Expansion of Electricity Distribution Networks Under Uncertainty

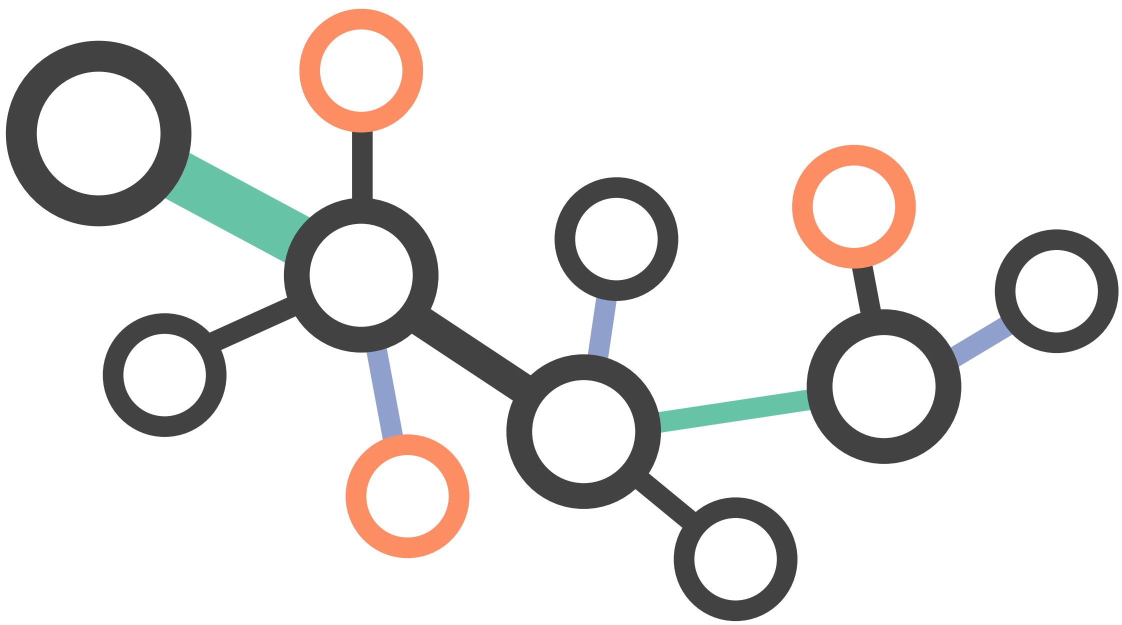

Author
Alex de Beer

Project Partner
Anton Aish

Supervisors
Prof. Andy Philpott
Dr. Tony Downward

Part IV Project Report
Deptartment of Engineering Science

Faculty of Engineering
The University of Auckland

Abstract

The electricity sector is experiencing growing uncertainty in future demand for electricity. In addition, non-wires alternatives, such as batteries and load management systems, are increasingly being considered as a means of deferring the need for large, expensive line upgrades and facilitating the development of flexible, decentralised networks. High costs are associated with upgrading and maintaining networks; there is a need, therefore, for effective planning mechanisms in this new and uncertain environment. This report describes the development of GENE.jl, a Julia package that uses stochastic programming to find least-cost expansion plans for distribution networks that maintain desired levels of network reliability, while accounting for uncertainty in future demand explicitly. GENE leverages an additional Julia package, JuDGE.jl, to solve the resulting models using Dantzig-Wolfe decomposition. GENE represents a significant improvement on previous modelling that has been conducted in this area, containing improved methods for modelling batteries and network reliability, improved visualisation capabilities, and the ability to illustrate the value of accounting for uncertainty in future demand explicitly, rather than planning to a deterministic demand forecast. It is believed that GENE is capable of providing useful insights to decision-makers in the electricity sector. Several case studies have been undertaken, using data on a distribution network in Auckland, New Zealand, to illustrate the capabilities of GENE. In addition, computational testing has been conducted to determine the effectiveness of solving a network expansion problem using Dantzig-Wolfe decomposition, relative to solving the equivalent large-scale mixed-integer program. Several factors that influence the relative performance of each type of model have been identified and discussed.

Declaration of Contribution

Initial work on the model formulation was done by both Anton and myself. I built on this work to add batteries that could be moved at a cost, and the quantification of reliability using SAIDI and chance constraints.

Other significant contributions I made included the design of the spreadsheet-based input system for GENE, the design of the majority of the structures to parse the input data into a suitable form for the optimisation, and the code to formulate the deterministic, highest-growth model corresponding to a given stochastic model. No code for the model formulation or any of these components of GENE was available from previous years (though most of the data for the test network used were available from a previous part IV project by Dylan Dong).

I made a small contribution during the initial development of the visualisation tool, but this was almost entirely the work of Anton.

Acknowledgements

I would like to thank Professor Andy Philpott and Dr. Tony Downward for their suggestions and guidance throughout this project. I enjoyed our discussions over the course of the year and have learnt a great deal.

I would also like to thank my project partner, Anton Aish. It has been a pleasure working together.

Contents

List of Figures

List of Tables

Acronyms

AC	alternating current
DC	direct current
GENE	generalised electricity network expansion
JuDGE	Julia decomposition for generalised expansion
MIP	mixed integer program
NWA	non-wires alternative
RMP	restricted master problem
SAIDI	system average interruption duration index
SP	subproblem
ZSS	zone substation

1 Introduction

Electricity providers are facing unprecedented levels of uncertainty in future demand for electricity (Vector, 2021). In New Zealand, a primary driver of this uncertainty is the response to climate change, including electrification of the transport system.

In addition to conventional line upgrades, electricity providers are increasingly investigating the benefits of decentralised network design, through the use of non-wires alternatives (NWAs) including batteries and methods for load management. Such investments improve the ability of the network to operate reliably during a line failure, and are flexible, helping to defer the need for more expensive investments until some uncertainty has been resolved and there is an improved understanding of whether they are required.

Given the significant uncertainty in future demand for electricity and the new technologies available for expansion, the problem of determining how an electricity provider should upgrade their infrastructure at least cost, while ensuring a reliable service, is difficult. Upgrading and maintaining electricity networks is expensive; Vector, New Zealand's largest provider of electricity, has spent $2.3 billion on its Auckland network over the past decade (Vector, 2021). Effective planning, therefore, has the potential to result in significant savings.

Nevertheless, existing expansion planning procedures for electricity distribution networks are often manual in nature, and consider only a handful of possible plans (Singh, 2006). Given the large number of possibilities, it is unlikely that the best possible plan is identified, and there is no way to measure the goodness of a selected plan relative to the best possible plan. Furthermore, planning is often done to a deterministic, highest-growth demand forecast; thus, it fails to explicitly account for the possibility of future demand being lower than anticipated, and, therefore, the value of planning that is flexible in the face of uncertainty.

The central aim of this project is to develop a software package that uses stochastic programming to find, for a given electricity distribution network, a least cost expansion plan that maintains a desired level of reliability, while explicitly accounting for uncertainty in future demand. Such a tool could provide valuable insights to electricity providers.

The remainder of this report is laid out as follows. Chapter 2 presents background information on electricity distribution networks and models of capacity expansion under uncertainty. Chapter 3 describes the mathematical model for distribution network expansion that was derived during this project, and introduces the corresponding software package, GENE, that was developed. Chapter 4 outlines several case studies that demonstrate the functionality of GENE, and presents computational results that demonstrate the scale of problems GENE is capable of solving. Finally, Chapter 5 presents the main conclusions drawn from this project, and improvements that could be made.

2 | Background

This chapter presents background information on expansion planning for electricity networks. The key components of electric power systems are described, and the idea of network reliability is introduced. Subsequently, a description of how operations research methodology has been applied to capacity expansion problems, with a focus on electricity distribution, is presented. Finally, several key studies this project aims to build on are discussed.

2.1 Electric Power Systems

The electric power system is composed of several key components. Once generated, power is transported over long distances, by means of a transmission network, to a set of subtransmission networks. Each subtransmission network provides the link between the transmission network and a set of distribution networks, which supply power to customers (Singh, 2006).

The key characteristic that differentiates each of these components of the system is the voltage at which they operate. The electric power system uses alternating current (AC), rather than direct current (DC). This means that transformers can be used to vary the voltages over the system, enabling the simultaneous use of an efficient, high-voltage transmission network, and safe, low-voltage distribution networks (Kiszka and Wozabal, 2021). Distribution networks are the focus of this project, though it is possible to generalise the methodology used to other parts of the electric power system as well.

2.1.1 Distribution Networks

A typical distribution network is composed of a number of components. Power enters the network through one or more *zone substations*, which form the connections between distribution and subtransmission networks. It then travels, through a series of *lines* and *junctions*, to *distribution substations*, where it is transferred to individual consumers.

At a given moment in time, a subset of the lines of a distribution network are active and distributing power, while the remainder are inactive. The activity of each line within the network is typically controlled by a pair of *switches*; one at either end of the line. A key characteristic of the networks studied in this project is that they must operate in a *radial* configuration; this means that there must always be a single path through active lines between any two locations in the network. Figure 2.1 illustrates this type of configuration. A subset of the lines are active during routine use of the network, forming the standard operating configuration. Within this subset, a line that is connected to a zone substation,

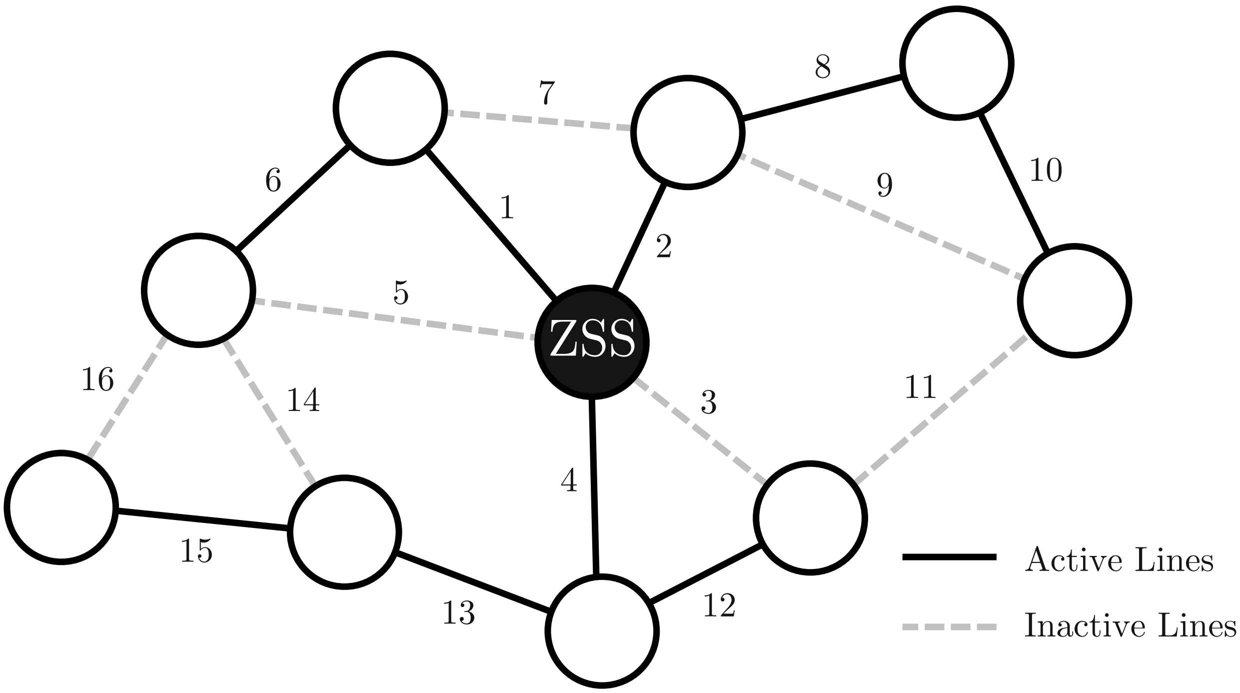

Figure 2.1. An example of a small distribution network operating in a radial configuration. The zone substation (ZSS) is at the centre of the network and supplies power to all other locations, which represent junctions or distribution substations.

and the lines downstream of it, collectively form what is referred to throughout this report as a *feeder*. The remainder of the lines are typically inactive and serve to reroute power in the event of a line within a feeder failing; these are referred to as *backfeeders*. The feeders of the network in Figure 2.1 are $\{1, 6\}$, $\{2, 8, 10\}$, and $\{4, 12, 13, 15\}$. The backfeeders are $\{3, 5, 7, 9, 11, 14, 16\}$.

2.1.2 Survivable Network Design

An important consideration when designing a distribution network is its *survivability*; this describes whether the network has sufficient excess capacity to reroute flow to customers in the event of one or more lines failing (Singh, Philpott, and Wood, 2004). Such outages may be planned or unplanned, and can last for anywhere between several hours and several weeks. In these situations, the switches on either side of the failed lines must be opened, to isolate them from the network, and the switches associated with one or more backfeeders must be closed, such that flow is rerouted through an alternative sequence of lines to affected customers while ensuring that a radial operating configuration is maintained. This is illustrated in Figure 2.2. The ideal network is $N-1$ *secure*; this means that the configuration of the network can be adapted, such that it continues to meet all demand, in the event of any one of its N lines failing.

Among the researchers to investigate the modelling of network survivability were Singh,

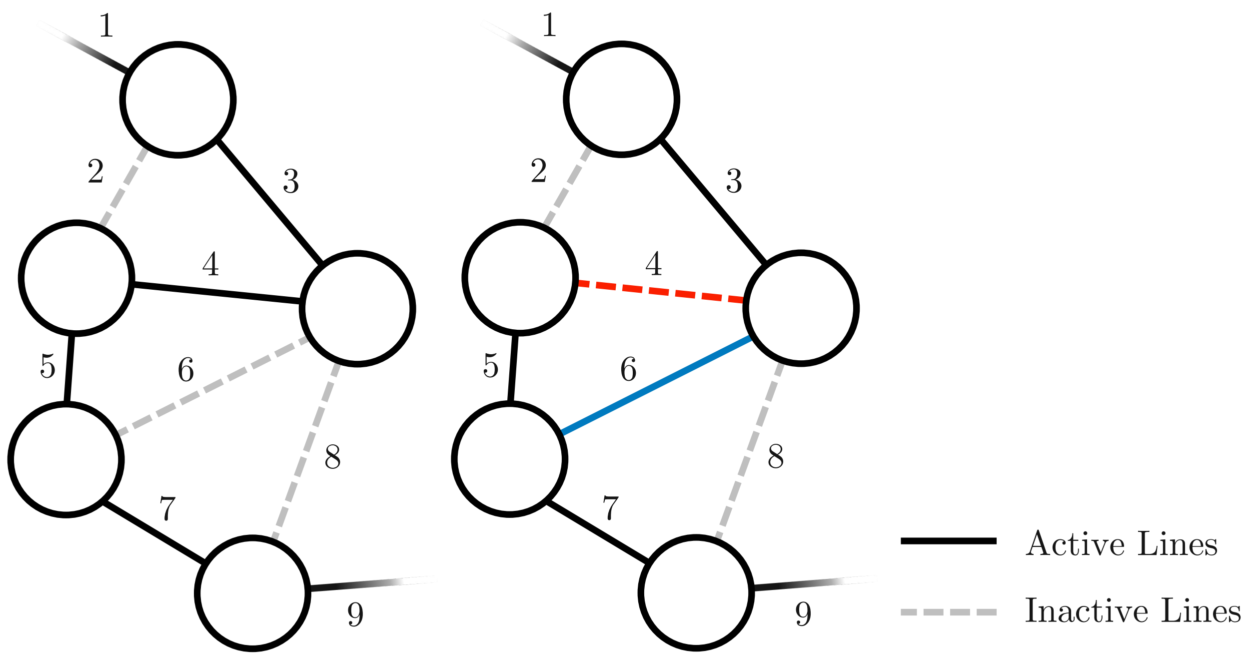

Figure 2.2. A section of a distribution network operating in a radial configuration before (left) and after (right) the failure of line 4, which forms part of a feeder. Line 6, a backfeeder, becomes active so a radial configuration is maintained.

Philpott, and Wood (2004), who derived a two-stage stochastic programming formulation for the design of survivable networks over a time horizon on the order of one year. The first stage decisions within this formulation determine where additional capacity should be allocated within the network to improve its survivability, while the recourse decisions determine how the network should be configured in response to various line failures. Unmet demand is penalised with a cost; the magnitude of this cost influences the reliability of the resulting network design. In the limiting case, as the cost per unit of unmet demand approaches infinity and all lines of the network are considered as potential candidates to fail, the problem becomes one of determining how to design a least-cost $N-1$ secure network.

2.2 Capacity Expansion Problems

Capacity expansion problems have been studied extensively, through the lens of operations research, since the 1950s. Such problems are typically concerned with determining the times, sizes, types, and locations of expansions for a system, taking into account additional complicating factors such as economies of scale and uncertainty in future demand, such that the expected discounted cost of expansion over a period of time is minimised (Luss, 1982). Early work on capacity expansion problems was largely theoretical in nature; among the first to develop case studies of expansion planning for complex, real-world systems was Manne (1967), who investigated capacity planning for the aluminium, caustic soda, cement, and nitrogenous fertiliser industries in India. Since then, case studies of expansion planning have been conducted in a wide range of areas, including transportation, water distribution, telecommunications, and, importantly, electric power distribution.

A key drawback Manne identified when describing his early case studies was the way in which the future was treated; despite the assumptions underpinning much of the modelling being highly uncertain, the future was treated as though it was known with certainty, to simplify the analysis. Since the introduction of stochastic programming by Dantzig (1955), much effort has been put into modelling uncertainty explicitly and examining the differences between stochastic and deterministic models.

2.2.1 Uncertainty and Stochastic Programming

Real-world capacity expansion problems invariably contain some form of uncertainty. In the context of electricity distribution, future demand is typically the most important stochastic parameter to consider, though other sources of uncertainty, including the future cost and availability of upgrades, are present as well.

Common practice when dealing with uncertainty in mathematical programming is to form the *expected value problem*, by replacing any stochastic parameters with their expectations. The associated deterministic model can be solved using the simplex method, and sensitivity analysis can be conducted on the optimal solution to ascertain its stability with respect to changes in the stochastic parameters. A more complex approach involves constructing a range of deterministic scenarios, each with varying values for the stochastic parameters, and using common decisions within the optimal solutions to generate a policy (Birge and Louveaux, 1997); this is referred to as *scenario analysis*. There are, however, issues with these approaches; the structure of the solutions to deterministic models are often fundamentally different to those found using stochastic programming, where uncertainty is incorporated explicitly. This is, in part, because a deterministic model does not place any value on the option of delaying an investment until uncertainty is revealed, nor an investment having flexibility with regard to potential future outcomes (Wallace, 2000).

The differences between deterministic and stochastic programming are similar to those between *discounted cash flow* analysis and *real options* methodology, which is seeing increased use in investment planning in energy systems (Schachter and Mancarella, 2016). Criticisms of discounted cash flow analysis include that it uses a deterministic view of the world, and that it assumes a static policy, in which all decisions are made at the start of the planning horizon; this results in options that are adaptable in their ability to respond to changing circumstances being undervalued. By contrast, real options methodology is able to quantify the value of flexibility in the face of uncertainty. It is worth noting, however, that while real options methodology can only be applied to options that have already been established, stochastic programming begins with a set of decisions, and is therefore capable of constructing options (Wallace, 2010).

Typically, the distributions of the uncertain parameters in a stochastic programming problem are continuous; this is certainly true, for example, of future demand for electricity. Apart from in trivial cases, however, the solution methods for such problems require discrete distributions with limited support. A common method to approximate the distributions of

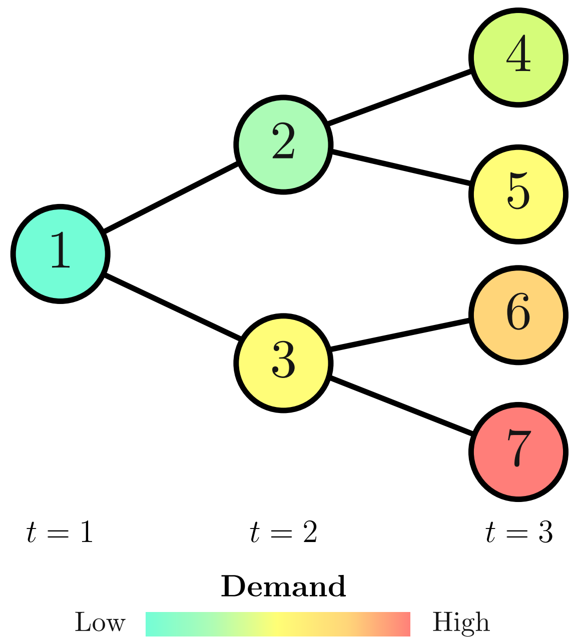

Figure 2.3. A scenario tree containing three stages and four scenarios, with varying levels of demand. Each path from the root node, $n = 1$, to a member of the set of leaf nodes, $\{4, 5, 6, 7\}$, forms a scenario.

a set of stochastic parameters is the use of a *scenario tree* (see e.g. Defourny, Ernst, and Wehenkel, 2011). Figure 2.3 shows an example of a scenario tree, in which the stochastic parameter being modelled is future demand. The scenario tree can be thought of as the aggregation of a number of deterministic scenarios, which are broken into a number of *stages* at which decisions may be made and the values of the stochastic parameters may change. At each stage, scenarios within the tree that are indistinguishable up until that point are grouped together; the scenario tree, therefore, represents the branching process that occurs as uncertainty in the stochastic parameters is revealed. Each node of the tree is assigned a probability; this is the combined probability of the scenarios it is part of occurring.

The scenarios that form the scenario tree approximate the true stochastic process under consideration, so can be expected to have a significant impact on the quality of the final solution to the problem (Kaut and Wallace, 2007). In this sense, the use of a large number of scenarios is desirable. However, there also exists a tradeoff between the size of the scenario tree and the tractability of the resulting stochastic programming model; if too many scenarios are used, the resulting model will become impossible to solve. Therefore, care must be taken when constructing a scenario tree to ensure that the limited set of scenarios used approximates the stochastic parameters well. Several techniques to generate representative sets of scenarios exist, including sampling methods and the matching of moments between the discrete and continuous distributions.

2.2.2 Solution Methods

A stochastic capacity expansion model typically requires integer variables; for example, in the context of electricity distribution, decisions to upgrade lines are typically binary, and the requirement that the network operates in a radial configuration is enforced using a set of binary variables. In theory, it is possible to solve such a model by formulating it as a single, large mixed-integer program (MIP) and solving it using branch and bound; this is often referred to as solving the *deterministic equivalent* problem. This, however, tends to become intractable as the number of scenarios under consideration grows, necessitating a form of decomposition.

A form of decomposition commonly used in stochastic programming is *progressive hedging* (Rockafellar and Wets, 1991), which uses a Lagrangian relaxation to disaggregate the scenarios that make up the problem. The models corresponding to these scenarios are then solved individually, in an iterative process, until an optimal solution to the problem is converged to, in which the decisions made in states of the world common to multiple scenarios are consistent. Progressive hedging has become increasingly popular in recent years, due, in part, to the development of efficient, parallelised versions of the algorithm. However, an alternative form of decomposition that has been shown to be effective in the context of stochastic capacity expansion problems, and can also be parallelised, is *Dantzig-Wolfe decomposition*.

Dantzig-Wolfe decomposition (Dantzig and Wolfe, 1960) splits a large linear program into a number of smaller, more tractable *subproblems* (SPs), coordinated by a *restricted master problem* (RMP). The subproblems and master problem are solved in an iterative process until an optimal solution to the overall problem is converged to.

The use of Dantzig-Wolfe decomposition requires a specific problem structure; a number of sets of constraints must be identified, such that the variables present in the constraints in a given set are not present in any other set. Each of these sets, and associated variables, form the constraints and variables for a subproblem. The remainder of the constraints, which contain variables from multiple subproblems, form the constraints of the master problem. Within the master problem, the decision variables associated with each subproblem are re-written as a convex combination of the extreme points of the subproblem's feasible region. The associated convex multipliers become the master problem decision variables. Solving the master problem, therefore, amounts to finding a least-cost convex combination of extreme points from the subproblems that satisfy the master problem constraints.

Generating a complete enumeration of the extreme points of each subproblem quickly becomes intractable as the size of the problem grows. Instead, the master problem begins with a subset of extreme points from each subproblem, and new points are generated as needed. When the master problem is solved, the optimal values of the dual variables are used to adjust the objective coefficients of each subproblem, such that solving each subproblem amounts to finding a new extreme point with a negative reduced cost. These extreme points are added to the master problem, which is then re-solved; this process is repeated until no extreme points with a negative reduced cost exist, at which point the solution to the master problem is optimal for the entire problem.

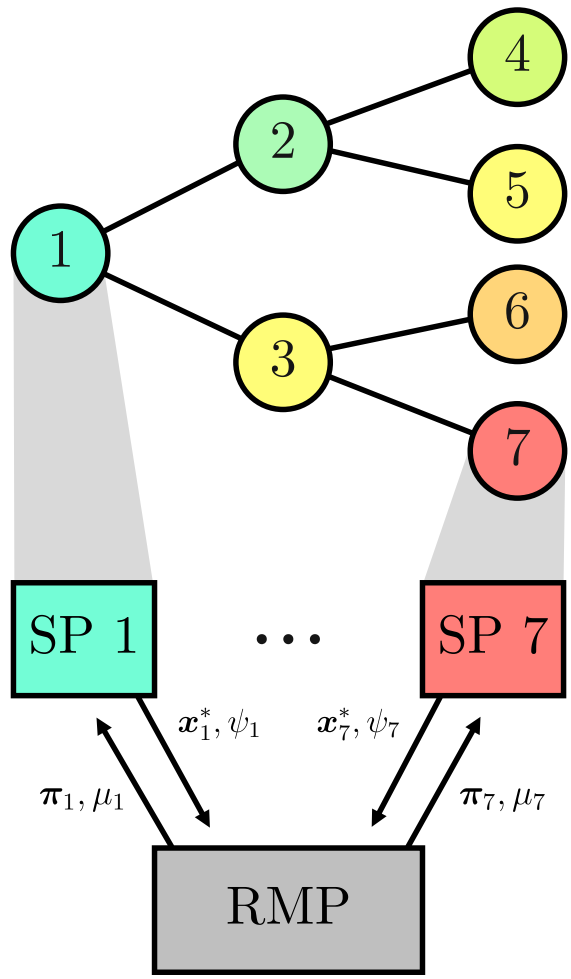

Figure 2.4. An illustration of the decomposition developed by Singh, Philpott, and Wood (2009). The extreme points generated by subproblem n, and associated costs, are denoted using $\boldsymbol{x}_n^*$ and ψ_n respectively. The dual variables associated with the constraints in the master problem corresponding to subproblem n are denoted using $\boldsymbol{\pi}_n, \mu_n$.

A key study that demonstrated the use of Dantzig-Wolfe decomposition in the context of capacity expansion for electricity distribution networks was carried out by Singh, Philpott, and Wood (2009). In contrast to progressive hedging, the decomposition described by Singh, Philpott, and Wood is a *nodal decomposition*; this structure is illustrated in Figure 2.4. Each node in the scenario tree forms a subproblem, and contains the constraints associated with operating the network.

To increase the strength of the decomposition, a *split-variable* reformulation of the problem is used; the variables associated with line upgrades are split into a set of master problem variables that indicate whether each upgrade is purchased at each node in the scenario tree, and a set of subproblem variables that indicate whether each upgrade is being utilised at each node. This gives the master problem greater flexibility; an upgrade does not have to be utilised at a node even if it has been purchased previously.

It is possible for the optimal solution to the master problem to be fractional. If this occurs, the master problem can be solved as a MIP in the hope of obtaining a high quality solution using the extreme points that have been generated, or a branch and price method can be used to generate additional extreme points and converge to the optimal solution (see e.g. Appelgren, 1971). However, Singh, Philpott, and Wood found that, in their own tests using a real-world distribution network, the optimal solution to the master problem tended to be naturally integer, attesting to the strength of the decomposition.

2.3 Previous Modelling Studies

During this project, it is of interest to investigate the use of stochastic programming for capacity planning under uncertainty for Vector, New Zealand's largest distributor of electricity and gas. Several prior studies have considered similar problems for Vector; this project aims to build on these to develop tools for Vector that are useful in the present day.

Singh (2006) investigated distribution network expansion using a section of Vector's network in Penrose as a case study. Singh first considered survivable network design and capacity expansion under uncertainty separately, before combining these concepts into a single model that was solved using Dantzig-Wolfe decomposition. Though Singh described the potential of incorporating distributed generation into this modelling framework, line upgrades were the sole type of expansion modelled.

More recently, van Helsdingen (2021) applied a similar methodology to the Wellsford planning area of Vector's network, developing a model that incorporated elements of survivable network design and capacity expansion under uncertainty. In addition to line upgrades, batteries and load management systems were added to the modelling framework as potential upgrades. van Helsdingen solved the model using Dantzig-Wolfe decomposition, implemented using a Julia package, JuDGE (Downward, Baucke, and Philpott, 2020), and developed a tool to visualise the expansion plans that were generated.

There were, however, some limitations associated with the work of van Helsdingen. Some aspects of the modelling that was conducted were unrealistic; for instance, batteries could be moved between locations in the network once purchased, but no cost was imposed on this. The data for the Warkworth distribution network that were used did not contain information on junctions or distribution substations downstream of those connected directly to the associated subtransmission network; therefore, the network that was modelled was a highly simplified version of the real network.

This project will focus on overcoming these limitations. Emphasis will be placed on modelling batteries and other components of the system in realistic ways, and demonstrating the capabilities of the software developed on a network of realistic size. It is also of interest to extend the work by considering alternative methods for quantifying network reliability (both Singh and van Helsdingen penalised unmet demand with a fixed cost per unit, but other possibilities exist), developing improved visualisation tools, and investigating the differences that exist between solving stochastic and deterministic network expansion models.

3 | Methods

This chapter outlines how the techniques described in Chapter 2 were used to model capacity expansion for distribution networks. Key model components and simplifications are discussed, and the central model formulation that was derived is presented. The corresponding software package that was developed and the network used to illustrate its capabilities are also described.

3.1 Model Components and Simplifications

Distribution networks are complex systems; care must be taken when modelling them to incorporate appropriate detail, such that the model is capable of generating useful insights while remaining computationally tractable. This section details the key considerations made during the modelling process.

3.1.1 Network Representation

A graph was used to represent the distribution network under consideration. A single vertex of the graph was used to represent the zone substation, which was assumed to always remain capable of supplying enough power to the network to meet all demand. All networks studied in this project contained a single zone substation; however, in the case of a network where power enters though multiple zone substations, the supply of these could simply be aggregated together. The remainder of the vertices were used to represent junctions and demand substations. From a modelling perspective, junctions were equivalent to demand substations with no demand, so were represented in this way. The edges of the graph represented the lines connecting each location.

Power losses in the lines of the network were assumed to be negligible. A direct current approximation of the network was used; at each vertex, the amount of power entering was required to be equal to the amount leaving. In reality, the typical distribution network operates using alternating current, but modelling this requires consideration of the reactive power in each line and the voltage angles at each location (Kiszka and Wozabal, 2021). This introduces nonlinearities into the problem, making it significantly more difficult to solve.

3.1.2 Uncertainty

The central source of uncertainty considered when formulating the model was uncertainty in future demand. A scenario tree was used to approximate the distribution of future demand over the planning horizon.

The demands worked with throughout this project were peak demands; if a network has sufficient capacity to operate reliably during periods of peak demand, it is capable of operating reliably during all other periods.

3.1.3 Upgrades

At each node of the scenario tree, decisions could be made to invest in upgrades to allow the network to meet increased demand. Several types of upgrades were modelled; line upgrades, load management, and batteries.

A line upgrade was used to increase the capacities of one or more lines by a set of specified amounts. Only existing lines were considered as candidates for expansion in this project. However, the framework developed allowed for the construction of new lines, by specifying a line with an initial capacity of zero, and a corresponding upgrade with the capacity of the proposed line. A load management upgrade decreased the peak demands at one or more demand locations by a specified set of percentages.

Batteries were modelled as power sources that could be placed at a subset of the demand locations in the network. Each battery was modelled as being able to contribute to meeting peak demand, both during routine operation of the network and in the event of a line failure. It was assumed that there would be sufficient time during off-peak hours for each battery to recharge completely. Also modelled was the ability for a previously purchased battery to move to a new demand location at each stage of each scenario, at a cost. The model allowed for restrictions to be placed on the number of batteries that could be purchased over the time horizon, and the maximum number of batteries that could be placed at each demand location.

All upgrades were assumed to have a lifespan of no less than the time horizon under consideration. An option to enforce a delay between a line upgrade or load management upgrade being purchased and available for use was modelled. Batteries were assumed to be relatively quick to install, so a delay between the purchase of a battery and its availability was not modelled. The maintenance costs associated with each upgrade were assumed to be negligible in comparison to the relevant capital costs.

3.1.4 Reliability

A key modelling consideration was the tradeoff between the cost required to upgrade the distribution network under consideration, and the resulting reliability of the network. Ideally, upgrading a distribution network should require as little investment as possible, and the resulting network should be able to continue to meet as much demand as possible in the event of a line failure. However, it is clear that these two objectives conflict; an increase in the reliability of a network typically requires a larger investment.

At each node of the scenario tree, a two-stage stochastic program similar to that derived by Singh, Philpott, and Wood (2004) was used to quantify the reliability of the network. In addition to the *base instance*, in which all lines of the network were operational, a set of *failure instances*, in which a single line in the network was unable to be used for a short period of time, were specified. It was assumed that no two failure instances would occur simultaneously. In the base instance, the network was required to meet all demand. The volume of unmet demand in each failure instance was used to evaluate the reliability of the network.

Several measures of reliability were modelled. An option was provided to penalise unmet demand with a fixed cost per kilowatt hour. This is common practice when modelling the expansion of distribution networks, and is convenient because it means that the two objectives are commensurate, so can be combined.

An alternative measure of network reliability commonly used by electricity providers is SAIDI (system average interruption duration index). The SAIDI of a network measures the amount of time per year a customer can expect to experience power outages for. An option was provided to specify a bound on the expected SAIDI of the network that could not be exceeded at any node of the scenario tree. In practice, it was found that this often led to significant variation in the expected SAIDI at different locations of the network; some locations experienced significantly more outages than average. To address this potential inequity, an additional option to specify a bound on the expected SAIDI experienced at each individual demand location within the network was modelled.

A final measure of reliability that was considered was the expected probability of a shortage within the network occurring at any point in time. This was enforced using chance constraints; options were provided to specify the maximum allowable probability of a shortage occurring over the whole network, or at individual locations of the network, at any point in time at each node of the scenario tree.

3.2 Model Formulation

The model was formulated within the Dantzig-Wolfe decomposition framework developed by Singh, Philpott, and Wood (2009); each node of the scenario tree formed a subproblem. The stochastic network reliability problem at each node of the tree was formulated as a MIP.

The following sections outline the notation and formulations of the restricted master problem and subproblems within the decomposition.

3.2.1 Notation

Sets and Indices

$n \in \mathcal{N}$	Set of nodes in the scenario tree. The root node is given by $n = 1$.
$h \in \mathcal{P}_n$	Set of predecessors of node n (including node n itself); that is, the path of nodes from node n to the root node of the tree. Note that the parent of node n is denoted using n_-.

$j \in \mathcal{J}_n$	Set of columns in the restricted master problem generated by subproblem n.
$i \in \mathcal{I}$	Set of locations in the network. The zone substation is given by $i = 1$. The remainder of the locations are modelled as demand points.
$i \in \mathcal{I}_b$	Set of demand locations at which batteries may be installed.
$i \in \mathcal{I}_{b'}$	Set of demand locations at which batteries may not be installed. Note that $\mathcal{I} = \{1\} \cup \mathcal{I}_b \cup \mathcal{I}_{b'}$.
$i \in \mathcal{I}_{>0}$	Set of demand locations with non-zero demand.
$k \in \mathcal{K}$	Set of lines.
$k \in \mathcal{K}_a$	Set of lines that are part of feeders.
$e \in \mathcal{E}$	Set of line expansions.
$m \in \mathcal{M}$	Set of load management upgrades.
$s \in \mathcal{S}$	Set of network instances. The base instance, in which no lines have failed, is denoted by $s = 1$. The remainder of the instances are failure instances.
$k \in \mathcal{K}_s$	Set of lines that fail in instance s. Note that $\mathcal{K}_1 = \varnothing$.

Restricted Master Problem Variables

λ^{nj}	Convex multiplier associated with column j generated by subproblem n.
w_e^n	1 if expansion e is purchased at node n, and 0 otherwise.
x_m^n	1 if load management upgrade m is purchased at node n, and 0 otherwise.
y_i^n	Number of batteries in use at location i at node n.
u^n	Number of batteries purchased at node n.
v_i^n	Number of batteries moved from location i between nodes n_- and n.

Subproblem Variables

p_e^n	1 if expansion e is used at node n, and 0 otherwise.
q_m^n	1 if load management upgrade m is used at node n, and 0 otherwise.
r_i^n	Number of batteries in use at location i at node n.
f_{ks}^n	Flow (kW) in line k in instance s at node n.
g_{ks}^n	Dummy flow in line k in instance s at node n.
o_{is}^n	Amount of shortage (kW) at demand location i in instance s at node n.
z_{ks}^n	1 if line k is active in instance s at node n, and 0 otherwise.

Restricted Master Problem Parameters

ϕ^n	Probability of the state of the world corresponding to node n occurring.
ρ^n	Discount factor at node n.
κ_e	Cost (\$) of line expansion e.
χ_m	Cost (\$) of load management upgrade m.
τ	Cost (\$) of purchasing a battery.

ν	Cost (\$) of moving a battery.
$\hat{p}_e^{nj}$	1 if line expansion e is used in extreme point j generated by subproblem n, and 0 otherwise.
$\hat{q}_m^{nj}$	1 if load management upgrade m is used in extreme point j generated by subproblem n, and 0 otherwise.
$\hat{r}_i^{nj}$	Number of batteries in use at location i in extreme point j generated by subproblem n.
Δ^{hn}	1 if a line upgrade purchased at node $h \in \mathcal{P}_n$ can be used at node n, and 0 otherwise.
Θ^{hn}	1 if a load management upgrade purchased at node $h \in \mathcal{P}_n$ can be used at node n, and 0 otherwise.

Subproblem Parameters

ϕ^n	Probability of the state of the world described by node n occurring.
ρ^n	Discount factor at node n.
ξ_i^n	Peak demand (kW) at location i at node n.
A_{ik}	1 if line k enters location i, -1 if line k leaves location i, and 0 otherwise.
c_k	Capacity of line k (kW), prior to any line expansions occurring.
$\bar{c}_k$	Maximum possible expanded capacity (kW) of line k.
E_{ke}	Additional capacity (kW) gained by line k if line expansion e is purchased.
M_{im}	Percentage reduction in peak demand at location i if load management upgrade m is purchased.
σ	Capacity (kW) of a single battery.
β	Maximum number of batteries that may be installed at a single location.
B	Maximum number of batteries that may be installed across the entire network.
ψ_s	Time (minutes per year) for which failure instance s is expected to occur.
Λ	Maximum allowable expected SAIDI of each demand location.

3.2.2 Restricted Master Problem Formulation

The restricted master problem (RMP) formulation is as follows. Note that dual variables for constraints are denoted using square brackets adjacent to the constraint formulation.

$$\text{RMP:} \quad \min_{\lambda,u,v,w,x,y} \quad \sum_{n\in\mathcal{N}} \rho^n \phi^n \left(\sum_{e\in\mathcal{E}} \kappa_e w_e^n + \sum_{m\in\mathcal{M}} \chi_m x_m^n + \tau u^n + \sum_{i\in\mathcal{I}_b} \nu v_i^n \right) \tag{1}$$

$$\text{s.t.} \quad \sum_{j\in\mathcal{J}_n} \hat{p}_e^{nj} \lambda^{nj} \leq \sum_{h\in\mathcal{P}_n} \Delta^{hn} w_e^h \qquad [\pi_e^n] \qquad \forall\, n \in \mathcal{N},\, e \in \mathcal{E} \tag{2}$$

$$\sum_{j\in\mathcal{J}_n} \hat{q}_m^{nj} \lambda^{nj} \leq \sum_{h\in\mathcal{P}_n} \Theta^{hn} x_m^h \qquad [\omega_m^n] \qquad \forall\, n \in \mathcal{N},\, m \in \mathcal{M} \tag{3}$$

$$\sum_{j \in \mathcal{J}_n} \hat{r}_i^{nj} \lambda^{nj} \le y_i^n \qquad [\gamma_i^n] \qquad \forall\, n \in \mathcal{N},\, i \in \mathcal{I}_b \tag{4}$$

$$\sum_{i \in \mathcal{I}_b} y_i^1 \le u^1 \tag{5}$$

$$\sum_{i \in \mathcal{I}_b} y_i^n - \sum_{i \in \mathcal{I}_b} y_i^{n-} \le u^n \qquad \forall\, n \in \mathcal{N} \setminus \{1\} \tag{6}$$

$$y_i^{n-} - y_i^n \le v_i^n \qquad \forall\, n \in \mathcal{N} \setminus \{1\},\, i \in \mathcal{I}_b \tag{7}$$

$$\sum_{j \in \mathcal{J}_n} \lambda^{nj} = 1 \qquad [\mu^n] \qquad \forall\, n \in \mathcal{N} \tag{8}$$

$$\lambda^{nj} \ge 0 \qquad \forall\, n \in \mathcal{N},\, j \in \mathcal{J}_n \tag{9}$$

$$w_e^n \ge 0 \qquad \forall\, n \in \mathcal{N},\, e \in \mathcal{E} \tag{10}$$

$$x_m^n \ge 0 \qquad \forall\, n \in \mathcal{N},\, m \in \mathcal{M} \tag{11}$$

$$y_i^n \ge 0 \qquad \forall\, n \in \mathcal{N},\, i \in \mathcal{I}_b \tag{12}$$

$$u^n \ge 0 \qquad \forall\, n \in \mathcal{N} \tag{13}$$

$$v_i^n \ge 0 \qquad \forall\, n \in \mathcal{N},\, i \in \mathcal{I}_b \tag{14}$$

The objective, (1), is to minimise the expected discounted cost of expansion of the network over the entire planning horizon. This is composed of the cost of line upgrades, load management upgrades, battery purchases, and battery movements (note that the extreme points generated by each subproblem do not have any associated costs).

Constraints (2) and (3) ensure that if a line expansion or load management upgrade is utilised at a node of the scenario tree, it must have been purchased at an earlier node. The amount of time between an expansion being purchased and available for use can be changed by altering the values of the Δ^{hn} and Θ^{hn} parameters.

Constraint (4) ensures that the number of batteries in use at each location at each node is recorded correctly. Constraint (5) ensures that the number of batteries that have been purchased at the root node is recorded correctly, while constraint (6) ensures that battery purchases at all other nodes are recorded correctly. Constraint (7) ensures that the number of batteries that have been moved from each location between each node (besides the root node, at which no movement costs are ever incurred) and its parent is recorded correctly.

Constraint (8) ensures that the convex multipliers associated with the set of extreme points from each node sum to 1. Constraint (9) ensures that the convex multipliers are non-negative, and constraints (10), (11), (12), (13), and (14) ensure that the line upgrade variables, load management variables, battery utilisation variables, battery purchase variables, and battery movement variables are non-negative.

3.2.3 Subproblem Formulation

The subproblem (SP) formulation for an arbitrary scenario tree node n is as follows. The formulation outlined imposes a bound on the expected SAIDI at each demand location in the network; this is the formulation used in the majority of the models described throughout this report. The formulations that model reliability using a SAIDI bound over the entire network, shortage costs, and chance constraints are similar in structure to this formulation; a software implementation of each of these can be found in the compendium that accompanies this report.

$$\mathrm{SP}_n: \quad \min_{f,g,p,q,r,o,z} \quad -\sum_{e\in\mathcal{E}} \pi_e^n p_e^n - \sum_{m\in\mathcal{M}} \omega_m^n q_m^n - \sum_{i\in\mathcal{I}_b} \gamma_i^n r_i^n - \mu^n \tag{15}$$

$$\text{s.t.} \quad \xi_i^n \leq \sum_{k\in\mathcal{K}} A_{ik} f_{ks}^n + \xi_i^n \sum_{m\in\mathcal{M}} M_{im} q_m^n + \sigma r_i^n + o_{is}^n \qquad \forall\, i\in\mathcal{I}_b,\ s\in\mathcal{S} \tag{16}$$

$$\xi_i^n \leq \sum_{k\in\mathcal{K}} A_{ik} f_{ks}^n + \xi_i^n \sum_{m\in\mathcal{M}} M_{im} q_m^n + o_{is}^n \qquad \forall\, i\in\mathcal{I}_{b'},\ s\in\mathcal{S} \tag{17}$$

$$r_i^n \leq \beta \qquad \forall\, i\in\mathcal{I}_b \tag{18}$$

$$\sum_{i\in\mathcal{I}_b} r_i^n \leq B \tag{19}$$

$$o_{is}^n \leq \xi_i^n \qquad \forall\, i\in\mathcal{I}\setminus\{1\},\ s\in\mathcal{S} \tag{20}$$

$$o_{i1}^n = 0 \qquad \forall\, i\in\mathcal{I}\setminus\{1\} \tag{21}$$

$$f_{ks}^n \leq c_k + \sum_{e\in\mathcal{E}} E_{ke} p_k^n \qquad \forall\, k\in\mathcal{K},\ s\in\mathcal{S} \tag{22}$$

$$-f_{ks}^n \leq c_k + \sum_{e\in\mathcal{E}} E_{ke} p_k^n \qquad \forall\, k\in\mathcal{K},\ s\in\mathcal{S} \tag{23}$$

$$f_{ks}^n \leq \bar{c}_k z_{ks}^n \qquad \forall\, k\in\mathcal{K},\ s\in\mathcal{S} \tag{24}$$

$$-f_{ks}^n \leq \bar{c}_k z_{ks}^n \qquad \forall\, k\in\mathcal{K},\ s\in\mathcal{S} \tag{25}$$

$$z_{ks}^n = 0 \qquad \forall\, s\in\mathcal{S},\ k\in\mathcal{K}_s \tag{26}$$

$$z_{ks}^n = 1 \qquad \forall\, s\in\mathcal{S},\ k\in\mathcal{K}_a\setminus\mathcal{K}_s \tag{27}$$

$$\sum_{k \in \mathcal{K}} A_{ik} g^n_{ks} = 1 \qquad \forall\, i \in \mathcal{I} \setminus \{1\},\ s \in \mathcal{S} \tag{28}$$

$$g^n_{ks} \leq (|\mathcal{I}| - 1)\, z^n_{ks} \qquad \forall\, k \in \mathcal{K},\ s \in \mathcal{S} \tag{29}$$

$$-\, g^n_{ks} \leq (|\mathcal{I}| - 1)\, z^n_{ks} \qquad \forall\, k \in \mathcal{K},\ s \in \mathcal{S} \tag{30}$$

$$\sum_{k \in \mathcal{K}} z^n_{ks} = |\mathcal{I}| - 1 \qquad \forall\, s \in \mathcal{S} \tag{31}$$

$$\sum_{s \in \mathcal{S} \setminus \{1\}} \psi_s \frac{o^n_{is}}{\xi^n_i} \leq \Lambda \qquad \forall\, i \in \mathcal{I}_{>0} \tag{32}$$

$$p^n_e \in \{0, 1\} \qquad \forall\, e \in \mathcal{E} \tag{33}$$

$$q^n_m \in \{0, 1\} \qquad \forall\, m \in \mathcal{M} \tag{34}$$

$$r^n_i \in \mathbb{Z}_{\geq 0} \qquad \forall\, i \in \mathcal{I}_b \tag{35}$$

$$o^n_{is} \geq 0 \qquad \forall\, i \in \mathcal{I},\ s \in \mathcal{S} \tag{36}$$

$$z^n_{ks} \in \{0, 1\} \qquad \forall\, k \in \mathcal{K},\ s \in \mathcal{S} \tag{37}$$

The subproblem objective, (15), is to minimise the cost associated with utilising line upgrades, load management upgrades, and batteries. These costs are adjusted using the optimal values of the dual variables from the master problem prior to each solve. If the optimal objective is negative, the corresponding extreme point has a negative reduced cost, so is added to the restricted master problem.

Constraints (16) and (17) ensure that if the net amount of power supplied to a demand location is less than the amount required (which may be reduced by the use of a load management upgrade), a shortage is recorded. Constraint (16) is used for locations at which batteries can be installed, while constraint (17) is used for locations at which batteries cannot be installed.

Constraint (18) ensures that the number of batteries in use at a given location does not exceed the allowable number. Constraint (19) ensures that the total number of batteries in use over the entire network does not exceed the number available.

Constraint (20) ensures that the volume of shortage at each location does not exceed demand. Constraint (21) ensures that there are no shortages in the base instance.

Constraints (22) and (23) ensure that the flow in each line does not exceed its (potentially expanded) capacity. Constraints (24) and (25) ensure that if a line is inactive, no flow may travel through it. The coefficient associated with the variable that determines whether the line is active is set to the largest possible expanded capacity of the line, tightening these constraints to the greatest degree possible.

Constraint (26) ensures that in each failure instance, the lines that have failed are inactive. Constraint (27) ensures that lines that are part of the base configuration of the network are always active, unless they have failed.

Constraints (28)-(31) ensure that the network always operates in a radial configuration. These constraints require the use of a set of dummy flows. Constraint (28) ensures that the net dummy flow into each demand location is equal to one unit during all instances. Constraints (29) and (30) ensure that there is no dummy flow in lines that are inactive. These constraints, therefore, enforce the requirement that there must always be a path through closed lines between each pair of locations in the network. Constraint (31) ensures that the number of active lines is always equal to one less than the number of nodes in the network; this means that the path between each pair of locations is unique.

Constraint (32) ensures that the expected SAIDI of each demand location does not exceed the specified bound. In each failure instance, the expected duration of outages experienced by the average customer at each demand location is given by the product of the time the failure is expected to occur for and the proportion of unmet demand at the location. The expected SAIDI of each demand location is simply the sum of the expected outage durations over all failure instances.

Constraints (33) and (34) ensure that decisions to utilise line upgrades and load management upgrades are binary, and constraint (35) ensures that the number of batteries in use at each demand location where batteries may be installed is integer. Constraint (36) ensures that shortages are non-negative, and constraint (37) ensures that the variables that control whether each line is active are binary.

3.3 Software Implementation

The preceding model formulation and associated variations were implemented as part of a Julia package, GENE.jl (Generalised Electricity Network Expansion). GENE leverages another Julia package, JuDGE.jl (Julia Decomposition for Generalised Expansion), to formulate and solve problems using Dantzig-Wolfe decomposition.

3.3.1 JuDGE

A potential deterrent associated with the use of Dantzig-Wolfe decomposition to solve capacity expansion problems is the difficulty associated with defining a suitable master problem and set of subproblems, and coordinating the solving of these in an iterative process. JuDGE (Downward, Baucke, and Philpott, 2020) aims to make the use of this decomposition more accessible to practitioners.

JuDGE implements the nodal decomposition described by Singh, Philpott, and Wood (2009). A user of JuDGE provides the structure of a scenario tree and specifies a suitable subproblem formulation using the syntax of the JuMP modelling language (Dunning, Huchette, and Lubin, 2017). JuDGE uses this information to automatically construct the associated restricted master problem, and solves the overall problem using the Dantzig-Wolfe algorithm, performing a branch and price procedure if required. The master problem and subproblems can be solved using a range of common solvers.

In addition to solving problems using the Dantzig-Wolfe algorithm, JuDGE is also capable of constructing the deterministic equivalent problem given a master problem and set of subproblem formulations. This allows for comparisons between the two types of models to be made.

3.3.2 GENE

GENE has been designed to provide an intuitive interface to model and visualise the results of capacity planning for distribution networks. Figure 3.1 illustrates the structure of GENE.

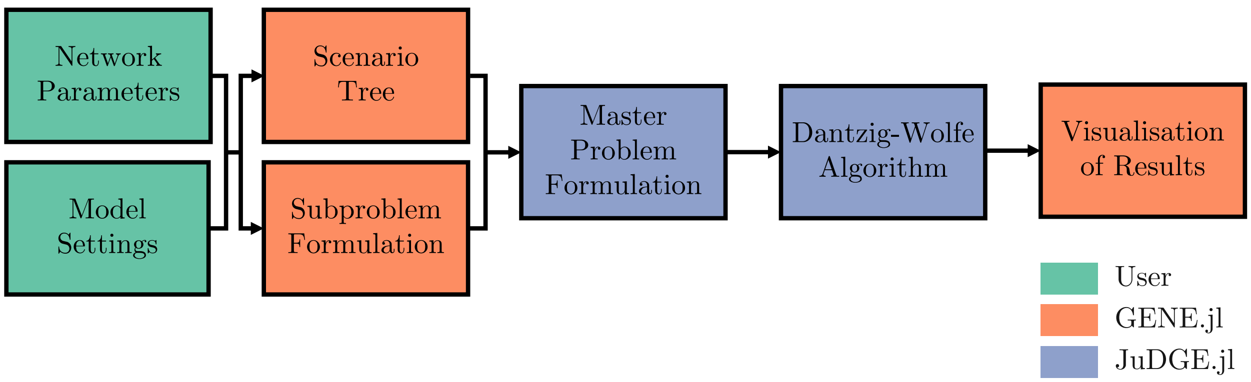

Figure 3.1. The structure of GENE, and its interaction with JuDGE.

GENE uses a spreadsheet-based input system that allows for easy input and modification of the parameters of the network under consideration, and the desired settings, which include the parameters of the scenario tree and the reliability measure to be used. The reader is referred to Appendix A for an overview of the input system. Given the information passed in, GENE generates the appropriate scenario tree and subproblem formulation. JuDGE is then used to construct the corresponding master problem formulation and solve the resulting model.

Once GENE has run a model, it is capable of visualising the model results. The visualisation module in GENE builds on the work of Dong (2018) and van Helsdingen (2021). It uses Leaflet, a JavaScript library, to generate a series of interactive maps showing characteristics of the expansion plan specified by the model solution. In addition to displaying the upgrades made to the network at each node of the scenario tree, and the lines which are active under each failure instance, which were also visualised by van Helsdingen, GENE is capable of visualising the expected SAIDI at each location in the network, and the change in the amount of flow in each line of the network during each failure instance. The reader is referred to Appendix B for a demonstration of these capabilities, and to Aish (2022) for a description of the development of the visualisation module.

3.4 Case Studies and Computational Testing

To demonstrate the functionality of GENE, and to evaluate the effectiveness of Dantzig-Wolfe decomposition when solving the types of network design problems studied in this project, a set of case studies and computational tests were undertaken.

3.4.1 Test Network

All case studies and computational tests were conducted using data for Vector's distribution network located in Clendon, Auckland. The dataset for the network was adapted from that used by Dong (2018). These data only contained information on lines that formed part of feeders; to allow the survivability of the network to be investigated, a set of backfeeders were added. The augmented network contained a single zone substation, 25 junctions, 189 demand substations, eight feeders composed of a total of 214 lines, and 10 backfeeders. This was significantly larger than the network studied by van Helsdingen, which contained fewer than 20 locations and 50 lines. Figure 3.2 shows a map of the network generated using GENE.

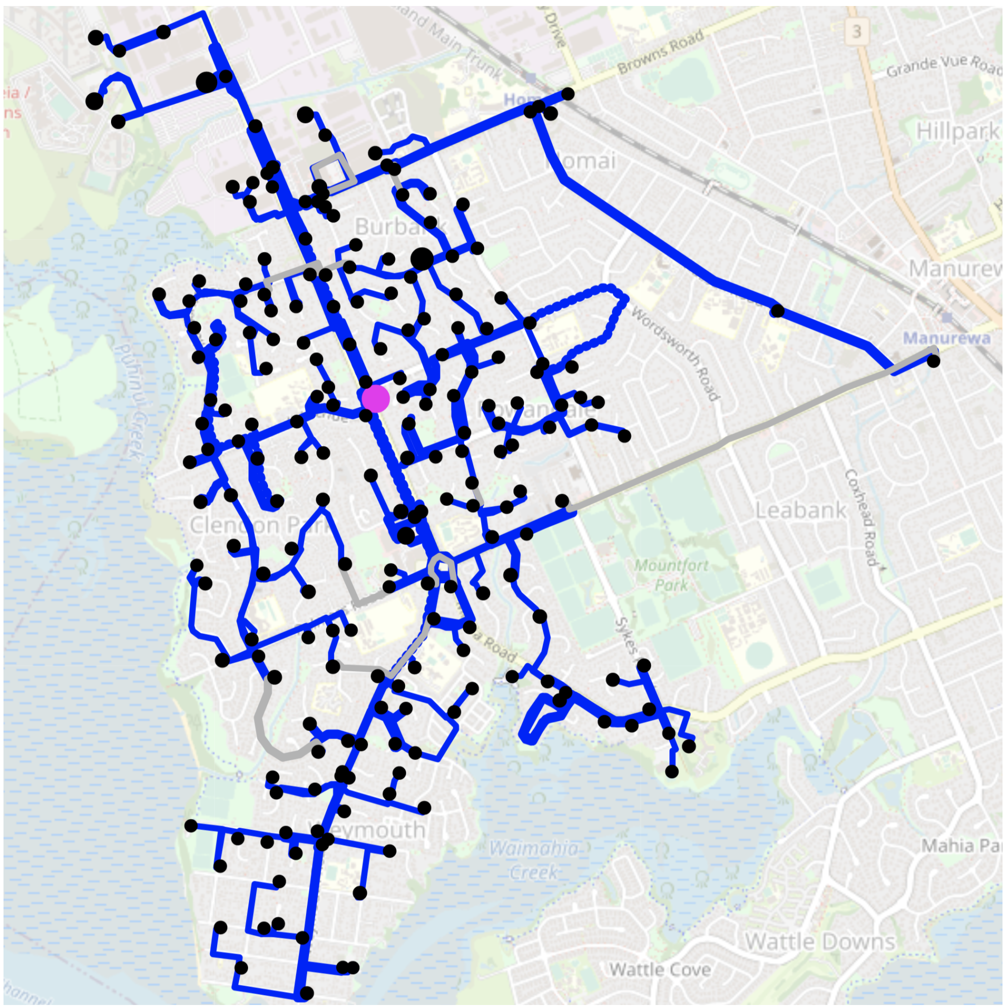

Figure 3.2. The test network. The Clendon zone substation is denoted using a pink circle, and junctions and demand substations are denoted using black circles. Lines that form part of feeders are blue, and backfeeders are grey.

Each line was associated with a corresponding upgrade, the capacity and cost of which were proportional to the initial capacity of the line. The additional capacity granted by each line upgrade ranged between 400 kW and 2500 kW, and the cost of each upgrade ranged between \$800,000 and \$2.5 million. A set of 20 representative locations were specified as locations at which batteries could be installed. Batteries had a capacity of 150 kW, and cost \$75,000. Moving a battery between two demand locations cost \$40,000. No information was available on the cost or reduction in peak demand associated with load management upgrades; these upgrades, therefore, were not included in the dataset.

A set of 12 key lines were specified as candidates to fail. Each of these corresponded to a failure instance. The times associated with each failure instance were sampled from a uniform distribution with bounds of 30 minutes per year and 90 minutes per year. Present day peak demands at each location were specified such that no location had a SAIDI of more than 75 minutes per year prior to any expansions being made.

3.4.2 Computational Testing

A set of computational tests were run to determine the effectiveness of the use of Dantzig-Wolfe decomposition to solve the network expansion problems studied in this project, relative to the use of the deterministic equivalent formulation.

Gurobi (Gurobi Optimization, 2022) was used as the solver for both types of models. When solving the master problem as part of a decomposed model, a barrier method was used; in their own computational tests, Singh, Philpott, and Wood (2009) observed that the master problem tended to suffer from dual degeneracy, and found this to be the best method to overcome it. If the master problem returned a fractional solution for five consecutive iterations of the Dantzig-Wolfe algorithm, it was solved as a MIP (a relatively inexpensive operation) in the hope of finding a new incumbent solution.

All computational tests were conducted on a laptop with a 2.4 GHz quad-core Intel Core i5 processor and 8 GB of RAM.

4 | Results and Discussion

A number of case studies were conducted throughout this project to illustrate the functionality of GENE. This chapter presents a subset of these. In addition, testing was conducted to evaluate the use of Dantzig-Wolfe decomposition to solve the network design problems considered in this project. These results are presented and the differences between the two types of models are discussed.

4.1 Case Studies

The case studies presented in this report investigate the differences in optimal network design with and without batteries as an investment option, the tradeoff between the desired reliability and cost of expansion of a network, and the benefits of using stochastic programming as a solution method for network expansion problems.

During all but the final case study, a scenario tree with a depth of two and degree of two was used (these are the same dimensions as the tree in Figure 2.3); the *depth* of a scenario tree is one less than the number of stages, while the *degree* refers to the number of nodes at the next stage each non-leaf node branches into. All scenarios had equal probabilities of occurring. The length of each stage was assumed to be five years; therefore, the planning period under consideration was 15 years. The rate of increase of demand was modelled as a Markov process; at each stage, it increased uniformly over the network by a factor of either 10% or 20% depending on the branch of the tree taken. A lag of one stage was enforced between a line upgrade being purchased and available for use; batteries, however, could be used at the stage at which they were purchased.

4.1.1 Planning Using Batteries

GENE can be used to gain insight into the value of different investments. To demonstrate the value of using batteries when expanding the test network, two models were developed; one in which up to two batteries could be purchased over the planning horizon (`Clendon-B2`), and another in which no batteries were available (`Clendon-B0`). In each model, the SAIDI at each demand location was constrained such that it did not exceed 90 minutes per year.

Figure 4.1 shows the investments that are made during the first stage, and during the second stage under a growth rate of 20%, in the optimal expansion plan found for `Clendon-B0`. At the first stage, six line upgrades are invested in, and at the second, eight are invested in. The expected cost of expanding the network over all scenarios is approximately $48 million.

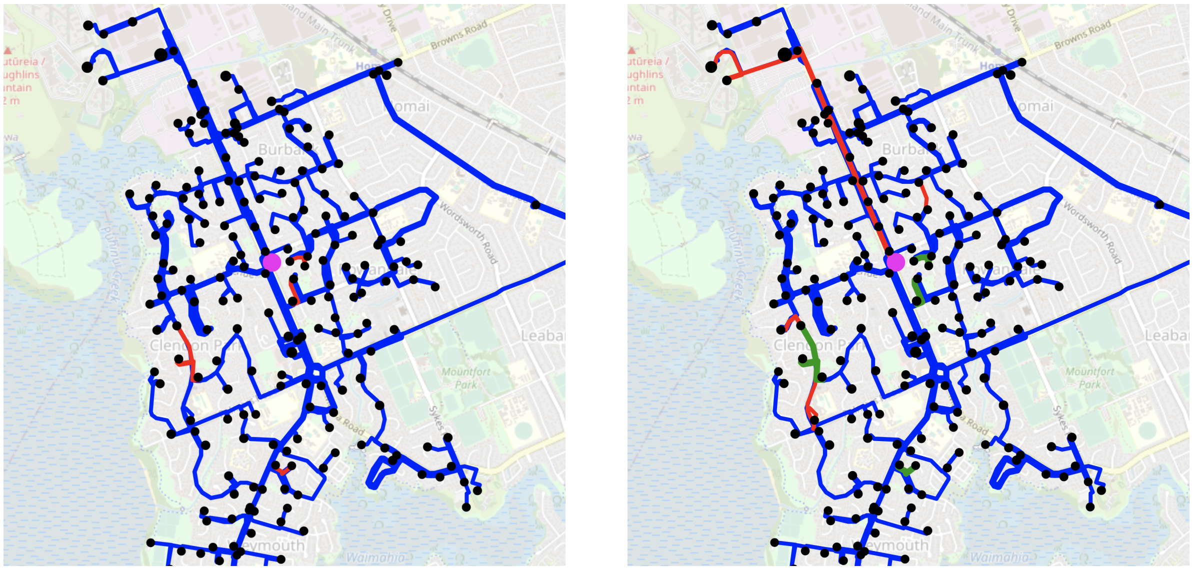

(a) First stage **(b)** Second stage

Figure 4.1. Line upgrades purchased (red) and in use (green) during the first stage, and the second stage under a growth rate of 20%, in the optimal plan found for `Clendon-B0`.

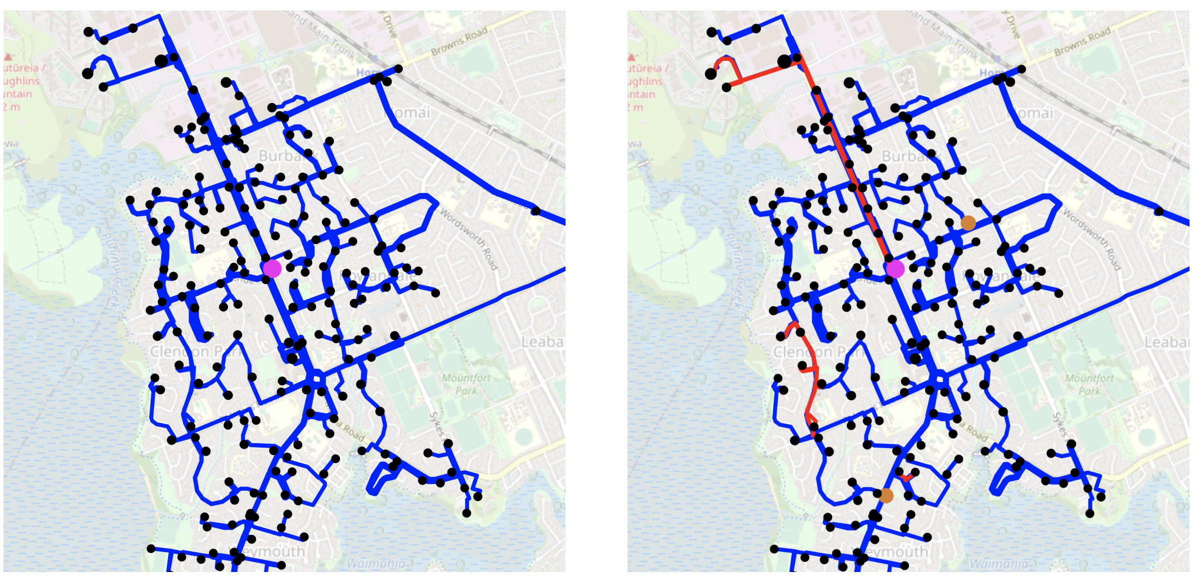

(a) First stage **(b)** Second stage

Figure 4.2. Line upgrades purchased (red), and batteries in use (orange) during the first stage, and the second stage under a growth rate of 20%, in the optimal plan found for `Clendon-B2`.

Figure 4.2 shows the investments made in the optimal expansion plan associated with `Clendon-B2`. This set of investments differs significantly to those made in the plan found for `Clendon-B0`; no upgrades are purchased during the first stage, while both batteries, and 10 line upgrades, are purchased during the second stage under a growth rate of 20%. Many of the line expansions made in the second stage of this model are made during the first stage of `Clendon-B0`, and the expected cost of expanding the network over the complete planning horizon is significantly reduced; $24 million. These results demonstrate the ability of batteries to delay the need for larger, more expensive investments.

It is also interesting to observe how batteries are moved throughout the network in the expansion plans found using GENE. Figure 4.3 shows the investments made at the second and third (final) stage of the optimal expansion plan found for `Clendon-B2`. At the second stage, one battery is placed in the eastern part of the network, while the other is placed in the southern part. However, at the third stage, after several line upgrades in the southern part of the network become active, the battery that was previously in this region is moved to the eastern part of the network, to accommodate increasing demand there. This further demonstrates the flexibility and utility of batteries.

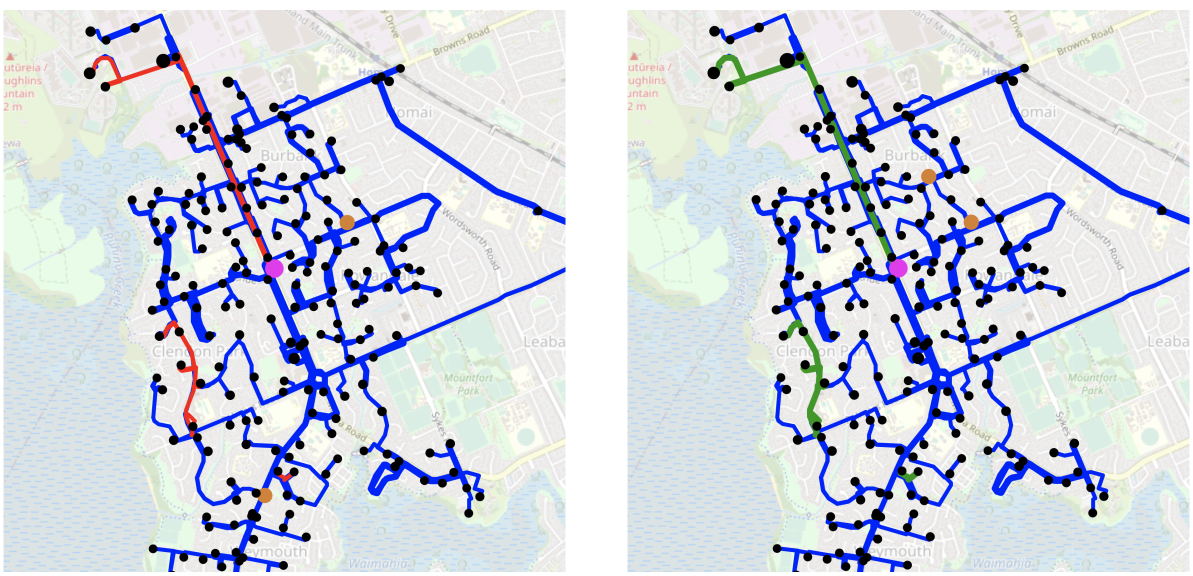

(a) Second stage **(b)** Third stage

Figure 4.3. Line upgrades purchased (red) and in use (green), and batteries in use (orange) during the second and third stages of the highest growth scenario, in the optimal plan found for `Clendon-B2`.

4.1.2 Expansion Cost and Network Reliability

GENE can be used to illustrate how the optimal expansion plan differs under varying desired levels of network reliability. To demonstrate this, a range of models with SAIDI tolerances at each demand location of between 80 and 100 minutes per year were run. The expected cost of the optimal set of investments associated with each model was plotted; these results are shown in Figure 4.4.

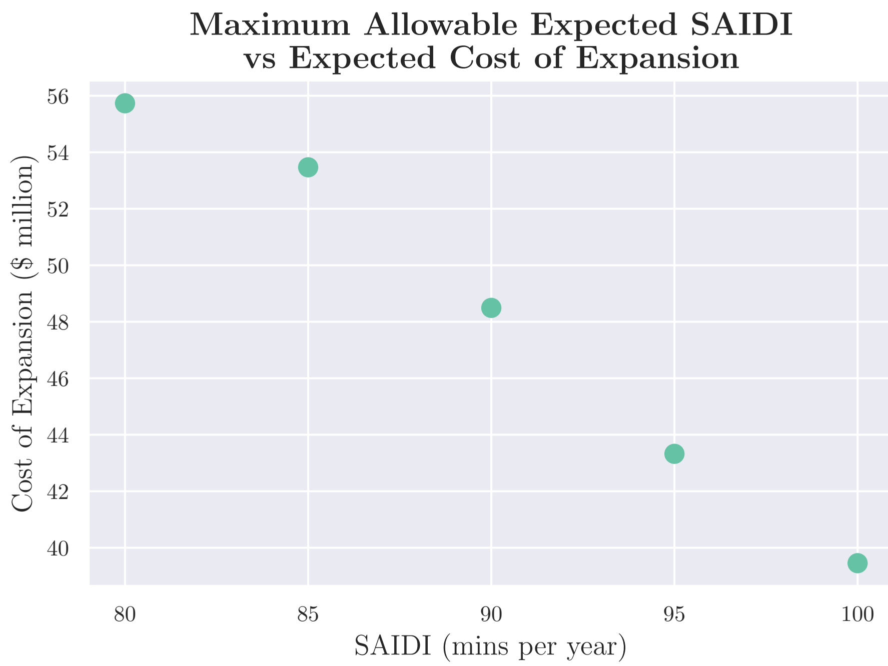

Figure 4.4. The expected cost of the optimal expansion plan for the test network under a range of SAIDI tolerances.

As the allowable SAIDI of a network increases, outages become more frequent; that is, the network becomes less reliable in the event of a failure. Figure 4.4 shows that as the desired level of reliability of the test network decreases, so does the required investment in expansions; relaxing the allowable SAIDI at each demand location from 80 minutes per year to 100 minutes per year results in a reduction of over $15 million in the required amount of investment. This illustrates the tradeoff between cost and reliability that must be considered when developing an expansion plan.

4.1.3 The Value of Stochasticity

GENE is capable of investigating the differences between an optimal expansion plan found using a stochastic program and one found using a deterministic program, and, therefore, the value of accounting for uncertainty explicitly during the planning process. In contrast to previous case studies, this case study investigates a two-stage problem in which the increase in demand between the two stages varies spatially, rather than increasing uniformly across the network. Two batteries were available, and a SAIDI tolerance of 90 minutes per year was enforced at each demand location.

Two equally likely scenarios were considered. In both of these, the location which was expected to have the highest increase in demand was specified. Demand at this location was expected to increase by 50% after the first stage. The increase in demand at all other locations was modelled as decaying exponentially the further the location was from the central location, such that the demand at a location 1.5 km away from the centre increased

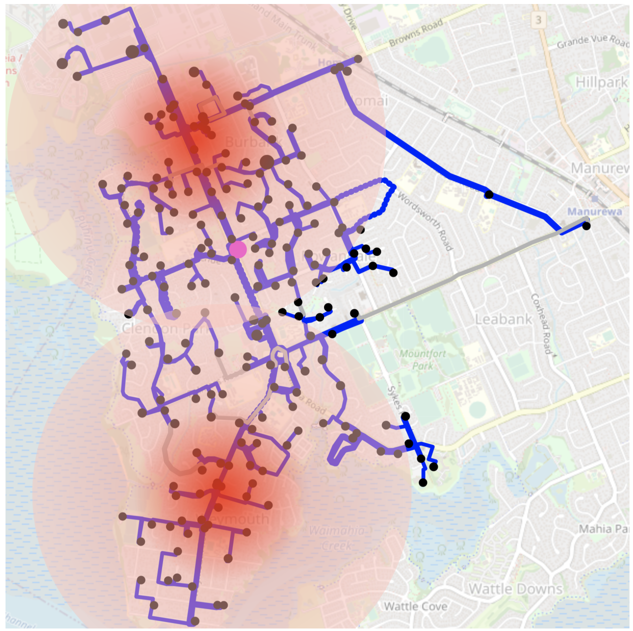

Figure 4.5. The centres of demand growth in the two scenarios considered. The edge of each red circle indicates the radius within which demand at each location increased at a rate that was at least half as much as at the centre.

by 25%. Figure 4.5 illustrates the centres of demand growth in each scenario.

As opposed to using a stochastic model, a network planner might instead choose to solve a simpler model that aims to account for both potential scenarios simultaneously. This was modelled using a two-stage deterministic model where, at the second stage, the increase in demand at each location was the largest increase in either of the previously modelled scenarios. Once the deterministic model was solved, the investments made at the first stage were fixed, and the model was re-run to determine the optimal investments to make at the second stage depending on the scenario from the stochastic problem that eventuated.

The investments made in the optimal solution to the stochastic problem are shown in Figure 4.6. Six lines are upgraded at the first stage. The majority of these are in the south of the network; the north of the network is left largely untouched. Two batteries are purchased at the second stage. In both scenarios, one battery is placed close to the zone substation. However, the other is placed in the far north or south of the network, depending on which scenario eventuates.

The investments made in the optimal solution to the deterministic problem are shown in Figure 4.7. Five line upgrades are made at the first stage. Similarly to the stochastic problem, the southern side of the network is the focus of these expansions, though the level of reinforcement made is greater, and the area that is reinforced differs slightly. At the second stage, one battery is placed close to the zone substation. In the scenario in which demand increases in the north of the network, an additional battery is placed in the north. However, in the scenario in which demand increases in the south, no additional battery is required.

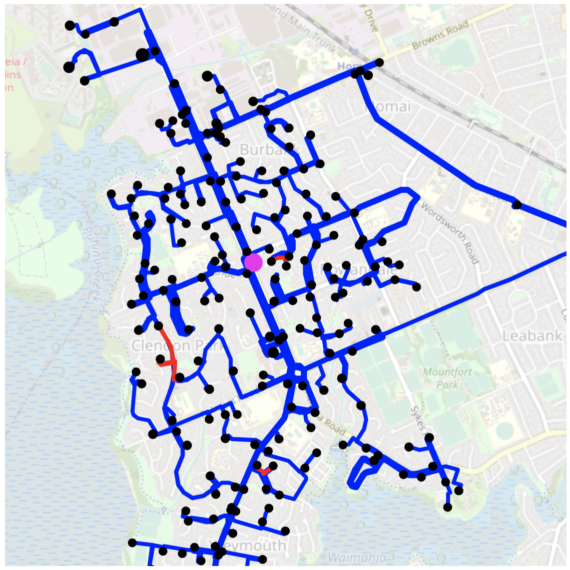

(a) First stage

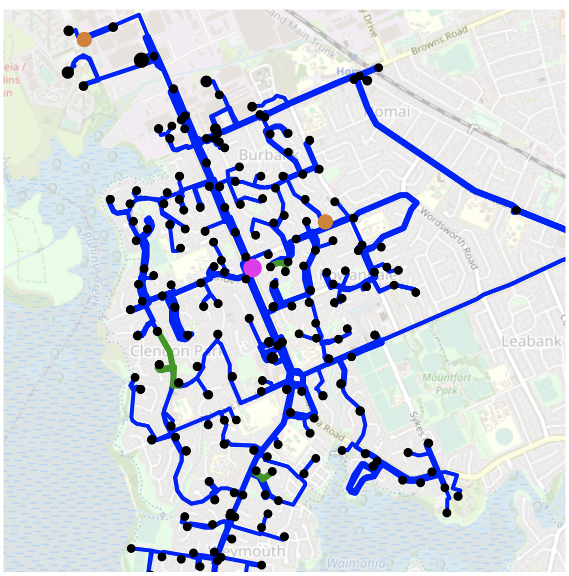

(b) Northern demand growth

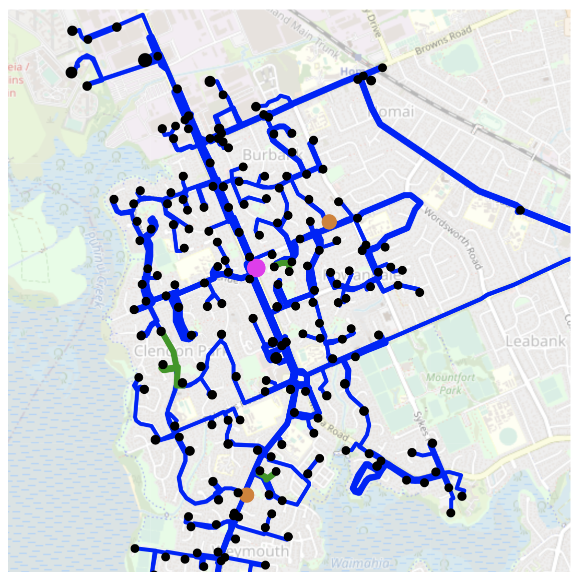

(c) Southern demand growth

Figure 4.6. Line upgrades purchased (red) and in use (green), and batteries in use (orange), during the first stage, and second stage under high growth in the north or south of the network, in the optimal plan found for the stochastic model.

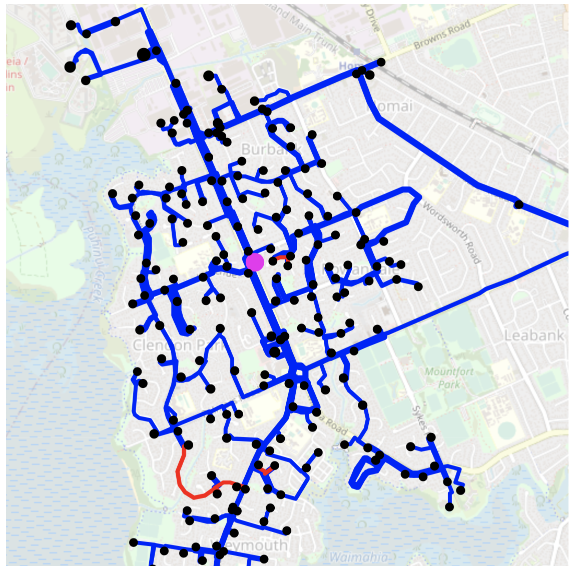

(a) First stage

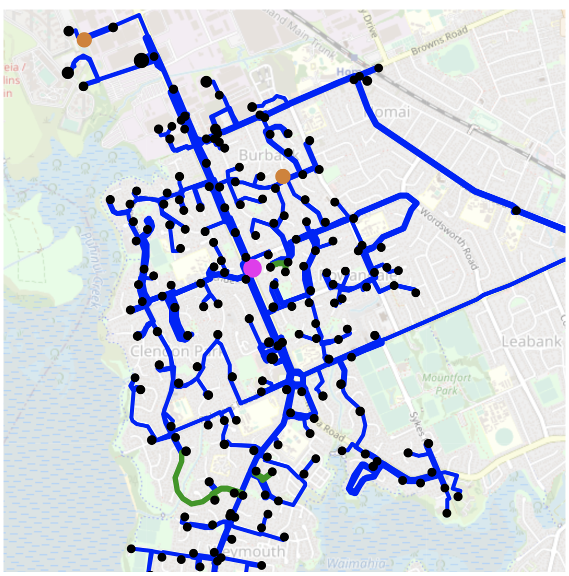

(b) Northern demand growth

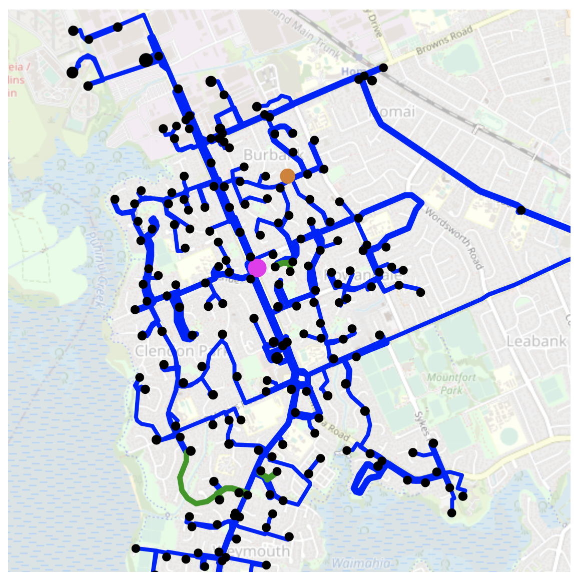

(c) Southern demand growth

Figure 4.7. Line upgrades purchased (red) and in use (green), and batteries in use (orange), during the first stage, and second stage under high growth in the north or south of the network, in the optimal plan found for the deterministic model.

The differences between the first stage decisions associated with the stochastic and deterministic models are the result of the stochastic model considering both demand scenarios independently at the first stage; it is known that it is possible to use batteries at the second stage to accommodate, in part, either demand scenario that could eventuate. Therefore, both the northern and southern parts of the network can be upgraded partially. However, it is not possible to use batteries to partially meet demand in both the north and south of the network simultaneously; in the deterministic model, the southern part of the network is upgraded entirely using line expansions, such that is does not rely on power from batteries at the next stage. The deterministic model, therefore, over-invests in capacity in the southern part of the network that may be unnecessary at the second stage.

This is reflected in the expected cost associated with each investment plan; the expected cost of implementing the plan found by solving the stochastic model is \$18.4 million, while the expected cost of the plan found by solving the deterministic model is \$22.5 million. This case study, therefore, illustrates the value of using stochastic programming to account for uncertainty explicitly, rather than attempting to account for all potential scenarios using a single, deterministic model.

4.2 Computational Testing

To evaluate the performance of Dantzig-Wolfe decomposition for solving the network design problems studied in this project, two sets of computational tests were conducted. In the first set (`S-60`), the expected SAIDI across the entire network was constrained such that it did not exceed 60 minutes per year, and in the second (`S-30`), it was constrained such that it did not exceed 30 minutes per year.

Within each set of tests, a number of models with scenario trees of varying sizes were run. In each model, five batteries were available. In models with a scenario tree of degree three, demand increased uniformly over the network by 5%, 10%, or 15% at each stage. In models with a tree of degree two, demand increased uniformly over the network by 5% or 15% at each stage.

Results of the `S-60` tests are provided in Table 4.1. It appears that, for models with small scenario trees, solving the deterministic equivalent formulation is more effective than using Dantzig-Wolfe decomposition. However, as the size of the scenario tree increases, the decomposed model begins to solve faster than the deterministic equivalent. The problems within this set of tests that the decomposition is capable of solving to near-optimality are very large; the largest test problem, which contains 81 scenarios, is solved to within 0.1% of optimality within two hours. The deterministic equivalent, by contrast, is unable to find an incumbent solution in this timeframe.

Results of the `S-30` tests are provided in Table 4.2. These are markedly different to the results of the `S-60` tests; the solve times for all problems are significantly greater. There appears to be a tradeoff between the desired reliability of the network and the time required to solve both types of model; as the required reliability of the network increases, so does the time required to solve the resulting models to near-optimality. The deterministic equivalent remains more effective when solving the problems with smaller scenarios trees. As the size

Table 4.1. Results for the `S-60` set of test problems. A dash indicates that no incumbent was found after two hours.

				1%		0.1%	
Degree	Depth	Nodes	Scenarios	Decomposed	DetEq	Decomposed	DetEq
2	2	7	4	43 s	9 s	43 s	9 s
2	3	15	8	116 s	242 s	325 s	242 s
2	4	31	16	415 s	1130 s	809 s	1149 s
3	2	13	9	226 s	27 s	226 s	27 s
3	3	40	27	732 s	1824 s	1662 s	1832 s
3	4	121	81	3233 s	–	6051 s	–

Table 4.2. Results for the `S-30` set of test problems. A dash indicates that no incumbent was found after two hours. In the event an incumbent was found but the model did not solve to the desired bound gap after two hours, the bound gap after two hours is provided in bold.

				1%		0.1%	
Degree	Depth	Nodes	Scenarios	Decomposed	DetEq	Decomposed	DetEq
2	2	7	4	103 s	34 s	103 s	34 s
2	3	15	8	227 s	374 s	369 s	374 s
2	4	31	16	**4.0%**	**8.5%**	**4.0%**	**8.5%**
3	2	13	9	334 s	68 s	334 s	68 s
3	3	40	27	4062 s	2877 s	**0.2%**	2877 s
3	4	121	81	**51.9%**	–	**51.9%**	–

of the problem grows, however, the decomposition obtains better results in some cases, while the deterministic equivalent obtains better results in others. In the largest test problem, neither model appears to be particularly effective. Although the decomposition is able to find an incumbent solution (in contrast to the deterministic equivalent), the bound gap after two hours remains higher than 50%.

Throughout these tests, the solve times of the decomposed models are dictated primarily by the subproblems. For example, when a scenario tree of depth three and degree three is used, the time spent solving the subproblems makes up approximately 98% of the solve time, regardless of the SAIDI tolerance. However, the subproblems appear to become more difficult to solve as the required reliability of the network increases; when a SAIDI tolerance of 60 minutes per year is used, solving the full set of subproblems at each iteration takes approximately 20 seconds on average, but when a tolerance of 30 minutes per year is used, this increases to 35 seconds.

The increase in solve times for both types of model as the required reliability of the network increases is theorised to be due, at least in part, to the increased number of expansions required. If all other parameters are held constant, increasing the desired level of reliability of a network typically necessitates the use of a greater number of expansions over the planning horizon. This is likely to cause an increase in the number of non-zero, potentially fractional expansion variables during the optimisation. The resulting problem, therefore, will be more difficult to solve, regardless of whether it is solved as a large-scale MIP or decomposed into a set of smaller problems.

Both sets of results suggest that the decomposition becomes increasingly competitive with the deterministic equivalent as the size of the scenario tree grows, and can be expected to outperform it once the tree reaches a certain size. However, the point at which this happens appears to be related to the difficulty associated with solving the subproblems; an increase in the time the subproblems of the decomposition take to solve results in an increase in the size of the scenario tree required for the benefits of the decomposition to be realised. If the subproblems are sufficiently difficult to solve, it is possible that once the scenario tree becomes large enough that the decomposition outperforms the deterministic equivalent, neither model can be solved in an acceptable timeframe.

There is a high level of sensitivity between the parameters of the problem under consideration and the relative performance of each type of model. Given a particular set of model parameters and a desired scenario tree size, it appears difficult, in general, to know which will be more effective a priori. Experimentation is likely required; this may involve partially solving the full formulation using both models, or solving a set of subproblems to ascertain their difficulty and considering the possible implications of this on the solve time of the decomposition.

An additional observation that can be made from both sets of results is that the decomposition appears to suffer from a *tailing-off effect* (see e.g. Desaulniers, Desrosiers, and Solomon, 2005); that is, as the bound gap decreases, convergence becomes increasingly slow. Figure 4.8 shows how the relative bound gap changes over time for both types of model, when a scenario tree of depth three and degree three is used and the SAIDI tolerance is 60 minutes per year. The decomposed model spends over half the solve time reducing the

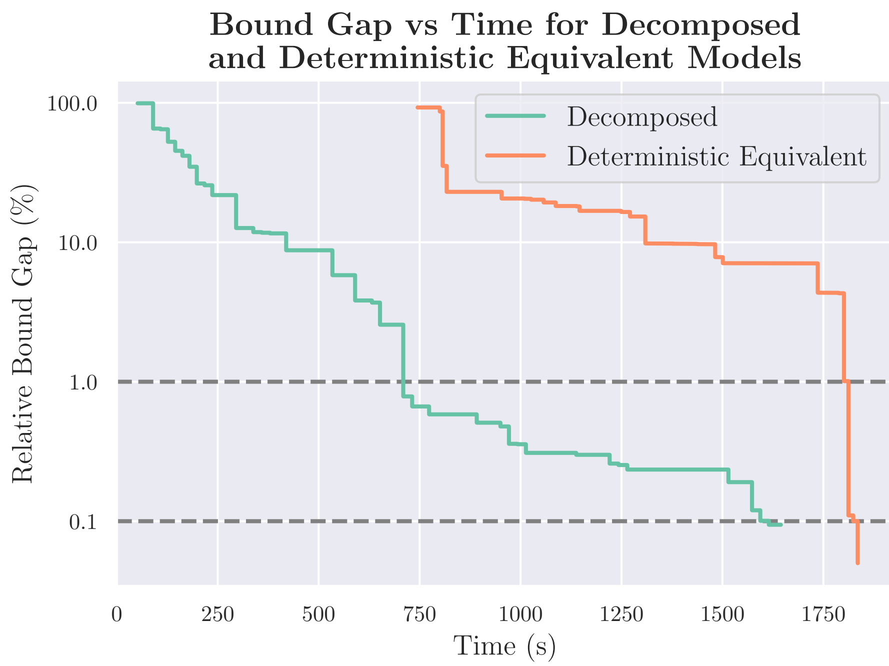

Figure 4.8. The relative bound gap of the decomposed and deterministic equivalent models over time, for the test problem with a scenario tree of depth three and degree three, where the SAIDI tolerance is 60 minutes per year. The dashed lines indicate relative bound gaps of 1% and 0.1%.

bound gap from 1% to 0.1%. The deterministic equivalent problem, by contrast, appears to spend very little time reducing the bound gap from 1% to 0.1%. A similar pattern can be observed across almost all test problems of a reasonable size.

To address this tailing-off effect, a hybrid solve was trialled for several of the larger problem instances in which the decomposed model reached a bound gap of 1% prior to the deterministic equivalent. The decomposed model was solved to within 1% of optimality, and the incumbent solution was used to seed the deterministic equivalent model, which was then solved to within 0.1% of optimality. This, however, did not result in an improvement in solve times. Seeding the deterministic equivalent model with an incumbent solution allows for the pruning of nodes with poor objective values from the resulting branch and bound tree, potentially speeding up the solve. However, in the cases trialled, these benefits were insignificant in comparison to the time spent solving the decomposed model to obtain an incumbent. It appears likely that, when it is not provided with an incumbent solution, the time the deterministic equivalent model spends solving nodes with poor objective values is minimal in comparison to the time spent carrying out other tasks which are largely unaffected by having an incumbent solution, such as solving nodes with objectives that are better than the incumbent, and adding valid inequalities to the problem to reduce fractionality.

5 | Conclusions and Future Work

The central objective of this project was to develop software capable of using stochastic programming to find least cost expansion plans for distribution networks that maintained desired levels of network reliability, while accounting for uncertainty in future demand explicitly. This project aimed to build on the work of van Helsdingen (2021) to develop an accurate and informative tool capable of generating useful insights for electricity providers. While the project was largely successful in meeting these objectives, several improvements could be made.

5.1 Conclusions

To achieve the objectives of this project, a Julia package, GENE, has been developed. GENE leverages an additional package, JuDGE, to solve problems using Dantzig-Wolfe decomposition, and has been designed to generalise to any distribution network. GENE addresses key shortcomings that are often associated with conventional planning methods; it accounts for uncertainty in future demand explicitly, and is capable of quantifying the goodness of a given plan it produces relative to the best possible plan. The functionality of GENE has been demonstrated using several case studies on Vector's distribution network in Clendon, which is significantly larger than the network studied by van Helsdingen (2021). GENE also contains a number of key features that represent improvements on the work carried out by van Helsdingen. These include

- the ability to model batteries that can be moved between locations at a cost;
- the ability to quantify network reliability using multiple metrics, including SAIDI, and to consider the reliability of both individual locations and the network as a whole;
- the ability to visualise additional characteristics of an expansion plan, including the expected SAIDI at each demand location and the change in flows in each line resulting from a failure; and
- the ability to formulate, for a given stochastic model, a deterministic model that plans for all possible demand scenarios simultaneously, and compare the optimal expansion plans obtained using each model.

This project has also provided insight into the differences between using Dantzig-Wolfe decomposition, relative to using the deterministic equivalent formulation, to solve network expansion problems. It has been found that as the size of the scenario tree grows, the decomposition becomes increasingly competitive with the deterministic equivalent. However, the size of the tree required for the decomposition to begin to outperform the deterministic equivalent is sensitive to small changes in the model parameters. A key parameter appears to be the desired reliability of the network; a stricter reliability requirement results in an increase of the size of the tree at which the benefits of the decomposition begin to be realised.

5.2 Future Work

There are a number of improvements and extensions that could be made to the modelling that was conducted as part of this project. A key limitation associated with this project was the dataset for the test network, which was lacking a number of components. In addition, modifications could be made to the models developed as part of this project, to make them more representative of reality, or to increase their tractability.

5.2.1 Network Data

During this project, no data on the cost or volume of demand reduced by load management upgrades were available; consequently, although GENE is able to model these upgrades, they were not included in any of the case studies conducted. The cost and capacity of batteries that were likely to be used when upgrading the network were unknown, so had to be estimated. No data on the frequency or duration of line failures were available; instead, these parameters required estimation as well. Though these estimates were sufficient when conducting case studies to illustrate the functionality of GENE, for GENE to generate insights that are useful to network planners, these data need to be available.

No forecasts for the change in demand expected over the planning horizon were provided for the test network; as a result, in the case studies conducted, the rate of increase of demand at each stage was typically modelled in a simple manner, as a Markov process. If a forecast for the distribution of future demand for electricity was available, more sophisticated methods to generate suitable demand scenarios could be used; for example, sampling methods, or the matching of moments between the forecast distribution and the discrete approximation.

Finally, no data were available on the typical variation in demand at each location in the network over the course of a day, nor the rate at which a typical battery charges. Therefore, it is difficult to determine whether the level of demand met by batteries in an expansion plan found by GENE is sustainable; that is, whether each battery will be able to operate at a steady state without running out of charge, as was assumed to be the case. If these data were available, the validity of this assumption could be tested, and the model adjusted to reduce the level of demand met by each battery if necessary.

5.2.2 Model Components

Several additional assumptions were made when modelling batteries that are likely to result in their value being overstated. Each battery was assumed to be able to operate over the entire planning horizon, at 100% efficiency. In reality, however, batteries are likely to have a lifespan that is less than the time horizon under consideration, and will lose efficiency over time. During this project, it was found that it was possible to model these characteristics. However, it proved difficult to do so in a way that maintained the tractability of the resulting model, due to the need to track the movements of individual batteries, rather than inferring the total number of movements at each stage without regard to the locations each movement was between. Future research could focus on modelling these characteristics effectively.

Another key element of the modelling conducted was the use of a DC relaxation to model the network under consideration. The use of this approximation meant the operation of the network could be modelled using a set of linear constraints, which enforced that the amount of real power entering each location was equal to the amount leaving. In reality, however, the typical distribution network uses alternating current; modelling the flow of power in these networks requires consideration of both the active and reactive power in each line, as well as the voltage angles at each location (Kiszka and Wozabal, 2021). The use of a DC relaxation may result in operating configurations being obtained that are, in reality, infeasible.

To investigate the effects of using a DC relaxation, after the conclusion of the optimisation, a set of power flow equations could be solved to determine, for the base instance and each failure instance at each node of the scenario tree, whether the associated operating configuration is feasible. It is unclear, however, what should be done if one or more configurations are found to be infeasible. An alternative approach involves modelling the network as an AC network; this, however, would introduce nonlinearities into the problem, which would be expected to make it significantly more difficult to solve. Future research could investigate each of these approaches further.

The sole source of uncertainty considered as part of the modelling process was uncertainty in future demand for electricity. However, additional sources of uncertainty exist. For instance, batteries are a relatively new technology; there is uncertainty around their future cost and availability. In addition, there is likely to be some uncertainty around the future frequency and duration of line failures, as a result of extreme weather events caused by climate change. Future research could investigate the effects of incorporating these additional forms of uncertainty as part of the modelling process.

Finally, future research could further investigate the differences between solving the deterministic equivalent problem, and the use of Dantzig-Wolfe decomposition; it would be of interest to conduct testing to gain insight into additional factors that influence the relative performance of the two types of model. Ways in which models developed using GENE can be solved more efficiently could also be investigated. For instance, these models could be decomposed further; it is possible to formulate a set of subproblems that correspond to each failure instance at each node of the scenario tree (Singh, 2006). To increase the efficiency of the Dantzig-Wolfe algorithm during each iteration of the solve, each subproblem could be solved to a non-zero bound gap. This may, however, result in the Dantzig-Wolfe algorithm taking longer to converge due to the columns generated being of decreased quality; to address this, the bound gap of each subproblem could be reduced as the overall bound gap decreases. If any of these modifications were found to make significant improvements to solve times, this would provide avenues for more accurate modelling to be carried out.

A | GENE: User Input

GENE uses a spreadsheet-based input system, which aims to make the input of network and model parameters a straightforward process. Figure A.1 shows an example of the sheet that records information about the lines of the network under consideration. Similar sheets are used to record information on the locations that make up the network, the expansions that are available, and the frequency with which each line of the network is expected to fail. General model parameters, including the dimensions of the scenario tree and the process by which demand changes at each stage, are recorded using another sheet.

	A	B	C	D	E	F
1	Line ID	Name	From Node ID	To Node ID	Initial Capacity (kW)	Base Line
2	1	Subtree 4 Main Line	1	43	3363.36	1
3	2	Subtree 1 Main Line	1	44	3535.84	1
4	3	Subtree 6 Main Line	1	67	2371.6	1
5	4	Subtree 5 Main Line	1	103	2975.28	1
6	5	Subtree 2 Main Line	1	126	3913.14	1
7	6	Subtree 7 Main Line	1	182	2587.2	1
8	7	Subtree 8 Main Line	1	207	3913.14	1
9	8	Subtree 3 Main Line	1	213	3535.84	1

Figure A.1. A section of the sheet used to record data for the lines of the network.

	A	B
1	Setting	Value
2	Formulation	SAIDI Bound (Individual Locations)
3	Formulation Type	Deterministic Equivalent
4	Solver	Gurobi
5	Solve Type	Standard
6	Relative Tolerance (%)	0
7	SAIDI Bound (mins/year)	75

Figure A.2. A section of the sheet used to record model settings.

A final sheet is used to record information on the model formulation to be used, and additional solve settings; this is shown in Figure A.2. The desired measure of network reliability can be altered by changing the formulation type. Similarly, the type of solve to conduct can be changed; for example, it is possible to solve a set of models with varying levels of reliability, or to formulate and solve a deterministic model in which, at each stage, the demand at each location is assumed to be the maximum of the demand at the same stage of any scenario in the corresponding stochastic model.

B | GENE: Visualisation Module

The visualisation module in GENE communicates a number of important characteristics of an expansion plan, by way of several interactive maps. A representation of the scenario tree is also displayed; selecting a node in the tree updates each map to display information about the expansion plan in the associated state of the world.

Figure B.1 shows the scenario tree for an example model, alongside the upgrades map, which displays the status of each network component at the node of interest. Lines are coloured differently depending on whether they have been upgraded, and demand locations are coloured differently depending on whether they are housing any batteries.

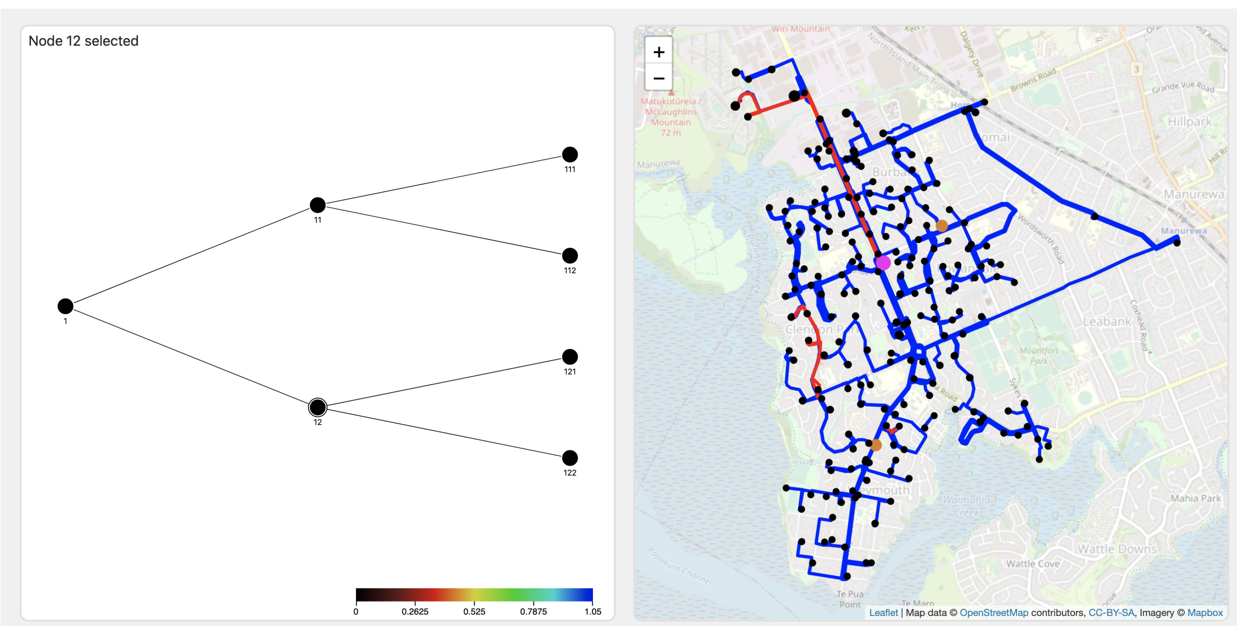

Figure B.1. A visualisation of the upgrades made at a node in the scenario tree as part of an expansion plan. Selecting a different node of the tree results in the investment plan in the associated state of the world being displayed.

The expected SAIDI at each location of the network can also be visualised (even if the formulation in use does not place any explicit restrictions on SAIDI, the SAIDI at each demand location can still be calculated). Figure B.2 shows an example of this visualisation. In the corresponding model, a bound on the SAIDI at each demand location was imposed;

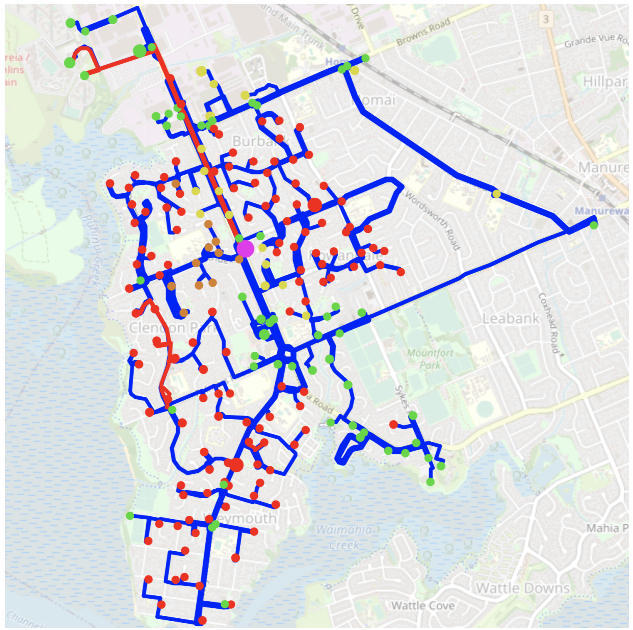

Figure B.2. A visualisation of the expected SAIDI at each location in the network. Red indicates a high SAIDI, while green indicates a low SAIDI.

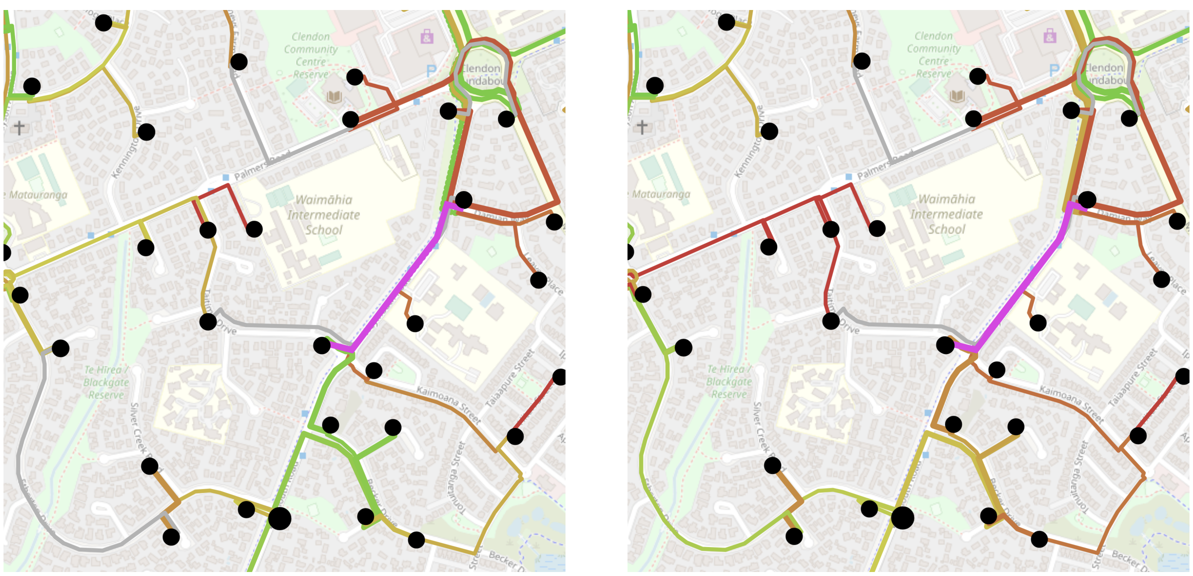

(a) Before failure **(b)** During failure

Figure B.3. The flows in a section of the network during normal operation (left) and while the purple line has failed (right). Lines coloured green are transporting large volumes of power, while lines coloured red are transporting volumes amounts of power. Inactive lines are grey.

it can be seen that the majority of locations in the network have a SAIDI that is relatively close to this tolerance, but that some areas in the far north and south east of the network experience comparatively few outages.

Finally, the operating configuration of the network in the base scenario, and under each failure scenario, can be visualised. Figure B.3 shows the flows in the network during normal operation of the network, and during a failure. The backfeeder in the bottom left of the picture is switched on to provide an alternative path for power to reach the affected locations in the bottom right. However, the re-routed flows to these locations are slightly lower than during normal operation of the network, so some outages will be experienced.

References

Aish, A. (2022). "Using JuDGE for Electricity Distribution Network Planning". University of Auckland. Honours Thesis.

Appelgren, L. H. (1971). "Integer Programming Methods for a Vessel Scheduling Problem". In: *Transportation Science* 5.1, pp. 64–78.

Birge, J. R. and Louveaux, F. (1997). *Introduction to Stochastic Programming*. New York, NY, USA: Springer-Verlag.

Dantzig, G. B. (1955). "Linear Programming under Uncertainty". In: *Management Science* 1.3-4, pp. 197–206.

Dantzig, G. B. and Wolfe, P. (1960). "Decomposition Principle for Linear Programs". In: *Operations Research* 8.1, pp. 101–111.

Defourny, B., Ernst, D., and Wehenkel, L. (Jan. 2011). "Multistage Stochastic Programming: A Scenario Tree Based Approach to Planning under Uncertainty". In: *Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions*.

Desaulniers, G., Desrosiers, J., and Solomon, M. (2005). *Column Generation*. Cahiers du GERAD. Springer US.

Dong, D. (2018). "Battery Optimization for Electricity Distribution Networks". University of Auckland. Honours Thesis.

Downward, A., Baucke, R., and Philpott, A. B. (2020). *JuDGE.jl: a Julia package for optimizing capacity expansion*. Electric Power Optimization Centre. URL: `optimization-online.org/?p=16779`.

Dunning, I., Huchette, J., and Lubin, M. (2017). "JuMP: A Modeling Language for Mathematical Optimization". In: *SIAM Review* 59.2, pp. 295–320.

Gurobi Optimization (2022). *Gurobi Optimizer Reference Manual*. URL: `www.gurobi.com`.

Kaut, M. and Wallace, S. W. (2007). "Evaluation of Scenario-Generation Methods for Stochastic Programming". In: *Pacific Journal of Optimization* 3.2, pp. 257–271.

Kiszka, A. and Wozabal, D. (2021). *Stochastic Dual Dynamic Programming for Optimal Power Flow Problems under Uncertainty*. Optimization Online. URL: `optimization-online.org/?p=18329`.

Luss, H. (1982). "Operations Research and Capacity Expansion Problems: A Survey". In: *Operations Research* 30.5, pp. 907–947.

Manne, A. S. (1967). *Investments for Capacity Expansion: Size, Location and Time-Phasing*. Woking and London: Unwin Brothers Ltd.

Rockafellar, R. T. and Wets, R. J. (1991). "Scenarios and Policy Aggregation in Optimization under Uncertainty". In: *Mathematics of Operations Research* 16.1, pp. 119–147.

Schachter, J. and Mancarella, P. (2016). "A critical review of Real Options thinking for valuing investment flexibility in Smart Grids and low carbon energy systems". In: *Renewable and Sustainable Energy Reviews* 56, pp. 261–271.

Singh, K. J. (2006). "Multistage stochastic capacity planning of survivable electricity distribution networks". University of Auckland. PhD Thesis.

Singh, K. J., Philpott, A. B., and Wood, R. K. (2004). *Column-Generation for Design of Survivable Electricity Distribution Networks*. URL: `researchgate.net/publication/235017288_Column-Generation_for_Design_of_Survivable_Electricity_Distribution_Networks`.

Singh, K. J., Philpott, A. B., and Wood, R. K. (2009). "Dantzig-Wolfe Decomposition for Solving Multistage Stochastic Capacity Planning Problems". In: *Operations Research* 57.5, pp. 1271–1286.

van Helsdingen, A. (2021). "Optimizing Electricity Distribution Networks Using JuDGE". University of Auckland. Honours Thesis.

Vector (2021). *Vector Electricity Asset Management Plan 2021-2031*. URL: `blob-static.vector.co.nz/blob/vector/media/vector2021/vec224-amp-2021-3031_310321.pdf`.

Wallace, S. W. (2000). "Decision Making Under Uncertainty: Is Sensitivity Analysis of Any Use?" In: *Operations Research* 48.1, pp. 20–25.

Wallace, S. W. (2010). "Stochastic programming and the option of doing it differently". In: *Annals of Operations Research* 177.1, pp. 3–8.