Final Scientific / Technical Report

Reporting Period: 10/1/2020 – 6/30/2024

Prepared for:

U.S. Department of Energy

Submitted by:
PACCAR Inc.
Bellevue, WA.
DUNS Number: 048-341-267
Prime Recipient Type: Private Company

DE-FOA-0002197: FY 2020 Vehicle Technologies Program Wide Funding Opportunity Announcement

Co-optimization of Vehicle and Routes (CoVaR) to Improve Commercial Transportation System Efficiency

DOE Program Award Number: DE-EE0009207
Award Type: Cooperative Research and Development Agreement

Prepared by:
Nick Hertlein
Principal Investigator
nick.hertlein@paccar.com

Contents

1. Executive Summary

The PACCAR Co-optimization of Vehicle and Routes (CoVaR) to Improve Commercial Transportation System Efficiency has completed the planned research, development and demonstration work for the project. This work includes demonstrating a 25% reduction in brake-specific energy per freight ton-mile (kWh/ton-mi) relative to a measured baseline on a fleet engaged in active commercial transportation in United States. Due to the recent developments in electric powertrains for commercial vehicles, the tools developed under the program are designed to be powertrain agnostic to benefit both diesel and electric powertrain options while demonstration focused on both technologies given that most existing fleets today operate with Diesel powertrain vehicles.

PACCAR followed a structured approach, evaluating new technology concepts early in the program with state-of the art simulation tools supported by verification on a local test fleet. Verification of concepts on a local test fleet was vital to ensuring minimal impact on our partners' business operations. As a result, the team was able to reach high maturity levels on new technologies earlier with less iterations. Due to the application of the technology on customer vehicles in another area of the country, it was especially vital to minimize rework loops as implementation of the changes was not able to be performed remotely in all cases.

The CoVaR program included development of a low-cost high-resolution telematics system, an Intelligent Driver Assistance (IDAS) system to provide information to the driver, a Fleet Management System (FMS) to allow the fleet manager to interface with the vehicles, and a Vehicle Specification Optimization Tool to determine which vehicle and powertrain specification will provide the best freight efficiency for the customers use cases. The program focused on developing technologies with high commercialization potential, keeping complexity and cost to a minimum while ensuring a robust solution. As a result, many of the technologies developed are now included in PACCAR's technology roadmap and production development programs.

The program has improved freight efficiency by 32% using the Vehicle Specification Optimization System for diesel powered vehicles. This powertrain optimization potential is regarded as the best-case scenario, optimizing fleet vehicles for specific routes and applications. PACCAR's experience, confirmed by the data-collection under this program, shows that fleet owners tend to be conservative in choosing powertrains to ensure a wider application coverage and will realistically achieve a range of 10-15% efficiency improvement. Much larger benefits are possible switching from Diesel to BEV (up to 380% improvement), but these results will be disregarded in the final numbers for this program as the original aim was to focus on existing fleet powertrain technologies.

Additional freight efficiency benefits are supported by the other fleet efficiency features developed under this program. A 15% freight efficiency improvement was demonstrated using the Intelligent Driver Assistance System (10% from eco-routing and 5% from eco-coaching). When combining new powertrain selection with the IDAS system, a total reduction in Brake-specific energy per freight ton-mile of 25-30% for diesel powertrains was demonstrated.

2. Acknowledgements

This material is based upon work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Advanced Vehicle Technologies Office Award Number DE-EE0009207.

The author of this report greatly acknowledges the support of the Department of Energy and the support of Prasad Gupte (Initial Technology Manager), Melissa Rossi (Technology Manager) and Kimberly Nuhfer (Project Officer) throughout this program. Your continued support greatly contributed to its success in meeting the goals of the Department of Energy and the internal goals for PACCAR.

The author of this report wants to thank the (sub-recipient) partners under this program who contributed to its success: Kenworth Truck Company, Kopius, Esri, Ohio State University, and especially the National Renewable Energy Laboratory (NREL).

3. List of Acronyms

AOI	Area of Interest
API	Application Programming Interface
BEV	Battery Electric Vehicle
CAN	Controller Area Network
CES	Custom Evaluator Script
DOE	Department of Energy
DoE	Design of Experiments
DSF	Design Space Filter
EEMS	Energy Efficient Mobility Systems
EERE	Energy Efficiency and Renewable Energy
ETL	Extract Transform Load
EWRPD	Energy Weighted Relative Percent Difference
FMS	Fleet Management System
FOA	Funding Opportunity Announcement
FRS	Fleet Recommendation System
GCVW	Gross Combination Vehicle Weight
GPS	Global Positioning System
GUI	Graphical User Interface
HVI	Human Vehicle Interaction
IDAS	Intelligent Driver Assistance System
IoT	Internet of Things
ML	Machine Learning
NN	Neural Network
NREL	National Renewable Energy Laboratory
NRMSE	Normalized Root Mean Square Error
O-D	Origin-Destination
OSU	Ohio State University
OTA	Over-the-Air
POC	Proof of Concept
PRS	Powertrain Recommendation System
PTC	PACCAR Technical Center
RF	Random Forest
RMSE	Root Mean Square Error
SDK	Software Development Kit
TRL	Technology Readiness Level
VRP	Vehicle Routing Problem

4. Introduction

In 2020, the U.S. Department of Energy (DOE) Office of Energy Efficiency and Renewable Energy (EERE) Issued a FOA for research projects focusing on efficient transportation of goods. The opportunity to harness the creativity of commercial vehicle OEMs and industry partners to improve transportation system efficiency through better utilization based on fleet connectivity was fully embraced by PACCAR, who put together a team of leading industry and research partners.

The Improving Transportation System Efficiency Through Better Utilization AOI under the FOA required that OEM's cost share must be at least 20% of the total allowable cost. The period of performance of the PACCAR Co-optimization of Vehicle and Routes (CoVaR) to Improve Commercial Transportation System Efficiency program was 39 months with a 6 month no-cost extension authorized later due to initial challenges of finding a fleet partner required to collect real-world data on commercial vehicles.

The FOA states, the objective of the new mobility technologies priority is to leverage advances in connectivity and automation to dramatically improve transportation system-level energy efficiency, energy productivity, and affordability. PACCAR accepted the challenge of developing innovative solutions to address the barriers related to the key levers to improve vehicle utilization and transportation efficiency (productive travel miles per unit of energy), particularly:

- Increased vehicle occupancy (e.g., payload)
- Improved repositioning to reduce deadheading (miles driven with no payload) for taxi, freight, and transportation network companies
- Improved vehicle routing for energy efficiency
- Energy efficient improvements for mobility

Over the period 2020 through 2024, PACCAR succeeded in developing a suite of tools that work together to leverage vehicle telematics, cloud computing, machine learning, and customer data to demonstrate a > 25% reduction in brake-specific energy per freight ton-mile.

These accomplishments created a knowledge base of practical engineering and operational information to assist in developing production energy efficient routing, driver coaching, fleet management, and vehicle optimization tools to meet the energy saving and business needs of the North American freight industry. These successes were in no small part due to the ingenuity and persistence of each of the program partners: PACCAR (prime), Kenworth Truck Company, Kopius, Esri, OSU, and NREL.

The fundamental basis of this transportation system efficiency demonstration is that a global (fleet) system-optimization approach encompassing the entire customer use case of a commercial vehicle as a single performing unit, can yield significantly better performance improvement than more piecemeal approaches looking at optimizing individual systems.

The intent of this report is to highlight the innovative technologies developed and demonstrated by PACCAR, and how they may benefit future production commercial vehicles and fleet operators to support the overall energy efficiency goals from the DOE. The concepts were designed to be applicable to both medium and heavy-duty commercial vehicles with traditional or electric powertrains for maximum benefit and future applicability.

5. Program Schedule & Milestones

The CoVaR program focused on Proof-of-concept (POC) integrated technology demonstrations at regular milestones. The gap in performance between program goals and those demonstrated in the POC will be used to drive (new) requirements for the technologies under development for the length of the program. Figure 1 shows the CoVaR technology development/integration process. All POC evaluations were performed at the PACCAR Technical center by leveraging internal product development trucks before final deployment on fleet operated vehicles.

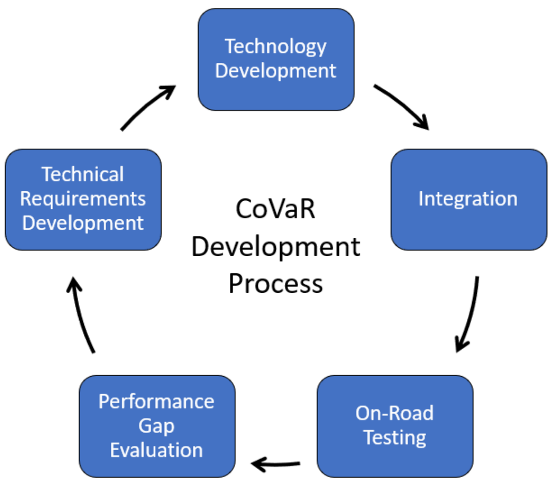

Figure 1: CoVaR technology development/ Integration process

The program was completed over three budget periods as outlined in the following sections.

Budget Period 1: Technology Development

Phase 1 will focus on selecting and evaluating a commercial trucking fleet appropriate to the goals of the program. This phase will then focus on selection, evaluation, and integration of advanced telematics hardware and sensors on fleet and developing a vehicle-to-cloud architecture to cost effectively store, retrieve, analyze and archive data. In this phase, fleet baseline testing equipped with telematics hardware will be completed. In-depth analysis of baseline fleet performance telematics and resource-planning strategy will be performed to develop detailed technical requirements for CoVaR. The potential of vehicle specification optimization, driver assistance, and fleet management to improve freight efficiency will be evaluated using simulation and analysis. Using this analysis, an on-premise machine-learning algorithm will be trained to generate recommendations for fleet management, driver assistance, and vehicle specification optimization. By utilizing a user-centered design approach, an effective UX/UI for delivering eco-driving assistance and fleet management recommendations will be developed. The key Phase 1 deliverable will be a recommendation for a CoVaR strategy that best meets the program goals including vehicle specification optimized for the fleet use case, and on-premise analytics engine to generate recommendations for driver assistance, and fleet management.

Milestone	Type	Description
Select and On-Board Commercial Fleet Partner	Go/No Go	Project team will make decision with about the commercial fleet for baseline data and project demonstration based on selection criteria. Report will be submitted to DOE.
Select Telematics Architecture and Key Vehicle Performance Indicators	Technical	The data transmit rate, consolidation, and on-board pre-processing will be optimized for available data bandwidth. Analysis report including hardware recommendations submitted to DOE.
Acquire Telematics Hardware	Technical	Acquire telematics hardware, complete R&D fleet durability testing, and fit fleet trucks with hardware and transmitting data for use during active commercial delivery runs. A telematics upgrade report will be submitted to the DOE.
Assess the Impact of the Intelligent Driver Assistance System on Freight Efficiency	Technical	Determine if the Intelligent Driver Assistance System for eco-driving can meet the 5% or 10% freight efficiency improvement target. An analysis report, including simulation results, will be submitted to DOE.
Assess the Impact of the Fleet Management System on Freight Efficiency	Technical	Determine if the Fleet Management System can meet the 5% freight efficiency improvement target by reducing deadheading and increasing utilization. An analysis report, including simulation results, will be submitted to DOE.
Develop UX/UI, Cloud Architecture, and Analytics Prototype for Intelligent Driver Assistance System	Technical	The team will develop a prototype for the UX/UI, cloud architecture, and analytics prototype for the Intelligent Driver Assistance System. A summary of design review results will be submitted to DOE.
Develop UX/UI, Cloud Architecture, and Analytics Prototype for Fleet Management System	Technical	The team will develop a prototype for the UX/UI, cloud architecture, and analytics prototype for the Fleet Management System. A summary of design review results will be submitted to DOE.
Assess the Impact of the Vehicle Configuration Optimization System on Freight Efficiency.	Technical	Determine if the Vehicle Configuration Optimization System can meet the 8% freight efficiency improvement target. An analysis report, including Pareto optimal solutions, will be submitted to DOE.

Complete Baseline Testing on Commercial Fleet	Go/No Go	Complete baseline testing and analysis. Simulate CoVaR technology on a commercial fleet to determine if a 25% freight efficiency improvement over the baseline is achievable and manufacturable at scale with a 3 year payback period. An analysis report, including a detailed energy breakdown, will be submitted to DOE.

Table 1: Budget Period 1 Milestones

Budget Period 2: Technology Implementation

Phase 2 will focus on implementing the technologies on the fleet while mitigating real-world risk. An alpha-release of the on-premise ML models for driver assistance and fleet management will be deployed to the cloud and recommendations will be delivered to the Program team internal fleet managers, and R&D fleet drivers using web interface and in-cab dash/voice notifications respectively. Feedback from the fleet managers, drivers, and continued analysis of data from the fleet will be used to iterate on the design and back-end ML models for a beta-release to demonstration fleet managers and drivers. This phase will also include compilation of demonstration fleet specific optimized vehicle specification recommendation and a workshop with the fleet partner to build a business case for fleet upgrade investment. The key Phase 2 deliverable of this phase is the successful deployment of prototype of each technology to the demonstration fleet.

Milestone	Type	Description
Deploy Prototype Intelligent Driver Assistance System on Project Team Internal Fleet	Technical	The prototype of the Intelligent Driver Assistance System will be deployed on the project team's internal fleet and driver feedback will be collected. A deployment and risk mitigation report will be submitted to the DOE. This process will be iterated with a second prototype and a second deployment report will be submitted to the DOE.
Deploy Prototype of the Fleet Management System at Fleet Location	Technical	A prototype of the Intelligent Driver Assistance System will be deployed at the internal fleet and fleet manager feedback will be collected. A deployment will be submitted to the DOE.
Host Workshop with Fleet Partner to Review the Vehicle Specification Optimization Recommendation	Technical	Host a workshop for the partners where the project team will present the methodology and optimization of the fleet vehicle upgrade investment. The results of the workshop will be reported to DOE.

Complete Testing of Prototype Technology, Cloud Architecture, and Code Database Design	Go/No Go	Using a simulation, demonstrate that the CoVaR strategy will cost-effectively meet the required freight efficiency improvement over baseline during active commercial transportation.

Table 2: Budget Period 2 Milestones

Budget Period 3: Testing and Validation

Phase 3 will include a systematic introduction of each technology into the demonstration fleet and the result of the technology intervention will be systematically characterized. The results of the data from the optimized vehicle specification will be used to improve the back-end ML models. This phase of the program will encompass refinement of the ML models for the Fleet Management and Intelligent Driver Assistance systems to co-optimize routing, scheduling, and driving. Once all the technologies have been integrated, the freight efficiency improvement of the combined technologies will demonstrated over several months of operation to account for variations in routes, weather, and drivers.

Milestone	Type	Description
Assess The Impact of the Fleet Management System on Freight Efficiency	Technical	Determine if the Fleet Management System improves freight efficiency by the target 5%. A fleet test and analysis report will be submitted to DOE.
Assess The Impact of the Vehicle Specification Optimization System on Freight Efficiency	Technical	Determine if the Vehicle Specification Optimization System improves freight efficiency by the target 8%. A fleet test and analysis report will be submitted to DOE.
Assess The Impact of the Intelligent Driver Assistance System on Freight Efficiency	Technical	Determine if the Intelligent Driver Assistance System improves freight efficiency by the target 15%. A fleet test and analysis report will be submitted to DOE.
Demonstrate the Impact of The Final Co-Optimized Fleet on Freight Efficiency	Technical	Demonstrate the cumulative impact of the Fleet Management System, Vehicle Specification Optimization System, and Intelligent Driver Assistance System and determine if the co-optimized system improves freight efficiency by the target 25%. A fleet test and analysis report will be submitted to DOE.
Complete Final Report	Technical	Conduct final review with DOE on the Fleet Management System, Intelligent Driver Assistance System, and Vehicle specification Optimization System based on the fleet test results.

Table 3: Budget Period 3 Milestones

6. Technology Definition and Implementation

Advanced Telematics Hardware

PACCAR focused on developing a flexible in-cab, low-cost and modular telematics and edge compute hardware architecture. This technology development should be regarded as the key enabler for the program to have a large-scale fleet telematics deployment with high data resolution and computational performance at low cost. The architecture defined provides the ability to modularize functions which can be addressed using commercial hardware with standard interfaces. Figure 2 shows the simplified architecture draft for the in-cab telematics systems. The architecture consists of standard truck hardware consisting of multiple vehicle Communication Area Networks (CAN) and specific hardware serving the functions of CAN data aggregation, integration of high-resolution GPS data, edge computing device for data processing, transmitting data to cloud, and finally an in-cab display to deliver routing and driver coaching recommendations.

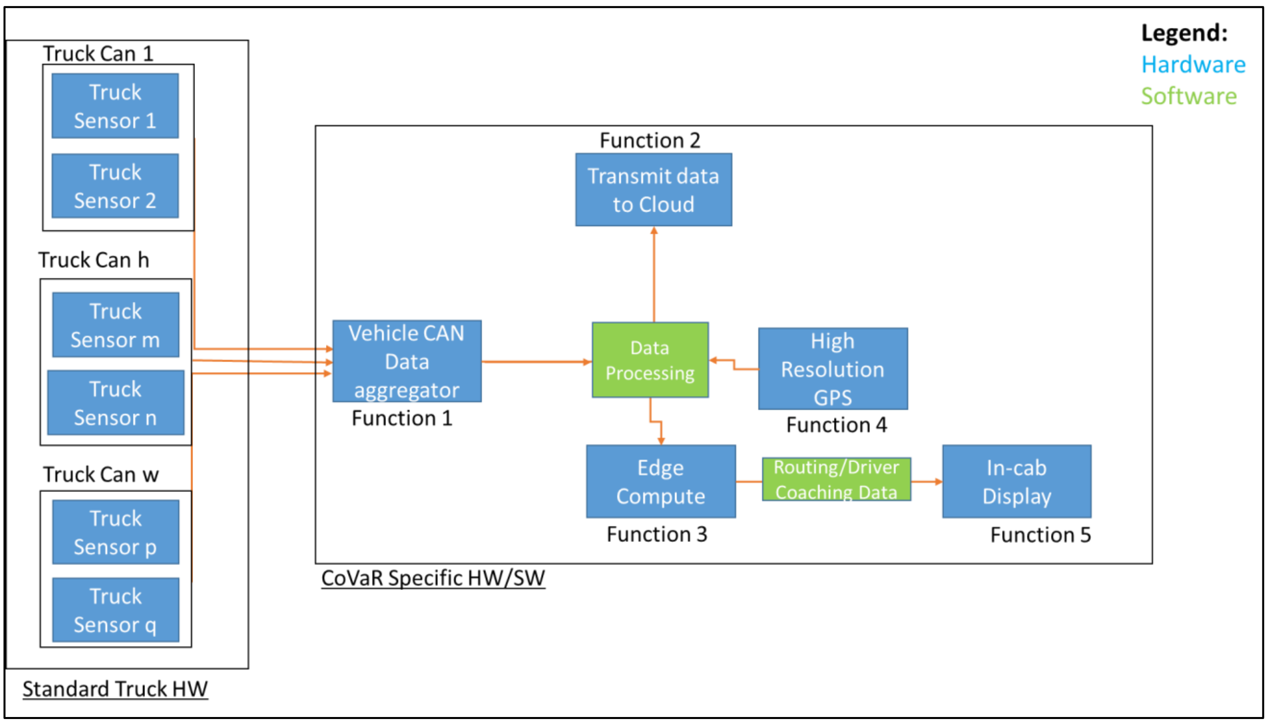

Figure 2: Simplified architecture draft for advanced telematics in-cab hardware

Table 4 shows the various supplier configurations that were compared for the five functions identified in the simplified architecture based on total cost, performance, and scalability. Performance refers to combined capability of the system to log large amounts of on-board signals at high data rates, perform near real-time data processing, and ability to execute advanced machine learning algorithms at the edge. Scalability refers to the ability of a given configuration to be reliably scaled to pilot operations in customer use.

Hardware Configuration option						Metric [1-10] (higher score is better)		
	Vehicle CAN Data Aggregator	Transmit to Cloud	Edge Compute	High Res GPS	In-cab Display	Total cost	Performance*	Scalability**
1	Aptiv EP-250	Aptiv EP-250	Intel Everfocus eVIP 5600 /Nvidia Jetson Nano	Integrated	Galaxy Tab Active2 8.0" 16GB (Wi-Fi)	1	6	10
2	Innomatix	Innomatix	Intel Everfocus eVIP 5600 /Nvidia Jetson Nano	Integrated	Galaxy Tab Active2 8.0" 16GB (Wi-Fi)	3	6	8
3	Vector	Cradlepoint Router	Intel Everfocus eVIP 5600 /Nvidia Jetson Nano	Integrated	Galaxy Tab Active2 8.0" 16GB (Wi-Fi)	1	10	6
4	Kvaser	Intel/Nvidia Jetson Nano	Intel Everfocus eVIP 5600 /Nvidia Jetson Nano	Integrated	Galaxy Tab Active2 8.0" 16GB (Wi-Fi)	5	10	5
5	New Eagle GCM48	Intel Everfocus eVIP 5600	Intel Everfocus eVIP 5600 /Nvidia Jetson Nano	Integrated	Galaxy Tab Active2 8.0" 16GB (Wi-Fi)	7	10	4
6	Axiomatix router	Audesse	Intel Everfocus eVIP 5600 /Nvidia Jetson Nano	Integrated	Galaxy Tab Active2 8.0" 16GB (Wi-Fi)	5	10	6
7	Raspberry Pi w PiCAN2	Raspberry Pi/Everfocus eVIP 5600 /Nvidia Jetson Nano	Intel Everfocus eVIP 5600 /Nvidia Jetson Nano	Integrated	Galaxy Tab Active2 8.0" 16GB (Wi-Fi)	10	6	1
*Performance refers to combined capability of the system to data-log large number of on-board signals at high rates, perform real-time data processing, and ability to execute advanced machine learning algorithms at the edge. **Scalability refers to the ability of a given configuration to be reliably scaled to pilot operations in customer use.								

Table 4: Comparison of various advanced telematics configurations

Based on the hardware configuration option metrics, option 7 utilizing a Raspberry Pi as the CAN data aggregator was selected due to its low cost, computing power, excellent documentation, and large community software support. A POC prototype system utilizing the Raspberry Pi was developed and successfully tested on a vehicle at PTC, proving the low-cost solution would meet the program's needs. A JSON text-based data output structure was selected for the telematics system due to its ease of parsing and wide software support on most systems.

The POC prototype was critical for ensuring the Raspberry Pi and general architecture was capable of recording data from the truck and broadcasting it to the cloud. The next step in the process was refining the system to reduce cost and complexity with additional prototyping and testing at PTC. Hardware changes were made by removing the Intel Everfocus eVIP 5600 vehicle computer due to high cost and lack of technical benefit. Testing determined the Raspberry Pi could aggregate the CAN data as well as performing the data processing functionality for the program needs. A Samsung S7 tablet was selected as the in-cab display and cellular modem for further system simplification.

Other areas of refinement to the system were related to cloud storage and software. The initial POC prototype broadcast data to a Microsoft Azure Event Hub Cloud architecture, however, Microsoft Azure IoT Central architecture was selected as the path forward due to its enhanced security, device provisioning services, and improved data extraction capabilities which aligned with program goals.

Software development for the advanced telematics device was particularly challenging. The software went through many iterations related to Android functionality changes with new releases and the requirement to update to the new releases to comply with PACCAR IT security protocols. As part of the PACCAR IT security protocol the Samsung tablets were managed with Microsoft InTune to enable OTA updates, “kiosk mode” to prevent drivers from using the tablets incorrectly, encrypted telemetry data transfer and provide intrusion protection. Through InTune the OTA updates were executed from the Google Play store which made remote updates possible without engineering tools.

The Java Android application used to interface with the Raspberry Pi went through a few major iterations due to security updates and Android functionality changes. Initially all cloud telemetry interface software was developed and deployed on the Raspberry Pi while the tablet was intended to be used as a simple pass-through mechanism to send telemetry received via Wi-Fi to the cloud via cellular. Due to restrictions in the Android OS the use of Wi-Fi was not allowed without user interaction, which would distract the driver and was determined an unacceptable risk.

To remove the Wi-Fi connection the tablet was connected to the Raspberry Pi through an Ethernet to USB-C adaptor. Due to this architecture change, the telemetry processing had to be moved from the Raspberry Pi to the Android tablet. This change required porting much of the C++/Linux code into Java/Android code which was a monumental task with many challenges. In many cases the SDK available for Android was not nearly as robust as the packages available for C++ and Linux. Another issue that kept plaguing the development team was the changing of Android sleep mode functionality with new releases which caused the tablets to go to sleep and stop broadcasting data. A GPS puck was added to the system instead of using the one in the Android tablet as initially planned due to the tablet pausing GPS output to conserve power and leading to inaccurate drive cycle information.

Despite these issues, PACCAR was able to successfully develop an advanced telematics solution that was secure and safe to use in the vehicle without causing a driver distraction. The final telematics architecture can be seen in Figure 3.

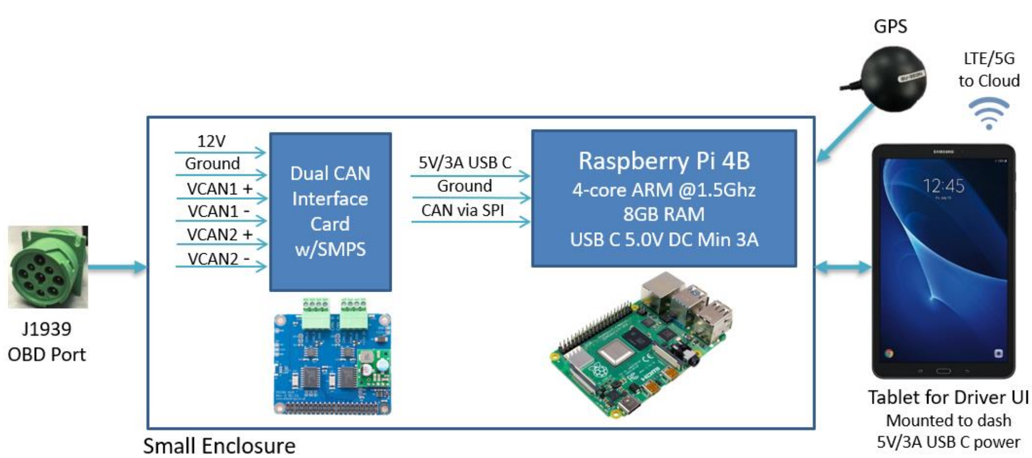

Figure 3: Final architecture for advanced telematics in-cab hardware

Deploying the telematics in the field was more difficult than anticipated. Across the fleets there was a mix of heavy-duty and medium-duty CAN architectures and baud rates. Vehicles with the same architecture had to be tested to resolve the issues we found in the field by including in the functionality auto detection features for baud rate and configuration to know which bus the various CAN signals were on. On older architectures the CAN message from the vehicle related to the ignition switch was different, causing issues with the tablet staying awake to broadcast and receive data for example.

One issue encountered we could not fix through software was the drivers unplugging the device. The telematics unit relied on receiving 12V power through a lighter socket, which was the only power outlet that could meet the current demand of charging the tablet. Using the lighter socket proved to cause issues with drivers removing the adaptor to utilize the socket themselves or unplugging the telematics device to utilize the USB cable for charging their own electronics. For future applications, the hardware and wiring will be behind the dash to prevent tampering.

Figure 4: Picture of CoVaR IDAS system components

Figure 5: IDAS system installed in customer vehicle

Advanced Telematics Cloud Infrastructure

CoVaR partner Esri developed the Advanced Telematics Cloud Infrastructure with input from all major stakeholders to ensure the system would meet all the requirements. The architecture defined for the program is complex, utilizing many different databases and tools combined to determine the most optimal route and output the necessary information to the IDAS and FMS systems. Figure 6 shows the simplified draft cloud architecture diagram created in a separate Microsoft Azure tenant and subscription created for this program to allow for continued rapid prototyping and implementation with stable, well documented, and supported cloud capabilities. Microsoft Azure infrastructure was selected as the infrastructure and services allow for rapid implementation of the architecture by commercial fleets, vehicle manufacturers and OEMs in the future.

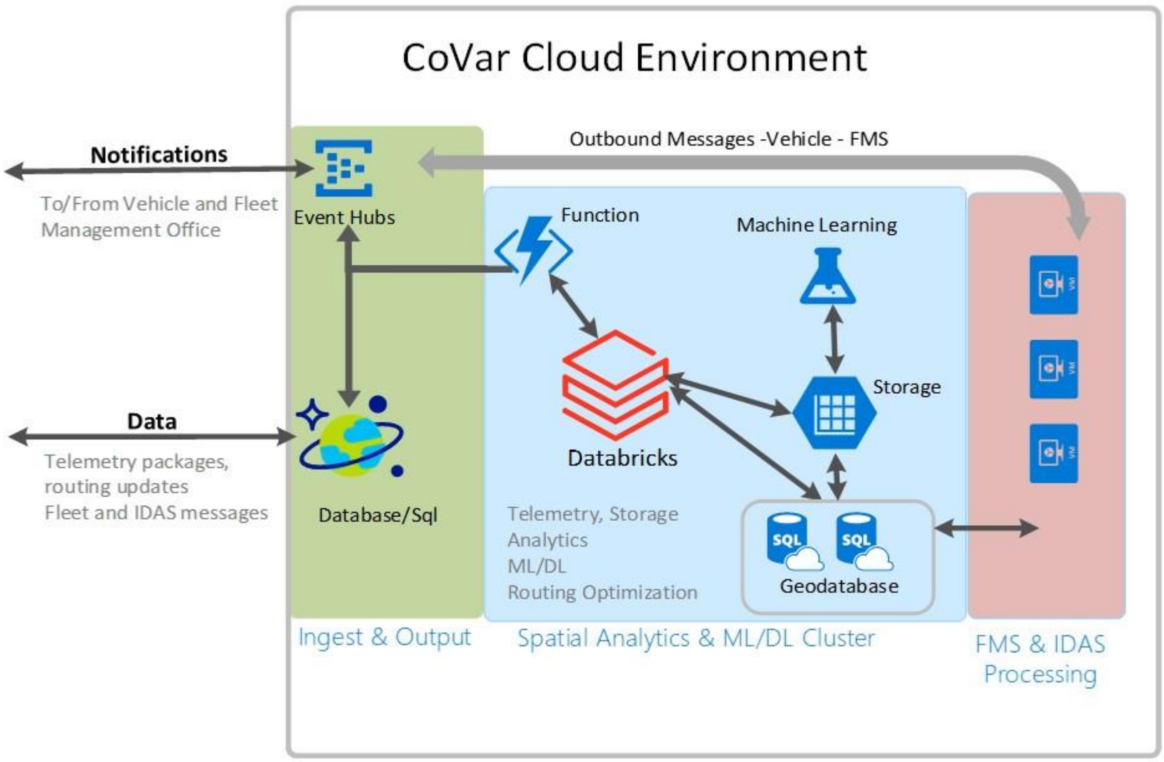

Figure 6: Simplified draft architecture diagram

After the draft architecture was established, work began to create large batch updates of the Road Network contained within the Geodatabase to add and update cost attributes and weighting to be used by the Route Solver to create routes optimized for vehicle energy consumption in addition to the traditional time and distance tradeoff. Detailed discussions between Esri and NREL began to integrate the Esri's Enterprise Geodatabase Road Network with the estimated road segment energy consumption values from NREL's RouteE-powertrain model web API to enable this input to the Route Solver.

Once the initial POC cloud infrastructure was complete and capable of ingesting data from the vehicle and determining an optimal route utilizing energy consumption as an input, the team focused on refining and optimizing the architecture. A major change to the architecture was moving from Microsoft Azure Event Hub Cloud architecture to Microsoft Azure IoT Central. This change required significant modifications to the ingest of telemetry to IoT Central as well as to the data storage and first tier ETL.

One advantage of Azure IoT Central is the ability to remove Databricks from the workflow and transition the analytics portion of ETL code to Azure Data Factory to streamline the process and remove unnecessary computation and complexity.

Another large aspect of the improvement and refining of the architecture was the efficient integration of the NREL RouteE energy estimation into the Road Network and ArcGIS Pro, which is utilized to visualize, modify, and quality check the Road Network data. One of the critical inputs to the NREL RouteE model is road grade values, which required enhancing the accuracy of road segment grades through TomTom to the whole United States road network. In addition to more accurate road gradients, historical vehicle speed profiles and commercial vehicle restrictions were added to the road network as additional variables taken into consideration to avoid traffic delays and dangerous situations.

Enhancements were also continually made to the route solver to incorporate the energy estimation into the optimal routing decision. The route solver was modified to add a weighting value to the time, distance, and energy cost variables so fleets could custom tune the optimization to focus on the aspects that are most important to that fleet knowing a one size fits all approach would not cover all applications and scenarios. An ArcGIS Pro add-in was developed to acquire the energy cost from the NREL web API and integrate it into the road network, so the information was made available as an input to the route solver. A Python script was developed to edit the road segment attributes in bulk to apply the energy cost estimation to the road network.

To speed up the application for use within the vehicle in real time, some significant changes were made to the working prototype. The Custom Evaluator Script (CES) moved the route definition and constraints from within the routing service to customizable code in a portable format to routing services used by the IDAS telematics device. A major change was removing the NREL web API to get predicted energy consumption per road segment, which did not work with the historical vehicle speed input and was too slow. The web API was replaced with a lookup data structure that was made from flattening the ML model used for energy prediction. This change was essential to deliver predicted energy to the routing algorithm at runtime.

To maximize efficiency gain from the intelligent routing service based on our fleet partners use case, a single vehicle-multiple stops (SVMS) routing solution was created. This routing solution is beneficial for situations where a single vehicle makes multiple stops in a route and needs to determine which order to make the stops in based on time of day, vehicle mass, etc. Based on fleet partner feedback, this functionality would fit their use case as each truck would make multiple stops per day at various clients.

Providing trip analytics in the FMS required implementing Azure Synapse and the ArcGIS GeoAnalytics Engine into the ETL. This is where Actual Route Driven geometry is created and injected into the Geodatabase used for calculating trip analytics compared to the recommended routing. The SqlServer instance used by Kopius to host routing and telemetry information for the FMS was expanded to hold spatial data related to the route recommended, and the route driven in addition to the vehicle telemetry to make the data accessible for the FMS dashboard in the same location.

Figure 7 shows the final streamlined cloud architecture for the program with brief explanations related to what contribution each component is providing.

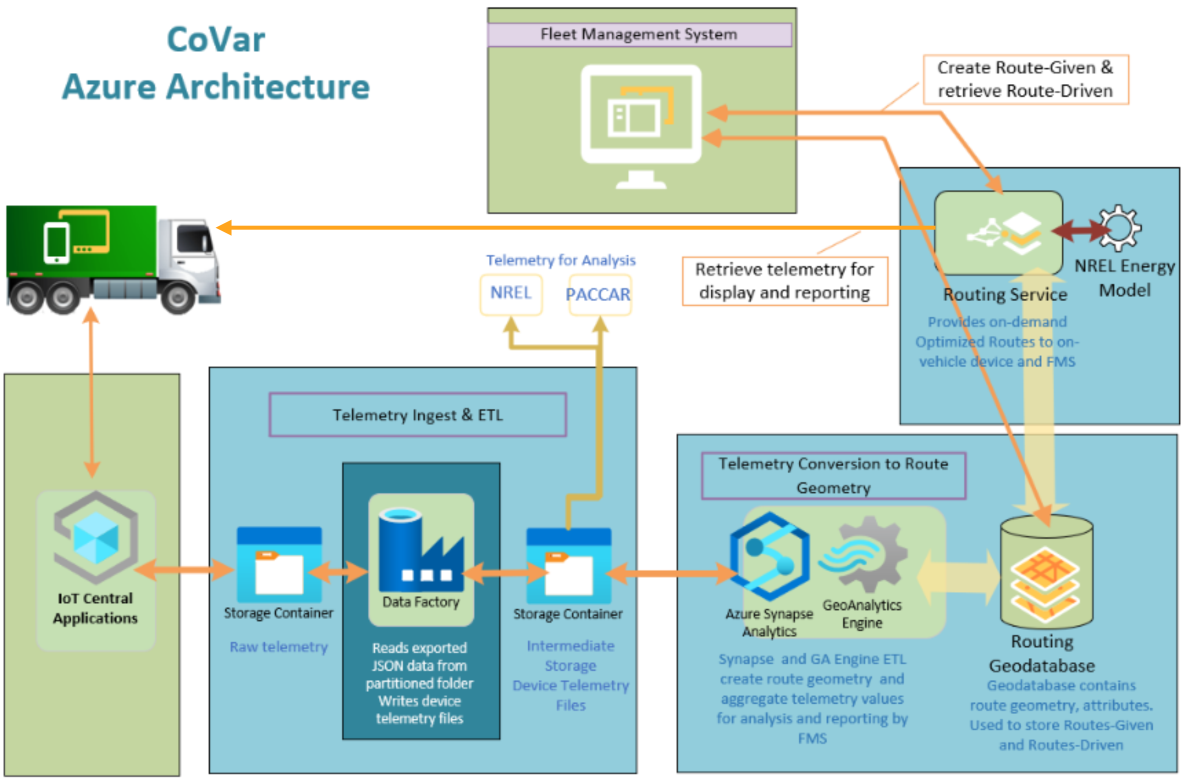

Figure 7: Final cloud architecture diagram

Energy Model Development

NREL's Route Energy Prediction Model (RouteE) tool estimates energy required by vehicles to travel on proposed routes, which was the tool selected to provide energy consumption information to the Route Solver as mentioned previously. To utilize the tool for the CoVaR program, energy models for heavy duty vehicles required development as the existing energy models were for light duty vehicles and were not representative. To develop an energy estimation model for heavy duty vehicles NREL utilized telematics data for six unique trucks from PACCARs development fleet.

Data from the development fleet was collected in two different batches focused on transit and grade drive cycles. The transit batch contained 2,345 miles of driving from 82 trips, and the grade batch contained 3,208 miles of driving from 395 trips. The trips in the datasets are relatively short, like what would be seen in local and regional pick-up and delivery class 8 applications. Histograms of trip distance for both transit and grade batches are shown in Figure 8.

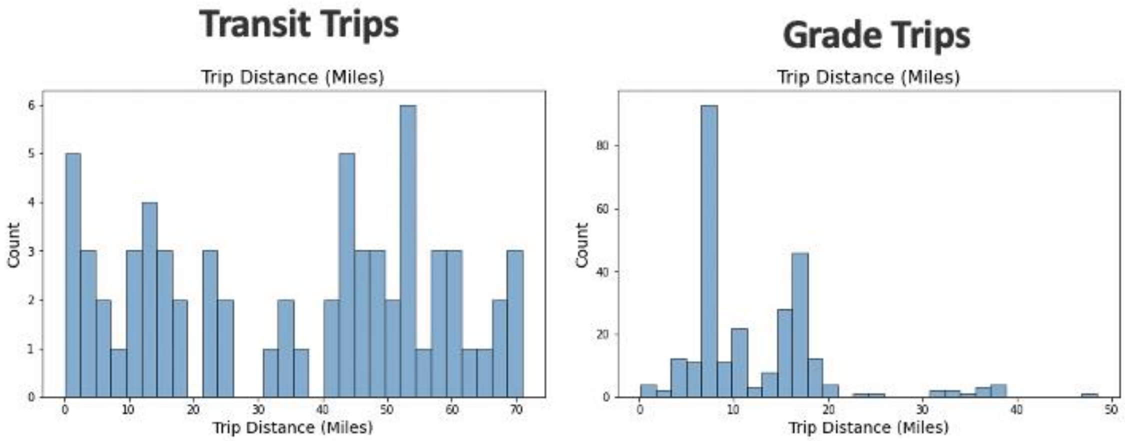

Figure 8: Histograms of transit trip (left) and grade trip distance (right).

The telematics data also contained vehicle estimated mass and fuel consumption rate. The distribution of vehicle masses, which is typically a conservative (higher than actual) estimate made by the brake controller, is shown in Figure 9. The dataset contains vehicles that fluctuate roughly between 28,000 and 38,000 kilograms with a handful of outlying trips. The inclusion of the mass dimension is critical to ultimately inform the RouteE-powertrain models of the impact of gross vehicle weight on fuel consumption rate. Leveraging the vehicle CAN signals to obtain the vehicle mass enables optimization of routes without weighing the vehicle multiple times throughout the delivery route.

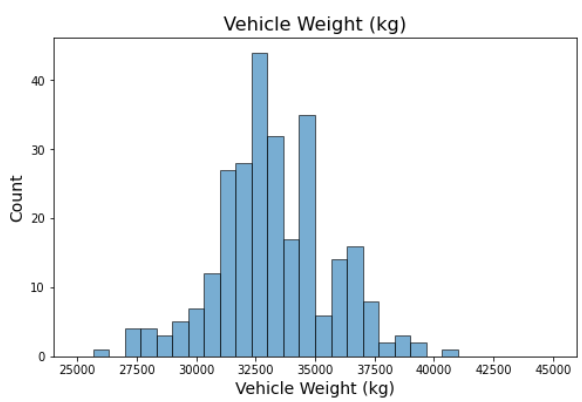

Figure 9: Distribution of vehicle weight in dataset

The GPS and telematics data were processed through a data pre-processing pipeline developed for the program outlined in Figure 10. First, the data was filtered and cleansed to ensure valid GPS, speed, and weight signals. The series of tests included checking for signal presence, positive speed values, and non-zero weight signals that did not exceed 45,000 kilograms.

The dataset, originally collected at 10 Hz, was first down sampled to 1 Hz. Next, road grade was derived from a smoothed and filtered (low pass) version of the CAN elevation signal. If no elevation signal is present, NREL's GradeIT software can also be used to provide elevation values. The geospatial points are then mapped to the TomTom Multinet road network using a map-matching algorithm. Finally, the 1 Hz points data was spatially aggregated by road network link. The aggregation included computing average grade and speed per link as well as total distance and fuel used per link. The aggregate link-level driving data is then used as training data for the RouteE-powertrain mesoscopic vehicle energy model.

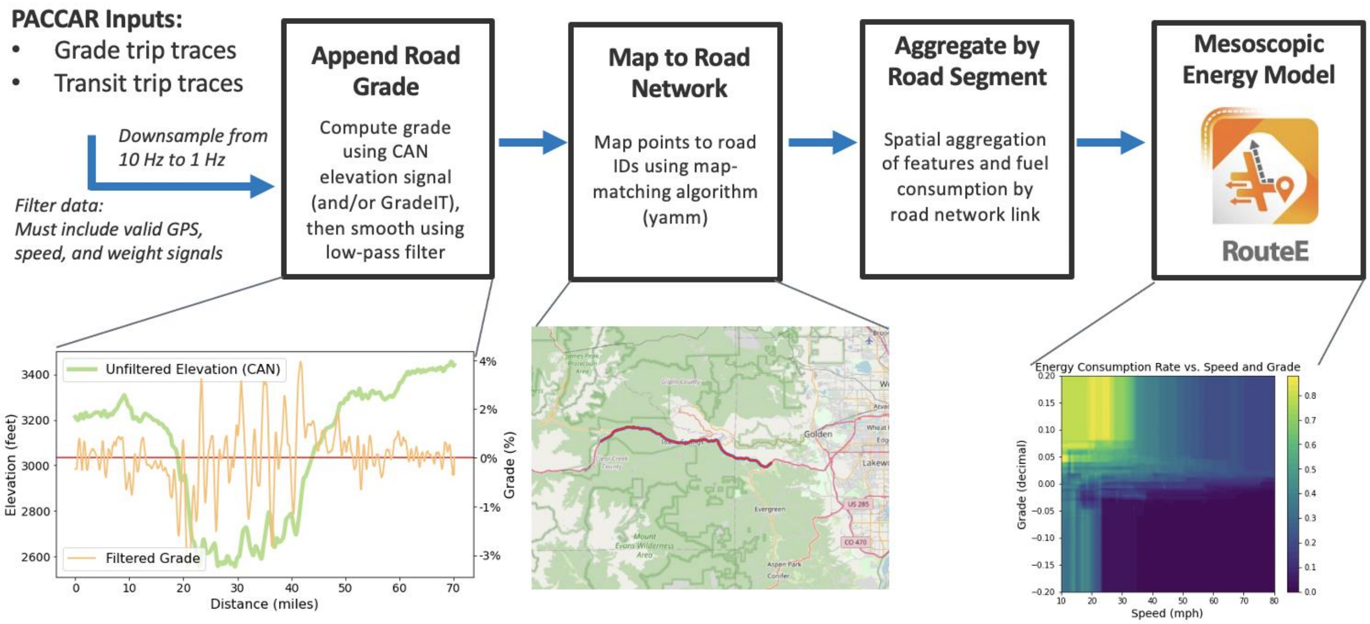

Figure 10: Data pre-processing pipeline workflow

This pipeline leveraged NREL's high performance computing (HPC) resources to enable fast and scalable computation for map-matching, grade appending, link aggregation, and RouteE model training. Accurate map matching is a critical component to ensure the quality of trained RouteE-powertrain models and correct application of the model in real-world deployment. Similarly, reliable elevation and road grade information are also critical and reliant on data processing. An example of the derived and filtered road grade compared to the unfiltered CAN elevation can be seen in Figure 11.

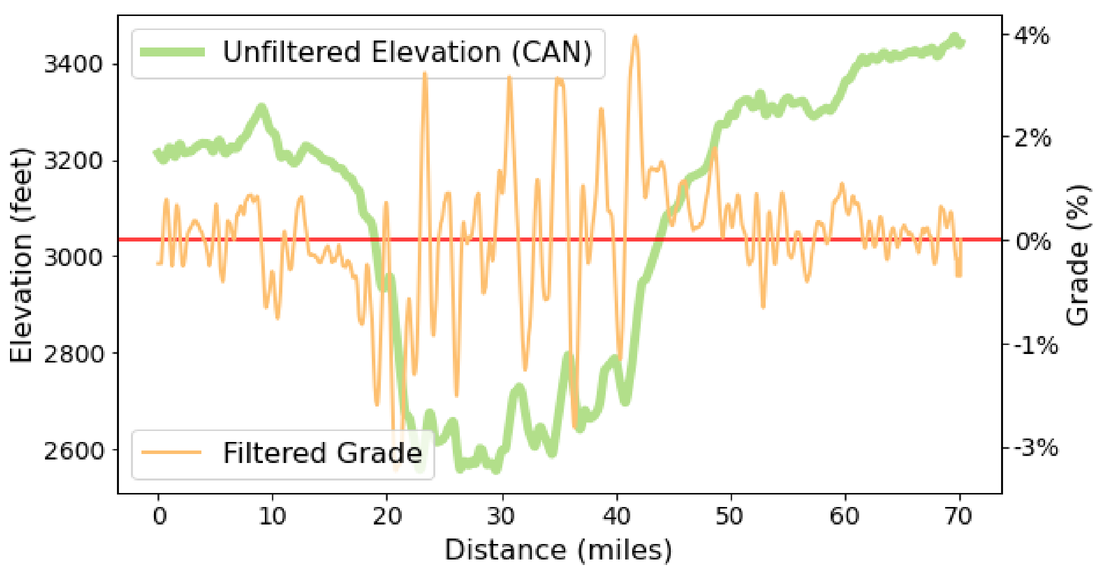

Figure 11: Computing Grade using CAN elevation signal and smoothing using low-pass filter

As mentioned, vehicle mass was (indirectly) provided in the dataset, and this is a critical variable for accurately predicting vehicle energy consumption. Based on the dataset, two separate mass categories we defined for RouteE training, one for trucks that weighed less than the median truck weight (<33,000 kg) and another for those that weighed more than the median truck weight (>33,000).

Aggregate link data that included average speed, road grade, and distance in addition to total fuel consumption over the link was fed as input training data for each vehicle mass category to the RouteE-powertrain software. The result is two RouteE-powertrain models, one for truck trips with a gross vehicle weight less than 33,000 kg and the over for truck trips with a gross vehicle weight greater than 33,000 kg.

A similar process was performed for the BEV trucks as was done for the diesel vehicles. At the time the analysis was performed in-use data from BEV class-8 trucks was not readily available so physics based simulation results from the Ohio State University (OSU) were used to train the RouteE energy models for this segment. Simulations were run on 2,286 miles of driving from 56 trips were used as the dataset.

Based on inputs provided by PACCAR, the simulations were conducted for 3 types of electric motors, 2 types of final gear ratio values, and 3 types of battery sizes, resulting in 18 configurations. The impact of battery size on vehicle weight was not considered as part of this analysis.

For each truck configuration, a RouteE model was trained. For all 18 truck configurations, the resulting trip normalized root mean square error (NRMSE) was between 9.6% to 14.1%, trip energy weighted relative percent difference (EWRPD) was between 7.1% to 10.7%. This result is expected because battery electric trucks can recover energy through regenerative braking. Interestingly, there appears to be a lot of variability in the operating state, switching between depleting and regenerating between -2% and 2% grade when disaggregating only by speed and grade. This is believed to be caused by the longer link lengths on highways combined with switching between depleting and regenerative electric vehicle operating states.

Fleet Baseline Testing and Energy Estimation

Freight Efficiency Metric

One of the initial tasks related to fleet baseline testing was determining the metric that would be used to measure efficiency improvement. It was decided to align with other DOE programs (mainly SuperTruck) and to express the improvement taking into consideration both the useful freight moved and the energy used:

$$Freight\ Energy\ Efficiency = \frac{Distance\ Travelled}{Energy\ Consumed} * Freight\ Mass\ Moved\ \left[\frac{Ton\ Mile}{KWh}\right]$$

Freight mass moved can be calculated using a post processing vehicle model to estimate total vehicle mass or utilizing Gross Combination Vehicle Weight (GCVW) if available on the vehicles CAN network. The remaining parameters in the formula will be derived from the J1939 required CAN network data collected by the vehicles.

Optimal Routing Criteria Selection

Another initial task was determining an approach for choosing the most optimal route. Many implementations of vehicle routing algorithms utilize shortest path algorithms, such as Dijkstra's shortest path algorithm, meant to find the path along a network graph from origin node to destination node that minimizes an objective (accumulation of edge weights). Typically, either driving distance or travel time are the edge weights that are minimized in vehicle routing problems using single objective optimization. This approach works well for passenger car applications but is not suited for commercial vehicles where the vehicle mass changes and significantly impacts energy utilization e.g., as a function of elevations and vehicle speeds across the route. Utilizing NREL's RouteE-powertrain energy consumption integrated into the road network allows for an energy edge weight to be utilized by the route solver as a potential minimization objective. The addition of energy weights to a network makes a multiple objective optimization become desirable to balance energy consumption with another variable. This allows the system to make a more desirable solution for the end user based on factors that are most important to their operation.

Multi-objective optimization is non-trivial to implement and can present performance issues given the added computational complexity. The generalized form of the composite weight will look like the equation below, where *A, B*, and *C* are constant weights and α, β, and γ are exponential factors that allow for a nonlinear combination.

$$W = A[travel\ time]^{\alpha} + B[energy]^{\beta} + C[distance]^{\gamma}$$

The routing tool takes the composite weight as the edge weight of road networks and leverages shortest path algorithms to find the path from origin node to destination node. This equation was implemented in ArcGIS and verified the combined weight metric could be successfully applied to a system originally designed for single objective optimization as shown in Figure 12.

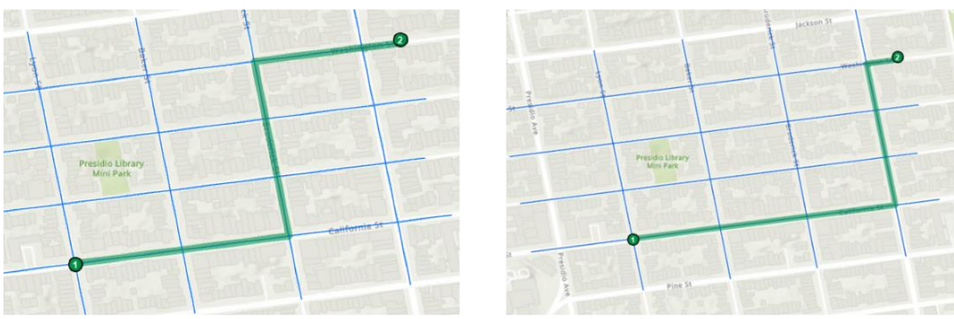

Least energy route Least time route

Figure 12: Routing tool test on prototype road network

Multi-Objective Optimization Validation

Further validation across various scenarios was needed to ensure the routing tool was functioning correctly with the composite weight method before accepting it as the optimal routing solution. To efficiently use the routing tool in Esri's ArcGIS Pro for large-scale routing analysis a Python-based workflow was developed. Trip origin-destination (O-D) information, road network, and parameters in the composite weight function for road network edge cost can be specified in a Python script and fed into ArcGIS Pro via ArcGIS API for Python. Using this Python-based workflow, the repeated tasks of finding optimal routes for hundreds of thousands of O-D pairs with different weight parameter values was automated.

To understand how the parameter values in the composite weight function impact the optimal route solutions and quantify the potential eco-routing energy savings opportunity, extensive numerical experiments were performed. 9450 synthetic O-D pairs are generated with 63 ports in Washington State as origins, 84 car dealerships and 66 Walmart store locations as destinations to approximate typical Class 8 routes. Origins and destinations used in the analysis are shown in Figure 13.

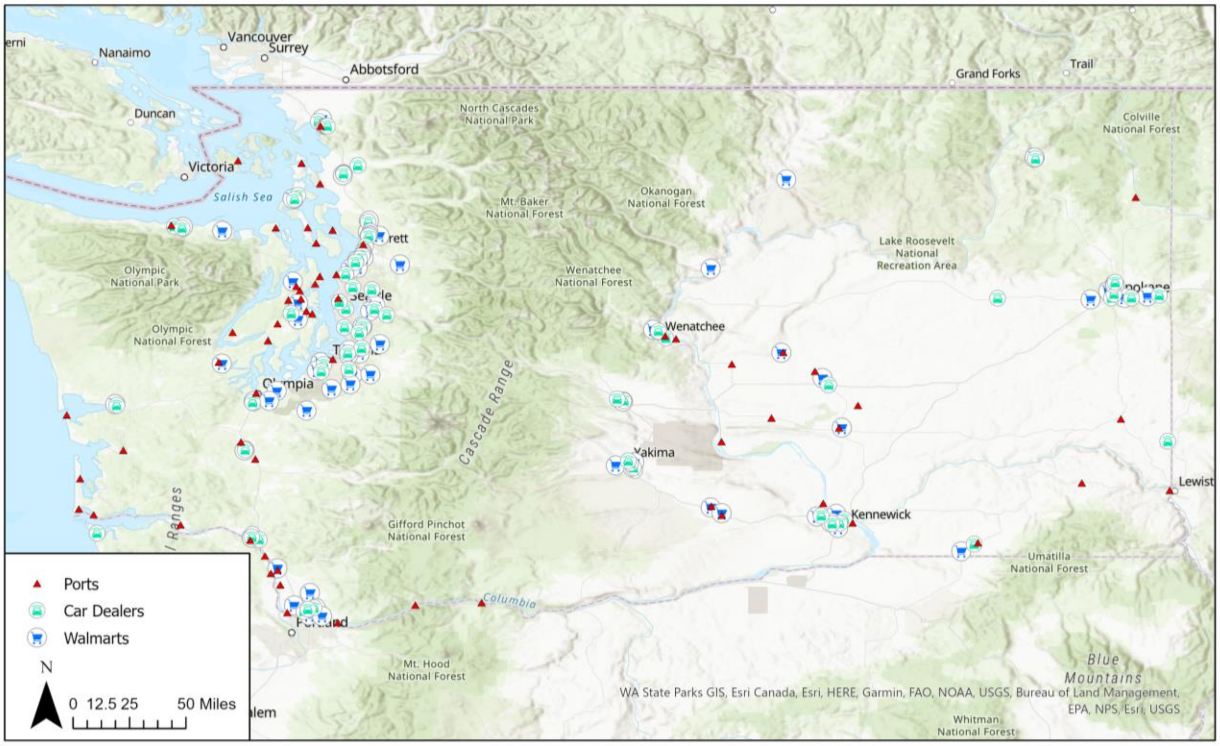

Figure 13: Synthetic O-D pairs for routing analysis

For this analysis the trade-offs between travel time and energy consumption were explored. The analysis determines the least-time and least-energy routes using the ArcGIS routing tool. Figure 14 shows the least-time routing results in the trip time distribution, and least-energy routing results are then used in the energy and distance distributions.

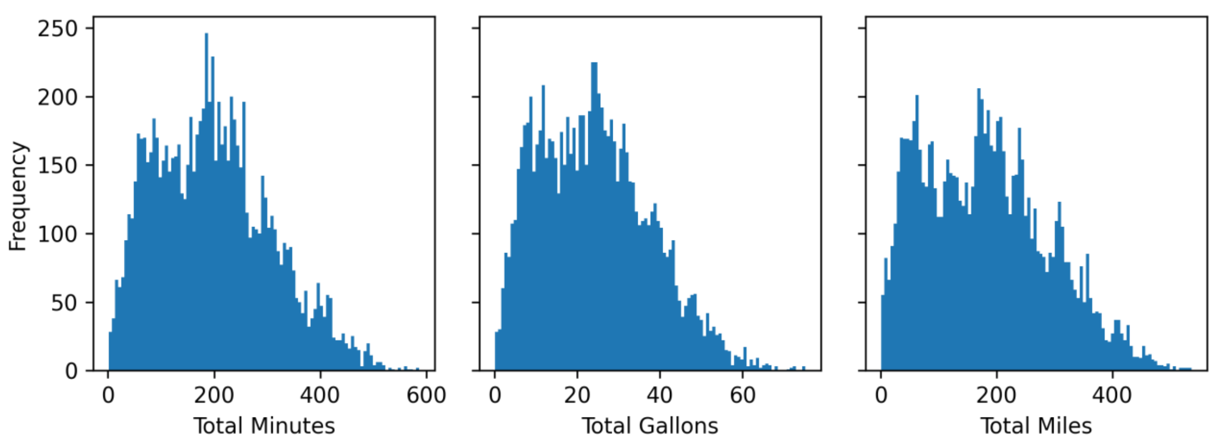

Figure 14: Trip time, energy, and distance distribution

Multi-Objective Optimization Benefit

A least energy route should always have less-than or equal-to energy consumption compared to its corresponding least time route. The heatmap in Figure 15 shows the energy saving potential by trip length, where the "Max Energy Difference" is defined as (energy for least time route – energy for least energy route)/ (energy for least time route). One can observe from Figure 10 that, although most O-D pairs have <5% "Max Energy Difference", many O-D pairs have the potential to save 5% to 35% fuel by taking the least energy routes instead of the least time routes.

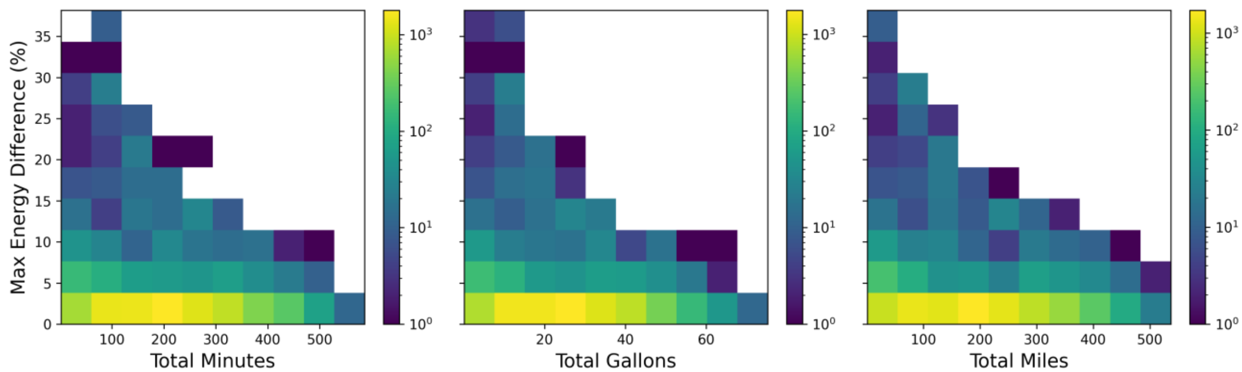

Figure 15: Heatmap showing the energy saving potential by trip length

801 O-D pairs were then extracted whose "Max Energy Difference" is greater than or equal to 5%. For those 801 O-D pairs, we obtain another 99 alternative routes with time weight ranging from 1 to 99 with a step size of 1 and energy weight being (100-time weight). There are usually not 100 distinct alternative route options between one O-D pair (i.e., many weight combinations result in the same optimal route). Figure 16 shows the distribution for the number of optimal alternative routes between each O-D pair as the time and energy weights are varied. One can observe from Figure 16 that, although we have 101 different weight combinations, the number of alternative optimal routes between each O-D pair is less than 13. Note that each O-D pair has at least 2 routes because "Max Energy Difference">=5% ensures that the least time route and the least energy route cannot be identical.

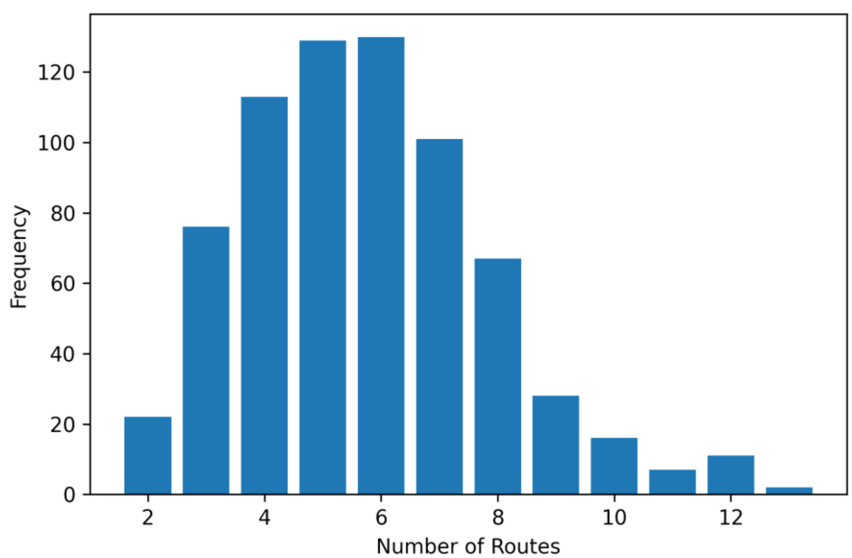

Figure 16 Distribution for the number of alternative optimal routes between each O-D pair with >=5% "Max Energy Difference"

The routing results for one representative O-D pair from "Port of Benton" to a "Car Dealer" is shown in Figure 17. The figure on the left shows three alternative optimal routes, where the "18% Time 82% Energy" route is a time-energy balanced route with the time weight being 18 and energy weight being 82. One can observe from Figure 20 that the balanced route combines the first half of the least time route with the last half of the least energy route to create a time-energy balanced route. Table 5 further reports the trip time, energy, and distance for the three highlighted routes in Figure 17. One can observe from Table 5 and the right figure in Figure 17 that, compared to the "Least Time" route, the "Least Energy" route can save 0.847 gallons of fuel with a 3.762-minute trip time increase, while the "18% Time 82% Energy" route can achieve a comparable 0.719 gallons of fuel saving with only a 0.49-minute trip time increase proving energy savings are possible without dramatic increases in trip time.

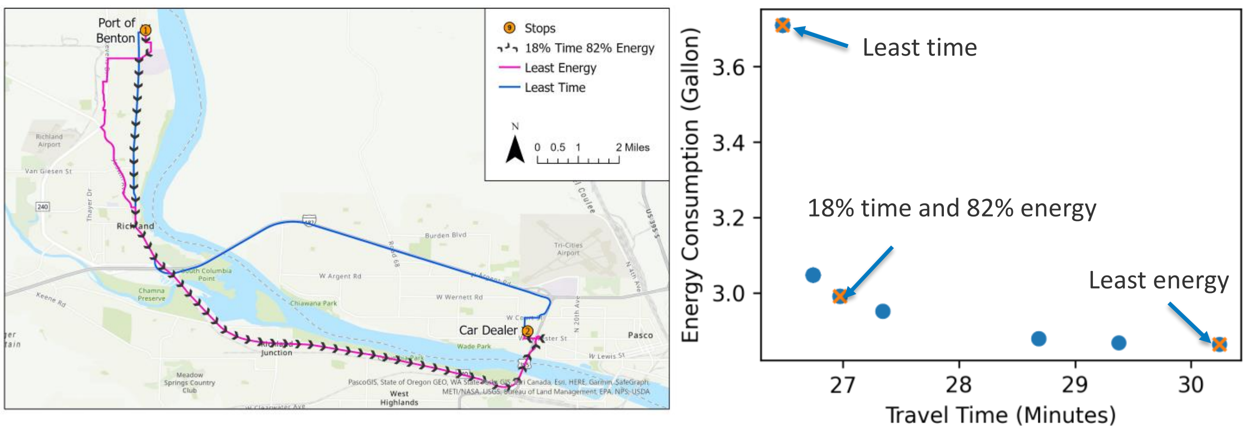

Figure 17: Routing results for one representative O-D pair

Route	Time (Minutes)	Energy (Gallons Diesel)	Distance (Miles)
Least Energy	30.243	2.864	15.821
Least Time	26.481	3.711	14.658
18% Time and 82% Energy	26.971	2.992	14.844

Table 5: Trip time, energy, and distance for one representative O-D pair

The economic feasibility of the routing results was then assessed to confirm there was a business benefit to justify the multi-objective optimization. According to the U.S. DOT 2016 guideline, the average value of travel time for truck drivers is $27.20/hour in 2015 dollars, which is equivalent in purchasing power to $32.36/hour in 2022 when the analysis was performed. If we further consider a $4.00/gallon fuel price, then the one gallon of fuel will be equivalent in monetary value to 7.41 minutes of travel time. With a $4.00/gallon fuel price, we call a route "economically feasible" if its trip time increase is less than 7.41 times the fuel saving compared to the least time route, i.e., the increased trip time monetary value is less than the fuel saving monetary value. For $3.00/gallon and $5.00/gallon fuel prices, the corresponding economic ratio between time and fuel are 5.56 and 9.26 min/gallon, respectively. Figure 18 shows the increased travel time and fuel saving results for all routes of the 801 O-D pairs which had greater than or equal to 5% max energy difference between least time and least energy routes. For a given fuel price, only the route points below the corresponding line are economically feasible. This economic tradeoff method allows for a normalized comparison of time and energy differences between routes to inform selection.

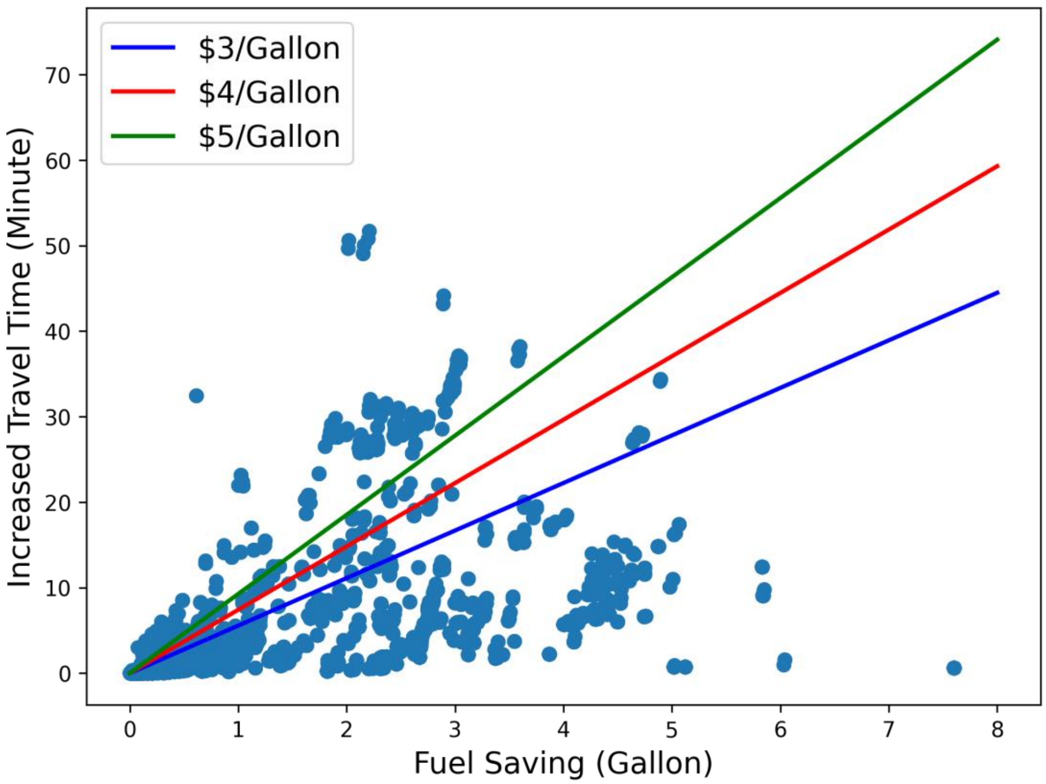

Figure 18: Increased travel time and fuel saving results for all routes

The economic tradeoff analysis helped illustrate the weight values for trip time and energy can have physical meanings. The road network edge cost can be considered as the total monetary cost for using the edge, which might include travel time monetary cost, fuel monetary cost, and distance-based monetary cost. Therefore, the weight values of time, fuel, and distance can be considered as the corresponding prices. Given the value of travel time for truck drivers and the fuel price, the optimal route for each O-D pair with minimum monetary travel cost can be obtained. The value of travel time for truck drivers is set to $32.37/hour and 10 groups of fuel prices from 1.50$/gallon to $6.00/gallon are considered with a step size of $0.50/gallon. With these 10 groups of weight values, the routing algorithm for all 9450 O-D pairs is executed. Figure 19 shows the distribution for the number of optimal alternative routes between each O-D pair. For those O-D pairs with only one route, their least time and least energy routes are identical, meaning there is no room to further reduce the fuel, time, or total monetary cost.

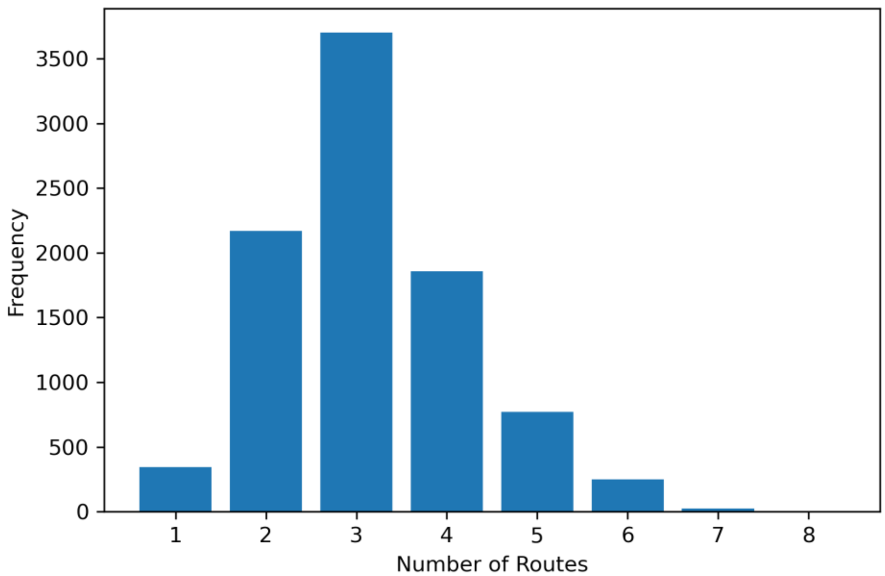

Figure 19: Increased travel time and fuel saving results for all routes

Considering a $4.00/gallon fuel price, we show the routing results for the same O-D pair from “Port of Benton” to a “Car Dealer” in Figure 20 and Table 6, where the [Time*Salary+Fuel*Price] route is the route with the time weight being the value of travel time for truck drivers ($32.37/hour) and energy weight being the fuel price ($4.00/gallon). One can observe from Figure 15 that the [Time*Salary+Fuel*Price] is identical to the “18% Time 82% Energy” route in Figure 20. One can observe from the last column in Table 6 that the [Time*Salary+Fuel*Price] route has the least monetary cost in comparison with the “Least Time” and “Least Energy” routes.

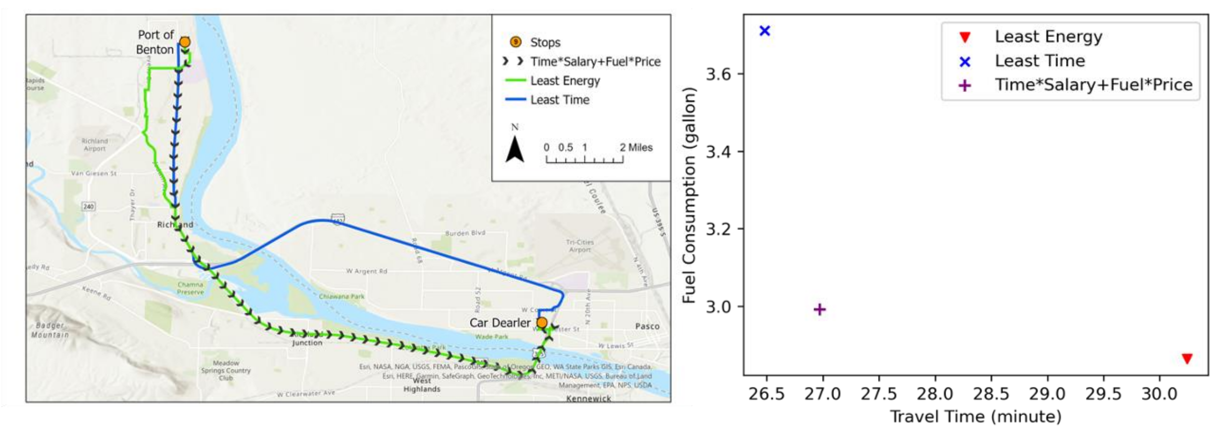

Figure 20: Price-based weighted routing results for one representative O-D pair

Route	Time (Minutes)	Energy (Gallons Diesel)	Distance (Miles)	Monetary Cost ($)
Least Energy	30.24	2.86	15.82	27.79
Least Time	26.48	3.71	14.66	29.14
18% Time and 82% Energy	26.97	2.99	14.84	26.53

Table 6: Trip time, energy, distance, and monetary cost for one representative O-D pair

Compared to the method of enumerating 100 groups of weight value combinations and then evaluating and comparing the routing results, the price-based weighted routing method is simpler and more elegant. If the value of travel time and the fuel price are known, the routing algorithm only needs to run once to obtain the economically optimal routes. Figure 21 compares the total trip costs between the routing results with 100 weights enumeration and those based on the prices of time and fuel under difference fuel prices. One can observe from Figure 21 that the price-based weighted routing methods always find the economically optimal routes with minimum monetary trip costs.

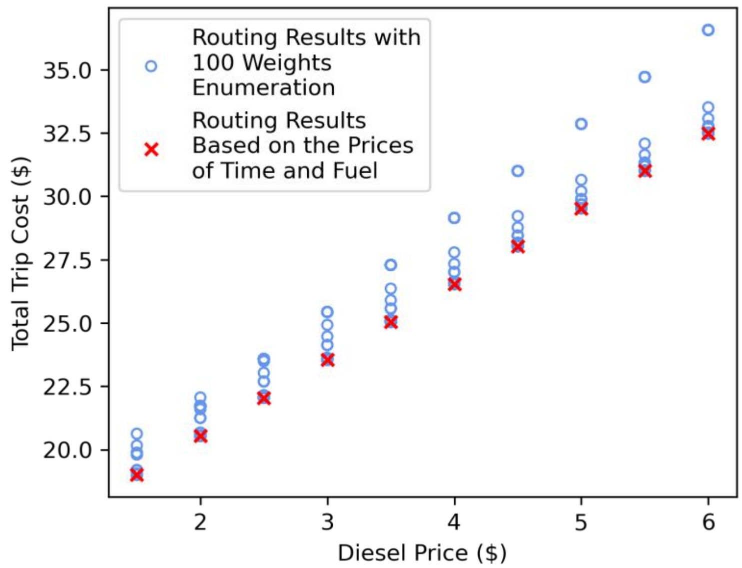

Figure 21: Monetary trip costs comparison between the weight enumeration method and the price-based method

Figure 22 further demonstrates the effectiveness of the price-based weighted routing methods. Each routing solution for a given fuel price is guaranteed to be below the corresponding threshold boundary line and therefore economically feasible.

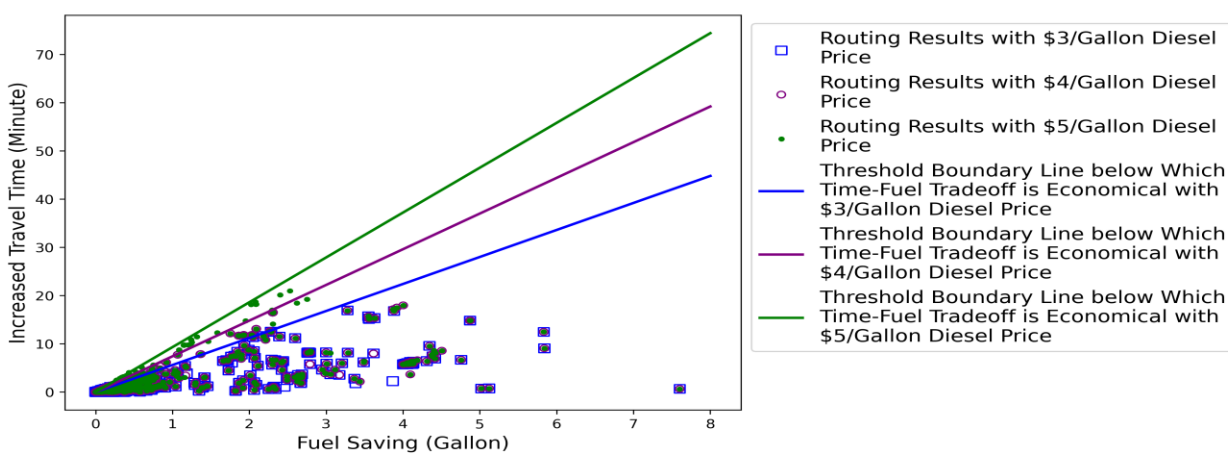

Figure 22: Increased travel time and fuel saving results for price-based routes

The next step is to evaluate the potential of eco-routing in fuel saving and freight efficiency. Based on a fixed value of travel time for truck drivers at $32.37/hour and a fuel price of $4.00/gallon, we have 6539 O-D pairs with positive fuel savings (compared to the least time route). Figure 23 shows the distribution of fuel saving percentage. Among the 6539 O-D pairs, we have 478 O-D pairs with ≥5% fuel savings. Figure 24 further shows the distribution of freight energy efficiency improvement percentage. Among the 6539 O-D pairs, we have 503 O-D pairs with ≥5% freight energy efficiency improvement. Note that the Freight Energy Efficiency is calculated using the formula agreed to for the program, where the “Distance Traveled” in the equation is always set to the distance value for the least time route, serving as a reference distance for all route alternatives. For all least time routes for those 6539 O-D pairs, the total fuel consumption is 168077.2 gallons while the total monetary cost is $1408394.9. For all price-based weighted routes for those 6539 O-D pairs, the total fuel savings are 1905.2 gallons, representing a 1.13% fuel savings, while the total net monetary cost reduction (fuel savings minus additional trip time cost) are $5156.6. Table 7 reports the sensitivity analysis results with fixed value of travel time for truck drivers at $32.37/hour and fuel prices ranging from $1.50/gallon to $6.00/gallon with a step size of $0.5/gallon. One can observe from Table 7 that with the increase of fuel price, all metrics of interest are increasing. This result is expected because with the increase of fuel price, trucks have increasing potential of reducing monetary trip costs by switching from least time routes to price-based economically optimal routes.

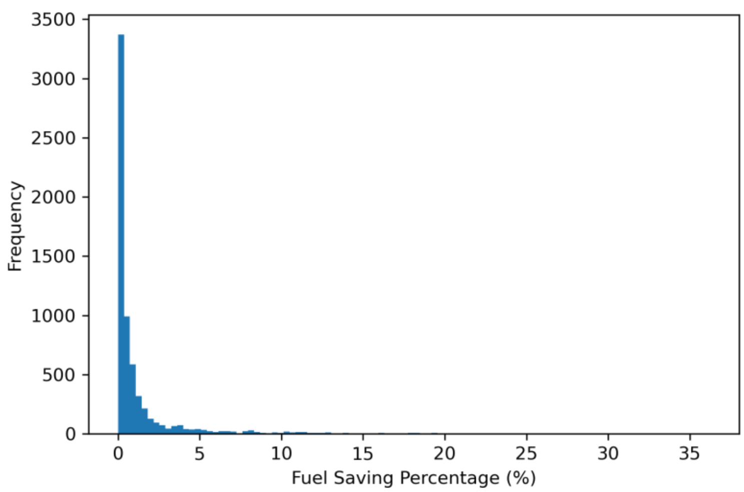

Figure 23: Distribution of the fuel saving percentage

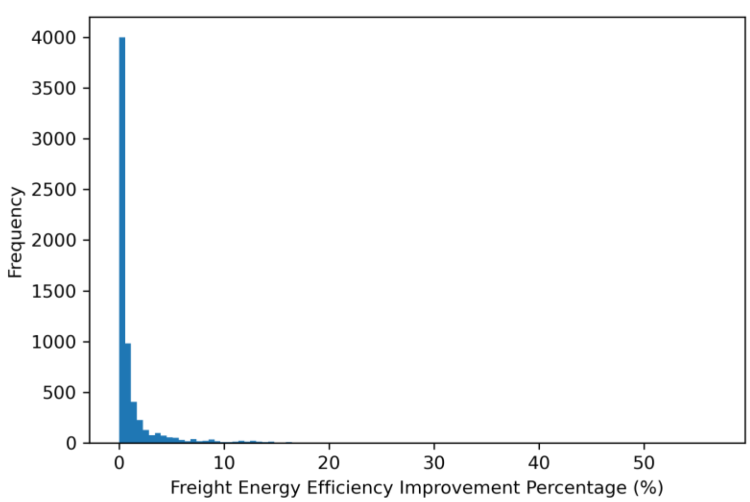

Figure 24: Distribution of the freight energy efficiency improvement percentage

Fuel Prices	Number of O-D pairs with ≥ 5% fuel savings	Number of O-D pairs with ≥ 5% freight efficiency improvement	Total fuel savings (gallons diesel)	Total percentage fuel savings (%)	Total monetary cost reduction ($)
1.5	318	330	1281.75	0.92	1096.02
2.0	349	363	1379.72	0.94	1770.40
2.5	422	447	1625.19	1.05	2503.86
3.0	445	470	1727.94	1.07	3345.82
3.5	462	487	1833.42	1.12	4219.56
4.0	478	503	1905.18	1.13	5156.59
4.5	481	505	1925.99	1.14	6105.81
5.0	493	523	2019.99	1.18	7090.78
5.5	503	533	2050.47	1.19	8105.17
6.0	521	552	2126.75	1.22	9140.91

Table 7: Sensitivity analysis results with different fuel prices

Fleet Level Vehicle Routing

Esri's ArcGIS Pro has a Vehicle Routing Problem (VRP) tool to optimize the delivery schedule for a fleet of vehicles. The original VRP tool was only able to consider the objective of minimizing the total travel time. To deal with multiple criteria in the vehicle routing problem, the composite weight function that considers time, distance, and energy for road network edge cost has also been implemented in the VRP tool. Below we use an example with one depot and five customers (as shown in Figure 25) to demonstrate the effectiveness of the tool. Suppose we have a fleet of delivery trucks with a capacity of 100 units. The pick-up and delivery requirement for each customer is reported in Table 8.

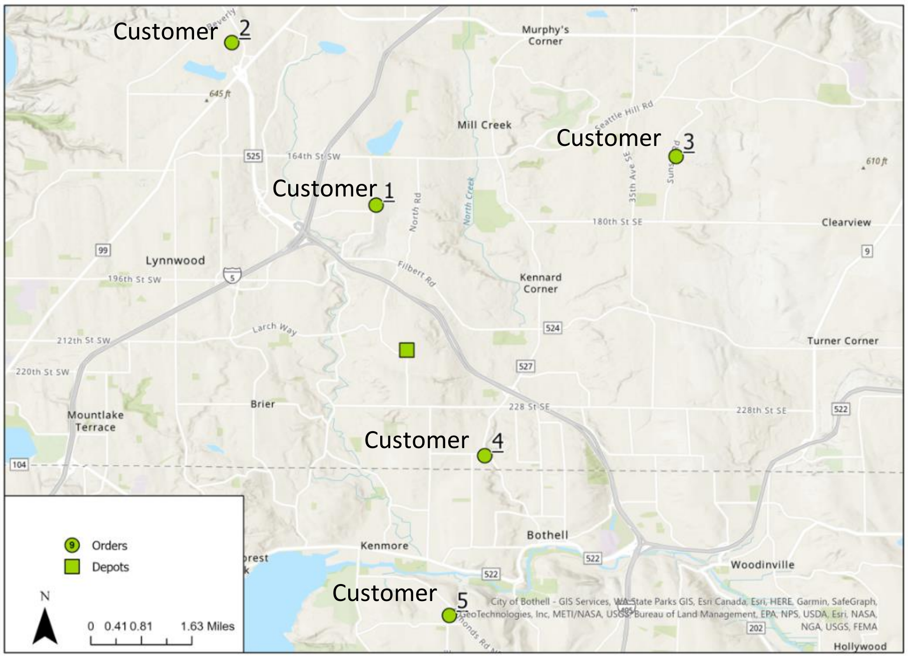

Figure 25: An example vehicle routing problem with one depot and five customers

Customer	Pickup Units	Delivery Units
1	70	0
2	0	70
3	0	30
4	70	0
5	0	100

Table 8: Pick-up and delivery requirement for each customer in the example

Using the VRP tool, we first consider the objective of minimizing the total travel time. Figure 26 shows the VRP results. One can observe from Figure 26 that two trucks are used to serve the 5 customers, with truck 1 serving customers 1 to 3 along the route: depot -> customer 3 -> customer 2 -> customer 1 -> depot and truck 2 serving customers 4 and 5 along the route: depot -> customer 5 -> customer 4 -> depot. Note that at least two trucks are needed for the VRP tool to find a feasible solution for the example VRP problem. Without the VRP tool, fleet operators might schedule the delivery- and pickup-tasks based on trial-and-error experience. Under the worst-case scenario, a fleet operator might use 5 trucks to serve the 5 customers in the example problem, with each truck serving each customer independently.

Another less-extreme scenario is that a fleet operator always tries to serve the nearest customer first. We name this scenario "greedy-strategy-based scenario". Under this scenario, 4 trucks are needed to serve the 5 customers, with truck 1 serving customer 4 (pickup), truck 2 serving customer 1 (pickup), truck 3 serving customer 5 (delivery), truck 4 serving customers 2 and 3 (delivery). Table 9 compares the VRP results from the VRP tool with the worst-case scenario and the greedy-strategy-based scenario. One can observe from Table 9 that, compared to the worst-case scenario and the greedy-strategy-based scenario, the results from the VRP tool can significantly reduce the total and deadheading travel time, fuel use, and distance. Compared to the worst-case scenario and the greedy-strategy-based scenario, the results from the VRP tool can improve the freight energy efficiency by 35.2% and 15.7%, respectively. Note that the "Distance Traveled" in the Freight Energy Efficiency equation is always set to the distance value for the results from the VRP tool and the "Freight Mass Moved" is assumed to be a constant value.

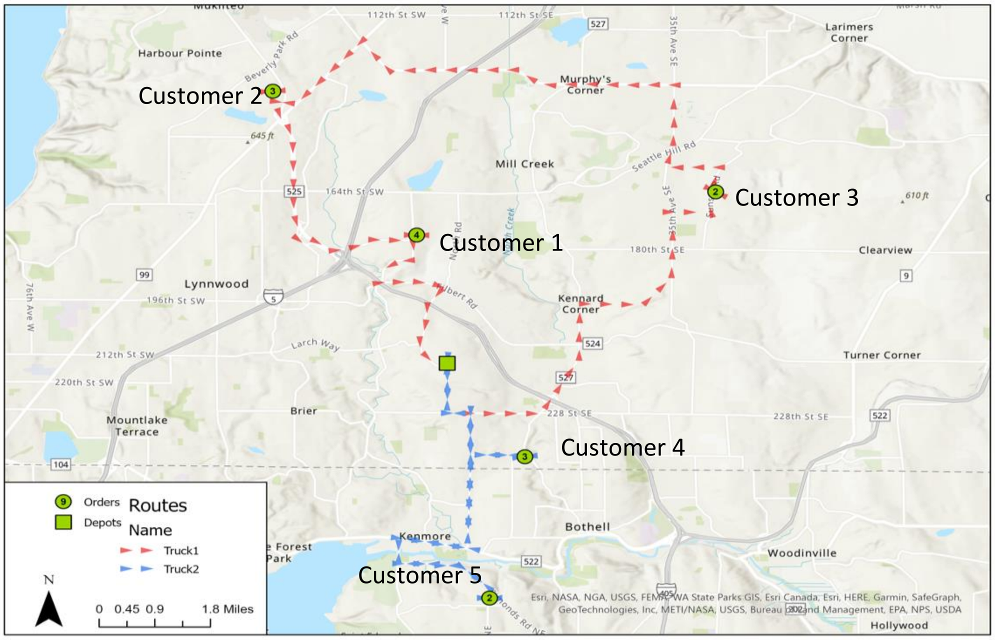

Figure 26: VRP results for the example network with the objective of minimizing the total travel time

Scenario	**Number of Trucks used**	**Total travel time (minute)**	**Deadheading travel time (minute)**	**Total fuel use (gallon)**	**Deadheading fuel use (gallon)**	**Total distance (mile)**	**Deadheading distance (mile)**
Worst-case scenario	5	110.31	54.88	13.94	5.92	56.12	28.80
Greedy-strategy-based scenario	4	98.32	42.04	11.93	4.64	47.88	19.74
VRP tool	2	78.39	19.58	10.31	1.99	37.31	9.36

Table 9: Results from the VRP tool, the worst-case scenario, and the greedy-strategy-based scenario

The VRP problem can also be solved with the objective of minimizing the total fuel consumption or total monetary cost, as shown in Figure 27 and Figure 28. Note that we only consider the monetary cost of time and fuel based on a fixed value of travel time for truck drivers at $32.368/hour and a fuel price of $4/gallon. One can observe from Figure 27-Figure 28 that those three VRP results are different from each other in terms of customer service order and/or routes token between two locations. Table 10 reports the total travel time, total fuel consumption, total travel distance, and the corresponding deadheading values for those three VRP results. As expected, the least time VRP result has the smallest total travel time, the least energy VRP result has the lowest fuel use, and the least monetary cost VRP result has the least total monetary cost.

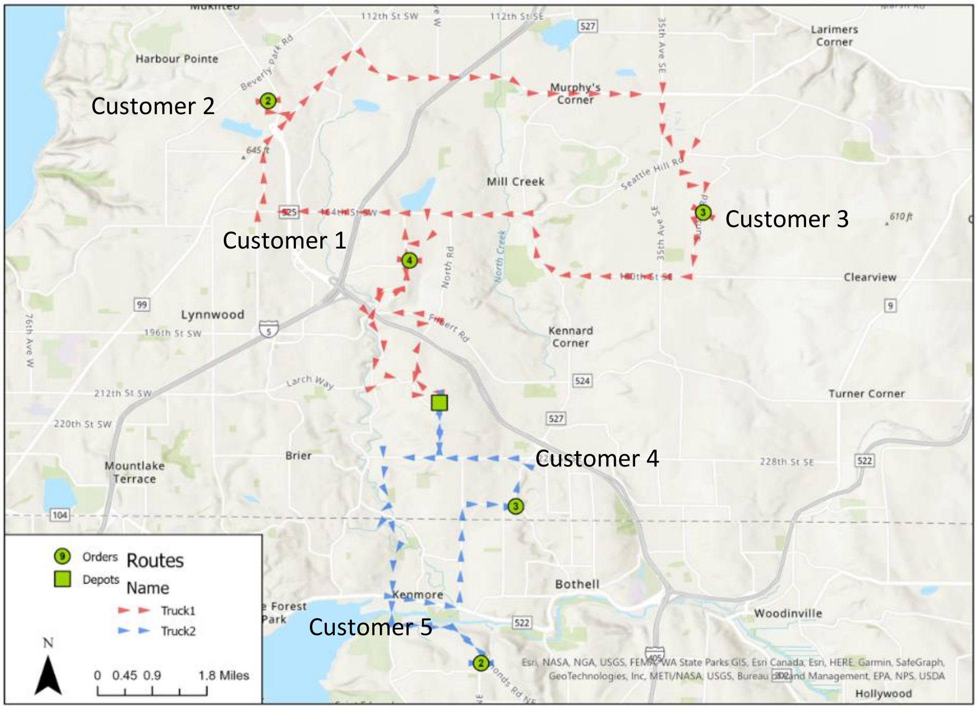

Figure 27: VRP results for the example network with the objective of minimizing the total fuel consumption

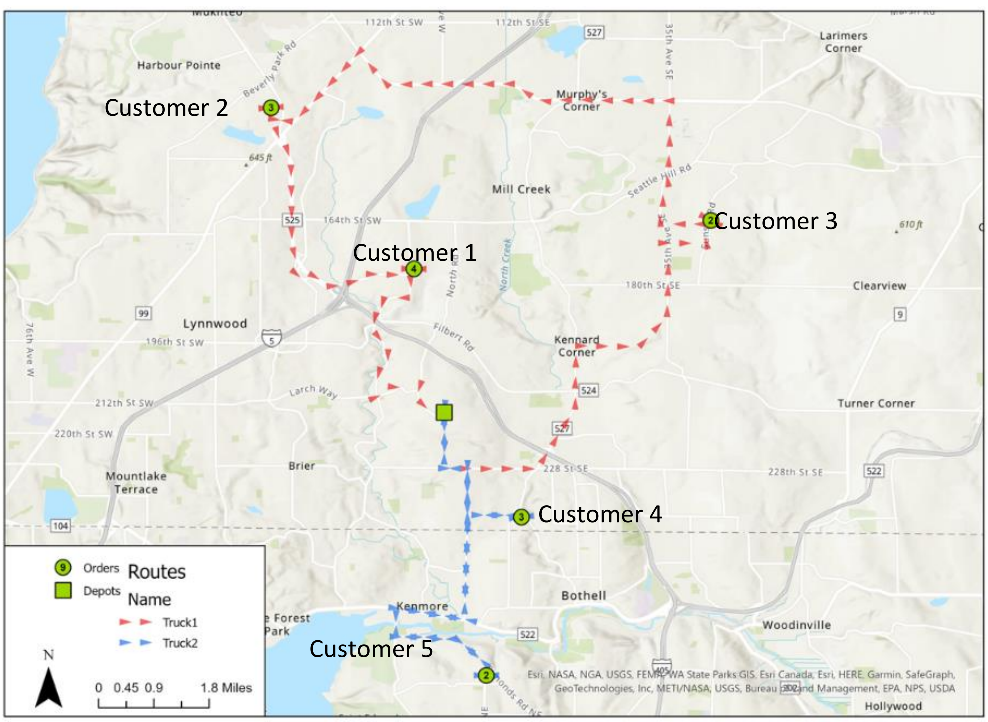

Figure 28: VRP results for the example network with the objective of minimizing the total monetary cost

Objective	Total monetary cost ($)	Total travel time (minute)	Deadheading travel time (minute)	Total fuel use (gallon)	Deadheading fuel use (gallon)	Total distance (mile)	Deadheading distance (mile)
Minimizing the total travel time	83.57	78.39	19.58	10.31	1.99	37.31	9.36
Minimizing the total fuel consumption	81.47	83.68	24.81	9.07	2.33	39.96	8.58
Minimizing the total monetary cost	81.32	80.29	19.58	9.49	1.99	38.52	9.36

Table 10: Detailed VRP results for the example network

This example study showed the potential benefits of utilizing the VRP, however, it was not utilized in the program because it did not benefit the partner fleets commercial operation. The method was developed and could benefit many different types of operations, so it was reported on as an example of the benefits.

Fleet Management System and Intelligent Driver Assistance System Development

Human factors engineering aspects related to the development of the Intelligent Driver Assistance System (IDAS) and Fleet Management System (FMS) was based on an initial group Human Vehicle Interaction (HVI) workshop and published research. Some key findings from the workshop which defined the development direction included drivers resist real time feedback on things like braking or acceleration while driving but are open to post trip summary feedback on how fuel efficiently they drove during a given trip. Additionally, driver prerequisites make accounting for color blindness largely unnecessary, but average driver age necessitates larger, easier to read fonts, icons, and indicators.

Requirements for the IDAS and FMS were refined though an agile process involving wireframe prototypes, driver feedback, program partner feedback, and architecture developments. One major change as part of the process was switching from a native Android application to a browser-based Blazor platform to allow support on more devices than Android and allow easier update deployment. Early wireframe examples are shown in Figure 29.

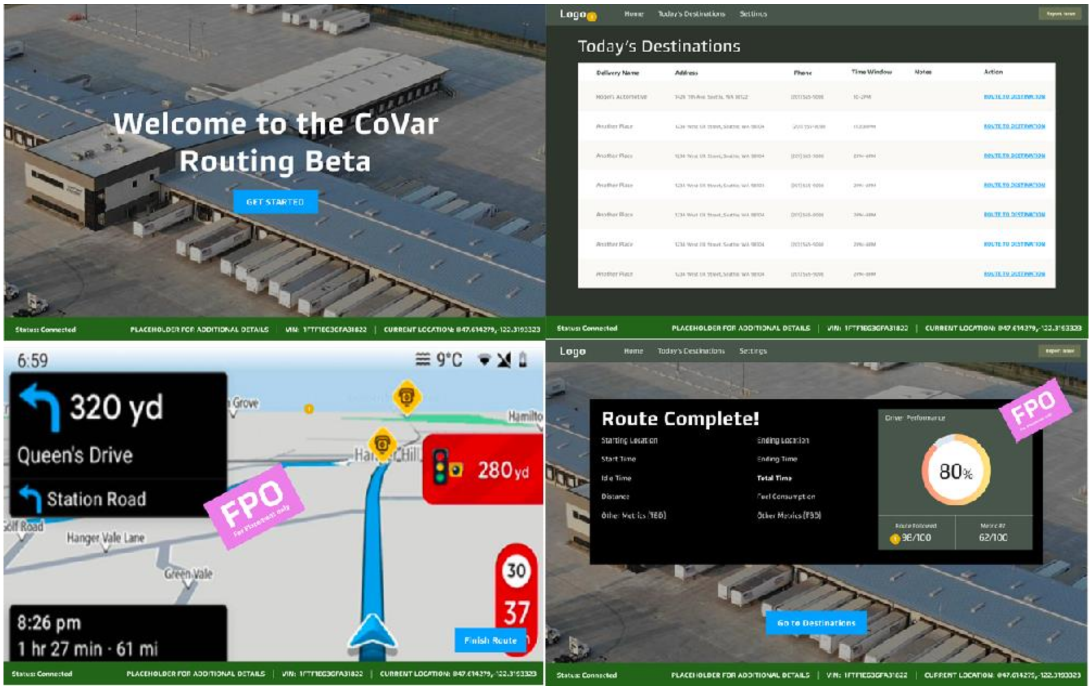

Figure 29: Initial IDAS wireframe concept

Another key change during development was to separate the FMS and IDAS into two applications from the initial single application used for prototyping. The change allowed each component to focus on its intended audience (fleet managers and drivers) by streamlining content to what was needed for the particular use case. Data is shared between the two applications through an API. Code refactoring was completed to move the applications database from Cosmos DB to a SQL database and creating a unified positioning interface within the application. These changes simplified the design to expedite development as well as improve the accuracy and consistency of the application.

To receive data from the fleet vehicles, Kopius and Esri created an Azure database to populate with CAN data from the vehicles and merge it with data fields populated from the routing engine used to populate the FMS dashboard.

To enhance the user experience, iterative changes were made over time based on feedback from fleet partner PacLease and engineers testing the system. One of the initial pieces of user feedback was to remove the need for the driver to log into the tablet application. To remove this feature but keep the connection secure, the IDAS application compared the tablet number from the telemetry data to a list of tablets registered in the database to ensure the device was part of the program. Voice navigation was added to the application so the driver could hear turn-by-turn directions without having to look at the screen to increase safety. Off-route detection was added to reroute with a most optimal route if the vehicle is more than a specified amount from the recommended route.

The final FMS layout is intuitive and easy to navigate as shown in the brief overview. The main dashboard home page displays the trips that are assigned and to which vehicle as shown in Figure 30. The vehicles that can be assigned to trips are entered from the Vehicles tab as shown in Figure 31. All vehicles marked in service will be available to assign to a trip as shown in Figure 32, and the trip will show up on the IDAS in the vehicle. Locations are required to be entered as shown in Figure 33 to be start, middle, or end point to be entered in the trip information.

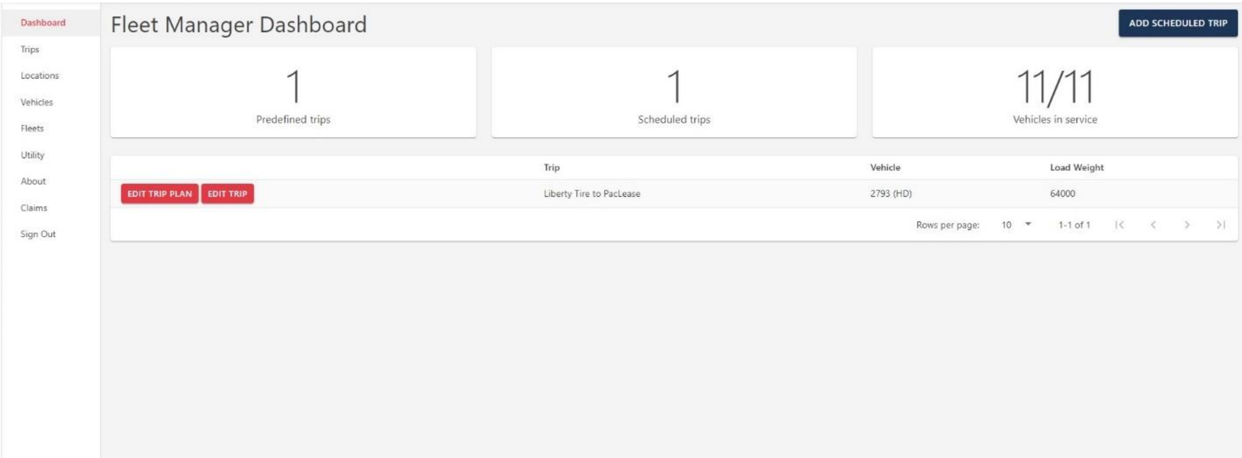

Figure 30: Fleet manager dashboard (home screen)

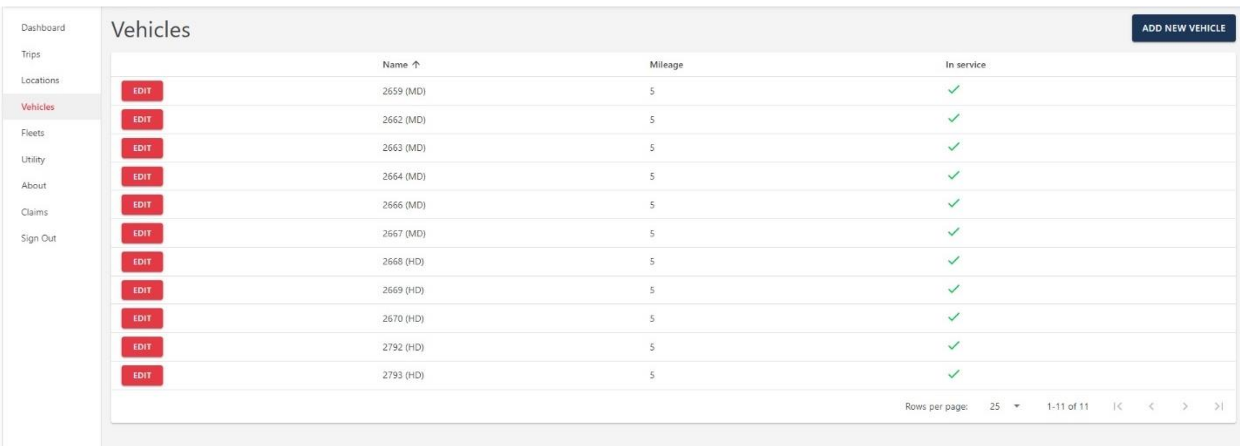

Figure 31: Vehicle management

Figure 32: Assigning a vehicle and an estimated GCVW to a trip

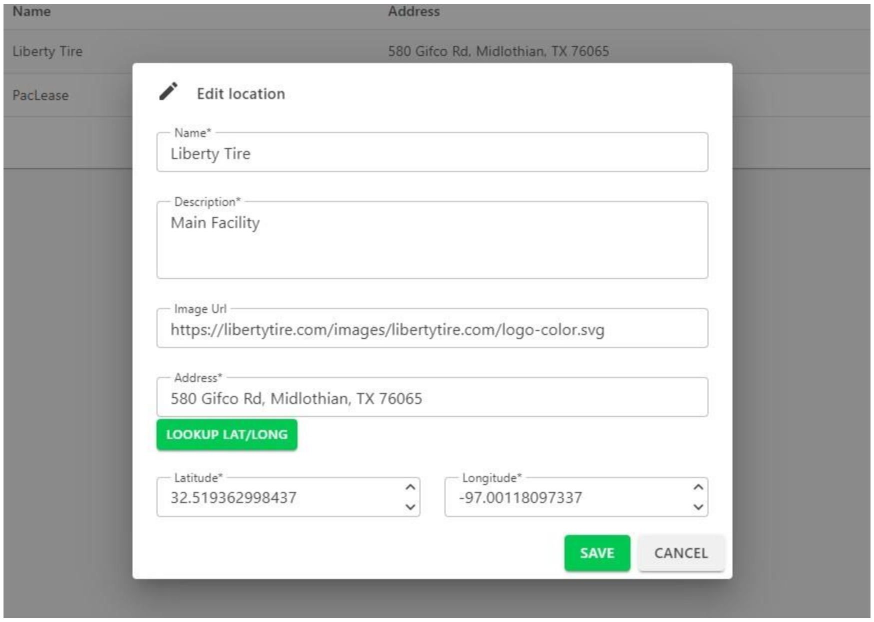

Figure 33: Location information

Vehicle Specification Optimization

Under the CoVaR program, OSU developed a Powertrain Recommendation System (PRS) to recommend vehicle configurations that meet the customer's performance requirements based on data recording of current application and are Pareto optimal (multiple objectives are considered simultaneously). The PRS system performs the following three functions as shown in Figure 34:

1. Find a subset of available powertrain configurations that meet the performance requirements of the customer called 'feasible powertrains'
2. Predict the energy efficiency of the feasible powertrains
3. Provide the trade-off relationship between different powertrains for objectives under consideration

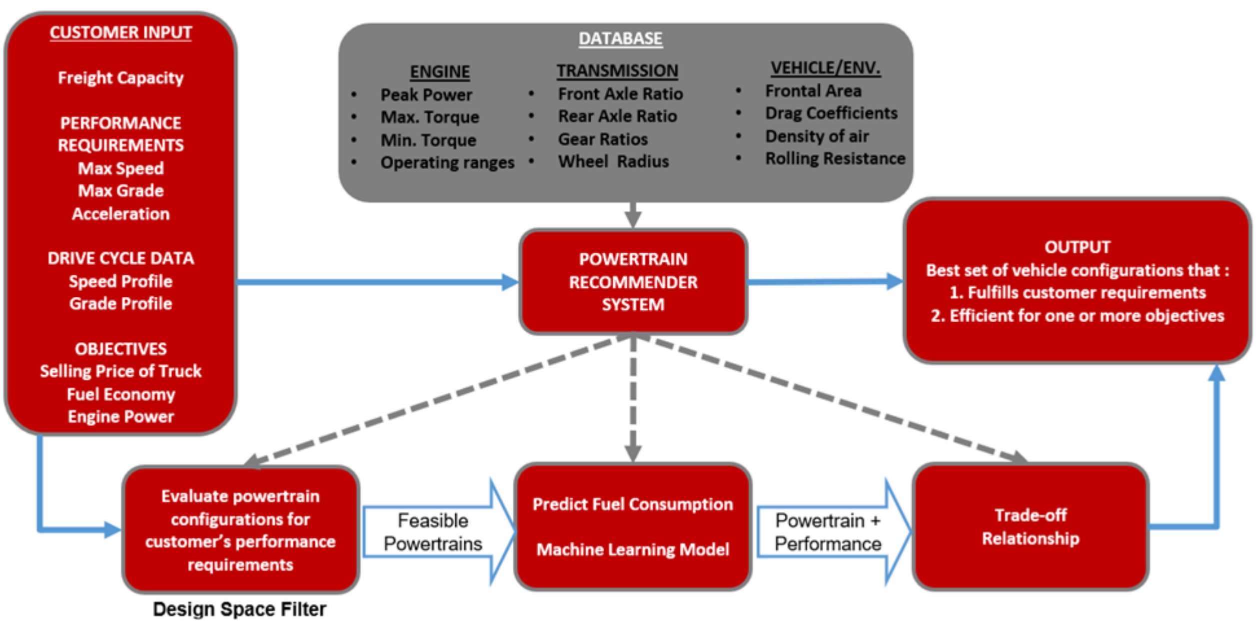

Figure 34: Powertrain Recommendation System data driven simulation design flow chart

A similar architecture is used to recommend a vehicle from a fleets portfolio that can complete a set of tasks in the most energy efficient manner possible in a tool called the Fleet Recommendation System (FRS). This system focuses on optimizing the cargo-weight assignment to the available fleet vehicles to maximize freight efficiency considering constraints such as cargo volume, cycle time, and operational cost. The FRS system flow chart to determine the optimal fleet was defined by the program partners as seen in Figure 35. This functionality was prototyped as part of the program but was not put into practice as part of the final results due to lack of variation in the fleet partners vehicles.

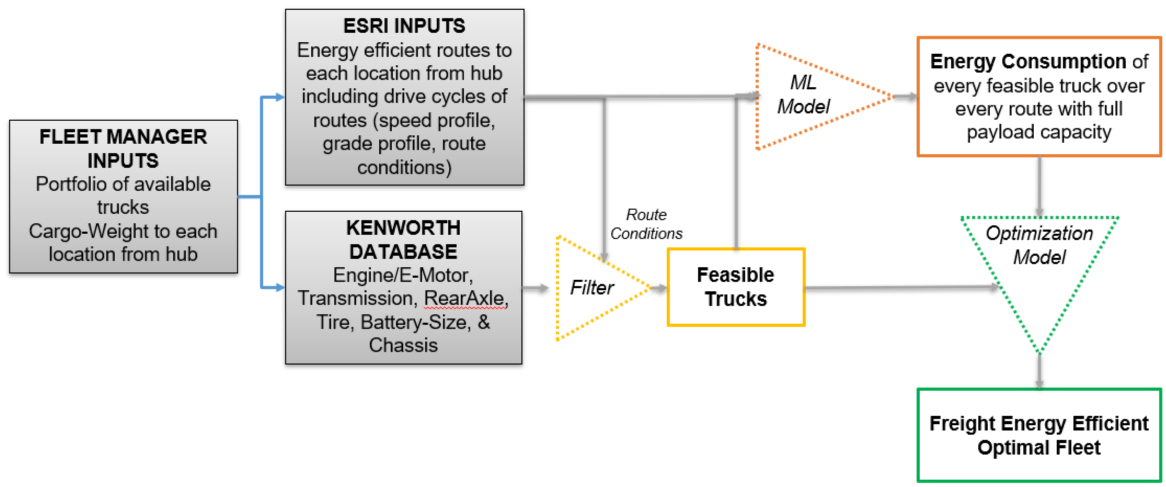

Figure 35: Fleet Recommendation System design flow chart

Conventional Powertrain Digital Twin

Development of a class 8 conventional powertrain vehicle digital twin was required to have the ability to calculate the energy consumption of different powertrain configurations on a drive cycle consisting of a speed versus time profile. The modeling activities were based on three diesel powertrain configurations: linehaul, vocational, and delivery. The modeling activities included:

- A machine learning model to predict the engine fuel consumption map from the engine test bench data
- Fitting the torque converter data to the Kotwicki torque converter model
- Analysis of transmission power losses
- Implementation of a gear shifting logic based on engine speed and road grade

The digital twin is a forward-looking energy-based model that relies on performance maps. The simulation parameters are divided into powertrain parameters (engine fuel consumption maps, transmission gear ratio, torque converter data, final gear ratio) and vehicle design parameters (mass, frontal area, drag coefficient, tire rolling resistance coefficient). The variables for the simulation are smooth speed versus time profiles, road grade, smooth grade versus time profile, and gross vehicle weight estimates.

Simulations were performed for trucks configured for linehaul, vocational and delivery with simulated data compared against the real-world data. Figure 36 shows the sample simulation data for a linehaul route compared to the recorded data. The simulation data and the recorded data have good correlation.

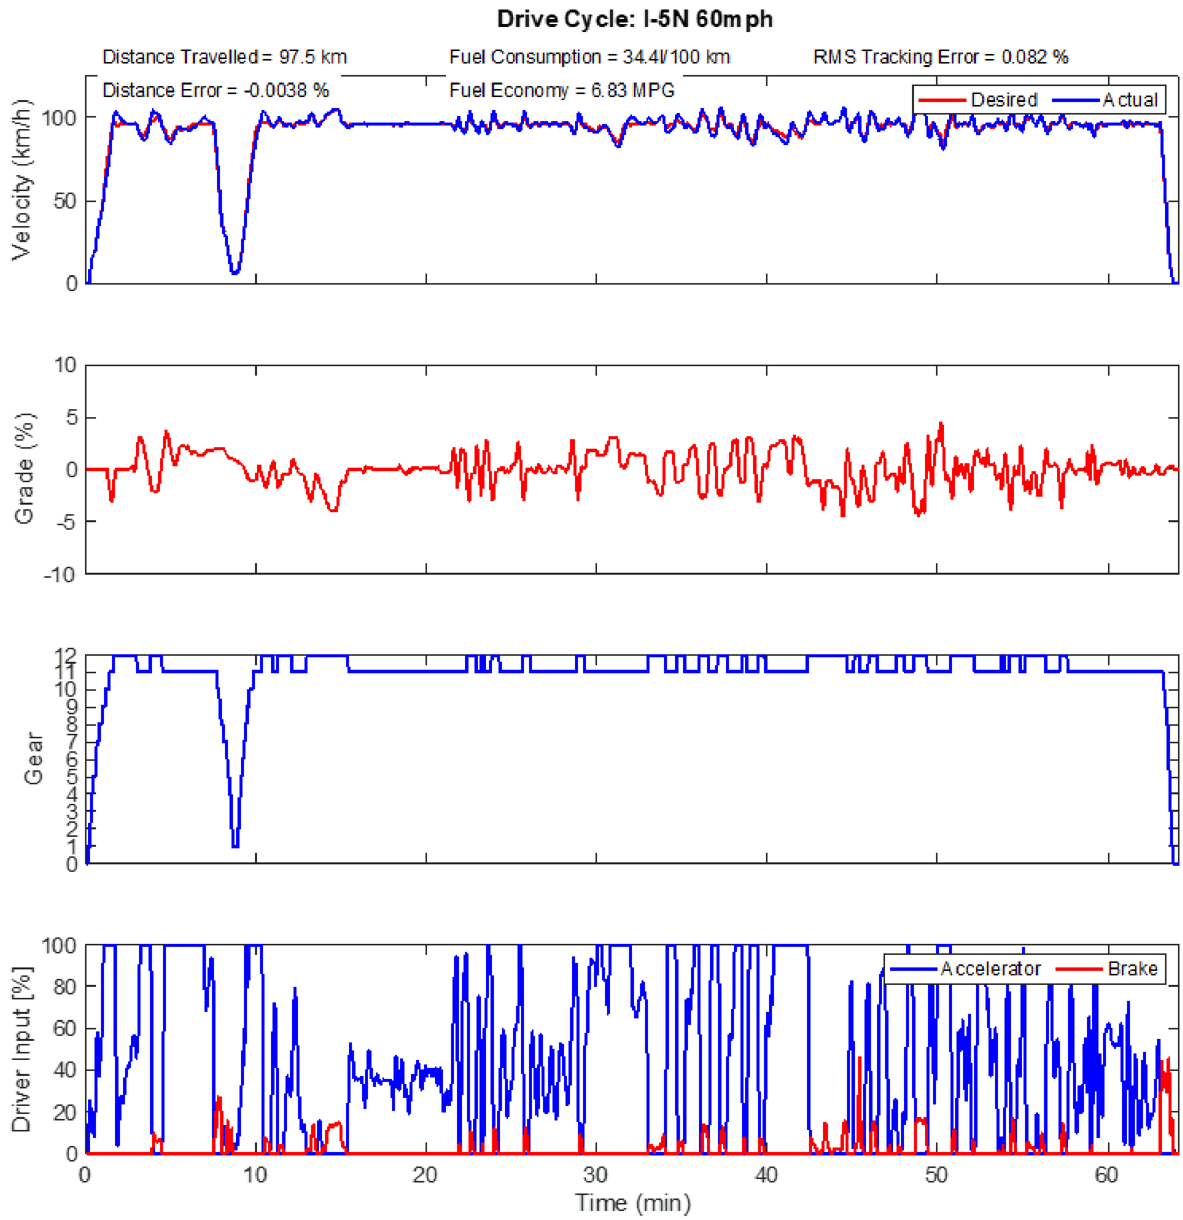

Figure 36: Comparison of simulated versus real world data for a linehaul cycle in the velocity plot

A Design of Experiments (DoE) was conducted for the three available powertrain configurations on 10 selected speed profiles from real-world driving data to ensure the results from the simulator were rational. Figure 37 shows the results of the DoE simulations. The first and second speed profile represent a complete interstate route, third to sixth represent highway operation (not necessarily starting or ending with vehicle at rest), seventh to tenth represent uphill climbing operation (not necessarily starting or ending with vehicle at rest). The results confirm that the vehicle specified for linehaul achieves the best fuel economy in highway/interstate operation. The results also show that the vehicle specified for delivery cannot be used in highway driving cycles demanding a speed above 65 mph (max allowed by this configuration's powertrain). It also appears that the grade in the last driving cycle is too demanding for the vocational configuration (more final gear speed reduction would help).

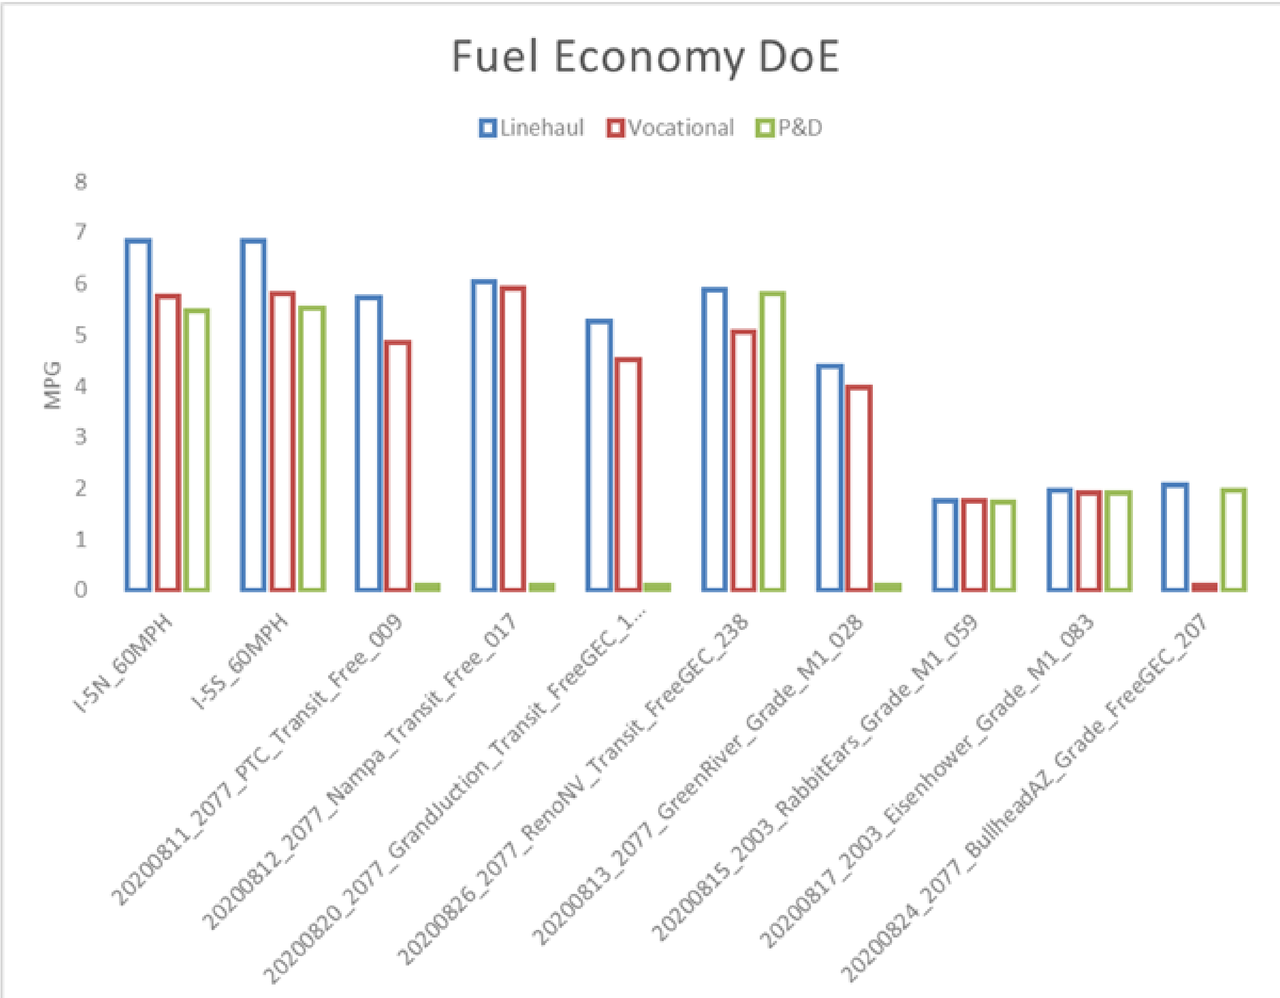

Figure 37: Comparison of fuel economy for a simulated design of experiments

Design Space Filter

Due to the very large number of class 8 diesel powertrain configurations available through Kenworth (97,198 included in the study) a Design Space Filter (DSF) was developed to reduce the number of available powertrain configurations, hence the required computational costs, to a subset of configurations known as ‘feasible powertrains’ that meet the specific customer performance requirements. As shown in Figure 38, the evaluation of powertrain configurations is done using physics-based longitudinal dynamics and road load equations that determine the engine power and transmission parameters required to meet core customer performance requirements such as the ability to maintain maximum vehicle speed on a maximum road grade, sustain continuous maximum vehicle speed, and starting acceleration on a flat road. The reduction to ‘feasible powertrains’ in the design space is crucial to speed up processing by removing configurations easily determined to not meet the minimum requirements.

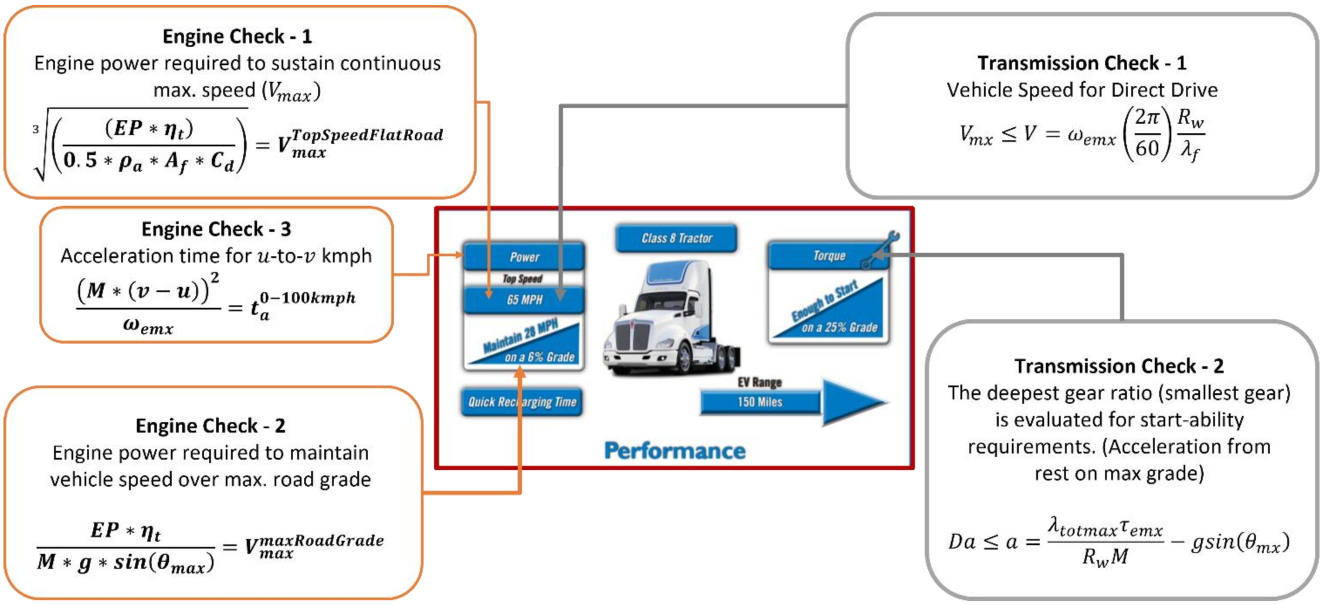

Figure 38: Performance evaluation and design space filter methods

Machine Learning Models

Multiple data-driven models to predict the energy consumption of the powertrain configurations were developed using supervised machine learning algorithms namely Neural Network (NN) and Random Forest (RF). The model was trained using real-world data for six powertrain configurations and was validated using simulation results for eleven powertrain configurations developed using the Kenworth design space for which real-world data was not available.

The relevant data features of vehicle specifications and powertrain performance to predict energy consumption were identified with the help of the Kenworth sales team and confirmed with data analysis conducted by OSU. Figure 39 shows the ML model training methodology and the data features that were used for model training.

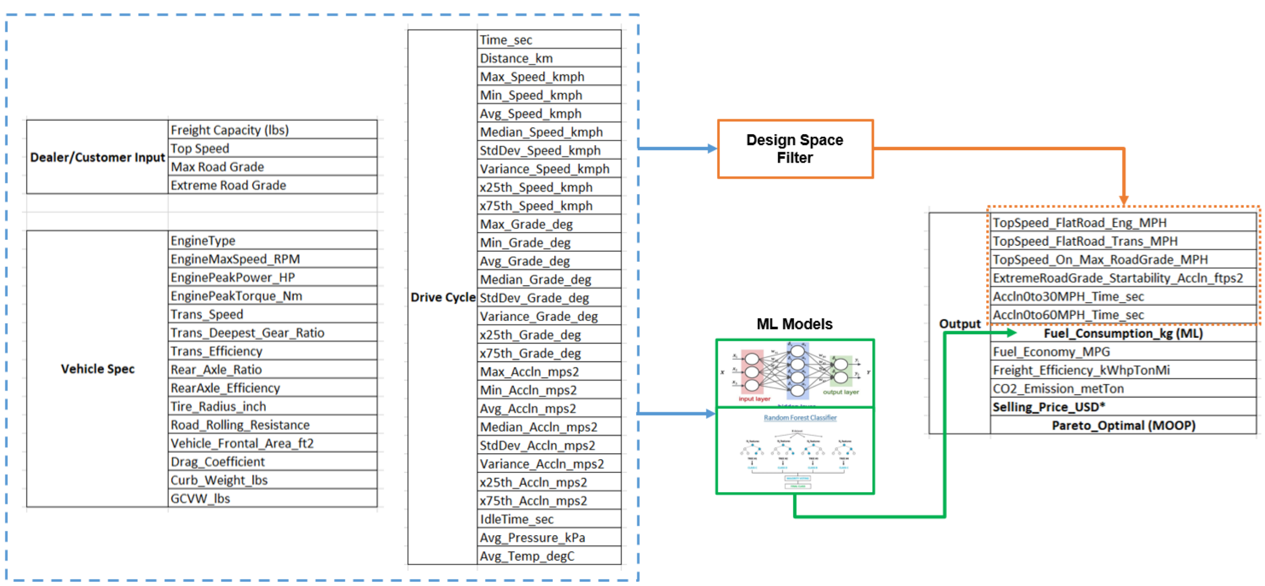

Figure 39: Machine Learning model training and relevant data features

RF models are better suited for limited data and where extrapolation is not critical. NN models are better suited when the dataset is sufficiently large, and extrapolation is critical. Figure 40 shows the model accuracy of the RF and NN models. The RF model outperforms the NN model for current training and test data due to limited datapoints and the test data is within interpolation limits. However, in the future, NN models will be preferred to leverage large scale telematics data when available.

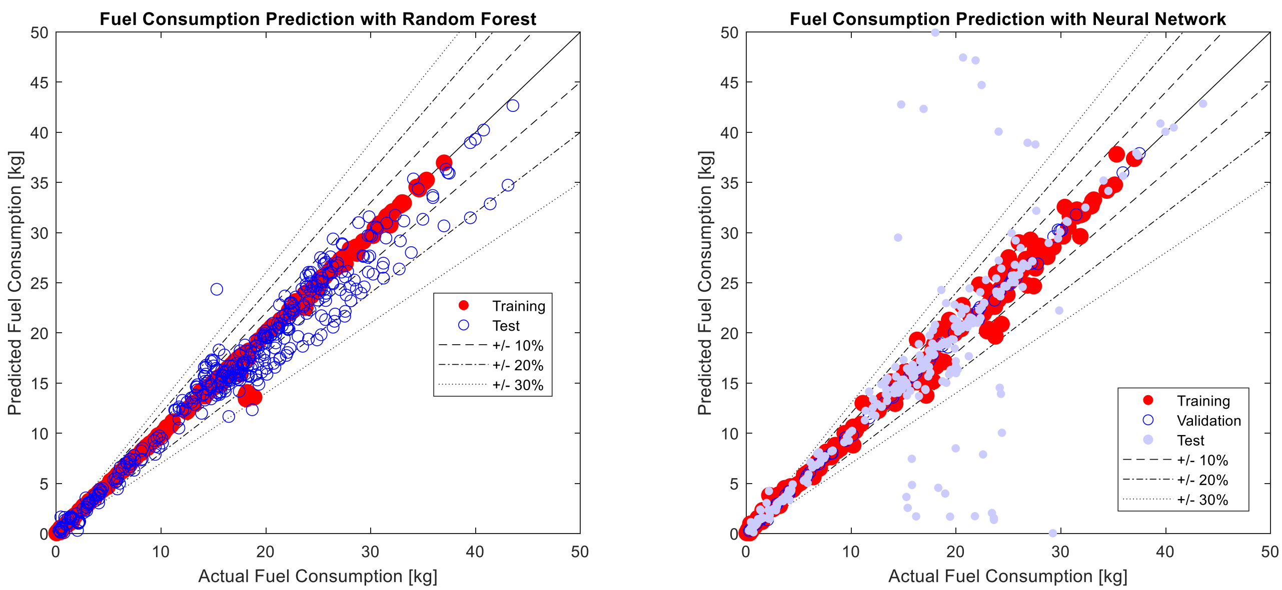

Figure 40: Machine Learning results: model accuracy of RF (left) and NN (right) models

Battery Electric Powertrain Digital Twin

The creation of the class 8 battery electric powertrain digital twin was a collaboration between OSU, Kenworth, PTC, and Meritor (now Accelera) who is the supplier of the EV axles utilized in the current PACCAR heavy duty EV's. The forward-looking model is based on the conventional digital twin and includes driver, controller, eAxle(s), battery, and vehicle longitudinal dynamics. A block diagram of the model is shown in Figure 41. The reference vehicle is the Kenworth T680E, whose specs are reported in Table 11.

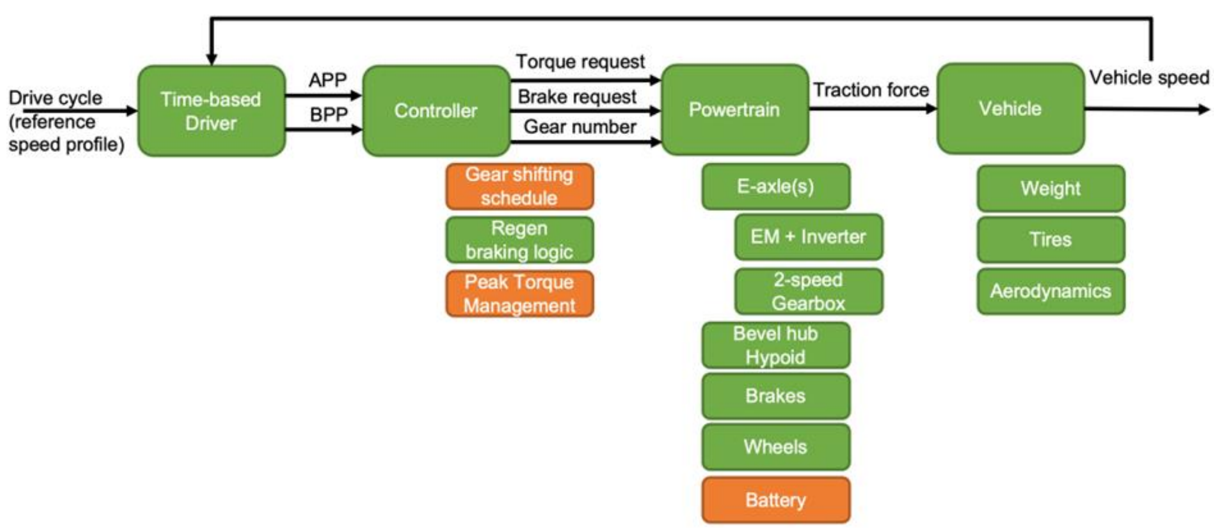

Figure 41: Block diagram of forward-looking model for battery electric powertrain

Specification	Value
Continuous Power	400 kW
Peak Power	500 kW
Continuous Torque	2,200 Nm
Battery Capacity	396 kWh
Range	150 miles (240 km)

Table 11: Kenworth T680E specifications example.

To obtain a variety of powertrains and performance within the category of battery electric trucks, three different e-Motors, two final gear ratios and three battery pack capacities have been modeled. The total number of possible powertrain configurations is 18.

A DoE with 18 configurations over 65 drive cycles each, for a total of 1170 simulations. Drive cycles come from the real-world operation of PACCAR conventional trucks on the routes shown in Figure 42. The simulation results have been analyzed and selected by using the following criteria:

- Vehicle velocity root mean square error (RMSE) must be less than 2 km/h over the drive cycle
- Battery state of charge (SOC) must not fall below 20% before the end of the drive cycle
- Battery SOC must be lower than 98% by the end of the drive cycle

If one of the above criteria is not met, the powertrain configuration cannot meet the speed, road grade, load, and range requirements of the drive cycle, and is discarded.

Figure 42: Real-world routes used to generate drive cycles (speed and grade profiles) for simulation

Based on the simulation the results show a 52% reduction in energy intensity between conventional diesel trucks and BEV trucks. From the kWh/mi figure, the average range for the three tested battery sizes was derived:

- 250 kWh battery: 95 miles
- 300 kWh battery: 114 miles
- 400 kWh battery: 152 miles

A similar approach to predicting the energy consumption of conventional diesel powertrains was used for BEV powertrain configurations. Using supervised machine learning algorithms, two models were developed: Neural Network (NN) and Random Forest (RF). The input dataset used was the same one previously created from simulation models from 18 battery electric truck configurations.

Figure 43 depicts the dataflow and relevant features considered in the performance evaluations and energy consumption prediction of the BEV powertrain configurations. For the machine learning models, vehicle specifications and drive cycles variables are considered for the predictions.

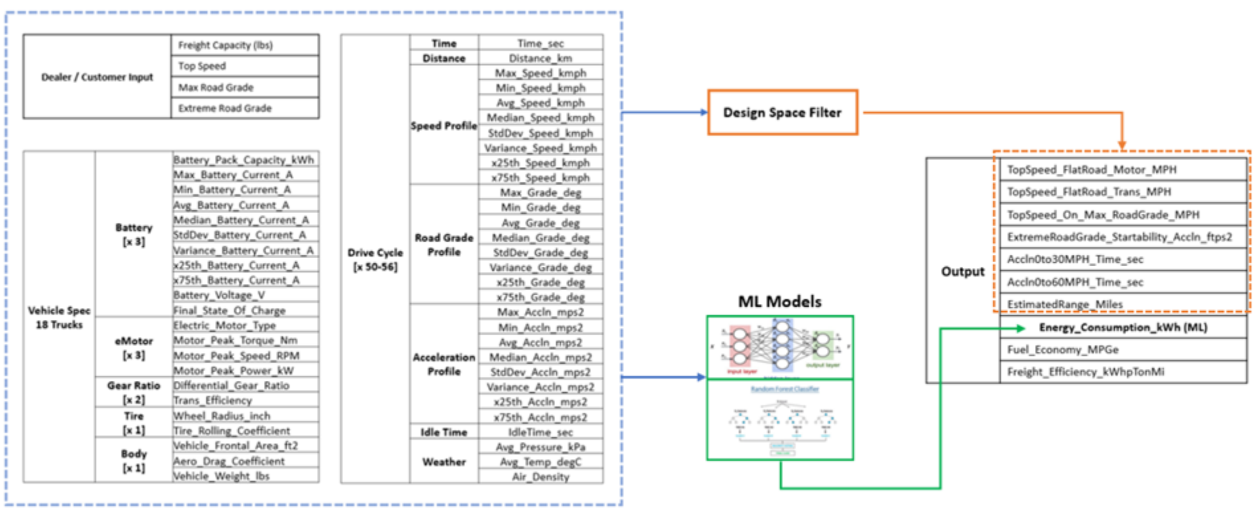

Figure 43: BEV powertrain machine learning model training and relevant data features

Similar trends were observed in the BEV machine learning models as compared to the conventional vehicle learning models. In situations where there is limited data and extrapolation is not significant, RF models perform better. Figure 44 provides RF and NN model performance for BEV trucks. As observed, the RF is performing better with an accuracy of >95% while the NN is > 90%. However, the RF model here outperforms NN because of the limited datapoints and the data it encounters in the test dataset is within the interpolation range, almost to the point of overfitting. The NN model provides good accurate values and does not show overly, almost absolute, accurate values as RF, indicating good generalizability tendencies and has the leverage to perform better when it encounters datapoints which are to be extrapolated.

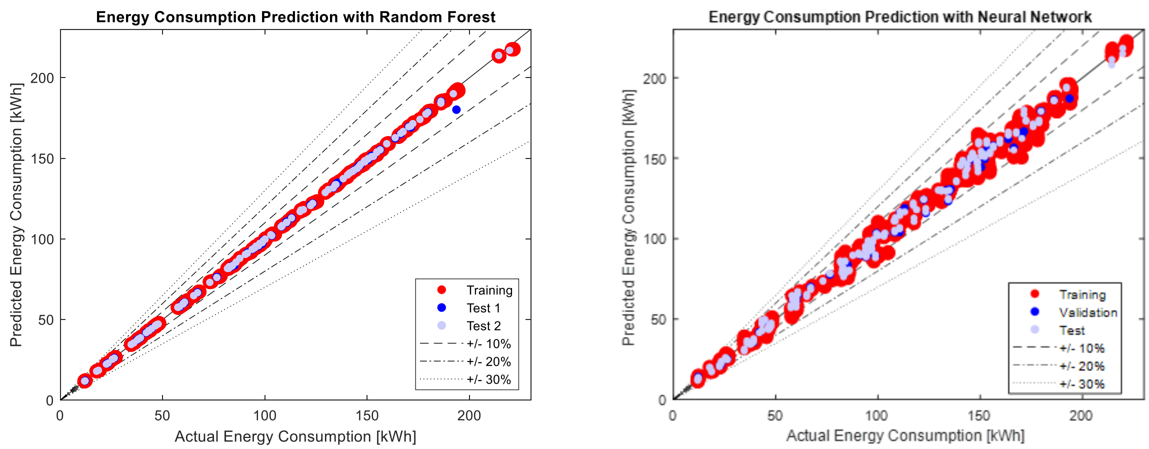

Figure 44: Electric powertrain machine learning results: model accuracy of RF (left) and NN (right) models

Powertrain Recommendation Tool Impact Assessment

For an initial assessment, a model configuration was selected which had one of the highest fuel economies (Miles per Gallon [MPG]), satisfactory freight energy efficiency [Ton Miles/kWh] and was utilized in all the driving scenarios. To keep the analysis uniform for all the recommendations, the client preferences or user constraints such as the maximum freight weight, maximum vehicle speed, maximum road grade, etc., are kept constant. For the impact assessment, real world drive cycles from PACCAR were used, and processed into micro-trips, with each file being representative of the individual regional route drive cycle. The micro-trip conversion was done to reduce data overloading and to retain coherent statistical data. These files are provided as inputs along with the available vehicle configurations to the PRS to obtain feasible configurations which meet the customer performance requirements and constraints.

The key indicators for understanding the improvement of the target were the number of efficient configurations which had better freight performances compared to the baseline and among them which provided the highest efficiency improvement. Figure 45 illustrates the configuration count which is based upon pooling similar efficient recommended configs together. The plot in Figure 45 signifies the initial estimate of improvement for the given route, with the better performing configurations highlighted in the Region of Interest (RegOI) and similar baseline performing configurations to the left of the figure.

The improvement in the freight ton efficiency (FTE) is analyzed with the following equation:

$$\Delta\, Improvement\ Ratio\ = \frac{FTE_{highest} - FTE_{baseline}}{FTE_{baseline}} \times 100\ \%$$

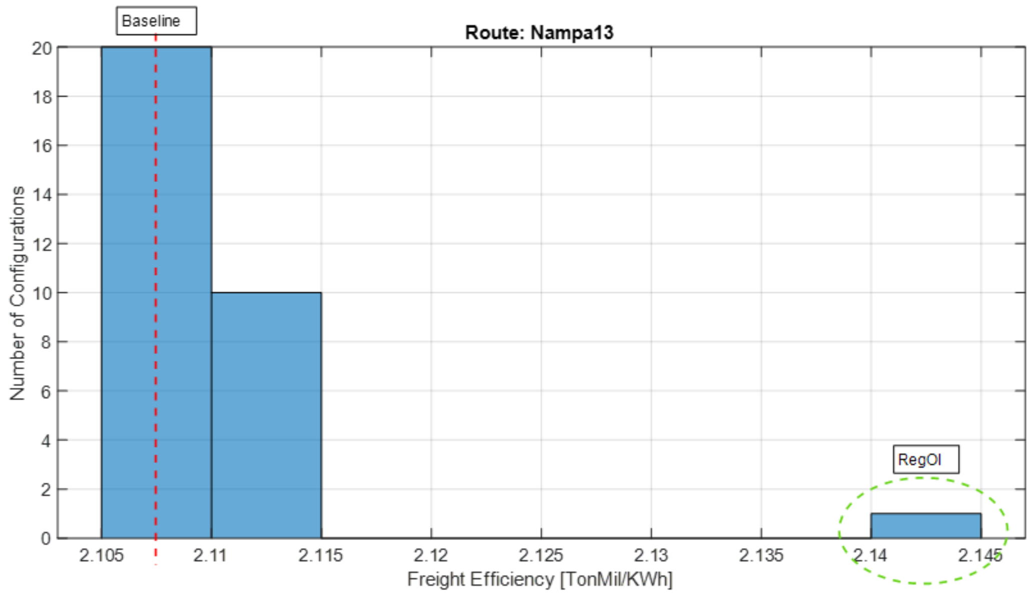

Figure 45: Histogram of configurations with their freight efficiencies for a given route

Figure 45 allows visualization of how the configurations are performing with each other for individual routes, with different recommendations providing the highest efficiency for each route. For example, the first block has an efficiency range of 2.105 to 2.11 with 20 configurations within it. The baseline configuration has an efficiency of 2.107 highlighted by the red dashed line but within that block. There are still 16 configurations performing better. To consider any real improvement in efficiency, we need to consider those configurations as highlighted in the RegOI area. There is only a single configuration for that route which shows any actual improvement in freight efficiency. For other regional routes the freight efficiency improvement was also found to be in the range ~1.5-2.2%, considering the configuration with the highest efficiency in that route against the baseline.

Application Integration

The powertrain recommender system application provides a graphical user interface (GUI) for the user to enter their requirements and get energy-efficient powertrain recommendations for a desired route. The PRS application is developed in MATLAB and takes in the customer inputs along with the design space and required drive cycle. It filters unfeasible configurations based on the user requirements through DSF and uses the developed machine learning models to predict energy consumption and provides a ranked recommendation based on freight efficiency metric.

Initially, the application would only consider conventional configurations and provide the powertrain configuration options best suited to the input parameters. Due to the inclusion of BEVs, the application was modified to take the BEV design space along with the conventional design space into consideration and uses both ML models to predict the energy consumption for both powertrains. Only one drive cycle file is required as the PRS is designed to provide conventional and BEV truck recommendations for the route and these results will be displayed under the respective tabs for the two corresponding vehicle architectures. If the length of the route is greater than what the BEV configurations can support, no electric vehicle is recommended to avoid confusion. This approach can be changed in the future as the charging infrastructure develops for opportunity charging. Figure 46 shows the completed application interface after the integration of BEV trucks. Figure 47 shows the application visualization of the processed drive cycle trace.

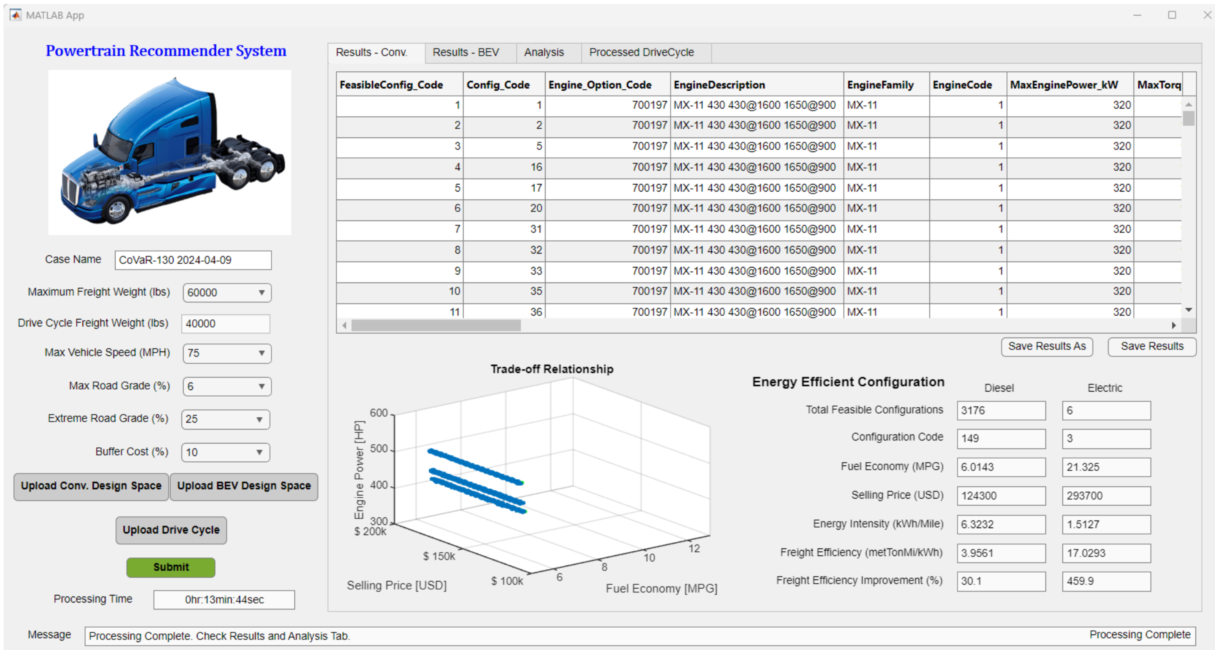

Figure 46: PRS application interface

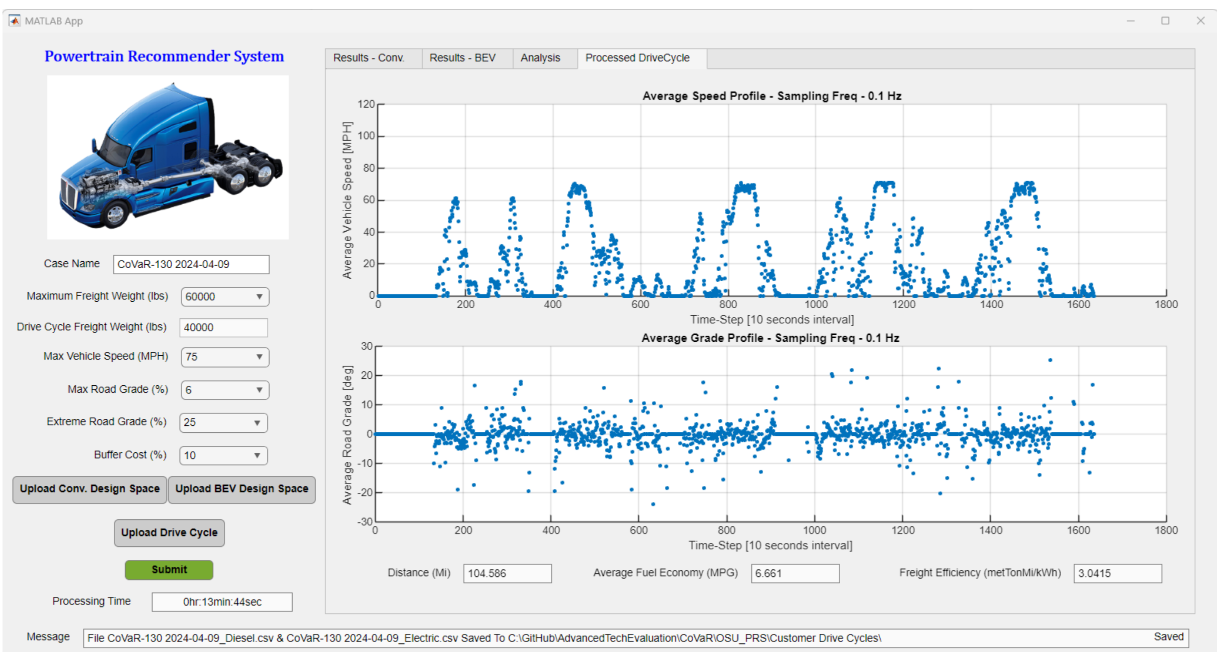

Figure 47: PRS drive cycle visualization

Fleet Partner Selection

Fleet partner selection was an area of the program that turned out to be more challenging than originally anticipated. As the fleet partners are operating a business, they understandably do not want to risk interfering with their normal operations nor share sensitive data. In addition, as time evolved from the application of this DOE funded program and where the technology stands today, many commercial vehicle fleets have adopted some sort of fleet management system offered by 3rd party suppliers, and they did not want to adapt another feature. This observation of changes in technology

maturity readiness slightly changed the direction of the program as we utilized as much live customer data and feedback as possible from the PacLease a division of PACCAR. The number of lease vehicles equipped with the telematics and cloud-communication system was similar to the originally scheduled fleet partner deployment but was characterized by a larger variation of use cases / applications.
The ideal criteria for a fleet partner in the program determined is listed in Table 12.

Priority	Fleet Partner Requirements
1	Interest in partnering with OEM on co-development of technologies
2	Interest in working with OEM to optimize fleet for Use-Case
3	Interest in introducing energy efficient technologies
4	Fleet size
5	Service (Mix of city and highway driving)
6	Interest in investing in Electrification/Hybridization in the future
7	Age of fleet (representative to market average)

Table 12: Requirements for CoVaR fleet partner selection

PacLease helped partnering with multiple fleets to utilize the CoVaR technology and record real-world in use data. This partnership allowed the use of multiple use cases, duty cycles, fleet operations, and more varied data than working with a singular fleet. The program utilized 2 fleets and 32 units with a mix of heavy duty and medium duty vehicles as shown in Table 13.

Fleet	Operation	Truck Class	Quantity
Partner 1	Air freight, LTL delivery	Medium duty (class 7)	11
		Heavy duty (class 8)	6
Partner 2	Used tire pickup and recycle	Medium duty (class 7)	6
		Heavy duty (class 8)	9

Table 13: Requirements for CoVaR fleet partner selection

7. Technology Evaluation

The fleet partners provided installation opportunities for the in-house developed telematics and high resolution cloud connectivity to collect data to support energy efficiency improvement analysis while utilizing and providing validation of the different applications developed under the program. Utilizing the data collected from the fleet partners enabled the use of realistic scenarios and best represents the program goal of a live demonstration.

Real-World Fleet Baseline Analysis

NREL analyzed the data collected during the program to characterize the baseline fleet operations for the Class 8 heavy-duty and Class 6 medium-duty vehicles, leveraging its FleetREDI (National Renewable Energy Laboratory) platform for supplemental data. The collected data consisted of engine-related parameters such as engine speed, engine fuel rate, and engine torque, and GPS parameters including location, speed, altitude, and route information for each vehicle class. Figure 48 provides an overview of vehicle operations based on the percentage of days analyzed across the 43 daily heavy-duty vehicle days.

Some of the key findings for the heavy-duty fleet are listed below:

- Daily Distance: Average daily distance is approximately 95 miles, with a maximum distance reaching up to 404 miles, however, most days are less than 200 miles.
- Average Driving Speed: Most daily average driving speeds are below 41 MPH (75th percentile), indicating a greater proportion of more transient city driving versus sustained highway speeds.
- Idle Time: The heavy-duty trucks spend a significant portion of their daily trips idling, with an average idle fraction of about 31%. Idle time is defined as periods when the vehicle speed is 0 for at least 45 consecutive seconds. This high idle time suggests potential opportunities for reducing energy consumption.

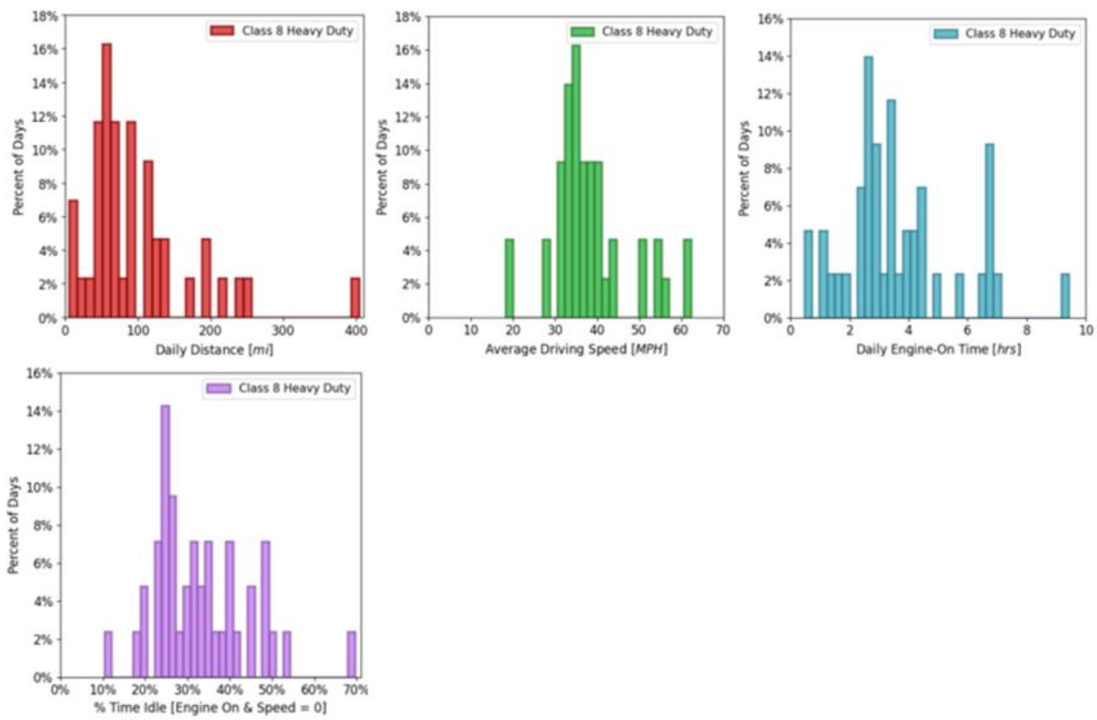

Figure 48: Histograms illustrating the distribution of metrics across all analyzed days for the heavy-duty truck dataset

Vehicle operations for the medium-duty vehicles are shown in Figure 49. The medium-duty dataset was limited, consisting of 12 vehicle days. This led to sparse distributions, making it challenging to draw conclusions or identify clear patterns within the data.

Key findings for the medium-duty fleet are listed below:

- Daily Distance: Average daily distance is approximately 52 miles, with a maximum distance reaching up to 187 miles.
- Average Driving Speed: In contrast to the heavy-duty trucks, the medium-duty trucks tend to spend more time driving in the city than on the highway, with average speeds typically below 43 MPH (75th percentile).
- Idle Time: The medium-duty trucks spend less time idling compared to the heavy-duty, with an idle fraction average of 23%.

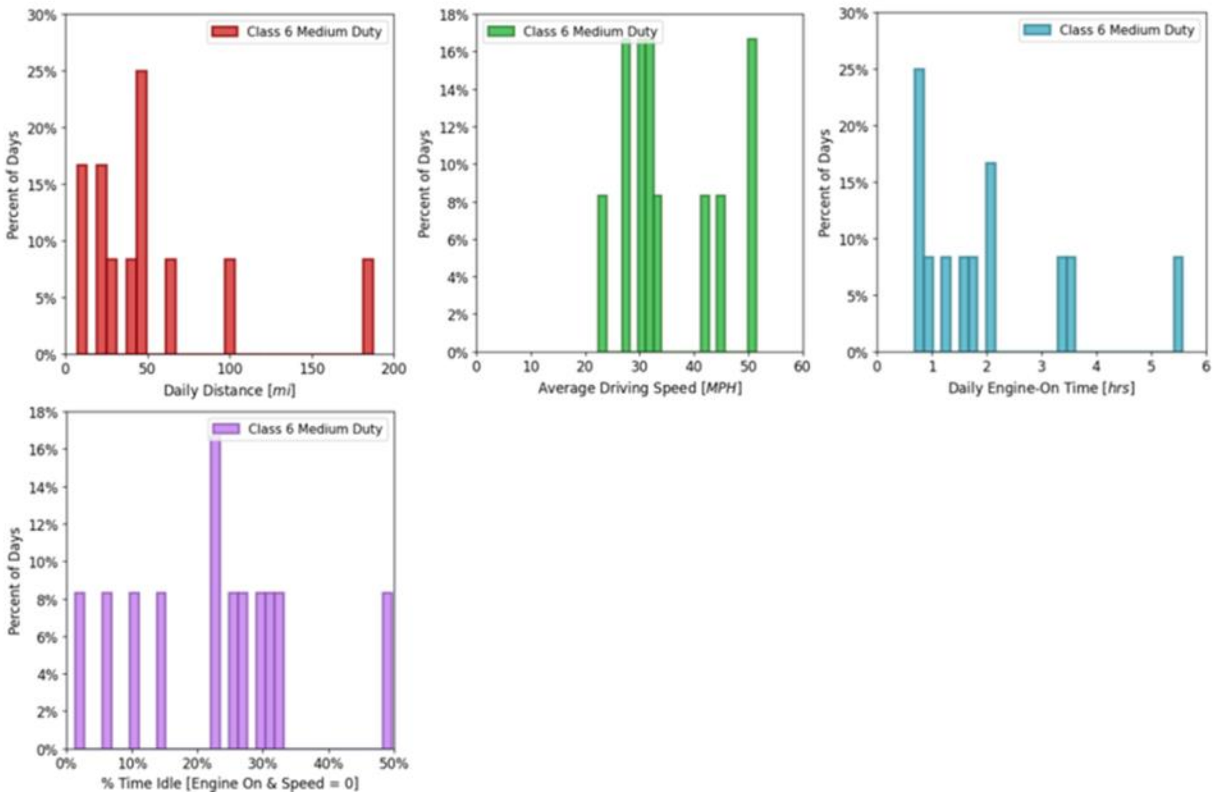

Figure 49: Histograms illustrating the distribution of metrics across all analyzed days for the medium-duty truck dataset

These metrics, in conjunction with the raw duty cycle data, were used to calibrate the Future Automotive Systems Technology Simulator (FASTSim) vehicle model by adjusting the model parameters to match the measured fuel consumption from the conventional trucks. This calibration finely tunes the model by adjusting the parameters that affect the road load equation until the error between predicted and actual energy consumption is minimized. Once the vehicle model was developed, it was run across the modified drive cycles to assess energy savings and provide driver feedback.

Vehicle Specification Optimization System Analysis

The heavy-duty data collected from the customer fleet was analyzed using the PRS designed by OSU. Drive cycles which were within the range that would qualify for BEV use were provided which BEV powertrain would be the most suitable along with the diesel powertrain for comparison. The diesel powertrain and the BEV powertrain recommended for all the drive cycles were the same, which makes sense as the powertrains currently in use are the same as well. The diesel powertrains utilized in the analysis had a PACCAR MX-13 455hp 1650lb-ft engine, with a PACCAR TX12 automated manual transmission and a 3.08 or 2.85 rear axle ratio, and the recommended replacement powertrain was a PACCAR MX-11 430hp 1650lb-ft engine, with a PACCAR TX12 automated manual transmission and a 2.6 rear axle ratio. This recommendation makes sense as the trucks are not running a maximum allowable GCVW, Texas is mostly flat without steep grades, and the speed limits for commercial vehicles are high compared to most other states. The recommended BEV powertrain for all drive cycles was a 14Xe 360kW 1030Nm set of eAxles and a 2.67 rear axle ratio.

Table 14 shows the comparison between the existing vehicle, the recommended diesel vehicle, and the recommended electric vehicle on the cycles. The recommended diesel powertrain provided an average of 31.8% freight efficiency improvement, and the BEV powertrain provided an average of 382.6% freight efficiency improvement.

Drive Cycle	Distance	GCVW	Diesel Improvement	BEV Improvement
1	72.1 Mi	60,000 lbs.	32.6%	507.2%
2	104.6 Mi	60,000 lbs.	30.1%	459.9%
3	99.8 Mi	60,000 lbs.	33.2%	435.2%
4	69.1 Mi	60,000 lbs.	44.2%	183.8%
5	102.0 Mi	60,000 lbs.	22.1%	257.1%
6	86.6 Mi	65,000 lbs.	28.8%	452.2%

Table 14: Recommended powertrain freight efficiency improvements

Trips that covered long distances at freeway speeds were excluded from the results as freight efficiency calculated was too high to be rational. The training data used by OSU to develop the ML model did not have vehicle speeds as high as the vehicles in Texas had. The average highway vehicle speed for trucks in Texas was around 70mph, which is higher than most states allow for commercial vehicles. It is believed the extrapolation of the ML model to estimate energy consumption at these higher speeds was not accurate as aerodynamic load has a large impact on energy consumption and is not taken into account with recorded data for correlation.

Fleet Management System Analysis

The fleet partners had a fleet of trucks with similar specs for heavy-duty and medium-duty. Because of the homogeneous fleet composition, the original intent of the FMS to select the more optimal truck for the route was not utilized. Instead, we focused on utilizing the FMS in the method that would best align with our partners, allowing them to assign the stops to the correct truck and allow the system to optimize the order of the stops and the best route between stops. As this system was not able to be deployed live, we were not able to quantify the benefit of the SVMS functionality built into the routing optimization as we could not verify the suggested stop order would work for the business purpose and be a valid energy savings.

A function that was added to the FMS to aid the fleet manager in assigning the best vehicle for the job was to include the option to select if the route is ideal for a BEV truck. It was determined to use a manual selection versus a calculated selection based on route parameters because there are so many variables related to BEV utilization. Based on experience the fleet managers would know which routes made sense to utilize a BEV truck, and the reminder is there to ensure if one is available, it should be deployed here. An example of a trip that has been identified to work well with a BEV truck is shown in Figure 50.

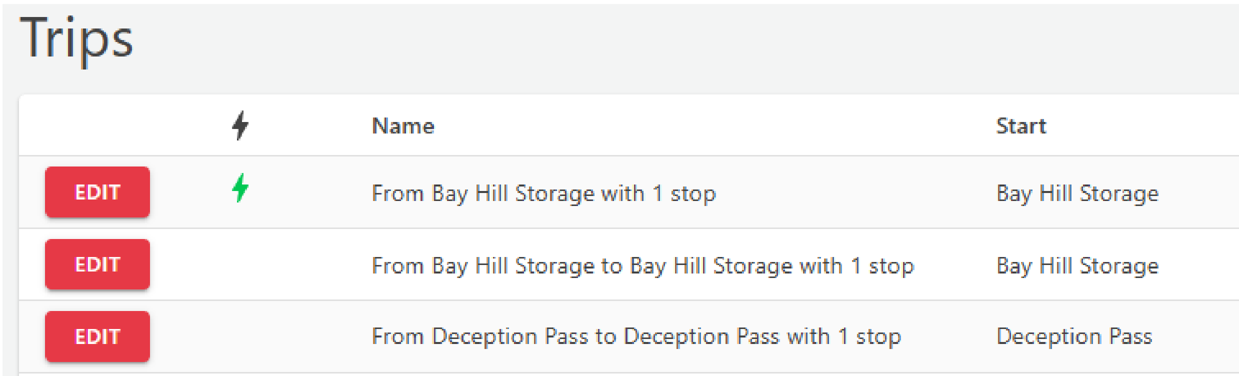

Trips

	⚡	Name	Start
EDIT	⚡	From Bay Hill Storage with 1 stop	Bay Hill Storage
EDIT		From Bay Hill Storage to Bay Hill Storage with 1 stop	Bay Hill Storage
EDIT		From Deception Pass to Deception Pass with 1 stop	Deception Pass

Figure 50: Identifying a trip as suitable for a BEV truck

NREL analyzed the baseline data to determine the percentage of trips that were suitable for BEV use based on a conservative range estimate compared to advertised range for both heavy-duty and medium-duty vehicles PACCAR currently offers. As shown in Figure 51 76.7% of the heavy-duty drive cycles and 91.7% of the medium-duty drive cycles used in the baseline analysis could be performed with a BEV truck. The analysis shows even at the moderate range commercial BEV offerings currently have; they are suitable for many fleets operations.

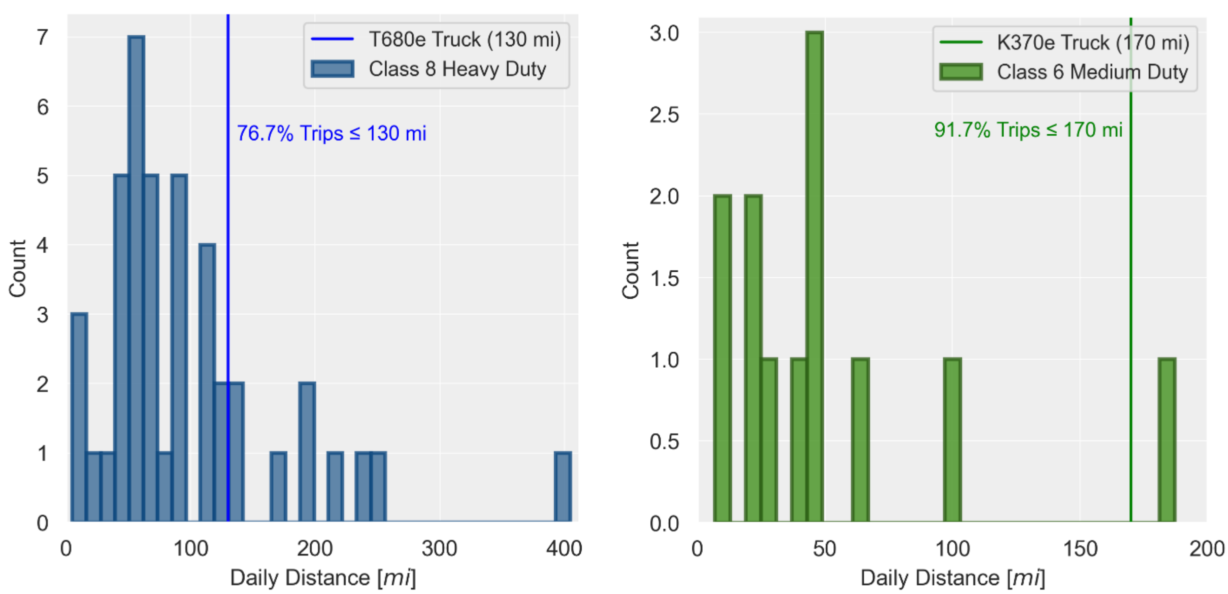

Figure 51: Histograms showing the percentage of baseline drive cycles suitable for BEV trucks

As this functionalities freight efficiency savings could not be quantified due to the lack of BEV trucks in the program no estimate for the FMS benefit is provided. It is important to note the FMS is the key enabler for the other functionalities like sending the necessary vehicle information to the routing optimization, so its impact on the overall freight efficiency improvement should not be downplayed.

Intelligent Driver Assistance System Analysis

The IDAS system is responsible for providing the driver with optimal routing (eco-routing) during the trip, and driver coaching (eco-Coaching) after the trip where it was determined the driver would be more receptive to the feedback. Due to the system performing two distinct functions, the analysis was separated into those functions for clarity.

Eco-Routing Analysis

To quantify the potential for energy savings from alternative vehicle routing, NREL used on-road data collected from medium and heavy-duty vehicles operated by the fleet partner in Texas using their RouteE-Compass tool. This data included a total of approximately 12 vehicle-days from 3 unique medium-duty vehicles and 44 vehicle-days from 6 unique heavy-duty vehicles containing nearly 300 microtrips (microtrips are separated by idle events of 3 minutes or longer). The methodology used for the eco-routing analysis is summarized in Figure 52.

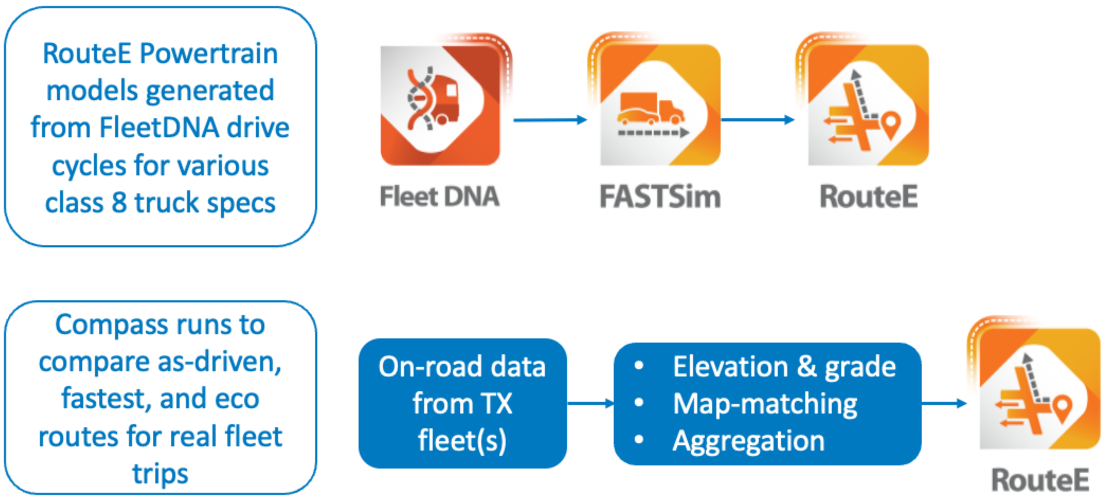

Figure 52: Summary of the eco-routing energy savings analysis methodology

RouteE-Compass was utilized to compare optimal routes considering multiple objectives (e.g., time and energy). The results of this comparison for all heavy-duty microtrips are shown in Table 15. The average energy savings and time penalty are shown for the population of microtrips with and without "long trips" which are longer than 75 miles. There are 5 "long trips" in the data with very high energy savings (~50%), but also impractically high time penalties (~40%), so results with and without these trips are presented.

Objective	Heavy-Duty (no long trips)	Heavy-Duty (all trips)
Average Energy Savings	11.8%	27.7%
Average Time Penalty	12.3%	20.7%

Table 15: Average energy savings and time penalties for the set of all heavy-duty microtrips

The real-world data collected from the medium-duty vehicle fleet is almost an order of magnitude smaller and the energy savings and time penalty results were highly sensitive to which microtrips were included in the analysis, so it was concluded that the data quantity was insufficient for this analysis. An average energy savings of 1.9% and time penalty of 5.4% was seen for the limited medium-duty trips. However, given the data sample size, no comparison is able to be confidently made between the energy savings opportunity from eco-routing for medium-duty vs. heavy-duty vehicles.

The distribution of energy and time implications of alternate paths for all heavy-duty microtrips are shown in Figure 53. Key takeaways from the analysis of this are:

- 58 trips (27%) had no alternate path with energy savings –the fastest path was already the most energy efficient
- 158 trips (73%) had a more efficient alternate path
- Among the 158 trips with a more efficient alternate path, the average energy savings are 11.2% (median = 4.1%) and the average time penalty is 14.7% (median = 8.0%)

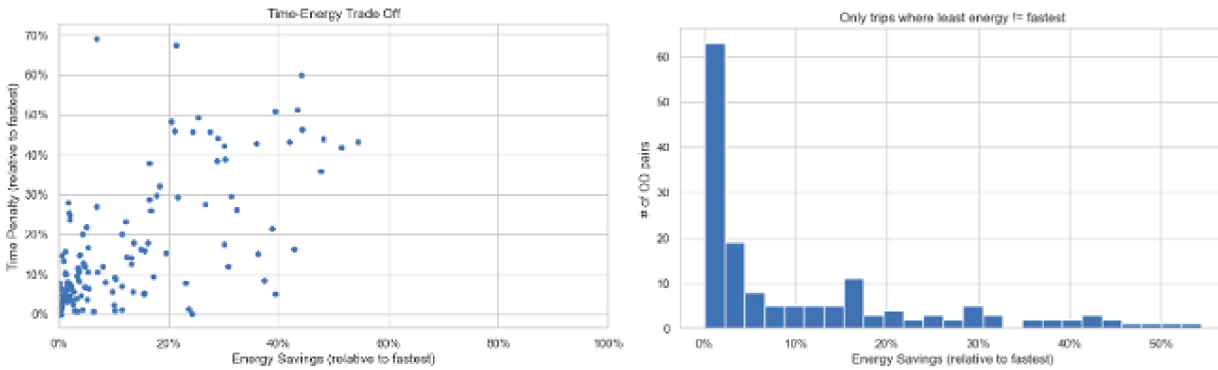

Figure 53: The scatterplot on the left shows the energy savings vs. travel time penalty for each microtrip with a viable alternative path. The histogram on the right shows the distribution of possible energy savings for the least energy consumption path compared to the fastest path

Figure 54 shows two examples of route alternatives resulting from RouteE-Compass. The time optimal paths are shown in purple and energy optimal in green. The tables show trip distance, energy savings, and travel time penalties for the energy optimal paths. The example on the left routes the truck along a lower speed combination of state routes and arterials, which result in more efficient cruising speeds than the time optimal interstate path. The example on the right shows an energy optimal path with a shorter distance than the time optimal path and cruise speeds are also slightly reduced.

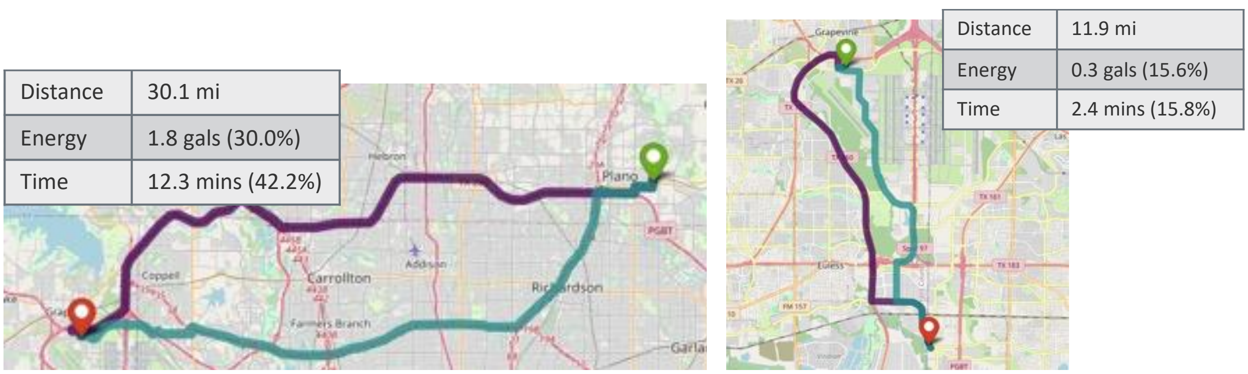

Distance	30.1 mi
Energy	1.8 gals (30.0%)
Time	12.3 mins (42.2%)

Distance	11.9 mi
Energy	0.3 gals (15.6%)
Time	2.4 mins (15.8%)

Figure 54: Two examples of time optimal (purple) and energy optimal (green) routes resulting from RouteE Compass for the heavy-duty dataset

It is important to emphasize that even these heavy-duty results are using a relatively small data set specific to two fleets and other fleets will have different energy savings opportunities and travel time penalties depending on application. Additionally, more can be done to identify alternate routes that achieve some or most of these potential energy savings with smaller time penalties by exercising the multi-objective optimization capability of RouteE Compass. Practical or economic thresholds of travel time penalties that fleets are willing to accept to achieve certain energy savings can also be imposed on this optimization to account for a particular fleet's labor, fuel, and maintenance costs. However, despite the data limitations, energy savings on the order of 10% through eco-routing have been demonstrated as possible for the two real-world partner fleets considered in this analysis.

Eco-Coaching Analysis

To determine the potential impact eco-coaching could have on the drive cycles collected from the fleet partners NREL used FASTSim to develop vehicle models for the Class 8 heavy-duty and Class 6 medium-duty trucks. FASTSim requires the following key inputs:

- A vehicle "database" in CSV format to input vehicle specifications (ex. type of powertrain, drag coefficient, frontal area, rolling resistance coefficient, number of wheels, max fuel convertor power, etc.)
- grade = Road Grade (%/100)
- mps = Vehicle Speed (m/s)
- time_s = Relative time in cycles (seconds)

The calibration process involved comparing simulated fuel consumption with the actual fuel consumption data from 1 Hz drive cycles for each vehicle class. To minimize discrepancies, the key variables in the road load equation were adjusted using a genetic algorithm, which found the optimal values for each parameter with the specified constraints for the vehicle "database." Any minor differences between the measured and simulated fuel consumption were acknowledged as limitations of the modeling process. Once the vehicle model is calibrated, the measured drive cycles are run in FASTSim to obtain several timeseries outputs, including the simulated fuel consumption and distance traveled. The unmodified simulated drive cycles from FASTSim serve as the baseline for comparison against the modified drive cycles to assess potential energy savings.

To simulate changes in driving behavior, the drive cycles were adjusted based on parameters that affect overall energy consumption that were identified during the eco-score development, with the following modifications:

- Maximum Speed Limit (Reducing by 5%, 10%, 15% for each drive cycle)
- Enforcing Speed Limit across all drive cycles if applicable (60 MPH, 65 MPH, 70 MPH)
- Maximum Acceleration Limit (Reducing by 10%, 15%, 20% for each drive cycle)
- Idle Reduction (Reducing occurrences in descending order by duration, applying reductions of 25%, 50%, 75% for each drive cycle)

The maximum speed in each daily drive cycle was reduced by 5%, 10%, and 15%. Figure 55 shows an example of a modified drive cycle (vs. original) with the cumulative distance traveled. To ensure the cumulative distance remained consistent for the modified drive cycle, speed values were added as necessary to compensate for reduced speeds, while ensuring the set speed limits were not exceeded. For the example shown, the original maximum speed was 70 MPH, and was reduced by 15%, with an adjusted maximum speed of 60 MPH. This adjustment increased the total driving time from 206.5 minutes to 208.5 minutes, confirming that lower speeds require more driving time to cover the same distance.

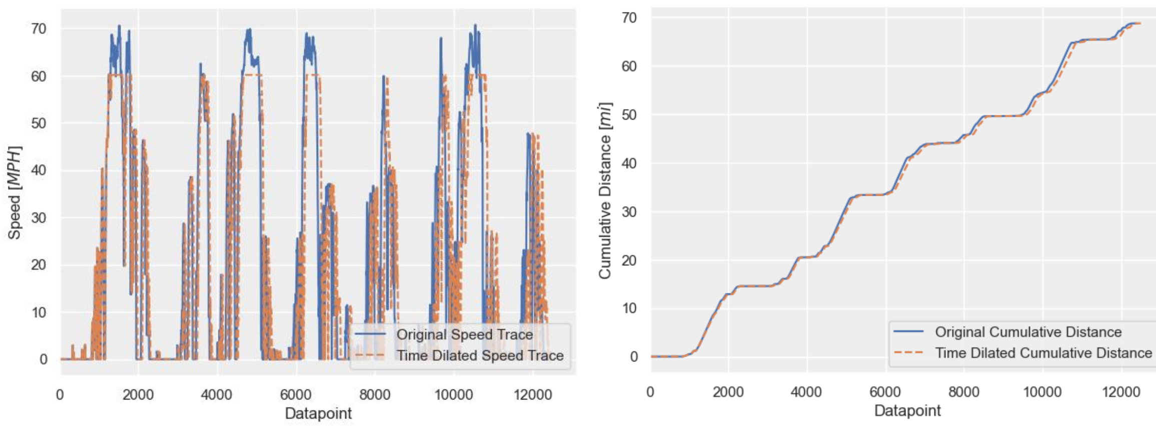

Figure 55: Example of reducing the maximum speed by 15% for the respective drive cycle

For assessing the effects of idle reduction, the idle windows were defined as periods where the vehicle's speed is 0 for at least 45 consecutive seconds (the minimum idle time can be adjusted). These idle windows are then sorted in descending order by duration, and reductions are made based on the specified percentage, prioritizing the removal of longer idle periods first. This approach assumes that during these periods the engine will be turned off, thus reducing overall energy consumption. Figure 56 shows the original drive cycle on the left and a modified drive cycle with the longest 50% of the idle occurrences removed.

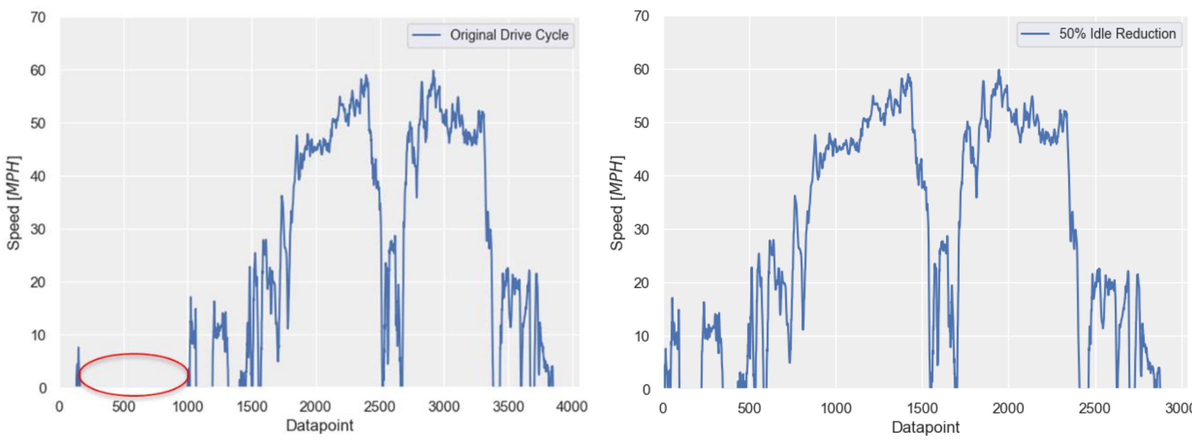

Figure 56: Example of the original drive cycle vs. idle reduction by 50% and assuming vehicle is off during those occurrences

The section below evaluates the impact each parameter had on the fuel consumption for each drive cycle. The relative energy savings are determined by running the modified drive cycles through the calibrated vehicle model and comparing it against the respective simulated fuel consumption from the unmodified drive cycles.

The results for reducing the maximum speed by a certain percentage for each drive cycle are presented below for heavy-duty and medium-duty trucks, with Figure 57 showing the relative energy savings and their respective counts for each speed reduction case. For a 15% reduction in maximum speed, the average relative savings were 2.3%, with a max of 9.2% for the heavy-duty trucks. Similarly, for a 15% reduction in maximum speed, the medium-duty had an average relative savings of 2.7%, with a max of 6.1%. Energy savings are relative to how often each drive cycle reaches their maximum speed limit.

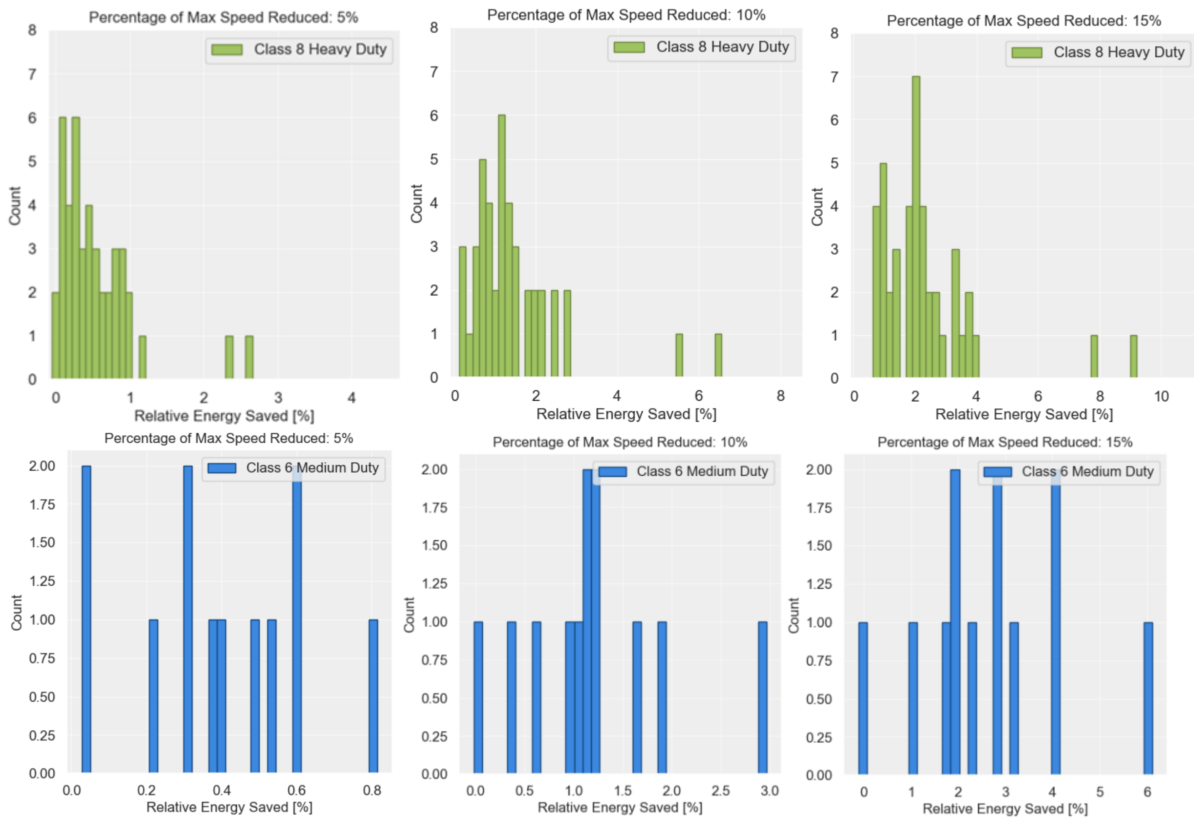

Figure 57: Relative energy savings (%) vs. count for each percentage reduction case in maximum speed. Top charts for the heavy-duty trucks and bottom charts for the medium-duty trucks

Figure 58 contains two scatter plots illustrating the energy savings achieved by reducing the maximum speed in each drive cycle by a specific percentage, compared to the relative time increase for heavy-duty and medium-duty trucks, respectively. As the percentage of maximum speed reduction increases, so does the driving time. This is dependent on the number of instances where the speed exceeds the limit and requires modification.

Figure 58: Relative time increase (%) vs. relative energy saved (%) for reducing maximum speed by a respective percentage for the heavy-duty (left) and medium-duty trucks (right)

Figure 59 shows the relative energy savings for the drive cycles impacted by the following enforced speed limits: 60 MPH, 65 MPH, and 70 MPH. This analysis uses a similar methodology to the maximum speed reduction by percentage but applies a uniform speed limit across all drive cycles that exceed it (the number of affected drive cycles for each speed limit is shown in Table 16). Only the drive cycles impacted by these speed limits are included to avoid skewing the histograms, which are shown in Figure 59.

Objective	60 MPH Speed Limit	65 MPH Speed Limit	70 MPH Speed Limit
Class 8 Heavy-Duty	40/43 Drive Cycles	40/43 Drive Cycles	37/43 Drive Cycles
Class 6 Medium-Duty	11/12 Drive Cycles	11/12 Drive Cycles	8/12 Drive Cycles

Table 16: Number of drive cycles affected by the given enforced speed limits for each vehicle class

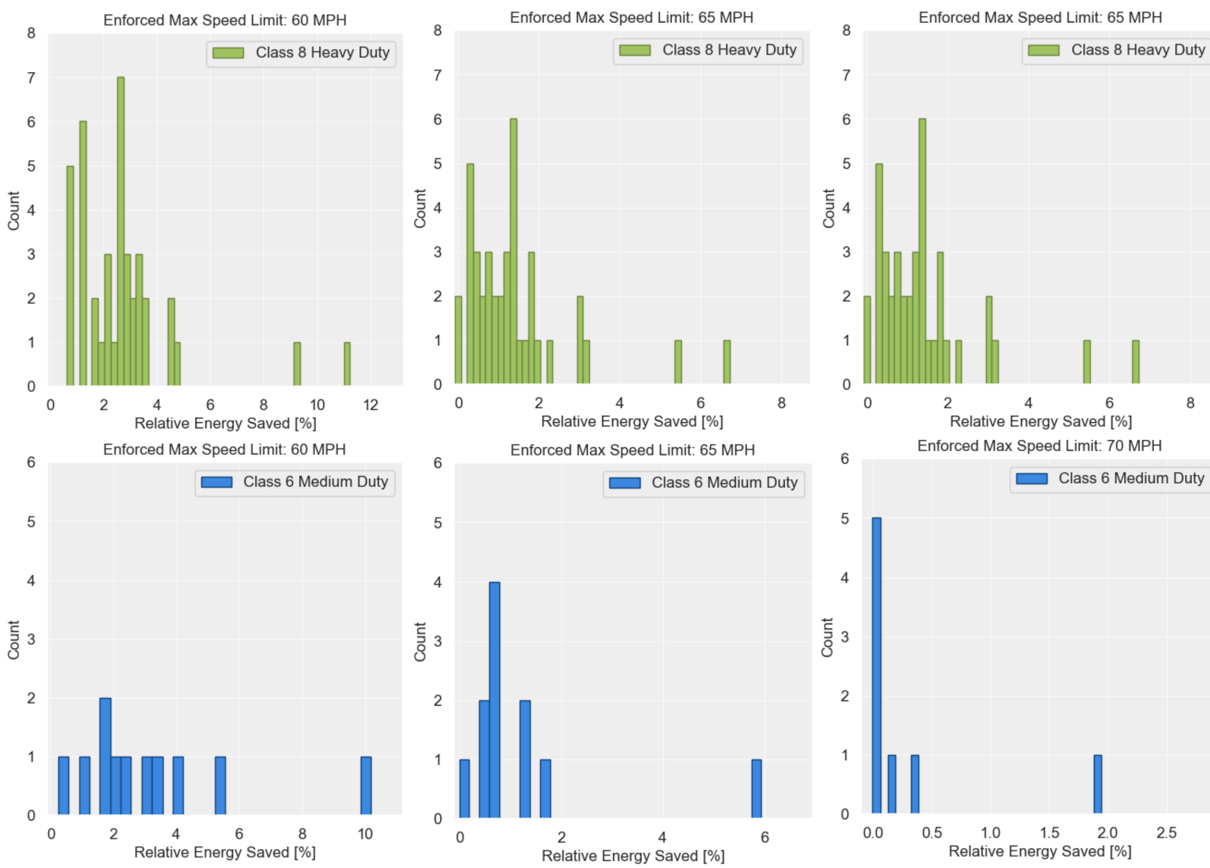

Figure 59: Relative energy savings (%) vs. count for each enforced speed limit. Top, heavy-duty trucks and bottom, medium-duty trucks

Figure 60 presents the results for idle reduction, demonstrating that minimizing idle time provides greater potential for energy savings compared to maximum speed reduction cases. This is likely because the effectiveness of speed reduction relies on how often the maximum speed limit is reached during each drive cycle, which directly impacts total driving time. In contrast, turning off the engine to reduce idle time can lead to consistent energy savings without increasing driving time, and is independent of driving conditions.

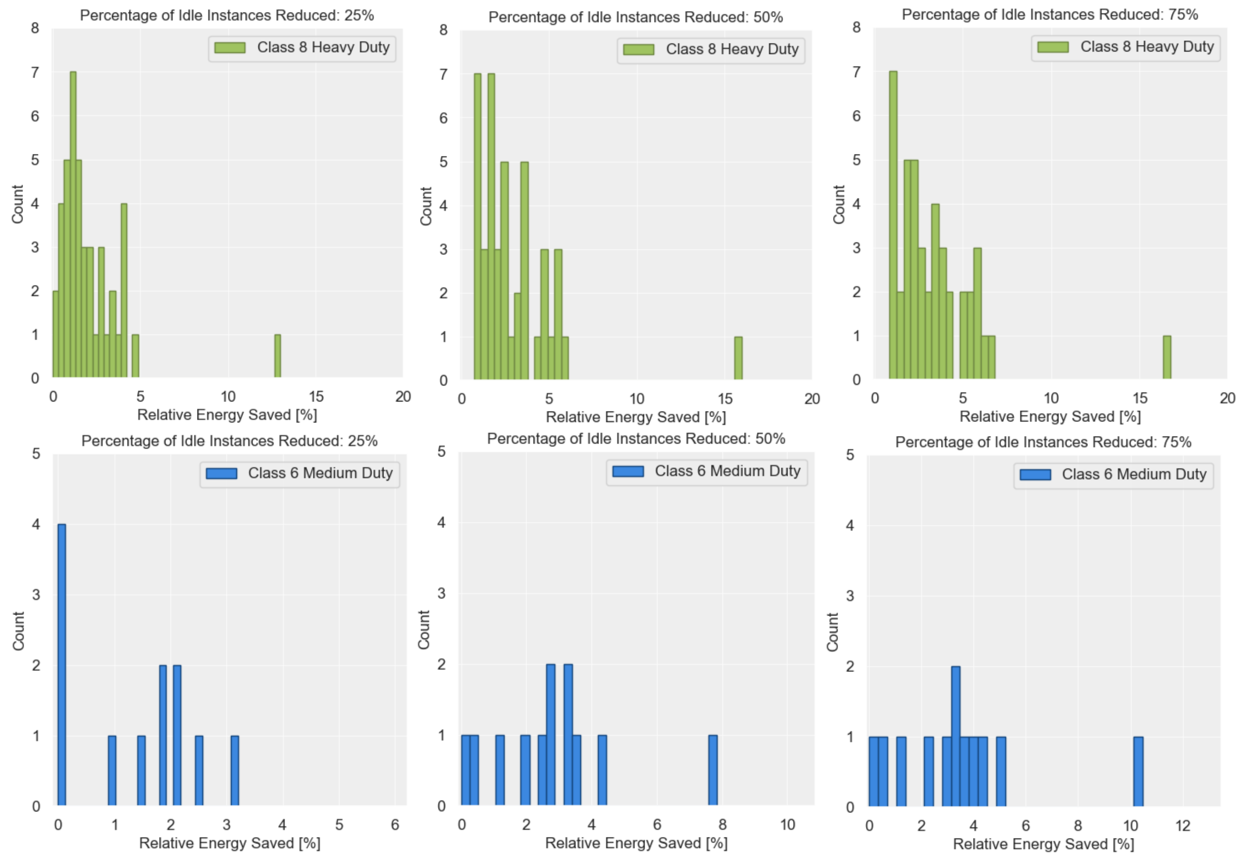

Figure 60: Relative energy savings (%) vs. count for each percentage reduction in idle occurrences. Top, heavy-duty trucks and Bottom, medium-duty trucks

Reducing maximum acceleration has a minimal impact on energy consumption, as shown in Figure 61, where the results remain close to those of the unmodified drive cycle. This limited effect occurs because changes to maximum acceleration primarily affect only a few acceleration events within each cycle. Since these events constitute a small portion of overall driving behavior, energy consumption remains largely unchanged (almost 0).

When maximum acceleration is capped, it controls the rate at which speed increases, requiring adjustments to maintain the same distance. This could result in quicker acceleration at different points in the drive cycle to compensate (without exceeding the given acceleration limit), potentially leading to slight increases in fuel consumption. In the future, implementing a speed-based acceleration limit may be beneficial because it would enable adjustments to acceleration based on the specific speed of the vehicle and the driving conditions at that moment. For example, instead of applying a fixed maximum acceleration limit, a speed-based limit could adjust dynamically: if the vehicle is traveling at a lower speed, the acceleration limit could be set lower to promote gradual increases in speed.

Additionally, changes in the maximum acceleration have little effect on driving time compared to changes in maximum speed (as shown in Figure 66), as high acceleration events typically occur at lower speeds and are less frequent in highway driving. Thus, while reducing acceleration can affect aspects of vehicle performance, its overall impact on energy consumption and driving time is relatively small for the selected reduction percentages.

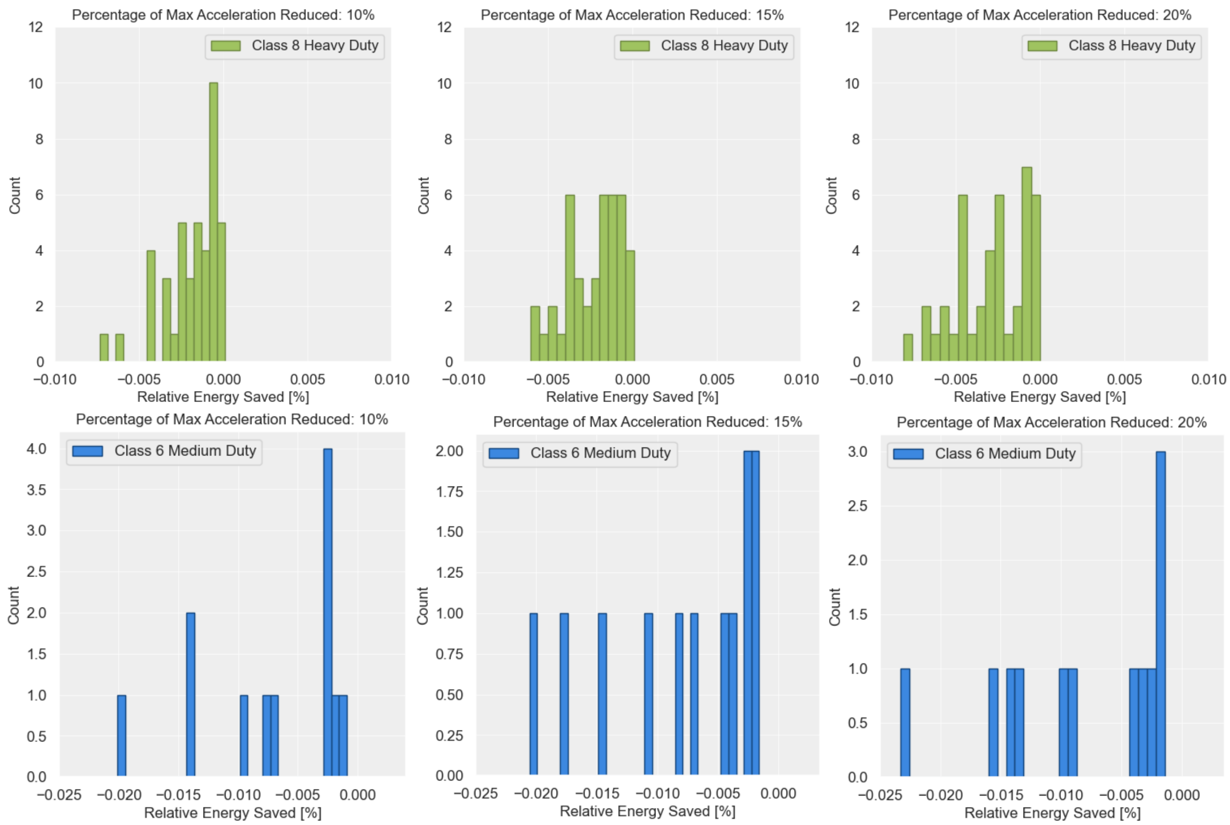

Figure 61: Relative energy savings (%) vs. count for each percentage reduction case in maximum acceleration for the heavy-duty trucks (top) and medium-duty trucks (bottom), respectively

Figure 62 combines both speed and acceleration effects in a single plot to compare their relative impact on energy and driving time. This highlights that reducing speed leads to a more substantial increase in travel time compared to reducing acceleration, but also allows for greater energy savings.

Figure 62: Relative time increase (%) vs. relative energy saved (%) for reducing maximum speed and maximum acceleration for the heavy-duty trucks (left) and medium-duty trucks (right), respectively.

As the dataset for the medium-duty trucks was limited, the combined analysis focuses on the heavy-duty trucks. Since idle and speed demonstrated the most potential for energy savings, they were combined to show mid-case and best-case scenarios. The scenarios included are as follows:

- Reduce maximum speed by 10% and reduce 50% of idle occurrences for each drive cycle.
- Reduce maximum speed by 15% and reduce 50% of idle occurrences for each drive cycle.
- Reduce maximum speed by 15% and reduce 75% of idle occurrences for each drive cycle.

Figure 63 shows the results for each scenario. The mid-case case scenario, which was 15% maximum speed and 50% idle reduction, had an average savings of approximately 5%. The best-case scenario, which was 15% maximum speed and 75% idle reduction, had an average energy savings of approximately 7% and a max of 18%. As mentioned, the savings are relative to how often the speed limit is reached, the number of idle occurrences, and the respective durations in the specific drive cycle. Both interventions can be implemented in the fleet with minimal impacts on the drivers.

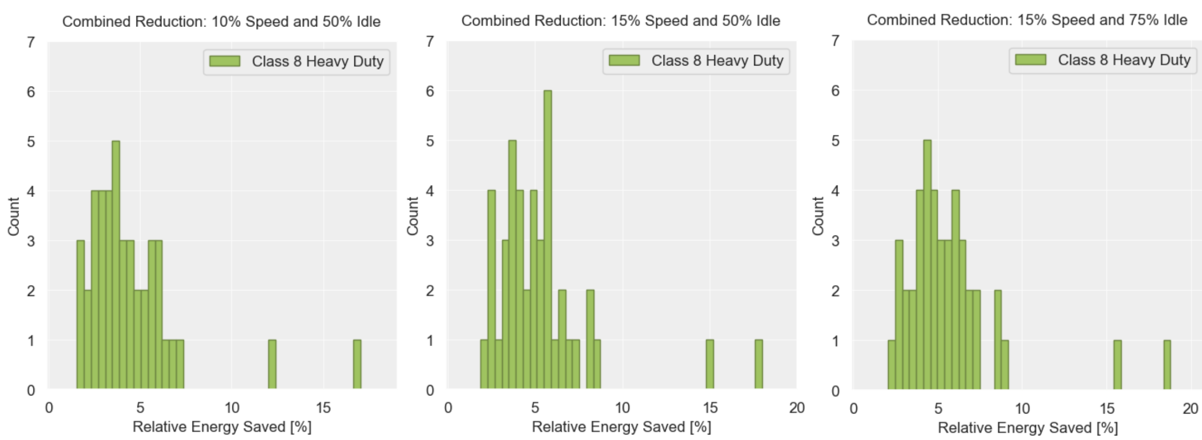

Figure 63: Maximum speed and idle reduction combined cases for heavy-duty trucks

These findings highlight the potential energy savings that can be achieved by implementing strategies such as limiting speed, limiting acceleration, and reducing idle time. The analysis demonstrates that even modest adjustments in driving behavior, such as establishing a maximum speed limit and minimizing idling, can lead to significant energy savings. While capping maximum acceleration showed minimal impact on energy savings, incorporating a speed-based acceleration limiter in the future may offer potential savings. It's important to note that the impact of different interventions on energy savings is dependent on the operations of the fleet. Reducing idle may have the biggest impact for one fleet, whereas lowering the maximum driving speed may be the most effective change for another. By adopting eco-driving techniques, drivers can increase fuel efficiency, which can contribute to cost savings and environmental sustainability. This is particularly important for fleets that cannot transition to electric due to infrastructure limitations, high costs, and the demands of long-haul routes, but aim to reduce emissions and increase fuel efficiency.

To communicate the best methods to conserve energy to the driver the final eco-score and individual component scores for idle, acceleration, braking, and top speed are displayed on the IDAS at the end of the trip. In addition to the eco-score, relevant trip metrics and methods to improve future scoring based on the components that most significantly impacted the overall score in a negative way are displayed. Based on the results from the NREL investigation into the impact of the various components on energy consumption scalars were applied to vary the impact on each component on the overall score. The scalars chosen for the example are 50% for idle, 10% for acceleration, 10% for braking, and 30% for top speed. An example of the driver display for the eco-coach is shown in Figure 64 for one of the drive cycles recorded.

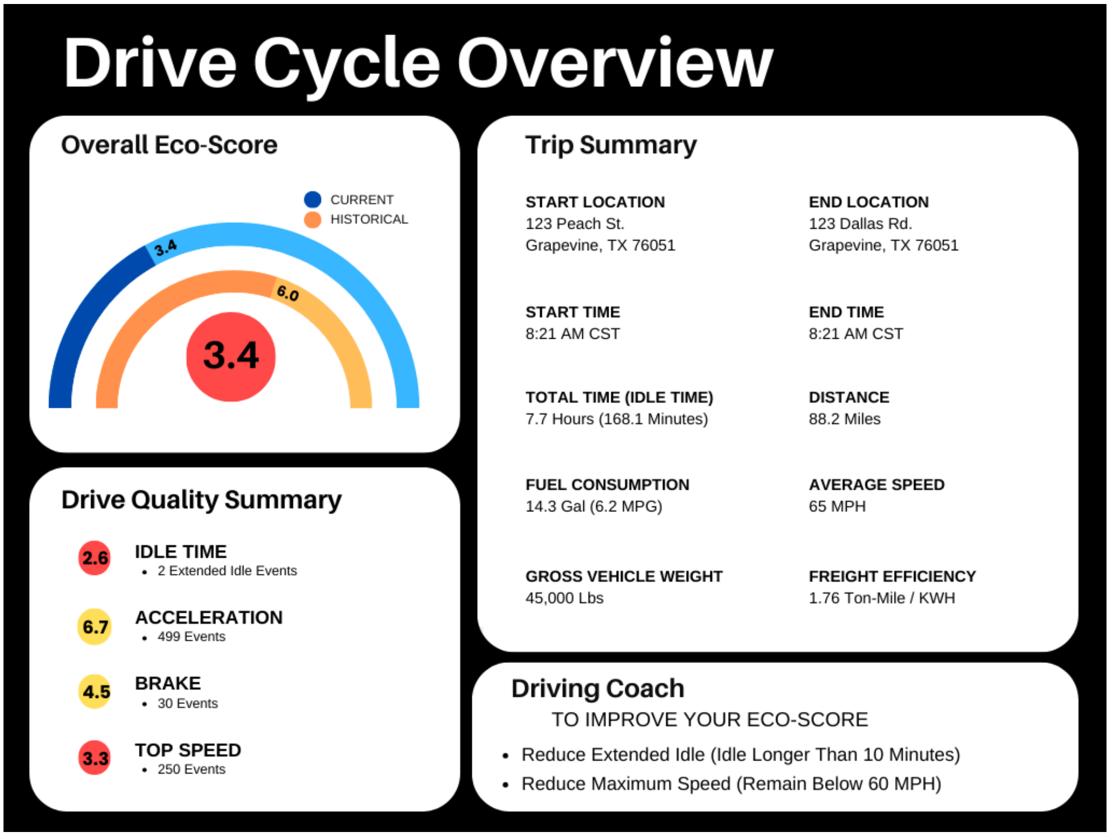

Figure 64: Eco-Coach example from drive cycle CoVaR-130 2024-05-03

Overall Energy Savings

The program has improved freight efficiency by 32% using the Vehicle Specification Optimization System for diesel powered vehicles. This powertrain optimization potential is regarded as the best-case scenario, optimizing fleet vehicles for specific routes and applications. PACCAR's experience, confirmed by the data-collection under this program, shows that fleet owners tend to be conservative in choosing powertrains to ensure a wider application coverage and will realistically achieve a range of 10-15% efficiency improvement. Much larger benefits are possible switching from Diesel to BEV (up to 380% improvement), but these results will be disregarded in the final numbers for this program as the original aim was to focus on existing fleet powertrain technologies.

Additional freight efficiency benefits are supported by the other fleet efficiency features developed under this program. A 15% freight efficiency improvement was demonstrated using the Intelligent Driver Assistance System (10% from eco-routing and 5% from eco-coaching). When combining new powertrain selection with the IDAS system, a total reduction in Brake-specific energy per freight ton-mile of 25-30% for diesel powertrains was demonstrated.

8. Technology Transfer

The technology developed under this program will be utilized in future commercialization and investigative projects. One of the technologies that will be utilized in future programs is the low-cost telematics solution to enable the large-scale collection of data. Experience working with the cloud database will enable easier data transfer, as well as provide access to cloud computational power to support ML analysis. Future versions of the telematics will be designed to utilize a cellular router for reliability and will only use a tablet for a visual display or optional interface to prevent tampering and software incompatibility between the Android operating system and the logging device.

The Android applications developed under this program enabled us to build a skillset that will be applied to future research projects that utilize a mobile device or tablet. Android is also the base operating system of many automotive applications that run on infotainment displays, which can be useful to prototype and get feedback on future development programs.

As part of this program an eco-score algorithm was developed for the BEV trucks leveraging the logic used in the diesel vehicles. The idle component of the score was removed and logic for regen utilization was added in its place as this is a critical component to energy utilization on the BEV trucks. As no BEV data was collected as part of the program results from this feature were not provided in the analysis, however, the BEV eco-score algorithm will be improved as part of PACCAR’s SuperTruck 3 (DE-EE0009861) program.

One development from the program which has shown good potential for commercial application is the PRS tool. The tool will be enhanced as part of the SuperTruck 3 program with available BEV truck data, and there is interest in leveraging the concept within the sales tool to help guide customer powertrain selection and identify if an EV would work for their application and which specification is the most suitable.

Currently 3rd party suppliers offer a range of fleet management systems which were not available when the program started. The energy saving potential for the applications as shown in this program provide justification to prioritize integrating available software into fleet offerings.