A dynamic optimal control system for complex compressed air networks

S.W. van Heerden

24046612

Thesis submitted in fulfilment of the requirements for the degree *Philosophiae Doctor* in Computer and Electronic Engineering at the Potchefstroom Campus of the North-West University

Promoter: Dr R. Pelzer

May 2016

Abstract

Title: A dynamic optimal control system for complex compressed air networks

Author: S.W. van Heerden

Supervisor: Dr R. Pelzer

Degree: Ph.D of Engineering (Computer/Electronic)

Keywords: dynamic compressor control, dynamic compressor system, DSM, energy management, mine compressor, compressed air ring, compressed air supply, compressor control

Mines use large compressed air networks to supply shafts and processing plants with compressed air. These networks can be complex where multiple compressors are located at different locations. To add to the complexity of the network, each end user of compressed air is a separate business entity – each following its own schedule and usage requirements. Some mines have general guidelines controlling these schedules.

Most mines still use static compressor control on compressed air networks. Advanced control strategies are based on simulation results of typical usage patterns. These static controls only work when all the end users use compressed air according to the data on which the control strategy was devised. If one end user deviates from this plan, the strategy becomes non-optimal. This happens almost on a daily basis.

Previous work into dynamic control of compressed air networks was only based on basic networks where compressors were stationed close together. As soon as compressors are stationed further apart, there is a noticeable pressure drop. Due to this effect, the controller could select compressors too far away from the demand and the system would not provide a viable solution. The Dynamic Compressor Controller (DCC) discussed in this thesis solves this problem.

The DCC accomplishes this by calculating multiple compressed air set points – one for each individual compressor. These set points take the location and demand of the compressed air network into account. The operating and trimming compressor are selected dynamically. In order to reduce cycling of compressors, the future airflow is predicted to ensure sufficient compressed air supply.

The above-mentioned factors are combined to simulate the compressed air network state and propose an optimal solution for controlling the network. The solution prescribes optimal operating compressor schedules as well as pressure set points for all compressors. The prescribed pressure set point is the minimum supply pressure needed to supply the entire network with required air pressure. Due to this, the DCC will lower the running cost of the compressed air network and ensure a more stable compressed air supply by eliminating the oversupply of compressed air.

The DCC was tested at two different mines – one mining platinum and the other mining gold. Both mines have large compressed air networks. However, the operating conditions and the requirements of the mines differed.

If implemented, the DCC will be able to reduce the electricity consumption of the gold mine by up to 86 MWh per day. This can be extrapolated as a yearly reduction of R17 million in cost. The electricity consumption of the platinum mine could only be reduced by 0.5 MWh per day as it already had an optimised control schedule due to the previous implementation of a dynamic compressed air controller. This can be extrapolated as a yearly reduction of R650 000.

In South Africa, mines consume 16% of the total electricity produced by Eskom, with gold and platinum mines accounting for 80% of that. The amount of electricity consumed by compressed air generation ranges from 25% in gold mines to 40% in platinum mines. This can be extrapolated to 6% of the total electricity usage of South Africa being consumed by compressed air generation. This can further be extrapolated to stating that the DCC has the potential to reduce the total electricity consumption of South Africa by up to 1.%.

Acknowledgements

I would like to thank the Lord Jesus Christ personally for my talents and His grace, without which this thesis would never have been possible.

I would also like to thank the following persons:

- Prof. E.H. Mathews and Prof. M. Kleingeld, I would like the thank you both for the opportunity to further my studies at CRCED, Pretoria.
- My parents, Naude and Marie van Heerden. Thanks for all the support you have given me during all my previous years of study.
- My supervisor, Dr Ruaan Pelzer, thank you for all the guidance you have given me during the writing of this dissertation.
- Dr Gerhard Bolt, thank you for all the extra information and help with the writing of this dissertation.
- Dr Johan du Plessis, thank you for all the extra help with coding of this thesis as well as the writing of this thesis.
- Dr Kobus van Tonder, thank you for the help in the development of this dissertation.
- Mattheus van Niekerk, thank you for the help in the implementation of this dissertation.
- To everyone in my life not mentioned here, thank you for all the support, friendship and prayer you have given me during this time. Your efforts were not in vain.

Lastly, I want to thank TEMM International (Pty) Ltd and Enermanage (Pty) Ltd for providing funding for the research and development of this dissertation.

Table of contents

List of figures

List of tables

List of equations

Abbreviations

CO_2	Carbon dioxide
CPI	Consumer price index
CPU	Central processing unit
csv	Comma separated values
DA	Data acquisition
DCC	Dynamic Compressor Controller
DCS	Dynamic Compressor Selector
DSM	Demand-side management
EMS	Energy Management System
GUI	Graphical user interface
km	Kilometre, 1×10^3 metre
kPa	Kilopascal, 10^3 pascal
MW	Megawatt, 1×10^6 watt
MWh	Megawatt-hour, 1×10^6 kilowatt-hour
OPC	Open platform communication
R	Rand
REMS-OAN	Real-time Energy Management System – Optimised Air Networks
RISC	Reduced instruction set computer
SCADA	Supervisory control and data acquisition
SD	Secure digital
SMS	Short message service
SP	Set point
USB	Universal serial bus
VFD	Variable frequency drive
VSD	Variable speed drive
µm	Micrometre, 1×10^{-6} metre

Nomenclature

Baseload compressor	Compressor that continuously supplies compressed air during the day.
Bernoulli’s principle	Fluid dynamics principle.
Bit	Unit of information – can only be 0 or 1.
Black box system	System featuring only inputs and outputs with no indication to its inner workings.
Centrifugal compressor	Type of compressor, uses centrifugal forces.
Compressed air ring	Interconnected network of multiple compressors and demand points.
Compressor cycling	Process of starting and stopping compressors continuously.
Compressor house	Building that houses one or more compressors.
Decision tree	Tree-like graph or model representing decisions and consequences.
Dynamic compressor controller (DCC)	Compressor controller that automatically changes schedule due to changes in the compressed air network.
Dynamic compressor selector (DCS)	Dynamic compressor controller developed by Van Heerden [1]
Evolutionary prototyping	Method of software development, see Appendix A.1 – Software prototyping.
Flownex	Air simulation software.
Fossil fuel	Natural fuel formed via decaying organic material over many years.
Genetic algorithm	Search heuristic; see Appendix A.2 – Genetic algorithms.
Greenhouse gas	Gas that absorbs and emits radiation in the thermal infrared range.
Header pressure	Pressure at the head or output pressure of the compressor house.
Hill climbing	Mathematical optimisation technique; see Appendix A.3 – Hill climbing.
Kelvin (K)	SI unit for temperature.
Kilogram (kg)	SI unit for weight.
Kilowatt-hour (kWh)	Unit of energy, equivalent to 3.6×10^{6} joule.
K-loss factor	Factor used to represent resistance to flow in compressed air pipelines caused by bends, narrowing etc.

KYPipe2014	Air simulation software.
Load compressor	Process of connecting a running compressor to the compressed air network to supply compressed air to the network.
Load shifting	Type of demand-side management project.
Localised compressors	Compressors located near the demand for compressed air.
Metre (m)	SI unit for length.
Pascal (Pa)	SI unit for pressure.
Peak clipping	Type of DSM project.
Rand (R)	Currency of South Africa, ZAR.
Roughness factor	Factor used to specify the roughness of the inside of a pipe used to transfer compressed air, specified in µm.
Second (s)	SI unit for time.
Static compressor controller	Compressor controller that controls according to a fixed schedule; requires human intervention to change schedule.
Trimming compressor	Compressor that is used to control the exact supplied compressed air; usually does not run at full capacity.
Turbomachine	Machine that transfers energy between a fluid and a rotor.
Unload compressor	Process to disconnect compressor from the network while still keeping the compressor running.
Watt (W)	SI unit for power.

1. Introduction

1.1. Background

1.1.1. Energy reduction

Globally, there is a movement to reduce energy consumption. This movement is driven by a few factors: reduction in greenhouse gases, the decrease of fossil fuel reserve levels, rising electricity demand and rising electricity prices. The extent to which each factor influences the movement to reduce electricity differs from country to country.

The Kyoto Protocol [2] and the Doha Amendment [3] are international treaties that were signed to reduce greenhouses gases. These two treaties established legal binding limits for governments regulating the reduction of greenhouse gases and carbon dioxide (CO_2) emissions [4]. The Kyoto Protocol was signed in December 1997 while the Doha Amendment was signed in 2012.

Greenhouse gases occur naturally but in recent years industries have increased the levels of greenhouse gases in the atmosphere [5]. Greenhouse gases trap heat in the atmosphere by absorbing longwave radiation [6]. The higher the concentration of greenhouse gases in the atmosphere, the more heat is absorbed and the higher the temperature of earth becomes.

According to Vitousek [7] and Hughes [8], the consequences of this are numerous and atrocious. CO2 is one of the main greenhouse gases [9]. According to Spalding-Fecher and Matibe, South Africa generated 90% of its electricity from coal power plants at the end of 1999 [10]. By burning fossil fuels such as coal, CO_2 is released into the atmosphere. By reducing energy consumption, energy supply can be reduced and this can be seen as an environmental advantage.

According to Eisenberg and Nocera, roughly 80% of the total energy generation in 1998 was from fossil fuels [11]. According to Shafiee and Topia this will increase to 84% in 2030 [12]. The three main fossil fuels are oil, gas and coal. These are estimated by Shafiee and Topia to be depleted after 35, 37 and 107 years respectively [12]. Coal is currently used in limited production of synthetic fuel [13], but it emits more CO_2 than non-synthetic fuel.

According to Winkler , South Africa currently generates 93% of its electricity using fossil fuel power plants, which includes our natural gas turbine power plants [14]. South Africa is building two new coal-fired plants, Kusile and Medupi, which will increase the percentage of electricity generated from coal power plants [15]. Although coal is estimated to last for a few more years, it is still a finite resource [12].

By reducing the demand for electricity, finite energy resources such as coal can last longer. South Africa has a very high CO_2 produced per capita as can be seen from Figure 1, with about double the world average [16]. South Africa is also the world's sixth-largest consumer of coal [17].

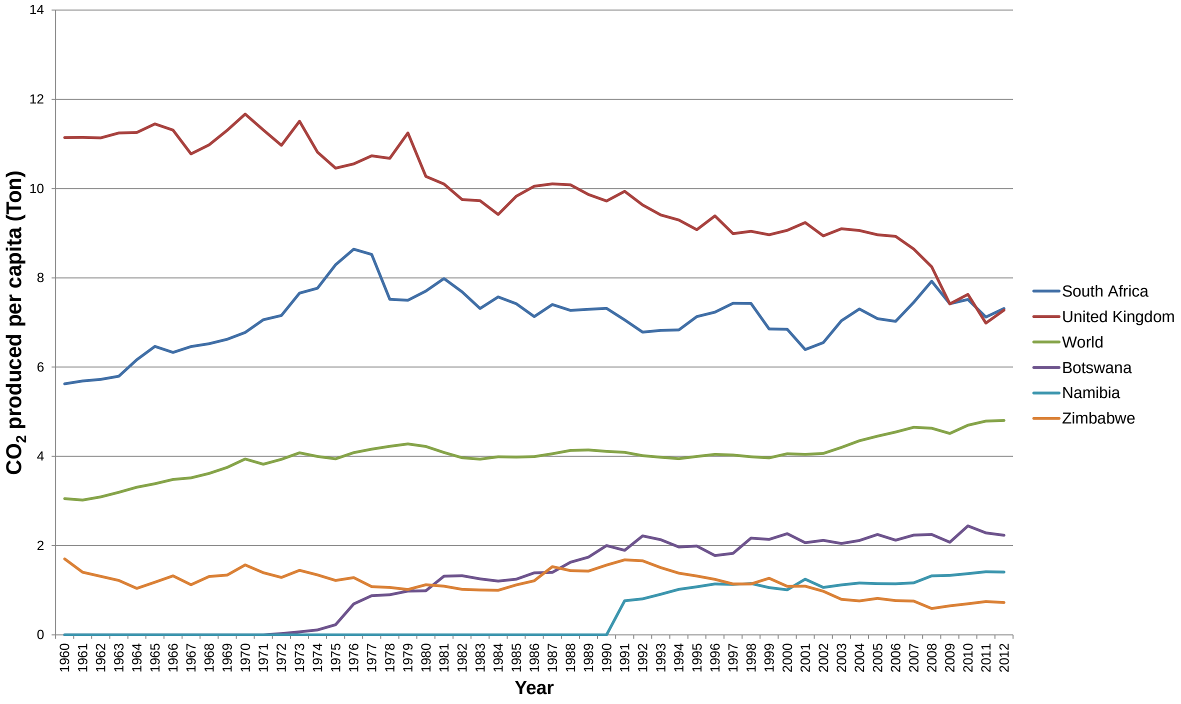

Figure 1: CO_2 per capita

In 1999, only a third of South Africa had electricity [18]. Since 1994, the government has been trying to ensure that everybody in the country had access to electricity [19]. This led to an aggregate electricity demand increase of 4% per annum [19]. Since 2007, this has resulted in load shedding occurring periodically due to the lack of reserve capacity [20].

South Africa's energy demand in the early 21th century was mostly driven by the mining industry [21], but since then other heavy industries have also increased the demand for electricity. Heavy industry in South Africa currently consumes about 45% of the total supply [21]. Between 1993 and 2006, the mining sector was the third-highest energy consumer in South Africa [22].

As expected, electricity prices increase over time due to inflation. South Africa's electricity price has always been relatively inexpensive compared with the rest of the world[23]. This has led to undesirable usage [24]. Over the last few years, the annual electricity price increases in South Africa have surpassed inflation by a large margin. High energy costs have always been offset by high production. But, the cost of electricity has increased by such a large margin that production cannot always be the first priority anymore.

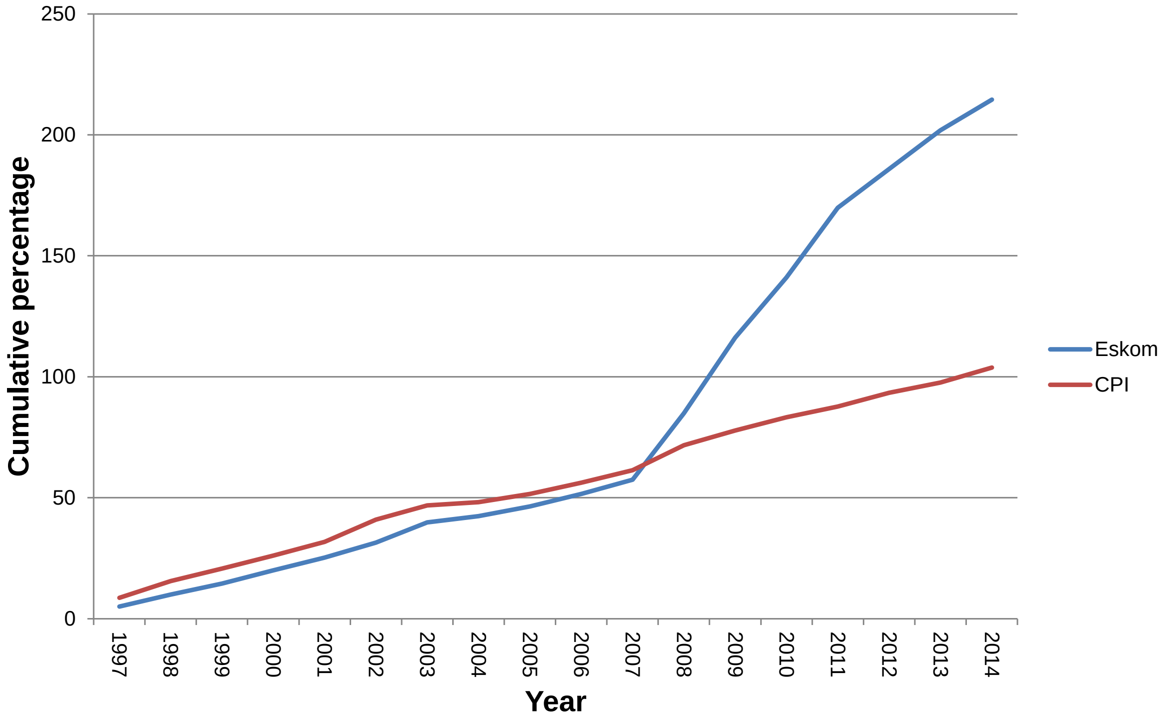

Figure 2: Cumulative Eskom price increase versus cumulative CPI

From Figure 2, it can be seen that the average price of electricity has increased by more than 100% over the consumer price index (CPI) during the period from 1997 to 2014 [25]. The South African economy relies on energy intensive industries [26]. Hence, the rising cost of electricity will have a negative impact on the cost of business for energy intensive industries.

The global movement for the reduction of energy consumption not only makes sense from an environmental perspective but also from an economic perspective. If it is possible to reduce electricity consumption without affecting production negatively, a business will be able to increase its profit.

1.1.2. Compressed air electricity consumption

South Africa has a very large mining sector, which forms an essential part of the economy [27]. Mining is a relatively energy intensive industry, especially from an electricity consumption point of view [28]. In 2011, mining consumed 16% of South Africa’s total electricity [29]. With this in mind, it is clear that the mining industry is an obvious target for electricity consumption reduction. Gold and platinum mines accounted for 80% of total energy use among mines [30], [31].

According to Fraser, compressed air generation on mines consumes from 25% for gold mines to 40% for platinum mines [32]. This is approximately 9% of the total electricity consumption of South Africa’s industrial sector [33]. This means that compressed air generation on mines can also be seen as a significant target for electricity consumption reduction and should, therefore, be prioritised.

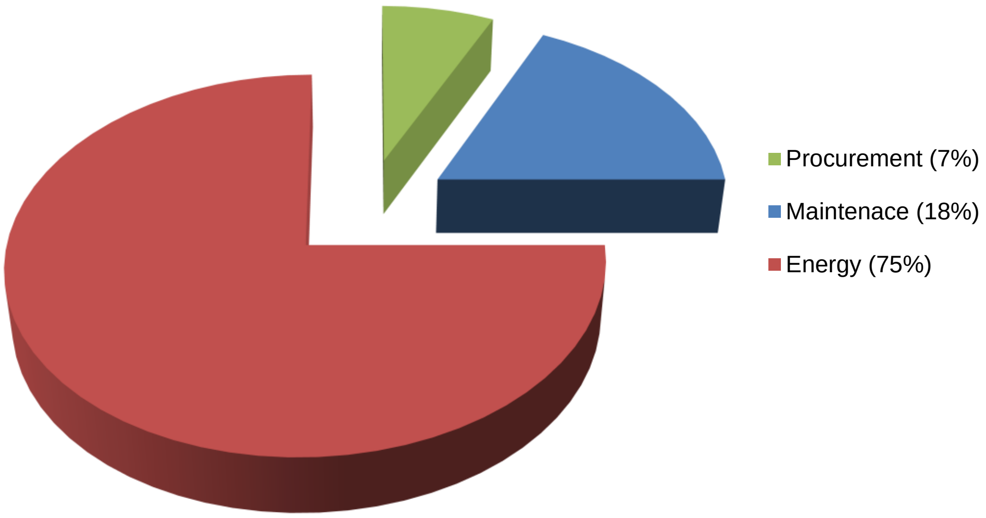

Figure 3: Compressor lifetime cost comparison

Air compressors used at mines are large (1 MW up to 15 MW) and can be very expensive to install and operate. Over and above the large investment cost, the total lifetime cost of a compressor over 10 years will be made up of approximately 75% electricity cost [34]. This is shown in Figure 3 [34]. It can thus be stated that operating a compressor for 10 years costs more than 10 times the initial procurement cost. This can be explained by the cost of electricity and the size of the compressors.

1.2. Mining compressed air systems

Compressed air is typically supplied by local compressors. These compressors are located close the demand, which simplifies the supply-demand matching of compressed air. This is typically how most factories and plants use compressed air. Most South African gold and platinum mines supply compressed air to shafts through large compressed air networks called compressed air rings.

These compressed air rings supply the shafts and processing plants of mines. Compressed air is pumped into the rings from either single or multiple locations, but each shaft or processing plant does not have its own compressed air source. These compressed air rings are usually very large, consisting of up to 75 km of piping [35].

Compressed air rings have the advantage of being able to supply surplus compressed air. This means that if a compressor is out of commission due to breakage or maintenance, another compressor can be used to supply the required compressed air. If localised compressors were used, each shaft or processing plant would need its own surplus compressor. With a ring, only one or two compressors might be required, depending on the size of the ring.

Localised compressors have the advantage that they do not require long lengths of compressed air piping to supply compressed air. Compressed air rings can easily leak a large amount of compressed air thus resulting in wastage [36]. Long lengths of piping introduce pressure losses as well [37]. This makes localised compressors more efficient by reducing pipeline losses, but this can be offset by using larger and more efficient compressors in rings [38].

Underground mining in South Africa is usually divided into three shifts: drilling, blasting and cleaning [39]. The drilling shift is when miners use drills to drill holes for explosives. During the blasting shift, these explosives are used to blow out parts of rocks. These rocks will then be cleaned up during the cleaning shift and brought up to surface for processing.

During the drilling and cleaning shifts, large amounts of compressed air is used [40]. Due to South African mining regulations, a positive pressure is always required underground [41]. This means that even during the blasting shift, compressed air is required although there is no personnel underground. This positive pressure is used to ensure that refuge bays are kept clear of potential toxic gases and that the air remains breathable [42].

Compressed air is used on mines rather than other forms of energy because of safety concerns [43]. South African mines tend to use manual labour to drill holes for blasting, as reefs are narrow and unsuitable for mechanised mining. Drilling underground with handheld pneumatic drills creates the added benefit of generating a cooling effect due to the compressed air depressurising [44].

Pneumatic drills have a lower efficiency than hydraulic or electric drills. Pneumatic drills are noisier [45], but this could change as hydraulic drills and electric drills have only recently started to come of age [46]. Most South African mines are too old to be redesigned for hydraulic or electric drills as they already have the required infrastructure for pneumatic drills in place. Maintenance will also be an issue since pneumatic drills are easier to service and repair.

Mines use compressed air for various applications [47] but pneumatic rock drills are the primary user of compressed air. The entire infrastructure was built around this fact. On mines, compressed air is used by the following:

- Pneumatic rock drills: Used for drilling holes in the rock face.
- Pneumatic cylinders: Used to open cutes and doors throughout the mine.
- Pneumatic loaders: Used to load rock into mine carts.
- Processing plants: Agitation of sediment and instrumentation require compressed air.
- Ventilation and cooling: Refuge bays require positive air pressure.

A supply pressure profile can be generated depending on the combination of equipment, the number of the compressed air requiring equipment and the specific times each of the mine’s three shifts is scheduled. These pressure profiles show the supply pressure for each time of the day. Each mine has a personalised pressure profile with a noted maximum and a minimum pressure. If the pressure is too high, it could damage equipment; if it is too low, equipment seizes to work.

Each end user on a compressed air ring has its own compressed air requirements. If all end users on the compressed air ring are combined, a total pressure profile can be generated. However, if one end user changes its requirements, the pressure profile of the ring will change. This makes the pressure profile of a compressed air ring a complicated problem.

Compressed air rings on mines supply shafts, processing plants and workshops with compressed air. Their pressure requirements vary vastly. Processing plants require a constant high pressure and low flow. Shafts require varied usage from low pressure and low flow to high pressure and high flow. This can sometimes be exploited in splitting the network into a high pressure and a low pressure ring during certain times of the day [48].

Compressor selection for a compressed air ring is usually done on a static basis for the complete ring [49]. The largest compressors are usually preferable to run as it is assumed that they are more efficient than smaller ones. Large compressors are known as baseload compressors. These run continuously and smaller compressors are used to increase flow [50].

Mines attempt to select compressors to run on a schedule. These selected compressors are chosen to match the schedule optimally. Because mines routinely deviate from normal schedules [51], these compressor selections become obsolete. This causes unstable flow delivery [52], waste air and wasted energy.

Unstable flow can cause problems with the control valves on the network as these incur problems operating at the correct pressure due to the oscillating flow delivery from the compressors.

1.3. Compressor characteristics

1.3.1. Compressor types

Some compressors are turbomachines because they convert electrical energy into mechanical energy using rotor blades [53]. Non-turbomachine compressors are classified as positive displacement machines [54]. Compressors are divided into two main groups: intermittent and continuous flow compressors [55]. Due to mine requirements, most compressors found on mines are centrifugal compressors. The classification of the main types of compressors can be seen in Figure 4 [36].

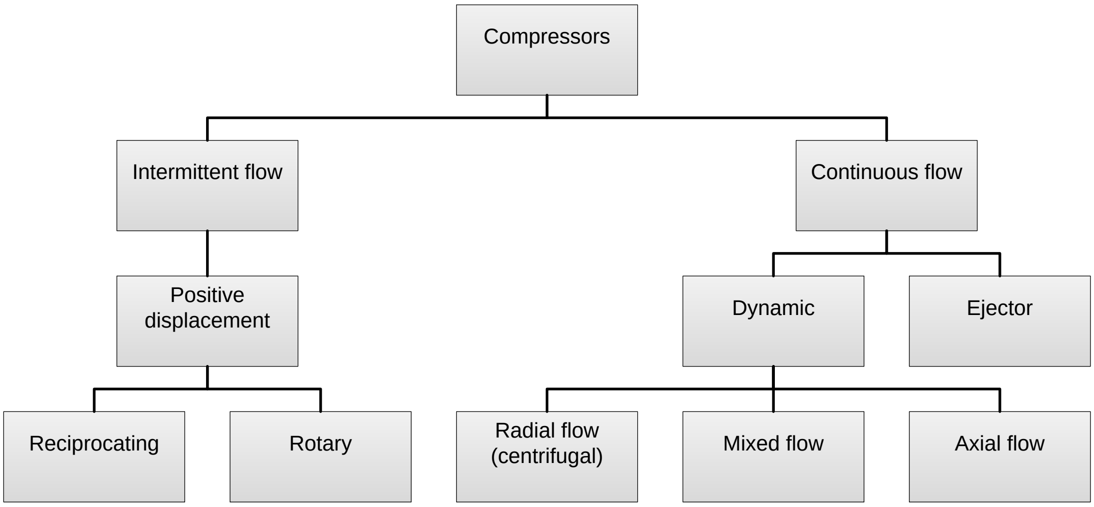

Figure 4: Classification of compressors by types

1.3.2. Centrifugal compressor basics

Centrifugal compressors consist of the following five elements: rotor, guide vanes (blades), shaft, volute and diffuser [53]. The layout of these components can be seen in Figure 5 [56].

Rotor: This is the rotating blade of the compressor. Energy is conveyed into the air from the rotor via mechanical rotation

Guide vanes: The guide vanes preswirl the air coming into the compressor. The guide vane angles are fixed on older and less expensive compressors, but on newer machines, the angle of the guide vanes can be varied.

Shaft: The pipe on which the rotor is mounted.

Volute: The housing of the compressor. The volute guides the air into the diffuser.

Diffuser: This passage converts kinetic energy conveyed onto the air from the rotor into a static pressure.

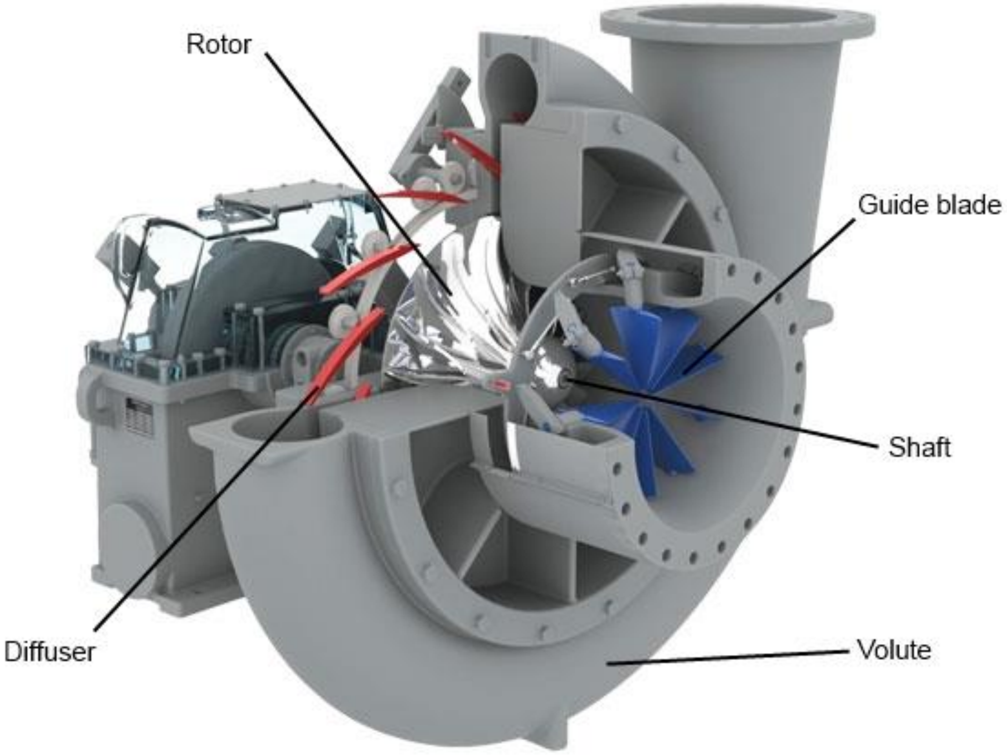

Figure 5: Centrifugal air compressor

Centrifugal compressors are the least complicated compressors, can supply a large quantity of compressed air, are relatively efficient, are reliable and have relatively low vibration [57]. Most mines use electric motors on centrifugal compressors to supply rotation to the shaft. This is because of the efficiency of electric motors and the size of most compressors used on mines [58].

On centrifugal compressors, air enters the compressors via the shaft. The impeller increases the velocity of the air and changes its direction by 90 degrees. The diffuser changes the direction of the air again and converts the velocity of the air into pressure. After leaving the diffuser, the air then travels in the volute to the outlet of the compressor. This can be seen in Figure 6 [59] and Figure 7 [59]. Figure 7 is the exact same as Figure 6 except viewed from the side.

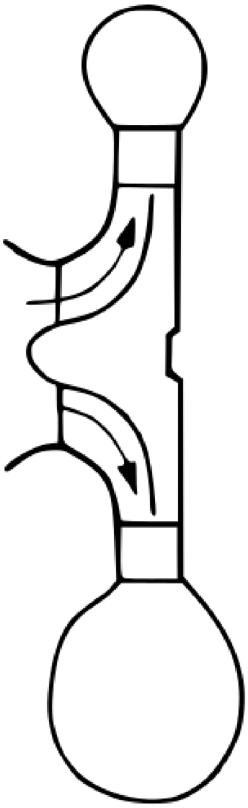

Figure 6: Compressor inlet flow

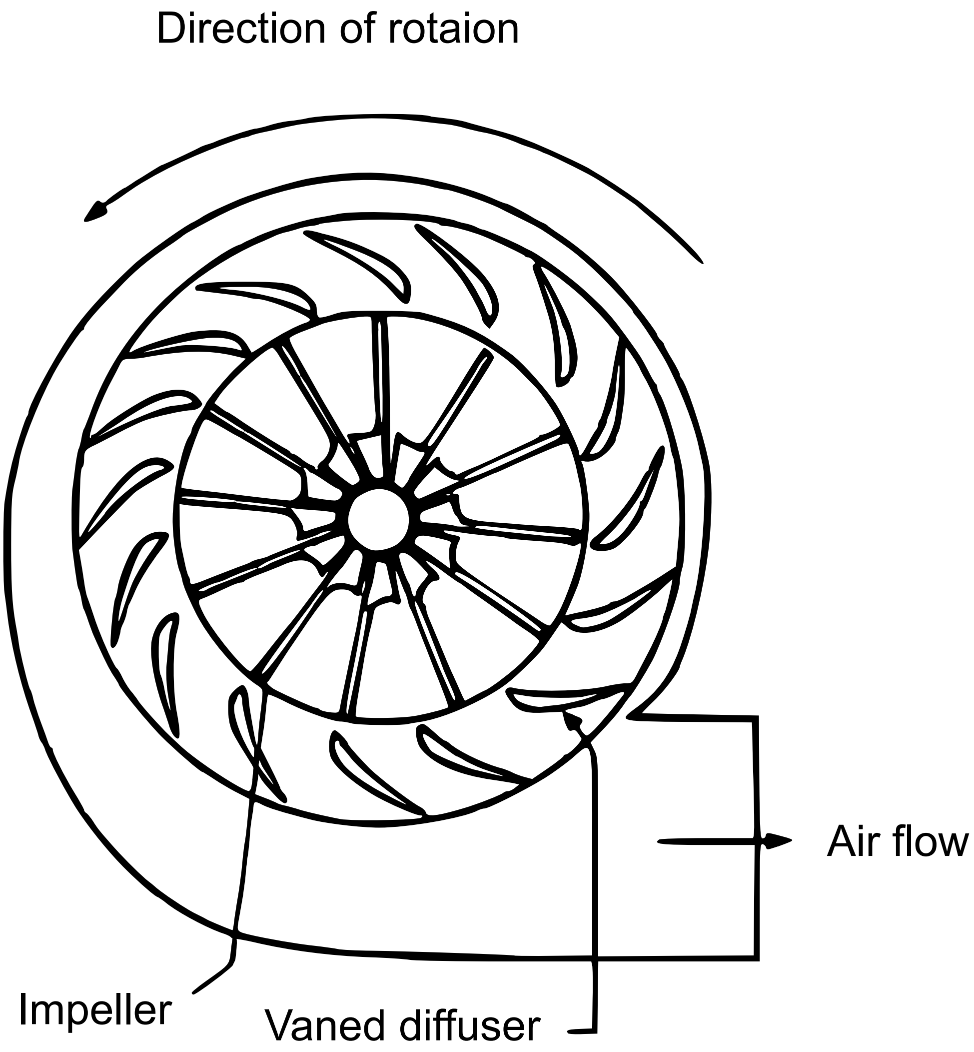

Figure 7: Compressor impeller flow

As the output pressure ratio of a centrifugal compressor is limited to 2.5, multi-stage compressors are used when higher pressures are required [60]. Multi-stage centrifugal compressors operate on the same principle as single-stage centrifugal compressors do, but they contain more than one centrifugal compressor connected in series.

The first stage has a normal inlet while each subsequent stage has the inlet connected to the previous stage's outlet. This can be seen in Figure 8 [61]. Figure 9 [62] shows how the inlets and outlets are combined in a six-stage centrifugal compressor.

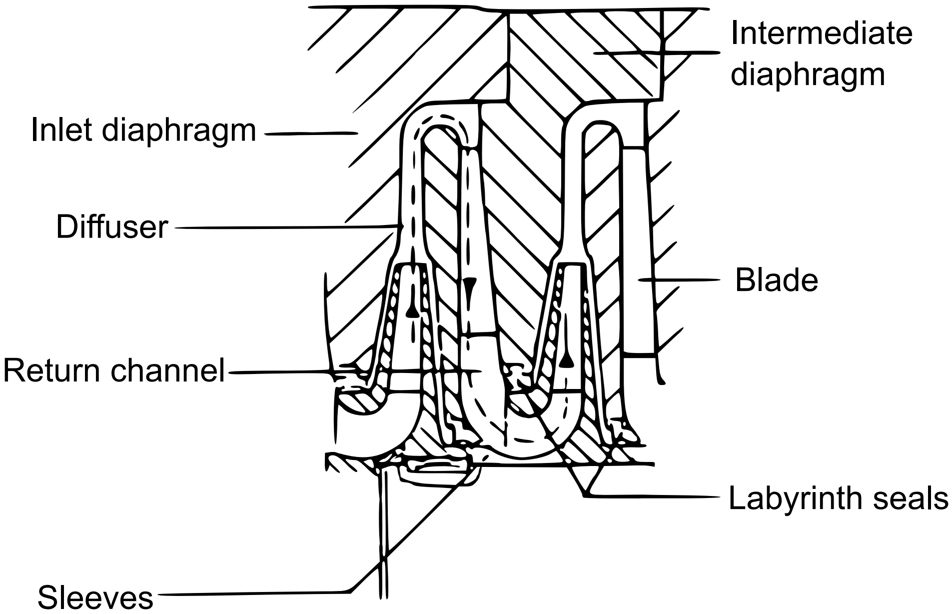

Figure 8: Multi-stage centrifugal compressor airflow

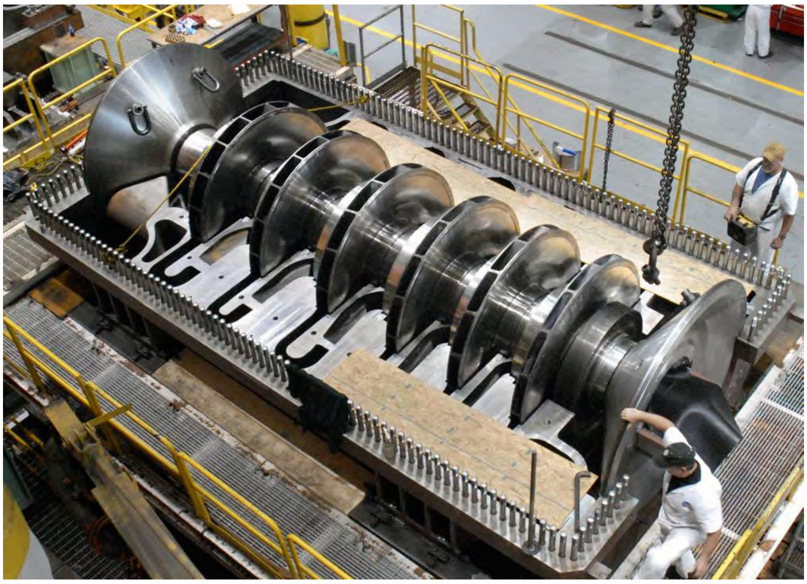

Figure 9: six stage centrifugal compressor

1.3.3. Centrifugal compressor characteristics

Centrifugal compressor performance in comparison to flow

Figure 10 [63] will be used to describe the relationship between pressure ratio and flow through the compressor. A theoretical scenario will be used where the flow increases through the compressor. At α, the airflow through the compressor is 0. As the airflow increases through the compressor, the diffuser will build up pressure and it will rise to a maximum efficiency at γ.

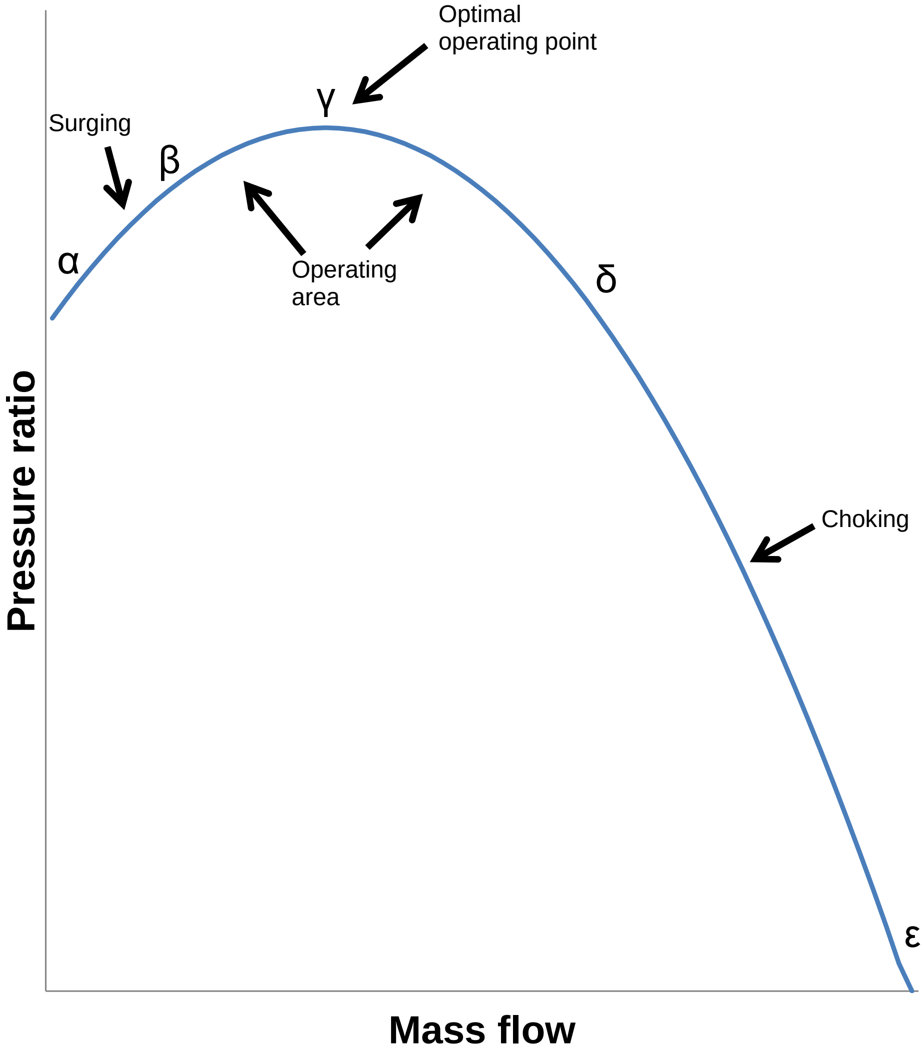

Figure 10: Pressure ratio compared to mass flow

As the airflow through the compressor keeps increasing, it will eventually reach ε due to internal friction. Although α might be reachable in practice, the area between α and β is not usable due to a phenomenon known as surging. δ is the point where the compressors blades will start to stall. In practice, a compressor can only operate between β and δ as well as on α.

Surge

Surging is a phenomenon that causes severe vibrations in compressors [64]. Surging is usually combined with loud noises due to the heavy vibrations. Surging occurs in low flow scenarios and can severely damage the compressor. In Figure 10, surging occurs in the area between α and β.

If the flow decreases from β, the pressure inside the compressors will also decrease. As the pressure decreases in the compressor, a scenario occurs where the pressure inside the compressor is lower than the pressure downstream. This causes negative flow through the compressor as airflow always takes place from high pressure to low pressure.

As the downstream pressure drops, the pressure inside the compressor will again be higher than the pressure downstream of the compressor. This causes the flow to be positive through the compressor again. This is repeated continuously and with the continuous change of direction of the air through the compressor, violent vibrations are caused.

Choke

Choking is the abrupt decrease in performance of a compressor as the air though the compressor reaches sonic conditions [65]. Choking occurs at δ in Figure 10. The location of δ on a real compressor depends on the design of the compressors, thermodynamic properties of the fluid through the compressor and the operating conditions.

Stall

Stalling is a phenomenon which occurs when the flow through the blades of the compressor is non-uniform [55]. Figure 11 [63] shows a blade that is stalling. Blade 2 will cause the angle of air into Blade 1 to be increased while simultaneously decreasing the angle of the air into Blade 3.

This causes the stalling of Blade 2, which in turn allows the air into Blade 2 to recover. As Blade 2 recovers from the stall, Blade 1 will stall. The stall will travel through all the blades in the opposite direction of the movement of the blades. This also increases vibrations, which can damage the compressor.

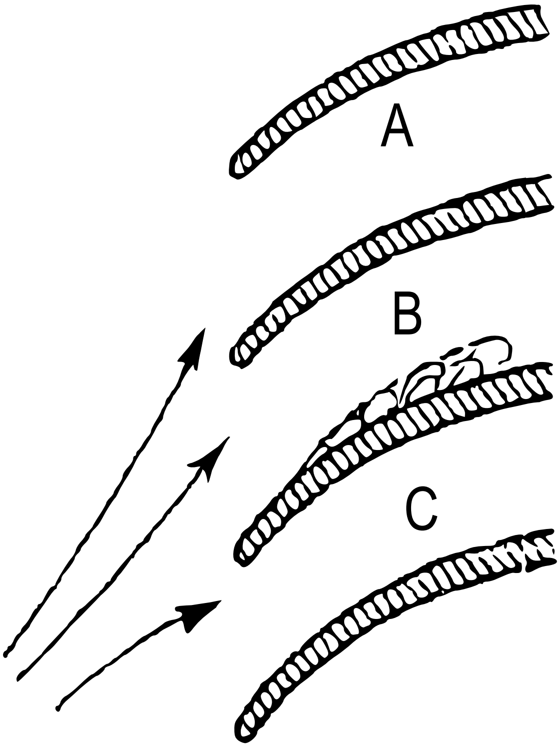

Figure 11: Stalling of compressor blades

Compressor map

The above-mentioned phenomena form a compressor curve. A compressor can be installed with guide vanes and a variable speed drive (VSD) or a combination of the two to obtain different curves, which are similar to Figure 10. When these curves are combined, they form what is known as a compressor map – an example can be seen in Figure 12 [66]. The compressor map in Figure 12 was obtained from a car turbocharger, but this is irrelevant as car turbochargers are small centrifugal compressors.

The left line on Figure 12 forms the surge line. If the combination of pressure and airflow delivery is left of this line, the compressor will surge. The right line forms the choking line where, if the pressure airflow delivery is to the right of this line, the compressor will start to choke.

The top and bottom lines are the maximum and minimum speeds of the impeller. This is only applicable if the compressor can change its impeller speed. The compressor changes its set points on the compressor map via a rotational governor, guide vane or control valve [67]. Guide vanes are found most often on mines due to the size of their compressors.

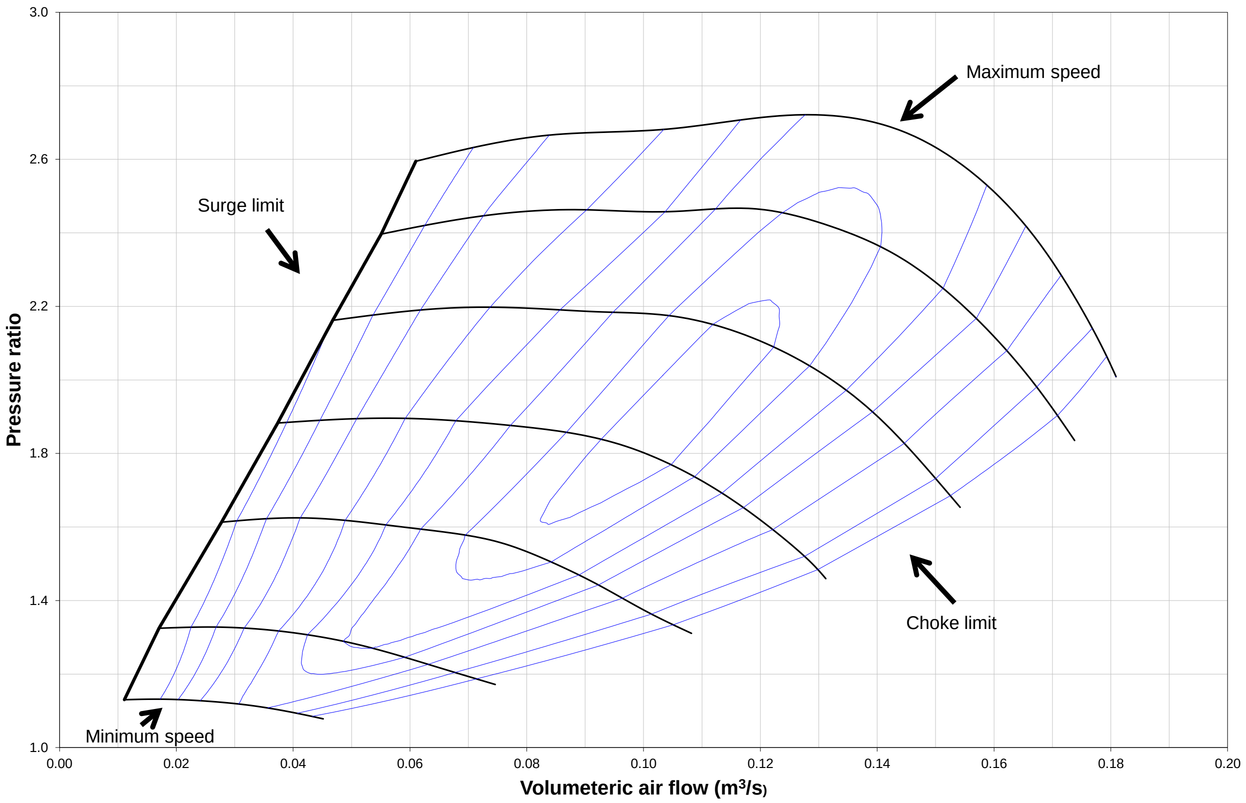

Figure 12: Compressor map

1.4. Need and contributions of this study

Background

South African deep shaft mines use large compressed air rings to supply compressed air to each individual shaft connected to the rings. Most of the time, these shafts are independent business entities. The compressors supplying the rings are run as distinct own business entities that invoice each shaft independently for its compressed air usage.

Each shaft has its own pressure requirement profile, which is determined by its usage of compressed air during a 24-hour period. The compressed air must be supplied to each shaft thus ensuring that the needs of each individual shaft are met. This becomes a complex problem when considering the location of each shaft and that the compressed air ring might have compressors situated in more than one location.

It has already been proven that dynamic compressed air controllers are feasible and that they reduce electrical energy consumption [1]. This thesis aims to address the serious shortcomings of the Dynamic Compressor Selector (DCS) as mentioned by Van Heerden [1].

These are:

- It has to be more stable.
- It has to be able to consider compressor locations.
- It has to use actual efficiency curves with the selection of compressors as there are notable difference between each other.

Contributions

All of the novel contributions of this thesis are given in this section. Each contribution is discussed in four points. Firstly, what must be done to accomplish the novel contribution. Secondly, how this is currently done in the industry. Thirdly, why the present method is not sufficient. Lastly, how the proposed novel contribution solves the issues defined with the present method.

1. Individual compressor set point pressure calculations for dynamic control

• **What must be done?** Develop and implement a practical method to assign optimal multiple pressure set points dynamically to individual compressors.

• **How is it done currently?** – Existing dynamic control systems have one network set point while static control systems have multiple static pressure set points.

• **Why is this not sufficient?** – Compressors are located in different locations. Because the distance between compressors can sometimes cause significant pressure losses, each compressor needs a different set point depending on its location and the demand of the network.

• **How does this study solve this problem?** – This study develops and implements a system to assign an optimal pressure set point dynamically to each compressor. This pressure set point is calculated while taking the location of compressors and end users into consideration.

2. Dynamic optimal operating compressor selection

• **What must be done?** – Develop and implement a practical method to select the optimal operating compressors dynamically.

• **How is it done currently?** – Currently, most networks have static controllers that select compressors based on a static priority list. Dynamic controllers make assumptions based on the size of the compressors as well as that they are most efficient when running on full power.

- **Why is this not sufficient**? – The largest compressor in a network is not always the most efficient compressor. Compressors are also not at their most efficient when running at full power. Neither static nor dynamic controllers take the efficiency curves of compressors into account. The assumption is made that fewer compressors running at full power will always be more efficient than more compressors running at reduced power.

- **How does this study solve this problem?** – The study develops and implements a system to select the optimal operating compressors dynamically, taking the varying efficiencies of these compressors into account.

3. Dynamic optimal trimming compressor selection

- **What must be done?** – Develop and implement a practical method to select the optimal trimming compressor dynamically.

- **How is it done currently?** – Static controllers have static priorities for compressors while dynamic controllers feature dynamic prioritisation of compressors. The trimming compressors are selected from the priority list.

- **Why is this not sufficient?** – Static controllers cannot dynamically adapt to the network. Although dynamic controllers can dynamically change the priorities of the compressors, they are not optimal and location aware. Neither the location of the compressors relative to the demand or the efficiency curve of the compressors is considered.

- **How does this study solve this problem?** – This study develops and implements a system to order the running compressors dynamically to ensure optimal trimming compressor selection.

4. Future flow prediction of demand in network

- **What must be done?** – Develop and implement a system to predict the future demand for the compressed air ring.

- **How is it done currently?** – Future airflow demand is not estimated.

- **Why is this not sufficient?** – In order to control compressors optimally and to prevent cycling of compressors, a way is needed to estimate the required future flow demand of a network. If no indication of the future demand of the network is known, a compressor could be started that is sufficient for the current demand of the network. However, the compressor could be too small to supply the required flow of the compressed air network in the immediate future, thus requiring a larger compressor to be started.

- **How does this study solve this problem?** – The study develops an integrated practical system to estimate the required future flow of the network. This is done in order to anticipate the future demand when a compressor is selected.

5. Location attentive

- **What must be done?** – Develop and implement an integrated dynamic control system to be able to control the network on a two-dimensional level while considering location.

- **How is it done currently?** – Existing dynamic control systems only work on a one-dimensional level with total input versus total demand.

Figure 13: One-dimensional compressed air network

- **Why is this not sufficient?** – To be able to control the network optimally, a controller is required to look at the complete two-dimensional network and be aware where each demand and supply are located. The difference can clearly be seen when comparing Figure 13, which shows a one-dimensional network, with Figure 14, which shows an example of a basic two-dimensional network. With a two-dimensional network, each supply and demand point have individual locations.

- **How does this study solve this problem?** – The study develops a method to give the controller the ability to be location attentive.

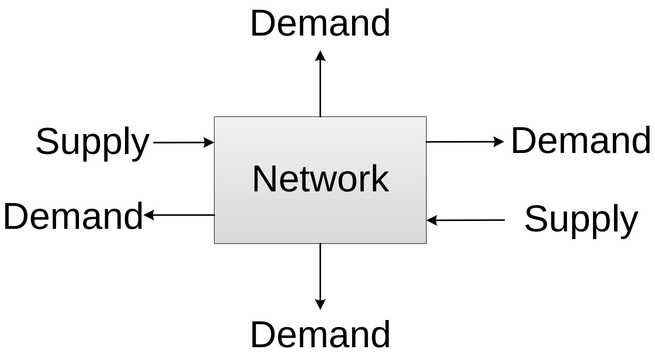

Figure 14: Two-dimensional compressed air network

6. Real-time compressed air network configuration adaption

- **What must be done? –** Develop and implement an application to adapt to changes in the compressed air network in real time when, for example, pipes are locked and closed for compressed air.

- **How is it done currently?** – Existing control systems cannot adapt to changes in the network without manual intervention.

- **Why is this not sufficient?** – One of the most used reasons for changing the network layout is to split the network into high pressure and low pressure compressed air networks during certain times. Splitting the network can save energy due to processing plants and shafts requiring vastly different pressure profiles.

- **How does this study solve this problem? –** This study develops and implements a system that can adapt to changes in the compressed air network in real time. Manual static control intervention will not be required when the network is split into two or more separate networks.

7. Integrated comprehensive compressed air solution

- **What must be done?** – Develop and implement an application to include all the novel contributions mentioned previously in one complete comprehensive solution.

- **How is it done currently?** – Current control systems do not include the novel contributions mentioned previously.

- **Why is this not sufficient? –** Existing control systems either do not work successfully on compressed air networks or they cannot control the network efficiently.

- **How does this study solve this problem?** – By developing and implementing a single comprehensive system that incorporates all seven of the novel contributions, it is ensured that all the problems are addressed.

1.5. Overview of this document

This thesis documents the development of a novel dynamic compressor controller. The goal of the dynamic controller is to select the optimal number of compressors supplying a compressed air network. This will allow the compressor controller to reduce the amount of electrical energy consumed as well as smooth out the delivery of compressed air.

This thesis comprises the following chapters:

Chapter 2: Literature study

This chapter will focus on the shortcomings of existing control systems on mine compressed air networks. It will compare these shortcomings to the new proposed controller.

Chapter 3: System design

The total design of the system will be laid out in this chapter. Requirements as well as their theoretical results will be discussed. A detailed design of each component of the complete controller will also be documented.

Chapter 4: Implementation and results

This chapter will contain all documentation on the implementation of the compressor controller on actual mine compressed air networks. It will also discuss the results obtained from testing the compressor controller on mine compressed air networks in detail.

Chapter 5: Conclusion

The conclusion will be discussed in this chapter. Possible future work and improvements will also be laid out in this chapter.

Appendix A: Software procedures

All software procedures used in the design will be covered in a short overview and explained in this chapter.

Appendix B: Extra compressor running graphs

This chapter includes a few extra compressor running graphs of case study 1. Each of the graphs represents one extra full day.

2 Study of compressor controllers

2.1 Preamble

This chapter focuses on a literature study of the Dynamic Compressor Controller (DCC). Techniques used to control compressors will be researched. The functionality and working of the existing dynamic control will be discussed in detail. Software design methods will be explored in Appendix A: Software procedures. Lastly, the proposed DCC will be compared with existing systems and their limitations.

2.2 Overview compressor control

2.2.1 History of compressed air control

Compressor control can be divided into demand-side control and supply-side control. Most of the control of compressed air takes place on the supply side since this is where the actual compressors are present. The perfect control system for a compressor should control both the demand and supply side – matching the supply to the demand. This would allow for optimal control.

Due to relatively inexpensive electricity prices in the past, undesirable electricity usage behaviour occurred. An example of this is that compressors ran at full capacity for 24 hours a day [68]. This was done because compressed air operators feared decreased production output due to low air pressures and because the cost of the electricity was offset easily by the income from production.

The first efforts to reduce compressed air on South African mines were peak clipping projects [69]. These projects reduced the pressure during Eskom peak times [70]. In a peak clipping demand-side management (DSM) project, the electrical load is reduced during peak times. This differs from a load shifting project in that energy is not shift to other periods [71]. The difference between peak clipping and load shifting can be seen in Figure 15 [72] and Figure 16 [73].

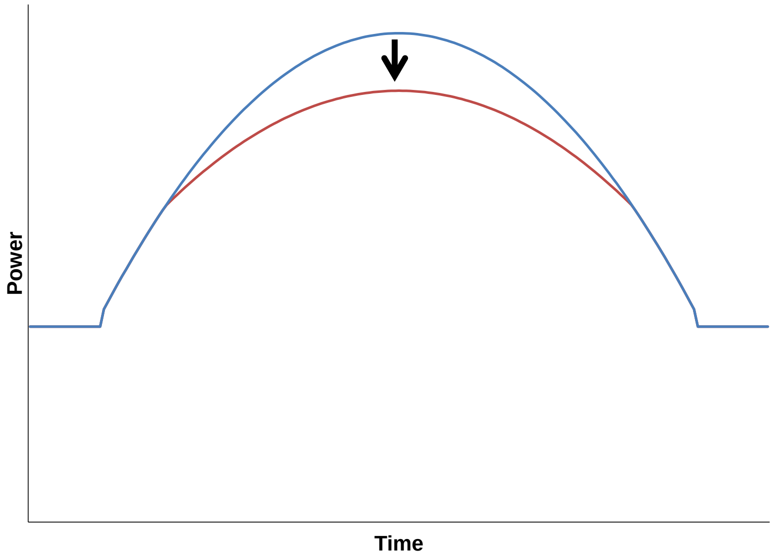

Figure 15: Peak clipping

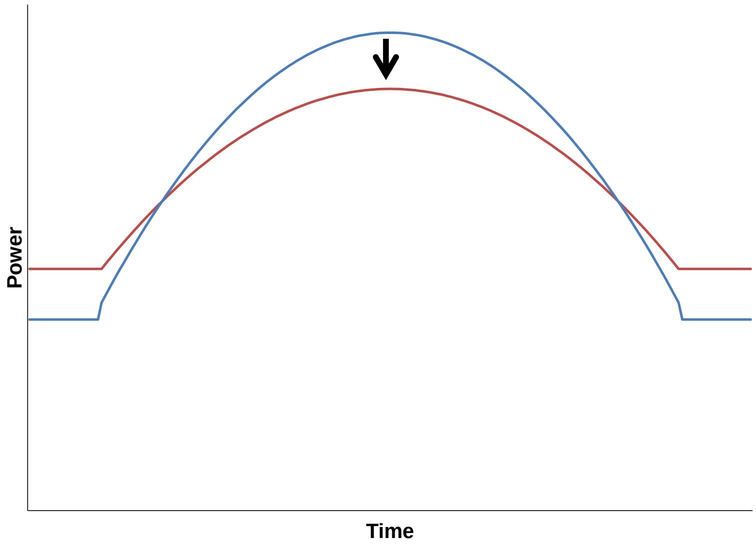

Figure 16: Load shifting

The next step in reducing compressed air usage was to attempt to reduce the compressed air pressure set point to match demand. A typical mine compressed air set point profile is shown in Figure 17 [74], from this it can clearly be seen how it is attempted to match set point to the demand. The compressed air pressure set point is the pressure set point at which the compressed air is supplied to the compressed air ring. These were energy efficiency DSM projects [75]. An example of an energy efficiency DSM project can be seen in Figure 18 [76].

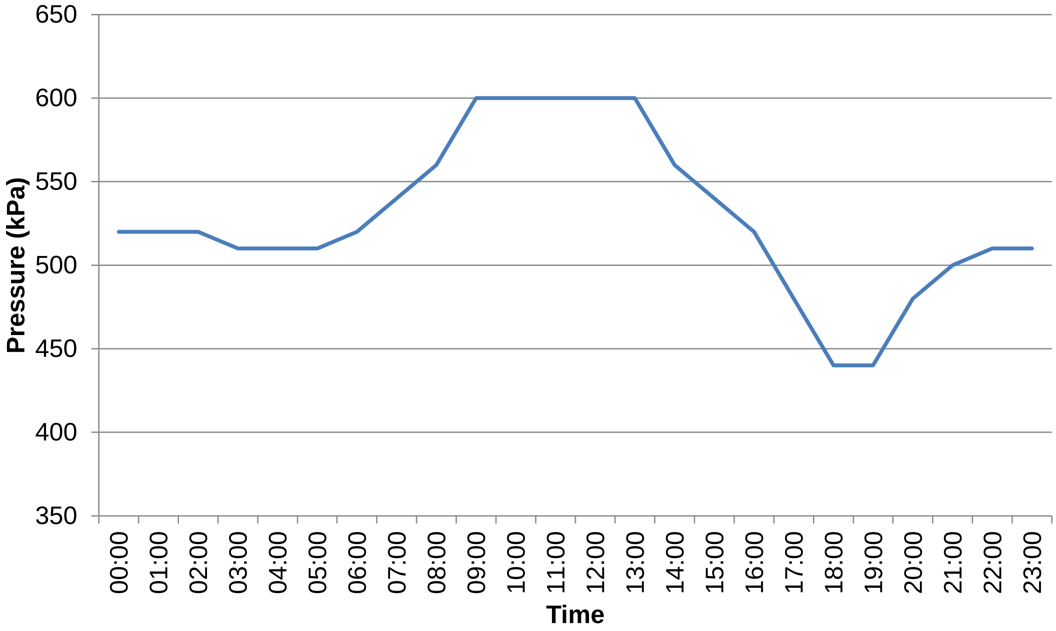

Figure 17: Typical mine pressure set point

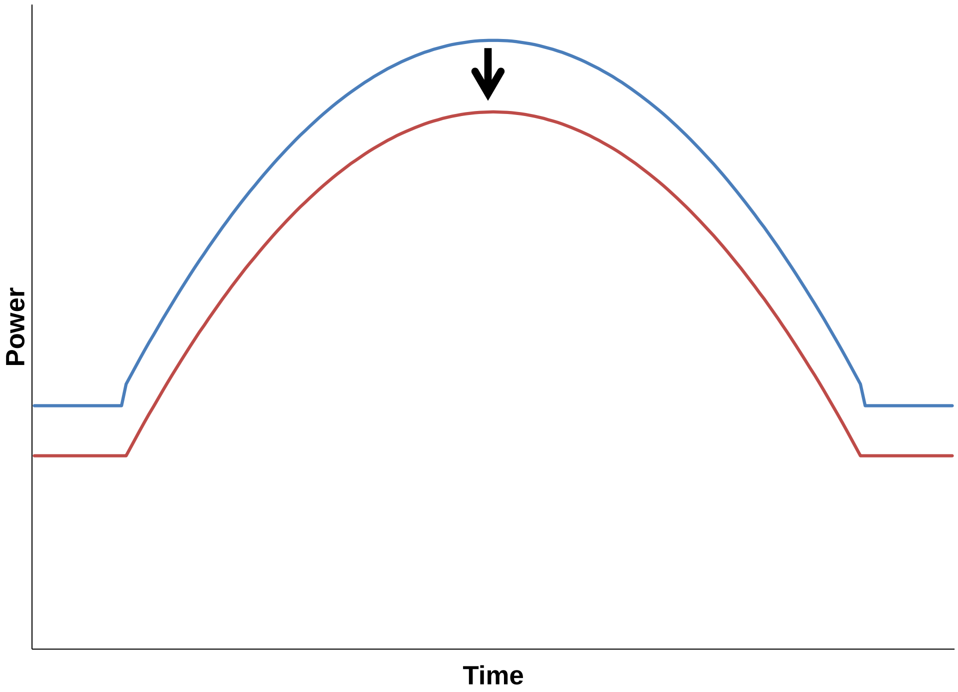

Figure 18: Energy efficiency baseline

Together with these reduced pressure set points in off-peak periods, compressed air networks were split into high and low pressure sub-networks [77]. Gold processing plants typically require high pressure during the entire day but do not require high flow [78]. This forces the ring pressure to be unnecessarily high even during off-peak periods.

To reduce the energy consumption of a compressed air network further, the focus shifted from the supply side to the demand side. The demand-side approach entails fitting pressure regulating valves for each end user on the compressed air ring [79]. By being able to restrict the flow through the valve, the pressure downstream of the valve is limited to a preset value. This allows for a further reduction in the pressure set point by prohibiting the oversupply of a compressed air user.

By fitting each end user with a pressure valve, each user can be assigned a unique pressure set point [80]. In an attempt to optimise these end user set points, the compressed air network is simulated with commercial flow simulation programs such as KYPipe2014 and Flownex® [81]. By simulating the compressed air network, accurate set points for each user can be calculated. These individual set points for each of the end users allow the supply to be matched to the demand.

The next logical step in compressor control was combining simulation with control to simulate in real time, as discussed by Venter [49]. By simulating in real time, the compressor controller can match the supply and demand accurately and continuously. While normal simulation can also match the supply and demand, they do not allow for transition periods or for deviations in normal operating conditions.

Compressor control can be divided into two distinct types: local and network control. Compressors without local control cannot function because local control gives input to the compressor to deliver a specific pressure and flow. Network control determines when a compressor should start, stop, unload and load. It also determines the pressure needed from each compressor.

2.2.2 Local control

According to Bloch, a local controller on a compressor has the following objectives [82]:

- **Performance control:** Ensure that the compressor delivers a certain pressure and flow.
- **Surge protection:** Prevent the compressor from surging.
- **Limiting control:** Keep the compressor away from limiting variables such as temperature.
- **Loop decoupling:** Minimise adverse interactions between control functions.
- **Event sequencing:** Automate processes such as start-up and shutdown.
- **Host communication:** Integrate communication to a supervisory control system such as a Supervisory Control and Data Acquisition (SCADA) system.

There are ways to control the performance of a compressor to match a required pressure or flow requirement [83]:

- **Variable speed control (VSD):** Speed control is the most efficient way of controlling a compressor [83]. But, it is also a very expensive because of the implementation cost due to the size of the compressors [35]. Speed control uses VSDs – also known as variable frequency drives – to increase and decrease the speed at which the impeller rotates. This in turn causes pressure and flow to decrease or increase.
- **Suction valve throttling:** Suction valve throttling is done by restricting airflow into the inlet of the compressor [84]. This is done by installing a control valve upstream of the compressor inlet. Flow is reduced by restricting the airflow into the compressor.
- **Discharge valve throttling:** Discharge valve throttling is the opposite of suction valve throttling and is done by restricting the airflow out of the compressor [83]. This method is not widely used – it is very inefficient because the compressor has to overcome backpressure.
- **Variable inlet guide vane control:** Guide vanes are blades located at the inlet of the compressor [85]. These blades change the velocity of air through the rotor. The energy transferred into the air is relative to the velocity of the air relative to the rotor. On centrifugal compressors, only the first stage can be fitted with guide vanes. Variable inlet guide vanes offer a better alternative to suction and discharge valve control in terms of energy efficiency. At a full open guide vane position, a compressor does not necessarily run at full efficiency [86].
- **Discharge control valve:** Discharge valve control is the most inefficient way of controlling compressor output. A control valve – also known as a blow-off valve or bleed valve – is used to discharge compressed air into the atmosphere. This method is widely used on mines due to its very inexpensive installation costs. It is also a cheap control method to avoid surge.

All the local control methods listed above are used to control how a compressor delivers only a certain set point of compressed air. Many of these methods are also used in surge protection and limiting control.

2.2.3 Network control

Network control covers a much higher level of compressed air control than local control. A network controller will not control the start/stop sequence or the inlet guide vanes of a compressor to supply compressed air at a certain pressure set point for example. A network compressed air controller will be used to supply instructions to each local controller of a compressor.

A network controller has the following goals:

- Reduce energy consumption
- Reduce compressor cycling
- Start, stop, load and unload compressors (instructions to do so)
- Set performance set points for local controllers

Reducing energy consumption of a compressed air ring is accomplished by assigning the correct start and stop priorities, issuing correct start and stop commands and setting correct set points for the local controllers. By issuing start and stop commands, compressors can be removed from the ring and be shut down to eliminate their energy consumption.

When not in use but running, a compressor can be unloaded by removing the compressor from the compressor air ring while still keeping the compressor running [87]. This is accomplished by opening the blow-off valve completely and running the compressor while it idles. This reduces the energy consumption while still keeping the compressor running. This allows the compressor to start pumping compressed air into the ring immediately if more compressed air is desired.

Compressor cycling is caused when a compressor is started and stopped excessively within a short period of time [88]. The time period differs depending on the size of the compressor – the larger the compressor, the larger the time period. Cycling causes unnecessary wear on the compressor, which shortens its lifespan and increase the frequency of required maintenance. When starting and stopping compressors, this extra cost needs to be offset by the energy savings obtained [35].

Most network controllers issue performance set points to local controllers based on static inputs. These inputs are usually in the form of static priorities for starting and stopping compressors, as well as static set points for average ring pressure. The network controller issues individual instructions to each of the local controllers based on these inputs.

2.3 Evaluation of existing control systems

Table 1 (adapted from [1]) shows a brief overview of all the key features for the below listed compressed air systems.

Table 1: Comparison of compressor controllers

	Integrated local controller	Automated control	Manual override	Static priorities	Dynamic priorities	Historical data	Monitoring	SMS alarms	Compressor limit	Simulation	Valve control	Number of dynamic set points	Predict future demand	Compressor priority ques	Location aware	Network configuration adaption
PL4000	X	X	-	-	X	X	X	-	∞	-	-	0	-	1	-	-
airtelligence provis 2.0	X	X	-	-	X	X	X	-	16	-	-	0	-	1	-	-
Hiprom	X	X	X	X	-	-	X	-	∞	-	-	0	-	1	-	-
Atlas Copco ES	X	X	X	X	-	-	X	X	∞	-	-	0	-	1	-	-
Murphy Centurion PLUS$^{\text{TM}}$	X	X	-	X	-	X	X	-	∞	-	-	0	-	1	-	-
CompAir SmartAir Master	-	X	-	X	-	X	X	-	12	-	-	0	-	1	-	-
EMS	-	X	X	X	-	X	X	X	∞	-	-	0	-	1	-	-
REMS-OAN	-	X	X	X	-	X	X	X	∞	-	X	0	-	1	-	-
KYPipe2014	-	-	-	-	-	-	X	-	∞	X	-	0	-	1	-	-
DCS	-	X	X	X	X	X	X	X	∞	X	X	1	-	1	-	-
DCC	-	X	X	X	X	X	X	X	∞	X	X	∞	X	∞	X	X

Legend: (X) – contains the mentioned feature; (-) – feature not available in the controller, (∞) – infinite number

PL4000

Pneu-Logic developed the PL4000 as a black box compressed air controller [89]. The controller can automatically stop and start compressors to ensure that a static preset pressure set point is maintained. A memory device can be installed inside the control unit of the PL4000 to allow historical logging of data via text files

The PL 4000 has a very minimalistic interface and is only used to ensure that the controller operates as intended. The intended market for the PL4000 is small factory compressors rather than large mine compressed air networks. It cannot control the valves of a compressed air network to optimise airflow.

airtelligence provis 2.0

BOGE compressed air systems designed the airtelligence provis 2.0 to also be a black box system [90]. The airtelligence provis 2.0 controller can also control fans and dryers. It is a more comprehensible control system than the PL4000. The controller allows the user to access all historical information from a web browser interface.

The shortfall of airtelligence provis 2.0 is similar to PL4000 because it is also geared towards small factory compressors – it cannot optimise a compressed air network by controlling valves. Another large limitation of the controller is that it can only control up to 16 compressors. This might not be an issue depending on the size of the compressed air network, the demand for compressed air and the size of the compressors installed.

Hiprom controller

Hiprom has been a division of Rockwell Automation since 2011. The Hiprom controller was specifically designed for mines [49]. This controller is a local controller with network controller capabilities. The network controller has a few basic features such as it starts, stops, unloads and loads compressors according to preset static pressure set points.

The Hiprom controller does not have alarm capabilities or the ability to log any historical values. It also has the shortfall that it can only control from static compressor priorities and static compressor set points.

Atlas Copco ES

The ES range of compressor controllers is made by Atlas Copco [91]. These controllers offer dynamic prioritisation but offer static and manual override priorities as well. ES controllers control compressors on a pressure band and keep the network at that specific pressure band. Short message service (SMS) alarm communication is also offered by the controllers.

Lack of dynamic set point control is one of the shortfalls of the ES range of controllers. The objective of ES controller controls is to keep the number of compressors to a minimum. This is not always economical because full power is not always the most efficient setting and it will sometimes be more economical to run two compressors rather than one. The ES range of controllers is also designed for factories rather than mines.

Murphy Centurion PLUS^{TM}

The Centurion PLUS^{TM} controller by Murphy is a compressor controller primarily designed for engine-driven compressors [92]. The controller can stop and start compressors based on custom inputs such as pressure. The controller can also control other equipment via custom outputs. Historical data is available via csv (comma separated values) files that can be downloaded from a universal serial bus (USB).

A shortfall of the Centurion PLUS^{TM} is that it only features static priorities and set points. If the Centurion PLUS^{TM} is required as a local controller, the local controller must be programmed on the Centurion PLUS^{TM} itself.

CompAir SmartAir Master

The SmartAir Master controller by CompAir is a compressor controller that can prioritise compressors dynamically to sustain a pressure set point of a network [93]. The controller features removable secure digital (SD) card storage. Historical data will be stored on the SD card, which can then be removed and inserted into an SD card reader.

The controller has much the same shortfalls as the Centurion PLUS^{TM}, except that it only supports up to 12 compressors. The pressure set points are also static.

Energy Management System (EMS)

The EMS is a general controller developed by Du Plessis that was extended to include a network compressor controller [94]. The controller is designed to work from a machine that has the Microsoft Windows® operating system. EMS communicates instructions to the SCADA via open platform communications (OPC). EMS features the ability to start, stop, load and unload compressors according to set points.

The ability to send out custom alarms via SMS and email are included in the EMS; these alarms are created by the user for certain conditions. Historical logging is also a feature where data is logged in two-minute intervals. This logged data can include any SCADA OPC tag the user desires. The EMS is built upon an energy management platform. This energy management platform supplies the OPC communication, data logging and alarms.

The shortfalls of the EMS is that it only controls via preset static compressor set points. The compressor priorities are also static. The controller will start or load a compressor when the current delivery pressure drops below the set point and stop or unload a compressor when the delivery pressure rises above the set point. This ensures that the minimum number of compressors is used.

Real-time Energy Management System – Optimised Air Networks (REMS-OAN)

The REMS-OAN [95] was designed as an evolution of the EMS, it is developed from the same platform as the EMS but features new unique components. The REMS-OAN includes all the features of the EMS, including the static compressor set points and static priorities. The REMS-OAN added the ability to control network valves on the compressed air network. In an effort to optimise the network, all valves were controlled on the mine up to shaft level. This allowed the airflow to be directed optimally to where it was needed.

The REMS-OAN included the same shortfalls in compressor control as the EMS. Although the REMS-OAN could reduce the demand on the compressed air ring by reducing wastage, it still only had static priorities and pressure set points. Due to this, demand and supply could not be optimally matched.

KYPipe2014

KYPipe2014 is a simulation software package developed by KYPipe [96]. KYPipe2014 is not a compressor controller, but only a simulation package to simulate compressed air networks. KYPipe2014 was used in simulations of compressed air networks in an attempt to match supply and demand of the networks at certain times [81]. This allowed for better efficiency through the day combined with static compressed air controllers.

The shortfall of this combined approach is that a trained and experienced operator must simulate certain times and shifts on the compressed air network. This approach can only match the supply and demand of certain times and for this to be effective at all, the simulations have to be performed periodically. If the demand differs from the supply in any way, the simulated results are void and new simulation parameters must be updated by a trained and experienced operator.

Dynamic Compressor Selector (DCS)

The DCS controller was designed by Van Heerden [1] as a complete compressor controller based on work done by Venter [49]. This controller included all the features of the EMS and the REMS-OAN as it shares the same platform as the EMS and the REMS-OAN. This allows the controller to send out alarms, supply historical logging via csv files and control network valves.

The controller can calculate compressor priorities and compressor set points dynamically based on the demand of the network. This allows the controller to match supply and demand on the compressed air ring more successfully than previous controllers.

The shortfalls of the DCS controller are that it could not properly anticipate future demand for compressed air and it could not prioritise compressors based on their location. Supply and demand matching was done one-dimensionally as seen Figure 13, which inhibited the controller from being location aware. This shortfall stops the controller from being able to control all networks as it could select a compressor that could not supply the required compressed air due to distance, pipe diameter etc.

Set points are also limited to only one per network, which inhibits the efficiency of the controller. The controller could only generate one dynamic compressor queue of priorities. This list was also generated on the assumption that larger compressors were more efficient. As mentioned by Van Heerden, the controller also suffered from stability issues that inhibited the controller from actively controlling the network without human intervention [1].

2.4 Focus on dynamic control

2.4.1 Introduction

This subsection of this chapter will discuss the DCS developed by Van Heerden [1] in detail. The DCS system was developed as a dynamic compressor air controller. It can be seen as the precursor to the DCC. The Simulations section (Section 2.4.3) will discuss formulas used in the simulation. The System section (Section 2.4.4) will discuss the workings of the controller.

The controller designed by Van Heerden [1] was based on work done by Venter [49] on dynamic compressor control. Venter made the following assumptions [49]:

- The compressed air network will continuously be in steady state.
- The airflow through the network is one-dimensional, incompressible and isothermal.
- Each section of pipe has the same roughness throughout.
- All historical logged data is correct.
- Flow losses due to air leaks are negligible.
- The entire compressed air ring is at the same height.
- The viscosity of air at 316 K is constant at 3.0134×10^{-5} kg /m•s.

2.4.2 Simulation formulas

In this section, all the basic formulas used will be discussed.

Fluid density (ρ)

Fluid density is the density of a fluid per volume and can be expressed as kilogram per cubic metre. The formula for fluid density is as follows [97]:

$$\rho = \frac{P}{RT} \qquad \text{Equation 1}$$

ρ : Fluid density [kg/m^3].
R : Gas constant [J•K^{-1} kg^{-1}].
T : Temperature [K].
P : Pressure [Pa].

Reynolds number

The Reynolds number is a dimensionless number used in fluid mechanics [98]. It is defined as the ratio of inertial forces to viscous forces and categorises flow as either laminar or turbulent flow. Due to the mass flow of compressed air rings on mines, the flow will always be turbulent. The formula for calculating the Reynolds number of a fluid is as follows:

$$Re = \frac{\rho v D}{\mu} \qquad \text{Equation 2}$$

Re: Reynolds number.
v : Fluid velocity [m/s].
D : Pipe diameter [m].
μ : Viscosity [kg/(sec•m)].

Bernoulli's principle

Bernoulli's principle is a fluid dynamics principle that states that for an inviscid flow of a non-conducting fluid, an increase in the speed of the fluid occurs simultaneously with a decrease in pressure or a decrease in the fluid's potential energy [99]. Bernoulli's principle can be expressed as a formula as follows:

$$\frac{\rho v^2}{2} + pgz\ + P\ = constant \qquad \text{Equation 3}$$

ρ : Fluid density [kg/m^3].

v : Fluid velocity [m/s].

g : Gravitational acceleration [m/s^2].

z : Measured height [m].

P : Pressure [Pa].

When this principle is applied to a pipe with two opposite ends, the formula will be:

$$\frac{\rho_1 v_1^2}{2} + \rho_1 g z_1\ + P_1\ = \frac{\rho_2 v_2^2}{2} + \rho_2 g z_2\ + P_2 \qquad \text{Equation 4}$$

ρ : Fluid density [kg/m^3].

v : Fluid velocity [m/s].

g : Gravitational acceleration [m/s^2].

z : Measured height [m].

P : Pressure [Pa].

Because the altitude on the entire compressed air network is constant and the fluid density will be the same through the pipe, the following holds true:

$$p_1 g z_1\ = p_2 g z_2 \qquad \text{Equation 5}$$

g : Gravitational acceleration [m/s^2].

z : Measured height [m].

P : Pressure [Pa].

Thus, Bernoulli's principle can be simplified to:

$$\frac{\rho v_1^2}{2} + P_1\ = \frac{\rho v_2^2}{2} + P_2 \qquad \text{Equation 6}$$

ρ : Fluid density [kg/m^3].

v : Fluid velocity [m/s].

P : Pressure [Pa].

Because a pipe will have frictional losses, these can be included as follows:

$$\frac{\rho v_1^2}{2} + P_1 = \frac{\rho v_2^2}{2} + P_2 + P_{loss} \qquad \text{Equation 7}$$

ρ : Fluid density [kg/m^3].

v : Fluid velocity [m/s].

P : Pressure [Pa].

P_{loss} : Cumulative friction losses between P_1 and P_2.

Mass flow

In fluid dynamics, mass flow is defined as the mass of a fluid which passes per unit of time [100]. This is usually defined as kilogram per second. The formula for mass flow through a pipe can be calculated as follows:

$$m = \rho v A = \rho Q \qquad \text{Equation 8}$$

m : Mass flow [kg/s].

ρ : Fluid density [kg/m^3].

v : Average pipe velocity [m/s].

A : Area of the pipe [m^2].

Q : Volume flow [m^3/s].

Darcy friction factor

The Darcy friction factor is a dimensionless number used to describe the frictional losses in a pipe [101]. It uses the Darcy–Weisbach equation. The Colebrook-White equation is used to solve the Darcy friction factor for turbulent flow [102]:

$$\frac{1}{\sqrt{f}} = -2 \bullet \log(\frac{\frac{\varepsilon}{D}}{3.7} + \frac{2.51}{Re\sqrt{f}}) \qquad \text{Equation 9}$$

f : Darcy friction factor.

ε : Pipe roughness [μm].

D : Pipe diameter [m].

Re : Reynolds number.

Using the Swamee–Jain approximation, the Darcy friction factor can be computed as [103]:

$$f = \frac{0.25}{\left[\log\left(\frac{\varepsilon}{3.7D} + \frac{5.74}{Re^{0.9}}\right)\right]^2} \qquad \text{Equation 10}$$

f : Darcy friction factor.

ε : Pipe roughness [μm].

D : Pipe diameter [m].

Re : Reynolds number.

2.4.3 Simulation

Van Heerden's simulation system is based on nodes [1]. A limitation of the system was that each node could only support one, three or four pipes. An example node can be seen in Figure 19. With this system, you could dynamically build most compressed air networks. By connecting each node to pipes, a network could be built by having each node represent a junction.

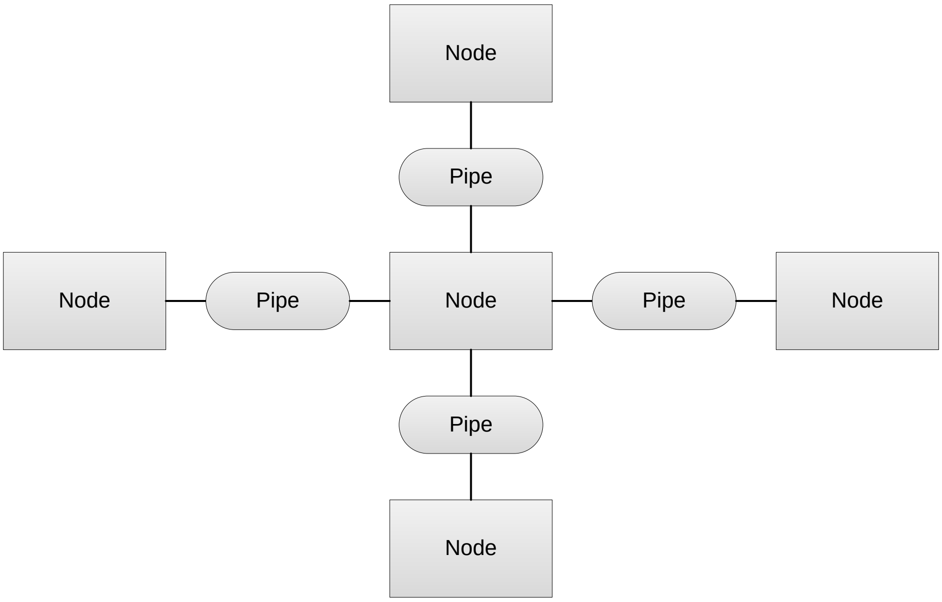

Figure 19: Node layout

A flow balance is done on each node. The total airflow into the node is equal to the total flow out of the node. Van Heerden then estimated the pressure of the centre node [1]. Using Bernoulli's principle, the mass flow in each pipe is calculated. If it is assumed that flow into the centre node is positive and flow out of the node is negative, the sum of all the pipe flows is calculated. An example network can be seen in Figure 20.

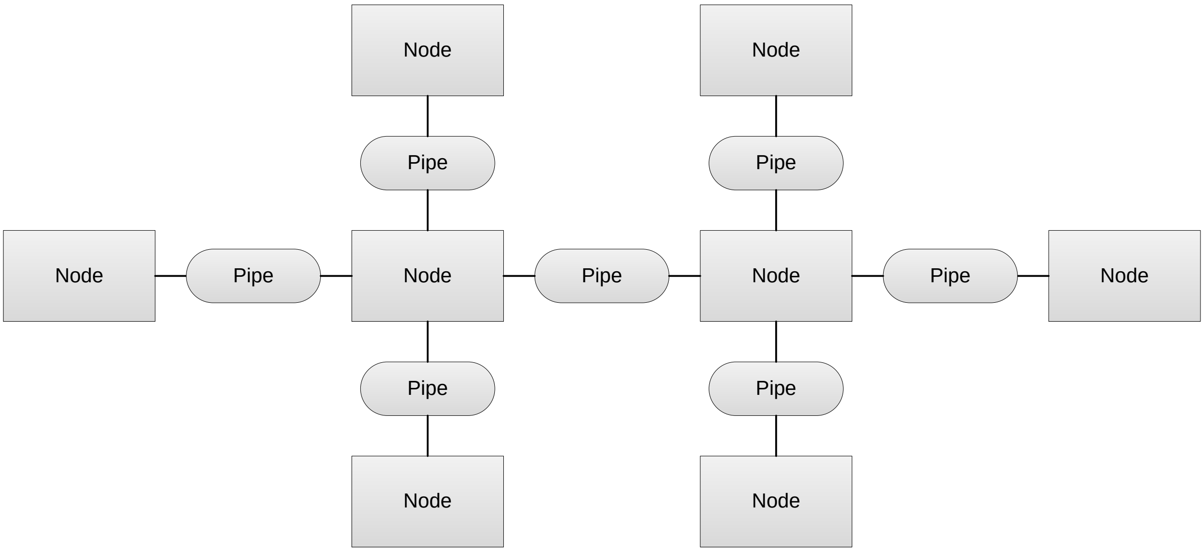

Figure 20: Example network

The pressure of the centre node will be iterated until the sum of all the pipe flows is equal to 0. The node will then be seen as balanced and the pressure as known. This process will be repeated for all nodes that have three or four pipes connected to them. The system will then rebalance the first centre node because the pressures of the neighbour nodes have been changed since the node has been balanced.

This process will be repeated until the rate of change between the pipe flows of two neighbouring nodes dropped close to 0. Van Heerden proved that the accuracy of this simulation technique was comparable to that of commercial airflow simulation packages [1].

In an attempt to increase the calculation speed, Van Heerden designed the system with a feedback loop. The previous simulation's results were used as the starting point of the new simulation. This method improved the calculation speed of all follow-up simulations. However, this method proved to be very unreliable for finding a solution because if an erroneous output was calculated, the next simulation would be based on this.

This caused the system to end in an endless calculation loop. In an attempt to reduce this, Van Heerden added filters to stop erroneous input into reaching the simulation. This improved the stability but did not solve it completely as erroneous input still reached the simulation and it would then still end up in an endless loop.

2.4.4 System

The DCS system designed by Van Heerden worked with six different components [1]. All of the components and the relationship between them can be seen in Figure 21 [1]. The Air Pipe component was used to represent a pipe. All properties of a pipe were inserted into the component, while the Air Node component was used to represent a junction or end point of the compressed air network. A network can be built by combining these two components.

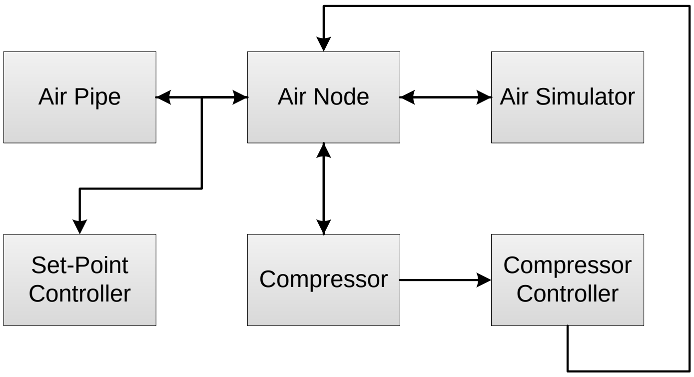

Figure 21: DCS system

The Air Simulator component was used to simulate the network. This component saved all the results back into the Air Node and Air Pipe components after the simulation. The Set point Controller was then used to calculate a single pressure set point for the network. To calculate the set point, an input pressure for all compressor houses was estimated and iterated until all end users had their required air pressures. It also calculated using the simulation results from the Air Simulator.

The Compressor component was used to represent compressors. This component stored all the information of the relevant compressor. The Compressor Controller component was used to calculate priorities of the compressors and issue stop, start, load and unload commands to the compressors.

2.5 Summary

This chapter described compressor control on mines as well as how compressor control evolved to be more efficient. It also investigated how DCC compares with other compressed air systems. The previous DCS system designed by Van Heerden [1] was explained in detail.

3 System design

3.1 Preamble

This chapter will focus on the design of the DCC system and its requirements. The design of each component and its inner workings will be discussed in detail. Theoretical results of the system will be tested in this chapter as well.

3.2 Design requirements

The design requirements for the DCC was based on requirements identified by Van Tonder [35] and requirements used by Van Heerden [1]. It was assumed that since the DCC is a network controller, it would work in conjunction with a local controller thus fulfilling some of the requirements. The requirements for the DCC were:

- Control any industrial compressed air network.
- Communicate with systems via OPC.
- Create dynamic individual compressor set points.
- Prioritise compressors dynamically.
- Estimate future state of the compressed air network.
- Issue start, stop, load and unload commands to all compressors.
- Control demand valves on the compressed air network.
- Include a reporting ability or support for a reporting program.
- Include the ability to split networks dynamically in a high pressure and low pressure networks.
- Include compressor location and network compressor air leakage into compressed air simulations.
- Include the ability to limit user access control.
- Display results via a graphical user interface (GUI).

These requirements are discussed in more detail in the section that follows.

Control any industrial compressed air network

To be able to control any industrial compressed air network, the user must be able to recreate any compressed air network dynamically. This requirement forces the program to be component-based. This allows the user to place components and link them dynamically.

This requirement is only relevant if the network can be controlled remotely and has the required automation and communication in place.

Communicate with systems via OPC

OPC was created to be an industry standard method of communications between devices by different manufacturers in real time. By supporting OPC, the program can communicate in real time to SCADA or data acquisition servers. These, in, turn communicate with devices.

Create dynamic individual compressor set points

Van Heerden's controller was designed to create only one dynamic compressor pressure set point for the entire compressed air network [1]. This does not allow the controller to consider location. It has the added disadvantage of only allowing certain networks to be controlled. By creating an individual compressor pressure set point, it is possible to control all networks while also being able to split the network while still controlling all sub-networks.

Prioritise compressors dynamically

Because a compressed air network's flow and pressure requirements can change [51], compressors priorities cannot be the same at all times. By prioritising compressors dynamically, demand and supply can be matched at all times regardless if normal schedules are followed or not. This allows for maximum efficiency through matched flow delivery.

Estimate the future state of the compressed air network

By estimating the future state of the compressed air network, it is possible to reduce cycling of compressors and match flow to the future need if the need is set to increase. Van Heerden established the need for future state estimation but the idea of using future set points proved not to work [1]. The DCC will base future state estimations on historically logged demand flow.

Issue start, stop, load and unload commands to all compressors

To control the compressors of the network at maximum efficiency, the DCC must be able to issue commands to the compressors. This ensures that the most efficient possible combination of compressors are running at any moment while also reducing cycling. This requirement is in addition to the requirement to prioritise compressors dynamically.

Control demand valves on the compressed air network

By being able to control demand valves on the compressed air network, the DCC gains the ability to reduce the set point pressure of each user independently. By being able to match the required pressure to the supplied pressure of each end user, the DCC should be able to reduce average ring pressure by reducing wastage of compressed air [104].

Include a reporting ability

A reporting feature is required by the DCC to be able to review its decisions and report on electricity usage. By including the ability to log historical data on the DCC, the user can use a reporting system such as the one developed by Goosen [105].

Include the ability to split networks dynamically

As mentioned by Joubert, Bolt and Van Rensburg, compressed air rings are sometimes split into high pressure and low pressure sub-rings [48]. This allows electrical energy savings by reducing the compressed air set point of one of the rings. The controller developed by Van Heerden cannot control two networks simultaneously and thus cannot control all networks continuously [1]. By being able to handle split networks dynamically, the controller can remain operational on all compressed air networks continuously.

3.3 Simulation process

The DCC simulations are done on the same principle as the simulations done by Van Heerden [1]. The principle of nodes and pipes is used in the same way as Van Heerden proved that the design can be used to create any network dynamically. This is done by forming “AirNodes”, with each AirNode simulating a node with more than one pipe interconnecting with it. A complete network can be built by linking AirNodes.

While this network principle and the use of Bernoulli’s theorem are the same, the simulations differ on key points. Van Heerden’s simulations were based on having a supply pressure and a demand flow [1]. The current network state was also simulated to be able to compute a required input pressure of the network.

The DCC simulations are based on the idea of calculating the ideal network state and not the current network state as was done by Van Heerden [1]. This is done because information about the current state of the network is not required to find the optimum state of the network and only results in unnecessary processing requirements.

The new simulations are done by using only the flow of the input and the output of the network at each network endpoint. Because the compressed air network can potentially leak compressed air, the measured demand flow at each end point is adjusted upwards. This is done because the simulations use flow and also to match supply and demand. Equation 11 is used to adjust the demand flow. The measured flow is obtained from flow meters on site.

$$m_{adjusted} = m_{measured} \bullet \left(1 + \frac{\sum m_{total\ supply}}{\sum m_{total\ demand}}\right) \qquad \text{Equation 11}$$

$m_{adjusted}$: Mass flow after scaling [kg/s].

$m_{measured}$: Mass flow measured by meter [kg/s].

$m_{total\ supply}$: Total measured mass flow at supply [kg/s].

$m_{total\ demand}$: Total measured mass flow at demand [kg/s].

Because a pressure is also required to simulate the network at the correct pressure, the pressure of the driven node is used as well. The driven node is the node that determines the minimum set point pressure of the compressed air ring that will supply each demand node with its own required air pressure.

The simulations of the DCC are done on the same AirNode flow balance as was done by Van Heerden [1]. Each pipe into a node is classified as either a “connecting pipe” or an “end pipe”. Connecting pipes are air pipes that connect two AirNodes, while end pipes connect an AirNode and an end user. Figure 22 shows the same network as in Figure 20 as an example.

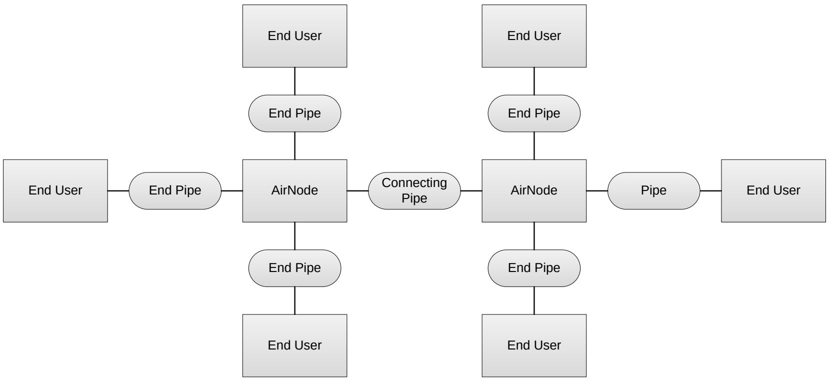

Figure 22: Example network with terms

The flow balance of an AirNode is done on the principle shown in Equation 12– the total sum of airflow through each pipe into the node is equal to the total sum of the airflow through each pipe out of the node. The flow through each pipe is calculated based on what type of pipe it is. The pressures of connecting pipes are known on both sides, while for end pipes, the flow is known but not the pressure at one end.

$$\sum m_{flow\ into\ node} = \sum m_{flow\ out\ of\ node} \qquad \text{Equation 12}$$

m : Mass flow [kg/s].

When the pipe is an end pipe, Equation 13 is used to calculate the end user pressure.

$$P = P_n - \left[\left(\frac{v^2}{2}\right) \bullet \rho \left(f\frac{L}{D} + k\right)\right] \qquad \text{Equation 13}$$

P : Pressure [Pa].
P_n: Pressure of other node [Pa].
v : Fluid velocity [m/s].
ρ : Fluid density [kg/m^3].
f : Darcy friction factor.
L : Measured length of pipe [m].
D : Pipe diameter [m].
k : K-loss factor.

When the pipe is a connecting pipe, the flow through the pipe is calculated by using

Equation 14.

$$m = A \bullet \rho \bullet \left[\sqrt{\frac{2(P_n - P)}{\left[\rho\left(f\frac{L}{D} + k\right)\right]}}\right] \qquad \text{Equation 14}$$

m : Mass flow [kg/s].
A : Area [m^2].
P : Pressure [Pa].
P_n: Pressure of other node[Pa].
v : Fluid velocity [m/s].
ρ : Fluid density [kg/m^3].
f : Darcy friction factor.
L : Measured length of pipe [m].
D : Pipe diameter [m].
k : K-loss factor.

Because the centre pressure of an AirNode is unknown, the centre pressure is estimated. All pipes are then calculated and a flow balance is done. The centre pressure is adjusted until the flow through the node balances. A balanced node is then acquired. This process is done for all nodes. Because the AirNodes change pressure after each flow balance, the process has to be repeated. As was proved by Van Heerden, the simulation converges to a point [1].

If a pipe is connected to a “DrivingNode”, which is the end user node driving the set point pressure of the compressed air network, the end pressure and flow through the pipe are known. This allows the calculation of the AirNode pressure without the need for guessing the pressure. At the start of the simulation, the node with the highest pressure is chosen as the DrivingNode.

After the simulated network values are converged, all values are known for the network. The set point pressures for each of the end users are compared with their calculated pressures. If the calculated pressure is not higher or equal to the set point pressure, a new DrivingNode is chosen and the process is repeated. This is repeated until all end users have pressures that are higher or equal to their individual set point pressures.

After this is done, all values of the optimal state of the network are known. These include, most importantly, the set-point pressures and the input flow of the supply nodes. These two values are used to determine the input pressure of the network as well as the combination of compressors that should be run.

The functionality of controlling on a two-dimensional level is provided in this step. When a compressor is further away from the required demand for compressed air, the compressor would be required to supply compressed air at a higher pressure than a compressor that is closer. This is to compensate for losses in the network. A compressor running at a higher compressed air set point will require more electricity and thus result in an inferior solution.

However, it should finally be noted – as Van Heerden [1] and Venter [49] both mentioned – that simulations of fluid dynamics can expect a 10% accuracy loss due to all the assumptions made before the simulations are done [106].

3.4 Detail component design

3.4.1 System layout

The complete system design is component-based. This allows for flexibility in adapting the control system to any existing compressed air network. Because this is a continuation of the work by Van Heerden, many of the components are based on versions of the components used in DCS by Van Heerden [1].

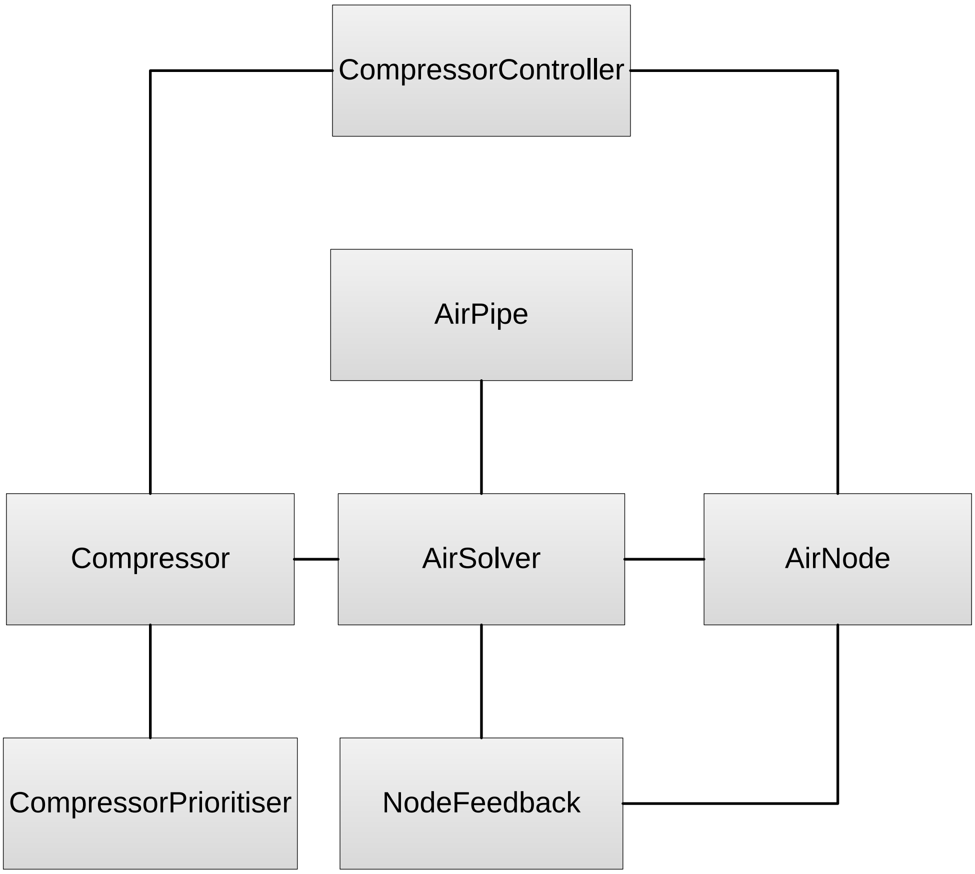

Figure 23: System layout

Figure 23 shows the components as well as the relationship between them. The AirSolver component completes all the simulation work. Since DCC is a network controller, the CompressorController component will translate and transmit all instructions to the local controller to be executed.

All of the communication between the components and the actual network is done via OPC. All the measured data is received via an OPC connection to meters installed on site at various locations depending on the data measured. These meters are calibrated according to OEM specifications. If meter are not installed, they are required to be installed before the DCC can control the network.

3.4.2 AirNode component

General

An AirNode component represents either an end point of the compressed air network or a junction connecting more than one AirPipe component. An AirNode component can be one of three distinct types: Node, Supply or Demand. Figure 24 shows the icons of the three types in the same order as mentioned. “P” in the icon represents pressure while “SP” represents set-point pressure. These values can be displayed in Pa or kPa, and are shown in gauge pressure. The icons are all at minimum width in the figure and will increase their width automatically if required.

Figure 24: AirNode component

When the AirNode is of type Node, the components need no input values. When the AirNode is op type Supply, the component needs to be linked to a CompressorController component. The user can supply the header pressure as well. This will increase the simulation speed as the AirSolver can then use the pressure as a starting point.

Lastly, when the AirNode is of type Demand, the component needs to be supplied with the airflow out of the network, as well as the set-point pressure of the specific end user. The user can also supply the pressure reading of the end user pressure valve. This also allows the AirSolver to increase the simulation speed.

The AirNode component supplies the AirSolver component with the flow, pressure and set-point pressure of each node. The flow out of the network will be the estimated future flow of the end user or the corrected current flow of the end user. The highest flow number between the two will be supplied to the AirSolver. The corrected flow is acquired from the NodeFeedback component and will be discussed in Section 3.4.6.

Future flow

The future flow estimation for the AirNode component is based on historical averages. The AirNode component will log the current flow on a resolution of 30 minutes. This data is then used to compile a historical average flow per day. Each day will be logged independently to allow for the fact that compressed air networks have different profiles each day.

The data is stored in daily csv files and is logged as “total amount logged, count”. The data is logged every two minutes. When the data is required, the data is simply read and the average is calculated. To allow flow to be adjusted over time and to stop the count from reaching large values, the count is limited to 650 readings.

When 650 readings are reached, a new data entry will be created. A new average is created to allow averages to be used continuously. Equation 15 is used to create a new total flow. The new count will be 200. This still allows the average to be used while giving precedence to newer flow readings over time.

$$Total\ flow_{new} = \left(\frac{Total\ flow_{old}}{650}\right) \bullet 200$$

Equation 15

$Total\ flow$: Mass flow [kg/s].

The future flow supplied to the AirSolver will be either the estimated flow in 30 minutes or an hour. The number chosen will depend on the highest number of the two.

Logging

The AirNode component will log the data through it on a two-minute resolution. The data logged will be the following: pressure upstream, pressure downstream, pressure set point, flow and its type. The upstream and downstream pressures refer to the upstream and downstream of the end user valve. The downstream value will be the read value while the upstream valve will be the calculated value.

3.4.3 AirPipe component

General

The AirPipe component represents an air pipe in the compressed air network. The AirPipe component does not require any dynamic variables via OPC. Each pipe has to be individually represented with an AirPipe component. The user can specify the length, diameter, roughness and K-loss of each individual pipe.

The K-loss factor represents the twists, bends and narrowing of each pipe. Figure 25 shows the AirPipe component icon. The flow is indicated in the centre of the icon. The unit is kg/s, which is the SI unit for mass flow. The icon is currently at minimum width and will increase its width automatically when required.

Figure 25: AirPipe component

Locking

The pipe component can be supplied by a variable to indicate if the pipe is locked and thus blocking the airflow through it. This variable can be changed externally. This allows AirSolver to establish which pipes are locked so as not to allow compressed air to flow through them. How the AirSolver handles a locked pipe will be discussed in more detail in Section 3.4.8.

Logging

The data of the AirPipe component is logged on a resolution of two minutes. Because most of the variables of the air pipe are fixed, only the flow and locked status are logged as these are the only variables that change.

3.4.4 Compressor component

General

The Compressor component represents an actual compressor. Figure 26 shows the Compressor component icon and its three main states of operation. A running compressor can either be loaded (represented by the “On” icon) or unloaded (represented by the “Unloaded” icon). The “Off” icon represents a compressor that is off.

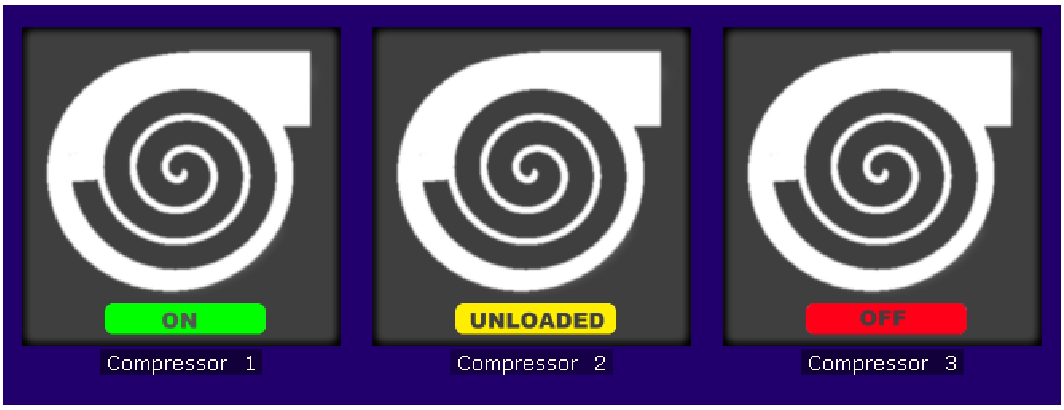

Figure 26: Compressor icons

When a compressor is loaded and running, the compressor is actively supplying compressed air to the network. A compressor in the unloaded state is a compressor that is running with the blow-off valve open and not connected to the network, thus not supplying compressed air to the network. During the off stage, the compressor is switched off and disconnected from the network.

The Compressor component is used to convey instructions to the local controller of each compressor. The component will receive general instructions, for example, to start, stop, load or unload. The component will then translate these instructions to the local controller if possible given certain parameters such as availability and maximum number of stops per day.

Each compressor is also allowed to be a reserve compressor, a baseload compressor or a normal compressor. A reserve compressor is a compressor that should ideally not be run as these compressors usually have problems starting or stopping and are only used when absolutely required. A baseload compressor is usually a large compressor that is always in the on state, as the client requires these compressors to run continuously.

Performance monitoring

Each compressor component will continuously and actively monitor its own performance and log these results for use in the simulations. These values are obtained from meters installed on the compressor. The monitoring is split into performance and efficiency. The performance is logged on a flow range per supplied pressure. The efficiency is logged as the power used to supply a certain flow at a certain pressure.

These results are logged continuously because the operating conditions of a compressor change continuously. Examples of this can be inlet filters that get clogged up, and summer and winter temperatures. When power values are not found in the results, the power is estimated. Marais stated that a 10% reduction in pressure should reduce the electricity consumption by 16% to 18% [47].

For simplification, pressure data is grouped together in 20 kPa intervals while flow data is grouped together in 1 kg/s intervals. Data logged will be rounded to the nearest 20 kPa or 1 kg/s before logged data is saved. This data is used when simulating to estimate the delivery flow and the power usage.

Logging

All the data of each compressor is logged on a resolution of two minutes. The logged data includes, but is not limited to running status, load status, availability, delivery pressure, delivery volume, electricity usage and guide vane position. From the log file, it is possible to retrieve the complete status of each compressor. This data is then used to generate reports for clients [105].

3.4.5 CompressorController component

General

The CompressorController component encapsulates the network controller for each compressor house. The compressor allows the linking of individual Compressor components to it. This allows a user to configure compressor houses dynamically with each CompressorController component representing a house.

Compressor controller
Compressor : -
Action : -
Reason : -
Implement : -
CompressorController 1

Figure 27: CompressorController icon

Control philosophy

The control philosophy is based on the work by Van Heerden [1] and Du Plessis [94]. Figure 28 [1] gives an overview of the control philosophy decision tree. The decision tree will only be used after a delay period has passed. The delay period allows the network to stabilise before a change can be made, thus ensuring that the network is again in its stable state after previous changes have been implemented.

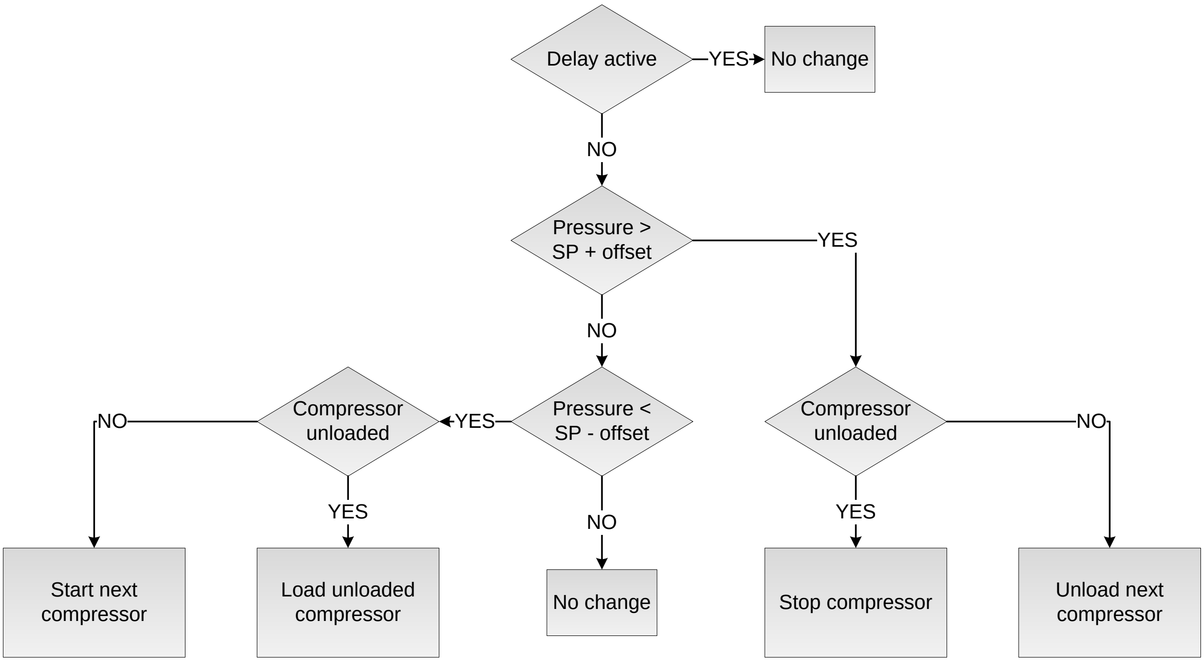

Figure 28: Control philosophy

The philosophy is based on the idea of keeping compressed air network pressure within certain bounds. These bounds are supplied by the control offset. If the pressure rises above the control offset, a compressor will first be unloaded and will then be shut down if needed. When the pressure drops below the control offset, a compressor will be loaded if possible, then one will be started. Figure 29 [94] shows an example.

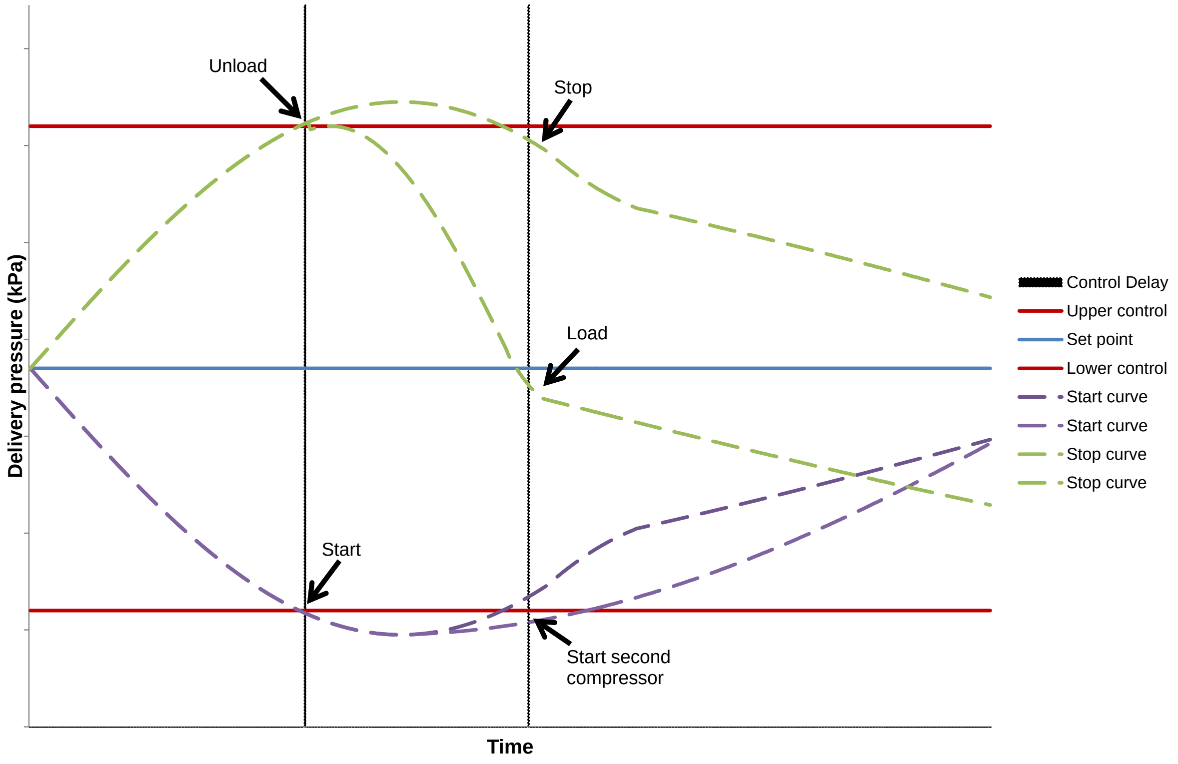

Figure 29: CompressorController offsets

The chosen compressor for all these decisions is based on the priorities of the compressors. When a compressor is started, the lowest priority compressor that is off will be used. When a compressor is stopped, the highest priority compressor that is off will be used.

Logging

The CompressorController component also incorporates historical data logging on a two-minute resolution. This log only logs the current running compressor count. The component also features an action log. This log only logs when the controller takes an action. This log can be used to review decisions made by the controller.

3.4.6 NodeFeedback component

General

The NodeFeedback component is used as a GUI component. The component displays current pressure set point, set point offset, total supply flow, total demand flow, flow loss percentage, driven node, name of the node with the highest set point pressure and the name of the node with the highest demand flow. Figure 30 shows the icon of the component.

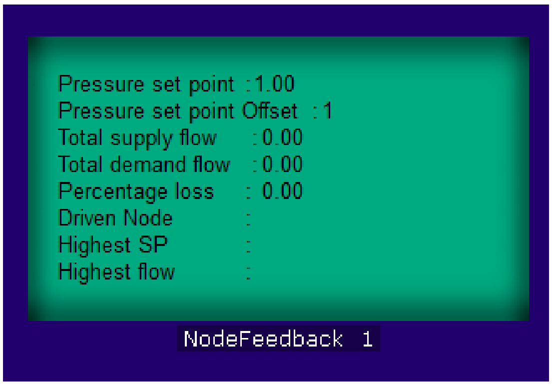

Pressure set point : 1.00
Pressure set point Offset : 1
Total supply flow : 0.00
Total demand flow : 0.00
Percentage loss : 0.00
Driven Node :
Highest SP :
Highest flow :

NodeFeedback 1

Figure 30: NodeFeedback icon

The set point offset allows the user to specify a negative or positive offset to be used to adjust a final set point pressure before writing out to a local compressor controller. The driven node concept will be discussed in Section 3.4.8.

Loss percentage

Venter assumed that flow losses could be neglected [49]. This, however, cannot be assumed as the flow losses may be very high (as is shown in Figure 31). The NodeFeedback component calculates a loss percentage, which is used by the AirNode components to scale the measured airflow from the flow meter.

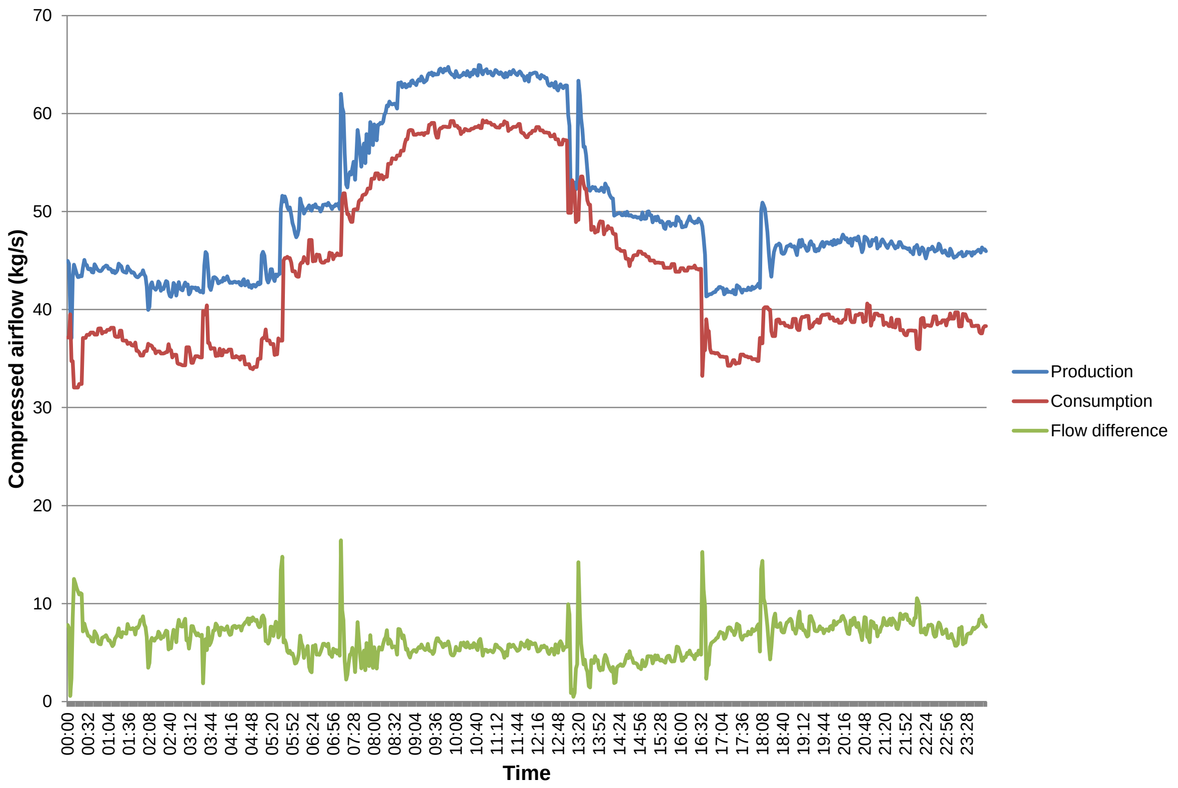

Figure 31: Flow difference

This scaling is done using Equation 16.

$$flow_{scaled} = \left(\frac{total\ flow_{supply}}{total\ flow_{demand}} \right) \bullet flow_{measured\ actual} \qquad \text{Equation 16}$$

$flow$: Mass flow [kg/s].

Logging

The historical logging of the NodeFeedback component is done on the same two-minute resolution as the other components. This historical log includes total supply flow, total demand flow, loss percentage, driven node name, highest set point pressure node name and the highest flow node name. This log allows the user to establish which nodes are responsible for the set point as well as how many compressed air losses the network has.

3.4.7 CompressorPrioritiser component

General

The CompressorPrioritiser component will generate a priority for each compressor component based on the results of the simulation. Compressors can either be combined in one large single priority queue or be combined in multiple priority queues – one for each compressor house. The component is also capable of assigning set point offsets. This component does not include any historical data logging.

Figure 32: CompressorPrioritiser icon

Compressor set point offsets

The offsets for the compressors work by assigning offsets. These offsets are then added to the compressor house set point as is calculated by the AirSolver. This allows each compressor to have an individual set point. Figure 33 shows an example of an offset.

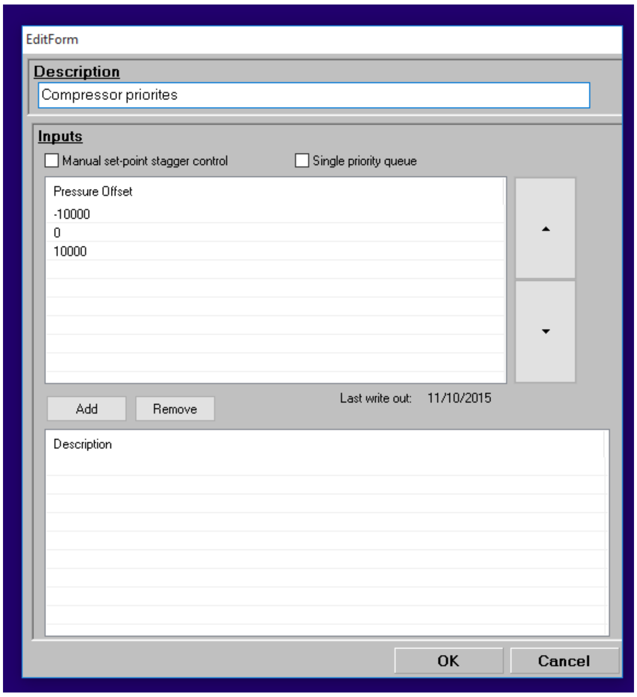

Description: Compressor priorites

☐ Manual set-point stagger control ☐ Single priority queue

Pressure Offset
-10000
0
10000

Last write out: 11/10/2015

Figure 33: CompressorPrioritiser example

In this example, compressors with the highest priority will have their set points reduced by 10 kPa. Compressors with priorities just one lower than the highest priority will have the same set points as their compressor houses have. Lastly, all other compressors will have their set point pressures increased by 10 kPa.

The offsets are used to allow certain compressors to run at full capacity while running other compressors at reduced flows. This is done to prevent all the compressors from running at reduced flow and thus preventing shutting a compressor down.

3.4.8 AirSolver component

General

The AirSolver component completes all simulations for the network. The component will calculate the required set point pressure, optimal compressor selection and their positions on their own control ranges. The AirSolver only features an icon and an Options menu. The icon can be seen in Figure 34. The component does not feature any logging, as all relevant data is logged by other components.

Figure 34: AirSolver icon

Figure 35 shows the simulation layout. The AirSolver compiles a solution that runs on its own thread. Each solution contains a number of individual worker threads. Each worker thread creates its own network to simulate. This network consists of nodes and pipes. These nodes and pipes are not linked to the actual AirNodes and AirPipes to stop concurrency problems [107].

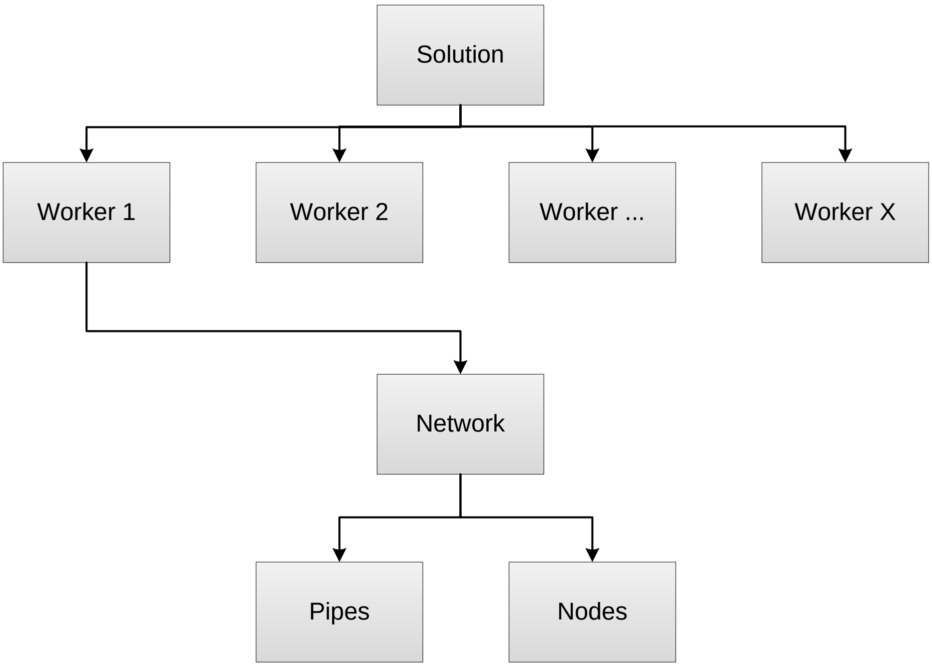

Figure 35: Simulation layout

Solutions

The solutions of the AirSolver are compressor combinations. These solutions are created from adding or removing running compressors. By taking inspiration from genetic algorithms, a solution can be represented by a string of binary characters. Genetic algorithms need to be evolved and compared, but with these simulations, the comparison factor is the power used by the solution. This is only known after simulation, thus a true genetic algorithm cannot be used.

Because each compressor is represented by a binary bit and the number of compressors on a compressed air ring is finite, the number of solutions will also be finite. It will thus be possible to test all possible evolved solutions. A solution tree can then be built from the starting point that contains all possible solutions. Figure 36 shows an example of a solution tree.

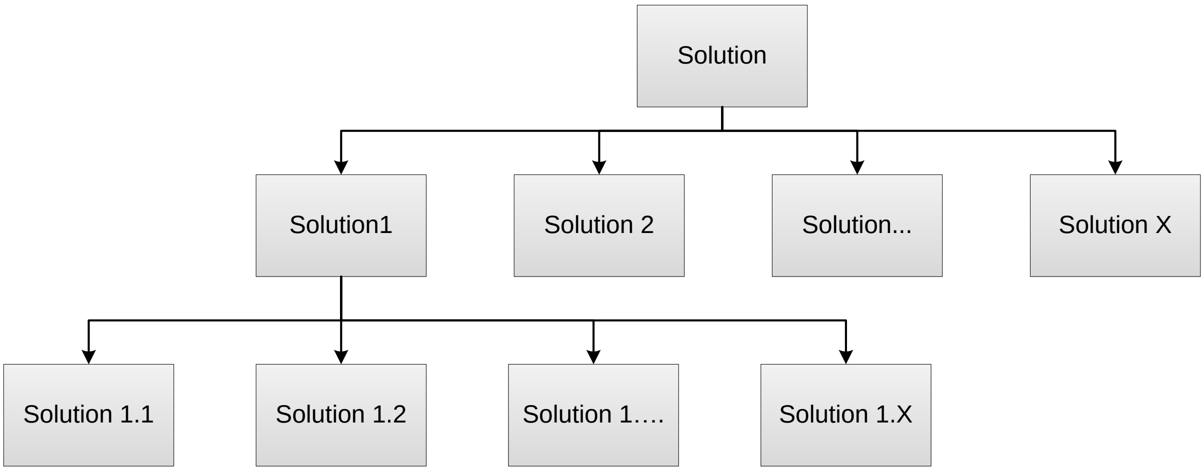

Figure 36: Solution tree

Before simulation, all the solutions are filtered. Some solutions are known to be unusable before a simulation is done. An example of an unusable solution is a solution that does not have enough running compressors that can supply the required flow to all the demand nodes. Solutions are filtered by minimum possible supplied flow being more than the required flow, maximum possible supplied flow being less than the required flow and contain all the baseload compressors.

To reduce the cycling of compressors, the first solution is always the current set of running compressors. This is then evolved in degrees of evolution. A degree of evolution is the number of changes from the main parent solution. This evolution can be either positive or negative. Figure 37 shows one degree of evolution. Each of the compressor states is changed from the parent and a new child is formed. In this figure, 1 is “on” while 0 is “off”.

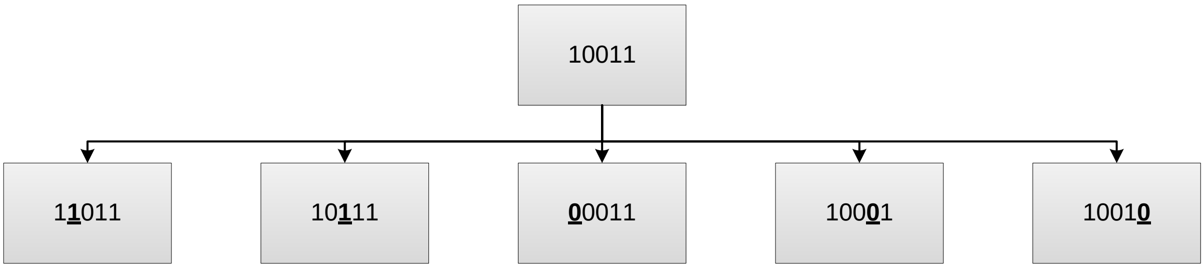

Figure 37: Solution evolution

These solutions are then filtered. If no possible solutions remain, two degrees of evolutions are then performed. This means that each of the child solutions is evolved by one degree. The exact same process as seen in Figure 37 is performed for each child, thus creating a much larger solution tree. Each new child solution now has two changes as opposed to the original parent solution that only has one. The degrees of evolution are increased until no possible solutions remain after filtering. A child is also eliminated if it is a duplicate of a previous child solution.

Worker creation

When the worker threads are created, a finite number of possible solutions remain. Each solution only contains a finite number of running compressors. Each compressor has its own unique efficiency curve. Finding the optimal running point for the compressors presents the same problem as for the solutions – the power usage will only be known after the simulation.

But, with the running compressor flow set points, the number of possible solutions are infinite and it is thus not possible to explore all possible solutions. Thus, the shotgun hill climbing technique is used to find an optimal solution [108]. Each time, only a local maximum is found but only after the simulation is it possible to determine the best local maximums.

Figure 38 shows the flow diagram for the hill climbing process. By doing this process for each worker thread, each worker thread obtains a unique network with each compressor supplying different amounts of compressed air.

Network simulation

Each worker will simulate its own network. The simulation process is described in section 3.3. Because each worker and solution run on an independent thread, the GUI will never be influenced by processor-intensive simulations. Because each worker thread is unique and contains its own resources, this method of simulation capitalises on modern processor design by utilising pipelines and multithreaded architecture. This allows for maximum usage of resources.

To be able to anticipate future demand and not make unnecessary changes, the current running solution is simulated on current flow values, while the other network solutions are simulated with future values. This ensures that the current solution will be kept as is and changes are not made prematurely. However, as soon as changes in the number of running compressors are required, the future state of the network is considered.

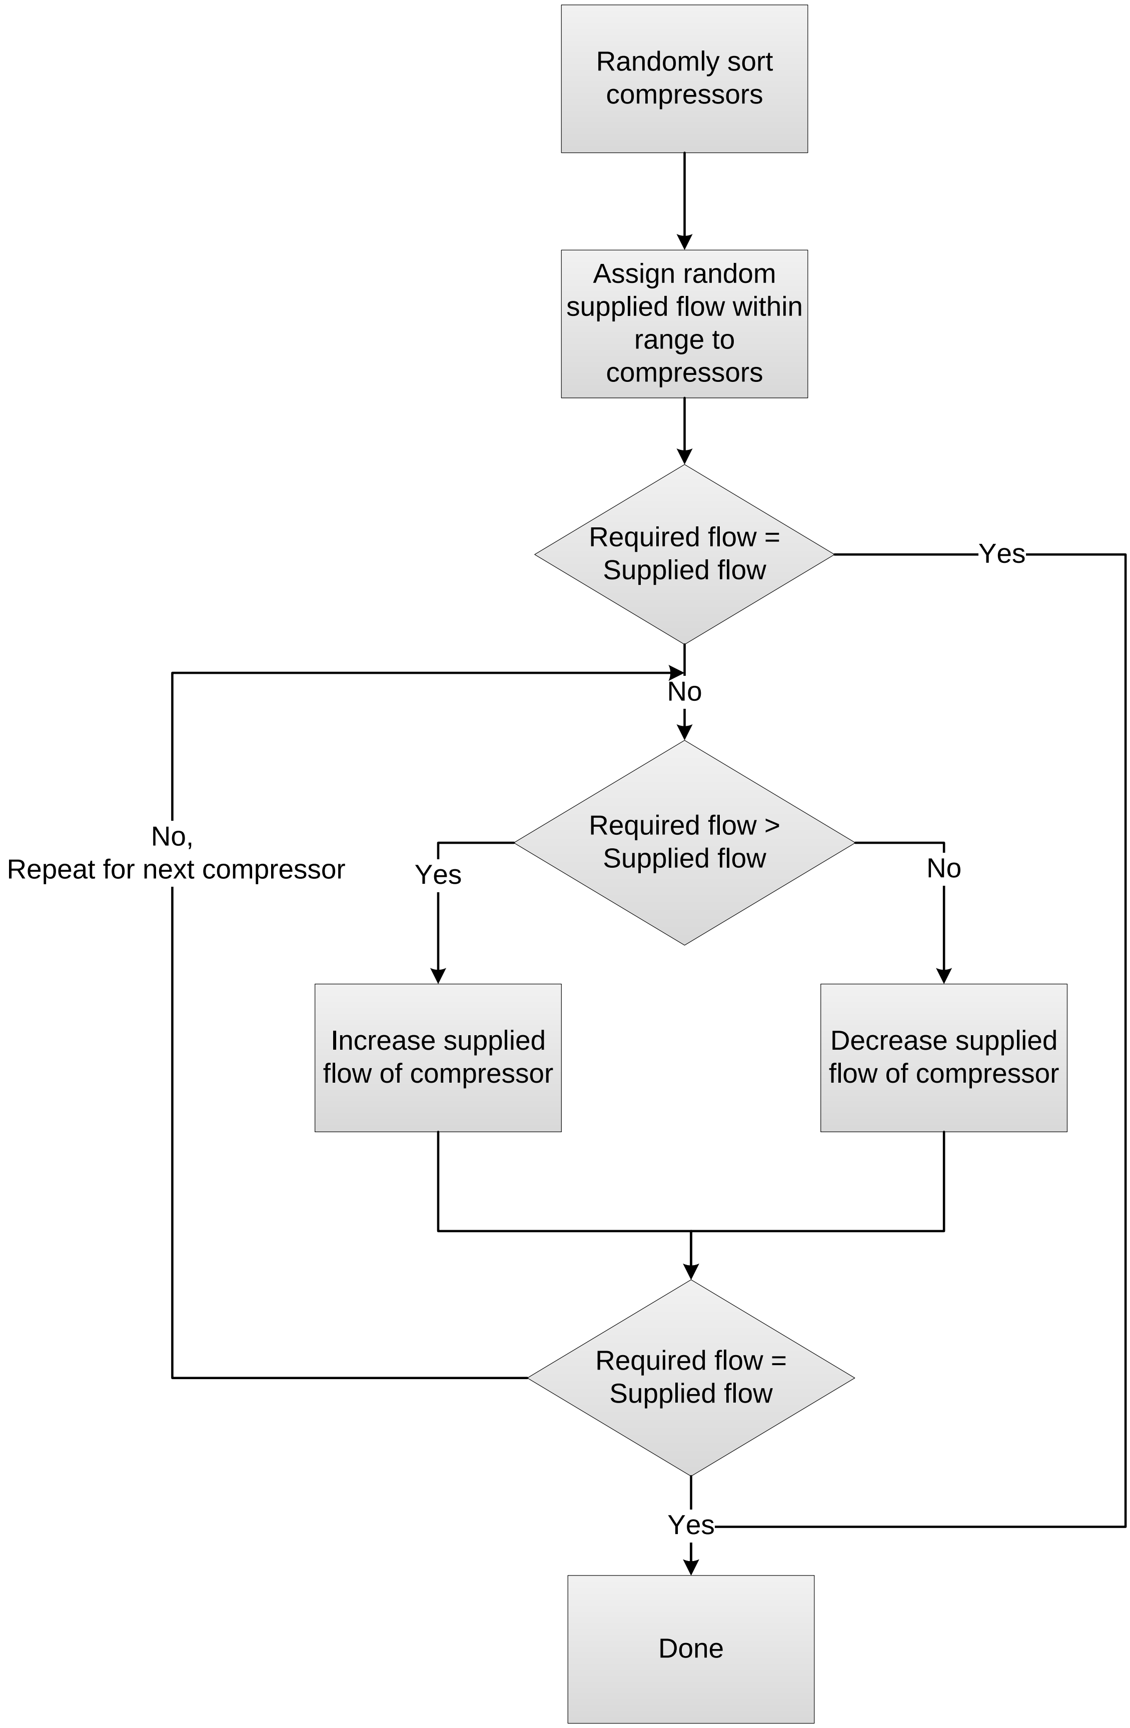

Figure 38: Hill climbing search flow diagram

Driven node

The driven node concept is that the one AirNode will be the reason that the set point has to be as high as it has been calculated. If that AirNode set point pressure is reduced, the compressed air set point can be decreased. This driven node is guessed as the AirNode with the highest set point pressure.

If this is incorrect, not all nodes will have an upstream pressure that is greater than their set point pressure after simulation. If this is the case, a new driven node has to be selected and the simulation process repeated. The second and subsequent guesses for the driven node are the AirNodes with the highest pressure differences between their upstream pressures and set point pressures, given that the nodes have not yet been chosen as driven nodes and that the upstream pressures are less than the set point pressures.

Set-up

Figure 39 shows the set-up window of the AirSolver component. This window allows the user to set up safety barriers for the results and fine-tune the simulations. The lists allow the user to inspect the created network to establish if it was created fault-free. The Max Thread count is the number of worker threads created per solution.

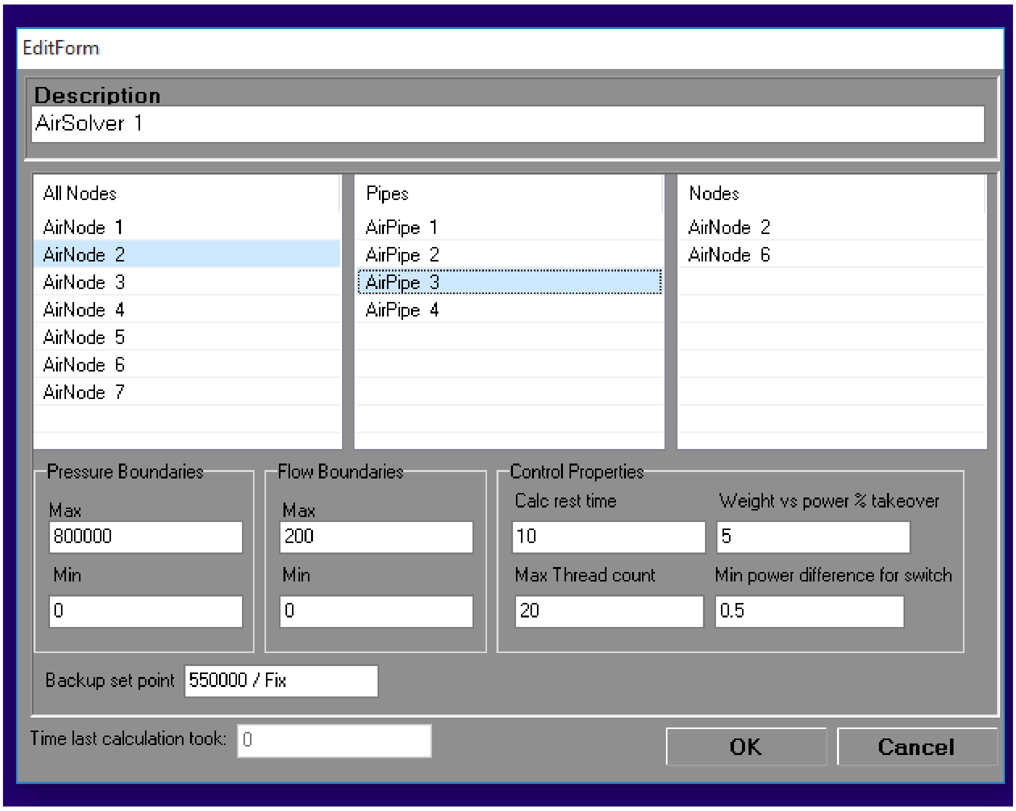

Description: AirSolver 1

All Nodes	Pipes	Nodes
AirNode 1	AirPipe 1	AirNode 2
AirNode 2	AirPipe 2	AirNode 6
AirNode 3	AirPipe 3	
AirNode 4	AirPipe 4	
AirNode 5		
AirNode 6		
AirNode 7		

Pressure Boundaries
Max: 800000
Min: 0

Flow Boundaries
Max: 200
Min: 0

Control Properties
Calc rest time: 10
Weight vs power % takeover: 5
Max Thread count: 20
Min power difference for switch: 0.5

Backup set point: 550000 / Fix

Time last calculation took: 0

Figure 39: AirSolver set-up window

Weight vs power % takeover is used to specify the power percentage bracket two solutions should fall in for the compressor weight to be taken into consideration. Min power difference for switch is used as an anti cycling safety factor, the solution must reducing electricity usage by this constant before being considered. The Backup Set point field is used when the AirSolver fails to find a solution or the input/output boundaries are exceeded as these will inhibit the AirSolver from simulating.

3.5 Theoretical results and design verification

3.5.1. Node balance

The purpose of this test was to verify if a simulation could do a flow balance on an AirNode. The simulation was done on the network in Figure 40. The node represented by the x is supply node. Only the minimum required pressure for each demand node and the flow out of the network to the demand nodes are known.

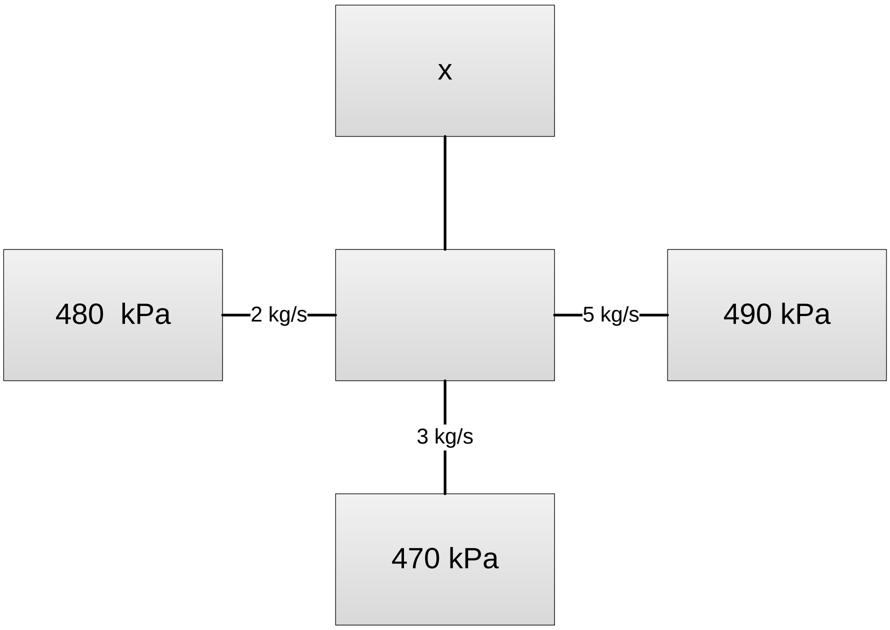

Figure 40: Test node

On this node, the driven node will be chosen as the AirNode with a pressure of 490 kPa. Figure 41 shows the results after simulation. As can be seen in the figure, a minimum supply pressure and an upstream pressure for each of the demand nodes were calculated. AirNode 4 in Figure 41 can be identified as the driven node because it has the lowest upstream pressure. Its upstream pressure is also equal to its required set point pressure.

Because this test only includes a single AirNode, there are no iterations required. The pressure for AirNode 4 and the flow through all the AirPipes are known, thus the pressure of the AirNode and the other nodes can easily be calculated using straightforward equations.

For simplification reasons, all the pipes in this test were seen as identical. They were all 0.6 m in diameter, had a length of 1 000 m and an internal roughness factor of 45 µm.

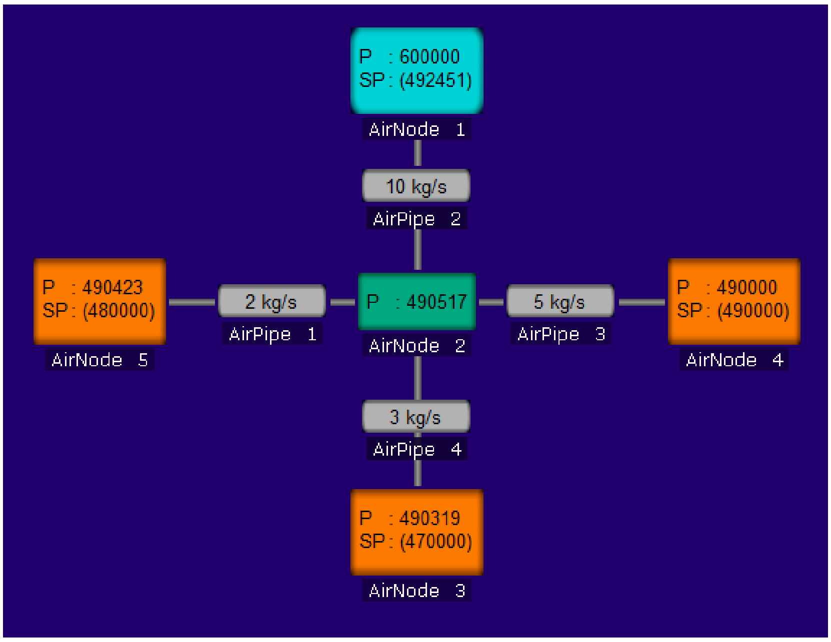

Figure 41: Completed test node

A second test was done with a larger network. This test forced the simulation to run iterations between the AirNodes, because if an AirNode is not connected to a driven node, the pressure of the AirNode cannot be computed and it has to be iterated. The applicable network can be seen in Figure 42.

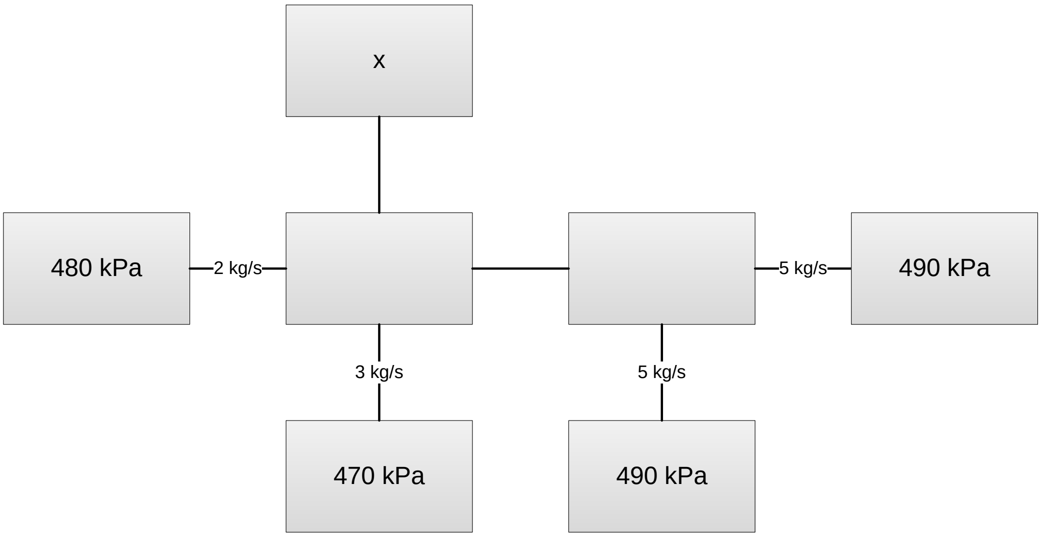

Figure 42: Test network

The results of the simulation can be seen in Figure 43. On this network, all pipes were 0.6 m in diameter, 1 000 m in length and had an internal friction of 45 µm, except for AirPipe 6, which had a length of 1 500 m, to help so how the driven node works. As can be seen from the results, this caused AirNode 7 to be driven node although AirNode 7 and AirNode 4 had the exact same flow and set point pressure requirements.

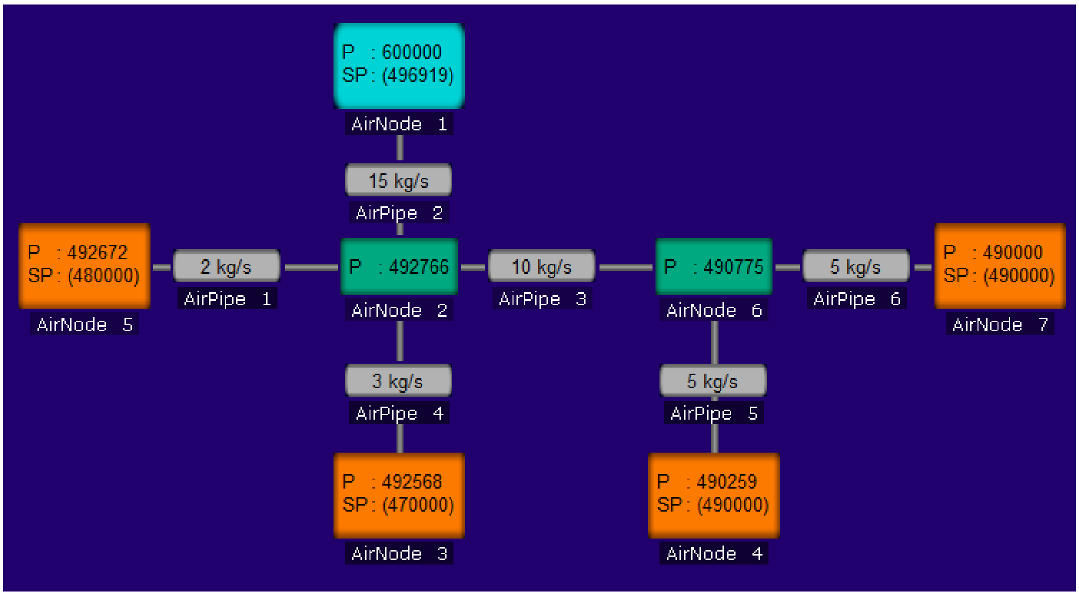

Figure 43: Completed test network

Figure 44 shows the centre pressure of AirNode 2 as it is iterated. The pressure is always iterated positively first. This is because the difference between the maximum pressure and the correct pressure is always less than 0 Pa pressure is to the correct pressure. The AirNode reached a convergence twice.

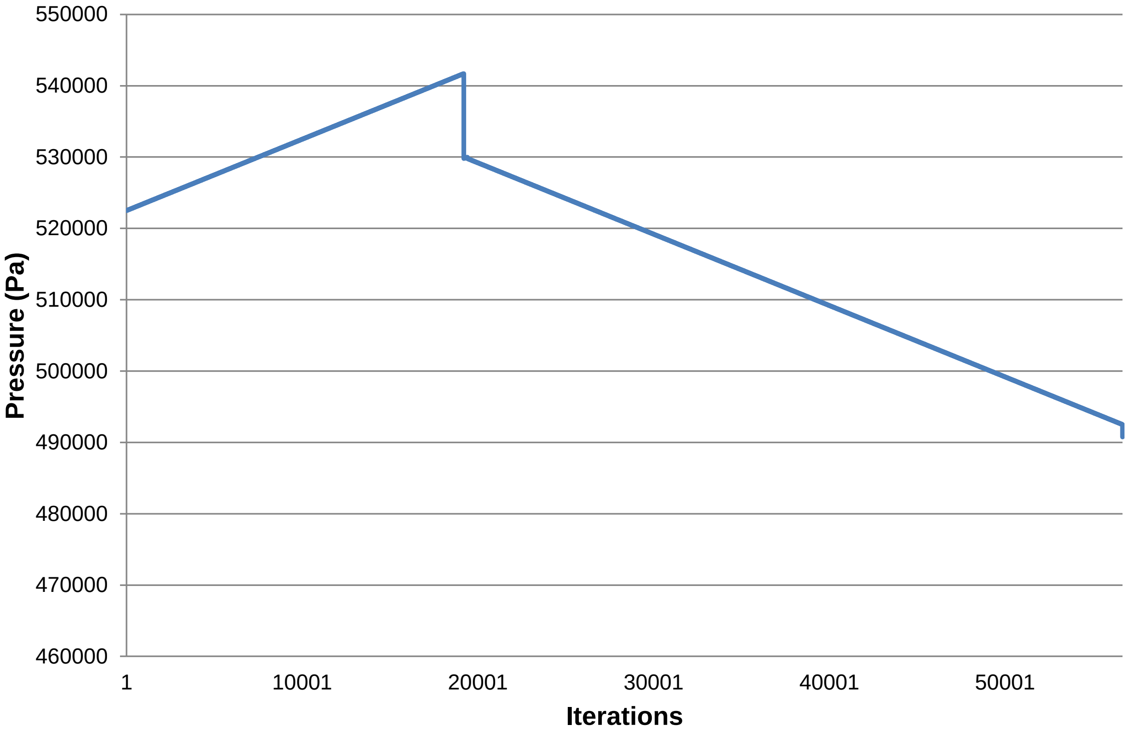

Figure 44: Pressure convergence

From Figure 45 it can be seen how the flow converges to 0. As the centre pressure of the AirNode reaches its true value, the flow difference between the pipes using Equation 12 converges into 0. Figure 45 also shows two iterations. The simulation on the test network in Figure 42 required two iterations of the nodes as can be seen in Figure 44 and Figure 45.

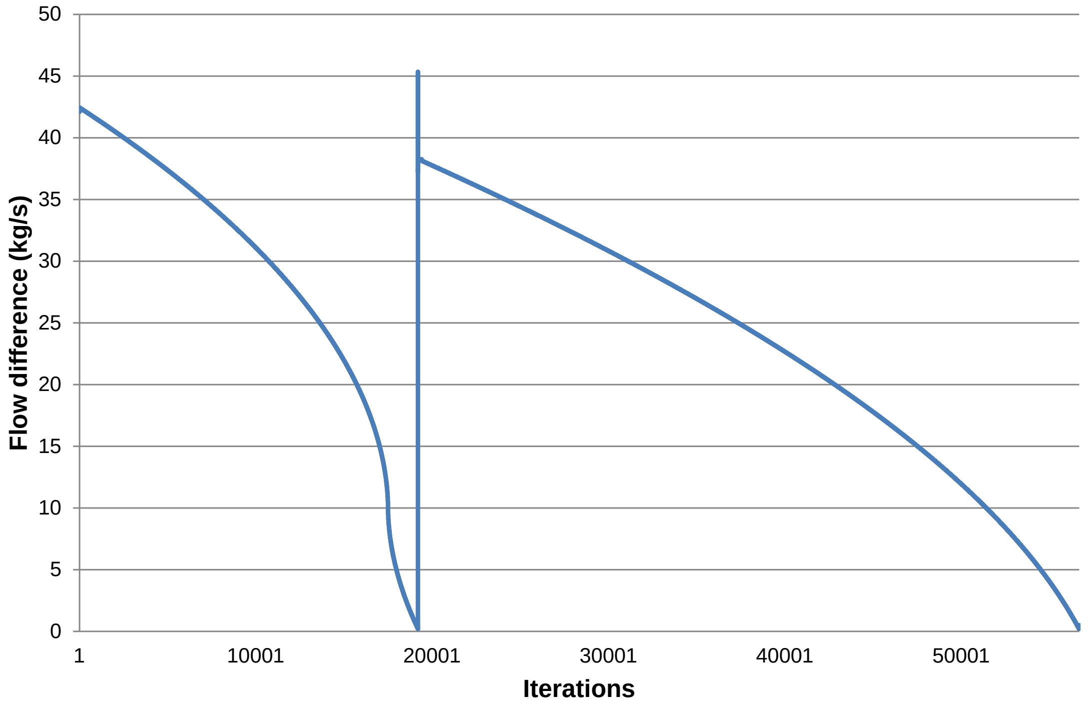

Figure 45: Flow difference

Figure 46 shows the trend of the pressure versus flow. From this graph the direct correlation between the pressure and the flow converging to 0 can be seen. Each time the flow difference reaches 0 the pressure stops.

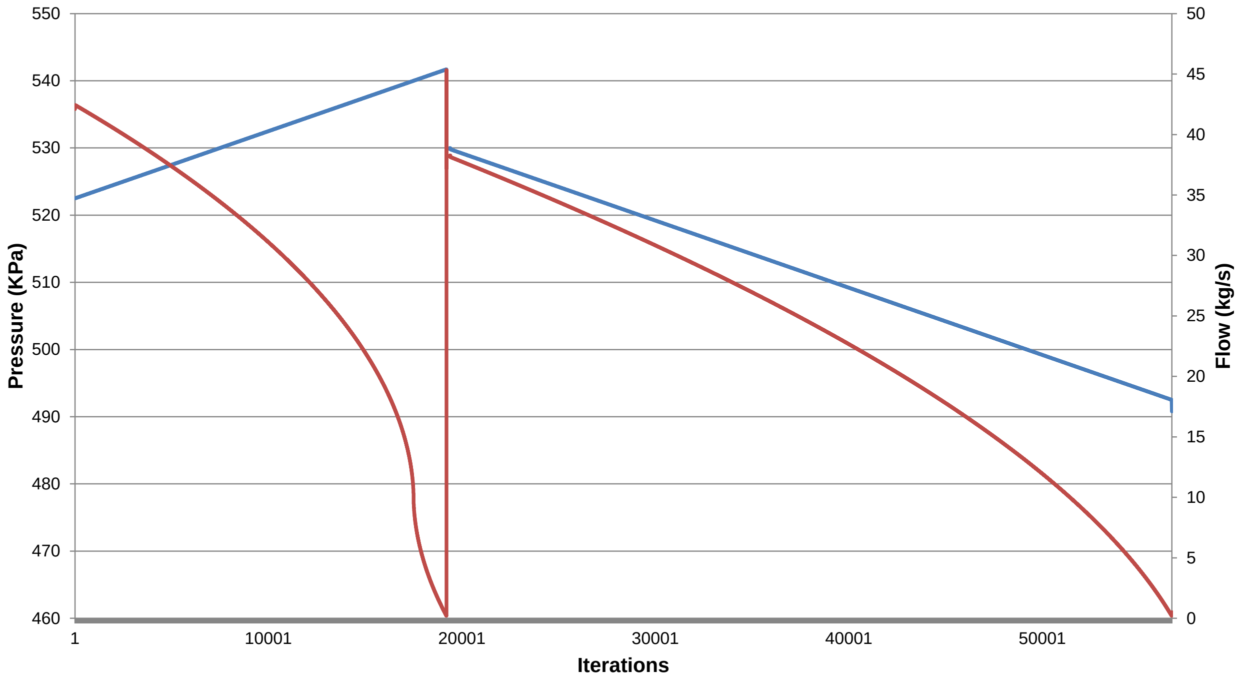

Figure 46: Pressure versus flow trend

3.5.2. Network simulation

The test below was done to calculate the accuracy of the simulation against the accuracy of the DCS [1] and KYPipe. The tested network can be seen in Figure 47 that shows the completed simulated network done in KYPipe. Van Heerden favourably compared DCS to KYPipe [1]. The same simulated network in DCS can be seen in Figure 48 [1].

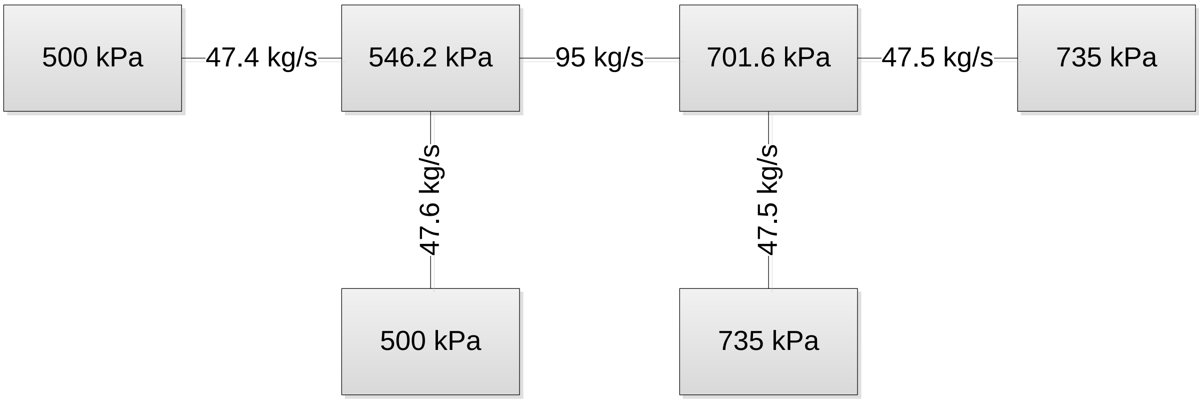

Figure 47: KYPipe simulation results

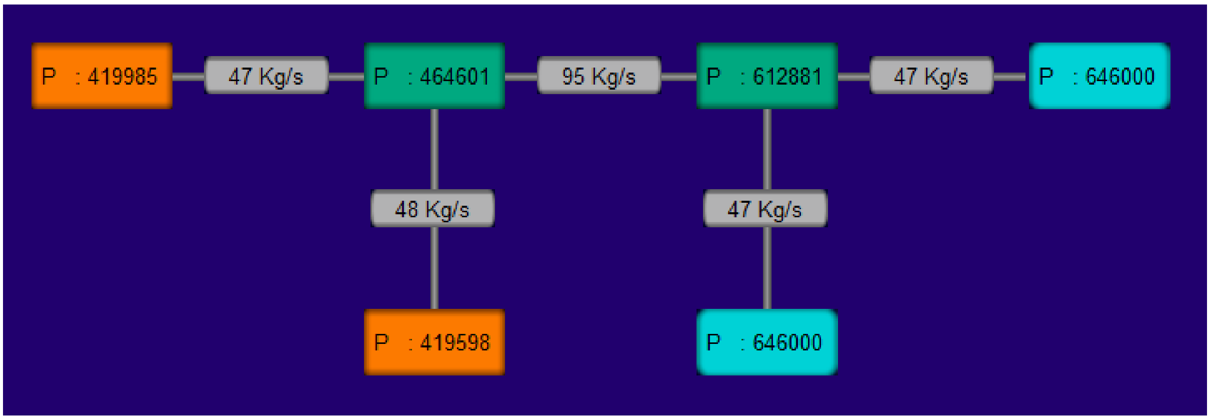

Figure 48: DCS simulation results

This network was recreated in the DCS. The results will differ somewhat as in the DCC it is not possible to stipulate the specified supplied pressure from a compressor house. The controller will automatically choose a supply flow and pressure according to the results of the hill climbing search [108]. The simulated DCC network can be seen in Figure 49.

Both the DCS and the DCC work on gauge pressure where KYPipe works on actual pressure. Gauge pressure is defined as the pressure above atmospheric pressure. On the test results of the DCS and the DCC, atmospheric pressure was defined as 89 kPa. For a comparison to be performed between DCC, DCS and KYPipe, the pressure readings of the DCC and the DCS had to be converted to actual pressure.

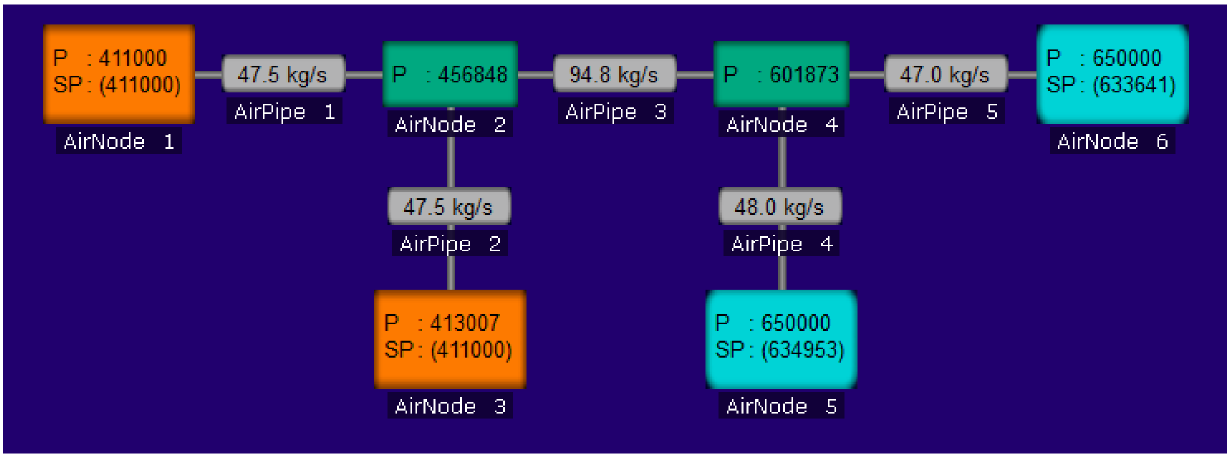

Figure 49: DCC simulation results

Table 2 and Table 3 show the results of the simulations on all three systems. The differences between the systems are at most 1.7%. The standard deviation was below 1%. The simulation results show that the DCC simulations are accurate and comparable to that of other systems.

Table 2: DCC simulation pressure accuracy comparison

KYPipe value	**DCS value**	**DCC value**	**Difference: K-pipe vs DCC**	**Similarity: K-pipe vs DCC**	**Difference: DCS vs DCC**	**Similarity: DCS vs DCC**
500.0 kPa	508.8 kPa	500.0 kPa	0.0 kPa	100.00%	8.8 kPa	98.27%
500.0 kPa	508.4 kPa	500.0 kPa	0.0 kPa	100.00%	8.4 kPa	98.35%
546.2 kPa	553.4 kPa	545.8 kPa	0.4 kPa	99.93%	7.6 kPa	98.63%
701.6 kPa	701.8 kPa	690.9 kPa	10.7 kPa	98.47%	10.9 kPa	98.45%
735.0 kPa	735.0 kPa	724.0 kPa	11.0 kPa	98.50%	11.0 kPa	98.50%
735.0 kPa	735.0 kPa	722.6 kPa	12.4 kPa	98.31%	12.4 kPa	98.31%
Average difference			5.75 kPa	99.20%	5.75 kPa	99.20%
Standard deviation			5.64 kPa	0.78%	5.64 kPa	0.78%

Table 3: DCC simulation flow accuracy comparison

KYPipe value	**DCS value**	**DCC value**	**Difference: K-pipe vs DCC**	**Similarity: K-pipe vs DCC**	**Difference: DCS vs DCC**	**Similarity: DCS vs DCC**
47.5 kg/s	47.4 kg/s	47.5 kg/s	0.0 kg/s	100.00%	0.1 kg/s	99.79%
47.5 kg/s	47.6 kg/s	47.5 kg/s	0.5 kg/s	98.95%	0.6 kg/s	98.74%
95.0 kg/s	94.8 kg/s	94.8 kg/s	0.2 kg/s	99.79%	0.0 kg/s	100.00%
47.5 kg/s	47.5 kg/s	48.0 kg/s	0.5 kg/s	98.95%	0.0	98.95%
47.5 kg/s	47.5 kg/s	47.0 kg/s	0.5 kg/s	98.95%	0.5 kg/s	98.95%
Average difference			0.34 kg/s	99.33%	0.34 kg/s	99.33%
Standard deviation			0.2 kg/s	0.47%	0.2 kg/s	0.47%

An actual mine network was simulated in Figure 50. The pressure set points and the flow requirements of shafts were taken from the actual site. From Figure 50, it can be seen that all the flows, pressure and pressure set points were calculated. The simulation used the actual data obtained from site logging to choose the running compressors to supply the airflow from each compressor house. The flow from the one compressor house is 0 because the compressor house is not actively suppling compressed air.

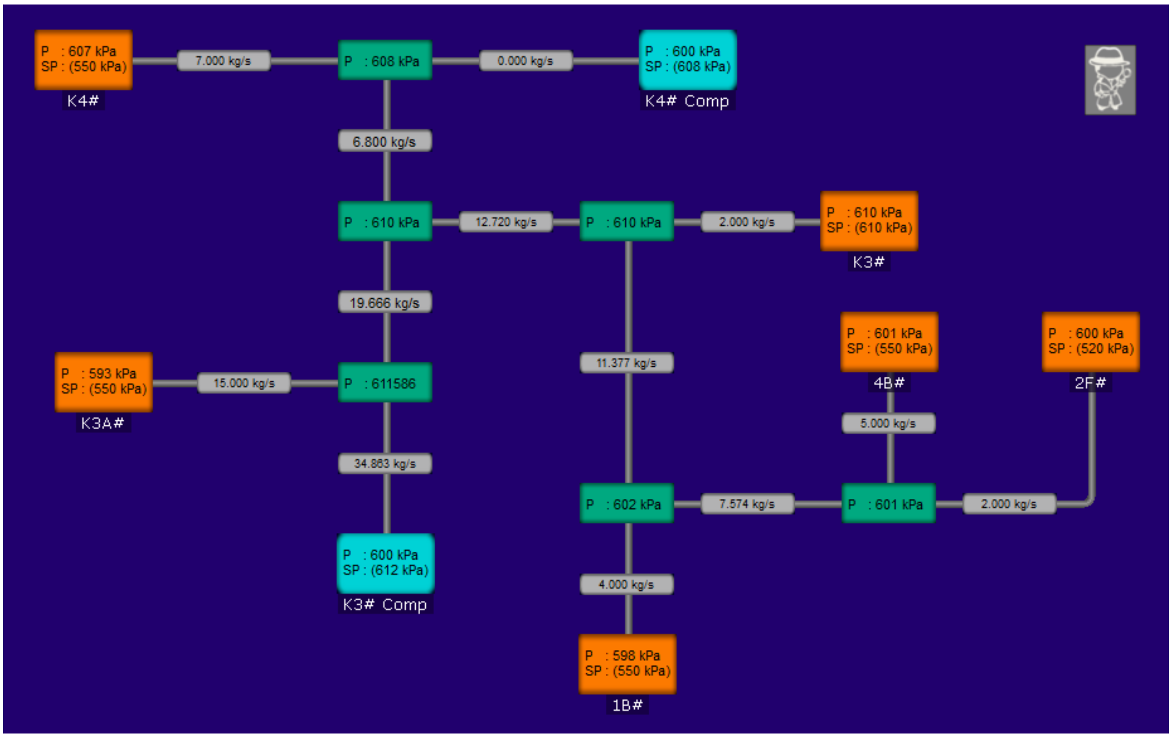

Figure 50: Actual mine network simulation

The test in Figure 50 was repeated, but the simulation was modified by locking the one air pipe. This forced the controller to choose the compressors in such a way as to supply airflow to each of the separate networks formed by the split. The result of this test can be seen in Figure 51. From the figure, it can be seen how the flow through one of the air pipes was locked at 0.1 kg/s. This was done to simulate small losses through the locked air pipe.

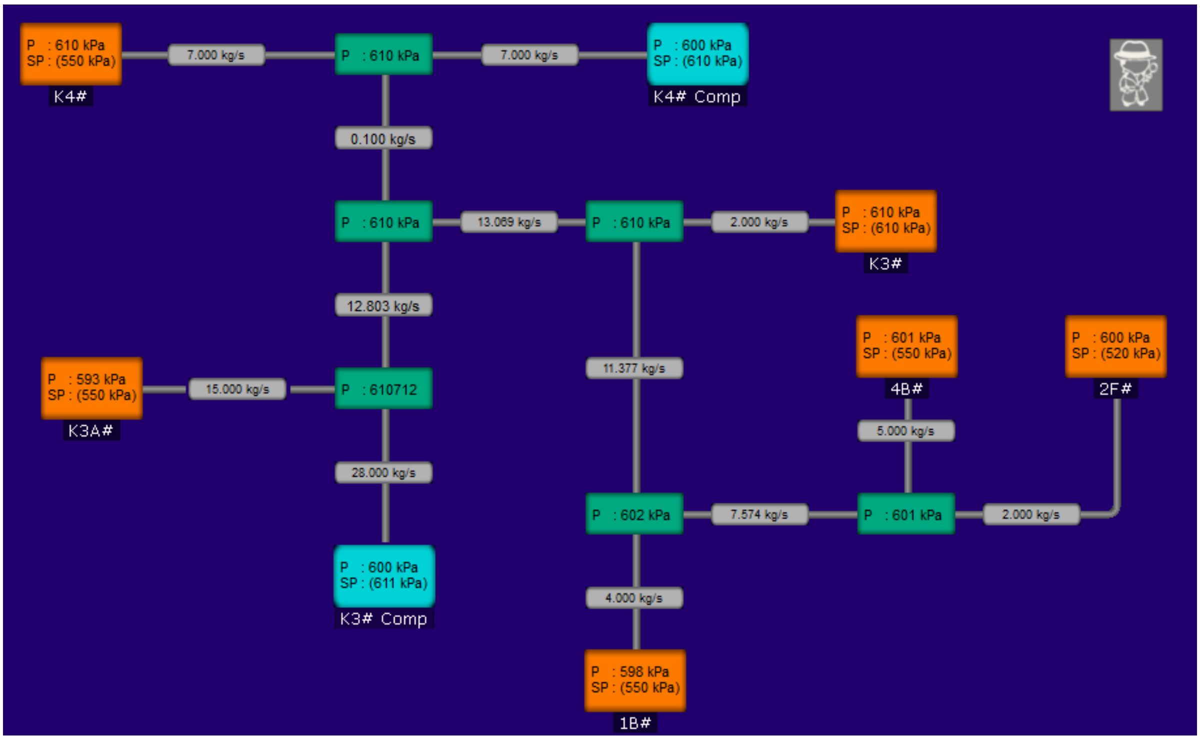

Figure 51: Actual mine network simulation with locked pipe

Due to the use of multithreading in the simulation, calculation speed is faster than calculation speed of the DCS [1]. Figure 52 shows a screenshot of Windows® Task Manager after a simulation was completed. In the screenshot, it can be seen how the usages of all four threads increase to 100%. The usage drops again after the simulation was completed. This test was performed on an Intel® Core ™ i5 laptop processor.

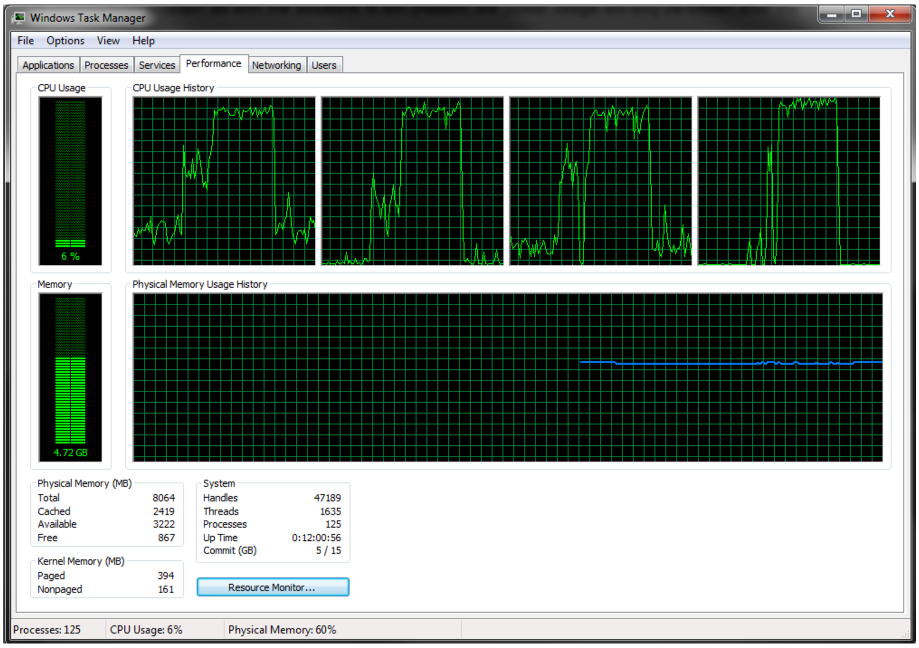

Physical Memory (MB)	
Total	8064
Cached	2419
Available	3222
Free	867

Kernel Memory (MB)	
Paged	394
Nonpaged	161

System	
Handles	47189
Threads	1635
Processes	125
Up Time	0:12:00:56
Commit (GB)	5 / 15

Processes: 125 | CPU Usage: 6% | Physical Memory: 60%

Figure 52: Central processing unit (CPU) usage – laptop

Figure 53 shows the simulation running on a server with more threads. It can be seen that it is still using all available threads, thus increasing simulation speed. It can thus be said that the DCC successfully scaled to the processor design.

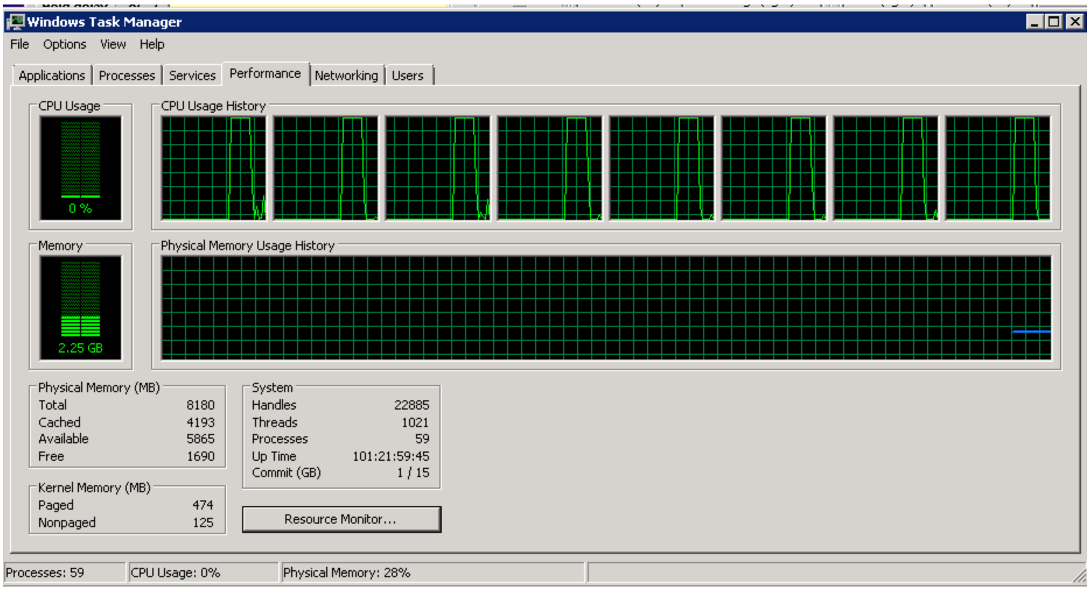

Physical Memory (MB)	
Total	8180
Cached	4193
Available	5865
Free	1690

Kernel Memory (MB)	
Paged	474
Nonpaged	125

System	
Handles	22885
Threads	1021
Processes	59
Up Time	101:21:59:45
Commit (GB)	1 / 15

Processes: 59 CPU Usage: 0% Physical Memory: 28%

Figure 53: CPU usage – server

The test in Figure 54 was done to confirm if the DCC would consider the location of the compressor and select the nearest compressor. In this test, the compressors located at AirNode 5 and AirNode 6 are identical when considering their output and power consumption. All the compressed air pipes in this network had a diameter of 0.6 m and a length of 1 000 m, except the pipe connecting AirNode 5 and AirNode 4. Its length was increased to 10 000 m. As can be seen from Figure 54, the DCC selected the closest compressor correctly.

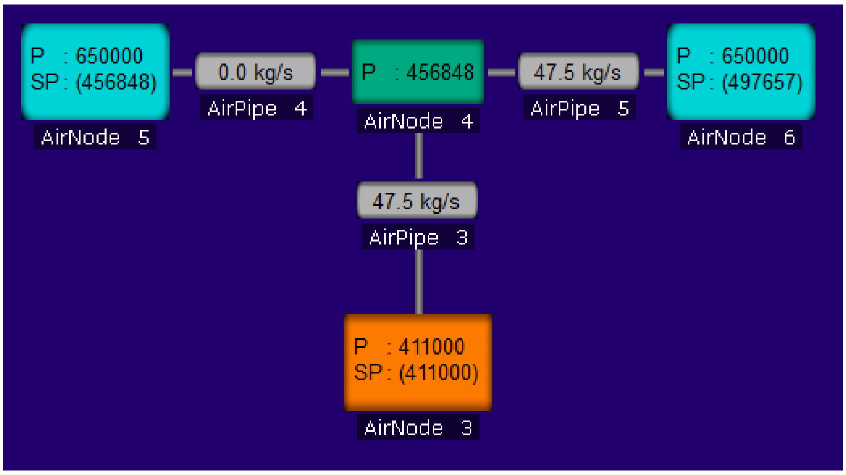

Figure 54: Location test

3.5.3. Design verification

Table 4 lists a short summary of the design requirements (Section 3.2) and whether the requirements were met or not.

Table 4: Design requirements

Requirements	**Result**
Control any industrial compressed air network	Passed
OPC communication	Passed
Dynamic individual compressor set points	Passed
Prioritise compressors dynamically	Passed
Estimate future state of the network	Passed
Issue start, stop, load and unload commands	Passed
Control demand valves on the network	Passed
Logging ability	Passed
Split networks dynamically	Passed
Include compressor location in simulations	Passed
Compensate for air leakage of network	Passed
Include the ability to limit user access control	Passed
GUI	Passed

As the DCC was built upon the DCS, various components were carried over. The OPC communication was achieved by using the same external library component, which is dOPC made by Kassl GmbH. Figure 55 shows the results of the dynamic compressor prioritisation. Because the DCS was built upon REMS-OAN, the same control for demand control valves is also present in the DCC system.

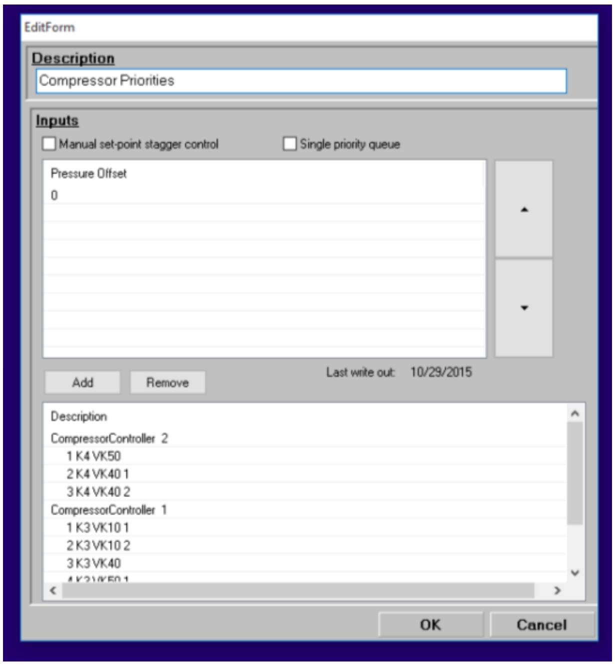

Description
Compressor Priorities

Inputs
☐ Manual set-point stagger control
☐ Single priority queue

Pressure Offset
0

Last write out: 10/29/2015

Description
CompressorController 2
1 K4 VK50
2 K4 VK40 1
3 K4 VK40 2
CompressorController 1
1 K3 VK10 1
2 K3 VK10 2
3 K3 VK40

Figure 55: Dynamic compressor prioritisation

Compressor locations are considered when simulations are done because compressors that are further away from the required demand would require a higher pressure output due to the increased distance. This causes the simulation to rate the solution with a higher energy consumption and as a less efficient solution.

DCC uses the same user access as DCS. Table 5 [1] summarises the user access rights.

Table 5: The DCC and the DCS user access rights

Action	**Viewer**	**Operator**	**Supervisor**	**Administrator**
Connect/reconnect OPC	Yes	Yes	Yes	Yes
Log in	Yes	Yes	Yes	Yes
Switch modes (auto/manual)	No	Yes	Yes	Yes
Save	No	Yes	Yes	Yes
Backup	No	Yes	Yes	Yes
Change component settings	No	No	Yes	Yes
User manager	No	No	No	Yes
Contacts	No	No	No	Yes
Alarms	No	No	No	Yes
Options	No	No	No	Yes
Tags	No	No	No	Yes
OPC options	No	No	No	Yes
Idle/edit	No	No	No	Yes

3.6 Summary

The design requirements for the DCC system together with the reasoning behind the resulting design were discussed in this chapter. A basic overview of the design components used in the design of the DCC can be viewed in Appendix A: Software procedures.

The inner workings of the simulations and the formulas used were discussed in detail in Section 3.3. Each component as well as the functions of each component were examined in detail in Section 3.4 as well as the functions of each component. Lastly, the theoretical results were discussed in Section 3.5.

4 Implementation and results

4.1 Preamble

This chapter focusses on the implementation of the DCC. As part of the previous chapter the system was tested in a local and isolated environment. This chapter will discuss details surrounding the practical implementation of the DCC on actual compressed air networks. Prerequisites regarding the DCC will be mentioned as well as their cost the implementation. Two case studies and their results will be discussed as well.

4.2 Requirements for potential installations

4.2.1. Introduction

During this subsection, details and requirements for a successfully implemented DCC project will be discussed. These are separated into three main sections: Control, Network and Compressor. The payback period for an implementation of a DCC will be determined by the level of current control implemented on the network as well as the existing hardware on the network.

A network with more advanced control would potentially see reduced energy savings but would require less automation and hardware and thus have a reduced initial installation cost. The opposite of this is also true.

Increasing the number of compressor starts and stoppages increases the amount of wear on the compressors. According to Van Tonder, the increase in maintenance is negligible when compared to the energy savings obtained [35].

4.2.2. Control

The DCC was designed as a network controller and thus cannot directly control a compressor – it can only issue commands. For the DCC to function, it requires a SCADA or data acquisition (DA) system to translate instructions to a local controller for each compressor.

Figure 56 shows the relationship between the DCC, SCADA system, local controller and the compressors. The relationship between a meter, SCADA and the DCC will be the same as the relationship between the DCC, SCADA, local controller and compressor. The same is true for all the component values read by the DCC.

The communication between the DCC and the SCADA system will be supplied via an OPC connection. This is done to ease the installation of the DCC and ensure broad compatibility with many SCADA systems. The rest of the communication is done with propriety technology and falls outside the scope of the thesis.

Each end user of compressed air requires a pressure profile for the DCC to be able to calculate the most efficient control schedule for the network. An example of such a profile can be seen in Figure 57 [52]. The profile will depend on how and when the end user consumes compressed air. These will be customised for each individual end user.

The DCC controller requires a REMS license because it was designed for the REMS platform. A hardware server and a Microsoft Windows® operating system licence are also required in order for the DCC to be able to function.

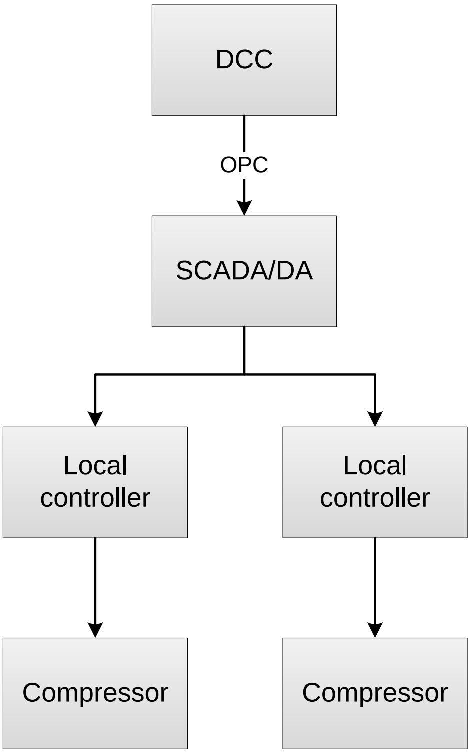

Figure 56: Control layout

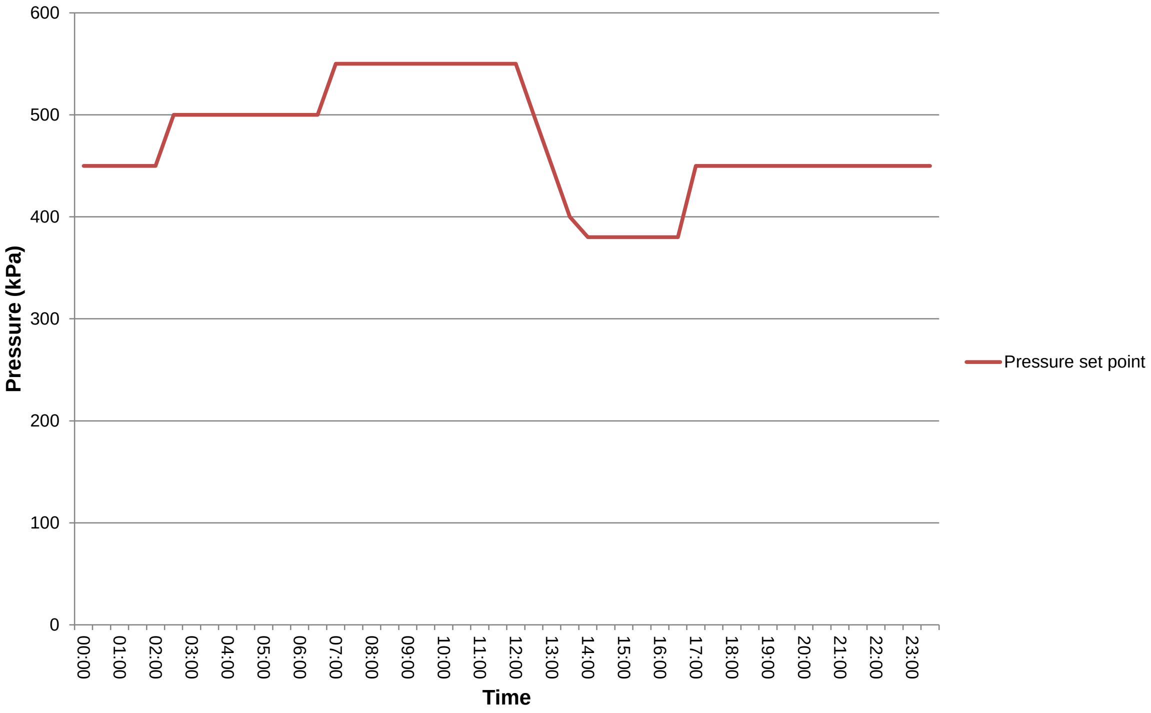

Figure 57: End user pressure profile example

4.2.3. Compressed air network

On the network side, the DCC controller requires a complete network layout. Satellite imagery is used to map the entire compressed air network layout. The satellite map will be used to calculate estimated distances of each of the compressed air pipes. The width, type and K-loss factor of the each pipe are also required for accurate simulation. K-loss factor value is dependent on the number of turns, narrowing and obstructions on and in the pipe.

Each end user on the network also requires a control value that can limit the compressed air pressure to the end user. An example of such a valve is pictured in Figure 58 [49]. This valve also includes a bypass line. A bypass line is used during low flow to restrict wear on the valve in the main line. Each of these control valves also needs to be fitted with a flow and pressure sensor. The valve limits the amount of airflow to match a preset pressure set point to the pressure sensor downstream.

Figure 58: Control valve with bypass line

4.2.4. Compressors

Each compressor requires a local controller as mentioned in Section 4.2.2. This controller must be able to control the compressor via pressure inputs from the DCC. Basic automation needs to include the ability to start and stop the compressors remotely from a centralised control location.

To be able to supply correct commands to each local controller the DCC requires measured data as OPC tags of the following in real time:

- Power usage
- Discharged pressure
- Discharged airflow
- Compressor guide-vane angle
- Compressor discharge valve position
- Compressor running status (on/off)
- Compressor loaded status (loaded/unloaded)

The local controller will in turn interpret the command signals from the DCC and ensure that each compressor runs as the DCC desires.

4.3 Case Study 1: Platinum mine

4.3.1. Compressed air network

The compressed air network of the platinum mine used as Case Study 1 operates nine compressors with a total peak electrical capacity of 34 MW. These compressors are all shown in Table 6 with their flow and power ratings. The compressors are smaller than compressors at other mines thus allowing multiple different combinations of compressors. It also allows for the compressed air delivery to be controlled more precisely.

Table 6: Platinum mine compressors

Name	Flow rating (kg/s)	Rated capacity (MW)	Location
C2_VK50_1	11.3–14.9	5	C2 compressor house
C2_VK50_2	10.4–15.7	5	C2 compressor house
C2_VK50_3	9.8–15.4	5	C2 compressor house
C1_VK50	10.0–15.0	5	C1 compressor house
C2_VK40	6.4–9.1	4	C2 compressor house
C1_VK40_1	6.5–9.8	4	C1 compressor house
C1_VK40_2	6.7–9.8	4	C1 compressor house
C2_VK10_1	2.2–2.3	1	C2 compressor house
C2_VK10_2	2.1–2.2	1	C2 compressor house

The case study mine is the same mine as was used by Van Heerden [1] and Van Niekerk [52] to test and implement the DCS. This means that the control should have already been optimised and the electrical energy saving should be small because the control was adjusted following input from the DCS. But, the payback period should be very low, as there is no hardware required and only a software upgrade must be done.

Because the mine’s shaft set point pressure was out of date and inaccurate, real set points could not be used as part of the test. All the tests were performed using the actual measured pressures at the shafts and it was assumed that these were the requirement pressures. From Figure 59 the actual pressures and required set point pressures can be compared. It can be seen that the actual pressures are below the requirements, thereby indicating outdated and incorrect requirement set point pressures.

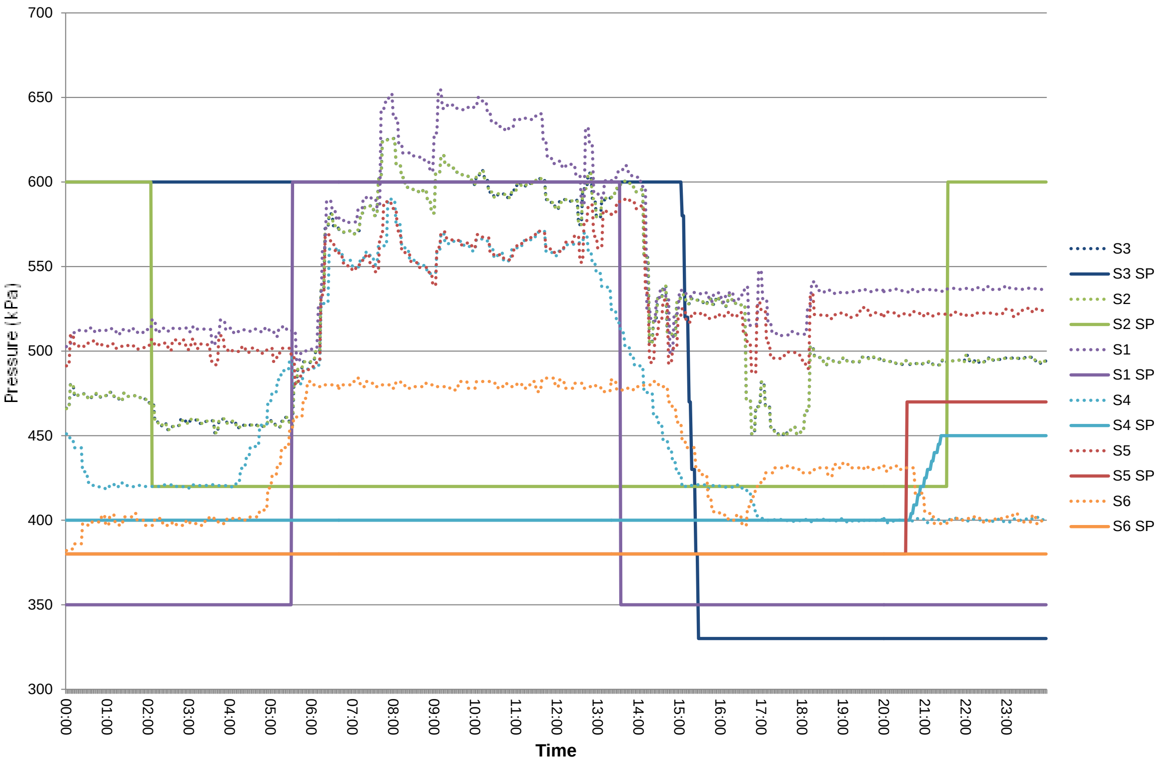

Figure 59: Platinum mineshaft pressures and set points

The DCS system is not currently operational. The current control is only optimised by simulation. As can be seen from the graphs in Section 4.3.3, optimised by simulation control can be efficient most of the time, but it cannot always supply efficient control as the network can deviate from the control [51]. Looking at the results obtained from Van Niekerk [52], it should be possible to reduce the energy consumption more by optimising the network set points further.

The DCC system was only installed and run in parallel to the existing control to evaluate the output of the system. Currently the original mine compressed air controller is active due to stability issues encountered by the DCS system.

There are two compressor houses located on the network with all compressors split between the two compressor houses. Figure 60 shows the location of the compressor houses and all of the mining shafts on the network. The shafts are all named S*x*, while the compressors houses are named C*x*. This network has a total pipeline length of 9.4 km. The network, as rebuilt in the DCC, can be seen in Figure 61.

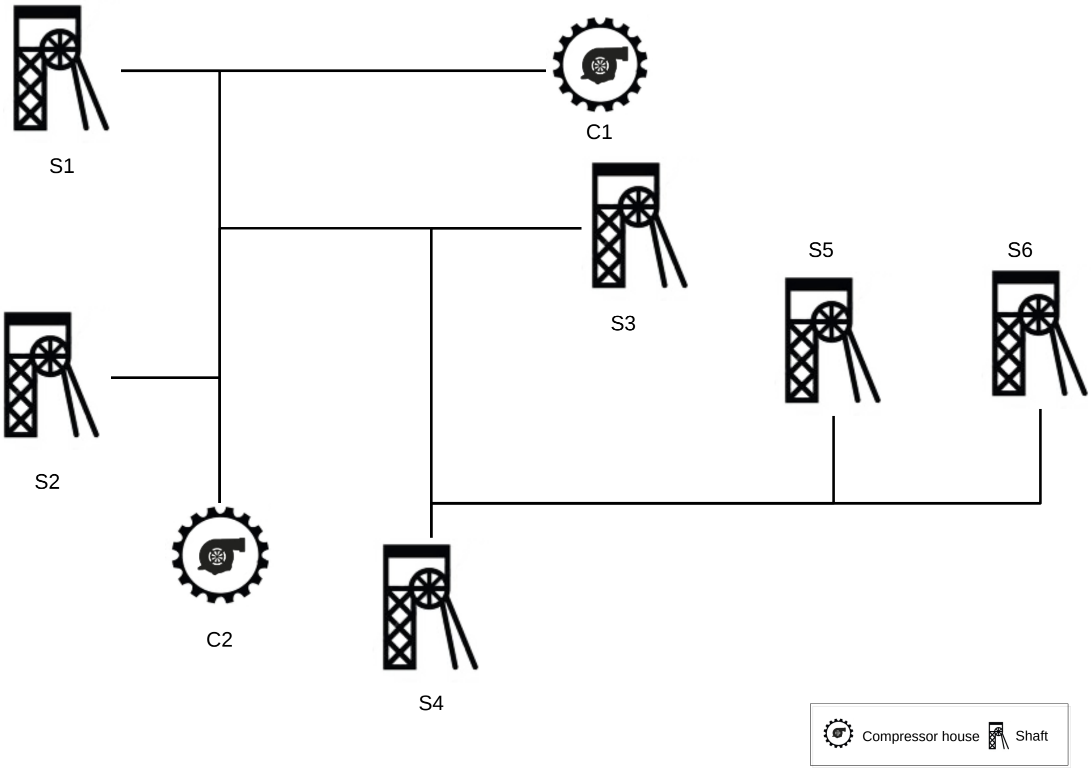

Figure 60: Platinum mine (not according to scale)

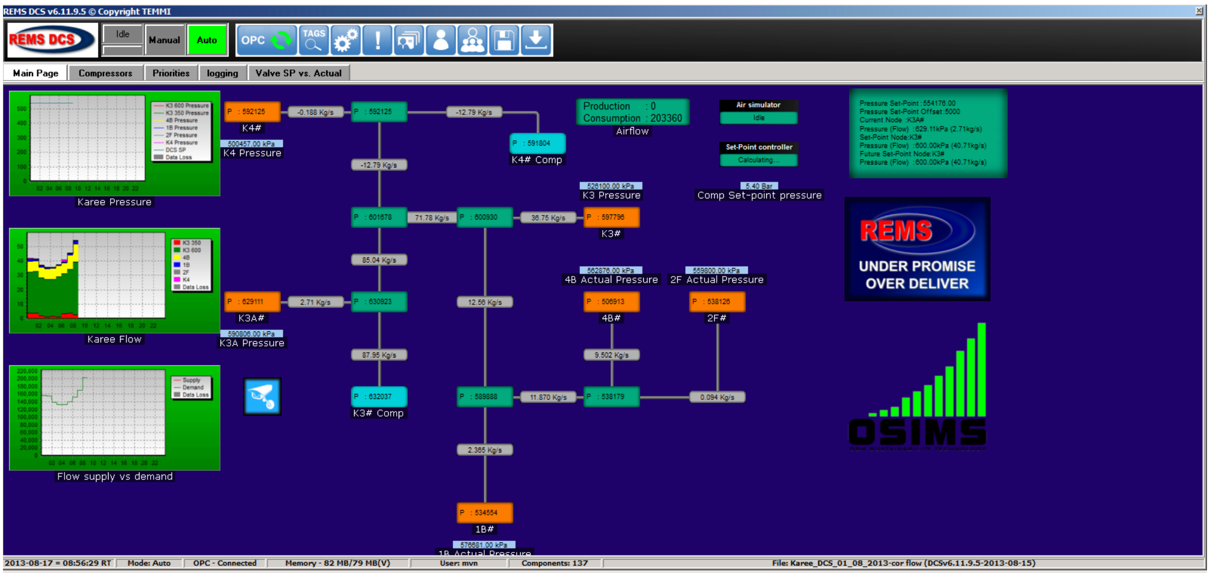

Figure 61: DCC platinum mine platform

The mining shafts on the network have set point pressures that oscillate between 350 kPa and 600 kPa. The set points of the shafts can be seen in Figure 59. The airflow to the mining shafts oscillates from just above 0 kg/s to over 45 kg/s on the main production shaft. The platinum mineshaft airflows can be viewed in Figure 62.

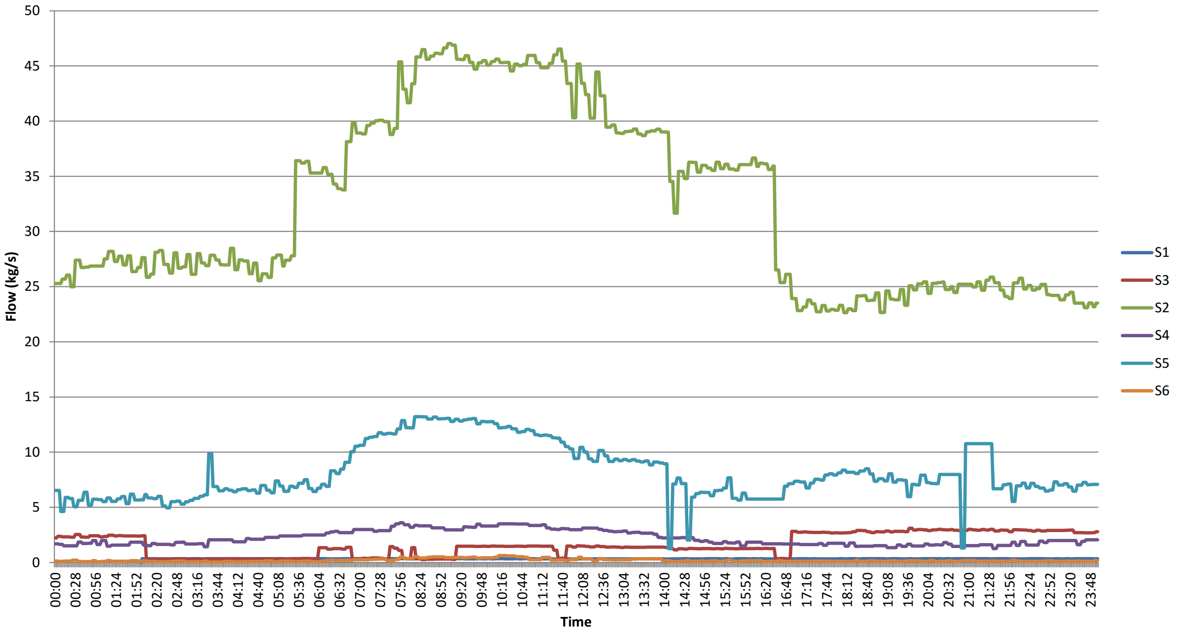

Figure 62: Platinum mineshaft airflows

4.3.2. Data logging

Data was logged in real time using an OPC connection to a SCADA computer. The values obtained from the SCADA system were obtained from data loggers installed on the machines and compressed air network as used by the mine. These loggers were calibrated and used by the mine.

Data used was from the same period in order to compare weekly and daily data. Most of the data was shown using a daily resolution in order to compare the results on a finer resolution. The data was logged on a resolution of two minutes. Additional graphs are shown in Appendix B.

All simulated values were based on the actual data logged. The DCC always use the newest available logged data as it constantly updates its own data. How estimations are made was discussed in 3.4.4.

4.3.3. Compressor priorities

Figure 63 shows the actual running compressors, while Figure 64 shows the compressors that the DCC desired to run. From the results, it can be seen that the DCC schedule differs only slightly from the actual running compressors. The results measured show the data from a day where actual usage closely followed the scheduled usage. This caused the installed current static control system to be very stable.

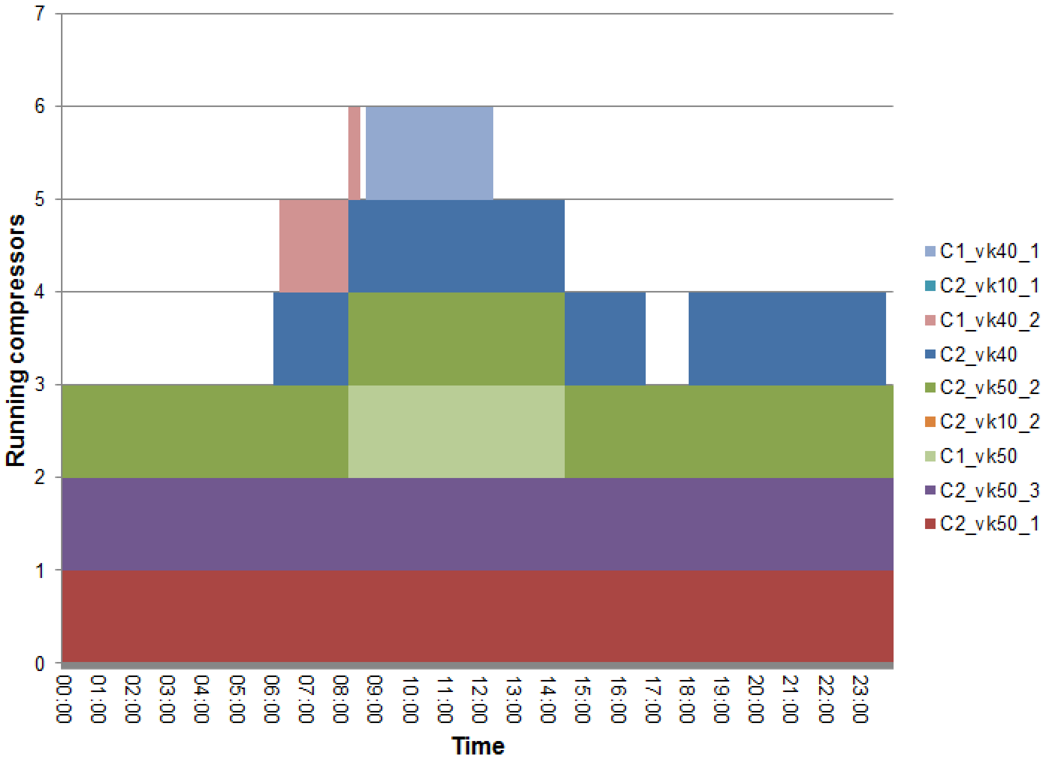

Figure 63: Actual running compressors

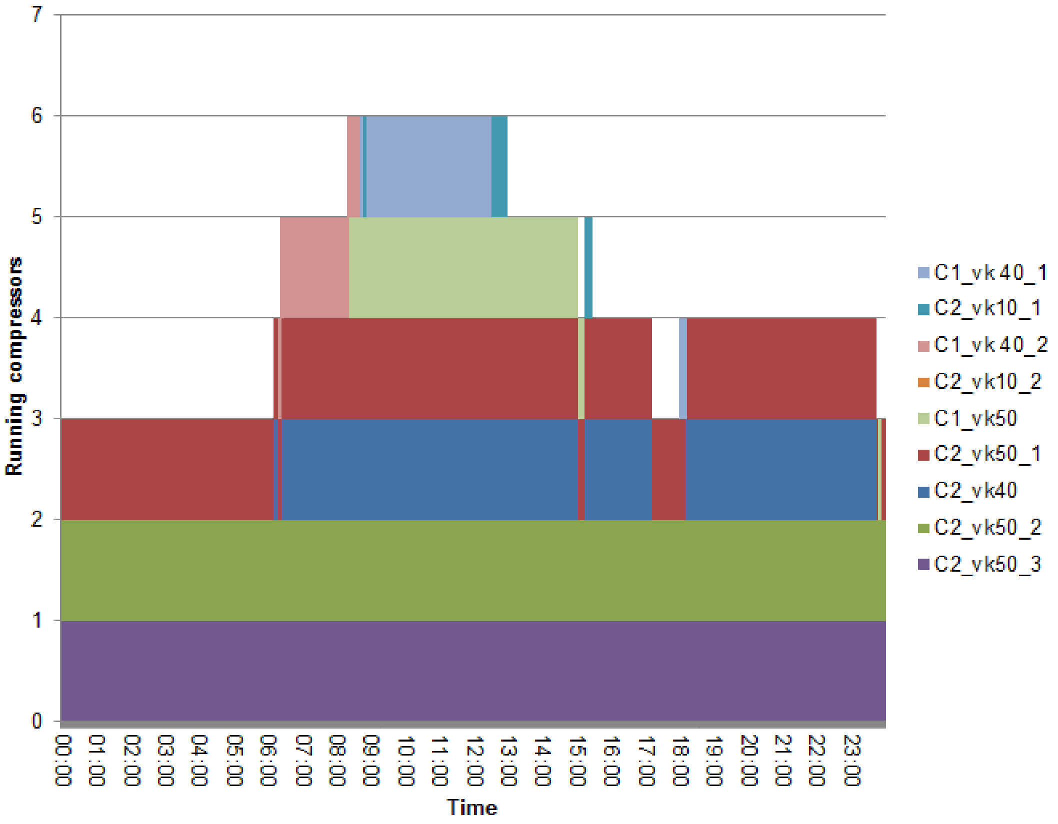

Figure 64: Simulated running compressors

Figure 65 shows the actual running compressors from a day where the actual usage did not follow the scheduled usage. It can be seen that a static controller could not react to changes in the network. This led to cycling and unstable control of the network.

From the results, it appears as if the DCC cycled compressors more. Van Heerden encountered a similar problem when simulating results [1]. This can be attributed to the DCC reacting to the actual compressor controller on the mine. In some situations, the DCC would have substituted compressors or issued a stop command to the compressor.

Because the actual compressor did not stop, the DCC perceived that it was running and considered this during the simulation. Thus, it might decide to keep it running the next time the DCC makes a decision. This is because the DCC ranks energy efficiency first and only takes action if the compressors reduce electricity consumption by a certain margin.

This margin was set at 0.5 MW for this test. This meant that if the solution that was chosen previously did not reduce the electricity consumption by more than 0.5 MW, the current operating solution would be chosen. This was only relevant when the current operating solution had not been eliminated before simulation because it was unusable.

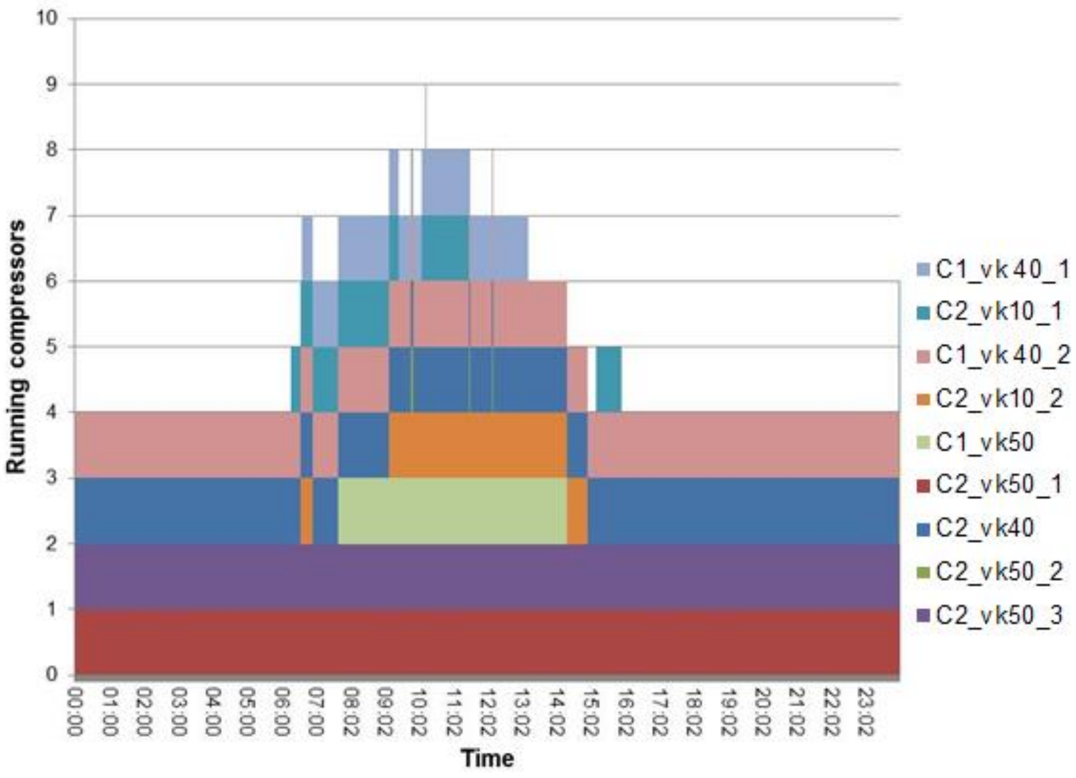

Figure 65: Actual operating compressors on a bad day

Results from more simulated days are given in Appendix B: Case study 1- extra compressor running graphs.

Figure 66 shows the measured supplied flow compared with the simulated supplied flow. It can be seen that the simulated supply flow closely mimicked the actual supplied flow. The simulated flow was smoothed by a filter before use to ensure that the spikes in the flow readings did not influence decision making.

These spikes in the flow readings were caused by ineffective control, compressor starts or stoppages. When a compressor starts, the local controller ramps up the compressor to near full power before the controller realises that the compressor does not need to run at full power. This causes the flow to spike as the local controller throttles the compressor back to its actual required output. When a compressor is stopped, the inverted occurs and the flow drops.

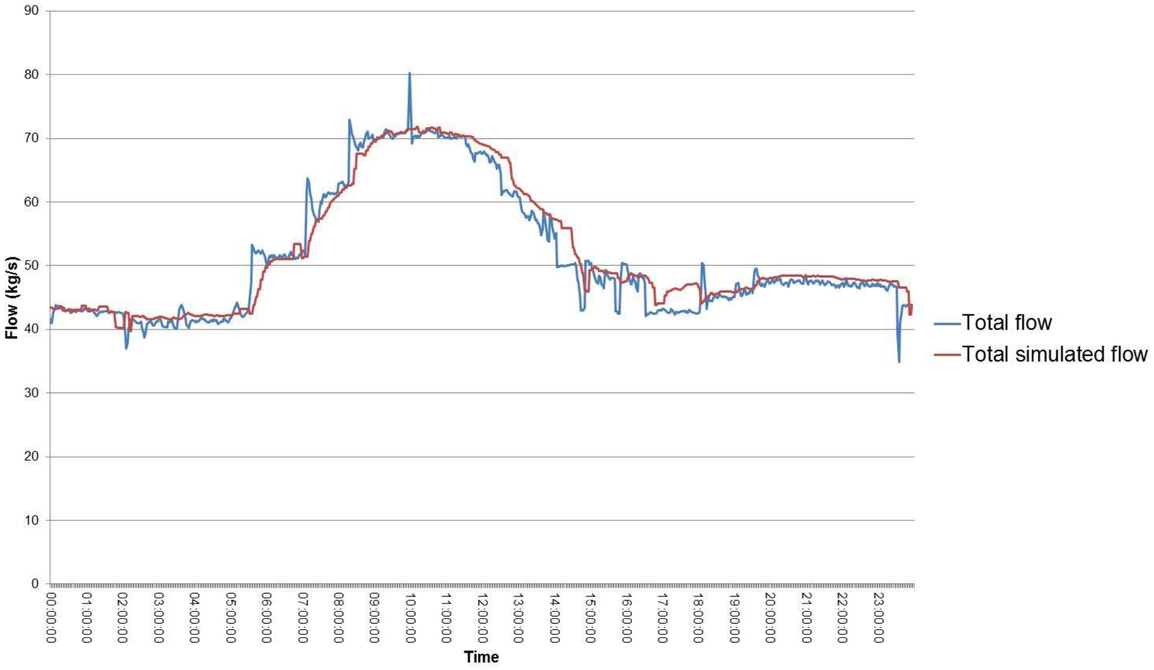

Figure 66: Actual delivered flow versus simulated delivery flow

Figure 67 shows the priorities of Compressor House 1. The black dotted line represents the number of running compressors. All priorities below or on the line can be seen as compressors that are running while priorities above it are compressors that are considered off. From the graph, it can be seen how the DCC dynamically changed the trimming compressor to ensure optimal power usage.

The compressors that were run on the black dotted line were used as trimming compressors. These compressors were run at reduced output. On the static controller, the smallest compressor was used as the trimming compressor. The DCC system selected the most efficient compressor to be used as the trimming compressor.

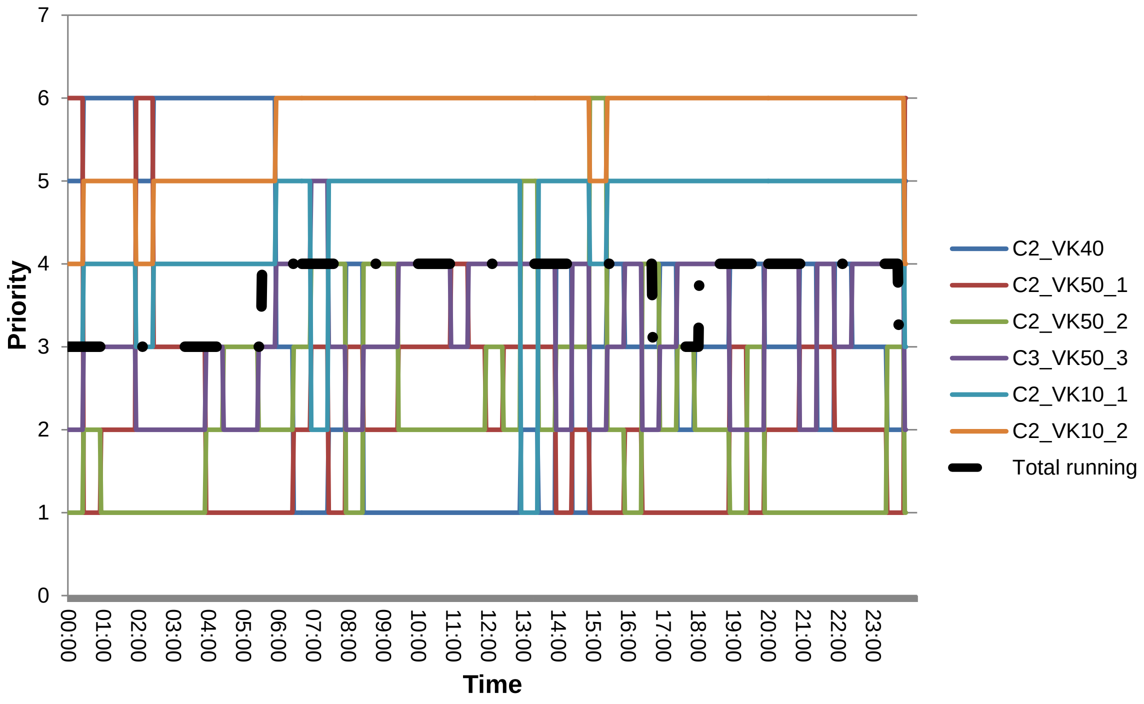

Figure 67: Priorities of compressors

4.3.4. Compressor set points

Figure 68 and Figure 69 compare the calculated required set point pressure and the actual recorded pressure readings at the respective compressor houses. In order to test the set point pressure accurately, the required pressure for each shaft was set to the actual recorded pressure for that specific shaft.

From the figures, it can be seen that the calculated set point pressure followed the pressure of each compressor house. The differences between the two can be attributed to different operating compressor combinations. It can thus be said that the DCC simulation calculated the set point pressure for this compressed air network accurately.

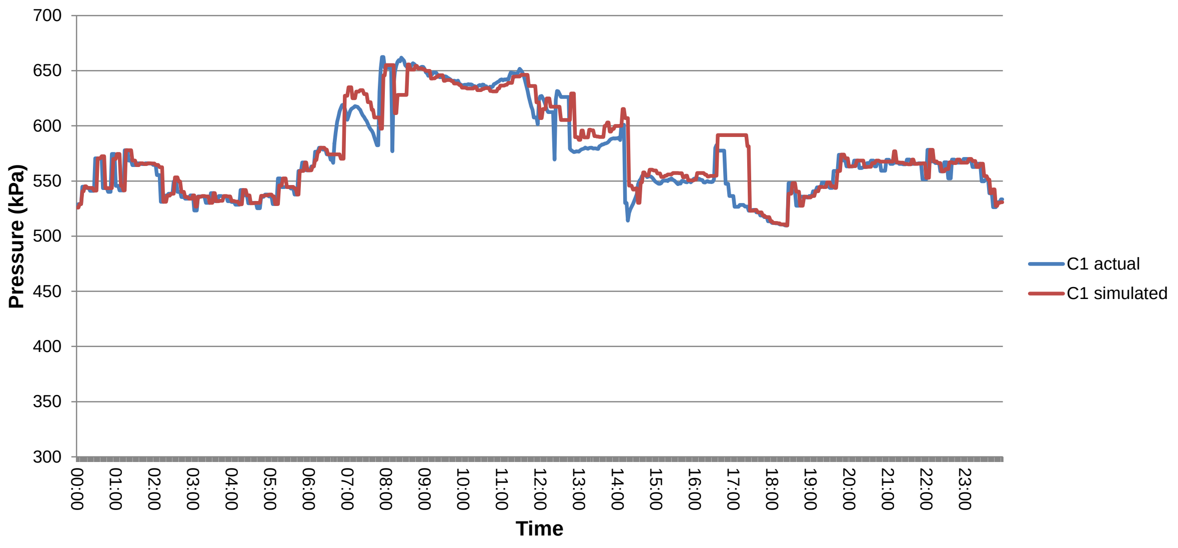

Figure 68: Actual pressure profile and calculated pressure profile of C1

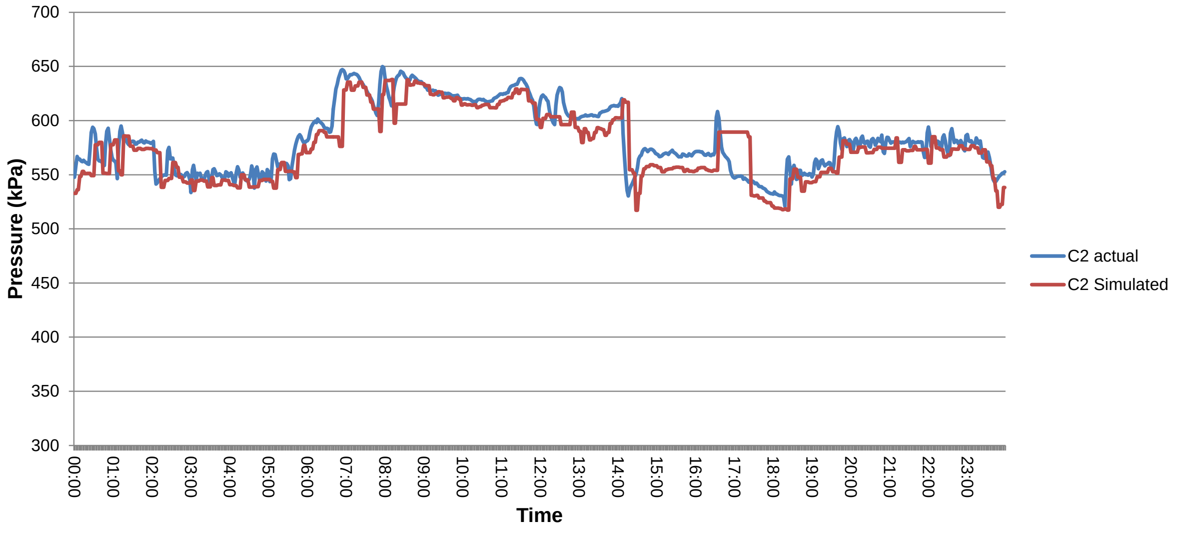

Figure 69: Actual pressure profile and calculated pressure profile of C2

4.3.5. Electricity usage

Figure 70 shows the actual power consumed by the compressors as well as the simulated compressor power usage. The simulated power's calculations where based from the logged power consumption as explained in Section 3.4.4. Figure 71 shows the daily electricity usage comparison. From Figure 70, Figure 71 and Figure 72 it can be seen that the simulated compressor used less electricity than the actual compressors, which represented a saving.

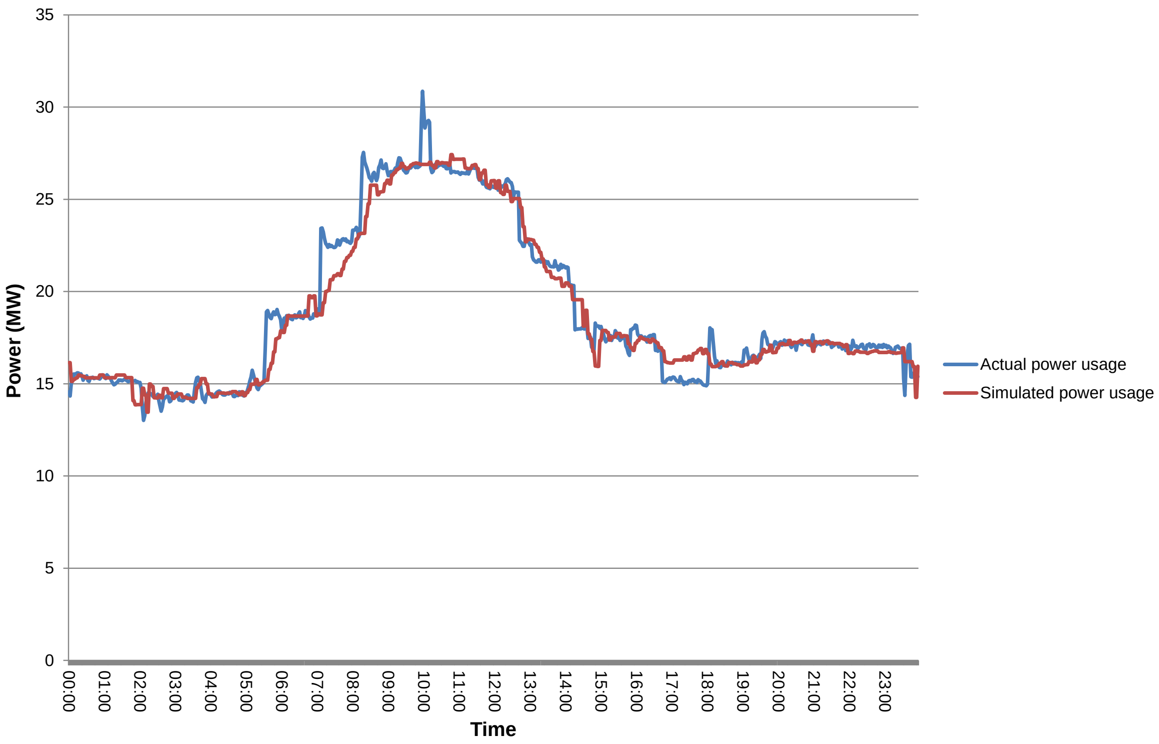

Figure 70: Actual power usage compared with simulated power usage

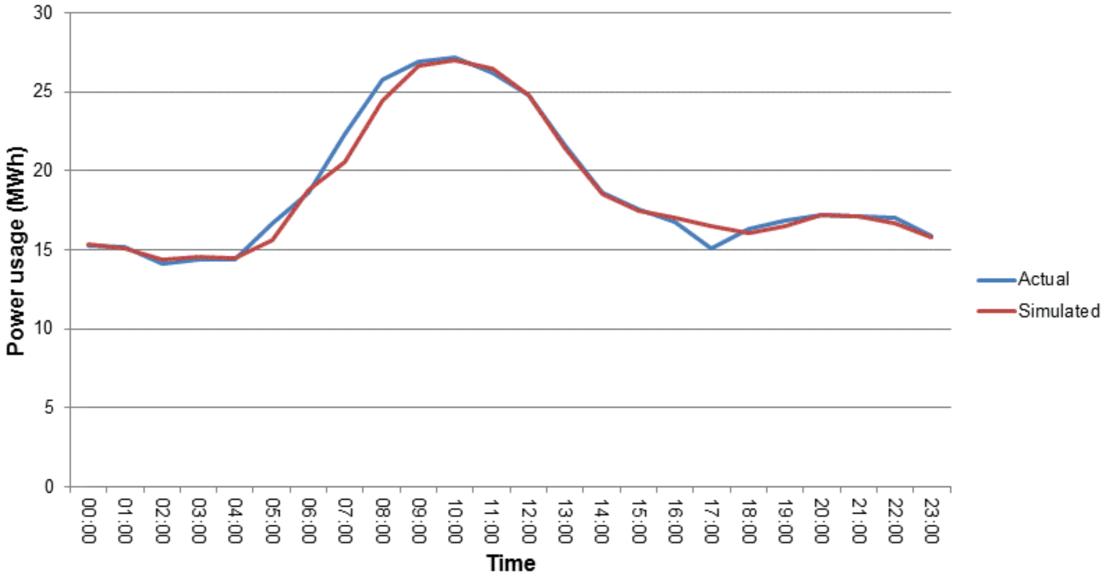

Figure 71: Accumulative power usage of actual compressors versus simulated compressors

Figure 71 shows the power usage of a single day. Figure 72 shows the accumulative power usage over a five-day week. From the graph, it can be seen that the simulated DCC solutions required less power than the actual running compressors. On average, the DCC reduced the energy consumption by 0.5 MWh over the entire week.

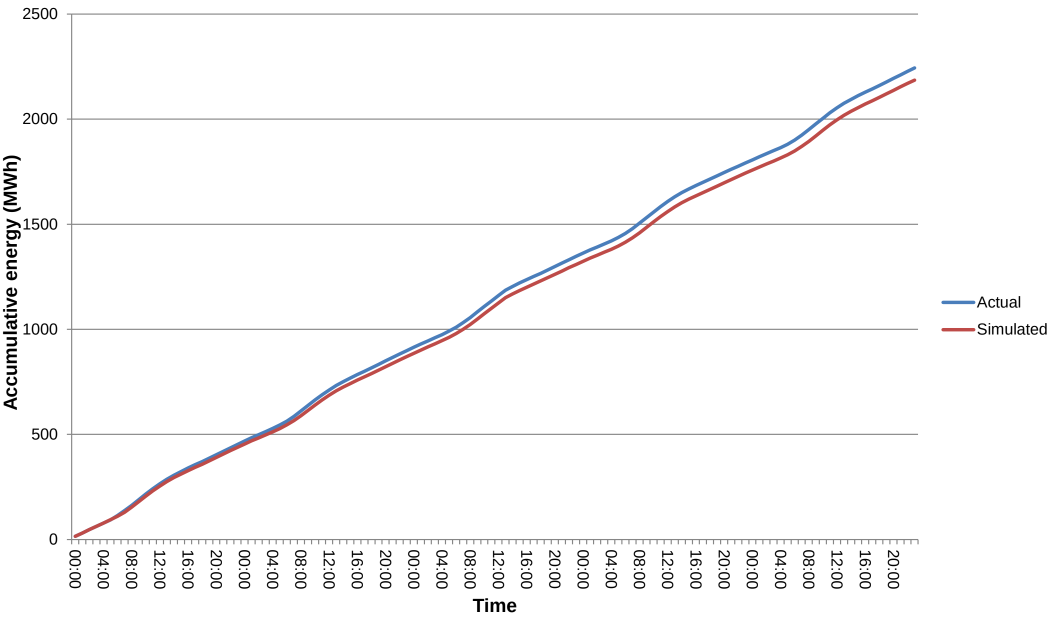

Figure 72: Accumulative energy over a week

4.3.6. Synopsis

It can be argued that the DCC system represents an improvement over the system designed by Van Heerden based purely on the performance obtained with the simulations. Although the current installed control strategy does not include the DCS, the control schedule was influenced by the DCS as it showed where control schedules had to be updated.

This system also does not suffer from the same stability problems as the DCS as outlined by Van Heerden [1]. This is also the reason that the DCS system is not currently being used by the mine as the stability issues compromise the usability.

Using Eskom 2014–2015 tariffs, the yearly amount saved on electricity cost would amount to R650 000 each year. Because this case study is done on the same mine as the DCS [1] and implementing the DCC would not incur any additional cost, the payback period would be minimal as the only cost that would be incurred is man-hours, but it is well within the period specified by the mine The exact calculations would require classified data such as billing hours, salaries, overloads etc. and can therefore not be disclosed.

If proper set point pressures are given to each of the shafts and proper control is implemented to allow the shafts to follow those set point pressures, the performance savings should increase even more. The obtained performance increase can largely be accredited to proper prioritisation of compressors.

On the current control, the smallest compressor is used as the trimming compressor even though it might not be the most efficient running compressor to use for trimming. However, the DCC ensures that it selects the most efficient trimming compressor.

4.4 Case Study 2: Gold mine

4.4.1. Compressed air network

The gold mine used as Case Study 2 has a compressed air network that features eight compressors with a total peak electrical capacity of 44.3 MW. The compressors are summarised in Table 7 with their flow and power ratings. This mine had already fitted automated control valves and all of the required pressure and flow meters to all the shafts. Thus, the mine already had a very advanced control scheme for the network although it was still static.

Table 7: Gold mine compressors

Name	Flow rating (kg/s)	Rated capacity (MW)	Location
C1_1	27	15.0	C1 compressor house
C1_2	27	15.0	C1 compressor house
C2_1	16	8.6	C2 compressor house
C2_2	11–16	4.8	C2 compressor house
C3_1	11–18	5.9	C3 compressor house
C3_2	11–18	5.9	C3 compressor house
C3_3	11–18	5.9	C3 compressor house

Each mining shaft has its own compressor house with the compressors divided between them. Figure 73 shows the location of the compressor houses as well as the locations of all the mining shafts on the network. Shafts were named *Sx*, compressor houses *Cx*, workshops and plants *Wx*. This network has a total pipeline length of 23 km. The network, as rebuilt in the DCC, can be seen in Figure 74.

Compressor C1_2 was installed but never used. The compressor was planned to be decommissioned but at that stage, the compressor had not yet been decommissioned. At the time, Compressor C2_1 was being recommissioned and would have been operational in the near future.

The mine had a network split installed to split the shafts from the processing plants [48]. The split was removed as the compressors were decommissioned. By using the network adaptability of the DCC, it could potentially increase electricity savings by splitting the network once more during off-peak hours.

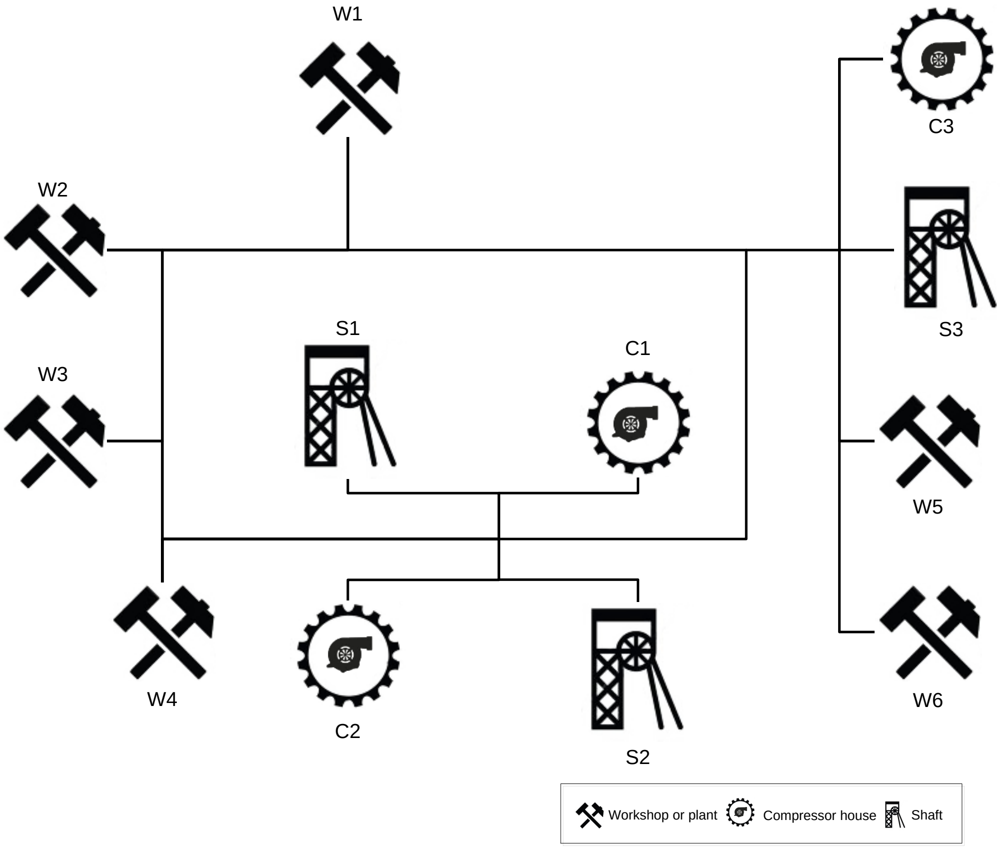

Figure 73: Gold mine layout (not according to scale)

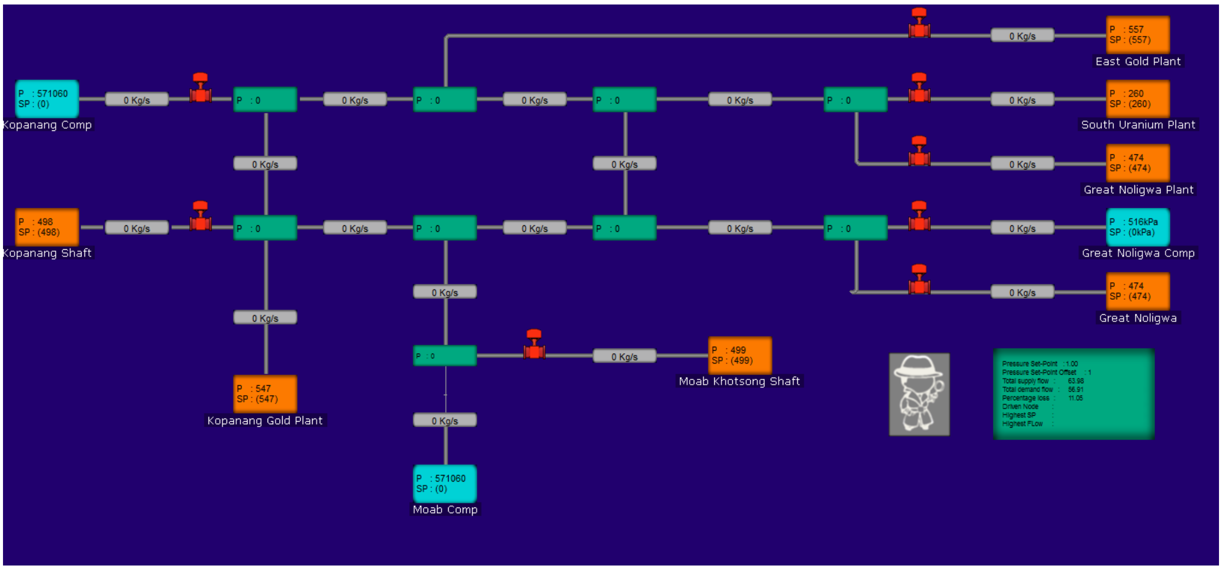

Pressure Set-Point : 1.00
Pressure Set-Point Offset : 1
Total supply flow : 63.98
Total demand flow : 56.91
Percentage loss : 11.05
Driven Node :
Highest SP :
Highest FLow :

Figure 74: Gold mine DCC layout

Due to time constraints, the DCC was never run on the actual mine and only simulated results could be obtained from investigations into a possible project. Because no project had been signed off by the client, three possible projects were proposed. All three of the proposed projects were to be funded via an Eskom DSM project. The first possible project was installing only the DCC and using it to control the compressors.

The second project was installing and controlling the compressors via the DCC while also installing an additional pipeline directly connecting C2 and C3. This can be seen in Figure 75. The reason for this pipeline was to enable the mine to increase the pressure obtainable at S2 as at that stage they desired a higher pressure at S2.

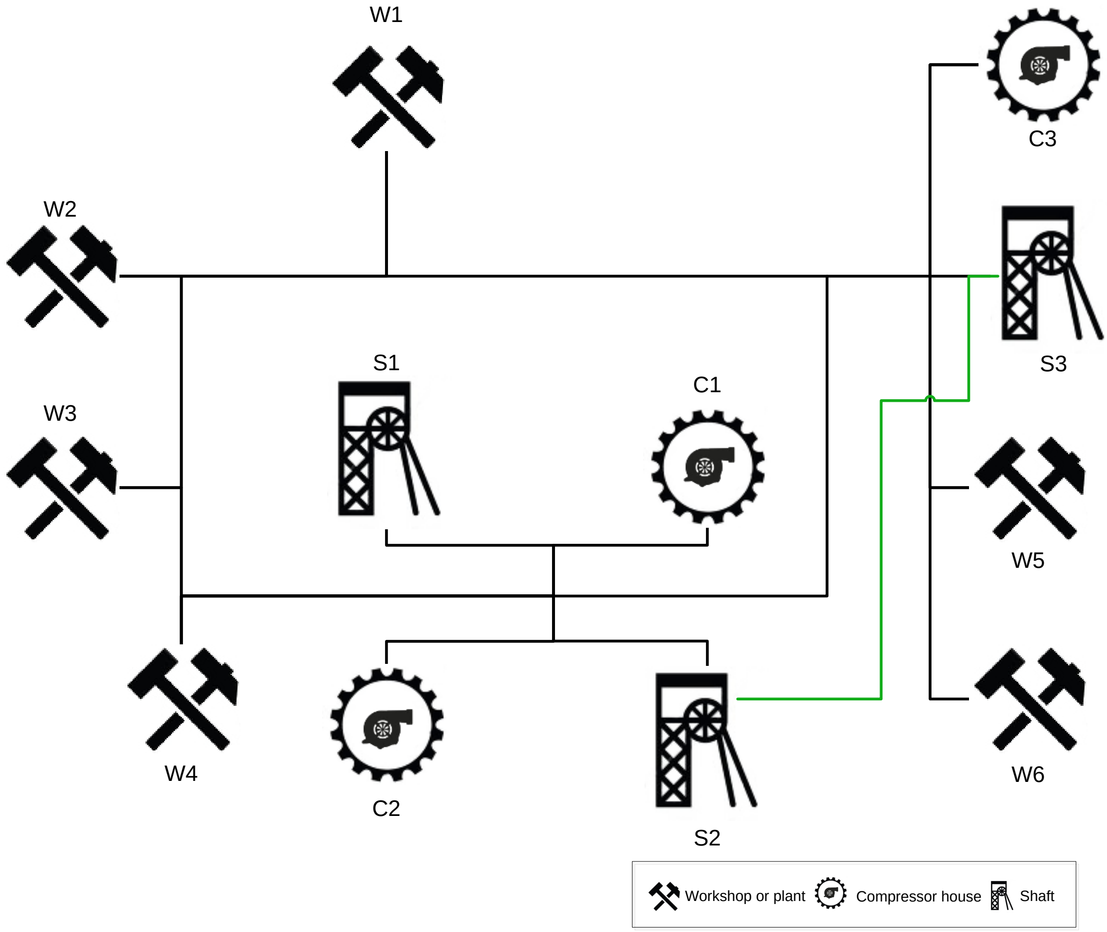

Figure 75: Gold mine with new proposed pipeline (not according to scale)

Lastly, the third project was decommissioning compressor C1_1, which was used as a baseload compressor by the client, recommissioning compressor C2_1 and controlling the compressors via the DCC.

The mining shafts on the network have pressures that oscillate between 420 kPa and 530 kPa. The pressures of the shafts can be seen in Figure 76. The airflow to the mining shafts oscillates from 8 kg/s to over 18 kg/s. The actual airflows can be viewed in Figure 77.

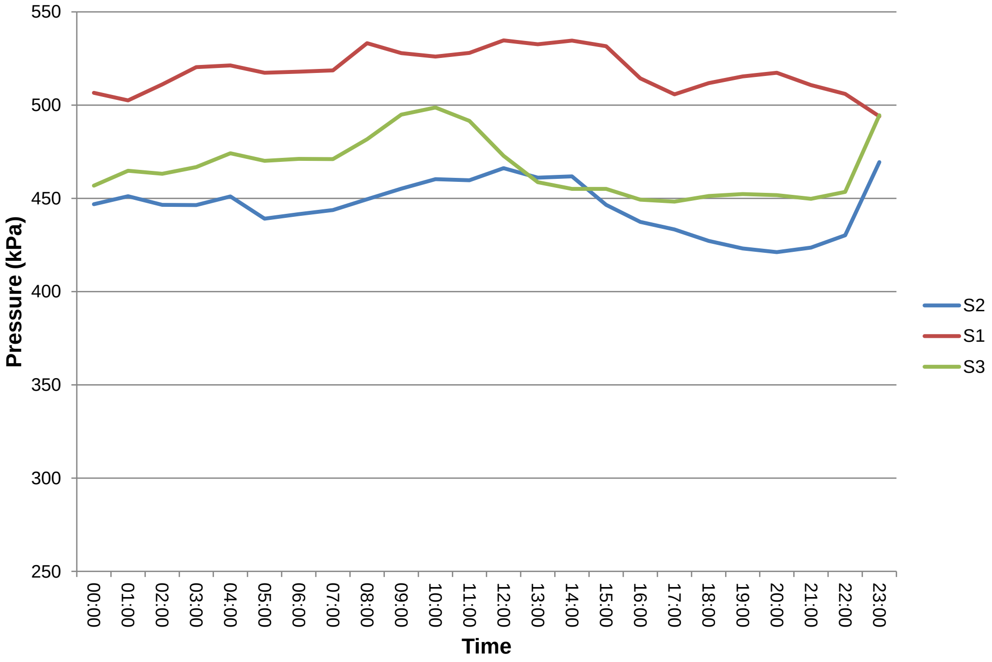

Figure 76: Gold mine actual shaft pressures

Figure 77: Gold mine actual shaft flows

The pressures and flow of the processing plants can be seen in Figure 78 and Figure 79 respectively. From the figures it can be seen that the processing plants used low flow but at high pressures. The shafts used higher flow but at lower pressures. This presents an opportunity to split the network. W2 and W3 were workshops with less airflow and they were thus omitted.

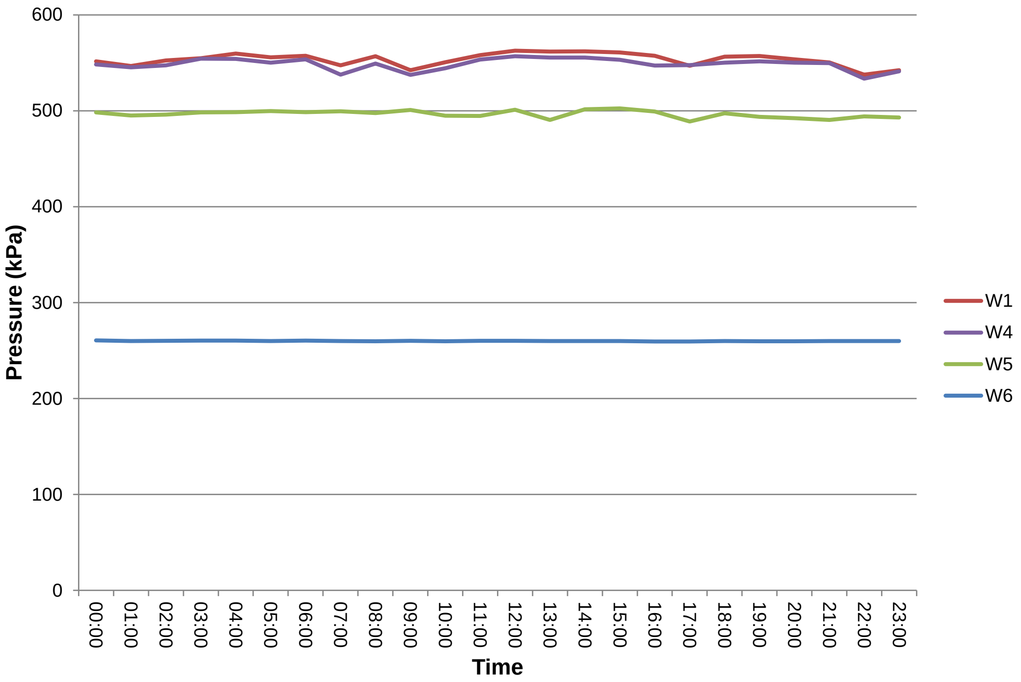

Figure 78: Gold plants actual pressure profiles

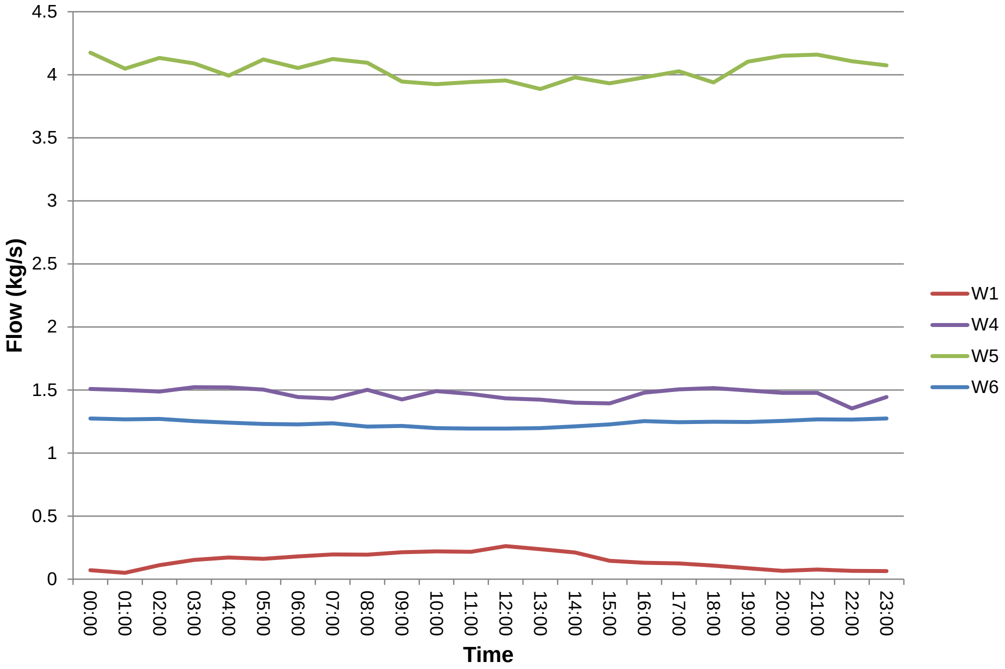

Figure 79: Gold plant actual flow profiles

From the above-mentioned graphs, it can be seen that the existing control was very consistent, but from Figure 77 it can be seen that the airflow dropped during the evening peak. This presented an opportunity to drop the pressure during the evening peak period. All three proposed projects were based on this fact to increase the electricity savings. An automated control system such as DCC enables the user to do this easily without the risk of interfering with production.

4.4.2. Data logging

All data was logged in real time using an OPC connection to a SCADA computer. The values obtained from the SCADA were obtained from data loggers installed on the machines and compressed air network as used by the mine. These loggers were calibrated by a contracted metering company and used by the mine.

All simulated power values were based on the third-party simulation application used successfully by Van Niekerk [52] and Deysel [109] to simulate and compare actual results on compressed air networks. In both cases, the simulated results could be seen as conservative as the actual results obtained were lower.

The simulated results were obtained by simulating the operating compressors as was selected by the DCC. The other simulated values were obtained directly from the logs generated by DCC during simulated runtime. Only the first proposed project was simulated because the result would be similar for each of the projects with only minor differences.

4.4.3. Compressor priorities

Figure 80 shows the simulated running compressors of the mine. The detailed simulated results shown in Figure 80 were only for first proposed project. The cycling that was present could be attributed to three factors. Firstly, the individual compressor data was not complete and only summarised delivery of flow and power were inserted in the DCC.

Secondly, the simulations were done without future flow estimations for the mining shafts and processing plants as these were never logged. Lastly, if this case occurred, the cycling period would be very small and would not occur because the compressor would be loaded and unloaded as explained in Section 3.4.5.

This can be compared to the graph in Figure 81, which shows the simulated airflow. From Figure 80 and Figure 81 , the running compressors can be seen to match the demand for compressed air closely. Figure 81 also shows how the simulated flow and compressed air supplied by the simulated running compressors compared.

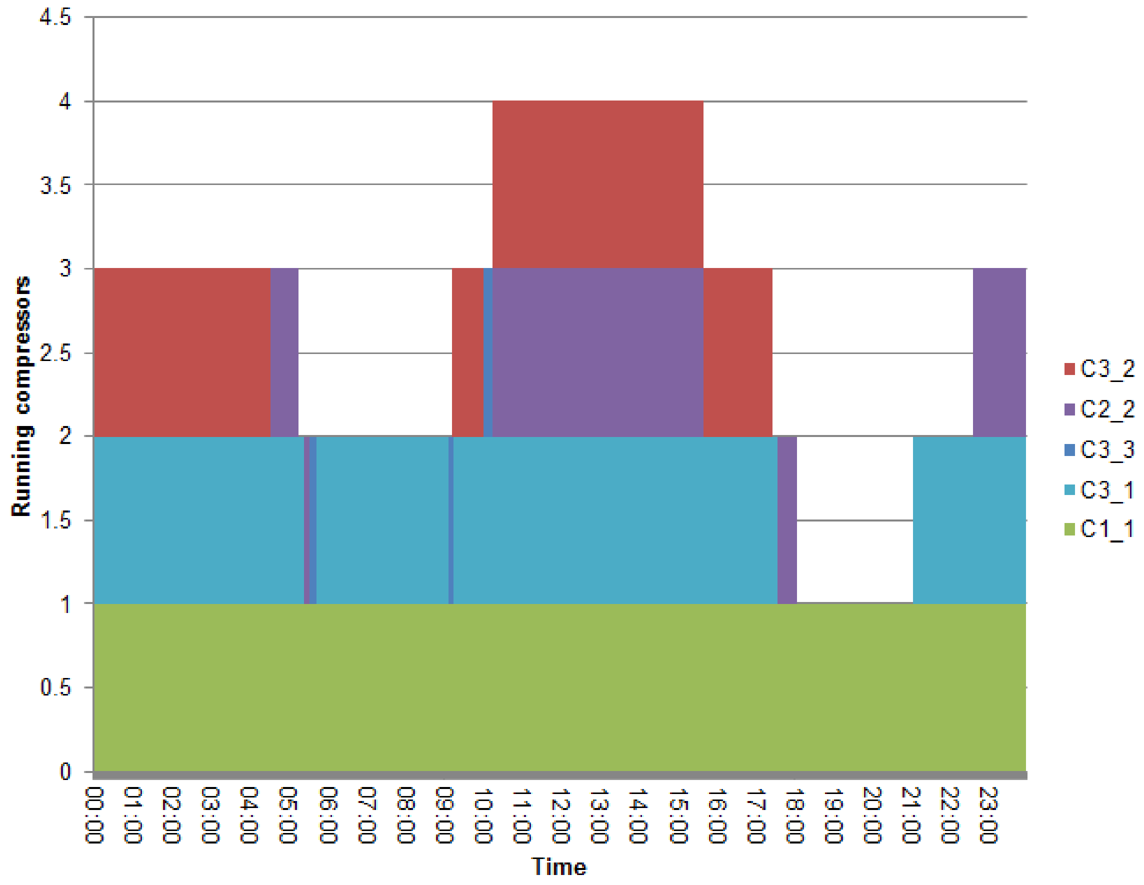

Figure 80: Gold mine simulated running compressors

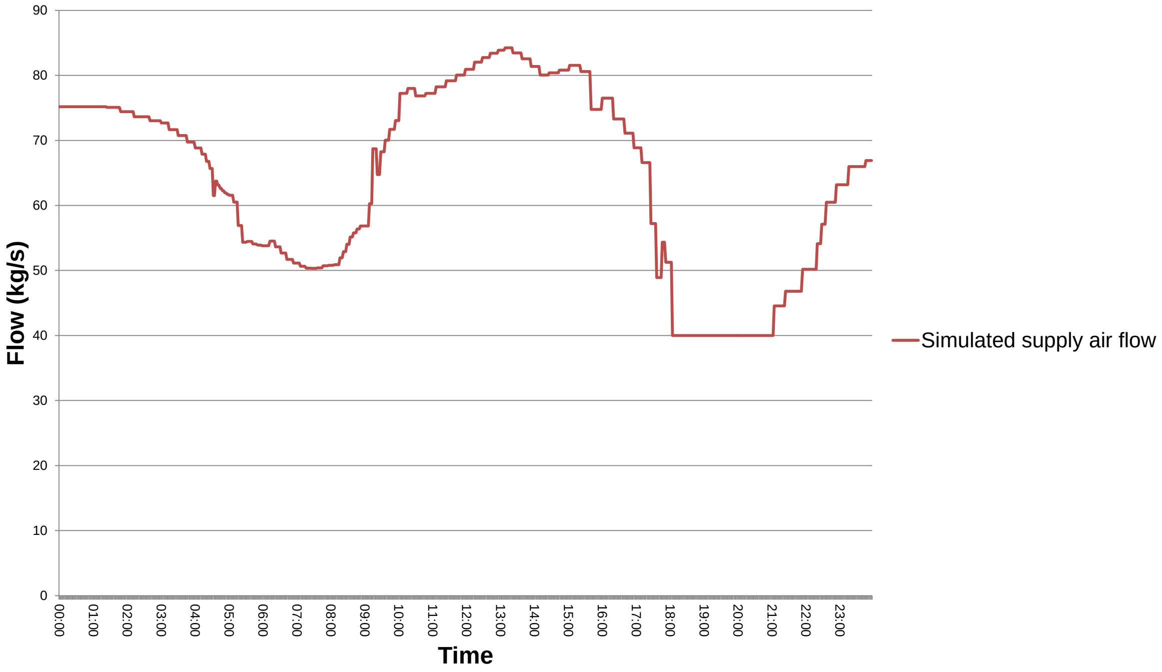

Figure 81: Gold mine simulated airflow

4.4.4. Compressor set points

In Figure 82, the simulated compressor pressure set points can be seen. This detailed result is only shown for the first proposed project.

This corresponds to the actual shaft pressures shown in Figure 76 and the actual workshop pressures shown in Figure 78. Because the required demand pressures did not vary much, the simulated supply pressures also stayed relatively constant although they did change over time as can be seen from Figure 82.

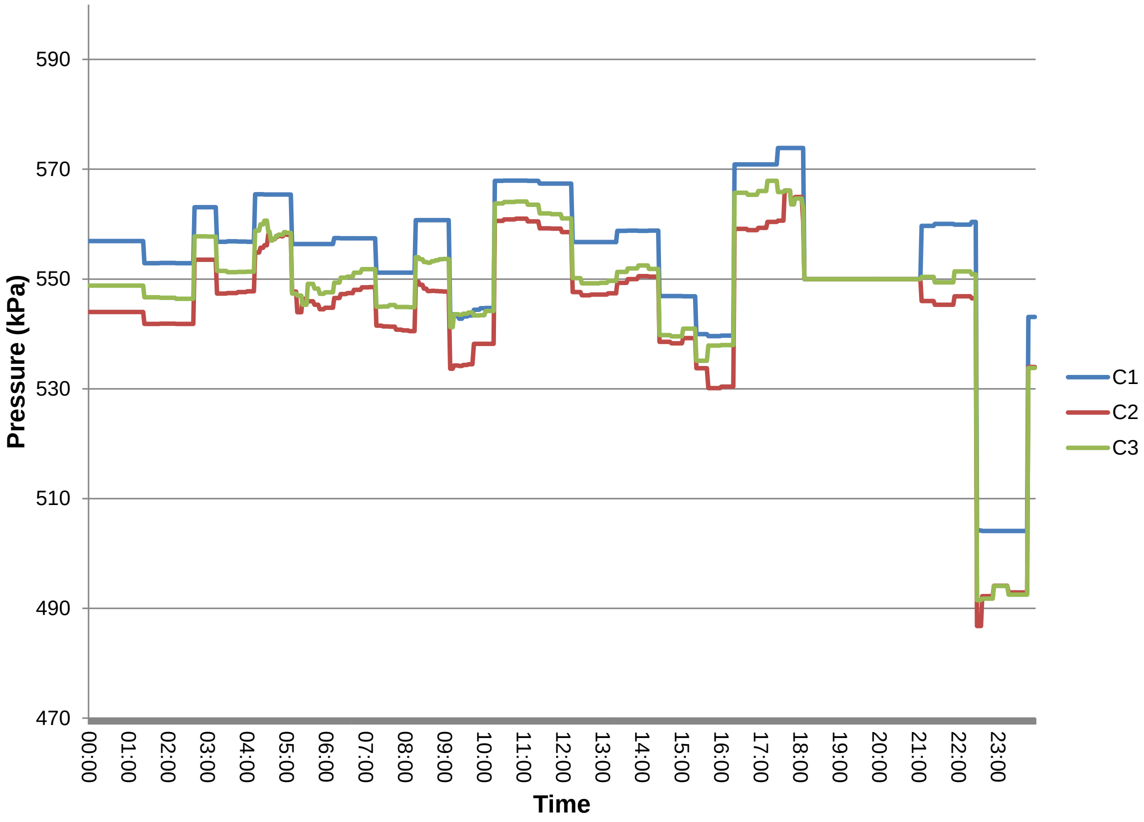

Figure 82: Compressor pressure set points

4.4.5. Electricity usage

Using the compressor schedules that the DCC calculated for the proposed projects, power profiles were simulated for each of the proposed projects. For the first proposed project, the simulated power profile can be seen in Figure 83. In this scenario, only the DCC was installed and allowed to control the network. The DCC reduced the energy over the period by 32 MWh. The energy reduction was aided by a peak clipping implementation.

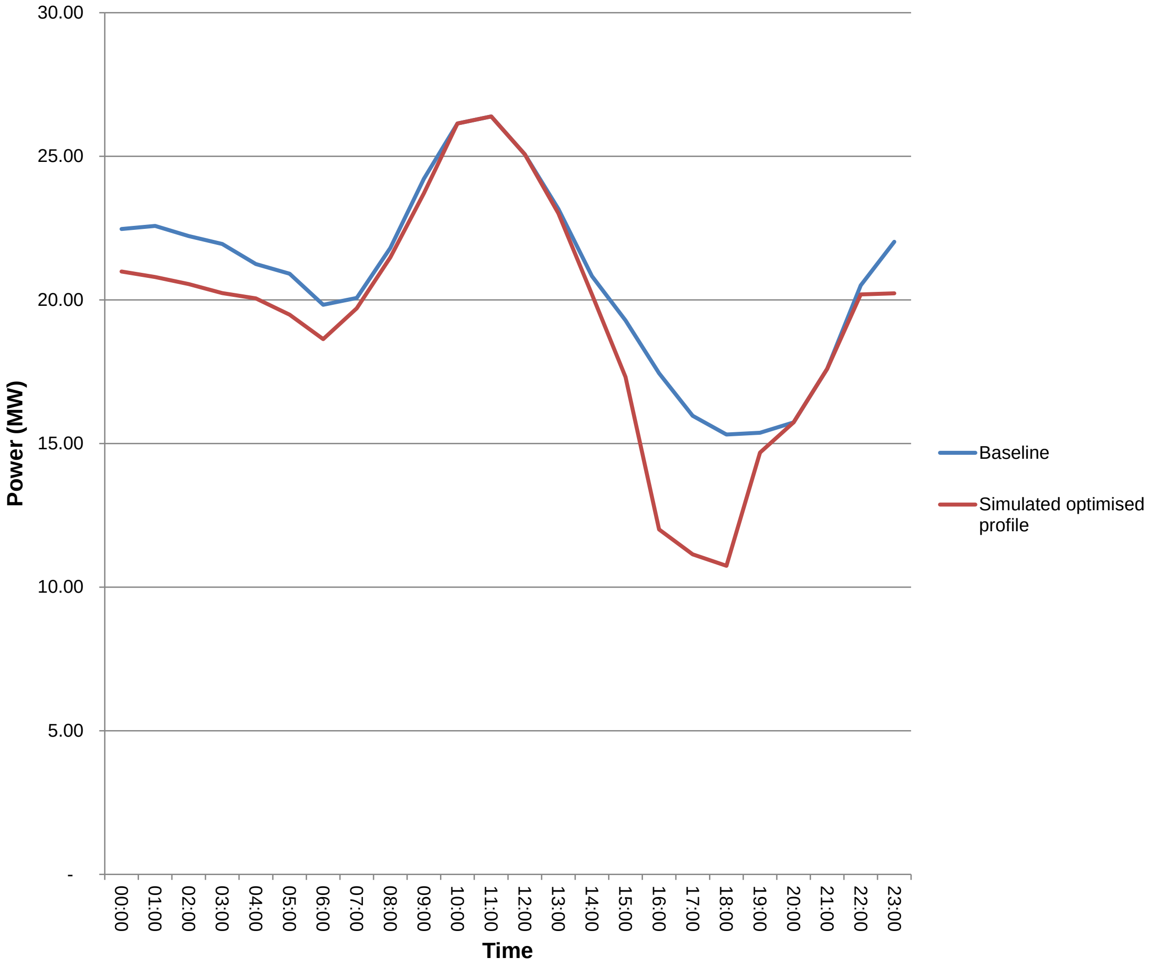

Figure 83: Gold mine simulated power profile 1

In the second proposed project, the DCC was allowed to control the network but an additional pipeline was installed to directly connect S2 and S3. The resulting power profile can be seen in Figure 84. In this scenario, the DCC reduced the energy by 18 MWh over the period. This proposed project countered some of the savings obtained from the first proposed project, but it increased the pressure supply to S2.

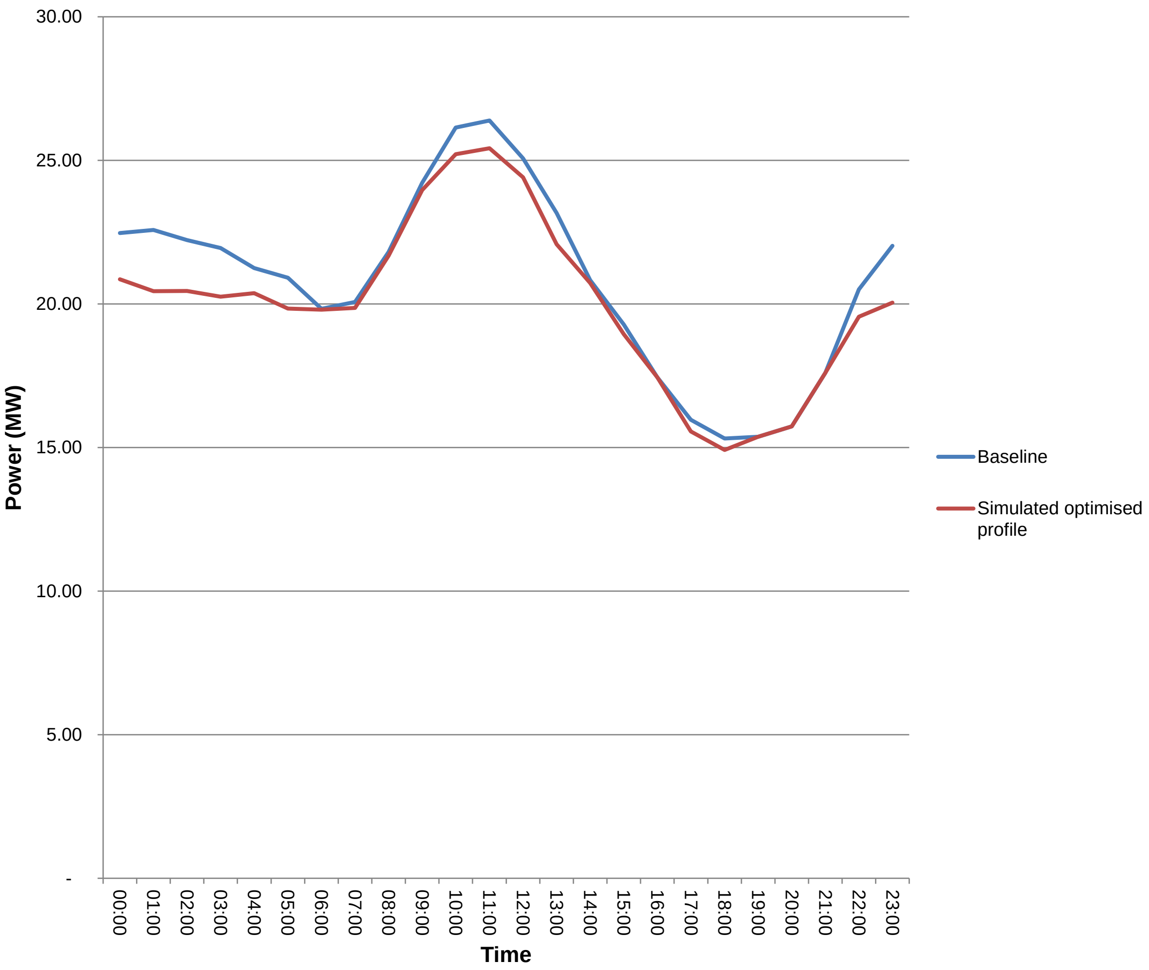

Figure 84: Gold mine simulated power profile 2

In the third proposed project, the baseload compressor C2_1 was decommissioned and compressor C2_1 was recommissioned. All the compressors in the network were then also controlled by the DCC. The resulting power profile can be seen in Figure 85. In this scenario, the DCC reduced the energy by 86 MWh over the period. This project would cost the most money to implement but would give the biggest energy reduction of all three proposed projects.

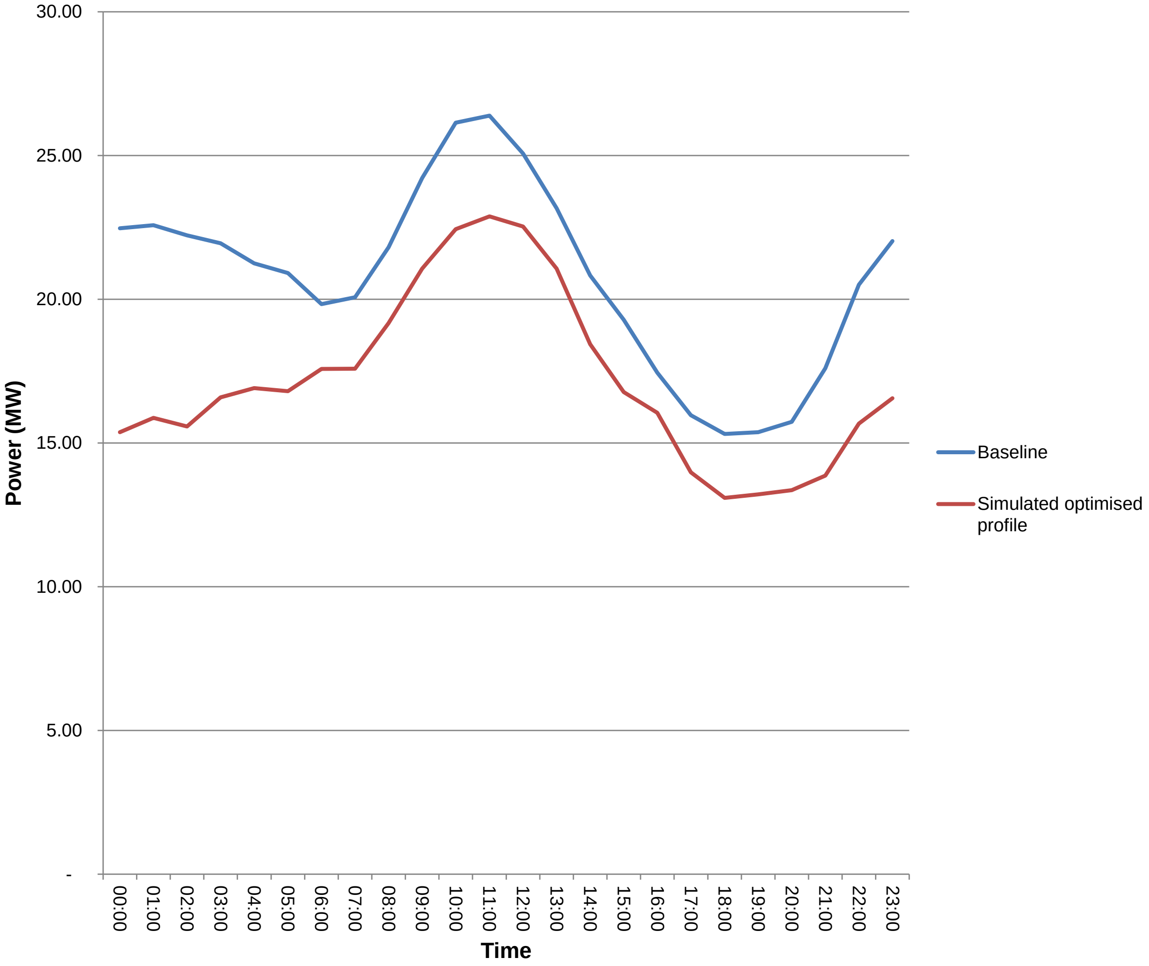

Figure 85: Gold mine simulated power profile 3

4.4.6. Synopsis

The DCC controller was simulated in three different proposed projects with the help of an external simulation package. In all three cases, the DCC reduced the energy consumption of the compressed air ring. The payback periods for each of the three different cases would be different. This is because the installation cost and required equipment differ between the three projects

Due to time constraints at the time of writing this document, a specific project has not yet been selected and implemented. The specific project selected would depend on the amount of funds supplied by the Eskom DSM project and the amount the client would be willing to spend.

With the first proposed project, the payback period would be minimal as very few costs would be incurred. The cost savings could start from just after implementation. If the first project was implemented, the cost saving would amount to R4 500 000 per annum.

With the second proposed project, the payback period would also be minimal as very few costs would be incurred. The cost of the pipeline has not been taken into account as the mine wanted to install the pipeline irrespective of an implemented project or not. The cost savings could start from just after implementation. If the second project was implemented, the cost saving would amount to R1 900 000 per annum.

The cost saving of the second proposed project is less than that of the first proposed project. The second proposed project was only done at the insistence of the client. The client wanted to install the additional pipeline because the pressure achievable at S2 was not at a desirable level and the client wanted the pressure to be higher.

With the third proposed project, the payback period would be higher as the cost of recommissioning a compressor needs to be considered. The cost savings could start from just after implementation but the incurred cost would have to be recovered before a net saving can be made. If the third project was implemented, the cost saving would amount to R17 800 000 per annum.

Depending on the project implemented, the DCC could save the client between R1 900 000 and R17 800 000 per annum in electricity savings. By combining this with the dynamic control of the DCC, it can be argued that the controller can be successful if it was implemented at the gold mine, as the controller will save the mine financially as well as improve their compressor control.

4.5 Novel contributions verification

This section serves as a conclusion to the verification of all the novel contributions. Each contribution is discussed independently as well as where the detailed results were discussed in this document.

1. Individual compressor set point pressure calculations for dynamic control

Current static controllers have pressure profiles to determine their pressure set points. These pressure set points are determined by simulations run by qualified personnel. In order to ensure optimal control, these simulations have to be repeated continuously. Current dynamic control only has the option to calculate a single pressure set point per compressed air ring.

The DCC simulated a set point pressure for each of the compressor houses. When this is combined with the CompressorPrioritiser component, each compressor will have its own individual compressed air pressure set point. Based on the simulations in Section 4.4.4 and

Section 4.3.4, it can be argued that the DCC simulated the required pressure set points accurately in order to satisfy the demand on a compressed air network.

2. Dynamic optimal operating compressor selection

Currently, static controllers control the operating compressors directly from the priority list. When the compressed air pressure from the ring drops below a certain point, a compressor is started. When the pressure increases above a certain point, a compressor is stopped. With dynamic control, compressors are chosen dynamically based on their size with the assumption that larger compressors are more efficient.

The DCC selected compressors dynamically based on their logged electricity efficiency. While the simulations were similar to the static control, small changes were still prevalent. Because of the similarity, it can only be assumed that the current static control is up to date. The simulations done in Section 4.3.2 and Section 4.4.2 proved that the DCC could select operating compressors dynamically. It can be argued that the DCC dynamically chooses optimal operating compressors.

3. Dynamic optimal trimming compressor selection

Static controllers currently only have fixed priorities that sometimes will differ based on time, but they still do not change based on changes of the compressed air ring. They have to be changed manually. Dynamic controllers can change priorities dynamically, but these are only based on the size of the compressor with the assumption that the largest compressor is the most efficient one. The trimming compressor is then selected as the compressor with the lowest running priority.

The DCC changed the compressors priorities dynamically based on the status and demand for compressed air from the ring. The simulations done in Section 4.3.2 and 4.4.2 proved that the DCC changes the priorities of the compressors dynamically. These priorities were used to select the trimming compressors. From the results obtained in Section 4.3.5 and 4.4.5, it could be argued that the DCC successfully prioritised compressors to select trimming compressors to reduce the power consumption.

4. Future flow prediction of demand in network

Currently, none of the available controllers does future flow prediction of the demand for compressed air in the ring. The DCC successfully predicted the airflow of the compressed air ring to start larger compressors at the start of peak times. This can be seen in Section 4.3.3. From the results, it can be seen that the DCC does not select small compressors only to select larger compressors later on. The DCC selects the larger compressors from the start to be able to supply the required flow in the future. The design and working of future flow prediction were discussed in Section 3.4.2.

5. Location attentive

Current dynamic controls do not consider the location of the compressors as well as the location of the demand for compressed air. The DCC successfully considered the location of the compressors and demand for compressed air in a test specifically designed for this in Section 3.5.2. The location attentiveness is incorporated into the design of the DCC and compressor selection. The design and working of the location awareness were discussed in Section 3.3.

6. Real-time compressed air network configuration adaption

Because current network controllers do not consider the location of the compressors, they also cannot adapt dynamically to changes in the network configuration. By being location attentive and possessing the ability to control the network from a two-dimensional angle, the DCC can close certain parts of the network dynamically and still continue to control the network optimally.

Real-time network configuration was proved in Section 3.5.2 by closing a compressed air pipe. The DCC successfully adapted and reconfigured the required compressors to supply the required compressed air continuously.

7. Integrated comprehensive compressed air solution

All of the novel contributions were met and they all functioned as part of a single compressed air controller known as the DCC. It can, therefore, be said that the DCC supplies a single integrated comprehensive compressed air solution.

4.6 Summary

This chapter provided results obtained from actual mine simulations. The results obtained were compared to how the actual mine performed. The simulations that were run on the mines were used to establish the validity of the network simulations. The simulations also proved that it is possible to save electricity by using a dynamic control system that continuously optimises the control schedule.

Lastly, this section added a small section to revisit the novel contributions to summarise the results obtained for each contribution. The location in the document of each specific test was also mentioned.

5 Conclusion and recommendations

5.1 Conclusion

The DCS was designed and created as the first dynamic compressor controller for mine compressed air networks [1]. Van Heerden mentioned a clear need for an improved system and the DCC was designed and created to improve upon these limitations [1].

The DCC proved during testing that it could run for extended periods without the need for manual user input that the DCS required. The DCC proved that it could control the compressed air networks dynamically in such a way as to reduce cycling and compressed air usage.

Existing dynamic compressed air controllers can only calculate one set point pressure for the entire compressed air network. During testing and verification, the DCC proved that it could calculate multiple set point pressures – one for each of the compressor houses situated in the compressed air network.

The DCC proved that it could select the optimal operating compressors dynamically, which could supply the compressed air network with enough compressed air to function as intended. The solutions provided by the DCC required less electricity to operate than the DCS.

While existing dynamic compressor controllers such as the DCS changed the priority of compressors dynamically based on the size of the compressors, the DCC successfully calculated the priorities based on recorded performance data. Combined with the dynamically selected operating compressors, the DCC reduced the electricity consumption of the compressed air network.

The DCC successfully predicted the future flow for use in simulation. This was proved by the fact that the DCC never had to start larger compressors soon after starting another compressor because the previously selected compressor could not supply enough compressed air. This made it a dynamic compressed air controller that could be used to predict the compressed air demand of industrial compressed air networks.

While the DCS could simulate compressed air networks, it never considered the effect of the location of a compressor when making a decision. During tests specifically designed to determine the effect on the location of the compressor, the DCC successfully made the correct decisions when selecting operating compressors.

By controlling the network dynamically, the DCC would reduce the required user input in order to control a compressed air network successfully. The DCC also achieved reduced compressor energy usage on sites where the control was considered optimal by dynamically adjusting to changes on the network.

On the one case study, the DCC simulated an electricity reduction of 0.5 MWh by only improving the trimming compressor selection. On the other case study, savings of between 32 MWh and 86 MWh were simulated. This varied by the type of project to be implemented. The large electricity reduction difference between the two case studies can be explained by the fact that the DCS was operated in a limited way on the first case study – through that it helped the mine to optimise its compressed air network schedule.

In South Africa, mines consume 16% of the total electricity with gold and platinum accounting for 80% of this. The amount of electricity consumed by compressed air generation differs from 25% in gold mines to 40% in platinum mines. This translates to 6% of the total electricity usage of South Africa being consumed by compressed air generation. This can further be extrapolated to the DCC having the potential to reduce the total electricity consumption of South Africa by up to 1%.

The yearly cost saving for the one case study was around R650 000 per annum. For the other case study, this was between R1 900 000 and R17 800 000 per annum. The payback period for the installation cost of the DCC will vary depending on the number of already installed automated control equipment present on the compressed air ring.

The DCC could potentially increase the maintenance required on compressors by increasing the number of times compressors are started and stopped. But, the financial impact of the increase in maintenance is negligible when compared with financial savings obtained by reducing the energy consumed [35]. On Case Study 1, the maintenance should not increase at all due to roughly the same number of starts and stops between the current controller and the DCC.

Comparing the difference in electricity savings between the two projects, the effect of a previously installed network, an optimised network or a combination of the two can be seen. On Case Study 2, there has never been a dynamic compressed air controller installed. If the control is optimised, the potential gains will be very small as was proved by Case Study 1. Theoretically, the network can be maintained by trained personnel and it stays optimised. But even for trained personnel it is hard to establish when the network is not optimised anymore.

While the DCC was never allowed to run fully autonomously, it successfully ran in parallel with an existing control system. One final test for the DCC would be to run completely autonomously in control of a compressed air network.

5.2 Suggestions for future work

5.2.1 Introduction

In this section, future research in dynamic compressor control is discussed. The future research can be divided into two main areas – simulation efficiency and automated network configuration.

5.2.2 Simulation efficiency

The DCC was designed to take advantage of modern computer processor architecture by doing simulations in a multithreaded environment. The simulation process used in this thesis takes thousands of iterations to complete. Although the simulations were completed in less time than is required for compressor control, these simulation times can be optimised.

Optimisations can be made in the way initial pressure values are seeded for each node. If the initial seed is closer to the target and hill climbing is used when searching for the pressure, the simulations should be completed in a shorter time. This will also reduce the processor requirement for the DCC.

5.2.3 Automatic network configuration

While DCC can automatically adapt to changes in the network where compressed air pipes are closed off and forming separate networks, this is still a manual process. The DCC can only react to manual intervention or automatically adjusted valves with regard to network configuration.

If the DCC can simulate the effect of closing off sections and splitting the network, this process could also be optimised and automated. By doing this, it should be possible to reduce compressed air usage further and thus reduce the electricity required by the air network.

By implementing automatic network configuration, the compressed air network would require more simulations before finding a solution. Thus, in order for automatic network configuration to be effective, the simulation efficiency also needs to be increased to compensate for the increase in the number of simulations required.

References

[1] S. W. van Heerden, "Developing a dynamic control system for mine compressed air networks," M. Eng. Dissertation, Computer and Electronic Engineering, North-West University, Potchefstroom, 2014.

[2] United Nations Framework Convention on Climate Change, *Kyoto Protocol*. Kyoto, Japan: United Nations, 1997.

[3] United Nations Framework Convention on Climate Change, *Doha Amendment*. Doha, Qatar: United Nations, 2012.

[4] C. Breidenich, D. Magraw, A. Rowley, and J. W. Rubin, "The Kyoto Protocol to the United Nations Framework Convention on Climate Change," *Am. J. Int. Law*, vol. 92, no. 2, p. 315, Apr. 1998.

[5] M. R. Raupach, G. Marland, P. Ciais, C. Le Quéré, J. G. Canadell, G. Klepper, and C. B. Field, "Global and regional drivers of accelerating CO2 emissions.," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 104, no. 24, pp. 10288–93, Jun. 2007.

[6] A. Raval and V. Ramanathan, "Observational determination of the greenhouse effect," *Nature*, vol. 342, no. 6251, pp. 758–761, Dec. 1989.

[7] P. M. Vitousek, "Beyond global warming: Ecology and global change," *Ecology*, vol. 75, no. 7, p. 1861, Oct. 1994.

[8] L. Hughes, "Biological consequences of global warming: Is the signal already apparent?," *Trends Ecol. Evol.*, vol. 15, no. 2, pp. 56–61, Feb. 2000.

[9] M. M. Halmann and M. Steinberg, *Greenhouse Gas Carbon Dioxide Mitigation: Science and Technology*. Boca Raton, FL: Lewis Publishers, 1999.

[10] R. Spalding-Fecher and D. K. Matibe, "Electricity and externalities in South Africa," *Energy Policy*, vol. 31, no. 8, pp. 721–734, Jun. 2003.

[11] R. Eisenberg and D. G. Nocera, "Preface: Overview of the forum on solar and renewable energy.," *Inorg. Chem.*, vol. 44, no. 20, pp. 6799–801, Oct. 2005.

[12] S. Shafiee and E. Topal, "When will fossil fuel reserves be diminished?," *Energy Policy*, vol. 37, no. 1, pp. 181–189, Jan. 2009.

[13] M. Höök and K. Aleklett, "A review on coal-to-liquid fuels and its coal consumption," *Int. J. Energy Res.*, vol. 34, no. 10, pp. 848–864, Aug. 2010.

[14] H. Winkler, "Renewable energy policy in South Africa: Policy options for renewable electricity," *Energy Policy*, vol. 33, no. 1, pp. 27–38, Jan. 2005.

[15] A. Eberhard, "The future of South African coal: Market, investment, and policy challenges," *Progr. Energy Sustain. Dev.*, 2011.

[16] World Resources Institute, "World Resources Institute CAIT country GHG emmisions." [Online]. Available: http://www.wri.org/resources/data-sets/cait-historical-emissions-data-countries-us-states-unfccc. [Accessed: 09-Jul-2015].

[17] U.S. Energy Information Administration, "International energy statistics," 2013. [Online]. Available: http://www.eia.gov/cfapps/ipdbproject/iedindex3.cfm?tid=1&pid=1&aid=2&cid=all,&syid=2008&eyid=2012&unit=TST. [Accessed: 18-Jan-2015].

[18] E. Ziramba, "The demand for residential electricity in South Africa," *Energy Policy*, vol. 36, no. 9, pp. 3460–3466, Sep. 2008.

[19] H. Amusa, K. Amusa, and R. Mabugu, "Aggregate demand for electricity in South

Africa: An analysis using the bounds testing approach to cointegration," *Energy Policy*, vol. 37, no. 10, pp. 4167–4175, Oct. 2009.

[20] R. Inglesi, "Aggregate electricity demand in South Africa: Conditional forecasts to 2030," *Appl. Energy*, vol. 87, no. 1, pp. 197–204, Jan. 2010.

[21] H. Winkler, O. Davidson, A. Kenny, G. Prasad, J. Nkomo, D. Sparks, M. Howells, and T. Alfstad, "Energy policies for sustainable development in South Africa," *Energy Sustain. Dev.*, vol. 11, no. 1, pp. 26–34, Mar. 2007.

[22] R. Inglesi-Lotz and J. N. Blignaut, "South Africa's electricity consumption: A sectoral decomposition analysis," *Appl. Energy*, vol. 88, no. 12, pp. 4779–4784, Dec. 2011.

[23] R. Christie, *Electricity, Industry and Class in South Africa*. Albany, NY: State University of New York: New York Press, 1984.

[24] Department of Minerals and Energy, "No. 31741," *Gazette*.

[25] P. Craven, "COSATU's comments to Nersa on Eskom's Multi Year Price Determination 3 (MYPD3) Application," COSATU, Johannesburg, 03-Nov-2012.

[26] H. Winkler and A. Marquand, "Changing development paths: From an energy-intensive to low-carbon economy in South Africa," *Clim. Dev.*, vol. 1, no. 1, pp. 47–65, Mar. 2009.

[27] Q. Bronkhorst, "The biggest companies in South Africa," *BusinessTech*, 2014.

[28] K. Fisher-Vanden, G. H. Jefferson, H. Liu, and Q. Tao, "What is driving China's decline in energy intensity?," *Resour. Energy Econ.*, vol. 26, no. 1, pp. 77–97, Mar. 2004.

[29] City of Cape Town, "No TitleState of energy and energy futures report," Cape Town, 2011.

[30] Departmnent of Energy, "Draft 2012 integrated energy planning report," South Africa, 2012.

[31] J. R. N. Tshisekedi, "Energy consumption standards and costs in South African gold and platinum mines," University of Witwatersrand, 2008.

[32] P. Fraser, "Saving energy by replacing compressed air with localized hydropower systems: A 'half level' model approach," *South. African Inst. Min. Metall.*, 2008.

[33] J. H. Marais, M. Kleingeld, and R. Pelzer, "Increased energy savings through a compressed air leakage documentation system," *Ind. Commer. Use Energy, Cape Penins. Univ. Technol.*, 2009.

[34] Cameron compression Systems, "Turbo-air 3000 centrifugal compressor." Houston, Texus, 2007.

[35] A. J. M. van Tonder, "Automation of compressor networks through a dynamic control system," PhD thesis, Electrical Engineering, North-West University, Potchefstroom, 2014.

[36] U.S. Department of Energy, *Improving compressed air system performance: a sourcebook for industry*. Washington, DC: U.S. Department of Energy, 1998.

[37] C.-Y. Yang, "Friction characteristics of water, R-134a, and air in small tubes," *Microscale Thermophys. Eng.*, vol. 7, no. 4, pp. 335–348, Jan. 2003.

[38] W. Jiang, J. Khan, and R. A. Dougal, "Dynamic centrifugal compressor model for system simulation," *J. Power Sources*, vol. 158, no. 2, pp. 1333–1343, Aug. 2006.

[39] L. Padachi, G. D. Bolt, M. Kleingeld, J. H. Marais, and E. H. Matthews, "Energy

efficiency through the integration of compressed air supply control with air network demand control," *Ind. Commer. Use Energy, Cape Penins. Univ. Technol.*, 2009.

[40] A. J. Schutte, M. Kleingeld, and J. Vosloo, "Various procedures to reduce a mine's compressed air usage," *Ind. Commer. Use Energy, Cape Penins. Univ. Technol.*, 2011.

[41] A. J. Schutte and M. Kleingeld, "Case studies of optimised c ompressed air usage on gold mines," *Ind. Commer. Use Energy, Cape Penins. Univ. Technol.*, 2009.

[42] R. Brake and G. Bates, "Criteria for the design of emergency refuge stations for and underground metal mine," *Australas. Inst. Min. Metall. AusIMM Proc.*, vol. 304, no. 2, 1999.

[43] R. Saidur, N. A. Rahim, and M. Hasanuzzaman, "A review on compressed-air energy use and energy savings," *Renew. Sustain. Energy Rev.*, vol. 14, no. 4, pp. 1135–1153, May 2010.

[44] A. Whillier, "Hydraulic rock drills and their effect on the underground environment," *SAIMM conf.*, 1975.

[45] J. I. Phillips, P. S. Heyns, and G. Nelson, "Rock drills used in South African mines: a comparative study of noise and vibration levels.," *Ann. Occup. Hyg.*, vol. 51, no. 3, pp. 305–10, Apr. 2007.

[46] D. Ndaba, "Electric rock drills invade precious-metal mines," *Mining Weekly*, South Africa, 2007.

[47] J. H. Marais, "An integrated approach to optimise energy consumption of mine compressed air systems," PhD thesis, Electrical Engineering, North-West University, Potchefstroom, 2012.

[48] H. P. R. Joubert, J. F. van Rensburg, and R. Pelzer, "Energy efficiency opportunities resulting from splitting a compressed air ring," *Ind. Commer. Use Energy, Cape Penins. Univ. Technol.*, 2011.

[49] J. Venter, "Development of a dynamic centrifugal compressor selector for large compressed air networks in the mining industry," M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2012.

[50] J. W. Lodewyckx, "Investigating the effects of different DSM strategies on a compressed air ring," M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2007.

[51] G. P. Heyns, "Challenges faced during implementation of a compressed air energy savings project on a gold mine," Nort West University, 2014.

[52] M. H. P. van Niekerk, "The implementation of a dynamic air compressor selector system in mines," M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2014.

[53] S. A. Korpela, *Principles of Turbomachinery*. Hoboken, NJ: John Wiley, 2011.

[54] D. Mills, V. Agarwal, and M. Jones, *Handbook of Pneumatic Conveying Engineering*, vol. 165. New York City, NY: CRC Press, 2004.

[55] P. Hanlon, *Compressor Handbook*. New York City, NY: McGraw-Hill, 2001.

[56] General Electric, "General Electric, Roots API-617 and API-672 IGCH centrifugal compressor." [Online]. Available: http://www.ge-energy.com/products_and_services/products/compressors/roots_igch_centrifugal_compressor.jsp.

[57] R. N. Brown, *Compressors: Selection and Sizing*. Houston, TX: Gulf Professional Publishing, 1997.

[58] S. Campanari, G. Manzolini, and F. Garcia de la Iglesia, “Energy analysis of electric vehicles using batteries or fuel cells through well-to-wheel driving cycle simulations,” *J. Power Sources*, vol. 186, no. 2, pp. 464–477, Jan. 2009.

[59] M. P. Boyce, *Centrifugal Compressors: A Basic Guide*, 1st ed. Tulsa, OK: Penwell, 2003.

[60] Atlas Copco, “Driving centrifugal compressor technology.” Atlas Copco, Cologne, Germany, 2011.

[61] J. M. Sorokes, “Selecting a centrifugal compressor.” American Institute of Chemical Engineers: CEP, pp. 44–51, Jun., 2013.

[62] Elliot Group, “Multi stage centrifugal compressors.” Elliot Group, Jeannette, PA, 2001.

[63] H. Cohen, G. F. C. Rogers, and H. I. H. Saravanamuttoo, *Gas Turbine Theory*, 4th ed. Harlow Essex, United Kindom: Longman Group, 1996.

[64] W. E. Forsthoffer, *Forsthoffer’s Rotating Equipment Handbooks Volume 3: Compressors*. Oxford, United Kingdom: Elsevier Ltd, 2005.

[65] R. S. R. Gorla and A. A. Khan, *Turbomachinery Design and Theory*. New York City, NY: Marcel Dekker, 2003.

[66] BorgWarner, “Performance Turbocharger Catalog.” BorgWarner, Auburn Hills, MI, 2014.

[67] A. Bhatia, *Compressed air systems*. FairFax, VA: PDH Center, 2012.

[68] W. Booysen, J. N. du Plessis, and J. F. van Rensburg, “Development of an energy management solution for mine compressor systems,” *Ind. Commer. Use Energy*, 2010.

[69] L. A. Els, “Load management on a municipal water treatment plant,” M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2014.

[70] H. Neser, “energy savings through the automatic control of underground compressed air demand.,” M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2008.

[71] J. W. Rautenbach, “Engineering a novel automated pump control system for the mining environment,” PhD thesis, Mechanical Engineering, North-West University, Potchefstroom, 2007.

[72] A. Botha, “Optimising the demand of a mine water reticulation system to reduce electricity consumption,” M. Eng. Dissertation, Electrical and Electronic Engineering, North-West University, Potchefstroom, 2010.

[73] A. L. Meek, “A systems engineering approach to improve the measurement and verification process of energy services companies,” M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2013.

[74] A. de Coning, “Sustained energy performance on compressed air systems for expanding gold mines,” M. Eng. Dissertation, Computer and Electronic Engineering, North-West University, Potchefstroom, 2013.

[75] C. Swart, “Optimising the operation of underground mine refrigeration plants and ventilation fans for minimum electricity cost,” Potchefstroomse Universiteit vir Christelike Hoer Onderwys, 2003.

[76] A. N. Hasan, “Maximising load shifting results with minimal impact on water quality,” M. Eng. Dissertation, Electricall Engineering, North-West University, Potchefstroom, 2009.

[77] H. P. R. Joubert, G. Bolt, and J. F. van Rensburg, “Energy savings on mining compressed air networks through dedicated process plant compressors,” *Ind. Commer. Use Energy, Cape Penins. Univ. Technol.*, 2012.

[78] R. Joubert and L. Liebenberg, “Cost and time efficient DSM on mine compressed air systems,” *Ind. Commer. Use Energy, Cape Penins. Univ. Technol.*, 2010.

[79] C. F. Scheepers, “Implementing energy efficiency measures on the compressed air network of old South African mines.,” M. Eng. Dissertation, Electrical Engineering, North-West University, Potchefstroom, 2011.

[80] J. I. G. Bredenkamp, “Reconfiguring mining compressed air networks for cost savings,” M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2014.

[81] W. F. van Niekerk, “The value of simulation models for mine DSM projects,” M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2012.

[82] H. P. Bloch, *A Practical Guide to Compressor Technology*. Hoboken, NJ: John Wiley, 2006.

[83] T. Giampaolo, *Compressor Handbook: Principles and Practice*. London, United Kingdom: Taylor & Francis, 2010.

[84] C. Hansen, “Dynamic simulation of compressor control systems,” M. Bsc., Aalborg University, Denmark, 2008.

[85] M. T. Gresh, “Preface,” in *Compressor Performance*, 2nd ed., Waltham, MA: Elsevier, 2001.

[86] R. C. Pampreeen, “The use of variable guide vanes for automotive gas turbine engine augmentation and load control,” *Soc. Automot. Eng.*, 19736.

[87] J. N. du Plessis and R. Pelzer, “Development of an intelligent control system for mine compressor systems,” *Ind. Commer. Use Energy, Cape Penins. Univ. Technol.*, 2011.

[88] P. N. Ananthanarayanan, *Basic Refrigeration and Air Conditioning*. New Delhi, India: McGraw-Hill Professional, 2005.

[89] Pneu Logic, “PL4000: Compressed air master control system,” *Pneu Logic*, Portland, OR, 2010.

[90] BOGE Compressed Air Systems, “Easy to use, hard to beat: airtelligence provis 2.0 master control.” BOGE compressed air systems, Bielefeld, Germany, 2014.

[91] Atlas Copco, “Control solutions: ES - the latest in control technology.” Atlas Copco, Stockholm, Sweden, 2012.

[92] Murphy, “Centurion^{TM} PLUS Full Featured Controller.” Murphy, Tulsa, OK, 2010.

[93] CompAir, “SmartAir Master Compressed Air Management System.” CompAir, Princeton, IL, 2010.

[94] J. N. du Plessis, “Development of an energy management solution for mine compressor systems,” M. Eng. Dissertation, Computer and Electronic Engineering, North-West University, Potchefstroom, 2010.

[95] F. W. Schroeder, “Energy Efficiency Opportunities in Mine Compressed Air Systems,” M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom,

2009.

[96] KYPipe, “Pipe2014.” KYPipe, Cary, NC, 2014.

[97] R. Shelquist, *An Introduction to Air Density and Density Altitude Calculations*. Longmont, CO: Shelquist Engineering, 2015.

[98] J. A. Schetz and A. E. Fuhs, *Fundamentals of Fluid Mechanics*. New York City, NY: John Wiley, 1999.

[99] D. R. Durgaiah, *Fluid Mechanics and Machinery*. New Delhi, India: New Age International, 2002.

[100] J. F. Wendt, J. D. Anderson, E. Dick, G. Degrez, R. Grundmann, J. Degroote, and J. Vierendeels, *Computational Fluid Dynamics*. Berlin, Germany: Springer, 2009.

[101] J. Kiijarvi, “Darcy friction factor formulae in turbulent pipe flow,” *Lunowa Fluid Mech.*, 2011.

[102] C. F. Colebrook, “Turbulent flow in pipes, with particular reference to the transition region between the smooth and rough pipe laws.,” *J. ICE*, vol. 11. pp. 133–156, 1939.

[103] S. E. Haaland, “Simple and explicit formulas for the friction factor in turbulent pipe flow,” *Journal of Fluids Engineering*, vol. 105. p. 89, 1983.

[104] A. J. M. van Tonder, “Sustaining compressed air DSM project savings using an air leakage management system,” M. Eng. Dissertation, Electric and Electronic Engineering, North-West University, Potchefstroom, 2010.

[105] P. Goosen, “Efficient monitoring of mine compressed air savings,” M. Eng. Dissertation, Computer and Electronic Engineering, North-West University, Potchefstroom, 2013.

[106] B. R. Munson, D. F. Young, and T. H. Okiishi, *Fundamentals of fluid mechanics*. Hoboken, NJ: J. Wiley & Sons, 2006.

[107] P. J. Deitel and H. M. Deitel, *Java, How to Program*. Upper Saddle River, NJ: Prentice Hall, 2007.

[108] M. Gendreau and J. Potvin, *Handbook of Metaheuristics*. London, United Kingdom: Springer, 2010.

[109] J. A. Deysel, “DSM strategies to reduce electricity costs on platinum mines,” M. Eng. Dissertation, Mechanical Engineering, North-West University, Potchefstroom, 2014.

[110] A. M. Davis, “Software prototyping,” *Adv. Comput.*, vol. 40, 1995.

[111] “Software development process.” [Online]. Available: http://lauyin.com/NPD/model/prototype.html. [Accessed: 01-Jan-2015].

[112] M. F. Smith, *Software Prototyping*. London, United Kingdom: McGraw-Hill, 1991.

[113] S. P. Overmyer, “Revolutionary vs. evolutionary rapid prototyping: Balancing software productivity and HCI design concerns,” *Int. Conf. Human-Comput.er Interact, Stuttgart, Ger.*, 1991.

[114] S. J. Russell, P. Norvig, and E. Davis, *Artificial intelligence*. Upper Saddle River, NJ: Prentice Hall, 2010.

[115] E. Burke and G. Kendall, *Search Methodologies*. Nottingham, United Kingdom: Springer, 2005.

[116] G. Gagn, “Additional searching strategies,” 2014. [Online]. Available: http://people.westminstercollege.edu/faculty/ggagne/fall2014/301/chapters/chapter4/in

dex.html. [Accessed: 03-Feb-2015].

[117] S. H. Jacobson and E. Yücesan, “Analyzing the performance of generalized hill climbing algorithms,” *J. Heuristics*, vol. 10, pp. 387–405, 2004.

[118] A. Clements, *Principles of Computer Hardware*. New York City, NY: Oxford University Press, 2006.

Appendix A: Software procedures

A.1. Software prototyping

Software prototyping is a method of software development whereby small prototypes of software are made – each testing different elements of the final software [110]. This allows the software to undergo changes while not changing the complete software. Each prototype is tested individually and modified thereafter. The process can be seen in Figure 86 [111].

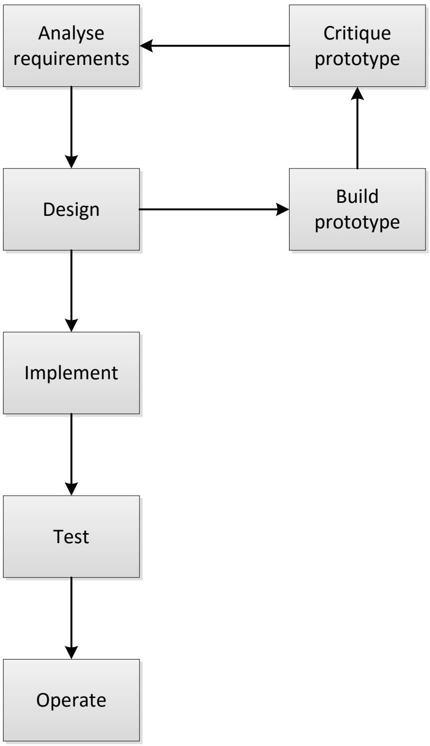

Figure 86: Software prototyping

Evolutionary prototyping is a form of software prototyping [112]. This differs from normal software prototyping in that there is only one prototype. This prototype begins with the design of the core of the software and expands to other less known features. As features are added, they are tested and refined. Features that are well understood are added and refined first to limit changes later on in the software.

Prototyping is used because it reduces the cost and time used to develop software [113]. This is done by improving the quality of the requirements. By improving the requirements, changes can be made faster thus, limiting the changes needed later on in the life of the software.

A.2. Genetic algorithms

Genetic algorithms are used to perform searches and are inspired by natural selection and genetics [114]. Genetic algorithms work by mutating a solution from a parent to create a new child offspring. By mutating, the idea is to generate a new better solution than the parent by combining the best properties of the parents.

Genetic algorithms generate a population of possible solutions. The best solutions are then chosen from all possible solutions. These solutions would then be mutated with each other to generate a new population of solutions. If the population generated is too small the genetic algorithm will never find the optimal solution. But, if the population generated is too large, the algorithm can occupy unnecessary computational time.

A key limitation in using a genetic algorithm to find a solution is that it must be possible to compare solutions. The steps for genetic algorithms are the following:

- Initialisation: An initial population is generated to start the process.
- Evaluation: Solutions are evaluated and assigned a fitness value.
- Selection: Select the solutions with the highest fitness number to be the new parents.
- Recombination: Combine parents to generate new offspring.
- Mutation: Add random new changes to the offspring to facilitate more change in the offspring.
- Replacement: Replace the original population with the new population.

Apart from the initialisation, these steps are repeated until an end condition is met. There are many different methods to do recombination, mutation and replacement. For example, in replacement, the entire population can be replaced, only certain solutions can be replaced or none can be replaced. Figure 87 [115] shows an example of three crossover methods.

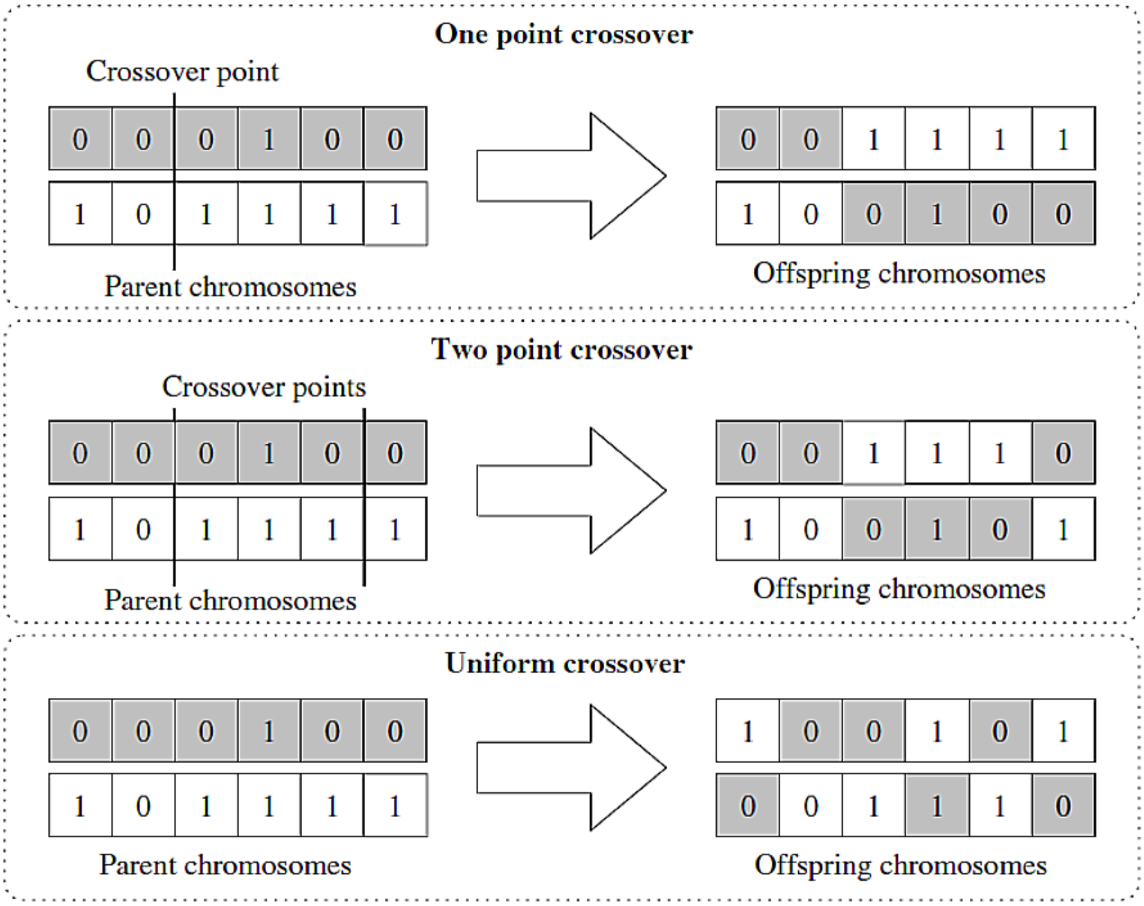

Figure 87: Genetic algorithms crossover methods

A.3. Hill climbing

Hill climbing search is a search algorithm that searches for an increasing value and tries to find a peak value [108]. One of the many problems relating to hill climbing is local maximums. The highest value might only be a local maximum and not a global maximum. Hill climbing has to overcome problems such as local maximums, shoulders and flat local maximums to find the global maximum. Figure 88 [116] shows these problems.

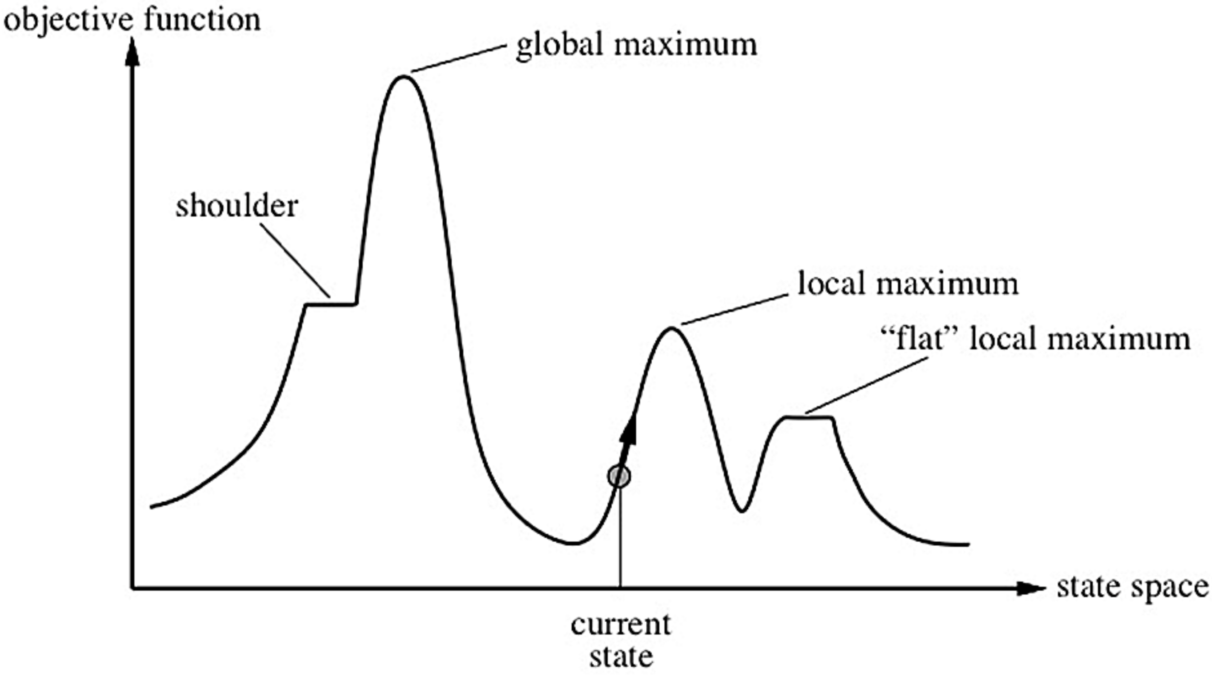

Figure 88: Hill climbing

According to Jacobson and Yücesan a random-restart hill climbing algorithm can quickly find a near optimal solution while using limited resources [117]. Random-restart hill climbing or shotgun hill climbing [114] restarts the search at a random local to try and find a different local maximum.

A.4. Processor pipelines

Reduced Instruction Set Computer (RISC) processors work with the bases of a pipeline [118]. The idea behind pipelining is that each instruction is split into certain parts. In the most basic pipeline example, this can be explained by each instruction having: instruction fetch (IF), instruction decode (ID), operand fetch (OP), execute (E), operand store (OS). Figure 89 [118] shows this visually.

Figure 89: RISC pipeline

By using a pipeline and splitting each instruction into separate instructions, throughput can be increased. If two instructions were to be executed on a CPU without pipelining, the CPU would take 10 clock cycles to complete the instructions. A CPU with pipelining would only take a minimum of six and a maximum of 10 clock cycles. This is because while the CPU is decoding the first instruction, it is already fetching the second instruction. Figure 90 shows the example visually.

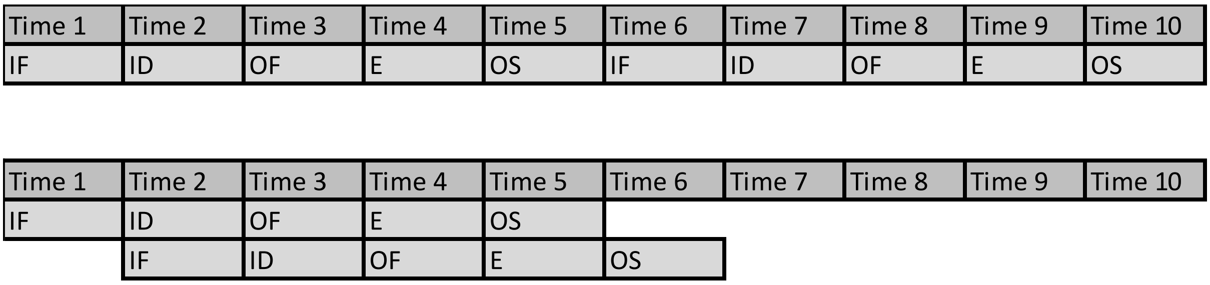

Time 1	Time 2	Time 3	Time 4	Time 5	Time 6	Time 7	Time 8	Time 9	Time 10
IF	ID	OF	E	OS	IF	ID	OF	E	OS

Time 1	Time 2	Time 3	Time 4	Time 5	Time 6	Time 7	Time 8	Time 9	Time 10
IF	ID	OF	E	OS					
	IF	ID	OF	E	OS				

Figure 90: Pipeline comparison

A problem with pipelining comes with branching and dependence. Going back to the above-mentioned example, if the second instruction required the data of the first, then the pipeline would stall on the operand fetch while the first instruction completes. Branching stops a pipeline at the instruction fetch. This is because the CPU does not know which instruction to fetch.

Modern CPUs have branch prediction to circumvent this. With branch prediction, the CPU will predict which way instruction to load. If the branch prediction was correct, the CPU would not have lost any time. If the branch prediction was incorrect, the CPU would have to drop the current instruction and reload the correct instruction. This would give the same result as if the CPU would have waited for the instruction to complete and only load the correct instruction.

Appendix B: Case study 1- extra compressor running graphs

This appendix includes extra daily graphs that show the running compressors in detail as was simulated on the mine mentioned in Case Study 1. As mentioned in Section 4.3.3, Van Heerden also encountered this problem [1]. The cycling present in the graphs can be explained by the fact that the DCC reacts to the actual compressor controller on the mine. In some situations, the DCC would have substituted compressors, closed or issued a stop command to the compressor.

Because the actual compressor did not stop, the DCC sees that it is running and considers this when simulating. Thus, it may decide to keep it running the next time the DCC makes a decision. This is because the DCC prefers energy efficiency and only takes action if the compressors reduce electricity consumption by a certain margin.

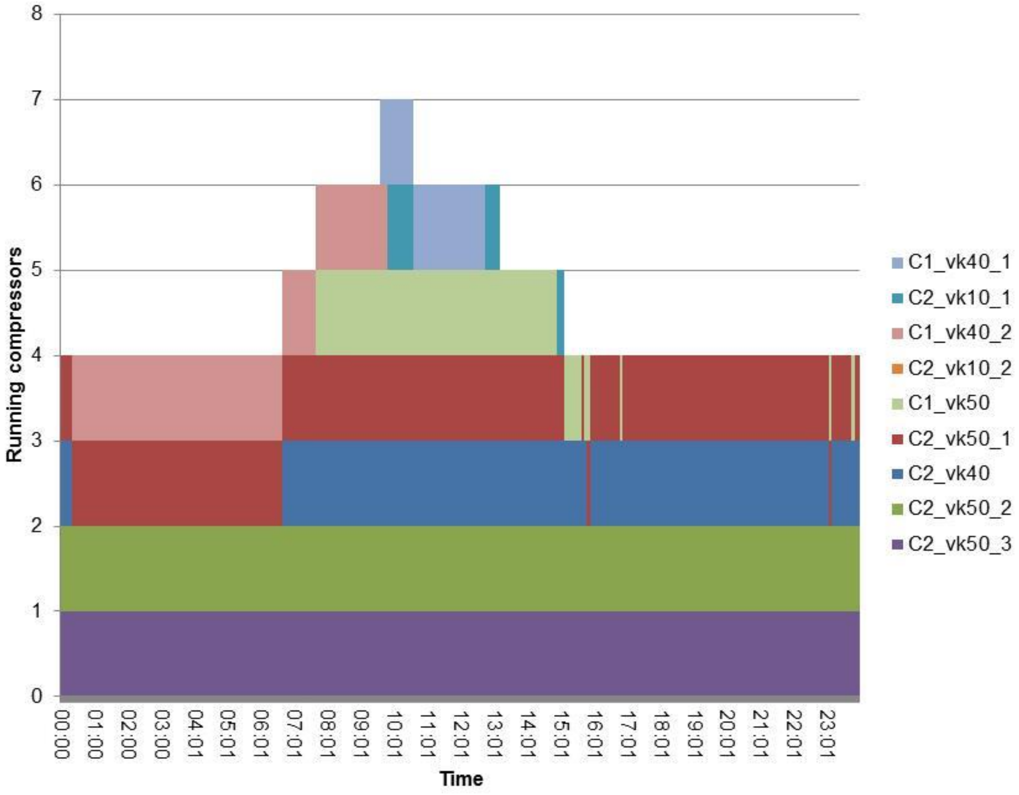

Figure 91: Simulated running compressors for day 1

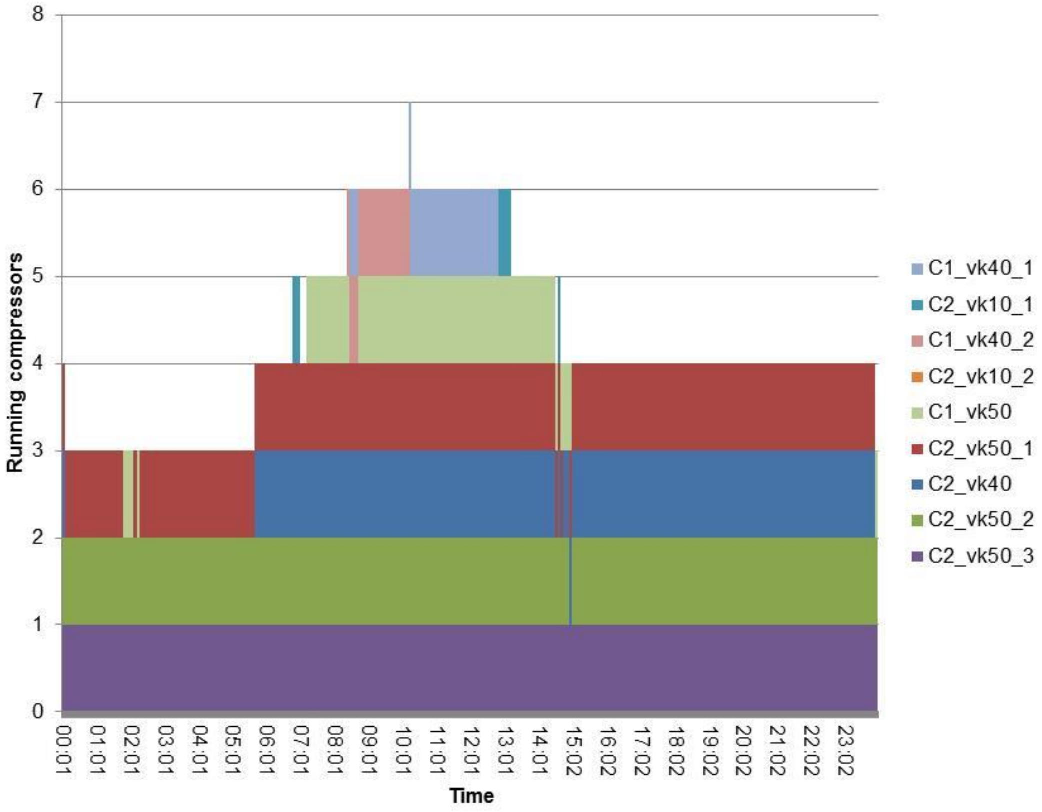

Figure 92: Simulated running compressors for day 2

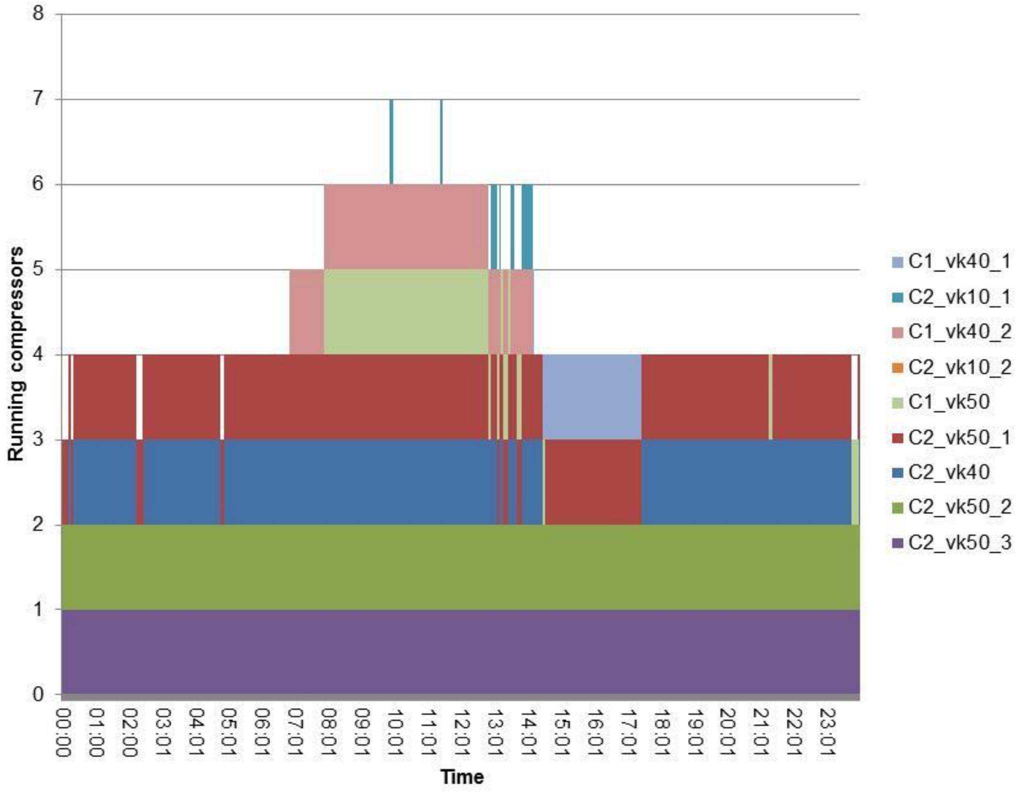

Figure 93: Simulated running compressors for day 3

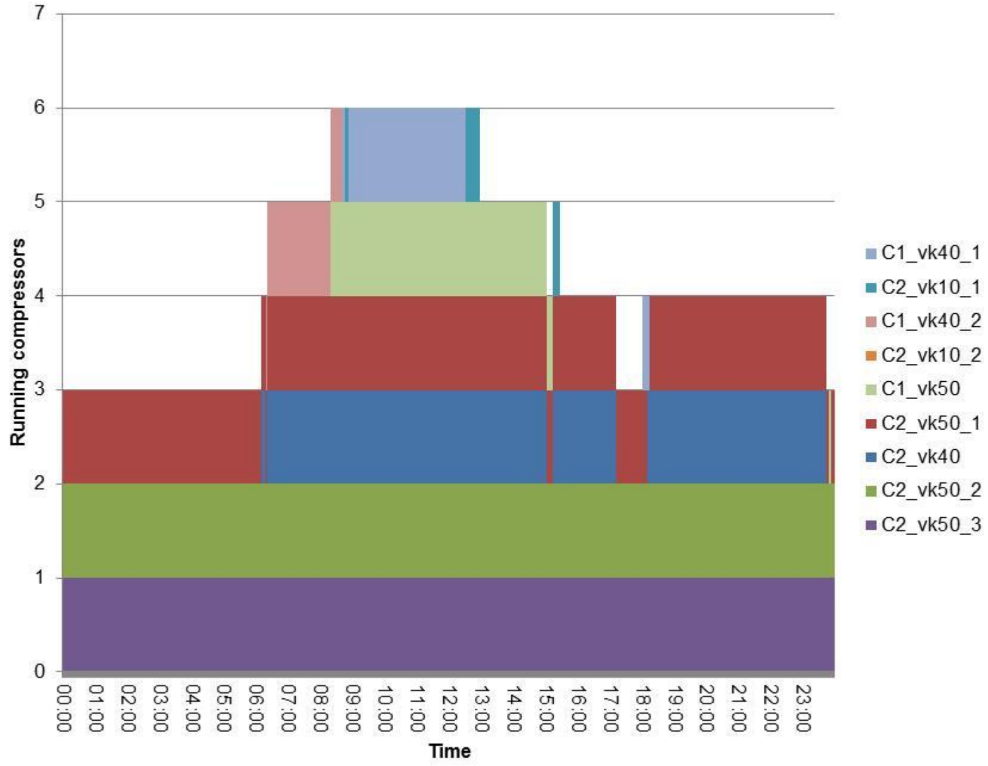

Figure 94: Simulated running compressors for day 4

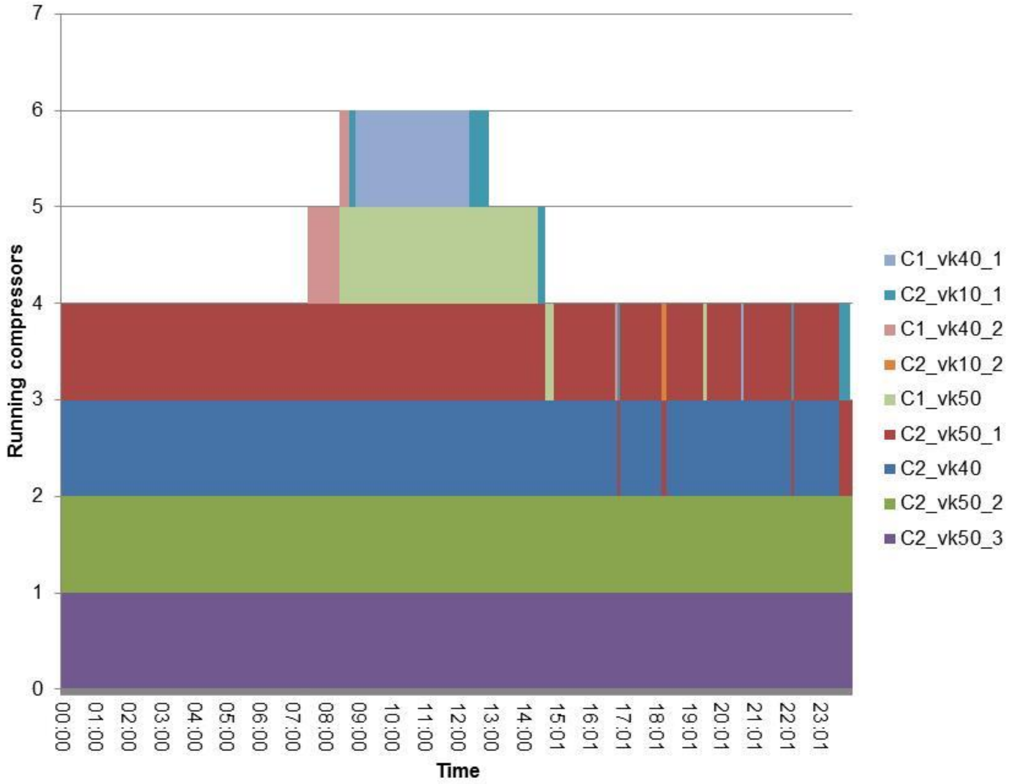

Figure 95: Simulated running compressors for day 5

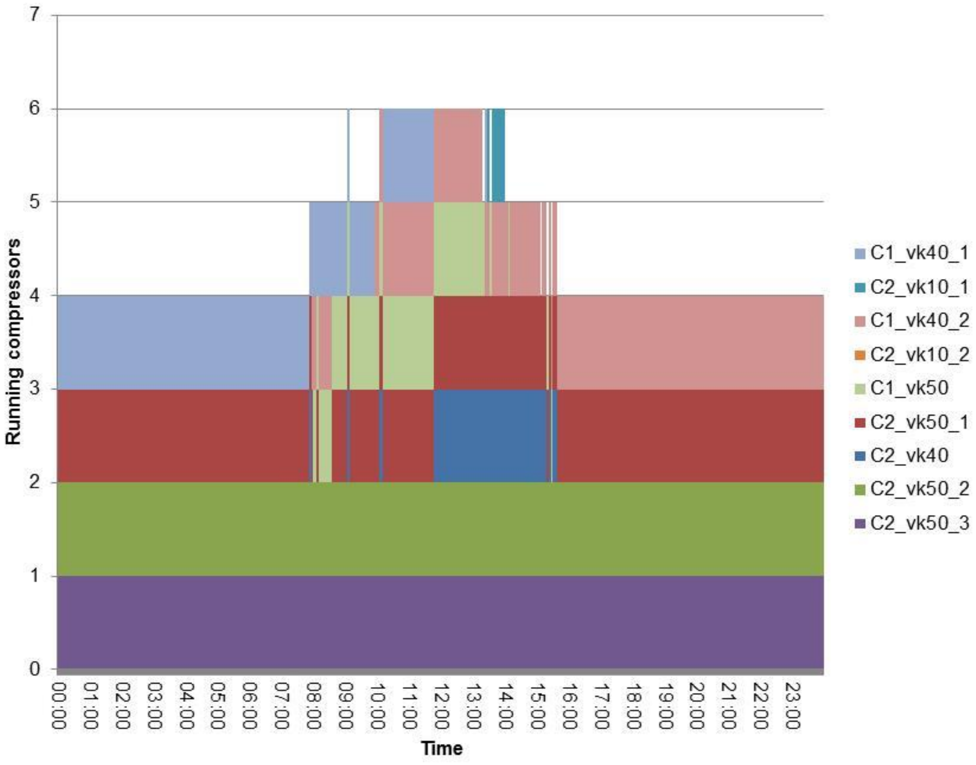

Figure 96: Simulated running compressors for day 6

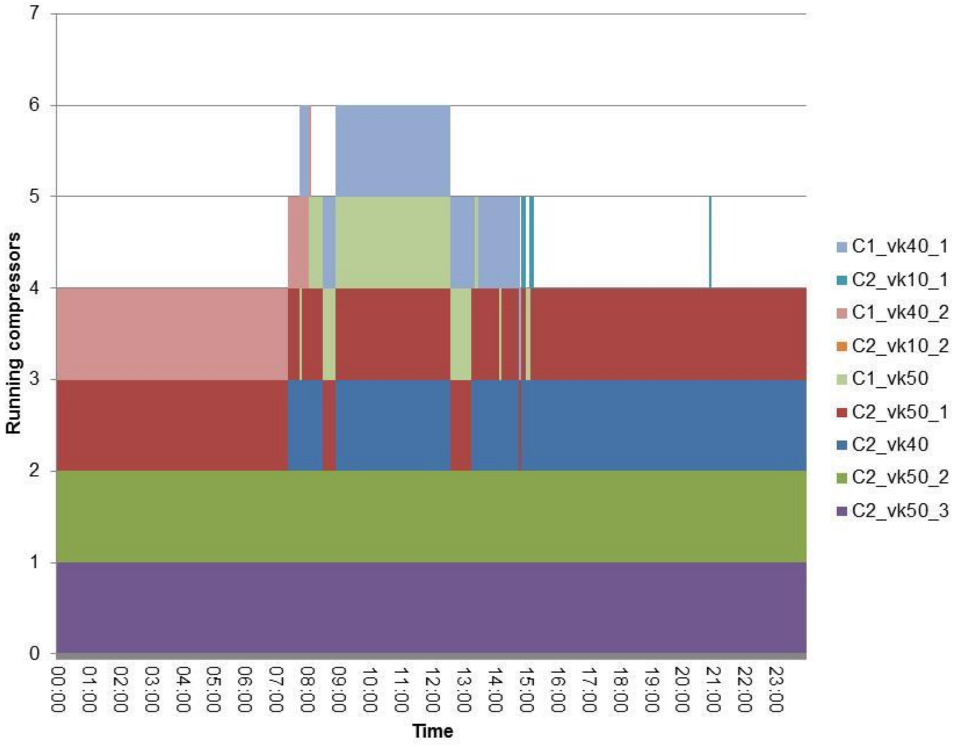

Figure 97: Simulated running compressors for day 7