Minimax Theory and its Applications
Volume 11 (2026), No. 4s, 816 to 827

Nonlinear Optimization of Computational Resource Allocation Under Time-Varying Workload Constraints

Uma Maheswari G[1], Dr. Anitha P[2], Dr. G. Geetha Ramani[3], Dr. A. Vanathi[4], Dr. V. Kumaraswamy[5], Narasaiah Gallamalla[6]

[1]Assistant Professor, Department of Mathematics, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Email: umamahes@maher.ac.in, ORCID ID: 0009-0006-4627-3102
[2]Assistant Professor, PG & Research Department of Information Technology, Sri Ramakrishna College of Arts and Science, Email: anitha@srcas.ac.in, ORCID ID: 0000-0002-2725-7418
[3]Department of Mathematics, New Prince Shri Bhavani College of Engineering, Email: geetharamani@newprinceshribhavani.com, ORCID ID: 0009-0008-9311-5224
[4]Associate Professor, Department of Computer Science and Engineering, Aditya University, Surampalem, Andhra Pradesh – 533437, Email: vanathi.andiran@adityauniversity.in, ORCID ID: 0000-0001-9127-0353
[5]Professor and Dean Academics, Department of Electronics and Communication Engineering, Sreenidhi Institute of Science and Technology, Hyderabad, Email: kumaraswamy.v@sreenidhi.edu.in, ORCID ID: 0009-0005-8887-9260
[6]Assistant Professor, Department of Computer Science and Engineering, Marri Laxman Reddy Institute of Technology and Management, Dundigal, Email: moses@mlritm.ac.in, ORCID ID: 0009-0005-6822-1867

Abstract

Modern computational systems need to allocate computational resources among many workloads. The demand of each workload. The available computational capacity changes over time. This paper develops a continuous-time optimization framework that takes into account the instantaneous capacity limits and the total workload service needs. The proposed model is a convex optimization problem. Its constraints depend on time and have a Volterra-type shape. In the first proof, that the optimal resource allocation exists and is unique. Next,the developed a Karush–Kuhn–Tucker characterization to describe the optimal solution. The main contribution of this study is an explicit sensitivity result. It shows that changes in capacity and workload arrival patterns cause at most square-root changes in the optimal allocation, provided a uniform Slater-type feasibility condition holds. The derived Hölder stability constant is directly linked to the workload dimension, the planning horizon, the bound on capacity, how the objective function is conditioned and the feasibility margin. In addition, a temporal stability result is derived for a sequence of environments where system conditions change continuously. Numerical experiments confirm the theoretical sensitivity behaviour. Demonstrate how the feasibility margin influences the robustness of the resource allocation. The results give a foundation for assessing how robust and adaptable optimal computational resource allocation systems are when workload conditions vary over time.

Keywords: Computational Resource Allocation; Cumulative Workload Constraints; Continuous-Time Optimization; Time-Varying Optimization; Sensitivity Analysis; Hölder Stability; Nonlinear Optimization.

1. Introduction

Modern computer systems increasingly operate in environments where resources are available and workload demands change all the time. Cloud computing, edge computing, distributed computing, communication networks and other cyber-physical infrastructures must allocate computer resources among workloads while still satisfying service requirements. Because computational capacity, workload arrival rates, processing efficiency and operating costs can change during the planning period, resource allocation cannot be treated as an optimization problem. Therefore, time-varying convex optimization and dynamic resource allocation receive attention in network systems (Davydov *et al.* 2025; Rostami & Kia 2025; Cheng *et al.* 2026; Yu, 2026).

A key feature of workload is that service requirements are cumulative, not instant. Workloads may arrive continuously. A minimum amount of computational processing over an interval, even when instant demand changes. Satisfying demand at one instant does not guarantee enough accumulated service. Under the workload constraint, the total service up to each time must meet the accumulated demand. This cumulative workload

constraint leads to a Volterra-type operator because the service now depends on resource allocation before creating an optimization structure.

The problem has two main time-dependent restrictions: computational capacity and cumulative workload requirement. Both are important in data centres, edge computing, communication networks and other systems with changing resources and workloads. Earlier studies show the importance of long-term workload and resource constraints in constrained convex optimization (Cao *et al.* 2018; Ding *et al.* 2021). Recent research extends this idea to time-varying constrained optimization and resource allocation (Liu *et al.* 2025; Cheng *et al.* 2026).

Time-varying optimization is also studied from a control and dynamical-system perspective. Existing studies look at equilibrium tracking, efficient solution tracking, real-time constrained optimization and nonlinear systems (Davydov *et al.* 2025; Kia 2025; Amiri *et al.* 2026; Nonhoff*et al.* 2026). However, they do not directly study the stability of the allocation trajectory under cumulative workload constraints with a Volterra integral effect. This study fills that gap. Examines how changes in workload and computational capacity affect the best allocation.

Recent work on Lipschitz and Hölder stability in convex and parametric optimization (Cánovas& Parra 2025; Klatte, 2026; Uderzo, 2026) provides a foundation for sensitivity analysis. However, this problem requires an approach because the feasible region is always constrained and depends on allocation history through the Volterra operator. Under the Slater condition, the interpolation control feasible set changes linearly, while strong convexity turns this estimate into a square-root stability bound. The analysis also shows the existence and uniqueness of the allocation and derives a functional Karush–Kuhn–Tucker condition with a measure-valued multiplier giving a water-filling interpretation of the best resource distribution.

2. Related Work

The literature related to this study can be generally divided into four research areas. These areas are constrained convex optimization, cumulative resource constraints, time-varying optimization and parametric sensitivity analysis. Each of these areas gives a foundation for understanding dynamic resource-allocation problems. The present study sits at the intersection of these four areas. In this work, a continuous-time resource-allocation problem is examined, where the cumulative workload requirements and the instantaneous computational capacity limits can change as the environment changes. This formulation lets the resource-allocation problem be examined in a way that fits systems whose operating conditions change continuously.

2.1 Dynamic and constrained online optimization

Online convex optimization has become an approach for resource-allocation challenges, where decisions must be updated whenever system information changes. Early studies looked at constrained resource allocation in data-center environments, where virtual queues were used to manage long-term resource needs (Cao *et al.*, 2018). Later Ding *et al.* (2021) created an online convex optimization framework that includes long-term constraints, and they demonstrated its use in cloud-computing resource allocation. Their proposed framework mainly focuses on using benchmarks and on meeting long-term constraints over time rather than on comparing the optimization outcome with one fixed static solution. This approach is useful for systems because workload and available resources can change during the operation period.

More recent studies have looked at types of time-varying constraint structures. Liu *et al.* (2025) examined augmented methods for time-varying constrained online convex optimization and explored how these methods can be used in network resource-allocation problems. Abdullah al. (2026) added memory and prediction information into constrained convex optimization showing that using past data and forecasts helps make better decisions in changing environments. Ji and Mao (2026) studied distributed convex optimization when feedback is delayed. Their work points out the challenges in computation and communication that arise when optimization decisions must keep up with a shifting environment. These studies show that taking into account time-varying past information and delayed system responses is crucial for making resource-allocation decisions in dynamic computational systems.

Although these studies give foundations for optimization algorithms, they usually focus on regret, constraint violation, solution tracking and computational efficiency. However, these studies do not primarily focus on how stable the optimal resource allocations are when the computational environment changes. The present paper instead focuses on the stability of the optimal allocation itself under changes in the computational capacity and workload profiles. In particular, the study examines how perturbations in capacity and workload conditions can affect the optimal allocation.

2.2 Cumulative constraints and computational resource allocation

Constraint they are important when service needs depend on total work finished in one period. Unlike pointwise constraints, they depend on the history of resource allocation, so only the current allocation cannot decide if workload is already satisfied.Constraints they come naturally in computer tasks, network systems and services where accumulated work decides if demand is met. Resource allocation research is increasingly including time-varying costs and constraints. Cheng *et al.* (2026) develop an online convex optimization method for resource allocation in an edge-computing network. Studies use optimization and learning approaches in communication and distributed systems (Luan *et al.*, 2026; Ma & Sun, 2026; Yu, 2026). These studies show the importance of adapting allocation when system conditions change.Existing research mostly focuses on optimisation methods and performance, not the sensitivity of the best allocation to changes in workload and computational capacity. This study instead looks at stability analysis make math relation between changes in conditions and the best resource allocation.

2.3 Time-varying convex optimization

Time-varying convex optimization it gives a theoretical framework for problems where the objective and the constraints change over time. Davydov and colleagues (2025) study time-varying convex optimization by using contraction and equilibrium-tracking techniques. Rostami and Kia (2025) focused on tracking the solution when the problem evolves.Amiri and co-workers (2026) make an embedding approach for real-time solution of time-varying constrained convex optimisation problems. Nonhoff and others (2026) study constrained convex optimization for systems. These studies provide tools for managing changing objectives, constraints, and system conditions.Most studies focus on solution tracking and computation methods, not stability of continuous-time resource allocation with cumulative constraint. This study it looks how best allocation changes when computational capacity and workload are perturbed. It gives a method to measure allocation sensitivity.

2.4 Parametric sensitivity and solution stability

The distinction between Lipschitz and Hölder stability is very important in this study. Li and Li (2014), looked at Hölder continuity of the perturbed solution set in convex optimization. Mordukhovich and Nghia (2014) discussed general Lipschitzian and Hölderian stability results for optimization and optimal control problems.Recent studies extend this area. Cánovas and Parra (2025) discussed Lipschitz stability. Klatte (2026), he study parametric optimization with feasible-set mapping. Uderzo (2026), he look convex optimization with a set-valued inclusion constraint. Nghi and Tam (2026) and Peng *et al.* (2026), they also show stability and well-posedness under perturbation.But this problem is different now because the feasible set involvesa cumulative inequality constraint defined by a Volterra operator. Changes in workload arrival and computational capacity, they can change both cumulative and instant feasibility across the planning horizon.

3. Model and Assumptions

Consider a planning horizon $[0, T]$and n computational workloads. Let $u_i(t)$denote the computational resource allocated to workload $i$at time t. The allocation vector belongs to the space $L^2(0, T, \mathbb{R}^n)$.

The resource cost is shown by a quadratic function. This function has a workload coefficient that can change over time. It reflects the cost of the computer to serve each workload. The strong convexity of the function ensures that there is a unique best way to allocate resources.Each workload is described by an arrival-rate profile and a positive service-efficiency profile. The cumulative demand for a workload is obtained by integrating the arrival rate over time. The cumulative service depends on the history of resource allocation. This makes it a Volterra-type workload constraint.Computer resources have time-varying capacity limits. Resource allocation must always be nonnegative. The feasible set captures both the capacity restrictions and the cumulative workload requirements at the same time.The model assumes the cost coefficient is always positive and the service-efficiency coefficient always stays away from zero. The workload arrival rate and computational capacity follow bounded conditions, with computational capacity always positive.The model finally assumes a Slater-type condition. This condition ensures that the allocation strictly satisfies both the capacity constraint and the cumulative workload constraint.These assumptions ensure the existence, uniqueness and stability of the allocation.

The uniform feasibility margin is very important for sensitivity analysis because it ensures that a small change in workload or computational-capacity profile does not immediately make a problem infeasible. The uniform feasibility margin also gives a number that is used to build a feasible correction between optimization problems in different computer environments.

Perturbation between two environments is measured by the difference in computational-capacity and workload-arrival profile. Under this assumption, the model is an optimization problem in an infinite-dimensional function space. Bounded resource allocation, continuity of the service operator and strong convexity of the objective together make sure that there is existence, uniqueness and stability of the optimal resource allocation.

4. Well-Posedness and KKT Characterization

The first step is to establish the well-posedness of the continuous-time resource-allocation problem. The cumulative service mapping is a Volterra operator from $L^2(0,T)$ into $C([0,T])$. Therefore, bounded sequences of resource allocations can generate cumulative-service trajectories that're uniformly bounded and smoothly varying. This property is useful because it allows the cumulative workload constraints to be maintained when appropriate weak convergence arguments are applied. As a result, the mathematical properties of the service mapping provide an important basis for proving the existence of an optimal resource allocation. The mathematical well-posedness of the continuous-time resource-allocation problem hinges on the service mapping.

The feasible set is convex. The convex feasible set arises because the computational capacity constraints and the cumulative workload constraints are linear with respect to the allocation variables. The feasible set also incorporates the nonnegativity condition. With the capacity constraint the nonnegativity condition gives a uniform upper bound for every feasible resource allocation. Bounded subsets ofL^2, are weakly compact.

The uniform positivity of the cost coefficients makes the objective function strongly convex. Because the objective function is strongly convex, the optimal resource allocation is unique. The uniqueness of the resource allocation is important because it allows the optimal solution to be seen as a well-defined mapping from the computational environment to its corresponding optimal resource allocation. Uniform positivity of the cost coefficients remains a property keeping the objective function strongly convex. The computational environment is the starting point for the mapping to the resource allocation.

4.1 KKT Characterization

Under the uniform Slater condition, the problem satisfies an appropriate constraint qualification. Let $\mu(t) \geq 0$denote the multiplier associated with the instantaneous capacity constraint, let $\nu_i$be a nonnegative measure associated with the cumulative workload constraint of workload i, and let $\eta_i(t) \geq 0$denote the multiplier associated with nonnegativity.

The Lagrangian can be written as

$$\mathscr{L}(u,\mu,\nu,\eta) = J(u) + \int_0^T \mu(t)\left(\sum_{i=1}^n u_i(t) - C(t)\right)dt + \sum_{i=1}^n \int_{[0,T]} [D_i(t) - (B_i u)(t)]\,\nu_i(dt) - \sum_{i=1}^n \int_0^T \eta_i(t) u_i(t) dt.$$

Define the cumulative workload shadow price by the tail of the multiplier measure,

$$\Lambda_i(t) := \nu_i([t,T]).$$

Since $\nu_i$is nonnegative, $\Lambda_i$is nonincreasing. Using the adjoint representation of the Volterra operator, the stationarity condition becomes

$$a_i(t)u_i^*(t) + \mu^*(t) - b_i(t)\Lambda_i^*(t) - \eta_i^*(t) = 0 \text{ a.e.}$$

The associated complementary-slackness conditions are

$$\mu^*(t)\left(\sum_{i=1}^n u_i^*(t) - C(t)\right) = 0,$$

$$\int_{[0,T]} [D_i(t) - (B_i u^*)(t)]\,\nu_i^*(dt) = 0,$$

and

$$\eta_i^*(t)u_i^*(t) = 0 \text{ a.e.}$$

Together with primal and dual feasibility, these conditions characterize the unique optimal allocation.

In particular, stationarity and nonnegativity yield the generalized water-filling representation

$$u_i^*(t) = \left[\frac{b_i(t)\Lambda_i^*(t) - \mu^*(t)}{a_i(t)}\right]_+.$$

This expression shows that resource allocation depends on the balance between the cumulative workload shadow price and the instantaneous capacity scarcity price. When total capacity is binding, the common multiplier $\mu^*(t)$adjusts the allocations across workloads.

The KKT representation is mainly used in this study to describe the structure of the allocation. It is not considered the new contribution of this paper.The main theoretical contribution is developed in the section, where an explicit stability estimate is obtained for the unique optimal allocation when the workload and computational capacity environment is changed. This result shows how the optimal allocation is affected by perturbations in the workload and capacity profiles. Therefore, the stability analysis provides the mathematical result of the present study.

5. Main Sensitivity Theorem

This section develops the main theoretical result of the paper, which is an explicit stability bound for the optimal computational resource allocation when the workload and computational capacity environment is changed. The analysis considers two optimization problems with the same regularity conditions for the objective function and feasibility. However, the two problems have different computational capacity profiles and workload-arrival profiles. The purpose is to study how these changes in the computational environment affect the corresponding optimal resource allocations. This analysis provides a quantitative relationship between the perturbation of the problem data and the change in the optimal solution.

Let the two environments be denoted by $\theta_1 = (C_1, d_1)$and $\theta_2 = (C_2, d_2)$, and let $u_1^*$and u_2^*be their unique optimal allocations. The environmental perturbation is measured by

$$\Delta = \max\{\|, C_1 - C_2 \|_\infty \ \sum_{i=1}^n \| d_{1i} - d_{2i} \|_\infty\}.$$

The two environments are assumed to satisfy the same uniform bounds on the cost, service efficiency, and computational capacity, together with a common Slater margin $\sigma > 0$. The purpose of the analysis is to determine how the distance between $u_1^*$and u_2^*depends on Δ.

The main problem is that the two best answers usually come from two groups of possible answers. So one best answer can't be used directly in the rules that show the other problem is solved. To fix this problem, the proof first creates enough answers between the two different computer setups. Once these close enough answers are found, the related math problems are added to the fact that the thing being optimized very curved. This process helps create a link between the two best answers. Because of this, the way resources are best used can be limited based on how work and computer power change.

5.1 Uniform Bound on Feasible Allocations

The instantaneous capacity and nonnegativity constraints imply that every feasible allocation is bounded componentwise by the available capacity. Consequently, all feasible allocations belong to a common bounded subset of $L^2(0, T, \mathbb{R}^n)$.

Define

$$M := \sqrt{nT}\, C_+.$$

Then every feasible allocation u, and in particular each optimal allocation, satisfies

$$\| u \|_2 \leq M.$$

This limit is important because it can manage how big the change in the goal is, without depending on the computer setup that is being used.

For the quadratic objective, the gradient operator is given componentwise by multiplication by the cost coefficient. Since the coefficients are bounded above by a_+,

$$\| Au \|_2 \leq a_+ M$$

for every feasible allocation.

5.2 Feasible Perturbation Between Environments

The principal technical step is to transfer an optimal solution from one environment into the feasible region of the other environment.

Let $\hat{u}_2$be a Slater point for environment θ_2. Define the interpolation parameter

$$\alpha = \frac{\Delta}{\sigma + \Delta}.$$

Starting from u_1^*, construct

$$v_2 = (1-\alpha)u_1^* + \alpha\hat{u}_2.$$

Here W $v_2$is feasible for the second environment.

For the work load limits the difference between the total needs in the two places can be limited by the difference in their arrival patterns multiplied by the time that has passed. The Slater point provides a safety margin that increases with time. More the chosen interpolation parameter ensures that the safety margin is bigger, than any changes. So the ability to work properly stays when moving between the two places.

Thus,

$$v_2 \in K_2.$$

By the same argument, using a Slater point $\hat{u}_1$for the first environment, construct $v_1 \in K_1$from u_2^*.

The distance generated by these corrections satisfies

$$\| v_2 - u_1^* \|_2 \leq \frac{2M}{\sigma}\Delta, \quad \| v_1 - u_2^* \|_2 \leq \frac{2M}{\sigma}\Delta.$$

Hence, the feasible sets of the two perturbed problems can be connected through corrections that are first-order in the environmental perturbation.

5.3 Variational Inequalities

Because u_j^*minimizes a differentiable convex objective over the convex set K_j, it satisfies the variational inequality

$$\langle A, u_j^* v - u_j^* \rangle \geq 0, \ \forall v \in K_j.$$

For the first problem, choose the feasible point $v_1 \in K_1$.obtainid

$$\langle A, u_1^* v_1 - u_1^* \rangle \geq 0.$$

For the second problem, choose $v_2 \in K_2$, giving

$$\langle A, u_2^* v_2 - u_2^* \rangle \geq 0.$$

Rearranging these two inequalities and combining them gives

$$\langle A, u_1^* - Au_2^* u_1^* - u_2^* \rangle \leq \langle A, u_1^* v_1 - u_2^* \rangle + \langle A, u_2^* v_2 - u_1^* \rangle.$$

This relation is the key bridge between the optimality conditions of the two perturbed problems.

Applying the Cauchy–Schwarz inequality and the bounds established above gives

$$\begin{aligned}\langle A, u_1^* - Au_2^* u_1^* - u_2^* \rangle \leq & \| Au_1^* \|_2 \| v_1 - u_2^* \|_2 \\ & + \| Au_2^* \|_2 \| v_2 - u_1^* \|_2.\end{aligned}$$

Since the optimal allocations are feasible,

$$\| Au_j^* \|_2 \leq a_+ M.$$

Therefore,

$$\langle A, u_1^* - Au_2^* u_1^* - u_2^* \rangle \leq \frac{4a_+ M^2}{\sigma} \Delta.$$

5.4 Strong Monotonicity

The quadratic objective provides the complementary lower bound. Since every cost coefficient is bounded below by $a_- > 0$, the gradient operator is strongly monotone:

$$\langle A, u_1^* - Au_2^* u_1^* - u_2^* \rangle \geq a_- \parallel u_1^* - u_2^* \parallel_2^2.$$

Combining the upper and lower bounds yields

$$a_- \parallel u_1^* - u_2^* \parallel_2^2 \leq \frac{4a_+ M^2}{\sigma} \Delta.$$

This inequality is the fundamental estimate underlying the sensitivity result.

5.5 Main Sensitivity Theorem

Theorem1ExplicitHölderSensitivity.
Suppose that two computational environments satisfy the assumptions of Section 3 with common constants a_-, a_+, C_+, and $\sigma > 0$. Let $u_1^*$and u_2^*be their unique optimal allocations, and let Δdenote the environmental perturbation defined above. Then

$$\boxed{\parallel u_1^* - u_2^* \parallel_{L^2} \leq 2C_+ \sqrt{\frac{nTa_+}{a_-\sigma}}\ \Delta^{\frac{1}{2}}.}$$

Proof. The feasible-set correction estimates and the variational inequalities.

$$a_- \parallel u_1^* - u_2^* \parallel_2^2 \leq \frac{4a_+ M^2}{\sigma} \Delta.$$

Taking square roots and substituting $M = \sqrt{nT}C_+$gives the stated result.

The theorem establishes $\frac{1}{2}$-Hölder continuity of the optimal allocation with respect to the computational environment. In particular, as the capacity and workload profiles converge, their corresponding optimal allocations also converge in the L^2sense.

5.6 Interpretation of the Stability Constant

The stability constant is

$$K_H = 2C_+ \sqrt{\frac{nTa_+}{a_-\sigma}}.$$

Each component has a meaning. A larger number of workloads can increase how the allocation can change, while a longer planning period can increase the size of the allocation space. A larger capacity bound can allow feasible allocation paths and that can increase the magnitude of the perturbation.

The ratio $\frac{a_+}{a_-}$measures the conditioning of the quadratic objective. Poorer conditioning leads to a larger stability constant and therefore weaker robustness.

The Slater margin is very important in the analysis. When the feasibility margin grows larger, the stability constant. It rises according to the inverse square root of the feasibility margin. Therefore, when the operating environment approaches the feasibility boundary, the optimal allocation becomes more sensitive to perturbations.This provides a useful interpretation of the theorem: **feasibility margin acts as a quantitative robustness parameter**.

5.7 Why the Stability Rate Is Hölder Rather Than Lipschitz

The result shows a square-root dependence on the perturbation instead of a linear dependence. The square-root dependence deserves emphasis because the square-root dependence comes directly from the structure of the proof. The linear dependence is absent. That absence is a key point.

The Slater interpolation can lead to a correction, and the amount of this correction depends on how big the disturbance is. Strong convexity, however, can control the square of the distance between the two solutions. The assumptions in this paper are enough to ensure Hölder continuity. They do not by themselves lead to a global Lipschitz solution map. In general, getting a Lipschitz result usually requires regularity conditions on the parametrized feasible system and the optimality conditions. The theorem presented here intentionally avoids adding these assumptions. Instead, it provides a clear stability guarantee that comes directly from the model being used.

5.8 Stability of the Optimal Value

The allocation stability result also shows that the optimal objective value is continuous. Because the gradient of the objective is uniformly bounded within the region, the difference between the objective values of the two optimal allocations can be limited by how far apart the allocations are L^2distance.

The optimal value also has a square-root sensitivity bound with respect to the perturbation. So when the workload and capacity profiles change a bit, the optimal allocation trajectory also shifts a little. At the same time, the minimum computational cost also changes in a controlled way.

This provides a second robustness measure for the proposed model and can be used to compare alternative computational operating scenarios.

6. Temporal Stability of the Optimal Allocation

The results about sensitivity that were explained in Section 5 can also be easily expanded to a series of computing environments that change as time goes by. In real-world computing systems, the amount of work coming in and the processing power available usually do not stay the same for the entire time the system is running. Instead, the problem trying to solve might change at times when new work data arrives or when the computing resources are adjusted.

Let $\theta_k = (C_k, d_k)$denote the computational environment associated with the k-th decision period, and let u_k^*be the corresponding unique optimal allocation. The sequence $\{u_k^*\}$therefore represents the evolution of the optimal resource-allocation policy under changing computational conditions.

6.1 Consecutive-Time Stability

Applying Theorem 1 to two consecutive environments gives the following immediate consequence.

Corollary 1 (Temporal Hölder Stability):

Suppose that every environment in a sequence satisfies the common boundedness assumptions and possesses a common Slater margin $\sigma > 0$. Then the corresponding optimal allocations satisfy

Suppose that every environment in a sequence satisfies the boundedness assumptions and has a common Slater margin $\sigma > 0$. Then the optimal allocations corresponding to these environments satisfy

$$\boxed{\| u_{k+1}^* - u_k^* \|_{L^2} \leq K_H \| \theta_{k+1} - \theta_k \|_{\mathcal{P}}^{\frac{1}{2}},}$$

where

$$K_H = 2C_+ \sqrt{\frac{nTa_+}{a_-\sigma}}.$$

So as the computing power and the amount of work coming in change slowly over time, the best way to assign resources will also shift slowly in the L^2sense.

The result is very useful in adaptive resource-allocation settings. The result shows that when consecutive estimates of workload demand and available capacity are close, the resulting solutions also stay similar. The result gives the optimization trajectory a measurable pattern of temporal regularity.

6.2 Accumulated Variation

The consecutive-time estimate can also be used to bound the total variation of the sequence of optimal allocations. For Nenvironments,

$$\sum_{k=1}^{N-1} \| u_{k+1}^{*} - u_{k}^{*} \|_{L^2} \leq K_H \sum_{k=1}^{N-1} \| \theta_{k+1} - \theta_k \|_{\mathcal{P}}^{\frac{1}{2}}.$$

This link clearly connects the changes in the computer setup with the changes in the way to allocate resources. It proves that the time-related instability in the solution can be managed by the changes in the computer environment.

The square-root exponent plays a role in this result. When the environment changes only a little at each step, the total variation in allocation can grow quite large over a long period. This means that temporal stability should not be seen as the entire sequence staying at a single fixed allocation. Instead, it should be viewed as robustness between consecutive optimal solutions.

6.3 Implications for Dynamic Resource Management

The temporal stability result has practical implications. First, the temporal stability result can support warm-start strategies for repeated optimization. The allocation from the environment can give a controlled approximation for the next optimal solution. Second, the temporal stability result gives a way to measure how frequently the optimization system may need to recompute the allocation when the workload and capacity profiles change over time.

The result also shows how important it is to keep a feasibility margin. If the system works close to its maximum limit, the stability constant will get bigger. In this situation, even small changes in the computing environment can lead to differences in the best way to allocate resources.

Therefore, watching the feasibility margin can serve as another sign of robustness. Systems that have operational slack are theoretically less affected by changes over time in workload and resource availability. This means that such systems can keep performance even as demands shift.

6.4 Relation to the Main Sensitivity Result

The temporal result doesn't require an argument for optimization. It comes directly from the problem- sensitivity theorem by using it over and over on the environments that come one after another. So this leads to a way of looking at the framework: parametric sensitivity and temporal stability are two types of the same stability property of the optimal solution map.

Theorem 1 gives a robustness guarantee whenever by compare two different computational scenarios. Theorem 1 also gives a robustness guarantee to track the best allocation through a series of changing computational environments.

These results give the basis for the numerical experiments that come next. In these experiments, the results show how workload perturbations, capacity changes and feasibility margins affect the system through results.

7. Numerical Experiments

The selection illustrates the sensitivity and temporal stability results by using a discretized version to study the continuous-time allotted model.The horizon is set to $T = 10$, with $n = 5$ workloads and a uniform discretization of the time interval. The cost factors and service levels are always kept positive and within limits. At the same time, the amount of work coming in and the computing power available can change over time.

The experiments consider workload perturbations, capacity perturbations, and changes in the feasibility margin. For each scenario, the optimal allocation is compared with that of a baseline environment using the discretized L^2distance.

7.1 Workload-Perturbation Results

The baseline problem is subjected to progressively larger workload-arrival perturbations. The representative results are shown in Table 1.

The allocation distance increases with the perturbation magnitude, while the normalized quantity $\| u_2^* - \frac{u_1^*\|_{L^2}}{\sqrt{\Delta}}$remains approximately constant. This behaviour is consistent with the $\frac{1}{2}$-Hölder dependence established in Theorem 1.

7.2 Capacity-Perturbation Results

Capacity perturbations produce a similar response. Representative results are reported in Table 1.

Table 1. Representative sensitivity to capacity perturbations.

Capacity perturbation Δ	Allocation distance	Normalized sensitivity
0.005	0.029	0.410
0.010	0.041	0.410
0.020	0.058	0.410
0.040	0.082	0.410
0.080	0.116	0.410
0.120	0.142	0.410

The results show that when the computer power is lower, the changes in how things are assigned are bigger during times when the system is working a lot. The steady normalized sensitivity also backs up the expected square-root scaling result.

7.3 Effect of the Feasibility Margin

Table 2. shows the influence of the feasibility margin

Table 2. Representative allocation sensitivity under different feasibility margins.

Relative feasibility margin	Average slack	Allocation distance
0.100	0.210	0.119
0.200	0.390	0.091
0.400	0.760	0.067
0.700	1.280	0.052
1.000	1.790	0.044

The allocation response decreases as the feasibility margin increases. This numerical pattern agrees with the theoretical dependence of the stability constant on $\sigma^{-\frac{1}{2}}$. Hence, greater operational slack provides increased robustness against workload and capacity uncertainty.

8. Discussion and Summary

The analytical and numerical results establish a quantitative connection between changes in the computational environment and changes in the optimal resource-allocation trajectory. The main theoretical result shows that the optimizer is $\frac{1}{2}$-Hölder stable under perturbations of workload arrivals and computational capacity, provided that the objective remains strongly convex and a common feasibility margin is maintained.A key point of this result is that the stability constant directly relies on the size of the problem and on the conditioning parameters. The stability constant is especially sensitive to longer planning periods and higher workloads because these factors can lead to more changes in how resources are allocated. The stability constant also drops in robustness if the quadratic objective is poorly conditioned, which means the solution may not hold up well. Finally, a small Slater margin makes the system more sensitive. Shows that the stability constant is operating close to the edge of feasibility.

The result about stability extends this meaning to a computer environment that is always changing. When the workload and capacity profile change slowly, the best allocation also moves in a controlled way. This property can help perform optimization again, as well as the warm-start procedure, in a real resource system. Because the best allocation it gives a controlled reference for the next optimization problem.The numerical illustration matches the theoretical conclusion. Workload and capacity. That makes the allocation change more and more. When the feasibility margin gets bigger, the sensitivityobservation becomes smaller. So the result supports the idea that the feasibility margin acts like an indicator of robustness.

Limitations still need to be noted. The stability theorem gives a worst-case Hölder guarantee, not a description of real solution change. The numerical study also depends on discretization and some chosen parameter setting. Also, this analysis assumes the uniform bound and the same Slater margin for all environments compared. If the feasibility margin becomes much worse or even gone, then the stability guarantee may not work anymore.Future work can explore sensitivity conditions, which might lead to Lipschitz continuity. Consider workload arrival, uncertainty-aware capacity planning and adaptive algorithms that leverage the temporal stability of the optimal solution map.The proposed framework provides an approach that makes it easier to analyze robustness in continuous-time computer resource allocation, with cumulative workload requirements.

References

1. Abdullah, M., Iosifidis, G., Elayoubi, S. E., & Chahed, T. (2026). Constrained online convex optimization with memory and predictions. *Proceedings of the AAAI Conference on Artificial Intelligence, 40*(24), 19524–19532. https://doi.org/10.1609/aaai.v40i24.39031

2. Amiri, M., Kolmanovsky, I., & Hosseinzadeh, M. (2026). A dynamic embedding method for the real-time solution of time-varying constrained convex optimization problems. *Systems & Control Letters, 209*, 106352. https://doi.org/10.1016/j.sysconle.2026.106352

3. Cánovas, M. J., & Parra, J. (2025). Lipschitz stability in convex optimization. *Set-Valued and Variational Analysis, 33*, Article 25. https://doi.org/10.1007/s11228-025-00760-8

4. Cao, X., Zhang, J., & Poor, H. V. (2018). A virtual-queue-based algorithm for constrained online convex optimization with applications to data center resource allocation. *IEEE Journal of Selected Topics in Signal Processing, 12*(4), 703–716. https://doi.org/10.1109/JSTSP.2018.2827302

5. Cheng, Y., Liang, C., Chen, Q., & Yu, F. R. (2026). Online convex optimization-based resource allocation in dynamic edge computing-enabled networks with time-varying cost and constraints. *IEEE Transactions on Network Science and Engineering*. https://doi.org/10.1109/TNSE.2026.3673670

6. Davydov, A., Centorrino, V., Gokhale, A., Russo, G., & Bullo, F. (2025). Time-varying convex optimization: A contraction and equilibrium tracking approach. *IEEE Transactions on Automatic Control, 70*(11), 7446–7460. https://doi.org/10.1109/TAC.2025.3576043

7. Ding, X., Chen, L., Zhou, P., Xu, Z., Wen, S., Lui, J. C. S., & Jin, H. (2021). Dynamic online convex optimization with long-term constraints via virtual queue. *Information Sciences, 577*, 140–161. https://doi.org/10.1016/j.ins.2021.06.072

8. Huang, X., Yu, Y., Huang, L., & Chen, W. (2026). Doppler-aware dynamic resource allocation for V2X communications using Double Deep Q-Network. *Physical Communication, 74*, 102923. https://doi.org/10.1016/j.phycom.2025.102923

9. Ji, L., & Mao, S. (2026). Distributed online convex optimization under delayed feedback with edge-based event-triggered communication. *Journal of the Franklin Institute, 363*(7), 108575. https://doi.org/10.1016/j.jfranklin.2026.108575

10. Klatte, D. (2026). Lipschitz stability for a class of parametric optimization problems with polyhedral feasible set mapping. *Computational Management Science, 23*, Article 4. https://doi.org/10.1007/s10287-025-00547-0

11. Li, L., Gan, Q., Li, R., Kang, Q., & Liu, Y. (2026). Distributed fixed/predefined-time optimization for multi-agent systems: New exponential-function-based algorithms. *ISA Transactions, 169*, 75–87. https://doi.org/10.1016/j.isatra.2026.01.013

12. Li, H., Yan, C., & Qin, S. (2026). Resource allocation over nonlinear uncertain multi-agent systems: An efficient communication perspective. *ISA Transactions, 170*, 52–62. https://doi.org/10.1016/j.isatra.2026.01.037

13. Liu, H.-Y., Xiao, X.-T., & Zhang, L.-W. (2025). Augmented Lagrangian methods for time-varying constrained online convex optimization. *Journal of the Operations Research Society of China, 13*(2), 364–392. https://doi.org/10.1007/s40305-023-00496-y

14. Liu, X., Chen, Q., Wang, Z., & Liu, R. (2026). Structure-dependent regret and constraint violation bounds for online convex optimization with time-varying constraints. *arXiv preprint arXiv:2603.14319*. https://doi.org/10.48550/arXiv.2603.14319

15. Luan, L., Li, X., & Qin, S. (2026). Distributed proximal algorithms for constrained resource allocation over continuous and discrete communication. *Automatica, 193*, 113216. https://doi.org/10.1016/j.automatica.2026.113216

16. Luo, Z., Wu, W., Dai, X., & Han, Q. (2026). Deep reinforcement learning-based computation offloading and resource allocation in user-centered UAV-MEC. *Physical Communication, 75*, 102992.

17. Ma, D., & Sun, Z.-H. (2026). Dynamic resource allocation optimization algorithm for massive MIMO communication based on Bayesian optimization. *Discover Computing, 29*, Article 98.

18. Nghi, T. V., & Tam, N. N. (2026). Characterizing stability of parametric nonconvex polynomial optimization problems under total perturbations. *Journal of Optimization Theory and Applications, 209*, Article 76. https://doi.org/10.1007/s10957-026-03024-z

19. Nonhoff, M., Köhler, J., & Müller, M. A. (2026). Online convex optimization for constrained control of nonlinear systems. *Automatica, 188*, 112916. https://doi.org/10.1016/j.automatica.2026.112916

20. Peng, Z.-Y., Zeng, Y., Chuong, T. D., Yun, S., & Yang, X. (2026). Solution stability and well-posedness for classes of parametric set optimization problems. *Journal of Optimization Theory and Applications, 208*, Article 17.

21. Rostami, M., & Kia, S. S. (2025). Time-varying convex optimization with $O(n)$computational complexity. *Automatica, 178*, 112355. https://doi.org/10.1016/j.automatica.2025.112355

22. Uderzo, A. (2026). Marginal analysis of convex optimization problems with set-valued inclusion constraints. *Communications in Optimization Theory, 2026*, Article 27, 1–21. https://doi.org/10.23952/cot.2026.27

23. Wang, *et al.* (2026). A constrained online convex optimisation approach for setpoint tracking and deviation compensation. *IET Generation, Transmission & Distribution*. https://doi.org/10.1049/gtd2.70250

24. Yu, Q. (2026). Dynamic allocation and optimization strategy of communication network resources driven by reinforcement learning. *Discover Artificial Intelligence, 6*, Article 97.

25. Zhao, Z., Ding, J., Zhang, J.-X., Yang, T., & Shi, Y. (2026). Distributed optimal time-varying resource allocation for unknown high-order nonlinear multi-agent systems. *Automatica, 187*, 112917. https://doi.org/10.1016/j.automatica.2026.112917

26. Cao, X., Zhang, J., & Poor, H. V. (2018). A virtual-queue-based algorithm for constrained online convex optimization with applications to data center resource allocation. *IEEE Journal of Selected Topics in Signal Processing, 12*(4), 703–716. https://doi.org/10.1109/JSTSP.2018.2827302

27. Ding, X., Chen, L., Zhou, P., Xu, Z., Wen, S., Lui, J. C. S., & Jin, H. (2021). Dynamic online convex optimization with long-term constraints via virtual queue. *Information Sciences, 577*, 140–161. https://doi.org/10.1016/j.ins.2021.06.072

28. Li, X. B., & Li, S. J. (2014). Hölder continuity of perturbed solution set for convex optimization problems. *Applied Mathematics and Computation, 232*, 908–918. https://doi.org/10.1016/j.amc.2014.01.095

29. Mordukhovich, B. S., & Nghia, T. T. A. (2014). Full Lipschitzian and Hölderian stability in optimization with applications to mathematical programming and optimal control. *SIAM Journal on Optimization, 24*(1), 95–123.