The Digital Controller Interface Standard (DCIS) – An Update on The Topographical, Configuration and Controller Data Standards Activities

Authors

Andy Graham, White Willow Consulting Ltd.
Mark Pleydell, PTC Ltd.
Peter Routledge, IRC Ltd
Jon Wade, Interoperate Ltd.

Abstract

While the safety and core functions traffic signals controllers are defined through European and UK standards and the UK has standards for network and systems operation, currently there are no UK equivalents for the disposition of signalised sites, the site-specific control information, nor live data from controllers. The DCIS project is addressing these areas.

MAPEM is a defined European standard for site topography. However, due to the nature of the source material, implementing MAPEM data files is not simple. Standardising site plans can substantially ease automated MAPEM generation and improve confidence in their accuracy.

Similarly, methods for site-specific configuration of signal-controlled sites are manufacturer dependent. The creation and management of a standard for the presentation and publication of this data supports many use cases and provides a structure for the live reporting of controller data for consumption by network services and related users up to and including autonomous vehicles.

1. Introduction

1.1. DCIS an overview

While most traffic engineers will be aware of TOPAS 2500, the traffic signal controller standard, and TSRGD, the Statutory Instrument governing, among other things, traffic signalling, and some will have an awareness of BS EN 50556 and BS EN 12675, the standards for traffic signalling system safety and controller safety respectively, reflection will show that there is currently no standard way of representing a) the locations and types of signal controlled sites, b) the topography of sites nor c) the configuration of the controllers.

The DCIS project was brought into being by the DfT through the Transport Technology Forum (TTF) to address these three gaps.

The objectives of the DCIS project are

I. to develop and deploy standards for these three areas,
II. put in place the processes for creating and managing the content, and
III. test this with the local authority and traffic signals manufacturing communities.

The deliverable standards are:

1.2. Intersection Summary Document (ISD)

This is a list of sites, locations, types and other high-level descriptors presented in a standard JSON format. Importantly, the ISD contains the intersection IDs for each site, in the UK these are known as the Simmonite numbers (in recognition of the contribution that the late Brian Simmonite made to the traffic control industry). The ISD is well-defined and a JSON schema exists to verify the content. With the classification of sites traffic streams are managed within the Simmonite number through a parent/child relationship tying all child streams to their parent controller.

The DCIS team have also trialled a process for implementing and deploying ISDs.

The scope of the ISD standard is structured to accommodate short-term deployment or portable or temporary traffic signals and these types of signal-controlled sites are also in scope for the MAPEM and E2500 standards.

The ISD has been presented several times since its definition and guidance published. Accordingly, it will only be referenced below where it assists in understanding the other elements of DCIS.

1.3. MAPEM – The Site Topography

MAPEM files are fully defined in European C-Roads standards. These are live and evolving so the DCIS team elected to use these. MAPEM files define the spatial extents

of the conflict zone, stop lines, lanes into and out of the site, including pedestrian pavements, and crossings, cycle lanes, bus lanes (by time of day) and so on. Importantly MAPEM also identifies which signals are in control of which traffic movements. Each MAPEM file contains the Simmonite number for that site. This is the cross-link between the ISD list and the specific site topography and signal disposition, and the site's controller configuration. The status and progress relating to MAPEM is discussed in more detail below.

1.4. E2500 and DCAG

This third deliverable has two parts. The first part is a representation of the static data that makes up the configuration of the site's traffic signal controller. It identifies the streams, stages, phases, phase types, conflict and inter-green tables and many other characteristics of the configuration. The E2500 is addressed in more detail below

The second part, DCAG, the (Dynamic Controller Automated Generation) is a file with the same structure as the E2500 within which the dynamic content of the signal controller may be reported by the controller to other systems.

Both the E2500 and DCAG files also include the Simmonite number as the unique identifier that allows systems consuming DCIS file content to associate information across the set of files correctly.

Development work on E2500 is also exposing areas within TOPAS 2500 where there is scope for better definition, or whether there are inconsistencies that need to be resolved. This is feeding into the TOPAS processes.

2. Contextual Considerations for The Generation, Storage, Security and Use of DCIS Files

The following sections consider the environment in which DCIS is being created and will be used, the factors that are being considered and how they are being addressed.

2.1. Resilience, Usefulness and File Validation

As with any data source that is made available to parties beyond the originators, there has to be confidence in the content and the formatting of that content, and of its usefulness and longer-term relevance.

To maximise relevance the development team are aware of and work to achieve extensibility and resilience in the design of the information content. They also work to limit duplication in the file structure.

Content formatting is checked in three ways.

- JSON parsing checks the syntax

- The JSON Schema checks the structure and declared constants
- Separate semantic and cross-field validation check internal consistency

Because the schema cannot use parametric data from within the file to test other parametric data, a number of dedicated context-specific unit tests, (pieces of code that read the file, extracts parameters and use those to test other characteristics), have been identified to 'sanity-check' files (e.g. do the dimensions of conflict matrices equal the numbers of phases?).

The schema and unit tests are applied to the files during their development, uploading, and subsequent publication.

2.2. Management of Signal Controller Data

As noted above in part II of the overview section 1.1 above DCIS is also addressing the information pipeline by which current data sources are converted into their corresponding DCIS formats and the structure and content is verified before the files are uploaded to an accessible platform where they are curated and from which they may be accessed by bona-fide users. The importance of file validation has already been mentioned.

More widely this task includes an assessment of the balance between the benefits of making the data available and the possible increase in risk to the network from such exposure. It should be stressed that there is no control pathway in the data structures, the flow is away from the roadside and more often away from systems that are distinct from the roadside and not connected to it, to other systems. That is there are no messages flowing towards the controller, which reduces the opportunities for malicious intervention. Further, as is discussed below, the processes for the generation, uploading, and downloading of DCIS files includes a number of security checks (as well as validity checks). These security checks will provide reassurance that data is legitimate, valid, timely, and authoritative.

2.3. Use Cases for DCIS Content

While ostensibly the DfT is seeking preparedness for the deployment of autonomous vehicles on the UK's road networks, work has been done within the project to look at other use cases for the DCIS file content, and those authorities that have kindly contributed DCIS thus far have been invaluable in offering suggestions and ideas. Two examples of situations where standardised data will simplify the extraction of useful information are provided below.

Having a standard format for the data that defines the road layout and the control applied to that site provides opportunities for faster and more extensive modelling and for closer alignment between the model and the real network control methods.

Having details of phases in stages and influences for stage changes allows a better understanding of stage-to-stage changes and timings.

2.4. Gaps and Areas Not Covered

As a part of the initial requirements for the project DCIS does not address the road links between signalised locations. While MAPEM is sufficient to define the links the C-Roads manual also does not address this part of the network, nor present it in any of the examples it offers. These are literal gaps in the file/data.

Many signalised locations also include regulatory signs that constrain traffic movements. These constraints are the subject of Traffic Regulation Orders and are digitised under the funding TRO project, and more broadly are covered by work in CEN and ISO on METR (Management of Electronic Traffic Regulations)

The following parts of this paper provide more insight into the MAPEM and E2500/DCAG files.

3. MAPEM

3.1. What Are MAPEM Files

Many different sectors have approached the process of standardising geographical data for their own purposes, and whole ISO and CEN working groups give time to this activity. Each user of geographic data places their own requirements on accuracy, features, readability, etc. For autonomous vehicles the trajectories they should follow to hold their lane (or pavement) position as they move through a network have to be defined to better than 10cm accuracy.

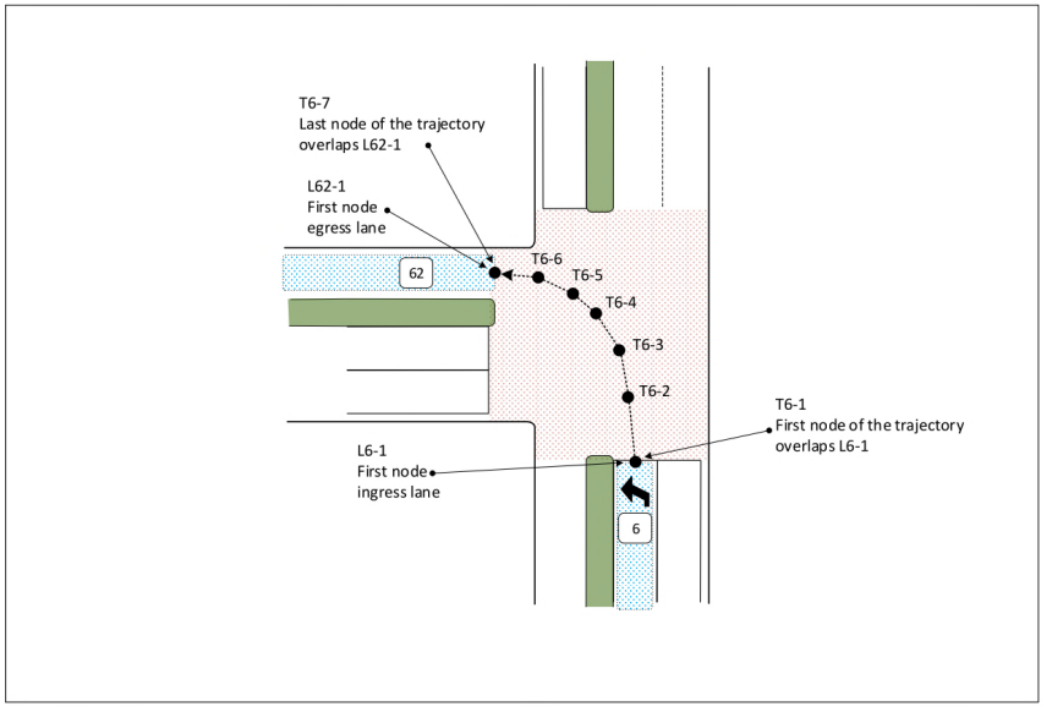


Figure 1 Diagrammatic representation of MAPEM content for a permitted (Right Hand Drive) trajectory

For this purpose the European C-Roads project[1] maintains a live set of standards for representing the topography of the ingress lanes, stop-lines, conflict zones and egress lanes of signalised locations on the road network and for the placement and identification of traffic signal hardware to control the specific traffic movements. MAPEM allows characterisation by permitted vehicle type(s), times of day, and other common attributes. A User Manual[2] is provided that ably leads a prospective user through the structure and content of MAPEM files.

The DCIS team elected to use MAPEM as its preferred standard for site topography.

Having identified the preferred file type we have to consider how best to implement those files using what is known to be available data from current signal-controlled sites.

3.2. What Existing Data is Needed and How is it Used

The main challenge we face is that authorities hold records of their sites in various formats from (for older sites and less well-resourced Authorities) hand marked up printouts of pdf files through to well-structured and layered CAD files (from recent sites and better provided for authorities). Each authority holds many such records, and the conversion of this content into MAPEM is not simple. We consider that, unless conversion can be automated, it is unlikely that these sites will ever be rendered in MAPEM without a very substantial injection of resource. Accordingly, the DCIS team

[1] https://www.c-roads.eu/
[2] https://releases.c-roads.eu/?sdm_process_download=1&download_id=327

placed a project with Imperial College London's Centre for Transport Studies as a Masters' student team project to implement an automated conversion tool.

The students were provided with six instances each from of site data each from Leeds City Council and Bath & North East Somerset Council covering two crossings, two simple intersections and two more complex sites.

The team broke the problem into stages. They reviewed the required criteria for MAPEM and classified each one. These are the data that have to be present in or inferable from the source data to ensure a MAPEM is realisable.

They concluded that creation of MAPEM from the source material is most reliable if there is a good mapping of elements in the source material to elements in the MAPEM file. If the originating source is not complete then the resulting MAPEM will necessarily be flawed. A measure of the completeness of source data helps to define the corrective actions needed to the source data to arrive at a sufficient set to fully enable a MAPEM outcome.

That being said, this is a necessary but not sufficient measure of the likelihood of generating the output file.

While we note the risk of conflicting content, access to and consumption of more than one source document will support each MAPEM element imbuing the conversion process with resilience. For each MAPEM element there is also a hierarchy of sources from which that datum can be drawn and an understanding of which sources are most reliable, allowing implicit weighting of these different sources.

Having made that assessment of completeness the team parsed each source document to extract and classify the quality of the source data items that would be used. Where data are not present or are of poor quality a report is generated to that effect.

This parsed data provides the input to a MAPEM file assembly process.

A clear outcome from the project is that a disciplined approach to the use of layers in CAD significantly simplifies MAPEM conversion and minimises the risk of errors creeping in through the conversion of features that resemble the required features, for example lining and marking can look a lot like kerb lines. Placing the boundaries of the carriageway and pavements on a CAD layer free from any other clutter would significantly assist in extraction of approach and egress lanes.

For new sites and reviews of existing sites where CAD exists there may be benefits in introducing additional layers that (for example) show ingress lane to egress lane moves.

Any MAPEM that is generated has to be assessed for its quality, both in completeness of coverage of a site and accuracy of the content that it provides. These assessments in

turn refine the definition of CAD layers and the content that they should include and possibly more importantly, not include.

3.3. Consequences for Existing Processes

From the work already undertaken it is clear (and unsurprising) that the best way to maximise the quality of MAPEM files is to minimise the sources of noise and error in the input data. The flattening of all CAD layers onto PDF files significantly complicates the task of extracting those features required for MAPEM without also picking up erroneous content. Where CAD files are available the ability to filter out layers that do not contain topographic content is greatly enhanced but this requires consistent CAD layering discipline. This discipline will extent to layer naming, so that, for example new and existing kerbs etc appear in a layer that is always called 'kerbing' for all users, and in the content of each layer, so kerbing will only include kerb details.

It will be apparent that applying a discipline in CAD file layer definition and specifically including one layer that ONLY has lining and marking and one that has kerb lines and one layer that only has pole locations and signal equipment. Keeping cabling and ducting, frames, version control, BoMs cabling schedules, horticulture and planting, etc. on other layers eases their exclusion from the file conversion process.

To anticipate the likely guidance: If in doubt use more layers and separation of content not less and always name layers consistently.

3.4. How will MAPEMs be generated, validated, stored, curated & accessed?

Subject to getting a reliable automation process running (and to the best of our knowledge there is no open MAPEM generation facility currently available in Europe), the UK generation of MAPEM (and E2500) files would likely be related to instances of Government funded projects, where MAPEM creation could be a deliverable under the conditions of funding. In this way, over a rolling period, authorities would create this content.

Once created the MAPEM requires validating both for syntactic correctness, but also in so much as can be confirmed, for geographical accuracy. Schemas and semantic validation checks can validate syntactic correctness and internal consistency. Geographic accuracy offers a greater challenge in that, if we had that data we would either preferentially be using it to generate MAPEM or it would render MAPEM redundant. Notwithstanding that constraint, feedback from usage will provide a valuable after-the-event way of checking accuracy and completeness to allow marking of files as suspect or workable for the use case in question. Regardless of the extent of the solution a checking process will be a part of the uploading of MAPEM files to the proposed storage. Any files that contain detectable errors at the time of uploading will be flagged to the authority uploading them for review and correction.

As noted above, the Simmonite number provides a means of linking the MAPEM site information to the ISD and to the E2500.

It is expected that, ultimately the generation of MAPEM would be by means of a tool written to work alongside CAD tools and which would generate the MAPEM for a site when the as-built drawing is signed off.

3.5. MAPEM File Usefulness

Every item of stored data begins to age from the moment it is captured. The reality can only diverge from the frozen-in-time data snapshot. Accordingly, the age of a file, or rather the age of the content used to create the file can give an indication of the likelihood of there being errors. Site layouts tend to change radically but infrequently, so a MAPEM file with an age of six months would be deemed likely to be accurate whereas a file at six years would be open to question. As well as getting feedback from MAPEM file use, the presence of short-term signal deployments at or close to the site (something discoverable from the ISD) would provide an indication of an intervention and therefore in all likelihood a change to the topography. That is, the presence of an intermediate set of works is also beneficial for users as it provides continuity of DCIS structured information to users as well as flagging a likely site change.

4. E2500 The Controller's Configuration Data

4.1. Overview

UK Traffic signal controllers operate according to a hierarchy of requirements, starting at the top with BS EN 50556 and BS EN 12675, both of which address many safety elements of traffic control and are called up in TSRGD. At the next level down, TOPAS 2500, the standard for traffic signal controllers ensures that these products deliver the modes of control (FT, VA, External control, CLF, EVP, Manual Control, etc.) and control structures (streams/stages/phases) for different types of sites. Then at the site-specific level we have the configuration data that defined how influences affect modes and timings and stage progression.

While the top two layers are open and generic, however the processes by which a traffic signal controller is configured for a specific site are bespoke to each signal controller manufacturer. This has several consequences, but the main effects are that there is no accessible way of converting existing configurations between these differing bespoke formats. This in turn substantially increases the effort involved in using the configuration data for other applications. Currently simulation models struggle to give an accurate representation of the controller deployed on street.

The E2500 part of the DCIS project implements a standardised form for configuration data. A preliminary draft of the JSON was published to UK traffic signal controller

manufacturers in 2025. At the time of writing the schema is substantially complete except for areas where the differences in signal controller company approach are substantially different and will require liaison to resolve and as noted below.

For clarity: The E2500 is NOT another version of that data input tool (colloquially 1827s or 2500s).

These forms provide the human interface to load new site design information.

Once the site configuration data has been entered into the controller tool suite it goes through a series of steps to verify that values are in permitted ranges, that data is complete and consistent, etc. Once those checks have been applied then a file is generated that is loaded to the controller. The E2500 is a parallel file to the one loaded to the controller. It is generated at the same step in the process, but it is made available for other permitted users to consume.

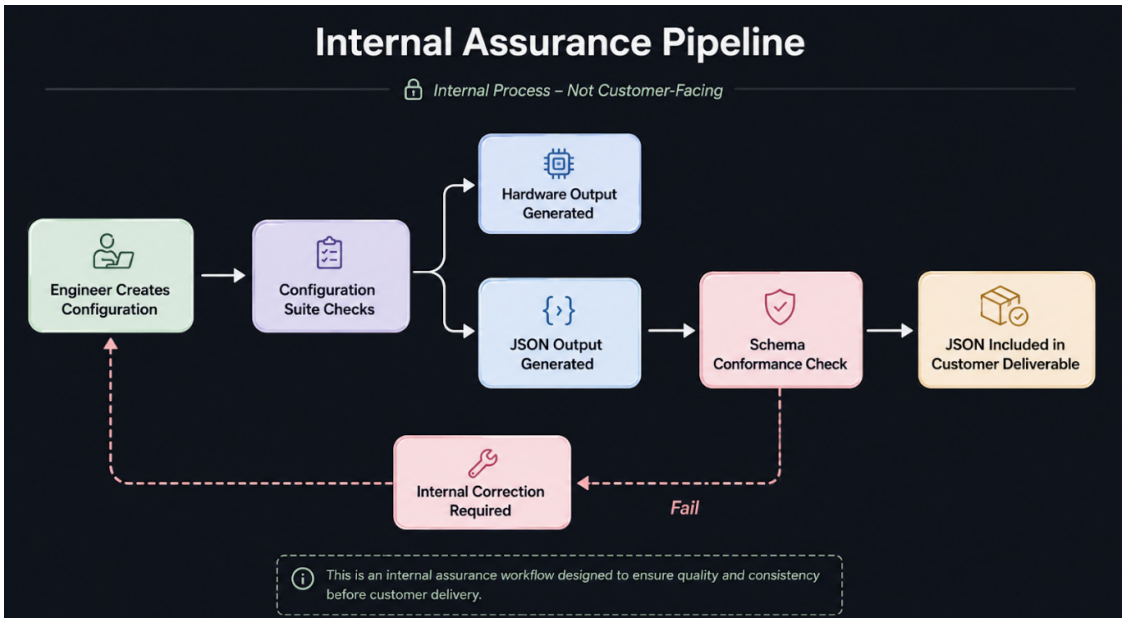


Figure 2 File assurance process

The main UK traffic signal controller manufacturers have each arrived at their current products through a rolling process of upgrades and evolution, so each carry and exhibit a lot of history. Put another way, in different controllers the same operational problem is provided in different ways so requires different parameters to define how that problem should be implemented on any specific site. This situation is exacerbated by the diversity of hardware platforms and operating systems of choice in controller CPUs and even more in the firmware architectures and information management that is used.

Accordingly, development of the E2500 standard has required the application of two principles, granularity and discoverability. Granularity is best defined by example. Consider where one product specifies a characteristic at the stage level and another at

the phase level. The E2500 requires that the characteristic is included in the JSON at the more granular phase level.

Discoverability is the principle that a value may be included where it is explicitly present in, or deterministically derivable from the configuration data and the derivation can be stated and reproduced.

For example, one controller manufacturer might define timetabling using a start date/time and elapsed intervals (during which events run), while another may use a series of specified dates/times (at which events start or stop). Either format allows the E2500 timetabling section to be 'discovered' and populated. Because there are notable differences in 'how' configuration information is implemented by different manufacturers the generation of a standard form of configuration file content does require an intermediate step from the signal companies to extract data from their own way of doing things, i.e. they also have to apply granularity and discoverability.

To summarise; the E2500 is a standardised representation of the uploadable file. Having that data in a standard form will support:

a) Generic representation of signal-controlled sites. Any tools or services wanting to represent a site can use the E2500 JSON
b) Interoperability: In principle an E2500 file from a controller made by manufacturer A can be used by manufacturer B to create a substantial part of the configuration file for their controller for the same site without going back to the original design forms and data.

The ability to understand what a signal controller is doing and will do next allows the processes running in vehicles to better adjust their approach to a signal-controlled location to meet a green signal (SPAT). However, having a standard form for the configuration also allows tools that predict timings and/or choices of stage-to-stage transitions to be implemented, it enables simpler implementation of controllers in simulation models AND more faithful simulation.

Turning to the content, what does an E2500 file include?

4.2. What is an E2500 File

A typical E2500 file comprises distinct blocks of content. The first block includes the site preamble, version and issue control, and the Simmonite number (the intersectionID), the subsequent blocks contain information about stream/stage/phase data, with types, conflicts, intergreens (absolute limits, handset limits, etc.,) all identified. There are blocks for I/O and the I/O processes including inversion, DFM, logical combination one to many mappings. Further blocks cover mode priority tables and processes and parameters related to those modes, CLF plans, entry points, etc,

timetabling is captured, and most challengingly special conditioning is, as a minimum, re-produced (see below).

The content is presented as a JSON file to optimise both machine and human readability. The data structure attempts to minimise the file size, avoid duplication of definition and partition the information into logical groupings. The structure is defined in the associated E2500 JSON Schema. Schemas are tools to confirm that data takes the right format, contains permitted values and information of the right type in the right place. In this respect the standard is largely composed of the schema.

In so far as different implementations of controllers allow, a complete E2500 contains everything that defines the operation of a particular site.

4.2.1. Special Conditioning

The different signal controller companies largely use special conditioning for two purposes. One is to implement generic functions that are not present in their controller firmware, the other is as a tool to address site-specific problems. As generic functions are largely addressed within the E2500 these areas of special conditioning lend themselves to a standardised interpretation. The site-specific facilities tend to be composed for logical functions, counters and timers and use data from within the controller. As such they also may render into discoverable relationships that can be presented in a standard form. Accordingly, it is anticipated that much special conditioning will be rendered into standardised forms. Exploring this idea further, the rationale for this proposal is that special conditioning starts and ends with real accessible influences (inputs, times of day, etc.) and outputs, (mode changes, stage changes, outputs being driven, etc.) and therefore the signal controller is able deterministically to process the special conditioning statement. If the signal controller can do this then the special conditioning functions can be standardised.

4.3. What is DCAG And How does DCAG differ from E2500

As noted in previous sections, the data relating to traffic signal-controlled location and the logic of their content is time sensitive. For MAPEM and topography changes may occur in periods measured in months and years, so the confidence in data accuracy decays on this timescale. Controller configurations and handset changes can happen quite often, particularly in the early days of a new site being installed so for these data so the confidence in the accuracy of E2500 decays on a week to month scale. This does not recognise the last level of data, the immediate state of the controller's inputs, outputs and internal processes. Controllers typically work on a 100ms tick, although many processes may happen a lot faster, and accordingly there is a set of data that changes at these timescales. Many systems that interact with signal controllers, particularly control systems (UTC, MOVA, etc.) rely on and require access to data in this timescale. While some systems accommodate timestamped data this is a

compensation for limits in transmission speed, data loss, latency etc in the information chain. DCIS is working to address the availability of immediate controller status through the provision of DCAG (dynamic controller automated generation) data in a standard form. This form will follow the data structures of the E2500 with current input and output states and internal processes being associated with the corresponding static identifiers and processes in the E2500 JSON. By way of example, where the JSON classifies all the inputs to the controller, the processes applied to them and the destinations of that data, the DCAG will also report the values of the data at those same points on a moment-by-moment basis. Using the static structure of the site's configuration allows machine consumption of the dynamic content to know in advance where to look to find what it needs.

Setting a DCAG standard provides a development target for signal controller developers that allows publication of this real-time content with a scope and completeness sufficient for diverse users.

DCAG does however introduce a whole new set of data security and validation challenges. The DCIS team recognise this challenge.

4.4. How will E2500s and DCAGs be generated, validated, stored, curated & accessed

It is anticipated that instances of these two file types will have very different routes to creation. This is expanded on below.

4.4.1. Validating E2500

All the information needed to control a particular site is held in the configuration tools set and files associated with the site. Pragmatically the best point at which to generate an E2500 file is simultaneously with the generation of the configuration file for the controller. In that way there is the least possibility of any divergence or discrepancy in the content. This is also a good place to validate the file, before it gets into wider circulation. Prudently the E2500 schema and semantic validation checks will be applied at the point of generation and at the point of uploading to the curated site and again on downloading from the site by users. Unit tests were outlined above. They assess self-consistency; for example, they verify information where content from one location in the E2500 (e.g. number N of phases) can be used to verify content in another part: (e.g. are all N phases represented in the conflict table).

A representation of the schema derived from data for a typical medium size site is presented graphically in Figure 2

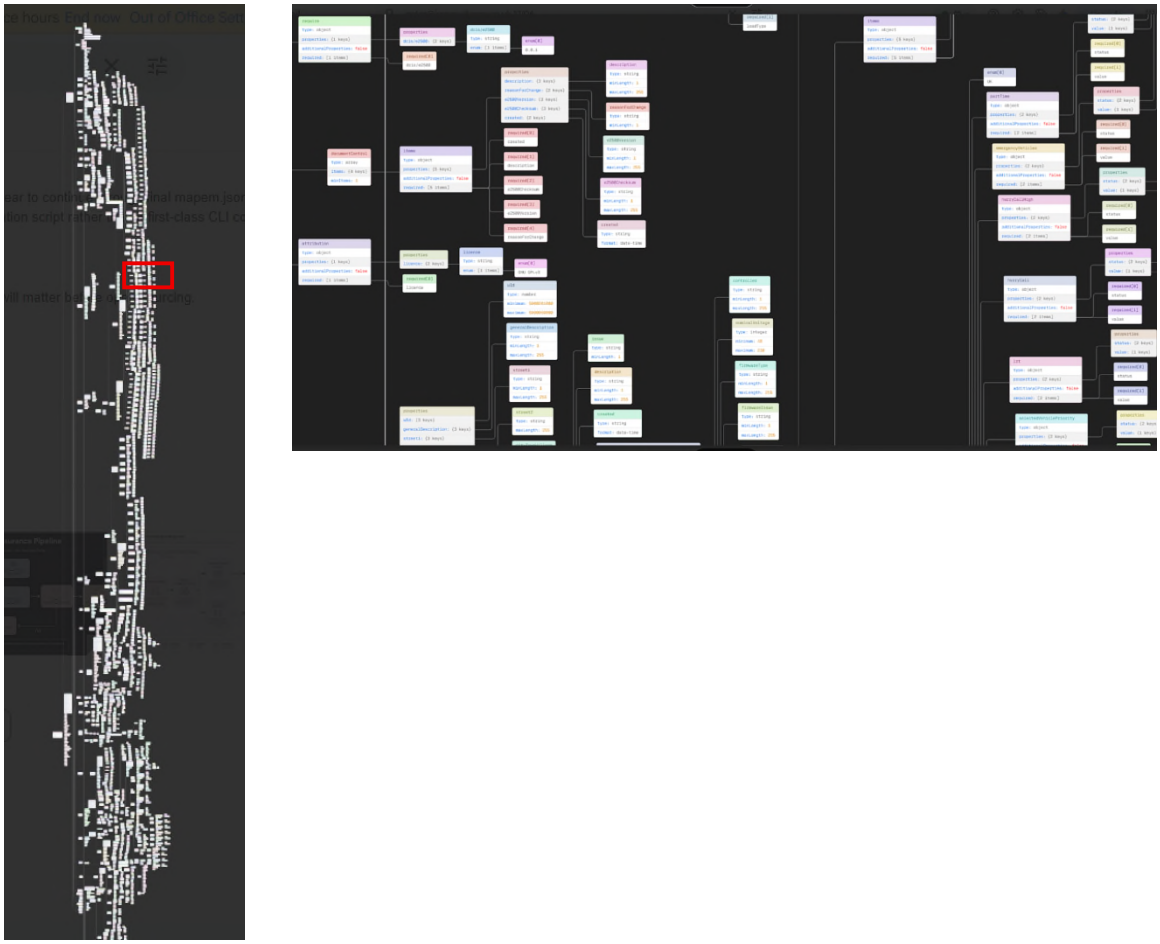


Figure 3 The E2500 JSON schema: A graphic representation (l) and detail of a section (r)

Configurations, and therefore E2500s are static files. While they are likely to change more often than MAPEM files they still change relatively infrequently. A subjective assessment of their validity can be made from their age, with hours and days being more trustworthy than months or years.

4.4.2. DCAG

DCAG content is intended to represent the instantaneous state of the controller and as such can only meaningfully be generated by the controller or a surrogate of it. There is no intention to centrally store or curate this content. It is instead intended for immediate consumption or local aggregation either at or close to the traffic signal cabinet. Because of the structural relationship between the E2500 and the DCAG it is likely that a subset of the E2500 generation tools may provide the basis for a controller-located publishing service. On-the-fly validation of DCAG may present challenges, so reliance on validation of the E2500, and off-line checking against a DCAG schema may go some way to alleviating processing overheads in publishing DCAG while retaining sufficient confidence in the veracity of the content.

4.4.3. Upload, storage, curation & access

It is anticipated that information on the E2500 standard will be aligned with the TOPAS management of TOPAS 2500, much as the MAPEM standard is owned by C-Roads.

Management of the individual E2500 files is a separate activity. These are likely to be transferred to a publication service alongside the DTROs mentioned above and other road network data.

While the DCIS team have a structure, work on this process also relies on other parties and will be reported when it is better defined.

5. A User's Perspective

Understandably users will ask about the benefits to them and obligations that will be placed on them. The next section looks at the benefits, and subsequent sections address what may be required of users.

5.1. Use Cases

An early activity in the DCIS project was an assessment of the range of use cases for the DCIS content.

Many potential uses were identified. A simplified summary of a selection is presented below.

5.1.1. Stage Movements and Intergreen Timings

Two use cases have been identified related to stage changes. Both the choice of next stage and the interval during stage-to-stage movements can vary according to influences on the controller e.g. the presence or otherwise of pedestrian demands leading to moves to stages to serve pedestrians or moves to stages that serve vehicle demands. Typically, currently a user has to wait for the stage move, or at best calculate the range of intergreens. The availability of data from the controller towards the end of the current stage can allow accurate forward prediction of the next move the controller will make. This predictability is further nuanced by the reporting of red signal faults within DCAG and any extensions that these cause.

5.1.2. Modelling

In the context of traffic control and network management, models represent the vehicles and traffic, they represent the road layout, but often do not have access to the particular control strategies being applied at signalised locations in that model. Some signal companies publish content that allows interfacing, but currently this is understandably bespoke. Widespread availability of MAPEM and E2500 will support faster and more accurate generation and updating of network models, and it will better support faster than real time (FTRT) control tools for network optimisation. The addition

of DCAG to FTRT tools has the potential to further reduce divergence between the model and reality.

5.1.3. Junction and Network Optimisers

As alluded to in section 4.3 above, the availability of live data from the roadside with levels of content that are substantially greater than content reported by current system-wide network tools (e.g. the bits transmitted in the UG405 data set) opens up opportunities for better understanding of the network and thereby better tuning of performance, regardless of the tools chosen. In effect DCAG in particular brings together all the data that the controller is using or even can use into one consistently structured place.

5.1.4. Autonomous Devices

As noted at the outset of this paper, autonomous devices using the highway network need the best available information about the current and future state of the network and the assets on it. At one end of the spectrum signalised crossings are a real barrier for delivery robots. At the other, coordination of autonomous vehicles and platoon movements with timings to pass platoons through signals offers a way to optimise the traffic to meet the available time, i.e. the saturation flow values used in optimisers and split calculations could be markedly higher.

5.2. User Resources and Commitments

Beyond these use cases are the understandable questions relating to the impact of DCIS on local highway authorities, devolved governments and other network operators. How will implementation of DCIS be resourced and how can the benefits be accessed. As noted above it is likely to require general funding for network operations to include specific requirements to bring DCIS into being and manage its orderly transition into a business-as-usual process.

Figure 4 shows how the responsibilities for generation of E2500 will fall across those parties currently owning and operating the road network. (The process for MAPEM differs mainly in the labels on the boxes, the sequence of actions is substantially the same, with, for example the as-built CAD and the MAPEM being generated outputs in the blue box). Local Highway Authorities are the custodians and owners of the source data but often delegate much of the operational activities to contractors. It is plausible to envisage that specialist contractors will implement software tools with which to service the needs of their local highway authority clients.

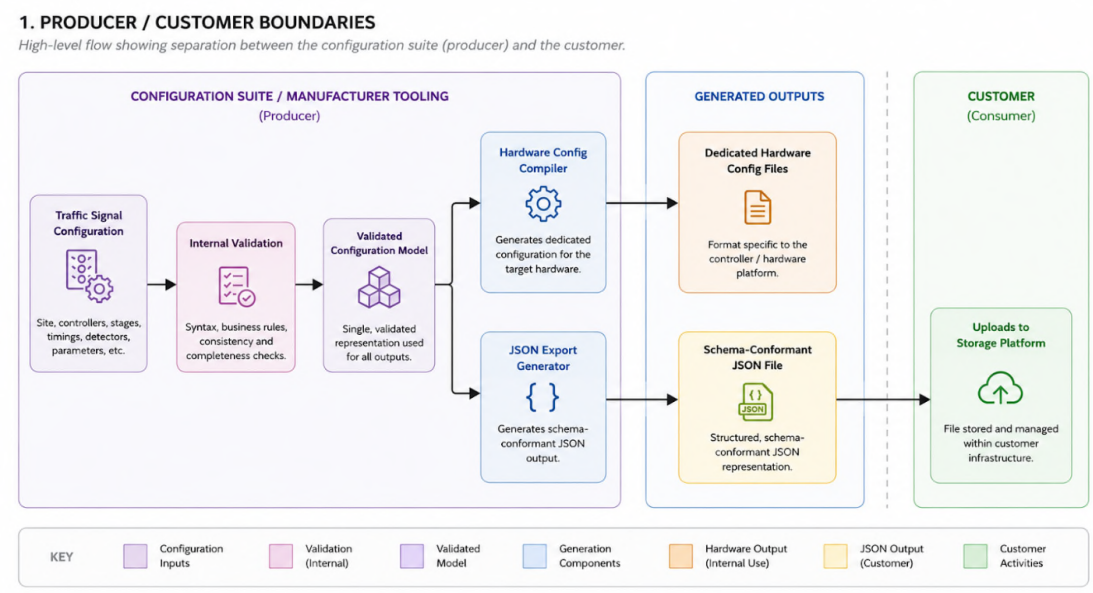


Figure 4 The operational processes for E2500 creation

This approach is deemed to be minimally disruptive to existing contractual relationships so most likely to operate smoothly.

Guidance and support will be provided for all parties in the chain from roadside to stored and selectively published data.

6. Summary

This is a report on the current position of the Digital Controller Interface Standard (DCIS) project and as such focuses on the prevailing areas of work.

The sections above have presented the reasons for, context and current position of the project. The long-term objective is to support autonomous vehicles on the UK's road networks. In the near-term standardised data allows many network services and operations to be rationalised and glean greater value. The project has four deliverable elements, a structured list of signalised sites (the ISD) falling within the boundaries of each local highway authority (this file also serves as a primary index of sites against which the following files are aligned). A MAPEM file conforming to C-Roads standards which defines the topography of each signalised site, lanes in and out, permitted movements, signals in control of each traffic (and pedestrian) movement, etc. The specific control processes are covered by the E2500 file, a standardised form of the site parameters from the controller configuration. The last element of the work is the DCAG file, which resembles elements of the E2500 in structure but can be used to report the real-time status of variables from the controller to other applications such as junction or network optimisation tools.

Processes for the generation, validation, storage and publication of the DCIS content have been described.