PipeTwin: A Scalable Framework for Digital Twin Deployment in Mineral Suspension Pipelines

P. Xie [a], E. Zheng [b], D. Nguyen [a], A. Chryss [b] and L. Hetherton [a]

[a] CSIRO Data61, Melbourne, Australia
[b] CSIRO Mineral Resources, Melbourne, Australia
Email: lachlan.hetherton@data61.csiro.au

Abstract:

The transport of mineral suspensions through pipelines is central to many mining and mineral processing operations, yet it presents significant challenges due to the complex, multiphase nature of suspension flow. Real-time monitoring and modelling are essential to ensure efficient operation, detect anomalies, and optimise performance. Digital twin technologies, when integrated with live instrumentation, offer a powerful means to address these challenges. We introduce PipeTwin, a flexible and scalable framework developed on the Workspace platform, designed to integrate real-time data from a diverse array of pipeline instruments (including rheometers, pressure transmitters, flowmeters and density meters) into digital twins. Recognising the heterogeneity of instrumentation and control systems across laboratory, pilot and industrial environments, PipeTwin is built to be protocol-agnostic, supporting both standard and custom communication methods. In our current lab-scale deployment, data acquisition is performed via VPN or LabVIEW network communication to stream live signals. All incoming data, regardless of source or protocol, is unified and managed within a cloud-native infrastructure using Google Firestore, allowing consistent time-series handling and scalable data processing. Once collected, the data is streamed to the web-based application, which collects user input parameters and send together with the data to Workspace platform, where it is interpreted through an integrated pipeline transport model to generate real-time performance insights. This paper presents the architecture of the PipeTwin system, highlighting its adaptability to various instrumentation setups, its ability to support different communication protocols, and its scalability from lab to industrial environments. By enabling flexible, robust integration between instruments and digital twins, PipeTwin demonstrates a practical foundation for adopting digital twin technologies in the mining and processing industry. The underlying integration framework of PipeTwin also offers broader applicability to other complex infrastructure systems, such as modern urban sewer networks.

Keywords: *Digital Twin, pipeline, mineral suspension, workspace*

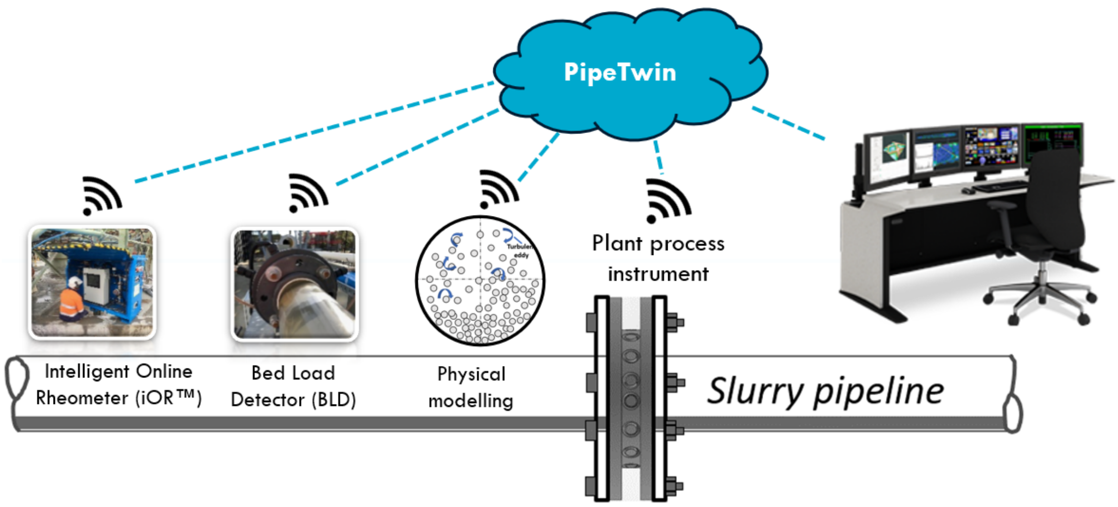

Figure 1. PipeTwin – a cloud-based platform for real-time monitoring of suspension pipelines

1. INTRODUCTION

Pipelines are used to convey suspensions in a wide range of industrial processes, including mining and mineral processing, where solid–liquid mixtures are transported using hydraulic principles. High-concentration suspensions, typically defined as solid-liquid mixtures containing about 55-65% solids by weight or higher, such as those encountered in thickened tailings transport - represent some of the most complex and operationally demanding cases. Accurate monitoring and control in such pipelines necessitate real-time insights into flow conditions, which are often limited by current instrumentation and modelling practices. Due to the multiple interactions present in a high concentration suspension, only through reconciliation of modelling and instrument data can accurate knowledge of the pipeline conditions exist, and appropriate control actions taken.

To address this challenge, the PipeTwin project combines multiple simultaneous measurements and feeds them through a Digital Twin for real time monitoring and control. The algorithm will turn the "underdetermined problem" that always exists with single measurements of complex flows into a determination of the most probable physical conditions that exist in the pipeline. The digital twin is built on the CSIRO Workspace platform. In this project, real-time data from proprietary instruments such as the CSIRO developed intelligent Online Rheometer (iOR) and Bed Load Detector (BLD) are used, as well as operational site data and numerical results from the CSIRO pipeline transport model. The iOR (Figure 2a) is an automated capillary rheometer that takes a side stream of the process fluid of interest and provides the real-time rheology information of the suspension (Chryss et al., 2019). The BLD (Figure 2b) is a thermal sensor based, pipeline solids level detection device. It is intended for suspension and tailings applications to detect the onset of solids settling and monitoring bed height (Chryss and Zheng, 2024). The integration of iOR and BLD measurements into pipeline transport models is expected to provide a scalable framework for digital twin applications across a wide range of suspension pipeline systems - including, but not limited to, tailings transport - while ensuring compatibility with different instrument types.

Figure 2. The CSIRO developed (a) intelligent Online Rheometer (b) Bed Load Detector integrated into a pipeline

2. RELATED WORK

Current practices in mineral suspension pipeline design and monitoring often rely on simplified modelling assumptions and a limited set of standalone instruments. For industrial pipelines, such as those used in thickened tailings transport, systems are often designed to operate in a turbulent regime, with average velocities exceeding the transitional threshold. Based on estimated transitional and deposition velocities, a minimum allowable transport velocity (MATV) is defined to ensure pipeline operation within a safe margin (Javadi et al., 2020).

MATV is generally estimated using semi-empirical models validated against small- and large-scale pipe loop tests and operational data - for example, the methods proposed by Slatter and Wasp (Slatter and Wasp, 2000), and later refined by Wilson and Thomas (Wilson and Thomas, 2006). However, experience with suspension systems, such as those used for iron ore tailings, shows that these correlations have limited applicability and there is no "one-size-fits-all" solution due to the high variability in suspension concentration and rheology (Ortiz et al., 2017). Rheological properties for a suspension sample are typically measured using bench-top rheometers. While there have been efforts to trial online rheology or yield stress measurement for real-time process control (Kealy et al., 2017), such approaches have not been widely adopted (Sofrà and Bhattacharjee, 2021). Most current instruments rely on side-stream sampling for offline laboratory analysis, which may not accurately reflect in-situ conditions due to time lags and sampling bias. As a result, decision-making is often based on outdated or incomplete data, limiting the ability to achieve true real-time process understanding.

Furthermore, existing pipeline models are rarely integrated with live data and are often calibrated in isolation, lacking the feedback mechanisms needed to capture the dynamic behaviour of dense, multiphase flows. There is little or no synchronisation between measurement tools and modelling frameworks, making it difficult to resolve the inherent uncertainties present in complex pipe flow systems. This fragmented approach underscores the need for an integrated, real-time solution that combines diverse data sources within a predictive framework - precisely what the proposed Digital Twin aims to deliver. Several studies have demonstrated the use of CSIRO's Workspace platform to implement modular, reproducible scientific workflows in digital twin systems (Watkins et al., 2021) and applications (Cummins et al., 2023). These examples highlight the capability of Workspace to integrate live data streams with advanced computational models for operational decision support. However, such approaches have not yet been widely applied to suspension pipeline monitoring, where the combination of live rheological measurements, multi-device integration, and predictive modelling remains relatively unexplored.

3. MAIN COMPONENTS OF THE PIPETWIN SYSTEM

Our PipeTwin architecture combines physical instruments, cloud data infrastructure, computational modelling tools and a web-based interface into a cohesive platform for slurry pipeline monitoring and simulation. This section describes each major component, outlining its role in data acquisition, processing and user interaction, as shown in Figure 3.

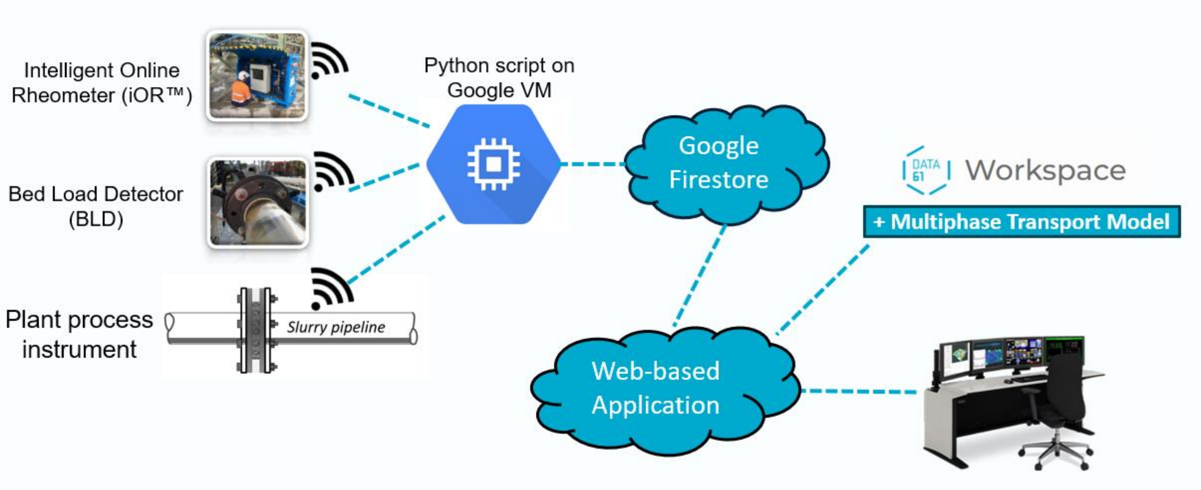

Figure 3. Main components of the PipeTwin system

3.1. Web-based application

Our system incorporates a comprehensive web-based application, combining both the frontend (React) and web server (Node.js) into a single cohesive subcomponent. The React interface enables users to input parameters, controlling and retrieving data from instruments like iOR, BLD and other pipeline instrumentation. The Node.js server underpins this interface by validating inputs, and orchestrating data exchanges with the underlying data processing infrastructure. Together, they form the primary access point for users to interact with the PipeTwin system.

3.2. Data processing & backend infrastructure

The backend for data processing integrates both the physical and computational aspects of the system:

- Mineral instruments installed along slurry pipelines continuously generate live operational data. Python scripts collect and push this data into database, ensuring centralized, time-stamped storage.
- A dedicated Workspace environment (CSIRO's scientific workflow platform) performs computational tasks such as advanced rheology calculations and bed height. Workspace communicates with the Node.js

server using ZeroMQ (ZMQ), a high-performance asynchronous messaging library, enabling robust, asynchronous execution of data-intensive jobs.

3.3. Database

Our system uses Google Firestore as a central, cloud-based database to store live data from mineral instruments. Its real-time data syntonisation and automatic indexing support responsive dashboards and efficient retrieval for simulations. This ensures reliable storage, easy scalability and a strong foundation for future analytics.

The system can be considered a digital twin because it maintains a dynamic, bidirectional connection between the physical process and its virtual representation. Instrument data such as rheology measurements are transmitted in real time to the digital environment through the Node.js server and ZMQ messaging layer. These data continuously update the computational models executed in Workspace, ensuring that the virtual system reflects the current state of the physical process. In turn, model predictions such as Reynolds number for a fluid, can be fed back to inform laboratory experiments and operational decisions. This ongoing synchronisation of data and feedback distinguishes the system from a stand-alone simulation and establishes it as a genuine digital twin.

4. IMPLEMENTATION

4.1. Google Firestore

We selected Google Firestore as the core database due to its cloud-native scalability, flexible document-based schema and real-time data synchronization capabilities. These features are well suited for handling diverse, time-stamped sensor data from multiple mineral instruments. Unlike traditional relational databases, Firestore allows us to easily adapt to varying data structures produced by different instruments, while automatically managing indexing and replication to support low-latency queries and future growth.

Python-based scripts run on a VM in the CSIRO Bowen research cloud, connecting to each device to fetch live measurements. The data is then stored in Firestore, organized into collections by device type and unique identifier. Each data point is stored as a document with a precise timestamp, enabling efficient retrieval of individual device histories over custom time ranges. This database design supports both real-time monitoring and retrospective analysis without additional restructuring.

The web application interfaces directly with Firestore to retrieve the latest device data. Firestore acts as a robust bridge between the physical monitoring infrastructure and the broader digital twin environment, ensuring consistent and reliable data availability. Its built-in support for real-time updates also positions the system to easily incorporate live dashboards, event-triggered processing or automated alerts in future extensions.

4.2. Workspace plugin

CSIRO Workspace is a workflow software that allow users to implement or reuse their code as Workspace plugins. The PipeTwin library based on a multiphase pipeline transport model developed by CSIRO was then ported to Workspace, along with its newly developed features both in C++ and Python, is now called the PipeTwin plugin. It provides the following capabilities:

- Operations that implement the data processing and calculations:

These operations are the processing units in Workspace. They are implemented in C++. The user can define their inputs and outputs. They can implement their own logic, or can call the PipeTwin library functions, thus reusing them from the legacy implementation. They can also run a Python script, and we use this capability to interface with our Firestore database. The flow of data traversing the operations is seamlessly and automatically taken care of by Workspace. Some examples of a PipeTwin operations are: FluidPressureDrop to calculate the Reynolds number for a fluid, given a fluid velocity and the length of the pipe; FirestoreTransactions that runs a Python script to add, update or delete data in the Firestore database.

- Workflows to orchestrate the processing logics:

The operations are put together in a workflow to execute some logics. The outputs of an operation can be fed into the inputs of one or many other operations. Not only PipeTwin plugin's operations are used, but workflows often include many operations from the rich library of Workspace plugins, e.g.: large data processing and graph plotting functions. Figure 4. illustrates an example of a PipeTwin workflow for receiving and processing data from connected devices (iOR/BLD/Pipeline instrumentation).

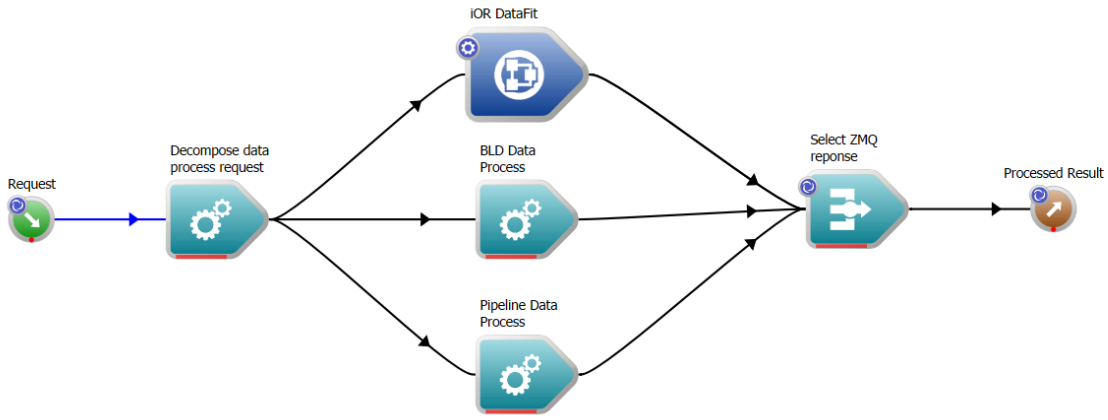

Figure 4. Pipetwin workflow for device data integration and processing

- Interface with web-application backend via ZMQ:

The PipeTwin plugin communicates with the web application using the ZeroMQ protocol. This is done by using the existing ZMQ operations in Workspace's Distributed plugin. For example, PipeTwin uses a ZMQ Response Loop operation in one of its data processing workflows to asynchronously receive data from the devices, then pass on the data to another workflow to process it. The output of this second workflow is the processing result that will be sent to the web-application.

4.3. Web application frontend + backend

We chose React for developing the frontend of our web application due to its component-based architecture, strong ecosystem, and efficient state management. React enables us to create dynamic, interactive pages that respond instantly to user inputs without requiring full page reloads, improving usability for users who frequently adjust simulation parameters. It has large ecosystem of libraries and built-in support for integrating with REST APIs, which also makes it easier to maintain and extend the interface as new device types or computational features are added. When a user modifies parameters or initiates a task, the frontend communicates these requests to Node.js server via the HTTP protocol. It also monitors data update from database via WebSocket.

The backend web server is built with Node.js, chosen for its non-blocking event-driven model, which is well-suited for handling multiple simultaneous user requests and orchestrating asynchronous interactions with the database and Workspace.

The Node.js server functions as the core intermediary layer, handling both data orchestration and system logic. It first retrieves live data from Google Firestore based on the requested time range from React; this data may include historical records or user-defined parameter presets necessary for accurate computation. Once the Node.js server gathers the necessary input data, it transmits the information to the Workspace using ZMQ messaging protocol. The Workspace performs data processing for devices (iOR/BLD/PipeLine) with specific operations. Upon completion, the processed results are sent back to the Node.js server via ZMQ. The Node.js server then forwards the output to the frontend, where it is presented to the user for interpretation, monitoring or further decision-making. This architecture supports real-time interaction, reproducibility and efficient data flow across all system layers.

The application interface offers a structured menu for users to select from different instruments like iOR or BLD, configure device-specific parameters, and submit these to the backend server for processing. In addition, we designed a dedicated Pipeline page to integrate data streams from all connected devices, presenting a unified view of system states and enabling derived calculations relevant to suspension pipeline performance, such as Sliding bed in-line concentration or Fluid pressure gradient. Live charts are drawn based on user-selected curve types to visualize data. Users can then not only manage individual devices but also analyse holistic pipeline behaviour directly from the web interface, bridging the gap between local sensor readings and system-level insights.

The web application is deployed in a VM on the CSIRO Bowen research cloud and there are plans to migrate to Google VM.

4.4. Example: iOR integration

As a concrete example, consider the integration of an iOR instrument providing raw measurements which are processed within the PLC in a step referred to as "Calculate" and output in the "Results" states. For each running cycle, data points like timestamp and shear rate are communicated to FireStore.

A lightweight Python script runs on a Google VM connected to the iOR's VPN network, periodically polling the device and packaging the readings into structured JSON documents. Each document includes the parameters needed for both "Calculate" and "Results", along with a precise timestamp and unique device identifier. This data is immediately pushed to the appropriate collection in Google Firestore, ensuring it is available in near real-time for downstream processes (Figure 5).

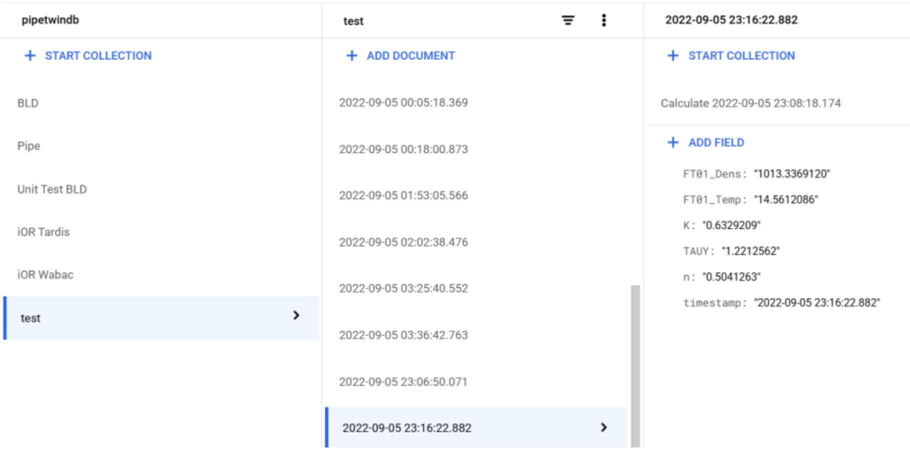

Figure 5. iOR data communicated to Google Firestore for management

The web application backend node.js aggregates the most recent iOR data from Firestore and sends it to the Workspace via ZMQ. Workspace then executes corresponding calculation by calling a processing module to perform rheology data validation and data fitting, returning results to Node.js via ZMQ.

The web application displays the rheogram showing both the current and previous rheology data fitting curves, alongside the corresponding measurement data for comparison (Figure 6). Key rheological parameters derived based on a non-Newtonian Herschel-Bulkley model - τ_y (yield stress), k (consistency), and n (flow index) - are shown, with their historical trends visualised on the right-hand side.

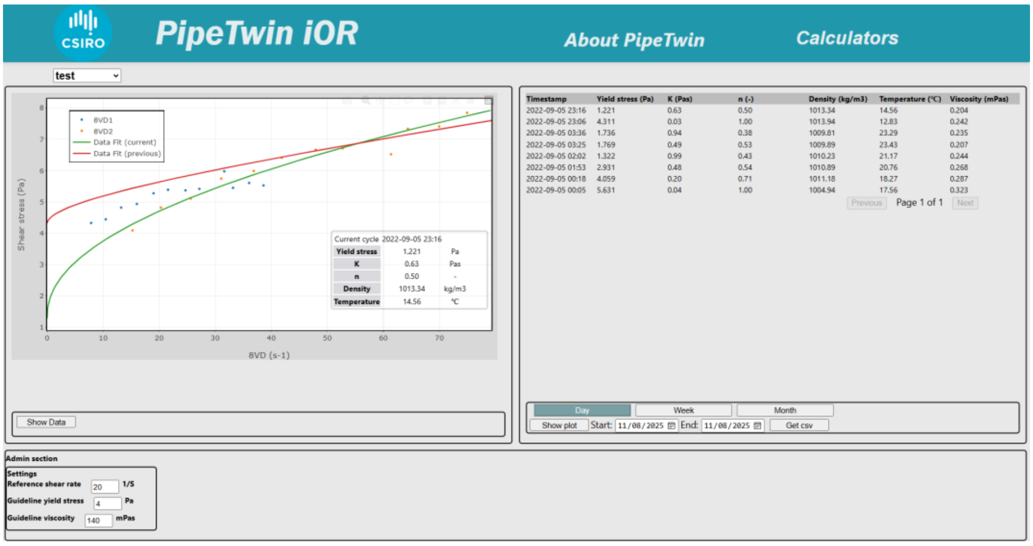

Figure 6. Web application UI for iOR devices

Both the graph and data table views can be toggled independently, it also allows users to switch between parameters such as τ_y, k, n, and density for trend analysis. Additionally, the table provides user-configurable reference and calibration inputs, including values like reference yield stress and shear rate, to support customised interpretation and calibration workflows. These rheological parameters are then utilized by the integrated transport model to perform real-time pipeline calculations. In this way, we overcome limitations associated with relying solely on bench-top rheometer measurements, which often fail to reflect actual process conditions. This serves as one example of how live data integration in PipeTwin enhances model accuracy - similar approaches are applied to other instruments integrated within the framework.

5. CONCLUSION

This paper presents the PipeTwin system, developed for real-time monitoring and control of suspension pipeline operations, enabling a shift from reactive to proactive management to prevent blockages and minimise downtime. By integrating live data from various mineral sensing devices with a web-based application, the system enables scientists and engineers to configure device parameters, visualize data, and coordinate operational decisions. The use of modern technologies - React for the frontend, Node.js for the backend, Google Firestore for scalable data storage, and Workspace for computational workflows - provides a modular and flexible architecture suitable for industrial applications. As future work, we plan to deploy the system to a Google Cloud Virtual Machine to enable secure web hosting and provide external access for customers. Additional development will focus on incorporating predictive analytics and expanding integration across more industrial environments.

REFERENCES

Chryss A, Monch A and Constanti-Carey K (2019) Online rheology monitoring of a thickener underflow. Paste 2019: Proceedings of the 22nd International Conference on Paste, Thickened and Filtered Tailings. Australian Centre for Geomechanics, 495-504.

Chryss A and Zheng E (2024) Monitoring bed formation in a pipeline: a comparative study of two measurement methods. Paste 2024: Proceedings of the 26th International Conference on Paste, Thickened and Filtered Tailings. Australian Centre for Geomechanics, 431-440.

Cummins SJ, Thomas DG, Sellers E, Fathi Salmi, E and Cleary, PW (2023) Building a simulation application for blasting: Case study of software component re-use and customisation in a workflow framework. 25th International Congress on Modelling and Simulation. Darwin, NT, Australia.

Javadi S, Pirouz B and Slatter P (2020) Paste and thickened tailings transportation design aspects overview. Paste 2020: 23rd International Conference on Paste, Thickened and Filtered Tailings. Gecamin Publications.

Kealy T, Bhattacharjee P and Griffin P (2017) Assessment of on-line rheology as a control tool for thickened tails. Paste 2017: Proceedings of the 20th International Seminar on Paste and Thickened Tailings. University of Science and Technology Beijing, 50-56.

Ortiz A, Rosewall A and Goosen P (2017) Pumping system design challenges for high-density iron ore tailings with highly variable slurry rheology. Paste 2017: Proceedings of the 20th International Seminar on Paste and Thickened Tailings. University of Science and Technology Beijing, 105-114.

Slatter PT and Wasp E (2000) Laminar/turbulent transition in large pipes. Zeszyty Naukowe Akademii Rolniczej we Wroclawiu. Konferencje (Poland).

Sofrà F and Bhattacharjee P (2021) Online yield stress measurement for real-time process control. Paste 2021: Proceedings of the 24th International Conference on Paste, Thickened and Filtered Tailings. Australian Centre for Geomechanics, 119-130.

Watkins D, Bolger M, Hetherington L, Zhang Z, Xie P, Harrison S and Cohen R (2021) Using workspace to implement digital twins in the mixed reality lab. MODSIM2021, 24th International Congress on Modelling and Simulation.

Wilson K and Thomas A (2006) Analytic Model of Laminar-Turbulent Transition for Bingham Plastics. The Canadian Journal of Chemical Engineering, 84, 520-526.