Deciding which instruments to use to balance gas flows is not easy. *Gido Brouns* and *Alexander Boogert* discuss how to achieve gas portfolio optimisation by integrating this with the various gas transportation options available

Gas portfolio and transport optimisation

★ The transport of natural gas has received significant attention in the last year, with the large price spikes in the UK seen during sudden cold weather and the stopping of supply flow from Russia to Ukraine – and more recently Belarus. Transport is a necessity in a world where gas sources are far removed from the gas demand, and also in a marketplace in which a gas portfolio easily spans several countries.

Meanwhile, the range of options within a gas portfolio is growing, with a rising number of instruments and increasing international gas trading. This has led to a situation where decisions have become non-trivial. The objective of this article is to describe how to undertake an integrated approach to gas portfolio and transport optimisation.

In general, an energy company with a gas portfolio is faced with gas deliveries at various locations, gas consumers at other locations and a grid of pipelines connecting them. While supply and demand change over time, the energy company must balance the flows at all times. In practice, the energy company has many instruments available in order to make the flows balance – the difficulty lies in deciding which ones to use and how best to utilise these. In this article we put an emphasis on costs and consider the central question of how to balance the gas network such that the operational costs are kept as low as possible.

In the next section of this paper, we first identify the different instruments constituting a gas portfolio and transport system. Then we describe the costs associated with the utilisation of these instruments, and consider a basic optimisation model. Next we give an example of how the model works, and address the complexity of the model. Finally, we discuss how traders and other professionals can use the model in practice, and conclude with some directions for the application of this within future research.

Instruments in a gas network

To illustrate the setting of this problem, we provide a schematic overview of the northwest European gas market in figure 1. The figure shows that the gas network consists of an interacting system of source locations (where gas is supplied), demand locations (where gas is to be delivered), markets (where gas is traded), storages (where gas can be stored for future use), and liquified natual gas (LNG) (where gas is liquefied) and gas pipes connecting one location to another.

A crucial constraint in a gas portfolio is that the network should be balanced in every hour of each day of the year. There are several market instruments that can be used for balancing, for example trading, swing and storage. We will discuss these in more detail below.

Other instruments include line packing, tolerance and the use of LNG. Line packing means that the volume in a gas pipe is increased (usually within a day). A gas pipe then acts like a small storage area. Tolerance is a service that allows one to have small volumetric deviations from a pre-agreed schedule and is mainly used for real-time balancing. An upcoming instrument is LNG, which can be seen as a special type of supply. On the one hand, an LNG ship can change its course, creating an optionality in delivery place. On the other hand, LNG can be stored, creating an optionality in delivery time.

Trading

Trading is only possible at major hubs. In northwest Europe we distinguish three such hubs: NBP in the UK, Zeebrugge in Belgium, and TTF in The Netherlands. At these markets, only flat daily profiles can be attained; the hourly gas market is quite illiquid. This means that for each day, the hourly positions within that day are all equal. There are other trading locations, but those are very illiquid.

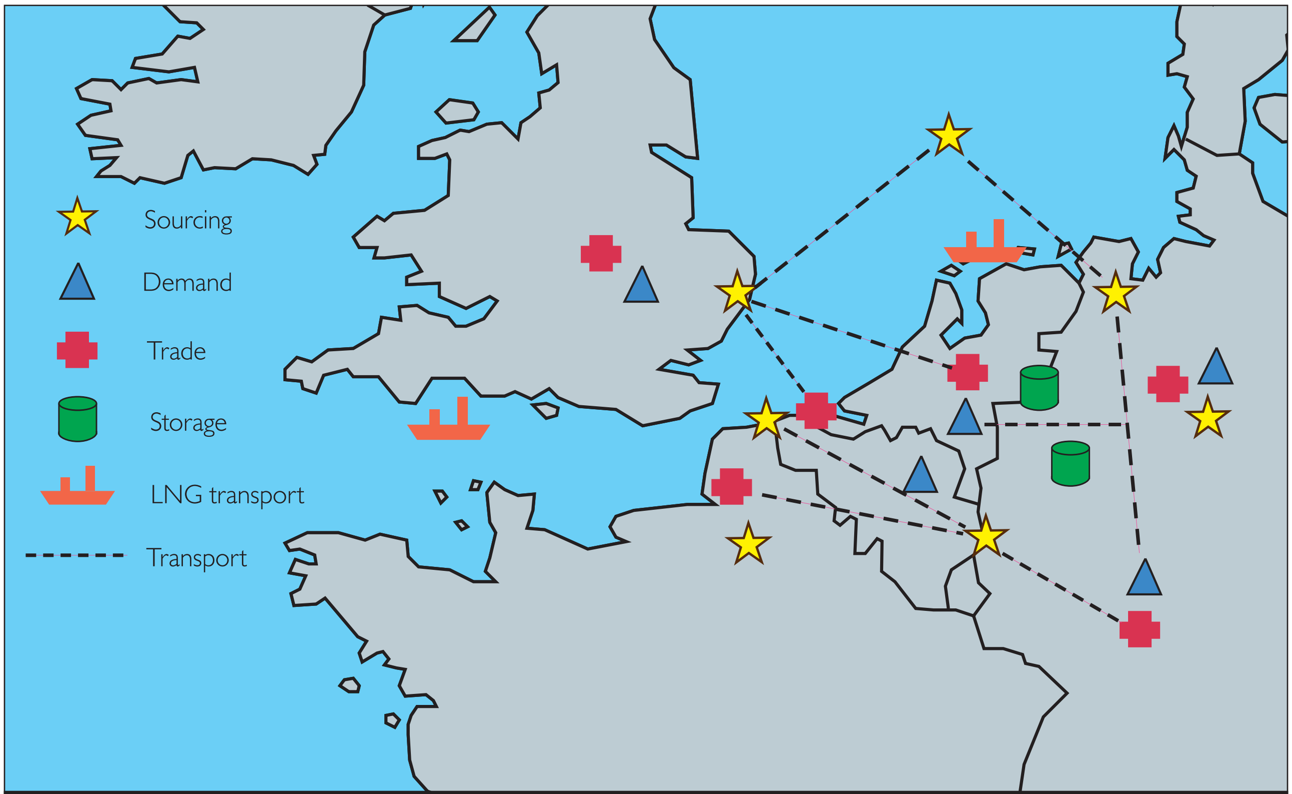

F1. Northwest European gas market

Gas can be bought or sold day-ahead, but also bought forward in weekly, monthly and yearly blocks. For our model, we will assume there exists a daily forward curve, which provides us with a spot price for each specific day in the future. For an illustration of historic spot prices at the three hubs, see figure 2. In the figure we can see there are potential arbitrage opportunities in time: buy gas now, store it, and sell it later at a higher price. In the figure we can also see potential arbitrage opportunities in location. For example, it may be profitable to buy gas at TTF and then transport it to Zeebrugge to sell it there. The depth of the market naturally restricts the extent to which a potential arbitrage possibility can be exploited. Above certain volumes the market will react unfavorably to the executed trades.

Swing

In conjunction with trading, swing contracts can be utilised. These are flexible supply contracts that can be used for daily as well as hourly balancing. Furthermore, they can be optimised dependent on price levels: if prices go up, more gas can be taken, which can be sold at the market at a higher price than the contract price.

Besides the price per unit of gas taken, swing contracts typically consist of a pre-agreed base profile with rights and limitations regarding the actual annual, daily and hourly amounts of gas taken – all of these should fall within some pre-determined bandwidth.

Storage

Another instrument that can be used for daily and hourly balancing is storage. In physical storage facilities, gas can be stored in large quantities for future use. They also serve an economic purpose: storage allows us to buy gas in the summer, store it until winter, and then sell it at a higher price, either at the market or to customers. The main advantage of storage is that it is more flexible than swing; storage limitations are commonly less rigid than swing contract limitations.

But there are also drawbacks. First, there are physical limitations, such as limited injection and withdrawal rates, which can be both time-dependent and volume-dependent. Second, storages are subject to outages and maintenance, restricting their use. Financial storage contracts contain less physical limitations.

Pipeline system

Finally, to be able to make use of the instruments described above, and hence to actually physically transport gas across the network, sufficient amounts of capacity are required for each of the pipes used for transport. Usually, the energy

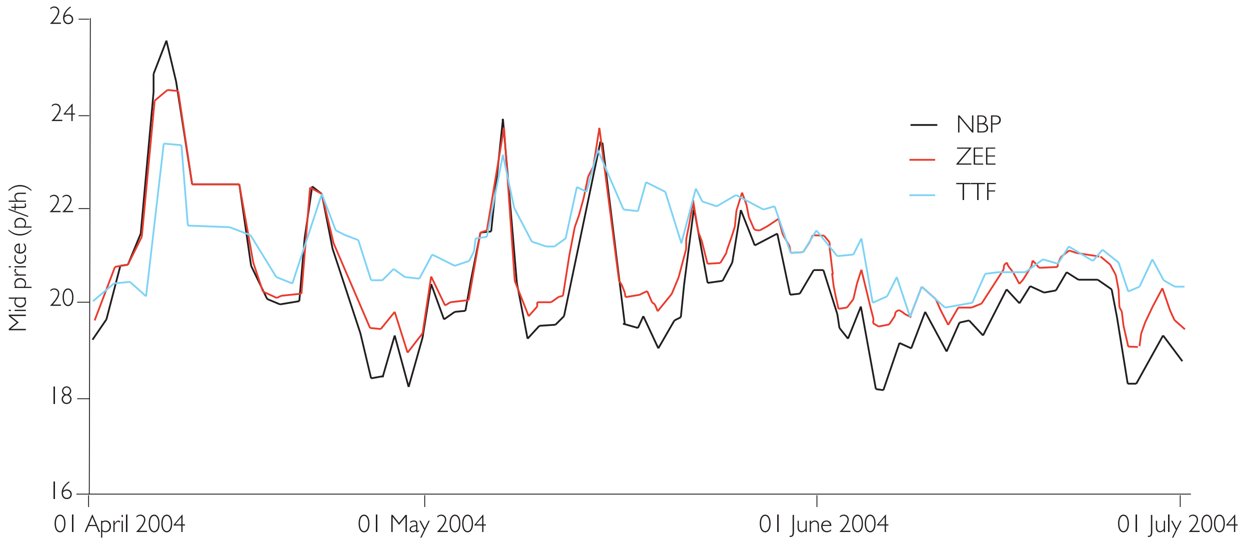

F2. Historic day-ahead spot prices

company will have already acquired capacity on several pipelines, but could purchase additional capacity for balancing or arbitrage purposes. This renders pipeline capacity a market instrument as well.

There are several problems with capacity. First, there are different transport systems: entry-exit in the UK and The Netherlands, and point-to-point in Belgium and Germany. Second, there are different price formation mechanisms: bilateral in Belgium and Germany, auction in the UK, and standard tariff according to the first-come-first-serve principle in The Netherlands. Furthermore, capacity contracts can have different durations and there are different types of gas: it can be of low and high caloric value.

Finally, there can be non-availability of capacity due to a variety of reasons such as not being able to buy more capacity on a particular pipeline or interruption of a pipeline because of an outage or scheduled maintenance.

Profits and costs

The operational costs in the network consist of factors associated with the different ways of balancing the network. The main profits and costs are related to the selling or buying of gas at one of the hubs and the costs of buying gas pipe capacity, tolerance and LNG. Besides, there are costs related to the utilisation of swing contracts and the injection and withdrawal of gas into and from storage. We note that in our basic optimisation model we will only consider existing swing and storage contracts. The decision to acquire new swing or storage contracts is beyond the scope of this paper.

Another type of costs is related to real-time imbalances. If we are not able to ship the correct amount of gas to or from each of the locations in the network in real-time, penalties will be incurred. These penalties can occur due to a mismatch in the long-term planning or due to short-term uncertainty surrounding, for example, customer demand. In our model we optimise under the constraint that for each location and for each hour, demand must equal supply (assuming this is feasible). We do not explicitly consider short-term uncertainty, although a kind of reserve margin could be included.

Model choice

Given the network structure and the different instruments available to balance the network, let us now focus on the central question of keeping the total operational costs as low as possible. It is not straightforward to decide how to set up a formal approach for this problem.

Our main idea is to formulate the problem in terms of a linear program (LP). This is a standard mathematical technique, which offers a high degree of flexibility to incorporate additional model features. It allows for a large number of decision variables, and there is a wide range of industrial solvers available to compute the optimal solution. This reduces the development and calculation time considerably. For an introduction to linear programming and related algorithms, we refer to, for example, Bazaraa *et al* (2004).

The main drawback of an LP formulation is that the optimisation is carried out based on one scenario (usually a prediction) of future prices and demand curves. In fact, all input is assumed to be deterministic and potential changes in the future are not taken into account at the time of optimisation. This means that in principle the influence of uncertainty is not captured and that robustness might be an issue.

If we were analytically and IT unconstrained, then we would opt for a dynamic programming solution in order to capture the full value of our flexible instruments. In practice, we see that solving this problem can be time-consuming and potentially even impossible due to memory constraints, and a full dynamic programming solution for the overall optimisation problem will not be achievable.

In this article we will present our LP approach, which includes storage and swing as part of the linear program.

LP optimisation

In this section we give a generic formulation of our optimisation problem in terms of an LP. An LP is composed of a cost function and a set of constraints. Both the cost function and the constraints are functions of input data and decision variables. Constraints arise because of physical or economical limitations. We will consider the optimisation problem to be a cost minimisation problem, where revenues are considered to be negative costs. Revenues are made when gas is sold at a hub.

Cost function and decision variables

Our objective is to keep the costs as low as possible. More formally, we want to minimise the total operational costs over the entire planning period by taking optimal decisions. There are essentially two types of decisions that can be made:

1. Which capacity contracts to engage in;
2. How much gas in each hour we send from one location to another.

It is easily seen that a third decision, the daily amounts of gas we buy/sell at any of the markets, is implied by the hourly flows through the network. As mentioned before, we assume there exists a daily forward curve. This curve is derived from the traded forward blocks and assumed to be known in advance.

Generic LP

One may state the LP in the following generic form:

minimise

$$\sum_{\text{gas pipes}} \text{capacity costs} + \sum_{\text{time}} \Bigg(\sum_{\text{hubs}} (\text{purchasing costs} - \text{revenues}) + \sum_{\text{swing}} \text{swing costs} + \sum_{\text{storages}} (\text{injection costs} + \text{withdrawal costs}) \Bigg)$$

such that:

★ The network is balanced in each hour;
★ Daily positions at the hubs meet the hourly flows;
★ Amounts of gas sent through a gas pipe do not exceed the contracted capacity;
★ Swing contract rules are followed;
★ Physical storage and gas pipe limitations are not violated.

Formal LP

Above we have formulated a generic LP for our problem. In this paper we refrain from a full analytical LP formulation due to the considerable space it would take. Instead we present below a formal LP formulation for a smaller problem: daily optimisation of a storage facility directly connected to a market. This is a known problem in energy to which also different solution techniques have appeared. See, for example, Boogert & De Jong (2006), which presents a dynamic programming solution. Such a solution is better able to capture the embedded optionality, but does not remain attainable for a large network, including many instruments as we discuss in this paper. Swing contracts can be treated similarly, but with different constraints.

In this case, our network consists of only two locations (a market and a storage facility), which are connected via a direct pipeline. We denote the market price at day t by $S(t)$, which is assumed to be known in advance. Let decision variable $f(t)$ denote the flow from the market through the pipe to the storage at day t. This means we buy $f(t)$ at the market and then inject it into the storage. At a later stage we can then withdraw gas from the storage and transport it to the market. Withdrawal can be seen as a negative injection, that is, in case of withdrawal $f(t)$ will take a negative value.

Over a total horizon of T days, our actions will result in a cost of

$$\textstyle\sum_{t=1}^{T} f(t)S(t)$$

The flow $f(t)$ is limited by the operational constraints from the storage facility. In this example, we illustrate two common types of operational constraints. In the first place the daily volume change $f(t)$ has to stay within the technical boundaries Δ^{min} and Δ^{max}. In the second place, the volume $\upsilon(t)$ has to stay between a minimum volume υ^{min} and a maximum volume υ^{max}. If we let $\upsilon(0)$ be the initial volume, then the volume in storage at the end of day t is

$$\upsilon(0) + \textstyle\sum_{u=1}^{t} f(u).$$

Combining the above we arrive at the following formal LP formulation:

$$\min_{f(t)} \sum_{t=1}^{T} f(t)S(t)$$

such that

$$\Delta^{min} \le f(t) \le \Delta^{max} \quad \text{for all } t = 1, \ldots, T$$

$$\upsilon^{min} \le \upsilon(0) + \sum_{u=1}^{t} f(u) \le \upsilon^{max} \quad \text{for all } t = 1, \ldots, T$$

The formal LP formulation for the full problem can be attained by expanding the above LP. One has to introduce a second decision variable for capacity, expand the network and translate the different constraints from the instruments into linear equations. Note that although we have considered a

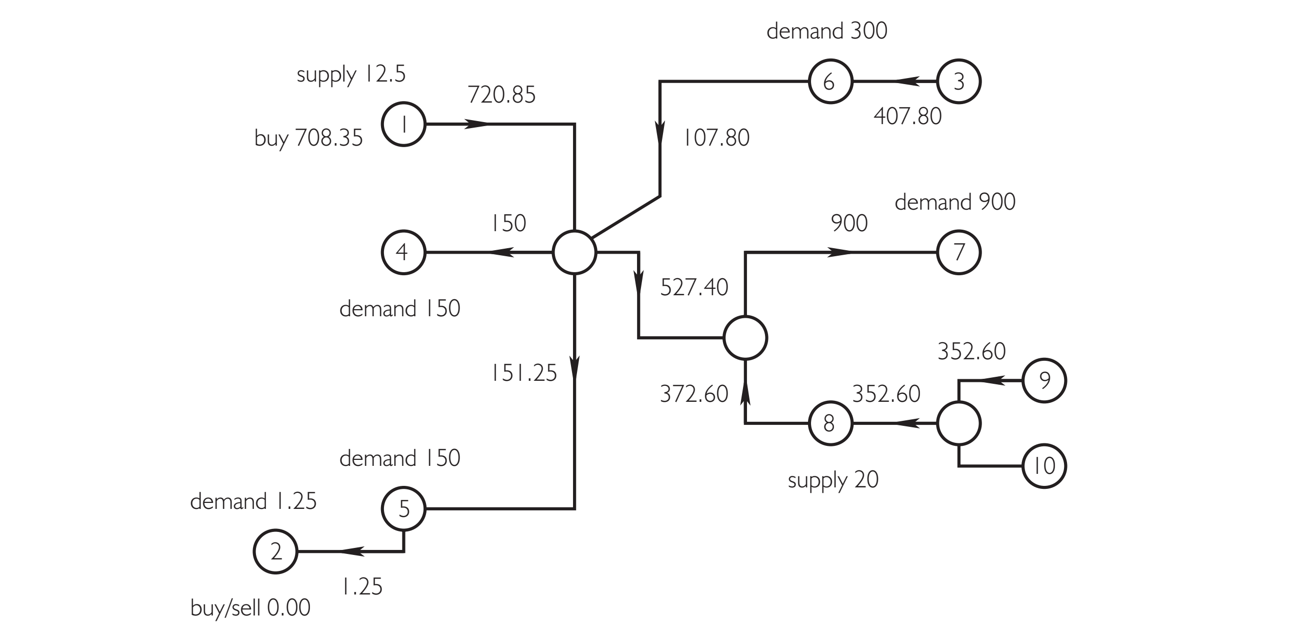

F3. Sample network

fairly straightforward example, it already requires quite a lot of equations ($4T$ to be precise) to be fully expressed.

The example above illustrates how cost elements and constraints can be incorporated in the model. Each cost element or constraint is a building block, which interacts with other building blocks, which together form the optimisation problem. Sometimes linearisation poses a problem, but this can often be avoided by introducing step functions to closely approximate non-linear expressions by linear functions.

Numerical example

To give an example of how the algorithm works in practice, consider the network shown in figure 3. We took a relatively small problem size of two hubs (locations 1 and 2), one hourly swing contract (location 3), five supply/demand locations (locations 4 to 8), and two storages (locations 9 and 10). There are also three artificial connectors, which have an hourly demand/supply of 0. The results shown correspond to the first hour of an optimisation period of five days.

We can make several observations, for example:

★ It is easily verified that the network is completely balanced in hour 1.

★ Apparently, it is optimal to buy an amount of:
$708.35 \times 24 = 17000.40$ at hub 1 on the first day.
Hub 1 has a supply of 12.5 in hour 1. So, an amount of:
$708.35 + 12.5 = 720.85$ is shipped out of the hub in hour 1.

★ To fulfill the demands in hour 1, gas is brought in from hub 1, the hourly swing contract, source location 8, and one of the two storages. The other storage is not used in hour 1.

With our model one can also run what-if scenarios. For example, suppose storage 9 is suddenly out of order. Then, we can rerun the optimisation to find an adjusted solution that takes this effect into account. The algorithm will then come up with an alternative transport scheme that satisfies all the constraints.

Complexity

An important aspect of the implementation of the LP is its complexity. Clearly, the algorithm should be able to cope with realistically sized gas networks and planning horizons. For example, for a realistic network with 80 locations and a planning horizon of a year, the number of variables in our model already reaches roughly 1.5 million, and the number of constraints reaches roughly 2 million. This illustrates that the complexity of the problem can grow very large in practice. We have carried out several tests to see how the performance of the algorithm depends on the length of the planning horizon. The tests indicate that the number of iterations needed to find the optimal solution increases linearly, whereas the solution time grows slightly more than quadratically.

Implementation of the model

Following its development, as well as a sequence of fine-tuning steps to enhance its performance, the model was implemented in an integrated decision support system to serve as its mathematical core. Based on market data that is updated daily, this system allows for daily decision-making. It supports traders by providing trading advice on physical arbitrage possibilities, given the current portfolio and transport limitations, and proposes which positions to take. Furthermore, from a risk management perspective it is important to know which contracts are needed to at least be able to serve the demand. To this end, the system shows detailed flow information and will detect any future capacity problems, such as bottlenecks in the network. The model can also be used to run what-if scenarios. For example, one can evaluate the impact of network extensions or determine the value of transmission capacity from a portfolio perspective. As such, the system can be used to confirm, from a quantitative point of view, a trader's vision of the market. All in all, it offers valuable insight into the dynamics of a gas portfolio and transport system.

Future directions

Besides the fact that the model presented in this paper has immediate practical use, it can also serve as a generic building block for even more sophisticated optimisation problems encountered in gas networks. We see various ways to move on from here.

One direction is to extend the structure of the network, for example, to include different instruments, such as LNG. Another direction is the modelling and integration of stochastic input parameters, so that we can incorporate uncertainty into our decision model. In this respect, one should distinguish between stochastic data that does not influence the feasibility of the optimisation problem and stochastic data that does.

An example of the first category is stochastic market prices. These can be incorporated by implementing several price sets each occurring with a certain preset probability. An example of the second category is stochastic demands. Incorporation of stochastic demands results in a solution that balances the gas network only in expectation, and real-time imbalances will occur inevitably. Hence, we would include imbalance costs as a part of the optimisation in addition to the other costs mentioned previously. More examples of the second category are stochastic supplies, random storage outages and random pipeline unavailability.

In any case, for all implementations we want to stress the importance of coping with complexity as the problem scale increases, since the model is clearly of little practical use if the optimisation does not run. As the network, time horizon or availability of market instruments grows, the algorithm will eventually fail, because the problem grows too large to solve or even too large to keep in memory. Consequently, intelligent modelling and fine-tuning are required to ensure that the algorithm will continue to produce good solutions in reasonable time. ER

Gido Brouns, Essent Energy Trading, Risk Management Department
Email: Gido.Brouns@essent.nl
Alexander Boogert, Birkbeck College, Commodities Finance Centre, University of London and Essent Energy Trading, Risk Management Department
Email: Alexander.Boogert@essent.nl

References

Bazaraa M, J Jarvis and H Sherali (2004)
Linear programming and network flows
Third edition, Wiley

Boogert, A and C de Jong (2006)
Gas storage valuation using a Monte Carlo method
Working paper Birkbeck College, Commodities Finance Centre, University of London