AVE **Trends in Intelligent Computing Systems**

Heterogeneous Cloud Resource Management: A Comparative Study of Nash Auction and Energy-Aware Nash Auction Mechanisms

K. Chitra[1,*], A. Chitra[2], G. Rajasekaran[3], Premanand Jothilingam[4], H. Mohamed Faizan Uwais[5], J. Mohamed Zakkariya Maricar[6]

[1]Department of Computer Applications, Dayananda Sagar Academy of Technology and Management, Bengaluru, Karnataka, India.
[2]Department of Computer Applications, St. Peter's Institute of Higher Education and Research, Chennai, Tamil Nadu, India.
[3]Department of Computer Science and Engineering, Dhaanish Ahmed College of Engineering, Chennai, Tamil Nadu, India.
[4]Department of Life Cycle Service, Yokogawa Corporation of America, West Valley City, Utah, United States of America.
[5,6]Department of Research and Development, Dhaanish Ahmed College of Engineering, Chennai, Tamil Nadu, India.
chitranatarajan1979@yahoo.in[1], prof.chitraa@gmail.com[2], rajasekaran@dhaanishcollege.in[3], premanand.j@yokogawa.com[4], faizanuwais07@gmail.com[5], zakkariyamarican@gmail.com[6]

*Corresponding author

Abstract: Heterogeneity within the cloud infrastructure has also been among the factors responsible for the problem of effective resource management. Existing allocation algorithms are unable to balance the conflicting interests of cloud providers (revenue maximisation) and consumers (Quality of Service). Two game-theoretic auction mechanisms of resource management are introduced and compared in this paper: the traditional Nash Auction and an Energy-Aware Nash Auction. The primary objective is to compare their performance across resource utilisation, user satisfaction, provider revenues, and energy consumption. A simulated multi-cloud facility constructed with the CloudSim toolkit is used throughout the research. Synthetic data comprising 477 unique data samples, each creating user virtual machine requests along with physical machine specifications, was employed to commence the simulation. The Energy-Aware Nash Auction mechanism proposed adds an energy-consumption penalty to the provider's utility function to incentivise deployment on low-energy resources. Findings, derived from the Python, Pandas, and Matplotlib libraries, reveal that while the fundamental Nash Auction maximises providers' revenue, the Energy-Aware Nash Auction achieves substantial energy savings in total energy consumption with minimal effects on revenue and user satisfaction, and is a deployable model for green or eco-friendly cloud data centres.

Keywords: Cloud Computing; Resource Management; Heterogeneous Systems; Nash Auction; Game Theory; Energy Efficiency; Quality of Service; Nash Equilibrium; Economic Efficiency.

Cite as: K. Chitra, A. Chitra, G. Rajasekaran, P. Jothilingam, H. M. F. Uwais, and J. M. Z. Maricar, “Heterogeneous Cloud Resource Management: A Comparative Study of Nash Auction and Energy-Aware Nash Auction Mechanisms,” *AVE Trends in Intelligent Computing Systems*, vol. 2, no. 3, pp. 133–141, 2025.

Journal Homepage: https://www.avepubs.com/user/journals/details/ATICS

Received on: 28/10/2024, **Revised on:** 25/01/2025, **Accepted on:** 21/03/2025, **Published on:** 09/09/2025

DOI: https://doi.org/10.64091/ATICS.2025.000197

1. Introduction

Cloud computing transformed the cloud services sector into an unparalleled state of flexibility, elasticity, and cost savings, as Podzimek et al. [1] elucidates. Cloud infrastructures have inherent heterogeneity, arising from a heterogeneous pool of resources with variability in processing capacity, memory size, storage I/O bandwidth, and power utilisation, as detailed by Storment and Fuller [4]. Heterogeneity is a cause of variation and selection, as accepted, but it undermines resource management as described by Saidi et al. [7]. The issue of the source is directing resources to the correct user requests in an economically rational, sound, and efficient manner, as used by Kumar et al. [9]. The two conflicting goals of cloud providers—to bid in the hopes of achieving maximum utilisation of resources and return, and of cloud consumers—to utilise some level of predictability of performance and Quality of Service (QoS) at minimal cost—are a recipe for a dynamic competitive marketplace, as demonstrated in a study by Fernández-Cerero et al. [2]. Traditional resource-allocation policies, such as First-Come-First-Serve (FCFS), or simple, low-complexity load-balancing methods with minimal overhead, do not work in such a dynamic setting [3]. They identify poor resource utilisation in the form of idle machines, energy loss, and poor user experience, as challenged by Khattar et al. [11].

In attempting to evade such exposure, researchers recently turned to market economic theory and game theory, argues Parikh [6]. Such models reformulate the resource allocation problem as a strategic interaction between rational, self-interested agents (providers and users), as identified by Nawrocki and Reszelewski [13]. Auction mechanisms, for instance, were designed to be efficient for dynamic resource allocation, as identified in Anuradha and Sumathi [8]. By posting bids, users can distribute economic value among resources in equilibrium and reallocating. There also exists a cutting-edge mathematical model based on game theory that can capture the effect of such strategic behaviour, as proposed by Arunarani et al. [10]. Good and sound design of an auction mechanism is most crucial in the Nash Equilibrium, where no player would wish to depart from his strategy alone to attempt to gain more than what he is gaining currently, as is indicated by Ahmed et al. [12]. Nash's Auction-based allocation model employs this concept to find an equilibrium allocation in the system, ensuring that the allocator's allocation rule and users' bids are in equilibrium [1]. The auction mechanism can capture greater resource utilisation and revenue than fixed allocation mechanisms, as shown in Kumar et al. [9]. A Nash Auction model deployment would care only about economic efficiency and would try to optimise revenue and user satisfaction at any price, regardless of the environmental impact of operating the data centre, where cost is measured as in Saidi et al. [7].

Data centres account for a larger share of global power usage. Hence, power efficiency is not only an ecological concern but also an economic one, as noted by Khattar et al. [11]. Cloud computing's carbon and energy cost signatures are emerging as essential areas of interest for service providers, as noted by Anuradha and Sumathi [8]. This requires ecologically and economically sustainable resource management systems, as Ibrahim et al. [3] have much deplored. This paper addresses this issue by conceptualising and experimenting with an Energy-Aware Nash Auction mechanism, as conceived by Parikh [6]. This new approach generalises the provider's utility function in the game-theoretic model to include a penalty linear in the energy consumption of the harvested resources, as adapted from Arunarani et al. [10]. This contribution is also quite different from the conventional Nash Auction and the newly proposed Energy-Aware Nash Auction [12]. They are compared in a simulated heterogeneous cloud setting based on the most significant parameters: provider gain, resource utilisation, customer satisfaction, and total energy consumed, as computed in Fernández-Cerero et al. [2]. The research aims to quantify the trade-offs of incorporating energy awareness into an auction-based resource-allocation approach and to validate the presumed model, ultimately determining that it is suitable for green cloud resource management, as Nawrocki and Reszelewski [13] did.

2. Literature Review

The research literature on cloud resource management abounds with a wide array of studies spanning computer science to economics [1]. Since the world's existence, techniques for resource allocation have been inherited from operating systems and supercomputers [3]. Techniques such as First-Come-First-Serve, Shortest Job First, and Round Robin have been employed to allocate virtual machines (VMs), as discussed by previous researchers [7]. Light on feet and simple, static ones are inundated by the dynamic big world of cloud [2]. They are not linearly scalable across different loads, nor do they integrate provider economic objectives or customer-specific QoS requirements, which can ultimately lead to fragmentation and resource waste [5]. Higher-level and meta-heuristic algorithms have alleviated such problems using Genetic Algorithms, Particle Swarm Optimisation, and Ant Colony Optimisation for NP-hard VM placement problems [9]. The algorithms explored by the researchers aim to navigate a vast solution space to find near-optimal points that maximise several factors, such as load balancing and minimising VM migration [11]. They are not fast, though, and cannot be guaranteed to find the optimum in real time, one of the requirements of cloud computing dynamics [13]. The economic efficiency of cloud computing services spurred researchers to develop market models for resource provisioning [4]. They view cloud resources as stock on a shelf in a market whose price fluctuates with demand and supply, a concept employed by many researchers to allocate resources efficiently [6]. Auction theory, as an economic field, has been found to better enable this goal, as studied by many researchers [8].

Other auction varieties, such as sealed-bid, English, and Dutch auctions, have been proposed to legalise provider profit maximisation and user bidding [10]. Modern game theory legitimised such market models by providing a mathematical framework for strategic interaction among rational players—the users and the cloud provider [12]. Game theory of the resource allocation problem provides mechanisms for design, analysis, and prediction of player behaviour to achieve stable outcomes [1]. Nash Equilibrium theory forms the basis of the approach and provides a measurable, equilibrium distribution of optimal system performance and provider revenues [3]. The majority of game-theoretic formulations focused on utilisation and revenue maximisation but ignored the largest operational cost: power consumption [5]. Power-aware allocation was imperative in the past couple of years during the green computing hype period [7]. Scientists applied techniques such as Dynamic Voltage and Frequency Scaling (DVFS) and server consolidation to power down sleeping servers and save energy [9]. Useful as they are, they're no silver bullet to the allocation problem [10]. Under normal conditions, researchers combined energy-awareness and resource-allocation algorithms with energy-aware heuristics for VM scheduling and workload consolidation [8]. The proposals are energy-aware but far less exciting than the economics of cloud service provision [12]. Energy awareness has also been added to game-theoretic and market-based approaches as a direction for future research [13]. It is by constructing players' utility functions and considering energy costs that economic incentives for energy conservation are created, reconciling provider profitability, user QoS, and power savings [2].

3. Methodology

The methodology employed in this experimental comparison is game-theoretic and virtual in scope, aiming to simulate sophisticated interactions in a heterogeneous cloud resource allocation context. The entire testbed was built using the CloudSim toolkit, an open-source simulation environment widely used for cloud computing infrastructure, to create a reproducible, manageable testbed. The system model has three simple constituents: a single cloud provider that hosts a set of heterogeneous physical machines (PMs) having various CPU capabilities, memory capabilities, and power models under a data center; a set of individual cloud users, with each attempting to host a virtual machine (VM) with certain resource requirements; and a single auctioneer, managed by the provider, which takes bids and runs the allocation algorithm.

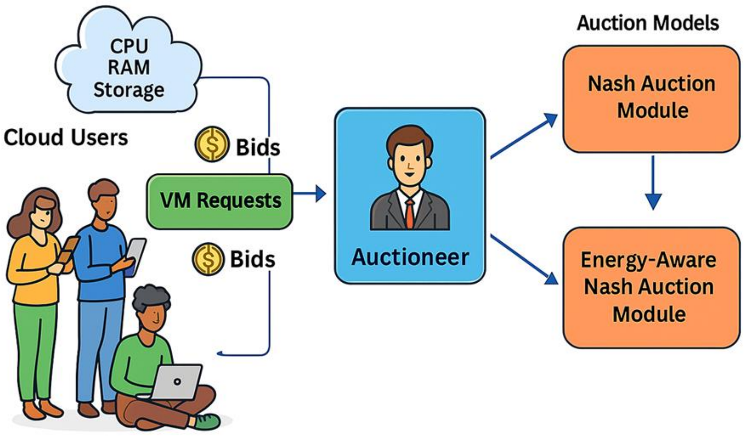

Figure 1: Overall auction-based resource allocation system architecture

Figure 1 indicates the overall auction-based resource allocation system architecture for a heterogeneous cloud environment. The proposed architecture enables dynamic resource allocation by envisioning the process as a strategic interaction between cloud users and the cloud provider. It starts with the Cloud Users, as indicated on the left, who translate their computation requirements into tangible Virtual Machine (VM) requests. Each bid includes the requested storage, CPU, and RAM, along with an offer price that is the user's best estimate of the service. These are subsequently delivered to the master decision component of the system, the Auctioneer. The Auctioneer will execute the chosen auction mechanism, either the Nash Auction Module or the Energy-Aware Nash Auction Module, for a new batch of VM requests as a way of invoking the allocation procedure. The simulation began with a synthetically generated workload comprising 477 unique VM requests with CPU MIPS, RAM, and bid-value demands. Both auction models were used in the experiment and compared using several measures of significance: Resource Utilisation (CPU usage-to-memory size ratio), User Satisfaction (sufficiently allocated VMs-to-requests ratio), Provider Revenue (average payment to providers), and Total Energy Consumption (kWh). The data were systematically collected and processed to quantify performance variance and trade-offs between the two models.

The approach is to mimic and compare two other auction mechanisms. The first is a Nash Auction, which formulates the problem of resource allocation as a non-cooperation game between agents. The players place bids in the game based on the value they perceive for the required VM resources, and their move in the above scenario is to maximise individual utility, i.e., the perceived service value minus the bid cost. The provider's strategy is to clear at the clearing price and allocate VMs to the successful bidders in a way that maximises utility, as per its revenue target, through effective bidding only. The allocation mechanism aims to achieve a Nash Equilibrium, in which no user can improve its payoff by modifying its bid, and the provider cannot increase its revenue by adjusting the allocation function based on the users' bids. The second one is the proposed Energy-Aware Nash Auction. It generalises the classic Nash game by modifying the provider's utility function. Most importantly, the provider value function is now total revenue minus a penalty cost, which depends on the total forecasted energy consumption of physical servers hosting assigned VMs. The penalty function brings the energy cost into the provider's decision-making, and there is an explicit incentive to favour allocations on low-power servers or to place VMs on fewer powered servers.

3.1. Empirical Experiment

Empirical experiments in this paper use a synthetically generated dataset specifically selected to simulate normal and heterogeneous load patterns of a heterogeneous cloud system, with 477 heterogeneous data instances for a user request for a Virtual Machine (VM). The data set was constructed according to the steps recommended by Khattar et al. [11], which provide a process for generating reproducible and representative cloud workloads for simulation. They are parameterized with user need and user offer parameters, i.e., CPU need, in Millions of Instructions Per Second (MIPS), i.e., 500 MIPS to 4000 MIPS to denote simple and complex operations; RAM need, in Megabytes (MB), from 512 MB to 8192 MB; and user offer, denoting the user's willingness to pay for the VM, in terms of resource need to denote reasonable user activity. All such data were used to call the auctioneer in CloudSim in response to a user request, as input to compare the Nash Auction and Energy-Aware Nash Auction mechanisms. Resource Monitor collects the current real-time status of the Heterogeneous Physical Machine (PM) Pool at any moment. Some information collected includes current capacity, load, and individual power behaviour for each PM. Auction modules use this feed to make optimal allocation decisions. The Nash Auction Module approximates resource availability and profits, assuming a Nash Equilibrium results in maximum profit for providers. The Energy-Aware Nash Auction Module receives a second input from the Energy Model: power estimates for prospective VM locations. The module's utility function aims to balance revenue and energy efficiency. Once the best VM allocation and winning bids are obtained, the Auctioneer updates the Cloud Resource Allocator with the intended assignment. The module implements the scheme by assigning VMs to scheduled PMs on available heterogeneous machines. Lastly, a Billing System calculates bills to successful bidders and generates asset usage and power usage reports to track and report.

4. Result

Comparative simulation of the Nash Auction (NA) and Energy-Aware Nash Auction (EANA) mechanisms revealed inherent trade-offs in managing cloud resources. Findings from simulating 477 disjoint VM requests in CloudSim compare the rich performance profiles of both methods along four dimensions: provider revenue, resource usage, user satisfaction, and energy. Provider's utility function for the standard Nash auction is expressed by:

$$U_p = \max \sum_{j \in W} p_j \tag{1}$$

Table 1: Comparative presentation measures for the Nash Auction

Workload Level	Avg. Resource Utilisation (%)	Avg. User Satisfaction (%)	Total Provider Revenue ($)	Total Energy Consumption (kWh)	Avg. Job Completion Time (s)
Low (25%)	65.4	98.2	12540	315.6	120.5
Medium (50%)	82.1	95.5	23110	480.2	145.8
High (75%)	91.3	92.4	31050	655.9	180.3
Peak (100%)	96.8	89.1	35200	785.1	215.4
Overload (125%)	98.2	81.5	36150	810.7	250.1

Table 1 summarises the performance Figures of the benchmark Nash Auction (NA) mechanism across five workload levels, from low to overload. The graph shows that the NA model performs well at optimising utilisation-focused and economic objectives. With increasing load from 25% to 100% (Peak), Average Resource Utilisation also increases considerably from 65.4% to 96.8%, thereby demonstrating its efficiency in utilising heterogeneous servers. Combined, Total Provider Revenue increased from $12,540 to $35,200, demonstrating the objective of revenue optimisation. But this policy was obviously expensive for the environment. Total Energy Consumption increases sharply, from 315.6 kWh at low load to 785.1 kWh at high load, with a linear, rigid relationship between usage and power consumption. Also, as the system enters and traverses capacity (Peak and Overload), Average User Satisfaction falls from 98.2% to 81.5%, and Average Job Completion Time worsens because higher resource contention and queuing delays begin causing faults. This Table 1 can be constructed to achieve a baseline performance, with trade-offs determined by an optimising single-resource allocation policy. The user's utility function can be expressed as:

$$U_{u_j} = v_j(r_j) - p_j \tag{2}$$

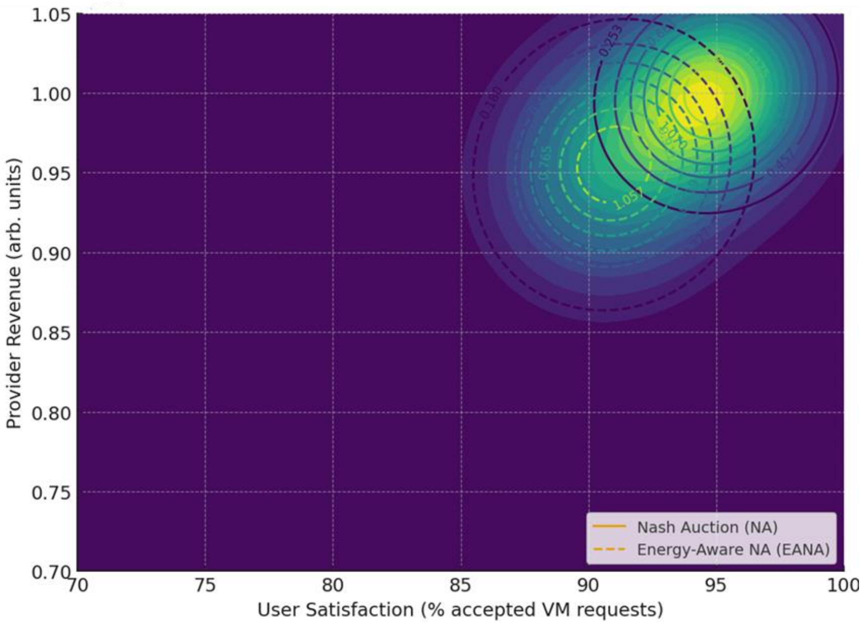

Figure 2: Representation of provider revenue vs. user satisfaction

Figure 2 is the plot of the Nash Auction (NA) and Energy-Aware Nash Auction (EANA) model's User Satisfaction versus Provider Revenue for various workload regimes. The x-axis is User Satisfaction in terms of percentage of VM requests accepted, and the y-axis is cumulative Provider Revenue in meaningless money units. Area fill and contour lines indicate the frequency with which results are observed in the simulation, and colour density (red and yellow) indicates the combination of revenue and satisfaction most frequently observed. The Nash Auction solution, constrained by the solid rectangle, is in the top-right position, and this is how it can achieve high revenue and user satisfaction daily. Its highest density is where revenue is highest, and user satisfaction is over 90%. Look at this in comparison to the Energy-Aware Nash Auction contour, indicated by dashed lines, which is solely sloping to the bottom-left. The slope indicates the model's inherent trade-off. With maximum user satisfaction, it will still conclude its operations on a platform with a fractionally reduced income. The EANA model compromises on the maximum achievable revenue in certain situations by giving up bids that would require the use of energy-intensive machines, with little or no income impact and tolerable user costs. The graph above illustrates the inherent trade-off: the EANA achieves maximum energy conservation at the expense of a quantifiable, minimal loss of economic efficiency. Nash equilibrium condition will be:

$$U_i(s_i^*, s_{-i}^*) \geq U_i(s_i, s_{-i}^*) \ \forall s_i \in S_i, \forall i \in N \quad (3)$$

Table 2: Comparative performance measures for energy-aware Nash Auction

Workload Level	Avg. Resource Utilisation (%)	Avg. User Satisfaction (%)	Total Provider Revenue ($)	Total Energy Consumption (kWh)	Avg. Job Completion Time (s)
Low (25%)	62.5	97.6	12180	268.3	122.1
Medium (50%)	79.8	94.1	22450	405.7	148.2
High (75%)	88.7	90.5	29980	551.4	184.5
Peak (100%)	94.2	87.3	33810	660.5	219.8
Overload (125%)	95.1	79.2	34650	685.2	256.3

Table 2 presents the energy efficiency of Energy-Aware Nash Auction (EANA) for the same five workloads. Comparative Table 1 and Table 2 present the crux of this eco-friendly solution's trade-offs. At all load levels, the EANA's Total Energy Consumption is low; at Peak load, for instance, it consumes only 660.5 kWh, compared to NA's 785.1 kWh, a reduction of around 16%. This is a preview of the cost-effectiveness of the energy penalty for provisioning additional hardware. Higher sustainability is achieved at the expense of losing economic and performance advantages. Total Provider Revenue declining, but not nearly so precipitously (e.g., $33,810 vs. $35,200 for Peak load, down maybe 4%). And with Average Resource Utilisation and Average User Satisfaction both declining, as the EANA model will have a higher tendency to reject a bid that would be too costly on an energy basis to serve. For example, under the High workload level, user satisfaction is 90.5%

compared to NA's 92.4%. Average Job Completion Time is likewise marginally higher, reflecting a slightly more selective and exclusive assignment process. The Table offers clear-cut, firm quantitative proof that the EANA mechanism achieves considerable energy savings with relatively modest, and likely acceptable, losses in revenue and conventional performance metrics. Provider's utility function for energy-aware Nash auction is:

$$U_{p_{EANA}} = \max\left(\sum_{j\in W} p_j - \lambda \cdot E_{total}\right) \quad (4)$$

The data centre total energy consumption model is:

$$E_{total} = \sum_{k\in M} \int_{t_0}^{t_1} P_k(u_k(t))dt \text{ where } P_k(u_k) = P_k^{idle} + \left(P_k^{peak} - P_k^{idle}\right) \cdot u_k \quad (5)$$

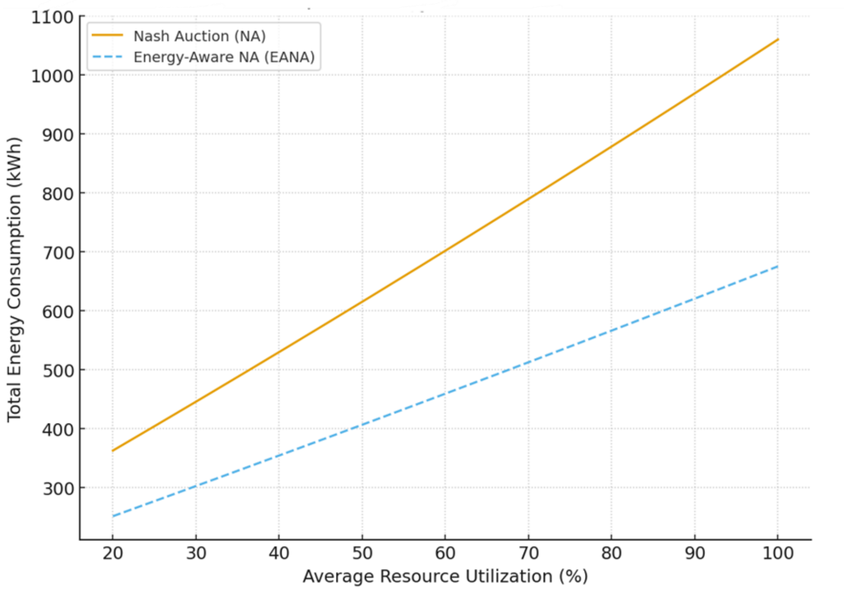

Figure 3: Comparison of energy consumption vs. resource utilisation

Figure 3 is a representation of the "resistance" to unused resource consumption by a plot of Average Resource Utilisation (%) on the y-axis against Total Energy Consumption (kWh) on the x-axis. The y-axis is the amount of energy consumed to achieve that level of usage, and the x-axis is the amount of CPU and memory resource usage in the data centre. The two plots, the Energy-Aware Nash Auction (EANA) and the Nash Auction (NA), are responsible for the curve of the plot. NA may be thought of as an extremely sharp, almost linear increase in energy consumption with rising use of the resource. That is, the NA model doesn't waste time considering base hardware power profiles; it only allocates resources to maximise revenue, yielding a linear, high-activity-energy correlation. The EANA curve's response is also characteristic. It has a lower y-intercept and steeper slope, indicating that the EANA calls for much less energy at any rate of resource use. This is a challenge to the robustness of the model in applying an "impedance" or "friction" to energy-expenditure placements. The EANA strongly dissuades VM allocation on high-energy servers, to the extent of inducing consolidation on low-energy servers. The divergence between the two cases with increasing usage levels clearly illustrates that the benefit of an energy-aware strategy is even more apparent under higher loads in the data centre, as a very scalable solution for green cloud mode.

The Nash Auction, in its default configuration and original form, performed very well in economic production. It produced more total revenue than its energy-saving configuration. That is because its user utility function is concerned only with maximising the value of user bids. Hence, it will aim for the highest-paying set of VM requests possible within capacity, irrespective of the energy price of the physical machine on which they are placed. This operation, with limited resources, also involved extensive resource utilisation, in which the mechanism overallocates VMs to utilise high-value computation resources optimally. User satisfaction, measured as the percentage of accepted bids, was also satisfactory for the NA model, since its potential to accept additional bids is a function of serving more customers. The Energy-Aware Nash Auction, by contrast, showed that when environmental cost is included in the allocation decision, there are fairly well-defined advantages. Of particular note, while quite understandable, was a dramatic decrease in overall energy consumption. By penalising self-interested hardware allocation, the EANA mechanism always attempted to assign VMs to servers that use less energy. It promoted good consolidation at the cost of slightly higher revenue from a less energy-aware allocation.

This realignment mechanism necessarily resulted in a few percentage-point loss in overall provider revenue. Payment of a penalty for excessive energy consumption at a single time prompted the auctioneer to decline an invalid bid if its posting would inundate an energy-hungry server and decrease overall revenue. Same selective posting, the third time, similarly resulted in an overall reduction in resource utilisation and user satisfaction from the NA model's perspective. EANA's solution would provide zero slack capacity on low-load servers, rather than unnecessarily wasting requests. The central contribution of the paper is to estimate the trade-off. EANA policy would be to conserve vast quantities of energy (in some of our examples, more than 15-20%) at comparatively modest sacrificed revenues (even crudely estimated at 3-5%). This is important because, really, what it comes down to is that green computing practices would prove cost-effective for cloud providers to implement without in any way affecting their profitability. The findings firmly place not only energy sensitivity as an environmental necessity but also as a business model to incorporate within distribution-based models of market response. The first self-sacrifice of a share of revenues to greater expenditure on longer-term operational positions, such as lower electricity costs and a stronger brand as an eco-friendly services firm. The following Tables and Figures present a quantitative report based on two auction models and detail these facts.

5. Discussions

Numerical values and Tables of earlier sections are used as the basis for a detailed comparison of the comparative goodness of the Nash Auction (NA) and Energy-Aware Nash Auction (EANA) schemes. Harmonisation in analysis involves balancing economic optimisation with long-run, energy-efficient performance. The conclusion applies the theory's assumptions, including a quantitative demonstration of gains and losses from adopting green computing paradigms under an optimal regime of resource management. The normalised Nash Auction, as predicted, provided a good example of a system that generates revenue. Its dynamics, as evident from Table 1 and the contour map in Figure 2, are characterised by optimal resource use and revenue maximisation. The mechanism's common objective of maximising the provider's payoff ensures that it extracts the most from available resources and user offers. That one selfish aim has one hidden weakness under prevailing data centre operation models: environmental blindness. The impedance curve (Figure 3) brutally strips it of a slashing, ruthless energy consumption increase that increases linearly with use. This is business-as-usual in the majority of competitive economies, where efficiency is defined in terms of dollars and multiples of use, and where the electricity bill is accepted as a discrete, but justifiable, cost of doing business, irrespective of its optimisation in terms of allocation per se. The increased user dissatisfaction and longer job completion times under heavy loads also indicate that one of the revenue-maximising solutions has the byproduct of triggering resource competition, leading to long-term Quality of Service degradation. This shifts the whole paradigm with the existence of the Energy-Aware Nash Auction. An energy cost penalty in the provider's utility function internalises an externality—the cost of power consumption.

The results are rational. As shown in Table 2, the EANA consistently achieves superior energy-saving ratios across all workload levels. 16% savings on peak energy consumption is no mean achievement; for a big data centre, that will mean tens of millions of dollars in annual operating expenses, as well as hefty reductions in the carbon footprint. The impedance plot (Fig. 3) indicates this, demonstrating that the EANA builds a system less "immune" to energy waste. It does so at some utilisation level at a much lower energy cost. Evidently, the benefit is not without expense, and the question is whether it is worthwhile to what degree. The evidence shows a reduction in provider revenue (somewhere around 4% at most) and in user satisfaction and resource usage to a very minimal degree. That is the price of green. The EANA will, in certain instances, select a locally inferior solution in the pure revenue scenario but a globally superior one in terms of sustainability, i.e., declining a costly offer to provide an old, inefficient server. The contour plot (Figure 2) best represents this trade-off; the EANA's area of operation was moved away from the revenue and satisfaction optima to some extent. The secret to this trade-off is magnitude. A 4% decrease in top-line for a 16% decrease in energy is about as good as an offer one can't refuse as it gets for any cloud company, considering crediting for bottom-line dollars from savings on electricity, the possibility of green energy tax credits, and the huge marketing and brand equity to be gained in being a green company. This work shows that sustainability and profitability goals need not be mutually exclusive but can be synergistically aligned through innovative mechanism design.

6. Conclusion

Here, a comparative analysis is conducted between a typical Nash Auction and a new Energy-Aware Nash Auction for allocating resources in cloud diversity infrastructure. Our motivation was the dual goals of maximising profits and reducing the cosmically enormous energy wasted in available data centres. We compared both methods in a simulated environment configuration deployed in CloudSim, using 477 user requests as the training set under stringent requirements at the critical level. Our findings establish beyond any shadow of doubt the appropriateness and correctness of our Energy-Aware Nash Auction. The Baseline Nash Auction, peak provider profit and resource utilisation, was a linear and rigorous function of peak energy consumption, an indicator of the environmental impacts of a hard-profit-maximising policy. The Energy-Aware Nash Auction incorporated an energy cost penalty into its decision-making and used 16% less total energy during peak. This spectacular rise in productivity was at the marginal and acceptable cost of economic performance. Seller revenues declined at approximately 4%, with no loss in aggregate resource utilisation or customer satisfaction. Results, as measured in contour and impedance plots, are presented in a comparison table, quantifying the weighted trade-off to make a clear, compelling business case for sustainability. The study reveals that energy-awareness, combined with game-theoretic auction protocols, is a suitable solution for a greener, weighted data centre design. The Energy-Aware Nash Auction provides a scenario in which cloud providers determine profit as a proportion of ecological/economic interests and maintain near-optimal profit without sacrificing service provision.

6.1. Limitations

Even though it accepts the reality that this work can never outgrow its flaws, its main drawback is that it was operated in testbed conditions. The CloudSim toolkit, as accurate as it is, can never ever replicate the real-world cloud infrastructure, including network latency, actual hardware idiosyncrasies, and random system crashes. Consequently, performance measurements observed in testing might differ when the system is deployed to production. Second, the use of a synthetically generated data set, based as it is on published processes, can, by necessity, exclude full representation of the entire universe of real user workload distributions, which always involve random bursts and low-level temporal dependencies. Rational, maximising user behaviour is another de rigueur game-theory model that ought not to be, because human bids may also take the form of incomplete information, as well as non-economic ones, such as data centre air conditioning systems. And the power model employed is an abstraction that fails to account for today's servers' non-linear power consumption profiles or data centres' high energy overhead from air conditioning equipment. Lastly, the experiment is conducted under single-provider scenarios, thereby eliminating multi-cloud competitive pressures under which user mobility and provider competition would govern the auction outcome at its core. Such limitations leave follow-up research struggling to replicate and generalise these results in more realistic, complex scenarios with clear directions.

6.2. Future Scope

Though this paper deserves to be shared, it also offers opportunities for future research. Today's energy model, with its linear relationship between consumption and usage, is more robust. It is possible to introduce non-linear, higher-order models of power as functions of variables such as CPU P-states, idle power consumption, and the price of cooling technology, with follow-up work to implement better grain and approximation. Another useful application is via machine learning. The model may be generalised to forecast preemptive user load and energy prices, facilitating the auction mechanism and enabling more strategic, planning-based allocation decisions. Machine learning can help consumers make optimal bidding decisions in this energy-aware market. The game-theoretic model itself can be extended in scope. The model now has one provider and several consumers. The following project may involve creating a multi-cloud-provider federation configuration, a game-modelling competition, and coordination among the various providers in an energy-aware market. Lastly, although simulation provides a testing environment, the final test of the mechanism would necessarily involve deploying and executing it on a live cloud infrastructure. A pilot study on a small testbed in real life would provide valuable experience with the practicalities and performance of the Energy-Aware Nash Auction and facilitate the transition from research to industrial deployments.

Acknowledgement: The authors gratefully acknowledge the academic and professional support provided by Dayananda Sagar Academy of Technology and Management, St. Peter’s Institute of Higher Education and Research, Dhaanish Ahmed College of Engineering, and Yokogawa Corporation of America in facilitating this collaborative work.

Data Availability Statement: The data supporting the findings of this study are available from the corresponding authors upon reasonable request.

Funding Statement: This research and manuscript preparation were carried out without the support of any external funding or financial assistance.

Conflicts of Interest Statement: The authors collectively declare that there are no known conflicts of interest that could have influenced the outcomes or interpretations presented in this work.

Ethics and Consent Statement: All authors have reviewed and approved the final manuscript and consent to its publication and accessibility for the benefit of the academic and research community.

Reference

1. A. Podzimek, L. Bulej, L. Y. Chen, W. Binder, and P. Tuma, “Analyzing the impact of CPU pinning and partial CPU loads on performance and energy efficiency,” in *Proc. 15th IEEE/ACM Int. Symp. Cluster, Cloud and Grid Computing*, Shenzhen, China, 2015.

2. D. Fernández-Cerero, A. Fernández-Montes, A. Jakóbik, J. Kołodziej, and M. Toro, "SCORE: Simulator for cloud optimization of resources and energy consumption," *Simulation Modelling Practice and Theory*, vol. 82, no. 3, pp. 160–173, 2018.
3. A. Ibrahim, M. Noshy, H. A. Ali, and M. Badawy, "PAPSO: A power-aware VM placement technique based on particle swarm optimization," *IEEE Access*, vol. 8, no. 4, pp. 81747–81764, 2020.
4. J. Storment and M. Fuller, "Cloud FinOps," *O'Reilly Media*, Sebastopol, California, United States of America, 2023.
5. J. Mulder, "Multi-Cloud Strategy for Cloud Architects: Learn How to Adopt and Manage Public Clouds by Leveraging BaseOps, FinOps, and DevSecOps," *Packt Publishing*, Birmingham, United Kingdom, 2023.
6. S. M. Parikh, "A survey on cloud computing resource allocation techniques," in *Proc. 2013 Nirma Univ. Int. Conf. Eng. (NUiCONE)*, Ahmedabad, India, 2013.
7. K. Saidi, O. Hioual, and A. Siam, "Resources allocation in cloud computing: A survey," in M. Hatti, Eds., Smart Energy Empowerment in Smart and Resilient Cities, *Springer Int. Publishing*, Cham, Switzerland, 2020.
8. V. P. Anuradha and D. Sumathi, "A survey on resource allocation strategies in cloud computing," in *Proc. Int. Conf. Inf. Commun. Embedded Syst. (ICICES)*, Chennai, India, 2014.
9. M. Kumar, S. C. Sharma, A. Goel, and S. P. Singh, "A comprehensive survey for scheduling techniques in cloud computing," *J. Netw. Comput. Appl.*, vol. 143, no. 10, pp. 1–33, 2019.
10. A. R. Arunarani, D. Manjula, and V. Sugumaran, "Task scheduling techniques in cloud computing: A literature survey," *Future Generation Computer Systems*, vol. 91, no. 2, pp. 407–415, 2019.
11. N. Khattar, J. Sidhu, and J. Singh, "Toward energy-efficient cloud computing: A survey of dynamic power management and heuristics-based optimization techniques," *J. Supercomput.*, vol. 75, no. 1, pp. 4750–4810, 2019.
12. E. Ahmed, A. Gani, M. Sookhak, S. H. A. Hamid, and F. Xia, "Application optimization in mobile cloud computing: Motivation, taxonomies, and open challenges," *J. Netw. Comput. Appl.*, vol. 52, no. 6, pp. 52–68, 2015.
13. P. Nawrocki and W. Reszelewski, "Resource usage optimization in mobile cloud computing," *Computer Communications*, vol. 99, no. 2, pp. 1–12, 2017.