Refactoring the iRiverMetrics algorithm

Thiaggo de Castro Tayer

The University of Western Australia

Cataloguing data

This publication (and any material sourced from it) should be attributed as: Tayer, T.C. 2025, Refactoring the iRiverMetrics algorithm. The University of Western Australia, Perth, WA., Report to the Resilient Landscapes Hub of the Australian Government's National Environmental Science Program. Jun 2025. CC BY 4.0.

This publication is available at https://nesplandscapes.edu.au/resources/

The Resilient Landscapes Hub is funded by the Australian Government under the National Environmental Science Program. The hub is hosted by The University of Western Australia.

Acknowledgements

The authors thank the Queensland Department of the Environment, Tourism, Science and Innovation (DETSI) Team for discussions on the direction of the project.

Acknowledgement of Country

We acknowledge the Traditional Owners of Country throughout Australia and their continuing connection to and stewardship of land, sea and community. We pay our respects to them and their cultures and to their Ancestors, Elders and future leaders. Our Indigenous research partnerships are a valued and respected component of National Environmental Science Program research.

Contents

Figures

Tables

1. Introduction

Accurately describing the hydrology of intermittent rivers is a critical step in improving our understanding and management of these systems, as it enables the assessment of their natural states, the evaluation of anthropogenic impacts, and the monitoring of ecological recovery. However, this characterisation is often not straightforward. Traditional methods, such as gauged discharge data, offer little information once flow ceases and fails to provide insights into the location, morphology, or resilience of intermittent river pools. Additionally, these approaches are often costly due to the need for ongoing field maintenance. As a result, many large rivers in remote and arid regions suffer from sparse or non-existent hydrological gauging data (Callow and Boggs, 2013).

Remote sensing technologies provide a cost-effective and rapid means of surveying extensive geographical areas (Feyisa et al., 2014; Huang et al., 2018; Johary et al., 2023). For instance, satellite imagery can capture entire river basins in a single pass, facilitating comprehensive water mapping. Remote sensing allows for frequent updates to water maps, such as annual or biennial assessments, which are crucial for regions experiencing significant interannual variability in hydrological conditions or undergoing rapid environmental changes, such as drying trends (Tayer et al., 2023a). High-resolution satellite data provide detailed insights into water distribution, while historical imagery enables the analysis of temporal changes, allowing researchers to track changes in water bodies over time. Multispectral images can be analysed to generate attributes for surface water (e.g., pool area, length, perimeter) along the length of a river or for specific sections within a river, which can then be used to derive metrics (e.g., pool complexity, pool size) that provide insight into ecological functioning.

The iRivermetrics algorithm, an open-source Python package developed under NESP Northern Hub Project 1.3.3, leverages this approach to study surface water dynamics in intermittent rivers. Unlike pixel-based methods, it analyses river sections to generate metrics on persistence and morphology, with a focus on persistent pools. Applied successfully to the Fitzroy River in Western Australia, iRivermetrics effectively captures seasonal patterns, showing how surface water features—such as size, complexity, elongation, and fragmentation—change as conditions shift from wet to dry (Tayer et al., 2023b, 2023a). However, its current version lacks performance profiling, making it difficult to identify inefficiencies in execution time or memory use. Additionally, it was not designed for compatibility with Dask protocols, which enable parallel computation for large datasets, nor fully integrated with frameworks like the Open Data Cube (ODC) or SpatioTemporal Asset Catalog (STAC) to enhance geospatial processing. These enhancements would significantly enhance computational efficiency and scalability, enabling broader and faster analyses.

Optimising iRivermetrics for robust analysis of large-scale surface water datasets is particularly relevant for rivers like the Gilbert River in Queensland's Carpentaria Gulf region, where pronounced wet and dry seasons affect water availability and ecosystem health (Gardner et al., 2021). Understanding these dynamics is essential for effective water resource management, particularly considering climate change and increasing anthropogenic pressures on water systems (Ellis et al., 2024; Marshall et al., 2024). The application of automated algorithms, such as iRivermetrics, can facilitate the scalability and performance of surface water analysis, ultimately contributing to the ecological integrity of the rivers (Garrido-Rubio et al., 2020).

This report addresses task 15, part of phase 3 (milestone 6) for NESP project 2.1, *Assessing risks to the environment from water-resource development in northern Australia, using north Queensland as a case study*. This task aims to enhance the iRivermetrics algorithm by incorporating code profiling to identify and resolve performance bottlenecks, integrating Dask for parallel processing, and ensuring compatibility with STAC protocols. We will test the refactored algorithm on a section of the Gilbert River to streamline data generation and reduce resource demands. These improvements will strengthen iRivermetrics as a tool for analysing large-scale surface water datasets via remote sensing, supporting the sustainable management of persistent pools and informing water resource planning in the Gulf of Carpentaria region.

2. Aim

This report aims to improve the computational efficiency and scalability of the iRivermetrics algorithm, enabling robust analysis of large-scale surface water datasets. Originally developed under the NESP Northern Hub Project 1.3.3, iRivermetrics uses remote sensing to generate metrics on surface water patterns, focusing on river sections rather than individual pixels. The algorithm offers insights into the persistence and morphology of surface water, with a particular emphasis on the dynamics of long-term persistent pools. By enhancing its performance, we seek to streamline data generation for intermittent rivers.

3. Approach

3.1 Study area

The Gilbert River, located in the Carpentaria Gulf region of north-eastern Australia, is an intermittent river part of a largely undisturbed catchment area that includes the Mitchell and Flinders rivers, which are essential for maintaining the ecological integrity of the region (Broadley et al., 2020). This river experiences a distinct wet-dry tropical climate, characterised by pronounced seasonal variations. The wet season typically spans from November to April, while the dry season extends from May to October, marked by significantly reduced precipitation and higher evaporation rates (Keller et al., 2019; Warfe et al., 2011). During the wet season, river flows increase dramatically, leading to enhanced ecological productivity and habitat availability for aquatic species (Godfrey et al., 2022; Warfe et al., 2011). Conversely, the dry season is characterised by lower water levels and higher water temperatures, which can affect fish assemblages and habitat dynamics (King et al., 2015; Pusey et al., 2016). The ecological significance of the Gilbert River is underscored by its role in supporting water-dependent ecosystems, which are increasingly threatened by climate variability and human-induced changes (Mcgregor et al., 2018).

This study was conducted in a limited reach of the Gilbert River, provided by the Queensland Department of the Environment, Tourism, Science and Innovation (DETSI), which encompasses a 100 km stretch of the main river channel, buffered by 1 km on each side (Figure 1). For testing algorithm performance, the study area was further divided into 50 sections, each approximately 2 km in length and similar in area.

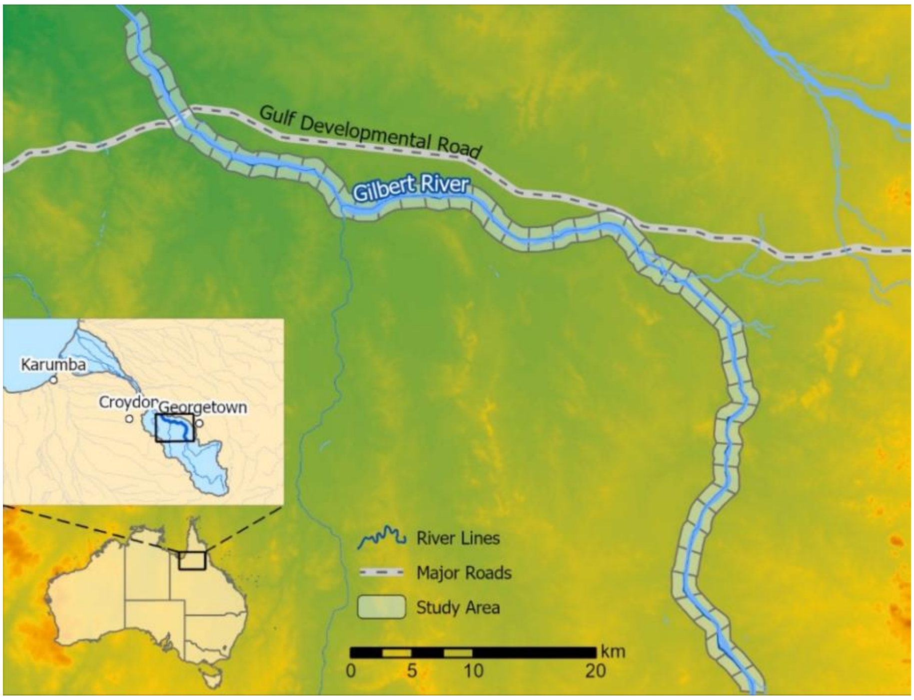

Figure 1. Location of the Gilbert River study area in north-eastern Australia. The map highlights a 100 km stretch of the main river channel, with a 1 km buffer on each side, and the tested sections, as provided by the

Queensland Department of Environment and Science (DES). An inset map shows the broader regional context of the study area within the Gilbert River catchment and Australia.

1.1 The algorithm

iRivermetrics is an open-source algorithm developed to analyse the surface water dynamics of intermittent rivers using multispectral satellite imagery. Its modular design comprises two primary components, allowing users to employ only the modules relevant to their specific projects—one for efficient water detection and the other dedicated to computing ecohydrological metrics, which is the focus of this report.

3.1.1 Inputs

To calculate surface water metrics, the algorithm requires several key inputs. The most crucial of these is da_wmask, which comprises a binary water mask time series derived from satellite imagery, indicating the presence or absence of water. These masks can be provided either as a directory path to image files (such as .tif files) or as an xarray.DataArray. All water masks must share consistent spatial resolutions and coordinate reference systems (CRS), and their filenames must include dates in the format "yyyy-mm-dd" or "yyyy_mm_dd".

Another important input is rcor_extent, which is an optional shapefile or a geodataframe that defines the Area of Interest (AOI). This typically represents specific river sections for which metrics will be calculated. If rcor_extent is not provided (i.e., it defaults to None), the boundaries of the input water masks will be used as the AOI, and in this scenario, the outdir (output directory) parameter must be explicitly specified to ensure results are saved.

The process of defining these river sections for rcor_extent is critical for meaningful ecohydrological analysis, and the specific methodology depends on the study's aims. To ensure that each river section maintains both consistent length and approximately the same area (many metrics are area-dependent), a common approach involves generating points along the vectorised river's main channel drainage line (e.g., using a GIS tool). These points are then used to split the drainage line into a predefined number of equal-length segments. The actual river sections are subsequently defined as a buffer (e.g., 1 km wide on each side) around these drainage line segments. Finally, a set of topology rules (for instance, ensuring polygons do not overlap) is applied to check for inconsistencies and guarantee that the generated sections are spatially distinct and coherent. Note that while this method is one way to define sections, the entire catchment could also be used as a single, unique section to calculate metrics at a broader, catchment level, or any other set of rules could be applied depending on the specific objectives of the analysis.

Other optional inputs include outdir, which specifies the destination directory for results (defaulting to a directory adjacent to the rcor_extent file if rcor_extent is provided and outdir is not specified); section_length, an optional parameter defining the length (in kilometres) for river sections when calculating metrics (if None, the algorithm relies on the predefined sections in rcor_extent); section_name_col, an optional parameter to specify the column in rcor_extent that contains the names of the river sections; min_pool_size, which sets the minimum detectable water pool size in pixels (defaulting to 2 pixels); img_ext, the file extension of the water mask images (defaults to '.tif'); export_shp, a boolean indicating whether to export detailed shapefiles of the analysed river sections (defaults to False); export_PP, a boolean to control whether a pixel persistence raster is exported (defaults to False); and fill_nodata, a boolean to determine if nodata values in the input data should be filled (defaults to True).

3.1.2 Processing

Once inputs are provided, the iRivermetrics algorithm initiates its processing workflow. The first step involves validating the data to ensure compatibility with the algorithm's requirements. Subsequently, data is clipped to the defined Area of Interest (AOI), any existing nodata values are filled, and all spatial data is reprojected to Universal Transverse Mercator (UTM) for consistency.

Following these preparatory steps, the algorithm calculates a comprehensive array of metrics for each designated river section. These include core measures such as total wetted area, number of pools, pool area, perimeter, and wetted length, as well as more complex indicators, including the Area-Weighted Mean Shape Index (AWMSI), Area-Weighted Elongation Ratio (AWRe), Area-Weighted Mean Pixel Area (AWMPA), Area-Weighted Mean Pool Length (AWMPL), and Area-Weighted Mean Pool Width (AWMPW). It also computes persistence measures (e.g., pixel persistence, refuge area) and fragmentation indicators (e.g., pool fragmentation, longitudinal fragmentation). In total, the iRiverMetrics tool generates 16 ecohydrological metrics for each time step of the water-mask time series (e.g. WOfS), as described in (Tayer et al., 2023b). These metrics are broadly categorised into surface water morphology and hydrological resilience (Table 1). Collectively, they capture four essential aspects of riverine systems: (1) the quantity of surface water in the channel, (2) its morphology, (3) its persistence, and (4) its fragmentation.

Table 1. Ecohydrological metrics calculated in this study, organised into two main categories: morphology (left) and resilience (right).

Morphology	Resilience
• Section area • Total wetted area • Total wetted perimeter • Wetted length • Area-weighted Mean Shape Index (AWMSI) • Area-weighted Elongation Ratio (AWRe) • Area-weighted Mean Pixel Area (AWMPA) • Area-weighted Mean Pool Length (AWMPL) • Area-weighted Mean Pool Width (AWMPW)	• Number of pools • Wetted Area Percentage of Section (APSEC) • Wetted Length Percentage of Section (LPSEC) • Pool fragmentation (PF) • Pool longitudinal fragmentation (PLF) • Pixel persistence (PP) • Refuge area (RA)

These metrics are derived by identifying connected water pixels, skeletonising river centrelines, and measuring distances within water bodies. Key hydrological attributes, such as pool area, length, perimeter, width, and count, are extracted by labelling connected water pixels and measuring distances within groups of connected pixels, which are then combined into the specified metrics.

Specifically, pools are defined as clusters of connected pixels using an eight-neighbourhood connectivity approach, meaning adjacent pixels in all eight directions are considered. Metrics are computed either for the entire river reach or for specific sections during testing. For instance, the total wetted area is determined by counting water pixels within the reach at each timestep. Area-Weighted Mean Pool Area (AWMPA) involves averaging the areas of all identified pools, weighted by their respective sizes. Similarly, Area-Weighted Mean Pool Length (AWMPL) and Area-Weighted Mean Pool Width (AWMPW) are derived by measuring the lengths and widths of each pool and calculating their area-weighted means, respectively. Pool Fragmentation (PF) is calculated as the ratio between the total number of pools and the overall wetted area within a section, with higher PF values indicating increased isolation of pools. Detailed equations and further explanations of each metric's behaviour are available in Tayer et al. (2023b, 2023a).

3.1.3 Outputs

The iRivermetrics algorithm generates a comprehensive set of outputs designed for detailed ecohydrological analysis. The primary output is a CSV table containing a time series of all calculated metrics for each input river section. Additionally, a pixel persistence raster is generated, which visually illustrates the temporal persistence of water across the AOI, aiding in the identification of refuge areas. Users also have the option to receive spatial layer outputs, such as detailed shapefiles for visualisation and further geospatial analysis. These can include shapefiles depicting their extent (polygons), length (polylines), and key point locations (start, end, and midpoints of pools). These outputs collectively provide a robust foundation for analysing water dynamics at various scales, from individual pixels to entire river sections and distinct pool units, enabling flexible spatiotemporal assessments. For more detailed documentation, users are encouraged to refer to the iRivermetrics GitHub page.

3.2 Spatial Data

The spatial analysis of the river's wetted area and pool characteristics was conducted using water masks derived from the Water Observations from Space (WoFS) dataset (Mueller et al., 2016). These water masks were accessed via the Geoscience Australia servers using the pystac_client and odc_stac libraries and processed as a Dask dataset. We utilised the Landsat Water Observations Collection 3 from August 16, 1986, to February 1, 2024, with a spatial resolution of 30 metres. All available images within the specified timeframe and study area bounding box were included, provided they contained at least 40% valid pixels and excluded those with contiguity issues. Pixel values were assigned as -1 for no data, 0 for other values, and 1 for water.

3.3 Refactoring the algorithm

The water masks dataset was loaded into memory as an Xarray Dataset to ensure compatibility with the original version of iRivermetrics, which was not fully compatible with Dask. To establish a performance baseline, we systematically analysed the algorithm's runtime behaviour using a combination of timers (e.g., timeit) and deterministic profilers (e.g., cProfile, line_profiler, memory_profiler) to measure execution time and resource usage. This approach helped identify critical areas for optimisation, such as excessive memory usage, inefficient CPU utilisation, and suboptimal data layouts that increased latency. Each function within the algorithm was evaluated to determine those requiring improvement for enhanced efficiency and Dask compatibility. We also assessed the necessity of extra preprocessing steps to ensure high-quality input data. The refactoring process was compartmentalised to guarantee that modifications were tested against the established baseline, maintaining result integrity. For the final efficiency test, we compared the average execution time over ten runs in two scenarios: (1) processing the study area as a single 100 km section and (2) dividing the study area into 50 equal-length sections of approximately 2 km each (Figure 1). All benchmarks were conducted on a workstation equipped with an Intel Core i7-13700K (16 cores: 8 performance + 8 efficient), 64 GB RAM, an NVIDIA RTX 4080 GPU, and a Samsung 990 Pro NVMe SSD. Using Dask's out-of-the-box configuration to leverage all 16 cores, we report wall-clock times rather than aggregate CPU-time sums—wall-clock measurements most accurately reflect end-user performance in our parallel pipeline.

4. Results

4.1 Algorithm Refactoring

4.1.1 Performance review

The refactored iRivermetrics algorithm demonstrated significant improvements in performance and resource efficiency compared to the original version. By incorporating Dask for parallel processing and lazy evaluation, the execution time for processing the Gilbert River dataset—which spans 38 years and includes 505 valid timesteps—was reduced by approximately 95%, decreasing from 65 minutes and 7 seconds to 3 minutes and 11 seconds for the entire study area analysed as a single 100 km section (Figure 2). Additionally, memory usage was optimised, with peak consumption reduced by 28%, enabling the algorithm to handle larger datasets without exceeding system limitations.

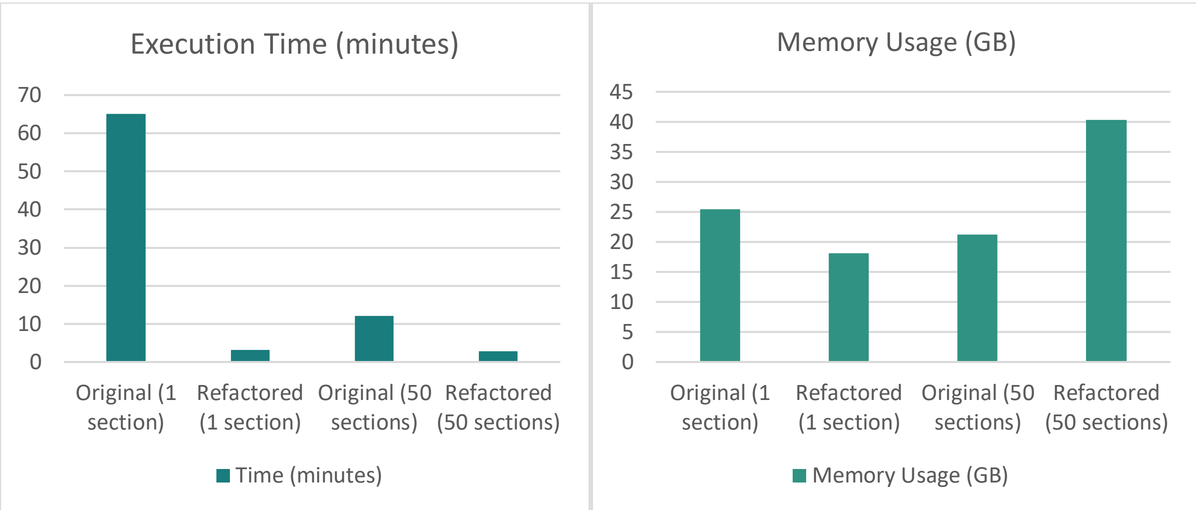

Figure 2. Comparative execution times (in minutes) and peak memory usage (in Gb) for the original and refactored iRivermetrics algorithm, tested on the 38-year Gilbert River dataset across different configurations.

When the study area was divided into 50 sections of approximately 2 km each, the refactored algorithm continued to outperform the original version, reducing execution time by around 77%—from 12 minutes and 1 second down to 2 minutes and 54 seconds. Although memory usage increased from 21.1 GB to 40.3 GB due to the overhead of parallel processing and concurrent task handling, this increase is not inherently detrimental. The benefits of reduced execution time and enhanced scalability outweigh the additional memory requirements, especially since the use of delayed functions effectively manages computation timing, preventing memory crashes.

4.1.2 Enhanced pre-processing and quality assurance

In addition to performance optimisations, the refactored algorithm incorporated enhanced data preprocessing steps to ensure high-quality input data. First, we excluded timesteps that contained only 'no data' values within the study area boundaries, ensuring that only relevant data contributed to subsequent analyses. After masking the data array, we calculated the percentage of valid pixels within the study area and excluded timesteps with less than 70% valid data. Missing values were addressed using a forward fill method over time, replacing 'no data' with the next valid observation or the second valid observation if necessary. A backward fill method was applied for values at the end of the time series, or, if forward filling was insufficient, it was propagated up to two observations before the missing value. If none of these conditions were met, the missing value remained as 'no data'. After filling no data values, we re-evaluated the data and excluded any time steps with less than 95% valid pixels, providing an additional layer of quality assurance. These quality control measures are critical, especially for calculations related to connectivity, fragmentation, and the number of pools, as layers with a high proportion of no data can falsely break connectivity, generating misleading conclusions and diminishing the reliability of the results.

4.1.3 Challenges and solutions

During the development and testing phases of the refactored iRivermetrics algorithm, we encountered several challenges that were systematically addressed to enhance scalability, reliability, and overall efficiency. One of the primary challenges was scalability and memory management. The original algorithm's approach of loading the entire raster dataset into RAM led to significant scalability issues and frequent memory crashes when processing large datasets, particularly on machines with limited memory resources. To overcome this limitation, the refactored algorithm utilises Dask for parallel processing and lazy evaluation, allowing more efficient memory management. Vectorised functions, utilising Xarray, NumPy and Pandas, enable the processing of entire arrays or data frames in a single operation, significantly speeding up numerical computations compared to traditional loop-based approaches. Additionally, we employed batch processing techniques to handle large time series in manageable chunks, further optimising overhead, memory usage and processing speed. These process-efficient functions collectively contributed to mitigating the risk of memory crashes and enabled the algorithm to handle larger datasets and longer time series without compromising stability.

Another critical challenge was managing the overhead associated with parallel processing. The initial parallel approach generated an excessive number of small tasks, which degraded performance due to high task management overhead. Conversely, the original algorithm processed polygons in a serialised manner, which was more efficient for handling many small polygons by reducing task management overhead. To mitigate this, the algorithm was optimised by grouping several time steps into single tasks and ensuring that each task was substantial enough to benefit from parallel processing without incurring excessive scheduling overhead. Profiling specific functions also helped in lowering overhead. This balance between parallelisation and task size was crucial in enhancing both execution speed and scalability without leading to significant increases in memory usage.

Limited input flexibility was also a concern, as the original algorithm only accepted data arrays for water masks, restricting integration with diverse data sources and workflows. The refactored algorithm addressed this by expanding input capabilities to include folder directories, Xarray, and Dask DataArrays and Datasets. This enhancement facilitates seamless integration with long-time series data and cloud sources, such as the Digital Earth Australia catalogue, thereby increasing the algorithm's versatility and usability.

After profiling the original algorithm and evaluating specific functions, a major constraint was the computational intensity involved in calculating pool start and endpoints, as well as pool length, which relied heavily on pixel-level operations. The original implementation utilised image processing techniques, including the Scikit-image morphology module and OpenCV's convolution filtering. This process encompassed several complex steps, such as labelling and skeletonising the water mask and calculating the Euclidean distance between grouped pixel pairs (Tayer et al., 2023b). Refactoring this component was essential to enhancing the overall efficiency of the algorithm.

The original code applied a custom kernel to calculate pool start and endpoints to the binary water mask. This operation resulted in a maximum possible pixel value of 18 if a water pixel was surrounded by eight other water pixels and a minimum value of 0 if surrounded by eight non-water pixels. We observed that the most common scenarios for start and endpoint pixels in the skeletonised image had either one or three neighbouring water pixels, corresponding to pixel values of 11 or 13, respectively. For each unique pool's potential start and endpoints, we calculated all possible combinations and the Euclidean distances between these pairs. Subsequently, we selected the five most distant pairs for further analysis using the Scikit-image graph module 'MPC Geometric' to find the least-cost path from two points through a given cost array (van der Walt et al., 2014), utilising the river centreline as the path. The most distant path between the selected start and end points was designated as the final segment, representing the pool length. However, this approach was both computationally costly and prone to errors, limiting the algorithm's scalability and reliability.

In contrast, the refactored code leverages graph theory and the igraph library, enabling skeletonised pixels to be represented as nodes within a graph structure. This methodology streamlines the identification of the longest path within each pool segment by implementing a breadth-first search (BFS) algorithm, significantly reducing calculation complexity and enhancing process reliability. As a result, the refactored method improves both accuracy and efficiency. Overall, the refactored code simplifies the logic of calculating pool length, allowing the algorithm to manage larger datasets and longer time series with greater ease and reliability.

Although this work focused on the technical optimisation and parallelisation of the iRivermetrics pipeline, users must recognise that the choice of river section definitions directly shapes the ecological meaning of per-section metrics. In our tests, we employed uniform sections to illustrate performance; however, any partitioning, ranging from a single catchment-wide polygon to geomorphologically informed reaches, can be applied to fit specific objectives (Tayer et al., 2025). Metrics are inherently sensitive to boundary placement—minor shifts can split or merge pools across adjacent sections, introducing artificial variability (Tayer et al., 2023a, 2025). We therefore recommend selecting segmentation rules that align with the ecological question at hand.

The refactored iRivermetrics algorithm demonstrated significant improvements in performance and resource efficiency compared to its original version. By addressing identified bottlenecks and optimising memory management, the algorithm’s efficiency and scalability have markedly increased, enabling it to support larger datasets more effectively and integrate seamlessly with modern data handling frameworks for enhanced geospatial data processing. Through comprehensive profiling and targeted refactoring, the algorithm not only became more efficient but also more reliable and scalable. These advancements collectively position the refactored iRivermetrics algorithm as a robust tool for assessing morphological and resilience metrics related to the hydrology of intermittent rivers, thereby contributing to sustainable water resource management in the Gulf of Carpentaria and beyond.

4.1.4 Summary of challenges and solutions

Table 2. Key challenges encountered during the refactoring of the iRivermetrics algorithm, the original approaches employed, and the corresponding refactored solutions.

Challenge	Original Approach	Refactored Solution
Scalability with Large Datasets	Loaded entire dataset into RAM, causing memory crashes.	Integrated Dask for parallel processing and lazy evaluation, enabling the handling of larger datasets by processing data in chunks and optimising memory usage.
Performance with Numerous Small Tasks	Serialisation using for loops led to inefficient task handling, especially with many small tasks.	Introduced batch processing of sections, enhanced task scheduling with delayed functions, and optimised function execution to reduce overhead and improve performance.
Computational Intensity of Pool Calculations	Relied on pixel-level operations, which were computationally expensive and time-consuming.	Leveraged graph theory with the igraph library, implemented BFS algorithm for pool calculations and utilised NumPy to efficiently handle distance calculations, reducing computational complexity.
Memory Usage Trade-offs	Minimal overhead but high memory usage, limiting scalability.	Slight increase in memory usage due to parallel processing, but significant speed improvements and better resource management.
Data Quality Control	Applied backward and forward fill methods to handle missing values.	Implemented robust quality control measures, including the exclusion of timesteps with insufficient valid data, and applied backward and forward fill methods to handle missing values more effectively.
Limited Input Flexibility	Accepted only specific input types (e.g., data arrays), limiting integration with other data sources.	Expanded input flexibility to include folder directories, Xarray, and Dask DataArrays and Datasets, ensuring compatibility with a wider range of data sources and workflows and enabling lazy evaluation for optimised resource management.

References

Broadley, A., Stewart-Koster, B., Kenyon, R.A., Burford, M.A., Brown, C.J., 2020. Impact of water development on river flows and the catch of a commercial marine fishery. Ecosphere 11, e03194. https://doi.org/10.1002/ecs2.3194

Ellis, E.A., Allen, G.H., Riggs, R.M., Gao, H., Li, Y., Carey, C.C., 2024. Bridging the divide between inland water quantity and quality with satellite remote sensing: An interdisciplinary review. WIREs Water 11, e1725. https://doi.org/10.1002/wat2.1725

Feyisa, G.L., Meilby, H., Fensholt, R., Proud, S.R., 2014. Automated Water Extraction Index: A new technique for surface water mapping using Landsat imagery. Remote Sens. Environ. 140, 23–35. https://doi.org/10.1016/J.RSE.2013.08.029

Garrido-Rubio, J., Calera, A., Arellano, I., Belmonte, M., Fraile, L., Ortega, T., Bravo, R., González-Piqueras, J., 2020. Evaluation of Remote Sensing-Based Irrigation Water Accounting at River Basin District Management Scale. Remote Sens. 12, 3187. https://doi.org/10.3390/rs12193187

Godfrey, P.C., Pusey, B.J., Pearson, R.G., Arthington, A.H., 2022. Predictable hydrology, habitat and food resources determine fish recruitment dynamics in an incised tropical Australian river. Ecohydrology 15, e2457. https://doi.org/10.1002/eco.2457

Huang, C., Chen, Y., Zhang, S., Wu, J., 2018. Detecting, Extracting, and Monitoring Surface Water From Space Using Optical Sensors: A Review. Rev. Geophys. 56, 333–360. https://doi.org/10.1029/2018RG000598

Johary, R., Révillion, C., Catry, T., Alexandre, C., Mouquet, P., Rakotoniaina, S., Pennober, G., Rakotondraompiana, S., 2023. Detection of Large-Scale Floods Using Google Earth Engine and Google Colab. Remote Sens. 2023 Vol 15 Page 5368 15, 5368. https://doi.org/10.3390/RS15225368

Keller, K., Allsop, Q., Brim Box, J., Buckle, D., Crook, D.A., Douglas, M.M., Jackson, S., Kennard, M.J., Luiz, O.J., Pusey, B.J., Townsend, S.A., King, A.J., 2019. Dry season habitat use of fishes in an Australian tropical river. Sci Rep 9. https://doi.org/10.1038/s41598-019-41287-x

King, A.J., Gawne, B., Beesley, L., Koehn, J.D., Nielsen, D.L., Price, A., 2015. Improving Ecological Response Monitoring of Environmental Flows. Environ. Manage. 55, 991–1005. https://doi.org/10.1007/s00267-015-0456-6

Marshall, A., Wohl, E., Iskin, E., Zeller, L., 2024. Interactions of Logjams, Channel Dynamics, and Geomorphic Heterogeneity Within a River Corridor. Water Resour. Res. 60, e2023WR036512. https://doi.org/10.1029/2023WR036512

Mcgregor, G.B., Marshall, J.C., Lobegeiger, J.S., Holloway, D., Menke, N., Coysh, J., 2018. A Risk-Based Ecohydrological Approach to Assessing Environmental Flow Regimes. Environ. Manage. 61, 358–374. https://doi.org/10.1007/s00267-017-0850-3

Mueller, N., Lewis, A., Roberts, D., Ring, S., Melrose, R., Sixsmith, J., Lymburner, L., McIntyre, A., Tan, P., Curnow, S., Ip, A., 2016. Water observations from space: Mapping surface water from 25 years of Landsat imagery across Australia. Remote Sens. Environ. 174, 341–352. https://doi.org/10.1016/J.RSE.2015.11.003

Pusey, B.J., Kennard, M.J., Douglas, M., Allsop, Q., 2016. Fish assemblage dynamics in an intermittent river of the northern Australian wet–dry tropics. Ecol. Freshw. Fish 27, 78–88. https://doi.org/10.1111/eff.12325

Tayer, T.C., Beesley, L., Douglas, M., Bourke, S., Meredith, K., McFarlane, D., 2023a. Identifying intermittent river sections with similar hydrology using remotely sensed metrics. J. Hydrol. 626, 130266. https://doi.org/10.1016/j.jhydrol.2023.130266

Tayer, T.C., Beesley, L.S., Douglas, M.M., Bourke, S.A., Callow, J.N., Meredith, K., McFarlane, D., 2023b. Ecohydrological metrics derived from multispectral images to characterize surface water in an intermittent river. J. Hydrol. 617, 129087. https://doi.org/10.1016/j.jhydrol.2023.129087

Tayer, T.C., Beesley, L.S., Stuart-Koster, B., Bond, N., Douglas, M.M., Rossi, J., McGregor, G., Marshall, J., 2025. Mapping Resilience: A framework for analysing persistent pool dynamics in non-perennial rivers using remote sensing, discharge, and rainfall data [In preparation].

Warfe, D.M., Pettit, N.E., Davies, P.M., Pusey, B.J., Hamilton, S.K., Kennard, M.J., Townsend, S.A., Bayliss, P., Ward, D.P., Douglas, M.M., Burford, M.A., Finn, M., Bunn, S.E., Halliday, I.A., 2011. The 'wet–dry' in the wet–dry tropics drives river ecosystem structure and processes in northern Australia. Freshw. Biol. 56, 2169–2195. https://doi.org/10.1111/j.1365-2427.2011.02660.x