Ranking Stranded Interconnection Capacity for Complementary Generation under FERC Surplus Interconnection Service

Dr Stuart Woolley
Full Stack Energy

August 2026

Abstract

A renewable interconnection is built and consented to a fixed export limit, yet renewable output reaches that limit only intermittently. For much of the year a large fraction of the connection sits idle. This idle headroom is stranded infrastructure: grid access already built and paid for, but not earning. United States interconnection reform now lets complementary generation use it, a gas unit exporting through the headroom the renewable plant leaves idle. This paper defines a compact, transparent method for measuring that stranded capacity and ranking candidate sites by it, cheaply enough for geographic screening across many locations. The method deliberately avoids mixed-integer programming, power-flow analysis, and any solver, so a result is both defensible to an engineer and explainable to a chief financial officer. A worked example, an economic unit-sizing rule, a validation against the exact constrained dispatch optimum, and the method's limitations are given. The method is deliberately elementary: it is the least quantitative work that will honestly separate promising sites from hopeless ones, enough to justify a prototype and to point a far heavier investigation at the sites that merit it. That deeper programme is set out in closing.

1 Introduction and Regulatory Basis

The commercial opportunity this paper screens for is created by regulation, so the regulation is the right place to start. Under *Surplus Interconnection Service* (SIS), an additional generating facility may use the unused portion of an existing interconnection customer's approved interconnection service, sited behind a single point of interconnection, provided the combined output never exceeds the approved service level. SIS was established by FERC Order No. 845 [1] in 2018 and materially expanded by FERC Order No. 2023 [2] in 2023, which requires transmission providers to open the SIS process earlier and eases the co-location of resources behind one connection. The service is technology-neutral: although the reforms were driven by renewables and storage, nothing in SIS restricts the complementary unit to a particular technology, so a gas turbine, a battery, or any dispatchable resource may in principle occupy the headroom.

One feature of the service is worth stating plainly: surplus interconnection service is typically non-firm and subordinate, so the surplus generator is interruptible when the original customer needs its capacity and gains no incremental deliverability without a further study. It is spare capacity, not firm capacity, which suits a low-capital merchant addition rather than a plant that must be guaranteed its output.

The consequence is a screenable class of siting opportunities. An existing solar or wind plant whose interconnection is rarely saturated carries latent value, because a complementary unit placed behind the same connection can export through the headroom the renewable plant leaves idle, without triggering the multi-year queue and network upgrades that a fresh interconnection would require. The screening problem is to rank candidate sites by that latent value cheaply enough to run across many locations at once, over geographic layers, without the cost and opacity of a full optimisation. The method below answers that problem with a single composite score in the range nought to one hundred, in which every quantity has a physical or commercial meaning. It is intended as a first rung and no more: its job is to earn a prototype the right to exist and to aim the deeper investigation of Section 9 at the sites worth the effort.

2 The Headroom Metric

Each candidate site is described by four quantities. The complementary unit is written as a gas unit for concreteness, but the arithmetic is agnostic to its technology.

- the interconnection limit P_{ic} in MW, the maximum export the connection permits;
- the installed renewable AC capacity P_{ren} in MW;
- a time-varying renewable capacity factor $CF_{\text{ren}}(t)$ in the range nought to one; and
- a candidate complementary unit size P_{gt} in MW.

Renewable output at hour t is the installed capacity scaled by its capacity factor,

$$P_{\text{ren,out}}(t) = P_{\text{ren}}\, CF_{\text{ren}}(t). \tag{1}$$

The spare interconnection headroom is whatever export capacity the renewable plant leaves unused, floored at zero so that a saturated connection contributes nothing,

$$H(t) = \max\big(0,\ P_{\text{ic}} - P_{\text{ren,out}}(t)\big). \tag{2}$$

The complementary unit can export into that headroom up to its own rating, so the usable export at hour t is the smaller of the candidate unit size and the available headroom,

$$P_{\text{gt,use}}(t) = \min\big(P_{\text{gt}},\ H(t)\big). \tag{3}$$

Summing the usable export over a representative year at hourly resolution gives the core value metric, the annual energy the unit could export through spare interconnection capacity,

$$H_{\text{usable}} = \sum_{t=1}^{8760} P_{\text{gt,use}}(t). \tag{4}$$

This single figure, in MWh per year, captures both how much headroom exists and how often it is available.

Optimality of the closed form. Equation (4) is not merely a heuristic. Consider maximising the complementary unit's annual exported energy $\sum_t x(t)$ subject only to the unit rating and the interconnection limit, that is $0 \le x(t) \le \min(P_{\text{gt}}, H(t))$ for every hour. The problem is separable across hours with no inter-temporal coupling, so each $x(t)$ is maximised independently at its own upper bound $\min(P_{\text{gt}}, H(t))$, and the optimum is exactly H_{usable}. The closed form is therefore the exact energy-capture optimum in the absence of dispatch constraints such as minimum stable load or minimum run time. Section 6 quantifies how far real dispatch constraints move the answer.

What the metric measures. H_{usable} is the most energy the headroom could pass in a year. It is the value captured directly by a must-run or PPA-contracted unit paid per MWh; a merchant unit dispatched on price would export less energy but at higher value, a refinement deferred to Section 9. The screen ranks this envelope, which bounds either case and needs no price data.

3 Scoring and Ranking

To rank sites against one another the annual energy is normalised across the *feasible* candidate set, the sites at which a complementary unit could actually be built. With $H_{\max} = \max_j H_{\text{usable},j}$ taken over feasible sites j (feasibility is the gate G defined below), site i receives

$$S_{H,i} = \frac{H_{\text{usable},i}}{H_{\max}}. \tag{5}$$

The score lies in the range nought to one, and the strongest feasible site scores one. Normalising over feasible sites only stops a site the method itself rejects from setting the ceiling the others are judged against.

Gas proximity. A site with abundant headroom is worthless if a gas unit cannot economically be connected to fuel. Let d_{gas} be the distance from the site to the nearest gas pipeline. Two penalty forms are offered, both yielding a factor S_G in the range nought to one. The linear form, with a maximum sane distance $D_{\max}$ such as ten kilometres, is the easier to explain,

$$S_G = \max\left(0,\ 1 - \frac{d_{\text{gas}}}{D_{\max}}\right), \tag{6}$$

while the exponential form, with a decay constant k, is smoother and never reaches zero,

$$S_G = e^{-k\, d_{\text{gas}}}. \tag{7}$$

For a fuel-free resource such as a battery this term is dropped and the weights renormalised.

Grid sanity. Where additional grid intelligence exists, such as a known congestion flag or a node marked unsuitable, it enters as a multiplicative modifier S_{grid} in the range nought to one, where one means no known issue and nought means do not touch. Where no such data exists the modifier is set to one for every site, leaving the ranking unchanged.

Composite score. The additive terms rank sites by quality, but they cannot express infeasibility: a site with no viable gas connection should be excluded, not merely penalised. Feasibility is therefore carried on a separate axis, a multiplicative gate $G \in \{0, 1\}$ that folds the grid modifier together with a hard maximum economic gas distance,

$$G = S_{\text{grid}} \cdot \mathbf{1}[\, d_{\text{gas}} \le D_{\max} \,]\,. \tag{8}$$

The composite score is the quality score gated by feasibility, on a nought-to-one-hundred scale,

$$S = 100\,\big(w_H\, S_H + w_G\, S_G\big)\, G. \tag{9}$$

Headroom is the dominant quality term and gas distance the secondary one, so the recommended weights are $w_H = 0.7$ and $w_G = 0.3$. Keeping feasibility and quality on separate axes is deliberate. The additive S_G still expresses how much better a closer site is among feasible candidates, while the gate removes sites that cannot be built at all: a site beyond $D_{\max}$ from gas, or flagged unusable on the grid, scores zero regardless of its headroom.

4 Worked Example

The method is applied to a set of six candidate sites, each pairing an interconnection limit with an installed solar plant and a candidate complementary unit. The worked unit is a must-run gas turbine, chosen for arithmetic simplicity because its value is simply its exported energy; a flexible, price-dispatched unit fits the non-firm service better and is valued in the deeper rung of Section 9. The renewable profile $CF_{\text{ren}}(t)$ is a synthetic clear-sky photovoltaic series computed from solar geometry at latitude 35°N, scaled to a representative annual capacity factor of 0.24; its summer-noon value is 0.91 and its winter-noon value is 0.49. The method is agnostic to the profile, and a metered series would be substituted in practice.

Take Site A: $P_{\text{ic}} = 100$ MW, $P_{\text{ren}} = 80$ MW of solar, a candidate $P_{\text{gt}} = 50$ MW unit, $d_{\text{gas}} = 2$ km. At summer noon the solar plant produces $80 \times 0.91 = 73.1$ MW, leaving $H = 26.9$ MW of headroom, so the unit is throttled to 26.9 MW. At night the solar output is zero, headroom is the full 100 MW, and the 50 MW unit runs flat out. Summed over the year, $H_{\text{usable}} = 421$ GWh, a slot capacity factor of 0.96: the headroom is so persistently large that a 50 MW unit could run near continuously. Normalised against the strongest feasible site (Site B, 500 GWh) this gives $S_H = 0.843$. The gas penalty is $S_G = 1 - 2/10 = 0.80$; the site is within gas range and grid-clear, so the feasibility gate is $G = 1$, giving

$$S = 100\,(0.7 \times 0.843 + 0.3 \times 0.800) \times 1 = 83.0.$$

Table 1 gives the full candidate set and their scores. Figure 1 shows Site A's headroom as a duration curve, the shaded area being exactly H_{usable}.

Site	P_{ic} MW	P_{ren} MW	P_{gt} MW	d_{gas} km	S_{grid}	G	H_{usable} GWh/yr	S_H	S_G	S
C	100	95	50	1	1.0	1	402	0.805	0.900	83.4
A	100	80	50	2	1.0	1	421	0.843	0.800	83.0
B	150	140	60	6	1.0	1	500	1.000	0.400	82.0
F	100	130	50	4	1.0	1	353	0.706	0.600	67.4
D†	200	90	80	12	1.0	0	701	–	0.000	0.0
E‡	120	100	50	3	0.0	0	426	–	0.700	0.0

Table 1: Candidate sites ranked by the composite score S, with S_H normalised over the feasible set (strongest feasible site: B, 500 GWh). †Site D lies 12 km from gas, beyond $D_{\max} = 10$ km, so the feasibility gate $G = 0$ excludes it despite having the largest raw headroom. ‡Site E is gated out by the grid modifier. Both are excluded from the normalisation and score zero. The leading three sit within 1.4 points (Section 7).

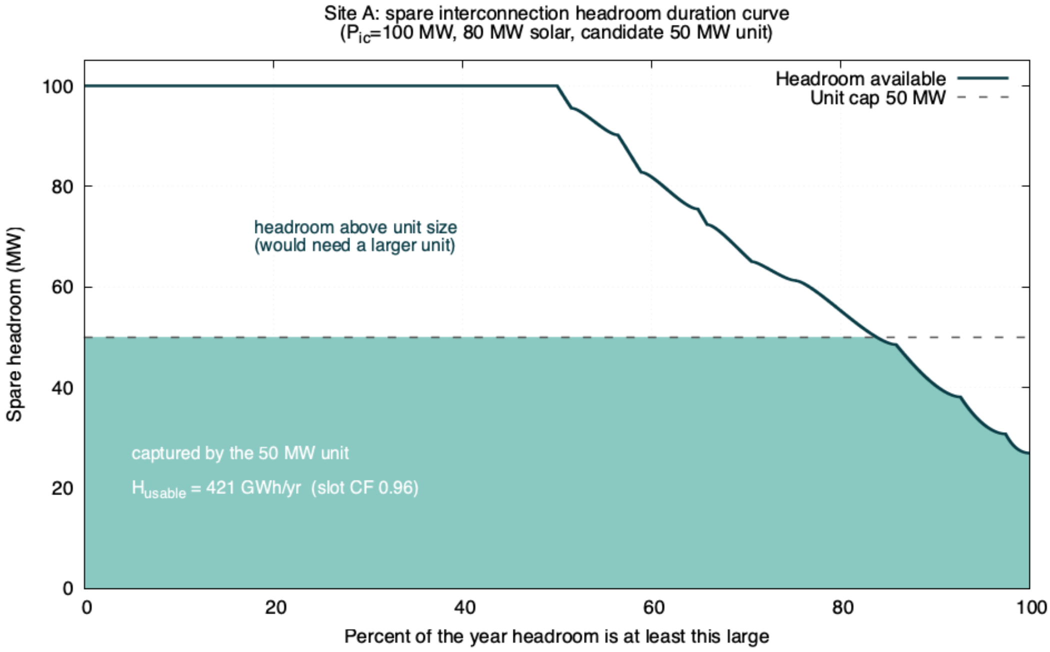

Figure 1: Site A spare-headroom duration curve. The teal area is the energy a 50 MW unit captures, equal to H_{usable}; the curve above the dashed unit cap is headroom too large for a 50 MW unit to use.

5 Unit Sizing and Indicative Economics

The screen so far takes the unit size P_{gt} as given. It need not: the same headroom series fixes the economically optimal size. Raising the unit by one MW captures extra energy only in the hours the headroom already exceeds the current size, so the marginal annual energy per MW is $N(P) = |\{t : H(t) > P\}|$, the exceedance hours read straight off the duration curve of Figure 1. With an indicative margin (spark spread) s in dollars per MWh and an annualised capacity cost a in dollars per MW-year, annual profit is $s\,H_{\text{usable}}(P) - a\,P$, maximised where the marginal MW just pays for itself,

$$s\,N(P^\star) = a \qquad \Longleftrightarrow \qquad N(P^\star) = a/s. \tag{10}$$

The optimal size is simply the headroom level exceeded for a/s hours a year. Figure 2 plots the margin curve for Site A under an indicative energy-only spark spread $s = \$20$ per MWh, as in a market such as ERCOT, and $a = \$130$k per MW-year, so $a/s = 6500$ h, peaking at $P^\star = 62$ MW. In a capacity market such as PJM a capacity payment would add revenue the screen does not price.

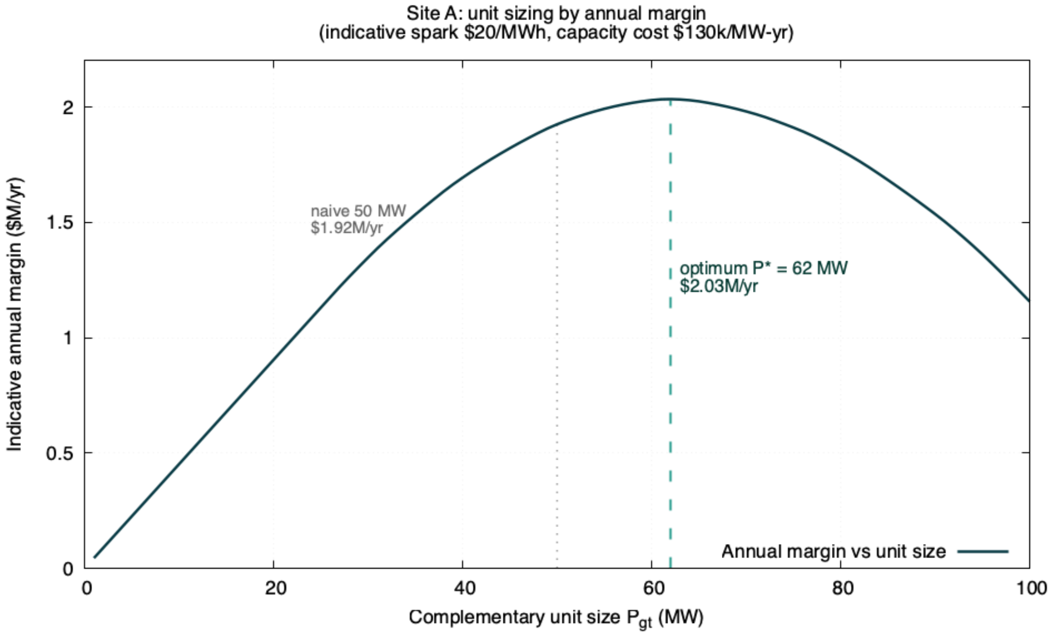

Figure 2: Site A annual margin against unit size. The interior optimum at $P^\star = 62$ MW follows from Equation (10); the 50 MW unit assumed earlier is below it.

Site	$P^\star$ (MW)	$H_{\text{usable}}(P^\star)$ (GWh/yr)	Indicative margin ($M/yr)
A	62	505	2.03
B	83	661	2.43
C	55	436	1.57
F	38	278	0.63

Table 2: Optimal unit size and indicative annual margin for the feasible sites, at $s = \$20$/MWh and $a = \$130$k/MW-yr.

Two things follow. The uniform 50 MW assumed earlier is wrong in both directions: too small for the fat-headroom sites A and B, and too large for F, whose oversized solar leaves thinner headroom and wants only 38 MW. And the margins are indicative only: they price the must-run energy of Section 2 at a flat spread, which a price-taking merchant unit would not achieve, so a real revenue figure needs the price-aware dispatch of Section 9. The sizing *rule*, by contrast, is robust, since it depends only on the ratio a/s and the headroom duration curve, not on the absolute price.

6 Validation against the Constrained Optimum

The closed form is the exact optimum only when the complementary unit dispatches freely. A real thermal unit has a minimum stable load and minimum run times, so the recoverable energy is lower. To bound the gap, each site's hourly headroom series was passed to an exact dynamic program that maximises annual export subject to a minimum stable load of 40% of P_{gt}, a minimum up time of four hours, and a minimum down time of two hours. The dynamic program returns the true constrained optimum, and as a check it reproduces H_{usable} exactly when those constraints are relaxed.

Site	Closed form (GWh)	Constrained optimum (GWh)	Recoverable
A	421	421	100.0%
B	500	497	99.4%
C	402	395	98.3%
D	701	701	100.0%
E	426	426	100.0%
F	353	345	97.8%

Table 3: Closed-form usable energy against the exact constrained-dispatch optimum. The proxy captures between 97.8% and 100% of the recoverable energy, and the site ranking is identical under both measures.

Two results matter. First, the proxy is robustly near-exact. It recovers between 97.8% and 100% of the constrained optimum across the set, and even a deliberately harsh stress case, 180 MW of solar behind a 100 MW connection with the unit held to a 50% minimum stable load and a six-hour minimum run, still recovers 96.0%. The reason is structural: solar headroom arrives in long overnight blocks that comfortably outlast any minimum run time, so unit-commitment constraints rarely bind on the question of *how much* energy can be captured. The screen therefore needs no dispatch model. Dispatch sophistication earns its keep on a different question, *when* to run for profit against volatile prices, taken up in Section 9. Second, and more important for a screening tool, the ranking of sites is identical whether measured by the cheap closed form or the exact optimum, so the screen ranks correctly even where it slightly overstates absolute energy.

7 Sensitivity of the Ranking

A ranking is only as firm as its weights and thresholds, so both were varied.

Weights (w_H, w_G)	Feasible ranking (score)
0.6/0.4	C 84.3 > A 82.6 > B 76.0 > F 66.4
0.7/0.3	C 83.4 > A 83.0 > B 82.0 > F 67.4
0.8/0.2	B 88.0 > A 83.4 > C 82.4 > F 68.5

Table 4: Feasible ranking under three headroom/gas weight splits. The leading cluster reorders; only F's position is stable.

Under the recommended 0.7/0.3 split the order is C, A, B, but the three sit within 1.4 points and reorder under reasonable variation: at 0.8/0.2, which rewards raw headroom, Site B leads. The feasibility threshold matters too. Widening $D_{\max}$ from 10 to 12 km readmits Site D, which then tops the table on headroom alone. The honest reading is that the screen reliably separates feasible from infeasible and strong from weak, Site F and the gated D and E are clearly distinguished, but it does not finely rank the leading cluster. C, A, and B should advance to detailed study together, not be ordered against one another on a one-point difference. A screen that produces a defensible shortlist is doing its job; one that claims a precise ranking from judgemental weights is overreaching.

8 Limitations

- **Feasibility is a gate, not a weight.** Gas proximity plays two roles that must not be conflated: within economic range it is a quality factor, carried by the additive S_G, while beyond $D_{\max}$ it is a hard infeasibility, carried by the gate G of Equation (8). Separating them stops a site that cannot be connected from surviving on headroom alone, as Site D shows. The threshold $D_{\max}$ is a judgement and should reflect real pipeline tie-in economics rather than a round number.
- **The proxy is an upper bound on dispatch.** H_{usable} ignores minimum-load and run-time constraints, overstating recoverable energy by a few per cent (Section 6). It is a screen, not a dispatch model.
- **The renewable profile is synthetic.** The worked example uses a clear-sky photovoltaic proxy, not metered data. A real screen would use measured or reanalysis capacity factors, which carry cloud and temperature effects the proxy omits.
- **The output is a shortlist, not a fine ranking.** The weights and $D_{\max}$ are judgemental, and Section 7 shows the leading cluster reorders under them. Read the screen as separating strong from weak, not as a precise order.

- **The economics are indicative.** The margins price must-run energy at a flat spread (Sections 2 and 5); a real revenue figure needs a price-aware dispatch. The sizing *rule* is unaffected.

9 From Screen to System

The screen answers one question cheaply: which sites are worth a closer look. A committed build must answer many harder ones, each a distinct body of quantitative work for which this screen is only the entry point.

1. **Unit commitment.** The dynamic program of Section 6 is the single-unit, energy-only special case. A real dispatch is a mixed-integer unit-commitment problem with start-up costs, ramp rates, and minimum run times across the whole plant.
2. **Price-aware, stochastic dispatch.** Headroom is worth most when prices are high, and both are uncertain. Co-optimising export against day-ahead and real-time prices under uncertainty is a multistage stochastic optimisation, for which stochastic dual dynamic programming is the standard tool.
3. **Network reality.** A single interconnection limit is a scalar proxy for a network. Confirming that the combined output clears under load and contingency requires an AC optimal-power-flow model of the local system.
4. **Statistical structure of the drivers.** The renewable and price series are neither independent nor identically distributed. Their long-memory and multifractal structure, measured through the Hurst exponent and wavelet spectra, governs how often large headroom coincides with high prices, which is what actually pays.
5. **Real data at scale.** Synthetic profiles give way to metered capacity factors, published interconnection-queue positions, pipeline geometry, and nodal congestion, screened across thousands of sites.
6. **Economics and risk.** Usable energy becomes revenue through a merchant price or spark spread, then net present value net of capital and connection cost, with the distribution of outcomes mattering as much as the mean.

Each rung is a step up in mathematical and engineering depth. The screen is the doorway; the value is in what it opens onto, and in building the rungs above it.

10 Status

The method is complete and specified in full. The screen, the worked example, and the validation are reproducible from the accompanying model. Deployment as a geographic screening application over interconnection and pipeline layers, and validation against metered site data, are the outstanding steps.

11 Related Full Stack Energy Analyses

Full Stack Energy has published work on several of the rungs above.

- S. Woolley, “The Dam in Your Battery”, Full Stack Energy Blog, June 2026. Stochastic dual dynamic programming for storage arbitrage.
- S. Woolley, “Wavelet-Based Dispatch for Mixed Energy Fleets”, Full Stack Energy Blog, April 2026. Fleet-scale dispatch by wavelet decomposition.
- S. Woolley, “A Dam Builder’s Yardstick in the Lone Star State”, Full Stack Energy Blog, April 2026. The Hurst exponent applied to ERCOT prices.
- S. Woolley, “Dispatch Is a Design Choice”, Full Stack Energy Blog, February 2026. A multifractal battery-arbitrage method.

References

[1] Federal Energy Regulatory Commission, *Order No. 845: Reform of Generator Interconnection Procedures and Agreements*, Docket No. RM17-8-000, 2018. Establishes Surplus Interconnection Service.

[2] Federal Energy Regulatory Commission, *Order No. 2023: Improvements to Generator Interconnection Procedures and Agreements*, Docket No. RM22-14-000, 2023. Expands access to Surplus Interconnection Service and co-location behind a single point of interconnection.