Published at
http://www.sciencedirect.com/science/article/pii/S095183201630076X
doi:10.1016/j.ress.2016.05.015

Assessing Survivability to Support Power Grid Investment Decisions

Anne Koziolek[a], Alberto Avritzer[b], Sindhu Suresh[b], Daniel S. Menasché[c,*], Morganna Diniz[d], Edmundo de Souza e Silva[c], Rosa M. Leão[c], Kishor Trivedi[e], Lucia Happe[a]

[a]*Department of Informatics, Karlsruhe Institute of Technology, Germany.*
[b]*Siemens Corporation, Corporate Technology, Princeton, NJ, USA.*
[c]*Federal University of Rio de Janeiro, RJ, Brazil.*
[d]*Federal University of the State of Rio de Janeiro, RJ, Brazil.*
[e]*Duke University, Durham, NC, USA.*

Abstract

The reliability of power grids has been subject of study for the past few decades. Traditionally, detailed models are used to assess how the system behaves after failures. Such models, based on power flow analysis and detailed simulations, yield accurate characterizations of the system under study. However, they fall short on scalability.

In this paper, we propose an efficient and scalable approach to assess the survivability of power systems. Our approach takes into account the phased-recovery of the system after a failure occurs. The proposed phased-recovery model yields metrics such as the expected accumulated energy not supplied between failure and full recovery. Leveraging the predictive power of the model, we use it as part of an optimization framework to assist in investment decisions. Given a budget and an initial circuit to be upgraded, we propose heuristics to sample the solution space in a principled way accounting for survivability-related metrics. We have evaluated the feasibility of this approach by applying it to the design of a benchmark distribution automation circuit. Our empirical results indicate that the combination of survivability and power flow analysis can provide meaningful investment decision support for power systems engineers.

1. Introduction

The reliability of power distribution systems has been widely studied for decades [1, 2, 3, 4, 5]. The fundamental problem consists of determining how the system behaves when faced with disruptions, and is generally tackled using detailed simulations and power flow analysis [6, 7, 8]. Different system characteristics, such as the workload and the availability of backup sources, are taken into account.

*Corresponding author: Federal University of Rio de Janeiro (UFRJ). POBOX 68.530. Zip code: 21941-590. Rio de Janeiro/RJ - Brazil. sadoc@land.ufrj.br (+55-21-998471550)

The result of detailed power flow analysis and simulations of power systems is an accurate assessment of how the system will behave under the considered configurations. Although the assessment is very precise, it falls short on scalability. The high computational costs preclude the analysis of a large number of configurations, and practitioners have to focus on the most likely or promising setups. Our goal is to propose metrics, models and heuristics to explore the state space in a principled, scalable and effective way that can guide equipment upgrades. The outcome of our analysis is a set of promising circuits, which should then be subject to more detailed investigation.

In this paper, we focus on survivability-related metrics and models to address the challenge of determining promising upgrades. Survivability-related metrics are computed assuming that the system starts in a failure state. They account for the phased-recovery of the system, i.e., how the system behaves since failure up to full recovery. By assuming that the system starts in a failure state, survivability-related metrics do not need to account for failure rates, which are typically orders of magnitude smaller than repair rates. This way, the proposed models to compute survivability-related metrics are amenable to easy solution, as we do not deal with stiff problems. In addition, the proposed survivability models do not need to capture the detailed state of all the system components. To increase scalability, we aggregate the sections of the power system in groups, depending on whether they are connected to a power source, and show how this simplifies analysis.

The proposed survivability-related models allow us to predict how the system behaves after changes. Such changes might be due to failures or due to investments. Consider a power utility company which has a given budget to invest in its circuits. The investments must account for equipment costs, as well as for the gains in terms of reliability and stability. To this end, we suggest heuristics which make use of survivability-related models and metrics to issue recommendations for investments.

This paper extends [9] by providing (1) a combined availability and survivability model that covers the whole failure and recovery behavior of the studied systems (Sections 3 and 4), (2) a formalization of the optimization problem under consideration (Sections 6.1–6.3), and (3) a detailed description of the heuristics used to solve the optimization problem (Section 6.4).

This work is based on our previous work on survivability assessment of power grids. In [10] we presented an analytical model to assess the survivability of distributed automation power grids. Then, we investigated the application of such model to scenarios with multiple failures [11], using historical data to parametrize the proposed model [12], and studied algorithmic aspects related to the network upgrade optimization problem [13]. Compared to our previous work, the focus of [9] and of the paper at hand is on the *combination of power flow analysis and survivability modeling* to achieve the optimal design of the distribution automation grid. The main contributions of the paper are summarized in two groups as follows.

Survivability model: We propose survivability-related metrics and models to capture the phased-recovery of the system from failure up to recovery. Power flow algorithms are used to parameterize the model. The model allows us to predict how the system will behave after failures and investments.

Survivability improvement: We leverage the predictive power of the proposed survivability model to issue recommendations on investments based on survivability-

related metrics. Given a budget, we consider heuristics to sample the solution space in a principled way, accounting for equipment costs and survivability gains. The efficient and scalable exploration of the solution space may be followed by detailed analysis of the most promising circuits.

The outline of this paper is as follows. In Section 2 we motivate the use of survivability metrics for the assessment of distributed automation grids. Then, in Section 3 we present real data on availability and survivability from a Brazilian utility. Section 4 derives a combined availability and survivability model and highlights its limitations, thus motivating the need for the aggregated phased recovery model described in Section 5. We present the optimization model used in this paper in Section 6. The analysis of our empirical results is presented in Section 7. Section 8 presents our conclusions and suggests future research.

2. Design Methodology

In this section we introduce an overview of the system and models considered in this paper.

2.1. System Overview

The power distribution network comprises a set of substations, renewable resources (e.g., Wind Power, Solar), load management (e.g., Demand Response), and devices associated with power distribution (e.g., lines, tap-changing transformers, capacitor banks, etc). The power distribution network is set up to guarantee that supply will equal demand, and that stability conditions are met. However, demand might go beyond predicted bounds, which might lead to instabilities. This occurs, for instance, due to failures of devices, incorrect load management, intentional attacks, or weather conditions (e.g., disruptions due to hurricane Sandy in the US [14, 12]). In this paper our focus is on the latter. Our goal is to issue investment recommendations to mitigate instabilities.

2.2. Terminology

Next, we introduce some basic terminology.

Availability model: captures failure and repair of resources. It accounts for the rate at which different components of the system fail.

Survivability model: phased-recovery model (or simply recovery model) characterizing the time-varying system behavior from failure up to full recovery [15].

Performance model: characterizes the performance of the system at different states. In this paper, performance is measured through the expected energy not supplied per time unit, and is captured through reward rates associated to each state.

Performability model: combination of performance model with availability and/or survivability model.

Power flow model: receives as input a set of load points and a circuit, and generates as output the angles and voltages (active and reactive power) associated to each section.

Violation matrices: matrices indicating for each section and for each load point whether the angle or voltages are beyond expected limits.

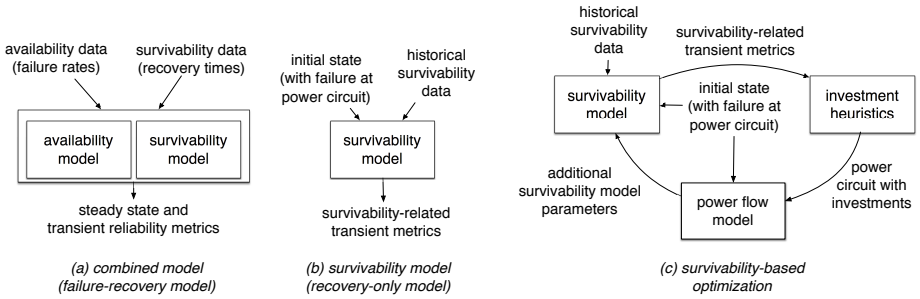

Figure 1: Methodology overview. In Sections 4, 5 and 6 we discuss the two models and the optimization approach described in Figures 1(a), 1(b) and 1(c), respectively.

2.3. Overview of Models

In this section we introduce the different models and optimization methodologies considered throughout this paper. We briefly describe the reliability model, survivability model and survivability-based optimization (Figure 1) and point out the fundamental motivations and goals associated to each of them.

The combined availability and survivability model presented in Figure 1(a) receives as input historical data about failure rates and recovery times of different system components. It yields steady state and transient reliability metrics, and is flexible to capture the interdependencies between failure and recovery of different system components. However, its solution might involve working with a stiff system as failure rates are typically orders of magnitude smaller than repair rates.

In Figure 1(b) we consider the survivability model. It characterizes the system from failure up to recovery, and does not involve failure rates. Instead, it receives as input the initial state of the system, presumed to be a failure state. It produces transient survivability metrics, from failure to recover. This is in contrast to the combined availability and survivability model, which can generate both transient and steady state metrics.

The proposed optimization framework makes use of the survivability model, which is amenable to efficient solution. The framework is depicted in Figure 1(c). Given the survivability-related metrics, investment recommendations are issued. Such recommendations are used to generate new power circuits, which are subject to power flow analysis. The results of the power flow analysis, in turn, are used to parameterize the survivability model, and the cycle repeats until either a) the budget is met or b) the survivability target is reached.

In Sections 4, 5 and 6 we further discuss the three models described in Figures 1(a), 1(b) and 1(c), respectively.

3. Availability and Survivability Data

In this section we report data collected from a Brazilian utility company, Light S.A., between 2009 and 2013. We present both availability related metrics, namely the number and type of failures, as well as survivability related metrics, namely the duration of the failures and the time required to fix the system. Our aim is to illustrate

the kind of data that is collected by real utilities, which can be used to parameterize the models described in the remainder of this paper.

The region considered is a rural neighborhood of the city of Rio de Janeiro, wherein reliability indices are worse than in the center of the city. The region is divided into three sections. Sections 1, 2 and 3 have lengths of 34.3 Km, 74.3 Km and 26 Km, and serve 1305, 482 and 131 clients, respectively.

3.1. Contingency Plan

The contingency plan adopted by Light S.A. in case of failure of the feeder consists of manually activating one of the two additional available feeders. In case of partial failures, the failed sections are isolated using a tie switch until they are fully repaired. In Section 4 we describe a model that captures this contingency plan accounting for multiple failures.

3.2. Availability Data

For each section, the utility company provides the number of failures, the time at which each failure event occurred and the time at which it was resolved. A summary of the data is presented in Table 1 (the section names are not identified due to privacy agreements). In Section 4 we will use data presented in Table 1 to parameterize the proposed availability model. Note that the survivability model presented in Section 5 does not make use of this data, as it only captures the dynamics of the system from repair up to recovery, not accounting for failure rates.

Table 1: Availability data provided by Light S.A. (number of failures per section)

year	section 1	section 2	section 3
2009	3	17	2
2010	13	28	8
2011	8	28	9
2012	11	42	19
2013	9	11	8
clients	1305	482	131
average failure rate (failures/year)	8.8	25.2	9.2

3.3. Survivability Data

Among survivability metrics, the utility provides the following ones, illustrated in Figure 2,

1. *PT (preparation time)*: time since the failure event is reported by the client until an utility team is prepared to visit the failure event place;

2. *DT (dislocation time)*: time taken by the team to go from the office to the failure event place;

3. *MT (mobilization time)*: sum of the two aforementioned times;

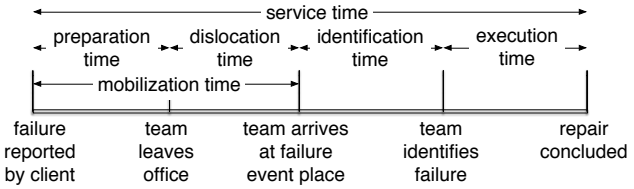

Figure 2: Recovery steps.

4. *IT (identification time)*: time taken to identify the failed components, once the team reaches the failure event place;

5. *ET (execution time)*: time taken to conclude the repair;

6. *ST (service time)*: sum of the three aforementioned times.

Table 2 illustrates the survivability metrics provided by Light S.A. for 12 out of the 17 failures that occurred between 2010 and 2011 in section 3 of the neighborhood of interest.

Table 2: Survivability data provided by Light S.A. Each line corresponds to a failure event. Time is measured in minutes.

year	PT	DT	MT	IT	ET	ST	cause
2010	524	23	547	0	20	567	falling tree
2010	15	37	52	0	25	77	falling branch
2010	75	3	78	67	0	145	falling branch
2010	163	40	203	0	25	228	falling branch
2010	13	92	105	2	22	129	falling branch
2010	16	0	16	0	14	30	falling tree
2011	0	1	1	0	0	1	unknown cause (quick repair)
2011	5	20	25	0	155	180	falling tree
2011	0	0	0	0	298	298	component failure
2011	47	5	52	2	18	72	unknown cause (slow repair)
2011	29	12	41	3	27	71	falling branch
2011	0	5	5	2	593	600	emergency maintenance

The survivability model presented in Section 5 is partially inspired by data provided in Table 2. According to Table 2, the mean service time was 199.80 minutes. In Sections 4 and 5, we will use this value to parameterize the proposed availability and survivability models. Using additional data from Table 2 to fine tune the survivability model is subject for future work.

4. General Availability and Survivability Model

In this section we present a combined availability and survivability model. The combined model encompasses both availability-related metrics such as the mean time

between failures, as well as survivability-related metrics, such as the mean time to recover after each failure.

4.1. Model Description

The power distribution network is divided into N sections, and each section can be in one of three states: active (A), failed (F) and not supplied due to a failure in another section (U). We assume that a backup feeder is always available, therefore energy is not supplied to a section if and only if the section is failed or if it is unreachable by the main and backup feeders. Note that each section can be in one of three possible states but once the failed sections are determined, the state of the others is fully determined (see Figure 3). Therefore, the cardinality of the state space equals 2^N (in the example of Figure 3, four out of the eight states are illustrated).

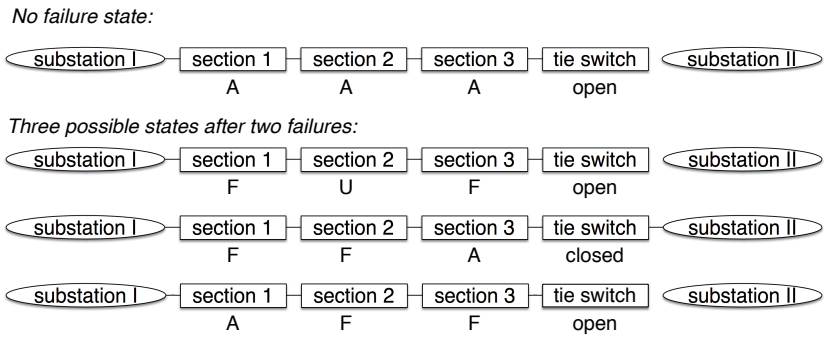

Figure 3: Availability model. Available, failed, and not supplied sections are denoted by A, F and U. From the initial state, after two sections fail there are three possible states. In the middle state, one subsection is fed by substation II (backup substation) through the closed tie switch.

The model is implemented in a modular fashion, using the Tangram-II tool [16]. Each substation, tie switch and section is characterized using an object. Objects can be easily replicated to construct a topology and they interact with their neighbors by sending and receiving messages, which are triggered by failure and repair events (refer to [16] for details and the source code of the model used in this Section).

Let ω_i be the expected energy not supplied at state i. Let $\mathrm{ENS}(t)$ be the energy not supplied at time t. Then,

$$\mathbb{E}[\mathrm{ENS}(t)] = \sum_{i=1}^{|\Omega|} \pi_i(t)\omega_i \tag{1}$$

where $\pi_i(t)$ is the probability that the system is at state i at time t, and Ω denotes the state space.

4.2. Model Parameterization

Next, we consider a simple line topology comprising three sections to capture the setup introduced in Section 3. Section 1 is connected to the main feeder, and section 3 is connected to the backup feeder. Section 2 has connections to sections 1 and 3. In

the setup of Section 3, the model requires eight parameters (the loads and the failure rates of the three sections, as well as the repair rate, which is assumed to be the same across all sections). The maximum load of section 3 was 0.2 MVA in the summer of 2011/2012. As the maximum load of the other sections was not provided by the utility company, we used the load of section 3 to approximate the loads of sections 1 and 2, assuming that the loads are proportional to the number of clients per section. The loads of sections 1, 2 and 3 were approximated as 2.0 MVA, 0.7 MVA and 0.2 MVA, respectively. The failure rates of sections 1, 2 and 3 are 8.8, 25.2 and 9.2 failures/year, respectively (see Table 1). The mean repair time is approximated as 4 hours (see Section 3.3).

4.3. Numerical Results

Next, we present steady state and transient metrics obtained using the proposed model. Our goals are to a) illustrate the impact of initial conditions on the transient metrics, b) numerically indicate that as failure rates are much lower than repair rates, the transient metrics in the horizon of one week might be very different than the steady state ones and c) motivate the use of survivability-related metrics.

To illustrate the importance of accounting for transient metrics, consider the accumulated average energy not supplied during the first week after a failure of section 1 in the model introduced above. In steady state, the expected energy not supplied during one week is 3.19 MVA. In contrast, assuming that the system starts from a failure state, where section 1 has failed, the expected energy not supplied during the week is 11.09 MVA. If all sections have initially failed, the expected energy not supplied during the week is 15.12 MVA.

To further illustrate the importance of provisioning the network accounting for the transient impact failures, Figure 4 (obtained with the proposed availability model) shows how the expected energy not supplied per hour varies as a function of the observation time, for different initial conditions. The initial condition of the system after a failure, in turn, critically depends on the investments applied to the power grid. The expected energy not supplied per unit time in steady state is negligible (in Fig. 4, it is represented by an horizontal line, almost indistinguishable from the one corresponding to a system that initiates without failures). It takes more than seven hours for the energy not supplied to converge to its steady state value. Focusing on the transient analysis, one can obtain more insight about how the system behaves when investments are more impactful, helping designers to issue investment recommendations.

4.4. Features and Limitations

The model proposed in this section accounts for failures and repairs in an integrated fashion, allowing for failures to occur while repair is taking place. Being inspired by the contingency plan introduced in Section 3.1, the model captures how the system evolves over time when it is functioning and under failures. In particular, it naturally accounts for multiple failures.

Nonetheless, the model has some limitations concerning the size of the state space and the accuracy of its numerical solution. The size of the state space grows exponentially with respect to the number of devices in the topology, as the state of each device

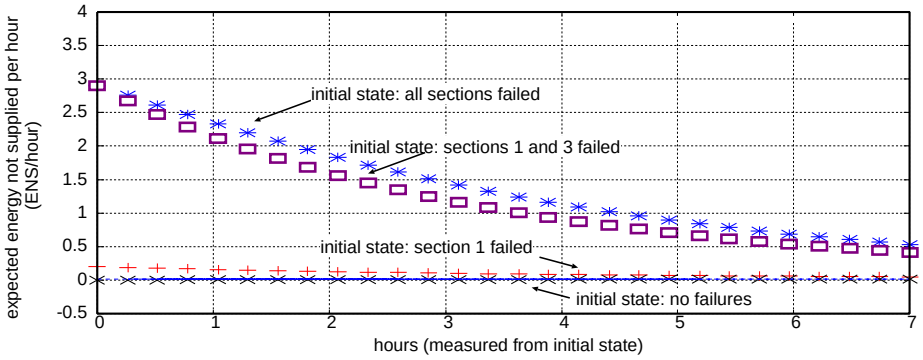

Figure 4: Availability results obtained with combined availability-survivability model: depending on the initial condition, it takes more than 7 hours for the expected energy not supplied per hour to approach the steady state value of 0.017 MVA (Mega-volt ampere).

is characterized by a separate state variable (see Section 4.1). In addition, as failure rates and repair rates have different orders of magnitude, the solution of the model might involve working with a stiff system.

Note that investments on the power grid can significantly impact the initial condition of the system after a failure. As indicated in Figure 4, to distinguish between investment options it is imperative to analyze the system during the first hours after a failure-event occurs, which further motivates the use of survivability-related models and metrics. Whereas survivability-related metrics might drastically vary as a function of investments, average metrics computed on steady-state may show small or negligible differences as a function of investments, being dominated by the very small failure rates. The lines corresponding to a system that initiates without failures and one that initiates in steady state are almost indistinguishable in Figure 4.

In the upcoming section we present an aggregated survivability model, henceforth referred to simply as *survivability model*. In contrast to the model presented in this section, the survivability model has constant state space size, irrespective of the topology. In addition, as the survivability model accounts for events that occur between a failure and a repair, it solely captures repair rates, avoiding numerical instabilities and allowing for a more refined characterization of the recovery process. In Section 6 we use the simplified survivability model for distribution automation optimization.

5. Survivability Model

Next, we describe the survivability model and metrics that we use as part of our methodology for power grid optimization. Our approach is very general, and can be coupled to any phased recovery model. We describe the particular instantiation of the phased recovery model from [10], illustrated in Figure 5.

5.1. Phased-Recovery Model

Recall that the power distribution network is split into sections. After section i is isolated, the system goes to one of three states, depending on whether there is enough

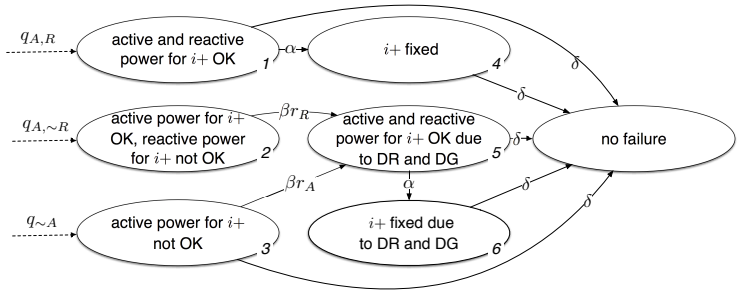

Figure 5: Phased recovery model

active and reactive backup power to supply energy for the sections that were indirectly affected by the failure (dotted lines in Figure 5). Such sections are referred to as the upstream sections of section i, and also denoted as $i+$. Note that the topology of a power distribution network is typically a tree, therefore the downstream and upstream of section i are well defined. In contrast, this is not necessarily the case, for instance, for power transmission networks.

After a failure occurs, the system goes to states 1, 2 or 3 depending on whether there is enough active and reactive power to supply for the upstream sections. In state 1 there is enough active and reactive power, in state 2 there is enough active but not reactive power and in state 3 there is not enough active power to supply for the upstream sections. At state 1, the system is amenable to automatic recovery, which occurs with rate α. After automatic recovery, the system transitions to state 6. Otherwise, demand response application needs to be activated in order to reduce the load in the upstream sections. The demand response application takes mean time $1/\beta$ min to be activated, and is effective to reduce the active and reactive loads with probabilities r_A and r_R, respectively. In case the demand response application effectively reduces the load, the system transitions to state 4 and is amenable to automatic recovery. At all states the system is amenable to manual repair, taking mean time $1/\delta$ hours. Note that although upstream sections may be automatically repaired, we assume that a truck is needed to fix the failed section i. After full recovery, the system transitions to state 0 (Figure 5).

We let $1/\delta = 4$ hours which roughly corresponds to the mean manual repair time of 3.33 hours presented in Section 3.3. In the remainder of this paper, based on expert knowledge we let $\alpha = 2\ \mathrm{min}^{-1}$ and $1/\beta = 15$ min.

Let $q_{A,R}$, $q_{A,\sim R}$ and $q_{\sim A}$ be the probabilities that the system transitions to states 1, 2 and 3, respectively, after a failure. These probabilities play a key role in our methodology, and depend on the circuit topology, on the amount of investment in distributed generation and on the load.

To each state state s_j, $j = 1, 2, \ldots, 6$, we associate its corresponding reward rate ω_j. In this paper, the reward rate associated to state j characterizes the energy not supplied at that state per time unit. Solving the Markov chain model, survivability related metrics such as the AENS in kWh can be computed [10, 17].

Let $\mathrm{AENS}^{(\tau)}$ be the accumulated energy not supplied by time τ after a failure. Recall that $\pi_j(t)$ is the probability that the system is at state j at time t. Then,

$\text{AENS}^{(\tau)} = \int_{t=0}^{\tau} m_j(t)dt$, where $m_j(t)$ is the instantaneous energy not supplied at time t, and its expected value is given by

$$\mathbb{E}[\text{AENS}^{(\tau)}] = \int_{t=0}^{\tau} \sum_{j=1}^{|\Omega|} \pi_j(t)\omega_j dt \tag{2}$$

Let AENS be the accumulated energy not supplied between failure and full system recovery. Its expected value is easily obtained from (2) letting $\tau = \infty$,

$$\mathbb{E}[\text{AENS}] = \sum_{j=1}^{|\Omega|} \tilde{\pi}_j \omega_j \tag{3}$$

where $\tilde{\pi}_j$ is the fraction of time the system remains at state j between failure and full recovery. The expected accumulated energy not supplied between failure and full recovery is the key survivability-related metric considered in this paper.

In this paper, we assume that the phased-recovery from a failure can be modeled as a homogeneous continuous time Markov chain. Extensions to accommodate non-exponential sojourn times have been considered in related work [12, 18].

5.2. Parameterization of the Survivability Model

Recall that the key input parameters of the survivability model are (1) the probabilities that the backup and distributed generation sources suffice to supply active and reactive energy for the sections that are affected by a failure, $q_{A,R}$, $q_{A,\sim R}$,$q_{\sim A}$, (2) the probabilities that demand response is effective, r_A and r_R, and (3) the reward rates at each of the model states. The methodology to obtain these three sets of parameters is described in sections 5.2.3, 5.2.4, and 5.2.5, respectively, and is a function of the load data and violation matrices that are described next in sections 5.2.1 and 5.2.2.

5.2.1. Workload

The inputs to our model parameterization are:

1. load data, e.g., a 24 hour load profile at 15-minute intervals (i.e., 96 load points);
2. a model of the power circuit topology after the failure and after fault isolation;
3. the voltage and angle at each section, obtained from the power flow model.

Load data. We consider a time series which characterizes how load varies for different times of the day over different sections. We assume that the day is divided into 15 minutes time slots. Each slot of 15 minutes is referred to as a given *time of the day* or *load point*. The expected active power load at section i and load point j is then denoted by $L_a(i, j)$. The corresponding expected reactive power load is denoted by $L_r(i, j)$. Let N be the number of sections in the circuit. Let t be the number of load points used to parametrize the survivability model. Except otherwise noted, let $t = 96$. We denote a matrix of scalars, of N lines and t columns, by $\mathbb{R}^{N\times t}$. Then, the workload is characterized by two matrices $L_a \in \mathbb{R}^{N\times t}$ in kW and $L_r \in \mathbb{R}^{N\times t}$ in kVar (see Table 4 for notation).

Power circuit model. The power circuit model comprises the topology of the circuit and the active and reactive power generated at each section, for each time of the day. Let $G_a(i, j)$ and $G_r(i, j)$ be the maximum expected active and reactive power generated at section i and load point j. Then, the the network capacity is characterized by two matrices $G_a \in \mathbb{R}^{N\times t}$ in kW and $G_r \in \mathbb{R}^{N\times t}$ in kVar.

Voltages and angles. Given the load data and the power circuit model, power flow equations [19] are used to obtain voltages and angles at each section, for each time of the day. Voltages and angles yield violation matrices, as described in the upcoming section.

5.2.2. Violation Matrices

The *violation matrices* $M_{\sim A}$, $M_{A,\sim R}$ and $M_{A,R}$ are determined by solving the power flow equations and characterize the chances that, after the failure of a tagged section (or a set of tagged sections), active and reactive power will suffice to supply the remaining sections. A violation matrix is a matrix $M \in \{0, 1\}^{N\times t}$.

In what follows, we will describe how to compute $M_{\sim A}$. Each entry of $M_{\sim A}$ characterizes whether there is a violation of active power at a given section at a given instant of the day, after the failure of a tagged section. The element in line i and column j of the violation matrix $M_{\sim A}$ is defined as a function of the output of the power flow algorithm. $M_{\sim A}(i, j) = 1$ if there is a violation of active power, i.e., the angle at section i and load point j is beyond its expected limits by more than 10%, and 0 otherwise. Matrix $M_{A,\sim R}$ (resp., $M_{A,R}$), where each entry characterizes whether there is a violation of reactive power at a given section at a given instant of the day while active power has no violation (resp., of there being no violation at all), can be defined and computed similarly.

5.2.3. Parametrization of Sufficiency of Backup Power

Let $\rho_{\sim A}$ be the number of section-load point pairs for which the load is large enough to cause the angle to be beyond desired bounds (Table 3),

$$\rho_{\sim A} = \sum_{i=1}^{N}\sum_{j=0}^{t} M_{\sim A}(i, j) \tag{4}$$

In this paper, we assume that the fraction of section-load point pairs for which the violation matrix entry is equal to 1 is a surrogate for the probability that there will be a disruption in the network. Then,

$$q_{\sim A} = \rho_{\sim A}/(Nt) \tag{5}$$

$q_{A,R}$ and $\rho_{A,R}$ (resp., $q_{A,\sim R}$ and $\rho_{A,\sim R}$) can be similarly defined and computed as a function of $M_{A,R}$ (resp., $M_{A,\sim R}$). Note that

$$q_{A,R} + q_{A,\sim R} + q_{\sim A} = 1 \tag{6}$$

$$\rho_{A,R} + \rho_{A,\sim R} + \rho_{\sim A} = Nt \tag{7}$$

$$\tag{8}$$

Table 3: Angle and voltage violations

Variable	Description	Comments
$\rho_{\sim A}$	number of angle violations (active power violations)	favors increase in active power (P)
$\rho_{A,\sim R}$	number of voltage violations (reactive power violations)	favors increase in reactive power (Q)

5.2.4. Parametrization of Effectiveness of Demand Response

To compute the probabilities r_A and r_R that demand response is effective to cope with a lack of active and reactive power, respectively, we reduce the load by the amount amenable to demand response. The loads amenable to demand response in section i at load point j are denoted by $D_a(i,j)$ in kW for the active load and $D_r(i,j)$ in kVar for reactive load. Let D_a and D_r denote the resulting matrices in $\mathbb{R}^{N\times t}$. Then, the new load after demand response has been called for is denoted by $L_a^{(dr)} = L_a - D_a$ for active load and $L_r^{(dr)} = L_r - D_r$ for reactive load. For these new load data, we solve the power flow again and obtain new violation matrices $M_{\sim A}^{(dr)}$, $M_{A,\sim R}^{(dr)}$, and $M_A^{(dr)}$. Let $\rho_{\sim A}^{(dr)}$ be the number of load points and sections that still have violations, after demand response is called for. $\rho_{\sim A}^{(dr)}$ is computed using (4), replacing $M_{\sim A}$ by $M_{\sim A}^{(dr)}$. Similarly, $\rho_{A,\sim R}^{(dr)}$ is computed using (4), replacing $M_{\sim A}$ by $M_{A,\sim R}^{(dr)}$.

The effectiveness of demand response equals one minus the ratio of the number of scenarios for which the circuit is unstable after reducing the load when demand response is called for, divided by the number of scenarios at which the circuit was unstable at first place,

$$r_A = 1 - \rho_{\sim A}^{(dr)}/\rho_{\sim A}, \quad r_R = 1 - \rho_{A,\sim R}^{(dr)}/\rho_{A,\sim R} \tag{9}$$

5.2.5. Parametrization of Reward Rates

The reward rates at each of the model states are obtained from the load data and the state of the individual sections. We refer to the average energy supplied (resp., not supplied) per time unit as $\overline{ES}$ (resp., $\overline{ENS}$), and consider a failure at section i.

In *states 1-3*, $\overline{ES} = 0$ and $\overline{ENS}$ equals the sum of the active load of all sections. In *state 4*, $\overline{ENS}$ decreases by the amount of load amenable to demand response. In *state 5*, upstream sections $i+$ have been recovered while the demand response program is active, so $\overline{ES}$ is the active load of sections $i+$ reduced by the amount of load amenable to demand response and $\overline{ENS}$ is the active load of section i. In *state 6*, all upstream sections have been recovered by the backup substation, so $\overline{ES}$ equals the active load of sections $i+$ and $\overline{ENS}$ equals the active load of section i. Finally, in *state 0* the main substation has been recovered as well, so that $\overline{ES}$ is the sum of the load of all sections and $\overline{ENS} = 0$.

Note that $\overline{ES}$ and $\overline{ENS}$ are active power metrics, and do not directly capture the effects of reactive power. This is because active power is what is billed to customers. So, the loss of revenue in case of failure is, as a first order approximation, the active energy not supplied. Nonetheless, reactive power is also taken into account in our

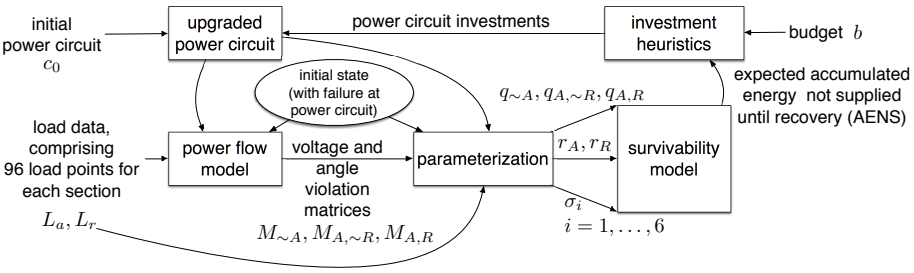

Figure 6: Optimization approach

survivability model. This is because the power flow algorithms capture the fact that without reactive power the circuit will be unstable and no active power will be supplied.

5.2.6. Model Solution and Scalability

The expected energy not supplied from failure up to recovery (see eq. (3)) can be computed in $O(|\Omega|^3)$ for a model with state space cardinality $|\Omega|$, using standard techniques such as the Grassman-Taksar-Heyman (GTH) solution method [20]. Let Ω_S and Ω_C be the state space of the survivability model proposed in this section and of the combined model described in Section 4.1, respectively. Note that $|\Omega_S|$ is independent of the topology (e.g., in Figure 5 we have $|\Omega_S| = 7$). The cardinality of the state space of the model proposed in Section 4.1, in contrast, equals $|\Omega_C| = 2^N$. Whereas the survivability model captures topology through the parameterization of reward rates, in the combined model the topology is captured directly in the state space characterization, which allows for a more detailed but less scalable analysis.

6. Optimization Approach

This section describes the proposed optimization problem (Section 6.1), its solution space (Section 6.2), formulation (Section 6.3) and the optimization algorithm (Section 6.4) for holistically supporting investment decisions for distribution automation network designs. This section shows how the presented combination of survivability analysis and power flow analysis can help to make well-informed investment decisions.

6.1. Optimization Problem

The optimization goal is to find a set of optimal investments for this initial distribution grid. Let Γ be the set of possible investment types. In this work, the investment types considered are (1) the addition of distributed generators of various types (such as hydro, biomass, or wind generators), which generate active and reactive power, (2) the investment into demand response infrastructure to achieve a more flexible demand, and (3) the procurement of DStatcoms, which provide reactive power. In the remainder of this paper, we let $\Gamma = \{\text{biomass, wind, solar, hydro, demand response, DStatcom}\}$. Table 5 summarizes the different available alternatives.

Each device v is of type $\gamma(v) \in \Gamma$. An investment in a device of type γ can be implemented using different models of devices, that have different costs and capacities.

Table 4: Table of notation

Load data	
N	number of sections in the circuit
t	number of load points per section
$G_a(i,j)$	maximum active power generated at section i and load point j
$G_r(i,j)$	maximum reactive power generated at section i and load point j
$L_a(i,j)$	active load at section i and load point j
L_a	$N \times t$ matrix of active loads
L_r	$N \times t$ matrix of reactive loads
$L_a^{(dr)}$	$N \times t$ matrix of active loads, after demand response is called for
$L_r^{(dr)}$	$N \times t$ matrix of reactive loads, after demand response is called for
Violation matrices and metrics	
$M_{\sim A}$	$N \times t$ matrix of active power violations
$M_{A,\sim R}$	$N \times t$ matrix of reactive power violations
$M_{A,R}$	$N \times t$ matrix of no violations
$\rho_{\sim A}$	number of load points (over all sections) for which load yields active power violation
$q_{\sim A}$	fraction of load points (over all sections) for which load yields active power violation
Quantities involved in the optimization	
Variable	Description
c	initial circuit
(v,i)	investment of deploying device v in section i
σ	$\sigma = \{(v_1,i_1) \cup \ldots (v_m,i_m)\}$ is the multiset of investments
$c' = c \cup \sigma$	current circuit
Device data	Description
Γ	device types, {biomass, wind, solar, hydro, dem/resp, dstatcom}
θ_γ	number of options of devices of type $\gamma \in \Gamma$
$\gamma(v)$	type of device v
$\theta(v)$	option of device v
$\boldsymbol{g}_a(v,j)$	expected generated active power, at load point j, $j = 1,\ldots,t$
$\boldsymbol{g}_r(v,j)$	expected generated reactive power, at load point j, $j = 1,\ldots,t$
$\boldsymbol{l}_a(v,j)$	expected maximum reduction of active load, due to demand response, at load point j, $j = 1,\ldots,t$
$\boldsymbol{l}_r(v,j)$	expected maximum reduction of reactive load, due to demand response, at load point j, $j = 1,\ldots,t$
$\textsc{Cost}(v)$	cost of device v
Parameter	Description
$\hat{S}$	(section constraint) maximum number of devices of given type per section
$\hat{G}$	(global constraint) maximum number of devices of given type per circuit
τ	survivability requirement (maximum tolerable accumulated ENS)
Metric of interest	Description
$\text{AENS}(c)$	expected accumulated energy not supplied, since failure up to full repair, for circuit c
$\textsc{Cost}(c)$	cost of circuit c

Table 5: Cost for active and reactive power investment

Equip. Type γ	Biomass		Wind		Solar		Small Hydro		Demand Response		DStatcom	
Power Type	active + reactive		active + reactive		active + reactive		active + reactive		active + reactive		reactive	
Equip. Options	Power in kW	Total inv. cost in $	Power in kW	Total inv. cost in $	Power in kW	Total inv. cost in $	Power in kW	Total inv. cost in $	Power in kW	Total inv. cost in $	Power in kVar	Total inv. cost in $
1	2	10000	2	16000	2	2600	2	4804	2	330	2	110
2	5	25000	5	40000	5	5500	5	12010	5	825	5	275
3	10	50000	10	60000	10	14000	10	24020	10	1650	10	550
4	30	150000	30	180000	30	57000	30	72060	30	4950	30	1650
5	100	500000	100	350000	100	200000	100	240200	100	16500	100	5500
6	250	750000	250	875000	250	350000	250	600500	250	41250	250	13750
7	500	1500000	500	1750000	500	300000	500	1201000	500	82500	500	27500
8	1000	3000000	1000	2000000	1000	800000	1000	2402000	1000	165000	1000	55000

Let θ_γ be the number of options available for devices of type γ. Let $\theta(v)$ be the option associated to device v. In Table 5, we have $\theta_\gamma = 8$ for all $\gamma \in \Gamma$. Each option is characterized by five attributes. Recall that we divide the day into t intervals, and each interval is associated to a load point. The first four attributes are $1 \times t$ vectors. Vectors are denoted with bold faces. The four vectors are (i) the expected generated active power $\boldsymbol{g}_a(v)$, (ii) the expected generated reactive power $\boldsymbol{g}_r(v)$, (iii) the expected maximum reduction of active load $\boldsymbol{l}_a(v)$, and (iv) the expected maximum reduction of reactive load $\boldsymbol{l}_r(v)$; the fifth attribute (v) is a scalar $\textsc{Cost}(v)$ that corresponds to the device cost (Table 5). Note that the five attributes exclusively depend on the considered device. To simplify presentation, we assume that none of the attributes depend on the place where the device is installed.

Let $\boldsymbol{g}_a(v,m)$ be the m-th entry of $\boldsymbol{g}_a(v)$, $1 \leq m \leq t$ ($\boldsymbol{g}_r(v,m)$, $\boldsymbol{l}_a(v,m)$ and $\boldsymbol{l}_r(v,m)$ can be similarly defined). Let $\overline{\boldsymbol{g}}_a(v)$ (resp., $\overline{\boldsymbol{g}}_r(v)$) be the expected active (resp., reactive) power generated averaged over all load points,

$$\overline{\boldsymbol{g}}_a(v) = \frac{\sum_{m=1}^{t} \boldsymbol{g}_a(v,m)}{t}, \quad \overline{\boldsymbol{g}}_r(v) = \frac{\sum_{m=1}^{t} \boldsymbol{g}_r(v,m)}{t},$$

$\overline{\boldsymbol{l}}_a(v)$ and $\overline{\boldsymbol{l}}_r(v)$ can be similarly defined.

To model the effects of an investment, consider an upgrade which consists of adding device v at section i. The active (resp., reactive) power generated at section i and load point j are increased by $\boldsymbol{g}_a(v,j)$ (resp., $\boldsymbol{g}_r(v,j)$),

$$G_a(i,j) \leftarrow G_a(i,j) + \boldsymbol{g}_a(v,j) \tag{10}$$
$$G_r(i,j) \leftarrow G_r(i,j) + \boldsymbol{g}_r(v,j) \tag{11}$$

for $j = 1, 2, \ldots, t$. In addition, the active and reactive load amenable to demand response are updated as follows

$$D_a(i,j) \leftarrow \max(L_a(i,j), D_a(i,j) + \boldsymbol{l}_a(v,j)) \tag{12}$$
$$D_r(i,j) \leftarrow \max(L_r(i,j), D_r(i,j) + \boldsymbol{l}_r(v,j)) \tag{13}$$

for $j = 1, 2, \ldots, t$. This leads to an updated load and generation profile and thus to a new candidate distribution grid model.

6.2. Solution Space

Let v_j and i_j be the j-th installed device and the section at which the device is installed, respectively. Together, (v_j, i_j) is referred to as the j-th investment. A *distribution grid model* is a multiset of investments, $\{(v_1, i_1), ..., (v_n, i_n)\}$, which are applied to an initial model c. Note that the order of investment options has no effect.

Given an initial circuit c the problem solution space is denoted by Ω_c. An element $\sigma \in \Omega_c$ is a distribution grid model. Let $\mathcal{J}_c(\gamma, i, \sigma)$ be the number of devices of type $\gamma \in \Gamma$ at section i of circuit $\sigma \in \Omega_c$,

$$\mathcal{J}_c(\gamma', i, \sigma) = |\{v|(v, i) \in \sigma \text{ and } \gamma(v) = \gamma'\}| \tag{14}$$

Let $\mathcal{J}_c(\gamma, \sigma)$ be the total number of devices of type $\gamma \in \Gamma$ over all sections at circuit $\sigma \in \Omega_c$,

$$\mathcal{J}_c(\gamma, \sigma) = \sum_{i=1}^{N} \mathcal{J}_c(\gamma, i, \sigma) \tag{15}$$

Power engineers can set a limit on the number of devices of a given type that can be deployed at each section ($\hat{S}$) and globally ($\hat{G}$). Thus, limitations on the availability of generator types e.g., due to geography or resident acceptance, can be reflected. To account for such constraints, an element $\sigma \in \Omega_c$ must satisfy,

$$\mathcal{J}_c(\gamma, i, \sigma) \leq \hat{S},\ \forall \gamma \in \Gamma, i = 1, \ldots, N \tag{16}$$

$$\mathcal{J}_c(\gamma, \sigma) \leq \hat{G},\ \forall \gamma \in \Gamma \tag{17}$$

The state space cardinality $|\Omega_c|$ is derived in Appendix A.

6.3. Problem Formulation

Based on this model, we formulate a multi-criteria optimization problem as follows. The two objectives of our optimization are:

1. Maximize survivability by minimizing the expected accumulated energy not supplied between failure and recovery. We denote the expected accumulated energy not supplied (AENS) of a candidate circuits as $\text{AENS}(c)$.
2. Minimize costs. The cost $\text{COST}(c)$ of a proposed circuit c is the sum of the costs of the applied devices $v_1, ..., v_n$ in c: $\text{COST}(c) = \sum_{j=1}^{n} \text{COST}(v_j)$.

While trading-off survivability-related gains and costs, $\text{AENS}(c)$ and $\text{COST}(c)$ are treated as incommensurable objectives, i.e., the goal of the multiobjective optimization is to find *Pareto efficient solutions*. Let b be the overall cost budget. Then, we define the multi-criteria optimization problem to be solved considering upper bound of possible investment budget b as

$$\text{MULTIOBJECTIVEPROBLEM}: \quad \min\left(\text{AENS}(c'), \text{COST}(c')\right) \tag{18}$$

$$\text{subject to } \text{COST}(c') < b \tag{19}$$

$$c' \in \Omega_c \tag{20}$$

where $\min(\text{AENS}(c'), \text{COST}(c'))$ is the set of ordered pairs $(\text{AENS}(c'), \text{COST}(c'))$ in the Pareto front.

The complexity of the optimization problem depends on (1) the complexity of the solution space and (2) the complexity of the evaluation functions. Here, the survivability evaluation function $\text{AENS}(c')$ (including power flow solution and survivability model solution) is a set of equations that is commonly solved by iterative algorithms [21]. The cost evaluation function $\text{COST}(c')$ is a simple additive function.

Alternatively, the heuristics proposed in Section 6.4 also yield approximate solutions to the following single objetive optimization problem,

$$\begin{aligned} \text{SINGLEOBJECTIVEPROBLEM}: \quad & \min \text{AENS}(c') && (21)\\ & \text{subject to } \text{COST}(c') < b && (22)\\ & c' \in \Omega_c && (23) \end{aligned}$$

6.4. Optimization Algorithm and Heuristics

To find approximated solutions to the aforementioned optimization problem, we use heuristic approaches based on the power flow results, as sketched in Figure 6 and shown in more detail in algorithm 1. The key goal of the heuristics is to sample the state space in a principled way, to find good candidates for detailed analysis. To summarize, the algorithm starts with the current circuit as the initial circuit candidate. It guesses the best applicable investment based on heuristics and applies it to the current distribution grid model. If an investment option does not improve the survivability metric AENS (i.e., the survivability objective function), it chooses another investment option. If an investment option improves the survivability, it is accepted and the next investment options are tried until the survivability requirement is met or the budget is used up.

The algorithm uses three types of heuristics.

6.4.1. Deciding Between Active and Reactive Power Investment

The type of power to add is selected based on the number of violations in the power flow results. If more voltage violations are observed, investments for reactive power should be chosen. If more angle violations are observed, investments for active power should be chosen. Given a candidate circuit c,

$$\text{POWERTYPEHEURISTIC}(c) = \begin{cases} \text{active}, & \text{if } \rho_{\sim A}(c) \geq \rho_{A,\sim R}(c) \\ \text{reactive}, & \text{otherwise} \end{cases} \quad (24)$$

6.4.2. Deciding In Which Section to Invest

The algorithm decides in which section to invest based on the number of voltage or angle violations in the power flow. The section with most violations is selected. If there is a tie, one of the sections with maximum number of violations is randomly selected.

Given a candidate circuit c, while targeting power improvements of type p ($p \in \{active, reactive\}$), let $\text{SECTIONHEURISTIC}(c, p)$ be the section in which to invest. If $p = active$ (resp., $p = reactive$), $\text{SECTIONHEURISTIC}(c, p)$ is a section selected uniformly at random from the set of sections which maximize $\sum_{1 \leq j \leq t} M_{\sim A}(i, j)$ (resp., $\sum_{1 \leq j \leq t} M_{A,\sim R}(i, j)$).

6.4.3. Deciding Which Device to Buy

The device to place in the selected section is selected based on the efficiency of the investments. Here, we formulated three heuristics which reflect investment strategies in the real world. For the heuristics, let I denote the current set of devices to choose from, and let p denote the selected power type. Furthermore, let $\textsc{Gain}(v, p)$ be the gain in terms of the amount of power of type p generated and the amount of power of type p that can be potentially reduced by investing on device v, i.e.,

$$\textsc{Gain}(v,p) = \begin{cases} \overline{\boldsymbol{g}}_a(v) + \overline{\boldsymbol{l}}_a(v), & \text{if } p = \text{active} \\ \overline{\boldsymbol{g}}_r(v) + \overline{\boldsymbol{l}}_r(v), & \text{if } p = \text{reactive} \end{cases} \tag{25}$$

Let I_p denote the set of available devices that improves the selected power type p, defined as

$$I_p = \{v \in I \,|\textsc{Gain}(v,p) > 0\}$$

1. The first strategy is to invest in the *cheapest* available device option that improves the selected power type p. Then,

$$\textsc{DeviceHeuristic}(I,p) = \min_{v \in I_p} \textsc{Cost}(v)$$

2. The second strategy is to invest in the most *efficient* device option in terms of the offered power in kW or kVAr and cost ratio.

$$\textsc{DeviceHeuristic}(I,p) = \max_{v \in I_p} \frac{\textsc{Gain}(v,p)}{\textsc{Cost}(v)}$$

3. The third strategy is to always invest in the most *powerful* option available in terms of offered power in kW or kVAr.

$$\textsc{DeviceHeuristic}(I,p) = \max_{v \in I_p} \textsc{Gain}(v,p)$$

Based on these heuristics, we implemented a basic greedy hill climbing optimization algorithm in Matlab, sketched in Algorithm 1. With the three investment heuristics DEVICEHEURISTIC, we get three variants of the algorithm that we denote *greedy/cheapest*, *greedy/efficient*, and *greedy/powerful* [1].

Algorithm 1 takes the initial circuit model c_0, the set of devices I, the budget constraint b, and the survivability requirement τ as an input. The outputs are a proxy to the Pareto front, which is an approximate solution to the MULTIOBJECTIVEPROBLEM, C^*, and an approximate solution to the SINGLEOBJECTIVEPROBLEM, c^*. Helper variables are initialized in line 1.

[1] We denote the algorithms as greedy because they accept the first neighbor solution that is better than the current solution. This procedure is also often called simple hill climbing. Note that most algorithms are deterministic, as the hill climbing is started from the initial circuit candidate just once and most heuristics are deterministic as well. Only the section heuristic does a random selection if several sections with the maximum number of violations are available.

Initially, the unavailable investments are determined based on the optimization constraints (line 3). The function APPLYCONSTRAINTS(I, U, c^*) receives as input a set of available devices, a set of previously determined unavailable investments (initially empty) and the current circuit. It returns the set of investments that do no meet the local and global constraints (16) and (17). This set of unavailable and unsuited investments U is updated during the optimization (lines 20 and 23).

In the main loop of Algorithm 1 (lines 4–26), it first determines a neighbor solution to investigate based on the three types of heuristics mentioned above (power type, section, and device) (lines 5-11). Together, these heuristics suggest a new investment (v, i) (line 11). The selected neighbor solution $c = c^* \cup \{(v, i)\}$ is built (line 12) and its survivability value is determined (line 13). If the survivability value is inferior to the current solution c, the neighbor solution c is discarded and stored as unsuitable (line 23). If the survivability value is superior, the neighbor solution is accepted as the new current solution (line 14–15), i.e., an uphill move is made. The spent amount of money is increased by the costs of the selected device (line 16). To make the search more efficient, a device that cannot be afforded any more at this point is removed from the set of devices I (line 19). Additionally, the set of unavailable investments is updated again as described above (line 20). The algorithm terminates when all investment options have been tried or discarded, the investment budget is met, or the survivability value reaches a defined target value (line 4).

The three investment heuristics could complement each other in finding good solutions. Thus, as a second optimization algorithm, we also developed a filtered steepest-ascent hill climbing approach. In each iteration of the optimization, instead of having one option suggested by one investment heuristic (line 11 in algorithm 1), we query each investment heuristic for a suggested next investment option. Then, a candidate solution is generated and evaluated for each suggested investment option. If any of the suggested investment options improves the survivability (in terms of AENS), the option with the highest improvement is selected. We denote this variant of the optimization algorithm as the *steepest-ascent* variant.

7. Evaluation

In this section, we present the application of our approach to a realistic case study system. Section 7.1 describes the case study system and the parametrization of the survivability model for this case study. Section 7.2 presents the evaluation results.

7.1. Case Study Setup

The case study system is inspired on a simple radial which has been suggested as a distribution automation benchmark by Rudion et al. [22]. The considered network is shown in Figure 7.

We consider a failure in section 1 as shown in Figure 7. This failure initially causes the left circuit to be unpowered. Then, the recloser can be opened to isolate the failure, and the tie switch can be closed to power sections 2 to 12 from the backup substation. Our power flow analyses will show in which cases the resulting circuit would be stable, i.e., the system can quickly recover by using distribution automation.

input : initial circuit model c_0
set of devices I
budget b
survivability requirement τ
output: Trace of accepted models, C^* (a proxy to the Pareto front which solves the MULTIOBJECTIVEPROBLEM), and
the best model found by our heuristics, c^* (yields an upper bound to the solution of the SINGLEOBJECTIVEPROBLEM)

$c \leftarrow \{c_0\}$; $c^* \leftarrow \{c_0\}$; $o \leftarrow 0$; $C^* \leftarrow \emptyset$;
// U stores unsuccessful and infeasible investments
$U \leftarrow$ APPLYCONSTRAINTS$(I, \emptyset, c^*)$;
while $(\text{AENS}(c^*) > \tau) \wedge (o \leq b) \wedge (I \neq \emptyset)$ **do**
  $p \leftarrow$ POWERTYPEHEURISTIC(c^*); // select power type p
  $i \leftarrow$ SECTIONHEURISTIC(c^*, p); // select section where to invest
  while *additional devices can still be deployed at section* i **do**
    // select device
    $v \leftarrow$ DEVICEHEURISTIC$(I \setminus \{v|(v,i) \in U\}, p)$;
    $c \leftarrow c^* \cup \{(v,i)\}$; // apply option
    if $\text{AENS}(c) < \text{AENS}(c^*)$ **then**
      $c^* \leftarrow c$;
      $C^* \leftarrow C^* \cup \{c^*\}$;
      $o \leftarrow o +$ COST(v);
      // remove options from I that cannot be afforded anymore
      $I \leftarrow \{v \in I : o +$ COST$(v) \leq b\}$;
      $U \leftarrow U \cup$ APPLYCONSTRAINTS(I, U, c^*);
      break; // leave inner while loop
    else $U \leftarrow U \cup \{(v,i)\}$;
  end
end
return C^* **and** c^*
// C^* and c^* are proxies to solutions of the MULTIOBJECTIVEPROBLEM and the SINGLEOBJECTIVEPROBLEM, respectively

Algorithm 1: Simple hill climbing optimization algorithm.

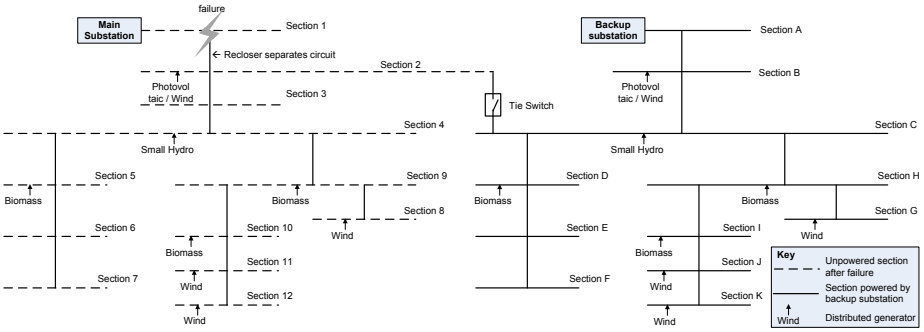

Figure 7: Case study circuit after failure at main substation (adapted from [19], details in Section 7.1).

Representative profile characteristics are based on the real-world conditions as described in [19]: the load profiles used in this case study were taken from the association of the electrical energy industry in Germany (BDEW), the generation profiles were taken from EnBW (Germany), and the wind profile used here was taken from E.on Netz (Germany). The data represents average values for the load and generation over the course of a day. The total active and reactive loads, averaged over the 96 load points, is equal to 13,906.00 KW and 4,570.80 KVar, respectively. We furthermore assume that, initially, 10% of the load in each section is amenable to reduction by demand response mechanisms in this system.

Investment Options and Constraints

We consider eight different types of investments, namely biomass plants, wind plants, solar plants, small hydro plants, demand-response investment, and DStatcom investment (the latter being a device to provide reactive power). For each type of investment, we assume 8 options with varying price and provided power as summarized in [9, Table I]. We derived the costs of the different options from [23, 24, 25] as detailed in [9]. All active power investments also generate reactive power. Reactive power generation is calculated based on the German code for distributed generators by BDEW as described in [19]. Reactive power investments (i.e., DStatcom) only add reactive power to the system.

The costs constraint is set to $b = 2$ million. To reflect additional constraints on investment options, we limit the number of selections as follows. Per section, only one generator per type can be added, e.g., one DStatcom (i.e., $\hat{S} = 1$). Globally, only three investment options of the same type can be used (i.e., $\hat{G} = 3$).

7.2. Results

Table 6 presents the statistics obtained by evaluating the four optimization heuristics on an IBM Thinkpad with two Intel Core 2 CPUs at 2GHz. The greedy/powerful variant is fastest because it has the lowest number of power flow and survivability evaluations. The greedy/cheapest variant is the slowest because more evaluations are required when the heuristic uses the cheapest options first.

Table 6: Statistics of the optimization algorithm runs

Variant	Number of evaluated candidate models	Duration in min
Greedy/Efficient	65	24
Greedy/Cheapest	118	46
Greedy/Powerful	27	10
Steepest-Ascent	85	27

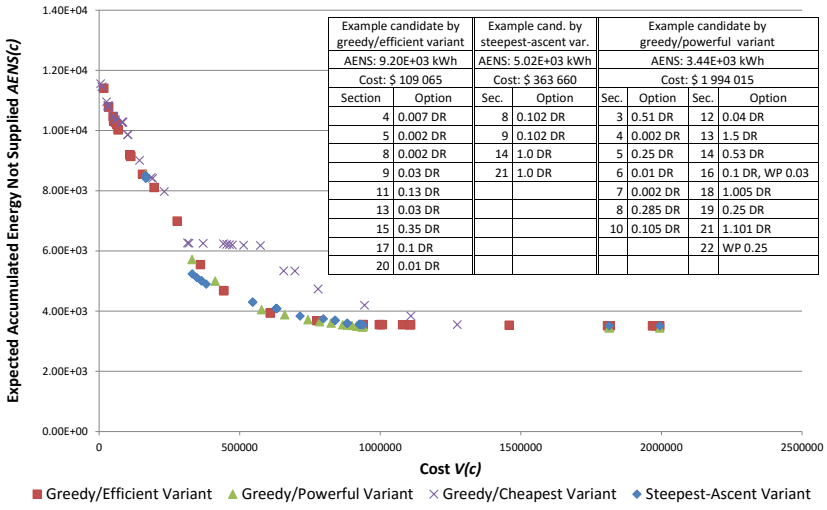

Example candidate by greedy/efficient variant		Example cand. by steepest-ascent var.		Example candidate by greedy/powerful variant			
AENS: 9.20E+03 kWh		AENS: 5.02E+03 kWh		AENS: 3.44E+03 kWh			
Cost: $ 109 065		Cost: $ 363 660		Cost: $ 1 994 015			
Section	Option	Sec.	Option	Sec.	Option	Sec.	Option
4	0.007 DR	8	0.102 DR	3	0.51 DR	12	0.04 DR
5	0.002 DR	9	0.102 DR	4	0.002 DR	13	1.5 DR
8	0.002 DR	14	1.0 DR	5	0.25 DR	14	0.53 DR
9	0.03 DR	21	1.0 DR	6	0.01 DR	16	0.1 DR, WP 0.03
11	0.13 DR			7	0.002 DR	18	1.005 DR
13	0.03 DR			8	0.285 DR	19	0.25 DR
15	0.35 DR			10	0.105 DR	21	1.101 DR
17	0.1 DR					22	WP 0.25
20	0.01 DR						

Figure 8: Comparison of optimization algorithms. Also shows table with three example solutions. Investment options are represented by generated power (in mW or mVAr) and investment type key. Key: DR = Demand response, WP = Wind Power

Figure 8 shows the results of each of the four variants. Each data point marks a candidate that improved the survivability of its parent. The candidates are plotted by their costs $\textsc{Cost}(c)$ on the x axis and survivability-related metric $\mathrm{AENS}(c)$ on the y axis.

The initial distribution grid is the starting point for all four variants at cost 0 and AENS 11,862 kWh. From there, the heuristics gradually succeed to reduce the AENS while inevitably resulting in higher costs.

The table in Figure 8 shows three example solution candidates showing the AENS, the costs, and the options selected for each section. The last example shown is the final candidate found by the heuristic which selects the most greedy/powerful choice at each iteration. It achieves the lowest AENS with 3,440 kWh (improvement of 8,422 kWh), but at high cost of almost \$1,994,015. The second example candidate is a solution found at around the knee point of the trade-off curve by the steepest-ascent heuristic and represents an investment of \$363,660 and an AENS improvement of 6,850 kWh when compared against the original circuit. It uses only demand response investments. Finally, the first candidate is found by the heuristic which selects the most greedy/efficient choice at each iteration, and represents a comparably inexpensive solution. It improves AENS by 2,660 kWh with investment costs of \$109,065.

The shape of the trade-off curve is also visible in Figure 8. Starting at costs \$0, first a steep improvement of AENS can be achieved by small investments. Later, the curve flattens and the same improvement is only achieved with higher investments.

We observe that the different algorithm variants perform differently in different regions of the trade-off front. For small investments, the greedy/cheapest variant and the greedy/efficient variant sample the solution space better and provide many solutions with small investment yet considerable AENS gain. In a slim middle region of \$330k to \$550k, the greedy/powerful and the steepest-ascent variants find the best solutions. In particular, the solutions found by the greedy/cheapest variant are inferior in this region, as they achieve higher AENS with the same budget. Finally, in the high investment region above \$550k, the greedy/powerful variant is most successful and finds superior solutions, whereas the greedy/efficient variant still provides more solutions of almost the same quality.

Next, we illustrate a sensitivity analysis of the steepest-ascent algorithm by considering how the algorithm performs when the DStatcom option is not available. The dotted lines in Figure 9 show how the proposed heuristic approximates the Pareto front when the DStatcom is available (dotted line with circles, also shown in Figure 8) and when it is not available (dotted line with squares). The similarity between the two curves suggests robustness in this particular example. It is interesting to note that at some points the availability of a DStatcom option slightly degraded the approximation of the Pareto front. This is fruit of the heuristic nature of the proposed algorithm.

Figure 9 also shows how the steepest-ascent algorithm compares against Pareto Archived Dynamically Dimensioned Search (PA-DDS) [26]. PA-DDS is an efficient multi-objective optimization algorithm that in previous work has presented good performance when contrasted against other alternatives, such as NSGA-II [27, 28]. PA-DDS easily handles discrete decision variables, and its Matlab implementation is publicly available at [29]. Using such an implementation (check [30] for details), we obtained the results shown in Figure 9.

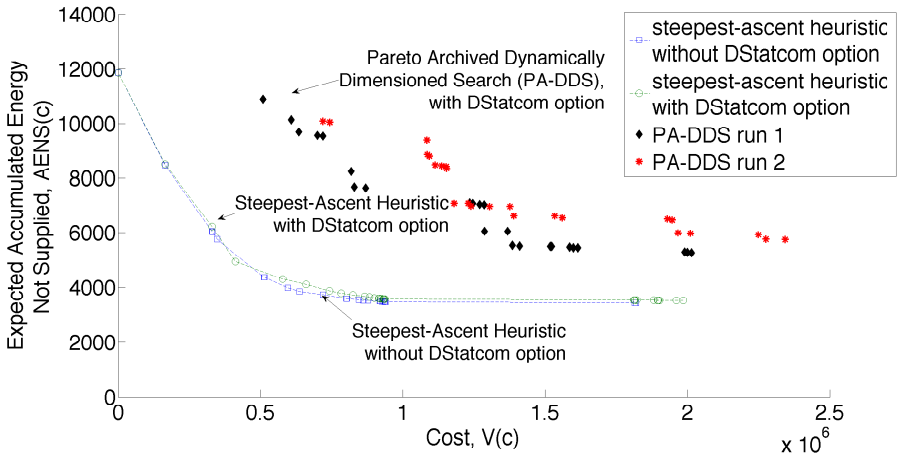

Figure 9: Comparison of proposed steepest-ascent heuristic against PA-DDS [26].

The encoding of the decision variables plays a key role when using PA-DDS. Next, we detail the encoding used in our numerical investigations. Recall that we consider a circuit with $N = 22$ sections, and each section can receive at most $\hat{S} = 1$ devices of each type. Each type comprises $\theta = 8$ options (Table 5). There are $\Gamma = 6$ device types (Table 4), each of which can be instantiated at most $\hat{G} = 3$ times at a given circuit. Therefore, we considered $\Gamma\hat{G} = 18$ decision variables. Let $\boldsymbol{d}$ be the matrix of decision variables, where $0 \leq d(k,l) \leq \hat{G}\theta$, for $k \in \{1,2,3\}$ and $l \in \{1,\ldots,\Gamma\}$. The element $d(k,l)$ characterizes the k-th placement of a device of type l. Let $d(k,l) = (v,i)$ if such placement occurs at section i taking device option v among the θ available options. Otherwise, let $d(k,l) = (0,0)$ if the k-th placement of a device of type l does not occur. To account for the constraint $\hat{S} = 1$, we associate infinite cost to combinations of decision variables which do not satisfy the constraint. Given such an encoding, PA-DDS can be readily applied.

We ran PA-DDS for two runs, with 1000 objective function evaluations per run, and marked the points found in the first and second run with diamonds and stars, respectively. In Figure 9 we show the results obtained with PA-DDS under the crowding source option [31] (the random option produced similar results [30]). Note that the approximation of the Pareto front obtained with PA-DDS is dominated by that found with the proposed heuristics. In a MacBook Pro Retina (late 2013 model), the former took roughly 2 hours per run to find the points marked with diamonds and stars, while the latter took less than 1 hour (see also Table 6 for running times of our algorithms in a different environment). This indicates that the heuristics, by making use of domain knowledge, perform significantly better than mainstream optimization algorithms. The performance of PA-DDS depends on the way the decision variables are encoded, as well as on initial conditions. In addition, PA-DDS requires fine tuning to take advantage of constraints. The budget constraint of $b = 2$ million was not encoded in our PA-DDS problem formulation, and solutions that violated the constraint are also pre-

sented in Figure 9. Our purpose in this section is to illustrate how different algorithms compare against each other. Future work consists of combining such approaches for better performance.

Discussion. Our tool can be used for return on investment computation (survivability/cost). We have found that an investment between $330K and $550K leads to good operating points as the improvement in survivability for higher investment values (greater than $550k) is marginal. The empirical results presented in this paper provide insights into the relation of investment and survivability and can thus support engineers to plan investments for the distribution grid.

The candidate circuits obtained by our algorithms are samples of the solution space. We envision that the efficient and principled sampling of the solution space as proposed in this paper will precede detailed and computationally intensive simulations of the most promising candidate circuits.

8. Conclusion

In this paper, we presented metrics, models and optimization heuristics to improve the reliability of power distribution networks. The basic building blocks of our contribution are survivability-related metrics, such as the expected accumulated energy since failure up to recovery (AENS). We proposed models that yield AENS as a function of the power circuit, its workload and corresponding power flow estimates. Then, we indicated that a model which accounts for failure rates might naturally take into account fine level details of the system, as well as multiple failures in parallel. However, the computational cost to solve it hampers its scalability. For this reason, we focused on the survivability-only model, and used it to issue recommendations on power circuit upgrades. We applied the proposed methodology to the design of a benchmark distribution automation circuit, indicating that the combination of survivability analysis and power flow can provide meaningful investment decision support for power systems engineers. We envision that the strategies presented in this paper can be used to efficiently explore the space of promising circuits, allowing engineers to consider options which would have been otherwise neglected.

References

[1] P. M. Duque and K. Sörensen, "A grasp metaheuristic to improve accessibility after a disaster," *OR spectrum*, vol. 33, no. 3, pp. 525–542, 2011.

[2] P. Maya Duque, S. Coene, P. Goos, K. Sörensen, and F. Spieksma, "The accessibility arc upgrading problem," *European J. of Operational Research*, vol. 224, no. 3, pp. 458–465, 2013.

[3] Z. Anwar, M. Montanari, A. Gutierrez, and R. Campbell, "Budget constrained optimal security hardening of control networks for critical cyber-infrastructures," *International Journal of Critical Infrastructure Protection*, 2009.

[4] I. Waseem, *Impacts of Distributed Generation on the Residential Distributed Network Operation*. M.Sc. Thesis, Virginia Polytechnic Institute, 2008.

[5] K. Moslehi and R. Kumar, "A reliability perspective of the smart grid," *IEEE Transactions on Smart Grid*, pp. 57–64, june 2010.

[6] D. Mazauric, S. Soltan, and G. Zussman, "Computational analysis of cascading failures in power networks," in *Proceedings of the ACM SIGMETRICS/international conference on Measurement and modeling of computer systems*. ACM, 2013, pp. 337–338.

[7] D. Ernst, M. Glavic, G.-B. Stan, S. Mannor, and L. Wehenkel, "The cross-entropy method for power system combinatorial optimization problems," *Proceedings of the 2007 Power Tech*, 2007.

[8] M. Poloczeki and M. Szegedy, "Randomized greedy algorithms for the maximum matching problem with new analysis," *FOCS*, 2012.

[9] A. Koziolek, A. Avritzer, S. Suresh, D. S. Menasché, K. Trivedi, and L. Happe, "Design of distribution automation networks using survivability modeling and power flow equations," in *ISSRE*, 2013.

[10] D. S. Menasché, R. M. Meri Leão, E. de Souza e Silva, A. Avritzer, S. Suresh, K. Trivedi, R. A. Marie, L. Happe, and A. Koziolek, "Survivability analysis of power distribution in smart grids with active and reactive power modeling," *SIGMETRICS Performance Evaluation Review*, vol. 40, no. 3, pp. 53–57, Jan. 2012.

[11] D. S. Menasché, A. Avritzer, S. Suresh, R. M. Leão, E. de Souza e Silva, M. Diniz, K. Trivedi, L. Happe, and A. Koziolek, "Assessing survivability of smart grid distribution network designs accounting for multiple failures," *Concurrency and Computation: Practice and Experience*, vol. 26, no. 12, pp. 1949–1974, 2014.

[12] A. Avritzer, L. Carnevali, L. Happe, A. Koziolek, D. S. Menasché, M. Paolieri, and S. Suresh, "A scalable approach to the assessment of storm impact in distributed automation power grids," in *QEST*, 2014.

[13] A. Avritzer, L. Happe, A. Koziolek, D. S. Menasché, S. Suresh, and J. Yallouz, "Scalable assessment and optimization of power distribution automation networks," in *Principles of Performance and Reliability Modeling and Evaluation*, L. Fiondella and A. Puliafito, Eds. Springer, 2016, pp. 321–340.

[14] Businessweek, "Hurricane Sandy boosts the generator makers," Bloomberg Businessweek Magazine, 2012, http://www.businessweek.com/articles/2012-11-05/hurricane-sandy-boosts-the-generator-makers.

[15] Y. Liu and K. Trivedi, "Survivability quantification: The analytical modeling approach," *International Journal of Performability Engineering*, vol. 2, no. 1, pp. 29–44, 2006.

[16] E. de Souza e Silva, D. R. Figueiredo, and R. M. M. Leão, "The TANGRAM-II integrated modeling environment for computer systems and networks," *SIGMETRICS Perform. Eval. Rev.*, vol. 36, no. 4, pp. 64–69, 2009, Tangram II availability model at http://www.dcc.ufrj.br/~sadoc/survivability.

[17] K. S. Trivedi, *Probability and Statistics with Reliability, Queuing and Computer Science Applications*, 2nd ed. John Wiley & Sons, 2001.

[18] A. Horváth, M. Paolieri, L. Ridi, and E. Vicario, "Transient analysis of non-markovian models using stochastic state classes," *Performance Evaluation*, vol. 69, no. 7, pp. 315–335, 2012.

[19] V. Janev, *Implementation and evaluation of a distribution load flow algorithm for networks with distributed generation*. Semester Thesis, Swiss Federal Institute of Technology, Zurich, 2009.

[20] W. K. Grassmann, M. I. Taksar, and D. P. Heyman, "Regenerative analysis and steady state distributions for markov chains," *Operations Research*, vol. 33, no. 5, pp. 1107–1116, 1985.

[21] G. Anderson, "Modelling and analysis of electric power systems," Lecture 227-0526-00, ITET ETH Zurich, 2008.

[22] K. Rudion, A. Orths, Z. Styczynski, and K. Strunz, "Design of benchmark of medium voltage distribution network for investigation of dg integration," in *Power Engineering Society General Meeting, 2006. IEEE*, 0-0 2006, p. 6 pp.

[23] US Gov. EPA, "Renewable energy cost database," http://www.epa.gov/cleanenergy/.

[24] F. Alvarado, B. Borissov, and L. Kirsh, "Reactive power as an identifiable ancilary server," Transmission Administrator of Alberta, 2003.

[25] F. E. R. Commission, "Building the dsr roadmap," Proceedings of the PJM Symp. on Demand Response II, 2008, http://www.ferc.gov/legal/staff-reports/12-08-demand-response.pdf.

[26] M. Asadzadeh and B. A. Tolson, "A new multi-objective algorithm, Pareto archived DDS," in *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers*. ACM, 2009, pp. 1963–1966.

[27] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *Evolutionary Computation, IEEE Transactions on*, vol. 6, no. 2, pp. 182–197, 2002.

[28] M. Asadzadeh, B. A. Tolson, and D. H. Burn, "A new selection metric for multiobjective hydrologic model calibration," *Water Resources Research*, vol. 50, no. 9, pp. 7082–7099, 2014.

[29] M. Asadzadeh, "PA-DDS website," 2016, http://home.cc.umanitoba.ca/~asadzadm/software.html.

[30] A. Koziolek, A. Avritzer, S. Suresh, D. S. Menasché, M. Diniz, E. de Souza e Silva, R. M. Leão, K. Trivedi, and L. Happe, "Source code of survivability models and optimization solutions presented in this paper," 2016, http://www.dcc.ufrj.br/~sadoc/survivability.

[31] M. Asadzadeh and B. Tolson, "Hybrid pareto archived dynamically dimensioned search for multi-objective combinatorial optimization: application to water distribution network design," *Journal of Hydroinformatics*, vol. 14, no. 1, pp. 192–205, 2012.

[32] H. S. Wilf, *Generatingfunctionology*. Academic Press, Inc, 1994.

Appendix A. State Space Cardinality

Next, we compute the cardinality of the solution space of the optimization problem. For a given section, we need to identify devices to be allocated to the specific section by considering the number of available investment options. Recall that an investment in a device of type $\gamma \in \Gamma$ can be implemented using different models of devices, that have different costs and capacities, and that θ_γ is the number of options for a device of type γ. Recall also that the maximum number of devices of a given type that can be allocated to a section is limited by a constant $\hat{S}$. Therefore, the maximum number of devices in a section is $\hat{S}|\Gamma|$. Accordingly, the global number of devices of a given type that can be allocated in the whole circuit is limited by $\hat{G}$.

First, we consider $\hat{G} = N\hat{S}$. As the maximum number of devices of a given type, per section, is $\hat{S}$, letting $\hat{G} = N\hat{S}$ is equivalent to $\hat{G} = \infty$. Let $|\Omega_c^{(\infty)}|$ be the cardinality of the solution space in this setting. Then,

$$|\Omega_c^{(\infty)}| = \prod_{\gamma\in\Gamma} \left(\left(\sum_{0\leq k\leq \hat{\mathrm{S}}} \left(\!\!\binom{\theta_\gamma}{k}\!\!\right) \right)^N \right) \tag{A.1}$$

where $\left(\!\binom{m}{n}\!\right)$ denotes the multiset binomial coefficient, $\left(\!\binom{m}{n}\!\right) = \binom{m+n-1}{n}$ and $\left(\!\binom{m}{0}\!\right) = 1$.

Let $k_{n,\gamma}$ be the number of devices of type γ at section n, $k_{n,\gamma} = 0, 1, \ldots, \hat{S}$. As there are θ_γ options of devices of type γ, we can allocate the θ_γ options in $\left(\!\binom{\theta_\gamma}{k_{n,\gamma}}\!\right)$ ways at section n. As $k_{n,\gamma}$ varies between $0, 1, \ldots, \hat{S}$, the number of possible allocations at a given section is given by the inner summation in (A.1). To account for all sections and all device types, the outer exponent and product in (A.1) are included.

In Section 7, $N = 22$, $|\Gamma| = 6$, $\theta_\gamma = 8$ for all types $\gamma \in \Gamma$, and $\hat{S} = 1$. Replacing these values into (A.1), the cardinality of the state space without the global constraint $\hat{G}$ is given by $|\Omega_c^{(\infty)}| = (1+8)^{22\cdot 6} \approx 9.12 \times 10^{125}$.

We now consider $\hat{G} < \infty$. Note that $0 \leq \hat{G} \leq N\hat{S}$. Let $|\Omega_c|$ be the cardinality of the solution space. Then,

$$|\Omega_c| = \prod_{\gamma \in \Gamma} \sum_{0 \leq k \leq \hat{\mathrm{G}}} \mathcal{L}(k; N, \hat{S}, \theta_\gamma) \tag{A.2}$$

where $\mathcal{L}(k; N, \hat{S}, \theta_\gamma)$ is the number of ways of allocating k devices of type γ into N sections, where each section can have up to $\hat{S}$ devices of type γ and each device of type γ is of one of θ_γ options, $\gamma = 1, 2, \ldots, |\Gamma|$.

Let $\mathcal{L}^\star(x; N, \hat{S}, \theta_\gamma)$ be the generating function of $\mathcal{L}(k; N, \hat{S}, \theta_\gamma)$. $\mathcal{L}(k; N, \hat{S}, \theta_\gamma)$ is given by the coefficient of x^k in $\mathcal{L}^\star(x; N, \hat{S}, \theta_\gamma)$. A closed-form for the generating function follows from [32, Chapter 1], and is given by,

$$\mathcal{L}^\star(x; N, \hat{S}, \theta_\gamma) = \left(\sum_{m=0}^{\hat{S}} \left(\binom{\theta_\gamma}{m} \right) x^m \right)^N \tag{A.3}$$

Note that if $\hat{G} = \infty$ (or, equivalently, $\hat{G} = N\hat{S}$) then $|\Omega_c| = |\Omega_c^{(\infty)}|$. To see this, note that if $\hat{G} = \infty$ the inner sum in (A.2) is the sum of all coefficients of x^m in (A.3). The sum of all coefficients of x^m in (A.3) can be obtained by setting $x = 1$ and equals the term between brackets in (A.1). This implies that if $\hat{G} = \infty$ then (A.1) and (A.2) are equal.

In the example of Section 7, we have $\hat{S} = 1$. Then, it follows from (A.3) that $\mathcal{L}^\star(x; N, 1, \theta_\gamma) = (1 + \theta_\gamma x)^N$ and $\mathcal{L}(k; N, 1, \theta_\gamma) = \binom{N}{k} \theta_\gamma^k$. We also have $N = 22$, $|\Gamma| = 6$, $\theta_\gamma = 8$ and $\hat{G} = 3$. Replacing these values into (A.2), the cardinality of the state space is given by $|\Omega_c| \approx 2.69 \times 10^{35}$. Note that we are not accounting for the budget constraint, which might reduce the number of options taken into account when looking for the optimal one.

The large solution space cardinality motivates the use of heuristics to do the state space exploration. The complexity of the algorithm proposed in Section 6 is the complexity to execute the power flow algorithm multiplied by the number of steps required to meet the requirements or the budget.