

1 **Parallel SnowModel (v1.0): a parallel implementation of a** 2 **Distributed Snow-Evolution Modeling System (SnowModel)**

3 Ross Mower^{1,2}, Ethan D. Gutmann¹, Glen E. Liston³, Jessica Lundquist², Soren Rasmussen¹

4 ¹The NSF National Center for Atmospheric Research, Boulder, Colorado, USA

5 ²Department of Civil and Environmental Engineering, University of Washington, Seattle, Washington, USA

6 ³Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins, Colorado, USA

7 *Correspondence to:* Ross Mower (rossamower@ucar.edu)

8 **Abstract.** SnowModel, a spatially distributed, snow-evolution modeling system, was parallelized using Coarray Fortran for
9 high-performance computing architectures to allow high-resolution (1 m to 100s of meters) simulations over large, regional
10 to continental scale, domains. In the parallel algorithm, the model domain was split into smaller rectangular sub-domains that
11 are distributed over multiple processor cores using one-dimensional decomposition. All the memory allocations from the
12 original code were reduced to the size of the local sub-domains, allowing each core to perform fewer computations and
13 requiring less memory for each process. Most of the subroutines in SnowModel were simple to parallelize; however, there
14 were certain physical processes, including blowing snow redistribution and components within the solar radiation and wind
15 models, that required non-trivial parallelization using halo-exchange patterns. To validate the parallel algorithm and assess
16 parallel scaling characteristics, high-resolution (100 m grid) simulations were performed over several western United States
17 domains and over the contiguous United States (CONUS) for a year. The CONUS scaling experiment had approximately
18 70% parallel efficiency; runtime decreased by a factor of 1.9 running on 1800 cores relative to 648 cores (the minimum
19 number of cores that could be used to run such a large domain because of memory and time limitations). CONUS 100 m
20 simulations were performed for 21 years (2000 – 2021) using 46,238 and 28,260 grid cells in the x and y dimensions,
21 respectively. Each year was simulated using 1800 cores and took approximately 5 hours to run.

22 **1 Introduction**

23 The cryosphere (snow and ice) is an essential component of Arctic, mountain, and downstream ecosystems, Earth's surface
24 energy balance, and freshwater resource storage (Huss et al., 2017). Globally, half the world's population depends on
25 snowmelt (Beniston, 2003). In snow-dominated regions like the Western United States, snowmelt contributes to
26 approximately 70% of the total annual water supply (Foster et al., 2011). In these regions, late-season streamflow is
27 dependent on the deepest snow drifts and therefore longest-lasting snow (Pflug and Lundquist, 2020). Since modeling snow-
28 fed streamflow accurately is largely dependent on our ability to predict snow quantities and the associated spatial and
29 temporal variability (Clark and Hay, 2004), high-temporal and -spatial resolution snow datasets are important for predicting
30 flood hazards and managing freshwater resources (Immerzeel et al., 2020).

31 The spatial and temporal seasonal snow characteristics also have significant implications outside of water resources.
32 Changes in fractional snow-covered area affect albedo and thus atmospheric dynamics (Liston, 2004; Liston and Hall, 1995).
33 Avalanches pose safety hazards to both transportation and recreational activities in mountainous terrain; the prediction of
34 which requires high-resolution (meters) snow datasets (Morin et al., 2020; Richter et al., 2021). Additionally, the timing and
35 duration of snow-covered landscapes strongly influence how species adapt, migrate, and survive (Boelman et al., 2019;
36 Liston et al., 2016; Mahoney et al., 2018).

37 To date, the primary modes for estimating snow properties and storage have come from observation networks, satellite-based
38 sensors, and physically derived snow algorithms in land surface models (LSMs). However, despite the importance of
39 regional, continental, and global snow, estimates of snow properties over these scales remain uncertain, especially in alpine
40 regions where wind, snow, and topography interact (Boelman et al., 2019; Dozier et al., 2016; Mudryk et al., 2015).
41 Observation datasets used for spatial interpolation of snow properties and forcing datasets used in LSMs are often too sparse
42 in mountainous terrain to accurately resolve snow spatial heterogeneities (Dozier et al., 2016; Renwick, 2014). Additionally,
43 remotely sensed products have shown deficiencies in measuring snowfall rate (Skofronick-Jackson et al., 2013), snow-water
44 equivalent (SWE), and snow depth (Nolin, 2010), especially in mountainous terrain where conditions of deep snow, wet
45 snow, and/or dense vegetation may be present (Lettenmaier et al., 2015; Takala et al., 2011; Vuyovich et al., 2014).
46 However, LSMs using high-resolution inputs, including forcing datasets from regional climate models (RCMs), have
47 demonstrated realistic spatial distributions of snow properties (Wrzesien et al., 2018).

48 Many physical snow models have been developed either in stand-alone algorithms or larger LSMs with varying degrees of
49 complexity based on their application. The more advanced algorithms attempt to accurately model snow properties at high
50 resolution especially in regions where snow interacts with topography, vegetation, and/or wind. Wind-induced snow
51 transport is one such complexity of snow that represents an important interaction between the cryosphere and atmosphere. It
52 occurs in regions permanently or temporarily covered by snow, influences snow properties (e.g. heterogeneity, sublimation,
53 avalanches, and melt timing), and has been shown to improve simulated snowpack distribution (Bernhardt et al., 2012;
54 Freudiger et al., 2017; Keenan et al., 2023; Quéno et al., 2023). Models that have incorporated wind-induced physics
55 generally require components to both develop the snow mass balance and incorporate atmospheric inputs of the wind field.
56 Additionally, these models typically require high resolution grids (1 to 100 m) as the redistribution components of the model
57 become negligible at larger spatial discretizations (Liston et al., 2007). However, there often exists a trade-off between the
58 accuracy of simulating wind-induced snow transport and the computational requirements for downscaling and developing
59 the wind fields over the gridded domain (Reynolds et al., 2021; Vionnet et al., 2014). Therefore, simplifying assumptions of
60 uniform wind direction has been applied in models like Distributed Blowing Snow Model (DBSM) (Essery et al., 1999; Fang
61 and Pomeroy, 2009). More advanced models have utilized advection-diffusion equations, like Alpine3D (Lehning et al.,
62 2006) or spatial distributed formulations like SnowTran-3D (Liston and Sturm, 1998). Finite volume methods for more
63 efficiently discretizing wind fields have been applied to models such as DBSM (Marsh et al., 2020). The most complex
64 models consider nonsteady turbulence which utilize three-dimensional wind fields from atmospheric models to simulate

65 blowing snow transport and sublimation; for example, SURFEX in Meso-NH/Crocus (Vionnet et al., 2014; Vionnet et al.,
66 2017), wind fields from the atmospheric model ARPS (Xue et al., 2000) being incorporated into Alpine3D (Mott and
67 Lehning, 2010; Mott et al., 2010; Lehning et al., 2008), and SnowDrift3D (Prokop and Schneiderbauer, 2011). Incorporating
68 wind-induced physics into snow models is computationally expensive; thus, parallelizing the serial algorithms would likely
69 be beneficial to many models.

70 For several decades, a distributed snow-evolution modeling system (SnowModel) has been developed, enhanced, and tested
71 to accurately simulate snow properties across a wide range of landscapes, climates, and conditions (Liston and Elder, 2006a;
72 Liston et al., 2020). To date, SnowModel has been used in over 200 refereed journal publications; a short listing of these is
73 provided by Liston et al. (2020). Physically derived snow algorithms, as used in SnowModel, that model the energy balance,
74 multilayer snow physics, and lateral snow transport are computationally expensive. In these models, the required
75 computational power increases with the number of grid cells covering the simulation domain. Finer grid resolutions usually
76 imply more grid cells and higher accuracy resulting from improved representation of process physics at higher resolutions.
77 The original serial SnowModel code was written in Fortran 77 and could not be executed in parallel using multiple processor
78 cores. As a result, SnowModel's spatial and temporal simulation domains (number of grid cells and time steps) were
79 previously limited by the speed of one core and the memory available on the single computer. Note that a "processor" refers
80 to a single central processing unit (CPU) and typically consists of multiple cores, each core can run one or more processes in
81 parallel.

82 Recent advancements in multiprocessor computer technologies and architectures have allowed for increased performance in
83 simulating complex natural systems at high resolutions. Parallel computing has been used on many LSMs to reduce compute
84 time and allow for higher accuracy results from finer grid simulations (Hamman et al., 2018; Miller et al., 2014). Our goal
85 was to develop a parallel version of SnowModel (Parallel SnowModel) using Coarray Fortran (CAF) syntax without making
86 significant changes to the original SnowModel code physics or structure. CAF is a Partitioned Global Address Space
87 (PGAS) programming model and has been used to run atmospheric models on 100,000 cores (Rouson et al., 2017).

88 In parallelizing numerical models, a common strategy is to decompose the domain into smaller sub-domains that get
89 distributed across multiple processes (Dennis, 2007; Hamman et al., 2018). For rectangular gridded domains (like
90 SnowModel), this preserves the original structure of the spatial loops and utilizes direct referencing of neighboring grids
91 (Perezhogin et al., 2021). The parallelization of many LSMs involve "embarrassingly parallel" problems requiring minimal
92 to no processor communication (Parhami, 1995); in this case, adjacent grid cells do not communicate with each other (an
93 example of this would be where each grid cell represents a point, or one-dimension, snowpack model that is not influenced
94 by nearby grid cells).

95 While much of the SnowModel's logic can be considered "embarrassingly parallel", SnowModel also contains "non-trivial"
96 algorithms within the solar radiation, wind, and snow redistribution models. Calculations within these algorithms often
97 require information from neighboring grid cells, either for spatial derivative calculations or for horizontal fluxes of mass
98 (e.g., saltating or turbulent-suspended snow) across the domain. Therefore, non-trivial parallelization requires implementing

99 algorithm changes that allow computer processes to communicate and exchange data. The novelty of the work presented
100 here includes 1) the presentation of Parallel SnowModel, high-resolution (100 m) distributed snow datasets over CONUS,
101 and an analysis of the performance of the parallel algorithm; 2) demonstrating how a simplified parallelization approach
102 using CAF and one-dimensional decomposition can be implemented in geoscientific algorithms to scale over large domains;
103 and 3) demonstrating an approach for non-trivial parallelization algorithms that involve spatial derivatives and fluxes using
104 halo-exchange techniques.

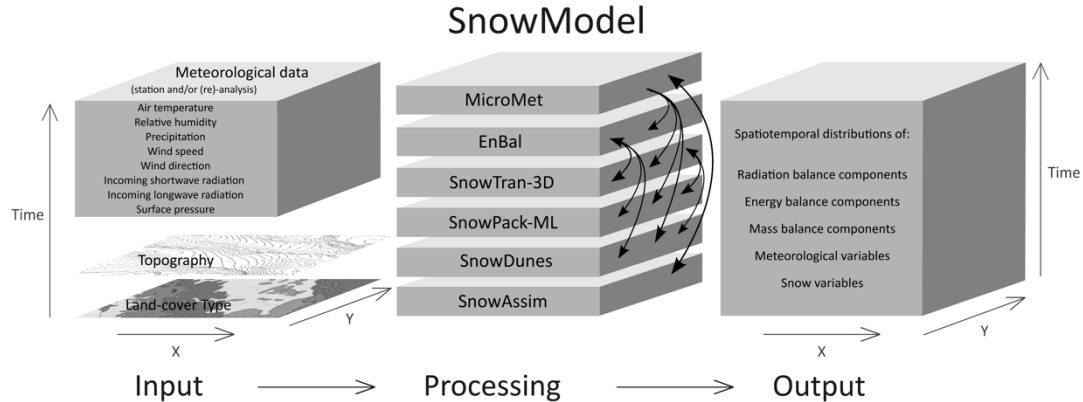
105 In Sect. 2, we provide background information on SnowModel, parallelization using CAF, data and domains used in this
106 study, and a motivation for this work. In Sect. 3, we explain our parallelization approach using CAF and introduce the
107 simulation experiments used to demonstrate the performance of Parallel SnowModel through strong scaling metrics and
108 CONUS simulations. In Sect. 4, we provide results of the simulation experiments introduced in Sect. 3. Lastly, we end with a
109 discussion in Sect. 5 and a conclusion in Sect. 6.

110 **2 Background**

111 **2.1 SnowModel**

112 SnowModel is a spatially distributed snow-evolution modeling system designed to model snow states (e.g., snow depth,
113 SWE, snow melt, snow density) and fluxes over different landscapes and climates (Liston and Elder, 2006a). The most
114 complete and up-to-date description of SnowModel can be found in the Appendices of Liston et al. (2020). While many
115 snow modelling systems exist, SnowModel will benefit from parallelization because of its ability to simulate snow processes
116 on a high-resolution grid through downscaling meteorological inputs and modelling snow redistribution. SnowModel is
117 designed to simulate domains on a structured grid with spatial resolutions ranging from 1 to 200 m (although it can simulate
118 coarser resolutions, as well) and temporal resolutions ranging from 10 m to 1 d. The primary modeled processes include
119 accumulation from frozen precipitation; blowing-snow redistribution and sublimation; interception, unloading, and
120 sublimation within forest canopies; snow-density and grain-size evolution; and snowpack ripening and melt. These processes
121 are distributed into four, core interacting submodules: MicroMet defines the meteorological forcing conditions (Liston and
122 Elder, 2006b), EnBal describes surface and energy exchanges (Liston, 1995; Liston et al., 1999), SnowPack-ML is a
123 multilayer snowpack sub-model that simulates the evolution of snow properties and the moisture and energy transfers
124 between layers (Liston and Hall, 1995; Liston and Mernild, 2012), and SnowTran-3D calculates snow redistribution by wind
125 (Liston et al., 2007). Additional simulation features include SnowDunes (Liston et al., 2018) and SnowAssim (Liston and
126 Hiemstra, 2008), which model sea-ice applications and data assimilation techniques, respectively. Figure 1 shows a
127 schematic of the core SnowModel toolkit. Additionally, the initialization submodules that read in the model parameters,
128 distribute inputs across the modeled grid, allocate arrays, etc., include PreProcess and ReadParam. Outputting arrays is
129 contained within the Outputs submodule. SnowModel incorporates first-order physics required to simulate snow evolution

130 within each of the global snow classes [e.g., Ice, Tundra, Boreal Forest, Montane Forest, Prairie, Maritime, and Ephemeral;
131 (Sturm and Liston, 2021; Liston and Sturm, 2021)].



132

133 **Figure 1: The original figure from Pedersen et al. (2015) was modified for the present paper, providing an example of possible**
134 **inputs, core submodules, and outputs of SnowModel.**

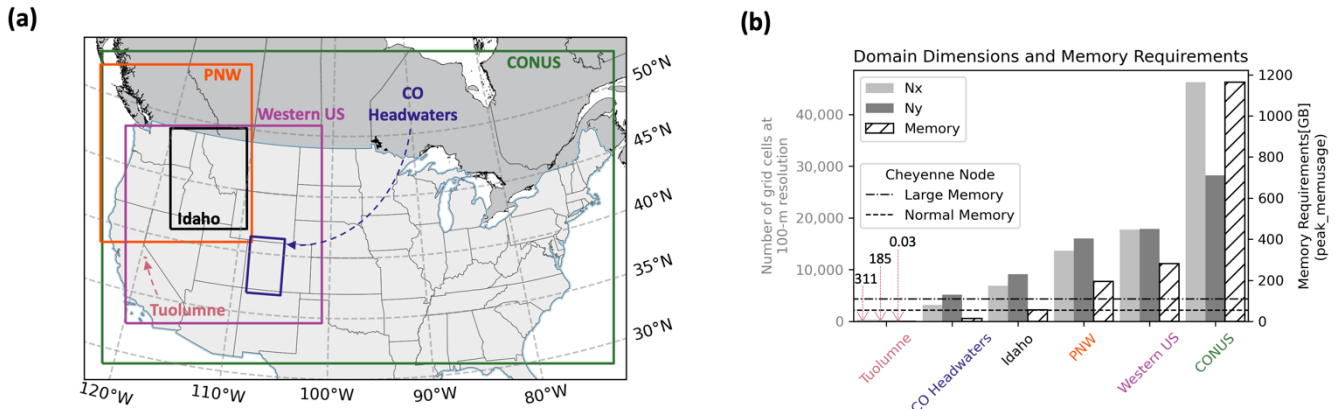
135 2.2 Coarray Fortran

136 CAF, formerly known as F-, (Iso/Iec, 2010; Numrich and Reid, 1998; Numrich et al., 1997) is the parallel language feature
137 of Fortran that was used to parallelize SnowModel. CAF is like Message Passing Interface (MPI) libraries in that it uses the
138 Single Program Multiple Data (SPMD) model where multiple independent cores simultaneously execute a program. SPMD
139 allows for distributed memory allocation and remote memory transfer. However, unlike MPI, CAF uses the PGAS parallel
140 programming model to handle the distribution of computational tasks amongst processes (Coarfa et al., 2005). In the PGAS
141 model, each process contains local memory that can be accessed directly by all other processes. While CAF and MPI syntax
142 often refers to processes as images or ranks, for consistency, we will continue to use the term “process”. Ultimately, CAF
143 offers a high-level syntax that exploits locality and scales effectively (Coarfa et al., 2005). For simulation comparisons, we
144 used OpenCoarrays, a library implementation of CAF (Fanfarillo et al., 2014) utilized by the gfortran compiler; intel and
145 cray compilers both have independent CAF implementations.

146 2.3 Model Domains, Data, and Computing Resources

147 The required inputs for SnowModel include 1) temporally varying meteorological variables of precipitation, wind speed and
148 direction, air temperature, and relative humidity taken from meteorological stations or atmospheric models and 2) spatially
149 distributed topography and land-cover type (Liston & Elder, 2006a). The following inputs were used for the experiments
150 introduced in Sect. 3: USGS National Elevation Dataset (NED) for topography (Gesch et al., 2018), The North American
151 Land Change Monitoring System (NALCMS) Land Cover 2015 map for vegetation (Homer et al., 2015; Jin et al., 2019;
152 Latifovic et al., 2016), and forcing variables from either the North American Land Data Assimilation System (NLDA-2)
153 (Mitchell, 2004; Xia, 2012a, b) on a 1/8 degree (approximately 12 km) grid or a high-resolution Weather Research Forecast

154 (WRF) model from the National Center for Atmospheric Research (NCAR) on approximately a 4 km grid (Rasmussen et al.,
 155 2023). The high-performance computing architectures used include NCAR’s Cheyenne supercomputer, which is a 5.43-
 156 petaflop SGI ICE XA Cluster featuring 145,152 Intel Xeon processes in 4,032 dual-socket nodes and 313 TB of total
 157 memory (Laboratory, 2019) and The National Aeronautics and Space Administration’s (NASA) Center for Climate
 158 Simulation (NCCS) Discover supercomputer with a 1,560-teraflop SuperMicro Cluster featuring 20,800 Intel Xeon Skylake
 159 processes in 520 dual-socket nodes and 99.84 TB of total memory. Simulation experiments were conducted over six domains
 160 (Tuolumne, CO Headwaters, Idaho, PNW, Western US, and CONUS) throughout the United States at 100 m grid resolution.
 161 The spatial location, domain dimensions (e.g., number of grids in the x and y dimensions), and memory requirements,
 162 derived from the peak_memusage package (https://github.com/NCAR/peak_memusage), for the simulation experiments are
 163 highlighted in Fig. 2.



164
 165 **Figure 2:** (a) Spatial location of simulated domains on WRF’s Lambert conformal projection (Rasmussen et al., 2023) and (b)
 166 corresponding grid dimensions (N_x – number of grids in x dimension; N_y – number of grids in y dimension) and memory obtained
 167 from peak_memusage package required for single-layer SnowModel simulation experiments. For reference, the dashed lines represent
 168 the normal and large memory thresholds (55 and 109 GB) for Cheyenne’s SGI ICE XA cluster.

169 2.4 Parallelization Motivation

170 The answers to current snow science, remote sensing, and water management questions require high-resolution data that
 171 covers large spatial and temporal domains. While modeling systems like SnowModel can be used to help provide these
 172 datasets, running them on single-processor workstations imposes limits on the spatiotemporal extents of the produced
 173 information. Serial simulations are limited by both execution time and memory requirements, where the memory limitation
 174 is largely dependent on the size of the simulation domain. Up to the equivalent of 175 two-dimensional and 10 three-
 175 dimensional arrays are held in memory during a SnowModel simulation, depending on the model configuration. In analyzing
 176 the performance of the Parallel SnowModel (Sect. 4), serial simulations were attempted over six domains throughout the
 177 United States at 100 m grid resolution (Fig. 2) for the 2018 water year (1 September 2017 to 1 September 2018). Only the
 178 Tuolumne domain could be simulated in serial based on the memory (109 GB for a large memory node) and time (12 h wall-
 179 clock limit) constraints on Cheyenne. The CO Headwaters and Idaho domains could not be simulated in serial due to time

180 constraints, while the three largest domains (Pacific Northwest (PNW), Western U.S. and CONUS) could not be executed in
181 serial due to both exceedances of the 12 h wall-clock limit and memory availability. Furthermore, we estimate that using a
182 currently available, state of the art, single-processor workstation, would require approximately 120 d of computer time to
183 perform a 1 y model simulation over the CONUS domain. SnowModel is regularly used to perform multi-decade
184 simulations, for trend analyses, climate change studies, and retrospective analyses (Liston and Hiemstra, 2011; Liston et al.,
185 2020; Liston et al., 2022). If this 1 y, 100 m, CONUS domain was simulated for a 40 y period (e.g., 1980 through present), it
186 would take approximately 4800 d, or over 13 y, of computer time. Clearly such simulations are not practical using single-
187 processor computer hardware and software algorithms.

188 **3 Methods**


189 In parallelizing SnowModel and distributing computations and memory over multiple processes, we demonstrate its ability
190 to efficiently run regional to continental sized simulations. Some of the model configurations were not parallelized for
191 reasons including ongoing development in the serial code base and limitations to the parallelization approach. These
192 configurations are further discussed in Appendix A. This section introduces the syntax and framework used to parallelize
193 SnowModel and the simulation experiments used to assess the performance of the parallel algorithm.

194 **3.1 Parallel Implementation**

195 Changes to the SnowModel logic were made through the parallelization process and included the partitioning algorithm,
196 non-trivial communication via halo-exchange, and file input and output (I/O) schemes.

197 **3.1.1 Partitioning Algorithm**

198 The partitioning strategy identifies how the workload gets distributed amongst processes in a parallel algorithm. The
199 multidimensional arrays of SnowModel are stored in row-major order, meaning the x dimension is contiguous in memory.
200 Additionally, dominant wind directions and therefore predominant snow redistribution occurs in the east-west direction as
201 opposed to south-north directions. Therefore, both the data structures and physical processes involved in SnowModel justify
202 a one-dimensional decomposition strategy in the y dimension, where the computational global domain $N_x \times N_y$ is separated
203 into $N_x \times l_{ny}$ blocks. If N_y is evenly divisible by the total number of processes (N), $l_{ny} = N_y / N$. If integer division is
204 not possible, the remaining rows are distributed evenly amongst the processes starting at the bottom of the computational
205 domain. Figure 3 demonstrates how a serial domain containing 10 grid cells in the x and y dimensions would be
206 decomposed with four processes using our partitioning strategy.



207

208 **Figure 3: Example 10 x 10 global domain and partitioning for (a) serial simulation and (b) parallel simulation using four processes.**

209

3.1.2 Non-trivial Parallelization

210

Each process has sufficient information to correctly execute most of the physical computations within SnowModel.

211

However, there are certain subroutines where grid computations require information from neighboring grid cells (e.g., data

212

dependencies) and therefore information outside of the local domain of a process. For SnowModel, these subroutines

213

typically involve the transfer of blowing snow or calculations requiring spatial derivatives. Furthermore, with our one-

214

dimensional decomposition approach, each grid cell within a process local domain has sufficient information from its

215

neighboring grid cells in the x dimension but potentially lacks information from neighboring grid cells in the y dimension. As

216

a regular grid method, SnowModel lends itself to process communication via halo-exchange where coarrays are used in

217

remote calls. Halo-exchange using CAF involves copying boundary data into coarrays on neighboring processes and using

218

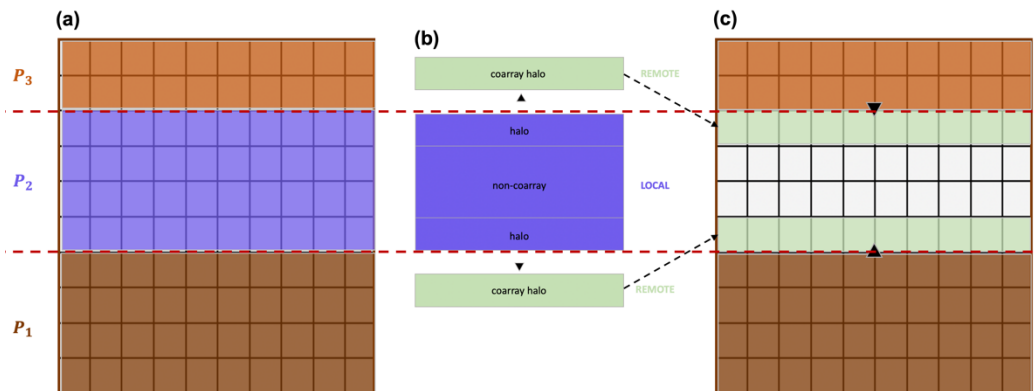
information from the coarrays to complete computations (Fig. 4). Although the entire local array could be declared a coarray

219

and accessed by remote processes more directly, some CAF implementations, (e.g. Cray) impose additional constraints upon

220

coarray memory allocations that can be problematic for such large allocations.

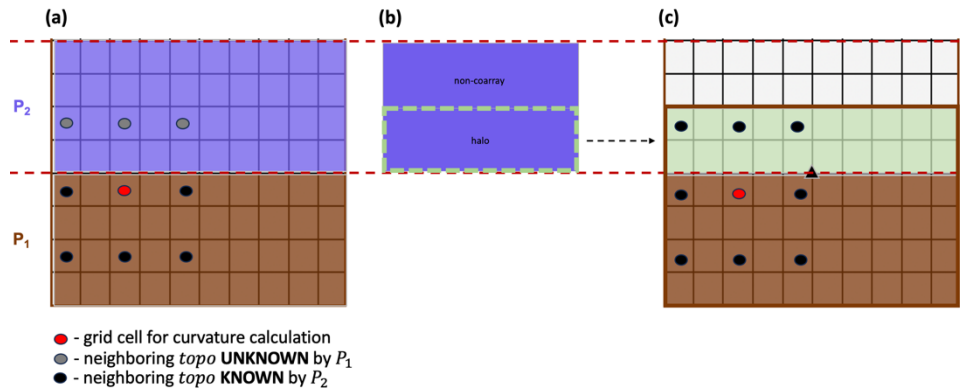


221

222 **Figure 4: Schematic showing halo-exchange using coarrays. The steps include: (a) initial gridded representation of local arrays for**
 223 **three processes, (b) P_2 copying boundary data into coarrays for remote access, (c) neighboring processes (P_1 and P_3) stitching coarray**
 224 **to local domains.**

225 **3.1.2.1 Topography – Wind and Solar Radiation Models**

226 The wind and solar radiation models in MicroMet require information about surrounding surface topography (Liston and
 227 Elder, 2006b). The wind model requires surface curvature, and the solar radiation model requires surface slope and aspect.
 228 These vary at each timestep as snow accumulates and melts because the defined surface includes the snow surface on top of
 229 the landscape. The surface curvature, for example, is computed at each model grid cell using the spatial gradient of the
 230 topographic elevation of eight neighboring grid cells. Using the parallelization approach discussed above, processes lack
 231 sufficient information to make curvature calculations for the bordering grid cells along the top and/or bottom row(s) within
 232 their local domains. Note that the number of row(s) (*inc*) is determined by a predefined parameter that represents the
 233 wavelength of topographic features within a domain. Future work should permit this parameter to vary spatially to account
 234 for changes in the length scale across the domain. For example, all grid cells along the top row of P_1 will be missing
 235 information from nearby grid cells to the north and require topographic elevation (*topo*) information from the bottom
 236 row(s) of the local domain of P_2 to make the calculation (Fig. 5a). Halo-exchange is performed to distribute row(s) of data to
 237 each process that is missing that information in their local domains (Fig. 5b). Processes whose local domains are positioned
 238 in the bottom or top of the global domain will only perform one halo-exchange with their interior neighbor, while interior
 239 processes will perform two halo-exchanges. By combining and appropriately indexing information from the process local
 240 array and received coarrays of topographic elevation, an accurate curvature calculation can be performed using this parallel
 241 approach (Fig. 5c).



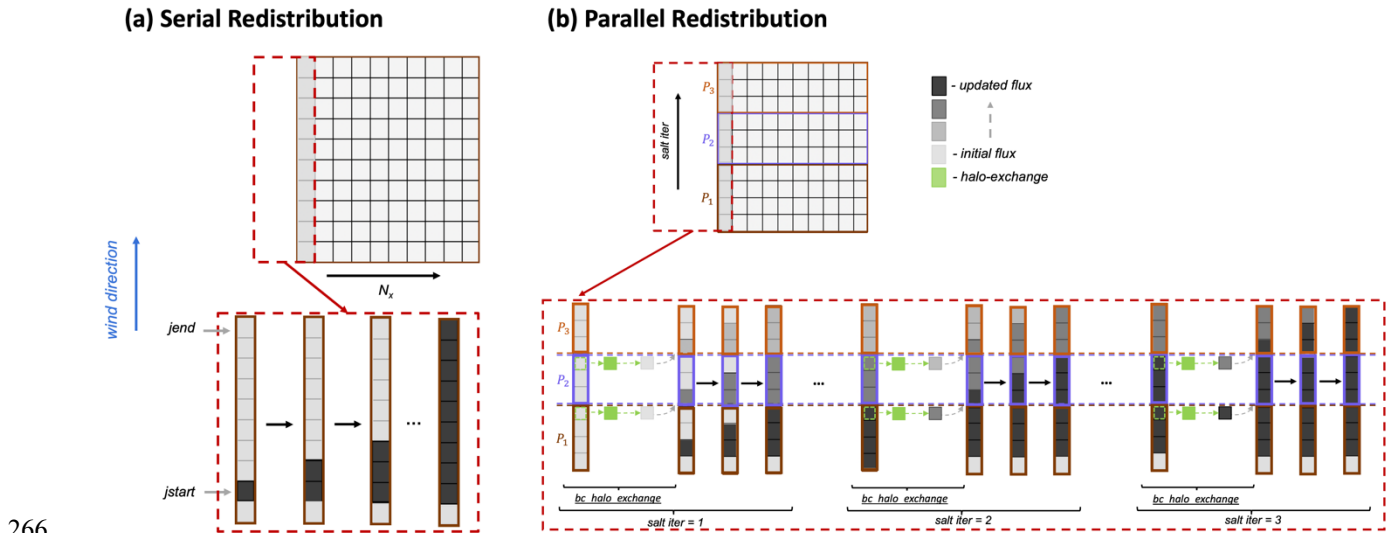
242

243 **Figure 5: Schematic for halo-exchange used in the curvature calculation by P_1 , where $inc = 2$. (a) Prior to halo-exchange, P_1**
 244 **contains insufficient information to perform the curvature calculation, (b) grid cells (halo) within the local domain of P_2 are (c)**
 245 **transferred to P_1 via coarrays. At this point, P_1 has sufficient information to make the curvature calculation.**

246 **3.1.2.2 Snow Redistribution**

247 Wind influences the mass balance of the snowpack by suspending and transporting snow particles in the air (turbulent-
 248 suspension) and by causing snow grains to bounce on top of the snow surface (saltation). In SnowModel, the saltation and
 249 suspension algorithms are separated into northerly, southerly, easterly, and westerly fluxes based on the u and v components
 250 of wind direction for each grid cell. Figure 6 shows a simplified schematic for the saltation flux from a southerly wind. In the

251 serial algorithm (Fig. 6a), SnowModel initializes the saltation flux based on the wind speed at that time step (*initial*
 252 *flux*). To calculate the final saltation flux (*updated flux*), SnowModel steps through regions of continuous wind
 253 direction (delineated by the indices: *jstart* and *jend*), updates the change in saltation fluxes from upwind grid cells and
 254 the change in saltation flux from the given wind direction, and makes adjustments to these fluxes based on the soft snow
 255 availability above the vegetation height (Liston and Elder, 2006a). Similar logic is used for the parallel implementation of
 256 the saltation and suspension fluxes with an additional iteration (*salt iter*) that updates the boundary condition for each
 257 process via halo-exchange. This allows the fluxes to be communicated from the local domain of one process to another. To
 258 minimize the number of iterations, *salt iter* was provided a maximum bound that is equivalent to snow being
 259 transported 15 km via saltation or suspension. This number was chosen based off prior field measurements (Tabler, 1975)
 260 and simulation experiments. It is possible that in other environments an even larger length may be required. To be
 261 guaranteed to match the serial results in all cases, the number of iterations would have to be equal to the number of
 262 processes; however, this would result in no parallel speed up and has no practical benefit. A schematic of the parallel
 263 calculation of the change in saltation due to southerly winds is illustrated in Fig. 6b. The *bc_halo_exchange* represents
 264 a halo-exchange of grid cells from upwind processes, allowing the saltation flux to be transported from one process local
 265 domain to the next.

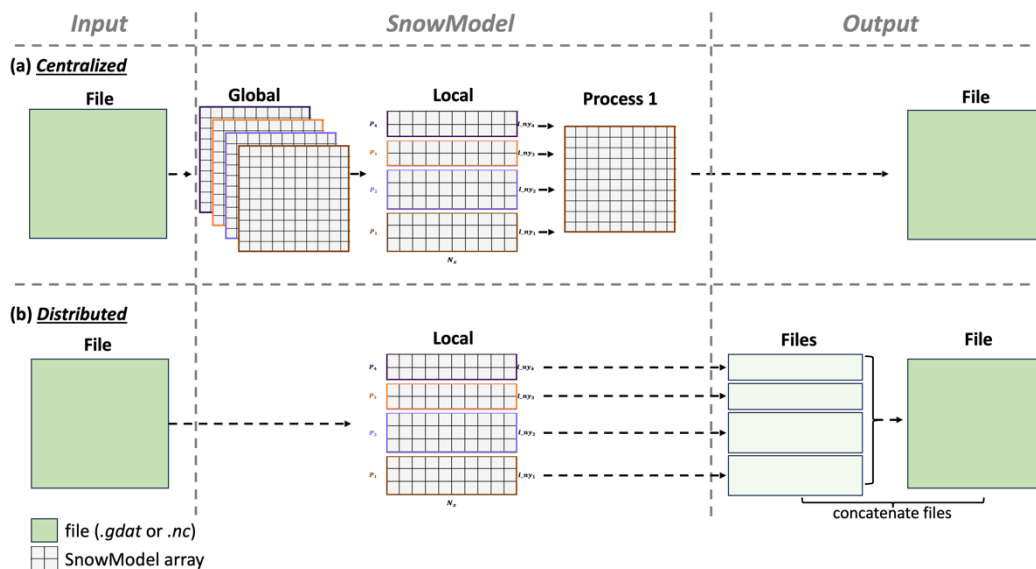


266
 267 **Figure 6:** (a) Schematic of the serial and (b) parallel redistribution algorithm showing the change in saltation flux due to southerly
 268 winds over a gridded domain for $N_x = 1$. The parallel schematic demonstrates how three processes (P_1 , P_2 , P_3) use an additional
 269 iteration (*salt iter*) to perform a halo-exchange (*bc_halo_exchange*) and update the boundary condition of the saltation flux.

270 3.1.3 File I/O

271 File I/O management can be a significant bottleneck in parallel applications. Parallel implementations that are less memory
 272 restricted commonly use local to global mapping strategies, or a **Centralized** approach for file I/O (Fig. 7a). This approach
 273 requires that one or more processes stores global arrays for input variables and that one process (Process 1; Fig. 7a) stores

274 global arrays for all output variables. As the domain size increases, the mapping of local variables to global variables for
 275 outputting creates a substantial bottleneck. To improve performance, *Distributed* file I/O can be implemented, where input
 276 and output files are directly and concurrently accessed by each process (Fig. 7b).



277

278 **Figure 7: (a) Schematic of global to local mapping for file I/O using a Centralized approach with four processes, and (b) Distributed**
 279 **file I/O where each process reads and writes data corresponding to its local domain.**

280 SnowModel contains static spatial inputs that do not vary over time (e.g., topography and land cover) and dynamic spatial
 281 inputs (e.g., air temperature and precipitation) that vary spatially and temporally. The static inputs are of a higher resolution
 282 compared to the dynamic inputs (cf., topography is on the model grid, while atmospheric forcing is almost always more
 283 widely spaced). To balance performance and consistency with the serial logic of the code, we used a mixed parallel file I/O
 284 approach. A goal of this work was to maintain nearly identical serial and parallel versions of the code in one code base that
 285 can be easily maintained and utilized by previous, current, and future SnowModel users with different computational
 286 resources and skills. Therefore, we wanted to maintain both the *Centralized* and *Distributed* file I/O approaches. However,
 287 for optimal parallel performance over larger simulation domains, file input (reading) is performed in a *Distributed* way for
 288 the static inputs and in a *Centralized* way for dynamic inputs, while file output (writing) is performed in a *Distributed* way,
 289 as described further below. This permits the new version of the code to be a drop in replacement for the original serial code
 290 without requiring users to install new software libraries or manage hundreds of output files, while enabling users who wish
 291 to take advantage of the parallel nature of the code to do so with minimal additional work and no changes to the underlying
 292 code.

293 3.1.3.1 Parallel Inputs

294 As noted above, SnowModel has two primary types of input files, temporally static files such as vegetation and topography
 295 and transient inputs such as meteorological forcing data. While acceptable static input file types include flat binary, NetCDF,

296 and ASCII files for the serial version of the code, optimizing the efficiency of Parallel SnowModel requires static inputs
297 from binary files that can be accessed concurrently and directly subset by indexing the starting byte and length of bytes
298 commensurate to a process local domain. Therefore, each process can read its own portion of the static input data. For very
299 large domains, the available memory becomes a limitation when using the centralized approach. For example, the CONUS
300 simulation could not be simulated using a centralized file I/O approach because each process would be holding global arrays
301 of topography and vegetation in memory, each of which would require approximately 5.2 GB of memory per process.
302 Reading of meteorological forcing variables (wind speed, wind direction, relative humidity, temperature, and precipitation)
303 can be performed in parallel with either binary or NetCDF files. Depending on the forcing dataset, the grid spacing of the
304 meteorological variables typically ranges from 1 to 30 km and therefore often requires a smaller memory footprint than static
305 inputs for high-resolution simulations. For example, the resolution of NLDAS-2 meteorological forcing has a grid of
306 approximately 11 km, while the high-resolution WRF model used has a 4 km grid. At each timestep, processes read in the
307 forcing data from every station within the domain into a one-dimensional array, index the nearest locations for each
308 SnowModel grid, and interpolate the data to create forcing variables over the local domain. All processes perform the same
309 operation and store common information; however, since the resolutions of the forcing datasets are significantly coarser than
310 the model grid for high-resolution simulations, the dynamic forcing input array size remains comparable to other local arrays
311 and does not impose significant memory limitations for simulations performed to date. While more efficient parallel file
312 input schemes could improve performance, we decided to keep this logic in part to maintain consistency with the serial
313 version of the code and minimize code changes.

314 **3.1.3.2 Parallel Outputs**

315 To eliminate the use of local to global mapping commonly used to output variables (Fig. 7a), each process writes its own
316 output file (Fig. 7b). A postprocessing script is then used to concatenate files from each process into one file that represents
317 the output for the global domain. Modern high-performance computing architectures have highly parallelized storage
318 systems making file output using a distributed approach significantly faster than the centralized approach. Therefore, file
319 output in this manner reduces time and memory requirements. Future work could leverage other established parallel I/O
320 libraries at the cost of additional installation requirements.

321 **3.2 Simulation Experiments**

322 Parallel SnowModel experiments were conducted to both evaluate the effectiveness of the parallelization approach used in
323 this study (Sect. 3.1) and to produce a high-resolution snow dataset over CONUS. All experiments were executed with a 100
324 m grid increment, a 3 h time step, a single-layer snowpack configuration, and included the primary SnowModel modules
325 (MicroMet, EnBal, SnowPack, and SnowTran-3D). These experiments are further described below, with results provided in
326 Sect. 4.

327 Validation experiments comparing output from the original serial version of the code to the parallel version were conducted
328 continuously throughout the parallel algorithm development to assess the reproducibility of the results. Additionally, a more

329 thorough validation effort was performed at the end of the study that compared output from the serial algorithm to that of the
330 parallel algorithm, while varying the domain size, the number of processes, and therefore the domain decomposition. Results
331 from all of these validation experiments produced root mean squared error (RMSE) values of 10^{-6} , which is at the limit of
332 machine precision, when compared to serial simulation results. See Appendix B for more details on the validation
333 experiments. The serial version of SnowModel has been evaluated in many studies across different snow classes (Sturm and
334 Liston, 2021; Liston and Sturm, 2021), time periods, and snow properties. Evaluations ranged from snow cover (Pedersen et
335 al., 2016; Randin et al., 2015), snow depth (Szczypta et al., 2013; Wagner et al., 2023), SWE (Freudiger et al., 2017;
336 Hammond et al., 2023; Morteza pour et al., 2020; Voordendag et al., 2021), and SWE-melt (Hoppinen et al., 2023; Lund et
337 al., 2022), using field observations, snow-telemetry stations, and remote sensing products. A full comparison of the Parallel
338 SnowModel simulations presented here with observations across CONUS is beyond the scope of the present work.
339 Incorrectly simulated SWE could affect the scaling results and CONUS visualizations presented in Sect. 3.2.1.1, 3.2.1.2, and
340 3.2.2; for example, if zero SWE were incorrectly simulated in many locations, processing time would be less than if SWE
341 had been simulated and tracked. However, based on the scale of these analyses and the fact that SnowModel has been
342 previously evaluated in a wide range of locations, we believe the impacts of this limitation on the computational results
343 presented here are minimal.

344 **3.2.1 Parallel Performance**

345 In high performance computing, scalability attempts to assess the effectiveness of running a parallel algorithm with an
346 increasing number of processes. Thus, scalability can be used to identify the optimal number of processes for a fixed domain,
347 understand the limitations of a parallel algorithm as a function of domain size and number of processes, and estimate the
348 efficiency of the parallel algorithm on new domains or computing architectures. Speedup, efficiency, and code profiling
349 were tools used to assess the scalability and performance of Parallel SnowModel on fixed domains. Speedup $[S(N); \text{Eq.}$
350 $(1)]$, a metric of strong scaling, is defined as the ratio of the serial execution time, $T(1)$, over the execution time using N
351 processes, $T(N)$. Optimally, parallel algorithms will experience a doubling of speedup as the number of processes is
352 doubled. Some reasons why parallel algorithms do not follow ideal scaling include the degree of concurrency possible and
353 overhead costs due to communication. Synchronization statements have an associated cost of decreasing the speed and
354 efficiency of an algorithm due to communication overhead and requirements for one process to sit idle while waiting for
355 another to reach the synchronization point. Furthermore, speedup tends to peak or plateau at a certain limit on a given
356 computing architecture and domain because either the overheads grow with an increasing number of processes, or the
357 number of processes exceeds the degree of concurrency inherent in the algorithm (Kumar and Gupta, 1991). For large
358 domains, where serial simulations cannot be performed either due to wall-clock or memory limitations, relative speedup,
359 $[\hat{S}(N); \text{Eq. (2)}]$, is commonly used. Relative speedup is estimated as a ratio of the execution time, $T(\hat{P})$, of the minimum
360 number of processes, (\hat{P}) , that can be simulated on a given domain over $T(N)$. An additional speedup metric, approximate

361 speedup [$\hat{S}(N)$; Eq. (3)], is introduced to estimate S by assuming perfect scaling from \hat{P} to a single process. While this is
362 only an approximation, it is helpful to compare the \hat{S} across the different domains on a similar scale. Additionally, efficiency
363 [$E(N)$; Eq. (4)], and approximate efficiency [$\check{E}(N)$; Eq. (5)] are the ratios of S to N and \hat{S} to N , respectively. A simulation
364 that demonstrates ideal scaling, would have 100% efficiency. Additionally, code profiling evaluates the cumulative
365 execution time of individual submodules (e.g. Preprocess, Readparam, MicroMet, Enbal, SnowPack, SnowTran-3D, and
366 Output) as a function of the number of processes. Together, code profiling and strong scaling can be used to understand
367 locations of bottlenecks in the algorithm and how changes to the code enhance performance.
368

$$S(N) = \frac{T(1)}{T(N)} \quad \text{Eq. 1}$$

$$\hat{S}(N) = \frac{T(\hat{P})}{T(N)} \quad \text{Eq. 2}$$

$$\check{S}(N) = \frac{T(\hat{P})}{T(N)} * \hat{P} \quad \text{Eq. 3}$$

$$E(N) = \frac{S}{N} * 100\% \quad \text{Eq. 4}$$

$$\check{E}(N) = \frac{\check{S}}{N} * 100\% \quad \text{Eq. 5}$$

369 3.2.1.1 Parallel Improvement

370 To better understand how changes to the Parallel SnowModel code have affected its performance, speedup and code
371 profiling plots were assessed for simulations using three distinct versions of the code. These versions represent snapshots of
372 the algorithms development and quantify the contributions of different types of code modifications to the final performance
373 of the model. These versions were identified by different GitHub commits (Mower et al., 2023) and can be summarized as
374 follows. The first or baseline version represents an early commit of Parallel SnowModel, where file I/O is performed in a
375 *Centralized* way, as described in Sect. 3.1.3. Each process stores both a local and global array in memory for all input
376 variables, makes updates to its local arrays, distributes that updated information into global arrays used by one process to
377 write each output variable. The embarrassingly parallel portion of the physics code has been parallelized, but the snow
378 redistribution step is not efficiently parallelized, it has a larger number of synchronizations and memory transfers. Therefore,
379 this approach has significant time and memory constraints. The *Distributed* version represents an instance of the code where
380 distributed file I/O (Sect. 3.1.3) had first been implemented. In this version, each process reads and writes input and output
381 variables for its local domain only. Global arrays and the communication required to update these variables are no longer
382 needed; this alleviates memory constraints and shows the value of parallelizing I/O in scientific applications. Lastly, the
383 *Final* version represents the most recent version of Parallel SnowModel, (at the time of this publication) where the snow
384 transport algorithm had been optimized to run efficiently. This was done by reducing unnecessary memory allocations,

385 reducing the transfer of data via coarrays, and optimizing memory transfers to reduce synchronization calls. This shows the
386 value of focused development on a single hotspot of the code base. The simulations were executed on the CO Headwaters
387 domain (Fig. 2) using 1, 2, 4, 16, 36, 52, 108, and 144 processes, outputted only a single variable, and were forced with
388 NLDAS-2 data from 23-24 March 2018. While 2-days is a short period to perform scaling experiments, a significant amount
389 of wind and frozen precipitation was observed over the CO Headwaters domain during the simulation to activate some of the
390 snow redistribution schemes in SnowTran-3D. Furthermore, to avoid disproportionately weighing the initialization of the
391 algorithm, we removed the timing values from the ReadParam and Preprocess submodules from the total execution time
392 used in the speedup analysis. Results from these experiments are provided in Sect. 4.1.

393 **3.2.1.2 Strong Scaling**

394 Strong scaling experiments of Parallel SnowModel were evaluated by comparing the approximate speedup and efficiency (\bar{S}
395 and \bar{E}) over six different size domains across the United States, all with a 100 m grid spacing [Tuolumne, CO Headwaters,
396 Idaho, PNW, Western U.S., and CONUS] (Fig. 2). These experiments use the *Final* version of the code according to Sect.
397 3.2.1.1. The simulations were forced with NLDAS-2 data for 2928 timesteps from 1 September 2017 to 1 September 2018
398 and output one variable (SWE). The number of processes used in these simulations varied by domain based on the 12 h wall-
399 clock and memory constraints on Cheyenne. Results from these experiments are provided in Sect. 4.2.

400 **3.2.2 CONUS Simulations**

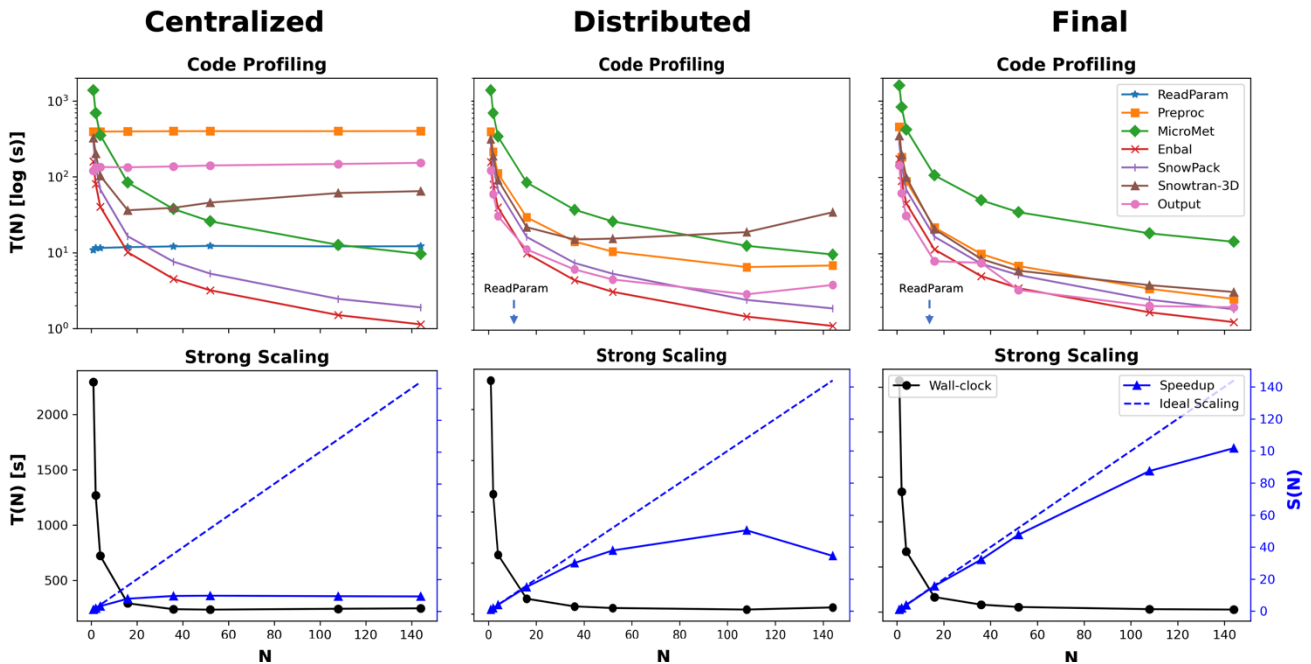
401 A primary goal of this work was to run Parallel SnowModel simulations for 21 years (2000 – 2021) over the CONUS
402 domain (Fig. 2) on a 100 m grid, while resolving the diurnal cycle in the model physics and creating a daily dataset of snow
403 properties, including snow depth, SWE, melt rate, and sublimation. Future work will analyze results from these simulations.
404 The CONUS domain contained 46,238 and 28,260 grid cells in the x and y dimensions, respectively. Simulations were
405 performed on a 3 h time step and forced with the WRF dataset. All simulations were executed on Discover using 1800
406 processes with a total compute time of approximately 192,600 core hours, or approximately 5 wall-clock hours per year.

407 **4 Results**

408 **4.1 Parallel Improvement**

409 Figure 8 demonstrates how the scalability of Parallel SnowModel evolved, as shown through code profiling (top row; Fig. 8)
410 and speedup (bottom row; Fig. 8) plots at three different stages (*Centralized*, *Distributed*, and *Final*) of the code
411 development. The code profiling plots display the cumulative execution time of each submodule ($T(N) [\log(s)]$) as a
412 function of the N . The strong scaling plots show the total execution time ($T(N) [s]$) and the speedup [$S(N)$; Eq. (1)] as a
413 function of N on the primary y-axis and secondary y-axis, respectively. As mentioned previously, the initialization timing
414 was removed from these values. The speedup of the *Centralized* version of the code quickly plateaus at approximately 10

415 processes. While the Enbal, SnowPack, and MicroMet subroutines scale with the number of processes (execution time
416 decreases proportional to the increase in the number of processes), the ReadParam, Preprocess, and Output subroutines,
417 which all perform file I/O or memory allocation, require a fixed execution time regardless of the number of processes used,
418 and the execution time of the SnowTran-3D submodule increases beyond 16 processes. This highlights the large bottleneck
419 that often occurs during the file I/O step in scientific code and the importance of code infrastructure outside of the physics
420 routines. In contrast, all the submodules in the *Distributed* version of the code, scale up to 36 processes, at which point the
421 inefficient parallelization of the SnowTran-3D submodule causes a significant slowdown, an increase in execution time as
422 the number of processes increases. This results in a speedup that plateaus at 52 processes and decreases beyond 108
423 processes. In the *Final* version of the code, scalability is observed well beyond 36 processes, with a maximum speedup of
424 100 observed using 144 processes. The execution time of all the submodules decreases as the number of processors
425 increases. This work highlights the value of going beyond the rudimentary parallelization of a scientific code base by
426 profiling and identifying individual elements that would benefit the most from additional optimization. This is a well-known
427 best practice in software engineering but often underappreciated in high-performance scientific computing. In Parallel
428 SnowModel, the improvement of these communication bottlenecks is primarily attributed to utilizing a distributed file I/O
429 scheme and minimizing processor communication by limiting the use of coarrays and synchronization calls. Ultimately,
430 without these improvements, the CONUS domain could not be simulated using Parallel SnowModel.




431
432 **Figure 8:** Code profiling (top row) and strong scaling (bottom row) results demonstrating the progression of Parallel SnowModel,
433 which includes a version of the code with centralized file I/O (*Centralized*; first column), a version of the code with distributed file
434 I/O (*Distributed*; second column), and a final version of the code at the time of this publication (*Final*; third column). These
435 versions can be found as different commits within the GitHub repository (Mower et al., 2023). The code profiling plots display the
436 cumulative execution time of each submodule on a logarithmic scale as a function of the number of processes (N). The arrow in the

437 code profiling plots of *Distributed* and *Final* indicates the ReadParam timing is below the y-axis at approximately 0.3 seconds and
438 0.003 seconds, respectively. The strong scaling plots show the total execution time ($T(N)$) against N on the primary y-axis and the
439 speedup (S) against N on the secondary y-axis.

440 4.2 Strong Scaling

441 In addition to the parallel improvement analysis, strong scaling was also performed on six domains for the 2018 water year
442 to better understand how Parallel SnowModel scales across different domain sizes and decompositions. Figure 9 displays the
443 approximate speedup [$\hat{S}(N)$; Eq. (3)] of Parallel SnowModel for three local/state domains (Tuolumne, CO Headwaters, and
444 Idaho) and three regional/continental domains (PNW, Western US, and CONUS). Additionally, Table 1 contains information
445 about the minimum and maximum number of processors (\hat{P} and P^* , respectively) simulated on each domain and their
446 corresponding execution time ($T(N)$ [m]), relative speedup [$\hat{S}(N)$; Eq. (2)], approximate speedup [$\check{S}(N)$; Eq. (3)], and
447 approximate efficiency [$\check{E}(N)$; Eq. (5)]. As mentioned previously, simulations were constrained by both the 12 h wall-clock
448 and 109 GB of memory per node on the Cheyenne supercomputer. In strong scaling, the number of processes is increased
449 while the problem size remains constant; therefore, it represents a reduced workload per process. Local-sized domains, e.g.,
450 Tuolumne, likely do not warrant the need for parallel resources because they have small serial runtimes (e.g., using 52
451 processes, Tuolumne had an \check{E} of 38%; Table 1). However, state, regional, and continental domains stand to benefit more
452 significantly from parallelization. The CONUS runtime decreased by a factor of 3 running on 3456 processes relative to 648
453 processes. Based on our approximate speedup assumption, we would estimate a CONUS \check{S} of 1690 times on 3456 processes
454 compared to one process, with an \check{E} of 49%. The Western US and PNW domains display very similar scalability results (Fig.
455 9), which is attributed to the similar number of grid cells in the y dimension (Fig. 2 and Table 1) and thus parallel
456 decomposition for each domain. Furthermore, these domains may also have a similar proportion of snow-covered grid cells.
457 While the PNW likely has more terrestrial grid cells that are covered by snow for a longer period throughout the water year,
458 it also has a significant number of ocean grid cells where snow redistribution would not be activated.

Parallel SnowModel - Strong Scaling



459

460 Figure 9: The left panel displays approximate speedup as a function of the number of processes (N) for local and state sized
 461 simulations (Tuolumne, CO Headwaters, and Idaho), while the right panel shows \hat{S} for the regional and continental sized domains
 462 (PNW, Western US, and CONUS).

463

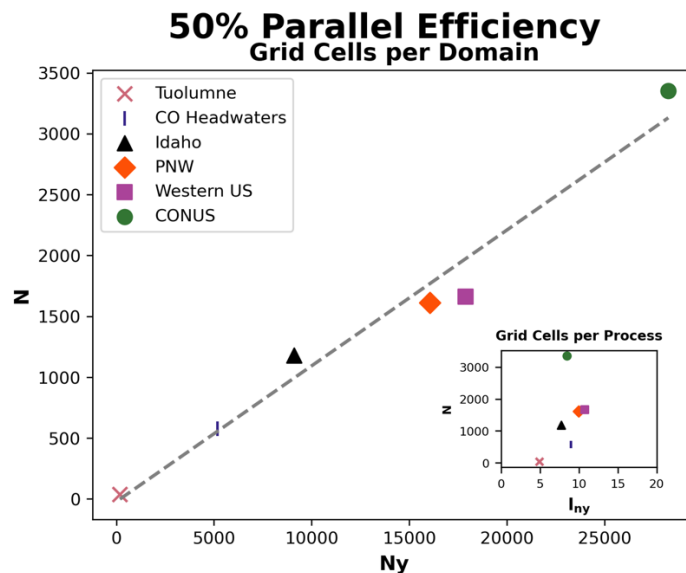
Domain	N_x	N_y	\hat{P} or P^*	Number of	Execution	Relative	Approximate	Approximate
				Processes	Time [m]	Speedup	Speedup	Efficiency
				N	$T(N)$	$\hat{S}(N)$	$\hat{S}(N)$	$\hat{E}(N)$
Tuolumne	311	185	\hat{P}	1	13	---	---	100
			P^*	52	1	20	20	38
CO Headwaters	3166	5167	\hat{P}	8	934	---	8	100
			P^*	576	24	39	308	53
Idaho	6916	9107	\hat{P}	27	1068	---	27	100
			P^*	1296	48	22	605	47
PNW	13677	16058	\hat{P}	84	1173	---	84	100
			P^*	2304	105	11	941	41
Western US	17737	17878	\hat{P}	120	1187	---	120	100
			P^*	3456	135	9	1058	31
CONUS	46238	28260	\hat{P}	648	1196	---	648	100
			P^*	3456	459	3	1690	49

464

465 Table 1: Parallel SnowModel strong scaling results containing grid dimensions (N_x and N_y), execution time [m], relative speedup,
 466 approximate speedup, and approximate efficiency for simulations executed with the minimum and maximum number of processes
 467 (\hat{P} and P^* , respectively) on the Tuolumne, CO Headwaters, Idaho, PNW, Western US, and CONUS domains. Values of the
 468 timing, speedup, and efficiency variables are rounded to the nearest integer.

469 Strong scaling analysis is useful for I/O and memory bound applications to identify a setup that results in a reasonable
 470 runtime and moderate resource costs. Based on these scaling results, Fig. 10 contains the relationship between the number of
 471 processes (N) at which each domain is estimated to reach 50% \hat{E} (using linear interpolation) with the total number of grid

472 cells in the y dimension (N_y) and the average number of grid cells in the y dimension per process (l_{ny} ; inset Fig. 10). At
 473 this level of efficiency, it is notable the consistency of both the linear relationship between N_y and N (8.7:1 ratio) and the
 474 values of l_{ny} (5 to 11) for these year-long simulations that vary in both domain size and the proportion of snow-covered
 475 area. Similar relationships (Fig. 10) can be used to approximate the scalability of Parallel SnowModel on different sized
 476 domains and can be adjusted for the desired level of efficiency. For example, we decided to run the CONUS simulations
 477 (Sect. 4.3) using 1800 processes based on its 70% approximate efficiency.

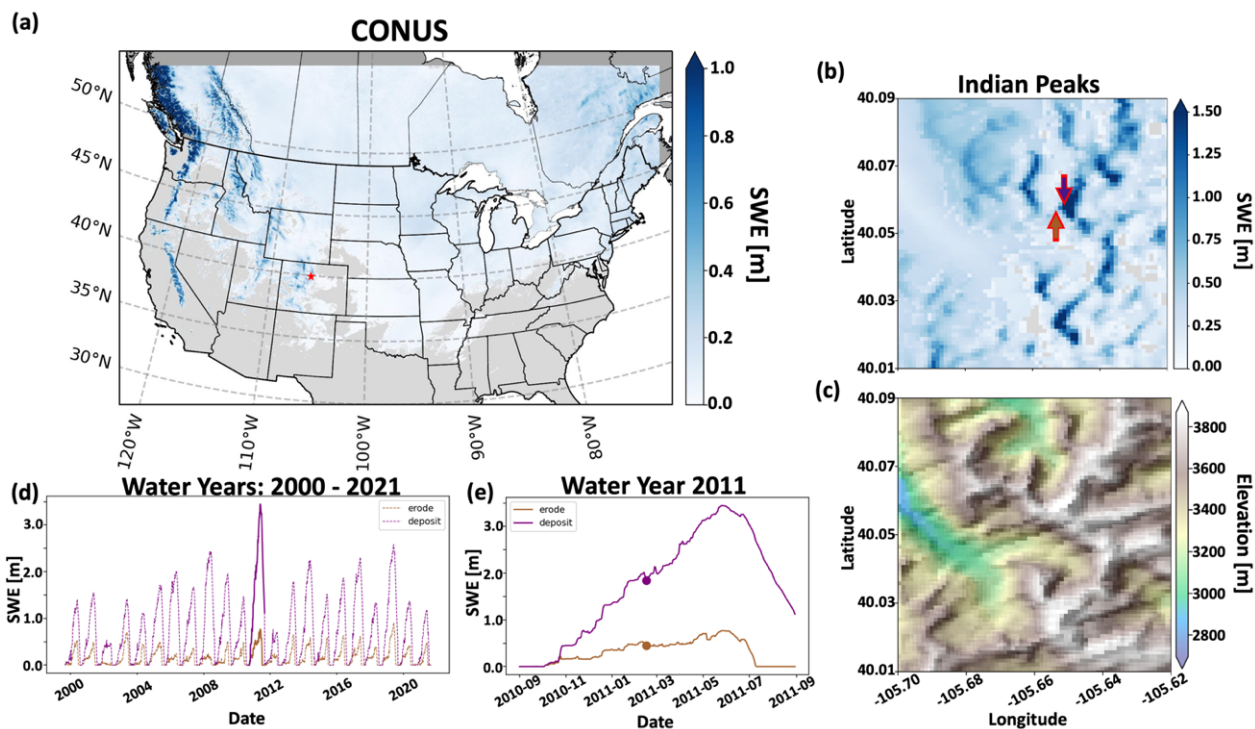


478
 479 **Figure 10: Relationship between the number of grid cells in the y dimension (N_y) and the number of processes (N) for each domain**
 480 **at which 50% approximate efficiency is estimated using the strong scaling analysis. The dashed line represents the best fit line for**
 481 **this relationship using OLS regression. The inset figure displays a similar relationship but compares N to the average number of**
 482 **grid cells in the y dimension per process (l_{ny}), instead of N_y .**

483 4.3 CONUS Simulations

484 Spatial results of SWE on 12 February 2011 over the CONUS domain and a sub-domain located in the Indian Peaks west of
 485 Boulder, Colorado are displayed in Fig. 11. On this date, simulated SWE was observed throughout the northern portion of
 486 the CONUS domain with the largest values concentrated in the mountain ranges (Fig. 11a). The Indian Peaks sub-domains of
 487 distributed SWE (Fig. 11b) with reference topography (Fig. 11c) underscores the ability of the large dataset to capture snow
 488 processes in a local alpine environment. It is important to note that while SnowModel does simulate snow redistribution, it
 489 does not currently have an avalanche model, which may be a limitation of accurately simulating SWE within this sub-
 490 domain. Additionally, Fig. 11b highlights two grid cells located 200 m apart on a peak. Figures 11d and 11e display the
 491 SWE evolution of these two grid cells over the entire dataset (water years 2000 – 2021) and the 2011 water year,
 492 respectively, further demonstrating the ability of Parallel SnowModel to capture fine-scale snow properties even when
 493 simulating continental domains. The upwind (western) grid cell is scoured by wind, and snow is transported to the downwind

494 (eastern) grid cells where a snow drift forms. The information and insight available in this high-resolution dataset will have
 495 important implications for many applications from hydrology, to wildlife and ecosystems, to weather and climate, and many
 496 more.



497

498 **Figure 11: Simulation results of Parallel SnowModel over CONUS using the WRF projection. (a) Spatial patterns of SWE over the**
 499 **CONUS domain for 12 February 2011, (b) highlighting the SWE distribution (c) and topography with an applied hillshade of a sub-**
 500 **domain near Apache Peak in the Indian Peaks west of Boulder, CO. (d) Time series of SWE from 2000-2021 and (e) over the 2011**
 501 **water year for grid cells (“erode” and “deposit”) identified in panel (b). The “erode” and “deposit” grid cells highlight areas of similar**
 502 **elevation but significant differences in SWE evolution resulting from blowing-snow redistribution processes.**

503 5 Discussion

504 Parallelizing numerical models often involves two-dimensional decomposition in both the x and y dimensions. While many
 505 benefits have been demonstrated by this approach, including improved load balancing (Dennis, 2007; Hamman et al., 2018),
 506 it comes with increased complication of the parallel algorithms, including the partitioning algorithm, file I/O, and process
 507 communication. The demonstrated speedup (Fig. 9) suggests Parallel SnowModel scales effectively over regional to
 508 continental domains using the one-dimensional decomposition approach. The added benefits obtained from two-dimensional
 509 decomposition strategies might not outweigh the costs of development, testing, and minimizing changes to the code structure
 510 and logic for applications such as SnowModel. Ultimately, our simplified parallelization approach can be implemented by
 511 other geoscience schemes as a first step to enhance simulation size and resolution.

512 Simulation experiments were conducted using Parallel SnowModel to validate the parallel logic, interpret its performance
513 across different algorithm versions and domain sizes, and demonstrate its ability to simulate continental domains at high-
514 resolution. Code profiling and speedup analyses over the CO Headwaters domain helped identify bottlenecks in file I/O and
515 processor communication in SnowTran-3D during the development of the parallel algorithm (Sect. 4.1). Corrections to the
516 referred bottlenecks allowed Parallel SnowModel to scale up to regional and continental sized simulations and highlights the
517 value of optimizing scientific code. For Parallel SnowModel scalability is primarily dependent on the number of grid cells
518 per process (N_x and l_{ny}) but is also affected by the proportion of snow-covered grid cells with sufficient winds and soft
519 snow available to be redistributed (Sect. 3.1.2.2). The scalability analyses showed similar results across domains with
520 significant differences in size (N_x and N_y), topography, vegetation, and snow classifications (Sturm et al., 1995; Sturm and
521 Liston, 2021) (Sect 4.2), highlighting the effectiveness of Parallel SnowModel for running state, regional, and continental-
522 sized domains. Furthermore, results from this analysis can be used to estimate the number of processors required to simulate
523 domains outside of the ones used in this study with a desired level of parallel efficiency (Fig. 10).
524 Additionally, these experiments emphasize the relationships among speed, memory, and computing resources for Parallel
525 SnowModel. A common laptop (~ 4 processes) has sufficient CPUs to run local sized domains within a reasonable amount
526 of time, but likely does not have sufficient memory for state-sized simulations. Similarly, the minimum memory (1160 GB;
527 Fig. 1) required to run the CONUS domain, could be simulated on a large server (~ 128 processes) with one process per
528 node. However, extrapolating from our scaling results on Cheyenne (Fig. 9), we estimate it would take over 2.5 days to run a
529 CONUS simulation for one water year with this configuration. In contrast, it took approximately 5 hours for CONUS to run
530 on the Discover supercomputer using 1800 processes. Therefore, by the time it took the large server to complete a CONUS
531 simulation for one water year, 12 water years could have been simulated on a supercomputer. Lastly, results from the
532 CONUS simulation highlight the ability of Parallel SnowModel to run high-resolution continental simulations, while
533 maintaining fine-scale snow processes that occur at a local level (Sect. 4.3).
534 SnowModel can simulate high-resolution outputs of snow depth, density, SWE, grain size, thermal resistance, snow strength,
535 snow albedo, landscape albedo, meltwater production, snow-water runoff, blowing snow flux, visibility, peak winter SWE,
536 snow-season length, snow onset date, snow-free date, and more, all produced by a physical model that maintains consistency
537 among variables. While several snow data products exist, few capture the suite of snow properties along with the spatio-
538 temporal extents and resolutions that can benefit a wide variety of applications. For example, current snow information
539 products include the NASA daily SWE distributions globally for dry (non-melting) snow on a 25 km grid (Tedesco and
540 Jeyaratnam, 2019), a NASA snow-cover product on a 500 m grid (Hall et al., 2006) that is missing information due to clouds
541 approximately 50% of the time (Moody et al., 2005), and the Snow Data Assimilation System (SNODAS) daily snow
542 information provided by the National Oceanic and Atmospheric Administration (NOAA) and the National Weather Service
543 (NWS) National Operational Hydrologic Remote Sensing Center (NOHRSC) on a 1 km grid (Center, 2004), which is itself
544 model derived and has limited geographic coverage and snow properties. The Airborne Snow Observatory (ASO) provides
545 the highest resolution data with direct measurements of snow depth on a 3 m grid, and derived values of SWE on a 50 m grid

546 (Painter et al., 2016), but has limited spatio-temporal coverage and a high cost of acquisition. Furthermore, there are many
547 fields of study that can benefit from 100 m resolution information of internally consistent snow variables, including wildlife
548 and ecosystem, military, hydrology, weather and climate, cryosphere, recreation, remote sensing, engineering and civil
549 works, and industrial applications. The new Parallel SnowModel described here permits the application of this modeling
550 system to very large domains without sacrificing spatial resolution.

551 **6 Conclusions**

552 In this paper, we present a relatively simple parallelization approach that allows SnowModel to perform high-resolution
553 simulations over regional to continental sized domains. The code within the core submodules (EnBal, MicroMet, SnowPack,
554 and SnowTran-3D) and model configurations (single-layer snowpack, multi-layer snowpack, binary input files, etc.) was
555 parallelized and modularized in this study. This allows SnowModel to be compiled with a range of Fortran compilers,
556 including modern compilers that support parallel CAF either internally or through libraries, such as OpenCoarrays
557 (Fanfarillo et al., 2014). Additionally, it provides the structure for other parallelization logic (e.g., MPI) to be more easily
558 added to the code base. The parallel module contains a simple approach to decomposing the computational domain in the y
559 dimension into smaller rectangular sub-domains. These sub-domains are distributed across processes to perform
560 asynchronous calculations. The parallelization module also contains logic for communicating information among processes
561 using halo-exchange coarrays for the wind and solar radiation models, as well as for snow redistribution. The scalability of
562 Parallel SnowModel was demonstrated over different sized domains, and the new code enables the creation of high-
563 resolution simulated snow datasets on continental scales. This parallelization approach can be adopted in other
564 parallelization efforts where spatial derivatives are calculated or fluxes are transported across gridded domains.

565 **Appendix A**

566 Some of the configuration combinations were not parallelized during this study for reasons including ongoing development
567 in the serial code base and limitations to the parallelization approach. These include simulations involving tabler surfaces
568 (Tabler, 1975), I/O using ASCII files, Lagrangian seaice tracking, and data assimilation.

569 **Appendix B**

570 Validation SnowModel experiments were run in serial and in parallel over the Tuolumne and CO Headwaters domains (Sect.
571 4.1) using the RMSE statistic. Important output variables from EnBal, MicroMet, SnowPack, and SnowTran-3D
572 demonstrated similar, if not identical values, when compared to serial results for all timesteps during the simulations; RMSE
573 values were within machine precision ($\sim 10^{-6}$) regardless of the output variable, domain, or number of processes used. The

574 validated output variables include albedo [%], precipitation [m], emitted longwave radiation [$W * m^{-2}$], incoming longwave
575 radiation reaching the surface [$W * m^{-2}$], incoming solar radiation reaching the surface [$W * m^{-2}$], relative humidity [%],
576 runoff from base of snowpack [$m * timestep$], rain precipitation [m], snow density [$kg * m^{-3}$], snow-water equivalent melt
577 [m], snow depth [m], snow precipitation [m], static-surface sublimation [m], snow-water equivalent [m], air temperature
578 [$^{\circ}C$], wind direction [$^{\circ}$], and wind speed [$m * s^{-1}$]. Ultimately, we feel confident that Parallel SnowModel is producing the
579 same results as the original serial algorithm.

580 **Code, data availability, and supplement**

581 The Parallel SnowModel code and the data used in Sect. 4 is available through a public GitHub repository (Mower et al.,
582 2023). For more information about the serial version of SnowModel, refer to Liston and Elder (2006a). The data includes
583 figures and SnowModel output files that contain the necessary information to recreate the simulations. The gridded output
584 variables themselves are not included due to storage limitations. Pending approval, we will submit our code to get a DOI.

585 **Author contribution**

586 EDG and GDL conceived the study. RM, EDG, GDL, and SR were integral in the code development. RM, EDG, and JL
587 were involved in the design, execution, and interpretation of the experiments. All authors discussed the results and
588 contributed to the final version of the draft.

589 **Competing interests**

590 The contact author has declared that none of the authors has any competing interests.

591 **Disclaimer**

592 Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims in published maps and
593 institutional affiliations.

594 **Financial support**

595 This material is based upon work supported by the NSF National Center for Atmospheric Research, which is a major facility
596 sponsored by the U.S. National Science Foundation under Cooperative Agreement No. 1852977. The authors would like to
597 acknowledge that this work has been performed under funding from NASA Earth Science Office (ESTO) Advanced
598 Information Systems Technology (AIST) Program (grant no. 80NSSC20K0207), support by the University of Washington's

599 College of Engineering Fellowship, and computational support from NSF NCAR Computational and Information Systems
600 Lab (CISL) and NASA High-End Computing (HEC) Program through the NASA Center for Climate Simulation (NCCS) at
601 Goddard Space Flight Center.

602

603 **Acknowledgements**

604 We acknowledge Alessandro Fanfarillo in his help during the early stages of the Parallel SnowModel code development. We
605 are also grateful for the feedback from various team members involved in the AIST project, including Carrie Vuyovich,
606 Kristi Arsenault, Melissa Wrzesien, Adele Reinking, and Barton Forman.

607 **References**

608 Beniston, M.: Climatic Change in Mountain Regions: A Review of Possible Impacts, *Climatic Change*, 59, 5-31,
609 10.1023/A:1024458411589, 2003.

610 Bernhardt, M., Schulz, K., Liston, G. E., and Zängl, G.: The influence of lateral snow redistribution processes on snow melt
611 and sublimation in alpine regions, *Journal of Hydrology*, 424-425, 196-206, <https://doi.org/10.1016/j.jhydrol.2012.01.001>,
612 2012.

613 Boelman, N. T., Liston, G. E., Gurarie, E., Meddens, A. J. H., Mahoney, P. J., Kirchner, P. B., Bohrer, G., Brinkman, T. J.,
614 Cosgrove, C. L., Eitel, J. U. H., Hebblewhite, M., Kimball, J. S., LaPoint, S., Nolin, A. W., Pedersen, S. H., Prugh, L. R.,
615 Reinking, A. K., and Vierling, L. A.: Integrating snow science and wildlife ecology in Arctic-boreal North America,
616 *Environmental Research Letters*, 14, 010401, 10.1088/1748-9326/aaec1, 2019.

617 Center, N. O. H. R. S.: Snow data assimilation system (SNODAS) data products at NSIDC, 2004.

618 Clark, M. P. and Hay, L. E.: Use of Medium-Range Numerical Weather Prediction Model Output to Produce Forecasts of
619 Streamflow, *Journal of Hydrometeorology*, 5, 15-32, 10.1175/1525-7541(2004)005<0015:Uomnwp>2.0.Co;2, 2004.

620 Coarfa, C., Dotsenko, Y., Mellor-Crummey, J., Cantonnet, F., El-Ghazawi, T., Mohanti, A., Yao, Y., and Chavarria-
621 Miranda, D.: An evaluation of global address space languages: co-array fortran and unified parallel c, *Proceedings of the*
622 *tenth ACM SIGPLAN symposium on Principles and practice of parallel programming*, 36-47,

623 Dennis, J. M.: Inverse space-filling curve partitioning of a global ocean model, 2007 IEEE International Parallel and
624 Distributed Processing Symposium, 1-10,

625 Dozier, J., Bair, E. H., and Davis, R. E.: Estimating the spatial distribution of snow water equivalent in the world's
626 mountains, WIREs Water, 3, 461-474, <https://doi.org/10.1002/wat2.1140>, 2016.

627 Essery, R., Li, L., and Pomeroy, J.: A distributed model of blowing snow over complex terrain, Hydrological Processes, 13,
628 2423-2438, [https://doi.org/10.1002/\(SICI\)1099-1085\(199910\)13:14/15<2423::AID-HYP853>3.0.CO;2-U](https://doi.org/10.1002/(SICI)1099-1085(199910)13:14/15<2423::AID-HYP853>3.0.CO;2-U), 1999.

629 Fanfarillo, A., Burnus, T., Cardellini, V., Filippone, S., Nagle, D., and Rouson, D.: OpenCoarrays: open-source transport
630 layers supporting coarray Fortran compilers, Proceedings of the 8th International Conference on Partitioned Global Address
631 Space Programming Models, 1-11,

632 Fang, X. and Pomeroy, J.: Modeling blowing snow redistribution to prairie wetlands, Hydrological Processes, 23, 2557-
633 2569, 10.1002/hyp.7348, 2009.

634 Foster, J. L., Hall, D. K., Eylander, J. B., Riggs, G. A., Nghiem, S. V., Tedesco, M., Kim, E., Montesano, P. M., Kelly, R. E.
635 J., Casey, K. A., and Choudhury, B.: A blended global snow product using visible, passive microwave and scatterometer
636 satellite data, International Journal of Remote Sensing, 32, 1371-1395, 10.1080/01431160903548013, 2011.

637 Freudiger, D., Kohn, I., Seibert, J., Stahl, K., and Weiler, M.: Snow redistribution for the hydrological modeling of alpine
638 catchments, WIREs Water, 4, e1232, <https://doi.org/10.1002/wat2.1232>, 2017.

639 Gesch, D. B., Evans, G. A., Oimoen, M. J., and Arundel, S.: The National Elevation Dataset, in, edited by: United States
640 Geological Survey, E. R. O. a. S. E. C., American Society for Photogrammetry and Remote Sensing, 83-110, 2018.

641 Hall, D., Riggs, G., and Salomonson, V.: MODIS/Terra Snow Cover 5-Min L2 Swath 500m, Version, 5, 2011167.2011750,
642 2006.

643 Hamman, J. J., Nijssen, B., Bohn, T. J., Gergel, D. R., and Mao, Y.: The Variable Infiltration Capacity model version 5
644 (VIC-5): Infrastructure improvements for new applications and reproducibility, Geoscientific Model Development, 11, 3481-
645 3496, 2018.

646 Hammond, J. C., Sexstone, G. A., Putman, A. L., Barnhart, T. B., Rey, D. M., Driscoll, J. M., Liston, G. E., Rasmussen, K.
647 L., McGrath, D., Fassnacht, S. R., and Kampf, S. K.: High Resolution SnowModel Simulations Reveal Future Elevation-
648 Dependent Snow Loss and Earlier, Flashier Surface Water Input for the Upper Colorado River Basin, Earth's Future, 11,
649 e2022EF003092, <https://doi.org/10.1029/2022EF003092>, 2023.

650 Homer, C., Dewitz, J., Yang, L., Jin, S., Danielson, P., Xian, G., Coulston, J., Herold, N., Wickham, J., and Megown, K.:
651 Completion of the 2011 National Land Cover Database for the conterminous United States—representing a decade of land
652 cover change information, *Photogrammetric Engineering & Remote Sensing*, 81, 345-354, 2015.

653 Hoppinen, Z. M., Oveisgharan, S., Marshall, H.-P., Mower, R., Elder, K., and Vuyovich, C.: Snow Water Equivalent
654 Retrieval Over Idaho, Part B: Using L-band UAVSAR Repeat-Pass Interferometry, *The Cryosphere Discussions*, 2023, 1-24,
655 2023.

656 Huss, M., Bookhagen, B., Huggel, C., Jacobsen, D., Bradley, R. S., Clague, J. J., Vuille, M., Buytaert, W., Cayan, D. R.,
657 Greenwood, G., Mark, B. G., Milner, A. M., Weingartner, R., and Winder, M.: Toward mountains without permanent snow
658 and ice, *Earth's Future*, 5, 418-435, <https://doi.org/10.1002/2016EF000514>, 2017.

659 Immerzeel, W. W., Lutz, A. F., Andrade, M., Bahl, A., Biemans, H., Bolch, T., Hyde, S., Brumby, S., Davies, B. J., Elmore,
660 A. C., Emmer, A., Feng, M., Fernández, A., Haritashya, U., Kargel, J. S., Koppes, M., Kraaijenbrink, P. D. A., Kulkarni, A.
661 V., Mayewski, P. A., Nepal, S., Pacheco, P., Painter, T. H., Pellicciotti, F., Rajaram, H., Rupper, S., Sinisalo, A., Shrestha,
662 A. B., Viviroli, D., Wada, Y., Xiao, C., Yao, T., and Baillie, J. E. M.: Importance and vulnerability of the world's water
663 towers, *Nature*, 577, 364-369, 10.1038/s41586-019-1822-y, 2020.

664 ISO/IEC: Fortran Standard 2008; Technical report, Geneva, Switzerland, 2010.

665 Jin, S., Homer, C., Yang, L., Danielson, P., Dewitz, J., Li, C., Zhu, Z., Xian, G., and Howard, D.: Overall methodology
666 design for the United States national land cover database 2016 products, *Remote Sensing*, 11, 2971, 2019.

667 Keenan, E., Wever, N., Lenaerts, J. T. M., and Medley, B.: A wind-driven snow redistribution module for Alpine3D v3.3.0:
668 adaptations designed for downscaling ice sheet surface mass balance, *Geosci. Model Dev.*, 16, 3203-3219, 10.5194/gmd-16-
669 3203-2023, 2023.

670 Kumar, V. and Gupta, A.: Analysis of scalability of parallel algorithms and architectures: A survey, *Proceedings of the 5th*
671 *international conference on Supercomputing*, 396-405,

672 Laboratory, C. a. I. S.: Cheyenne, 10.5065/D6RX99HX, 2019.

673 Latifovic, R., Homer, C., Ressler, R., Pouliot, D., Hossain, S. N., Colditz, R. R., Olthof, I., Giri, C. P., and Victoria, A.: 20
674 North American Land-Change Monitoring System, *Remote sensing of land use and land cover*, 303, 2016.

675 Lehning, M., Löwe, H., Ryser, M., and Raderschall, N.: Inhomogeneous precipitation distribution and snow transport in
676 steep terrain, *Water Resources Research*, 44, 2008.

677 Lehning, M., Völkisch, I., Gustafsson, D., Nguyen, T., Stähli, M., and Zappa, M.: ALPINE3D: A detailed model of mountain
678 surface processes and its application to snow hydrology, *Hydrological Processes*, 20, 2111-2128, 10.1002/hyp.6204, 2006.

679 Lettenmaier, D. P., Alsdorf, D., Dozier, J., Huffman, G. J., Pan, M., and Wood, E. F.: Inroads of remote sensing into
680 hydrologic science during the WRR era, *Water Resources Research*, 51, 7309-7342,
681 <https://doi.org/10.1002/2015WR017616>, 2015.

682 Liston, G., Reinking, A. K., and Boleman, N.: Daily SnowModel Outputs Covering the ABoVE Core Domain, 3-km
683 Resolution, 1980-2020, 10.3334/ORNLDAAAC/2105, 2022.

684 Liston, G. E.: Local advection of momentum, heat, and moisture during the melt of patchy snow covers, *Journal of Applied*
685 *Meteorology and Climatology*, 34, 1705-1715, 1995.

686 Liston, G. E.: Representing Subgrid Snow Cover Heterogeneities in Regional and Global Models, *Journal of Climate*, 17,
687 1381-1397, 10.1175/1520-0442(2004)017<1381:Rsschi>2.0.Co;2, 2004.

688 Liston, G. E. and Elder, K.: A distributed snow-evolution modeling system (SnowModel), *Journal of Hydrometeorology*, 7,
689 1259-1276, 2006a.

690 Liston, G. E. and Elder, K.: A Meteorological Distribution System for High-Resolution Terrestrial Modeling (MicroMet),
691 *Journal of Hydrometeorology*, 7, 217-234, 10.1175/jhm486.1, 2006b.

692 Liston, G. E. and Hall, D. K.: An energy-balance model of lake-ice evolution, *Journal of Glaciology*, 41, 373-382, 1995.

693 Liston, G. E. and Hiemstra, C. A.: A simple data assimilation system for complex snow distributions (SnowAssim), *Journal*
694 *of Hydrometeorology*, 9, 989-1004, 2008.

695 Liston, G. E. and Hiemstra, C. A.: The changing cryosphere: Pan-Arctic snow trends (1979–2009), *Journal of Climate*, 24,
696 5691-5712, 2011.

697 Liston, G. E. and Mernild, S. H.: Greenland freshwater runoff. Part I: A runoff routing model for glaciated and nonglaciated
698 landscapes (HydroFlow), *Journal of Climate*, 25, 5997-6014, 2012.

699 Liston, G. E. and Sturm, M.: A snow-transport model for complex terrain, *Journal of Glaciology*, 44, 498 - 516, 1998.

700 Liston, G. E. and Sturm, M.: Global Seasonal-Snow Classification, Version 1 [dataset], 2021.

701 Liston, G. E., Perham, C. J., Shideler, R. T., and Chevront, A. N.: Modeling snowdrift habitat for polar bear dens,
702 Ecological Modelling, 320, 114-134, <https://doi.org/10.1016/j.ecolmodel.2015.09.010>, 2016.

703 Liston, G. E., Winther, J.-G., Bruland, O., Elvehøy, H., and Sand, K.: Below-surface ice melt on the coastal Antarctic ice
704 sheet, Journal of Glaciology, 45, 273-285, 1999.

705 Liston, G. E., Haehnel, R. B., Sturm, M., Hiemstra, C. A., Berezovskaya, S., and Tabler, R. D.: Simulating complex snow
706 distributions in windy environments using SnowTran-3D, Journal of Glaciology, 53, 241-256, 2007.

707 Liston, G. E., Polashenski, C., Rösel, A., Itkin, P., King, J., Merkouriadi, I., and Haapala, J.: A distributed snow-evolution
708 model for sea-ice applications (SnowModel), Journal of Geophysical Research: Oceans, 123, 3786-3810, 2018.

709 Liston, G. E., Itkin, P., Stroeve, J., Tschudi, M., Stewart, J. S., Pedersen, S. H., Reinking, A. K., and Elder, K.: A Lagrangian
710 snow-evolution system for sea-ice applications (SnowModel-LG): Part I—Model description, Journal of Geophysical
711 Research: Oceans, 125, e2019JC015913, 2020.

712 Lund, J., Forster, R. R., Deeb, E. J., Liston, G. E., Skiles, S. M., and Marshall, H.-P.: Interpreting Sentinel-1 SAR
713 Backscatter Signals of Snowpack Surface Melt/Freeze, Warming, and Ripening, through Field Measurements and
714 Physically-Based SnowModel, Remote Sensing, 14, 4002, 2022.

715 Mahoney, P. J., Liston, G. E., LaPoint, S., Gurarie, E., Mangipane, B., Wells, A. G., Brinkman, T. J., Eitel, J. U.,
716 Hebblewhite, M., and Nolin, A. W.: Navigating snowscapes: scale-dependent responses of mountain sheep to snowpack
717 properties, Ecological Applications, 28, 1715-1729, 2018.

718 Marsh, C. B., Pomeroy, J. W., Spiteri, R. J., and Wheeler, H. S.: A Finite Volume Blowing Snow Model for Use With
719 Variable Resolution Meshes, Water Resources Research, 56, e2019WR025307, <https://doi.org/10.1029/2019WR025307>,
720 2020.

721 Miller, P., Robson, M., El-Masri, B., Barman, R., Zheng, G., Jain, A., and Kalé, L.: Scaling the isam land surface model
722 through parallelization of inter-component data transfer, 2014 43rd International Conference on Parallel Processing, 422-
723 431,

724 Mitchell, K. E.: The multi-institution North American Land Data Assimilation System (NLDAS): Utilizing multiple GCIP
725 products and partners in a continental distributed hydrological modeling system, J. Geophys. Res., 109, D07S90, 2004.

726 Moody, E. G., King, M. D., Platnick, S., Schaaf, C. B., and Gao, F.: Spatially complete global spectral surface albedos:
727 Value-added datasets derived from Terra MODIS land products, *IEEE Transactions on Geoscience and Remote Sensing*, 43,
728 144-158, 2005.

729 Morin, S., Horton, S., Techel, F., Bavay, M., Coléou, C., Fierz, C., Gobiet, A., Hagenmuller, P., Lafaysse, M., Ližar, M.,
730 Mitterer, C., Monti, F., Müller, K., Olefs, M., Snook, J. S., van Herwijnen, A., and Vionnet, V.: Application of physical
731 snowpack models in support of operational avalanche hazard forecasting: A status report on current implementations and
732 prospects for the future, *Cold Regions Science and Technology*, 170, 102910,
733 <https://doi.org/10.1016/j.coldregions.2019.102910>, 2020.

734 Mortezapour, M., Menounos, B., Jackson, P. L., Erler, A. R., and Pelto, B. M.: The role of meteorological forcing and snow
735 model complexity in winter glacier mass balance estimation, Columbia River basin, Canada, *Hydrological Processes*, 34,
736 5085-5103, <https://doi.org/10.1002/hyp.13929>, 2020.

737 Mott, R. and Lehning, M.: Meteorological Modeling of Very High-Resolution Wind Fields and Snow Deposition for
738 Mountains, *Journal of Hydrometeorology*, 11, 934-949, <https://doi.org/10.1175/2010JHM1216.1>, 2010.

739 Mott, R., Schirmer, M., Bavay, M., Grünewald, T., and Lehning, M.: Understanding snow-transport processes shaping the
740 mountain snow-cover, *The Cryosphere*, 4, 545-559, 10.5194/tc-4-545-2010, 2010.

741 Mower, R., Gutmann, E. D., and Liston, G. E.: Parallel-SnowModel 1.0 [code], [https://github.com/NCAR/Parallel-](https://github.com/NCAR/Parallel-SnowModel-1.0)
742 [SnowModel-1.0](https://github.com/NCAR/Parallel-SnowModel-1.0), 2023.

743 Mudryk, L. R., Derksen, C., Kushner, P. J., and Brown, R.: Characterization of Northern Hemisphere Snow Water
744 Equivalent Datasets, 1981–2010, *Journal of Climate*, 28, 8037-8051, 10.1175/jcli-d-15-0229.1, 2015.

745 Nolin, A. W.: Recent advances in remote sensing of seasonal snow, *Journal of Glaciology*, 56, 1141-1150,
746 10.3189/002214311796406077, 2010.

747 Numrich, R. W. and Reid, J.: Co-Array Fortran for parallel programming, *ACM Sigplan Fortran Forum*, 1-31,

748 Numrich, R. W., Steidel, J. L., Johnson, B. H., Dinechin, B. D. d., Elsesser, G., Fischer, G., and MacDonald, T.: Definition
749 of the F— Extension to Fortran 90, *International Workshop on Languages and Compilers for Parallel Computing*, 292-306,

750 Painter, T. H., Berisford, D. F., Boardman, J. W., Bormann, K. J., Deems, J. S., Gehrke, F., Hedrick, A., Joyce, M., Laidlaw,
751 R., and Marks, D.: The Airborne Snow Observatory: Fusion of scanning lidar, imaging spectrometer, and physically-based
752 modeling for mapping snow water equivalent and snow albedo, *Remote Sensing of Environment*, 184, 139-152, 2016.

753 Parhami, B.: SIMD machines: do they have a significant future?, *ACM SIGARCH Computer Architecture News*, 23, 19-22,
754 1995.

755 Pedersen, S. H., Liston, G. E., Tamstorf, M. P., Schmidt, N. M., and Abermann, J.: Linking vegetation greenness and
756 seasonal snow characteristics using field observations, SnowModel, and daily MODIS imagery in high-Arctic Greenland,
757 *AGU Fall Meeting Abstracts*, GC42A-07,

758 Pedersen, S. H., Liston, G. E., Tamstorf, M. P., Westergaard-Nielsen, A., and Schmidt, N. M.: Quantifying Episodic
759 Snowmelt Events in Arctic Ecosystems, *Ecosystems*, 18, 839-856, 10.1007/s10021-015-9867-8, 2015.

760 Perezhugin, P., Chernov, I., and Iakovlev, N.: Advanced parallel implementation of the coupled ocean–ice model FEMAO
761 (version 2.0) with load balancing, *Geoscientific Model Development*, 14, 843-857, 2021.

762 Pflug, J. M. and Lundquist, J. D.: Inferring Distributed Snow Depth by Leveraging Snow Pattern Repeatability: Investigation
763 Using 47 Lidar Observations in the Tuolumne Watershed, Sierra Nevada, California, *Water Resources Research*, 56,
764 e2020WR027243, <https://doi.org/10.1029/2020WR027243>, 2020.

765 Prokop, A. and Schneiderbauer, S.: The atmospheric snow-transport model: SnowDrift3D, *Journal of Glaciology*, 57, 526-
766 542, 10.3189/002214311796905677, 2011.

767 Quéno, L., Mott, R., Morin, P., Cluzet, B., Mazzotti, G., and Jonas, T.: Snow redistribution in an intermediate-complexity
768 snow hydrology modelling framework, *EGUsphere*, 2023, 1-32, 10.5194/egusphere-2023-2071, 2023.

769 Randin, C. F., Dedieu, J.-P., Zappa, M., Long, L., and Dullinger, S.: Validation of and comparison between a semidistributed
770 rainfall–runoff hydrological model (PREVAH) and a spatially distributed snow-evolution model (SnowModel) for snow
771 cover prediction in mountain ecosystems, *Ecohydrology*, 8, 1181-1193, <https://doi.org/10.1002/eco.1570>, 2015.

772 Rasmussen, R. M., Liu, C., Ikeda, K., Chen, F., Kim, J.-H., Schneider, T., Gochis, D., Dugger, A., and Viger, R.: Four-
773 kilometer long-term regional hydroclimate reanalysis over the conterminous United States (CONUS), 1979-2020, *Research*
774 *Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory*
775 [dataset], 10.5065/ZYY0-Y036, 2023.

776 Renwick, J.: MOUNTerrain: GEWEX mountainous terrain precipitation project, *GEWEX news*, 24, 5-6, 2014.

777 Reynolds, D. S., Pflug, J. M., and Lundquist, J. D.: Evaluating wind fields for use in basin-scale distributed snow models,
778 *Water Resources Research*, 57, e2020WR028536, 2021.

779 Richter, B., Schweizer, J., Rotach, M. W., and van Herwijnen, A.: Modeling spatially distributed snow instability at a
780 regional scale using Alpine3D, *Journal of Glaciology*, 67, 1147-1162, 10.1017/jog.2021.61, 2021.

781 Rouson, D., Gutmann, E. D., Fanfarillo, A., and Friesen, B.: Performance portability of an intermediate-complexity
782 atmospheric research model in coarray Fortran, *Proceedings of the Second Annual PGAS Applications Workshop*, 1-4,

783 Skofronick-Jackson, G. M., Johnson, B. T., and Munchak, S. J.: Detection Thresholds of Falling Snow From Satellite-Borne
784 Active and Passive Sensors, *IEEE Transactions on Geoscience and Remote Sensing*, 51, 4177-4189,
785 10.1109/TGRS.2012.2227763, 2013.

786 Sturm, M. and Liston, G. E.: Revisiting the global seasonal snow classification: An updated dataset for earth system
787 applications, *Journal of Hydrometeorology*, 22, 2917-2938, 2021.

788 Sturm, M., Holmgren, J., and Liston, G. E.: A Seasonal Snow Cover Classification System for Local to Global Applications,
789 *Journal of Climate*, 8, 1261-1283, 10.1175/1520-0442(1995)008<1261:Assccs>2.0.Co;2, 1995.

790 Szczypta, C., Gascoin, S., Houet, T., and Fanise, P.: Impact of climate versus land-use changes on snow cover in Bassiès,
791 Pyrenees, *International Snow Science Workshop Grenoble 2011 Chamonix Mont-Blanc*, 1278-1281,

792 Tabler, R. D.: Estimating the transport and evaporation of blowing snow, *Great Plains Agric Counc Publ*, 1975.

793 Takala, M., Luojus, K., Pulliainen, J., Derksen, C., Lemmetyinen, J., Kärnä, J.-P., Koskinen, J., and Bojkov, B.: Estimating
794 northern hemisphere snow water equivalent for climate research through assimilation of space-borne radiometer data and
795 ground-based measurements, *Remote Sensing of Environment*, 115, 3517-3529, <https://doi.org/10.1016/j.rse.2011.08.014>,
796 2011.

797 Tedesco, M. and Jeyaratnam, J.: AMSR-E/AMSR2 Unified L3 Global Daily 25 km EASE-Grid Snow Water Equivalent,
798 Version 1, Boulder, Colorado USA, NASA National Snow and Ice Data Center Distributed Active Archive Center, 2019.

799 Vionnet, V., Martin, E., Masson, V., Lac, C., Naaim Bouvet, F., and Guyomarc'h, G.: High-resolution large eddy simulation
800 of snow accumulation in Alpine terrain, *Journal of Geophysical Research: Atmospheres*, 122, 11,005-011,021, 2017.

801 Vionnet, V., Martin, E., Masson, V., Guyomarc'h, G., Naaim-Bouvet, F., Prokop, A., Durand, Y., and Lac, C.: Simulation of
802 wind-induced snow transport and sublimation in alpine terrain using a fully coupled snowpack/atmosphere model, *The*
803 *Cryosphere*, 8, 395-415, 2014.

804 Voordendag, A., Réveillet, M., MacDonell, S., and Lhermitte, S.: Snow model comparison to simulate snow depth evolution
805 and sublimation at point scale in the semi-arid Andes of Chile, *The Cryosphere*, 15, 4241-4259, 2021.

806 Vuyovich, C. M., Jacobs, J. M., and Daly, S. F.: Comparison of passive microwave and modeled estimates of total watershed
807 SWE in the continental United States, *Water Resources Research*, 50, 9088-9102, <https://doi.org/10.1002/2013WR014734>,
808 2014.

809 Wagner, C., Hunsaker, A., and Jacobs, J.: UAV and SnowModel Estimates of Wind Driven Snow in Eastern USA
810 Avalanche Terrain, Copernicus Meetings, 2023.

811 Wrzesien, M. L., Durand, M. T., Pavelsky, T. M., Kapnick, S. B., Zhang, Y., Guo, J., and Shum, C. K.: A New Estimate of
812 North American Mountain Snow Accumulation From Regional Climate Model Simulations, *Geophysical Research Letters*,
813 45, 1423-1432, <https://doi.org/10.1002/2017GL076664>, 2018.

814 Xia, Y.: Continental-scale water and energy flux analysis and validation for North American Land Data Assimilation System
815 project phase 2 (NLDAS-2): 1. Intercomparison and application of model products, *J. Geophys. Res.*, 117, D03109, 2012a.

816 Xia, Y.: Continental-scale water and energy flux analysis and validation for North American Land Data Assimilation System
817 project phase 2 (NLDAS-2): 2. Validation of model-simulated streamflow, *J. Geophys. Res.*, 117, D03110, 2012b.

818 Xue, M., Droegemeier, K. K., and Wong, V.: The Advanced Regional Prediction System (ARPS) – A multi-scale
819 nonhydrostatic atmospheric simulation and prediction model. Part I: Model dynamics and verification, *Meteorology and*
820 *Atmospheric Physics*, 75, 161-193, [10.1007/s007030070003](https://doi.org/10.1007/s007030070003), 2000.

821