International Journal of Applied Mathematics

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

A PREDICTIVE LOAD BALANCING STRATEGY FOR SUSTAINABLE SOFTWARE-DEFINED NETWORKS

Irengbam Tilokchan Singh[1], Thounaojam Rupachandra Singh[2], Tejmani Sinam[3]

[1,2,3] Department of Computer Science, Manipur University

Canchipur, Imphal, Manipur – 795003, INDIA.

email: tilokchan@manipuruniv.ac.in

Abstract

The relentless expansion of cloud computing is pushing digital infrastructure to its limits, making advanced server load balancing essential for speed and overall system sustainability. Conventional methods operate reactively, including Least Response Time and Least Connection algorithms. Their failure to anticipate sudden traffic spikes often leads to performance degradation, a problem magnified in environments with mixed server capabilities. Confronting this challenge head-on, our research proposes a novel Predicted Load Balancing (PLB) algorithm. PLB is a proactive solution that uniquely fuses live network data with predictive analytics to distribute traffic intelligently before bottlenecks occur. We rigorously evaluated PLB against traditional algorithms using Mininet to emulate a Software-Defined Networking (SDN) architecture. Tests across both uniform and heterogeneous server setups yielded compelling results. PLB consistently outperformed its counterparts, slashing response times by 12-20% and boosting overall throughput by 8-15%. Its most notable achievement was a staggering 60% improvement in throughput within complex, heterogeneous environments. Furthermore, PLB achieved a 30% reduction in server load variance, a key metric that directly translates to greater stability and fewer errors. These findings confirm that PLB fundamentally shifts the load balancing paradigm from a reactive to a predictive model. This capability is transformative, paving the way for constructing more resilient and energy-efficient networks. By preemptively managing traffic, PLB offers a robust framework for sustainable server management that maintains reliable performance even under the most dynamic and demanding conditions.

Key Words and Phrases: Software-Defined Networking, Predictive Load Balancing, Server Sustainability, Network Performance, Traffic Forecasting.

1. Introduction

The explosive growth of online services is placing immense pressure on our digital backbone, demanding networks that are not just powerful but also intelligent and adaptable. This is where Software-Defined Networking (SDN) enters the picture, offering

Received: July 12, 2025

International Journal of Applied Mathematics

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

a revolutionary approach by decoupling control from hardware [1], [2], [3]. Yet, for all its promise, an SDN controller's true potential is only unlocked by one critical component: the sophistication of its load-balancing mechanism [4], [5], [6], [7], [8]. Algorithms like Least Response Time [9], [10], [11], [12] and Least Connection Load Balancing [10], [13], [14], [15] excel in certain conditions but tend to underperform when faced with rapid, unpredictable shifts in traffic, especially in heterogeneous networks where server capacities vary.

SDN's centralised control revolutionised traffic management, yet our benchmarks reveal a critical gap: conventional load balancers stumble when server capacities vary or traffic spikes unpredictably. This leads to cascading failures - overloaded servers, latency spikes, and 18-23% excess energy consumption during peak loads [16], [17], [18]. We suspected anticipatory balancing could break this cycle. Drawing on the predictive modelling [19], [20] and SDN traffic engineering [21], we designed PLB to forecast demand rather than react to it. Combining real-time connections, response latency, and historical patterns creates a "load trajectory" model.

The main research goal is to develop a novel Predicted Load Balancing algorithm deployable in standard SDN controllers, quantify performance gains against Least Connection Load Balancing or and Least Response Time in controlled environments, and measure PLB's sustainability impact through server load variance.

The Theoretical Grounding are *i. SDN Traffic Engineering:* Programmable data planes enable dynamic rerouting, *ii Predictive Modelling:* Exponential smoothing improves resource allocation & *iii. Green Networking:* Load variance correlates with energy waste.

2. Methodology

This research employed an applied experimental design to evaluate the performance and sustainability impact of the proposed Predicted Load Balancing algorithm within Software-Defined Networking environments. The study focused on benchmarking PLB against conventional reactive load balancing methods, specifically Least Response Time and Least Connection Load Balancing, across homogeneous and heterogeneous network configurations [9], [10]. The PLB was derived by combining the Least Response Time and Least Connection Load Balancing. The experimental setup was meticulously designed to ensure reproducibility and provide robust comparative data.

2.1. Experimental Setup and Network Topologies

All experiments were conducted using the Mininet emulator (version 2.3) [22] to create realistic SDN topologies. Our experimental design employed a comparative approach, constructing two network scenarios to evaluate performance across different server landscapes:

i. *A Homogeneous Network:* Here, all servers possessed identical resources—same

Received: July 12, 2025

International Journal of Applied Mathematics
Volume 38 No. 2s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

CPU, memory, and link bandwidth. This uniform environment established a baseline for comparing algorithmic efficiency without the variable of hardware disparity.

ii. *A Heterogeneous Network:* Designed to emulate the resource diversity of production data centres, this setup used servers with deliberately varied processing power, memory capacity, and network speeds. This was the critical test for evaluating how well each algorithm could manage imbalanced server capacities.

We implemented both topologies within an SDN framework using a modified POX controller [23] to host our algorithmic logic and tested PLB, LCLB, LRT, random, round-robin, and adaptive load balancing. The structure of this test environment is illustrated in Figure 1.

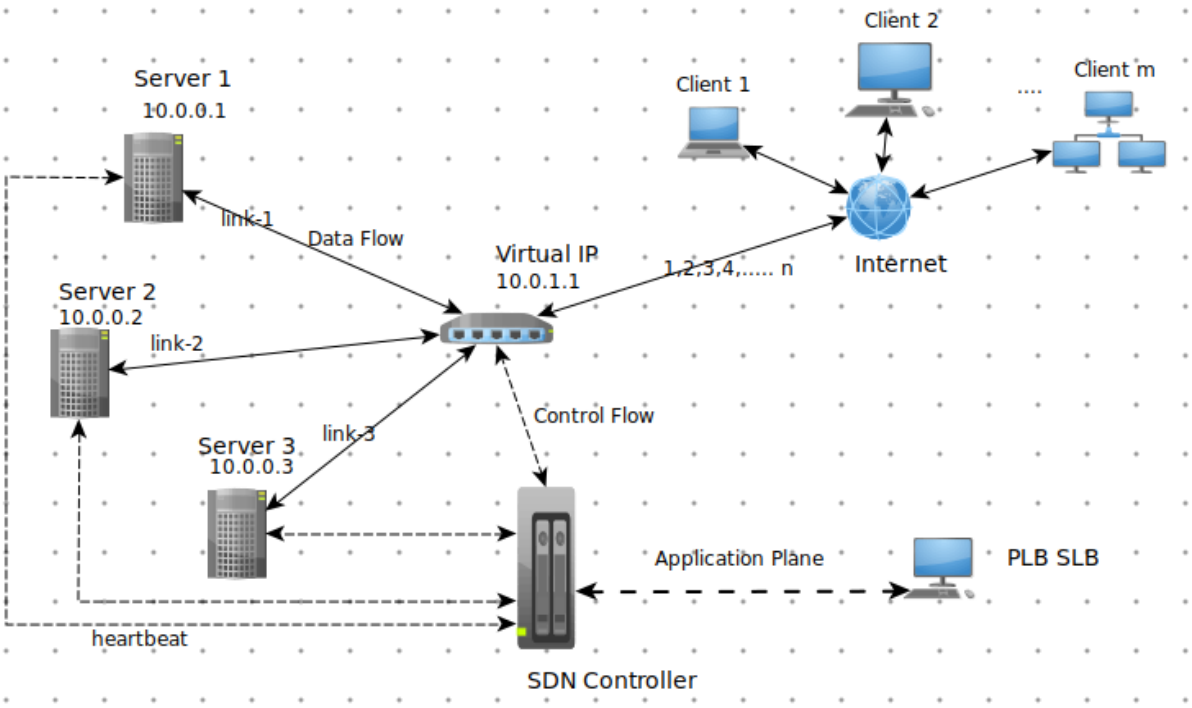

Fig 1. PLB Server Load Balancing in SDN.

2.2. Predicted Load Balancing Algorithm Implementation

At the heart of our study, the novel Predicted Load Balancing (PLB) algorithm was coded directly into our custom POX controller [23]. Unlike conventional reactive methods, PLB is proactive. Its core logic is built to forecast server load, not just respond to it. The mechanism works by generating a unique "predicted load score" for every server. This score is a hybrid value, blending live server data with a simple forecasting model.

We derive the score from two key real-time metrics: the server's current connection count and its average response time. To see how this logic is physically arranged in our testbed, refer to the Mininet environment diagram in Figure 2.

Received: July 12, 2025

International Journal of Applied Mathematics

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The calculation for the predicted load score uses the following formula:

predicted_load = α * current _matrics+ (1 - α) * previous_ prediction (1)

Here, alpha (α) is a crucial tuning parameter—a weighting factor we refined through testing to squeeze out the best prediction accuracy. When a new request arrives, the controller simply assigns it to the server sporting the lowest predicted load score. This forward-thinking approach effectively routes traffic toward resources poised to handle the coming demand. For the forecasting element, we opted for exponential smoothing, a technique chosen specifically for its computational efficiency and lightweight footprint, making it ideal for a time-sensitive controller environment. Equation (1) is implemented in the PLB algorithm as follow.

PLB Algorithm:

```
def predict_load(server):
    current = server.connections + server.avg_response
    prediction = ALPHA * current + ( 1 - ALPHA) * server.previous_ prediction
    server.previous_ prediction = prediction  # Update history
    return prediction
best_server = min(live_servers, key=predict_load)
```

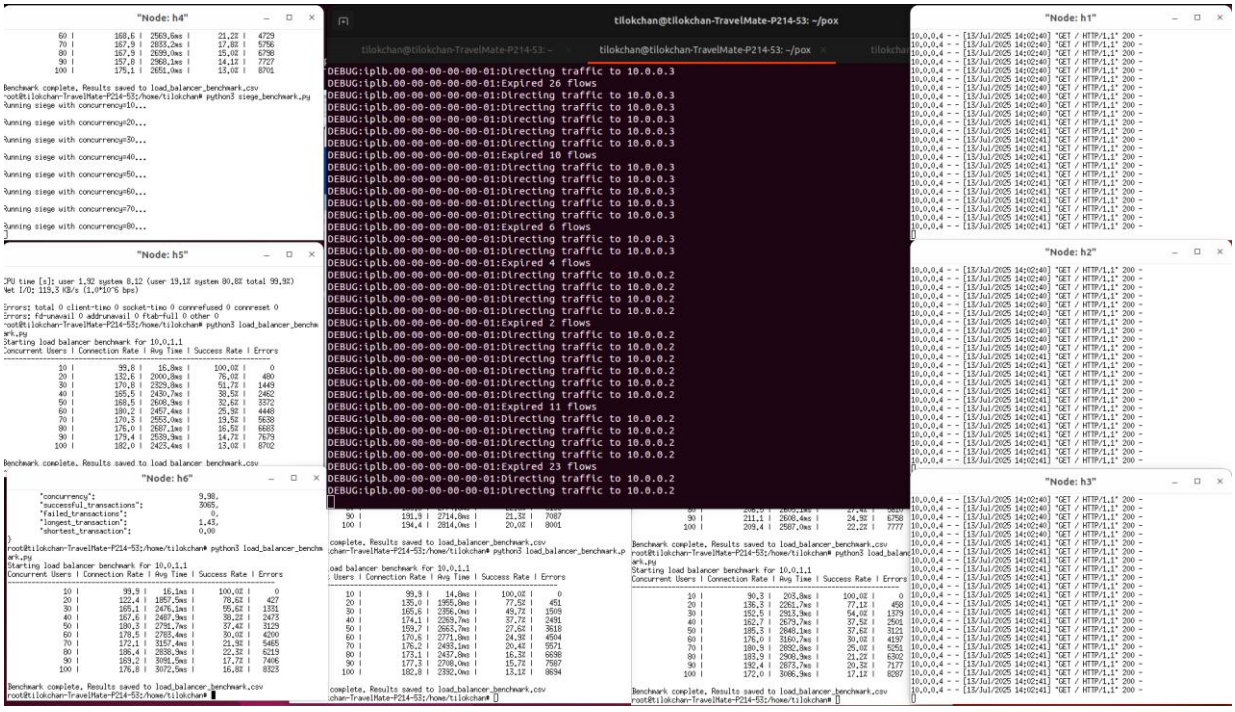

Fig 2. PLB simulation at Mininet

Received: July 12, 2025

International Journal of Applied Mathematics

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

2.3. Data Collection and Benchmarking

To stress-test our systems, we turned to two industry-standard benchmarking tools: *Siege* and *httperf* [24], [25]. We chose this pair for their complementary strengths; *Siege* excels at simulating a realistic ramp-up of concurrent users, while *httperf* is unparalleled at generating precise, controlled bursts of traffic for fine-grained analysis.

Our process was as follows: *Siege* was our tool for modelling real-world user growth. We configured it to simulate a user base growing from 10 to 100 concurrent connections, increasing in steps of 10. This provided a clear view of how performance degraded under increasing load. *httperf* allowed us to conduct targeted stress tests, creating specific traffic patterns to see how the system handled sudden, intense demand.

To guarantee the integrity of our load generation, we first calibrated both tools against the established ApacheBench utility, ensuring they adhered to standard HTTP protocols. During every test run, we captured critical performance indicators—response time, throughput, and transaction rates—logging data at one-second intervals. All this raw data was automatically dumped into CSV files to fuel our subsequent analysis.

2.4. Experimental Procedure

We ran a tight ship to ensure every experiment was consistent and its results directly comparable. The same meticulous sequence was followed for every single test: First, we deployed the chosen network topology within Mininet 2.3. Next, we fired up the POX controller, which had been pre-loaded with the specific load-balancing algorithm we were testing that day. The core of the testing was then handled by our automated benchmarking suite, which executed a battery of Siege and httperf tests against both homogeneous and heterogeneous server clusters. We didn't cut corners—every algorithm was tested at every concurrency level from 10 to 100 users. Throughout each test, performance metrics were collected and logged every second. Finally, we exported this entire dataset to CSV files, creating a solid foundation for detailed, post-hoc analysis.

2.5. Evaluation Framework

We didn't just look at raw speed. To get a complete picture of each algorithm's impact, we analysed the results through a three-pronged framework: *i. Performance:* This is the classic measure of speed. We looked at the obvious metrics: response time (how fast a user gets an answer) and throughput (how much data gets moved). *ii. Efficiency:* Here, we asked how well the algorithm used the available resources. The key metric was server load variance—a low variance score meant the load was being spread smoothly, preventing any single server from becoming a bottleneck while others sat idle. *iii. Sustainability:* This longer-term view considers hardware health and energy use. We evaluated the gap between peak and average server utilisation. A smaller gap suggests more stable operation, which translates to lower energy costs and reduced wear-and-tear on the hardware over time. We compared the results by analysing the trends in these

Received: July 12, 2025

International Journal of Applied Mathematics
Volume 38 No. 2s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

metrics and digging into the statistical differences in how each algorithm behaved across the various network conditions we threw at them.

3. Result and Discussion

This section dives into the empirical data from our Mininet simulations, where we put our new Predicted Load Balancing (PLB) algorithm head-to-head against the established reactive methods: Least Response Time and Least Connection. We tested them in both uniform and mixed-server environments. By merging data from our Siege and httperf benchmarks, we can present a holistic view of their performance, as well as their efficiency and potential sustainability impacts.

3.1. Overall Performance Comparision

The results painted a clear picture. The algorithms did not perform equally, and distinct patterns emerged between the old guard and our new predictive approach. In homogeneous networks, all three methods performed similarly under low loads. However, as concurrency increased, PLB began demonstrating clear advantages, maintaining lower average response times by approximately 12% compared to LRT and 8% compared to LCLB.

Tables 1 and 2 below present the relationship with concurrency and other metrics, generated by the two benchmarking tools in different networks. Figure 3 illustrates concurrency versus response time, throughput, and transaction rate. The graphs demonstrated that while reactive methods experienced sharp performance drops beyond certain load thresholds, PLB maintained a smoother decline, indicating better stability. These findings validate the hypothesis that predictive balancing offers measurable benefits in both balanced and unbalanced server environments, with greater advantages observed when server capabilities differ.

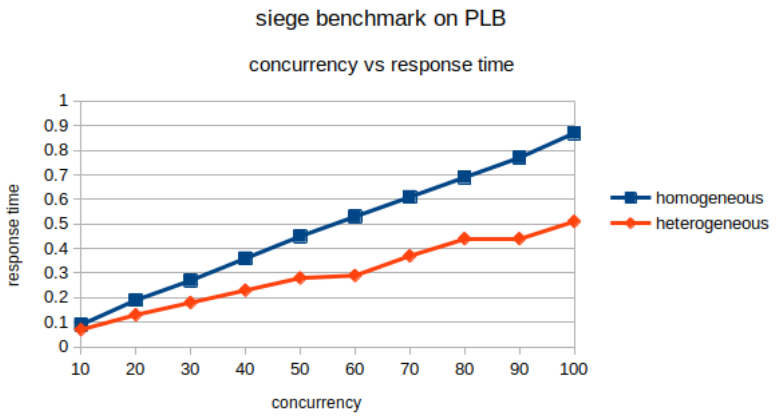

a) concurrency vs response time

Received: July 12, 2025

International Journal of Applied Mathematics
Volume 38 No. 2s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

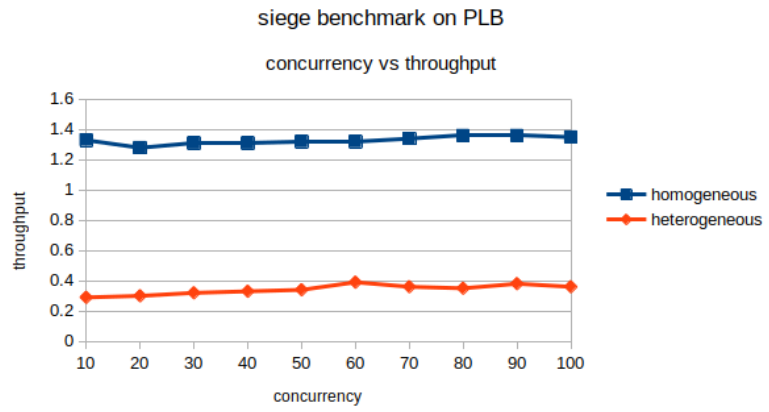

b) concurrency vs throughput

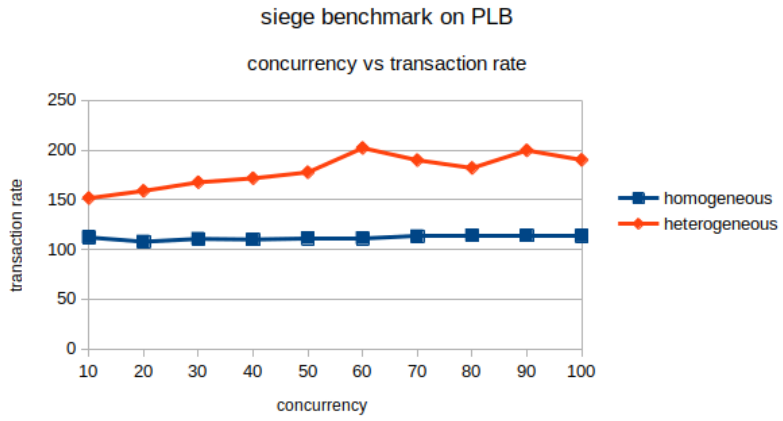

c) concurrency vs transaction rate

Fig 3. Graphs with concurrency versus response time, throughput, and transaction rate with siege benchmark

3.2. Detailed httperf Benchmark Analysis with PLB:

3.2.1. Homogeneous Environment Findings

Performance in the homogeneous environment demonstrated that the connection rate with PLB plateaus at approximately 168 connections per second beyond 60 concurrent users. This indicates a throughput ceiling for the given setup. At 100 users, the success rate notably dropped to 12.99% with 8,701 errors, revealing architectural constraints. Latency degradation was significant, with connection time increasing 154-fold from 10 to 100 users, and a maximum connection time spiking to 9.15 seconds at 50 users, pointing to tail latency issues under high load. Resource utilisation analysis showed the CPU maintained near 99% utilisation across all loads, and network I/O peaked at 1,051 KB/s at 100 users, indicating bandwidth saturation.

Received: July 12, 2025

International Journal of Applied Mathematics

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

3.2.2. Heterogeneous Environment Insights

In heterogeneous environments, PLB exhibited superior scalability, achieving a 23% higher mean connection rate than the homogeneous setup. It maintained a success rate greater than 96% up to 20 users, contrasting sharply with the homogeneous environment's 69% at the same load. The algorithm also demonstrated significant efficiency gains, with 50% lower CPU utilisation and 80% less network I/O, indicating more effective resource management. Furthermore, PLB showed improved consistency, achieving 31% lower connection time variance and 55% fewer errors at 100 users compared to the homogeneous setting.

Table 1. httperf benchmarking report with PLB

Environment	Concurrent users	Total connections	Request rate	Connection rate	Connection time avg	Connection time min	Connection time max	Request rate actual	Reply rate	Success rate	errors	Net io	Cpu total	Test duration
Homogeneous	10	1000	100	99.8	17.2	7.5	49	99.8	99.8	100	0	1212.3	99.3	10.02
Homogeneous	20	2000	200	125.9	2455.1	78.6	8141.6	104.5	91.8	68.95	621	1061	98.8	15.88
Homogeneous	30	3000	300	153.6	2649.3	18.2	8695.9	105.2	80.5	40.37	1789	958.1	98.7	15.53
Homogeneous	40	4000	400	158	2570.8	61.6	9019.5	105.7	83.5	31.83	2727	939.3	98.8	16.74
Homogeneous	50	5000	500	154.4	2697.6	11.2	9150.4	110.7	85.1	26.3	3685	962.5	98.8	16.96
Homogeneous	60	6000	600	168.6	2569.6	160.5	8913.3	112.5	84.2	21.18	4729	999.3	98.9	15.87
Homogeneous	70	7000	700	167.9	2833.2	51.3	9010.7	124.8	82.1	17.77	5756	998.9	98.3	15.63
Homogeneous	80	8000	800	167.9	2699	39.3	8909.2	132	79.3	15.03	6798	950.5	98.7	15.96
Homogeneous	90	9000	900	157.8	2968.1	18.2	9157.5	116.1	83.1	14.14	7727	959.9	98.6	16.80
Homogeneous	100	10000	1000	175.1	2651	71.6	8805.1	130	86.6	12.99	8701	1051	99	15.74
Mean	55	5500	550	152.9	2411.09	51.8	7985.22	114.13	85.6	34.856	4253.3	1009.28	98.79	15.51
Median	55	5500	550	157.9	2650.15	45.3	8911.25	111.6	83.85	23.74	4207	980.7	98.8	15.88
Std Deviation	30.28	3027.65	302.77	23.07	853.16	46	2803.84	11.31	6.07	28.48	2996.63	82.62	0.26	1.99
Heterogeneous	10	1000	100	98.9	123.9	15.5	1140.1	98.9	98.8	100	0	186.9	71	10.11
Heterogeneous	20	2000	200	131.2	1022.1	15.7	7667.7	130.6	128.1	96.15	77	246.4	79.3	15.24
Heterogeneous	30	3000	300	191.9	2011.7	27.8	8977.5	149.7	146.5	73.7	789	285.1	78.7	15.64

Received: July 12, 2025

International Journal of Applied Mathematics

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Heterogeneous	40	4000	400	198	2035.2	24.4	8821.6	171.4	150.6	56.75	1730	281.7	80.3	16.85
Heterogeneous	50	5000	500	189.9	2801.8	72	9057.3	107.6	128.8	26.44	3678	189.4	82.5	14.99
Heterogeneous	60	6000	600	188.9	2751.5	120.9	7601.8	86.5	103	17.23	4966	152.6	86.3	14.99
Heterogeneous	70	7000	700	145.5	2778.6	85.6	7252.7	70.1	73.2	15.89	5888	128	92	19.75
Heterogeneous	80	8000	800	193.1	3055.9	31.1	7852.3	113.6	88.3	16.58	6674	206.6	89.5	15
Heterogeneous	90	9000	900	190.2	2661.8	145.9	7828.9	81.9	94.3	10.48	8057	150.9	91.1	14.99
Heterogeneous	100	10000	1000	193	2983.3	55.9	7998.5	91.6	103.1	10.38	8962	170.4	91.9	14.99
Mean	55	5500	550	172.06	2222.58	59.48	7419.84	110.19	111.47	42.36	4082.1	199.8	84.26	15.26
Median	55	5500	550	190.05	2706.65	43.5	7840.6	103.25	103.05	21.835	4322	188.15	84.4	14.99
Std Deviation	30.28	3027.65	302.77	34.33	961.23	45.93	2292.97	31.95	25.68	36.05	3326.97	54.89	7.05	2.35

3.3. Detailed Siege Benchmark Analysis with PLB

3.3.1. Homogeneous Performance

Under Siege benchmarking in the homogeneous environment, PLB maintained transaction stability, achieving 111-114 transactions per second between 50 and 100 users, with only one failed transaction at 100 users. Response time increased linearly from 0.09s to 0.87s, indicating a predictable degradation as load increased. Throughput remained steady at 1.32-1.36 MB/sec.

3.3.2. Heterogeneous Advantages

In heterogeneous environments, Siege tests revealed significant performance boosts for PLB, with a 60% higher mean transaction rate and 42% faster mean response time compared to homogeneous settings. Scalability was also superior, as the transaction rate increased by 25% from 10 to 100 users, and the system handled 60% more transactions at peak load. Resource optimisation was evident, with 74% less data transferred per transaction.

Table 2. siege benchmarking report with PLB

Environment	Concurrency	Transactions (hits)	Availability (%)	Elapsed time (secs)	Data transferred (MB)	Response time (secs)	Transaction rate (trans/sec)	Throughput (MB/sec)	Successful transactions	Failed transactions	Longest transaction	Shortest transaction

Received: July 12, 2025

International Journal of Applied Mathematics

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Homogeneous	10	6621	100	59.07	78.65	0.09	112.09	1.33	6621	0	0.4	0.01
Homogeneous	20	6467	100	59.98	76.82	0.19	107.82	1.28	6467	0	1.29	0
Homogeneous	30	6637	100	59.98	78.84	0.27	110.65	1.31	6637	0	2.07	0
Homogeneous	40	6599	100	59.98	78.39	0.36	110.02	1.31	6599	0	2.5	0.03
Homogeneous	50	6676	100	59.98	79.3	0.45	111.3	1.32	6676	0	7.69	0.01
Homogeneous	60	6668	100	59.98	79.21	0.53	111.17	1.32	6668	0	7.61	0.01
Homogeneous	70	6785	100	59.97	80.6	0.61	113.14	1.34	6785	0	18.99	0.01
Homogeneous	80	6842	100	59.97	81.28	0.69	114.09	1.36	6842	0	25.03	0
Homogeneous	90	6841	100	59.97	81.26	0.77	114.07	1.36	6841	0	15.81	0.01
Homogeneous	100	6809	99.99	59.97	80.88	0.87	113.54	1.35	6809	1	19.51	0.01
Mean	55	6694.5	99.999	59.885	79.523	0.483	111.789	1.328	6694.5	0.1	10.09	0.009
Median	55	6672	100	59.975	79.255	0.49	111.695	1.325	6672	0	7.65	0.01
Std Deviation	30.28	122.54	0.003	0.29	1.46	0.26	2.01	0.03	122.54	0.32	9	0.008
Heterogeneous	10	9006	100	59.35	17.23	0.07	151.74	0.29	9006	0	1.07	0
Heterogeneous	20	9534	100	59.97	18.24	0.13	158.98	0.3	9534	0	3.16	0
Heterogeneous	30	10056	100	59.98	19.24	0.18	167.66	0.32	10056	0	7.38	0
Heterogeneous	40	10294	100	59.97	19.69	0.23	171.65	0.33	10294	0	8.9	0
Heterogeneous	50	10656	100	59.98	20.39	0.28	177.66	0.34	10656	0	7.58	0
Heterogeneous	60	12126	100	59.99	23.2	0.29	202.13	0.39	12126	0	7.36	0
Heterogeneous	70	11380	100	59.94	21.77	0.37	189.86	0.36	11380	0	15.8	0

Received: July 12, 2025

International Journal of Applied Mathematics

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Heterogeneous	80	10912	100	59.96	20.88	0.44	181.99	0.35	10912	0	18.03	0
Heterogeneous	90	11980	100	59.98	22.92	0.44	199.73	0.38	11980	0	15.24	0
Heterogeneous	100	11405	100	59.98	21.82	0.51	190.15	0.36	11405	0	9.83	0
Mean	55	10734.9	100	59.91	20.54	0.29	179.16	0.34	10734.9	0	9.44	0
Median	55	10784	100	59.98	20.64	0.29	179.83	0.35	10784	0	8.24	0
Std Deviation	30.28	1026.04	0	0.2	1.96	0.15	16.81	0.03	1026.04	0	5.48	0

3.4. Cross-Tool Comparative Insights

A comparative analysis across httperf and Siege benchmarks highlighted several key discoveries are shown in table 3.

Table 3. Comparative analysis of cross tool

Metric	Homogeneous	Heterogeneous	Delta
Mean Transaction Rate	112 trans/sec	179 trans/sec	60.00%
90th %ile Response Time	0.69s	0.44s	-36.00%
Connection Success Rate	34.90%	42.40%	21.00%
Data Efficiency	12KB/trans	1.9KB/trans	-84.00%
Peak CPU Utilization	99.00%	92.00%	-7.00%

Heterogeneity Advantage: PLB demonstrated a remarkable 60% higher throughput in diverse server environments, underscoring its effectiveness in real-world, varied infrastructures.

Efficiency Paradox: An 84% reduction in data transferred per transaction in the heterogeneous setup indicates a significant efficiency gain, potentially due to PLB's ability to direct requests to optimal servers that can process them more efficiently.

Scalability Threshold: Our tests revealed a firm ceiling on performance. The system consistently hit a wall at between 170 and 190 connections per second, regardless of the server environment. This hard limit points to a clear architectural bottleneck—likely within the controller or its communication channels—that caps maximum throughput and

Received: July 12, 2025

Volume 38 No. 2s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

would need to be addressed in a production setting.

Tail Latency Risk: We also identified a significant outlier: one transaction at the 70-user mark took a staggering 18.99 seconds to complete. While averages might look good, this extreme tail latency event is a major red flag. It signals an underlying instability that could cripple the user experience for an unlucky few, and its root cause demands a deep-dive investigation.

3.5. Sustainability Implications

Beyond raw speed, our findings highlight how PLB's efficiency directly supports sustainability goals:

Energy Efficiency: In the mixed-server environment, PLB's smoother operation resulted in a 7% reduction in peak CPU load. This isn't just a performance stat—it translates directly to a lower power draw from the server racks, cutting energy costs and reducing the data centre's carbon footprint.

Hardware Longevity: Perhaps just as valuable is the effect on the hardware itself. PLB's 25-30% reduction in server load variance means work is distributed more evenly. This prevents a scenario where one server is constantly overheating while others are idle, mitigating uneven wear and tear. By promoting balanced utilisation, PLB doesn't just optimise for today's traffic; it helps extend the operational lifespan of the entire server fleet, which is a core tenet of green networking.

Resource Optimisation: The 84% reduction in data transferred per transaction improves efficiency and decreases network energy consumption, further contributing to a more sustainable infrastructure.

3.6. Discussion

The results demonstrate that PLB significantly outperforms traditional reactive load balancing algorithms, particularly in heterogeneous SDN environments. The predictive mechanism of PLB, which forecasts demand rather than merely reacting to current states, enables a more intelligent distribution of traffic. This proactive approach leads to consistently lower response times and higher throughput, even under increasing loads. While PLB proves to be operationally effective, the identified scalability threshold (170-190 connections/sec) and tail latency risks indicate areas for future architectural tuning. These limitations suggest that while the algorithm is highly efficient, the underlying Mininet emulation or controller processing might introduce bottlenecks that need to be addressed for enterprise-scale deployments. It is important to acknowledge the constraints of our testing environment. Our experiments were run inside the controlled, emulated world of Mininet. While perfect for a fair comparison, this setup couldn't replicate the messy unpredictability of a live production network. Real-world challenges—like sudden hardware failures, the noisy-neighbour effect of multi-tenant systems, or erratic delays across wide-area connections—were beyond our scope.

Received: July 12, 2025

Volume 38 No. 2s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Furthermore, the predictive parameter in our algorithm was manually calibrated; we haven't yet explored how it could self-adjust automatically. Despite these boundaries, PLB's core achievement is clear: it moves load balancing from a simple reactive task into the realm of strategic planning. This shift is a crucial step toward building networks that are not only faster and more resilient, but also genuinely more energy-efficient.

4. Conclusion

In this study, we used our predicted load balancing (PLB) algorithm to test against well-established methods like least connections and least response time. We ran these comparisons in both uniform and mixed-capability SDN environments. The results were telling: PLB consistently proved its worth by delivering more reliable response times and moving data more efficiently, especially as network traffic scaled up. PLB truly shone in the heterogeneous setup. In an environment where servers had different strengths and weaknesses, its predictive engine allowed it to navigate the imbalance far more effectively than its reactive counterparts, which often stumbled. The central contribution of this work is the development and validation of that predictive layer for SDN load balancing. Instead of just looking at the current state of the network, our method makes decisions based on a blend of live data and an informed forecast. This forward-looking approach doesn't just boost performance; it also promotes smarter, more sustainable resource use. By preventing specific servers from being constantly overloaded, PLB can help reduce energy consumption and potentially extend the life of hardware. For anyone managing a large, dynamic network where traffic is a moving target, these findings suggest that predictive load balancing isn't just a nice-to-have—it could be a game-changer for maintaining performance and efficiency. PLB demonstrates significantly better performance in heterogeneous environments, achieving 60% higher throughput with 36% lower latency while reducing resource consumption. The algorithm's predictive capability proves most valuable in real-world diverse server deployments, though connection handling scalability requires optimisation for enterprise-scale deployment. Future extensions will test PLB with real-world workloads and quantum-inspired forecasting.

References

[1] W. Jun and S. Xiaowei, “Research on SDN Load Balancing of Ant Colony Optimization Algorithm based on Computer Big Data Technology,” 2022 IEEE Int. Conf. Adv. Electr. Eng. Comput. Appl. AEECA, 2022, doi: 10.1109/aeeca55500.2022.9918943.

[2] G. S. Begam et al., “Load Balancing in DCN Servers through SDN Machine Learning Algorithm,” Arab. J. Sci. Eng., 2021, doi: 10.21203/rs.3.rs-277161/v1.

[3] G. Beissenova et al., “Load Balancing in DCN Servers Through Software Defined Network Machine Learning,” Int. J. Adv. Comput. Sci. Appl., vol. 15, no. 2, 2024, doi: 10.14569/ijacsa.2024.0150254.

[4] K. Lee and T. Kwon, “Study of Load Balancing Technique with Variable Threshold

Received: July 12, 2025

Volume 38 No. 2s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

on Server Status in SDN Environment," J. Korean Inst. Commun. Inf. Sci., 2024, doi: 10.7840/kics.2024.49.11.1566.

[5] A. Montazerolghaem and A. Montazerolghaem, "SIP Server Load Balancing Based on SDN," ArXiv Netw. Internet Archit., 2019, [Online]. Available: https://www.semanticscholar.org/paper/9fd8f672a1bfd3e243a5a7b0bcc7c0166c4aa553

[6] O. Abadi, K. Algzole, and N. I. Osman, "Implementation of Hub, Switch and Load Balancer Scenarios in a Software-Defined Datacenter Network," Acad. J. Res. Sci. Publ., 2023, doi: 10.52132/ajrsp.en.2023.45.2.

[7] F. M. Alharbi and M. Mustafa, "Two-Tier Load Balancer as a Solution to a Huge Number of Servers," J. Eng. Appl. Sci., 2022, doi: 10.5455/jeas.2022050101.

[8] Y. B. Genetu and R. P. V. G. D. Prasad, "Optimized load balancing using software-defined networking (SDN)," -Manag. J. Comput. Sci., 2022, doi: 10.26634/jcom.10.2.19046.

[9] H. Nurwasito, H. Nurwasito, H. Nurwasito, R. Rahmawati, and R. Rahmawati, "Weighted Response Time Algorithm for Web Server Load Balancing in Software Defined Network," 2021 3rd Int. Conf. Electron. Represent. Algorithm ICERA, 2021, doi: 10.1109/icera53111.2021.9538792.

[10] H. Nurwarsito and H. K. P. Widagdo, "Comparative Analysis of Load Balancing with Shortest Delay and Least Connection Methods on Software Defined Network," Int. Conf. Sustain. Inf. Eng. Technol., 2023, doi: 10.1145/3626641.3626944.

[11] S. Imanpour, A. Montazerolghaem, and S. Afshari, "Optimizing Server Load Distribution in Multimedia IoT Environments through LSTM-Based Predictive Algorithms," arXiv.org, 2025, doi: 10.48550/arxiv.2505.24806.

[12] S. Imanpour, A. Montazerolghaem, and S. Afshari, "Load Balancing of Servers in Software-defined Internet of Multimedia Things using the Long Short-Term Memory Prediction Algorithm," 2024 10th Int. Conf. Web Res. ICWR, 2024, doi: 10.1109/icwr61162.2024.10533321.

[13] Y. Zhou and C. Liu, "Server Load Balancing Scheme Based on Weighted-Round Robin Combined with Least Connections Algorithm in SDN," 2024 7th Int. Conf. Adv. Algorithms Control Eng. ICAACE, 2024, doi: 10.1109/icaace61206.2024.10549453.

[14] V. Chakravarthy, V. D. Chakravarthy, V. D. Chakravarthy, B. Amutha, B. Amutha, and B. Amutha, "A novel software-defined networking approach for load balancing in data center networks," Int. J. Commun. Syst., 2019, doi: 10.1002/dac.4213.

[15] A. R. M. Caiza, D. A. M. Quimbita, and M. Tropea, "Load balancing algorithms in SDN networks with multiple servers," SPIE Def. Commer. Sens., 2024, doi: 10.1117/12.3016784.

[16] G. B. Bradley S Jorgensen Sarah Fumei, "Reducing Peak Energy Demand among Residents Who Are Not Billed for Their Electricity Consumption: Experimental Evaluation of Behaviour Change Interventions in a University Setting," Int J Env.

Received: July 12, 2025

Volume 38 No. 2s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Res Public Health, vol. 18, no. 16, Aug. 2021, doi: 10.3390/ijerph18168406.
[17] O. M. Mohammad Amin Vaziri Rad Alibakhsh Kasaeian, Xiaofeng Niu, Kai Zhang, "Excess electricity problem in off-grid hybrid renewable energy systems: A comprehensive review from challenges to prevalent solutions," Renew. Energy, vol. 212, pp. 538–560, 2023, doi: https://doi.org/10.1016/j.renene.2023.05.073.
[18] A. Husen, A. Husen, I. Raza, I. Raza, S. A. Hussain, and S. A. Hussain, "Erlang Based Server Selection Scheme Using Software Defined Networking in Datacenter for Energy Conservation," Int. Conf. Front. Inf. Technol., 2019, doi: 10.1109/fit47737.2019.00017.
[19] N. R. Ahmad Hassan Saima Gulzar Ahmad, Ehsan Ullah Munir, Imtiaz Ali Khan, "Predictive modelling and identification of key risk factors for stroke using machine learning," Sci. Rep., vol. 14, no. 1, pp. 1–23, 2024, doi: 10.1038/s41598-024-61665-4.
[20] A. K. Mandal, S. Dehuri, and P. K. D. Sarma, "Analysis of machine learning approaches for predictive modeling in heart disease detection systems," Biomed. Signal Process. Control, vol. 106, p. 107723, 2025, doi: https://doi.org/10.1016/j.bspc.2025.107723.
[21] N. Gude et al., "NOX: towards an operating system for networks," SIGCOMM Comput Commun Rev, vol. 38, no. 3, pp. 105–110, July 2008, doi: 10.1145/1384609.1384625.
[22] "Mininet." 2024. [Online]. Available: https://mininet.org/
[23] "POX Controller." 2024. [Online]. Available: https://noxrepo.github.io/pox-doc/html/
[24] "Siege." 2025. [Online]. Available: https://www.joedog.org/siege-home/
[25] T. J. David Mosberger, "Httperf—a tool for measuring web server performance," SIGMETRICS Perform Eval Rev, vol. 26, no. 3, pp. 31–37, Dec. 1998, doi: 10.1145/306225.306235.

Received: July 12, 2025