




RADemics

Edge Computing and Fog Computing Paradigms for Ultra Low Latency IoT Based Smart City Applications

R Thiyagarajan, M Arivamudhan
Government College of Engineering,

Edge Computing and Fog Computing Paradigms for Ultra Low Latency IoT Based Smart City Applications

¹R Thiagarajan, Associate Professor, Dept of ECE, Government College of Engineering, Dharmapuri. thiyagucdm@gmail.com

²M Arivamudhan, Professor, Dept of ECE, Government College of Engineering, Dharmapuri.

aumaei@gmail.com

Abstract

Edge and fog computing paradigms have emerged as pivotal enablers for meeting the ultra-low-latency requirements of Internet-of-Things (IoT) applications in contemporary smart city environments. This chapter presents a systematic examination of hierarchical and hybrid orchestration architectures, highlighting the roles of edge nodes and fog clusters in distributing computational workloads to minimize end-to-end response times. Detailed analysis of distributed consensus mechanisms and adaptive microservice partitioning underscores the importance of fault tolerance and dynamic resource allocation in maintaining sub-10 millisecond latencies under variable network conditions. Standardized benchmarking methodologies are introduced to evaluate control-plane overhead, while empirical case studies spanning real-time traffic signal control and public-safety video analytics demonstrate the performance gains achievable through strategic edge-fog deployments. Critical challenges, including scalability, security, and interoperability, are delineated, and future research directions are proposed to guide the evolution of resilient, high-performance IoT infrastructures.

Keywords: Edge Computing; Fog Computing; Ultra-Low Latency; IoT Orchestration; Smart Cities; Fault Tolerance

Introduction

Edge and fog computing paradigms have transformed the landscape of smart city infrastructures by pushing computation closer to data sources and providing intermediate processing layers between edge devices and centralized clouds [1]. The exponential growth of Internet-of-Things (IoT) deployments ranging from autonomous traffic systems and environmental sensors to public-safety surveillance has generated a deluge of real-time data requiring immediate processing [2]. Cloud-centric models, while offering scalable resources, often fail to meet the stringent latency demands imposed by critical applications such as collision-avoidance in intelligent transportation systems or emergency-alert analytics [3]. Edge computing addresses this gap by relocating computational tasks to the network periphery, thereby reducing round-trip delays and conserving bandwidth [4]. Edge nodes alone can be constrained by limited processing power, heterogeneous capabilities, and lack of coordination at scale. Fog computing emerges as a complementary

paradigm, introducing a hierarchical layer of fog nodes that aggregate, filter, and orchestrate workloads across distributed edge devices and the cloud. Together, these paradigms form an edge–fog–cloud continuum capable of delivering sub-10 millisecond response times, ensuring reliability, and enabling context-aware services in ultra-low-latency IoT applications [5].

A critical aspect of this continuum involves architectural designs that support dynamic workload distribution, resource management, and fault tolerance across geographically dispersed nodes [6]. Microservice-based architectures, containerization technologies, and software-defined networking (SDN) offer the flexibility and programmability needed to implement fine-grained task offloading and service migration [7]. By decomposing monolithic applications into smaller, independently deployable services, microservice partitioning enables precise placement of functionalities such as data aggregation, analytics, or machine-learning inference at optimal network layers [8]. SDN and network slicing further enhance end-to-end quality-of-service (QoS) by prioritizing latency-sensitive traffic and allocating dedicated network resources [9]. Meanwhile, container orchestration frameworks such as Kubernetes, extended to support edge and fog environments, automate deployment, scaling, and healing of services in response to fluctuating workloads and node statuses. These advancements collectively underpin the agility required to meet the evolving demands of smart city IoT applications [10].

Ensuring consistent system state and service continuity in the face of failures or network partitions constitutes another foundational challenge [11]. Distributed consensus mechanisms ranging from Paxos and Raft to Byzantine Fault Tolerant (BFT) protocols facilitate agreement on critical configuration and state transitions across clusters of edge and fog nodes [12]. Conventional consensus algorithms often incur prohibitive communication overhead in large-scale deployments; thus, research is focused on lightweight and adaptive variants that minimize control-plane latency while preserving fault tolerance [13]. Coupling consensus-based state replication with self-healing strategies enables automated detection of node failures, dynamic reallocation of services, and rapid restoration of system coherence [14]. This synergy is vital for maintaining ultra-low-latency guarantees in mission-critical applications such as real-time traffic signal control or public-safety video analytics [15].