil, Jeremy Blackburn, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Jordi Luque Serrano, and Gianluca Stringhini. 2019. "You Know What to Do": Proactive Detection of YouTube Videos Targeted by Coordinated Hate Attacks. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 207 (nov 2019), 21 pages. https://doi.org/10.1145/3359309

[30] Anastasia Powell, Adrian J Scott, and Nicola Henry. 2020. Digital harassment and abuse: Experiences of sexuality and gender minority adults. *European Journal of Criminology* 17, 2 (2020), 199–223. https://doi.org/10.1177/1477370818788006

[31] Afsaneh Razi, Karla Badillo-Urquiola, and Pamela J Wisniewski. 2020. Let's talk about sext: How adolescents seek support and advice about their online sexual experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13. https://doi.org/10.1145/3313831.3376400

[32] Mohammad Hammas Saeed, Shiza Ali, Jeremy Blackburn, Emiliano De Cristofaro, Savvas Zannettou, and Gianluca Stringhini. 2022. TROLLMAGNIFIER: Detecting state-sponsored troll accounts on reddit. In *IEEE Symposium on Security and Privacy (SP)*. IEEE, 2161–2175. https://doi.org/10.1109/SP46214.2022.9833706

[33] Ion-Alexandru Secara. 2020. Zoombombing – the end-to-end fallacy. *Network Security* 2020, 8 (2020), 13–17. https://doi.org/10.1016/S1353-4858(20)30094-5

[34] Sectigo Store. 2023. How to Secure Zoom: 7 Important Zoom Security Settings. https://sectigostore.com/blog/how-to-secure-zoom-7-important-zoom-security-settings/ Accessed: 2024-07-18.

[35] Kamran Siddiqui and Shabir Ahmad. 2023. Zoombombing: causes and preventions. In *E3S Web of Conferences*, Vol. 371. EDP Sciences, 05026. https://www.e3s-conferences.org/articles/e3sconf/pdf/2023/08/e3sconf_afe2023_05026.pdf

[36] Matthew Skiles, Euijin Yang, Orad Reshef, Diego Robalino Muñoz, Diana Cintron, Mary Laura Lind, Alexander Rush, Patricia Perez Calleja, Robert Nerenberg, Andrea Armani, et al. 2022. Conference demographics and footprint changed by virtual platforms. *Nature Sustainability* 5, 2 (2022), 149–156. https://doi.org/10.1038/s41893-021-00823-2

[37] Peter Snyder, Periwinkle Doerfler, Chris Kanich, and Damon McCoy. 2017. Fifteen minutes of unwanted fame: detecting and characterizing doxing. In *Proceedings of the 2017 Internet Measurement Conference* (London, United Kingdom) *(IMC '17)*. Association for Computing Machinery, New York, NY, USA, 432–444. https://doi.org/10.1145/3131365.3131385

[38] Kate Starbird. 2017. Examining the Alternative Media Ecosystem Through the Production of Alternative Narratives of Mass Shooting Events on Twitter. *Proceedings of the International AAAI Conference on Web and Social Media* 11, 1 (May 2017), 230–239. https://doi.org/10.1609/icwsm.v11i1.14878

[39] Kate Starbird, Ahmer Arif, and Tom Wilson. 2019. Disinformation as Collaborative Work: Surfacing the Participatory Nature of Strategic Information Operations. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 127 (nov 2019), 26 pages. https://doi.org/10.1145/3359229

[40] Gianluca Stringhini, Christopher Kruegel, and Giovanni Vigna. 2010. Detecting spammers on social networks. In *Proceedings of the 26th Annual Computer Security Applications Conference* (Austin, Texas, USA) *(ACSAC '10)*. Association for Computing Machinery, New York, NY, USA, 1–9. https://doi.org/10.1145/1920261.1920263

[41] Gianluca Stringhini, Gang Wang, Manuel Egele, Christopher Kruegel, Giovanni Vigna, Haitao Zheng, and Ben Y. Zhao. 2013. Follow the green: growth and dynamics in twitter follower markets. In *Proceedings of the 2013 Conference on Internet Measurement Conference* (Barcelona, Spain) *(IMC '13)*. Association for Computing Machinery, New York, NY, USA, 163–176. https://doi.org/10

.1145/2504730.2504731

[42] Tangila Islam Tanni, Mamtaj Akter, Joshua Anderson, Mary Jean Amon, and Pamela J. Wisniewski. 2024. Examining the Unique Online Risk Experiences and Mental Health Outcomes of LGBTQ+ versus Heterosexual Youth. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 867, 21 pages. https://doi.org/10.1145/3613904.3642509

[43] Kurt Thomas, Devdatta Akhawe, Michael Bailey, Dan Boneh, Elie Bursztein, Sunny Consolvo, Nicola Dell, Zakir Durumeric, Patrick Gage Kelley, Deepak Kumar, Damon McCoy, Sarah Meiklejohn, Thomas Ristenpart, and Gianluca Stringhini. 2021. SoK: Hate, Harassment, and the Changing Landscape of Online Abuse. In *2021 IEEE Symposium on Security and Privacy (SP)*. 247–267. https://doi.org/10.1109/SP40001.2021.00028

[44] Colin G Walsh, Kim M Unertl, and Jon S Ebert. 2021. Rapid supportive response to a traumatic "zoombombing" during the COVID-19 pandemic. *Academic Medicine* 96, 1 (2021), e6–e7. https://doi.org/10.1097/ACM.0000000000003739

[45] Noel Warford, Tara Matthews, Kaitlyn Yang, Omer Akgul, Sunny Consolvo, Patrick Gage Kelley, Nathan Malkin, Michelle L. Mazurek, Manya Sleeper, and Kurt Thomas. 2022. SoK: A Framework for Unifying At-Risk User Research. In *2022 IEEE Symposium on Security and Privacy (SP)*. 2344–2360. https://doi.org/10.1109/SP46214.2022.9833643

[46] Janith Weerasinghe, Bailey Flanigan, Aviel Stein, Damon McCoy, and Rachel Greenstadt. 2020. The Pod People: Understanding Manipulation of Social Media Popularity via Reciprocity Abuse. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) *(WWW '20)*. Association for Computing Machinery, New York, NY, USA, 1874–1884. https://doi.org/10.1145/3366423.3380256

[47] Chao Yang, Robert Chandler Harkreader, and Guofei Gu. 2011. Die Free or Live Hard? Empirical Evaluation and New Design for Fighting Evolving Twitter Spammers. In *Recent Advances in Intrusion Detection*, Robin Sommer, Davide Balzarotti, and Gregor Maier (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 318–337. https://doi.org/10.1007/978-3-642-23644-0_17

[48] Sarah Young. 2021. Zoombombing Your Toddler: User Experience and the Communication of Zoom's Privacy Crisis. *Journal of Business and Technical Communication* 35, 1 (2021), 147–153. https://doi.org/10.1177/1050651920959201

[49] Dong Yuan, Yuanli Miao, Neil Zhenqiang Gong, Zheng Yang, Qi Li, Dawn Song, Qian Wang, and Xiao Liang. 2019. Detecting Fake Accounts in Online Social Networks at the Time of Registrations. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 1423–1438. https://doi.org/10.1145/3319535.3363198

[50] Savvas Zannettou, Tristan Caulfield, William Setzer, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2019. Who Let The Trolls Out? Towards Understanding State-Sponsored Trolls. In *Proceedings of the 10th ACM Conference on Web Science* (Boston, Massachusetts, USA) *(WebSci '19)*. Association for Computing Machinery, New York, NY, USA, 353–362. https://doi.org/10.1145/3292522.3326016

[51] Savvas Zannettou, Mai Elsherief, Elizabeth Belding, Shirin Nilizadeh, and Gianluca Stringhini. 2020. Measuring and Characterizing Hate Speech on News Websites. In *Proceedings of the 12th ACM Conference on Web Science* (Southampton, United Kingdom) *(WebSci '20)*. Association for Computing Machinery, New York, NY, USA, 125–134. https://doi.org/10.1145/3394231.3397902

[52] Zoom. 2016. How to Keep Uninvited Guests Out of Your Zoom Event. https://blog.zoom.us/keep-uninvited-guests-out-of-your-zoom-event/.

[53] Zoom Video Communications, Inc. 2023. 4 Ways Zoom Works to Protect User Privacy. https://www.zoom.com/en/blog/4-ways-zoom-works-to-protect-user-privacy/ Accessed: 2024-07-18.

[54] Zoom Video Communications, Inc. 2023. Which Zoom Security Features Are Best for Your Industry? https://www.zoom.com/en/blog/which-zoom-security-features-are-best-for-your-industry/ Accessed: 2024-07-18.

A INTERVIEW QUESTIONS

0. Get to know the victim.

- Please provide your demographic information.

1. User habits.

- How often do you use online meeting tools?
- How many Zoombombing experiences do you have?
- What meeting tools do you use?

- What role did you play in the attacked meeting (i.e., host or attendees)?

2. Meeting settings.

- When was the meeting?
- How big was the audience of the meeting?
- What was the purpose of the meeting?
- How long was the meeting purposed to last?
- What precautions (if any) did the host take to secure the meeting?

3. Delivery of the attack.

- How was the Zoombombing attack launched?
- What kind of actions did the attackers perform?
- Was the attack targeted at one or more participants?
- What was the purpose of the attack?
- How long did the attack last?
- How many attackers?
- Did you see/hear the attackers?
- Did you/other participants of the meetings know any of the attackers?
- Please describe the attackers.
- Please describe the Zoombombing attack in details.
- Do you have any ideas of why the meeting being bombed?

4. Mitigation of the attack.

- What did the host do to stop the attack?
- How did the attackers react to these mitigations?

5. Attack aftermath.

- How did the attack make the you feel?
- Are you now more reluctant to attend online meetings?
- Did you talk with others about the Zoombombing attack?
- Did the attack motivate them to learn more about security?
- Is the participant now taking more precautions when hosting online meetings?
- Did the host report the attack to their IT, the meeting platform, or law enforcement?
- Did you consider changing the meeting tool?
- What is your attitude towards the online meeting tools?
- Have you perceived any improvement in security measures of meeting tools? Do you prefer online or in-person meetings, or hybrid?

ID	Year	Platforms	Reported by	Size	Type	Scheduled length	Precaution
M1	2020	Zoom	Attendee	10	Research Seminar	>45min	NA
M2	2022	Zoom	Host	10	Hobby Gathering	>45min	No
M3	2021	Zoom	Attendee	30	Educational Talk	>45min	NA
M4	2022	Zoom	Host	180	Hobby Gathering	>45min	Password
M5	2020	Zoom	Attendee	10	Hobby Gathering	>45min	Registration
M6	2022	Zoom	Host &Attendee	30	Research Seminar	>45min	Password
M7	2022	Webex	Attendee	10	Research Seminar	>45min	NA
M8	2021	Zoom	Attendee	20	Research Seminar	>45min	Default muted & camera off & password
M9	2020	Zoom	Host	20	Work Meeting	>45min	Password & authenticated account
M10	2022	Zoom	Attendee	50	Research Seminar	>45min	NA
M11	2020	Zoom	Attendee	40	Hobby Gathering	>45min	NA
M12	2021	Zoom	Host	130	Working Meeting	>45min	No
M13	2021	Zoom	Host	5	Educational Talk	>45min	No
M14	2022	Zoom	Host	70	Research Seminar	>45min	No

Table 2. Information on zoombombed meetings ($N = 14$).

ID	Gender	Seniority	Race	Place of interview	Occupation	Security-related
P1	F	Mid	White	US	Researcher/Teacher	Y
P2	M	Senior	White	US	Engineer	N
P3	M	Junior	Black	US	Engineer	N
P4	M	Senior	White	US	Film editor	N
P5	F	Mid	White	CA	Anti-racism Activist	N
P6	M	Mid	White (Jewish)	US	Researcher/Teacher	Y
P7	M	Mid	White	US	Engineer	Y
P8	M	Junior	Black	US	Student	Y
P9	F	Mid	White	US	Researcher/Teacher	N
P10	M	Junior	Asian	US	Student	Y
P11	F	Mid	White	CA	Researcher	N
P12	F	Mid	White	UK	Office Administrator	N
P13	M	Senior	White	US	Company Executive	N
P14	F	Mid	White	FR	Researcher/Teacher	N
P15	M	Senior	Asian	US	Researcher/Teacher	N

Table 3. Socio-demographic profile of the participants to this study.

Variables	Fraction
Number of attackers	
Single	2/14
Multiple	12/14
Known attacker	
Yes	0/14
No	11/14
Unclear	3/14
Targeted	
	9/14
Attack length	
Less than 5 minutes	11/14
Over 5 minutes	3/14
Goal of attack	
Racism	9/14
Disruption	4/14
Fun	3/14
Sexism	1/14
Unclear	2/14
Attack formats	
Impersonation	3/14
Re-entering and attack	7/14
Audio:	
yelling, shouting hate speech	8/14
play audio	5/14
Video:	
play video	3/14
turn the camera on	3/14
Desktop:	
screensharing, annotating hateful language	3/14
change the profile photo into nasty image	2/14
Post offended text in chat	5/14
Host	
Kicking the attackers off	14/14
Mute the attacker	1/14
End the meeting for all	7/14
Initiate another meeting	6/14
Others	4/14
Looking for support from organizations	
	5/14
Report to the platforms or law enforcement	
	5/14

Table 4. Details of the delivery of Zoombombing attacks ($N = 14$).

Variables	Fraction
Feeling reluctant to attend online meetings	
	2/15
Talking about Zoombombing experience with others	
	9/15
Considering changing tools	
	2/15
Attitudes towards online meeting tools	
Positive	8/15
Neutral	4/15
Negative	3/15
Perceived improvements in the security measurements	
	4/15
Applying more precaution in future meetings	
	15/15
Preferred meeting format	
Online	5/15
In-person	5/15
Hybrid	5/15

Table 5. Harm caused by Zoombombing as reported by participants ($N = 15$)