
Multi-horizon stochastic programming

Michal Kaut and Kjetil T. Midthun and Adrian S. Werner and Asgeir
Tomasgard and Lars Hellemo and Marte Fodstad

SINTEF Technology and Society, Trondheim, Norway
and

Norwegian University of Science and Technology, Trondheim, Norway

in: Computational Management Science. See also `BIBTEX` entry below.

`BIBTEX`:

```
@article{KautEA12,  
  author = {Michal Kaut and Kjetil T. Midthun and Adrian S. Werner and Asgeir  
Tomasgard and Lars Hellemo and Marte Fodstad},  
  title = {Multi-horizon stochastic programming},  
  journal = {Computational Management Science},  
  year = {2014},  
  volume = {11},  
  number = {1--2},  
  pages = {179--193},  
  note = {Special Issue: Computational Techniques in Management Science},  
  doi = {10.1007/s10287-013-0182-6}  
}
```

© Springer-Verlag 2013.

The original publication is available at www.springerlink.com.

Multi-horizon stochastic programming

Michal Kaut¹, Kjetil T. Midthun¹, Adrian S. Werner¹, Asgeir
Tomasgard^{1,2}, Lars Hellemo¹, and Marte Fodstad¹

¹*SINTEF Technology and Society, Department for Applied Economics and Operations
Research, Trondheim, Norway*

²*Norwegian University of Science and Technology (NTNU), Department of Industrial
Economics and Technology Management, Trondheim, Norway*

July 2012; updated June 2013

Infrastructure-planning models are challenging because of their combination of different time scales: while planning and building the infrastructure involves strategic decisions with time horizons of many years, one needs an operational time scale to get a proper picture of the infrastructure's performance and profitability. In addition, both the strategic and operational levels are typically subject to significant uncertainty, which has to be taken into account. This combination of uncertainties on two different time scales creates problems for the traditional multistage stochastic-programming formulation of the problem due to the exponential growth in model size.

In this paper, we present an alternative formulation of the problem that combines the two time scales, using what we call a *multi-horizon* approach, and illustrate it on a stylized optimization model. We show that the new approach drastically reduces the model size compared to the traditional formulation and present two real-life applications from energy planning.

Keywords: stochastic programming, multistage, energy planning, scenario tree construction

1 Introduction

Infrastructure-planning models typically focus on long-term investment strategies with time horizons of years or even decades. For many such models it is important to acknowledge how operational aspects will affect (and will be affected by) these strategic decisions: in order to find robust, flexible, and profitable solutions, it is necessary to assess how the infrastructure can be used, at what cost, and how it can respond to varying conditions. Hence, the goal is often to

design infrastructure in a way that maximizes the net present value of constructing, maintaining, and using the infrastructure, while still allowing the users to satisfy their operational targets. These aspects are on very different time scales: construction and maintenance of the infrastructure are strategic decisions with time scales of months or years, while the operational targets and costs are driven by operational decisions with time scales of days, hours, or even minutes. The main topic of this paper is how to address these complementary time scales in one optimization model.

To further complicate the matter, these planning processes are subject to substantial uncertainty: on the strategic level, there is uncertainty about future demands, prices, technology development, etc. The uncertainty on the strategic level will influence the operational decisions, and in addition there may be uncertainty in the short term. The sources of the uncertainties vary between different types of applications.

One of the standard tools for solving optimization problems with uncertainty is stochastic-programming models, in particular their scenario-tree-based formulations. With these models, it is not obvious how to combine the two time scales, without the exponential growth in model size described in Section 2. Depending on the objective function and purpose of the analysis, this problem can be addressed in different manners. We can have separate strategic and operational models and run them alternatively in a loop, adjusting the strategic model based on the results from the operational model. An example of such an approach is provided in Myklebust (2010), where a strategic and an operational model are run in sequence, as two separate entities.

There are also papers that consider aspects related to handling both time scales in one model and, partially, uncertainty on the different time scales. One example is provided by Schütz et al. (2009) who include short-term variations in a strategic model for the Norwegian meat industry. The model is a two-stage stochastic model where a supply chain is designed in the first stage and then operated under uncertain demand in the second stage. The uncertainty in demand comes from both short-term variations and long-term trends. De Jonghe et al. (2011) use an equilibrium model to study the expansion of electricity generation capacity. They integrate the short-term demand response in their strategic model by way of a representative profile and discuss the effects on flexibility of the generation capacity. Their approach, however, considers only a one-period deterministic model. Sönmez et al. (2013) analyze strategic investment decisions in liquefied natural gas transport and discuss the impact of using a stochastic model (simulation) for the throughput, i.e., at the operational level. They show that, when deciding about technology and capacity choices, also operational flexibility is important in order to cope with short-term variations and has a significant impact on profitability. Finally, Singh et al. (2009) describe a multi-stage capacity-planning problem with potentially stochastic operational ‘submodels’ and show how to solve it using Dantzig-Wolfe decomposition with variable splitting. This paper focuses on the solution method and does not describe the structure of the operational uncertainty.

In this paper, we describe a model and scenario-tree structure that allows for using operational decisions to evaluate the quality of the strategic decisions, while having uncertainty on both the strategic and the operational levels. This is done by using a ‘multi-horizon’ optimization model, where we embed important operational features directly into the strategic decision model, thus allowing an immediate evaluation of potential investment solutions. An application of the same scenario-tree structure has been presented in Hellemo et al. (2013). That paper, however, focuses

on a mathematical model developed for the natural gas industry, without an in-depth discussion of the tree structure for the considered application. In contrast, our contribution lies in a detailed discussion of the structure and its comparison to standard multi-period trees. We also provide guidance on how to populate such scenario trees and describe two applications from the area of energy planning.

The rest of this paper is organized as follows: in Section 2, we show why the standard approach to building multistage stochastic models is not well suited for models with both strategic and operational uncertainties. Then, in Section 3, we present our new multi-horizon structure that resolves the problem and illustrate its use on a stylized optimization problem. Finally, in Section 4, we describe the two example applications. We also discuss other areas where it could potentially be useful, before we conclude the paper.

2 Limitations of the standard approach

In this section, we explain why the standard approach for building scenario trees is not suitable for models that combine long- and short-term uncertainties. We start by introducing the necessary terminology for multistage stochastic-programming problems and the related scenario trees: a *period* denotes the time interval between two consecutive time-discretization points. A *stage* starts at a time point where new information is received. Per definition, the first time point in the tree starts a stage.

A *node* represents a time point where decisions are made. We define two node types, strategic nodes (illustrated in the figures with ‘●’) for long-term decisions and operational nodes (illustrated with ‘■’) for short-term decisions. A leaf node represents either a time point only or the start of the last time period. An example of the former is the evaluation of a portfolio’s value at a later point of time, while the latter could be decisions about operating a system in the last hour. In this paper, all leaf nodes start a last operational period. In the figures, this last period is represented by ‘1’. Note that we do not include a corresponding illustration of the last strategic period, since the succeeding operational nodes make it clear that the last strategic node is not momentary.

2.1 Models with one strategic stage

We start with the case where the standard approach works well: the classical two-stage stochastic programming model. There, the first stage represents the strategic decisions about the infrastructure and the second stage the operational decisions, see cases (a) and (b) in Figure 1. The only difference between the two cases is that (b) includes several operational periods; since these are deterministic (no branching there), this is still a two-stage problem.

The last case (c) is different in the sense that it is a three-stage problem, since there is one extra operational stage. This means that unlike the previous two cases that modelled operating the infrastructure under several deterministic scenarios, this one actually models its operation under uncertainty. In this way, it is a more complex model, but still relatively easy to solve as it has only one strategic period, so we avoid mixing the strategic and operational time steps.

Since this class of problems is relatively easy to both model and solve, we have to understand




Figure 1: Example of scenario trees with one strategic stage

the assumptions behind using just one strategic stage. Firstly, having only one strategic stage means that all the strategic decisions have to be found *now*, there is no way of postponing a strategic decision until we have learned more. In other words, the *option* to wait does not exist in such a model—see Christiansen and Wallace (1998) or Fleten et al. (1998) for the connection between options and stochastic optimization.

Furthermore, with only one strategic node, the model does not take into account the time it takes to implement the strategic decisions—all the strategic decisions are implemented during the first stage and used in the rest of the model.

There are many applications where this approach is sufficient. As an example, consider the case where the strategic decision concerns routes and/or timetables for some form of transport and the operational periods represent daily usage of those routes under different conditions (scenarios). Such models, using the scenario-tree structure from Figure 1(b), can be found, for example, in King et al. (2012); Lium et al. (2009); Thapalia et al. (2012a,b).

2.2 Models with multiple strategic stages

With multiple strategic stages, things become much more complicated, as illustrated in Figure 2. We see that going from the deterministic case (a) to a case with uncertainty at the strategic level (b), the number of scenarios increases, though in a manageable manner. The problem is that we still have no operational uncertainty—and how realistic is it to assume that we know all parameters three years ahead? When we add the operational uncertainty, we get a scenario tree similar to case (c). We see that the number of scenarios grows from four to 32, for a model with three strategic periods and two branches per period for both the strategic and operational nodes. Clearly, such an approach is not practical for real-life problems: if we increase the number of branches from two to ten, which might still be too few for many practical problems, we would need 100 000 scenarios. Hence, the question is whether we can come up with a way of modelling these two uncertainties that avoids this explosion in the number of scenarios; we present one such a way in the next section.

3 Multi-horizon scenario trees

Our solution to the problem from the previous section is based on the observation that strategic decisions typically do not depend directly on any particular operational scenario, but rather on the overall operational performance during the time since the previous strategic decision. This

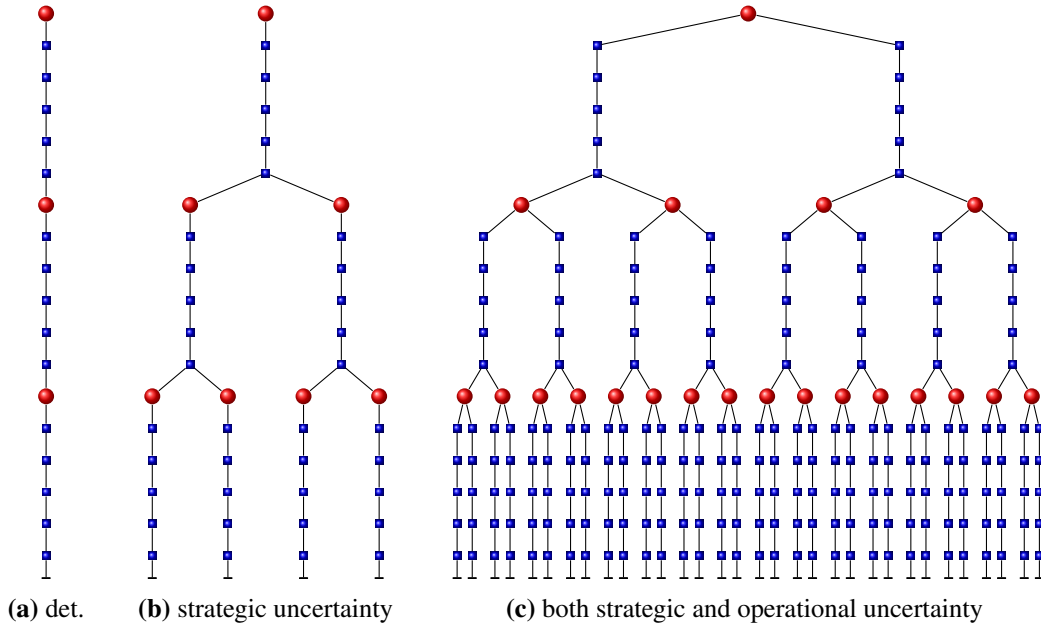


Figure 2: Models with multiple strategic stages, using standard scenario trees. This illustrates the increase of the tree size from a deterministic model (a), through a model with strategic uncertainty (b), to a model with uncertainty on both the strategic and operational level (c).

implies that, in the strategic model, it is enough to branch only between the strategic stages; the operational nodes can be seen as embedded into (or attached to) their respective strategic nodes. They are there, in a way, to check the feasibility and profitability of the decisions made in their respective strategic nodes.

This is presented in Figure 3. In this figure, tree (a) includes the same nodes as the tree in Figure 2(b), except that we interpret the nodes differently: in the new structure, the tree is constructed solely out of the strategic nodes, which have the associated operational nodes ‘embedded’ in them. The second tree in the figure corresponds directly to the tree from Figure 2(c). Finally, Figure 3(c) demonstrates that the operational part does not need to consist of a set of deterministic scenarios (or profiles), but can itself be a multistage stochastic problem. This allows for an evaluation of the infrastructure under more complex stochastic operational conditions. In this sense, Figure 3(c) provides a multi-strategic-stage equivalent of Figure 1(c).

We can see from the trees in Figure 2(c) and 3(b) that the new multi-horizon structure brings a dramatic reduction in the size of the resulting tree. How large this difference will be depends on the numbers of operational and strategic periods and the number of successors in the branchings. Table 1 illustrates the size of this reduction, where the different base in the exponential function is the main driver. This transfer to the problem size in terms of number of variables and constraints, since variables and constraints in a period are duplicated for each node in that period.

Since the new structure significantly reduces tree and problem sizes, it is natural to ask about the cost of this reduction. The answer depends on the structure of the optimization model.

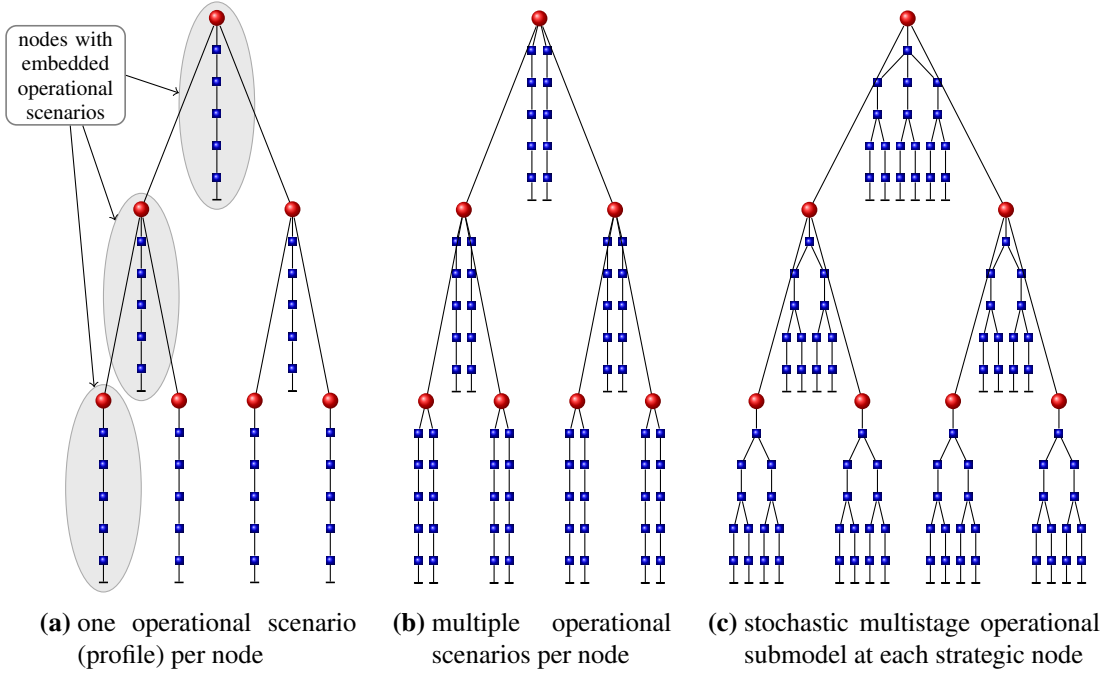


Figure 3: Multi-horizon scenario trees with strategic uncertainty.

Table 1: Numbers of strategic and operational nodes for the standard tree in Figure 2(c) and the multi-horizon tree in Figure 3(b). P_s is the number of strategic periods, P_o the number of operational periods per strategic period. B_s and B_o are the numbers of strategic and operational successors in each strategic and operational branching, respectively. We assume that these numbers are constant throughout the tree and that $B_s > 1$.

	Standard tree	Multi-horizon tree
# strategic nodes	$\frac{1-(B_o B_s)^{P_s}}{1-B_o B_s} \approx (B_o B_s)^{P_s}$	$\frac{1-(B_s)^{P_s}}{1-B_s} \approx B_s^{P_s}$
# operational nodes	$P_o B_o \frac{1-(B_o B_s)^{P_s}}{1-B_o B_s} \approx P_o B_o (B_o B_s)^{P_s}$	$P_o B_o \frac{1-(B_s)^{P_s}}{1-B_s} \approx P_o B_o B_s^{P_s}$

The new approach is exact if the following conditions are satisfied: firstly, strategic uncertainty must be independent of the operational uncertainty and the strategic decisions must not depend on any particular operational decisions. Without these requirements, we would not be able to have a single strategic decision following two or more operational scenarios. Secondly, the first operational decision in a strategic node cannot depend on the last operational decision from the previous period—in the proposed structure; there is no connection between operational scenarios of two consecutive strategic nodes.

The first requirement is quite natural and can be expected to be fulfilled in many situations, especially if the difference between the strategic and operational time scales is big (years vs. hours). The second requirement, on the other hand, is harder to fulfil exactly; for example, if we have strategic periods coinciding with calendar years, together with one-hour operational periods, we require that the operational decisions at 00:00 on January 1 do not depend on the operational decisions from 23:00 on December 31 from the previous year. It follows that the proposed structure will, in most cases, be an approximation of the ‘standard trees’ from the previous section. How good an approximation will be very much case-dependent. As an example, consider a power producer with hydro power plants; clearly, the water reservoirs introduce a memory aspect to the operational model, invalidating the second requirement above. However, at least in regions with cold winters, it can be expected that the reservoirs will be at their minimum level by the end of winter, in most scenarios. Hence, if we place the strategic decisions at the end of winter, the approximation error will be very small.

In the application discussed in Section 4.1, EnRiMa, the links between the operational decisions and the following investment periods are weak, and we can rely on such a scenario tree structure. On the other hand, in the Ramona model presented in Section 4.2, the decisions in the operational periods may influence the decision space for the following investment periods, such that the tree structure will be an approximation.

3.1 Representative sub-periods

Even if the multi-horizon trees from Figure 3 are significantly smaller than the standard trees from Figure 2, they have the same number of operational periods. With yearly strategic decisions and hourly operational resolution, each operational scenario consists of $365 \times 24 = 8760$ periods—this might make our model intractable. Our solution is to split the interval between strategic decisions into several sub-periods. For the ease of explanation, we describe the idea on the case of yearly strategic decisions and hourly operational periods.

Let us first assume that we have a system that is being actively operated only during the day and reverts to some ‘default state’ during the night. In such a case, there would be no link between two consecutive days, i.e. the operational model could be run for each day separately. For the scenario-tree structure from Figure 3(a), this would mean that a one-year scenario with 8760 hourly periods could be equivalently interpreted as 365 daily scenarios with 24 hourly periods. In our idealistic example with no links between days, this restructuring provides an equivalent model formulation. If there are links between consecutive days, usually in some form of storage, then we get only an approximation. In such cases, one may split where the links between consecutive operational periods are the weakest. For example, there might be strong links between weekdays, but much weaker during weekends; hence, we might consider splitting

the one-year interval into 52 weeks instead.

This restructuring does not increase the problem size when using the multi-horizon scenario tree. With the restructured tree we can now employ different scenario reduction techniques to reduce the number of operational scenarios (see, for instance, Römisch (2009) for an overview), resulting in a suitable number of representative operational sub-periods.

3.2 Generating values for multi-horizon scenario trees

Once we have decided the form of the scenario tree and the decomposition of the operational periods, we are ready to build the full multi-horizon scenario tree and fill it with values. We can start by generating the tree of the strategic nodes: this is a standard multistage scenario tree, so one can use any standard method—see, for example, Dupačová et al. (2000) or Høyland and Wallace (2001). The only non-standard aspect is that, because of the very long time span of these trees, one often cannot use historical data and has to rely on some prediction methods and/or expert opinions to get the parameters of the distributions.

Next, we have to generate values for the operational scenarios. We distinguish between operational parameters that have the same distributions in all strategic periods and parameters that change over time (either deterministically or stochastically). An example for the former is weather, at least when we ignore climate-change effects. For these parameters, we can use the same set of operational scenarios in all the strategic nodes. Parameters that evolve also on the strategic scale such as spot market energy prices require a different treatment: we have to separate the long-term trends from the short-term uncertainty. One way of doing this is to express the operational values in terms relative to the corresponding value in the strategic node. We can then use the same operational scenario in all the strategic nodes (given that the distribution itself is independent of the long-term trends).

3.3 Illustrative example

We illustrate the approach on a small two-stage model, inspired by the EnRiMa model presented in the next section: we consider installing photovoltaic (PV) panels on a building and want to find out what capacity we should install and when. To be able to evaluate the value of the panels, we need to model how they help to cover the electricity demand on an hourly basis. We make the strategic periods one year long, so we have to have operational profiles to model how the production of the PV panels changes throughout the year. In our simple example, we model this using three profiles: *winter* and *summer* representing 90 days each, and one for *the rest* of the year.

The model itself is very simple; we have only two decision variables in each strategic node n : x_n denoting the capacity to be installed and y_n the total installed capacity, both measured in kW. In the operational profiles, we have, in addition, variables $z_n^{p,t}$ denoting the amount of purchased electricity at time t of profile p , in kWh. The objective is to minimize the expected cost, which




Figure 4: Example: the strategic tree, its parameters, and the complete multi-horizon scenario tree. Note that we show only the first four operational periods in (c).

gives the following model:

$$\min \sum_{n \in \mathcal{N}} Pr_n \left(CI_n x_n + \sum_{p \in \mathcal{P}} W^p \sum_{t \in \mathcal{T}} CE_n^{p,t} z_n^{p,t} \right) \quad (1)$$

$$\text{s. t.} \quad y_n = y_{\text{Pa}(n)} + x_n \quad \forall n \in \mathcal{N} \quad (2)$$

$$\Delta_n^{p,t} R_n^{p,t} y_n + z_n^{p,t} \geq D_n^{p,t} \quad \forall n \in \mathcal{N}, p \in \mathcal{P}, t \in \mathcal{T} \quad (3)$$

In the objective function (1), Pr_n is the probability of node n , CI_n is the PV installation cost in €/kW, W^p is the weight of the profile (90 for winter and summer, 185 for ‘rest’), and $CE_n^{p,t}$ is the electricity cost at time t of profile p , in €/kWh. Constraints (2) keep track of the installed capacity. There, $\text{Pa}(n)$ denotes the parent node of n ; since the root node does not have a parent, $y_{\text{Pa}(n)}$ is defined as zero there.¹ Finally, constraints (3) ensure that we have enough power to satisfy the demand $D_n^{p,t}$, given in kWh, at each period, profile and node. There, $\Delta_n^{p,t}$ is the duration of the operational periods, in our case 1 hour, and $R_n^{p,t}$ is a factor specifying what percentage of the nominal capacity the panel actually produces at the given hour.

Generating values for the stochastic parameters

To keep the model as simple as possible, we have only three strategic nodes: the root node representing ‘now’ and two nodes representing two different scenarios one year ahead—see Figure 4(a). When we add the three operational profiles, we get the multi-horizon tree in Figure 4(c).

We start by generating values for the strategic stochastic parameters. In our case, there is only one such parameter, namely the PV installation costs CI ; its values in the three strategic nodes are presented in Figure 4(b).

However, we need additional strategic parameters for modelling operational parameters that evolve on the strategic time scale, as described at the end of Section 3.2. The first such a parameter is the electricity price CE , which we want to be stochastic also on the strategic scale. This is accomplished by introducing additional parameters SCE modelling the long-term averages of the price—see Figure 4(b). The actual operational values are then computed as

$$CE_n^{p,t} = SCE_n \times ICE^{p,t},$$

¹Actually, $y_{\text{Pa}(n)}$ in the root node represents the currently installed capacity—which we assume to be zero.

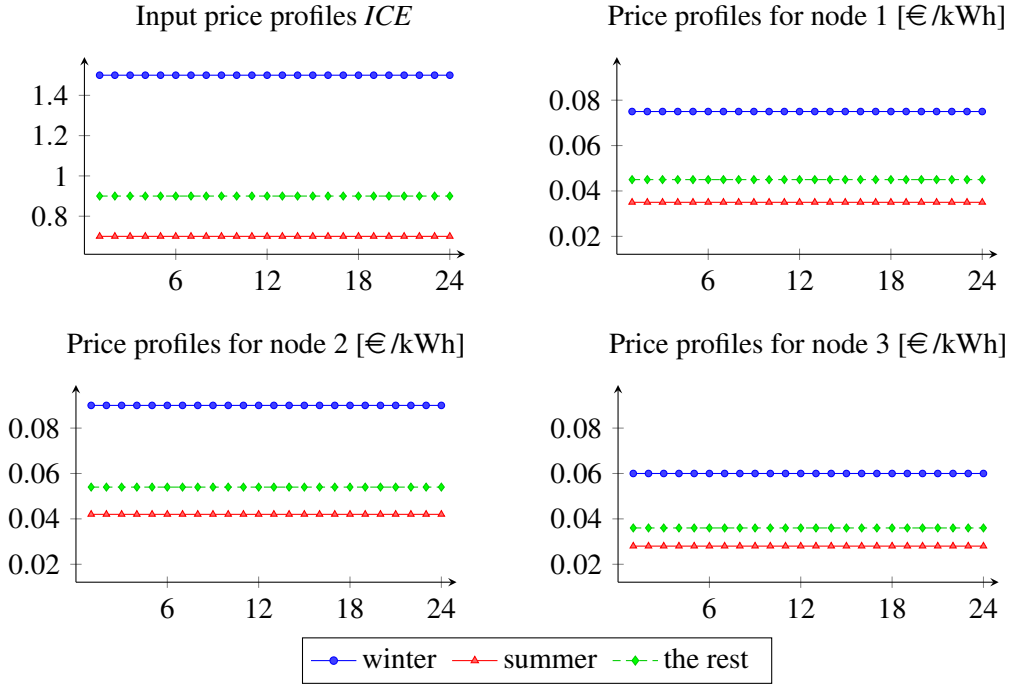


Figure 5: Operational profiles for electricity prices. The first plot shows the input profiles ICE , given as relative values, the remaining three plots show the resulting profile values CE .

where $ICE^{p,t}$ are dimensionless multipliers modelling the daily price profiles for each $p \in \mathcal{P}$. In our case, we let the price profiles be constant throughout the day but varying in the course of the year, with the highest prices in winter—see Figure 5. This corresponds to the situation in Norway, where electricity prices typically do not depend on the hour, but change often throughout the year.

For the demand values D , we simplify the situation slightly by assuming that the long-term development is known (deterministic). This is modelled by a parameter SD^τ , where τ is the strategic time period. In our case, we have used $SD^\tau = 10$ kWh. The operational values are, therefore, computed as

$$D_n^{p,t} = SD^{\text{Per}(n)} \times ID^{p,t},$$

where $\text{Per}(n)$ denotes the strategic period of node n and $ID^{p,t}$ are, again, the multipliers from the operational profiles. Unlike the electricity prices, we let the demand vary throughout the day, as shown in Figures 6.

Finally, the PV-production factors R are assumed to be constant in the long term. Since these are relative values by definition, we get simply

$$R_n^{p,t} = IR^{p,t},$$

so we do not need any extra strategic parameter. The profiles are presented in Figure 7—the values correspond to Bergen, Norway, as provided by the PVWatts™ calculator.²

²The PVWatts™ calculator was developed by the National Renewable Energy Laboratory and is available from

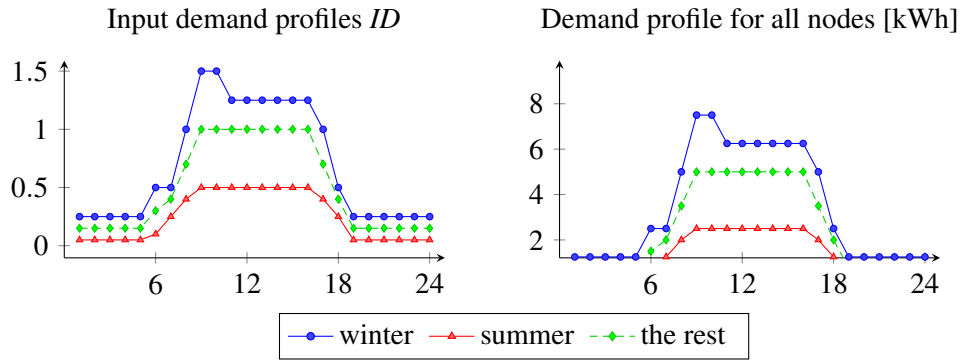


Figure 6: Demand profiles. The left plot shows the input profiles ID , given as relative values, the right plot the resulting profile values D in kWh. The long-term demand is given as $SD=10$ kWh.




Figure 7: Profiles for the fraction of PV production (IR). This is given as data and repeated in all the strategic nodes.

Numerical results

With the given data, the optimal solution is to not install any PV panels and buy all the electricity. This is hardly surprising, given the assumptions of the model: the time horizon is too short to make the panels profitable for the considered location. We also do not allow for the electricity to be sold and ignore end-of-horizon issues, i.e. we do not assign any value to the PV panels at the end of the last period. To create a more realistic investment model, we would also have to include aspects such as discount rates and depreciation. Instead, the presented model is meant as a demonstration of how to populate a multi-horizon tree—something that would be less transparent with a more complex model.

<http://www.nrel.gov/rredc/pvwatts/>.

4 Applications

4.1 Efficient energy usage – EnRiMa

EnRiMa (Energy Efficiency and Risk Management in Public Buildings) is a research project funded by the European Commission via the 7th Framework Programme (FP7). In the context of our paper, we are interested in the strategic decision-support system (DSS) developed as part of the project. This DSS considers retrofits, investments into new equipment and decommissions of obsolete installations, for a given public building. During the course of the project, the EnRiMa DSS shall be implemented at several test sites in Austria and Spain.

As the above example demonstrates, the situation fits nicely into the multi-horizon framework. We have strategic decisions with a long time horizon (ten years or more) and want to include the option to postpone some of the decisions until later, due to the expected development of some of the involved technologies: should we install PV panels now, or wait a couple of years until their efficiency and price improve—or should we perhaps install only a few now and wait with the rest? For this, we need multiple strategic periods.

At the same time, we need operational periods in order to evaluate operational costs, efficiency, and robustness of the installed portfolio of equipment under different scenarios: with the strategic period length set to one year, we need to test the performance of the installed infrastructure in both summer and winter, using both the representative and the extreme load scenarios. All the operational scenarios are for one day, with hourly resolution. For the representative operational periods, a few profiles representing seasonal variation appear sufficient. The critical periods are modelled using CVaR values from historical data, implying that we do not require the system to handle the most extreme cases—it is most likely acceptable if temperature in an office building goes outside of the specified comfort zone on the very warmest or coldest days.

The most important stochastic parameters on the strategic level are the long-term development of electricity and gas prices, the development of price and efficiency of different technologies and, finally, regulations such as government subsidies or new electricity tariffs. For the operational periods, the most important stochastic parameters are the energy loads of the buildings. In the model, these are calculated as a function of weather (temperature, humidity, and wind), occupancy of the building, and building characteristics (Groissböck et al., 2011), where the latter might be dependent on some of the strategic decision variables. In addition, there can also be uncertainty about the electricity prices in the case of real-time pricing; this provides an example of a parameter that can be stochastic on both the strategic and the operational level and, therefore, requires the special treatment described in Section 3.2.

Model size and comparison to standard scenario trees

The model sizes depend on the number of technologies and energy types, but even a small realistic example will have about 100 binary variables in each strategic node. Furthermore, if we have 10 operational profiles with one-hour resolution for 24 hours ahead, this gives about 30000 continuous variables per strategic node. For a three-stage model that plans 10 years ahead, with 10 branches in the fifth period, the multi-horizon model will have 65 strategic nodes, which means about 6500 binary and 2 million continuous variables. A standard stochastic-

programming model, as presented in Figure 2(c), would need over 10^{11} strategic nodes to model the same situation, clearly an impossible task to handle.

4.2 Natural gas transport infrastructure planning – Ramona

The Ramona project (‘Regularity and uncertainty analysis and management for the Norwegian gas processing and transportation system’) was funded by the Norwegian Research Council and ran from 2008 till 2011. The principal objective was the development of new theory, methods, and tools to optimize regularity and capacity utilization in gas production, processing, and transportation systems. Part of the project was concerned with developing decision support system for the design of robust and flexible processing and transport infrastructure from fields (reservoirs) to markets, allowing reliable and profitable operations under various, and also adverse situations.

New investments in natural gas transport network infrastructure such as platforms, pipelines, compressors, or processing facilities should work well with existing and future infrastructure. Rather than evaluating these investment options in isolation and independently of the total system, their interactions with and effects on other infrastructure elements need to be taken into account. The timing aspect is important with respect to satisfying production obligations and developing new production fields, thus re-using infrastructure.

Increased focus on production assurance and security of supply makes it paramount to also evaluate how the design solutions will perform during daily operations and what their financial effects (costs and revenue) will be. For example, would a new pipeline allow to better satisfy delivery contracts in critical times or to route gas not bound in contracts to the most profitable markets—and how would it affect the gas flows in other pipelines? System effects in the pipeline network mean that the pressure and flow in one part of the network may influence the capacity in other parts (Midthun et al., 2009). Line-pack and other storage options require a multi-period approach to fully appreciate their value for the system (Midthun, 2007). Consequently, finding a robust and flexible design of natural gas transport networks requires, in addition to economics, also considering operational aspects such as physical processes and day-to-day gas routing decisions.

At the same time, decision makers face various kinds of uncertainty. Some uncertain parameters such as gas composition and volumes in undeveloped reservoirs, discoveries of new reservoirs, and long-term changes (trends) in price and demand levels, refer to the strategic model horizon. Other uncertain parameters may vary from day to day. Examples are prices and demands at the markets or nominations in long-term delivery contracts. Another kind of short-term uncertainty is unplanned events (e.g., network outages) that can cause problems for the security of supply in the system by drastically reducing capacity in parts of the network, if only for a short time. Considering only average values for these uncertain parameters may conceal important details. For example, delivery bottlenecks occurring during peak demand will not be visible when aggregating and using average demand levels. Brief outages of critical infrastructure may seriously affect the security of supply; using averaged values would completely disguise them.

A detailed description of the developed model is presented in Hellemo et al. (2013).

Model size

The investment analysis is typically performed over a time horizon of between twenty and fifty years while the operational analysis is carried out with daily time resolution. A typical case instance would contain about 200 network elements. For a three-stage model with 12 strategic periods and daily operational profiles with 10 branches, the multi-horizon model will then have approximately 80 million continuous variables and 20 million binary variables, already a very large model. In contrast, the corresponding model with a standard stochastic programming formulation as in Figure 2(c) would be two orders of magnitude larger with around 7 billion continuous variables and almost 2 billion binary variables.

4.3 Other applications

The presented multi-horizon structure is useful for strategic models where dealing with operational uncertainty is an important aspect for the strategic decisions and there are many potential applications.

In the energy-planning sector, an example is the design of power networks capable of dealing with fluctuating production from wind farms and other non-dispatchable energy sources—a problem that will become even more important in the coming decades, with an increasing share of renewable energy sources. It can, for example, be expected that this will cause the model to suggest the installation of short-term energy-storage solutions—while these do not have any value if we do not consider variability on the operational level. Such a model could also be extended to include other energy carriers, such as natural gas, to take advantage of the interaction between them.

For the design of supply chains, both strategic and operational uncertainty can be of significant importance, as shown in the work by Schütz et al. (2009) mentioned in the Introduction. Another example of a two-stage model with both strategic and operational nodes is presented in Pérez-Valdés et al. (2012), where the design of an industrial park is considered. Extending these models from two to multiple strategic decision points would enable the user not only to optimize the static design of the supply chain, but also to optimize the timing of the strategic design decisions. It is well known from real options theory that optimal timing is highly affected by uncertainty, for instance, through the value of postponing a decision until more information becomes available. And we are sure there are many other situations where the multi-horizon structure will be useful.

5 Conclusions

In this paper, we have discussed a multi-horizon structure for multistage stochastic programs and their associated scenario trees. The structure allows one to model and solve problems that need to combine strategic (long term) and operational (short term) uncertainty, without the explosion in the problem size that would follow from using a standard multistage model. We have discussed conditions under which the new structure is equivalent to the standard approach, and also provided guidelines for generating values for the multi-horizon scenario tree.

We have illustrated the proposed approach on a stylized optimization problem and also presented two real-world examples from the energy sector, one concerning climate control of public

buildings and the other gas pipelines and related infrastructure in the Norwegian and North Sea.

Acknowledgements

The research presented in this paper has been supported by the project ‘Energy Efficiency and Risk Management in Public Buildings’ (EnRiMa), funded by the European Commission via the 7th Framework Programme (FP7), project number 260041. Part of the presented work also builds on research performed in the Ramona project (The Research Council of Norway, project number 175967) on production assurance and security of supply for natural gas transport.

References

- D. S. Christiansen and S. W. Wallace. Option theory and modeling under uncertainty. *Annals of Operations Research*, 82:59–82, 1998.
- C. De Jonghe, B. Hobbs, and R. Belmans. Integrating short-term demand response into long-term investment planning. Cambridge working papers in economics 1132, Faculty of Economics, University of Cambridge, 2011.
- J. Dupačová, G. Consigli, and S. W. Wallace. Scenarios for multistage stochastic programs. *Annals of Operations Research*, 100:25–53, December 2000.
- S.-E. Fleten, T. Jørgensen, and S. W. Wallace. Real options and managerial flexibility. *Telettronikk*, 94(3/4):62–66, 1998.
- M. Groissböck, M. Stadler, T. Edlinger, A. Siddiqui, S. Heydari, and E. Perea. The first step for implementing a stochastic based energy management system at campus Pinkafeld. Technical Report C-2011-1, Center for Energy and innovative Technologies, Hofamt Priel, Austria, 2011.
- L. Hellemo, K. Midthun, A. Tomasgard, and A. Werner. Multi-stage stochastic programming for natural gas infrastructure design with a production perspective. In H. I. Gassmann, S. W. Wallace, and W. T. Ziemba, editors, *Stochastic Programming: Applications in Finance, Energy, Planning and Logistics*, World Scientific Series in Finance. World Scientific, 2013.
- K. Høyland and S. W. Wallace. Generating scenario trees for multistage decision problems. *Management Science*, 47(2):295–307, February 2001.
- A. J. King, S. W. Wallace, A.-G. Lium, and T. G. Crainic. *Service Network Design*, chapter 5. Springer Series in Operations Research and Financial Engineering. Springer, 2012. doi: 10.1007/978-0-387-87817-1_5.
- A.-G. Lium, T. G. Crainic, and S. W. Wallace. A study of demand stochasticity in stochastic network design. *Transportation Science*, 43(2):144–157, 2009. doi: 10.1287/trsc.1090.0265.

- K. T. Midthun. *Optimization models for liberalized natural gas markets*. PhD thesis, Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, Trondheim, Norway, 2007. URL <http://urn.kb.se/resolve?urn=urn:nbn:no:ntnu:diva-1943>.
- K. T. Midthun, M. Bjørndal, and A. Tomasgard. Modeling optimal economic dispatch and system effects in natural gas networks. *Energy Journal*, 30:155–180, 2009.
- J. Myklebust. *Techno-economic modelling of value chains based on natural gas – with consideration of CO2 emissions*. PhD thesis, Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, 2010.
- G. Pérez-Valdés, M. Kaut, V. Nørstebø, and K. Midthun. Stochastic MIP modeling of a natural gas-powered industrial park. *Energy Procedia*, 26:74–81, 2012. doi: 10.1016/j.egypro.2012.06.012. Proceedings of the 2nd Trondheim Gas Technology Conference.
- W. Römisch. Scenario reduction techniques in stochastic programming. In *Stochastic Algorithms: Foundations and Applications. Lecture Notes in Computer Science*, volume 5792, pages 1–14. Springer, 2009.
- P. Schütz, A. Tomasgard, and S. Ahmed. Supply chain design under uncertainty using sample average approximation and dual decomposition. *European Journal of Operational Research*, 199(2):409–419, 2009. doi: 10.1016/j.ejor.2008.11.040. URL <http://www.sciencedirect.com/science/article/pii/S037722170801031X>.
- K. J. Singh, A. B. Philpott, and R. K. Wood. Dantzig-Wolfe decomposition for solving multi-stage stochastic capacity-planning problems. *Operations Research*, 57(5):1271–1286, 2009. doi: 10.1287/opre.1080.0678.
- E. Sönmez, S. Kekre, A. Scheller-Wolf, and N. Secomandi. Strategic analysis of technology and capacity investments in the liquefied natural gas industry. *European Journal of Operational Research*, 226(1):100–114, 2013. doi: 10.1016/j.ejor.2012.10.042.
- B. K. Thapalia, T. G. Crainic, M. Kaut, and S. W. Wallace. Single-commodity network design with random edge capacities. *European Journal of Operational Research*, 220(2):394–403, 2012a. doi: 10.1016/j.ejor.2012.01.026.
- B. K. Thapalia, T. G. Crainic, M. Kaut, and S. W. Wallace. Single source single-commodity stochastic network design. *Computational Management Science*, 9(1):139–160, 2012b. doi: 10.1007/s10287-010-0129-0. Special issue on ‘Optimal decision making under uncertainty’.