Caving 2026 – A van As, D Cumming-Potvin & J Wesseloo (eds)

doi:10.36487/ACG_repo/2645_26

Real-time simulation projection of block cave production rate

Colin Eustace [a,*], Peter Kontos [a], Kevin Nguyen [a]

[a] *Deswik, Australia*

Abstract

Real-time dispatch optimisation is used successfully in many of the world's operating block caves, including Newmont's Cadia Valley Operations since 2015, and more recently at PT Freeport Indonesia and Rio Tinto Oyu Tolgoi. Optimisation of production loader allocation to available drawpoints is effective in both increasing production output and improving draw compliance. However, as the plan is continually optimised rather than fixed at the start of the shift, visibility of the operations plan is limited. The lack of forward visibility can make it difficult to work around production activities and to realise the benefits from optimising the production sequence.

Simulation can be used to provide forward visibility of production operations. Initialised in the current operating state, and calling the ORB optimisation engine for dispatch decisions, a simulation function is used to provide an accurate projection of future operations. In addition to providing forward visibility, the simulation model is used to test the potential effects of changes such as in equipment availability, secondary break allocation or drawpoint order priorities. The current implementation is focused on providing visibility for the next 12–24 hours of operations, but the approach is also extendable to longer time horizons.

This paper provides an overview of the simulation process used to provide forward visibility of production operations when using ORB for dispatch optimisation.

Keywords: *block cave, dispatch optimisation, simulation*

1 Introduction

1.1 Dispatch for block cave production levels

Dispatch of load–haul–dump (LHD) equipment to drawpoints on the production level is a complex planning exercise which needs to consider multiple system level constraints and objectives, while coordinating activities across the production fleet. A combination of the following operations planning objectives and constraints presents a challenge for the daily management of production operations:

- draw compliance (tonnes available at each drawpoint, draw priority, maximum draw rates)
- secondary break priorities and drive/panel closures
- potential for congestion at tipples and with operations in adjacent drives
- balancing of longer and shorter production cycles with downstream haulage and materials handling system capability
- moisture limits and blending requirements
- production fleet and secondary break availability.

Many alternative block cave production scheduling methods have been developed for setting long-term and medium-term production targets (Khodayari & Pourrahimian 2015). These long-term targets are core

[*] Corresponding author. Email address: colin.eustace@deswik.com

objectives for daily operations planning. Effective coordination of dispatch across the cave is essential for consistently achieving high production output while maintaining draw compliance to caving constraints and long-term draw profile objectives. This requires frequent re-evaluation of the draw pattern across the whole footprint.

1.2 Limitations of manual dispatch methods

Historical management of block cave production operations typically used a manual planning process of allocating LHD units to extraction drives at the start of the shift. Then, treating LHD units independently, each LHD unit was assigned a fixed drawpoint sequence using a simple prioritisation heuristic. This process did not typically accommodate consideration for equipment interaction delays from shared ore passes and red-light delays from materials handling systems. Independent planning for each LHD unit resulted in a system that was not coordinated, with conflicts between both LHD activities, and inconsistencies between planned activities and overall constraints and objectives.

The complex constraints and system level objectives of block cave production level operations present a good opportunity for optimisation to improve both draw compliance and overall production rate, while also automating the planning process.

1.3 Dispatch optimisation and autonomous short interval control

Significant benefits from using dispatch optimisation have been documented including for the Argyle mine (Donaldson et al. 2020) and Panel Cave 2 at Cadia Valley Operations (Donaldson et al. 2021). For these operations, following initial rollout, productivity improvements of 7 and 20%, respectively, were realised. Draw compliance for these operations, using a weighted average measure, also improved by 15 and 17%, respectively. Productivity improvements are primarily achieved through the transition from fixed daily to monthly progression targets (Argyle) or ratio bogging targets (Cadia Valley Operations) and the corresponding removal of the planning reset and associated operational inefficiencies at the start of each day. Draw compliance improvements from the introduction of dispatch optimisation have played a crucial role in controlling grade and maintaining cave health for these operations.

In addition to tighter draw compliance and increased production, dispatch optimisation also provides a means to react quickly and consistently to changes in the system state, promotes consistency for production dispatch across crews and supervisors, and allows a shift in focus from operational to tactical and longer-term planning.

1.4 Use of simulation for real-time projection of operational performance

While the benefits of real-time dispatch optimisation are well established, operating with a plan that is continuously optimised rather than fixed at the start of a shift limits visibility of the forward plan and expected performance over the shift. The lack of forward visibility can make it difficult to coordinate secondary break with production activities and to fully realise the benefits from optimising the production sequence.

Simulation can be used to provide forward visibility of how the decisions of the dispatch optimisation engine and production operations are likely to play out over time. Projections of block cave operations are run over a defined period by initialising exactly in the current state and then running through the operations in detail.

Simulation is used extensively for strategic analysis of production capacity for block cave operations (Eustace et al. 2020, 2024; Morrison & Labrecque 2020; Mariager et al. 2022). Models used for strategic analysis of block cave production level performance often employ a detailed representation of production level operations (e.g. draw constraints, equipment movement and interactions, secondary break activities and access). However, even if a sufficient level of operation detail is represented in the model, strategic models generally do not employ the following 2 key features required for the model representation to be applicable to real-time operations:

1. A real-time data interface and functionality to initialise the model in a specific state aligned with the actual system state.
2. Where a dispatch optimisation engine is used for operations purposes, an ability for the simulation model to call the dispatch engine directly.

To provide forward visibility of production operations, existing block cave simulation software has been adapted to support real-time representation of production level operations by initialising the model in the actual cave state and then directly calling the ORB dispatch optimisation engine for each drawpoint order. The simulation model plays out each cycle of the set of drawpoint orders allocated by the ORB dispatch engine. On completion of each order or following any changes to the system state (such as drawpoint hang-up, secondary break clearance or equipment failure) that require a new drawpoint order, the model sends another request to the ORB dispatch engine. Following this iterative process, illustrated in Figure 1, the simulation model progressively develops a forward projection of block cave production operations, using the same dispatch engine that is used by the actual production system.

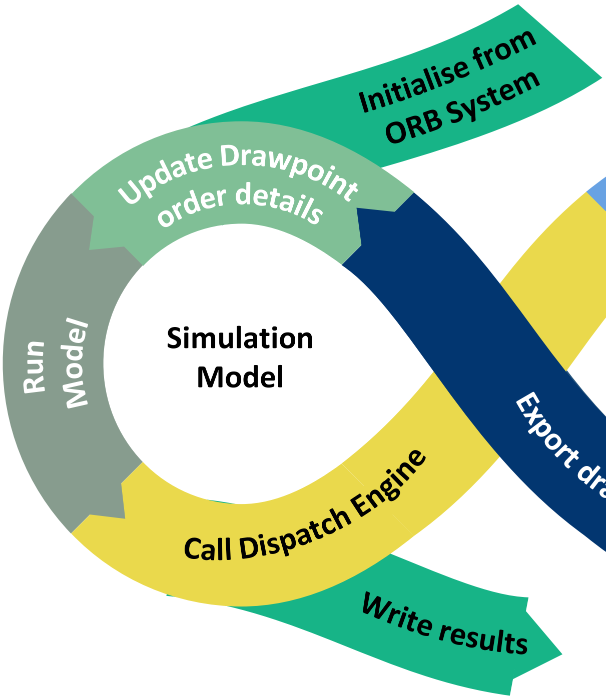

Figure 1 Iteration of simulation and dispatch optimisation functions to develop a forward projection for block cave operations

The current implementation is focused on development of accurate and reliable estimates of system performance over the next 24 hours to support operations planning and testing of operational alternatives. However, the process is not limited to 24-hour forecasts and can be set to run for longer periods. The core simulation logic used for 24-hour duration runs remains the same for longer-duration runs.

Performance estimates from shorter-duration simulation runs benefit from regular synchronisation with system data, such that the state of the system (e.g. tonnes drawn from each drawpoint) is relatable to the actual system. Stochastic events, such as drawpoint hang-up and equipment failures, result in increasing differences between the estimated state and actual state over time, even if the simulation model provides a good representation of the overall rate of stochastic events. While longer-term runs can provide accurate estimates of performance of the overall system, the simulation model representation drifts from the actual state at an individual drawpoint level and is no longer useful for operations support.

Run time can also be a limitation for longer-duration runs if the simulation model uses the ORB engine for dispatch decisions. Calling the ORB dispatch optimisation engine for each simulation dispatch decision provides an accurate representation of the actual system dispatch process but is computationally expensive. Longer simulation runs (e.g. >1 week) can be configured to use a simulation heuristic dispatch engine to approximate the ORB optimisation engine decisions. This results in a marginal decrease in the accuracy of performance estimates but significantly improves execution speed for longer runs.

A further requirement for longer-term simulation runs is the need to transition hang-up frequency with draw column height and draw call as the draw profile changes from month to month. This is supported by the model but not used for the current implementation.

2 System architecture

Real-time simulation of production level performance requires both initialisation of the simulation model in a specific state and the ability to directly call the ORB dispatch engine to replicate the drawpoint orders dispatched in practice. The simulation model and ORB system use a common cloud framework. A simplified representation of the cloud framework is provided in Figure 2. The framework facilitates transfer of data used to define the current operating state from the ORB system, supports dispatch requests from the simulation model to the ORB dispatch engine, and allows embedded reporting dashboards for forward projections of production level performance.

2.1 Initialisation of production level representation

A simulation run can be initialised from within the ORB cloud application. The data required to initialise a simulation run are exported from the ORB cave state database and include:

- static data for the block cave configuration (drawpoint, EXD, orepass/tipple, ID location, groupings and other panel configuration details)
- dynamic data (current number, location and order assignments for production LHD units; state, draw call (tonnes available), material attribute and draw compliance data for each drawpoint; cycle time data updates, current secondary break crew configuration and allocations).

These data are pushed to a cave state store for the simulation model and can represent either the current snapshot in time or the operational state at a specified time in the past. Historical snapshots are primarily used for model testing and validation purposes. The static and dynamic state data are augmented with additional parameterised operations configuration data stored in the simulation model application database (e.g. drawpoint hang-up and equipment failure rate assumptions). Operations configuration and cave state data can be accessed through the simulation model web application after importing from the ORB system, if required, and modified to create alternative scenarios for testing purposes.

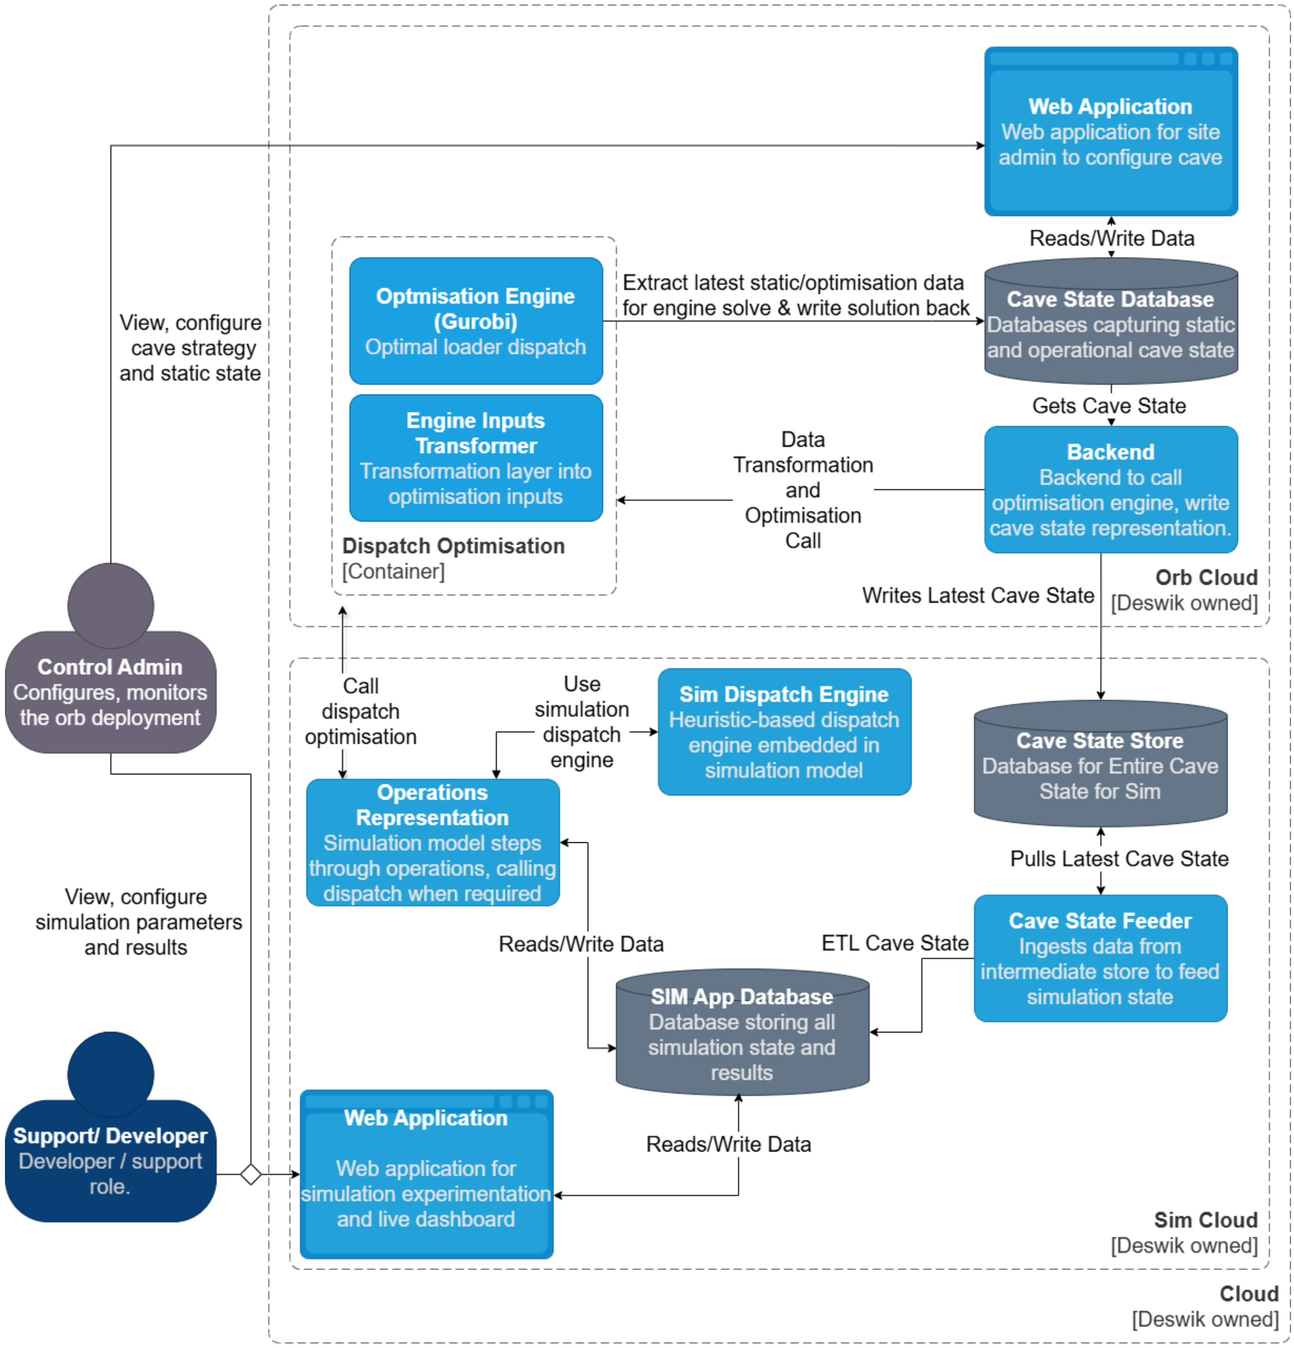

Figure 2 ORB simulation architecture

2.2 Execution of simulation runs

The simulation model is set to run automatically at regular intervals but can also be triggered manually from the ORB application (e.g. if actuals start to deviate from simulation projections) or run directly from the simulation web application user interface. Automated runs are typically configured to run through the next 24 hours of operations, using the ORB dispatch engine for drawpoint allocation decisions. Operations projections are published to an operations control room dashboard. The simulation model can be run in alternative configurations through the simulation model application interface to support different analysis objectives, including scenario configurations:

- from the current time for performance projections or from a historical starting point to support comparisons with historical performance for model validation and testing purposes

- in a reference state or with modifications to test the production level system response to changes (e.g. drive closures or opening additional panels, alternative secondary break strategies)
- using the ORB dispatch optimisation engine (more accurate, but computationally expensive) or a SIM Dispatch Engine (less accurate, but faster and suited to longer runs/many replications)
- with or without a representation of stochastic events, such as hang-up events, equipment failures (upper bound versus expected average or performance range).

The system architecture also supports initialisation and execution of many concurrent simulation runs, publishing of simulation results for in-progress scenario runs and multi-user access.

3 Model testing and scenario configuration and use cases

Real-time simulation is currently being tested at selected sites where ORB is already used for dispatch optimisation. The testing is currently focused on 24-hour projections, reviewing accuracy, reliability and root causes for differences between predicted and actual performance, including:

- expected versus actual LHD cycle times for each drawpoint order
- drawpoint, panel availability and changes in cave state over the shift/day
- loader availability for production across the shift
- actual versus expected secondary break performance.

An indicative operations centre dashboard is provided in Figure 3, with a comparison between target (green), simulated throughput (magenta) and actual throughput (blue) for a 24-hour period. The model currently uses known cycle times for each drawpoint and orepass/tipple combination to predict production rates. Drawpoint-specific cycle times are based on records of recent historical performance maintained and regularly updated by the ORB system. Additional system metrics are configured on a site-specific basis to indicate system state and to highlight areas requiring further investigation. The intention of these metrics is to highlight areas where actual performance is deviating from expected (simulated) to prompt further investigation.

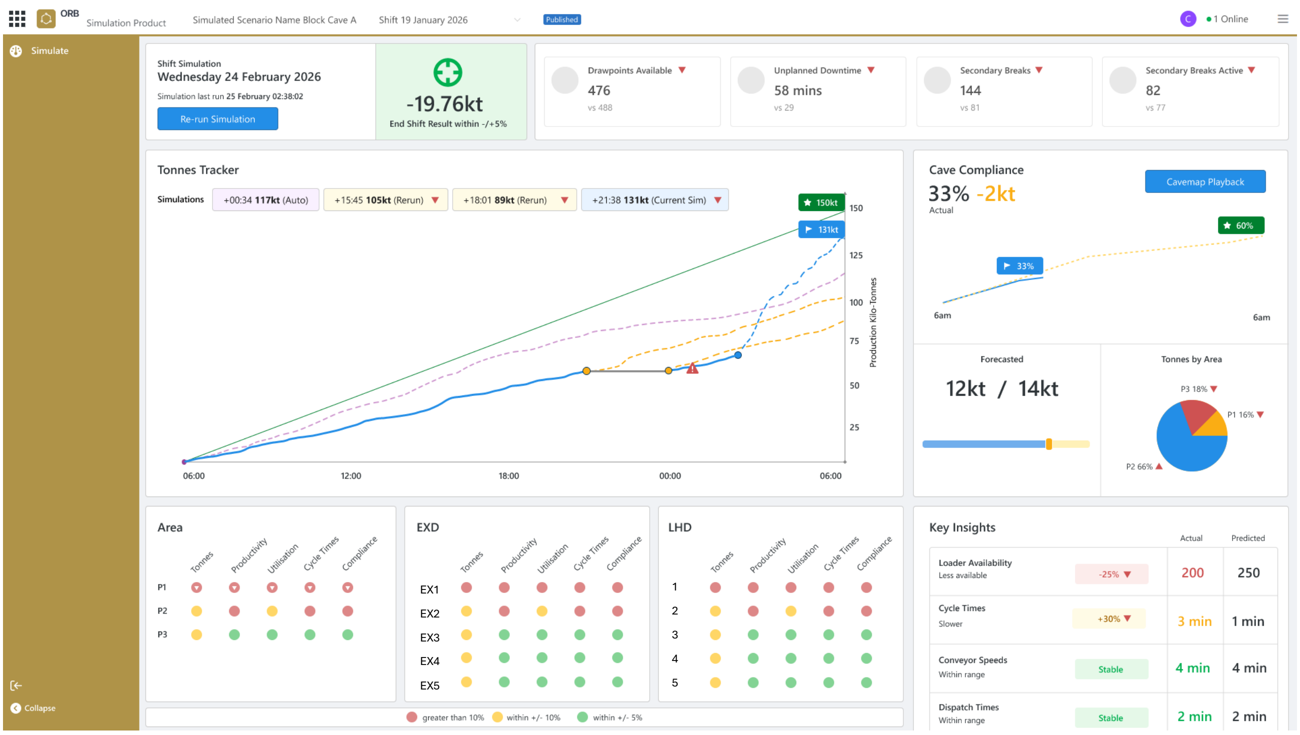

Figure 3 Operations centre dashboard with production forecast shown in comparison to daily targets

3.1 Representation of planned system changes

The simulation model uses data descriptions of planned system changes to modify the system state from the initialised state during the course of a simulation run. Some of the planned system changes, such as planned drive/panel closures, are provided with the initialisation data exported from the ORB system. Other planned changes, such as planned production LHD and secondary break fleet availability, are currently configured directly in the simulation model interface.

3.2 Representation of unplanned events

The simulation model can be run with and without unplanned stochastic events. Deterministic runs, excluding unplanned events, can be used to provide an upper bound representation of production performance. Inclusion of unplanned events (e.g. drawpoint hang-ups and equipment failure rates) can be used to provide estimates of likely ranges for performance, which considers both the current cave state objectives and constraints as well as background failure rate.

Including consideration for hang-up and oversize predictions allows an improved estimate of how long production activities can operate in a panel/extraction drive before secondary break is required. This has the potential to change the way production and secondary break activities are planned during the shift/day.

As the model is initialised exactly in the current state, the accuracy of the model representation of the state of the system tends to degrade over time as the influence of unplanned changes from the starting point (e.g. drawpoint hang-ups, early or late completion of production and secondary break activities) has increasing influence. Progressive refinement of both the representation of operations fundamentals and frequency and duration of unplanned events will improve the duration over which the model can reliably produce accurate operational performance projections in normal day-of-operations circumstances. Ongoing improvements in drawpoint-specific hang-up and oversize predictions and equipment availability will help to improve the accuracy of forward projections, particularly for medium-term and longer-term scenario runs.

3.3 Management of moisture and ore attributes

The ORB dispatch engine can be configured to manage drawpoint moisture, fragmentation and ore grade by issuing orders that blend ore from multiple drawpoints. For short duration runs, drawpoint attributes are assumed to remain constant for the duration of the simulation run. For longer-term simulation runs, expected changes to drawpoint moisture, fragmentation and ore grade over time may be required to improve the simulation representation of the system over time. The simulation model does not currently provide a mechanism for representing long-term changes in drawpoint material attributes.

3.4 Current use cases

The primary use case for the 24-hour look-ahead provided by the simulation is to inform production versus secondary break planning for each panel or extraction drive. Understanding the likely state of each panel at the end of each shift can influence secondary break or production priorities if some areas of the cave are predicted to diverge from draw compliance targets. Scenario testing can be used to evaluate whether alternative allocations are likely to provide a better outcome.

In addition to influencing operational decisions, the simulation functionality also helps to develop confidence in the automated ORB optimisation engine, which can sometimes be viewed as a ‘black box’, by showing how ORB dispatch decisions are likely to play out over time.

The short-term operations simulation can also be leveraged to refine long-term performance estimates by creating an operationally relevant validation state. Repeated simulation of each day of operations and comparison with the actual state of the system helps to highlight areas where the operations simulation can be refined and these refinements can also be reflected long-term simulation runs.

4 Future use cases for real-time simulation

Ongoing improvements in functionality are intended to broaden the capability of the model and outputs to support identification of causes for departures of actuals from expected performance. Beyond the current short-term use case, further enhancement will also support the seamless transition from current state to long-term performance projections.

4.1 Short-term performance

Short-term simulation projections are intended to provide forward visibility of production level operations based on the known cave state production objective and operational constraints. This provides an advance indicator of requirements to make adjustments and can also show the likely effects of typical unexpected events. Ongoing enhancements to the functionality of the tools are intended to support a number of different use cases for operations planning, including:

1. Visibility of expected performance. Fine-tuning accuracy and reliability of performance estimates and proving detailed data for comparison of actual against expected performance.
2. Root cause analysis for shortfall relative to targets. Ongoing enhancements to allow operators to drill down into performance data (e.g. by drawpoint or by LHD unit). There is also potential for automated identification of root causes for variances.
3. Manual testing of alternative options. The current implementation of the simulation model supports scenario analysis of alternative scenario configurations (e.g. test production output for an alternative drive closure program). Ongoing enhancements are planned to provide broader functionality for testing of alternative scenario configurations.
4. Automated testing of a limited set of alternative scenario configurations. This testing could also be used to identify the likely success of system configuration changes, such as panel closure sequences, to present options for operator selection.

4.2 Extension to medium-term and long-term performance estimates

Although simulation scenario processing times increase with the duration of the run, simulation run times can be extended to provide long-term (through to life of mine) production output and draw compliance performance results. To maintain manageable simulation run times, a SIM dispatch engine is used in place of the ORB dispatch optimisation engine. The SIM dispatch engine follows draw constraints but provides fast rather than optimal dispatch decisions. This allows simulation execution speed to be suitable for life of mine (or remainder of life) runs. Refinements to the SIM dispatch engine to provide comparable performance to the ORB dispatch engine are ongoing.

5 Conclusion and future enhancements

Real-time dispatch optimisation for block cave operations is well established and has been demonstrated to provide improvements in both draw compliance and productivity, but the continuous optimisation process limits visibility of expected shift operations. To provide forward visibility, simulation functionality has been integrated with the ORB dispatch application. The simulation model can be initialised in the current cave state and is configured to call the ORB dispatch engine in response to changes in cave state that require new dispatch decisions.

Currently being tested at selected sites where ORB is already used for dispatch of production LHD units, the simulation functionality is used to provide a rolling forecast for production level operations over the next 24 hours. The tool is expected to provide improved visibility of expected performance and can be used to test the effects of changes to the system (such as panel closures) on draw compliance and production output.

References

Donaldson, SN, Harney, W & Cox, T 2020, 'Operational and real-time LHD dispatch at Rio Tinto's Argyle Diamond Mine', in R Castro, F Báez & K Suzuki (eds), *MassMin 2020: Proceedings of the Eighth International Conference & Exhibition on Mass Mining*, University of Chile, Santiago, pp. 693–702, https://doi.org/10.36487/ACG_repo/2063_47

Donaldson SN, Hollis BL, Wynn K, Potts L & Evans, M 2021, 'Real tine LHD dispatch optimisation with ratio targets at Newcrest's Cadia Valley Operations', *Proceedings of the Underground Operator Conference*, Australasian Institute of Mining and Metallurgy, Melbourne, pp. 206–213.

Eustace, C, Lagace, D & Bobbermen, L 2020, 'Simulation capacity analysis for the Carrapateena block cave', in K-H Bae, B Feng, S Kim, S Lazarova-Molnar, Z Zheng, T Roeder & R Thiesing (eds), *Proceedings of the 2020 Winter Simulation Conference*, Institute of Electrical and Electronics Engineers, New York.

Eustace, C, Hynard, K, Mariager, S & Fargher, M 2024, 'Hybrid configurations for block cave operations', in Daniel Johansson & Håkan Schunnesson (eds), *MassMin 2024: Proceedings of the International Conference & Exhibition on Mass Mining*, Luleå University of Technology, Luleå, pp. 1522–1536, https://doi.org/10.36487/ACG_repo/2435_R-01

Khodayari, F & Pourrahimian, Y 2015, 'Mathematical programming applications in block-caving scheduling: a review of models and algorithms', *International Journal of Mining and Mineral Engineering*, vol. 6, pp. 234–257.

Mariager, S, Zhang, Q, Eustace, C, Hegarty-Cremer, S & Nguyen, K 2022, 'A simulation comparison of operating strategies for electrified block cave mines', in Y Potvin (ed.), *Caving 2022: Proceedings of the Fifth International Conference on Block and Sublevel Caving*, Australian Centre for Geomechanics, Perth, pp. 315–330, https://doi.org/10.36487/ACG_repo/2205_21

Morrison, DM & Labrecque, PO 2020, 'High-capacity production systems for Block Caving mines', in R Castro, F Báez & K Suzuki (eds), *MassMin 2020: Proceedings of the Eighth International Conference & Exhibition on Mass Mining*, University of Chile, Santiago, pp. 703–715.