Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
April 23, 2026
A Regional Conference of the Society for Industrial and Systems Engineering

A Methdology to Assess Energy Resilience and Economic Tradeoffs at Military Installations

Jake Moffat[2], Wyatt Cyprow[1], Sebastian English[1], Stuart York[1], Trenton Hogan[3], James Grymes[2], David Sang[4], Karoline Hood[2], Scott Katalenich[3], and F. Todd Davidson[1]

[1]Department of Mechanical and Aerospace Engineering, United States Military Academy, West Point, New York 10996

[2]Department of Mathematical Sciences, United States Military Academy, West Point, New York 10996

[3]Department of Civil and Environmental Engineering, United States Military Academy, West Point, New York 10996

[4]United States Space Force

Corresponding author's Email: frederick.davidson@westpoint.edu

Author Note: The authors would like to thank the members of the offices of the Assistant Secretary of the Army for Installations, Energy, and the Environment and the West Point Directorate of Public Works for their time and expertise. The views expressed herein are those of the authors and do not reflect the position of the United States Military Academy, the Department of the Army, or the Department of War.

Abstract: Reliable and resilient energy is essential for military installations, yet the supply remains vulnerable to a wide range of disruptions, from adversarial attacks to natural disasters. This paper presents a methodology that integrates an hourly energy dispatch model with a Genetic Algorithm to generate microgrid architectures capable of sustaining critical infrastructure during grid outages. The framework enables decision makers to evaluate tradeoffs between resilience and cost across alternative system designs. Six energy technologies are considered: solar photovoltaics, wind turbines, diesel generators, pumped storage hydropower, nuclear reactors, and lithium-ion batteries, alongside grid-supplied electricity. The methodology is demonstrated using the US Army Garrison West Point as a case study. Results show that battery backup and diesel generators are essential for near-term resilience, while long-duration outage risks highlight the value of fuel-independent sources such as nuclear, solar, and batteries, supporting installation-level planning for contested environments.

Keywords: Energy Security, Energy Resilience, Microgrid Islanding, Nuclear Reactors, Military Installations

1. Introduction

Energy infrastructure is a critical component of mission readiness. The military of the 21st century relies on power for nearly every mission activity, from running airfields to powering soldier equipment. Military bases serve as essential nodes in the national defense network and must remain functional under all conditions. Title 10 U.S. Code § 2920 governs energy resilience and security measures on military installations (United States Code, 2026). By fiscal year 2030, every installation is required to be able to provide 100% of the energy needed to maintain a 99.9% minimum level of availability necessary to support critical mission requirements. A critical mission is defined as a task critical to the successful performance of the installation in meeting strategic national defense measures such as, but not limited to, communications in headquarters facilities, computer systems within airfield control towers, assembly lines in munitions plants, and radar systems for missile defense. The code also states that installations should promote the use of diverse sources of energy while favoring on site sources such as modular generation and microgrids. Installations must also determine how long they can support critical missions if cut off from the commercial grid and what future plans they can implement to increase this duration (United States Code, 2026). A resilient energy system ensures continuous power to critical facilities, enabling sustained readiness and rapid response capabilities. Without such resilience, disruptions in energy supply could compromise operational effectiveness, jeopardize mission execution, and threaten the safety of personnel and assets, thereby preventing the Army from projecting force worldwide and accomplishing its mission of fighting and winning the Nation's wars. The model, Decision Algorithm for Viable Energy (DAVE) helps determine technologies that should be utilized in installation microgrids. The user inputs specific information about an installation, and DAVE runs a genetic algorithm that outputs a series of potential microgrid compositions. The model provides a quantitative framework to evaluate how certain infrastructure design choices affect performance, cost, and resilience. It allows users to simulate real-world conditions, such as extreme weather or grid failures. Moreover, the model tests different energy configurations to assess performance under long duration outage events. The model helps bridge the gap between theoretical energy optimizations and practical constructability.

2. Literature Review

2.1. Defining Resiliency

Defining resilience is a crucial first step to constructing an effective energy resilience model. Different energy related institutions vary in how they choose to define resilience, but this project largely adopts the definition for energy resilience as presented in Title 10 U.S. Code § 101 as "the ability to avoid, prepare for, minimize, adapt to, and recover from anticipated and unanticipated energy disruptions in order to ensure energy availability and reliability sufficient to provide for mission assurance and readiness, including mission essential operations related to readiness, and to execute or rapidly reestablish mission essential requirements" (United States Code, 2024). In contrast, Hashimoto, Stedinger, and Loucks (1982) define resilience as "how quickly a system is likely to recover or bounce back from failure once the failure occurred". These definitions are distinctly in that Hashimoto sees resiliency as limited to once a system has failed while U.S. Code § 101 defines resiliency as also being an ability to avoid and minimize failures. Jesse *Et al.* reviewed more than 130 papers in an attempt to narrow down a definition for resilience. They found definitions tended to revolve around the key concepts of prepare, adapt, recover, and absorb (Jesse, Kramer, & Koning, 2024). These key concepts certainly apply to military installations. The Army cannot fail in its mission, therefore resiliency should not be defined in terms of failure recovery, but rather it should be defined in terms of preventing critical mission failures. The model herein calculates resiliency based on how well a given microgrid can meet an installation's critical load and for what percentage of the outage event it can sustain the critical load. This means energy resilience is an installation's ability to prepare for and respond to a grid failure in order to ensure it has the necessary power to meet its critical mission tasks.

2.2. Cost Basis

The Energy Information Administration's (EIA) "Capital Cost and Performance Characteristics for Utility-Scale Electric Power Generating technologies" governs cost analysis (LaRose, Diefenderfer, & Namovicz, 2024). The report, published in January 2024, represents costs in 2025. EIA commissioned an external consultant to develop up-to-date cost and performance estimates for utility-scale electric generating plants for the 2025 Annual Energy Outlook (AEO2025). This information allows EIA to compare the costs of different electric generation technologies on a standardized basis (LaRose et al., 2024). The EIA presents the total cost a developer would expect to incur during the construction of a project, excluding financing costs. These costs are split into three groups: capital, fixed operations and maintenance, and variable operations and maintenance. A parametric evaluation of costs from actual and planned projects with similar scopes and configurations to the generation technology estimated the capital cost. Dollar values given in the report represent a baseline cost for a set sized plant. This cost was turned into a cost per MW-hr and linearly interpolated to account for changes in size. A labor rate factor was calculated to account for changes in labor costs based on location. A baseline labor cost was found by averaging labor costs in 30 major cities from across the country. The country was split into 25 electricity market model regions to reflect changes in local labor costs. Each region's labor factor was made using published labor rates for each location and comparing them to the 30-city average. Labor rates were taken from RS Means Labor Rates for the Construction Industry 2023 edition. Costs were then added to cover social security, workers' compensation, and federal and state unemployment insurance. From there the regional labor productivity factor from Compass International Global Construction Costs Yearbook, 2022 edition were factored in (LaRose et al., 2024).

3. Methodology

3.1. User Interface

The DAVE model is designed with the recognition that the user may not be experienced with MATLAB, a commercial software programming language. In order to reduce the learning curve for the user and make the model more effective at delivering the best possible recommendations, the model implements a graphical user interface (GUI) through the use of an Excel spreadsheet linked to the MATLAB code. In addition, the technical manual provides step-by-step instructions for how to use the code. These two products create a path that allows any user to understand what data inputs are required and provides the user with a clear location for the input.

3.2. Initialization of the Code

In the DAVE model, the initialization of the code consists of MATLAB pulling user-defined inputs and assumptions from the GUI. It then formats the data into matrices and arrays for use in the code. While the variables seen in the inputs section change depending on location, there are a multitude of quantities that remain constant across all regions and situations. For example, a 2 MW Caterpillar diesel generator is assumed to produce the same power regardless of its geographical location. These values were determined through technical reports and communication with industry experts.

3.3. Model Operation

The genetic algorithm works by first generating a set of solutions. It then evaluates those solutions by discarding the poorest performers and keeping those closest to the Pareto frontier. The model then uses the solutions that are kept as parents, combining them to create a new round of solutions before repeating the process again (MathWorks, n.d.). Latin hypercube sampling (LHS) initializes the algorithm, ensuring balanced exploration of our problem space. This was chosen over a random initialization to ensure the entirety of the data space was explored. LHS splits the feasible range of each of the parameter values and forces the initial population set to cover the entire space evenly, leading to a faster convergence time in the model (Grymes, Newman, Zolan, & Mehta, 2025). Each microgrid configuration within the initial population then undergoes stochastic failure modeling, where each system can experience a failure. The failure probability is assigned from an exponential distribution with each technology having a unique failure rate, Equation 1 (Wales, Zolan, Hamilton, Newman, & Wagner, 2023) (all mathematical notation is defined in Table 1). The duration and severity of failure are also randomly assigned. From there, the model runs a genetic algorithm for different combinations of technologies, testing their performance under the determined failure conditions. Unlike typical commercial sizing tools, DAVE embeds mission-centric dispatch rules that prioritize critical-load survival and explicitly models discrete and non-linear technology choices (including nuclear modules).

$$F(t) = 1 - e^{-\lambda_j t} \tag{1}$$

$$\max \left\{ \sum_j -c_j X_j, \sum_j R_j \right\} \quad \text{s.t.} \quad b_j^l \leq X_j \leq b_j^u \tag{2}$$

Table 1: Model Notation and Definitions.

Symbol	Description	Units
Sets		
$j \in J$	Energy technology $j \in \{$Solar, Diesel, Wind, Battery, Hydro, Nuclear$\}$	[–]
Parameters		
c_j	One-year cost of technology j	[$]
b_j^u	Upper bound for technology j	[MW], [MWh], [#], [m^2]
b_j^l	Lower bound for technology j	[MW], [MWh], [#], [m^2]
λ_j	Failure rate of technology j	[1/year]
t	15-minute time interval index	[–]
Decision Variables		
X_j	Number of units of technology j	[#]
R_j	Resilience contribution of technology j	[%]

The genetic algorithm assigns the configurations a fitness value based on how well they performed in cost and resilience, Equation 2. It computes resilience outcomes, numerically defined as the number of hours of critical load without coverage

ISBN: 97819384962-5-7

measured as a percentage of time critical load is unmet, and then ranks or optimizes portfolios to identify the best cost-resilience tradeoffs. The fitness value drives the algorithm to determine which solutions are kept and which will undergo crossover. The algorithm is tuned through crossover, mutation rate, and elitism. Crossover is how the population mates to create the next generation of microgrids, while mutation is a small change in a component of that configuration that allows exploration of more options. The model uses the scattered crossover technique. Each solution is given a random boolean mask over the 14 decision variables. A value of one indicates taking the first parent's value whereas zero indicates the other. Lastly, elitism is the portion of configurations that are kept in the next generation due to having the highest fitness values. Beyond the initialization technique, a distinguishing feature of this algorithm is how these tuners are implemented. The crossover method used is preferred when the solution requires combinations of variables in different regions of the search space (Brownlee, n.d.). Elitism is set to 0.05, retaining the top five percent, which balances convergence speed and diversification. The mutation probability was altered to be adaptive; in the early generations, the value is larger (starting at 0.05). As the algorithm progresses, both the number of generations and the fitness values will trigger a decrease in mutation rate. This ensures a sufficient exploration phase while also reducing variation during the exploitation phase. To terminate the algorithm, one of two criteria must be met. Either the algorithm reaches the maximum number of generations set by the user, or the solutions only improve by $10,000.00 with a new generation. Once the algorithm terminates, it produces a final set of solutions that represent the cost-resilience tradeoff space for candidate microgrids.

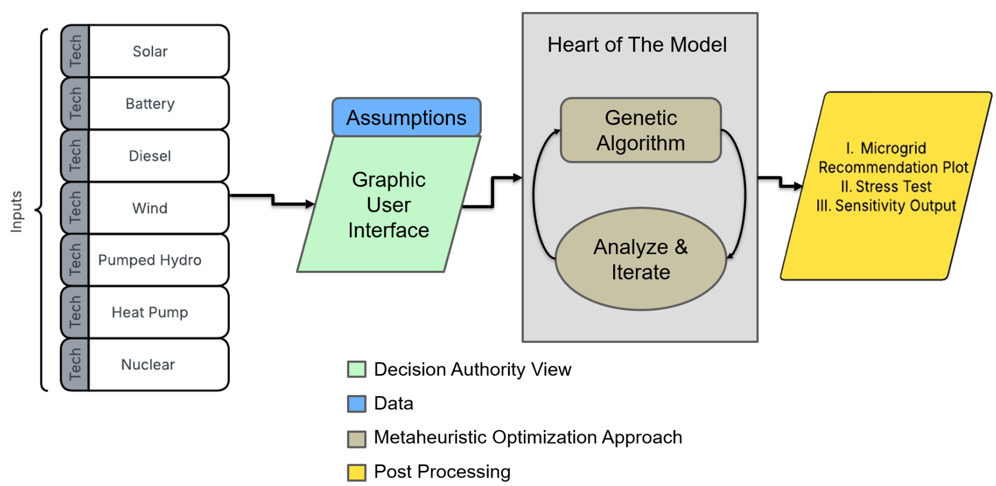

Figure 1: DAVE Model System Workflow Diagram and Result.

3.4. Sensitivity Analysis

The model performs sensitivity analysis through stress-testing and solution robustness analysis. Stress tests involve running the model with a 20% increase in critical load demand and comparing results against baseline configurations. A simulation-based robustness analysis is conducted to evaluate solution performance under uncertain operating conditions. The final solution set is forced back through the failure simulation, and its performance under multiple iterations is stored and returned. The final output consists of minimum, mean, maximum, 10th percentile, and 90th percentile resilience and cost values.

3.5. Nuclear Siting Feasibility

When the DAVE model identifies nuclear generation as feasible, the analysis shifts to siting. Installation constraints such as terrain, access, and constructability determine where development can occur. This section applies a quantitative, multi-criteria framework to create an Excel evaluation tool. The method documents key site attributes, scores each location across nine criteria, and uses a pairwise comparison matrix to generate normalized weights tied to installation priorities. These weights combine with site scores to produce a defensible ranking. While the results shown are from West Point sites selected to illustrate the model's usefulness, the framework itself is installation-neutral and exportable, enabling consistent, priority-driven site selection across Army installations.

Site Scoring & Ranking: (5 – Strongly favorable / Minimal mitigation required) (3 – Workable with moderate mitigation or cost) (1 – Major mitigation / High risk but not absolute disqualifier)

Site ID	Site Name	Fatal Flaw?	Security & Force Protection (1-5)	Resilience Contribution (1-5)	Speed to Field (1-5)	Electrical Integration (1-5)	Environmental & Permitting Complexity (1-5)	Constructability(1-5)	Logistics & Transportation Access (1-5)	Public Interface / Acceptance Risk (1-5)	Cost Drivers (1-5)	Weighted Score	Rank	Include?
1	PX	N	5	4	2	5	3	2	4	1	1	2.46	1	YES
2	Motorpool	N	2	3	2	2	4	4	5	5	3	2.17	2	YES

Figure 2: Multi-criteria scoring and ranking of candidate microreactor sites at West Point. Sites are evaluated across nine criteria on a 1–5 scale. Results identify the PX site as the top-performing location and demonstrate a transparent, exportable siting framework.

4. Results

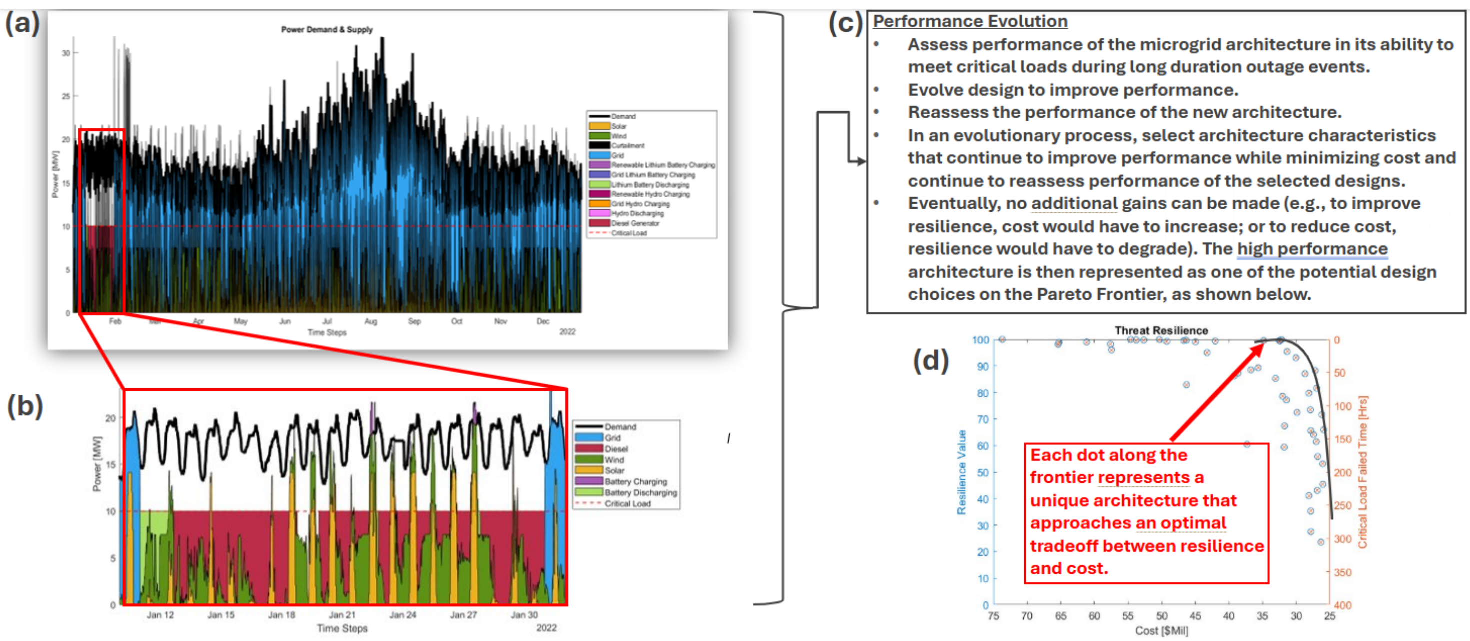

Figure 3: The model observes the performance of a microgrid architecture, as shown in (a), over the course of 8760 hours of a year. A more detailed look at a long duration outage event can be seen in (b). The model then evolves the given architecture, as shown in (c) in an attempt to further improve resilience while minimizing cost. High-performance architectures are then displayed in the final output of the model as one of the dots on the Pareto frontier, as shown in (d). The model plots each microgrid by their cost, decreasing on the x-axis, and resilience score, increasing on the y-axis.

The model outputs a Pareto frontier with each point representing a microgrid, providing the decision maker a variety of options for improving energy resilience. By having multiple options, decision-makers can balance the objective criteria of cost and resilience with subjective influences unable to be represented within the model. In a Pareto frontier, all dominated results will be excluded; if there were two points with the same resilience value, the more expensive will be dropped and the solutions will generate towards the frontier with the best results appearing in the top right corner as seen in Figure 3.The composition of each microgrid is outputted in a table format as seen in Table 2. Each microgrid is evaluated for resilience through a stochastically occurring grid outage event. In Figure 3 the demand is met primarily through the grid shown in blue. The rest is served by the other technologies shown in the legend. However, at the beginning of the year, there exists a gap, signifying a failure. A close-up of the failure is shown in section (b) of the figure. Here, the grid fails and the critical load of 10MW is primarily met by wind and discharging the battery, with some contribution from solar. The battery is unable to sustain the critical load infinitely and around January 13th, the diesel generator begins to operate. Around January 20th the system fails to meet the critical load for a period of time, which will reduce its resilience metric. When the system exceeds the regular demand in an outage, the battery is charged to later supplement the diesel in meeting the critical load.

ISBN: 97819384962-5-7

Table 2: Technology Configuration Results (Population 10 after 8 Generations NY)

Solar		**Wind**	**Li-Ion Battery**			**Diesel**				**Pumped Hydro**		**Demand**		**Results**	
Tracking	Panel Size (m^2)	Turbines (#)	Strategy	Energy (Mwh)	Power (Mw)	Strategy	Energy (MW)	Generators (#)	Size (MW)	Strategy	Selection	EV	Heat Pump	Cost ($M)	Resilience (%)
2	169000	19	6	190	20	1	1200	9	2	4	2	0	0	49	100
1	0	0	6	0	0	1	0	0	1	1	1	0	0	21	44
2	133000	19	6	570	16	1	1400	7	3	4	2	0	1	57	99
1	45000	10	6	66	2	2	670	4	3	1	3	1	0	28	95
1	88000	10	6	140	13	1	790	8	2	3	3	0	0	36	99

5. Conclusions and Recommendations

A pattern across all regions is that the cheapest solutions are not the most resilient. The plot displays a variety of solutions that were generated based on three different regional pricing scales. Each region results in the same trend. In both NY and TX, the lowest-cost option is about $21M, but resilience drops to about 44%. This suggests a minimalist system can save money up front, but it does not provide acceptable backup capability if resilience is a priority. The Pareto frontier offers a digestible visual from which different stakeholders can select. Their selections will vary depending on their personal budget, resilience needs, and risk attitude. Providing them with the frontier allows them to effectively map out their preferences to meet the demands of their installation.

Related to the above conclusion regarding expense, the results of this work found unique value in battery energy storage systems and diesel generators for providing backup to support critical loads during an outage event. However, modeling of long duration outage events and subsequent energy dispatch behavior identified fuel supply lines as a clear vulnerability for diesel generators. Technologies without fuel supply lines, namely solar photovoltaics and nuclear reactors, were shown to provide unique value when outage events stretched into numerous days. In general, architectures that mitiage the fuel supply line risk also require greater upfront capital cost commitments, another example of the cost and resilience tradeoffs that leadership of military installations need to wrestle with.

There exist other academic and commercial models and simulation software for assessing the performance of microgrids, but many of these lack important elements required to analyze the priorities for a military installation. Other software such as HOMER® are capable of modeling the failure of a set microgrid; however, these commercial solutions often lack the capability to model how a microgrid will respond to a stochastic failure of the grid. Furthermore, many of these software products do not account for the potential of modular nuclear reactors. Conversely, the DAVE model is able to incorporate and assess the value of nuclear power to support resilience. In addition, unlike the deterministic outputs of many commercial solutions, the DAVE model provides a Pareto frontier solution giving Garrison Commanders a decision support tool with multiple options that other models do not always provide.

References

Brownlee, J. (n.d.). *Algorithm afternoon.* `https://algorithmafternoon.com/about`. (Accessed: Apr. 8, 2026)

Grymes, J., Newman, A., Zolan, A., & Mehta, D. (2025). A matheuristic for design and dispatch of a utility-connected distributed energy system. *Journal of Heuristics*, *31*. doi: 10.1007/s10732-025-09567-0

Hashimoto, T., Stedinger, J. R., & Loucks, D. P. (1982, Feb). Reliability, resiliency, and vulnerability criteria for water resource system performance evaluation. *Water Resources Research*, *18*(1), 14–20. doi: 10.1029/WR018i001p00014

Jesse, B.-J., Kramer, G. J., & Koning, V. (2024). Characterization of necessary elements for a definition of resilience for the energy system. *Energy Sustainability and Society*. doi: 10.1186/s13705-024-00478-9

LaRose, A., Diefenderfer, J., & Namovicz, C. (2024, Jan). *Capital cost and performance characteristics for utility-scale electric power generating technologies.*

MathWorks. (n.d.). *What is the genetic algorithm? MATLAB & simulink.* `https://www.mathworks.com/help/gads/what-is-the-genetic-algorithm.html`. (Accessed: Oct. 15, 2025)

United States Code. (2024). *10 u.s.c. § 101.* `https://www.law.cornell.edu/uscode/text/10/101`. (Accessed: 2024-05-22)

United States Code. (2026). *10 u.s.c. § 2920.* `https://www.law.cornell.edu/uscode/text/10/2920`.

Wales, J., Zolan, A., Hamilton, W., Newman, A., & Wagner, M. (2023). Combining simulation and optimization to derive operating policies for a concentrating solar power plant. *OR Spectrum*, *45*. doi: 10.1007/s00291-022-00688-7