Check for updates

Sinfonia: Cross-tier orchestration for edge-native applications

Mahadev Satyanarayanan[1]*, Jan Harkes[1], Jim Blakley[1], Marc Meunier[2], Govindarajan Mohandoss[2], Kiel Friedt[2], Arun Thulasi[3], Pranav Saxena[3] and Brian Barritt[3]

[1]School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, United States, [2]ARM Inc., San Jose, CA, United States, [3]Meta Inc., Menlo Park, CA, United States

The convergence of 5G wireless networks and edge computing enables new *edge-native applications* that are simultaneously bandwidth-hungry, latency-sensitive, and compute-intensive. Examples include deeply immersive augmented reality, wearable cognitive assistance, privacy-preserving video analytics, edge-triggered serendipity, and autonomous swarms of featherweight drones. Such edge-native applications require network-aware and load-aware orchestration of resources across the cloud (Tier-1), cloudlets (Tier-2), and device (Tier-3). This paper describes the architecture of *Sinfonia,* an open-source system for such cross-tier orchestration. Key attributes of Sinfonia include:

- support for multiple vendor-specific Tier-1 roots of orchestration, providing end-to-end runtime control that spans technical and non-technical criteria;

- use of third-party Kubernetes clusters as cloudlets, with unified treatment of telco-managed, hyperconverged, and just-in-time variants of cloudlets;

- masking of orchestration complexity from applications, thus lowering the barrier to creation of new edge-native applications.

We describe an initial release of Sinfonia (https://github.com/cmusatyalab/sinfonia), and share our thoughts on evolving it in the future.

KEYWORDS

edge computing, cloud computing, mobile computing, pervasive computing, IoT, Kubernetes, edge AI, 5G

1 Introduction

The roots of Edge Computing reach back over a decade (Satyanarayanan et al., 2009). It was motivated by the observation that the *consolidation* of cloud computing has negative consequences. Consolidation lengthens network round-trip times (RTT) from mobile users and increases cumulative ingress bandwidth demand from Internet of Things (IoT) devices. These negative consequences stifle the emergence of new real-time, sensor-rich applications such as *deeply-immersive augmented reality (AR)* and *streaming video analytics.* Edge computing creates the 3-tier architecture (Satyanarayanan, 2017; Satyanarayanan et al., 2019) shown in Figure 1. The *network proximity* of Tier-2 to Tier-3,

FIGURE 1
Three-tier system architecture (adapted from (Satyanarayanan et al., 2019)).

FIGURE 2
Associating devices (Tier-3) with cloudlets (Tier-2).
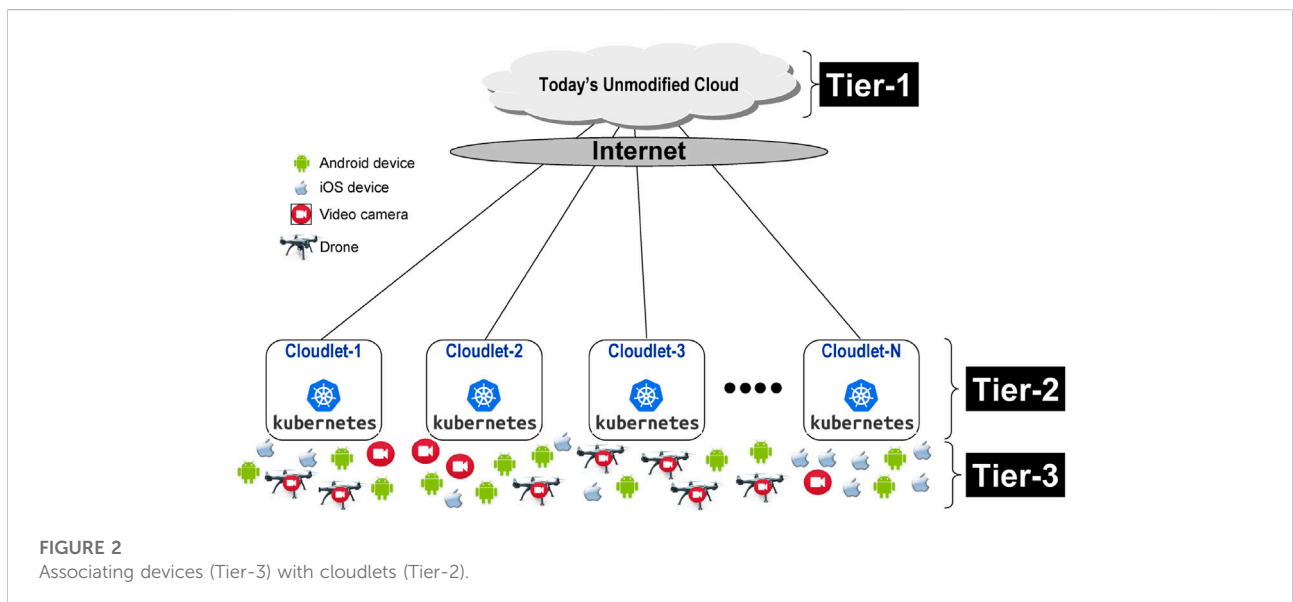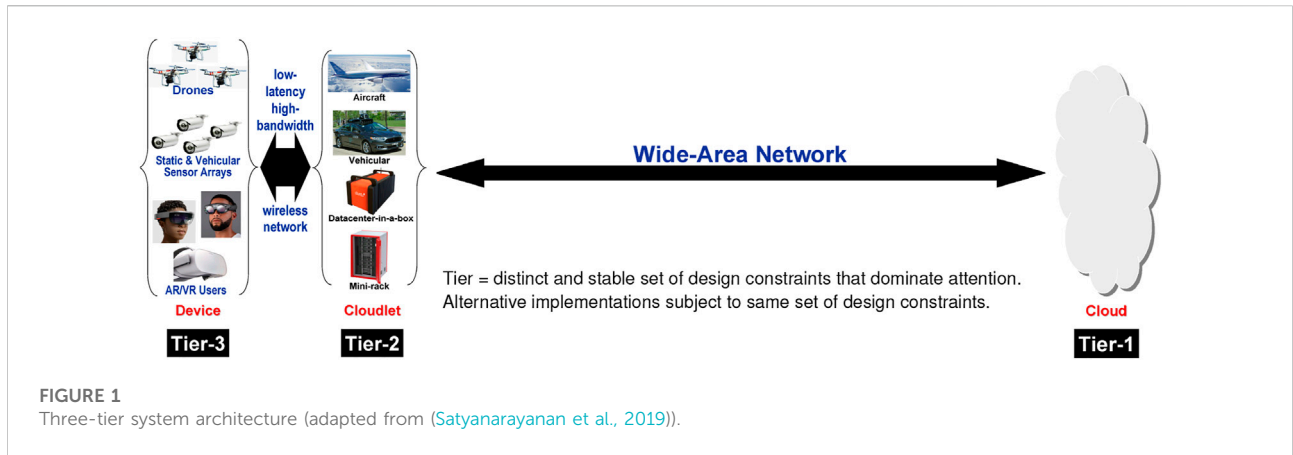
and the cloud-like computing resources available at Tier-2, together enable new *edge-native applications* that are simultaneously bandwidth-hungry, latency-sensitive, and compute-intensive. By enabling offloading of compute-intensive operations at very low latency to Tier-2, edge computing helps Tier-3 to overcome stringent design constraints such as weight, size, battery life, and heat dissipation. It also improves bandwidth scalability by avoiding excessive bandwidth demand anywhere in the system.

Sinfonia is an open source system that enables an app launched on a Tier-3 device to find and dynamically associate with its software back-end on a Tier-2 cloudlet. This association is transient, and it may involve dynamic provisioning and launching of the back-end software on the chosen cloudlet. The association is typically stable for periods ranging from a

few minutes to a few hours, and may be broken for many reasons: e.g., the app is terminated, the device moves by a large distance, the cloudlet becomes overloaded, *etc.* Sinfonia can then be used to find a new cloudlet. It can also be proactively invoked to prepare for a seamless handoff across cloudlets.

2 Problem statement

The discovery and association problem addressed by Sinfonia is framed by Figure 2. Myriad Tier-3 devices such as drones, video cameras, and wearable or hand-held devices are widely dispersed at planet scale. These are typically owned and self-managed by individuals or corporate entities. A large number of cloudlets, deployed by various service providers using diverse

business models, are also widely dispersed. For simplicity, Sinfonia assumes that every cloudlet is a Kubernetes cluster. In contrast to serverless functions, such as those supported by AWS GreenGrass (https://aws.amazon.com/greengrass/), Sinfonia is meant for settings where the back-end may involve non-trivial state. The debate between stateless and stateful implementations of distributed systems is an old one, going back over 30 years (Ousterhout, 1991). In Sinfonia, the one-time cost of cloudlet discovery and provisioning is amortized over many operations that are performed during the lifetime of the association. The cloudlet discovery problem can be simply stated as follows:

> For an app launched on a Tier-3 device, which is the optimal cloudlet for offloading operations in the near future?

This simple question has a complex answer because "optimal" is a slippery term. First, it is important to distinguish between *network proximity* and *physical proximity*. As Figure 1 indicates, edge computing requires network proximity (i.e., low latency and high bandwidth) between Tier-3 and Tier-2. Physical proximity is neither necessary nor sufficient to ensure network proximity. At the speed of light in fiber, 1 ms translates to 200 km of physical distance. With a 5G first hop below 5 ms one-way, an end-to-end RTT below 15 ms can be achieved with a cloudlet that is physically quite far away. This is low enough for deeply immersive AR. In fact, this physically distant cloudlet may be a better choice than a nearby cloudlet with a heavily-loaded ingress network. That said, it is typically the case that more distant cloudlets are reached over more network hops. Each network hop is a potential bottleneck, and adds some queueing delay. Hence, there is a weak correlation between physical and network proximity.

Second, network proximity alone cannot ensure choice of an optimal cloudlet. What matters is total end-to-end offloading performance. In AR jargon, this is referred to as *motion-to-photon latency (MTPL)*. It is a metric consisting of both transmission and processing components, and is measured at a Tier-3 device by benchmarks such as OpenRTiST (George et al., 2020). In this benchmark, MTPL consists of (a) pre-processing and encoding of a video frame captured at a Tier-3 device; (b) uplink transmission of the frame from Tier-3 to Tier-2, including any delays due to contention on a shared wireless network; (c) processing at Tier-2, including queuing delays and cache/TLB interference on shared cloudlet resources such as CPUs, GPUs and memory; (d) downlink transmission of the transformed video frame, including possible network contention delay; (e) decoding and rendering at the Tier-3 device. Items (a) and (e) are fixed for a given Tier-3 device. Network proximity only affects items (b) and (d); it does not impact item (c). The relative importance of (c) vs. (b)+(d) depends on how much processing needs to be done, vs. the amount of data that needs to be transferred. This is clearly application-dependent. In some cases (e.g., recognition of all faces in an image), it is also highly data dependent; in other cases (e.g., inference for scene classification by a deep neural network (DNN)), it is not. The available hardware resources at Tier-2 also matter. For certain algorithms (e.g., DNN inferencing), a GPU makes a huge difference; for others (e.g., a Viola-Jones face detector (Viola and Jones, 2001)), a GPU is irrelevant. In summary, MTPL is a very complex dynamic metric with no simple mapping to static metrics of processing or networking. This can lead to counter-intuitive outcomes. As an example, for DNN inferencing on the Android device at the bottom extreme left of Figure 2, Cloudlet-2 rather than the more powerful and closer Cloudlet-1 may be the better choice at the moment because Cloudlet-1's GPU is heavily overloaded. However, for a different app that does not use a GPU, Cloudlet-1 may indeed be the better choice.

Third, non-technical factors may influence the definition of "optimal." For example, app creator A may have a business relationship with cloudlet provider B. As a result, it may be more profitable for A to offload its app to a cloudlet from B rather than to one from C or D. As long as B's cloudlet can provide an adequate QoE for A's user (e.g., meet the acceptable MTPL bound for the app), technical optimality of the alternative cloudlets is irrelevant. In general, non-technical factors may span a wide range of business, legal and societal considerations. Consider, for example, the futuristic scenario presented in Figure 3. In this scenario, Ron sees no reduction in QoE but the AR gamers suffer poorer QoE. This resource allocation policy is the moral equivalent of reserved parking spaces for the disabled in everyday life. It is an example of *differentiated QoE* across users. More generally, the ability to offer differentiated QoE enables new business models. Higher-paying customers can be offered a better definition of "optimal cloudlet" than lower-paying customers.

3 Solution requirements

Based on the discussion in Section 2, as well as current trends in edge and cloud computing, we list below the requirements of a good cloudlet discovery mechanism.

First and foremost, such a mechanism must allow end-to-end control. Edge-native applications are diverse, complex, and stateful, with behavior and resource needs that defy simple characterization. New releases of an app may change its behavior and resource needs. Attempting to characterize such complexity in declarative form is unlikely to be successful. The logic to choose from available options is best expressed as code provided by the app developer. Sinfonia should center the discovery process around the invocation of such app-specific code.

Second, emerging edge computing trends should be leveraged. Clusters based on Kubernetes orchestration are

Ron is a young veteran with combat-related brain injury. He is extremely forgetful, often unable to recognize friends and relatives, and frequently neglects daily tasks. He is helped by "Clarence," a wearable edge-native assistive system. Ron takes the bus this morning to meet his friends. Clarence recognizes his bus number, and alerts him as it approaches. On the bus, Ron is able to smile and greet acquaintances because Clarence promptly whispers their names. As his stop nears, Clarence reminds him to get off, and guides him to the coffee shop where he is meeting his friends. A large group of teenagers are there, playing a deeply immersive multi-user AR game. Ron's arrival overloads the coffee shop's cloudlet. Although Clarence continues to work well, the gamers receive an alert message that they will be receiving lower-quality video and lowered crispness of game controls. Neither the gamers nor Ron are aware of how much happens under the covers to effect these changes. Ron's friends arrive soon, and one of them (Pam) is also assisted by Clarence. With Pam's arrival, the cloudlet in the coffee shop is so oversubscribed that its only recourse is to shed load. All the gamers receive a message that their game will have to be briefly interrupted. Within one minute, their shared game state has been seamlessly migrated to the cloud. Execution is resumed there with game immersiveness settings changed to reflect the much higher end-to-end latency, and the teenagers resume their AR game. After visiting for an hour, Ron and his friends leave. The gamers' back-end state is seamlessly migrated back to the coffee shop cloudlet, and they receive an alert about the improvement in immersiveness that they will shortly see. Soon they are back to enjoying their AR game in full fidelity …

FIGURE 3
An everyday AR scenario from 2030.

emerging as a *de facto* standard for cloudlets. While alternatives such as OpenStack are available, there is value in focusing an initial implementation on a single underlying abstraction. In the Kubernetes ecosystem, the Prometheus resource monitoring subsystem (https://github.com/prometheus-operator/kube-prometheus) has gained traction. Monitoring by Sinfonia for within-cloudlet orchestration should leverage Prometheus, if possible.

Third, a diversity of Tier-2 business models is emerging, with no clear winners yet in terms of specific models or companies. This may continue for a number of years, before there is clarity in this space. Major Tier-1 players themselves embrace this diversity regarding Tier-2. AWS Wavelength (https://aws.amazon.com/wavelength/), for example, extends Tier-1 abstractions to Tier-2 sites that are physically operated by telco entities such as Vodafone and Verizon (Vodafone Press Release, 2021; Ashraf, 2022). In contrast, AWS Snowball Edge (https://docs.aws.amazon.com/snowball/index.html) and Microsoft Azure Stack Hub (https://azure.microsoft.com/en-us/products/azure-stack/hub/) are physically located on customer premises. Sinfonia should embrace this diversity of deployment models, and minimally constrain current and future business models in the open-source space. In particular, it should enable participation by any entity that wishes to stand up a Tier-2 presence, even if it has no Tier-1 or Tier-3 presence. Exactly how their business models might work remains an open question. Blockchain-based micropayment solutions may play a role here. To accelerate the roll-out of edge computing and thereby catalyze the creation of new edge-native applications, Sinfonia should aim to simplify participation by any Tier-2 entity. It should support the case of a Tier-3 device owner installing a standalone cloudlet on her private wireless network, and seamlessly enabling its discovery

as the optimal cloudlet. Canonical's "Ubuntu Orange Box" was an early example of such a hyperconverged cloudlet (Vaughan-Nichols, 2014). Although this product is now discontinued, we posit that similar just-in-time cloudlets will be widely used in the future (Section 4.4).

Finally, independent of Sinfonia, many proprietary ecosystems for edge computing are emerging. Examples include AlefEdge (https://www.wearealef.com/), Equinix (https://www.equinix.com/), Nodeweaver (https://nodeweaver.eu/), and StackPath (https://www.stackpath.com/). These *walled gardens* fragment the edge ecosystem. At the same time, their presence benefits edge computing as a whole. Especially in the early days of edge computing, when coverage by major players is poor, it is valuable to have many service providers with different coverage areas and service niches. For the largest market and easiest deployability, edge-native applications should be usable with any of these walled gardens. Unfortunately, this complicates the apps, making their initial release and maintenance harder. Sinfonia has the opportunity to simplify life for the app developer. By subsuming the complexity of these diverse walled gardens and providing a single unified interface, Sinfonia can lower the barrier to entry of new edge-native applications. This is analogous to printer drivers in an operating system that allow diverse printers, each with its unique proprietary technology, to coexist; a document can be easily printed on any of them. If a developer writes code using the Sinfonia API, the resulting app should be able to use any supported walled garden with no code modifications. In this situation, Sinfonia can treat the entire orchestration process of the walled garden as a black box. It need only act as an authorized agent of the app.

FIGURE 4
Cross-tier workflow of cloudlet discovery and association.
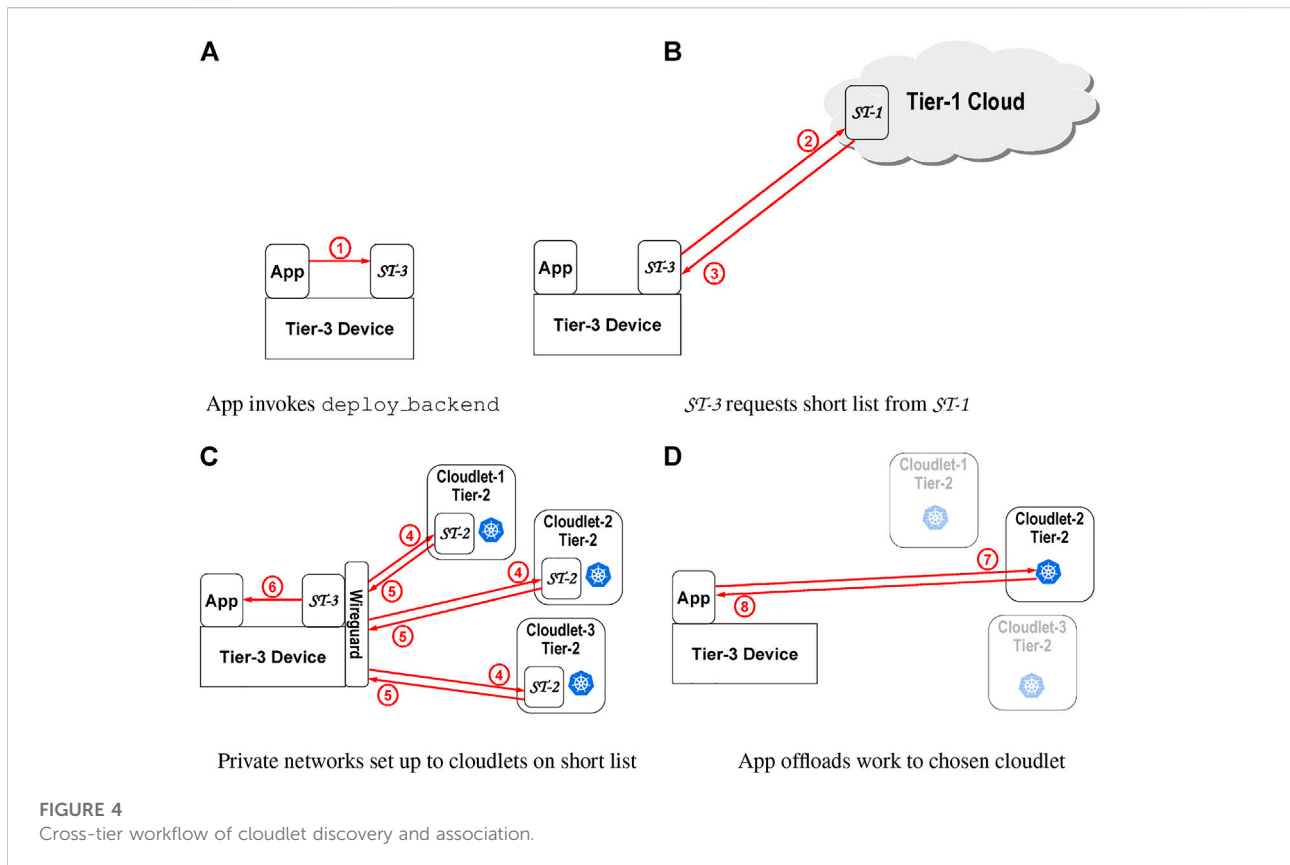
4 Sinfonia design and implementation

Befitting its role as a cross-tier mechanism, Sinfonia has code components that reside at Tier-1, Tier-2 and Tier-3. We refer to these as ST-1, ST-2 and ST-3 respectively. The workflow begins when an app at Tier-3 is launched, and requests ST-3 on its device for a backend on a cloudlet. ST-3 returns a short list of plausible targets, much like hostname lookup in DNS. We interpret "short" to mean "3 or less; " in many cases, it may just be one. Each target is a VPN endpoint to access the private IP address space where the backend is deploying. Sinfonia views all choices in the short list as "good enough" to meet the app's stated requirements.

The narrowing of the short list to one element may be done in many ways. A simple-minded app may always pick the first element. A slightly more sophisticated app may randomize its choice, with the goal of load balancing. An app that is very particular about the end-to-end properties of offloading can conduct runtime measurements (e.g., pings, test data transfers, and test operations) on each target and use that information to make the final choice. Providing this level of flexibility provides true end-to-end control. It recognizes the fact that cross-tier orchestration at Internet scale is necessarily best-effort rather than truly optimal. The constraints of latency and scalability

induce uncertainty in ST-1's knowledge of current system-wide conditions. In addition, current or past conditions may not be indicative of future conditions. Thus, Sinfonia's decision making is inherently fraught with uncertainty. In the face of this uncertainty, the next three sections describe how Sinfonia components work together to generate a short list of targets. Figure 4 illustrates the steps involved in this workflow.

4.1 Workflow at Tier-3

Linux and Android implementations of ST-3 are provided. The Sinfonia workflow begins when an app on a Tier-3 device invokes the `deploy_backend` operation on that device's ST-3 (Figure 4A). This operation could equally well be named `find_cloudlet`, because it returns a short list of "optimal" cloudlets with already-launched backends that are ready for use by this app. We focus on the discovery aspect in this section, deferring to Section 6 the discussion of alternative ways of provisioning the chosen cloudlets with the exact backend needed by this app.

An input parameter provided by the app to the `deploy_backend` call is the URL of the ST-1 to use as the root of orchestration. The developer of the app embeds this URL

in its source code. There is thus true end-to-end control of the orchestration process—it is the app developer who has ultimate control over which code will be used for the orchestration process. We expect that each major software vendor (e.g., Electronic Arts, Microsoft, Adobe, Meta, *etc.*) will have its own ST-1 at Tier-1. However, the mechanism is open-ended and flexible enough to support cross-organization sharing of a single ST-1, if that is desired. This may be a future business opportunity for Internet registrars such as NameCheap and GoDaddy.

Also provided by the app in the `deploy_backend` call is a UUID that uniquely identifies this specific version of the app. ST-3 forwards the app's request to ST-1 after adding device-specific details such as its current geo-location, its current networking environment (e.g., IP address), and device hardware attributes. Thus, ST-1 receives all the contextual Tier-3 knowledge needed for making a good decision (Figure 4B). Section 4.2 describes how ST-1 obtains the Tier-2 knowledge needed. The output returned by ST-1 to ST-3 is a short list of public IP addresses of cloudlets. ST-3 uses the Wireguard VPN mechanism (Donenfeld, 2017) to bind a private network to a target cloudlet (Figure 4C), and associates the VPN with the application. The application can perform any runtime performance tests it desires. It can then try another target until it finds the optimal one. From that point onwards, the app directs offloading requests to the chosen backend (Figure 4D). The unused backends, if any, are asynchronously garbage collected.

4.2 Workflow at Tier-1

There is considerable flexibility in how ST-1 is implemented. For simplicity, our description here assumes a monolithic implementation in the cloud, using well-known techniques for scale-out, fault-tolerance and high availability. We expect that many hyperscalers will provide such implementations as part of their standard cloud services. Not discussed here is a possible implementation of ST-1 as a collection of software-defined networking (SDN) applications that are located on the northbound interface of SDN controllers to benefit from a graph-based workflow of a Tier-2 topology. Such an SDN-based implementation would allow ST-1 to estimate network proximity from live measurements rather than relying on the geolocation of ST-3 and the weak correlation between physical and network proximity.

To handle incoming cloudlet discovery requests, ST-1 maintains a large data structure called "CloudletTable" (Figure 5). As its name implies, this is a two-dimensional array in which each row corresponds to a cloudlet known to ST-1 and identified by a UUID. The columns correspond to attributes of that cloudlet. Some of these attributes may be static. Examples include owner/operator of the cloudlet, its hostname or public IP address, its public key, its hardware configuration

(including acceleration hardware such as GPUs), and (usually) its geolocation. In the case of a cloudlet on a moving platform (e.g., truck, ship or aircraft), the geolocation will be a dynamic rather than static attribute. Other examples of dynamic attributes include current load, utilization of CPU cores and GPUs, memory utilization, and ingress and egress bandwidth utilization. Software resources (e.g., presence or absence of a large ML model) may also be viewed as attributes; in a system that caches such state, these would be dynamic rather than static attributes.

We expect the actual algorithm that produces a short list of cloudlets from CloudletTable to be a topic of intense experimentation and empirical exploration. The literature is full of theoretical results on optimal node selection for offloading in a distributed system. Few of these algorithms have real-world relevance. To simplify experimentation and easy customization of this algorithm in live deployments, it is invoked as an upcall from ST-1 to external code. The rest of ST-1 views this algorithm as a black box; all that matters is that it produce a short list of cloudlets that would be good offloading sites for the app that initiated this request. Parameters of the upcall include the CloudletTable, as well as the app and device information provided by ST-3. The default selection algorithm provided by Sinfonia uses CPU/GPU utilization and ingress/egress bandwidth as the primary variables of interest. We expect this default algorithm to evolve significantly in the light of actual usage experience with new edge-native applications.

Given the importance of CloudletTable, maintaining it becomes one of the primary responsibilities of ST-1. This maintenance is based on a push model: i.e., ST-1 relies on ST-2s to announce the availability of their cloudlets for service, and to periodically provide dynamic attributes of the cloudlets. These periodic updates effectively serve as keepalives. If ST-1 does not hear from an ST-2 for longer than a pre-determined timeout period, it assumes that the cloudlet is no longer available and deletes that row from CloudletTable. When ST-1 hears from a new cloudlet (possibly one that was dropped earlier because of a timeout), it extends CloudletTable with this new entry. This onboarding of a brand new cloudlet may involve additional steps to insure integrity and security, as discussed in Section 4.3.

4.3 Workflow at Tier-2

After a cold start of ST-1, its CloudletTable is empty. The filling of this table occurs incrementally, as ST-2s contact this ST-1 and each indicates that it is able and willing to serve as a cloudlet. An ST-1 is free to accept or decline the services of an ST-2. Conversely, an ST-2 is free to concurrently offer its services to many ST-1s. It is not obligated to offer its services to any specific ST-1. The relationship between an ST-1 and an ST-2 thus resembles a free market. We expect that all ST-1s will be widely known (e.g., possibly listed in a global web page), but

FIGURE 5
Conceptual view of CloudletTable at ST-1.
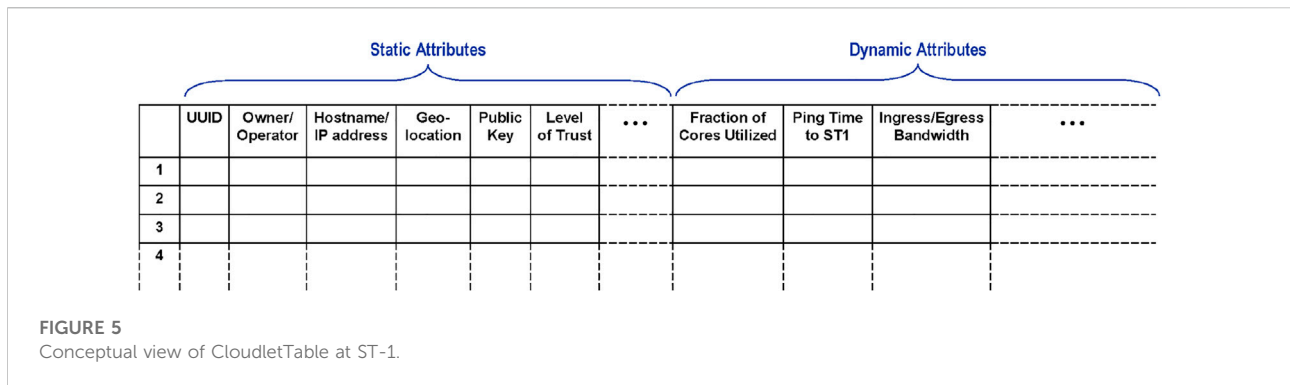
cloudlets may be more regional in ownership, management, and coverage.

If an ST-1 welcomes an ST-2, it creates an entry in its CloudletTable for that cloudlet and fills its static attributes. It also fills initial values of the cloudlet's dynamic attributes. Relevant billing-related initialization is also performed at this point. This is the point at which mutual trust is established. The procedure for establishing mutual trust may evolve as Sinfonia is developed, and may include use of mechanisms such as signed certificates and trusted enclaves. The level of trust that an ST-1 has in a cloudlet may be a static attribute in its CloudletTable entry for that cloudlet.

This dynamic, decentralized, and bottom-up binding of cloudlets to roots of orchestration draws inspiration from the early evolution of the World Wide Web. A critical catalyst to that explosive early growth was the fact that anyone, anywhere could set up a web server and publish pages on it. No centralized permission or authorization was necessary. In like vein, we wish to enable any entrepreneurial entity to easily set up a cloudlet if it sees a business opportunity. The push model of cloudlets marketing their services to ST-1s is consistent with this vision. Various ST-1s may or may not choose to do business with a specific cloudlet. Such policy decisions are best left to each ST-1, and do not have to be centralized.

We believe that Sinfonia's approach will scale well. We expect relatively few ST-1s worldwide, perhaps on the order of 10^2–10^3. This figure is consistent with the published estimate of roughly 10^5 independent software vendors (ISVs) in existence worldwide today. A few are large ISVs (e.g., Microsoft, Adobe, Amazon, Meta, Apple, Electronic Arts, *etc.*) that may each run their own ST-1. But many more are very small, and are likely to outsource their ST-1s to shared registrars. Since we are only in the early stages of edge computing rollout, it is hard to predict how many cloudlets there will be worldwide. Our current best guess is that there will eventually be on the order of 10^5–10^6 cloudlets worldwide. Since the inhabited area of the earth is estimated to be about 25×10^6 square miles, this suggests a coverage area between 25 and 250 square miles per cloudlet. This estimate may be easily off by an order of magnitude, since there may be competing cloudlets covering the same area and their distribution will almost certainly be clumpy rather than uniform. In spite of this uncertainty, the estimate of 10^5–10^6 cloudlets seems reasonable from two other viewpoints. First, there are on the order of 10^4 cities worldwide, and we expect at least a few cloudlets per city eventually. Second, there are estimated to be 6×10^9 smartphones worldwide today. If 10% are used as Tier-3 devices for edge-native applications, a fan-in of 10^2 devices per cloudlet is plausible.

After initial registration, an ST-2 is expected to periodically send ST-1 a keepalive with current values of dynamic attributes. The frequency of these keepalives is negotiated at registration. As mentioned earlier, if an ST-1 does not hear from a registered ST-2 beyond some timeout period, it deletes that ST-2 from its CloudletTable.

4.4 Just-in-time cloudlets

We envision many use cases in which a cloudlet is set up for temporary use at a work site. Examples include military settings, disaster recovery settings, outdoor construction sites, and drone operations for infrastructure maintenance. In this case, the owner of Tier-3 devices (e.g., drones) is also providing the cloudlet to be used by them. Sinfonia supports such just-in-time (JIT) cloudlets by having them configured to be on the same private wireless network as the Tier-3 devices. Using a Zeroconf discovery protocol analogous to Bonjour, ST-3s on these devices can discover the JIT cloudlet. This discovered JIT cloudlet can be added by an ST-3 to the short list returned by ST-1, most likely as the very first element of the list. Simple apps that just use the first entry in the list will always select the JIT cloudlet. More sophisticated apps can treat it as just one candidate in their final selection process. It is also possible to configure ST-3 so that it completely avoids contacting ST-1 if a JIT cloudlet is discovered. This leads to a form of static scoping: a JIT cloudlet is deemed to always be a better choice than any infrastructure cloudlet.

From the viewpoint of an app, use of a JIT cloudlet is identical to its use of any other cloudlet. No configuration or software changes to the app are necessary. It may be helpful in some use cases that the JIT cloudlet is in the same trust domain as the Tier-3 device. Especially during the early evolution of edge computing, when coverage by service providers is patchy and limited, the ability to easily setup a JIT cloudlet will be a valuable enabler for edge-native applications. For example, a JIT cloudlet in the home may enable an elderly person to use an assistive AR app on a wearable device. That elderly person can benefit from the app long before 5G coverage or telco-provided edge computing appears in his or her neighborhood.

5 Related work

The 3-Tier model of computing was only introduced in 2019 (Satyanarayanan et al., 2019), and the specific problem of cross-tier orchestration inspired by it is even more recent. We are not aware of any system prior to Sinfonia framing this problem along the lines described in this paper, and describing a system architecture to address it. Specifically, we are not aware of any prior orchestration system that specifically embodies these novel attributes of Sinfonia:

- Support for multiple vendor-specific Tier-1 roots of orchestration, providing end-to-end runtime control that spans technical and non-technical criteria.
- Use of third-party Kubernetes clusters as cloudlets, with unified treatment of telco-managed, hyperconverged, and just-in-time variants of cloudlets.
- Masking of orchestration complexity from applications, thus lowering the barrier to creation of new edge-native applications.

While Sinfonia frames and solves a new problem, it builds upon 30 years of prior work. The penalty imposed by mobility on hardware devices due to weight, size, thermal, and energy constraints was first pointed out in 1993 (Satyanarayanan, 1993). The specific innovation of overcoming this penalty by remote execution on a server was first demonstrated in 1997 by the Odyssey system (Noble et al., 1997). An important aspect of that work was adaptive selection of the optimal execution site (local, remote or split) based on runtime factors such as current network bandwidth. Flinn and Satyanarayanan (1999) broadened this concept from a narrow focus on bandwidth to also include energy. Abstracting and generalizing from these efforts, the concept of "cyber foraging" was introduced in 2001 (Satyanarayanan, 2001). That work intoduced the metaphor of mobile devices foraging for resources to overcome the mobility penalty. Today, we use the term "offloading" instead of "cyber foraging":

> *Cyber foraging,* construed as "living off the land", may be an effective way to deal with this problem. The idea is to dynamically augment the computing resources of a wireless mobile computer by exploiting wired hardware infrastructure.

Building on this foundational work, the period from 2001 to 2008 saw vigorous research activity. A detailed account of these efforts is provided in the excellent survey by Flinn (2012). The theme of cyber foraging was expanded by Balan et al. (2002). Addressing the data access requirements of remote execution for a highly mobile user were investigated in the context of *fluid replication* and *waystations* (Noble et al., 1999; Cox and Noble, 2001; Kim et al., 2002). Dynamic tradeoffs in energy and performance through self-tuned remote execution were investigated by Flinn et al. (2001); Flinn et al. (2002). Balan et al. showed how execution plans that combine local and remote execution steps could be statically identified in a declarative form called *tactics*, and then dynamically selected at runtime to optimize some desired performance-energy tradeoff. Narayanan et al. showed how history-based prediction of resource consumption by mobile applications could be combined with runtime resource monitoring in a dynamic solver to select an optimal execution strategy (Narayanan and Satyanarayanan, 2003). Goyal and Carter (2004), Su et al. (Ya-Yunn and Flinn, 2005), Ok et al. (2007) and Kristensen (2008) described different approaches to implementing infrastructure for cyber foraging. On the question of how difficult it is to manually partition legacy applications into local and remote components, Balan et al. showed that the declarative approach of tactics can simplify matters considerably (Balan et al., 2007). The emergence of cloud computing in 2006 introduced the abstraction that eventually became Tier-1 in the 2019 taxonomy. Remote execution from Tier-3 to Tier-1 now became possible. MAUI (Cuervo et al., 2010) in 2010, Odessa (Ra et al., 2011) and CloneCloud (Chun et al., 2011) in 2011, and COMET (Gordon et al., 2012) in 2012 are examples of work that explored Tier-3 offload to Tier-1 using programming language and virtual machine support.

Complementing this large body of experimental work have been many theoretical investigations into optimal selection of remote execution sites. Different conditions of network bandwidth, server load, client capabilities and other resource constraints have been explored. These efforts have exploded in the past 5 years, since the emergence of edge computing (Satyanarayanan, 2017). There has been little experimental validation of this work; typically, simulation has been used to confirm theoretical insights. Examples of this genre of work include Pang et al. (2017), He et al. (2018), Chen et al. (2019), Cheng et al. (2019), and Liu

et al. (2021). Sinfonia can be viewed as a platform that will enable experimental validation of these theoretical investigations.

6 The road ahead

At the time of writing this document, an initial version of Sinfonia has been released on GitHub. This is a bare-bones implementation, sufficient for the workflow described here but without many optimizations that we envision adding later. The goal of this release is to "fail early" — i.e., to help us gain hands-on experience with the cross-tier orchestration approach of Sinfonia, and to expose flaws in our thinking as soon as possible. Some of the key assumptions and mechanisms that we hope to validate through hands-on usage include:

- the value of end-to-end control and multiple roots of orchestration.
- the value of making the cloudlet selection algorithm a field-replaceable external module.
- the ease with which new cloudlets can be added, and the practical issues involved in sustaining a market-oriented approach to Tier-2.
- the simplification provided by Sinfonia to developers of new edge-native applications, and the QoE that it helps them achieve.
- the ability to integrate multiple walled gardens into Sinfonia, and the unique opportunities and challenges of using diverse types of cloudlets.
- implementation of cloudlet selection algorithms published in the literature, and experimental validation of their claims to optimality.

Evolution beyond the bare-bones implementation will depend on our learnings from hands-on usage, as well as the broader evolution of the ecosystem for edge-native applications. We see a number of long-term research thrusts and development initiatives arising from this work. We list these below in no particular order, mainly as placeholders to capture these thoughts for the future.

6.1 Provisioning

Section 4.1 is silent on the exact details of how provisioning of the backend for a specific app occurs on a chosen cloudlet. In practice, there are many approaches that could work well. Only usage experience will tell which of these approaches strikes the optimal balance between reliability, speed, flexibility and ubiquity. The simplest approach is static provisioning, in which the presence of

the precise backend for an application is one of the cloudlet's static attributes. A slightly more general approach is for provisioning details to be part of the app-specific knowledge possessed by ST-1, and for these details to be conveyed to the ST-2s of the chosen cloudlets. The latter may fetch the relevant bits if they are not already present. This cache state would be a dynamic attribute of the cloudlet. The most general approach is *Just-in-Time Provisioning* as described by Ha et al. (2013). In this approach, the app provides (by value) a bag of bits called an *overlay* that can be applied to a *base image* to rapidly synthesize a bit-exact backend at runtime. It pushes the end-to-end aspect of Sinfonia to its extreme, and may have particular value for a JIT cloudlet that has no connectivity to the Internet. Such a disconnected JIT cloudlet may have been discovered *via* a Zeroconf protocol.

6.2 Privacy

Implicit in the offloading of processing from Tier-3 to Tier-2 is a widening of the privacy perimeter of data captured at a device. Since Tier-1 (*via* ST-1) only participates in the control plane of offloading and not in its data plane, the privacy perimeter does not expand to Tier-1; it only expands to Tier-2. But even this limited off-device expansion of the privacy perimeter may be unsettling for some users, or illegal in some jurisdictions. One possible solution is to carefully choose the cloudlets permissible for sensitive applications, by using the static attributes of CloudletTable. A different approach is to use the concept of *denaturing* (Simoens et al., 2013; Davies et al., 2016; Wang et al., 2017), in which real-time policy-guided transformation of raw data is performed to preserve privacy. For example, the wearable devices of the gamers in the scenario of Figure 3 could black out all human faces before transmitting video to their cloudlets. ASIC hardware for face detection is cheap and fast enough for this today. By doing this, bystanders in the coffee shop are not exposed by the gamers to Tier-2. Unfortunately, this approach does not help Ron in Figure 3, because he needs help to recognize the faces of people who greet him. Fortunately, there is empirical evidence that this is indeed an acceptable tradeoff in the context of assistive technologies. Siewiorek et al. (2012), for example, report that "people will trade privacy for enhanced capability and the ability to maintain independence." The bottom line is that privacy is a huge and sensitive topic in the context of edge-native applications, and will need careful thought and attention to be paid to the policies and mechanisms that are developed.

6.3 Micropayments and trust

The "free market" relationship between Tier-1 and Tier-2 sketched in Section 4.3 suggests many opportunities for developing blockchain-based micropayment business models for edge offload. In the networking world, companies like Helium ((https://www.helium.com/)) have pioneered such business models. However, safeguarding the integrity of offloading is a challenging problem. Freedom and open-endedness bring with them questions of trust. How does an ST-1 gain confidence that a particular cloudlet has adequate integrity to return as a possible target? How does it guard against Byzantine cloudlets? One approach could possibly be for ST-1 to conduct attestation-based validation of a target cloudlet before it returns control to ST-3. Mechanisms such as secure enclaves become especially valuable in this context. The problem is even more challenging if a JIT cloudlet with no Internet connectivity is discovered *via* a Zeroconf protocol. ST-3 is then on its own; there is no ST-1 to help. Such situations may prevail in disaster recovery and military settings. Clearly, there are many research questions to be addressed in this space before the vision of a safe micropayment-based ecosystem for offloading becomes viable.

6.4 Network protocols for edge offload

The application-level protocols between Tier-3 and Tier-2 represent a unique opportunity. Today's application protocols have been developed for the Internet, where worst-case bandwidth may be low and latency high. However, the whole point of edge computing is to ensure network proximity. New edge-native applications thus have an opportunity to rethink from first principles how the assumptions of low latency and high bandwidth can benefit the protocols that they use for offloading.

In closing, the convergence of 5G wireless networks and edge computing enables many new edge-native applications that are simultaneously bandwidth-hungry, latency-sensitive, and compute-intensive. Such applications require network-aware and load-aware orchestration of resources across tiers. We have created Sinfonia, an open-source system for such cross-tier orchestration. This paper has described its driving vision, high-level assumptions and initial implementation. Sinfonia opens the door to many opportunities for experimental research and business innovation.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

This work describes collaborative research in which all the authors contributed creative insights and ideas over an extended period of time. The initial system implementation was done primarily by the second author (JH). The writing was mostly done by the first author (MS). All authors reviewed early drafts and gave valuable feedback for the final submission.

Funding

Conflict of interest

Authors MM, GM, KF was employed by ARM Inc. AT, PS, BB was employed by Meta Inc.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Ashraf, C. (2022). Verizon and AWS expand edge computing to 19 U.S. metro areas. Available at: https://www.verizon.com/about/news/verizon-aws-expand-edge-computing-19-us-metro-areas.

Balan, R., Flinn, J., Satyanarayanan, M., Sinnamohideen, S., and Yang, H. (2002). "The case for cyber foraging," in Proceedings of the 10th ACM SIGOPS European Workshop, Saint-Emilion, France, September 2002.

Balan, R., Gergle, D., Satyanarayanan, M., and Herbsleb, J. (2007). "Simplifying cyber foraging for mobile devices," in Proceedings of the 5th International Conference on Mobile Systems Applications and Services, San Juan, Puerto Rico, June 2007.

Chen, Q., Zheng, Z., Hu, C., Wang, D., and Liu, F. (2019). "Data-driven task allocation for multi-task transfer learning on the edge," in 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS).

Cheng, S., Chen, Z., Li, J., and Gao, H. (2019). "Task assignment algorithms in data shared mobile edge computing systems," in 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS).

Chun, B. G., Ihm, S., Maniatis, P., Naik, M., and Patti, A. (2011). "CloneCloud: Elastic execution between mobile device and cloud," in Proceedings of EuroSys 2011, Salzburg, Switzerland, April 2011.

Cox, L. P., and Noble, B. D. (2001). "Fast reconciliations in fluid replication," in Proceedings of the The 21st International Conference on Distributed Computing Systems.

Cuervo, E., Balasubramanian, A., Cho, D. K., Wolman, A., Saroiu, S., Chandra, R., et al. (2010). "Maui: Making smartphones last longer with code offload," in Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services, San Francisco, CA, June 2010.

Davies, N., Taft, N., Satyanarayanan, M., Clinch, S., and Amos, B. (2016). "Privacy mediators: Helping IoT cross the chasm," in Proc. of ACM HotMobile 2016, St. Augustine, FL, February 2016.

Donenfeld, J. A. (2017). "WireGuard: Next generation kernel network tunnel," in Proceedings of the 24th Annual Network and Distributed System Security Symposium, (NDSS), February 2017.

Flinn, J. (2012). Cyber foraging: Bridging mobile and cloud computing via opportunistic offload. San Rafael, California: Morgan & Claypool Publishers.

Flinn, J., Narayanan, D., and Satyanarayanan, M. (2001). "Self-tuned remote execution for pervasive computing," in Proceedings of the 8th IEEE Workshop on Hot Topics in Operating Systems, Schloss Elmau, Germany, May 2001.

Flinn, J., Park, S., and Satyanarayanan, M. (2002). "Balancing performance, energy conservation and application quality in pervasive computing," in Proceedings of the 22nd International Conference on Distributed Computing Systems, Vienna, Austria, July 2002.

Flinn, J., and Satyanarayanan, M. (1999). "Energy-aware adaptation for mobile applications," in Proceedings of the 17th ACM Symposium on Operating Systems and Principles, Kiawah Island, SC, December 1999.

George, S., Eiszler, T., Iyengar, R., Turki, H., Feng, Z., Wang, J., et al. (2020). OpenRTiST: End-to-End benchmarking for edge computing. IEEE Pervasive Comput. 19, 410–418. doi:10.1109/mprv.2020.3028781

Gordon, M. S., Jamshidi, D. A., Mahlke, S., Mao, Z. M., and Chen, X. (2012). "COMET: Code offload by migrating execution transparently," in 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12), Hollywood, CA, October 2012.

Goyal, S., and Carter, J. (2004). "A lightweight secure cyber foraging infrastructure for resource-constrained devices," in Proceedings of the 6th IEEE Workshop on Mobile Computing Systems and Applications.

Ha, K., Pillai, P., Richter, W., Abe, Y., and Satyanarayanan, M. (2013). "Just-in-Time provisioning for cyber foraging," in Proceedings of MobSys 2013, Taipei, Taiwan, June 2013.

He, T., Khamfroush, H., Wang, S., Porta, T. L., and Stein, S. (2018). "It's hard to share: Joint service placement and request scheduling in edge clouds with sharable and non-sharable resources," in 2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS).

Kim, M., Cox, L. P., and Noble, B. D. (2002). "Safety, visibility, and performance in a wide-area file system," in Proceedings of the 1st USENIX Conference on File and Storage Technologies, Monterrey, CA.

Kristensen, M. D. (2008). "Execution plans for cyber foraging," in MobMid '08: Proceedings of the 1st Workshop on Mobile Middleware, Leuven, Belgium.

Liu, C., Li, K., and Li, K. (2021). A game approach to multi-servers load balancing with load-dependent server availability consideration. IEEE Trans. Cloud Comput. 9 (1), 1–13. doi:10.1109/tcc.2018.2790404

Narayanan, D., and Satyanarayanan, M. (2003). "Predictive resource management for wearable computing," in Proceedings of the 1st international conference on Mobile systems, applications and services, San Francisco, CA.

Noble, B., Fleis, B., Kim, M., and Zajkowski, J. (1999). "Fluid replication," in Proceedings of the Network Storage Symposium, Seattle, WA, October 1999.

Noble, B., Satyanarayanan, M., Narayanan, D., Tilton, J., Flinn, J., and Walker, K. (1997). "Agile application-aware adaptation for mobility," in Proceedings of the 16th ACM Symposium on Operating Systems Principles, Saint-Malo, France, October 1997.

Ok, M., Seo, J. W., and Park, M. S. (2007). "A distributed resource furnishing to offload resource-constrained devices in cyber foraging toward pervasive computing," in Network-based information systems. Vol. 4658 of lecture notes in computer science. Editors T. Enokido, L. Barolli, and M. Takizawa (Berlin, Heidelberg: Springer).

Ousterhout, J. K. (1991). "The role of distributed state," in In CMU Computer Science: a 25th Anniversary Commemorative.

Pang, A. C., Chung, W. H., Chiu, T. C., and Zhang, J. (2017). "Latency-Driven cooperative task computing in multi-user fog-radio access networks," in 2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS).

Ra, M., Sheth, A., Mummert, L., Pillai, P., Wetherall, D., and Govindan, R. (2011). "Odessa: Enabling interactive perception applications on mobile devices," in Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys 2011), Bethesda, MD, June 2011.

Satyanarayanan, M., Bahl, P., Caceres, R., and Davies, N. (2009). The case for VM-based cloudlets in mobile computing. IEEE Pervasive Comput. 8, 4. doi:10.1109/mprv.2009.64

Satyanarayanan, M., Gao, W., and Lucia, B. (2019). "The computing landscape of the 21st century," in Proceedings of the 20th International Workshop on Mobile Computing Systems and Applications (HotMobile '19), Santa Cruz, CA.

Satyanarayanan, M. (1993). Mobile computing. Computer 26, 81–82. doi:10.1109/2.231283

Satyanarayanan, M. (2001). Pervasive computing: Vision and challenges. IEEE Pers. Commun. 8, 10–17. doi:10.1109/98.943998

Satyanarayanan, M. (2017). The emergence of edge computing. Computer 50, 30–39. doi:10.1109/mc.2017.9

Siewiorek, D., Smailagic, A., and Dey, A. (2012). Architecture and applications of virtual coaches. Proc. IEEE 100, 2472–2488. doi:10.1109/jproc.2012.2200560

Simoens, P., Xiao, Y., Pillai, P., Chen, Z., Ha, K., and Satyanarayanan, M. (2013). "Scalable crowd-sourcing of video from mobile devices," in Proceedings of the 11th International Conference on Mobile Systems, Applications, and Services (MobiSys 2013), June 2013.

Vaughan-Nichols, S. (2014). Canonical's cloud-in-a-box: The ubuntu Orange box. Available at: https://www.zdnet.com/article/canonicals-cloud-in-a-box-the-ubuntu-orange-box/.

Viola, P., and Jones, M. (2001). "Rapid object detection using a boosted cascade of simple features," in Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001).

Vodafone Press Release (2021). Vodafone uses AWS Wavelength to launch first Multi-access Edge Computing services in European region. Available at: https://newscentre.vodafone.co.uk/press-release/partnership-aws-wavel-ength-launch-first-multi-access-edge-computing-services-in-europe/.

Wang, J., Amos, B., Das, A., Pillai, P., Sadeh, N., and Satyanarayanan, M. (2017). "A scalable and privacy-aware IoT service for live video analytics," in Proceedings of ACM Multimedia Systems, Taipei, Taiwan, June 2017.

Ya-Yunn, S., and Flinn, J. (2005). "Slingshot: Deploying stateful services in wireless hotspots," in Proceedings of the 3rd International Conference on Mobile systems, Applications, and Services.