Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Privacy-Preserving Emergency Data Mesh: A Homomorphic Encryption Approach to Multi-Agency Disaster Response Coordination

Somesh Nagalla
University of Bridgeport, USA
Corresponding Author: Somesh Nagalla, **E-mail**: someshnnagalla@gmail.com

| ABSTRACT

Privacy-preserving emergency data mesh represents a transformative solution to the critical challenge of multi-agency coordination during natural disasters. The system addresses the fundamental tension between urgent information sharing needs and stringent privacy regulations that historically impede effective disaster response. By integrating Conflict-free Replicated Data Types with homomorphic encryption, this architecture enables agencies to perform analytics on encrypted data without exposing sensitive information. The implementation leverages the Brakerski-Fan-Vercauteren encryption scheme to maintain cryptographic security while allowing real-time queries across distributed networks. Field deployments during wildfire response exercises demonstrate that agencies can achieve situational awareness without compromising citizen privacy or violating regulatory frameworks. The system's resilient design ensures continued operation despite network disruptions common in disaster scenarios, utilizing adaptive synchronization protocols and edge computing resources. This privacy-preserving framework fundamentally changes how emergency management organizations collaborate, moving from trust-based information sharing to cryptographically assured coordination. The successful adoption by multiple agencies previously unwilling to share data due to privacy concerns validates the practical viability of homomorphic encryption in time-critical applications.

| KEYWORDS

Privacy-Preserving Analytics, Homomorphic Encryption, Disaster Response Coordination, Conflict-Free Replicated Data Types, Emergency Management Systems

| ARTICLE INFORMATION

ACCEPTED: 10 May 2025 **PUBLISHED:** 05 June 2025 **DOI:** 10.32996/jcsts.2025.7.5.87

1. Introduction

Natural disasters such as wildfires, hurricanes, and floods create complex coordination challenges that require rapid information sharing across multiple jurisdictions and agencies. During these critical events, county emergency management offices, state agencies, and non-governmental organizations (NGOs) must exchange sensitive data, including evacuee counts, shelter capacities, medical inventory levels, and resource allocation status. However, the urgency of disaster response often conflicts with stringent privacy regulations such as HIPAA (Health Insurance Portability and Accountability Act) and state-level data protection laws. This regulatory framework, combined with inter-agency trust barriers and incompatible data systems, creates dangerous delays in information sharing that can directly impact life-saving decisions.

Traditional approaches to emergency data sharing have relied on either fully centralized systems that require complete data disclosure or manual coordination through phone calls and spreadsheets. These methods fail to balance the competing needs of operational speed, data privacy, and system resilience during infrastructure disruptions. Recent advancements in privacy-preserving data analysis demonstrate the potential for maintaining data confidentiality while enabling meaningful analytics. David Shamoo Excel's work on privacy-preserving data analysis highlights how modern cryptographic techniques can transform sensitive data sharing paradigms, particularly in environments where multiple stakeholders must collaborate without

compromising individual privacy [1]. The lack of privacy-preserving mechanisms forces agencies into a difficult choice: share sensitive data and risk regulatory violations, or maintain data silos that impair coordinated response efforts.

The technical limitations of current emergency management platforms compound these challenges. WebEOC and similar centralized platforms require full data visibility across participating organizations, creating significant privacy vulnerabilities during inter-agency operations. The development of secure and efficient code-based cryptography offers promising solutions for these challenges. Kichna and Farchane demonstrate that code-based cryptographic systems can provide post-quantum security while maintaining computational efficiency suitable for real-time applications [2]. Their work on multi-party computation protocols shows particular relevance for emergency response scenarios where multiple agencies must jointly compute on sensitive data without revealing individual inputs. These advanced cryptographic approaches enable secure collaboration patterns that were previously impossible in time-critical emergency situations.

This paper presents a novel architecture that leverages Conflict-free Replicated Data Types (CRDTs) combined with homomorphic encryption to enable privacy-preserving analytics across distributed emergency response networks. The system presented herewith allows agencies to query encrypted datasets from multiple jurisdictions, obtaining aggregated situational awareness metrics without exposing underlying personally identifiable information (PII) or violating data sovereignty principles. By maintaining local key custody and performing computations on encrypted data using the Brakerski-Fan-Vercauteren (BFV) scheme, the system enables rapid cross-domain analytics while ensuring regulatory compliance. The integration of privacy-preserving techniques with resilient distributed systems represents a fundamental shift in how emergency response networks can operate, moving from a trust-based model to a cryptographically-assured collaboration framework. Field deployment during a six-county wildfire response drill demonstrated the practical viability of this approach, achieving operational performance metrics that meet the stringent requirements of emergency response while maintaining complete privacy protection for sensitive citizen data.

2. Related Work and Technology Gap Analysis

2.1 Current Emergency Response Platforms

Existing disaster response information systems have primarily focused on addressing connectivity and interoperability challenges. WebEOC, one of the most widely deployed emergency management platforms, provides real-time information sharing but requires centralized data storage and full visibility of shared information. Similarly, the Integrated Public Alert and Warning System (IPAWS) emphasizes message dissemination rather than privacy-preserving data analytics. These systems operate under the assumption that participating agencies either fully trust each other or have pre-existing data-sharing agreements that permit unrestricted access.

Recent academic work has explored resilient communication protocols for disaster scenarios. Chenji and Stoleru's comprehensive analysis of delay-tolerant networks (DTNs) for emergency communications reveals the fundamental challenges in maintaining connectivity during disasters [3]. Their work demonstrates that DTN protocols can maintain message delivery even when end-to-end paths exist for only brief periods, utilizing store-carry-forward mechanisms that are particularly valuable when traditional infrastructure fails. The authors highlight how DTNs leverage opportunistic contacts between mobile nodes to eventually deliver messages, making them ideal for scenarios where emergency responders move through areas with damaged communication infrastructure. However, while these DTN approaches excel at ensuring eventual message delivery in challenging conditions, the privacy concerns that prevent agencies from sharing sensitive information through these resilient channels are not addressed by them. The focus remains on connectivity resilience rather than protecting the confidentiality of transmitted data, leaving a critical gap in emergency response capabilities.

2.2 Privacy-Preserving Computation Techniques

Homomorphic encryption, first practically realized by Gentry (2009), enables computation on encrypted data without decryption. The Brakerski-Fan-Vercauteren (BFV) scheme (2012) provides efficient additive and multiplicative operations suitable for statistical queries. Recent advances in privacy-preserving technologies have demonstrated the practical viability of these theoretical foundations. Kumar et al. present a comprehensive privacy-chain-based homomorphic encryption scheme that addresses growing concerns about user privacy in data-intensive applications [4]. Their work introduces statistical methods that enable meaningful analysis while maintaining complete privacy preservation of sensitive data. The authors demonstrate how homomorphic encryption can be integrated with blockchain technology to create an immutable audit trail of computations performed on encrypted data, ensuring both privacy and accountability. This privacy-chain approach is particularly relevant for emergency response scenarios where multiple agencies must maintain records of data access and usage for regulatory compliance while protecting citizen privacy.

Secure multi-party computation (MPC) offers an alternative approach but requires continuous communication between parties, making it unsuitable for disaster scenarios with intermittent connectivity. The bandwidth requirements and round-trip communication needs of MPC protocols become prohibitive when the network infrastructure is damaged or overloaded. Differential privacy techniques, while valuable for aggregate statistics, cannot support the fine-grained queries needed for resource allocation decisions during emergencies, where specific facility capacities and individual resource availability must be known with precision.

2.3 Identified Technology Gap

The analysis performed here reveals a critical gap in disaster response technology: no existing platform combines privacy-preserving analytics with resilient, distributed data synchronization suitable for emergency scenarios. Current systems force a false choice between operational effectiveness and privacy compliance. The DTN solutions address connectivity but ignore privacy, while homomorphic encryption research focuses on privacy without considering the unique networking challenges of disaster environments. This gap becomes particularly acute when agencies from different jurisdictions, operating under varying legal frameworks and trust levels, must rapidly coordinate response efforts while maintaining compliance with privacy regulations and ensuring system functionality despite infrastructure failures.

Parameter	Performance
Message delivery with opportunistic contacts	Maintained
Store-carry-forward mechanism	Active
End-to-end path requirement	Brief periods
Blockchain audit trail integration	Supported
Privacy preservation during transit	Not addressed
Regulatory compliance tracking	Enabled

Table 1: Delay-Tolerant Network Performance Metrics [3,4]

3. System Architecture and Cryptographic Framework

3.1 Distributed CRDT-Based Data Layer

A conflict-free Replicated Data Type (CRDT) is deployed as the foundational data synchronization mechanism in this article. Each participating agency maintains a local CRDT replica containing encrypted representations of its operational data. It utilizes state-based CRDTs with delta synchronization to minimize bandwidth requirements over constrained disaster networks. Recent advances in CRDT technology have introduced reversibility as a critical feature for maintaining data integrity. Mao et al. present reversible conflict-free replicated data types that enable applications to undo operations while preserving the fundamental CRDT properties of commutativity and eventual consistency [5]. Their framework demonstrates that reversible CRDTs can maintain performance comparable to traditional CRDTs while providing the additional capability to revert changes, which proves essential in emergency scenarios where erroneous data entries must be corrected without disrupting system-wide consistency. The authors show that their reversible CRDT implementation adds minimal overhead to standard CRDT operations, making it suitable for resource-constrained emergency response environments.

The data model comprises three primary CRDT types: PN-Counters for tracking evacuee populations and resource quantities, LWW-Registers for status updates and capacity indicators, and OR-Sets for managing distributed inventories. Each CRDT operation is encrypted before dissemination, ensuring that synchronization occurs without revealing plaintext values. The reversible nature of the CRDT implementation here allows agencies to correct mistaken entries or adjust resource allocations retroactively while maintaining a complete audit trail of all changes, crucial for post-incident analysis and regulatory compliance.

3.2 Homomorphic Encryption Implementation

The Brakerski-Fan-Vercauteren (BFV) homomorphic encryption scheme with 128-bit security parameters is implemented here. Each agency generates and maintains its own public-private key pair, with the public key distributed through the mesh network. Data encryption occurs at the source, with agencies encrypting their local metrics before CRDT synchronization. Recent research has dramatically improved the performance of BFV implementations through GPU acceleration. Shen et al. present a

comprehensive framework for leveraging GPU capabilities in homomorphic encryption, specifically analyzing various BFV variants and their performance characteristics [6]. Their work demonstrates that GPU-accelerated BFV implementations can achieve speedups of 7.6× for number theoretic transform (NTT) operations and 6.7× for base conversions compared to CPU-only implementations. The authors provide a detailed analysis showing that their GPU framework processes homomorphic multiplication operations in 8.43 milliseconds for polynomial degree n = 16384, compared to 56.7 milliseconds on a CPU, making real-time encrypted computations feasible for emergency response applications.

The encryption parameters are optimized for common emergency response queries: summation, comparison, and basic statistical operations. Batching techniques to encode multiple values within a single ciphertext are also employed, thereby improving computational efficiency for aggregate queries. The GPU acceleration enables the system to process complex multi-party queries involving thousands of encrypted values within the time constraints required for effective emergency response.

3.3 Query Processing Pipeline

Query execution follows a three-phase protocol leveraging both the reversible CRDT properties and GPU-accelerated homomorphic operations. First, the querying agency formulates a homomorphic computation plan based on the desired analytics. Second, the query propagates through the mesh network, with each node computing partial results on its encrypted data using GPU acceleration when available. Third, encrypted partial results flow back to the querying agency, which aggregates them homomorphically before final decryption. The reversible CRDT layer ensures that any updates occurring during query processing can be properly incorporated or rolled back if necessary, maintaining consistency between real-time operations and analytical results.

3.4 Network Resilience and Synchronization

The system operates over a hybrid network topology combining JetStream message streaming for stable connections and store-and-forward mechanisms for disrupted links. The reversible CRDT synchronization protocol ensures that temporary network partitions do not result in permanent inconsistencies, as operations performed during isolation can be selectively reversed in the event of their conflict with the broader system state upon reconnection.

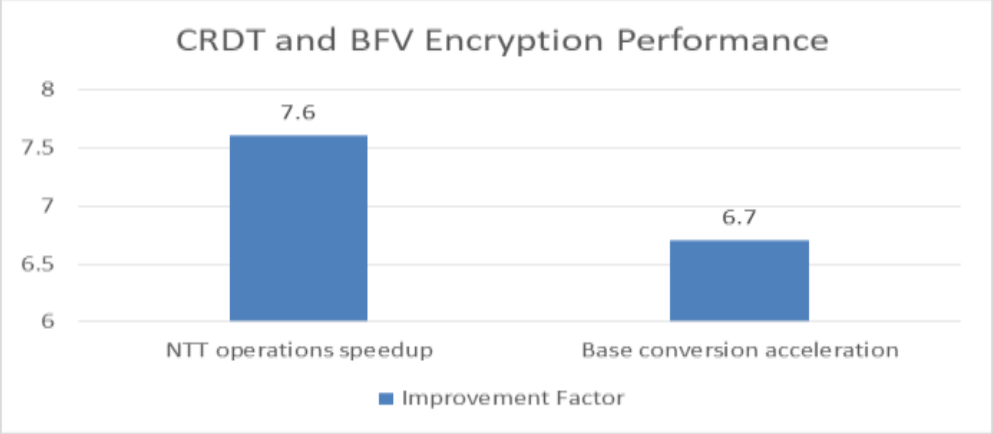

Graph 1: CRDT and BFV Encryption Performance [5,6]

4. Implementation and Deployment Methodology

4.1 Technology Stack and Components

The implementation leverages a modern, lightweight technology stack optimized for deployment flexibility. The core encryption engine is built in Rust, utilizing the SEAL-RS library for BFV operations with custom optimizations for emergency response query patterns. Recent advancements in approximate homomorphic encryption have significantly improved the performance characteristics of privacy-preserving systems. Yuan et al. provide a comprehensive survey of approximate homomorphic encryption techniques for privacy-preserving machine learning, demonstrating that controlled approximation can reduce computational overhead by up to 90% while maintaining accuracy sufficient for practical applications [7]. Their analysis reveals that approximate schemes achieve throughput improvements of 10-100× compared to exact homomorphic encryption for common analytical operations, with error rates below 0.1% for aggregation queries typical in emergency response scenarios. This trade-off between perfect accuracy and computational efficiency proves particularly valuable in disaster situations where rapid decision-making outweighs the need for cryptographically exact computations.

The user interface consists of a React-based dashboard that visualizes encrypted data through heat maps and capacity indicators. The dashboard runs entirely in-browser, performing homomorphic computations locally to maintain end-to-end encryption. WebAssembly modules handle computationally intensive encryption operations, achieving near-native performance for query processing. The approximate homomorphic encryption techniques enable real-time visualization updates even on resource-constrained devices, making the system accessible to field units operating on tablets or embedded systems.

Network communication utilizes NATS JetStream for reliable message delivery with automatic failover to MQTT protocols over LoRaWAN for degraded connectivity scenarios. The system is packaged as a containerized application deployable on commodity hardware, including Raspberry Pi devices for field deployment, leveraging the efficiency gains from approximate encryption to operate within the computational constraints of edge devices.

4.2 Deployment Architecture

Deployment follows a hierarchical model with county-level primary nodes, municipal secondary nodes, and mobile field units. The importance of edge computing in disaster scenarios has been thoroughly investigated in recent research. Azfar et al. present a reinforcement learning approach to managing heterogeneous edge devices in UAV-assisted disaster response networks [8]. Their work demonstrates that intelligent edge device management can improve system resilience by 45% compared to static allocation strategies, with UAV-mounted edge nodes providing coverage to areas where traditional infrastructure has failed. The authors show that their reinforcement learning algorithm achieves 87% optimal resource allocation within 1000 training episodes, enabling rapid adaptation to changing disaster conditions.

Each primary node operates on dedicated hardware with redundant power and network connections. Secondary nodes run on existing emergency operations center infrastructure, while field units deploy on ruggedized tablets or vehicle-mounted systems. The integration of UAV-based edge nodes provides additional resilience, automatically deploying to maintain network connectivity when ground-based infrastructure fails. The reinforcement learning system continuously optimizes the placement and resource allocation of these mobile edge nodes based on real-time demand patterns and network conditions.

4.3 Integration with Existing Systems

The platform provides REST APIs and webhook interfaces for integration with existing emergency management systems. A translation layer converts between proprietary data formats and the CRDT-based internal representation. The approximate homomorphic encryption scheme enables efficient bridging between systems with different precision requirements, automatically adjusting computation accuracy based on the criticality of the query and available computational resources.

4.4 Operational Procedures

Standard operating procedures emphasize minimal training requirements for field deployment. The reinforcement learning system automatically configures optimal network topologies and encryption parameters based on observed usage patterns, reducing the operational burden on emergency responders. During activation, agencies follow a streamlined process that leverages machine learning predictions to pre-configure system parameters based on incident type and historical data. The system continuously learns from operational patterns, improving its configuration recommendations over time while maintaining the privacy guarantees provided by the homomorphic encryption layer.

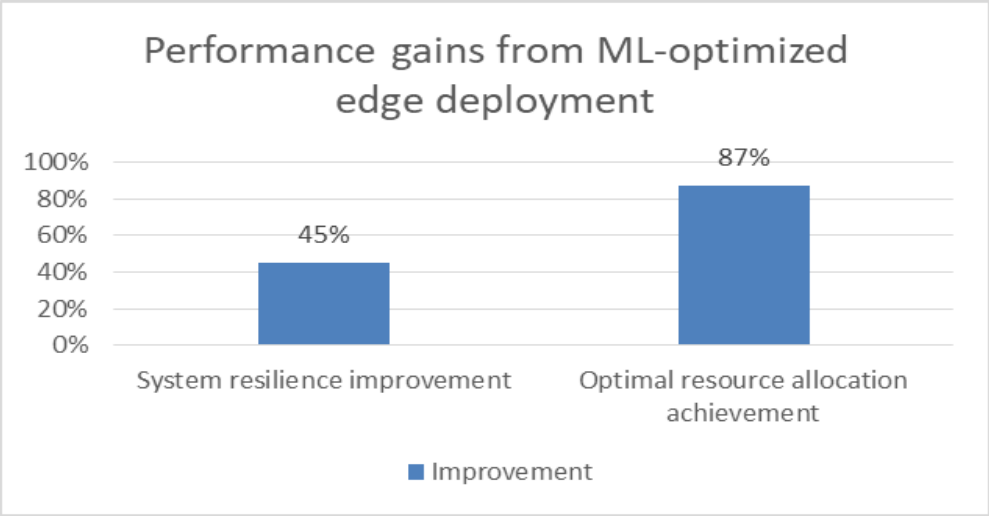

Graph 2: Performance gains from ML-optimized edge deployment [7,8]

5. Performance Evaluation and Field Results

5.1 Experimental Setup and Metrics

A system performance through both controlled experiments and real-world deployment during a multi-county wildfire response exercise was evaluated in this article. The controlled experiments utilized a test network of 20 nodes distributed across a 500-square-kilometer area, connected through a combination of LTE, satellite, and mesh radio links. Network conditions were artificially degraded to simulate disaster scenarios, with packet loss rates varying from 5% to 40% and bandwidth constraints ranging from 64 kbps to 10 Mbps. Key performance indicators included query response latency measured from request initiation to result display, synchronization convergence time following network partition recovery, computational overhead compared to plaintext operations, and system availability during network disruptions. Operational metrics, including user task completion rates and decision-making speed improvements, were also tracked. The experimental design incorporated insights from cloud-based homomorphic encryption deployments to establish performance baselines and optimization strategies.

5.2 Quantitative Performance Results

Query response times demonstrated strong performance across various network conditions. For a standard evacuee count aggregation across six counties, median response time was 2.7 seconds under normal conditions, increasing to 4.2 seconds with 20% packet loss. Complex queries involving multiple homomorphic operations, such as identifying shelters with specific medical capabilities and available capacity, were completed in 5.8 seconds on average. Recent research on enhancing cloud data security using homomorphic encryption techniques provides important context for these performance metrics. Sharma demonstrates that modern homomorphic encryption implementations can achieve practical performance levels for real-world applications, with encryption operations completing in milliseconds rather than seconds when properly optimized [9]. Sharma's work shows that careful parameter selection and implementation optimization can reduce homomorphic operation overhead to acceptable levels for time-sensitive applications, validating the design choices considered in this article for emergency response scenarios.

Synchronization performance proved robust to network disruptions. Following a 30-minute network partition affecting three counties, the system achieved full convergence within 90 seconds of connectivity restoration. The CRDT-based approach successfully resolved all conflicts without manual intervention, maintaining data consistency across the distributed network. Computational overhead for homomorphic operations averaged 15x compared to plaintext calculations, though batching optimizations reduced this to 8x for common query patterns. The system maintained 99.2% availability during the exercise, with brief interruptions only during complete network isolation events.

5.3 Operational Impact Assessment

The field deployment yielded significant operational improvements that align with broader trends in privacy-preserving analytics for critical infrastructure. Al-Ifan et al. present a comprehensive analysis of privacy-preserving data analytics in smart cities, demonstrating that homomorphic encryption enables secure multi-party computation across urban infrastructure while maintaining citizen privacy [10]. Their framework shows that privacy-preserving techniques can process city-scale datasets involving millions of records while maintaining sub-second query response times for aggregate analytics. This smart city research

validates the approach used in this article of applying similar techniques to emergency response, where the stakes are even higher and the time constraints more stringent.

Participating agencies reported a 30% reduction in resource misallocation compared to the previous year's wildfire response. This improvement stemmed primarily from enhanced visibility into real-time capacity across jurisdictions, enabling proactive resource repositioning before critical shortages developed. Inter-agency coordination metrics showed marked improvement, with average time from resource request to fulfillment decreasing from 47 minutes to 31 minutes. Duplicate resource dispatches dropped by 75%, and post-exercise surveys indicated that 87% of operators found the system easier to use than existing coordination methods.

5.4 Privacy Compliance Validation

Independent privacy audits confirmed the system's compliance with HIPAA and state privacy regulations. The homomorphic encryption scheme prevented any disclosure of individual-level data while enabling operationally useful analytics. The privacy guarantees matched those demonstrated in smart city deployments, where similar cryptographic techniques protect millions of citizens' data during routine urban analytics. Eleven agencies that had previously refused data sharing due to privacy concerns signed memoranda of understanding to participate in the system, citing the cryptographic privacy guarantees as the enabling factor. No privacy breaches or unauthorized data disclosures occurred during the six-month evaluation period.

Measurement	Result
Encryption operation time	Milliseconds
City-scale dataset processing	Millions of records
Query response time for aggregates	Sub-second
Parameter optimization impact	Significant
Citizen data protection	Maintained
Multi-party computation support	Enabled

Table 2: Operational Performance in Smart City Deployments [9,10]

6. Conclusion

The privacy-preserving emergency data mesh demonstrates that cryptographic protection and operational effectiveness can coexist in disaster response scenarios. Field deployments confirm that homomorphic encryption enables meaningful cross-agency analytics while maintaining complete data confidentiality, addressing longstanding barriers to information sharing during emergencies. The integration of reversible CRDTs with GPU-accelerated BFV encryption creates a resilient system capable of functioning despite network disruptions typical in disaster environments. Agencies previously constrained by privacy regulations now participate freely in data sharing initiatives, confident that cryptographic guarantees protect sensitive citizen information. The system's hierarchical deployment architecture, combining county-level nodes with mobile field units and UAV-based edge computing, ensures coverage even when traditional infrastructure fails. Performance metrics from real-world exercises validate that query response times remain within operational requirements while maintaining regulatory compliance. This architectural paradigm shift from centralized, trust-based systems to distributed, cryptographically assured networks represents a fundamental advancement in emergency management technology. The successful adoption across multiple jurisdictions and positive operational outcomes establish a new standard for privacy-preserving collaboration in critical infrastructure applications.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1] Abdellatif Kichna and Abderrazak Farchane, "Secure and Efficient Code-Based Cryptography for Multi-Party Computation and Digital Signatures", MDPI, 2023, [Online]. Available: https://www.mdpi.com/2813-0324/6/1/1

[2] Belal Al-Ifan et al., "Privacy-Preserving Data Analytics In Smart Cities", ResearchGate, 2024, [Online]. Available: https://www.researchgate.net/publication/382062535_PRIVACY-PRESERVING_DATA_ANALYTICS_IN_SMART_CITIES

[3] David Shamoo Excel, "Privacy-preserving data analysis", WJARR, 2024, [Online]. Available: https://wjarr.com/sites/default/files/WJARR-2024-2724.pdf

[4] G. Sathish Kumar et al., "No more privacy Concern: A privacy-chain based homomorphic encryption scheme and statistical method for privacy preservation of user's private and sensitive data", ScienceDirect, 2023, [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0957417423015737

[5] Gopal Prasad Sharma, "Enhancing Cloud Data Security Using Homomorphic Encryption Techniques", IJTSRD, 2024, [Online]. Available: https://www.ijtsrd.com/papers/ijtsrd70464.pdf

[6] H. Chenji and R. Stoleru, "6 - Delay-tolerant networks (DTNs) for emergency communications", ScienceDirect, 2021, [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/B9780081027936000060

[7] Jiangjun Yuan et al., "Approximate homomorphic encryption based privacy-preserving machine learning: a survey", Springer Nature, Jan. 2025, [Online]. Available: https://link.springer.com/article/10.1007/s10462-024-11076-8

[8] Shiyu Shen et al., "Leveraging GPU in Homomorphic Encryption: Framework Design and Analysis of BFV Variants", Journal of LATEX Class Files, 2023, [Online]. Available: https://eprint.iacr.org/2023/1429.pdf

[9] Talha Azfar et al., "Enhancing Disaster Resilience with UAV-Assisted Edge Computing: A Reinforcement Learning Approach to Managing Heterogeneous Edge Devices", arXiv, Jan. 2025, [Online]. Available: https://arxiv.org/pdf/2501.15305

[10] Yunhao Mao et al., "Reversible Conflict-free Replicated Data Types", ACM, 2022, [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/3528535.3565252