

Resilient by Design – Active Inference for Distributed Continuum Intelligence

Praveen Kumar Donta*, Alfreds Lapkovskis*, Enzo Mingozzi†, and Schahram Dustdar‡

*Department of Computer Systems and Sciences, Stockholm University, Stockholm 164 25, Sweden.

†Department of Information Engineering, University of Pisa, Pisa, Italy.

‡Distributed Systems Group, TU Wien, Vienna 1040, Austria and ICREA Barcelona 08002, Spain

{praveen, alfreds.lapkovskis}@dsv.su.se, enzo.mingozzi@unipi.it, dustdar@dsg.tuwien.ac.at

Abstract—Failures are the norm in highly complex and heterogeneous devices spanning the distributed computing continuum (DCC), from resource-constrained IoT and edge nodes to high-performance computing systems. Ensuring reliability and global consistency across these layers remains a significant challenge, particularly for AI-driven workloads that require real-time, adaptive coordination. This work-in-progress paper introduces a Probabilistic Active Inference Resilience Agent (PAIR-Agent) to achieve resilience in DCC systems. PAIR-Agent performs three core operations: (i) constructing a causal fault graph from device logs, (ii) identifying faults while managing certainties and uncertainties using Markov blankets and the free energy principle, and (iii) autonomously healing issues through active inference. Through continuous monitoring and adaptive reconfiguration, the agent maintains service continuity and stability under diverse failure conditions. Theoretical validations confirm the reliability and effectiveness of the proposed framework.

Index Terms—Distributed Computing Continuum; Active Inference; Resilience; Free energy principle;

I. INTRODUCTION

NO one can deny that computing infrastructure has transformed dramatically with Generative AI, driving unprecedented demand for high-performance resources spanning sensor nodes, edge devices, network elements, and server clusters. In response, the Distributed Computing Continuum (DCC) paradigm has emerged, orchestrating complex tasks across the full device continuum, allocating smaller, latency-sensitive tasks locally to reduce energy use and preserve privacy, while offloading resource-intensive tasks to high-performance servers [1], [2]. Such dynamic task distribution optimizes resource use, balances energy demand, and supports scalable operations. As the number and diversity of devices increase, data volumes surge, and workloads become more dynamic, system complexity concurrently grows, intensifying failure risks [3]. In general, task processing failures in the DCC occur due to factors such as user-initiated aborts, incorrect inputs, insufficient allocated resources, execution errors, node crashes, or communication failures; In DCC, additional failures include the unpredictability brought by device heterogeneity, mobile and battery-dependent nodes, dynamic network and resource

shifts, and large-scale consistency challenges associated with hundreds or thousands of edge devices.

Resilience mechanisms are crucial for modern computing systems, since failures are the norm in these infrastructures. The term resilience is chosen over mere fault tolerance in DCC because these systems must not only tolerate discrete faults but also effectively manage and adapt to a wide range of inherent uncertainties [4], [5]. In DCC, resilience reflects the system's capability to sustain continuous operation through proactive adaptation, localized recovery, and graceful degradation across diverse cloud, fog, edge, and IoT layers, where unpredictable workload fluctuations, device mobility, connectivity disruptions, and energy variability are common [6], [7]. To achieve these, this paper designs a novel probabilistic active inference resilience agent (PAIR-Agent), with the following contributions.

- PAIR-Agent builds a probabilistic Causal Fault Graph (CFG) using Bayesian Network Structure Learning (BNSL) through various log data from the devices.
- Next, it uses the Markov blanket to identify potential faults from CFGs. Further, PAIR-Agent applies the free energy principle for uncertainty-aware fault inference.
- Finally, PAIR-Agent autonomously selects corrective actions, minimizing expected free energy and enabling self-healing through active inference.

II. SYSTEM MODEL

The DCC ($\mathcal{S}$) are composed of n devices including one or c cloud or data centers $\mathcal{C} = \{C_1, C_2, \dots, C_c\}$, a set of f Fog nodes $\mathcal{F} = \{F_1, F_2, \dots, F_f\}$, a set of e Edge nodes $\mathcal{E} = \{E_1, E_2, \dots, E_e\}$, m mobile devices $\mathcal{M} = \{M_1, M_2, \dots, M_m\}$, a set of ι IoT devices $\Gamma = \{I_1, I_2, \dots, I_\iota\}$, and a set of ς sensor nodes (SNs) $\Psi = \{\psi_1, \psi_2, \dots, \psi_\varsigma\}$. We assume that Ψ are data producers and they are not capable of computing any data. Therefore, they transmit their data to the nearest base station, i.e., Γ , $\mathcal{M}$, or $\mathcal{E}$ using wireless communication through Bluetooth Low Energy, Zigbee, Cellular, or Wi-Fi channels [8].

A set of AI-related computational tasks is defined as $\mathcal{T} = \{T_1, T_2, \dots, T_t\}$, where each task $T_i = \langle \omega_i, \chi_i, \phi_i, \delta_i \rangle$

denotes workload ω_i , input data dependencies χ_i , a computational mapping or model function $\phi_i : \chi_i \mapsto y_i$, and an execution deadline δ_i . The DCC graph $\mathcal{S} = (\mathcal{V}, \mathcal{R})$ consists of nodes $\mathcal{V} = \mathcal{C} \cup \mathcal{F} \cup \mathcal{E} \cup \mathcal{M} \cup \Gamma \cup \Psi$ and communication relationships $\mathcal{R} \subseteq \mathcal{V} \times \mathcal{V}$, where each node $v_j \in \mathcal{V}$ executes a subset of tasks $\mathcal{T}(v_j)$ such that $\bigcup_{v_j \in \mathcal{V}} \mathcal{T}(v_j) = \mathcal{T}$. Each task execution $T_i(v_j)$ follows a three-phase structure $\langle \alpha_i, \beta_i, \gamma_i \rangle$, representing task initiation, computation, and completion or checkpointing, respectively. In general, these complex AI workflows, such as training, are composed of sequential subtasks, e.g., $T_{\text{train}} = T_{\text{data-load}} \triangleright T_{\text{forward}} \triangleright T_{\text{backward}} \triangleright T_{\text{update}}$, where each subtask can complete and save progress independently to support partial recovery. Each task establishes a stable checkpoint $\zeta_i = \langle s_i, D_i, \theta_i \rangle$, recording the intermediate state s_i , processed dataset segment D_i , and model parameters θ_i , enabling restoration and continuation upon failure.

III. RESILIENCE IN ACTION

The proposed PAIR-Agent iteratively performs the following steps (as depicted in Figure 1): (a) identifying changes in the system logs and collecting data to generate a CFG; (b) performing fault inference by determining certain and uncertain faults using the Markov blanket and free energy principle [9]; and (c) healing the faults through active inference to achieve resilience, as discussed in more detail in the subsequent subsections.

A. Collect and Process Logs

Initially, the PAIR-Agent coordinates distributed log acquisition across the $\mathcal{S}$ using remote procedure calls (RPCs) to query each device ($v_i \in \mathcal{S}$) for its most recent checkpoint ($\zeta_i^t = \langle s_i^t, D_i^t, \theta_i^t \rangle$), which encapsulates the local state, processed dataset segment, and learned parameters at analysis round t . These checkpoints act as temporal anchors that partition system observations into discrete analysis intervals, allowing the agent to process only the log entries generated after ζ_i^t . This design ensures synchronization across heterogeneous nodes, prevents redundant ingestion of log segments, tracks round-based progress, and maintains a globally consistent causal timeline of system activity [10]. Each new round ($t + 1$), therefore, begins with a synchronized collection of incremental log updates ΔL_i^{t+1} , forming a time-evolving dataset that supports continual inference.

After collection, raw logs are parsed and normalized into a structured representation. PAIR-Agent extracts features corresponding to potential fault interpretation or indicators, including user-aborted tasks, resource allocation failures (due to device mobility, computational capacity, or energy level), node crashes, execution-time anomalies, data inconsistencies, and communication disruptions, among others. Unlike traditional distributed systems, PAIR-Agent considers each fault as a probabilistic event f_i embedded within a continuous spatiotemporal context rather than an isolated deterministic failure. Each f_i is characterized by a conditional

probability distribution $P(f_i \mid \mathbf{x}_i, t)$, where $\mathbf{x}_i$ represents contextual features extracted from the normalized log matrix and t denotes the checkpoint time window. This probabilistic modeling allows the system to capture uncertainty and latent dependencies among hardware, software, and environmental factors influencing fault occurrence across heterogeneous nodes.

To uncover causal dependencies among probabilistic fault variables, PAIR-Agent constructs a time-evolving CFG $\mathcal{G}_{\text{fault}}^{t+1} = (\mathcal{V}_f^{t+1}, \mathcal{E}_f^{t+1})$ using a BNSL. Here, each node $v_i \in \mathcal{V}_f^{t+1}$ represents a stochastic variable corresponding to a specific fault or contextual feature extracted from the normalized feature matrix X^{t+1} , and each directed edge $e_{ij} \in \mathcal{E}_f^{t+1}$ encodes a conditional dependency $P^{t+1}(v_j \mid \text{Pa}(v_j))$, where $\text{Pa}(v_j)$ denotes the parent set of v_j . This structure is inferred incrementally at each checkpoint interval $(t, t + 1]$, allowing PAIR-Agent to integrate newly observed evidence $\Delta \mathcal{D}^{t+1}$ while retaining prior knowledge from earlier analysis rounds. This is achieved through Hill-Climbing with Bayesian Dirichlet equivalence scoring [11], combined with probabilistic regularization to handle noise and partial observability in DCC data. In scenarios with strong nonlinearity or non-stationary correlations, PAIR-Agent employs variational Bayesian learners to refine causal relations across checkpoints adaptively [12]. Each updated Bayesian network $\mathcal{G}_{\text{fault}}^{t+1}$ represents the posterior belief over causal structure given new observations, expressed as $P^{t+1}(\mathcal{G}_{\text{fault}} \mid \mathcal{D}^{t+1}) \propto P(\mathcal{D}^{t+1} \mid \mathcal{G}_{\text{fault}}) P^t(\mathcal{G}_{\text{fault}})$, where $P^t(\mathcal{G}_{\text{fault}})$ encodes the prior structural belief accumulated from previous rounds.

B. Fault Inference

Building upon the Bayesian CFG $\mathcal{G}_{\text{fault}}^{t+1}$, PAIR-Agent performs probabilistic fault inference by exploiting the local Markov properties of the network. For each fault variable $f_i \in \mathcal{V}_f^{t+1}$, the agent defines its Markov blanket $\mathcal{M}(f_i) = \text{Pa}(f_i) \cup \text{Ch}(f_i) \cup \text{Pa}(\text{Ch}(f_i))$, where $\text{Pa}(f_i)$ and $\text{Ch}(f_i)$ denote the parent and child sets of f_i , respectively [13]. This blanket represents the minimal sufficient set of variables that render f_i conditionally independent from the rest of the network. By monitoring the probabilistic interactions within $\mathcal{M}(f_i)$, the agent computes the posterior probability of fault activation or propagation, i.e.

$$P^{t+1}(f_i \mid \mathcal{M}(f_i)) \propto P^{t+1}(f_i \mid \text{Pa}(f_i)) \times \prod_{c_j \in \text{Ch}(f_i)} P^{t+1}(c_j \mid \text{Pa}(c_j)).$$

At each checkpoint interval $(t, t + 1]$, the agent continuously updates these beliefs using newly collected evidence $\Delta \mathcal{D}^{t+1}$. Since the exact computation of posteriors is generally infeasible for complex distributions, PAIR-Agent integrates free energy minimization with the Bayesian CFG and Markov blanket structure to infer approximate posterior beliefs. It utilizes the free energy principle (FEP) [13] to align its




Fig. 1: The proposed PAIR-Agent for resilience in the DCC collects and normalizes logs to form a CFG, infers faults via Markov blanket and FEP analysis, and heals through active inference to restore system stability.

internal model with observed system behavior continuously. At each checkpoint $t+1$, the agent computes an approximate posterior belief $Q^{t+1}(\mathbf{f})$ by minimizing the variational free energy, defined as

$$F^{t+1} \triangleq \mathbb{E}_{Q^{t+1}(\mathbf{f})} \left[\ln Q^{t+1}(\mathbf{f}) - \ln P^{t+1}(\mathbf{f}, \mathbf{x}^{t+1}) \right],$$

where $Q^{t+1}(\mathbf{f})$ approximates the posterior belief over the faults $\mathbf{f} = \{f_i\}_i$, given the contextual features $\mathbf{x}^{t+1} = \{\mathbf{x}_i^{t+1}\}_i$. Minimizing F^{t+1} updates these beliefs, refining the agent's estimates under model uncertainty and capturing dependencies among faults.

To further refine fault-type attribution, PAIR-Agent partitions the inferred fault space into hardware-related and software-related components, $\mathcal{V}_f^{t+1} = \mathcal{V}_{\text{HW}}^{t+1} \cup \mathcal{V}_{\text{SW}}^{t+1}$. The hardware set $\mathcal{V}_{\text{HW}}^{t+1}$ captures physical and low-level operational indicators, such as temperature variations [14], power fluctuations, and connectivity degradation, while the software set $\mathcal{V}_{\text{SW}}^{t+1}$ represents task failures, model divergence, resource contention, and execution-time anomalies. Using this decomposition, PAIR-Agent performs conditional marginalization to obtain $P^{t+1}(f_i | \mathcal{V}_{\text{HW}}^{t+1})$ and $P^{t+1}(f_i | \mathcal{V}_{\text{SW}}^{t+1})$, which enables the agent to differentiate whether a detected fault most likely originates from hardware-level disturbances or software-driven inconsistencies. Importantly, this stage performs only *fault inference*; the subsequent step handles all decision-making and recovery actions based on these inferred fault states.

C. Active Inference for Healing

Once fault inference identifies faults, PAIR-Agent employs active inference [15] to autonomously plan and execute corrective actions. The agent evaluates a set of possible

actions $a_t \in \mathcal{A}$ and selects the one that minimizes the expected free energy $G(\cdot)$:

$$a_t^* = \arg \min_{a_t} G(a_t).$$

The expected free energy for action a_t is given by

$$G(a_t) \triangleq D_{\text{KL}}[Q^{t+1}(\mathbf{x} | a_t) \| P^*(\mathbf{x})] + \mathbb{E}_{Q^{t+1}(\mathbf{f}|a_t)}[H[P(\mathbf{x} | \mathbf{f})]],$$

where $P^*(\mathbf{x})$ encodes the preferred contextual features, $Q^{t+1}(\mathbf{f} | a_t)$ denotes the predicted posterior belief about the faults under action a_t , $P(\mathbf{x} | \mathbf{f})$ is the likelihood of observing the contextual features, given faults, and $Q^{t+1}(\mathbf{x} | a_t) \triangleq \sum_{\mathbf{f}} Q^{t+1}(\mathbf{f} | a_t) P(\mathbf{x} | \mathbf{f})$ is the predictive distribution over future contextual features under action a_t . By minimizing $G(a_t)$, PAIR-Agent selects actions that steer the system toward its preferred operational state while explicitly accounting for uncertainty in both hardware-level indicators ($\mathcal{V}_{\text{HW}}^{t+1}$) and software-level behavioral variables ($\mathcal{V}_{\text{SW}}^{t+1}$). After executing the selected action, PAIR-Agent observes the resulting system state, updates posterior fault probabilities $P^{t+1}(f_i | \mathbf{x}_i^{t+1})$ with new evidence, and advances the checkpoint ζ^{t+1} to record the updated configuration and belief state. This closed-loop cycle of *perception–inference–action–update* enables continual adaptation and resilience across analysis rounds. During the healing process, PAIR-Agent not only redistributes tasks but also performs device-level recovery through automated reinitialization, firmware reloading, or adaptive configuration tuning. If such self-healing measures fail to restore stability, the agent isolates the faulty node and applies corrective actions at the continuum level $\mathcal{S}$, including workload redistribution across fog or cloud layers, rerouting through redundant network paths, managing dynamic load or thermal

management, and escalation to human operators for physical interventions.

IV. THEORETICAL RESULTS AND DISCUSSIONS

This section presents the theoretical foundations showing how PAIR-Agent achieves resilient operation in the DCC through principled probabilistic inference and cautious, uncertainty-aware action selection. The results highlight four key properties: the locality and scalability of fault inference, the quality of the variational posterior approximation, the safety of the agent's corrective actions, and the robustness of belief updates under missing or delayed logs.

Result 1 (Locality and Scalability of Fault Inference). For any fault variable f_i , PAIR-Agent only needs the variables in its Markov blanket $\mathcal{M}(f_i)$ (its parents, children, and co-parents) to update the belief $Q(f_i)$. Thus, the time and memory per fault update depend on $|\mathcal{M}(f_i)|$, and not on the total number of variables in the CFG. If the average blanket size is B , updating all d faults in one round costs $O(d \cdot B)$, which stays manageable as the DCC grows, provided local neighborhoods stay small. This guarantees that the pipeline in Fig. 1 scales due to inference being local by design.

Result 2 (Best Possible Approximate Posterior). At checkpoint $t+1$, PAIR-Agent selects $Q^{t+1}(\mathbf{f})$ by minimizing the variational free energy $F^{t+1}(Q)$. Since, $F^{t+1}(Q)$ differs from $D_{KL}[Q(\mathbf{f}) \parallel P^{t+1}(\mathbf{f} \mid \mathbf{x}^{t+1})]$ only by a constant independent of Q (i.e., $-\ln P^{t+1}(\mathbf{x}^{t+1})$), the update $Q^{t+1} = \arg \min_{Q \in \mathcal{Q}} F^{t+1}(Q)$ is equivalently the closest distribution to the exact Bayesian posterior within the chosen variational family $\mathcal{Q}$. Therefore, given the current CFG and log data, PAIR-Agent computes the best posterior approximation it can represent, ensuring principled and internally consistent fault inference.

Result 3 (Conservative Action Selection [Never Worse Than Doing Nothing]). Assume that the action set includes a baseline "do nothing" action $a^\emptyset$, which keeps the current system configuration. Then, because PAIR-Agent selects $a_t^* = \arg \min_{a_t} G(a_t)$, the chosen action never has higher expected free energy than $a^\emptyset$. Conversely, when an action is predicted to move $Q(\mathbf{x} \mid a)$ closer to preferences and/or reduce uncertainty, it will be selected. This guarantees that the agent will not take an action that it predicts will diverge the system from its operational objectives and not increase its model certainty.

Result 4 (Robustness to Missing or Delayed Logs). At checkpoint $t+1$, the belief update for any fault f_i depends only on variables in its Markov blanket $\mathcal{M}(f_i)$. Therefore, missing or delayed logs for variables outside $\mathcal{M}(f_i)$ do not change the updated belief $Q^{t+1}(f_i)$. Since telemetry in DCC is often partial or out-of-order, this property makes PAIR-Agent's inference insensitive to unrelated data gaps and supports asynchronous, scalable updates: only logs inside the small local neighborhood $\mathcal{M}(f_i)$ can affect the diagnosis of f_i .

V. CONCLUSIONS

This paper introduced the PAIR-Agent as a unified framework for achieving resilience in DCC systems. The PAIR-Agent agent performs: (i) constructing a CFG through device log analysis, (ii) identifying faults while managing uncertainties using Markov blankets and the FEP, and (iii) autonomously healing issues via active inference. Through its closed-loop cycle of observation, inference, and action, PAIR-Agent ensures adaptive stability, self-healing capability, and sustained operational continuity, marking a significant step toward autonomous and resilient DCC systems. Theoretical results confirm its scalability, principled and accurate fault inference, safe decision-making, and robustness to incomplete telemetry. Future work will extend this research with detailed and distributed implementation, empirical validation, and evaluation under a real-time testbed.

REFERENCES

- [1] S. Dustdar, V. C. Pujol, and P. K. Donta, "On distributed computing continuum systems," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 4, pp. 4092–4105, 2023.
- [2] P. K. Donta, B. Sedlak, I. Murturi, V. Casamayor-Pujol, and S. Dustdar, "Human-based distributed intelligence in computing continuum systems," *IEEE Internet Computing*, vol. 29, no. 2, 09 2025.
- [3] P. K. Donta, B. Sedlak, V. Casamayor Pujol, and S. Dustdar, "Governance and sustainability of distributed continuum systems: a big data approach," *Journal of Big Data*, vol. 10, no. 1, pp. 1–31, 2023.
- [4] A. A. Wang, Y. Ma, Y. Zhang, Z. Zhao, Y. Cai, X. Qiu, B. Dong, and C. He, "Perturbation-resilient integer arithmetic using optical skyrmions," *Nature Photonics*, pp. 1–9, 2025.
- [5] L. Svobodova, "Resilient distributed computing," *IEEE Transactions on Software Engineering*, no. 3, pp. 257–268, 2009.
- [6] M. Sagor, A. Haroon, R. Stoleru, S. Bhunia, A. Altaaweel, M. Chao, L. Jin, M. Maurice, and R. Blalock, "DistressNet-NG: A resilient data storage and sharing framework for mobile edge computing in cyber-physical systems," *ACM Transactions on Cyber-Physical Systems*, vol. 8, no. 3, pp. 1–31, 2024.
- [7] H. Fang, P. Yu, C. Tan, J. Zhang, D. Lin, L. Zhang, Y. Zhang, W. Li, and L. Meng, "Self-healing in knowledge-driven autonomous networks: Context, challenges, and future directions," *IEEE Network*, vol. 38, no. 6, pp. 425–432, 2024.
- [8] O. Cheikhrouhou, "Secure group communication in wireless sensor networks: a survey," *Journal of Network and Computer Applications*, vol. 61, pp. 115–132, 2016.
- [9] M. Kirchhoff, T. Parr, E. Palacios, K. Friston, and J. Kiverstein, "The markov blankets of life: autonomy, active inference and the free energy principle," *Journal of The royal society interface*, vol. 15, no. 138, p. 20170792, 2018.
- [10] A. Saleh, R. Morabito, S. Tarkoma, A. Lindgren, S. Pirttikangas, and L. Lovén, "Agentic TinyML for intent-aware handover in 6G wireless networks," *arXiv preprint arXiv:2508.09147*, 2025.
- [11] M. Scutari, "Dirichlet bayesian network scores and the maximum relative entropy principle," *Behaviormetrika*, vol. 45, no. 2, pp. 337–362, 2018.
- [12] K. J. Friston, L. Harrison, and W. Penny, "Dynamic causal modelling," *Neuroimage*, vol. 19, no. 4, pp. 1273–1302, 2003.
- [13] K. Friston, "The free-energy principle: a unified brain theory?" *Nature reviews neuroscience*, vol. 11, no. 2, pp. 127–138, 2010.
- [14] A. Lapkovskis, B. Sedlak, S. Magnússon, S. Dustdar, and P. K. Donta, "Benchmarking dynamic slo compliance in distributed computing continuum systems," in *2025 IEEE International Conference on Edge Computing and Communications (EDGE)*, 2025, pp. 93–102.
- [15] N. Sajid, P. J. Ball, T. Parr, and K. J. Friston, "Active inference: demystified and compared," *Neural computation*, vol. 33, no. 3, pp. 674–712, 2021.