EDITOR: Schahram Dustdar, dustdar@dsg.tuwien.ac.at

DEPARTMENT: INTERNET OF THINGS, PEOPLE, AND PROCESSES

Connectivity Technology Selection and Deployment Strategies for IoT Service Provision Over LPWAN

Pantelis A. Frangoudis, Christos Tsigkanos, Schahram Dustdar, *Distributed Systems Group, TU Wien, 1040, Vienna, Austria*

The plethora of available IoT connectivity technologies makes selection of the most suitable and affordable offering challenging for IoT service providers. This is a typical problem when considering deployments over Low-Power Wide-Area Networks (LPWAN). Each technology has different implications as to how edge, fog, and cloud resources can be utilized to provide an end-to-end service. The connectivity decision directly affects the initial investment, operating expenditure, level of control desired over the underlying infrastructure, and the resulting managerial overhead. This article provides a generic cost model and framework to assist IoT service providers towards cost-aware LPWAN technology selection and dimensioning decisions. Our framework captures key cost factors and considers both network and compute infrastructure elements. Considering a monitoring IoT application as a reference, we analyze various deployment strategies enabled by two major IoT connectivity technologies, namely LoRaWAN and NB-IoT, quantifying their associated costs and distilling general deployment guidelines.

The Internet of Things (IoT) has ushered a revolution, bringing novel types of Internet computing systems integrating heterogeneous devices, computing infrastructure, and networking technologies. Connectivity technologies for IoT devices are themselves diverse and each is tailored to specific application classes, exhibits diverse characteristics and targets different scenarios. As such, selecting the most affordable and suitable offering is quite a challenge, a position that IoT service providers (SP) and system designers are often found in. Furthermore, each available technology has different implications as to how edge, fog, and cloud compute resources can be utilized to provide an end-to-end service; at the same time, system dimensioning factors have to be strongly appraised. Those entail design decisions that take into account initial and operating expenditure, level of control desired over the network and compute infrastructure that powers the IoT service and naturally, the associated managerial overheads.

In this space, low-power wide-area networking (LPWAN) has been established as the key connectivity solution for IoT scenarios where long range, very low energy consumption, and low cost matter. LPWAN protocols are designed to connect massive numbers of low-end battery-powered devices, such as sensors attached to micro-controller units for delay-tolerant applications that require low throughput per device. Typical examples can be found within environmental monitoring, smart agriculture and smart city services.

As multiple competing LPWAN technologies exist – each with different technical characteristics, business model, and supported deployment models – this article provides a cost-driven comparison of the two major IoT connectivity technologies, namely LoRaWAN and Narrowband IoT (NB-IoT), from the system designer's or IoT SP's viewpoint. To this end, we analyze different deployment scenarios and connectivity

1089-7801 © 2021 IEEE
Digital Object Identifier 10.1109/MIC.2020.3031104
Date of current version 3 February 2021.

options from a cost perspective, over a multi-level monitoring architecture which we use as our reference IoT service. We introduce a framework which is generic and captures the key cost-contributing factors, ranging from the end devices to networking and computation. For the latter, we draw from experimental results on the capabilities of typical edge and cloud compute infrastructure to support the target IoT applications, and use this information for compute resource dimensioning. After applying our framework to quantitatively compare the different scenarios, we conclude by distilling designer guidelines, useful for actionable decision making regarding LPWAN technology selection and IoT deployment strategies.

BACKGROUND AND RELATED WORK

LPWA networks are widely deployed worldwide, and are projected to reach more than 1.7 billion connections by 2023; LoRaWAN and NB-IoT, two major LPWAN technologies today, are expected to account for 86% of them.[1] LoRaWAN is an open standard that specifies the communication protocols on top of the proprietary LoRa physical layer. Typically, end-devices communicate with gateways, which relay messages to network servers responsible for MAC-layer operations (packet deduplication, downlink transmission scheduling, etc.), security functions, forwarding end-device data to applications, etc. LoRaWAN operates in unlicensed spectrum (Industrial, Scientific, and Medical Band or short ISM bands), thus allowing for the deployment of private networks without any operator's involvement. In this case, the SP is in charge of installing and operating all network equipment, including gateway devices. On the contrary, NB-IoT is a cellular-based alternative and works via traditional subscription models. It is standardized by the 3GPP and can coexist with 4G LTE and 5G networks. End devices attach to cellular base stations and have IP connectivity. For a deeper technical overview of the two technologies, the reader is referred to the works of Haxhibeqiri *et al.*[2] and Wang *et al.*[3]

The attractive features of LPWAN have recently received significant research attention. Gu *et al.*[4] present a survey of seven key LPWA technologies, but do not study cost-related aspects. Mekki *et al.*,[5] on the other hand, narrow down their scope to LoRa, NB-IoT, and Sigfox, and provide a technical comparison and a brief summary of different deployment models and cost aspects. The latter, however, are limited to a comparison of base station and end-device costs, and the wide variety of different deployment scenarios and their associated costs are not analyzed. Del Campo *et al.*[6] analyze technical, functional, and cost aspects to provide guidelines for selecting the most appropriate LPWAN technology, and apply their analysis to different IoT use cases. Interestingly, they compare 13 different technologies in terms of deployment cost, considering device, gateway, and connectivity expenses. On the contrary, while focusing primarily on NB-IoT and LoRaWAN, we delve more deeply into different pricing options and deployment scenarios that each technology enables, and provide a generalizable cost model that goes beyond connectivity aspects to also capture computation-related costs.

We should further note that our work addresses the full device-to-cloud continuum, and edge computing in particular. Gusev and Dustdar[7] discuss the evolution from cloud-based IoT service provision to one that takes advantage of edge computing resources. Edge computing enables data processing closer to the source, i.e., IoT devices, which reduces latency[8,9] and saves on backhaul network resources.[10] This can also enhance privacy[11,12] by on-device or on-premise data processing, multistaged filtering, and privacy-aware data and service placement. Notable developments have also taken place in the telco space, with the specification of a family of standards around Multi-access Edge Computing (MEC) driven by ETSI.[13] MEC is important as an enabler for IoT services.[14] IoT application components can be deployed at telco-operated edge data centers via ETSI MEC interfaces. Since MEC standards are driven by the telecom industry, it is typically assumed that traffic delivered to/from MEC applications originates from or terminates at end devices connected to the operator's cellular network. As such, MEC appears as a more natural fit for integration with NB-IoT. However, traditional telecom operators have started offering LoRaWAN connectivity plans,[15] which has motivated research on the integration of MEC with other LPWAN technologies.[16,17]

REFERENCE IOT SERVICE ARCHITECTURE

In a parallel strand of research, we have developed a modular application design tailored to edge-intensive IoT monitoring services. Our architecture builds around the concept of a *monitor*. The key feature of a monitor is its capability to verify temporal logic properties at runtime, applied to event streams originating from IoT devices. Furthermore, the output of a monitor can be fed as input to another monitor that is verifying a different property. This way we can compose complex runtime verification pipelines in a distributed way

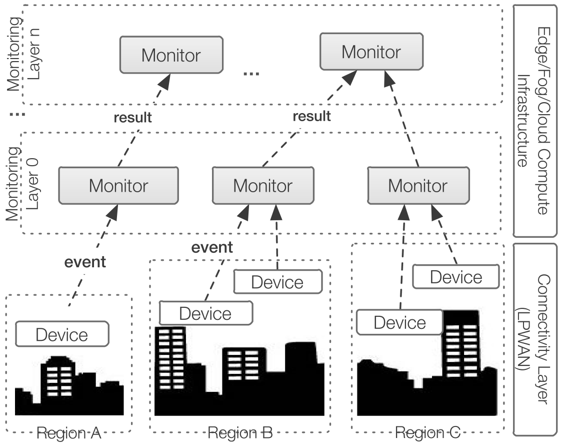

FIGURE 1. Reference IoT service architecture.

and with the versatility to capture very diverse applications. We have explored different use cases in the contexts of environmental monitoring, smart parking, and others. Depending on the application requirements and the infrastructure capabilities, a monitor can be can be deployed anywhere in the device-to-cloud continuum. The design is recursive, in the sense that with the monitor as a building block with well-defined and unified interfaces, we can compose arbitrary multilevel monitoring structures, where each level expresses a different level of abstraction. The application provider defines the properties to be verified, the monitor pipelines and hierarchy, and the monitor placement over edge/cloud/fog infrastructures.

Figure 1 presents our IoT service architecture applied in the context of smart city monitoring. IoT devices deployed across a city and organized in regions sense their environment and emit events using LPWAN technology. The events are delivered to edge monitors (level 0) which verify properties (e.g., related with air pollution levels) per region. The verification results of level-0 monitors are propagated as input to level-1 monitors, which can also be executed at edge, fog, or cloud compute infrastructures. In our implementation, each monitor exposes REST API endpoints to ingest property input (raw events or the output of lower level monitors) and for other control actions.

We select this architecture as a reference for the cost analysis we present in this article, arguing that it is fairly generic: It can express different topologies, such as service chains or hierarchy trees, while the core verification engine of the monitor can be replaced with other tools such as stream processing engines without affecting the generality of our design and methodology. Dimensioning decisions concerning the amount and type of compute resources to be allocated (e.g., edge computing devices, cloud VMs) can take advantage of workload profiling. This is well suited for predictable workloads, as is the case for many IoT applications where the number of IoT devices and the rate at which they generate events is known, controllable, or easy to estimate accurately. With this knowledge, the SP can carry out an initial planning step where, via experiments, it can measure the processing capacity of each compute unit of interest, for a given application workload. In our case, we have performed extensive experiments with different hardware technologies, including specific Single-Board Computers (SBC) and cloud VMs, to derive accurate figures of the workload that each monitor can handle when deployed on different hardware technologies. Then, knowing the event processing throughput of a monitor, we can calculate the amount of compute resources necessary and, in turn, the (monetary) cost incurred in a straightforward way.

CANDIDATE TECHNOLOGIES AND DEPLOYMENT SCENARIOS

We evaluate two candidate technologies for device connectivity, namely LoRaWAN and NB-IoT. Their different nature has implications on IoT service design. We identify the following key deployment choices, which we use to extract and evaluate specific technology selection and service deployment strategies, also accounting for different pricing models.

- Connectivity technology: LoRaWAN versus NB-IoT.
- Gateway ownership model: Service provider, community, or mobile network operator driven.
- Compute infrastructure use: Deployment on edge devices controlled by the SP, use of (telco- or SP-operated) edge clouds, or use of global cloud service providers.

Table 1 provides an overview of the different scenarios that we consider in this article and their distinctive characteristics along the dimensions of connectivity technology, network infrastructure ownership, and pricing.

LoRaWAN-Based Deployments

Due to its operation in ISM frequencies, LoRaWAN comes license-free. This gives flexibility to the SP in terms of network infrastructure deployment and use, and gives rise to various alternative configurations, each showing a different tradeoff between infrastructure control and setup and maintenance costs.

TABLE 1. Deployment scenarios and their properties.

Scenario	Connectivity		Gateway ownership			Netw. subscription model			Compute infrastructure		
	LoRaWAN	NB-IoT	SP	Community	MNO	Free	Flat rate	Monthly	Edge dev.	Cloud DC	Edge/ MEC DC
LoRaWAN-aio	✓		✓			✓			✓		
LoRaWAN-split gateway	✓		✓			✓			✓		
LoRaWAN-cloud based	✓		✓			✓			✓	✓	
LoRaWAN-community based	✓			✓		✓				✓	
LoRaWAN-managed stack	✓		✓					✓		✓	
LoRaWAN-MNO-cloud	✓				✓			✓		✓	
LoRaWAN-MNO-MEC	✓				✓			✓		✓	✓
NB-IoT-cloud-flat		✓			✓		✓			✓	
NB-IoT-MEC-flat		✓			✓		✓			✓	✓
NB-IoT-cloud-monthly		✓			✓			✓		✓	
NB-IoT-MEC-monthly		✓			✓			✓		✓	✓

SPs with the expertise to setup LoRaWAN network infrastructure (gateways and network stacks) and the willingness to handle the managerial overhead may opt for a solution with minimal dependence on external entities, making it possible to host everything in-house. They can dispense of the use of cloud infrastructure to host application components and deploy everything on top of SBCs or small edge clusters thereof. If, additionally, the target is to minimize the SBC equipment costs, an all-in-one installation is feasible, where (i) the LoRaWAN gateway hardware, (ii) the LoRaWAN network and application server stack, and (iii) the IoT application that consumes device-generated data, are all hosted on a single SBC device. Otherwise, the gateway device can be kept as lightweight as possible, where a dedicated device is used to host the gateway hardware, and the LoRaWAN network stack and IoT service components are deployed on top of different SBCs or cloud VMs. We denote the above three scenarios as "LoRaWAN-aio,""LoRaWAN-split gateway,"and "LoRaWAN-cloud based." In all three scenarios, there are no connectivity fees. Additionally, offerings have emerged where the devices and gateways are managed by the SP, but a third party offers the network stack on a monthly fee which depends on the number of devices (the higher the number of devices, the lower the per-device monthly cost). This setup is suitable for SPs who wish to control network coverage and device connectivity, but do not wish to cope with LoRaWAN network management. We term this model "LoRaWAN-managed stack."

An alternative path can be followed by the SP if the latter aims minimizing initial setup costs (for gateway procurement) and the operational cost of managing gateways. We have identified two options here. Either the SP connects its devices to community-operated gateways, as is the case for TheThingsNetwork (https://www.thethingsnetwork.org), or to commercial networks that are run either by Mobile Network Operators (MNOs) or other operators. A disadvantage of the former is potentially limited coverage, while the latter (if available) come with a per-message or monthly fee. These scenarios further limit the SP's deployment options regarding where to place IoT service components: While the full gateway ownership model allows the SP to instantiate its application anywhere in the device-to-cloud continuum (e.g., on-device, on-premise, or in edge or remote clouds), in the gateway-less options application components are typically placed in remote clouds. On the positive side, the SP need not care about tedious gateway

TABLE 2. Costs associated with LoRaWAN and NB-IoT deployments.

Network infrastructure costs	
LoRaWAN end device cost[a]	15 EUR
NB-IoT end device cost[a]	25 EUR
LoRaWAN gateway cost	200 EUR (setup) + 5 EUR/mo (operational expenses)
Communication costs	
LoRaWAN managed network stack	{100, 175, 250, 500} EUR/mo for up to {250, 500, 1000, 2500} devices, resp.
LoRaWAN MNO subscription	1 EUR/device/mo; unlimited data.
NB-IoT monthly subscription	1.24 EUR/device/mo for a 500 KB data plan (shared across all devices); 0.004836 EUR/KB excess fees.
NB-IoT SIM activation fee	2.5 EUR/device.
NB-IoT one-off flat fee	10 EUR/device (500 MB over 10 years; re-booking needed if data usage exceeded).
Gateway backhaul connectivity	5 EUR/mo; 4G access, 2 GB data plan; 0.01 EUR/KB excess fees.
Compute costs	
Cloud	18 EUR/vCPU/mo
MEC	36 EUR/vCPU/mo
SBC	50 EUR (setup) + 5 EUR/mo (operational expenses)

[a] *Includes the radio and MCU platform.*

configuration, health monitoring, and maintenance procedures. The gateway-less scenarios that we study are denoted as "LoRaWAN-community based," "LoRaWAN-MNO-cloud," and "LoRaWAN-MNO-MEC," the latter referring to a case where MEC infrastructure is available to SPs to instantiate IoT application components.

Narrowband (NB-) IoT-Based Deployments

NB-IoT is gaining momentum as a high-coverage cellular-based LPWAN connectivity solution. The IoT application provider equips end devices with NB-IoT radio modules and SIM cards and subscribes with an MNO. End devices communicate with base stations, thus there is no need for the SP to deploy and operate gateway hardware. This simplifies service management and drives setup costs down. On the other hand, initial per device activation fees (including SIM procurement costs) may drive setup costs up. Moreover, monthly fees may apply. The IoT application is then launched on centralized or edge cloud infrastructures. The latter is possible if the MNO provides MEC facilities and is more relevant for applications requiring low-latency and/or high reliability. For delay tolerant applications, this option may be less attractive, since MEC offerings are expected to be more expensive due to the inherent resource scarcity of MEC resources and the special features MEC offers, such as radio network and location awareness, and traffic offloading capabilities. Pricing-wise, any of the above deployment models can be coupled either with a flat up-front fee per device or with a monthly fee with a shared data plan across all devices. To recap, we study the following NB-IoT scenarios: "NB-IoT-cloud-flat" (centralized cloud deployment; flat rate), "NB-IoT-MEC-flat" (edge application components hosted at MEC servers; flat rate), "NB-IoT-cloud-monthly," and "NB-IoT-MEC-monthly" (cloud versus MEC deployment, respectively; monthly fee).

FRAMEWORK FOR DEPLOYMENT COST ESTIMATION

Cost Factors

We can already identify the key factors influencing the cost of such an IoT deployment: (i) end devices, (ii) LPWAN gateways, (iii) device connectivity, (iv) backhaul connectivity to transport data across IoT service tiers, and (v) physical or virtual compute resources to host gateways, the LPWAN stack, and IoT service components. Table 2 provides details on the values we used for each of these costs, based on offerings as of September 2020.

End Devices

A typical LPWAN-capable end device is composed of a microcontroller, a radio module including an antenna, a battery and the necessary sensing equipment. The latter is excluded from our analysis as it is application-specific. Due to the increased protocol and hardware complexity, and the need for a standard or an embedded SIM, NB-IoT modules are more expensive. Taking into account the cost figures reported by del Campo *et al.*[6] combined with current pricing information based on up-to-date hardware offerings, we set the device cost to 15 and 25 EUR for LoRaWAN and NB-IoT, respectively.

Gateways

This applies only to the LoRaWAN-based deployment models where the service provider owns and manages gateways. The availability of low-cost SBCs such as Raspberry Pis and open-source gateway software can drive the cost to well below 200 EUR, the figure we have used in our analysis. In fact, we have assembled such gateways for our implementation and experiments using Raspberry Pi 3 Model B devices (less than 30 EUR) and compatible LoRaWAN concentrator hardware (approximately 120 EUR).

Device Connectivity

Network operator-driven LoRaWAN offerings (LoRaWAN-MNO-cloud and LoRaWAN-MNO-MEC models) have recently emerged in some countries, following a pricing model similar to standard cellular offerings, i.e., a monthly per device cost. In our analysis, we have used the cost figures of a large global network operator, where there is no limit on the volume of data uploaded. However, strict duty cycling limits apply by the local regulator in each country; typically, an end device cannot be active more that 1% of the time. For the managed stack case, we have used the price figures from a dedicated LoRaWAN network server provider, which applies a stepwise pricing model as described in Table 2. There is a maximum number of devices supported; if exceeded, additional subscriptions need to be acquired.

Regarding NB-IoT, the monthly subscription model is straightforward. We have used tariff figures from a European ISP, which offers a shared data plan. The customer pays a flat fee of 1.24 EUR/device/month and for a maximum of 500 KB of traffic shared across all devices under the same data plan. Excess fees apply if this volume is exceeded. Also, it is typical to charge an initial activation fee for each device connection. An alternative model that has emerged—albeit with very limited examples currently—entails a one-off flat per device fee. In this case, the customer pays once for a multi-year period and for a specific traffic volume. If this is exceeded, the customer needs to re-book.

Gateway Backhaul Connectivity

When gateways are managed by the service provider, the collected data may need to be forwarded to remote IoT service components over metered connections such as 4G LTE. This applies to the LoRaWAN-cloud-based scenario, where we have set backhaul connectivity costs to 5 EUR/month for a data plan of 2 GB. Excess costs per KB apply. We further assume that for LoRaWAN-aio and LoRaWAN-split-gateway there are no backhaul costs (e.g., the gateway is connected via an unmetered high capacity link).

Compute Infrastructure

The cost of compute resources varies along the Cloud/Fog/Edge continuum, as is also the case for the respective processing capacity, required management overhead, and reliability levels. The LoRaWAN deployment schemes where the gateway is owned by the service provider allow for deployment of private networks where both the LoRaWAN network stack and IoT service components run on top of edge devices. In this case, SBCs can be deployed at locations where the service provider has access (e.g., lamp posts, municipal buildings) to deliver smart city services such as urban sensing. This burdens the SP with infrastructure operational costs; we have set these costs to 5 EUR/month/edge device, which includes electricity and other maintenance costs. We consider this a pessimistic estimate. The number of SBCs to deploy is a function of the deployment strategy and the compute resources necessary. For the LoRaWAN-aio model, if a single SBC can handle the load for both the LoRaWAN network stack and the IoT service component, a single gateway-capable edge device is adequate. Otherwise, additional SBCs are added (each at a cost of 50 EUR) to scale the network stack and/or the IoT service horizontally at the edge. The LoRaWAN-split-gateway model mandates that at least two SBCs are used; one to host the gateway hardware (200 EUR), and one to host the gateway stack (50 EUR). The necessary number of extra edge devices dedicated to handle the IoT service workload should then be added. Finally, the LoRaWAN-cloud-based model only needs a single SBC to operate as a gateway.

TABLE 3. Approximate processing capacity per compute unit (in events/s).

H/W class Component	SBC (RPi 3 Model B)	Cloud/MEC VM (1 vCPU)
LoRaWAN stack	25	75
Monitor	50	150
Stack+monitor bundle	15	45

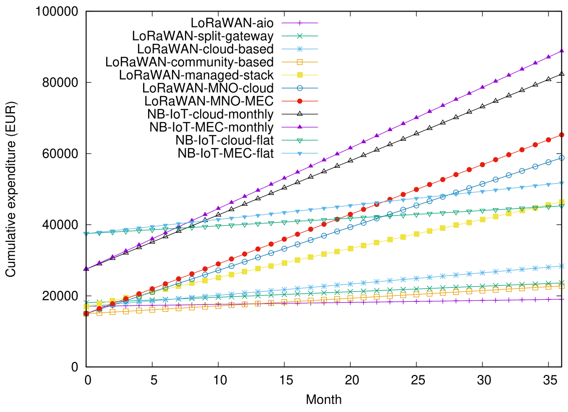

FIGURE 2. Cumulative expenditure including equipment, compute resource, and connectivity costs for different deployment strategies as time progresses.

The latter is a case for using cloud resources to host the LoRaWAN network stack and the IoT service. The price we have assumed per cloud vCPUs is 18 EUR/month, which corresponds to that of an Amazon EC2 A1 instance. We selected this type of instance in particular because it is suitable to execute ARM-based applications, typical of IoT services and also in line with our architecture prototype. This approach further enhances portability and simplifies dynamic deployment and migration of application components between cloud and edge hosts. We have assumed the same price settings for any LoRaWAN and NB-IoT-based model that consumes cloud resources. Unfortunately, as of this writing, actual MEC deployments are at their infancy, while full implementations of the ETSI MEC standard, commercial and open-source, have started emerging. Due to the expected scarcity of MEC resources and MEC's special features, it is reasonable to expect that they will be priced significantly higher than traditional cloud ones. In our analysis, we have set the monthly vCPU fees for MEC servers to double those of cloud ones, i.e., 36 EUR/month.

Compute Resource Dimensioning

As we have hinted already, accurate knowledge of the application workload and the processing capacity of each compute unit (vCPU, SBC) is key for effective resource dimensioning. To demonstrate how this information can assist our model, we have carried out a set of testbed experiments where we deploy our monitor software and LoRaWAN network stack on different classes of compute units and measure the event processing throughput of each software component. In LoRaWAN-based scenarios, each event generated by an IoT device is received by a gateway, processed by the LoRaWAN network server stack, and delivered to a level-0 monitor, where the runtime verification engine is invoked. We used the ChirpStack LoRaWAN network stack implementation (https://www.chirpstack.io), which we deployed on Raspberry Pi 3 Model B devices, as well as inside Linux VMs which we launched at our DC. The latter represents the case for a Cloud or MEC-based network stack deployment. We did the same for our IoT monitoring software and simulated event streams of increasing intensity, recording the maximum supported workload for (i) the LoRaWAN network stack in isolation, (ii) the monitor in isolation, (iii) both the monitor and the network stack bundled in the same compute unit (SBC or the same VM). Our findings are summarized in Table 3; combined with knowledge of the rate at which an end-device emits events and the number of devices per region, the minimum number of necessary edge SBCs or cloud/MEC vCPUs, and thus the respective monetary costs, can be estimated. The same applies for the higher IoT application tiers. The service provider can calculate the workload that higher layer monitors should handle based on the IoT service topology, and the number of necessary compute resources follows.

Deployment Model Comparison

Based on the above cost breakdown and empirical results, we are ready to apply our framework to compare the different deployment strategies. We assume a two-tier monitoring service instance, where there is a single level-0 monitor per region, and all level-0 outputs are aggregated at a single level-1 monitor instance. Each IoT device emits events with a frequency of 1 per hour, while each event message carries 20 bytes of payload. First, we study a case where there is a fixed number of 10 regions, with 100 IoT devices each. Figure 2 presents

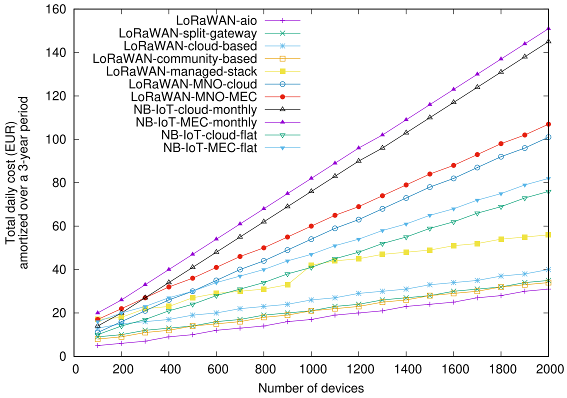

FIGURE 3. Total daily cost per scenario for increasing numbers of devices. The costs are amortized over a 36-month period.

the cumulative expenditure at each month in the first three years of the service operation. This accounts for all the costs described and for the cost settings summarized in Table 2. The cumulative operational expenses of subscription-based models and the end-device procurement costs weigh significantly more than the upfront investments necessitated by models where the network infrastructure is owned and managed by the SP. This is more clearly manifested the larger the deployment becomes. Figure 3 shows the total daily cost per scenario, amortized over a three year period, also factoring in gateway installation and management, as well as compute infrastructure costs. For a fixed number of regions (and thus gateways, if necessary) and for the given application settings, device procurement and LPWAN connectivity (when applicable) dominate other costs for large-scale deployments.

We should finally note that our quantitative results are sensitive to the specific price settings that we have applied. While these are grounded on up-to-date public information, the interested reader may wish to experiment with different settings, and even extend our model to account for more complex IoT service topologies and pricing structures. To this end, our model is available as open-source (https://github.com/pfrag/itechcmp).

DEPLOYMENT GUIDELINES AND CONCLUSION

Based on our results, LoRaWAN may appear more attractive price-wise. However, the service provider may not be willing to take on the overhead of installing and operating a full end-to-end private network, gateways included, or lack the necessary resources and expertise. NB-IoT often simplifies service provision, also taking into account that it brings IP connectivity to the IoT device, thus allowing it to directly address IoT service components in the Cloud. Notably, hidden costs and overheads may be associated with increased reliability requirements. For example, redundant LoRaWAN gateways may be necessary in order to increase the probability that a transmission is successful. In the same sense, redundant edge compute resources (e.g., SBCs) may be deployed as a fail-over solution for edge IoT service components. On the contrary, gateway-less and cloud-based models reduce maintenance effort and the time-to-repair in case of failures. As a result of our study, we draw the following general conclusions and deployment guidelines.

- Device connectivity fees typically associated with NB-IoT subscriptions quickly dominate other capital and operational expenses. This significantly increases costs, especially for very large IoT deployments. In such cases, it is worth exploring strategies based on full end-to-end private LoRaWAN networks.
- NB-IoT and gateway-less LoRaWAN options relieve the stress of managing network and compute infrastructure. As such, they are more suitable for SPs who wish to avoid the related managerial overhead.
- In the absence of MEC offerings by telcos, LoRaWAN-based alternatives where the network is managed by the SP allow better integration with edge computing resources. IoT service components can be directly deployed close to gateways or at on-premise edge servers, thus close to the data sources, without the need to go through an operator's core network or be transported to the Cloud.
- Fully-managed edge-centric LoRaWAN deployments are more suitable in challenged environments where the edge is characterized by low-throughput or intermittent Internet connectivity, there is lack of cellular network coverage, or when dependence to third parties for data transport, processing, and storage infrastructure should be minimized, as potentially dictated by strong privacy requirements.

A key takeaway is that the selection of the appropriate connectivity technology and deployment strategy has multiple facets. However, a major factor that drives this decision is the expenditure involved, and our study can offer SPs valuable insight and the tools to accurately estimate it. Importantly, our framework has the versatility to capture a wide range of potential deployment models, technologies, and service topologies, and accounts for both connectivity and computation costs.

ACKNOWLEDGMENTS

This work was supported by Austrian Science Foundation project M 2778-N "EDENSPACE."

REFERENCES

1. L. Ratliff, "Unlocking captive value: LPWAN enables emerging IoT applications," IHS Markit, Tech. Rep., Jun. 2019. [Online]. Available: https://lora-alliance.org/sites/default/files/2019-07/ihsmarkit_berlin_2019_0.pdf
2. J. Haxhibeqiri, E. D. Poorter, I. Moerman, and J. Hoebeke, "A survey of LoRaWAN for IoT: From technology to application," *Sensors*, vol. 18, no. 11, 2018, Art. no. 3995.
3. Y. E. Wang *et al.*, "A primer on 3GPP narrowband Internet of Things," *IEEE Commun. Mag.*, vol. 55, no. 3, pp. 117–123, Mar. 2017.
4. F. Gu, J. Niu, L. Jiang, X. Liu, and M. Atiquzzaman, "Survey of the low power wide area network technologies," *J. Netw. Comput. Appl.*, vol. 149, 2020, Art. no. 102459.
5. K. Mekki, E. Bajic, F. Chaxel, and F. Meyer, "A comparative study of LPWAN technologies for large-scale IoT deployment," *ICT Express*, vol. 5, no. 1, pp. 1–7, 2019.
6. G. del Campo, I. Gomez, G. Cañada, L. Piovano, and A. Santamaria, "13 - Guidelines and criteria for selecting the optimal low-power wide-area network technology," in *LPWAN Technologies for IoT and M2M Applications*, B. S. Chaudhari and M. Zennaro, Eds. New York, NY, USA: Academic, 2020, pp. 281–305.
7. M. Gusev and S. Dustdar, "Going back to the roots—The evolution of edge computing, an IoT perspective," *IEEE Internet Comput.*, vol. 22, no. 2, pp. 5–15, Mar./Apr. 2018.
8. C. Shu, Z. Zhao, Y. Han, G. Min, and H. Duan, "Multi-user offloading for edge computing networks: a dependency-aware and latency-optimal approach," *IEEE Internet Things J.*, vol. 7, no. 3, pp. 1678–1689, Mar. 2020.
9. R. O. Aburukba, M. AliKarrar, T. Landolsi, and K. El-Fakih, "Scheduling Internet of Things requests to minimize latency in hybrid fog–cloud computing," *Future Gener. Comput. Syst.*, vol. 111, pp. 539–551, 2020.
10. D. Kim, S. Kim, and J. H. Park, "A combined network control approach for the edge cloud and LPWAN-based IoT services," *Concurrency Comput. Pract. Experience*, vol. 32, no. 1, pp. 76–83, 2020.
11. C. Tsigkanos, C. Avasalcai, and S. Dustdar, "Architectural considerations for privacy on the edge," *IEEE Internet Comput.*, vol. 23, no. 4, pp. 76–83, Jul./Aug. 2019.
12. F. Pallas, P. Raschke, and D. Bermbach, "Fog computing as privacy enabler," *IEEE Internet Comput.*, vol. 24, no. 4, pp. 15–21, Jul./Aug. 2020.
13. Mobile Edge Computing (MEC), Framework and Reference Architecture, ETSI Group Specification MEC 003, V2.1.1, Jan. 2019.
14. L. Zanzi *et al.*, "Evolving multi-access edge computing to support enhanced IoT deployments," *IEEE Commun. Standards Mag.*, vol. 3, no. 2, pp. 26–34, Jun. 2019.
15. Orange IoT marketplace, 2020. Accessed: Oct. 23, 2020. [Online]. Available: https://iotmarket.orange.com/connectivity.html
16. R. Sanchez-Iborra, J. Sanchez-Gomez, and A. F. Skarmeta, "Evolving IoT networks by the confluence of MEC and LP-WAN paradigms," *Future Gener. Comput. Syst.*, vol. 88, pp. 199–208, 2018.
17. A. Ksentini and P. A. Frangoudis, "On extending ETSI MEC to support LoRa for efficient IoT application deployment at the edge," *IEEE Commun. Standards Mag.*, vol. 4, no. 2, pp. 57–63, Jun. 2020.

PANTELIS A. FRANGOUDIS is currently a University Assistant with the Distributed Systems Group, TU Wien, Vienna, Austria. His research interests include mobile and wireless networking, network softwarization, edge computing, network security, and Internet multimedia. From 2017 to 2019, he has been a Researcher with the Communication Systems Department, EURECOM, France, and from 2012 to 2017, he was with team DIONYSOS at IRISA/INRIA Rennes, France, which he originally joined under an ERCIM "Alain Bensoussan" postdoctoral fellowship. He received a Ph.D. degree in computer science from Athens University of Economics and Business, Athens, Greece, in 2012. He is the corresponding author of this article. Contact him at pantelis.frangoudis@dsg.tuwien.ac.at.

CHRISTOS TSIGKANOS is currently a Lise Meitner Fellow, TU Wien, Vienna, Austria. His research interests lie in the intersection of distributed systems and software engineering, and include dependable self-adaptive and cyber-physical systems, requirements engineering and formal verification. Formerly, he was a Postdoctoral Researcher with the Distributed

Systems Group, TU Wien, and Politecnico di Milano, Italy. He received the B.Sc. degree in computer science from University of Athens, Athens, Greece, and the M.Sc. degree in software engineering from the University of Amsterdam, Amsterdam, The Netherlands, and the Ph.D. degree in 2017 from the Politecnico di Milano, Milano, Italy. Contact him at christos.tsigkanos@dsg.tuwien.ac.at.

SCHAHRAM DUSTDAR is currently a Professor of computer science with the Distributed Systems Group, TU Wien, Vienna, Austria. From 2004 to 2010, he was an Honorary Professor of information systems with the University of Groningen, Groningen, The Netherlands, from 2016 to 2017, he was a Visiting Professor with the University of Sevilla, Seville, Spain, and in 2017, he was a Visiting Professor with the University of California at Berkeley, Berkeley, CA, USA. He is an elected member of the Academia Europaea, where he is Chairman of the Informatics Section. He was recipient of the ACM Distinguished Scientist Award (2009), the IBM Faculty Award (2012), and the IEEE TCSVC Outstanding Leadership Award (2018). He is the Co-Editor-in-Chief of the ACM Transaction on Internet of Things and the Editor-in-Chief of Computing (Springer). He is also an Associate Editor of the *IEEE Transaction on Services Computing*, the *IEEE Transaction on Cloud Computing*, the ACM Transaction on the Web, and the ACM Transaction on Internet Technology. He serves on the Editorial Board of IEEE Internet Computing and the IEEE Computer Magazine. He is a Fellow of an IEEE. Contact him at schahram.dustdar@dsg.tuwien.ac.at.