

Reactive Documents for Modelling and Simulation

G.D. Denehy

*Advanced Modelling & Simulation Group,
Weapons and Combat Systems Division,
Defence Science and Technology Organisation
Email: gregory.denehy@dsto.defence.gov.au*

Abstract: Traditionally we find ourselves using static documents for reporting on analysis of modelling and experimental data. However, static documents constrain the author to defined formats for presenting information and explaining concepts. An author's thinking must be condensed using text, charts, diagrams and images representing selected snapshots or narrow slices through large datasets. Such formats limit the ability of the author to effectively communicate the subtleties of their insights with the reader. Access to large datasets must be referenced in an external source.

Reactive documents employ dynamic, interactive elements and visualisations within the document to explore data sets and models. This allows the reader to play with the author's assumptions and analyses, and see the consequences. By providing direct access to interactive data visualisations the reader is able to explore and verify the author's claims.

The primary motivation for incorporating reactive documents into analysis reports is to improve the communication of concepts and ideas, and to increase the trust the reader places in the claims made by the author. Reactive documents enable and encourage active reading by providing the reader with an 'environment to think in'. The author can guide the reader through a cogent analysis of the data, where the reader can further explore the data for themselves and test the author's claims. Traditional reporting approaches require lengthy appendices or reference external data sources to provide the reader with such access to the data.

In order to evaluate the effectiveness of reactive documents custom data visualisations were developed for two separate analysis tasks; a study into the mobility of a convoy limited by sensor coverage, and a study exploring the performance of Time Management within constructive High Level Architecture (HLA) simulations. These were developed utilising standards-based, client-side web technologies, namely HTML5, CSS, and JavaScript. The use of web technologies allows the documents to be read by a wide audience. Numerous development resources are available with a significant influx in web-based data visualisations and associated libraries over recent years.

Our initial work into implementing reactive documents has yielded several insights. Firstly, the use of visualisation for reporting can often be more constrained than for initial analysis work, where an analyst is likely to require a much more flexible environment within which to explore the data. Secondly, interactive data visualisation does not inherently avoid introducing author biases, as every visualisation is designed with a particular view of the data in mind. Thirdly, custom visualisations require careful design, particularly where interactivity is concerned. Even carefully designed interactions may not be obvious, and it may be useful to provide textual help to direct the reader. Fourthly, existing business processes may make this approach difficult to adopt. The use of static documentation for reporting is entrenched in such processes, both internal to the organisation and for external publications. Finally, reactive documents require a new design language. There is currently a lack of authoring tools targeting reactive documents where the integrated and interactive nature of the document requires a new way of thinking about how the document's message is delivered.

This initial work has generated enough interest and it has been determined there is enough potential to warrant further exploration of this approach. It has been demonstrated that reactive document techniques can be an effective tool to improve communication of concepts and results.

Keywords: *Reactive documents, analysis, visualisation, web technologies*

1. INTRODUCTION

Traditionally we find ourselves using static documents for reporting on analysis of modelling and experimental data. However, static documents constrain the author to defined formats for presenting information and explaining concepts. An author's thinking must be condensed using text, charts, diagrams and images representing selected snapshots or narrow slices through large datasets. Such formats limit the ability of the author to effectively communicate the subtleties of their insights with the reader. Access to large datasets must be referenced in an external source.

In discussing 'Explorable Explanations' Victor (2011a) asserts that a static document presents the author's claims and nothing else. The reader is unable to answer further questions, explore alternatives or verify assumptions. He presents three concepts to overcome such limitations; explorable examples, reactive documents, and contextual information.

Explorable examples 'make the abstract concrete, and allow the reader to develop an intuition for how a system works'. Using interactive components integrated with the explanation of a system or model allows the reader to see the behaviour of the system, and how the behaviour changes when variables change. Explorable examples are typically used for complex systems, as described by Victor (2011b), but can also be very useful in describing algorithms, as demonstrated by Bostock (2014).

Reactive documents employ dynamic, interactive elements and visualisations within the document to explore data sets and models. This allows the reader to play with the author's assumptions and analyses, and see the consequences. By providing direct access to interactive data visualisations the reader is able to explore and verify the author's claims. Barry and Card (2014) implement this well by providing such a visualisation of the Boston subway, guiding the reader to data points of interest whilst giving the reader an opportunity to explore the data for themselves.

Contextual information provides the reader with additional material in place, removing the need for the reader to separately search for related material.

Our initial work has focused on using reactive documents to report findings from modelling and simulation based analyses. Future activities will expand the scope to include concepts from explorable examples and contextual information. This paper presents the approaches taken to incorporate reactive elements for two use cases and details the lessons learnt.

The primary motivation for incorporating reactive documents into analysis reports is to improve the communication of concepts and ideas, and to increase the trust the reader places in the claims made by the author.

As Victor (2011a) states, 'an active reader asks questions, considers alternatives, questions assumptions, and even questions the trustworthiness of the author'. Reactive documents enable and encourage active reading by providing the reader with an 'environment to think in'. The author can guide the reader through the data analysis, where the reader can further explore the data for themselves and test the author's claims. Traditional reporting approaches require lengthy appendices or reference external data sources to provide the reader with such access to the data.

The increased transparency that results from deeper access to data also assists the reader in establishing a greater level of trust in the author's claims. The reader can directly assess the assumptions and claims made by the author rather than relying upon the particular view of the data provided by the author.

2. APPROACH

In order to evaluate the effectiveness of reactive documents custom data visualisations were developed for two separate analysis tasks; a study into the mobility of a convoy limited by sensor coverage as detailed in section 3.1, and a study exploring the performance of Time Management within constructive High Level Architecture (HLA) simulations detailed in section 3.2.

This work is part of a wider research topic focused on the use of web-based technologies applied to modelling and simulation. Using standards-based, client-side web technologies, namely HTML5, CSS, and Javascript, allow the documents to be read by a wide audience. Browsers are ubiquitous and therefore the reader does not require additional software to be installed, including on mobile devices. Numerous resources are available and recent years have seen a significant influx in web-based data visualisations and associated libraries. Furthermore, the technologies provide an environment flexible enough to create custom data visualisations.

Other potential technologies have been identified, including the Computable Document Format, Mathcad, Tableau and Plotly, however these have not been evaluated at this time.

3. DISCUSSION

3.1. Mobility Study

The first of two use cases described in this paper is based on a study into the mobility of a convoy limited by sensor vehicle coverage. The convoy is required to travel a specified distance while remaining within continuous sensor coverage.

The initial stage of the study used preliminary simulation data based on an arbitrary set of input parameters, generating results for transit duration and average convoy speeds for each of 2500 scenarios. Each scenario also generated a time series dataset representing the state and position of each sensor at five-second intervals, over 20 million data points in total.

A study of this type would normally be analysed using existing tools, such as Excel, Matlab or R. The analysis report would make use of selected views of the data in tables or charts to support the conclusions of the analyst.

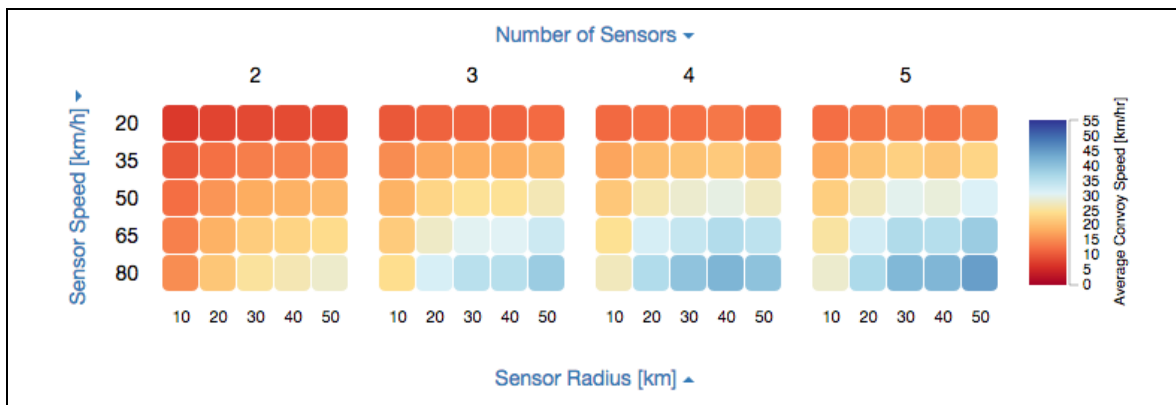


Figure 1. Multi-faceted heat map representation of results

A set of interactive visualisations was developed, providing the basis for both an analysis tool and a set of components for use in a reactive document. Figure 1 shows a multi-faceted heat map representation of the results. The variable represented by each of the three axes is selectable, as shown in Figure 2, allowing the reader to change the view of the data. The value of each heat map cell represents the mean of the data points matching the corresponding axis values.

The full set of scenario results are provided in a dynamic table, shown in Figure 3. The reader is able to selectively filter each of the scenario inputs. Selecting an individual heat map cell will set the table filters to match the corresponding axis values, providing an interface for the user to view the data points contributing to the cell value.

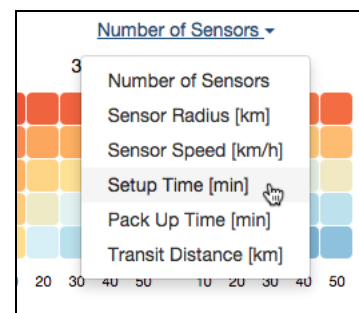


Figure 2. Heat map axis selection

Selection of an individual table row presents a view of the time series data for that particular scenario, shown in Figure 4. The view provides a timeline of each sensor’s state in addition to a depiction of the position of each sensor at a particular point in time. The view is animated, showing the sensors move across the diagram as time progresses. Furthermore, the reader is able to drag the handle across the timeline to move back and forth through time to more closely examine how the scenario executes.

The implementations of these three visualisation components are loosely coupled. Each component can be manipulated separately, and can be configured programmatically. This provides the author with the flexibility to guide the reader through views of the data that are of particular interest. For instance, hyperlinks in the report text can be used to reconfigure the components or selecting points in one component can be used to configure another.

Number of Sensors	Sensor Radius [km]	Sensor Speed [km/h]	Setup Time [min]	Pack Up Time [min]	Transit Distance [km]	Transit Duration [hr]	Average Convoy Speed [km/hr]
Filter: 5	Filter: 50	Filter: 80	Filter	Filter	Filter		
5	50	80	5	5	210	3.7939	55.352
5	50	80	5	8.75	210	3.9189	53.587
5	50	80	5	12.5	210	4.2528	49.379
5	50	80	5	16.25	210	4.4403	47.294
5	50	80	5	20	210	4.6278	45.378
5	50	80	8.75	5	210	4.2744	49.129

Showing 25 of 2500 results [Clear all filters](#)

Figure 3. Dynamic table of scenario results

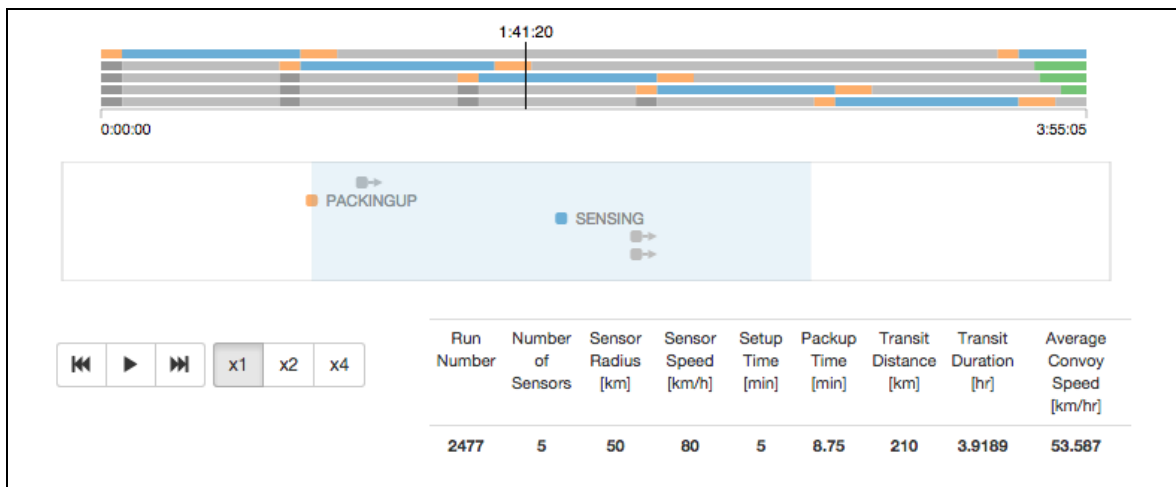


Figure 4. Time series view of an individual scenario run

3.2. RTI Performance

The second use case was supporting a study exploring the performance of Time Management within constructive High Level Architecture (HLA) simulations. The study assessed performance by measuring the wallclock time required to complete an individual simulation frame and the wallclock time to complete the entire simulation. A large number of simulation executions were configured using pairs of federates with different parameter values to identify key correlations between parameter combinations and clusters of performance measures.

As with the first use case, a dynamic visualisation was implemented to assist with the analysis of results and to provide a basis for exploring the use of reactive documents. This study also involved a multivariate dataset with the potential for performance results being dependent upon combinations of parameters. Initial analysis had already identified several clusters of measured performance, as can be seen in Figure 5, and had identified the Timestep parameter as the likely contributor to the variations in performance.




Figure 5. RTI performance results ordered by time

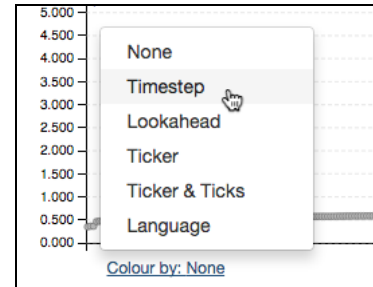


Figure 6. Colour selection

By providing the reader with the option to colour code the chart by a specific parameter, as seen in Figure 6, the Timestep parameter is clearly identified as a contributing factor. While this could have been achieved using a static chart, the use of a dynamic chart gives the reader the opportunity to confirm this by exploring other parameters.




Figure 7. RTI performance results coloured by timestep

A further option was added to group the results by a specific parameter. In conjunction with the colour coding this provides an interface to explore how combinations of parameters impact the performance results. As shown in Figure 8, this approach identifies that the Lookahead combined with the Timestep may also be contributing to the performance variations.

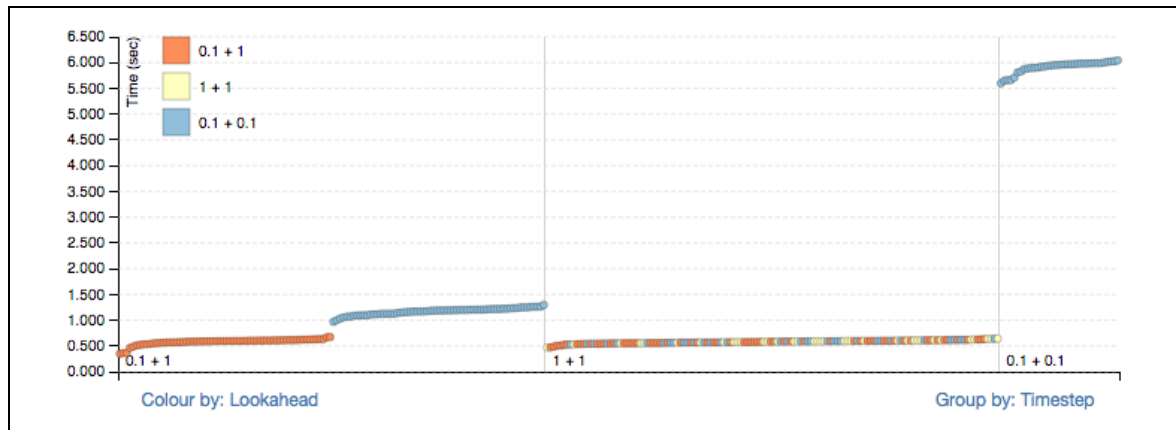


Figure 8. RTI performance results coloured by lookahead and grouped by timestep

While the primary motivation is to provide the reader with an environment to analyse the entire dataset, this approach to data visualisation also allows the author to replace multiple static charts with a single dynamic chart. Multiple static charts would be required to show the various data views supporting the author's claims, just as is necessary with this paper.

4. LESSONS

Our initial work into implementing reactive documents has yielded several insights.

Building a visualisation for analysis can be different to building a visualisation for reporting. The use of visualisation for reporting can often be more constrained than for initial analysis work. This is largely due to the document's space limitations combined with the need for intuitive interactions that do not overwhelm the reader. An analyst is likely to require a much more flexible environment within which to explore the data and need not be concerned with concise representations. The data visualisations used for reporting can be directed towards showing the report's findings while ignoring views and options that add no further value to the analysis.

Just as static charts can be presented with an author's bias, the choices made when designing and developing a visualisation can also be biased. Static charts present a particular view of the data and, unless care is taken, there is potential for that view to be biased. The role of reactive documents is to provide deeper access to the data to help avoid such biases, however, every visualisation is still designed with a particular view of the data in mind.

Intuitive data visualisation is hard. Custom visualisations require careful design, particularly where interactivity is concerned. Even carefully designed interactions may not be obvious, and it may be useful to provide textual help to direct the reader. The use of hyperlinks in the text can also be useful in guiding the reader, showing where the visualisation can be manipulated. Barry and Card (2014) use a combination of approaches to great effect, where hyperlinks are used to highlight or direct the reader, and additional instructional text is provided where appropriate.

Existing business processes may make this approach difficult to adopt. The use of static documentation for reporting is entrenched in such processes, both internal to the organisation and for external publications. The ability to generate static output from a reactive document may be useful to explore, particularly in the context of printed output.

Reactive documents require a new design language. Beyond the basic visualisations using charts and other standard components, there is currently a lack of authoring tools targeting reactive documents. Custom data visualisation requires dedicated development time, skills that analysts are unlikely to have. As stated by Goldstein (2015) 'explorable explanations are highly multidisciplinary. They draw on data visualisation, information design, programming, mathematics, [and] storytelling'. The integrated and interactive nature of the document requires a new way of thinking about how the document's message is delivered. Victor (2013) describes this as 'a new medium for understanding' and therefore we require a new language for describing how data is explored in this medium.

5. CONCLUSIONS AND RECOMMENDATIONS

This initial work has generated enough interest and it has been determined there is enough potential to warrant further exploration of this approach. Reactive document techniques can be an effective tool to improve communication of concepts and results, however, there are several areas that require further exploration.

We see a need to evaluate the connection between analysis tools and the associated reporting mechanisms. As identified above, the requirements for an analysis tool can be different from those for reporting in a reactive document. An evaluation of existing tools recognised in section 2 is necessary to identify synergies with current analysis tools, and to evaluate gaps that may make it necessary to build custom components that can be reused across analysis and reporting.

We intend to investigate the potential for embedding modelling and simulation directing into reactive documents, rather than relying upon the visualisation of a dataset. For cases like the mobility study detailed in section 3.1, it may be more efficient to embed the simulation in the document to generate the time series visualisation on demand rather than load the pre-executed results. This would not only reduce the data download requirements, but open the potential for the reader to explore beyond the boundaries of the existing dataset.

Finally, it is important to find an appropriate mechanism for authors to build visualisations into reactive documents. This extends beyond the tools available, as the authors need to understand the concepts behind reactive documents and how they can be used to better convey their ideas. Incorporating reactive components into a document requires more than using dynamic charts, and authors will need to develop an appreciation of what is effective.

ACKNOWLEDGEMENTS

The author would like to acknowledge Matthew Christie, Stefano Landi, Edward Dawson and Marcus Saunders for their time and effort incorporating this work in their mobility study, and to Anthony Cramp for the RTI Performance data.

REFERENCES

- Barry, M. and Card, B. (2014). Visualizing MBTA Data. <https://mbtaviz.github.io>
- Bostock, M. (2014). Visualizing Algorithms. <http://bost.ocks.org/mike/algorithms>
- Computable Document Format, <https://www.wolfram.com/cdf>
- Goldstein, M. (2015). Exploring “Explorable Explanations”. https://medium.com/@Max_Goldstein/exploring-explorable-explanations-92f865c8d6ba
- Mathcad, <http://www.ptc.com/product/mathcad>
- Plotly, <https://plot.ly>
- Tableau, <https://www.tableau.com>
- Victor, B. (2011a). Explorable Explanations. <http://worrydream.com/ExplorableExplanations>
- Victor, B. (2011b). Scientific Communication As Sequential Art. <http://worrydream.com/ScientificCommunicationAsSequentialArt>
- Victor, B. (2013). Media for Thinking the Unthinkable. <http://worrydream.com/MediaForThinkingTheUnthinkable>