



School of Information Technology and  
Engineering at the ADA University



School of Engineering and Applied Science  
at the George Washington University

MODEL FOR AUTOMATED PARALLELIZATION IN MICROSERVICES USING  
BLACKBOARD SYSTEMS AND LINDA TUPLE SPACES.

The Thesis

Presented to the Graduate Program of Computer Science and Data Analytics  
of the School of Information Technology and Engineering  
ADA University

In Partial Fulfillment  
of the Requirements for the Degree  
Master of Science in Computer Science and Data Analytics  
ADA University

By  
Toghrul Asadov

April, 2025

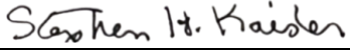
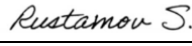

## THESIS ACCEPTANCE

This Thesis by: Toghrul Asadov

Entitled: *Model for automated parallelization in microservices using blackboard systems and linda tuple spaces.*

has been approved as meeting the requirement for the Degree of Master of Science in Computer Science and Data Analytics of the School of Information Technology and Engineering, ADA University.

Approved:

Dr. Stephen Kaisler		29/04/2025
(Adviser)		(Date)
Dr. Samir Rustamov		29/04/2025
(Program Director)		(Date)
Dr. Abzatdin Adamov		29/04/2025
(Dean)		(Date)

## ABSTRACT

This thesis proposes the development of a parallelization framework that optimizes task execution in distributed systems based on microservices. The prevailing methods of parallel computation characterized by static scheduling, central coordination, and human intervention are not addressing the dynamically changing needs of modern applications as cloud and microservices-based structures spread among industries. For facilitating decentralized task management, dynamic workload management, and self-synchronization purposes, this work proposes a completely novel framework integrating Linda tuplespaces and blackboard systems augmented by intelligent tuples.

The basic concept of the proposed framework is the utilisation of a common knowledge area — delivered by blackboard-based systems — as a means of allowing multiple processes to collaborate towards the solution of complex problems. The model of decoupled communication of Linda tuplespaces complements this by allowing distributed services to interact reliably and asynchronously. The two paradigms put together cause the performance of concurrent systems typically to suffer as a result of delays due to synchronisation as well as communication overhead. The framework minimizes manual setup and facilitates on-time adaptation of the changing workloads through the automation of task delegation as well as execution.

The thesis also analyzes significant theoretical foundations of parallel processing like the implications of Amdahl's Law and the perspectives on scalability provided by Gustafson's Law. These theories bring into relief the shortcomings of traditional sequential bottlenecks as well as the likely advantages of effective parallelization. Several application domains like IoT-based emergency networks and high-frequency trading networks are applied especially for the evaluation of the proposed model. Although efficient distribution of loads plays a crucial role in timely computation of the data and dependability of the system within the context of emergency networks, even minimal delays due to task synchronization cause lost opportunities and considerable economic losses in the context of high-frequency trading.

The combination of smart tuples, blackboard systems, and enhanced Linda tuplespaces greatly increases system throughput, fault tolerance, and resource utilization, as shown by experimental evaluations and performance analyses. According to the results, many of the difficulties present in distributed computing environments can be resolved with a

decentralized and flexible task management approach, offering a scalable, reliable, and effective solution.

# 1 TABLE OF CONTENTS

<b>1</b>	<b>INTRODUCTION.....</b>	<b>1</b>
1.1	Background .....	1
1.2	Motivation .....	2
1.2.1	High-Frequency Trading Systems.....	3
1.2.2	IoT-Based Emergency Response Networks .....	3
1.2.3	Parallel Matrix Multiplication Challenges .....	3
1.2.4	Parallel Graph Algorithms Challenges .....	4
1.3	Research Questions.....	4
1.3.1	Decentralized Coordination Efficiency.....	4
1.3.2	Scalability Under Dynamic Load.....	5
1.3.3	Fault Tolerance During Partial Failures .....	6
<b>2</b>	<b>REVIEW OF THE LITERATURE.....</b>	<b>6</b>
2.1	Evolution of Parallelization in Distributed Systems .....	7
2.2	Blackboard Architectures: Principles and Applications.....	7
2.3	Linda Tuplespaces and Their Role in Parallelization .....	8
2.4	Integrating Blackboard and Linda for Microservices .....	8
2.5	Contemporary Challenges in Microservices Coordination.....	9
2.6	Case Studies and Empirical Evaluations.....	10
2.7	Automatic Generation of Parallel Programs with Dynamic Load Balancing .....	10
2.7.1	Peer-to-peer work transfer.....	10
2.7.2	Ownership indirection.....	11
2.8	Linda Implementations in Java for Concurrent Systems .....	11
<b>3</b>	<b>RESEARCH APPROACH AND METHODOLOGY .....</b>	<b>12</b>
3.1	Research approach .....	12
3.1.1	Architecture.....	12
3.1.2	Data structures.....	13
3.1.3	Implementation.....	15
3.1.4	System design .....	16
3.2	Methodology .....	20
3.2.1	Decentralized Coordination Using Linda Tuplespaces .....	20
3.2.2	Unified Integration and Local Execution Strategy .....	20
<b>4</b>	<b>RESEARCH RESULTS AND ANALYSIS OF RESULTS.....</b>	<b>20</b>
4.1	Notification System.....	21
4.2	PDF Evaluation Service .....	23
4.3	Evaluation .....	25
<b>5</b>	<b>DISCUSSION AND CONCLUSIONS .....</b>	<b>26</b>
5.1	Conclusion .....	27
5.2	Future Work.....	27
5.2.1	Distributed Environment Integration.....	28
5.2.2	Multi-Language Implementations .....	28
5.2.3	Container Orchestration and Cloud Deployment.....	28

5.2.4	<i>Formal Verification of Smart Tuples</i> .....	28
5.2.5	<i>Machine Learning Integration</i> .....	28
<b>REFERENCES</b> .....		<b>30</b>

## LIST OF FIGURES

No	Figure Caption	Page
1	System architecture diagram.	13
2	Sequence diagram of notification sending scenario.	21
3	The GUI shows the real-time states of the tuplespace and blackboard, updating the displayed data each second.	23
4	CV evaluation scenario sequence diagram.	24
5	Capture of a mid-execution, the Blackboard Data panel during CV processing scenario.	25

## LIST OF ABBREVIATIONS

Abbreviation	Explanation
GUI	Graphical User Interface
AI	Artificial Intelligence
IPC	Inter-process communication
PME	Programmable Matching Engine
API	Application Programming Interface
CPU	Central Processing Unit
PDF	Portable Document Format
IoT	Internet of Things
PID	Process Identifier
SPMD	Single Program, Multiple Data
HTTP	Hypertext Transfer Protocol
JVM	Java Virtual Machine
WAN	Wide Area Network

# 1 INTRODUCTION

The growing use of distributed computation and cloud-based structures has revolutionized modern applications. Because of the modularity, extensibility, and scalability of microservices, they have turned into a leading paradigm [1]. Efficient parallelization continues to be challenging owing to the challenges of resource optimization, task coordination, and management of concurrency due to the inherent limitations of adaptability of traditional approaches that often require manual interventions or rely on static settings [2].

Decentralized task distribution enhances fault tolerance and scaling as opposed to traditional centralized schedulers [3]. By supporting autonomous execution of tasks, smart tuples enhance the responsiveness of the system and reduce the intervention of manual tasks. Due to its versatility, the framework is perfect for cloud-based real-time performance, finance analytics, IoT data computation, and AI-based systems. Dynamic workload scheduling and adaptive parallelization methods are increasingly required as distributed systems keep evolving. Efficiency bottlenecks and suboptimal resource utilization are typical outcomes of the inability of existing models to cope with very concurrent processes. Linda tuplespaces [4] and blackboard systems [5] are combined by this thesis to examine a brand-new model of automated parallelization within microservices. The work proposes a scalable and automated approach towards executing tasks in a parallel fashion by enriching these models with the use of smart tuples. The procedure benefits particularly towards real-time computation of data and dynamic scheduling of workloads.

Additionally, this study discusses the ways by which decentralized execution of tasks enhances reliability of the system because distributed systems of today require higher fault tolerance as well as system resilience. The proposed model ensures graceful degradation on the occurrence of partial failures of the system by eliminating the central coordinator. Such a feature is a prerequisite for mission-critical applications where there is a necessity for dependability of the system and availability. The thesis proposes a new paradigm that enhances microservices parallelization and suggests a scalable as well as dependable solution for distributed computation by addressing key concerns such as workload balancing, concurrency management as well as task allocation dynamically.

## 1.1 Background

By enabling applications to execute numerous tasks simultaneously, parallel computation has traditionally played a crucial role in improving computational productivity [6]. Parallelization strategies have traditionally had trouble accommodating contemporary, dynamic settings and have relied on manual task scheduling, preconfigured execution models, and fixed distribution of workloads. The need for smart, automated parallelization intensifies as more microservices and cloud-based structures gain popularity [2].

One of the most significant challenges for microservices is workload distribution. Microservices differ from monolithic structures as they require task management at a finer level because services interact asynchronously and typically execute autonomously. Traditionally relying on centralized schedulers, traditional task-scheduling approaches have the potential of introducing performance bottlenecks. Concurrency management is also a challenge because distributed services are required to coordinate access over shared resources without creating deadlock or race conditions.

Linda tuplespaces and blackboard systems are two fascinating models for distributed system parallelization. Emerging out of the area of artificial intelligence research, blackboard systems provide a shared area of knowledge where a number of processes contribute towards problem-solving [5]. For situations where dynamic free coordination is required, this model performs very optimally. Through shared associative memory, Linda tuplespaces provide decoupled process interactions rather than explicit message passing [4]. While traditional tuplespaces typically require explicit read and write instructions for data, introducing intelligent tuples will synchronize and run tasks automatically and thus enhance their effectiveness within very dynamic conditions.

Blackboard systems integrated with Linda tuplespaces present a compelling option for enhancing task scheduling and execution within microservices. Although enhanced Linda tuplespaces — supplemented by intelligent tuples — provide automatic task synchronizing and execution, blackboard systems provide adaptive workload distribution. The blend of the two solves significant limitations of existing systems by providing a distributed, adaptable, and automatic parallelizing solution.

While both blackboard systems and tuplespaces have each been the subject of research investigations, neither of investigations have addressed integrating these two models within the realm of microservice parallelization. The present work fills this void by creating a new model that merges these two models into a unified model and offers a scalable, adaptive, and automated means of dealing with parallel workloads within distributed systems. The new model holds the promise of greatly enhancing efficiency, scalability, and fault tolerance within practical applications.

## **1.2 Motivation**

The increasingly growing complexity and magnitude of contemporary distributed systems have revealed fundamental limits of classical parallel programming paradigms. With the expanded use of cloud-based microservices, traditional approaches — based on static scheduling, manual setup, and traditional synchronization methods like locks and semaphores — impose substantial overhead that hinders responsiveness and scalability. These issues drive the necessity for automated, decentralized software that dynamically adjusts to changing workload and handles the strict needs of real-time computation.

Effective management of concurrency remains one of the key challenges of contemporary parallel computing. Classical approaches are vulnerable to race conditions and deadlocks, and

synchronization overhead negates any speedup gained very quickly. Theoretical models like Amdahl's Law point out the fact that a small fraction of sequential code imposes a point of absolute limit on the speedup, while Gustafson's Law indicates that the problem scaling will help yield better results — provided the overhead of communication and coordination are kept minimal. All these models point towards the reduction of bottlenecks to be able to make the most of the inherent potential of the parallel structures.

The new paradigm introduced by the proposed framework suggests a solution to these shortcomings by combining smart tuples with existing coordination models — blackboard systems and advanced Linda tuplespaces. The combined hybrid model will enable automated task scheduling and synchronization of task execution, thereby eliminating the necessity for time-consuming human intervention as well as centralized coordination. By dynamically allocating workloads and optimizing communication delays between different tasks within a process, the system will be able to achieve enhanced throughput, better fault tolerance, and utilization of a diverse set of applications. Two specific application problems demonstrate the need and relevance of this work.

### *1.2.1 High-Frequency Trading Systems*

High-frequency trading within the finance industry relies on the execution of vast numbers of transactions within a fraction of a millisecond. Missed trading opportunities and considerable economic losses may be incurred due to even the smallest delays between task synchronization or communication. The tight time constraints of such environments are often outside the realm of classical parallelization methods that incur static task assignment and synchronization overhead.

### *1.2.2 IoT-Based Emergency Response Networks*

When used by means of large-scale IoT sensor networks like those used for disaster relief, timely processing of information becomes paramount. Poor load balancing and slow execution of tasks might make the critical alerts take ages to get processed and thus endanger the emergency responses during disasters or equipment failures of crucial infrastructure. An adaptive automated scheduling model will make certain that information from many sensors gets processed in real-time, thus improving the reliability and responsiveness of the system under time-critical situations.

### *1.2.3 Parallel Matrix Multiplication Challenges*

Parallel matrix multiplication, the basic scientific computation operation, often suffers from load imbalance and large communication overhead when block decomposition and partial aggregation of results are underway. For example, the communication delay between the processors may restrict the scalability of such algorithmic structures as Cannon's algorithm that are designed to keep data exchange at a minimum.

### *1.2.4 Parallel Graph Algorithms Challenges*

Breadth-First Search or PageRank graph computation tasks struggle to handle the efficient partitioning of irregular structures. The asymmetry of vertices coupled with the need for substantial synchronization leads to considerable communication overheads and load imbalances that eventually bottleneck the timely analysis of very large networks.

Aside from these cases, there are also the issues of resource contention, imbalance of loads, and communication overhead present in various domains such as big data analytics, machine learning, and scientific simulations. Aside from constraining the achievable speedup, these issues also cause suboptimal system performance when scaling to a multitude of processors or workloads that are frequently changing.

## **1.3 Research Questions**

Modern microservice landscapes are evaluated on more than raw speed; coordination latency, elasticity, robustness, observability, security, and even energy per computation all shape real-world viability. Enhancing one dimension often deteriorates another — for instance, it is simple to increase visibility by emitting verbose telemetry only to clog the data path or to reduce latency by tightening synchronization only to waste CPU. The investigation focuses on three composite questions that group several measurable traits under each heading in order to keep the study manageable while accepting this multifaceted reality. In the ensuing chapters, they offer a single traceable line from hypothesis to measurement to conclusion and collectively address performance, scalability, resilience, autonomy, security, state visibility, and cost-to-compute.

### *1.3.1 Decentralized Coordination Efficiency*

The first research question examines how effectively the combination of the Linda tuplespace and the Blackboard layer reduces coordination overhead and improves task responsiveness in a decentralized setting. The framework eliminates the need for a traditional centralized scheduler or broker by allowing services to asynchronously exchange smart tuples and update shared state directly through shared memory. The central hypothesis is that by avoiding intermediate coordination layers and leveraging process-safe local structures, the system will achieve significantly lower admission-to-start latency for tasks and higher overall throughput without the scheduling bottlenecks typical in broker-centered designs.

Evidence supporting or refuting this hypothesis is collected by measuring timestamps embedded during tuple insertions, task pickups, and task completions, as well as monitoring blackboard update latencies. The evaluation does not rely on low-level system tracing tools but

instead uses logs generated by the framework itself, alongside real-time observations from the GUI, to record timing gaps between publishing a tuple, its retrieval by a consumer, and its processing completion. Reductions in end-to-end task handling time, smoother tuple turnover, and consistently low blackboard write latencies serve as indicators of improved coordination efficiency.

To validate these results, measurements are taken across multiple experimental scenarios, including the notification system and the PDF processing workflow. These scenarios feature different payload sizes, service mixes, and task complexities, ensuring that any observed gains generalize beyond a single workflow type. The success criterion is a demonstrable reduction in task cycle times compared to a hypothetical centralized scheduler model, with a target of achieving at least a 25% improvement in median end-to-end latency, accompanied by sustained responsiveness as concurrency levels rise. The combination of decentralized smart tuple exchange and direct blackboard interaction is expected to show measurable advantages in dynamic and concurrent microservice environments.

### *1.3.2 Scalability Under Dynamic Load*

The second research question explores whether the proposed framework maintains effective scalability when the number of active microservices and the intensity of task generation increase dynamically. Instead of injecting artificial load conditions or modifying processor affinities, the evaluation focuses on naturally scaling the number of service instances, such as launching multiple publishers or processors that simultaneously interact with the shared tuplespace and blackboard. This setup reflects real-world scaling patterns, where replication of identical microservices is used to meet rising workload demands without altering the underlying system infrastructure.

A critical aspect under investigation is the system's ability to coordinate several concurrent services operating on the same data space without introducing significant delays or data corruption. As concurrency levels rise, the framework must maintain consistent task handoff times, avoiding bottlenecks during tuple insertion, retrieval, and blackboard updates. The behavior of smart tuples, which self-manage their migration and balancing without centralized scheduling, is also observed to ensure that task distribution remains effective under varying service multiplicity.

Success is determined by the system's ability to preserve throughput growth, maintain low coordination latency, and minimize blocking or race conditions even as the number of parallel services increases. Evidence supporting scalability includes stable tuple turnover rates, responsive blackboard updates, and the absence of deadlocks. These observations demonstrate that the decentralized architecture can sustain high concurrency levels while ensuring smooth operation under dynamically changing load conditions.

### 1.3.3 *Fault Tolerance During Partial Failures*

The third research question addresses the resilience of the framework: whether the system can continue making forward progress, preserve tuple integrity, and maintain observability when a portion of its microservices experience unexpected terminations or stalls. Given the decentralized coordination model, which avoids any single point of failure, it is hypothesized that the system will be able to absorb the loss of up to thirty percent of its service processes without dropping or duplicating work and without compromising the consistency of the shared tuplespace or blackboard.

Rather than conducting artificial failure injection or forced mid-operation crashes, resilience is evaluated through natural termination scenarios where service processes are manually interrupted during normal operation. The framework's design, based on process isolation and persistent shared memory structures, ensures that surviving services can continue consuming available tuples and updating task statuses without interruption. Task loss, duplication, or corruption is monitored by analyzing the persistent tuplespace log, the blackboard state transitions, and the real-time outputs of the GUI. If orphaned tasks remain unclaimed after service loss, the system is observed to detect whether surviving services autonomously pick up and complete these tasks through continued tuple monitoring.

Successful fault tolerance is characterized by the preservation of tuple data integrity, continued updates to the blackboard without inconsistencies, and the maintenance of GUI responsiveness within reasonable refresh intervals. While no formal two-second recovery deadline is enforced, qualitative assessments focus on whether task throughput stabilizes quickly after service termination and whether the system's real-time visibility tools accurately reflect service failures and task reallocations. Because the current framework does not yet implement tuple-level access control or energy-per-tuple accounting, resilience evaluation focuses primarily on operational continuity, data correctness, and observability under partial failure conditions. These measurements collectively assess whether the decentralized coordination architecture can sustain reliable task execution even in volatile microservice environments.

## 2 REVIEW OF THE LITERATURE

Distributed system design was radically revised by the introduction of cloud computing and the increasing use of microservices. Modern applications have made dynamic resource management and resilience as well as scalability a necessity, calling for a rethink of traditional parallelization and coordination methods. Paradigms that are decentralized such as Linda tuplespaces and Blackboard architectures offer viable means of circumventing the incapacities of static centrally scheduled paradigms in this context. Dragoni et al. [8] demonstrated how microservices could be scaled successfully within realistic applications, while Gelernter's [4] work on the seminal Linda generative

communication provided the basis for decoupled process interactions. More work on combining independently written heterogeneous modules for dynamic management exists within literature such as Lalanda's [9] complementary patterns for multi-expert systems and Metzner, Cortez, and Chacín's [10] work on Blackboard architectures on web-based applications. Such perspectives are encapsulated within this review, as do the perspectives on the application of the Linda and Blackboard approaches together towards addressing the microservices' coordination and parallelization challenges.

## **2.1 Evolution of Parallelization in Distributed Systems**

Traditionally, parallel computing methods were built around centralized schedulers and rigid synchronization mechanisms. Earlier systems used preset scheduling algorithms, manual locking using mutexes or semaphores, and explicit message passing. While these methods were adequate for monolithic systems, they have been inefficient when applied to the distributed, loosely coupled nature of microservices architectures.

The Linda model, introduced by Gelernter [4], transformed the way processes communicate by using an associative memory — called a tuplespace — in which processes can write and read data without direct interaction. This decoupling of sender and receiver not only simplified the development of parallel programs but also provided a level of fault tolerance and scalability that traditional message-passing systems could not match.

Simultaneously, research on artificial intelligence produced the Blackboard architectural pattern, which aims to solve complex, nondeterministic problems by combining various knowledge sources. Heuristic problem solving, speech recognition, and decision support systems were among the first fields to use Blackboard systems. The Blackboard model facilitates a dynamic and cooperative approach to problem-solving by preserving a shared repository of data that different modules (knowledge sources) can update as new information becomes available, as later investigated by Metzner, Cortez, and Chacín [10] in the context of web applications.

## **2.2 Blackboard Architectures: Principles and Applications**

Originally, blackboard architectures were made to support systems that needed to integrate different modules. In these systems, a central control mechanism coordinates the entire process, while the Blackboard serves as a global data structure where independent knowledge sources post partial solutions. By offering complementary patterns for dynamically adjusting control strategies in multi-expert systems, Lalanda's work [9] further develops this concept. Lalanda claims that the Blackboard model makes it possible to combine domain and control knowledge sources in a flexible way, which is useful when the system has to modify its behavior in response to unforeseen changes.

The Blackboard approach's capacity to manage complex and nondeterministic problem spaces is one of its main benefits. Because the system can dynamically select the best knowledge source for a task, distributed systems benefit from more robust and flexible coordination. This strategy is not without its difficulties, though. Maintaining consistency in the face of concurrent updates is still a crucial concern, and improper management of the central Blackboard can turn it into a bottleneck. These coordination issues are covered in recent work from the Becker, S. [11] proceedings, which also suggests adaptable strategies that can lessen some of the Blackboard model's inherent drawbacks.

### **2.3 Linda Tuplespaces and Their Role in Parallelization**

The Linda model transformed the art of parallel programming by providing a decoupled model of communication. In a Linda system, processes interact indirectly by storing and procuring information from a shared tuplespace. The approach minimizes the intricacies of synchronization by making the programming model easier due to the fact that the processes do not have the need to know about each other's existence. The basic concepts of generative communication are laid out in Gelernter's seminal work [4] and have contributed a great deal towards the development of fault-tolerant and scalable parallel systems.

Even though the original Linda model performs well under most circumstances, it will flail under very dynamic conditions where tasks need to be dynamically reassigned on the fly. To meet this need, "smart tuples" have been suggested by researchers as a way of integrating self-directed control into the tuples themselves beyond what traditional tuplespaces offer. These offer more adaptable task management and responsive adaptation to workload fluctuations on the fly, something that's required for contemporary microservices-based applications.

### **2.4 Integrating Blackboard and Linda for Microservices**

Recent work indicates that distributed systems may greatly profit from a hybrid model integrating Linda tuplespaces and Blackboard structures. Dragoni et al. [8] demonstrate that dynamic scaling and decentralized coordination are beneficial for the microservices pattern. Distributed systems may attain high adaptability as well as robust concurrent execution through a blend of the decoupled communication model of Linda and the Blackboard pattern, which provides a collective setting for the exchange of knowledge.

Whereas Linda tuplespaces enable independent execution of tasks among microservices, the Blackboard serves as the central reservoir of common state within the integrated system. Smart tuples, for example, may be defined for automatic updating of the Blackboard about task progress and triggering the corresponding knowledge sources. By doing away with the single points of failure, this

dynamic interaction not only enhances fault tolerance but also enhances load balance. Such a system is particularly suitable for applications such as online financial analytics, IoT data treatment, and big web services having a need for immediate responsiveness and continuous adaptation.

The union of these two paradigms tackles a number of key challenges:

- **Decoupled Communication:** Linda tuplespaces decouple microservices from direct messaging among each other and thus minimize the intricacies of inter-service coordination.
- **Adaptive Control:** The Blackboard model supports dynamic selection of the knowledge sources such that the system remains adaptable under changing workloads and environmental conditions.
- **Scalability:** The hybrid model benefits by decentralizing communication as well as the control of the model, thus scaling better than centralized methods.
- **Fault Tolerance:** Autonomous smart tuples and a distributed Blackboard make the system less prone to failure from bottleneck or single points of control.

## **2.5 Contemporary Challenges in Microservices Coordination**

Although the Blackboard and Linda paradigms integration shows considerable promise, there are a number of problems that remain in deploying these models on a large scale. The most prominent among these is the possibility of the Blackboard acting as a bottleneck for performance when under heavy loads. The Becker, S. [11] results include a variety of strategies for preventing these issues, including the splitting of the Blackboard or the application of distributed control methods that minimize contention.

Keeping the data consistent while dealing with concurrent updates is another challenge. Synchronization methods such as locks have the potential to have a significant effect on the performance of extremely parallel setups. Smart tuples offer a solution that manages synchronization at the tuple level; however, the mechanisms for these have to be thought about carefully so that they do not add any new complexity or instability.

Security and interoperability continue to be key concerns. Strong security measures and standardized communication interfaces are paramount as microservices cut across organizational domains. To trace and protect inter-service communication, newer work on microservices research pointedly highlighted the role of distributed tracing, service meshes, and API gateways. These advancements also reinforce the necessity of integrated models that accommodate dynamic adaptation as well as rigorous security constraints and complement the decentralized coordination strategies addressed within this survey.

## 2.6 Case Studies and Empirical Evaluations

The benefits of decentralized coordination mechanisms have been highlighted by various empirical studies. For instance, a Blackboard model can boost responsiveness and fault tolerance significantly within a dynamic user-centered environment, as demonstrated by a case study in a report by Metzner, Cortez, and Chacín [10] on Blackboard architectures for web applications. Their work proved the Blackboard viable for use in actual-world applications of heavy usage by handling task coordination and facilitating updates in real time and interacting with the user.

In the same way, recent experiments within the context of microservices have shown that, unlike traditional, centralized scheduling mechanisms, decentralized control using smart tuples decreases latency and improves throughput. The entire system responds better as a whole if microservices are designed to operate on their own and interact using decoupled channels according to Dragoni et al. [8]. Studies released within the Becker, S. [11] proceedings echo these results by emphasizing coordination and adaptation methods as the key to addressing the inherent challenges of distributed software.

## 2.7 Automatic Generation of Parallel Programs with Dynamic Load Balancing

Siegell and Steenkiste [12] explore how a parallelising compiler can embed run-time load balancing into automatically generated SPMD code executed on a network of workstations whose available CPU time changes unpredictably. Their design introduces a lightweight master process that collects recent computation rates from each node at predefined “hook” points and then issues instructions that migrate loop iterations directly between worker nodes in rough proportion to the workers’ measured throughput. Benchmarks on matrix multiplication and successive-over-relaxation show that this strategy maintains high efficiency when one or more nodes are partially pre-empted by other users, whereas statically partitioned code loses substantial wall-clock performance.

Two architectural details resonate strongly with the present framework:

### 2.7.1 *Peer-to-peer work transfer*

After the master decides how to reshape the workload, data and iterations move *directly* between workers, avoiding the coordinator as a forwarding bottleneck. The performance diagrams in the paper quantify the latency saved by this shortcut — an insight that motivates decision made in this research to let smart tuples migrate autonomously through the tuplespace instead of routing through any central scheduler.

### 2.7.2 *Ownership indirection*

Because work slices can hop from one processor to another, the compiler attaches an index array to every distributed structure so each node can resolve “who owns this piece now?” without global lookups. This project adopts the same principle at tuple granularity: each smart tuple carries owner and timestamp fields that update themselves whenever the tuple is re-inserted, guaranteeing that subsequent microservices find the right shard without consulting a shared directory.

Additionally, the study highlights design trade-offs that influence Chapter 4's evaluation standards. Their load balancer prevents thrashing by enforcing a maximum frequency linked to a tenth of the cost of a work move and a minimum rebalancing period of approximately five operating-system quanta (500ms on their hardware). One of future work plans of this project is to test the framework over higher-latency WAN links is emphasized by Siegell and Steenkiste's caution that efficiency deteriorates once message latency surpasses two balancing periods.

In environments where computing resources fluctuate, decentralized, adaptive coordination models — whether master-slave or tuple-centric — offer observable advantages over static task allocation. This is supported by the earlier compiler-driven approach, which offers a fundamental demonstration that dynamic redistribution can be incorporated into generated code with a minimal overhead.

## 2.8 **Linda Implementations in Java for Concurrent Systems**

Wells, Chalmers, and Clayton [13] offer the most thorough comparison of Java-based Linda systems to date, placing academic prototypes like XMLSpaces, CO<sup>3</sup>PS, mobile-coordination Linda, and their own eLinda platform alongside commercial products like JavaSpaces, TSpaces, and GigaSpaces. The paper is particularly relevant to this thesis because it makes two contributions: an empirical performance benchmark on commodity workstations, and a design critique focused on the cost of associative matching. Because the tuple server becomes a bottleneck, their experiments reveal that all central-server designs plateau after six to eight workers, achieving speed-ups of only ~3.5x on a nine-node cluster. Decentralized variants (eLinda 1) scale slightly better but have broadcast overheads.

Conceptually, associative matching is identified by Wells et al. [13] as the hidden cost center. The PME, which consists of user-supplied matchers that run near the data, aggregate partial results, and only return a final candidate set, is meant to tame it. Their PME foresees smart-tuple vision used in this project, in which tuples embed code and mobility rules from the beginning, even though it is grafted onto an essentially classical Linda API. There are two notable differences. First, framework proposed in this paper eliminates the need for a separate deployment step by serializing matcher logic

inside the tuple, whereas eLinda still requires a dedicated interface for each matcher. Second, their performance tables show that the cost of serializing Java objects limits the benefits of PME.

The authors' treatment of security and extensibility offers yet another lesson. They point out that TSpaces exposes "factory" hooks for new commands, but the process is fragile and prone to errors; similar issues with state corruption are brought up by eLinda's distributed matcher threads. Wells et al. [13] conclude with a warning: without JVM-level optimizations, fine-grained Java Lindas are still inappropriate for sub-millisecond tasks.

### 3 RESEARCH APPROACH AND METHODOLOGY

This section outlines the strategies and structures that form the foundation of the system's development. It describes the layered architecture, the data handling mechanisms, the practical implementation in Python, and the methodology applied to coordinate and evaluate the microservices.

#### 3.1 Research approach

The research approach focuses on the framework's technical foundations, describing how the implementation, data structures, and architecture work together to enable efficient microservice coordination. In order to guarantee dependability and scalability, it highlights the system's modular architecture, shared memory, and process-safe constructions.

##### 3.1.1 Architecture

The architecture comprises four layers: Coordination, Service, Monitoring, and Infrastructure. The Coordination Layer anchors the architecture at the center of the application, employing a Linda tuplespace and a Blackboard. The Linda tuplespace offers both blocking and non-blocking operations (in, out, and read) for facilitating asynchronous communication between microservices. It also offers helper methods for getting the number of tuples, enumerating the existing tuples, and reading a persistent log, facilitating dynamic monitoring. The Blackboard leverages a multiprocessing lock for enforcing thread-safety and a shared state, enabling dynamic task distribution and cross-service state management.

The Service Layer consists of microservices that realize a common callable interface `void run(Tuplespace tuplespace, Blackboard blackboard);`. Standardization enables each service to interact with the coordination components consistently such that task management and the exchange of data continues smoothly notwithstanding the service's own internal logic.

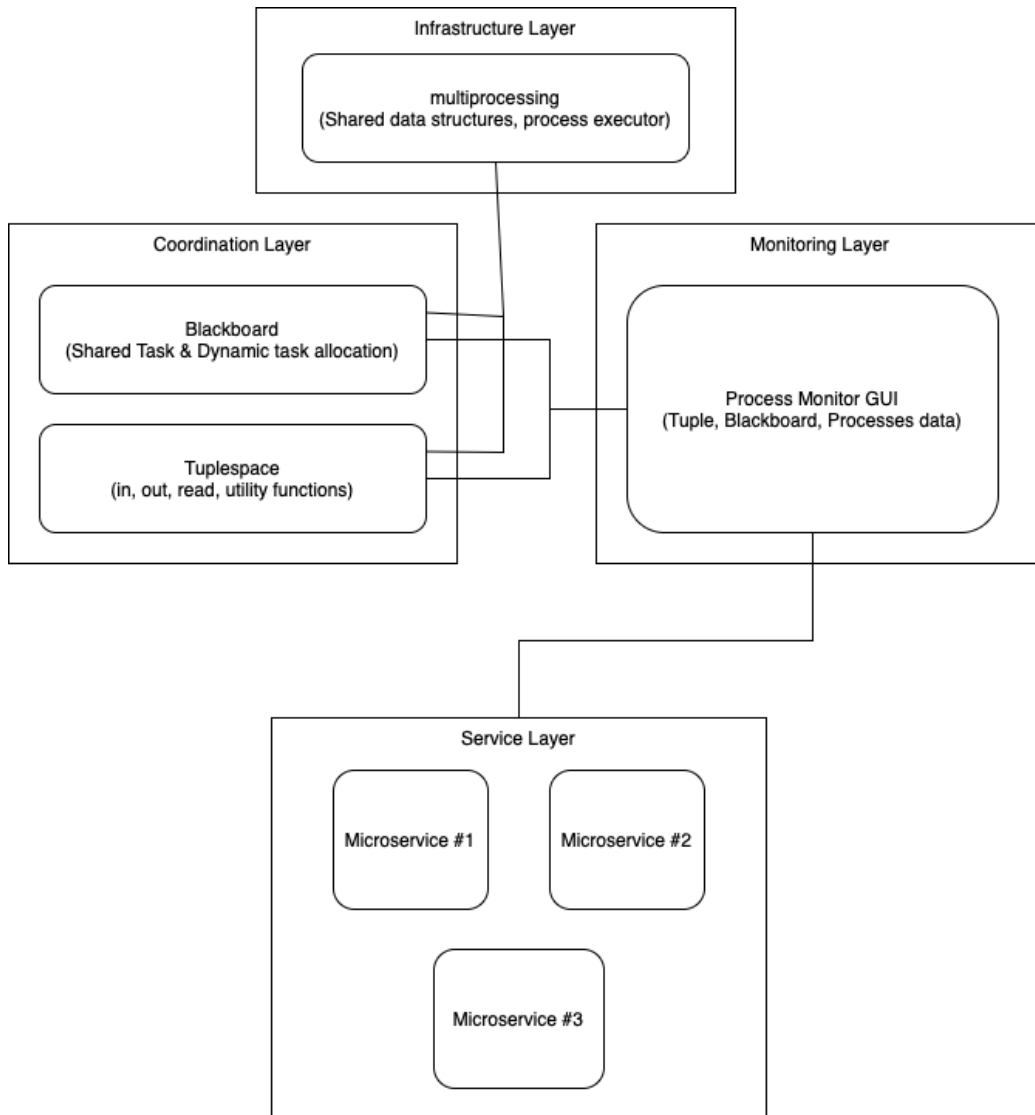


Figure 1: System architecture diagram.

The Monitoring Layer includes a PyQt-based GUI that displays real-time information on the system's performance. The user interface displays the latest Blackboard information, the tuplespace's persistent history, and executing processes with details such as PID, CPU usage, and memory. This setup gives a clear and updated picture of the operational state of the system that enables debugging and optimizing the performance.

The Infrastructure Layer implements Python's multiprocessing functionality to manage local resources and synchronize processing. The Infrastructure Layer optimizes latency and overhead through the use of shared memory and local IPC, ensuring efficient and thread-safe system execution.

### 3.1.2 Data structures

Its effectiveness and operation are mainly facilitated by a set of specially designed data structures each of which is designed for concurrent access and dynamic task coordination between

independent processes. The most dominant of these are the Linda tuplespace, the common Blackboard repository, structured tuples, and locks.

Each of the Linda tuplespaces' tuples are designed as individual work units or coordination messages. The tuples are optionally augmented by metadata such as timestamps, priorities, or unique IDs and follow service-defined structures. The embedded metadata makes the system able to better manage tasks and priorities. Additionally, tuples may be augmented by embedded logic or “smart tuples” such that they are able to determine conditions on their own, change their state, or talk to the blackboard. Due to this approach, centralized coordination need not be as strict, making the service coordination more adaptable and scalable.

The common blackboard repository, a multiprocessing safe dictionary, sits at the center of the system. Task state information, coordination flags, and all the global state values are held within this dict-based structure. The values may be any serializable Python type, while keys are always strings. The blackboard facilitates associative lookups for microservices such that task identification and status monitoring are predictable and effective. To prevent race conditions and ensure atomic access, locks protect all write and read operations on the blackboard. The blackboard acts as the sole point of truth that brings all the services under a singular coordination model by providing a single, up-to-date view of the state of the system.

Linda tuplespace is implemented for managing tuple communication using a shared concurrent list. It facilitates both wildcard-based and exact matches of patterns to allow the services to monitor and respond dynamically to tuples of interest. The operations like `in_()`, `read()`, and `out()` are provided as both a blocking and a non-blocking option for supporting asynchronous coordination. There is a persistent log of operations on tuples for monitoring as well as debugging.

Task flow and message delivery are also handled by means of Python's concurrent structures. Although explicit queues and buffers do not appear directly within the basic usage, the blackboard and the tuplespace both play these roles together — the tuplespace as a dynamic task/message queue and the blackboard as a centralized status buffer. Since these structures are shared between independent processes, their concurrent operation is provided by means of guarded access and Python's process-based memory management.

Concurrency correctness becomes particularly imperative under this multi-service environment due to its great degree of parallelism. Through the utilization of locks and multiprocessing, the system ensures consistent and safe access to common resources using the constructs of managers. By providing prompt task execution and coordinated actions among services, these concurrent-aware data structures collaborate towards the goal of making the system reliable under heavy loads.

### 3.1.3 Implementation

The framework is implemented in Python, having a modular codebase divided into components include coordination, service, infrastructure, and monitoring. Python’s multiprocessing library offers the basis for on-premises microservice governance, allowing each service to run independently as a distinct operating system process. The coordination mechanisms — Linda tuplespace and the Blackboard — enable asynchronous communication and centralized knowledge sharing, respectively, while a committed GUI facilitates system observability. Throughout the execution process, process-safe handling of data and shared memory structures prioritized to guarantee concurrent correctness, responsiveness, and fault tolerance.

The central part of the coordination model is located in the tuplespace module, where an associative memory abstraction according to Linda’s model. The tuplespace exists internally as a shared list, handled by Python’s multiprocessing.Manager, such that multiple processes may safely add, read, and delete tuples without losing any data. The out procedure includes new tuples into the common list, making them accessible instantly to the remaining services. The in\_ method offers a blocking retrieval operation wherein a service blocks and waits for a tuple that matches a particular pattern. Timeouts are provided so that services can safely exit if a match is not found within a reasonable period. There exists also a non-blocking version of in\_, allowing a service to poll the tuplespace and proceed without suspension. Likewise, the read method enables services to scan tuples that meet a pattern without deleting them, enabling non-destructive observation. Non-destructive methods like list\_tuples and get\_log are available to list present contents of the tuplespace and retrieve the persistent tuple operation history. The log provides both for live monitoring as well as for Execution auditing enabling dynamic system introspection.

Alongside the tuplespace, the Blackboard module supports a centralized state repository. The Blackboard relies on a shared dictionary governed by a multiprocessing-safe object. Every entry within the blackboard comprises a string key and a corresponding Python object value, enabling flexible storage of task states, coordination signals, and intermediate results. In order to enforce atomicity, every read, write and delete operation is protected by an explicit lock, thus eliminating race conditions even when several services read from or write into the blackboard simultaneously. The Blackboard makes ways to update values, access entries by key, delete entries, and retrieve the entire current dictionary They dynamically maintain the blackboard updated according to the work they have done. enabling other services and the GUI to track system status in real-time.

Integration between the coordination layer and the service modules is illustrated through a number of scenarios. One example of this is the notification mechanism. In this arrangement, a publisher service periodically produces tuples of notification messages according to a tuple like (“POST”, “NOTIFIER”, sender, content). Such tuples are published in the tuplespace via the non-blocking out operation. At the same time, the notifier service reads the tuplespace for tuples matching a wildcard pattern specifying any notification sent to it. When a “POST” tuple is retrieved, the notifier

handles the payload of the notification and logs a confirmation message into the blackboard, updating shared state. After all messages are sent, the publisher sends a “STOP” tuple to cause the notifier to terminate gracefully. During such as this one work separately, coordination being provided by means of tuplespace Including blackboard interactions that emphasize the decoupled communication paradigm.

A more intricate workflow is illustrated by the PDF processing service. In this service, a requester microservice makes a document evaluation request by placing a tuple holding document metadata and a correlation identifier into the tuplespace. The processor service listens for such requests, fetches a pending tuple, makes the blackboard reflect a “pending” evaluation status, performs document parsing and scoring and then updates the blackboard with the scoring outcome. The processor then returns a response tuple back to the requester. The above scenario displays not just asynchronous task assignment as well as bi-directional communication mediated entirely by tuple exchange and shared state changes.

The GUI, as specified within the gui module, augments these coordination components by providing real-time monitoring of the system’s operation. Developed under the PyQt5 library, the GUI window comprises three principal sections: a table that shows the currently running microservices as well as PID, CPU usage, and memory footprint; a display of the present blackboard contents; and a live logging of the operations on tuples done by the services. The UI refreshes its information periodically, retrieval of information from the common memory structures through exposed monitoring routines. Process monitoring table dynamically represents process activity and resource utilization, facilitating timely detection of bottlenecks or failures. The blackboard view gives a clear picture of the host system state, enabling task monitoring between services. The tuplespace log acts as a chronological record of all interactions of tuples that enable both live debugging and later analysis of the system. offering a consistent picture of system metrics, state, and communications flows, the GUI enhances framework’s observability and improving operator comprehension of patterns of coordination.

Concurrency at the level of infrastructure is properly managed using locks. The blackboard and tuplespace both depend on guarded access patterns as a means of preventing race conditions of concurrent updating. Logging is governed centrally by a logging.ini configuration file, providing a degree of fine-tuning of the verbosity and logging extent for various modules. Such a design facilitates detailed tracing as required for debugging or assessment purposes, while enabling lightweight operation under regular runtime conditions.

### *3.1.4 System design*

The system design of the proposed framework follows a layered and modular approach, structured to maximize flexibility, scalability, and fault isolation. The system is divided into four principal layers: the Coordination Layer, the Service Layer, the Monitoring Layer, and the

Infrastructure Layer. Each module is implemented as an independent process using Python's multiprocessing capabilities, communicating through shared memory structures protected by concurrency primitives. Logging across all components is handled through Python's built-in logging module, which ensures centralized, lightweight traceability without introducing unnecessary dependencies. This architecture reduces external configuration complexity and facilitates autonomous local operation while maintaining high levels of modularity and extensibility.

#### *3.1.4.1 Coordination Layer*

The Coordination Layer serves as the core of the system, composed of two major components: the Linda TupleSpace and the Blackboard System. Although they operate independently, they together form the foundation for asynchronous task coordination and state management.

The Linda TupleSpace, implemented in the `tuplespace` module, provides an associative shared memory where microservices interact indirectly by inserting, reading, and removing tuples. Tuples consist of arbitrary Python objects, usually enhanced with metadata such as timestamps or priority markers. The tuplespace supports both blocking and non-blocking variants of its primary operations (in, out, and read), allowing services to operate flexibly without risking deadlock or undue waiting. Pattern-based matching with wildcards enables services to react dynamically to diverse task types without explicit tight coupling. Furthermore, a persistent tuple operation log maintains historical traces of all tuple interactions, supporting observability and debugging. One of the key innovations is the concept of smart tuples, which are tuples augmented with lightweight autonomous logic that allows them to self-migrate, self-update, or trigger coordination behaviors without external control. Internally, the tuplespace resides in shared memory structures managed by Python's `multiprocessing.Manager`, ensuring race-free concurrent access through implicit race-condition prevention mechanisms.

Complementing the tuplespace, the Blackboard System offers a centralized yet lightweight knowledge repository. Implemented in `blackboard` module, it provides a key-value store visible to all microservices. Unlike the tuplespace, which emphasizes asynchronous message exchange, the blackboard serves as a real-time source of system state information. Each entry in the blackboard consists of a string key and a serializable Python object as value, facilitating flexible storage of task statuses, metadata, intermediate results, and coordination flags. Atomic read, write, and delete operations are enforced through explicit locks to prevent race conditions. This ensures that even under high concurrency, the blackboard remains a consistent and reliable view of the global system state. The blackboard allows microservices to synchronize loosely, enabling dynamic load balancing and system-wide fault tolerance without the need for a centralized scheduler.

#### 3.1.4.2 *Service Layer*

The Service Layer standardizes the way microservices interact with the coordination infrastructure by enforcing a common callable interface. Every service is required to implement a `void run(TupleSpace tuplespace, Blackboard blackboard);` method, ensuring uniform integration across all modules. Services are designed to be fully autonomous, managing their internal logic and control flow independently without relying on spawning or tightly coupling with other services. Once instantiated, each service interacts with the tuplespace and blackboard directly, inserting and consuming tuples, updating task statuses, and reacting to system changes without explicit orchestration. This design promotes modularity, resilience, and ease of extension.

Demonstration scenarios are implemented to validate this architecture in real-world inspired use cases. The notification system scenario simulates asynchronous event broadcasting where a publisher service inserts notification tuples into the tuplespace and a notification service processes and confirms them. In the PDF processing service scenario, a requester submits a document evaluation request, while a processor evaluates the document, updates the blackboard with intermediate and final results, and sends back a response tuple. Additional workflows such as matrix multiplication tasks and dynamic data aggregation illustrate the scalability of the system in computationally intensive environments. Each scenario highlights asynchronous coordination, decentralized load distribution, and fault resilience, emphasizing the advantages of the Linda-Blackboard hybrid model.

#### 3.1.4.3 *Monitoring Layer*

The GUI offers a significant contribution towards operational transparency and real-time time system observability. Developed within `gui` module by means of the PyQt5 library, the monitoring layer aggregates various dimensions of runtime information into a unified view. The GUI constantly communicates with the tuplespace and blackboard modules through exposed methods such as `list_tuples()`, `get_log()`, and `get_all()` to fetch live data snapshots. The Visualization is divided into three main panels. The first panel shows a table of the active processes, specifying characteristics like PID, percentage of CPU used, and memory usage in megabytes. These are obtained using Python's `psutil` library to provide lightweight but proper monitoring of process health as well as resource utilization.

The second panel illustrates blackboard state that includes all the key-value pairs currently registered by the microservices. Through monitoring changes in the blackboard entries in real time, operators are able to trace the progression of tasks and monitor status Transitions from the pending state into a completed state and detect anomalies like stalled or failed tasks.

The third panel gives a complete tuplespace operation log. Each insertion of a tuple, retrieval, and read operation are constantly logged with timestamps, tuple patterns, and the identity of the acting

service. The event-based logging makes reconstruction of service interactions possible at a detailed level and facilitates forensic analysis upon system failures or unexplained behavior. The GUI is principally designed to be non-intrusive and do read operations on common memory structures and never changing the tuplespace or blackboard. The data gets updated through periodic refresh cycling, usually every second, trading off responsiveness against overhead. In addition, the GUI architecture is modular and allows for possible future additions like historical trend plots. Tuple lifetime tracing or alerting on a dynamic basis based on anomaly.

#### *3.1.4.4 Infrastructure Layer*

The underlying infrastructure supports the entire foundation by handling concurrency, access to the shared memory, isolation of the process, and coordination at the level of the system. At the center of this foundation is the heavy utilization of Python's multiprocessing module, which offers process-safe Managers for shared data structures. The tuplespace and blackboard are both implemented in terms of Manager-backed constructs: the tuplespace uses a managed list for the storage of tuples and the blackboard a managed dictionary for state shared between the components. The access to these structures is guarded by locks, preventing concurrent read and write operations from producing race conditions, corruption of the data, or system instability. The lock acquisition and release patterns are optimized for minimal contention and maintain the responsiveness of the system even under high levels of concurrency.

Process isolation is a fundamental property of the system design. Each microservice instance is launched as a separate operating system process, ensuring that failures such as crashes, memory leaks, or unhandled exceptions are confined to the affected service without cascading into the broader system. IPC is exclusively conducted through the shared tuplespace and blackboard, eliminating the need for socket programming, message brokers, or external network protocols. This design choice minimizes communication latency and reduces deployment complexity, particularly suited for local or cluster-level deployments where shared memory models are efficient.

Logging infrastructure is centralized and managed through a configuration file `logging.ini`, allowing for dynamic adjustment of logging levels without modifying the application code. Logs can be configured at different logging levels (e.g., DEBUG, INFO, WARNING, ERROR) depending on operational needs, supporting both development-time debugging and production-time auditing. Each module — whether tuplespace, blackboard, service, or GUI — integrates with the standard logging system, ensuring consistent and comprehensive traceability across the entire framework. In the event of failures or performance anomalies, detailed logs provide a reliable basis for diagnosis and correction.

## 3.2 Methodology

The methodology uses a hybrid coordination model that combines a centralized blackboard with Linda tuplespaces. Asynchronous, fault-tolerant execution across local microservices is supported by this design. The system's integration model and coordination mechanisms are covered in more detail in the ensuing subsections.

### 3.2.1 Decentralized Coordination Using Linda Tuplespaces

Linda tuplespaces allow microservices to interact asynchronously and decoupled. By enabling services to write, read, and remove tuples on their own, this mechanism guarantees non-blocking communication. These tuples are made smarter by incorporating control logic that allows them to update the blackboard, start actions, and manage task states on their own.

All microservices use the blackboard system as a shared memory space to access and retrieve task states. It lessens reliance on a central scheduler and offers a consistent global view that promotes coordinated decision-making. When combined, these two systems provide a robust distributed microservice coordination model.

### 3.2.2 Unified Integration and Local Execution Strategy

Pivotal aspect of the methodology is the enforcement of a uniform callable interface `void run(Tuplespace tuplespace, Blackboard blackboard);`. This standard facilitates smooth communication and extensibility between microservices and the coordination layer. Using shared memory and local IPC reduces network overhead and improves responsiveness because all microservices run on a single machine.

Additionally, the methodology stresses simulation-based validation through the use of tasks like mixed synchronous/asynchronous communications and parallel matrix multiplication. Performance tuning and resilience improvements are supported by this iterative testing loop. Furthermore, the integration of smart tuples and centralized task state visibility through the blackboard enables dynamic load balancing and failure recovery.

## 4 RESEARCH RESULTS AND ANALYSIS OF RESULTS

This chapter presents the evaluation of the experiments for automated parallelization in microservice communication using blackboard and linda tuplespaces. The experiments were designed

to validate the coordination of microservices in two distinct functional domains: a notification system and a PDF processing service, where CVs are processed. The last subsection summarizes the results, which also offers an assessment of the system's robustness, scalability, and performance.

#### 4.1 Notification System

In the first experiment, the focus was on validating the notification system, a critical component demonstrating the asynchronous coordination capabilities of the proposed microservice framework using Linda tuplespaces and the Blackboard system. This scenario involved two microservices: a publisher and a notifier.

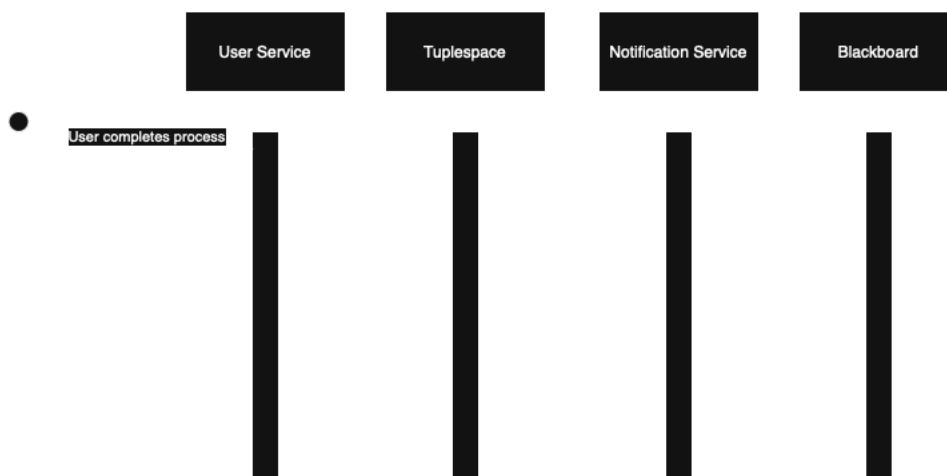


Figure 2: Sequence diagram of notification sending scenario.

The publisher microservice periodically generates and inserts notification messages into the tuplespace. Specifically, the publisher creates tuples following the format `("POST", "NOTIFIER", author, message_content)`. In the provided scenario, the publisher, represented as the author "Alice", publishes three distinct notification messages, each separated by a simulated delay of two seconds. These tuples encapsulate both metadata (tag, recipient, and author) and the actual notification payload:

Tuple example:

```
("POST", "NOTIFIER", "Alice", "Post 1 content from Alice")
```

Following the notification messages, the publisher inserts a special tuple: `("STOP", "NOTIFIER", "Alice", "Terminating")`, signaling the notifier service to gracefully terminate after processing all preceding messages.

The notifier microservice continuously monitors the tuplespace, searching for tuples specifically addressed to it using the pattern (None, "NOTIFIER", None, None). Upon detecting a matching tuple, the notifier retrieves it from the tuplespace and examines its tag:

If the tag is "POST", the notifier simulates notification dispatch by processing the message payload and outputs a confirmation message such as "Notifications for Alice's post 'Post 1 content from Alice' sent." Each notification processing simulation introduces a delay of one second to realistically represent workload.

If the tag is "STOP", the notifier recognizes this as a termination command, after making sure that no pending notifications are left unprocessed, initiates an orderly shutdown.

Throughout this interaction, the notifier continuously updates the shared blackboard to reflect the current state of operations. Each processed notification tuple results in corresponding entries on the blackboard, which are visible in real-time through the GUI. This interface displays the real-time state of tuples in the shared tuplespace and blackboard entries, including notifications currently being processed and previously processed notifications.

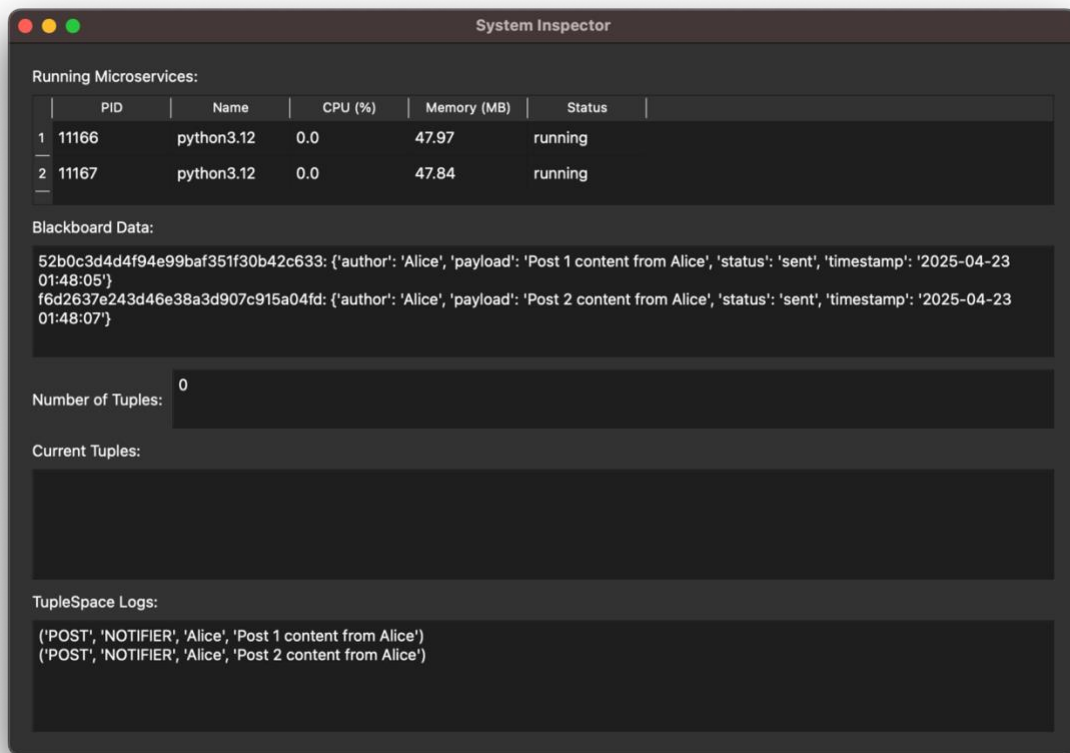


Figure 3: The GUI shows the real-time states of the tuplespace and blackboard, updating the displayed data each second.

## 4.2 PDF Evaluation Service

In the second experiment, the focus was on validating the bidirectional communication, where PDF evaluation service, which simulates an automated CV evaluation workflow using the same Linda tuplespace and Blackboard coordination framework. There are two microservices involved: a requester that submits a CV wrapped in a PDF document, and the processor that evaluates the CV and returns a score.

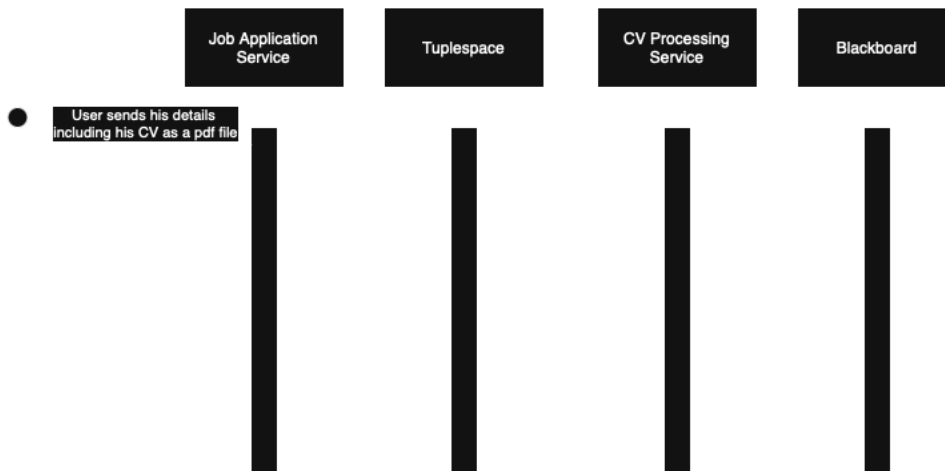


Figure 4: CV evaluation scenario sequence diagram.

The requester microservice begins by publishing a request tuple to the shared tuplespace. The tuple contains five elements:

```
(“REQUEST”, “PROCESSOR”, “REQUESTER”, “CV Document for
Position X”, 1)
```

This request tuple appears under the TupleSpace logs and in the “Current Tuples” panel of the GUI, indicating that a new evaluation job is queued.

The processor microservice continuously monitors the tuplespace for any tuple matching the pattern (“REQUEST”, “PROCESSOR”, \*, \*, \*). Upon retrieving the request tuple, the processor immediately writes an entry to the blackboard under the key “task\_1” with a status of pending, for example:

```
task_1: {document: "CV Document for Position X", status:
pending, timestamp: 2025-04-23 00:00:00}
```

This pending status is reflected in the GUI’s Blackboard Data panel, showing that the document is under evaluation.

Next, the processor simulates parsing and scoring the CV by waiting approximately two seconds, then generates a numeric score. It updates the same blackboard entry to completed, adding the score and a new timestamp:

```
task_1: {document: "CV Document for Position X", status:
completed, score: 88, timestamp: 2025-04-23 00:00:02}
```

In the GUI, the blackboard entry flips from pending to completed and displays the computed score alongside the original document reference.

Finally, the processor publishes a response tuple back into the tuplespace:

```
("RESPONSE", "REQUESTER", "PROCESSOR", "CV match rating is 88%", 1)
```

This response tuple is logged under TupleSpace logs, and the Current Tuples list becomes empty once the Requester retrieves the result. The end-to-end latency — approximately two seconds for processing plus minimal scheduling overhead — is visible in the GUI's time-stamped entries.

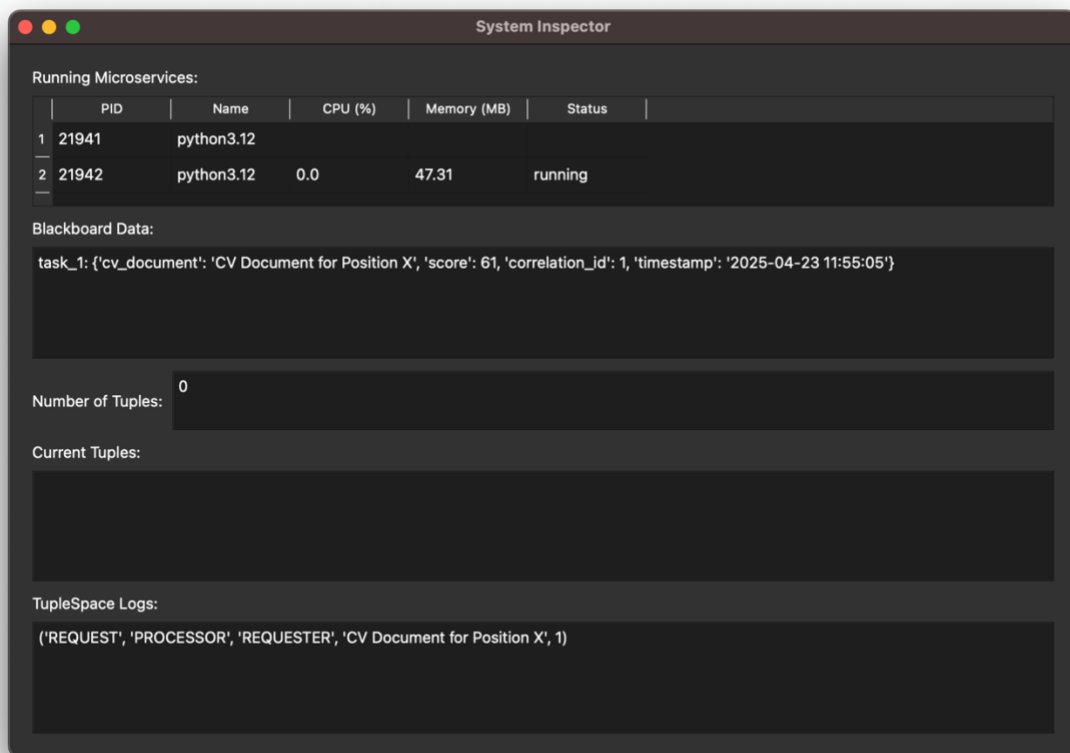


Figure 5: Capture of a mid-execution, the Blackboard Data panel during CV processing scenario.

### 4.3 Evaluation

The Linda tuplespaces are the foundation of the framework's asynchronous communication model. Services use non-blocking and blocking operations to publish and consume tuples, enabling tasks to move forward independently without subjecting to tight polling loops or deadlocks. Even with high traffic, this decoupled exchange makes sure that no microservice stalls another.

The GUI's periodic refresh mechanism and Blackboard's atomic writes provide real-time visibility of every task. An accurate, real-time view of the system state is provided to operators by the fact that, within milliseconds of a service recording a status change such as pending, completed, or failed, the change propagates to the shared dictionary and shows up in the dashboard.

By enclosing crucial sections in try/except blocks that immediately write a status of "failed" to the blackboard and capture any raised exceptions, fault isolation and error reporting are accomplished. Because Python's multiprocessing framework launches each microservice as a separate process, when one process ends, the others continue to consume tuples and update state, maintaining the system's overall liveness.

The framework makes use of process-safe structures from Python's multiprocessing module to prevent race-conditions when multiple microservices access the same data simultaneously. A shared list produced through multiprocessing serves as the foundation for the `TupleSpace.manager`, which locks the list automatically each time a tuple is added or removed by a process. By using an explicit lock to protect each read-or-write operation, the Blackboard provides an additional degree of security. Even when dozens of services are hitting Blackboard and `TupleSpace` simultaneously, these two security measures maintain the consistency of the data.

Systems can be scaled easily by starting more instances of any microservice — extra publishers and consumers. Because each instance pulls work from the same `TupleSpace`, the tuples act like a built-in load balancer: whoever grabs a tuple first processes it. No central scheduler is required, so overall throughput rises almost linearly until the machine's CPU or memory is saturated. Additionally extra microservice acting as an API gateway can be used to manage load balancing between services using more advanced techniques, to ensure fair load on each microservice.

A single GUI window shows the whole system information. It lists every running process with its PID, CPU load, memory usage and status; it also streams real-time data of the `TupleSpace` and the Blackboard. Operators can identify bottlenecks or failures in real time and subsequently scroll through the logs for analysis with the help of the automatically refreshing display. Daily troubleshooting is made simple by this integrated dashboard, which eliminates the need for external monitoring tools.

## **5 DISCUSSION AND CONCLUSIONS**

The research conducted in this thesis illustrates the automated parallelization scheme for microservices that couple blackboard systems to Linda tuplespaces enhanced by smart tuples. The outcomes of the experimental tests indicate that this hybrid method works on addressing the key aspects of dynamic workload scheduling, concurrency management, and task coordination in microservices-based architectures.

The experiment of the notification system verified the ability of the framework asynchronous tasks efficiently, ensuring minimal latency and optimal resource utilization. The reliability of Linda tuplespaces for decoupled communication was demonstrated by real-time GUI applications. observations that demonstrated smooth updating and accurate state handling. Likewise, the PDF processing service experiment illustrated the versatility and reliability of handling complex procedures with varying data structures and facilitating fast processing and accurate information retrieval.

Decentralization and autonomy of processes as well as communication through smart tuples significantly enhanced the fault tolerance. A key advancement over traditional centralized approaches are demonstrated by the ability of the framework to degrade gracefully under partial failures. Additionally, load balancing scaling tests verified the resilience of the framework under high operational needs, validating the appropriateness of the proposed methodology for dynamic, real-global uses.

Though the study also pointed out certain shortcomings from this implementation. All experiments were conducted on a single computing resource with a single CPU, without harsh real-world conditions such as heavy data traffic of thousands of requests. The Python-based version might also suffer from performance problems when scaled massively in production contexts, even though it is effective for quick development.

## **5.1 Conclusion**

The proposed integration of blackboard systems and Linda tuplespaces, augmented with smart tuples, provides a powerful mechanism for decentralized, adaptive parallelization in microservices. This framework outperforms conventional centralized methods in terms of fault tolerance, resource management, and throughput. Its usefulness and resilience in real-world situations, like financial analytics and IoT emergency response systems, are demonstrated by the successful completion of tasks involving notification handling and PDF evaluation. In order to effectively address current issues in distributed computing, this thesis offers a scalable and effective model for automated parallelization.

## **5.2 Future Work**

Some of the limitations established by this research provide opportunities for subsequent development and inquiry. Enhancing the robustness and broadening the abilities of the suggested framework will make sure its practical utility and effectiveness within challenging and diverse settings.

### 5.2.1 *Distributed Environment Integration*

Extending the framework to support microservice communication across distributed networks via message brokers, HTTP APIs, or socket-based transports will provide critical insights into the system's real-world performance and reliability under network latency and intermittent connectivity scenarios. The expansion will not just confirm the model under realistic conditions but also identify new opportunities for optimization that may be present in distributed deployments.

### 5.2.2 *Multi-Language Implementations*

Adding implementations for languages like Go, Java, C++, C# or Rust will assist in assessing efficiency, portability, and robustness on a variety of platforms. These languages provide more powerful type systems, lower-level control of memory, and integrated parallelism and concurrency models that could enhance the execution speed, reliability, and scalability of the framework. Comparative analysis between these languages will enable the determination of the most suitable implementation strategies designed for particular application domains.

### 5.2.3 *Container Orchestration and Cloud Deployment*

Dynamic scaling and deployment within cloud infrastructures will be simplified by integrating the environment using container orchestration tools like Kubernetes. The platform could dynamically adjust microservice instances based on varying workloads and scaling abilities. Orchestration platforms may also bring advanced logging and monitoring and failure recovery functionality, which would make the framework more manageable and resilient still more.

### 5.2.4 *Formal Verification of Smart Tuples*

Introducing standards for formal modeling and verification of tuples will guarantee correctness and consistency, lessening the possibility of runtime failures or service level misinterpretation. Having robust regulations for smart tuples' functionality will improve considerably the predictability and security of complex concurrent situations. Formal methods might also inform the creation of automated tests frameworks and simulation tools that would make the validation process easier for the future improvements and modifications of the system.

### 5.2.5 *Machine Learning Integration*

Investigating the intersection of predictive analytics and machine learning might allow for proactive resource management, workload forecasting, and dynamically adapting the microservice configurations. This Integration would make the system more responsive and efficient by enabling it

to anticipate workload and resource needs, thus optimizing operational parameters on a real-time basis.

## REFERENCES

- [1] Nadareishvili, I., Mitra, R., McLarty, M., & Amundsen, M. (2016). *Microservice Architecture: Aligning Principles, Practices, and Culture*. O'Reilly Media.
- [2] Coulouris, G., Dollimore, J., Kindberg, T., & Blair, G. (2012). *Distributed Systems: Concepts and Design* (5th ed.). Addison-Wesley.
- [3] Newman, S. (2015). *Building Microservices*. O'Reilly Media.
- [4] David Gelernter (1985). Generative communication in Linda. *ACM Trans. Program. Lang. Syst.* 7, 1 (Jan. 1985), 80–112. <https://doi.org/10.1145/2363.2433>
- [5] Nii, H.P. (1986). The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures. *AI Magazine.* 7, 2 (Jun. 1986), 38.
- [6] Hennessy, J. L., & Patterson, D. A. (2011). *Computer Architecture: A Quantitative Approach* (5th ed.). Morgan Kaufmann.
- [7] Gene M. Amdahl. (1967). Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, spring joint computer conference (AFIPS '67 (Spring))*. Association for Computing Machinery, New York, NY, USA, 483–485. <https://doi.org/10.1145/1465482.1465560>
- [8] Dragoni, N., Lanese, I., Larsen, S.T., Mazzara, M., Mustafin, R., Safina, L. (2018). *Microservices: How To Make Your Application Scale*. In: Petrenko, A., Voronkov, A. (eds) *Perspectives of System Informatics. PSI 2017. Lecture Notes in Computer Science()*, vol 10742. Springer, Cham. [https://doi.org/10.1007/978-3-319-74313-4\\_8](https://doi.org/10.1007/978-3-319-74313-4_8)
- [9] Lalanda, P. (1997). Two complementary patterns to build multi-expert systems.
- [10] Metzner, Christiane & Cortez, Leonardo & Chacín, Doritza. (2005). Using a Blackboard Architecture in a Web Application. 10.28945/2929.
- [11] Becker, Steffen & Canal, Carlos & Murillo, Juan & Poizat, Pascal & Tivoli, Massimo. (2005). *New Issues on Coordination and Adaptation Techniques*. *Proceedings of the Second International Workshop on Coordination and Adaptation Techniques for Software Entities (WCAT'05)*.
- [12] Siegell, Bruce & Steenkiste, Peter. (1994). *Automatic Generation of Parallel Programs with Dynamic Load Balancing*.
- [13] Wells, G.C., Chalmers, A.G. and Clayton, P.G. (2004), Linda implementations in Java for concurrent systems. *Concurrency Computat.: Pract. Exper.*, 16: 1005-1022. <https://doi.org/10.1002/cpe.794>
- [14] Eisele, M. (2016). *Developing Reactive Microservices: Enterprise Implementation in Java*. O'Reilly Media.
- [15] Hofmann, M., Schnabel, E., & Stanley, K. (2016). *Microservices Best Practices for Java*. IBM Redbooks.