

Benchmarking Realistic Synthetic Instances Against a Large-Scale District Heating Network: A Multi-Objective Optimization Study for Berlin

Annika Buchholz^a, Stephanie Riedmüller^a, Matthew Passage^a, Janina Zittel^a

^a*Zuse Institute Berlin, Berlin, Germany, buchholz@zib.de (CA), riedmueller@zib.de, passage@zib.de, zittel@zib.de*

Abstract:

Decarbonizing urban energy systems requires optimization approaches capable of handling the operational complexity of large-scale district heating networks. However, existing studies typically focus on a single, real-world network, which makes results difficult to compare and limits the transferability of insights to other systems. To overcome these challenges and enable a more systematic evaluation, realistic synthetic instances play an important role: they provide controlled, reproducible environments for testing optimization algorithms independent of a specific case study while still capturing essential structural and temporal characteristics of real systems. When designed carefully, such instances enable systematic benchmarking, facilitate methodological development, and support comparative studies across different algorithms and modeling choices.

In this work, we generate a suite of large-scale synthetic instances for multi-objective optimization of district heating systems. The instances are openly available in the form of the underlying network topology in a JSON format and also as a mixed integer program (MIP) as MPS files to enable modeling and algorithmic benchmarking. They are constructed using a transparent procedure that allows to recreate or extend the given instances or apply the same procedure to other network based problems. We apply this methodology to Berlin's district heating network, the most complex in Western Europe, formulating a tri-objective mixed-integer model focused on unit commitment over a horizon of up to 25 years with a 4-hour temporal resolution. The corresponding real-world instance serves as a reference point to demonstrate that the generated datasets reflect key behaviors of actual district heating operations.

A computational study provides a detailed comparison between the synthetic instances and the real-world Berlin data, showing under which conditions the generated instances reproduce realistic optimization characteristics. Furthermore, we investigate which features make the resulting models computationally challenging. The findings highlight how well-designed synthetic instances can support robust benchmarking practices and enable meaningful assessment of (multi-objective) optimization methods for large-scale district heating systems.

Keywords:

District heating, Test-suite generation, Unit commitment, Mixed-integer programming, Multi-objective optimization.

1. Introduction

While district heating is recognized as a cornerstone of the thermal transition, the field lacks a standardized, scalable benchmark for evaluating optimization methods across diverse network topologies and investment horizons. The highly complex mixed-integer optimization problems involved in planning and designing large-scale district heating networks feature strong temporal coupling and heterogeneous system components. Such models quickly become computationally challenging when integrating operational, structural, and investment decisions [1]. Despite methodological advances, empirical evaluation remains limited: most studies rely on single case studies based on proprietary data and simplified settings, restricting reproducibility and cross-study comparability [2]. At the same time, increasing model detail through sector coupling, high-resolution temporal modeling, and

multi-objective formulations further exacerbates computational difficulty [3]. Due to a lack of established benchmark libraries, systematic benchmarking for district heating and multi-energy systems is limited, hindering the analysis of structural drivers of problem difficulty [1, 4]. To address this gap, we generate structured synthetic instance families derived from a real-world reference system in Berlin. By varying structural and parametric properties while preserving key system characteristics, the approach enables controlled, reproducible, and realistic computational experiments. Our main contributions are

- the methodology to generate a meaningful set of instances based on a given instance of a multi-energy network, that vary a predefined set of parameters and structural properties, while maintaining key features,
- a computational study based on the district heating network of Berlin, including an evaluation of which properties make the problem difficult,
- the publication of the network topology of the generated instances and mixed-integer programs for a corresponding unit commitment problem (Zenodo project 19550701 [5]).

To support diverse research objectives, we provide network topologies that remain agnostic to specific model constraints; meanwhile, the MPS files represent a finalized optimization instance of the DHOT model [6], serving as a standardized testbed for benchmarking advanced solution heuristics.

2. Model base

In the following, we consider multi-energy systems based on a network structure. The system is represented as a directed, unweighted graph $G = (V, A)$, where V denotes the set of nodes and A the set of directed arcs connecting these nodes. The arcs represent feasible flows of resources such as heat, fuel, power, steam, and emissions between nodes, for example, via pipelines or transmission lines. The nodes correspond to system components (assets) that fulfill different functional roles: they may convert one set of resources into another (converters), store a resource (storage), enable the import or export of a resource (markets), collect and distribute flows (balance nodes), or specify the demand for a given resource. Fig. 1 illustrates such a network structure using the example of two production sites of the district heating network.

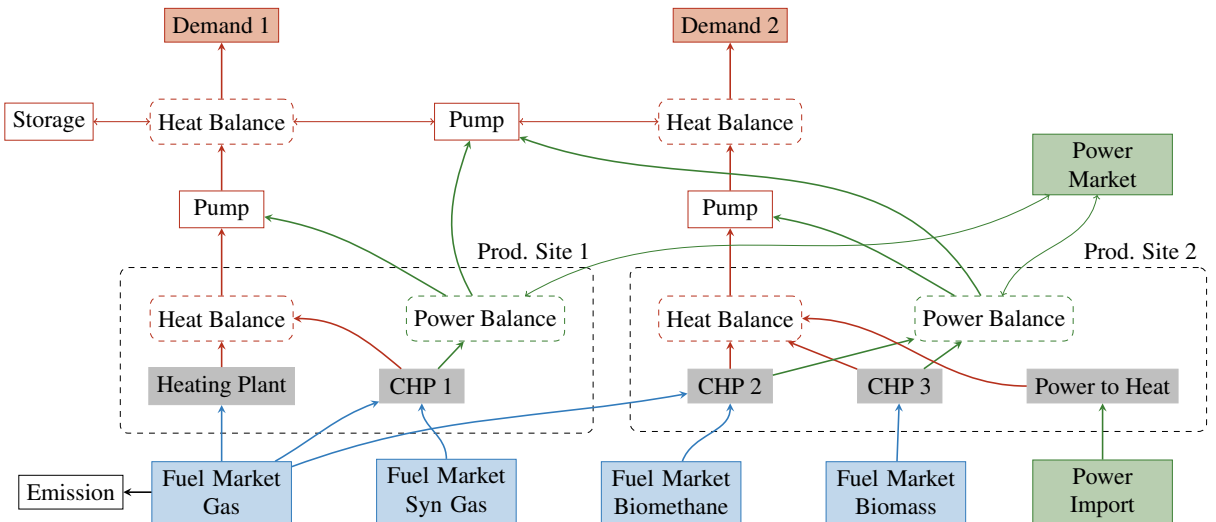


Figure 1: Simplified illustration of an instance with two production sites (fuel (blue), heat (red), power (green)), emission (black).

3. Instance generation

We aim to design the framework in a way that enables systematic scalability, controlled structural variation, and realistic parametrization. To achieve this, we introduce a transparent and reproducible

procedure for generating realistic synthetic multi-energy instances. This process defines component configurations, constructs network topologies, generates time-series data, and exports the resulting optimization models in standardized formats.

3.1. Generator parameters

The generator is controlled by configurable global and temporal parameters, network component parameters, structural parameters, and stochastic and scenario parameters, summarized in Table 1.

Table 1: Parameters of instance generation

Category	Parameter Group	Key Parameters & Features
Global & Temporal	Scope & Economics	Planning horizon, temporal resolution, constant discount rate, and constant inflation rate.
Network components	Demand	Number of nodes, associated heat demand time series, transport capacity/costs between heat demand nodes.
	Storage	Number of thermal storage units, connected to a local heat balance node through inflow/outflow heat arcs.
	Converters	Number of converters, fuel-to-heat, combined heat and power plant (CHP), and power-to-heat units, characterized either by a fixed conversion efficiency or by a piecewise-linear conversion function, each requires one or two specific input fuels.
	Markets	Number of fuel supply options (gas, oil, coal, biomass, etc.) with distinct price series, emission certificates, electricity import/export nodes.
Structural	Production Sites	Number of sites (urban districts), each conversion unit is assigned to exactly one site, local heat/power balance nodes enable cost-free internal heat exchange.
	Fuel & Power Transport	Stochastic capacity constraints (λ_{fuel}), stochastic transport costs (κ_{fuel}).
	Heat Transport	Site-to-demand connectivity, heat transport nodes, stochastic capacity (λ_{heat}), and power-driven transport (κ_{heat}).
Stochastic & Scenario	Variability	Conversion efficiencies, storage losses, transport costs, capacity levels, and ramping limits.
	Time Series	Emission factors/prices, synthetic temperature, demand, fuel, and power price profiles.

3.2. Generation of components

All technical parameters for conversion units and storage units, including capacities, efficiencies, losses, ramping limits, minimum up- and down-times, and activation costs are generated randomly for each node during instance generation and remain constant over the planning horizon.

The conversion of fuel to heat, fuel to heat and power, and power to heat is represented by a characteristic curve defined by a source series x and one or two target series y . These series specify the defining points of the piecewise-linear input-output relation. For one-to-one converters (e.g., fuel to heat and power to heat), the curve is constructed by using the beforehand randomly sampled minimum and maximum output capacities c_{min} and c_{max} . These capacities are scaled by random factors drawn from

a uniform distribution \mathcal{U} of a predefined interval $\alpha_{\min}, \alpha_{\max} \sim \mathcal{U}(\underline{\alpha}, \overline{\alpha})$, resulting in

$$x = \left[\frac{c_{\min}}{\alpha_{\min}}, \frac{c_{\max}}{\alpha_{\max}} \right], \quad y = [c_{\min}, c_{\max}].$$

For one-to-two converters (e.g., combined heat and power units), the second target series is generated by scaling c_{\min} and c_{\max} with two factors drawn from distinct distributions $\beta_{\min} \sim \mathcal{U}(\underline{\beta}_1, \overline{\beta}_1)$ and $\beta_{\max} \sim \mathcal{U}(\underline{\beta}_2, \overline{\beta}_2)$. For a user-defined subset of converters, conversion ratios are used instead of characteristic curves. The conversion ratio ρ is sampled independently from a predefined interval, $\rho \sim \mathcal{U}(\underline{\rho}, \overline{\rho})$, and the conversion is modeled as $x_t^{\text{out}} = \rho x_t^{\text{in}}$, where x_t denotes the flow variable at time $t \in T$. To ensure feasibility, the total heat production capacity is enforced to be sufficient to cover the aggregated heat demand.

Emissions are modeled using an emission factor r_f sampled from $r_f \sim \mathcal{U}(\underline{r}_f, \overline{r}_f)$. Emissions are linked to an emission certificate node with a certificate price $r_c \sim \mathcal{U}(\underline{r}_c, \overline{r}_c)$, resulting in emission costs $r_f \cdot x_{\text{fuel}} \cdot r_c$, where x_{fuel} denotes the consumed fuel.

Transport costs are modeled as nodes with technology-specific characteristic curves that map transported fuel flow (source series) to cost (target series), generated from predefined intervals. Transportation capacity limits are likewise represented as nodes, using values drawn from a uniform distribution over a specified interval. For the transportation of heat, heat pumps that require electricity are required. They follow a similar approach to converters, with characteristic curves defined analogously but using distinct parameter ranges.

3.3. Construction of network topology

The topology is generated algorithmically using the following procedure.

1. Converters are assigned to production groups by uniform random sampling. Consequently, the number of converters per group may vary substantially, reflecting the heterogeneous structure of real-world energy systems.
2. Each production group is connected randomly to exactly one demand node via heat balance nodes.
3. The heat balance nodes associated with a given demand are interconnected through transportation cost nodes and heat pumps. With probability $\lambda_{\text{heat}} \in [0, 1]$, a capacity constraint on a heat link is added, with probability $\kappa_{\text{heat}} \in [0, 1]$ the heat transport requires electricity input.
4. Each storage unit is connected to exactly one heat balance node through two directed edges representing loading and unloading.
5. Converters are connected to at most two fuel market nodes via transportation or capacity-limiting utility nodes, reflecting fuel specificity. To capture that some fuels (e.g., natural gas) are used by more technologies, market nodes are assigned different weights. Let M denote the set of market nodes and C the set of converters. Each converter receives at most two incoming market edges. Market nodes are ordered by fuel type (natural gas, synthetic gas, oil, coal, biomethane, biomass) and assigned strictly decreasing weights $w_i = |M| - i$, $i = 0, \dots, |M| - 1$. The total number of market-converter edges is bounded by $2|C|$, and edges are distributed across market nodes proportionally to w_i .
6. With probability $\lambda_{\text{fuel}} \in [0, 1]$, a capacity constraint modeled by an edge utility node is added between a fuel import market and a converter. With probability $\kappa_{\text{fuel}} \in [0, 1]$, a transport cost component is added.
7. Each fuel that produces emissions is connected to the CO_2 market node.
8. The power import node is connected to all converters that transform electricity into heat.
9. Power balance nodes are connected to the power export market, and each heat pump node is assigned an incoming edge of a power balance node.

Direct coupling of converters (e.g., fuel-to-steam followed by steam-to-heat) is not modeled explicitly. Instead, such chains are represented by a single conversion mapping. For example, a fuel-to-steam conversion φ followed by steam-to-heat conversion ϑ is contracted into a fuel-to-heat conversion $\vartheta \circ \varphi$ by a composition of their corresponding characteristic curves and conversion ratios.

3.4. Generation of temporal profiles

Temporal profiles are generated using methods tailored to their respective applications, as different types of time series exhibit distinct statistical properties and dynamics.

3.4.1. Temperature forecast profiles

Temperature profiles are generated by combining a seasonal baseline with stylized diurnal fluctuations, a technique that aligns with existing energy demand and meteorological models employing periodic harmonic components and stochastic noise [7, 8]. Let each time step $t \in T$ be associated with a month $s(t) \in S = \{1, \dots, 12\}$. For each month, a fixed average temperature level $\bar{\theta}_{s(t)}$ is specified. The temperature at time step t is then constructed as

$$\theta_t = \bar{\theta}_{s(t)} + \delta_t,$$

where the deviation term δ_t represents stylized intra-day fluctuations. Given the four-hour temporal resolution, one day is represented by six consecutive time steps. Thus, $k(t) = (t - 1) \bmod 6$ denotes the position of time step t within the corresponding day. The intra-day fluctuation is modeled as

$$\delta_t = \xi_t \left(-\cos\left(\frac{\pi}{3}k(t)\right) \right),$$

where the amplitude factor ξ_t is drawn independently for each time step from a uniform distribution, $\xi_t \sim \mathcal{U}(-4, 4)$. This construction produces a repeating sinusoidal temperature pattern over each day, centered around the monthly mean temperature and perturbed by random amplitude variations. The resulting time series, therefore, captures both seasonal temperature differences and stylized diurnal variability, while remaining lightweight and fully reproducible.

3.4.2. Heat demand profiles

Heat demand profiles are generated in two stages. First, a single aggregate heat demand profile d for the entire system is generated using a seasonally structured, correlated stochastic process. In a second step, this aggregate profile is split into multiple spatially distributed demand profiles $d_{i,t}$ for each demand node $i \in D$ representing the urban districts and each time step $t \in T$. This autoregressive structure captures short-term persistence and seasonal behavior, which is consistent with existing heat load modeling and forecasting literature [9].

For each month s , we specify a demand interval $[\underline{d}_s, \bar{d}_s]$, using the real-world Berlin model as reference. From this interval, a segment-specific reference level $\mu_s = \frac{\underline{d}_s + \bar{d}_s}{2}$ and dispersion $\sigma_s = \frac{\bar{d}_s - \underline{d}_s}{6}$ are derived. The raw aggregate demand trajectory is generated by a seasonally varying first-order autoregressive process. For all time steps $t \geq 2$,

$$d_t = \mu_{s(t)} + \phi(d_{t-1} - \mu_{s(t-1)}) + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_{s(t)}^2),$$

where $\phi \in (0, 1)$ is a fixed autocorrelation parameter controlling the smoothness of the profile and ε_t denotes the noise. The initial value d_1 is drawn uniformly from the admissible interval of the first segment. To ensure physically meaningful values and to strictly respect the user-defined seasonal ranges, all generated values are clipped to the admissible interval, $d_t = \min\{\bar{d}_s, \max\{\underline{d}_s, d_t\}\}$, for t in month s .

This aggregate profile d_t is subsequently split into $|D|$ individual demand profiles. To this end, a vector of structural shares $(\omega_1, \dots, \omega_{|D|})$ is sampled from a symmetric Dirichlet distribution with all parameters fixed to 0.3, yielding persistent heterogeneous demand sizes across nodes.

Small time-dependent perturbations are introduced to avoid perfectly fixed proportional allocations. For each time step t and each demand node i , a multiplicative noise factor $\eta_{i,t}$ is drawn independently from a Gamma distribution with parameter 50.0 and scale 1. The time-dependent shares are then computed as

$$\omega_{i,t} = \frac{\omega_i \eta_{i,t}}{\sum_{j=1}^{|D|} \omega_j \eta_{j,t}}.$$

Finally, the individual demand profiles are obtained by

$$d_{i,t} = \omega_{i,t} d_t, \quad i \in D, \quad t \in T.$$

This procedure yields structurally persistent large and small demand nodes, representing different districts, while still allowing for moderate temporal variability in their relative share.

3.4.3. Market profiles

Autoregressive models are also commonly used to represent temporal dependence in both electricity and fuel prices [10]. Fuel price profiles, as well as power import and export prices, are generated using the same seasonal process as the aggregate heat demand profile without the subsequent spatial splitting step. The generated instances include up to six fuel types: natural gas, synthetic gas, biomass, biomethane, oil, and coal. For each fuel, monthly admissible price intervals are specified based on reported European district heating fuel cost ranges [11]. Natural gas and synthetic gas are assigned wider intervals to capture stronger seasonal effects and higher short-term volatility, while biomass, biomethane, oil, and coal exhibit smoother variation. For electricity markets, prices can either be generated using the autoregressive model or sampled as a uniformly distributed ratio relative to a reference price level.

3.5. Calibration against the Berlin reference network

Although the procedure described in Section 3.2. - 3.4. is system-independent, it is calibrated against a specific real-world system to ensure that the generated instances remain realistic and transferable to proprietary datasets. In particular, the synthetic instances are calibrated to reproduce key operational and economic characteristics of real district heating systems, using a detailed reference model of the Berlin network. Calibration targets structural indicators, including the number of production sites, the share of conversion technologies, and the typical ratio of installed heat capacity to peak demand. Temporal profiles, such as seasonal demand patterns and fuel price variability, are generated to match the statistical properties observed in the reference system. While the generated parameters are not exactly aligned with the reference values, their sampling ranges are derived from the Berlin instance and extended (e.g., by approximately $\pm 10\%$) to introduce controlled variability around realistic baselines.

4. Solution approaches

For the purpose of evaluating the generated dataset, we optimize a multi-objective unit commitment problem, which is modeled as a mixed-integer program (MIP).

4.1. MIP formulation

Operation optimization of the multi-energy system is formulated as MIP, following the approach in [12]. The system consists of resources R , generating units I , storage units K , markets M , demand nodes D , and balance nodes B , represented by a network graph $G = (V, A)$ with $V = I \cup K \cup M \cup D \cup B$ and arcs A defined by the network topology. All inputs are discretized over a set of time steps T . The MIP captures unit commitment, production, and network flow decision variables:

- $z_{i,t} \in \{0, 1\}$: unit status,
- $s_{i,t} \in \{0, 1\}$: status changes,
- $x_{t,v}^r \geq 0$: incoming/outgoing flows of resource $r \in R$ at node $v \in V$,

- $h_{t,k}^r$: storage levels,
- $p_t^r, e_t^r \geq 0$: purchased and sold resources.

min (costs, emissions, – CHP heat) (MIP)

$$\sum_{i \in I} x_{t,i}^{r,out} + \sum_{k \in K} x_{t,k}^{r,out} + p_t^r = d^r + \sum_{i \in I} x_{t,i}^{r,in} + \sum_{k \in K} x_{t,k}^{r,in} + e_t^r \quad \forall r \in R \quad (1)$$

$$x_{t,i}^{r_2,out} = \varphi_{i,t}^{r_1,r_2} \left(x_{t,i}^{r_1,in} \right) \quad \forall i \in I, t \in T, r_1, r_2 \in R \quad (2)$$

$$s_{i,t} \leq z_{i,t}, s_{i,t} \leq 1 - z_{i,t}, s_{i,t} \geq z_{i,t} - z_{i,t-1} \quad \forall i \in I, t \in T \quad (3)$$

$$\sum_{\tau \in T_{i,t}^{up}} (s_{i,t} - z_{i,\tau}) \leq 0, \quad \sum_{\tau \in T_{i,t}^{down}} (s_{i,t} + z_{i,\tau} - 1) \leq 0 \quad \forall i \in I, t \in T \quad (4)$$

$$x_{t+1,i}^{r,out} - x_{t,i}^{r,out} \leq a_i^{up}, x_{t,i}^{r,out} - x_{t+1,i}^{r,out} \leq a_i^{down} \quad \forall i \in I, t \in T, r \in R \quad (5)$$

$$h_{t+1,k}^r = a_{t,k}^{loss} h_{t,k}^r + a_{t,k}^{load} x_{t,k}^{r,in} - a_{t,k}^{unload} x_{t,k}^{r,out} \quad \forall k \in K, t \in T, r \in R \quad (6)$$

$$h, x, p, e \leq c_{max}, \quad h, x, p, e \geq c_{min} \quad (7)$$

The objective minimizes operational cost and CO₂ emissions while maximizing heat production from combined heat and power plants. Key constraints include demand and resource balance (1), resource conversion via a piecewise-linear conversion map φ (2), technical unit constraints (3-5), storage dynamics (6), and capacity limits (7). For a detailed description of the model, see [6].

4.2. Algorithms

Multi-objective optimization can be approached using a variety of methods, including weighted sum, lexicographic ordering, and the ε -constraint method [13]. Dynamic variants of the ε -constraint approach have previously been applied to district heating production portfolio planning [14].

To ensure reproducibility and consistent comparison across different instances, we adopt a lexicographic optimization strategy. Objectives are prioritized sequentially: operational cost is minimized first, followed by CO₂ emissions, and finally the negative heat production of combined heat and power units. In each stage, the solution of the preceding objective is treated as a constraint, preserving optimality with respect to higher-priority objectives. This approach ensures a unique and well-defined solution direction across instances.

5. Computational study

We evaluate the performance of the proposed instance generation and solution approach under a controlled computational setup. All experiments are performed using the Gurobi Optimizer 13.0.0 [15] (build v13.0.0rc1, Linux 64-bit, Debian GNU/Linux 12) on an Intel Xeon Gold 6338 CPU at 2.00 GHz, featuring 64 physical cores (128 logical processors), with up to 32 threads utilized per run. Lexicographic optimization is applied with three objectives prioritized sequentially: operational cost, CO₂ emissions, and CHP heat production. Each objective is solved with a time limit of 10 hours, resulting in a total maximum runtime of 30 hours per instance. All solver parameters are kept at their default values, except for a MIP optimality gap of 10^{-6} . To accelerate convergence, warm starts are used between consecutive objectives. The additional lexicographic constraint is enforced with a tolerance of 10^{-1} and scaled by 10^{-3} to improve numerical stability.

5.1. Test instances

The test instances considered in this section comprise the real-world reference model and the set of synthetically generated benchmarks.

5.1.1. The real-world Berlin district heating model – B--VG2

Reference data of the real-world Berlin district heating network is provided by BEW Berliner Energie und Wärme GmbH to create instance B--VG2. The data involves the topology, parameters and time series data for the VG2 (Versorgungsgebiet 2) subgrid, which covers the East-Berlin section of the main interconnected network (Verbundnetz) in Berlin. For historical reasons, the district heating systems of East-and West-Berlin developed independently and were only later connected. Although today they are technically linked, the coupling remains limited, and the distinction between the two subgrids continues to play an operational role. The modeled subgrid supplies district heating to approximately 350,000 households. The time series dataset spans up to 25 years with a temporal resolution of 4 hours. The topology includes 5 substations with in total 5 heating plants, 3 combined heat and power plants (CHP), 2 combined cycle gas turbines (CCGT), 5 power-to-heat plants (P2H), 5 heat pumps, 5 pumps, and 1 geothermal power plant. Thermal energy storage is not included in this instance, as storage facilities are currently only available in the West Berlin subgrid. The individual heat generation assets provide capacities of up to 924 MW of thermal output and up to 502 MW of electrical power. Most units are predominantly operated with gas [16]. The data further incorporates a single demand for the whole grid and market representations for electricity, CO₂ emissions, and multiple fuels, including natural gas, synthetic gas, oil, biomass, biomethane, and coal.

5.1.2. Description of the synthetic instances – uc-- [. . .]

The synthetic instances are based on the Berlin district heating network, serving as a conceptual reference for feasible topologies and component compositions. Building upon this reference, the instances are systematically generated with variations across all stochastic parameters, enabling the exploration of a diverse set of potential configurations:

- length of the time horizon,
- number of demand nodes and production sites,
- number of converter units,
- inclusion or exclusion of storage units,
- number of fuel markets.

A total of 100 instances are generated. To ensure comparability, key global parameters are fixed across all instances, including the probability of transportation costs, the structure of network limiters, and the configuration of power markets, with exactly one import and one export market in each case.

We generate instances with a temporal resolution of four hours and horizons of either 10 or 25 years, corresponding to 21 900 and 54 750 time steps, respectively. The instances are organized into groups sharing the same structural configuration, resulting in a largely consistent topology within each group. For each configuration, 10 stochastic realizations are generated to capture variability in converter characteristics, time series (e.g., demand and prices), and other operational inputs. Table 2 summarizes the main characteristics of the benchmark instances. Each group (e.g., uc_000–uc_009) represents one configuration, while individual instances correspond to its stochastic realizations. The instances highlighted in gray serve as baseline references, both with and without storage. All other instances differ from one of these baselines by a single parameter, which is highlighted in bold. The instance B--VG2 represents the real-world Berlin reference model and serves as a baseline for comparison.

5.2. Data formats and availability

The generated dataset of 100 instances is available on Zenodo (Project 19550701) [5] in JSON format to facilitate use by energy modelers. In this format, each file encodes the topology of the multi-energy instance as a graph network, including its nodes, edges, associated parameters and time series. Additionally, MPS files are provided for direct use with optimization solvers; these files contain the full MIP formulation of each instance for each lexicographic optimization objective, along with the corresponding constraints.

Table 2: Features of the benchmark instances.

Instance	T	Demand	Prod. Sites	Converters	Storage	Fuel Markets
B-VG2	54750	3	5	20	0	2
uc_000 to uc_009	54750	3	5	20	0	2
uc_010 to uc_019	54750	1	5	20	0	2
uc_020 to uc_029	54750	3	3	20	0	2
uc_030 to uc_039	54750	3	5	10	0	2
uc_040 to uc_049	21900	3	5	20	0	4
uc_050 to uc_059	54750	3	5	20	1	2
uc_060 to uc_069	54750	5	5	20	1	2
uc_070 to uc_079	54750	3	5	20	1	4
uc_080 to uc_089	54750	1	5	20	1	6
uc_090 to uc_099	21900	3	5	20	1	2

5.3. Computational results

Fig. 2 illustrates the distribution of problem characteristics (e.g., number of variables, constraints, and branch-and-bound nodes) for the synthetic instances, with the real-world VG2 network shown as a diamond marker for reference. Presolve reductions are most effective for the first objective, while significantly fewer variables and constraints are eliminated for the second and third objectives. In lexicographic optimization, the use of a warm start when solving the MIP with an added constraint alters the problem structure; as a result, we can observe a reduction in the effectiveness of presolve. The ratio of discrete variables to the total number of variables (i.e., discrete plus continuous) remains stable across cases, but is lower for the second and third objectives. This indicates a higher share of continuous variables induced by the lexicographic constraint.

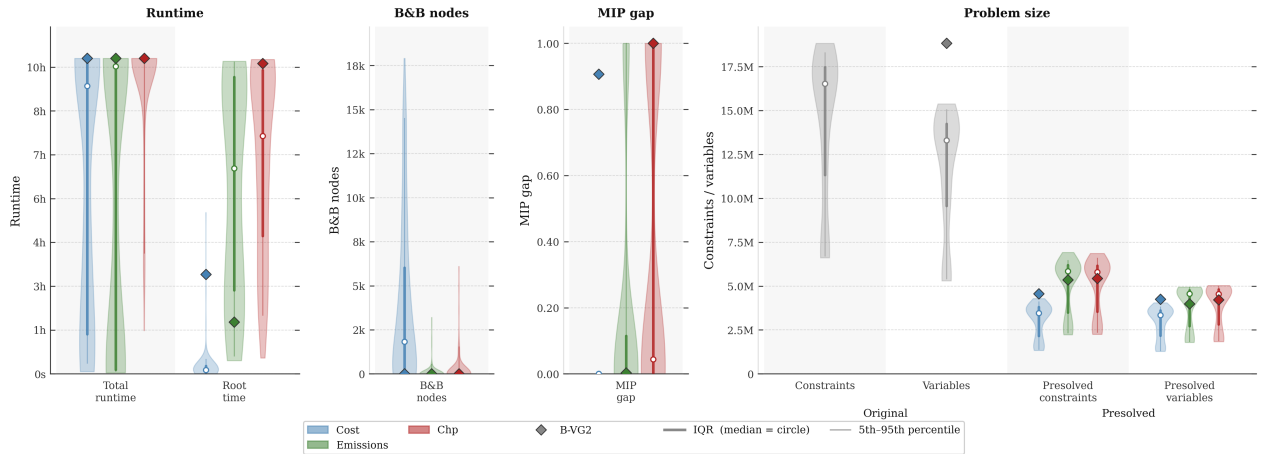


Figure 2: Distribution of characteristics over all instances for the three lexicographic objectives (1. cost, 2. emissions, 3. CHP).

From a computational perspective, the first objective yields significantly larger branch-and-bound trees, highlighting its combinatorial complexity. In contrast, the second and third objectives require more root LP iterations and exhibit a wider RHS range, indicating increased numerical difficulty. This shift is also reflected in runtime: while a large fraction of instances for the first objective (cost minimization) is solved to optimality within the 10-hour limit, this fraction decreases markedly for the subsequent objectives. Presolve time accounts for less than 1% of the total runtime across all instances. To mitigate numerical issues, the additional constraint is scaled by 10^{-3} ; without scaling, performance deteriorates substantially. Overall, the results indicate a transition from predominantly combinatorial complexity in the first objective to increased numerical difficulty in the lexicographic extensions.

Figure 3 shows the number of solved instances over time on a logarithmic scale. The two baseline

instance sets defined in Table 2 are depicted as dashed lines and serve as reference datasets. The results indicate that

- including storage increases computational effort: 30/50 instances without storage were solved within the time limit (mean: 8,100 s), compared to 22/50 with storage (mean: 13,745 s).
- instances with 10-year time series (uc_040–uc_049, uc_090–uc_099) are solved both faster and more consistently. This behavior can be attributed to the reduced number of variables and constraints. Overall, their average solving time decreases by 0.535% relative to the baseline with 25-year time series.
- reducing the number of production sites to three (uc_020–uc_029) or limiting the number of converters to ten (uc_000–uc_039) increases computational difficulty.
- changes in the number of fuel markets have mixed effects: four markets (uc_070–uc_079) slightly reduce solving time compared to the two-market baseline, while six markets (uc_080–uc_089) increase it. Overall, these effects are small and influenced by time series and stochastic parameters.
- varying the number of demands shows minor and partially counterintuitive effects: five demands (uc_060–uc_069) marginally ease the problem, while a single demand (uc_010–uc_019) increases difficulty, yet their solving times remain close to baseline values.




Figure 3: Runtime of first objective (cost) across each instance configuration.

5.3.1. Comparison with the Berlin model

The 10 baseline instances (uc_010–uc_019) exhibit structural and operational characteristics similar to the subgrid VG2 of Berlin, with a comparable network topology and differences primarily in temporal aspects (e.g., demand profiles and time-dependent parameters). In terms of nodes and edges, the Berlin instance falls within the size range of the synthetic instances. The number of variables and constraints is substantially higher due to extended modeling features and a different topological representation. However, presolve eliminates structurally implied or non-binding constraints, such that the reduced model size is consistent with that of the synthetic instances (cf. Fig. 2). Computational performance is similarly aligned: all three objectives remain unsolved within the 10-hour time limit, mirroring the behavior observed for a subset of the synthetic instances. This further supports the conclusion that the synthetic instances provide a realistic representation of the Berlin model.

6. Conclusion and outlook

We present a methodology for generating realistic synthetic benchmark instances of large-scale district heating systems, using a real-world model from Berlin as a reference. The framework is transferable to other cities and multi-energy networks. We provide three main contributions: a validated instance generation method, a computational study of factors affecting problem difficulty, a dataset of 100 topology files for reproducible benchmarking of multi-energy systems, and corresponding optimization formulations. The realistic synthetic instances exhibit wide variability in computational difficulty, demonstrating their suitability for benchmarking. The results show that computational complexity depends on temporal, structural, and formulation factors; finer temporal resolution and storage increase model size and runtime, while structural simplifications may increase combinatorial complexity. The proposed framework is particularly valuable for assessing whether an existing model, which is already challenging to solve, exhibits increased computational difficulty under topological or stochastic modifications or whether such changes have a negligible effect on computational effort.

Acknowledgments

The work for this article has been conducted in the Research Campus MODAL funded by the German Federal Ministry of Research, Technology and Space (BMFTR) (fund numbers 05M14ZAM, 05M20ZBM, 05M2025).

Nomenclature

Abbreviations

CHP combined heat and power plant
MIP mixed-integer program

Latin symbols

d, a, c, r_f, r_c, w parameter vector for demand, coefficients, capacities, emission factor and emission certificate price, weight of fuel assignment
 $G = (V, A)$ graph G with set of vertices V and set of arcs A
 p, e variables for purchased and sold resources
 R, I, K, M, D, B, C set of resources, generating units, storage units, markets, demands, balance nodes
 T, S set of time steps, months
 x, y source, target series for characteristic curve
 x, z, s, h variables for resource flow, operation, status and storage level

Greek symbols

α, β, ρ ratios
 δ, θ Temperature level, fluctuation
 ϕ, ε parameter of autocorrelation, noise
 κ, λ probabilities of capacity and cost nodes
 μ, σ temperature reference level, dispersion
 ω, η variable, noise of structural shares of the demand per district
 φ, ϑ resource conversion map
 ξ amplitude factor

Superscripts and subscripts

i/k generating/ storage unit
 in/out ingoing/ outgoing flow

r resource

t, s, k time step, monthly step, position of time step within a day

\mathcal{U} uniform distribution

References

- [1] Lambert J., Ceruti A., Spliethoff H. *Benchmark of mixed-integer linear programming formulations for district heating network design*. Energy, 308; 2024. DOI: 10.1016/j.energy.2024.132885
- [2] Sporleder M., Rath M., Ragwitz M. *Design optimization of district heating systems: A review*. Frontiers in Energy Research, 10:971912; 2022. DOI: 10.3389/fenrg.2022.971912
- [3] Guo T., Chen Y., Kriegel M. *Novel multi-level optimization of district heating systems: Managing spatial scale and equipment portfolio design*. Energy and Buildings, 317:114395; 2024.
- [4] Perea-Moreno A.J., Muñoz-Rodríguez D., Hernandez-Escobedo Q., et al. *Worldwide scientific landscape on district heating systems*. Energy Efficiency, 19:19; 2026. DOI: 10.1007/s12053-026-10419-4
- [5] Buchholz, A., Riedmüller S., Passage, M., Zittel J. *Realistic Synthetic Benchmark Instances for Multi-Objective Optimization of District Heating Systems* Available at: <https://zenodo.org/records/19550701> [accessed 13.04.2026].
- [6] Riedmüller S., Buchholz A., Zittel J. *Enhancing Multi-Energy Modeling: The Role of Mixed-Integer Optimization Decisions*. arXiv preprint arXiv:2505.14492; 2025.
- [7] Wilks D.S., Wilby R.L. *The weather generation game: a review of stochastic weather models*. Progress in Physical Geography 1999;23(3):329–357.
- [8] Hörsch J., Hofmann F., Schlachtberger D., Brown T. *PyPSA-Eur: An open optimisation model of the European transmission system*. Energy Strategy Reviews 2018;22:207–215.
- [9] Xue G., Pan Y., Lin T., Song J., Qi C., Wang Z. *District heating load prediction algorithm based on feature fusion LSTM model*. Energies 2019;12(11):2122.
- [10] Weron R. *Forecasting wholesale electricity prices: A review of time series models*. Financial Markets: Principles of Modelling, Forecasting and Decision-Making, FindEcon Monograph Series, WUL, Lodz 2008.
- [11] Henahan R. (IBS) *Publishable report on District heating cost modelling and advised billing cases*. Available at: https://www.relatedproject.eu/wp-content/uploads/2021/06/D4.6_-Publishable_R_DH_cost_modelling_and_advised_billing_cases_2_0.pdf[accessed 26.03.2026].
- [12] Clarner J.-P., Tawfik C., Koch T., Zittel J. *Network-induced Unit Commitment - A model class for investment and production portfolio planning for multi-energy systems*. ZIB-Report 22-16. Zuse Institute Berlin (ZIB), Takustr. 7, 14195 Berlin; 2022.
- [13] Ehrgott M. *Multicriteria Optimization*. Springer Berlin, Heidelberg; 2005. DOI: <https://doi.org/10.1007/3-540-27659-9>
- [14] Zittel J., Clarner J.-P., Tawfik C., Dykes M., Rivetta F., Riedmüller S. *A multi-objective optimization strategy for district heating production portfolio planning*. In: Proceedings of the 37th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2024); 2024.
- [15] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*; 2025. Available at: <https://www.gurobi.com>
- [16] Berliner Energie und Wärme GmbH *Erzeugungsanlagen* Available at: <https://www.bew.berlin/fernwaermesystem/erzeugungsanlagen/>[accessed 26.03.2026].