



Development of a multi-layer canopy model for E3SM Land Model with support for heterogeneous computing

Gautam Bisht^{a,*}, William J. Riley^b, Richard Tran Mills^c

^a Pacific Northwest National Laboratory, 902 Battelle Blvd, Richland, WA, 99354, USA

^b Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA, 94720, USA

^c Argonne National Laboratory, 9700 S Cass Ave, Lemont, IL, 60439, USA

ARTICLE INFO

Keywords:

Multi-layer canopy model

ABSTRACT

The vertical structure of vegetation canopies creates micro-climates. However, the land components of most Earth System Models, including the Energy Exascale Earth System Model (E3SM), typically neglect vertical canopy structure by using a single layer big-leaf representation to simulate water, CO₂, and energy exchanges between the land and the atmosphere. In this study, we developed a Multi-Layer Canopy Model for the E3SM Land Model to resolve the micro-climate created by vegetation canopies. The model developed in this study re-implements the CLM-ml_v1 to support heterogeneous computing architectures consisting of CPUs and GPUs and includes three additional optimization-based stomatal conductance models. The use of Portable, Extensible Toolkit for Scientific Computation provides a speedup of 25–50 times on a GPU relative to a CPU. The numerical implementation of the model was verified against CLM-ml_v1 for a month-long simulation using data from the Ameriflux US-University of Michigan Biological Station site. Model structural uncertainty was explored by performing control simulations for five stomatal conductance models that exclude and include the control of plant hydrodynamics (PHD) on photosynthesis. The bias in simulated sensible and latent heat fluxes was lower when PHD was accounted for in the model. Additionally, six idealized simulations were performed to study the impact of three environmental variables (i.e. air temperature, atmospheric CO₂, and soil moisture) on canopy processes (i.e. net CO₂ assimilation, leaf temperature, and leaf water potential). Increasing air temperature reduced net CO₂ assimilation and increased air temperature. Net CO₂ assimilation increased at higher atmospheric CO₂, while decreasing soil moisture resulted in lower leaf water potential.

1. Introduction

The vertical structure of canopies creates micro-climates and can affect ecosystem response to climate change. Wind speed, air temperature, and scalar concentrations, such as water vapor and CO₂, vary vertically within a vegetation canopy [1]. Under unstable atmospheric conditions and during daytime, CO₂ is higher above the canopy and decreases within the canopy due to the uptake of CO₂ by plants through photosynthesis and increases again near the soil surface due to soil respiration [2]. Under similar atmospheric conditions, the water vapor in the canopy and near the soil surface is higher than in the atmosphere due to evapotranspiration. Using paired measurements, De Frenne et al. [3] showed that vegetation cover decreases the maximum temperature within the understory, while simultaneously increasing the minimum temperature. They reported that the difference in temperature between above and within the canopy becomes larger as the above-canopy air temperature increases, and this temperature difference can be of greater

magnitude than the increase in land surface temperature over the last century.

The vertical structure of canopies is neglected in most current global Land Surface Models (LSMs), including the Energy Exascale Earth System Model (E3SM) Land Model (ELM), although a few next-generation models are addressing this shortcoming. The vegetation canopy in current LSMs is typically represented by a single big-leaf that, along with a surface soil layer, exchanges water, energy, and CO₂ fluxes with the atmosphere. The Monin–Obukhov similarity theory (MOST) is used to describe the vertical transport of momentum, energy, and water vapor between the canopy and overlying atmosphere. The variation of light is typically accounted for by including a sunlit and shaded big-leaf, both of which are assumed to be at the same canopy height. Plant hydrodynamics (PHD) has been included in many site-level and regional ecosystem models [4–9], but only recently has PHD been included in global LSMs [10,11]. Recently, multi-layer canopy

* Corresponding author.

E-mail addresses: gautam.bisht@pnnl.gov (G. Bisht), wjriley@lbl.gov (W.J. Riley), rtmills@anl.gov (R.T. Mills).

URLs: <https://www.gautambisht.org> (G. Bisht), <https://riley.lbl.gov> (W.J. Riley), <https://climatemodeling.org/~rmills> (R.T. Mills).

<https://doi.org/10.1016/j.jocs.2024.102366>

Received 8 January 2024; Received in revised form 23 April 2024; Accepted 15 June 2024

Available online 19 June 2024

1877-7503/© 2024 Argonne National Laboratory, Battelle Memorial Institute and Lawrence Berkeley National Laboratory. Published by Elsevier B.V.

models (MLCMs) have been added in LSMs [12–15]. MLCMs vertically resolve radiation (solar and longwave), photosynthesis, and leaf temperatures within canopies, though the canopy air space (CAS) can be represented with one or multiple layers. The standard boundary-layer flux–gradient relationships used in the MOST turbulence scheme are violated within the canopy and in the roughness sublayer (RSL) near the top of the canopy. MLCMs employ turbulence schemes that are valid for the turbulent transport of momentum, energy, and scalars within the canopy and RSL. The turbulence scheme of Massman and Weil [16] is used in ED2.2 [15] and ORCHIDEE-CAN v1.0 [17], while the [1,18] turbulence scheme is used in CLM-mlv0 [14] and CLM-ml v1 [19].

The vertical uptake of CO₂ via photosynthesis within a canopy needs to be modeled accurately to match observed profiles of CO₂; multiple stomatal conductance models (SCMs) have been developed to account for the stomatal control of photosynthesis. SCMs use semi-empirical relationships or optimization theory and may additionally account for PHD. A commonly used SCM is the Ball–Berry (BB) model that uses a semi-empirical approach for modeling the dependence of stomatal conductance on vapor pressure deficit by using fractional humidity at the leaf surface [20]. Leuning [21] modified the BB model by adjusting for the CO₂ compensation point and directly included dependence on vapor pressure deficit. Alternatively, optimization theory has also been used to model stomatal conductance, g_s , in which stomata maximize the carbon gain relative to a cost associated with the carbon gain. The commonly used objective functions in optimization-based SCMs are water-use efficiency ($WUE = \Delta A_n / \Delta E$) and intrinsic water-use efficiency ($iWUE = \Delta A_n / \Delta g_s$), defined as the marginal carbon gain (ΔA_n) with respect to water loss (ΔE) and stomatal opening (Δg_s), respectively [12,22–24]. A few semi-empirical and $iWUE/WUE$ SCMs additionally include the control of PHD on g_s [25–27]. Bonan et al. [12] developed an optimization-based g_s model that simultaneously satisfied two constraints of $iWUE$ and leaf water potential to be greater than a minimum threshold. An objective function other than WUE and $iWUE$ has also been used in optimization-based SCMs. Wang et al. [28] evaluated 10 optimization-based SCMs for seven criteria that tested for a valid mathematical solution and realistic biological stomatal responses to the environment that were consistent with empirical studies. They found that only four SCMs met all seven criteria. Two of those four models used the loss of xylem hydraulic conductivity as the marginal cost for carbon gain [28,29], while the other two models used non-stomatal limitations on photosynthesis as the marginal cost for carbon gain [30,31]. Stomatal conductance models that use optimization theory and include a hydraulic constraints have shown superior skill in reproducing observations [12,28,32,33].

Development of new process models within LSMs should include support for the upcoming heterogeneous high-performance architectures and have a flexible core to explore model structural uncertainty. High-performance computing architectures are transitioning from homogeneous CPUs-based systems to heterogeneous systems that include CPUs and GPU accelerators. Currently, there are many GPU programming models such as Kokkos [34], RAJA [35], CUDA [36], ROCm [37], and OpenCL [38]; and it is not feasible for LSMs to include native support for all of the programming models. The profusion of these programming models has led the Exascale Computing Project to develop performance portable programming models that aim to provide high-level abstractions for mapping code to appropriate hardware. The Portable Extensible Toolkit for Scientific Computation (PETSc) library, supported by the Exascale Computing Project, provides scalable numerical solvers [39] and allows application code (e.g. LSMs) to be agnostic of the underlying computer architectures [40]. Additionally, PETSc provides a flexible framework to assemble individual physics problems into a single tightly coupled multiphysics system via DM-Composite [41]. The DMComposite has been previously used in the open-source, multi-physics problem (MPP) library to flexibly couple multiple biophysics problems within multiple domains, e.g., the transport of water through the soil, root, and xylem; and the transport of water and energy in the subsurface [42].

Objectives of this work include developing a model that resolves micro-climate created by the vertical structure of the vegetation canopy and providing support for heterogeneous computing architectures. In this study, a standalone MLCM is developed for ELM, hereafter referred to as ELM-MLCMv1. The new model is based on CLM-ml v1 and includes three additional optimization-based SCMs as compared to CLM-ml v1. PETSc is used in ELM-MLCMv1 to provide numerical solutions of discretized equations, flexible model coupling, and performance portability. First, the numerical implementation of the model was verified by comparing results against CLM-ml v1 for a month-long simulation performed using data from the AmeriFlux station at the University of Michigan Biological (UMB). Second, the impact of SCM model structure on the simulation of biophysical processes was studied. Five SCMs, calibrated against site-level observations of sensible and latent heat fluxes, were then used to study the response of vegetation under three idealized climate change (CC) scenarios: (i) higher air temperature, (ii) higher atmospheric CO₂ concentrations, and (iii) drier soil moisture. Third, the performance portability of ELM-MLCMv1 was studied by performing simulations across a range of problem sizes on two GPU systems that used NVIDIA and AMD GPUs. The model developed in this study provides the foundation for ELM to undertake research to better understand the role of vertical canopy heterogeneity on the ecosystem climate sensitivity.

2. Model description

The ELM-MLCMv1 (Fig. 1) accounts for sunlit and shaded leaves at each canopy level and includes the following four sub-models: (i) shortwave and longwave radiation, (ii) stomatal conductance, (iii) RSL parameterization, and (iv) transport of heat and water vapor in the CAS. The radiation model and RSL parameterization lump sunlit and shaded leaves as one at each canopy layer, while the other two sub-models explicitly account for sunlit and shaded leaves. Similar to CLM-ml v1, it is assumed that water from the soil is transported to each leaf via unconnected xylems (Fig. 1e). The ELM-MLCMv1 includes the same parameterizations as those of the CLM-ml v1, whose model description is described in detail [12,14,43]. With respect to CLM-ml v1, ELM-MLCMv1 includes three additional optimization-based SCMs that are described in detail in Section 2.2. For the sake of completeness, we have summarized the ELM-MLCMv1 physics in a technical guide that is available at <https://mpp-lsm.github.io/mpp-doc/>, as well as been archived [44]. The vertical profiles of the leaf and stem of a plant are described using a beta distribution. The abbreviations used in this study are listed in Table 1.

2.1. Shortwave and longwave radiation sub-model

The shortwave radiation sub-model is based on Norman [45] and accounts for direct beam and diffuse radiation. At each canopy level, the sunlit and shaded leaves are first lumped into a single leaf. The lumped leaf intercepts a fraction of the beam and diffuse radiation incident on the leaf surface. The intercepted radiation is absorbed or scattered as diffuse radiation. The radiation at each canopy position thus consists of (i) non-intercepted beam radiation, (ii) non-intercepted diffuse radiation, and (iii) intercepted beam and intercepted diffuse radiation that is scattered upward and downward. The shortwave radiation model resolves radiation in the visible and near-infrared bands. For N canopy layers, the discretized model leads to a linear system of equations with $(N + 1) \times 2$ unknowns corresponding to upward and downward radiation at each canopy level and the soil surface. The simulated radiation absorbed by a lumped leaf at each canopy level is disaggregated for sunlit and shaded leaves. It is assumed that shaded leaves absorb diffuse radiation, while sunlit leaves absorb beam and diffuse radiation. Further details about the shortwave radiation model are provided in the technical guide.

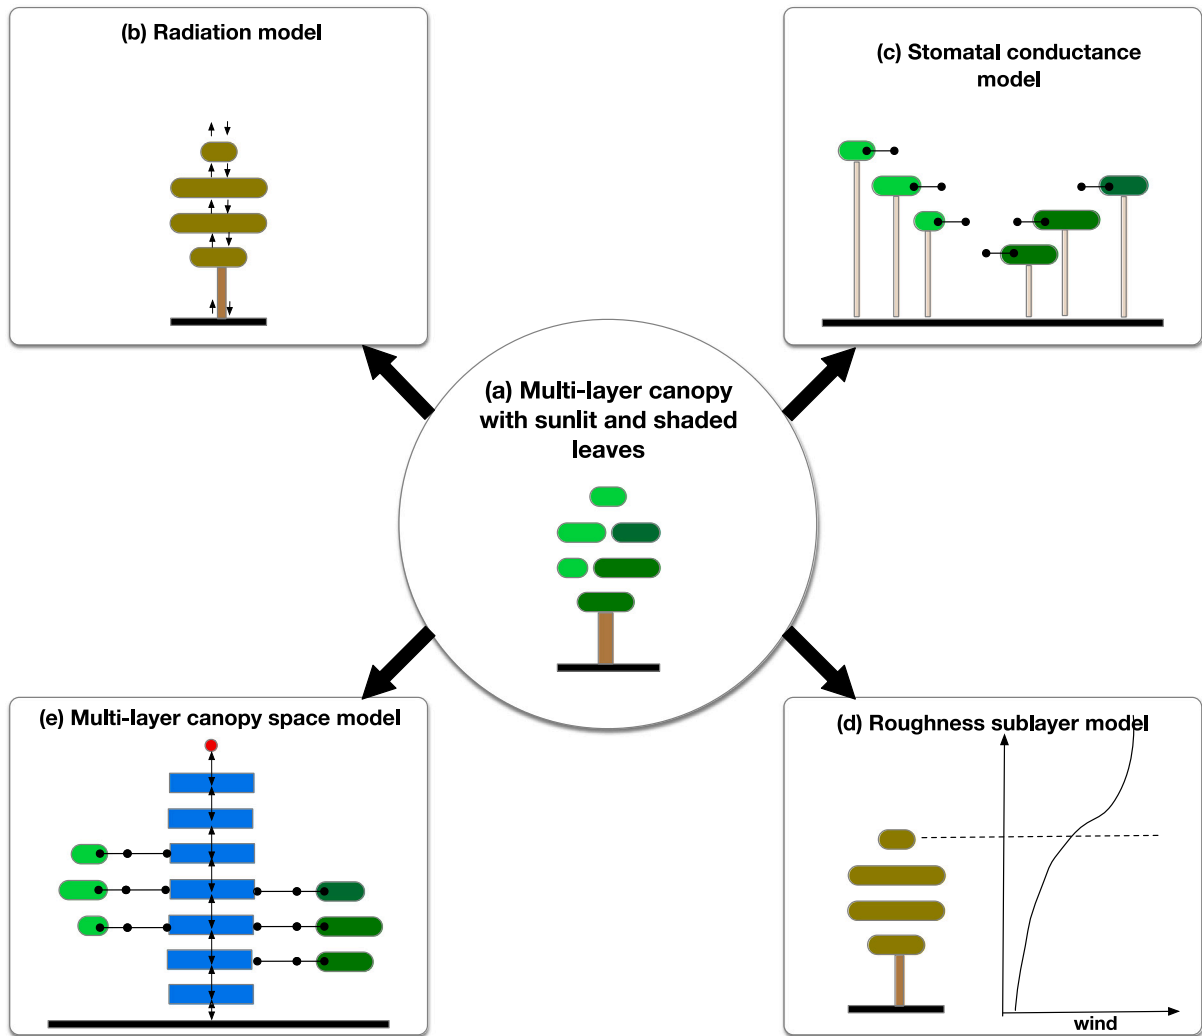


Fig. 1. Schematic representation of ELM-MLCMv1 that includes (a) leaves at multiple layers that are sunlit (light green) and shaded, (b) vertically-resolved models for shortwave and longwave radiation that combine sunlit and shaded leaves into a single leaf (brown) at each canopy level, (c) model for stomatal conductance of sunlit and shaded leaves, (d) roughness sublayer parameterization of turbulent flow within and above vegetation canopy, and (e) model for the transport of heat and water vapor through the canopy air space.

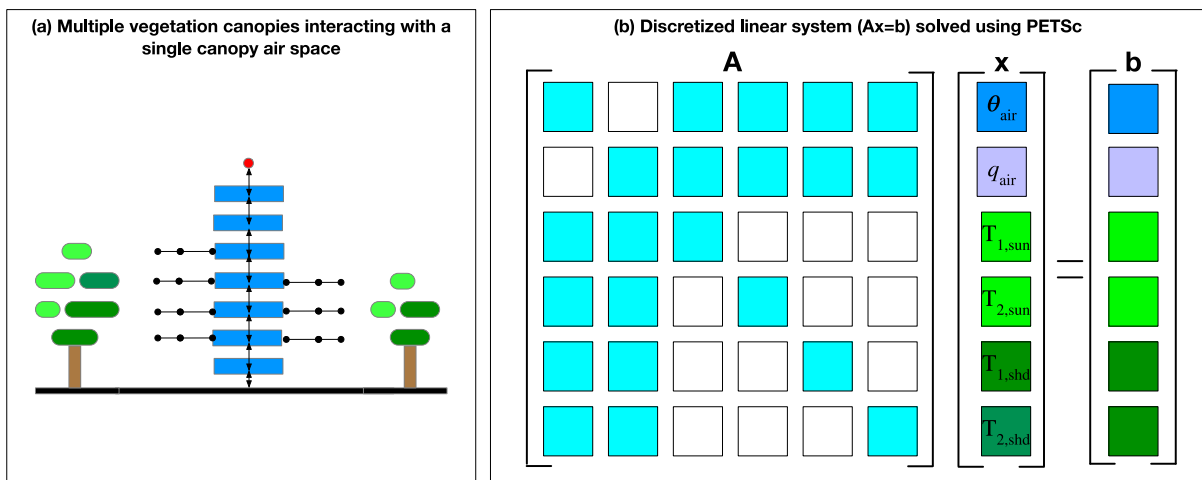


Fig. 2. (a) Schematic representation of two vegetation canopies that are interacting with a single canopy air space (CAS). The model configuration comprises of four governing equations and comprises of six unknown variables that include CAS temperature (θ_{air}), CAS water vapor (q_{air}), sunlit leaf temperatures ($T_{1,sun}$ and $T_{2,sun}$ corresponding to two trees), and shaded leaf temperatures ($T_{1,shd}$ and $T_{2,shd}$ corresponding to two trees). (b) The discretized linear system is flexibly assembled using PETSc's DMComposite and solved using PETSc's KSP solver.




Fig. 3. The vertical profile of sunlit (in light green) and shaded (in dark green) LAI at (a) 8:00 am and (b) 12:00 pm. The vertical layer thickness is 0.5 m.

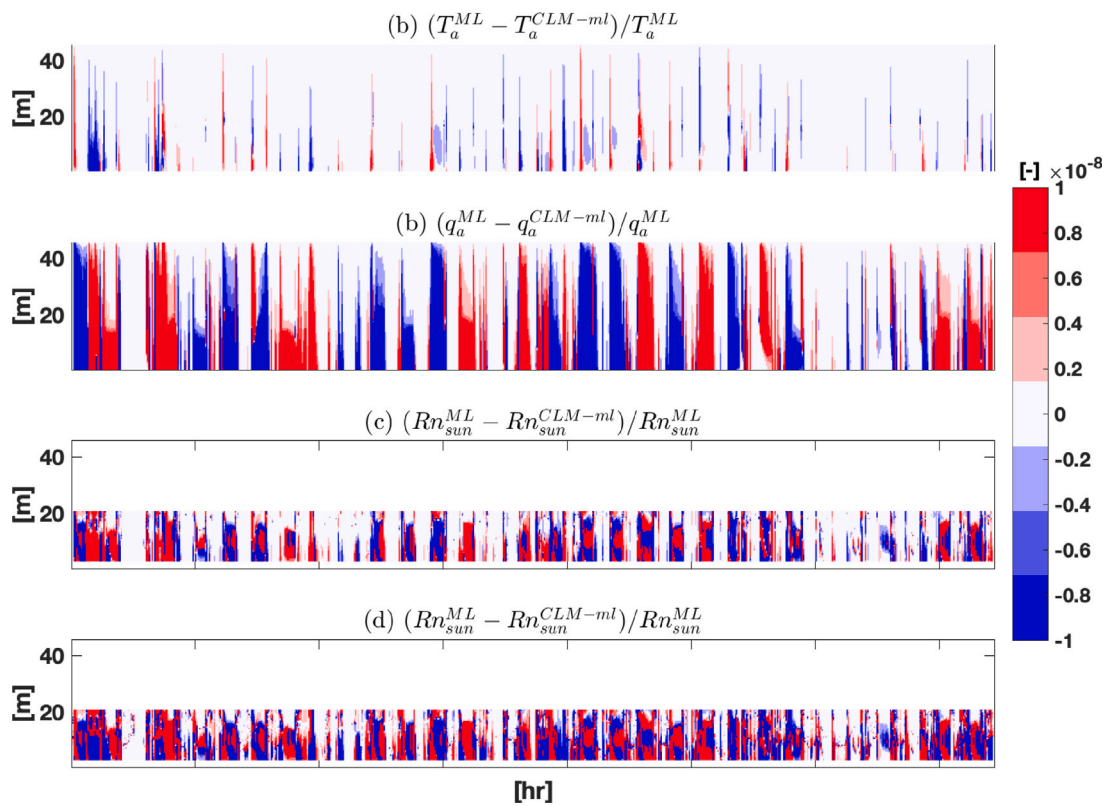


Fig. 4. Time series of the vertical profile of relative error for Medlyn SCM with respect to CLM-ml v1 for (a) air temperature, (b) water vapor, (c) net radiation for sunlit leaf, and (d) net radiation for shaded leaf. simulated net radiation, R_{net} . The results are for the Bonan14 SCM.

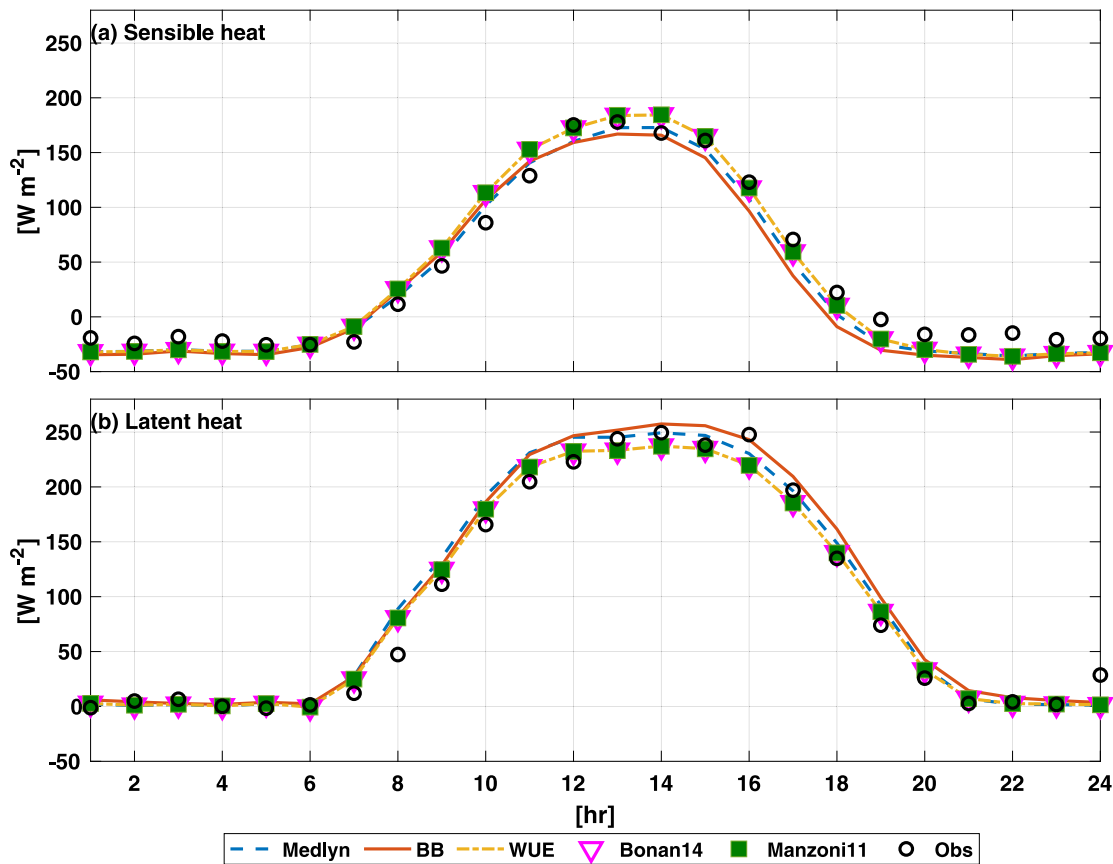


Fig. 5. Simulated monthly average diurnal cycle of (a) sensible heat flux, H , and (b) latent flux, LH , for five SCMs. The observations are shown in circle symbols.

The longwave radiation sub-model is also based on Norman [45]. Unlike the shortwave radiation model, the longwave radiation model only includes diffuse radiation. The emitted longwave radiation of sunlit and shaded leaves is weighted by their fractions at each canopy layer. The sunlight and shaded leaf absorb equal longwave radiation per unit leaf area. The discretized model also leads to a linear system of equations with $(N + 1) \times 2$ unknowns corresponding to upward and downward radiation at each canopy level and the soil surface. The details about the longwave radiation model are provided in the technical guide.

2.2. Stomatal conductance sub-model

Photosynthesis is modeled assuming the biological demand for CO_2 uptake is met by the diffusive transport of CO_2 from the CAS into the stomata. In such a modeling approach, there are two equations with three unknowns. The two equations describe the biological uptake and diffusive transport of CO_2 , and the three unknowns include net CO_2 assimilation, CO_2 concentration within the stomata, and stomatal conductance (see the technical guide for details). Thus, in order to close the system of equations, a third equation for stomatal conductance is needed. ELM-MLCMv1 includes five previously published SCMs that use semi-empirical relationship or optimization theory and may additionally account for PHD (see Table 2). All five SCMs result in a nonlinear equation that is solved for sunlight and shaded leaves at each canopy level. A detailed description of all supported SCMs and their numerical solutions are presented in the technical guide. The model for the transport of CO_2 from the CAS into the stomata additionally includes parameterizations for the leaf boundary layer.

Table 1
Abbreviations used in this study.

Abbreviation	Refers to
BB	Ball-Berry stomatal conductance model
Bonan14	Stomatal conductance model of Bonan et al. [12]
CAS	Canopy air space
CC	Climate change
CLM-ml v0	Community Land Model - multi layer version 0
CLM-ml v1	Community Land Model - multi layer version 1
E3SM	Energy Exascale Earth System Model
ELM	E3SM Land Model
GE	Governing equation
iWUE	intrinsic water-use efficiency
KGE	Kling-Gupta Efficiency
LAD	Leaf area density
LAI	Leaf area index
LSM	Land Surface Model
MLCM	Multi layer canopy model
MOST	Monn-Obukbov Similarity Theory
MPP	Multi-physics problem
NSE	Nash-Sutcliffe efficiency
OLCF	Oak Ridge Leadership Computing Facility
PETSc	Portable Extensible Toolkit for Scientific Computation
PHD	Plant hydrodynamics
RMSE	Root mean square error
RSL	Roughness sublayer
SAI	Stem area index
SCM	Stomatal conductance model
SoE	System of equations
UBM	University of Michigan Biological
WUE	Water-use efficiency

2.3. Turbulent transport sub-model

Unlike big-leaf canopy models, MLCMs require models to describe vertical profiles of scalars (i.e. wind, air temperature, and water vapor) within the canopy. ELM-MLCMv1 uses the model of Harman and




Fig. 6. Simulated monthly average vertical profile of (a) leaf water potential, ψ_{leaf} , (b) air temperature, T_{air} , and (c) air vapor pressure, q_{air} . The black line indicates the average of five SCMs, while the shading denotes the standard deviation among five SCMs.

Table 2

Description of SCMs supported in ELM-MLCMv1. The BB and Medlyn SCM are included within CLM-ml v1.

Name	Description	Reference
BB	Semi-empirical approach without accounting for PHD	[20]
Medlyn	Semi-empirical approach without accounting for PHD	[46]
WUE	Optimization theory without accounting for PHD	[24]
Bonan14	Co-optimization of two constraints including a PHD constraint	[12]
Manzoni11	WUE-based approach with down regulation based on PHD	[25]

Table 3

Simulation configurations for studying the impact on canopy processes from three environmental variables that include air temperature, atmospheric CO_2 , and soil moisture.

	Name	Air temperature	Atmospheric CO_2	Soil moisture
1	Control	Default	367 ppm	Default
2	2 K	Default + 2 K	367 ppm	Default
3	5 K	Default + 5 K	367 ppm	Default
4	467 ppm	Default	467 ppm	Default
5	567 ppm	Default	567 ppm	Default
6	sm85	Default	367 ppm	$0.85 \times \text{Default}$
7	sm70	Default	367 ppm	$0.70 \times \text{Default}$

Finnigan [1,18] that provides a unified theory of turbulent transport within and above the canopy. The scalars within RSL of the canopy follow an exponential profile, while the log profile above the canopy deviates from the MOST predicted profile. Continuity of scalar values is assumed at the canopy height which leads to a system of equations that is solved iteratively. A detailed description of the RSL model is presented the technical guide. Similar to the radiation models, the RSL model lumps sunlit and shaded leaves together at each canopy level.

2.4. Multi-layer canopy air space sub-model

ELM-MLCMv1 simulates the transport of energy and water vapor within the CAS using a one-dimensional vertical diffusion model. The conservation of energy for sunlit and shaded leaves is modeled explicitly at each canopy level. The sensible and latent heat fluxes from the leaves tightly couple CAS water and energy transport equations with leaf energy balance equations leading to a system of four tightly coupled equations. The water vapor and energy fluxes between the leaf and CAS take into account stomatal and leaf boundary conductances, while the vertical fluxes of water and energy account for aerodynamic turbulence within and above the canopy. The equations are solved using finite-volume spatial discretization and an implicit time integration scheme. Similar to CLM-ml v1, the non-linear source and sink terms of water and energy are linearized leading to a linear system of equations. A detailed description of the MLC sub-model is presented in the technical guide.

3. Numerical implementation

The discretized equations of ELM-MLCMv1's four sub-models result in linear and non-linear equations that are solved using PETSc's KSP

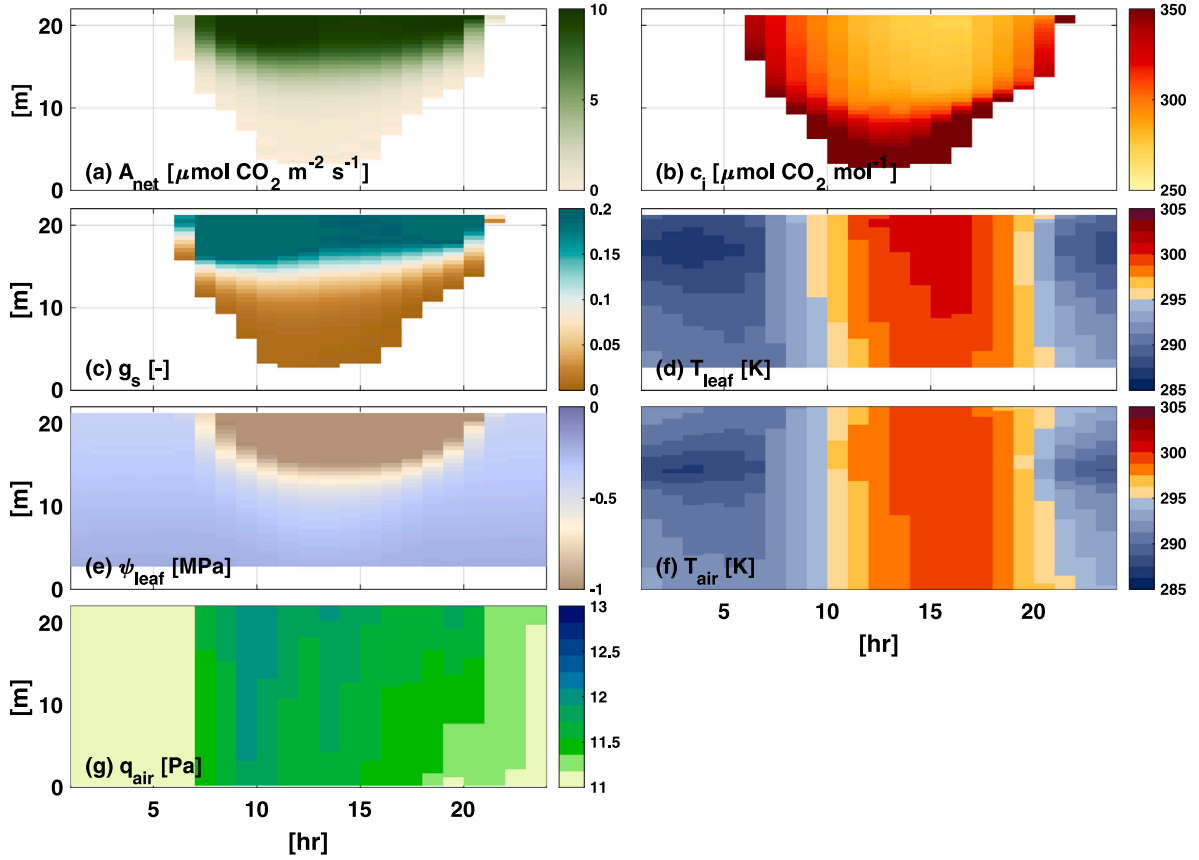


Fig. 7. Simulated vertical profile of average diurnal cycle of (a) net assimilation, A_{net} , (b) intercellular CO_2 , c_i , (c) stomatal conductance, g_s , (d) leaf temperature, T_{leaf} , (e) leaf water potential, ψ_{leaf} , (f) air temperature, T_{air} , and (g) water vapor, q_{air} , for the Medlyn SCM.

Table 4

Mean relative error between CLM-ml_v1 and ML for net radiation of sunlit ($Rn_{\epsilon_{sun}}$) and shaded ($Rn_{\epsilon_{sha}}$) leaves, stomatal conductance of sunlit ($g_{\epsilon_{sun}}$) and shaded ($g_{\epsilon_{sha}}$) leaves, air temperature (T_a), and water vapor (q_a).

SCM	$\Delta Rn_{\epsilon_{sun}}/Rn_{\epsilon_{sun}}^{ML}$	$\Delta Rn_{\epsilon_{sha}}/Rn_{\epsilon_{sha}}^{ML}$	$\Delta g_{\epsilon_{sun}}/g_{\epsilon_{sun}}^{ML}$	$\Delta g_{\epsilon_{sha}}/g_{\epsilon_{sha}}^{ML}$	$\Delta T_a/T_a^{ML}$	$\Delta q_a/q_a^{ML}$
BB	+3.06e-09	+6.96e-09	+7.69e-06	-4.12e-09	-1.17e-11	+3.47e-10
Medlyn	+1.25e-07	-1.69e-05	+1.35e-07	+2.85e-07	-3.30e-10	+2.83e-09

Table 5

Evaluation metrics for the control simulation for five SCMs. The metrics in the top row are for simulated monthly-averaged diurnal cycle of sensible and latent heat flux, while the metrics in the bottom row within parentheses are for simulated hourly heat fluxes.

SCM	Sensible heat flux					Latent heat flux				
	Bias	RMSE	R ²	KGE	NSE	Bias	RMSE	R ²	KGE	NSE
Medlyn	-7.3	13.5	0.99	0.80	0.97	6.8	16.2	0.99	0.92	0.97
	(-9.1)	(49.3)	(0.87)	(0.73)	(0.75)	(9.2)	(53.3)	(0.89)	(0.85)	(0.78)
BB	-10.3	18.0	0.98	0.73	0.94	10.4	16.4	0.99	0.88	0.97
	(-11.5)	(47.9)	(0.89)	(0.69)	(0.77)	(12.1)	(51.3)	(0.90)	(0.83)	(0.80)
WUE	-2.3	14.0	0.99	0.88	0.96	1.3	13.2	0.99	0.97	0.98
	(-3.8)	(47.3)	(0.88)	(0.85)	(0.77)	(3.3)	(50.6)	(0.90)	(0.86)	(0.80)
Bonan14	-2.3	14.0	0.99	0.88	0.96	1.3	13.2	0.99	0.97	0.98
	(-3.8)	(47.3)	(0.88)	(0.85)	(0.77)	(3.3)	(50.6)	(0.90)	(0.86)	(0.80)
Manzoni11	-2.2	14.0	0.99	0.88	0.96	1.2	13.2	0.99	0.97	0.98
	(-3.7)	(47.3)	(0.88)	(0.85)	(0.77)	(3.3)	(50.6)	(0.90)	(0.86)	(0.80)

and SNES solvers, respectively. Unlike CLM-ml v1, ELM-MLCMv1 allows for multiple vegetation canopies to simultaneously interact with a CAS (Fig. 2). The DMComposite of PETSc, a special type of the DM (“Domain Management”) abstraction in PETSc, is used to provide a flexible framework to solve equations for the transport of energy and water within CAS which are tightly coupled to the energy balance equation for sunlit and shaded leaves. Similar to the previous use of

DMComposite in the MPP library [42], the multi-layer CAS model includes four governing equation (GE) objects: (1) CAS temperature, (2) CAS vapor pressure, (3) sunlit leaf temperature, and (4) shaded leaf temperature. The four GE objects make up a system of equations SoE object that is used to obtain the numerical solution via DMComposite. The details about the use of GE and SoE is presented in [42]. A schematic representation of two vegetation canopies sharing a CAS is

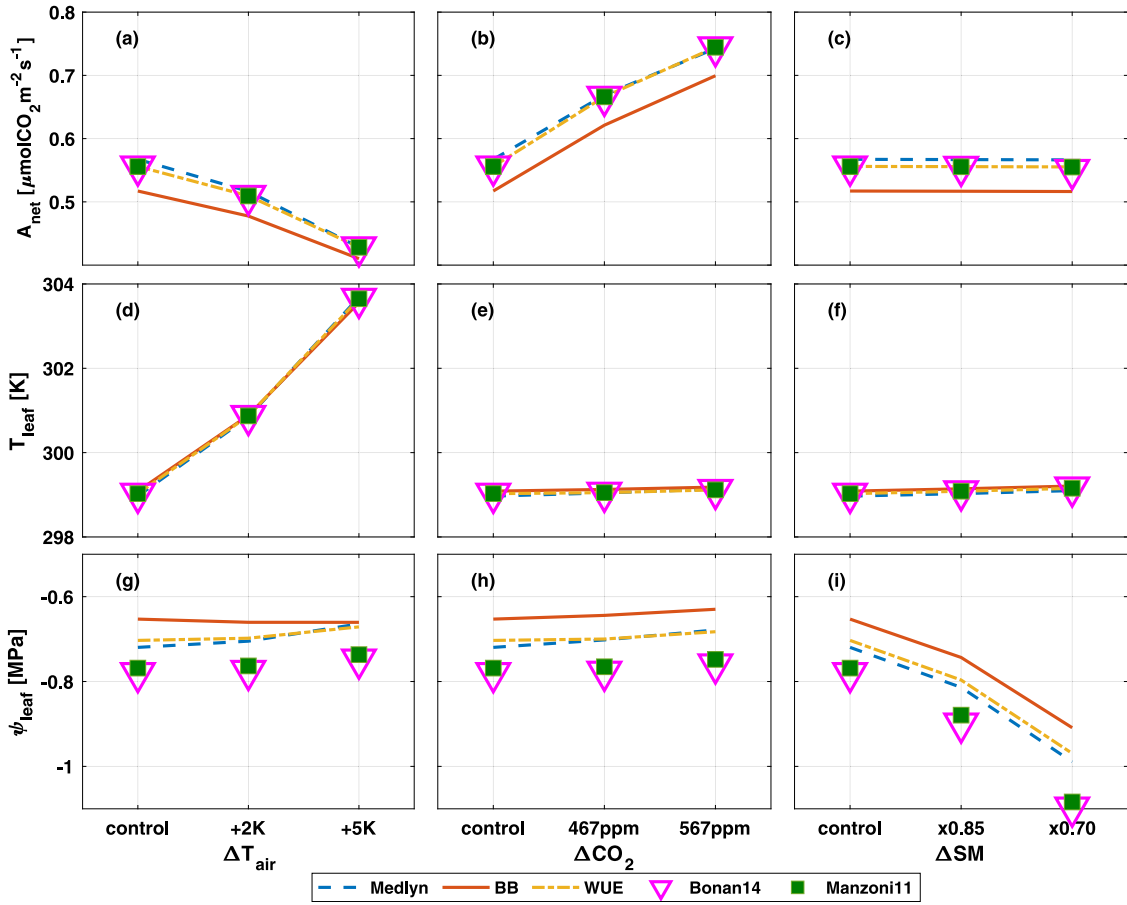


Fig. 8. Monthly-average net assimilation, A_{net} , due to changes in three environmental variables: (a) increase in temperature, ΔT_{air} , (b) increase in CO_2 , ΔCO_2 , and (c) decrease in soil moisture, ΔSM . Impact of the changes in the same three environmental variables on (d-f) monthly average leaf temperature, T_{leaf} , and (g-i) monthly average leaf water potential ψ_{leaf} . The results are presented for five different SCMs.

shown in Fig. 2a, while the resulting discretized linear system, $Ax = b$, with four GEs is shown in Fig. 2b. The entries of GE corresponding to air temperature and air water vapor within the x and b vector are shown in blue and magenta, respectively. The entries of GEs corresponding to the sunlit leaf (in light green) and shaded leaf (in dark green) includes unknown temperature values for the two vegetation canopies. The cyan colored sub-block of matrix A indicates non-zero values. The DMComposite allows the flexibility to grow the linear system when multiple vegetation canopies share a CAS.

An important detail associated with using DMComposite has a potentially large impact on performance. To avoid very high costs due to repeated memory allocation and copying during assembly, accurate preallocation – that is, giving PETSc an estimate of how many nonzero elements will be present in each row – for sparse matrices in PETSc is required. When a DM is being used to manage the computational domain, PETSc can usually use the information in the DM to automatically perform the appropriate preallocation, but when DMComposite is being used to glue many different subdomains together, there is not always a good way to determine, *a priori*, the nonzero pattern of the full coefficient matrix. It is important to accurately characterize this pattern since underestimating the number of nonzero elements can dramatically increase the time required for matrix assembly, and overestimating can easily exhaust available memory resources. We address this issue in ELM-MLCMv1 by using two passes through the routines that calculate the matrix elements: In the first pass, we perform no actual computation of matrix entries and “assemble” a “matrix” of type MATPREALLOCATOR, a dummy Mat type that is used simply to count nonzero elements. With the accurate counts made in this first phase, we then preallocate storage for the actual coefficient matrix,

which is assembled during the second pass. This simple optimization is absolutely necessary; without it, the matrix assembly is several orders of magnitude more time-consuming than it should be – so slow as to make using the code impracticable.

Besides providing data management tools such as DMComposite and implementations of many state-of-the-art nonlinear and linear solver algorithms, utilizing PETSc also provides a means to offload to GPU accelerator devices the most computationally expensive portion of the ELM-MLCMv1 model: the solution of the large linear system, which comprises at least half of the total execution time of the model when executing purely on CPUs. The design of PETSc makes this transparent to the application developer; changing from CPU to GPU execution simply requires specifying the appropriate runtime options. PETSc follows a delegation pattern approach to object-oriented design, in which each PETSc object is an instance of a class whose specific data structures and functionality are determined by choosing a delegated implementation type at runtime. GPUs can be used for linear algebra operations by simply choosing appropriate delegated types for the Vec and Mat objects that use the desired GPU back-end for executing vector and matrix operations, respectively. Because higher level components of PETSc (KSP linear solvers, SNES nonlinear solvers, TS time steppers, etc.) employ Vec and Mat operations for the bulk of their underlying computations, most of the computation in these solvers will occur on the GPU. The current back-ends supported by PETSc’s Vec and Mat classes include CUDA (NVIDIA’s native model), HIP (AMD’s “Heterogeneous Interface for Portability” model, which we use to target their ROCm GPU platform), and Kokkos (a performance portability framework that targets NVIDIA, AMD, and Intel GPUs as well as parallel CPU execution).




Fig. 9. The vertical profile of the diurnal cycle of changes in net CO₂ assimilation, ΔA_{net} , for the case (a) with an increase of 5 K in air temperature and (b) higher atmospheric CO₂ at 567 ppm. (c) The vertical profile of the diurnal cycle of changes in leaf temperature, ΔT_{leaf} , with an increase of 5 K air temperature, and (d) the vertical profile of the diurnal cycle of changes in leaf water potential, $\Delta \psi_{leaf}$, with reduction of soil moisture by a factor of 0.75. The changes A_{net} , T_{leaf} , and ψ_{leaf} are computed with respect to the control case.

4. Methodology

4.1. Experiment design for model verification

We verified the numerical implementation of ELM-MLCMv1 by comparing results against CLM-ml v1. Simulations were performed using atmospheric data from the Ameriflux US-University of Michigan Biological Station (US-UMB) [47], which is a temperate deciduous broadleaf forest site [48]. The simulations follow the model configuration of Bonan et al. [19] in which the 44 m vertical CAS is divided into 92 vertical layers with a uniform grid spacing of 0.5 m. The CAS is occupied by a single 22 m tall tree with a leaf area index (LAI) and stem area index (SAI) of 4.15 m² m⁻² and 0.90 m² m⁻², respectively (Fig. 3). Half-hourly atmospheric forcing data from July 2006 was used that included (1) downwelling direct and diffused solar radiation in the near-infrared and visible bands, (2) downwelling longwave radiation, (3) near-surface air temperature, (4) near-surface specific humidity, (5) near-surface atmospheric pressure, (6) near-surface wind speed, (7) near-surface atmospheric CO₂ concentration, and (8) near-surface atmospheric O₂ concentration. The time series of vertically-resolved soil moisture, soil temperature for the first soil layer, and surface albedo was prescribed as boundary conditions. Simulations were performed for the BB and Medlyn SCMs. The model used a 30-min time step to match the atmospheric forcing dataset. The ELM's default exponentially varying soil thicknesses were used for the vertical discretization of soil layers. A PHD model was not used in these simulations and instead, the ELM's default approach of computing hydraulic resistance based on an assumed root profile and a time-varying soil moisture was employed. The relative error for ELM-MLCMv1 with respect to CLM-ml v1 was computed at each model time step to verify the numerical implementation.

4.2. Experiment design for studying the impact of environmental variables on canopy processes

After establishing confidence in the numerical implementation of ELM-MLCMv1, a series of seven simulations (Table 3) were designed to investigate the impact of environmental variables on canopy processes and how these impacts vary across the different SCMs. First, a control simulation for each SCM was performed to calibrate model output to match observations of the monthly average diurnal cycle of sensible and latent heat fluxes using the same model setup as described in Section 4.1. The model performance was evaluated using the following metrics: (1) bias, (2) root-mean-square-error (RMSE), (3) correlation coefficient (R²), (4) the Kling-Gupta Efficiency (KGE) [49], and (5) the Nash-Sutcliffe efficiency (NSE) [50]. Next, additional simulations were performed for each SCM by changing the following three environmental variables: (1) air temperature, (2) atmospheric CO₂, and (3) soil moisture. Each environmental variable is changed twice while keeping the other variables unchanged. The simulated changes in net CO₂ assimilation, leaf temperature, and leaf water potential were studied.

4.3. Experiment design for studying the computational performance on heterogeneous computing architectures

The computational performance of the multi-layer CAS sub-model of ELM-MLCMv1 on CPU and GPU is studied by performing simulations for a range of problem sizes, using two platforms at the Oak Ridge Leadership Computing Facility (OLCF): The Summit Supercomputer, the seventh fastest supercomputer in the world as of November 2023 (<https://www.top500.org/lists/top500/2023/11/>), and the Crusher testbed

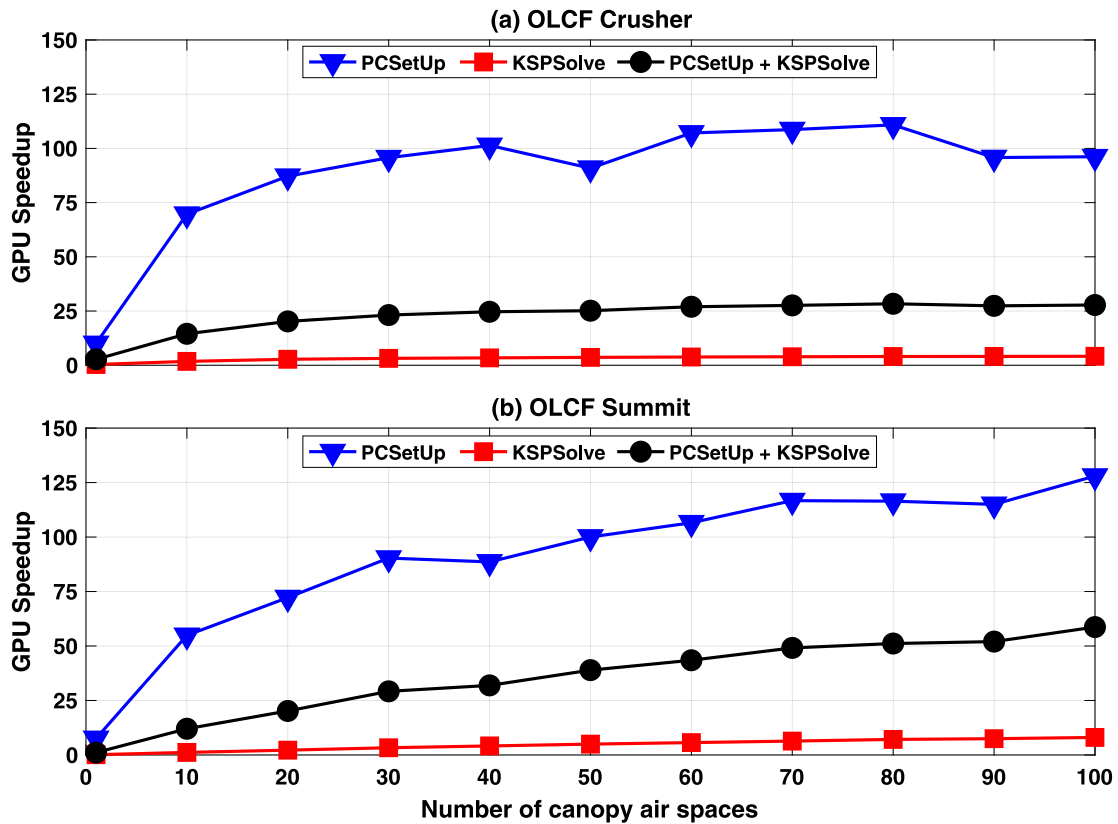


Fig. 10. Speedup of the PETSc solvers as a function of problem size (determined by number of CAS) when running on one GPU vs. one CPU on a node of the OLCF Summit (left) and Crusher (right machines), broken down into time spent in PCSetUp, KSPSolve, and the combination thereof.

system, which contains identical hardware and similar software as the DOE’s Frontier Exascale supercomputer, currently ranked as the fastest supercomputer in the world. These systems were chosen both because they represent the prior and current generation of leadership-class supercomputers, and because they rely on GPUs from different vendors and with different programming models. Each compute node of Summit has two 22-core IBM POWER9 CPUs and six NVIDIA Volta V100 GPUs, which uses NVIDIA’s CUDA programming model, while each compute node of Crusher has one 64-core AMD EPYC CPU and four AMD MI250X GPUs, which uses AMD’s ROCm programming model. We note that each AMD MI250X GPU consists of two Graphics Compute Dies (GCDs), each of which can be viewed by the programmer as a separate GPU; for the purposes of making consistent comparisons, we will use the term “GPU” in the rest of the paper to refer to one GPU on Summit or one GCD on Crusher.

We performed static-scaling experiments, in which simulations are run across a range of problem sizes, using a fixed level of parallelism (in our case, using a single CPU core or a single GPU). This analysis is in contrast to the more usual strong-scaling (in which the degree of parallelism is varied for a fixed problem size, and the marginal efficiency of each additional parallel resource is measured) or weak-scaling (in which the local problem size is fixed and the number of parallel resources is varied) analyses. Strong-scaling or weak-scaling analyses do not make sense for ELM-MLCMv1, which, although intended to be used in parallel ELM simulations, is essentially a serial code: an independent instance of the ELM-MLCMv1 model will be run by each process in a parallel ELM simulation, so the parallelism (from the viewpoint of ELM-MLCMv1) is fixed, using only a single CPU core or GPU (though the full number of GPU cores will be used by the PETSc solvers via one of the PETSc GPU back-ends). The relevant question to be answered is how a given compute platform will handle the range of problem sizes with which the ELM-MLCMv1 model will be run. By running over a wide range of problem sizes, we can answer questions

such as whether there is a minimum solution time below which the cost of the computation is dominated by GPU latency, or whether there is a problem size where suboptimal memory access becomes a dominant factor. Our aim is not to do a fair comparison between CPU and GPU platforms, which would require restructuring ELM-MLCMv1 to enable the use of thread-level parallelism within a process to use all available cores on a CPU: instead, our goal is simply to study how the work of a single-threaded process executing an ELM-MLCMv1 instance can be efficiently offloaded to an attached GPU device using the back-ends provided by PETSc.

We compared the model performance on a single CPU core is against one GPU, using the problem setup from Section 4.1, modified such that a CAS is shared by 1000 trees. Static scaling experiments were performed in which the number of CAS is varied between 1 to 100. It should be noted that there was no interaction between CASs. The numerical solution of the model uses PETSc’s KSP iterative linear solver component and the computational performance of the model in this work examines the two major steps of the solution process: (1) The PCSetUp event, in which the linear preconditioner data structures are set up and computed (in this case, and incomplete LU factorization with zero fill-in is computed) (2) The KSPSolve event, which occurs after PCSetUp and involves solution of the linear system by a preconditioned Krylov subspace iteration method (GMRES with a restart size of 30 here); this phase includes the application of the preconditioner (sparse triangular solve) but not the setup. The sum of the time spent in PCSetUp and KSPSolve comprises the total time spent inside PETSc solvers in an ELM-MLCMv1 simulation. We consider the time spent in PETSc solvers only for two reasons: First, we wish to evaluate how the model can be sped up by making use of GPU accelerators, and the PETSc solve is the most suitable part of the computation for GPU acceleration. Second, the time in the solvers represents fifty percent or more of the total time spent in the benchmark (when GPU acceleration

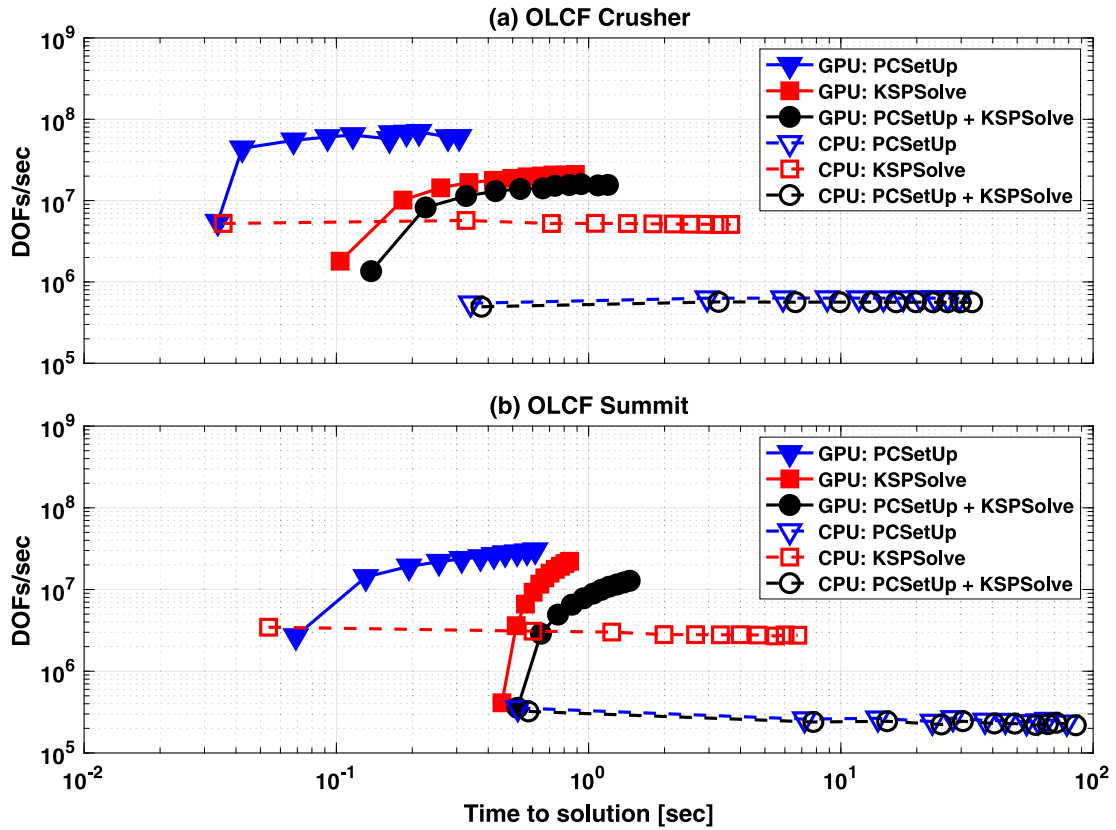


Fig. 11. A work-time spectrum view of the performance of PETSc solvers for our computational performance benchmark runs on the OLCF Summit and Crusher computers. The plot depicts the performance measurements for all of the different runs, though the quantity being controlled (the problem size, i.e., the number of CAS) is not plotted on any axis. This view has the advantage of allowing both latency and asymptotic throughput to be read directly from the plot.

is not used), and would represent an even larger fraction when ELM-MLCMv1 is deployed in an ELM simulation. In the benchmark code, the linear system is assembled and solved only once, but as part of an ELM simulation, much of the costly initial setup will be amortized across multiple solves; though the entries of the linear system matrix will be recomputed, the nonzero structure of the matrix will stay the same and the allocated data structures can be re-used.

The computational performance of the model was quantitatively evaluated using (1) speedup, and (2) work-time spectrum plots [51,52]. Speedup, a commonly employed metric in parallel computing, is computed here as the ratio of the execution time (as measured for a given PETSc event) for the CPU-only case to that of the GPU-enabled run. Speedup up gives a simple measure of how much faster the computation is when the GPU is employed. The work-time spectrum plot presents a less common but very informative view of model's computational performance by comparing the computation throughput (total degrees of freedom in the problem divided by the execution time) against the execution time. (Note that the quantity being controlled – the problem size – does not appear on any axis, but it increases from left to right for each curve.) The advantage of this view is that the latency and asymptotic throughput can be read directly from the plot: the former is represented simply by the smallest time that a given curve passes through, and the latter by the maximum throughput at the top of the curve. When a static scaling experiment runs over a wide enough range of problem sizes, performance curves on a work-time spectrum plot generally follow three regimes that proceed from left to right as: (1) regime with a positive slope where execution is close to the strong-scaling limit because the problem size is too small to allow full use of all of the parallel resources, (2) regime where slope levels off to a horizontal line that indicates a region of optimal scaling in which there is sufficient work for full utilization of all resources, and (3) regime of a negative slope where parallel efficiency drops due to memory subsystem effects or other causes.

5. Results

5.1. Model verification

The ELM-MLCMv1 model was able to accurately match the spatio-temporal results from the CLM-ml_v1 simulation. The two models show excellent agreement for vertically-averaged time series of air temperature, water vapor, and net radiation for shaded and sunlit leaves for Medlyn SCM (Fig. 4). The relative error in the time series of the vertical profile for the aforementioned variables for Medlyn SCM is $\mathcal{O}(10^{-9})$. The relative error between the two models for SCMs other than Bonan14 is similarly very small (Table 4). While the ELM-MLCMv1 is not validated against observations in this study, these results provide a high degree of confidence in the implementation of the ML model as it can accurately reproduce results of CLM-ml_v1 that has been validated against observational datasets [14].

5.2. Model validation

The calibrated ELM-MLCMv1 for each SCM was able to accurately simulate the monthly-averaged diurnal cycle of sensible and latent heat fluxes (Fig. 5) with low bias and RMSE, and high R^2 , KGE, and NSE (Table 5). The bias of the three SCMs that include PHD (i.e., WUE, Bonan14, and Manzoni11) is smaller than the bias of the two SCMs that do not include PHD (i.e., Medlyn and BB). However, all other evaluation metrics are comparable across the five SCMs. The model performance of all SCMs is lower at the hourly scale than compared to the monthly-averaged diurnal cycle (see metrics in parentheses in Table 5), though the KGE and NSE still remain larger than 0.69 and 0.77, respectively.

The simulated monthly-averaged vertical profiles of leaf water potential (ψ_{leaf}), air temperature (T_{air}), and air vapor pressure (q_{air}) show

vertical variation within the canopy (Fig. 6). The differences in the simulated variables across SCMs is small (as indicated by the cyan shading) and only limited to below the canopy height (as indicated by the dashed line). Given the strong agreement among SCMs, we only present temporally-varying, vertically-explicit profiles of variables simulated by ELM-MLCMv1 for the Medlyn SCM (Fig. 7). The top of the canopy has largest net assimilation, A_{net} (Fig. 7a) with corresponding lower intercellular CO_2 (Fig. 7b) and high stomatal conductance (Fig. 7c). The canopy buffers the variation of temperature by lowering the maximum leaf temperature of the lower canopy as compared to canopy top during mid-day and vice-versa during the night (Fig. 7d). The higher A_{net} at the top of the canopy results in higher water loss through transpiration, which results in lower ψ_{leaf} (Fig. 7e). The vertical profile of T_{air} shows similar spatio-temporal pattern as that of T_{leaf} with the upper canopy buffering temperature extrema in the lower canopy (Fig. 7f). The vertical gradient in q_{air} is only observed during the day with higher values at the top of canopy (Fig. 7g).

5.3. Impact of climatic factors on canopy processes

All five SCMs agree on the direction of the simulated changes (i.e., an increase or a decrease) in the three vertically-averaged canopy processes (i.e., A_{net} , T_{leaf} , and ψ_{leaf}) due to the changes in the three climatic factors (i.e., ΔT_{air} , ΔCO_2 , and ΔSM) (Fig. 8). The changes in a climate factors impacted canopy processes differently. For example, an increase in CO_2 resulted in higher A_{net} , but T_{leaf} and ψ_{leaf} remained unchanged (Fig. 8b,e,h). Only the changes in T_{air} impacts two canopy processes (i.e., A_{net} and T_{leaf}), and only A_{net} is impacted by two climatic factors (i.e., T_{air} and ΔCO_2). While the decrease in soil moisture lowers the ψ_{leaf} (Fig. 8i), the reduction in ψ_{leaf} is not large enough to limit the photosynthesis (Fig. 8c). With no changes in photosynthesis, the leaf temperature remains unchanged (Fig. 8f).

The spatio-temporal differences of the simulated changes in the canopy processes are studied next for the four cases (Fig. 8a,b,d,i) that showed a noticeable change due to the changes in climatic factors (Fig. 9). An increase in T_{air} (or a decrease in CO_2) leads to a decrease (or an increase) in A_{net} for all canopy layers (Fig. 9a,b). The change in A_{net} is largest at the top of the canopy and has a well-defined diurnal pattern with higher values during mid-day. The changes in T_{leaf} due to ΔT_{air} vary spatially and temporally (Fig. 9c). During nighttime, mid-canopy has lower ΔT_{leaf} compared to the top and the bottom of the canopy, and vice-versa during daytime. The ψ_{leaf} decreased due to ΔSM throughout the canopy with the largest changes at the top of the canopy and during late afternoon (Fig. 9d).

5.4. Computational performance

The speedup plots for both the Summit and Crusher machines follow the expected trends, with speedup generally increasing as the problem size grows to be large enough to allow full utilization of the large number of GPU compute elements present (Fig. 10). The speedup is most pronounced for the PCSetUp event, which involves a large amount of work that can saturate the GPU, whereas the KSPSolve event – which involves comparatively much less work and cannot keep the GPU as busy – displays much more modest speedup. On Summit, the speedup increases consistently with problem size, whereas on Crusher the speedup levels off somewhat before the maximum problem size is reached. At larger problem sizes, the speedups on Summit are consistently higher than those on Crusher, but the execution on Crusher is actually always faster (Fig. 11) as the speedups in each plot are measured relative to the performance of that machine's CPU.

The work-time spectrum plots for the CPU-only runs are close to horizontal (dashed lines in Fig. 11), indicating that, for this range of problem sizes, the CPU resources are saturated and no memory effects come into play. For the GPU-enabled runs (solid lines in Fig. 11), we see that the PCSetUp phase reaches the optimal scaling regime on

Crusher, and is perhaps approaching it on Summit. The KSPSolve phase contains less work and is expected to require a larger problem size to reach the optimal scaling regime. We observe that on Crusher KSPSolve appears to be approaching the optimal scaling phase at the largest problem sizes, whereas on Summit the KSPSolve seems to still be near the strong-scaling limit. On both machines, the GPU performance beats the CPU performance in all cases except for the KSPSolve with the smallest problem size (a single CAS); for this case, the amount of work done in the KSPSolve is too small to mask high kernel launch latencies for the GPUs, but because this does not hold for the PCSetUp, the total time spent in the PETSc solvers is still lower for the GPU. As expected, because the CPUs and GPUs in Summit are several years older than in Crusher, the Crusher GPU and GPU performance consistently beats that of its counterparts on Summit.

6. Caveats and future work

This study is a first step to including the representation of vertical canopy structure within ELM and we now briefly discuss possible future developments in the new model required for implementation in the ELM big-leaf scheme. First, the prescribed timeseries of soil moisture, surface temperature, and surface albedo as a boundary condition should instead use ELM's prognostic soil hydrology and soil temperature models. Second, the evaporation of canopy intercepted water needs to be accounted for in the energy balance within the CAS. Third, the radiation sub-model could be extended to simulate solar-induced induced chlorophyll fluorescence, which has emerged as a proxy for terrestrial photosynthesis and can be observed from satellites [53]. Fourth, a more realistic vertical structure of sunlit and shaded leaves should be included in which leaves at each canopy level are not directly connected to the soil via individual xylems but are connected to a single xylem [54]. Fifth, multi-year and multi-site simulations need to be performed to evaluate the performance of the model against observations. Sixth, a computational performance analysis of the model should be performed for a global ELM simulation. Seventh, the sensitivity of the model skill and computational cost with respect to the number of vertical layers need to be investigated in the future. These additional model developments will significantly improve ELM's capability to resolve micro-climate due to vertical canopy structure and enable support for upcoming heterogeneous high-performance architectures.

7. Summary

In this study, we developed a standalone ELM-MLCMv1 that resolves the micro-climate created by the vertical structure of the vegetation canopies. The new model includes sub-models for shortwave and long-wave radiation, stomatal conductance, turbulence scheme for modeling flow within the vegetation canopy, and transport of heat and water vapor within CAS. ELM-MLCMv1 uses PETSc to provide a numerically robust solution of discretized equations and includes support for heterogeneous computing architectures. The numerical implementation of ELM-MLCMv1 was verified by comparing results against CLM-ml_v1, an existing MLC model, for a month-long simulation using data from the Ameriflux US-UMB site.

We found that simulations using SCMs of various degrees of complexity are each able to accurately simulate observed sensible and latent heat fluxes upon model calibration. A set of idealized climate change simulations showed that all SCMs predict a consistent change in the canopy processes due to changes in environmental factors (i.e., air temperature, atmospheric CO_2 , and soil moisture). The differences in the canopy processes including net assimilation, leaf temperature, and leaf water potential vary vertically within the canopy in response to changes in environmental factors. While net assimilation was impacted by the changes in both air temperature and atmospheric CO_2 ; however, leaf temperature and leaf water potential were only impacted by changes

in air temperature and soil moisture, respectively. The use of PETSc enabled the same ELM-MLCMv1 source code to run on CPUs and GPUs without requiring any code modifications, as well as includes support of GPUs from multiple vendors (i.e., NVIDIA and AMD). Static scaling experiments for an idealized problem showed ELM-MLCMv1 achieved a speedup of 25–50 times on a GPU relative to a CPU. This work provides the first necessary model development to include the representation of vertical canopies in ELM and additionally includes the support for heterogeneous computing architectures.

Software and data availability

The ELM-MLCMv1 code is available under the BSD 3-clause license at <https://github.com/MPP-LSM/MPP> and the documentation is available at <https://mpp-lsm.github.io/mpp-doc/>. The version of the model used in this study and the corresponding documentation are archived at <https://zenodo.org/record/7809207> and <https://doi.org/10.5281/zenodo.8213108>, respectively. Furthermore, the scripts and the required input data to perform simulations presented in this study are available at <https://github.com/MPP-LSM/mlcm-simulation>.

CRedit authorship contribution statement

Gautam Bisht: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Conceptualization. **William J. Riley:** Writing – review & editing, Conceptualization. **Richard Tran Mills:** Writing – review & editing, Software.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The software and data are publicly available and detailed information in provided in Section “Software and Data Availability”.

Acknowledgments

This research was supported as part of the Energy Exascale Earth System Model (E3SM) project, funded by the U.S. Department of Energy, Office of Science, Office of Biological and Environmental Research Earth Systems Model Development Program area of Earth and Environmental System Modeling. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

References

- [1] I.N. Harman, J.J. Finnigan, A simple unified theory for flow in the canopy and roughness sublayer, *Bound.-Layer Meteorol.* 123 (2) (2007) 339–363.
- [2] T.M. Scanlon, J.D. Albertson, Turbulent transport of carbon dioxide and water vapor within a vegetation canopy during unstable conditions: Identification of episodes using wavelet analysis, *J. Geophys. Res.: Atmos.* 106 (D7) (2001) 7251–7262, <http://dx.doi.org/10.1029/2000JD900662>, arXiv:<https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2000JD900662>. URL: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2000JD900662>.
- [3] P. De Frenne, F. Zellweger, F. Rodríguez-Sánchez, B.R. Scheffers, K. Hylander, M. Luoto, M. Vellend, K. Verheyen, J. Lenoir, Global buffering of temperatures under forest canopies, *Nat. Ecol. Evol.* 3 (5) (2019) 744–749.
- [4] J.S. Sperry, F. Adler, G. Campbell, J. Comstock, Limitation of plant water use by rhizosphere and xylem conductance: results from a model, *Plant Cell Environ.* 21 (4) (1998) 347–359.

- [5] G. Bohrer, H. Mourad, T.A. Laursen, D. Drewry, R. Avissar, D. Poggi, R. Oren, G.G. Katul, Finite element tree crown hydrodynamics model (FETCH) using porous media flow within branching elements: A new representation of tree hydrodynamics, *Water Resour. Res.* 41 (11) (2005).
- [6] R. Grant, T. Black, D. Gaumont-Guay, N. Klujin, A. Barr, K. Morgenstern, Z. Nestic, Net ecosystem productivity of boreal aspen forests under drought and climate change: Mathematical modelling with ecosys, *Agric. Forest Meteorol.* 140 (1–4) (2006) 152–170.
- [7] X. Xu, D. Medvigy, J.S. Powers, J.M. Becknell, K. Guan, Diversity in plant hydraulic traits explains seasonal and inter-annual variations of vegetation dynamics in seasonally dry tropical forests, *New Phytol.* 212 (1) (2016) 80–95.
- [8] E. Agee, L. He, G. Bisht, V. Couvreur, P. Shahbaz, F. Meunier, C.M. Gough, A.M. Matheny, G. Bohrer, V. Ivanov, Root lateral interactions drive water uptake patterns under water limitation, *Adv. Water Resour.* 151 (2021) 103896.
- [9] M. Silva, A.M. Matheny, V. Pauwels, D. Triadis, J.E. Missik, G. Bohrer, E. Daly, Tree hydrodynamic modelling of the soil–plant–atmosphere continuum using FETCH3, *Geosci. Model Dev.* 15 (6) (2022) 2619–2634.
- [10] D. Kennedy, S. Swenson, K.W. Oleson, D.M. Lawrence, R. Fisher, A.C. Lola da Costa, P. Gentine, Implementing plant hydraulics in the community land model, version 5, *J. Adv. Modelling Earth Syst.* 11 (2) (2019) 485–513.
- [11] Y. Fang, L.R. Leung, B.T. Wolfe, M. Detto, R.G. Knox, N.G. McDowell, C. Grossiord, C. Xu, B.O. Christoffersen, P. Gentine, et al., Disentangling the effects of vapor pressure deficit and soil water availability on canopy conductance in a seasonal tropical forest during the 2015 El Niño drought, *J. Geophys. Res.: Atmos.* 126 (10) (2021) e2021JD035004.
- [12] G. Bonan, M. Williams, R. Fisher, K. Oleson, Modeling stomatal conductance in the earth system: linking leaf water-use efficiency and water transport along the soil–plant–atmosphere continuum, *Geosci. Model Dev.* 7 (5) (2014) 2193–2222.
- [13] J. Ryder, J. Polcher, P. Peylin, C. Otlé, Y. Chen, E. van Gorsel, V. Haverd, M.J. McGrath, K. Naudts, J. Otto, et al., A multi-layer land surface energy budget model for implicit coupling with global atmospheric simulations, *Geosci. Model Dev.* 9 (1) (2016) 223–245.
- [14] G.B. Bonan, E.G. Patton, I.N. Harman, K.W. Oleson, J.J. Finnigan, Y. Lu, E.A. Burakowski, Modeling canopy-induced turbulence in the earth system: a unified parameterization of turbulent exchange within plant canopies and the roughness sublayer (CLM-ml v0), *Geosci. Model Dev.* 11 (4) (2018) 1467–1496.
- [15] M. Longo, R.G. Knox, D.M. Medvigy, N.M. Levine, M.C. Dietze, Y. Kim, A.L.S. Swann, K. Zhang, C.R. Rollinson, R.L. Bras, S.C. Wofsy, P.R. Moorcroft, The biophysics, ecology, and biogeochemistry of functionally diverse, vertically and horizontally heterogeneous ecosystems: the ecosystem demography model, version 2.2 – part 1: Model description, *Geosci. Model Dev.* 12 (10) (2019) 4309–4346, <http://dx.doi.org/10.5194/gmd-12-4309-2019>, URL: <https://gmd.copernicus.org/articles/12/4309/2019/>.
- [16] W. Massman, J. Weil, An analytical one-dimensional second-order closure model of turbulence statistics and the Lagrangian time scale within and above plant canopies of arbitrary structure, *Bound.-Layer Meteorol.* 91 (1) (1999) 81–107.
- [17] Y. Chen, J. Ryder, V. Bastrikov, M.J. McGrath, K. Naudts, J. Otto, C. Otlé, P. Peylin, J. Polcher, A. Valade, et al., Evaluating the performance of land surface model ORCHIDEE-CAN v1. 0 on water and energy flux estimation with a single-and multi-layer energy budget scheme, *Geosci. Model Dev.* 9 (9) (2016) 2951–2972.
- [18] I.N. Harman, J.J. Finnigan, Scalar concentration profiles in the canopy and roughness sublayer, *Bound.-Layer Meteorol.* 129 (3) (2008) 323–351.
- [19] G.B. Bonan, E.G. Patton, J.J. Finnigan, D.D. Baldocchi, I.N. Harman, Moving beyond the incorrect but useful paradigm: reevaluating big-leaf and multilayer plant canopies to model biosphere-atmosphere fluxes—a review, *Agric. Forest Meteorol.* 306 (2021) 108435.
- [20] J.T. Ball, I.E. Woodrow, J.A. Berry, A model predicting stomatal conductance and its contribution to the control of photosynthesis under different environmental conditions, in: *Progress in Photosynthesis Research*, Springer, 1987, pp. 221–224.
- [21] R. Leuning, A critical appraisal of a combined stomatal-photosynthesis model for C3 plants, *Plant Cell Environ.* 18 (4) (1995) 339–355.
- [22] I. Cowan, F. GD, Stomatal function in relation to leaf metabolism and environment, 1977.
- [23] G. Katul, S. Manzoni, S. Palmroth, R. Oren, A stomatal optimization theory to describe the effects of atmospheric CO2 on leaf photosynthesis and transpiration, *Ann. Botany* 105 (3) (2010) 431–442.
- [24] T.N. Buckley, L. Sack, G.D. Farquhar, Optimal plant water economy, *Plant Cell Environ.* 40 (6) (2017) 881–896.
- [25] S. Manzoni, G. Vico, G. Katul, P.A. Fay, W. Polley, S. Palmroth, A. Porporato, Optimizing stomatal conductance for maximum carbon gain under water stress: a meta-analysis across plant functional types and climates, *Funct. Ecol.* 25 (3) (2011) 456–467.
- [26] B.O. Christoffersen, M. Gloor, S. Fauset, N.M. Fyllas, D.R. Galbraith, T.R. Baker, B. Kruijff, L. Rowland, R.A. Fisher, O.J. Binks, et al., Linking hydraulic traits to tropical forest function in a size-structured and trait-driven model (TFS v. 1-hydro), *Geosci. Model Dev.* 9 (11) (2016) 4227–4255.
- [27] G. Mirfenderesgi, G. Bohrer, A.M. Matheny, S. Faticchi, R.P. de Moraes Frasson, K.V. Schäfer, Tree level hydrodynamic approach for resolving aboveground water storage and stomatal conductance and modeling the effects of tree hydraulic strategy, *J. Geophys. Res.: Biogeosci.* 121 (7) (2016) 1792–1813.

- [28] Y. Wang, J.S. Sperry, W.R. Anderegg, M.D. Venturas, A.T. Trugman, A theoretical and empirical assessment of stomatal optimization modeling, *New Phytol.* 227 (2) (2020) 311–325.
- [29] J.S. Sperry, M.D. Venturas, W.R. Anderegg, M. Mencuccini, D.S. Mackay, Y. Wang, D.M. Love, Predicting stomatal responses to the environment from the optimization of photosynthetic gain and hydraulic cost, *Plant Cell Environ.* 40 (6) (2017) 816–830.
- [30] T. Hölttä, A. Lintunen, T. Chan, A. Mäkelä, E. Nikinmaa, A steady-state stomatal model of balanced leaf gas exchange, hydraulics and maximal source–sink flux, *Tree Physiol.* 37 (7) (2017) 851–868.
- [31] R. Dewar, A. Mauranen, A. Mäkelä, T. Hölttä, B. Medlyn, T. Vesala, New insights into the covariation of stomatal, mesophyll and hydraulic conductances from optimization models incorporating nonstomatal limitations to photosynthesis, *New Phytol.* 217 (2) (2018) 571–585.
- [32] C.B. Eller, L. Rowland, M. Mencuccini, T. Rosas, K. Williams, A. Harper, B.E. Medlyn, Y. Wagner, T. Klein, G.S. Teodoro, et al., Stomatal optimization based on xylem hydraulics (SOX) improves land surface model simulation of vegetation responses to climate, *New Phytol.* 226 (6) (2020) 1622–1637.
- [33] Y. Wang, P. Köhler, L. He, R. Doughty, R.K. Braghieri, J.D. Wood, C. Frankenberg, Testing stomatal models at stand level in deciduous angiosperm and evergreen gymnosperm forests using clima land (v0.1), *Geosci. Model Dev. Discuss.* 2021 (2021) 1–35, <http://dx.doi.org/10.5194/gmd-2021-154>, URL: <https://gmd.copernicus.org/preprints/gmd-2021-154/>.
- [34] H.C. Edwards, C.R. Trott, D. Sunderland, Kokkos: Enabling manycore performance portability through polymorphic memory access patterns, *J. Parallel Distrib. Comput.* 74 (12) (2014) 3202–3216.
- [35] D.A. Beckingsale, J. Burmark, R. Hornung, H. Jones, W. Killian, A.J. Kunen, O. Pearce, P. Robinson, B.S. Ryuujin, T.R. Scogland, RAJA: Portable performance for large-scale scientific applications, in: 2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in Hpc, P3hpc, IEEE, 2019, pp. 71–81.
- [36] NVIDIA, P. Vingelmann, F.H. Fitzek, CUDA, release: 10.2.89, 2020, URL: <https://developer.nvidia.com/cuda-toolkit>.
- [37] A. ROCm, A new era in open GPU computing, 2019.
- [38] J.E. Stone, D. Gohara, G. Shi, OpenCL: A parallel programming standard for heterogeneous computing systems, *Comput. Sci. Eng.* 12 (3) (2010) 66.
- [39] S. Balay, S. Abhyankar, M.F. Adams, S. Benson, J. Brown, P. Brune, K. Buschelman, E. Constantinescu, L. Dalcin, A. Dener, V. Eijkhout, W.D. Gropp, V. Hapla, T. Isaac, P. Jolivet, D. Karpeev, D. Kaushik, M.G. Knepley, F. Kong, S. Kruger, D.A. May, L.C. McInnes, R.T. Mills, L. Mitchell, T. Munson, J.E. Roman, K. Rupp, P. Sanan, J. Sarich, B.F. Smith, S. Zampini, H. Zhang, H. Zhang, J. Zhang, PETSc/TAO Users Manual, Technical Report ANL-21/39 - Revision 3.16, Argonne National Laboratory, 2021.
- [40] R.T. Mills, M.F. Adams, S. Balay, J. Brown, A. Dener, M. Knepley, S.E. Kruger, H. Morgan, T. Munson, K. Rupp, et al., Toward performance-portable PETSc for GPU-based exascale systems, *Parallel Comput.* 108 (2021) 102831.
- [41] J. Brown, M.G. Knepley, D.A. May, L.C. McInnes, B.F. Smith, Composable linear solvers for multiphysics, in: 11th International Symposium on Parallel and Distributed Computing, Munich, Germany, 2012.
- [42] G. Bisht, W.J. Riley, Development and verification of a numerical library for solving global terrestrial multiphysics problems, *J. Adv. Modelling Earth Syst.* 11 (6) (2019) 1516–1542.
- [43] G. Bonan, *Climate Change and Terrestrial Ecosystem Modeling*, Cambridge University Press, 2019.
- [44] G. Bisht, Documentation for the multi-layer canopy model (MLCMv1), 2023, <https://doi.org/10.5281/zenodo.8213108>.
- [45] J. Norman, in: B.J. Barfield, J.F. Gerber (Eds.), *Modeling the Complete Crop Canopy*, Modification of the Aerial Environment of Plants, in: ASAE Monogr, 1979, pp. 249–277.
- [46] B.E. Medlyn, R.A. Duursma, D. Eamus, D.S. Ellsworth, I.C. Prentice, C.V. Barton, K.Y. Crous, P. De Angelis, M. Freeman, L. Wingate, Reconciling the optimal and empirical approaches to modelling stomatal conductance, *Glob. Change Biol.* 17 (6) (2011) 2134–2144.
- [47] C. Gough, G. Bohrer, P. Curtis, AmeriFlux BASE US-UMB univ. of mich. Biological station, ver. 18-5, AmeriFlux AMP, (dataset), 2022, <https://doi.org/10.17190/AMF/1246107>.
- [48] H. Schmid, H.-B. Su, C. Vogel, P. Curtis, Ecosystem-atmosphere exchange of carbon dioxide over a mixed hardwood forest in northern lower Michigan, *J. Geophys. Res.: Atmos.* 108 (D14) (2003).
- [49] H.V. Gupta, H. Kling, K.K. Yilmaz, G.F. Martinez, Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling, *J. Hydrol.* 377 (1–2) (2009) 80–91.
- [50] J.E. Nash, J.V. Sutcliffe, River flow forecasting through conceptual models part I—A discussion of principles, *J. Hydrol.* 10 (3) (1970) 282–290.
- [51] J. Chang, K. Nakshatrala, M.G. Knepley, L. Johnsson, A performance spectrum for parallel computational frameworks that solve PDEs, *Concurr. Comput.: Pract. Exper.* 30 (11) (2018) e4401.
- [52] J. Chang, M.S. Fabien, M.G. Knepley, R.T. Mills, Comparative study of finite element methods using the time-accuracy-size (TAS) spectrum analysis, *SIAM J. Sci. Comput.* 40 (6) (2018) C779–C802.
- [53] R. Li, D. Lombardozi, M. Shi, C. Frankenberg, N.C. Parazoo, P. Köhler, K. Yi, K. Guan, X. Yang, Representation of leaf-to-canopy radiative transfer processes improves simulation of far-red solar-induced chlorophyll fluorescence in the community land model version 5, *J. Adv. Modelling Earth Syst.* 14 (3) (2022) e2021MS002747.
- [54] G. Mirfenderesgi, A.M. Matheny, G. Bohrer, Hydrodynamic trait coordination and cost–benefit trade-offs throughout the isohydric–aniso-hydric continuum in trees, *Ecohydrology* 12 (1) (2019) e2041.



Gautam Bisht received his Ph.D. in Environmental Engineering from the Massachusetts Institute of Technology, Cambridge. He is an Earth Scientist in the Atmospheric Sciences and Global Change Division at Pacific Northwest National Laboratory. His research interest includes developing computationally robust models for global terrestrial biophysical processes that efficiently use high-performance computing.



William J. Riley received his Ph.D. in Environmental Engineering from the University of California, Berkeley. He is a Senior Scientist in the Climate and Ecosystem Sciences Division at Lawrence Berkeley National Laboratory. His research interest includes modeling terrestrial ecosystems and their interactions with climate and climate change.



Richard Tran Mills received his Ph.D. in Computer Science from The College of William & Mary, Virginia. He is a Computational Scientist in the Mathematics and Computer Science Division at Argonne National Laboratory. His research interest includes developing and applying numerical methods and software to enable the solution of very computationally challenging problems that arise in the natural sciences.