Plant Water Uptake in OpenHydroQual

Transpiration sources and the demand-driven crop block

Abstract

OpenHydroQual offers two ways to take water out of a soil profile through plants. The first attaches a transpiration *source* to each soil layer and scales a reference evapotranspiration by a crop coefficient, a stress function and a per-layer root fraction that the user supplies. The second introduces a *crop block* that owns the atmospheric demand, and one *uptake link* per soil layer that derives its own root fraction from the geometry of the profile and redistributes demand away from layers that have dried out. This document sets out both formulations, states precisely what each can and cannot represent, and documents the components and their parameters. The models referenced throughout are in the repository and are linked as they arise.

Contents

1 Scope and the choice between the two models

Everything in this document concerns *transpiration*: water taken up by roots and lost through stomata. Evaporation directly from the soil surface is a separate process and is left to the `Evapotranspiration_*(Soil)` sources. The crop coefficients used here are FAO-56 *basal* coefficients K_{cb}, which exclude soil evaporation by definition, so the two can be added without double counting.

The two models differ in where the demand lives.

Model 1, transpiration sources

A source is attached to each soil layer. It computes that layer's transpiration directly from a reference evapotranspiration, a crop coefficient, a stress coefficient evaluated on the layer's own moisture state, and a root fraction that the user types in. The layers do not know about one another. If one dries out, its share of the transpiration is simply lost.

Model 2, crop block and uptake links

A `Crop` block computes a single potential transpiration for the whole plant. One `Root_uptake_link` per layer carries a share of that demand. The shares are derived from the profile geometry rather than entered, and a compensation mechanism lets layers that are still wet make up for layers that are stressed.

Model 1 is the lighter of the two and is adequate when the profile stays reasonably uniform or when the root distribution is genuinely known and fixed. Model 2 is what you want when the profile dries unevenly, when roots grow during a season, or when you care about the difference between potential and actual transpiration.

Both live in `resources/plants.json`.

2 Model 1: transpiration sources

2.1 Formulation

For a soil layer i,

$$T_i \;=\; ET_0\, K_{cb}\, K_s(\theta_i)\, f_i\, A, \tag{1}$$

where ET_0 is the reference evapotranspiration, K_{cb} the basal crop coefficient, K_s a water stress coefficient in $[0, 1]$, f_i the fraction of the root system in this layer and A the layer's plan area.

The stress coefficient is a linear ramp in effective saturation between the pressure heads corresponding to field capacity and permanent wilting point. In terms of the van Genuchten parameters that the `Soil` block already exposes,

$$K_s \;=\; \min\!\left(\frac{S_e - S_e(h_{\mathrm{wp}})}{S_e(h_{\mathrm{fc}}) - S_e(h_{\mathrm{wp}})},\, 1\right)^{+}, \qquad S_e(h) = \big(1 + (\alpha h)^n\big)^{-m}, \tag{2}$$

evaluated at $h_{\mathrm{fc}} = 3.3$ m and $h_{\mathrm{wp}} = 150$ m of suction. The superscript $+$ denotes clipping at zero.

2.2 Why the depletion fraction p is not available here

The full FAO-56 stress function introduces a crop-specific readily available water fraction p, so that $K_s = 1$ until depletion exceeds $p\,\mathrm{TAW}$ and falls linearly thereafter. That form is not expressible as a source in OpenHydroQual, and the reason is worth stating because it is a genuine structural limit rather than an oversight.

A source computes its contribution as

$$\text{value} \;=\; \text{coefficient} \times \text{timeseries} \times \text{rate},$$

where `coefficient` is evaluated *against the block* it is attached to, and `rate` and `timeseries` are evaluated *against the source*. A block expression can see the soil state but not the source's parameters; a source expression can see its own parameters but not the soil state. The FAO-56 K_s mixes the two inside a single $\min(\cdot)$, so it can be written in neither place. Setting $p = 0$ recovers the linear ramp above, which is what these sources implement.

Model 2 does not have this restriction, because a link expression can read both of its endpoints. The Feddes function in Section 3.4 is the full four-parameter form.

2.3 Components

Type	Purpose
`Transpiration_Time_Series (Soil)`	ET_0 supplied as a time series.
`Transpiration_Penman (Soil)`	ET_0 computed internally from solar radiation, air temperature, relative humidity and wind speed.

Both take `Kcb` and `root_fraction`. Attach one to each soil layer's `Evapotranspiration` property and make the root fractions sum to one yourself; nothing checks this for you.

2.4 Examples

- `Examples/Plant_Transpiration/rooted_soil_profile.ohq` — three soil layers with root fractions 0.5/0.3/0.2 and ET_0 from a time series.
- `Examples/Plant_Transpiration/rooted_soil_profile_penman.ohq` — the same profile with ET_0 computed by the Penman combination equation from half-hourly weather observations.

The folder is `Examples/Plant_Transpiration`.

3 Model 2: the crop block and root uptake links

3.1 Structure

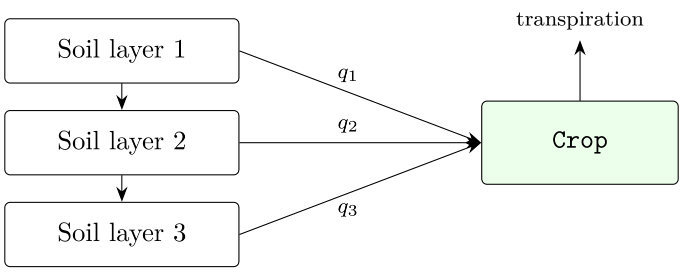

The crop holds a small storage that represents the plant water pathway. Uptake links fill it; transpiration empties it with a short residence time `K_plant`. The storage exists so that supply and demand balance dynamically rather than instantaneously; it is not meant to represent plant water content.

3.2 Atmospheric demand

$$T_p \;=\; ET_0\, K_{cb}\, f_c\, A, \qquad f_c \;=\; 1 - e^{-k\,\mathrm{LAI}}. \tag{3}$$

The cover fraction f_c is Beer's law with extinction coefficient k (about 0.6 for most field crops), so the transpiring fraction of the surface follows canopy development rather than being assumed. Leaf area index is

$$\mathrm{LAI} \;=\; \mathrm{LAI}_{\mathrm{constant}} + \mathrm{LAI}_{\mathrm{series}}, \tag{4}$$

so a static canopy needs only the constant and a seasonal canopy needs only the series. The rooting depth is composed the same way.

3.3 Root distribution

Root density falls off exponentially with depth,

$$\rho(z) \;\propto\; e^{-\lambda z/Z_r}, \qquad 0 \le z \le Z_r, \tag{5}$$

with Z_r the rooting depth and λ a shape parameter. Small λ approaches a uniform distribution; $\lambda \approx 2$ is typical.

Each uptake link integrates this density over its own layer. The link reads the depths to the top and bottom of the layer from the layer's `bottom_elevation` and `depth` and the crop's `surface_elevation`, clipping both at the root tip:

$$z_{\text{top}} = \min\big((E_s - E_b - d)^+,\; Z_r\big), \qquad z_{\text{bot}} = \min\big((E_s - E_b)^+,\; Z_r\big), \tag{6}$$

and then

$$\beta \;=\; \frac{e^{-\lambda z_{\text{top}}/Z_r} - e^{-\lambda z_{\text{bot}}/Z_r}}{1 - e^{-\lambda}}. \tag{7}$$

Two properties of (7) matter in practice. Because the denominator normalises the density over $[0, Z_r]$, the β of a set of contiguous layers spanning the whole root zone sums to one without the user doing anything. And because the clipping is per-link, a layer that the root tip reaches only partway into gets exactly the share of the root system that lies inside it. A layer entirely below the root tip gets $\beta = 0$ and drops out of its own accord.

Nothing here has to be entered by hand, and nothing has to be revised when a layer is added, a layer thickness is edited, or the roots grow during a season.

3.4 Water stress

Each link reduces its share by the Feddes stress function, evaluated on its own layer's pressure head h:

$$\alpha(h) \;=\; \left[\min\!\left(\min\!\left(\frac{h - h_4}{h_3 - h_4},\, 1\right),\, \frac{h - h_1}{h_2 - h_1}\right)\right]^+ . \tag{8}$$

The four heads are the anoxia limit h_1, the wet end of the optimal range h_2, the dry end of the optimal range h_3, and the wilting point h_4, all negative and ordered $h_1 > h_2 > h_3 > h_4$. Uptake is at its full share between h_2 and h_3, ramps down towards zero at the wilting point, and also ramps down towards zero as the layer approaches saturation, which is the waterlogging limb.

Symbol	Property	Default (m)	Meaning
h_1	`h_anoxic`	−0.1	uptake stops above this (anoxia)
h_2	`h_optimal_wet`	−0.25	wet end of the optimal range
h_3	`h_optimal_dry`	−5	dry end of the optimal range
h_4	`h_wilting`	−160	permanent wilting point

3.5 Compensation

Without compensation, layer i delivers $T_p\, \beta_i\, \alpha_i$ and the shortfall in a dry layer is lost. Actual transpiration is then $T_p \sum_i \beta_i \alpha_i$, which can be far below potential even when the profile as a whole holds plenty of water.

Real root systems do not behave that way: uptake shifts towards whatever part of the profile can still supply. SWAP and APEX both model this. The implementation here scales every layer by a common factor so that the total returns to the potential rate,

$$q_i \;=\; T_p\, \beta_i\, \alpha_i\, C, \qquad C \;=\; \frac{1}{\sum_j \beta_j \alpha_j}, \tag{9}$$

which gives $\sum_i q_i = T_p$ whenever any rooted layer can supply.

Obtaining the sum without knowing the layer count

C needs $\sum_j \beta_j\alpha_j$, a sum over the crop's links, but a block expression can only take the *average* of a quantity over its links, via the `.v` suffix. A sum would need the number of links, which the block does not know and which the user should not have to supply.

The way around it uses (7). Since the root fractions sum to one over a complete profile, $\overline{\beta} = 1/N$ and therefore $N = 1/\overline{\beta}$, so

$$\sum_j \beta_j\alpha_j \;=\; N\,\overline{\beta\alpha} \;=\; \frac{\overline{\beta\alpha}}{\overline{\beta}}, \tag{10}$$

which is two link averages and a division. In the template this is

```
sum_beta_alpha      = beta_alpha.v/(beta.v+0.000000001)
compensation_factor = 1/(compensation*_max(sum_beta_alpha;0.001)+(1-compensation))
```

and it self-corrects for however many uptake links are drawn. The `compensation` property switches the mechanism off by collapsing the factor to one.

The compensation is unlimited

C is not capped. As a profile dries, $\sum_j \beta_j\alpha_j$ falls and C grows without bound, concentrating the entire demand on whichever layer is still wet; values around ten are reached in the worked example below. This is the behaviour of full compensation as SWAP and APEX describe it, but some implementations impose a ceiling. If you need one, the practical levers are to raise h_4 so that layers reach $\alpha = 0$ sooner, or to switch compensation off.

3.6 Components

Type	Purpose
`Crop`	Plant block; ET_0 supplied as a time series.
`Crop_Penman`	Plant block; ET_0 computed internally by the Penman combination equation from solar radiation, air temperature, relative humidity and wind speed.
`Root_uptake_link`	Draw from a soil layer to a crop. Carries that layer's share of the demand.

Crop properties

Property	Default	Meaning
`area`	10000	ground area covered by the plant
`surface_elevation`	0	datum from which root depths are measured
`LAI_constant`	3	constant part of the leaf area index
`LAI_timeseries`	—	seasonal part, added to the constant
`k_extinction`	0.6	Beer's law canopy extinction coefficient
`Kcb`	1.0	FAO-56 basal crop coefficient
`root_depth_constant`	1.0	constant part of the rooting depth
`root_depth_timeseries`	—	seasonal part, added to the constant
`root_lambda`	2.0	root density shape parameter
`h_anoxic` ... `h_wilting`	see above	Feddes parameters
`compensation`	1	1 enables compensation, 0 disables it

Property	Default	Meaning
K_plant	0.01 d	residence time of the canopy pathway

Outputs worth plotting

Quantity	Meaning
Tp	potential transpiration, m^3/day
soil_stress	$\sum_j \beta_j \alpha_j$, the root-weighted water availability
relative_transpiration	actual over potential transpiration
cover_fraction	$1 - e^{-k\,\mathrm{LAI}}$
LAI, root_depth	the composed canopy and root state
beta, alpha, flow	on each link: root share, stress, uptake
z_top, z_bottom	on each link: the depth interval it integrated over

Use `soil_stress` as the crop stress index. With compensation on, total uptake equals T_p by construction, so `relative_transpiration` sits at one apart from the lag of the canopy reservoir, which makes it a noisy diagnostic; it is informative when `compensation` is zero. `soil_stress` carries no lag.

3.7 Two rules to observe

1. **Attach only uptake links to a crop block.** The compensation factor is built from `beta.v` and `beta_alpha.v`, which average over *all* of the block's links. Any other link type attached to a `Crop` corrupts the denominator.
2. **Do not wrap the soil layers in a composite.** A composite exposes one attachment point per link type per direction, so a composite containing the soil layers could accept only one uptake link. The crop and its layers must be wired individually.

3.8 Examples

Root water compensation, side by side

`Examples/Crop_Root_Uptake/crop_root_uptake.ohq` places two identical 1.5 m profiles under the same weather, differing only in the crop's `compensation` setting. Rain falls for the first 15 days and then stops for 135 days. At the end of the drought:

	compensation on	compensation off
soil_stress	0.084	0.587
$\sum q_i$ / T_p (m^3/day)	21.9 / 21.9	12.9 / 21.9
θ, layer 1 (0.0–0.3 m)	0.100	0.107
θ, layer 2 (0.3–0.7 m)	0.100	0.141
θ, layer 3 (0.7–1.5 m)	0.119	0.184
profile storage (m^3)	1655	2354

The compensated crop has drawn the whole profile down to wilting point and is still meeting its full demand, taking 98% of it from the deepest layer. The uncompensated crop is held at 59% of potential although its profile still holds 700 m^3 more water, because most of its root mass sits in layers that have dried out.

The rooting depth is 1.0 m against a 1.5 m profile, so the third layer is clipped: its `z_bottom` reports 1.0 m and its β is 0.129 rather than the 0.229 the full layer would carry. The three β still sum to 1.000000.

A corn season with measured weather

`Examples/Crop_Root_Uptake/crop_growing_season_penman.ohq` runs a four-layer profile under a `Crop_Penman` block from 1 April to 1 October 2010, with half-hourly observed solar radiation, temperature, humidity and wind. Leaf area index is driven from a crop calendar (0 at planting, 5.0 in early July, 0.5 at harvest) and the rooting depth grows from 0.2 m to 1.2 m over the same period.

Date	**LAI**	Z_r **(m)**	ET_0	T_p	$\sum \beta$
17 May	1.74	0.55	13.3	9.9	1.000000
16 Jul	4.97	1.20	14.3	15.6	1.000000
14 Sep	2.04	1.20	12.9	10.5	1.000000

Rates are instantaneous, in mm/day, sampled mid-afternoon. The root fractions sum to one at every stage while the rooting front advances, and the share of each layer migrates downwards as the roots grow: the topsoil link holds 0.600 of the root system in May and 0.328 in July, without any of that being entered. Drainage to the water table is negative through the summer, which is capillary rise supporting the crop from below.

A caution on the Penman radiation term. The Penman formulation in OpenHydroQual drives the radiation term with incoming shortwave radiation rather than net radiation, so it overestimates ET_0 by roughly a factor of two on clear days. This example sets `solar_scale_fact` to 0.5 to compensate; that property is the intended calibration knob and should be set from site data before the model is used quantitatively. The same caution applies to `Transpiration_Penman (Soil)` in Model 1 and to the stock `Evapotranspiration_Penman` sources.

Both models, the forcing data and a fuller commentary are in `Examples/Crop_Root_Uptake`.

4 Adding soil evaporation

Neither model evaporates water from the soil surface. To add it, attach an `Evapotranspiration_*(Soil)` source to the top layer. Because K_{cb} is a basal coefficient, the transpiration and evaporation terms are additive and no double counting arises.

One caveat applies to the time-series variant. The source value is formed as `coefficient` × `timeseries` × `rate`, and `Evapotranspiration_Time_Series (Soil)` defines its `rate` as `corr_fact*timeseries`, so the supplied series enters the product twice and the resulting flux is proportional to its square. Neither of the examples in Section 5.5 uses that source. If you need a time-series soil evaporation, the correct `rate` is `corr_fact` alone.

5 Summary of the components

Type	**Kind**	**Notes**
`Transpiration_Time_Series (Soil)`	source	Model 1; ET_0 from a series; user supplies `root_fraction`
`Transpiration_Penman (Soil)`	source	Model 1; ET_0 from weather
`Crop`	block	Model 2; ET_0 from a series
`Crop_Penman`	block	Model 2; ET_0 from weather
`Root_uptake_link`	link	Model 2; soil layer → crop
`Plant`	block	legacy soil–plant–atmosphere continuum block

Type	Kind	Notes
`Soil_to_plant_link`	link	legacy, used with `Plant`

All are defined in `resources/plants.json`.