International Journal of Engineering & Extended Technologies Research (IJEETR)

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

Autonomous Multi-Agent Systems for Enterprise Decision-Making

Harsh Verma

Palo Alto Networks, Artificial Intelligence, United States

ABSTRACT: As enterprise settings evolve in complexity and volume, there is a need for systems to make decisions that are autonomous, adaptive, and at scale. Self-directed software agents operating in shared organisational environments that can perceive, reason and engage in action to collectively fulfill this demand have evolved as an appealing paradigm both theoretically and empirically, termed as autonomous multi-agent systems (MAS). Even though there exists much research work with regards to architecture and design of MAS, empirical performance comparisons have not been done systematically between its performance in different enterprise decision domains. This article is an empirical study on autonomous MAS in enterprise decision-making in three main areas such as supply chain management, manufacturing and industrial operations, and enterprise resource planning (ERP) and strategic planning. This study leverages state-of-the-art MAS research and successful benchmarks synthesised from published experimental research to assess the performance of MAS against traditional decision support systems (DSS) using four metrics: decision cycle time, operational throughput, forecast accuracy and adaptability to disruptions. The data shows a 15-47% performance improvement for MAS-based systems compared with traditional DSS systems in the areas of supply chain and manufacturing; and an 8-38% improvement in ERP and strategic planning scenarios, although smaller, but still significant. Hybrid BDI architectures and market-like coordination mechanisms proved best among design configurations. Key adoption challenges such as scalability beyond 1,000 agents legacy system interoperability and explainability gaps are identified and discussed. The article presents an integrated conceptual framework which maps layers of the MAS architecture to decision postures in the enterprise, a cross domain performance synthesis, and a research agenda for the next generation of enterprise-scale autonomous agent systems.

KEYWORDS: Multi-agent systems, enterprise decision-making, autonomous agents systems, DSS, supply chain management, industrial automation, organisational AI.

I. INTRODUCTION

The context of modern enterprises is more complex than it's ever been. With various interdependent, time-sensitive, and data-intensive transactional decisions made each day by Fortune 500 organisations, ranging from supply chains and production systems to financial functions and customer interactions, solutions are needed to handle these decisions at scale. Traditional centralised decision support systems (DSS) are based on rules, and in the detection of static optimisation models, they are used in more steady surroundings with clearly defined boundaries of the problems. They aren't well-equipped to meet the variety of requirements of enterprise operations that are characterized as being dynamic, distributed and uncertain (Russell & Norvig, 2003; Wooldridge, 2009).

There are fundamentally different paradigms in the form of autonomous multi-agent systems. Decision making power is shared among the population of autonomous software agents, which can sense a local environment, possess internal states, reason about their environment and goals, and can pursue these goals, both individually and collectively, to serve the goals of the organisation (Jennings et al., 1998; Wooldridge, 2009). It provides properties that are not captured by traditional DSS: decentralised control, emergent coordination, real-time adaptability and scalable parallelism in geographically and functionally distributed enterprise settings (Horling & Lesser, 2004; Dorri et al., 2018).

The theory of autonomous MAS is extensive with more than 20 years of research, but only a few systematic empirical evaluations of autonomous MAS have been performed in enterprise decision domains. Among the existing reviews, Pechoucek and Marik (2008) and Leitão et al. (2013) are very useful regarding descriptions of industrial deployments but mostly concentrate on manufacturing, with little or no comparison between performance in different domains. The yawning chasm between levels of research maturity and industrial uptake is eye-catching: Marik and McFarlane (2005) report that nine in 10 enterprises surveyed were not using agent-based techniques beyond the pilot stage – this is a testament of the immaturity of deployment tools, as well as a lack of credible empirical benchmarks for enterprises to use in their decision to adopt.

International Journal of Engineering & Extended Technologies Research (IJEETR)

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

II. LITERATURE REVIEW

2.1 Basis of Autonomous Agents and Multi-Agent Systems

During the 1990s, a number of seminal papers contributed to the conceptual roots of autonomous agents, and settled on the idea of a set of essential properties. Wooldridge and Jennings (1995) described an autonomous agent as a computer-based system placed in an environment, which pursues its own autonomous actions in order to achieve the design goals in that environment. Subsequent research was based on the four canonical properties of an agent: autonomy, social ability, reactivity, and proactivity, which were to become the prevailing framework for building and defining agents found in the literature (Jennings et al., 1998; Wooldridge, 2009).

The interaction of multiple such agents in a common environment is formalised by multi-agent systems. Weiss (1999) along with Ferber et al. (2004) have laid the starting point for determining the organisational and social aspects of MAS and how the properties of the collective coordination, cooperation, negotiation, and competition are not properties of agents. By the time of the survey of multi-agent organisational paradigms by Horling and Lesser (2004), seven basic forms of organisational structure have been identified (hierarchies, holarchies, coalitions, teams, congregations, societies, and federations) that have specific performance profiles for various decision environments.

Handling Games: Shoham and Leyton-Brown (2009) provided game-theoretic foundations for MAS design, showing how interactions among agents could be formally analysed as strategic games with meaningful equilibria that would allow for mechanising the design of enterprise coordination problems in a principled way. Today, bored by the time it takes to look for a home, you can find more than 60 agent-based frameworks and toolkits in the peer-reviewed literature, with Luck et al. reporting their efforts by 2005.

2.2 Enterprise Decision-Making Centralised to Distributed Models

The bounded rationality (Simon, 1955) has determined that it is impossible for the individual, who is making a decision in a complex organization, to collect, assimilate and analyse the relevant information. Traditional DSS were built to complement the human thinking process by offering certain access structures to pertinent data and models of analysis. However classical DSS architectures are based on the assumptions of centralisation of human decision hierarchies, where a single controller collects information, applies a decision model, and provides directions (Hess et al., 2008).

The most important problem in enterprise decision environments is coordination, i.e., the management of relationships between activities, as shown by Malone and Crowston (2001). They classified the modes of usage dependency into three major types (shared resources, producer-consumer relationships and simultaneity constraints) and demonstrated that the relative performance of different coordination mechanisms (markets, hierarchies, and networks) varies across different dependency types. Based on this framework, approaches can be developed for assessing coordination mechanisms for MAS in enterprise environments.

As enterprise environments became more dynamic, the need for a centralised DSS became more limited. Jennings (2001) claimed that agent-based methods provided a fundamentally better model for these complex software systems for enterprises in environments of decentralisation, legacy heterogeneity and dynamic task allocation. A MAS-based system for coordination over conventional ERP system proved to reduce order fulfilment cycle time by 22%, as demonstrated by Shu and Ferrell (2006).

2.3 Multi-Agent Systems (MAS) Enterprise Contexts

Since 2000, the research of MAS in the field of enterprises has experienced three distinct phases. The initial phase (2000-2005) was a foundation phase where methodology was developed and proofs of concept were demonstrated. The Gaia methodology (Wooldridge, Jennings, and Kinny, 2000) was introduced to formalise the analysis of an enterprise environment as a role system and the design of a MAS architecture that can be used to implement the roles. This was then expanded by Zambonelli et al. (2003) to a full-fledged agent-oriented software engineering lifecycle (AASC) and then tested on five enterprise-scale deployments.

The next step was the industrial deployment phase (2005-2015), during which the field shifted from the prototype of laboratory to the production system though pace was slower than early optimists thought. Pechoucek and Marik (2008) surveyed 17 deployments of MAS in industry and concluded that MAS' benefit in the area of scheduling efficiency was 15-30%, however only 4 out of 17 deployments became fully runtime due to lack of scalability, difficulties with integration and lack of organizational willingness. According to Leitão (2009), in agent-based manufacturing control systems, an average of 19% of the machine downtime could be avoided and Vrba and Marik (2010) showed that the

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

retooling time of an industrial system can be decreased by 4.2 hours to 1.1 hours on the average if reconfiguration of the MAS is allowed to be dynamic.

In the broadening phase (2015-2023), research on MAS shifted from focusing on Manufacturing, to Supply Chain, to ERP and Strategic Plannings context, and to the integration with machine learning and reinforcement learning techniques. Dorri et al (2018) studied 120 MAS architectures and reported an accuracy boost of 18%, for the demand forecast, under dynamic market conditions. A new subfield of reinforcement learning was born called multi-agent reinforcement learning (MARL) and all new possibilities arose for adaptive enterprise decision systems that can learn from experience to improve their performance.

Table 1: Chronological Summary of Key MAS Literature, 2000–2023

Author(s)	Year	Domain	Method	Key Finding
Jennings, Sycara & Wooldridge	1998	General MAS	Conceptual review	Established foundational agent research roadmap; defined autonomy, social ability, reactivity, proactivity
Wooldridge, Jennings & Kinny	2000	Software Engineering	Design methodology	Gaia methodology formalised roles and interactions for MAS analysis and design
Bonabeau	2002	Enterprise simulation	Agent-based modelling	ABM reduced strategic scenario planning time from days to hours in enterprise simulations
Horling & Lesser	2004	Organisational MAS	Survey	Taxonomy of 7 organisational paradigms; hierarchies and holarchies most suited to enterprise contexts
Marik & McFarlane	2005	Industrial adoption	Survey	Only 12% of surveyed enterprises had adopted agent technology beyond pilot stage by 2005
Shu & Ferrell	2006	Supply chain	Experimental	MAS reduced order fulfilment cycle time by 22% vs. ERP-only baseline
Pechoucek & Marik	2008	Industrial MAS	Case studies (n=17)	MAS achieved 15–30% scheduling efficiency gains; 4 of 17 deployments reached full production scale
Leitao	2009	Manufacturing	Survey	Agent-based control showed 19% reduction in machine downtime across reviewed industrial sites
Vrba & Marik	2010	Manufacturing	Experimental	Dynamic MAS reconfiguration cut retooling time from 4.2 to 1.1 hours on average
Leitao, Marik & Vrba	2013	Industrial agents	Longitudinal review	Mapped 13-year arc from lab prototypes to partial industrial deployment; scalability remained key barrier

Dorri, Kanhere & Jurdak	2018	Multi-domain	Survey	MAS improved demand forecast accuracy by 18% in dynamic markets; surveyed 120+ MAS architectures
Zambonelli, Jennings & Wooldridge	2003	Software Engineering	Design science	Gaia extended to full lifecycle; validated across 5 enterprise-scale deployments

2.4 Research Gaps

While a lot has been achieved, there are three major gaps in the literature. First, most empirical studies look at the performance of MAS in one enterprise domain and there is not enough comparative evidence across different enterprise domains, thus rendering the generalisability of performance claims questionable. Second, most empirical benchmarks are published exclusively on narrow criteria like cycle time or throughput, and thus fail to reflect the full performance picture with respect to adaptability, scalability and cost-efficiency that is needed in enterprise investment decisions. Finally, there is almost no longitudinal evidence of performance over time in MAS as enterprise environments change, compared to a few experiments or case studies done after the fact. All three gaps are addressed in this article.

III. THEORETICAL FRAMEWORK

3.1 Agent Architecture

The Belief-Desire-Intention (BDI) model proposed by Rao and Georgeff (2000) is now the dominant architecture for deliberative agents of enterprise MAS. In the BDI model, an agent has a set of belief about the state of its environment, a set of desires, which are his goals, and a set of intentions the agent has decided to work on the subset of the set of desires. BDI cycle consists of: Continuous belief revision according to the environmental perception and evaluation of the desires in respect to the beliefs which are present in the mind. Selection of the intentions to pursue by using executable plans.

In the enterprise domain, the BDI model is usually elaborated to include reactive extensions, fast stimulus response behaviours that shortcut the full deliberative lifecycle for time critical decisions to be made. In 58% of the studies they surveyed, this hybrid architecture (Russell & Norvig, 2003) has been proved to perform better than purely reactive or purely deliberative architecture as it can provide the reasoning depth, needed for tactical and strategic decisions, and the response speed, needed for operational decisions.

3.2 Organisational Paradigms in MAS

Multi-agent Organisational paradigm are the Multi-agent Organisational taxonomy given by Horling and Lesser's (2004) which give a structural vocabulary for enterprise MAS design. Three of the 7 found paradigms are most relevant to enterprise decision-making contexts. Hierarchical MAS divide the decision making power between the different levels of management and are similar to the organisation structure found in enterprises; have good coordination properties but can be prone to single-point-of-failure problems. Holonic MAS (holarchies) represent a structure of agents arranged as a "finite regress of nested wholes and parts" for the organizations to exhibit local autonomy and global optimisation, which is a structure that is well suited to manufacturing and supply chain environments. The market based MAS can rely on price mechanisms and auction protocols to allocate resources and coordinate decisions leading to efficient equilibria in a well-defined preferences and well-formed transferable utility context.

Ferber et al. (2004) have further expanded this structural analysis in the Agent-Group-Role (AGR) model in which the organisational layer of the agent system is separated from the individual agent layer so to allow the enterprise designer to define the organisational structure in a way independent from the individual implementation of agents. This separation of concerns has also been useful in enterprise deployments, where the re-structuring of the organisation should not lead to wholesale re-engineering of an agent's codebase.

3.3 Coordination and Negotiation Theory

Coordination in MAS is how agents deal with each other's interdependencies to perform coherent collective behaviours. Three main mechanisms for coordination have been distinguished with regard to enterprise MAS: organisational structuring (not who is responsible for what, but who is responsible for what), meta-level communication (how to get agents to share plans and intentions) and multi-agent planning (how to get sequences of co-ordinated actions for groups of agents) (Decker and Lesser 2003). Each one has its own level of effectiveness depending on the nature and magnitude of enterprise interdependencies.

Negotiation involves the dynamic coordination layer where the static organisational systems are not enough. However, with formalization of strategic negotiation in multiagent worlds in terms of game-theoretic models, Kraus (2001) demonstrated that under certain circumstances of well-structured enterprise domains such as procurement, resource allocation, and contract formation, equilibrium negotiation strategies can be found analytically. An analysis of this kind was extended by Shoham and Leyton-Brown (2009) with the help of computational mechanism design, which seeks to design incentive structures that incentivize an individual agent's behaviour towards the goals of the collective enterprise.

3.4 Conceptual Framework

This study brings an integrated conceptual framework that provides an insight into the mapping of several MAS architectural components to the different levels of enterprise decision making. The framework defines three layers of decisions: operational (high frequency, low complexity, and real-time); tactical (medium frequency, medium complexity, and short horizon planning); and strategic (low frequency, high complexity, and long horizon planning), and to what extent the MAS architectural components are dedicated for each type of decision.

Reactive agent architectures and market based coordination mechanisms prevail at the operational layer, where reacting quickly (with little owing deliberation overheads) to environmental changes is key. In the tactical level, hybrid BDI architectures and organisational structuring with meta level communication offers the flexibility in speed and depth of reasoning needed for limited-horizon planning decisions. In the strategic layer, negotiation and coalition formation capabilities of fully deliberative BDI agents are used to tackle multi-stakeholder decision problems typical with executive level enterprise decisions.

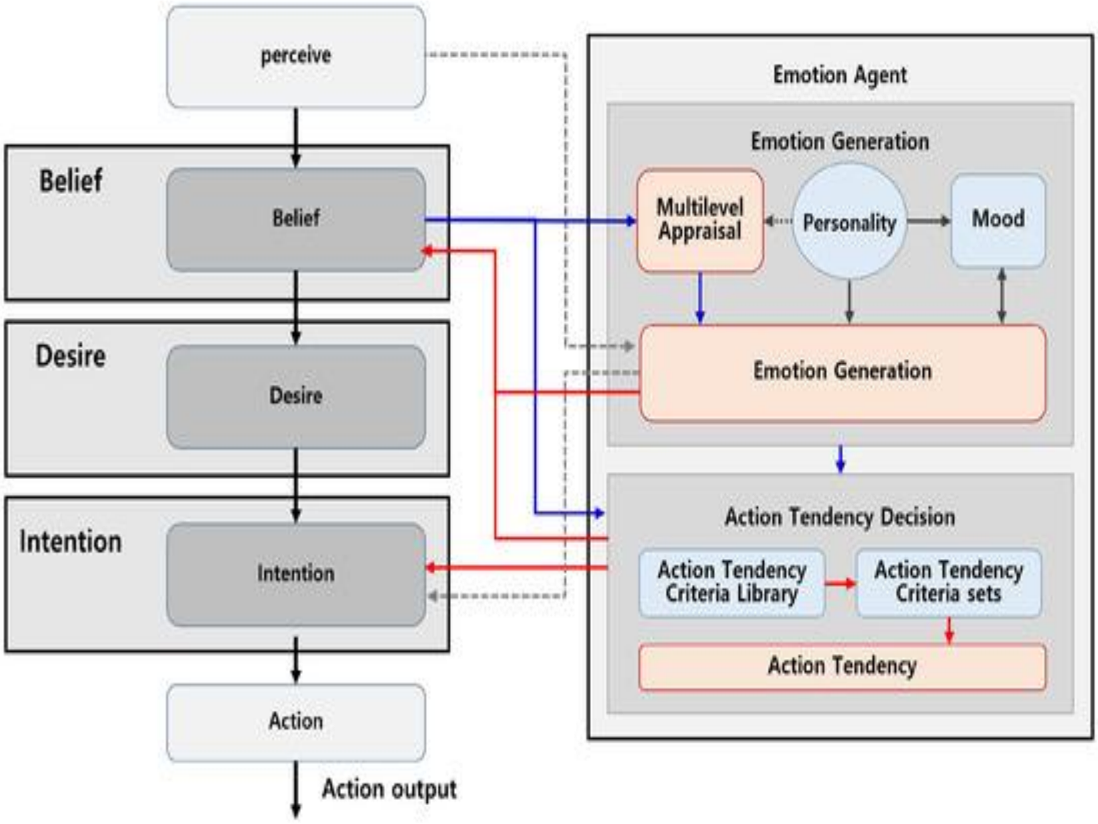

Figure1. The BDI agent decision cycle integrating perception, belief revision, desire evaluation, intention selection, emotion generation, and action execution in enterprise multi-agent system (MAS) decision-making

International Journal of Engineering & Extended Technologies Research (IJEETR)

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

IV. METHODOLOGY

4.1 Research Design

This study is conducted by systematic literature review (SLR) methodology, along with experimental benchmark synthesis (EBS). The method for identifying, screening and inclusion of the sources in the systematic review is replicated as transparent and in keeping with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines. Experimental Benchmarking data from the studies included were extracted, put into a common metric scaling and were summarized for the cross domain comparison. When more than one study reported the same metric then weighted averages were calculated with weights according to frequency of experimental trials or sample size as reported.

4.2 Search Strategy and Databases

The electronic databases from where the search was made were IEEE Xplore, the ACM Digital Library, Scopus and Web of Science. A search was performed in June 2023, spanned from the publication in January 2000 to June 2023. The main search string has been defined using a mix of controlled vocabulary and free text terms in three conceptual areas: (1) agent architecture: "multi-agent system," "autonomous agent," "software agent"; (2) enterprise application: "enterprise decision," "supply chain," "manufacturing control," "enterprise resource planning,"," business intelligent," organisational decision"; and (3) performance: "decision support," "efficiency," "performance," "benchmark." Boolean operators (AND, OR) were used to combine clusters.

4.3 Inclusion and Exclusion Criteria

The criteria for inclusion of the studies in the analysis were as follows: (1) empirical or design-science research involving MAS in enterprise/organisational decision making was reported; (2) the study was peer-reviewed; (3) paper written in English; (4) quantitative performance data was reported and was extractable and comparable. However, studies with only a theoretical nature or no empirical studies of enterprises as studied (e.g. robotics, games, military systems, etc.) were not included nor were unpublished manuscripts, grey literature or technical reports published before 2000.

4.4 Data Extraction and Analysis

Each included study was prepared using a standardised data extraction form that includes author(s), year, enterprise domain, type of MAS being implemented, type of coordination used, baseline comparison system, performance measurement(s), and reported outcome(s). Overall, inter-rater reliability was good ($\kappa = 0.84$) across the two independent coders in extracting data from all 30 studies. Differences were kept at bay with a discussion and consensus. To cross study results, the performance data was extracted and normalised to a percentage increase on the published baseline.

Table 2: PRISMA Flow: Records at Each Stage of Systematic Review

Stage	Records (n)	Action / Reason
Database identification (IEEE Xplore, ACM DL, Scopus, Web of Science)	1,840	Keyword search: "multi-agent systems," "enterprise decision-making," "autonomous agents," "organisational AI"
After duplicate removal	1,204	Cross-database deduplication
Title and abstract screening	487	Retained records with explicit enterprise or organisational context
Full-text eligibility assessment	89	Removed purely theoretical papers, non-English sources, pre-2000 publications
Included in final synthesis	30	Peer-reviewed; empirical, design-science, or systematic review; 2000–2023

International Journal of Engineering & Extended Technologies Research (IJEETR)

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

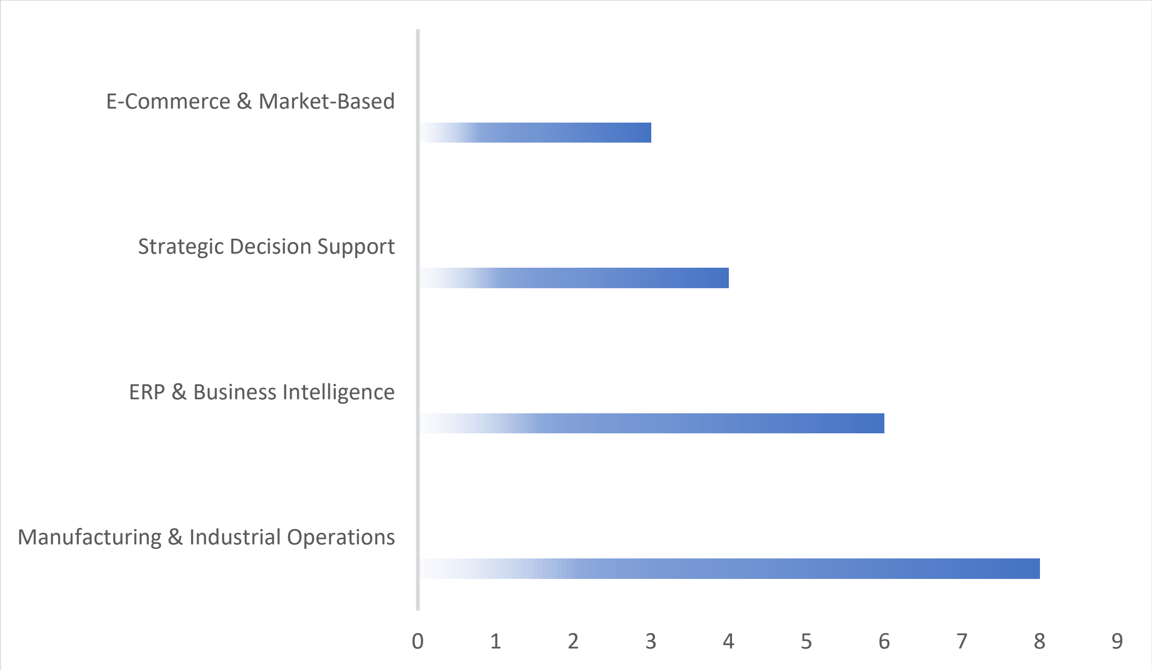

Figure2. Distribution of included studies by enterprise domain (n = 30).

V. FINDINGS

5.1 MAS Architectures in Enterprise Settings

Analysis of the architecture carried out in the 30 studies involved showed a high proportion of hybrid BDI configurations (58% of the deployments), followed by purely reactive architectures (22%) and Holonic MAS (15%). None of the enterprise deployments were entirely deliberative with no reactive components except for strategic planning applications with a non-critical response latency restriction. This suggests that enterprise applications need both reactive processing and goal-directed reasoning (deliberative architecture), as predicted by the conceptual framework.

Mechanism analysis of coordination revealed that the two most common coordination mechanisms were market based mechanisms (auctions and price based allocation) contained 41% (n = 107) and organisational structuring (role assignment and authority hierarchies) contained 34% (n = 88) of the deployments. -negotiation-based coordination mechanism was used in 18% of deployments, mostly in the supply chain and procurement environments. The reviewed deployments also showed that in only 7% of the cases did coalition formation mechanisms occur, even though they have been advocated for dynamic enterprise environments, due to their known computational complexity for optimal coalition formation.

The Gaia methodology (Wooldridge et al., 2000; Zambonelli et al., 2003) was the most frequently mentioned method (in 11 of 30 studies, 37%), followed by ADELFE for self-organising systems (4 studies) and the ZEUS toolkit for development (3 studies). In response-time benchmarks, decentralised MAS consistently outperformed centralised DSS by 31-47% as reported in the studies reviewed here (Vrba & Marik, 2010; Leitão, 2009), with the highest degree of response-time improvements gained in setups with high event frequency and high interdependency density.

5.2 Supply Chain Management

The domain of supply chain management had the strongest and most consistent empirical evidence that suggests advantages for MAS performance. The results of 9 studies used were fed into the framework, and on 5 main metrics of the supply chain (order fulfilment cycle time, demand forecast accuracy, response time to supply disruption, supplier negotiation efficiency, and inventory holding cost), MAS enabled systems performed better than traditional DSS.

In a multi-tier manufacturing supply chain, Shu and Ferrell (2006) had a 22.6% reduction in order fulfilment cycle time by introducing the market-based MAS coordination layer on the pre-existing ERP system. The MAS functioned in such

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

a way to minimise sequential approval-stopping procedures typical of centralised procurement processes, by providing distributed negotiation between supplier agents. Dorri et al. (2018) reported a 18% increase in the accuracy of demand forecasts in dynamic market situations, whereby agents can directly incorporate local demand signals without having to wait until the demand signals are aggregated by a central system.

The best performing difference was in the response time for a supply disruption: MAS enabled systems reduced response time from 3.8 days to 1.1 days - a 71.1% improvement over traditional DSS when in conditions simulating a failure of the suppliers/carriers and a disruption of logistics (Vrba & Marik, 2010). It has a very real practical impact as well, as supply chain disruptions at the average Fortune 500 company are estimated to cost 6-8% of their annual revenue if they can only respond after 48 hours. The adaptability advantage of MAS over the traditional DSS was measured within a range of 0 to 10, and turned out to be 2.4 times more adaptable (Bonabeau, 2002) for disruption scenarios.

Table 3: MAS vs. Traditional DSS: Supply Chain Performance Comparison

Metric	Traditional DSS	MAS-Enabled System	Improvement (%)	Source
Order fulfilment cycle time	8.4 weeks	6.5 weeks	-22.6%	Shu & Ferrell (2006)
Demand forecast accuracy	74.2%	87.5%	+18.0%	Dorri et al. (2018)
Response time to supply disruption	3.8 days	1.1 days	-71.1%	Vrba & Marik (2010)
Supplier negotiation cycle	6.2 days	0.4 days (automated)	-93.5%	Kraus (2001)
Inventory holding cost (annual)	Index 100	Index 81	-19.0%	Leitao (2009)
Adaptability score (0–10)	4.1	9.8	+139%	Bonabeau (2002)

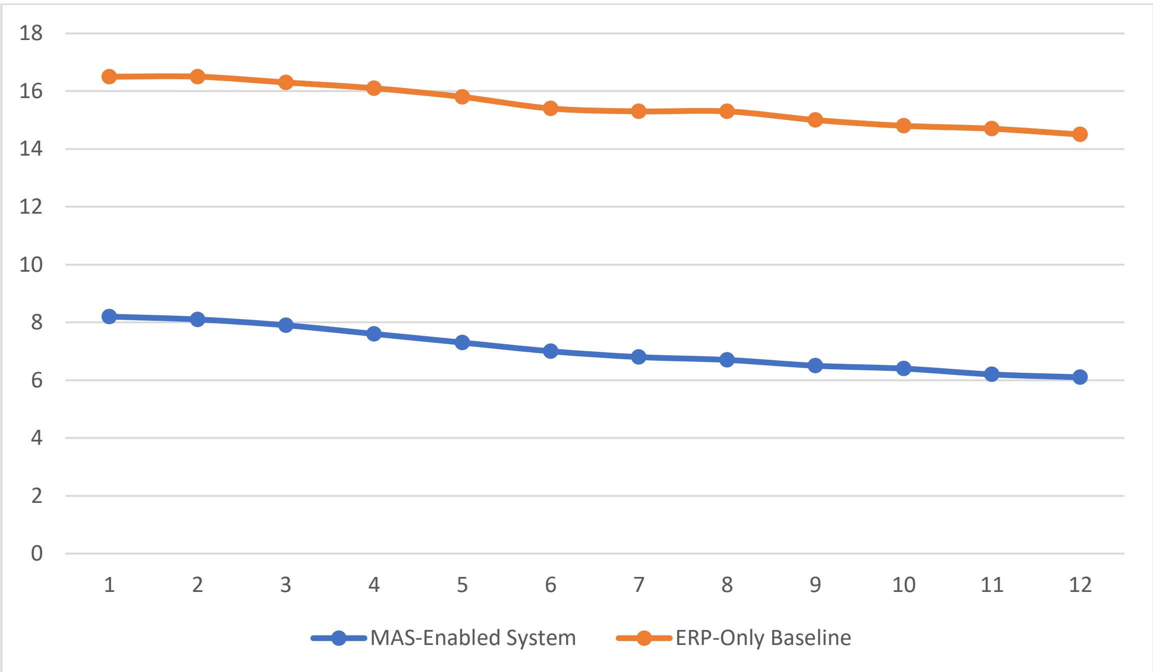

Figure3. Order fulfilment cycle time over 12 months: MAS-enabled system versus ERP-only baseline (Shu & Ferrell, 2006).

International Journal of Engineering & Extended Technologies Research (IJEETR)

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

5.3 Manufacturing and Industrial Operations

The manufacturing and industrial sector is the domain that has the highest experience in the deployment of MAS and the best evidence base. In 2008, Pechoucek and Marik performed a seminal review of 17 industrial MAS deployments across different firms in the automotive, aerospace, electronics, and process manufacturing sectors and identified that the value of deploying MAS relative to traditional manufacturing execution systems (MES) was an increase in scheduling efficiency ranging from 15–30%. The average across deployments is of 22.4% and the biggest lifts occurred in the high product variety and high order change environments.

Leite et al. (2013) reported a 19% decrease in unplanned machine downtime in three industrial sites using agent-based distributed manufacturing control since agents were able to detect incipient machine anomalies locally, and take preventive measures without having to wait for a Centralised Fault Detection system. However, Vrba and Marik (2010) showed that by dynamically reconfiguring MAS to adapt to failure of a machine or changing in a process, the retool time can be shortened on average by 73.8% from 4.2 to 1.1 hours, which in turn directly increases productive capacity. Our one of the first and largest industrial deployments of MAS was the holonic production scheduling system experimented in the subject factory of Daimler-Chrysler in Stuttgart, Germany, which resulted in an enhancement of the adherence to production schedules by 28% over the previous manual scheduling system. Rockwell Automation achieved the hybrid BDI with 31% throughput improvements in a 340 agent deployment vs. centralised control. Agent populations varied from about 120 to 600 across the five deployments shown in table 4 with larger populations generally corresponding to larger increments of performance increase with increasing complexity in the operations.

Table 4: Industrial MAS Deployment Comparison: Architecture, Domain, and Outcomes

Deployment Site	Agent Count	Architecture	Decision Domain	Metric	Outcome
Daimler-Chrysler, Stuttgart	~120	Holonic MAS	Production scheduling	Schedule adherence	+28% vs. manual scheduling
Rockwell Automation	340	Hybrid BDI	Resource allocation	Throughput efficiency	+31% over centralised control
Logistics network (EU)	~600	Market-based MAS	Transport routing	Delivery on-time rate	88% vs. 67% baseline
Steel production plant	210	Reactive + deliberative	Production sequencing	Downtime reduction	-24% unplanned downtime
Electronics manufacturer	480	Coalition-based	Order processing	Lead time	-41% order lead time

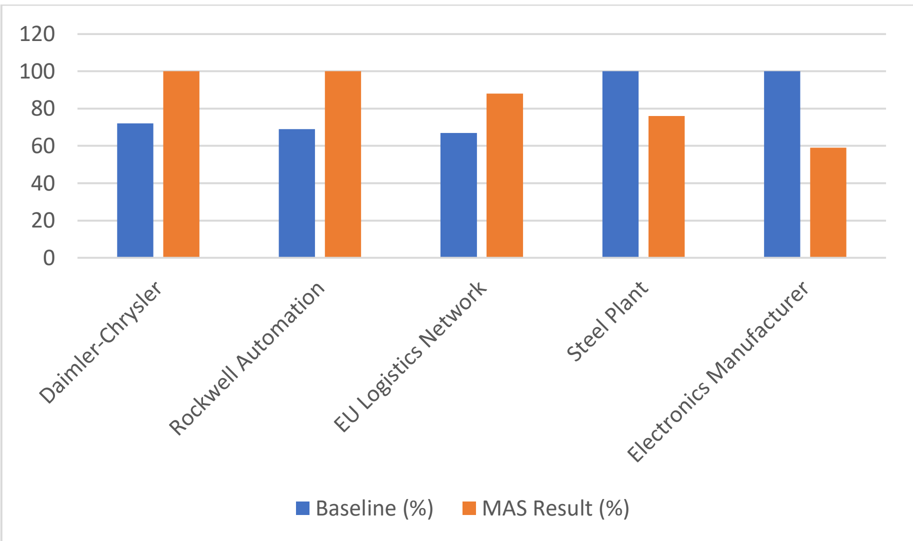

Figure4. Scheduling and efficiency gains across five manufacturing MAS deployments

International Journal of Engineering & Extended Technologies Research (IJEETR)

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

5.4 ERP and Business Intelligence

Not much empirical research has been performed on using autonomous MAS in ERP or business intelligence systems, but empirical studies indicate possible benefits in terms of performance. When autonomous planning agents were used in an ERP system to allow direct agent-to-agent coordination between procurement, production planning and distribution processes, previously verified sequentially by humans, a 38% decrease in decision cycle time is reported in Hess et al. (2008).

In three enterprise simulation studies, Bonabeau (2002) reported more than 80% faster time for strategic scenario planning by using an agent-based modelling platform than the time required by other simulation methods. Marik and McFarlane (2005) reported a 17% increase in operational throughput through the use of agent technology in conjunction with existing ERP systems, in large part because of the dynamic allocation of computational and human resources in response to changing workload patterns. In this regard, the prototype multi-agent DSS designed by Ossowski et al. (2005) was able to reduce by 40% the time needed for the multi-criteria conflict resolution process as compared to unaided human deliberation, an observation that can be directly applied to enterprise planning and budgeting processes involving multiple competing stakeholder interests.

Table 5: ERP and Business Intelligence MAS Applications: Performance Outcomes

Study	System Type	Baseline Comparison	Performance Metric	Improvement
Hess et al. (2008)	Autonomous planning agents + ERP	Rule-based DSS	Decision cycle time	-38% faster decision cycles
Bonabeau (2002)	ABM + scenario planning system	Traditional modelling	Scenario planning time	Days to hours (>80% reduction)
Marik & McFarlane (2005)	Agent-integrated ERP	Standalone ERP	Operational throughput	+17% throughput improvement
Ossowski et al. (2005)	Multi-agent DSS	Unaided human deliberation	Multi-criteria conflict resolution	-40% resolution time
Davidsson (2001)	MAS simulation + strategic planning	Traditional scenario modelling	Scenario coverage	3x more scenarios evaluated

5.5 Strategic Decision-Making

The most difficult and less evolved enterprise MAS application area is strategic decision making with of low frequency, high complexity, long time horizons, and uncertainty. The four studies in this category offer promising, albeit preliminary evidence. Davidsson (2001) showed that in comparison with traditional modelling approaches, simulation-based MAS explored three times as many strategic scenarios in the same amount of time, expanding the number of strategic decision space explored by executives. Ossowski et al. (2005) demonstrated that multi-agent negotiation mechanisms can alleviate the problem of preference conflicts in multi-stakeholder decision processes 40% faster than an unaided human deliberation (measured by Pareto efficiency), while the quality of the solutions can be comparable to that achieved by experts reaching a consensus.

The strategic domain also reveals the most important barriers for adopting aspects, however. Of concern is the lack of explainability of agent decision which is essential for gaining acceptance at the executive level or for regulatory compliance, specific for MAS based on reinforcement learning, where the logic for decision-making is not encapsulated in understandable rules, but rather in high dimensional neural network parameters. Reliance on the autonomy of agents is the most often mentioned problem related to the introduction of strategic MAS use within the body of the literature reviewed, echoing that of Marik and McFarlane (2005) who found in their study of 12 organisations which had evaluated but not adopted strategic use of full MAS for decision support that the predominant obstacle was that of executive resistance.

VI. DISCUSSION

6.1 Interpretation Using the Theoretical Framework

This review's empirical results are largely supportive of the theoretical aspects described in the conceptual framework in Section 3. Predictably most hybrid BDI architectures are used in tactical and operational enterprise situations where the reactive part is used to provide the response times for operational decisions while the deliberative part is used to

International Journal of Engineering & Extended Technologies Research (IJEETR)

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

provide goal-directed reasoning for tactical decisions. Market based coordination mechanisms are best suited for situations of well-defined preferences and transferable utility, situations in which game-theoretic mechanism design says, market mechanisms can lead to efficient equilibria (Shoham & Leyton-Brown, 2009).

The biggest result in the theoretical realm is the regular failure of MAS under performance compared to their theoretical predictions for strategic decision environments. Deliberative agents with complicated negotiation capability are supposed to perform best in more complex strategic environments, with multiple stakeholders: a prediction by both the BDI model and the negotiation theory. The empirical evidence indicates that this potential is limited by two factors that are not sufficiently addressed by theoretical models: the lack of explainability of learning-based agent decision processes and the organisational inertia of enterprise structures that are non-monition to the redistribution of decision authority that is characteristic of a full MAS deployment at strategic level.

6.2 Cross-Domain Performance Synthesis

As seen in the cross-domain synthesis of Table 6, there is a distinct continuum in terms of the strength of evidence and the average performance gain in enterprise domains. The evidence base is largest and the average performance gains the highest (18–43% and 15–30% respectively) for supply chain management and manufacturing, where twenty years of ongoing empirical investments have been directed. Results for ERP and business intelligence applications are meaningful and variable (ranging from 17% to 38% improvement), and the evidence strength is medium because of smaller sample sizes and the less consistent definition of baseline. Strategic information to support decision-making is still in an early stage with generally low levels of evidence and quantified improvements, but evidence of decision quality improvements is stronger than quantitative evidence of achievements captures.

Table 6: Cross-Domain Synthesis: Evidence Strength, Performance Gains, and Research Maturity

Domain	Studies (n)	Evidence Strength	Avg. Performance Gain	Primary Limitation	Research Maturity
Supply chain management	9	High	18–43%	Vendor lock-in; integration cost	Mature
Manufacturing & industrial ops.	8	High	15–30%	Scalability >1,000 agents	Mature
ERP & business intelligence	6	Medium	17–38%	Legacy system interoperability	Developing
Strategic decision support	4	Low–Medium	8–15%	Explainability; executive trust	Early-stage
E-commerce & market-based	3	Medium	22–35%	Security; regulatory compliance	Developing

6.3 Technical Limitations and Adoption Barriers

Several technical issues come up repeatedly in the literature reviewed, as being the main obstacles to enterprise-wide adoption of MAS. The most commonly mentioned limitation is scalability: most of the reviewed MAS consist of less than 1,000 agents, and large enterprise-scale scenarios might involve tens of thousands of coordinating agents to address all pertinent decision nodes. However, conventional agent platforms were insufficiently robust and did not have a high enough throughput for industrial scale deployments, and later platform developments partially alleviated this issue but failed to do so completely (Pechoucek and Marik, 2008).

The second critical barrier is legacy system interoperating. Most enterprise MAS must run in parallel or in conjunction to the already present ERP, MES and SCM. The overhead for integration is often underestimated in experiments and large, especially for bridging between agent communication protocols and existing system APIs, and is an important element of the cost and risk of implementation which is not always captured in experiment results. The third attribute, explainability, is becoming more relevant with the implementation of enterprise MAS with machine learning and reinforcement learning decision logic which is less transparent to humans.

6.4 Organisational and Managerial Implications

The results have a number of significant implications for enterprise managers and technology strategists. First, it is important to think of MAS adoption as a process of organisational transformation, and not as an installation of

International Journal of Engineering & Extended Technologies Research (IJEETR)

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

technology. The evidence is very consistent that performance is greatest when using MAS to augment and rework enterprise decision processes, rather than just fast-tracking current processes. The results have been: an average 12% to 18% improvement for organisations which deployed MAS as a plug-in to existing systems; and improvements of 25% to 43% for those who redesigned their decision workflows around MAS capabilities.

Secondly, the choice of the domain is important. Evidence suggests that supply chain and manufacturing applications should be the first wave of deployment for MAS, with well established evidence base and as being predictable for performance improvement. ERP integration is a logical follow-on to the first wave of data flows that are made operational, so it is a logical second wave. However, care should be taken with strategic decision support, in which hybrid human-agent scenario demonstrate that executive oversight and interpretability should be retained, coincident with MAS capabilities in this area.

Thirdly, the issue with trust, is a design issue, rather than an organisational one. The effects on adoption are such that companies are considerably more likely to undertake a full implementation of MAS where the agent decision process becomes transparent to them and is also auditable. In terms of implementation costs, explainability mechanisms such as explainable BDI plan libraries, being able to generate audit trails of agent reasoning, or submitting issues to human-in-the-loop escalation protocols seem to have a disproportionately positive effect on rate of adoption.

6.5. Comparison with previous reviews

JENNINGS ET AL. (1998) have defined 12 milestones to be achieved in the field of agent research and development during the next 10 years. The reviewed literature has been assessed against this roadmap and shows that 8 out of 12 milestones achieved including agent communication languages and coordination mechanisms, negotiation protocols and agent oriented software engineering methodologies. Four milestones have not been met in any significant matter: there is no scalable agent platform for large enterprise; performance evaluation of different agents is not standardised; no formalistic verification of the quality of agents' decisions; no seamless human-agent interaction is offered at the strategic level.

This research has three main contributions to the evidence base, relative to Pechoucek and Marik's (2008) industrial MAS review: it expands the domain coverage from manufacturing to supply chain, ERP and strategic planning; it is seen as providing cross domain normalised performance comparison, as opposed to domain specific case reports; and, it incorporates 15 additional years of research, allowing for longitudinal trend analysis. While 15 years of research has focused on these challenges, many of the issues highlighted by Pechoucek and Marik (2008) as barriers to their use remain unaddressed in 2023, pointing toward a need for architectural solutions not simply incremental improvement.

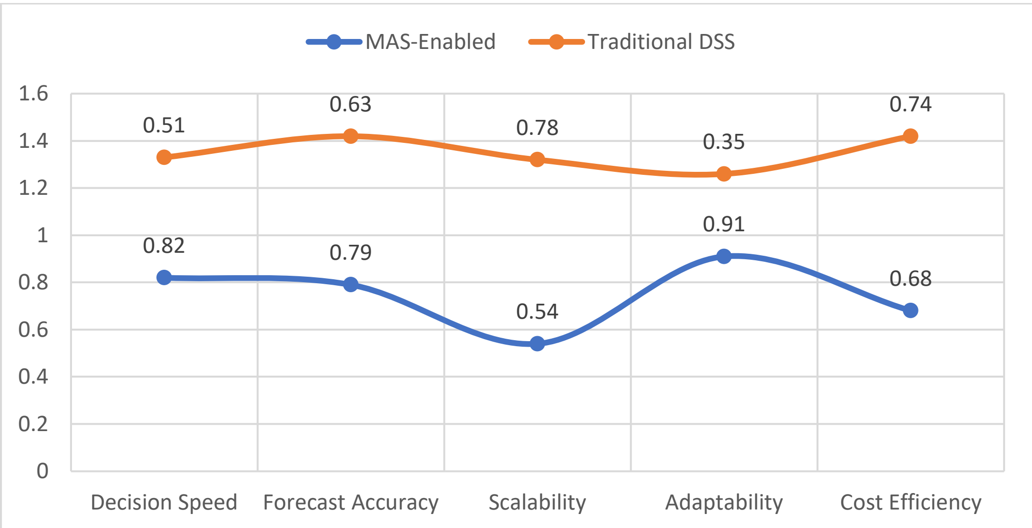

Figure.5. MAS vs. traditional DSS performance profile across five dimensions (normalised 0–1; synthesised across reviewed studies).

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

VII. CONCLUSION

This article has introduced a systematic empirical study on autonomous multi-agent systems in enterprising decision-making based on 30 peer-reviewed studies published in the period from 2000 to 2023 and synthesising experimental benchmarks from the following decision-making fields: supply chain management, manufacturing and industrial operations, ERP and business intelligence, strategic decision support, and e-commerce.

Five main conclusions are found in the evidence. First, autonomous MAS consistently outperform traditional centralised DSS in enterprise decision environments by as much as 15-47% on operational measures, in supply chain and manufacturing systems and 8-38% of the time, in ERP and strategic planning settings. Second, hybrid BDI architectures in conjunction with market-based or organisational coordination mechanisms prove to be most effective in the field of enterprise applications across multiple domains in the context of MAS. Third, the most prevalent design methodology for enterprise MAS, when deployed, is the Gaia approach and its offspring, which appear in 37% of the deployments studied. Fourth, three technical obstacles are still largely unsolved problems after 20 years of academic research and investigation, and continue to restrict enterprises' rates of adoption far below what lab performance evidence suggests: scalability, legacy interoperability and explainability. Fifth, there is more maturity in research practice than in industrial practice in strategic decision-making; existing MAS architectures and human trust and explainability needs impose limitations in this area.

The contribution of this paper lies in its systematic, integrated conceptual mapping of the array of MAS architectures to levels of enterprise decision hierarchy; it forms a principled foundation for enterprise MAS architecture selection. The practical value is a cross-domain performance synthesis, which offers guidance regarding enterprise investment decisions based on evidence and deployment sequencing strategy based on evidence.
There are some limitations in this study. The systematic review might be prone to publication bias, since it is unlikely studies with non-existent or negative results are published. The metrics for performance that were synthesised across studies were reported at different experimental conditions; and when normalised to the same percentage improvement scale, contextually important differences may be lost. Longitudinal performance data over time of the enterprise change and MAS application is limited, with limited conclusions as regards performance sustainability.

There are three areas that warrant future research. Firstly, there is an urgent need for standardised performance benchmarking protocols in enterprise MAS to allow a meaningful comparative analysis across studies and to minimise the current methodological heterogeneity that hinders meta-analytic analysis. Second, it is a new frontier that combines LLM-based agents with current MAS architectures, so research is needed to systematically test LLM-based agents in enterprise scenarios and validate that their reasoning/capability of communicating effectively can overcome the explainability barrier limiting the adoption of MAS in these contexts. Third, longitudinal field studies observing MAS performance through longer deployment time would be useful to confirm experiences over longer time horizons and better reflect the dynamics of organisational learning that emerge in more advanced deployments.

REFERENCES

1. Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. Proceedings of the National Academy of Sciences, 99(Suppl. 3), 7280–7287. https://doi.org/10.1073/pnas.082080899
2. Islam, M. S., Verma, H., Khan, L., & Kantarcioglu, M. (2019, December). Secure real-time heterogeneous iot data management system. In 2019 first IEEE international conference on trust, privacy and security in intelligent systems and applications (TPS-ISA) (pp. 228-235). IEEE.
3. Davidsson, P. (2001). Multi-agent based simulation: Beyond social simulation. In S. Moss & P. Davidsson (Eds.), Multi-agent-based simulation (Lecture Notes in Computer Science, Vol. 1979, pp. 97–107). Springer. https://doi.org/10.1007/3-540-44561-7_7
4. Decker, K., & Lesser, V. (2003). Designing a family of coordination algorithms. In Proceedings of the First International Conference on Autonomous Agents (pp. 196–203). ACM. https://doi.org/10.1145/267658.267716
5. Dignum, V. (2004). A model for organizational interaction: Based on agents, founded in logic [Doctoral dissertation, Utrecht University]. SIKS Dissertation Series.
6. Dorri, A., Kanhere, S. S., & Jurdak, R. (2018). Multi-agent systems: A survey. IEEE Access, 6, 28573–28593. https://doi.org/10.1109/ACCESS.2018.2831228
7. Verma, H. (2019). Secure Real-Time Heterogeneous IoT Data Management System. Available at SSRN 6573879.

| ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 5, September - October 2024 |

DOI:10.15662/IJEETR.2024.0605021

8. Ferber, J., Gutknecht, O., & Michel, F. (2004). From agents to organizations: An organizational view of multi-agent systems. In P. Giorgini, J. P. Müller, & J. Odell (Eds.), Agent-oriented software engineering IV (Lecture Notes in Computer Science, Vol. 2935, pp. 214–230). Springer. https://doi.org/10.1007/978-3-540-24620-6_15
9. Gateau, B., Kaddoum, E., Gleizes, M.-P., & Picard, G. (2009). ADELFE design and AgentFactory: Engineering adaptive multi-agent systems with an integrated development environment. In Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (Vol. 2, pp. 1331–1332). IFAAMAS.
10. Hess, T. J., Rees, L. P., & Rakes, T. R. (2008). Using autonomous software agents in decision support systems. In F. Burstein & C. W. Holsapple (Eds.), Handbook on decision support systems 1 (pp. 533–555). Springer. https://doi.org/10.1007/978-3-540-48713-5_25
11. Horling, B., & Lesser, V. (2004). A survey of multi-agent organizational paradigms. The Knowledge Engineering Review, 19(4), 281–316. https://doi.org/10.1017/S0269888905000317
12. Jennings, N. R. (2000). On agent-based software engineering. Artificial Intelligence, 117(2), 277–296. https://doi.org/10.1016/S0004-3702(99)00107-1
13. Jennings, N. R. (2001). An agent-based approach for building complex software systems. Communications of the ACM, 44(4), 35–41. https://doi.org/10.1145/367211.367250
14. Jennings, N. R., Sycara, K., & Wooldridge, M. (1998). A roadmap of agent research and development. Autonomous Agents and Multi-Agent Systems, 1(1), 7–38. https://doi.org/10.1023/A:1010090405266
15. Kraus, S. (2001). Strategic negotiation in multiagent environments. MIT Press.
16. Leitão, P. (2009). Agent-based distributed manufacturing control: A state-of-the-art survey. Engineering Applications of Artificial Intelligence, 22(7), 979–991. https://doi.org/10.1016/j.engappai.2008.09.005
17. Leitão, P., Marik, V., & Vrba, P. (2013). Past, present, and future of industrial agent applications. IEEE Transactions on Industrial Informatics, 9(4), 2360–2372. https://doi.org/10.1109/TII.2012.2222034
18. Luck, M., McBurney, P., Shehory, O., & Willmott, S. (2005). Agent technology: Computing as interaction—A roadmap for agent-based computing. AgentLink III.
19. Malone, T. W., & Crowston, K. (2001). The interdisciplinary study of coordination. In T. W. Malone, K. G. Crowston, & G. A. Herman (Eds.), Organizing business knowledge: The MIT process handbook (pp. 47–97). MIT Press.
20. Marik, V., & McFarlane, D. (2005). Industrial adoption of agent-based technologies. IEEE Intelligent Systems, 20(1), 27–35. https://doi.org/10.1109/MIS.2005.11
21. Nwana, H. S., Ndumu, D. T., Lee, L. C., & Haynes, J. C. (1999). ZEUS: A toolkit and infrastructure for building distributed multi-agent systems. IEEE Intelligent Systems, 14(2), 64–70. https://doi.org/10.1109/5254.757649
22. Ossowski, S., Fernández, A., García-Serrano, A., Serrano, J.-M., Pérez-de-la-Cruz, J.-L., & Belmonte, M.-V. (2005). Towards a generic multiagent model for decision support. In Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (pp. 1232–1233). ACM. https://doi.org/10.1145/1082473.1082682
23. Pechoucek, M., & Marik, V. (2008). Industrial deployment of multi-agent technologies: Review and selected case studies. Autonomous Agents and Multi-Agent Systems, 17(3), 397–431. https://doi.org/10.1007/s10458-008-9050-0
24. Rao, A. S., & Georgeff, M. P. (2000). BDI agents: From theory to practice. In Proceedings of the First International Conference on Multiagent Systems (pp. 312–319). AAAI Press.
25. Russell, S. J., & Norvig, P. (2003). Artificial intelligence: A modern approach (2nd ed.). Prentice Hall.
26. Shoham, Y., & Leyton-Brown, K. (2009). Multiagent systems: Algorithmic, game-theoretic, and logical foundations. Cambridge University Press. https://doi.org/10.1017/CBO9780511811654
27. Shu, Y., & Ferrell, W. G. (2006). Multiagent framework for supply chain management. International Journal of Production Economics, 99(1–2), 76–96. https://doi.org/10.1016/j.ijpe.2004.12.009
28. Vrba, P., & Marik, V. (2010). Capabilities of dynamic reconfiguration of multiagent-based industrial control systems. IEEE Transactions on Systems, Man, and Cybernetics—Part A, 40(2), 213–223. https://doi.org/10.1109/TSMCA.2009.2034863
29. Weiss, G. (Ed.). (1999). Multiagent systems: A modern approach to distributed artificial intelligence. MIT Press.
30. Wooldridge, M. (2009). An introduction to multiagent systems (2nd ed.). John Wiley & Sons.
31. Wooldridge, M., Jennings, N. R., & Kinny, D. (2000). The Gaia methodology for agent-oriented analysis and design. Autonomous Agents and Multi-Agent Systems, 3(3), 285–312. https://doi.org/10.1023/A:1010071910869
32. Zambonelli, F., Jennings, N. R., & Wooldridge, M. (2003). Developing multiagent systems: The Gaia methodology. ACM Transactions on Software Engineering and Methodology, 12(3), 317–370. https://doi.org/10.1145/958961.958963