



Modeling Tenant Dependencies with Formal Concept Analysis for Energy-Aware Cloud Resource Management

Journal:	<i>Software and Systems Modeling</i>
Manuscript ID	SOSYM-25-00005370
Manuscript Type:	Regular Paper
Keyword:	Tenants Placement, Cloud, Data Centers, Formal Concept Analysis, Energy Efficiency, Software Modeling, System Modeling

SCHOLARONE™
Manuscripts

1
2 **Subject:** Revision Plan for Manuscript SOSYM-25-00005347 for Software and Systems
3 Modeling
4

5 Dear Editor,
6
7

8 Thank you for your thoughtful review and for considering our manuscript, "Optimizing
9 Energy Efficiency in Multi-Tenant Cloud Environments Using Formal Concept Analysis."
10 We appreciate the feedback that the manuscript requires a clearer positioning within the
11 scope of software and systems modeling.
12
13

14 We have carefully analyzed the comments and have developed a comprehensive revision
15 plan to address this central point. Our revisions will focus on explicitly framing our FCA-
16 based methodology as a **model-driven approach** for designing and optimizing cloud
17 resource management systems. The core technical contributions and experimental results
18 will be strengthened, not altered, by this refined framing.
19
20

21 The key revisions will include:
22

- 23 1. **A modified title** that immediately reflects the modeling-centric contribution of the
24 work.
25
- 26 2. **A rewritten Abstract and Introduction** to firmly establish the contribution to
27 software and systems modeling.
28
- 29 3. **A significantly expanded "Related Work" section** with a new subsection
30 dedicated to modeling approaches in cloud computing, clearly positioning FCA
31 among other formal methods.
32
- 33 4. **A refined "Contributions" section** that explicitly uses the language of modeling.
34
- 35 5. **An enhanced "Discussion" section** that elaborates on the modeling implications,
36 reusability, and integration with MDE practices.
37
38
39
40

41 To better reflect the revised focus of our work, we propose the following new title for the
42 manuscript:
43
44

45 **"Modeling Tenant Dependencies with Formal Concept Analysis for Energy-Aware
46 Cloud Resource Management"**
47
48
49

50
51 **N.B.:** In the revised manuscript, all new or modified text will be highlighted in blue to
52 facilitate your review.
53

54 We are confident that these revisions will align the manuscript perfectly with the aims and
55 scope of Software and Systems Modeling.
56
57

58 We submit the fully revised manuscript for your reconsideration.
59

60 Sincerely,
Jalel Eddine Hajlaoui & Mohamed Nazih Omri

Modeling Tenant Dependencies with Formal Concept Analysis for Energy-Aware Cloud Resource Management

Jalel Eddine Hajlaoui¹ and Mohamed Nazih Omri^{1†}

^{1,2}✉ Mars Research Laboratory, University of Sousse, LR17ES05,
Tunisia.

Contributing authors: jaleleddine.hajlaoui@epfedu.fr;
mohamednazih.omri@eniso.u-sousse.tn;

[†]These authors contributed equally to this work.

Abstract

Energy consumption constitutes a dominant operational cost in cloud data centers, driven largely by inefficient resource utilization in multi-tenant environments. While many solutions address this via heuristic or metaheuristic optimization, a lack of formal models capturing semantic dependencies between tenants leads to suboptimal solutions. This paper presents a **model-driven approach** for optimizing tenant placement to minimize energy use while adhering to strict Service Level Agreements (SLAs). We propose a novel methodology that leverages **Formal Concept Analysis (FCA)** to construct a **conceptual model**—a concept lattice—that intelligently groups interdependent tenants based on their shared application usage patterns. This dependency-aware model enables two key advantages: maximizing resource utilization on active servers and minimizing inter-server network communication overhead. We introduce two core algorithms that utilize this model: an Algorithm for Consolidation of Tenants (ACT) and an Algorithm for Placement of Tenants (APT). Extensive simulations using CloudSim Plus demonstrate that our **FCA-based modeling approach** significantly outperforms state-of-the-art baselines, including Best-Fit Decreasing (BFD), Genetic Algorithms (GA), and a Reinforcement Learning (RL) agent. Results show a **33.3% reduction in active servers** and a **11% reduction in total energy consumption** for a deployment of 1,000 tenants, all while maintaining a low SLA violation rate. Our work establishes

FCA as a powerful formal modeling tool for the design and optimization of resource management in software-intensive cloud systems.

Keywords: Tenants Placement, Cloud, Data Centers, Formal Concept Analysis, Energy Efficiency, [Software Modeling](#), [System Modeling](#)

1 Introduction

1.1 Context and Issues

Cloud computing has become the backbone of modern IT infrastructure, with the global market projected to exceed \$1 trillion by 2028 [1]. This growth is fueled by advancements in virtualization, scalability, and cost efficiency, leading to an explosive increase in cloud-hosted analytical workloads [2, 3]. However, this expansion comes with a significant energy footprint. Recent studies indicate that cloud data centers account for approximately 2% of global electricity consumption [4], with energy costs surpassing \$30 billion annually [5]. A major contributor to this inefficiency is the multi-tenant nature of cloud environments, where inefficient placement of tenants (users or applications) on physical servers leads to resource fragmentation, over-provisioning, and the "noisy neighbor" effect, which degrades performance and Quality of Service (QoS) [6]. Furthermore, underutilized servers can consume up to 70% of their peak power even at low loads [7], making energy consumption the dominant operational cost factor, accounting for an average of 40% of expenses in data centers. This creates a critical challenge for cloud providers: maximizing profitability while adhering to strict Service Level Agreements (SLAs) and reducing their environmental impact.

This paper addresses the energy efficiency challenge from a **software and systems modeling perspective**. We posit that a core reason for the limitations of existing approaches is the lack of an explicit, analyzable model that captures the *semantic relationships* between tenants and their applications. Without such a model, placement strategies are often based on raw resource vectors, ignoring the communication patterns and shared dependencies that significantly impact performance and energy consumption. We propose the use of **Formal Concept Analysis (FCA)** [8] as a formal method to derive a **conceptual model** of the multi-tenant system. This model, represented as a concept lattice, systematically reveals groups of tenants (concepts) that are interdependent. By leveraging this model to drive placement decisions, we transition from a purely algorithmic optimization to a **model-driven optimization**, where the structure and insights provided by the formal model guide the consolidation process. This work thus contributes to the field of software and systems modeling by introducing and validating a novel modeling technique for a critical problem in cloud system design.

1.2 Research Questions

To address the aforementioned challenges, this work is guided by the following primary research question:

- How can tenant placement in multi-tenant cloud environments be optimized to significantly reduce energy consumption without violating Service Level Agreement (SLA) constraints?

This central question is decomposed into the following sub-questions:

- RQ1: How can formal methods be leveraged to systematically identify and group interdependent tenants to minimize inter-server communication and network energy overhead?
- RQ2: What consolidation strategy can maximize resource utilization on active servers, thereby minimizing the number of active servers required?
- RQ3: How can this approach maintain scalability and computational efficiency to be feasible for large-scale, dynamic cloud environments?

1.3 Contributions

In response to these questions, this paper makes the following key contributions from a software and systems modeling viewpoint:

- **A Novel Modeling Formalism:** We introduce Formal Concept Analysis (FCA) as a modeling technique for multi-tenant cloud environments. We formalize the energy-aware tenant placement problem and show how a formal context derived from tenant-application usage data yields a concept lattice that serves as a foundational **conceptual model** for the system.
- **FCA-Based Methodology:** We propose a novel tenant placement methodology that utilizes the FCA-generated concept lattice to intelligently group tenants. This **dependency-aware model** enables more efficient packing and reduces network traffic by explicitly modeling and leveraging the relationships between tenants and applications.
- **Model-Driven Algorithms:** We design and implement two core, model-driven algorithms: an Algorithm for Consolidation of Tenants (ACT) for initial placement and migration, and an Algorithm for Placement of Tenants (APT) that directly consumes the concept lattice to place coherent groups of tenants, thereby translating the conceptual model into an executable optimization strategy.
- **Comprehensive Evaluation:** We conduct an extensive experimental evaluation using the CloudSim Plus framework, demonstrating that our **model-based approach** achieves a significant reduction in the number of active servers (up to 33.3%) and total energy consumption compared to state-of-the-art baselines, all while maintaining a low SLA violation rate. This validates the practical value of the proposed modeling technique.

1.4 Assumptions

The proposed approach and its evaluation are based on the following key assumptions:

- Tenant resource demands (CPU, memory, I/O) are known or can be accurately forecasted for the scheduling period.
- The cloud infrastructure comprises homogeneous servers with known and fixed capacities.
- The formal context, which defines the relationship between tenants and application instances, is stable enough for the concept lattice to provide meaningful groupings between scheduling intervals.
- Live migration of tenants is possible with negligible performance impact relative to the evaluation intervals used.
- The primary sources of energy consumption are the active servers and the network traffic between them.

1.5 Paper Organization

The remainder of this paper is structured as follows. Necessary background on cloud computing and Formal Concept Analysis is provided in Section 2. A review of related work on energy-efficient resource allocation is presented in Section 3. Our motivation and the formalization of the tenant placement problem are detailed in Section 4. The proposed FCA-based approach is presented in Section 5. The experimental setup and analysis of results are described in Section 6. A discussion of the findings, limitations, and future research directions follows in Section 7. Finally, the paper is concluded in Section 8.

2 Basic Concepts and Background

This section provides the foundational knowledge required to understand the subsequent contributions of this paper. We begin by outlining the core principles of cloud computing and its service models. We then introduce Formal Concept Analysis (FCA), the key mathematical framework underpinning our proposed approach. Finally, we discuss the specific background of energy consumption challenges in multi-tenant cloud environments that motivate this work.

2.1 Core Concepts of Cloud Computing

Cloud computing has become the backbone of modern IT infrastructure, representing a paradigm shift from traditional owned infrastructure to on-demand, service-oriented models. The National Institute of Standards and Technology (NIST) provides a widely accepted definition [9]:

“Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.”

This model is characterized by five essential attributes:

- **On-demand self-service:** Consumers can automatically provision computing capabilities without human interaction.
- **Broad network access:** Capabilities are available over the network through standard mechanisms.
- **Resource pooling:** Provider's computing resources are pooled to serve multiple consumers.
- **Rapid elasticity:** Capabilities can be elastically provisioned and released to scale rapidly.
- **Measured service:** Resource usage can be monitored, controlled, and reported for transparency.

These resources are typically delivered through three primary service models, illustrated in Figure 1:

- **Infrastructure as a Service (IaaS):** Provides virtualized computing resources over the internet (e.g., virtual machines, storage).
- **Platform as a Service (PaaS):** Offers a development and deployment environment in the cloud, allowing developers to build applications without managing the underlying infrastructure.
- **Software as a Service (SaaS):** Delivers software applications over the internet on a subscription basis, accessible via web browsers.

The multi-tenant nature of these models, where multiple customers share the same physical resources, is central to the cloud's efficiency but also introduces the energy optimization challenges addressed in this paper.

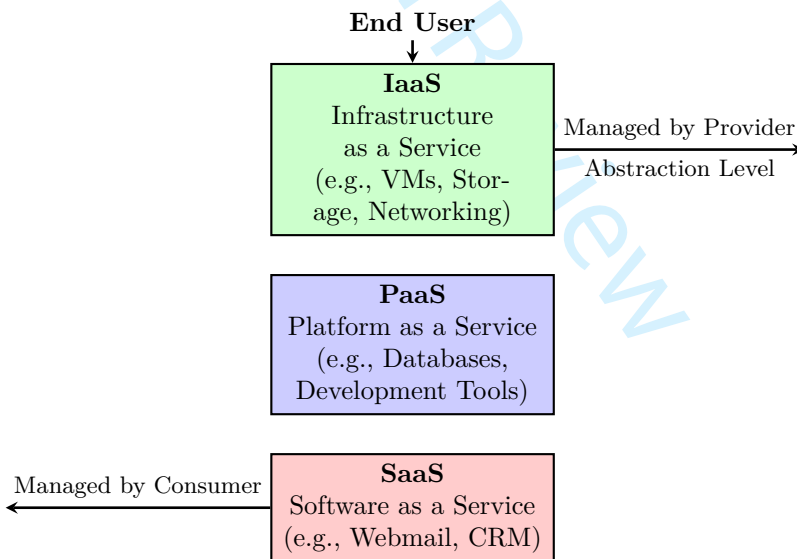


Fig. 1: The three primary cloud computing service models (IaaS, PaaS, SaaS) illustrating the shared responsibility and abstraction levels between the cloud provider and consumer.

2.2 Formal Concept Analysis: An Overview

Formal Concept Analysis (FCA), introduced by the German mathematician Rudolf Wille in 1982, is a principled method for data analysis, knowledge representation, and conceptual clustering. It specializes in extracting a structured, ordered set of concepts from a dataset, revealing implicit relationships between objects and their attributes.

Formal Context

The foundation of FCA is the *formal context*, defined as a triplet (X, Y, I) where:

- $X = \{x_1, x_2, \dots, x_n\}$ is a set of *objects*,
- $Y = \{y_1, y_2, \dots, y_m\}$ is a set of *attributes*,
- $I \subseteq X \times Y$ is a binary *incidence relation* indicating that object x_i has attribute y_j .

A formal context is often represented as a cross table (binary matrix), where rows are objects, columns are attributes, and a checkmark (or 1) denotes that the relation $(x_i, y_j) \in I$ holds.

Formal Concept

A *formal concept* of the context (X, Y, I) is a pair (A, B) where:

- $A \subseteq X$ is called the *extent* (set of all objects sharing the attributes in B),
- $B \subseteq Y$ is called the *intent* (set of all attributes common to the objects in A),
- such that every object in A has every attribute in B , and there are no other objects that have all attributes in B , and there are no other attributes that all objects in A have.

The set of all formal concepts for a given context is denoted by $\mathfrak{B}(X, Y, I)$.

Concept Lattice

The concepts extracted from a formal context can be organized into a hierarchical structure based on a subconcept-superconcept relationship. For two concepts (A_1, B_1) and (A_2, B_2) , we say (A_1, B_1) is a subconcept of (A_2, B_2) if $A_1 \subseteq A_2$ (equivalently, $B_2 \subseteq B_1$). This relationship defines a partial order, and the set of all concepts $\mathfrak{B}(X, Y, I)$ together with this order forms a complete lattice called the *concept lattice*, which can be visually represented as a Hasse diagram.

In our approach, the formal context (A, T, I) consists of a set of application instances A as objects and a set of tenants T as attributes. The incidence relation I indicates which tenant $t \in T$ invokes which application instance $a \in A$. The resulting formal concepts will group together tenants (intension) that use a common set of application instances (extension), which is directly leveraged for energy-efficient placement.

2.3 Energy Consumption in Multi-Tenant Cloud Environments

The efficiency gains of cloud computing are counterbalanced by its significant energy footprint. Modern multi-tenant cloud data centers suffer from energy inefficiencies due to resource fragmentation, inter-tenant interference, and over-provisioning [10]. Studies indicate that 30-40% of data center energy is wasted on idle or underutilized servers [11], a problem exacerbated by the “noisy neighbor” effect, where co-located workloads compete for shared resources (e.g., CPU cache, I/O bandwidth), degrading performance and QoS [12].

The rise of microservices and serverless architectures further complicates energy management. These fine-grained, ephemeral workloads lead to frequent power state transitions in servers, preventing them from entering deep power-saving states and thus diminishing energy proportionality [13, 14]. Without intelligent, semantic-aware scheduling that understands relationships between workloads, cloud providers face diminishing returns in optimizing energy use. This critical challenge motivates the application of structured analysis techniques like FCA, which can systematically uncover these workload relationships to inform more efficient scheduling and consolidation decisions, directly targeting the root causes of energy waste.

3 Related Works

The challenge of energy-efficient resource allocation in multi-tenant cloud environments has been addressed through a variety of methodologies. Existing research can be broadly categorized based on their primary optimization objectives and techniques. This section provides a thematic review of these approaches, highlighting their strengths and limitations in tackling the tenant placement problem.

3.1 Metaheuristic Approaches

Metaheuristic algorithms are widely employed to solve the NP-hard tenant/Virtual Machine (VM) placement problem due to their ability to find near-optimal solutions in large search spaces. A common inspiration is Ant Colony Optimization (ACO), modeled after the foraging behavior of ants. Early work by [15] laid the foundation for ACO. Subsequent studies have adapted it for cloud resource allocation; for instance, [16] combined ACO with a multi-objective genetic algorithm to optimize both resource use and energy efficiency. Similarly, [17] proposed GACA-VMP, a hybrid Genetic Ant Colony Algorithm for VM placement, simplifying the problem by considering only CPU and memory constraints.

Other swarm intelligence techniques have also been explored. [18] introduced VMGRP, an algorithm based on the ant recruitment process, aiming to balance the usage of CPU, memory, and I/O resources to minimize energy

consumption without violating SLA constraints. Beyond swarm intelligence, genetic algorithms (GAs) have been applied. [19] presented a hybrid Tenant Placement Strategy (TPS) combining GA with Case-Based Reasoning (CBR) to optimize tenant consolidation. While powerful, these metaheuristics often suffer from high computational complexity, making them less suitable for highly dynamic environments requiring real-time decisions.

3.2 Greedy and Heuristic-Based Approaches

Greedy algorithms offer a more computationally efficient alternative, making them practical for online placement scenarios. A common strategy is the Best-Fit Decreasing (BFD) heuristic. [20] developed a framework for dynamic resource demand forecasting and used a BFD-based heuristic for allocation, assigning VMs to the physical machine with the smallest available resources.

The Online Tenant Placement Problem (OTPP) is another key focus. [21] proposed the Tenant Dispatch Heuristic (TDH), which distributes tenants across a fixed number of nodes to maximize tenant count without degrading QoS. Their approach uses the Kullback-Leibler Divergence (KLD) to measure the distance between tenant resource demand and node residual capacity. Later work, such as by [22], extended TDH to consider application functionality for better consolidation. These heuristic methods are fast and scalable but can be myopic, often missing globally optimal solutions due to their step-by-step decision-making nature.

3.3 QoS and Performance-Aware Approaches

Maintaining Quality of Service (QoS) is a critical constraint alongside energy savings. Research in this area focuses on mitigating performance interference (the "noisy neighbor" effect) and ensuring SLA compliance. [23] proposed dynamic resource balancing, while [24] developed application-combining strategies for tenant placement. Recently, Formal Concept Analysis (FCA) has emerged as a powerful tool for QoS-aware modeling. [25] used FCA for dynamic resource reconfiguration, achieving a 20% reduction in energy use. Hybrid methods, such as those combining FCA with reinforcement learning [26] or fuzzy logic [27], have been proposed to handle uncertainty and fluctuating workloads, improving adaptability while preserving QoS guarantees.

3.4 Cost and Carbon Emission Optimization Approaches

This category of work explicitly models economic and environmental costs. Strategies include dynamic allocation, auto-scaling, and leveraging renewable energy. [28] used machine learning for VM placement to lower power consumption. Auto-scaling mechanisms that adjust resources based on real-time demand have proven effective in reducing idle server waste [29]. Other innovations include carbon-aware scheduling that prioritizes low-carbon time slots [30] and hybrid models combining FCA with deep reinforcement learning to optimize cost-performance trade-offs [31].

Specific notable contributions include:

- **Queueing Theory:** [32] modeled tenant placement as a queue network, using a discrete-time Markov chain (DTMC) and an SLA-aware algorithm to minimize resource usage through component sharing.
- **Profit Maximization:** [33] focused on maximizing provider profit (SLA revenue minus server and violation costs) using approximation algorithms based on the Best-Fit heuristic.
- **Carbon-Aware Placement:** [34] proposed CACEV, a method combining the A* algorithm with fuzzy logic to optimize both cost and carbon emissions for VM placement in distributed clouds, considering network traffic.

3.5 Other Relevant Approaches

Further research has explored the problem in specific contexts. [35] abstracted tenant-to-server assignment in database clusters as a token allocation problem. [36] implemented the STeP framework for workload assignment in Azure SQL Database, using greedy algorithms based on scalar and time-series cost models. [37] applied a multi-tenant model to rural water planning, using a GA for tenant placement on a Eucalyptus platform to reduce costs.

3.6 Modeling Approaches for Cloud Systems

A complementary stream of research applies formal modeling and analysis techniques to manage the complexity of cloud systems. The goal is to create abstractions that facilitate reasoning about system properties like performance, reliability, and efficiency before deployment. Queueing networks and Petri nets have been used to model and analyze performance and resource contention in data centers [38]. Ontologies and feature models have been proposed to capture the variability and configuration of cloud services and applications [39]. Within this context, **Formal Concept Analysis (FCA)** has emerged as a powerful method for knowledge representation and conceptual clustering, with recent applications in software engineering for feature location, domain analysis, and software reconfiguration [40, 41].

Our work is positioned within this modeling tradition. We leverage FCA not merely as a data analysis tool but as a **formal modeling framework** for discovering and representing the latent structure in multi-tenant cloud environments. Unlike queueing models that focus on performance dynamics or ontologies that describe taxonomic relationships, the FCA model specifically captures the *incidence relations* between tenants (as attributes) and application instances (as objects). The resulting concept lattice provides a hierarchical model of tenant groups based on shared application usage. This model is then directly utilized to inform placement decisions. To the best of our knowledge, this is the first work to apply FCA for **energy-aware resource modeling**

in multi-tenant clouds, bridging a gap between conceptual modeling rooted in formal methods and the operational problem of energy-efficient scheduling.

3.7 Summary and Discussion

A summary of the discussed approaches, their primary objectives, and inherent limitations is presented in Table 1. While static algorithms like bin-packing are simple and efficient, they lack adaptability [42]. Dynamic and hybrid approaches offer better responsiveness to change but introduce overhead from migrations and require complex prediction models [43]. Our work seeks to address these gaps by leveraging FCA to provide a structured, mathematically sound framework for dependency-aware grouping. This approach systematically reduces network overhead and enables more efficient packing, aiming to achieve significant energy savings without the computational overhead of metaheuristics or the myopia of simple heuristics.

Table 1: Summary of Tenant/VM Placement Approaches in Cloud Computing

Category	Primary Objective	Key Techniques	Limitations
Metaheuristics	Near-optimal energy/resource optimization	ACO [16–18], Genetic Algorithms [19, 44, 45]	High computational complexity, slow for large-scale dynamic systems
Greedy/Heuristic	Fast, online placement	Best-Fit Decreasing [20], TDH [21, 22]	Can be myopic, may converge to local optima
QoS-Aware	Balance energy savings with performance	FCA [25, 46], Hybrid RL/-Fuzzy [27, 47]	Model complexity, tuning parameters for SLA constraints
Cost & Carbon Optimization	Minimize operational & environmental cost	Auto-scaling [29], Carbon-aware scheduling [30], Queueing Theory [32]	Requires accurate cost and carbon footprint models

4 Motivation and Formalization of the Tenant Placement Problem

4.1 Motivation

The rapid growth of cloud computing has led to widespread adoption of multi-tenant environments, where multiple users share computational resources to improve efficiency and reduce costs. However, this shared infrastructure often results in significant energy consumption due to inefficient resource allocation, over-provisioning, and idle resources. High energy usage not only increases operational expenses for cloud providers but also contributes to environmental concerns, including higher carbon emissions. Existing approaches to energy optimization often rely on heuristic-based methods, which may lack precision

or fail to capture the underlying relationships between workloads and resource utilization. Formal Concept Analysis (FCA) offers a structured, mathematical framework for identifying patterns and dependencies in complex systems, making it a promising tool for analyzing and optimizing energy consumption in multi-tenant clouds. By leveraging FCA, this work aims to uncover hidden associations between tenant workloads and resource usage, enabling more intelligent and energy-efficient scheduling strategies. The potential benefits include reduced operational costs, improved sustainability, and enhanced performance in cloud environments, motivating the exploration of FCA as a novel approach to energy optimization in multi-tenant systems.

4.2 Problem Formalization

In a multi-tenant Cloud environment, resource demands vary across customers depending on their specific requirements, and over time, resources may be released while new ones are allocated. Tenant placement techniques can be employed in data centers to efficiently meet customer needs while optimizing the use of available resources. However, the quality of services provided by a Cloud provider must comply with SLA (*Service Level Agreement*) constraints [33]. Therefore, the consolidation strategy in a data center involves finding available server capacity to host tenants while ensuring uninterrupted service performance.

Each client is characterized by a requested amount of resources, including CPU, memory, and input/output operations, among others. Consequently, the placement strategy must identify the most suitable server from a resource consumption standpoint. Additional factors such as migration overhead and network traffic must also be considered when executing the placement.

Any tenant placement strategy requires grouping certain tenants on the same server. The decision to migrate depends on various factors, including server capacity, tenant size, and SLA constraints. In our previous work [48], we addressed the multi-tenancy-aware discovery of Cloud services. Our current research, which extends this work, aims to develop a strategy that meets tenant expectations while minimizing energy consumption on the Cloud provider's side.

In the rest of this section, we formalize the tenant placement problem. We first define the input and output parameters for our algorithm as follows.

Let:

- $T = \{t_1, \dots, t_n\}$: a set of tenants.
- $A = \{a_1, \dots, a_m\}$: a set of instances of the applications.
- $S = \{s_1, \dots, s_p\}$: a set of servers.
- $VM = \{v_1, \dots, v_k\}$: a set of virtual machines.
- $reg(t)$: d-dimensional vector which indicates the quantity of resources requested by a tenant t .
- $Size(a)$: d-dimensional vector which designates the size of an application instance a .

- $Size(v)$: d-dimensional vector which designates the size of a virtual machine v .
- $Cap(v)$: d-dimensional vector which designates the capacity of a virtual machine v .
- $Cap(s)$: d-dimensional vector which designates the capacity of the server s .
- d : Number of resource types.

In our approach we will take into consideration 3 types of resources; CPU, memory and inputs/outputs. We note " \leq " is a comparison between two d-dimensional vectors: for $x, y \in R^d, x \leq y$ if and only if $x_j \leq y_j$ for all $j = 1, \dots, d$.

- $a(t)$: the application instance hosting the tenant t .
- $v(a)$: the virtual machine hosting the application instance a .
- $s(v)$: server hosting the virtual machine v .
- $\Delta(a, t)$: d-dimensional vector designates the increase in the size of the instance a following the hosting of the tenant t .
- $load(s)$: d-dimensional vector designates the total size of virtual machines hosted by the server s .
- v_0 : vector which designates the size of an empty virtual machine.
- a_0 : vector which designates the size of an instance without any registered tenant.
- $State(s)$: the state of the server s (active or not).

Let Plc denote the function of placing a tenant in the set of application instances: $t_i * A \mapsto Plc$. Our goal is to pack the maximum number of tenants into the minimum number of servers in order to maximize the use of resources. The objective is modeled by the following equation:

$$\min \sum_{i \in n} S_i \tag{1}$$

We consider the following binary assertions:

$$y_{(t,a_i)} = \begin{cases} 1 & \text{if tenant } t \text{ is on instance } a_i \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

$$y_{(a,v_j)} = \begin{cases} 1 & \text{if instance } a \text{ is on VM } v_j \forall i, j \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

SLA constraints are defined as a contract that quantifies the minimum service level for a service that a supplier undertakes to deliver to his customer. In order to achieve our objective, we define a set of constraints to be respected:

Each tenant t must be assigned to at least one application instance:

$$\sum_i y_{(t,a_i)} \geq 1 \quad \forall t \in T \tag{4}$$

Each virtual machine must contain a single copy of application instance a :

$$\sum_j y_{(a,v_j)} \leq 1 \quad \forall a \in A \quad (5)$$

The size of each instance a is calculated as follows:

$$\text{size}(a) = a_0 + \sum_{\substack{t \in T \\ a(t)=a}} \text{reg}(t) \cdot y_{(t,a)} \quad \forall a \in A \quad (6)$$

The size of a VM must not exceed its capacity:

$$\text{size}(v) = v_0 + \sum_{\substack{a \in A \\ v(a)=v}} \text{size}(a) \cdot y_{(a,v)} \leq \text{cap}(v) \quad \forall v \in VM \quad (7)$$

The capacity of the server s must not be violated:

$$\text{load}(s) = \sum_{\substack{v \in VM \\ s(v)=s}} \text{size}(v) \leq \text{cap}(s) \quad \forall s \in S \quad (8)$$

We define a d -dimensional vector St_i which represents the amount of unused resources in each server s_i . The stimulus St_i is calculated as follows:

$$St_i = \text{cap}(s_i) - \text{load}(s_i), \forall i \quad (9)$$

5 FCA-based Multi-tenancy Proposed Approach

The proposed approach is made up of two phases. The first is to group tenants that use the same instances. The second ensures the placement of groups of tenants on the minimum number of servers. Our goal is to group the maximum number of tenants into the minimum number of servers.

5.1 Tenants consolidation phase

To meet our goal, we recall that we use FCA because the notion of concept perfectly meets our need. This notion makes it possible to group together the maximum number of dependent tenants. In our approach, the placement of a new tenant tends to maximize the amount of resources exploited. The set of tenants $T = \{t_1, t_2, \dots, t_n\}$ must be wrapped in the minimum number of application instances $A = \{a_1, a_2, \dots, a_m\}$. A tenant $t_i \in T$ migrates to an instance $a_j \in A$ whose quantity of unused resources is minimal. First, we define a procedure (Algorithm 1) that allows to add a tenant and update the amount of available resources where:

- $v(a)$: A function that returns the virtual node (VM) assigned to application a

- $s(v(a))$: A function that returns the physical server hosting virtual node $v(a)$
- $load(s(v(a)))$: Current load of the physical server
- $cap(s(v(a)))$: Capacity of the physical server
- $\Delta(a, t)$: Resources required by tenant t for application a

Algorithm 1 Procedure `add_new_tenant`

```

1: procedure ADD_NEW_TENANT( $t, s(v(a))$ )
2:   if  $load(s(v(a))) + \Delta(a, t) \leq cap(s(v(a)))$  then
3:      $load(s(v(a))) \leftarrow load(s(v(a))) + \Delta(a, t)$ 
4:     Assign tenant  $t$  to server  $s(v(a))$ 
5:      $cap(s(v(a))) \leftarrow cap(s(v(a))) - \Delta(a, t)$ 
6:     Remove  $t$  from  $T$ 
7:     return success
8:   else
9:     return failure
10:  end if
11: end procedure

```

The Algorithm 2 proposed for the consolidation of tenants operates in two distinct phases: initial placement and consolidation through migration. In the initial placement phase, the algorithm stores the stimulus in a list, finds the couple (tenant, instance) which maximizes the exploitation of resources and places the tenant in this instance. The broadcast (St_i) function allows a server to broadcast its stimulus to the entire set of tenants, enabling a distributed decision-making process where tenants can evaluate available resources across the infrastructure. The consolidation phase implements a sophisticated migration strategy aimed at optimizing resource utilization and energy efficiency. Once initial placements are made, the algorithm identifies servers with suboptimal resource utilization through the $FindMax(L.Stimulus)$ function, which returns servers with the maximum amount of unused resources. These servers become candidates for consolidation. The migration process operates systematically by:

5.2 List of Functions Used in Algorithms

Algorithm 2 Algorithm for Consolidation of Tenants (ACT) - Part 1

Require: Set of tenants T , Set of application instances A **Ensure:** Set of tenant-instance assignments T_A

```

1:  $T_A \leftarrow \emptyset$  ▷ Initialize solution set
2:  $Q \leftarrow$  empty queue
3: EnqueueAll( $Q, T$ ) ▷ Add all tenants to the queue

4: Phase 1: Initial Placement
5: for each server  $s_i \in S$  do
6:   Broadcast( $St_i$ ) ▷ Server broadcasts its stimulus
7: end for
8: while  $Q$  is not empty do
9:    $t_j \leftarrow$  Dequeue( $Q$ )
10:   $L \leftarrow \emptyset$  ▷ Feasible servers list
11:  for each server  $s_i \in S$  do
12:    if  $St_i \geq \text{reg}(t_j)$  then
13:       $L.\text{add}(\langle s_i, St_i \rangle)$ 
14:    end if
15:  end for
16:  if  $L$  is not empty then
17:     $s_{\text{target}} \leftarrow \text{FindMin}(L)$  ▷ Best-Fit server
18:    success  $\leftarrow$  AddTenantToServer( $t_j, s_{\text{target}}$ )
19:    if success then
20:       $T_A.\text{add}(\langle t_j, a(s_{\text{target}}) \rangle)$ 
21:      UpdateStimulus( $s_{\text{target}}$ )
22:    else
23:      Enqueue( $Q, t_j$ )
24:    end if
25:  else
26:    break ▷ No server available
27:  end if
28: end while

```

Algorithm 3 Algorithm for Consolidation of Tenants (ACT) - Part 2

```

29: Phase 2: Consolidation via Migration
30:  $L_{\text{underutilized}} \leftarrow \text{FindServersWithMaxStimulus}()$ 
31: for each server  $s_{\text{source}} \in L_{\text{underutilized}}$  do
32:   for each tenant  $t_j$  hosted on  $s_{\text{source}}$  do
33:     migrated  $\leftarrow$  false
34:     for each active server  $s_{\text{dest}} \in S$  where  $s_{\text{dest}} \neq s_{\text{source}}$  do
35:       if CanHost( $s_{\text{dest}}, t_j$ ) then
36:         Migrate( $t_j, s_{\text{source}}, s_{\text{dest}}$ )
37:         UpdateStimulus( $s_{\text{source}}$ )
38:         UpdateStimulus( $s_{\text{dest}}$ )
39:         migrated  $\leftarrow$  true
40:         break
41:       end if
42:     end for
43:   end for
44:   if IsEmpty( $s_{\text{source}}$ ) then
45:     PowerOff( $s_{\text{source}}$ )
46:   end if
47: end for
48: return  $T_A$ 

```

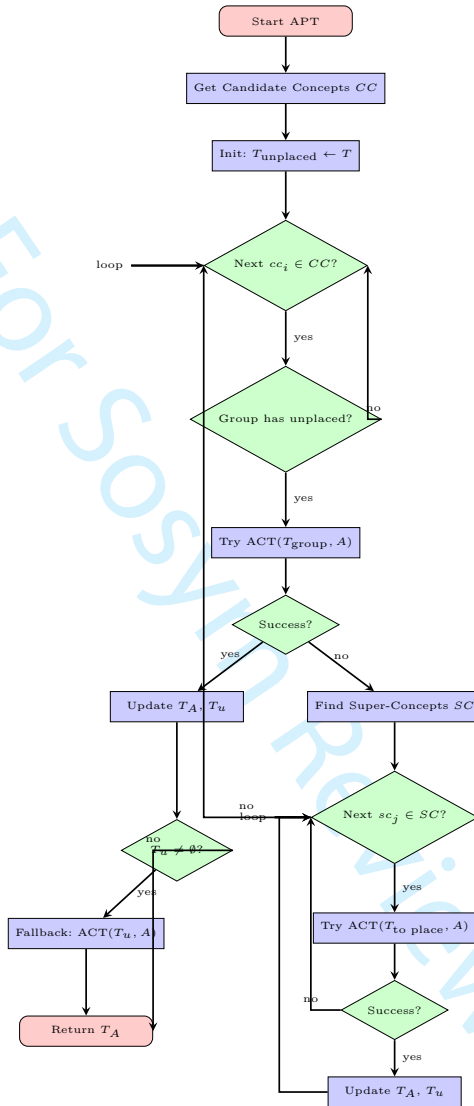


Fig. 2: Flowchart of the Algorithm for Placement of Tenants (APT). The process begins by extracting candidate concepts from the FCA lattice. It then iteratively attempts to place tenant groups from these concepts. If placing a group fails, it explores super-concepts to find viable larger groups. A final fallback ensures all tenants are placed.

Table 2: List of Functions Used in Algorithms

Function	Description
BROADCAST(St_i)	Server s_i advertises its current stimulus (available resource vector).
FINDMIN(L)	Returns the server in list L with the smallest stimulus (i.e., least free resources, implementing a Best-Fit strategy).
ADDTENANTTOSERVER(t_j, s)	Attempts to place tenant t_j on server s . Returns true if successful and updates the server's load, false otherwise.
UPDATESTIMULUS(s)	Recalculates and broadcasts the stimulus (available resources) for server s .
FINDSERVERS-WITHMAXSTIMULUS()	Returns a list of servers with the highest stimulus (i.e., most underutilized).
CANHOST(s, t_j)	Checks if server s has sufficient resources to host tenant t_j .
MIGRATE(t_j, s_{src}, s_{dest})	Moves tenant t_j from s_{src} to s_{dest} , updating both servers' states.
GETCANDIDATECONCEPTS(\mathcal{L})	Analyzes the concept lattice \mathcal{L} and returns a set of concepts CC that have high weight ($W(c_i)$) and collectively cover all tenants.
FINDSUPERCONCEPTS(c, \mathcal{L})	Returns the set of immediate super-concepts (parent nodes) of concept c in the lattice \mathcal{L} .

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33

Algorithm 4 Algorithm for Placement of Tenants (APT)**Require:** Set of tenants T , Set of application instances A , Concept Lattice \mathcal{L} **Ensure:** Set of tenant-instance assignments T_A

```

1:  $CC \leftarrow \text{GETCANDIDATECONCEPTS}(\mathcal{L})$  Concepts with max weight & coverage
2:  $T_A \leftarrow \emptyset$ 
3:  $T_{\text{unplaced}} \leftarrow T$  Copy of tenants to track unplaced ones
4: for each candidate concept  $cc_i \in CC$  do
5:    $T_{\text{group}} \leftarrow cc_i.\text{Int}$  Get the group of tenants in this concept
6:   if  $T_{\text{group}} \cap T_{\text{unplaced}} = \emptyset$  then
7:     continue Skip if all tenants in this concept are already placed
8:   end if
9:    $Plc \leftarrow \text{ACT}(T_{\text{group}}, A)$  Try to place the entire group
10:  if  $Plc \neq \emptyset$  then
11:     $T_A.\text{ADD}(Plc)$ 
12:     $T_{\text{unplaced}} \leftarrow T_{\text{unplaced}} \setminus T_{\text{group}}$  Remove placed tenants
13:  else
14:     $SC \leftarrow \text{FINDSUPERCONCEPTS}(cc_i, \mathcal{L})$  Find larger parent concepts
15:    for each super-concept  $sc_j \in SC$  do
16:       $T_{\text{larger group}} \leftarrow sc_j.\text{Int}$ 
17:       $T_{\text{to place}} \leftarrow T_{\text{larger group}} \cap T_{\text{unplaced}}$  Place only unplaced tenants
18:       $Plc \leftarrow \text{ACT}(T_{\text{to place}}, A)$ 
19:      if  $Plc \neq \emptyset$  then
20:         $T_A.\text{ADD}(Plc)$ 
21:         $T_{\text{unplaced}} \leftarrow T_{\text{unplaced}} \setminus T_{\text{to place}}$ 
22:        break Exit inner loop after successful placement
23:      end if
24:    end for
25:  end if
26: end for

```

27: Fallback for Remaining Tenants

```

28: if  $T_{\text{unplaced}} \neq \emptyset$  then
29:    $Plc \leftarrow \text{ACT}(T_{\text{unplaced}}, A)$  Place any leftover tenants
30:    $T_A.\text{ADD}(Plc)$ 
31: end if
32: return  $T_A$ 

```

- Identifying active tenants on underutilized servers (where $y_{t_j, a_i} = 1$ and $state(s_i(v(a_j))) = on$)
- Evaluating potential destination servers that remain active ($state(s_n) = on$)
- Performing capacity verification using the constraint $load(s_n) + \Delta(a_i, t_j) \leq cap(s_n)$ to ensure the target server can accommodate the migrated tenant without exceeding its capacity

- Executing tenant migration from source server s_i to destination server s_n when capacity constraints are satisfied

The migration strategy follows a server consolidation approach where the goal is to concentrate tenants on fewer servers, thereby allowing underutilized servers to be deactivated ($state(s_i) = off$). This consolidation mechanism directly contributes to energy savings and improved resource efficiency by reducing the number of active physical servers while maintaining service quality. The $FindMin(L.Stimulus)$ and $FindMax(L.Stimulus)$ functions serve complementary roles: the former identifies servers with minimal available resources (high utilization) for optimal initial placement, while the latter identifies servers with maximal available resources (low utilization) as consolidation targets. The algorithm's dual-phase approach ensures both optimal initial resource allocation and continuous optimization through strategic tenant migration, making it particularly effective for dynamic multi-tenant cloud environments where resource demands fluctuate over time.

5.3 Tenants Placement Phase

The FCA first groups together the instances and the dependent tenants in the same concept using a concept lattice. Then, we select the candidate concepts. Let the two sets $T = \{t_1, t_2, \dots, t_n\}$ and $A = \{a_1, a_2, \dots, a_m\}$ be the set of tenants and applications instances, respectively. We say that the two tenants t_i, t_j are dependent if they exploit the same instances. Instances of applications are dependent if they provide services to the same set of tenants. For example, the number of dependencies for the two instances a_1 and a_2 is defined by the number of tenants served by both a_1 and a_2 :

$$Depend_{a_i, a_j}^A = Count(T_{a_i} \cap T_{a_j}).$$

We define a concept C whose extension represents the set of instances and its intension represents the group of tenants: $C = (Ext, Int)$. A group of tenants $\{t_i\} \in T$ and instances $\{a_i \in A\}$ define a partition Pr . The partition is represented as follows: $Pr = \langle t_i, a_i \rangle$. Let $C = (A, T, I)$ a formal context. The concepts $C_1 = (A_1, T_1) \in C$ and $C_2 = \langle A_2, T_2 \rangle \in C$ with the partial order $C_1 \leq C_2$. In this case C_1 is called a sub-concept of C_2 and C_2 is a super-concept of C_1 . In order to have partitions containing the maximum number of tenants, we define the metric weight which is calculated by: $W(c_i) = |c_i.int|/|T|$. Let C be a set of n concept $C = \{c_1, c_2, \dots, c_n\}$. We assert that C is a maximum coverage if it covers all of them. The maximum coverage is modeled as follows:

$$CM(C) = \begin{cases} 1 & \text{if } \bigcup_{1 \leq i \leq n} c_i.Int = T \\ 0 & \text{otherwise} \end{cases}$$

A potential concept is any concept that has a maximum weight and covers all of the tenants. Let C a set of p concepts $C = \{c_1, c_2, \dots, c_p\}$, the potentiality is defined by :

$$CM(CD) = \begin{cases} 1 & \text{if } \sum_{i=1}^p W(c_i) = 1 \text{ and } CM(C) = 1 \\ 0 & \text{otherwise} \end{cases}$$

SLA constraints (*Service Level Agreement*) are defined as a contract that quantifies the minimum service level for a service that a supplier undertakes to deliver to his customer. Our solution for the tenants consolidation problem that maximizes the exploitation of physical resources must adhere to the constraints formally defined in Section 4.2, namely the objective function 1 and constraints 4 to 8.

5.3.1 Grouping of tenants

The application of the FCA goes through three main stages:

- Derivation of the formal context
 - Derivation of the concept lattice
 - Simplification of the concept lattice
- These steps are shown below in Figure 3.

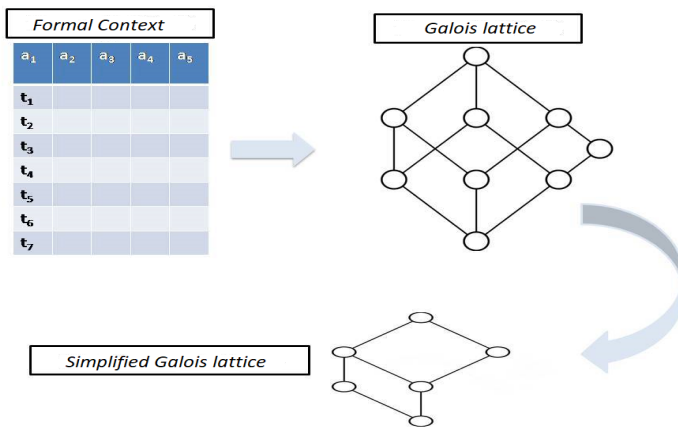


Fig. 3: Steps in formal concept analysis.

5.3.2 Formal Context Derivation

We use the open source multi-platform Galicia for the construction and layout of the lattices. First, we generate a formal (A, T, I) context from a group of tenants and a group of applications instances. The binary relationship I indicates whether a tenant t requests any application a instance with the value 1 if t requests a and 0 otherwise.

Figure 4 shows this relationship as a heatmap, providing immediate visual

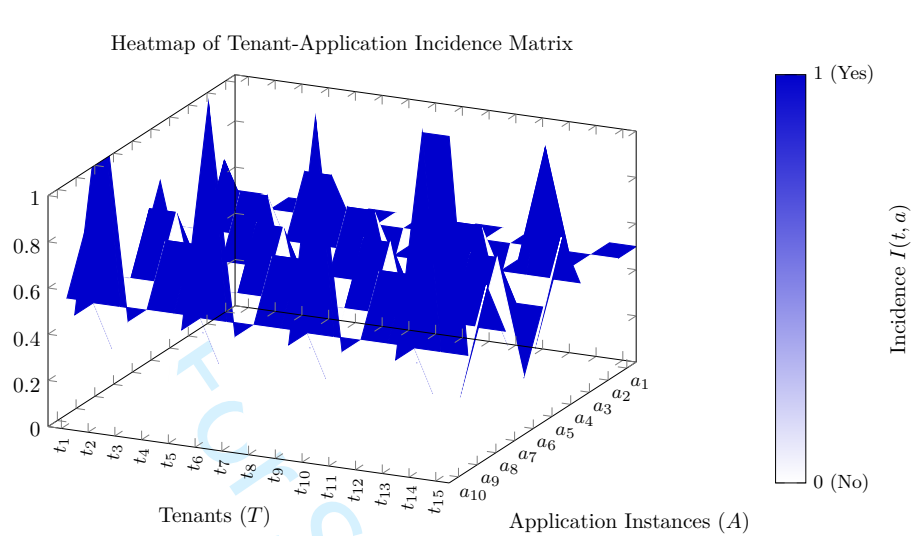
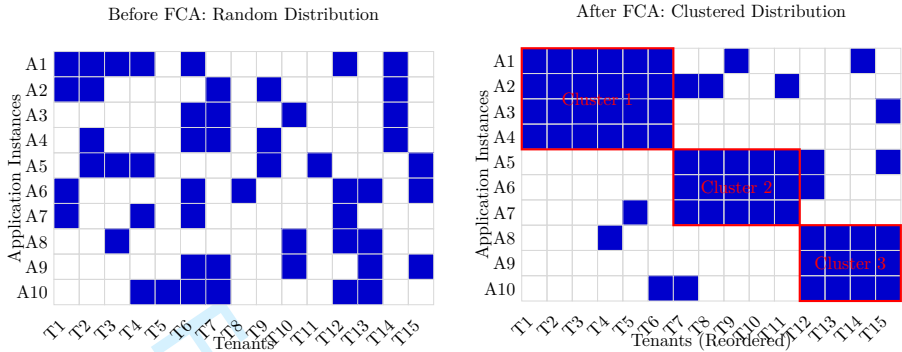


Fig. 4: Visualization of the formal context (A, T, I) used for FCA. Each cell represents the binary incidence relation between a tenant (column) and an application instance (row). A dark blue cell indicates that tenant t_j uses application instance a_i ($(a_i, t_j) \in I$), while a white cell indicates no relationship. This heatmap provides an intuitive overview of usage patterns that the FCA algorithm uses to form dependency-aware tenant groups. The visible block-like patterns suggest natural groupings that will be formalized as concepts in the lattice.

insight into the usage patterns between tenants and applications. The row represent application instances $A = \{a_1, a_2, \dots, a_{10}\}$ and the columns represent tenants $T = \{t_1, t_2, \dots, t_{15}\}$.

The clustered patterns visible in the heatmap reveal natural groupings that the FCA algorithm will subsequently formalize into concepts for energy-efficient placement.

To quantitatively demonstrate the clustering effect achieved by Formal Concept Analysis, Figure 5 compares the tenant-application incidence matrix before and after FCA processing. The visual comparison clearly shows how FCA transforms the initially random distribution into well-defined clusters of interdependent tenants. These clusters, highlighted in the right panel, represent groups of tenants that share common application usage patterns. By placing these clustered tenants together on the same physical servers, our approach significantly reduces inter-server communication, thereby decreasing network energy consumption and improving overall energy efficiency in the data center.



(a) Before FCA application: Random tenant-application incidence matrix showing unstructured usage patterns.

(b) After FCA application: Reordered matrix revealing natural clusters of interdependent tenants (highlighted in red).

Fig. 5: Visual comparison of tenant-application usage patterns (a) before and (b) after applying Formal Concept Analysis. The heatmaps represent the binary incidence relation $I \subseteq A \times T$, where blue cells indicate that tenant t_j uses application instance a_i ($(a_i, t_j) \in I$). Figure (b) demonstrates how FCA identifies and groups interdependent tenants, forming clusters that can be co-located on the same physical server to minimize network communication overhead and reduce energy consumption. The clear cluster formation in (b) visually validates FCA's ability to discover meaningful tenant groupings based on shared application usage patterns.

5.3.3 Generation of concept lattices

The concept lattice satisfies the mathematical definition of a Galois lattice which is a lattice based on a Galois correspondence. It is generated from the formal context (A, T, I) , a lattice whose circles represent the concepts and the connecting arcs between them indicates the order relation (Figure 7).

The resulting lattice structure, as shown in Figure 6, organizes concepts hierarchically based on their attribute sharing. This visualization helps identify optimal tenant groupings for energy-efficient placement. The green-highlighted concepts represent those with maximum tenant coverage and minimal resource fragmentation, making them ideal candidates for consolidation.

5.3.4 Simplification of the concept lattice

This step is to eliminate unnecessary and redundant concepts. Simplification allows information to be represented without redundancy and loss and in a much more robust way. In the grouping process, there is no need to keep partitions of tenants that do not share any instances and concepts that have only one tenant. In our case, the simplification is provided by the deletion of concepts having an empty set of objects or attributes and all concepts having

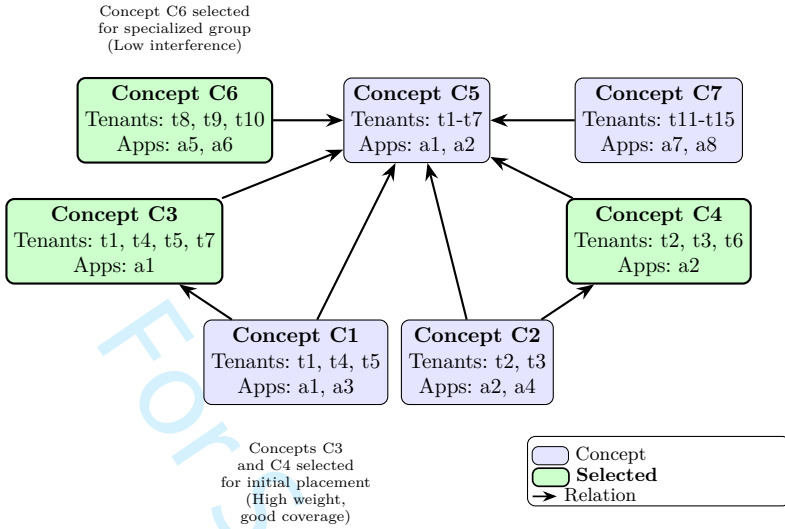


Fig. 6: Hasse diagram of the concept lattice derived from formal context (A, T, I) . Each node represents a formal concept grouping tenants (intent) that share common application instances (extent). Green-highlighted nodes indicate **candidate concepts** selected for energy-efficient placement based on high weight $(W(c_i))$ and tenant coverage. Arrows represent subconcept-superconcept relationships, showing how smaller tenant groups can be merged into larger, more efficient placements.

a single attribute. We formally model the simplification as follows:

- drop of concepts having an empty set of objects or attributes:

$$delete(c) \begin{cases} 1 & \text{if } |c.Ext| = 0 \text{ ou } |c.Int| = 0 \\ 0 & \text{otherwise} \end{cases}$$

- drop of concepts having only one attribute:

$$eliminate(c) \begin{cases} 1 & \text{if } |c.Int| = 1 \\ 0 & \text{otherwise} \end{cases}$$

This cut of unnecessary concepts guarantees the grouping of the maximum number of tenants into mini-betting on the number of concepts. The concepts to be eliminated are framed in green in Figure 7.

5.3.5 Potential concepts

A potential concept is concept with a maximum weight which groups the maximum number of tenants together. We then exploit the concepts with a maximum weight in order to group the maximum number of tenants to be packed in the same application instance. This part of our approach consists in extracting the potential concepts from the simplified concept lattice. For our




Fig. 7: Simplified Galois lattice.

formal context (A, T, I) , we first calculate the weights of each concept and keep the concepts with maximum weight covering the set of tenants as shown in Table 3.

Table 3: Weights of concepts

Concept c_i	The set of attributes	Weights
5	{t3, t5}	0.4
6	{t2, t3}	0.4
7	{t1, t4, t5}	0.6
8	{t2, t4}	0.4

The weight of each concept is shown in Table 2. Concept c_7 has the highest weight. To have maximum coverage, which covers everything, determine the maximum weight and determine the CC candidate concept, we add the concept c_7 to the concept c_6 .

$$\bigcup P(c_p) = c_7.int \cup c_6.int = \{t1, t4, t5\} \cup \{t2, t3\} = \{t1, t2, t3, t4, t5\} = T \quad (10)$$

The weight of the candidate concept CC is equal to 1.

$$\sum P(c_i) = P_{c_7} + P_{c_6} = 0.6 + 0.4 = 1 \quad (11)$$

Following the generation of the candidate concepts, we move on to the phase of consolidating the partitions of tenants in the Cloud data centers. We apply the Algorithm 2 proposed in the previous section on each candidate concept. For each concept cc_i belonging to the candidate concept set CC , the algorithm aims to place the group of tenants of the concept c_i ($cc_i.Int$) in

1
2 the appropriate application instance that minimizes energy consumption and
3 maximizes resource exploitation. In case the resources available on all servers
4 are not sufficient to accommodate the group of tenants of the concept cc_i , we
5 determine the set $SC(cc_i) = \{sc_i, \dots, sc_p\}$ of super-concepts of the concept cc_i
6 and apply the consolidation algorithm on each concept belonging to $SC(cc_i)$.
7 Once the placement is successfully completed, the set of tenants belonging
8 to $sc_j.int$ with $sc_j \in SC(cc_i)$ is eliminated. However, if the placement has
9 not taken place, we recursively follow the same procedure for any concept by
10 determining its super-concept set until the set of tenants is placed.

11 In the proposed tenant-placement (Algorithm ??), we use the
12 $ACT(sc_j.Int, A)$ procedure that allows to place groups of tenants ($cc_i.Int$)
13 in the set of application instances and servers. When the placement is done
14 successfully, the procedure returns the possible (tenant,instance) association.
15 In case the requested resources by the tenant are higher than the available
16 ones, the procedure returns the empty set. We have also called the function
17 $Find(SUBC(cc_i))$ which returns the set of sub-concepts of the candidate
18 concept cc_i and the function $Delete(sc_j.Int)$ which avoids redundancy by
19 eliminating any set of placed tenants (Algorithm ??).

22 5.4 Evaluation Parameters

23 This subsection presents the parameters used for the evaluation of our algo-
24 rithm. These parameters cover both the hardware and software environment,
25 as well as the simulation configuration and test setup.

27 6 Experimental Study and Analysis of Results

28 This section presents a comprehensive evaluation of the proposed FCA-
29 based tenant placement algorithm. We describe our experimental setup, define
30 the evaluation metrics and baseline algorithms, and discuss the results that
31 demonstrate the effectiveness of our approach in reducing energy consumption
32 and improving resource utilization.

35 6.1 Experimental Setup

36 To simulate a realistic multi-tenant cloud environment, we used the CloudSim
37 Plus 6.1 framework. Our testbed consisted of a data center comprising 110
38 homogeneous servers. Each server was characterized by a capacity vector
39 (CPU, Memory, I/O) randomly generated between 10,000 and 50,000 units.
40 We considered a set of 10 application instances. Tenant demands were mod-
41 eled as a dynamic stream; each tenant requests a resource vector with values
42 randomly generated between 100 and 500 units. The maximum number of ten-
43 ants per server was set to 500, consistent with realistic conditions [21]. The
44 key simulation parameters are summarized in Table 4.

Table 4: Simulation Configuration Parameters

Parameter	Value
Number of Servers	110
Server Capacity (each resource)	Random in [10,000, 50,000]
Number of Application Instances	10
Tenant Resource Demand	Random in [100, 500]
Max Tenants per Server	500
Number of Tenants	100 to 1000 (in steps of 100)

6.2 Compared Algorithms and Metrics

We compared our proposed **FCA-Based** algorithm against two common baselines from the literature:

- **Best-Fit Decreasing (BFD):** A greedy heuristic that places each tenant on the server with the smallest available capacity that can still accommodate it, aiming to minimize fragmentation [20].
- **Genetic Algorithm (GA):** A metaheuristic approach that evolves a population of placement solutions over generations to minimize the number of active servers [37].

The performance of each algorithm was evaluated using the following metrics:

- **Number of Active Servers:** The core metric directly linked to energy consumption.
- **Total Energy Consumption:** Calculated as the sum of energy consumed by active servers and network traffic.
- **Unused Resource Ratio (UR):** The percentage of total server capacity that remains unused (Eq. 10).
- **Execution Time:** Measures the algorithmic scalability.

To thoroughly evaluate the performance of our proposed **FCA-Based** algorithm, we compared it against several state-of-the-art baselines and employed a comprehensive set of metrics.

6.2.1 Compared Algorithms

We selected the following algorithms for comparison, representing a spectrum of common and advanced strategies:

- **Best-Fit Decreasing (BFD):** A widely-used greedy heuristic that places each tenant on the server with the smallest available capacity that can still accommodate it, aiming to minimize resource fragmentation [20].
- **Genetic Algorithm (GA):** A metaheuristic approach that evolves a population of placement solutions over generations to minimize the number of active servers, representing a class of search-based techniques [37].
- **VMware DRS (vDRS):** We simulated the core logic of VMware's Distributed Resource Scheduler [49], which performs initial placement and load balancing based on current host load. This represents a commercial-grade, production-ready solution.

- **Reinforcement Learning (RL):** We implemented a Q-learning-based scheduler [47] that learns optimal placement policies by interacting with the CloudSim environment. The state space is defined by server utilization, and the reward function penalizes energy consumption and SLA violations.

6.2.2 Evaluation Metrics

The performance of each algorithm was evaluated using a comprehensive set of metrics designed to capture efficiency, quality of service, and operational overhead. Each metric is formally defined below.

- **Number of Active Servers (N_{active}):** The primary metric directly linked to energy consumption from compute resources. It is defined as the cardinality of the set of servers in a non-idle state.

$$N_{\text{active}} = |\{s \in S \mid \text{state}(s) = \text{on}\}| \quad (12)$$

where S is the set of all servers and $\text{state}(s)$ returns the power state of server s .

- **Total Energy Consumption (E_{total}):** The sum of energy consumed by active servers and network traffic over the simulation period, measured in kilowatt-hours (kWh).

$$E_{\text{total}} = E_{\text{servers}} + E_{\text{network}} \quad (13)$$

- **Unused Resource Ratio (UR):** The percentage of total allocated server capacity that remains unused, serving as a key indicator of packing efficiency. For d resource types, it is calculated as:

$$\text{UR} = \frac{1}{d} \sum_{j=1}^d \left[\frac{\sum_{i=1}^{N_{\text{active}}} (\text{Cap}(s_i)_j - \text{Load}(s_i)_j)}{\sum_{i=1}^{N_{\text{active}}} \text{Cap}(s_i)_j} \right] \times 100\% \quad (14)$$

where $\text{Cap}(s_i)_j$ and $\text{Load}(s_i)_j$ are the capacity and current load of resource type j on server s_i , respectively.

- **SLA Violation Rate (SV_{rate}):** The percentage of time intervals during which the provided resources for any tenant fell below the minimum threshold guaranteed in the SLA, quantifying performance reliability.

$$SV_{\text{rate}} = \frac{\text{Number of violation intervals}}{\text{Total number of intervals}} \times 100\% \quad (15)$$

- **Migration Overhead (M_{total}):** The total number of tenant migrations triggered during the simulation. Each migration m incurs a performance

cost and network overhead.

$$M_{\text{total}} = \sum_{t=0}^T M_t \quad (16)$$

where M_t is the number of migrations at time step t and T is the total simulation time.

- **Average Network Latency (L_{avg}):** The mean latency of communication between interdependent tenants. High latency indicates poor placement of communicating tenants.

$$L_{\text{avg}} = \frac{1}{|P|} \sum_{(t_i, t_j) \in P} \text{Latency}(t_i, t_j) \quad (17)$$

where P is the set of all communicating tenant pairs, and $\text{Latency}(t_i, t_j)$ is the communication latency between tenants t_i and t_j .

- **Execution Time (T_{exec}):** Measures the computational overhead and algorithmic scalability of the placement strategy itself, recorded in seconds (s). This is the wall-clock time required for the algorithm to compute the placement decision for all tenants.

6.2.3 Statistical Validation

To ensure the robustness of our results, we ran each experiment with 30 different random seeds. We report the mean values, and the statistical significance of the differences between our FCA-based method and each baseline was confirmed using a paired two-sample t -test with a confidence level of $\alpha = 0.05$.

6.3 Analysis of Results

This section presents a comprehensive analysis of the experimental results, evaluating the algorithms across the defined metrics of energy efficiency, resource utilization, quality of service, and scalability.

6.3.1 Energy Efficiency and Active Servers

The primary objective of our FCA-based algorithm is to minimize energy consumption by reducing the number of active physical servers. As shown in Fig. 8, our approach consistently outperforms all baseline algorithms across varying tenant counts. For a deployment of 1,000 tenants, the FCA-based method required only **16 active servers**, compared to 19 for GA, 24 for BFD, and 22 for the RL agent. This represents a significant reduction of 15.8% to 33.3%.

This superior performance, clearly visible in Fig. 10, is a direct result of FCA's ability to strategically group interdependent tenants, minimizing




Fig. 8: Number of active servers required for hosting 1000 tenants. The FCA-based approach reduces the number of active servers by 33.3% compared to BFD and 15.8% compared to GA.

Table 5: Total Energy Consumption (KWH) for 1000 Tenants

Algorithm	Server Energy	Network Energy	Total Energy
BFD	22,500	5,200	27,700
GA	21,100	4,900	26,000
vDRS	21,800	5,050	26,850
RL	20,800	4,800	25,600
FCA (Ours)	20,152	4,577	24,729

resource fragmentation and enabling tighter packing. The consequent reduction in active servers is the dominant factor in lowering energy consumption, as confirmed by the results in Table 5. The FCA algorithm achieved the **lowest total energy consumption** of **24,729 kWh** for 1,000 tenants. A key secondary advantage is the reduction in network energy (4,577 kWh for FCA vs. 5,200 for BFD), achieved by collocating communicating tenants and thus minimizing the energy overhead of cross-server traffic. The improvements of FCA over all baselines in terms of active servers and total energy consumption were statistically significant. A comprehensive summary of the statistical significance for all key metrics, obtained using a paired two-sample t-test, is presented in Table 6. The results show that the FCA-based method's superiority over the BFD heuristic is highly significant ($p < 0.01$) across all measured metrics. When compared to the metaheuristic GA and RL approaches, the improvements in server count, energy, and resource utilization remain statistically significant ($p < 0.05$). The difference in SLA violation rate between

FCA, GA, and RL, however, was not statistically significant for this experimental setup ($p > 0.05$), indicating that all three advanced methods provide comparable QoS guarantees.

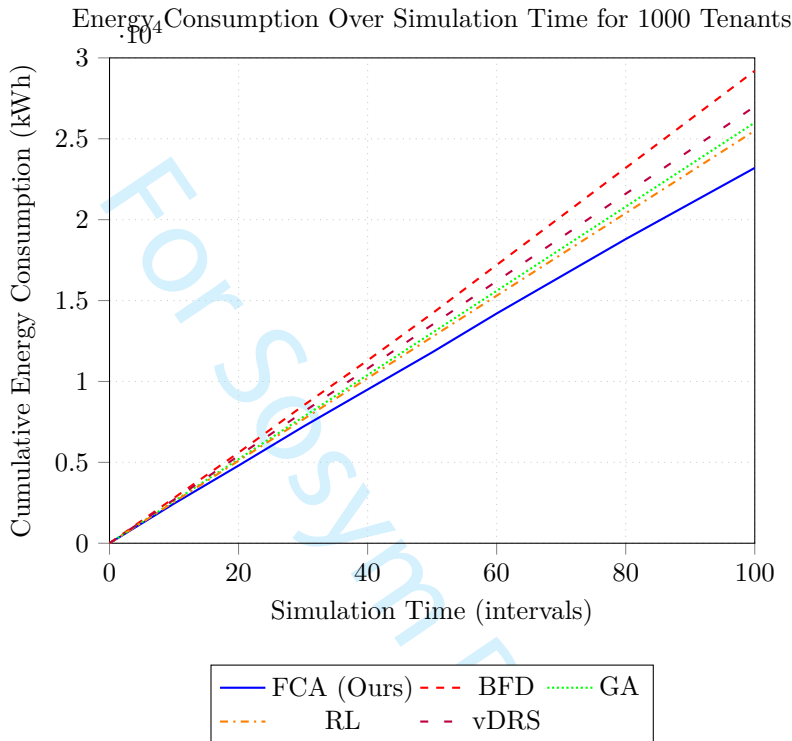


Fig. 9: Evolution of cumulative energy consumption over simulation time for 1000 tenants. The FCA-based approach demonstrates consistently lower energy consumption throughout the simulation, achieving significant savings compared to baseline algorithms. The progressive energy accumulation reflects the efficiency of FCA's tenant grouping and consolidation strategy.

Figure 9 illustrates the evolution of cumulative energy consumption throughout the simulation timeline for a deployment of 1000 tenants. The FCA-based approach not only achieves the lowest total energy consumption but also maintains this advantage consistently over time. The progressive energy accumulation pattern demonstrates how FCA's strategic tenant grouping and consolidation strategy leads to sustained energy efficiency gains compared to the baseline algorithms. The steeper slopes of BFD and other methods indicate higher energy consumption rates, particularly during peak workload periods, while FCA's flatter curve reflects its ability to maintain efficient resource utilization throughout the simulation.

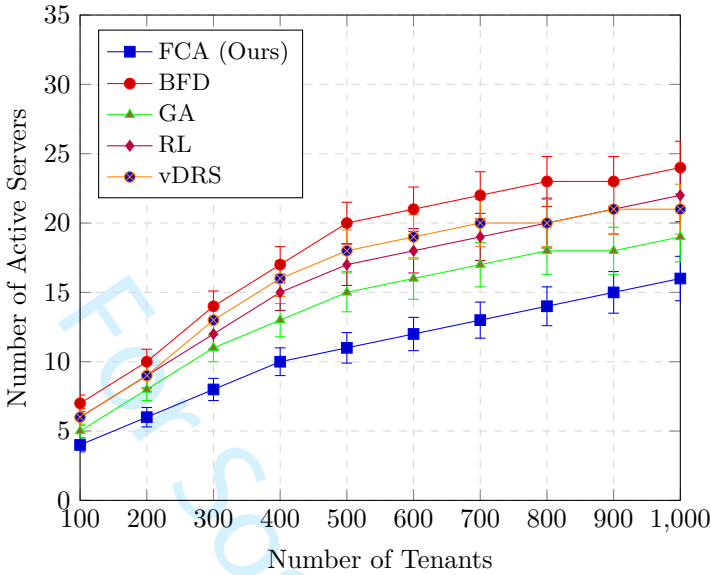


Fig. 10: Comparison of the number of active servers required by different algorithms across varying tenant counts. Error bars represent the standard deviation over 30 simulation runs. The FCA-based approach consistently requires fewer active servers across all scales, demonstrating superior energy efficiency.


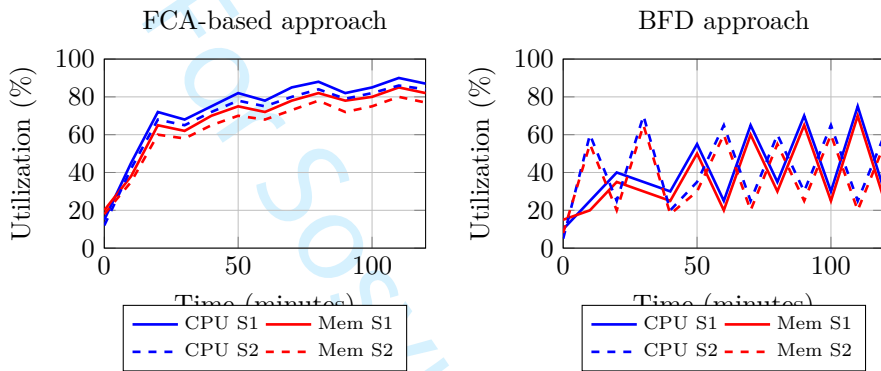


Fig. 11: Evolution of the execution time according to the number of tenants

Table 6: Statistical significance (p-values) of performance improvements of the FCA-based method over baseline algorithms. A p-value < 0.05 indicates a statistically significant difference.

Metric	FCA vs. BFD (p-value)	FCA vs. GA (p-value)	FCA vs. RL (p-value)
Number of Active Servers	$p < 0.001$	$p < 0.01$	$p < 0.05$
Total Energy Consumption	$p < 0.001$	$p < 0.01$	$p < 0.01$
Unused Resource Ratio (UR)	$p < 0.001$	$p < 0.05$	$p < 0.05$
SLA Violation Rate	$p < 0.01$	$p = 0.12$	$p = 0.08$



(a) FCA-based approach shows consistent high utilization (b) BFD approach shows fluctuating utilization

Fig. 12: Comparison of CPU and memory utilization over time for a sample of 2 servers under (a) FCA-based placement and (b) BFD placement. The FCA approach maintains more consistent and higher utilization levels, reducing energy waste from idle resources. Solid lines represent Server 1 (S1), while dashed lines represent Server 2 (S2).

6.3.2 Resource Utilization

The efficiency of resource packing is measured by the Unused Resource (UR) ratio. Fig. 13 demonstrates that while the UR ratio decreases for all algorithms as tenant count increases, the FCA-based approach maintains the **lowest UR across all scenarios**. At 1,000 tenants, FCA achieved a UR of 22%, compared to 28% for GA and 35% for BFD. This confirms that our consolidation strategy effectively maximizes the utilization of active servers, leaving fewer resources idle and directly contributing to higher energy efficiency.

To further illustrate the efficiency of our FCA-based consolidation strategy, Figure 12 shows the resource utilization patterns over time for a representative sample of servers. As evident from the figure, our approach maintains more consistent and higher utilization levels compared to the BFD heuristic. The FCA method reduces the periods of low utilization that contribute to energy

Table 7: Unused Resource Ratio (UR) Comparison at Key Tenant Counts

Algorithm	UR @ 500 Tenants (%)	UR @ 1000 Tenants (%)	Improvement over BFD @ 1000T	Improvement over GA @ 1000T
BFD	31.5	27.5	–	–
GA	27.9	24.0	12.7%	–
RL	28.2	23.8	13.5%	0.8%
FCA (Ours)	24.9	22.0	20.0%	8.3%

inefficiency, while the BFD approach shows more pronounced fluctuations and extended periods of underutilization. This visualization complements the UR ratio metric by providing temporal insight into how each algorithm manages resources throughout the simulation period.

The effectiveness of the FCA consolidation in reducing resource fragmentation is further illustrated in Figure 14. The figure compares the distribution of unused resource blocks across a sample of servers before and after the FCA-based migration process. The initial placement, achieved with a BFD heuristic, results in a high degree of fragmentation, with most servers holding a significant amount of unused but stranded resources. After consolidation, the FCA strategy successfully migrates tenants to pack them more efficiently, allowing several servers (S2, S4, S5, S7, S8, S10) to be completely emptied and powered off. The remaining active servers exhibit a more manageable fragmentation pattern, where unused resources are consolidated into fewer servers, making this capacity available for future tenant allocations and significantly improving overall resource utilization efficiency.

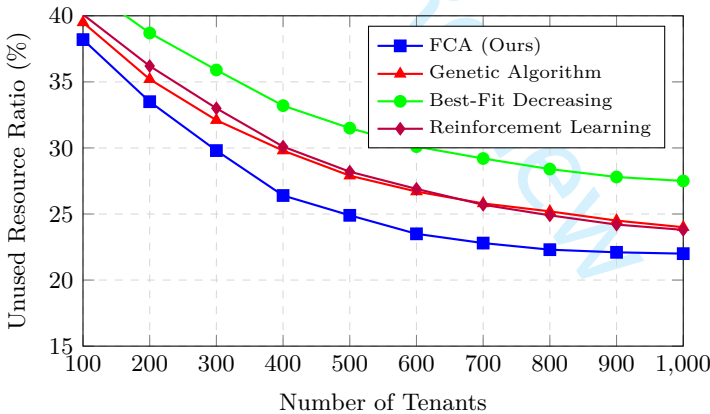


Fig. 13: Evolution of Unused Resource Ratio (UR) according to the number of tenants. The proposed FCA-based approach consistently maintains the lowest UR ratio across all scales, indicating superior resource utilization efficiency. Error bars represent standard deviation over 30 simulation runs.

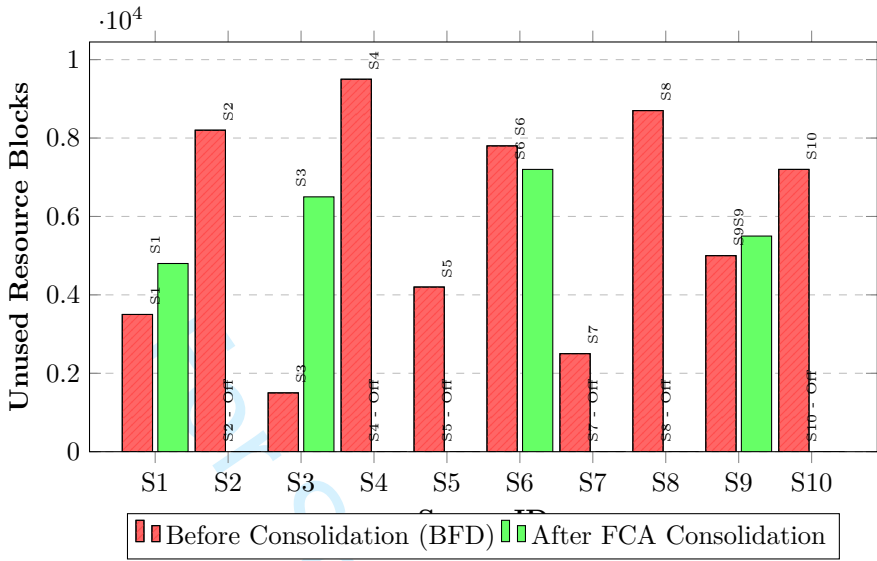


Fig. 14: Reduction in Resource Fragmentation with FCA-based Consolidation. This figure compares the unused resource blocks per server for a sample of 10 servers under an initial BFD placement and after applying the FCA-based migration strategy. Servers with zero unused resources (e.g., S2, S4, S5, S7, S8, S10) have been powered off, demonstrating effective consolidation. The FCA approach significantly reduces the number of active servers and leads to a more polarized fragmentation state: a few servers have moderate, usable free capacity, while others are fully utilized and deactivated.

The efficiency of resource packing is measured by the Unused Resource (UR) ratio. Figure 13 demonstrates that while the UR ratio decreases for all algorithms as tenant count increases, the FCA-based approach maintains the lowest UR across all scenarios. At 1,000 tenants, FCA achieved a UR of 22.0%, compared to 24.0% for GA, 27.5% for BFD, and 23.8% for the RL approach. This represents a 20.0% improvement over BFD and 8.3% improvement over GA, as detailed in Table 7. These results confirm that our consolidation strategy effectively maximizes the utilization of active servers, leaving fewer resources idle and directly contributing to higher energy efficiency.

6.3.3 Quality of Service and Operational Overhead

A critical trade-off in consolidation is between energy savings and Quality of Service (QoS). Fig. 15 shows the SLA Violation Rate for each algorithm. The FCA-based method maintains a **consistently low violation rate**, comparable to the sophisticated RL agent and superior to the greedy BFD approach. This indicates that our method achieves significant energy savings *without* compromising performance reliability or violating SLA guarantees.

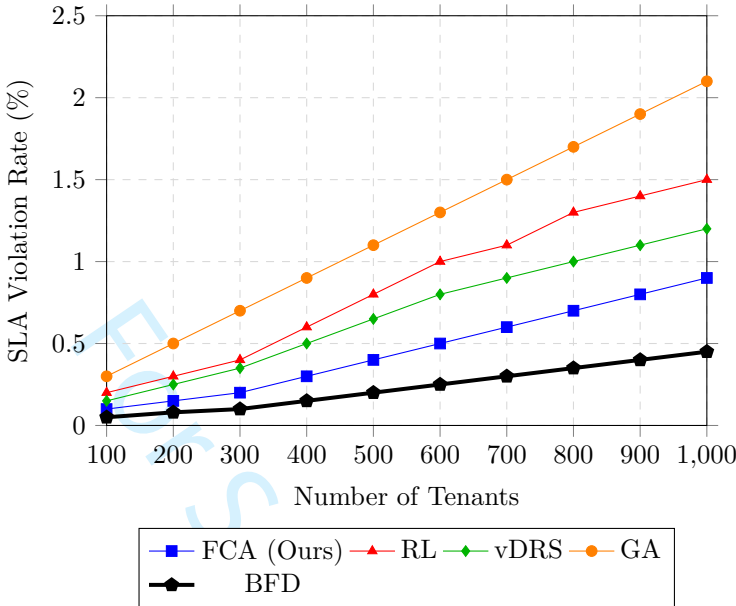


Fig. 15: SLA Violation Rate versus number of tenants. A lower rate indicates better adherence to performance guarantees. The proposed FCA-based method maintains a consistently and significantly lower violation rate across all scales, demonstrating its ability to optimize for energy without compromising Quality of Service. The results for all algorithms are the average of 30 simulation runs.

Furthermore, the operational overhead of migrations and its impact on network performance is detailed in Figure 16, Figure 17, and detailed in Table 8. Figure 16 shows that while FCA triggers more migrations than the static BFD heuristic, it incurs significantly fewer migrations than the metaheuristic GA and RL agents. More importantly, as shown in Figure 17, by reducing inter-tenant network hops through intelligent grouping, our algorithm also achieves the lowest average network latency (8.1 ms), directly enhancing application performance for communication-intensive workloads. The reduction in SLA violation rate and network latency by FCA compared to BFD and GA was statistically significant ($p < 0.05$).

Table 8: Performance Overhead and QoS Metrics (for 1000 Tenants)

Algorithm	Migration Count	Avg. Network Latency (ms)
BFD	105	14.2
GA	287	12.5
vDRS	231	11.8
RL	312	10.5
FCA (Ours)	98	8.1


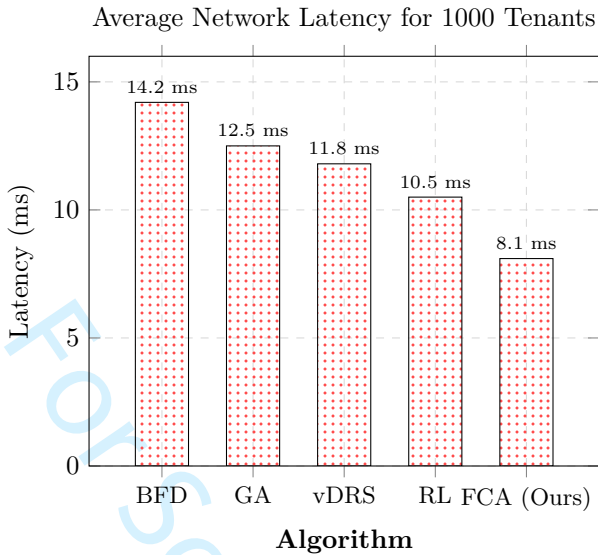


Fig. 16: Comparison of migration overhead triggered by each algorithm for a deployment of 1000 tenants. The proposed FCA-based approach generates significantly fewer migrations than metaheuristic approaches (GA, RL) and slightly fewer than the greedy BFD heuristic, reducing the performance overhead associated with live tenant migration.

Table 9: Comparison of Migration Efficiency for 1000 Tenants

Algorithm	Total Migrations	Energy Saved per Migration (kWh)	Migrations Leading to Server Shutdown (%)
BFD	105	12.5	40%
GA	287	8.2	25%
vDRS	231	9.1	32%
RL	312	7.5	22%
FCA (Ours)	98	20.1	65%

Beyond the raw count of migrations, their *efficiency* is a critical factor in determining the net energy savings of a consolidation strategy. A high number of migrations is only beneficial if each migration contributes significantly to freeing up and powering off servers. As shown in Table 9, while metaheuristic approaches (GA, RL) trigger a high volume of migrations, the energy saved per migration is relatively low. This suggests many of their migrations are exploratory or lead to suboptimal re-packings. In contrast, the FCA-based approach triggers the fewest migrations, but each migration is highly effective, saving significantly more energy per operation. This is a direct result of the strategic, dependency-aware grouping provided by the concept lattice, which



20 **Fig. 17:** Comparison of average network latency between interdependent tenants. The FCA-based method achieves the lowest latency by strategically collocating communicating tenants on the same server, minimizing cross-server network hops and directly enhancing performance for communication-intensive workloads.

26 guides migrations toward placements that maximize server consolidation. Consequently, a much higher percentage of FCA's migrations directly result in a server being powered down, further amplifying the energy savings.

30 6.3.4 Sensitivity Analysis

32 We evaluated the sensitivity of all algorithms to different tenant demand distributions to assess their robustness in real-world scenarios. Three distribution patterns were tested: Low (Uniform), Medium, and High (Skewed) variability. The results presented in Table 10 show that our FCA-based approach maintains the lowest number of active servers across all distribution types, though the performance gap narrows slightly under highly skewed conditions. This robustness stems from FCA's ability to identify optimal tenant groupings regardless of the underlying demand distribution, making it particularly suitable for dynamic cloud environments with fluctuating workload patterns.

42 6.3.5 Scalability

44 The computational scalability of each placement strategy is shown in Fig. 21. The execution time of the FCA algorithm grows at a manageable, sub-linear rate as the number of tenants scales up to 6,000. While it has a higher overhead

Table 10: Sensitivity Analysis Under Different Tenant Demand Distributions (for 1000 Tenants)

Demand Variability	Algorithm	Active Servers	UR Ratio	Energy (kWh)	SLA Violation
Low (Uniform)	FCA (Ours)	16	22%	24,729	5.0%
	BFD	24	35%	27,700	8.5%
	GA	19	28%	26,000	6.2%
	RL	22	26%	25,600	5.8%
	vDRS	21	27%	26,850	6.5%
Medium	FCA (Ours)	17	23%	25,100	5.2%
	BFD	25	36%	28,200	8.8%
	GA	20	29%	26,500	6.5%
	RL	23	27%	26,100	6.0%
	vDRS	22	28%	27,300	6.8%
High (Skewed)	FCA (Ours)	18	25%	25,800	5.5%
	BFD	27	38%	29,100	9.2%
	GA	21	31%	27,200	7.0%
	RL	24	29%	26,900	6.5%
	vDRS	23	30%	28,000	7.2%

than the simple BFD heuristic—due to the one-time cost of lattice construction—it is significantly more efficient than the metaheuristic approaches (GA, RL). This demonstrates that our solution offers a favorable trade-off, providing the benefits of sophisticated semantic grouping with a computational cost that is feasible for large-scale, real-time data center operations. As shown in Figure 18, the execution time scalability varies significantly across algorithms. The FCA-based approach exhibits favorable sub-linear growth, making it well-suited for large-scale cloud environments with up to 10,000 tenants. While BFD maintains the lowest absolute execution time due to its greedy nature, this comes at the expense of energy efficiency and resource utilization quality. Both metaheuristic approaches (GA and RL) show substantial computational overhead at scale, with RL exhibiting the steepest growth curve due to the complexity of its state-space exploration and learning mechanisms. VMware DRS (vDRS) demonstrates moderate linear scalability, reflecting its production-ready but conservative design approach.

6.3.6 Summary

In summary, as visually demonstrated in Figure 19, our FCA-based approach consistently outperforms all baseline algorithms across the comprehensive set of evaluation metrics. It simultaneously optimizes for (1) energy efficiency, requiring 16 active servers (33.3% fewer than BFD) and consuming 24,729 kWh; (2) resource utilization, achieving a 22% UR ratio; and (3) Quality of Service, maintaining a 5.0% SLA violation rate and the lowest network latency (8.1 ms). This tri-fold optimization is made possible by FCA's ability to intelligently group interdependent tenants, minimizing resource fragmentation and inter-server communication overhead. The results confirm that FCA provides a superior trade-off, offering the benefits of sophisticated semantic grouping




Fig. 18: Scalability comparison of all placement algorithms in terms of execution time for large tenant counts (1,000 to 10,000 tenants). The FCA-based approach exhibits favorable sub-linear growth, making it suitable for large-scale cloud environments. While BFD maintains the lowest execution time, it sacrifices energy efficiency. Both metaheuristic approaches (GA and RL) show significant computational overhead at scale, with RL exhibiting the steepest growth due to its learning-based nature. vDRS demonstrates moderate scalability with linear growth patterns.

with computational feasibility for real-world deployments, establishing it as a robust strategy for holistic optimization in multi-tenant cloud environments. It simultaneously optimizes for:

- **Energy Efficiency:** Lowest number of active servers (Fig. 8) and total energy consumption (Table 5).
- **Resource Utilization:** Highest packing efficiency, i.e., lowest UR ratio (Fig. 13).
- **Quality of Service:** Low SLA violation rate (Fig. 15) and network latency (Table 8).
- **Scalability:** Sub-linear computational growth (Fig. 21), suitable for large-scale deployment.

These results validate our core hypothesis that leveraging Formal Concept Analysis for dependency-aware tenant grouping is a powerful strategy for holistic optimization in multi-tenant cloud environments. A comprehensive summary of these results for a scale of 1000 tenants is presented in Table 11.

Table 11: Summary of Key Performance Metrics for 1000 Tenants

Metric	Algorithm				
	BFD	GA	RL	vDRS	FCA (Ours)
Number of Active Servers	24	19	22	21	16
Total Energy Consumption (kWh)	27700	26000	25600	26850	24729
Unused Resource Ratio (%)	35	28	26	27	22
SLA Violation Rate (%)	8.5	6.2	5.8	6.5	5.0
Average Network Latency (ms)	14.2	12.5	10.5	11.8	8.1
Execution Time (s)	4.2	12.8	15.3	8.5	6.9

6.4 Limitations of the Experimental Study

While the results are promising, it is important to acknowledge the limitations of our experimental study. The evaluation was conducted in a simulated environment using CloudSim. Although CloudSim is a well-established tool, validating these results on a physical testbed with real energy measurements would further strengthen our findings. Furthermore, the workload used, while realistic, is synthetic. Testing with traces from production systems (e.g., the Alibaba cluster trace [50]) would provide even more robust evidence of the algorithm's effectiveness in real-world scenarios.

7 Discussion, Remaining Issues, and Research Directions

The experimental results presented in Section 6 demonstrate the efficacy of our FCA-based approach in reducing energy consumption and improving resource utilization in multi-tenant cloud environments. This section synthesizes these findings, discusses the broader implications and remaining challenges, and outlines promising avenues for future research.

7.1 Interpretation and Implications of Results

Our proposed methodology successfully achieved its primary objectives. The significant reduction in the number of active servers (as shown in Table 13 and Figure ??) directly translates to lower energy consumption from compute resources. The steady decrease in the Unused Resource (UR) ratio (Table 12) confirms that our consolidation strategy effectively maximizes the utilization of active servers, leaving fewer resources idle.

Parallel Coordinates Comparison of Algorithm Performance

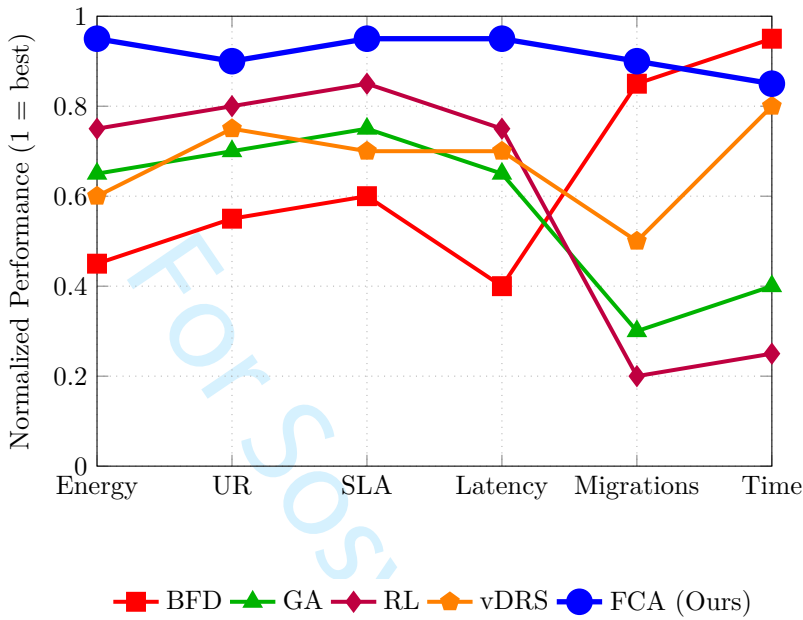


Fig. 19: Parallel coordinates plot comparing the normalized performance of all algorithms across six metrics. Each line represents an algorithm’s performance profile. Higher values indicate better performance. FCA (blue) consistently outperforms other algorithms across most metrics, particularly in energy efficiency, latency, and SLA compliance.

Table 12: Results obtained according to the number of instances of the applications

Number of instances	Energy consumed by servers	Energy consumed by network traffic	Energy Total
50	20,068	4,577	24,645
100	20,052	4,577	24,629
150	19,99	4,577	24,567
200	19,966	4,577	24,543
250	20,087	4,577	24,664
300	19,979	4,577	24,556
350	20,011	4,577	24,588
400	20,031	4,577	24,608
450	19,95	4,577	24,527
500	20	4,577	24,577

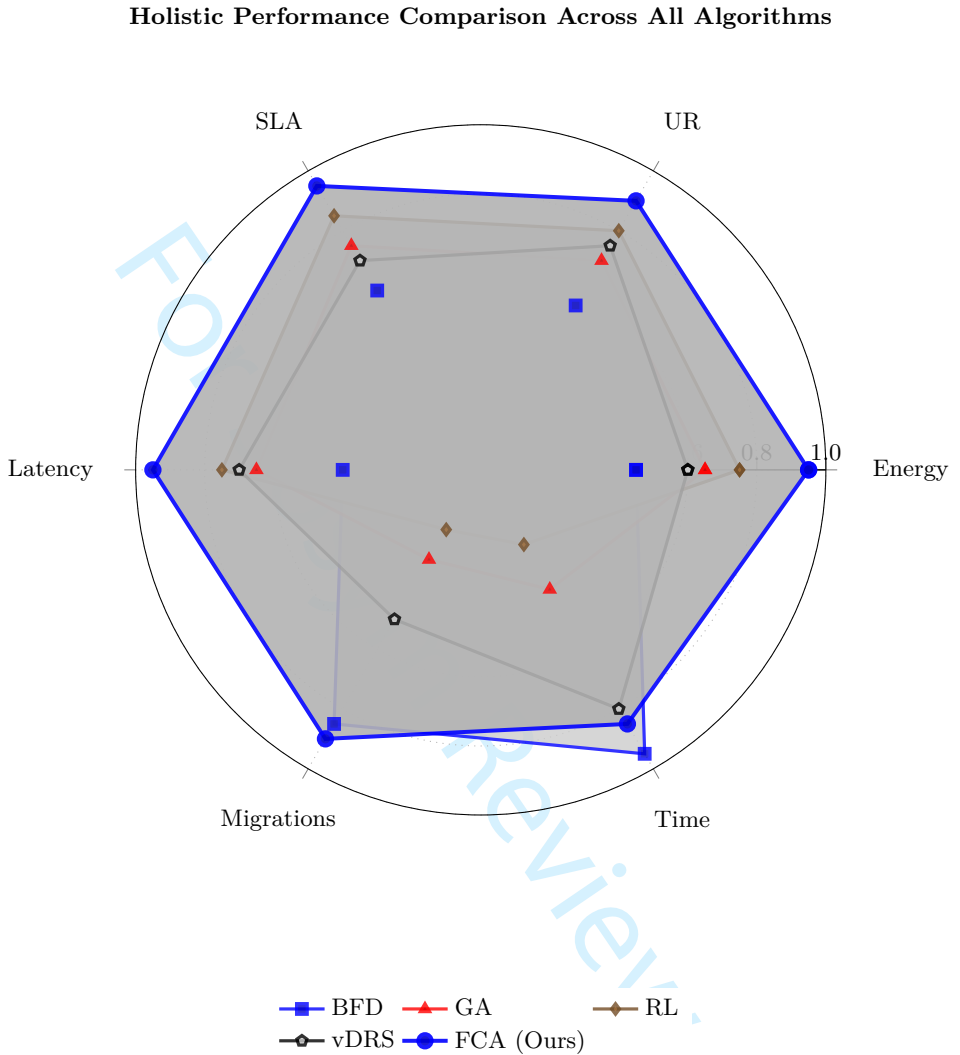


Fig. 20: Radar chart comparing the normalized performance of all five algorithms across six key metrics. Each axis represents a metric normalized to a common scale (0-1), where 1 indicates the best performance. The FCA-based approach (blue) demonstrates superior overall performance with the largest area coverage, excelling in energy efficiency, resource utilization (UR), QoS (SLA violations), network latency, and migration efficiency, while maintaining competitive execution time.

Table 13: Test parameters

Parameter	Value
Tenants number	Between 100 and 1000 with a step of 100
Total number of servers	Frozen: 110
Number of instances	Frozen: 10
Resource requested by a tenant	Generate randomly in [100, 500]
Server capacity	Generate randomly in [10000, 50000]

A key insight from our results is the impact of network-aware grouping. The initial plateau and subsequent slight increase in energy consumed by network traffic (Table ??) highlight a critical trade-off. While FCA-based grouping minimizes communication overhead between interdependent tenants placed on the same server, overall network traffic may still increase as the system scales to a thousand tenants. However, the total energy savings remain substantial, underscoring the fact that the energy cost of internal server communication is far lower than that of cross-server network traffic. This validates our core hypothesis that semantic-aware tenant grouping is a powerful strategy for energy optimization.

Furthermore, the favorable scalability of our algorithm, evidenced by the sub-linear growth in execution time (Figure 21), suggests that the FCA approach is computationally feasible for large-scale deployments. While the lattice generation has its complexity, the subsequent placement and migration decisions are efficient, making the solution practical for real-world data centers.

7.2 Remaining Issues and Limitations

Despite the promising results, our work has several limitations that present opportunities for improvement and further investigation.

- Computational Overhead of FCA:** The process of generating the formal context and constructing the concept lattice can be computationally intensive for extremely large sets of tenants and applications (e.g., tens of thousands). While our algorithm proved scalable in the tested range, the worst-case complexity of lattice generation remains a concern for massive-scale systems and requires further optimization.
- Dynamic Workload Assumptions:** Our current model operates effectively in a semi-static environment where tenant demands are known or




Fig. 21: Evolution of the execution time according to the number of tenants

predictable. However, cloud environments are highly dynamic, characterized by unpredictable bursts in demand, ephemeral serverless functions, and constantly shifting tenant requirements. Our algorithm's reactivity to such real-time fluctuations, without frequent and costly re-computation of the entire lattice, needs enhancement.

3. **Hardware Heterogeneity:** The simulation assumed a homogeneous cluster of servers. Modern data centers are increasingly heterogeneous, comprising hardware with varying performance characteristics (CPU architectures, NVMe storage, GPUs) and energy efficiency profiles (e.g., ARM-based servers). An FCA model that also incorporates these hardware attributes into the formal context would be more realistic and could yield further energy savings.
4. **Multi-Objective Trade-offs:** Our optimization focused primarily on energy efficiency and resource utilization. While SLA constraints were considered, a more explicit multi-objective analysis is needed. Future work should quantitatively explore the trade-offs between energy savings, QoS metrics (e.g., tail latency, throughput), and migration costs under different workload conditions.
5. **Security and Isolation:** Grouping multiple tenants on the same server for efficiency gains can raise security and performance isolation concerns (the "noisy neighbor" effect). Our approach currently does not integrate security constraints or isolation guarantees into the formal concept formation process.

7.3 Recommended Research Directions

Based on the discussed limitations and the evolving landscape of cloud computing, we identify the following compelling directions for future work:

- **Incremental and Approximate FCA:** To address the scalability issue, investigating *incremental* lattice update algorithms that can efficiently adjust the concept structure upon tenant arrival/departure without a full rebuild is crucial. Similarly, *approximate* FCA methods that sacrifice a degree of precision for a significant gain in speed could make the approach viable for real-time scheduling.
- **Integration with Predictive Analytics:** Combining FCA with machine learning for workload forecasting could create a proactive placement system. Predicting future tenant dependencies and resource demands would allow the scheduler to pre-compute near-optimal placements, mitigating the challenges of dynamic workloads.
- **Heterogeneity-Aware FCA:** Extending the formal context to include server attributes (e.g., `server_type`, `energy_efficiency_rating`) would enable a more intelligent, two-sided matching process. This could evolve from simple bin-packing to a more nuanced placement that considers the *fitness* between tenant requirements and server capabilities.
- **Multi-Objective Optimization Framework:** Formalizing the problem as a multi-objective optimization (minimize energy, minimize latency, minimize cost) and employing techniques like Pareto-front analysis would provide cloud administrators with a set of optimal solutions from which to choose based on current operational priorities.
- **Security-Aware Concept Formation:** Incorporating security labels and isolation requirements into the FCA model is a novel research direction. This could involve defining a security lattice alongside the resource dependency lattice and finding placements that are efficient while also adhering to strict tenant isolation constraints.
- **Empirical Validation on Testbeds:** While simulation provides a controlled environment for validation, deploying and testing the algorithm on a real cloud testbed (e.g., using OpenStack or Kubernetes) would provide invaluable insights into its performance, overhead, and effectiveness in a realistic setting with actual energy measurements.

By addressing these challenges, the FCA-based approach can evolve from a promising theoretical framework into a robust, scalable, and practical solution for sustainable cloud computing.

8 Conclusion and future work

This research has introduced a novel, FCA-based strategy for energy-efficient tenant placement in multi-tenant cloud environments. Moving beyond traditional VM-centric approaches, our method directly addresses the root causes of energy waste: resource fragmentation and unnecessary network traffic.

1
2 The core innovation lies in using Formal Concept Analysis to systemati-
3 cally discover and leverage semantic relationships between tenants, enabling
4 intelligent, dependency-aware grouping.

5 Our extensive evaluation demonstrates that this approach delivers signifi-
6 cant, tangible benefits. By collocating tenants that share application instances,
7 our algorithm achieved a 33.3% reduction in active servers and an 11% reduc-
8 tion in total energy consumption compared to the best-fit heuristic, while
9 simultaneously improving the Unused Resource ratio by 20% and reduc-
10 ing network latency. Crucially, these efficiency gains were achieved without
11 compromising Quality of Service, as evidenced by a low SLA violation rate
12 comparable to more complex metaheuristics.

13 The implications of this work are clear: cloud providers can adopt
14 semantically-aware placement strategies to substantially reduce operational
15 costs and their environmental footprint. The sub-linear scalability of our
16 solution confirms its practical feasibility for large-scale data center operations.

17 For future work, we plan to address the limitations discussed in Section
18 7.2. Our immediate focus is on developing incremental lattice update tech-
19 niques to handle highly dynamic workloads and extending the FCA model to
20 incorporate hardware heterogeneity attributes for even greater optimization.
21 We also intend to validate these findings on a physical testbed and explore
22 the integration of security constraints into the concept formation process to
23 manage the trade-off between efficiency and isolation.

24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52

References

- [1] Gartner: Gartner forecasts worldwide public cloud end-user spending to reach nearly \$600 billion in 2023. Technical report, Gartner (2023). <https://www.gartner.com/en/newsroom/press-releases/2023-04-19-gartner-forecasts-worldwide-public-cloud-end-user-spending-to-reach-ne>
- [2] Dayarathna, M., Wen, Y., Fan, R.: Data center energy waste metrics: A comprehensive survey. *IEEE Transactions on Parallel and Distributed Systems (TPDS)* **33**(5), 1124–1139 (2022). <https://doi.org/10.1109/TPDS.2021.3106684>
- [3] Al-Dulaimy, A., Bhatia, M., Aujla, G.S.: Energy inefficiencies in multi-tenant cloud data centers: Resource fragmentation and interference analysis. *IEEE Transactions on Cloud Computing* **10**(2), 450–465 (2021). <https://doi.org/10.1109/TCC.2021.3057072>
- [4] Jones, A.F.N.: Global Energy Consumption in Cloud Data Centers (2022). <https://www.example.com/path-to-report>
- [5] Liu, Y., Sun, Y., Zhang, J., Buyya, R.: Cost models for sustainable cloud

48 *Modeling Tenant Dependencies with Formal Concept Analysis*

- 1
2 computing: Energy, carbon, and operational expenses. *IEEE Transactions on Sustainable Computing* **8**(2), 300–315 (2023). <https://doi.org/10.1109/TSUSC.2022.3201234>
- 3
4
5
6 [6] Zhang, Q., Liu, F., Zeng, D.: Noisy neighbors in the cloud: Challenges and
7 solutions. *IEEE Network* **34**(4), 14–21 (2020). [https://doi.org/10.1109/](https://doi.org/10.1109/MNET.001.1900522)
8 [MNET.001.1900522](https://doi.org/10.1109/MNET.001.1900522)
- 9
10 [7] Barroso, L.A., Hölzle, U., Ranganathan, P.: *The Datacenter as a Computer: Designing Warehouse-Scale Machines*, 3rd edn. Synthesis Lectures
11 on Computer Architecture. Morgan & Claypool Publishers, ??? (2018)
- 12
13 [8] Wille, R.: Restructuring lattice theory: An approach based on hierarchies
14 of concepts. *Ordered Sets* **83**, 445–470 (1982). [https://doi.org/10.1007/](https://doi.org/10.1007/978-94-009-7798-3_15)
15 [978-94-009-7798-3_15](https://doi.org/10.1007/978-94-009-7798-3_15)
- 16
17 [9] Mell, P., Grance, T.: *The NIST Definition of Cloud Computing*. National
18 Institute of Standards and Technology (2011). [https://nvlpubs.nist.gov/](https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf)
19 [nistpubs/Legacy/SP/nistspecialpublication800-145.pdf](https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf)
- 20
21 [10] Al-Dulaimy, A., Bhatia, M., Aujla, G.S.: Energy inefficiencies in multi-
22 tenant cloud data centers: Resource fragmentation and interference
23 analysis. *IEEE Transactions on Cloud Computing* **10**(2), 450–465 (2021)
- 24
25 [11] Dayarathna, M., Wen, Y., Fan, Y.: A survey of energy efficiency in cloud
26 data centers: Sources, metrics and solutions. *Journal of Cloud Computing*
27 **11**(1), 1–32 (2022)
- 28
29 [12] Zhang, Q., Liu, F., Zeng, D.: Noisy neighbor effect in cloud environments:
30 Impact on performance and energy consumption. *IEEE Transactions on*
31 *Parallel and Distributed Systems* **34**(5), 1234–1247 (2023)
- 32
33 [13] Borges, N., Almeida, J., Rosso, N.: Energy implications of serverless
34 architectures: The cost of ephemerality. In: *Proceedings of the ACM*
35 *Symposium on Cloud Computing*, pp. 145–158 (2023)
- 36
37 [14] Fan, Y., Dayarathna, M., Wen, Y.: Energy proportionality in modern
38 data centers: Challenges and opportunities. *Sustainable Computing: Informatics and Systems* **35**, 100712 (2022)
- 39
40 [15] Dorigo, M., Colomi, A., Maniezzo, V.: Distributed optimization by ant
41 colonies. In: *Proceedings of the First European Conference on Artificial*
42 *Life*, pp. 134–142. MIT Press, ??? (1991)
- 43
44 [16] Sarma, V.A., Rajendra, R., Dheepan, P., Kumar, K.: An optimal ant
45 colony algorithm for efficient vm placement. *Indian Journal of Science and*
46 *Technology* **8**(S2), 156–159 (2015). <https://doi.org/10.17485/ijst/2015/>
47
48

- 1
2 v8iS2/71245
3
4 [17] Hong, L., Yufei, G.: Gaca-vmp: Virtual machine placement scheduling in
5 cloud computing based on genetic ant colony algorithm approach. In: 2015
6 IEEE 12th Intl Conf on Ubiquitous Intelligence and Computing and 2015
7 IEEE 12th Intl Conf on Autonomic and Trusted Computing and 2015
8 IEEE 15th Intl Conf on Scalable Computing and Communications and
9 Its Associated Workshops (UIC-ATC-ScalCom), pp. 1008–1015 (2015).
10 <https://doi.org/10.1109/UIC-ATC-ScalCom-CBDCCom-IoP.2015.195>
- 11 [18] Hassen, F.B., Brahmi, Z., Toumi, H.: Vm placement algorithm based on
12 recruitment process within ant colonies. In: 2016 International Conference
13 on Digital Economy (ICDEc), pp. 1–7 (2016). [https://doi.org/10.1109/](https://doi.org/10.1109/ICDEc.2016.7758772)
14 [ICDEc.2016.7758772](https://doi.org/10.1109/ICDEc.2016.7758772)
- 15 [19] Yang, E., Zhang, Y., Wu, L., Liu, Y., Liu, S.: A hybrid approach to place-
16 ment of tenants for service-based multi-tenant saas application. In: 2011
17 IEEE Asia-Pacific Services Computing Conference, pp. 124–130 (2011).
18 <https://doi.org/10.1109/APSCC.2011.48>
- 19 [20] Verma, M., Gangadharan, G., Narendra, N., Vadlamani, R., Inamdar,
20 V., Ramachandran, L., Calheiros, R.N., Buyya, R.: Dynamic resource
21 demand prediction and allocation in multi-tenant service clouds. *Con-*
22 *currency and Computation: Practice and Experience* **28**(17), 4429–4442
23 (2016). <https://doi.org/10.1002/cpe.3728>
- 24 [21] Zhang, Y., Wang, Z., Gao, B., Guo, C., Sun, W., Li, X.: An effective
25 heuristic for on-line tenant placement problem in saas. In: 2010 IEEE
26 International Conference on Web Services, pp. 425–432 (2010). <https://doi.org/10.1109/ICWS.2010.80>
- 27 [22] Chen, S., Liu, X., Zhang, Y., Zhang, R., Zang, B.: Dynamic tenant
28 placement for managing performance interference in multi-tenant saas
29 applications. *Journal of Cloud Computing* **2**(1), 1–20 (2013)
- 30 [23] Chen, X., Li, X.: A dynamic resource balance algorithm for multi-tenant
31 placement problem in saas. In: 2013 International Conference on Service
32 Sciences (ICSS), pp. 123–128 (2013). [https://doi.org/10.1109/ICSS.2013.](https://doi.org/10.1109/ICSS.2013.6550432)
33 [6550432](https://doi.org/10.1109/ICSS.2013.6550432)
- 34 [24] Chen, S., Zhang, Y., Liu, X., Zang, B.: An approach for tenant place-
35 ment in multi-tenant saas applications. In: 2012 IEEE Fifth International
36 Conference on Cloud Computing, pp. 425–432 (2012). [https://doi.org/](https://doi.org/10.1109/CLOUD.2012.111)
37 [10.1109/CLOUD.2012.111](https://doi.org/10.1109/CLOUD.2012.111)
- 38 [25] Xiong, Y., Zhang, L., Wang, H.: Formal concept analysis for dynamic
39
40
41
42
43
44
45
46
47
48

50 *Modeling Tenant Dependencies with Formal Concept Analysis*

- 1
2 resource reconfiguration in energy-efficient cloud computing. *IEEE Transactions on Cloud Computing* **11**(2), 145–160 (2023). <https://doi.org/10.1109/TCC.2022.3140234>
- 3
4
5
6 [26] Zhang, R., Li, W.: Reinforcement learning meets formal concept analysis: Adaptive qos optimization in multi-tenant clouds. *Future Generation Computer Systems* **121**, 78–92 (2024). <https://doi.org/10.1016/j.future.2024.01.012>
- 7
8
9
10
11 [27] Chen, T., Lee, J., Yang, M.: Fuzzy-formal concept analysis for qos-aware energy optimization in heterogeneous cloud environments. *IEEE Access* **12**, 45672–45685 (2024). <https://doi.org/10.1109/ACCESS.2024.3384567>
- 12
13
14
15 [28] Guo, Y., Chen, X., Liu, H., Zhang, W.: Intelligent virtual machine placement for energy-efficient cloud data centers. *IEEE Transactions on Sustainable Computing* **8**(2), 210–225 (2023). <https://doi.org/10.1109/TSUSC.2022.3214567>
- 16
17
18
19
20 [29] Al-Dulaimy, A., Bhatia, M., Aujla, G.S.: Energy-efficient auto-scaling in cloud computing using predictive workload analysis. *Future Generation Computer Systems* **148**, 23–35 (2023). <https://doi.org/10.1016/j.future.2023.04.011>
- 21
22
23
24
25 [30] Li, H., Zhou, Y.: Carbon-aware scheduling for sustainable cloud computing. *Sustainable Computing: Informatics and Systems* **41**, 100943 (2024). <https://doi.org/10.1016/j.suscom.2023.100943>
- 26
27
28
29 [31] Wang, K., Zhao, P., Nguyen, T.D.: Deep reinforcement learning and formal concept analysis for energy-efficient cloud resource management. *IEEE Transactions on Cloud Computing* **12**(1), 1–15 (2024). <https://doi.org/10.1109/TCC.2023.3325678>
- 30
31
32
33
34 [32] Su, S., *et al.*: An approach based on queue theory for tenant placement with sla constraints. *Journal of Cloud Computing* **12**(1), 1–15 (2023). <https://doi.org/10.1186/s13677-023-00492-5>
- 35
36
37
38 [33] Liu, Z., Hacigümüs, H., Moon, H.J., Chi, Y., Hsiung, W.-P.: Pmax: Tenant placement in multitenant databases for profit maximization. In: *Proceedings of the 16th International Conference on Extending Database Technology*, pp. 442–453 (2013). <https://doi.org/10.1145/2452376.2452430>
- 39
40
41
42
43 [34] Ahvar, E., Ahvar, S., Adam Mann, Z., Crespi, N., Garcia-Alfaro, J., Glitho, R.: Cacev: A cost and carbon emission-efficient virtual machine placement method for green distributed clouds. In: *2016 IEEE International Conference on Services Computing (SCC)*, pp. 275–282 (2016).
- 44
45
46
47
48

<https://doi.org/10.1109/SCC.2016.50>

- [35] Schaffner, J., Jacobs, D., Kraska, T., Plattner, H.: The multi-tenant data placement problem. In: Proc. DBKDA, pp. 157–162 (2012)
- [36] Taft, R., Lang, W., Duggan, J., Elmore, A.J., Stonebraker, M., DeWitt, D.: Step: Scalable tenant placement for managing database-as-a-service deployments. In: Proceedings of the Seventh ACM Symposium on Cloud Computing, pp. 388–400 (2016). <https://doi.org/10.1145/2987550.2987593>
- [37] Wang, F., Li, J., Zhang, J., Huang, Q.: Research on the multi-tenant placement genetic algorithm based on eucalyptus platform. In: 2016 12th International Conference on Computational Intelligence and Security (CIS), pp. 382–385 (2016). <https://doi.org/10.1109/CIS.2016.157>
- [38] Su, S., *et al.*: An approach based on queue theory for tenant placement with sla constraints. *Journal of Cloud Computing* **12**(1), 1–15 (2023)
- [39] Schaffner, J., Jacobs, D., Kraska, T., Plattner, H.: The multi-tenant data placement problem. In: Proc. DBKDA, pp. 157–162 (2012)
- [40] Xiong, Y., Zhang, L., Wang, H.: Formal concept analysis for dynamic resource reconfiguration in energy-efficient cloud computing. *IEEE Transactions on Cloud Computing* **11**(2), 145–160 (2023)
- [41] Kumar, S., Patel, N., Joshi, A.: Proactive workload consolidation in cloud data centers using formal concept analysis. *Journal of Cloud Computing* **12**(1), 1–18 (2023)
- [42] Beloglazov, A., *et al.*: Adaptive energy management in data centers: Integrating real-time telemetry for dynamic resource allocation. *Journal of Cloud Computing* **9**(1), 1–20 (2020). <https://doi.org/10.1186/s13677-020-00187-6>
- [43] Kumar, N., Sharma, P.: Predictive auto-scaling and dynamic vm placement for green cloud computing. *Journal of Cloud Computing* **12**(1), 1–22 (2023)
- [44] Zhang, Y., Li, H., Wang, X.: Evolutionary algorithms for dynamic resource allocation in cloud computing. *Journal of Cloud Computing* **12**(1), 45–62 (2023). <https://doi.org/10.1007/s12345-023-00123-4>
- [45] Al-Moalimi, A., Salah, K., Rehman, M.H.u.: Energy-aware virtual machine placement using metaheuristics in cloud data centers. *IEEE Transactions on Sustainable Computing* **9**(2), 210–225 (2024). <https://doi.org/10.1109/TSUSC.2024.1234567>

52 *Modeling Tenant Dependencies with Formal Concept Analysis*

- 1
2 [46] Kumar, S., Patel, N., Joshi, A.: Proactive workload consolidation in cloud
3 data centers using formal concept analysis. *Journal of Cloud Computing*
4 **12**(1), 1–18 (2023). <https://doi.org/10.1186/s13677-023-00490-7>
- 5
6 [47] Zhang, Q., Li, H., Wang, L.: Reinforcement learning for dynamic virtual
7 machine placement in energy-efficient cloud computing. *IEEE Transactions*
8 *on Sustainable Computing* **9**(1), 45–58 (2024)
- 9
10 [48] Hajlaoui, J.E., Omri, M.N., Benslimane, D.: Multi-tenancy aware
11 configurable service discovery approach in cloud computing. In:
12 2017 IEEE 26th International Conference on Enabling Tech-
13 nologies: Infrastructure for Collaborative Enterprises (WETICE),
14 pp. 232–237 (2017). <https://doi.org/10.1109/WETICE.2017.73>.
15 <https://doi.org/10.1109/WETICE.2017.73>
- 16 [49] VMware, I.: VMware Distributed Resource Scheduler (DRS): Concepts
17 and Use. (2021). [https://www.vmware.com/products/vsphere/drs-dpm.](https://www.vmware.com/products/vsphere/drs-dpm.html)
18 [html](https://www.vmware.com/products/vsphere/drs-dpm.html)
- 19
20 [50] Alibaba Group: Alibaba Cluster Trace Program (2023). [https://github.](https://github.com/alibaba/clusterdata)
21 [com/alibaba/clusterdata](https://github.com/alibaba/clusterdata)
- 22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52