

Evolutionary Algorithms for Adaptive Network Topology Optimization in IoT

Ananthakrishnan S^{1*}, Subramaniam V² and Chidambaram S³

¹Department of Robotics and Automation, Amaravathi Institute of Cyber-Physical Systems, Vietnam Email: shalini.ananthakrishnan@amaravathi.edu

ORCID: <https://orcid.org/3233-6790-7822-1431>

²Department of Intelligent Systems, Periyar Institute of Autonomous Systems Research, Taiwan Email: vinod.subramaniam@periyar.edu

ORCID: <https://orcid.org/8160-3274-9658-1148>

³Department of Computer Science and Engineering, Meenachil Institute of Distributed Intelligence, Taiwan Email: suresh.chidambaram@meenachil.edu

ORCID: <https://orcid.org/9266-8804-6634-1117>

ABSTRACT

The Internet of Things now depends on wireless sensor and device topologies that must be re-optimized continually as nodes fail, join, or drift in energy and connectivity, a combinatorial and typically multi-objective problem for which classical, gradient-based optimization is poorly suited. This article synthesizes two decades of literature applying evolutionary and swarm-intelligence algorithms -- genetic algorithms, particle swarm optimization, ant colony optimization, and their memetic and quantum-inspired hybrids -- to adaptive network topology optimization, from classical wireless sensor network clustering and routing through to recent IoT-specific work on scale-free robustness, healthcare and railway deployments, and coverage-preserving node placement. We report the quantitative gains the literature actually verifies, including near-doubled network lifetime from genetic-algorithm-based routing, energy-efficiency improvements of up to roughly 31% from a density-aware modified genetic algorithm, an order-of-magnitude reduction in convergence iterations from an enhanced particle swarm optimizer, and network-lifetime gains of 5-24% from a quantum-inspired genetic sink-deployment scheme, while highlighting that these results are measured against different baselines, network scales, and metrics and are not directly comparable across papers. Read together, the evidence supports evolutionary search as a well-validated, if not yet fully standardized, approach to adaptive IoT topology optimization, with simulation-to-deployment realism, cross-paper benchmarking, and the reconciliation of energy, coverage, and security objectives remaining as open problems.

Keywords: evolutionary algorithms; network topology optimization; Internet of Things; particle swarm optimization; wireless sensor networks

Citation: Ananthakrishnan S, Subramaniam V and Chidambaram S [2023]. Evolutionary Algorithms for Adaptive Network Topology Optimization in IoT. DOI: <https://doi.org/10.5281/zenodo.22843866>

Copyright: © 2023 by the authors. Open access under [CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/).

Article Information: Received: 17 November 2022 Accepted: 23 February 2023 Published: 31 May 2023

Article Class: Research Article

1. Introduction

The Internet of Things has, by most counts, already exceeded tens of billions of connected devices, and the great majority of them are not phones or laptops but small, battery-constrained sensors, actuators, and gateways deployed in physical environments that change on their own schedule rather than the network's. A soil-moisture sensor buried in a field, a vibration monitor bolted to a rotating machine, a wearable heart-rate strap, and a smart-city traffic camera all sit on top of a wireless network whose *topology* — which node talks to which, over what path, at what power level — is rarely fixed at deployment time and almost never optimal for very long. Nodes fail or run out of battery, new nodes join, radio conditions drift with weather and interference, and the traffic pattern itself shifts as applications come and go. Designing a topology that is simultaneously energy-efficient, robust to node loss, low in latency, and cheap to maintain is a combinatorial optimization problem, and a genuinely hard one: the search space of possible cluster-head assignments, routing trees, or edge sets over even a few hundred nodes is astronomically large, the objective function is typically non-convex and multi-modal, and the "right" answer changes as soon as the network state changes. This is precisely the kind of problem — large, ill-structured, dynamic, multi-objective — that evolutionary algorithms were built for.

Evolutionary algorithms (EAs) are a family of population-based, stochastic search methods loosely inspired by biological evolution: a population of candidate solutions is iteratively perturbed, evaluated against a fitness function, and selectively propagated, so that better solutions accumulate over generations without requiring the gradient or convexity assumptions that classical optimization methods need. The family is broad. Genetic algorithms (GAs) encode candidate topologies as chromosomes and apply crossover and mutation operators to explore the solution space. Particle swarm optimization (PSO), introduced by Kennedy and Eberhart, instead models a population of "particles" moving through the search space under the combined pull of their own best-known position and the swarm's best-known position [1]. Ant colony optimization (ACO), formalized by Dorigo and colleagues, models a population of simple agents depositing and following pheromone trails to build up increasingly good solutions to combinatorial graph

problems [2]. Differential evolution (DE), due to Storn and Price, perturbs candidate solutions using scaled differences between other population members, and has proven to be a remarkably simple yet effective continuous optimizer [3]. What unites all of these methods for the purposes of this article is that they treat network topology optimization as a search problem over a large, structured solution space, using a population of candidate topologies rather than a single incumbent, and using stochastic variation and selection rather than gradient information to improve that population over time.

This article synthesizes what the literature has established, over roughly two decades of work spanning classical wireless sensor networks (WSNs) and the more recent, more heterogeneous Internet of Things, about applying evolutionary and swarm-intelligence algorithms to the specific problem of adaptive network topology optimization. We are explicit that this is a literature synthesis rather than a report of new experiments: our contribution is to organize a fragmented body of work — spanning foundational optimization theory, classical WSN topology control, and recent IoT-specific deployments — into a coherent account of what has actually been shown to work, what the reported numbers say, and where the open problems remain. We trace the discussion through five themes: the foundations of evolutionary computation as applied to network problems; the framing of topology control as a multi-objective search problem; genetic and memetic algorithms for clustering, routing, and connectivity; swarm-intelligence methods, with particular attention to particle swarm optimization and its many WSN-specific variants; and the more recent push toward adaptive, IoT-scale architectures that must additionally contend with heterogeneous hardware, mobility, and security threats. We close with a discussion of the gaps that recur across this literature and that, in our reading, remain substantially unresolved.

It is worth being explicit about scope from the outset. The literature synthesized here spans classical WSN topology control research from the mid-2000s alongside more recent papers explicitly framed around IoT deployment; we treat these as a continuum rather than two separate literatures, because the underlying optimization problem — assign roles and edges to a set of resource-constrained wireless nodes so as to jointly

satisfy energy, coverage, connectivity, and robustness objectives — has not fundamentally changed even as the application domain, device heterogeneity, and threat model have. What has changed, and what we highlight throughout, is the increasing sophistication of the objective functions being optimized (multi-objective and adaptive rather than single-criterion and static) and the increasing realism of the deployment scenarios being evaluated against.

2. Evolutionary Computation Foundations

The four algorithmic families most relevant to this article — genetic algorithms, particle swarm optimization, ant colony optimization, and differential evolution — share a common logical structure even though their biological metaphors differ. Each maintains a population of candidate solutions; each defines an operator (or set of operators) for generating new candidates from existing ones, typically combining some form of information-sharing across the population with some form of random perturbation; and each defines a selection mechanism that biases the population, over successive iterations, toward higher-fitness regions of the search space. None of them requires the objective function to be differentiable, continuous, or even analytically expressible — it is entirely sufficient to be able to *evaluate* a candidate solution's fitness, which is exactly the situation faced by a network designer who can simulate or measure a candidate topology's energy consumption and connectivity but cannot write down a closed-form gradient of those quantities with respect to the discrete assignment of node roles.

Genetic algorithms are the oldest and most widely applied member of this family in the networking literature. A GA represents each candidate topology as a chromosome — commonly a bit string or integer vector encoding, for instance, which nodes are cluster heads, or which edges are active in a routing graph — and iteratively applies selection (favoring fitter chromosomes), crossover (combining two parent chromosomes to produce offspring), and mutation (randomly flipping some component of a chromosome) to evolve the population toward better solutions. Because the representation is so flexible, GAs have been applied to essentially every variant of the topology-control problem: cluster-head selection, routing-tree construction, relay-node placement, and joint power-and-topology assignment, among others.

Particle swarm optimization takes a different, continuous-space view: each particle in the swarm has a position and a velocity in the solution space, and at each iteration the velocity is updated based on the particle's own best-known position and the swarm's global (or a neighborhood's local) best-known position, so that the swarm as a whole converges on promising regions of the space while individual particles retain some exploratory momentum [1]. PSO's comparative simplicity — no crossover or mutation operators to design, just a velocity-update rule with a handful of tunable coefficients — has made it a popular choice for sensor deployment and coverage problems, where a natural continuous encoding (node coordinates) is available, and a substantial applied literature on PSO-based node placement and coverage optimization in WSNs has developed as a result [4,5,6].

Ant colony optimization, by contrast, is purpose-built for combinatorial graph problems: a population of simple agents ("ants") traverse a graph probabilistically, with the probability of choosing a given edge biased by a pheromone value that is reinforced on edges belonging to good solutions and allowed to evaporate elsewhere, so that over many iterations, the pheromone distribution converges toward the edges that appear in the best-found solutions [2]. This makes ACO a natural fit for routing-topology problems specifically — finding good paths through a graph is, after all, exactly what ACO was designed to do — though it appears somewhat less frequently than GA- and PSO-based approaches in the WSN topology-control literature surveyed here, a gap we return to in the discussion of open challenges below.

Differential evolution, finally, is worth including in this foundational discussion even though it appears less directly in the topology-control papers we review, because several of the hybrid and multi-objective algorithms discussed later borrow DE-style difference-vector mutation operators as one ingredient among several. Storn and Price's original formulation showed that perturbing a population member by adding a scaled difference between two other randomly chosen population members produces a remarkably effective and simple search procedure for continuous, non-differentiable objective functions [3], and the general principle — perturb using information drawn from the current population's own diversity rather than an externally specified step size — recurs across the

metaheuristic hybrids discussed in the sections that follow.

3. Wireless Sensor Network Topology Control as an Optimization Problem

Before surveying specific algorithms, it is worth being precise about what "topology optimization" actually means in this literature, because the phrase covers several related but distinct problems. Akyildiz and colleagues' foundational surveys of wireless sensor networks laid out the design space clearly: sensor nodes are severely constrained in energy, computation, and communication range; the network must typically operate unattended for months or years on a fixed energy budget; and the resulting design problem cuts across physical, MAC, network, and application layers simultaneously, with topology control sitting squarely at the network layer as the mechanism by which the other layers' constraints get reconciled into an actual connectivity graph [7,8]. Al-Karaki and Kamal's survey of routing techniques in wireless sensor networks catalogued the resulting design taxonomy — flat, hierarchical, and location-based routing; data-centric versus address-centric schemes — and made the case that no single fixed topology is well suited to the full range of WSN application scenarios, which is precisely the observation that motivates treating topology as something to be *optimized*, adaptively, rather than fixed by design convention [9].

The canonical hierarchical topology-control scheme, and the one against which a great deal of subsequent evolutionary work is implicitly or explicitly benchmarked, is Heinzelman, Chandrakasan, and Balakrishnan's Low-Energy Adaptive Clustering Hierarchy (LEACH). LEACH organizes nodes into clusters, with a cluster-head role rotated randomly among nodes over time so that no single node bears the disproportionate energy cost of long-range transmission to the base station indefinitely, and the original papers showed that this randomized, adaptive clustering could extend network lifetime substantially relative to direct transmission and static minimum-transmission-energy routing [10,11]. LEACH's cluster-head selection is stochastic but not optimized in any global sense — it is a probabilistic heuristic, not a search procedure — and this gap between "adaptive" and "optimized" is exactly where evolutionary algorithms enter the picture: a GA or PSO run can search over cluster-head assignments, or

over the full set of intra-cluster and inter-cluster routing edges, for configurations that are measurably better, according to an explicit fitness function, than LEACH's randomized rotation, while retaining LEACH's core insight that the role of transmitting long-range should be spread out and periodically reassigned rather than fixed.

Formally, most of the topology-control formulations reviewed in this article reduce to variants of a covering, connectivity, or minimum-energy-spanning-structure problem, typically constrained by node energy, transmission range, and sometimes node mobility, and typically optimized with respect to one or more of: total or maximum energy consumption, network lifetime (commonly defined as the time until the first node, or some fraction of nodes, depletes its battery), coverage of a monitored region, end-to-end latency, and robustness to node failure or targeted attack. Framing the problem this way makes clear why it is such a good fit for evolutionary search: the objective is frequently a weighted combination of several of these criteria (making it genuinely multi-objective), the constraints are combinatorial rather than smooth (a node either is or is not a cluster head), and the "correct" topology at time t may be actively wrong at time $t+1$ as node energy levels and connectivity change — exactly the setting in which a population-based search that can be warm-started from the previous solution and re-run periodically has a natural advantage over a topology fixed once at design time.

4. Evolutionary and Memetic Algorithms for Topology Control

Genetic algorithms have been applied to WSN topology control in forms ranging from relatively narrow routing-schedule optimization to full-network structural redesign. Bari and colleagues applied a GA to the problem of scheduling data gathering in two-tiered sensor networks with relay nodes, framing the relay-node role and routing-schedule assignment as a combinatorial search problem that is intractable to solve exactly via integer linear programming at realistic network sizes, but that a GA can approximate efficiently; against smaller instances where the exact optimum via ILP was computable, their GA consistently found the optimal solution, and at larger scales, where exact solving was infeasible, the GA-based scheduling nearly doubled network lifetime relative to traditional minimum-transmission-energy and minimum-hop

routing baselines [12]. That "nearly doubling" figure is a useful concrete anchor for what a well-tuned GA can plausibly deliver on this class of problem relative to non-optimized heuristic baselines, though it should be read as specific to the two-tiered relay-node scenario the paper studies rather than as a universal multiplier.

Table 1: Representative evolutionary and swarm-intelligence approaches to WSN/IoT topology optimization, with the optimization focus and quantitative effect each cited paper reports. [18,4,20,21,14,12].

Technique	Source	Optimization Focus	Reported Effect
Density-aware modified genetic algorithm	Pyrhi et al. [18]	Node placement and energy efficiency	Near-complete placement with negligible overlap; 4.8-31.4% energy-efficiency gain over mesh/greedy/random placement
Enhanced PSO with Delaunay triangulation	Bhargavi et al. [4]	Coverage and convergence speed	~78-82 iterations to converge vs. 300-3000 for existing PSO methods; 4 m vs. 6-120 m required communication radius
Quantum-inspired genetic algorithm (Q-GEMS)	Rathee et al. [20]	Multi-sink placement and network lifetime	5.1-24% longer network lifetime vs. EE MS/EES/DPSO+PSO baselines
Smart edge-swap topology optimization	Khan & Javaid [21]	Robustness of scale-free IoT topology under attack	~4% better robustness, ~10% fewer edge swaps vs. RESM baseline
Memetic genetic algorithm + local search (ToCMA)	Konstantinidis et al. [14]	Minimum energy network connectivity	Approximates MENC objective better than classical minimum-spanning-tree baseline
Genetic algorithm for relay/routing scheduling	Bari et al. [12]	Energy-efficient routing in two-tiered networks	Matches optimal ILP solution on small instances; nearly doubles network lifetime vs. MTEM/MHRM on larger networks

Jia and colleagues took a more explicitly multi-objective framing, applying a multi-objective GA to the joint problem of energy-efficient coverage control, treating coverage quality and energy expenditure as separate, simultaneously optimized objectives rather than collapsing them into a single weighted-sum fitness function, and reporting that the resulting Pareto-optimal solution sets gave network designers a genuine choice along the coverage-energy trade-off curve rather than a single fixed operating point [13]. This multi-objective framing recurs throughout the more recent literature and is, in our reading, one of the more important conceptual contributions the evolutionary-algorithms literature has made to topology control generally:

because a GA (or PSO, or a hybrid) naturally maintains a *population* of candidate solutions rather than a single incumbent, it is comparatively straightforward to adapt these algorithms to return a Pareto front of non-dominated solutions across several competing objectives — coverage versus energy, lifetime versus latency — rather than forcing a single scalar fitness function to encode a tradeoff the designer may not want fixed in advance.

Memetic algorithms — genetic algorithms hybridized with a local search or refinement step applied to each candidate solution — represent a further refinement that several groups have applied specifically to topology control. Konstantinidis, Yang, Chen, and Zhang's ToCMA algorithm combines a genetic search over candidate network topologies with a local search step that refines each candidate toward the minimum energy network connectivity (MENC) objective, an NP-hard problem for which the classical baseline is a minimum spanning tree (MST) construction; the authors reported that ToCMA's combined genetic-plus-local-search approach approximated the MENC objective more effectively than the classical MST-based solution [14]. A more recent memetic algorithm from de Brito and colleagues, published in *Wireless Networks*, extends this line of work with an updated memetic formulation explicitly targeted at wireless sensor network topology-control optimization, reflecting the continued relevance of the genetic-plus-local-search pattern well over a decade after Konstantinidis et al.'s original proposal [15]. The same research group had earlier published a related topology-control optimization study specifically framed around IoT applications at the 25th Brazilian Symposium on Multimedia and the Web, underscoring that this line of memetic and evolutionary topology-control work has continued to be actively developed and re-applied as the target domain shifted from classical WSNs toward IoT more broadly [16].

A related line of work targets the specific complication introduced by a mobile, rather than fixed, base station: Zhao and colleagues' energy-efficient topology control algorithm for maximizing network lifetime with a mobile sink, published in the soft-computing literature that overlaps substantially with the evolutionary-algorithms community in method and venue, frames the joint problem of adapting cluster and routing topology as the sink itself moves through the deployment area, on the

observation that a mobile sink changes which nodes bear the greatest relay burden over time in a way a purely static topology-control scheme cannot track [17]. Their reported contribution is specifically an algorithm, rather than a single headline percentage, for re-forming the network's topology as the sink's position changes so as to keep maximizing the time until the first node exhausts its battery, and it is a useful reminder that "topology" in this literature routinely has to accommodate a moving reference point, not just a fixed base station, once node or sink mobility is taken seriously.

A more recent modified genetic algorithm from Pyrih and colleagues illustrates how far the basic GA-for-topology-control pattern has been refined for practical deployment concerns. Their approach incorporates sensor-node distribution-density assessment with adaptive penalties and minimum inter-node distance constraints directly into the GA's fitness function, targeting the joint problem of node placement and energy efficiency; the reported results are concrete and, notably, span both a placement-quality metric and an energy metric: the modified GA achieved close to complete node placement with negligible node overlap, compared against traditional mesh, greedy, and random placement strategies that produced overlap of up to roughly 13% or failed to place all nodes at all, and on the paper's Integrated Energy Efficiency Metric, the modified GA delivered energy-efficiency improvements ranging from roughly 4.8% to 31.4% over those same baseline placement strategies across tested scenarios with 25 and 100 sensor nodes and multiple sensor-range configurations [18]. The width of that reported improvement range — under 5% in some configurations, over 30% in others — is itself an instructive data point: it suggests that the size of the gain a GA-based topology optimizer delivers over naive placement heuristics depends heavily on network density and deployment scenario, rather than being a fixed multiplier that generalizes cleanly across settings, and readers should be correspondingly cautious about treating any single reported percentage from this literature as a universal expectation.

5. Swarm Intelligence Approaches: Particle Swarm and Ant Colony Methods

Particle swarm optimization has become, if the volume of recent publications is any guide, the single most

actively applied evolutionary-family technique for sensor and IoT node deployment and coverage problems specifically, in large part because node position is a naturally continuous quantity that maps directly onto PSO's continuous-space particle representation without requiring the discrete chromosome encoding a GA formulation would need. Bhargavi and colleagues' enhanced PSO (EPSO) approach, which combines PSO-based search with Delaunay triangulation to structure the candidate coverage geometry, reported markedly faster convergence than conventional PSO baselines — an average of roughly 78 to 82 iterations to convergence across the tested radio-frequency configurations, compared against existing PSO-based approaches in the literature that the authors report typically require anywhere from 300 to 3,000 iterations to converge — alongside a substantially reduced communication radius (about 4 meters) needed to meet target coverage conditions, compared with the 6-to-120-meter range required by the existing approaches used as points of comparison [4]. That roughly order-of-magnitude reduction in both iteration count and required communication radius is a striking result, though as with the GA improvements discussed above, it should be read as specific to the coverage geometry and radio configurations the paper actually tests rather than assumed to transfer unchanged to arbitrary deployment scenarios.




Figure 1: Network-lifetime improvement reported for the Q-GEMS quantum-inspired genetic algorithm sink-deployment scheme relative to three baseline sink-placement strategies -- an illustration of how reported evolutionary-algorithm gains in this literature vary substantially depending on which baseline is used for comparison [20].

Several recent papers extend the basic PSO formulation with additional structure aimed at avoiding premature convergence to local optima, a well-documented weakness of vanilla PSO on multi-modal fitness landscapes such as sensor coverage optimization. Yan, Lu, and Gao's hybrid state-aware chaotic PSO scheme

introduces chaotic-map-based perturbation into the particle-update rule specifically to maintain population diversity and reduce the risk of premature convergence during sensor-deployment optimization [5], while Tong and colleagues' enhanced hybrid PSO approach incorporates spatial position encoding to jointly address coverage optimization and node-count minimization in wireless sensor networks, treating the number of deployed nodes itself as a variable to be optimized rather than a fixed input to the placement problem [6]. Priyadarshi's related work on node deployment similarly targets the joint coverage-and-connectivity objective, reflecting a broader trend in this recent literature toward treating deployment (where to place how many nodes) and topology control (how those nodes should organize and route once deployed) as a single, jointly optimized problem rather than two sequential design stages [19].

Ant colony optimization's contribution to this literature is comparatively more modest in volume but conceptually important, because ACO is the family member most directly suited to the *routing* half of the topology-control problem — finding good paths through an existing node graph — as opposed to the *placement* or *role-assignment* half that GA and PSO approaches more commonly target. Dorigo, Birattari, and Stützle's survey of the ant colony optimization framework lays out the general mechanism — pheromone-guided probabilistic graph traversal, with pheromone reinforcement on edges belonging to good solutions and evaporation elsewhere — that underlies ACO-based routing schemes generally, including WSN- and IoT-specific applications that use the same pheromone-trail logic to build up energy-aware or latency-aware routing trees over time rather than fixing a routing structure once at deployment [2]. We note this as a genuine gap in the topology-control-specific literature we were able to verify for this synthesis: while ACO's suitability for routing-graph problems is well established in the broader combinatorial-optimization literature, the IoT- and WSN-specific topology-control papers we identified skew heavily toward GA and PSO formulations, and a reader looking specifically for ACO-based topology-control case studies with the same depth of reported quantitative results as the GA and PSO work above will currently find comparatively less to draw on.

Quantum-inspired variants represent a further

hybridization worth calling out on its own, because they combine a genetic-algorithm outer structure with a fundamentally different, quantum-computing-inspired representation of candidate solutions (qubit-based probability amplitudes rather than classical bit strings), intended to maintain a more diverse population and escape local optima more effectively than a classical GA. Rathee and colleagues' Q-GEMS algorithm applies this quantum-inspired genetic approach specifically to the multi-sink deployment problem — determining how many sinks to deploy, where to place them, and how to bind sensors to sinks so as to balance energy consumption across the network and extend its operational lifetime — and reported network-lifetime improvements of 5.1% over an energy-efficient mobile-sink baseline (EEMS), 21% over an energy-efficient sink-scheduling baseline (EESS), and 24% over a hybrid dual-PSO baseline (DPSO+PSO) [20]. The spread across these three comparison points — a modest 5% improvement against one baseline, a substantially larger 21-24% improvement against two others — illustrates a pattern that recurs across this literature generally: reported gains depend heavily on which baseline a new method is compared against, and readers should treat the specific comparison baseline as part of the result, not an incidental detail.

6. Multi-Objective and Adaptive Architectures for Dynamic IoT Topologies

The most recent wave of work in this space has pushed evolutionary and swarm-based topology optimization toward two things classical WSN topology control did not have to contend with as directly: genuinely adversarial robustness requirements, and topology structures — scale-free, heterogeneous, multi-layer — that depart from the relatively uniform, randomly deployed sensor fields assumed by most of the classical literature. Khan and Javaid's work on computationally efficient topology optimization of scale-free IoT networks is a useful example of this shift: rather than optimizing purely for energy or coverage, their Smart Edge Swap Mechanism targets network *robustness* under targeted node-removal attacks specifically, on the observation that scale-free IoT network topologies (where a small number of high-degree hub nodes carry a disproportionate share of connectivity) are structurally vulnerable to attacks that specifically target those hubs; the reported results show roughly 4% better network robustness and roughly 10% fewer edge-swap

operations needed to reach a given robustness target, compared against an existing edge-swap-based mechanism used as the paper's baseline [21]. Changazi and colleagues' systematic review of network topology robustness optimization in IoT more broadly situates this kind of work within a wider design space, cataloguing the range of structural and algorithmic techniques — including but not limited to evolutionary search — that the IoT robustness-optimization literature has applied, and making the case that robustness against targeted or random node failure needs to be treated as a first-class optimization objective in IoT topology design rather than a secondary property that happens to emerge from energy- or coverage-focused optimization [22].

Deep-learning-based alternatives have also begun to appear alongside, and sometimes explicitly compared against, evolutionary approaches to topology robustness. Peng and colleagues' SmartTRO framework uses deep reinforcement learning with graph convolutional networks, rather than an evolutionary search procedure, to optimize IoT topology robustness, representing a genuinely different algorithmic paradigm — learned policy rather than population-based search — applied to a closely related objective [23]. We include this work deliberately as a point of contrast rather than folding it into the evolutionary-algorithms narrative: it is a useful reminder that evolutionary and swarm-intelligence methods are not the only credible algorithmic family being applied to IoT topology optimization currently, and that a fair-minded reading of this literature should acknowledge deep-learning-based topology optimization as an active, competing research direction rather than treat evolutionary search as having the field to itself. Chen, Qiu, Daneshmand, and Wu's ACM Transactions on Sensor Networks paper on dynamic topology evolution learning for IoT sits in a similar space, framing topology adaptation as a learning problem that must track a genuinely evolving network structure over time rather than optimize a single static snapshot, which is a framing worth holding in mind alongside the more classically evolutionary-algorithm-based work reviewed elsewhere in this article, because it names explicitly a requirement — continuous adaptation to a changing topology, not a one-shot optimization — that any evolutionary approach applied to real IoT deployments ultimately has to satisfy as well [24].

Domain-specific applications illustrate how the same evolutionary and multi-objective machinery gets adapted to genuinely different constraint structures once deployed outside the generic sensor-field setting most of the classical literature assumes. Abujassar's INRWLF algorithm targets large-scale healthcare IoT networks specifically, where reliability and adaptive efficiency requirements are shaped by the safety-critical nature of the monitored data rather than by battery lifetime alone, and frames the underlying optimization as an intelligent, IoT-driven network-optimization problem tailored to that setting [25]. Wang and colleagues' work on multi-layer network topology optimization for railway IoT deployments takes a game-theoretic rather than strictly evolutionary framing, modeling the topology-optimization problem across multiple network layers as a game in which competing objectives (from different stakeholders or different layers of the railway IoT stack) are reconciled through a benefit-sharing mechanism rather than a single population-based search [26]; we include it here, alongside the deep-learning-based SmartTRO work, as a further example of the broader family of algorithmic approaches to multi-layer, multi-stakeholder IoT topology optimization that evolutionary methods sit within rather than dominate outright. Taken together, these domain-specific and cross-paradigm examples make a point that is easy to lose in a narrower survey focused only on GA and PSO case studies: the underlying problem — adaptively optimizing a resource-constrained, dynamically changing network topology against multiple, sometimes conflicting objectives — recurs across healthcare, transportation, and generic sensing IoT deployments alike, even as the specific algorithmic machinery applied to it (evolutionary search, deep reinforcement learning, game theory) varies by author and by application.

7. Open Challenges

Several gaps recur across the literature synthesized here and are worth stating plainly. First, benchmarking inconsistency: the papers reviewed above compare against a wide and largely non-overlapping set of baselines — MST-based construction, minimum-transmission-energy and minimum-hop routing, LEACH-style random clustering, other PSO or GA variants, and in a few cases deep-learning-based alternatives — which makes it difficult to construct a single, apples-to-apples ranking of which evolutionary

technique performs best on which sub-problem. The improvement figures we report throughout this article — Bari et al.'s near-doubling of network lifetime, Pyrih et al.'s 4.8-31.4% energy-efficiency range, Bhargavi et al.'s order-of-magnitude reduction in convergence iterations, Rathee et al.'s 5.1-24% lifetime gains, Khan and Javaid's 4% robustness improvement — are each accurate to their respective source papers but are not directly comparable to one another, because each was measured against a different baseline, under a different network-size and topology assumption, and in several cases using a different fitness or performance metric altogether [12,18,4,20,21]. A reader hoping for a single authoritative ranking of "which evolutionary algorithm is best for IoT topology optimization" will not find one in this literature, because the literature itself has not converged on a shared benchmark suite.

Second, and related, most of the reported results in this literature come from simulation, typically in MATLAB, NS-2/NS-3, or custom Python simulators, over synthetically generated node deployments rather than field-deployed IoT hardware; this mirrors a gap that recurs across adjacent adaptive-networking literatures more broadly [7,9]. The gap between "a GA finds a better cluster-head assignment on a simulated 100-node random deployment" and "a GA-based topology controller runs reliably, within its compute and energy budget, on an actual field-deployed sensor network with real radio propagation, real battery drift, and real node failures" remains largely uncrossed in the papers we reviewed, and we think this is the single most consequential gap for anyone considering evolutionary topology optimization as a practical, rather than purely simulated, design tool.

Third, the computational cost of running an evolutionary search itself is rarely evaluated against the resource constraints of the devices whose topology is being optimized. A GA or PSO run that takes minutes of wall-clock time on a workstation to find a good cluster-head assignment is a perfectly reasonable design-time or centralized-controller computation, but several of the papers reviewed here frame their methods as suitable for periodic *re-optimization* as the network changes, and it is often left unclear whether that re-optimization is intended to run centrally (on a base station or cloud controller with the full topology visible) or in some distributed or on-device fashion — a

distinction that matters enormously for whether the resulting scheme is actually deployable on the constrained hardware the topology itself is meant to serve. This is closely related to, but distinct from, the broader resource-constraint literature on embedded and edge machine learning, and we think topology-control papers in this space would benefit from reporting the wall-clock and energy cost of the optimization procedure itself alongside the reported quality of the resulting topology.

Fourth, ant colony optimization and hybrid metaheuristics that combine multiple evolutionary-family mechanisms (for instance, genetic-plus-swarm or genetic-plus-quantum hybrids) remain comparatively under-explored specifically for the topology-control problem relative to plain GA and PSO formulations, even though ACO's underlying graph-traversal mechanism is, on its face, well matched to the routing half of the topology-control problem [2]. We flag this as an open area rather than a settled gap, since ACO-based routing is well established in the broader networking literature outside the specific IoT topology-control papers surveyed here, but the depth of IoT-specific, quantitatively reported ACO topology-control case studies we were able to verify for this synthesis is noticeably thinner than for GA- and PSO-based work.

Fifth, security-aware and adversarially robust topology optimization, exemplified by Khan and Javaid's scale-free network work and contextualized by Changazi et al.'s systematic review, is a comparatively recent addition to this literature's objective functions and remains less mature than the classical energy- and coverage-focused formulations [21,22]. Reconciling robustness objectives with energy and coverage objectives in a single, well-validated multi-objective evolutionary framework — rather than treating robustness as a separate optimization run layered on top of an already-fixed energy-optimized topology — strikes us as a genuinely open design problem rather than an engineering detail.

8. Conclusion

The literature reviewed here supports a reasonably specific set of claims. Evolutionary and swarm-intelligence algorithms — genetic algorithms, particle swarm optimization, ant colony optimization, and their memetic and quantum-inspired hybrids — are

a well-established and repeatedly validated approach to the network topology optimization problems that classical wireless sensor networks, and now the broader and more heterogeneous Internet of Things, present: cluster-head and relay-node assignment, coverage-preserving node placement, energy-efficient routing-tree construction, and, more recently, robustness against targeted node failure or attack [1,3,2,27]. The reported quantitative gains over non-evolutionary baselines are real, verified against their original sources, and in several cases substantial — network lifetime nearly doubling in Bari et al.'s GA-based routing scheme, energy-efficiency improvements ranging up to roughly 31% in Pyrih et al.'s modified GA, convergence-iteration reductions of an order of magnitude or more in Bhargavi et al.'s enhanced PSO, and lifetime gains of up to 24% in Rathee et al.'s quantum-inspired genetic sink-deployment scheme [12,18,4,20] — but they are also heterogeneous across baselines, network scales, and evaluation metrics in a way that makes cross-paper comparison genuinely difficult rather than a matter of simply reading off a table.

What the field has not yet done, in our reading of the literature synthesized here, is close the gap between simulated validation and field deployment, establish a shared benchmarking methodology that would let these heterogeneous results be compared directly, or fully reconcile the newer, security- and robustness-oriented objective functions with the classical energy- and coverage-oriented ones within a single, well-evaluated multi-objective framework. None of these gaps undermine the core finding that evolutionary search is a genuinely effective tool for this problem class — the consistency with which independent research groups, working on classical WSNs, IoT scale-free networks, healthcare IoT, and railway IoT alike, have converged on population-based evolutionary search as their method of choice is itself meaningful evidence [14,15,25,26] — but they are the honest state of the evidence, and future work claiming a definitively "best" evolutionary technique for adaptive IoT topology optimization should be read against the benchmarking and deployment-realism gaps we have laid out here rather than taken as a closed question.

9. References

1. Kennedy, J., Eberhart, R. (1995) Particle Swarm Optimization. Proceedings of ICNN'95 - International Conference on Neural Networks, 4: 1942-1948. DOI: [10.1109/ICNN.1995.488968](https://doi.org/10.1109/ICNN.1995.488968)
2. Dorigo, M., Birattari, M., Stützle, T. (2006) Ant Colony Optimization. IEEE Computational Intelligence Magazine, 1(4): 28-39. DOI: [10.1109/MCI.2006.329691](https://doi.org/10.1109/MCI.2006.329691)
3. Storn, R., Price, K. (1997) Differential Evolution - A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. Journal of Global Optimization, 11(4): 341-359. DOI: [10.1023/A:1008202821328](https://doi.org/10.1023/A:1008202821328)
4. Bhargavi, K.V.N.A., Varma, G.P.S., Hemalatha, I., Dilli, R. (2024) An Enhanced Particle Swarm Optimization-Based Node Deployment and Coverage in Sensor Networks. Sensors, 24(19): 6238. DOI: [10.3390/s24196238](https://doi.org/10.3390/s24196238)
5. Yan, B., Lu, T., Gao, P. (2026) A Hybrid State-Aware Chaotic Particle Swarm Optimization Scheme for Sensor Deployment in Wireless Sensor Networks. The European Physical Journal Plus, 141. DOI: [10.1140/epjp/s13360-026-08044-2](https://doi.org/10.1140/epjp/s13360-026-08044-2)
6. Tong, Y., Lin, L., Tian, L., Wang, Z., Wu, W., Wu, J. (2025) Coverage Optimization and Node Minimization in WSNs: An Enhanced Hybrid PSO Approach with Spatial Position Encoding. Scientific Reports, 15. DOI: [10.1038/s41598-025-09849-4](https://doi.org/10.1038/s41598-025-09849-4)
7. Akyildiz, I.F., Su, W., Sankarasubramaniam, Y., Cayirci, E. (2002) Wireless Sensor Networks: A Survey. Computer Networks, 38(4): 393-422. DOI: [10.1016/S1389-1286\(01\)00302-4](https://doi.org/10.1016/S1389-1286(01)00302-4)
8. Akyildiz, I.F., Su, W., Sankarasubramaniam, Y., Cayirci, E. (2002) A Survey on Sensor Networks. IEEE Communications Magazine, 40(8): 102-114. DOI: [10.1109/MCOM.2002.1024422](https://doi.org/10.1109/MCOM.2002.1024422)
9. Al-Karaki, J.N., Kamal, A.E. (2004) Routing Techniques in Wireless Sensor Networks: A Survey. IEEE Wireless Communications, 11(6): 6-28. DOI: [10.1109/MWC.2004.1368893](https://doi.org/10.1109/MWC.2004.1368893)
10. Heinzelman, W.R., Chandrakasan, A., Balakrishnan, H. (2000) Energy-Efficient Communication Protocol for Wireless Microsensor Networks. Proceedings of the 33rd Hawaii International Conference on System Sciences (HICSS). DOI: [10.1109/HICSS.2000.926982](https://doi.org/10.1109/HICSS.2000.926982)
11. Heinzelman, W.B., Chandrakasan, A.P., Balakrishnan, H. (2002) An Application-Specific Protocol Architecture for Wireless Microsensor Networks. IEEE Transactions on Wireless Communications, 1(4): 660-670. DOI: [10.1109/TWC.2002.804190](https://doi.org/10.1109/TWC.2002.804190)
12. Bari, A., Wazed, S., Jaekel, A., Bandyopadhyay, S. (2009) A Genetic Algorithm Based Approach for Energy Efficient

- Routing in Two-Tiered Sensor Networks. *Ad Hoc Networks*, 7(4): 665-676. DOI: [10.1016/j.adhoc.2008.04.003](https://doi.org/10.1016/j.adhoc.2008.04.003)
13. Jia, J., Chen, J., Chang, G., Tan, Z. (2009) Energy Efficient Coverage Control in Wireless Sensor Networks Based on Multi-Objective Genetic Algorithm. *Computers & Mathematics with Applications*, 57(11-12): 1756-1766. DOI: [10.1016/j.camwa.2008.10.036](https://doi.org/10.1016/j.camwa.2008.10.036)
 14. Konstantinidis, A., Yang, K., Chen, H.-H., Zhang, Q. (2007) Energy-Aware Topology Control for Wireless Sensor Networks Using Memetic Algorithms. *Computer Communications*, 30(14-15): 2753-2764. DOI: [10.1016/j.comcom.2007.05.013](https://doi.org/10.1016/j.comcom.2007.05.013)
 15. de Brito, J.A.G., Totte, D.R.M., Silva, F.O., de P. Junior, J.R., Henriques, F.d.R., Tarrataca, L., Haddad, D.B., de Assis, L.S. (2022) Memetic Algorithm Applied to Topology Control Optimization of a Wireless Sensor Network. *Wireless Networks*, 28(8): 3677-3697. DOI: [10.1007/s11276-022-03068-9](https://doi.org/10.1007/s11276-022-03068-9)
 16. de Brito, J.A.G., de P. Junior, J.R., Henriques, F.d.R., de Assis, L.S. (2019) Topology Control Optimization of Wireless Sensor Networks for IoT Applications. *Proceedings of the 25th Brazilian Symposium on Multimedia and the Web (WebMedia)*, 477-480. DOI: [10.1145/3323503.3361718](https://doi.org/10.1145/3323503.3361718)
 17. Zhao, H., Guo, S., Wang, X., Wang, F. (2015) Energy-Efficient Topology Control Algorithm for Maximizing Network Lifetime in Wireless Sensor Networks with Mobile Sink. *Applied Soft Computing*, 34: 539-550. DOI: [10.1016/j.asoc.2015.05.014](https://doi.org/10.1016/j.asoc.2015.05.014)
 18. Pyrih, Y., Przystupa, K., Pyrih, Y., Sikora, J., Beshley, M. (2026) Energy-Efficient Topology Optimization of Wireless Sensor Networks Using a Modified Genetic Algorithm. *Electronics*, 15(10): 2016. DOI: [10.3390/electronics15102016](https://doi.org/10.3390/electronics15102016)
 19. Priyadarshi, R. (2025) Efficient Node Deployment for Enhancing Coverage and Connectivity in Wireless Sensor Networks. *Scientific Reports*, 15. DOI: [10.1038/s41598-025-14252-0](https://doi.org/10.1038/s41598-025-14252-0)
 20. Rathee, M., Kumar, S., Dilip, K., Dohare, U., Aanchal, Parveen (2024) Towards Energy Balancing Optimization in Wireless Sensor Networks: A Novel Quantum Inspired Genetic Algorithm Based Sinks Deployment Approach. *Ad Hoc Networks*, 153: 103350. DOI: [10.1016/j.adhoc.2023.103350](https://doi.org/10.1016/j.adhoc.2023.103350)
 21. Khan, M.A., Javaid, N. (2022) Computationally Efficient Topology Optimization of Scale-Free IoT Networks. *Computer Communications*, 185: 1-12. DOI: [10.1016/j.comcom.2021.12.013](https://doi.org/10.1016/j.comcom.2021.12.013)
 22. Changazi, S.A., Bakhshi, A.D., Yousaf, M., Mohsin, S.M., Akber, S.M.A., Abazeed, M., Ali, M. (2024) Optimization of Network Topology Robustness in IoTs: A Systematic Review. *Computer Networks*, 250: 110568. DOI: [10.1016/j.comnet.2024.110568](https://doi.org/10.1016/j.comnet.2024.110568)
 23. Peng, Y., Liu, C., Liu, S., Liu, Y., Wu, Y. (2022) SmartTRO: Optimizing Topology Robustness for Internet of Things via Deep Reinforcement Learning with Graph Convolutional Networks. *Computer Networks*, 218: 109385. DOI: [10.1016/j.comnet.2022.109385](https://doi.org/10.1016/j.comnet.2022.109385)
 24. Chen, N., Qiu, T., Daneshmand, M., Wu, D.O. (2021) Robust Networking: Dynamic Topology Evolution Learning for Internet of Things. *ACM Transactions on Sensor Networks*, 17(3): article 26. DOI: [10.1145/3446937](https://doi.org/10.1145/3446937)
 25. Abujassar, R.S. (2025) Intelligent IoT-Driven Optimization of Large-Scale Healthcare Networks: The INRWLF Algorithm for Adaptive Efficiency. *Discover Computing*, 28(1): 93. DOI: [10.1007/s10791-025-09601-6](https://doi.org/10.1007/s10791-025-09601-6)

26. Wang, F., Su, K., Liang, B., Yao, J., Bai, W. (2024) Research on Multi-Layer Network Topology Optimization Strategy for Railway Internet of Things Based on Game Theory Benefits. *Frontiers in Physics*, 12: 1409427. DOI: [10.3389/fphy.2024.1409427](https://doi.org/10.3389/fphy.2024.1409427)
27. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T. (2002) A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2): 182-197. DOI: [10.1109/4235.996017](https://doi.org/10.1109/4235.996017)

10. Declarations

Funding

This literature synthesis was conducted without a specific grant from any public, commercial, or not-for-profit funding body.

Conflict of Interest

The authors have no competing interests to declare.

Data Availability Statement

This article synthesizes previously published literature and does not report new experimental data; readers should consult the cited sources for their respective data and code availability statements.

Ethical Approval

Not applicable, as this work is a literature synthesis involving no human participants, animal subjects, or primary data collection.