Evaluating Charging System Probability of Failure on Demand

To Improve Vehicle/Grid Integration Reliability

CHARGEX consortium

MARCH 2026

INL/RPT-26-91549

Argonne National Laboratory

INL Idaho National Laboratory

NATIONAL LABORATORY OF THE ROCKIES

Disclaimer

This information was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness, of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. References herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

This report was prepared for the U.S. Department of Energy under DOE Idaho Operations Office Contract No: AC07-05ID14517, U.S. Department of Energy M&O Contract No: DE-AC36-08GO28308, and U.S. Department of Energy Office of Science Laboratory Contract No: DE-AC02-06CH11357. Funding was provided by the Joint Office of Energy and Transportation

Author

D. M. Anand, Idaho National Laboratory

Executive Summary

Traditional reliability metrics often focus on component uptime or success rates. The end-user experience, defined by the expectation that charging will be successfully initiated "first time, every time," demands a more sophisticated, customer-centric metric. This report details a framework for quantifying Probability of Failure on Demand (PFD), a metric that measures the likelihood of a system failing at the precise moment a user requires its service.

The analytical foundation of this framework is the application of continuous-time Markov processes to model the charging system represented as a set of discrete operational states ($\mathbb{X}$) —some functioning, some degraded, and some failed. By defining the probabilistic rates of transition between them, it is possible to create a dynamic, quantitative representation of system behavior. This report elucidates this methodology, detailing the crucial distinction between the transition rate matrix ($\mathbb{A}$), which captures system failure pathways, and the reset matrix ($\mathbb{R}$), which quantifies the effectiveness of autonomous recovery strategies. This separation allows independent analysis and optimization of component robustness and architectural resilience. This formulation is particularly well suited to model complex failures where there is incomplete information about some sub-systems or components. In this report we consider three distinct charging architectures where the electric vehicle charging infrastructure would have to interact with the electric grid to control load and respond to grid signals. Architectures used for vehicle grid integration (VGI) are known to have complex failure modes that may be difficult to diagnose. Further, formal failure models are likely to benefit grid planners as they aim to integrate electric vehicle charging infrastructure more fully into their grid reliability portfolio.

Three distinct VGI architectures are considered in this report. The first is a model for a "Day-Ahead Pricing" system where charging loads respond to retail energy costs. The second is a generic “DER Definition” architecture commonly used to enroll a charging load into a grid service program such as Demand Response or Critical Peak Pricing. The third is a future looking architecture for "Transformer Protection" where fast acting load control is used to proactively manage transformer utilization —these distinct use cases were selected to demonstrate the practical value of the PFD metric. The analysis reveals that architectural choices can significantly improve reliability and help determine actionable design principles. Specifically, the findings indicate that systems designed with **more granular** state machines, **shorter cycles** (faster timeouts/loops), and **option paths** (redundant/failover routes) exhibit lower PFD.

Finally, the report outlines a forward-looking vision for a continuously learning reliability model. This iterative, data-driven process provides a roadmap for the industry to move from a state of high uncertainty (treating systems as "black boxes") to one of deep, predictive insight. By progressively incorporating field data, from simple failure counts to detailed internal state transitions, this framework culminates in the creation of an inference engine capable of advanced diagnostics and automated recovery. This approach not only promises to enhance the reliability of individual systems but also enables a powerful, pre-competitive model for industry-wide collaboration through the privacy-preserving sharing of aggregated reliability data.

Table of Contents

Executive Summary **4**
Table of Contents **6**
List of Figures **7**
List of Tables **7**
List of Abbreviations **8**
1. Introduction: Probability of Failure on Demand **9**
1.1 The Customer-Centric View of Charging Success 9
1.2 A Taxonomy of Vehicle Grid Integration Failure Modes 9
1.3 Formalizing The PFD Framework 11
2. Modeling the Dynamics of Charging Systems: The Markov Process Approach **12**
2.1 Representing System Health: States, Transitions, and the Markov Property 12
2.2 Rate of Inherent Failure: Transition Rate Matrix 12
2.3 Modeling Resilience: Autonomous Recovery and Reset Matrix 13
3. From Theory to Practice: Quantifying Architectural Reliability **16**
3.1 Translating Operational Logic into Analytical Models 16
3.2 Case Study 1: The “Day-Ahead Pricing: Architecture 17
3.3 Case Study 2: The “DER Definition” Architecture 18
3.4 Case Study 3: The “Transformer Protection” Architecture 20
3.5 Deriving Design Principles for Robust VGI Systems 23
4. The Future of VGI Reliability: Using a Learning Framework **24**
4.1 Strategic Implications and Recommendations for System Architects and Engineers 25
4.2 Strategic Implications and Recommendations for System Operators and Utilities 26
4.3 Strategic Implications and Recommendations for Policymakers and Standards Bodies 26

List of Figures

Figure 1.Taxonomy of failure modes that affect the probability that the next charge session will be successful 9
Figure 2. Example four-state Markov model with reset cycles 14
Figure 3. State probabilities over time for the Markov model 15
Figure 4: State machine for the Day-Ahead Pricing use case 17
Figure 5: Markov model for the Day-Ahead Pricing use case 18
Figure 6: State machine for the DER Definition use case 19
Figure 7: Markov model for the Day-Ahead Pricing use case 20
Figure 8: State machine for the Transformer Protection use case 21
Figure 9: Markov model for Transformer Protection use case 22

List of Tables

Table 1. VGI System Failure Modes 11
Table 2. Example Four-State Markov Model Definitions 14

List of Abbreviations

Abbreviation	Description
DER	Department of Environmental Regulations
EV	Electric Vehicle
IEEE	Institute of Electrical and Electronics Engineers
MRDI	Minimum Required Diagnostic Information
MREC	Minimum Required Error Codes
MTTF	Mean Time to Failure
PFD	Probability of Failure on Demand
PFMEA	Process Failure Mode and Effects Analysis
SCAR	Successful Charge Attempt Rate
TOU	Time of Use
VGI	Vehicle/Grid Integration

1. Introduction: Probability of Failure on Demand

The success of the transition to electric mobility hinges not only on the availability of charging infrastructure but, more critically, on its consistent and predictable reliability. As charging systems evolve from simple energy dispensers into sophisticated, grid-interactive assets, the definition and measurement of reliability must evolve in tandem. This section establishes the foundational concepts of a new reliability paradigm centered on the customer experience, defines the spectrum of failure modes unique to Vehicle/Grid Integration (VGI), and formalizes the Probability of Failure on Demand (PFD) as the principal metric for its quantification.

1.1 The Customer-Centric View of Charging Success

The ultimate arbiter of charging system quality is the end-user: the Electric Vehicle (EV) driver. For this user, reliability is not an abstract percentage of uptime measured over months but a discrete outcome at a specific moment in time: does the charging session start (and end) successfully when initiated? The ChargeX Consortium encapsulates this customer-centric imperative with the vision of "Any Driver, Any EV, Any Charger," working successfully the "First Time, Every Time".

PFD is introduced as a metric for this customer-centric imperative. It is defined from the user's perspective as the probability that the managed charging system will fail when needed. This definition fundamentally reframes the reliability question. Instead of asking, "What percentage of the time is the system operational?", it asks, "Given that a demand is placed on the system now, what is the probability of an irrecoverable failure?" This is a far more relevant question for a system that is used periodically, such as an EV charger, where the state of the system between demands is of little consequence to the user compared to its state during the critical interval of attempted use.

1.2 A Taxonomy of Vehicle Grid Integration Failure Modes

VGI charging systems, which involve interacting power electronics, embedded software, communication networks, and back-end cloud services, can present a diverse set of potential failure modes, each with different root causes and implications for system modeling. Some commonly used modeling frameworks are outlined in Figure 1 and Table 1.

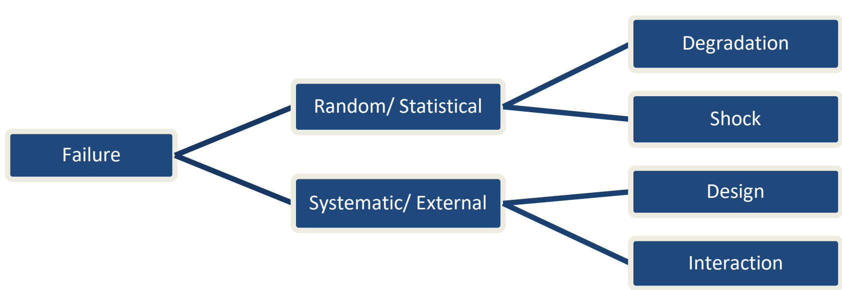

Figure 1.Taxonomy of failure modes that affect the probability that the next charge session will be successful

- **Systematic/External & Design:** These failures are inherent to the system's architecture or its interaction with the external environment. Design flaws, software bugs, or incompatibilities between components fall into this category. While a Process Failure Mode and Effects Analysis (PFMEA) is a traditional tool for identifying such modes, the sheer complexity of a modern VGI ecosystem makes an exhaustive enumeration impractical.

- **Degradation:** This category includes failures that arise from the gradual deterioration of physical or even digital components over time. This could manifest as wear on a charging connector, degradation of a power module, or accumulating errors in a database that eventually lead to system failure.

- **Random/Statistical and Shock:** These are failures that occur without prior warning, representing the low but finite probability that any given charging session could fail irrecoverably. Shock events are often the focus of probabilistic reliability analysis because they are, by nature, unpredictable for any single event but statistically predictable over a large population of events.

- **Interaction Errors:** This category is particularly salient for VGI systems. It encompasses failures that do not originate from a single faulty component but from a failed interaction between multiple, otherwise functional, components (e.g., the EV, the charger, and the central management system). A crucial aspect of these errors is the possibility of recovery. The system may employ autonomous strategies like retries, resets, or failovers. The outcome of such a recovery attempt is not guaranteed; it may fully restore the system to its optimal state, or it may leave it in a degraded state, potentially increasing its susceptibility to future failures until a more significant intervention occurs.

This last point on interaction errors and degraded recovery states acknowledges that failure in a complex system is not always a simple, binary event. A system can exist in a partially failed or degraded state where it "mostly works" but is less resilient or performant. A single error, even if seemingly recovered from, can therefore have a cascading impact on future reliability. A robust reliability framework must be capable of modeling not just the probability of an initial failure but also the long-term consequences of imperfect recovery mechanisms. This provides a more nuanced and realistic assessment of overall system health.

Table 1. VGI System Failure Modes

Failure Mode Category	Definition	VGI-Specific Example
Systematic/Design	Failures inherent in the system's architecture or logic.	A charger's firmware is incompatible with a new EV model's communication protocol, leading to a handshake failure.
Degradation	Failures resulting from the gradual worsening of components over time.	The insulation on a high-voltage cable wears down after multiple flex cycles, eventually causing a ground fault.
Random/Statistical (Shock)	Irrecoverable failures that occur with a low probability and without prior indication.	The failure of a defective electronic component causes the control board to crash irrecoverably during a session.
Interaction Error	Failures arising from the complex interplay of cyber and physical components.	A transient network packet loss causes a timeout in the communication between the charger and the backend server, interrupting a managed charging command. A software retry may or may not succeed.

1.3 Formalizing The PFD Framework

The PFD framework assumes that the charging system is used periodically, with demand for charging occurring within a time interval of length τ. The state of the system at any time during the n^{th} charging session where t ∈ $[(n-1)\tau,\ n\tau]$ is denoted by $X(t)$. The complete set of all possible system states is represented by a finite set $\mathbb{X}$ = {0, 1, . . . , r }.

This total state space, $\mathbb{X}$, is partitioned into two mutually exclusive subsets:

$\mathbb{B}$: The set of all functioning states, which may include optimal, degraded, or post-recovery operational states.

$\mathbb{F} = \mathbb{X} - \mathbb{B}$: The set of all failed states, where the system is unable to perform its required function.

With these definitions, the Probability of Failure on Demand for the n^{th} charging session is formally defined as the average probability of being in a failed state over the demand interval:

$$PFD(n) = \frac{1}{\tau}\int_{(n-1)\tau}^{n\tau} \Pr\left(X(t) \in \mathbb{F}\right) dt$$

The term $\Pr\left(X(t) \in \mathbb{F}\right)$ represents the instantaneous probability that the system is in a failed state at time t. By integrating this probability over the interval τ and then averaging, the PFD captures the likelihood of encountering a failed system at any random point within the window of time that a user might initiate a charge. This integral formulation is superior to a simple point-in-time probability because it accounts for the fact that failures can occur and resolve over time, providing a more holistic measure of the risk faced by the user.

2. Modeling the Dynamics of Charging Systems: The Markov Process Approach

To calculate the PFD, one must first be able to compute the probability of the system being in any given state (faulted or otherwise) at any given time. For a system as complex as grid-responsive charging, we consider the continuous-time Markov process to capture the dynamic, probabilistic nature of state transitions to be a well-established mathematical tool for modeling the evolution of stochastic systems. This section details this approach, from the foundational concepts of states and transitions to the specific matrices used to represent failure and recovery.

2.1 Representing System Health: States, Transitions, and the Markov Property

The core of the methodology is to model the charging system's lifecycle as a journey through a finite set of discrete states. A continuous-time Markov process is used to describe this journey. The defining characteristic of the Markov property is that the future evolution of the system depends only on its present state, not on the sequence of events that preceded it. Mathematically, this is expressed as:

$$\mathbf{Pr}\big[X(n\tau)|X\big((n-1)\tau,\ \cdots,X(0)\big)\big] = \mathbf{Pr}\big[X(n\tau)|X\big((n-1)\tau\big)\big]$$

This assumption of memorylessness, while not perfectly representative of all physical processes (e.g., some forms of connector degradation have memory), it is an effective memory-efficient approximation for systems dominated by random failures and state-dependent transitions.

Within this framework, the state space is partitioned into functioning states ($\mathbb{B}$) and failed states ($\mathbb{F}$). The failed states are often "absorbing," meaning that once the system enters such a state (e.g., a critical hardware failure), it cannot leave without an external intervention like a repair, which is outside the scope of the immediate operational model. The PFD calculation, therefore, becomes a matter of tracking the cumulative probability of the system entering one of these absorbing failure states over time.

2.2 Rate of Inherent Failure: Transition Rate Matrix

The dynamics of how a system moves between states due to inherent failure mechanisms are captured by a set of transition rate operators, denoted as λ_{ij}. This operator represents the instantaneous rate at which the system transitions from state x_i to state x_j.These individual rates are assembled into a transition rate matrix ($\mathbb{A}$).

$\mathbb{A}$ encodes the rates of degradation, the likelihood of random shocks, and the pathways by which minor issues can cascade into critical failures. The diagonal elements in $\mathbb{A}$ are typically negative and represent the total rate of leaving state x_i, while the off-diagonal elements are the positive rates of all transitions into x_i.

2.3 Modeling Resilience: Autonomous Recovery and Reset Matrix

A charging system is not merely a passive entity succumbing to failure; it is an active system designed with a degree of resilience. It can attempt to recover from transient faults through autonomous actions like communication resets, alternate protocols or session retries. These recovery mechanisms are a fundamentally different process from inherent failure and are therefore modeled separately using a reset/restoration transition matrix ($\mathbb{R}$).

$\mathbb{R}$ describes probabilities that the system successfully transitions to a functioning state given that it entered a recoverable error state and a reset was attempted. Consistent with the Markov modelling framework, $\mathbb{R}$ quantifies the "statistical effectiveness" of the system's built-in resilience strategies, abstracting away the specific technical details of the recovery methods themselves.

This conceptual separation of failure physics from recovery logic allows the two major aspects of reliability to be analyzed and optimized independently. Hardware engineers and component suppliers can focus on improving Mean Time To Failure (MTTF), which would reduce the transition rates in the failure matrix. Simultaneously, software and systems architects can focus on designing more robust retry loops and failover protocols, which would improve the success probabilities in the matrix. This decoupling provides a formal framework for evaluating critical design trade-offs. For instance, an organization can quantitatively assess whether it is more cost-effective to invest in a more reliable hardware component or in more sophisticated recovery scheme.

2.4 A Markov Failure Model: Simple Four-State Charging System.

To illustrate these concepts, Figure 2 shows a simple four-state Markov model of a managed charging system. This model serves as a concrete example of how the abstract matrices and state definitions are applied in practice.

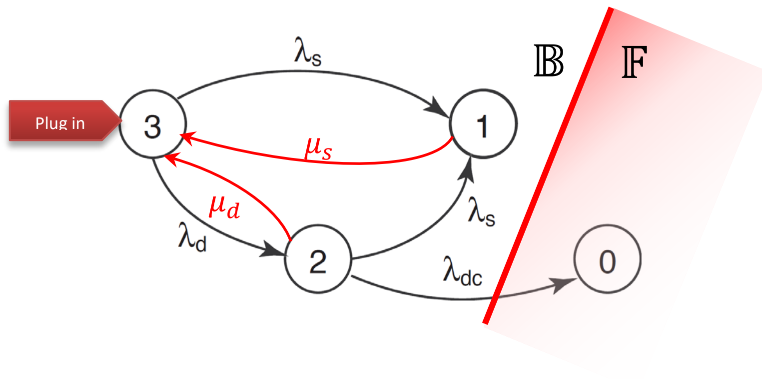

Figure 2. Example four-state Markov model with reset cycles

Table 2. Example Four-State Markov Model Definitions

Component	Description
State 3	System works perfectly, as designed.
State 2	Mostly worked; recovery, fallback, or degraded operation occurred.
State 1	Shock failure; random failure with no prior indication.
State 0	Degradation failure; degradation worsens over time to cause a full failure.
Rate Operator λ_s	Rate of failures caused by random shock.
Rate Operator λ_d	Rate of degradation failures (e.g., from State 3 to State 2).
Rate Operator λ_{dc}	Rate of degraded failures that become critical (e.g., from State 2 to State 0).
Rate Operator μ_s	Reset probability from a random shock failure.
Rate Operator μ_d	Recovery probability from a recoverable degraded operating state.

In this model, States 3 and 2 belong to the set of functioning states, while States 1 and 0 belong to the set of failed states. The transition rate matrix combines the failure matrix and reset matrix for this system as is given as:

$$\mathbb{A} * \mathbb{R} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & -\mu_s & 0 & \mu_s \\ \lambda_{dc} & \lambda_s & -(\lambda_{dc} + \lambda_s + \mu_d) & \mu_d \\ 0 & \lambda_s & \lambda_d & -(\lambda_s + \lambda_d) \end{pmatrix}$$

The rows and columns correspond to states 0, 1, 2, and 3. The first row of all zeros indicates that State 0 is an absorbing failure state - once entered, it cannot be left. The second row reflects a fault reset probability (between states 1 and 3) of μ_s. The third row shows that from the degraded state (State 2), the system can transition to the critical degradation failure state (State 0) at a rate of λ_{dc} or to the shock failure state (State 1) at a rate of λ_s. The fourth row shows that from the perfect operational state (State 3), the system can transition to the shock failure state (State 1) at a rate of λ_s or to the degraded state (State 2) at a rate of λ_d. This simple model provides all the necessary components to simulate the system's behavior over time and ultimately compute its PFD.

PFD for a Markov failure model is simply the probability of being in a failed state at time t: $P_{\mathbb{F}}(t)$. For the model in Figure 2, the probability of being in State 0, which is the only element in $\mathbb{F}$, is given by the relation $P_{\mathbb{F}}(t) = \frac{\lambda_{dc}}{\lambda_s+\lambda_{dc}}\left(1 - e^{-(\lambda_s+\lambda_{dc})t}\right)$. We can then integrate $P_{\mathbb{F}}(t)$ over the utilization interval τ to get average PFD.

Applying transition probabilities of $\lambda_s = .5$, $\lambda_d = .4$, $\lambda_{dc} = .2$, $\mu_s = .3, \mu_d = .6$ to the model in Figure 2, we can simulate the Markov model to analyze the state probabilities over time as illustrated in Figure 3. Here we assume that every user plugs into a charger which is in State 3 since we only consider failures on demand. The figure shows the dynamics of the Markov model over time demonstrating that for the given transition probabilities eventually the system will fail to State 0 since it is an absorbing state. We can define PFD as the average likelihood that a system will fail if a demand occurs randomly during a utilization interval τ= 20-time units. With these assumptions, PFD for this illustrative model is 32%.

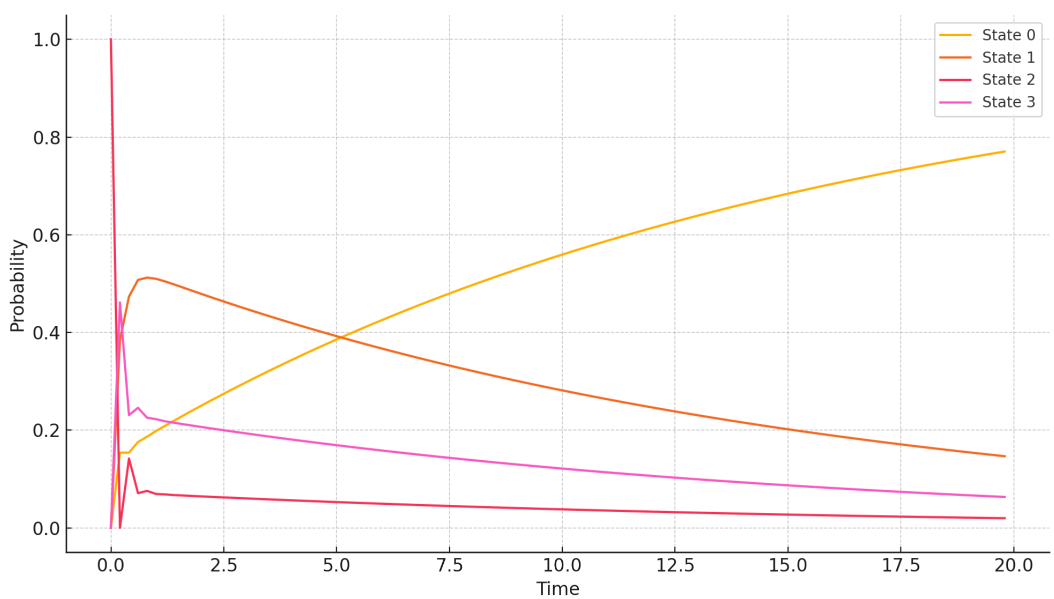

Figure 3. State probabilities over time for the Markov model

3. From Theory to Practice: Quantifying Architectural Reliability

The true value of the PFD framework lies in its ability to move beyond theoretical models and provide quantitative, actionable insights into the reliability of real-world system designs. It provides a structured methodology to translate the operational logic of a charging process into a mathematical model that can be analyzed and compared against alternatives. This section details this translational process and applies it to a case study comparing three different VGI system architectures, deriving key design principles for building more robust systems.

3.1 Translating Operational Logic into Analytical Models

The application of the PFD framework to a specific system architecture follows a four-step process that bridges the gap between practitioner knowledge and formal analysis.

1. **Sequence Diagrams:** The process begins by capturing the operational workflow of a VGI function in a sequence diagram. These diagrams are a standard tool in software engineering for visualizing the interactions between different system components over time. They are an effective way to document the knowledge and experience of practitioners regarding communication protocols, such as those required for responding to a day-ahead electricity tariff or setting up a Distributed Energy Resource (DER) for grid service or responding to transformer load alerts.

2. **State Machines:** The sequence diagram is then formalized into a state machine. This involves identifying the discrete states the system occupies during the workflow (e.g., "Idle," "Communication Setup," "Receiving Pricing Data") and mapping the transitions between them. This step transforms the procedural view of the sequence diagram into a structural graph of states and their connections.

3. **Markov Chains:** The state machine is converted into a quantifiable continuous-time Markov chain. This is achieved by assigning transition rates to the transitions that represent potential failures and recovery probabilities (e.g., "Comms Reset," "Session Reset"). The result is a complete mathematical model of the architecture, ready for PFD computation.

4. **Obtaining Transition Probabilities:** Clearly, the accuracy of transition probabilities/rates is essential to compute a meaningful PFD. However, even in cases where transition rates have not been empirically validated and average values are used, the PFD metric can provide useful insight when comparing different architectural designs by quantifying the relative impact of individual changes to system design.

We can now apply our analysis technique to three VGI use cases identified by ChargeX.

3.2 Case Study 1: The “Day-Ahead Pricing: Architecture

The first architecture analyzed represents a system designed to handle Time-of-Use (TOU) or day-ahead pricing signals. Its state machine, shown in Figure 4 and derived from a comprehensive sequence diagram developed but not included in this report, involves a sequence of states including "Communication Setup," "Transaction Started," "Receiving Pricing," "Optimization," and "Charging Adjustment".

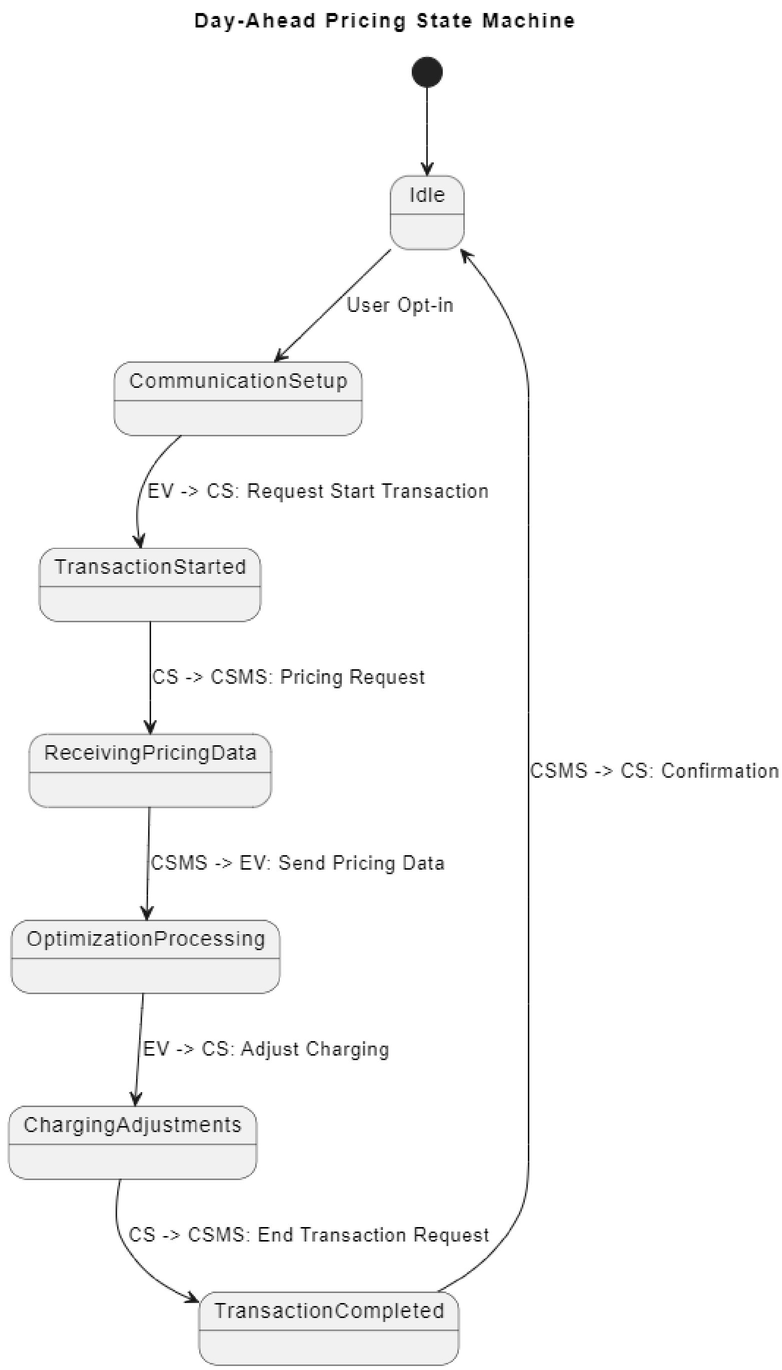

Figure 4: State machine for the Day-Ahead Pricing use case

The state machine is augmented with several recovery paths, such as a "Comms Reset" and a "Session Reset," which offer opportunities for the system to recover from transient faults to generate a Markov model shown in Figure 5.

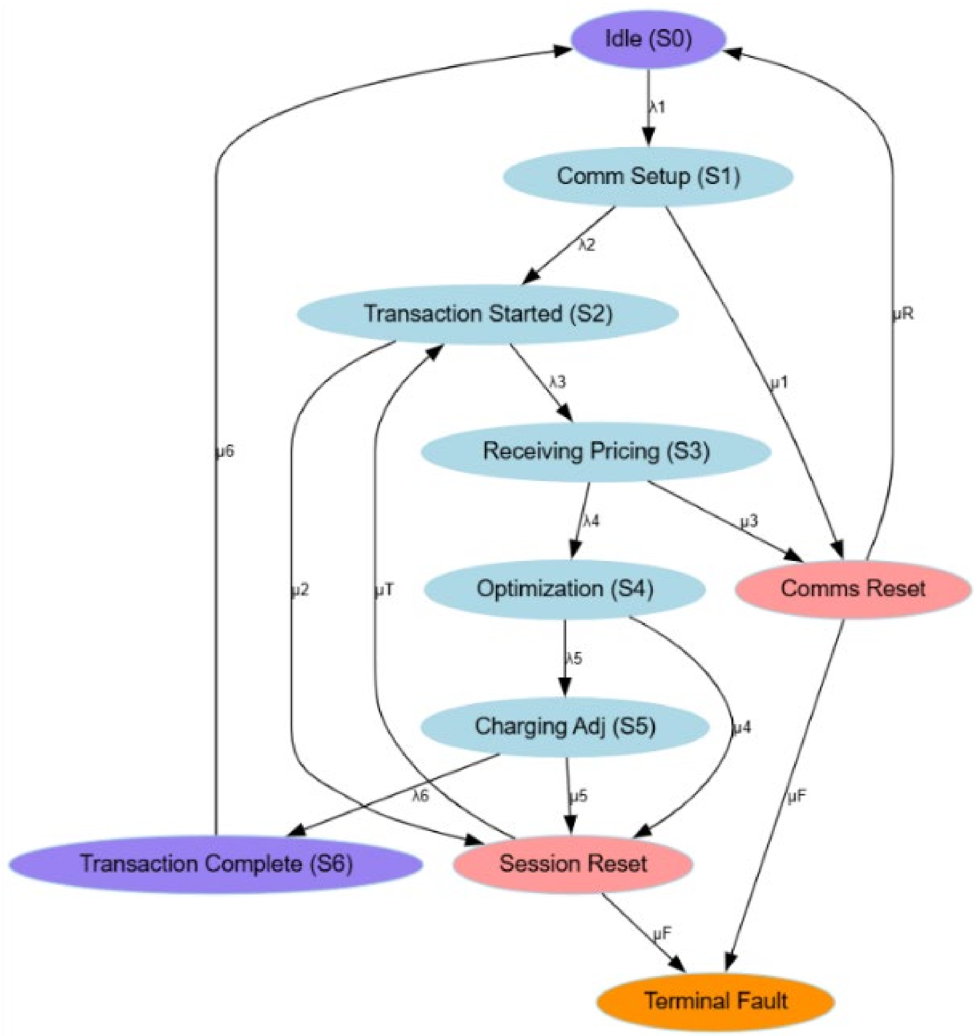

Figure 5: Markov model for the Day-Ahead Pricing use case

3.3 Case Study 2: The "DER Definition" Architecture

The second architecture focuses on the grid interface required to set up a managed charging capability, treating the EV charger as a DER. We can derive a state machine from the sequence diagram we developed (not included in this report) with states like "Fetching ProgramList," "Subscribing," "Fetching ControlList," that capture the major conditional operating modes used to properly configure charging system to provide DER services per the IEEE 2030.5 protocol. This state machine is shown in Figure 6.

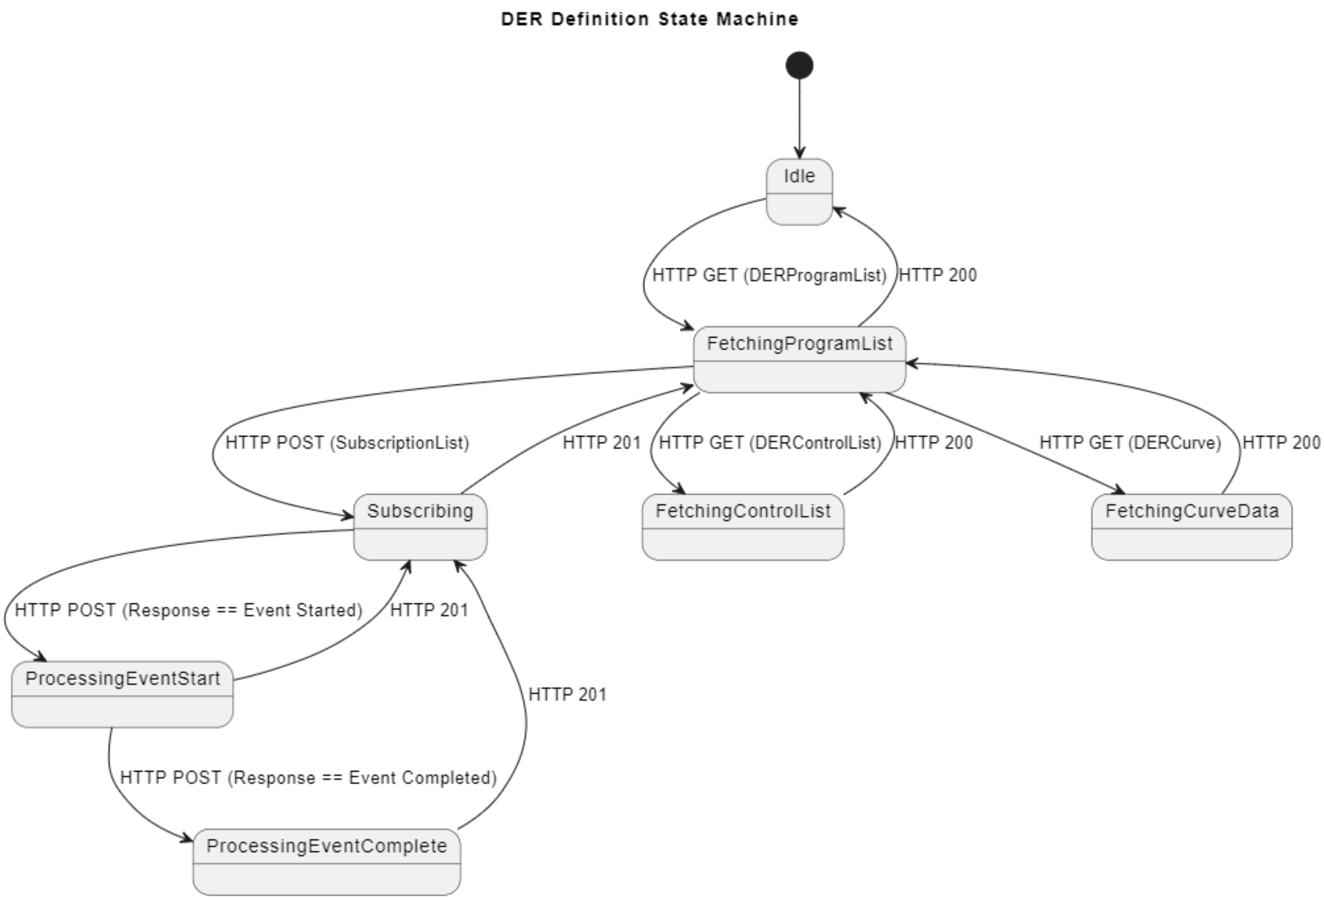

Figure 6: State machine for the DER Definition use case

Visually, this state machine appears to break the overall process into more discrete, independent steps compared to the first case study. The Markov model for this case study illustrates an architecture with multiple session and communications reset opportunities.

The same simulation methodology using normalized transition rate values was applied to this second case study. The PFD calculated over 1000 trials was approximately 1.84%. The key observation between case study 1 and 2 is the order of magnitude improvement in PFD demonstrating that even if component reliabilities are unknown, system architecture can greatly impact reliability.

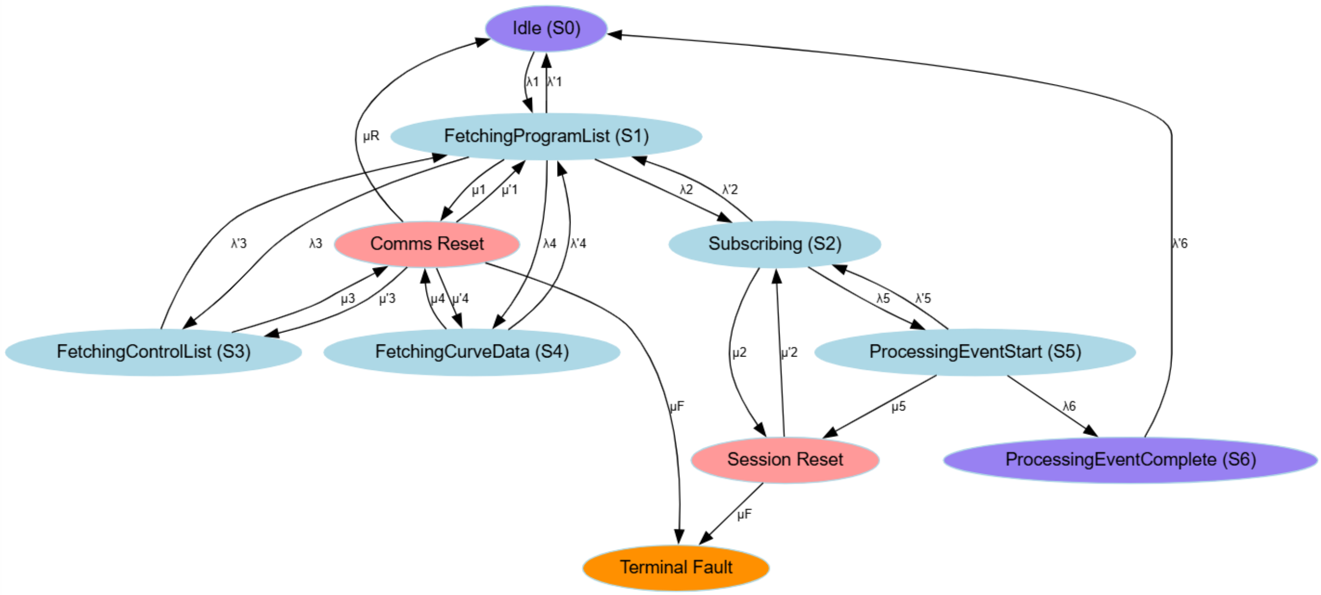

Figure 7: Markov model for the Day-Ahead Pricing use case

3.4 Case Study 3: The "Transformer Protection" Architecture

Finally, we consider a relatively complex architecture for a VGI system that can manage charging loads to avert transformer overloading. This case study is selected for study to highlight how the PFD calculation can be conducted from different starting states reflecting the needs of different stakeholders. In this example, PFD can be represented as the probability of failure that an EV is not able to initiate and complete a charge session as in previous case studies. The same Markov model can also be used to compute the probability that the charging load does not successfully respond to a transformer alert condition when demanded. The tradeoff in reliability between these two perspectives is obvious in the extreme case; if a transformer trips off to avert an overload state, the outcome is successful for the transformer but a failure for the charging system. The state machine and Markov model for this use case are shown in Figures 8 and 9, respectively.

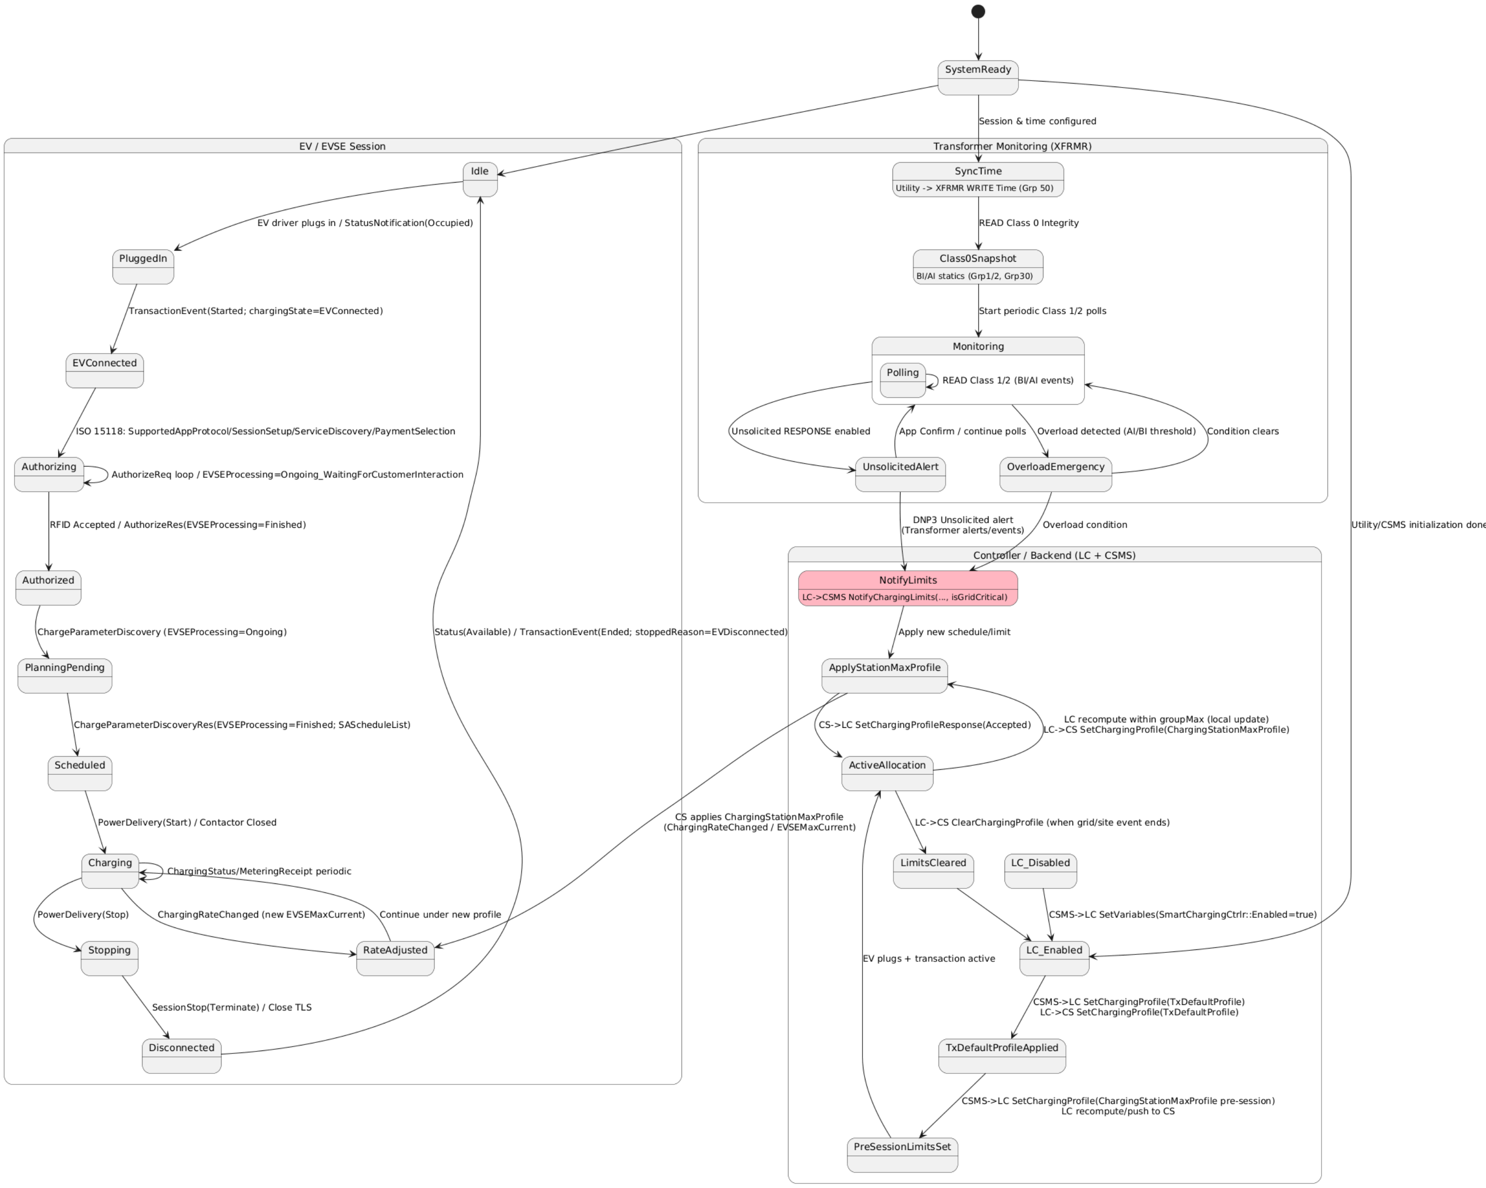

Figure 8: State machine for the Transformer Protection use case

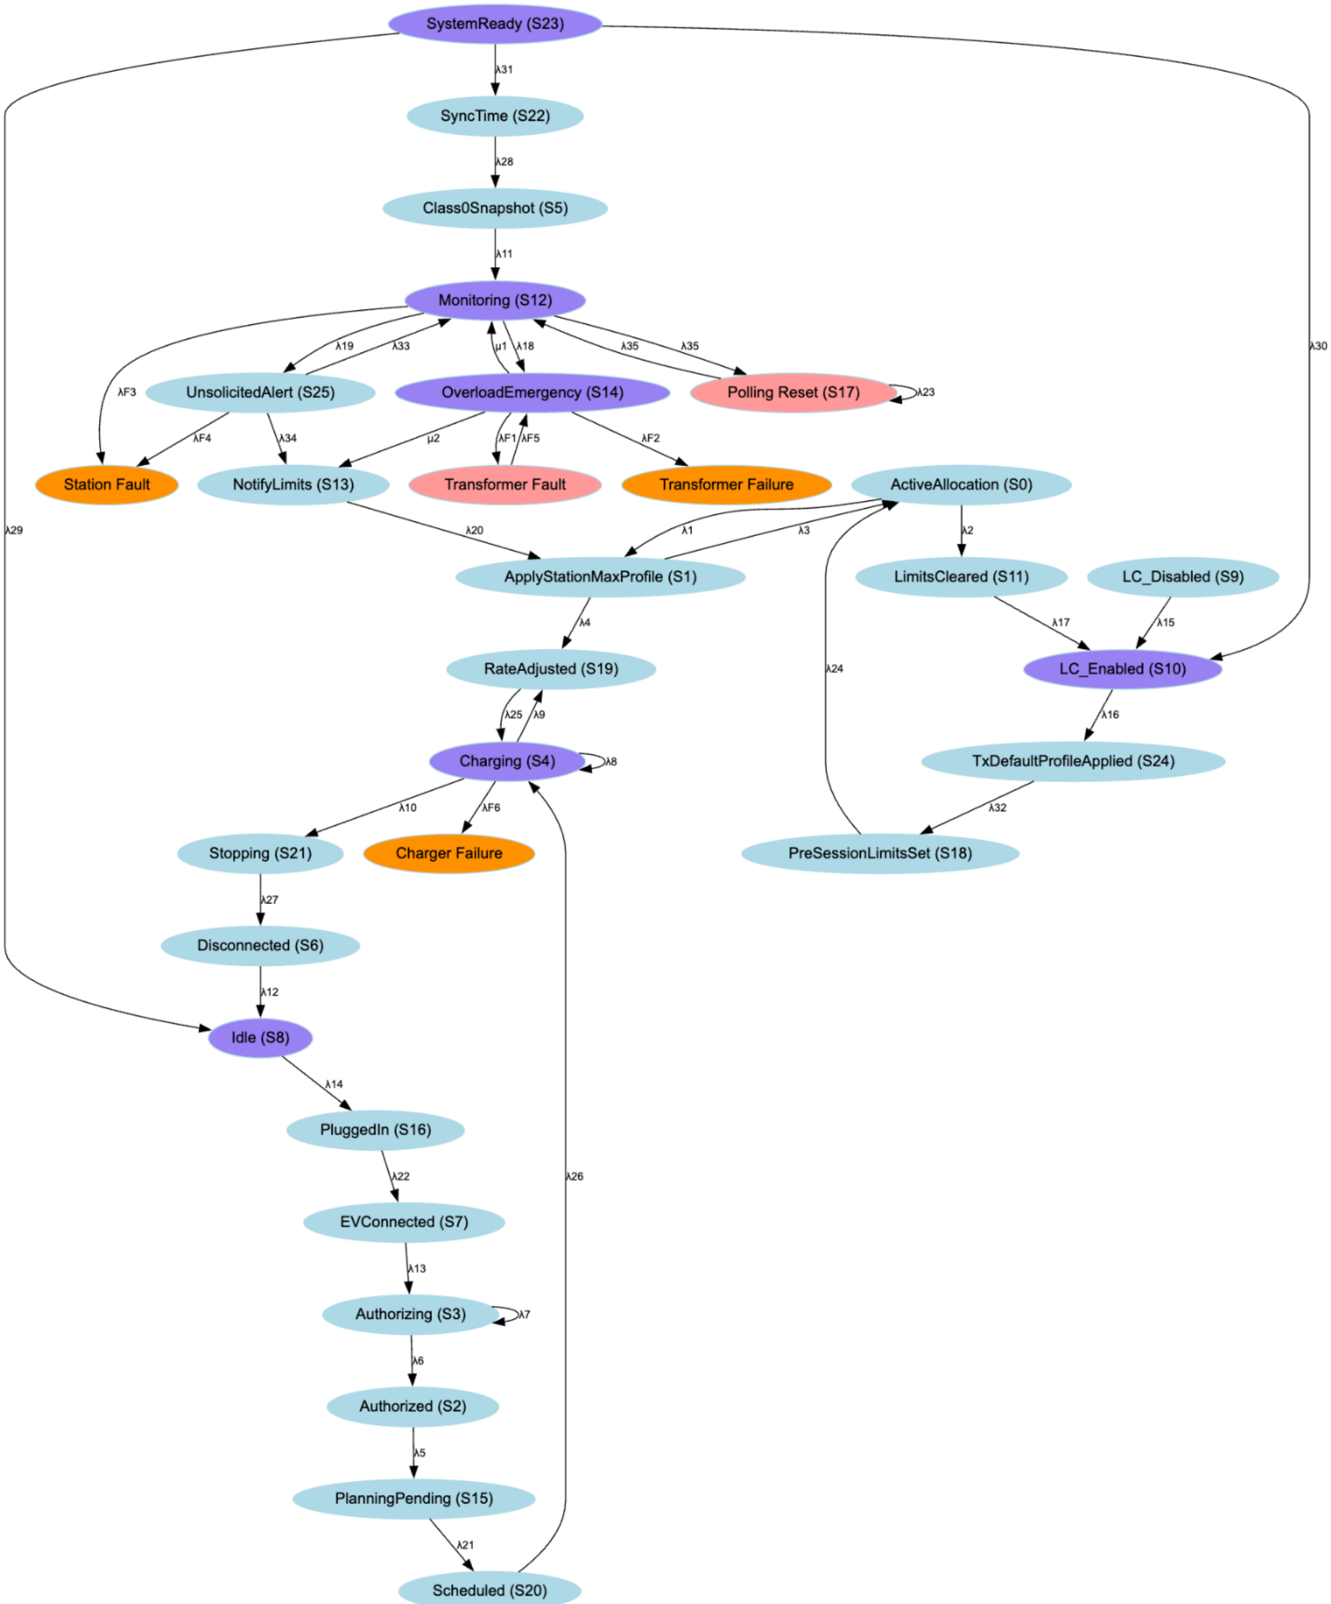

Figure 9: Markov model for Transformer Protection use case

Repeating the 1000 sample PFD analysis on this Markov model we can set the starting state to be S8 to represent the condition of the system prior to an EV plug-in event. In this case the computed PFD is 17.57%. However, performing the same analysis from the state S0 to focus on the control pathway between the transformer management system and the charging system the computed PFD is significantly higher at 32.22%. In effect, for this architecture, it is twice as likely that EV charging loads will fail to adequately respond to a transformer overload alert signal than the likelihood that the EV charging service will fully fail.

It is important to reiterate that all three of the analyses presented here are purely for illustrative purposes and none of the transition rates or architectural designs are based on or vetted against real systems. The selected case studies are intended to demonstrate the capabilities of the Markov model and PFD approach and to show how they may be applied to derive reliability insights about VGI systems.

3.5 Deriving Design Principles for Robust VGI Systems

The quantitative comparison of the three architectures allows for the extraction of generalizable design principles for creating highly reliable systems. The superior performance of the DER Definition architecture is not arbitrary; it is a direct consequence of its structural properties. The analysis explicitly identifies key characteristics of a more robust system architecture:

- **Granularity:** A more granular state machine is preferable. Breaking a complex, monolithic process into a larger number of smaller, simpler states can improve reliability.
- **Cycle Length:** Shorter operational cycles are better than longer ones. This suggests that processes that can be completed and verified quickly are less susceptible to failure than long-running, uninterruptible operations.
- **Optional Paths:** The inclusion of alternative paths and recovery loops is good for robustness. Providing multiple ways for a system to recover from an error state significantly reduces the probability of that error leading to a terminal failure.
- **Evaluating Multiple Subsystems:** For systems that have multiple independent subsystems, each with its own reliability objective improving PFD for one stakeholder may adversely impact another. Modeling multiple interacting subsystems as a Markov model enables analysis from multiple stakeholder/user perspectives.

The quantitative validation of these principles is a powerful result. It confirms a core tenet of resilient software and systems engineering: reliability is an emergent property of architectural structure, with modularity and fault isolation being key drivers. The DER Definition architecture, with its lower PFD, exemplifies this. By decomposing the complex task of grid integration into a series of smaller, independent, and individually recoverable steps (e.g., "Fetching ProgramList"), it creates a system where a failure in one step does not necessarily cascade to a total system failure. A fault in an early stage can be retried or trigger a specific recovery loop (like "Comms Reset") without terminating the entire managed charging session. In contrast, a more monolithic design is inherently more brittle, as a single point of failure within a long, uninterruptible state is more likely to lead directly to a "Terminal Fault." The PFD calculation, therefore, serves not just as a final score but as a quantitative tool that validates the profound impact of architectural choices on end-to-end system reliability.

4. The Future of VGI Reliability: Using a Learning Framework

While the Markovian PFD framework provides a powerful tool for analyzing and comparing static system designs, its ultimate potential is realized when it is transformed into a dynamic system that learns and adapts based on real-world field data. A forward-looking vision for this evolution is a staged, iterative process for progressively refining a reliability model, moving from a state of high uncertainty (a "black box") to one of deep, actionable, and predictive insight about the components in the system and by extension the transition rates in the Markov model.

At the outset, reliability metrics for the “black box” system such as the ‘Successful Charge Attempt Rate’ (SCAR) considered by bodies like the California Energy Commission could be used to update the normalized transition rates to ensure that the objective value of the PFD metric is benchmarked against real world statistics.

However, the Markov framework could go a step further by using real-time fault reports to inform a Bayesian inference model. By assuming that the failure outcomes follow a Beta distribution - a natural choice for modeling the probability of a binary event - one can continuously update a probabilistic model of the system's failure rate based on real-time error flags and system alerts. This approach is superior to simply calculating the raw failure rate because it yields a full probability distribution, which captures the uncertainty in the estimate. As more data is collected, this distribution becomes sharper and more precise, representing a "learned" model of the system's black-box failure probability that can be used to inform an initial PFD calculation.

The next stage of maturity involves moving beyond simple success/failure counts to understand the internal workings of the system. This is achieved by introducing "tracer signals" - rich data points from within the system, such as Minimum Required Error Codes (MRECs) and Minimum Required Diagnostic Information (MRDI) - and by mapping the sequences of data exchanges between components. This richer dataset allows for the creation of multivariate statistical models.

Artificial intelligence and machine learning methods can be applied to this stream of observational data to automatically learn the most likely values for the transition rates in the Markov model. Data-driven models may then be augmented with a crucial additional layer: documented diagnostic insights from human experts. Experienced engineers can identify and describe nuanced failure modes and their root causes that might be difficult to discern from telemetry data alone. Integrating this qualitative expert knowledge with the quantitative data-driven model creates a hybrid intelligence system that is more powerful than either approach on its own.

The culmination of a data-validated and expert-augmented Markov model is its value as operational tool. Used not only for offline analysis and prediction of system reliability but also for real-time applications such as predictive maintenance and even triggering automated recovery actions when the system is detected to be entering a high-risk state.

This evolutionary framework provides a practical roadmap for creating what is effectively a "Digital Twin" specifically for VGI system reliability. A Digital Twin is a virtual model of a physical asset, continuously updated with real-world data, used for simulation, prediction, and optimization.

Furthermore, this framework could potentially serve as a model for industry-wide collaboration to improve reliability. Raw operational data and event logs are often proprietary and contain sensitive business or customer information. However, transition rate matrices are aggregated statistical models of behavior. They represent the *rates* of failure and *probabilities* of recovery, not the specific details of any single event. This abstraction allows competing organizations - such as Charge Point Operators, EV Manufacturers, and Network Providers - to contribute their reliability data in the form of these anonymized matrices. This would enable industry-wide benchmarking, the identification of common, cross-platform failure patterns, and the collaborative development of best practices, without requiring any organization to reveal its proprietary operational data. It transforms reliability from a siloed, competitive struggle into a shared, pre-competitive challenge that the entire VGI ecosystem can work to solve together.

4.1 Strategic Implications and Recommendations for System Architects and Engineers

The individuals and teams responsible for designing the hardware, software, and communication protocols of VGI systems, the PFD framework offers a powerful new set of tools for design and validation.

Recommendation: Adopt PFD as a primary design metric alongside traditional metrics like cost, performance, and efficiency. Utilize the Markov modeling approach early in the design phase to conduct quantitative "what-if" analyses. By modeling different architectural choices and comparing their resulting PFDs, engineers can make data-driven decisions that optimize for reliability. This analysis provides a mandate to prioritize designs that are granular and modular, feature short operational cycles, and incorporate multiple, robust, and well-defined recovery paths.

4.2 Strategic Implications and Recommendations for System Operators and Utilities

The entities that own, operate, and manage networks of charging infrastructure, including utilities that rely on VGI for grid services, have a vested interest in the long-term, fleet-level reliability of these assets.

Recommendation: Champion the implementation of a data collection framework across deployed assets. The journey can begin simply, by systematically collecting basic success/failure data for all charging sessions to establish a baseline reliability model. Over time, operators should work with vendors to evolve their systems toward greater instrumentation, enabling the capture of the internal state transition data needed to generate high-fidelity Markov chains and transition matrices. The resulting models will be valuable for predictive maintenance and improve the speed and accuracy of fault diagnosis, reducing downtime and operational costs, and ultimately enhancing the reliability of the entire managed charging network.

4.3 Strategic Implications and Recommendations for Policymakers and Standards Bodies

Regulators and organizations that set industry standards play a crucial role in establishing the baseline expectations for reliability and interoperability in the VGI ecosystem. The PFD framework offers a path toward more sophisticated and effective reliability standards.

Recommendation: Evolve reliability standards beyond simple, static, pass/fail metrics like SCAR. While such metrics are a useful starting point, they lack the nuance to capture the dynamic and architectural aspects of system robustness. Policymakers should consider incorporating requirements for model-based reliability reporting, potentially leveraging the PFD framework as a standardized methodology. This would provide a much deeper and more forward-looking view of system reliability. Furthermore, these bodies should actively foster industry collaboration by supporting initiatives, like those proposed by the ChargeX Consortium, that create secure and privacy-preserving mechanisms for sharing aggregated reliability models. Such collaboration is essential for accelerating the learning curve of the entire industry and building a universally reliable charging experience.

About the ChargeX Consortium

The National Charging Experience Consortium (ChargeX Consortium) is a collaborative effort between Argonne National Laboratory, Idaho National Laboratory, National Laboratory of the Rockies, electric vehicle charging industry experts, consumer advocates, and other stakeholders. The ChargeX Consortium's mission is to work together to measure and significantly improve public charging reliability and usability by June 2025. For more information, visit chargex.inl.gov.

MARCH 2026

INL/RPT-26-91549

Argonne National Laboratory

Idaho National Laboratory

National Laboratory of the Rockies