Supporting Augmented Reality: Looking Beyond Performance

Lemuel Soh
University of California, Los Angeles
lemuel@cs.ucla.edu

Jeff Burke
University of California, Los Angeles
jburke@remap.ucla.edu

Lixia Zhang
University of California, Los Angeles
lixia@cs.ucla.edu

ABSTRACT

Recent years have witnessed a surge in augmented reality (AR) applications in various markets and verticals, together with emerging toolkits and platforms to support new developments. However, the vision of a pervasive augmented reality held by many still seems a distance away. Notwithstanding the many ongoing efforts to tackle AR performance challenges, we argue that much attention is needed to other research areas including network architecture, security, privacy, and the development of business cases. Similar to the Web, existing AR applications are built upon TCP/IP protocol stack and rely on cloud computation. To enable pervasive AR applications, we believe that new computing paradigms, new approaches to network communications, and new business models need to be explored. Edge computing paradigms, which utilize performance advantage of server class hardware within physical vicinity, could achieve the required low latency while protecting user privacy. We further argue that Named Data Networking (NDN), a proposed new internet architecture, can be an enabler for pervasive AR by supporting local resource discovery, offering built-in communication security, and enabling experimentation with new business models. We hope that this position paper spurs greater thinking beyond performance improvements to push AR forward.

CCS Concepts

• **Networks** → **Network architectures**; *Edge computing*;
• **Networks** → **Network protocols**; • **Human-centered computing** → **Ubiquitous and mobile computing**

Keywords

Pervasive Augmented Reality; Named Data Networking; Information Centric Networking

ACM Reference format:

Lemuel Soh, Jeff Burke, and Lixia Zhang. 2018. Supporting Augmented Reality: Looking Beyond Performance. In *Proceedings of VR/AR Network '18, Budapest, Hungary, August 24, 2018*, 6 pages. <https://doi.org/10.1145/3229625.3229627>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

VR/AR Network '18, August 24, 2018, Budapest, Hungary

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5913-9/18/08...\$15.00

<https://doi.org/10.1145/3229625.3229627>

1. MORE ACCESSIBLE INFORMATION VIA AUGMENTED REALITY

The Internet heralded a wave of mass communication not seen since the invention of Gutenberg's printing press. Unlike books, which allowed communication through static words or pictures, the Internet enabled real-time interaction, rich ways of expressing oneself, and drastically lowered the barrier in accessing content. Augmented Reality (AR), representing a new wave of Internet applications, has the potential to once again fundamentally change the way we interact with each other in physical and cyber spaces. By customizing content according to one's given context in real-time, AR could help enrich users' experience with the most relevant content, just at the time it is needed.

In this position paper we follow the general vision laid out by Grubert *et al.* for pervasive augmented reality as a "continuous, omnipresent, and universal augmented interface to information in the physical world" [8]. While appreciating the importance of performance enhancements to achieve this objective, we note that several other important questions related to computing paradigm, network architecture, security, privacy, and business models must also be addressed for pervasive AR to become a reality. We begin with describing our understanding of how AR applications are being prototyped today over TCP/IP network protocol stack. Next, we briefly examine Web applications and its Hypertext Transfer Protocol (HTTP), draw parallels to how AR is implemented today, and identify potential limitations. We then use an example of pervasive AR to highlight the functions needed and suggest new design patterns. Edge computing paradigms are important in accomplishing pervasive AR, for performance and privacy reasons. To support edge computing, we introduce Named Data Networking (NDN), a proposed future internet architecture, and discuss how NDN could address the requirements of resource discovery, trust management, multicast support for context-content exchange, and experimentation with new business and user experience models.

2. TODAY'S AR APPLICATIONS

Caudell, a Boeing research scientist, coined the term "Augmented Reality" (AR) in the early 1990's while trying to help workers assemble Boeing's aircraft more efficiently with less error. Following the general definition of AR given by Azuma [1] as combining real and virtual content, being interactive in real time, and registered in 3D, Billinghurst *et al.* [2] details many use cases in 2015. AR apps have sprung up in domains ranging from medicine to marketing, and even to building design. Such use cases are no doubt reflective of the need of different industries, and different communities.

One AR use case is Wayfinding, an American Airlines AR app (https://www.youtube.com/watch?v=X9PpUTUX_Kk).

Wayfinding helps passengers find anything that they might want at an airport terminal. Upon arriving at the airport, indoor positioning systems detect passengers' position and offer information on their mobile phone viewing stream. They are first directed to the security

checkpoint with the shortest wait time. Upon clearing security, Wayfinding directs passengers to the updated gate at the appropriate boarding time, providing the current occupancy and wait times of restaurants and shops along the way. The information about occupancy, waiting time and locations could be supplied by indoor Internet of Things (IoT) devices built into the airport environment, or inferred from other users' view streaming.

Generally speaking, an AR application needs *context* data to provide a user with customized content in a particular physical environment under given circumstances. As described by Burke [3], the context may include spatiotemporal information such as the user's location and time, field of view, preferences and interactive choices. For the Wayfinding application, the context includes the passenger location, what the passenger is seeing on the mobile device, mobile device orientation and up-to-date status of how busy different terminal locations are, provided through external service. Resulting customized content includes overlaying directions and highlighting the expected wait times for nearby amenities on the mobile view screen. AR can be most effective if an ongoing, streaming relationship exists between context from the user and the resulting customized content.

Today's AR apps, running over TCP/IP and often HTTP, provide the AR experience through client connectivity to cloud services. By collecting and processing all context information in the cloud, the relevant customized content is determined and sent to the user. For example, in apps like Wayfinding, based on users' location and mobile view, the cloud can provide sensor-collected location wait times in their view. Security is maintained by ensuring all devices are connected to trusted cloud servers and secured by Transport Layer Security (TLS). Such connection-centric security assumes that all passenger context, and airport sensor information are from trusted parties based on authenticated IP tunnels.

3. FOLLOWING WEB'S FOOTSTEP

Although the Internet's core network protocol suite (TCP/IP) was published in 1981, it was not until the invention of the World Wide Web (WWW) in the early 1990's that cyberspace information became widely accessible [26]. The Web protocol, HTTP, introduced a request/response communication semantics, enabling many Web applications not envisioned by Tim Berners-Lee to be built using HTTP. The app developers utilized the available HTTP to build apps utilizing client-server connections over TCP/IP protocol stack, without deviating too much from this pattern. Application semantics were used, and content interconnections were made at the application layer.

Since HTTP was built over TCP/IP, limitations of the TCP/IP protocol stack were inherited as well. Applications generally have to make use of point-to-point connections, even if the application semantics are different. IP's point-to-point communication model showed its limitation as the Web grew. Traffic growth led to content server overload. The one-to-many application semantics of content distribution thus led to the deployment of Content Delivery Networks (CDN). CDN servers are placed in various local places to deliver content with lower latency, relying on Domain Name Service (DNS) and other redirection techniques to map user requests to CDN nodes in users' proximity.

Figure 1. Cloud Computing Paradigm supported by the TCP/IP network architecture. AR apps employ centralized connectivity to the cloud, and connection-based security, for resource discovery, even locally. Red dots: known IP endpoints. Dotted lines: IoT sensor data upload. Solid lines: Request/Response data exchange with cloud.

3.1 TCP/IP, Cloud, and AR of the Future

Following web paradigms and protocols, today's AR applications are built on TCP/IP*, inheriting both its affordances and limitations. Though the existing AR apps bring utility in specific domains, their reliance on cloud service and connection-centric security may limit the potentials of AR.

One such consequence is that, because of connection-centric security and cloud-based business models, most AR apps today are vertically-integrated applications, termed as *stovepipe* systems [20]. Each such stovepipe application is designed to solve a specific problem and/or keep the user within a given content ecosystem. This design pattern leads to multiple AR apps individually requesting the same context data, such as user view, for different uses. For example, the Wayfinding app needs user view for overlaying a path to the boarding gate while another AR app might need user view to identify faces of friends in the airport crowd.

Another consequence is that latency-sensitive apps demanding computation power and memory resources cannot be built satisfactorily using cloud services, which can be many network hops away from user locations [6]. Measurements made by Ha et. al. [9] indicate that edge services can shorten the service latency by 80 to 200ms on average as compared to the cloud service.

Lastly, it is difficult to discover local resources in a secure way using TCP/IP. This is very relevant to AR, which is about content customized for local context. Multicast DNS has been developed to assist local resource discovery, however it does not address the security challenges in an environment with unverified mobile devices. In today's network practice, both rendezvous and security functions are provided by the cloud, in addition to computation service.

App developers usually follow available patterns, and the dominant pattern available today relies on client-server connections over IP's point-to-point packet delivery, interconnecting contents at the *application layer*. This leads to stovepipe AR apps running on

*AR implementations utilizing other transport protocols such as UDP or QUIC are not known.

cloud computing. These apps are restricted by network delays, lack of data sharing across apps, and a high barrier to entry from the requirement of offering the whole stovepipe to users. Relieving those restrictions could lead to the next wave of AR innovation and experimentation.

4. PERVASIVE AUGMENTED REALITY

As described by Grubert *et. al.*, pervasive AR is a “continuous and pervasive user interface that augments the physical world with digital information registered in 3D, while being aware of and responsive to the user context” [8]. With continuous use, pervasive AR should be primarily guided by user context to minimize the human intervention needed. Content and service provision, whether indoors or outdoors, should be provided seamlessly without the user needing to switch between devices or service providers. Through context-awareness, an AR app can augment the user’s experience with customized content with minimal input.

4.1 Carlos and Sally using Pervasive AR

To illustrate the vision of pervasive AR, we describe an example scenario which takes inspiration from the pervasive AR scenarios mentioned by Grubert *et. al.* [8] and the life of Sal mentioned by Weiser [21].

Being alerted by his voice assistant through his earphones that it was time to travel to meet Sally, Carlos closes his laptop and heads for the car. As soon as he steps into his self-driving car, his smart glasses overlay his field of view partially with a map of his destination and estimated time of arrival. At the end of the journey, Carlos steps out into the busy Grand Central Market of Los Angeles. Within seconds, a few faces are identified among the large crowd as his friends, one of whom is Sally. On meeting up with Sally, both Sally and Carlos’ glasses display walking directions overlaid on their visual field of the ground, and they make their way towards a local café fancied by them. Right at this time, Carlos and Sally both receive a simultaneous early warning of a major earthquake about to happen. Spots that are most structurally strong are highlighted to them within the building, and they run toward nearest spots. As the power goes out and parts of the building collapse, Carlos and Sally get separated, trapped in separate parts of the building, although their glasses exchange data via Bluetooth and notify them that neither is hurt. Previously invisible IoT sensors connect with their glasses (as programmed for an emergency) and are now visually displayed as places where first responders can access video and audio feeds. Both Carlos and Sally wait trapped under the rubble for the first responders to arrive.

The above scenario describes pervasive AR. Services are delivered by taking advantage of continually changing context, from Carlos leaving his home, to getting into the car, to meeting Sally, experiencing the earthquake, and reconnecting with her. The AR user interface of the smart glasses is context-aware and able to augment the physical world with relevant content. This content is customized based on context and delivered without user’s input. When the earthquake occurs, pervasive AR is able to function in spite of failed infrastructure and compromised local resources.

4.2 Requirements of Pervasive AR

A key technical requirement of this new generation of AR applications is fast information response time that is invariant as a function of the bandwidth demanded and infrastructure availability [20]. Resource discovery services, resource distribution services, intrinsic security, seamless mobility and scalable content distribution are also essential for such applications. This section

elaborates on some of the important requirements needed to achieve a pervasive AR experience as illustrated above.

Computing Resources within Physical Proximity: Computationally intensive tasks such as face recognition (to recognize Sally) need to be offloaded. Such offloading requires computing and memory resources that can be accessed securely on the edge in the busy market.

Resource Discovery Services in Unknown Environments: Secure local resource discovery is crucial in every aspect of Carlos’ and Sally’s pervasive AR experience. The display of personalized content on their smart glasses, such as friend recognition and walking directions, all required computational resources in their local environment. Furthermore, when the earthquake happened, and network infrastructure had been disrupted, Carlos’ smart glasses needed to discover what remaining resources were available for computation.

Intrinsic Trust and Security: After the earthquake, both Carlos and Sally had functioning smart glasses with Bluetooth. Though there was no surrounding network infrastructure or very low bandwidth, their smart glasses were able to verify the information that neither party was hurt.

Business Model Experimentation: In all the scenarios, both Carlos and Sally required computing, storage and content resources in physical proximity while at home, in commute, or at Grand Central Market. Cloud service providers, Internet Service Providers (ISPs), local community infrastructure, or even surrounding mobile devices may all provide such resources, pending appropriate business models to provide viability and incentives. For example, local community infrastructure requires maintenance that may be paid by its local users or passing commuters. However, the means of locally authenticating users and collecting fees do not currently exist, and consumers’ willingness to pay for resources provided by local infrastructure do not have an existing viable business model.

The AR scenario of disaster recovery (as an epilogue to our story of Carlos and Sally) would face all the challenges mentioned above, in particular the data-centric security required for communication among first responders using a variety of communication media. In addition, disasters such as a large-scale earthquake are likely to disrupt the electrical grid and damage network infrastructure. Under such adverse conditions, AR-based rescue applications could help first responders not only locate people trapped in collapsed buildings, but also assess their conditions through video and images to take best informed actions. However, AR applications likely have to run over edge devices and peer-to-peer networking for resilient information discovery and communication, together with stringent security measure. Fast response time and coordination across administrative boundaries are key to saving lives. We will discuss running such AR applications over the edge using peer-to-peer networking in the next two sections.

5. EDGE COMPUTING FOR PERFORMANCE AND PRIVACY

Pervasive AR is not only latency-sensitive, but also requires significant computational power and memory for many of its tasks such as face recognition. Furthermore, though mobile device hardware is improving, there is a persistent performance advantage of server hardware over the typical mobile [7], making task offloading essential in order to improve response time and reduce battery consumption. If offloading is performed, owing to the variable and sometimes high latency of distant cloud services [12],

proximity of the server matters greatly. As reported by Li *et al.* [12] and Satyanarayanan [16], roundtrip times (RTTs) to clouds are too long for satisfactory user experience.

Edge computing, as defined by Shi *et al.*, refers to the enabling technologies that allow computation to be performed at a network's edge, where the edge is along the path between data sources and the cloud servers [18]. In addition to providing computing, storage, and caching for mobile devices, the edge is also able to request and deliver services to the cloud. Satyanarayanan *et al.* argued that mobile computing and cloud computing are converging with cloudlets being an important architectural component [16]. Cloudlets are a specific manifestation of edge computing and can be thought of as the middle element of a design that includes mobile device, cloudlet, and cloud. It serves to "bring the cloud closer" and has shown promise of improved response times [22] and energy savings [9].

How to maintain user privacy within increasingly pervasive technologies is an important question to address for pervasive AR to bring good to society. Unfortunately, privacy is in worse shape today than at the turn of the century because of current practice of sending all user data to the cloud [6]. For pervasive AR, the collected user context information will be increasingly personal as more and more applications are built for all aspects of our life. Consequently, mobile devices are increasingly not just data consumers but content creators. How and where personal context and created content should be stored, protected, and disseminated will prove crucial to the adoption and success of pervasive AR.

Centralized data control by cloud providers could be detrimental to privacy. If, however, data is kept in the local vicinity where it is processed, the opportunities for massive breaches of privacy would be reduced since personal data is spread out over many different local providers of cloudlet services. Such local processing and storage would require new business models and will be explained briefly later. Together with possible legislation empowering individuals such as the personal data guardian [11] and expansion of the code of fair information practices [19], building AR support on edge computers can empower individuals to be in charge of their private data while improving app performance. Utilizing edge computing thus enables operating on local context (and content) locally, leading to both privacy and performance gains.

Despite much progress in edge computing, Satyanarayanan highlighted major challenges surrounding resource discovery, trust, and business models in a recent interview [6]. The concept of cyber-foraging [7] to discover the closest resources and establish trust cannot be easily achieved today without the cloud assistance. Furthermore, there is a widely open question regarding the business model for edge computing. The incentives of providing such resources must be well-thought out to spur deployment. We next examine a new internet architecture that could help address the challenges faced by edge computing when supporting pervasive AR.

6. ARCHITECTURAL SUPPORT FOR DISCOVERY, SECURITY AND BUSINESS

In this section, we investigate the use of Named Data Networking (NDN) as a means to address the above challenges in supporting pervasive AR over edge computing. NDN is a manifestation of the new information-centric networking paradigm [24]. From 10,000 feet, one could view the basic idea of NDN as shifting HTTP's request and response semantics to the network

Figure 2. Differences between TCP/IP and NDN

layer. NDN's requests (for named data) and responses (the data objects themselves) operate at network packet granularity. That is, requests are a single packet—one Interest, carrying the name of requested data; each Interest fetches one Data packet back. NDN forwards interest packets according to their names, and forwards data packets back to requesters by letting them reverse the paths of the corresponding interest packets. Some of other desired NDN properties include the feedback loop created by its packet granularity interest-data exchange, at every hop in the network. This enables feedback on congestion, overload, or failures. Built-in communication security is provided by signing and validating on all data packets. These features make NDN a solid architectural foundation, upon which we can build the future AR platforms and applications.

6.1 Resource Discovery

When entering a new environment, the user's AR mobile device needs to first discover relevant available resources, such as information, computing, and storage resources pertinent to its AR application. NDN enables information discovery at the network layer by using application-layer names to forward user interest packets [17]. Thus a user's AR application can discover available resources by requesting the named resources at the network layer. By allowing application resource discovery at the network layer, the mobile can learn about available data from surrounding IoT sensors, as well as computing and storage resources that are potentially provided by multiple parties, without relying on a mapping service. By removing the need for mapping between app name and IP addresses, NDN also simplifies the many-to-many communications required for efficient context-for-content exchanges with multiple providers simultaneously. However, such many-to-many data exchanges impose new security challenges, which we address below.

6.2 Establishing Security and Trust

Adding to the vibrant debate on the security and privacy implications of AR [15], NDN focuses on a different approach to security. Instead of securing data containers or communication channels, NDN secures data directly. Communications are secured by binding the name and content of each data packet through a cryptographic signature. Content confidentiality can be achieved by data encryption where key exchange is bootstrapped using per-packet signatures for authentication. The verifiability of individual data packets can be achieved through trust relations among different actors in the AR ecosystem, rather than relying on a channel back to trusted cloud servers [17]. Whether an application

Figure 3. Edge Computing Paradigm supported by NDN Network Architecture. AR apps use many-to-many connectivity with named data abstraction. Discovery of local edge resources enabled by requesting the resource directly via names. Security is intrinsic as signatures are enforced on all named data packets. Blue Cylinder: Edge resources including compute, memory, storage and content resources.

trusts a piece of received data is determined by the application's trust schema, which can leverage data naming to simplify verification [23]. This opens up many possibilities for establishing trust and accessing different AR services efficiently at a highly granular level.

NDN has been experimenting with various trust schemes that allow trust to be bootstrapped from entities trusted by users for *specific content types*. In our example scenario, this could mean that if Sally's application trusts location data signed by the Grand Central Market, and Carlos trusts location information signed by Sally, then Carlos's application could also choose to trust location information signed by the Market. NDN has also developed a number of supporting tools to enable usable security by automating the management of cryptographic keys for data signing and verification [23], as well as encryption/decryption [25]; the latter addresses data access control and privacy by ensuring that only Carlos can view Sally's packets and content.

6.3 Business Model Experimentation

Business models are not agnostic to network protocols. They are constrained by network design choices. Cloud computing is supported by the client-server application paradigm encouraged by TCP/IP's point-to-point communication model. Fueled by the economies of scale, cloud service has been the viable and dominant business model for distributed applications. Since the cloud is used for resource discovery and trust establishment, large companies with the resources to operate massive cloud computing centers dominate services. However, as we mentioned earlier, relying on cloud service for pervasive AR might result in unsatisfactory performance for local areas where a particular cloud provider is too far away. It is also likely to limit the development of diverse ecosystems of overlapping services that are all augmenting our

reality. For example, imagine Sally might want to switch to a different AR service provider for location services. This can be made possible using TCP/IP's communication paradigm if she always coordinates through a cloud rendezvous service. However, besides introducing delay, requiring connection to one or a small number of major rendezvous providers limits choice and flexibility.

It is expected that pervasive AR, with its innovative applications in local environments, will require new business models to sustain edge computing. NDN can facilitate the experimentation with different business models by enabling all players to interconnect their resources at network layer, making their services available through routing announcements, and letting end users choose through its request/reply communication using app names. Assuming the use of well-defined namespaces, an exploration of different solutions including cloudlets controlled by cloud provider, ISP-offered resources, user community resources, end user devices, or any combination of the above, can be effectively evaluated through market economics. A variety of different market scenarios may result. Users may subscribe to AR computation and storage services, in the same way that they subscribe to cellular, storage, content, mail, and other services. Vendors may also start selling personalized AR computational / storage boxes for people to use when they are performing sensitive and private AR tasks. If the past is any indication, advertisers will even support the computation if users are willing to receive targeted advertisements based upon their context.

In our example scenario, Carlos' computation resources can be provided by his car (personalized AR computational / storage box) while he is traveling. Once at Grand Central Market, his glasses can detect and switch to a free service provided by a cloudlet provider that has targeted advertising. However, once the earthquake occurs, his cloudlet is destroyed, and the glasses switch to a global wireless cloud service provider with a paid subscription that has relatively low bandwidth. Since TCP/IP requires a connection point, choices are made without reference to which local services can provide the best performance in real-time for the personalized content requested. NDN is thus more supportive of a variety of AR business models by enabling and encouraging all players (cloud providers, network service providers, advertisers, end users and user communities) to both collaborate and compete.

7. ENABLING PERVASIVE AR

This paper hopes to motivate the computing community towards bigger questions that surround pervasive AR and away from only battling of performance challenges [4, 5, 10, 13, 14]. Performance is important; the scenarios described in this paper require low latency for usability. Important architectural design questions, however, tend to remain on the wayside as they are difficult to address, even if they will have a major impact. Proposed answers to these questions are often more difficult to validate and communicate to the community, especially using a new technology such as NDN. Despite this, we should remember that architectural innovations, including the TCP/IP specification, the creation of HTTP and its associated hypertext markup language (HTML), enabled new functions and business models previously impossible.

We thus believe that many important questions and possibilities lie within how the communication abstraction for network applications on the edge are defined. In this paper, we have explored (briefly) how current AR apps are built with restrictions imposed by the TCP/IP protocol stack. With a pervasive AR vision, we identify some key challenges that include local computing resource availability, resource discovery, security and trust

establishment, and new business models. By proposing the use of Named Data Networking to support edge communication, we suggest how the technical needs and exploration of different business models can be enabled in this new era of edge computing.

We plan to gain experience through hands-on experimentation, and to evolve our research toward developing an AR platform, on top of which many new applications can be developed, as the Web platform has been. To ease deployment challenges, the usage of NDN can be explored at the edge, without changing the core communication network (<http://ice-ar.named-data.net>). We realize that the solutions to the challenges outlined in our paper may turn out to be very different from what we suggested. Nonetheless, it remains for all of us as an academic community to begin tackling broader problems beyond performance, experimenting with possible solutions, and paving the way forward for pervasive augmented reality with strong support for user control, security and privacy.

ACKNOWLEDGMENTS

This work is partially supported by the US National Science Foundation under award CNS-1719403.

REFERENCES

- [1] Azuma, R.T. 1997. A Survey of Augmented Reality. *Presence: Teleoperators and Virtual Environments*. 6, 4 (Aug. 1997), 355–385.
- [2] Billinghurst, M., Clark, A. and Lee, G. 2015. A Survey of Augmented Reality. *Foundations and Trends® in Human-Computer Interaction*. 8, 2–3 (2015), 73–272. DOI:<https://doi.org/10.1561/11000000049>.
- [3] Burke, J. 2017. Browsing an Augmented Reality with Named Data Networking. *Computer Communication and Networks (ICCCN), 2017 26th International Conference on* (2017), 1–9.
- [4] Chakareski, J. 2017. VR/AR Immersive Communication: Caching, Edge Computing, and Transmission Trade-Offs. (2017), 36–41.
- [5] Duanmu, F., Kurdoglu, E., Hosseini, S.A., Liu, Y. and Wang, Y. 2017. Prioritized Buffer Control in Two-tier 360 Video Streaming. (2017), 13–18.
- [6] Ebling, M.R., Watson, I.T.J. and Want, R. 2017. Satya Revisits “Pervasive Computing: Vision and Challenges.” (2017), 4.
- [7] Flinn, J. 2012. Cyber Foraging: Bridging Mobile and Cloud Computing. *Synthesis Lectures on Mobile and Pervasive Computing*. 7, 2 (Sep. 2012), 1–103. DOI:<https://doi.org/10.2200/S00447ED1V01Y201209MPC010>.
- [8] Grubert, J., Langlotz, T., Zollmann, S. and Regenbrecht, H. 2017. Towards Pervasive Augmented Reality: Context-Awareness in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics*. 23, 6 (Jun. 2017), 1706–1724. DOI:<https://doi.org/10.1109/TVCG.2016.2543720>.
- [9] Ha, K., Chen, Z., Hu, W., Richter, W., Pillai, P. and Satyanarayanan, M. 2014. Towards wearable cognitive assistance. (2014), 68–81.
- [10] Ju, R., He, J., Sun, F., Li, J., Li, F., Zhu, J. and Han, L. 2017. Ultra Wide View Based Panoramic VR Streaming. (2017), 19–23.
- [11] Kang, J., Shilton, K., Estrin, D., Burke, J. and Hansen, M. 2010. Self-Surveillance Privacy. *SSRN Electronic Journal*. (2010). DOI:<https://doi.org/10.2139/ssrn.1729332>.
- [12] Li, A., Yang, X., Kandula, S. and Zhang, M. 2010. CloudCmp: comparing public cloud providers. (2010), 1.
- [13] Mangiante, S., Klas, G., Navon, A., GuanHua, Z., Ran, J. and Silva, M.D. 2017. VR is on the Edge: How to Deliver 360° Videos in Mobile Networks. (2017), 30–35.
- [14] Ran, X., Chen, H., Liu, Z. and Chen, J. 2017. Delivering Deep Learning to Mobile Devices via Offloading. (2017), 42–47.
- [15] Roesner, F., Kohno, T. and Molnar, D. 2014. Security and privacy for augmented reality systems. *Communications of the ACM*. 57, 4 (Apr. 2014), 88–96. DOI:<https://doi.org/10.1145/2580723.2580730>.
- [16] Satyanarayanan, M., Chen, Z., Ha, K., Hu, W., Richter, W. and Pillai, P. 2014. Cloudlets: at the Leading Edge of Mobile-Cloud Convergence. (2014).
- [17] Shang, W., Wang, Z., Afanasyev, A., Burke, J. and Zhang, L. 2017. Breaking out of the Cloud: Local Trust Management and Rendezvous in Named Data Networking of Things. (2017), 3–13.
- [18] Shi, W., Cao, J., Zhang, Q., Li, Y. and Xu, L. 2016. Edge Computing: Vision and Challenges. *IEEE Internet of Things Journal*. 3, 5 (Oct. 2016), 637–646. DOI:<https://doi.org/10.1109/JIOT.2016.2579198>.
- [19] Shilton, K., Burke, J.A., Estrin, D., Govindan, R. and Hansen, M. 2009. Designing the Personal Data Stream: Enabling Participatory Privacy in Mobile Personal Sensing. *37th Research Conference on Communication, Information and Internet Policy (TPRC)* (2009), 20.
- [20] U.S. General Services Administration Information Resources Management Service 1993. Data Management Issues Associated With Stovepipe Systems. (Oct. 1993), 12.
- [21] Weiser, M. The Computer for the 21st Century. *Scientific American*. 265, 94–104.
- [22] Yi, S., Hao, Z., Qin, Z. and Li, Q. 2015. Fog Computing: Platform and Applications. (Nov. 2015), 73–78.
- [23] Yu, Y., Afanasyev, A., Clark, D., claffy, kc, Jacobson, V. and Zhang, L. 2015. Schematizing Trust in Named Data Networking. (2015), 177–186.
- [24] Zhang, L., Afanasyev, A., Burke, J., Jacobson, V., claffy, kc, Crowley, P., Papadopoulos, C., Wang, L. and Zhang, B. 2014. Named data networking. *ACM SIGCOMM Computer Communication Review*. 44, 3 (Jul. 2014), 66–73. DOI:<https://doi.org/10.1145/2656877.2656887>.
- [25] Zhang, Z., Zhang, H., Newberry, E., Mastorakis, S., Li, Y., Afanasyev, A. and Zhang, L. 2018. Security Support in Named Data Networking. *NDN Project, Technical Report*. NDN-0057, (Mar. 2018).
- [26] Zimmerman, M. 2000. Weaving the web: the original design and ultimate destiny of the world wide web by its inventor [Book Review]. *IEEE Transactions on Professional Communication*. 43, 2 (Jun. 2000), 217–218. DOI:<https://doi.org/10.1109/TPC.2000.843652>.