



Received: 07-11-2025
Accepted: 17-12-2025

ISSN: 2583-049X

A Conceptual Framework for Characterizing Fundamental Energy, Reliability, and Latency Trade-offs in Green Communication Paradigms for Next-Generation IoT

¹ William Asiedu, ² Robert Quainoo, ³ Alice Asiedu

¹ Iowa State University, USA

² Northeastern University, USA

³ Ashesi University, Ghana

Corresponding Author: **William Asiedu**

Abstract

The convergence of massive device populations, sustainability mandates, and mission-critical service requirements has placed energy consumption, communication reliability, and end-to-end latency at the centre of next-generation Internet of Things (IoT) design. These three objectives are mutually coupled, so that improvement along one axis routinely imposes a penalty along the others, yet the field continues to study them in fragmented, paradigm-specific silos. This paper develops a conceptual framework that characterizes the fundamental trade-offs among energy, reliability, and latency, referred to here as the ERL trilemma, within green communication paradigms intended for next-generation IoT. The three axes are formalized through a shared set of operational metrics and connected to information-theoretic anchors, drawing on classical capacity theory, the wideband energy-per-bit limit, and finite-blocklength coding bounds to explain why the coupling is fundamental rather than incidental. Achievable

operating points are represented as a Pareto frontier, so that any green technique can be located by the trade-off signature it produces, and the dominant green communication paradigms, including energy harvesting, duty-cycled and sleep-mode operation, low-power wide-area networking, mobile edge computing, large-scale antenna and reconfigurable-surface techniques, and learning-based orchestration, are positioned according to where each paradigm displaces the surface. The resulting mapping yields a set of design principles and an agenda of open challenges, including joint-objective optimization under finite blocklength, lifecycle-aware energy accounting, and the reliability cost of learning-based control. The framework is intended as an organizing lens rather than a single optimization recipe, offering researchers and system designers a consistent vocabulary for reasoning about trade-offs that have so far been treated in isolation.

Keywords: Green Communication, Internet of Things, Energy Efficiency, Ultra-reliable Low-latency Communication, Energy-reliability-latency Trade-off, Finite Blocklength, Pareto Frontier, Sustainable Networks

Introduction

The Internet of Things is projected to interconnect tens of billions of devices that sense, actuate, and exchange data across industrial, civic, agricultural, and consumer settings. Two forces are reshaping how such systems must be engineered. The first is sustainability. As connected device populations expand and as the carbon and operational cost of network infrastructure grows, energy has shifted from a deployment convenience to a primary design objective, and green communication, the systematic minimization of energy per useful bit and per useful task, has become a central concern rather than an afterthought (Chen *et al.*, 2011; Feng *et al.*, 2013; Mahapatra *et al.*, 2016) ^[11, 16, 20], with recent work extending predictive energy management to machine-learning-enabled control over IoT sensor networks (Kumuyi *et al.*, 2024) ^[19]. The second force is criticality. A widening class of IoT applications, including industrial control, vehicular coordination, remote health monitoring, and smart-grid protection, demands stringent guarantees on both reliability and latency, a regime now studied under the heading of ultra-reliable low-latency communication (URLLC) (Durisi *et al.*, 2016; Popovski *et al.*, 2018) ^[14, 27]. These two forces do not act independently. Reducing energy consumption frequently lengthens delay or weakens reliability, while tightening latency and reliability targets generally raises energy expenditure. The classical green-networking literature recognized several pairwise trade-offs early, most notably the relationships among energy efficiency, spectral efficiency,

deployment density, bandwidth, and delay (Chen *et al.*, 2011) [11]. The URLLC literature, in turn, established that reliability and latency are bound together through the statistics of short-packet transmission, where the law of large numbers no longer smooths out errors (Durisi *et al.*, 2016; Polyanskiy *et al.*, 2010) [14, 26]. What the field has lacked is an integrated treatment in which energy, reliability, and latency are characterized jointly, as three faces of a single constrained design problem, rather than as a sequence of two-dimensional trade-offs examined in separate papers.

This paper argues that next-generation IoT design is best understood through a three-way tension, the ERL trilemma, and it builds a conceptual framework for reasoning about that tension across the diverse paradigms that travel under the green communication banner. The framework is deliberately conceptual. It does not propose a new protocol or a single optimization algorithm. Instead, it offers a structured vocabulary and a set of analytical anchors that allow heterogeneous techniques, from energy harvesting to reconfigurable surfaces to learning-based orchestration, to be compared on common ground.

Three elements give the framework its structure. The three axes are first expressed through a shared set of operational metrics and connected to three information-theoretic anchors, namely Shannon capacity, the wideband energy-per-bit limit, and the finite-blocklength coding rate, which together indicate that the coupling among the axes is fundamental rather than an artifact of any particular system. The achievable combinations of the three objectives are then represented as a Pareto frontier within a triangular objective space, so that a design choice is understood as the selection of a position on a surface rather than a single number, and so that any technique can be seen either to move along that surface or to displace it outward. The dominant families of green communication techniques are finally classified by the direction in which each displaces the surface, which yields a compact mapping from which a set of design principles and open challenges follows.

Background and Related Work

Green communication paradigms

Green communication refers to the body of techniques whose explicit aim is to reduce the energy required to deliver a given communication outcome, whether that outcome is measured per transmitted bit, per completed task, or per unit of network lifetime (Chen *et al.*, 2011; Feng *et al.*, 2013; Mahapatra *et al.*, 2016) [11, 16, 20]. Early work concentrated on the radio access segment, where the power amplifier and the always-on components of base stations and devices dominate consumption. Subsequent work broadened the scope to the full protocol stack and to the network as a whole, encompassing sleep-mode and duty-cycling strategies, energy harvesting and wireless power transfer, low-power wide-area networking for sparse low-rate traffic, computation offloading toward the network edge, and, more recently, learning-based control that adapts operating points to traffic and channel conditions (Buzzi *et al.*, 2016; Mahapatra *et al.*, 2016; Mao *et al.*, 2017; Raza *et al.*, 2017) [10, 20, 21, 28]. Beyond the radio segment, the green agenda extends to the computing and physical infrastructure that supports connected systems, including energy-efficient virtual-machine placement in data centres (Ahmed & Odejobi, 2018) [2], climate-adaptive building envelopes that lower operational energy in tropical deployments (Uduokhai

et al., 2025a) [33], and the optimization of electrical load distribution for hybrid solar and solar-diesel installations that power devices in developing and underserved regions (Ijiga *et al.*, 2023; Sunday & Omoegun, 2019) [18, 31].

A recurring theme across these paradigms is that energy savings are rarely free. Sleeping a radio saves power but delays the next available transmission opportunity. Harvesting ambient energy removes a battery constraint but introduces an intermittent and uncertain supply that complicates timely and reliable delivery. Offloading computation reduces device energy but adds communication and queuing delay. The framework developed here treats this absence of free savings as the organizing principle rather than as a series of isolated observations.

The fundamental trade-off literature

The notion that wireless objectives trade against one another is well established. Chen and colleagues articulated four fundamental trade-offs for green wireless networks, relating energy efficiency to spectral efficiency, to deployment density, to bandwidth, and to delay (Chen *et al.*, 2011) [11]. The energy-efficiency and spectral-efficiency relationship in particular has a clear information-theoretic basis, since pushing toward the Shannon limit at high spectral efficiency requires disproportionately more energy per bit, while operating in the wideband regime approaches the minimum energy per bit at the cost of spectral efficiency (Shannon, 1948; Verdu, 2002) [30, 35]. Survey treatments have since catalogued energy-efficient techniques at every protocol layer and have formalized energy-efficiency optimization through fractional programming, where the ratio of throughput to power is maximized subject to quality-of-service constraints (Buzzi *et al.*, 2016; Mahapatra *et al.*, 2016; Zappone & Jorswieck, 2015) [10, 20, 36]. The broader optimization literature offers complementary machinery for navigating such trade-offs, including constraint-satisfaction models for network resource allocation (Ahmed *et al.*, 2019) [3] and approximation-complexity treatments of cloud optimization problems (Odejobi *et al.*, 2019) [23].

In parallel, the URLLC literature established the coupling between reliability and latency. When packets are short, as they typically are for control and sensing traffic, the achievable coding rate falls below the Shannon capacity by a penalty that grows as both the target error probability and the blocklength shrink (Durisi *et al.*, 2016; Polyanskiy *et al.*, 2010) [14, 26]. Because latency budgets cap the blocklength and reliability budgets cap the error probability, the two objectives are inseparable at the physical layer, and retransmission-based schemes that improve reliability do so by spending both time and energy (Avranas *et al.*, 2018; Popovski *et al.*, 2018) [7, 27].

Energy, reliability, and latency together

A smaller body of work has begun to treat all three objectives jointly. Studies of machine-to-machine and IoT scheduling have characterized how energy efficiency, latency, and packet loss move together under realistic fading and resource-allocation policies (Torkudzor *et al.*, 2021) [32], and multi-parameter schemes have been proposed that target green and reliable operation simultaneously in cellular IoT (Din *et al.*, 2018) [13]. Reviews of URLLC for IoT have classified methods by the quality-of-service dimensions they prioritize and have noted the energy implications of reliability-enhancing mechanisms (Parvez *et al.*, 2018; Sefati *et al.*, 2023) [25, 29]. Complementary insights come from the study of reliability and performance at scale in

large technology programs and public services, where applied performance optimization frameworks (Oteri & Edivri, 2024)^[24], predictive capacity planning and resource utilization forecasting (Edivri & Oteri, 2022)^[15], performance intelligence models for outcome measurement (Adelanwa *et al.*, 2023)^[1], and predictive maintenance scheduling for aging infrastructure (Uduokhai *et al.*, 2025b)^[34] address how dependable operation is sustained as systems grow. These contributions confirm that the three-way tension is real and consequential. What remains missing is a generalizing structure that abstracts away from any single scheduling policy or air-interface and lets disparate green paradigms be positioned relative to one another.

Read together, these strands leave three gaps. The literature remains largely pairwise, treating energy against spectral efficiency, energy against delay, and reliability against latency in separation, so that the joint surface these objectives define is rarely drawn explicitly. It is also largely technique-centric, reporting results for a particular harvesting model, sleep schedule, or offloading policy in a way that frustrates comparison across paradigms. And it offers no shared vocabulary for the displacement a technique produces, that is, for whether a technique trades one objective for another along a fixed surface or instead enlarges the achievable region. The work that follows responds to all three by unifying the axes, introducing a surface abstraction, and classifying paradigms by their displacement signature.

The Energy, Reliability, and Latency Trilemma

This section defines the three axes precisely enough to support the framework while remaining independent of any particular system. Each axis is given an operational meaning and a representative metric, after which the coupling mechanisms that bind them are described.

The three axes

Energy. The energy axis captures the resource cost of communication and the computation that supports it. At the link level it is naturally measured as energy per delivered information bit, which separates the efficiency of a scheme from the volume of traffic it carries. At the device level it is measured as average power and, through battery or harvested supply, as expected operational lifetime. At the network level it extends to aggregate consumption and to the embodied and operational carbon associated with infrastructure. A complete accounting therefore spans instantaneous power, energy per useful task, and lifetime, and a technique can improve one of these while worsening another.

Reliability. The reliability axis captures the probability that information is delivered correctly and on time. At the physical layer it is expressed as a decoding error probability for a given transmission. At the service layer it becomes availability, the fraction of time that the system meets its delivery guarantee, and is frequently specified at very small target error probabilities for mission-critical traffic. Reliability is meaningful only relative to a deadline, which is what links it to the latency axis.

Latency. The latency axis captures the time from the generation of information to its correct delivery and use. It comprises transmission time, propagation, queuing, retransmission, and any processing or offloading delay. For control and coordination traffic the relevant quantity is often a worst-case or high-percentile latency rather than a mean,

because a deadline missed is, for such traffic, indistinguishable from a delivery failure. Jitter, the variability of latency, is itself a design quantity for closed-loop applications. Queueing-theoretic and concurrency-aware models developed in adjacent domains, including throughput and capacity analysis of automated logistics chains (Akinola *et al.*, 2024)^[6], performance evaluation of multi-tenant cloud services under high concurrency (Odejebi & Ahmed, 2018)^[22], and machine-learning-based predictive scaling of cloud resources (Ahmed *et al.*, 2020)^[4], illustrate how latency and load are characterized and managed in practice.

Coupling mechanisms

The three axes are coupled through several concrete mechanisms that recur across systems. Transmit power couples energy to reliability, because raising power lowers the error probability at fixed rate and blocklength, while lowering power conserves energy at the cost of reliability. Retransmission couples all three axes, since each retransmission improves reliability but consumes additional time and energy (Avranas *et al.*, 2018; Popovski *et al.*, 2018)^[7, 27]. Blocklength couples latency to reliability, because a shorter packet reduces transmission delay yet enlarges the gap between the achievable rate and capacity, which must be repaid in either power or error probability (Durisi *et al.*, 2016; Polyanskiy *et al.*, 2010)^[14, 26]. Duty cycling and sleep scheduling couple energy to latency, since deeper or longer sleep saves power but postpones the next opportunity to transmit or receive. Computation offloading couples device energy to latency and to reliability, because moving work to the edge saves local energy but introduces communication delay and an additional point of failure (Mao *et al.*, 2017)^[21]. The framework treats these mechanisms as the edges of the trilemma, the channels through which a gain on one axis is paid for on another.

Characterizing the trade-offs

The framework has two components. The first is a set of analytical anchors that ground the trilemma in established theory and explain why the coupling is fundamental. The second is the trilemma surface, a geometric abstraction that turns the qualitative tension into a structure on which techniques can be located.

Analytical anchors

Anchor one: *Capacity and the energy-spectral-efficiency relationship.* Shannon's channel capacity, expressible for an additive white Gaussian noise channel as the bandwidth multiplied by the logarithm of one plus the signal-to-noise ratio, establishes the maximum reliable rate for a given power and bandwidth (Shannon, 1948)^[30]. Rearranging capacity in terms of energy per bit shows that operating at high spectral efficiency forces energy per bit to grow without bound, whereas spreading the same information over more bandwidth reduces energy per bit toward a fundamental minimum. This is the origin of the energy-efficiency and spectral-efficiency trade-off that the green-networking literature treats as foundational (Chen *et al.*, 2011; Verdu, 2002)^[11, 35].

Anchor two: *The wideband energy-per-bit limit.* In the wideband regime the minimum energy per bit required for reliable communication approaches a fixed limit determined by the noise spectral density, corresponding to the well-known minimum signal-to-noise ratio per bit of

approximately minus 1.59 decibels (Verdu, 2002) [35]. This anchor sets the floor of the energy axis. No green technique can deliver a bit below this energy without changing the channel itself, for example by improving the effective signal-to-noise ratio through beamforming or by harvesting energy from the environment so that the cost is borne elsewhere. The limit therefore distinguishes techniques that move along the surface from techniques that genuinely relocate the channel.

Anchor three: The finite-blocklength coding rate. For short packets the maximum coding rate is well approximated by the channel capacity reduced by a term proportional to the square root of the channel dispersion divided by the blocklength, scaled by the inverse Gaussian function of the target error probability (Polyanskiy *et al.*, 2010) [26]. This single expression binds the three axes at once. The blocklength is set by the latency budget, the target error probability is set by the reliability requirement, and the resulting rate, together with the bandwidth and power, determines the energy expended. The finite-blocklength anchor is the formal heart of the trilemma, since it shows that for the short transmissions characteristic of IoT control traffic, energy, reliability, and latency cannot be optimized separately even in principle (Durisi *et al.*, 2016; Popovski *et al.*, 2018) [14, 27].

The trilemma surface

The analytical anchors imply that, for a fixed channel and a fixed set of physical resources, the achievable combinations of energy, reliability, and latency form a bounded region. The efficient boundary of that region is a Pareto frontier: On the frontier, no axis can be improved without sacrificing another, while points strictly inside the frontier are dominated and points outside are infeasible. Figure 1 represents this region within a triangular objective space whose vertices denote the three axes.

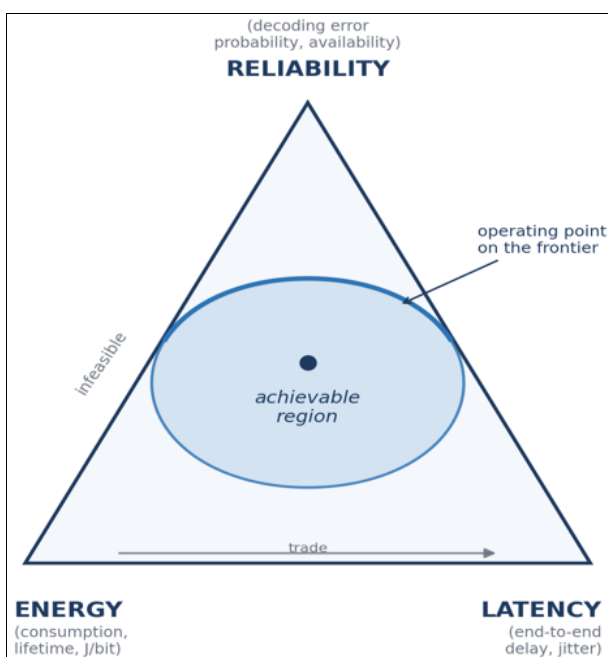


Fig 1: The energy, reliability, and latency (ERL) trilemma. The triangular objective space bounds the three competing objectives; the shaded achievable region is delimited by a Pareto frontier on which no objective can be improved without conceding another. A green technique either moves the operating point along the existing frontier or displaces the frontier outward

The surface abstraction yields two distinct kinds of design action, which the framework names explicitly. A traversal moves the operating point along the existing frontier, trading one objective for another without changing what is achievable; raising transmit power to lower error probability at fixed latency is a traversal. An expansion displaces the frontier outward, enlarging the achievable region so that better combinations of all three objectives become simultaneously reachable; adding spatial degrees of freedom through large antenna arrays, or improving the effective channel through a reconfigurable surface, is an expansion. This distinction is the principal analytical contribution of the framework, because it separates techniques that merely reallocate a fixed budget from techniques that change the budget itself.

Locating a technique

Any green communication technique can now be characterized by asking which axis it primarily targets, which axis or axes pay for the gain and thus which coupling edge is exercised, and whether the action is a traversal along the frontier or an expansion of it. The answers form a compact signature that places the technique within the framework and makes it directly comparable with others. The following section applies this lens to the dominant paradigms.

Mapping Green Paradigms onto the Trilemma

This section examines the principal families of green communication techniques for next-generation IoT and locates each within the framework. Table 1 summarizes the mapping; the discussion that follows explains the entries.

Energy harvesting and wireless power transfer

Harvesting energy from solar, radio-frequency, kinetic, or thermal sources, and transferring power wirelessly to devices, relaxes the energy axis by replacing or supplementing a finite battery. The cost falls primarily on latency and reliability, because harvested energy arrives intermittently and unpredictably, so a device may have to wait to accumulate sufficient charge before it can transmit reliably. In framework terms this is chiefly a traversal that buys energy at the expense of timeliness, although by removing the battery as a hard lifetime constraint it can also expand the feasible region for long-lived deployments.

Duty cycling, sleep modes, and discontinuous reception

Switching radios and subsystems off when idle is the most widely deployed energy-saving mechanism in IoT. It targets the energy axis directly by eliminating the consumption of always-on components, and it pays on the latency axis, since a sleeping receiver cannot react until it next wakes, and on reliability where missed wake windows cause loss. This is a clear traversal along the energy-latency edge, and the depth and schedule of sleep are the control variables that select the operating point (Bockelmann *et al.*, 2016; Mahapatra *et al.*, 2016) [9, 20].

Low-power wide-area networking

Low-power wide-area networks serve sparse, low-rate, delay-tolerant traffic over long ranges at very low power (Raza *et al.*, 2017) [28]. They achieve their energy and coverage characteristics by accepting low data rates and relaxed latency, and by using robust, low-rate modulation that trades spectral efficiency for energy efficiency and range. Within the framework these technologies occupy a corner of the surface that favours energy and coverage while

conceding latency and throughput, and they are well matched to applications whose reliability requirement is modest and whose deadlines are loose.

Mobile edge computing and computation offloading

Moving computation from constrained devices to nearby edge servers reduces device energy and can shorten the latency of compute-heavy tasks relative to distant cloud processing (Mao *et al.*, 2017) [21]. The trade is subtle. Offloading saves local energy but adds uplink and downlink communication, queueing at the edge, and a dependency on the availability of the edge resource, which affects reliability. Whether offloading is a traversal or an expansion depends on the task: For computation that is expensive locally but cheap to transmit, offloading can expand the achievable region across all three axes, whereas for light tasks the added communication delay makes it a traversal that is often not worthwhile. Edge architectures also host increasingly sophisticated workloads, such as deep-learning-driven classification for cloud-native microservices (Idika *et al.*, 2021) [17], which sharpens the trade between the energy saved by offloading and the additional latency, communication, and failure surface that offloading introduces.

Large-scale antennas and reconfigurable surfaces

Massive antenna arrays and reconfigurable intelligent surfaces improve the effective channel rather than merely reallocating a fixed budget. Large arrays concentrate energy spatially, raising the effective signal-to-noise ratio so that the same reliability and latency can be met at lower transmit energy, and they improve hardware and energy efficiency at scale (Bjornson *et al.*, 2017) [8]. Reconfigurable surfaces shape the propagation environment to strengthen the desired signal and suppress interference (Di Renzo *et al.*, 2020) [12]. Both are expansions in the framework sense: By changing the channel they push the entire trilemma surface outward, allowing simultaneous gains on more than one axis, at the cost of additional hardware, control overhead, and channel-state acquisition.

Learning-based orchestration

Data-driven control, including reinforcement learning for power, scheduling, and resource allocation, adapts the operating point to traffic and channel conditions far more responsively than fixed policies. Multi-agent formulations extend this idea to coordinated control across many devices, as in swarm-intelligence frameworks for autonomous fleet coordination (Akanbi *et al.*, 2025) [5], where distributed agents jointly manage objectives that no single device controls. Its principal value in the framework is twofold. It can keep the system near the Pareto frontier under changing conditions, reducing the dominated gap that static policies leave, and through better joint optimization it can sometimes approximate an expansion. The cost is a reliability concern of a new kind: A learned policy may behave unpredictably outside its training distribution, and the energy and latency of the learning and inference processes themselves must be counted. Learning-based orchestration therefore appears in the framework as both a frontier-tracking traversal and a source of new reliability questions.

Table 1: Mapping of green communication paradigms onto the ERL trilemma

Paradigm	Targets	Pays on	Action	Coupling edge
Energy harvesting and wireless power transfer	Energy, lifetime	Latency, reliability	Traversal, partial expansion	Energy to latency
Duty cycling, sleep, discontinuous reception	Energy	Latency, reliability	Traversal	Energy to latency
Low-power wide-area networking	Energy, range	Latency, throughput	Traversal	Energy to latency
Mobile edge computing and offloading	Device energy	Latency, reliability	Task dependent	Energy to latency
Massive antennas and reconfigurable surfaces	Energy, reliability	Hardware, overhead	Expansion	Channel improvement
Learning-based orchestration	All three	Robustness, compute	Traversal, approximate expansion	Frontier tracking

Design Principles and Open Challenges

Design principles

Several principles follow directly from the framework, stated as guidance rather than as proofs. Because the objectives lie on a surface, a requirement that fixes only one axis is underspecified, so a well-posed IoT design states a target for all three axes together and recognizes that the third is implied once two are fixed on the frontier. Techniques that improve the channel, such as spatial multiplexing or reconfigurable surfaces, relax the trilemma for every objective at once, and they are therefore worth considering before techniques that merely reallocate a fixed budget, subject to their hardware and overhead cost. Energy should be accounted for over the full task and lifecycle, since link-level energy per bit can mislead when retransmissions, protocol overhead, idle consumption, and the energy of supporting computation are excluded; the relevant quantity is energy per completed and timely task, measured over the device lifetime. Reliability should be treated as deadline-relative, because for mission-critical IoT a correct delivery after its deadline is a failure, so reliability targets belong with a latency budget and finite-blocklength effects should be modelled rather than assumed away. Finally, the system should be kept near the frontier under change, since static operating points drift into the dominated interior as traffic and channels vary, and adaptive control that tracks the frontier recovers the otherwise wasted margin, provided its own energy, latency, and robustness costs are counted.

Open challenges

The framework also clarifies where research is most needed. Joint-objective optimization under finite blocklength remains analytically hard, because the closed-form relationships that hold asymptotically become approximations that must be handled carefully when packets

are short and error targets are extreme (Durisi *et al.*, 2016; Polyanskiy *et al.*, 2010)^[14, 26]. Lifecycle-aware and carbon-aware energy accounting is immature, since most models still optimize operational link energy while ignoring embodied energy and the system-level consequences of device proliferation. The reliability of learning-based control is an open question of growing importance, as policies that perform well in training can fail under distribution shift, and the field lacks accepted guarantees for such controllers in mission-critical settings. Finally, the framework itself invites quantification: Turning the qualitative trilemma surface into measured frontiers for representative IoT scenarios, so that the displacement produced by each paradigm can be reported in comparable units, is a natural and valuable next step.

Limitations

This work is a conceptual contribution and carries the limitations of that genre. The framework is an organizing lens, not a predictive model; it does not by itself yield numerical operating points, which require system-specific models and measurement. The trilemma surface is a simplification that suppresses dimensions such as spectral efficiency, security, and cost, which interact with the three axes and which a fuller treatment would incorporate. The mapping in Table 1 records dominant tendencies rather than universal truths, since the precise signature of any technique depends on the deployment and traffic context. These limitations do not undermine the framework's purpose, which is to provide a shared structure for reasoning, but they mark the boundaries within which it should be applied.

Conclusion

Energy, reliability, and latency have become the three governing constraints of next-generation IoT, and they are fundamentally coupled rather than independently tunable. This paper has developed a conceptual framework that treats this coupling as a three-way tension, the ERL trilemma, and that grounds the tension in established theory through three analytical anchors: Shannon capacity, the wideband energy-per-bit limit, and the finite-blocklength coding rate. The framework introduces a trilemma surface on which any green technique can be located, and it distinguishes traversals, which reallocate a fixed budget, from expansions, which enlarge what is achievable. Applying the framework to the dominant green communication paradigms produced a compact mapping, a set of design principles, and an agenda of open challenges. The intent throughout has been to supply a common vocabulary and structure for a problem that the literature has largely addressed in fragments, so that future work on sustainable, dependable, and timely IoT can reason about its trade-offs on shared ground.

References

- Adelanwa A, Basnet A, Anene UN. Performance intelligence models for optimization and outcome measurement in large scale public services. Shodhshauryam, International Scientific Refereed Research Journal. 2023; 6(1):319-354.
- Ahmed KS, Odejebi OD. Resource allocation model for energy-efficient virtual machine placement in data centers. IRE Journals. 2018; 2(3):1-10.
- Ahmed KS, Odejebi OD, Oshoba TO. Algorithmic model for constraint satisfaction in cloud network resource allocation. IRE Journals. 2019; 2(12):516-532.
- Ahmed KS, Odejebi OD, Oshoba TO. Predictive model for cloud resource scaling using machine learning techniques. Journal of Frontiers in Multidisciplinary Research. 2020; 1(1):173-183.
- Akanbi O, Ganiu OS, Sunday EA. A multi-agent AI framework for swarm intelligence in autonomous mobile robot (AMR) fleet coordination. International Journal of Scientific Research in Humanities and Social Sciences. 2025; 2(1):81-109.
- Akinola AS, Adesanya OS, Okafor CM, Dako OF. Value-chain automation in beverage logistics: Throughput, capacity, and cost avoidance via queueing models. International Journal of Scientific Research in Computer Science, Engineering and Information Technology. 2024; 10(4):1112-1132.
- Avranas A, Kountouris M, Ciblat P. Energy-latency tradeoff in ultra-reliable low-latency communication with retransmissions. IEEE Journal on Selected Areas in Communications. 2018; 36(11):2475-2485.
- Bjornson E, Hoydis J, Sanguinetti L. Massive MIMO networks: Spectral, energy, and hardware efficiency. Foundations and Trends in Signal Processing. 2017; 11(3-4):154-655.
- Bockelmann C, Pratas N, Nikopour H, Au K, Svensson T, Stefanovic C, *et al.* Massive machine-type communications in 5G: Physical and MAC-layer solutions. IEEE Communications Magazine. 2016; 54(9):59-65.
- Buzzi S, I C-L, Klein TE, Poor HV, Yang C, Zappone A. A survey of energy-efficient techniques for 5G networks and challenges ahead. IEEE Journal on Selected Areas in Communications. 2016; 34(4):697-709.
- Chen Y, Zhang S, Xu S, Li GY. Fundamental trade-offs on green wireless networks. IEEE Communications Magazine. 2011; 49(6):30-37.
- Di Renzo M, Zappone A, Debbah M, Alouini M-S, Yuen C, de Rosny J, *et al.* Smart radio environments empowered by reconfigurable intelligent surfaces: How it works, state of research, and the road ahead. IEEE Journal on Selected Areas in Communications. 2020; 38(11):2450-2525.
- Din S, Ahmad A, Paul A, Rho S. MGR: Multi-parameter green reliable communication for internet of things in 5G network. Journal of Parallel and Distributed Computing. 2018; 118:34-45.
- Durisi G, Koch T, Popovski P. Toward massive, ultrareliable, and low-latency wireless communication with short packets. Proceedings of the IEEE. 2016; 104(9):1711-1726.
- Edivri J, Oteri O. Predictive capacity planning and resource utilization forecasting models for multi-program and multi-stakeholder IT portfolios. International Journal of Scientific Research in Computer Science, Engineering and Information Technology. 2022; 8(4):826-845.
- Feng D, Jiang C, Lim G, Cimini LJ, Feng G, Li GY. A survey of energy-efficient wireless communications. IEEE Communications Surveys and Tutorials. 2013; 15(1):167-178.
- Idika CN, Salami EO, Ijiga OM, Enyejo LA. Deep learning driven malware classification for cloud-native microservices in edge computing architectures.

- International Journal of Scientific Research in Computer Science, Engineering and Information Technology. 2021; 7(4).
18. Ijiga OM, Oladoye SO, Bamigwojo OV, Ogboji AJ. Techno-economic evaluation of hybrid solar-diesel microgrids for underserved communities using simulation-based load forecasting. *Engineering Science & Technology Journal*. 2023; 4:1-28.
 19. Kumuyi O, Uzoka E, Akeju B, Ozowara DE. Architecture for machine learning-enabled predictive energy management using IoT sensor networks. *International Journal of Advanced Multidisciplinary Research and Studies*. 2024; 4(6):3138-3149.
 20. Mahapatra R, Nijsure Y, Kaddoum G, Hassan NU, Yuen C. Energy efficiency tradeoff mechanism towards wireless green communication: A survey. *IEEE Communications Surveys and Tutorials*. 2016; 18(1):686-705.
 21. Mao Y, You C, Zhang J, Huang K, Letaief KB. A survey on mobile edge computing: The communication perspective. *IEEE Communications Surveys and Tutorials*. 2017; 19(4):2322-2358.
 22. Odejebi OD, Ahmed KS. Performance evaluation model for multi-tenant Microsoft 365 deployments under high concurrency. *IRE Journals*. 2018; 1(11):92-107.
 23. Odejebi OD, Hammed NI, Ahmed KS. Approximation complexity model for cloud-based database optimization problems. *IRE Journals*. 2019; 2(9):1-10.
 24. Oteri O, Edivri J. Applied performance optimization frameworks for managing high-demand enterprise technology programs and system rollouts. *Gyanshauryam, International Scientific Refereed Research Journal*. 2024; 7(1):188-210.
 25. Parvez I, Rahmati A, Guvenc I, Sarwat AI, Dai H. A survey on low latency towards 5G: RAN, core network and caching solutions. *IEEE Communications Surveys and Tutorials*. 2018; 20(4):3098-3130.
 26. Polyanskiy Y, Poor HV, Verdú S. Channel coding rate in the finite blocklength regime. *IEEE Transactions on Information Theory*. 2010; 56(5):2307-2359.
 27. Popovski P, Nielsen JJ, Stefanovic C, de Carvalho E, Strom E, Trillingsgaard KF, *et al.* Wireless access for ultra-reliable low-latency communications: Principles and building blocks. *IEEE Network*. 2018; 32(2):16-23.
 28. Raza U, Kulkarni P, Sooriyabandara M. Low power wide area networks: An overview. *IEEE Communications Surveys and Tutorials*. 2017; 19(2):855-873.
 29. Sefati S, *et al.* Ultra-reliability and low-latency communications on the internet of things based on 5G network: Literature review, classification, and future research view. *Transactions on Emerging Telecommunications Technologies*. 2023; 34:e4770.
 30. Shannon CE. A mathematical theory of communication. *Bell System Technical Journal*. 1948; 27(3):379-423.
 31. Sunday EA, Omoegun GO. Optimizing electrical load distribution for hybrid solar installations in developing economies. *International Journal of Scientific Research in Mechanical and Materials Engineering*. 2019; 3(6):27-47.
 32. Torkudzor MK, *et al.* Energy efficiency, latency and reliability trade-offs in M2M uplink scheduling. *IET Communications*, 2021.
 33. Uduokhai DO, Garba BMP, Sanusi AN, Nwafor MI. Computational modelling of climate-adaptive building envelopes for energy efficiency in tropical regions. *Global Journal of Engineering and Technology Review*. 2025a; 1(3):129-141.
 34. Uduokhai DO, Nwafor MI, Sanusi AN, Garba BMP. Predictive framework for optimizing maintenance schedules in aging public infrastructure systems. *Global Journal of Engineering and Technology Review*. 2025b; 1(3):142-152.
 35. Verdú S. Spectral efficiency in the wideband regime. *IEEE Transactions on Information Theory*. 2002; 48(6):1319-1343.
 36. Zappone A, Jorswieck E. Energy efficiency in wireless networks via fractional programming theory. *Foundations and Trends in Communications and Information Theory*. 2015; 11(3-4):185-396.