

Flexible Redundancy Generation for Virtual Network Embedding with an Application to Smart Grids

Allan Almeida Santos
allan.santos@eins.tu-darmstadt.de
Technical University of Darmstadt
Darmstadt, Germany

Amr Rizk
amr.rizk@uni-ulm.de
Ulm University
Ulm, Germany

Florian Steinke
florian.steinke@eins.tu-darmstadt.de
Technical University of Darmstadt
Darmstadt, Germany

ABSTRACT

Embedding applications consisting of interconnected logical function blocks, denoted Virtual Network Requests (VNR), onto physical compute and communication networks, denoted Substrate Networks (SN), allows for the automatic generation of a variable degree of redundancy. The need for this feature arises for instance in smart power distribution grids with many decentral devices. Their heterogeneous communication interconnections often feature low reliability and face frequently changing conditions. At the same time, high service reliability is required for critical applications such as voltage control. In this work, we show how to detect potential voltage violations in a medium voltage feeder with high probability at low monitoring cost. We employ a probabilistic power flow model to determine the time-dependent required reliability for the links of voltage monitoring VNRs and embed it onto a SN consisting of a mix of plausible smart grid communication technologies. We use a novel approach based on chance-constrained mixed integer linear programming to generate a minimal, but sufficient degree of redundancy. This allows for optimal resource usage of the SN, flexibility to adapt the embedding in case of changes of the VNR or SN parameters, and reduced design efforts in comparison to manual redundancy planning. Compared with a static redundancy setup, the operational communication costs can be more than halved in our simulation experiments.

CCS CONCEPTS

• **Hardware** → **Smart grid**; • **Networks** → **Network reliability**;
• **Mathematics of computing** → *Integer programming*.

KEYWORDS

Smart Grid, Network Reliability, Communication Redundancy, Mixed-Integer Linear Programming

ACM Reference Format:

Allan Almeida Santos, Amr Rizk, and Florian Steinke. 2020. Flexible Redundancy Generation for Virtual Network Embedding with an Application to Smart Grids. In *The Eleventh ACM International Conference on Future Energy Systems (e-Energy'20)*, June 22–26, 2020, Virtual Event, Australia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3396851.3397693>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

e-Energy'20, June 22–26, 2020, Virtual Event, Australia

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8009-6/20/06...\$15.00

<https://doi.org/10.1145/3396851.3397693>

1 INTRODUCTION

Network virtualization (NV) and the closely related field of software-defined networking (SDN) [3, 7, 18] are well-established in the field of communication systems. They are key technologies for the internet [19] and cloud infrastructures [11] that enable increased reliability of the provided services, efficient usage of the communication hardware, and the ability to perform quick adaptations in communication networks [8]. These properties make NV and SDN also well suited for smart grids, which integrate physical power networks and communication networks [26]. For example, NV enabled by SDN technology can provide scalable and efficient solutions for virtual power plants [26]. SDN can also support IEC 61850 implementations [17] and is an important tool for increasing the resilience [12] and cyber security [1] of smart grids.

At the core of many NV tasks lies the virtual network embedding (VNE) problem: Networks of inter-connected logical functional blocks, so-called Virtual Network Requests (VNR), are to be mapped for execution onto physical networks of compute resources and communication links, the so-called Substrate Networks (SN). VNE is a widely studied problem [6] where one line of work called survivable VNE [9] focuses on increasing the reliability of the embedded logic in case of SN failures. Some proposed mechanisms for this task are reactive, in that they redirect communication traffic to an alternative link or re-instantiate a compute job on another node only after a failure in the SN has occurred [24]. The availability of sufficient backup capacities for all kinds of failures should, of course, be ensured beforehand [13]. Other mechanisms act proactively, by keeping redundant backup resources online, even before any failure happens [2]. In this case immediate service continuation can be guaranteed.

In this paper, we formulate and solve a chance-constrained mixed integer linear program (MILP) for VNE that proactively generates redundancy for the VNR when embedded onto the SN. Unlike previous approaches such as [2], we do not provision a fixed degree of redundancy, e.g. two-fold instantiation, for all VNR elements. Instead, we flexibly embed each VNR element one or more times onto the SN, satisfying individual reliability requirements at minimal embedding cost. This reduces SN resource usage for uncritical elements while at the same time allowing critical elements to be secured with high degrees of redundancy. The automated redundancy generation process is well-suited for adapting the embedding and its redundancy degree when the reliability requirements of the VNR or the reliability guarantees of the SN change over time.

We then show how the proposed approach can suitably be employed to reduce communication costs in smart distribution grids. Renewable energy integration leads to potential voltage violations,

especially in medium voltage (MV) feeders [4], and various monitoring and control schemes have been proposed for this purpose, e.g., [15]. We construct a VNR for a voltage monitoring scheme that connects distributed on-load tap changer (OLTC) power transformers with the control center. We embed it into a SN consisting of different communication modalities such as power line communication (PLC), long range wide area network (LoRaWAN) and general packet radio service (GPRS) networks. Unlike the dedicated communication networks of transmission grid operators, the considered links are highly prone to failures and outages individually, requiring a certain extent of redundancy to guarantee reliable service execution. We compute for each link the required reliability given the current grid state and use our proposed approach to generate the minimal amount of redundancy for each generation. In comparison to a fixed dual redundancy, as is e.g. done in [16] using PLC and wireless mesh networks, we are able to reduce average operating costs by more than half in our simulation experiments.

The remainder of this paper is structured as follows: Section 2 defines important concepts and terms and formally states the problem of automated redundancy generation to be solved. Our proposed MILP formulation for this task is then presented in Section 3. Its features are first demonstrated with a simple example in Section 4. In Section 5 we then present the voltage control application and corresponding simulation experiments. We conclude in Section 6.

2 PROBLEM STATEMENT AUTOMATED REDUNDANCY GENERATION

In this section, we first define the key terms for the VNE problem, before discussing SN failures and the proposed redundancy generation problem.

2.1 Formal VNE Concepts

Let an SN be defined as an undirected multigraph $G = (\mathcal{V}, \mathcal{L})$ with substrate nodes \mathcal{V} and substrate links \mathcal{L} . Each substrate node $i \in \mathcal{V}$ has cost of allocation c_i and processing capacity p_i . Substrate links are identified by the triple $(i, j, n) \in \mathcal{L}$, with i and j being the end nodes of the link and n the index of the link between both nodes. We use this multigraph notation here to be able to represent heterogeneous networks below, where different communication modalities connect the same nodes. Each link has cost of allocation c_{ijn} and available bandwidth b_{ijn} . A VNR is defined as an undirected graph $\tilde{G} = (\tilde{\mathcal{V}}, \tilde{\mathcal{L}})$ with virtual nodes $\tilde{\mathcal{V}}$ and virtual links $\tilde{\mathcal{L}}$. A virtual node $u \in \tilde{\mathcal{V}}$ requires processing capacity p^u , and a virtual link $(u, v) \in \tilde{\mathcal{L}}$ requires bandwidth b^{uv} . A VNE is a mapping of each virtual node onto one substrate node and of each virtual link onto a path in the SN, such that the mapped elements have the same connectivity as the VNR and satisfy capacity and bandwidth constraints.

2.2 Reliability and Redundancy Generation

In a non-perfect SN, one or more nodes and links can fail at any time. Any such failure mode of the SN can be encoded by a vector of failure indicators, one for each network element. We do not consider correlations between failures in time and assume the failure modes to be independently and identically distributed according to a known probability distribution over this finite set of vectors. To

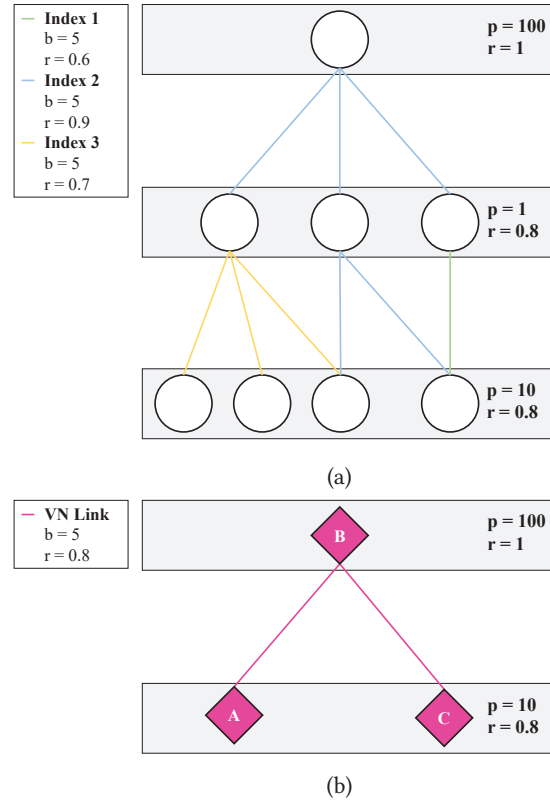


Figure 1: Exemplary embedding problem with reliability and capacity specifications for each element: (a) SN where each color represents a different index n of the substrate links (i, j, n) , (b) VNR.

simplify our outline we describe the failure distributions over the SN only with their marginal failure probabilities for each element, and assume independence between the failures of different elements. Specifically, we define for each substrate node a reliability r_i and for each substrate link a reliability r_{ijn} indicating that the node i and link (i, j, n) are available that fraction of time, respectively.

When a VNR represents a critical service, as is the case for many smart grid applications, it is necessary to specify suitable reliability requirements. We thus define a reliability requirement r^u for each virtual node u and a reliability requirement r^{uv} for each virtual link (u, v) , which determine the allowed minimum fraction of time each component of the request shall be available. Figure 1 shows an example of a SN and a VNR along with their capacity and reliability parameters.

In the case a VNR element has a greater reliability requirement than provided by the single SN element it is embedded onto, it may be helpful to embed a virtual element onto several redundant copies in the SN network. The probability that all these copies fail at the same time is reduced for most failure distributions of the SN, e.g., if the failures of individual SN elements are assumed to be fully independent. Given such independence and sufficient substrate

resources, it is possible to satisfy any reliability requirement with sufficiently many copies.

To represent the variable amount of redundancy for each VNR element, we formally determine up to K full virtual network embeddings, where $K \geq 2$ is a user-defined parameter. These embeddings, however, are not required to be disjoint. Whenever a virtual node is mapped onto the same substrate network node in several such VNEs, we assume that only one copy is executed in the SN environment and SN resources are used only once. Cost minimization then drives many copies to be located in the same place, implying a degree of redundancy smaller than K , while still satisfying the reliability requirement. The same holds for the embedding of virtual links onto SN paths. Here, it could additionally happen that the paths are only partly identical, i.e. they split or merge at certain nodes. In case of splitting, we assume that the software-defined network replicates the communication messages along both splitting paths. In case of path mergers, we assume that the data traffic from one incoming link is dropped, as long as the other one is available. These procedures enable operating with varying degrees of redundancy for each element of the VNR.

The problem to be solved in this work is to cost-optimally determine K feasible VNEs that jointly fulfill the described reliability requirements posed to the VNR. Additionally, it may be desirable that the VNE map changes only minimally in case of slight changes of the VNR or SN parameters that leave the topology unchanged.

3 MILP FORMULATION

In this section, we solve the proposed problem using mixed integer linear programming (MILP), which is a widespread modeling and optimization tool for constrained combinatorial problems. MILP comes with several high performance solvers that often yield good solutions in short time and allow for hard optimality guarantees during the optimization process.

3.1 Decision Variables for Node and Link Embeddings

Let a VNR $\tilde{G} = (\tilde{\mathcal{V}}, \tilde{\mathcal{L}})$, a SN $G = (\mathcal{V}, \mathcal{L})$, and a set of embedding indices $\mathcal{K} = \{1, \dots, K\}$ be given. We then define a binary decision variable $x_{i,k}^u$ for all nodes $u \in \tilde{\mathcal{V}}$, $i \in \mathcal{V}$, and $k \in \mathcal{K}$ as

$$x_{i,k}^u = \begin{cases} 1, & \text{if } u \text{ embedded onto } i \text{ in embedding } k, \\ 0, & \text{otherwise.} \end{cases}$$

We also define a binary decision variable $y_{ij,n,k}^{uv}$ for all links $(u, v) \in \tilde{\mathcal{L}}$, $(i, j, n) \in \mathcal{L}$, and $k \in \mathcal{K}$ as

$$y_{ij,n,k}^{uv} = \begin{cases} 1, & \text{if } (i, j, n) \text{ is part of the } k\text{-th embedding of } (u, v), \\ 0, & \text{otherwise.} \end{cases}$$

As our modeling assumes that several embeddings of one VNR element on the same SN element mean that only one instance is executed in the SN, we define the following auxiliary variables. $\hat{x}_i^u \in \{0, 1\}$ indicates whether VNR node $u \in \tilde{\mathcal{V}}$ is mapped onto SN node $i \in \mathcal{V}$ in any of the K mappings. $\hat{y}_{ijn}^{uv} \in \{0, 1\}$ indicates whether SN link $(i, j, n) \in \mathcal{L}$ is part of any of the K SN paths that VNR link $(u, v) \in \tilde{\mathcal{L}}$ is mapped to. Constraints (1) and (2) ensure

these definitions,

$$\hat{x}_i^u \geq x_{i,k}^u, \forall u \in \tilde{\mathcal{V}}, \forall i \in \mathcal{V}, \forall k \in \mathcal{K}, \quad (1)$$

$$\hat{y}_{ijn}^{uv} \geq y_{ijn,k}^{uv}, \forall (u, v) \in \tilde{\mathcal{L}}, \forall (i, j, n) \in \mathcal{L}, \forall k \in \mathcal{K}. \quad (2)$$

3.2 Objective Function

We aim to minimize the embedding cost considering only the actually executed instances of the virtual elements. Therefore, we define the following objective function,

$$\min \left(\sum_{u \in \tilde{\mathcal{V}}} \sum_{i \in \mathcal{V}} c_i \hat{x}_i^u + \sum_{(u,v) \in \tilde{\mathcal{L}}} \sum_{(i,j,n) \in \mathcal{L}} c_{ijn} \hat{y}_{ijn}^{uv} \right). \quad (3)$$

3.3 Connectivity and Capacity Constraints

With condition (4), we ensure that every virtual node is mapped onto exactly one substrate node,

$$\sum_{i \in \mathcal{V}} x_{i,k}^u = 1, \forall u \in \tilde{\mathcal{V}}, \forall k \in \mathcal{K}. \quad (4)$$

We also require that all virtual links are embedded onto a substrate path with minimum length 1 by using constraint (5). Constraint (6), also called the multi-commodity flow constraint [5], ensures a continuous substrate path for each virtual link embedding,

$$\sum_{(i,j,n) \in \mathcal{L}} y_{ijn,k}^{uv} \geq 1, \forall (u, v) \in \tilde{\mathcal{L}}, \forall k \in \mathcal{K}, \quad (5)$$

$$\sum_{j \in \mathcal{N}(i)} \left(y_{ijn,k}^{uv} - y_{jin,k}^{uv} \right) = x_{i,k}^u - x_{i,k}^v, \forall (u, v) \in \tilde{\mathcal{L}}, \quad (6)$$

$$\forall i \in \mathcal{V}, \forall k \in \mathcal{K}.$$

Here, $\mathcal{N}(i)$ is the set of neighbors of SN node i . Constraints (7) and (8) express capacity and bandwidth limitations in terms of the executed instances,

$$\sum_{u \in \tilde{\mathcal{V}}} p^u \hat{x}_i^u \leq p_i, \forall i \in \mathcal{V}, \quad (7)$$

$$\sum_{(u,v) \in \tilde{\mathcal{L}}} b^{uv} \hat{y}_{ijn}^{uv} \leq b_{ijn}, \forall (i, j, n) \in \mathcal{L}. \quad (8)$$

3.4 Constraints for Redundancy Generation

We now need to express for each virtual element the condition that the probability of all its K embeddings failing at the same time is less or equal than its maximal allowed failure level, which is one minus its required reliability level. To express these chance constraints in a MILP setting, we use a scenario-based approach [22]. We first present the constraints for ensuring VNR node reliability and then for VNR link reliability.

Let Ω be the set of all failure modes of the SN with an additional element for the undisturbed case, and π_ω be the probability of each scenario $\omega \in \Omega$. We define parameter $\xi_{i,\omega}$ to indicate whether node $i \in \mathcal{V}$ fails in a given scenario $\omega \in \Omega$,

$$\xi_{i,\omega} = \begin{cases} 1, & \text{if } i \text{ failed in scenario } \omega, \\ 0, & \text{otherwise.} \end{cases}$$

Then we let binary variable $\eta_{k,\omega}^u$ indicate whether the substrate node i that virtual node $u \in \tilde{\mathcal{V}}$ is mapped onto in embedding k fails in scenario ω . This is expressed by the following constraint,

$$\eta_{k,\omega}^u \geq \xi_{i,\omega} x_{i,k}^u, \forall u \in \tilde{\mathcal{V}}, \forall i \in \mathcal{V}, \quad (9a)$$

$$\forall k \in \mathcal{K}, \forall \omega \in \Omega.$$

In constraint (9b), we summarize over k to obtain ζ_ω^u which indicates whether all K embeddings of virtual node u fail in scenario $\omega \in \Omega$,

$$\zeta_\omega^u \geq \sum_{k \in \mathcal{K}} \eta_{k,\omega}^u - (K-1), \forall u \in \tilde{\mathcal{V}}, \forall \omega \in \Omega. \quad (9b)$$

The sum of the probabilities of such failure scenarios must be less or equal to the maximum allowed probability of failure of that virtual node. This is ensured by the following condition,

$$\sum_{\omega \in \Omega} \pi_\omega \zeta_\omega^u \leq 1 - r^u, \forall u \in \tilde{\mathcal{V}}. \quad (9c)$$

For virtual links we construct analogous constraints (10) using the link element failure indicators $\xi_{ijn,\omega}$ and the auxiliary embedding variables $\eta_{k,\omega}^{uv}$ and ζ_ω^{uv} to express the failure of one or all embeddings for a link in a given scenario, respectively. This results in

$$\eta_{k,\omega}^{uv} \geq \xi_{ijn,\omega} y_{ijn,k}^{uv}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \quad (10a)$$

$$\forall k \in \mathcal{K}, \forall \omega \in \Omega,$$

$$\zeta_\omega^{uv} \geq \sum_{k \in \mathcal{K}} \eta_{k,\omega}^{uv} - (K-1), \forall (u,v) \in \tilde{\mathcal{L}}, \forall \omega \in \Omega, \quad (10b)$$

$$\sum_{\omega \in \Omega} \pi_\omega \zeta_\omega^{uv} \leq 1 - r^{uv}, \forall (u,v) \in \tilde{\mathcal{L}}. \quad (10c)$$

The previous conditions imply that communication can still be routed through a node, even if it failed. This may sometimes be plausible, but we present the following "node-failure-implies-link-failure" conditions that can be used in addition to (10a) to ensure that a node failure implies also the failure of its adjacent links, or equivalently that a virtual link embedding fails if any of the nodes on its path fails,

$$\eta_{k,\omega}^{uv} \geq \xi_{i,\omega} y_{ijn,k}^{uv}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \quad (11a)$$

$$\forall k \in \mathcal{K}, \forall \omega \in \Omega,$$

$$\eta_{k,\omega}^{uv} \geq \xi_{j,\omega} y_{ijn,k}^{uv}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \quad (11b)$$

$$\forall k \in \mathcal{K}, \forall \omega \in \Omega.$$

3.5 Adaptation Costs

When we re-optimize the VNE problem several times to adapt it to slightly changing parameters of the VNR or the SN, it is desirable that the embedding changes as little as possible. This avoids the replacement of the embedded logical functions or the reshaping of the communication paths. Let $\hat{x}_i^{u'}$ and $\hat{y}_{ijn}^{u'v'}$ be the values of \hat{x}_i^u and \hat{y}_{ijn}^{uv} from a previous mapping. We then define the auxiliary variables $\bar{x}_i^u \in \{0, 1\}$ and $\bar{y}_{ijn}^{uv} \in \{0, 1\}$ that indicate whether the current node and link mapping remains the same, respectively. This

yields

$$\bar{x}_i^u \geq \hat{x}_i^u - \hat{x}_i^{u'}, \forall u \in \tilde{\mathcal{V}}, \forall i \in \mathcal{V}, \quad (12)$$

$$\bar{x}_i^u \geq \hat{x}_i^{u'} - \hat{x}_i^u, \forall u \in \tilde{\mathcal{V}}, \forall i \in \mathcal{V}, \quad (13)$$

$$\bar{y}_{ijn}^{uv} \geq \hat{y}_{ijn}^{uv} - \hat{y}_{ijn}^{u'v'}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \quad (14)$$

$$\bar{y}_{ijn}^{uv} \geq \hat{y}_{ijn}^{u'v'} - \hat{y}_{ijn}^{uv}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}. \quad (15)$$

Thereafter, we add the term Θ to the objective function, making large changes in the two consecutive mappings more costly,

$$\Theta = \sum_{u \in \tilde{\mathcal{V}}} \sum_{i \in \mathcal{V}} \lambda^V \bar{x}_i^u + \sum_{(u,v) \in \tilde{\mathcal{L}}} \sum_{(i,j,n) \in \mathcal{L}} \lambda^L \bar{y}_{ijn}^{uv}.$$

Here, λ^V and λ^L parametrize the adaptation cost for node and link mappings, respectively.

4 DEMONSTRATIVE EXAMPLE

In this section we demonstrate our redundancy generation algorithm on the exemplary VNE problem shown in Figure 1. The MILP formulation is implemented in Python, utilizing PuLP, and is solved with CPLEX using its default configuration. We first show the effects of the node and link reliability constraints separately, before considering all such constraints and non-zero adaptation costs.

The set of scenarios Ω is generated by considering, apart from the normal operation mode, the failure of any single element $\sigma \in \Sigma$, i.e. a node or link, of the SN. The probability π_ω of scenario $\omega \in \Omega$ is defined with the help of the failure indicators $\xi_{\sigma,\omega}$ of element σ as

$$\pi_\omega = \frac{1}{Z} \prod_{\sigma \in \Sigma} [(1 - r_\sigma) \xi_{\sigma,\omega} + r_\sigma (1 - \xi_{\sigma,\omega})]. \quad (16)$$

Since we do not consider all possible failure modes, e.g., no two-element failures, the probabilities are normalized by a suitable constant Z . Note that by this normalization step, the marginal reliability of the elements of the SN is increased compared to the specified values. If not mentioned otherwise, the capacity and reliability parameters are taken as in Figure 1. The cost for embedding c_i for all substrate nodes is set to 1. The cost for link embedding c_{ijn} is 5, 10, and 20 for all the links with index 1, 2 and 3, respectively. $K = 2$ is sufficient here to satisfy all requirements.

Node reliability: We first embed the VNR onto the SN considering only the node reliability constraints (9). The left diagram of Subfigure 2 (a) shows the resulting embedding. The VNE program uses the cheapest links to embed the request, as, in this case, the link reliability constraints (10) are not considered. Increasing the reliability requirement of node A from $r^A = 0.8$ to $r^A = 0.95$, the program automatically generates another copy of node A , along with its adjacent links, as seen in the right diagram of Figure 2 (a).

Link reliability: Next, we consider embeddings using only the link failure constraints (10). The left diagram of Subfigure 2 (b) shows the resulting embedding for default parameters. The program embeds the virtual links onto substrate links of indices 2 and 3, as the reliability of substrate links with index 1 is below the required reliability level of the VNR and thus would require a double embedding at increased cost. Increasing the reliability requirement of virtual link (A, B) to $r^{AB} = 0.97$, that virtual link is mapped onto two disjoint paths in the SN, see the right diagram of Figure 2 (b).




Figure 2: The VNE problem of Figure 1 is solved with three different sets of redundancy constraints (a)-(c) and with non-zero adaptation costs (d). The left graph always shows the embedding with the default parameters, while the right one is obtained by increasing a reliability requirement of the VNR. Dashed lines denote redundant link embeddings while nodes with dashed outline denote virtual redundant nodes. (a) Use of node reliability constraints (9) and increase of reliability requirement for node A from $r^A = 0.8$ to $r^A = 0.95$. (b) Use of link reliability constraints (10) and increase of reliability requirement of link (A, B) from $r^{AB} = 0.8$ to $r^{AB} = 0.99$. (c) Use of the reliability constraints (9), (10), and (11) and increase of reliability requirement of node A from $r^A = 0.8$ to $r^A = 0.95$ and of link (A, B) from $r^{AB} = 0.8$ to $r^{AB} = 0.9$. (d) Same as (a) but with adaptation costs $\lambda^V = \lambda^L = 100$.

Node and link reliability: We now consider embeddings with all the conditions (9), (10), and (11). The resulting embedding shown in the left diagram of Subfigure 2 (c) is computed with default parameters. The VNE program generates redundancy for nodes A and C , along with their respective paths to B . This is due to constraint (11) that relates a node failure to the failure of adjacent links. This can be seen as follows for the redundancy of A : the summed probabilities of the scenario with a failure of the leftmost communication link of index 3 and the scenario with a failure of the leftmost intermediate node are larger than the allowed maximum failure level of the virtual link A to B . Thus another copy of A with another link is generated. However, fully disjoint paths between primary and secondary mapping are not necessary given the default reliability requirements for the links. Increasing the reliability requirement of link (A, B) to $r^{AB} = 0.9$, the program then generates a disjoint redundancy for that link. The resulting embedding is shown in the right diagram of Subfigure 2 (c).

Adaptation: The previous simulations did not consider adaptation costs when updating parameter values, e.g., when changing the reliability requirement of node A from $r^A = 0.8$ to $r^A = 0.95$ in Subfigure 2 (a). We now repeat this experiment with $\lambda^V = \lambda^L = 100$, obtaining the right diagram in Subfigure 2 (d). In contrast to the right diagram of Subfigure 2 (a), both mappings of nodes A and C and their respective paths towards node B now remain constant in comparison to the left diagram, and the only difference is the addition of node redundancy for A and its path towards B due to the increased reliability requirement.

5 SMART GRID VOLTAGE CONTROL

The increasing amount of Distributed Energy Resources (DERs) can cause critical voltage rises throughout the electric distribution grid [4]. In radial low voltage networks these problems can be solved mostly locally by using OLTC transformers at the connection point to the MV level. At the MV level, a coordinated effort

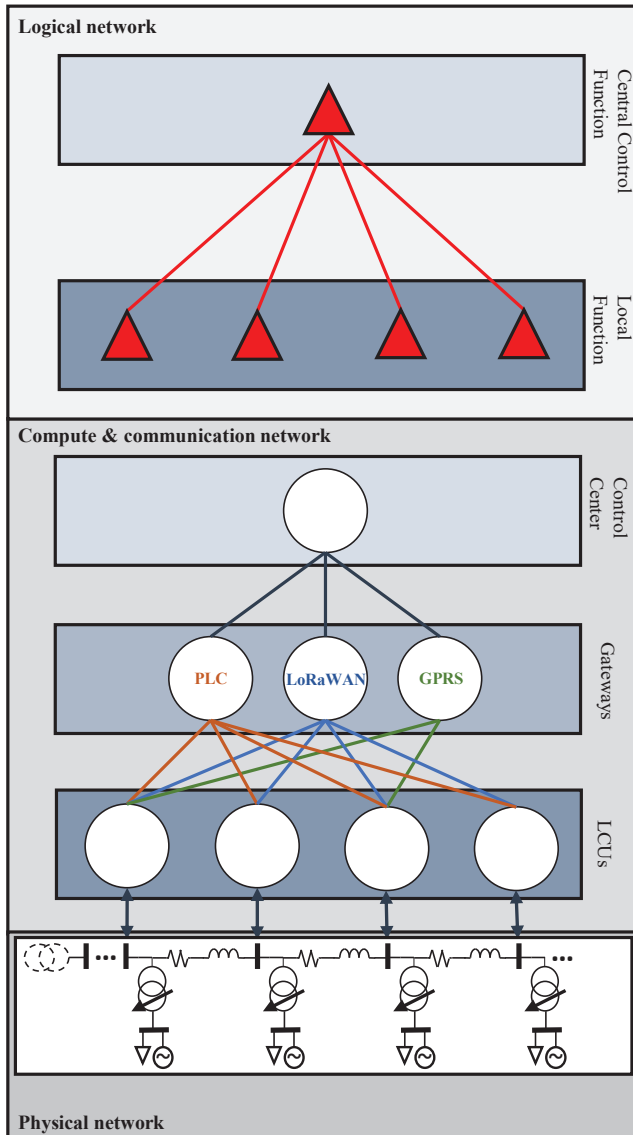


Figure 3: Overview of the simulated setup with one MV feeder as the underlying physical network. The compute and communication network includes one LCU per OLTC power transformer that connects underlying low voltage grids. The three considered communication technologies have one gateway each. The logical network has a central control function connected to a decentral logic in each LCU.

with communication is needed [15]. In this section, we integrate a probabilistic power flow model with our proposed flexible redundancy generation approach, to detect at low operational cost the possible over-voltages at the interconnected OLTC transformers with a defined, high reliability, despite potentially error prone communication channels.

Smart grids are an instance of the general principle of cyber-physical systems [21]. When discussing VNE in this context we

Table 1: Parameter values of compute & communication network.

	cost [US\$/MB]	average data rate [bps]	reliability $r \in [0, 1]$
PLC	0	10k	0.85
LoRaWAN	0	5.36	0.66
GPRS	5	9k	0.94

have to clearly differentiate between three types of networks: the *physical network* comprises the tangible components of the distribution grid that are part of the control scheme. The *compute & communication network* consists of the compute capabilities to run the required software, both on field devices and in the control center, and different communication media and their infrastructure, such as gateways. This network acts as the SN in our case. The *logic network* holds the interconnected function blocks that implement the voltage control mechanism. They act as the VNR.

We will use this terminology to present the considered setup in the following, see also Figure 3. Next, we discuss a probabilistic model of the time-dependent power flow in the grid, which determines the probability of over- or under-voltage at any node and thereby the required reliability of its communication link. We then present simulation results and discuss the benefits of our flexible redundancy generation approach for this setting.

5.1 Simulated Setup

Physical Network: We consider one MV feeder with 15 OLTC transformers connecting underlying low voltage grids.

Compute & Communication Network: Three different communication technologies are employed for connecting the local control units (LCU) associated with each OLTC transformer to a gateway. We assume that the communication link between the gateways and the control center via a wired backbone network is perfectly reliable and has no data capacity constraints, allowing us to model the path between gateways and control center as a single link with maximum reliability and infinite data capacity.

Table 1 lists the relevant properties for the considered technologies PLC, LoRaWAN, and GPRS. Narrowband PLC links have near zero operating cost, as they operate on existing power lines and are directly owned by the grid operator. Maximum data rates vary between 10-500 kbps [14], which is sufficient for the considered control application. However, narrowband PLC does not have the necessary reliability to be deployed as the only communication channel for a critical control mechanism such as voltage control. Its packet delivery ratio is only around 85% in similar setups [10]. LoRaWAN allows for long range communication (up to tens of kilometers), making it ideal for connecting decentral OLTC transformers in long feeders, e.g., in rural areas. Similar to narrowband PLC, LoRaWAN can be directly operated by the grid operator with almost zero operating cost, due to the usage of unlicensed sub-gigahertz frequency bands. The drawbacks of this technology are two-fold. First, LoRaWAN can have reliability as low as 66% depending on the distance from the gateway, line of sight and spreading




Figure 4: Voltage levels at 13:00 of the 15 buses with 95% (marginal) confidence interval. Below the diagram we show for each bus the communication technology utilized for connecting the associated LCU to the control center.

factor [20], which in our setup was considered as the largest possible value for range maximization. Also, capacity constraints arise due to low data rates and the maximum allowed duty cycle, which in Europe is 1% in most of the frequency bands. By using the equations that rule the LoRa modulation and considering maximum spreading factor, a packet with 12 bytes of payload and 13 bytes of header would have transmission time of around 1500 ms and the time interval between sending two packets in the same sub-band is about 150 s. Therefore, the theoretical data rate of each channel is approximately 1.34 bps. In this setup, we assume the usage of 4 channels, resulting in a combined data rate of 5.36 bps. Due to the capacity constraints of LoRaWAN, we include GPRS as a third alternative into the setup. This technology provides maximum data rates around 9 kbps and has 94% reliability in similar scenarios [23]. For grid operators it incurs high operating costs, since third-party communication infrastructure needs to be paid for. We assume a price range of US\$5-10/MB [25]. Furthermore, we assume that only half of the LCUs have GPRS transceivers, as these lead to additional costs. For our experiments, we assigned GPRS transceivers to LCUs at buses with an even index.

Logical Network: The considered voltage control application consists of a central control function connected to a decentral logic in each LCU. Note that the node mapping is fixed in this setup. Considering a monitoring and/or control message each 5 minutes and a packet of 25 bytes size, the application requires at least 0.67 bps for each link between an LCU and the control center. The reliability requirement of the virtual links is determined by the probability of a voltage violation given the current grid state and the desired detection rate.

5.2 Probabilistic Power Flow Model

We consider a MV feeder with line topology and buses $\{1, \dots, N\}$. Bus 0 connects to the high voltage grid and acts as the slack node.

The active power in-feed at bus i at time t is denoted as $p_i(t)$. It consists of the load $d_i(t)$ and the photovoltaic generation $\phi_i(t)$ connected in the underlying low voltage grids. While the loads are modeled independent of each other, i.e. $d_i(t) \sim N(\mu_{d,i}(t), \sigma_{d,i}^2)$, the photovoltaic production is assumed to depend on a joint irradiation value for the whole region, i.e. $\phi_i(t) = c_i \phi_0(t)$ where $\phi_0(t) \sim N(\mu_\phi(t), \sigma_\phi^2(t))$ and c_i the installed capacity at bus i . Denoting $\mathbf{p} = [p_1, \dots, p_N]^T$, we obtain $\mathbf{p} \sim N(\boldsymbol{\mu}_p, \boldsymbol{\Sigma}_p)$ for each time step t where $\boldsymbol{\mu}_p = -\boldsymbol{\mu}_d(t) + \mu_\phi(t)\mathbf{e}$ and $\boldsymbol{\Sigma}_p = \sigma_\phi^2 \mathbf{1}\mathbf{1}^T + \mathbf{e}\mathbf{e}^T \sigma_\phi^2$, with \mathbf{e} , \mathbf{c} , $\mathbf{1}$ being a vector of ones, the vector of photovoltaic capacities, and the identity matrix of appropriate dimension, respectively.

Using a linearized power flow model, the active power flow f_i from bus i to its upstream neighbor $i-1$ is then given by

$$f_i = \sum_{j=i}^N p_j, \forall i \in \{1, \dots, N\}, \quad (17)$$

which can be written in matrix form as $\mathbf{f} = \mathbf{A}\mathbf{p}$ for an appropriately defined matrix \mathbf{A} . With u_i denoting the voltage amplitude at bus i in p.u. and assuming a power factor close to one, we have

$$y_i(u_i - u_{i-1}) = f_i, \forall i \in \{1, \dots, N\}, \quad (18)$$

with $y_i = \mathcal{R}(Y_i)U_0$ where Y_i is (the real part of) the admittance of the line and U_0 the nominal voltage. We assume that $u_0 = 1$. This equation can be written with appropriate definition of \mathbf{B} and \mathbf{v} as $\mathbf{B}\mathbf{u} + \mathbf{v} = \mathbf{f}$. Together we obtain

$$\mathbf{u} = \mathbf{B}^{-1}\mathbf{A}\mathbf{p} - \mathbf{B}^{-1}\mathbf{v} \sim N(\boldsymbol{\mu}_u, \boldsymbol{\Sigma}_u), \quad (19)$$

where $\boldsymbol{\mu}_u = \mathbf{B}^{-1}\mathbf{A}\boldsymbol{\mu}_p - \mathbf{B}^{-1}\mathbf{v}$ and $\boldsymbol{\Sigma}_u = \mathbf{B}^{-1}\mathbf{A}\boldsymbol{\Sigma}_p\mathbf{A}^T (\mathbf{B}^{-1})^T$.

Suppose the probability of an undetected voltage violation at bus i shall be kept below ϵ . This probability is the product of $P(u_i < \underline{U} \vee u_i > \overline{U})$, with \underline{U} and \overline{U} being the lower and upper limits of the allowed voltage range, and the probability of a failure of the

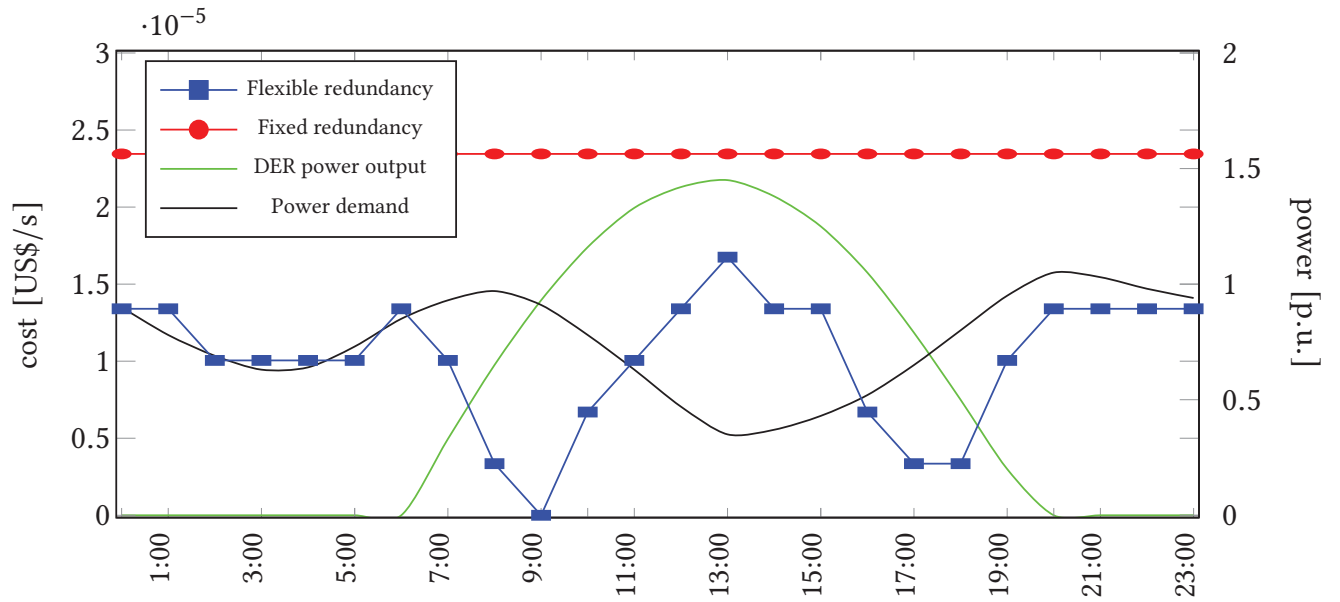


Figure 5: Cost comparison of the embeddings generated by flexible and two-fold fixed redundancy VNE programs as a function of the DER power output and power demand profiles throughout a day.

communication link between the LCU associated to bus i and the control center, upper bounded by $1 - r_{0,i}$ where $r_{0,i}$ is requested reliability of the link. This allows to assign reliability requirements to the virtual links in the VNR as

$$r^{0,i} \geq 1 - \frac{\epsilon}{P(u_i < \underline{U} \vee u_i > \overline{U})}. \quad (20)$$

In our simulation experiment, we use nominal voltage $U_0 = 33$ kV and set the real part of the admittance between two buses constant equal to $Y_i = 0.122$ S. The upper and lower limits of the allowed voltage range are 1.1 and 0.9 p.u., respectively. Powers are determined in p.u. relative to $P_{base} = 5$ MVA. The installed capacity c_i was set to 1 p.u. for every bus i in the feeder. The employed means $\mu_{d,i}(t)$ and $\mu_\phi(t)$ are shown in Figure 5, while we set $\sigma_{d,i}^2 = 0.1$ throughout and $\sigma_\phi^2 = 0.5$. The maximum probability of undetected voltage violations ϵ is set to 0.005.

Note that the modeling that was presented for a line topology here, see e.g. (17), can straight-forwardly be extended to radial grids that are commonly found at the distribution level.

5.3 Experimental Results

Figure 4 shows the voltage levels of the 15 buses at 13:00 during maximum photovoltaic in-feed. Since the photovoltaic production exceeds the demand in expectation, the voltage levels rise on average over the length of the feeder. The exact voltage levels for each bus depend on the exact realization of the uncertain demands and the irradiation factor. The resulting voltage uncertainty increases over the length of the feeder as well. We observe that the first two buses do not require communication link redundancy since the probability of over-voltage is too low here. This changes with larger distance from the slack node and requires more reliable communication, i.e. redundancy. Due to bandwidth limits of the

LoRaWAN gateway, not all nodes can use this cheap technology. Instead, where available, some nodes need to use the GPRS service as a backup.

Figure 5 presents the specific operational communication costs for the developed flexible redundancy approach and two-fold fixed redundancy over the course of one day. Similarly, to Figure 4, a high demand for redundant backup results in times where the photovoltaic demand is likely to be low and the demand high, however, this time due to potential under-voltages. In the transition times, when demand and photovoltaic production are likely to balance less redundancy is required. Compared to the two-fold redundancy method, the flexible redundancy VNE algorithm produces a cheaper embedding for every time step, with average savings of 57% of the embedding cost when using the fixed redundancy approach. This value corresponds to a reduction of US\$ 27/day or US\$ 9,855/year of the operating cost for each MV feeder in the distribution grid.

6 DISCUSSION & CONCLUSION

In this paper, we have developed an approach for VNE that generates a flexible degree of redundancy for each VNR element. To this end, we consider both the provided reliability of the SN and the reliability requirement of the VNR elements. The algorithm is formulated as a mixed integer program utilizing scenario-based chance constraints. Large changes in consecutive embeddings with slightly different SN and/or VNR parameters are avoided by adding an adaptation cost to the objective function.

The developed program is integrated into a voltage control scheme to reduce communication costs while guaranteeing high reliability services. The reliability requirement of the virtual links were set accordingly to a probabilistic model of the feeder's power flow. The experiment shows that the proposed program yields

cheaper mappings than VNE with fixed degree of redundancy, by generating just enough redundancy for the requirements.

The described use case presents a timely application of smart grids, as more and more renewables need to be integrated into the distribution grids and intelligent, communication-enhanced OLTC transformers are becoming more common. This is especially true for rural areas with long feeders, lots of renewable energy production, and often only patchy availability of cellular communication infrastructure. In this context, redundancy may be required for safe grid operation, as there is a high probability of violation of the voltage limits and the few available communication links are often unreliable.

While we applied our approach to voltage control in this paper, many more applications of VNE with soft or hard reliability requirements could benefit from this flexible redundancy generation approach. Candidates are substation automation in power systems or other CPS applications beyond power. Flexible redundancy generation enables efficient usage of the underlying communication networks and adapts well in dynamically changing environments, thus saving manual configuration efforts.

Future work might combine this proactive communication protection scheme with reactive measures, deciding automatically on the best mix for the given specifications.

ACKNOWLEDGMENTS

This work has been funded by the German Research Foundation (DFG) as part of the projects X2 and B4 within the Collaborative Research Center (CRC) 1053 – MAKI.

REFERENCES

- [1] Rajat Chaudhary, Gagangeet Singh Aujla, Sahil Garg, Neeraj Kumar, and Joel JPC Rodrigues. 2018. SDN-enabled multi-attribute-based secure communication for smart grid in IIoT environment. *IEEE Transactions on Industrial Informatics* 14, 6 (2018), 2629–2640.
- [2] Shihabur Rahman Chowdhury, Reaz Ahmed, Md Mashrur Alam Khan, Nashid Shahriar, Raouf Boutaba, Jeebak Mitra, and Feng Zeng. 2016. Dedicated protection for survivable virtual network embedding. *IEEE Transactions on Network and Service Management* 13, 4 (2016), 913–926.
- [3] Xinshu Dong, Hui Lin, Rui Tan, Ravishankar K Iyer, and Zbigniew Kalbarczyk. 2015. Software-defined networking for smart grid resilience: Opportunities and challenges. In *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security*. ACM, 61–68.
- [4] Nicholas Etherden and Math HJ Bollen. 2014. Overload and overvoltage in low-voltage and medium-voltage networks due to renewable energy—some illustrative case studies. *Electric Power Systems Research* 114 (2014), 39–48.
- [5] Shimon Even, Alon Itai, and Adi Shamir. 1975. On the complexity of time table and multi-commodity flow problems. In *16th Annual Symposium on Foundations of Computer Science (sfcs 1975)*. IEEE, 184–193.
- [6] Andreas Fischer, Juan Felipe Botero, Michael Till Beck, Hermann De Meer, and Xavier Hesselbach. 2013. Virtual network embedding: A survey. *IEEE Communications Surveys & Tutorials* 15, 4 (2013), 1888–1906.
- [7] Natasha Gude, Teemu Koponen, Justin Pettit, Ben Pfaff, Martín Casado, Nick McKeown, and Scott Shenker. 2008. NOX: Towards an Operating System for Networks. *ACM SIGCOMM Computer Communication Review* 38, 3 (July 2008), 105–110.
- [8] Bo Han, Vijay Gopalakrishnan, Lusheng Ji, and Seungjoon Lee. 2015. Network function virtualization: Challenges and opportunities for innovations. *IEEE Communications Magazine* 53, 2 (2015), 90–97.
- [9] Sandra Herker, Ashiq Khan, and Xueli An. 2013. Survey on survivable virtual network embedding problem and solutions. In *International Conference on Networking and Services, ICNS*, 99–104.
- [10] Augustine Ikpehai, Bamidele Adebisi, and Khaled M Rabie. 2016. Broadband PLC for clustered advanced metering infrastructure (AMI) architecture. *Energies* 9, 7 (2016), 569.
- [11] Raj Jain and Subharthi Paul. 2013. Network virtualization and software defined networking for cloud computing: a survey. *IEEE Communications Magazine* 51, 11 (2013), 24–31.
- [12] Dong Jin, Zhiyi Li, Christopher Hannon, Chen Chen, Jianhui Wang, Mohammad Shahidehpour, and Cheol Won Lee. 2017. Toward a cyber resilient and secure microgrid using software-defined networking. *IEEE Transactions on Smart Grid* 8, 5 (2017), 2494–2504.
- [13] Md Mashrur Alam Khan, Nashid Shahriar, Reaz Ahmed, and Raouf Boutaba. 2016. Multi-path link embedding for survivability in virtual networks. *IEEE Transactions on Network and Service Management* 13, 2 (2016), 253–266.
- [14] Murat Kuzlu, Manisa Pipattanasomporn, and Saifur Rahman. 2014. Communication network requirements for major smart grid applications in HAN, NAN and WAN. *Computer Networks* 67 (2014), 74–88.
- [15] Xiaohu Liu, Andreas Aichhorn, Liming Liu, and Hui Li. 2012. Coordinated control of distributed energy storage system with tap changer transformers for voltage rise mitigation under high photovoltaic penetration. *IEEE Transactions on Smart Grid* 3, 2 (2012), 897–906.
- [16] Pin Lv, Xudong Wang, Yang Yang, and Ming Xu. 2013. Network virtualization for smart grid communications. *IEEE Systems Journal* 8, 2 (2013), 471–482.
- [17] Elias Molina, Eduardo Jacob, Jon Matias, Naiara Moreira, and Armando Astarloa. 2015. Using software defined networking to manage and control IEC 61850-based systems. *Computers & Electrical Engineering* 43 (2015), 142–154.
- [18] Open Networking Foundation. 2012. Software-Defined Networking: The New Norm for Networks. *ONF White Paper* (2012).
- [19] Panagiotis Papadimitriou, Olaf Maennel, Adam Greenhalgh, Anja Feldmann, and Laurent Mathy. 2009. Implementing network virtualization for a future internet. In *20th ITC Specialist Seminar on Network Virtualization Proceedings*.
- [20] Juha Petajarvi, Konstantin Mikhaylov, Antti Roivainen, Tuomo Hanninen, and Marko Pettissalo. 2015. On the coverage of LPWANS: range evaluation and channel attenuation model for LoRa technology. In *2015 14th International Conference on ITS Telecommunications (ITST)*. IEEE, 55–59.
- [21] Ragnathan Rajkumar, Insup Lee, Lui Sha, and John Stankovic. 2010. Cyber-physical systems: the next computing revolution. In *Design Automation Conference*. IEEE, 731–736.
- [22] Daniel Reich. 2013. A linear programming approach for linear programs with probabilistic constraints. *European Journal of Operational Research* 230, 3 (2013), 487–494.
- [23] Timo Ruohonen, Leena Ukkonen, Mikael Soini, I. Sydanheimo, and Markku Kivikoski. 2004. Quality and reliability of GPRS connections. In *First IEEE Consumer Communications and Networking Conference, 2004. CCNC 2004*. IEEE, 268–272.
- [24] Nashid Shahriar, Reaz Ahmed, Shihabur Rahman Chowdhury, Aimal Khan, Raouf Boutaba, and Jeebak Mitra. 2017. Generalized recovery from node failure in virtual network embedding. *IEEE Transactions on Network and Service Management* 14, 2 (2017), 261–274.
- [25] Saravut Yaiparaj, Fotios Harmantzis, and Vinoth Gunasekaran. 2008. On the economics of GPRS networks with Wi-Fi integration. *European Journal of Operational Research* 187, 3 (2008), 1459–1475.
- [26] Jianchao Zhang, Boon-Chong Seet, Tek-Tjing Lie, and Chuan Heng Foh. 2013. Opportunities for software-defined networking in smart grid. In *2013 9th International Conference on Information, Communications & Signal Processing*. IEEE, 1–5.