

Shaddi Hasan *Virginia Tech, VA, USA* **Amar Padmanabhan** *Databricks, CA, USA*
Bruce Davie *Systems Approach, Melbourne, Vic, Australia* **Jennifer Rexford** *Princeton, NJ, USA*
Ulas Kozat *Meta, CA, USA* **Hunter Gatewood** *Independent*
Shruti Sanadhya, Nick Yurchenko, Tariq Al-Khasib, Oriol Batalla, Marie Bremner *Meta, CA, USA*
Andrei Lee *Independent* **Evgeniy Makeev** *Meta, CA, USA*
Scott Moeller *Independent* **Alex Rodriguez, Pravin Shelar, Karthik Subraveti** *Meta, CA, USA*
Sudarshan Kandi *Independent* **Alejandro Xoconostle, Praveen Kumar Ramakrishnan** *Meta, CA, USA*
Xiaochen Tian *Independent*

Editors: Nicholas D. Lane and Xia Zhou

BUILDING FLEXIBLE, LOW-COST WIRELESS ACCESS NETWORKS WITH MAGMA

Excerpted from “Building Flexible, Low-Cost Wireless Access Networks with Magma,” 2023.
20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23).



Good Internet connectivity has become a basic necessity all over the world. Although more than one-third of the global population does not have access to the Internet [46], commercial network operators claim that today’s Internet has reached the user footprint that seems commercially viable to serve [22]. Reaching these users requires reducing the cost of providing Internet access to enable actors beyond traditional, large-scale commercial operators to build sustainable, scalable network infrastructure. Operators need effective ways to reduce costs through less expensive equipment and software and less reliance on highly skilled network administrators. At the same time, providers need ways to manage their limited network resources effectively to enable sustainable network operation – typical network policies in today’s commercial networks would be “rate limit customer C to X Mbps until they have sent Y GB in interval t_1 , then limit to Z Mbps for interval t_2 .”

Most people today connect to the Internet wirelessly using mobile devices, and two technology standards are predominant – WiFi and cellular, as specified by the 3GPP standards. WiFi networks allow even inexperienced network operators to run simple, low-cost networks on their own, with low barriers to entry thanks to use of unlicensed spectrum and a best-effort service model. However, WiFi access points operating

on unlicensed spectrum cannot generally provide efficient coverage to large geographic regions (e.g., sparsely populated rural areas), nor do WiFi networks typically offer fine-grained policies to manage resources. In contrast, cellular base stations can offer wide coverage, support many users, and connect to core networks that support more flexible policies. Yet, although cellular networks scale up to even nation-scale networks,

they do not scale *down* well: a small cellular deployment still has significant up-front costs, relying on expensive core equipment, complex protocols, and skilled administrators.

More fundamentally, we observe that choosing to use a cellular radio access network (RAN) today *forces* a network operator to make a series of decisions that deeply impact their network operations that are not inherently related to their choice



of access network technology. Indeed, the two main classes of radio access technology used today emerged as extensions to existing wireline networks – WiFi for IP networks, and cellular for legacy phone networks. These wireline networks represented vastly different design philosophies [18], the legacy of which can still be seen in today’s wireless access networks. In particular, the choice to use cellular radios binds a network operator to: (i) a specific network architecture – namely the 3GPP defined arrangement of interfaces for network management and on-path devices for policy enforcement, (ii) an ecosystem of vendors that has largely evolved to meet the needs of massive-scale telecom operators, (iii) a particular set of radio frequencies and associated regulatory requirements, and (iv) a reliance on highly skilled staff to manage specialized cellular infrastructure.

The differences between WiFi and cellular are *not* fundamental. The building blocks of network policies are common in each; what is missing is architectural support. Software-

defined networking can help address these gaps by enabling network-wide control over a distributed infrastructure, and adopting “scale out” techniques based on commodity components can reduce cost. In short, adopting and extending successful Internet and cloud approaches to scalability and management can make it possible to create a wireless access network that is both flexible and affordable.

The *Magma* project (magmacore.org) aims to achieve this by creating an open source, carrier-grade wireless networking platform that supports a wide range of deployment scenarios. Specifically, Magma’s design follows from the principle that the choice of radio access technology should not drive network architecture. Magma deployments can leverage whatever radio access technology is available and appropriate for their density of subscribers or deployment scenario. Magma achieves this goal through *access gateways* that terminate the radio-specific protocols as close to the

radios as possible. As a result, Magma allows carriers to augment an existing cellular deployment with WiFi hotspots in popular locations (e.g., athletic venues), or use LTE base stations to serve homes in rural areas, using a single core network and management platform. Ideally, new deployments could start small and grow over time. Magma achieves a “scale as you go” design through horizontal scaling of software components that run on commodity hardware, as is common in cloud-computing environments. Magma also leverages open source software components (e.g., Open vSwitch, gRPC, Kubernetes, Prometheus) commonly used in cloud settings to reduce cost, and it simplifies management by adopting software-defined networking concepts, so that a central point of control can be used to set network policies and manage subscribers. Magma adopts a hierarchical control plane to improve scalability. Magma supports only the essential features for efficient Internet access (e.g., authentication, accounting, and

per-user policies), and forgoes some complex features – including aspects of mobility.

We have seen cost savings in one deployment of 43% compared to traditional approaches due to lower operational, hardware, and software costs. Our deployment experience also illustrates how Magma scales both up and down, with one deployment supporting more than 9000 eNodeBs (base stations) in 49 American states as of August 2023 – with more than 5000 sites deployed in just five months.

MAGMA'S KEY DESIGN CHOICES

In building Magma, we took a clean-slate approach to the network core, inspired by work from software defined networking (SDN), community cellular networks, and modern cloud software architecture. Here, we briefly describe choices we made in the Magma project to achieve our goal of enabling lower-cost access networks. A more complete treatment of Magma's design is available in our recent NSDI '23 paper [26].

Magma cannot overcome the shortcomings of existing solutions simply by reimplementing a standard, 3GPP-compliant mobile core. Instead, Magma terminates the radio-specific protocols as close to the radio as possible, in *access gateways* (AGWs) connected directly to the radio access network, as shown in Figure 2. These access gateways are instrumental in handling a variety of radio technologies in a single design. AGWs are under the control of a centralized *orchestrator*, which is the central point of control for the system and maintains authoritative state related to system-wide configuration. This design helps to achieve the scaling goals of Magma, in allowing a small minimum footprint (scaling down) as well as

scaling up. A minimal Magma deployment would be a single AGW and an orchestrator. The orchestrator is typically three virtual machine instances in a cloud, while the AGW itself is a small (4-core) x86 commodity server. This is dramatically less hardware than a conventional cellular packet core.

The unit of scaling in Magma is the AGW itself: by co-locating the core network function with RAN elements, *the RAN is the bottleneck* for performance on a per-site basis. A typical eNodeB, for example, supports 96 simultaneously active users and radio channels of at most 20MHz; this channel capacity, in turn, corresponds to a peak aggregate throughput of 126Mbps [13] under ideal conditions, for a typical cell site maximum capacity of 378Mbps. Our benchmarking shows that a Magma AGW deployed even on low-end hardware (such as an Intel J3160 quad-core 1.6GHz CPU with 8GB of RAM and an Intel I210 1Gbps NIC) can easily handle this throughput requirement. Higher loads, as might be seen in a cRAN-style deployment, can easily be handled by server-class hardware.

Each AGW is a small fault domain, ensuring that the failure or upgrade of any one component affects relatively few users. In this way, Magma's architecture is similar to modern cloud systems designed to run on low-cost hardware that is prone to failure [20]. Magma adopts other ideas from cloud architectures, including the use of gRPC for communication among components, a “desired state” model for state synchronization, and a software-based, programmable data plane. While common in cloud computing deployments, these decisions deviate significantly from the way typical 3GPP networks are designed and managed.




FIGURE 1. An early Magma deployment with a small rural ISP in Peru (their first cellular site). Components (top to bottom) include (a) point-to-point wireless backhaul, (b) LTE radio and antenna, (c) ruggedized embedded PC serving as Magma AGW, (d) and solar power and battery backup for site.

In addition, Magma is able to federate with other mobile networks to support capabilities such as roaming or otherwise extend existing networks. Magma supports this use case via a *Federation Gateway* (FeG), a component analogous to the AGW but interfacing with legacy 3GPP network cores, rather than RAN elements. The FeG is a centralized, on-path device, but this serves a practical purpose: traditional MNOs prefer a single point of interconnection between their sensitive core network and “extension” networks [24].

We next take a closer look at four key design choices we made below.

Key idea #1: Abstracting the Radio Access Technology. The details of the radio access technology traditionally “leak” into the core network. To counter this, Magma identifies a core set of functions that

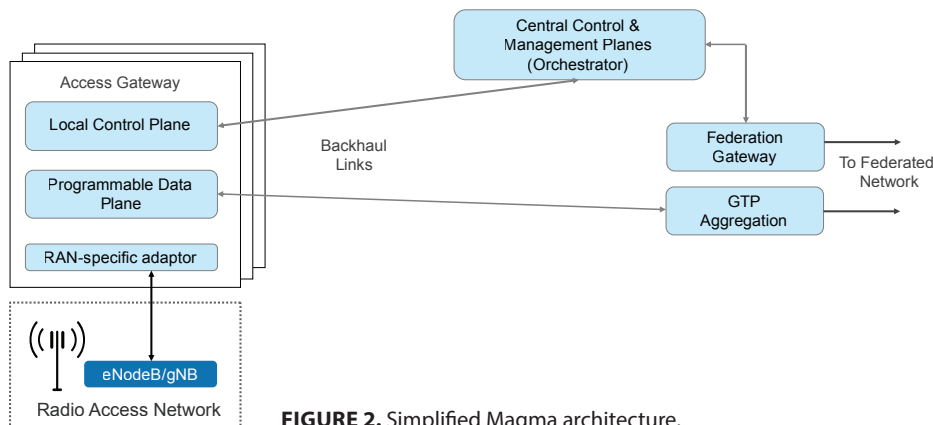
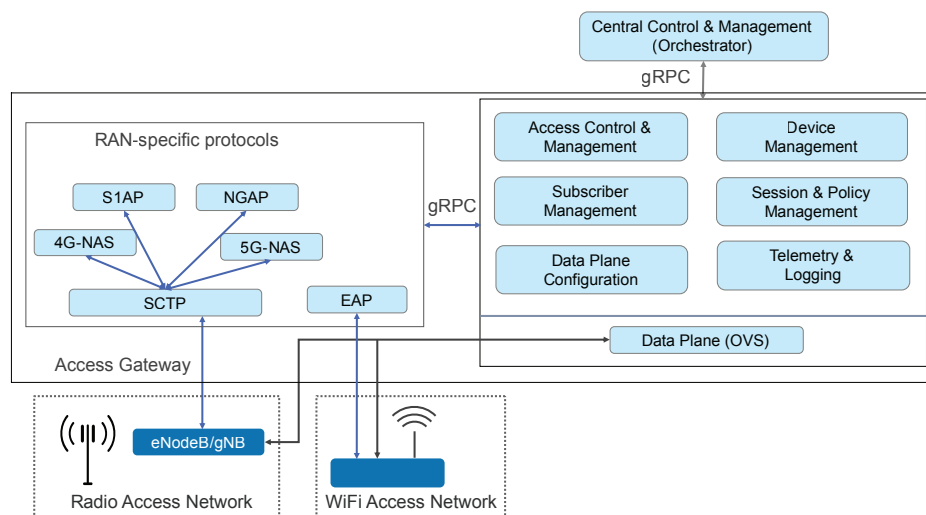


FIGURE 2. Simplified Magma architecture.

TABLE 1. Magma abstractions vs. RAN-specific versions.

Magma	LTE	5G	WiFi
Access Control/Management	MME	AMF	RADIUS AAA
Subscriber Management	HSS	UDM/AUSF	RADIUS AAA
Session/Policy Management	MME/PCRF	SMF/PCF	RADIUS AAA
Data Plane Configuration	SGW/PGW	SMF	WiFi data plane
Data Plane	SGW/PGW	UPF	WiFi data plane
Device Management	per-box configuration		
Telemetry and logging	no equivalent defined		

**FIGURE 3.** Common functions and RAN-specific protocols in the Magma architecture.

the AGW must implement for any radio technology (e.g., finding the appropriate policy for a given subscriber) and provides them in an access-technology-independent way. These functions form the heart of an AGW, as illustrated on the right side of Figure 3. Control protocols, which are specific to a given radio technology, are terminated early in technology-specific modules close to the radio. These modules, on the left of the figure, communicate with the generic functions (e.g., subscriber management, access control and management) on the right using messages that are RAN-agnostic.

Table 1 shows how the various components of 4G, 5G, and WiFi are all mapped onto a common set of Magma abstractions. The key observation here is that there are a certain set of functions that need to be performed to authenticate users, establish session state, control the data plane, and so on. Additionally, Magma adds some

generic functions that are not part of the 3GPP standards: device management and telemetry. Coupled with the SDN architecture, this simplifies the management of a large number of devices spread over a wide geographical area. We have found that considering device management and telemetry as first-class responsibilities of Magma significantly reduces the operational complexity and cost of operating access networks (See “Fixed Wireless Hotspots”). We do not claim that Magma’s decomposition of functionality (Figure 3) is fundamental, but our operational experience shows that it is useful both from an engineering perspective and for a wide range of use cases.

Key idea #2: Hierarchical SDN Control Plane Magma leverages SDN concepts to reduce operational complexity and minimize reliance on skilled staff. Rather than configuring a distributed collection of devices, providers specify network-wide

policies at the orchestrator. The orchestrator provides a central point of control and exposes a northbound API for integration with other systems (e.g., for metrics, alerting, and monitoring). However, running the entire control plane in a central controller would impose limits on the scalability of the system, so Magma adopts a hierarchical control plane, like other practical SDN systems [30,36].

In a hierarchical control-plane design, elements of the control plane that have network-wide scope are placed in the central controller. For example, the long-lived information about a subscriber is network-wide information that is created and maintained by the central controller. Conversely, runtime state associated with a UE is localized to a single AGW. For example, upon becoming active, a UE’s session state is created and managed by the local control plane of the AGW to which it is connected. Thus, much of the control plane is able to scale out with increasing numbers of base stations and subscribers, rather than by scaling up the central controller.

This division between central and local control planes roughly corresponds to the timescale of changes to the control-plane state.

In Magma, state can take one of three forms, for which Magma makes different guarantees:

- 1. Runtime state** is associated with a UE and its network activity. Runtime state is encapsulated within the AGW, and is both ephemeral and recoverable in the case of failure: a UE can simply reconnect.
- 2. Configuration state** is associated with the configuration of an entity such as an AGW or a subscriber. This state is managed centrally, and pushed asynchronously to the AGW. Examples include classes of network policy to be applied to classes of user or radio configuration to be applied by an AGW to connected RAN equipment. This state is stored durably in the orchestrator and generally changes on human timescales (i.e., minutes or hours)
- 3. Metrics state** is telemetry from Magma elements, which is captured on a best-effort basis and collected by the orchestrator.

The decision to place local control-plane functions on the AGWs, while beneficial for scalability, does introduce trade-offs. For

example, while Magma supports mobility across radios served by a common AGW, seamless mobility between AGWs would require communicating some control-plane state from one AGW to another. While many use cases can be supported without this feature, we expect to add it in the future.

Key idea #3: Fault Tolerance Via Small

Fault Domains Building a low-cost solution influences our approach to fault tolerance. Low-cost hardware is prone to failure, and so it is expected that individual components will fail. A failure of a component must affect as few users as possible (i.e., fault domains must be small) and must not affect other components. This is in stark contrast to traditional 3GPP implementations.

In a typical 3GPP implementation, the runtime state of a UE is spread among several large components (e.g., the PGW, SGW, and MME in the LTE case). In contrast, Magma localizes the runtime state of a UE to a single AGW, simplifying failure handling. Runtime state is checkpointed regularly and may be copied to a backup instance of the AGW running as a cloud service. When an AGW fails, the backup cloud instance is brought into service, and can manage connections for the affected set of UEs until the primary AGW is restarted. An AGW may continue to establish sessions to UEs even when disconnected from the orchestrator.

While it is common for a traditional cellular packet core to serve millions of subscribers, each Magma AGW is a fault domain that holds state for a relatively small number of UEs served by a small number (typically less than ten) of base stations. The failure of a single AGW would impact the set of UEs currently served by the attached base stations, but has no impact on the rest of the network or its customers. This contrasts with the much larger fault domains typical of a standard mobile core implementation.

Key idea #4: Use a programmable, software

data plane. In Magma, the data plane is responsible for (i) recognizing the flows for active sessions (traffic to and from active UEs); (ii) collecting statistics for those flows; (iii) adding and removing tunnel headers; and (iv) enforcing policies such as rate limits per subscriber. Magma uses Open vSwitch (OVS) [40], which provides a programmable data plane controlled by OpenFlow [33]. In

TABLE 2. Cost breakdown of active RAN equipment for a typical Magma deployment. Excludes site-specific passive infrastructure and backhaul costs.

Item	Unit Cost (US)	Qty	Total (US)	Notes
LTE eNodeB	\$4,000	3	\$12,000	Baicells Nova 223: 1W, 3.5GHz, 96 user, 2x2 MIMO.
AGW	\$450	1	\$450	Same as used in experiments
Accessories	\$450	3	\$1,350	18dBi sector antenna, RF cables, connectors, grounding.
RAN CapEx (per site)			\$18,760	

our experience, OVS offers entirely adequate performance in software on commodity hardware for Magma's deployments, which are typically bottlenecked by RAN capacity or backhaul, not our data plane. While OpenFlow and OVS are convenient implementation choices, other data planes could be used, assuming they are programmable and implemented entirely in software.

MAGMA IN PRODUCTION

Microbenchmarking demonstrates the viability of Magma's distributed core approach for supporting a wide range of deployment scenarios [26], and further evidence supporting Magma's approach comes from its commercial adoption. As of February 2022, twenty commercial networks were operating using Magma across eight countries in Africa, Asia, North America, and South America. These networks support a range of access modalities and policies, including providing backhaul for WiFi hotspots, fixed wireless broadband to homes and businesses, "carrier" WiFi to extend a traditional mobile operator's service to indoor WiFi, and traditional mobile broadband service. More than 250 committers have contributed to the Magma codebase.

To demonstrate how Magma is used, we worked with one of the largest commercial entities, FreedomFi, that provides support to operators deploying Magma. FreedomFi provided data to characterize two significant deployments they help operate. This data was provided to the authors in de-identified form, and only operational data (not user data) was used in our analysis.

Fixed Wireless Hotspots

One of FreedomFi's first commercial deployments was AccessParks [1], a US-based operator that provides public WiFi hotspot networks in large outdoor areas; their

deployment locations require multiple WiFi access points (APs) to provide consistent service. With the availability of CBRS spectrum, AccessParks sought to use LTE to provide backhaul to their WiFi hotspots in some of their larger deployments. End users connect to AccessParks's WiFi access points via traditional WiFi mechanisms and an existing captive portal system, and the UEs in the Magma network are fixed wireless modems that connect the WiFi APs to the Internet via Magma.

AccessParks's deployment began in December 2020 with a ten site pilot to evaluate Magma. Today, the network consists of fourteen sites providing backhaul to over 200 access points, with plans to continue expanding. Figure 4 depicts active subscribers and hourly throughput of the network.

Operational complexity. AccessParks's original Magma pilot was motivated in part by their poor experiences with the operational complexity of other commercial and open-source cellular core software in their previous two years of deployment. Although operational complexity is subjective, one quantifiable way in which it manifests is in an operator's labor costs: simpler systems should require less staff time and support to manage. Table 3 shows the results of this comparison for AccessParks. For identical access network infrastructure, AccessParks achieved a 43% reduction in per-site deployment costs using Magma compared to traditional architectures, largely driven by a reduction in support costs and engineering time for site configuration and planning.¹

¹ Unfortunately, we do not have data on ongoing maintenance costs from AccessParks; however, AccessParks' decision to use Magma for future deployments suggests it compared favorably.

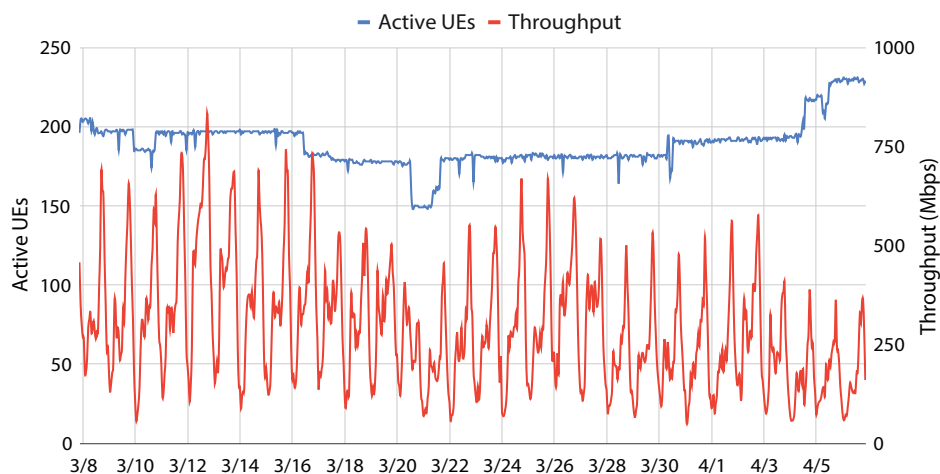


FIGURE 4. Per-hour AccessParks usage during Mar.-Apr. 2022.

TABLE 3. Comparison of per-site installed costs for AccessParks's traditional cellular system compared to Magma. Total cost per site decreased by 43%, driven primarily by Magma's reduction in operational complexity for deployment.

RAN	\$7,950	\$7,950	-	Identical RAN backup power
Core HW	\$1,200	\$300	-\$900 (-75%)	Same as used in experiments
Core SW	\$2,000	\$600	-\$1,400 (-70%)	Licenses/support
Field Eng.	\$200	\$200	-	Installation
LTE Eng.	\$5,000	\$330	-\$4,670 (-93%)	Planning, core config.
Cost/Site	\$16,350	\$9,380	-\$6,970 (-43%)	

Franchised MNO Extension

A second (and, to our knowledge, the largest) deployment of Magma is an early-stage deployment to provide a franchised, neutral host network.² This network is unique in that the physical deployment of network infrastructure is not managed by any single network operator. Instead, “micro network operators” (which include individuals, small ISPs, and enterprises) deploy LTE and 5G RAN equipment alongside Magma AGWs that have been customized by FreedomFi to support their proprietary traffic accounting and settlement system.

Services and Policy. The neutral host network is operated by FreedomFi and allows customers of incumbent MNOs to

² A neutral host network describes a business model in which a mobile network is operated by an entity for the sole purpose of providing wholesale capacity to third-party retail MNOs and MVNOs; the neutral host network operator does not have its own users, but instead enables users of its customers to use the neutral host network on a shared basis.

use this network for service. The core “policy” supported by this network is tunnelling all user traffic back to the appropriate MNO; a user’s MNO, in turn, applies their standard network policies for billing, charging, and throttling within their existing core network. The FreedomFi network provides access on a best-effort basis, with each micro network operator leveraging shared CBRS [14] spectrum (as done in the previous deployment). This service requires integrating the thousands of distributed AGWs with a partner MNO’s centralized core network, leveraging the federation capabilities described in “Magma’s Key Design Choices.”

Scale. As of this writing, this network is still in early testing, so it does not have significant user traffic. However, it still provides a useful example of how the Magma *control plane* scales with network size: even without users, Magma still manages device configuration, network monitoring, and supports interconnection with partner MNO core networks.

MAGMA CREATES AN OPEN SOURCE, CARRIER-GRADE WIRELESS NETWORKING PLATFORM THAT SUPPORTS A WIDE RANGE OF DEPLOYMENT SCENARIOS

The FreedomFi network began initial deployments in November 2021, and as of August 2023 consists of 5,011 AGWs and 9,385 eNodeBs. At peak expansion, the network added on average 150 new AGWs per week, all of which were deployed on an ad-hoc basis by micro-network operators; these AGWs are deployed in 49 states across the United States.³ FreedomFi spends about \$4,000 per month on hosting fees for their orchestrator infrastructure (including FeG) to support this entire network.

We view the rapid deployment of this network as cautious evidence for Magma’s ability to support largescale networks with unique business models. We hope to further investigate the operational dynamics of this network in future work.

RELATED WORK

Open-source LTE/5G core networks: Several projects share our goal of creating an open-source LTE/5G cellular core network [4,9,8,12]; these were preceded by similar efforts to build open 2G and 3G networks [10,11]. With the exception of OpenBTS [10] (a GSM-to-VOIP bridge), each of these focuses on implementing traditional, 3GPP-compliant, core networks.⁴ Aether [2,37] is an open-source 5G-connected edge platform, which brings together 5G connectivity and edge-cloud servers. Like Magma, Aether adopts cloud design principles, but does not refactor the core network to decouple the radio access technology.

³ The network only operates in the United States for regulatory reasons.

⁴ We note that the Magma AGW’s LTE-specific portion was originally based upon OpenAirInterface [8], as it was the most mature open-source core available at the inception of Magma’s development.

Expanding connectivity access: Many efforts have proposed or described novel solutions for expanding Internet access to under-served people [38,23,39,19,48,41,47]. Similarly, small(er)-scale network operators have a rich history providing service to especially rural communities [21], such as community networks [5,6,3,17] and small ISPs [25]. Of this extensive literature, Magma is most closely related to work on community cellular networks [27,15,44].

NextG cellular core architecture: The networking research community is actively rethinking the design of next generation network cores, including refactoring of state [43], support for public cloud deployment [34], and elastic control and data planes [42,16]. Although these works all focus on (logically) centralized core networks, the techniques described are complementary to Magma.

Other work takes a more “clean slate” approach to reimagining the cellular core. CellBricks [32] and LOCA [31] contemplate a highly federated cellular network with location privacy, and are implemented as extensions to Magma. dLTE [29] makes 4G networks more like WiFi through a decentralized design, and SoftCell [28] uses SDN principles to improve the scalability and flexibility of the packet core network. CCM [24] is a distributed cellular 2G core that enables semi-disconnected operation over unreliable rural backhaul connections; this work served as an early inspiration for Magma. Similarly, CoLTE [45] provides a lightweight core which – like an AGW – is co-located with RAN elements, but unlike Magma focuses on small, independent community networks. Magma draws on this body of work for inspiration while maintaining a backwards-compatible, standards-compliant edge to facilitate production deployment.

Open radio access networks: Several recent initiatives focus on opening up the radio access network (RAN). For example, the OpenRAN project [49] and the O-RAN alliance [35,7] develops standards that disaggregate 3GPP RANs, with open interfaces between the layers. These efforts are complementary to Magma, as they focus on the cellular interface – the part of the network before reaching Magma’s access gateway.

⁵ <https://github.com/magma/magma>

CONCLUSION

We have presented our experiences in designing and deploying Magma, an open-source platform for building access networks. The most important design decision was to terminate the RAN-specific protocols in access gateways close to the radio. This simple design decision brings many benefits: supporting diverse radio technologies, tolerating disruptions in backhaul links, using a low-cost software data plane, and scaling naturally with a hierarchical SDN control plane. In line with Magma’s goal to enable practical networks, we demonstrated that Magma can support typical deployment scenarios and discussed two large-scale commercial networks that use Magma. Importantly, Magma also scales *down*, with a small minimum footprint that supports incremental

deployment, thus filling a gap between traditional WiFi and cellular. All software artifacts for Magma are available on GitHub.⁵

Magma was designed with the primary goal of reaching under-served communities by supporting heterogeneous radio and backhaul technologies and reducing capital and operational cost. We believe that Magma is a good fit for other deployment scenarios, including enterprise 5G networks. We look forward to extending the Magma code base, and the community of contributors to the software, so the platform can evolve to serve more users. ■

Acknowledgements

We thank our NSDI shepherd Ranveer Chandra and the anonymous reviewers for their helpful feedback. We thank Boris Renski, Matthew Mosesohn, and Joey Padden

REFERENCES

- [1] AccessParks. <https://accessparks.com/>. Retrieved 7/2022.
- [2] Aether. <https://www.aetherproject.org/>. Retrieved 7/2022.
- [3] AlterMundi. <https://altermundi.net/>. Retrieved 7/2022.
- [4] free5GC. <https://www.free5gc.org/>.
- [5] Guifi.net. <https://guifi.net/>. Retrieved 7/2022.
- [6] NYCMesh. <https://www.nycmesh.net/>. Retrieved 7/2022.
- [7] O-RAN Alliance. <https://www.o-ran.org/>.
- [8] Open Air Interface. <https://openairinterface.org/>.
- [9] Open5Gs. <https://open5gs.org/>. Retrieved 7/2022.
- [10] OpenBTS. <http://openbts.org>. Retrieved 7/2022.
- [11] Osmocom. <https://osmocom.org/projects/cellular-infrastructure>. Retrieved 7/2022.
- [12] srsRAN. <https://www.srsran.com/>.
- [13] 3GPP: Evolved universal terrestrial radio access (E-UTRA); physical layer procedures. <https://port.al.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2427>.
- [14] CBRS Alliance. <https://www.cbrsalliance.org>.
- [15] Abhinav Anand, Veljko Pejovic, Elizabeth M. Belding, and David L. Johnson. VillageCell: Cost effective cellular connectivity in rural areas. 2012. *International Conference on Information and Communication Technologies and Development*, 180–189.
- [16] Arijit Banerjee, Rajesh Mahindra, Karthik Sundaresan, Sneha Kasera, Kobus Van der Merwe, and Sampath Rangarajan. 2015. Scaling the LTE control-plane for future mobile access. *ACM SIGCOMM CoNEXT Conference*, 1–13.
- [17] Luca Belli, Sarbani Banerjee Belur, Peter Bloom, Anriette Esterhuysen, Nathalia Foditsch, Maureen Hernandez, Erik Huerta, Mike Jensen, Meghna Khaturia, Michael J. Oghia, et al. December 2017. *Community Networks: The Internet by the People, for the People*. FGV Direito Rio.
- [18] David Clark. The design philosophy of the darpa internet protocols. 1988. *Symposium Proceedings on Communications Architectures and Protocols*, 106–114.
- [19] Michaelanne Dye, David Nemer, Josiah Mangiameli, Amy S. Bruckman, and Neha Kumar. 2018. El Paquete Semanal: The Week’s Internet in Havana. *CHI Conference on Human Factors in Computing Systems*, 1–12.
- [20] Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, and Paul Gauthier. 1997. Cluster-based Scalable Network Services. *ACM Symposium on Operating Systems Principles*, 78–91.
- [21] Hernan Galperin and François Bar. 2006. The Microtelco opportunity: Evidence from Latin America. *Information Technologies and International Development*, 3(2).
- [22] GSMA. 2016. Unlocking Rural Coverage: Enablers for commercially sustainable mobile network expansion, 7. <https://www.gsma.com/mobilefordevelopment/resources/unlocking-rural-coverage-enablers-commercially-sustainable-mobile-network-expansion/>.
- [23] S. Guo, M.H. Falaki, E.A. Oliver, E.A. Oliver, S. Ur Rahman, S. Ur Rahman, A. Seth, M.A. Zaharia, and S. Keshav. 2007. Very low-cost internet access using KioskNet. *ACM SIGCOMM Computer Communication Review*, 37(5):95–100.
- [24] Shaddi Hasan, Mary Claire Barela, Matthew Johnson, Eric Brewer, and Kurtis Heimerl. 2019. Scaling community cellular networks with CommunityCellularManager. *USENIX Symposium on Networked Systems Design and Implementation*, 735–750.
- [25] Shaddi Hasan, Yahel Ben-David, Max Bittman, and Barath Raghavan. 2015. The challenges of scaling WISPs. *Annual Symposium on Computing for Development*, 3–11.
- [26] Shaddi Hasan, Amar Padmanabhan, Bruce Davie, Jennifer Rexford, Ulas Kozat, Hunter Gatewood, Shruti Sanadhya, Nick Yurchenko, Tariq Al-Khasib, Oriol Batalla, et al. 2023.

for their assistance gathering deployment data for this paper. We also thank the Magma developer and user community for their important contributions, as well as Meta Connectivity for supporting the early development and deployments of Magma.

Shaddi Hasan is an assistant professor of Computer Science at Virginia Tech, VA. Prior to this, Shaddi was a software engineer at Facebook Connectivity. His research interests are at the intersection of networking, human-computer interaction, and connectivity policy. He received his PhD from UC Berkeley.

Amar Padmanabhan is a networking lead at Databricks, CA, working on scaling networks for data platforms. Prior to this, he was a lead on the Magma project at Facebook and has been involved in software networking for more than a decade. He holds a masters from the University of Wisconsin, Madison and an undergraduate degree from IIT Kharagpur, India.

Bruce Davie is the co-founder of Systems Approach, Tucson, AZ, where he produces open source books on computer networking. He is the former CTO for VMware's Asia Pacific region, the author of several Internet standards, and a Fellow of the ACM. He received his PhD from the University of Edinburgh.

Jennifer Rexford is the Gordon Y.S. Wu Professor of Engineering at Princeton University. Her research focuses on computer networking, with the larger goal of making future networks worthy of the trust society increasingly places in them.

Ulas Kozat, Hunter Gatewood, Shruti Sanadhya, Nick Yurchenko, Tariq Al-Khasib, Oriol Batalla, Marie Bremner, Andrei Lee, Evgeniy Makeev, Scott Moeller, Alex Rodriguez, Pravin Shelar, Karthik Subraveti, Sudarshan Kandji, Alejandro Xocostle, Praveen Kumar Ramakrishnan, Xiaochen Tian and Anoop Tomar are all engineers who enjoy building systems towards the mission of making real-world impact. They worked on Magma as part of Meta Connectivity and are now pursuing their passions in different roles across the industry.

APPENDIX

Acronym	Definition
MME	Mobility Management Entity
HSS	Home Subscriber Server
PCRF	Policy and Charging Rules Function
SGW	Serving Gateway
PGW	Packet Gateway
AMF	Access and Mobility Function
SMF	Session Management Function
PCF	Policy Control Function
UDM	Unified Data Management
AUSF	Authentication Server Function
S1AP	S1 Access Protocol
NGAP	Next Generation Access Protocol
SCTP	Stream Control Transmission Protocol
NAS	Non-Access Stratum
RAN	Radio Access Network
LTE	Long Term Evolution
3GPP	Third Generation Partnership Project
UE	User Equipment (a phone or other cellular client)
eNodeB	The "access point" for an LTE network
gNodeB	The "access point" for an 5G network
AGW	Access Gateway
AAA	Authentication, Authorization, and Accounting
RADIUS	Remote Authentication Dial-In User Service

Building flexible, {low-cost} wireless access networks with magma. *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, 1667–1681.

- [27] Kurtis Heimerl, Shaddi Hasan, Kashif Ali, Eric Brewer, and Tapan Parikh. 2013. Local, sustainable, SmallScale cellular networks. *International Conference on Information and Communication Technologies and Development (ICTD '13)*, ACM, 2–12, Cape Town, South Africa.
- [28] Xin Jin, Li Erran Li, Laurent Vanbever, and Jennifer Rexford. 2013. SoftCell: Scalable and flexible cellular core network architecture. *ACM SIGCOMM CoNEXT Conference*, ACM, 163–174.
- [29] Matthew Johnson, Spencer Sevilla, Esther Jang, and Kurtis Heimerl. 2018. dLTE: Building a more WiFi-like cellular network (Instead of the other way around). *ACM Workshop on Hot Topics in Networks*, ACM, 8–14.
- [30] Teemu Koponen, Keith Amidon, Peter Bolland, Martin Casado, Anupam Chanda, et al. 2014. Network virtualization in multi-tenant datacenters. *USENIX Symposium on Networked Systems Design and Implementation*, 203–216.
- [31] Zhihong Luo, Silvery Fu, Natacha Crooks, Shaddi Hasan, Christian Maciocco, Sylvia Ratnasamy, and Scott Shenker. LOCA: A location-oblivious cellular architecture. 2023. *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, 1621–1646.
- [32] Zhihong Luo, Silvery Fu, Mark Theis, Shaddi Hasan, Sylvia Ratnasamy, and Scott Shenker. 2021. Democratizing Cellular Access with CellBricks. *ACM SIGCOMM*.
- [33] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. 2008. OpenFlow: Enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74.

- [34] Binh Nguyen, Tian Zhang, Bozidar Radunovic, Ryan Stutsman, Thomas Karagiannis, Jakub Kocur, and Jacobus Van der Merwe. 2018. ECHO: A reliable distributed cellular core network for hyper-scale public clouds. *Annual International Conference on Mobile Computing and Networking*, ACM, 163–178.
- [35] O-RAN Alliance. O-RAN: Towards an Open and Smart RAN, October 2018. White paper, https://assets-global.website-files.com/60b4ffd4ca081979751b5ed2/60e5afb502810a0947b3b9d0_O-RAN%2BWP%2BFInal%2B181017.pdf
- [36] OVN Architecture. 2021. <https://www.ovn.org/support/dist-docs/ovn-architecture.7.pdf>.
- [37] Guru Parulkar. Aether: An open source platform for private 5G connected edge cloud-as-a-service. 2020. Keynote presentation at ONF Spotlight on 5G Connected Edge Cloud for Industry 4.0 Transformation, <https://www.youtube.com/watch?v=Zn7FZyiw5KM&t=4s>.
- [38] Rabin Patra, Sergiu Nedeveschi, Sonesh Surana, Anmol Sheth, Lakshminarayanan Subramanian, and Eric Brewer. 2007. WiLDNet: Design and implementation of high performance WiFi based long distance networks. *USENIX Symposium on Networked Systems Design and Implementation*.
- [39] Alex (Sandy) Pentland, Richard Fletcher, and Amir Hasson. 2004. DakNet: Rethinking connectivity in developing nations. *Computer*, 37(1):78–83.
- [40] Ben Pfaff, Justin Pettit, Teemu Koponen, Ethan Jackson, Andy Zhou, Jarno Rajahalme, Jesse Gross, Alex Wang, Joe Stringer, Pravin Shelar, et al. 2015. The design and implementation of open vSwitch. *USENIX Symposium on Networked Systems Design and Implementation*, 117–130.
- [41] Thomas Pötsch, Salman Yousaf, Barath Raghavan, and Jay Chen. 2018. Zyxt: A network planning tool for rural wireless ISPs. *ACM SIGCAS Conference on Computing and Sustainable Societies*, 1–11.

- [42] Zafar Ayyub Qazi, Phani Krishna Penumarthi, Vyas Sekar, Vijay Gopalakrishnan, Kaustubh Joshi, and Samir R. Das. 2016. KLEIN: A minimally disruptive design for an elastic cellular core. *ACM Symposium on SDN Research*. ACM, 1–12.
- [43] Zafar Ayyub Qazi, Melvin Walls, Aurojit Panda, Vyas Sekar, Sylvia Ratnasamy, and Scott Shenker. 2017. A high performance packet core for next generation cellular networks. *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pages 348–361. ACM.
- [44] Rhizomatica. <http://rhizomatica.org/>, 2013. Retrieved 4/2013.
- [45] Spencer Sevilla, Matthew Johnson, Pat Kosakanchit, Jenny Liang, and Kurtis Heimerl. 2019. Experiences: Design, implementation, and deployment of CoLTE, a community LTE solution. *ACM MobiCom*, 1–16.
- [46] Statista. Internet usage worldwide—statistics & facts. <https://www.statista.com/topics/1145/internet-usage-worldwide/#dossierKeyfigures>.
- [47] Sonesh Surana, Rabin K. Patra, Sergiu Nedeveschi, Manuel Ramos, Lakshminarayanan Subramanian, Yahel Ben-David, and Eric A. Brewer. 2008. Beyond pilots: Keeping rural wireless networks alive. *USENIX Symposium on Networked Systems Design and Implementation*, volume 8, 119–132.
- [48] William Waites, James Sweet, Roger Baig, Peter Buneman, Marwan Fayed, Gordon Hughes, Michael Fourman, and Richard Simmons. 2016. RemIX: A distributed internet exchange for remote and rural networks. *Workshop on Global Access to the Internet for All*, 25–30.
- [49] Mao Yang, Yong Li, Depeng Jin, Li Su, Shaowu Ma, and Lieguang Zeng. 2013. OpenRAN: A SoftwareDefined RAN Architecture via Virtualization. *ACM SIGCOMM Computer Communication Review*, 43(4):549–550.