TERRAIN INTEROPERABILITY IN LARGE FEDERATIONS: CORRELATION AND CONSISTENCY IN MC02

Steven D. Prager, Ph.D. and Dale D. Miller, Ph.D. and Durwood Gafford
Lockheed Martin Information Systems – Advanced Simulation Center
Bellevue, Washington

ABSTRACT

This paper details the complexities associated with creating an SNE terrain database that is consistent and correlated across a diverse set of interests, ranging from members of the simulation federation to C4I and planning systems.

In simulation federations, a number of players such as JSAF use CTDB data that is typically a polygonal terrain surface and is constructed in a real world (curved earth) spatial reference frame (SRF). In MC02 a critical federation member is JCATS, which typically uses a surface derived from bilinear interpolation of gridded elevation data in augmented UTM space. In order to achieve high levels of correlation, a series of developments in both the terrain database production process and in the simulation environment were required. A set of data products integrating the requirements for JCATS and CTDB correlation were produced, a regularly triangulated terrain surface for the CTDB and a series of elevation posts in the same SRF for JCATS, and near perfect terrain correlation was achieved. Special methodologies were developed to ensure that the representation of features in both data sets were as consistent as possible, further enhancing interoperability. The specially produced area was then integrated with the larger CTDB for the entire MC02 play box, and a seamless database with its special high correlation insert has become the foundational data product for the federation.

As there are numerous MC02 participants who consume standard NIMA products, a series of surrogate DTED data were produced from the combined terrain surface. This process ensures that terrain representation is consistent across participants, and users outside of JCATS and JSAF remain highly correlated.

This paper discusses the technical details and compromises associated with achieving the abovementioned terrain correlation and feature consistency, including specific discussion of SRF issues and the overall process of building the correlated and consistent terrain database.

ABOUT THE AUTHORS

Steven Prager is the lead engineer involved with the MC02 terrain database production process. Steven has over ten years of experience with spatial data automation and other spatial data technologies. Steven has been a critical resource in the development of a number of environmental data models and related modeling technologies as well as a key engineer in the development of a number of GOTS software products produced by the Lockheed Martin Advanced Simulation Center.

Dale Miller has led numerous projects involving geospatial data and environmental database technology, including the development of over ten integrated 3-D environmental databases for modeling and simulation applications as well as tools for automating the production of environmental databases. He currently leads the STRICOM OneSAF Environmental Data Modeling project, where he is developing EDMs for a number of legacy M&S and C4ISR systems, deriving EDMs for a number of NIMA vector products, and working with OneSAF SMEs to develop an EDM for the OneSAF Objective System. From 1994 to 1998, Dale led the Dynamic Terrain and Objects program and the environmental database Generation and Technology programs for the DARPA-sponsored Synthetic Theater of War (STOW). He is the author of over 35 papers in these fields and holds four patents.

Durwood Gafford has over 15 years of experience in the simulation database field. Durwood is an expert in the field of projections and coordinate systems and was one of the principle engineers behind the design and development of the methodologies required to achieve highly correlated terrain.

TERRAIN INTEROPERABILITY IN LARGE FEDERATIONS: CORRELATION AND CONSISTENCY IN MC02

Steven D. Prager, Ph.D. and Dale D. Miller, Ph.D. and Durwood Gafford
Lockheed Martin Information Systems – Advanced Simulation Center
Bellevue, Washington

INTRODUCTION

Modeling and simulation technologies have advanced to a point where large federations of many different simulation systems can interoperate in spite of the often disparate nature of each of the individual simulation systems. A key element of this interoperation is the presence of a synthetic natural environment (SNE) that is consistent and correlated across each of the simulation systems.

The Millennium Challenge 2002 (MC02) is "the nations primer integrating event." Sponsored by United States Joint Forces Command (USJFCOM), MC02 examines a "high-end, small-scale contingency that has the potential to escalate to a major theater war." The diversity of objectives associated with the MC02 exercise dictate the use of a variety of modeling and simulation, C4I, and Common Operational Picture (COP) systems. In that any number of systems require custom tailored configurations of exercise specific SNE data, special attention must be paid to the creation and distribution of the required terrain database.

Creating Consistent And Correlated Terrain

Shortly after the initiation of the database requirements analysis process, an extensive effort was carried out to capture the details regarding the manner in which the terrain data were to be utilized in the MC02 federation. As was quickly discovered, there were any number of simulation systems associated with the then current federation, and many required datasets that were not formally specified as components of the project's formal terrain database requirements.

A substantial systems engineering effort ensued, and was premised on the notion that each simulation and C4I system must have as accurate a representation of the data as possible, and that that representation must be as consistent and correlated as possible with each of the other simulation and C4I systems. In that many of the systems that were capable of ingesting feature data had the ability to consume data in the open ESRI® shapefile format, there was only a modest investment on both the production and consumption sides to accommodate feature consistency.

Terrain correlation was somewhat more problematic than feature consistency. Due to specific requirements of the exercise, the basis data used in constructing the polygonal terrain surface was substantively different than the equivalent NIMA products (DTED0) of the same areas. Whereas any particular simulation or C4I system may have ordinarily consumed a standard NIMA product for terrain, the requirement for high levels of correlation dictated that data consumers must use terrain information that matched the polygonal terrain surface in JointSAF as closely as possible. To this end, sets of surrogate DTED products were constructed for the exercise. These data were sampled directly from the polygonal terrain surface used to create the CTDB (Compact Terrain Database) for JSAF and served to represent the polygonal surface in grid form in a manner as correlated as possible.

Perhaps the most challenging aspect of achieving terrain correlation was with JCATS (Joint Conflict and Tactical Simulation). JCATS is an entity level simulation capable of representing terrain, features and entities with very high levels of fidelity. Whereas a simulation such as JSAF utilizes a database containing an irregularly triangulated polygonal representation of the terrain surface, JCATS dynamically interpolates a terrain surface based on a series of gridded elevation posts. Furthermore, at the onset of the project JSAF and JCATS utilized two very different spatial reference frames – GCS using the Geotile Reference System (GTRS) and augmented UTM, respectively. Needless to say, a cooperative spirit prevailed between the LMIS and JCATS teams and the most optimal solution set was arrived at through significant methodological adaptation and software engineering efforts by both teams.

In MC02, federation members JCATS and JSAF were responsible for the majority of the simulation entities. As a great deal of the exercise occurs simultaneously in both JCATS and JSAF, it was of the utmost importance that activities in JCATS and JSAF be as consistent and correlated as possible. The specific efforts required to achieve terrain correlation between JCATS and JSAF, as well as the compromises with the chosen methodologies are detailed throughout this document.

THE GLOBAL COORDINATE SYSTEM

The Global Coordinate System (GCS) is a spatial reference frame (SRF) developed in support of JSAF exercises requiring large databases. GCS defines multiple, local, geospecific Cartesian coordinate systems, tiling the surface of the Earth into cells. GCS includes well-defined transformations among cell-local coordinates and other coordinate systems in common use (e.g., geocentric and geodetic). GCS allows very large terrain databases to be represented with no positional distortion introduced by flat Earth projections. It allows simulation algorithms based on local Cartesian coordinate systems to be used with little modification.

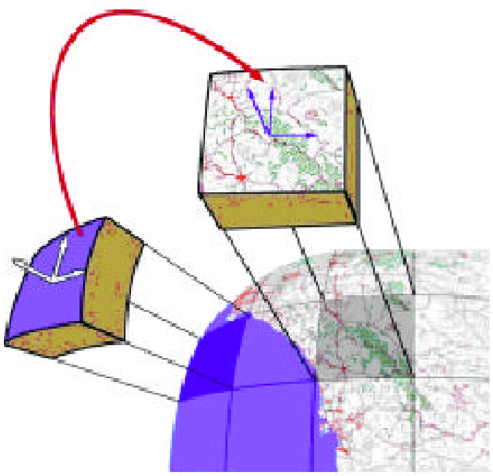

Figure 1. Each GCS cell has a unique coordinate system.

As illustrated (Figure 1), each 1°x1° area is referenced via a unique coordinate system. In cases where interactions in the SAF occur across tile boundaries, there are well defined coordinate conversion methodologies to account for coordinate changes between GCC, GCS, and geodetic. As all federation members interface with the federation network via GCC coordinates, however, only JSAF is required to be GCS aware.

The only exception to the GCS awareness requirement was the Joint Conflict and Tactical Simulation. JCATS had a set of very specific requirements in order to achieve maximum terrain correlation. The following section outlines how the GCS coordinate system was extended to assist in enhancing JSAF/JCATS terrain correlation.

GCS and JCATS

A critical component of the MC02 exercise is taking place within the 6°x6° inset area controlled by the JCATS system. Ordinarily, JCATS uses an augmented (i.e., with Z values) UTM projection for its simulation exercises. The UTM projection is based on a functional form using X and Y parameters to project spherical geodetic coordinates to the flat, orthogonal UTM coordinate system. The consequence of this projection/flattening of the data is that points are differentially projected as a function of their distance from the central meridian of the appropriate UTM zone.

In contrast to an augmented UTM projection, a real-world coordinate system preserves geometric relationships in three-dimensional space, allowing for a consistent representation of the terrain surface characteristics with minimal distortion to distance, elevation, intervisibility and other simulation factors (Toms & Birkel 1999). Being a "real-world" coordinate system, GCS is not a projection, but rather a transformation based on X, Y and Z parameters and a consistent spatial reference model. As mentioned above, each 1°x1° GCS tile exactly preserves the surface geometry (to within numerical precision), although a vector in the Z-direction will differ slightly from the gravity vector.

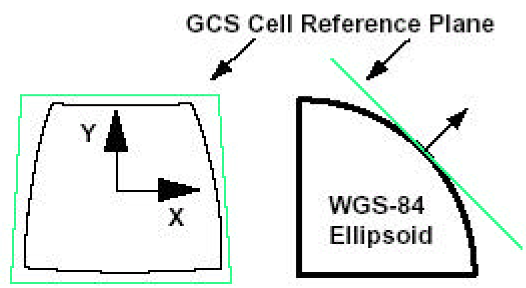

Figure 2. The GCS cell forms a curvilinear tangent plane.

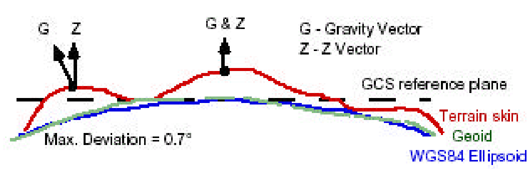

Figure 3. Variations in the gravity vector occur near the perimeter of each cell.

Importantly, the origin of each GCS cell is in the center of the tile, serving to equally distribute the effects of the tilt in the gravity vector. The maximum deviation occurs near the edge of the tile and does not exceed 0.7°. The multi-cell approach of GCS serves to minimize the net deviation, though the deviation associated with the "single cell GCS" was deemed acceptable and can be computationally compensated for if required (as it is a function of the angle between the gravity and Z vectors).

In that JCATS required as high a correlation as possible with JSAF, they chose to build their terrain based on a modified GCS spatial reference frame rather than the more typical UTM projection. To accomplish this without requiring JCATS to support multi-cell GCS, the GCS spatial reference frame was supplemented with an extended coordinate system capable of containing the entire 6°x6° JCATS play box. This coordinate system, an extended GCS cell known as the "single cell GCS" (although, strictly speaking, not truly GCS), facilitated the production of a database that would afford the greatest possible correlation between the JSAF and JCATS system while minimizing changes to JCATS software.

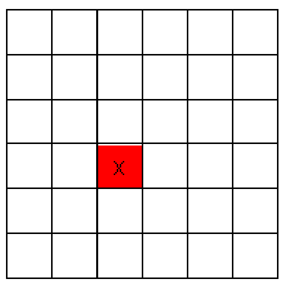

Figure 4. The home cell for the JCATS GridASCII database.

The single cell GCS coordinate space is a 6°x6° area that is treated as one, albeit large, regular GCS cell. Additionally, whereas the origin of the standard GCS coordinate system is the center of the plane for a 1°x1° GCS tile, the origin for the 6°x6° single cell remains the center of the home GCS cell for computational reasons.

The advantage of the single cell GCS plane is that it provided a systematic basis for the remainder of the terrain database production process to translate coordinates and elevation values between the JSAF and JCATS database production streams. The following section outlines the production efforts required to produce the terrain database, discusses how the production of the evenly sampled JCATS terrain posts worked in conjunction with the JSAF database, and how the single cell GCS for the 6°x6° inset area facilitated the entire process.

BUILDING THE CORRELATED TERRAIN

One of the most challenging elements of achieving terrain correlation was the development of methodology and means by which to achieve precise representation of elevation values common to both JCATS and JSAF. In a typical scenario, JCATS utilizes bilinear interpolation between each series of three grid posts (usually based on DTED data) to compute a "virtual" surface on an as-needed basis for mobility and line of sight calculations. JSAF, on the other hand, stores an irregularly triangulated surface (i.e., a TIN), that unambiguously defines elevation values for all points on the terrain through the interpolation along the faces of each irregular individual triangle making up that terrain.

In that JCATS requires a series of regularly spaced points from which to interpolate its elevation values, the irregular nature of the TIN terrain makes it unusable for JCATS. Additionally, the difference in coordinate systems as detailed in the previous section would pose significant difficulties in consistently correlating X, Y and Z values, particularly when Z values are interpolated for locations with no grid post (the majority of the database). After settling on the use of the GCS coordinate system, a methodology was developed wherein the interpolated triangles in JCATS would be physically represented in JSAF. The process of building the derivative elevation grid for JCATS, the regularly triangulated polygonal database for JSAF, and the implementation of a means to properly attribute the regularly triangulated database were all critical elements in achieving terrain consistency and correlation.

Building the JCATS Terrain File

The process of building the correlated JCATS/JSAF terrain began with the development of the "single cell GCS" concept detailed above. Based on this concept, a series of sample locations at 1000m intervals in both the X and Y directions were determined in the extended GCS space. For each sample point in the regular grid, the GCS Z was determined through an iterative process that converged on the geodetic latitude, longitude and elevation. Once the geodetic elevation was determined, from a DTED formatted file (produced from modified raster data derived from the Terrain Scenario Generation & Archiving data repository), it was converted back to GCS to complete the GCS XYZ triad. For purposes of providing the elevation data to JCATS, a simple "GridASCII" file format was used wherein a header detailed the origin and sample spacing of the orthogonal 6°x6° grid, and each value in the body of the file represented one elevation post in the extended single cell GCS space.

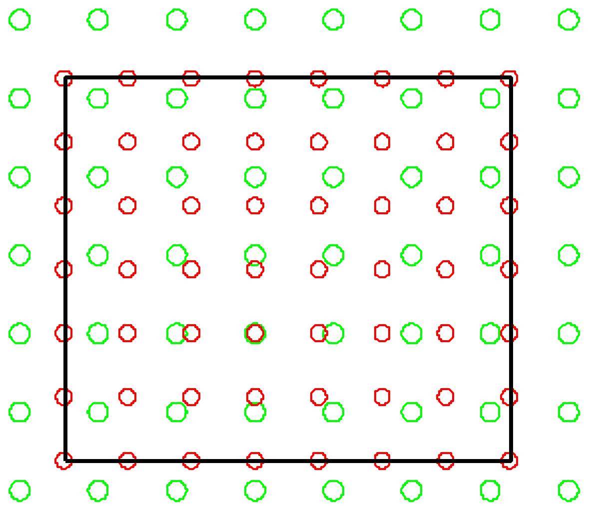

Figure 5. Offset sampling and interpolation is required to build the corresponding surfaces.

The above image illustrates a simplified example of the sampling concept. The red circles and the box represent the full extent and the 30 arc second sample posting of the input DTED elevation data. As illustrated, the 1000m sample posting (green, larger circles) requires interpolation to approximate the Z value for locations between the samples in the 30 arc second posting of the input data. The end result is an orthogonal set of samples that slightly exceeds the actual 6°x6° area, insuring full coverage in GCS space of the 6°x6° area in geodetic space.

Building the S1000 Database

The process chosen for building the exercise CTDB utilizes an S1000 database for key steps in the data processing procedure. For this reason, the S1000 database format served as an ideal platform from which to build a JCATS-correlated polygonal terrain surface. The Landtool, a system designed and maintained by Lockheed Martin specifically to build gridded S1000 terrain databases, had already been extended to accommodate the GCS SRF and was a logical choice to extend to accommodate the newly designed extended single cell GCS SRF.

In that the 6°x6° area in JSAF requires a "true" GCS representation comprised of 36 GTRS GCS cells, a multi-step process is required to convert the JCATS elevation posts to terrain polygons in the appropriate target cell. First, each set of four adjacent elevation posts that potentially fall within the target cell are used to derive two polygons, each with three vertices (using a lower-left/upper-right diagonal rule). These polygons are compared to the geodetic bounding box (plus some tolerance to account for the difference between the curved earth and the orthogonal GCS plane. Any polygons not falling completely within the bounding box are transformed to the geodetic coordinate space, clipped to the geodetic geotile boundary, and converted back to GCS space.

Once the set of elevation posts are converted to polygons in their appropriate target cells and clipped to the geodetic boundaries, the polygons are then clipped to a unit known as the "load module" boundary. The load modules are a functional unit for handling terrain areas in JSAF, and evenly divide the GCS cell into a series of smaller rectangular regions. In that the chosen load module size (4000m) coincides with every fourth set of terrain polygons, and the axes of the polygons are perfectly cardinal in the home cell, no clipping was required (Figure 6).

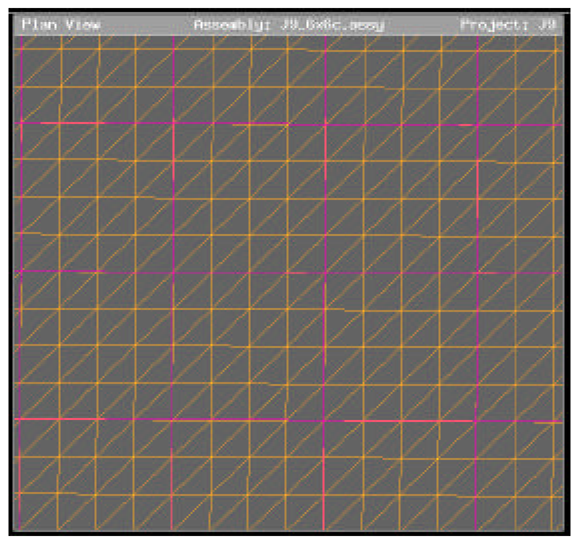

Figure 6. A zoomed view of the "home cell."

In contrast to the home cell, when polygons outside of the home cell are converted to their target cell, they do not align with the load module boundaries of their target cell. Figure 7 illustrates how the load module lines (which are oriented orthogonal to the X & Y axes) are canted relative to the regular terrain polygons.

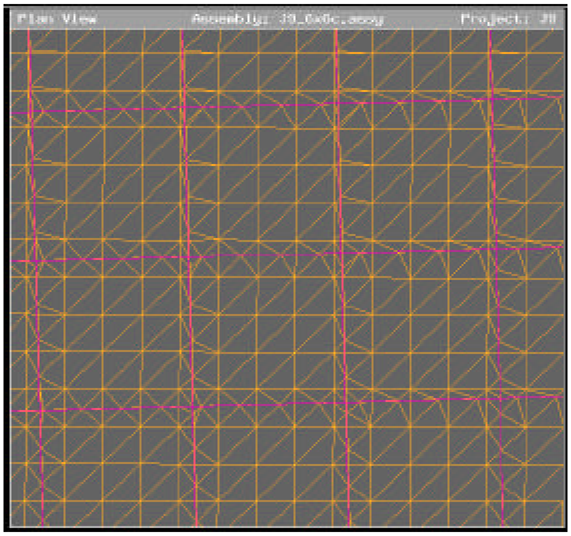

Figure 7. Canted load module lines show obvious offset along the periphery of the inset area.

Importantly, neither clipping to the geodetic tile boundary nor clipping to the load module lines modified the morphology of the terrain skin. Figure 8 depicts a shaded relief view of the wire frame illustrated in Figure 7. As shown, the planarity (as evidenced by consistent slope) of the terrain polygons associated with the gridded JCATS terrain file is preserved.

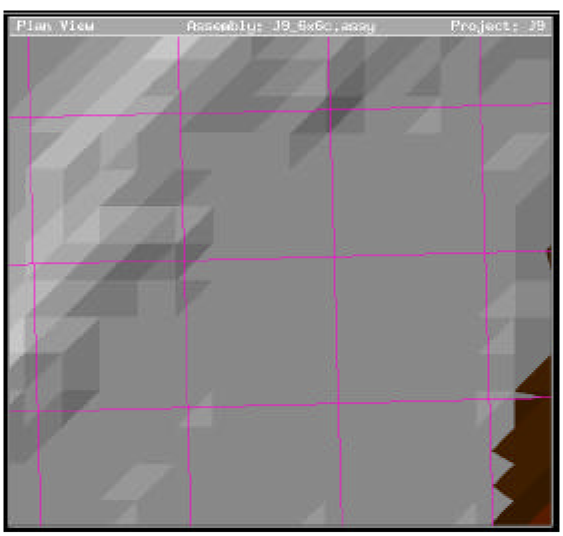

Figure 8. Shaded relief illustrates consistent surface morphology in clipped terrain polygons.

Clipping Features into the S1000 Assembly

Once the terrain polygons are constructed in a manner correlated with the JCATS terrain file, they must be attributed so as to reflect their proper SIMNET soil type and other surface feature characteristics (e.g., attribution defined by the representative Environmental Data Model). The attributes defining the surface characteristics are assembled during the GIS data processing phase into a data set known as the "Surface Area" data. These data contain a continuous manifold coverage comprising the entire surface area of the database. The surface area data are made of a series of shapefiles, one for each feature type (e.g., cropland, ocean, built up area, etc.). A series of preprocessing steps merge the single feature class shapefiles into a single shapefile (one per geotile) that, in turn, serves as the primary input for the remainder of the process clipping and attribution process.

The standard S1000-based route to CTDB production integrates surface area feature information directly into the TIN as it is produced. This automatic incorporation of the surface area feature data is only possible, however, with irregular triangulation and does not permit either the generation or preservation of terrain surface morphology that meets the JCATS correlation requirement. Consequently, an alternative strategy was developed wherein the terrain polygons would be attributed in a post hoc manner that preserved their geometric characteristics and, consequently, JCATS correlation.

The first round of polygon attribution (in support of one of the early data deliverables) was performed using a methodology wherein the shapefile feature data associated with the approximate centroid of each terrain polygon determined the attribute of that terrain polygon. This process was already supported by the Landtool and was used to provide a preliminary set of attributed terrain polygons.

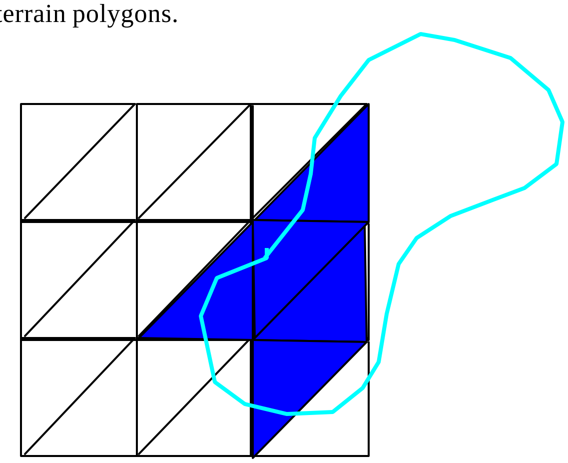

Figure 9. A lake abstract used to attribute terrain polygons.

In this simple application of the polygon attribution process, each terrain polygon associated with the lake (shaded triangles as shown in Figure 9) is completely unchanged with the exception of its new attribution. This process results in a characteristic known as the "jaggies" and is a cause of feature inconsistency between TDB representations. In JSAF the abstracts are used to illustrate terrain features such as areal lakes and rivers are imported through a process that uses the original shapefiles to determine the feature geometry. As illustrated above, however, the shape of the abstract and the shape of the underlying polygons that represent the abstract are distinctly different, and could cause unexpected behavior during an actual simulation (in JSAF, the lake abstract would be visible, but the unseen polygon attribution would determine vehicle mobility).

The alternative strategy (as employed for the later deliverables) requires clipping each terrain polygon to the boundaries of each of the surface area feature. Using an algorithm known as the "Weiler Clipper," (Weiler 1980) this clipping process preserves the surface morphology of the original JCATS compliant triangulation. Each parent triangle is subdivided (if necessary) in to smaller triangles that accurately reflect the shape of the intersecting surface area feature.

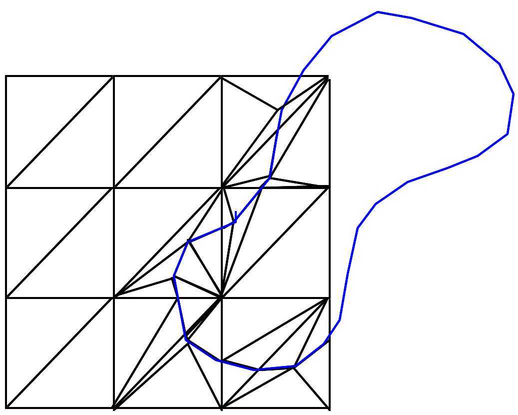

Figure 10. The abstract used to clip geometry in X & Y.

Once clipped, the new set of terrain polygons are attributed, and the attributed areas properly reflect the geometry of the features in question. The same centroid based attribution mechanism used to attribute earlier releases of the database is still usable. Now, however, the terrain polygons precisely align with the feature boundaries, and no jaggies occur.

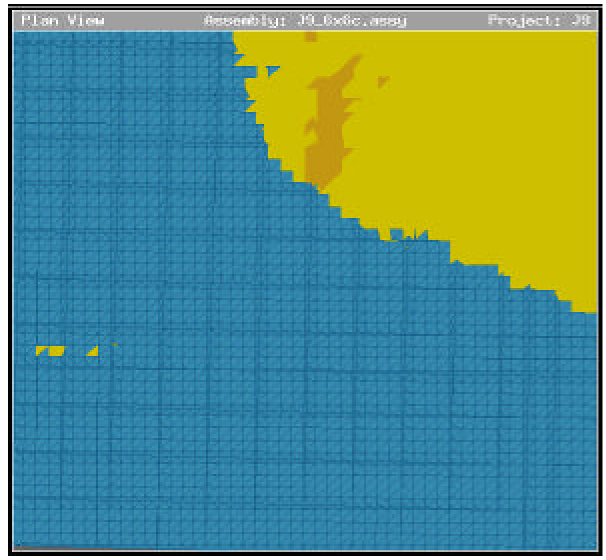

Figure 11. Attribution without clipping results in the jaggies.

Figure 12. Attribution after clipping results in proper attribution.

The above images illustrate two snapshots of the S1000 database produced just before (Figure 11) and just after (Figure 12) the clipping. As illustrated, the clipping process completely eliminates the jaggies, greatly enhancing terrain consistency. Importantly, both of the above methodologies preserve correlation with the terrain surface geometry produced from the JCATS terrain file. The latter methodology, while more expensive from both the computational and engineering cost perspectives, provides a significantly higher quality database product with more predictable SAF performance and improved JCATS feature consistency as well as consistency between the features and the terrain surface polygons within JSAF.

Integrating JCATS-correlated Terrain with JSAF Specific Terrain

Given the regular triangulation associated with the JCATS correlation requirement, the production process associated with compiling the 6°x6° is fundamentally different than that used to compile the terrain for the surrounding area. A consequence of this difference is that extra steps must be taken to integrate the two portions of the database resulting from the two different production streams. These extra steps are required as the edges of tiles in the 6°x6° JCATS inset area and the edges of the abutting tiles are triangulated using different methodologies. Such edge mismatch poses significant problems in the JSAF arena, and requires specific remedial steps to enforce continuity along the perimeter of the 6°x6° inset.

The process of enforcing database continuity in the transition area involves building a series of three dimensional reference lines that exactly replicate the vertical cross-section of the 6°x6° boundary. In that JCATS correlation is the paramount edict for the 6°x6° area, the abutting terrain is preferentially adjusted to match the 6°x6° inset.

The process of constructing the vertical cross-sections (known as "3D edges") is relatively straightforward. The tiles in the 6°x6° area are sampled in a manner constrained to capture terrain polygon vertices along all four edges of each tile. In turn, the X, Y and Z (in the GCS SRF) are stored as features in a 3D shapefile.

The vertices on the inner edge of the tiles along the perimeter of the 6°x6° area (as outlined by the outer ring above) must be forced to reflect the vertices associated with the inset area. A procedure is required to replace the shared edges generated for the perimeter tiles with the edges generated for the inset tiles. Once the edges are built, this process becomes the spatial analogue of a "cut & paste" operation.

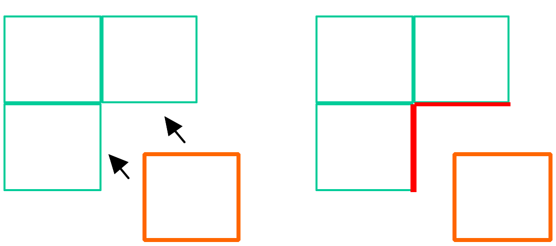

Figure 13. Transposition of inset edges to corresponding terrain.

As the above image illustrates, the edges taken from the inset area are superimposed on the edges captured for the surrounding geotiles. Once the set of hybrid edges are constructed, they serve as inputs into a process where the TIN for each tile in the outer ring is recomputed, but now with an enforced line of constraint based on the content of the 3D edge files. Enforcing the 3D edges insures that the outer perimeter of the inset area will match the inner perimeter of the ring, and the outer perimeter of the ring will continue to match its adjacent tiles in the remainder of the database.

While methodologically clear-cut, the process of executing the above operation is somewhat complicated by the fact that the inset area coupled with its outer ring crosses two megatiles (the basic unit for compiling an S1000 assembly). This overlap required replication of the re-TINing process for a number of tiles, and careful management of all operations involving tiles on both sides of the transition boundary. Once the individual S1000 assemblies were reworked with the new edge data, they could simply be fed to the CTDB compiler for construction of the runtime database.

QUALITY, COMPROMISE AND DERIVATIVES

Correlation and consistency are the key requirements that drove the database production process. Without maintaining quality, however, the produced data would be of limited utility in the exercise federation. To this end, extensive quality checks (both internal and external) were performed. Furthermore, in spite of the suitability and appropriateness of the production methodology, internal constraints of the simulation systems coupled with other database requirements (e.g., resolution), resulted in a state wherein even the most consistent and correlated terrain possible had limitations. Finally, though the majority of the simulated entities are generated by JCATS and JSAF, a great deal of activity occurs in simulation and C4I systems ancillary to JCATS and JSAF. These other systems also required consistent and correlated data and had to be considered appropriately. Derivative data products in support of non-JCATS and non-JSAF federation members needed to maintain the same level of quality and reliability as the more visible data products if the true benefits of the correlation edict were to be fully realized.

SEE-IT Analysis of Terrain Products

In that much of the software to produce the above mentioned data was developed in house by LMIS and was based on the extension of a number of legacy systems, a great deal of time was spent closely examining the minutia of the terrain data, the associated geospatial features, and the associated attribution. The consequence of this "data intimacy" is something of a problem akin to the adage of "not seeing the forest through the trees." When one is worrying whether floating point coordinates associated with two particular vertices are "equal" to 12 or more significant decimal digits, it is easy to lose sight of the bigger picture. The value of a third party analysis is that it serves to leverage a different set of expertise and to lend new perspectives from which to view the data. For this reason, among others, the Institute for Defense Analysis (IDA) was enlisted by J9 to evaluate the quality of the delivered data products.

The Synthetic Environment Evaluation-Identification Tool (SEE-IT) is a software system developed by IDA in support of evaluating the quality of SNE terrain databases intended for the DoD M&S community. Among its key capabilities are its ability to detect and report horizontal and vertical tears in terrain polygons, terrain polygon overlaps (i.e., where multiple polygons are collocated), "T" vertices (i.e., an improperly triangulated terrain area), and slivers (i.e., long, thin terrain polygons).

For purposes of analyzing the MC02 terrain data, IDA extended the capabilities of SEE-IT to natively read in the JCATS terrain file, the intermediate S1000 assembly, and the end product CTDB. This collective capability enabled IDA to compare each of the terrain database products to check for cross-product correlation (particularly between the JCATS GridASCII and the CTDB), and to examine each product for any of the above listed problems.

The IDA analysis quickly confirmed that the datasets that were expected to be consistent and correlated met expectations. Importantly, IDA also found a number of holes (polygon dropouts), T-vertices, slivers and tears in the production database. Their analysis quickly pointed out problems with the production methodology along the 1°x1° cell boundaries.

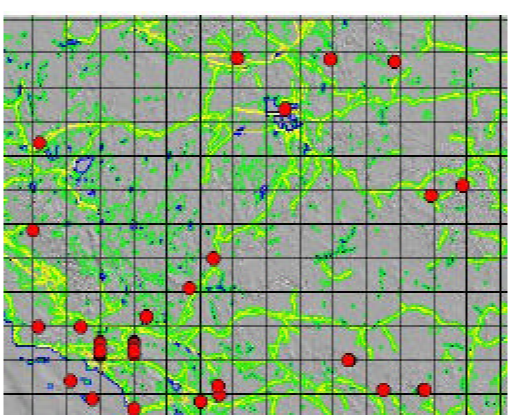

Figure 14. The SEE-IT analysis illustrates the locations of tears in the final CTDB (red dots).

While the SEE-IT analysis was very useful in pointing out terrain anomalies, it was also useful in describing other database characteristics as well. SEE-IT is capable of performing extensive line-of-sight analysis as well as terrain and feature cross-section analysis. At several points during the terrain analysis iteration, SEE-IT was able identify potential database traits that might cause problems with JSAF behaviors. While in many instances the traits SEE-IT was found were not adversely affecting the SAF, there is significant potential to develop CTDB/SAF optimized SEE-IT-based analytical processes. This approach would ensure that the analysis performed on CTDB and related data was specific to the requirements, characteristics and behavior of the JSAF runtime environment.

Overall, through the use of SEE-IT, IDA determined that the terrain file produced for JCATS and the CTDB for JSAF were very highly correlated. With minor exception, the two databases offered nearly identical discrete elevation values and line of site characteristics.

Terrain Resolution and Coastline Issues

Early in the development of the exercise scenarios and corresponding database requirements, it was determined that low resolution terrain (i.e., equivalent to DTED0 in sample spacing) would serve to optimize the size of the CTDB relative to the load of JSAF entities it was expected to support. An advantage of lower resolution terrain is the reduction in polygon count for terrain polygons in the CTDB. Not only does a reduced polygon count allow for the creation of a smaller CTDB, it also simplifies the CTDB production process (particularly on databases as large as that required for MC02).

Unfortunately, while reduced resolution terrain does facilitate a physically smaller database, it does result in the potential obfuscation of possibly important landforms. Furthermore, in that the polygons that were produced for JCATS compliance were regularly triangulated, standard procedures such as coastline clamping were not possible. Ordinarily, the coastline feature is integrated during the TIN process to ensure that all its vertices have elevation at 0m. The purpose of this action is to define a clear boundary between the bathymetric data and the terrain elevation data. Clamping the coastline also serves to ensure that all polygons falling within the ocean abstract are below 0m MSL and all terrestrial polygons adjacent to the coastline are above 0m MSL.

For this database, the area outside of the 6°x6° inset was processed in the normal manner and the interface between the water and land was clamped to a 0m line of constraint. The requirement of maintaining the regular triangulation in the inset area superseded the clamping requirement, however, and therefore introduced some elevation anomalies along the land-water interface. The problem originates from the fact that the 1000m sample spacing could occur anywhere in relation to the actual coastline (and, hence, a known 0m MSL elevation value). In that a substantial amount of elevation change can occur over the 1000m sample distance, there is a significant likelihood that, through interpolation, positive elevation values can fall in the water and negative elevation values could fall on land.

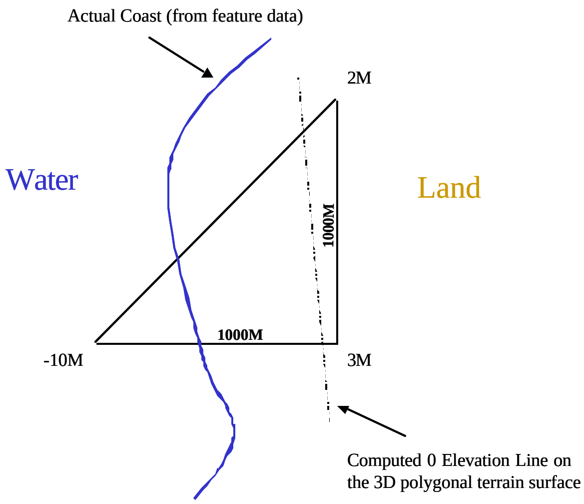

Figure 15. The zero-elevation line can substantively differ from the actual vector coastline.

As illustrated , if the regular triangulation is maintained, there are cases where a significant portion of any given terrain polygon may be above or below 0M MSL, even though the feature type would indicate otherwise. In the above case, the portion of the terrain polygon that is left of the computed 0 elevation line would return a negative elevation value even though is in the land area. This phenomenon was limited to the inset area and was determined to by less significant than the miscorrelation between JSAF and JCATS that would be introduced if the coastline were clamped as per standard production methodology.

Interestingly, this problem affects JCATS and JSAF in different ways. JCATS assumes that when you enter the ocean (as defined by the polygonal ocean feature), the height of the terrain surface (i.e., the ocean) is 0M MSL. Ordinarily, this would match the JSAF assumptions about the ocean surface. JSAF, however, has a multifaceted representation of the ocean that includes both an ocean surface and a bathymetric ocean bottom. In cases where positive elevations occurred on the ocean side of the coastline, this bathymetric surface was above the surface of the ocean. This condition resulted in initially puzzling inconsistencies between the two simulations but was later resolved with simple code changes in JSAF.

Derivative Data Products

Though terrain correlation between JSAF and JCATS required the development of special methodology and custom tailored terrain database production strategies, the correlation mantra was required to extend to all of the MC02 federation. In that there were numerous participants, both simulation and C4I systems, that required a consistent view of the world, the production strategy was extended to build datasets that would serve as surrogates to the standard NIMA data that these systems ordinarily consumed.

The primary surrogate NIMA data was a set of DTED0 and DTED1 files produced for use by the C2W Analysis and Targeting Tool (CATT) and Distributed Information Warfare Constructive Environment (DICE) which model enemy air defense radar and need precise terrain correlation. In addition to the various models and simulations, the many C4ISR systems also consumed the surrogate NIMA products: the Global Command & Control System (GCCS) and the Automated Deep Operations Coordination System (ADOCS), among others, were able to utilize the surrogate products to provide an operational picture and analytical environment in agreement with the simulation taking place. These data were closely compliant with the NIMA DTED performance specification, but were custom built via sampling the terrain polygons that made up the JSAF database. Interestingly, the custom constructed DTED also helped to mitigate a federation issue wherein one of the standard RTI connections was expecting data referenced to the ellipsoid instead of the geoid (MSL). In that the simulations that were affected by this issue were DTED consumers, an additional set of surrogate DTED referenced to the ellipsoid in lieu of MSL were produced. This allowed the affected simulations to simply load an alternative dataset rather than worry about constantly performing the geoid correction that would otherwise be required.

CONCLUSION

The success of Millennium Challenge 2002 is prefaced on the quality of the input data supporting the exercise scenario and the federation members. In that the terrain and feature data characterizing the Synthetic Natural Environment are a critical component of the input data, it is fair to suggest that the quality of the exercise hinges, at least in part, on the quality of the “dirt.” Arguably, good “dirt” can, and should, be a background phenomenon. The extent to which the transparency of the terrain minimizes the need for special workarounds, is a direct measure of the quality of that terrain. While the data produced for MC02 are not completely without caveat (some of which is directly attributable to the manner in which the database was specified), the success of the test spirals to this point is indicative of the overall success and appropriateness of the production process.

REFERENCES

Ceranowicz, A., Torpey, M., Helfinstine, B., Evans, J., & Hines, J. "Reflections on Building the Joint Experimental Federation", I/ITSEC, Orlando FL, (2002).

Lukes, G. E. & Birkel, P. A., Eds. "Synthetic Environments Final Technical Report", DARPA, (1998)

Richbourg, R. F., Graebener, R. J., Stone, T. & Green, K. "Verification And Validation (V & V) of Federation Synthetic Natural Environments", I/ITSEC, Orlando FL, (2001).

Toms, R. M. & Birkel, P. A. "Choosing A Coordinate Framework For Simulations", 99F-SIW-183, Proceedings of the Spring Simulation Interoperability Standards Organization Conference, Orlando FL, (1999).

Weiler, K. "Polygon Comparison Using a Graph Representation", SIGGRAPH 80, pp 10-18 (1980).

ACKNOWLEDGEMENTS

Many people were responsible for the overall success of this project. LMIS had an opportunity to establish a close working relationship with JFCOM-J9. Working initially with LtCol John Kelly and then Mr. Dan Snyder, numerous problems and issues were realized, defined and quickly resolved.

LMIS also had the privilege to work closely with the JCATS team at Lawrence Livermore National Laboratory. Establishing a strong working relationship with Tom Kelleher, Jeff Pimper and the remainder of the JCATS team was crucial to the success of the high levels of correlation achieved for the inset portion of the database.

Finally, the team at IDA proved to be a crucial element of the project. Bob Richbourg, Tim Stone and Robert Graebner are extremely knowledgeable regarding any number of terrain database and derivative formats, and were significant contributors to the overall focus on a quality database and sound production methodology.