Distribution Network Transition Problem: A Planning Knowledge Model Capturing Structural Constraints

Francesco Percassi[1], **Sandra Castellanos-Paez**[1], **Mauro Vallati**[1] **Marie-Cécile Alvarez-Hérault**[2]

[1]School of Computing and Engineering, University of Huddersfield, Huddersfield, UK
[2]Univ. Grenoble Alpes, CNRS, Grenoble INP, G2Elab, F-38000 Grenoble, France
F.Percassi@hud.ac.uk, S.castellanos@hud.ac.uk, M.Vallati@hud.ac.uk, marie-cecile.alvarez@g2elab.grenoble-inp.fr

Abstract

Power grid long-term planning considers the problem of transforming a power distribution network from an initial configuration to a target one while preserving structural validity throughout the transition. We formalise this task as the Distribution Network Transition Problem (DNTP), a constrained graph transformation problem in which all intermediate networks must satisfy global topological properties, including radiality and robustness to failures. To address this task, we propose a planning knowledge model of DNTP in which network modifications are represented as actions and global constraints are captured declaratively. The resulting formulation is modular, separating local transition dynamics from global invariants, and ensures that only valid intermediate configurations are reachable. The proposed approach is evaluated on representative instances to analyse its computational behaviour and the impact of enforcing different classes of constraints. The results indicate a trade-off between compliance guarantees and solvability, and suggest that relaxed models often suffice in practice, paving the way for approaches that generate candidate solutions efficiently and validate them a posteriori.

Introduction

Electricity distribution networks are evolving due to the increasing penetration of renewable energy sources, the electrification of demand, and the deployment of smart grid technologies (Gönen, Ten, and Mehrizi-Sani 2024; Fang et al. 2011). As a result, power grid long-term planning has become a central task for Distribution System Operators (DSOs), who must determine not only suitable target network configurations but also how to transform existing infrastructures to reach them while maintaining operational reliability (Conejo, Carrion, and Morales 2010; Khator and Leung 1997; Sultana et al. 2016).

Power grid long-term planning is inherently a multi-stage and multidimensional process that is currently performed manually by domain experts. In practice, DSOs define a target network configuration based on forecasts of demand, generation, and quality-of-supply requirements, and then determine a sequence of infrastructure interventions, such as line constructions, removals, and switch reconfigurations, to transition from the current topology to the target over a multi-year horizon. Each intermediate network configuration remains in operation for extended periods and must therefore satisfy stringent operational requirements. In particular, all customers must remain supplied throughout the transition process: under normal conditions the network operates in a radial configuration, while under contingencies it must remain reconfigurable to restore supply after failures (Guo et al. 2016; Castelli et al. 2024).

Despite the multidimensional nature of long-term planning, which involves economic and electrotechnical aspects, the feasibility of any transition ultimately rests on a fundamental combinatorial problem: transforming one network topology into another through a sequence of admissible modifications, while preserving global structural invariants at every intermediate step. These invariants, most notably radial operation and robustness to single failures, define the space of permissible configurations and determine whether a transition is viable in the first place. If this discrete layer is not explicitly characterised, higher-level reasoning about economic or electrical feasibility necessarily operates over an incomplete or invalid transition space.

In this paper, we formalise this discrete core as the *Distribution Network Transition Problem* (DNTP) and propose a classical planning approach to address it. Our method models the transition process as a planning problem, capturing structural constraints inherent to distribution network topology while ensuring topological feasibility of the resulting plans. We present an empirical evaluation analysing the computational behaviour of the approach across a range of realistic scenarios, highlighting its potential as a structured tool for reasoning about the evolution of power distribution infrastructures.

Classical Planning with Axioms

Automated planning concerns the synthesis of action sequences that transform an initial state into a goal state under a formal transition model (Ghallab, Nau, and Traverso 2016). In its classical formulation, planning assumes a fully observable and deterministic environment where states are represented compactly by propositional variables. Over time, the classical framework has been extended with expressive modelling features (Speck, Seipp, and Torralba 2025), including conditional effects (Nebel 2000), state-dependent action costs (Speck et al. 2021), and axioms (Thiébaux, Hoffmann, and Nebel 2005). In this work, we adopt classical planning with axioms (Speck et al. 2019).

Axioms introduce derived variables whose truth values are defined declaratively from base variables. This enables complex properties to be encoded without simulating them through additional actions or variables. Without axioms, such properties would require cumbersome encodings that substantially increase model size and may incur exponential compilation overhead (Thiébaux, Hoffmann, and Nebel 2005).

A *classical planning problem with axioms* is defined as a tuple $\Pi = \langle F, D, X, A, I, G\rangle$. F is a finite set of *base variables*. A *state* is a set $s \subseteq F$ representing the variables that are true in that state. D is a finite set of *derived variables* such that $F \cap D = \emptyset$. Let $L(F \cup D) = \{p, \neg p \mid p \in F \cup D\}$ be the set of literals over base and derived variables. An *axiom* $r \in X$ is a rule of the form $r := d \leftarrow \psi$, where $d \in D$ is the head of r and ψ is a propositional formula over literals in $L(F \cup D)$, referred to as the body. Axioms define the truth values of derived variables that are not part of the state as a function of the current state and the axiom set X. The set A is a finite set of *actions*. Each action $a \in A$ is defined as a pair $a = \langle \mathit{pre}(a), \mathit{eff}(a)\rangle$, where $\mathit{pre}(a) \subseteq L(F \cup D)$ specifies the preconditions of a and $\mathit{eff}(a) \subseteq L(F)$ specifies its effects. The positive and negative effects of a are defined as $\mathit{eff}^+(a) = \{p \mid p \in \mathit{eff}(a)\}$ and $\mathit{eff}^-(a) = \{p \mid \neg p \in \mathit{eff}(a)\}$, respectively.

The initial state is given by a set $I \subseteq F$. The goal condition is a set of literals $G \subseteq L(F \cup D)$.

We assume that the set of axioms X is *stratified*. This assumption guarantees that, for every state s, the truth value of each derived variable is uniquely determined (Apt, Blair, and Walker 1988; Thiébaux, Hoffmann, and Nebel 2005). Given a state s, the *closure* of s under X, denoted by $\mathscr{C}_X(s) \subseteq F \cup D$, is obtained by evaluating the axioms layer by layer according to the stratification. The closure $\mathscr{C}_X(s)$ is called the *extended state* of s.

Given a state s and a literal $\ell \in L(F \cup D)$, we write $\mathscr{C}_X(s) \models \ell$ iff $\ell = p$ and $p \in \mathscr{C}_X(s)$, or $\ell = \neg p$ and $p \notin \mathscr{C}_X(s)$. For a set of literals $L \subseteq L(F \cup D)$, we write $\mathscr{C}_X(s) \models L$ if $\mathscr{C}_X(s) \models \ell$ holds for every $\ell \in L$. An action a is *applicable* in a state s if $\mathscr{C}_X(s) \models \mathit{pre}(a)$. Applying an applicable action a in s yields the successor state $\gamma(s, a) = (s \setminus \mathit{eff}^-(a)) \cup \mathit{eff}^+(a)$. The corresponding extended state is $\mathscr{C}_X(\gamma(s, a))$.

A *plan* for Π is a finite sequence of actions $\pi = \langle a_1, \ldots, a_n\rangle$. The plan π is *valid* if, starting from the initial state $s_0 = I$, there exists a sequence of states $\langle s_1, \ldots, s_n\rangle$ such that for all $i = 0, \ldots, n-1$, $\mathscr{C}_X(s_i) \models \mathit{pre}(a_{i+1})$ and $s_{i+1} = \gamma(s_i, a_{i+1})$, and $\mathscr{C}_X(s_n) \models G$. The cost of a plan π is defined as $|\pi|$. A plan is optimal if, for any other plan π', $|\pi| \leq |\pi'|$.

Hereinafter, to simplify the notation, $\Pi = \langle F, D, X, A, I, G\rangle$ will be denoted more compactly as $\Pi = \langle F, X, A, I, G\rangle$, leaving the set of derived variables implicitly defined as $D = \{\mathit{head}(r) \mid r \in X\}$, where $\mathit{head}(r)$ denotes the head of rule r.

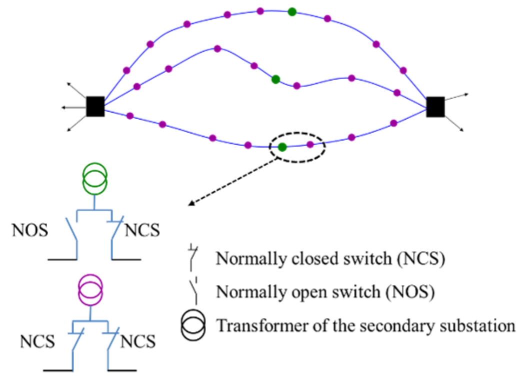

Figure 1: Medium-voltage distribution network supplied by two primary substations (black squares) feeding interconnected feeders (blue lines) and secondary substations; purple circles represent secondary substations equipped with two normally closed switches (NCS), while green circles denote secondary substations equipped with one NCS and one normally open switch (NOS) used for radial operation and reconfiguration. In this work, we treat both types of secondary substations uniformly, and represent their operational differences through the status of the network edges (open, close). (Figure adapted from (Castellanos, Alvarez-Herault, and Lalanda 2023).)

Research Context

Urban electrical distribution networks are commonly designed are commonly designed according to a *secured feeder* architecture (see Figure 1) (Gönen, Ten, and Mehrizi-Sani 2024). In this architecture, primary substations act as the supply points of the medium-voltage network, injecting power into the system through main feeders that interconnect substations and supply multiple secondary substations. These secondary substations serve local areas and include distribution transformers that step down the medium-voltage level to low voltage for end users and are equipped with switching devices that enable operational flexibility.

Each secondary substation is connected to a feeder through two switches: normally closed switches (NCS), which allow power flow under normal operating conditions, and normally open switches (NOS), which remain open to enable network reconfiguration when contingencies occur. The secure feeder architecture operates according to a radial structure under normal conditions, enhancing supply reliability. This radial configuration is ensured by maintaining exactly one normal open switch (NOS) along each feeder, with all other switches closed (e.g.. each secondary substation is supplied from exactly one primary substation through the active network). When a fault occurs, switching operations allow the affected section to be isolated and service to be rapidly restored to downstream customers via alternative feeder paths during repair operations. These structural principles constrain admissible network topologies.

Infrastructure modifications in distribution grids are expensive and difficult to reverse. The cost of a single power line may range from 50 K€ to 200 K€ per Kilometre, depending on technology (overhead vs. underground), cable type, and local conditions (rural, semi-urban, or urban). At the scale of a national distribution system such as the French network, which extends over approximately 1.4 Million Kilometres, even minor changes can result in significant investments. Consequently, DSOs aim to minimise deviations from the existing infrastructure, favouring incremental and targeted interventions (Conejo, Carrion, and Morales 2010).

To ensure that intermediate network states remain operational during the development phase and to prevent inefficient resource allocation, grid operators typically follow established planning rules. Based on expert interviews with French DSOs and secured feeder design principles, Castellanos, Alvarez-Herault, and Lalanda (2023) identified the following structural and operational rules that govern admissible intermediate networks:

- R1 There is no closed loop between two primary substations; that is, a normally open switch is required to ensure radial operation, meaning that each secondary substation in the network is supplied by exactly one primary substation.
- R2 The network must be able to withstand a failure of a single component, ensuring that a radial configuration, in which all secondary nodes remain supplied, can always be restored via switching operations (opening and closing lines). This is referred to as the $N-1$ reliability criterion (Commission 2017).
- R3 For urban areas, the number of connections of secondary substations is between 2 and 3.
- R4 Only lines belonging to the target network can be created, and only lines not belonging to the target network can be removed. These operations are irreversible, e.g., a line cannot be removed and later re-added.
- R5 There is only one connection between two nodes (distribution/source).

Problem Formalisation

We model the discrete core of long-term distribution network planning as a constrained graph transformation problem. Starting from an initial network, the goal is to reach a target one through admissible edits while preserving structural validity at every intermediate step.

Definition 1 (Distribution Graph). *A* distribution graph *is a tuple $\mathcal{G} = (V, E, \sigma)$, where $V = V_P \cup V_S$ is the set of substations, partitioned into primary substations V_P and secondary substations V_S, with $V_P \cap V_S = \emptyset$ and $|V_P| \geq 2$; $E \subseteq \binom{V}{2}$ is the set of undirected lines; and $\sigma : E \to \{\top, \bot\}$ labels each line as closed ($\top$) or open ($\bot$). The* active distribution graph *induced by $\mathcal{G}$ is $\mathcal{G}^\top = (V, E^\top)$, where $E^\top = \{e \in E \mid \sigma(e) = \top\}$.*

For $v \in V$, let $\deg_{\mathcal{G}}(v)$ denote its degree in $\mathcal{G}$. We write $u \rightsquigarrow_{\mathcal{G}} v$ when u and v are connected in $\mathcal{G}$, and $u \rightsquigarrow_{\mathcal{G}^\top} v$ when they are connected in $\mathcal{G}^\top$.

Definition 2 (Radial and $N-1$ Compliant Distribution Graph). *A distribution graph $\mathcal{G}$ is* radial *iff $\mathcal{G}^\top$ is a forest and every connected component of $\mathcal{G}^\top$ contains exactly one primary substation.*

Let $f \in V_P \cup E$ be a fault. If $f \in V_P$, the faulty primary and all its incident lines are removed; if $f \in E$, only that line is removed. The resulting graph is denoted by $\mathcal{G}^{(f)}$.

A distribution graph $\mathcal{G}$ is $N-1$ compliant *iff for every fault $f \in V_P \cup E$ and every $v \in V_S$, there exists a primary substation u such that $v \rightsquigarrow_{\mathcal{G}^{(f)}} u$, with $u \neq f$ when $f \in V_P$.*

We consider three graph editing operations on a distribution graph: (i) addition of a new open line, (ii) removal of an existing open line, and (iii) double-switching, which opens a closed line and closes an open incident line at the same secondary substation. All operations are assumed to be well-defined.

Definition 3 (Edit Path). *An* edit path *on $\mathcal{G}$ is a finite sequence $\varphi = (\mathsf{op}_1, \ldots, \mathsf{op}_m)$ of well-defined editing operations. It induces a sequence of graphs $\mathcal{G}_{(0)} := \mathcal{G}$ and $\mathcal{G}_{(i)} := \mathsf{op}_i(\mathcal{G}_{(i-1)})$ for $i \in \{1, \ldots, m\}$. We write $\varphi(\mathcal{G}) := \mathcal{G}_{(m)}$.*

Definition 4 (DNTP). *A* Distribution Network Transition Problem (DNTP) *is a pair $\Phi = (\mathcal{G}, \mathcal{G}^\star)$ of distribution graphs over the same set of substations. A solution is a feasible edit path φ such that $\varphi(\mathcal{G}) = \mathcal{G}^\star$ and every intermediate graph $\mathcal{G}_{(i)}$ satisfies:*

- *C1.* $E_{(i)} \subseteq E \cup E^\star$*;*
- *C2. lines removed from E cannot be reintroduced, and lines added from $E^\star$ cannot be removed;*
- *C3.* $2 \leq \deg_{\mathcal{G}_{(i)}}(v) \leq 3$ *for every $v \in V_S$;*
- *C4. $\mathcal{G}_{(i)}$ is radial;*
- *C5. $\mathcal{G}_{(i)}$ is $N-1$ compliant.*

An optimal edit path for Φ is a feasible one that minimises $|\varphi|$.

DNTP as Classical Planning

In this section, we encode a DNTP instance as an equisatisfiable planning task, i.e., one admitting a valid plan iff the original instance admits a feasible edit path, with the same cost. The encoding is modular. We first introduce a core formulation capturing the local constraints (C1)–(C3) using classical planning, and then extend it with axioms and derived variables to model the global constraints (C4)–(C5).

We define the auxiliary sets $E^{\text{build}} = E^\star \setminus E$, $E^{\text{rem}} = E \setminus E^\star$, and $E^{\text{all}} = E \cup E^\star$, representing, respectively, lines to be built, lines to be removed, and all lines appearing in either the initial or target network.

To obtain a compact and symmetry-less representation, we assume a fixed total order $\prec$ over V such that for any distinct $v, u \in V$, exactly one of $v \prec u$ or $u \prec v$ holds. Moreover, $v \prec u$ always holds when $v \in V_S$ and $u \in V_P$. This allows us to define a canonical representation of binary predicates of the form $x\{v, u\} = x(\min(v, u), \max(v, u))$. Since we are dealing with an undirected graph, this allows us to avoid, for example, keeping $conn(v, u)$ and $conn(u, v)$ to represent the connection between v and u, and the consequent duplication in the preconditions and effects of actions. The same convention applies for action, i.e., $\mathsf{action}\{v, u\} = \mathsf{action}(\min(v, u), \max(v, u))$.

Local Formulation

Let $\Phi = \langle G, G^\star \rangle$ be a DNTP instance. We encode it as a classical planning task $\Pi^\Phi = \langle F, A, I, G \rangle$ capturing the local transition dynamics.

State Representation. The planning state represents the current network topology using three types of variables: (i) $conn(v, u)$ indicates whether a line is present, (ii) $\mathsf{c}(v, u)$ indicates whether a present line is closed, and (iii) $full(v)$ tracks whether a secondary substation has reached its maximum degree. Together, these variables provide a compact representation of intermediate network configurations.

Actions. Actions mirror the graph editing operations defining DNTP. In particular, we consider: (i) addition of new lines, (ii) removal of existing lines, and (iii) switching operations that reconfigure connections locally. Their preconditions enforce local feasibility conditions (e.g., degree bounds and line availability), while their effects reproduce the corresponding modifications of the network.

Initial and Goal States. The initial state encodes the input network G, while the goal condition enforces that the resulting state matches the target network $G^\star$ in terms of both topology and line status.

Correctness. This encoding faithfully captures the local constraints of DNTP. In particular, feasible edit paths correspond to valid plans, and vice versa, preserving both feasibility and cost.

Global Formulation

The local encoding captures constraints (C1)–(C3), but not the global constraints (C4) and (C5), which depend on the network as a whole. Given a DNTP instance Φ and its local planning task $\Pi^\Phi = \langle F, \emptyset, A, I, G \rangle$, we define the corresponding global planning task as

$$\Pi^\Phi_{\text{FULL}} = \langle F, X_{\text{FULL}}, A_{\text{FULL}}, I, G \cup \{\textsc{rad}, \textsc{rest}\} \rangle. \quad (1)$$

Π^Φ_{FULL} extends Π^Φ with two derived variables, RAD and REST, capturing radiality and failure-restoration ($N-1$ compliance) requirements, respectively, through the axioms in X_{FULL}. To enforce these constraints throughout the transition, each action $a \in A$ is lifted to an action in A_{FULL} whose preconditions include $\{\textsc{rad}, \textsc{rest}\}$, and the same requirement is added to the goal. After each action, the axioms recompute both derived variables; if either becomes false, the resulting state is a dead end. In this way, constraints (C4) and (C5) are enforced at every step of the plan.

Handling Radiality Radiality is enforced by introducing a derived reachability predicate $\textsc{rc}\{v, u\}$, which holds if v and u are connected through closed lines. This predicate is defined recursively using axioms capturing direct connections and path composition.

From this reachability relation, radiality is encoded by deriving a predicate RAD that enforces two conditions: (i) every secondary substation is connected to at least one primary, and (ii) no two distinct primaries are connected through closed lines. Formally:

$$\textsc{rad} \leftarrow \bigwedge_{v \in V_S} \bigvee_{p \in V_P} \textsc{rc}\{v, p\} \;\wedge \bigwedge_{\substack{p_1, p_2 \in V_P \\ p_1 \prec p_2}} \neg \textsc{rc}\{p_1, p_2\}. \quad (2)$$

The first condition ensures that each secondary is connected to the network, while the second guarantees that connected components contain at most one primary, preventing cycles between feeders.

These axioms are integrated into the planning task by requiring RAD to hold before and after every action. As a result, any state violating radiality becomes a dead end and cannot be extended into a valid plan.

This encoding captures radiality in the transition setting: feasible transitions correspond to plans in which RAD holds throughout execution.

Handling $N-1$ Compliance To enforce $N-1$ compliance, we introduce a derived predicate REST, which holds when every secondary substation remains connected to some primary after the failure of any single primary substation or line. We decompose this condition into two parts:

$$\textsc{rest} \leftarrow \textsc{rest-p} \wedge \textsc{rest-e},$$

where REST-P captures primary failures and REST-E captures line failures.

For primary failures, we define a derived reachability predicate $\textsc{rco-ex-s}(v, u, v_\dagger)$ expressing that v remains connected to a primary u while excluding the failed primary $v_\dagger$. Reachability is evaluated over existing lines, independently of their open/closed status, and recursion is restricted so that the failed primary cannot be traversed. Using this relation, we derive

$$\textsc{rest-p} \leftarrow \bigwedge_{v_\dagger \in V_P} \bigwedge_{v \in V_S} \bigvee_{u \in V_P \setminus \{v_\dagger\}} \textsc{rco-ex-s}\{v, u, v_\dagger\}, \quad (3)$$

which states that every secondary must remain connected to some surviving primary after the failure of any primary substation.

For line failures, we analogously define a derived predicate $\textsc{rco-ex-e}(v, u, v_\dagger, u_\dagger)$ expressing connectivity between v and u while excluding the failed line $\{v_\dagger, u_\dagger\}$. This yields the rule

$$\textsc{rest-e} \leftarrow \bigwedge_{\{v_\dagger, u_\dagger\} \in E^{\text{all}}} \left(conn\{v_\dagger, u_\dagger\} \Rightarrow \bigwedge_{v \in V_S} \bigvee_{u \in V_P} \textsc{rco-ex-e}\{v, u, v_\dagger, u_\dagger\} \right), \quad (4)$$

ensuring that every secondary remains connected to a primary after the failure of any present line.

These axioms are combined into the derived predicate REST and enforced throughout execution by requiring REST before and after every action. Hence, any state violating $N-1$ compliance becomes a dead end. This encoding captures the intended restoration requirement.

Full Encoding The complete planning model is obtained by combining the local formulation with the axioms enforcing radiality and $N-1$ compliance. The resulting task augments each action with preconditions requiring both RAD and REST to hold, and enforces the same conditions on the goal. This yields a unified encoding in which local transition dynamics are captured by actions, while global structural constraints are enforced through derived predicates at every step of the plan. The encoding is equisatisfiable and cost-preserving: a DNTP instance admits a feasible edit path if and only if the corresponding planning task admits a valid plan of the same length. Detailed correctness proofs are beyond the scope of this paper and will be provided in an extended version.

Experimental Analysis

We evaluate the proposed approach along two dimensions: (i) the representational and computational cost of enforcing different global constraints, and (ii) the compliance of the generated plans under progressively stronger encodings.

Experimental Settings

Experiments were run on Linux on Intel Xeon Gold 6140M CPUs at 2.30GHz, with 8GB RAM and a timeout of 1,800 seconds per run.

All encodings were implemented in lifted PDDL and solved using FAST DOWNWARD (Helmert 2006), with LAMA (Richter and Westphal 2010) in satisficing mode (first plan found) and SYMK (Speck, Seipp, and Torralba 2025) in optimal mode.

The benchmark includes 11 real-world DNTP instances derived from anonymised French distribution network data, with sizes ranging from 6 to 150 secondary substations.

We compare the following encodings: Π^{Φ} (local only), Π^{Φ}_{RAD} (radiality), $\Pi^{\Phi}_{\text{REST-P}}$ (primary-failure compliance), $\Pi^{\Phi}_{\text{REST-E}}$ (line-failure compliance), $\Pi^{\Phi}_{\text{RAD+REST-P}}$ (radiality + primary-failure compliance), and Π^{Φ}_{FULL} (full model).

Bottleneck Analysis

Table 1 and Table 2 show the impact of enforcing individual global constraints. The main representational bottleneck is $\Pi^{\Phi}_{\text{REST-E}}$, which causes a sharp growth in the number of grounded axioms and leads to memory exhaustion from medium-size instances onward. By contrast, Π^{Φ}_{RAD} and $\Pi^{\Phi}_{\text{REST-P}}$ remain compact and scale much more smoothly.

The runtime results follow the same trend. For LAMA, Π^{Φ} and Π^{Φ}_{RAD} remain inexpensive, $\Pi^{\Phi}_{\text{REST-P}}$ adds a moderate overhead, and $\Pi^{\Phi}_{\text{REST-E}}$ quickly becomes prohibitive. The effect is even stronger for SYMK, where line-failure compliance becomes intractable already on relatively small instances. Overall, explicit modelling of line-failure robustness is the main source of computational difficulty.

Compliance Analysis

Table 3 reports plan validation results for encodings that incrementally combine global constraints. As expected, the local encoding may violate radiality and reliability constraints. Enforcing radiality removes radiality violations, while adding primary-failure compliance removes violations of that class. The full encoding is the only one that guarantees complete compliance by construction.

A key practical result is that the relaxed encoding $\Pi^{\Phi}_{\text{RAD+REST-P}}$ already performs very well. In the satisficing

$(\|V\|,\|E\|)$	Number of axioms			
	Π^{Φ}	Π^{Φ}_{RAD}	$\Pi^{\Phi}_{\text{REST-P}}$	$\Pi^{\Phi}_{\text{REST-E}}$
(8,11)	0	58	58	1,369
(12,18)	0	94	93	5,296
(12,18)	0	95	96	5,489
(18,24)	0	138	147	19,657
(22,34)	0	184	188	38,238
(30,31)	0	238	254	97,300
(42,43)	0	358	386	295,961
(52,54)	0	469	499	MO
(70,102)	0	601	655	MO
(102,115)	0	729	784	MO
(152,176)	0	1,117	1,187	MO

Table 1: Real-world DNTP instances. Number of grounded axioms generated by the FAST DOWNWARD translator per instance and encoding. "MO" indicates that the approach ran out of memory.

$(\|V\|,\|E\|)$	LAMA-first			
	Π^{Φ}	Π^{Φ}_{RAD}	$\Pi^{\Phi}_{\text{REST-P}}$	$\Pi^{\Phi}_{\text{REST-E}}$
(8,11)	0.42	0.41	0.42	0.74
(12,18)	0.47	0.47	0.56	2.76
(12,18)	0.45	0.43	0.50	2.64
(18,24)	0.51	0.53	1.00	12.37
(22,34)	0.60	0.66	1.68	29.34
(30,31)	0.89	1.06	3.29	128.87
(42,43)	1.62	2.14	7.39	707.46
(52,54)	3.50	6.31	19.12	MO
(70,102)	6.34	8.20	47.88	MO
(102,115)	15.62	22.93	35.56	MO
(152,176)	47.49	81.10	103.56	MO
$(\|V\|,\|E\|)$	SYMK			
	Π^{Φ}	Π^{Φ}_{RAD}	$\Pi^{\Phi}_{\text{REST-P}}$	$\Pi^{\Phi}_{\text{REST-E}}$
(8,11)	0.76	1.05	1.09	4.84
(12,18)	0.93	1.25	1.33	60.07
(12,18)	0.93	1.24	1.30	63.22
(18,24)	1.17	1.63	2.84	788.00
(22,34)	1.57	2.13	4.15	TO
(30,31)	21.27	11.34	980.24	MO
(42,43)	TO	TO	MO	MO
(52,54)	TO	TO	TO	MO
(70,102)	TO	TO	MO	MO
(102,115)	16.58	33.35	41.77	MO
(152,176)	71.29	TO	TO	MO

Table 2: Real-world DNTP instances. CPU runtime (seconds) per instance and encoding for LAMA-FIRST (suboptimal) and SYMK (optimal). "MO" and "TO" mean memory out and time out.

$(\lvert V\rvert, \lvert E\rvert)$	LAMA-first							
	Constraint Compliance				CPU runtime			
	Π^{Φ}	Π^{Φ}_{RAD}	$\Pi^{\Phi}_{\text{RAD+REST-P}}$	Π^{Φ}_{FULL}	Π^{Φ}	Π^{Φ}_{RAD}	$\Pi^{\Phi}_{\text{RAD+REST-P}}$	Π^{Φ}_{FULL}
(8,11)	×	✓	✓	✓	0.42	0.41	0.50	0.83
(12,18)	✓	✓	✓	✓	0.47	0.47	0.57	2.45
(12,18)	×	✓	✓	✓	0.45	0.43	0.58	2.79
(18,24)	×	✓	✓	✓	0.51	0.53	1.05	14.21
(22,34)	×	✓	✓	✓	0.60	0.66	1.61	33.54
(30,31)	×	✓	✓	✓	0.89	1.06	3.84	122.15
(42,43)	×	◇	◇	✓	1.62	2.14	7.93	574.67
(52,54)	×	✓	◇	na	3.50	6.31	24.70	MO
(70,102)	×	✓	✓	na	6.34	8.20	47.86	MO
(102,115)	✓	✓	✓	na	15.62	22.93	43.70	MO
(152,176)	✓	✓	✓	na	47.49	81.10	128.88	MO

$(\lvert V\rvert, \lvert E\rvert)$	SymK							
	Constraint Compliance				CPU runtime			
	Π^{Φ}	Π^{Φ}_{RAD}	$\Pi^{\Phi}_{\text{RAD+REST-P}}$	Π^{Φ}_{FULL}	Π^{Φ}	Π^{Φ}_{RAD}	$\Pi^{\Phi}_{\text{RAD+REST-P}}$	Π^{Φ}_{FULL}
(8,11)	✓	✓	✓	✓	0.76	1.05	1.13	5.37
(12,18)	✓	✓	✓	✓	0.93	1.25	1.55	69.93
(12,18)	△	△	✓	✓	0.93	1.24	1.53	73.21
(18,24)	✓	✓	✓	✓	1.17	1.63	2.62	828.48
(22,34)	✓	✓	✓	na	1.57	2.13	12.11	TO
(30,31)	×	✓	✓	na	21.27	11.34	209.52	MO
(42,43)	na	na	na	na	TO	TO	TO	MO
(52,54)	na	na	na	na	TO	TO	TO	MO
(70,102)	na	na	na	na	TO	TO	TO	MO
(102,115)	✓	✓	✓	na	16.58	33.35	76.89	MO
(152,176)	×	na	na	na	71.29	TO	TO	MO

Table 3: Real-world DNTP instances. Constraint satisfaction per instance and encoding for LAMA-FIRST and SYMK. *Left*: compliance results (×: radiality violation; △: $N-1$ primary violation; ◇: $N-1$ line violation; ✓: all constraints satisfied). *Right*: CPU runtime (seconds).

setting, it solves all 11 real-world instances, and 9 of the returned plans are validated as fully compliant. In the optimal setting, it solves 7 out of 11 instances, again often producing fully compliant plans. By contrast, Π^{Φ}_{FULL} provides full guarantees but at a substantially higher computational cost, returning fully compliant plans only on smaller instances.

Overall, the experiments highlight a clear trade-off between explicit compliance guarantees and solvability. While the full encoding is often too expensive, enforcing radiality and primary-failure robustness already provides a strong and practically useful approximation, often yielding plans that are fully compliant even without explicitly modelling line-failure robustness.

Related Work

Artificial Intelligence techniques have been applied to power systems from different perspectives, addressing different decision-support and control problems. Early symbolic and knowledge-based approaches investigated expert systems for power system operation and operator assistance (Wollenberg and Sakaguchi 1987). Planning-based methods have been applied to problems such as power substation voltage control (Bell et al. 2009) and supply restoration after faults (Bertoli et al. 2002; Thiébaux et al. 2013), where the focus lies on short-term operational recovery. More recently, automated planning has also been explored in the context of long-term planning for smart grids (Castellanos-Paez, Alvarez-Herault, and Lalanda 2022; Castellanos, Alvarez-Herault, and Lalanda 2023). The present work builds upon this line of research by further distilling the DNTP into an independent combinatorial abstraction and by providing a formal encoding into classical planning using axioms.

Finally, other AI paradigms have been explored for power system decision support, including evolutionary computation techniques such as genetic algorithms for energy management (Arabali et al. 2012), as well as data-driven approaches based on reinforcement learning for control and decision-making tasks (Glavic, Fonteneau, and Ernst 2017; Vázquez-Canteli and Nagy 2019). From a graph-theoretic perspective, the DNTP can be related to graph transformation and graph edit distance (Gao et al. 2010), but with significant differences due to the restricted operations and the strong structural constraints. While effective in their respective domains, these approaches typically focus on operational optimisation or learning-based control and do not address the formal structure of long-term planning.

Conclusions

In this paper, we introduced the Distribution Network Transition Problem (DNTP) as the discrete core of long-term distribution network planning. We model DNTP as a constrained graph transformation problem and encode it in classical planning, where editing operations correspond to actions and global topological constraints are captured declaratively through axioms. The resulting formulation is modular and preserves the structure of feasible edit paths.

Our results highlight a clear trade-off between expressiveness and scalability. While radiality and primary-substation $N-1$ compliance can be handled with moderate overhead, line-failure $N-1$ compliance constitutes the main bottleneck due to the size of the resulting encoding. At the same time, relaxed models often yield fully compliant solutions after validation, suggesting that partial constraint enforcement provides a practical compromise.

Future work will focus on improving the scalability of the approach while retaining its declarative structure. In particular, this includes exploring incremental enforcement of global constraints, where relaxed models are refined only when violations are detected, as well as decomposition or hybrid approaches that combine a discrete planning layer with specialised validation procedures for global properties. Together, these directions point toward more flexible and scalable methods for reasoning about network transitions within a planning framework.

References

Apt, K. R.; Blair, H. A.; and Walker, A. 1988. Towards a Theory of Declarative Knowledge. In *Foundations of Deductive Databases and Logic Programming*, 89–148. Morgan Kaufmann.

Arabali, A.; Ghofrani, M.; Etezadi-Amoli, M.; Fadali, M. S.; and Baghzouz, Y. 2012. Genetic-algorithm-based optimization approach for energy management. *IEEE Transactions on Power Delivery*, 28(1): 162–170.

Bell, K.; Coles, A.; Coles, A.; Fox, M.; and Long, D. 2009. The role of AI planning as a decision support tool in power substation management. *Ai Communications*, 22(1): 37–57.

Bertoli, P.; Cimatti, A.; Slaney, J. K.; and Thiébaux, S. 2002. Solving power supply restoration problems with planning via symbolic model checking. In *Proceedings of ECAI*, 576–580.

Castellanos, S.; Alvarez-Herault, M.-C.; and Lalanda, P. 2023. Decision support tool for the development of power distribution networks based on AI planning. In *27th International Conference on Electricity Distribution (CIRED 2023)*, volume 2023, 1425–1429. IET.

Castellanos-Paez, S.; Alvarez-Herault, M.-C.; and Lalanda, P. 2022. Automated Planning to Evolve Smart Grids with Renewable Energies. In *Artificial Intelligence for Knowledge Management, Energy, and Sustainability*, 141–155. Springer.

Castelli, A. F.; Pilotti, L.; Monchieri, A.; and Martelli, E. 2024. Optimal design of aggregated energy systems with (N-1) reliability: MILP models and decomposition algorithms. *Applied Energy*, 356: 122002.

Commission, E. 2017. Commission Regulation Establishing a Guideline on Electricity TSO - 2017/1485 - EN - EUR-Lex. https://eur-lex.europa.eu/eli/reg/2017/1485/oj/eng. [Accessed 27-02-2026].

Conejo, A. J.; Carrion, M.; and Morales, J. M. 2010. *Decision Making Under Uncertainty in Electricity Markets*. Springer.

Fang, X.; Misra, S.; Xue, G.; and Yang, D. 2011. Smart grid—The new and improved power grid: A survey. *IEEE communications surveys & tutorials*, 14(4): 944–980.

Gao, X.; Xiao, B.; Tao, D.; and Li, X. 2010. A survey of graph edit distance. *Pattern Analysis and applications*, 13(1): 113–129.

Ghallab, M.; Nau, D.; and Traverso, P. 2016. *Automated planning and acting*. Cambridge University Press.

Glavic, M.; Fonteneau, R.; and Ernst, D. 2017. Reinforcement learning for electric power system decision and control: Past considerations and perspectives. *ifac-Papersonline*, 50(1): 6918–6927.

Gönen, T.; Ten, C.-W.; and Mehrizi-Sani, A. 2024. *Electric power distribution engineering*. CRC press.

Guo, Z.; Chen, R. L.-Y.; Fan, N.; and Watson, J.-P. 2016. Contingency-constrained unit commitment with intervening time for system adjustments. *IEEE Transactions on Power Systems*, 32(4): 3049–3059.

Helmert, M. 2006. The Fast Downward Planning System. *J. Artif. Intell. Res.*, 26: 191–246.

Khator, S. K.; and Leung, L. C. 1997. Power distribution planning: A review of models and issues. *IEEE Transactions on Power Systems*, 12(3): 1151–1159.

Nebel, B. 2000. On the compilability and expressive power of propositional planning formalisms. *Journal of Artificial Intelligence Research*, 12: 271–315.

Richter, S.; and Westphal, M. 2010. The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks. *J. Artif. Intell. Res.*, 39: 127–177.

Speck, D.; Borukhson, D.; Mattmüller, R.; and Nebel, B. 2021. On the compilability and expressive power of state-dependent action costs. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 31, 358–366.

Speck, D.; Geißer, F.; Mattmüller, R.; and Torralba, Á. 2019. Symbolic Planning with Axioms. In *ICAPS*, 464–472. AAAI Press.

Speck, D.; Seipp, J.; and Torralba, Á. 2025. Symbolic Search for Cost-Optimal Planning with Expressive Model Extensions. *J. Artif. Intell. Res.*, 82.

Sultana, B.; Mustafa, M.; Sultana, U.; and Bhatti, A. R. 2016. Review on reliability improvement and power loss reduction in distribution system via network reconfiguration. *Renewable and sustainable energy reviews*, 66: 297–310.

Thiébaux, S.; Coffrin, C.; Hijazi, H.; and Slaney, J. 2013. Planning with MIP for supply restoration in power distribution systems. In *Proceedings of IJCAI*, 2900–2907.

Thiébaux, S.; Hoffmann, J.; and Nebel, B. 2005. In defense of PDDL axioms. *Artificial Intelligence*, 168(1-2): 38–69.

Vázquez-Canteli, J. R.; and Nagy, Z. 2019. Reinforcement learning for demand response: a review of algorithms and modeling techniques. *Applied Energy*, 235: 1072–1089.

Wollenberg, B. F.; and Sakaguchi, T. 1987. Artificial intelligence in power system operations. *Proceedings of the IEEE*, 75(12): 1678–1685.