A Concept for Creating, Storing, Updating and Delivering Variable Resolution Terrain Models.

Jan T. Bjørke[1,2], Morten Hæger[1], Knut Landmark[1] and Stein Nilsen[1]

[1] Norwegian Defence Research Establishment,
PO Box 115, N-3191 Horten, Norway
jtb@ffi.no, moh@ffi.no, kla@ffi.no, sni@ffi.no
[2] Norwegian University of Life Sciences,
PO Box 5003, 1432 Ås, Norway

Abstract. The NVRG/Sarepta is an implementation of a terrain model server based on a global variable resolution grid. Access to the system is through a web service, which makes it essentially platform independent. The grid tiles are quadrilaterals of nearly homogeneous size across the globe. Finer resolution levels are generated by regular subdivision. The terrain data is organized in a file structure that supports easy updating and maintenance. Users can run local copies of NVRG/Sarepta on stand alone machines or on separate networks, and they can import their own data into the system.

Keywords: Global terrain models, web services

1 Introduction

When comparing terrain elevation measurements obtained from multiple sources, the most noticeable difference between data sets is usually the resolution. For example, echo sounders mounted on ships or underwater vehicles are used for measuring the topography of the seafloor. The measurement density depends on the distance from the transducer to the seafloor, and the characteristics of the echo sounder, such as the number of beams and the ping rate. In a seafloor model compiled from a combination of single beam and multibeam echo sounder data, we can have areas where the sample distances are varying from centimeter to kilometer level. Variable resolution is therefore a key question in the design of digital terrain models.

What we propose in this paper, is a concept for creating and managing terrain models where the resolution of the original measurements is reflected in the models. The concept is implemented at the Norwegian Defence Research Establishment in a system called NVRG/Sarepta. This system is capable of creating, storing, updating and delivering global digital terrain models. Access to the data is through a web service.

Research in digital terrain models has resulted in a multitude of interpolation methods and data structures. The different methods have strong and weak sides, and it is probably impossible to find a digital terrain model implementation that in all aspects outperforms every other alternative. Classification of digital terrain models can be based on how the terrain domain is partitioned and how the interpolation is defined. Commonly used partitions of the domain are either based on rectangular cells (raster models) or triangles. The obvious advantage of triangle models compared with raster models, is their flexibility in how the domain is partitioned; the drawback is their complex construction.

Triangle models were introduced in digital terrain modeling in the 1970's, see for example [McL76], [Yoë77] or [MR80]. As the application of triangle models became more popular, focus was brought to their time complexity, [Bjø88]. The introduction of adaptive triangle models at the end of the 1980's, e.g., [Hel90], opened a new branch of research.

A multiresolution database based on the Delaunay pyramid is described by Ware and Jones,[WJ92]. The database is accessed via the triangle structure and topographic objects. This allows features necessary for characterizing the shape of the terrain surface and topologically structured objects such as lines, to be integrated into the pyramid.

De Floriani et al. [DFP95] have focus on the shape of the terrain surface and show how a system of hierarchical triangle models can be constructed. VARIANT, [DFMP00], is a terrain database manager based on the idea of multiresolution and the application of triangulated irregular networks. The authors admit that the present implementation of VARIANT is a prototype which needs some restructuring to become effective and efficient at global scale. One of the problems seems to be how a small part of the model can be loaded from the disk into the main memory, i.e., how the model can be split into parts that are multi-triangulations. Moreover, updating the models each time new data are made available, is stated by the authors as a non-trivial problem.

A method for merging Delaunay networks is presented in [Mid93]. The author shows how triangle models can be partitioned into disjoint blocks, and merged on the fly at application time. The size of the blocks can be optimized based on the capacity of computer resources like network, main memory and disks.

A parallel to the diversity of methods in the field of digital terrain models can be found in multidimensional access methods. Gaede and Günther, [GG98], point out that in this field vendors tend to resort to structures that are easy to understand and implement. We adopted the same pragmatic view for the design of a global digital terrain model. Due to the complexity of irregular triangulated networks, our construction is based on a matrix representation. This concept is used to develop variable resolution terrain models with transparent addressing scheme and robust methods for splitting the models into segments of appropriate size. In this way global terrain models that are both easy to understand and simple to implement are obtained. The structure is called the NVRG – the 'N' Variable Resolution Grid. The underlying tiling scheme has grid lines that are parallels and segments of meridians, and each tile has approximately the same shape and size all over the globe; rectangular with height equal to one minute of arc, and a latitude dependent width close to one nautical mile. Using regular subdivision of the tiles gives us infinitely many resolution levels.

Since we are faced with the problem of handling data sets from different sources, our implementation must provide a solution for how to handle overlapping data sets with different accuracy and resolution. To get a coherent model, resampling of one or more of the data sets might be a strategy, but this procedure will most likely introduce artifacts in the model. In addition, the variance of the data will in many cases be reduced, [BN07].

These problems can be avoided by allowing data sets with differing properties to live side-by-side, but this strategy leads to several new challenges. Without a common representation, we must consider: Which model and representation is valid in a domain of overlap? How should data sets be adjusted and imported into existing models? Can the data be represented efficiently? And, do common GIS tools have the ability to visualize the data?

An NVRG model is defined as a collection of terrain data files that are either (1) prepared for special purposes, (2) compiled by another organization, or (3) generated from our own survey data. We prefer not to merge data sets obtained from different sources into one common model, but this can be done upon request. The procedures for constructing coherent models are application dependent, and outside the scope of the present study. When it comes to representation and visualization of variable coverage/resolution data sets, there are several viable candidates. One open source alternative is the Python configurable OpenEV.

Sarepta is a web service built on top of the NVRG. Together they constitute the NVRG/Sarepta system. The storage component of the system is not implemented as a data base, but as a file collection. There are several reasons for making this choice: Maintenance is easy and requires little training. This is important for the deployment of the system to military, as well as civilian, mobile units. Many of the features present in spatial data bases, such as interpolation, point evaluations and coordinate transformations, we prefer to implement ourselves, and spatial indexing is taken care of by the NVRG. Since our terrain models are essentially images, meaning that they are not meant to be updated or edited by users, there is not much to be gained, even from a performance perspective, by using a relational data base.

The NVRG/Sarepta system contains variable resolution terrain models with fixed resolution levels. The natural hierarchy of the grid makes XML a suitable and efficient format for describing the models. The XML files list every grid block covered by a model, the resolution level of the coverage, and the filepaths. The Sarepta service locates models by querying the XML documents. This solution not only decouples the service from the storage component, but also increases the flexibility of the system. For instance, a model can be accessed directly from a CD without the need for installation as long as an XML model file is made available to the service.

The NVRG, which is a refinement of the tiling scheme presented in [BGHN03] and [BN04], is described in detail in the next section. We then give an overview of the NVRG/Sarepta system, before we discuss the different components.

2 NVRG

2.1 Tiling Scheme

The WGS84 ellipsoid is divided into bands and sectors. Band heights differ, but sectors are always 30°. The cell structure inside each sector is identical. The part of a band that lies between two sector lines is called an intersection-block, or i-block for short. In general, any collection of cells of rectangular shape will be referred to as a block. We use the convention that the southern and western boundaries always belong to a domain. Consequently, the south-west corner uniquely identifies a cell/block of a given size.

The specification of the grid is given in Fig. 1. Columns two and three show the exact latitudes of the band boundaries, expressed in minutes of arc and in degrees/minutes. The next columns show the number of cells vertically and horizontally in each i-block. The last column lists the exact widths of the cells, measured in minutes of arc. Figs. 2 and 3 show the structure, and how the cells change across the band boundaries.

2.2 Request and Resolution levels

The fundamental unit in the grid is the cell. All cells have roughly the same vertical and horizontal dimensions: One minute of arc by (approximately) one nautical mile. We define different resolution levels inside the cells by using regular subdivision. There is thus no limit to the resolution the grid can support. Resolution level zero is a whole cell – at level one there are four subcells. In general, at level n there are $2^n \times 2^n$ subcells.

There are several ways of grouping cells into larger units. A user might request data covering an area delimited by particular angles of latitude and longitude. Up to 74° the grid supports blocks covering $5' \times 5'$ regions; up to 88°, $15' \times 15'$ regions. A common unit for partitioning the globe, is the $1^\circ \times 1^\circ$ block, which is supported by the NVRG up to 89°. We restrict the request areas to cells, subcells, and certain aggregations of cells: These are the *request levels*. We use the notation `Cn` to describe cell and subcell request levels; e.g., `C0` is a fundamental cell, and `C8` is a subcell covering $(1/2^8) \times (1/2^8)$ of a cell, i.e., an area with approximate dimensions 8m by 8m. For larger areas we use the notation `Am`. Request level `A1` is a block covering 5' by 5'; `A2` covers 15' by 15'; `A3` is 1° by 1°, `A4` is a 6° wide subsector of an i-block, and `A5` is a whole i-block.

The *resolution level* refers to the size of the pixels in a request area. The coarsest resolution level defined is 0; this corresponds to a pixel size equal to one cell. Resolution level `n` means that there are $2^n \times 2^n$ pixels inside a cell.

bdry #	latitude (mm) / (ddmm)		band height	cells in sector	cell width
0	0m	0d 0m	1800 cells	1800 cells	0d 1m 0s
1	1800m	30d 0m	1200 cells	1440 cells	0d 1m15s
2	3000m	50d 0m	900 cells	1080 cells	0d 1m40s
3	3900m	65d 0m	540 cells	720 cells	0d 2m30s
4	4440m	74d 0m	240 cells	480 cells	0d 3m45s
5	4680m	78d 0m	240 cells	360 cells	0d 5m 0s
6	4920m	82d 0m	240 cells	240 cells	0d 7m30s
7	5160m	86d 0m	120 cells	120 cells	0d15m 0s
8	5280m	88d 0m	60 cells	60 cells	0d30m 0s
9	5340m	89d 0m	30 cells	30 cells	1d 0m 0s
10	5370m	89d30m	15 cells	15 cells	2d 0m 0s
11	5385m	89d45m	7 cells	10 cells	3d 0m 0s
12	5392m	89d52m	4 cells	4 cells	7d30m 0s
13	5396m	89d56m	2 cells	2 cells	15d 0m 0s
14	5398m	89d58m	1 cell	1 cell	30d 0m 0s
15	5399m	89d59m	1 cell	1/3 cell	90d 0m 0s
16	5400m	90d 0m			

Fig. 1. Definition of the NVRG. The ellipsoid is divided into 32 bands, and 12 sectors. Sectors are 30° wide. Measured in geographical coordinates, all grid cells within a band have the same width. The number of bands, and the location of the band boundaries, are chosen so that the physical size and shape of the cells is as homogeneous as possible across the globe. An arbitrary "level one" subcell will fit inside every grid cell; there are four "level one" subcells in a cell.

2.3 Cell and subcell contents

Values assigned to cells and subcells are area values, not point values. The error/variance/uncertainty associated with a point value only applies at the location of the point. A corresponding number associated with an area value applies to the whole area. Point values are usually the result of interpolation. The density of a grid of interpolated point values does not necessarily reflect the resolution of the underlying data set. An area value has an implicit resolution given by the dimensions of the area.

Area values should only be assigned to cells where the sample density supports the calculation of that value; there should be no interpolation. A central idea behind the NVRG/Sarepta system is processing and storage of terrain data with the exact same resolution as the data collected during the survey. Coarser models can be distributed, but smoothing and interpolation, either for visualization or for populating areas where no values are assigned, belongs at the application level.

It is not required that values are assigned at every level of resolution. Inhouse models are consistent across resolution levels, but this need not be the case for multiresolution models created by other organizations. The solution is to base the NVRG "instance" of the model on data from the finest level, and use the NVRG projection operators to generate the coarser levels; see next section.

If possible, we assign three values to each subcell: `cell_mean`, `cell_maximum`, and `cell_minimum`. A vital criterion for a variable resolution grid, is that values need not be assigned to every subcell at a given resolution level. At some levels, matrices might be populated predominantly by `no_value`-identifiers. This will perhaps, at first, seem undesirable from a user's standpoint, but how the data is used is a problem that belongs at the application level, and not at the data collection, managing, or distribution levels. In the end, having data sets where deficiencies are highlighted, should allow for more judicial use of the data.

In summary: The NVRG is a global grid model with banded structure, and grid cells of nearly uniform size and shape. To each cell is associated a sequence of matrices. The size of a matrix is $2^n \times 2^n$ subcells. The values assigned to subcells are area values and not point values. Not all levels need to be present, and matrices can be populated by mostly `no_value`-identifiers. A user can easily import, e.g., classified resolution levels, into the system in the form of a "special purpose" model.

3 The NVRG/Sarepta system

The NVRG/Sarepta is a system for creating, storing, updating and serving terrain models. It is not designed for performing complicated searches or data analysis. The system supports several different models, but individual model files are not meant to be modified by users.

The processes and components of NVRG/Sarepta are shown in Fig. 4. We describe how to generate model files from survey data, as this is the most elaborate model generation scheme.

The event chain is as follows: The NVRG/Sarepta reads output files from the bathymetric post processing system in the form of corrected height/depth values, (`xyz`-values). Model source files are then generated for each affected cell. Depending on coverage and sounding density, model files of various request and resolution levels are generated.

If possible, only one model file is created for each area. The resolution should be as high as possible, and will thus depend on both the depth and the collection equipment. Coarser levels should preferably be generated on the fly. An important point here, is that the procedures that map block mean/min/max values from fine to coarser levels, are projection operators. As long as there is full coverage at some level, coarser scale values are uniquely defined. Whether projecting level by level, or directly, the resulting value is the same.

3.1 Preprocessing

The NVRG/Sarepta system takes filtered and tide corrected echo sounder data as input. The first processing step is data binning, where each measurement is assigned to an NVRG cell. Data from a survey will thus be split across several NVRG cells, and a cell can contain soundings from several surveys. During the next step, the xy-coordinates are quantized, which in effect maps every point to a standard local coordinate system. The resulting data sets are stored in so-called `.ijz` files. In order to increase the efficiency of the modeling algorithms, the local coordinates are in this file format represented by an index that allows subcell assignment using bitshifts.

3.2 Modeling

The coverage and resolution of the `.ijz`-files determines the block size and the resolution of the model files. Multiple resolution levels can be built covering the same area, and different block sizes can be used to model a given resolution level. However, in order to simplify searches and minimize storage overhead, for each resolution level, we only allow one model file to cover a given area. In the actual model building, the crucial step is the transformation from point to area values that takes place at the finest scale. A simple solution is to use the observed mean, but more robust methods are sometimes required. The coarser scales can be generated level-by-level by averaging over 2×2 subcells, allowing one or perhaps two of the subcells to have undefined values.

3.3 Web service

The most common way to access the NVRG models is through the Sarepta web service. Web services (WS) is essentially a means to achieve interoperability over a network between software applications running on different computers. It is thus similar to other technologies for distributed computing such as CORBA and Java RMI (Remote Method Invocation).

Although the definition of WS is broad and encompasses several types of applications, [W3C04], we here use the concept of a service as similar to an ordinary application programming interface, i.e., as a set of callable methods. The service provider will implement a set of functions and publish them on a server. These functions may then be called remotely in software running on other machines. WS uses the SOAP protocol, [W3C03], for exchanging messages between computers, and the Web Services Description Language (WSDL), [W3C01], for describing the service. SOAP messages are transmitted on top of a network protocol such as HTTP.

Both SOAP and WSDL are open, XML-based standards, which ensures that interoperability using WS is inherently platform and language independent. In the NVRG/Sarepta system, data processing and messaging is implemented in C# .NET, but the Sarepta service may be accessed using Matlab or Java or in principle any other language. In fact, one of the advantages of standardization is that most of the messaging code can be generated automatically for the application programmer. As a result NVRG/Sarepta may easily be incorporated into programs that need terrain data input. A concrete example of this is an acoustic propagation model running at the Norwegian Defence Research Establishment.

In the default mode, the Sarepta service offers methods to extract point data, terrain profiles or arrays of gridded data from single models. Various interpolation schemes can be applied. Output data may be generated at a fixed resolution or using the best resolution available for an area. Due to the simple, hierarchical structure of the NVRG, it is also straightforward to extract a subset of the database (an area of interest) and run a copy of the NVRG/Sarepta system on a stand-alone computer or on a different network.

4 Discussion and Conclusions

The NVRG/Sarepta system is presented and its design criteria are reviewed. Since the system is based on matrix representation of the earth's surface domain, the design requirement of "a terrain model which is easy to understand and implement", is met. The tiling scheme, which is termed the NVRG, supports infinitely many resolution levels. A set of request areas is also defined. These definitions enable search for data inside user defined windows at specified resolutions. The fundamental unit in the grid is the cell. The values assigned to the cells and subcells, are area values like `cell_mean`, `cell_maximum`, and `cell_minimum`. Other kinds of area values can be defined. This representation is a suitable basis for various products. Users with different data quality requirements can either process true-resolution models themselves, or provide algorithms that can be implemented as part of Sarepta. The NVRG/Sarepta system thus supports several applications from one single architecture.

NVRG/Sarepta is not implemented as a relational data base system, but as a file collection. If future applications of the system will benefit from a data base implementation, the partitioning of the data volume imposed by the NVRG tiling scheme, supports such a migration.

By basing the interface on web service technology, the system can operate on different platforms. In addition to the system installed at the Norwegian Defence Research Establishment, copies of the NVRG/Sarepta system that allow users to import their own models into the system, can run on stand-alone computers or on different networks.

Acknowledgments

This research is supported by the FFI-project POSEIDON.

References

[Bjø88] Jan T. Bjørke. Quadtrees and triangulation in digital elevation models. In *International Archives of Photogrammetry and Remote Sensing*, pages 38–44. International Society for Photogrammetry and Remote Sensing, Committee of the 16th International Congress of ISPRS, 1988. Volume 27, part B4, Commission IV.

[BGHN03] Jan T. Bjørke,John K. Grytten, Morten Hæger and Stein Nilsen. A Global Grid Based on Constant Area Quadrilaterals. *Proceedings of ScanGIS 2003*, 2003.

[BN04] Jan T. Bjørke and Stein Nilsen. Examination of a Constant Area Quadrilateral Grid in Representation of Global Digital Terrain Models. *International Journal of Geographical Information Science*, 18(7):653–664, 2004.

[BN07] Jan T. Bjørke and Stein Nilsen. Computation of Random Errors in Digital Terrain Models. *GeoInformatica*, online, 2007.

[DFMP00] Leila De Floriani, Paola Magillo, and Enrico Puppo. Variant: A system for terrain modeling at variable resolution. *GeoInformatica*, 4(3):287–315, 2000.

[DFP95] Leila De Floriani and Enrico Puppo. Hierarchical triangulation for multiresolution surface description. *ACM Transactions on Graphics*, 14(4):363–411, October 1995.

[GG98] Volker Gaede and Oliver Günther. Multidimensional access methods. *ACM Computing Surveys*, 30(2):170–231, 1998.

[Hel90] Martin Heller. Triangulation algorithms for adaptive terrain modelling. In *Proceedings of the 4th International Symposium on Spatial Data Handling*, pages 163–174, July 1990.

[McL76] D.H. McLain. Two dimensional interpolation from random data. *The Computer Journal*, 19(2):178–181, 1976.

[Mid93] Terje Midtbø. *Spatial Modelling by Delaunay Networks of Two and Three Dimensions*. PhD thesis, The Norwegian Institute of Technology, Department of Surveying and Mapping, N-7034 Trondheim, February 1993.

[MR80] M.J. McCullagh and C.G. Ross. Delaunay triangulation of a random data set for isarithmic mapping. *The Cartographic Journal*, 17(2):93–99, des 1980.

[W3C01] W3C. W3C note: Web services description language (wsdl) 1.1. http://www.w3.org/TR/wsdl, 15 March 2001.

[W3C03] W3C. W3C recommendation: Soap version 1.2 part 1. http://www.w3.org/TR/soap12-part1, 24 June 2003.

[W3C04] W3C. W3C working group note: Web services architecture. http://www.w3.org/TR/ws-arch, 11 February 2004.

[WJ92] Mark Ware and Christopher B. Jones. A multiresolution topographic surface database. *International Journal of Geographical Information Systems*, 6(6):479–496, 1992.

[Yoë77] P. Yoëli. Computer executed interpolation of contours into arrays of randomly distributed height-points. *The Cartographic Journal*, 14(2):103–108, des 1977.

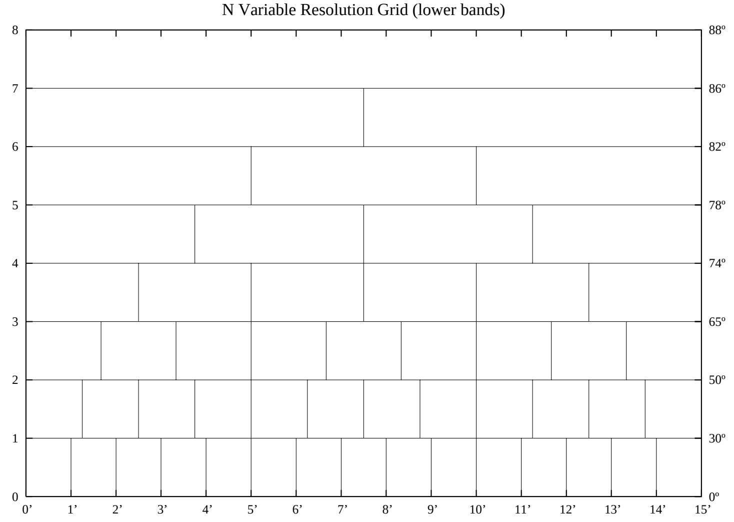

Fig. 2. NVRG tiling scheme, lower bands. The vertical axis/horizontal lines show the location of the band boundaries, and the band height. The horizontal axis/vertical lines show cell widths. The unbroken vertical lines from 0° to 74° illustrate that the grid supports 5′ blocks up to 74°.

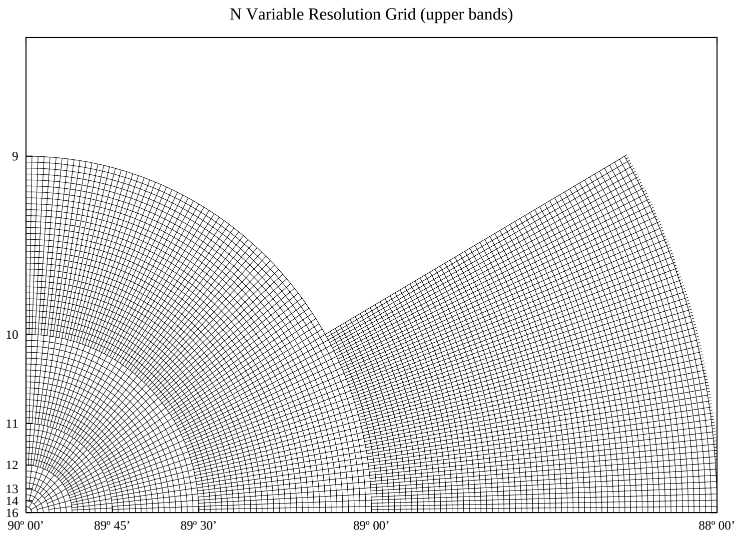

Fig. 3. NVRG tiling scheme. Cell structure of upper bands.

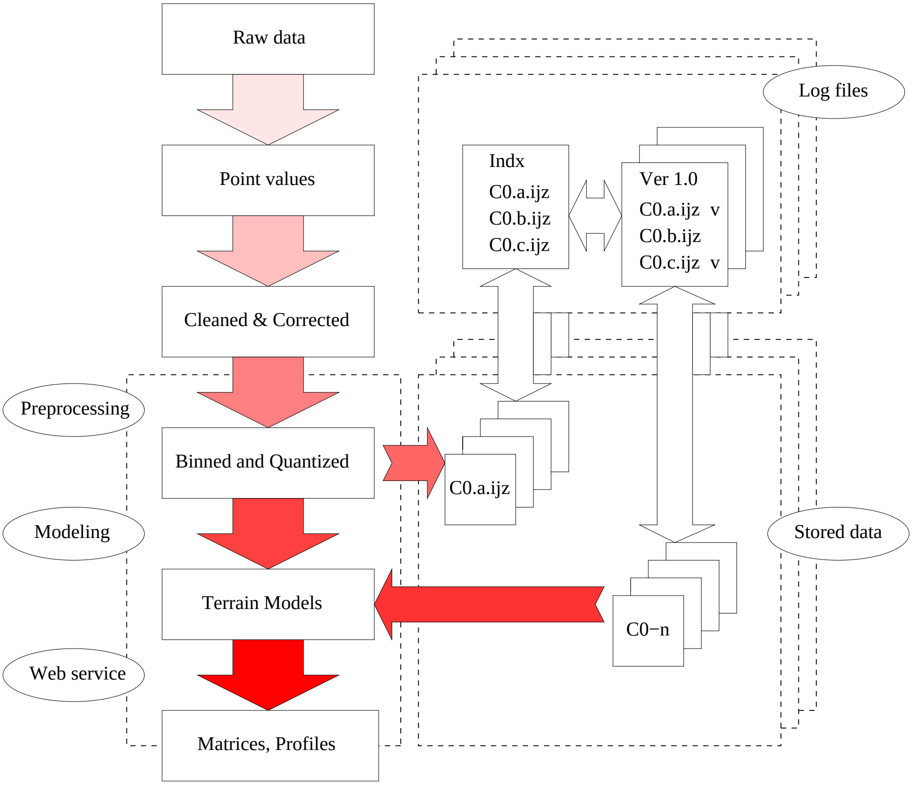

Fig. 4. Components and processes of the NVRG/Sarepta system.