The Application of Planning to Power Substation Voltage Control

Keith Bell, Andrew Coles, Maria Fox, Derek Long and Amanda Smith
Strathclyde Centre for Applied Intelligent Systems,
University of Strathclyde, Glasgow, G1 1XH, UK
email: firstname.lastname@cis.strath.ac.uk

Abstract

The management of power substations is a complex control task involving the coordination of interactions between customer demand and substation components. In this paper we look at the efficient management of substations in a *pre-fault* state through the use of planning, considering the supply criteria, fault management and implied operational costs of control decisions. We discuss the modelling of the problem in PDDL and the performance of current state-of-the-art planners, followed by details of an optimisation system for the problem based around iterative refinement.

1 Introduction[1]

The control of electricity distribution is a challenging problem, where the infrastructure has to be able to supply a desired voltage to customers in the context of demand that varies significantly throughout the day. The workhorse at the heart of this is the power substation, where large distribution voltages are stepped down to lower voltages and, ultimately, to the domestic supply voltages. To maintain system stability and customer voltage requirements, the voltages on the circuits leaving a substation must lie within a specified voltage range. To attain these, a reactive system at the substation is used to meet voltage target within this range, set for each time of the day. The reactive system makes use of the adjustable substation components (Wood & Wollenberg 1996): multi-tap transformers and mechanically-switched capacitors.

The problem of setting voltage targets is an interesting multi-objective optimisation problem: the voltages must be maintained within the security standards, but wear-and-tear on substation components must be minimised while achieving this. At present, voltage targets are set by power systems engineers, with a great deal of concern being given to the former of these criteria but with little consideration being given to the latter.

In this paper, we discuss how we have used AI planning techniques to build a system for defining substation voltage targets. First, we will describe the problem in more detail, before considering how it can be modelled in PDDL (Fox & Long 2003). Then, we discuss how state-of-the-art planners fare when faced with the problem of defining voltage targets for the Grendon substation, near London, England. Based on these results, we discuss the system we constructed to exploit what we found to be possible, where a planner is used as a kernel of an optimisation system. Finally, we discuss the efficacy of our system in this domain, and future directions for the work.

[1]Some of the material presented here also appears in a paper currently under review for publication in *AI Communications*.

2 Substation Voltage Control

In the regulation of the distribution of electricity through a power grid, the output voltage on a circuit leaving a power substation must be within a certain range of the ideal voltage. These minimum and maximum voltages are known as *security standards* (National Grid PLC 2004). In the UK, supply companies are subject to regulatory penalties if they fail to meet these standards. Other countries apply similar controls to ensure electrical supply meets the required standards. Exceeding maximum voltages can cause damage to equipment, while failing to meet minimum voltages can cause “brown outs” or even black outs, where consumers observe interruptions in supply.

A substation serves as a *node* within the power distribution network where the circuit voltage can be measured and regulated in order to meet these security standards. The circuits leaving the substation, called *branches*, extend either to consumers or to other substations (for further regulation). Within this network configuration, the minimum circuit voltage limit (as measured at the node) must be respected to ensure the supply of electricity to customers is adequate and, as the branches may extend to other substations, attaining the minimum voltage will help guarantee the stability of the power distribution system as a whole. The maximum circuit voltage limit must be respected in order to avoid driving too much current through consumer appliances, and to avoid the breakdown of the electrical insulation protecting power distribution components: a major safety hazard, due to the high voltages exposed.

At the Grendon substation, the ideal output circuit voltage is $132kV$ and the security standards specify a tolerance of 5%: the voltage must lie in the range $[125.4, 138.6]kV$. For convenience, the measurement of circuit voltage is normalised to be measured in *per-unit volts* by dividing the actual voltage by the ideal voltage. Hence, in per-unit volts, the output voltage must lie in the range $[0.95, 1.05]$. The lower bound of 0.95 is a hard constraint. That is, the voltage must *never* fall below this limit under normal operations. The upper bound of 1.05 is desirable but not *as* critical: ideally, the system should be maintained at a voltage below this level, but an output voltage of up to 1.1 is acceptable for periods of up to 10 minutes (electrical insulation can withstand over-voltage for this length of time before overheating and breaking down).

Without taking any action to maintain the voltage at an appropriate level, the changes in customer demand would cause the circuit voltage to vary dramatically, and to lie well outside these ideal levels at all but the point of mean demand. Increased outgoing energy demand causes the circuit voltage to decrease and decreased demand causes the voltage to increase. The voltage at the substation is therefore subject to continuous fluctuating change. The Grendon substation employs two of the available methods for controlling the circuit voltage (Wood & Wollenberg 1996) in response to this demand change. First, the substation has banks of mechanically-switched capacitors (MSCs) which allow response to large demand change: these can be switched on or off to increase/decrease the voltage on the output circuit. Second, for finer grained control of the output voltage, banks of multi-tap transformers are used, and changing of the tap settings on these controls the ratio between the input voltage to the substation and the output circuit voltage. At Grendon, each transformer has 30 tap positions, allowing a finer level of control than the on–off behaviour of the MSCs.

As shown in Figure 1 (taken from (National Grid PLC 2007)) the electricity demand from consumers in the UK follows a roughly predictable pattern throughout the day, although this will vary slightly on a day-to-day basis depending on temperature and peoples' activities. The graph also shows that there are significant seasonal changes: in the UK, winter power usage is higher than in the summer due to the need for heating. In contrast, in some countries summer power usage can be higher due to the demand for air conditioning. Further extraneous factors, not shown on the graph, can affect the voltage in the circuit. Line losses, either incoming lines from power stations or outgoing lines to customers, can decrease or increase (respectively) the voltage on the circuit. Faults occur rarely, but the costs incurred in terms of potential damage to equipment when these faults do occur is potentially very high. Whilst faults are unpredictable, they are more likely under certain conditions, such as during adverse weather. In addition to these unpredictable events, planned maintenance operations occur, during which substation components are taken offline, input or output lines are disconnected, or nearby substations are taken out of service. Clearly, these operations dramatically affect the circuit voltages at the substation and appropriate actions must be taken.

At present, the voltages in the circuits are controlled in two phases. *Voltage targets* are set by hand for different times of day and then a *reactive* controller is used that responds to the current state of the circuit in order to maintain the supply at the prescribed voltage. The reactive controller interfaces with the transformers and MSCs, controlling their behaviour. The transformers in use are *auto-transformers*:

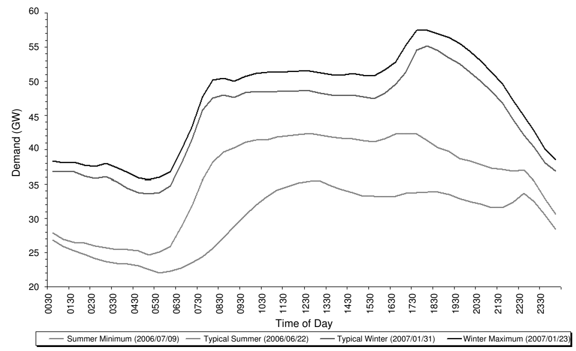

Figure 1: Model of Variation in Demand of Electricity Consumers in the United Kingdom Over 24 Hours (Copyright (National Grid PLC 2007))

the transformers automatically change tap settings to cater for small changes in circuit voltage. For larger, or more rapid changes in voltage, the reactive system will switch on or off mechanically-switched capacitors (MSCs) through the control of isolation circuit breakers: connecting an MSC to a circuit increases its voltage, and provides a larger change than a single transformer step, and in less time. Thus, the reactive control system acts analogously to a thermostatic control on a domestic heating system, while the voltage targets are analogous to a planned profile of settings for the thermostat over a day.

The reactive system could hypothetically be configured to maintain a voltage target of 1.0 per-unit volts throughout the day, performing as many transformer step alterations or capacitor switching operations as were necessary to support this. However, doing so would result in substantial wear-and-tear on components, and hence shorten asset life: both transformers and MSCs are mechanical systems with moving parts. Hence, to strike a balance between wear on components and the supply voltage level, the voltage targets differ at different times of the day. Another important reason to set targets other than 1.0 per-unit volts during the day is that large anticipated changes in demand might be best accommodated by buffering the change, setting the voltage target in the preceding period to a level that keeps the anticipated change within the required range. The targets represent the voltage that the reactive system should aim to achieve at each specified time. The reactive system will follow these targets, taking appropriate action to keep the voltage within 2% of the target voltage for any given time point and targets are set so that the circuit output voltage as a whole will lie in the range $[0.95, 1.05]$ per-unit-volts under nominal conditions. A well-crafted set of voltage targets for a day will take advantage of this acceptable voltage range to give an adequate supply, whilst reducing the component wear when compared to simply using a voltage target of 1.0.

3 Modelling in PDDL

Before considering how we model the problem in PDDL, it is worth considering whether planning is even an appropriate technology for tackling this problem. Although the engineers themselves speak of "planning voltage targets", it is not clear that this is the same sense of 'planning' as is generally considered in the AI Planning research community. Indeed, at one level, it may appear that this is a resource scheduling problem. Furthermore, the uncertainty and integration of the behaviour with real-time control might appear to militate against the use of classical planning models and technology. However, an examination of the current practice shows that the engineers rely on the reactive controllers to manage the uncertainty and they base their planning on predicted customer demand without concern for the stochastic behaviour it will exhibit in practice. An important motivation for exploiting planning in this setting is the tight integration with broader operational decision making, such as maintenance planning. For reasons of space, however, our focus in this paper will be on modelling the voltage control problem, ignoring maintenance, in PDDL and the use of the resulting model with planners.

3.1 Basic Modelling Approach

An interesting decision is at what level actions should be captured in the model. The power engineers can decide to raise or lower the voltage target levels, so one possibility is to model these actions and to plan a series of such activities across the day. However, the purpose of the plan is to achieve the voltage targets while minimising equipment wear and this requires consideration of the actions that it is anticipated will be taken by the reactive controller in response to the changing targets and demand and supply levels. We found that this is made easier if the planner works with a model of the control actions themselves. That is, the planner is used to construct a plan that reflects how the reactive controller could respond to the goals and demand and supply *if it were able to plan*. We then take this plan and use it to generate the voltage targets that would actually achieve the behaviour we want.

Across the day, four factors affect the voltage on a circuit. Including their impact (based on figures from full reactive power models) these are:

1. Changing transformer settings. Changing a transformer up or down by one step increases or decreases the voltage on the circuit to which it is attached by 0.008 per-unit-volts.
2. Changing MSC settings. Switching in or out an MSC increases or decreases the voltage on a circuit to which it is attached by by 0.02 per-unit-volts.
3. The fluctuating demand. A demand increase or decrease of 14GW (UK level) decreases or increases the circuit voltage (at Grendon) by 0.01 per-unit-volts.
4. Line-loss faults. A line-loss fault can increase or decrease the circuit voltage by 0.02 per-unit-volts.

The first two of these are controllable, and we model each directly with PDDL actions to increase or decrease circuit voltage by the requisite amount. These are also bound by the restrictions of the equipment: a transformer has 30 variable steps, end-to-end, while an MSC can only be on or off. To model the fluctuating demand, and hence the passage of time, we have two actions:

- 'Check circuit', which can only be applied if the voltage on that circuit is within the range $[0.95, 1.05]$, and which then applies the voltage change of the next 30 minutes
- 'Advance time', which requires all circuits to be checked and resets the number checked to 0.

These actions maintain an additional house-keeping proposition 'locked' — deleted by 'Check Circuit' and added by 'Advance Time' — required as a precondition of the transformer and MSC actions.

The result is a numeric, non-temporal planning model. A problem in this domain can then be specified in terms of the changes in demand across each half-hour window, the initial component configurations and the goal to reach the end of the day. The plan produced will then perform the necessary transformer and MSC stepping actions to achieve the goal. This plan, in terms of control actions, indirectly leads to a series of voltage targets for the day: by executing the plan, using the plan validator (Howey, Long, & Fox 2004), the values of the circuit voltage levels form the voltage targets for the period modelled in the plan.

3.2 Capturing Costs

So far, we have shown how the PDDL model captures the necessary conditions that *must* be achieved in any solution. Beyond this, we augment the model with cost information, maintaining a numeric value — *total-cost* — used to capture the notional cost of the plan. This cost is updated by actions as follows:

- Stepping a transformer has a fixed cost of 0.1;
- Switching in/out an MSC has a fixed cost of 0.1;
- Checking a circuit has one of three possible costs:
  - 'Check circuit, low', if the circuit voltage is in the range$[0.95, 0.97)$, has a cost of 0.5. If a voltage-lowering line-loss fault occurred at this point, the voltage would fall below the lower bound of 0.95.
  - 'Check circuit, nominal', for the range $[0.97, 1.03]$, has zero cost;
  - 'Check circuit, high', for the range $(1.03, 1.05]$ has a cost of 0.2. If a voltage-increasing line-loss fault occurred at this point, the voltage would exceed the upper bound of 1.05. As mentioned in Section 2, it is acceptable for the voltage to reach 1.1 for short periods, by which point remedial action can be taken if the fault persists.

4 Experiment 1 — Initial Investigation

Using a problem file for a typical winter's day at Grendon (see Figure 1), we explored the performance of several numeric (optionally temporal) planners. Three planners were able to solve the problem[2]: Metric-FF (Hoffmann 2003), MIPS-XXL (Stefan Edelkamp 2006) and CRIKEY 3 (Coles *et al.* 2008b). All of these planners are underpinned by

[2]Neither SGPlan 5 or 4, LPG-td, LPG-1.2, or Sapa were able to solve the problem.

the approach taken in Metric-FF. Since the problem is non-temporal, the behaviour of CRIKEY 3 is very similar to that of Metric-FF. In the absence of temporal actions, the TRPG heuristic is equivalent to the RPG heuristic, and the plan produced is sequential (in terms of instantaneous atemporal actions). Hence, the two planners behave almost equivalently.

MIPS-XXL, an optimising planner, works as a wrapper around Metric-FF, augmenting the problem and domain with additional information. Here, in the absence of PDDL3 language features, no additional information is added. To produce an optimised plan, MIPS-XXL begins by passing this unmodified plan to Metric-FF, giving an upper-bound on plan quality. It then cycles, adding a goal at each iteration that the total-cost be less than last time. In the case of the problem in question, the first solution has a total cost of 30.2, much of it accrued by advancing time when the voltage is within the fault risk boundaries. Hence, after the first iteration, the goal `(< (total-cost) 30.2)` is added, before invoking Metric-FF once again. The injection of this metric goal poses substantial difficulties to search for Metric-FF. Under the metric RPG heuristic, the negative numeric effects of actions are ignored, so that when considering a goal for a value to be *below* a specified level actions that *increase* that value are ignored. Thus, in this example, the heuristic ignores action costs and will happily promote entry into any state where the total cost is less than 30.2. This leads to a situation where a near-goal state is reached (where the accrued cost is just below 30.2) but the application of a necessary further action takes the cost over this threshold. At this point, the heuristic detects the goal is unreachable and declares the state to be a dead-end. Unfortunately, these dead ends are detected too late in search and Metric-FF resorts to an expensive and often ineffective best-first search.

5 Targeted Goal Injection

Based on these results, we can observe that a key barrier faces the application of AI planning in this setting: the ability of the planner to effectively consider the plan quality metric. MIPS-XXL was the only optimising planner able to solve problems in this domain but its ability to optimise these plans is limited in practice due to difficulties with the heuristic. Examination of the plans produced by Metric-FF reveals that a large source of cost is setting voltage targets below 0.97 per-unit-volts. Although not optimal, Metric-FF attempts to minimise the number of actions used, so generally the number of transformer and MSC actions is reasonable. However, only one action is used to check a circuit irrespective of its voltage level. Hence, we turned our attention to the impact of the 'Check Circuit, Low' action upon plan quality. One possibility is to remove this action from the domain entirely, but due to component limitations or potential demand it is impossible to guarantee that the voltage can always be set to at least 0.97. Instead, we adopt an iterative approach, similar to MIPS-XXL, but with the addition of domain-specific goals to prevent circuit checking with a low voltage at specific times.

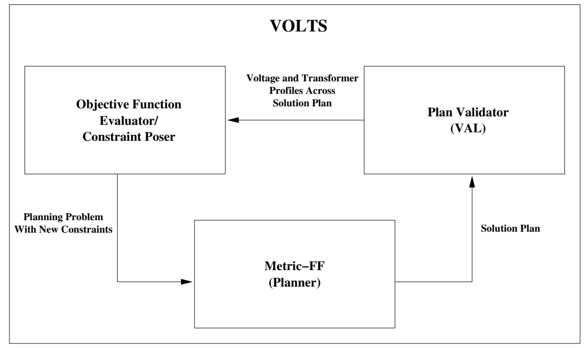

Figure 2: Overview of the VOLTS System

5.1 Overview

An overview of our system for this problem, VOLTS (Variable On-line Transformer Scheduling), is shown in Figure 2. The input to the system is the domain model and the problem described in Section 3. Initially the only requirements are that the voltage remains between 0.95 and 1.05 per-unit volts, the *hard* constraints of the problem. The *Objective-Function/Constraint Poser* module passes this problem straight to Metric-FF, which generates a solution. This solution is fed into a plan validator, VAL (Howey, Long, & Fox 2004), which verifies plan correctness and computes the voltage trace across the day given this plan. The output voltage trace is then fed back to the *Objective-Function/Constraint Poser* which adds extra constraints to the problem before returning the problem to the planner. This process iterates until a termination criterion is satisfied (a plan is found that is robust to faults).

5.2 Injecting Constraints

Metric-FF solves the initial problem efficiently, in under 20 seconds on a modest desktop computer (3GHz Pentium CPU, 1GB of RAM). This provides an *adequate* plan: one that ensures that the voltage is maintained within the range 0.95 to 1.05 per-unit volts. The constraints selected for injection at each subsequent iteration are generated automatically from the solution plan to the previous iteration. We use the plan validation tool, VAL, to produce a trace of the voltages across the plan that can be used to identify the time at which the voltage is the lowest. This is the point at which the plan is at its least secure: recall that an outgoing line loss causing the voltage to drop below 0.95 per-unit volts is the most serious type of fault. If the voltage at its lowest point is less than 0.97 per-unit volts then, were an outgoing line fault to occur at this point (reducing the voltage by 0.02 per-unit volts), the voltage could go outside the required range. We may choose to accept this risk (as it is small, despite the cost being high) if avoiding it were to involve many transformer or MSC switches, or if there simply exists no solution that avoids it. Remaining 0.02 per-unit volts above the minimum boundary line is a *soft* constraint, to be traded off against others in the objective function. The iterating behaviour gives a useful benefit: during each iteration the aim is to make the

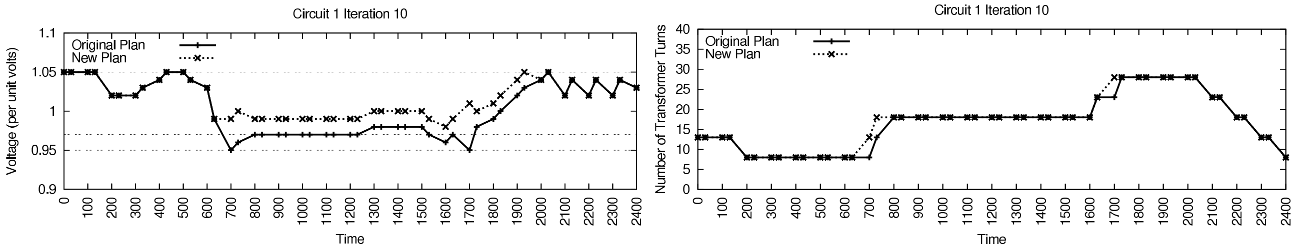

(a) Voltages

(b) Transformer Steps

Figure 3: Comparing the Initial Plan to the Final Plan after Completion of the Cycling Process.

plan more *secure*. However a more secure plan may yield a worse value of the overall objective function, as it might require more transformer steps and MSC switches. We can maintain the best plan found so far which might not be the plan we discover on the final iteration. Often, however, it is the case that the final plan is the one with the best total cost value, as fault tolerance is a factor in its calculation (see Section 3.2).

We add each *soft* constraint to the problem by treating it as a *hard* goal. The strength of the planner is in quickly finding a solution that satisfies the goals, not in optimising the objective function. The optimisation is achieved by the process of increasingly constrained iteration. We cannot pose all the soft constraints as hard constraints at the outset, since it might not be possible to satisfy them all. For instance, if demand is on the borderline of what the system is capable of delivering at a certain time of day, it may not be possible to have the voltage above 0.97 per-unit volts. Further, introducing soft constraints as hard constraints introduces a trade-off. That is, the solution may require more transformer stepping or MSC switches operations to be carried out. Starting with an easier problem and then injecting harder constraints has the benefit that the planner can find an *adequate* solution to the problem quickly which can be gradually refined to improve the quality.

6 Experiment 2 — VOLTS

In this section we demonstrate the effectiveness of VOLTS and discuss the time it takes to solve problems. Figure 3a shows the initial and final voltage targets produced by VOLTS running over the course of a typical winter's day. The initial voltage targets (shown by the solid line) do not represent a secure plan, that is, the voltage would drop out of the desired range were a fault to occur. However, by the time the iterative refinement process is complete the final plan is secure to faults on outgoing lines at any time. Our final plan has several important desirable properties that overcome the weaknesses of the current system.

First, the voltage targets have been generated automatically, without the need for human intervention. Second, engineers face the problem that the bounds imposed on the voltage targets throughout the day encourage a *boundary-hugging* behaviour. Observing the actual values of the voltage targets we can see that VOLTS does not suffer from the boundary-hugging phenomenon encountered when relying on the reactive system to control the voltage. The decision about the time to change target voltages is based on the best strategy to minimise component wear-and-tear and anticipation of future expected demand, whereas the reactive system behaves myopically, waiting until the voltage reaches the extreme of the range before reacting. VOLTS maintains the voltage targets at appropriate levels without hugging the boundary, intelligently choosing the times for voltage targets to avoid unnecessary auto-transformer activity.

Perhaps the most impressive result in generating these voltage targets is that VOLTS has successfully built a secure plan — that is, one that is tolerant to faults — that uses the same number of transformer steps as the initial plan. This means that by reasoning about minimising the objective function through the addition of constraints, we have successfully made a secure plan that has no additional execution costs when compared to the original plan. Figure 3b shows the transformer turns for a transformer, and their timings, that VOLTS considers necessary in order to achieve the voltage targets throughout the day. These are not prescriptive — the reactive system is not controlled directly by this plan — but they do give an indication that these voltage targets can be achieved in practice, and indeed could be used as guidance if desired. Observing this figure it is clear that by shifting the timing of voltage targets (i.e. shifting predicted transformer stepping actions) it is possible to improve security of the voltage targets without increasing the number of transformer steps, something that would remain difficult, even for a human operator, in the general case.

VOLTS produces the original plan in 12.5 seconds on a modest desktop computer. This is a very encouraging result: currently, in order to avoid over-sensitivity to short-lived spikes in voltage, the underlying reactive system only responds to voltage changes after 10 seconds have elapsed. VOLTS is able to produce a plan for a 24 hour period in 12.5 seconds, so were an unexpected fault to occur later during the day it could be possible to produce some initial voltage targets within 10 seconds (planning over a shorter horizon requires less time) before the time at which the reactive system would kick in. This has the advantage that the reactive system can then react according to a revised set of voltage targets. The initial plan often does not produce voltage targets that are as successful at minimising cost as the final

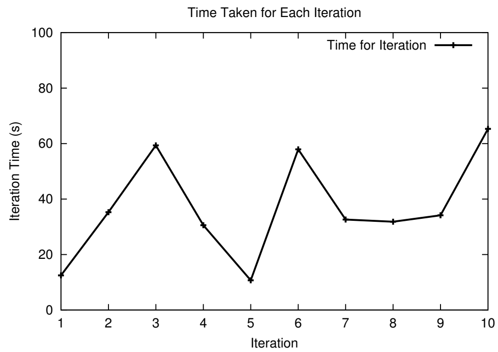

(a) Time Taken for Each Iteration.

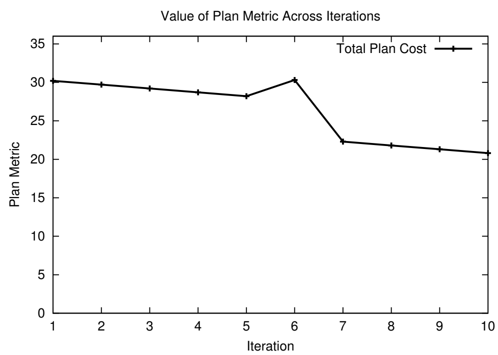

(b) Improvement in the Quality Metric.

Figure 4: Time and Quality at Successive Iterations

plan: iteration is required to ensure security and to optimise the objective function. However, the initial plan gives an estimation of the best course of action given the predicted future demand which is often better than the greedy approach used by the recovering reactive system. Overall, the difficulty of finding a plan depends on the combination of constraints to be solved. Figure 4a shows the time taken for each of the 10 iterations used to find a secure solution to the problem. It is clear that the time taken to solve the problem does not increase monotonically with the number of constraints injected. Rather, the interaction between the constraints can make some more constrained problems easier to solve. The total time taken to solve the problem, summed across all iterations, is 370 seconds. Also noteworthy is that no single iteration takes longer than 65 seconds to complete.

7 Conclusions and Discussion

In this paper we have shown how power substation voltage control can be performed using AI planning. Although no planner available was able to find high-quality plans directly, in this cost-critical domain, a planner serves as a powerful kernel within VOLTS and, indeed, it allowed rapid development of an appropriate system.

There are several interesting directions for future research. We intend to extend the model to include several substations within a region. At present, we consider only a single substation. Although this is useful, nearby stations interact with their neighbours and considering this interaction can improve system reliability. By way of example, if the circuit voltage is in the range $[0.95, 0.97)$ at Grendon, and the circuit voltage at a downstream substation is over 1.00, then decreasing the voltage there will increase the voltage at Grendon. The challenge is to find a set of interacting voltage targets that minimise predicted operational cost and the impact of faults. Due to the strong numeric interaction in this problem, we intend to investigate applying the LP-RPG heuristic (Coles *et al.* 2008a), and the role of the quality metric in the LP during RPG solution extraction.

We also intend to extend the model to consider the management of hydro-electric and wind power. At present, we are considering power-distribution, meeting the demand of consumers by managing components. In managing power distribution, the challenge is to maximise profits through production of the most electricity when it is likely that electricity will cost the most. For example, generation with wind-power must be considered alongside the wind-speed forecast to determine how many turbines to have running, how many to have on standby, and how many to have turned off. In effect, the components are managed to meet economic and weather forecasts, rather than a demand forecast.

References

Coles, A. I.; Fox, M.; Long, D.; and Smith, A. J. 2008a. A Hybrid Relaxed Planning Graph–LP Heuristic for Numeric Planning Domains. In *Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling (ICAPS 2008)*.

Coles, A. I.; Fox, M.; Long, D.; and Smith, A. J. 2008b. Planning with Problems Requiring Temporal Coordination. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI 08)*.

Fox, M., and Long, D. 2003. PDDL2.1: An Extension of PDDL for Expressing Temporal Planning Domains. *Journal of Artificial Intelligence Research* 20:61–124.

Hoffmann, J. 2003. The Metric-FF Planning System: Translating "Ignoring Delete Lists" to Numeric State Variables. *Journal of Artificial Intelligence Research* 20:291–341.

Howey, R.; Long, D.; and Fox, M. 2004. VAL: Automatic Plan Validation, Continuous Effects and Mixed Initiative Planning using PDDL. In *The 16th IEEE International Conference on Tools with Artificial Intelligence*, 294–301.

National Grid PLC. 2004. GB Security and Quality of Supply Standard. https://www.nationalgrid.com/uk/Electricity/Codes/gbsqsscode/DocLibrary/.

National Grid PLC. 2007. National Grid Electricity Transmission, 2007 Seven Year Statement. http://www.nationalgrid.com/uk/Electricity/SYS/.

Stefan Edelkamp. 2006. MIPS-XXL. Booklet of the Fifth International Planning Competition.

Wood, A. J., and Wollenberg, B. F. 1996. *Power Generation, Operation and Control*. Wiley.