Caving 2022 - Y Potvin (ed.)

doi:10.36487/ACG_repo/2205_60

Oyu Tolgoi and Rio Tinto partnership with Palantir Technologies to provide effective geotechnical risk management

TH Moorcroft *Rio Tinto, UK*

K Simanjuntak *Rio Tinto, Mongolia*

O Dorjsuren *Rio Tinto, Mongolia*

M Sanaakhorol *Rio Tinto, Mongolia*

E Enkhtaivan *Rio Tinto, Mongolia*

G Watt *Rio Tinto, Australia*

V Eickhoff *Palantir Technologies, Australia*

L Cerny *Palantir Technologies, UK*

C Deasy *Palantir Technologies, Australia*

T Zimmermann *Palantir Technologies, Australia*

Abstract

Oyu Tolgoi, a copper and gold mine in Mongolia, is one of the deepest and largest block cave mines in the world. The depth and size, combined with demanding rock conditions, brings new technical and operational challenges in providing effective geotechnical risk management. At Oyu Tolgoi, Rio Tinto has partnered with Palantir Technologies, the enterprise software company, to support geotechnical risk management. This paper describes how effective geotechnical risk management practices are facilitated through a data-driven operations platform that connects data, analytics, and business teams to a common foundation. Through data integration, and a strong data foundation, Oyu Tolgoi engineers can monitor and act upon the mine's aggregated data asset against their trigger action response plans and production data with cave management plan rules.

Data-driven cave monitoring workflows improve visibility into operations at the cave, support incident prevention, provide transparency for optimal collaboration between geotechnical and production engineers and inform decisions on cave construction sequencing and draw strategy. This amounts to a new and stronger footing from which to provide effective geotechnical risk management.

Keywords: *geotechnical monitoring, instrumentation, block caving, automation, data model, engineering, Palantir Technologies, Oyu Tolgoi, Rio Tinto, Mongolia, copper, gold*

1 Introduction to Oyu Tolgoi

1.1 General context

Oyu Tolgoi, in the South Gobi region of Mongolia, is one of the largest known copper and gold deposits in the world. It is also one of the most modern, safe, and sustainable operations in the world. Open pit mining began at Oyu Tolgoi in 2013. In January 2022, underground operations commenced. This step unlocks the most valuable part of the mine, with sustainable underground production expected in the first half of 2023.

1.1.1 Mine location

Oyu Tolgoi is in Mongolia, south of the capital, Ulaanbaatar, in the South Gobi Desert, about 80 km north of the Chinese border.

1.1.2 Geologic background

The underground operation at Oyu Tolgoi is currently focused on the Hugo North deposit. The Hugo North deposit is a porphyry deposit associated with quartz monzodiorite (QMD) intrusions. The intrusions are below the Devonian and Carboniferous sequence of sedimentary and volcanic rocks. The host rocks for the Hugo North deposit dip steeply to the east and are volcanic and quartz monzodiorite (QMD). Dacite tuff and a breccia (IGN) sequence overlay a set of basalt flows and minor volcaniclastic strata (Va). This is the lowest level of the stratigraphy. A group of younger sub-vertical dikes crosscut the deposit. These dikes are rhyolitic, hornblende biotite andesite, dacite, and basalt (David & Kavalieris 2012; Khashgerel et al. 2008).

1.1.3 Commercial background

Oyu Tolgoi is one of the deepest and largest block cave mines in the world. The mine is a partnership between Rio Tinto, Turquoise Hill Resources, and the Government of Mongolia. Oyu Tolgoi has one of the world’s largest deposits of copper and gold. Within a decade, Oyu Tolgoi is widely expected to be in the top five of the world’s largest copper mines, generating billions of dollars of revenue for the Government of Mongolia over its lifetime. By 2028 Rio Tinto expects Oyu Tolgoi to triple production from 2021 levels to around 500,000 t of copper per year from the open pit and underground combined (Rio Tinto 2022).

1.2 Lessons learned from other Rio Tinto operations

To fully manage the stress loading and convergence on the extraction level, a comprehensive excavation stability monitoring system is called for. This system is comprised of a dense array of automated sensors installed in the rock mass, regularly documented inspections, and an analysis system to indicate when damage to excavations may occur. The results are used by subject matter experts when considering whether to enact a repair or ‘use draw’ to alleviate loading. The resultant actions are ultimately aimed to increase the recovery of the orebody. The system reaches into all excavations on the mine to ensure that safe access to the orebody is maintained**.**

Undercutting and drawbelling phases bear a significantly increased risk of major operational failure, as there is a higher risk of detrimental events, such as air blasts, strain bursts and material in rush, that could compromise large parts of the infrastructure and pose a serious personal safety risk. Whilst ad hoc alarming in response to major events, e.g. major seismicity, is well handled by existing systems at Oyu Tolgoi, preventative workflows that monitor defined risk factors over time, and allow for cross-functional decision making, that were constrained by the scale and complexity of data sources as mentioned above.

In response to these major risks, the geotechnical engineers devised a cave management plan (CMP) that is compliant with underground safety standards and safety risk assessment guidelines CMP and includes 25 trigger action response plans (TARPs), which are aimed at detecting risks as early as possible to prevent natural and operational incidents from happening. By integrating these TARPs into the digital cave monitoring workflows and stable data pipelines, geotechnical engineers will be able to assess accurately the cave’s current and prospective risk picture and take preventative action, such as decreasing the local draw rate to prevent excessive airgap formations.

2 Concepts and operational practices

2.1 Excavation risk rating

Excavation risk rating (ERR) is based on the exposure to personal and operational criticality. This provides a method for identifying those excavations which pose the highest safety and production risks. The ERR category assesses the current exposure, driving the applicable inspection frequency and TARP limits.

The excavation ERR designation has been broken down into four categories, ERR-1 through to ERR-3. Each tag corresponds to a particular risk tolerance both from a safety and production standpoint, with ERR-1 having high safety and high production risk and ERR-2 through to four de-escalating from this.

ERR-1 (excavation risk rating type 1): These excavations are used for areas where personnel congregate, including, but not exclusively: offices, training rooms, mechanics shops, lunchrooms, refuge chambers and shaft stations. This class also includes other mine critical installations such as crusher chambers, conveyor transfers, electrical installations, shafts, fan installations and magazines.

ERR-2 (excavation risk rating type 2): This classification is used for drives that undergo significant stress changes as a result of mining development or production activities, including, but not exclusively: excavation drives, drill drives, return airways, fresh airways, haulage drives and active development headings. Personnel is not customarily expected to visit/operate in these excavations.

ERR-3 (excavation risk rating type 3): Typically, this classification is applied to areas where personnel regularly travel but do not dwell in one specific place. The area is not subject to expected significant stress increases. Such regions include off footprint accessways, ventilation drives, conveyor drives/declines and service declines, and their associated crosscuts, stockpiles and utility bays.

2.2 Change to the use of an excavation

Any change to the use of an excavation, or the planned lifespan, requires management of change (MOC) to ensure that it remains fit for purpose to perform the new function. This ensures that excavations designed for muck bays or other temporary usage are not converted to offices or areas where people may congregate, or critical infrastructure may be relocated without upgrading the ground support. Excavations used for offices or critical housing infrastructure require a higher Factor of Safety due to the increased personnel exposure (Stegman et al. 2018). Changing the uses of excavation may result in a change to its ERR classification.

2.3 Trigger action response plans

To increase excavation reliability and mitigate impacts and risks to the operation, actions are taken at the earliest opportunity to prevent further deterioration. Oyu Tolgoi manages these actions using (TARPS). The trigger level of the excavation is re-evaluated following rehabilitation work. In most cases, the levels are reset to zero as the majority of rehab at Oyu Tolgoi involves the installation of new primary support elements, resulting in a support capacity and Factor of Safety being returned to the design levels. The excavation performance is measured through damage mapping and instrumentation to correspond to four levels of damage. Each monitoring method and ERR location has its own unique set of TARPs. All, however, follow the general structure outlined below.

Level 0 (normal): State in which convergence behaviour and visual status are acceptable and considered sustainable long-term. The response is to carry on monitoring and operating as per normal activities.

Level 1: This level will be reached when approximately 1% of primary support element strain is observed. In response, the ground support's degradation and deterioration will be identified and its extent quantified. Additional monitoring or support or scaling of loose material may be required.

Level 2: Further deterioration develops beyond level 1 when approximately 2% of the primary support element strain is observed. Ground support work orders shall take other remedial measures to control the ground stability when an excavation reaches this level. Access through the area affected may be restricted based on the severity, extent of the damage and how rapidly the damage may be escalating.

Level 3**:** Further deterioration develops beyond level 2 when approximately 5% of the primary support element strain is observed. As an immediate response, the area will be barricaded to restrict travel, and a rehabilitation plan will be required to ensure the drive's longevity.

3 Mine data asset in Palantir Foundry

3.1 Palantir Foundry platform at Oyu Tolgoi

Oyu Tolgoi's operations are supported by Palantir's Foundry platform to remove the friction between back-end data management and front-end data analysis.

The front-end tools have allowed users to tap into the power of Oyu Tolgoi's data. By using:

- A central data foundation to drive collaboration and discovery across functions.
- A common ontology to turn a complex data landscape into an intelligible representation of the entire organisation.
- Datasets and analyses that feedback into the platform to allow users to build on one another's work.
- Legible data lineage to connect insights to the data and logic that feed them.
- Diverse analytical tooling to supercharge traditionally non-technical functions and accelerate advanced analytical initiatives.

On the back-end, a suite of best-in-class capabilities for data integration runs on data and business logic in tandem. This involves:

- Versioning semantics to keep data and business logic in sync.
- Dynamic, systemwide security and access control to replace unreliable one-off policies.
- Branching of code, analyses, and reports to enable safe experimentation.
- A microservice architecture that constantly incorporates new technologies.
- Open APIs and data formats to interoperate with an organisation's entire data ecosystem.
- Flexible data protection frameworks to keep up with evolving regulations (such as the General Data Protection Regulation) and industry best practices.

3.2 Integration of mine data

3.2.1 Data connection and data integration capabilities

Foundry provides tools for connecting to and interacting with virtually any type of third-party system, regardless of the scale, structure or schema of the data exchanged. The platform's graphical data integration interfaces allow both technical and non-technical users to bring data into the platform, as well as push it to external systems, as needed. These interfaces control ready-made connectors for industry-standard source types, such as SQL databases, external file-based systems, cloud-hosted databases and data streams. If required, the platform also supports using custom connectors for sources that are not supported out of the box. These ready-made and custom connectors ensured that Foundry slots into the Oyu Tolgoi information technology landscape with minimal disruption to existing systems and processes.

3.2.2 Data integration at Oyu Tolgoi

Both production and geotechnical information are fed into Foundry by Oyu Tolgoi to provide the foundation for analysis and operational workflows. The data is ingested from many different source systems (Batkhuu & Johnson 2020; Savage et al. 2018).

The following types of data are integrated into the platform:

- Geotechnical inspections: damage mappings.
- Instrument readings: convergence arrays, multipoint borehole extensometers (MPBXs), smart markers, cave trackers, open boreholes, time-domain reflectometers (TDRs), mobile scanning.
- Seismic events.
- Installed ground support design and rehabilitation work orders.
- Ground support quality control: shotcrete thickness, bolt pull tests, spacing and encapsulation.
- Development blasting quality control: overbreak and underbreak (percentage and volume), round length.
- Undercut information: design of the undercut level, ring position and statuses.
- Ring firing on the undercut level.
- Production mucking data (tons and buckets).

3.2.3 Data processing at Oyu Tolgoi

Once the data has been integrated, it needs to be prepared before it is made available to the end-users in the Decisions and Operations layers. Those data transformation steps include cleaning datasets, joining datasets together, enriching the data. This sequence of transformations composes a data pipeline. Each data source is usually prepared in its own pipeline.

After the initial preparation steps, processed data from different sources is often passed together as input for further processing and analysis in other data pipelines.

Figure 1 shows a simplified version of the cave management pipeline. The left side shows the source systems which the data is ingested from. In this case, mine design data is ingested from Deswik, mucking and drilling data is ingested from Pitram and a version-controlled set of undercut management rules is available as input as well.

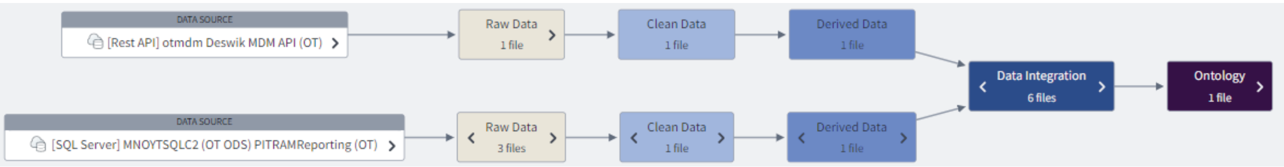

Figure 1 Simplified version of the cave management pipeline

The integration stage joins the data together and performs the required logic to transform the raw data into information regarding the compliance state of the undercut to the rules laid out in the CMP. This returns a set one for each of the cave management rules, the resulting datasets are then mapped to a flexible data model called 'ontology'. Objects and their relationships to each other form the core of the ontology, the digital copy of the underground at Oyu Tolgoi. This is the common platform for downstream workflows, and aims to make data understandable to humans, not in terms of rows and tables but rather as tangible items which relate to the real world (Palantir Technologies 2022a, b, c). Ontology concepts are covered in detail in Section 4.

3.3 Data quality, integrity, and transparency at Oyu Tolgoi

As data quality and data integrity are the foundations for any data project, it is instrumental that data engineers are first empowered with the tools and transparency needed to build a quality data pipeline. Below, some of the unique tooling and capabilities of Palantir Foundry that help ensure data quality are highlighted.

3.3.1 Leveraging Foundry data lineage tools

Understanding where the data has been and where it is going allows data engineers to trace potential data quality issues across the platform, ensuring verified data is available to every applicable team on the ground. For instance, what is the root source of the data? When was it updated? What transformations, filters, or enrichments have been made on this data?

Exploring upstream inputs: Foundry allows data engineers and users at Oyu Tolgoi to trace back, investigate, and resolve data quality issues at their source. Oyu Tolgoi has an approach to integrate many different instrument types into a unified TARP and inspection process as discussed in more detail in Section 2. When a section of the mine was unexpectedly marked as ‘out of compliance with the required inspections’, data engineers were immediately able to look upstream at every parent dataset that fed in that transform, tracing every change and merge all the way back to the source. The engineers were able to inspect how the data has been transformed over time, providing critical debugging capabilities. Using these tools, they were then able to determine that convergence data had been incorrectly entered in meters rather than mm, resulting in a 1,000-fold change in the measured tunnel size. The data was then corrected at the source and a new data health check was added in the pipeline in Foundry.

Tracking downstream effects: Engineers at Oyu Tolgoi use the same data lineage tooling to proactively trace where data quality issues may impact downstream artifacts. This has proved most useful when migrating between systems or carrying out major updates. This tooling has enabled these users to search along the pipeline for where there might be issues with how the data was prepared with code. Some common issues include exploding joins, dropping or filtering data prematurely, or missing data. The users can zero in on where in the pipeline this regression is introduced and remediate. Data lineage allows data engineers to proactively alert downstream consumers and fix the data flow before issues compound.

3.3.2 Ensuring the health of mine data asset

Foundry automates checks designed to detect aberrations in the data – whether timeliness of data updates, completeness, consistency or even identify missing contents – to ensure robust data quality at scale. The pipelines can also listen for these checks and prevent the propagation of the data downstream if any of the checks fail.

For instance, users can define a target data schema, such as the column names and expected content types, and establish validations that ensure that changes in the data model are flagged early. If validation fails, Foundry can trigger alerts for review and remediation. In addition to ingestion, Foundry also performs data pipeline health checks that trigger alerts all along the data pipeline. Some examples of data health checks include those listed below: data staleness, data shape, uniqueness, null percentage, and expected values.

The geotechnical team at Oyu Tolgoi has 100s of data health checks spread thought out the data pipelines. In some areas of the mine, tunnel displacement is measured with tape extensometers. The tape extensometer data is stored in a database. When that data is ingested into Foundry a check is run to ensure that column names and data types match the expected values. Later in the data processing pipeline checks are run to ensure that the absolute displacement on any convergence array is less than -5 mm or more than 400 mm. There are checks in place to ensure that data remains fresh in the system and others which ensure that new data keeps being added. If any of these fail, then a member of the geoscience data management team is informed, at which point they will inspect the data and then take the matter up with the appropriate engineer or department. Having the data contently checked has led to a major reduction in false-positive TARP triggers and is improving the credibility of the geoscience team.

3.4 Security and permissioning at Oyu Tolgoi

Data security is a central pillar of Foundry's design. This allows all of the mine's data and every action to be protected by industry-leading security tooling and governed in accordance with Rio Tinto's cybersecurity policy.

One of the primary divers for Oyu Tolgoi's adoption of Foundry was security capabilities. These sit across every layer of the platform. As such, they govern all operations without exception. Oyu Tolgoi data administrators can configure security policies centrally instead of having to manage multiple policies across separate systems. This helps avoid both erroneous breaches and overprotection.

Oyu Tolgoi protects its production and exploration grade data and drillhole information while still making it available to appropriate, approved, authorised staff. For example, production sampling data can be accessed via the platform by the scheduling and geology team who are carrying out reconciliation. The engineers and geologists can access and manipulate the data using the tooling within the platform however if they were to share their work with a 3rd party who lacked the correct access to any of the upstream datasets, the 3rd party would be unable to view the source data or the analyses.

Seismic data is not as sensitive as grade but should still be handled with care and viewed in context as laid out in the seismic hazard management plan. If seismic data is not correctly interpreted and well explained, then an inaccurate impression could be gained regarding the risk to personal safety. This could be very concerning for the workforce in general, who for the most part are not seismic experts. As a result Oyu Tolgoi has access controls in place to ensure that only trained staff can access and use this data asset.

3.5 Data processing and representation

3.5.1 MPBX displacement contours

The platform allows complex multi-dimensions operations to be carried out in a controlled, approved, repeatable and auditable fashion. Data from MPBXs is transformed to be represented as a set of displacement contours. These contours are made available in geojson format. The input dataset has a relatively coarse spacing at ~50 m between arrays. To produce smoother contours, the displacement is projected onto a much finer mesh grid. The transform described above is performed using grid data from Scipy (Virtanen et al. 2020) and contour from Matplotlib (Hunter 2007). The final stage of contouring is completed with QuadContourSet.

3.5.2 Lithology for level plans

It is tempting to think that all spatial data in an underground mine should be viewed in 3D. The Oyu Tolgoi and Palantir teams have challenged this preconception. In some cases, viewing data in 3D clearly adds value. However, when many sets of data are combined and the data scale grows, performance becomes a consideration. There comes a tipping point between the advantage of 3D visualisation for allowing intuitive leaps based on complex multi-level relationships versus the speed and simplicity of a well-chosen section or level plan.

Geological models of the Oyu Tolgoi orebody are maintained and updated by the geology team. The models are stored in Deswik MDM. To ensure that all plans and sections have up-to-date geology shown on them, the data feeds update this triangulation many times a shift and construct 2D polylines. This is powered by the python library Trimesh. The resulting polygons are then clipped against each other based on their stratigraphic order to remove any overlap.

3.5.3 Representation of damage mappings as polylines

The damage mapping system used at Oyu Tolgoi provides measurements on a 1 m spacing, this results in over 100,000 data points if the whole mine is visualised at any given time. While visualising so many data points is not an impossible task, it is a challenge to maintain high performance and interactivity. Oyu Tolgoi significantly reduces the computational challenges by combining neighbouring points which share the same damage mapping values into polylines. This improves the performance of the application while maintaining the same granularity of information relayed to the end-user.

3.5.4 Geospatial grouping of monitoring data

Not all monitoring methods have the same geographical coverage. For example, an MPBX gives you displacement at a point in space, while damage mapping provides quantitative data spread over a large area. Oyu Tolgoi has many MPBX deployed across the production footprint; these are read automatically every hour. A simple grid (Figure 2) is used to deal with the disparity in coverage both in space and time. A geospatial join is carried out, allowing nearby instrumentation to be linked to the grid cell. The maximum TARP level of any instrument within a grid cell is used to set the required inspection of all the monitoring within that location. The grid used at Oyu Tolgoi is composed of hexagons 17.5 m wide and 35 m long; this sizing was chosen as a factor of the production drive spacing. The grid has been rotated by 55 degrees to align with the production drive and the principal stress direction. Figure 2 shows how data from point locations such as MPBXs and convergence arrays are mapped into the grid.

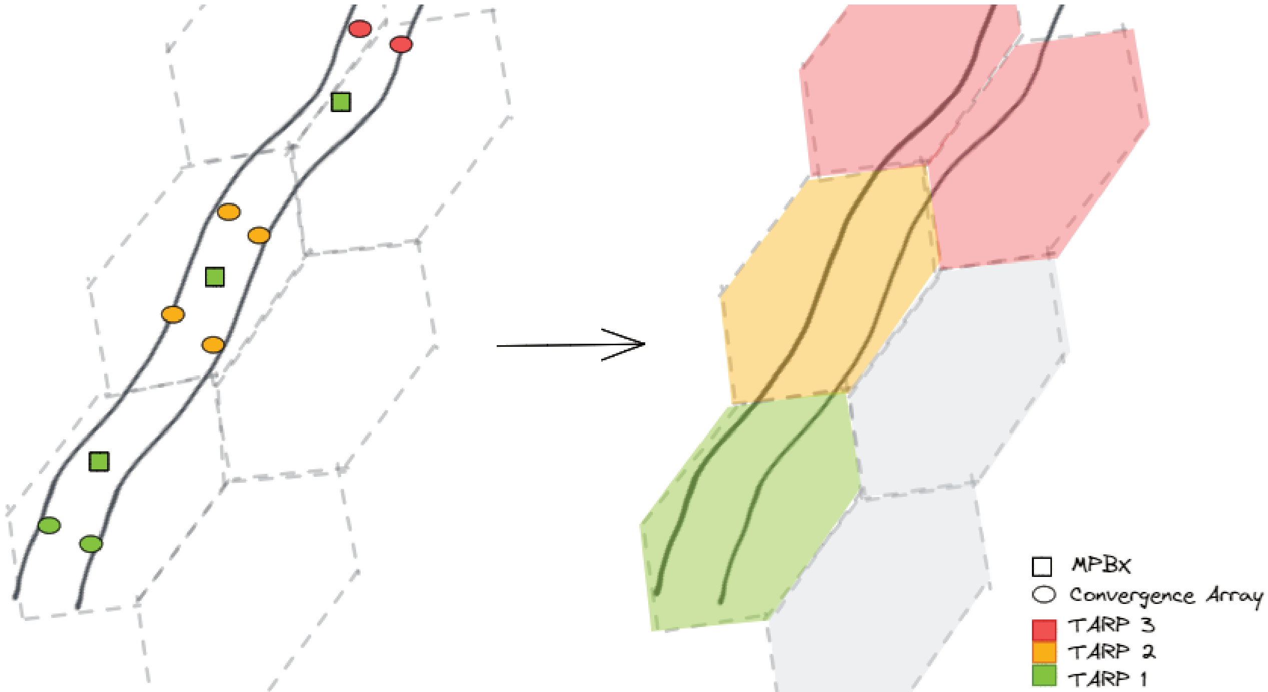

Figure 2 Mapping point locations into the grid

4 Building a common operating picture

4.1 Oyu Tolgoi ontology in Foundry

Once data has been integrated into Foundry, it is mapped to a flexible data model or 'ontology'. Ontologies in Foundry are designed to reflect accurately how organisations actually operate and to evolve over time as workflows and needs change. Most people do not speak the language of data (rows, tables, primary keys, etc.) but instead communicate and think in terms of objects, relationships, and actions (e.g. loader, MPBX instrument, drawbell).

These are the building blocks of ontologies in Foundry. Fitting data to these familiar concepts empowers users to interact and leverage data in a more intuitive fashion and to make better, faster, and more data-driven decisions.

By making the ontology the central focal point of data and analysis, users gain several operational advantages. Most notably, everyone operates on a single common operating picture. Foundry integrates sensor, production, and geotechnical data and ensures users are made aware of the freshness of data as well as its providence. This helps ensure the data is trusted and allows teams to work more collaboratively with greater speed and accuracy.

In the context of underground mining, there are a diverse set of users including geotechnical and production engineers, and managers, Their primary points of engagement with the mine data asset is through the Oyu Tolgoi ontology via Foundry applications that leverage the ontology.

The Oyu Tolgoi ontology provides a rich and integrated data asset, which both production and geotechnical teams leverage to answer ad hoc questions and expanding their ability to troubleshoot beyond the TARPs laid out in the CMP. Senior geotechnical engineers especially are able to perform deep-dive analyses to find hidden root causes of alerts and subsequently take action based on a differentiated risk picture in a timely and effective manner. The Oyu Tolgoi ontology serves as the common data foundation for geotechnical and production workflows (described in Section 5), allowing team collaboration and supporting operations and data-driven decision making on the ground.

Figure 3 shows a small section of the Oyu Tolgoi ontology which shows the interrelationship of the undercut management rules with monitoring data. A user can start their exploration at any point in the ontology and navigate via the links between object types to explore the digital world. For example, a geotechnical engineer can use the ontology to quickly investigate a non-compliance with the rules (e.g. the local lead/lag rule): by starting at the local lead/lag TARP object of interest and traversing its links to rings, the engineer would quickly see which rings have been recently blasted that cause that non-compliance, and from there the engineer can easily navigate to the monitoring instrumentation which is installed in that area (e.g. the MPBX instrument object type) to see if the non-compliance had an impact on the instrumentation readings and stability of the mine.

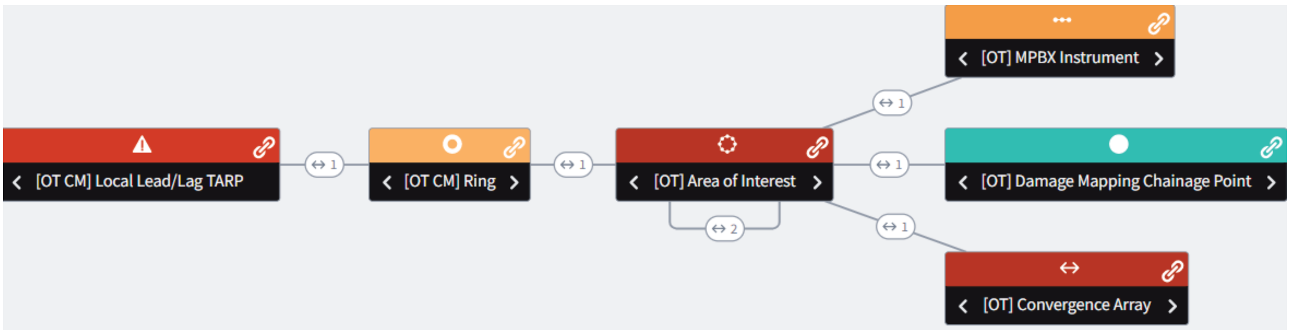

Figure 3 A small section of the Oyu Tolgoi ontology

4.2 Combining data

4.2.1 Combining production and geotechnical data

Figure 4 shows a plot of how ground movement is detected by MPBX instruments – a clear correlation can be seen between undercut ring firing and ground movement. This is an example of how integration of production data with geotechnical monitoring data and sharing a common ontology enables collaboration between geoscience and production teams.

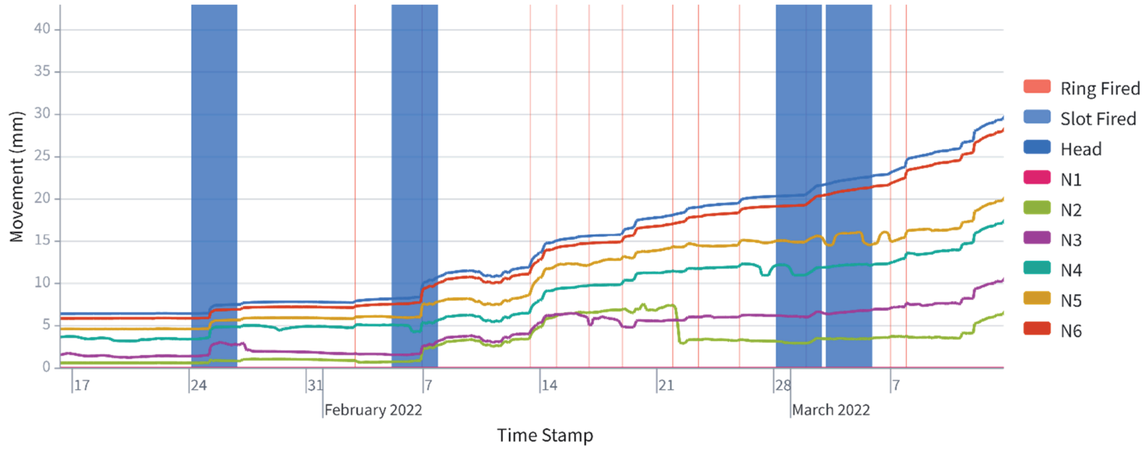

Figure 4 Rings and slots versus MPBX nodes

4.2.2 Combining instrument data

There is likely to be various types of monitoring installed and data collected in any area, given the strengths and weaknesses of each monitoring method in terms of accuracy, repeatability, and coverage. It is considered unlikely that the TARP levels for the different systems will directly align with both varying timeframes for data collection and repeatability. It is also probable that two instruments of the same type in the same location may have different TARP levels: particularly MPBX instruments installed in an array pattern in the back and the ribs of the excavation. Using the grid cells described in Section 3.5.4 it is now possible to surface the most alarming TARP state on an instrument and have it inform required inspections on those around it. This unlocks several workflows.

Consider the following hypothetical example. When the primary crusher complex is at TARP level zero, it should be inspected by a geotechnical engineer every three months. Six weeks after its last inspection an MPBX shows an increased rate of movement, triggering TARP level one. At TARP level one the inspection frequency is four weeks. The location is now out of compliance, geotech is informed and carries out the required inspection. Appropriate actions are then taken.

From a geotechnical viewpoint having inspections driven by levels of ground support damage and the criticality of the location is intuitive, however, the implementation of this in a systematic way has proven to be extremely difficult prior to the use of geospatial aggregation.

4.2.3 Combining development blasting and MPBX data

Having all this data available in a single platform allows some traditionally time-consuming workflows to be optimised in terms of time and effort required. It is now possible for everyone to view the same up-to-date data, with no manual manipulation required.

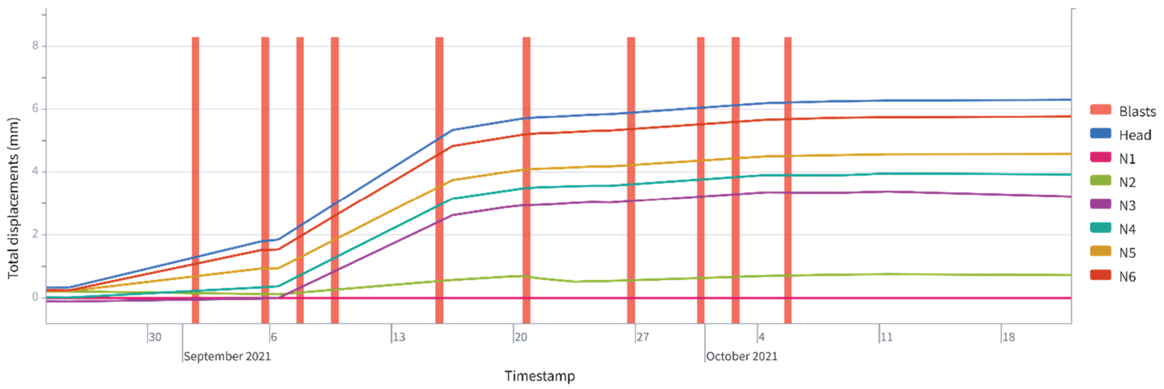

Figure 5 MPBX versus development blasts

Figure 5 shows MPBX data taken from a location proximal to a newly open drawpoint. The dataset containing MPBX measurement is combined with the blasting date of the drawpoint. The impact of the blasting can be clearly seen at the MPBX when the data is viewed together. Prior to the use of the platform both the MPBX and blasting data would have to be exported and manipulated manually. This processing makes it very difficult to replicate the process and track the providence of the source data, ultimately reducing its value.

5 Data-driven operations and decision making

5.1 Effective risk management

Oyu Tolgoi has increased the safety and assurance of the Oyu Tolgoi underground mine by implementing a proactive and an effective geotechnical risk management system. Geotechnical risk is measured using multiple metrics based on instrument readings from physical inspections and various sensor data, all of which are disparate and stored across various systems. By using Foundry, the engineers have a real-time risk picture of the mine to ensure all activities underground are tightly controlled and in compliance with safety standards that minimise the risk for workers and operations.

A single interface (Figure 6) allows users to view production and geotechnical information in the same place. This enables subject matter experts to see the full picture in one environment, rather than having to jump between many different systems to understand what is occurring within the cave.

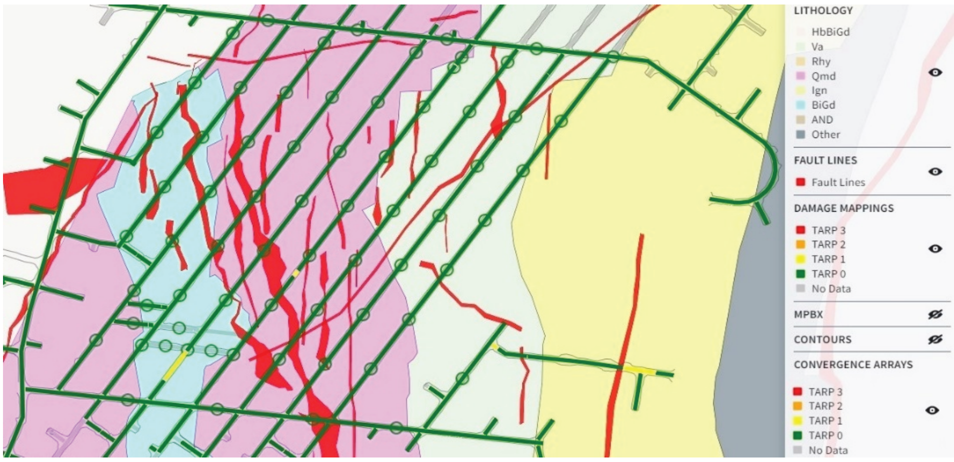

Figure 6 Map interface showing the Oyu Tolgoi footprint

5.1.1 Inspection timing and uncertainty

It is impractical for every location in the mine to be continually inspected. A matrix of ERR and TARP levels is used to determine which area prioritises inspection. Locations with low ERR and high TARP levels are being regularly inspected. The combination of three ERR and four TARP levels results in 12 different inspection regimes throughout the mine, with different scopes and cadence of inspection. Although this may seem complex, it is not significantly different to the traditional behaviour of the geotechnical team, where areas of damage or matters critical to operations would be inspected more frequently.

It is important to prioritise risk management attention to different areas based on geotechnical risk. However, simple metrics like ‘days till out of compliance’ are not effective for cross-comparison across regimes, as there is a significant difference in the inspection frequency and purposes of inspection regime. For example, an ERR-1 location with a TARP value of two would require daily inspection, whilst an ERR-3 with a TARP of zero would only need to be inspected once a year.

Instead, the Oyu Tolgoi team has developed the concept of uncertainty risk which prioritises and compares different locations with different ERR and TARP values. Uncertainty risk is a simple way to quantify the likelihood that conditions have changed since the data was collected. The scale of uncertainty ranges from zero to one. Uncertainty risk centres around the idea that the moment the inspection is carried out, there is no uncertainty (uncertainty = 0), and the measurement is still valid as time progresses. Then, the level of surety in that observation decreases. When the uncertainty reaches 0.8 it is considered out of compliance. This is deliberately not the value of one to allow for some differentiation in the levels of not compliant should readings not be taken at the required frequency. Once a task is out of compliance, uncertainty rapidly increases to one.

5.2 Cave management undercut monitoring

Both cave initiation and cave propagation are crucial steps in the development of the mine. These processes also bear a significant risk to geotechnical safety if compliance targets are neglected, or sequence rules are breached. For example, excessive drawing can expose the underlying structures to the risk of air blasts. This can cause caving rocks to fall into insufficiently buffered drawpoints and lead to loss of machinery or fatal human injuries. Alternatively, blasting drawbells within the correct rate, but outside a planned sequence, can cause imbalances to the local caving process and equally result in dangerous cave instability. To coordinate these critical steps safely, the geotechnical and planning teams have laid out cave management rules and TARPs, which are based on research and risk models. These rules must therefore be followed, monitored and actioned in the case of non-compliance, and based on reliable and integrated information to support well-informed and speedy decision making.

The following rules are laid out in the Oyu Tolgoi undercut management plan and have been implemented in Foundry.

- Undercut advance time distribution: This rule is designed to control the number of rings blasted per drill drive over a specific time period.
- Undercut advance rate/distribution pre-drill drive per month (m/drill drive/month): This rule ensures that the total number of blast events per 30 days is controlled.
- Undercutting rate (sqm/month): Calculated as the total number of rings blasted in the last 30 days (moving window) across the full undercut level.
- Local lead/lag: This applies to the north edge of the cave. This is the maximum allowed lead/lag between two connected adjacent drives calculated from west to east.
- Instantaneous lead/lag: This applies to the southern face of the undercut. After the cave is connected there is a maximum of 30 days grace period to respect the prescribed lead/lag.
- Undercut length: Measures the linear distance between the two undercut cave front extremes from west to east.
- Undercut front direction: The undercut front direction rule controls the relationship between the maximum stress direction and the undercut front.
- Veranda length: This is measured from the most recently blasted drawpoint to the edge of the undercut.

The cave management application provides an overview of the current state of the cave. The application is available to both engineers on the ground as well as caving subject matter experts throughout the Rio Tinto group. It allows for timely identification of non-compliance to the cave management rules, empowers users to investigate root causes and then see the impact on the mine stability via instrumentation readings, as well as to act, all in the same place.

6 Conclusion

The partnership between Oyu Tolgoi and Palantir Technologies is providing a solution that supports the teams in effectively managing geotechnical risk. By democratising data access, we make it easy for different subject matter experts to view and access the information they need. As a result of the collaboration, it is now possible for engineers to carry out complicated, auditable, and transparent transformations on mine data. Data governance components in place help ensure high quality of the data, traceability, security, appropriate permissions, and workflow stability. This has allowed Oyu Tolgoi engineers and geologists to collaborate more effectively, supporting more informed operations and data-driven decisions.

Acknowledgement

The authors thank Mr Christopher Carlson and Liam Johnson for their geotechnical engineering input on this project.

The authors also thank the Palantir team members who are part of this partnership for their essential and numerous contributions: Conor Neilson, Kelvin Lau, Emily Deng, Cormac Biggar, Viktoria de La Rochefoucauld, David Hack, Ayberk Arseven, Riccardo Ceccarelli, Monty Evans, Davìd Markey, Michael Shafer, Lisa Xie, and Xavier Durand-Smet.

References

Batkhuu, B & Johnson, LA 2020, 'Ground support strategy for optimisation of time schedule, Oyu Tolgoi underground mine', in J Wesseloo (ed.), *UMT 2020: Proceedings of the Second International Conference on Underground Mining Technology*, Australian Centre for Geomechanics, Perth, pp. 367–376, doi.org/10.36487/ACG_repo/2035_19

David, C & Kavalieris, I 2012, 'Geologic overview of the Oyu Tolgoi porphyry Cu-Au-Mo deposits, Mongolia', in JW Hedenquist, M Harris & F Camus, *Geology and Genesis of Major Copper Deposits and Districts of the World: A Tribute to Richard H. Sillitoe*, The Society of Economic Geologists, Littleton, pp. 187–213.

Hunter, JD 2007, 'Matplotlib: A 2D graphics environment', *Computing in Science & Engineering*, vol. 9, pp. 90–95.

Khashgerel, B, Imants, K, Hayashi, K 2008, 'Mineralogy, textures, and whole-rock geochemistry of advanced argillic alteration: Hugo Dummett porphyry Cu–Au deposit, Oyu Tolgoi mineral district, Mongolia', *Mineralium Deposita*, vol. 43, no. 8, pp. 913–932.

Palantir Technologies 2022a, *Palantir Blog*, blog.palantir.com

Palantir Technologies 2022b, *Foundry Technical Overview*, viewed 7 July 2022, palantir.com/assets/xrfr7uokpv1b/54mwrnjeu6Y55Lj24RS1Sx/2e2b7b7b4fe2f0ad9ee6a95e976c03c6/FfB_Technical_Overview_v4.pdf

Palantir Technologies 2022c, *Palantir Foundry Launch*, palantir.com/palantir-foundry/21-launch/

Rio Tinto 2022, *Oyu Tolgoi Partners Reach Comprehensive Agreement and Approve Commencement of Underground Mining Operations*, press release, viewed 7 July 2022, riotinto.com/news/releases/2022/Oyu-Tolgoi-partners-reach-comprehensive-agreement-and-approve-commencement-of-underground-mining-operations

Savage, E, Moorcroft, T & Talu, S 2018, 'Using a structured data management system to improve mine development efficiencies', in Y Potvin & J Jakubec (eds), *Caving 2018: Proceedings of the Fourth International Symposium on Block and Sublevel Caving*, Australian Centre for Geomechanics, Perth, doi.org/10.36487/ACG_rep/1815_61_Moorcroft

Stegman, C, van As, A & Peebles, E 2018, 'Past learnings focus innovative solutions to future cave mining', in Y Potvin & J Jakubec (eds), *Caving 2018: Proceedings of the Fourth International Symposium on Block and Sublevel Caving*, Australian Centre for Geomechanics, Perth, pp. 37–42, doi.org/10.36487/ACG_rep/1815_0.3_Stegman

Virtanen, P, Gommers, R, Oliphant, TE ... SciPy 1.0 Contributors 2020, 'SciPy 1.0: Fundamental algorithms for scientific computing in Python', *Nature Methods*, vol. 17, pp. 261–272.