An Information Model for Geospatial and Temporal References

November 2011

Author(s): Chris Partridge

QINETIQ

An Information Model for Geospatial and Temporal References

Aim

To review the current practices and artefacts used for the information modelling of the Geospatial and Temporal References (G&TR) in the surface ship community and investigate the ways in which the use of ontologies can help to improve them.

Summary

Within the surface ship community, there is a significant amount of content describing G&TR; Def Stan 22-61 is the prime example. However, from an information modelling perspective, much of this is unstructured. And where it is structured, it reflects a data implementation rather than an information modelling perspective; in MDA terms, a PSM rather than a CIM perspective. It aims at describing how G&TR data should be used, but has not been so directly aimed at describing what G&TR is – the target of an information model. This has led to a situation where a solid foundation for the information model is missing and there is no clear articulation of the fundamental components for the information model.

The analysis below provides a sketch that can be developed into an information model that would provide the foundation for the 'how' model in Def Stan 21-66. The combined models would provide a better, more accurate, overall model.

The brief review has identified a number of areas where an ontological approach would help:

- It could make explicit a simple, common semantic foundation for G&TR terms that would form the basis for a common understanding at the right level of precision; reducing significantly the problems that arise from misunderstanding.
- It could provide a G&TR information model based upon general patterns that would simplify the architecting process, in particular helping to separate the design and implementation concerns; leading to simpler, improved architectures.
- It could provide a systematic, objective test of the semantic quality of the current G&TR standards; identifying actual and potential sources of error and suggesting ways of resolving them.

Background

This is part of a larger project to move towards information models for surface ship combat systems. This draws extensively on previous QinetiQ work, including the JTADIS project.

Substantial work has been done in the G&TR area as it is a vital part of the combat system. However, this has addressed issues of how G&TR is to be used and little effort has been expended on what G&TR is. This has been in part because there have been few tools to investigate this but also because the benefits of doing have not been investigated and articulated – something we try to do here.

Approach

The prime source for this analysis was Def Stan 21-66. Other supplementary sources are listed in Appendix A). A key source of prior ontological G&TR models is the JTADIS project.

The analysis will use the BORO approach, which has a number of key features; it:

- Enables the extraction of the ontology from existing assets. This reuse of existing assets not only harvests the original investment them, reducing the analysis effort and improving the quality of the results but helps to ensure consistency and interoperability with the existing assets.
- Identifies the objects represented by the model. This provides the basis for semantic quality assurance (QA). Where the object identified by the model cannot be unambiguously identified there is an issue. This provides an objective, systematic process for semantic QA.
- Identifies the dependences between the objects represented. This helps makes clear what needs to be (and does not need to be) included in the information model. For example, if a coordinate frame is dependent upon a reference frame, where the model contains a coordinate frame its foundation reference frame should be included as well.
- Identifies the underlying general patterns in a domain. This enables the domain to be understood and modelled more easily – as fewer simpler patterns are involved.

Overview of the Task

This is intended to be a very small task – it has been allocated 1.5 days. Within this timeframe we have produced a broad brush sketch of an information model for G&TR to provide an impression of what it would look like and also the kind of benefits that would arise. This is not an information model, but it can provide the foundation for one should it need to be developed at a later stage.

Key Aspects of the Analysis

The analysis identified several key aspects, which all reflect the underlying requirement for a clearer picture of the G&TR domain.

The analysis has started to identify the common G&TR components that get re-used in multiple combinations in different composite structures. This illustrates how this kind of 'conceptual' re-use can aid a common understanding as well as simplify the models.

The analysis has identified a dependency structure; the top layer for position coordinates is shown graphically below. Making this structure explicit helps simplify the modelling process, as it informs the modeller where he or she needs to incorporate elements in the model. For example, when modelling position coordinates, coordinate systems need to have worldline reference frames – as they are dependent upon them.

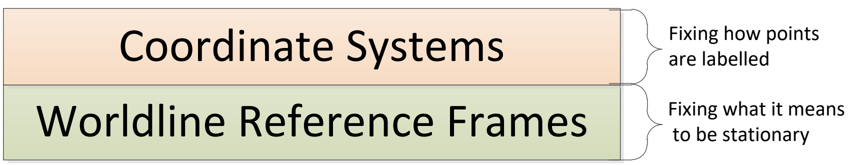

Figure 1 - Position Coordinate Dependencies

The analysis has provided a way to treat 'change over time' though the use of a 4D ontology. This enables a consistent treatment of 'change over time'. It enables one to answer exactly questions such as whether 52° N is the same on the 1^{st} and 2^{nd} December.

The analysis has provided a number of illustrations of how labels should be separated from the things they label. This is particularly acute in G&TR where the numerical label is often confused with the quantity that it is labelling. So, for example, 52° N is clearly a label, the analysis clearly identifies both the label and what it labels.

The analysis has provided a number of illustrations of the need to be clear about identity. One example in the text is whether the identity of reference frames should be determined by the physical datums that are used to identify them or by the worldlines that characterise them.

Analysis

This analysis is intended to provide a sketch of a G&TR Information Model by unearthing objects and the patterns sufficient to characterise the G&TR domain.

Structure of Spacetime

There is a need to decide upon a structure for spacetime. For our purposes the right trade-off between accuracy and simplicity is Galilean spacetime[1].

Stationary Reference Frame

The foundation for the framework is a notion of a stationary reference frame. This is needed to differentiate between two objects that are at rest with respect to one another from two objects that are moving with respect to each other. These are typically associated with a physical object (such as a platform or the Earth).

This notion needs to be translated into an object. The first stage of this is to identify worldline objects. We select a physical object. At a point in time, the space that the physical object occupies can be considered to be composed of points. We assume that it is rigid; in other words, it does not change shape (more specifically, the distances between identified points on the object are unchanged) over time. Over time, each notional point in the object traces out a line in 4D spacetime – what is conventionally called a worldline. Another way of looking at this is that the notional points are parts of the worldline – and that the worldline is the (mereological) sum of these parts.

The worldlines are ultimately defined by and depend upon physical objects. Without them as a reference, it is impossible to determine how the worldline unfolds in time. As physical objects can come into and go out of existence (this is more common with platforms than planets, such as the Earth), the worldlines that depend upon them also come into and go out of existence over time.

Worldline Reference Frames

The physical object's worldlines can be used to identify the object that exactly corresponds to the notion of being at rest (or stationary). One can collect all a physical object's (point's) worldlines. As the object is (assumed to be) rigid, so over time not only do these worldlines not overlap, but they maintain a consistent relative distance between each other. Provided the object is extended in the three spatial dimensions, this is enough to fix an extended set of worldlines. This is done by picking out a point outside the object and identifying the unique worldline that meets the criterion of maintaining relative distance with the initial worldlines at each of the timeslices. There is no obvious (non-arbitrary) place to stop extending, so we continue until we include all of space. We propose

[1] A description of the mathematic structure of Galilean spacetime and the reasons for choosing it are given in Appendix C.

considering this set of worldlines a reference frame. For neatness's sake, we assume that the worldlines in a reference frame set all start and stop at the same time – in the limit case, this will be at plus and minus infinity.

Timeslices

It is important to have the notion of same time – to know when two (instantaneous) events occur at the same time and when they occur at different times. The object that corresponds to this notion is a timeslice. Galilean space-time has an absolute notion of simultaneity and using this one can define a unique timeslice – corresponding to each 'time'. We start with a 4D point and pick out all the other points that are at the same time as it; this gives us a 3D space. And if we were to repeat the procedure for any point in the 3D space, we would get the same 3D space. This 3D space (at a point in time) is a timeslice. If two events are in the same timeslice, they happen at the same time; if they are not in the same timeslice, they happen at different times.

Worldlines have the interesting property that where there is an intersection between a timeslice and the worldline, the intersection is a point: in other words, the worldline only intersects the timeslice once. Another way of looking at this is that for every point along a worldline, there is a unique timeslice.

Reference Frame Datum

The physical markings that are used to identify a reference frame are called a datum. Typically, the datum is also used to characterise a co-ordinate frame – a topic covered later in the paper.

The simplest way to do this is by marking three non-collinear points on a physical object with fixed separations from one another. These three points (for their lifetime) trace out three worldlines, which are sufficient to characterise a reference frame. This kind of datum is a set of three worldlines which are a sub-set of the reference frame (which is also a set of worldlines).

Each datum uniquely identifies a reference frame. But a reference frame can have a number (potentially infinite) of datums – as any three non-collinear worldlines within the overall reference frame set can constitute a datum for that reference frame.

In practice, Royal Navy surface ship engineers conventionally mark two planes (rather than points) on the platform and used these to identify the reference and co-ordinate frames – the exact mechanism for this is discussed later.

Reference Frames and Worldlines

The same worldline can be in a number of reference frames, as a single worldline by itself does not determine the object's angular velocity – whereas the reference frame (set) does. Here is an example that illustrates this. Consider the centre of a globe spinning on the Earth's surface and take its worldline. This is stationary relative both to the globe and the Earth's surface. Now consider a reference frame based upon the spinning globe and another based upon the Earth upon which it is spinning. These share the globe centre worldline, but different worldlines at all other points. This illustrates how a reference frames (a set of worldlines) fix the angular velocity.

Plainly any two objects whose relative distance changes over time (and so exist at the same time) will give rise to different worldlines and so reference frames. Indeed, some of the worldlines in one reference frame will intersect some of the worldlines in another, breaching the relative distance criterion for worldlines within a reference frame.

Celestial Reference Frames and Worldlines

Celestial objects, such as the Earth, are often used as reference frames. However, finding an accurate datum to identify the reference frame is a difficult technical problem. This has been well studied in relation to the Earth and for most maritime systems an idealised ellipsoid is used as the notional datum for the reference frame – the most commonly used ellipsoid is the WGS84 reference ellipsoid. These techniques are, in principle, generalizable to other celestial objects.

Non-Rigid Datum Object Worldline Reference Frames

There are cases where the identification of the worldline reference frame depends upon objects which are not rigidly separated. An example is a reference frame for a platform that is stabilised to the earth – this involves two objects which are not rigidly separated; the platform and the earth. The result though is the same, a reference frame set of worldlines.

In this case, one selects a single worldline from the base object (in the case of a platform, typically its Common Reference Point) and then aligns the other worldlines using the stabilising reference frame. In effect, this eliminates any angular velocity that the base object may have relative to the stabilising reference frame, leaving the linear velocity. This can be seen as constructing a notional object that tracks the movements of the platform (or, at least, its Common Reference Point) but is aligned to be stable relative to the (spinning) Earth.

A Taxonomy of Worldline Reference Frames

It may be helpful to illustrate the types of reference frame in a taxonomy – as done graphically below.

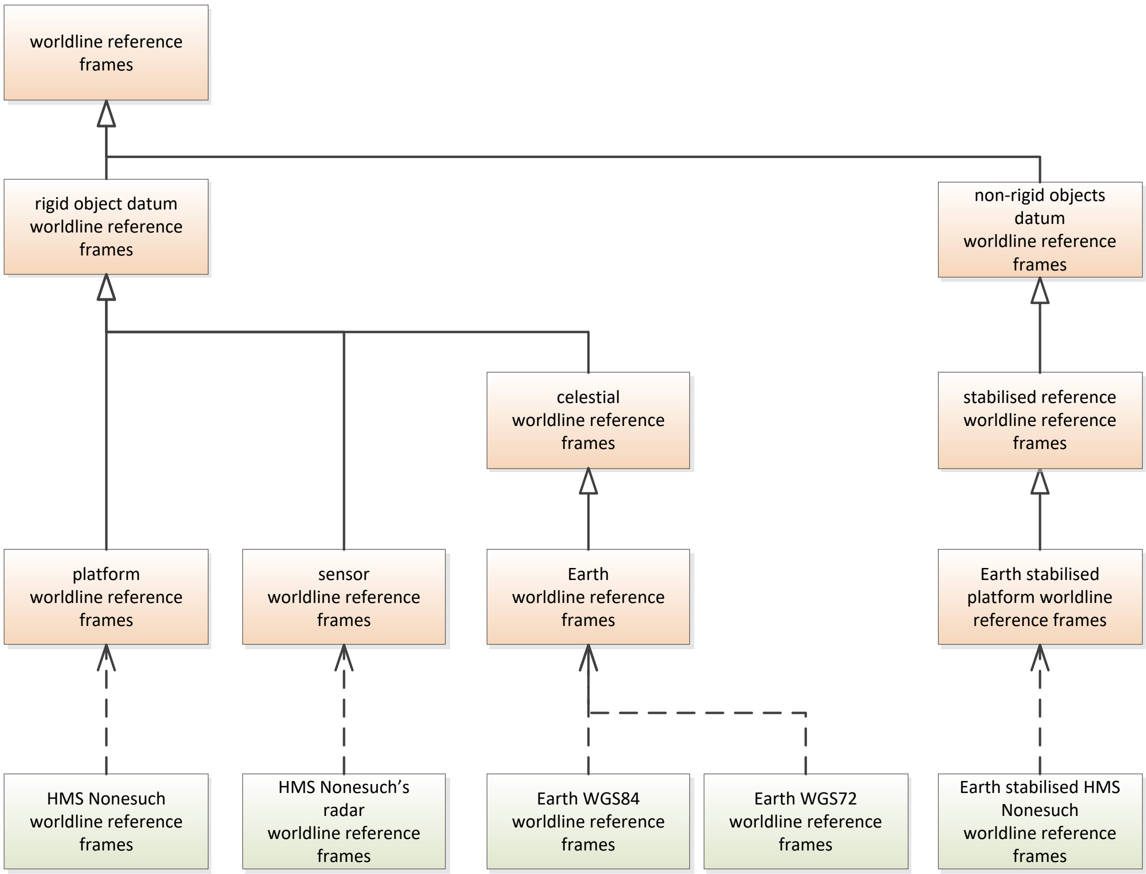

Figure 2 - Taxonomy of worldline reference frames

This taxonomy makes the simplifying assumption (also made in Def Stan 21-66) that there is a single reference frame for an object – rather than (maybe) a number over time. The possibility that there may be combat system session reference frames – that come into and go out of existence with the session – needs to be investigated.

Sensor can be the manufactured or the installed unit. Where 'Sensor' is a manufactured (serial numbered) unit, which can be de-installed from a platform and re-installed on another platforms, then it has a different worldline reference frame from the platform. Where 'Sensor' refers to the installed component on the platform, then the Sensor will typically be aligned to the platform, so that, for example, a radar antenna rotates about the platform vertical and not at a tilt, and it is mounted pointing 'forwards' so that it can measure bearings correctly. However this alignment does not guarantee that the Sensor and the Platform share the same reference frame[2]. Platforms bend with the force of the waves, and differential heating causing different parts of the platform to expand and contract.

Labelling Positions Using a Coordinate System

A coordinate system can be divided into a number of parts, which have various dependencies. This is shown graphically below. Typically there are a number of different options for each of the parts.

[2] See Def Stan 21-66, Section 6.5.1.

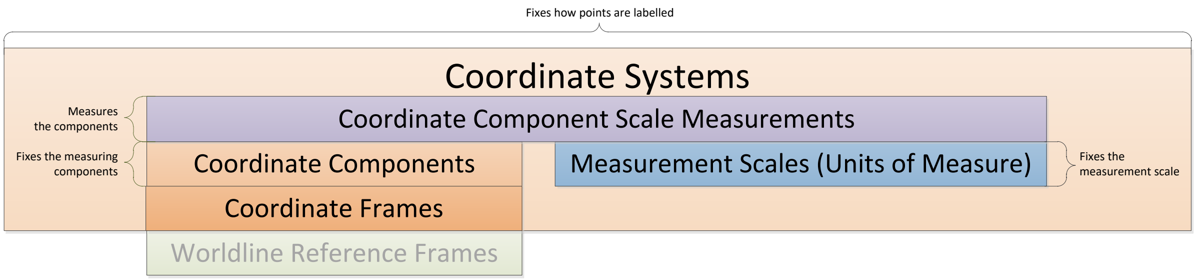

Figure 3 - Coordinate system parts and dependencies

The Situation

The origin-position displacement

There is a position, a (4D) point. This can be identified using a displacement (a line segment) from a known point (the origin) to the position point at a point in time (in other words, both points intersect the same timeslice).

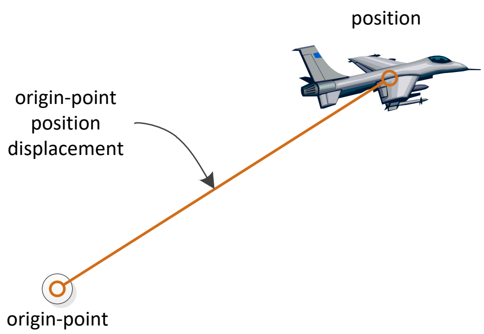

Figure 4 - Origin-point position displacement

Note that the bare displacement is a line segment that is only dependent upon the origin and the position end points (these are sufficient to identify the straight line joining them); it does not depend upon a reference frame.

The position coordinate system

A coordinate system is a mechanism for uniquely labelling positions using algorithms based upon displacements from an origin – typically with a numerical label that can be used in calculations. It aims to decompose the position into components that can be given a single numerical (or almost numerical[3]) label – and then constructs a composite label as an ordered list of the numerical labels.

All the coordinate systems we are interested in separate the temporal and spatial components of the position. They take advantage of the fact that within a reference frame each point can be uniquely identified as the intersection of a timeslice and a specific worldline by providing separate mechanisms for labelling the timeslices and the reference frame worldlines and combining the two labels into a composite label for their intersection. The timelines can be easily labelled numerically,

[3] E.g. 52°N, which has a numerical component tagged with an alphabetic code component.

but the spatial elements need to be broken down further to get components that can also be easily labelled numerically.

The use of worldlines (within a reference frame) provides the basis for regarding points on the same worldline at different times as being, in some sense, the 'same' point. And this is reflected in the use of the same component label in the composite label for the various points. So, for example, the two position labels (52°N, 21°W, 100m, t_1) and (52°N, 21°W, 100m, t_2) have different temporal component labels (t_1 and t_2), but share the spatial component label (52°N, 21°W, 100m) which refers to a worldline in the associated reference frame. The two position labels refer to points on this worldline – in that sense they refer to the 'same' point at a different time.

Another perhaps more unfamiliar but better way of looking at this, is to start with the worldline rather than the point. From this perspective, one component of the coordinate system mechanism uniquely labels each worldline within its reference frame. Another component labels each timeslice. Using these, one can uniquely label a point with a composite label containing the two component labels.

The position coordinate frame

A position coordinate system uses a co-ordinate frame to set the orientation of the spatial and temporal dimensions. It is the framework within which the coordinate components are situated so that they can uniquely identify each worldline within its reference frame.

The position coordinate components

The position coordinate components uniquely characterise worldlines relative to a reference frame. This, in association with the characterisation of a timeslice, identifies a point-position.

The construction of these position coordinate components starts as a 3D structure for identifying points within a timeslice relative to a reference frame. The 3D structure is then extended into a 4D structure by taking its worldlines. Effectively this creates a 4D structure for identifying worldlines within a reference frame, prior to labelling them.

This provides a structure for uniquely identifying a point within each time (timeslice) which is, relative to the reference frame, invariant over time. Hence the same structure applied at different times will identify different points on the same worldline.

Basic Position Coordinate Frames

A coordinate frames fixes the origin, temporal and spatial dimensions. It consists of a temporal component upon which a spatial component is built. A coordinate frame with only the temporal component is called an origined coordinate frame. One with both temporal and spatial components is called a position coordinate frame.

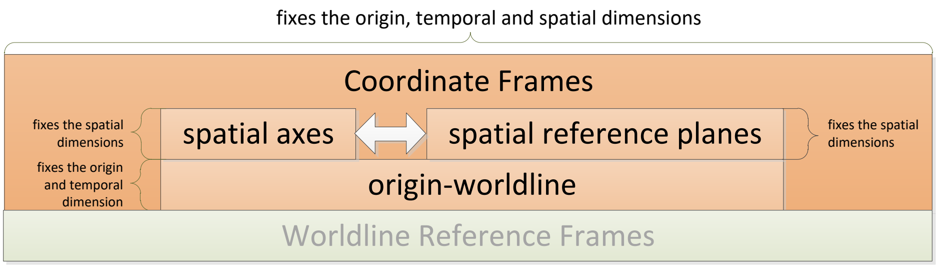

Figure 5 – Coordinate frame components and dependencies

The origin-worldline and origined coordinate frame

Given a reference frame, one can identify a coordinate frame with just a temporal component; with no spatial component. All this involves is identifying one of the worldlines in the reference frame as the origin-worldline; then this (the reference frame and the selected worldline) are an origined coordinate frame. This becomes, in effect, the time axis.

Given a reference frame, any worldline in the reference frame's set is a candidate origin-worldline – so a reference frame can have an infinite number of origined coordinate frames based upon it. However, it is traditional to pick some kind of central point in the physical object that is the basis of the reference frame. Selecting the origin-worldline can be regarded as selecting the time axis of the coordinate frame.

These are typically specific points in a physical object. Examples are the Common Reference Point (CRP) in a platform and the notional centre of a celestial object, such as the Earth.

Where different origined coordinate frames share the same underlying reference frame, the two origin-worldlines are at rest relative to one another. As noted above, a worldline can belong to a number of different reference frames. Hence, a worldline can be the origin-worldline for a number of different reference frames. This reinforces the notion that worldlines are, by themselves, insufficient to determine rest; as they cannot discriminate angular velocity.

As we will discuss below, an origined coordinate frame is sufficient to describe speeds and velocities.

The spatial axes and planes and position coordinate frame

The general strategy here is to identify a 3D spatial frame in a particular timeslice and then (using the reference frame's worldlines) extend it across timeslices along the origin-worldline into a spatio-temporal frame. The spatial frame picks out the spatial dimensions using both axes and planes, where there is an interdependency between these such that one uniquely determines the other. It simplifies things if the selected axes (or planes) are orthogonal to one another, though strictly the only requirement is that none of the axes (or none of the planes) are parallel to any other axis (plane).

One can identify the origin-worldline explicitly in advance of the spatial frame, or use the intersection of the spatial axes (or planes) to identify an origin-point that the associated reference frame associates with a worldline.

Assuming one has identified a reference frame and one of its worldlines as the origin-worldline. To construct a coordinate frame, one then takes a timeslice that intersects the origin worldline and hence the point where they intersect, which is the origin-point in that timeslice. One then selects three orthogonal straight lines through the origin in the 3D timeslice to act as axes –labelled, say, these x, y and z. From these one can identify the associated coordinate planes as two lines uniquely identify a plane: so we have the x-y-plane, the x-z-plane and the y-z plane. Each of these coordinate planes is orthogonal to the axis that is not used to identify it; for example, the x-y-plane is orthogonal to the z axis. In this case, we say the x-y-plane is the z axis's orthogonal coordinate plane.

Once these 3D objects are in place we can extend them by taking the worldlines for each of the points they contain as parts. The 3D lines extend to 4D worldsheets; maybe better named as worldsheet-axes. The 3D planes extend to 4D plane world-volumes or world-hyper-planes.

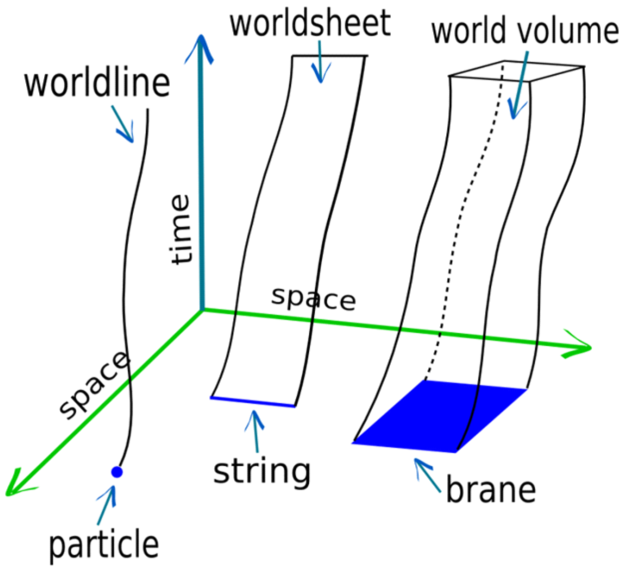

Figure 6 - Moving from 3D to 4D

We can also reverse this and take the intersection of a specific timeline and position coordinate frame to get a 3D spatial frame for the timeslice. The 3D spatial frame is independent of the reference frame, and by association with different reference frames can be extended to identify different position coordinate frames. This turns out to be an important feature for platform reference frames, which we explain later.

For simplicity's sake, we make no distinction between the axes and places and consider the position coordinate frame to be the set of both the world-axes and world-planes – though each by itself (or suitable combinations) is sufficient to identify the frame.

The platform frame

In the UK maritime domain, there are two platform datum planes (the Master Training Datum (MTD) and the Master Level Datum (MLD)) that are used to identify its coordinate frame. The origin, the Combat System Common Reference Point (CRP), is a physical location along the intersection of the MLD and MTD that is defined as a part of the platform design.

These can be used to identify the axes as shown below (this is a copy of Figure 1 from Def-Stan 21-66).

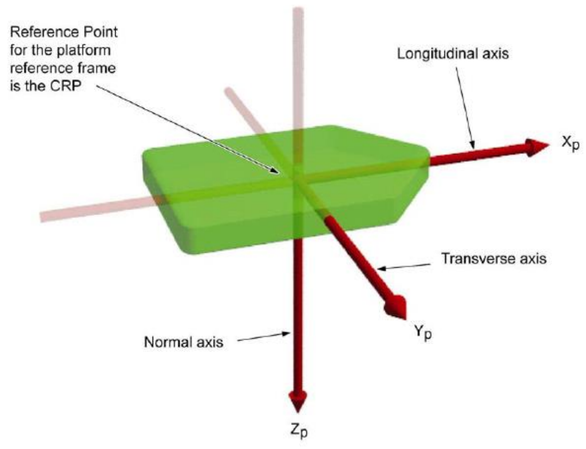

Figure 7 - Platform axes

They can also be used to identify the coordinate planes, which are known as reference planes and individually named; basic, longitudinal and transverse – these are shown graphically below (this is a copy of Figure 2 from Def-Stan 21-66).

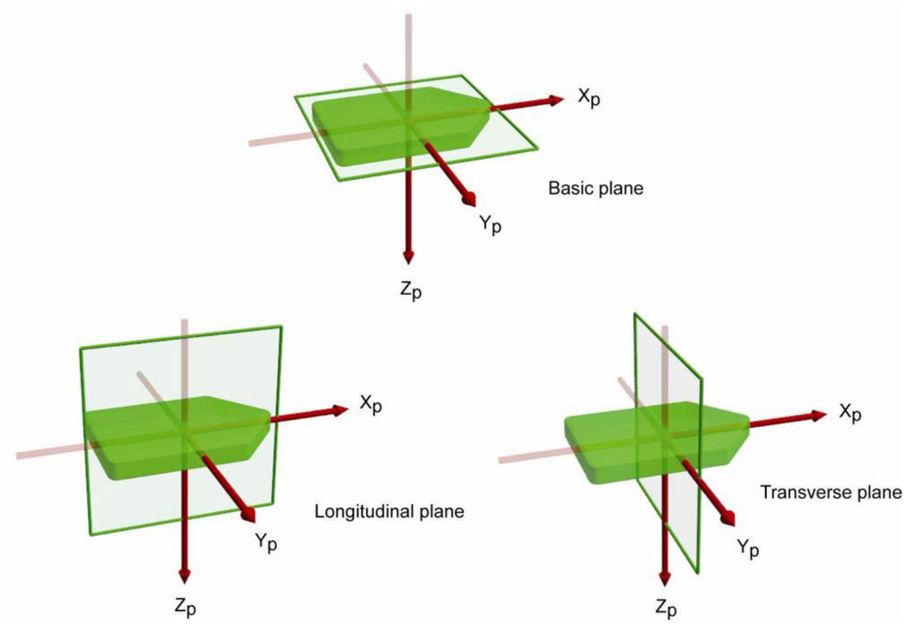

Figure 8 - Platform reference planes

The celestiodetic coordinate frame extension

Where measurements are made in relation to the Earth's surface, a celestiodetic coordinate frame extension – based upon a celestiodetic reference frame - can be useful. This marries a coordinate frame with to the idealised geometric figure selected as an approximation to the shape of the celestial body. In the case of the Earth, the idealised figure is typically an ellipsoid (the WGS84 standard defines it in this way) – and the extension called a geodetic coordinate frame extension.

It is normal to take an idealised axis of the Earth's spin as one of the axes and the origin-worldline as the notional centre of mass of the Earth. These datums are sufficient to align the idealised figure with a coordinate frame.

Position Coordinate Components

The origin-point position displacement can be decomposed into components that uniquely identify it in a number of ways. The two most common decomposition techniques (which can be combined) are Cartesian linear decomposition and polar angular decomposition. The Cartesian displacements technique decomposes the displacement into component (Cartesian) displacements[4] along the three axes. The polar angular technique provides angular geometric objects that capture the direction from the origin in which the displacement is situated, which when combined with the length of the original displacement identifies the position.

Another common decomposition technique is the celestiodetic (in maritime domains this is normally applied to the earth, as is known as geodetic). This is based upon an idealised geometric figure that approximates to the shape of the celestial body.

Cartesian linear decomposition

For any position in a coordinate frame, it's displacement from the origin-point uniquely projects onto each of the three axes giving three orthogonal displacements. The original displacement can be regarded as the sum of three component orthogonal displacements.

Geometrically, each component axis displacement projection is structured as follows. There is a unique plane that is orthogonal to the axis and intersects the position. This plane also intersects the axis at a unique point – the projection point. The line segment from the origin for the projection point is the component displacement. In the limiting case, the projection point will be the origin and there will be no displacement.

The identification of the position from the component displacements is essentially the reverse of this process. For each axis take the displacement (which may be zero) and take the orthogonal plane at that point. This will result in three orthogonal planes which will intersect at a point – the position.

The axis projections are extended into 4D by the usual procedure of taking the worldlines of all their component points. These 4D objects are then labelled in the coordinate system.

Polar angular decomposition

In a polar angular decomposition, the original displacement is broken down into two components; a radial component which characterises the length of the displacement and an angular component which characterise the direction of the displacement.

The radial component is the unique sphere centred on the origin which intersects the position. The angular component is characterised using one of the axes – the polar axis and its orthogonal plane, the equatorial plane. There are two sub-components associated with this; the elevation angular component and the azimuth angular component.

[4] They are not yet Cartesian coordinates as they have not yet been assigned a numerical label.

The elevation angular component is the unique (right circular) cone that has its apex at the origin, its axis as the polar axis and intersects the position. This cone can be consider composed of lines though the apex all at the same angle either to the polar axis or the equatorial plane. In the two limiting cases, the cone collapses into a half-line along the polar axis.

The azimuth angular component is the unique half-plane whose edge is the polar axis and intersects the position. Together these uniquely identify the position. In either of the two elevation angular limiting cases, there is no azimuth angular component.

There is an illustration below of the spherical coordinate components for a position marked with a black sphere. The vertical axis is the polar axis. The red sphere is the radial component. The blue cone is the elevation angular component. The yellow half-plane is the azimuth angular component.

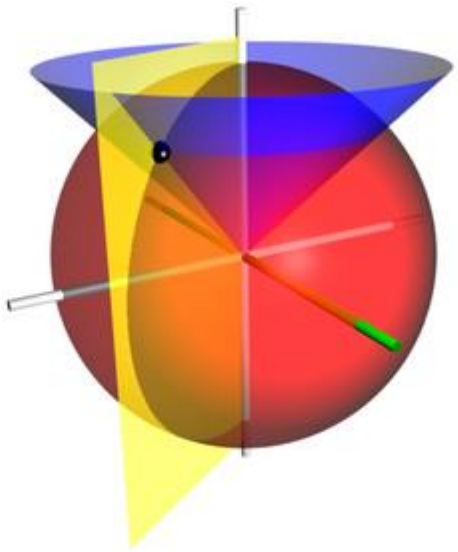

Figure 9 - Spherical Coordinate Components

The components are extended into 4D by the usual procedure of taking the worldlines of all their component points. These 4D objects are then labelled in the coordinate system.

Celestiodetic decomposition

This is similar to the polar angular decomposition, with the difference that the angular components are defined relative the extended celestiodetic frame.

As in the polar angular decomposition, there is a polar axis and equatorial plane. In principle, any axis could be the polar axis but in practice it is usually the (idealised) axis of rotation of the celestial body. The azimuth angular component (known as longitude) is defined in a similar way – it is the unique half-plane whose edge is the polar axis and intersects the position.

The elevation angular component (known as latitude) is defined in a different way. Instead of the cone being constructed with its apex at the origin, it is constructed so that at its intersection with the idealised figure it is at a normal to the surface. In the case of the WGS84 Ellipsoid with the standard polar axis, the result is a cone whose apex is then still on the polar axis but the other side of the Equatorial plane. To ensure uniqueness, the section of the cone on the other side of the equatorial plane is ignored.

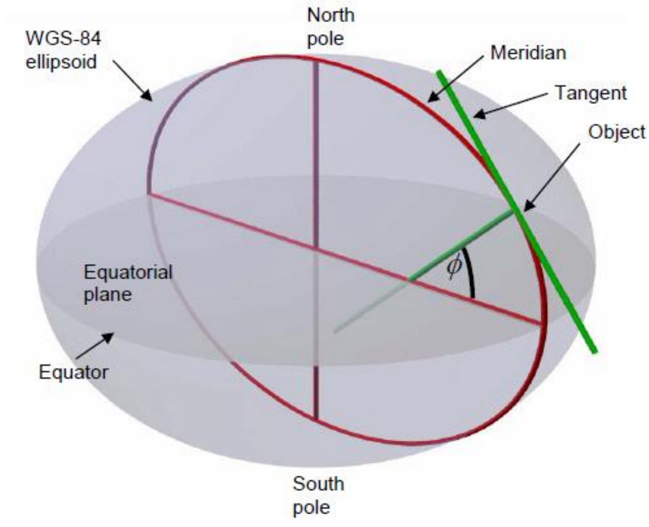

Figure 10 - Longitude construction

The radial component is an extruded version of the idealised figure, which maintains a constant distance to the original figure. In the case of WGS84, this will be an extruded ellipsoid.

Each point can be uniquely identified by a combination of one of each of these components.

The components are extended into 4D by the usual procedure of taking the worldlines of all their component points. These 4D objects are then labelled in the coordinate system.

Celestiodetic-dependent coordinate frame

Where measurements are typically made in relation to the Earth's surface, a celestiodetic-based coordinate frame extension can be useful. In the maritime domain, there are two coordinate frame that are dependent upon the celestiodetic-based coordinate frame extension; the platform stabilised coordinate frame and the celestiodetic local coordinate frame. In 3D these frames are constructed in the same way, what differs is the reference frame used to extended them into 4D.

What makes these two frames interesting is within each timeslice, the stabilised platform (reference) coordinate frame and the local coordinate frame for the platform frame's origin-point instantaneously align. Hence, at that point in time, every position will have the same coordinate components and so the same label in equivalent coordinate systems – making translation simple. However, as the coordinate frames are based upon different reference frames, the velocities and accelerations of objects in one frame will not be the same as the other.

The 3D celestiodetic-based coordinate frame

Given a celestiodetic coordinate frame extension, a timeslice and a point-position in the timeslice, one can identify the frame whose orientation is based upon its position. We take the point-position as the origin-point. We identify as an axis the line that is normal to the celestiodetic figure and goes through the origin. From this a horizontal plane can be identified – as shown in the figure below.

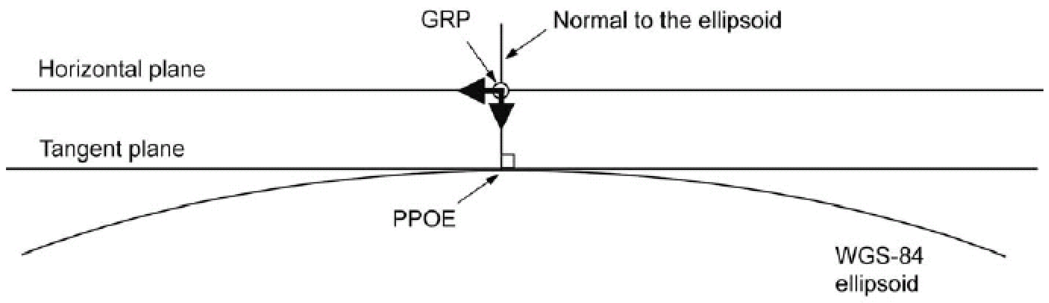

Figure 11 – The horizontal plane

Another axis is the line in the horizontal plane that points 'true north' through the origin[5]. Geometrically, this line runs from the origin to the intersection of the horizontal plane with the polar axis. The third and final axis is the line though the origin that is orthogonal to the other two axes. These axes are sufficient to identify the full set of coordinate planes in a 3D timeslice.

Platform stabilised coordinate frame

This frame marries the platform stabilised reference frame with a celestiodetic coordinate frame. The origin of the frame is the platform CRP. The 3D axes and planes are extended to 4D world-axes and world-planes by the worldlines in the platform stabilised reference frame. This identifies a coordinate frame that always has the platform CRP as its origin.

The celestiodetic local coordinate frame

All the frames we have been looking at so far are single global coordinate frames. However, one can have local frames, where each worldline has its own frame. This local frame has at each point a 3D celestiodetic-based coordinate frame that is extended to 4D using worldlines from the celestial reference frame, and is always a rest relative to the celestial object.

Measurement scales (units of measure)

The coordinate system extends the coordinate frame with algorithms for labelling the components. Typically the labels are chosen to facilitate calculations.

The labelling is typically done through the use of an appropriate measurement scale (often called a Unit of Measure) for the component. There are usually a number of different possible scales for each component.

In the case of the Cartesian axis displacement components the appropriate measurement scale is a length scale. There are a number of these; metres, feet, miles, etc. For simplicity, the same unit of measure is usually applied to each axis – though this is not always the case. For measurement simplicity, the origin is normally given a scale measurement of zero. This means that one half-line needs to have a positive scale and the other a negative scale. It is at this stage of attaching the scale that the axis acquires a direction.

The measurement of the Cartesian displacement is direct. This is not so in all cases. For example, the measurement of the radial sub-component in a polar angular component is based upon the radius of the sphere – this is a proxy measurement for the sphere. In the case of the celestiodetic coordinate

[5] Note that at this stage we do not need to consider direction. The measurement scale adds the direction.

frame it can be the measure of the distance from the original reference ellipsoid to the extruded ellipsoid. In other cases, it can be even more indirect. To avoid ambiguity, in these indirect cases the measurement proxy needs to be specified explicitly.

Coordinate components scale measurements

Applying the measurement scale usually involves some choices. For example, applying the length scale to the Cartesian axes will involve a choice of direction – which side of the origin is to be positive and which negative. This is an important decision for the coordinate system.

Once the application choices have been made, the component to be measured and the result is usually numerical label. For each coordinate system, the composite label for a point-position will be the labels for the coordinate components in an agreed order. For example, a common choice of coordinate label order for geodetic systems is; latitude, longitude, elevation and time.

Comparison with the Def Stan 21-66 standard

A key input to the analysis was the Def Stan 21-66 standard. However, the analysis has resulted in a different set of classifications and general patterns.

A key difference is the meaning of the term of ‘reference frame’. In Def Stan 21-66 this is initially defined as “A reference frame is a fixed relationship between reality and a mathematical representation of it.” This is reflected in the list of reference frames in the standard’s Table 2 - Reference Frames and Coordinate Systems – shown below.

Reference frame	Coordinate systems	
Sensor	Cartesian	Spherical polar
Platform	Cartesian	Spherical polar
Geographic	Cartesian	Spherical polar
Geodetic (WGS-84)	Geodetic coordinates (Latitude, Longitude, Altitude)	

Table 1 - Def Stan 21-66’s Reference Frames and Coordinate Systems

It is subsequently defined thus: “The definition of each reference frame consists of the identification of a number of physical datums, the definition of an initial coordinate system based on measurements from those datums, and then the definition of any additional coordinate systems by means of mathematical transformations.”

In this ontological analysis, one Def Stan 21-66 reference frame may consist of a number of coordinate systems and their associated reference frames; where the coordinate systems are further broken down. This is done to reveal the finer dependence structure and the common general patterns. It is also done to make clear what the coordinate numerical labels are referring to.

This is an example of a general difference in architectural approach between the two structures. Def Stan 21-66 aims to simplify by offering a small number of ‘monolithic’ complete options, with some exceptions. Whereas, the ontological analysis aims to simplify by identifying general re-usable components and how these are composed into complete solutions.

There are smaller differences between the two. There is one example that illustrates the ontological analysis’s concern with identity. The Def Stan 21-66 reference frame is associated with specific “physical datums” – they are part of its identity. Whereas the description of the worldline reference frames recognises that different "physical datums" can identify the same worldline reference frames. If one is interested in capturing the notion of stationary, then differentiating between reference frames that are stationary relative to one another, but have been constructed from different physical datums is not necessary.

Example – Velocity and Speed

A key requirement in this analysis is the discrimination between the G&TR geometric objects that are being measured and the measurements expressed in numerical form, and describing the results in general patterns that are repeated across the analysis. Position is central to G&TR, but there are a significant number of other G&TR geometric objects. We look at two examples, velocity and speed.

Velocity

The velocity of an object is relative to a reference frame; in other words, it depends upon the reference frame. When one measures the velocity of an object, one is (in effect) postulating the trajectory it would have and would follow if it were travelling constantly at that velocity. If one considers the object as a point at its centre of mass, then this trajectory is a 4D worldline, straight in the associated reference frame, with an origin-point at the measurement point. As there is no non-arbitrary place to start or end the line, for simplicity it is considered infinite.

The velocity is relative to the reference frame. This is made clear when one identifies the velocity in a different reference frame. The velocity object is a straight worldline in the alternative reference frame, but it is a different worldline.

The velocity is normally labelled in terms of a measure of the gradient of the velocity line. Typically the measure will identify the ratio between the spatial and temporal components – this remains constant as the velocity (gradient) is constant – relative to the reference frame. Measuring the gradient requires the selection of both a coordinate system and a velocity scale; in other words, it is dependent upon them.

Let us chose as our coordinate system an appropriate coordinate frame and a Cartesian decomposition (into a velocity vector) and a velocity scale of metres per second. To see how the measure works, consider a specific measurement. We measure the line by starting at the origin-point and taking its worldline (the origin worldline). We take a displacement along the origin worldline of one second; the one-second-origin – the ratio is simplest if we take a unit in the temporal scale. We then take the timeslice at that point and find the position of the intersection of the timeslice with the velocity line – the one-second-velocity-point. This is the point that object will arrive at after one second if it continues at the constant velocity. We then take the displacement from the one-second-origin on the timeline to the one second velocity point. We then use the standard Cartesian decomposition process on this displacement relative to the selected coordinate frame. This gives us the Cartesian displacements which we measure in metres – which give us the values of the coordinates in the velocity vector. This is a characterisation of the direction of the velocity worldline. Note that most of the stages here use the same patterns as position coordinate systems.

Speed

The speed of an object is also relative to a reference frame. The speed is the position an object would have travelled if it moves at the constant speed in a straight line. For each point in time, there will be a sphere, whose surface contains all the points that can be reached when travelling constantly at the measured speed. If you map these spheres in spacetime, the result is a 4D double (right circular) cone, whose apex is at the measurement point, whose axis is an origin worldline based upon the apex and whose intersection with timeslices is a sphere, whose centre is on the origin worldline.

Labelling and measurement of the speed is done in a similar way to velocity. One difference is there is no need for a coordinate system, just for a speed measurement scale; as the speed has no direction, it does not need a coordinate system. If we use metres per second as the speed scale[6], then we again take the displacement of one second along the origin-worldline. The intersection of the speed cone with one second timeslice is a sphere with the one second point as its centre. We measure the radius of the sphere. This gives us the metres per second measure of the speed. This is shown graphically for a 3D space in the figure below – 4D spacetime is challenging to show on a 2D page.

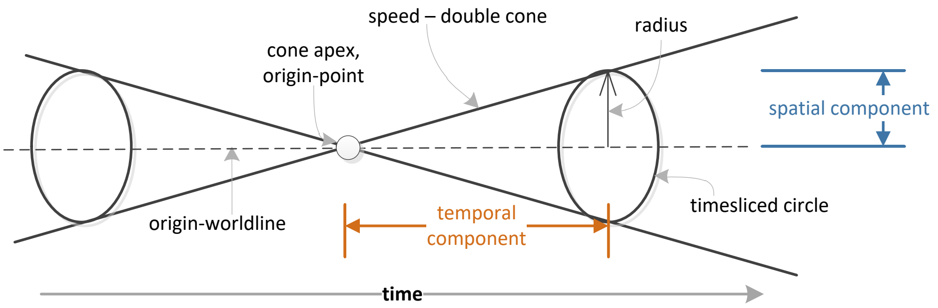

Figure 12- 3D Speed diagram

It should be noted that though the velocity and speed scales appear superficially similar – both are 'metres per second'. They are measured in different ways and so are not the same.

The pattern of dependency for speed – shown graphically below - while containing many of the same components as position, is different.

[6] Though the speed and velocity scales in our example use the same component scales, they are not the same – one is the gradient of a line (and has direction) the other is the gradient of a cone (and has no direction).

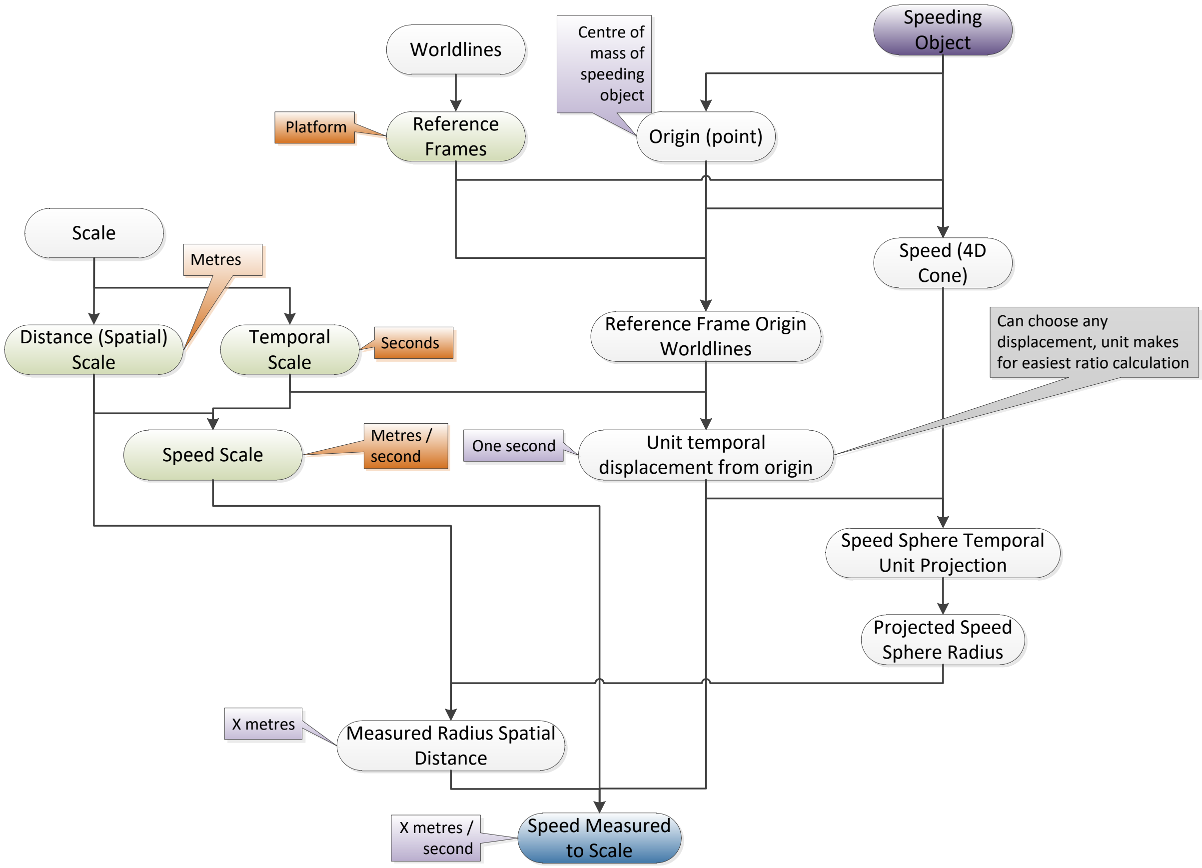

Figure 13 - Speed dependency map

Semantic Quality Assurance

The process of identifying the G&TR geometric objects that are being measured from the measurements expressed in numerical form and can be used as a systematic, objective test of the semantic quality. Where it is not clear what object is being measure, this is an indication that there is a semantic issue. Often, in context, there are general patterns that suggest what it might be. In this way, the ontological analysis process provides a systematic, objective test of the semantic quality, which not only identifies actual and potential sources of error and but suggests ways of resolving them.

Position, velocity and speed are common and well understood, so the ambiguity has been driven out of them. However, there are less common terms in the maritime domain what have not yet had all the ambiguity driven out of them. 'Track Made Good' is an example.

Track Made Good (TMG)

The Def-Stan 22-61 definition is:

> "6.8.3.6.2 Track Made Good (TMG) (from Def-Stan 09-100) - The direction of movement of a platform's CRP between consecutive fixes with smoothing applied over a period specified by the user."

Supplemented by:

> "NOTE Platform velocity and acceleration express the geodetic movement of the platform (represented by the CRP) in relation to the Earth's surface."

It is unclear what object this is. The platform's CRP will, relative to a geodetic reference frame, trace out a curve in spacetime. Consider a situation where smoothing due to changes in direction is not required; where the platform is travelling in one 'direction'. In the geodetic reference frame, the actual path is not a straight line. One could draw a straight line from one CRP position to another, but is this what is intended? Also, how does measure the angle of? Is it in the platform stabilised reference place at the start or the end of the measurement. The prior analysis has provided the components out of which the unambiguous definition can be crafted.

Future work – next stage

The sketch given above is intended to give an impression of what an information model would look like. The next stage is to build the information model.

The scope of the sketch is the core fundamental components of G&TR. The full information model will need to be extended to handle trajectories, error estimates, error chain management, corrections, extended objects, time intervals of validity, and other matters. From a project management perspective, this should be approached in stages.

Appendix A – Source Material

Def Stan 21-66 Issue 1 - Ministry of Defence - Defence Standard 21-66 - Common References Standard - Issue 1 Publication Date 16 April 2010 [UC].

Joint Tactical Air Defence Information Systems (JTADIS) Project.

Appendix B – Selecting the Right Structure for Spacetime

There is a need to decide upon a structure for spacetime. For our purposes the right trade-off between accuracy and simplicity is Galilean spacetime, so this is taken as the structure. The exact structure is currently a matter of scientific research. However, this is dealing with a fineness of detail that is far outside the scope of combat systems.

Spacetime Type	Velocity	Simultaneity	Distance
Aristotelian spacetime	Absolute	Absolute	Absolute
Galilean spacetime	Relative	Absolute	Absolute
Lorentzian spacetime	Relative	Relative	Relative
Pure Euclidean 4D spacetime	Relative	Arbitrary	Absolute

In spacetimes with absolute simultaneity (such as Aristotelian and Galilean space), the timeslice is not dependent upon the worldline of the point being used to timeslice.

Appendix C – Glossary and Abbreviations

Glossary

Term	Description
Ontology	In this context, the set of things whose existence the information model commits itself to. This includes all the things that is represents, but is often wider as it implicitly commits to the existence of some things. (1)
Master Training Datum (MTD)	A plane defined as part of the platform design. Designated by the physical siting of three datum plates. May or may not be collocated with the MLD and need not lie on the physical centre-line of the platform.
Master Level Datum (MLD)	A plane defined as part of the platform design. Designated by the siting and levelling of the level datum surface. May or may not be collocated with the MTD and need not lie on the physical centre-line of the platform.
Worldline	
Plane	A flat, two-dimensional surface.
Line (straight)	A straight curve. A straight one-dimensional figure having no thickness and extending infinitely in both directions.
Curve	A continuous one-dimensional figure having no thickness and extending infinitely in both directions.
Semantics	In this context, the relationship between the icons in the information model and the things they represent. This is what gives the icons meaning.

(1) This is closely based upon this kind of more general description of ontology as "the set of things whose existence is acknowledged by a particular theory or system of thought." E. J. Lowe in the Oxford Companion to Philosophy.
(2) A more general description of semantics is "… an important relation of words to objects – or better – of words to other objects, some of which are not words – or even better, of objects some of which are words to objects some of which are not words." Nelson Goodman. In the Introduction to Quine's lectures published as Roots of Reference.

Abbreviations

The following abbreviations are used in this report.

Abbreviation	Description
CRP	Common Reference Point
MLD	Master Level Datum
MTD	Master Training Datum
GRP	Geographic Reference Point
G&TR	Geospatial and Temporal Reference
TMG	Track Made Good
MDA	(OMG's) Model Driven Architecture
PSM	(MDA's) Platform Specific Model
CIM	(MDA's) Computation Independent Model
OMG	Object Management Group