The Driverless Alumina Refinery

Nicholas Salmon
Process Engineer, Worley, Brisbane, Australia
Corresponding author: nicholas.salmon@worley.com

Abstract

Since the 1950s, process industries of all varieties have gradually adopted automation to achieve simple process goals. In recent years, industries like mining have extended this principle to the complete replacement of human operators with technology. With the rapid advancement in artificial intelligence and machine learning, it appears likely that this trend of automation will continue, with computers gaining the capacity to perform complex tasks which previously required the watchful eye of skilled operators. Within the alumina industry, the benefits of "going driverless" are clear: an operator-free refinery offers improved safety outcomes, higher labour productivity and better equipment efficiency. On the other hand, it presents an enormous technical challenge, demands more intense capital investment and risks community backlash. This paper imagines an operator-free alumina refinery, and contemplates which economic and technical conditions, if any, would make it a favourable investment.

Keywords: Autonomous systems, Alumina, Labour.

1. Introduction

The advent of simple programming techniques in the 20th century brought with it an industrial drive to create automated systems. These automated systems are able to interpret simple inputs and can actualise a response from a set of pre-defined options to achieve well-defined goals. In some applications, these automated systems can respond more quickly, more consistently and more effectively than a human tasked with achieving the same goals.

As a consequence, this trend of automation has brought with it both improved performance and a gradual improvement in the labour productivity of all industries, including mineral processing and alumina refining.

The revolution of the 21st century will move beyond the *automated* system to the *autonomous* system. Unlike automated systems, autonomous systems interpret more than simple inputs; they collect and interpret data from whole systems, and then devise complex responses which consider the impacts upon the entirety of that system. In so doing, they offer immense potential to expand the scale, and hence the benefits, of automation.

The instance of technological autonomy with which the public is most familiar is the driverless car. It is the most prominent among many examples of systems where the entirety of decision-making authority is left to the hands of a computer.

But, to state the obvious, an alumina refinery is not like a car. Its capital cost is in the order of 100 000 times greater, and its average life span is much higher. And thus, the cost of a "driverless" refinery is much greater than that of a car, although its benefits can be reaped for far longer. On the flipside, a non-autonomous car bought today may still be useful for the entirety of its ~10 year life span; a non-autonomous alumina refinery built today risks looking antiquated well before it reaches middle age.

Naturally, for a system to be autonomous, it also needs to be automated to a significant extent. This paper does not seek to provide detailed answers to the question of how one might go about automating every task in a refinery that an operator might perform. Rather, it considers the characteristics of autonomous systems in order to identify the opportunities and challenges facing the alumina industry as it looks toward a driverless future.

2. Industrial Automation

2.1. The Journey to Autonomy

Many of the world's large industries are already making the transition to autonomous systems. The transport industry is a clear example. 100 years ago, when trucks were first used for material transport, they required expertise and system understanding to handle the truck safely, and perform the frequent maintenance required. Modern trucks demand less from their drivers – from bifocal mirrors to automated gearboxes, satellite navigation and rear-view cameras, the task of driving and maintaining a truck is far easier than it used to be. The industry has now reached a point where the truck itself is so sophisticated that none of the driver's knowledge is required at all to maintain normal, safe operation.

In other words, for processes in which any form of decision making is required, there is a gradual progression from "art" to "science". Nascent processes tend to be artistic, meaning decision making is entirely left at the hands of an operator, who interprets information and makes decisions based largely on personal experience. By contrast, in well developed, scientific industries, the task of gathering information and formulating responses has been well codified, and more responsibility can be left in the hands of a computer. The journey from art to science has been made for challenges as outlandish as space travel and as mundane as baking a cake.

With autonomous decision making yet to be present to a significant extent in most industries, humans are still considered a necessary part of the safe and effective operation of scientific processes. This is because humans can be trusted to provide a system with resilience. In an autonomous system, we must instead trust the computer to provide such resilience. This industrial progression is summarized in Figure 1.

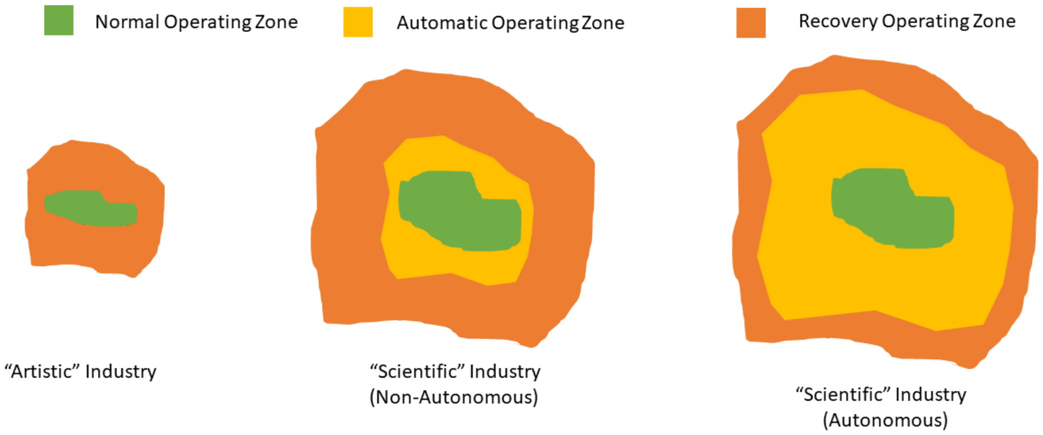

Figure 1. Diagrammatic Representation of industries at various developmental stages.

In Figure 1, the green area represents the operating conditions of a stable system operating normally; yellow represents conditions in which automatic controls can be trusted to return the system to its normal state, and dark-orange represents an area in which complex decision making is required urgently to prevent loss of system control, here referred to as the "recovery" zone.

Actions in this zone cannot always be rooted in procedure, as the conditions may be unique and unforeseen. The area outside the dark-orange boundary represents an out of control system.

In an artistic industry, the normal operating zone is small. Whenever the system leaves a normal operating zone, an operator must intervene to improve performance. Early examples of security filtration are a good example of such an industry. They required a well-controlled feed from the settling area, and frequent maintenance. Any adverse solids readings would require operator intervention to correct; perhaps to hose out the filter, or change out its cloth.

Even non-autonomous scientific industries are more resilient than their artistic counterparts, with a greater scope for normal operating conditions, and more automatic control to handle abnormal but predictable conditions. For instance, in the case of security filtration, a modern vertical pressure filter can normally operate at higher alumina supersaturations than a sand filter, represented by the increase in size of the green area. Meanwhile, the predictable, but abnormal, operations of the filter cycle (i.e. turnback and mud dump) are controlled by automatic systems, as shown by the yellow area.

For the non-autonomous scientific industry, the dark-orange area which requires operator action is still necessary. Part of their role is predictable, non-automatic tasks that do not require significant critical thinking (e.g. swinging valves to put the filter feeds on caustic wash at regular intervals, or replacing filter cloths) The remainder of their role consists of non-predictable tasks which require critical thinking (e.g. diagnosing which filter is the source of high solids in the filtrate).

Thus, a two-fold development is required to move a system from non-automated to automated with equal or improved overall system stability. Firstly, an operator's predictable tasks need to be able to occur automatically. In the example of security filtration, this would incorporate automatic valves to enable a filter to be caustic cleaned without manual intervention. On Figure 1, this is represented by expanding the scope of the yellow area. Secondly, to offer the same overall system resilience, the computer needs to be able to perform the same complex problem solving as the operator, which is why the dark orange area is still present.

Notably, this does not require the expansion of the green area representing normal operation. While the expansion of this area is always welcome, and may result in improved process outcomes, no system will always operate entirely within this green zone. Confidence in autonomous systems requires confidence that they will make the right decision in all complex circumstances it is faced with. To do that we need to increase the size of the dark orange area (i.e. make the computer smarter), and increase the size of the yellow area (i.e. reduce the frequency with which the computer needs to be smart).

2.2. A Diagnostic for Autonomy

Based on the above, we propose four characteristics of an industry that make it suitable for autonomous operation:

- Technical Maturity;
- Technical Complexity;
- Measurability; and
- Process Risk.

The first two criteria, technical maturity and complexity, match the two requirements identified in Section 2.1. Technical Maturity refers to the ease with which the size of the yellow shape in

Figure 1 can be increased; technical complexity refers to the ease with which with size of the dark orange shape can be increased.

A technically mature process can more easily be made autonomous, because the frequency with which complex problem solving is required is low. Most of the complications should be predictable using past experience, meaning it is likely that all tasks are (or can be) represented by procedures which are to be followed in all but the most extreme of circumstances. Indeed, where these procedures are well drafted, even an unintelligent computer may do a better job than a human, as it is unlikely to miss a step, and may be able to execute tasks more quickly than a human. Haulage trucks in mining are a good example of a technically mature process which has been identified by many companies as suitable for autonomous operation.

The second characteristic, technical complexity, tends to complicate the task of making a system autonomous. Technically complex systems are characterized by strong interrelations between subprocesses. Accordingly, to return the system from an abnormal to a normal state may involve consideration of a wide range of options which consist of many steps with unpredictable outcomes. It may also be more difficult to make all of the various subprocesses automatic if there is a large number of options for a computer to choose from. For instance, manufacturing processes tend towards autonomy when the number of options available to an operator is low, and when an incorrect decision has limited impact on the remainder of the process [1] By contrast, continuous processes with a variety of unit operations and internal feedback loops typically are complex due to the significant interrelations of operations.

Measurability is included in the set of characteristics, as it is a prerequisite to good decision making. If the computer cannot detect the state of the system, then it cannot take appropriate action to correct the state if it is undesirable. Alternatively, it may take unnecessary actions even when the system is operating well. This facet includes quantitative measurements (both online and offline) and qualitative measurements. Qualitative assessments are often used by operators to identify systems that are not in an appropriate state (e.g. a motor making grinding noises may be in need of maintenance), or to identify the source of a problem (e.g. quantitative data may indicate a disc filter is not performing; visual inspection may be required reveal the cause of the failure to be poor cake thickness due to cloth condition). Existing autonomous systems incorporate both qualitative and quantitative measurements. For instance, some Australian coal seam gas companies operate their wells autonomously, with performance monitoring by both quantitative data, sent to a central control room, and qualitative data collected by drones.

The final characteristic considered is process risk; it is included here because it is such a significant driver of corporate decision-making that it is likely to impact decisions on system autonomy. Process risk refers to the possibility of widespread harm occurring to plant, people and environments because of loss of control of the system. It may also include significant loss of production. Even where the likelihood of incident is very low, it becomes very hard to justify removing the physical presence of a skilled operator where the consequences of failure are great. Even though modern planes have sophisticated autopilot systems, there are very few who would be willing to board a driverless plane – and even fewer who would be willing to offer it as a commercial venture.

3. The Feasibility of Automated Alumina

Having identified the traits of an industry which is suitable for automation, we here analyse the extent to which the alumina industry – specifically, Bayer alumina – possesses those traits.

Technical Maturity – the Bayer process is extremely technically mature, having been in use since the beginning of the 20th century. The vast majority of the processes within the industry are well understood, and there is an abundance of data that enables the quantification of processes to predict their outcomes. It would be possible to construct an autonomous refinery which entirely made use of proven technology.

Technical Complexity – Although the Bayer process is technically mature, it is also technically complex. Because the process operates on a closed loop, actions taken in any part of the refinery may impact any other part of the refinery. Process goals are often conflicting, such as balancing hydrate production, hydrate size and hydrate strength in the precipitation train. Because of the presence of corrosive liquors, scaling liquors and abrasive slurries, equipment is prone to poor performance, which makes it challenging to predict the outcome of various actions that may be taken by operators. Different bauxites lead to different liquor compositions, which can make it difficult to identify the source of problems on new refineries or when changing bauxite supply or co-processing different types of bauxites, even where there is a wealth of data available from existing installations.

However, while this technical complexity is likely to create teething issues for an autonomous refinery, the computer operator may also be able to demonstrate improved performance relative to a human operator, particularly one who is new to the industry. These benefits are discussed in more detail in Section 3.1.

Measurability – Measurement in the Bayer process is challenging. Scaling liquors and settling slurries makes operating online instrumentation challenging. Lab sampling is required on a frequent basis to verify the results of online instrumentation, and to provide data where instrumentation is not available. The quality of instrumentation is a major barrier in the progression of Alumina to an autonomous industry.

Process Risk – The process risk inherent to the Bayer process is low, particularly on low temperature plants. In comparison to, for instance, the oil and gas industry, it does not make extensive use of flammable materials that have the capacity to create fire and explosion. The bulk materials in wide use on the refinery are non-toxic. There are risks of vessel overpressure, but the majority of refineries already endeavor to use entirely automatic systems to prevent accidental overpressure.

On the flipside, the majority of safety risk from the Bayer process occurs to personnel onsite because of the presence of moving equipment, high temperatures and corrosive liquors. This risk would be substantially reduced by reducing the number of people onsite.

3.1. Two Possible Automated Refineries

Here, we do not imagine technology so far advanced that the refinery can be constructed and left to produce alumina without any human intervention for half a century. Instead, we consider the most realistic next steps in the journey to autonomous operation, which is the replacement of a site's operators with technology. It is unlikely that maintenance and technical staff will be entirely replaceable in the near-term future. Therefore, in order to realise the goal of having no people onsite, the refinery maintenance strategy needs to favour short, intense periods of maintenance rather than rolling maintenance as is the norm on existing sites.

This may involve a campaign style digestion approach, which would operate without onsite staff for a period of ~ 2 months. At the end of the campaign, a team of contractors would be brought to site to conduct intensive maintenance. Such a refinery could operate in an isolated location – for instance, close to a bauxite mine – without requiring people to live in an isolated location. It would either require significant sparing in many areas (Such that intense maintenance could be conducted without a shutdown) or technology which would allow for rapid start-up and shut-down of the refinery.

Alternatively, for a refinery in a relatively centralized location, the automated operator would raise purchase orders for contractors to come onsite on an as needed-basis. Personnel would only be present where maintenance was specifically required, and the refinery would usually be unmanned at nights, weekends and on public holidays, unless there was a particular emergency that required more maintenance than the norm. Such a model would rely on locating the refinery in an industrialized area, such that maintenance staff would be available who could be mobilized at short notice.

Regardless of which model of maintenance is selected, there are a number of common traits between both models. Automatic angle valves would need to be installed. All processes which need to occur with high frequency (e.g. filter cloth replacement, caustic cleaning of precipitation tanks, clearing of screen boxes) would need to be fully automatic. Some devices to make these processes automatic already exist – for instance, descaling robots are already in common use – others may need to be developed. For small scale maintenance, such as thermocouple replacement, drones or simple robots would be used. Drones may also be used for the collection of sample liquors for return to the lab, although an autonomous refinery would benefit from moving mostly towards online instrumentation. For other instrument issues where the hardware cannot be replaced, software solutions may be implemented – for instance, sophisticated algorithms may be able to return reasonable data even in the face of instrument drift.

Refinery data would be passed continually to central control rooms by a combination of online instrumentation (quantitative) and drone surveillance (mostly qualitative). This could be accessed by technical staff, and a relatively small staff of operators who would continue to operate the refinery from a distance. Estimates from the mining industry indicate that the number of operators is ~1/8th of the number of operators required on a non-automated site [2]. The number of operators is likely to decline over time as the challenges associated with moving to autonomous systems are resolved. The operators may be able to oversee either multiple areas within individual autonomous refineries, or single types of areas at multiple autonomous refineries.

It is likely that retrofitting an existing refinery to operate autonomously would be hugely expensive, as the operating philosophy is entirely different, and substantial amounts of equipment would need to be scrapped and replaced with automatic counterparts. The automated refinery is likely to be a greenfield installation.

3.2. Process Control

Existing alumina refineries typically use automatic controls, which normally serve to achieve process goals within a specific area. The operator will select the set point consistent with refinery-wide goals. In some instances, the set point is set not by the operator, but via cascade from another controller. Cascade control is, in general, superior because the control loops target refinery-wide goals more directly.

Some refineries go further, and make use of model predictive control [3][4]. In model predictive control, the controller collects data from a range of instruments, and uses a model to predict the best route from the refinery's current state to its target state by adjusting all the variables which it can manipulate (i.e. control valves and variable speed drives). The "best" route is the one that achieves the controller's goals most quickly without compromising process stability. This is a more complex mode of control, but it is superior to cascade control because it can target several goals at once. Even when used to control single variables, the predictive model, if suitably designed, often also returns more stable results than a traditional PID controller (even one with feed-forward mode).

An autonomous refinery at the very least would need to make widespread use of model predictive control. A refinery will not operate safely or effectively if it consists only of a series of mostly independent control loops. Realistically, the refinery's control system will also require machine learning algorithms. Machine learning algorithms enable the control system to recognize conditions similar to those it has encountered before, and use the results of that encounter to inform the decisions it makes. In other words, the predictive model constantly updates itself by using plant data. While predictive model control is well developed in a number of industries, machine learning is still relatively novel technology, although it is advancing rapidly.

Machine learning is an important enabler of a driverless refinery. It is unlikely that the first day of autonomous operation will produce results as good as those from standard refineries; however, with machine learning, the refinery can improve itself over time, and eventually outperform a standard operator.

4. Cost-Benefit Comparison

4.1. Advantages of Autonomy

Given that there are technical challenges – some of which are significant – of moving to an operator-free refinery, the benefits of the transition need to be significant in order to demonstrate a business case for the change. We discuss four benefits: Safety, Labour Cost, Maintenance Cost and Plant Performance.

Safety – The most substantial and indisputable benefit is to operator safety. There are many inherent risks to working in a Bayer process environment. The most effective way to guarantee that all operators are safe is to remove them from that environment. In the foreseeable future, maintenance will still require human interactions which expose people to the typical risks of an alumina refinery, but these occur at a much lower frequency.

Labour Cost – The second clear benefit is a reduction in labour costs of operators. In some jurisdictions, this reduction will be very small in comparison to the total operating cost. However, in other areas such as Australia, labour costs are substantial, and the benefits to be realized are more significant. Where there are pressures placed on the operating costs of a refinery, reducing the cost of labour is often the only pathway available – even the best refinery operating the modern Bayer process must use energy, bauxite and caustic, but in the future the best refinery may be able to go without operating labour.

Maintenance Cost – A well-designed autonomous refinery is likely to have a lower maintenance budget than a non-autonomous one. On mines, the use of autonomous technology reduces the likelihood of running equipment to failure, by detecting equipment which is giving indication of failure in the near future, and by collating residual life data to optimize the maintenance schedule. Artificial Intelligence, which can interpret huge volumes of data, is more suited to this task than a human. This reduces the need for emergency maintenance and enables a better scheduled and more predictable maintenance workload.

There are countless opportunities for similar application in the alumina industry. Rake failures in conical vessels can be predicted and prevented by collecting and appropriately collating time series data on power usage, rake torque and solids depth. Measurements of noise and vibration can detect motors in need of maintenance. With predictive mass balance models, instrument failures can be identified and corrected more quickly. All of these changes and many more can reduce maintenance budgets, and the requirement for maintenance staff to complete dangerous onsite work.

Plant Performance – Benefits to plant performance may materialize as increased production, or reduced use of raw materials and energy. Again, the list of opportunities for autonomous performance improvement is long. Refineries already track the gradual decline in heat transfer coefficient between cleaning cycles in digestion heaters; an automated refinery would be able to take that information and optimize the cleaning cycle in real time, making it more dynamic than a schedule which is beholden to operator shifts. Computers are far less likely to be siloed operators of individual plant areas; they are more capable of optimizing whole of refinery behavior. Computers are less prone to idiosyncratic changes of behavior between shifts. To the extent that more measurements are required to support autonomy, the autonomous operator can make use of those measurements to detect deviations from normal behavior and design suitable responses. In short, a plethora of options become available to the engineer to optimize the process that were previously not possible because of the realities of day-to-day refinery operation.

Admittedly, computers do not yet possess the many skills discussed above, and progress towards a well-operated autonomous refinery will no doubt be gradual. In the interim period, the small number of operators managing the refinery from the central control room are likely to be the most skilled operators available. The use of model predictive controllers with machine learning algorithms, as discussed in Section 3.1.1, will gradually improve performance. For refineries in isolated areas, or in countries without an existing alumina industry, these operators are likely to achieve better performance than the operators who may have been available onsite.

4.2. Risks of Autonomy

Despite the prospective benefits, there are a number of other risks associated with autonomous operation that may jeopardise both autonomous process industries generally and the alumina industry specifically. Three risks are considered here: Cyberattack, Community Backlash and Capital Intensity.

Cyberattack – The most common risk identified by industry associated with autonomous operation is cyberattack. The volume of information which must leave the plant boundary on an autonomous refinery is significantly larger than that which must leave the refinery for manned operation. Moreover, the consequences of cyberattack are for obvious reasons much greater. Such cyberattacks could result in the theft of valuable data or the intentional maloperation of the refinery. However, with the growing risk of cyberattack comes a growing industry in cybersecurity. Other industries are likely to drive solutions to cyberattack, and the alumina industry ought to be able to benefit from their efforts.

Community backlash – Autonomous refineries will face greater challenges to obtain social license to operate. Alumina refineries can be loud, dusty and odourous. They release visible emissions. They consume finite bauxite resources – often resources which belong to the country in which the refinery is located. At present, all operating refineries still produce bauxite residue and store it in tailings dams, which do not always meet public expectations for environmental protection. Refineries typically rely on a local water source to supply their needs.

At the moment, local communities tolerate these shortcomings because of the huge employment opportunities from refineries. These opportunities will be significantly reduced, or entirely unavailable, from autonomous refineries. Even if a small number of jobs are retained, the perception of their loss may be of significant concern. In response to this community concern, it is possible that future governments will impose more substantial barriers to regulatory approval and higher taxation rates on an alumina industry that is autonomous. The refineries themselves may also be a target of vandalism and theft – which is far harder to prevent with no personnel on site. A carefully managed public relations strategy is essential for a driverless alumina refinery to become reality.

Capital Intensity – The autonomous alumina refinery will have higher capital costs than their non-autonomous counterparts. This will partially originate from a need to automate various components of the process, which naturally has associated capital costs. It may also require more substantial sparing or surge tankage in order to allow for maintenance to be conducted on an intense, sporadic basis rather than consistently by a relatively small team. In return for expending greater capital, the refinery reduces its operating costs.

Alumina has always been a capital-intensive industry, and this change will make it more so. The capital pressure poses a risk to overall plant performance, as attempts to reduce upfront costs may lead to shortcuts which ultimately cost the quality of installed refinery or equipment. These refineries are therefore more likely to be viable in jurisdictions with low interest rates that justify a capital-intensive industry.

Although these risks are significant, and cannot be underestimated, they are also not unique. They are common to all industries that are making the transition to autonomy. The alumina industry will not need to pioneer solutions to these challenges; they are already being pioneered by industries across the world, from whom alumina can borrow and adapt the right answers.

5. Conclusions

There are significant opportunities to be realized in refinery performance by going autonomous. At the same time, there are many technical barriers to overcome before a such a refinery and its benefits can be realized.

Overcoming those barriers may well be a necessity for the alumina and aluminium industry. The trend towards autonomous operation is not unique to this industry; it is common across the global economy. Aluminium’s competitors – other light metals, steel, and light composite materials like carbon fibre – will begin the move to autonomy with or without the alumina industry. Some of them, like carbon fibre, are more suited to autonomy than alumina – increasingly, the processes to produce them are mature, without the technical complexity and measurability challenges faced in the alumina industry. To stay relevant in comparison to those industries, autonomy may be essential.

Greenfield alumina refineries built today will need to be able to operate competitively in the autonomous future. Retrofitting those refineries to make them autonomous would be an enormous task. However futuristic a driverless alumina refinery may sound, it is a future the industry needs to pursue today.

6. References

1. Angel Sanchez-Salas, A framework to support automation in manufacturing through the study of process variability, *PhD Thesis submitted to Loughborough University*, 2016.
2. Drew Bellamy and Luka Pravica, Assessing the impact of driverless haul trucks in Australian surface mining, *Resources Policy.* Vol 36, (2011), 149-158.
3. Michael Missala et al., How Digitalization can further improve plant performance and product quality, *Proceedings of Light Metals 2018*, Phoenix, USA, 2018, 105 - 115
4. Michael Barnes, David Koffler, and Cullen Bowels, Digital Maturity in Alumina Refining, *Proceedings of the 36th International ICSOBA Conference,* Belem, Brazil, 29 October – 1 November 2018, Travaux 47, 241 - 254