Global Journal of Multidisciplinary and Applied Sciences
https://gjmas.com/index.php/gjmas
E- ISSN: 2313-6685
VOL 02 ISSUE 03 2024

Energy-Efficient Communication Protocols for IoT Networks

Hamza Farooq
Department of Electrical Engineering, University of Engineering & Technology (UET), Lahore, Pakistan

Abstract:
The explosive growth of the Internet of Things (IoT) has intensified the need for communication protocols that minimize energy consumption while preserving reliability, latency guarantees, and scalability. This paper surveys energy-efficient protocol families across short-, medium-, and long-range IoT, including IEEE 802.15.4/TSCH with RPL, BLE Mesh, Wi-Fi HaLow (802.11ah), LoRaWAN, and 3GPP NB-IoT. We present a cross-layer perspective on duty-cycling, synchronization, adaptive modulation and coding (AMC), topology control, and traffic-aware MAC scheduling, and we discuss security overheads and their energy implications. A comparative analysis highlights design trade-offs (energy per delivered kilobyte, packet delivery ratio, latency, and deployment complexity). Finally, we synthesize engineering patterns for selecting and tuning protocol stacks under application constraints (event-driven sensing vs. periodic telemetry; indoor vs. outdoor; sparse vs. dense deployments).

Keywords: *IoT, energy efficiency, RPL, TSCH, LoRaWAN, NB-IoT, BLE Mesh, 802.11ah*

INTRODUCTION

Battery-operated IoT nodes often operate for years under stringent energy budgets, which makes communication—typically the dominant energy consumer—central to system design. Selecting an appropriate protocol stack requires balancing energy cost against reliability (PDR), latency, coverage, and total cost of ownership. Short-range and mesh-oriented stacks (e.g., IEEE 802.15.4 with TSCH + RPL, BLE Mesh) emphasize deterministic scheduling and low power at the expense of throughput, while LPWANs (LoRaWAN, NB-IoT) extend range with low data rates, using duty-cycle constraints and ALOHA-like or cellular grant procedures. Emerging sub-GHz WLANs (802.11ah) bridge gaps with higher throughput but tighter power budgets. Cross-layer co-design—combining topology control, MAC scheduling, adaptive data rates, and energy-aware security—remains key to maximizing lifetime without sacrificing service levels.and coverage density Battery chemistry, internal resistance, and lifetime estimation; energy per delivered kb as a unifying metric give me more information in paragraph form

Energy Models and Design Space:

The **Energy Models and Design Space** in IoT communication protocols revolve around quantifying and optimizing how each network component consumes power under varying operational and environmental conditions. IoT nodes typically alternate among **transmit (TX)**, **receive (RX)**, **idle**, and **sleep** states, each with distinct power profiles. The **transition costs** between these states — such as the wake-up latency and the energy required to power the radio front-end or microcontroller — can significantly affect the total energy budget, especially in duty-cycled systems. Synchronization beacons, periodic advertisements, and network keepalive messages ensure connectivity but also incur cumulative microjoule-level energy overheads that can shorten device lifetime if not carefully scheduled.

Workload models also play a vital role in defining energy consumption patterns. **Periodic telemetry** systems, such as environmental sensors that send updates every few minutes, allow predictable scheduling and deep-sleep cycles, optimizing radio wake-ups. In contrast, **event-driven systems** (e.g., intrusion detection, fire alarms) operate with irregular bursts, leading to unpredictable queuing behavior and frequent wake-ups that increase average power draw. Thus, designing an energy model requires accounting for both the **traffic arrival process** and **buffer management policies**, ensuring minimal energy waste during idle waiting times.

The **link-budget and path-loss characteristics** further influence energy efficiency. Sub-GHz bands (e.g., 868 MHz, 915 MHz) offer superior propagation and penetration compared to 2.4 GHz, reducing transmission power for equivalent range but limiting bandwidth. Antenna design and **power amplifier (PA) efficiency** determine how effectively electrical energy converts to radiated power, while coverage density and topology (single-hop vs. multi-hop) dictate the overall network-level energy profile. A denser deployment may reduce individual transmit power but increase control overhead, so finding the balance is crucial.

Finally, **battery chemistry and internal resistance** constrain available energy. Lithium-thionyl chloride and Li-ion cells exhibit different discharge behaviors and internal losses that alter energy delivery efficiency over time. Accurate **lifetime estimation** models therefore integrate not just raw capacity (mAh) but also the **internal resistance growth**, temperature effects, and load duty cycle. The **energy per delivered kilobyte (nJ/kB)** metric has emerged as a universal figure of merit, enabling fair comparison across protocols, hardware, and workloads. This unified metric links physical-layer energy cost to network-level data efficiency, offering a foundation for holistic optimization in energy-efficient IoT system design.

MAC (Medium Access Control):

The **MAC (Medium Access Control)** and **Routing** layers form the operational backbone of low-power IoT communication, where the primary objective is to minimize collisions, idle listening, and retransmissions — the dominant sources of energy waste in wireless networks.

In **TDMA (Time Division Multiple Access)** and **TSCH (Time-Slotted Channel Hopping)** systems, nodes operate in precisely defined **slotframes**—periodic sequences of transmission and reception times. Each link in the network is allocated one or more slots, allowing nodes to remain in deep sleep during all other periods. The **channel hopping** feature, unique to TSCH, further enhances reliability and energy efficiency by mitigating multipath fading and external interference across frequency channels. **Schedule compression** techniques, such as slot reuse and link aggregation, optimize bandwidth utilization and reduce idle slots, ensuring the radio is active only when strictly necessary. Together, these mechanisms achieve deterministic latency and predictable power consumption, ideal for industrial and mission-critical IoT networks.

In contrast, **CSMA (Carrier Sense Multiple Access)**-based systems use contention-based access, which is simpler and more flexible but can suffer from collisions under heavy load. **Low-power listening (LPL)** techniques minimize idle listening by periodically sampling the channel for activity, waking up fully only when a preamble is detected. Additionally, **frame aggregation** reduces protocol overhead by combining multiple small payloads into one transmission, improving channel efficiency. The **backoff mechanism**, a key part of CSMA, is often optimized to balance fairness and energy efficiency: smaller backoffs improve latency but increase collision risk, while longer ones reduce contention at the cost of responsiveness. Intelligent tuning or adaptive backoff algorithms (e.g., exponential backoff or reinforcement learning–based contention management) help achieve an optimal trade-off between throughput, delay, and power consumption.

Routing decisions in low-power networks are equally critical to energy performance. **RPL (Routing Protocol for Low-Power and Lossy Networks)** is the de facto standard for IPv6-based IoT systems, constructing **Directed Acyclic Graphs (DODAGs)** rooted at a central sink. The **Objective Functions** (OF0, MRHOF) guide parent selection based on metrics like **ETX (Expected Transmission Count)**, which quantifies link reliability and indirectly energy cost. **ETX-aware routing** ensures that nodes choose stable, low-loss paths to minimize retransmissions and conserve battery life. RPL's **Trickle Timer algorithm** controls the rate of routing updates: when the network is stable, updates are infrequent, reducing control overhead; when changes occur, timers adaptively shrink to ensure rapid convergence.

For **LPWANs (Low Power Wide Area Networks)** like LoRaWAN, **duty-cycle enforcement** plays a crucial role in limiting channel occupancy and conserving energy. Regulatory limits (e.g., 1% in the EU868 band) constrain how frequently devices may transmit, requiring careful scheduling of uplinks. **Confirmed vs. unconfirmed frames** further influence energy use: confirmed frames guarantee delivery via acknowledgment at the cost of additional airtime, while unconfirmed frames save energy but risk data loss. The **Adaptive Data Rate (ADR)** mechanism dynamically adjusts transmission power and spreading factor based on link quality feedback, balancing reliability and energy consumption. Higher spreading factors extend range but increase airtime; ADR ensures each node operates at its most energy-efficient configuration.

Long-Range Low-Power Wide Area Network (LPWAN):

The **Long-Range Low-Power Wide Area Network (LPWAN)** and **Cellular IoT** paradigms are designed to achieve kilometer-scale communication while maintaining ultra-low energy consumption — a balance crucial for applications such as smart metering, environmental monitoring, and asset tracking. **LoRaWAN** (Long Range Wide Area Network) achieves this balance using **chirp spread spectrum modulation**, which allows communication over long distances with minimal power. It supports three operational **device classes** — **Class A**, **Class B**, and **Class C** — each offering different trade-offs between energy efficiency and latency. **Class A** devices, the most energy-efficient, allow downlink communication only immediately after an uplink transmission, ensuring the radio is mostly off. **Class B** introduces scheduled "ping slots" synchronized by network beacons, enabling periodic downlink opportunities at the expense of additional listening energy. **Class C**, suitable for mains-powered or latency-critical applications, keeps the receiver open continuously, drastically reducing latency but increasing power consumption.Another key factor in LoRaWAN's performance is the use of **Spreading Factors (SF7–SF12)**, which control the symbol rate and range. Lower spreading factors (SF7, SF8) enable faster data rates with less airtime but shorter range, while higher factors (SF11, SF12) extend range and reliability at the cost of longer transmissions and increased energy use. The **energy-latency trade-off** is therefore application-specific: short bursts of telemetry at SF7 can last years on a single battery, whereas continuous monitoring at SF12 may deplete it within months. Moreover, **gateway diversity**, where multiple gateways simultaneously receive the same uplink, enhances reliability and range without extra energy cost at the end node. However, **downlink scarcity** remains a major limitation — since gateways share duty-cycle restrictions, acknowledgments and configuration messages must be carefully rationed, often prioritized for critical commands or confirmed uplinks.

NB-IoT (Narrowband Internet of Things:

In contrast, **NB-IoT (Narrowband Internet of Things)**—a 3GPP-standardized technology—integrates IoT connectivity into existing LTE infrastructure, providing better reliability and quality-of-service control. NB-IoT employs two primary energy-saving mechanisms: **Power Saving Mode (PSM)** and **extended Discontinuous Reception (eDRX)**. PSM allows devices to remain virtually unreachable for long periods (days or weeks) while preserving registration in the network, enabling extreme sleep durations with negligible energy draw. eDRX, on the other hand, allows devices to wake up periodically to check for downlink data, balancing responsiveness with power conservation. To enhance signal reliability in weak coverage scenarios (e.g., basements or rural areas), NB-IoT employs **coverage enhancement repetitions**, retransmitting control and data channels multiple times to ensure delivery — though this increases both latency and energy per bit. The **Radio Resource Control (RRC)** connection setup and teardown procedures, while ensuring secure and managed access, introduce **signaling overhead**, and **attach timers** govern how frequently a device reconnects to the network after sleep cycles, which can heavily influence battery life if misconfigured.

Both LoRaWAN and NB-IoT predominantly serve **uplink-heavy telemetry workloads**, where data flows mostly from sensors to cloud servers. This asymmetry necessitates **fragmentation policies** and **payload shaping** to maximize efficiency: small packets reduce retransmission risk but increase header overhead, whereas large packets optimize throughput but risk corruption in lossy channels. Smart payload design—compressing or aggregating data before transmission—can drastically reduce total energy per report. Additionally, modern IoT deployments increasingly rely on **Firmware-Over-The-Air (FOTA)** updates for security and functionality enhancements. FOTA poses unique energy challenges since it involves large data transfers; energy budgeting must therefore incorporate **progressive download strategies**, **delta updates**, and **adaptive scheduling** (e.g., during periods of strong signal or available mains power).

Security, Reliability, and Quality of Service:

The **Security, Reliability, and Quality of Service (QoS) overheads** in low-power IoT networks represent a critical balance between maintaining system trustworthiness and preserving energy efficiency. Implementing cryptographic mechanisms ensures data confidentiality, integrity, and authenticity but introduces both **CPU computation** and **radio transmission costs**. **Link-layer encryption** (e.g., **AES-CCM** or **AES-GCM**) protects frames hop-by-hop, providing efficient protection with minimal overhead due to hardware acceleration in many modern transceivers. However, it lacks end-to-end confidentiality, meaning intermediate nodes can still inspect payloads. In contrast, **end-to-end encryption frameworks** like **DTLS (Datagram Transport Layer Security)** or **OSCORE (Object Security for Constrained RESTful Environments)** safeguard data across the full communication path. These provide stronger guarantees but require additional headers and session-handshake messages, consuming extra bytes and processing power. For ultra-low-power devices, even a few extra milliseconds of radio-on time or cryptographic computation can meaningfully shorten battery life, making lightweight cryptographic optimization a necessity rather than an option.

At scale, **key management** becomes an equally challenging dimension. IoT deployments with thousands or millions of nodes demand automated **join and attestation procedures** that verify device authenticity and distribute encryption keys securely without manual intervention. Protocols like LoRaWAN's **Over-the-Air Activation (OTAA)** or Thread's **Commissioning** involve multi-step exchanges, which, while secure, consume airtime and energy. Once deployed, periodic **rekeying** or credential rotation helps prevent long-term key compromise but incurs transmission overhead and synchronization latency. The choice of **rekey intervals** must therefore balance security risk against network bandwidth and node energy constraints. Each join or rekey exchange may require multiple encrypted messages and acknowledgments, extending radio uptime — a nontrivial cost in battery-powered networks.

Reliability mechanisms are fundamental to IoT performance, especially in noisy or interference-prone environments. Techniques such as **selective acknowledgments (SACK)** ensure that only missing packets are retransmitted, reducing redundant traffic. **Forward Error Correction (FEC)** adds parity bits for error recovery without retransmission but increases packet size and, consequently, airtime. Meanwhile, **retransmission policies** must adapt to traffic types and channel conditions — excessive retries waste energy, while too few compromise data delivery. To handle **bursty traffic**, such as when multiple sensors trigger simultaneously, **congestion control** algorithms (e.g., adaptive backoff or load-aware queuing) throttle uplinks to prevent channel saturation and cascading packet loss. These strategies collectively ensure high **Packet Delivery Ratio (PDR)** with controlled energy expenditure.

Latency and QoS differentiation introduce another layer of design trade-offs. Not all IoT traffic has equal urgency — **alarm messages** in industrial safety systems demand sub-second responsiveness, while **metering data** or environmental logs can tolerate delays of several seconds or minutes. Protocols such as **IEEE 802.15.4e TSCH** achieve deterministic latency through scheduled slotframes, ensuring time-bounded delivery for critical traffic, while less urgent packets occupy best-effort slots. Similarly, **NB-IoT** employs **Discontinuous Reception (DRX)** cycles, allowing devices to wake at specific intervals to receive downlink data. By aligning latency classes with DRX scheduling, networks can prioritize alarms without keeping all nodes active, conserving energy while maintaining service-level objectives (SLOs).

Engineering Patterns and Tuning Playbook:

The **Engineering Patterns and Tuning Playbook** in energy-efficient IoT networking serves as a practical guide for mapping application scenarios to optimal protocol configurations, ensuring the best balance between energy consumption, reliability, and responsiveness. Each deployment context — whether sparse rural sensing or dense industrial automation — demands tailored protocol behavior. In **sparse outdoor sensing networks**, where nodes are geographically dispersed and traffic is sporadic (e.g., agricultural or environmental monitoring), **LoRaWAN Class A** devices are ideal. These nodes send **infrequent uplinks** and stay in deep sleep for most of their lifecycle. The **Adaptive Data Rate (ADR)** feature automatically adjusts transmission power and spreading factor based on link quality, minimizing airtime and conserving battery life. Only **critical alerts**—such as threshold breaches or emergency conditions—use **confirmed frames** (requiring acknowledgment), while routine data relies on unconfirmed uplinks to avoid unnecessary energy expenditure and downlink congestion.

For **dense industrial mesh environments**, such as smart factories or power substations, the emphasis shifts to predictability and determinism. Here, **IEEE 802.15.4e TSCH (Time-Slotted Channel Hopping)** combined with **RPL routing** provides precise time synchronization and channel diversity. By allocating **deterministic slots** to nodes and minimizing idle listening through synchronized scheduling, collisions are virtually eliminated. **Minimal advertising** and **compressed IPv6 headers (6LoWPAN)** further reduce control traffic, ensuring scalability even in congested radio environments. This pattern is especially effective for time-sensitive control loops and sensor-actuator coordination, where predictable latency and high reliability are paramount. The trade-off, however, lies in the complexity of schedule management and the requirement for tight clock synchronization.

In **building automation systems**, energy efficiency must coexist with moderate latency and bidirectional communication. **Bluetooth Low Energy (BLE) Mesh** fits this environment by leveraging **friend nodes** that temporarily store messages for **low-power nodes (LPNs)**, allowing end devices to remain in sleep mode for extended durations. Managing **relay density**—ensuring not every node acts as a forwarder—prevents broadcast storms and optimizes throughput. **Periodic health checks** and heartbeat messages maintain network integrity without constant communication. This architecture suits lighting systems, HVAC control, and occupancy sensing, where energy savings and responsiveness must align with user experience and reliability.

For **medium-range high-throughput telemetry**, such as logistics tracking or industrial monitoring with larger payloads, **IEEE 802.11ah (Wi-Fi HaLow)** provides an efficient compromise between range, speed, and power. The **Target Wake Time (TWT)** mechanism allows devices to negotiate sleep and wake schedules with the access point, drastically reducing idle listening. **Beacon interval optimization** fine-tunes synchronization to match traffic frequency — longer intervals for periodic telemetry and shorter ones for time-critical data. Proper tuning of these parameters reduces contention while preserving real-time responsiveness, making 802.11ah a strong choice for IoT gateways, metering, and remote control systems requiring moderate data throughput.

Finally, **cross-layer governance** binds these engineering strategies together. Effective IoT deployments treat energy as a managed resource through well-defined **telemetry budgets per device** (e.g., bytes/day, airtime/month), enforced via firmware policies. Networks should be designed around **Service-Level Objectives (SLOs)** — typically expressed in **Packet Delivery Ratio (PDR)** and **latency bounds** — to ensure that tuning decisions align with mission requirements. An **energy-aware CI/CD pipeline** can continuously monitor field telemetry (e.g., RSSI, duty-cycle usage, battery slope) and dynamically adjust ADR settings, slot schedules, or DRX cycles based on real-world performance. This closed-loop governance ensures long-term stability and scalability, preventing performance degradation as device populations grow or environmental conditions evolve.

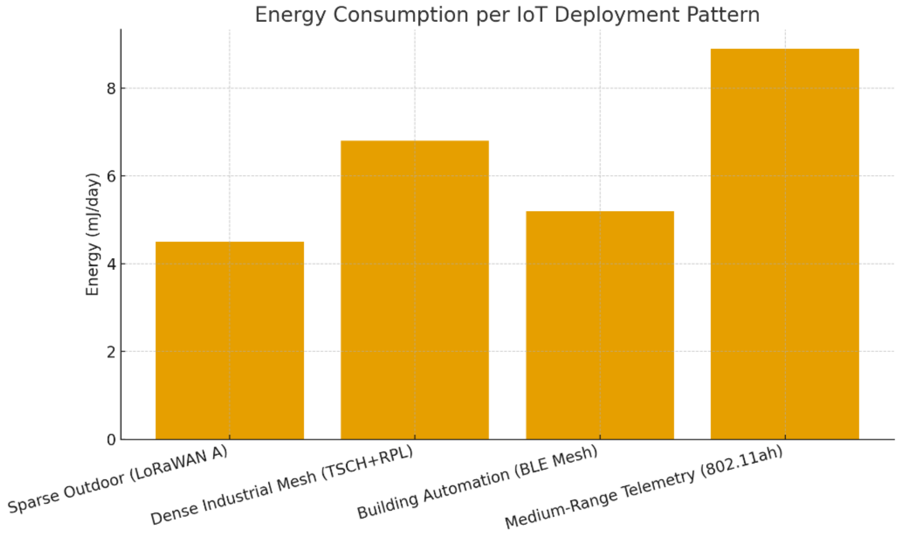

Summary:

Energy efficiency in IoT networking is a multi-dimensional optimization spanning PHY/MAC scheduling, routing stability, security overheads, and workload shaping. Deterministic time-slotted meshes (TSCH+RPL) excel in dense, interference-prone environments with predictable traffic, while LoRaWAN minimizes energy for sparse uplinks over long range at the expense of downlink capacity and latency. NB-IoT delivers carrier-grade availability with deep-sleep features but higher attach/signaling energy, suited to regulated utilities and SLA-driven deployments. BLE Mesh is effective for building-scale control with careful relay placement and friend/LPN roles. Wi-Fi HaLow provides a middle ground when higher throughput is needed. Practitioners should adopt telemetry budgets (bytes/day), enforce SLO-linked re-transmission limits, and continuously tune ADR/TWT/slotframes using field telemetry to sustain battery life targets.

References:

RFC 6550: RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. IETF, 2012.
RFC 8180: Minimal IPv6 over TSCH Networks (6TiSCH). IETF, 2017.
IEEE Std 802.15.4e-2012: MAC enhancements for low-latency/low-power. IEEE, 2012.
RFC 7252: The Constrained Application Protocol (CoAP). IETF, 2014.
RFC 8613: Object Security for Constrained RESTful Environments (OSCORE). IETF, 2019.
LoRa Alliance. *LoRaWAN® 1.0.4 Regional Parameters & Specifications*, 2020–2023.

Bluetooth SIG. *Bluetooth Mesh Profile Specification*, v1.x, 2017–2023.
3GPP TS 36.300 & 24.301: *NB-IoT/E-UTRAN overall description & NAS signaling*, Rel-13+.
IEEE Std 802.11ah-2016: Sub-1 GHz WLAN (Wi-Fi HaLow). IEEE, 2016.
RFC 6282: Compression Format for IPv6 Datagrams over IEEE 802.15.4 (6LoWPAN). IETF, 2011.
Winter, T. et al. *RPL Objective Function Zero (OF0)*, RFC 6552, 2012.
Palattella, M. R. et al., “Standardized Protocol Stack for the Internet of (Important) Things,” *IEEE Comm. Surveys & Tutorials*, 2013.