



Cambium Documentation: Version 2020

Pieter Gagnon, Will Frazier, Elaine Hale, and Wesley Cole

National Renewable Energy Laboratory

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-6A20-78239
November 2020



Cambium Documentation: Version 2020

Pieter Gagnon, Will Frazier, Elaine Hale, and Wesley Cole

National Renewable Energy Laboratory

Suggested Citation

Gagnon, Pieter, Will Frazier, Elaine Hale, and Wesley Cole. 2020. *Cambium Documentation: Version 2020*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-6A20-78239. <https://www.nrel.gov/docs/fy21osti/78239.pdf>.

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-6A20-78239
November 2020

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Strategic Priorities & Impact Analysis Office and Building Technologies Office. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Acknowledgments

We gratefully acknowledge the many people whose efforts contributed to this report. We thank Jamy Bacchus, Shanti Pless, Chioke Harris, Jared Langevin, Stephanie Ross, and Ryan Hledik for their helpful perspectives while this document was being written, and Andrew Mills, Andy Satchwell, Matt Irish, Aven Satre Meloy, and Jesse Cohen for their comments during drafting. Finally, we are grateful to Mike Meshek for his diligence in editing. This report was funded by the DOE Office of Energy Efficiency and Renewable Energy under contract number DE-AC36-08GO28308. Any errors or omissions are the sole responsibility of the authors.

List of Abbreviations and Acronyms

AC	alternating current
BA	balancing area
CES	clean energy standard(s)
CSP	concentrating solar power
LRMER	long-run marginal emission rate
NGCT	natural gas combustion turbine
NREL	National Renewable Energy Laboratory
PRM	planning reserve margin
PV	photovoltaics
RAZ	reliability assessment zone(s)
REC	renewable energy credit
RPS	renewable portfolio standard(s)
RTE	round-trip efficiency
SRMC	short-run marginal costs
SRMER	short-run marginal emission rate

Table of Contents

1	Cambium Overview	1
1.1	ReEDS	1
1.2	PLEXOS	2
2	User Guidance, Caveats, and Limitations of Cambium Databases	4
3	Cambium Metric Definitions	7
3.1	Busbar and End-Use Values	7
3.2	Time and Geographic Identifiers	7
3.3	Generation and Capacity Metrics	9
3.4	Emission Metrics	10
3.5	Cost Metrics	13
3.6	Interregional Transmission Metrics	16
3.7	Load Metrics	17
3.8	Operational Metrics	18
3.9	Portfolio Standard Metrics	22
3.10	Marginal Generators and Marginal Energy Sources	23
4	Cambium Methods	25
4.1	Technologies represented in Cambium	25
4.2	Coloring Power Flows	26
4.3	Calculating Long-Run Marginal Emission Rates	29
4.4	Identifying a Region’s Marginal Generator	35
4.5	Identifying the Energy Source When an Energy-Constrained Generator is on the Margin	39
4.6	Calculating Time-Varying Distribution Loss Rates	44
4.7	Calculating Hourly Marginal Capacity Costs	46
4.8	Calculating Marginal Portfolio Costs	50
4.9	Operating Reserve Requirements and Limitations	52
5	References	55

List of Figures

Figure 1. Balancing areas and reliability assessment zones used in Cambium, ReEDS, and PLEXOS.....	9
Figure 2. Simple network for illustrating power flow coloring	27
Figure 3. Charge and discharge patterns of an electrical battery	41
Figure 4. Example of an “opportunity window” for an electric battery.....	42
Figure 5. Examples of histograms of hourly distribution loss rates (average and marginal), drawn from the Mid-case of the 2020 Standard Scenarios	46

List of Tables

Table 1. Cambium Technologies	25
Table 2. In-Region Generation and Transmission Flow Between Two Nodes.....	30
Table 3. Generation Allocated to Two Nodes.....	31
Table 4. Efficiency Adjustments.....	43
Table 5. Operating Reserve Requirements	52
Table 6. Technology Flexibility Parameters	53

1 Cambium Overview

Cambium is a tool that assembles structured data sets of simulated hourly cost and operational data for modeled futures of the U.S. electric sector with metrics designed to be useful for long-term decision-making. It was built to expand the metrics reported in the National Renewable Energy Laboratory's (NREL's) Standard Scenarios—an annually released set of projections of how the U.S. electric sector could evolve across a suite of different potential futures, looking forward through 2050 (Cole et al. 2020). The Cambium data sets for the Standard Scenarios can be viewed and downloaded at <https://cambium.nrel.gov/>.

In this documentation, we define the metrics reported in Cambium databases (Section 2) and document the Cambium-specific methods for calculating those metrics (Section 3). As this document is intended to cover multiple Cambium data releases, we do not document specific scenarios here—for individual data releases, readers should look for an accompanying report that describes the assumptions and data that underlie that specific release. For Cambium databases that are built on NREL's Standard Scenarios, for example, reports describing the scenarios can be found on the Standard Scenarios web site (<https://www.nrel.gov/analysis/standard-scenarios.html>).

Cambium draws primarily from the outputs of two models:¹

- The Regional Energy Deployment System (ReEDS) model, which uses a least-cost framework to project structural changes in the U.S. electric sector under different possible futures (Brown et al. 2020).
- PLEXOS, which is a commercial production cost model that we use to simulate the hourly operation of the future electric systems projected by ReEDS (Energy Exemplar 2019).

Cambium transforms the outputs produced by these models into a structured data set that conveys future electric grid (grid) operations and value streams.

1.1 ReEDS

The first of two models that Cambium draws from is ReEDS (Brown et al. 2020).² ReEDS is an NREL-developed, publicly available mathematical programming model of the electric power sector. Given a set of input assumptions such as fuel costs, technology costs, and policies, ReEDS models the evolution of generation and transmission assets, solving a linear program to make investment and operational decisions to minimize the overall cost of the electric system. The model has been used to explore how the evolution of the electric sector is impacted by a range of technology and policy scenarios.

¹ We briefly summarize ReEDS and PLEXOS in this section, and we refer readers to the literature cited in each of those sections for documentation of how each model functions.

² Brown et al. (2020) documents the 2019 version of ReEDS. More information about ReEDS, including the most up-to-date documentation, open-source access, and other publications can be found at <https://www.nrel.gov/analysis/reeds/>.

The conterminous United States (i.e., the lower 48 states and the District of Columbia) is represented in ReEDS as 134 model balancing areas (BAs), which are connected by a representation of the transmission network. The network starts with existing transmission capacity and can be expanded as part of ReEDS' decision-space. Likewise, the model starts with representations of all existing generation capacity and announced future builds for each BA, and it can choose to build new capacity or retire old capacity to meet demand at the lowest cost. Historical patterns are used as a starting point for assumptions about end-use electricity demands, and assumptions about the evolution of that demand vary by scenario.

The linear program for balancing supply and demand within ReEDS includes a representative set of time-slices that are meant to capture seasonal and diurnal trends. A submodule, Augur, is used to calculate key parameters from hourly data, where the time-slice representation is too coarse. The ability of variable generators (e.g., wind and solar) and storage to contribute firm capacity, curtailment of variable generators, and the energy arbitrage value of storage generators are examples of parameters derived in Augur.

The linear program that forms the core of ReEDS makes investment and retirement decisions for bulk power system assets. For behind-the-meter solar photovoltaics (PV), the model imports projections from NREL's Distributed Generation Market Demand Model (dGen, [Sigrin et al. 2016]).³

1.2 PLEXOS

The ReEDS reduced-form dispatch, aided by Augur's parameterization, aims to capture enough operational detail for realistic bulk power system investment and retirement decisions, but it does not have the temporal resolution that is desired for Cambium databases. To obtain more-detailed simulations of the electric systems projected by ReEDS, NREL developed utilities to represent a ReEDS capacity expansion solution in the second of the two models that Cambium draws from: PLEXOS (Energy Exemplar 2019).⁴

PLEXOS is a commercial production cost model that can simulate least-cost hourly dispatch of a set of generators with a network of nodes and transmission lines. It incorporates representations of unit-commitment decisions, detailed operating constraints (e.g., maximum ramp rates and minimum generation levels), and operating reserves; and it can be run with nested receding horizon planning periods (e.g., day-ahead and real-time) to simulate realistic electric system operations.⁵

For representing a ReEDS solution as a PLEXOS model, the spatial resolution from ReEDS is retained: the 134 BAs in ReEDS are represented as transmission nodes in PLEXOS, and the connections between them are modeled using the line capacities and loss rates in the ReEDS

³ (Sigrin et al. 2016) is the most recent documentation of the dGen model. More information about dGen, such as the most up-to-date documentation, open-source access, and other publications can be found at <https://www.nrel.gov/analysis/dgen>.

⁴ The ability to represent a ReEDS-modeled solution in PLEXOS has also been used to study ReEDS-built systems for other NREL analyses (Frew et al. 2019; Cole et al. 2020; Cole et al. 2019).

⁵ Separately from ReEDS, PLEXOS has been used extensively by NREL for analyses of the electric sector, such as the Western Wind and Solar Integration Study (Lew et al. 2013) and the Eastern Renewable Grid Integration Study (Bloom et al. 2016).

aggregated transmission representation. Generation capacity at each node is, however, converted from aggregate ReEDS capacity to individual generators using a characteristic unit size for each technology. For consistency, ReEDS cost and performance parameters are used when possible and reasonable, but values derived from previous NREL studies (e.g., Lew et al. [2013]) are used when parameters are unavailable from ReEDS or are available but unreasonable because of structural differences between the models.⁶

Once the ReEDS solution is converted to a PLEXOS database, the hourly dispatch of the grid can be simulated for a full year. For Cambium databases, we run PLEXOS as a mixed integer program, with day-ahead unit commitment and dispatch (without any real-time adjustments for subhourly dispatch or forecast error). For each modeled year, generators have constant heat rates, short-run marginal costs (SRMC), and maximum generator output. Supply and demand are balanced at the busbar level, and distribution losses are captured in data pre- and post-processing, as described in Section 4.5. Inter-BA transmission is represented as pipe flow with constant loss rates, with no intra-BA transmission losses. Generator outages are represented as discrete events, where planned outages are scheduled dynamically by PLEXOS and forced outages are random events based on the outage rates used by ReEDS. Three operating reserves are represented—regulation, flexibility, and spinning reserves—as detailed in Section 5.2.

We draw from these simulated results—from both ReEDS and PLEXOS—to calculate the metrics reported within Cambium databases, with varying degrees of post-processing, as described in the remainder of this document.

⁶ Generator outages, for example, are represented in ReEDS by derating a technology’s available capacity, due to the relatively coarse temporal resolution. In PLEXOS, in contrast, both planned and unplanned outages are represented as discrete events.

2 User Guidance, Caveats, and Limitations of Cambium Databases

When projecting the expansion and operation of the U.S. electric system in coming decades, it is necessary to make various simplifications. Here, we list some important limitations and caveats that result from these simplifications:

- **The Metrics in Cambium Databases Have Not Been Validated for Every Possible Application:** The models used by Cambium are necessarily simplifications of reality. The metrics produced by those models, therefore, deviate from reality to some degree. Given the breadth of possible uses, it is not feasible to validate and summarize each metric’s accuracy for every conceivable application. We encourage users to critically evaluate the suitability of Cambium metrics for new analyses by reviewing existing literature and comparing Cambium data against empirical data where possible.
- **Cambium’s Metrics are Derived From System-Wide, Cost-Minimizing Optimization Models:** The models that Cambium draws from take system-wide, cost-minimizing perspectives that do not necessarily reflect the decision-making of individual actors, whose actions may not align with system-wide cost-minimization because of differing incentives or information deficits.
- **The Spatial and Temporal Resolution of the Underlying Models is Coarse:** Though the models that Cambium draws from have high spatial and temporal resolution for models of their scope, they do require simplifications and aggregations along those dimensions.⁷ Perhaps most critically, the United States is represented as 134 “copperplate” balancing areas (BA). This lack of transmission losses and constraints within BAs tends to produce lower marginal costs than what is observed in practice.
- **Cambium Reports Marginal Costs, Which Can Differ from Market Prices:** Cambium databases contain estimates of various marginal costs (i.e., how much the costs of building and operating the power sector increases with an increase in demand). Importantly, market prices in practice can deviate from marginal costs due to market design, contract structures, cost recovery for non-variable costs, and bidding strategies. We strongly encourage users to read the descriptions of each marginal cost metric reported in Cambium (Section 3.6) for an understanding of the limitations of each metric.
- **Cambium’s Marginal Costs are Not Estimates of Retail Rates:** The marginal costs in Cambium should not be directly used as estimates of retail electricity prices because (1) retail rates typically include cost recovery for administrative, distribution infrastructure, and other expenses that are not represented in Cambium databases and (2) retail rates are often set through a ratemaking process that, while sometimes reflecting temporal patterns in the marginal costs of electricity, are generally not priced directly at marginal costs but rather seek to balance cost-recovery, equity, and cost-causation.
- **The Full Range of Uncertainty is Not Captured:** The models that Cambium draws from compute deterministic least-cost solutions for a particular set of assumptions—each

⁷ See (Cole et al. 2017) for a multi-model analysis that, in part, explores the impact of spatial and temporal resolution in long-term planning models.

scenario, therefore, does not fully reflect the uncertainties in the underlying assumptions and data. Cambium, through the Standard Scenarios, tries to address this by providing a suite of possible futures, although the full range of possible outcomes can never be fully captured; for example, no scenario includes a severe economic depression as one of many possible futures that are not modeled.

- **Cambium is Not Designed to Assess Grid Reliability or Resource Adequacy:** Although the models that Cambium draws from can recognize dropped load when insufficient capacity is available to meet demand, these runs should not be considered as assessing grid reliability or resource adequacy, because, among other reasons, (1) only a specific set of conditions (weather, load, and renewable resource quality) is simulated, (2) important temperature effects on generator efficiencies and transmission losses are not represented, (3) transmission line outages are not represented, and (4) intra-BA transmission line constraints are not captured.
- **ReEDS' Capacity Expansion Decisions Have Limited Foresight:** Except when it runs with intertemporal optimization, ReEDS has limited long-term foresight and therefore model decision-making in a particular modeled year does not account for anticipated changes to markets or policies in future years. For example, when running without intertemporal optimization, ReEDS would not anticipate or react to the upcoming expiration of an incentive program. Unless otherwise specified, scenarios in Cambium databases are not run with intertemporal optimization.
- **PLEXOS' Production Cost Modeling Does Not Have Forecast Error:** Although PLEXOS is capable of representing load and variable generator forecast error, we do not deploy this feature in the runs that Cambium draws from.
- **Cambium Databases Do Not Contain Elasticity Data:** There are no estimates of the elasticity of the metrics reported in Cambium databases (i.e., how much a metric's value would change if load or generation changed). Though many metrics may change little across a wide range of intervention sizes, some metrics may be highly elastic, particularly at certain points in time. Hours with a marginal energy cost of zero, for example, could become nonzero with small increases in demand. In general, the larger an intervention being analyzed, the more likely it is that the Cambium values are not accurate for that specific intervention.
- **Flexible Loads are Not Currently Represented in Cambium:** Although the models that Cambium draws from have the ability to represent flexible load, this capability is not currently used for Cambium. Grid-responsive buildings and intelligent charging of electrical vehicles are two of many potential examples of electric loads that may play a meaningful role in operation of the grid in the future—for example, by absorbing otherwise-curtailed energy or shifting load away from high-demand periods.
- **The Project Pipeline and Retirements Data is Likely Incomplete:** Although ReEDS incorporates data of planned or under-construction projects, these data are unlikely to include *all* projects in progress, and it is possible that some planned projects included in ReEDS will not be finished. Similarly, scheduled near-term retirements are represented, but may be incomplete or reversed.

- **Only the U.S. Electric Sector is Represented in Cambium:** The models that Cambium draws from only represent the electric sector of the coterminous United States, not adjacent sectors nor the global energy economy. For example, competing uses of natural gas or financial capital across sectors and countries are not dynamically represented.
- **A Single Year of Weather Data is Used for Most Cambium Metrics:** The PLEXOS runs that Cambium draws from use weather data from a single year, 2012. Therefore, most metrics in Cambium databases (e.g., the marginal cost and variable generator patterns) are not “expected values,” although they do capture realistic weather-induced hourly variations.

To help users better understand some of these limitations, we have provided modeled data for 2018 in the Mid-case of the 2020 Standard Scenarios data release so that users can compare observed historical market and operational data with the data provided in Cambium databases.

3 Cambium Metric Definitions

In this section, we briefly define all the metrics in Cambium databases. The outputs from ReEDS and PLEXOS are the starting point for Cambium’s processing; some of the metrics are direct reports from those models, but others involve extensive post-processing. We describe the Cambium-specific post-processing methods in Section 5.

3.1 Busbar and End-Use Values

Many metrics in Cambium databases are reported at both the busbar and end-use level. Busbar refers to the point where bulk generating stations connect to the grid, whereas end use refers to the point of consumption. Analyses of bulk generators would typically use busbar values, whereas analyses of electricity consumers would typically use end-use values.

Busbar and end-use values differ by the loss rates between the two points. The models that Cambium draws from represent two types of losses: distribution losses and transmission losses between balancing areas (BA). The models treat each of the 134 BAs as a “copper plate” and therefore do not capture intra-BA transmission losses.

Therefore, for a given BA within a Cambium database, busbar and end-use values are related by distribution loss rates. For a generic “average” metric X (e.g., the average CO₂ rate of in-region generation), the two values are related by the average distribution loss rate α . For a generic “marginal” metric Y (e.g., the marginal energy cost), the two values are related by the marginal distribution loss rate μ .

$$\text{Averages: } X_{end-use} = \frac{X_{busbar}}{(1 - \alpha)}$$

$$\text{Marginals: } Y_{end-use} = \frac{Y_{busbar}}{(1 - \mu)}$$

Hourly α and μ values are given in Cambium databases as *distloss_rate_avg* and *distloss_rate_marg* respectively. See Section 5.3 for our approach and assumptions for calculating these metrics.

3.2 Time and Geographic Identifiers

Metric Family: timestamp

Metric Names: *timestamp* and *timestamp_local*

The *timestamp* metric is the time in Eastern Standard Time. The *timestamp_local* variable is the time in the local Standard Time. If no *timestamp_local* variable is in a file, the data are in Eastern Standard Time (and therefore indicated with the *timestamp* variable).

Both timestamp variables are hour-beginning, meaning a 1:00 timestamp indicates data for 1:00–2:00. Neither timestamp variable includes the effects of Daylight Savings Time. Every year in a

Cambium data set has 8,760 hours; leap days are skipped in the timestamps during leap years, although the 7-day weekday/weekend pattern is preserved.⁸

Every time series—and the underlying data—in a Cambium data set starts on a Sunday, regardless of the actual day of the week for January 1 of that year.⁹ This keeps the weekend/weekday patterns and hour positions consistent between years in the data, which facilitates analysis that spans across multiple years.

Metric Family: time zone

Metric Name: *tz*

The *tz* variable in the metadata indicates the time zone used for the *timestamp_local* variable. For regions that contain multiple time zones, the data are reported using the time zone where the majority of the load is located.

Metric Family: ReEDS model balancing area (BA)

Metric Names: *r*

The balancing area (*r*) is the finest geographic unit for which Cambium data are reported (Figure 1). There are 134 BAs, which are used as the nodes for balancing supply and demand in both the ReEDS and PLEXOS models that Cambium draws from.

⁸ If users wish to represent the additional 24 hours in leap years, we recommend they copy the data from the third day from each year (a Tuesday in Cambium) and add it to the end of the year's time series and then rename the date-times to incorporate February 29. If, instead, 24 hours of data are added between February 28 and March 1, the weekday/weekend pattern of the time series would be disrupted.

⁹ Cambium data starts on a Sunday because the underlying weather and load data are from 2012, which also start on a Sunday.

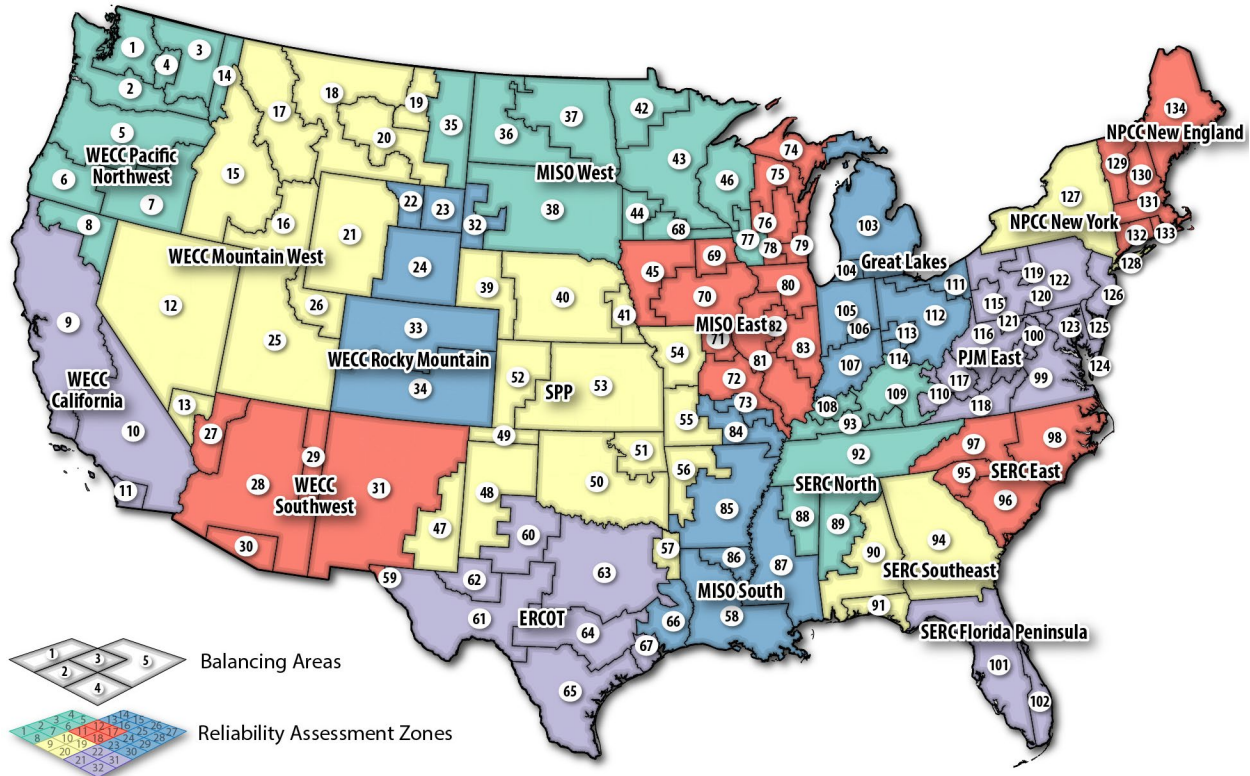


Figure 1. Balancing areas and reliability assessment zones used in Cambium, ReEDS, and PLEXOS

MISO: Midcontinent Independent System Operator; NPCC: Northeast Power Coordinating Council; SERC: SERC Reliability Corporation; SPP: Southwest Power Pool; WECC: Western Electricity Coordinating Council

Metric Family: reliability assessment zone

Metric Name: *raz*

The reliability assessment zones (RAZ) are 18 regions that ReEDS uses for various calculations, such as determining the capacity credit of variable generators. They are also used for enforcing operating reserve requirements in PLEXOS. They are shown in Figure 1.

3.3 Generation and Capacity Metrics

Metric Family: total generation

Metric name: *generation*

Units: MWh_{busbar}

The *generation* metric reports the total generation from all generators within a region. It includes generation from storage (e.g., batteries or pumped hydro storage). It does not include curtailed energy. If there are net imports or exports from a region, generation will not match load.

Behind-the-meter PV generation is included in the *generation* metric and is reported as the equivalent amount of busbar generation (i.e., it is increased to reflect the assumption that it does not incur distribution losses).

Metric Family: variable generation

Metric Name: *variable_generation*

Units: MWh_{busbar}

The *variable_generation* metric reports the total generation from all variable generators within a region, which includes PV, concentrating solar power (CSP) without storage, and wind. It does not include curtailed energy. Behind-the-meter PV generation is included, and as with *generation*, is reported as the equivalent amount of busbar generation.

Metric Family: generation by technology

Metric Name: *technology_MWh*

Units: MWh_{busbar}

These metrics report the total generation within a region from each of the 18 technologies listed in Table 1 (page 25). These generation values do not include curtailed energy. Generation from behind-the-meter PV, which is assumed to occur at the point of end use, is reported as an equivalent amount of busbar generation. The McIntosh compressed air energy storage plant in McIntosh, Alabama, has its generation reported as one-third natural-gas combined cycle generation, and two-thirds battery generation.

Metric Family: nameplate capacity by technology type

Metric Name: *technology_MW*

Units: MW

These metrics report the total nameplate capacity within a region from each of the technologies listed in Table 1 (page 25) (except for Canadian imports). Behind-the-meter PV is reported as the AC inverter capacity—it is not adjusted to a busbar equivalent capacity, unlike generation from the same technology. The capacities of wind and solar generation are reported at their original nameplate capacities when they were installed (i.e., their reported capacity is not reduced over time by degradation). Both planned and unplanned outages are represented in Cambium as discrete events, not as a derating of capacity. The McIntosh compressed air energy storage plant in McIntosh, Alabama, has its capacity reported as one-third natural-gas combined cycle capacity and two-thirds battery capacity.

3.4 Emission Metrics

Metric Family: average emission rates of in-region generation

Metric Name: *co2_rate_avg_gen*

Units: kg CO₂/MWh_{generation}

The *co2_rate_avg_gen* metric is the average emission rate of all generation within a region, in kilograms (kg) of CO₂ emissions per megawatt-hour (MWh) of busbar generation. No adjustment is made for imported or exported electricity. Start-up and shut-down emissions are not included.

Metric Family: average emission rates of generation induced by a region’s load

Metric Name: *co2_rate_avg_load_enduse*

Units: kg CO₂/MWh_{enduse}

The *co2_rate_avg_load_enduse* metric is the average emission rate of the generation that is induced by a region’s end-use load in kilograms of CO₂ emissions per megawatt-hour of end-use load. Unlike the average emission rate of in-region generation, this metric includes the effects of imported and exported power. For example, if the power for a region’s load is being supplied in equal parts from in-region generation with a CO₂ emission rate of 400 kg/MWh and out-of-region generation with a CO₂ emission rate of 1,000 kg/MWh, the average CO₂ rate for generation supplying the power for that load would be 700 kg/MWh. If distribution losses were 5%, this metric would report the average CO₂ rate as 737 kg/MWh of end-use load.

Cambium allocates generation, and therefore the emissions from generation, by assuming perfect mixing through nodes. For a description of the method employed in Cambium for power flow accounting, and its limitations, see Section 5.2.

The PLEXOS runs that Cambium draws from do not restrict which generators can provide power to which locations. In practice, some states have restrictions on the types of out-of-state power that can be imported, such as California’s limits on long-term contracts for out-of-state coal power. Additionally, utilities often contract with specific suppliers that may justify them claiming power mixtures that are different from the estimates produced by the perfect-mixing approach implemented in Cambium. Cambium does not currently capture these state or utility accounting effects; it just assumes perfect mixing through all nodes. This assumption, and therefore this metric, may not be appropriate for some analyses.

Similarly, this metric does not include the effects of any type of energy attribute accounting, such as unbundled renewable energy credit (REC) purchases. End users could claim that the emission rate induced by their load is different from this metric if they retired RECs to cover some or all of their demand.

Metric Family: short-run marginal emission rates for a region’s load

Metric Name: *co2_srmer_enduse*

Units: kg CO₂/MWh_{enduse}

The *co2_srmer_enduse* metric is the short-run marginal emission rate (SRMER) for end-use load, which is the rate of emissions that would be induced by a marginal increase in a region’s load at a specific point in time. The value is the emission rate of whichever generator would have served the marginal increase in load, modified by any relevant transmission, distribution, and efficiency losses.¹⁰ For an overview of the method that Cambium uses to interpret PLEXOS results and estimate which generator was on the margin at every point in time, see Section 5.4.

¹⁰ If the marginal generator is not the initial source of energy (e.g., when the marginal generator is a battery), the marginal emission rate is derived from the emission rate of the actual marginal source of energy and is further modified by the efficiency of the energy-storing generator. For a discussion of how Cambium determines what the marginal energy source is in these circumstances, see Section 5.5.

SRMER does not capture the structural changes to the electric sector that may be induced by increased demand, and it is therefore typically unsuitable for assessing the emissions impact of a change in demand that is expected to persist for several years or longer. Instead, the long-run marginal emission rate (LRMER) is a more suitable metric for persistent changes in demand.

For every balancing area and every hour, the Cambium method identifies a single marginal generator, although multiple regions can have the same marginal generator if they are connected by a partially utilized transmission line. In practice, which generator is on the margin can (and typically does) switch much more frequently than the 1-hour resolution of the Cambium data sets.

SRMER depends on proper identification of the marginal generator (and energy source). As we discuss in Section 5.4, identification of marginal generators and energy sources is an ongoing area of research. We encourage any researcher working with this metric to approach it with a critical eye.

Cambium’s modeled SRMER values are not appropriate for real-time operational decision-making. The primary intended use of this metric is to inform research questions that depend on anticipating the patterns of SRMER in potential futures (e.g., what the patterns of SRMER might look like in a future with high variable generator deployment).¹¹

Metric Family: long-run marginal emission rates for a region’s load

Metric Name: *co2_lrmer_enduse*

Units: kg CO₂/MWh_{enduse}

The long-run marginal emission rate (LRMER) is the emission rate of the mixture of generation that would serve a persistent change in end-use demand, taking into account any structural changes to the grid in response to the change in demand. The units are kilograms of CO₂ per megawatt-hour of end-use load. For a description of the methodology used to calculate LRMER, see Section 5.7.

LRMER is designed to inform decisions that would result in a persistent change in load. The “long-run” in LRMER does not refer to a specific length of time, but instead refers to the equilibrium least-cost solution when the option for structural changes to the grid are taken into account. Therefore, we recommend LRMER be used for analyzing changes in load that are expected to persist for several years or more, although the actual length of time it takes for an intervention to influence the structure of a grid depends on the particularities of the local planning processes.

As with all marginal values in Cambium databases, LRMER can be applied to either load increases or decreases. Load increases could be estimating the electric-sector emissions that would be induced by increased electric vehicle charging or replacing a natural gas furnace with a heat pump. Load decreases might be estimating the emissions avoided by installing a more efficient cooling technology.

¹¹ For real-time data meant to inform operational decisions, we refer to the work of groups such as WattTime (Callaway et al. 2017).

3.5 Cost Metrics

The metrics in this section are estimates of the marginal costs induced by an increase in demand (or avoided costs from a decrease in demand). In some instances, it may be appropriate to use these values as approximations of market prices for corresponding electric services, but it is important for users to understand the limitations of using marginal costs from least-cost optimization models as estimates of market prices. We strongly recommend users read Section 2, as well as the methods sections that discuss each cost metric in more detail than do the brief summaries given in this section.

In the Mid-case of the 2020 Standard Scenarios data release, we provide data for 2018 so users can compare the marginal costs in Cambium with observed historical price patterns and can thus better understand the limitations of the data.

All dollar values are in real terms for a constant dollar year. For the Standard Scenarios data sets, the dollar year is the year preceding the release (e.g., the 2020 Standard Scenarios dollar values are in 2019 dollars).

Metric Family: marginal energy costs

Metric Names: *energy_cost_busbar* and *energy_cost_enduse*

Units: \$/MWh_{busbar} and \$/MWh_{end-use}

The *energy_cost_busbar* and *energy_cost_enduse* metrics report the short-run marginal costs of providing the energy for a marginal increase in load, in dollars per megawatt-hour of either busbar or end-use load. These metrics are derived using the shadow price off of an energy constraint in the PLEXOS model. They include short-run costs that vary as a function of load (fuel and variable costs), but they do not reflect other operational costs that are fixed or vary as “steps,” such as start-up costs or fixed operation and maintenance costs.

These metrics are conceptually similar to a day-ahead locational marginal price, given the limitations discussed in Section 2. Specifically, the coarse geographic resolution, lack of temperature effects on generator heat rates and transmission losses, and fact that these are derived from the shadow prices out of a system-wide least-cost optimization model all contribute to these marginal costs tending to be less variable than observed prices in energy markets.

These marginal costs include the effects of generator short-run marginal costs, inter-BA transmission losses, and inter-BA transmission congestion. In the case of *energy_cost_enduse*, distribution loss effects are also included. Cost-recovery for start-up costs are not reflected in these values, as these are marginal costs and start-up costs are step changes. Debt service and fixed operations and maintenance costs are likewise not reflected in these marginal costs.

As a least-cost optimization model, PLEXOS can sometimes find solutions that result in exceptionally high marginal costs. For example, PLEXOS will sometimes drop a small amount of a reserve product and incur the associated penalty rather than incur the costs of starting up a generator that could have provided those reserves. This results in the marginal energy cost being set by the \$/MWh penalty for dropping the reserve product in that hour. Though this is a technically correct description of the least-cost solution, we feel these marginal cost spikes are not useful descriptions of the situation and are not generally helpful for the types of analyses for

which Cambium data are used. Therefore, for each BA and each time-step, we cap the marginal energy costs at the short-run marginal cost (plus an adder for partial start-up cost recovery) of the least-expensive natural-gas combustion turbine (NGCT) that (1) has available capacity in that BA. In the infrequent instances of a BA not having or being connected to an NGCT with remaining capacity, the SRMC plus start-up adder of the most expensive NGCT in the conterminous United States is used as a cap. As stated in Section 2, the models that Cambium draws from are not set up to assess resource adequacy or reliability, and the implementation of these price caps reflects that limitation.

Metric Family: marginal capacity costs

Metric Names: *capacity_cost_busbar* and *capacity_cost_enduse*

Units: \$/MWh_{busbar} and \$/MWh_{end-use}

The *capacity_cost_busbar* and *capacity_cost_enduse* metrics report the long-run cost of additional capital investment necessary to maintain a target planning reserve margin when demand is increased. An annual marginal capacity cost is derived from the shadow price off of the capacity constraint in the ReEDS model, which is set by the least-cost option for obtaining a marginal increase in firm capacity within each BA. The increase in firm capacity can be achieved by building new generation capacity, by holding on to existing generation capacity that would otherwise have been retired, or by building new inter-BA transmission capacity, whichever is the least-cost solution.

The annual shadow price is then increased by the planning reserve margin and allocated to the highest net-load hours within the year. The use of net-load is a heuristic for identifying the hours with the highest loss of load probability, and therefore the hours in which increased demand would induce a need for more firm capacity. For a detailed discussion of these methods and their limitations see Section 5.1.

A region can have a marginal capacity cost of zero for a year if the capacity constraint was not binding for that year (i.e., the available firm capacity exceeded what was required by the region's planning reserve margin and peak load).

The annual planning reserve margin and shadow price on the capacity constraint are also provided in Cambium databases as the metrics *prm* and *capacity_shadow_price* respectively. The quantity of firm capacity set by the planning reserve margin is reported as *planning_capacity_MW*.

Metric Family: marginal operating reserve costs

Metric Names: *op_res_cost_busbar* and *op_res_cost_enduse*

Units: \$/MWh_{busbar} and \$/MWh_{end-use}

The *op_res_cost_busbar* and *op_res_cost_enduse* metrics report the marginal costs of procuring the additional operating reserves that would be required by a marginal increase in the busbar and end-use load respectively.

Three operating reserves are represented in Cambium's PLEXOS runs: spinning reserves, flexibility reserves, and regulation reserves. Of these three products, the required quantity of both spinning reserves and regulation reserves are assumed to increase with busbar load, with spinning reserves being held at 3% of load and regulation reserves held at 1% of load. Flexibility

reserves are assumed to be a function of wind generation and online solar capacity, and not a function of load. These quantities are a simplification—in practice, different areas of the country have different requirements for reserve quantities. For details about the representation of operating reserves in these model runs, see Section 4.2.

Because spinning and regulation reserves are the only reserves influenced by an increase in load, *op_res_cost_busbar* and *op_res_cost_enduse* are the combined cost of procuring the required amount of those two reserves, as induced by a marginal increase of load. The marginal cost for each service is calculated by multiplying the shadow price for that service’s constraint by the rate at which reserves must be held for a given amount of load. If the marginal cost of procuring a megawatt of spinning reserves is \$10/MW, for example, the cost induced by a marginal increase in busbar load is \$0.3/MWh of busbar load.

The hourly shadow price and demand quantity for each service is also available in Cambium databases (the *spin_MW*, *flex_MW*, *reg_MW*, *spin_shadow_price*, *flex_shadow_price*, and *reg_shadow_price* metrics). The shadow price values can be used to calculate the individual spinning and regulation reserve marginal costs if a breakdown is desired.

Importantly, the marginal operating reserve costs in Cambium are often zero because of (1) the assumption that there are no direct operating costs for generators to provide these reserves and (2) the step changes in available capacity that result from the commitment of a generator. As a result, there are only nonzero shadow prices for these reserves when there is an opportunity cost associated with the capacity that would provide a marginal increase in the reserve product. Because most hours tend to have at least some online capacity not already dedicated to energy generation and reserve provision, these shadow prices are often zero. In practice, market prices for operating reserves are typical non-zero, and therefore the actual cost to the electric system to maintain operating reserves would tend to be higher than the values given here.

Metric Family: marginal costs of renewable and clean energy portfolio standards

Metric Names: *policy_cost_busbar* and *policy_cost_enduse*

Units: \$/MWh_{busbar} and \$/MWh_{end-use}

The *policy_cost_busbar* and *policy_cost_enduse* metrics report the marginal cost of staying in compliance with a state’s portfolio standard policies—both renewable portfolio standards (RPS) and clean energy standards (CES)—when end use demand is increased. For Cambium databases, unless otherwise noted in a scenario, enacted state-level RPS and CES are included.

For example, if a noncompliant technology (e.g., a natural gas generator for an RPS) is on the margin during a particular hour, additional consumption during that hour would require an increase in compliant generation at another point in the year for the standard to still be met. This cost reflects the cost of obtaining the required generation or credits through either operations or purchase.

In contrast, if a compliant technology is on the margin (e.g., a curtailing solar photovoltaic generator under most portfolio standards), there would be a value (i.e., negative cost) to additional consumption during that hour, as additional consumption would create credits from the otherwise-curtailed-generator, decreasing the need to acquire them through other means.

These costs are zero if either there are no portfolio standard policies or the policies that exist are not binding at that point in time.

For a discussion of how these metrics are calculated for Cambium databases, see Section 5.6. For a discussion of which policies are represented, and how they are represented in ReEDS, see Brown et al. (2020).

Cambium databases also include the annual shadow prices on the policy constraints (see *rps_all_shadow_price*, *rps_solar_shadow_price*, *rps_wind_shadow_price*, and *ces_shadow_price*), as well as the fraction of end-use load covered by each policy (see *rps_all_f*, *rps_solar_f*, *rps_wind_f*, and *ces_f*).

Metric Family: total marginal cost

Metric Names: *total_cost_busbar* and *total_cost_enduse*

Units: \$/MWh_{busbar} and \$/MWh_{end-use}

The *total_cost_busbar* and *total_cost_enduse* metrics are the sum of energy, capacity, operating reserve, and portfolio costs. These are only the costs that are currently included in Cambium databases, and they do not include costs for distribution capacity, administrative and general expenses, and other electric sector expenses. Therefore, this metric does not capture all the costs of building and operating the electric system. If the intervention being analyzed would influence costs beyond the ones currently included in the Cambium database, those additional costs may need to be estimated through other methods for a complete analysis.

Additionally, we emphasize that these costs are estimates of the costs incurred by the bulk power system by marginal consumption, and they are not estimates of retail electricity prices. Retail prices typically include cost recovery for other expenses and are often set by ratemaking methods designed to collect target revenue amounts from various customer classes, instead of adhering strictly to marginal cost pricing.

3.6 Interregional Transmission Metrics

These transmission metrics include only transmission between BAs, not within BAs. They also do not include Canadian imports and exports, which are represented as generation and end-use loads in the respective border regions.

Metric Family: total imports and total exports

Metric Names: *imports* and *exports*

Units: MWh_{busbar}

The *imports* and *exports* metrics capture the total imports and exports into and out of a region through interregional transmission lines, in megawatt-hours of energy at the busbar level.

This value is the energy sent along the transmission lines, and it is not netted by transmission losses. Transmission losses (reported in Cambium databases as additional load in the *trans_losses* metric) are allocated equally between the sending and receiving regions. For example, if 100 MWh of energy is transmitted between two regions while incurring 5 MWh of losses, the load in both the sending and receiving regions would increase by 2.5 MWh. In effect,

the receiving region would receive a net of 97.5 MWh of energy while the burden on the sending region would be higher by 102.5 MWh. This would be reported as 100 MWh of *imports* in the receiving region, 100 MWh of *exports* in the sending region, and 2.5 MWh of *trans_losses* in both the sending and receiving regions.

3.7 Load Metrics

Metric Family: total load at the busbar

Metric Name: *busbar_load*

Units: MWh_{busbar}

The *busbar_load* metric reports the total electric load in a region, in megawatt-hours of busbar load. It includes the load from end uses (including the busbar equivalent of end-use load that is served by behind-the-meter PV), load incurred through transmission losses, and load from storage generators charging.

Metric Family: end-use load

Metric Names: *enduse_load* and *busbar_load_for_enduse*

Units: MWh_{enduse} and MWh_{busbar}

The *enduse_load* metric reports the amount of electricity consumed at the point of end use within a region, including end-use load that is served by behind-the-meter PV. The metric *busbar_load_for_enduse* reports the quantity of load consumed at the busbar level to meet that end-use load. Therefore, *busbar_load_for_enduse* is larger because it is prior to incurring distribution losses whereas *enduse_load* is smaller because it is after incurring distribution losses.

Neither of these metrics includes transmission losses or storage load, which are both loads induced at the busbar.

In border regions, Canadian exports are included in the end-use load metric.

Metric Family: load from transmission losses

Metric Name: *trans_losses*

Units: MWh_{busbar}

The *trans_losses* metric reports the amount of energy that is lost due to inter-BA transmission losses. The losses are represented as an additional load at the busbar level, split equally between the sending and receiving BA.

Metric Family: load from storage generators that are charging

Metric Names: *battery_charging* and *phs_charging*

Units: MWh_{busbar}

The metrics *battery_charging* and *phs_charging* report the busbar load caused by the charging of electric battery storage and pumped hydro storage, respectively. The *battery_charging* metric includes charging from the McIntosh compressed air energy storage plant in McIntosh, Alabama.

Metric Family: net load
Metric Name: *net_load_busbar*
Units: MWh_{busbar}

The metric *net_load_busbar* reports the *busbar_load_for_enduse* minus *variable_generation*.

3.8 Operational Metrics

Metric Family: distribution loss rates
Metric Names: *distloss_rate_avg* and *distloss_rate_marg*
Units: MWh_{losses}/MWh_{busbar_load}

The metric *distloss_rate_avg* is the average distribution loss rate (i.e., the rate of losses incurred in the distribution of electricity to end uses in a region). The metric *distloss_rate_marg* is the marginal distribution loss rate (i.e., the rate of losses incurred in the distribution system by a marginal increase in the end-use load in the region). Marginal loss rates, and therefore average loss rates, increase as the end-use load in a region increases.

The average loss rate (α) is defined as losses (L) per busbar load consumed for end use (D):

$$\alpha = \frac{L}{D}$$

For example, if 100 MWh of energy were consumed at the busbar for end uses, and 5 MWh were lost in distribution, the total consumption at the end use would be 95 MWh and the average loss rate would be 5%. Similarly, the marginal loss rate (μ) is defined as the increase in losses per marginal increase in busbar load consumed for end use.

See Section 5.3 for our approach and assumptions for calculating average and marginal distribution loss rates.

Metric Family: planning reserve margin
Metric Name: *prm*
Units: MW_{firm}/MW_{peak}

The *prm* metric reports the planning reserve margin (PRM) used within ReEDS. Utilities, regulators, and system operators use the PRM as a heuristic for procuring sufficient firm capacity to achieve a desired level of resource adequacy, where resource adequacy is defined as “the ability of supply- and demand-side resources to meet the aggregate electrical demand” (NERC 2020).

The PRM is defined as the fraction of firm capacity above peak demand:

$$PRM = \frac{\text{Firm Capacity} - \text{Peak Demand}}{\text{Peak Demand}}$$

For example, in a region with a peak demand of 100 MW and a PRM of 0.15, the planned capacity would be 115 MW.

The PRMs applied in ReEDS are based on reserve margin requirements for North American Electric Reliability Corporation reliability subregions (NERC 2010).¹²

Metric Family: planning capacity

Metric Name: *planning_capacity_MW*

Units: MW_{firm}

The *planning_capacity_MW* metric reports how much firm capacity is called for in each region to meet the planning reserve margin, where firm capacity is defined as capacity that can reliably contribute to meeting the region's peak demand. For documentation of how ReEDS assesses the ability of different generators to contribute firm capacity, see Brown et al. (2020) and the forthcoming documentation for the 2020 version of ReEDS.

The sum of the BA-level planning capacities will exceed the maximum amount of firm capacity available in the conterminous United States, because peak demand periods are noncoincident across the country, and therefore capacity trading can reduce the total capacity needed to below the sum of the BA's requirements.

Metric Family: shadow price on the capacity constraint

Metric Name: *capacity_shadow_price*

Units: \$/MW_{firm}

The *capacity_shadow_price*, which is an annual value from the ReEDS model, is the marginal cost of procuring another megawatt of firm generation capacity. It is used in the calculation of the marginal capacity cost.

The shadow price off of this constraint is the \$/MW-year marginal cost for obtaining additional firm capacity. The model will find the least-cost option through three possible decision variables within the model:

- **New Generation Capacity:** Referred to as net cost of new entry (net CONE), the shadow price of the capacity constraint may be set by the annualized revenue needed to recover the costs of the generator that can provide firm capacity at the lowest cost, minus any revenue that the generator could obtain by providing other services (e.g., energy or operating reserves). This is often a natural gas combustion turbine plant, although in certain regions, it can also be variable resources like wind and solar, if their generation is well-aligned with peak demand.¹³

¹² The PRM is constant year-over-year for all regions except ERCOT. Because ERCOT was below its NERC-recommended level in 2018, the ERCOT PRM is set to the observed level of 10.9%, and the 2019 value is set at the 2019 projected level of 8.5%. The 2020 value is set at the average of 8.5% and the NERC-reference level of 13.75%, and years from 2021 onward are set at 13.75%.

¹³ ReEDS assesses the ability of variable generators (wind and solar) to provide firm capacity through a net load duration curve approach. Doing so tends to result in variable generators being able to provide firm capacity in the nearer term, which eventually goes to zero as net load peaks shift away from times of peak variable generation. See (Brown et al. 2020) and the forthcoming documentation of the 2020 version of ReEDS for a more detailed discussion of this approach.

- **New Transmission Capacity:** If a neighboring BA has excess generation capacity, the shadow price of the capacity constraint may be set by the annualized cost of building additional transmission capacity, minus the revenue the line would obtain from transmitting energy or operating reserve products.
- **Delayed Retirement:** ReEDS will choose to retire generation capacity if the capacity is not providing sufficient value to the system to cover its fixed costs (amplified by a multiplier to represent the “stickiness” of retirement). When this is happening, the shadow price of the capacity constraint can be set by the revenue that would have been required to keep that capacity online, minus the revenue it would have received for any other services.

Because of the prevalence of retiring generators, and the ability of wind and solar to contribute firm capacity, Cambium results for the 2020s often show capacity shadow prices that are substantially lower than what they would be if the shadow price were only being set by the net CONE of a natural gas combustion turbine. For some analyses, it may be appropriate for an analyst to substitute Cambium’s shadow price with a different estimate of the marginal cost of capacity (e.g., annualized cost of a new combustion turbine) if there is reason to believe such an estimate describes the behavior of the region being analyzed better than the solution found by Cambium.

Metric Family: technology of the marginal generator and marginal energy source

Metric Names: *marg_gen_tech* and *marg_es_tech*

The *marg_gen_tech* and *marg_es_tech* are the technologies of the short-run marginal generator and the short-run marginal energy source for a given location and time. These metrics only refer to the short-run marginal generator. In the long run, a marginal increase in demand is typically served by a mixture of generators.

The marginal generator is the generator that would provide the power to cover an increase in load, precisely when the load is increased. However, we differentiate between the marginal generator and the marginal energy source because some generators are energy-constrained and would therefore be unable to create new electrical energy if they were the marginal generator. We call these generators energy-constrained generators, and they include both generators that never have the ability to create new energy (e.g., batteries) and generators that have a limited budget of energy that they would dispatch entirely under expected conditions (e.g., dispatchable hydropower).

If an energy-constrained generator provides power, a different generator—one that can create new electric energy—must increase its generation at a different time. The generator that would ultimately increase its generation in response to the energy-constrained generator providing power as a marginal generator is the generator we consider to be the marginal energy source.

Each balancing area can only have a single marginal generator and marginal energy source during each time-step, although multiple balancing areas can share the same generators through transmission.

For a discussion of how the marginal generators and marginal energy sources are identified for Cambium databases, see Sections 5.4 and 5.5.

Metric Family: curtailment

Metric Name: *curt_solar* and *curt_wind*

Units: MWh

The *curt_solar* and *curt_wind* metrics report the amount of curtailed energy from nondispatchable solar and wind generators, respectively. Curtailed energy is energy that could have been produced but was not for economic or operational reasons. For example, energy could be curtailed when it is not needed (e.g., the amount of possible variable generation exceeds the amount of load at a point in time), or for operational cost-reduction (e.g., a small amount of energy may be curtailed from a variable generator to avoid an expensive shutdown and start-up of a thermal generator).

Due to degeneracy in the least-cost solution, the relative curtailment from wind and solar when both are generating is not meaningful. In those circumstances, it is best to look at the total amount of curtailment, understanding that it could have come from either wind or solar generators.

We encourage caution about analysis that draws conclusions from abundant curtailed energy. Our models do not currently have flexible load that might adapt to absorb curtailed energy. In futures with large amounts of curtailed energy, there may be meaningful quantities of flexible load that could respond to the low prices generally induced by curtailment. This does not prevent useful analyses that use estimates of curtailed energy, but it should be kept in mind when interpreting results.

Metric Family: operating reserve requirements

Metric Names: *spin_MW*, *flex_MW*, and *reg_MW*

Units: MW

The *spin_MW*, *flex_MW*, and *reg_MW* metrics are the requirements in megawatts for three types of operating reserves: spinning reserves, flexibility reserves, and regulation reserves.

Generators that are turned on but not dedicating their full capacity to energy generation can provide these reserves. Generators can provide both spinning and flexibility reserves from the same capacity if they meet the relevant requirements. Nuclear, PV, CSP without storage, and wind are not eligible to provide these reserves. Electric batteries and pumped hydro storage can provide these reserves if the generators have enough stored energy that they could discharge at full capacity for at least one hour.

For a discussion of these reserves, which generators can provide them, and the limitations of our representation, see Section 4.2.

Metric Family: shadow price on operating reserve constraints

Metric Names: *spin_shadow_price*, *flex_shadow_price*, and *reg_shadow_price*

Units: \$/MW

The *spin_shadow_price*, *flex_shadow_price*, and *reg_shadow_price* metrics are the shadow prices on each of the operating reserve constraints from PLEXOS. They represent the marginal cost of procuring another megawatt of each reserve.

Importantly, these shadow prices are often zero in Cambium because of (1) the assumption that there are no direct operating costs for generators to provide these reserves and (2) the step changes in available capacity that result from the commitment of a generator. As a result, there are only nonzero shadow prices for these reserves when there is an opportunity cost associated with the capacity that would provide a marginal increase in the reserve product. Because most hours tend to have at least some online capacity not already dedicated to energy generation and reserve provision, these shadow prices are often zero.

In practice, market prices for operating reserves are typical non-zero. Though we provide these values as indications of how our models view the marginal costs of holding these reserves, we recommend against using these costs as estimates for the market prices for similar services.

For a discussion of these reserves, which generators can provide them, and the limitations of our representation, see Section 4.2.

3.9 Portfolio Standard Metrics

Metric Family: shadow price on portfolio standard constraints

Metric Names: *rps_all_shadow_price*, *rps_solar_shadow_price*, *rps_wind_shadow_price*, and *ces_shadow_price*

Units: \$/MWh_{credit}

These four metrics are the shadow prices on four portfolio standard constraints from the ReEDS model. Unless otherwise specified in a scenario's description, the ReEDS runs that Cambium draws from represents both renewable portfolio standards (including separate representations of wind and solar carveouts, where applicable) and clean energy standards.

These metrics are annual values, and they represent the marginal cost of procuring another megawatt-hour of generation (or an unbundled credit from eligible generation, where allowed) that is eligible to satisfy the requirements of the policy.

Although these metrics are conceptually similar to the price of policy credits (e.g., RECs), important limitations to our representation mean these values are unlikely to be good estimates of the market price of these credits in practice. These limitations include the lack of inter-year credit banking, imperfect representations of policies, no representation of other consumers of policy credits, and the fact that these are long-run not short-run values (i.e., the shadow price reflects the option of building additional capacity to generate more credits).

For a discussion about how these values are used in calculating the marginal portfolio costs, see Section 5.6. For documentation of how these policies are represented in ReEDS, see Brown et al. (2020).

Metric Family: portfolio standard fractions

Metric Names: *rps_all_f*, *rps_solar_f*, *rps_wind_f*, and *ces_f*

Units: $MWh_{\text{credit}}/MWh_{\text{end-use}}$

These four metrics are the requirements of four portfolio standard constraints from the ReEDS model. Unless otherwise specified in a scenario's description, the ReEDS runs that Cambium draws from represent both renewable portfolio standards (including separate representations of wind and solar carveouts, where applicable) and clean energy standards.

All policies are represented as requiring a certain amount of eligible generation (or unbundled credit from eligible generation, where allowed) as a fraction of end-use load. Most renewable portfolio standards follow this same convention in practice, where the required credits are determined as a function of retail sales.

These fractions are the average requirement for the end-use load within the region covered by the policy. Defined in this way, these fractions are frequently lower than the nominal top-line number of the policy they represent, as many policies exclude certain types of load from being covered. For example, many states exempt utilities below a certain size from their policies. Such exemptions would result in a fraction that is lower than the nominal RPS goal.

For a discussion about how these values are used in calculating the marginal portfolio costs, see Section 5.6. For documentation of how these policies are represented in ReEDS, see Brown et al. (2020).

3.10 Marginal Generators and Marginal Energy Sources

In Cambium databases, we differentiate between marginal *generators* (the generator that would provide the power to cover an increase in load at the moment when the load is increased) and marginal *energy sources*. The distinction is needed because some generators, such as electric batteries, cannot create new electrical energy. We call these generators energy-constrained generators, and they include both generators that never have the ability to create new energy (e.g., batteries) and generators that have a limited budget of energy that they would dispatch entirely under expected conditions (e.g., dispatchable hydropower).

If an energy-constrained generator provides power, a different generator—one that can create new electric energy—must increase its generation at a different time. The generator that would ultimately increase its generation in response to the energy-constrained generator providing power as a marginal generator is the generator that we consider to be the marginal energy source.

Energy-constrained generators are conceptually equivalent to transmission lines, except they connect marginal energy sources to load over time instead of over space. In the same way that it can be insufficient to say a transmission line is providing the marginal unit of power at some moment in time, it also can be insufficient to say that an energy-constrained generator is providing the marginal unit of power.

Identifying the marginal energy source is relevant for some metrics in a Cambium database because the characteristics of the marginal energy source determines the impacts of increasing demand. The short-run marginal emission rate, for example, is driven by the emission rate of the marginal energy source, and modified by relevant transmission and efficiency losses, if the marginal generator is an energy-constrained generator.

We describe how we identify marginal generators in Section 5.4 and how we identify marginal energy sources in Section 5.5.

4 Cambium Methods

Cambium draws from the outputs of ReEDS and PLEXOS when creating its output database. Some outputs from the models are passed through without additional processing; for example, the installed capacity output from ReEDS is directly reported without any additional processing. Many other outputs require further processing, however. This section documents the Cambium-specific post-processing methods.

4.1 Technologies represented in Cambium

Data are reported for 18 technology groups in Cambium (Table 1). The actual number of discrete technologies in both the ReEDS and PLEXOS runs is greater, but the data are grouped to reduce the size of the database. The ReEDS and PLEXOS technologies that are within each of Cambium’s technology groups is given in Table 1. Within each ReEDS and PLEXOS technology, there can be generators with varying performance characteristics (e.g., heat rates), based on improvement over time, but we do not consider those as distinct technologies. For a detailed discussion of how each technology is represented, see Brown et al. (2020).

Unless otherwise specified in a project or scenario description, Cambium databases are based on NREL’s Standard Scenarios. See Cole et al. (2020) for a description of the technology cost and performance assumptions for each scenario, most of which draw from NREL’s Annual Technology Baseline (NREL 2020).

The McIntosh compressed air energy storage plant in McIntosh, Alabama, has both its generation and capacity reported as one-third natural gas combined cycle and two-thirds battery to align with its representation in the PLEXOS model runs.

Table 1. Cambium Technologies

Technology in Cambium	Technology Name in Cambium Database	Technologies in ReEDS and PLEXOS
Behind-the-meter PV	distpv	Behind-the-meter PV
Biomass	biomass	Biopower and landfill gas
Canadian imports	canada	Canadian imports
Coal	coal	Coal (scrubbed and unscrubbed, integrated gasification combined cycle, and biomass cofired technologies)
Coal with carbon capture and storage	coal-ccs	Coal with carbon capture and storage
Concentrating solar power	csp	Concentrating solar power (with and without thermal energy storage)
Electric batteries	battery	Electric batteries, two-thirds of McIntosh compressed-air energy storage plant
Geothermal	geothermal	Geothermal (hydrothermal, near-field enhanced geothermal, and deep enhanced geothermal systems)

Technology in Cambium	Technology Name in Cambium Database	Technologies in ReEDS and PLEXOS
Hydropower	hydro	Hydropower (existing and undiscovered, dispatchable and nondispatchable)
Natural gas combined cycle	gas-cc	Natural gas combined cycle, one-third of McIntosh compressed-air energy storage plant
Natural gas combined cycle with carbon capture and storage	gas-cc-ccs	Natural gas combined cycle with carbon capture and storage
Natural gas combustion turbine	gas-ct	Natural gas combustion turbine
Nuclear	nuclear	Nuclear
Offshore wind	wind-ofs	Offshore wind
Oil-gas-steam	o-g-s	Oil-gas-steam
Onshore wind	wind-ons	Onshore wind
Pumped hydro storage	phs	Pumped hydro storage
Utility-scale PV	upv	Utility-scale PV and distributed-utility-scale PV

4.2 Coloring Power Flows

When calculating the characteristics of the generation induced by load at a certain point—such as the average emission rate of the generators serving end-use consumption—the composition of the source generation must be determined, including the contribution of generators in different regions that may be sending power through transmission lines. Therefore, we need a method for allocating the generation from each generator to loads—or, from the other perspective, finding where the power for a given node’s end use originally came from. To do so, we take the network of nodes and transmission flows in each PLEXOS solution that Cambium draws from, and assume each node is a “perfect mixer” (i.e., that any electricity consumed or exported from a node is a perfect mixture of the electricity being supplied to the node).

Consider the network in Figure 2:

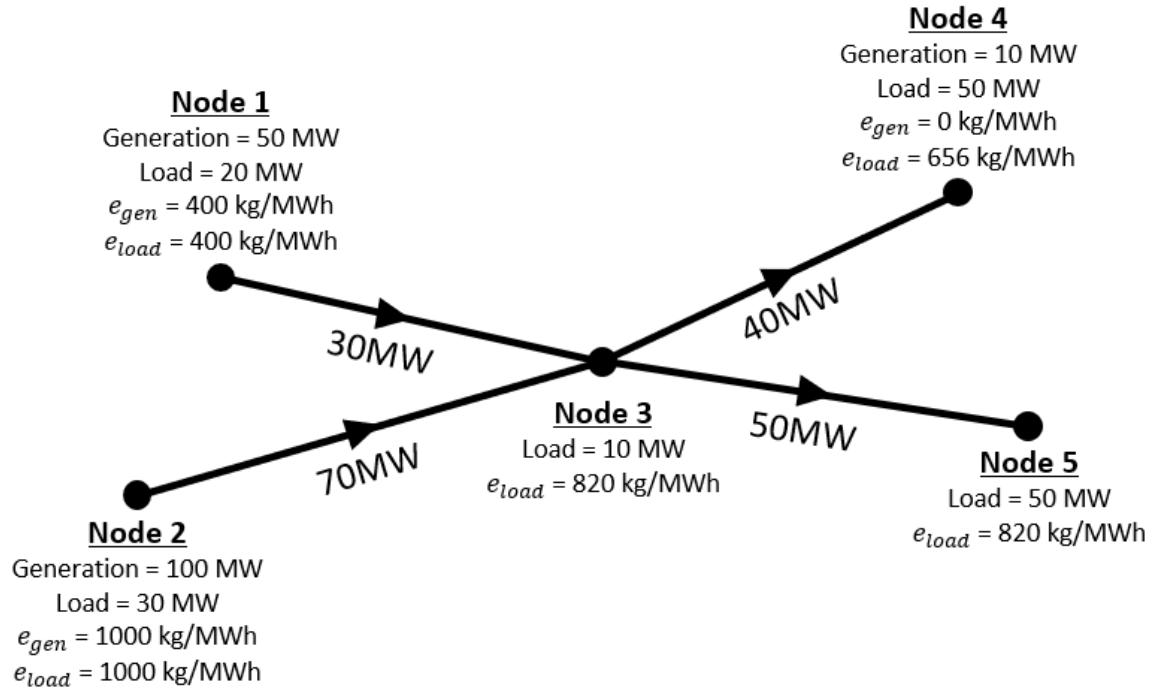


Figure 2. Simple network for illustrating power flow coloring

In this toy example shown in Figure 2, we have a system with five nodes ($N_1 - N_5$), connected by four transmission lines. Only nodes N_1 , N_2 , and N_3 have generation, with emission rates for in-region generation (e_{gen}) of 400, 1,000, and 0 kg/MWh respectively. For this example, we assume there are no transmission losses.

The question we are trying to answer is: What is the emission rate that you could ascribe to each of the five nodes' load (e_{load})?

For N_1 and N_2 , the only power flowing into the node is from their own generation (they are not importing any power), and therefore we consider the emission rates induced by their load to be the rates of their in-region generation, which are 400 kg/MWh and 1,000 kg/MWh respectively.

For N_3 , we see that 30% of the power it is receiving is coming from N_1 , and 70% is coming from N_2 . The weighted average of those two sources is 820 kg/MWh, and therefore we take that as the emission rate induced by the load at N_3 .

Given the assumption of perfect mixing through N_3 , we assume the power that both N_4 and N_5 are receiving from N_3 must be identical, and of the same character as the power that was consumed within N_3 itself. Therefore, the transmission from N_3 to N_4 and from N_3 to N_5 is assumed to have an emission rate of 820 kg/MWh.

N_4 is receiving 40 MW of power from N_3 at 820 kg/MWh and 10 MW of power from its own generation at 0 kg/MWh. The result is an emission rate of 656 kg/MWh ascribed to the load in N_4 .

N_5 is only importing power from N_3 , and therefore its load is ascribed the emission rate of 820 kg/MWh.

The example above is a trivial network: to calculate the contribution of each BA's generation to each BA's loads for the 134 BAs in Cambium (which are the nodes in our models), we use the downstream-looking algorithm from Bialek (1996), which we summarize next for our application here.

For each BA and each time-step, we take the generation in the node (g_i), total imports into the node ($f_{in,i}$), and total exports from the node ($f_{out,i}$), and derive the load (l_i):¹⁴

$$l_i = g_i + f_{in,i} - f_{out,i}$$

We then calculate the nodal through-flow P_i as the sum of the node's load and outflow.

$$P_i = l_i + f_{out,i}$$

We then calculate the values for the downstream distribution matrix (\mathbf{A}_d), which is a square matrix whose length and width is the number of nodes. The (i, l) element of \mathbf{A}_d is:

$$[\mathbf{A}_d]_{i,l} = \begin{cases} 1, & i = l \\ -|f_{i-l}|/P_i, & l \in \alpha_i^{(d)} \\ 0, & otherwise \end{cases}$$

Where $\alpha_i^{(d)}$ is the set of nodes that are directly supplied by node i and f_{i-l} is the total flow into node l directly from node i .

We then take the inverse of \mathbf{A}_d to obtain \mathbf{A}_d^{-1} . Using each (i, l) element of \mathbf{A}_d^{-1} , we calculate the amount of generation from source BA i that can be allocated to the load in destination BA l (g_{i-l}):

$$g_{i-l} = g_i * l_l * [\mathbf{A}_d^{-1}]_{i,l}/P_i$$

Lastly, we use this allocation of generation to determine the weighted average of the characteristics of the generation sources supplying each BA. For example, for generic attribute X_l :

$$X_l = \frac{\sum_{i=1}^n g_{i-l} * X_i}{l_l} \quad i = 1, 2, \dots, n$$

This is the general form of the method we described above for our trivial network in Figure 2.

¹⁴ We derive the load, instead of using the output from PLEXOS, to avoid violations of Kirchhoff's current law that could arise from rounding errors in the outputs.

In Cambium, we iterate through every time-step, calculating the allocation factors g_{i-l} between every pair of the 134 BAs (i.e., what fraction of the generation from one BA is allocated to the load in another BA). With these allocation factors, we can calculate the characteristics of the generation that is supplying every BA, at every time-step.

Caveats and Limitations

Allocating generation based on the assumption of perfect mixing through every node is not always appropriate. If load in one BA contracted to have generation of a certain type produced in a neighboring BA and shipped to them (e.g., with a power purchase agreement), it may be appropriate to assign all that generation to the BA that contracted for that specific type of energy. Similarly, some states, like California, have restrictions on importing power from certain fuel types (e.g., coal), and these restrictions would not necessarily be respected with this perfect-mixing implementation.

Relatedly, some state renewable portfolio standards allow unbundled RECs to be purchased from other regions and retired to satisfy the requirements for a utility's load. Depending on the question an analyst is trying to answer, it may be appropriate to take these unbundled REC transfers into account when reporting an emission rate associated with a particular node's load. Our approach does not currently have a representation of unbundled REC transfers; instead, it essentially assumes all credit for renewable energy is attached to the flow of the renewable power itself.

Lastly, because transmission losses are represented as load in our models, we are not accurately capturing their impact. Our current approach slightly dilutes the emission rate ascribed to end-use load, because it is not separated from the load induced by transmission losses. In future iterations of Cambium, we intend on improving this algorithm to properly reflect the effects of transmission losses, which our current implementation does not do.

4.3 Calculating Long-Run Marginal Emission Rates

The long-run marginal emission rate (LRMER) is the emission rate of the generation that a marginal increase in load would be served by, when you include the possibility that the new load could induce structural changes to the grid (e.g., building new generation or transmission capacity).¹⁵ This is in contrast to the short-run marginal emission rate, which is the emission rate that would serve a marginal increase in load, but with the capital assets of the grid being fixed.

Because utilities periodically invest in new assets in response to changes in load, we generally consider the LRMER to be more suitable than the SRMER for estimating the emissions that are induced by a marginal increase in load, if the load is going to be a persistent increase for more than several years.¹⁶ The differences between the two metrics can be significant, with the

¹⁵ The terms “short-run” and “long-run” do not refer to specific lengths in time, but instead are just referring to whether the equilibrium solution is evaluated with fixed capital assets (short-run), or by allowing capital assets to vary as part of the solution (long-run).

¹⁶ Published research on marginal emission rates generally focuses on short-run marginal emission rates (Siler-Evans et al. 2012; Thind et al. 2017). Though the literature has addressed long-run marginal emission rates (Hawkes 2014; 2010), no method for calculating hourly long-run marginal emission values has been published previously, to our knowledge.

SRMER often not fully capturing the effects of wind and solar generation induced by load increases, and therefore generally being higher than the LRMER.

For Cambium, we estimate the LRMER by comparing the mixture of generation that serves each balancing area’s (BA) load in a given year and the mixture of generation that serves that BA’s load 2 years out from that given year (e.g., we compare the 2030 grid mixture to the 2032 grid mixture).¹⁷ Because there are changes in demand between those two years, and our models determine the least-cost capacity expansion to serve that new demand, we interpret the observed changes in the grid mixture as same mixture that would have served a further increase in demand; in other words, we see what the model used to serve an increase in demand between years and then assume an additional amount of load would have been served by the same mixture.

We walk through a simple example of this below. We discuss the limitations of our current method at the end of this section.

A Simple Example for Calculating a Long-Run Marginal Emission Rate

To illustrate our method for calculating a LRMER, we consider a simple example: two nodes (N_1 and N_2), connected by transmission, that are served by different mixtures of generation between two modeled years (t_1 and t_2). Our method evaluates each hour independently—the example we give here is just for a single hour compared between the two years. The in-region generation for each node, during this single hour, is given in Table 2. N_2 is exporting energy to N_1 during both years.

Table 2. In-Region Generation and Transmission Flow Between Two Nodes

	CO ₂ Emission Rate (kg/MWh)	Node 1			Node 2		
		Year 1 (MWh)	Year 2 (MWh)	Change (MWh)	Year 1 (MWh)	Year 2 (MWh)	Change (MWh)
Load	—	110	130	20	90	130	40
Natural gas	400	50	60	10	50	90	40
Coal	1,000	50	20	-30	0	0	0
Wind	0	0	0	0	50	90	40
Imports	—	10	50	40	-10	-50	-40

We assume the installed capacity and mixture of generation for each node in each year represents the long-run equilibrium position. That is what happens in our models, although in practice, utility assets are likely never in long-run equilibrium.

¹⁷ We calculate the LRMER in 2-year steps because Cambium draws from biennial model runs. The 2-year step is not meant to restrict the value to applying to interventions with a two-year lifespan; if an intervention is expected to have a longer life, we recommend the LRMER values from the entire expected lifetime be used, either as discrete year-over-year values or by calculating a single hourly or annual value by levelizing the LRMER over the lifetime of the intervention.

Taking these nodes and mixtures of generation, we allocate the energy from each generator based on the power flow coloring method we describe in Section 4.1 (i.e., we allocate the generation from each generator by assuming all generation is perfectly mixed in each BA). For this toy example, the result is trivial, as a portion of N_2 's generation is allocated to N_1 in both years (e.g., in the first year, the generation mixture of N_2 is 50% natural gas and 50% wind, and N_1 imports 10 MWh from N_2 , therefore N_1 is considered to be importing 5 MWh of generation from natural gas and 5 MWh of generation from wind).

The resulting allocations of generation to each node in each year is given in Table 3. This is the generation we assume serves each node during this hour; for a discussion of the limitations of this allocation method, see Section 4.1.

Table 3. Generation Allocated to Two Nodes

	CO ₂ Emission Rate (kg/MWh)	Node 1			Node 2		
		Year 1 (MWh)	Year 2 (MWh)	Change (MWh)	Year 1 (MWh)	Year 2 (MWh)	Change (MWh)
Natural gas	400	55	85	30	45	65	20
Coal	1,000	50	20	-30	0	0	0
Wind	0	5	25	20	45	65	20

It may seem appropriate to simply divide the difference in total CO₂ emissions by the difference in total allocated generation (or, equivalently, the change in load) for each node to obtain a LRMER. However, if we do that for N_1 in this example, we see there was a decrease of CO₂ emissions of 18 metric tons and an increase in generation of 20 MWh, which would produce a value of -900 kg/MWh (i.e., such an approach would suggest a marginal *increase* in demand of 1 MWh would result in 900 kg *less* emissions). In this example, the negative value is driven by the decrease in coal generation. Technically, it could be possible that further increases in load would drive further decreases in coal generation (e.g., one could imagine a policy that imposes this result); but we consider it more reasonable to assume decreases in generation from a particular technology between years were not induced by any load increases but were rather induced by another effect (e.g., retirement). Therefore, we assume further increases in load would *not* induce further reductions in coal generation.

To reflect this assumption in our accounting, we calculate the LRMER by dividing the change in emissions by the change in allocated generation *only for the technology groups whose allocated generation has increased*. In both N_1 and N_2 in our example, that is both natural gas and wind but not coal.

In our toy example, the amount of natural gas generation allocated to N_1 increased by 30 MWh and the wind generation allocation increased by 20 MWh—in other words, 60% of new generation was served by natural gas and 40% was served by wind, for a weighted average emission rate of 240 kg/MWh. We assume any marginal increases in load would be served by the same mixture of generation, and therefore the LRMER is 240 kg/MWh. What makes this a *long-run* marginal instead of a *short-run* marginal is that the increase in any of these generation types could have occurred by the addition of new capacity. The increase in wind generation, in

particular, would likely have come from new wind capacity. If a short-run marginal had been calculated for this hour during the first year, it would likely have indicated that a marginal increase in load would have been served by either natural gas or coal generation, which would have resulted in a higher marginal emission rate. Instead, we see that a persistent increase in load would likely have been served by a mixture of natural gas and wind generation.

Repeating the same process for Node 2, we see that 50% of its new generation came from natural gas and 50% came from wind, for a LRMER of 200 kg/MWh.

When implementing this in Cambium, we aggregate generation by the technology groups listed in Table 1. Because technology groups can be composed of generators with slightly different heat rates (and therefore different emission rates), we calculate the change in emissions using the change in generation and the average CO₂ emission rate of each technology group in the second year, instead of directly comparing the CO₂ emissions from each technology groups between the two years. Doing so only has a small effect on the reported values, but it avoids some situations where negative LRMER values were being reported because of heat rate changes within a technology group, which we considered to be an unreasonable interpretation of those results, as further increases in demand would not decrease the total emissions from that technology group.

For a given pair of years (t_1 and t_2 in our example), we assign the calculated LRMER value to the first year. For Cambium, we repeat these calculations for every hour of the year and then step forward in time, such that the second year becomes the first year, and repeat the calculations with the new pair of years. We do this for all the available modeled years for each scenario, which gives us a set of hourly LRMER for each region for each modeled year.¹⁸

The simple example given here represents how most generators are accounted for, but certain energy-constrained generators need to be treated slightly differently. We explain the special treatment of those generators in the following section.

Accounting for Dispatchable Energy-Constrained Generators When Calculating LRMER

The simple example above only had dispatchable generators without energy constraints (natural gas and coal) and nondispatchable energy-constrained generators (wind). Some technologies, such as hydropower and CSP with storage, can be dispatchable but energy-constrained. These technologies present a complication for our method because of differences in generation patterns between years. For example, a given hydropower generator might dispatch relatively more in the afternoon in 2030 and relatively more in the evening in 2032, even if the total amount of generation was constant between the two years. Because our method works by identifying any increases in generation while treating each hour independently, this type of difference would look like an increase in hydropower generation in the evening, even though the energy budget did not vary between the years (and therefore the generation in the evening came at the expense of generation during the morning).

To mitigate this issue, we calculate the difference in total monthly generation for each of these generators between solve years. We then calculate the sum of the individual hourly increases,

¹⁸ In the 2020 Standard Scenarios data set release, the 2050 data are a duplicate of the 2048 data, because our underlying projections did not extend to 2052.

which will always be equal to or greater than the actual increase. We then calculate the fraction of the nominal hourly increases that is actually new generation; for example, if generation for a hydropower generator increased 10 MWh in January between solve years, but the nominal increase we observed by summing all the individual hourly increases was 100 MWh, we would say that only 10% of those nominal increases were actually new generation, while the other 90% was just redispatched energy from other hours. We ascribe this non-new generation an emission rate that is equal to the monthly average LRMER of all generators that do not fall into this dispatchable-but-energy-constrained category, for the BA in which hydropower generator is located.

Conceptually, we are acknowledging that a portion of the observed increase in hydropower generation was not actually new generation but must have been drawn from a different hour. This must have necessitated other generation covering the redispatched hydropower generation. By assigning the non-new hydropower generation the emission rate of the monthly average LRMER rate for that BA, we are approximating the actual new generation that was induced by the redispatching of the non-new portion of the hydropower generation.

Calculating LRMER When Load Decreases Between Modeled Years

For Cambium databases, we take advantage of the fact that we are modeling the hourly operation of the grid biennially, which naturally includes changes in load within each BA between years. However, in some instances, load decreases for particular hours between the two model solutions we are comparing. If the load decreases, two situations can happen, and we handle them differently.

First, it is often the case that even when load within a BA decreases between two model solves, there was still an increase of at least one generation type. Coal generation might decrease by 20 MWh, while wind generation increased by 10 MWh, for example. In situations like this, no special action is required; we can treat these regions the same as we treat the regions where there was positive load growth. The interpretation is the same: we observed an increase of a particular mixture of generation, so we assume any marginal increases in demand would have been served by that same mixture.

Second, it is possible that load decreases between years and that no generation of any type increased. Because the method we describe above would produce a null value if there were no generation increases, we fill in these null values with the month-hour average LRMER for that BA (e.g., if a noon hour in January in 2020 has this occurrence, we fill it in with the average LRMER for other noon hours in January in 2020 from the same BA).

Approximating “Expected Values” from a Single Weather-Year

For many applications, it would be preferable to derive a LRMER that combines results from a range of possible futures, including probability-weighted weather and load pattern variations. These would ideally be collapsed into single “expected values.” Because of the high computational costs of these model runs, we only run a single weather year for the Cambium data set. Doing so, among having other effects, creates undesirable hour-to-hour variations in our projected generation from variable generators; for example, a particular region might show a very high solar contribution in one hour and a very low solar contribution in the next because

of weather variations in our inputs, whereas the actual expected solar contribution in the two hours is probably somewhere between the two.

To partially mitigate this issue, we calculate a rolling average. Each hour's value is averaged with the values from the same hour during the preceding and the following 14 days, and this averaged value is what is reported in Cambium databases, approximating the effects of load and weather variations.

Caveats and Limitations

This method is best understood as an estimate, relying on several assumptions. It is a tractable way to interpret the data that produces a reasonable result, although it has several known limitations:

- **Power Flow Allocation:** As discussed in the power flow coloring section (Section 4.1), our power flow allocation method just assumes perfect mixing of power through each node and therefore does not capture relevant restrictions or modifications of the actual allocation one might give to the power flows. For example, restrictions on the amount of generation from coal plants that can be imported into California would not necessarily be respected. Cross-state contracts for electricity bundled with RECs, for renewable portfolio standard compliance, might also not be captured.
- **Year-Over-Year Values:** For Cambium databases, we calculate these LRMER values for every other year. Our assumption is that analysts would apply these year-over-year values for the duration of the intervention they are analyzing (e.g., an analysis of an asset with a 10-year lifetime starting in 2020 would incorporate the five values from 2020 through 2028). Given that this is an evolving area of research, we do not yet know how reasonable it is to estimate the emissions impact of an intervention with these year-over-year values.
- **Unbundled REC Tracking:** Similar to the power flow accounting issue, some state RPS policies can be at least partially satisfied through unbundled RECs purchased from out-of-state. For some purposes, it may be reasonable to include these RECs when estimating the emissions induced by load in a region (i.e., if unbundled RECs were taken into account, the emissions induced by load would be considered less than if the RECs were ignored).
- **Interdependency of Hours:** We treat each hour as independent, but the equilibrium build-out of the power sector is influenced by the combined shape of increasing demand across hours, and operational constraints of the electric sector means the dispatch in one hour can influence the dispatch in another hour. The generation that serves a particular hour's demand would likely be different if, for example, that increase was just for that hour, relative to whether it was part of a span of increased demand that included neighboring hours.
- **Transmission Losses:** Our method does not currently accurately capture the effects of transmission losses, because those losses are represented as a load, which dilutes the actual emission rate induced by an increase in load.

- **Policy Constraints:** Our method implicitly assumes the marginal increase in demand would face the same policy constraints as the observed increases. For policies whose effects are proportional to load, like most RPS and CES, this is likely reasonable; however, a policy whose effects are not proportional to load (e.g., a mandated amount of capacity of a specific generator type) would have its effects incorrectly incorporated into the LRMER calculation.
- **Outages:** The PLEXOS runs we draw from includes both scheduled and unscheduled outages. These can happen at different times between the two years we are comparing, creating distortions in the calculation of LRMER during the hours with the outages. If a nuclear generator had a forced outage during the first year but was online at that time during the second year, our method would suggest a large amount of nuclear generation came online at that point and would incorporate that into the LRMER calculation, although we would not actually expect further load to induce more nuclear generation.
- **Linearization Effects:** This method implicitly assumes any observed increases in generation would increase linearly with any additional marginal increases in demand. In certain circumstances, this might not be an appropriate assumption. For example, a coal plant might go from 80% to 100% utilization between solve years but not have further capacity additions in the long-run equilibrium solution. Our method would naively see the increase in coal generation and assume marginal increases in load would induce more coal generation even though the plant is already at full capacity. The same effect could be caused by transmission line utilization, variable renewable energy supply curves, siting restrictions, and several other possible reasons for a linearization of observed changes to be inaccurate.

Many of these issues would be difficult to avoid, given our objective of producing a set of LRMER values that can be applied to the analyses of different types of interventions. Others, like the unbundled REC tracking, could likely be addressed with further development of the post-processing methods. This is an ongoing area of research, and we expect improvements to these methods to continue for some time.

4.4 Identifying a Region's Marginal Generator

The short-run marginal generator, for a particular location and time, is the generator whose output would increase if there were a marginal increase in demand at that location and time.¹⁹ For several of the metrics reported in Cambium databases (e.g., the short-run marginal emission rate), it is necessary to identify which generator is the marginal generator.²⁰

¹⁹ Note that, in Cambium, we differentiate between the marginal *generator* and the marginal *energy source*. The marginal generator is the generator that would provide the power to cover an increase in load, at the moment when the load is increased. If the marginal generator does not have the ability to create energy (e.g., an electric battery), a different generator must ultimately increase its generation at a different time for the battery to be the marginal generator. The marginal energy source refers to that generator. See Section 4.3 for a more information.

²⁰ Much of the published research on short-run marginal emission rates takes an empirical approach (Siler-Evans et al. 2012), often deriving marginal emission factors based on the changes in generation mixtures between sequential hours in data from system operators. We take the approach of identifying the marginal generator from our simulations, to maintain consistency with other metrics being reported by Cambium. Though we do maintain

Unfortunately, the marginal generator is not a native output of the PLEXOS runs that Cambium draws from. It is therefore necessary for us to analyze the PLEXOS results to make reasonable judgements as to which generator was likely the marginal generator for each node during each time-step. In this section, we describe our method for doing so, which follows these five steps:

1. Identify balancing areas (BA) that share a marginal generator (T-regions)
2. Identify T-regions with dropped load
3. Evaluate non-energy-constrained generators
4. Evaluate energy-constrained generators
5. Reevaluate non-energy-constrained generators, with relaxed filters.

Step 1: Identifying Regions That Share a Marginal Generator (T-Regions) in Each Hour

We run PLEXOS with inter-BA transmission represented as pipe flow with fixed loss rates. Given this, and knowing that the solution is a least-cost optimization, we assume any set of BAs that are connected by partially utilized transmission lines share a marginal generator during that hour. Working with this assumption, we identify all partially utilized transmission lines during each time-step, and we then identify all groupings of BAs that are connected by those lines.²¹ Because we weren't feeling particularly creative on the day we first did this, we gave these groupings of BAs the uninspired name of "transmission connected regions," using the shorthand T-regions.

The BAs that make up T-regions—and therefore the BAs we assume share a marginal generator—shift every time-step. A single BA can often be its own T-region, although it is also common for them to be large, covering dozens of BAs.

Step 2: Identifying T-Regions with Dropped Load

After identifying T-regions, we find which ones have dropped load. These T-regions will not have a marginal generator, so we label them as such.

Step 3: Evaluating Non-Energy-Constrained Generators

Having identified BAs we assume share a marginal generator for a given time-step, we try to estimate which generator in those BAs is the marginal generator. We first identify all the generators that were generating in those BAs at that point in time. We then filter out generators that are:

- At their minimum generation levels.
- At their maximum generation levels²²

consistency, the simulation-based method is highly sensitive to the accuracy and peculiarities of the dispatch model being used.

²¹ For our implementation we use the undirected graph capabilities of the python-based networkx package.

²² For variable generators like wind and solar, this is the maximum output that they can generate during that hour, given the weather conditions – not their nameplate capacity. A variable generator would only be below its maximum generator level if it is curtailing.

- Ramping their energy generation as fast as they can, relative to either the previous or the following hour
- Energy-constrained generators, which we evaluate later in step 4.²³

We assume any generator that passes these filters is more likely to be the marginal generator at that time than the generators that were filtered out. Often, only a single generator passes the filters, which is consistent with a least-cost dispatch of the grid. In that case, we select that single generator as the marginal generator for that T-region.

If multiple generators are curtailing,²⁴ we select the curtailing generator with the greatest noncurtailed generation during that time-step, and we designate that generator as the marginal generator. If there are multiple generators and none of them is curtailing, we find the one whose short-run marginal cost (SRMC) is closest to the average marginal energy cost across that T-region, and we designate that generator the marginal generator for that T-region.

Step 4: Evaluating Energy-Constrained Generators

In many instances (often in around half the T-regions), no generators make it through the filters listed in Step 3. In these T-regions, the marginal generators must be an energy-constrained generator: either an energy-constrained generator that is discharging (in which case it could discharge more to serve a marginal increase in demand) or an energy-constrained generator that is charging (in which case it could charge less to free up power to serve a marginal increase in demand).

For T-regions where no marginal generator was identified in Step 3, we estimate the short-run marginal cost (SRMC) of all the energy-constrained generators that are actively charging or discharging—but not at their maximum rates. Because the energy-constrained generators cannot produce additional power, but rather only shift it, their SRMC is estimated by finding a non-energy-constrained generator that could have increased its output in a different time-step to allow the energy-constrained generator to have more available energy during the time-step being evaluated. The SRMC of the energy-constrained generator is set by the SRMC of the non-energy-constrained generator, modified by any relevant transmission and efficiency losses. For a detailed explanation of how this is done, see Section 5.5.

Energy-constrained generators can form “chains” (e.g., a battery can charge from another battery that is ultimately charging from a thermal generator). To capture this potential, we iterate over Step 4 several times to try to identify pathways where energy-constrained generators are chaining with other energy-constrained generators.

²³ Technologies such as batteries and pumped hydro storage (which cannot ever create new electricity, only shift it around) are always energy-constrained. Technologies like dispatchable hydropower CSP with thermal energy storage – which have fixed budgets of energy flowing into them – are typically energy-constrained (if, in our PLEXOS dispatch, they dispatched all the energy that was available to them), but also can be classified as non-energy-constrained (in the more rare occurrences where they did not expend all of their available energy).

²⁴ This can happen because the solution is degenerate when curtailment is happening.

Once we have estimated the SRMC of all the energy-constrained generators in each of the T-regions being evaluated in this step, we find the one with the SRMC closest to the average marginal energy cost in each T-region, and we designate that as the marginal generator.

Step 5: Filling in Any Remaining Regions by Relaxing Filters

Once we have evaluated both energy-constrained and non-energy-constrained generators, and we have identified T-regions with dropped load, we find all the T-regions in which we have not identified a marginal generator. This is always the case for at least a small number of T-regions at this point.

For these T-regions, we repeat Step 3 (evaluating non-energy-constrained generators), but without filtering out generators at their minimum generation levels. If there are multiple such generators, we select the one whose SRMC is closest to the average marginal energy cost in the T-region. In this way, we are selecting a generator that was committed and generating at that time, and which has generation capacity remaining.

For any remaining T-regions that still do not have a designated marginal generator, we again repeat Step 3, but without filtering out any generators for minimum or maximum generation levels or for ramping. Doing so identifies a generator that is committed and generating.

Caveats and Limitations

In addition to the general caveats discussed in Section 2, our current approach for identifying marginal generators has three significant limitations:

- Because the marginal generator is not a native output of a solution to a production cost model, our method relies on trying to post-process the results and identify which generator would most likely have been the marginal generator at any point in time. Although it is clear what the marginal generator is at some points in time, it is less clear at many other points, particularly when energy-constrained generators are involved. Our method is likely not perfectly accurate in finding the marginal generator.
- Even when the marginal generator for the PLEXOS solution is correctly identified, whether the result matches real-world marginal generator patterns depends on how well the PLEXOS solutions match real unit commitment and dispatch decisions. Because we run PLEXOS as a system-wide least-cost optimization without forecast error, the PLEXOS dispatch is likely deviating from dispatches in practice, potentially in important ways. Often, for example, PLEXOS leverages energy-constrained generators to avoid starting up thermal generators in a way that is potentially too precise and would not be realized in practice. Characterization of the results of this method against real-world marginal generator practices is an ongoing effort.
- Relative to other metrics that we report, the identification of the marginal generator is highly sensitive to changes in demand. Therefore, these marginal generator patterns are likely inappropriate for analyses that assume there are significant quantities of load that is being shifted in reaction to what generator is on the margin; for example, if tens of megawatts of electric vehicle charging was timed to try and minimize how much charging was done when coal was on the margin, that would likely be enough to meaningfully change the patterns of which technologies are on the margin at what times.

Altogether, as we have mentioned, we recommend analysts approach these marginal generator patterns with a critical eye, as we work to improve our understanding of the quality and usefulness of these modeled results.

4.5 Identifying the Energy Source When an Energy-Constrained Generator is on the Margin

For some analyses, we are interested in identifying the effects of marginally increasing demand at a particular location and time. For example, the short-run marginal emission rate tells us what the increase in short-run emissions would be if demand were marginally increased. The first step in this process is identifying the marginal generator, as we discussed in Section 3.3.

For most generators that are the marginal generator, we have the information we need. If a natural gas generator is the marginal generator, for example, a marginal increase in demand would induce more generation from that generator, and we can calculate the metrics we are interested in.

For generators that are energy-constrained—meaning they cannot create their own energy (e.g., batteries or pumped hydro storage) or they have a constrained energy budget (e.g., dispatchable hydropower or CSP with thermal energy storage)—the treatment is more complicated.²⁵ Because such generators cannot create new energy, any actions by them must induce a different generator (one that can create new energy) to increase its generation at a different point in time. If we wish to know the effects of increasing demand when an energy-constrained generator is on the margin, we must therefore also identify the non-energy-constrained generator that would be induced to increase its generation as a result of the increase in demand.

In Cambium, we use the terms *marginal generator* and *marginal energy source* to describe these two generators. In this section, we describe how we try to identify the marginal energy source when the marginal generator is an energy-constrained generator.

Consider, for example, trying to determine the short-run emissions impact of increasing demand when an electric battery is on the margin. Because the electric battery cannot create new energy, but can only shift energy, we know that our increased demand from the battery must result in a the need for a different generator—one capable of creating energy—to increase its generation in a different hour, to enable the electric battery to be a marginal generator during the hour we are increasing demand. If, by demanding more energy from the battery, a coal plant would increase its generation in a different hour, and that would clearly lead to a different emissions impact than if a natural gas plant increased its generation.

This is one example of a general situation: If an energy-constrained generator is on the margin, we must find out which source-energy generator would increase its generation, at a different point in time, to enable the energy-constrained generator to increase generation. The

²⁵ Technologies such as batteries and pumped hydro storage (which can never create new electricity, but only shift it) are always energy-constrained. Technologies like dispatchable hydropower CSP with thermal energy storage, which have fixed budgets of energy flowing into them, are typically energy-constrained (if, in our PLEXOS dispatch, they dispatched all the energy that was available to them), but they also can be classified as non-energy-constrained (in the more rare occurrences where they have not expended all their available energy). Therefore, whether a generator is energy-constrained is not an immutable characteristic of the generator; it also depends on how it was dispatched.

characteristics of the source-energy generator, modified by relevant transmission and efficiency factors, define the implications of increasing our demand during the hour when the energy-constrained generator is on the margin.

For Cambium, we developed a method for identifying the source-energy generator that would most likely increase its generation, if an energy-constrained generator is the marginal generator. Our method is specifically designed to interpret the results from a production cost model: post-processing a given pattern of unit commitment and dispatch to identify the marginal source-energy generator for every energy-constrained generator. Although many of the concepts here likely transfer to similar situations in real-world dispatch, we discuss only the interpretation of simulated dispatches in this section.

Our method for identifying the source energy generator for an energy-constrained marginal generator is:

1. Identify the span of time the energy-constrained generator could have obtained more energy: its “opportunity window”
2. Reduce the span of the window if it extends beyond what is reasonable for the scheduling and forecasting assumptions of the run under consideration (e.g., restrict the window to +/- 24 hours from the time-step being analyzed)
3. Remove all time-steps where the energy-constrained generator is already charging fully
4. Remove all time-steps when no generator is available that could increase its own generation, to either charge the energy-constrained generator or cover its reduced discharge
5. From the remaining time-steps, calculate efficiency and transmission adjustments, to determine the energy-constrained generator’s SRMC if it drew from that time-step
6. Select the time-step (and associated generator) with the lowest resulting SRMC
7. Calculate derivative values, such as marginal emission rates.

Here we explain our method for this with a toy example: an electric battery that is charging and discharging over a 20-hour period. The battery’s charging and discharging patterns, and its state-of-charge, are shown in Figure 4. The battery has a maximum charge and discharge rate of 1 MW and a maximum energy storage level of 2 MWh. The battery has a round-trip efficiency of 80%.²⁶

For this example, we evaluate the battery’s behavior assuming it was identified as the marginal generator during the 11th hour, as indicated with the shaded area in Figure 3 (page 42). Because the battery itself cannot create energy, we want to identify the source-energy generator that would have increased its generation in a different time-step, in order for the battery to have the energy required to be the marginal generator during the 11th hour. Doing so allows us to calculate

²⁶ For the sake of simplicity in this toy example, we apply all the losses during charging, and we treat the charge and discharge limits as limits to the rate of change of the battery’s energy level. Increasing the stored energy level by 1 MWh requires 1.25 MWh of consumed energy, for example, and is shown as a 1-MW rate of charging in the figures.

the implications of a marginal increase in demand during the 11th hour, such as the short-run marginal emission rate.

For the battery to marginally increase its output during the 11th hour, one of two things must happen: the battery must either enter the 11th hour with a marginal amount more energy or exit the 11th hour with a marginal amount less energy. These actions would then necessitate the battery either charging more or discharging less during a different hour. For the method we are discussing here, we assume this action could have happened either before or after the 11th hour; in other words, we assume the increased demand (relative to the original system dispatch) during the 11th hour was anticipated and system operators could have planned accordingly. If we wanted to analyze the implications of an unexpected increase in demand, this method would have to be modified, but we do not explore that scenario here.

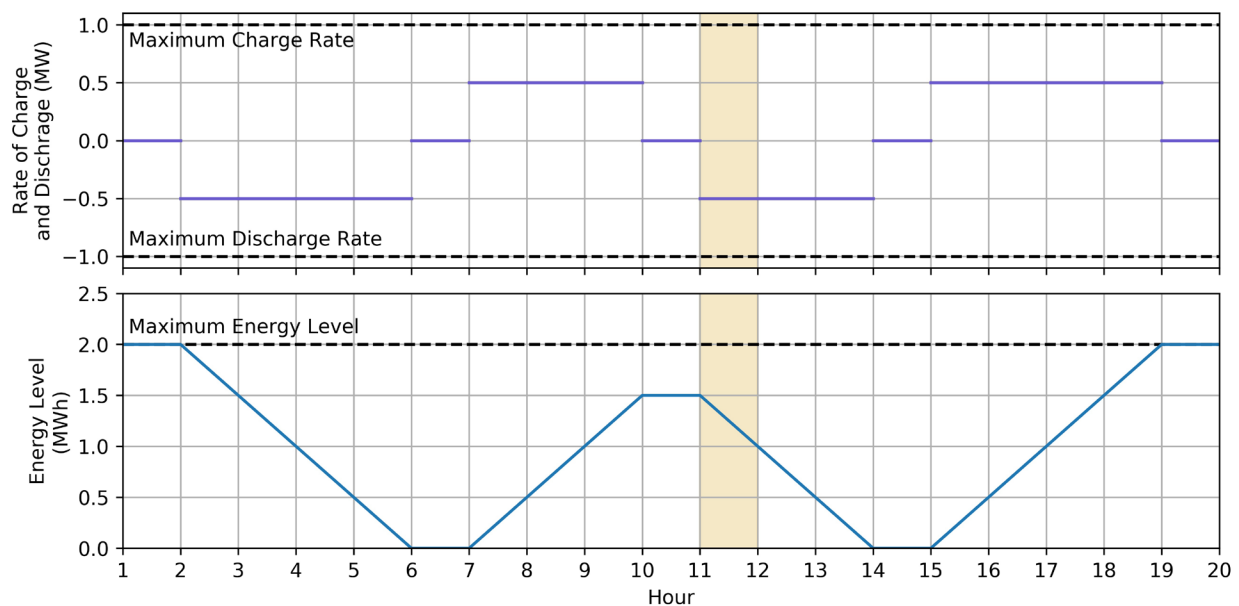


Figure 3. Charge and discharge patterns of an electrical battery

Looking before the 11th hour, we see that the battery could have obtained a marginal increase of energy at any point after the end of the 1st hour. Before the end of the 1st hour, however, it could not have obtained more energy and held it until the 10th hour, because the battery was already full during the 1st hour. A similar approach could be taken looking after the 11th hour: the battery could exit the 11th hour with an energy deficit and make it up at any point before the beginning of the 14th hour. However, once the 14th hour is reached, the battery becomes depleted, and it therefore could not hold the deficit beyond that point.²⁷

We illustrate these bounds in Figure 4. In the top panel of the figure, the 11th hour shows a marginal discharge, because that is the hour in which the battery is assumed to be on the margin.

²⁷ The assessment of these bounds illustrates a fundamental assumption of our approach: we make only marginal adjustments to the original dispatch of the battery. Clearly, the battery could be entirely redispached differently to extend these bounds beyond the 1st and 14th hours; however, we assume that, if it were not cost-optimal to dispatch the battery in that manner initially, it would also not be cost-optimal to redispach it in that manner for a marginal increase.

Starting at the beginning of the second hour and until the end of the 13th, the shaded area shows the span in which the battery would have been able to increase its stored energy, to enable the increased discharge during the 11th hour.

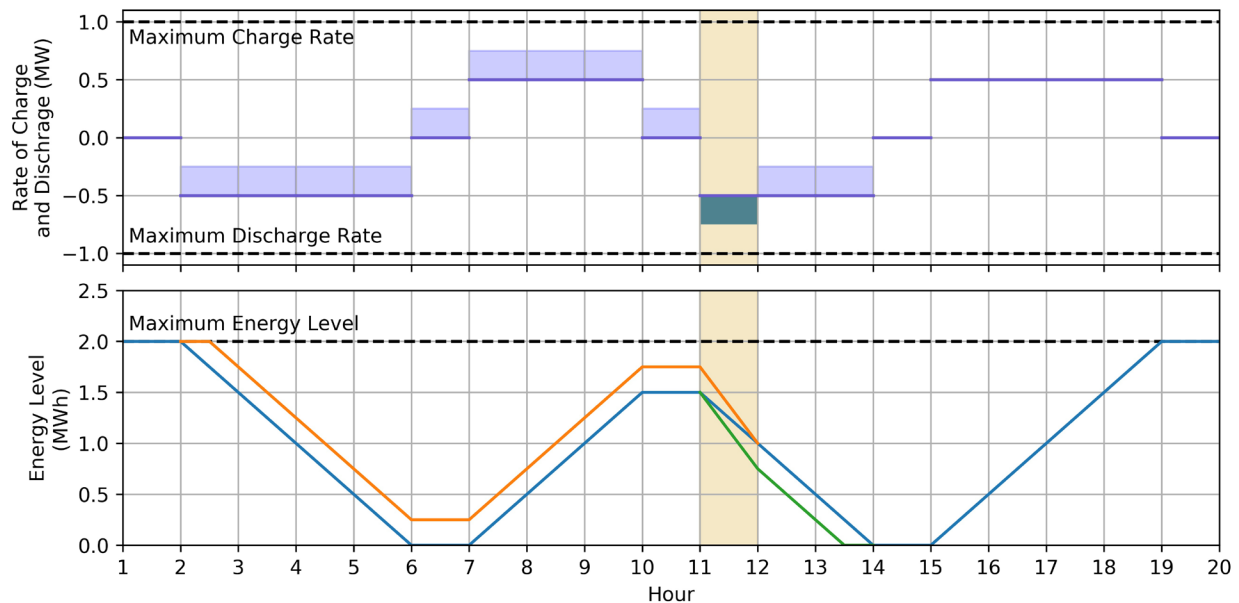


Figure 4. Example of an “opportunity window” for an electric battery

At first thought, it may seem like we are looking only for opportunities for the battery to charge (i.e., obtain more energy that it could then discharge during the 11th hour). However, this is not the case, because the increased energy does not always need to pass through the battery. For example, during the 2nd through the 5th hours and during the 12th through the 13th hours, the battery could simply have discharged less, thus reserving more energy for the 11th hour. For this to happen, a different generator would still have to increase its generation—but importantly for our purposes, there is no round-trip efficiency penalty for the scenarios where the battery discharges less, because the energy is not actually passing through the battery.

In the lower panel of Figure 5, we can visually see the general rule: the battery’s range of opportunity goes back in time as far as the most recent time that the battery was completely full, and as far forward in time as the next time that the battery is completely empty.²⁸ This establishes the bounds of an “opportunity window”—Step 1 in our list of steps given above (page 40)—which is the span of time in which the battery could either charge more or discharge less. In theory, this window could extend indefinitely. For Cambium databases, however, we restrict the windows to +/- 24 hours from the hour in which the battery is the marginal generator, to reflect practical limits to dispatch scheduling (Step 2).

²⁸ This is only for the situation where the battery is discharging during the hour that it is on the margin (i.e., we are considering the right-hand marginal). When the battery is charging (i.e., the left-hand marginal), the bounds are inverted: the prior bound is the most recent time the battery was empty, and the following bound is the next time the battery is full. We therefore have an interesting situation where the right-hand and left-hand marginals can be different. All Cambium values are currently generated for right-hand marginals.

Having identified the opportunity window, we then filter out the hours during which the battery could not have actually drawn more energy (Steps 3 and 4).

First, we remove the hours where the battery was already charging at its maximum rate. If the battery were already charging as much as it could have, it clearly could not charge more.

Next, we remove the hours where we have not yet identified a marginal generator that could have increased its generation to either to charge the battery or cover the battery’s reduced discharge. For a detailed discussion of the process by which we identify marginal generators, see Section 5.4.²⁹

After applying those two filters, we end up with a set of hours where the battery is technically capable of charging and there is a generator that can either charge the battery or cover its reduced contribution. To identify which hour, and therefore which generator, would be called on, we calculate the short-run marginal cost (SRMC) of drawing from each hour.

When calculating the SRMC, four possible situations can occur, and each has different implications for applying an adjustment for the effects of efficiency. The four situations, shown in Table 4, are defined by what the energy-constrained generator is doing in (1) the time-step in which it is on the margin (the “anchor time-step”) and (2) the time-step in which it would induce an increase in another generator (the “point time-step”).

Table 4. Efficiency Adjustments

Anchor Time-Step Behavior	Point Time-Step Behavior	Description	Efficiency Adjustment
Charging	Charging	Energy-constrained generator reduces its charging during the anchor time-step and increases its charging during the point time-step.	1.0
Charging	Discharging	Energy-constrained generator reduces its charging during the anchor time-step and reduces its discharging during the point time-step.	RTE
Discharging	Charging	Energy-constrained generator increases its discharge during the anchor time-step and increases its charging during the point time-step.	1/RTE
Discharging	Discharging	Energy-constrained generator increases its discharge during the anchor time-step and reduces its discharge during the point time-step.	1.0

RTE is the round-trip efficiency of the energy-constrained generator.

To better understand the contents of Table 4, and to show how the efficiency adjustment is applied (Step 5), we return to our toy example. We assume the only hours that had available generators were the 8th hour and the 12th hour. Because the battery is discharging during the

²⁹ Typically, an hour will be missing a marginal generator if the marginal generator in that hour is itself an energy-constrained generator – for this reason, we must iterate through the steps we are describing here, to find the situations where energy-constrained generators chain off of each other.

anchor time-step, these hours correspond to the discharging-charging and discharging-discharging situations respectively. To put some numbers on our example, we assume the battery was called on to discharge 0.25 MWh during the 11th hour, that the SRMC of the marginal generator in the 8th hour was \$10/MWh, and the SRMC of the generator in the 12th hour was \$11/MWh. As mentioned before, the round-trip efficiency of the battery is assumed to be 80%.

If the battery increased its stored energy during the 8th hour, it would have to increase its rate of charging. Therefore, the marginal generator in that hour would have to generate 0.3125 MWh more, for a cost of \$3.125. Because of losses in the battery, the delivered 0.25 MWh would have an effective SRMC of \$12.5/MWh, which is the SRMC of the marginal generator in the point time-step multiplied by the efficiency adjustment of 1/RTE.³⁰

Looking to the 12th hour, we see that because the battery was already discharging, it could increase its stored energy by simply discharging less, not charging more. Therefore, the 0.25 MWh of energy could be provided to the 11th hour at a cost of \$2.75/MWh. This is an effective cost of \$11/MWh, which is the SRMC of the marginal generator in the point time-step multiplied by the efficiency adjustment of 1; in other words, there was no efficiency penalty, because the additional energy did not pass through the battery.

Of the two hours being considered, the costs are lower if the battery draws from the 12th hour, and therefore we assume it would do so (Step 6). Note that the SRMC of the marginal generator in the 12th hour was greater than the SRMC of the marginal generator in the 8th hour, but because the energy from the 8th hour would have had to take an efficiency penalty, the 12th hour was the lower-cost solution.

If the marginal generator in the hour that was selected is a source-energy generator, we have the information we want: the characteristics of the source-energy generator will allow us to calculate the impacts of a marginal increase in demand during the 11th hour (after applying the same transmission and efficiency adjustments that we previously used for the SRMCs).

If, however, the marginal generator is another energy-constrained generator, we have a chain, and we must follow the chain until we ultimately reach a source-energy generator. When deriving values for Cambium databases, we iterate over this step several times, to identify these chains.

4.6 Calculating Time-Varying Distribution Loss Rates

Both ReEDS and PLEXOS balance load and generation at the busbar level (i.e., before distributing electricity to end users). However, Cambium databases include end-use values, so to get end-use metrics from busbar metrics, we need time-varying distribution loss rates.

Our method for calculating both average and marginal hourly distribution loss rates draws primarily from Borenstein and Bushnell (2019). As in their work, we assume 25% of annual

³⁰ If transmission losses would occur as a result of the energy-constrained generator drawing from this hour, the SRMC of the marginal generator in that hour should also be modified by the transmission losses. For this example, we assume there are no transmission losses.

distribution losses are fixed losses that do not vary with load (e.g., losses in transformers), and 75% are resistive losses that scale with the square of the flow on a line.

To remain consistent with ReEDS, we assume the annual average distribution loss rate is 5% within each balancing area (BA), meaning that if 100 GWh of busbar energy were sent to end uses over the course of the year, 5 GWh would be lost during distribution.³¹ We start by calculating the total annual fixed ($L_{f,a}$) and variable ($L_{v,a}$) losses for each BA as a function of the annual busbar load consumed for end uses ($Q_{b,a}$), the aforementioned no-load loss fraction (π) and annual loss rate (r_a).

$$L_{f,a} = Q_{b,a} * r_a * \pi$$

$$L_{v,a} = Q_{b,a} * r_a * (1 - \pi)$$

We then calculate an annual variable loss factor $f_{v,a}$:

$$f_{v,a} = L_{v,a}/Q_{b,a}^2$$

We then calculate each BA's hourly variable losses ($L_{v,h}$), using the annual variable loss factor ($f_{v,a}$) and the hourly busbar load consumed for end uses ($Q_{b,h}$):

$$L_{v,h} = f_{v,a} * Q_{b,h}^2$$

The total hourly losses ($L_{t,h}$) are then the sum of the hourly variable losses and one hour's worth of the fixed no-load losses:

$$L_{t,h} = L_{v,h} + L_{f,a}/8760$$

We can then calculate each hour's average distribution loss rate (α_h):

$$\alpha_h = L_{t,h}/Q_{b,h}$$

and each hour's marginal distribution loss rate (μ_h) as the derivative of the square of the hour's busbar load times the annual variable loss factor:

$$\mu_h = Q_{b,h} * 2 * f_{v,a}$$

³¹ This 5% value likely overstates the losses in the distribution system. The estimated grid gross loss for the entire United States reported by eGrid for 2018 was 4.87% (EPA 2020). The U.S. Environmental Protection Agency's grid gross loss estimate includes both transmission and distribution losses.

Examples of the hourly average and marginal loss rates are shown in Figure 5. These distributions are drawn from the Mid-case of the 2020 Standard Scenarios.

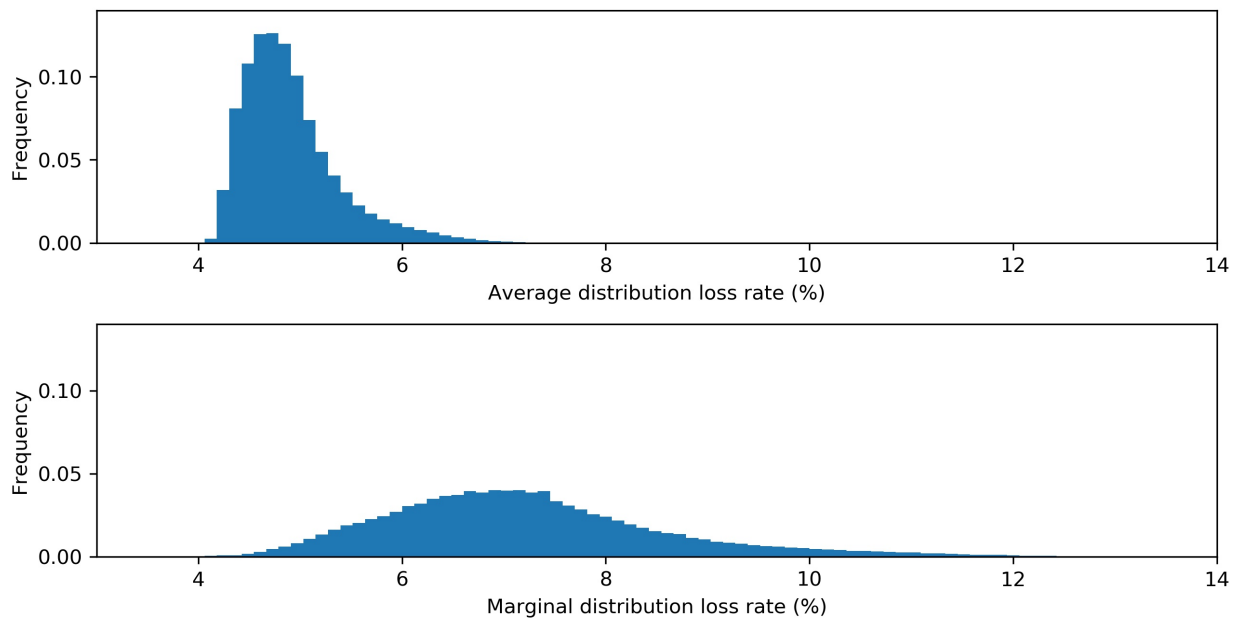


Figure 5. Examples of histograms of hourly distribution loss rates (average and marginal), drawn from the Mid-case of the 2020 Standard Scenarios

4.7 Calculating Hourly Marginal Capacity Costs

The marginal capacity costs in Cambium are estimates of the costs of acquiring sufficient firm capacity to meet a system’s planning reserve margin if there is a marginal increase in peak demand. The annual marginal cost of firm capacity is determined by our ReEDS model, which is then allocated to the highest net-load hours to produce an hourly marginal capacity cost pattern. We first explain how ReEDS determines the annual marginal cost of firm capacity, and we then explain how we allocate that value to specific hours.

Calculating an Annual Marginal Cost of Firm Capacity with ReEDS

ReEDS has a constraint that requires sufficient firm capacity to be procured in each balancing area (BA) to exceed a year’s peak demand by a planning reserve margin (e.g., if the peak busbar demand in a BA is 100 MW and the planning reserve margin is 0.15, ReEDS will require 115 MW of firm capacity).³²

³² Planning reserve margins are heuristics for the amount of capacity required to maintain a desired level of reliability in the electric system. Probabilistic resource assessments and the associated metrics, like loss-of-load probabilities, can give a more accurate assessment of the reliability of an electric grid. Cambium relies on planning reserve margins, however, because of shortcomings in the integration of these more sophisticated methods into our capacity expansion models, particularly in the presence of large amounts of variable generation and storage generators.

The shadow price off of this constraint is the \$/MW-year marginal cost for obtaining additional firm capacity. ReEDS will find the least-cost option through three possible decision variables within the model:

- **New Generation Capacity:** Referred to as net cost of new entry (net CONE), the shadow price of the capacity constraint may be set by the annualized revenue needed to recover the costs of the generator that can provide firm capacity at the lowest cost, minus any revenue that generator could obtain by providing other services (e.g., energy or operating reserves). This is often a natural gas combustion turbine plant, although in certain regions it can also be variable resources like wind and solar, if their generation aligns well with peak demand.³³
- **New Transmission Capacity:** If a neighboring BA has excess generation capacity, the shadow price of the capacity constraint may be set by the annualized cost of building additional transmission capacity, minus the revenue that the line would obtain from transmitting energy or operating reserve products.
- **Delayed Retirement:** ReEDS will choose to retire generation capacity if the capacity generator is not providing sufficient value to the system to cover its fixed costs (amplified by a multiplier to represent the “stickiness” of retirement). When this is happening, the shadow price of the capacity constraint can be set by the revenue that would have been required to keep that capacity online, minus the revenue it would have received for any other services.

Because of the prevalence of retiring generators, and the ability of wind and solar to contribute firm capacity, Cambium results in the 2020s often show capacity shadow prices that are substantially lower than what they would be if the shadow price were only being set by the net CONE of a natural gas combustion turbine.

If the capacity constraint in ReEDS is not binding, the shadow price on the constraint will be zero.

Allocating the Annual Shadow Price to Individual Hours

Having obtained an estimate of the annual marginal cost of additional firm capacity (i.e., the shadow price on the capacity constraint from ReEDS), we want to allocate that value to individual hours. Our method follows these steps:

³³ ReEDS assesses the ability of variable generators (wind and solar) to provide firm capacity through a net load duration curve approach. Doing so tends to result in variable generators being able to provide firm capacity in the near term, which eventually goes to zero as net load peaks shift away from times of peak variable generation. See (Brown et al. 2020) and the forthcoming documentation of the 2020 version of ReEDS for a more detailed discussion of this.

1. Obtain each BA's shadow price (ρ_{BA}) off of the annual capacity constraint in ReEDS.³⁴
2. Multiply the annual shadow price by (1 + planning reserve margin), to obtain the marginal cost of procuring the firm capacity that would be required by an increase in peak busbar load (δ_{BA}):³⁵

$$\delta_{BA} = \rho_{BA} * (1 + prm_{BA})$$

3. Calculate the hourly net load ($\eta_{raz,h}$) for the 18 reliability assessment zones (RAZ) shown in Figure 1 (page 9). The net load in Cambium is given by *net_load_busbar* and is the *busbar_demand_for_enduse* less generation from nondispatchable wind and solar generators.³⁶
4. Determine a threshold MW value ($\eta_{raz,threshold}$) for each RAZ that is either the net load during the 41st greatest net-load hour ($\eta_{raz,41}$) or 95% of the RAZ's annual peak net load ($\eta_{raz,1}$), whichever is lower:

$$\eta_{raz,threshold} = MIN(\eta_{raz,41}, \eta_{raz,1} * 0.95)$$

5. Calculate the total amount of each RAZ's net load that exceeds its threshold value (N_{raz}):

$$N_{raz} = \sum_{h=1}^{8760} MAX(\eta_{raz,h} - \eta_{raz,threshold}, 0)$$

6. Calculate a weight for each hour ($w_{raz,h}$) whose net load exceeds the threshold value, defined as the amount that hour's net load exceeds the threshold value divided by the total amount of load exceeding the threshold in that RAZ in that year; the weights will sum to 1.

$$w_{raz,h} = \frac{MAX(\eta_{raz,h} - \eta_{raz,threshold}, 0)}{N_{raz}}$$

7. Allocate each BA's annual marginal capacity cost (δ_{BA}) from Step 2 using the hourly weights of the RAZ that that BA is in:

$$c_{ba,h} = w_{raz,h} * \delta_{BA}$$

³⁴ ReEDS is not allowed to endogenously retire generators in the first nonhistorical year (i.e., only announced retirements are allowed), because imperfections in our representation sometimes produce unrealistically large retirements in the near term. For this reason, the capacity constraint in the first nonhistorical year is often zero, which we do not consider realistic. Therefore, if the shadow price is zero in the first nonhistorical year, we replace it with the shadow price of the second nonhistorical year.

³⁵ For example, if the planning reserve margin is 15%, 1 MW more of peak busbar demand will require 1.15 MW more firm capacity. If the shadow price for firm capacity was \$10/MW-year, the capacity cost per MW of additional peak load would be \$11.50/MW-year.

³⁶ Load from storage generators charging is not included in the net load, with the idea being that that load is flexible and in most instances the charging could be reduced (without impacting reliability) if there were a capacity shortage. However, it is possible that in certain futures there could be situations where storage would need to charge during certain periods of time for reliability reasons (e.g., charging during the day after providing required firm capacity during a morning peak, in anticipation of being needed during an evening peak). We do not capture that possibility.

Caveats and Limitations

Our method of calculating marginal capacity costs has three important limitations.

First, *our method relies on heuristics*. The use of a planning reserve margin, and the subsequent allocation of the annual capacity cost based on a net-load threshold, are heuristics. They are meant to approximately capture the costs induced by increased demand during high-net load hours, but they do not represent the most sophisticated techniques for resource adequacy assessment.

The use of the 41st-hour/95%-peak threshold, in particular, is only an approximation, although it is similar to the top-hour counts used by other models. The 2020 version of ReEDS uses the top 10 net-load hours in each of 4 seasons for assessing the capacity credit of variable resources. Hale, Stoll, and Mai (2016) used the top 100 net-load hours for estimating the capacity value of flexible resource in the Resource Planning Model. In the publicly available Avoided Cost Calculator (Energy+Environmental Economics 2016), which was developed for use in California, the default values for 2020 has 334 hours with nonzero weights for generation capacity costs, although 80% of the weight was in the top 40 hours and 99% in the top 90 hours.³⁷

We selected this approach was selected because it has three attractive features:

- The hours with the highest net loads have the greatest marginal capacity costs.
- The marginal capacity costs phase out, instead of cutting out sharply at a threshold (which would occur if costs were allocated evenly to a number of top hours).
- The sum across a year's hourly capacity costs will equal the \$/MW-year shadow price, amplified by the PRM.

A more technically sophisticated approach could involve the derivation of hourly probabilistic loss metrics (e.g., loss of load probability). These probabilistic metrics could be used to assign weights to individual hours, or to directly calculate a marginal cost of capacity by using the rate of change of the loss metrics as a function of increased load and a cost of lost load. Currently, however, we do not have a method for calculating hourly probabilistic loss metrics that can be deployed coherently with our ReEDS model. Given that the ReEDS model makes capacity investment decisions based on a combination of a planning reserve margin and net load duration curve techniques, more-sophisticated assessments of the reliability of the systems that ReEDS builds could produce nonsensical results—for example, consistently showing negligible loss of load probabilities (if ReEDS tends to over-build the electric system), and therefore showing negligible capacity costs. The use of heuristics to allocate the ReEDS-derived shadow price is a tractable solution to a complex problem.

A second limitation of our method of calculating marginal capacity costs is that *Cambium capacity costs can be low*. ReEDS often has a capacity shadow price that is significantly lower than the annualized cost of a natural gas combustion turbine (NGCT), which is sometimes taken

³⁷ The default generation capacity value allocation factors in the Avoided Cost Calculator were based on loss-of-load-probability calculations within the Renewable Energy Capacity Planning, or RECAP, model developed by Energy + Environmental Economics.

as a benchmark value for additional firm capacity in practice. These low values are often driven by the prevalence of retiring generators and the ability of variable generators to contribute firm capacity, especially in the 2020s and early 2030s. These low capacity costs may not be suitable estimates for all analyses. If an analyst has a better estimate of the marginal cost of firm capacity for the region they are studying (e.g., if they know that additional NGCT capacity will have to be procured because of relevant market structures, regardless of the contribution of variable generators), replacing Cambium’s estimate with the different estimate may be suitable.

Third, as with all marginal costs in Cambium databases, we do not provide elasticities for these marginal capacity costs. Large interventions (e.g., widespread electrification of transportation) could change load patterns sufficiently to change these marginal capacity costs: the annual capacity cost, the hourly pattern, or both. If an analysis includes a large intervention, we encourage analysts to consider directly calculating changes in peak net load using the *net_load_busbar* values in Cambium, instead of relying on these marginal price-taking values.

4.8 Calculating Marginal Portfolio Costs

Marginal portfolio costs are the costs associated with staying in compliance with renewable portfolio standards (RPS) and clean energy standards (CES), when end-use demand is increased. Treating this in post-processing is necessary for Cambium because RPS and CES are represented in our ReEDS runs but not our PLEXOS runs, which means the marginal energy costs do not have these costs embedded in them (i.e., the marginal unit of energy may not be in compliance with a state’s RPS or CES policy, necessitating remedial action to stay in compliance). In this section, we discuss how we calculate the cost of that remedial action, when necessary.

Depending on the scenario, Cambium databases can include state RPS (including technology-specific carveouts), national RPS, and CES policies. They are all handled the same way, however, so we just generically refer to policies in this section.

Calculating a Marginal Portfolio Cost

Each policy in Cambium databases is represented in ReEDS as a constraint. The shadow price on that constraint is the dollar cost of obtaining one more credit for the policy; for example, a \$10/credit shadow price on a RPS constraint is conceptually equivalent to a price of \$10 per 1 MWh of renewable energy credits (REC); however, it is different in important ways, as we explain at the end of this section.³⁸ If the policy is not binding in given state and year, the shadow price would be zero, and therefore the marginal cost would also be zero.

For documentation of how the various policies are represented in ReEDS, see Brown et al. (2020).

Using the shadow price for each policy, we calculate the cost of staying in compliance with each of the policies as end-use demand is increased. The marginal cost can change from hour to hour, depending on whether the marginal generator at that point in time can contribute to the policy

³⁸ We use the term “credit” to refer to the mechanism by which policy compliance is tracked, although it should be noted that different policies use various terms and have various tracking mechanisms in practice.

(a generator needs to be an eligible technology and either be in a location that within the region covered by the policy or be able to trade credits with the region covered by the policy).

If the marginal generator is unable to contribute to the policy (because it is either not an eligible technology or it is not in a location that can trade credits with the region covered by the policy), the marginal cost of policy n (C_n) in \$/MWh of end-use demand is given by:

$$C_n = f_n * p_n$$

where f_n is the fraction of end-use demand that must be covered by generation from an eligible technology, and p_n is the annual shadow price for policy n .

We emphasize that, when calculating f_n , we calculate the average fraction for the region covered by the policy. That can result in f_n values that deviate from the nominal top-line numbers used to describe a policy, as some load within a region is often excluded from the policy. Many states, for example, exempt utilities below a certain size from their RPS. This would result in a f_n that is lower than the nominal RPS goal.

The equation given above is straightforward: if a state has a policy where, say, 30% of end-use load must be covered by a credit, and the shadow price on credits is \$10/MWh, a 1-MWh increase in end-use demand when a noneligible generator is on the margin means 0.3 credits must be obtained for a marginal cost of \$3/MWh of end-use demand.

Notably, if the marginal generator can contribute to the policy (i.e., it is both an eligible technology and in a region that can trade with the region covered by the policy), the marginal cost can be negative (i.e., there is actually a marginal benefit to increasing consumption):

$$C_n = -\left(\frac{1}{1 - \mu_h} - f_n\right) * p_n$$

where μ_h is the marginal distribution loss rate at that point in time, as explained in Section 5.3.

Conceptually, the reason the cost can be negative is because the additional consumption can create more credits than are required by the policy to cover that additional consumption. Because the credits have value, the excess credits count as a benefit.

Consider a situation where an eligible technology is on the margin (e.g., an in-region solar generator is currently curtailing), the marginal distribution loss rate is 5%, the policy covers 30% of end-use demand, and the shadow price is \$10/MWh. In this case, an increase in end-use demand of 1 MWh increases generation from the solar plant by 1.053 MWh (amplified slightly because of the distribution losses). However, only 0.3 MWh of credits are needed to cover the additional load, so 0.753 MWh of credits remain and they are valued at \$10 per credit. Therefore, there is a marginal benefit of \$7.53/MWh of end-use demand.

For busbar marginal costs, we modify the end-use marginal costs by the marginal distribution loss rate.

Caveats and Limitations

Our representations of marginal portfolio costs have two important limitations:

- **Incomplete Policy Representations:** The shadow prices used to calculate marginal portfolio costs are driven by the representations of the policies in ReEDS, which can be incomplete. Though significant effort is put into correctly categorizing technology eligibility, fractions of load covered, and trading restrictions, there are still missing components (e.g., inter-year REC banking and technology-specific multipliers).
- **Shadow Prices are Long-Run Values:** Because the annual shadow prices on each policy come from ReEDS, which solves for the long-run equilibrium position, the shadow prices themselves are long-run values. In other words, they incorporate the option of building new capacity to generate credits. Where credits are traded in practice, their prices would potentially be better described as a short-run prices, although the ability of banking RECs better years makes the distinction less clear. Altogether, we recommend the policy shadow prices not be used directly as forecasts of future prices of these credits, as the markets for the credits likely deviate in significant ways from our modeled representations.

4.9 Operating Reserve Requirements and Limitations

Cambium represents three operating reserve products, which are simplified representation of the reserves held to ensure supply-side flexibility is sufficient to address fluctuations in the balance of supply and demand. In order of fastest to slowest response time, the three types of reserves are:

- *Regulating reserves* are reserves held to provide rapid response for restoring system frequency. They may be deployed after an event and are also used to address normal random short-term fluctuations in the balance of supply and demand.
- *Spinning reserves* are held to supply power if there are failures in generators or transmission lines.
- *Flexibility reserves* are an emerging and evolving reserve product that is used to address relatively “slow” variations in net load, such as those caused by variability in net load from wind and solar generators. They are also known as *load-following* or *ramping* reserves.

The amount of each reserve product that is required in each hour is shown in Table 2.

Table 5. Operating Reserve Requirements

Reserve Product	Load Requirement (% of load) ^a	Wind Requirement (% of generation) ^b	PV Requirement (% of capacity) ^b	Time Requirement to Ramp (minutes)
Regulation	1%	0.5% ^c	0.3% ^c	5
Spinning	3%	—	—	10
Flexibility	—	10%	4%	60

^a See Lew et al. (2013), Section 5.3.4

^b Reserve requirements for wind and PV are derived from the outcomes of Lew et al. (2013). The flexibility requirement for wind is estimated as the ratio of the change in the reserve requirement to

the change in wind generation from the Lew et al. High Wind scenario; the requirement for PV was similarly estimated using the Lew et al. High Solar scenario.

° The estimated regulation requirements (0.5% wind generation and 0.3% PV capacity) are based on incremental increases in regulation reserves across all scenarios in Lew et al. (2013).

Flexibility and regulation reserve products are influenced by PV capacity, but only when the PV capacity is generating. Capacity is used instead of generation because these reserves are especially important around sunset and sunrise, when PV generation is lower.

In the PLEXOS runs that Cambium draws from, the quantity of operating reserves is determined for 18 reliability assessment zones (RAZ), shown in (Figure 1, page 9). Any eligible generator within a RAZ can contribute to the reserves required for the RAZ: there are no transmission capacity requirements. This is a simplification. It would be more realistic to have the reserve products defined for each BA and require bilateral trading within a RAZ. For the reserve product quantities in Cambium databases, we report the amount of reserve product induced by the load and generation within each BA, although it should be noted that the PLEXOS runs balance the operating reserves at the RAZ level.

Unless otherwise specified in a particular scenario’s description, nuclear, wind, PV, and CSP without storage cannot provide operating reserves. All other generators can provide reserves, limited by their availability capacity (their online capacity minus any capacity being used for energy generation) and their ramp rates. Table 6 (page 53) shows the ramp rates used for different technologies in Cambium’s PLEXOS runs, and the resulting amount of capacity that each generator can dedicate toward each reserve product. For example, natural gas CTs (NGCT) can ramp 8% of their capacity per minute, and the regulation reserve product requires a 5-minute ramp, therefore NGCTs can dedicate 40% of their capacity to regulation reserves.

Energy storage technologies can provide reserves, as long as they have at least one hour’s worth of stored energy. Generators are allowed to provide both flexibility and spinning reserves with the same capacity.

Table 6 actually overstates the capacity available for reserves for most technologies, because minimum generation constraints force generators to dedicate a fraction of their capacity to energy generation if they are online. For example, because NGCTs have a minimum generation level of 55%, they can actually only ever dedicate 45% of their capacity to flexibility and spinning reserves.

Table 6. Technology Flexibility Parameters

	Assumed maximum ramp rate (%/min)	Maximum Capacity Available for Reserves ^a		
		Flexibility (60- min. ramp requirement)	Spinning (10-min. ramp requirement)	Regulation (5-min, ramp requirement)
Natural Gas CT	8	100	80	40
Natural Gas CC	5	100	50	25
Coal	4	100	40	20

	Maximum Capacity Available for Reserves ^a			
	Assumed maximum ramp rate (%/min)	Flexibility (60- min. ramp requirement)	Spinning (10-min. ramp requirement)	Regulation (5-min, ramp requirement)
Geothermal	4	100	40	20
CSP with Storage	10	100	100	50
Biopower	4	100	40	20
Landfill gas	4	100	40	20
Oil-gas-steam	4	100	40	20
Hydro	100	100	100	100
Storage	100	100	100	100

^a (% of online capacity) = Ramp Rate (%/min) * Ramp Requirement (min)

Caveats and Limitations

Our representation of operating reserves has several important caveats and limitations:

- **Simplified Representations:** These are stylized, simple representations of reserve products that are meant to capture the need to hold capacity in reserve for reliable operation of the grid. Actual markets may have different products, or they have similarly named products with different definitions.
- **Lack of Direct Operating Costs:** We do not assign an operating cost to generators for providing these reserves, which means these values do not reflect any direct costs that would be incurred by a marginal increase in reserve provision.
- **Generator Capacity is Committed in Blocks:** Because the PLEXOS runs that Cambium draws from represent the commitment of generators as an integer decision, each generator's capacity is committed as a block. As a result, most hours and regions tend to have at least some online capacity that is not already dedicated to energy generation and reserve provision. When combined with the previously mention absence of direct operating costs, this tends to result in the shadow price on the operating reserve constraints being zero in Cambium databases. Therefore, the shadow prices for operating reserve constraints in Cambium are not good estimates of market prices for similar reserves, which tend to be non-zero in practice.
- **Incomplete Representation of Services:** Not all operating reserves are represented. Frequency-response reserves, which operate on a much faster timescale than our PLEXOS runs, are not represented. Other reliability services, such as black-start and voltage support, are also not captured by these products.
- **Reserve Requirements for Some Futures are Poorly Understood:** Some of our scenarios are significantly different from the current grid (e.g., large amounts of variable generators and storage). The types and quantities of reserves needed to sustain reliable grid operation in these future scenarios is an ongoing area of research, and the quantities represented here may be too much or too little.

5 References

- A. D. Hawkes. 2010. “Estimating Marginal CO₂ Emissions Rates for National Electricity Systems.” *Energy Policy* 38 (10): 5977–87.
- . 2014. “Long-Run Marginal CO₂ Emissions Factors in National Electricity Systems.” *Applied Energy* 125: 197–205.
- Bloom, Aaron, Aaron Townsend, David Palchak, Joshua Novacheck, Jack King, Clayton Barrows, Eduardo Ibanez, et al. 2016. “Eastern Renewable Generation Integration Study.” NREL/TP-6A20-64472. National Renewable Energy Lab. (NREL), Golden, CO (United States). <https://doi.org/10.2172/1318192>.
- Borenstein, Severin, and James Bushnell. 2019. “Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency.” Working Paper Energy Institute WP 294R. Energy Institute at HAAS: Energy Institute at HAAS.
- Brown, Maxwell, Wesley Cole, Kelly Eurek, Jon Becker, Dave Bielen, Ilya Chernyakhovskiy, Stuart Cohen, et al. 2020. “Regional Energy Deployment System (ReEDS) Model Documentation: Version 2019.” NREL/TP-6A20-74111. Golden, CO: National Renewable Energy Laboratory. <https://doi.org/10.2172/1505935>.
- Callaway, Duncan S., Meredith Fowlie, and Gavin McCormick. 2017. “Location, Location, Location: The Variable Value of Renewable Energy and Demand-Side Efficiency Resources.” *Journal of the Association of Environmental and Resource Economists* 5 (1): 39–75. <https://doi.org/10.1086/694179>.
- Cole, W., B. Frew, T. Mai, Y. Sun, J. Bistline, and G. Blanford. 2017. “Variable Renewable Energy in Long-Term Planning Models: A Multi-Model Perspective.” NREL/TP-6A20-70528. EERE, NREL, EIA, EPA ... <https://www.osti.gov/scitech/biblio/1416124-variable-renewable-energy-long-term-planning-models-multi-model-perspective>.
- Cole, Wesley, Sean Corcoran, Nathaniel Gates, Daniel Mai, Trieu, and Paritosh Das. 2020. “2020 Standard Scenarios Report: A U.S. Electricity Sector Outlook.” NREL/TP-6A20-77442. Golden, CO: National Renewable Energy Laboratory. <https://www.nrel.gov/docs/fy21osti/77442.pdf>.
- Cole, Wesley, Nathaniel Gates, Trieu Mai, Daniel Greer, and Paritosh Das. 2019. “2019 Standard Scenarios Report: A U.S. Electricity Sector Outlook.” NREL/TP-6A20-74110. Golden, CO: National Renewable Energy Laboratory. <https://doi.org/10.2172/1481848>.
- Cole, Wesley, Daniel Greer, Jonathan Ho, and Robert Margolis. 2020. “Considerations for Maintaining Resource Adequacy of Electricity Systems with High Penetrations of PV and Storage.” *Applied Energy* 279 (December).
- Energy+Environmental Economics. 2016. “Avoided Cost Calculator User Manual.” Energy+Environmental Economics.

Energy Exemplar. 2019. *PLEXOS Integrated Energy Model*. Energy Exemplar. <https://energyexemplar.com/solutions/plexos/>.

EPA. 2020. “EGRID 2018 Gross Grid Loss Estimates.” Washington, D.C.: EPA. <https://www.epa.gov/egrid/download-data>.

Frew, Bethany, Wesley Cole, Paul Denholm, Will Frazier, Nina Vincent, and Robert Margolis. 2019. “Sunny with a Chance of Curtailment: Operating the US Grid with Very High Levels of Solar Photovoltaics.” *IScience* 21 (November): 436447.

Hale, Elaine, Brady Stoll, and Trieu Mai. 2016. “Capturing the Impact of Storage and Other Flexible Technologies on Electric System Planning.” *Renewable Energy*, 91.

Janusz Bialek. 1996. “Tracing the Flow of Electricity.” *IEEE Proceedings - Generation, Transmission, and Distribution* 143 (4): 313–20.

Lew, D., G. Brinkman, E. Ibanez, B. M. Hodge, M. Hummon, A. Florita, and M. Heaney. 2013. “The Western Wind and Solar Integration Study Phase 2.” NREL/TP-5500-55588. National Renewable Energy Lab. (NREL), Golden, CO (United States). <https://doi.org/10.2172/1095399>.

NERC. 2010. “2010 Long-Term Reliability Assessment.” North American Electric Reliability Corporation. https://www.nerc.com/files/2010_ltra_v2-.pdf.

———. 2020. “Glossary of Terms Used in NERC Reliability Standards.” North American Electric Reliability Corporation. https://www.nerc.com/files/glossary_of_terms.pdf.

NREL. 2020. “2020 Annual Technology Baseline.” Golden, CO: National Renewable Energy Laboratory. <https://atb.nrel.gov/>.

Sigrin, Benjamin, Michael Gleason, Robert Preus, Ian Baring-Gould, and Robert Margolis. 2016. “The Distributed Generation Market Demand Model (DGen): Documentation.” NREL/TP-6A20-65231. Golden, CO: National Renewable Energy Laboratory. <http://www.nrel.gov/docs/fy16osti/65231.pdf>.

Siler-Evans, Kyle, Inês Lima Azevedo, and M. Granger Morgan. 2012. “Marginal Emissions Factors for the U.S. Electricity System.” *Environmental Science & Technology* 46 (9): 4742–48. <https://doi.org/10.1021/es300145v>.

Thind, Maninder P. S., Elizabeth J. Wilson, Inês L. Azevedo, and Julian D. Marshall. 2017. “Marginal Emissions Factors for Electricity Generation in the Midcontinent ISO.” *Environmental Science & Technology* 51 (24): 14445–52. <https://doi.org/10.1021/acs.est.7b03047>.