Balancing Privacy and Serendipity in Cyberspace

Mahadev Satyanarayanan
Carnegie Mellon University

Nigel Davies
Lancaster University

Nina Taft
Google

Abstract

Unplanned encounters or casual collisions between colleagues have long been recognized as catalysts for creativity and innovation. The absence of such encounters has been a negative side effect of COVID-enforced remote work. However, there have also been positive side effects such as less time lost to commutes, lower carbon footprints, and improved work-life balance. This vision paper explores how serendipity for remote workers can be created by leveraging IoT technologies, edge computing, high-resolution video, network protocols for live interaction, and video/audio denaturing. We reflect on the privacy issues that technology-mediated serendipity raises and sketch a path towards honoring diverse privacy preferences.

CCS Concepts • **Human-centered computing → Ubiquitous and mobile computing systems and tools**; **Collaborative and social computing systems and tools**; • **Security and privacy → Social aspects of security and privacy**; **Usability in security and privacy**; **Privacy protections**.

ACM Reference Format:
Mahadev Satyanarayanan, Nigel Davies, and Nina Taft. 2022. Balancing Privacy and Serendipity in Cyberspace. In *The 23rd International Workshop on Mobile Computing Systems and Applications (HotMobile '22), March 9–10, 2022, Tempe, AZ, USA*. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3508396.3512873

1 Triggering the Creative Spark

Unexpected encounters and unplanned interactions are among the biggest joys of resuming in-person work after 18 months of COVID-enforced isolation. This is the essence of *serendipity,* which is defined as "the faculty or phenomenon of finding valuable or agreeable things not sought for" [24]. Its importance to innovation has been widely recognized by many [20, 22, 27], including Steve Jobs who is quoted by biographer Walter Isaacson as saying:

> "There's a temptation in our networked age to think that ideas can be developed by email and iChat — that's crazy. Creativity comes from spontaneous meetings, from random discussions. You run into someone, you ask what they're doing, you say 'Wow,' and soon you're cooking up all sorts of ideas." [18]

To facilitate serendipity and collaboration, new corporate campuses have been created in the past decade. Yet, a year of COVID-enforced remote work has also revealed many benefits. Less time is lost to commutes, and carbon footprints are lower. Remote work also offers a better work-life balance for many people. Can we preserve these benefits without sacrificing serendipity?

A recent paper [40] suggests that the lack of serendipity leads to siloing. Based on an analysis of interactions between Microsoft employees before and after the COVID shutdown, it states:

> "Our results show that firm-wide remote work caused the collaboration network of workers to become more static and siloed, with fewer bridges between disparate parts. Furthermore, there was a decrease in synchronous communication and an increase in asynchronous communication. Together, these effects may make it harder for employees to acquire and share new information across the network."

Many companies that have resumed in-person work implement "de-densification" strategies. Examples include staggered lunch schedules in cafeterias, and randomized in-person attendance on three days a week. Such measures have been shown to have negative impact on informal communication among employees [25]. In addition to being a catalyst for innovation, serendipity is also valuable because it can mitigate siloing. It is often through unplanned interactions that you learn of important work-related developments.

In this vision paper, we ask *"Can we create technology-mediated serendipity for coworkers who are not colocated?"* We reflect on how such a capability could be implemented, and what attributes would needed for such a system to succeed in the real world. In spite of the speculative nature of this paper, the vision it describes is not science fiction. The technological building blocks for this vision are already available. These include large displays, high-resolution video cameras, highly accurate and fast face recognition, edge computing, last-mile fiber networks for low latency and high bandwidth, and protocols (e.g., Zoom) for live interactions over the Internet. In Section 3, we sketch how these can be integrated into a system that triggers privacy-controlled chance encounters.

Technology-mediated serendipity raises serious privacy issues. We believe that it is important to think through these issues up front, before attempting a system implementation. With that goal in mind, we discuss many facets of the privacy-serendipity tradeoff space and sketch a path towards honoring the privacy preferences of diverse stakeholders. This paper makes the following contributions:

- It proposes a new form of technology-mediated serendipity for remote workers.
- It highlights the complex options and requirements for implementing this concept via edge computing, networking, face recognition, personal data management, user controls, and cross-organization policies.
- It explores the privacy challenges that arise in this context, and identifies the diverse stakeholders who need to be involved in any workable solution.
- It proposes a novel solution to bystander privacy.

2 Background and Related Work

Blending physical and virtual presence to support new forms of working has a rich research history dating back to the mid 1980's and Xerox PARC's work on Media Spaces [37]. Early research focused on creating AV links between remote spaces – perhaps most famously by Galloway and Rabinowitz in their art installation "Hole-In-Space" [14] in which large displays in two shop windows, one in New York and one in Los Angeles, were linked with a persistent AV connection. Throughout the late 1980s and 1990s multiple projects

HotMobile '22, March 9–10, 2022, Tempe, AZ, USA

ACM ISBN 978-1-4503-9218-1/22/03.
https://doi.org/10.1145/3508396.3512873

explored long-lived connections to facilitate remote collaboration, e.g. Xerox's Media Spaces, EuroPARC's RAVE [15] and Bellcore's Video Window and Cruiser systems [12].

Privacy issues in the context of media spaces have been extensively explored, often focusing on the inevitable tension between the desire for awareness and the risk of privacy invasion. Hudson and Smith [17] explored how to balance awareness and privacy in systems that support distributed work groups. They observed that there is a dual tradeoff between privacy and awareness on the one hand, and between awareness and disturbance on the other. Boyle et al [8] present a comprehensive framework and associated vocabulary for analysing media spaces in terms of privacy. The framework highlights many of the challenges in providing users with privacy controls, even when technology such as blur filters are provided. This is primarily because of the difficulty in offering users sufficiently fine-grained control, and because of the need to apply these controls to a range of channels beyond video (e.g. audio must also be filtered). Studies of privacy concerns in media spaces have been conducted at scale, notably by Friedman et al [13] who explored user reactions to a public media space deployment. As in previous work, the study illustrated the multiple factors that influence users' perceptions of privacy.

Technologies such as Skype, Zoom, Teams and Alexa (via Drop Ins) have made high-quality video calls an integral part of daily life. Their widespread has given rise to privacy concerns [19] and interest in new techniques for addressing these concerns [23]. Since these video calling technologies are not bound to physical spaces, they lack the semi-public nature of media space prototypes. Further, since video calls are explicitly initiated by the user, there is no serendipity in their use. If an autonomous mechanism for triggering serendipitous interactions were to be created (such as the one described in this paper), the resulting video/audio session could be implemented using one of these technologies.

Social apps such as WeChat allow a user to voluntarily reveal his or her presence to others in physical proximity. The intended purpose is to invite messaging with colocated strangers. A similar capability has been leveraged by contract-tracing apps for COVID.

The use of public displays to support social interactions in workplaces has been explored in systems such as the Notification Collage [16] and the Aware Community Portal [33] that sought to increase awareness of coworkers activities. Brignall and Rodger's Opinionizer [9] extended the notion of awareness, and sought to actively catalyze discussion between physically colocated colleagues.

Our vision differs from prior work in three key ways:

- Unlike early persistent media links, we envision transient links being established on demand when the presence of suitable pairs of viewers are detected. Crucially, we capitalize on organizations' rich knowledge of the collaboration networks and activities of individual workers to trigger connections.
- In contrast to workplace awareness systems, we do not seek to build connections between physically colocated workers. Rather, we aim to strengthen links between coworkers who are too far apart to experience in-person serendipity.
- We leverage recent work in *privacy mediation* [11] that allows live video and audio links to be stripped in real time of information that would leak privacy. This greatly simplifies bystander privacy.

> Alice is at the elevator, with her watch set to OPT-IN for Pomme. She hears a gentle chime. On the large wall display, she sees her co-worker Bob at a coffee machine. He works at another site of Alice's company. Bob hears a chime too, and sees Alice on the large display near him. They smile and wave at each other, and start talking. During their conversation, people entering and exiting the elevator near Alice appear as solid blobs on Bob's display. Bob can tell that they are people, but nothing more. Similarly, Alice sees people around Bob as solid moving blobs. Bob can only hear Alice, and vice versa; they can't hear the conversations of blobbed-out people. It is soon clear that their latest projects have a lot of synergy. Alice suggests a follow-on meeting for a deeper discussion. Energized by their discussion, Alice and Bob say goodbye and return to work.

Figure 1. FreeZones at Work

3 A World With Remote Serendipity

3.1 Overview

Imagine a system called *Pomme* that implements remote serendipity in the workplace. Pomme consists of many small privacy-controlled physical spaces called *FreeZones* in which it triggers remote encounters. A FreeZone might be located in the lobby of an office building, in front of an elevator, in a lunch room at work, in a printer room, near the water cooler, and so on. The boundaries of a FreeZone are clearly marked, and there are large signs indicating that you are in a FreeZone. End-to-end latency limits the acceptable dispersion of FreeZones. Today, with fiber connectivity and edge computing, FreeZones can easily be separated by a hundred kilometers or more. That is large enough to encompass a major city such as London or New York, as well as its outermost exurbs.

The scenario in Figure 1 illustrates the kind of user experience that we envision from FreeZones in a work setting. Considerable implementation complexity is hidden in this simple scenario. For example, if Alice is already talking on her phone or to another person in the FreeZone, she should not be selected for a trigger. How best to achieve such context sensitivity is an open question. We discuss some possible approaches in Section 3.2.

Real-life serendipity is not restricted to the workplace. It is technically feasible to create a FreeZone at home. However, the privacy challenges are daunting, especially if children are present. Hence, we focus exclusively on workplace settings in this paper.

3.2 Design Alternatives

Many implementation variants can be used to achieve the functionality sketched in Figure 1. For example, Alice's presence in the FreeZone could be detected via face recognition. Computer vision could also be used to detect that she is alone, and appears to be interruptible. Alternatively, a device worn or carried by Alice (e.g., smartwatch or smartphone) could use deep personal knowledge of its owner and sensor inputs to dynamically indicate to the FreeZone whether whether she is opted-in. It could also indicate whether this is an opportune moment to deliver a serendipity trigger to Alice.

A key tradeoff in this design space is between privacy control and frictionless user experience that closely emulates real-life serendipity. For some individuals, the need to always carry or wear a device

in order to have serendipitous interactions may be annoying. For others, this may be a trivial requirement. This tradeoff extends to many aspects of the scenario in Figure 1. For example, that description uses a chime from a FreeZone to indicate the triggering of a new encounter. A different implementation could use Alice's and Bob's personal devices to do the triggering and acceptance. In that case, the large display would share live video and audio of the other party only after both parties have accepted a trigger. Some Pomme implementations may rely solely on opt-in by personal devices, and avoid face recognition completely. In many cases, pilot usage experience and feedback from the user community may be needed to identify the best choices for a specific deployment.

Socio-cultural and institutional norms may play significant roles in these optimal choices for a Pomme community. System-initiated triggering can lead to social awkwardness such as *snubbing*, in which one party eagerly accepts the trigger but the other party declines. We discuss this further in Section 6.1.

3.3 Diverse Workplace Arrangements

Figure 1 depicts a traditional pre-COVID work setting, in which Alice and Bob go to work at different buildings that are owned or leased by the same corporation. With the exception of a few rare visitors, almost everyone in those buildings works for the same corporation. Anyone that Alice or Bob encounters in person is already likely to be a colleague. Although compartmentalization between groups may still pose concerns about information leakage, Pomme privacy is likely to be most tractable in this setting.

As we emerge from COVID, a growing number of alternative workspace arrangments are being explored. For example, pubs have been used as workspaces in the UK [6]. Companies such as *WeWork* (http://wework.com), *Distil* (https://distilcoworking.space) and others offer short-term physical office space. A private office, open cubicle, or desk can be yours for hours, days, weeks or longer. People still "go to work," but their destination is only a short commute away to a work location in their own neighborhood. Shared spaces such as lobbies, elevators, kitchens, and lunch rooms are frequented by employees of different companies.

FreeZones can be especially valuable in such neighborhood work settings, because the serendipity they provide increases the cohesivenes of a dispersed organization. At the same time, the privacy bar is also raised. In the scenario of Figure 1, the people around Alice and Bob are more likely to be strangers than coworkers.

4 Architectural Considerations

To develop FreeZones, we can make use of existing technology building blocks. Specifically, we can build on mechanisms such as OpenFace on cloudlets for accurate face recognition [5]. More generally, when combined with edge computing, computer vision based on deep neural networks offers a solid foundation for situational awareness [30]. Commercial products such as Zoom deliver interactive video and audio at Internet bandwidths, and offer APIs for integration with external mechanisms. We can also leverage existing work on denaturing live video [38].

The decision on when to establish and terminate connections between FreeZones requires knowledge of which Pomme users are within a given set of FreeZones. There are at least two distinct sources for such knowledge, i.e. real-time face recognition or short-range wireless communication with a body-worn (e.g., wrist-watch) or hand-held (e.g., smartphone) device on the occupant. The tradeoffs between these approaches have been well explored in other ubicomp/IoT scenarios (e.g. [26]), but in either case low-latency identification is critical as users may only be within a FreeZone for a very short period of time. Prior to establishing a connection Pomme must also understand the preferences and relationships between users to know whether a connection is indeed appropriate.

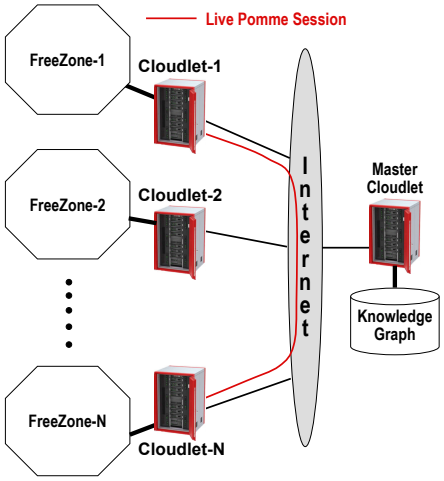

Figure 2. Pomme System Architecture

Figure 2 illustrates a plausible approach to integrating these building blocks into a working system. We do not claim that this design is optimal or unique, but present it only to ground our discussion and to provide context. Most importantly, a FreeZone has to track reality fast enough to *denature* an ongoing Pomme video and audio session. As described in earlier work [11, 35, 38], denaturing is the process of modifying a sensor stream to preserve privacy in accordance with a specified external policy. In the example of Figure 1, the moving blobs seen by Alice and Bob are the results of denaturing. The elimination of all voices except their own is another form of denaturing.

Denaturing imposes severe bandwidth and processing demands, while the user experience of live interaction imposes tight latency constraints. Since each 4K camera generates a roughly 32 Mbps video stream, a FreeZone with multiple cameras can easily generate upwards of 100 Mbps continuously. Other sensor streams such as audio will add to this. This bandwidth demand is only scalable today via edge computing, in which processing is done on a *cloudlet* that is the middle tier of a 3-tier compute hierarchy (IoT device – cloudlet – cloud) [31, 32]. In Figure 2, physical dispersion is bounded by the largest tolerable end-to-end latency between FreeZones. The bandwidth between a FreeZone and its cloudlet needs to be high enough to sustain continuous video transmission from all the high-resolution cameras in that FreeZone. The bandwidth between cloudlets, and the processing demands of multi-tenant denaturing, determine the maximum number of concurrent Pomme sessions.

Each FreeZone cloudlet can share its situational awareness with the Master Cloudlet shown in Figure 2. The bandwidth demand for this is modest, most likely just a few tens or hundreds of kbps, and scales with the number of occupants detected in that FreeZone. However, the latency has to be low, preferably no more than 100 ms. If a cloud data center is located close enough to the FreeZones, it could implement the Master Cloudlet functionality.

The shared real-time situational awareness across FreeZone and Master cloudlets represents the control plane of Pomme. As shown in Figure 2, the Master Cloudlet has access to a Knowledge Graph that represents Pomme's understanding of the social relationships among users. In this paper, we do not elaborate on exactly what the contents of the Knowledge Graph are, and how the static information from it is combined with situational awareness. These are clearly crucial topics for any real implementation of Pomme to address, and are likely to be the source of important research questions. Further, this decision process represents power, and may be the source of considerable social angst. An unscrupulous Pomme service could market access to high-value individuals via serendipity. Conversely, it could punish individuals by rarely delivering valuable triggers to them. Establishing trust in the fairness of the triggering process will be crucial to the acceptance of Pomme by a community. How best to do this is an open question at this time.

From time to time, based on this unspecified decision process, the Master Cloudlet initiates a Pomme session between two specific individuals in two different FreeZones. Such a session can be viewed as part of Pomme's data plane. The chimes that Alice and Bob hear in Figure 1 indicate the creation of such a session.

By construction, denaturing on a Pomme session removes all video and audio data that might leak privacy for anyone other than its participants. Pomme sessions are terminated when there is no further interaction on them. In the example of Figure 1, this happens some time after Alice leaves. On a different occasion, Alice or Bob may ignore the chime because they do not wish to interact. The other party can see that the invitation is being ignored, and can try to get attention by waving or saying something. But the uninterested party can simply ignore all of these efforts and depart. The fruitless Pomme session will be garbage collected later.

5 Workplace Stakeholders

Typical formulations of privacy problems involve two parties: a user and a service provider. In contrast, Pomme deployments in workplaces typically involve multiple stakeholders:

- *Physical Space Owners:* FreeZones inherently require physical spaces in which to be deployed. In practical terms this means that the owner (or operator) of the space must agree to the physical installation of hardware (cameras, displays etc.) and the use of the FreeZone technology in the space.
- *FreeZone System Owners:* While a FreeZone owner is typically the same as the owner of the physical space in which it is deployed, that may sometimes not be the case. For example, a company may offer FreeZones as a service to building owners, in order to increase the rental value of the building. Parallels exist with technologies such as WiFi or digital signage in which the owner/operator of the system is distinct from the owner of the space in which they are deployed.
- *Employers and Employees:* In this paper we have focused on the use of FreeZones in work-related scenarios. This means that the decision to engage with FreeZones is likely to live with employers and in particular their legal, HR and IT departments. Where an organization decides to embrace FreeZones it may be possible for individual employees to opt out. In that case, one key issue will be the level of control offered to employees (as opposed to their employers).
- *Bystanders:* In contrast to many other technologies, FreeZones also have privacy consequences for bystanders who simply happen to be within a FreeZone.

The simplest scenario is one in which FreeZones are deployed within a single organization, on premises they own and behind physical access control systems that ensure only employees are present. In this case, the number of stakeholders and complexity of interactions are significantly reduced. However, in the more likely scenario in which FreeZones emerge in a wide range of workplace settings, resolving issues relating to differing organizational policies or national legislation become much more complex.

6 Managing Privacy

A service like Pomme requires careful consideration of privacy issues. We start with the premise that total avoidance of all undesirable encounters is not the goal. This strategy is not even achieved in the physical world today. A fairer criterion is that (a) Pomme's positive consequences far outweigh its negatives, and (b) its privacy provisions inspire trust.

A system like Pomme will have some traditional privacy threats. One such threat to Pomme users is that the service itself could obtain a great deal of personal information, collected for the purpose of smartly generating encounter suggestions. Plus, this information could in turn be shared or sold to other companies, or retained for excessively long periods of time. In addition, a remote serendipity service surfaces other less evident privacy issues. Pomme users could potentially have their conversation with a remote colleague overheard by a broader set of people than in a typical hallway conversation. Importantly, bystanders who walk through FreeZones, and are not signed up with Pomme, can have their privacy impacted if they are captured on video or audio streaming. Also due to the involvement of multiple parties, there is a chance of privacy leakage due to inconsistent privacy policies across organizations.

In the discussion below, we outline options and trade-offs to these issues. Due to lack of space, we do not address important and related issues such as security and trustworthiness of the physical space owners (assumed herein).

6.1 Privacy for enrolled users

To give users agency over the collection of personal data, careful attention needs to be given to the issue of consent at multiple levels. Clearly, general consent to use the system would occur at initial signup. If an enterprise offers its employees portable devices with such a service preloaded, then it should be configured by default to an OPT-OUT setting. Beyond having the ability to opt-out of the service, it is preferable to allow users to opt-out of specific encounters. For example, in Figure 1, Alice could have set her wristwatch to OPT-OUT for an afternoon and would have completely avoided the remote encounter with Bob. In contrast, there is nothing Alice can do to avoid the encounter if it occurs in person: e.g., Bob visits Alice's work site, and runs into her at the elevator. Thus remote serendipity can be more controllable than in-person serendipity.

More fine-grained consent can be obtained through the use of privacy controls that allow users to set privacy preferences. There is a clear privacy trade-off pertaining to the quality of triggering decisions that Pomme makes versus the amount of personal data gathered. The timing, frequency, and perceived value of triggers

is highly dependent on the amount of contextual knowledge that Pomme has about the users involved in the trigger. User-specific sources of knowledge are needed to wisely select a person with whom to propose an encounter. Alice's contact list is one source of potential people she might enjoy meeting in a serendipitous encounter. Similarly her calendar might indicate times she would be open to receiving suggestions. Her Zoom and phone call logs can help ensure that triggers are not attempted with people that Alice has talked to recently. Alice can further express her privacy preferences by indicating in advance which time-of-day or location she is open to suggestions for encounters. The use of all such data sources should be in privacy controls managed by the user. These privacy controls effectively obtain consent on a per-data-stream, or per-sensor level. Other sources of knowledge might include the org chart of Alice's employer, which might suggest people working on projects of relevance to Alice. Both consent from the employer and consent from Alice may be needed in this case.

It would be important to also be mindful of how users are notified, or triggered, when Pomme suggests an encounter. One approach is that captured in Figure 1, where Alice and Bob heard chimes and saw the other in the display without giving explicit permission for this encounter. This emulates real life serendipity. It is frictionless but affords a user less control if they would have preferred to deny the encounter. A more privacy-controlled approach could first prompt Alice on her device about the possible encounter with Bob, and vice versa. For a user to make an informed decision about accepting a trigger, the identity of the other party has to be revealed. This can result in social awkwardness when one party declines, while the other accepts. In real life, plausible deniability can be used to avoid this situation. For example, phone calls can be screened without revealing one's own state of availability. Similarly, an undesirable in-person encounter can be terminated by saying that you have to rush off to a meeting. Such excuses are not credible with a smart triggering mechanism.

Although spontaneous face-to-face human encounters seem effortless, they make use of substantial non-verbal signalling and are often more structured than apparent. For example, Kendon and Ferber's analysis [21] reveals a multi-stage protocol for typical human greetings, which involves negotiation and possible avoidance of the encounter at low social cost. Attempting to precisely mirror human protocols like this in technology-mediated serendipity is a difficult task, particularly since the human protocols may shift in order to adapt to this new setting. Usage experience is needed to strike the right balance between privacy and effortless serendipity, and to possibly develop a new protocol that can be learned for this new class of interactions.

The issue of data sharing beyond the first party service should be clearly articulated in Pomme's privacy policy. Clearly, the policy would also cover the standard elements of privacy policies[1, 2]. These include statements about what data is collected, the purpose and how it is used, with whom data is shared or sold, provisions for EU members under the GDPR, and how a user can obtain a copy of any data collected about him or her [3].

The use of edge computing and cloudlets enables Pomme to handle data retention issues better than many services today. The raw sensor data used by Pomme for situational awareness does not ever need to leave the cloudlet associated with the FreeZone. (Indeed a key privacy benefit of using cloudlets is that they make it easier to follow the data minimisation principle.) Even on that cloudlet, it only has to be retained for a few seconds, while the cloudlet determines if this is a currently opted-in user. The knowledge that a currently opted-in user was at this FreeZone at this point in time is shared with the master cloudlet, which then determines whether to issue a trigger. To guide future suggestions, Pomme only remembers the history of suggestion attempts: timestamp, users involved, FreeZones where those users were located, and whether the trigger was successful. If a user is currently opted-out, even this limited knowledge is not shared with the master cloudlet.

Finally, users should have no expectation of confidentiality in their Pomme conversations. As with a physical encounter, anyone near them can hear what is being said, and see facial expressions and body language. If greater privacy is desired, a private meeting can be set up at the encounter.

6.2 Bystander Privacy

In the scenarios we describe, there could indeed be other people in our FreeZones that have not signed up for Pomme service. Bystander privacy arises when there are cameras or microphones in shared spaces. Privacy is a challenge as bystanders do not engage in any consent process. Bystander privacy issues occur in a variety of services such as IoT applications in smart homes [41]), public face recognition systems [29], and services that assist the visually impaired [4]. Example solutions proposed include approaches to deceive [34, 39] or obfuscate identification in facial recognition [28]. Some of these solutions come with the need to wear additional equipment, such as specialized goggles [29] that may be impractical. We believe that denaturing [38] offers a compelling alternative for the bystander privacy problem.

We imagine two ways address to bystander privacy in always-on cameras. In one scenario, the database used in face recognition would only identify users enrolled in Pomme. All other users would be labeled as "stranger." Denaturing of a live video stream would make strangers appear as solid blobs in Pomme sessions. By construction, Pomme never retains any record of the identity of strangers. Another version of denaturing, is to have the entire video image be blobbed out except for the person using Pomme. An advantage of this latter case, is that Pomme is less dependent upon the speed of denaturing; however this may lead to a less natural encounter as the Pomme participants do not see the FreeZone. In both these scenarios, no bystanders are visible to the remote person involved in a serendipitous encounter. The key to using such a video denaturing capability is speed. Wang et al's work on denaturing live video [38] confirms that it is possible to keep up with a 30 frames per second HD video stream, even when using modest cloudlet hardware. Wang et al also describe optimizations to reduce the chances of visual privacy leaks. Denaturing of audio has also been shown to be feasible [10].

Emerging legal restrictions on face recognition software pose a challenge for video denaturing. More than two dozen cities, starting with San Francisco in 2019, have banned the use of face recognition software [36]. A more workable approach would be to use enrollment in Pomme as implicit permission to include that user in the database used for face recognition. There would be no information about any other users in the database, and they would be labeled as "stranger" by face recognition. As further reassurance, video is discarded immediately after use — there is no long-term retention.

6.3 Privacy across multiple parties

The serendipity service we describe in a shared office space, such as WeWork, would involve many stakeholders each of which would need a privacy policy. One possible instantiation might have WeWork as the physical space owner, Pomme as the FreeZone system owner, and one or more employers each offering Pomme to their employees. WeWork would likely be responsible for posting signs in Freezones indicating the presence of cameras. Pomme's privacy policy would follow all the elements outlined above. However, the employer would also have to explain their contractual agreement with Pomme, that could influence, for example, whether calendar and contact information can be shared with Pomme. Contradictions could arise, such as how notice about the always on cameras is delivered, how long data is retained, and for default settings. Although beyond the scope of this paper, we highlight that any such contradictions would have to be resolved.

7 Towards a Real Implementation

In this paper, we have shown how technology-mediated serendipity can be achieved in an organization that embraces remote work. The privacy challenges are tractable if "remote work" is not interpreted as working from home, but as working from a neighborhood facility using a WeWork-like model. This may emerge as a sweet spot for work modality in the post-COVID era, as evidenced by the October 2021 announcement of *Saksworks* to repurpose obsolete Saks department stores as neighborhood workspaces [7].

While we have focused on workplaces, Pomme could be also valuable in other settings where the value of social contact far outweighs privacy concerns. For example, it could provide spontaneous encounters for residents of residential care settings who have limited freedom of mobility due to physical or cognitive decline. It could also bring the benefits of serendipity to visually-impaired and hearing-impaired people, who are often unable to benefit from real-life serendipity.

We have focused on privacy in this paper because it is the biggest obstacle to real-world deployment of Pomme-like systems. However, there are also other challenges that will need to be addressed in areas such as usability, scalability, and efficacy. The last of these poses the biggest unanswered question. If we create a Pomme-like system, will people use it? Will they reap the expected benefits from it? Will they accept the privacy compromises? Or will Pomme fall far short of expectations? The only way to answer these questions is by implementing Pomme and deploying it for real use. Only such a deployment can provide the validation, hands-on experience, and insights to advance our vision.

Throughout history, serendipity has played an out-sized role in stimulating breakthrough innovations and creative insights. A year and a half of COVID-induced remote work has helped us to realize what we lose by giving up chance encounters with colleagues. An extended period of remote work has the potential to result in siloing of teams, as documented by the Microsoft study mentioned earlier [40]. The message is clear: we cannot afford to sacrifice serendipity as we experiment with diverse modalities of work. Some mechanism along the lines described in this paper will need to be an essential component of our future workplace arrangements.

Acknowledgements

We thank Suman Banerjee, our shepherd, and the anonymous reviewers for their guidance in improving this paper. We also thank Jim Blakley, Wei Gao, Scott Hudson, Bobby Klatzky, Babu Pillai, and Junjue Wang for reviewing early versions of this paper and giving us valuable feedback for improving it. This work was supported in part by the National Science Foundation (NSF) under grant number CNS-2106862, and in part by the EPSRC under the auspices of grant EP/T022574/1. Additional support was provided by gifts from Intel, Vodafone, Deutsche Telekom, CableLabs, Crown Castle, InterDigital, Seagate, Microsoft, VMware and the Conklin Kistler family fund. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the view(s) of their employers or the above funding sources.

References

[1] Privacy Policy Guidance. https://developers.google.com/assistant/console/policies/privacy-policy-guide. [Online; accessed 1-Oct-2021].

[2] Sample Privacy Policy Template. https://www.privacypolicies.com/blog/privacy-policy-template/. [Online; accessed 1-Oct-2021].

[3] What Privacy Information Should We Provide. https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/the-right-to-be-informed/what-privacy-information-should-we-provide/. [Online; accessed 5-Oct-2021].

[4] T. Ahmed, A. Kapadia, V. Potluri, and M. Swaminathan. Up to a Limit? Privacy Concerns of Bystanders and Their Willingness to Share Additional Information with Visually Impaired Users of Assistive Technologies. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2, Sept. 2018.

[5] B. Amos, B. Ludwicznik, and M. Satyanarayanan. OpenFace: A general-purpose face recognition library with mobile applications. Technical Report CMU-CS-16-118, School of Computer Science, Carnegie Mellon University, June 2016.

[6] J. Bartholomew. What's the Most Productive Workspace? Might Be the Pub, November 2 2020. https://www.wsj.com/articles/whats-the-most-productive-workspace-might-be-the-pub-11604350266.

[7] G. Bellafante. Meet Me in My Office, in Men's Underwear on 5. *New York Times*, October 15 2021. https://www.nytimes.com/2021/10/15/nyregion/saksworks-coworking.html.

[8] M. Boyle, C. Neustaedter, and S. Greenberg. *Privacy Factors in Video-Based Media Spaces*, pages 97–122. Springer London, 2009.

[9] H. Brignull and Y. Rogers. Enticing People to Interact with Large Public Displays in Public Spaces. In *Proceedings of the IFIP International Conference on Human-Computer Interaction*, pages 17–24, September 2003.

[10] R. R. Choudhury. Earable Computing: A New Area to Think About. In *Proceedings of HotMobile 2021*, 2021.

[11] N. Davies, N. Taft, M. Satyanarayanan, S. Clinch, and B. Amos. Privacy Mediators: Helping IoT Cross the Chasm. In *Proc. of ACM HotMobile 2016*, St. Augustine, FL, February 2016.

[12] R. S. Fish, R. E. Kraut, R. W. Root, and R. E. Rice. Evaluating Video as a Technology for Informal Communication. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1992.

[13] B. Friedman, P. H. Kahn, J. Hagman, R. L. Severson, and B. Gill. The Watcher and the Watched: Social Judgments about Privacy in a Public Place. *Journal of Human-Computer Interaction*, 21(2):235–272, May 2008.

[14] K. Galloway and S. Rabinowitz. Hole-In-Space. http://www.ecafe.com/getty/HIS/ [Last accessed: October 2012], 1980.

[15] W. Gaver, T. Moran, A. MacLean, L. Lovstrand, P. Dourish, K. Carter, and W. Buxton. Realizing a Video Environment: EuroPARC's RAVE System. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1992.

[16] S. Greenberg and M. Rounding. The Notification Collage: Posting Information to Public and Personal Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2001.

[17] S. E. Hudson and I. Smith. Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work*, Boston, MA, 1996.

[18] W. Isaacson. *Steve Jobs.* Simon & Schuster, 2011.

[19] D. Kagan, G. F. Alpert, and M. Fire. Zooming Into Video Conferencing Privacy and Security Threats, 2020.

[20] M. Kamprath and T. Henike. Serendipity and Innovation. In *The Routledge Companion to Innovation Management*, chapter 17, pages 343–360. Taylor and Francis Group, 2019.

[21] A. Kendon and A. Ferber. A description of some human greetings. In R. Michael and J. Crooks, editors, *Comparative Ecology and Behavior of Primates*, pages 591–668. Academic Press, 1973.

[22] R. Koch and G. Lockwood. *Superconnect.* W.W. Norton and Co., 2010.

[23] S. Lin, A. Ryabtsev, S. Sengupta, B. Curless, S. Seitz, and I. Kemelmacher-Shlizerman. Real-Time High-Resolution Background Matting. *arXiv*, pages arXiv–2012, 2020.

[24] Merriam-Webster. Dictionary and Thesaurus. https://www.merriam-webster.com/. Last accessed on 2021-07-21.

[25] J. R. Methot, A. S. Gabriel, P. Downes, and E. Rosado-Solomon. Remote Workers Need Small Talk, Too. *Harvard Business Review*, March 2021. https://hbr.org/

2021/03/remote-workers-need-small-talk-too.

[26] M. Mikusz, P. Shaw, N. Davies, P. Nurmi, S. Clinch, L. Trotter, I. Elhart, M. Langheinrich, and A. Friday. A longitudinal study of pervasive display personalisation. *ACM Transactions on Computer-Human Interaction*, 28(1), Jan. 2021.

[27] A. Pentland. *Social Physics: How Good Ideas Spread-The Lessons from a New Science*. The Penguin Press, 2014.

[28] A. J. Perez, S. Zeadally, L. Y. M. Garcia, J. A. Mouloud, and S. Griffith. Facepet: Enhancing bystandersâĂŹ facial privacy with smart wearables/internet of things. *Electronics*, 2018.

[29] A. J. Perez, S. Zeadally, S. Griffith, L. Y. M. Garcia, and J. A. Mouloud. A User Study of a Wearable System to Enhance Bystanders' Facial Privacy. *IoT*, 1(2):198–217, 2020.

[30] M. Satyanarayanan. Edge Computing for Situational Awareness. In *Proceedings of the 23rd IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN 2017)*, Osaka, Japan, June 2017.

[31] M. Satyanarayanan. The Emergence of Edge Computing. *IEEE Computer*, 50(1), 2017.

[32] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies. The Case for VM-Based Cloudlets in Mobile Computing. *IEEE Pervasive Computing*, 8(4), 2009.

[33] N. Sawhney, S. Wheeler, and C. Schmandt. Aware Community Portals: Shared Information Appliances for Transitional Spaces. *Personal Ubiquitous Computing*, 5(1):66–70, Jan. 2001.

[34] M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.

[35] P. Simoens, Y. Xiao, P. Pillai, Z. Chen, K. Ha, and M. Satyanarayanan. Scalable Crowd-Sourcing of Video from Mobile Devices. In *Proceedings of the 11th International Conference on Mobile Systems, Applications, and Services (MobiSys 2013)*, Taipei, Taiwan, June 2013.

[36] T. Simonite. Face Recognition Is Being Banned — but It's Still Everywhere. *Wired*, December 2021. https://www.wired.com/story/face-recognition-banned-but-everywhere/.

[37] R. Stults. Media Space. Technical report, Xerox PARC Tech. Report, 1986.

[38] J. Wang, B. Amos, A. Das, P. Pillai, N. Sadeh, and M. Satyanarayanan. A Scalable and Privacy-Aware IoT Service for Live Video Analytics. In *Proceedings of ACM Multimedia Systems*, Taipei, Taiwan, June 2017.

[39] T. Yamada, S. Gohshi, and I. Echizen. Privacy visor: Method based on light absorbing and reflecting properties for preventing face image detection. *2013 IEEE International Conference on Systems, Man, and Cybernetics*, pages 1572–1577, 2013.

[40] L. Yang, D. Holtz, S. Jaffe, S. Suri, S. Sinha, J. Weston, C. Joyce, N. Shah, K. Sherman, B. Hecht, and J. Teevan. The effects of remote work on collaboration among information workers. *Nature Human Behavior*, September 2021. https://doi.org/10.1038/s41562-021-01196-4.

[41] Y. Yao, J. R. Basdeo, O. R. Mcdonough, and Y. Wang. Privacy Perceptions and Designs of Bystanders in Smart Homes. In *Proceedings of the ACM on Human-Computer Interaction*, 2019.