Options for the Use of Machine Learning in Non-Daily Metered Gas Allocation

Contents

Background

The Unidentified Gas Task Force

Following the implementation of the Project Nexus suite of process and systems changes in June 2017, the arrangements for daily allocation of gas changed, introducing the new concept of daily UIG as the balancing factor in each LDZ.

In early 2018, gas industry participants were still experiencing considerable issues due to the level and variability of daily UIG, which could vary significantly from day to day, across LDZs and between Nominations (forecasts) and Allocations (actuals). There was growing concern about the levels and costs of UIG, as well as uncertainty about the level of “final” UIG that would be seen at Line-in-the-Sand after (almost) all meter points had received a meter point reconciliation.

In May 2018 Total Gas & Power raised UNC Modification 0658 (Urgent): “CDSP to identify and develop improvements to LDZ settlement processes”. The Modification was approved by Ofgem under Urgent arrangements with effect from 6th July 2018. The Modification introduced a new Service Line into the Data Services Contract (which governs Xoserve’s role a Central Data Services Provider – the “CDSP”) with the following outputs: “Investigate and report on the causes of and contributors to Unidentified Gas; and suggest to industry measures by which the levels and/or volatility of Unidentified Gas might be reduced

Following approval of the UNC Modification and Change Proposal, Xoserve promptly set up an Unidentified Gas Task Force, which was organised into two main streams: Issues Analysis and Advanced Analytics. The Advanced Analytics stream brought in an external Analytics Consultancy partner with in-depth experience in **Machine Learning** to undertake high-volume complex data investigations.

What is Machine Learning?

Machine learning is a branch of artificial intelligence which involves complex computer systems teaching themselves how to perform tasks without being explicitly programmed to do so, usually based on input data and observed outcomes. The process of providing the input data to the Machine Learning system so that it can “learn” the relationships between historic data and results is often referred to as “training” and that term is used in this paper to describe the “learning” phase of the process.

Advanced Analytics Stream

The Advanced Analytics stream undertook a wide range of investigations into the effectiveness of the industry Algorithm that predicts daily consumption for Non-Daily Metered supply points (“the NDM Algorithm”). The findings were published on Xoserve.com as the project progressed, along with any specific recommendations. A number of recent gas industry initiatives, such as the introduction of additional End User Categories and the use of additional weather data in the Composite Weather, can be traced back to the outputs of the Machine Learning analysis.

The Advanced Analytical stream used past data on individual sites' gas consumption and possible influencing factors (including a wide range of weather data items) to try to build predictive models for gas demand at site level that could give better results than the current NDM Algorithm.

The Role of the NDM Algorithm

Although there are now several million Smart gas meters in operation in the GB market, very few daily Smart meter reads find their way into the gas allocation processes. The vast majority (more than 99%) of GB gas meter points are Non-Daily Metered ("NDM"), meaning that they are only expected to submit an actual meter reading into central settlement processes after one, six or twelve months.

As part of the daily gas allocation processes, all those sites receive an estimated share of daily gas used on the network. This estimate is calculated using a formula (the NDM Algorithm). The algorithm has over 500 separate models each year which are based on a simple regression analysis between historic observed reactions at a sample of individual sites and factors like geographical location (one of 13 regions), weather, day of the week and other generic factors. There is no sub-division by business type (e.g. school, office, hospital etc.) as the gas industry central systems do not hold that information. The models build a straight-line relationship between daily gas demand and weather (in the form of the Composite Weather Variable, which uses a number of daily weather measurements and other data items to improve the relationship to NDM demand).

For each gas day, the NDM Algorithm predicts gas usage at Shipper, End User Category (EUC) and LDZ level. The algorithm uses the relevant model for each EUC, the actual weather observations for that LDZ for the day, and the total AQ of that Shipper's customer base in that End User Category/LDZ combination. The model is used in forecast mode before the gas day, to produce NDM Nominations (a prediction of the Shipper's energy requirements for the following day) and in actual mode after the gas day, to determine the Shipper's share of gas usage for the day – the "NDM Allocation". Differences between Nominations and Allocations can result in costly gas buying errors, so it is important to the gas industry to have a good alignment between the two processes.

Whilst the models may work well in aggregate, especially for large numbers of typical domestic sites, they will not always work well at individual level. They may also fail to deal with non-weather related events, such as short-notice Bank Holidays, unseasonal weather or pandemics. Because daily gas usage must be balanced in total there is a concept of Unidentified Gas which is the balancing figure in each region (Local Distribution Zone – "LDZ"). As the NDM sector is far larger than the Daily Metered sector, any modelling inaccuracies or variations can cause large swings in Unidentified Gas. This adds significant uncertainty to the industry and makes daily and longer term gas buying decisions harder.

For these reasons, the Advanced Analytics stream focused on the NDM Algorithm and on investigating ways to improve its accuracy.

UIG Task Force Findings

The UIG Task Force findings were reported to the industry in fortnightly updates during the first phase of “Sprint” activity and then on a monthly basis. The Advanced Analytics stream tested a number of different Machine Learning techniques to assess whether they could reduce the overall levels of UIG and/or the volatility, i.e. the level of day-to-day variation, both of which were significant areas of industry concern.

During the first phases the analysis was focused on End User Category (“EUC”) 01 (sites with an Annual Quantity of 73,200 kWh or less). This category accounts for over 70% of the NDM market (in usage terms), and has far more available historic site level usage data.

The headline result was that the use of Neural Networks (a branch of machine learning which mimics the human brain by creating “networks” of results to identify patterns or connections) for EUC01 could give an **annual UIG reduction of 70% on average across a year, and a 20% reduction in UIG volatility**, using the measure of Standard Deviation. However it is important to highlight that those reductions are not consistent across all days, and that for some days the simulated UIG would have been higher than the actual outturn.

After the initial focus on EUC01, the Advanced Analytics stream began to look at the other EUCs for sites with large AQs. EUCs are defined by AQ and location only (not by type of business, as previously mentioned) and the higher EUCs typically have far fewer sites in them, so far less daily sample data is available. The current NDM models tend to aggregate data for multiple LDZs for modelling purposes. Machine Learning tends to work best with large volumes of data from which to learn, so small sample data pools present real issues.

Initial experiments with developing **national (GB) level models** did not give any improvement against the current models.

Tests which used the total quantity of energy input to an LDZ did give an improvement, but this approach would only work for gas Allocations after the Gas Day. As it relies on knowing the actual quantity of gas delivered to the LDZ it could not be used before the Gas Day in forecast mode, and a different model would be needed for forecasts (Nominations). As this would create more volatility between Nominations and Allocations, this is not considered a viable approach.

The final sets of investigations used the results of the Machine Learning models for EUCs 01 and 02 as an input to the modelling for the remaining seven EUCs, through a process of “transfer learning”. The best results came through using Sequential Neural Networks for EUC 01 and 02, and Functional Neural Networks for EUC 03 to 09.

A more complete summary of the outcomes of both streams of the UIG Task Force (Advanced Analytics and Issue Analysis) were presented to the DSC Contract Management Committee in March 2020 in the Change Completion Report[1].

---

[1] https://gasgov-mst-files.s3.eu-west-1.amazonaws.com/s3fs-public/ggf/2020-03/2.4%20CCR%20UIG%20Task%20Force_Mod%200658_XRN4695%20%28Part%201%20of%202%29.docx?2Nsp4__d_5Qdu6_GJybKV7e2CHWQwYkM=

Options for using Machine Learning in the NDM Algorithm

The UIG Task Force summarised the following four options for the use of Machine Learning in the NDM Algorithm in a presentation to the DSC Contract Management Committee in March 2020[2]:

1. "No Change" – retain the current arrangements to produce the annual deliverables for the NDM Algorithm (the Annual Load Profiles and Daily Adjustment Factors – the "ALPs and DAFs")
2. "Machine Learnt Annual Profiles" – use Machine Learning to develop the same annual deliverables as under the current arrangements (the ALPs and DAFs)
3. "Machine Learnt Annual Direct Outputs" – use Machine Learning to predict daily gas usage before and after the day, based on an annual training process, but without producing any ALPs and DAFs
4. "Continual Machine Learning with Direct Outputs" – use Machine Learning to predict daily gas usage before and after the day, but re-train the models more frequently, e.g. monthly, without producing any ALPs and DAFs

The four Options are compared in the next section.

Appendix 2 reviews how well each of the options would have coped with the disruption to gas demand patterns caused by the COVID-19 pandemic and mandatory lockdown. Although this is not part of any formal decision-making process, it may be useful to consider whether any of the options would have been better than others at coping with the impacts.

[2] Presentation "Machine Learning Options Overview" available on Joint Office of Gas Transporters website:
https://gasgov-mst-files.s3.eu-west-1.amazonaws.com/s3fs-public/ggf/2020-03/2.4%20UIG%20TaskForce%20%28supporting%20Information%20Machine%20Learning%20Options%20Overview%20%28paper%202%20of%202%29.pdf?s5uZ8gYyjfTRAU0rmMBKmXDQnjBOJ71B=

Comparison of Options for using Machine Learning in the NDM Algorithm

Option	Likely Benefits	Likely Industry Impacts	Implementation considerations
1. No Change	Demand Estimation process continues to produce the same well understood outputs. No changes required to CDSP or Shipper/Transporter systems.	No incremental improvement to UIG levels or volatility. Several initiatives are already under way, including the additional End User Categories for different Market Sectors and payment types for Gas Year 2019, and the use of solar radiation in the Composite Weather Variable from Gas Year 2020 onwards. These have not been in use long enough to assess the benefits to NDM Allocation and UIG.	No lead time to implement this option (no change option).

Comparison of Options for using Machine Learning in the NDM Algorithm

Option	Likely Benefits	Likely Industry Impacts	Implementation considerations
2. Machine Learnt Annual Profiles	Likely to improve UIG levels and/or volatility, based on previous analysis Harnesses benefits of Machine Learning, by using wider range of inputs Should not require any changes to core UKLink or Gemini systems Demand Estimation process will still produce a set of ALPs and DAFs for each Gas Day for each End User Category for use in daily Gas Allocation – these outputs are well understood and should not require any changes to Shipper/Transporter systems	Changes required to CDSP Demand Estimation modelling systems (but not to core UKLink or Gemini) The output of the Demand Estimation processes (the ALPs and DAFs) will be harder for industry parties to replicate or explain, compared to the current process which uses a simple regression analysis approach.	Requires (as a minimum) a change to the NDM Demand Estimation methodology document, which is owned by the UNC Demand Estimation Sub-Committee. Requires at least a 12 to 18 month lead time for implementation, assuming that the change would go live at the start of a Gas Year: time required to agree methodology and data items, develop, test and train models. If a parallel running phase is required, this would add at least 12 further months. Wider industry engagement also required to explain the change and the impacts. Probably does not require changes to UKLink or Gemini systems.

Comparison of Options for using Machine Learning in the NDM Algorithm

Option	Likely Benefits	Likely Industry Impacts	Implementation considerations
3. Machine Learnt Annual Direct Outputs	Likely to improve UIG levels and/or volatility, based on previous analysis. Harnesses benefits of Machine Learning, by using wider range of inputs, likely to give additional benefit compared to “Machine Learnt Annual Profiles” option, as it would not be constrained by the need to produce ALPs and DAFs.	A major change to daily NDM energy calculations – expectation is that all or part of these calculations would need to take place outside of Gemini systems with results being passed to passed to Gemini for inclusion in daily balancing Demand Estimation process will no longer produce ALPs and DAFs – these will not be available to industry parties for them to use in their own modelling/forecasting systems. The new separate forecasting/ allocating system could include a facility for industry parties (especially Shippers) to interact with the system to run simulations of future allocations under different scenarios.	Requires changes to UNC Section H, which sets out the NDM Algorithm, and to the NDM Demand Estimation methodology document, which is owned by the UNC Demand Estimation Sub-Committee. As changes would be needed to core CDSP systems, the lead times would be longer than Option 2, in addition to the time taken to develop and test a modelling system. Timescales would need to fit around the existing Gemini programme of change. Wider industry engagement also required to explain the change and the impacts.

Comparison of Options for using Machine Learning in the NDM Algorithm

Option	Likely Benefits	Likely Industry Impacts	Implementation considerations
4. Continual Machine Learning with Direct Outputs	Likely to improve UIG levels and/or volatility, based on previous analysis. Harnesses benefits of Machine Learning, by using wider range of inputs, likely to give additional benefit compared to “Machine Learnt Annual Profiles” option, as it would not be constrained by the need to produce ALPs and DAFs. Likely to give a relatively small incremental improvement over and above Option 3, assuming that consumer behaviours do not usually change very suddenly.	Consistent with Option 3, i.e. a major change to daily NDM energy calculations. ALPs and DAFs would not be available to industry parties for them to use in their own modelling/forecasting systems. The new separate forecasting/ allocating system could include a facility for industry parties (especially Shippers) to interact with the system to run simulations of future allocations under different scenarios. More frequent updates to the models might mean that Shippers would also want to update their simulations more frequently.	As for Option 3, this requires changes to UNC, and NDM Demand Estimation Methodology. The same changes would also be required to CDSP core systems. Implementation is likely to take several years.

Appendix 1 shows high level context diagrams for each option, to demonstrate which systems would be impacted, and where key calculations would take place.

Cost v Benefit implications of the Options

Each of Options 2 to 4 aims to achieve more accurate NDM Allocation than the current processes, and so reduce the average level of UIG and the day-to-day variation in UIG ("volatility"). These options will not reduce the total physical quantity of gas consumed each day, so there will be no overall reduction in GB carbon footprint. If UIG is lower as a result, this means that NDM Allocation on average will be higher. Depending on each Shipper's customer portfolio and the sharing rules for UIG, some Shippers might need to buy more gas on a day, others less. There may be benefits in lower marginal prices for gas bought each day, but they would depend on the liquidity of the market each day.

Overall, it is difficult to identify a positive financial cost benefit at industry level for any changes to the NDM Algorithm. Individual Shippers could see positive or negative financial benefits. If the NDM Algorithm becomes more accurate, less energy would be moving between Shippers due to meter point reconciliation, which should help Shippers with their forecasting and budgeting processes.

Other considerations

The NDM Demand models flow data into a number of other processes, which also need to be considered, e.g.

- AQ calculation, which uses a weather-adjusted version of the Annual Load Profiles to convert from a non-365 read period under actual weather conditions to an annual view of consumption under seasonal normal conditions
- Calculation of NDM Peak Load values (usually referred to as SOQs – System Offtake Quantities) which uses a standard Peak Load Factor which is also produced by the NDM Demand Modelling processes

Proposed Next Steps

The Uniform Network Code includes an obligation on the Demand Estimation Sub-Committee ("DESC") to review the NDM Algorithm every three years (H2.2.2). The current NDM Algorithm was introduced in June 2017 with the Project Nexus suite of changes and has not been reviewed since then.

DESC will plan its workload for the coming Autumn/Winter period in October 2020 and should include the possible use of Machine Learning when reviewing the suitability of the current NDM Algorithm.

The NDM Algorithm is critical to industry processes, such as daily gas allocation, AQ calculation and Peak Load calculations. We recommend that the wider industry is engaged in the review of the NDM Algorithm and the possible use of Machine Learning, either through an initial consultation phase, to understand the appetite for change (and what scale of change) or through a UNC Review Proposal, to allow wider engagement. DESC is an open meeting, but its work is often seen as too complex or specialist for non-members. The Committee needs to break down this perception and explain the proposals in layman's terms to ensure a wide engagement.

APPENDIX 1: CONTEXT DIAGRAMS

High level context diagrams are shown for each of the four options, to aid understanding.

Option 1 – No Change (ALPs and DAFs developed each year based on regression analysis)

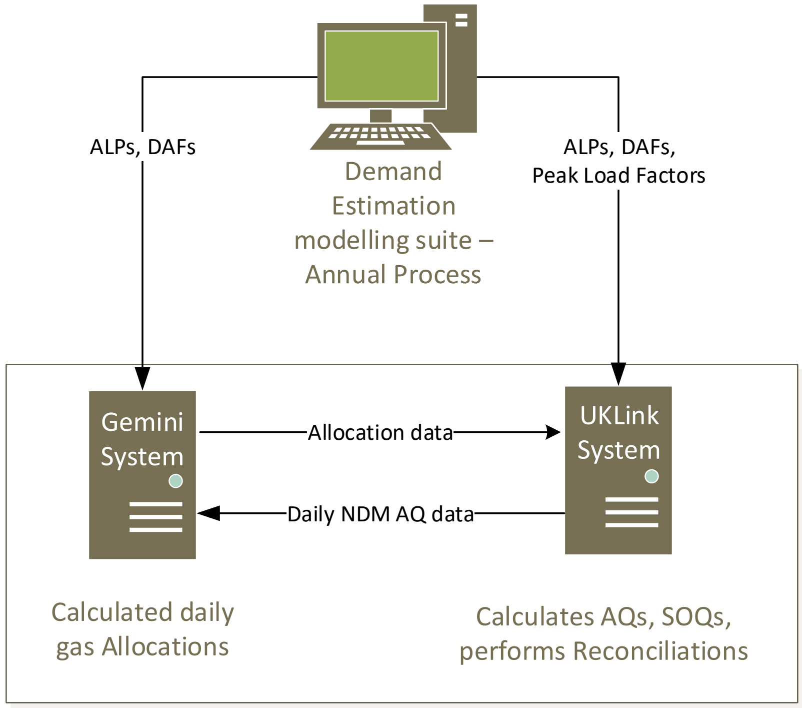

Option 2 – Machine Learning development of ALPs and DAFs each year

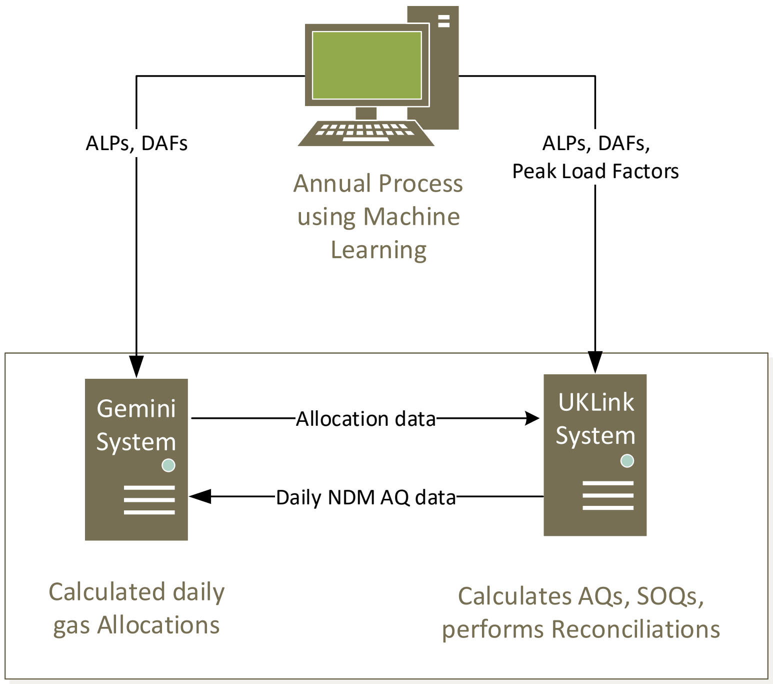

APPENDIX 1: CONTEXT DIAGRAMS

Option 3 – Annual Machine Learning with direct outputs into the Allocation Systems

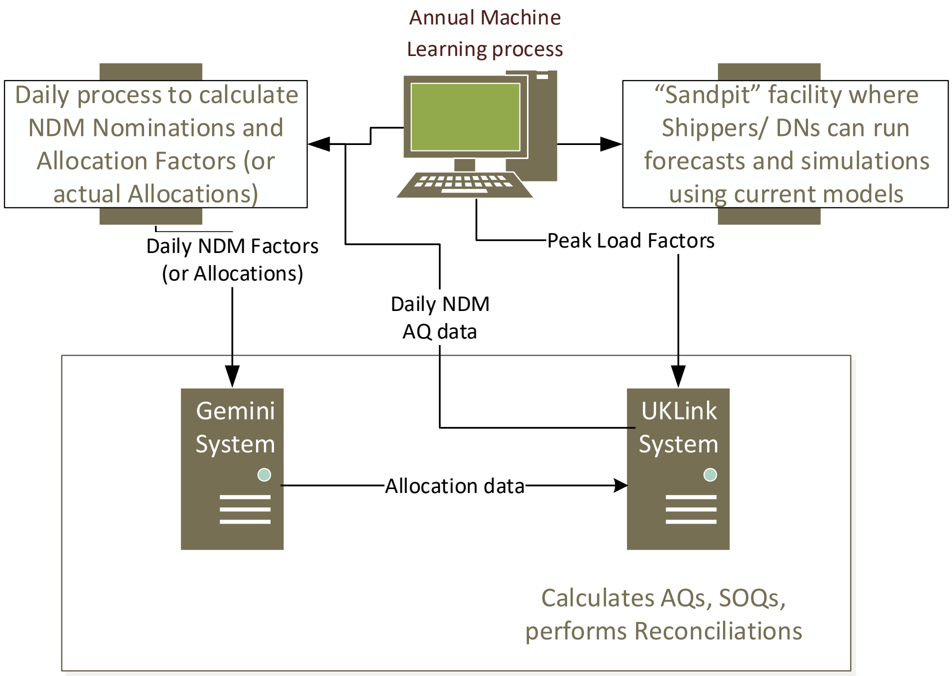

Option 4 – Continual Machine Learning with direct outputs into the Allocation Systems

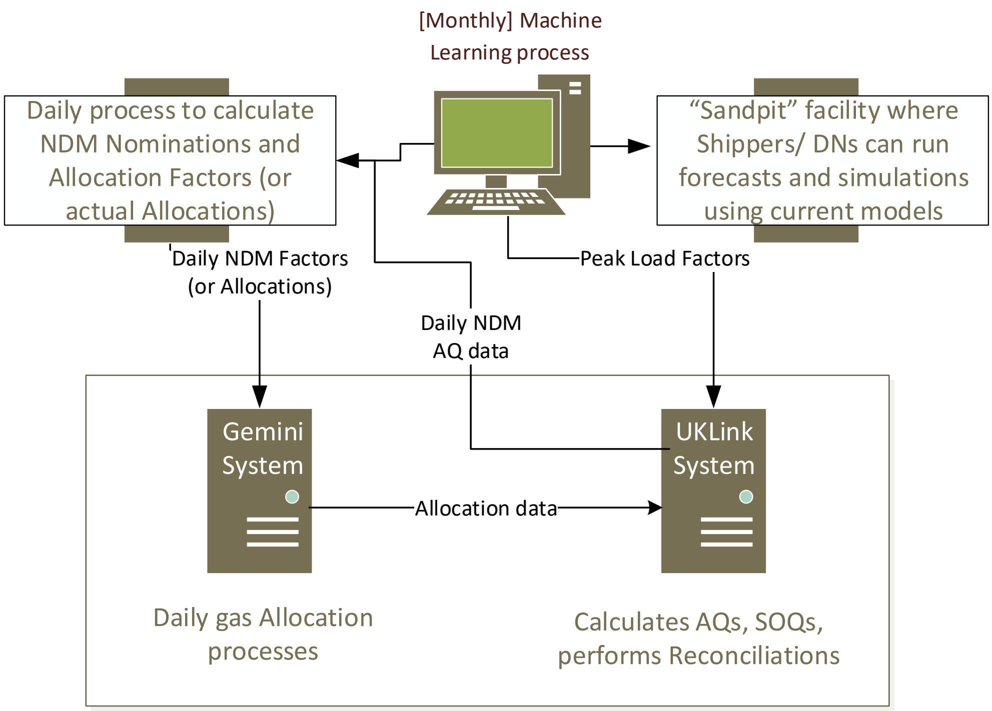

APPENDIX 2: How would these Options have coped with the COVID-19 pandemic and the effects of the GB Lockdown?

The unexpected COVID-19 pandemic in 2020 and the associated unprecedented mandatory lockdown of many GB businesses (as well as travel restrictions) caused a noticeable reduction in gas usage, which was not built into the NDM Allocations models, resulting in unusually negative levels of Unidentified Gas for several months. Although this was a unique and unprecedented event, it is worth reviewing how each of the Options might have coped with the event, and whether the levels of NDM Allocation (and therefore UIG levels) might have been different.

Option 1 – No Change (ALPs and DAFs developed each year based on regression analysis)

The NDM profiles for Gas Year 2019/20 were developed during the Spring of 2019, before the novel Coronavirus strain was identified. The models have no concept of the recent lockdown, so they continued to allocate “normal” levels of gas usage to all NDM Industrial/ Commercial (I&C) sites, regardless of whether they had to stay closed or had actually stepped up production to meet consumer needs (e.g. PPE). Because the gas End User Categories don’t include the type of business, any ad-hoc updates to the models in March or April would have been very broad brush, e.g. a 20% reduction across the board. Meter point reconciliation corrects the position at individual site level, once valid reads have been accepted into settlement.

Option 2 – Machine Learning development of ALPs and DAFs each year

The NDM profiles would still be developed before the Gas Year, so would still have had no concept of the recent lockdown, so would have continued to allocate “normal” levels of gas usage to all NDM I&C sites, regardless of the actual usage patterns. This would have created negative UIG. As the gas industry central systems do not hold information about the type of business, any emergency mid-year changes would still have been very broad brush. Meter point reconciliation would correct the position at individual site level, once valid reads have been accepted into settlement.

Option 3 – Annual Machine Learning with direct outputs into the Allocation Systems

Under this option the NDM profiles would still be developed before the Gas Year, so would still have had no concept of the recent lockdown, and would have continued to allocate “normal” levels of gas usage to all NDM I&C sites, regardless of the actual usage patterns, creating negative UIG. Once again, only broad brush emergency mid-year changes (e.g. 20% reduction to daily allocation for certain End User Categories for a short period) would be possible. Meter point reconciliation would correct the position at individual site level, once valid reads have been accepted into settlement.

APPENDIX 2: How would these Options have coped with the COVID-19 pandemic and the effects of the GB Lockdown?

Option 4 – Continual Machine Learning with direct outputs into the Allocation Systems

Under this option the NDM profiles would be developed before the Gas Year, but would update throughout the year, using more recent data, possibly monthly. In normal circumstances, the within-year change is likely to be minor. In the case of the COVID-19 pandemic, there would be sudden large changes in usage at some sites. However without a separate initiative to capture and maintain information about the nature of each business (e.g. school, shop, office, care home) and build separate usage models for different business types, the Machine Learning tool would only have been able to apply average updates to the models, e.g. a [20%] reduction to I&C models. Due to the lag between receipt of new data and updates to the models, any changes would probably not have kept pace with the latest industry position: it would take a while for the models to react and by then, many sectors would be coming out of lockdown.

Once again, meter point reconciliation would correct the position at individual site level, once valid reads have been accepted into settlement.

Summary of COVID-19 impacts

None of the proposed models for improving the NDM Algorithm could have predicted the impacts on the lockdown on gas usage. Only Option 4 could have reacted to the sudden changes in gas usage, and even then, only at an average level for each end user category. There would also be a time lag between demand changes being observed and the impacts being used to update the NDM models for allocation purposes.

The use of a sample of live actual data from Smart meters or AMR devices to estimate NDM usage for a Gas Day might have allowed NDM Allocations to react more quickly to COVID-related usage changes – but only in aggregate, not necessarily at site level. None of the four Options would make use of live actual data. That is not feasible at present due to the concept of Shipper Hub in the Gas industry, and the lack of direct CDSP (Central Data Services Provider, i.e. Xoserve) access to the DCC (Data Communications Company – central provider of access to some Smart meters). Use of live usage data in NDM Allocations could create bigger variations between NDM Nominations and Allocations, and one of the industry's stated objectives was to improve that alignment rather than worsen it.