Geoscientific
Model Development

Upscaling with the dynamic two-layer classification concept (D2C): TreeMig-2L, an efficient implementation of the forest-landscape model TreeMig

J. E. M. S. Nabel[1,2,a]

[1]Department of Environmental Systems Science, Swiss Federal Institute of Technology ETH, 8092 Zürich, Switzerland
[2]Dynamic Macroecology, Landscape Dynamics, Swiss Federal Institute for Forest, Snow and Landscape Research WSL, Zürcherstrasse 111, 8903 Birmensdorf, Switzerland
[a]now at: Max Planck Institute for Meteorology, Bundesstrasse 53, 20146 Hamburg, Germany

Correspondence to: J. E. M. S. Nabel (julia.nabel@mpimet.mpg.de, jemsnabel@gmail.com)

Abstract. Models used to investigate impacts of climatic changes on spatio-temporal vegetation dynamics need to balance required accuracy with computational feasibility. To enhance the computational efficiency of these models, upscaling methods are required that maintain key fine-scale processes influencing vegetation dynamics. In this paper, an adjustable method – the dynamic two-layer classification concept (D2C) – for the upscaling of time- and space-discrete models is presented. D2C aims to separate potentially repetitive calculations from those specific to single grid cells. The underlying idea is to extract processes that do not require information about a grid cell's neighbourhood to a reduced-size non-spatial layer, which is dynamically coupled to the original two-dimensional layer. The size of the non-spatial layer is thereby adaptive and depends on dynamic classifications according to pre-specified similarity criteria.

I present how D2C can be used in a model implementation on the example of TreeMig-2L, a new, efficient version of the intermediate-complexity forest-landscape model TreeMig. To discuss the trade-off between computational expenses and accuracy, as well as the applicability of D2C, I compare different model stages of TreeMig-2L via simulations of two different application scenarios. This comparison of different model stages demonstrates that applying D2C can strongly reduce computational expenses of processes calculated on the new non-spatial layer. D2C is thus a valuable upscaling method for models and applications in which processes requiring information about the neighbourhood constitute the minor share of the overall computational expenses.

1 Introduction

Impact studies of climatic changes on spatio-temporal vegetation dynamics are often conducted with so-called dynamic vegetation models (DVMs). DVMs are mainly implemented as time- and space-discrete models, simulating ecological processes that are key to vegetation dynamics, such as establishment, growth and mortality, usually under consideration of biotic and abiotic influences (see e.g. Snell et al., 2014). As all models do, DVMs need to balance accuracy with computational feasibility and parametrisation requirements (Huntley et al., 2010; He et al., 2011). Modelled processes and their level of detail vary among DVMs, with a close link to the trade-off between spatial resolution and spatial extent of the simulation area. DVMs simulating small-scale processes with a fine spatial resolution ($< 1\,\mathrm{km}^2$) often have large computational expenses. Therefore, they can only operate on much smaller spatial extents than models with coarser resolution that neglect small-scale heterogeneity (see examples listed in Snell et al., 2014). Although there is a steady increase in the spatial extents that small-scale models can be applied on, due to increasing computational capacities and cost reductions via pure computational methods (e.g. parallelisation techniques), there is still a gap in what can be studied with small- and with large-scale models. One of the main problems is that the spatial resolution on which a process is simulated can markedly influence simulation results. Using a coarser spatial resolution to enable an increase in the spatial extent therefore risks introducing strong biases, such as replacing rare with dominant forest types (He et al., 2011), or

overestimating dispersal distances and population sizes (Bocedi et al., 2012). Thus, there is a need to develop upscaling methods that maintain the required fine resolution for key small-scale processes (Bocedi et al., 2012).

Many upscaling methods have been proposed and applied in the context of ecological modelling (see e.g. Urban et al., 1999; Lischke et al., 2007; Auger et al., 2012). Most of these methods seek to aggregate small-scale information to a coarser scale. Aggregations can thereby be temporal, spatial or with regards to the representation of the modelled entities, i.e. the simulated processes and state variables. When the fine spatio-temporal resolution should be maintained (as recommended by Bocedi et al., 2012), remaining possibilities for cost reductions thus can only change the representation of the modelled entities. One example finding broad application in DVMs is the aggregation of individuals with similar properties into cohorts (e.g. Scheller and Mladenoff, 2004; Smith et al., 2014). With cohorts solely one representative calculation, instead of multiple replicate calculations, needs to be conducted. Another more comprehensive example is the approximation of entire grid cell populations. Height-structured distributions (Lischke et al., 1998), as well as size- and age-structured approximations (Moorcroft et al., 2001), can be applied to upscale individual-based, stochastic models. By aggregating individuals of forest patches influenced by small-scale stochastic processes, such approximations allow one to dramatically reduce the number of calculations required to determine the vegetation dynamics in a grid cell. Both examples, cohorts and structured approximations, utilise similarities for within grid cell aggregations. The upscaling method presented in this paper – the dynamic two-layer classification concept (D2C) – is based on a more extensive similarity approach aiming to combine different grid cells with similar properties.

The motivation for D2C stems from the assumption that spatially independent locations experiencing similar environmental conditions likely accommodate similar vegetation compositions. Similar vegetation compositions, in turn, can lead to similar, potentially redundant calculations in simulations with DVMs. Especially time- and space-discrete DVMs will entail replicate calculations when different grid cells share similar state variables and drivers for long periods of time. D2C aims to avoid replicate calculations while still allowing for diverging and novel vegetation compositions. For this purpose, grid cells are dynamically classified into groups with similar properties, for which subsequently only one representative calculation is conducted. The fact that time- and space-discrete DVMs can entail replicate calculations for similar vegetation compositions has already been used to decrease computational expenses in finite state individual based models with a simple age based succession (Yang et al., 2011) and for spatially explicit models without spatial linkage (Nabel and Lischke, 2013). These two cases can be regarded as restricted application cases and will be picked up in the discussion.
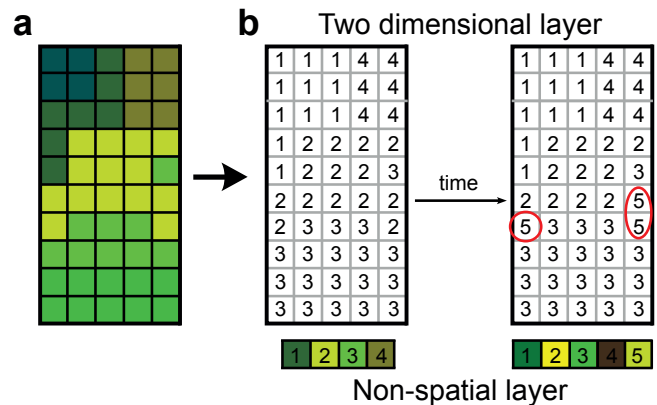

Figure 1. Visualisation of the dynamic two-layer classification concept. **(a)** Pre-defined similarity criteria (on abiotic drivers and biotic state variables) are used to classify similar grid cells to distinct types. Cells of the same type are coloured with the same shade of green. **(b)** For each type of cells one element is required on the non-spatial layer (coloured row, bottom). Cells on the two-dimensional layer (2-D-grid, top) are associated with these elements (numbers). Species' compositions can change over time due to processes simulated on both layers (changes in the coloured row). Furthermore, associations between the layers can change (changed numbers highlighted by red circles), for example when processes simulated on the spatial-layer (e.g. seed dispersal) lead to differences among cells associated with the same element that violate one of the similarity criteria.

In the following I will first outline the basic principles of D2C. Afterwards, I present how D2C can be applied on the example of TreeMig-2L, a new, efficient implementation of the intermediate-complexity DVM TreeMig (Lischke et al., 2006). Subsequently, the usefulness of this D2C implementation is examined by means of two different application scenarios.

2 The dynamic two-layer classification concept (D2C)

The target models for D2C are complex, two-dimensional, spatially linked time- and space-discrete DVMs. The aim of D2C is to increase the computational efficiency of a DVM through identification and elimination of redundant calculations. Redundant calculations can occur in simulations with a DVM because grid cells with comparable species' composition and abundances, i.e. comparable values in the cells' state variables, tend to follow the same successional paths, provided that (1) their abiotic drivers follow the same temporal pathways and (2) none of the grid cells are subject to any cell-specific deviations (e.g. caused by immigration). To identify and eliminate redundant calculations in a DVM, D2C uses two different layers (Fig. 1) onto which the processes of the DVM are divided. The first layer is the original two-dimensional layer containing all grid cells. On this layer only those processes of the model are simulated that use in-

formation on the spatial position of a cell relative to other cells in the simulation area. One example for such a spatial process is seed dispersal using source and sink positions. The second layer is a new associated non-spatial layer, on which all other processes are simulated, for example competition, growth and seed production. Each element on the non-spatial layer represents a certain type of cells on the two-dimensional layer, characterised by prevailing abiotic conditions and momentary species' composition and abundances. Each cell on the two-dimensional layer is thus associated with one element on the non-spatial layer (visualised in Fig. 1 with numbers). The size of the non-spatial layer is adaptive and depends on a dynamic classification according to similarity criteria on abiotic and biotic properties. These similarity criteria need to be pre-specified and are key to D2C because they define a discretisation of the abiotic drivers and the biotic state space. Spatial processes simulated on the two-dimensional layer can lead to violations of the similarity criteria because they can entail cell-specific deviations from otherwise similar grid cells, for example when a new species is dispersed to some but not all cells represented by one element. Thus, spatial processes can lead to changes in associations and can necessitate adding new elements to the non-spatial layer (see e.g. Fig. 1). Elements on the non-spatial layer, on the other hand, can be merged when they satisfy all similarity criteria.

To apply D2C, the processes of a model and its state variables need to be separated onto the two layers. Additionally, the exchange of status information between cells on the two-dimensional layer and their associated elements on the non-spatial layer needs to be specified. The assignment of processes and the definition of the interface between the two layers are critical steps because the information directed from different cells to their associated element is used to test if the biotic similarity criteria are violated, which would necessitate a split of the element.

The potential for reductions of computational expenses with D2C will be influenced by three aspects: (1) the non-reducible base load due to processes simulated on the two-dimensional layer; (2) the ratio of cells on the two-dimensional layer and elements on the non-spatial layer, which is strongly influenced by the specified similarity criteria; (3) the overhead introduced for managing elements on the non-spatial layer, associations between the two layers and the exchange of status information between the layers.

3 Methods

3.1 Implementation of TreeMig-2L

In this section I outline how D2C can be applied on the example of TreeMig-2L, a two-layer implementation of the forest-landscape model TreeMig. Further, and more detailed information on the implementation of TreeMig-2L is given in Supplement Sect. S1.

3.1.1 TreeMig

TreeMig is an intermediate-complexity DVM simulating local stand dynamics of multiple competing tree species. In addition to local dynamics, TreeMig allows for the explicit simulation of tree species' migration, with the rare advantage of including seed production, seed dispersal and subsequent regeneration processes (Thuiller et al., 2008). TreeMig simulations require time series of three different annual bioclimate drivers: the minimum winter temperature, the sum of daily mean temperatures above $5.5\,^{\circ}\mathrm{C}$, and an index denoting the severity of droughts (Lischke et al., 2006; Nabel et al., 2014). These drivers are derived from monthly averaged temperatures and monthly precipitation sums (see e.g. Supplement Sect. S2). TreeMig's state variables are population densities of tree species in a constant number of height classes per grid cell (Lischke et al., 2006). Additionally, TreeMig stores the density of currently available seeds per tree species, representing the seed bank of a cell.

The TreeMig two-layer implementation described in the following – TreeMig-2L – is implemented in Fortran and based on TreeMig-Netcdf 2.0 (Nabel et al., 2014). TreeMig-Netcdf 2.0 includes all processes described in the original TreeMig version (Lischke et al., 2006) with the growth response curve amendments described in Rickebusch et al. (2007), the minimum population density thresholds described in Nabel et al. (2013) and the climate extrapolation method preserving spatial autocorrelation described in Nabel et al. (2014).

3.1.2 Assigning processes to the two layers

Most of the processes simulated with TreeMig do not require information on the position of the cell relative to other cells on the grid: competition for light, growth, mortality and production of seeds in a cell only use information that is independent of other cells. Therefore, these processes can be simulated on the new non-spatial layer (see Fig. 2). In TreeMig-Netcdf 2.0, the only process which requires information on the position of a cell is seed dispersal. Other processes that would require information on the neighbourhood are, for example spatially connected disturbances, such as snow avalanches (e.g. Zurbriggen, 2013). In TreeMig-Netcdf 2.0, however, spatially connected disturbances are not represented. Thus, seed dispersal is the only process in TreeMig-2L that has to be simulated on the two-dimensional layer. Yet, to enable efficient simulations on two layers, associations between cells on the two-dimensional layer and elements on the non-spatial layer need to be as long lived as possible. This means that re-merging of elements that resulted from a very recent split needs to be prevented, and splits should only be conducted if they entail actual changes in
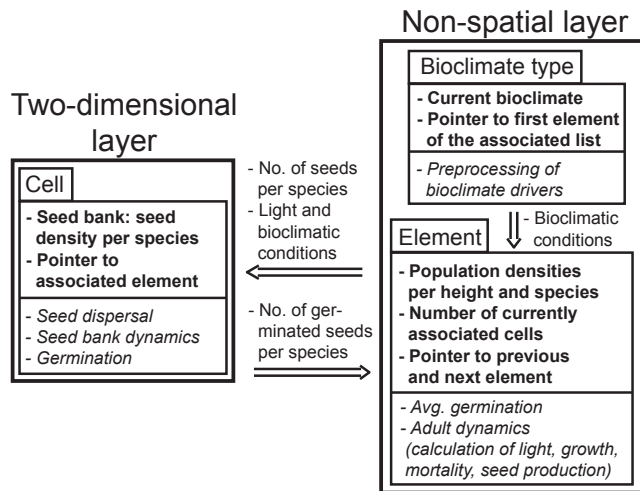
Two-dimensional layer

Figure 2. Outline of the architecture of TreeMig-2L. The two-dimensional layer consists of single cells whose state variables are seed densities per species in the seed bank. Additionally, each cell has a pointer – a data type allowing for direct access – to the element on the non-spatial layer with which the cell is currently associated. The non-spatial layer consists of bioclimate types and linked lists of associated elements (for more information see Supplement Sect. S1.1). State variables are printed in bold, processes in italic type. Items listed on the arrows represent information exchanged between the layers and between bioclimate types and elements of their associated list.

species' compositions. The availability of seeds of a species does not necessarily have to lead to changes in the species' composition, because a species might not be able to regenerate in a cell. Therefore, the entire regeneration process was assigned to the two-dimensional layer (Fig. 2) in order to only induce splits when newly dispersed seeds actually establish.

3.1.3 Architecture of TreeMig-2L

When designing the architecture of TreeMig-2L the two main requirements were an efficient organisation of the elements on the non-spatial layer and a fast exchange of status information between the two layers.

One of the challenges for an efficient organisation of the elements on the non-spatial layer was that the number of elements is not known in advance, because it is an emergent property of the simulation. To allow for an arbitrary number of elements on the non-spatial layer, the elements are stored in linked lists, instead of using an array structure with a pre-determined fixed size, as it is usually done in space-discrete DVMs. However, using one large linked list, and comparing all elements with each other, would be very inefficient and would lead to a large organisational overhead. To reduce the organisational overhead required for the comparison of elements during runtime, TreeMig-2L uses several linked lists, and only elements in the same list are compared. To define

the number of linked lists, the fact that the bioclimate drivers are an input to TreeMig is used, and that these drivers can thus be utilised to pre-structure a simulation area (see Fig. 3 for an example, and Supplement Sect. S1.2 for additional information). To pre-structure a simulation area, grid cells are classified into "bioclimate types". Each bioclimate driver is thereby discretised according to a set of pre-defined "bioclimate bins" (see e.g. Table 2). To limit the number of possible bioclimate types and to save computation time, the discretisation is only applied for temporal averages of the bioclimate drivers on a pre-defined set of "supporting periods". For each cell, one average per driver and supporting period is calculated. Two cells are classified into the same bioclimate type, if their averages fall into the same bioclimate bin for each driver and supporting period, whereby the bioclimate bins can differ in between supporting periods. For each existing bioclimate type, i.e. each type with which at least one cell is associated, an own linked list is used. The bioclimate drivers for a bioclimate type are calculated as the averages of all associated cells.

An element can be associated with a changing number of cells, therefore, an element only tracks how many but not which cells are currently associated with it (see Fig. 2). The information exchange between the layers is induced by the cells. Each cell on the two-dimensional layer accesses required information from its currently associated element. For seed dispersal calculations the density of produced seeds per species is required. For regeneration, additionally, the current bioclimatic conditions and the light distribution of the lowest height class need to be accessed, because they influence the density of newly germinated seeds (see Lischke et al., 1998, 2006). Newly germinated seeds are used to determine if splits are necessary (see Sect. 3.1.4). After possibly required dynamic changes in associations between the layers, cells push their germinated seed densities to the currently associated element and each element on the non-spatial layer calculates the average of the germinated seed densities received from its associated cells. Averaged densities of germinated seeds falling below a pre-specified presence threshold are thereby set to zero.

3.1.4 Dynamic associations

In order to account for spatial processes resulting from seed dispersal, associations between the two-dimensional and the non-spatial layer need to be dynamic. In a TreeMig-2L simulation, elements on the non-spatial layer can be split up or merged. For details on the execution sequence of a TreeMig-2L simulation see Supplement Sect. S1.3.

Splitting, and thus introduction of new elements on the non-spatial layer is required as soon as the density of germinated seeds among any two cells associated with the same element is considered not similar enough. In TreeMig-2L, similarity in the density of germinated seeds can be defined by a set of thresholds. The simplest possible set consists of a sin-
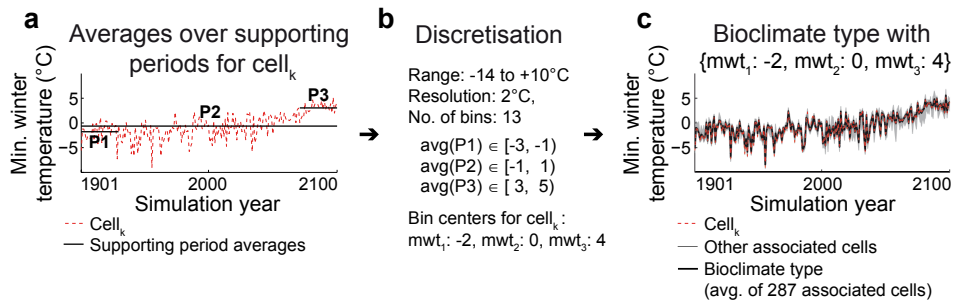
Figure 3. Visualisation of the pre-structuring into bioclimate types on the example of the minimum winter temperature (min. winter temperature), one of TreeMig's three bioclimate drivers. **(a)** The minimum winter temperature driving cell$_k$ is averaged for each of the supporting periods (here P1: 1901–1930, P2: 1901–2100, P3: 2071–2100). **(b)** The range of the minimum winter temperature (-14 to $+10\,^\circ$C) is discretised into 13 bins with a resolution of $2\,^\circ$C (as e.g. done for E3, the set of bioclimate bins with the coarsest resolution – see e.g. Table 2). The averages of the supporting periods P1–P3 for cell$_k$ are classified according to these bins. **(c)** Cells whose averages fall into the same bin in each of the supporting periods for each of the bioclimate drivers are classified into the same bioclimate type. The bioclimate driver (here the minimum winter temperature) of the bioclimate type is calculated as the average (black line) of its 287 associated cells (grey lines).

gle threshold that defines presence, i.e. species with a germinated seed density below this threshold are assumed to not have germinated. However, other thresholds are also possible, for example dividing sparse occurrences from more frequent ones. If the density of germinated seeds of any two cells associated with the same element fall on different sides of any of these thresholds for a single species, a split is required. One determinant for the efficiency of a TreeMig-2L simulation is assumed to be the number of considered splits. When $k - 1$ thresholds are specified and n species are simulated, k^n different combinations could possibly entail splits. In TreeMig-2L this trade-off between accuracy and possible splits is approached by not considering splitting for all species, but only for a set of species previously specified as species to be tracked; i.e. for tracked species differences in the number of germinated seeds among grid cells associated with the same element can lead to splits. For untracked species, in contrast, the number of germinated seeds is not compared among grid cells and deviations, thus, cannot lead to splits. Apart from the splitting step, all species are treated the same, in particular the number of germinated seeds on an element is calculated as the average over all associated cells for each species, no matter if tracked or untracked. As a consequence, untracked species might be represented less accurately, which can feedback on other species via competition. Therefore, species that are expected to change their distribution in the course of the simulation should be defined as tracked species.

Merging of elements is possible as soon as two elements belonging to the same bioclimate type are similar enough. The state variables stored in the elements on the non-spatial layer are population densities per species for a fixed set of height classes. In TreeMig-2L, two elements belonging to the same bioclimate type are regarded similar enough, when deviations between each pair of their population densities do not exceed a similarity threshold pre-specified for each of

the height classes. To avoid immediately re-merging of elements that were recently split up, newly germinated seed densities are also compared. As opposed to splitting, merging of elements on the non-spatial layer potentially decreases accuracy. Moreover, there is a trade-off between computational expenses involved with merging and the reduction of repetitive calculations caused by similar elements. Therefore, merging is only performed after a pre-defined number of iterations in TreeMig-2L.

3.2 Simulations with TreeMig-2L

3.2.1 Expected benefit of the D2C implementation

Previous TreeMig simulations were conducted with spatial resolutions ranging from cell side lengths of 25 m to 1 km and spatial extents ranging from single cells up to $77\,000\,\text{km}^2$. The number of simulated interacting tree species ranged from one up to 31 species and simulation areas had differing spatio-temporal complexity (see e.g. Epstein et al., 2007; Nabel et al., 2013; Nabel and Lischke, 2013; Zurbriggen et al., 2014). Such differences in the simulation settings are expected to influence the benefit of the D2C implementation, because they control two of the three aspects assumed to be most important (see Sect. 2): (1) the non-reducible base load due to processes simulated on the two-dimensional layer, and (2) the ratio of grid cells and elements on the non-spatial layer, whereby the latter is strongly influenced by the number of bioclimate types.

The number of bioclimate types is, for example expected to be large when the simulation area has a high spatio-temporal complexity, i.e. a high heterogeneity in its bioclimate drivers. The number of bioclimate types will also be influenced by the spatial extent of a simulation area: a larger extent, up to a certain point, potentially leads to a larger number of types; however, there might be a threshold beyond which the number of already contained bioclimate types ex-
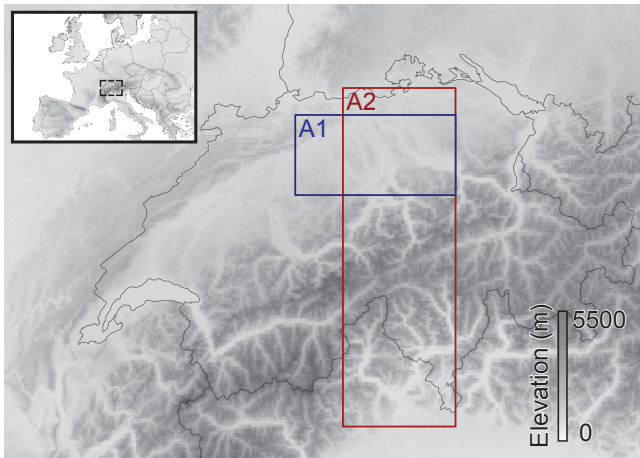
Figure 4. Location of the two simulation areas in Switzerland and in Europe. The map illustrates the differences in elevational heterogeneity (digital elevation model by Jarvis et al., 2008) between the simulation area of scenario A1 and of scenario A2.

ceeds the number of newly added types. A similar effect is expected for the spatial resolution; bioclimate drivers in a cell are always averages and with finer resolution fewer bioclimate extremes might be smoothed out, leading to a larger number of required bioclimate types. However, especially in areas with a homogeneous bioclimate, a finer resolution might entail a larger number of similar grid cells, increasing the expected benefit of a D2C implementation. A finer resolution, on the other hand, will increase the computational expenses for seed dispersal. In TreeMig, seed dispersal is simulated from the perspective of the source cell, providing seeds to sink cells according to a pre-calculated truncated probabilistic density function (see Supplement of Lischke et al., 2006). The number of sink cells thereby depends on the spatial resolution. A fine resolution implies a large number of sink cells, causing large computational expenses on the two-dimensional layer, and therefore a large base load. A fine resolution thus can diminish the benefit of a D2C implementation. Further aspects that can influence the benefit of the D2C implementation are the number of tracked species, the splitting and merging thresholds, and the number of iterations after which merging is considered.

3.2.2　Application scenarios

TreeMig-2L simulations were conducted for two different application scenarios (A1 and A2), pertaining to different regions of Switzerland. Figure 4 shows the location of their simulation areas. Scenario A1, in the flatter areas of northern Switzerland, stems from a study with a preliminary TreeMig-2L version without dynamic associations between the layers (Nabel and Lischke, 2013). Scenario A2, a north–south transect across the Swiss Alps, stems from a study investigating the influence of interannual bioclimate variability in simu-

lations of the northwards migration of *Ostrya carpinifolia* Scop. (European Hop Hornbeam) with TreeMig-Netcdf 1.0 (Nabel et al., 2013). An in-depth description of the scenarios can be found in Supplement Sect. S2.

The two scenarios were selected because they strongly differ in their simulation settings (see Table 1 for their main characteristics). Application scenario A1 has a more homogeneous simulation area and a finer spatial resolution than A2. Therefore, on one hand, a more beneficial ratio between the number of cells on the two-dimensional layer and the number of bioclimate types is expected for A1. On the other hand, due to the finer spatial resolution a higher base load is expected for A1 than for A2.

The tracked species in scenario A2 is naturally the investigated migrating species (*O. carpinifolia*). The other 21 competing species simulated in this application scenario were not tracked. For A1 4 out of the 31 simulated competing species were selected as tracked species: *Quercus pubescens*, *O. carpinifolia*, *Larix decidua* and *Pinus sylvestris*. These species were selected because they have the highest drought tolerance indices in TreeMig. With increasing drought severity (and increasing temperatures), these species are therefore expected to extend their spatial distributions, whilst the other species might be less effected or have declining distributions. For both scenarios, merging was considered every 100 simulation years and the same splitting and merging thresholds were used (for further details and sensitivity tests see Supplement Sect. S3)

Both application scenarios were driven by bioclimate time series derived from SRESA1B (Nakicenovic et al., 2000) scenario projections, though from different models and downscaled with different observational data (see Supplement Sect. S2). To cover the entire simulation time span of the examples, a stochastic extrapolation method accounting for the spatial correlation of bioclimate fluctuations was used (Nabel et al., 2014).

3.2.3　Pre-structuring of the simulation areas

Bioclimate types for both application scenarios were derived with the same sets of bioclimate bins: E1, E2 and E3 (Table 2). For both scenarios three supporting periods were averaged: the first 30 (A1: 1961–1991; A2: 1901–1931) and the last 30 years (both: 2071–2100), as well as the whole time span (A1: 1961–2100; A2: 1901–2100).

3.2.4　Conducted simulations

Simulations were conducted with all three sets of bioclimate types (Table 2). To account for the stochasticity involved with the extrapolation of the bioclimate drivers multiple repetitions of the simulations were performed. Simulations with application scenario A1 were repeated 5 times. Simulations with scenario A2, having much less grid cells and requir-

Table 1. Main characteristics of the two application scenarios A1 and A2. Scenario A1 stems from a study with a preliminary TreeMig-2L version (Nabel and Lischke, 2013), scenario A2 from a study on the influence of interannual bioclimate variability on the simulated migration of *Ostrya carpinifolia* (Nabel et al., 2013). For an in-depth description of the scenarios see Supplement Sect. S2.

	Spatial complexity	Spatial resolution[a]	Spatial extent	Number of grid cells (stockable cells[b])
A1	Rather homogeneous	200 m	5000 km^2	125 000 (110 789)
A2	Very heterogeneous	1 km	14 700 km^2	14 700 (12 230)

	Competing species	Tracked species	Simulated time span	Number of repetitions
A1	31	Four most drought resistant	1400–2500	5
A2	22	*Ostrya carpinifolia*	1400–3000	100

[a] Cell side length. [b] Stockable cells denote grid cells in which trees can grow in a TreeMig simulation (trees cannot grow in cells covered with water and cells with solid rock surfaces).

Table 2. Three sets of bioclimate bins (E1, E2 and E3) were used to pre-structure the simulation areas. The bioclimate bins resulted from discretising the bioclimate variables with different granularity (E1: fine, E2: moderate, E3: coarse). Minimum and maximum values of the bioclimate variables cover the bioclimate ranges of both application scenarios.

Bioclimate variable	Range		Resolution (number of bioclimate bins)		
	min	max	E1	E2	E3
DDsum$_{>5.5\,°C}$ [a]	0	4200	50 (85)	100 (43)	200 (22)
Min. witemp.[b]	−14	10	0.5 (49)	1.0 (25)	2.0 (13)
Drought index	0.0	0.7	0.025 (29)	0.05 (15)	0.1 (8)

[a] DDsum$_{>5.5\,°C}$: sum of daily mean temperatures above 5.5 °C.
[b] Min. witemp.: minimum winter temperature [°C].

ing only one-tenth of the CPU time of scenario A1[1], were repeated 100 times. To disentangle the effects of the pre-structuring into bioclimate types and of dynamic associations between the layers, simulations of four model versions with increasing complexity were compared (Table 3). (1) 1L-ORG: the original one-layer approach with the original bioclimate driver. Simulations with this version were conducted to obtain reference values for the comparisons. (2) 1L-PB: a model version in which bioclimate types were used to derive averaged bioclimate drivers, but in which all processes were still simulated on one layer. Simulations with this version thus isolate the loss in accuracy due to the averaging of the bioclimate drivers. (3) 2L-NDA: a model version using bioclimate types and two layers but no dynamic associations, i.e. in this version each bioclimate type has only one element. Because dynamic associations are switched off, differences in the number of germinated seeds between cells associated with the element do not lead to splits, and a species introduced in one of the cells will thus occur in all associated cells. Therefore, simulations with this version can elu-

cidate if there is a necessity to track species. Furthermore, having no costs for splitting and merging, this version can be used to derive the maximum CPU time reduction that can be achieved. (4) 2L: the full TreeMig-2L model with two layers and dynamic associations.

3.2.5 Applied performance measures

The performance of the D2C implementation was assessed by two means: an accuracy measure and the required CPU time[2]. To measure accuracy, different output variables were compared with a similarity coefficient (SC – Eq. 1) already used for intra-model comparisons in previous studies (e.g. Lischke et al., 1998; Nabel et al., 2013).

$$SC_y = 1 - \frac{\sum_i^{cell} |D_{sum_i}|}{\sum_i^{cell} S_{sum_i}} \tag{1}$$

SC_y for a year y ranges from 0.0 (no similarity) to 1.0 (identical output) and is reciprocally dependant on the ratio of the sum of an output variable of two simulations S_{sum_i} and their differences $|D_{sum_i}|$ summed over all cells of the simulation

---

[1] The average CPU time with the original one-layer approach (1L-ORG) was 50 353 s for A1 (average of five simulations with $\sigma = 784$ s) and 5195 s for A2 (average of 100 simulations with $\sigma = 38$ s).

[2] CPU time was measured with the intrinsic Fortran procedure CPU_TIME. All simulations were conducted on 2.8 GHz AMD Opteron CPUs.

Table 3. Simulations of four model versions with increasing complexity were compared. The model versions differ in whether they (1) use the pre-structuring to bioclimate types, (2) do simulations on two layers and (3) apply dynamic associations between the two layers, i.e. allow for splitting an merging of elements on the non-spatial layer; 1L-ORG refers to the original one-layer approach with the original bioclimate driver; 1L-PB introduces the averaging of the bioclimate drivers according to the bioclimate types, but still runs on one layer; 2L-NDA applies the averaging and runs on two layers, albeit without dynamic associations; 2L, finally, refers to the full TreeMig-2L model with two layers and with dynamic associations.

	Bioclimate types	Two layers	Dynamic associations
1L-ORG	no	no	–
1L-PB	yes	no	–
2L-NDA	yes	yes	no
2L	yes	yes	yes

area. The SC was used to compare different variables among simulations with the 1L-ORG model version and model versions using bioclimate types (Table 3). Each comparison was thereby conducted with simulations using the same pseudorandom number stream to extrapolate the bioclimate driver. For both application scenarios, the SC for the sum of the biomass of all species (SC_{sum}) and the SC of the biomass per species (SC_{spec}) were calculated, for A1 every century and for A2 every 50 years. For A2 simulations, furthermore, the SC of the biomass of *O. carpinifolia* (SC_{OC}) was calculated for each year.

Two further measurement were taken to assess the applicability of D2C in view of the three aspects listed at the end of Sect. 2. First, the maximum number of elements reached during a simulation was tracked, here refereed to as the peak element-cell ratio. Second, the ratio of the time spent on different processes was profiled with callgrind, a tool that records function calls in a program during runtime (Weidendorfer, 2008).

4   Results and discussion

4.1   Pre-structuring of the simulation areas

For both application scenarios the number of bioclimate types resulting from the pre-structuring was considerably smaller than the number of grid cells. However, the resulting number of types differed in absolute as well as relative terms (Table 4). The simulation area of application scenario A1 always contained fewer bioclimate types than A2 (E1: ∼ factor 2; E2 and E3: ∼ factor 4). Moreover, the resulting ratio of bioclimate types to grid cells was much smaller (E1: ∼ factor 20; E2 and E3: ∼ factor 40). As suggested in

Table 4. Number of bioclimate types resulting from the three different sets of bioclimate bins (E1–E3 – Table 2). For both application scenarios (A1 and A2), absolute numbers of bioclimate types and percentages relative to the number of stockable cells of the simulation area (see Table 1) are listed.

	Application scenario A1		Application scenario A2	
	No. of bioclimate types	%	No. of bioclimate types	%
E1	3460	3.1 %	7941	64.9 %
E2	798	0.7 %	3424	28.0 %
E3	213	0.2 %	884	7.2 %

Sect. 3.2.1, this is due to the higher spatial homogeneity, the finer spatial resolution and the smaller spatial extent (and thereby smaller bioclimatic range) of the simulation area of scenario A1 compared to the simulation area of A2. A2's simulation area is divided in a larger number of bioclimate types, with two-thirds as many bioclimate types as cells for the set of bioclimate bins with the finest resolution (E1). For this set, thus, more than half of the cells end up with their own bioclimate type. For an example of the distribution of numbers of cells to bioclimate types see Supplement Fig. S3.3.

Averaging of the bioclimate drivers of associated cells to obtain the drivers for the bioclimate types led to small but visible deviations (e.g. Fig. 3c and Supplement Sect. S3). These deviations in the bioclimate driver entail deviations in the simulation results when comparing to results from one-layer simulations, i.e. simulations with the original grid cell-based bioclimate driver (see column 1L-PB in Table 5 and Sect. 4.2). Using bioclimate bins with a coarser resolution thereby led to larger deviations in the driving bioclimate and, thus, to larger deviations in the simulation results (Table 5). In the current implementation of TreeMig-2L these deviations can only be reduced using a finer discretisation into bioclimate bins or more supporting periods. The dynamic association of cells to elements on the non-spatial layer cannot reduce this error, because the bioclimate driver is not processed on the single elements but on the bioclimate types. The pre-structuring of the simulation area is fundamental for the efficiency of TreeMig-2L, because the thereby obtained static structure is used for the organisation and maintenance of the elements on the non-spatial layer (see Sect. 3.1.3). However, averaging of the bioclimate in the pre-structuring step could potentially be replaced by a dynamic approach. In a dynamic approach, pre-processing of the bioclimate drivers could be re-transferred to the two-dimensional layer. In each time step of the simulation the bioclimate drivers of each element could then be calculated by averaging the bioclimate of the presently associated cells. Whilst this could reduce deviations in the simulation results compared to original one-layer simulations, it would also require more computations and thus could lead to less CPU time reductions (see Sect. 4.2.2).

Applying bioclimate types in their current static form is reminiscent of the stratified sampling methods used

Table 5. Performance measures for simulations with both application scenarios (A1 and A2) and with all three sets of bioclimate types (E1–E3). For both scenarios, the table lists the mean similarity coefficients (SC) comparing results of the last simulation year from the three model versions (1L-PB, 2L and 2L-NDA) to 1L-ORG simulations for different target variables (biomass sum over all species, biomass per species and, for scenario A2, biomass of *Ostrya carpinifolia*). In addition to the SCs, average peak element-cell ratios and mean CPU time reductions relative to 1L-ORG simulations are listed. All means for A1 stem from five repetitions, for A2 from 100 repetitions. SCs from 2L simulations were always very close to SCs from 1L-PB simulations, indicating that the dynamics in 2L simulations follow the dynamics in 1L-ORG simulations and that most of the deviations are due to the averaging of the bioclimate drivers for the bioclimate types. SCs from 2L-NDA are smaller, in particular for SC_{OC}, underlining the importance to track migrating species.

| | | Application scenario A1 (5 repetitions) | | | | Application scenario A2 (100 repetitions) | | | | |
| | | Avg. SCs[a] | | Avg. peak | Avg. CPU time[c] | Avg. SCs[a] | | | Avg. peak | Avg. CPU time[c] |
		SC_{sum}	SC_{spec}	e-c ratio[b]	reduction [%[d]]	SC_{sum}	SC_{spec}	SC_{OC}[e]	e-c ratio[b]	reduction [%[d]]
	1L-PB	0.98	0.90	–	∅	0.99	0.95	0.97	–	∅
E1	**2L**	**0.98**	**0.89**	**56.5** %	**52.4** %	**0.99**	**0.95**	**0.96**	**71.2** %	**32.6** %
	2L-NDA	0.97	0.85	3.1 %	59.6 %	0.98	0.92	0.71	64.9 %	33.7 %
	1L-PB	0.97	0.86	–	∅	0.96	0.88	0.89	–	∅
E2	**2L**	**0.97**	**0.85**	**39.9** %	**56.0** %	**0.96**	**0.87**	**0.88**	**38.3** %	**65.6** %
	2L-NDA	0.96	0.82	0.7 %	61.0 %	0.96	0.85	0.67	28.0 %	68.3 %
	1L-PB	0.96	0.78	–	∅	0.94	0.82	0.83	–	∅
E3	**2L**	**0.96**	**0.78**	**29.4** %	**57.8** %	**0.94**	**0.81**	**0.82**	**18.1** %	**84.7** %
	2L-NDA	0.96	0.75	0.2 %	61.4 %	0.94	0.80	0.63	7.2 %	87.0 %

[a] SC comparisons were conducted for the last year of the simulations, i.e. for the simulation year 2500 in case of A1 simulations and 3000 in case of A2 simulations. Simulations of the listed model versions (1L-PB, 2L, 2L-NDA – see Table 3) were compared with 1L-ORG simulations using the same pseudo-random number streams to extrapolate the bioclimate driver. Examples of the temporal development of the SCs are given in Fig. 5, Fig. 6 and in Supplement Sect. S3. Standard deviations for A1 simulations were always smaller than 0.02. Standard deviations for A2 simulations with the model version 2L-NDA were much larger and reached up to 0.04.
[b] Peak ratio between the number of elements on the non-spatial layer and the number of stockable cells in the simulation area of the application scenario (Table 1).
[c] Standard deviations of CPU times were always smaller than 1.5 %; therefore, only average CPU times are shown.
[d] Reduction relative to the average CPU time required for 1L-ORG simulations (see footnote 1).
[e] OC: *Ostrya carpinifolia*, the tree species whose northwards migration is simulated in application scenario A2.
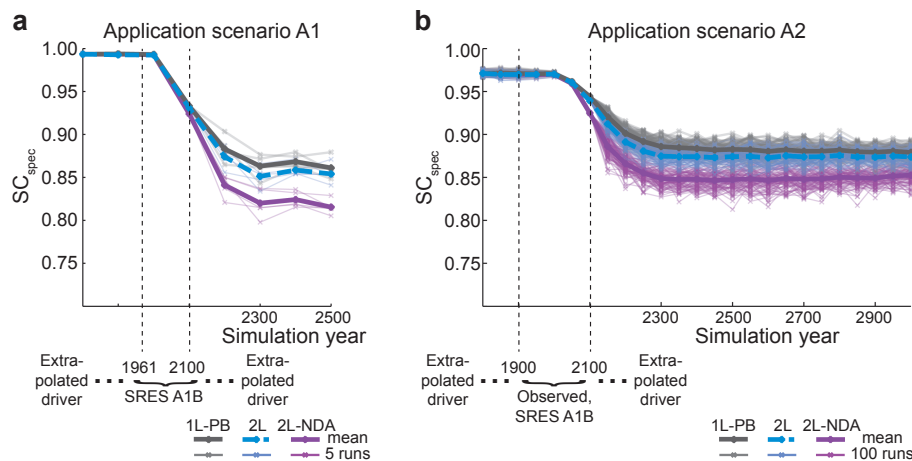

Figure 5. Temporal development of SC_{spec} comparing the biomass per species. Depicted is the comparison of 1L-ORG simulations to two-layer simulations with and without dynamic associations (2L and 2L-NDA, respectively) and simulations on one layer but with the bioclimate of the associated bioclimate types (1L-PB). For both application scenarios (A1: **a**, A2: **b**) bioclimate types were derived with the set of bioclimate bins with moderate resolution (E2; for other bioclimate types see Supplement Sect. S3). SC_{spec} values for A1 were calculated every 100 years and values for A2 every 50 years. For each setting several repetitions were compared, which were simulated using different pseudo-random number streams to extrapolate the bioclimate driver before 1961 (A1) or 1901 (A2) and after 2100. Five simulations were conducted for A1 and 100 for A2. Single simulations and their means are printed half-transparent and bold, respectively. All depicted SC_{spec} time series decline in the transient phase of climate change and continue to decline for about 200 years after which SC_{spec} stabilises. Over the whole simulated time span SC_{spec} values of 2L simulations are much closer to SC_{spec} values from 1L-PB simulations than to SC_{spec} values from 2L-NDA simulations, indicating that the deviations are mainly due to the averaging of the bioclimate drivers.

in explicit spatial upscalings of single site models (Bugmann et al., 2000), and is comparable to the ecoregions used in the forest-landscape model LANDIS (see e.g. Mladenoff and He, 1999; Scheller and Mladenoff, 2004). Similarly to cells associated with the same bioclimate type in TreeMig-2L, cells associated with the same ecoregion in LANDIS share important process rates influencing establishment and biomass development (e.g. Scheller and Mladenoff, 2004). Like bioclimate types in TreeMig-2L, ecoregions in LANDIS do not need to be contiguous but can be distributed in space (He et al., 1999). Not demanding contiguousness is an important advantage over other upscaling methods, which are often based on local spatial aggregations, such as naive upscalings that decrease the spatial resolution of a simulation area (as described in Bocedi et al., 2012). The possibility for bioclimate types, which are defined over similarity and not over spatial proximity, to be arbitrarily distributed in space, potentially reduces the number of required types and, even more important, prevents errors involved with inappropriate averaging of neighbouring grid cells. Thus, as opposed to local spatial aggregations, bioclimate types have the advantage that they conserve the spatio-temporal heterogeneity to a large degree. Supplement Fig. S3.6 gives an example for the conservation of spatial variability; Figs. 3, S3.4 and S3.5 give examples for the conservation of temporal variability.

4.2　Performance of TreeMig-2L simulations

To evaluate the performance of TreeMig-2L, different performance measures (Sect. 3.2.5) were used to compare two-layer simulations using different bioclimate discretisations (Table 2) to one-layer simulations using the original bioclimate. To assess to what extent different approximations involved with D2C led to performance decreases, further simulations were conducted using two different pre-stages of the 2L implementation (Table 3). In addition, sensitivity test for splitting and merging thresholds, merging intervals and tracked species were conducted (see Supplement Sect. S3).

4.2.1　Accuracy

The accuracy of TreeMig-2L was evaluated by comparing biomass distributions from simulations using different model versions with a similarity coefficient (SC) ranging from 0.0 (no similarity) to 1.0 (identical biomass distributions in space). Comparisons generally led to SCs in the upper range for all output variables (ranging from about 0.8 to about 1.0 – Table 5) and for both application scenarios. The level of the SC was thereby largely determined by the resolution of the applied set of bioclimate bins; the coarser the resolution is, the smaller the SCs for all variables (differences in the SCs of up to 0.14 – Table 5).

To be able to assess the relevance of deviations from 1.0 in the SCs, results from 1L-ORG simulations using different pseudo-random number streams to extrapolate the bioclimate
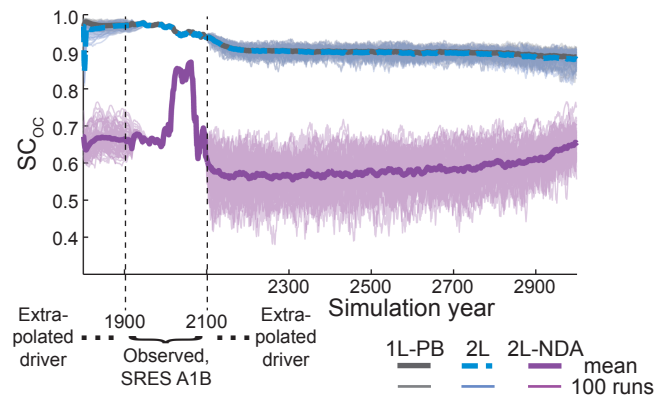

Figure 6. Temporal development of SC_{OC} comparing the biomass of *Ostrya carpinifolia* in simulations of application scenario A2. Depicted SC_{OC} values stem from comparisons of 1L-ORG simulations to two-layer simulations with and without dynamic associations (2L and 2L-NDA, respectively) and simulations on one layer but with the bioclimate of the associated bioclimate types (1L-PB). Bioclimate types were derived with the set of bioclimate bins with moderate resolution (E2; for other bioclimate types see Supplement Sect. S3). For each setting 100 repetitions were compared, which were simulated using different pseudo-random number streams to extrapolate the bioclimate driver before 1901 and after 2100. Single simulations and their means are printed half-transparent and bold, respectively. Comparable to SC_{spec} values (Fig. 5), SC_{OC} values of comparisons between 1L-ORG and 2L simulations are very close to SC_{OC} values from comparisons with 1L-PB simulations over the whole simulated time span. These SC_{OC} time series thus reflect differences due to the averaging of the bioclimate driver. SC_{OC} values from comparisons of 1L-ORG simulations with 2L-NDA simulations, in contrast, mainly reflect the reaction of *O. carpinifolia* to climatic changes on the whole simulation area (see Supplement Sect. S3.2), since 2L-NDA simulations were conducted without dynamic associations. 2L-NDA simulations get closer to 1L-ORG simulations towards the end of the simulation due to the northwards spread of *O. carpinifolia* in 1L-ORG over time.

driver were compared with each other, i.e. the deviation from 1.0 in the SC due to interannual bioclimatic variability was calculated (see Supplement Sect. S3.2). The only SCs that were markedly larger than the SCs resulting from these 1L-ORG intra-comparisons were SCs from comparisons with E3 bioclimate types (compare Table 5 and Table S3.1).

In all cases, SCs from comparisons between 1L-ORG and 2L simulations were very close to SCs from comparisons of 1L-ORG and 1L-PB simulations (differences in the SCs ≤ 0.01). This indicates that the deviation in 2L simulations are mainly due to the averaging of the bioclimate drivers. SCs comparing 2L and 1L-ORG were, in particular, larger than SCs comparing 2L-NDA and 1L-ORG (differences in SC_{spec} up to 0.04 and in SC_{OC} up to 0.21), indicating the importance to track migrating species. The four tracked drought tolerant species in application scenario A1 and *Ostrya carpinifolia* in application scenario A2 seem to be good indicators to test for required splits. Sensitivity tests showed that fewer

tracked species increased the error (up to 0.03 smaller SC_{spec} values) and more species only slightly decreased it (only ~ 0.01 larger SC_{spec} values – see Supplement Fig. S3.13). The sensitivity tests further showed that the choice of merging and splitting thresholds had only limited impact on the SCs. Strong differences in the SC were only observed for simulations that tested elements for merging after a decade instead of after a century (see Supplement Sect. S3).

The temporal development of SC_{spec} was comparable among simulations with the three sets of bioclimate types (E1–E3) for all model versions (2L, 1L-PB and 2L-NDA) and for both application scenarios. SC_{spec} decreased in the transient phase of climate change (from around 2000 on) and stabilised after a few centuries (Fig. 5 and Supplement Sect. S3). The stabilisation on a lower level is mainly due to a stronger impact of differences in the drought index between bioclimate types and single cells due to overall larger drought indices in the second half of the 20th century (see Supplement Figs. S3.4 and S3.5). A comparable effect resulted for inter-comparisons of 1L-ORG simulations (see Supplement Sect. S3.2). While the temporal development in SC_{spec} was comparable among all model versions, trajectories resulting from 2L-NDA simulations differed from those resulting from 2L and 1L-PB simulations for SC_{OC} (Fig. 6). Having no dynamic associations, SC_{OC} for 2L-NDA simulations mainly reflects changes in the bioclimate over time (see Supplement Sect. S3.2 for details).

Being a single index for the whole simulation area, the SC is a rough indicator for the similarity between two simulations and could in particular conceal biases in the spatial biomass distribution. Furthermore, depending on the research question, other aspects could be more important than absolute biomass values per grid cell, especially when studying the migration of a species. Therefore, simulations of application scenario A2 were additionally compared focussing on the migration of the tracked species. Comparisons of the spatial spread of *O. carpinifolia* showed a reasonable approximation of 1L-ORG simulations by 2L simulations with all bioclimate discretisations (Fig. 7 and Supplement Fig. S3.11).

In summary, SCs comparing 2L and 1L-ORG simulations were within the magnitude of deviations due to interannual bioclimatic variability and deviations where smaller than reported for previous upscalings (e.g. Lischke et al., 1998). Furthermore, comparisons of the spatial spread of *O. carpinifolia* between 1L-ORG and 2L simulations did not indicate spatial biases and comparisons with 2L-NDA simulations underlined the necessity to track migrating species.

4.2.2 Computational expenses

Simulations with two layers led to considerable reductions in CPU time for both application scenarios (Table 5). When considering the ratio between the number of bioclimate types and the number of cells in the simulation areas (Table 4),
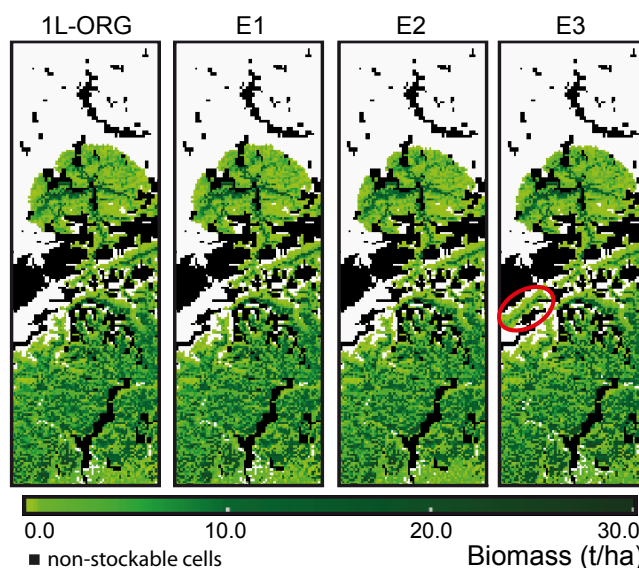

Figure 7. Spatial spread of *Ostrya carpinifolia* in simulations with application scenario A2. Depicted are maps of the biomass distribution of *O. carpinifolia* from the last simulation year (3000) of a 1L-ORG simulation and 2L simulations with the three sets of bioclimate types (E1–E3). The maps are hardly visually discernible and the only notable difference between the depicted maps is one valley in the west of the transect, which is inhabited by *O. carpinifolia* in E3 (red circle), but not in the other depicted simulations. The absence of notable spatial biases indicates that a large share of the differences in the biomass of *O. carpinifolia* between 2L and 1L-ORG simulations (i.e. SC_{OC} values) stem from small variations in the local biomass. All depicted maps stem from simulations using the same pseudo-random number streams to extrapolate the bioclimate driver and led to SC_{OC} values (E1: 0.97, E2: 0.88, E3: 0.82) close to the means (Table 5). Maps of the biomass of *O. carpinifolia* were created with Paraview (Ahrens et al., 2005).

larger CPU time reductions could have been expected for scenario A1 than for A2. Yet, the actual reductions are not systematically larger for A1 (E1: 52.9 %, E2: 56.0 %, E3: 57.8 %) than for A2 (E1: 32.6 %, E2: 65.6 %, E3: 84.7 %). To a small part this is due to increased dynamics in the number of elements on the non-spatial layer in A1 compared to A2 (Fig. 8), leading to a comparable magnitude in the peak element-cell ratio for both application scenarios (Table 5). However, sensitivity tests for the application scenario A1 showed that changes in the number of tracked species and in the merging interval leading to large changes in the peak element-cell ratio did not lead to notable changes in CPU time reductions (see Supplement Sect. S3.3.1). Moreover, the small difference in CPU time reductions among A1 simulations with the three sets of bioclimate types in combination with notable differences in the peak element-cell ratio (Table 5) also indicates that the peak element-cell ratio is not the main cause for the limitation in the CPU time reduction for A1.

The actual main cause is revealed when looking at the percentage of executed instructions for the different processes (Table 6): the percentage of instructions spent on seed dispersal was simply much larger for A1 than for A2. 1L-ORG simulations for application scenario A1 spent about 45 % of executed instructions on seed dispersal and only about 50 % on adult dynamics, i.e. on processes simulated on the non-spatial layer. A2 1L-ORG simulations, in contrast, only spent about 6 % of the executed instructions on seed dispersal and about 86 % on adult dynamics (Table 6). Differences in the execution ratios were already expected (see Sect. 3.2.1) because of the increased number of sink cells considered in seed dispersal calculations for the grid with 200 m cell side length in A1 compared to the grid with 1 km cell side length in A2. As a consequence, the base load that cannot be reduced with D2C was more than 7 times larger for A1 than for A2.

For A2 differences in CPU time reductions were strongly driven by the applied set of bioclimate types. Sensitivity tests showed that splitting and merging thresholds as well as the merging interval were far less influential (Supplement Sect. S3.3.2).

4.3   Applicability of D2C

In Sect. 2 three different aspects were hypothesised to influence the benefit of a D2C implementation: (1) the non-reducible base load, (2) the element-cell ratio, and (3) the organisational overhead. The implementation and the applications of TreeMig-2L confirmed the importance of these aspects. Differences in the benefit comparing application scenarios A1 and A2 demonstrated the key role of the first aspect – the non-reducible base load due to time spent with spatially linked processes (Table 6). The key role of this aspect is also underlined when comparing simulations including spatial linkages in this study to the simulations with a predecessor of TreeMig-2L in a study without spatial linkages by Nabel and Lischke (2013): simulations without spatial linkages entailed much larger reductions in CPU times. The second aspect – the number of elements on the non-spatial layer – was, on the one hand, controlled by the number of bioclimate types derived in the pre-structuring (particularly for A2) and, on the other hand, by the number of tracked species (for A1). The number of bioclimate types, as well as the number of tracked species, influenced the reductions in CPU time via the number of elements on the non-spatial layer. Compared to the first aspect, the second aspect played a minor role. The third aspect – the organisational overhead – only had a small contribution in TreeMig-2L (see Table 6), which was mainly possible due to the developed architecture: the pre-structuring into bioclimate types and the asymmetric communication between cells and associated elements enabled an efficient maintenance of the dynamic non-spatial layer and the pointers kept in each cell enabled the direct access of the associated elements.
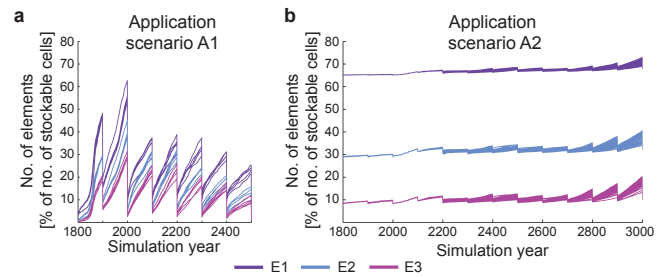

Figure 8. Ratio of the number of elements on the non-spatial layer to cells in the simulation area over time from simulations with all three sets of bioclimate types (E1–E3) and for both application scenarios (A1: **a**, A2: **b**). For A1 five repetitions using different pseudo-random number streams to extrapolate the bioclimate driver are depicted; for A2 100 repetitions are shown. The number of elements in A1 simulations increased strongly in between the merging intervals with smaller increases in the course of the simulation. For A2 simulations, in contrast, the number of elements increased faster towards the end of the simulation time, due to the growing perimeter of the migration front of the tracked species *Ostrya carpinifolia* (see Supplement Fig. S3.11).

The three aspects can be used to assess the applicability of D2C for other DVMs. The main determinant for the non-reducible base load is the number and complexity of processes requiring information on the neighbourhood of a grid cell, i.e. spatially linked processes. The number and complexity of spatially linked processes might also influence the number of required elements due to induced splits and will thereby also influence the organisational overhead for the maintenance of the non-spatial layer, not only due to the required splits but also due to the required information exchange between the layers. Thus, D2C might be less suitable for models with many complex and interacting spatially linked processes, such as LANDIS-II (Scheller et al., 2007), if used with several spatially linked extensions (e.g. Sturtevant et al., 2012). It should, however, be applicable for most models with few and simple spatially linked processes, such as TreeMig (provided a relatively coarse spatial resolution is used) and for spatially independent one-dimensional DVMs (sensu Fisher et al., 2010), such as ED (Moorcroft et al., 2001; Fisher et al., 2010) and most implementations of LPJ-Guess (e.g. Hickler et al., 2012; Smith et al., 2014). An upscaling of a one-dimensional DVM with D2C could perhaps free computational resources for the inclusion of a simple seed dispersal algorithm, which would be an important step towards the explicit simulation of migration. Simulating migration explicitly would be highly desirable (Neilson et al., 2005; Thuiller et al., 2008), which is also underlined by the results of this study comparing 1L-ORG and 2L-NDA simulations (see e.g. Supplement Fig. S3.11).

Besides the number and complexity of the spatially linked processes, also properties of the local processes will influence the applicability of D2C: the concept will be suitable for

Table 6. Percentage of instructions executed for selected computation tasks. Shown are results from a callgrind (Weidendorfer, 2008) profiling of 1L-ORG simulations and of 2L simulations with all three sets of bioclimate types (E1–E3). The measured percentage of instructions executed for selected computation tasks is a performance measure comparable to the relative amount of CPU time spent with the tasks. Measures stem from simulations using the same pseudo-random number stream to extrapolate the bioclimate driver.

| | Application scenario A1 | | | |
	Seed dispersal	Bioclimate prep.[a]	Adult dynamics	Overhead[b]
1L-ORG	44.31 %	3.75 %	49.61 %	–
2L: E1	83.71 %	0.21 %	14.66 %	0.11 %
2L: E2	88.40 %	0.05 %	10.19 %	0.10 %
2L: E3	91.80 %	0.01 %	6.84 %	0.10 %
	Application scenario A2			
	Seed dispersal	Bioclimate prep.[a]	Adult dynamics	Overhead[b]
1L-ORG	6.03 %	4.80 %	85.97 %	–
2L: E1	9.02 %	4.62 %	83.04 %	0.15 %
2L: E2	17.69 %	3.84 %	74.16 %	0.24 %
2L: E3	39.52 %	2.20 %	51.42 %	0.47 %

[a] Bioclimate prep.: pre-processing of the bioclimate driver, which involves reading of the current values from a NetCDF file and derivation of species-specific coefficients for later calculations with bioclimate dependencies, such as growth, mortality and establishment.
[b] Overhead: organisation of dynamic associations between the two-dimensional and the non-spatial layer.

a model in which local processes are rather deterministic, as in TreeMig, or in which the stochasticity in local processes is realised as patch replicates calculated and averaged for each iteration within a grid cell, as for example done in LPJ-Guess (Smith et al., 2001). For models with such properties, grid cells with similar values in the climatic drivers will tend to also have comparable species' compositions, and these grid cells can thus be represented by the same element. D2C will be less suitable for a model in which stochasticity in the local processes is realised on the level of the grid cell, as in LANDCLIM (Schumacher et al., 2004), because stochasticity on the cell level entails diverging species' compositions and thus frequent splits of elements.

The implementation of D2C in a specific model will ultimately depend on the modelled processes, the model drivers and its state variables. Besides the assignment of the processes to the two layers and the specification of the exchanged information, various similarity criteria need to be specified controlling the composition and the dynamics of the non-spatial layer. For merging, for example criteria need to be specified to determine when the state variables that are stored for an element are similar enough. Because TreeMig's state variables are real-valued population densities on a constant number of height classes, a set of similarity thresholds for the density in each height class was required for merging. A model with, for example continuous height values for each individual (or cohort) and a discrete but arbitrary number of individuals would require the specification of similarity thresholds on both the continuous heights and the discrete individuals. For the special case of a model with bounded discrete state variables, Yang et al. (2011) showed a very efficient technical solution with hash maps – data structures with

unique key-value pairs enabling a fast lookup of associations – having a similar approach as D2C. Because the state variables are discrete and bounded, this method can aggregate cells with identical states (identical keys) and thus, no similarity criteria need to be specified. This method, however, will not be applicable whenever any state variable is continuous or unbounded because possible states cannot be uniquely represented with a finite set of elements on the non-spatial layer. An infinite number of possible states necessitates the specification of similarity criteria and an active management of the associations between the two layers as provided by D2C.

5 Conclusions

The implementation of TreeMig-2L and the example simulations demonstrated that D2C can be applied to strongly reduce computational expenses for processes which do not require information on the spatial position of a grid cell relative to other grid cells. D2C is adaptable regarding the criteria used to define similarity for the drivers of a model and for its state variables. The implementation and application with TreeMig-2L indicated that these similarity criteria can be used to adjust the resulting discretisation errors. In the applications, no strong spatial biases were detected. Differences between the original one-layer model and the D2C implementation were in the magnitude of differences among simulations with the original model using different pseudo-random number streams to extrapolate the bioclimate driver, i.e. differences due to interannual bioclimate variability. A large advantage of D2C as opposed to static upscaling

techniques is the possibility to track a migrating species and to account for novel vegetation compositions. Finally, D2C maintains the spatio-temporal resolution of the driver and the simulated processes (as recommended by Bocedi et al., 2012).

D2C is applicable for a broad range of DVMs and under certain conditions probably also for time- and space-discrete models simulating other organisms. D2C might, for example be applicable for models simulating (predominantly) sessile organisms: organisms for whom spatial processes take place on a coarser timescale than local processes (e.g. only in certain life stages or once per year), or organisms with a few migration corridors.

The expected benefit when using D2C to upscale a model depends on different model properties, in particular the ratio of spatially linked to spatially unlinked processes, but also the scale on which the model applies stochasticity and the numerical properties of the state variables. With regards to currently applied DVMs an implementation of D2C together with an efficient seed dispersal algorithm could strongly improve the extent-resolution trade-off, enabling new applications on larger extents or greater numbers of stochastic repetitions to better capture important uncertainties.

Code availability

A tar file containing the source code of TreeMig-2L and the configuration and data files required to run the simulations described in this paper can be requested from jemsnabel@gmail.com or heike.lischke@wsl.ch (please also refer to http://www.wsl.ch/fe/landschaftsdynamik/projekte/TreeMig/index_EN).

The Supplement related to this article is available online at doi:10.5194/gmd-8-3563-2015-supplement.

References

Ahrens, J., Geveci, B., and Law, C.: ParaView: an end-user tool for large-data visualization, in: Visualization Handbook, Elsevier Butterworth–Heinemann, Oxford, UK, 717–732, 2005.

Auger, P., Poggiale, J., and Sánchez, E.: A review on spatial aggregation methods involving several time scales, Ecol. Complex., 10, 12–25, 2012.

Bocedi, G., Pe'er, G., Heikkinen, R. K., Matsinos, Y., and Travis, J. M.: Projecting species' range expansion dynamics: sources of systematic biases when scaling up patterns and processes, Meth. Ecol. Evolut., 3, 1008–1018, 2012.

Bugmann, H., Lindner, M., Lasch, P., Flechsig, M., Ebert, B., and Cramer, W.: Scaling issues in forest succession modelling, Climatic Change, 44, 265–289, 2000.

Epstein, H., Yu, Q., Kaplan, J., and Lischke, H.: Simulating future changes in Arctic and Subarctic vegetation, Comput. Sci. Eng., 9, 12–23, 2007.

Fisher, R., McDowell, N., Purves, D., Moorcroft, P., Sitch, S., Cox, P., Huntingford, C., Meir, P., and Woodward, F. I.: Assessing uncertainties in a second-generation dynamic vegetation model caused by ecological scale limitations, New Phytol., 187, 666–681, 2010.

He, H. S., Mladenoff, D. J., and Crow, T. R.: Linking an ecosystem model and a landscape model to study forest species response to climate warming, Ecol. Model., 114, 213–233, 1999.

He, H. S., Yang, J., Shifley, S. R., and Thompson, F. R.: Challenges of forest landscape modeling – simulating large landscapes and validating results, Landscape Urban Plan., 100, 400–402, 2011.

Hickler, T., Vohland, K., Feehan, J., Miller, P. A., Smith, B., Costa, L., Giesecke, T., Fronzek, S., Carter, T. R., Cramer, W., Kühn, I., and Sykes, M. T.: Projecting the future distribution of European potential natural vegetation zones with a generalized, tree species-based dynamic vegetation model, Global Ecol. Biogeogr., 21, 50–63, 2012.

Huntley, B., Barnard, P., Altwegg, R., Chambers, L., Coetzee, B. W. T., Gibson, L., Hockey, P. A. R., Hole, D. G., Midgley, G. F., Underhill, L. G., and Willis, S. G.: Beyond bioclimatic envelopes: dynamic species' range and abundance modelling in the context of climatic change, Ecography, 33, 621–626, 2010.

Jarvis, A., Reuter, H., Nelson, A., and Guevara, E.: Hole-filled SRTM for the globe Version 4, available from the CGIAR-CSI SRTM 90m Database at: http://srtm.csi.cgiar.org (last access: 9 January 2012), 2008.

Lischke, H., Löffler, T. J., and Fischlin, A.: Aggregation of individual trees and patches in forest succession models: capturing variability with height structured, random, spatial distributions, Theor. Popul. Biol., 54, 213–226, 1998.

Lischke, H., Zimmermann, N. E., Bolliger, J., Rickebusch, S., and Löffler, T. J.: TreeMig: a forest-landscape model for simulating spatio-temporal patterns from stand to landscape scale, Ecol. Model., 199, 409–420, 2006.

Lischke, H., Löffler, T. J., Thornton, P. E., and Zimmermann, N. E.: Model up-scaling in landscape research, chapt. 16, in: A Changing World. Challenges for Landscape Research, Springer, the Netherlands, 259–282, 2007.

Moorcroft, P. R., Hurtt, G. C., and Pacala, S. W.: A Method for Scaling Vegetation Dynamics: The Ecosystem Demography Model (ED), Ecol. Monogr. 71, 557–586, 2001.

Mladenoff, D. J. and He, H. S.: Design, behavior and application of LANDIS, an object-oriented model of forest landscape disturbance and succession, in: Spatial Modeling of Forest Landscape Change: Approaches and Applications, Cambridge University Press, Cambridge, UK, 125–162, 1999.

Nabel, J. E. M. S. and Lischke, H.: Upscaling of spatially explicit and linked time and space discrete models studying vegetation dynamics under climate change, in: Environmental Informatics and Renewable Energies – 27th International Conference on Informatics for Environmental Protection, edited by: Page, B., Fleischer, A. G., Göbel, J., Wohlgemuth, V., Shaker Verlag, Germany, ISBN 978-3-8440-1676-5, 842–850, available at: http://enviroinfo.eu/sites/default/files/pdfs/vol7995/0842.pdf (last access: 19 June 2015), 2013.

Nabel, J. E. M. S., Zurbriggen, N., and Lischke, H.: Interannual climate variability and population density thresholds can have a substantial impact on simulated tree species' migration, Ecol. Model., 257, 88–100, 2013.

Nabel, J. E. M. S., Kirchner, J. W., Zurbriggen, N., Kienast, F., and Lischke, H.: Extrapolation methods for climate time series revisited – spatial correlations in climatic fluctuations influence simulated tree species abundance and migration, Ecol. Complex., 20, 315–324, 2014.

Nakicenovic, N., Alcamo, J., Davis, G., de Vries, B., Fenhann, J., Gaffin, S., Gregory, K., Grübler, A., Jung, T. Y., Kram, T., La Rovere, E. L., Michaelis, L., Mori, S., Morita, T., Pepper, W., Pitcher, H., Price, L., Riahi, K., Roehrl, A., Rogner, H. H., Sankovski, A., Schlesinger, M., Shukla, P., Smith, S., Swart, R., van Rooijen, S., Victor, N., and Dadi, Z.: IPCC Special Report on Emissions Scenarios, Cambridge University Press, Cambridge, UK, 2000.

Neilson, R., Pitelka, L., Solomon, A., Nathan, R., Midgley, G., Fragoso, J., Lischke, H., and Thompson, K.: Forecasting regional to global plant migration in response to climate change, BioScience, 55, 749–759, 2005.

Rickebusch, S., Lischke, H., Bugmann, H., Guisan, A., and Zimmermann, N. E.: Understanding the low-temperature limitations to forest growth through calibration of a forest dynamics model with tree-ring data, Forest Ecol. Manag., 246, 251–263, 2007.

Scheller, R. M. and Mladenoff, D. J.: A forest growth and biomass module for a landscape simulation model, LANDIS: design, validation, and application, Ecol. Model., 180, 211–229, 2004.

Scheller, R. M., Domingo, J. B., Sturtevant, B. R., Williams, J. S., Rudy, A., Gustafson, E. J., and Mladenoff, D. J.: Design, development, and application of LANDIS-II, a spatial landscape simulation model with flexible temporal and spatial resolution, Ecol. Model., 201, 409–419, 2007.

Schumacher, S., Bugmann, H., and Mladenoff, D. J.: Improving the formulation of tree growth and succession in a spatially explicit landscape model, Ecol. Model., 180, 175–194, 2004.

Smith, B., Prentice, I. C., and Sykes, M. T.: Representation of vegetation dynamics in the modelling of terrestrial ecosystems: comparing two contrasting approaches within European climate space, Global Ecol. Biogeogr., 10, 621–637, 2001.

Smith, B., Wårlind, D., Arneth, A., Hickler, T., Leadley, P., Siltberg, J., and Zaehle, S.: Implications of incorporating N cycling and N limitations on primary production in an individual-based dynamic vegetation model, Biogeosciences, 11, 2027–2054, doi:10.5194/bg-11-2027-2014, 2014.

Snell, R. S., Huth, A., Nabel, J. E. M. S., Bocedi, G., Travis, J. M. J., Gravel, D., Bugmann, H., Gutiérrez, A. G., Hickler, T., Higgins, S. I., Scherstjanoi, M., Reineking, B., Zurbriggen, N., and Lischke, H.: Using individual based forest models to simulate species range shifts, Ecography, 37, 1184–1197, 2014.

Sturtevant, B. R., Miranda, B. R., Shinneman, D. J., Gustafson, E. J., and Wolter, P. T.: Comparing modern and presettlement forest dynamics of a subboreal wilderness: does spruce budworm enhance fire risk?, Ecol. Appl., 22, 1278–1296, 2012.

Thuiller, W., Albert, C., Araújo, M. B., Berry, P. M., Cabeza, M., Guisan, A., Hickler, T., Midgley, G. F., Paterson, J., Schurr, F. M., Sykes, M. T., and Zimmermann, N. E.: Predicting global change impacts on plant species' distributions: future challenges, Perspect. Plant Ecol., 9, 137–152, 2008.

Urban, D., Acevedo, M., and Garman, S.: Scaling fine-scale processes to large-scale patterns using models derived from models: meta-models, chap. 4, in: Advances in Spatial Modeling of Forest Landscape Change: Approaches and Applications, Cambridge University Press, Cambridge, UK, 70–98, 1999.

Weidendorfer, J.: Sequential performance analysis with callgrind and kcachegrind, in: Tools for High Performance Computing, Springer, the Netherlands, 93–113, 2008.

Yang, J., He, H. S., Shifley, S. R., Thompson, F. R., and Zhang, Y.: An innovative computer design for modeling forest landscape change in very large spatial extents with fine resolutions, Ecol. Model., 222, 2623–2630, 2011.

Zurbriggen, N.: Avalanche disturbance and regeneration in mountain forests under climate change: experimental and modeling approaches, PhD thesis, Swiss Federal Institute of Technology Zurich, Switzerland, doi:10.3929/ethz-a-009933944, 2013.

Zurbriggen, N., Nabel, J. E. M. S., Teich, M., Bebi, P., and Lischke, H.: Explicit avalanche-forest feedback simulations improve the performance of a coupled avalanche-forest model, Ecol. Complex., 17, 56–66, 2014.