

Context-Driven Cooperative Intelligent Control for Distributed Cyber-Physical Actuation Platforms Using CTDE Multi-Agent Reinforcement Learning

P.Joshua Reginald

Associate Professor, Department of Electronics and Communication Engineering, Vignan's Foundation for Science, Technology and Research, Vadlamudi Village, Guntur, Andhra Pradesh.

KEYWORDS:

Cyber-Physical Systems,
Cooperative Control,
Context-Aware Intelligence,
Multi-Agent Reinforcement Learning,
CTDE,
Intelligent Environments

ARTICLE HISTORY:

Submitted : 21.07.2025
Revised : 17.08.2025
Accepted : 10.09.2025

DOI:

<https://doi.org/10.17051/RANGWCS/01.01.06>

ABSTRACT

Modern ubiquitous and intelligent environments rely on distributed cyber-physical actuation platforms, which are difficult to control because of tight integration between actuators and partial observability as well as would not be achievable to net out centralised and fixed notions of control in dynamical environments. Centralised strategies are lacking in scalability as well as robustness limits whereas independent controllers fail to reduce cross-actuator interference. This paper suggests using Centralised Training and Decentralised Exercises (CTDE) Multi-Agents reinforcement learning (MARL) context-based control framework to propose a cooperative framework of intelligent control. Training the agents use globally available context information such as couplings strength and environmental conditions to acquire coordinated control policy during training, though execution is completely decentralised based on local observations only. The agent state representation incorporates the idea of context embedding to facilitate adaptive decision making and decision making that is aware of interference. The results of extensive simulations indicate that the proposed framework can significantly enhance control precision and reduce any interference cross-actuator in addition to ensuring stable convergence at alternative coupling levels in comparison with centralised and non-cooperative baselines. These findings demonstrate the efficiency of CTDE based cooperative learning of scalable and robust control and demonstrate its applicability to next generation ubiquitous computing and intelligent cyber-physical spaces.

Author's e-mail: drpjr_ece@vignan.ac.in

How to cite this article: Reginald PJ, Context-Driven Cooperative Intelligent Control for Distributed Cyber-Physical Actuation Platforms Using CTDE Multi-Agent Reinforcement Learning, Recent Advances in Next-Generation Wireless Communication Systems , Vol. 1, No. 1, 2026 (pp.43-50).

1. INTRODUCTION

Smart environments ubiquitous computing environments are increasingly based on distributed platforms of cyber-physical actuation to make it possible to combat smartness in uses like industrial automation, smart buildings, and autonomous robotics. Such environments combine high density of spatially dispersed sensing, communication and actuation assets that are required to work jointly with strict resource and reliability considerations. We have already noticed in the literature regarding wireless

sensor networks (WSNs) and intelligent environments the significance of energy-efficient communication, dynamic data collection, and scalable coordination strategies to help such distributed systems [1], [4], [7]. With the cyber-physical infrastructures approaching greater autonomy and closer connexion with the physical world, the proper management of distributed actuators will become a fundamental issue in the need to preserve the overall performance and strength of the system.

In spite of these developments, the distributed actuation presents a number of underlying issues such

as strong coupling and interference between actuators, the global systems states can only be partially observed, and, as the system size increases, scaling is problematic. It can occur as a result of interference caused by the use of common physical or communication resources that can very easily cause control inaccuracies and stability to be compromised unless they are dealt with [2], [5]. In addition, the actuators are mostly local sensing in nature and have inadequate communication resulting in a lack of full situational awareness and response time. The conventional data-driven coordination and clustering solutions, which have been offered to manage WSNs, can solve certain scalability challenges but are not adequate to tightly coupled control activities in dynamic cyber-physical settings [3], [11]. These issues require control strategies that are more adaptive and cooperative and able to react intelligently to both environmental and system conditions.

Such distributed cyber-physical actuation platforms are not always compatible with existing centralised and heuristic control strategies. The major drawbacks of centralised schemes include high-communication overheads, lack of scalability, and susceptibility to single point failures especially when it comes to large-scale intelligent environments [6], [9]. Conversely, the heuristic or fixed control policies are not adaptive and do not take advantage of contextual information hence yielding suboptimal results in dynamic interference and changes in workload [8], [12]. Whereas machine-learning-based methods have been applied to network-level protocols and routing choices, they have not been widely applied to cooperative actuation control, particularly in environments involving coordination as well as decentralised performance [10].

This paper attempts to solve these shortcomings by developing a situational-dependent framework of cooperative intelligent control on distributed platforms of cyber-physical actuation implemented under Centralised Training and Decentralised Execution (CTDE) Multi-Agent Reinforcement Learning (MARL). The proposed architecture makes actuators train cooperative control policies with the help of globally available contextual information and provides fully decentralised and scalable interaction at the run-time. Through incorporating contextual clues with respect to coupling strength and environmental dynamics into the process of learning, the framework attains the adaptive interference mitigation and accuracy in control. The three primary contributions of this work are as follows: (i) it presents a context-aware cooperative architecture of distributed actuation, (ii) it presents a formulation of CTDE-based MARL models specific to the dynamics of cyber-physical systems with interference coupling, and (iii) it evaluates the effectiveness of the present methodology on a large-scale basis by providing a

quantitative performance assessment of the architecture (in terms of accuracy, stability and robustness) in the face of dynamic coupling conditions.

2. RELATED WORK

Distributed control known as cyber-physical systems has undergone great insight with the focus of coordination, stability, and scalability. Instructed classical, methods like consensus-based control, adaptive control, and distributed model predictive control (MPC), have proved to be useful in coordinating various agents and have structured communication assumptions. Such methods are common in networked control systems because they give theoretical guarantees as well as being interpretable. Nevertheless, their operation is closely tied to performance in precise system models and predetermined structures of interactions and restricts their use in dynamic and interference-prone systems. Uncertainty, variability of coupling, and communication constraint are other factors present in large-scale intelligent infrastructures and they significantly weaken purely model-based distributed control strategies.

Similar to the embodiment of control theory, intelligent environment research has been personally directing towards context-aware control, in which decisions are modified on the basis of context beyond the environmental, functional, or network-level. The context modelling of sensed information has been utilised to enhance energy consumption, flexibility and system life of wireless sensor and ubiquitous computing networks [1], [3], [7]. Sensors based adaptation and cross-layer optimization techniques have been suggested to respond to the changing environment like changes in traffic load, interference, or node failures [4], [5], [9]. Although such methods are more responsive, they are mainly dependent on heuristic tasks or constant services of adaptation and do not explicitly treat the cooperative control between multiple actuators in a case of intense connexions.

Recent Multi-agent reinforcement learning (MARL) MARL is one of the promising directions in distributed decision-making in cyber-physical systems that allow agents to learn control policies as they act in the environment and not depending on explicit models. Current MARL methods are generally of the type of fully centralised and fully decentralised learning schemes. Centralised learning provides better convergence speed and has the disadvantageous aspect of scalability and bottlenecks in communication unlike the decentralised learning which provides better scale at the expense of non-stationary and instability. An alternative paradigm known as the Centralised Training and Decentralised Execution (CTDE) has been suggested to strike a balance

between them in order to exploit global information during training and retain local decision-making during execution. In the literature, CTDE-based MARL has demonstrated great capacity of operating in a complex environment through cooperation control but there is little to no context-aware application to the interference-coupled, cyber-physical actuation.

A comparative overview of the representative previous work of distributed control, context-aware systems and learning-based coordination is represented in (Table 1). As pointed out in the table, the currently elected methodologies lack cooperation ability in

learning, or are not mindful of the dynamics of contextual interference, or adhere to centralised implementation that limits scalability. Specifically, particular lack of frameworks, which collectively incorporate context awareness, cooperative multi-agent learning, and decentralised execution to distributed actuation platforms that are subject to dynamic coupling, is evident. It is this gap that drives the current work where a context-based cooperative control framework utilising CTDE-MARL is presented to clearly deal with interference-conscious coordination and scalable implementation in ubiquitous intelligent systems.

Table 1. Comparative Summary of Related Work

Approach Category	Control Paradigm	Context Awareness	Cooperation Level	Scalability
Consensus / Distributed MPC	Model-based, distributed	Limited (static models)	Moderate (rule-based coordination)	Limited in large-scale CPS
Heuristic Context-Aware Control	Rule-based adaptation	Partial (predefined context)	Low (independent decisions)	Moderate
Centralized Learning-Based Control	Data-driven, centralized	High (global observation)	High	Poor (single-point bottleneck)
Decentralized MARL	Learning-based, fully decentralized	Low (local-only context)	Limited (non-cooperative learning)	High
Proposed CTDE-MARL Framework	Learning-based, cooperative	High (embedded context)	High (cooperative policies)	High (decentralized execution)

3. System Model and Problem Formulation

A distributed cyber-physical actuation platform, which is a set of such spatially distributed actuators networked together, with physical processes and communication links, is what we take to be shown in (Figure 1). The actuators have the responsibility of controlling a local physical variable which could be a force, a position, a temperature or flow although they are all connected to a global system goal. The actuators are driven in a networked environment where they have limited bandwidth of communication and depend on local sensing to monitor system states. The actuators are allowed to communicate either in peer to peer fashion or by using a light weight coordination layer which allows information exchange in learning without losing the decentralised execution. This distributed architecture is an indication of practical deployment situations in intelligent environments, centralized control is commonly not possible in intelligent environments because of scale and reliability considerations.

One of the most important features of the system under consideration is the existence of a coupling and interference between actuators due to the common physical dynamics, struggle of resources, or electromagnetic interaction. This type of coupling means one actuator may affect the performance of the other actuators which have cross-interference resulting in a loss of overall control accuracy unless effectively managed. There is an assumption that the

strength of coupling changes after time because of external conditions or working modes so that the static or independent designed controllers could not be applicable. Such a structure coupled by interference is the source of the necessity of cooperative control mechanisms, which are able to adjust to dynamical interaction patterns, as opposed to depend on fixed coordination rules.

The proposed framework takes context representation as its component of the system model to allow adaptive decision-making. Context information entails the environmental aspects in the form of loads, interference intensity, and system operational states, which contender the behaviour of the actuators and systems dynamics. Training is assumed to have a global context view which is accessible in terms of centralised aggregation so that the agents learn to learn coordinated strategies using complete system information. Each actuator is however, based on locally available context and state only, during execution, as realistic constraints of deployment. Such a difference in the observability of local and global context is probably natural in line with the paradigm of centralised training and decentralised execution, which is followed in this work.

The general control goal is to develop a cooperative policy, so that every actuator can choose actions in control which can reduce local control error and also coordinate with other actuators to suppress cross-actuator interference and deliver system-wide

stability. Besides accuracy, the control approach should be used to enhance coordination between the many actuators to prevent conflicting actuators and oscillatory behaviour due to dynamic coupling conditions. The objectives are established as a joint optimization problem, with the individual actuators being the learning agents, which aim at maximising a common long-term payoff that represents the accuracy, interference suppression and stability. This statement forms the basis on which CTDE based multi-agent reinforcement learning can be used as a means to get scalable and context-aware management in distributed cyber-physical conditions.

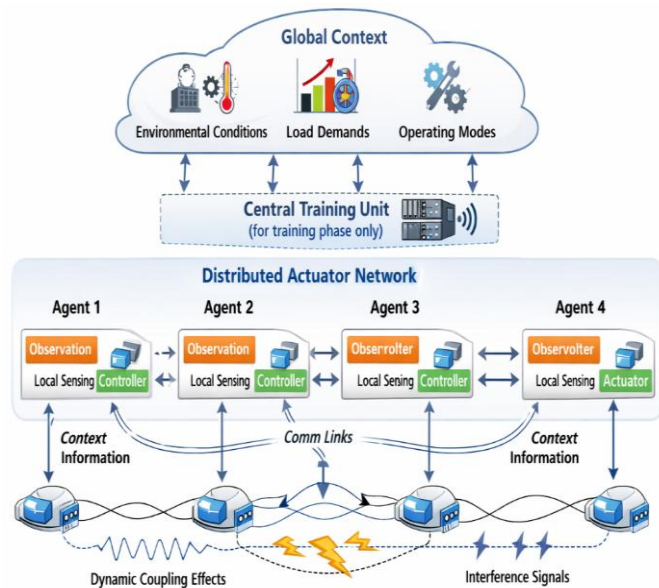


Fig. 1: System Architecture of Distributed CPS Actuation Platform

4. CTDE-Based Cooperative Intelligent Control Framework

The suggested model of cooperative control follows the Centralised Training and Decentralised Execution (CTDE) paradigm to overcome the issue specificity of distributed cyber-physical actuation. The fundamental justification of CTDE consists in the utilisation of the information available worldwide in the training phase to stabilise the learning process and provide the ability to organise it and provide a decentralised operation at the deployment stage. With the information of global context that is shown in (Figure 2), a centralized training unit becomes the source of the information containing state of actuator coupling and environmental conditions to optimize the policies. Conversely, in execution, the actuators work independently with only locally measurable conditions and contexts, which removes the overheads of runtime communication, as well as the single-point-of-failures of centralised control.

In this context, actuators are modelled as autonomous learning agents with local sensing, computational and actuation capabilities. The agent state space is developed to pull together local physical states (e.g. tracking error or actuator output) and embedded context aspects that reflect interference levels, load conditions or operating modes. This state enrichment allows the agents to reason about their dynamics as well as those that they engage in with their neighbour's actuators. The action space is the control inputs that are performed on the physical process and it can be continuous or discrete based on the actuation means of operation. The structure of the framework enables scalability to large populations of actuators by having a regular agent structure throughout the network, and does not need explicit coordination rules.

Another important feature of cooperative behaviour is an effectively designed reward exercise. The reward formulation is a collaborative approach to achieving a variety of goals: Minimising local control error to make sure some accuracy is achieved, penalising cross-actuator interference to encourage coordination, including regularisation terms to discourage excessive control effort or a sudden change of actions. The presented multi-objective reward scheme motivates agents to discover smooth and interference sensitive control policies which trade-off individual performance with system-level stability. Common reward signal in training causes agents to act in an unselfish but cooperatively instead of selfish manner, despite the implementation of execution being highly decentralised.

To learn, the Framework uses a Multi-agent reinforcement learning algorithm that is compatible with CTDE, where Multi-Agent Proximal Policy Optimization (MAPPO) was chosen based on being more stable and efficient in continuous-control cooperative tasks. Every agent has a local observation that is mapped with a local policy (actor), and a centralised value function (critic) utilises the information of global state and context information during training to minimise variation and deal with non-stationary. The training process progressively refines and trains both the actor and critic networks with experience being gathered until convergence. After training, the centralised critic is no longer utilised, and instead, the learned actor policies are deployed and allow the fully decentralised execution with the incorporation of local inference consistent with the collaborative behaviour acquired during centralised training.

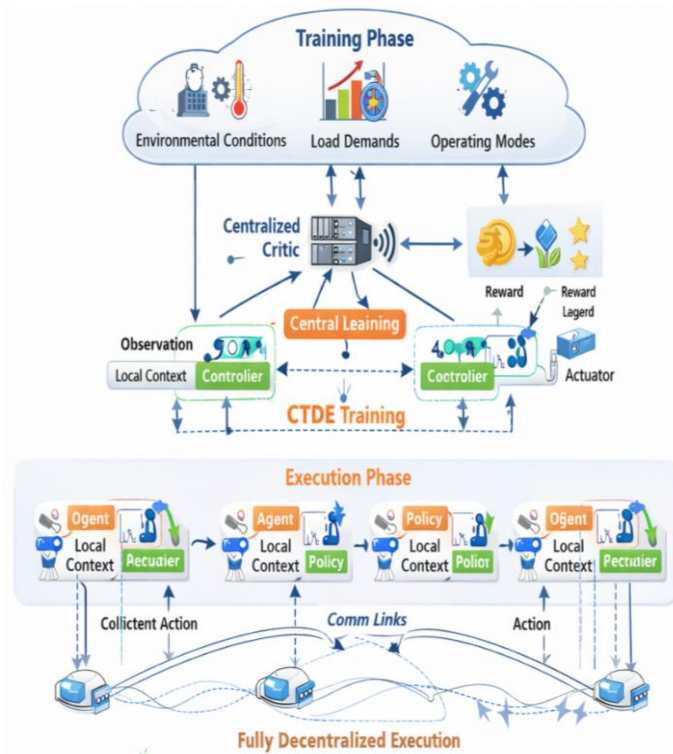


Fig. 2: CTDE Training and Decentralized Execution Workflow

5. Experimental Setup

A simulation-based assessment of the proposed CTDE-based cooperative control framework is conducted in terms of a custom simulation environment, which is a distributed cyber-physical actuation platform, dynamic in nature in terms of coupling. The cyber model of the simulated system comprises several actuators coupled by a common physical dynamics within which each actuator controls a local control variable, and by which each actuator interacts with neighbouring ones via an effect of interference. There is a variable number of actuators to investigate scalability, and where there is weak to strong interference among various coupling situations. Time-varying load conditions, as well as modes of operation, are added in order to simulate the real-world intelligent environment conditions, and to measure the adaptability of the proposed approach in the case of non-stationary conditions.

The proposed CTDE-MARL framework is compared to various representative baseline control schemes in order to obtain the fair and comprehensive comparison. The independent reinforcement learning (RL) method is used in which control policies are learned independently by each actuator by using local observations, without having direct coordination. A centralised deep reinforcement learning (DRL) scheme is also taken into consideration, where the state of a global system is observed by a single controller, and all the actuator control actions are calculated by the controller. Moreover, a consensus-based/adaptive

control method is available, as a classical distributed baseline, and it is based on coordination rules that are present or local adaptation mechanisms. A combination of these baselines represents the full range of centralised, decentralised and heuristic control strategies that are commonly used in distributed cyber-physical systems.

The evaluation of performance is made on a number of measures reflecting the quality of local control and the coordination of the system activity on the whole. The main measure that is used to determine the accuracy of the tracking and the performance of regulation in individual actuators is control error. The measures of the cooperative behaviour are the interference suppression ratio, which is a measure of how well each control scheme reduces the cross-actuator interference, at selected coupling strengths. These measurements alone give an idea of how the balance between the personal performance of the controllers and the overall objective of the system is balanced.

Besides the measures that measure accuracy, learning efficiency and robustness are measured in convergence time and stability in dynamic conditions. Convergence time measures how quickly learning-based controllers perform stable in the course of training, whereas stability analysis measures how each scheme is capable of maintaining controlled control error when coupling intensity and context variables vary suddenly. The experimental setup that evaluates accuracy, mitigating interference, efficiency in learning, and robustness together allows evaluating the proposed CTDE-based cooperative intelligent control framework thoroughly and in an unbiased manner.

6. Results, Performance Analysis, and Discussion

The proposed CTDE-based cooperative intelligent control scheme is initially tested regarding the control accuracy, where the coupling is varied, and the second case of the scheme is occurring with the weakening of interference between actuators. The control error of all schemes increases with the intensity of coupling as (Figure 3) has shown, although the suggested method has the lowest error in all cases. With moderate strength or strong degree of coordination, independent reinforcement learning is degraded quickly because of the lack of coordination, whereas centralised DRL has sluggish adaptation and responsiveness to global state variations. However, the CTDE-MARL model, in contrast, uses cooperative learning in the training process to predict the interference effect leading to a more fluent error behaviour and greater robustness to highly coupled conditions.

In addition to accuracy, cross-actuator interference can be suppressed and this is considered a vital sign of

cooperative behavior in distributed actuation systems. The quantitative outcomes summarised in (Table 2) indicate that the proposed method presents much higher ratio of the interference suppression at a far better rate than the centralised and non-cooperative baselines. Independent controllers do not help reduce interference effectively, in fact tend to enhance disturbances by conflicting actions whereas consensus-based control does not help much since coordination principles are constant. This better performance of CTDE-MARL underscores the value of policies based on learning of interference behaviour with explicit control over the dynamics of coupling based on shared reward maximisation as opposed to traditional heuristic policies or centralised coordination.

The dynamics and stability of learning on the proposed framework are also analysed by convergence of training and policy robustness analysis. Unlike other non-centralised algorithms, (Figure 3) indicates that CTDE-MARL has a stable and monotonic convergence, which is the advantage of centralised critics that decrease gradient variations and eliminates non-stationary in the training process. CRDs converge more gradually than independent RL and exhibit oscillatory behaviour because of increasing the complexity of joint action space because the latter displays unstable

learning in dynamic context variations. When deployed, the trained policies in the form of decentralised control are found to have a steady control performance despite the sudden change of coupling intensity and operating conditions, which makes the CTDE-trained policies robust during execution.

Scalability analysis demonstrates that the advantages in the performance of the proposed framework are also maintained with increasing the number of actuators. Whereas the centralised implementation of DRA faces significant performance decreases and computational costs with system size, decentralised workload (CTDE) and localised inference (MARL) are fundamental to CTDE. This scaling is achieved at the expense of increased training overhead since centralised critics must have access to global context through learning, however this trade-off is compensated by the removal of runtime communication and resultant performance of the execution. In terms of practical implementation, the findings suggest that the use of CTDE-based cooperative control may provide the possible and efficient solution to the large-scale intelligent environment, where versatility, resilience, and decentralisation must be the key characteristics.

Table 2. Quantitative Performance Comparison across Control Schemes

Performance Metric	Independent RL	Centralized DRL	Adaptive / Consensus Control	Proposed CTDE-MARL
Average Control Error ↓	High	Moderate	Moderate	Low
Interference Suppression Ratio ↑	Low	Moderate	Low-Moderate	High
Convergence Time ↓	Fast but Unstable	Slow	Fast	Moderate & Stable
Stability Under Dynamic Coupling ↑	Poor	Moderate	Moderate	High
Scalability (Actuator Count) ↑	High	Poor	Moderate	High

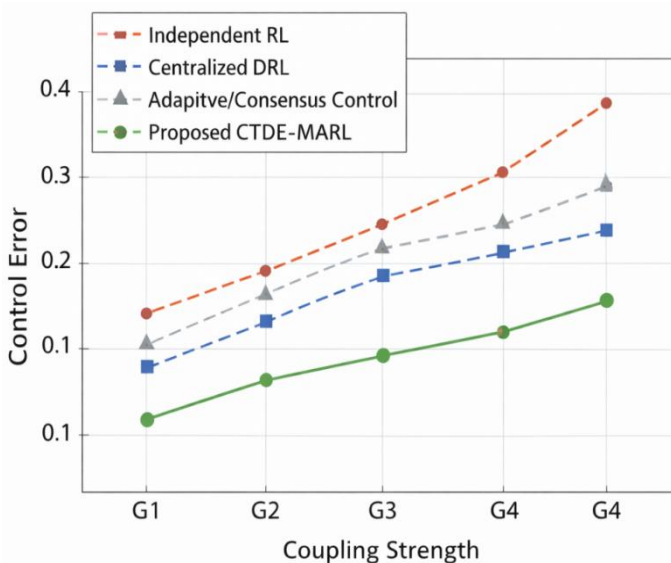


Fig. 3: Control Error vs Coupling Strength

7. Limitations and Future Work

Although the research has proven benefits of the suggested cooperative control scheme based on the CTDE, there are a range of constraints. First, the cost of training in the case of centralised training scales with the number of actuators in the network because the centralised critic needs to accept global state and context information of all the agents. Although this is a cost that is only borne during the offline training process and has no impact on runtime execution systems it can be a bottleneck with very large-scale cyber-physical systems of hundreds or even thousands of actuators. It is also a significant issue to reduce the complexity of training by sharing a parameter or using hierarchical learning structure.

The second weakness is the reliance on the correct context information during training. The proposed framework presupposes that it is possible to observe or estimate the global context in the strength of coupling

and environmental conditions in the training phase. Context measurements, in practise, can be noisy and/or delayed and partially unavailable, and this can interfere with policy generalisation when implementing a policy in a decentralised setting. The learned policies are resilient to the variations in context but when a distribution between training contexts is substantially different between the training and deployment context, then the performance of the learned policy can plummet.

Another weakness of the framework is the small depth of non-stationary coupling conditions, including sudden topological changes or the very dynamic interference patterns not adequately modelled during training. Although CTDE-MARL has the ability of eliminating non-stationary by the use of centralised critics, swift and sustained non-stationary changes in the dynamics of coupling may undermine learned policies. This restriction presents the necessity of adaptive mechanisms such that the agents may modify their policies or refreeze them on-line without disrupting system functioning.

These limitations result in several potential directions of future research. The graph-based MARL with incorporation can explicitly be used to model dynamic coupling-topology and allow agents to evolve to changing interaction-topology. The transfer learning and meta-learning methods can be used to learn faster about changing policies to different environments or actuator layouts with less overhead about the retraining. Lastly, the framework needs to be extended to hardware-in-the-loop validation to evaluate the performance in terms of real-time, safety and robustness of intelligent environment at the real implementation level, where the gap between simulation and practise is closed.

CONCLUSION

The current paper discussed a situational cooperative intelligent control in the context of distributed actuation platforms cyber-physical based on the centralised training and decentralised execution (CTDE) paradigm of multi-agent reinforcement learning. The proposed approach can effectively solve major issues of coupling, partial observability and scalability problems in distributed actuation systems by introducing contextual information into the cooperative policy learning and allowing decentralised execution. Substantial simulation findings indicated steady enhancement of the control accuracy, successful cross-actuator interference rejection, and convergence stability in any condition of dynamic coupling exercises when contrasted with centralised, independent and heuristic control schemes. These results indicate the promise of a CTDE-based cooperative learning as a scalable and robust control scheme, which has a role in the next generation of

ubiquitous and intelligent cyber-physical environments that requires adaptive, reliable, and decentralised control.

REFERENCES

1. Afroz, F., & Braun, R. (2020). Energy-efficient MAC protocols for wireless sensor networks: A survey. *International Journal of Sensor Networks*, 32(3), 150-173.
2. De Figueiredo, F. A., Jiao, X., Liu, W., Mennes, R., Jabandžić, I., & Moerman, I. (2018). A spectrum sharing framework for intelligent next generation wireless networks. *IEEE Access*, 6, 60704-60735.
3. Ding, Q., Zhu, R., Liu, H., & Ma, M. (2021). An overview of machine learning-based energy-efficient routing algorithms in wireless sensor networks. *Electronics*, 10(13), 1539.
4. Guravaiah, K., Kavitha, A., Velusamy, R. L., & Yellampalli, S. S. (2020). Data collection protocols in wireless sensor networks. *Wireless Sensor Networks-Design, Deployment and Applications*.
5. Gurewitz, O., Shifrin, M., & Dvir, E. (2022). Data gathering techniques in wsn: a cross-layer view. *Sensors*, 22(7), 2650.
6. Hussien, M., Taj-Eddin, I. A., Ahmed, M. F., Ranjha, A., Nguyen, K. K., & Cheriet, M. (2023). Evolution of MAC protocols in the machine learning decade: A comprehensive survey. *arXiv preprint arXiv:2302.13876*.
7. Lin, D., Wang, Q., Min, W., Xu, J., & Zhang, Z. (2020). A survey on energy-efficient strategies in static wireless sensor networks. *ACM Transactions on Sensor Networks (TOSN)*, 17(1), 1-48.
8. Nayak, N. K. S., & Bhattacharyya, B. (2021). Machine learning-based medium access control protocol for heterogeneous wireless networks: A review. *2021 Innovations in Power and Advanced Computing Technologies (i-PACT)*, 1-6.
9. Sadeq, A. S., Hassan, R., Sallehudin, H., Aman, A. H. M., & Ibrahim, A. H. (2022). Conceptual framework for future WSN-MAC protocol to achieve energy consumption enhancement. *Sensors*, 22(6), 2129.
10. Shukla, R., Kumar, A., & Niranjana, V. (2022, December). A Survey: Faults, Fault-tolerance & Fault Detection Techniques in WSN. In *2022 5th international conference on contemporary computing and informatics (IC3I)* (pp. 1761-1766). IEEE.
11. Wang, X., & Chen, H. (2022). A Reinforcement Learning-Based Dynamic Clustering Algorithm for Compressive Data Gathering in Wireless Sensor Networks. *Mobile Information Systems*, 2022(1), 2736734.
12. Zheng, Z., Jiang, S., Feng, R., Ge, L., & Gu, C. (2023). Survey of reinforcement-learning-based MAC protocols for wireless ad hoc networks with a MAC reference model. *Entropy*, 25(1), 101.
13. Kavuluri, H. V. R. (2025). SERVERLESS DATABASES: FUTURE TRENDS IN CLOUD DATABASE MANAGEMENT AND COST OPTIMIZATION. *Archives for Technical Sciences*, 2(33), 307-321.
14. Kavuluri, H. V. R. (2025). Enhancing Disaster Recovery with Multi-Region Data Replication in Public Sector Databases. *Indian Journal of Information Sources and Services*, 15(3), 370-380.

15. Kavuluri, H. V. R., Avula, S. B., & Sirimalla, A. (2025). Performance Tuning for Cloud-Based Databases Oracle/Postgres: Analyzing Query Optimization Techniques.
16. Kavuluri, H. V. R., & Avula, S. B. (2024). Database Engineering for Microservices: PostgreSQL as the Foundation. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 5(4), 20-28.
17. Kavuluri, H. V. R. (2024). Oracle vs PostgreSQL: Cost Analysis for Enterprises Adopting Cloud-Native Strategies. *International Journal of Emerging Research in Engineering and Technology*, 5(3), 22-31.
18. ReddyKavuluri, H. V. (2024). Advanced Role-Based Access Control Mechanisms in Oracle Databases. *International Journal of AI, BigData, Computational and Management Studies*, 5(3), 24-32.
19. Kavuluri, H. V. R. (2024). PostgreSQL vs. Oracle: A Comparative Study of Performance, Scalability, and Enterprise Adoption. *International Journal of Emerging Trends in Computer Science and Information Technology*, 5(2), 33-41.
20. Kavuluri, H. V. R., & Sirimalla, A. (2023). Enhancing Oracle database security: A deep dive into TDE, VPD, and Audit Vault implementation. *The SIJ Transactions on Computer Science Engineering & its Applications (CSEA)*, 11(1), 49-54.
21. Keshireddy, S. R., & Kavuluri, H. V. R. (2021). Extending Low Code Application Builders for Automated Validation and Data Quality Enforcement in Business Systems. *Journal of Scalable Data Engineering and Intelligent Computing*, 1(1), 7-11.
22. Keshireddy, S. R., Kavuluri, H. V. R., Mandapatti, J. K., Jagadabhi, N., & Gorumutchu, M. R. (2021). Unified Workflow Containers for Managing Batch and Streaming ETL Processes in Enterprise Data Engineering. *Transactions on Secure Communication Networks and Protocol Engineering*, 1(1), 6-10.