Statement: As TCP/IP is to the Web, ICN is to the...?

Jeff Burke
jburke@remap.ucla.edu
UCLA REMAP
Los Angeles, California, USA

ABSTRACT

Recent developments in the public imagination of future networked environments provide an opening for the ICN research community to make its case, metaphors to explain the value, and a motivation for specific areas of research. The opportunity is proposed and future research directions are briefly discussed.

CCS CONCEPTS

• **Networks** → **Network architectures**; • **Computing methodologies** → **Mixed / augmented reality**; • **Computer systems organization** → *Peer-to-peer architectures.*

ACM Reference Format:
Jeff Burke. 2022. Statement: As TCP/IP is to the Web, ICN is to the...?. In *9th ACM Conference on Information-Centric Networking (ICN '22), September 19–21, 2022, Osaka, Japan.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3517212.3558094

1 INTRODUCTION

An avatar and a bot walk into a bar, followed by a browser tab. They sit down. Relaxing to the live music, the human avatar and AI bot finish a task on the tab and close it. They make small talk with others at a nearby table. The avatar moves closer to hear the band more clearly. The bot texts a human friend, who is nearby and likes jazz. The friend joins them. The three take a selfie. They post it next to other photos in a corner. Though separated for a bit when the friend's device runs out of charge, they all keep listening. Later, they replay the last third of the song together, solo'ing the saxophonist, who was very impressive. As they leave, the bot shares more about the band. It mentions they'll play in the same place again next week.

Is this another metaverse fantasy, just missing a sprinkling of crypto (for marketing) before it can be pitched? Perhaps. But if the friend, bar, some patrons, and most of the band are "in real life" (IRL) and the avatar and bot are not, perhaps the author's intent is clearer - an allegorical tale of interoperability, persistence, trust, context awareness, real-time communication (RTC), edge-supported rendering, resource discovery, and AI-assisted integration of virtual and physical space. Each short sentence is individually demonstrable today. Even the whole paragraph's blended reality is somewhat conceivable within a walled garden using a curated palette of applications, services, and devices. But how far can current approaches go? What "app" do our protagonists use to walk into the bar? The same as at the coffee shop across the street? How do they know what content to trust in the bar? What if each person's device is affiliated with a different walled garden?

2 LIMITS AND ICN OPPORTUNITY

Accompanying the hype of metaverses are substantial real shifts in how we communicate, create, and interact. Emerging approaches to shared virtual worlds and augmentation of the real world generally build on the current technical, economic, and socio-political structures of the Internet, which is centralized around a small group of dominant providers and services. Yet, public interest in applying technologies like blockchain to such distributed extended reality (XR) applications, even naively, signals an increasing frustration with the concentration of control [1] and (perhaps) renewed engagement with the ethos of a decentralized Internet.

So far, most considerations of technology for metaverses seem to assume the use of the current Internet architecture, modern web protocols, and cloud service providers. For example, the Khronos Group, an industry consortium, promotes interoperability of hardware platforms through OpenXR [3] and asset standardization through its glTF format [4] and other efforts. WebAR [9] uses browser-based media technologies to implement AR on the web, offering a standards-focused vision for accessing metaverses. These approaches inherit the advantages of the Internet protocols' ubquity, but also current limitations with respect to mobility, security, and other areas discussed at length in the ICN community. Broad assumptions around the use of current cloud technologies and existing web protocol security, for example, suggests that rhetoric of decentralized metaverses is imagined with minimal thought to whether typical network stacks of today align with the application goals.

ICN offers a ground-up interoperability story for metaverses, in which common network infrastructure directly supports scaling of XR content distribution and decentralized interaction across mobile publishers and consumers. The ICN research community should articulate perspectives, design approaches, and create examples that apply information-centric principles to address the challenges emerging from building large, decentralized XR systems. Below, the paper explores this a bit further, using Named Data Networking (NDN) as an example of a specific ICN architecture when needed.

3 FUTURE RESEARCH DIRECTIONS

Whether one is a true believer or not, the "metaverse" offers a shared mental model–a loose one, at least–in which the motivations for deploying ICN are reasonably approachable. Most metaverse concepts involve large-scale, persistent, shared virtual worlds accessible over the network, presented primarily via real-time rendering of visual, auditory, and haptic media. Shared views of these worlds are synchronized between participants at latencies supporting interactivity.

ICN '22, September 19–21, 2022, Osaka, Japan

ACM ISBN 978-1-4503-9257-0/22/09.
https://doi.org/10.1145/3517212.3558094

The notion of a "host" is abstracted further away than in today's web. Client devices fade into the background in favor of the data supplanting or augmenting the real world. Thus, a metaverse consists of *information* not associated with the physical world unless it needs to describe or provide interaction with it. The experiential semantics are viscerally information-centric, which helps to motivate the ICN research opportunities discussed below.

Persistence. The information forming a metaverse persists across sessions and users. Research on support for persistent storage in ICN networks is thus vital to the domain. NDN's foundation of immutable data and its support for intrinsic storage[1] suggest the network layer itself, or ubiquitous services immediately above it, could support such persistence of virtual objects. Such support would more closely align the stack's semantics with this key property of metaverses. (It would also make ICN's story more competitive when compared, perhaps unfairly, to most basic cloud services.) If persistence is consistently supported network-wide, there is a straightforward basis upon which to design decentralized persistence at the application layer. If the network standardized object-level storage and allowed providers to sell it, for example implementing the WORM-like model discussed in [11], NDN would have a substantially improved value proposition.[2]

"Content" and Interoperability. Research is needed to design the relationships among metaverse-layer objects and the named packets that an ICN network moves and stores. For example, formats promoted by members of the Metaverse Standard Group [5]–such as the glTF runtime 3D asset format mentioned above–could be examined or extended to use NDN's packet-level naming to enable applications to efficiently choose which elements to access within objects or streams as required by a given interactive view (for example). Most common real-time platforms download entire 3D objects before rendering them and largely conceive of assets as static content made interactive by programmatic behaviors.[3] Careful design of how objects are serialized for NDN networks offers opportunities for better organization of assets and granular (pre)fetching of only what is needed by a given renderer for a given view.[4] More compellingly, it opens up the opportunity for efficient services that progressively stream object elements just-in-time. From this author's perspective, this is more consistent with metaverses' real-time emphasis than the domain's current interoperability focus on file-like objects.

Naming and Spatial Organization. The name components that describe these interoperable, granularly-accessible data packets are likely on the right hand side (RHS) of a longer NDN name in which the parent components (LHS) provide higher-level organization, e.g., describing "where" and "in which metaverse" the object was published. The best LHS naming strategies for spatially-organized data is an area of research that points to a long-running challenge to the ICN community: how to best integrate knowledge from research in databases and related fields where these challenges have been considered for decades. This does not necessarily require new Interest/Data semantics in the case of NDN. Instead, research could explore how higher-level, transport-like protocols like Sync could support range-based queries on spatially named data.[5] There are many other interesting challenges in both LHS naming (typically network, administrative, and content organization) and RHS naming (object organization, versioning, segmentation, etc.). For example, how hierarchical scene graphs could be represented in hierarchical namespaces–and whether this provides value to application developers–is a potentially interesting topic.

Trust, Provenance, and Transactions. Critically, ICN can enable metaverse objects to be disentangled from the security provided by a source or a given channel of communication, with the named data representation secured at the time of publication instead. The author proposes that this could provide an *actual solution* for the typically espoused vision of an interoperable metaverse–not by solving all aspects of trust and provenance in the domain, but by allowing metaverse objects to be considered independently of their immediate source. The primary needs from the ICN research community continue to be simple and accessible security tools for name-based authentication and confidentiality.

A variety of interesting domain-specific security opportunities exist, as well. Spatial names might be leveraged to enable trust based on shared regions in the metaverse, in addition to other properties. Mutual trust in spatial name components could allow the timely exchange of decryption keys for audio streams only for listeners occupying a space "near" the jazz band in the example above. Integration of provenance (whether in an envelope or external ledger) for transformed data objects could give application developers consistent means to manage ownership and related transactions. Distributed ledger, smart contract, and other transaction support for a named data approach to spatially organized data appear to be some other necessary pieces required for commercial viability and, more generally, a shared, virtual world with an evolving history.

4 CONCLUSION

This statement paper encourages ICN research that considers how humans *will* interact with information in the future. Shared virtual worlds are emerging as a key component of tomorrow's experiences. Their information-centric nature is a semantic fit for ICN. It can simplify the network stack used by metaverses to realize the decentralized interoperability and security envisioned by many.

Finally, performance requirements are intentionally not discussed above. This is partially for brevity, but also because their consideration alone is unlikely to aid adoption of ICN for metaverse applications. While ICN may offer performance benefits at scale and should perform on par with IP solutions to be competitive, there are more immediate, fundamental, and cross-layer benefits in simplifying stack semantics for data exchange and security in metaverses.[6]

[1]Here, storage is meant to include both caching as typically discussed, as well as the possibility of standardized persistent storage for NDN objects.
[2]It is out of scope here, but some consideration is needed for how the network (and metaverse applications) should support the socially important act of forgetting.[10]
[3]In the popular press last year, [7] breaks down they journey to sub-asset-level streaming needed for open world gaming or metaverses, and is interesting to read with ICN (and the author's company's solution [2]) in mind.
[4]We have shown aspects of this elsewhere with 2D video[6]. The opportunity appears more significant for 3D scenes or objects that are already semantically organized.
[5]An early example of this is [8].
[6]That said, a nice short-term opportunity may be to demonstrate robustness of ICN-based XR applications on mobile networks in the wild or using multiple network interfaces, both cases where deployed IP-based approaches struggle with performance.

5 ACKNOWLEDGEMENTS

The author is grateful to the anonymous reviewers for their insightful comments and encouragement to include more details in this short paper, and the paper's shepherd, Susmit Shannigrahi, for feedback and guidance on the revision.

REFERENCES

[1] Jana Arbanas, Allan V. Cook, Kevin Westcott, and Chris Arkenberg. 2022. The metaverse and Web3: The next internet platform. https://www2.deloitte.com/us/en/insights/industry/technology/web3-and-metaverse-the-future-of-the-internet.html.
[2] Concurrents. 2021. The Future of Discovery and Sharing in the Creator Economy. https://www.concurrents.com/about.
[3] Khronos Group. 2019. OpenXR Overview. https://www.khronos.org/openxr/.
[4] Khronos Group. 2022. glTF Overview. https://www.khronos.org/gltf/.
[5] Metaverse Standards Group. 2022. Overview. https://metaverse-standards.org/.
[6] Peter Gusev, Jeff Thompson, and Jeff Burke. 2019. Data-centric video for mixed reality. In *2019 28th International Conference on Computer Communication and Networks (ICCCN)*. IEEE, 1–11.
[7] Barry Jenkins. 2021. The Journey To Open-World Gaming. https://www.forbes.com/sites/forbestechcouncil/2021/07/21/the-journey-to-open-world-gaming/?sh=15863a19126b.
[8] Philipp Moll, Selina Isak, Hermann Hellwagner, and Jeff Burke. 2021. A Quadtree-based synchronization protocol for inter-server game state synchronization. *Computer Networks* 185 (2021), 107723.
[9] Rock Paper Reality. 2022. The Main Components of WebAR Development. https://rockpaperreality.com/web-ar/webar-development-main-components/.
[10] Katie Shilton, Jeffrey A Burke, Kimberly C Claffy, and Lixia Zhang. 2016. Anticipating policy and social implications of named data networking. *Commun. ACM* 59, 12 (2016), 92–101.
[11] Christian Tschudin. 2018. End-to-end encrypted scalable abstract data types over ICN. In *Proceedings of the 5th ACM Conference on Information-Centric Networking*. 88–94.