SimBuild
2024

Eleventh National Conference of IBPSA-USA
Denver, Colorado
May 21 – 23, 2024

A Comparative Case Study of Heuristic Optimization of an Affordable Housing Development

Robert E Best[1], Ananth Iyengar[1], Melinda Lai [1], Sara Tepfer[1], Terence Caulkins[1]
[1]Arup US, Inc., New York, United States of America

Abstract

A comparison of an affordable housing complex in Visalia, California was performed using a genetic algorithm to optimize energy conservation measures (ECMs) from a prior low-energy design by design practitioners. Proposed ECMs were developed and evaluated with a cloud-based workflow using NSGA-II to create combinations of ECMs tested against carbon, cost, and Time Dependent Valuation. Numerous lower carbon solutions were identified, with the greatest saving 9%. The biggest determinant of savings was the selected HVAC system, with passive strategies yielding limited savings. Despite this magnitude of savings, the approach yields positive payback and financeable solutions with greater stability in energy cost against future climate changes.

Introduction

Affordable housing can be among the most challenging sectors to cost-effectively decarbonize given the capital cost constraints of developers, limited mechanisms for cost recovery from tenants, and the premium of efficient and all-electric operations. Portfolio affordable housing developers have identified cost-effective strategies to achieve high performance within the constrained capital confines of these developments. However historically, these approaches have not enabled net-zero or grid-interactive buildings, thereby risking leaving behind a critical population of buildings and residents in the transition to a decarbonized economy. Simultaneously, a report from the American Council for an Energy Efficiency Economy indicated that decarbonizing low and moderate income housing could avoid 140 million metric tonnes of operational CO_2 annually (York 2022). They continued that a combination of efficiency and electrification was critical to achieving this goal (York 2022). Identifying means to achieve this reduction is essential given that globally, operational carbon emissions from buildings are responsible for 28% of total global carbon emissions (IEA 2024). Given that highly efficient buildings with on-site energy resources often have co-benefits such as increased resilience, reduced risk of energy price inflation, and improved thermal comfort, it is paramount to identify means of supporting affordable housing in the move toward decarbonized buildings.

Optimization techniques such as genetic algorithms have offered a means to cost-effective decarbonization pathways and highly efficient building design. However, the merits of these strategies applied to affordable housing designs is limited, as is the comparison to a sector which has already developed efficient designs. Given the perceived high cost of a parametric optimization approach by practitioners (Barber and Krarti 2022), understanding the value optimization provides to affordable projects is helpful in trading off the value of added design complexity and downstream performance. This paper presents a workflow to deploy a genetic optimization strategy in a replicable, cost-effective fashion within the context of a typical practitioner’s energy modeling workflow. It further provides a comparison of an affordable housing complex in Visalia, California aiming for deep decarbonization between a portfolio affordable housing developer’s best performing project and a design set generated via optimization with the same geometry and site layout via the algorithm. A genetic algorithm (GA) was utilized to optimize various energy conservation measures (ECMs) and compare to the team’s prior low-energy design developed without the benefit of the GA. Given the high-performing starting point of the baseline project, the design was also compared to current California Energy Code to determine the potential of optimized affordable housing designs for broader energy savings in the industry.

Background

Building energy optimization using metaheuristic algorithms has been increasingly applied to identify pathways for deep energy savings. Among the algorithms commonly used, genetic algorithms (GA) are commonly applied as they have been shown to offer fast convergence to a stable solution across competing objectives (Wortmann 2017). Case studies have shown the effectiveness of approaches such as genetic algorithms (GA) in significantly reducing energy use across a range of building types (Ciardello 2020). These have shown the value of optimizing both geometric form and building construction and operation characteristics such as envelope composition, heating, ventilation, and air conditioning (HVAC) system type, setpoints, and equipment.

Several studies have applied GA to residential buildings. One study in Oman found that a GA was able to effectively identify a range of optimal solutions with 16-29% variation in energy performance compared to capital cost (Al-Saadi 2020). Many of these studies benchmark against typical construction characteristics rather than high performing code baselines.

Furthermore, few studies compare optimized designs to energy efficient strategies determined by professional designers. One reason a comparison of professional solutions compared to computational designs is not available is provided by Barber and Krarti, who found that data-driven design and optimization is often used for specific subsystems such as renewable energy and storage rather than for entire designs (2022). A prior study of a GA applied to a grocery store retrofit found that the highest performing solution of the GA performed 9% better than the initial solution identified by the professional design team (Best, et al. 2018).

The perceived higher cost and complexity of deploying optimization techniques for building energy modelling in practical projects coupled with the absence of rigorous studies comparing best-in-class practitioner created designs with optimized outcomes creates a need for understanding how computational optimization can support high performing design and improve designs relative to the code minimum energy performance.

Methods/Simulation

Leveraging a genetic algorithm to compare a previous design against a computer optimized form required constructing and parameterizing an initial baseline energy model. The baseline model was created in EnergyPlus to represent the prior precedent building design of the architecture and engineering team for the affordable housing developer.

The baseline model in EnergyPlus was generated from the permitted prior all-electric design deployed in Tulare, CA, by the development team. This was noted as the latest iteration of their affordable housing village, the development of which spanned several generations of improvements driven both by California Energy Code (Title 24, Part 6) and efforts from the design team's experience to reduce energy consumption and cost through engineering and architectural decisions. The village, shown in Figure 1, consists of multiple buildings, each of which represents the same prototype of a 12 unit structure connected side by side with a shared wall.

Figure 1 Visalia Village Layout

To simplify the modeling, a single 12 unit structure was used for analysis rather than the entire site. This basic unit comprises the residential units and the central stairwell and common hallways and services for all units. This unit is shown in Figure 2.

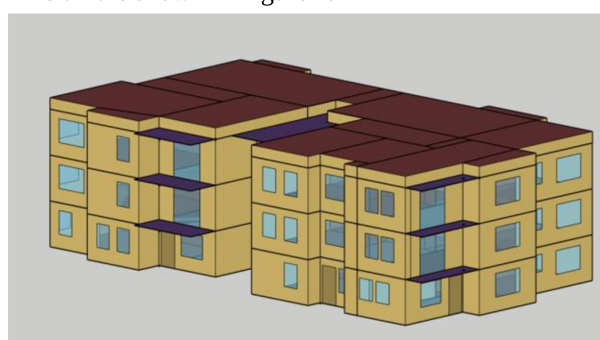

Figure 2 Single building unit as modeled in EnergyPlus

To confirm that the selection of functional unit for analysis was appropriate, the team tested the unit in various orientations and rotations within the site and determined that the variation in energy performance was less than 5% for all iterations of shared wall, rotation, and orientation. While past studies have identified orientation and massing as important characteristics for optimization toward low energy buildings (e.g., Ascione 2019), for this study massing was treated as fixed given the cost and layout optimization inherent in the proposed unit configuration and design.

The baseline building was simulated using the Visalia TMYx meteorological year weather file. Where operating schedules were known based on past project information the project developer, these were adopted in the model; otherwise, schedules from the Pacific Northwest National Laboratory Residential Prototype Building Models were utilized (Mendon and Taylor 2014). Using these inputs, the baseline energy model when simulated produced an energy consumption of 27 kBtu/SF/yr.

Parameterization of Model Inputs

To develop the potential solution space, envelope, lighting, equipment, and HVAC inputs were identified that had feasible alternate solutions for the project. These were developed in a team-wide charrette to ensure that the relevant owner, architect, contractor, and engineers agreed that solutions proposed could be feasible for the project. Each selected system was modeled as a potential energy conservation measure (ECM) that could be included in combination with other measures unless the same type of measure would conflict with another (e.g., multiple lengths of window overhang). Geometry variations were limited to roof and window overhangs to reduce solar gain; no massing variations were studied at this stage of the project. A future refinement to include massing variations alongside envelope and systems measures could potentially be considered for improved savings. 27 total ECMs were identified by the team that were deemed feasible, covering the following categories:

- HVAC: 5 measures
- Ventilation: 1 measure
- Envelope: 11 measures
- Refrigerants: 1 measure
- Shading: 8 measures
- Water Heating: 1 measure

On-site energy generation and storage were not considered as part of the algorithmic optimization, though they are intended to be included as part of the final construction project. Both were evaluated separately using the EnergyPlus model results. A full list of the ECMs studied is presented in Table 1.

Table 1 Energy Conservation Measures (ECM) considered

HVAC	**Ventilation**	**Envelope**	**Refrigerants**	**Shading**	**Water Heating**
• **Ductless split system** • **Building level 4-pipe ASHP with radiant floors** • **Unit level 2-pipe ASHP with radiant floors** • **Building level 4-pipe ASHP with FCU** • Unit level 2-pipe ASHP with FCU	• Energy recovery ventilator (ERV)	• Improved roof R • Wall Insulation (1" continuous mineral wool board) • Wall Insulation (Staggered stud 2x4 at 16" oc w fiberglass) • Wall Insulation (2" continuous XPS) • Improved glazing U and SHGC • Improved glazing (Alpen Windows Zenith Series ZR-5) • Thermal mass wall • Thermal mass floor • Reduced infiltration. • Phase change materials • High reflectance TPO roofing	• Low-GWP refrigerant swap in unitary heat pumps	• 3ft roof overhang • 5ft roof overhang • 11in recessed windowsill • 15in recessed windowsill • 2'6" window overhang • 3' window overhang • 3'6" window overhang • 4' window overhang	• HPWH in unit

Each measure was provided with an EnergyPlus model fragment that could be substituted into the main model to simulate the component. To enable multi-objective optimization, capital costs of each measure were provided by the contractor against the baseline system.

Optimization via Genetic Algorithm

NSGA-II was selected as the algorithm to be deployed for the optimization, as this algorithm has been shown in prior studies to converge relatively quickly to stable solutions for building energy modeling (Ciardello 2020). The NSGA-II genetic algorithm was implemented with a Python and JavaScript software and a cloud-based computational workflow. The workflow takes in the baseline building specification and potential ECMs, performs the genetic algorithm, stores results on an ongoing basis, and creates a final visualization to evaluate the results. Each building specification was stored in an EnergyPlus readable Intermediate Data Format (idf) file that can be manipulated using text-based addition, deletion, and modification. The ECMs were inputted as coded instructions that the software could use to make the appropriate updates to the idf file to prep it for simulation. If many ECMs were chosen, the software sequentially modified the baseline idf file in a pre-defined order. While the number of ECMs for the Visalia case study was limited as noted previously, this approach was developed as it is scalable to a number of ECMs and allows for ready parameterization of common inputs such as lighting performance, insulation thickness, and other values which can be used to create a library of common replicable ECMs.

The software utilizes the Distributed Evolutionary Algorithms in Python (DEAP) Python package to carry out the specific calculations of the genetic algorithm and the scoop module to run tasks concurrently. The scoop module allowed for the thousands of simulations to be created and processed asynchronously, which provided significant reduction in computational time to fit within a typical schematic design timeline. Python packages that manipulate EnergyPlus input files directly such as eppy were considered for use but ultimately discarded as the necessary manipulations for the purposes of the Visalia Village use case required more involved changes to each individual building specification. Multiple chosen ECMs required a combination of changes that would not have been feasible using the eppy package alone. In future use cases, it may be beneficial to leverage the foundation and abstraction of the eppy package to simplify the parametric approach, provide more out-of-the-box functionality, and streamline further the start-up costs of base functionality.

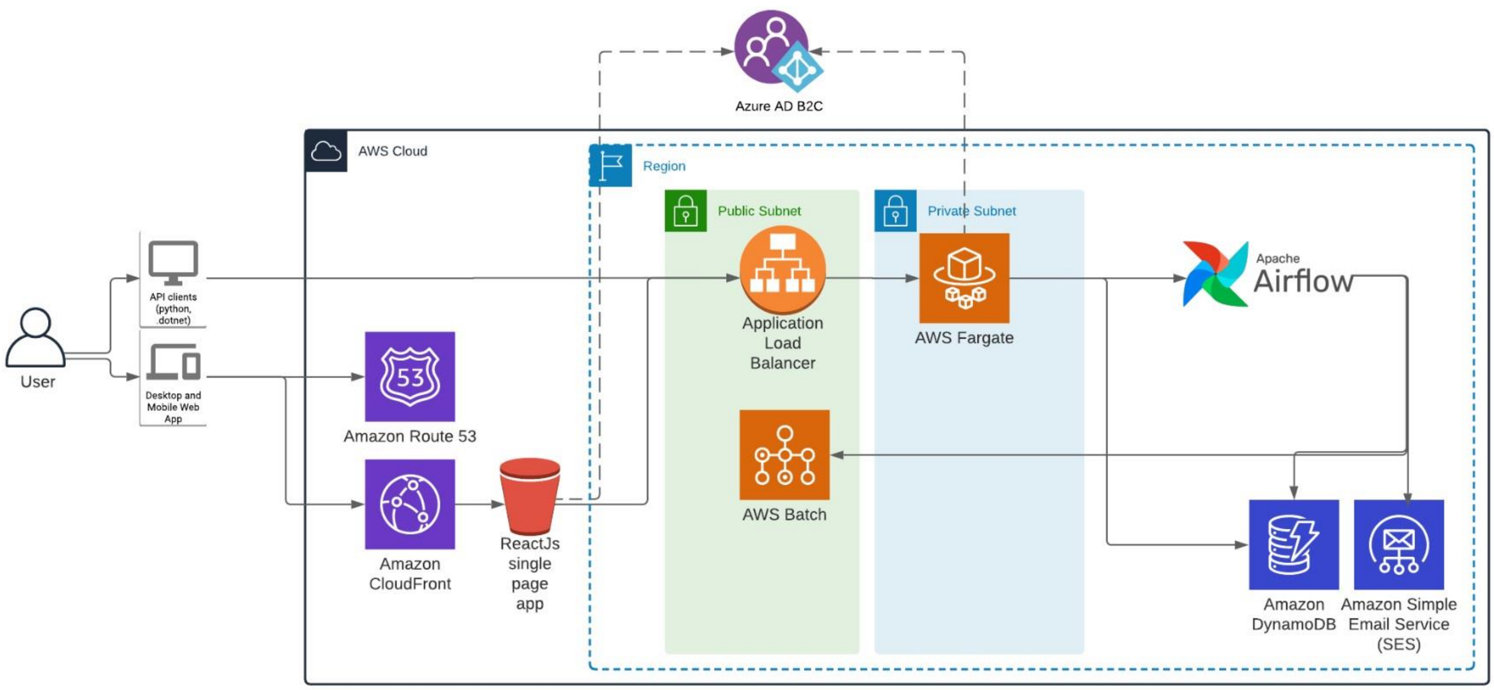

Figure 3 Cloud-Based Energy Modeling Solution Architecture

Utilizing the Python-based backend of the software, each idf file was generated, pre-processed, and transferred up to a custom, AWS-based solution to perform large quantities of energy simulations. Figure 3 shows this full cloud-based energy modeling software architecture.

This cloud-based approach to performing the energy simulations allowed for rapid evaluation of results. Leveraging Elastic Load Balancers, AWS Batch Compute, and Apache Airflow for process scheduling, models were processed via a localized version of EnergyPlus in a streamlined, sequential manner. All results were saved into a DynamoDB NoSQL database that could be queried by the genetic optimization codebase for scoring. In tandem with the use of the DEAP python package, energy simulations were run asynchronously and in batches, significantly minimizing the time required for processing models. As multiple thousands of models would need to be run over many different executions of the genetic optimization, any savings on computation time was necessary to meet deliverable timelines.

The backend of the software was designed to handle test cases as well as the full NSGA-II genetic algorithm. All test cases posed were to evaluate the feasibility and functionality of each ECM individually combined with the baseline building specifications. Distinctive modules ingested the given ECM specifications and modified the baseline idf file accordingly. All ECM-incorporated idf files were sent via the python SDK to the cloud-based solution, and the backend would await completion of the simulations. After all simulations were completed, results and errors were compiled and presented back to the user. From there, energy modelers could evaluate the effectiveness of standalone ECMs and address whether the results reflected were accurate or required further programmatic intervention. Once all ECMs were validated and confirmed, the genetic algorithm could be fully implemented. This step provides a higher startup for initial testing, but ultimately validates the database of ECMs for potential replication on additional buildings.

The initial step in the algorithm was to create a randomized population of building specifications based on the ECMs provided. In order to get a statistically near-random sample of buildings, Latin Hypercube Sampling was utilized. Each individual building design (or model) within the population was given a randomized set of ECMs. Data associated with the model were critical for assessing its estimated performance but also for use in the analysis synthesis. Each individual was given a unique identifier that could also be utilized for deeper analysis and debugging.

Due to restrictions in the expected final building design and with inter-ECM compatibility, not all randomized combinations of ECMs were appropriate for analysis. Before modifying each file and performing the simulations, each individual was checked, and if a suboptimal number of them were deemed valid, a new randomized population would be constructed. This check would occur until enough models were valid or an acceptable breaking point was reached.

Table 2 Energy conservation measure bundle parameters

Optimized	Operational Carbon ($kgCO_2e$)[1]	Annual operational carbon emissions, calculated with the annual energy usage and the 2022 grid emissions factors for CAMx
Optimized	Capital Cost (USD)	Sum of cost scores applied to each individual ECM in a bundle. Cost scores are rough order of magnitude estimates that represent the cost delta between the ECM and the baseline system. Cost numbers were determined in conjunction with design team.
Tracked	Total Annual Energy (MBtu)	Annual site energy
Tracked	TDV Energy (MBtu)[2]	Time Dependent Valuation (TDV) is a metric that adjusts energy use depending on the impact of that use during a given hour of the year. CZ13 TDV (residential 30 year) values are used.
Tracked	2050 Operational Carbon ($kgCO_2e$)[3]	Annual operational carbon emissions, calculated with the annual energy usage and the 2050 grid emissions factors for CAMx
Tracked	Operational Cost (USD)[4]	Calculated based on current Southern California Edison TOU-D-PRIME rate schedules
Tracked	Embodied Carbon	Sum of embodied carbon scores applied to each individual ECM in bundle. Embodied carbon scores are qualitative and ranked on a scale from 1 to 5. 1 represents high embodied carbon savings, 3 represents no change from the baseline, and 5 represents high added embodied carbon.
Tracked	Replicability	Average of replicability scores applied to each individual ECM in bundle. Replicability scores are qualitative and ranked on a scale from 1 to 5. 1 represents ECMs that are hard to replicate on future projects, and 5 represents ECMs that are easy to replicate on future projects. Replicability numbers were determined in conjunction with design team.

[1] http://scenarioviewer.nrel.gov/

[2] https://wcec.ucdavis.edu/resources/what-is-time-dependent-valuation-tdv/

[3] http://scenarioviewer.nrel.gov/

[4] https://www.sce.com/residential/rates/Time-Of-Use-Residential-Rate-Plans

After batch computation, results were postprocessed to evaluate embodied carbon accounting, energy costing, operational carbon adjustments, and Time Dependent Valuation (TDV) consideration per California's Title 24 2022 energy code. Embodied carbon considered only change relative to the options studied and was calculated by summing the embodied carbon in elements that varied as a result of the ECMs. Structure and façade elements common to all scenarios were not evaluated. Calculations were based on standard materials databases for carbon content of materials. Table 2 provides the definitions of all objectives and tracked, post-processed variables used in the study.

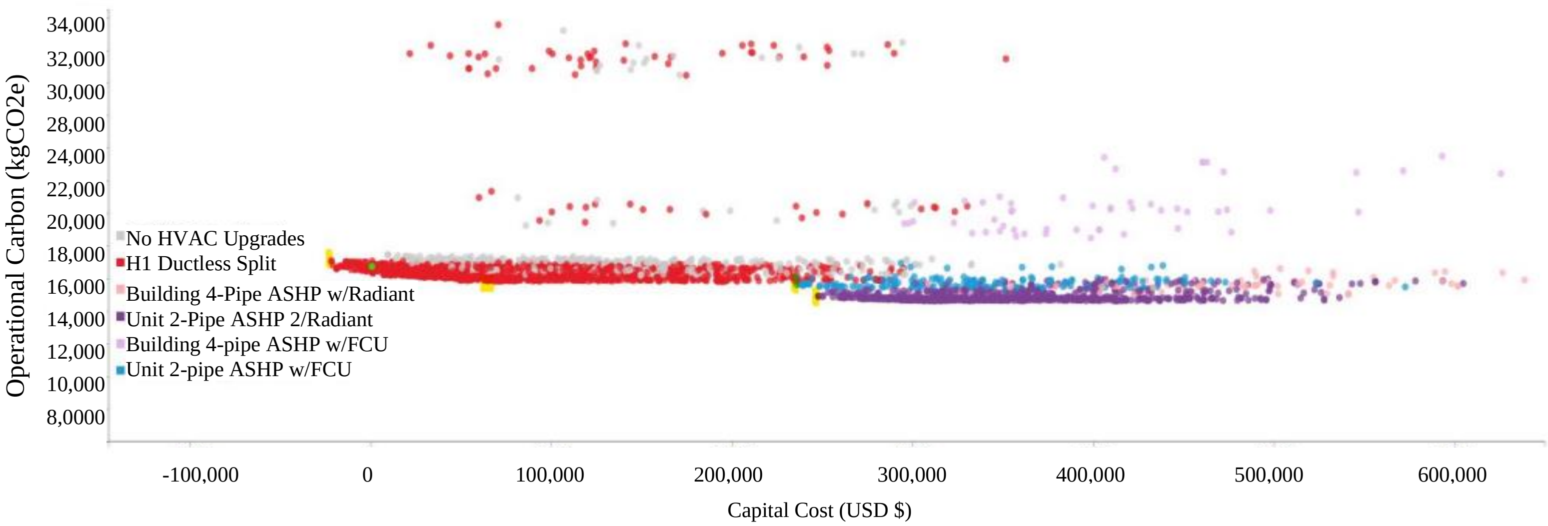

Figure 4 Results of genetic algorithm optimization

Utilizing this post-processed, calculated data, each model's profile could be evaluated against the predetermined fitness function. In the case of Colegio Village affordable housing community, operational carbon and capital costs were deemed the most relevant for guiding the evolutionary algorithm. These two metrics represented the steering parameters for scoring each model, however, other metrics were saved to provide additional points of analysis for evaluating a holistic design recommendation. From this point, the GA continued through multiple generations until a generation-based stopping criteria was achieved. An average fitness improvement metric was tested as well, but ultimately did not out-perform a generation-based stopping criteria, so the latter was used for simplicity.

The final component of the software solution developed for this genetic algorithm was a JavaScript based visualization tool that could be used to evaluate the results and compare ECMs against one another. This tool allowed for designers and engineers to compare different pertinent metrics against one another to provide a recommendation. Evaluators could isolate individual generations to see how drastically the convergence was happening from generation to generation. Additionally, the tool allowed for individual ECM isolation, so designers could assess which design decisions might be consistently the most impactful in the overall building's performance.

Future Climate Modeling

All solutions in the GA optimization were performed utilizing the Visalia, CA TMYx weather file, similar to the baseline case. However given the benefit of energy efficient designs in mitigating extreme weather conditions, the final Pareto optimal solutions were tested as well against future climatic conditions using fTMY files generated from the WeatherShift platform. This approach uses a morphing approach to transmute TMY3 weather data to future conditions based on downscaling of coarse resolution climate model predictions (Belcher 2005). These shifts are based on projected changes in monthly averages of several climatic variables developed for the Intergovernmental Panel on Climate Change (IPCC) Fifth Assessment Report (AR5). To develop these shifts, the AR5 report defines several emissions scenarios for the 21st century and projects changes in climate variables for each of them. Weathershift has data available for two scenarios, Representative Concentration Pathway (RCP) 4.5, which represents an emissions stabilization scenario, and RCP 8.5, which represents 'business as usual.' This study used the more pessimistic scenario, Representative Concentration Pathway 8.5, for this analysis.

Conditions were tested for a base year of 2050 and a base year of 2090, referred to as "Mid-Century" and "Late-

Century," respectively, to identify changes in operational energy performance over time as warmer climates in Visalia prevail.

Results

Genetic Optimization Results

The GA was allowed to run over 20 generations, exploring 3,427 of a possible 1.1 million ECM combinations. As noted above, the primary focus was on the tradeoff of capital cost and annual operational carbon emissions. Capital cost was chosen as the appropriate cost tradeoff metric given the typical constraints on financing and per unit cost for affordable housing.

With each simulation expressed as a point on a scatter plot of these two objectives, Figure 4 shows the results of this simulation. Coloration in the graph below indicates the different HVAC systems evaluated, which was identified as one of the primary drivers of energy performance and cost.

Based on the results of the genetic algorithm, four ECM bundles along the Pareto-optimal front were selected to be compared to the baseline. The overall lowest cost bundle was chosen, which was a bundle with ductless split system HVAC. The ductless split system bundle with the most operational carbon savings was also chosen to be evaluated. Lower carbon solutions evaluated included unit-level 2-pipe air source heat pumps; centralized 4-pipe heat pumps with radiant flooring were not preferred for cost, and the improvement in operational carbon was marginal relative to the highest performing 2-pipe heat pump alternate. Additionally, a bundle with unit-level 2-pipe air source heat pumps with fan coil units were deemed to be equivalent in energy but lower in cost then building-scale 4-pipe heat pumps with fan coil units.

Table 3 shows these selected results combined into a table with characteristics of ECMs selected in the alternate, capital cost increase relative to the baseline cost, and operational carbon in kg CO_2 per year. Carbon performance was evaluated using an hourly carbon emissions profile for the California energy grid, allowing for inclusion of time of use carbon impacts driven by changes in shifting generating resources on the grid. This is important when evaluating high-performing designs to study grid-interactivity and ensure that carbon can be abated properly during high grid-carbon periods through on-site energy generation and storage.

Table 3 Selected Capital Cost and Carbon Packages

Model	HVAC System	ECMs	Capital Cost Increase ($)	Operational Carbon (kg CO_2)
Baseline	Ducted split system	None	$0	18,550
Lowest capital cost	Ductless split system	- HPWH in unit - Low-GWP refrigerant in unitary heat pumps	$-21,650	19,062 *+3%*
Lowest operational carbon w/ ductless split system	Ductless split system	- Improved roof R-value - Thermal mass floor - Reduced infiltration - 15" recessed window sill	$66,550	17,833 *-4%*
Lowest operational carbon w/ 4-pipe ASHP with radiant	2-pipe ASHP and radiant floors	- HPWH in unit	$247,950	16,901 *-9%*
Lowest operational carbon w/ 4-pipe ASHP and FCU	2-pipe ASHP with FCU	- Improved roof R-value - 15" recessed window sill - Shading (4' window overhang)	$255,250	17,828 *-4%*

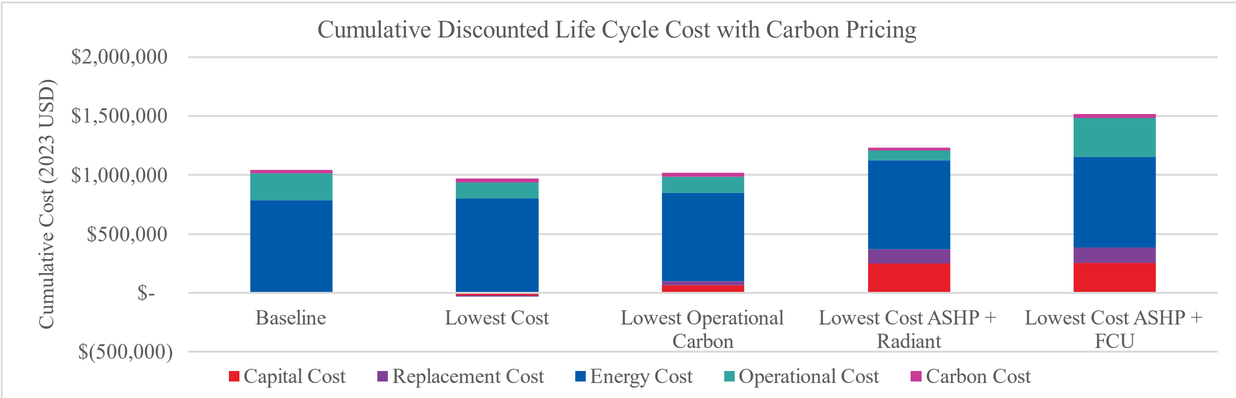

Figure 5 Life cycle cost results of the selected high performing cases on the Pareto-optimal front

As noted in Table 3, the highest performing option from a carbon perspective reduced emissions by 9% from the baseline case, primarily by impactfully reducing peak energy and emissions during typically high emissions periods. However, it does so at a premium of $250,000 per 12-unit building for the project. This premium is primarily driven by the high cost of a radiant flooring system and the air to water heat pump required to drive the heating and cooling. By comparison, ductless split systems are significantly lower cost. This results in an inability of the saved energy and operational cost overcome the capital premium over a 30-year life cycle, as noted in Figure 5.

By contrast, the lowest operational carbon alternate with a ductless split requires $66,550 more up front per building but has a positive life cycle savings of $24,803 for the building, which may allow for financing of the higher capital cost. However this alternate limits carbon savings relative to the baseline to only 4%.

Operational energy costs in the above were determined using Southern California Edison's residential time of use rate TOU-D-PRIME. This rate provides a higher energy cost during the 4 pm – 9 pm period in both summer and winter.

In addition to carbon emissions and cost, in California, the Title 24 Part 6 Energy Code requires evaluation of Time Dependent Valuation, a metric which utilizes a blend of hourly energy cost and hourly carbon emissions on the grid. Evaluation against the hourly performance of an energy model is required to determine compliance. When using this metric, the savings even for the lowest carbon alternate (2-pipe ASHP with radiant heating and cooling) is more marginal, improving TDV by only 4.7% relative to the baseline.

Comparison to Code Minimum Baseline

As noted in the Methods section, the baseline utilized for this study was a precedent design created by the affordable housing developer and a design team that had iterated through several versions of an efficient affordable product. Therefore, the starting point can be assumed to outperform a minimum building under California Energy Code, Title 24 Part 6. To test the optimized building against a code-based comparison, a minimally compliant building energy model with the same geometry was created in EnergyPlus. Simulated for the same weather file, and prior to application of on-site energy generating and storing resources, the highest performing solution (2-pipe air source heat pump with radiant floors) was found to reduce carbon emissions by 19%. This indicates also that the baseline building against which all options were evaluated already outperforms a code minimum performance by 10%.

Future Climate Results

Utilizing the EnergyPlus idf files for the four Pareto-optimal solutions and the baseline, future climate scenarios were tested against the Mid-Century and Late-Century weather files selected and noted above. The results of this analysis are shown in Table 4.

As noted earlier, the benefits of the optimized package compared to the baseline were shown to be marginal when comparing carbon emissions and TDV. However the added envelope measures and higher efficiency HVAC systems allow these savings measured as EUI, carbon emissions, and operational cost to grow over time to the end of the century. This finding suggests a claim that the optimal solutions with deeper efficiency and improved short-term carbon performance are also more resilient to the impacts of a warming climate. In effect they provide a hedge against rising energy consumption and cost resulting from the changing climate.

Table 4 Annual energy, carbon, and operational cost per square foot for each time horizon and scenario

Model	Time Horizon	Operational Energy (kBtu/ft^2-yr)	Operational Carbon ($kgCO_2e/ft^2$-yr)	Operational Cost ($/ft^2-yr)
Template design (baseline)	Historical	25.2	1.5	$2.45
	Mid-Century	26.1	1.5	$2.53
	Late-Century	27.0	1.5	$2.62
Lowest capital cost	Historical	25.9 *+3%*	1.5 --	$2.50 *+2%*
	Mid-Century	26.3 *+1%*	1.5 --	$2.57 *+1%*
	Late-Century	26.7 *-1%*	1.5 --	$2.61 *-1%*
Lowest operational carbon w/ ductless split system	Historical	24.1 *-4%*	1.4 *-4%*	$2.36 *-3%*
	Mid-Century	24.6 *-6%*	1.4 *-5%*	$2.41 *-5%*
	Late-Century	25.1 *-7%*	1.5 *-6%*	$2.45 *-7%*
Lowest operational carbon w/ 4-pipe ASHP with radiant	Historical	23.7 *-6%*	1.3 *-8%*	$2.36 *-3%*
	Mid-Century	24.2 *-7%*	1.4 *-9%*	$2.40 *-5%*
	Late-Century	24.5 *-9%*	1.4 *-11%*	$2.44 *-7%*
Lowest operational carbon w/ 4-pipe ASHP and FCU	Historical	25.1 *-1%*	1.4 *-1%*	$2.46 *+1%*
	Mid-Century	26.0 --	1.5 --	$2.55 *+1%*
	Late-Century	27.0 --	1.5 --	$2.65 *+1%*

The cause of this is reduced heating and cooling need in the peak heating and cooling seasons, as shown in Figure 6 where the bars show the change from the base year (solid) to mid-century (mid-opacity) and late-century (low-opacity). The graph indicates that changes in the peak energy consumption by month are lowest in the models which demonstrate the lowest short-term operational carbon savings, driven by improved internal thermal performance and higher efficiency heating and cooling units.

Discussion

The Results section noted that leveraging the metaheuristic optimization to identify low carbon, high-efficiency strategies for the Visalia Village affordable housing development generated less than 10% annual carbon emissions savings prior to the application of on-site energy generation and storage resources. Even when compared to a code baseline, savings were capped at 19%. Several drivers are contemplated as the cause of this limited improvement in performance.

One key driver is the restricted ability to vary massing and geometry. As noted in the Background section, prior studies have indicated that optimizing form of the building alongside internal and envelope parameters provides the greatest opportunity to identify low-carbon designs. Given the limited ability to modify form on the project site at the time this study was performed, potentially deeper savings generated by changing massing were not evaluated. Part of the reason for limiting evaluation of form as well was the timing of the study related to a real-world project. Within the design

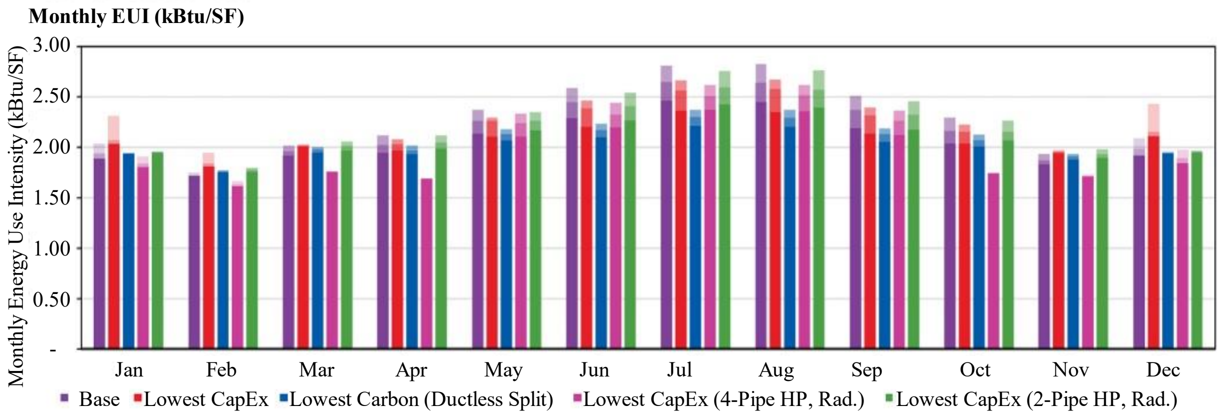

Figure 6 Monthly variations in energy between tested scenarios in base year, mid-century, and late-century climatic conditions

process, energy modeling is typically not undertaken until schematic design when form is largely set by either site constraints or architectural intent. Large changes in massing are unlikely at this point, leaving little latitude for geometric moves beyond small changes, overhangs, roof depths, and wall impacts. Modeling for this case study aligned with this typical process and was undertaken in schematic design.

An additional understanding of why performance improvement appears to be limited despite the application of computational optimization is the large contribution of interior equipment to the baseline energy consumption profile. Figure 7 shows the breakdown of energy consumption for the lowest carbon emissions scenario.

As can be noted in the graph, over 50% of remaining electricity is from interior equipment. This is notable since interior equipment were not part of the ECMs studied for two reasons. First, the appliances deployed in each unit were already EnergyStar compliant in the base building and therefore only marginal energy performance was anticipated from studying in detail the appliances. Second, operating characteristics and details of interior equipment are largely determined by occupants and are difficult to regulate or manage from a building owner and operator perspective. Decisions during design and construction other than selection of appliances are unlikely to impact this portion of energy consumption. Additionally, equipment are typically considered to be unregulated loads by California Energy Code, requiring in code compliance simulations for equipment loads to be constant between options evaluated as these are typically the purview of the tenant/resident or are governed federally and not subject to local codes.

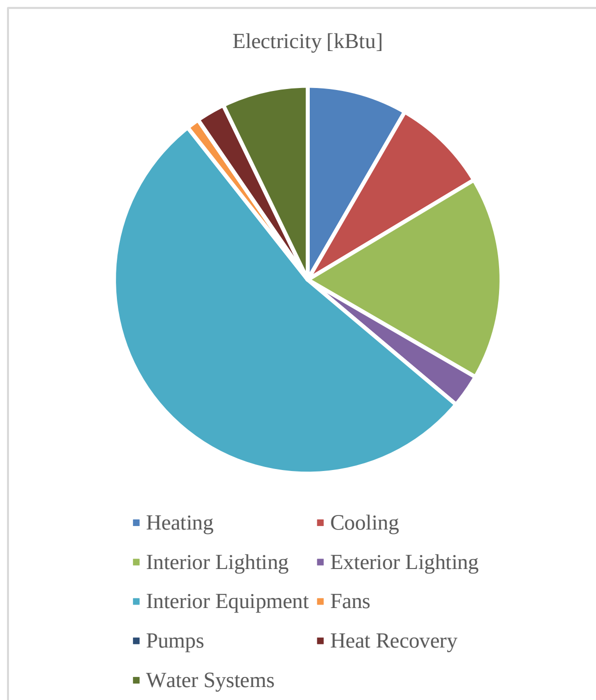

Figure 7 Electricity use breakdown for simulated low carbon building

A third main driver of the limited upside performance from applying optimization techniques is the high performance of the starting baseline. Many prior studies evaluating the potential of genetic algorithms to improve energy efficiency in building design are self-referential, meaning that they study the performance of the best option either to the initial baseline created or to against the worst performing option in the resulting solution set. For existing buildings, performance is usually against the prior energy consumption which may be high relative to new construction due to older standards and failure to properly maintain energy systems over time. Even when studies compare to building codes, often reference starting points are versions of ASHRAE 90.1 or International Energy Conservation Code (IECC). The starting version of the selected code as described in the standard year has a large impact on the potential savings available from the baseline due to continuous cycles of improvement in these standards.

For this evaluation, the starting baseline was a scenario that already outperformed the incumbent energy code by 10%. This therefore limited the availability of higher performing solutions. Furthermore, the proposed baseline featured no natural gas on-site, which also limited the depth of carbon emissions accessible to the optimization model by eliminating fuel switching as a strategy. Additionally, the California Energy Code, the baseline for all buildings in the study, is known to be a high performing standard that meets or exceeds the latest ASHRAE 90.1 standard.

When evaluated compared to other studies comparing optimization to the best design or standards in practice, the findings are well aligned. As noted in the Background section, prior studies identified 9% savings for an optimized design relative to a practitioner led design for a grocery store retrofit, and in Oman, computationally optimized housing performance exceeded code by an average of approximately 20%; both benchmarks align with the findings of this study.

This indicates that as codes and standards improve to further reduce energy performance for all buildings governed by the standards, the marginal improvement available to even computational standards is limited. However, as approaches to computational design improve, for instance using reusable libraries of measures and results as noted in the methodology applied to this case study, the marginal cost of computational design also decreases. This tradeoff will likely govern the rate of adoption of computational optimization for building energy efficiency within the industry.

Additionally, further studies should be performed to evaluate where practitioner intuition and standard approaches to energy modeling-driven decision-making in schematic design fall shortest compared to computational optimization. Identifying the sectors where computational design has the greatest benefit relative to the current standard of practice would be a benefit to the current literature and potentially identify the greatest opportunity for metaheuristic techniques in building decarbonization.

Conclusion

This case study demonstrated that for an affordable housing project in Visalia, CA, application of a genetic algorithm-based optimization to evaluate carbon emissions relative to capital cost identified solutions with up to a 9% improvement in emissions relative to a high-performing, practitioner-led baseline design. Relative to the baseline energy code, savings were 19%. While lower overall in savings compared to some prior published energy and carbon reductions using metaheuristic optimization, these indicate that prior art in the industry combined with leading codes and standards is capable of leading to high-performing designs, and that there is value in doubling the potential magnitude of energy savings through computational techniques. These findings match limited prior studies comparing computational techniques to practitioner performance, though more study is needed to identify the sectors in which optimization provides the greatest savings against current design practice.

Finally, the comparison evaluated in this study demonstrating carbon savings of 20% relative to a code baseline through a comparative approach has broader applicability to the affordable housing development industry. Identifying the combination of water heater and HVAC measures that reliably led to lower performance and, when coupled with envelope improvements, to resilience against rising temperatures creates a replicable baseline for ensuring that at-risk populations are included in the energy transition and decarbonization movement equitably and that the value of computational techniques in decarbonization is applicable to even capital-constrained development scenarios.

Acknowledgments

This work was supported through a grant from the California Energy Commission via GFO-20-305 – The Next EPIC Challenge: Reimagining Affordable Mixed-Use Development in a Carbon-Constrained Future.

References

Al-Saadi, S.N. and Al-Jabri, K.S. (2020) Optimization of envelope design for housing in hot climates using a genetic algorithm (GA) computational approach. *Journal of Building Engineering, 32*.

Barber, K.A., and Krari, M. (2022). A review of optimization-based tools for design and control of building energy systems. *Renewable and Sustainable Energy Reviews 160.*

Belcher, S.E., et al. (2005). Constructing design weather data for future climates. *Building Services Engineering Research and Technology 26*.

Best, R. and Levine, E. (2018). An Innovative Approach to Evaluating Energy Efficiency Measures for Zero Net Energy Supermarkets. *Proceedings of the ACEEE Summer Study 2018*.

Ciardello A., et al. (2020). Multi-objective approach to the optimization of shape and envelope in building energy design. *Applied Energy 280*.

International Energy Agency (2024) Building. Online Resource, Accessed February 5, 2024.

Wortmann, T., et al. (2017) Are Genetic Agorithms Really the Best Choice in Building Energy Optimization? *Proceedings of the Symposium on Simulation for Architecture and Urban Design 2017*.

York, D. ,et al. (2022) Building Decarbonization Solutions for the Affordable Housing Sector. *Report of the American Council for an Energy Efficiency Economy*.