Implementation of a Scheduling and Pricing Model for Natural Gas

W. Pepper, B.J. Ring, E.G. Read, and S.R. Starkey

Abstract Since 1999, the Australian state of Victoria has operated a natural gas spot market to both determine daily prices for natural gas and develop an optimal schedule for the market based on an LP (Linear Programming) approximation to the underlying inter-temporal nonlinear aspects of the gas flow optimization problem. This market employs a dispatch optimization model and a related market clearing model. Here we present the model employed for both the operational scheduling and price determination. The basic dispatch optimization formulation covers the key physical relationships between pressure, flow, storage, with flow controlled by valves, and assisted by compressors, where flow and storage are measured with respect to energy rather than in terms of mass. But we also discuss a range of sophisticated mathematical techniques which have had to be employed to create a practical dispatch tool, including iterating between piecewise and successive linearization; iterating between barrier and simplex algorithms to manage numerical accuracy and solution speed issues, and special methods developed to deal with scheduling flexibility. The market clearing model is a variation on the dispatch optimization model which replaces the gas network with an infinite storage tank with unlimited transport capacity. We address the performance of the model including accuracy and run time.

The authors wish to thank the Australian Energy Market Operator (AEMO) for providing information used in this chapter and for review of the content.

W. Pepper
ICF International, Fairfax, Virginia, USA
e-mail: wpepper@icfi.com

B.J. Ring
Market Reform, Melbourne, Australia
e-mail: brendan.ring@marketreform.com

E.G. Read (✉) • S.R. Starkey
University of Canterbury, Christchurch, New Zealand
e-mail: grant.read@canterbury.ac.nz; egread@attglobal.net; stephen.starkey@pg.canterbury.ac.nz

A. Sorokin et al. (eds.), *Handbook of Networks in Power Systems II*, Energy Systems,
DOI 10.1007/978-3-642-23406-4_1,

Keywords Linear Programming (LP) • Linearization • Market • Natural gas • Optimization • Pipelines • Prices

1 Introduction

Since 1999 a gas market has been operating in the Australian state of Victoria, based on a "common carriage"[1] model, similar to that employed in many electricity markets. The motivation for introduction of that market concept has been discussed by Read et al. [1], who demonstrate that, conceptually, such a market can be based on an LP formulation analogous to that employed in electricity markets. That paper develops and interprets an LP formulation based on a textbook representation of the gas flow equations, and does not deal fully with some real-life complexities, such as the potential non-convexity of some functional relationships. Thus that formulation is not the one actually employed for dispatch purposes in the Victorian gas market. Further, while the gas market concept developed by Read et al., like the electricity market concepts on which it was based, would produce different prices for each system injection/extraction point in each trading interval, market participants in Victoria have preferred to trade on a much simpler basis, which captures some of the inter-temporal variation, but treats all gas, anywhere in the system, as interchangeable, within each day. While Read et al. [1] describes a conceptual model this paper describes the model implemented in the working market and operational dispatch models as previously presented by Pepper and Lo [20].

The Victorian market produces both the physical dispatch schedules and the market price are determined in accordance with the Market Clearing Logic described by Ruff [2] which called for optimization of the physical gas dispatch in the context of market driven supply and demand to form a realistic constrained schedule of hourly injections and withdrawals of natural gas into/from the pipeline system, alongside calculation of another hypothetical unconstrained schedule to determine the daily market prices and schedules used in the settlement of the market. These features of the Market Clearing Logic were implemented within the Market Clearing Engine (MCE) developed by ICF International [3].

The MCE can be run to produce a physical gas dispatch, the so called 'Operational Model', or to determine market prices and trading schedules, the so called 'Market Model'. Both of these models are run by the gas system operator, originally VenCorp, now absorbed into the Australian Energy Market Operator (AEMO). The Operational Model includes a detailed representation of the physical gas system and can optimize supply, based on supplier injection bids, purchaser withdrawal bids,

[1] The use of common carriage is due to the similarities with electricity markets. Victoria actually use the term "market carriage" to represent its pool based market, contrasting with "contract carriage" which relates to more traditional scheduling under bilateral contracts for access to each individual transmission pipelines.

and non-price sensitive uncontrollable withdrawal, by hour and location (node) over the gas day.

The Operational Model represents network storage (linepack) constraints, constraints on minimum and maximum allowable pressures along the pipeline system, network devices such as compressors, regulators, and check valves, as well as physical and operational constraints on supply and demand side network entry and exit points (reflecting the fact that there can be multiple participants trading through these points so that individual participants cannot manage these limits). The Operational Model produces an operational schedule which is the primary basis for scheduling gas. All gas flow and storage quantities are represented in terms of energy rather than mass, since trading mass is inadequate in a market context if different gas sources have different energy content. The basic underlying equations, as derived in the Appendix, are based on kg and kg/s, but all calculated quantities are converted to GJ and GJ/h before use in the optimization. The Operational Model also produces indicative nodal prices for natural gas, and these could be used as a basis for market trading, as discussed by Read et al. On the other hand, the Operational Model has also proved to be a highly effective dispatch tool in its own right. Thus this kind of LP formulation represents a viable approach to gas system dispatch optimization, anywhere, and irrespective of any market developments.

The Market Model is functionally very similar to the Operational Model except that the gas network is represented by an ‘infinite tank’. In effect, no storage limits are represented and negative storage inventories are even allowed during the middle of the gas day. Gas is assumed to be able to move from any point in the network, at a given point in time, to any other point in the network at any other time within the gas day. This has the effect that gas prices are the same throughout the network in all time periods. This price is used to settle the market, with constrained-on payments funded through an uplift charge on participants, being made when constraints force higher cost gas to be supplied or lower value demand to occur in the operational schedule. There are no constrained-off payments made.

From 1999 to 2007 the Operational Model was run at the start of the gas day, and as required during the gas day to provide updated schedules for the remainder of that day. The Market Model was run to produce indicative prices whenever the Operational Model was run, but settlement energy prices were based on an ex post run of the Market Model after the gas day, with demand based on the actual on-the-day conditions that had occurred. Thus there was one price for the entire gas day and this was based on actual gas flows. This concept of ex post pricing is used in numerous markets, and is described in more detail in the context of the New Zealand electricity market by.

Since 2007, the gas day has been divided into five intervals with a different price applying in each interval. But prices are not determined for all five intervals when the model is first run for each day. The first interval commences at 6 a.m. and lasts 24 h. The Operational Model and the Market Model are both run for that interval, with the Market Model’s ex ante price being used to settle the market, based on the gas scheduled for the entire 24 h (not gas as eventually flowed). The second interval runs from 10 a.m. for 20 h. Both models are run again, with a new schedule and a new ex ante price determined. This price is used to settle differences between actual gas flows and scheduled gas flows from 6 a.m. to 10 a.m., and settles all changes in scheduled gas flows for the 20 h from 10 a.m.. In effect, the 6 a.m. schedule provides a form of forward market for the 10 a.m. schedule. Similar processes are repeated at 2 p.m., 6 p.m. and 10 p.m., the last schedule covering the remaining hours from 10 p.m. to 6 a.m. the next day. The Operational Model can also be run at other times if required.

The MCE seems unique in many ways. Aspects of its function are captured in other models, however. Zhu et al. [4] note that 'the literature on gas pipeline control is rather sparse', yet they do make note of two commercial software packages for dynamic gas pipeline simulation, the Hyprotech PYPESYS and the Gregg Engineering WinTran model. Johnson et al. [5] discuss the use of WinTran model for the Tennessee Gas Pipeline which consists of over 15,000 miles of pipeline and 70 compressor stations. These models determine the (operational) gas flows, and WinTran optimizes compressor fuel use, yet they have no explicit concept of costs or economics. ICF International has two other models of gas pipeline markets that have a much different focus than the MCE. The ICF Gas Market Model [6] represents the North American Gas Market focusing on finding a market equilibrium and uses costs and economics but the daily pipeline capacities are fixed. Similarly, the ICF RIAMS model provides detailed regional modeling of pipeline capacity and determines a market equilibrium using costs and economics but again uses fixed estimates of daily pipeline capacity. The Scheduling and Pricing Engine (SPE) employed in AEMO's Short Term Trading Model (STTM) [7] for gas trading at hubs in Adelaide and Sydney (Australia) can be thought of as a form of Market Model which determines a single price and schedule for the day, but only represents a simple capacity limit on the ability of transmission pipelines to supply a hub.

Reviews of recent literature may be found in Read et al. [1] and Zheng et al. [8]. The latter note optimization models being applied to gas production, and gas pipeline network development and operations, including optimization of gas compressor operation to minimize fuel use, which may be seen as a sub-problem of the application discussed here. Both Peretti and Toth [19] and Rios-Mercado [21] use Dynamic Programming to optimize fuel use. In particular, Wu et al. [9] discuss the issues in considerable detail, although they take a different convexification approach than we do. Zheng et al also discuss formulation of the "least cost gas purchasing problem", which deals with the same physical flow requirements as our problem. De Wolf and Smeers [10] suggest a solution strategy broadly similar to our own, based on successive application and refinement of a piece-wise linearization approach. But their approach to piece-wise linearization [11] is different from ours. Several more recent papers describe practical models developed for the Norwegian Gas Industry. Somewhat like Read et al. [1], Tomasgard et al. [12] and Midthun et al. [13] both discuss linearization strategies to deal with non-linear pipe flow losses using Taylor's expansions, and note that this strategy "makes analyses computationally feasible even for large networks".

But the linearization approach we have adopted here uses the "lambda", or "convex combination" method, as discussed by Martin et al. [14]. For a recent PhD thesis on gas flow linearization techniques see Van der Hoeven [15].

The current paper differs from all of the above in that it describes a Market Clearing Engine that has been specifically implemented for the purpose of determining operational schedules, and clearing short term gas market trading on a routine basis. After describing the physical gas system involved, we present an outline of the Operational Model formulation. This model was developed on a different and more pragmatic basis than the formulation of Read et al. [1], partly because it was unclear how easy it would be to tune the more theoretical formulation of Read et al to produce accurate results in an acceptable computation time. Thus the priority for the MCE was that it had to be developed commercially in a limited timeframe to meet specific performance requirements. But the MCE, as implemented in its "Operational Model" form, is capable of producing hourly nodal prices, as discussed by Read et al. On the other hand, we later describe the "Market Model" form of the MCE, which determines the much more aggregated prices actually used for trading by solving a highly simplified special case of the Operational Model.

2 The Physical Gas Transmission System

The Victorian Gas Market is a spot market developed to facilitate the trade of gas between privately owned wholesale suppliers of gas and privately owned retailers and industrial users of gas across the Principal Transmission System (PTS), the core gas transmission network.[2] As shown in Fig. 1, this is a meshed network rather than a single long pipeline. The major market participants are both suppliers of gas and purchasers of gas so are only exposed to spot prices to the extent that their supply and usage are out of balance.

At its commencement, the market was characterized by a large demand around Melbourne supplied by a single gas production facility, Longford (represented by ESSO at the Pre-Longford node), approximately 160 km from Melbourne. The capacity of the Longford to Melbourne pipeline is small enough that the pipe's throughput could be constrained within a day. If this occurred, typically during peak demand winter days, then the only alternative supply source was from an LNG facility, with limited storage, at the Dandenong City Gate (DCG Inlet) within Melbourne. Since market commencement, the PTS has become interconnected with South Australia (SEAGas Pipeline connected at Pre-Iona), New South Wales (EAPL Pipeline connected at Pre-Culcairn) and, via the VicHub (at the Pre-Longford node), with pipelines (not shown) linking to Tasmania (Tasmanian Gas Pipeline) and New South Wales (Eastern Gas Pipeline). An underground storage

[2] As of the Victorian gas market recently coming under Australia' National Gas Rules the PTS is now called the DTS, or Determined Transmission System.

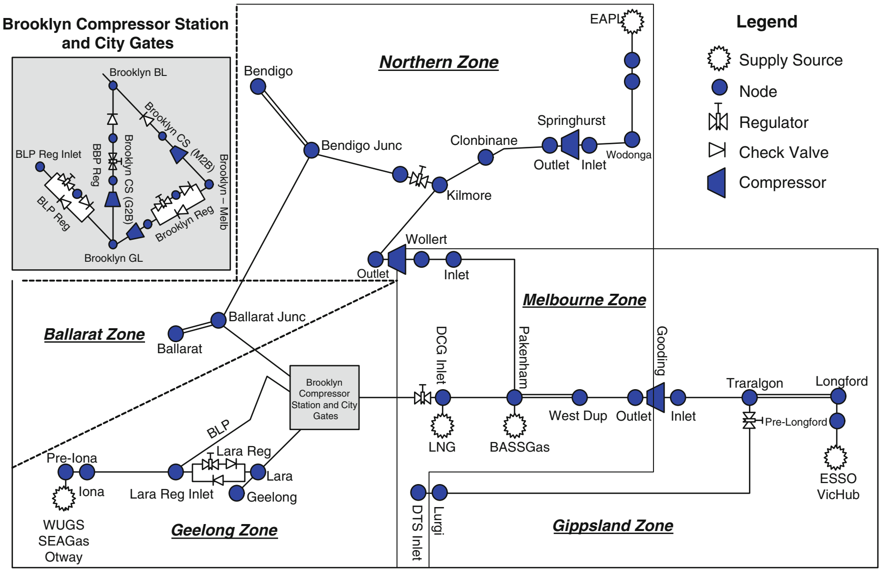

Fig. 1 Victorian gas network schematic (Source: AEMO)

field (WUGS connected at pre-Iona) which takes in gas at times of low prices and releases it at times of high price has also been established. Gas fields off the south cost of Victoria have also connected (Ottway at Pre-Iona and BassGas at Pakenham) to the PTS. The major supply node point is still Pre-Longford, though.

Major compressor stations, which effectively pump gas between regions of the PTS, are installed at Brooklyn, Wollert, Springhurst, and Gooding. These compressors are powered by a small quantity of the gas withdrawn from the network. According to Wu et al. [9] "compressor stations typically consume about 3–5% of the transported gas". In this case, though, total compressor gas usage is relatively small. For a 1,200 TJ peak demand scenario with all compressor stations running no more than 3 TJ, or 0.25% of daily demand, is consumed by compressors.

Since the commencement of the market, the demand for gas has increased, particularly in the form of exports to surrounding jurisdictions and increased use of gas fired power generation. Gas powered generators (GPGs) are located at Traralgon (Jeerlang and Valley Power GPGs), Brooklyn-Melbourne (Newport GPG), Brooklyn-Geelong (Laverton GPG) and Wollerton Outlet (AGL Somerton GPG). With all generators running, the GPG demand is about 5–10% of daily demand, and 20% of demand during the critical evening period, on a winter peak demand day.[3] AEMO is both a system and market operator. It runs the market and

[3] Based on information supplied by AEMO for winters 2007 to 2009 inclusive

has the duty to manage the system so as to maintain secure operation of the gas system. Principally this entails managing gas stored within the PTS so as to ensure that supply and demand match. The core issues which AEMO variously deal with are forecasting demand, dealing with the characteristics of various supply subsystems, managing the use of the transmission system, and determining an appropriate end of day storage pattern.

Participants in the Victorian Gas Market are typically retailers who have contracts with physical suppliers of gas. When scheduled in the Victorian Gas Market, these participants must nominate gas inflow levels from physical suppliers in accordance with their contracts. Operators of gas processing plants and interconnected pipelines typically like to keep gas flows relatively constant over the day, but may allow gas flow rates to be changed several times during the day. This means that AEMO must match relatively inflexible supply rates with diurnally varying gas demand in Melbourne. The gas day begins at 6 a.m. with supply exceeding demand, causing the pressures in the PTS to rise through the morning. During the late afternoon, demand typically outstrips supply, and remains higher than supply rates until late in the evening, causing pressures to drop. Hence AEMO must manage the system to ensure that enough gas is supplied early enough to meet demand, and to keep pressures above minimum pressure levels especially during the evening peak period.

Demand in the Victorian Gas System is quite variable. During the summer months the peak demand can be as low as 400 TJ per day, unless gas fired power generators are operating. At the peak of winter, demand can be in the region of 1,200 TJ per day. A sudden and significant weather change can produce large changes in demand. If enough gas has not been scheduled to flow to Melbourne from the various supply sources in advance of the change, then the time may not be available to schedule new supplies from producers or interconnected pipelines before minimum pipeline pressures are reached. Such 'surprise' events can force AEMO to call on LNG to meet demand. While typical gas prices are in the region of \$3/GJ – \$4/GJ, LNG may cost \$10/GJ or as much as \$800/GJ, which is the maximum market price. Hence the use of LNG imposes significant costs on the market. If LNG becomes exhausted, or cannot be vaporized fast enough,[4] then demand curtailment of industrial users must be employed. If LNG runs out then the (constrained) market may be in shortage situations on peak days, and curtailment is priced at \$800, hence the highest prices for LNG tend to approach \$800. AEMO forecasts the demand for gas (aided by knowledge of participant forecasts, for which participants are financially accountable) at the start of the day, and updates that forecast throughout the day as the situation changes.

[4] Storage of LNG is limited, and it is possible to use a significant proportion of the storage quite quickly if LNG is used too freely. Peak LNG vaporization for 1 h can take about 1 day to replace at the (emergency) maximum rate of LNG production, though more typically takes several days to replace.

Sections of the PTS can become constrained depending upon the diurnal schedule and location of injections and withdrawals. These constraints are typically pressure related and can reduce both the ability to supply consumers and the ability to take injections from suppliers. The minimum pressure of a pipeline reflects the pressure associated with minimum linepack storage. These limits cannot readily be breached – in 1998 an explosion at the Longford gas processing plant required the system to be shut down for several weeks to protect the system's minimum pressures. The maximum pressure of a pipeline is related to the physical capabilities of the pipeline, though lower limits may be set for broader system operation reasons.

Gas flows are primarily driven by the pressure differences across a pipeline, where a pipeline may be tens or hundreds of kilometers long. If gas is to flow in the desired direction, at the desired rate, then a specific pressure difference must be achieved. If demand is too great, relative to the scheduled supply, and no operator actions are taken, then consumption can literally suck gas from the pipeline at a rate which lowers the outlet pressure to the point where the outlet pressure hits its minimum level. If demand is too low, relative to the specific pressure difference, then the pressure at the demand end will rise and the pressure difference will drop. This will reduce the ability for gas to flow away from the injection point and may limit the ability to inject gas into the PTS. This can make it difficult to move enough gas through the system in time to meet increased consumption later.

AEMO use gas powered compressor stations to pump gas around the network. The PTS is divided into zones (see Fig. 1), and compressors allow gas to be transferred from low to high pressure zones, effectively making gas flow against the natural direction of flow. They can also be used to 'pump up' a zone during off-peak periods so that gas will flow from it to Melbourne during peak periods. Pipelines are often configured with other useful fittings, such as regulator valves which are designed to either restrict flow, or adjust pressure differentials. Check valves may also be used in a pipe to restrict the gas flow to a single direction.

At the commencement of the market, transmission constraints were rare, impacting only a few days per year and creating relatively small costs. The incidence and severity of constraints grew as demand grew, particularly from gas fired generation. More recently, network augmentation has increased the storage capacity of the PTS, thus increasing the ability of the system to meet demand at peak times, and significantly reduced the incidence and severity of constraints.

A major consideration for AEMO in scheduling the system is the end of day storage. There is a relatively small range of end storage levels within which AEMO seeks to finish the day. A minimum storage must be maintained for system security reasons. Too high a storage level equates to high pressures around Melbourne and this may limit the ability of gas to be supplied from locations like Longford on the next day.

In summary, then, AEMO must forecast demand and manage the scheduling of gas to manage both the daily peak demand and the end of day conditions so as to minimize the risk of surprise events requiring the use of LNG and to manage the timing of gas delivery and the positioning of stored gas to minimize transmission congestion.

3 Operational Model Formulation

Read et al. [1] develop a general formulation based on a theoretical representation of gas mass moving through adjacent pipe-line segments, described as nodes. They also discuss the treatment of practical complexities including bi-directional flows, fitting losses, compressor behavior, and flow continuity through junctions. Here we describe the actual model, which addresses these issues in a generally more pragmatic fashion.

Any operational system must balance a range of computational and practical limitations. Read et al. explicitly define the physics within pipe segments by layers of interconnected equations. This is useful in terms of providing a variety of insights, especially for engineers and pipeline operators, when looking at specific issues, such as flow losses in a given pipe segment, or compressor losses in a given zone. The challenge here, though, is not to micro-manage specific system issues, but to economically allocate the gas across the entire PTS, whilst ensuring compliance with key physical and operational pipeline and market constraints. Thus the major innovation of this formulation is to use rather longer pipe segments than those envisaged by Read et al., and to approximate the key pressure flow relationships over each pipeline segment, or system element, by fitting a convex linearization to what is essentially empirical performance data, as discussed later in Sect. 4.

A comprehensive formulation for the Operational Model is beyond the space available in this paper. However, many Operational Model constraints are unrelated to the core issues, instead implementing constraints on the operation of particular injection or withdrawal points, or market rule requirements. Here we focus on the core gas flow equations, with some nonlinear relationships depicted generically, and assuming generic forms of bid constraints, and pipeline constraints. Notationally, this formulation is rather different from that of Read et al, but the basic relationships are essentially the same, and we will focus mainly on points of real difference.[5]

This formulation represents each pipeline segment as an arc, not a node. The nodes in this formulation occur at either end of a length of pipe. This change in the construction of the formulation allows us to more readily deal with a multiple interconnected network of meshed pipes and loops. It also allows key pressure/flow relationships to be defined over the length of each pipeline arc. Flows into a node include injections (from supply bids) and flows into that node from pipe segments. The flow into the node must balance flows out of the node, into pipelines, and withdrawals from the node, as in Fig. 2.

Thus the mass conservation equation (1) is expressed in a different way from that of Read et al. Specifically, gas moves through a pipe segment from an origin node

[5] In this formulation, variables are generally upper case, and constants lower case, whereas Read et al. [1] use lower case for variables, and upper case for constants. Upper and lower limits are still represented by over and under bars, respectively. Unless otherwise stated, all variables in this formulation are positive.

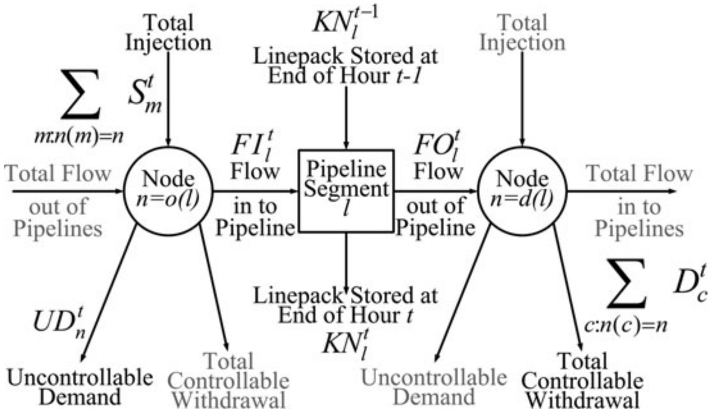

Fig. 2 Mass balance representation

of the pipe segment l, identified by $o(l)$, to the corresponding destination node $d(l)$. We denote the flow into and out of pipe segment l, in period t, as FI_l^t and FO_l^t respectively. S_m^t is gas scheduled to be supplied according to injection offer m, in period t while D_c^t is gas scheduled to be withdrawn according to withdrawal bid c. Hence the operational MCE defines the pipe network mass balance as[6]:

$$0 = \sum_{m:n(m)=n} S_m^t - \sum_{c:n(c)=n} D_c^t - UD_n^t + \sum_{l:d(l)=n} FO_l^t - \sum_{l:o(l)=n} FI_l^t + slack \quad : \forall n, t \qquad (1)$$

Here the expressions $n(m)$ and $n(c)$ respectively indicate the node associated with injection offer m and withdrawal bid c. Given each pipe segment has an origin node and a destination node, the MCE can model pipe segments as allowing flow to only move from the origin to the destination; the MCE can also allow for reversible flows by allowing the origin and destination pressures to be reversed. Although not evident in (1), the formulation also differs from that in Read et al. in that it is expressed in terms of energies not gas masses.

Extra variables, UD, are added to represent a base load of "uncontrollable" demand. This accounts for a large proportion of total market demand. Here, and elsewhere in the formulation, the "slack" term indicates high priced violation variables, which ensure that the optimization can still solve, if all constraints on the problem cannot be satisfied.[7] This allows the model to produce useful solutions

[6] *slack* is a variable, yet it is also a vector, and as such it is stated in lower case.

[7] Although this terminology is common in such models these are not really "slack" variables in the traditional sense. They do indicate how far the final solution point is from the constraint, but it lies outside the feasible region, not inside. Thus they have the form of generalized slack variables, which are positive when the constraint is slack, but negative when it is violated.

in that situation, giving the dispatch schedule that most nearly meets all requirements, and indicating the location and extent of potential violation.

Read et al. introduce specific gas flow relationships for segments within a pipeline, describing pressure/flow relationships, and also the way in which those relationships will change over time. It can be seen that, through a series of substitutions, flows along any length of pipeline will be related to the change in pressure across the length of the pipe. For the operational formulation; the flow into pipe segment l, in period t, is represented as a nonlinear function, Q_o, of the origin and destination pressures, $PB^t_{o(l)}$ and $PB^t_{d(l)}$ respectively, for that pipe segment. Thus we implicitly define the pipe inflow as:

$$FI^t_l = Q_o(PB^t_{o(l)}, PB^t_{d(l)}) \quad : \forall l, t \tag{2}$$

The flow out of pipe segment l is defined in a similar fashion, being dependent again on the origin and destination pressures. The function Q_d differs from Q_o in that it also accounts for compressor fuel usage for pipe segments that have compressors, and the net change in useable linepack in the pipeline from period *t-1* to *t*. If we let $KN^{t:t=0}_l$ be the linepack in pipe segment l measured (or extrapolated) at the start of the optimization, we define flow out of the pipe segment via:

$$FO^t_l = Q_d(PB^t_{o(l)}, PB^t_{d(l)}) - KN^t_l + KN^{t-1}_l \quad : \forall l, t \tag{3}$$

FO^t_l will tend to differ from FI^t_l because of changes in linepack stored within the pipe segment or, in the case of compressors, because of fuel used in the segment. FI^t_l and FO^t_l will be non-negative for directional pipe segments, but may be negative for reversible pipe segments. Useable linepack in pipe segment l can also be expressed as a nonlinear function of the origin and destination pressures. If a pipe segment has a pressure regulator, then the origin pressure is modified by a throttle variable, TR^t_l. Regulators are modeled such that they are situated between an origin node and the corresponding pipe segment l, and are subject to a minimum regulator outlet pressure limit, $\underline{\mathrm{op}}_l$. If the pressure at the node is $PB^t_{o(l)}$ then this is the inlet pressure to a regulator, but the outlet pressure of the regulator, and hence the inlet pressure of the attached pipeline is $PB^t_{o(l)} - TR^t_l$. With TR^t_l set to zero for pipe segments without regulators the general equation for linepack on a pipe segment is

$$KN^t_l = M(PB^t_{o(l)} - TR^t_l, PB^t_{d(l)}) \quad : \forall l, t \tag{4}$$

In addition, if "*reg*" denotes the set of pipe segments with regulators then for those pipe segments that have regulators the following limit applies.

$$\underline{\mathrm{op}}_l \leq PB^t_{o(l)} - TR^t_l + slack \quad : l \in reg \tag{5}$$

To manage linepack, we allow pipe segments to be assigned to zones, which may overlap, thus allowing us to state upper/lower linepack limits for each zone, including an aggregate system linepack limit for the whole system, Zone 0. We can add "slack" variables to allow targeted end of day minimum linepack constraints for a given zone, $\underline{m}_z$, to be violated at a high penalty cost, and also a variable ML_z that discourages, but does not prevent, end-of-day linepack exceeding targeted levels – it is non-negative and has a small penalty cost, ps, applied to it. So we have:

$$\underline{m}_z = \sum_{l \in z} KN_l^T + \sum_{t,l \in z} (FI_l^t - FO_l^t) - ML_z + slack \quad : \forall z \tag{6}$$

Read et al. [1] discuss a number of further constraints that may be imposed on pressures in pipeline segments, or at the end of segments, or on pressure differences (and hence flow velocity) across segments. Here we simply represent these constraints generically, in terms of the underlying pressure variables, as:

$$PC(PB_{o(l)}^t; PB_{d(l)}^t; TR_l^t; slack) = 0 \quad : \forall l, t \tag{7}$$

A market bidding model is overlaid on the physical gas flow representation, using bids and offers essentially as described by Read et al. We let SP_{ms} be the amount of gas scheduled from step m[8] of injection bid s, at price c_{ms}, and D_{cd} be the amount of gas scheduled from step c of withdrawal bid d, at price b_{cd}. While uncontrollable demand, UD_n^t is taken to be non-price responsive, it can be curtailed by the Operational Model at the 'value of lost load', ν, which is currently set to \$800/GJ. This defines an effective bid price for uncontrollable withdrawal. Finally $Penalty(\underline{Slack})$ is a generic function defining the penalty for using "slack" (i.e. violation) variables on each constraint. Penalties are discussed further in Sect. 9. Thus the objective function may be stated as follows:

Minimise

$$\sum_{m,s} c_{ms} S_{ms} - \sum_{c,d} b_{cd} D_{cd} - \sum_{n,t} \nu UD_n^t + \sum_z psML_z + Penalty(\underline{Slack}) \tag{8}$$

Note that while Read et al. discuss the possibility of trading end-of-day linepack, it is not explicitly represented in this objective because it was decided that such trading was not appropriate. There is, in fact, only a small range of end-of-day linepack levels that will leave the system in a state from which the next day's supply requirements can reliably be met. Thus linepack requirements are specified by the System Operator, via constraint (6).

Further, the offer/bid terms here have been expressed in a very generic way, not necessarily tied to particular hours as in Read et al. In fact, when run to determine market prices, using the simplified infinite tank assumption as in Sect. 8, the same

8 Market participants submit supply offers to the market, these 'bids' are composed of up to 10 price and quantity tranches, where we call each individual tranche a 'step'.

offers and bids are assumed to persist over the entire remaining gas day, and corresponding (pseudo-)dispatch schedules will be produced. Conversely, in reality, injections and withdrawals at particular points in the PTS may be constrained by a number of constraints not discussed by Read et al. These include physical constraints, such as maximum hourly quantities, ramp limits, and restrictions on the frequency and timing of changes to gas flow rates, and constraints implied by market rules. For example, where possible, the rules require pro-rating the schedules of tied offers or tied bids.[9]

Thus we describe supplier offers and consumer bids by their aggregated injection or withdrawal price and quantity step combinations. But these must also be related to the hourly injection/withdrawal variables employed in the gas flow modeling representation above. Rather than detail those relationships here, we will express them in a very general way, using a set of generic "Market Constraints" (*MC*), defining relationships between uncontrollable demand in each period, UD_n^t, the supply and demand scheduled to be dispatched in each period from offers and bids, S_m^t and D_c^t, and the aggregate quantities cleared from those bids and offers, S_{ms} and D_{cd}[10]:

$$MC(UD_n^t \forall n, t; S_m^t \forall m, t; D_c^t \forall c, t; D_{cd} \forall c, d; S_{ms} \forall m, s; slack) = 0 \tag{9}$$

4 Linearized Flow and Linepack Representation

The Operational Model is built around a physical gas flow model, as was the conceptual formulation developed by Read et al. [1]. But Read et al. basically only stated that the nonlinear gas flow/linepack relationships should be linearized, without specifying how that should be done in practice. Read and Whaley [16] did explore a linearization approach based on analytical differentiation of the theoretical gas flow equations, but the implemented MCE described here adopted a very different, and more pragmatic, approach. In this section we discuss linearization of the basic gas flow model without considering specialized equipment like compressors. We begin by developing the relationship between flow rate and the pressure change across a pipe segment, a relationship which is then employed within both piece-wise and successive linear representations of the gas flow problem.

The pipeline system is divided into a number of pipe segments where the physical characteristics of the pipe segments and constraints on the pipe segment are relatively uniform within the pipe segment. This includes the pipe

[9] The "slack" variables applied to such constraints, have very small penalties, so as to encourage, but not force, tie-breaking.

[10] Parameters associated with a number of these constraints limits/bounds may also be scaled in a pre-processing step so as to resolve conflicts between quantities of gas previously scheduled and quntities actually observed, for example.

segment diameter and minimum and maximum constraints on pressure in the pipe segment. The flow of gas along a pipe segment is a direct function of the average pressure of the gas in the pipe segment and the pressure decline from one end of the pipe segment to the other, and hence of the inlet and outlet pressures, as indicated by Eq. 2 and 3 for FI and FO above. But that relationship is nonlinear, and approximately related to the square root of the average pressure times the pressure decline.

The allowable pressure drop along a pipe is actually nonlinear, but our aim is to develop an LP based formulation. For a short enough pipe segments, and time intervals, an approximately linear loss function could be considered accurate enough, as suggested by Read et al. [1], and this strategy is employed by iterative pipeline simulation models such as Dorin and Toma-Leonida [17]. But our model of the Victorian Gas System use a piece-wise linearization of performance over longer pipe segment lengths, and has been calibrated to within acceptable tolerances against a more general physical gas flow model. Similar experience was recently reported by Midthun et al. [13], whose linearised approximation, a Taylor series expansion about multiple-points, was deemed acceptable for work undertaken for North Sea gas system users and its operator, Gassco.

The gas flows within the pipe segments are derived assuming isothermal conditions, uniform pipeline diameter and surface conditions, and a constant altitude (sea level). While not described in this paper, the MCE includes pre-processing steps to translate measured pressures and pipeline pressure limits at pipeline altitude to sea-level equivalents. It scales flow and linepack equations to reflect the higher density of gas at sea level, and then converts pressures derived from the optimization back to be applicable at the relevant altitude. The Appendix contains the detailed equations used to estimate flow rates and linepack, given specified values of the (sea-level) pressure at both ends of a pipe segment. The MCE went through a rigorous calibration exercise during its development, and the model results for pressures closely match measured pressures from actual gas days given the initial conditions on the pipeline system, hourly injections and withdrawals, and final linepack on a given gas day.

Pipelines, valves, and compressors control the flow of gas between injection/withdrawal nodes. Additionally, a pipe segment stores gas, so a key parameter is the linepack stored at the end of each hour and the change in linepack over the hour, which impacts the flow from the pipe segment. The gas storage in a pipe segment is primarily a function of the pressures in the pipe segment and its length and diameter (See Appendix).

For pipe segments in which the flow direction is specified, this pressure delta must be non-negative and less than a maximum pressure delta. But, for most pipe segments, the pressure decline required to produce maximum allowable flow rates is significantly less than the difference between the maximum allowed inlet pressure and the minimum allowed outlet pressure. So the MCE allows for a user specified maximum pressure decline on each pipe segment set, so that the estimated flow rate corresponding to that decline will be no less than the maximum flow rate expected at any time on the pipe segment. Limiting the allowable pressure decline in this way allows the MCE to provide a more accurate linear representation of the nonlinear pressure flow relationship, by linearizing over a narrower pressure range. For bi-directional pipe segments the same logic applies, but the roles of the origin and destination nodes are expanded to allow for the possibility that the pressure conditions will be reversed.

4.1 Piece-Wise Linearization Using Convex Combinations

An optimization based on this continuous feasible region would need to directly handle the nonlinearity of gas flow and linepack equations. To avoid this complexity, a piece-wise linearization is used. A classic piece-wise linear scheme would use separate pressure variables for each block and sum these to form the actual pressure, however the approach we used here was a convex weighting approach, similar to that subsequently described by Martin et al. [14]. To implement this method, feasible combinations of inlet pressure and pressure deltas are first chosen to form a discrete "grid", as shown for a uni-directional flow pipeline in Fig. 3. In principle this method will work for any randomly chosen set of "grid" points, with the feasible region being implicitly determined by the convex hull of those points. In practice, though, a (non-rectangular) grid is formed using "families" of points, arranged in rows. Four 'families' of points are shown. The family forming the top row (circular points) corresponds to maintaining the maximum inlet pressure across different pressure deltas, while the family at the bottom (square points) corresponds to maintaining minimum inlet pressures. The intermediate families (triangular points) are evenly distributed between them.

Typically, the range of pressure deltas is very large and the rate at which the flow rate changes declines significantly as the pressure decline increases. Thus the MCE

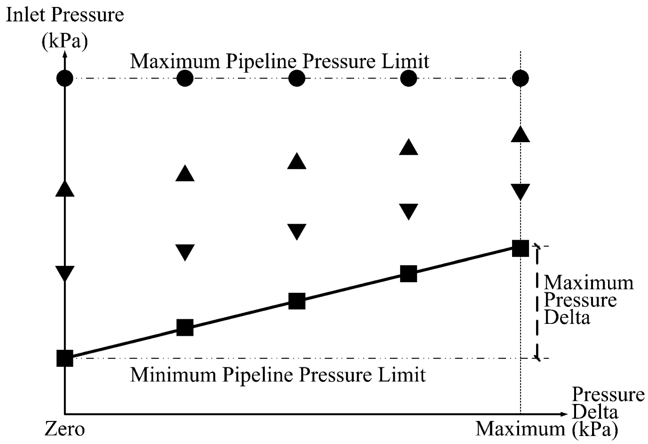

Fig. 3 Discrete representation of pressures and pressure deltas

uses 5 pressure deltas ($d = 0$ to $4 = PS$[11]) for each family specified, or 20 points in total, so that the differences in pressure deltas increase as the pressure delta increases. Obviously, it is desirable to allocate those points to the most strongly curved regions, thus giving greater accuracy in that region. This is achieved by placing point d of each row at a level $1\text{-}((PS\text{-}d)/PS)^{1.75}$ of the maximum pressure delta. For 5 pressure deltas, this implies pressure deltas at 0%, 9%, 30% and 60% and 100% of the maximum pressure delta.

For each pipeline segment, j, we have an origin pressure $P_{O(j)}$ and a pressure delta ΔP_j. Figure 3 presents a discrete grid of points, each defined by a specific value $P_{O(j)}$ and for ΔP_j. But we can represent any other point in the convex hull of these grid points as a weighted sum of grid points.[12] by weighting point $(P_{O(j)}, \Delta P_j)_i$, with a variable, α_{ji}, and require that these weights be positive, and collectively sum to 1.[13] Mathematically, any general point, $(P_{O(j)}, \Delta P_j)$, in the implicit feasible region for pipe segment j can then be expressed as:

$$P_{O(j)} = \sum_i \alpha_{ji} P_{O(j)i} \tag{10}$$

$$k \tag{11}$$

where $\alpha_{ji} \geq 0 \; \forall j, i$ and $\sum_i \alpha_{ji} = 1$

But note that other characteristics, such as the destination pressure, the steady state gas flow (FO and FI), and the linepack stored in the pipe (KN) can be determined as functions of P_O and ΔP, as described by the Appendix. Thus, while the location of each point shown in Fig. 3 is defined with respect to only two variables $(P_{j(o)}, \Delta P_j)$, it may have associated with it a vector giving the values for any number of characteristics that can be calculated as a function of those two variables. We denote these characteristics by x_j^k for $k = 1,\ldots K$. We can then create an effective piece-wise linear representation of any one of these characteristics by applying the same weights used to represent $(P_{O(j)}, \Delta P_j)$ in terms of the underlying grid point values.

$$x_j^k = \sum_i \alpha_{ji} x_{ji}^k \quad \forall j, k \tag{12}$$

Thus the generalized functions stated in Eqs. 2, 3 and 4 for flow in and out of a pipe segment, and linepack in that segment, are actually defined by essentially empirical value determined for a set of grid points. And the fundamental LP variables in this representation are not actually flows, pressures etc, but the

[11] PS indicates the number of “pressure states”, although actually there are $PS + 1$, including state 0.

[12] For a 2 dimensional representation like this, every point can actually be defined as a weighted sum of three particular grid points. But, in order to define a convex feasible region, we need to leave the model to determine which grid points it prefers to use.

[13] Together, these two constraints actually mean all weights must lie in the range [0,1].

weighting variables, α_{ji}. This allowed the model to be implemented quickly, and readily tuned to match observed empirical relationships. Subsequently the same approach was used to improve modeling of losses on the HVDC inter-island link in the New Zealand electricity market clearing engine, and to model losses on all transmission lines in the Singapore market clearing engine. While it is not immediately obvious whether this type of formulation will increase or decrease computational times, it has proven at least competitive with more traditional piece-wise linearization in the latter context. And here it has the advantage of increasing the length of pipeline segments, and hence reducing their number, relative to using explicitly defined gas flow equations of the type proposed by Read et al.[14]

4.2 Successive Linearization to Deal with Non-convexities

The piece-wise linearization approach works by creating a convex LP feasible region, and implicitly assumes that the underlying physical equation set also forms a convex feasible region. Unfortunately the real problem here is non-convex, even if the LP defines a convex feasible region. This is because the physical equations require that the solution lie on the boundary of the LP feasible region, and that boundary is itself a non-convex set. In many cases this does not matter because the objective function makes it desirable for the optimal solution to lie on the physically feasible boundary of the feasible region anyway. Thus piece-wise linearization is often applied without problems to optimization problems where the physical feasible region is actually a non-convex boundary set. That is not always the case here, though.

Figure 4 highlights the kind of issue that can be associated with the pressure delta, although there can be a similar problem with respect to the relationship between flow and inlet pressure. In the diagram, the flow should be on the top part of the curve (at the point marked Realistic) but the LP formulation allows for any flow in the shaded ‘Initial Solution Area’. If the marginal value of natural gas is less at the pipe segment outlet than at the inlet, the model will initially solve in the shaded area or on the bottom line (at the point marked Not Realistic), because it then becomes desirable to retard flow as much as possible.

Cases where the MCE will want to solve in the shaded area or on the bottom line are not uncommon. They can occur when the pipeline system will be constrained during an hour later in the day and it would be beneficial to send gas to some part of

[14] Another potential advantage of the convex combination approach to piece-wise linearization is that, if and when convexity issues arise, they could be dealt with by employing an integer type formulation, employing Special Ordered Sets (SOS2) as in Martin et al. [14], to force the model to apply weights only to adjacent grid points. A similar strategy is mentioned by DeWolf and Smeers (2000a), but dismissed on computational efficiency grounds. Still, that approach has been applied successfully in modeling HVDC link losses in the New Zealand electricity market model.

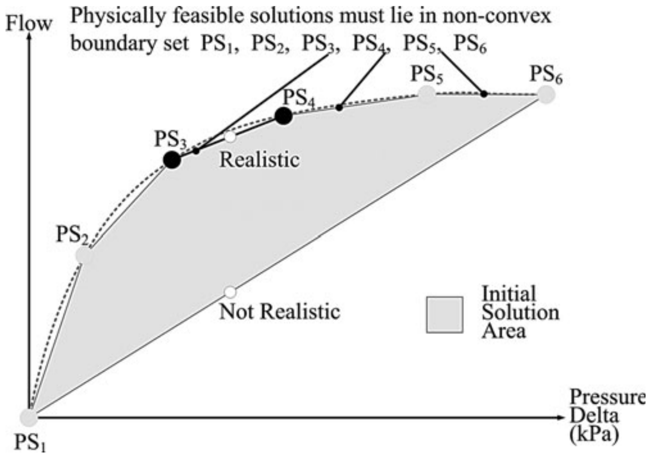

Fig. 4 Non-feasible flow solution

the system earlier in the day so that it is available as linepack during the constrained period. Or they can occur when the system is reaching minimum pressures at one point on the system but the pressure/flow equations imply that upstream gas will be diverted to another direction due to pressure differentials in that part of the system.

Cases where the MCE will want to solve in the shaded area or on the bottom line are not uncommon. They can occur when the pipeline system will be constrained during an hour later in the day and it would be beneficial to send gas to some part of the system earlier in the day so that it is available as linepack during the constrained period. Or they can occur when the system is reaching minimum pressures at one point on the system but the pressure/flow equations imply that upstream gas will be diverted to another direction due to pressure differentials in that part of the system.

There can also be an issue when negative prices arise, but this is not common. They occur when either the bid or pipeline constraints force too much gas into the system, or where the system is constrained during the day, perhaps during the evening peak. If the bid and pipeline constraints are treated as being inflexible, they can end up forcing more gas into the PTS during the night, thus producing end-of-day linepack greater than AEMO would want, and has specified. Compressor fuel use can seem desirable to the optimization during this period, simply to use up negatively priced gas, but the impact is usually limited by constraints on pressures and flows.

Whatever the cause may be, the MCE implements a successive iteration process to correct the problem whenever the LP optimization recommends solutions that are not physically feasible. At the first iteration, the problem is solved using piece-wise linearization, with no special features added. This might produce the iteration one solution shown in Fig. 4, where the modeled flow is ‘Not Realistic’, and less than

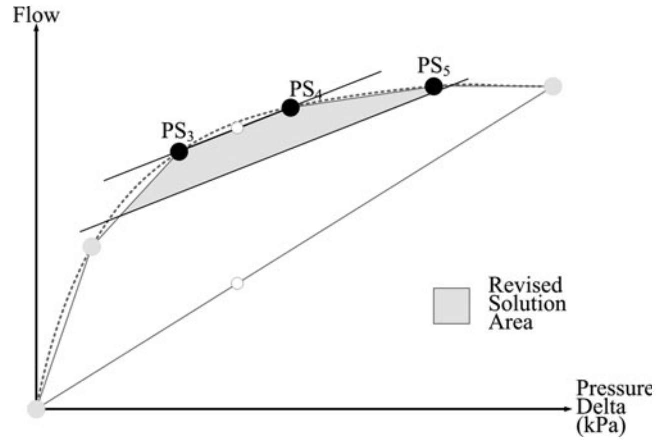

Fig. 5 Revised solution at second iteration

the flow that would actually occur, given that pressure delta. In successive iteration mode, the MCE will then do three things:

- Restrict the inlet pressures used in the model from PS_4 to PS_3 based upon the three inlet pressures that are closest to the solution
- Restrict the pressure deltas included in the model from PS_5 to PS_3 based upon the three pressure deltas closest to the solution
- Create a slope line that restricts the flow as a function of the pressure

All of these are illustrated in Fig. 5 below, where the choice of three inlet pressures and three pressure deltas allows flexibility of the model to adjust if one of the constraints (let's say pressure delta) on one pipeline segment causes the inlet pressure required on another pressure delta to adjust up or down, possibly making the constraints on that pipeline segment infeasible. Small deviations from the constraint line are allowed without penalty. But larger deviations are subject to a small cost penalty. Additional system wide constraints limit the total deviation allowed across all pipe segments. These constraints can be relaxed at subsequent iterations if found to be overly restrictive. The stopping condition is that the deviation is small enough, or that a maximum of 14 iterations are completed.[15] Also, in each successive iteration, the three inlet pressures and three pressure deltas are reviewed to see if they are still appropriate, and alternatives chosen if required.

[15] Early experience with the model showed that most cases solved within 14 iterations. Also, cases which did not solve showed negligible improvement in convergence after 14 successive approximations.

5 Modeling Compressors

Pipelines move gas from node to node, and through the process of transporting the gas there are pressure losses to the gas. Unlike gas pipeline segments, compressors move gas from a lower pressure state to a higher pressure state, while consuming gas in the process. Depending on the inlet gas pressure, a compressor can be run over a range of speeds (rpm) to achieve various combinations of gas flow and pressurization (pressure delta), and hence of outlet pressures and gas consumption. We can think of the feasible region as being defined by three primary (LP) variables: Flow, Inlet Pressure, and Outlet Pressure, or equivalently, Pressure Delta.

Figure 6 simplifies this representation by projecting the feasible region down on to only two dimensions, Flow and Pressure Delta, with Inlet Pressure and/or Outlet pressure being only implicit. This representation allows us to define the maximal feasible region over which the compressor could operate, in terms of the extreme low/delta combinations that can be achieved at various rpm, if the pressure (and hence the density) of gas flowing through the compressor were also allowed to vary all the way from its maximum (typically limited by the maximum output pressure $\bar{p}_{out}$) down to its minimum (typically limited by the minimum input pressure $\underline{p}_{in}$). Our approximation methodology requires identifying performance at a number of key points, both within the (3-dimensional) range of variations, and at its corner points. Figure 6 shows some of the key characteristics of each compressor, in terms of their projection onto this 2-dimensional representation:

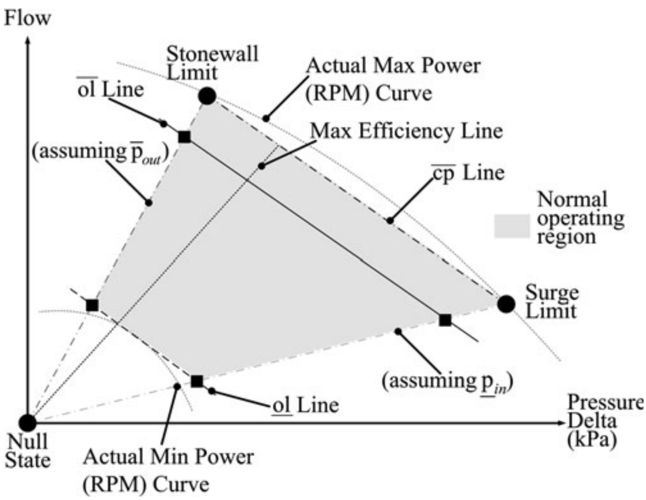

Fig. 6 Compressor operational envelope[16]

[16] In practice, it is usual for the axes to be reversed to match compressor manufactures data relating to efficiency curves, however we present them in this orientation for consistency

- Minimum and Maximum Compressor Power curves ($\underline{cp}$ and $\overline{cp}$)
- A Null State (*NS*) corresponding to no pressure delta, and no flow[17]
- A "Surge Limit" line, modeled as a ray passing through the maximum pressure delta that can be achieved at $\overline{cp}$, and $\underline{p}_{in}$
- A 'Stonewall Limit' line, modeled as a ray passing through the minimum pressure delta that can be achieved at $\overline{cp}$ and $\overline{p}_{out}$.
- A maximum efficiency line, assumed to be a ray passing through the most efficient operating point on the Maximum Compressor Power curve.
- Maximum ($\overline{ol}$) and Minimum ($\underline{ol}$) operating lines and their associated, actual, nonlinear feasible operating curves are also shown.

Since the work required from the compressor to achieve a fixed flow is a function of the compression ratio (Outlet pressure divided by inlet pressure), the maximum, minimum, and most-efficient pressure deltas are scaled to estimate the performance when solving for inlet pressures greater than the minimum inlet pressure. We model the compressor in the LP using the same convex combination approach as described in relation to Eqs. 10, 11 and 12 above. In this case, we use six operational states to represent combinations of the maximum, minimum, and most efficient pressure deltas with the minimum inlet pressure and maximum outlet pressure, all assuming the maximum compressor operating level. As above, we can require that all solution weights lie between 0 and 1, and that they must sum to 1. This would limit the solution to lie within the feasible region defined by the (3-dimensional) convex hull of this set of points at which performance parameters such as gas consumption have been assessed.

Figure 6 shows the limits of feasible compressor operation as being linearised, and this was deemed to be a good enough approximation in this case. But the upper limit could be replaced with a piecewise frontier at points selected along the nonlinear maximum power curve. Further points within the running region could also be modeled as in Martin et al. [14]. The lower limit looks more problematic, because it is non-convex, but this linearization is less restrictive than it may appear.

There are also two null states, one for minimum inlet pressure and maximum outlet pressure, both representing zero flow, zero pressure delta, and zero operating level. So, if we ignore the minimum running curve and allow non-zero weights to be placed on the null operational states, the model can choose to run the compressors at zero, if that seems optimal. But it also allows the model to choose low operating states that are not physically feasible.[18] But, as the solution is refined, and as real time approaches, we can impose conventional minimum or maximum constraints

[17] There are actually two null operational states one for the minimum inlet pressure, and the other for the maximum outlet pressure, although they appear as one point in Fig. 6

[18] In theory, operating states between zero and the minimum operating level could actually be achieved, on average, by operating efficiently for only part of the time. If so, the convex approximation to performance in that region could actually be valid. Hydro generators, for example, can be validly operated, and represented, in that way. In this case, though, the savings in operational efficiency would probably be outweighed by increased startup/shutdown costs.

on the LP variables, representing Maximum ($\overline{\text{ol}}$) and Minimum ($\underline{\text{ol}}$) compressor operating levels for any given hour. These can be used to force the compressor to be in the null operational state in some periods, or above its minimum running level in others. Or they can simply "trim" the feasible region to have linearized boundaries, as shown in Fig. 6. These constraints could have been imposed on the pressure/flow variables formed by the convex combination process, but are actually implemented by restricting the sum of the weights applied in the convex combination approximation to a narrower range than [0,1], as follows.

Consider a case where the inlet pressure is at the minimum, and the null state for the minimum inlet pressure is assigned a solution weight, $NS_{\underline{p}_{in}}$, of say 0.25. This means that the weight being placed on the points involving maximum compressor power must be 0.75, and that implies that the flow rate and pressure delta values will be both 0.75 of the values that would be produced at the maximum compressor power. Now, at a constant inlet pressure, the work required by the compressor is approximately proportional to the flow rate multiplied by the pressure delta. So the compressor power required to achieve this flow/delta combination must be approximately $0.75^2 = 0.5625$ of the maximum compressor power. This approximation also works well when weight is put on states at the maximum outlet pressure, or on any linear combination of minimum inlet pressure and maximum outlet pressure. So we can generalize this approach to apply across the full set of inlet/outlet pressures and pressure deltas, and to apply to upper and lower operating limits, both defined in terms of the power output required by the compressor.[19] Thus we have:

$$\frac{\underline{\text{ol}}}{\overline{\text{cp}}} \leq (1 - NS_{\underline{p}_{in}} - NS_{\bar{p}_{out}})^2 \leq \frac{\overline{\text{ol}}}{\overline{\text{cp}}} \tag{13}$$

Or, to make this constraint linear in the LP variables:

$$\left(\frac{\underline{\text{ol}}}{\overline{\text{cp}}}\right)^{0.5} \leq (1 - NS_{\underline{p}_{in}} - NS_{\bar{p}_{out}}) \leq \left(\frac{\overline{\text{ol}}}{\overline{\text{cp}}}\right)^{0.5} \tag{14}$$

In the LP, the compressor operating level (i.e. its power requirement and hence its gas consumption), is determined by linear interpolation between the values determined by the set of points for which it has been pre-computed. Once the LP solution is known, though, the operating level can be better approximated using the quadratic formula:

$$(1 - NS_{\underline{p}_{in}} - NS_{\bar{p}_{out}})^2 \times \overline{\text{cp}} \tag{15}$$

[19] This lower limit here defines the minimum running boundary of the feasible region if we force the compressor to be "on", thus eliminating the null states. As shown, this boundary is actually non-convex, but it can reasonably be approximated by a straight line, at the cost of eliminating a relatively small set of operating states that are seldom utilized, in practice.

6 Modelling Pressure Regulator Valves and Check Valves

The gas pipeline system in Victoria uses pressure regulation valves at a number of key points in the system, and check valves to regulate flow and optimize the daily operation of the system.

A pressure regulator controls the decline in pressure across it, while maintaining the rate of flow consistent with that for the corresponding supplying and receiving pipe segments.[20] A pressure regulator often has a maximum pressure at its outlet, which can be pre-set at the valve, and in most cases this value is also the desired pressure to be achieved. The MCE models the pressure regulator and the pipe segment at the outlet of the regulator as a single entity with both a maximum and minimum pressure at the outlet of the regulator. The pipe segment inlet pressure is allowed to be less than the inlet node pressure, but must be between the maximum and minimum outlet pressure of the regulator.

Check valves generally constrain flow across them to a single direction, their operation depends on their being a positive pressure difference between corresponding inlet and outlet pipeline segment pressures. Check valves are often installed with by-pass lines, or can be manually and temporarily by-passed, the model includes such by-passes to allow reverse flow when the pressure difference is in the opposite direction. Although such an abnormal-operation attracts a high penalty cost. The modeling of a check valve is more complex. We model this by determining the flow and pressure delta combinations that occur at each combination of maximum and minimum inlet and outlet pressures and corresponding flow, or lack of, across the check valve. This results in there being five potential limiting operational points, as shown in Table 1. The four points labeled *F* allow the MCE to cover the range of possible forward flows and pressure combinations, while the three points labeled *C* allow the MCE to represent zero flow at any combination of pressures where the inlet pressure is less than the outlet pressure (that is, where the flow is “checked”).

The convex combination approach to piece-wise linearization is then applied, as above. The problem in modeling check valves is that some linear combinations of the weights applied to pressure flow states can produce physically infeasible results,

Table 1 Check valve point modeling

		Outlet pressure			
		Zero flow		Maximum flow	
		Minimum	Maximum	Minimum	Maximum
Inlet pressure	Minimum	F,C	C	F	
	Maximum		F,C		F

[20] In other words, the pressure can decline significantly without implying the increase in flow rate one would otherwise expect for such a decline.

as for the pipe segment flow equations discussed earlier, but much more severe. We address this problem, if necessary, using the successive iteration logic discussed earlier.

7 Modeling Injection and Off-take Restrictions

Participants in the Victorian Gas Market can make injection bids to supply gas, or withdrawal bids to buy gas, or specify uncontrollable withdrawal, being gas which will be purchased at any price. Injection bids are associated with supply points, and withdrawal bids and uncontrollable withdrawal are associated with withdrawal points, represented as "nodes" in the network. AEMO can over-ride the aggregate uncontrollable withdrawal if the cumulative participant forecast differs significantly from its own. AEMO also controls how the aggregate uncontrollable withdrawal is profiled across the network and across the day. These profiles can be varied during a day as weather conditions or observed demand patterns change. In combination, the following features allow the Operational Model to more closely represent how gas is physically injected into the system over the gas day. Collectively they define the constraint set referred to as *MC* in Eq. 9.

Bids include ten price-quantity steps as well as a response time, an expiration time, and ramp limits. While participants can freely submit revised price-quantity steps between daily scheduling intervals, these other parameters must be approved by AEMO and are only changed infrequently. There may be multiple participants trading at a supply point. For instance, there are multiple participants bidding to supply gas to the market at the Longford supply point. This supply point represents a physical gas production facility. A set of constraints can be imposed on aggregate supply or demand at a point, and these may over-ride constraints on individual bids at those points (with logic included to resolve conflicts). Supply and demand point constraints can impose minimum hourly flow rates, maximum hourly flow rates, minimum daily flow rates, maximum daily flow rates, response times, expiration times, and ramp limits.

Ramp limits restrict the rate of increase and decrease in schedules between hours. A response time can be set for each hour during the gas day and indicates how long before that time a schedule must be issued for the participant to match that schedule. Thus if response times are 2 h for all hours after 10 a.m., and a schedule is issued at 6 a.m., but then revised at 10 a.m., the 10 a.m. schedule will still match the schedule issued at 6 a.m. for the first 2 h after 10 a.m., with the schedule only changing to match the 10 a.m. market-clearing dispatch schedule from noon on. The expiration time is a time during the day after which the participant will no longer respond to re-schedules. A participant with an expiration time of 8 a.m. could be scheduled for the day at 6 a.m., but would then retain that same schedule for the remainder of the day, irrespective of how prices, and other parties' dispatch schedules, may change when the market is subsequently re-cleared during the day.

The Victorian Gas Market also has a form of capacity right called Authorized Maximum Daily Quantity (AMDQ). This reflects an amount of gas that a participant can supply at one point, or receive at one point, without being deemed to have "caused" any constraints.[21] Withdrawal bids or injection bids within a participant's AMDQ limits are given priority in scheduling if tied with withdrawal bids or injection bids with the same price that are not covered by AMDQ. The market model also allows for directional flow constraints. These can be imposed across a supply point and withdrawal point so as to constrain the net flow. This can be useful in modeling gas storage fields which can only either take in gas or release gas each day. The directional flow constraint allows limits to be imposed on the net supply from a group of points. Minimum hourly net flow rates, maximum hourly net flow rates, minimum daily net flow rates, and maximum daily net flow rates can be implemented.

A problem that existed in the early history of the market was that the Operational Model was not fully reflecting how gas facility operators managed the flow of gas. They typically manage gas flow at a constant hourly rate, changing rates several times during the day. The Operational Model imposes small costs on ramping to encourage flat scheduling, but this feature failed to account for the specific times at which schedules could be changed. These could be over-ridden when strong economic incentives existed to encourage more varied scheduling. While it might be thought desirable to have more variable flow rates, the operators of gas production and supply facilities are not generally participants in the market, and ignoring their physical operating characteristics in scheduling has the potential to create costs for participants which are beyond their control (e.g. for not following their schedule). Thus the MCE could produce schedules which were not implementable in the real-world, and in extreme cases could create system security risks.

This problem has been largely resolved in recent years by allowing bids and supply and demand point constraints to have defined flexibility limits. If flow at a bid or pipeline point is specified to be completely inflexible, then the optimization will determine an hourly flow rate which is constant over the time window for which inflexibility applies. If an aggregate supply point flow is inflexible, but the injection bids there are flexible, then the total supply through that pipeline point must maintain a constant rate, but the participant schedules at that pipeline point can vary, while maintaining the same total flow.

Inflexibility restrictions are relaxed for the times at which major constraints, such as minimum and maximum hourly flow limits or directional flow constraints, change, and for a sufficient period after that, so as to allow prior schedules to be followed within the response time, and then for the dispatch schedule to be ramped to minimum or maximum hourly schedule quantities, given the applicable ramp limits. In practice, most ramp rates are quite fast, so these intervals of flexibility can be quite short.

---

[21] This being important because it determines who faces penalties and/or receives constrained-on/off payments when the actual dispatch schedule differs from the market trading schedule.

8 The Market Model Formulation

Each time it is run, the Market Model determines the single market price to be applied to trades (including deviations from trades cleared earlier) over the remainder of the gas day. It also determines the set of injection and off-take schedules that would occur if there were no limits on the ability of the gas transmission system to move gas to, or from, any node at any time within the optimization horizon. The differences between these unconstrained schedules and the schedules produced by the Operational Model runs, over the same optimization horizon, are used to determine compensation for the costs created by the constraints.

The Market Model LP is structured similarly to that for the operating schedule with a few rather significant exceptions:

- All injections and withdrawals are treated as if they were at a single node so that there are no constraints modeled on the ability to flow gas between locations. Consequently there is no representation of pipelines, compressors, check valves or regulators, although pipeline point constraints and directional flow constraints – which are applied at injection and off-take points, not pipe segments – still apply.
- With no pipeline pressure-flow relationships modeled, successive iterations is not required.
- As pipe segements are not modeled A system wide total initial linepack is used rather than pipe segment specific linepack
- The market schedule imposes no upper or lower limits on linepack levels during the day, other than an end-of-day minimum linepack constraint.
- The end-of-day minimum linepack constraint differs slightly from that used in the operational schedule formulation. Compressors are modelled in the Operational Model, and consume gas, but are not modeled in the Market Model. To ensure that each model has the same change in linepack over the day, and hence the same pricing in an unconstrained case, the compressor fuel usage determined in the Operational Model is subtracted from the minimum linepack limits used in the Operational Model but not from the minimum linpeack limits in the Market Model. As both models include cost penalties to discourage linepack above minimum, these minimum levels are effectively target end-of-day linepack levels. When the Operational Model schedules gas supplies run compressors then it must supply correspondingly less gas to achieve the end-of-day linepack target, ensuring that total supply over the day matches that in the Market Model.

9 Solution Methods

The MCE is currently executed as a set of LP problems solved with CPLEX version 9.1. Early versions of the MCE developed in the late 1990s solved the problem using the Simplex method, as this is well suited to reliably and accurately determining prices in markets. The MCE solved the problem up to 14 times in executing success iterations for the Operational Model. Even using an advanced basis for each iteration, the Simplex method was found to perform too slowly on some problems, taking several hours using the technology of the day, so a switch was made to using the barrier method.

The problem then encountered was that the barrier method had difficulty maintaining adequate numerical accuracy to cope with the range of numerical values used in a typical MCE problem. This is because the MCE employs a hierarchy of penalties; which have been adjusted over the years to minimise inappropriate interactions between them. Within this hierarchy there are many very small penalties in implementing tie-breaking, penalizing infeasible flow rates during successive iterations, and encouraging some other desirable outcomes. Since the highest market price is $800/GJ, all infeasibility penalties are set greater than this. Penalties on operator configurable constraints, which can be modified if infeasible, typically have penalties in the region of $3,000/GJ and $4,000/GJ.There are also some very large penalties to address physical infeasibilities. For example, the mass balance constraint, Eq. 8, has a violation penalty value of $9,999.9/GJ. As at 2010, the smallest penalty used is $\$2 \times 10^{-6}$/GJ while the largest used is $10,099.9/GJ, a variation of nine orders of magnitude.

To resolve the precision problem, while using the barrier method, an additional three iteration phases were introduced to both the Operational Model and the Unconstrained Market Model. Broadly, the first phase involves solving the problem with only the large penalties activated, and all small adders removed. This determines physically feasible aggregate quantities of gas to schedule over the day. The second phase fixes the aggregate schedules over the day, thus not requiring large penalties, but uses small penalties to allocate the gas flows optimally across time. Both phases 1 and 2 are performed for each successive iteration. Finally, Phase 3 is only performed at the completion of the problem – it involves effectively fixing the problem to a tight region around the optimal solution, with some minor penalty costs set to zero, and re-solving the problem using the simplex method to determine nodal prices (for the Operational Model) and the single market clearing price (for the Unconstrained Market Model).

The market timelines require that the software be able to reliably produce a schedule in the time between the window for bids and offers closing and the deadline for publishing schedules. While this time frame is about 1 h, both the Operational Model and the Unconstrained Market Model must be solved, and it may be necessary to solve multiple schedules in that time to allow the operators of the system to modify compressor commitments and other constraints under their control to correct issues seen in prior solutions. At the commencement of the market the Operational Model could be solved, including data input and output processing, in approximately 15 min using the simplex approach and earlier versions of CPLEX. The Unconstrained Market Model, which was then solved after the Operational Model only took seconds to solve.

Since 2007 the Unconstrained Market Model has been solved before running the Operational Model, providing the operators with insights about the schedule before needing to determine compressor commitments to be used in the Operational

Model. Today, using CPLEX 9.1 on faster workstations, and despite a significantly more complicated PTS than existed in the late 1990s, the Operational Model can be solved within 3 min while the Unconstrained Market Model solves well within 1 min.

10 Conclusions

The Market Clearing Engine described here has been used in the Victorian Gas Market since 1999, and demonstrates the practical use of optimization techniques to schedule dispatch, and determine prices, for a complex gas market. In principle it could be used to support market trading based on determine hourly nodal prices, as originally proposed by Read and Whaley [16], and explained by Read et al. [1]. In practice this has not eventuated and the market has evolved along somewhat different lines. This has occurred for various reasons, but it should be clear that it is not because it proved impossible to determine hourly nodal prices, since the MCE model does, in fact, determine such prices. Nor is it because the prices determined by the MCE model are always the same, everywhere in the network, as may be seen from the example discussed by Read et al. Thus while, in Victoria, the full nodal version of the MCE has proved most successful as a dispatch optimization tool, a model of this type could also be used to clear markets and support trading, and this paradigm may well prove more beneficial elsewhere, where congestion is more prevalent, and there is greater economic value at stake.

Appendix: Detailed Pressure Flow Equations for Flow Rates and Linepack

In this Appendix we present the key equations for deriving natural gas flows and pressures in the Victorian pipe network. The derivation is based on six initial equations described in Eqs. 16, 17, 18, 19, 20, 21) below.

$$P_l = \rho_l \mathrm{RTz} \tag{16}$$

The ideal gas law equation 1 describes pressure, P_l at a point l along a pipeline of length $0 < l < \mathrm{L}$. R is the ideal gas constant in units of kPa*m^3/(°K × kg), T is the temperature of the pipeline in °K, and z is the supercompressibility of gas, while ρ_l is the density in units of kg/m^3.

$$q_l = \rho_l v_l \mathrm{A} \tag{17}$$

The gas flow rate q_l at a point l, in kg/s, is described in (17). Where A is the pipe cross-sectional area in units of m^2, v_l is the velocity of gas in m/s in the pipe at point l, at the point l measured at sea level.

$$\frac{dP}{dl} \times 1000 = -\frac{f}{2\mathrm{D}} \rho_l v_l^2 \tag{18}$$

The Fanning equation 18 describes the rate of change of pressure at position l along a pipeline. D is the pipe diameter in units of m while f is the Fanning friction factor.

$$Re = \frac{\rho_l v_l \mathrm{D}}{\mu} \tag{19}$$

The Reynolds Number, Re, is defined by (19). In this equation μ is the viscosity of gas measured in kg/ms.

$$f = \frac{0.316}{\eta} \left(\frac{1}{Re} + \left(\frac{fturb}{0.316} \right)^4 \right)^{0.25} \tag{20}$$

The Fanning friction factor formulation in (20) utilizes the Blasius formulation [18], where η is the pipeline efficiency (a fraction). It describes friction as gas flows along a pipe. This makes use of a turbulent friction factor, *fturb*.

$$fturb = 0.0053 + 0.1662 \left(\frac{\mathrm{e}}{1000\mathrm{D}} \right)^{0.35} \tag{21}$$

Smooth pipelines have no turbulence while rough pipelines have more turbulence. Turbulence becomes more important as the Reynolds Number increases. Given e, a measure of the roughness of the pipeline in mm, then the relative roughness of a pipeline can be defined with respect to its diameter as (e/D). Using empirical data the form of the turbulent friction factor shown in (21) was developed by gas system engineers working on the development of the MCE. Now, using Eq. 17 to substitute for v_l in (18) we get:

$$\frac{dP}{dx} = -\left[\frac{\mathrm{f}}{\mathrm{D}} \times \frac{q_l^2}{\mathrm{A}^2 \rho_l \times 2} \right] \div 1000 \tag{22}$$

Equation 22 can be further refined by using Eq. 16:

$$\frac{dP}{dx} = -0.5\mathrm{G_f} \frac{1}{P_l} q_l^2 \tag{23}$$

Where:

$$G_f = \frac{fRTz}{1{,}000DA^2} \tag{24}$$

By integrating (23) with respect to l and observing that for $l = 0$, $P_l = P_o$, the origin pressure of the pipeline gives:

$$P_l = \left(P_0^2 - G_f q_l^2 \times l\right)^{0.5} \tag{25}$$

Using (25) to define the value of P_l where $l = L$, i.e. the destination pressure, and assuming a constant flow ($q_l = q$), friction factor, and supercompressibility, then the steady state flow rate can be derived as:

$$q = \left(\frac{P_0^2 - P_L^2}{G_f L}\right)^{0.5} \tag{26}$$

Here Eq. 26 is closely related to the Weymouth panhandle equation referred to by Zheng et al. [8] and Midthun et al. [13]. It is used later to define friction factors as a function of flow rate. However, we can derive another flow rate equation by assuming no friction arises from turbulence ($fturb = 0$), and substituting the Reynolds Number from (20) into (19) and the resultant equation for f into (22) to give:

$$P_l \times \frac{dP_l}{dl} = -0.5G \times q_l^{1.75} \tag{27}$$

Where:

$$G = \frac{0.316\mu^{0.25}RTz}{1{,}000\eta Dia^{1.25}A^{1.75}} \tag{28}$$

Integrating (27) with the assumption that the flow is constant along the pipe segment and the requirement that for $l = 0$, $P_l = P_o$ (the origin pressure) we get:

$$P_l = \left(P_0^2 - Gq^{1.75} \times l\right)^{0.5} \tag{29}$$

Using (29) to define the value of P_l where $l = L$, i.e. destination pressure, then for a non-constant friction factor and constant supercompressibility the steady state flow rate can be derived as:

$$q_l = \left(\frac{P_0^2 - P_L^2}{GL}\right)^{(1/1.75)} \tag{30}$$

Equations 6 and 30 describe flow rates under different assumptions about friction factors. These are used later to define more general flow rate equations. However, before exploring that, it is necessary to consider the linepack equations. The linepack, *I*, in a pipe segment can be calculated by integrating the volume of gas in each slice of pipeline along the length of the pipeline:

$$I = \int_0^{L} \rho_l A dl = \int_0^{L} \frac{P_l A}{RTz} dl \tag{31}$$

Assuming a constant friction factor this can be rewritten as:

$$I = \frac{A}{RTz} \times \int_0^{L} P_l dl \tag{32}$$

We also have[22]:

$$dF(l) = P_l dl \tag{33}$$

Substituting this into (32), we derive:

$$I = \frac{A}{RTz} \times \{F(L) - F(0)\} \tag{34}$$

Substituting the expression for P_l from (25), which was based on a constant fraction factor, into (33), then integrating over *l,* we get, for a non-zero flow rate:

$$dF(l) = P_l dl = \left(P_0^2 - G_f q_1^2 \times l\right)^{0.5} dl \tag{35}$$

Given the rule $(f(x))^n$ differentiated by *x* gives $n \times df/dx \times f(x)^{n-1}$ then we can integrate dF(*l*) by the reverse transformation to give:

$$F(l) = \frac{\left(P_0^2 - G_f q^2 \times l\right)^{1.5}}{-1.5 G_f q^2} \tag{36}$$

Evaluation of (36) for F(0) and for F(L), and substituting these into (34) gives:

$$I = \frac{A}{RTz} \frac{\left(P_0^3 - P_L^3\right)}{1.5 G_f q^2} \tag{37}$$

[22] Energy = Work = Force × Displacement and Pressure = Force÷Area; so combining these two results in Pressure × Length = Force × Displacement÷Area which then equals the flow of energy past a point, which is Work÷Area = Energy÷Area

Further, substituting for q from (26) gives linepack for non-zero flow of:

$$I = \frac{\mathrm{AL}}{1.5\mathrm{RTz}} \frac{\left(P_0^3 - P_\mathrm{L}^3\right)}{\left(P_0^2 - P_\mathrm{L}^2\right)} \tag{38}$$

Where the flow rate is zero, then $P_l = P_0$ for all l, and (16) implies

$$I = P_0 \frac{\mathrm{AL}}{\mathrm{RTz}} = \frac{\mathrm{AL}}{\mathrm{RT}} \frac{P_0}{\mathrm{z}} \tag{39}$$

Equations 38 and 39 can be represented in terms of an average pressure on the pipeline P_a:

$$I = \frac{\mathrm{AL}}{\mathrm{RT}} \frac{P_a}{\mathrm{z}} \tag{40}$$

Hence

$$P_a = \frac{2}{3} \frac{\left(P_0^3 - P_\mathrm{L}^3\right)}{\left(P_0^2 - P_\mathrm{L}^2\right)} \quad : \text{if} \quad P_0 > P_\mathrm{L} \tag{41}$$

Or:

$$P_a = P_0 \quad : \text{if} \quad P_0 = P_\mathrm{L} \tag{42}$$

Given user defined values of typical low and high pressures in the system, P_{low} and P_{high}, and corresponding supercompressibility values z_{low} and z_{high} it is possible to define an average supercompressability z_a as:

$$z_a(P_a) = z_{low} + \frac{(P_a - P_{low})}{(P_{high} - P_{low})} \times (z_{high} - z_{low}) \tag{43}$$

To take advantage of the linear relationship between linepack and pressure in (40), the MCE formulation uses this equation to compute linepack for all cases, including the case when the pipeline inlet and outlet pressure values are different. The term (P_0/z) is replaced by the average value of the supercompressibility-adjusted-pressures at the inlet and outlet nodes. Further adjustments are made to these equations to allow for altitude. The model also combines the results of (26) and (30) to determine flows which address the impact of a varying friction factor, non-zero *fturb*, and varying supercompressibility along the pipeline, for given pressures at the origin and destination of the pipe. All of these values are refined using an iterative approach that converges quickly and tests have demonstrated that a further iteration past the current stopping point would typically impact final flows by less than 0.2%.

References

1. Read EG, Ring BJ, Starkey SR, Pepper W (2012) An LP formulation for a natural gas market. A. Sorokin et al. (eds.), Handbook of Networks in Power Systems II, Energy Systems, Springer-Verlag Berlin Heidelberg
2. Ruff LE (1997) Victorian Gas Market Clearing Logic Version MCL 1.2. Released by Energy Projects Division, Department of Treasury and Finance, 30 October.
3. Pepper W (1999) AEMO market clearing engine: mathematical formulation. Version 1.20.4 Draft Report by ICF Consulting to VENCorp (now AEMO)
4. Zhu G, Henson MA, Megan L (2001) Dynamic modeling and linear model predictive control of gas pipeline networks. J Process Control 11(2):129–148
5. Johnson AT, Marquart BC, Istre ML, Walloopillai RK (2000) Integrating an expert system and pipeline simulation to enhance gas pipeline operation, profitability and safety. www.psig.org/papers/2000/0007.pdf. Accessed 23 Nov 2010
6. Hogan WW (1992) Contract networks for electric power transmission. J Reg Econ 4 (3):211–242. ICFI (2010) ICFI gas market model. www.icfi.com/markets/energy/doc_files/nangasweb.pdf. Accessed 28 Oct 2010
7. AEMO (2010) National gas rules, Part 20 at http://www.aemc.gov.au/Gas/National-Gas-Rules/Current-Rules.html. Accessed 2 Sep 2010
8. Zheng QP, Rebennack S, Iliadis N, Pardalos PM(2010) Optimization models in the natural gas industry. In: Rebennack IS, Pardalos PM, Pereira MVF, Iliadis NA (eds) Handbook of power systems I. Springer-Verlag Berlin Heidelberg, pp 121–148
9. Wu S, Rios-Mercado RZ, Boyd EA, Scott LR (2000) Model relaxations for the fuel cost minimization of steady-state gas pipeline networks. Math Comput Model 31:197–220
10. De Wolf D, Smeers Y (2000a) The gas transmission problem solved by an extension of the simplex algorithm. Manag Sci 46:1454–1465
11. De Wolf D, Smeers Y (2000b) The simplex algorithm extended to piecewise-linearly constrained problems. http://citeseerx.ist.psu.edu/viewdoc/summary?. doi=10.1.1.14.8090. Accessed 23 Nov 2010
12. Tomasgard A, Rømo F, Fodstad M, Midthun KT (2007) Optimization models for the natural gas value chain, In: Hasle G, Lie KA, Quak E (eds), Geometric modelling, numerical simulation, and optimization: applied mathematics at SINTEF, Springer-Verlag Berlin Heidelberg, pp 521–558
13. Midthun KT, Bjørndal M, Tomasgard A (2009) Modeling optimal economic dispatch and system effects in natural gas networks. Energy J 30(4):155–180
14. Martin A, Moller M, Moritz S (2006) Mixed integer models for the stationary case of gas network optimization. Math Program Ser B 105:563–582
15. Van der Hoeven, T. (2004), Math in gas and the art of linearization, PhD thesis International Business School and Research Center for Natural Gas, Groningen, The Netherlands
16. Read EG, Whaley R (1997) A gas market model for Victoria: dispatch/pricing formulation. Report by Putnam, Hayes & Bartlett – Asia Pacific Ltd. to VENCorp (now AEMO)
17. Dorin B, Toma-Leonida D (2008) On modelling and simulating natural gas transmission systems (part I). J Contr Eng Appl Infor 10(3):27
18. Hager W (2003) Blasius: a life in research and education. Exp Fluids 34(5):566–571
19. Peretti A, Toth P (1982) Optimization of a pipeline for the natural gas transport. Eur J Oper Res 11:247–254
20. Pepper W, Lo G (1999). Application of linear program to gas market and pipeline model in the state of Victoria, Australia. In: INFORMS conference, Cincinnati, May 1999
21. Rios-Mercado RZ (2002) Natural gas pipeline optimization. In: Pardalos PM, Resende MGC (eds) Handbook of applied optimization. Oxford University Press, New York, pp 813–826