ISSN 1745-8587

Birkbeck
UNIVERSITY OF LONDON

School of Economics, Mathematics and Statistics

Birkbeck Working Papers in Economics & Finance

BWPEF 0615

Gas Portfolio and Transport Optimization

Gido AJF Brouns
Essent Energy Trading

Alexander F Boogert
Birkbeck, University London

March 2006

▪ Birkbeck, University of London ▪ Malet Street ▪ London ▪ WC1E 7HX ▪

Gas Portfolio and Transport Optimization

Gido A.J.F. Brouns*   Alexander F. Boogert*,†

March 1, 2006

1 Introduction

The transport of natural gas has received significant attention in the last months with the large price spikes in the UK facing sudden cold weather and the flow stop from Russia to Ukraine. Transport is a necessity in a world where gas sources are far removed from the gas demand, and in which a gas portfolio easily spans several countries. Meanwhile, the range of options within a gas portfolio is growing with an increasing 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 the construction of an integrated approach for gas portfolio and transport optimization.

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 the supplies and demands 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 and should decide which ones to choose and how to utilize them. 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 we first identify the different instruments constituting a gas portfolio and transport system. Then we describe the costs associated with the utilization of these instruments, and consider a basic optimization model. Next we give an example of how the model works, and address the complexity of the model. Finally we discuss how the model can be used in practice by traders and other professionals, and conclude with some directions for future research.

2 Instruments in a gas network

To illustrate our problem setting we provide in Figure 1 a schematic overview of the Northwest European gas market. 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), LNG (where gas is liquefied) and gas pipes connecting one location to another.

*Essent Energy Trading, Risk Management Dept., PO Box 689, 5201 AR 's-Hertogenbosch, The Netherlands
†Birkbeck College, Commodities Finance Centre, University of London

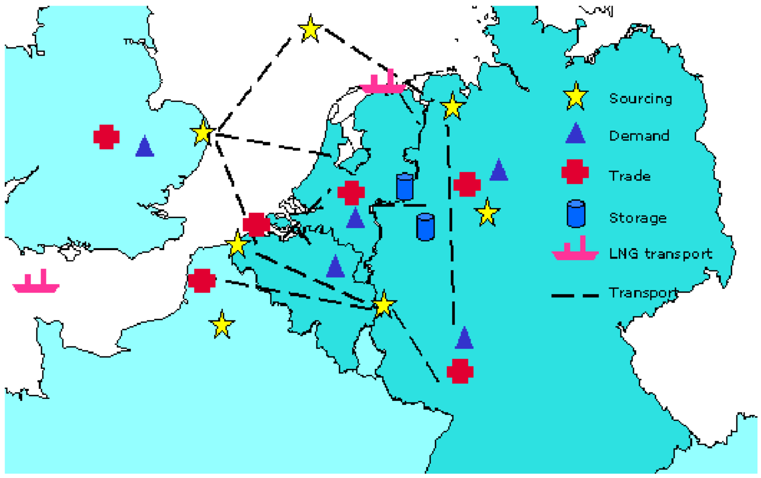

Figure 1: Northwest European gas market

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 LNG. Line packing means that, in the order of a day, the volume in a gas pipe is increased. A gas pipe then acts like a small storage. 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 of 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.

Gas can be bought or sold day-ahead, but also forward in weekly, monthly and yearly blocks. For our model we will assume there exists a daily forward curve, which provides us with a price for each specific day in the future. The price curve is typically highly volatile. This is illustrated by Figure 2, which shows the historic spot price at the three hubs during Q2 of 2004. The high volatility indicates that there are possibly arbitrage opportunities in time: buy gas now, store it, and sell it later at a higher price.

In Figure 2 we 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. In this article we focus on physical transport, i.e., we do not consider paper trading between hubs. The extent to which a potential arbitrage possibility can be exploited is naturally restricted

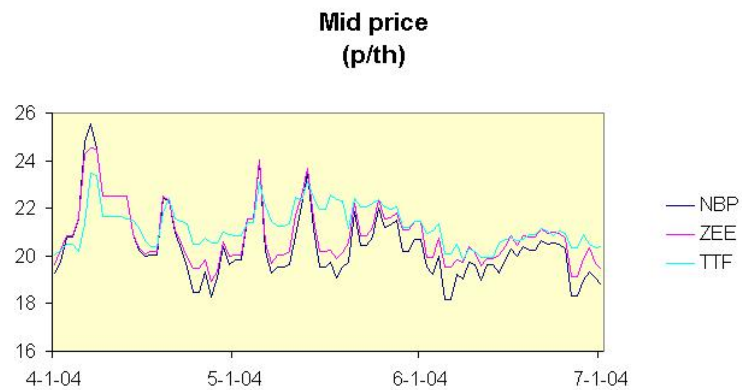

Figure 2: Historic day-ahead spot prices

by the depth of the market. Above certain volumes the market will react unfavorably to the executed trades.

Swing

In conjunction with trading, swing contracts can be purchased. These are flexible supply contracts that can be used for daily as well as hourly balancing. Furthermore, they can be optimized 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 as regards the actual annual, daily and hourly amounts of gas taken, which 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 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.

3 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 utilization of swing contracts and to injection and withdrawal of gas into and from storage.

Another type of costs is related to real-time imbalance. 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, e.g., customer demand. In our model we optimize under the constraint that for each location and for each hour demand must equal supply. We do not explicitly consider short-term uncertainty, although a kind of reserve margin could be included.

4 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, e.g., Bazaraa *et al.* [1].

The main drawback of an LP formulation is that the optimization is carried out based on one scenario (usually a prediction) of future prices and demand curves. In fact, all input is assumed to be constant and known in advance. This means that the influence of uncertainty is not captured and that robustness might be an issue. However, some degree of uncertainty can be incorporated by implementing several input sets each occurring with a certain preset probability. This approach has been shown to work in Doege *et al.* [2] in the case of a hydro storage. Furthermore, we see that solving the problem can be time-consuming and potentially even impossible due to memory constraints. However, other techniques that aim for an accurate solution would suffer from this as well.

An alternative approach would be to carry out a sequential optimization. In a first step we would optimize the individual building blocks like storage and swing. These can be described well in a stochastic framework, and allow for a careful valuation on a stand-alone basis. In the next step we would take the flows from these building blocks as given, and optimize over the remaining model parts. It is hard to estimate the accuracy of the solution obtained from this approach, compared to the optimal solution. For example, economic exercise of a swing contract might lead to a situation where we require extra transport capacity. If this capacity is highly expensive, the fancied profit on the swing contract could vaporize.

A related field is the pricing of electricity in integrated networks (see, e.g., Kristiansen [3]). This field focuses on the pricing of transmission capacity, and congestion plays an important role. The main differences are due to the storability of gas (non-storability of electricity) and the diversity of market structures (uniform in integrated electricity networks).

5 LP optimization

In this section we give a generic formulation of our optimization problem in terms of an LP. We will consider the optimization problem to be a cost minimization problem, where revenues are considered to be negative costs. Revenues are made when gas is sold at a hub.

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. To achieve a linear program, these functions should all be linear. Constraints arise because of physical or economical limitations. Together, the constraints form a polyhedron from which the values of the decision variables are to be selected. Including all instruments involves setting up many sets of constraints, each containing many individual equations. In this article we will restrict ourselves to some key ingredients to explain how the LP is set up.

Cost function and decision variables

Our objective is to keep the costs as low as possible. More formally, we want to minimize 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. in each hour, the amount of gas 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.

Generic LP

One may state the LP in the following generic form:

minimize

$$\sum_{\text{gas pipes}} \text{capacity costs} \quad + \quad \sum_{\text{time}} \Bigg(\sum_{\text{hubs}} (\text{purchasing costs} - \text{revenues}) \quad +$$

$$\sum_{\text{swing}} \text{swing costs} \quad + \quad \sum_{\text{storages}} (\text{injection costs} + \text{withdrawal costs}) \Bigg)$$

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*

The constraints are all equations made up of input data and decision variables. As an example, we will focus on the constraint that gas pipes have maximum capacities which may not be exceeded, and demonstrate how to formulate the equations corresponding to this constraint. To do this, we first need some notation.

The network consists of locations (such as hubs and demand locations) and gas pipes. Let (i,j) denote the gas pipe between locations i and j, if one exists, and let N denote the set of all gas pipes in the network. Let decision variable $f_{ij}(t)$ denote the flow through pipe $(i,j) \in N$ in hour t. For simplicity, let us assume that for each pipe there is only one type of capacity contract, which is valid for the entire duration of the planning horizon. The costs of buying capacity on (i,j) are γ_{ij} per unit per hour, and decision variable c_{ij} denotes the amount of capacity bought on (i,j). If one has a capacity of u on a certain pipe, then one is entitled to send up to u units of gas through that pipe in each hour. Besides contracted capacity also physical capacity restricts the amount of gas one may send through a pipe. Events such as scheduled maintenance cause the physical capacity to be time-dependent. Let $\phi_{ij}(t)$ denote the physical capacity of pipe (i,j) in hour t.

Now we can make the capacity costs and gas pipe constraints explicit. First, the costs of buying capacity are equal to

$$\sum_{\text{gas pipes}} \text{capacity costs} \quad = \quad \sum_{(i,j)\in N} \gamma_{ij} c_{ij}$$

Second, in each hour the flow through a pipe may not exceed the contracted capacity nor the physical capacity of that pipe, nor may it be negative, which also holds for the amounts of capacity bought, i.e.,

$$\begin{array}{rrcl}
\forall_t\ \forall_{(i,j)\in N} & f_{ij}(t) & \leq & c_{ij} \\
\forall_t\ \forall_{(i,j)\in N} & f_{ij}(t) & \leq & \phi_{ij}(t) \\
\forall_t\ \forall_{(i,j)\in N} & f_{ij}(t) & \geq & 0 \\
\forall_{(i,j)\in N} & c_{ij} & \geq & 0
\end{array}$$

Incorporating multiple capacity contracts requires some more effort, but is very well possible. In that case there is an extra summation, over all possible contracts one can engage in with respect to a certain pipe. Note that the gas pipe constraints, which are fairly straightforward, already require quite a lot of equations to be fully expressed.

The example above shows 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 optimization problem. Sometimes linearization poses a problem, but this can often be avoided by introducing step functions to closely approximate non-linear expressions by linear functions.

6 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 2 hubs (locations 1 and 2), 1 hourly swing contract (location 3), 5 supply/demand locations (locations 4 to 8), and 2 storages (locations 9 and 10). There are also 3 artificial connectors, which have an hourly demand/supply of 0.

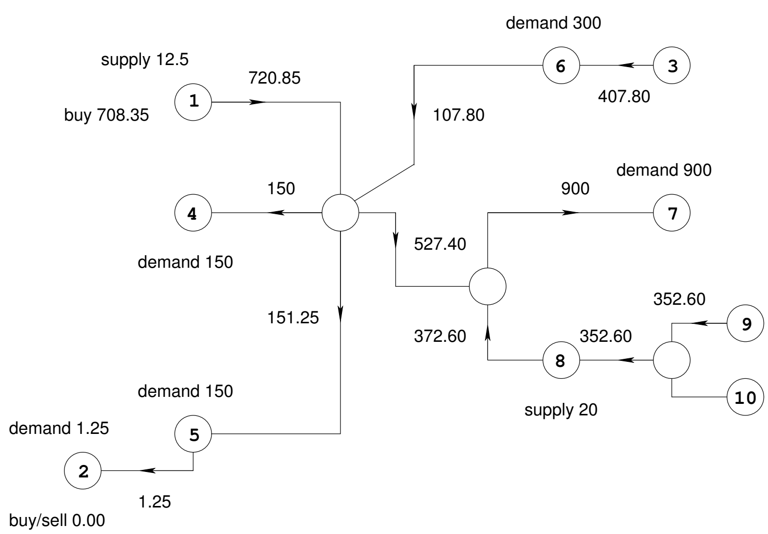

Figure 3: Sample network

The results shown correspond to the first hour of an optimization period of 5 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 17000.40 at hub 1 on the first day, or $17000.40/24 = 708.35$ per hour. Hub 1 has a supply of 300 on day 1, or $300/24 = 12.5$ per hour. 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 optimization 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.

7 Complexity

An important aspect of the implementation of the LP is the complexity of the problem. 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 problem size 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. As input we took the network of Figure 3, with random price and demand curves. All calculations were performed on a 2.8 GHz machine with 2 GB of memory, and the results are shows in Figure 4.

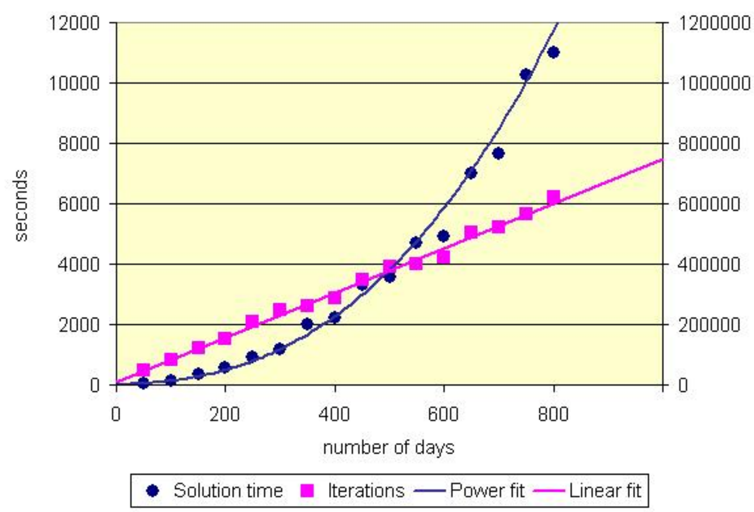

Figure 4: Performance test

The solution time, indicated by the left-hand side vertical axis, is the total time needed to read the data, generate the LP model in memory, and then solve it. We found that the number of iterations needed to find the optimal solution increases linearly, whereas the solution time grows slightly more than quadratically. If we increase the planning horizon further than 800 days, the algorithm runs into memory difficulties. This means we can optimize our network up to two years on an hourly granularity.

8 Implementation and usage 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 renewed on a daily granularity, 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.

9 Future directions

In this article we have shown a method to integrate gas portfolio and transport optimization. Besides the fact that the presented model has immediate practical use, it can also serve as a generic building block for even more sophisticated optimization problems encountered in gas networks. We see various ways to move on from here. One direction is to extend the structure of the network, e.g., to include different instruments, such as LNG. Another direction is the modelling and integration of stochastic input parameters, so that we can incorporate 'true' uncertainty into our decision model. 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 optimization 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.

References

[1] Bazaraa, M.S., J.J. Jarvis and H.D. Sherali, *Linear Programming and Network Flows*, 3rd ed., Wiley, 2004.

[2] Doege J., P. Schiltknecht and H.-J. Lüthi, *Risk management of power portfolios and valuation of flexibility*, *OR Spectrum*, 2006, to appear.

[3] Kristiansen, T., *Risk Management in Electricity Markets Emphasizing Transmission Congestion*, PhD thesis, Norwegian University of Science and Technology, Department of Electrical Power Engineering, 2004.