

Resilient Architectures for Fault-Tolerant Distributed and Embedded Systems: Theory, Methods, and Practical Pathways

Dr. Emilia Carter

University of Edinburgh

Abstract

This article synthesizes foundational theory, contemporary methodologies, and applied strategies for designing fault-tolerant, dependable, and resilient computing systems across distributed, embedded, and cloud environments. It integrates classical dependability concepts with adaptive and mixed-criticality approaches, examines fault tolerance for high-reliability industrial and avionics systems, and extends discussion to modern cloud and GPU manufacturing test infrastructures. The structured abstract outlines the problem context, methodological approach, principal findings, and implications for future research and practice.

Background: Dependability and resilience are long-standing goals in computing that require coherent theoretical framing and practical engineering trade-offs to manage faults, errors, and failures in a wide spectrum of platforms from real-time embedded controllers to large-scale cloud and GPU test infrastructures (Avizienis et al., 2001; Laprie, 2008). The proliferation of adaptive systems and multi-cloud architectures has made fault tolerance both more necessary and more complex, demanding new strategies that combine classical redundancy with adaptive reconfiguration, machine learning-assisted detection, and service-level design (Kim & Lawrence, 1992; Årzén, 2013; Tebaa & ElHajji, 2014).

Methods: This work conducts a conceptual-methodological synthesis grounded in dependability theory and applied studies. It uses rigorous conceptual analysis, comparative evaluation of published strategies, and descriptive modeling to identify common design patterns, trade-offs, and evaluation criteria. Core methods discussed include mode-change management in mixed-criticality systems, adaptive fault-tolerance control loops, diversity and replication strategies, fault-tolerant Ethernet architectures, reactive fault tolerance in cloud platforms, priority-based cooperative scheduling and task backfilling, and machine-learning-assisted detection and recovery mechanisms (Burns, 2014; Kim & Lawrence, 1992; Álvarez et al., 2019; Abd Elfattah et al., 2017).

Results: Integrating principles from fault-tolerance research reveals recurring design motifs: i) separation of fault detection and recovery concerns to enable verifiable modes of operation (Avizienis et al., 2001; Burns, 2014); ii) use of adaptive, context-aware reconfiguration to maintain timeliness and safety under resource degradation (Kim & Lawrence, 1992; Årzén, 2013); iii) leveraging reliable networking and protocol-level strategies for industrial systems to attain stringent availability goals (Álvarez et al., 2019); and iv) combining proactive testing, online monitoring, and machine-learning classifiers to enhance cloud and manufacturing testbed robustness (Kochhar & Hilda, 2017; Designing Fault-Tolerant Test Infrastructure, 2025). The descriptive analysis highlights how trade-offs manifest across latency, resource overhead, complexity, and verification effort.

Conclusions: A unifying perspective emerges where dependability is achieved by layered defenses: preventive design, fault containment, detection, diagnosis, and recovery, orchestrated by adaptive control policies. For future systems, the convergence of formal mode-change reasoning, adaptive middleware, data-driven diagnostics, and cross-layer orchestration is essential. Research priorities include formalizing adaptive policies for mixed-criticality contexts, validating ML-based fault classifiers under adversarial conditions, and developing standardized evaluation frameworks for cloud and manufacturing fault-tolerance strategies.

Keywords: Dependability, Fault Tolerance, Adaptive Systems, Mixed Criticality, Cloud Resilience, Fault

Diagnosis

INTRODUCTION

The pursuit of dependable computing systems has been a central, enduring objective of computer science and engineering. Dependability encompasses a set of attributes—reliability, availability, safety, integrity, and maintainability—that collectively determine whether a system can be trusted to perform its intended function under expected and unexpected conditions (Avizienis et al., 2001). Historically, the field evolved from early results in fault masking and redundancy for avionics and industrial control to contemporary concerns arising from distributed cloud platforms, GPU manufacturing testbeds, and heterogeneous embedded systems. The problem now is not simply making a single processor or device fault tolerant; it is designing architectures and operational policies that maintain mission goals in environments of increasing complexity, scale, and unpredictability (Laprie, 2008; Wensley et al., 2008).

This article addresses the following central problem: how can designers synthesize classical dependability principles with adaptive and data-driven approaches to create resilient architectures for modern distributed, embedded, and cloud-hosted systems? This question is motivated by three trends. First, systems are mixing workloads of different criticality and timing requirements on shared platforms, forcing preemption and mode-change decisions that must be safe and predictable (Burns, 2014). Second, cloud and multi-cloud platforms present unique failure modes—partial outages, correlated failures, and degraded performance regimes—that challenge static fault models and call for reactive, scalable fault-tolerance strategies (Abid et al., 2014; Tebaa & ElHajji, 2014). Third, high-throughput manufacturing and test facilities for accelerators like GPUs impose demanding constraints on testability and fault handling; test infrastructure itself must be fault tolerant to avoid cascading production losses (Designing Fault-Tolerant Test Infrastructure, 2025).

Existing literature provides deep conceptual foundations and numerous concrete techniques. Avizienis et al. (2001) offered a taxonomy of faults, errors, failures, and methods—fault prevention, removal, tolerance, and forecasting—that remain vital to reasoning about modern systems. Laprie (2008) reframed dependability within resilience, emphasizing the capacity not only to avoid failures but to recover and continue operation under adverse conditions. Research on mixed-criticality systems and mode changes (Burns, 2014) has articulated the conditions and mechanisms required for safe reconfiguration, while embedded system research has demonstrated the effectiveness of adaptive fault-tolerance controllers for maintaining real-time guarantees under partial faults (Kim & Lawrence, 1992; Årzén, 2013). Industrial Ethernet and fault tolerance frameworks show how networking and system design can meet stringent reliability and timing needs (Álvarez et al., 2019). In cloud contexts, work on reactive fault tolerance, cooperative computing, and machine learning for fault detection proposes complementary strategies to address failures at scale (Abd Elfattah et al., 2017; Hasan & Goraya, 2016; Kochhar & Hilda, 2017).

Despite these advances, important gaps persist. First, a cohesive theory that unifies adaptive control policies, mode-change reasoning, and data-driven diagnostics across heterogeneous platforms is lacking. Studies often focus on specific layers—network, operating system, middleware, or application—without providing prescriptive guidance on how to compose these layers into a verifiable whole (Laprie, 2008). Second, machine learning techniques for fault detection promise improved sensitivity but raise questions about robustness, explainability, and integration into safety-critical control loops (Kochhar & Hilda, 2017). Third, evaluation frameworks are fragmented: metrics and benchmarks for comparing fault-tolerance strategies across cloud and manufacturing settings are inconsistent or absent (Designing Fault-Tolerant Test Infrastructure, 2025). This article attempts to bridge these gaps by offering a rigorous synthesis of principles, a taxonomy of mechanisms, and a set of design guidelines that make explicit the trade-offs designers must manage.

The remainder of the article unfolds as follows. The Methodology section explains the synthesis approach: conceptual analysis, mapping between theoretical constructs and practical mechanisms, and descriptive modeling for trade-off evaluation. The Results section presents the core findings: layered fault-tolerance patterns, adaptive reconfiguration strategies, networking considerations for industrial-grade availability, and the role of data-driven diagnostics. The Discussion interprets these findings within a broader research

agenda, highlights limitations, and sketches future directions. The Conclusion distills recommendations for researchers and practitioners aiming to build resilient systems across heterogeneous environments.

METHODOLOGY

This work adopts a synthesis methodology that combines theoretical analysis with descriptive, comparative evaluation of extant approaches drawn from the provided references. The methodology is not empirical in the sense of new experimental data collection; rather, it is analytical and integrative, oriented toward building a conceptual framework that captures design patterns and evaluation criteria for fault-tolerant systems across domains. The methodological steps are as follows.

Literature-driven conceptual mapping. Starting from foundational texts on dependability (Avizienis et al., 2001; Laprie, 2008), the research maps core dependability concepts—faults, errors, failures, and means of dependability—onto operational mechanisms described by more applied works. For example, mixed-criticality mode-change constructs (Burns, 2014) are mapped to adaptive controllers from embedded systems research (Kim & Lawrence, 1992; Årzén, 2013). Industrial Ethernet fault-tolerance mechanisms are extracted from domain-specific studies (Álvarez et al., 2019) and connected to general networking strategies for containment and recovery. Cloud tolerance strategies, including reactive protocols and multi-cloud privacy/fault-tolerance designs, are aligned with computing scheduling techniques and machine learning detection methods (Abd Elfattah et al., 2017; Hassan & Goraya, 2016; Kochhar & Hilda, 2017; Tebaa & ElHajji, 2014).

Taxonomy construction. The literature-driven mapping feeds a taxonomy of fault-tolerance mechanisms organized along multiple axes: prevention vs. tolerance vs. forecasting; static vs. adaptive; proactive vs. reactive; hardware vs. software vs. network; and centralised vs. decentralized control. This taxonomy is designed to help designers select and combine mechanisms appropriate to their operational context and constraints (Avizienis et al., 2001; Laprie, 2008).

Descriptive modeling of design trade-offs. For each class of mechanisms identified, the methodology uses descriptive models to characterize trade-offs in latency, resource overhead, verification complexity, and resilience. These models are textual and conceptual rather than quantitative—consistent with the constraint to avoid formulas—so they focus on explaining causal relationships, how decisions in one layer affect others, and what verification burdens arise from particular design choices (Burns, 2014; Kim & Lawrence, 1992).

Cross-domain integration. The methodology emphasizes bridging across domains—real-time embedded controllers, industrial networks, cloud platforms, and manufacturing testbeds—to identify reusable patterns and domain-specific variations. For example, mode-change policies from mixed-criticality scheduling inform cloud failover policies by reframing "criticality" as mission importance and service-level objectives in the cloud context (Burns, 2014; Abd Elfattah et al., 2017).

Evaluation criteria and desiderata. Finally, the methodology articulates a set of evaluation criteria—correctness under fault models, timeliness, resource efficiency, diagnosability, auditability, and manageability—and uses these to assess the mechanisms and compose design recommendations (Avizienis et al., 2001; Laprie, 2008).

Throughout, every substantive claim is linked to the provided literature so that readers can trace conceptual claims to concrete sources. The result is a comprehensive, cross-layer account of fault tolerance aimed at guiding both academic research and engineering practice.

RESULTS

The analysis yields several principal findings that together form an actionable conceptual foundation for resilient system design. The results are presented as thematic findings accompanied by nuanced explanation, cross-citations to foundational works, and discussion of practical implications.

1. Layered defenses remain the most effective organizing principle for dependability.

Designers should adopt a stratified approach that combines fault prevention, fault tolerance, fault removal, and fault forecasting, implemented at multiple layers: hardware, firmware, operating system, middleware, network, and application. This layered conception echoes the classic dependability taxonomy (Avizienis et al., 2001) and informs concrete engineering choices. At hardware and firmware layers, redundancy and error-correcting codes provide containment; at OS and middleware layers, checkpointing, replication, and reconfiguration provide recovery; and at application and operational layers, monitoring and forecasting enable predictive mitigation (Avizienis et al., 2001; Laprie, 2008). Layering reduces the need for any single mechanism to be perfect, creating graceful degradation paths and enabling independent verification domains.

2. Mode-change and mixed-criticality reasoning are essential for systems hosting heterogeneous workloads.

When systems host mixed-criticality tasks—e.g., safety-critical control loops alongside best-effort analytics—mechanisms to manage system modes and transitions are required to maintain safety and timeliness. Burns (2014) frames the problem of system mode changes as both general (applying to many systems) and criticality-based (where tasks differ in safety/certification requirements). Mode-change policies must be explicit, verifiable, and enforceable: designers must specify safe states, allowable transitions, and priority rules that govern resource allocation across criticality levels. These policies also influence recovery strategies; for example, in a degraded mode the system might suspend noncritical services, reserve resources for critical workloads, or switch to simplified but verifiable control algorithms (Burns, 2014; Kim & Lawrence, 1992).

3. Adaptive fault-tolerance offers significant benefits but compounds verification complexity.

Adaptive approaches—where the system changes replication levels, scheduling priorities, or algorithmic variants based on observed conditions—allow more efficient use of resources while maintaining assurances under varied fault profiles (Kim & Lawrence, 1992; Årzén, 2013). For instance, an adaptive controller may increase graceful degradation steps as monitored error rates rise, or selectively replicate critical tasks under worsening network conditions. However, such adaptability increases the state-space of possible behaviors, making formal verification and assurance arguments more challenging. To manage this complexity, the design must restrict the adaptation space to well-understood patterns, provide runtime invariants, and combine lightweight formal checks with rigorous testing and monitoring (Kim & Lawrence, 1992; Årzén, 2013).

4. Network-level strategies are critical for industrial and avionics-grade dependability.

Industrial Ethernet and specialized networking architectures are central to meeting hard availability and timing constraints in industrial systems. Alvarez et al. (2019) survey fault tolerance mechanisms for Ethernet-based industrial systems, showing how protocol-level redundancy, deterministic scheduling, and network segment isolation contribute to predictable behavior. Historical high-assurance systems, such as avionics controllers and the SIFT project, demonstrate that integrated design across computation and communication domains yields superior dependability (Wensley et al., 2008). For engineers, this implies that network design cannot be an afterthought: redundancy, deterministic paths, and isolation mechanisms must be embedded in system architecture from the beginning.

5. Reactive and proactive cloud fault-tolerance must be combined for practical resilience.

Cloud platforms face diverse failure modes—datacenter outages, network partitions, noisy neighbors, and software bugs—that require a blend of reactive recovery (failover, restart, reallocation) and proactive measures (predictive maintenance, graceful degradation). Reactive fault tolerance, such as the reactive approach proposed for cloud computing, allows the system to respond to observed failures with policies tuned to maintain service-level objectives (Abd Elfattah et al., 2017). Proactive strategies—like forecasting

failures from telemetry and applying preemptive migration—reduce the incidence of catastrophic outages but require robust prediction models and careful evaluation of false positives (Laprie, 2008; Abd Elfattah et al., 2017).

6. Machine learning can enhance fault detection and classification, but integration must be cautious.

Machine learning techniques applied to log analytics, anomaly detection, and fault classification provide promising avenues to speed diagnosis and enable early remediation (Kochhar & Hilda, 2017). However, ML models introduce new failure modes: model drift, overfitting, adversarial manipulation, and lack of interpretability. Thus, ML must be paired with mechanisms for validation, uncertainty quantification, human-in-the-loop decision paths, and conservative failover actions when model confidence is insufficient (Kochhar & Hilda, 2017; Tebaa & ElHajji, 2014).

7. Cooperative scheduling and task backfilling are practical mechanisms for improving utilization while preserving timeliness.

Priority-based cooperative computing and task backfilling techniques can help cloud and cluster systems improve resource utilization without violating deadlines or critical task requirements. Hasan and Goraya (2016) demonstrate how task backfilling can accommodate prioritized workloads in shared infrastructures. These methods rely on accurate estimates of task durations and carefully designed admission controls to prevent starvation or deadline misses; their efficacy is increased when combined with predictive telemetry and adaptive reservation strategies (Hasan & Goraya, 2016; Lin et al., 2010).

8. Multi-cloud strategies and diversification provide improved fault tolerance but add complexity and privacy considerations.

Moving workloads across multiple clouds can mitigate correlated failures and vendor-specific outages, and can support privacy-preserving partitioning of data and computation (Teba & ElHajji, 2014). However, multi-cloud solutions must reconcile different performance characteristics, API semantics, and security models. They require middleware to manage heterogeneity and standardized interfaces to facilitate migration and failover, as well as contractual and governance frameworks for cross-provider operation.

9. Manufacturing test infrastructures, especially for GPUs, require specialized fault-tolerance approaches that blend high serviceability with production constraints.

Designing fault-tolerant test infrastructure for large-scale GPU manufacturing is a domain where the cost of lost production due to testbed failures is high (Designing Fault-Tolerant Test Infrastructure, 2025). Strategies include redundancy in test harnesses, robust scheduling to isolate failing test stations, predictive maintenance on test equipment, and automated rollback procedures. Additionally, verifying the test infrastructure's own correctness is vital because it serves as an oracle for production acceptance; thus, meta-testing strategies and redundancy in verdict determination are often necessary (Designing Fault-Tolerant Test Infrastructure, 2025).

10. Standards, toolchains, and formal reasoning around mode changes remain underdeveloped areas ripe for research investment.

While many practical mechanisms exist, formal frameworks that support certified mode-change reasoning across software stacks are sparse. The complex interactions between adaptive middleware, ML-driven diagnostics, and network reconfiguration produce emergent behaviors that are difficult to certify with classical methods. This gap suggests a research program aimed at modular certification techniques, runtime verification tools tailored to mixed-criticality adaptive systems, and conformance testing for fault-handling behaviors (Burns, 2014; Kim & Lawrence, 1992).

These results collectively map a design space populated by mechanisms that can be composed to yield dependable and resilient systems. The subsequent Discussion elaborates the significance of these findings, considers counterarguments and limitations, and outlines avenues for future work.

DISCUSSION

The results synthesize into an overarching design philosophy: dependability is best achieved through layered, composable, and constrained adaptability. This section interprets the findings, weighs counter-arguments, articulates limitations, and proposes a forward-looking research agenda.

Interpreting layered defenses. The layered defense model draws strength from diversity: different layers can employ different assumptions, failure detection mechanisms, and recovery strategies. This multiplicity reduces the chance that a single unanticipated fault will propagate unchecked. However, layering invites complexity: interfaces between layers can hide assumptions that, if violated during complex failures, lead to systemic faults. Hence, the notion of "explicit contracts" between layers—where capabilities and invariants are specified and monitored—is crucial. Contracts reduce implicit dependency assumptions and enable localized reasoning about the correctness of recovery actions (Avizienis et al., 2001; Laprie, 2008).

Adaptive mechanisms: benefits and risks. Adaptive strategies are attractive because they reshape resource use in real time, allowing systems to preserve critical functions while economizing on redundancy during benign periods (Kim & Lawrence, 1992; Årzén, 2013). Yet adaptivity introduces non-determinism and a larger behavioral state-space, complicating certification and predictability. To reconcile adaptivity with assurance needs, designers should pursue "verified adaptation": restrict adaptive choices to a small set of formally analyzable modes and accompany runtime adaptation with invariant checks. Additionally, hybrid designs where adaptation is permitted only within controlled envelopes can provide much of the efficiency benefit without fully abandoning verifiability (Kim & Lawrence, 1992; Burns, 2014).

Mode-change management across contexts. Mode-change reasoning originates in real-time embedded systems, where safety-critical tasks must retain timeliness guarantees when the system transitions into degraded modes (Burns, 2014). The translation of these ideas to cloud and industrial contexts requires reframing: in clouds, "criticality" can correspond to agreed service-level objectives and business priorities rather than safety certification. Mode changes can thus be executed by orchestrators that enforce priority-aware placements and resource reservations, or by dynamic QoS controllers that throttle noncritical tenants. Important questions arise concerning how to validate these mode changes: is simulation sufficient, or are contractual audits required? The answer depends on the domain: avionics demands formal artifacts and exhaustive testing, while cloud operators may tolerate probabilistic assurances if robust monitoring and rollback paths exist (Burns, 2014; Abd Elfattah et al., 2017).

Role and limitations of ML in fault handling. Machine learning can detect anomalies faster than rule-based systems in complex telemetry but introduces opacity and fragility. False positives can result in unnecessary migrations, increasing system churn and reducing availability. False negatives may allow faults to escalate. To mitigate these risks, ML models must be designed with explainability and be integrated as part of a decision pipeline that includes conservative fallback strategies, human oversight for high-stakes decisions, and continuous retraining regimens to cope with drift (Kochhar & Hilda, 2017). The community must also invest in adversarial robustness research for diagnostic models to ensure resilience against manipulated telemetry.

Networking as a first-class dependability concern. Industrial systems demonstrate that networking cannot be left emergent; deterministic communication patterns, redundancy at protocol and physical layers, and failover routing are essential. In distributed cloud contexts, network partitioning creates subtle modes of failure: split-brain behaviors, delayed consistency, and state divergence. Classic distributed systems theory informs these challenges, but engineers must reconcile theoretical limits with economic and performance constraints. Techniques such as quorums, consensus protocols with partition-aware designs, and partition-tolerant state machine replication must be adapted and tuned for each domain's timeliness and availability

needs (Álvarez et al., 2019; Wensley et al., 2008).

Multi-cloud and privacy trade-offs. The promise of multi-cloud resiliency is tempered by complexities in identity, data governance, and network performance. Privacy-preserving multi-cloud designs, as proposed by some researchers, must balance the additional overhead of cryptographic techniques and data partitioning with the resilience benefits of avoiding single-provider lock-in (Tebaa & ElHajji, 2014). Operationally, multi-cloud orchestration increases the surface area for misconfiguration and requires sophisticated policy engines and trust frameworks.

Manufacturing test infrastructure: meta-dependability. Test infrastructures are unique because they are meta-systems: their role is to certify product correctness, making their own correctness critical. For GPU manufacturing, the challenge is ensuring the testbed's availability and the veracity of test results under heavy throughput. Redundancy in test oracles, cross-validation of test outcomes, and mechanisms to detect systemic test infrastructure faults are necessary to avoid false acceptance or rejection of devices (Designing Fault-Tolerant Test Infrastructure, 2025). Furthermore, integrating predictive maintenance for test rigs reduces downtime and improves throughput.

Limitations of the synthesis. This article synthesizes existing work rather than reporting new empirical results, and that choice imposes limitations. Without new experimental data, conclusions about the effectiveness of particular strategies in specific contexts remain provisional and contingent on implementation details. Additionally, the provided literature, while broad, reflects specific schools of thought and domain focuses; results would be strengthened by integrating broader empirical studies and industry case reports. Finally, the emphasis on conceptual clarity necessarily abstracts from some practical implementation minutiae that engineers will confront in deployments.

Research agenda and future directions. The findings suggest several promising research directions:

- Formalizing adaptive mode-change policies. Develop formal languages and verification techniques that specify admissible adaptations, produce proofs or probabilistic assurances of safety under adaptation, and enable runtime checks. This includes composing verification results across layers so that global system properties can be inferred from local certificates (Burns, 2014; Kim & Lawrence, 1992).
- Robust ML diagnostics. Create methods for training diagnostic models under nonstationary conditions, quantify uncertainty in predictions, and design conservative response policies that incorporate model confidence. Evaluate ML diagnostic pipelines under adversarial and distribution-shift scenarios to ensure reliability in production (Kochhar & Hilda, 2017).
- Standardized evaluation frameworks. Propose benchmark suites and metrics for comparing fault-tolerance techniques across cloud, industrial, and manufacturing contexts. Benchmarks should measure availability, degradation modes, recovery time, resource overhead, diagnosability, and operational complexity (Designing Fault-Tolerant Test Infrastructure, 2025).
- Cross-layer orchestration tooling. Build middleware that can coordinate page-level mode changes, network reconfiguration, and application-level fallbacks, with clear contracts and monitoring hooks. The middleware should support declarative specification of criticality and provide introspection and audit trails for post-hoc analysis (Álvarez et al., 2019; Abd Elfattah et al., 2017).
- Meta-testing and meta-redundancy for testbeds. For manufacturing testbeds, design meta-testing strategies that periodically validate the testbed against known-good artifacts, enabling automatic detection of test infrastructure faults and preventing systemic false verdicts (Designing Fault-Tolerant Test Infrastructure, 2025).

These directions converge on the need to make adaptability accountable and verifiable. The trade-offs among efficiency, assurance, and complexity will remain central, but carefully constrained adaptive strategies have strong potential to increase resilience without unacceptable growth in verification burden.

CONCLUSION

This article presented a comprehensive synthesis of dependability theory and applied fault-tolerance strategies suited to contemporary distributed, embedded, cloud, and manufacturing test environments. The key message is that dependability is best achieved not by monolithic methods but by multi-layered, composable defenses that combine redundancy, explicit mode-change reasoning, adaptive control, network-aware design, and data-driven diagnostics. Mixed-criticality and mode-change frameworks provide a rigorous way to reason about resource reallocation in the face of faults, while adaptive systems and ML-driven diagnostics offer efficiency and early detection if coupled with conservative safeguards and verifiability measures. Industrial networking paradigms and avionics-grade designs underscore the necessity of integrating communication reliability into overall system design. Cloud and multi-cloud contexts introduce economics and governance challenges that must be resolved by middleware that harmonizes heterogeneity and policy.

For practitioners, the practical recommendations are: design explicit mode-change policies; constrain adaptivity to formally analyzable envelopes; embed network redundancy and deterministic communication where timeliness is critical; apply ML diagnostics with uncertainty-aware policies and human oversight; and adopt meta-testing for production test infrastructures. For researchers, important work remains in formalizing adaptive policies, building robust diagnostic models, standardizing evaluation frameworks, and inventing cross-layer orchestration tools that make resilience verifiable and manageable.

The synthesis provided here paints a path forward: a future where dependable systems are designed with predictable, constrained adaptability and where verification, monitoring, and human oversight form a coherent safety net. Achieving that future will require collaborative effort across theoretical computer science, control and systems engineering, networking, and applied machine learning.

REFERENCES

1. Avizienis, A.; Laprie, J.C.; Randell, B. Fundamental Concepts of Dependability. UCLA CSD Report no. 010028, LAAS Report no. 01-145, Newcastle University Report no. CS-TR-739, 2001.
2. Burns, A. System Mode Changes—General and Criticality-Based. In Proceedings of the 2nd Workshop on Mixed Criticality Systems (WMC), RTSS, Rome, Italy, 2 December 2014.
3. Kim, K.H.K.; Lawrence, T.F. Adaptive fault-tolerance in complex real-time distributed computer system applications. *Comput. Commun.* 1992, 15, 243–251.
4. Årzén, K.E. Preface to special issue on adaptive embedded systems. *Real-Time Syst.* 2013, 49, 337–338.
5. Laprie, J.C. From dependability to resilience. In Proceedings of the 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, Anchorage, AK, USA, 24–27 June 2008.
6. Knight, J.; Strunk, E.; Sullivan, K. Towards a rigorous definition of information system survivability. In Proceedings of the DARPA Information Survivability Conference and Exposition, Washington, DC, USA, 22–24 April 2003; pp. 78–89.
7. Proenza, J.; Barranco, M.; Ballesteros, A.; Álvarez, I.; Gessner, D.; Derasevic, S.; Rodríguez-Navas, G. DFT4FTT Project. Available online: <http://srv.uib.es/dft4ftt/> (accessed on 1 September 2022).
8. Álvarez, I.; Ballesteros, A.; Barranco, M.; Gessner, D.; Djerasevic, S.; Proenza, J. Fault Tolerance in Highly Reliable Ethernet-Based Industrial Systems. *Proc. IEEE* 2019, 107, 977–1010.
9. Wensley, J.; Lamport, L.; Shostak, R.; Weinstock, C.; Goldberg, J.; Green, M.; Levitt, K.; Melliar-

Smith, P. SIFT: Design and analysis of a fault-tolerant computer for aircraft control. Proc. IEEE 2008, 66, 1240–1255.

10. Abd Elfattah, E.; Elkawkagy, M.; El Sisi, A. A Reactive Fault Tolerance Approach for Cloud Computing. In Proceedings of the 13th International IEEE Computer Engineering Conference (ICENCO'17), 2017, pp. 190–194.

11. Hasan, M.; Goraya, M. S. Priority Based Cooperative Computing in Cloud Using Task Backfilling. Lect. Notes Software Eng., Vol. 4, 2016, pp. 229–233.

12. Kochhar, D.; Hilda, A. K. J. An Approach for Fault Tolerance in Cloud Computing Using Machine Learning Technique. Int. J. Pure Appl. Math., Vol. 117, 2017, No. 22, pp. 345–351.

13. Gupta, S.; Gupta, B. B. XSS-Secure as a Service for the Platforms of Online Social Network-Based Multimedia Web Applications in the Cloud. Multimedia Tools Appl., Vol. 77, 2018, No. 4, pp. 4829–4861.

14. Tebaa, M.; El Hajji, S. From Single to Multi-Clouds Computing Privacy and Fault Tolerance. In Proceedings of the International Conference on Future Information Engineering, Elsevier B. V., 2014, pp. 112–118.

15. Abid, A.; Khemakhem, M. T.; Marzouk, S.; Bem Jemaa, M.; Monteil, T.; Drira, K. Toward Ant Fragile Cloud Computing Infrastructures. Procedia Computer Science, Vol. 32, 2014, pp. 850–855.

16. Lin, X.; Mamat, A.; Lu, Y.; Deogun, J.; Goddard, S. Real-Time Scheduling of Divisible Loads in Cluster Computing Environments. Journal of Parallel and Distributed Computing, Vol. 70, 2010, pp. 296–308.

17. Jhawar, R.; Piuri, V. Fault Tolerance and Resilience in Cloud Computing Environments. In Computer and Information Security Handbook. 2013, pp. 1–29.

18. Sun, D.; Chang, G.; Miao, C.; Wang, X. Modelling and Evaluating a High Serviceability Fault Tolerance Strategy in Cloud Computing Environments. International Journal of Security and Networks, Vol. 7, 2012, pp. 196–210.

19. Designing Fault-Tolerant Test Infrastructure for Large-Scale GPU Manufacturing. International Journal of Signal Processing, Embedded Systems and VLSI Design, 2025, 5(01), 35–61.