

An Architecture for Interspatial Communication

Anil Madhavapeddy
Computer Laboratory
University of Cambridge
avsm2@cl.cam.ac.uk

KC Sivaramakrishnan
Computer Laboratory
University of Cambridge
sk826@cl.cam.ac.uk

Gemma Gordon
Computer Laboratory
University of Cambridge
gg417@cl.cam.ac.uk

Thomas Gazagnaire
Tarides
Paris, France
thomas@tarides.com

Abstract—Digital infrastructure in modern urban environments is currently very Internet-centric, and involves transmitting data to physically remote environments. The cost for this is data insecurity, high response latency and unpredictable reliability of services. In this paper, we lay out a software architecture that inverts the current model by building an operating system designed to securely connect physical spaces with extremely low latency, high bandwidth local-area computation capabilities and service discovery. We describe our early prototype design OSMOSE, which is based on unikernels and a distributed store.

I. INTRODUCTION

In his classic essay on “Computing for the 21st Century” [1], Mark Weiser observed that:

The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.

Since then, there have been tremendous advances in mobile and sensing technologies, and there is an ongoing rapid deployment of “smart” digital infrastructure that augments the physical environment. Consider the following scenario for a next-generation building:

“A group of people are meeting in an coworking space to discuss a project they are working on together. They are the first to arrive at the building on a cold winter morning, and are directed to a meeting room automatically by audio directions individually projected to them. As they walk into the room the lights are already on and the heating has been preset to a comfortable level. They wave to a wall where a display appears, and the building recognises the gesture and projects a shared folder of the project. It begins to record and encrypt the conversation with a group secret keyed to the participants. One of the participants receives a tight-beam audio notification informing them that they are needed elsewhere briefly. When they leave the room the shared recording is immediately paused within milliseconds to allow other participants to casually chat among themselves. Subsequently, the encrypted recording is transcribed, and made available in the datastores of each participant under the terms of their shared disclosure agreement for their project.”

This scenario illustrates the distance between Weiser’s 1999 vision of ubiquitous computing and where we currently are almost two decades on. The above narrative features *no wearables or mobiles* among the human participants – instead of an array of individual smart phones, watches and laptops, the situated environment of the building is providing secure,

shared, multi-tenant infrastructure for audio and visual communications among the human participants. This is analogous to the shift from individual computers to multi-tenant cloud infrastructure in the past decade [2], but applied to the sharing of physical places and devices.

Another aspect to consider in the scenario is the seamless, low-latency nature of the interactions that were described between humans and electronic devices. When a human takes an action the environment is able to react within milliseconds to take action immediately, rather than delaying for seconds and causing the awkward delays that we have become used to with modern Internet-connected mobile devices. By providing immediacy, the technology can complement social interactions rather than interrupting them [3].

Encouragingly, much of the hardware required to physically assemble the described building is available off the shelf. E-ink display wallpaper, parametric audio speakers for directional broadcast, gesture recognition and smart lighting and heating are all available or relatively easy to construct. The missing link is the foundational software that manages, coordinates and secures the distributed hardware – the operating system for digitally connecting physical spaces together.

In this position paper, we begin to bridge this gap between available hardware and missing software. We construct an architectural model for *interspatial networking* – an operating system for the dense, interconnected and shared urban spaces that most humans live in. It aims to shift us away from the wide-area, mobile-oriented devices that are permeating early deployments of smart infrastructure, and towards a sustainable digital model that is far more similar to a conventional utility such as electricity or gas. Our system aims to make it far simpler and more secure to introduce, manage and rely on digital devices during day-to-day life in urban environments.

There is also an exciting new generation of upcoming applications that demand vast amounts of bandwidth and low-latency responses. Augmented and virtual reality, environmental e-ink displays, parametric directional sound, and robotic appliances simply do not work well if the latencies of interactions with them rise above a certain point. This paper proposes an operating software architecture aimed at fully supporting these new applications on modern hardware platforms. We will next describe some of the challenges in more detail (§II), discuss the design principles behind interspatial applications (§III), then present the design of our OSMOSE prototype OS (§IV), and finally discuss examples and some implications (§V).

II. THE PROBLEM WITH EXISTING MOBILE SOFTWARE

Given that the hardware is all available, why is the above scenario difficult to implement? The answer lies in the traditional operating systems and network architecture that power the current generation of smart devices. Because of the rise of cloud computing, they are typically built around communication to centralised Internet services. For example, consider what happens when we speak into an Apple Watch in order to retrieve some information. For this to work, the watch must be connected to a mobile phone, which in turn needs to establish a wireless connection to the Internet, where Apple voice recognition services will dispatch a query to the Google search engine. If any one of the services in this chain breaks (for example, the common case of the phone signal being “trapped” by a wifi authentication page), then the user experience is broken. Even when it does work, it can take seconds to respond to the voice query, and with very variable latency. Once the response comes through, handoff to other devices is also difficult unless they are owned by the user.

A. Reliability

Services deployed in physical environments need to work all the time, and be locally debuggable when they do fail. How do we shift from a light switch that “almost works” after being pressed several times to seamless voice or gesture-driven services that are always tuned and available in a given environment? They also need to work independently of Internet connectivity, so that every building can be an island of digital services even when offline.

B. Security

The amount of sensitive data being captured in these environments is tremendous, and much of it should not leave the confines of the physical space without explicit permission from all parties involved. This is extremely difficult to police given the amount of Internet-wide coordination used in existing devices, but can be fixed if the local environment provides a structured mechanism for handling such storage securely with respect to local environmental policies.

C. Latency

Interactive services require response times beneath the uncanny valley of human perception. For the scenario above to really feel seamless, we need a new “latency first” application architecture that makes data and computation capacity available physically near the human users, with scheduling tolerances for responses in the milliseconds rather than seconds.

Solving these problems is difficult to do piecemeal across individual parts of the software stack, since they are currently general-purpose and loosely coupled. A typical IoT device might run its own copy of the Linux kernel, with an embedded userspace, a VPN into a centralised management server operated by the vendor for updates, and a mobile application for the user to manage it. There is little synergy or dependence on shared infrastructure to assist with the process.

III. TOWARDS INTERSPATIAL APPLICATIONS

Before proposing a solution, we consider some design principles to underpin how applications in physical spaces operate. When dealing with physical devices, the latency of response to external events is paramount. Safety critical systems often require “hard realtime” operation, or more often attempt to provide soft guarantees about when they respond. As we move further up the stack to modern mobile applications, there are no such guarantees. Pressing a button on an iOS or Android device goes through many layers of scheduling and network connectivity before a response is generated.

Our interspatial architecture needs to allow us to deploy “latency-first apps” into physical environments – ones that are designed to run on local devices with response budgets in the milliseconds, and with minimal external connectivity needs. This implies that there is simply no time for doing conventional operations such as network scanning or service discovery in serial – by the time the user has asked a question, the physical environment should already have the appropriate resources established. We thus need to rearrange the programming models around building interspatial applications to make them suitable for such an environment.

A. Incremental Networking

While it is currently possible to spin up services on demand within milliseconds [4] it is difficult to transmit a response for complex clustered services within a time budget of a few milliseconds. A device that is starting “cold” will need to establish a network connection to the local node, most commonly via TCP and subsequently a secure transport via TLS. The number of hops required to establish a secure connection have been recognised as a problem on the wider Internet, and new low-latency protocols such as QUIC [5] are being deployed in browsers. However, even these newer protocols only solve part of the problem, since they do not integrate closely with the application layer. Once the connection has been established, we still need to authenticate the user(s), negotiate security keys, perhaps perform version negotiation between devices, and usually interrupt the user at an inconvenient time for a software security update.

The reason that every device has to currently do all this work is partly due to the end-to-end principle that guides the design of Internet protocols. Every IP node is a “dumb” router that knows how to forward protocol packets, but lacks higher-level application information. Existing system interfaces (such as the BSD sockets API, or DNS resolution) implement this network stack, which is in turn baked into existing application logic. Each device needs to repetitively perform the same actions to establish connectivity, with the surrounding network environment being unable to assist.

An alternative interface in a physical environment is to move away from just-in-time connectivity towards a model of establishing connectivity incrementally as the opportunity arises. When sufficient information from the environment arises to establish a partial connection (for example, upon entering a building but before making any service requests)

the application needs to perform the operations that it can do so then, such as negotiating security keys or ensuring that software versions match. When a service request does finally happen, the previously derived connection information can be used to perform a single hop. Incremental connectivity requires fundamental programming interface changes in order to combine information across the traditional operating system and application stacks, and for us to move away from the venerable sockets API [6].

B. Spatial Interfaces

Incremental networking lets an individual network connection to a node be established quickly. In a physical environment it is also necessary to mesh wired sensors (such as motion detectors or cameras) with mobile ones (humans carrying location beacons). The wired systems need to be establishing connectivity with services required by applications, and the wireless ones need to be searching for nearby network nodes via Zigbee [7] or Bluetooth [8]. This connectivity mesh is necessary to ensure that all of the distributed components that comprise a physical environment can all communicate with predictable and very low latency.

Internet protocols such as TCP have not traditionally not handled connection handoff well – extensions such as multipath TCP are effective [9] but have struggled with extending existing interfaces such as the sockets API with the new routing semantics available to applications [10].

For our interspatial applications, a building needs to be able to manage all the local resources (include bandwidth or storage) just as it does so with other utilities such as electricity. In return, individual applications do not need to manage their own networking, storage and authentication needs as they are provided by the surrounding environment. This requires a change in the traditional programming model to allow applications to multiplex their connectivity needs around a networking environment that is multipath and dynamic as humans move around the physical space.

C. Native Hardware

Applications are traditionally built for a specific device profile in mind – for example, a mobile phone or a desktop application. The proliferation of multiple display form factors has resulted in a new mode of “responsive” design [11] that can adapt an interface to multiple resolutions and sizes.

With interspatial applications, we wish to eliminate the need for importing wearable and mobile devices into a building as the sole mechanism of interaction with a user. Hardware present in the environment such as parametric speakers for 3D positional audio [12] or wallpaper projections [13] should be able to be used by local users, rather than being constrained to the single form factor mobile devices (e.g. a watch or a mobile phone) that we carry around. For this to work, interspatial applications needs to be designed to have responsive user interfaces, but also a suitable trust model to let users establish secure connections to environmental hardware that is not directly owned by them [14].

D. A Cross-Layer Unikernel Software Architecture

The three design principles for interspatial applications programming interfaces (incremental networking, spatial interfaces and use of native hardware) require cross-cutting changes across traditional operating system and application programming interfaces. What would a software solution to operating embedded physical infrastructure might look like, if designed with a clean slate in mind? The first step to making this a practical prospect is to adopt the discipline of deploying *unikernels* [15] on the hardware, and replacing the traditional layered OS stack.

Unikernels are specialised operating systems that are compiled together with application source code and configuration, resulting in a specialised binary that can boot in milliseconds [16] and eliminates traditional runtime layers in favour of optimised build-time assembly. Unikernels were originally developed from the concept of library operating systems [17] and applied to cloud computing. We now find that the same approach is a perfect fit to driving the world of resource-constrained embedded hardware found in physical environments.

In a library operating system, the functionality that is conventionally found in a monolithic operating system kernel is broken out into software libraries that are available for use by the application in exactly the same way as other higher-level functionality currently is. All of the software libraries are linked together with a small boot layer – this includes traditionally kernel-based interfaces such as hardware device drivers, which can now run in the same privilege level as the application driving them. This model is ideal for embedded devices, since it can result in direct low-latency and energy-efficient access to the hardware.

Crucially, the approach also allows the application to be tailored to the operational model of the embedded hardware. For example, the same application source code could be compiled to a tiny embedded processor without an MMU, or also be recompiled into a conventional Unix process. One downside of this specialisation approach is that traditional multi-user (e.g. UNIX processes) operation is more difficult in a single device once the unikernel has been deployed (we address this in §IV-B). An upside is that the baseline operating system can be highly stripped down and easily substituted – for example, a conventional Linux kernel can be replaced by a small, formally verified microkernel such as seL4 without having to rewrite the rest of the application code [18].

The unikernel approach is thus important to the long term sustainability of designing software for physical environments. It lets us close the gap between the requirements of application programming interfaces and the diverse requirements of the embedded hardware. Instead of forcing vendors to squeeze an ever-growing operating system stack for trivial tasks (such as driving a lightswitch), adopting unikernel-based interfaces means that the same unified software codebase can be tailored across the variety of hardware that we can expect to see in a typical situated environment in the next few decades.

IV. THE INTERSPATIAL OPERATING SYSTEM

We now discuss the design of a prototype interspatial operating system based on a unikernel architecture, dubbed OSMOSE. Its scope is to drive all the hardware in a single physical space – such as a building – including the thousands of sensors and actuators that may be present. As background requirements, we assume that:

- we have a reliable physical topology available of the situated environment. Details are out of scope for this position paper – there are a number of indoor localisation technologies available [19].
- users and devices have an out-of-band method to authenticate to the operating system – this might be biometric or face recognition or a password or hardware token [20], but this has to happen as they enter the physical space.
- there is ample hardware available locally for applications – for example an array of embedded devices in each room – and the hardware is all connected via room-local reliable wired networks, with wireless at the last hop only.

Since OSMOSE needs to run across a variety of hardware with differing resource constraints, the traditional model of booting a single image no longer holds. Instead of application binaries, the input to the system is a collection of declarative application fragments that represent processing logic and how it ties together. This is then combined with physical topology information and security policies to compile a set of unikernels that are booted on the building’s hardware. This cycle of compilation happens continuously and incrementally – as applications are introduced or policies change, the unikernels are recompiled in real-time to all the devices.

Applications do not communicate freely between each other over the network with conventional TCP/IP – instead, the OS sets up explicit channels that provide a high-level streaming interface suited to the application needs (for example, a low-bandwidth control channel *vs* a high-bandwidth video processing ring buffer). The resulting set of communication circuits is tracked as a dynamic dataflow graph, with the building hardware representing nodes and the application logic as edges. All communications in the system are dispatched through a distributed storage system that immutably stores interactions for audit and debugging purposes.

Figure 1 illustrates the overall architecture of OSMOSE. It implements these components using the MirageOS¹ unikernel framework, written in the OCaml programming language.

A. Declarative Layer

Applications are partitioned up into declarative source code fragments that represent units of computation, and then scheduled together by via a structured event stream using the Capnp² RPC framework. This is becoming a popular architecture on container-native computing infrastructures via the Serverless movement [21], since it allows applications to run on distributed virtual machines that can be rescheduled depending on

load and hardware availability. In an interspatial deployment, the application logic is compiled to run directly on the physical hardware as a unikernel [22], with appropriate network channels established depending on the available circuits (e.g. a local point-to-point wireless link).

The services are composed together via the domain-specific language of combinators provided by MirageOS. For example, a voice recognition application can request an audio stream, which can then be mapped onto a feature extraction module. The application does not know the exact source of the audio stream at this stage. In the case of our prototype design, we use the MirageOS device driver model to define many of these sources and sinks. The resulting programming model is event-driven and reactive, and encourages the programmer to provide ongoing incremental updates so that user interfaces can be updated rapidly.

B. Builder and Artefact Layer

Once we have all the source code (in the form of application logic, topology and configuration code), we need to compile it into executable unikernels that can run on the available hardware. Unlike a conventional operating system that runs on a single hardware host, the code is compiled into a heterogenous set of bootable kernels that are specialised to the various devices available in the cluster. For example, a portion of the application logic may be compiled to run on a GPU or FPGA, and another piece to run on an embedded ARM CPU without an MMU. This is accomplished via an incremental build service that can rapidly link and specialise the source code into unikernels.

The build service tracks all of the source code that via a branching datastore that is similar to Git. Every single line of application, operating system and configuration logic is stored in the same place, and so the builder can easily track precisely what code is running on the array of hardware. We use the Irmin [23] datastore in our prototype design. Irmin is a unikernel implementation of the Git model for MirageOS, and provides the facility to handle complex distributed datastructures just like source code.

Irmin is structured as a series of OCaml libraries that expose increasingly complex storage functionality. Applications request the minimal functionality they need for their purposes, and the appropriate storage layer is crafted from the combination of requirements. For example, one application (a button) might just need a key-value store, whereas another (a projector) might want a read-only filesystem. Importantly, all of the storage is backed by the same Merkle-tree-based storage, granting us the ability to precisely track the behaviour of every hardware node.

C. Scheduler Layer

The scheduler is responsible for triggering the builder, deploying the resulting unikernels to hardware, and establishing network circuits between devices. It receives an event stream from the hardware (such as network topology changes, or node failures). This layer is the most radical departure from a

¹See <https://mirage.io>

²See <http://capnproto.org>

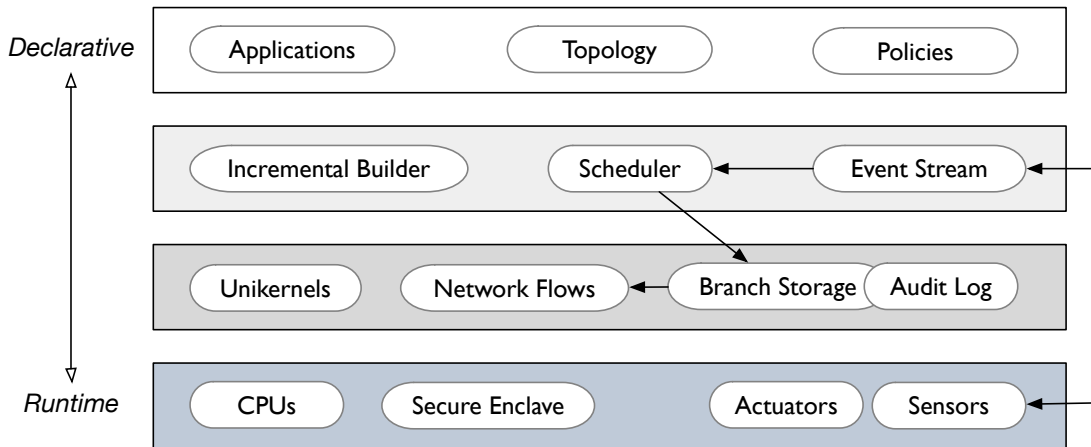


Fig. 1. Architecture of the interspatial operating system. The system is configured declaratively via a set of application source code, policies and network topology, and is then incrementally built and deployed onto hardware. All communications happens via the Merkle-tree based branch storage.

conventional operating system model. Instead of a process-based system, a deployment can be viewed as a graph of distributed hardware nodes, through which event stream data from the environment flows.

Every coordination operation between nodes is piped through the Irmin datastore, allowing the complete state of the system to be tracked with strong provenance. We assume that there is a large amount of persistent local storage available in the deployment, but this can run in a separate processing node from the low-power embedded hardware. The basic model for the Irmin-based communication has been validated over the past decade via a series of shared-memory implementations in the Xen hypervisor [24]. In order to keep latency down despite the use of Merkle-tree-based coordination, Irmin can expose a shared-memory endpoint to each node, and local operations between two nodes are asynchronously reconciled to a remote store via a set of short-lived branches.

When applications need access to the outside world, they do so via RPC calls to the scheduler. We use the Capnp RPC framework for this purpose, which effectively acts as the system call layer. Capnp features a very efficient low-level serialisation mechanism (important when communicating between embedded devices), but also a compiler for protocol schemas that integrates secure capabilities [25]. These capabilities can be passed around nodes as opaque references and used to authenticate access to different parts of the system. Since a capability can only be generated and communicated through the Irmin store, this means that tracing distributed function calls through a deployment are available “for free” by inspecting and reconstructing the storage.

Interspatial applications require a lot of immediate interaction with the physical environment. This could involve reading sensor input (e.g. gesture recognition, audio inputs, pressure sensors) and actuating outputs to sensors (e.g. heating, lighting, speakers, displays). The scheduler layer thus maintains multipath network connections that track a *physical network topology* that describes the containment relationships present

in a physical environment – a chair is on a floor beside a wall within a room that is on the first floor of a building in Cambridge in the United Kingdom.

As users interact with their environment, network connections are set up as they move *in parallel* with their actions. In our earlier scenario, when the user first enters the building they are authenticated with the building systems and given a session key that is used to subsequently track them. As they walk towards the meeting room, their preferences are reconciled with other participants. Once they enter the room, they are allocated a network slot inside that room’s network, and their voiceprint details uploaded to local embedded processors. Every time a new participant enters, their collection of identities are aggregated to generate temporary encryption keys for any shared communication. Video and voice are encrypted in real-time using these keys and saved to secure storage enclaves within the walls, available for immediate processing by locally connected computation. Once a participant leaves a space, their session keys are deleted and the hardware resources such as displays are freed, ready for re-use by other humans.

This design inverts the conventional model of establishing on-demand connections to remote services with a model that incrementally establishes nearby circuits for applications, with local data processing and resource allocation done as the humans move around the situated environment in real time. Since connections are established incrementally as humans move around, when an action actually has to be taken it can be done so with a single-hop packet, thus minimising round-trip times and error-prone connection establishment protocols that would normally add latency variance to an interaction.

D. Hardware Layer

The unikernels stored in the artefact layer are regularly booted on the various pieces of embedded hardware throughout the deployment. There is no need for dynamic network discovery since static topology information is baked into every booting kernel. If the network environment changes, then the scheduler will trigger a recompile via the branch store and

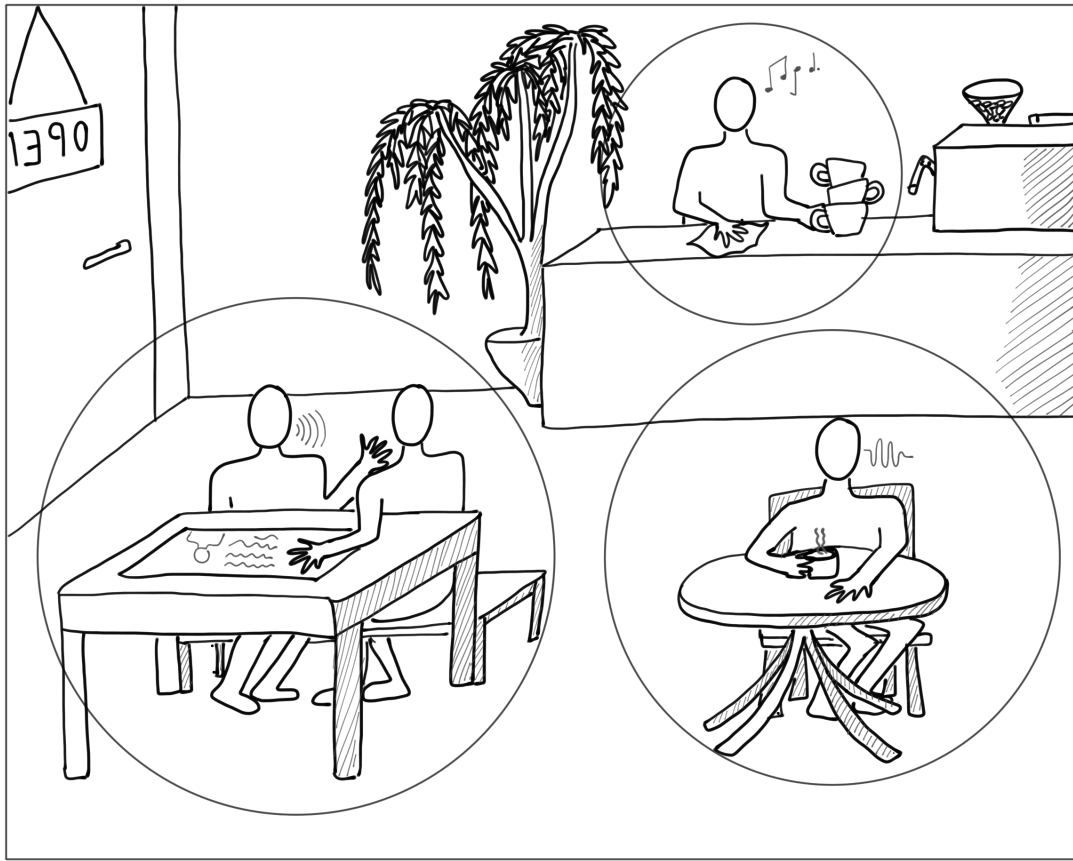


Fig. 2. A shared co-working space where OSMOSE is deployed. There are two users collaborating using a shared display on the table, while another user is listening to a personalised audio stream. The barista is receiving music and audio notifications while working.

perform a rolling upgrade. This is a similar discipline to a Continuous Integration / Deployment pipeline commonly adopted in the cloud, except applied to embedded firmware via unikernels. It is a model that is easily possible as a result of the pervasive use of the Irmin datastore.

At an individual device level, it is now common to find domain-specific hardware. The most notable is the presence of a secure enclave (e.g. SGX [26] or Trustzone [27]) that can be used to store private user data [28]. All of the data that is captured from sensors can only be committed to the local Irmin store, and so it can be directly encrypted via a secure enclave before ever being stored persistently. Although this may seem expensive from a resource usage perspective, the data is not stored for long periods of time, and nor does it need to be transmitted externally. It can, however, be buffered for long enough for interested local applications to process it directly (e.g. transcribe a voice recording to a smaller, but still encrypted, textual equivalent).

The design of OSMOSE that we have described in this section is still quite a low-level one, and is concerned with establishing the basic primitives required to boot and run unikernel code on an array of embedded hardware. However, the use of a schema-based RPC compiler for inter-component communication means that it is very easy to extend the system

to accommodate new device drivers and data processing models. The MirageOS application framework that we are using has been under development since 2007, and has a growing set of library-based abstractions for networking, storage and coordination that work well for cloud-based deployments. We anticipate that as OSMOSE matures, there will be many equivalent abstractions (e.g. real time video processing) contributed for interspatial device drivers.

V. DISCUSSION

A. An Example Deployment

Figure 2 illustrates what an OSMOSE deployment might be in a shared co-working space. Instead of the usual crowded array of laptops, mobile phones and headphones, the users benefit from use of the hardware present in the situated environment. The two users on the bottom left can collaborate using a touchscreen built into the table, and their conversation transcribed and encrypted via local microphones that can filter out noise from other tables due to their proximity to the user. Meanwhile, the lone worker on the bottom right can listen to an encrypted podcast via parametric audio speakers that project the sound directly from the ceiling, without a need for the user to hook up headphones and a mobile device. Finally, the barista at the top right can work while listening to their

music selection with personalised notifications, and also pause it as soon as a customer requires the barista’s attention.

Behind the scenes, OSMOSE is providing important operating system services to this scenario: each of the users has been authenticated upon entry to the building and could install their interspatial applications using the local wireless network. They each have their own personal policies (e.g. audio sources and conversation transcription) that have been installed into the building. Each of the users need to be tracked by the building with low latency to ensure that they can immediately interact with the devices in front of them, while also enforcing isolation of their data across the shared hardware in the building such as the microphones and parametric audio speakers. As they purchased their caffeinated beverages and took their seats, other services such as payments could also be potentially performed securely using this infrastructure.

Our next task is to design user studies of particularly interactive local applications (with some inspiration drawn from the field of ambient intelligence [29]) and evaluate them in OSMOSE versus a conventional cloud-driven deployment.

B. Security and Privacy

Cyberphysical security is vital to creating a trustworthy digital future for our smart environments. Accordingly, every layer of OSMOSE is engineered with this in mind. At the lower levels, the unikernel approach is designed to eliminate unnecessary code and promote the use of safer programming languages throughout the software stack. At the hardware level, unikernels can compile flexibly and efficiently enough to make use of the limited resources in CPU-based secure enclaves [28]. At the storage level, the use of capabilities and distributed ledgers (aka “blockchain”) to track provenance throughout the system means that there is accountability builtin from the ground up.

The broader concern about systems like OSMOSE are the drawbacks that arise from pervasive data recording. Any such highly personalized and context-aware system builds brings up societal and cultural concerns about individual privacy [30]. This is mainly a concern due to data leakage that arises from third parties (e.g. cloud providers) gaining access to this private data. Our approach of primarily keeping data local to the building (due to the need to operate without external Internet connectivity) provides a natural defence against data leakage. To be practical however, the applications running on OSMOSE need to provide the same or better levels of data analysis as cloud-based recommender systems currently operate. There is complementary research ongoing to house data silos near the user [31], and combining this local data with broader global data from online social networks [32]. The data capture systems in OSMOSE can replace the role currently taken by mobile devices [33], and also provide a corresponding increase in accuracy due to not requiring the data to be transmitted remotely before being processed.

The European General Data Protection Regulation [34] that is being introduced from mid-2018 means that any interspatial deployment also needs to engineer *deletion* for all the

tracked data. OSMOSE is designed to ensure that deletion is an expensive but reliable operation – since every bit of data is tracked, a distributed garbage collector can traverse everything and rewrite history on all nodes to remove any individual data that should be excised. The system also regularly runs garbage collection on data to compact or delete it once the immediate processing needs are met. Most obviously, OSMOSE is designed to move computation to the data instead of uploading the data to a remote cloud. This is the biggest improvement for security and privacy for individuals – they can benefit from the futuristic augmented reality interactions while knowing that their data is kept physically on the premises.

VI. CONCLUSION

We have presented an early design for a distributed operating system designed for deploying a new generation of low-latency, interactive applications in urban environments. Our design inverts the existing model of funnelling data to the cloud, and instead provides the infrastructure for rapidly processing data locally. In return, this will provide a foundation for sustainably and securely managing the trillions of embedded devices that form the emerging smart cities movement.

Our OSMOSE prototype design brings the traditional benefits of an operating system to the distributed array of hardware that comprises a physical building — resource scheduling, hardware isolation and a userlevel programming interface. The programming model is aimed at building real-time, interactive interfaces that can do complex data processing local to the user, without being forced to ship data remotely to the cloud.

Any operating system is only worth the applications that run on it. While our application interfaces are intentionally not backwards compatible with existing frameworks to give us room to experiment, their underlying programming model is a familiar one to those who program cloud-based infrastructure [35]. We are hoping that the micro service-like application model combined with automated deployments on embedded devices [36] will be compelling to developers.

Our next steps are to create a fuller specification of the design outlined in this position paper and evaluate realistic next-generation interactive applications on this platform, and to build a real physical implementation of our design in a building environment. All of the source code of the constituent components discussed here such as MirageOS and Irmin are available under a BSD-style license from <https://github.com/mirage>.

ACKNOWLEDGEMENTS

The authors would like to thank Mark Wormald from Pembroke College for coining the term “interspatial”. We also thank the MirageOS development team for their contributions that made OSMOSE possible, and in particular to Thomas Leonard for writing the Capnp MirageOS library. The ideas in this paper have also benefited from discussions with Ian Leslie, Allison Randal, Zahra Tarkhani, Jon Crowcroft, David Scott, Thomas Hagggett, Sadiq Jaffer and Richard Mortier.

REFERENCES

- [1] M. Weiser, "The computer for the 21st century," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 3, no. 3, pp. 3–11, Jul. 1999. [Online]. Available: <http://doi.acm.org/10.1145/329124.329126>
- [2] B. Hayes, "Cloud computing," *Commun. ACM*, vol. 51, no. 7, pp. 9–11, Jul. 2008.
- [3] J. Deber, R. Jota, C. Forlines, and D. Wigdor, "How much faster is fast enough?: User perception of latency," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ser. CHI '15. New York, NY, USA: ACM, 2015, pp. 1827–1836.
- [4] A. Madhavapeddy, T. Leonard, M. Skjegstad, T. Gazagnaire, D. Sheets, D. Scott, R. Mortier, A. Chaudhry, B. Singh, J. Ludlam, J. Crowcroft, and I. Leslie, "Jitsu: Just-in-time summoning of unikernels," in *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*. Oakland, CA: USENIX Association, 2015, pp. 559–573.
- [5] A. Langley, A. Riddoch, A. Wilk, A. Vicente, C. Krasic, D. Zhang, F. Yang, F. Kouranov, I. Swett, J. Iyengar, J. Bailey, J. Dorfman, J. Roskind, J. Kulik, P. Westin, R. Tenneti, R. Shade, R. Hamilton, V. Vasiliev, W.-T. Chang, and Z. Shi, "The QUIC transport protocol: Design and internet-scale deployment," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '17. New York, NY, USA: ACM, 2017, pp. 183–196. [Online]. Available: <http://doi.acm.org/10.1145/3098822.3098842>
- [6] G. V. Neville-Neil, "Whither sockets?" *Commun. ACM*, vol. 52, no. 6, pp. 51–55, Jun. 2009. [Online]. Available: <http://doi.acm.org/10.1145/1516046.1516061>
- [7] D. Gislason, *Zigbee Wireless Networking*, pap/onl ed. Newton, MA, USA: Newnes, 2008.
- [8] M. Dideles, "Bluetooth: A technical overview," *Crossroads*, vol. 9, no. 4, pp. 11–18, Jun. 2003. [Online]. Available: <http://doi.acm.org/10.1145/904080.904083>
- [9] C. Paasch and O. Bonaventure, "Multipath TCP," *Commun. ACM*, vol. 57, no. 4, pp. 51–57, Apr. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2578901>
- [10] B. Hesmans and O. Bonaventure, "An enhanced socket api for multipath TCP," in *Proceedings of the 2016 Applied Networking Research Workshop*, ser. ANRW '16. New York, NY, USA: ACM, 2016, pp. 1–6. [Online]. Available: <http://doi.acm.org/10.1145/2959424.2959433>
- [11] B. Gremillion, *Responsive Web Design: Getting the New Baseline in Web Design Right*. Smashing Magazine, 2013.
- [12] N. Muramatsu, K. Ohshima, R. Kawamura, O. C. Wei, Y. Sato, and Y. Ochiai, "Sonoliards: Rendering audible sound spots by reflecting the ultrasound beams," in *Adjunct Publication of the 30th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '17. New York, NY, USA: ACM, 2017, pp. 57–59. [Online]. Available: <http://doi.acm.org/10.1145/3131785.3131807>
- [13] W. Sun, I. Sobel, B. Culbertson, D. Gelb, and I. Robinson, "Calibrating multi-projector cylindrically curved displays for "wallpaper" projection," in *Proceedings of the 5th ACM/IEEE International Workshop on Projector Camera Systems*, ser. PROCAMS '08. New York, NY, USA: ACM, 2008, pp. 1:1–1:8. [Online]. Available: <http://doi.acm.org/10.1145/1394622.1394624>
- [14] S. Ravi, A. Raghunathan, P. Kocher, and S. Hattangady, "Security in embedded systems: Design challenges," *ACM Trans. Embed. Comput. Syst.*, vol. 3, no. 3, pp. 461–491, Aug. 2004. [Online]. Available: <http://doi.acm.org/10.1145/1015047.1015049>
- [15] A. Madhavapeddy, R. Mortier, C. Rotsos, D. Scott, B. Singh, T. Gazagnaire, S. Smith, S. Hand, and J. Crowcroft, "Unikernels: Library operating systems for the cloud," *SIGPLAN Not.*, vol. 48, no. 4, pp. 461–472, Mar. 2013.
- [16] F. Manco, C. Lupu, F. Schmidt, J. Mendes, S. Kuenzer, S. Sati, K. Yasukata, C. Raiciu, and F. Huici, "My VM is lighter (and safer) than your container," in *Proceedings of the 26th Symposium on Operating Systems Principles*, ser. SOSP '17. New York, NY, USA: ACM, 2017, pp. 218–233. [Online]. Available: <http://doi.acm.org/10.1145/3132747.3132763>
- [17] D. R. Engler, M. F. Kaashoek, and J. O'Toole, Jr., "Exokernel: An operating system architecture for application-level resource management," in *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles*, ser. SOSP '95. New York, NY, USA: ACM, 1995, pp. 251–266.
- [18] K. Elphinstone, A. Zarrabi, K. Mcleod, and G. Heiser, "A performance evaluation of rump kernels as a multi-server os building block on sel4," in *Proceedings of the 8th Asia-Pacific Workshop on Systems*, ser. APSys '17. New York, NY, USA: ACM, 2017, pp. 11:1–11:8. [Online]. Available: <http://doi.acm.org/10.1145/3124680.3124727>
- [19] Z. B. Tariq, D. M. Cheema, M. Z. Kamran, and I. H. Naqvi, "Non-gps positioning systems: A survey," *ACM Comput. Surv.*, vol. 50, no. 4, pp. 57:1–57:34, Aug. 2017.
- [20] F. Stajano, "Pico: No more passwords!" in *Proceedings of the 19th International Conference on Security Protocols*, ser. SP'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 49–81. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-25867-1_6
- [21] A. Kalso and A. Youssef, "Serverless: Beyond the cloud," in *Proceedings of the 2Nd International Workshop on Serverless Computing*, ser. WoSC '17. New York, NY, USA: ACM, 2017, pp. 6–10.
- [22] S. Kuenzer, A. Ivanov, F. Manco, J. Mendes, Y. Volchikov, F. Schmidt, K. Yasukata, M. Honda, and F. Huici, "Unikernels everywhere: The case for elastic CDNs," *SIGPLAN Not.*, vol. 52, no. 7, pp. 15–29, Apr. 2017.
- [23] B. Farinier, T. Gazagnaire, and A. Madhavapeddy, "Mergeable persistent data structures," in *Vingt-sixièmes Journées Francophones des Langages Applicatifs (JFLA 2015)*, D. Baelde and J. Alglave, Eds., Le Val d'Ajol, France, Jan. 2015.
- [24] T. Gazagnaire and V. Hanquez, "Oxenstored: An efficient hierarchical and transactional database using functional programming with reference cell comparisons," in *Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming*, ser. ICFP '09. New York, NY, USA: ACM, 2009, pp. 203–214.
- [25] T. A. Linden, "Operating system structures to support security and reliable software," *ACM Comput. Surv.*, vol. 8, no. 4, pp. 409–445, Dec. 1976.
- [26] F. McKeen, I. Alexandrovich, I. Anati, D. Caspi, S. Johnson, R. Leslie-Hurd, and C. Rozas, "Intel software guard extensions (Intel SGX) support for dynamic memory management inside an enclave," in *Proceedings of the Hardware and Architectural Support for Security and Privacy 2016*, ser. HASP 2016. New York, NY, USA: ACM, 2016, pp. 10:1–10:9. [Online]. Available: <http://doi.acm.org/10.1145/2948618.2954331>
- [27] A. M. Azab, P. Ning, J. Shah, Q. Chen, R. Bhutkar, G. Ganesh, J. Ma, and W. Shen, "Hypervision across worlds: Real-time kernel protection from the ARM TrustZone secure world," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '14. New York, NY, USA: ACM, 2014, pp. 90–102. [Online]. Available: <http://doi.acm.org/10.1145/2660267.2660350>
- [28] S. Arnautov, B. Trach, F. Gregor, T. Knauth, A. Martin, C. Priebe, J. Lind, D. Muthukumar, D. O'Keefe, M. L. Stillwell, D. Goltzsche, D. Eyers, R. Kapitza, P. Pietzuch, and C. Fetzer, "Score: Secure Linux containers with Intel SGX," in *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI'16. Berkeley, CA, USA: USENIX Association, 2016, pp. 689–703.
- [29] F. Sadri, "Ambient intelligence: A survey," *ACM Comput. Surv.*, vol. 43, no. 4, pp. 36:1–36:66, Oct. 2011. [Online]. Available: <http://doi.acm.org/10.1145/1978802.1978815>
- [30] D. Wright, S. Gutwirth, M. Friedewald, E. Vildjiounaite, and Y. Punie, *Safeguards in a World of Ambient Intelligence (The International Library of Ethics, Law and Technology)*, 1st ed.
- [31] A. Chaudhry, J. Crowcroft, H. Howard, A. Madhavapeddy, R. Mortier, H. Haddadi, and D. McAuley, "Personal data: Thinking inside the box," *Aarhus Series on Human Centered Computing*, vol. 1, no. 1, p. 4, 2015.
- [32] S. S. Rodríguez, L. Wang, J. R. Zhao, R. Mortier, and H. Haddadi, "Personal model training under privacy constraints," *CoRR*, vol. abs/1703.00380, 2017.
- [33] S. A. Ossia, A. S. Shamsabadi, A. Taheri, H. R. Rabiee, N. D. Lane, and H. Haddadi, "A hybrid deep learning architecture for privacy-preserving mobile analytics," *CoRR*, vol. abs/1703.02952, 2017.
- [34] P. Voigt and A. v. d. Bussche, *The EU General Data Protection Regulation (GDPR): A Practical Guide*, 1st ed. Springer Publishing Company, Incorporated, 2017.
- [35] N. Savage, "Going serverless," *Commun. ACM*, vol. 61, no. 2, pp. 15–16, Jan. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3171583>
- [36] P. Persson and O. Angelsmark, "Kappa: Serverless IoT deployment," in *Proceedings of the 2nd International Workshop on Serverless Computing*, ser. WoSC '17. New York, NY, USA: ACM, 2017, pp. 16–21.