Australasian Transport Research Forum 2023 Proceedings
29 November – 1 December, Perth, Australia
Publication website: http://www.atrf.info

Adapting ADAPT: A service-based accessibility toolkit

Tristan W. Reed[1,2], Doina Olaru[2], Chao Sun[2], Sharon Biermann[2]

[1]School of EECMS, Curtin University, Kent St, Bentley, Western Australia

[2]PATREC, University of WA, Stirling Hwy, Crawley, Western Australia

Email for correspondence (presenting author): tristan.reed@curtin.edu.au

1. Background

Calculating accessibility is a common task undertaken during transport planning and analysis, which enables the analyst to determine areas of relative advantage and disadvantage. Understanding relative accessibility is important for urban planning to better ensure equity of opportunities between areas, balancing transport demand (reducing journey length, contributing to positive environmental outcomes) and helping to ensure a just society. Accessibility, while having differing definitions, refers to in this case the number of opportunities (such as jobs, hospitals or recreation facilities) that can be reached from a region, utilising a particular mode (or modes) of transport in a given time frame.

Most tools to calculate these metrics are either bespoke (designed for a specific urban area and/or transport mode) or difficult to use (requiring a large amount of work in configuration or data to set up a simple analysis). Some of these tools are open source, however, documentation is limited or outputs are geared towards a particular use case. At times, large amounts of proprietary-formatted data are required, with the user expected to understand the intricacies of the data formats, alongside the scope of the data required itself. This further creates a burden on users (such as practitioners and researchers), who often spend more time gathering and processing data, than undertaking modelling tasks.

Extending the goal of the authors' previous work (Sun et al., 2017), the generic Automatic Dynamic Accessibility Planning Tool (ADAPT, herein referred to as the 'monolithic tool') has been further enhanced as part of the Australian Transport Research Cloud (ATRC) project. The aim of the ATRC project was to provide an integrated modelling environment consisting of multiple disparate 'components' (herein referred to as the 'individual components'), sourced and developed by different transport research groups throughout Australia. These components were designed to be linked together in such a way that an output of one individual component became the input of another individual component. This interoperability was achieved through a common architecture, where each individual component executed one 'task' and conformed to a common schematic architecture defining the inputs and outputs of each 'task'.

This was complemented by packaging each individual component in a containerised execution architecture, including operating system and software dependencies. The overall project contributed to a common execution (modelling) environment, where each individual component could operate on its own, or in concert with other individual components.

Such an arrangement introduced enhanced flexibility in the modelling workflows, enabling the more extensive reuse of components originally developed for ADAPT, thus mitigating the limitations of its monolithic nature. The original goal of an open-source accessibility tool is maintained, offering generic analyses for any region or sub-region, if data can be provided in open standards-compliance or de facto standard formats used for this purpose, obtained through many different external tools or provided by third parties.

This paper describes the work undertaken with regard to the ADAPT tool as part of the ATRC project, which aims to achieve two objectives: firstly, to further improve the ability to generically calculate accessibility metrics and secondly, to show how the services (individual components, known as 'tools' within the ATRC platform) provided by the existing monolithic tool, once split out from it, can be used as part of a wider modelling environment. Finally, this paper will discuss further work still ongoing to further achieve these objectives.

2. Methods

The original ADAPT developed by Sun et al. (2017) was first analysed to determine its constituent components that could be decoupled from the monolithic tool. This single package already provided separate interfaces for the Travel Time Matrix Generator and Accessibility Analytics Service, hence separating these two individual components from the tool was the first concern. Next, each component was studied to determine what input data was required for their use and in what format, as well as what output data was generated and in what format to enumerate the requirements of the software.

Subsequently, software was developed such that each individual component was separated into its own script file, including only the necessary dependencies required for that particular component. Additional scaffolding was provided for containerisation and input schematisation (utilising the YAML markup language), as well as the error checking of inputs, per the ATRC requirements to ensure interoperability. This architecture was implemented where each individual component functioned as a 'single run'; that is, a set of input parameters (conforming to the schema) are specified, checked for validity, a process is then undertaken and then an output is generated (assuming no errors).

Software code was primarily written using Python 3, although some code exists in Java to bind together external dependencies, alongside shell script and Dockerfile for containerisation. Complexity within each individual component, which may require further transformations and iteration, is hidden from the user to enable the widest range of end users to directly employ these individual software 'packages'.

Improvements both in performance and features were implemented by studying the process undertaken by the software, examining and refactoring the software code already written (including dependency updates) and stakeholder engagement. This yielded a third component being split out, the Centroid Generator, which was improved to provide additional options for users. Each of these refined components (including their input and output options) is described in detail in the sub-sections below.

2.1. Travel Time Matrix Generator

The purpose of the Travel Time Matrix Generator is to generate a matrix of zone-to-zone travel times by a specified mode (or modes) of transport. This is primarily achieved through a wrapper around the OpenTripPlanner (n.d.) tool, although investigations were undertaken to utilise the Conveyal (2023) tool instead. However, this proved far more difficult when the same could be achieved through upgrading the version of OpenTripPlanner used (compared to the original monolithic ADAPT). The OpenTripPlanner component calculates a travel time between two points; hence, the travel time between the centroids of individual sub-regions serves as a representative approximation of the travel time between sub-regions. This travel time is then applied within the Accessibility Analytics Service, however, it can also be used with other tools both within and outside the ATRC. The travel time is calculated through routing between an origin and destination over the road, rail, cycle and pedestrian networks (utilising OpenStreetMap), as well as including the public transport timetables (supplied in GTFS format), where appropriate. Users can supply the ESRI shapefile of the sub-regions, an OSM file of the transport network in the area that is covered by the sub-regions, the GTFS formatted public transport timetables, the name of the field that identifies each sub-region in the ESRI shapefile and for public transport, the date and time when the travel time needs to be calculated.

As part of the ATRC, the Travel Time Matrix Generator component was split out from the monolithic tool (where it was joined to the Accessibility Analytics Service) and additional features and configuration options were added, which are accessible and provided through standardised input schematisation as well as being described in documentation. Additional travel modes were able to be specified, creating a new option in the configuration, utilising the features of OpenTripPlanner. The underlying code was re-written to include the newer version of OpenTripPlanner, increasing performance. The calculation of the centroids, the generation of the region's graph (an intermediate step combining the input data, which is required for generating the travel time matrix) and the actual calculation were split out as three options of the same component (which could further be split into individual tools) and additional options introduced to consider this. Finally, filtering of input files, specification of precision (to improve speed) and resource allocation options were also provided, alongside the ability to apply an alternate mode of calculation, which is faster and allows for smaller sub-regions to have accessibility calculated for them in a reasonable timeframe.

2.2. Accessibility Analytics Service

The Accessibility Analytics Service is the second service extracted from the monolithic ADAPT service; this service aims to generate metrics and visualisations of these accessibility metrics. The metrics calculated by this service are cumulative opportunity access metrics, using a travel time threshold function (although ongoing work described below aims to expand this to alternate metrics and functions). These metrics quantify the number of opportunities accessible from one sub-region, by comparing the travel time from this sub-region to every other sub-region and if the travel time is below the threshold, the destination sub-region is deemed 'reachable'. Each sub-region has a specified number of opportunities allocated and associated with it; these are accumulated for each reachable sub-region from the initial sub-region to determine the number of accessible opportunities.

Initially, this feature of ADAPT allowed the user to select 'directionality of accessibility', calculating opportunities inbound (with access calculated *from* each region) or outbound (with access calculated *to* each region), the travel time threshold, as well as allowing the user to specify the ESRI shapefile of the sub-regions, the travel time matrix (generated either from the above tool or externally to it) specifying the travel time between sub-regions and the opportunity file (specifying the number of opportunities in each sub-region). Users also supplied the identifier for the sub-region zones (in both the shapefile and opportunity file), the identifier for the column containing the opportunity values, as well as the desired output format (either an attribute table, a colour or monochrome PDF or PNG map, or a CSV file).

As part of the ATRC project, this process was separated from the monolithic tool (as it requires far less resources than the travel time matrix generator), inputs formally schematised, introducing error checking on the inputs and was extended to generate additional outputs (e.g. a GeoJSON format spatial data file for use in other tools or applications). The tool was also completely re-written from using a variety of frameworks in Python 2 to primarily utilise GeoPandas and Python 3, to further improve performance and ensure accuracy of outputs. This entailed the use of vectorised functions to undertake spatial operations within GeoPandas DataFrames, rather than the use of built-in data and programming structures in Python 2. Source code can be provided upon request to the corresponding author. Documentation was also provided for ease of use by others.

A newly added feature allows the users to specify whether relative or absolute accessibility (i.e., % of overall accessibility – percentage of opportunities that can be accessed from each sub-region) or a traditional volume-based accessibility (count of opportunities that can be accessed from each sub-region) was generated instead. An example of the output generated by this tool is provided below as Figure 1; in this case, it is a colour PNG map, as generated through the ATRC workflow runner, which is described below in the 'Case Study' section of this paper. This map shows the accessibility of jobs within 30 minutes throughout Perth, outbound, by private vehicles, at SA2 level.

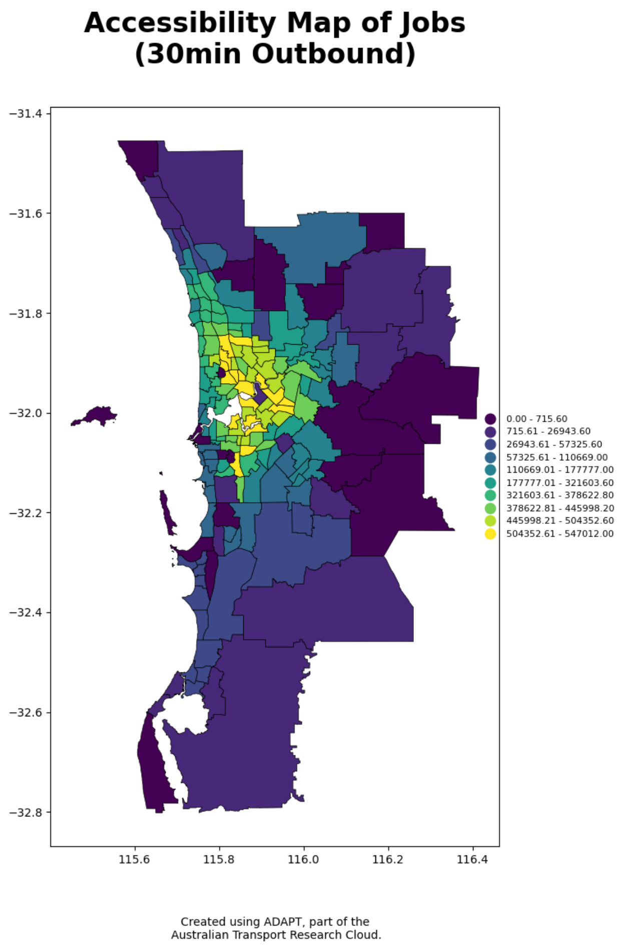

Figure 1. Output of Accessibility Analytics Service

2.3. Centroid Generator

The original monolithic version of ADAPT utilised a purely geometric process to determine the centroids of sub-regions – that is, finding the 'middle' of a region in terms of its minimum and maximum latitude and longitude values. This naïve method can work well in small, dense areas, however, in (geographically) larger sub-regions that have been drawn based upon population counts, this is unfitting. For example, for Australian Bureau of Statistics (ABS) Statistical Area, Level 2 (SA2) regions which form sub-regions of a larger city or regional area, this method calculates travel times from centroids that may be located in remote or uninhabited areas (the 'middle of nowhere'), rather than from regions where the majority of population or opportunities are generally concentrated. Hence, this approach may use travel times that are not representative of the experience of actual residents or for business opportunities.

As part of the ATRC project, the Centroid Generator was developed to provide an alternative to using direct geometric centroids for the Travel Time Matrix Generator. The centroids are calculated using a weighted (sub-)sub-region methodology, in the example above using ABS Mesh Blocks, as described below in relation 1:

$$L_D = \frac{\sum_{r=0}^{n}(P_r \times D_r)}{n} \tag{1}$$

where L_D is the weighted centroid value in one dimension (either easting or northing), the Mesh Blocks of $0 \dots n$ are denoted as r, the population of the Mesh Block r is denoted as P_r and the dimension (either easting or northing) of the Mesh Block r is denoted as D_r. Coordinates are converted between cartesian (easting / northing) and geodetic (longitude / latitude) as required.

The Centroid Generator requires users to supply an ESRI shapefile of the SA2s and Mesh Blocks, a CSV file that maps the ID of the Mesh Blocks to a weight value – such as the population, as well as the names of the ID and weighting columns of the Mesh Blocks. An output consisting of the SA2s and their centroids is generated.

3. Case study

The second objective of this work was to decouple and integrate the services (individual components) within the ATRC project to enable users to undertake analyses provided by the monolithic tool - cumulative opportunity accessibility analysis and (supporting this) travel time matrix generation as well as undertaking new feature development to further provide for the generation of desirable outputs in a simple and easy-to-use manner, as described above.

The three individual components described above are packaged as Docker containers, containing the analysis scripts and software dependencies required to run the code on any system architecture that runs the Docker hypervisor - which is the most common computer system currently in use. This means that these components can be run as part of the GitHub Actions modelling (workflow) environment provided as part of the ATRC project, utilising the Linux-based Docker image provided by the project on GitHub Actions (nested containers) defined using YAML, locally on the machine of user (if enough computational resources are available, utilising a shell script) or other systems such as supercomputers or cloud virtual machines (such as Amazon Web Services Elastic Cloud Compute instances). As indicated, each individual component has a simple and specified set of options (including data formats) to be supplied, described both in a machine-readable schema and human-readable documentation. These simple options supplied by the user map to more complex options 'behind the scenes' within the scripts and take away the complexity of the user employing the tool. In alignment with this approach, error checking is undertaken for each individual component. This involves assessing these values for reasonableness (in type and value) prior to proceeding with the main analysis.

Together with other test and integration exercises, two Development Case Studies were undertaken with the other ATRC partners using (in the first instance) an area of Melbourne surrounding Melbourne University (Carlton) and in the second instance, the overall Greater Melbourne area (where applicable, in SA2 regions). While the functions of both the Travel Time Matrix Generator and Accessibility Analytics Service were tested individually, the individual components were only tested together with each other, with a travel time matrix generated by the eponymous individual component as input, within the now separated Accessibility Analytics Service. The two case studies separately proved that the tools can be used within the wider ATRC platform and that these (and other) individual components can be orchestrated manually (in the first case study) and automatically from a script (in the second case study) together, in more complex modelling workflows. As more components are being developed and deployed within the platform, more complex workflows can be undertaken; an example of this would be using a GTFS cleaner developed by a different ATRC partner as input to the Travel Time Matrix Generator.

4. Discussion and further research

This paper has presented work undertaken by the authors to translate and extend the existing monolithic ADAPT generic accessibility tool into a new format consisting of smaller individual components, interoperable with other tools, and allowing better separation of functionality, while at the same time implementing new features. Each component of the monolithic tool can be run independently and with devoted enough computational resources for its purpose, only when required and in use, when instantiated following the easy-to-use instructions and generic data requirements. As part of the ATRC project, each of the individual components can be coupled with other components to enable new workflows to be undertaken, further improving the ability to solve problems with new analyses, such as cleaning GTFS data before using it to calculate a travel time matrix, as seen in the case study described above.

Further work can be undertaken to further split these components into even smaller independent units, such as each ‘mode’ of the Accessibility Analytics Server and the Travel Time Matrix Generator Service being its own module. Ongoing research by the authors is focused on incorporating a broader array of impedance functions as options to the user, moving beyond traditional accessibility thresholds. This expansion would enable, for example, the possibility to choose power or exponential functions, as well as to generate an additional ‘competitive accessibility metric’, where destination opportunities are shared among origin sub-regions, as described by Shen (1997) and expanded by Levinson & King (2020, p51-52). The authors look forward to sharing additional details about this further work currently underway at the conference. Integration of additional work already undertaken as part of other projects to develop distance-based metrics as part of the Travel Time Matrix.

Acknowledgement

This work was supported by the Australian Transport Research Cloud (ATRC) project, which received investment from the Australian Research Data Commons (ARDC) and technical support from the Australian Urban Research Infrastructure Network (AURIN), both of which are funded by the National Collaborative Research Infrastructure Strategy (NCRIS).

References

Conveyal, LLC. 2023. *Conveyal: Evaluate changes to your public transport system*. Available from: https://conveyal.com/.

Levinson, D., & King, D., 2020. *Transport Access Manual: A Guide for Measuring Connection between People and Places*. Committee of the Transport Access Manual, University of Sydney. https://ses.library.usyd.edu.au/handle/2123/23733

OpenTripPlanner. n.d. *OpenTripPlanner*. Available from: https://docs.opentripplanner.org/en/v1.5.0/.

Shen, Q., 1997. Location Characteristics of Inner-City Neighborhoods and Employment Accessibility of Low-Wage Workers. *Environment and Planning B: Planning and Design*, *25*(3), 345–365. https://doi.org/10.1068/b250345

Sun, Y., McCarney, G., Sabri, S., Amirebrahimi, S., Agunbiade, M., Chen, Y., Keshtiarast, A., Kalantari, M., Bright, C., Olaru, D., Reynolds, M., Rajabifard, A., & Biermann, S., 2017. A flexible accessibility analysis tool for enhanced urban analytics. *Australasian Transport Research Forum 2017 Proceedings*, 1–7.