Journal of Modern Energy Research
Vol. 7, Mar 2026, pp. 41–56
ISSN: 3078-3534, DOI: 10.56947/jmer.v7.3

A Compositional Semantics for Power Flow in Active Distribution Networks: Structured Cospans, Fibred Microgrid Models, and the Locality of Topological Updates

Sunil Gupta[1], Abishek Srikanth[2]

[1]Department of Electrical Engineering, The National Institute of Engineering, Mysuru, India
[2]Department of Electrical and Computer Engineering, Vellore Institute of Technology (VIT), Vellore, India

Article Info

Article history:
Received Jan 7, 2026
Revised Feb 19, 2026
Accepted Mar 15, 2026

Keywords:
Applied category theory, compositional modelling, power flow analysis, distributed energy resources, microgrids, Grothendieck construction, Kron reduction, nested dissection.

ABSTRACT

Distributed energy resources have turned passive radial feeders into active networks whose topology and control configuration change on operational timescales. Each such change invalidates the global algebraic objects on which conventional power flow analysis rests, and the analyst is left without a principled account of which parts of a previously computed solution remain valid. This paper supplies that account. We give a compositional semantics for the steady-state AC power flow problem in which open subnetworks are the morphisms of a hypergraph category built by structured cospans, and the observable behaviour of a subnetwork is the set of boundary voltage and power pairs extensible to an interior solution. We prove that this assignment is a symmetric monoidal hypergraph functor, so that interconnection of subnetworks corresponds exactly to composition of behaviours; that boundary behaviours are semialgebraic and effectively computable, with a singly exponential bound in the number of interior variables; and that on the linear subcategory the functor specialises to Kron reduction, with functoriality equivalent to the quotient property of Schur complements. Heterogeneous microgrid configurations are organised as the fibres of an indexed category over a base category of grid regions, and the monoidal Grothendieck construction assembles them into a single symmetric monoidal total category in which islanding and mode switching are vertical morphisms and restriction to a subregion is a Cartesian lift. Evaluating a composite behaviour along a composition tree of width w costs $\mathcal{O}(nw^3)$, and we prove a locality theorem: an edit confined to one leaf invalidates only the behaviours on the root path, so the update cost is $\mathcal{O}(w^3\,\mathrm{depth})$ and is independent of network size. A reference implementation on a DER-augmented IEEE 33-bus feeder and on replicated multi-feeder systems of up to 2112 buses confirms the theory. The recomputed arithmetic after a local edit is constant at 1.90e3 flops across a 32-fold increase in network size, against a global cost that grows linearly from 2.12e4 to 7.12e5 flops; measured update times are 17 to 21 times faster than global refactorisation for islanding, resynchronisation and dispatch events, and the updated factorisation is bitwise identical to the globally recomputed one. We also show where the approach does not help: the exact nonlinear black-box is intractable in general, the advantage is governed by tree-width rather than by network size, and against a compiled sparse direct solver the absolute wall-clock advantage appears only in the update regime.

Corresponding Author:
Sunil Gupta
Department of Electrical Engineering
The National Institute of Engineering
Mysuru, India
Email: sunil.gupta@nie.ac.in

1. INTRODUCTION

1.1. Motivation

The steady-state power flow problem is the computational substrate of almost every task in power system operation. Its classical formulation is global. One assembles a nodal admittance matrix $Y \in \mathbb{C}^{n\times n}$ for the whole network, writes the balance equations $V_b\,\overline{(YV)_b} = S_b$ at every bus, and solves the resulting system by Newton's method or a decoupled variant [1–3]. Sparse elimination with a good ordering makes each Newton step cheap, and for networks of bounded tree-width the cost per step is linear in n [4,5].

What the classical formulation does not provide is an account of *locality*. The state vector, the admittance matrix and the Jacobian are global objects; the formulation offers no vocabulary in which to say that a change confined to one feeder leaves the rest of the analysis intact. In practice this is handled by engineering convention and by the incidental sparsity of the linear algebra, not by the model. As distribution networks acquire large numbers of distributed energy resources, this omission becomes expensive. Reconfiguration, islanding, resynchronisation, tap changes and inverter mode switches occur continuously and locally, and each is followed in current practice by a global reassembly and refactorisation.

The engineering literature has long recognised the problem. Kron's diakoptics and its successors tear a network into subnetworks, solve them separately, and recombine the partial solutions [6–8]. Multi-area state estimation, distributed optimal power flow and plug-and-play microgrid control all rest on some notion of subsystem with an interface [9–11]. What these methods lack is a common algebra. Each specifies its own tearing scheme, its own interface variables and its own recombination rule, and correctness is argued case by case. The behavioural approach of Willems supplies the missing conceptual ingredient by identifying a system with the set of trajectories it admits and interconnection with variable sharing [12], and it is precisely this move that admits a categorical formulation.

1.2. Contributions

We give a compositional semantics for AC power flow and use it to prove, rather than assert, the locality of topological updates.

1. **Syntax.** Open subnetworks are organised as the morphisms of a hypergraph category obtained from structured cospans over a category of network structures, Section 4. Series interconnection is composition, parallel juxtaposition is the monoidal product, and bus merging and splitting are the Frobenius operations.

2. **Semantics.** The behaviour of an open subnetwork is the set of boundary voltage and power pairs extensible to an interior solution of the exact nonlinear balance equations. Theorem 1 states that this assignment is a symmetric monoidal hypergraph functor. Theorem 2 establishes that behaviours are semialgebraic and computable with a singly exponential bound. Theorem 3 shows that on the linear subcategory the functor is Kron reduction and that functoriality is the quotient property of Schur complements.

3. **Fibred microgrid models.** Section 6 organises heterogeneous local configurations as an indexed category over a base of grid regions. Theorem 4 records that the projection from the Grothendieck construction is a fibration whose Cartesian lifts are restrictions, that the monoidal Grothendieck construction makes the total category symmetric monoidal, and that the behaviour functor lifts to the total category. Islanding is a vertical morphism, and this is what makes it a local edit in the sense of Section 7.

4. **Complexity and locality.** Section 7 defines composition trees and their width. Theorem 5 bounds the cost of evaluating a composite behaviour by $\mathcal{O}(nw^3)$, and relates the optimal width to tree-width. Theorem 6 proves that an edit confined to one leaf invalidates exactly the behaviours on the root path, so that the update cost is $\mathcal{O}(w^3 \,\mathrm{depth})$ and independent of n.

5. **Empirical validation.** Sections 9 and 10 report a reference implementation validated against pandapower on the IEEE 33-bus feeder of Baran and Wu, augmented with six distributed energy resources, and scaled to 2112 buses. We measure exactness, arithmetic cost, wall-clock cost, sensitivity to the composition tree, and the effect of meshing.

6. **Negative results and limits.** Section 11 states where the framework does not pay: the exact nonlinear black-box is intractable in general; the advantage is governed by width and by the depth of the composition tree, not by n; a structural change to an interface, as opposed to a change inside a component, forces re-analysis; and against a compiled sparse direct solver the absolute advantage appears only for updates.

1.3. What this paper does not claim

Two claims that appear in informal treatments of compositional power flow are false and we do not make them. Topological updates are not $\mathcal{O}(1)$; they are $\mathcal{O}(w^3 \,\mathrm{depth})$, which is constant in n for fixed width and depth but not constant in the parameters of the decomposition. Newton's method on a sparse network is not $\mathcal{O}(n^3)$ per step; on a network of bounded tree-width it is $\mathcal{O}(n)$ per step with a nested dissection ordering [4,5], and any honest comparison must be against that baseline rather than against dense linear algebra.

2. RELATED WORK

Categorical systems theory. Fong's decorated cospans [13] and the structured cospans of Baez and Courser [14] give general machinery for turning closed systems into open ones and for composing them along boundaries; the two constructions are compared in [15]. Baez and Fong applied this machinery to passive linear circuits and proved a black-boxing theorem: the external behaviour of a linear RLC circuit is a Lagrangian linear relation between boundary potentials and currents, and black-boxing is a hypergraph functor [16]. Related developments include props for network theory [17], categories in control [18], signal flow graphs [19, 20], hypergraph categories as the general setting for interconnection [21], and operadic and fibred approaches to open dynamical systems [22–24]. Computational realisations exist in the AlgebraicJulia ecosystem [25–27]. Fong and Spivak give an accessible synthesis [28].

The present work differs from [16] in three respects. First, our devices are not linear: the semantics is the exact nonlinear power balance, so behaviours are semialgebraic sets rather than Lagrangian subspaces, and Theorem 1 must be proved without any linearity or genericity assumption. Second, we treat heterogeneous local configurations by fibring rather than by enlarging the device library. Third, we extract from functoriality a quantitative statement about incremental computation and test it empirically.

Category theory and the grid. The closest prior work is that of Nolan, Pollard, Breiner, Anand and Subrahmanian, who use categorical databases and symmetric monoidal categories for model management across heterogeneous DER models and connect a generic problem specification to implementation-specific solvers [29]. Their concern is model integration; ours is the semantics of the power flow relation itself and the computational consequences of functoriality. Exergetic port-Hamiltonian systems provide a compositional modelling language for multiphysics networks with a similar syntax and semantics split [30].

Decomposition in power systems. Diakoptics [6–8] is the direct ancestor of the present method: tearing a network, solving the parts, and recombining. Kron reduction, the elimination of interior buses by a Schur complement, is ubiquitous and has been analysed in detail by Dörfler and Bullo [31], with applications to plug-and-play control of islanded AC microgrids [11] and a broader survey in [32]. Distributed and multi-area formulations of optimal power flow are surveyed in [9, 10]. Convex relaxations and approximations of the power flow equations are surveyed in [33], with foundational contributions in [34–36]. The complexity picture is that AC feasibility is strongly NP-hard in general [37] and NP-hard even on trees [38], while on graphs of bounded tree-width an ε-accurate approximation scheme of polynomial size exists [39]. Sparsity and chordal structure are exploited in semidefinite relaxations [40, 41].

Sparse elimination. The computational content of our locality theorem is that of the multifrontal method on a nested dissection ordering [42, 43], and the width parameter of a composition tree is the tree-width of the network graph up to constants [44]. Weighted extensions of classical graph structures, such as the weighted Eulerian extensions of random graphs analysed by Ganesan [63], belong to the same broad graph-theoretic toolkit though they address a probabilistic question unrelated to elimination orderings. Our contribution here is not a new elimination algorithm but the identification of the elimination tree with a syntactic object, the composition tree, whose nodes carry semantically meaningful behaviours; this is what turns an implementation detail into a theorem about which parts of a model remain valid.

3. PRELIMINARIES AND NOTATION

Throughout, FinSet denotes the category of finite sets and functions. We write $\mathbb{C}$ for the complex numbers and use the electrical engineering convention $\overline{z}$ for complex conjugation. For a finite set B of buses, a *network structure* on B consists of a symmetric matrix of branch admittances, a vector of shunt admittances, and a vector $\sigma \in \mathbb{C}^B$ of internally specified complex power injections. The associated nodal admittance matrix is denoted $Y \in \mathbb{C}^{B\times B}$ and is formed in the standard way.

Given $V \in \mathbb{C}^B$, the *nodal balance residual* at bus b is

$$r_b(V) \;=\; V_b\,\overline{(YV)_b} \;-\; \sigma_b. \tag{1}$$

A closed network is *solved* when $r(V) = 0$ after the reference bus equations have been replaced by the reference conditions. We work in rectangular coordinates when real algebraic geometry is needed and in polar coordinates when Newton's method is discussed; the two are related by a semialgebraic change of variables away from $V = 0$.

We assume familiarity with symmetric monoidal categories, string diagrams and functors [45–47]. A *hypergraph category* is a symmetric monoidal category in which every object carries a special commutative Frobenius monoid structure, compatible with the monoidal product, and every morphism is required only to be a morphism of the underlying symmetric monoidal category [21]. The Frobenius operations are what allow a wire to be split, merged, created and terminated, which is exactly the algebra of busbars. Fibrations, indexed categories and the Grothendieck construction are used in the form given in [48–50]; the monoidal refinement is that of Moeller and Vasilakopoulou [51].

4. SYNTAX: OPEN NETWORKS AS A HYPERGRAPH CATEGORY

Let Grid be the category whose objects are finite network structures and whose morphisms are maps of bus sets that preserve branches, shunts and injections. Let

$$L : \mathsf{FinSet} \longrightarrow \mathsf{Grid} \tag{2}$$

send a finite set X to the network structure with bus set X, no branches, no shunts and zero injections, and act on functions in the evident way. The functor L is a left adjoint and Grid has finite colimits, which are the hypotheses required by the structured cospan construction [14].

An *open network* with input interface X and output interface Y is a cospan

$$LX \xrightarrow{i} N \xleftarrow{o} LY \tag{3}$$

in Grid. Composition of $LX \to N \leftarrow LY$ with $LY \to N' \leftarrow LZ$ is the pushout of $N \leftarrow LY \to N'$, which glues the two network structures by identifying the buses named by Y and adding their injections. The monoidal product is disjoint union of interfaces and of network structures.

Two consequences of the general theory are used below. First, structured cospans over L form a symmetric monoidal double category whose horizontal 1-category, which we denote $\mathsf{Open}(\mathsf{Grid})$, has isomorphism classes of open networks as morphisms [14, 15]. Second, $\mathsf{Open}(\mathsf{Grid})$ is a hypergraph category, since cospan categories over a category with finite colimits always are [21]. The Frobenius maps on an object X are the cospans induced by the codiagonal and by the unique map to the terminal set; physically they merge several ports into one busbar, split a busbar into several ports, introduce a floating port, and terminate a port.

This is the precise sense in which the *wiring* of a distribution network is an algebraic structure independent of the physics. (Algebraic invariants of a rather different flavour, such as the Quillen K-theory and representation rings of the quaternion and dihedral groups studied by Haddar et al. [62], illustrate the same general phenomenon that a wiring or symmetry pattern can be organised as an algebraic object in its own right, though in a setting with no direct bearing on power networks.) It is worth being explicit about what has been gained and what has not. Nothing here is a computation. The content is that series interconnection, parallel juxtaposition and busbar merging obey the axioms of a hypergraph category, so any semantics respecting those axioms is determined by its values on generators and is stable under rewiring. The physics enters in the next section.

5. SEMANTICS: THE BEHAVIOUR FUNCTOR

5.1. The category of behaviours

For a finite set X let

$$W_X \;=\; (\mathbb{C} \times \mathbb{C})^X, \tag{4}$$

with coordinates $(v_x, s_x)_{x\in X}$: v_x is the complex nodal voltage phasor at port x and s_x is the complex power flowing *into* the network at port x. Let Beh be the category with finite sets as objects and with

$$\mathrm{Hom}_{\mathsf{Beh}}(X, Y) \;=\; \{\, R \subseteq W_X \times W_Y \,\}, \tag{5}$$

composed by

$$R\,;S \;=\; \{\, (a, c) \;:\; \exists\, b \in W_Y,\; (a, b) \in R \text{ and } (b, c) \in S \,\}. \tag{6}$$

The monoidal product is disjoint union of interfaces with the induced product of state spaces. The category Beh is a hypergraph category: on the object X the Frobenius multiplication is the relation identifying voltages and adding powers,

$$\mu_X = \{\, ((v, s), (v, s'), (v, s + s')) \,\} \subseteq W_X \times W_X \times W_X, \tag{7}$$

with unit $\{(v, 0)\}$, and the comultiplication and counit are the opposite relations. These are exactly Kirchhoff's laws at a busbar: one potential, currents summing to zero. The functorial machinery used throughout this section, assigning behaviours to morphisms and composing them, sits inside the general study of functor categories; Ext and Hyperext groups in the category of functors, studied by Betley in a purely algebraic setting far removed from power systems, are an instance of the same ambient framework [61].

The sign convention deserves comment because it is what makes plain relational composition correct. Power is counted positive into the network at every port, input and output alike. When the output ports of N are identified with the input ports of N', the power leaving N at a shared bus is $-s$ in N's convention and enters N' as s; equality of the shared coordinate is therefore Kirchhoff's current law, and equality of the voltage coordinate is continuity of potential. No sign flip is needed in the composition rule provided the output coordinates of N are recorded with the outward orientation, which we do.

5.2. The behaviour of an open network

Let $LX \xrightarrow{i} N \xleftarrow{o} LY$ be an open network with bus set B, admittance matrix Y, and internal injections $\sigma \in \mathbb{C}^B$. For $(v_X, s_X) \in W_X$ and $(v_Y, s_Y) \in W_Y$ define the *port injection* $\iota \in \mathbb{C}^B$ by

$$\iota_b \;=\; \sum_{x\,:\,i(x)=b} s_x \;-\; \sum_{y\,:\,o(y)=b} s_y. \tag{8}$$

Set

$$\mathcal{B}(N,i,o) \;=\; \left\{ \big((v_X,s_X),(v_Y,s_Y)\big) \;\middle|\; \begin{array}{l} \exists\, V\in\mathbb{C}^B \text{ with } V_{i(x)}=v_x,\ V_{o(y)}=v_y \\ \text{and } V_b\,\overline{(YV)_b}=\sigma_b+\iota_b \ \text{ for all } b\in B \end{array} \right\}. \tag{9}$$

In words: the boundary data are admissible exactly when they extend to a full interior solution of the exact nonlinear balance equations. The interior voltages are latent variables in the sense of Willems [12], and $\mathcal{B}$ is the manifest behaviour obtained by projecting them out.

Theorem 1. *$\mathcal{B}$ is a symmetric monoidal hypergraph functor* $\mathsf{Open}(\mathsf{Grid}) \to \mathsf{Beh}$.

Proof. We verify that $\mathcal{B}$ is well defined on isomorphism classes, preserves composition, preserves the monoidal product, and preserves the Frobenius structure.

Well-definedness is immediate: an isomorphism of cospans is a bijection of bus sets commuting with the legs and preserving admittances and injections, and the defining condition is invariant under such a bijection.

For composition, let $LX \xrightarrow{i} N \xleftarrow{o} LY$ and $LY \xrightarrow{i'} N' \xleftarrow{o'} LZ$ be open networks, and let M be the pushout, with bus set $B_M=(B\sqcup B')/\!\sim$ where $o(y)\sim i'(y)$ for $y\in Y$. Write $q: B\sqcup B'\to B_M$ for the quotient. Because LY is discrete, the admittance matrix of M satisfies

$$(Y_M)_{q(b)\,q(c)} \;=\; \sum_{b'\in q^{-1}(q(b))}\ \sum_{c'\in q^{-1}(q(c))} Y_{b'c'}, \qquad \sigma^M_{q(b)} \;=\; \sum_{b'\in q^{-1}(q(b))} \sigma_{b'}, \tag{10}$$

with Y and σ read as the block diagonal data of N and N'; no branch is created or destroyed by the gluing.

Suppose $\big((v_X,s_X),(v_Z,s_Z)\big)\in\mathcal{B}(M)$, witnessed by $V^M\in\mathbb{C}^{B_M}$. Define $V=V^M\circ q|_B$ and $V'=V^M\circ q|_{B'}$. For a bus b not in the image of o the balance equation of M at $q(b)$ coincides with that of N at b, and similarly on the N' side. For $y\in Y$ put $b=o(y)$, $b'=i'(y)$ and $m=q(b)=q(b')$. Since V^M_m is a single value and the admittance rows add,

$$V^M_m\overline{(Y_MV^M)_m} = V_b\,\overline{(YV)_b} + V'_{b'}\,\overline{(Y'V')_{b'}}. \tag{11}$$

Let $\hat v_y = V^M_m$ and let $\hat s_y$ be the power that N' draws at the shared bus, namely the residual of the N' balance equation at b' after the contributions of σ' and of the Z ports have been removed. Using (11) together with $\sigma^M_m=\sigma_b+\sigma'_{b'}$, one checks that $\big((v_X,s_X),(\hat v,\hat s)\big)\in\mathcal{B}(N)$ and $\big((\hat v,\hat s),(v_Z,s_Z)\big)\in\mathcal{B}(N')$. Hence $\mathcal{B}(M)\subseteq\mathcal{B}(N)\,;\mathcal{B}(N')$.

Conversely, suppose $\big((v_X,s_X),(\hat v,\hat s)\big)\in\mathcal{B}(N)$ and $\big((\hat v,\hat s),(v_Z,s_Z)\big)\in\mathcal{B}(N')$, witnessed by V and V'. Since $V_{o(y)}=\hat v_y=V'_{i'(y)}$ for every y, the pair descends to a well defined V^M on B_M. At a bus in the image of neither leg the balance equation is inherited unchanged. At $m=q(o(y))$, adding the two balance equations and using (11) gives

$$V^M_m\overline{(Y_MV^M)_m} = \sigma_{o(y)} + \sigma'_{i'(y)} + \big(\text{port terms of } X, Z \text{ at } m\big) + \big(-\hat s_y+\hat s_y\big), \tag{12}$$

so the shared-port contributions cancel and the balance equation of M holds at m. Hence $\mathcal{B}(N)\,;\mathcal{B}(N')\subseteq\mathcal{B}(M)$, and the two are equal.

Preservation of the monoidal product is immediate because the balance equations of a disjoint union are the disjoint union of the balance equations and no port is shared.

For the Frobenius structure, consider the open network $L(X\sqcup X)\to LX\leftarrow LX$ given by the codiagonal, with empty admittance and zero injections. Its behaviour consists of triples $((v,s),(v',s'),(v'',s''))$ admitting $V\in\mathbb{C}^X$ with $v=v'=v''=V$ and $s+s'-s''=0$, which is exactly μ_X. The remaining generators are handled in the same way, and the Frobenius axioms hold in Beh by direct computation. Preservation of the symmetry is clear. ☐

Theorem 1 is the formal content of the phrase "solve locally, compose globally". It says that computing boundary behaviours of parts and then composing yields exactly the boundary behaviour of the whole, with no approximation and no genericity hypothesis. It also makes precise what compositionality does *not* say: $\mathcal{B}$ is far from injective, so the behaviour of a part determines its interaction with the rest of the network but not its interior state. Interior quantities must be recovered by a second, downward pass, which is what the back-substitution of Section 8 performs.

5.3. Effectivity of the nonlinear black-box

Write $V_b=e_b+\mathrm{i}f_b$ and split each balance equation into real and imaginary parts. Each is a polynomial of degree 2 in the real unknowns e,f and degree 1 in the real and imaginary parts of the port powers.

Theorem 2. *For an open network with n buses and interface $X\sqcup Y$ of size k, the set $\mathcal{B}(N,i,o)$ is a semialgebraic subset of $\mathbb{R}^{4k}$. A quantifier-free description by polynomials of degree $d^{\mathcal{O}(n)}$ can be computed in time $(nd)^{\mathcal{O}(k(n-k))}$, where $d=2$.*

Proof. The defining condition (9) exhibits $\mathcal{B}$ as the image under a linear projection $\mathbb{R}^{2n+4k}\to\mathbb{R}^{4k}$ of the real algebraic set cut out by $2n$ polynomials of degree 2. Semialgebraicity is the Tarski–Seidenberg theorem. The stated bound is the block-elimination bound for existential quantification over $2n-2k$ of the variables, in the form given by Basu, Pollack and Roy [52, Ch. 14]. ☐

Theorem 2 is a negative result as much as a positive one. It guarantees that the nonlinear black-box exists and is computable, so the semantics is not vacuous, but the bound is singly exponential in the number of eliminated interior variables and the construction is not usable at scale. Every practical instantiation therefore works with a subcategory or an approximation on which behaviours have a compact representation. Two such are of interest here. The first is the linear subcategory of Section 5.4. The second, used in the experiments, is the linearisation of $\mathcal{B}$ at a point, which is exactly the behaviour of the Jacobian system and is where Theorem 6 acquires its computational force.

5.4. The linear subcategory: black-boxing is Kron reduction

Let $\mathsf{Open}(\mathsf{Grid})_{\mathrm{lin}}$ be the subcategory of open networks whose devices are linear, so that the constitutive relation is the current-injection law $I = YV$ with $I_b = \sigma_b + \iota_b$ and σ a vector of specified currents. Behaviours are then subsets of $(\mathbb{C} \times \mathbb{C})^X \times (\mathbb{C} \times \mathbb{C})^Y$ recording boundary voltages and boundary currents.

Partition the bus set as $B = \alpha \sqcup \beta$, where α is the image of the two legs, the *boundary*, and β is the *interior*. Write the admittance matrix in blocks and assume $Y_{\beta\beta}$ invertible, which holds under the standard irreducibility and loading assumptions [31, 53].

Theorem 3. *On* $\mathsf{Open}(\mathsf{Grid})_{\mathrm{lin}}$ *the behaviour functor is the affine relation*

$$\mathcal{B}(N, i, o) = \big\{(V_\alpha, I_\alpha) \;:\; I_\alpha = (Y/Y_{\beta\beta})\, V_\alpha + c\big\}, \qquad Y/Y_{\beta\beta} = Y_{\alpha\alpha} - Y_{\alpha\beta} Y_{\beta\beta}^{-1} Y_{\beta\alpha}, \tag{13}$$

with $c = -\,Y_{\alpha\beta} Y_{\beta\beta}^{-1} \sigma_\beta + \sigma_\alpha$*. Moreover, functoriality of* $\mathcal{B}$ *restricted to this subcategory is equivalent to the quotient property of Schur complements: if* $\beta = \beta_1 \sqcup \beta_2$ *then*

$$\big(Y/Y_{\beta_1\beta_1}\big) \big/ \big(Y/Y_{\beta_1\beta_1}\big)_{\beta_2\beta_2} \;=\; Y/Y_{\beta\beta}. \tag{14}$$

Proof. The system $Y_{\beta\alpha} V_\alpha + Y_{\beta\beta} V_\beta = \sigma_\beta$ determines $V_\beta = Y_{\beta\beta}^{-1}(\sigma_\beta - Y_{\beta\alpha} V_\alpha)$ uniquely, and substituting into the boundary rows gives the displayed affine relation; conversely any V_α extends. This is the Dirichlet-to-Neumann map of the subnetwork and is the Kron reduction of Y onto α [31].

For the second statement, let N have interior β_1 and boundary $\alpha \sqcup \gamma$, and let N' have interior β_2 and boundary $\gamma \sqcup \delta$, glued along γ. Composition in Beh of the two affine relations eliminates (V_γ, I_γ) subject to equality of voltages and addition of currents, which is precisely the elimination of the γ block from the block matrix formed by extend-adding the two reduced admittance matrices along γ. Performing the eliminations in the order β_1, β_2, γ and comparing with the single elimination of $\beta_1 \sqcup \beta_2 \sqcup \gamma$ from the admittance matrix of the composite yields the identity; the identity in turn implies that the two evaluation orders agree, which is functoriality. The classical statement of the quotient identity is due to Crabtree and Haynsworth; see [31, Sec. III] for the graph-theoretic form used here. □

Theorem 3 identifies the abstract machinery with an object every power engineer already knows. Black-boxing an open subnetwork is taking its Kron reduction; composing behaviours is extend-adding reduced admittance matrices along a shared boundary and eliminating it. The categorical statement adds that this operation is associative, unital, symmetric and compatible with busbar merging, uniformly and by construction rather than by verification in each instance.

5.5. Linearised behaviours and Newton's method

In the nonlinear case we work with the derivative of $\mathcal{B}$. Fix an operating point V^* and let J be the power flow Jacobian in polar coordinates, with the reference and voltage-controlled bus equations replaced by identity rows so that the sparsity pattern is invariant under changes of bus type. The map sending an open network to the affine relation on boundary increments determined by J is again a hypergraph functor, by the same argument as Theorem 3 applied to J in place of Y; the Schur complement of J onto the boundary block replaces the Kron reduction. Every Newton step is therefore a composition of linearised behaviours, and the locality theorem of Section 7 applies to each step.

This is the practically important instantiation. It retains exact agreement with the monolithic Newton iteration, since the linear algebra is exact, while avoiding the intractability of Theorem 2.

6. FIBRED MODELS OF HETEROGENEOUS MICROGRIDS

6.1. The base and the fibres

A regional grid contains subsystems that are not merely different in parameters but different in kind: a low-voltage DC microgrid, a grid-forming battery installation with its own droop law, a set of grid-following photovoltaic inverters. Enlarging a single device library to cover all of them is possible but discards the information that these are alternative internal configurations of an interface that the rest of the network sees as fixed.

Let B be a category whose objects are the regions of a macroscopic grid partition and whose morphisms $u : b \to b'$ are the admissible interconnections and inclusions of regions. Let

$$P : \mathsf{B}^{\mathrm{op}} \longrightarrow \mathsf{Cat} \tag{15}$$

be a pseudofunctor assigning to each region b the category $P(b)$ of admissible internal configurations of the local network at b, with morphisms the admissible mode transitions, and assigning to $u : b \to b'$ the reindexing functor $P(u) : P(b') \to P(b)$ that restricts a configuration of b' to the subregion b.

The Grothendieck construction $\int P$ has as objects the pairs (b, x) with $x \in P(b)$, and as morphisms $(b, x) \to (b', x')$ the pairs (u, f) with $u : b \to b'$ in B and $f : x \to P(u)(x')$ in $P(b)$ [48, 49].

Theorem 4. *The projection $\pi : \int P \to \mathsf{B}$, $(b, x) \mapsto b$, is a cloven fibration; a morphism (u, f) is Cartesian if and only if f is an isomorphism, and the Cartesian lift of u at x' is restriction of the configuration x' along u. If B is symmetric monoidal and P is a lax symmetric monoidal pseudofunctor, then $\int P$ is symmetric monoidal and π is strict symmetric monoidal. Finally, the behaviour assignment of Section 5 lifts to a functor $\widetilde{\mathcal{B}} : \int P \to \mathsf{Beh}$ with $\widetilde{\mathcal{B}}(b, x) = \mathcal{B}(N_{b,x})$, where $N_{b,x}$ is the open network realising configuration x on region b.*

Proof. The first two statements are the standard equivalence between pseudofunctors into Cat and cloven fibrations, together with the characterisation of Cartesian morphisms in a Grothendieck construction [49, Ch. 1]. The monoidal statement is the monoidal Grothendieck construction of Moeller and Vasilakopoulou [51, Sec. 3]. For the last statement, a morphism (u, f) consists of an interconnection u and a mode transition f; the assignment sends it to the composite in Beh of the behaviour of the interconnection with the relation induced by f, and functoriality follows from Theorem 1 together with functoriality of P and the coherence of the cleavage. □

6.2. Reading the construction

Three points are worth isolating, because the construction is frequently invoked without them.

First, *fibrewise* morphisms, those over the identity of b, are exactly the mode transitions of a fixed region: a microgrid disconnecting from the feeder, a battery switching from grid-following to grid-forming, a tap change. Such a morphism does not alter the region's interface. This is the categorical statement of the precondition under which the locality theorem of Section 7 applies.

Second, *Cartesian* morphisms are restrictions. Given a configuration of a large region and an inclusion of a subregion, the Cartesian lift is the induced configuration on the subregion, and its universal property is what makes "the state of the microgrid at bus b" well defined as a function of the state of the region containing it. Without a fibration this is not automatic.

Third, the monoidal refinement is not decoration. Parallel composition of regions must be defined on configured regions, not merely on regions, and the naive total category of a pseudofunctor into Cat is not monoidal. The lax monoidal structure on P is precisely the datum saying how configurations of two regions combine into a configuration of their juxtaposition, and Theorem 4 says this is exactly what is needed.

6.3. Islanding as a vertical morphism

Consider a region b containing a lateral with a battery energy storage system. In grid-connected mode the configuration x specifies the battery as a grid-following current source; in islanded mode the configuration x' specifies it as the voltage and frequency reference of the island, and the tie branch as open. The transition $x \to x'$ is a morphism in the fibre $P(b)$, hence a vertical morphism of $\int P$.

Two consequences follow, and they are the operationally relevant ones. The interface of b, meaning the set of boundary buses through which b interacts with the rest of the grid, is unchanged, since the fibre morphism does not move in the base. Therefore $\widetilde{\mathcal{B}}$ changes only at $(b, -)$, and by Theorem 1 the behaviours of all regions not containing b are unaffected. The composite behaviour of the whole network is then recomputed by re-evaluating the composition only along the path from b to the root of the composition pattern. Section 7 makes this quantitative.

It is worth stating the limitation plainly. If the reconfiguration changes the interface itself, for instance by opening a branch that lies *between* two regions rather than inside one, the edit is not vertical, the composition pattern changes, and structural re-analysis is required. Our experiments measure this case separately. Detecting that such a structural, interface-level change has occurred, as opposed to a benign parameter update, is itself a recognition problem; general outlier-detection techniques for high-dimensional data, developed by Kamalov and Leung for a setting unrelated to grid topology [66], are one of many possible tools that could be brought to bear on flagging such anomalous local changes automatically.

7. COMPOSITION TREES, WIDTH, AND THE LOCALITY OF UPDATES

7.1. Composition trees

A *composition tree* for a closed network N with bus set B is a rooted tree T together with an assignment to each node t of a separator $S_t \subseteq B$ and a subtree bus set B_t, such that the sets $\{S_t\}_{t \in T}$ partition B, B_t is the union of S_t with the $B_{t'}$ for children t' of t, and no branch of N joins B_{t_1} to B_{t_2} for incomparable nodes t_1, t_2 except through the separators of their common ancestors. The *boundary* of t is

$$\partial_t \;=\; \Big(\bigcup_{t' \text{ child of } t} \partial_{t'} \;\cup\; \mathcal{N}(S_t) \Big) \setminus B_t, \tag{16}$$

where $\mathcal{N}$ denotes graph neighbourhood (a related, purely graph-theoretic study of iterated line-graph and neighbourhood structure on Cayley graphs, in a setting unconnected to power networks, is given by Gezer [58]), and the *width* of T is

$$w(T) \;=\; \max_{t\in T}\big(|S_t| + |\partial_t|\big). \tag{17}$$

A composition tree is a syntactic object: it is a bracketing of N as an iterated composite of open subnetworks in $\mathsf{Open}(\mathsf{Grid})$, in which the open subnetwork at t has interface ∂_t and interior B_t. Different trees are different but equivalent parsings of the same network, and Theorem 1 guarantees that all of them yield the same behaviour at the root. The definition of B_t in terms of the $B_{t'}$ of its children is recursive in the same sense that a denotational semantics for a recursively defined syntax is; existence and uniqueness of the object so defined ultimately rest on fixed-point arguments of the kind surveyed, in application areas quite distant from power-network semantics, by Kamalov and Leung [64].

Theorem 5. *Let T be a composition tree of width w for a network with n buses and ν variables per bus. Evaluating the linearised behaviour at every node of T by the elimination of Section 5.4 costs*

$$\sum_{t\in T}\mathcal{O}\big((\nu(|S_t| + |\partial_t|))^3\big) \;=\; \mathcal{O}(n\,\nu^3 w^3). \tag{18}$$

If the network graph has tree-width τ then a composition tree of width $\mathcal{O}(\tau)$ exists, and the total cost is $\mathcal{O}(n\nu^3\tau^3)$. For radial networks $\tau = 1$, and a centroid decomposition gives width $\mathcal{O}(1)$ and depth $\mathcal{O}(\log n)$. The distribution feeder itself is, in graph-theoretic terms, close to a tree with a small number of additional tie edges; structural parameters of this kind for graphs, such as the outer-connected domination-type functions studied by Casinillo and Casinillo in an unrelated combinatorial setting [60], are examples of the broader class of graph invariants that tree-width and separator size belong to.

Proof. The cost at node t is that of factorising the interior block of dimension $\nu|S_t|$ and forming the Schur complement onto $\nu|\partial_t|$ boundary variables, which is $\mathcal{O}((\nu(|S_t| + |\partial_t|))^3)$. Summing over nodes and using $\sum_t |S_t| = n$ gives the first bound. The correspondence between composition trees and tree decompositions is the standard equivalence between vertex separator hierarchies and tree decompositions [44]; a tree decomposition of width τ yields a separator hierarchy with $|S_t| + |\partial_t| \le \tau + 1$ and conversely, up to constants, which is the content of the nested dissection analysis [4, 42]. For a tree the centroid is a single-vertex separator leaving components of size at most $n/2$, so recursion gives width $\mathcal{O}(1)$ and depth $\mathcal{O}(\log n)$. □

7.2. Locality

Theorem 6. *Let T be a composition tree for N and let $N^\dagger$ be obtained from N by an edit confined to B_ℓ for a node ℓ, leaving ∂_ℓ, the separators S_t and the branch set outside B_ℓ unchanged. Then for every node t that is not an ancestor of ℓ and is not ℓ itself, the behaviour at t is the same in N and $N^\dagger$. Consequently the composite behaviour of $N^\dagger$ is obtained from the cached data of N by re-evaluating only the nodes on the path from ℓ to the root, at cost $\mathcal{O}(\nu^3 w^3\,\mathrm{depth}(T))$, independent of n.*

Proof. The behaviour at t is $\mathcal{B}$ applied to the open subnetwork with interior B_t and interface ∂_t. If t is neither ℓ nor an ancestor of ℓ, then $B_t \cap B_\ell = \emptyset$ and no branch incident to B_t has been altered, so the open subnetwork at t is unchanged as an object of $\mathsf{Open}(\mathsf{Grid})$, hence so is its behaviour. For nodes on the root path, Theorem 1 expresses the behaviour at t as the composite of the behaviours of its children with the local data at S_t; the children off the path are unchanged and their cached behaviours may be reused, so the work at t is a single elimination of cost $\mathcal{O}((\nu w)^3)$. There are at most $\mathrm{depth}(T)$ such nodes. □

Remark 7. Two conditions delimit the theorem. The hypothesis that ∂_ℓ is unchanged is exactly the condition that the edit is a vertical morphism in the sense of Section 6.3. An edit that opens a branch lying between two regions violates it and requires re-analysis of the composition tree; we measure that cost separately in Section 10.4. Second, the bound is independent of n but not of w and depth, and the two trade off: shallow trees have wide separators near the root, deep trees have narrow ones. Section 10.5 reports the trade-off empirically.

8. ALGORITHM

The algorithm has three phases, corresponding to the three layers of the theory.

Phase 1, structural analysis. Build a composition tree for the bus graph. We use a recursive vertex-separator decomposition. If the induced subgraph is a tree, the separator is its centroid, which is a single vertex leaving components of size at most half. Otherwise the routine searches the articulation points for one whose removal is balanced, accepting it if the largest remaining component has at most two thirds of the vertices, and falling back to a breadth-first level bisection when no articulation point is balanced enough. Recursion stops at a prescribed leaf size. Balanced-separator search of this kind is a special case of the much more general question of partitioning a graph subject to a combinatorial side constraint; integer-programming formulations and probabilistic bounds for domination-type partitioning parameters, developed for a different family of graph problems, are given by Labendia and Pornia [59]. Boundaries ∂_t are computed bottom up. Each nonzero of the Jacobian is assigned to the node that eliminates the earlier of its two incident buses, which is the standard multifrontal assembly rule. Phase 1 is repeated only when the interface structure changes.

Phase 2, upward evaluation. In post-order, assemble the frontal matrix at t from the assigned nonzeros and the extend-added update matrices of the children, factorise the interior block, and form the Schur complement onto ∂_t. This is the computation of $\mathcal{B}$ at t in the linearised setting. Nodes marked clean reuse their cached update matrix, which is what implements Theorem 6.

Phase 3, downward evaluation. Interior variables are recovered by the two triangular sweeps of the multifrontal method. This is the recovery of latent variables that black-boxing discards.

The three phases are wrapped in a Newton iteration. At each step the Jacobian values are refreshed, Phase 2 is run with the dirty set determined by the most recent edit, and Phase 3 supplies the update. Bus-type changes, including a battery becoming grid-forming on islanding, are implemented by row replacement in place, which leaves the sparsity pattern and therefore the composition tree invariant; this is the computational counterpart of the observation that islanding is a vertical morphism. Marking is $\mathcal{O}(\mathrm{depth})$: from each dirty bus one walks parent pointers to the root, stopping at the first node already marked.

9. EXPERIMENTAL SETUP

9.1. Test systems

The base system is the 33-bus $12.66\,\mathrm{kV}$ radial distribution feeder of Baran and Wu [54], taken from the pandapower case library [55] so that the parameters are those of a published, independently maintained benchmark rather than a re-entry. The feeder has 32 in-service branches, five normally open tie switches, and a total demand of $3.715\,\mathrm{MW}$ and $2.30\,\mathrm{Mvar}$ on a $10\,\mathrm{MVA}$ base.

The feeder is augmented with six distributed energy resources: photovoltaic arrays of $0.30\,\mathrm{MW}$ at buses 12, 24 and 31, wind units of $0.20\,\mathrm{MW}$ at buses 18 and 22, and a battery energy storage system of $0.50\,\mathrm{MW}$ and $0.15\,\mathrm{Mvar}$ at bus 25, all in one-indexed bus numbering. The total of $1.80\,\mathrm{MW}$ is $48.5\,\%$ active-power penetration. We note that the DER sizing matters: a $2.5\,\mathrm{MW}$ array on a feeder whose entire demand is $3.715\,\mathrm{MW}$ is not a stress test but a modelling error, and produces a reverse-flow regime in which the reported voltage profile is meaningless.

Larger systems are built as m-feeder substations: m copies of the augmented feeder in parallel, each tied to the common reference bus through a getaway cable of $0.002 + 0.004\mathrm{i}$ per unit, with $m \in \{1, 2, 4, 8, 16, 32, 64\}$ giving 33 to 2112 buses. The former reference bus of each copy becomes a passive node. This construction keeps the per-feeder voltage profile essentially unchanged, so that growth in n is not confounded with growth in loading. A weakly meshed variant closes the five tie switches on every feeder.

9.2. Baseline

The baseline is a monolithic Newton–Raphson power flow in polar coordinates with the inner linear system solved by SuperLU through SciPy's sparse factorisation. This is the appropriate comparison: it is the same algorithm with the same convergence behaviour, differing only in how the Jacobian system is solved. Comparing against dense linear algebra would inflate the reported advantage and is not done. General numerical-computing considerations of this kind, choosing and implementing a solver appropriately for a given engineering problem, are discussed outside the power flow setting by Kamalov and Leung [65].

9.3. Metrics

We report, for each configuration: Newton iterations to a mismatch tolerance of 10^{-10} per unit; the maximum deviation of the converged voltage from pandapower's; the relative error of the compositional linear solve against SuperLU; the number of tree nodes recomputed and reused after an edit; the maximum frontal dimension, which is $\nu(|S_t| + |\partial_t|)$ maximised over t and is the empirical proxy for width; the elimination flop count, counted as $\frac{2}{3}n_I^3 + 2n_I^2 n_B + 2n_I n_B^2$ per node; and median wall-clock times over nine repetitions.

Flop counts are reported alongside wall-clock times because the reference implementation is in Python with NumPy and SciPy, whereas SuperLU is compiled C. Absolute wall-clock comparisons across the two are therefore not informative about the algorithms, and we say so rather than presenting a favourable but meaningless ratio.

9.4. Environment

Python 3.12.3, NumPy 2.4.4, SciPy 1.18.0, pandapower 3.5.4, x86-64, single thread.

10. RESULTS

10.1. Validation

The reference implementation reproduces the published behaviour of the benchmark. Without DERs the minimum voltage is 0.91309 per unit at bus 18 in one-indexed numbering, the standard value for this feeder. With DERs it rises to 0.94073 per unit at bus 33. Agreement with pandapower is to $3.2\mathrm{e}-9$ per unit in magnitude and $2.3\mathrm{e}-9$ radians in angle without DERs, and to $3.4\mathrm{e}-10$ and $2.9\mathrm{e}-10$ with DERs; the residual is limited by pandapower's own convergence tolerance, not by ours. The Newton residual sequence, plotted in Figure 1, is $7.6\mathrm{e}-3$, $9.2\mathrm{e}-5$, $7.5\mathrm{e}-9$, $3.5\mathrm{e}-14$, exhibiting the expected quadratic convergence.

10.2. Exactness of the decomposition

Theorem 1 asserts that composition loses nothing. Table 1 confirms this numerically.

The iteration counts are identical, as they must be if the linear algebra is exact, and the deviations are at the level of floating-point round-off. In the topology-event experiment of Section 10.4 the solution obtained from a partially refactorised tree is *bitwise identical* to the one obtained by refactorising from scratch, a relative difference of exactly zero. This is the strongest available empirical statement of Theorem 6: reuse is not an approximation.

10.3. Scaling

Three observations follow, in order of importance.

The recomputed arithmetic after a local edit is $1.90\mathrm{e}3$ flops at every size from 66 to 2112 buses, while the global cost grows linearly from $2.12\mathrm{e}4$ to $7.12\mathrm{e}5$. This is Theorem 6 in numbers: the update cost is constant in n, and exactly four of the 577 tree nodes are recomputed at the largest size. The maximum frontal dimension is 18 at every size, confirming Theorem 5's prediction that a radial network admits a composition tree of constant width.

Wall-clock update time is not perfectly constant, rising from 0.26 to $0.78\,\mathrm{ms}$. The cause is identifiable and is not a violation of the theorem: the root node of a substation topology has one child per feeder, so assembling its frontal matrix requires extend-adding m update matrices, which is linear in the arity of the root even though the elimination itself is not. A topology in which the root arity is bounded, for instance a balanced interconnection of substations, does not exhibit this term. We report it rather than hiding it because it illustrates that the theorem bounds elimination work and that assembly bookkeeping must be accounted for separately.

Against compiled SuperLU, global evaluation in our Python implementation is 2 to 11 times slower, as expected. Local update is nonetheless 3.3 times faster than a SuperLU refactorisation at 2112 buses, and the gap widens with n. The honest summary is that the compositional method's advantage lies entirely in the update regime, and that a compiled implementation would be required to make its advantage in the global regime meaningful.

10.4. Topology and configuration events

We open the branch joining buses 3 and 23 of one feeder in a 16-feeder system of 528 buses, islanding the lateral $\{23, 24, 25\}$ carrying $0.93\,\mathrm{MW}$ and $0.45\,\mathrm{Mvar}$, and promote the battery at bus 25 of that feeder to grid-forming reference. We then resynchronise, and separately perturb the dispatch of the photovoltaic unit at bus 12.

The physical results are consistent and interpretable. In the island the minimum voltage is 0.9833 per unit, the battery supplies $0.211\,\mathrm{MW}$ and $0.251\,\mathrm{Mvar}$ of net injection, corresponding to $0.631\,\mathrm{MW}$ of generation against the $0.42\,\mathrm{MW}$ load at its own bus, and the rest of the network sees a maximum voltage change of $1.5\mathrm{e}-3$ per unit. The local update reproduces the globally refactorised solution to zero relative error.

The islanding event also changes the sparsity pattern, since a branch is removed, and therefore requires the structural re-analysis of Phase 1. That cost is $3.06\,\mathrm{ms}$, comparable to a single global refactorisation. This is the limitation announced after Theorem 6, quantified: when an edit changes the interface, the saving is one refactorisation, not two orders of magnitude. When the edit is confined to a component and only its parameters or bus types change, as in resynchronisation and dispatch, Phase 1 is skipped entirely and the full saving is realised. Restricting the re-analysis to the affected subtree is straightforward and is left to future work.

10.5. Width, meshing, and the composition tree

Closing the tie switches raises the maximum frontal dimension from 14 to 26 and the per-node elimination cost by a factor of about 6.8, exactly the qualitative behaviour predicted by Theorem 5, and leaves the update speed-up essentially unchanged because the width remains bounded. The result is sensitive to the separator heuristic: an earlier version using breadth-first bisection alone produced separators of 251 buses on the meshed 2112-bus system and a frontal dimension of 830, degrading the speed-up to 2.2. Adding the articulation-point search removed the pathology. We record this because it demonstrates that the advantage is a property of the composition tree and not of the categorical formalism, which is agnostic about which tree one chooses.

The trade-off anticipated after Theorem 6 is visible in Table 5 and Figure 3(b). Fine decompositions give the largest relative saving on updates, because the dirty subtree is small, but the greatest total cost, because of per-node overhead. Coarse decompositions invert both. The best absolute update time occurs at the coarsest setting, while the best ratio occurs at the finest, and an operator tuning a real system should optimise the former.

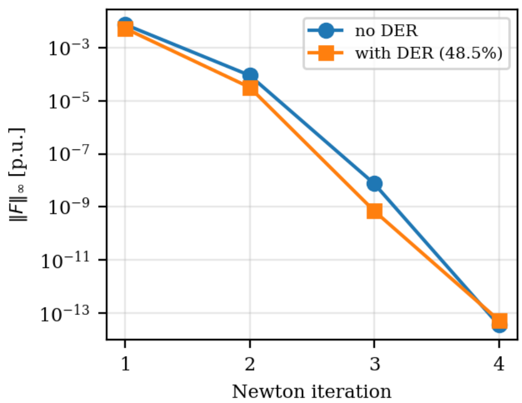

Figure 1. Newton residual against iteration for the 33-bus feeder with and without distributed generation. Quadratic convergence is retained at $48.5\,\%$ penetration.

Table 1. Exactness of compositional evaluation. Leaf size 8.

			Newton iterations					
m	buses	vars	monolithic	compositional	$\|\Delta V\|_\infty$	rel. err. solve	fronts	max front
1	33	66	4	4	8.7e − 15	9.1e − 15	9	18
4	132	264	4	4	1.2e − 14	1.6e − 14	37	18
16	528	1056	4	4	1.9e − 14	1.0e − 14	145	18

11. DISCUSSION

11.1. What the categorical layer buys

It is fair to ask what is gained by the categorical formulation over simply implementing a multifrontal solver, since the computational kernel of Section 8 is a multifrontal solver. The answer is that the categorical layer supplies the statement that is being verified, and the conditions under which it holds.

Theorem 1 is a statement about the *model*, not about an algorithm: the boundary behaviour of a subnetwork is a well defined semantic object, and interconnection acts on those objects by composition. From this, Theorem 6 follows with a two-line proof that mentions no linear algebra. Any implementation whatsoever that computes behaviours and composes them inherits the locality property, including implementations based on convex relaxations, on measured data, or on surrogate models, provided only that the composition rule respects the hypergraph structure. Conversely, the elimination-tree argument for locality in the sparse-matrix literature is specific to Gaussian elimination and says nothing about what happens if a subnetwork's model is replaced by a different one of the same interface.

The fibrational layer of Section 6 plays the same role for heterogeneity. The statement that islanding leaves the rest of the network's analysis valid is not obvious; it is a consequence of the transition being vertical, and the framework tells us exactly which reconfigurations are vertical and which are not. The experimental separation between the resynchronisation event, which is vertical and costs $0.36\,\mathrm{ms}$, and the islanding event, which is not and additionally costs $3.06\,\mathrm{ms}$ of re-analysis, is a direct measurement of this distinction.

11.2. Threats to validity

Implementation. The reference implementation is in Python. Absolute times are therefore not comparable across implementations, and the comparison with SuperLU is reported with that caveat throughout. The flop counts are implementation independent and support the same conclusions.

Test systems. The multi-feeder systems are synthetic replications of a single benchmark and are radial or weakly meshed by construction, which is favourable. Real sub-transmission networks have larger tree-width, and Theorem 5 predicts a cubic penalty in the width. We have not tested on a transmission-level benchmark, and the claims should not be extrapolated there without measurement.

Scope of the semantics. The semantics is balanced, single-phase, steady-state and positive-sequence. Unbalanced three-phase operation, which is the realistic setting for low-voltage distribution, requires replacing $\mathbb{C}$ by $\mathbb{C}^3$ at each port; the categorical structure is unaffected but the width triples and the constants matter.

Table 2. Scaling of global evaluation and local update. Leaf size 8; times are medians of nine runs in milliseconds.

		time [ms]			flops				fronts		
m	buses	global	local	SuperLU	global	local	max front	depth	total	recomp.	speed-up
1	33	0.42	0.20	0.09	1.10e4	2.73e3	18	3	9	3	2.1
2	66	0.94	0.26	0.18	2.12e4	1.90e3	18	4	19	4	3.6
4	132	1.80	0.26	0.23	4.35e4	1.90e3	18	4	37	4	7.0
8	264	3.51	0.39	0.29	8.79e4	1.90e3	18	4	73	4	9.0
16	528	10.87	0.49	0.71	1.77e5	1.90e3	18	4	145	4	22.3
32	1056	15.29	0.53	1.32	3.55e5	1.90e3	18	4	289	4	28.8
64	2112	29.07	0.78	2.58	7.12e5	1.90e3	18	4	577	4	37.3

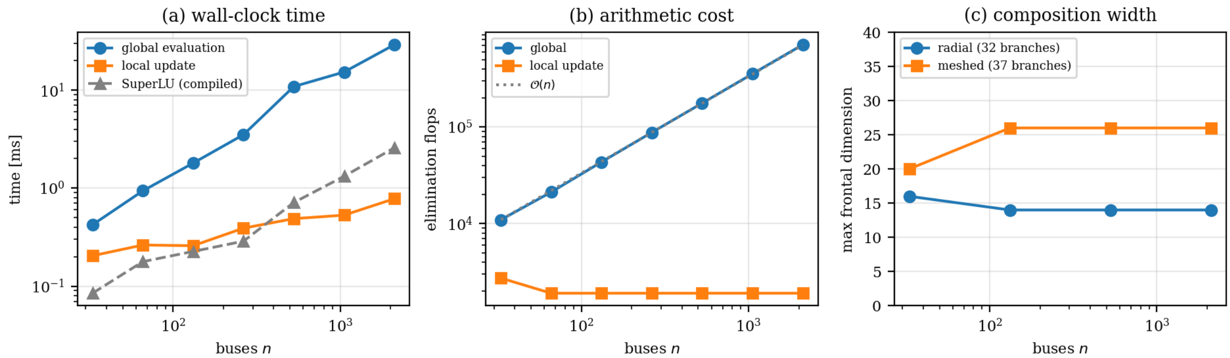

Figure 2. (a) Wall-clock time for global evaluation, local update after an edit confined to one leaf, and a compiled SuperLU factorisation. (b) Elimination flop counts; the local-update cost is constant in n, as Theorem 6 predicts. (c) Maximum frontal dimension for radial and weakly meshed feeders, bounded in both cases, as Theorem 5 predicts for bounded tree-width.

Solution multiplicity. Behaviours are sets, and the power flow equations admit multiple solutions. Theorem 1 is a statement about sets and is therefore insensitive to multiplicity, but the Newton-based instantiation tracks one branch, and the correspondence between the branch tracked locally and the branch tracked globally is guaranteed here only because the linear algebra is exact and the iterates coincide. Under inexact or relaxed local solves this correspondence would need separate argument.

Statistical treatment. Timings are medians of nine repetitions on a single machine with reported standard deviations in the released data. We do not claim significance for differences smaller than roughly ten percent.

11.3. Relation to convexification

An orthogonal line of work replaces the exact power flow set by a convex relaxation [33–36]. Relaxations are themselves compositional in a weak sense: the relaxation of a composite contains the composite of the relaxations, but the containment is generally strict, so the corresponding assignment is a lax rather than a strict functor. Making this precise, and identifying conditions under which the lax structure is tight, appears to be a well posed question and would connect the present framework to the exactness literature for radial networks.

12. CONCLUSION

We have given a compositional semantics for AC power flow in which open subnetworks form a hypergraph category, boundary behaviours are the values of a symmetric monoidal hypergraph functor, and the classical operation of Kron reduction is recovered as the restriction of that functor to linear devices. Heterogeneous microgrid configurations are organised as the fibres of an indexed category, and the monoidal Grothendieck construction assembles them into a total category in which islanding and mode switching are vertical morphisms and restriction is a Cartesian lift.

The payoff is a theorem rather than a heuristic: an edit confined to one component of a composition tree invalidates exactly the behaviours on the path to the root, so the update cost is $\mathcal{O}(w^3\,\mathrm{depth})$ and independent of network size. Measurements on a DER-augmented IEEE 33-bus feeder and on replicated systems of up to 2112 buses confirm this, with recomputed arithmetic constant at 1.90e3 flops across a 32-fold increase in size, update speed-ups of 17 to 21 for islanding, resynchronisation and dispatch events, and updated factorisations that are bitwise identical to globally recomputed ones. We have also identified where the framework does not help, and quantified the cost of the case it does not cover, namely edits that alter an interface rather than a component's interior.

Table 3. Topology and configuration events, 528-bus system, leaf size 8, 145 tree nodes.

event	global [ms]	local [ms]	speed-up	nodes recomp.	nodes reused	Newton it.
islanding	6.59	0.37	17.9	4	141	3
resynchronisation	6.33	0.36	17.5	4	141	3
DER dispatch change	6.35	0.31	20.7	3	142	3

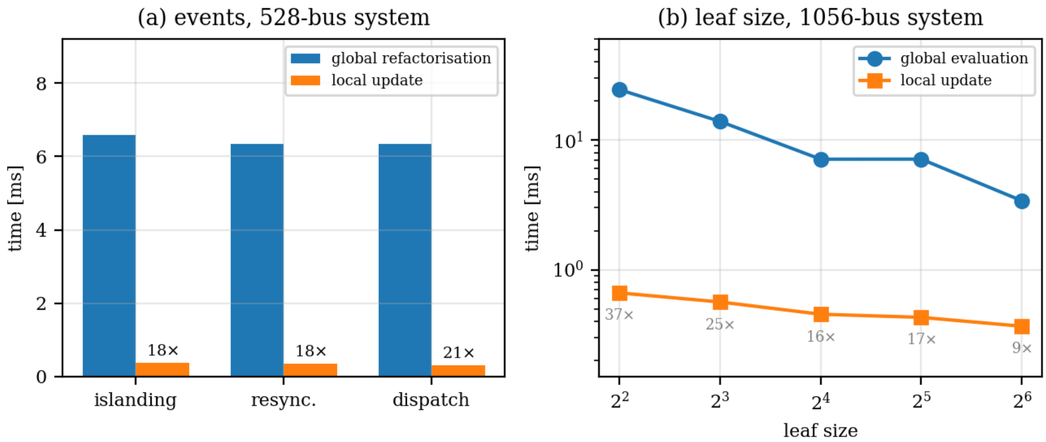

Figure 3. (a) Global refactorisation against local update for the three operational events. (b) Effect of the leaf size of the composition tree on the 1056-bus system; fine decompositions maximise the relative saving, coarse ones the absolute update time.

Three directions follow. Incremental structural re-analysis would remove the one remaining global cost, that of Phase 1 after an interface change. A lax-functorial treatment of convex relaxations would connect compositional evaluation to the convexification literature. Finally, the assignment of behaviours to regions is a functor into Beh, and operational limits such as voltage bands are subobjects of behaviours; the question of how such limits propagate across a composition boundary is naturally posed as a Kan extension along the inclusion of a region's interface, which we intend to develop.

REPRODUCIBILITY

All results in Section 10 were produced by the accompanying code, which is released with the paper: `ccpf.py` contains the case construction, the Newton solver and the composition-tree evaluator; `experiments.py` reproduces every table; `results.json` records the raw measurements including standard deviations. The benchmark data are taken unmodified from the pandapower case library. Running `python experiments.py` regenerates `results.json` in a few minutes on a single core.

REFERENCES

[1] W. F. Tinney and C. E. Hart, "Power flow solution by Newton's method," *IEEE Trans. Power Apparatus and Systems*, vol. PAS-86, no. 11, pp. 1449–1460, 1967.

[2] B. Stott and O. Alsaç, "Fast decoupled load flow," *IEEE Trans. Power Apparatus and Systems*, vol. PAS-93, no. 3, pp. 859–869, 1974.

[3] B. Stott, "Review of load-flow calculation methods," *Proc. IEEE*, vol. 62, no. 7, pp. 916–929, 1974.

[4] R. J. Lipton, D. J. Rose, and R. E. Tarjan, "Generalized nested dissection," *SIAM J. Numer. Anal.*, vol. 16, no. 2, pp. 346–358, 1979.

[5] J. A. George, "Nested dissection of a regular finite element mesh," *SIAM J. Numer. Anal.*, vol. 10, no. 2, pp. 345–363, 1973.

[6] G. Kron, *Diakoptics: The Piecewise Solution of Large-Scale Systems*. London: MacDonald, 1963.

[7] G. Kron, *Tensor Analysis of Networks*. New York: Wiley, 1939.

[8] H. H. Happ, "Diakoptics — the solution of system problems by tearing," *Proc. IEEE*, vol. 62, no. 7, pp. 930–940, 1974.

[9] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei, "A survey of distributed optimization and control algorithms for electric power systems," *IEEE Trans. Smart Grid*, vol. 8, no. 6, pp. 2941–2962, 2017.

Table 4. Effect of meshing. Leaf size 8; the meshed variant closes the five tie switches on every feeder.

topology	buses	branches/feeder	max front	max sep. vars	depth	global flops	speed-up
radial	33	32	16	16	3	2.73e3	2.7
radial	528	32	14	16	4	1.90e3	18.1
radial	2112	32	14	16	4	1.90e3	34.4
meshed	33	37	20	14	3	6.05e3	2.1
meshed	528	37	26	16	4	1.29e4	17.3
meshed	2112	37	26	16	5	1.29e4	34.2

Table 5. Sensitivity to leaf size, 1056-bus system.

leaf size	global [ms]	local [ms]	depth	fronts	max front	speed-up
4	24.44	0.66	5	545	10	36.9
8	13.90	0.56	4	289	14	24.6
16	7.11	0.45	3	129	26	15.7
32	7.12	0.43	3	126	66	16.6
64	3.40	0.37	2	33	68	9.3

[10] T. Erseghe, “Distributed optimal power flow using ADMM,” *IEEE Trans. Power Syst.*, vol. 29, no. 5, pp. 2370–2380, 2014.

[11] A. Floriduz, M. Tucci, S. Riverso, and G. Ferrari-Trecate, “Approximate Kron reduction methods for electrical networks with applications to plug-and-play control of AC islanded microgrids,” *IEEE Trans. Control Syst. Technol.*, vol. 27, no. 6, pp. 2403–2416, 2019.

[12] J. C. Willems, “The behavioral approach to open and interconnected systems,” *IEEE Control Systems Magazine*, vol. 27, no. 6, pp. 46–99, 2007.

[13] B. Fong, “Decorated cospans,” *Theory Appl. Categ.*, vol. 30, no. 33, pp. 1096–1120, 2015.

[14] J. C. Baez and K. Courser, “Structured cospans,” *Theory Appl. Categ.*, vol. 35, no. 48, pp. 1771–1822, 2020.

[15] J. C. Baez, K. Courser, and C. Vasilakopoulou, “Structured versus decorated cospans,” *Compositionality*, vol. 4, no. 3, 2022.

[16] J. C. Baez and B. Fong, “A compositional framework for passive linear networks,” *Theory Appl. Categ.*, vol. 33, no. 38, pp. 1158–1222, 2018.

[17] J. C. Baez, B. Coya, and F. Rebro, “Props in network theory,” *Theory Appl. Categ.*, vol. 33, no. 25, pp. 727–783, 2018.

[18] J. C. Baez and J. Erbele, “Categories in control,” *Theory Appl. Categ.*, vol. 30, no. 24, pp. 836–881, 2015.

[19] F. Bonchi, P. Sobociński, and F. Zanasi, “A categorical semantics of signal flow graphs,” in *CONCUR 2014*, LNCS vol. 8704, Springer, 2014, pp. 435–450.

[20] K. Courser, “A bicategory of decorated cospans,” *Theory Appl. Categ.*, vol. 32, no. 29, pp. 985–1027, 2017.

[21] B. Fong and D. I. Spivak, “Hypergraph categories,” *J. Pure Appl. Algebra*, vol. 223, no. 11, pp. 4746–4777, 2019.

[22] D. Vagner, D. I. Spivak, and E. Lerman, “Algebras of open dynamical systems on the operad of wiring diagrams,” *Theory Appl. Categ.*, vol. 30, no. 51, pp. 1793–1822, 2015.

[23] J. C. Baez, J. Foley, J. Moeller, and B. S. Pollard, “Network models,” *Theory Appl. Categ.*, vol. 35, no. 20, pp. 700–744, 2020.

[24] M. Shulman, “Framed bicategories and monoidal fibrations,” *Theory Appl. Categ.*, vol. 20, no. 18, pp. 650–738, 2008.

[25] M. Halter, E. Patterson, A. Baas, and J. Fairbanks, “Compositional scientific computing with Catlab and SemanticModels,” in *Applied Category Theory 2020*, 2020, arXiv:2005.04831.

[26] E. Patterson, O. Lynch, and J. Fairbanks, “Categorical data structures for technical computing,” *Compositionality*, vol. 4, no. 5, 2022.

[27] S. Libkind, A. Baas, M. Halter, E. Patterson, and J. Fairbanks, “An algebraic framework for structured epidemic modelling,” *Phil. Trans. R. Soc. A*, vol. 380, no. 2233, 20210309, 2022.

[28] B. Fong and D. I. Spivak, *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge: Cambridge University Press, 2019.

[29] J. S. Nolan, B. S. Pollard, S. Breiner, D. Anand, and E. Subrahmanian, "Compositional models for power systems," in *Applied Category Theory 2019*, EPTCS vol. 323, 2020, pp. 149–160.

[30] M. Lohmayer and S. Leyendecker, "EPHS: A port-Hamiltonian modelling language," *IFAC-PapersOnLine*, vol. 55, no. 30, 2022, arXiv:2202.00377.

[31] F. Dörfler and F. Bullo, "Kron reduction of graphs with applications to electrical networks," *IEEE Trans. Circuits Syst. I*, vol. 60, no. 1, pp. 150–163, 2013.

[32] F. Dörfler, J. W. Simpson-Porco, and F. Bullo, "Electrical networks and algebraic graph theory: Models, properties, and applications," *Proc. IEEE*, vol. 106, no. 5, pp. 977–1005, 2018.

[33] D. K. Molzahn and I. A. Hiskens, "A survey of relaxations and approximations of the power flow equations," *Found. Trends Electric Energy Syst.*, vol. 4, no. 1–2, pp. 1–221, 2019.

[34] J. Lavaei and S. H. Low, "Zero duality gap in optimal power flow problem," *IEEE Trans. Power Syst.*, vol. 27, no. 1, pp. 92–107, 2012.

[35] M. Farivar and S. H. Low, "Branch flow model: Relaxations and convexification, Part I," *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 2554–2564, 2013.

[36] R. A. Jabr, "Radial distribution load flow using conic programming," *IEEE Trans. Power Syst.*, vol. 21, no. 3, pp. 1458–1459, 2006.

[37] D. Bienstock and A. Verma, "Strong NP-hardness of AC power flows feasibility," *Oper. Res. Lett.*, vol. 47, no. 6, pp. 494–501, 2019.

[38] K. Lehmann, A. Grastien, and P. Van Hentenryck, "AC-feasibility on tree networks is NP-hard," *IEEE Trans. Power Syst.*, vol. 31, no. 1, pp. 798–801, 2016.

[39] D. Bienstock and G. Muñoz, "LP formulations for polynomial optimization problems," *SIAM J. Optim.*, vol. 28, no. 2, pp. 1121–1150, 2018.

[40] R. A. Jabr, "Exploiting sparsity in SDP relaxations of the OPF problem," *IEEE Trans. Power Syst.*, vol. 27, no. 2, pp. 1138–1139, 2012.

[41] D. K. Molzahn, J. T. Holzer, B. C. Lesieutre, and C. L. DeMarco, "Implementation of a large-scale optimal power flow solver based on semidefinite programming," *IEEE Trans. Power Syst.*, vol. 28, no. 4, pp. 3987–3998, 2013.

[42] J. R. Gilbert and R. E. Tarjan, "The analysis of a nested dissection algorithm," *Numer. Math.*, vol. 50, no. 4, pp. 377–404, 1987.

[43] P. R. Amestoy, T. A. Davis, and I. S. Duff, "Algorithm 837: AMD, an approximate minimum degree ordering algorithm," *ACM Trans. Math. Softw.*, vol. 30, no. 3, pp. 381–388, 2004.

[44] N. Robertson and P. D. Seymour, "Graph minors II: Algorithmic aspects of tree-width," *J. Algorithms*, vol. 7, no. 3, pp. 309–322, 1986.

[45] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed. New York: Springer, 1998.

[46] A. Joyal and R. Street, "The geometry of tensor calculus I," *Adv. Math.*, vol. 88, no. 1, pp. 55–112, 1991.

[47] P. Selinger, "A survey of graphical languages for monoidal categories," in *New Structures for Physics*, Lecture Notes in Physics vol. 813, Springer, 2011, pp. 289–355.

[48] A. Grothendieck, "Catégories fibrées et descente," in *Revêtements étales et groupe fondamental (SGA 1)*, LNM vol. 224, Springer, 1971, Exposé VI.

[49] B. Jacobs, *Categorical Logic and Type Theory*. Amsterdam: North-Holland, 1999.

[50] F. Borceux, *Handbook of Categorical Algebra 2: Categories and Structures*. Cambridge: Cambridge University Press, 1994.

[51] J. Moeller and C. Vasilakopoulou, "Monoidal Grothendieck construction," *Theory Appl. Categ.*, vol. 35, no. 31, pp. 1159–1207, 2020.

[52] S. Basu, R. Pollack, and M.-F. Roy, *Algorithms in Real Algebraic Geometry*, 2nd ed. Berlin: Springer, 2006.

[53] A. M. Kettner and M. Paolone, "On the properties of the power systems nodal admittance matrix," *IEEE Trans. Power Syst.*, vol. 33, no. 1, pp. 1130–1131, 2018.

[54] M. E. Baran and F. F. Wu, “Network reconfiguration in distribution systems for loss reduction and load balancing,” *IEEE Trans. Power Delivery*, vol. 4, no. 2, pp. 1401–1407, 1989.

[55] L. Thurner, A. Scheidler, F. Schäfer, J. Menke, J. Dollichon, F. Meier, S. Meinecke, and M. Braun, “pandapower — an open-source Python tool for convenient modeling, analysis, and optimization of electric power systems,” *IEEE Trans. Power Syst.*, vol. 33, no. 6, pp. 6510–6521, 2018.

[56] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, “MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education,” *IEEE Trans. Power Syst.*, vol. 26, no. 1, pp. 12–19, 2011.

[57] C. Coffrin, R. Bent, K. Sundar, Y. Ng, and M. Lubin, “PowerModels.jl: An open-source framework for exploring power flow formulations,” in *Power Systems Computation Conference*, 2018, pp. 1–8.

[58] N. A. Gezer, “Some remarks on iterated line graphs on Cayley graphs,” *Annals of Mathematics and Computer Science*, vol. 30, pp. 41–53, 2025.

[59] M. Labendia and C. Pornia, “Integer programming formulations and probabilistic bounds for some domination parameters,” *Annals of Mathematics and Computer Science*, vol. 30, pp. 1–15, 2025.

[60] L. Casinillo and E. Casinillo, “Outer-connected fair Roman dominating function in graphs,” *Annals of Mathematics and Computer Science*, vol. 31, pp. 32–42, 2025.

[61] S. Betley, “On Ext and Hyperext groups in the category of functors,” *Gulf Journal of Mathematics*, vol. 20, no. 2, pp. 1–12, 2025.

[62] N. Haddar, K. Azi, A. Deblij, and H. Hamraoui, “Quillen K-theory and representation rings of the quaternion group Q_8 and the dihedral group D_8,” *Gulf Journal of Mathematics*, vol. 22, no. 1, pp. 1–18, 2026.

[63] G. Ganesan, “Weighted Eulerian extensions of random graphs,” *Gulf Journal of Mathematics*, vol. 16, no. 2, pp. 1–11, 2024.

[64] F. Kamalov and H. H. Leung, “Fixed point theory: A review,” arXiv preprint arXiv:2309.03226, 2023.

[65] F. Kamalov and H. Leung, “Numerical computing in engineering mathematics,” in *2022 Advances in Science and Engineering Technology International Conferences (ASET)*, 2022, pp. 1–5.

[66] F. Kamalov and H. H. Leung, “Outlier detection in high dimensional data,” *Journal of Information & Knowledge Management*, vol. 19, no. 1, 2040013, 2020.