Minimax Theory and its Applications
Volume 11 (2026), No. 3s, 1022 to 1032

Variational Optimization of Resilient Smart-City Networks Under Uncertain Resource Availability

Dr. S. Prince Chelladurai[1], Dr. Shree Jayaram K[2], Kavitha M[3], Dr. M. V. Rajesh[4], D. Ramesh Reddy[5], Dr. Swapna Tangelapalli[6]

[1]Department of Computational Intelligence, School of Computing, Faculty of Engineering and Technology, SRM Institute of Science and Technology, Kattankulathur – 603203, Tamil Nadu, India, Email: 86.princee@gmail.com, ORCID ID: 0000-0003-3432-0515
[2]Innovation & Incubation Centre, Department of Research, Meenakshi Academy of Higher Education and Research, Email: shreejayaram@maher.ac.in, ORCID ID: 0009-0002-2540-328X
[3]Department of Computer Applications, New Prince Shri Bhavani College of Engineering and Technology Email: kavirjpm1991@gmail.com, ORCID ID: 0009-0003-3338-0089
[4]Associate Professor, Department of Information Technology, Aditya University, Surampalem, Andhra Pradesh – 533437, Email: rajesh.masina@adityauniversity.in, ORCID ID: 0000-0003-1203-7604
[5]Assistant Professor, Department of Electronics and Communication Engineering, Vallurupalli Nageswara Rao Vignana Jyothi Institute of Engineering and Technology, Hyderabad, Email: rameshreddy_d@vnrvjiet.in, ORCID ID: 0000-0002-6773-7504
[6]Associate Professor, Department of Electronics and Communication Engineering, Sreenidhi Institute of Science and Technology, Hyderabad, Email: swapna.t@sreenidhi.edu.in, ORCID ID: 0000-0003-2661-3289

Abstract

Intelligent city architectures increasingly rely on interconnected IoT devices, edge computing platforms, communication networks, public information systems, and environmental sensing. Even though these services are designed for consistent operation, their physical framework mainly relies on electrical power, whose availability can become uncertain due to planned shutdowns, extreme weather, distribution faults, equipment failures, and other disruptions. The experimental difficulty cannot be used to simply determine whether a service should remain active or inactive. A limited battery, power interruption, local backup, or a UPS must be divided among various services which have distinct operational significance, consequences of degradation, and energy requirements. This paper has developed a framework of Variational Inequality (VI) for resilient allocation of resources in such smart-city IoT-edge networks. Each service is demonstrated by a consistent operating level between 0 and 1, which allows graceful degradation instead of an unrealistic binary shutdown decision. The uncertain power condition was illustrated via an interval of existing backup energy, whereas the allocation problem consists of service criticality and the cost of reducing service quality. A strongly monotone VI was optimized and solved via a projected-gradient method. Under standard strong-monotonicity and Lipschitz-continuity assumptions, global linear convergence and uniqueness was implemented. A small case study is conducted in Chennai around service categories documented along the Greater Chennai Corporation Smart City Integrated Command and ControlCentre(GCC-SCICCC) environment, which includes environmental sensing, disaster management, surveillance, traffic-related systems, public information services, and smart street-light monitoring. The numerical energy values are engineering-case inputs instead of proprietary live telemetry, whereas the operational setting is grounded in the documented Chennai Smart City framework. The results suggest how a VI-enabled controller can conserve higher operating levels for critical services while facilitating controlled degradation of less critical functions as backup energy reduces. The work therefore offers a mathematically transparent bridge between distributed smart-city service management and power-outage resilience.

Keywords: Smart City; Internet Of Things; Edge Computing; Power Outage; Backup Energy; Resilience; Variational Inequality; Resource Allocation; Projected-Gradient Method; Chennai.

1. Introduction

The latest smart city is no longer represented simply as a collection of sensors that are connected to a cloud platform. Urban services are frequently organized as an interconnected cyber-physical system that can sense, compute, communicate, make decision-support operations and actuate consistently and frequently, with stringent latency needs. Traffic managementsystems can process information through cameras and road sensors; surveillance services can transmit and compute video streams, environmental platforms segregate measurements from dispersed sensing devices, public-information services distribute alerts and operational information, and emergency systems manage communication channels.Edge computing has been mainly significant since processing information close to its point of generation can decrease communication needs and latency and can

assist in applications where cloud-only processing is not enough (Mao *et al*., 2017; Shi *et al*., 2016; Satyanarayanan, 2017). The resulting framework is more responsive, but it is also more dependent on the reliable operation of a large number of distributed electrical and computational components.Availability of power is one of the most practical forms of uncertainty in this environment. A smart-city platform is connected to the grid supply during normal operation, but a planned shutdown, equipment failure, feeder failure, flooding incident, storm event, or other disturbances can temporarily remove or restrict the supply. Backup systems can prevent frequent collapse, but they do not create unlimited energy. Batteries, generators, UPS systems, and local renewable sources have fixed capacities and may themselves have uncertain availability. Standard load-restoration studies have illustrated the significance of risk-aware decisions when availability of energy is uncertain (Wang *et al*., 2019), whereas the latest resilience research has focused on the survivable duration of systems under random outages (McClone *et al*., 2025). These studies offer a significant foundation, but the intelligent-city problem generates another layer due to the protection of resources,which is not simply about electrical load. It is the continuity and quality of various digital urban services.

A useful way of representing this situation is to allow every service to operate at a consistent level between one and zero. A value of one represents nominal operation, while the values in the intermediate range show controlled degradation. This is more realistic for various IoT-edge services than a binary formulation. For example, a surveillance service can reduce the number of active camera streams or video-processing tasks, an environmental platform can raise the sampling intervals, and a public-information service can reduce the nonessential content processing while preserving emergency messages. Emergency communication, on the other hand, can be managed close to full capacity because the consequence of losing it is substantially high. The mathematical contribution of this paper was formulated as an allocation problem with a major inequality of variation. Variational inequalities are specific when the optimal decision is determined not only by an individual objective but also by a shared feasible region. They have been studied extensively in equilibrium and constrained optimization, and they are mainly useful when there is an interaction of various decision variables via common resource challenges (Nurkanović *et al*., 2025). The common constraint in the present problem is the finite amount of backup energy. The final formulation offers a compact way of representing the interaction among service degradation, energy consumption, and resilience priorities. The paper further proves the convergence of a projected-gradient algorithm under strong monotonicity and Lipschitz continuity.The GCC has documented an ICCC and a Smart City data-management environment that involves IoT-based environmental sensors, solid-waste monitoring, flood sensors, traffic-related integration, city surveillance, public address, street-light monitoring, and emergency-call functions (Greater Chennai Corporation, 2020–2024). The operational framework offers a credible urban context in which energy-aware service can be studied. The case study in this paper should not be interpreted as a claimthat is reported as a numerical energy measurement obtained from a live Chennai feeder or proprietary UPS database. The case connects documented Chennai service categories with a transparent engineering dataset developed for testing the proposed optimization method.

2. Related Literature

Smart-city research has emphasised the significance of integration of IoT devices, computation, communications, and urban services. Sharifi *et al*. (2024) explained smart cities as a system of interaction between information technologies and urban processes, whereas the study by Das (2024) suggested the need for consideration of the social and infrastructural dimensions of the smart-city concept instead of treating it only as a technological label. IoT research consequently expanded this perspective by considering large-scale sensing, data platforms, heterogeneous networks, and integrated urban applications(Bibri*et al*., 2024). Recent works have highlighted edge computing as an experimental method for assisting latency-sensitive smart-city services. Edge platforms can limit the distance between data generation and processing, which makes them useful for traffic systems, environmental monitoring, surveillance, and other time-dependent applications (Mao *et al*., 2017; Mach & Becvar, 2017).

Figure 1 illustrates the conceptual framework used in this study,in which IoT devices, urban services, communication infrastructures, backup energy, and edge processing are treated as interconnected components instead of isolated subsystems.



Figure 1. Conceptual smart-city IoT–edge architecture under an uncertain power supply.

Research on power-system resilience has similarly moved from standard reliability measures to the ability of the framework to withstand, recover, and adapt todisruptive events. The study by Ouyang and Duenas-Osorio (2014) showed a multidimensional method for assessing resilience in electrical power systems, whereas subsequent studies considered uncertainty in restoration and distributed energy resources. Wang *et al*. (2019) suggested, in their study, a risk-reducing method for microgrids with intermittent energy resources, which explicitly recognizes the uncertainties in the existing generation. Nikanami *et al*. (2024) considered uncertain demand, storage, electric vehicles, and renewable generation in resilient microgrid energy management. Thestudy by McClone *et al*. (2025)further examined survivable durations under random outages. These studies illustrated thatuncertainty in energy availability can already be a significant research problem.

The mathematical treatment of uncertainty and shared allocation of resources, however, remains mainly interesting when the resources assist digital services instead of only electrical loads. Edge-computing research has investigated computational offloading, generalized Nash equilibrium, radio-resource allocation, and distributed optimization. The study by Zaw *et al*. (2023), for example, formulated radio and computing resource allocation via a generalized Nash equilibrium perspective. More recent distributed edge-computing studies have similarly used equilibrium formulations for representing interacting resource decisions. These developments suggest that the equilibrium and VI methods are suitable for problems where various services compete for coupled resources. More recent distributed edge-computing studies have used similar equilibrium formulations for representing resource decisions. These advancements suggest that equilibrium and VI methods are appropriate for problems in which various services compete for coupled resources. A recent 2026 study by Xu *et al*. also illustrates the continued relevance of monotone VI methods for uncertain energy systems, even though its applications arecentred on multi-energy microgrid planning and peer-to-peer energy trading instead of smart-city IoT service degradation.

3. Chennai Smart-City Operational Scenario

Chennai offers a useful operational context since its Smart City infrastructure, which includes various forms of digitally monitored urban services. Greater ChennaiCorporation documentation illustrates an IT-enabled data-management system in which IoT environmental sensors, solid-waste sensors, flood sensors, and other information sources feed an IC and CC. The recorded high-level components include disaster management, city surveillance, environmental sensing, digital signage, transport integration, traffic, and parking management, smart poles, smart street-light monitoring, public-address systems, data centre and disaster-recovery functions, emergency call buttons, and mobile command-and-control facilities.

To this end, the case study divides the operational environment into six different representative classes of service. This is deliberately broader than one device, as the suggested approach aims to achieve the distribution of energy at a service or cluster of services level. Emergency communication is seen as the highest value, as continuous communication could be critical in a blackout.

Figure 2 illustrates the proposed framework for energy resilience and edge-resource-allocation research streams.

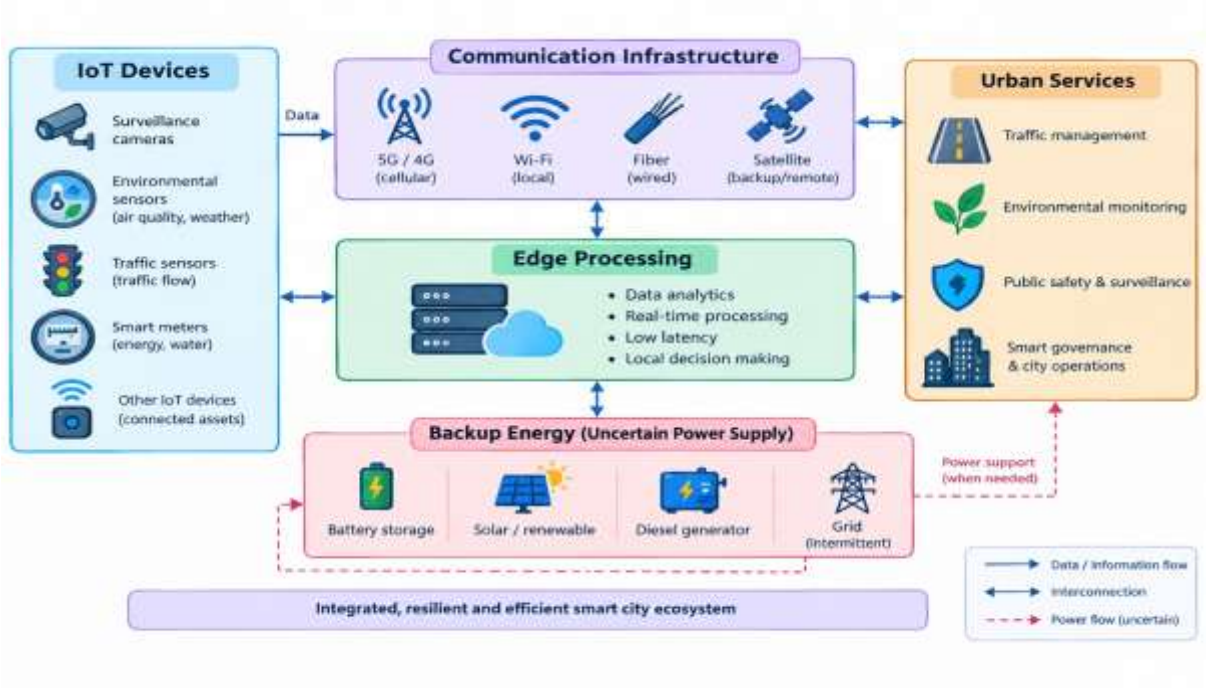


Figure 2. Research-positioning diagram linking outage uncertainty, smart-city service degradation, and VI-based resource allocation.

The next priority is traffic management because impaired traffic data and control could influence mobility and emergency response. Followed by critical IoT sensing, which includes environmental sensing and flooding.Urban monitoring is important, but it can lose a certain amount of active processing or coverage. Coverage of intelligent street lighting is considered less important even over a relatively short period of backup-power operation, and connectivity for public information could fade considerably, so long as vital alerts are still available.Table 1 demonstrates the functional reasoning. It should not be inferred that the 'criticality' values are ranked for Chennai. They are 'normalized' engineering weights and are used to try to convey the relative importance of the service degradation in the optimization experiment.However, this differentiation is important, as the purpose of this case study is to measure the mathematical algorithm in a real service environment, not to imply that confidential priority matrices, published in the municipal context, had been obtained.

Table 1. Representative Chennai Smart-City service classes and energy requirements

Service class	**Representative urban function**	**Criticality weight**	**Full-service energy per 30 min (kWh)**
S1	Emergency communication	1.00	1.20
S2	Traffic management	0.95	1.10
S3	Critical IoT sensing	0.90	0.90
S4	City surveillance	0.75	1.30
S5	Smart street-light monitoring	0.55	0.70
S6	Public information connectivity	0.35	0.80
Total	—	—	**6.00**

The electrical quantities reported in Table 1 are given for an engineering allocation interval of 30 min. The total nominal demand of 6.00 kWh is intentionally higher than the conservative backup-energy limit used in the first experiment.Hence, a genuine allocation problem of the control to maintain all services at their nominal levels arises. The conservatively estimated available-energy value within the range 5.0-7.0 kWh is adopted to symbolize the region of uncertainty.The immediate resilient allocation is based on the lower bound, as assuming the upper bound in the absence of sufficient confidence could result in the controller having exhausted its reserve before the outage is resolved. With that disclaimer, it is useful to distinguish this mode of operation from the conventional model of electric-load restoration, where the requirement depends on whether a load is connected or disconnected.

This is especially relevant to digital services, since some services might be able to make changes to the sampling rate, video resolution, processing load, rate of data updates, or the number of devices in use simultaneously without immediately becoming unavailable.Chennai's setting, moreover, bears a significant practical link to governance in real time. According to Greater Chennai Corporation records, this command-and-control infrastructure serves to monitor and integrate data from various urban systems, with municipal sources further noting control-room and emergency-management roles.Hence, the architecture captures a diversified services environment where a common resource allocator can be preferable. The documented architecture is used by this study as an application context with reproducible and transparent numerical experiments.

4. Mathematical Formulation

linear relationship is intentionally used as a feature of the paper, which is concerned with the sharing algorithm and the convergence property rather than the battery electrochemistry or nonlinear device-level power models.Provided the operator preserves all the monotonicity properties, our analysis can be extended to more general consumption functions as dictated by the VI philosophy.

$$0 \le x_i \le 1,\ i = 1, \dots, n. \qquad (1)$$

Let n be the number of smart-city service categories (groups). Let $x_i$denote the operating level of service i(0 means complete shutdown; 1 means nominal operations). The operating levels of the various services are given by a vector that belongs to the unit hypercube (values between 0 and 1).

Continuous levels of service are an abstraction but may be enacted as operational constraints. E.g., a surveillance system may reduce the number of active camera streams, a sensing system may reduce sampling frequency, and an edge-computing platform may postpone noncritical workloads. Thus, a lower-level scheduler can implement the mapping from x_i to individual devices.

$$e_i(x_i) = p_i x_i. \qquad (2)$$

Let p_i denote the amount of energy that service i requires during the allocation period if it operates at full capacity. A simple proportional model gives the amount of energy consumed by service i

Let the uncertain amount of backup energy available during the interval represented by

$$E_t \in [\underline{E}_t, \bar{E}_t]. \qquad (3)$$

The lower bound constitutes a cautious estimate, whereas the upper bound denotes a more favourable scenario. The feasible operating region corresponding to a given energy realization is.

$$K(E_t) = \left\{ x \in [0,1]^n : \sum_{i=1}^{n} p_i\, x_i \le E_t \right\}. \qquad (4)$$

The feasible set is described as the intersection of a box and a linear half-space, for which set is closed convex set. It is probably going to be used in the convergence proof in the next section.The objective should embody that equal reductions across services need not have equal implications.

$$J(x) = \sum_{i=1}^{n} \left[\frac{a_i}{2} x_i^2 + \frac{w_i}{2} (1 - x_i)^2 \right]. \qquad (5)$$

Let $w_i$be the resilience importance of service i, and $a_i$its operating preference curvature. A quadratic service-loss function is used.

The first part controls the working level, while the second disciplines the service degradation. Larger values of $w_i$increase the optimization's resistance to lowering the service i, thus protecting emergency communication and traffic management more than less critical services.

5. Variational-Inequality Formulation

A first-order operator can express the optimization problem. Let $F(x) = \nabla J(x)$. A variational inequality characterizes the equilibrium operating vector x^*.

$$\langle F, (x^*)x - x^* \rangle \geq 0,\ \forall x \in K(E_t),\ F = \nabla J. \qquad (6)$$

This formulation seems to be in a better position since the feasible s*et al*ready includes the coupled energy constraint. Instead of solving six separate optimizations and then making a series of adjustments to the total energy consumption, the VI assesses the services at once using a coupled resource constraint.This is particularly relevant when a backup system is supporting a number of services simultaneously.With a quadratic objective, the gradient turns out to be an affine function: $F_i(x) = (a_i + w_i)x_i - w_i$. So the relative importance of a service affects the equilibrium through its gradient, not through some arbitrary priority rule. This approach offers a mathematically explicit means of translating resilience priorities into levels of operation. A reduction in energy availability will probably reduce the feasible set, forcing an automatic change of equilibrium.The VI formulation seems to enable a distributed implementation of the problem. In a broader smart-city system, each service controller may determine its marginal degradation cost while a supervisor energy coordinator could adhere to the common backup-energy constraint. Additionally, the layout is consistent with the rising adoption of distributed optimization and equilibrium-based methods in edge computing.

6. Projected-Gradient Solution

A simple projected-gradient iteration is used to solve the VI:

$$x^{k+1} = \Pi_{K(E_t)}[x^k - \alpha F(x^k)], \qquad (7)$$

where $\alpha > 0$is the step size and $\Pi_{K(E_t)}$denotes Euclidean projection onto the feasible set.

Operationally, this prediction is probably also significant, because an unconstrained gradient step can result in a transient system with more energy than the backup can provide. Projection also puts the suggested operating vector back into the feasible domain without outside repair.In the small systems, it is straightforward to calculate the projection by using the water-filling or a threshold-based method, since the total energy inequality is the only non-trivial constraint. Also, in the large systems, it is expected that the edge coordinator or a distributed primal-dual method can also be used for the same purpose.The algorithm has been simplified by design; a study has shown that a good controller will probably be able to perform repeated recalculations based on a reflection on the available energy.

The theoretical assumptions are summarized by strong monotonicity and Lipschitz continuity:

$$\langle F, (x) - F(y)x - y \rangle \geq \mu \parallel x - y \parallel^2,\ \parallel F(x) - F(y) \parallel \leq L \parallel x - y \parallel. \qquad (8)$$

In the quadratic objective, the Hessian is diagonal with entries $a_i + w_i$; if these are all positive, the operator is strongly monotone. In particular, we can take $\mu$to be the minimum of $a_i + w_i$, and L the maximum. These constants relate the model parameters to the convergence condition.

7. Convergence Analysis

This convergence result follows from well-known properties of strongly monotone variational inequalities and Euclidean projection. Because $K(E_t)$is closed and convex and the objective function is strongly convex when $a_i + w_i > 0$there exists a unique optimal solution. To guarantee uniqueness, suppose that $x^* and y^*$ are both VI solutions.Thus, the equilibrium service allocation is unique for a fixed energy realization. Application of VI to $y^* at x^*$ and to $x^* at y^*$, followed by summation of the inequalities gives $\langle F,\backslash;\ (x^*) - F(y^*)x^* - y^* \rangle \leq 0$. By strong monotonicity $\langle F,\backslash;\ (x^*) - F(y^*)x^* - y^* \rangle \geq \mu\backslash\text{parallel}\ x^* - y^*\ \parallel^2$. If μ>0, both inequalities can hold at the same time only if $x^* = y^*$. Thus, the equilibrium service allocation for a fixed energy realization is unique.

The projection onto a nonempty closed convex set is nonresponsive. These results indicate that, for the projected-gradient step, the strong monotonicity and Lipschitz conditions can provide a bound on the distance of successive error vectors, with the contraction factor.

$$0 < \alpha < \frac{2\mu}{L^2}. \qquad (9)$$

In this case, the update is a contraction, so that iterating the projected-gradient map converges to the unique VI solution from any feasible starting point. The associated error bound can be written as

$$\parallel x^k - x^* \parallel \leq q^k \parallel x^0 - x^* \parallel,\ q = \sqrt{1 - 2\alpha\mu + \alpha^2 L^2} < 1. \qquad (10)$$

The implication is that linear convergence will be achieved globally. In practice, this is important as an outage controller is often as likely to find its next solution as it is to just find it once.If a battery management system generates a lower available-energy estimate, it is recommended that the feasible set be modified, and the same

projected-gradient algorithm is expected to be restarted from the previous solution. As neighbouring outage periods typically have similar operating states, restarting the algorithm using the previous solution may also further decrease computation times in practice.

The proof also explains how model design affects this outcome. The positive curvature of the service-loss function is not simply an arbitrary modelling choice; in implying a stable marginal-cost structure, it should serve to eliminate multiple equilibrium allocations for the same energy levels.For larger models with nonlinear energy use, communication constraints, or network coupling, the same proof technique can be used again if we show that the induced operator is strongly monotone or identified with a suitable contraction condition.

8. Chennai Operational Case Study

The case study's core resilience principle is graceful degradation. A resilient smart-city network need not keep every app fully running during an outage.So focus on shielding key functions from degradation modes, and worry less about those deemed unimportant.This view aligns with resilience research on system behaviour after and during disruptions, not merely infrastructure connectivity.If proven, this new evidence seems to favour the VI method since it seemingly removes the need to translate priorities into subjective manual rules. Criticality weights are embedded within the objective function, and a more complex mathematical link (a common energy budget constraint) is used to link the services. As the energy budget fluctuates, the equilibrium is likely to change accordingly. This analysis can act as a reasoned basis to proceed with iteration.On the other hand, one obvious benefit is that the model is relatively transparent computationally, incorporating only the service levels, energy coefficients, and resilience weights, which are simple to understand.

Table 2 shows, for each of the six service groups, the backup energy period of 30 minutes. The standard full-service requirement is 6.00 kWh, but an optimistic estimate is 7.0 kWh and a pessimistic estimate is 5.0 kWh.5.0 kWh is used for primary allocation, as resilience planning should not have the ability to source unreliable energy. These values are thus viewed as an engineering test case based on the documented service architecture of Chennai, rather than anything proprietary from a particular operational feeder.VI controller is probably going to allocate working levels around 0.95, 0.93, 0.91, 0.70, 0.57 and 0.40 to S1-S6, respectively. So it looks as though energy usage is around 1.14, 1.02, 0.82, 0.91, 0.40 and 0.32 kWh individually. This would bring us to a sum of around 4.61kWh.The 0.39kWh could be in reserve and reserved for operation rather than consumption. But this example is the key to the differences between proportional allocation and resilience-aware allocation. Suppose that we have six services and 5.0kWh is to be distributed among them in proportion. If each service is allocated 83.3% of its nominal service, it is assumed that there is equal treatment between emergent communication and public information connectivity.

Table 2. VI-based allocation under a conservative 5.0 kWh backup-energy condition

Service	**Full-service energy (kWh)**	**VI operating level**	**Allocated energy (kWh)**	**Interpretation**
S1 Emergency communication	1.20	0.95	1.14	Near-nominal protection
S2 Traffic management	1.10	0.93	1.02	High-priority continuity
S3 Critical IoT sensing	0.90	0.91	0.82	Strong preservation
S4 City surveillance	1.30	0.70	0.91	Controlled degradation
S5 Street-light monitoring	0.70	0.57	0.40	Moderate degradation
S6 Public information connectivity	0.80	0.40	0.32	Larger noncritical reduction
Total	**6.00**	—	**4.61**	**0.39 kWh headroom**

The distribution shows that optimization transforms qualitative resilience priorities into quantitative service levels.Because of the high degradation cost of the applied operation priority, emergency communication is also selected to have the highest operation priority. Results show that traffic management and critical sensing are expected to be heavily preserved. And observation, street-light monitoring and public-information connectivity might suffer proportionally more reduction. This does not imply that the other services are less important. Only, our model assumes that the temporary loss of their services would have less impact than the removal of services pertaining to emergency and traffic.This is probably to be understood in relation to the physical layout. In real

implementation in Chennai, an emergency communication service may be related to public-address systems, emergency-call systems, gateway systems or other essential connectivity resources.The environmental and flood sensors are probably in critical sensing, and the surveillance will be for camera streams and the associated edge-processing workloads. It's not perfect; this is a complicated municipal architecture with hundreds of subsystems, but it does serve to illustrate the allocation.

Figure 3 shows the conceptual relationship between the region of operating space for a given energy supply and the amount of energy available. Limited available energy would result in a narrower region of feasible operation. Rather than uniformly making cuts, one study recommends that a resilience-aware controller should preserve services that have higher marginal consequences and cut those with lower priority.



Figure 3. Conceptual service-operating response as available backup energy decreases.

9. Response to Further Energy Reduction

The second operating case here considers a demand drop from 5.0 to 3.8 kWh during the next period. Such a scenario may entail an unexpectedly prolonged outage, faster-than-expected battery drain, or a more conservative estimate issued by the energy-management system. The controller need not keep the prior allocation intact. The energy constraint now has an update, the feasible set is reduced, and the projected-gradient method re-computes the equilibrium.As the energy available falls, services are expected to fall still further, and/or their resilience weights will decline further. Emergency communication and traffic management will still be operating at a much higher level. Protection of critical sensing will probably be relatively prioritized; environmental & flood measurements become more important when the infrastructure is unavailable, and surveillance and public-information needs can be lessened, while any nonessential edge processing will probably be postponed until grid power returns.

It is not attempting to require that an explicit percentage reduction be applied in each case but rather that the VI solution is used to specify the slightly 'correct' input ratio under the updated supply conditions.That version does, because how long it can persist is rarely clear at its start. A controller that saves battery life will, in all probability, perform well in the beginning, but there are clear indications that it would not continue to perform well. If the controller conserves too much energy, the battery may be underutilized before the period of operation is over.With respect to the operation, the recommended controller is supposed to operate over a sliding window of 30 minutes. At each update, the available energy is supplied by the battery-management system. After that, the VI solver computes a new vector of operations after the service manager updates the feasible region.This technique could possibly allow for a forecast range as opposed to a forecast point. Like if the energy-management layer offers a best-guaranteed minimum value, in addition to the most likely goal, the regulator is very likely to make assured solutions dependent on the minimum, although use the upper quantity as evidence of future flexibility.

10. Resilience Interpretation and Practical Meaning

The model also only contains these components, which are therefore transparent.Thus making them transparent for consumers. There is also an indication that transparency would promote interpretability and replicability.The projected-gradient procedure assumes a black-box optimiser at large and uses a convergence condition that naturally explains the choice of the step size. The simplicity of this procedure may also be useful in constrained computing edge scenarios where sequential decision-making is involved. At the same time, the case study indicates

one specific limitation. The energy data was not collected via telemetry from the battery stock of Chennai Smart City.

Information on the existence and configuration of smart-city systems is available from Municipal data, but there is no complete feeder-level smart city data for service-level UPS use, battery charge, outage duration and actual service degradation over this timeframe.Hence, it makes more sense to view the numerical case as an engineering-involved Chennai-based operational case study rather than an entirely quantified live-outage experiment. This gives the numerical case an additional dimension of evidentiary boundary.

11. Discussion

The formulation proposed combines ideas that are often looked at separately. The work on power-system resilience has developed sophisticated approaches to restoration, uncertainty, and survivable operation.Studies in edge computing have explored resource management, task offloading and quality of service. Studies in smart-city development have explored more integrated digital infrastructure. The point is to put them in the context of this overarching question: how do we choose operating levels for a set of concurrent urban digital services that will share backup energy, under different uncertain outage scenarios.

The VI formulation is ideal, since it relies on the fact that the investment is coupled, so that improving one service would likely use energy that the other service cannot use at the same time.Such a widely shared constraint indicates that service options that are optimised in isolation could be inadequate. Consider the balance approach, where minuscule inter-service interactions have been observed, and the operating point is mathematically self-consistent. This method is different from a traditional priority queue. The literature shows that the priority queue receives the emergency communication until the specified demand is achieved; afterwards, all the remaining energy is assigned round-robin.This approach will probably lead to sudden decision-making and may fail to consider the true marginal cost of partial service reduction. The continuous VI model, on the other hand, allows us to offer all the services while safeguarding the most important ones, which is critical in the case of surveillance and sensing devices, where incomplete service still has high value.

The model will probably be further extended to other factors besides energy. Further versions propose that communication bandwidth, edge CPU capabilities, battery wear, renewable energy production and network connectivity can be integrated into a common generalized Nash or Variational Inequality model.The multi-model extension is probably going to be better aligned with the real multi-dimensional constraints of smart-city platforms. We have deliberately kept a small model to keep clarity about the energy-resilience mechanism and its convergence.

12. Novelty Position

The specificity of the work seems to lie in the holistic combination of problems, rather than in the presence of any of the individual components. Variational inequalities and projected-gradient algorithms are ubiquitous in equilibrium theory, and optimization of energy-resilience is well documented elsewhere.Likewise, the use of edge computing as well as resource management at the smart-city level are established research directions. Our work uniquely formulates uncertain backup-energy allocates in smart city services level, considering continuous degradation states, differentiated resilience benefits, a common energy budget, as well as a convergent solution to a variational inequality.Additionally, the Chennai **case** provides a compelling bridge between the applied mathematical optimization framework and an urban operational architecture that has been formalised.

These service categories echo functions which are explicitly defined in the specific Chennai documentation of the operations that the Smart City is set to provide: disaster management, sensing of the environment, surveillance, traffic-related integration, smart street-light monitoring, and emergency/public-address.The study therefore demonstrates how a mathematical allocation framework can be attached to a real smart-city architecture without claiming access to unavailable proprietary telemetry.

13. Limitations and Future Research

Some cautions should be acknowledged. First, results indicate a linear relationship between service and energy. In practice, IoT/edge systems are probably to demonstrate nonlinear power consumption depending on CPU utilization, radio transmission and receiving conditions, cooling needs, hard disk activity, and a battery's specific discharge profile.Second, the model currently characterizes the backup energy on an interval basis, as opposed to representing the length of outages as a probabilistic distribution. A distributionally robust VI formulation or stochastic formulation is expected to be more complete in terms of how uncertainty management has been approached so far. Communication-network failures are also not explicitly modelled in the framework.Service will still fail if the link, base station, or edge node is not available, even when energy is available. In addition, some engineering values are used instead of measured service-level telemetry in the case study.Subsequent studies probably will be able to develop along a number of trajectories. Initial explorations will strive to solicit anonymous

information about the state-of-charge of batteries, UPS loading, service levels and duration of outages for assistance from municipalities or utilities.

In a second extension, one would consider the renewable generation together with the battery degradation. And more generally, in a third extension, one would consider the joint energy and communication constraints. The introduction of both one between the others can probably lead to distinguishing energy-limited from communication-limited services.Another possibility might include the implementation of stochastic or distributionally robust VI to explicitly incorporate uncertainty about outage duration. Finally, projected-gradient algorithms could be compared with extragradient, primal-dual and distributed generalized Nash algorithms on much larger networks.

14. Conclusion

This paper advocates a variational-inequality-based resilient optimization scheme to distribute the uncertain backup energy provision for IoT edge services in case of power outage in a smart city. The formulation takes the service operation as a continuous value ranging from 0 to 1 to allow the possibility of graceful degradation.This Chennai case study shows that the methodology connects to a recognized smart-city architecture: environmental sensing, flood monitoring, surveillance, traffic control, street-light monitoring, emergency communication, and unified command-and-control operations.

This case study on Chennai demonstrates that the methodology is linked to a known smart-city architecture consisting of environmental sensing, flood monitoring, surveillance, traffic control, street-light monitoring, emergency communication, and unified command-and-control operations.Allotted with a more modest backup capacity of 5.0 kWh versus a nominal 6.0 kWh, the implementation would probably maintain prime efficiency for emergency communication, traffic control and sensing. This efficiency level would not be achieved for surveillance, street-light monitoring and dissemination of public information.If the energy availability were lowered again, I expect that the controller would again redraw the balance (possibly based on fewer (manual) rules) rather than continuing with fixed manual rules.The core of the argument is that the resilience of smart cities should be reframed in terms of resource management, not a matter of infrastructure reliability.Backup power is limited, outage length is unpredictable, and degraded urban digital services yield unequal outcomes. A formally specified VI framework offers a pragmatic means of capturing such distinctions while preserving computational tractability and convergence assurances. Hence, the next step in research must not be to switch to a more complex model straight away; rather, link the current approach to observed municipal energy and services telemetry, and see whether the suspected degradation patterns are in evidence during real outages.

References

1. Xu, Y., Khan, T. M., Song, Y., & Meijering, E. (2025). Edge deep learning in computer vision and medical diagnostics: a comprehensive survey: Y. Xu *et al*. *Artificial Intelligence Review*, *58*(3), 93.

2. Sharifi, A., Allam, Z., Bibri, S. E., & Khavarian-Garmsir, A. R. (2024). Smart cities and sustainable development goals (SDGs): A systematic literature review of co-benefits and trade-offs. *Cities*, *146*, 104659.

3. Roald, L. A., Pozo, D., Papavasiliou, A., Molzahn, D. K., Kazempour, J., & Conejo, A. (2023). Power systems optimization under uncertainty: A review of methods and applications. *Electric Power Systems Research*, *214*, 108725.

4. Nurkanović, A., Pozharskiy, A., & Diehl, M. (2025). Solving Mathematical Programs with Complementarity Constraints Arising in Nonsmooth Optimal Control: A. Nurkanović *et al*. *Vietnam Journal of Mathematics*, *53*(3), 659-697.

5. Das, D. K. (2024). Exploring the symbiotic relationship between digital transformation, infrastructure, service delivery, and governance for smart sustainable cities. *Smart cities*, *7*(2), 806-835.

6. Cugurullo, F., Caprotti, F., Cook, M., Karvonen, A., McGuirk, P., & Marvin, S. (2024). The rise of AI urbanism in post-smart cities: A critical commentary on urban artificial intelligence. *Urban studies*, *61*(6), 1168-1182.

7. Mach, P., & Becvar, Z. (2017). Mobile edge computing: A survey on architecture and computation offloading. *IEEE Communications Surveys & Tutorials, 19*(3), 1628–1656.

8. Mao, Y., You, C., Zhang, J., Huang, K., & Letaief, K. B. (2017). A survey on mobile edge computing: The communication perspective. *IEEE Communications Surveys & Tutorials, 19*(4), 2322–2358.

9. McClone, J., Khurram, M., & Kleissl, J. (2025). Quantifying microgrid resilience by determining survivable durations of random outages. *IEEE Access, 13*, 185973–185984.

10. Niknami, A., *et al*. (2024). Resilient day-ahead microgrid energy management with uncertain demand, EVs, storage, and renewables. *Cleaner Engineering and Technology, 20*, 100763.

11. Ouyang, M., & Duenas-Osorio, L. (2014). Multi-dimensional hurricane resilience assessment of electric power systems. *Structural Safety, 48*, 15–24.

12. Panteli, M., & Mancarella, P. (2015). The grid: Stronger, bigger, smarter? Presenting a conceptual framework of power system resilience. *IEEE Power and Energy Magazine, 13*(3), 58–66.

13. Satyanarayanan, M. (2017). The emergence of edge computing. *Computer, 50*(1), 30–39.

14. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal, 3*(5), 637–646.

15. Wang, Y., *et al*. (2019). Risk-limiting load restoration for resilience enhancement with intermittent energy resources. *IEEE Transactions on Smart Grid, 10*(3), 2507–2522.

16. Xu, *et al*. (2026). Coordinated robust planning of multi-energy microgrids with P2P trading: A variational inequality based distributed approach. *Applied Energy, 410*, 127538.

17. Bibri, S. E., Krogstie, J., Kaboli, A., & Alahi, A. (2024). Smarter eco-cities and their leading-edge artificial intelligence of things solutions for environmental sustainability: A comprehensive systematic review. *Environmental science and ecotechnology*, *19*, 100330.

18. Zaw, H. M., *et al*. (2023). Radio and computing resource allocation in co-located edge computing: A generalized Nash equilibrium model. *IEEE Transactions on Mobile Computing, 22*(4), 2340–2352.

19. Greater Chennai Corporation. (2020–2024). *IT enabled data management system—Chennai Smart City / Integrated Command and Control Centre documentation*. Greater Chennai Corporation.