Generic Tools For Generalising Ordnance Survey Base Scale Landcover Data

Patrick Revell
patrick.revell@ordnancesurvey.co.uk
Ordnance Survey®, Research, Romsey Road, Southampton, UK

Abstract

This paper presents a set of generic landcover generalisation tools. These were developed during a year long project generalising Ordnance Survey base scale data to 1:10 000 scale. A generic method for reclassifying landcover data is presented, which is most suited to specifications permitting combinations of landcover types. A suite of topological generalisation tools are described, which are applicable to any dataset representing a polygonal subdivision of the plane. The generalisation results are symbolised using a group of flexible landcover symbol placement tools.

KEYWORDS: Landcover, Generalisation, Reclassification, Amalgamation, Topology, Symbolisation.

1 Introduction

In 2001 Ordnance Survey created the OS MasterMap Topography Layer. This vector product forms a complete coverage of Great Britain at a high level of detail; being surveyed at 1:1250 scale in urban areas, 1:2500 scale in rural areas and 1:10 000 scale in mountain and moorland areas. OS MasterMap is supplied from a seamless national large scale database and is constantly kept up to date with the latest changes (Ordnance Survey 2007a). Ordnance Survey Research are developing tools which ultimately will allow all current and future products to be derived from this database with minimal manual intervention.

A two year research project investigating automated re-creation of the current 1:50 000 scale product completed at the end of 2005 (Revell et. al. 2006). During this project it became clear that generalisation tools developed to solve specific problems should strive to be as generic and flexible as possible. Developing such tools requires more initial effort, but in the long run saves time when the tools are reused for other purposes. During 2006 research work moved on to investigate 1:10 000 scale (Ordnance Survey 2007b), completing in the first quarter of 2007. This project had an even stronger focus on developing generic tools. In addition to the landcover generalisation tools presented in this paper, this project investigated generic techniques for displacing roads, generalising urban blocks and partitioning the entire country.

The landcover generalisation work was split into three main research areas. Section 2 describes a generic method for reclassifying Ordnance Survey landcover data. Section 3 details tools for performing topologically consistent model generalisation of polygon data. Section 4 covers a flexible method for symbolising landcover data. All programming was carried out in Java. Clarity generalisation software from 1Spatial (formerly Laser-Scan) was used as the research platform (1Spatial 2007). The research prototypes developed during the first and last work packages are not reliant upon Clarity APIs and could easily be ported to work with other systems.

2 Reclassification tools

2.1 The source and target data models

Ordnance Survey has a very detailed source landcover data classification, which distinguishes between 40 different *landcover types*. The data model permits combinations of these landcover types, for example “Boulders + Sand + Shingle”. Certain nonsensical combinations are not permitted, such as “Boulders + Orchard”. The total number of valid source *landcover combinations* is 470. In some rare cases a combination can involve up to six individual landcover types. There is a feature code for each landcover combination and this code is stored as an attribute on polygon data.

For the current Ordnance Survey 1:10 000 scale product there are 19 landcover types. These are listed in **Table 1**. Some source landcover types are simply not shown, or can be combined together in a composite landcover type. In some rare cases a single source type can correspond to two target types, for example “Sand Dunes” translates to “RoughGrass + Sand” in the target. Landcover combinations are permitted in the target specification, but with the restriction that a combination can contain no more than three landcover types. The total number of valid landcover combinations in the target specification is 93.

Group	Landcover Type	Abbreviation
Wooded Vegetation	Coniferous Trees	Con
	Scattered Coniferous Trees	ConScat
	Non Coniferous Trees	NonCon
	Scattered Non Coniferous Trees	NonConScat
	Scrub (i.e. bushes)	Scrub
	Orchard	Orchard
	Coppice or Osiers (types of traditionally managed woodland)	CoppOsier
Surface Vegetation	Heath	Heath
	Rough Grassland	RoughGrass
	Marsh, Reeds or Saltmarsh	MarshReeds
Rock	Boulders	Boulders
	Scattered Boulders	BouldScat
	Rock	Rock
	Scattered Rock	RockScat
	Scree	Scree
Coastal	Mud	Mud
	Sand	Sand
	Shingle	Shingle
Water	Inland, Tidal and Permanent Tidal Water	Water

Table 1. Target landcover types

2.2 A generic solution

From the above analysis, it is clear that both the landcover types and landcover combinations need to be reduced, to allow a translation from the source to the target. (Schylberg 1992a) presents a simple language developed using shells scripts, for reclassifying area features in a raster environment. However, no currently available software tools could be found to tackle this problem.

A basic approach to the problem would be to consider the 470 source combinations individually, and for each one manually select one of the 93 target combinations. This translation table could then be hard-coded as a parameter to reclassification software. However, this approach is very inflexible and would make changes to the source or target specifications difficult to implement. In addition, every time a new product was required, a new translation table would need to be set up manually. Instead, this research developed a tool which allows two landcover specifications to be connected and validated, with minimal manual effort. In addition the tool can be used for defining specifications of new products containing landcover information.

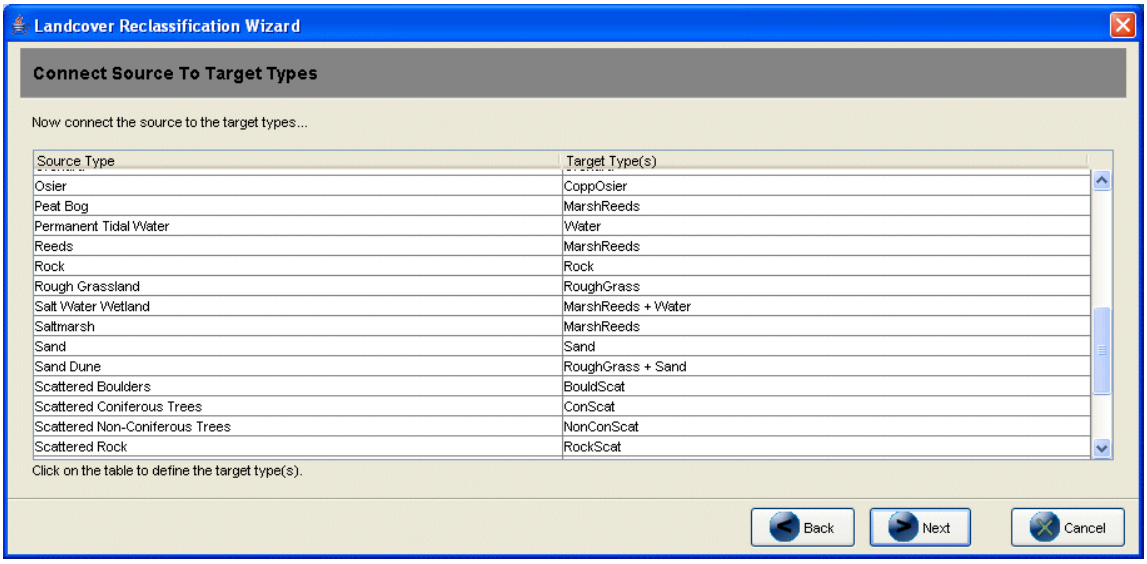

Figure 1. Connecting source and target landcover types

2.3 The landcover reclassification wizard

The tool was designed as a "wizard style" interface; i.e. a series of panels defining a sequence of steps, with a "Finish" button once the steps have been completed. The first panel asks if a new specification is to be defined or if an existing specification is to be loaded. An existing specification represents everything entered into the wizard from a previous session, and is held as XML.

The first four wizard panels load the source and target landcover types and combinations from text files. Once these source and target specifications have been defined, the next panel allows the source and target types to be linked, as shown in **Figure 1**. Clicking in the Target Type(s) column presents a pick-list allowing zero or more target types to be defined for each source type. Note that some of the translations are one-to-one rewordings, others are many-to-one, one-to-none or one-to-two.

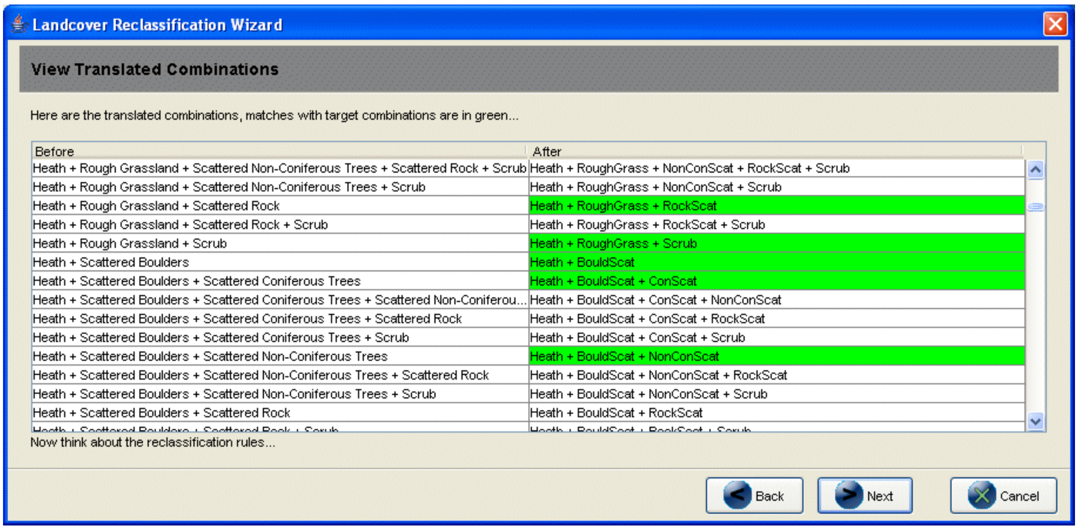

Figure 2. Translating the source combinations (valid target combinations in green)

The translation between source and target landcover types is then applied to the source combinations. In most cases this is just a simple rewording. Note that the one-to-none mappings reduce the number of landcover types in combinations, while the one-to-two mappings increase the number. Now the source combinations are *in the language of* the target specification, they are checked to see if they are present in the list of valid target combinations. The results are presented on a wizard panel, with the successful matches highlighted in green, as shown in **Figure 2**.

No further classification is applied to the landcover combinations highlighted in green. Non-highlighted rows require further classification to make them conform to the target specification. Reducing the number of landcover types in a combination is the usual method for this. In some instances it may be necessary to replace one landcover type with another, such as replacing "Rock" with "RockScat". Alternatively the interface allows non-conforming combinations to be added to the target specification.

During the research it was observed that there are many combinations to which the same reclassification rule could be applied. Therefore a means of automatically selecting such combinations was required. SQL was considered for this, but was rejected since the wizard is aimed at non-programmers. A simplistic query language was devised, with the main syntax being the logical operators AND, OR and NOT in conjunction with the names of the target landcover types. Thus the query "Heath AND Boulders" would select all non-conforming combinations containing both Heath and Boulders. Statements can be nested together with parentheses to form complex logical expressions.

During development it was noted that the same cumbersome expressions involving landcover types were required for more than one rule. Therefore *landcover type groups* were introduced. For example "Coastal" can be defined as shorthand for "Shingle OR Sand OR Mud", then used in as many rules as necessary. Sometimes a selection needs to be dependant on the number of landcover types within a combination. For this the COUNT function was introduced, along with the numerical comparison operators. For example "COUNT > 3" selects all combinations containing more than three landcover types. A *landcover type group* can be used in conjunction with the COUNT function, for example "COUNT Coastal = 2", selects all combinations containing two landcover types in the Coastal group.

Figure 3 shows the user interface used for defining these reclassification conditions. The "Matches" button reveals all combinations which match the current condition. The "Targets" button presents the list of valid combinations in the target specification.

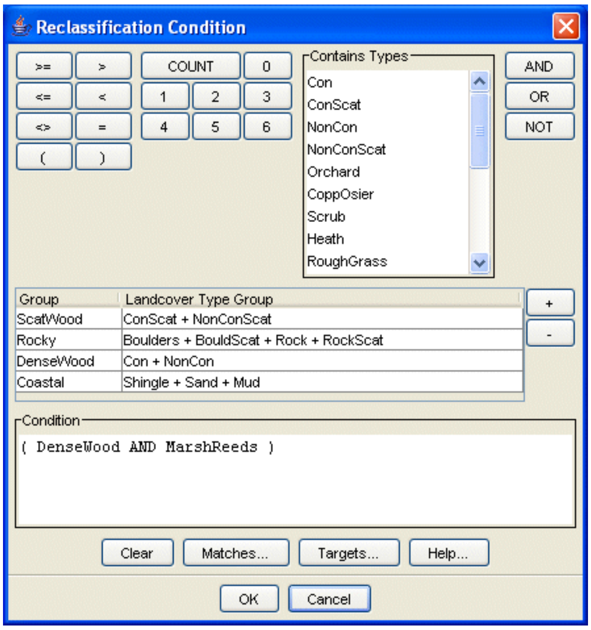

Figure 3. Reclassification Condition interface

Reclassification rules are formulated using a reclassification condition (the "If" statement), followed by a list of landcover types to remove or add. **Figure 4** shows the wizard panel in which reclassification rules are formulated. Clicking on "Add Rule" opens the Reclassification Condition interface and when OK is pressed a new row is added to the table. Clicking in the "Then Remove" column allows the definition of a list of types to be removed from combinations fulfilling the condition. Similarly clicking in the "Then Add" column allows landcover types to be added to combinations fulfilling the condition.

Note that the rules are applied sequentially, and once a landcover combination conforms to the target specification, no further rules are applied to it. This means that the order of the rules in the table matters, so "Up" and "Down" buttons are provided to allow re-ordering of the rules.

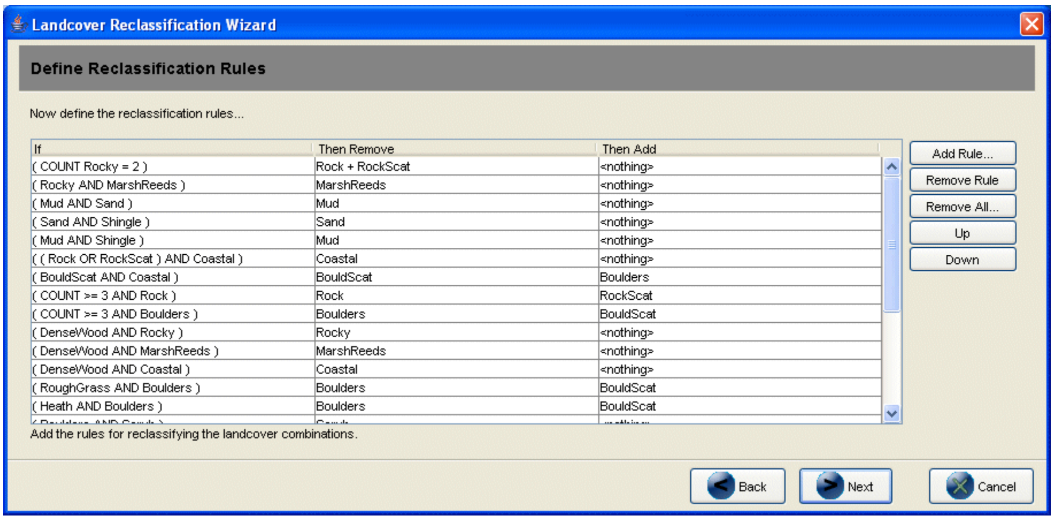

Figure 4. Defining reclassification rules

The next wizard panel shows the landcover combinations before and after the application of the reclassification rules. This is shown in **Figure 5**. Landcover combinations which are green in both columns are in the target specification and no reclassification rules were applied to them. Combinations which are white on the left and green on the right were successfully reclassified into the target specification using the reclassification rules. Combinations which are white on the left and orange on the right were reclassified, but they still do not conform to the target specification. Combinations which are white on both sides were not subjected to any rules.

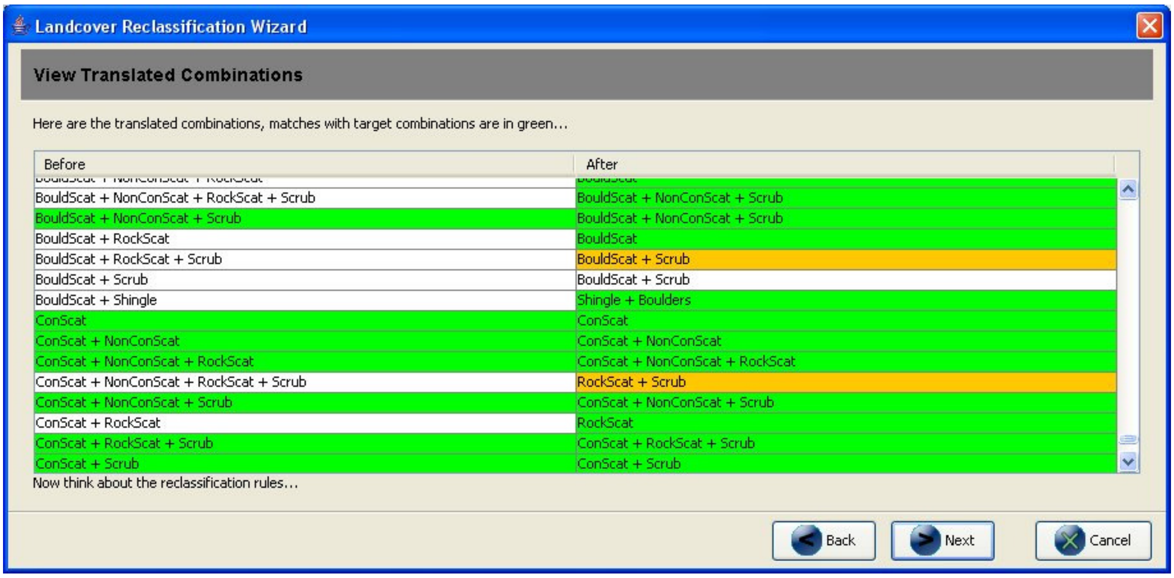

Figure 5. Viewing combinations before and after application of reclassification rules

It is not possible to move beyond this wizard panel until all of the combinations are green in the right hand column. The back button can be used to return to the reclassification rules panel, to do further work on the rules.

The final wizard panel presents all of the target combinations, and the number of source combinations which have been mapped to each of them, as shown in **Figure 6**. This is provided for the purposes of validation. Targets which correspond to a large number in the source should be treated with suspicion. Clicking on the numbers allows the source combinations to be viewed and verified.

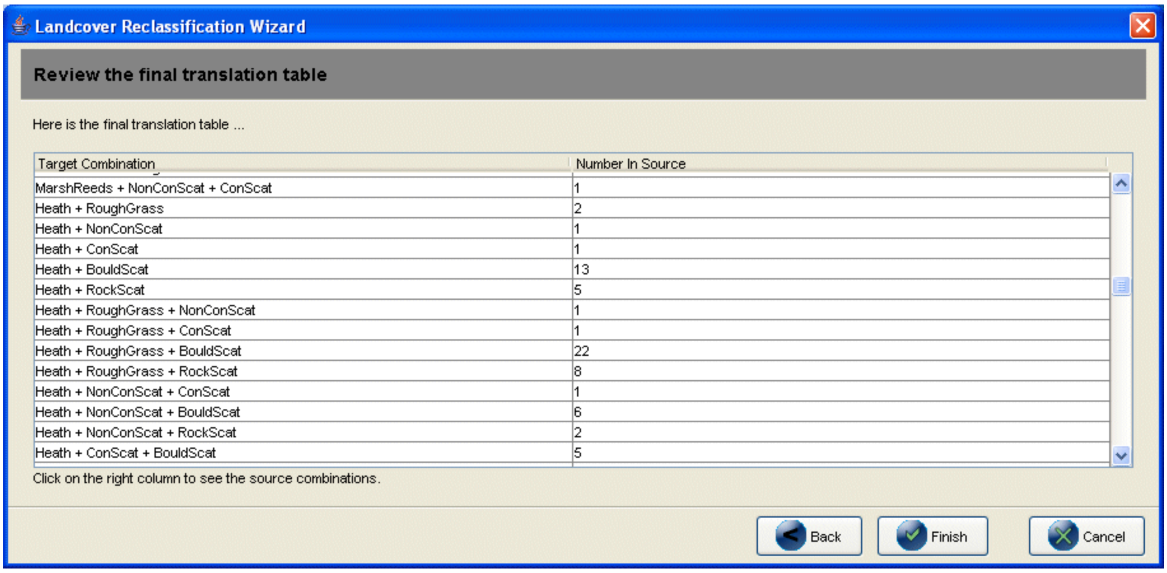

Figure 6. Reviewing the final translation

Clicking on the “Finish” button will allow all information entered into the wizard to be saved as an XML specification file for future editing. In addition a translation table, derived from the specification, can be saved to XML. This explicitly lists the mapping between source and target combinations. The translation table can then be loaded up by a landcover reclassification process. This process operates on the landcover polygon feature class and sets new attributes for the target landcover combinations.

2.4 Future extensions

If an ontology was available for the source landcover data, then developing reclassification rules could be much more automated and intuitive. For example currently the system has no understanding of the relationship between coniferous trees and scattered coniferous trees.

On the panel in **Figure 5**, it would be useful to be able to pick a landcover combination on the right side of the table and find out exactly which landcover rules have been applied to it. Adjusting the order of the rules can have quite dramatic effects on the classification, so it would be useful to have a way of visualising these effects. Also there is currently no way of knowing if a rule is redundant. When developing reclassification rules, a good knowledge of the target specification is assumed. It would be helpful to be able to select a landcover combination and for the system to automatically propose a list of suitable target combinations and suggest rules which would make it conform to them.

It would be useful for the system to summarise the reclassification results. For example it could list the number of features in each target combination, and the number of features in each source combination which have been reclassified into them.

3 Topological generalisation tools

3.1 Topology In Clarity

Topology is stored explicitly in Clarity using a link and node model. Data can either be structured directly when it is imported into Clarity, or alternatively it can be imported as spaghetti and structured as a separate process. The structured data avoids duplication by constructing real world object (**RWO**) geometry on the fly from the topological primitives. It is possible to query database references to go from an RWO to its links and conversely from a link to the RWOs which share it. Therefore it is possible to find adjacent features efficiently, avoiding the need for slow spatial querying.

Topology is also useful for modifying the data in a consistent way. It is possible to construct new line and area RWOs from existing links and nodes. In addition modifying a link geometry will automatically modify the RWOs which share it.

Ordnance Survey base scale topographic data represents a polygonal subdivision of the plane. Each polygon is totally surrounded by linear features. For example these can represent fences/walls/hedges, streams, edges of roads/tracks/paths or simply a change of landcover type. When generalising this data it is important to maintain topological consistency between the polygons and the lines.

3.2 Polygon Dissolve By Attribute

Following landcover reclassification, adjacent polygons can end up with the same attribute values. The dissolve by attribute tool is designed to search for such cases and merge them into single polygons. This extends the ArcToolbox Dissolve tool (ESRI 2007) by using specified line classes to block dissolving and deleting line features which have been dissolved across. Note that the ArcToolbox Dissolve tool additionally allows specified attributes for merged features to be summarised.

Inputs

- A class of polygon objects to process.
- A list of simple attribute types to use for the value matching (double, integer, string, boolean). If no attributes are provided, then the dissolving is unrestricted. This is useful for removing internal building divisions, for example.

- A list of line classes which cannot be dissolved across. For example fences and rivers can be used for blocking.
- An option for deleting lines which have been dissolved across.

Description

For each RWO in the class to process, the following steps are performed:

1. Values for the matching attributes are retrieved from the **current RWO**.
2. An empty set **Y** for RWOs to definitely merge with, is initialised.
3. The topology links are retrieved from the **current RWO**.
4. Two empty sets are initialised: **P** for the RWOs to potentially merge with and **N** for the RWOs not to merge with.
5. For each link, the RWOs are retrieved and for each of these RWO:
    1. If the RWO is the same as the **current RWO**, processing moves on to the next RWO.
    2. If the RWO is already in **Y** or **N**, processing moves on to the next RWO.
    3. The **class name** of the RWO is retrieved.
    4. If the **class name** is in the list of blocking line classes, all RWOs sharing the current link are added to **N** and removed from **P**. Thus a single blocking line will prevent dissolving into all adjacent polygons sharing that line. In this case processing moves on to the next link. This is illustrated in **Figure 7**.
    5. Merging only takes place within a single class. Therefore if the **class name** is not the same as the **current RWO**, the RWO is added to **N** and processing moves on to the next RWO. This can be seen in Figure 7.
    6. Merging only takes place if the specified attribute values match. Therefore if the attribute values for the RWO do not match with the **current RWO**, the RWO is added to **N** and processing moves on to the next RWO. This is shown in **Figure 7**.
    7. If the bottom of the loop is reached, the RWO is added to the potential set **P**, but may be removed later if another link blocks it.
6. Each RWO in the potential set **P** is now checked. If it is **not** already in the definite set **Y**, it is added. Each RWO in turn becomes the **current RWO** and steps 3 to 5 are repeated iteratively until all connected merge candidates are found. **Figure 7** shows two iterations.
7. If **Y** is empty at the end of the iteration, then there is nothing to merge and processing moves on to a new **current RWO**.
8. Otherwise a merge is required. This is performed topologically by first finding the **external links** (those used by only one object in **Y**) and the **internal links** (those used by more than one object in **Y**).
9. A new polygon object is formed from the **external links** and values for the matching attributes are set on it. All other attributes are discarded.
10. Optionally, any lines which use the **internal links** are deleted.
11. The connected polygon objects **Y** are deleted and are removed from the set of objects to process.

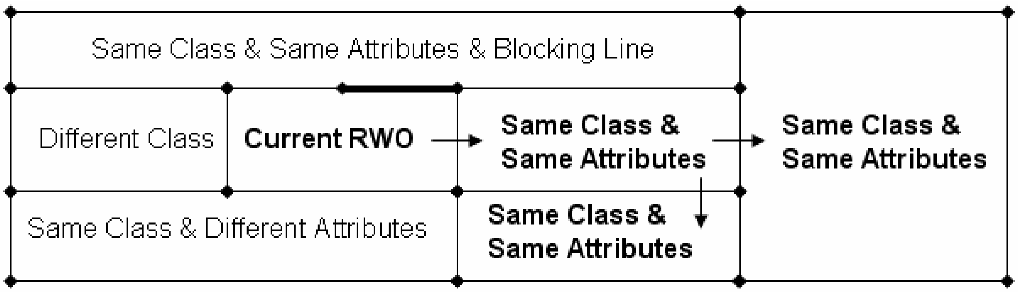

Figure 7. Dissolve by attribute algorithm. The thick black line blocks the dissolve.

Future Extensions

The tool currently works only on all RWOs in a particular class. It would be useful to extend it to operate on subsets, and prevent merging with RWOs outside the subset. This would allow the algorithm to be invoked on a partition-by-partition basis.

It is not easy to verify the results of the algorithm since it modifies the data as it goes along. For debugging it would be handy if the new area features were marked up and the deleted areas/lines were moved into new feature classes for inspection. It would also be useful for the algorithm to report on the number and total area of features which have been merged.

3.3 Polygon Dissolve Small Holes

This tool identifies and removes small holes from polygons, along with any features contained in the holes. Basic algorithms for removing holes from polygons are available in GIS, but they cannot be used to remove polygon and line features contained in the holes.

Inputs

- A class of polygon objects to process.
- The minimum hole size threshold.
- An option for deleting lines which have been dissolved across.

Description

For each RWO in the class to process, the following steps are performed:

1. The **current RWO** is tested for inner rings below size and any corresponding inner ring links are retrieved.
2. If there are no inner ring links, processing moves on to a new **current RWO**.
3. All links are retrieved from the **current RWO**, the below-size inner ring links are removed and a new RWO is constructed from this set of links.
4. All simple attributes are copied across to the new RWO.
5. Polygons and (optionally) lines sharing the inner ring links are identified and deleted. The set of objects to process is updated.

Future Extensions

The tool currently works only on all RWO in a particular class, so like the previous tool it would be useful to make it work on subsets. It would be useful to have an option which dissolves all holes, not just those below a specified threshold (the current workaround is to set a very high minimum hole size threshold). It would be useful for the algorithm to summarise how many holes have been removed and how many features have been affected.

3.4 Polygon Dissolve Small Areas

This tool identifies areas below a threshold size and dissolves them into their neighbours (possibly in other feature classes). This is similar to the ArcToolbox Eliminate tool (ESRI 2007), which only operates within a single feature class and cannot delete lines which have been dissolved across. However, the Eliminate tool does allow arbitrary selections of features to be dissolved, rather than just those below a threshold size.

Inputs

- A class of polygon objects to process.
- The minimum size threshold.
- A choice of merging with the neighbour of the largest area or the neighbour with the longest shared boundary.
- A list of polygon classes not to merge with. For landcover this will typically be manmade classes such as buildings. Similar rules are discussed in (Schylberg 1992b). In cases where all neighbours are blocked, this list can be optionally ignored and the merge forced.
- An option for deleting lines which have been dissolved across.

Description

For each RWO in the class to process, the following steps are performed:

1. If the **current RWO** is above the minimum size threshold, processing moves on to the next RWO.
2. Otherwise the topology links are obtained from the **current RWO**, and from these links a set of neighbouring RWOs is found.
3. If the feature class of a neighbouring RWO is present in the list of polygon classes not to merge with, then the neighbour is ignored. An example is shown in **Figure 8**. This rule is optionally overlooked if it would cause there to be no valid neighbours.
4. From the set of neighbouring RWOs a single neighbour must be chosen using one of the following procedures:
    - **Largest Area**. Each feature in the set of neighbouring RWOs is checked to find the neighbour with the largest area. This case is illustrated in **Figure 8**.
    - **Longest Shared Boundary**. Common links are found between the source feature and each of the neighbouring RWOs. The length of the common links is computed to find the neighbour with the longest shared boundary. This can be seen in **Figure 8**.
5. Now a merge is required between the source feature and the chosen neighbour. This is performed topologically by first finding the **external links** (those used by only one feature) and the **internal links** (those used by both features).
6. A new polygon feature is formed from the **external links** and all attributes are copied on to it from the chosen neighbour.
7. Optionally, any lines which use the **internal links** are deleted.
8. The **current RWO** and the chosen neighbour are deleted and are removed from the set of objects to process.

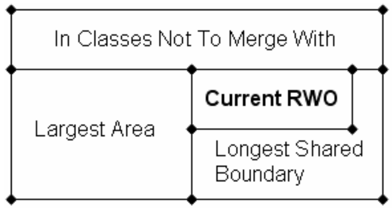

Figure 8. Dissolve small areas algorithm.

Future Extensions

The tool currently works only on all RWO in a particular class, so like the previous tool it would be useful to make it work on subsets.

An option could be added to dissolve every feature, not just those below threshold (the current workaround is to set a very high minimum size threshold). This would be useful for a global elimination of a feature class. It would be helpful if the algorithm could provide a summary of how many below size features have been merged and their total area.

It would also be useful to add a new option which divides up the features below threshold evenly between all surrounding neighbours, using a skeleton. An algorithm to perform this task was described by Bader and Weibel (1997).

3.5 Polygon Topological Simplification

This tool simplifies polygons in such a way that there are no slivers or overlaps between polygons and all lines remain consistent with the polygons. This is similar to the ArcToolbox Simplify Polygon tool (ESRI 2007), which applies either a POINT_REMOVE or BEND_SIMPLIFY algorithms to polygon data. The ESRI tool does preserve shared boundaries within a polygon class, but not with other polygon classes or between polygons and lines. Additionally it can collapse simplified zero-area polygons to points and resolve topological errors if required.

Inputs

- A class of polygon objects to process.
- The simplification tolerance.
- A list of area and line classes not to simplify. For example buildings.

Description

For each RWO in the class to process, the following steps are performed

1. The topology links are retrieved from the **current RWO**, and are processed one at a time.
2. If a link is shared by a RWO whose class is present in the classes not to simplify, then processing moves on to the next link.
3. The Douglas Peucker (1973) simplification algorithm is applied to the link. If the result self-intersects, the simplification is applied with half the tolerance. This is repeated until a non-self-intersecting result is obtained, but after ten attempts have been made this will abort.
4. The simplified geometry is now set on the link, which automatically updates all connected polygon and line features.
5. The topology update operation will fail if the modified link intersects with another link. In this case the Douglas Peucker algorithm is applied with half the tolerance and update is attempted again. This is repeated until the result will commit, but after ten attempts have been made this will abort.

Future Extensions

The tool currently works only on all RWO in a particular class, so like the previous tool it would be useful to make it work on subsets.

The approach could be extended by allowing other simplification algorithms to be selected instead of Douglas Peucker, for example Visvalingham-Whyatt or Weighted Effective Area (Zhou and Jones 2004). The tool would also be more flexible if a simplification tolerance could be specified for each feature class. Then for any given link, the minimum tolerance of all the features sharing the link would be selected. For validation it would be useful to record those links where simplification fails, or succeeds with a reduced tolerance. Some self-intersections could easily be resolved automatically by clipping off small closed loops on the invalid line, instead of resorting to a reduced tolerance.

3.6 Results

The overall results of applying the topological generalisation tools can be seen in **Figure 9**. Some polygons have been merged together following reclassification. Small holes and small polygons have been eliminated and polygon boundaries have been simplified. It would be useful in the future to have some summary statistics about the total area and number of features in each target landcover combination and how these correspond to the total area and number of features in each source combination.

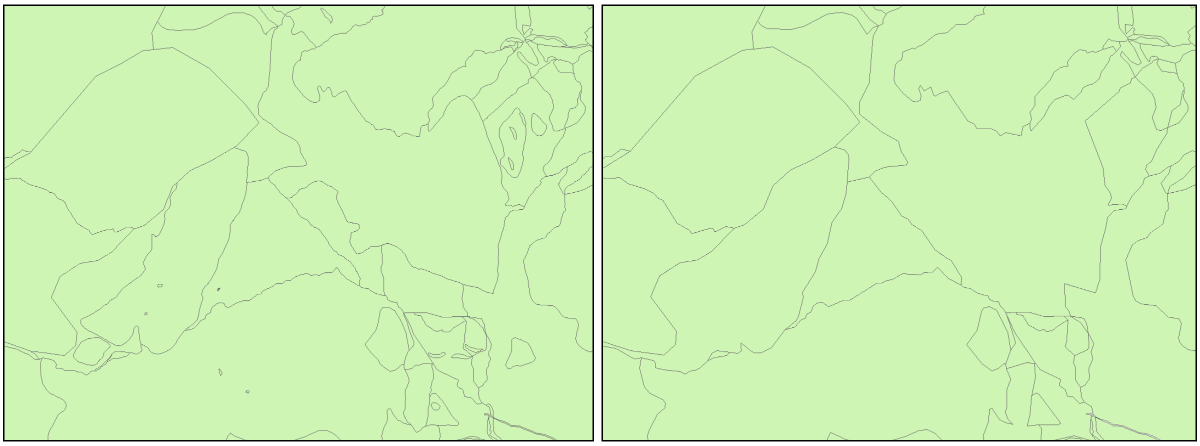

Figure 9. Source data (left) and the results from the topological generalisation tools (right).

3.7 Future Tools

During this research no tools to resolve polygon proximity conflicts were investigated. This is because the jump in scale from the source data to 1:10 000 scale is not that great. The tools could be used for generalising to 1:25 000 scale, which has a very similar specification, but with the added requirement that proximity conflicts must be resolved.

In some cases the proximity conflicts need to be resolved by local enlargement or displacement. An approach for this is described by Bader and Weibel (1997). This would be useful for widening long thin polygon features such as rivers. Collapsing to centrelines is another technique which could be applied to resolve such conflicts. In this approach a skeleton could be used, and the divided segments of the collapsed polygon would be shared amongst the neighbouring polygons (similarly to elimination). This technique was also presented by Jones et. al. (1995).

Bader and Weibel describe an amalgamation algorithm which also utilises a skeleton. When there is narrow portion of polygon between two amalgamation candidates, a skeleton is built inside the narrow polygon. The two amalgamation candidates are extended to meet the skeleton and are merged together. This type of amalgamation could be used for generalising woodland. During the 1:50 000 scale generalisation project a woodland amalgamation algorithm was developed using a triangulation (Revell 2005). This could be made more efficient by first applying a skeleton to the thin woodland path and track polygons which need to be amalgamated across.

4 Cartographic symbolisation of results

Once the model generalisation is complete, the final stage for the landcover generalisation is cartographic symbolisation. In the current 1:10 000 scale production system, a finishing process applies different vegetation and rock symbol “wallpaper” patterns for each of the 93 landcover combinations. Thus the specification is very difficult to change. The patterns do not look cartographically pleasing when the symbols are broken by the perimeters of the polygons. Small polygons can end up with no symbols at all, leading to ambiguity. The possibility of improving the landcover depiction was investigated by extending and adapting vegetation symbol placement algorithms developed for 1:50 000 scale (Harrie and Revell 2007).

Specifically there is a diagonal grid and a semi-random placement algorithm, which are both designed to avoid broken symbols around the polygon boundaries. The semi-random placement can handle up to four different types of symbol in a single polygon. Parameters can be stored in an XML *symbol placement pattern* file. This controls the symbol density, and in the semi-random case, the amount of deviation from a diagonal grid. Generic user interfaces were developed to allow the symbolisation to be easily customised. These interfaces are described in more detail in (Revell 2007).

A symbol library can be loaded from a Scalable Vector Graphics (SVG) file. Each landcover type is then associated with zero or more symbols from the library, along with a symbol scale factor controlling the symbol size, a symbol colour and a background fill colour. For example, coniferous trees are represented by black tree symbols on a green background. This *symbol type setup* is stored as XML.

A symbol placement pattern XML file is then associated with each landcover combination. The same placement pattern is frequently reused for many of the combinations. *The symbol combination setup* is also stored as XML. Finally for each landcover combination, the system automatically selects which landcover symbols to use and the pattern to employ for their placement.

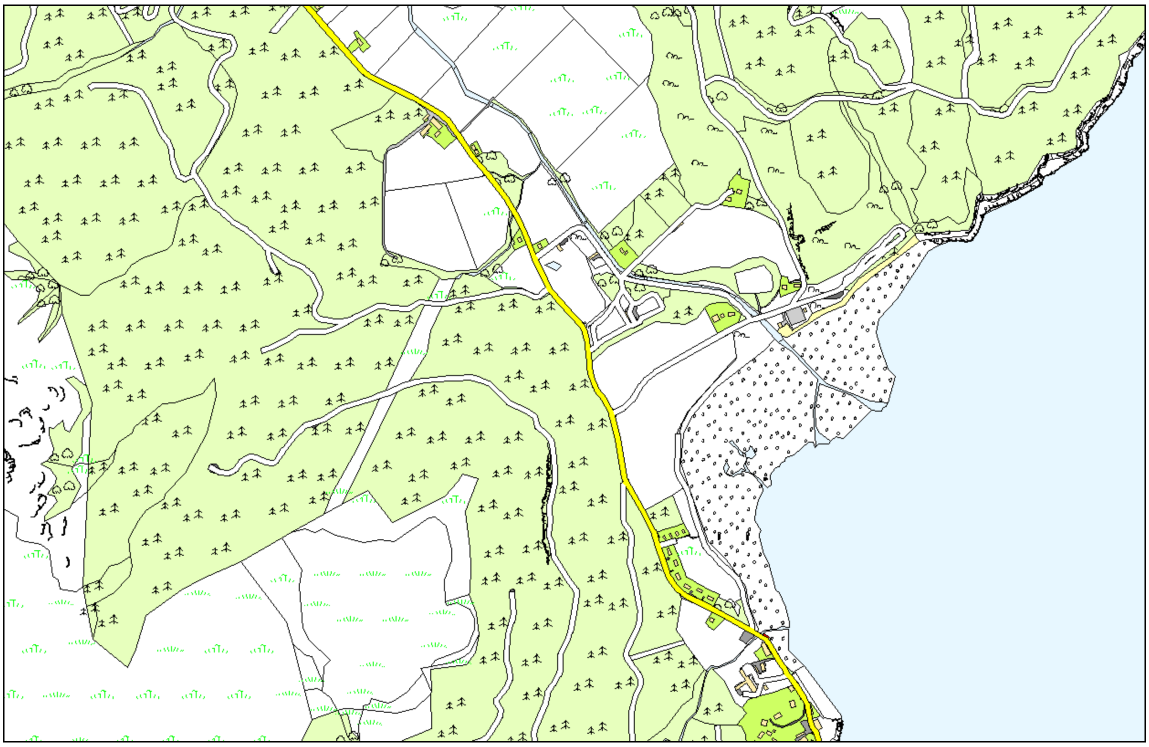

Figure 10. Results of automated generalisation and symbol placement.

The finished results are shown in **Figure 10** and **Figure 11**. The prominent black graphics represent hand-drawn rock detail, stored as a polygon layer. The symbol placement algorithms were tuned to avoid placing symbols on top of the hand-drawn rock. At the reclassification stage an automatic process prevented clashes between rock symbols and the hand-drawn rock. More detailed information about this is given in (Revell 2007).

Note that the symbol placement algorithms always force the placement of at least one symbol per polygon, even when it causes the symbol to protrude beyond the polygon boundary. It would be better to have the option of placing an alternative smaller symbol or no symbol at all.

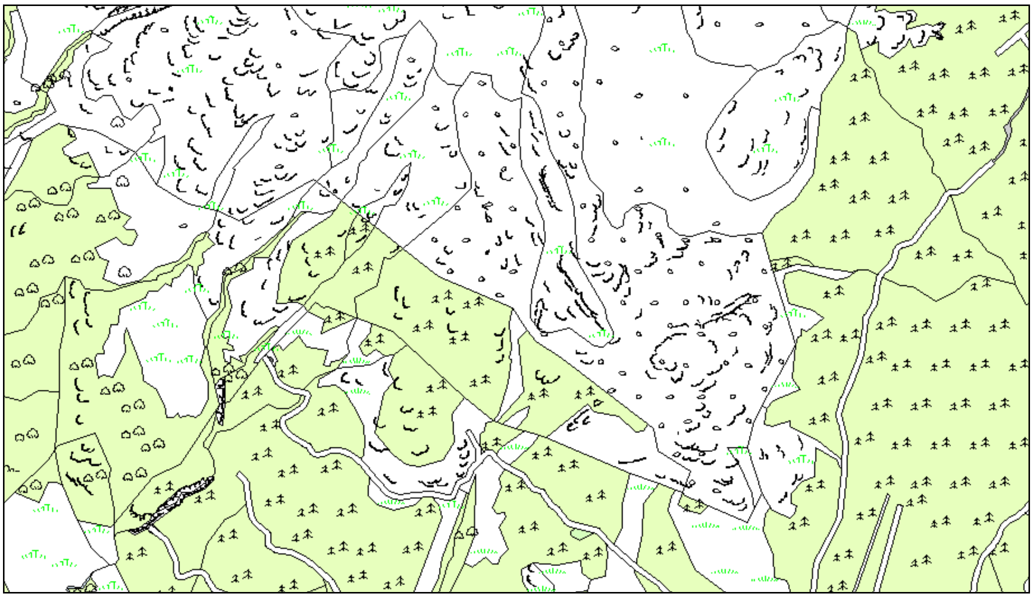

Figure 11. Results of automated generalisation and symbol placement.

5 Conclusions

A set of tools for generalising and symbolising landcover data have been presented. The tools were developed in the context of a 1:10 000 scale generalisation project, but were designed to be as flexible and generic as possible. The landcover reclassification wizard could be used for reclassifying any landcover data, but would be most useful for specifications which permit combinations of landcover types.

The topological generalisation tools are also generic and could be implemented on any database which supports a link-node topological model. The tools are applicable to any dataset representing a polygonal subdivision of the plane. They are especially useful for maintaining consistency when the polygons are bounded by line features. The results of the model generalisation are symbolised using a set of flexible automatic symbol placement tools, which could easily be adapted to work with other landcover data.

The work carried out by this project has demonstrated that it is possible to develop reusable tools, while working within the constraints of specific requirements. Development does take slightly longer, but benefits are reaped by subsequent projects. Research at Ordnance Survey is now moving towards long term development of a generalisation system which can create products defined by arbitrary specifications. At the same time research must respond to short term requirements, such as a current project investigating the replacement of mid-scale production systems. The challenge is to fulfil the short-term requirements without compromising the long term aims.

References

1SPATIAL (2007), *Radius Clarity – a rules based environment for automated generalisation*, 1Spatial website. http://www.1spatial.com/products/radius_clarity/

BADER, M. and WEIBEL, R. (1997), *Detecting and resolving size and proximity conflicts on the generalisation of polygonal maps*, Proceedings 18th International Cartographic Conference, Stockholm.

DOUGLAS, D., PEUCKER, T. (1973), *Algorithms for the reduction of the number of points required to represent a digitised line or its caricature*, The Canadian Cartographer, Vol. 10, No. 2, pp. 112-122.

ESRI (2007), *ArcGIS 9.2 desktop, ArcToolbox online help, Dissolve, Eliminate and Simplify Polygon*. http://webhelp.esri.com/arcgisdesktop/9.2/index.cfm?TopicName=Dissolve_(Data_Management) http://webhelp.esri.com/arcgisdesktop/9.2/index.cfm?TopicName=Eliminate_(Data_Management) http://webhelp.esri.com/arcgisdesktop/9.2/index.cfm?TopicName=Simplify_Polygon_(Data_Management)

FORREST D., PEARSON, A., COLLIER P. (1997), *The representation of topographic information on maps- the coastal environment*, The Cartographic Journal, Vol. 34, No. 2, pp. 77 - 85.

FORREST D., PEARSON, A., COLLIER P. (1998), *The representation of topographic information on maps- vegetation and rural land use*, The Cartographic Journal, Vol. 35, No. 2, pp. 191-197.

HARRIE L., REVELL, P. (2007), *Automation of vegetation symbol placement on Ordnance Survey 1:50 000 scale maps*, to be published in the Cartographic Journal.

JONES C., BUNDY G., WARE, M. (1995), *Map generalisation with a triangulated data structure*, Cartography and Geographic Information Systems, Vol. 22, No. 4, pp. 317-331.

ORDNANCE SURVEY (2007a), *OS MasterMap®: definitive digital map of Great Britain designed by Ordnance Survey*, Ordnance Survey website. http://www.ordnancesurvey.co.uk/oswebsite/products/osmastermap/

ORDNANCE SURVEY (2007b), *1:10 000 Scale Raster: mid-scale high-resolution detailed mapping*, Ordnance Survey website. http://www.ordnancesurvey.co.uk/oswebsite/products/10kraster/

REVELL, P. (2005), Seeing the wood from the trees: *Generalising OS MasterMap tree coverage polygons to woodland at 1:50 000 scale*, presented at the 8th ICA Workshop on Generalisation and Multiple Representation, La Coruña, Spain. http://ica.ign.fr/Acoruna/Papers/Revell.pdf

REVELL, P., REGNAULD, N., THOM, S. (2006), *Generalising and symbolising Ordnance Survey base scale data to create a prototype 1:50 000 scale vector map*, presented at the 9th ICA Workshop on Generalisation and Multiple Representation, Portland, Oregon. http://ica.ign.fr/Portland/paper/ICA2006-Revell.pdf

REVELL, P (2007), *Automated Generalisation and Representation of Ordnance Survey Polygonal Landcover Data at 1:10 000 Scale*, to be presented at the XXIII International Cartographic Conference, Moscow, Russia, 4th to10th August 2007.

SCHYLBERG, L., (1992a), *Knowledge based area generalisation of topographic raster data*.

SCHYLBERG, L., (1992b), *Cartographic amalgamation of area objects*, The XVIIth International Congress for Photogrammetry and Remote Sensing, International Archives of Photogrammetry and Remote Sensing, Washington, D.C., Commission IV, XXIX(B4), pp. 135 - 138.

ZHOU, S., JONES, C. (2005), *Shape-Aware Line Generalisation With Weighted Effective Area*, Developments in Spatial Data Handling, Editor Peter Fisher, Springer, pp. 369 - 380.