Supplement of Geosci. Model Dev., 15, 3603–3639, 2022
https://doi.org/10.5194/gmd-15-3603-2022-supplement

Geoscientific Model Development
Open Access

Supplement of

Explicitly modelling microtopography in permafrost landscapes in a land surface model (JULES vn5.4_microtopography)

Noah D. Smith et al.

Correspondence to: Noah D. Smith (nds211@exeter.ac.uk)

Supplementary material

Diagnosing the water table position in a partially saturated layer

The implemented lateral flow scheme requires the position of the local water table for a layer to be determined, here referred to as the 'layer head', or ψ_p. Here, we present a new method of determining the water table within a partially saturated layer in a way that is consistent with the assumptions of the cell-centred numerical method for solving the Richards equation used by JULES. The model procedure for finding the layer head is then described. JULES calculates the matric potentials of layers using the layer liquid water as a fraction of saturation, or fractional saturation, θ. Previously, the position of the water table in JULES has been calculated by assuming that θ refers to the saturation at the midpoint of the layer, θ_{mid}. This aligns with the usual calculation of vertical Darcy fluxes using the potential difference between the midpoint of each layer. Setting the potential of the water table to be 0, the soil moisture above the water table is assumed to have the equilibrium profile such that the increase in gravitational potential is balanced by a decreasing matric potential such that the total potential is constant. Using the Brookes-Corey (BC) relation at the layer midpoint,

$$\psi = \psi_s \theta_{mid}^{-b} = z_w \ , \tag{8}$$

where z_w (m) is the height above the water table, ψ_s (m) is the BC matric potential at saturation and b is the Clapp and Hornberger (1978) soil exponent. The result of this calculation is shown by the blue curve in Figure S 1. A similar calculation can be done using the Van Genuchten (VG) relation, for which we use an alternative formulation:

$$\psi = \psi_s \theta_{mid} (\theta_{mid}^{-b-1} - 1) = z_w \ , \tag{9}$$

with similar results to BC, giving the orange curve. This correctly places the water table at the midpoint when saturated, rather than being offset by ψ_s.

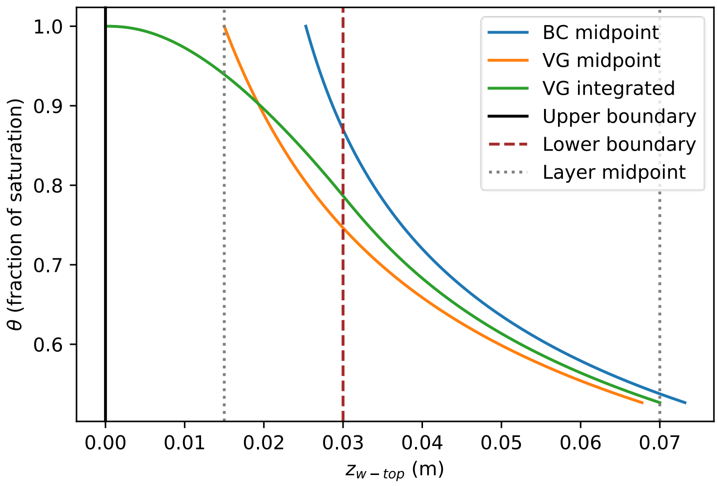

Figure S 1: The fractional saturation of the layer, $\boldsymbol{\theta}$, vs the distance of the water table below the top of the layer $\boldsymbol{z_{w-top}}$. Previously, the position of the water table has been calculated by assuming that $\boldsymbol{\theta}$ refers to the fractional saturation at the midpoint of the layer, resulting in the blue and orange curves, and the water table not exceeding the midpoint of the layer when the layer is saturated. However, in JULES, when considering how much water fits in a layer, $\boldsymbol{\theta}$ refers to the average saturation of the layer. Using this interpretation results in the green curve, where the water table is at the top of the layer when the layer is saturated. Making the simplification that the green curve is approximately linear while the water table is within the layer, the position of the water table within a partially saturated layer can easily be found from the average saturation of the layer.

However, in JULES, θ is also a prognostic variable tracking the average saturation of a layer and is updated according to how much water flows into or out of a layer. This contradicts the calculation of the water table depth and of the direct calculation of the matric potential at the midpoint from θ, as a saturated layer ($\theta = 1$) would have its water table at the top, and conversely the saturation at the midpoint could be 1 even when θ is not. We therefore need a way of calculating the water table from the value of the average saturation of the layer. We start by doing the reverse. If the water table distance below the top of the layer,

z_{w-top} (m), is known, the average fractional saturation of the layer, θ_{avg}, can be found by dividing the integration of the equilibrium VG profile for θ over the layer by the layer thickness, z_T (m), noting that below the water table is saturated.

z_{w-top} <= z_T:

$$\theta_{avg} = 1 - \frac{z_{w-top}}{z_T} + \frac{1}{z_T}\int_{z=0}^{z_{w-top}} \left(\left(\frac{z}{\psi_s}\right)^{\frac{b+1}{b}} + 1\right)^{\frac{1}{-b-1}} dz \,, \tag{10}$$

$z_{w-top} > dz$:

$$\theta_{avg} = \frac{1}{z_T}\int_{z=0}^{z_{w-top}} \left(\left(\frac{z}{\psi_s}\right)^{\frac{b+1}{b}} + 1\right)^{\frac{1}{-b-1}} dz \,, \tag{11}$$

This gives the green curve in Figure S 1. The water table is now at the top of the layer when the layer is saturated. The deviation from either the standard VG or BC calculation is small when the water table is below the bottom boundary of the layer, indicating that $\theta_{avg} \approx \theta_{mid}$ is a good approximation for most cases away from the saturated/unsaturated boundary. Accounting for the difference is therefore only of any importance in wetland environments. To invert the function to calculate the water table depth from θ_{avg} is non-trivial. However, within the ranges of b and ψ_s used at the studied sites, $z_{w-top}(\theta_{avg})$ is approximately linear while the water table is within the layer. We therefore calculate θ_{min}, the value of θ_{avg} when the water table is at the bottom of the layer ($zw_{top} = z_T$) and interpolate between (θ_{min}, z_T) and (1,0) to finally find z_{w-top} as a function of θ_{avg}:

$$\theta_{z_T} = \left(\left(\frac{z_T}{\psi_s}\right)^{\frac{b+1}{b}} + 1\right)^{\frac{1}{-b-1}} \,, \tag{12}$$

$$VG_{int}(\theta) = \int \psi_s \theta(\theta^{-b-1} - 1)\, dz = \psi_s\left(\frac{1}{-b+1}\theta^{-b+1} - \frac{\theta^2}{2}\right) , \tag{13}$$

$$\theta_{min} = \frac{1}{z_T}\left(VG_{int}(1) - VG_{int}(\theta_{z_T})\right) + \theta_{z_T} \,, \tag{14}$$

when $\theta_{avg} < \theta_{min}$:

$$z_{w-top} \approx \frac{z_T(\theta_{avg} - \theta_{min})}{1 - \theta_{min}} \tag{15}$$

The final method is simple to implement and compares well to the full numerical integration ('VG integrated' in Figure S 1), with fully accurate end-point values.

Now that the position of the water table within a partially saturated layer can be found, the model procedure for finding the local water table for a layer is as follows: if a layer is saturated, then each layer above the original layer is checked sequentially until a layer that is partially saturated is found. The position of the water table within this partially saturated layer is then determined, and the total height from the original layer to this is the layer head. If the original layer is partially saturated, the position of the water table within the original layer is determined only if the layer beneath is saturated. If a layer has frozen water present, then the layer head is zero for that layer. If, while sequentially checking the layers above the original layer, a layer which contains frozen water is found, the layer head is the height from the original layer to the bottom of that layer. If the saturation continues to the surface, then the layer head is the height from the original layer to the pond surface (if present).

Limiting soil under- and over-saturation

The standard methods JULES uses to avoid supersaturation or undersaturation as a result of the water flux calculation numerics can result in the unintended consequence of water being passed out of the soil column. This is particularly a problem for freezing saturated soils. Here, we present a method which avoids this problem, and which also integrates with the scheme for simulating lateral fluxes of water.

In some cases, the explicitly and implicitly calculated Darcy fluxes could lead to layers which have a fractional saturation outside the range of 0 to 1. To avoid this, JULES can be set to either adjust the water flux into the top of the layer (l_soil_sat_down = false) or to adjust the flux out of the bottom of the layer (l_soil_sat_down = true) so that the saturation does not go outside this range (Best et al., 2011). The soil water extraction and sub-surface runoff are not limited, as they cannot cause oversaturation, and should be absent in the case of under-saturation. However, both approaches can cause problems, as in both cases the choice of which flux to limit does not account for the direction of water flow. This can cause water to either be able to pass downwards through permafrost, or be ejected upwards out of the soil when water is drawn upwards towards a frozen surface layer when refreezing (Figure S 2 and Figure S 4). In addition, an approach is needed that can also limit the lateral fluxes when necessary.

Here, we solve this problem by only limiting the *incoming* fluxes when a layer will potentially be oversaturated (soilsat-updown). This is done by multiplying them by an appropriate factor such that the layer becomes exactly saturated (Figure S 3, Supplementary). A multiplicative factor was used so that all fluxes scale equally. For a potentially under-saturated layer, any outgoing fluxes (discounting qbase and extraction) are limited in a similar manner. This is necessarily an iterative process as for any layer either flux could be limited, potentially creating the need to adjust a flux in an adjacent layer. The iteration was found to be most efficient sweeping upwards from the lowest layer, and repeats until either there is no appreciable change, or the iteration limit (set at 1 greater than the number of layers) is reached. Generally, only a few of sweeps are required, and the iteration limit is never reached. As lateral flows can be limited, which changes the flows of the paired tile, the limiting code is run a second time for each box after it has been run for all boxes once.

For the most part, limiting the incoming fluxes to a saturated layer (soilsat-updown) results in a soil moisture profile very similar to that of simply limiting fluxes into the top of a saturated layer (l_soil_sat_down = false), and there are only a very few cases where flux out of the top of the soil is avoided (Figure S 2). This does however suggest that l_soil_sat_down = false is in general the more physically realistic scheme. This being said, the few instances where the top of the soil is saturated, and layers of soil exhibit an upward suction into saturated layers are of importance to freezing wetland environments and become more frequent with the inclusion of ponding. Best et al. (2011) point out that, considered globally, l_soil_sat_down = false generates too much runoff in permafrost environments due to correctly allowing a perched water table above a frozen layer. However, the perching of the water table above saturated permafrost is a key behaviour responsible for permafrost wetlands (Avis et al., 2011), so setting l_soil_sat_down = true is undesirable. Our addition of ponding will help reduce runoff, while the effect of macropores (Bechtold et al., 2019) should also be considered, so as to get the right infiltration in the right places for the right reasons. As a side note, an alternative, more physically realistic solution to water being drawn into already saturated frozen layers would be the inclusion of ice segregation.

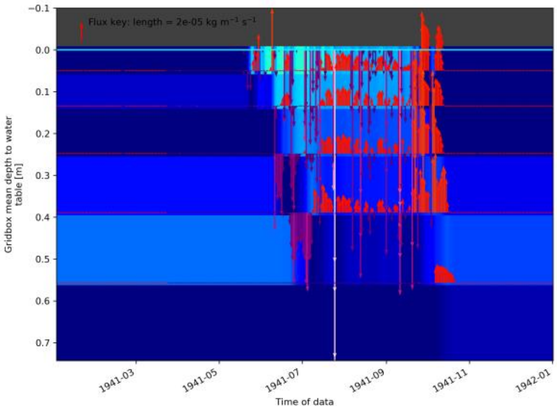

Figure S 2: l_soil_sat_down = false - the excess flux into a saturated layer at the beginning of winter is passed up and out of the soil, as the layer that would become saturated modifies the layer flux at its upper boundary to avoid this, passing the problem to the next layer.

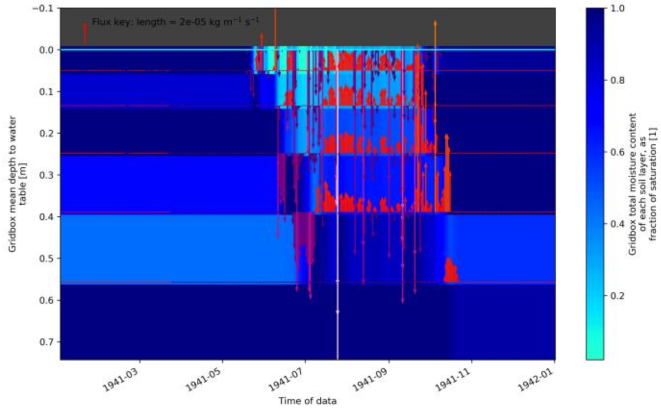

Figure S 3: Soilsat updown - the flux is not passed upwards, as in the case of oversaturation it is always the flux into a saturated layer which is limited, in this case, the flux in from the lower boundary.

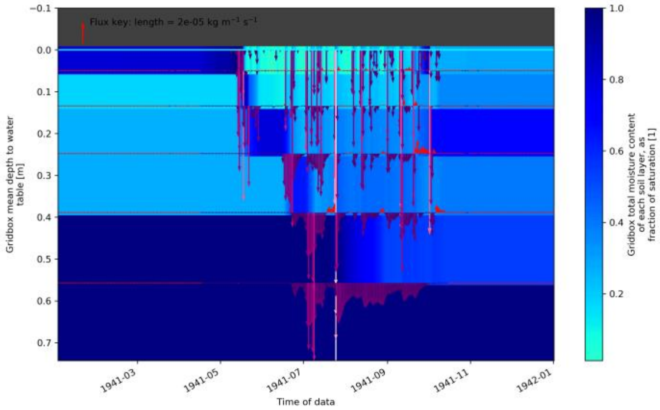

Figure S 4: l_soil_sat_down = true – the downwards flux into a saturated layer is mitigated by altering the water flux at the lower boundary of the saturated layer, causing the same problem in the layer below and so passing water down and out of the soil. The same effect is not seen in soilsat updown.

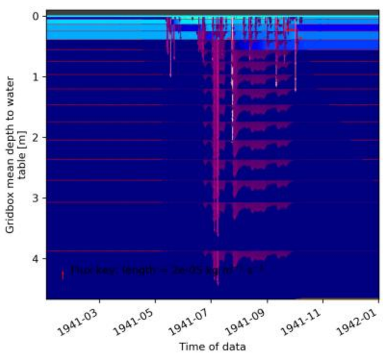

Figure S 5: l_soil_sat_down = true - water is able to fully infiltrate the frozen soil.

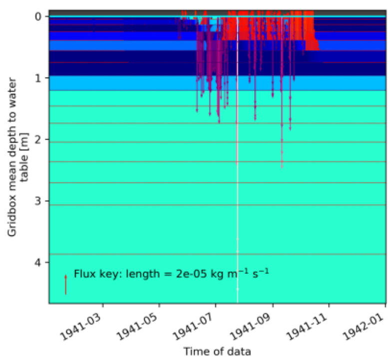

Figure S 6: l_soil_sat_down = false - water is unable to infiltrate a frozen soil.

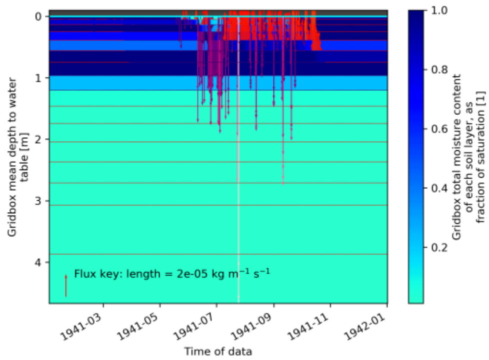

Figure S 7: Soilsat updown - note the similarity to soilsat up.

It is worth noting the differences and similarities between the schemes in Figure S 6, Figure S 7 and Figure S 7. For a fresh spinup, water is only able to infiltrate the frozen soil and saturate the lower soil layers in soilsat down. This means that care needs to be taken to set sthuf (the initial soil wetness) correctly for soilsat up and soilsat updown. Conversely, soilsat up and soilsat updown tend to result in wetter surface layers, though soilsat updown occasionally has a dry surface layer in winter where soilsat up would have a wet one, due to water being able to pass upwards through the saturated ice layer in soilsat up.

Reasoning for choosing horizontal flows over sloped flows

This section expands the discussion within the main text, providing reasoning as to why it was decided to connect layers horizontally, rather than having the connections sloping and connecting each layer to its corresponding layer in the other tile. The strengths and weaknesses of both schemes are also discussed.

Lateral flows of water were introduced into JULES using an approach mirroring the existing calculation of vertical fluxes, where fluxes are calculated based on the difference in matric potential between soil layers due to their level of saturation. Existing functions for calculating the layer matric potentials and hydraulic conductivities are used, and the fluxes interfaced into the existing code for the water balance for each layer. Unlike for the vertical fluxes, no implicit correction is used, as fluxes are assumed to be relatively small. Two options were considered for connecting laterally adjacent layers: a sloped scheme where layers are sequentially connected to their corresponding layer in the neighbouring tile, and a horizontal scheme where a layer is connected to any layers horizontally adjacent to it, taking into account the area of overlap of each connection (Figure 3). The sloped flow scheme was used by Heather Rumbold in JULES 3.2 in a gridded UK run (https://jules.jchmr.org/sites/default/files/Ashton_0.pdf, accessed 3rd June 2021). Here, this sloped flow code was adapted to take into account the areas and geometries associated with permafrost microtopography. The horizontal flow scheme is new to this study. In initial tests with small elevations (or large horizontal scales) both schemes appeared to generate very similar results. The horizontal scheme was eventually chosen and further extended for this work, adding horizontal exchange between the soil of the raised tile and ponded water on the low tile, and integration of the horizontal fluxes with JULES' oversaturation limiting code.

In JULES, the vertical water flux between layers, W $(kgm^{-2}s^{-1})$ is calculated using Darcy's law:

$$W = -K\left(\frac{\Delta\psi_m}{\Delta z}+1\right), \qquad (16)$$

where K $(kgm^{-2}s^{-1})$ is the hydraulic conductivity, Δz (m) is the distance between the centres of the layers, and $\Delta\psi_m$ (m) the difference in matric potential between layers. ψ_m and K are calculated using either the Brooks and Corey (BC) or the Van Genuchten (VG) relations. When calculating lateral flows, if flows are horizontal, the change in gravitational potential with distance $d\psi g / dL = 0$ and hence the 1 can be dropped, or for sloped flows, $d\psi_g/\mathrm{dl} = \sin(\phi)$, where ϕ is the angle of the slope.

Both sloped and horizontal flow schemes have their own strengths and weaknesses. For sloped flows each layer can be matched 1:1 with its corresponding layer. This means that it is simpler to keep track of where water is coming and going, and physically makes sense where landscape changes are continuous, and water flows primarily follow layers of higher conductivity. The weakness is that an imbalance of soil moisture is enforced, and the balancing of the water-table across both tiles becomes unachievable, in part due to the lack of a calculated positive pressure ψ_p in a saturated lower layer to balance the greater ψ_g of the raised layer. Figure S 8 A. illustrates the problem: if the water table is to balance across tiles via sloped flows, then water must flow from a saturated layer (pictured is the flow from layer 3), down the slope through another saturated layer which is further beneath its local water table, and up through further saturated layers to the water table. If layers have different conductivities, and indeed could be frozen, then it is non-trivial to solve the rate at which water flows from one tile to the other, leading to the need for iterative schemes. Furthermore, it cannot be done by the current method used by JULES to solve the Richards equation where ψ is calculated based on the layer saturation, as it requires a consideration of ψ_p. Conversely, for the horizontal flow scheme, a layer may overlap more than one horizontally adjacent layer (Figure S 8B.), meaning flows are harder to keep track of and correct, but there is no danger of flows being dominated by ψ_g. Although the rate may be incorrect, the water table can now balance across tiles, even without considering saturated to saturated flow paths, through horizontal flow from saturated to unsaturated layers. The VG relations are more appropriate for this scheme as they avoid a mismatched ψ_m when both layers are saturated.

The horizontal flow scheme raises the question of what to do with the upper layers of the raised tile, which may have no horizontally adjacent layer. If the water table in the elevated tile is above the surface of the lower tile, and above the level of the surface of any ponded water (if present), water will be able to laterally egress the soil (Figure S 8, Figure 3 C.). Again, this is not possible in the sloped connection scheme. Similarly, if a pond is present on the low tile and the surface of the pond is above the level of the water table in the high tile, water can flow from the pond laterally into the soil (Figure S 8 D.). For these flows it is therefore necessary to determine the level of the local water table. For the high tile, for a particular connection, the local pressure head in the high tile ψ_{ph} (m) is the height of the water table above the midpoint of the vertical overlap. The overlap and midpoint may be different to the layer thickness and midpoint due to the layer being partially saturated, or if only part of the layer is above the pond height (Figure S 8 C. and D.). For water egress, $\Delta\psi = \psi_{ph}$ as there is no matric suction from air. For water ingress, $\Delta\psi$ is the height of the pond above the midpoint of the connection, plus the matric suction of the layer the water is entering.

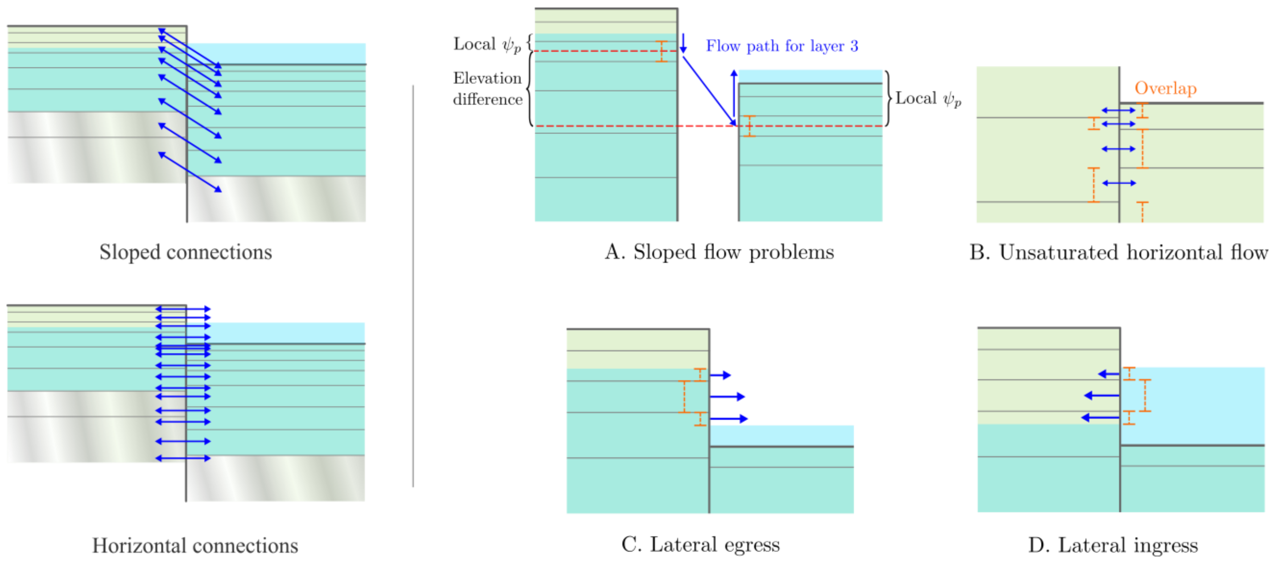

Figure S 8: Paths of flow.

While the local pressure head must be calculated in order for horizontal flows between the soil and the air or the pond and the soil to be modelled, flows are always from a saturated or partially saturated region to an unsaturated one. The model therefore does not implement saturated lateral flows, which are conceptually important for flows between polygons (Wales et al., 2020), meaning that advective flows of heat may not be properly represented. However, the model will still be able to act to balance the water table and moisture potentials between tiles, though the rate with which equilibrium is reached may be different. Any discrepancy in rate is however expected to be less than the usual timescales over which the water table changes, and therefore not a problem.

Additional references for supplement:

Avis, C. A., Weaver, A. J., and Meissner, K. J.: Reduction in areal 65 extent of high-latitude wetlands in response to permafrost thaw, Nat. Geosci., 4, 444–448, https://doi.org/10.1038/ngeo1160, 2011.

Clapp, R. B. and Hornberger, G. M.: Empirical equations for some soil hydraulic properties, Water Resour. Res., 14, 601–604, https://doi.org/10.1029/WR014i004p00601, 1978.