

Using Workspace to implement Digital Twins in the Mixed Reality Lab

Damien Watkins^a, Matt Bolger^a, Lachlan Hetherington^a, Zikai Zhang^a, Ping Xie^a, Simon M. Harrison^a and Raymond C. Z. Cohen^a

^a *CSIRO Data61, Clayton, Victoria, Australia*
Email: damien.watkins@data61.csiro.au

Abstract: The growing need for decision support, situational awareness and predictive analytics systems in real world iterative scenarios has given rise to the research field of “Digital Twins”; the combining of cyber-physical systems, artificial intelligence, computational processing platforms and mixed-reality visualisation. Digital Twins are living, digital replicas of a physical object, process or system and are comprised of three elements; a physical system; a digital model of that system; and the interactions that flow between these two systems.

In this paper we describe our approach to the ongoing task of building our Mixed Reality Lab (MRL) using our Scientific Workflow System (SWS) called Workspace and describe its characteristics that make it suitable for use in producing Digital Twins for both research and commercial applications. Digital Twins need to address many of the same issues that SWSs address, for example platform independence, handling large data sets, local and remote processing, algorithm integration, hardware integration, advanced visualisation, reproducibility and scalability. In general, we believe that the elements that make our SWS applicable to the creation of Digital Twins could also apply to other SWSs, subject to their capabilities in areas discussed in this paper.

Two prototypical use cases supported in our MRL, Defect Detection in Manufacturing scenarios and Human Movement (a.k.a. Digital Human) are outlined. Although these two use cases work at different scales, we see how both can be modelled in the MRL using the infrastructure constructed on top of Workspace.

Capability strengths of our SWS that are of particular importance to our Digital Twin use cases are discussed. These include: The use of both continuous execution (interactive) workflows with rich user interfaces as well as distributed, cross-platform, remote execution processing workflows; Advanced interactive visualisation capabilities ranging from 2D widgets and charts, OpenGL-based 3D rendering through to a real-time Virtual and Augmented Reality pipelines using third party rendering engines; a modular, extensible plugin based architecture that exposes capabilities from many permissive open source libraries combined with in-house developed capabilities; and a streamlined deployment and productisation process to support different distribution or commercialisation needs.

The MRL is very much a work in progress, optimisations to the current systems are ongoing and work to add extra facilities and infrastructure to support more scenarios is progressing.

Keywords: *Scientific Workflow Systems, Digital Twins, Mixed Reality*

1. INTRODUCTION

Underpinning the three aforementioned elements of Digital Twins is sophisticated information management and computational processing infrastructure that, in order for the twin to be truly ‘living’, permits timely information to flow in two directions: from the system to the model and vice versa. The Workspace team is involved in several research projects within CSIRO that use Digital Twins, with applications ranging from Industry 4.0 to Human Performance Engineering. Digital Twins can be relatively simple or extremely complex. For example, a global manufacturing supply chain would have producers in multiple countries, all using different software stacks and applying numerous software processes, whereas a human gait tracking system could use a laptop and a number of cameras. As part of the development and testing of these Digital Twins, we constructed the Mixed Reality Laboratory (MRL). The MRL allows us to validate the systems we create and to demonstrate the applicability of Digital Twins to different scenarios.

In this paper we describe our approach to the ongoing task of building our MRL using our Scientific Workflow System (SWS) called Workspace and describe its characteristics that make it suitable for use in producing Digital Twins for both research and commercial applications. In general, we believe that the elements that make our SWS applicable to the creation of Digital Twins could also apply to other SWSs, subject to their capabilities in areas discussed in this paper. We then describe the elements of our MRL and demonstrate some use cases within the MRL.

2. OUR APPROACH

Since 2005 our team has been developing a SWS called Workspace (Cleary *et al.* 2020), which we have chosen to employ as the information management and computational processing infrastructure underlying our Digital Twins. Workspace has been used in numerous scientific domains both inside and outside CSIRO. Examples include: ArcWeld (Thomas *et al.* 2019), Dive Mechanic (Cohen *et al.* 2017; 2020), Fractura (Kear *et al.* 2019), HelioSim (Potter *et al.* 2017), and Spark (Hilton *et al.* 2015)

While many SWSs aim to provide similar general functionality, all have some specific advantages and disadvantages in relation to their effectiveness to support Digital Twins. In the case of Workspace these advantages would include: 1) continuous execution mode with minimal execution paths and Remote Execution modes, 2) advanced interactive visualisation and 3) suitability for commercialisation.

2.1. Continuous execution mode with minimal execution paths

Most workflow engines execute a workflow serially from beginning to end, this is especially true of commercial workflows systems. Workspace workflows are comprised of a series of Operations (atomic tasks) with their outputs connected to the inputs of other Operations, forming a dependency graph. Workspace executes a workflow by inspecting the workflow’s required outputs, working out the minimal path to satisfy them and then executing that path. Workspace has a continuous execution mode where it remembers which Operations are up-to-date and only re-executes an operation when it is out-of-date and required to regenerate outputs, allowing intermediate results of complex sub-workflows to be re-used multiple times without re-execution. Operations are continuously monitored by Workspace and marked out-of-date when inputs are changed by user input, external stimuli (such as a notification from a sensor) or execution of upstream operations. This is particularly useful for Digital Twins, where multiple workflows often operate in concert, listening for input from sensors before re-executing.

2.2. Remote Execution modes

Workspace workflows are stored as XML, which permits workflows to be shared across boundaries, for example across different operating systems. For a workflow to execute on a machine, it requires Workspace to execute the workflow using plugins to provide all the required operations, datatypes, and widgets. Workspace can execute workflows in batch mode (e.g. run at a command prompt.) Workspace provides support for industry standard scheduling modes such as PBS and SLURM. Workspace provides its own Grid scheduling system, allowing an ad hoc network of machines to cooperate and coordinate remote execution of workflows. This Grid scheduling systems is the backbone of processing within the MRL.

2.3. Advanced interactive visualisation

Workspace allows inputs and outputs of operations anywhere along a workflow to be dynamically visualised, while the workflow is executing. As a result, intermediate inputs and outputs can be visualised and understood just as easily as the final results. This is achieved by connecting a widget to an input or output. This is a two-

way connection, allowing a widget to display updates to the current value and allowing the user to modify the value. In continuous execution mode, modifying the value via a widget then triggers re-execution of relevant portions of the workflow. Workspace widgets range from simple UI elements like integer spin boxes or string line edit elements through to complex interactive 2D charts (such as those based on NV3D or Plotly) or complex interactive 3D scenes using OpenGL. A particular focus of Workspace is interactive 3D visualisation and this capability has been further extended to technologies such as OpenVR and support for third-party visualisation platforms such as Unity (<https://unity.com>) and Unreal Engine (<https://www.unrealengine.com>). Visualisation using these third-party platforms is achieved via inter-process communication to an independent client developed in the engine. In the case of this inter-process visualisation, an update in a workflow can quickly manifest as an updated scene in the visualisation engine which may be running on a separate device such as a table or wireless augmented reality headset.

2.4. Suitability for commercialisation

One immediate requirement for suitability for commercialisation is licensing. Licensing must not tie developers into specific licensing models where developers are forced into requirements such as paying for expensive mandatory components or being forced to license source code under restrictive models (e.g., viral licenses): developers should be free to do either (or none). Workspace's license permits plugin developers to incorporate dependencies as necessary and according to their commercialisation requirements (Cleary 2017).

2.5. Other relevant advantages of Workspace

Workspace provides a plugin architecture that allows the integration of software libraries and components. These libraries and components can be created anew, or they can be pre-existing. For example, Workspace has several plugins that have been largely created in-house, such as Mesh Processing, Geospatial, and Particle-based Modelling, which provides higher level operations that utilise multiple underlying libraries. Other plugins are essentially wrappers for popular third-party libraries such as OpenCV, VTK and PCL. Workspace plugins can add Operations, Datatypes and Widgets – making functionality added by users virtually indistinguishable from Workspace's built-in capabilities.

To simplify end-user interaction, Workspace provides for the streamlined creation of standalone applications. Such applications can be considered as consisting of 1) Workspace, 2) one or more plugins, 3) one or more workflows and 4) a Graphical User Interface (GUI), Command-Line Interface (CLI) or web-based application or API. Examples of such applications include the aforementioned ArcWeld, Dive Mechanic and HelioSim. These applications permit the experience of the end-user to be tailored to their specific requirements. In the case of Digital Twins, this can be critical for feedback loops involving human operators that enable situational awareness.

3. MIXED REALITY LAB

In order to test and demonstrate our Digital Twins we created the MRL at CSIRO's Clayton site. The MRL is equipped with both physical devices and software systems. The physical devices include:

- Input Devices – under orchestration by Workspace, these devices feed the system with critical sensory data. Many such devices are cameras, including Nikon DSLRs, Ximea Industrial Cameras, Odos Swift Time of Flight (ToF) cameras and HD webcams.
- Device PCs –input devices are accessed via different endpoints. Some use Optical Fibre connections (e.g., Ximea XIB) or USB connections (e.g., Nikon), and for these a PC (or similar) is required proximate to the device.
- IoT and Edge compute nodes – Workspace workflows can be deployed and executed on Linux-based edge controller and compute devices such as Raspberry Pi and NVIDIA Jetson nodes. These are used in the MRL for functions such as projector control. Robot control integration and edge image processing are also under development using these nodes.
- GPU server – many data processing algorithms are GPU accelerated and so we have a high-performance GPU server cluster in order to maximise throughput for near real-time processing.
- Hybrid devices – some devices are input and output devices, such as our HoloLens units – the HoloLens is our main output device used to create the mixed reality environment in the MRL.

The software includes:

- Workspace and a suite of plugins for image and geometry processing, visualisation, data management and provenance. The core Workspace plugins are built on top of the Qt Toolkit for cross-platform capabilities and other utilities, and a number of embedded scripting languages and APIs such as OpenGL, Python, HDF5, ZeroMQ to name but a few.
- Dedicated Workspace extension plugins are also used extensively in the MRL for image and computational geometry processing based on popular C++ libraries including OpenCV, PCL and VTK.
- Workspace's Job System database and web dashboard (one of Workspace's several job scheduling systems) which is used to schedule, monitor and retrospectively analyse distributed jobs across the nodes in lab.
- The Unity 3D engine, including extensions such as Microsoft's Mixed Reality Toolkit, is used as part of Workspace's inter-process visualisation capability for live visualisation of workflow data in AR and VR devices as part of a world locked virtual augmentation of the real lab environment.
- Software Development Kits used to access specific devices, e.g., the Ximea industrial camera SDKs
- Generic Software Standards and Protocols, e.g., many industrial cameras support the GenICam interface. When different devices support such a standard then it is much easier to mix and match devices from different vendors.
- DevOps tools and systems for continuous integration and testing, deployment, automated system testing and monitoring.

3.1. Manufacturing Scenarios

There are a number of stages common to all of our manufacturing scenarios:

- **Manufacturing:** This stage is somewhat optional, either we can produce a component via our MakerBot plastic 3D printer in the lab, utilise other manufacturing capabilities at CSIRO Clayton (such as Lab22 metallic 3D printing facility) or an already constructed object can be brought into the MRL.
- **Data capture:** At this stage applications or workflows capture sensor data from all devices (e.g., images) and either 1) display these in a user interface and/or 2) write these to storage. Displaying can be thought of as a 'smoke test' of the system (e.g., can we communicate with all devices and receive data from them, are exposure and focus settings correct, etc.). Captured sensor data needs to be written to storage for distributed processing.
- **Sensor calibration:** As sensors can be repositioned in the MRL depending on the scenario under test, it is necessary to calibrate the particular setup (e.g., calculate positions for all the devices). Calibration involves the acquisition of several poses of a calibration target in different positions and post processing. Our current process involves the use of a ChArUco board, approximately 1 m² in area, photographed in numerous positions around the capture space. The intrinsic and extrinsic properties of the cameras are calculated using the operations in the OpenCV library in a manner like that described in An *et al.* (2018). Validation of the calibration data is achieved via two systems, a 3D visualisation of the cameras' locations including quality metrics of the solution (either on a desktop or via a HoloLens) and a more detailed calibration assessment application that analyses detailed aspects of the calibration quality providing charts and other reports. Alternative camera calibration solutions are also under investigation.
- **3D reconstruction** – these processes use different techniques to create point cloud reconstructions of the object in the field of view. Currently we support Structured Light Systems and CSIRO's patented Stereo Depth Fusion (Stainlay & Poropat 2005). The reconstruction process consists of synchronising the projection of patterns into the scene, capturing images from the cameras and storing them on a file server, scheduling and executing numerous background process/workflows and storing the resultant point clouds back on the file server including all associated metadata.
- **Surface mesh processing** – Multiple stages of post processing and merging are performed on the raw point clouds produced by the Photogrammetry Reconstruction stage depending on the target use case. These can include various types of point cloud filtering, segmentation (to remove background and foreground elements), alignment and stitching of individual patches of the target (seen by separate groups of sensors), down sampling, triangulate mesh generation and further mesh refinement. Workspace is ideally setup to support flexible experimentation of these (potentially complex) pipelines allowing different methods and parameters to be used and combined in a unified environment that can fully leverage the distributed compute and automatic provenance features available. Current workflows make extensive use of algorithms available in PCL and VTK for much of this processing.
- **Surface Comparison** – Similar to the Surface Mesh Processing stage, workflows and visualisations can be created to compare the resulting 3D reconstructed surface with reference data such as CAD models of the target part or past baseline results.
- **Backend tools and processes:** All these stages make use of shared tools and process such as:

- Sensor Controllers – applications that parse the network configuration of the MRL from a central database then dispatch a number workflows to the necessary machines in order to control the physical devices. This process uses IP addresses and dispatches platform-independent workflow descriptions (XML), thus the system is largely oblivious to aspects such as the underlying OS of any of the Device PCs.
- Distributed Workflow Scheduling and Monitoring
- Software Deployment – a system to deploy Workspace, plugins and workflow files on to all the machines in the MRL

3.2. Human Performance Engineering Scenarios

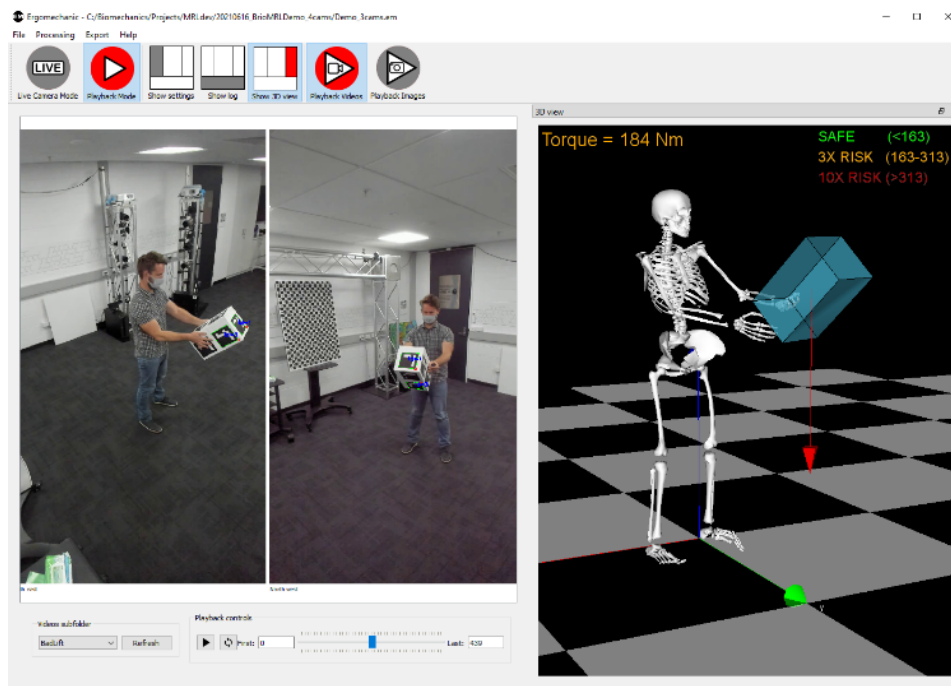


Figure 1. Digital Human in the Mixed Reality Lab, as rendered via the Digital Human Workspace Application

Many of the same types of processes can be used in similar ways to measure human movement and provide feedback on the outcomes of movement with the purpose of reducing injury risk and improving the performance of the body. Multiple Logitech Brio webcams are placed around the capture area in a manner that enables two or more cameras to see any one item of interest. The following processes are performed:

- System calibration: this involves calibrating the camera network of webcams using ChArUco boards as described above.
- Data capture: the system records a video stream from each of the webcams, buffering to RAM if practical and then writing to disc for postprocessing
- Image processing: a deep learning (DL) algorithm such as OpenPose (Cao et al. 2019) or Detectron2 (Wu et al. 2019) is used to identify body joints and other points of interest in the video streams.
- 3D reconstruction: points of interest from the 2D images are triangulated into 3D space using the calibration parameters from step (1). Included in this step is an assessment of the contribution of the 2D dataset from each camera. Because the DL algorithms can occasionally erroneously or inaccurately detect human features, it is best to test the variation in 3D projection with and without the data from each camera. Cameras which have 3D consistency between their 2D feature detections can be automatically given a higher trust weighting than cameras which are inconsistent with the rest. For this reason, it is best to use as many cameras as is practical.
- Human body scaling: a generic musculoskeletal model is scaled to the current participant using the average measured distance between key joints (such as the wrist to elbow and elbow to shoulder). This scaling includes the physical dimensions of the musculoskeletal model and the inertial properties (which includes the mass and moments of inertia) of each body segment (hand, forearm, upper arm, etc).

- Analytics and evaluation: depending on the application, the output metrics of the system may be related to the amount of movement or the movement speed (termed kinematics) or the loading on particular body parts (termed kinetics). First, an inverse kinematics algorithm, originally developed for robotics, is used to find the joint angles that result in the calculated joint movements. Second, a dynamic simulation is performed to calculate the load on each joint from gravity, inertia (i.e., from quick movements), and external objects including the ground. Third, the calculated kinematic and kinetic data are compared to known allowable bounds to enable decision making.
- Visualisation and decision making: the results of the 3D reconstruction, analytics, and evaluation can be presented to the user and/or participant. This can be as complex as an AR visualisation of the loading inside the participant's body superimposed on the view of the participant, and as simple as a traffic light indicator showing the participant that they are loading their body excessively. The system outputs can be used to make a decision, whether that be to change lifting technique for manual handling scenarios, or whether or not to change exercises for a rehabilitation application. An example visualisation is shown on the right.

3.3. Other uses cases in the MRL

These are by no means the only two usage scenarios for the development and demonstration of Digital Twins in the MRL. They are, however, representative of the breadth and scale of possible applications; from the manufacturing of large hard components through to the smaller soft scale of humans.

3.4. The Mixed Reality Lab in Action

The image below shows Elle in the MRL interacting with a virtual model of a space station. Above the screen you can see the un-augmented view while in the screen you can see the view through a HoloLens showing other AR elements such as camera calibration and field of view markers.




Figure 2. The Mixed Reality Lab in Action

4. FUTURE WORK

The MRL and the Digital Twin use cases it supports are very much a work in progress. One important aspect of the MRL is the ability to mix-and-match variations in technologies, for example, as mentioned before alternative COTS camera calibration solutions are currently under investigation. On the technology side, facilities such as Robotics are in the process of being added to the MRL. Before we can address higher level functionality such as situational awareness and predictive analytics, it has been necessary to create a core substrate of technology in the MRL.

5. DISCUSSION AND CONCLUSION

The MRL has allowed us to undertake the development of a unified platform for Digital Twins of systems, ranging from manufacturing to human movement. These systems require the integration of numerous elements, from hardware to software, with different technical limitations and licensing restrictions. The use of our SWS, Workspace, as the backbone of the system, has enabled us to develop demonstrator Digital Twins reusing and leveraging many of our existing capabilities and adding new capabilities as required. Features of a SWS like Workspace include providing a modular and extensible framework for software reuse; an ability to run in a distributed nature across different platforms and interface with many device types; flexible and customizing user-interfaces and visualizations for complex processes; unified job processing and provenance recording, have proven to align well with the requirements needed to develop Digital Twins in the example domains discussed. In summary, we believe that we have demonstrated that the elements that make our SWS applicable to the creation of Digital Twins could also apply to other SWSs.

REFERENCES

- An, G.H., Lee, S., Seo, M.-W., Yun, K., Cheong, W.-S., Kang, S.-J., 2018. Charuco Board-Based Omnidirectional Camera Calibration Method. *Electronics* 7, 421.
- Cao, Z., Hidalgo, G., Simon, T., Wei, S. E., and Sheikh, Y., 2019, OpenPose: Realtime multi-person 2D pose estimation using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1), 172–186.
- Cleary, Paul W. Thomas, David. Hetheron, Lachlan, Bolger, Matt, Hilton, James E., Watkins, Damien, 2020 *Workspace: A workflow platform for supporting development and deployment of modelling and simulation*, *Mathematics and Computers in Simulation*, Volume 175, 2020, Pages 25-61, ISSN 0378-4754
- Cleary, Paul W. Watkins, Damien, Hetheron, Lachlan, Bolger, Matt, Thomas, David, 2017 Opportunities for workflow tools to improve translation of research into impact presented at 22nd International Congress on Modelling and Simulation (MODSIM) 2017
- Thomas, David G., Murphy, Anthony B., Chen, Fiona F., Xiang, Junting, Feng, Yuqing 2019 ArcWeld: A case study of the extensibility of software applications built using Workspace architecture presented at 23rd International Congress on Modelling and Simulation (MODSIM) 2019
- Cohen, Raymond C. Z., Harrison, Simon M., Cleary, Paul W., Bolger, Matt, 2017 Dive Mechanic: Bringing 3D virtual experimentation to elite level diving using the Workspace workflow engine presented at 22nd International Congress on Modelling and Simulation (MODSIM) 2017
- Cohen, RaymondC.Z., Harrison, Simon M., Cleary, Paul.W., 2020, Dive Mechanic: Bringing 3D virtual experimentation using biomechanical modelling to elite level diving with the Workspace workflow engine. *Mathematics and Computers in Simulation* 175, 202–217.
- James Kear, Dane Kasperczyk, Lachlan Hetheron, Dang Quan Nguyen, Chris Rucinski, Abbas Movassagh, Saeed Salimzadeh, Xi Zhang, 2019, Fractura – Common modelling platform for multiple hydraulic fracture simulators presented at 23rd International Congress on Modelling and Simulation (MODSIM) 2019
- Hilton, James; Miller, Claire; Bolger, Matt; Hetheron, Lachlan; Prakash, Mahesh, 2015 An integrated workflow architecture for natural hazards, analytics and decision support. In: ISESS; Melbourne. Springer; 2014. 1.
- Stainlay, Maxwell L., Poropat, George V., (2005) A method for automated 3d imaging. WO2006024091A1 (<https://patents.google.com/patent/WO2006024091A8/3Den>)