

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 higher
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 and greatly influences snow heterogeneity, sublimation,
53 avalanches, and melt timing. Models that have incorporated wind-induced physics generally require components to both
54 develop the snow mass balance and incorporate atmospheric inputs of the wind field. However, there often exists a trade-off
55 between the accuracy of simulating wind-induced snow transport and the computational requirements for downscaling and
56 developing the wind fields over the gridded domain (Reynolds et al., 2021; Vionnet et al., 2014). Therefore, simplifying
57 assumptions of uniform wind direction has been applied in models like Distributed Blowing Snow Model (DBSM) (Essery
58 et al., 1999; Fang and Pomeroy, 2009). More advanced models have utilized advection-diffusion equations, like Alpine3D
59 (Lehning et al., 2006) or spatial distributed formulations like SnowTran-3D (Liston and Sturm, 1998). Finite volume
60 methods for more efficiently discretizing wind fields have been applied to models such as DBSM (Marsh et al., 2020). The
61 most complex models consider nonsteady turbulence which utilize three-dimensional wind fields from atmospheric models
62 to simulate blowing snow transport and sublimation; for example, SURFEX in Meso-NH/Crocus (Vionnet et al., 2014;
63 Vionnet et al., 2017), wind fields from the atmospheric model ARPS (Xue et al., 2000) being incorporated into Alpine3D
64 (Mott and Lehning, 2010; Mott et al., 2010; Lehning et al., 2008), and SnowDrift3D (Prokop and Schneiderbauer, 2011).

65 Incorporating wind-induced physics into snow models is computationally expensive; thus, parallelizing the serial algorithms
66 would likely be beneficial to many models.

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

79 Recent advancements in multiprocessor computer technologies and architectures have allowed for increased performance in
80 simulating complex natural systems at high resolutions. Parallel computing has been used on many LSMs to reduce compute
81 time and allow for higher accuracy results from finer grid simulations (Hamman et al., 2018; Miller et al., 2014; Sharma et
82 al., 2004). Our goal was to develop a parallel version of SnowModel (Parallel SnowModel) using Coarray Fortran (CAF)
83 syntax without making significant changes to the original SnowModel code physics or structure. CAF is a Partitioned Global
84 Address Space (PGAS) programming model and has been used to run atmospheric models on 100,000 cores (Rouson et al.,
85 2017).

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

93 While much of the SnowModel's logic can be considered "embarrassingly parallel", SnowModel also contains "non-trivial"
94 algorithms within the solar radiation, wind, and snow redistribution models. Calculations within these algorithms often
95 require information from neighboring grid cells, either for spatial derivative calculations or for horizontal fluxes of mass
96 (e.g., saltating or turbulent-suspended snow) across the domain. Therefore, non-trivial parallelization requires implementing
97 algorithm changes that allow computer processes to communicate and exchange data. The novelty of the work presented
98 here includes 1) the presentation of Parallel SnowModel, high-resolution (100 m) distributed snow datasets over CONUS,

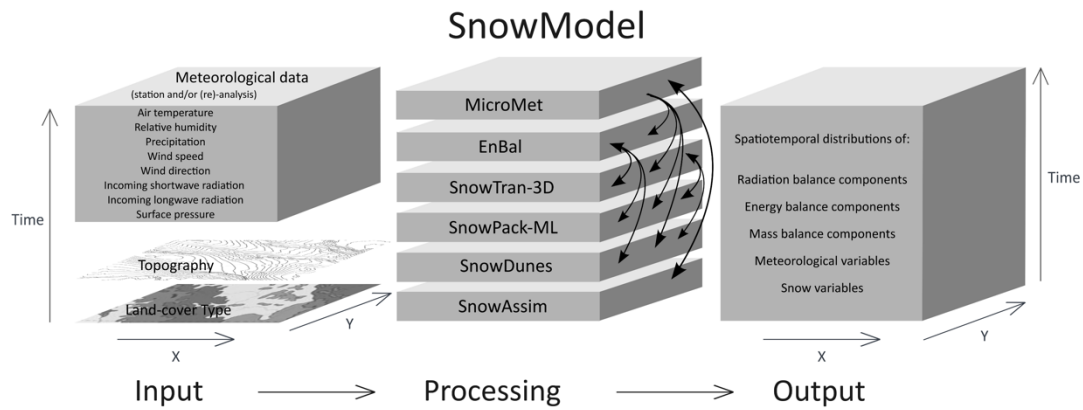
99 and an analysis of the performance of the parallel algorithm; 2) demonstrating how a simplified parallelization approach
100 using CAF and one-dimensional decomposition can be implemented in geoscientific algorithms to scale over large domains;
101 and 3) demonstrating an approach for non-trivial parallelization algorithms that involve spatial derivatives and fluxes using
102 halo-exchange techniques.

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

108 **2 Background**

109 **2.1 SnowModel**

110 SnowModel is a spatially distributed snow-evolution modeling system designed to model snow states (e.g., snow depth,
111 SWE, snow melt, snow density) and fluxes over different landscapes and climates (Liston and Elder, 2006a). The most
112 complete and up-to-date description of SnowModel can be found in the Appendices of Liston et al. (2020). While many
113 snow modelling systems exist, SnowModel will benefit from parallelization because of its ability to simulate snow processes
114 on a high-resolution grid through downscaling meteorological inputs and modelling snow redistribution. SnowModel is
115 designed to simulate domains on a structured grid with spatial resolutions ranging from 1 to 200 m (although it can simulate
116 coarser resolutions, as well) and temporal resolutions ranging from 10 m to 1 d. The primary modeled processes include
117 accumulation from frozen precipitation; blowing-snow redistribution and sublimation; interception, unloading, and
118 sublimation within forest canopies; snow-density and grain-size evolution; and snowpack ripening and melt. These processes
119 are distributed into four, core interacting submodules: MicroMet defines the meteorological forcing conditions (Liston and
120 Elder, 2006b), EnBal describes surface and energy exchanges (Liston, 1995; Liston et al., 1999), SnowPack-ML is a
121 multilayer snowpack sub-model that simulates the evolution of snow properties and the moisture and energy transfers
122 between layers (Liston and Hall, 1995; Liston and Mernild, 2012), and SnowTran-3D calculates snow redistribution by wind
123 (Liston et al., 2007). Additional simulation features include SnowDunes (Liston et al., 2018) and SnowAssim (Liston and
124 Hiemstra, 2008), which model sea-ice applications and data assimilation techniques, respectively. Figure 1 shows a
125 schematic of the core SnowModel toolkit. Additionally, the initialization submodules that read in the model parameters,
126 distribute inputs across the modeled grid, allocate arrays, etc., include PreProcess and ReadParam. Outputting arrays is
127 contained within the Outputs submodule. SnowModel incorporates first-order physics required to simulate snow evolution
128 within each of the global snow classes [e.g., Ice, Tundra, Boreal Forest, Montane Forest, Prairie, Maritime, and Ephemeral;
129 (Sturm and Liston, 2021; Liston and Sturm, 2021)].



130

131 **Figure 1: Schematic modified by Pederson et al. (2015) providing an example of possible inputs, core submodules, and outputs of**
 132 **SnowModel.**

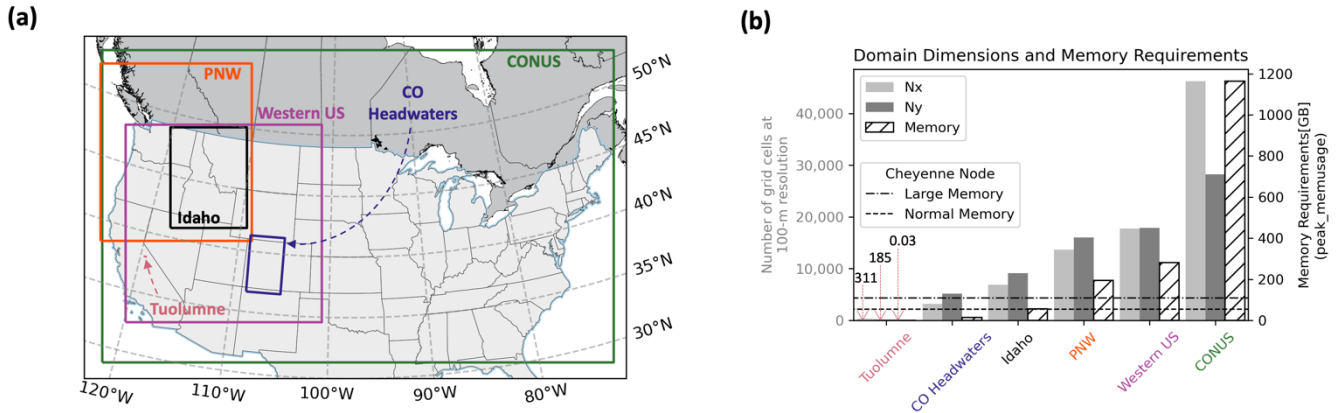
133 2.2 Coarray Fortran

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

144 2.3 Model Domains, Data, and Computing Resources

145 The required inputs for SnowModel include 1) temporally varying meteorological variables of precipitation, wind speed and
 146 direction, air temperature, and relative humidity taken from meteorological stations or atmospheric models and 2) spatially
 147 distributed topography and land-cover type (Liston & Elder, 2006a). The following inputs were used for the experiments
 148 introduced in Sect. 3: USGS National Elevation Dataset (NED) for topography (Gesch et al., 2018), The North American
 149 Land Change Monitoring System (NALCMS) Land Cover 2015 map for vegetation (Homer et al., 2015; Jin et al., 2019;
 150 Latifovic et al., 2016), and forcing variables from either the North American Land Data Assimilation System (NLDAS-2)
 151 (Mitchell, 2004; Xia, 2012a, b) on a 1/8 degree (approximately 12 km) grid or a high-resolution Weather Research Forecast
 152 (WRF) model from the National Center for Atmospheric Research (NCAR) on approximately a 4 km grid (Rasmussen et al.,
 153 2023). The high-performance computing architectures used include NCAR’s Cheyenne supercomputer, which is a 5.43-

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



162
 163 **Figure 2:** (a) *Spatial location of simulated domains on WRF’s lambert conformal projection (Rasmussen et al., 2023) and (b)*
 164 *corresponding grid dimensions (N_x – number of grids in x dimension; N_y – number of grids in y dimension) and memory obtained*
 165 *from peak_memusage package required for single-layer SnowModel simulation experiments. For reference, the dashed lines represent*
 166 *the normal and large memory thresholds (55 and 109 GB) for Cheyenne’s SGI ICE XA cluster.*

167 2.4 Parallelization Motivation

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

180 currently available, state of the art, single-processor workstation, would require approximately 120 d of computer time to
181 perform a 1 y model simulation over the CONUS domain. SnowModel is regularly used to perform multi-decade
182 simulations, for trend analyses, climate change studies, and retrospective analyses (Liston and Hiemstra, 2011; Liston et al.,
183 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
184 would take approximately 4800 d, or over 13 y, of computer time. Clearly such simulations are not practical using single-
185 processor computer hardware and software algorithms.

186 **3 Methods**


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

192 **3.1 Parallel Implementation**

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

195 **3.1.1 Partitioning Algorithm**

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

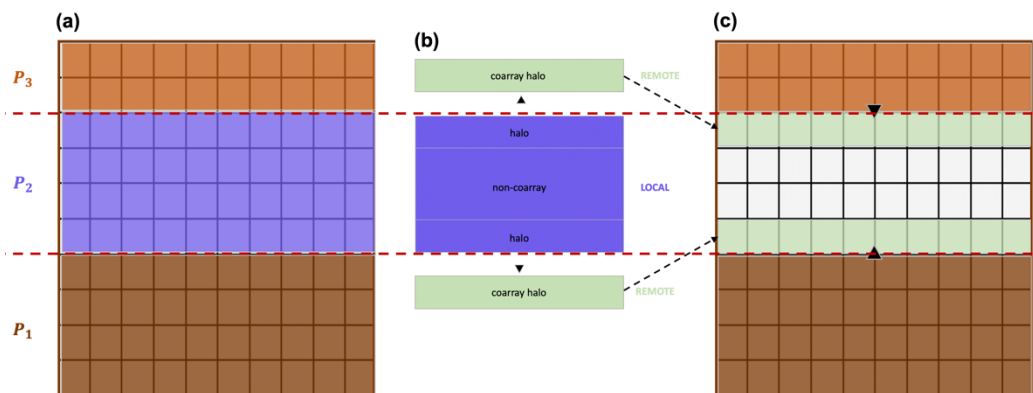


205

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

207 **3.1.2 Non-trivial Parallelization**

208 Each process has sufficient information to correctly execute most of the physical computations within SnowModel.
 209 However, there are certain subroutines where grid computations require information from neighboring grid cells (e.g., data
 210 dependencies) and therefore information outside of the local domain of a process. For SnowModel, these subroutines
 211 typically involve the transfer of blowing snow or calculations requiring spatial derivatives. Furthermore, with our one-
 212 dimensional decomposition approach, each grid cell within a process local domain has sufficient information from its
 213 neighboring grid cells in the x dimension but potentially lacks information from neighboring grid cells in the y dimension. As
 214 a regular grid method, SnowModel lends itself to process communication via halo-exchange where coarrays are used in
 215 remote calls. Halo-exchange using CAF involves copying boundary data into coarrays on neighboring processes and using
 216 information from the coarrays to complete computations (Figure 4). Although the entire local array could be declared a
 217 coarray and accessed by remote processes more directly, some CAF implementations, (e.g. Cray) impose additional
 218 constraints upon coarray memory allocations that can be problematic for such large allocations.

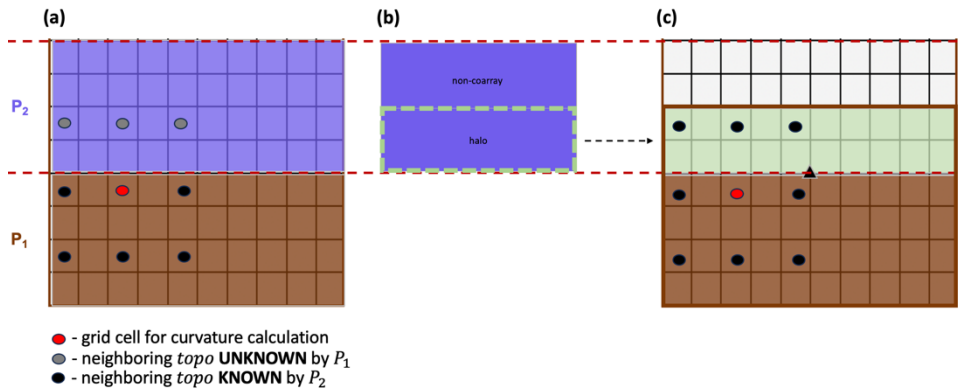


219

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

223 **3.1.2.1 Topography – Wind and Solar Radiation Models**

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

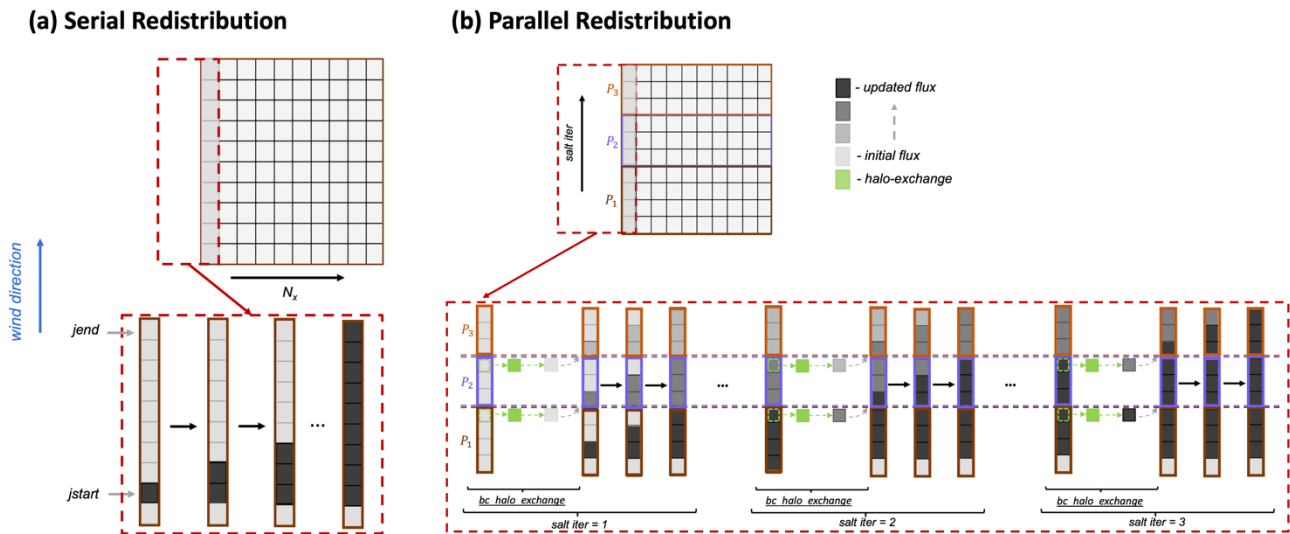


240
 241 **Figure 5: Schematic for halo-exchange used in the curvature calculation by P_1 , where $inc = 2$. (a) Prior to halo-exchange, P_1**
 242 **contains insufficient information to perform the curvature calculation, (b) grid cells (halo) within the local domain of P_2 are (c)**
 243 **transferred to P_1 via coarrays. At this point, P_1 has sufficient information to make the curvature calculation.**

244 **3.1.2.2 Snow Redistribution**

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

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



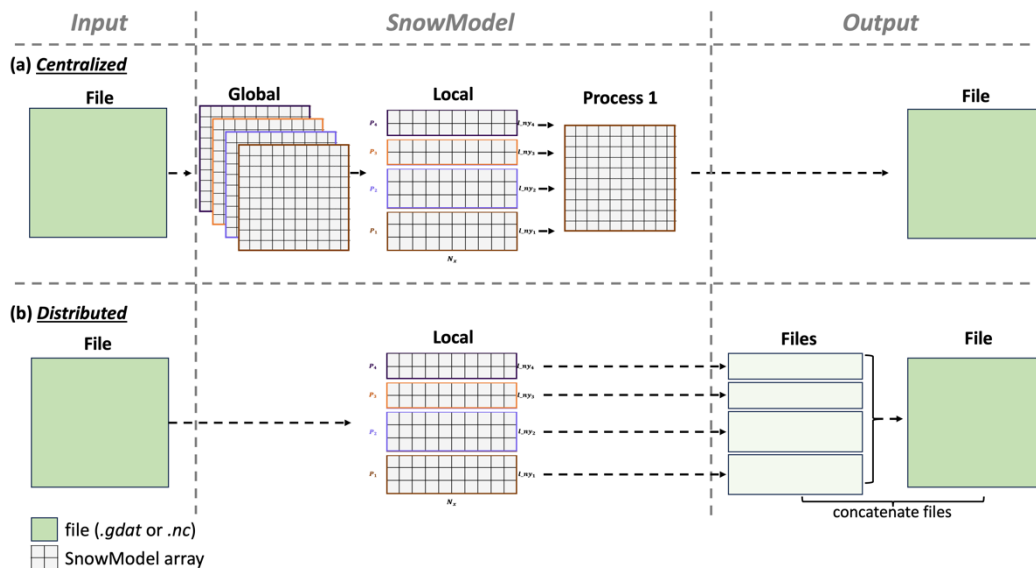
263

264 **Figure 6: (a) Schematic of the serial and (b) parallel redistribution algorithm showing the change in saltation flux due to southerly**
 265 **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**
 266 **iteration (*salt iter*) to perform a halo-exchange and update the boundary condition of the saltation flux.**

267 3.1.3 File I/O

268 File I/O management can be a significant bottleneck in parallel applications. Parallel implementations that are less memory
 269 restricted commonly use local to global mapping strategies, or a **Centralized** approach for file I/O (Figure 7a). This approach
 270 requires that one or more processes stores global arrays for input variables and that one process (Process 1; Figure 7a) stores
 271 global arrays for all output variables. As the domain size increases, the mapping of local variables to global variables for

272 outputting creates a substantial bottleneck. To improve performance, *Distributed* file I/O can be implemented, where input
 273 and output files are directly and concurrently accessed by each process (Figure 7b).



274

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

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

290 3.1.3.1 Parallel Inputs

291 As noted above, SnowModel has two primary types of input files, temporally static files such as vegetation and topography
 292 and transient inputs such as meteorological forcing data. While acceptable static input file types include flat binary, NetCDF,
 293 and ASCII files for the serial version of the code, optimizing the efficiency of Parallel SnowModel requires static inputs

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

311 **3.1.3.2 Parallel Outputs**

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

318 **3.2 Simulation Experiments**

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

324 Validation experiments comparing output from the original serial version of the code to the parallel version were conducted
325 continuously throughout the parallel algorithm development to assess the reproducibility of the results. Additionally, a more
326 thorough validation effort was performed at the end of the study that compared output from the serial algorithm to that of the

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

341 **3.2.1 Parallel Performance**

342 In high performance computing, scalability attempts to assess the effectiveness of running a parallel algorithm with an
343 increasing number of processes. Thus, scalability can be used to identify the optimal number of processes for a fixed domain,
344 understand the limitations of a parallel algorithm as a function of domain size and number of processes, and estimate the
345 efficiency of the parallel algorithm on new domains or computing architectures. Speedup, efficiency, and code profiling
346 were tools used to assess the scalability and performance of Parallel SnowModel on fixed domains. Speedup (S ; Eq. 1), a
347 metric of strong scaling, is defined as the ratio of the serial execution time, $T(1)$, over the execution time using N processes,
348 $T(N)$. Optimally, parallel algorithms will experience a doubling of speedup as the number of processes is doubled. Some
349 reasons why parallel algorithms do not follow ideal scaling include the degree of concurrency possible and overhead costs
350 due to communication. Synchronization statements have an associated cost of decreasing the speed and efficiency of an
351 algorithm due to communication overhead and requirements for one process to sit idle while waiting for another to reach the
352 synchronization point. Furthermore, speedup tends to peak or plateau at a certain limit on a given computing architecture and
353 domain because either the overheads grow with an increasing number of processes, or the number of processes exceeds the
354 degree of concurrency inherent in the algorithm (Kumar and Gupta, 1991). For large domains, where serial simulations
355 cannot be performed either due to wall-clock or memory limitations, relative speedup, (\hat{S} ; Eq. 2), is commonly used.
356 Relative speedup is estimated as a ratio of the execution time, $T(\hat{P})$, of the minimum number of processes, (\hat{P}), that can be
357 simulated on a given domain over $T(N)$. An additional speedup metric, approximate speedup (\check{S} ; Eq. 3), is introduced to
358 estimate S by assuming perfect scaling from \hat{P} to a single process. While this is only an approximation, it is helpful to

359 compare the \tilde{S} across the different domains on a similar scale. Additionally, efficiency (E ; Eq. 4), and approximate
360 efficiency (\tilde{E} ; Eq. 5) are the ratios of S to N and \tilde{S} to N , respectively. A simulation that demonstrates ideal scaling, would
361 have 100% efficiency. Additionally, code profiling evaluates the cumulative execution time of individual submodules (e.g.
362 Preprocess, Readparam, MicroMet, Enbal, SnowPack, SnowTran-3D, and Output) as a function of the number of processes.
363 Together, code profiling and strong scaling can be used to understand locations of bottlenecks in the algorithm and how
364 changes to the code enhance performance.

365

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

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

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

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

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

366 3.2.1.1 Parallel Improvement

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

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

390 **3.2.1.2 Strong Scaling**

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

397 **3.2.2 CONUS Simulations**

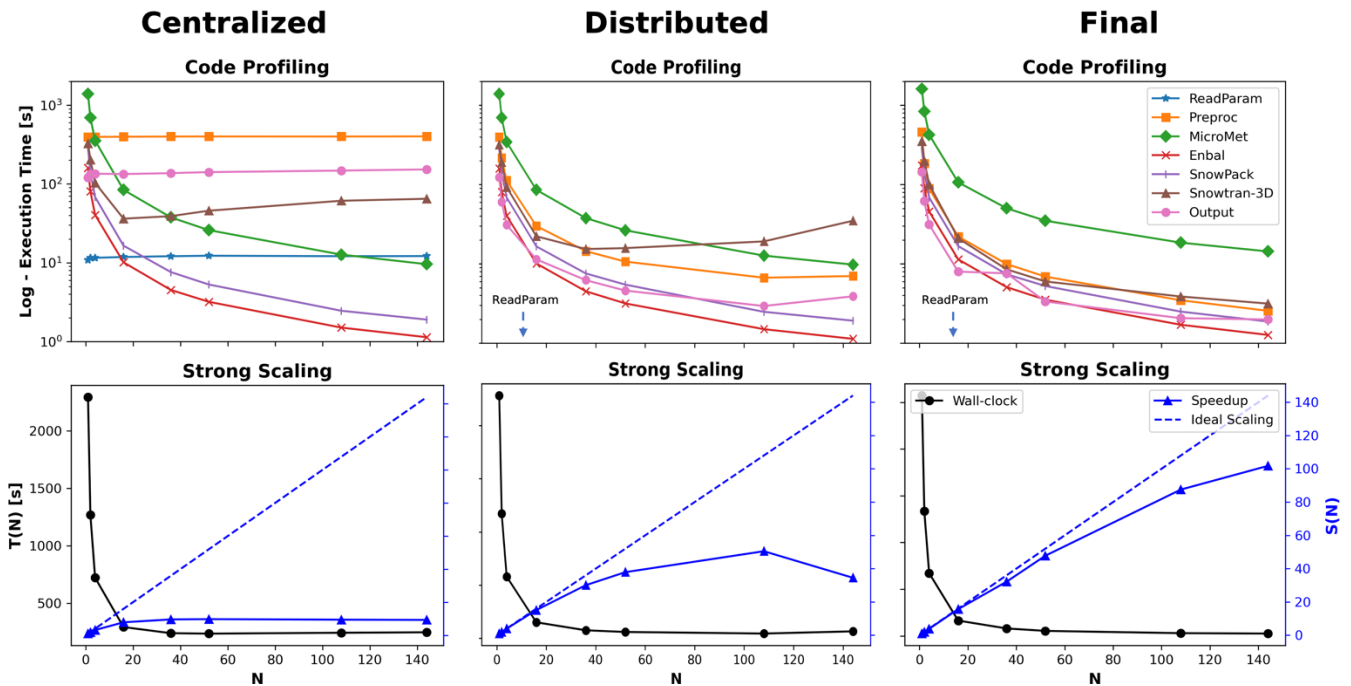
398 A primary goal of this work was to run Parallel SnowModel simulations for 21 years (2000 – 2021) over the CONUS
399 domain (Figure 2) on a 100 m grid, while resolving the diurnal cycle in the model physics and creating a daily dataset of
400 snow properties, including snow depth, SWE, melt rate,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

405 **4 Results**

406 **4.1 Parallel Improvement**

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

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




429
430 **Figure 8: Code profiling (top row) and strong scaling (bottom row) results demonstrating the progression of Parallel SnowModel,**
431 **which includes a version of the code with centralized file I/O (*Centralized*; first column), a version of the code with distributed file**
432 **I/O (*Distributed*; second column), and a final version of the code at the time of this publication (*Final*; third column). These**
433 **versions can be found as different commits within the GitHub repository (Mower et al., 2023). The code profiling plots display the**

434 cumulative execution time of each submodule on a logarithmic scale as a function of the number of processes (N). The arrow in the
435 code profiling plots of *Distributed* and *Final* indicates the ReadParam timing is below the y-axis at approximately 0.3 seconds and
436 0.003 seconds, respectively. The strong scaling plots show the total execution time ($T(N)$) against N on the primary y-axis and the
437 speedup (S) against N on the secondary y-axis.

438 4.2 Strong Scaling

439 In addition to the parallel improvement analysis, strong scaling was also performed on six domains for the 2018 water year
440 to better understand how Parallel SnowModel scales across different domain sizes and decompositions. Figure 9 displays the
441 approximate speedup (\tilde{S} ; Eq. 3) of Parallel SnowModel for three local/state domains (Tuolumne, CO Headwaters, and Idaho)
442 and three regional/continental domains (PNW, Western US, and CONUS). Additionally, Table 1 contains information about
443 the minimum and maximum number of processors (\hat{P} and P^* , respectively) simulated on each domain and their
444 corresponding execution time, relative speedup (\hat{S} ; Eq. 2), approximate speedup (\tilde{S} ; Eq. 3), and approximate efficiency (\tilde{E} ;
445 Eq. 5). As mentioned previously, simulations were constrained by both the 12 h wall-clock and 109 GB of memory per node
446 on the Cheyenne supercomputer. In strong scaling, the number of processes is increased while the problem size remains
447 constant; therefore, it represents a reduced workload per process. Local-sized domains, e.g., Tuolumne, likely do not warrant
448 the need for parallel resources because they have small serial runtimes (e.g., using 52 processes, Tuolumne had an \tilde{E} of 38%;
449 Table 1). However, state, regional, and continental domains stand to benefit more significantly from parallelization. The
450 CONUS runtime decreased by a factor of 2.6 running on 3456 processes relative to 648 processes. Based on our approximate
451 speedup assumption, we would estimate a CONUS \tilde{S} of 1690 times on 3456 processes compared to one process, with an \tilde{E} of
452 49%. The Western US and PNW domains display very similar scalability results (Figure 9), which is attributed to the similar
453 number of grid cells in the y dimension (Figure 2; and Table 1) and thus parallel decomposition for each domain.
454 Furthermore, these domains may also have a similar proportion of snow-covered grid cells. While the PNW likely has more
455 terrestrial grid cells that are covered by snow for a longer period throughout the water year, it also has a significant number
456 of ocean grid cells where snow redistribution would not be activated.

Parallel SnowModel - Strong Scaling



457


458 **Figure 9: The left panel displays approximate speedup (\bar{S} ; Eq. 3) as a function of the number of processes (N) for local and state**
 459 **sized simulations (Tuolumne, CO Headwaters, and Idaho), while the right panel shows \bar{S} for the regional and continental sized**
 460 **domains (PNW, Western US, and CONUS).**

Domain	N_x	N_y	\hat{P} or P^*	Number of Processes	Execution Time (m)	Relative Speedup	Approximate Speedup	Approximate Efficiency (%)
				N	$T(N)$	$\bar{S}(N)$	$\hat{S}(N)$	$\bar{E}(N)$
Tuolumne	311	185	\hat{P}	1	13	---	---	100
			P^*	52	0.7	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

461

462 **Table 1: Strong scaling results containing grid dimensions (N_x and N_y), number of processes, execution time, relative speedup,**
 463 **approximate speedup, and approximate efficiency of simulations executed with the minimum and maximum number of processes**
 464 **for the Tuolumne, CO Headwaters, Idaho, PNW, Western US, and CONUS domains.**

465 Strong scaling analysis is useful for I/O and memory bound applications to identify a setup that results in a reasonable
 466 runtime and moderate resource costs. Based on these scaling results, Figure 10 contains the relationship between the number
 467 of processes (N) at which each domain is estimated to reach 50% \bar{E} (using linear interpolation) with the total number of grid
 468 cells in the y dimension (N_y) and the average number of grid cells in the y dimension per process (l_{ny} ; inset Figure 10). At
 469 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
 470 values of l_{ny} (5 to 11) for these year-long simulations that vary in both domain size and the proportion of snow-covered
 471 area. Similar relationships (Figure 10) can be used to approximate the scalability of Parallel SnowModel on different sized
 472 domains and can be adjusted for the desired level of efficiency. For example, we decided to run the CONUS simulations
 473 (Sect. 4.3) using 1800 processes based on its 70% approximate efficiency.

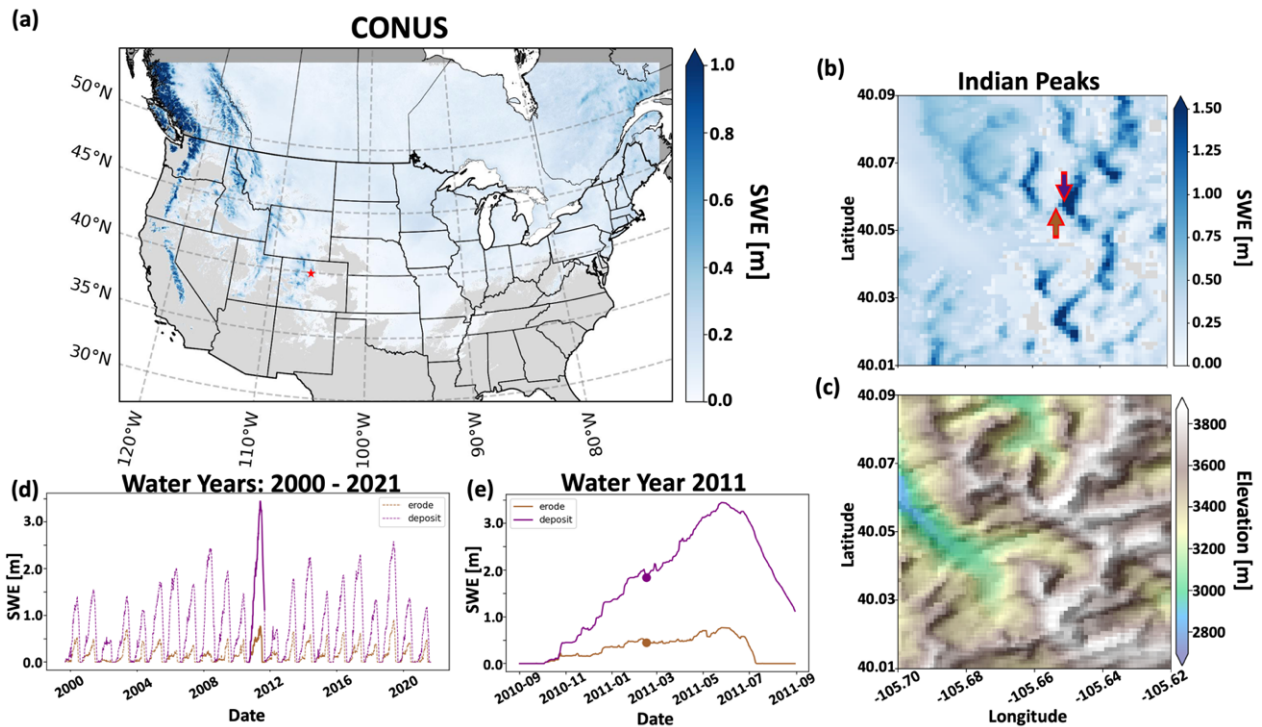


474

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

479 4.3 CONUS Simulations

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



493

494 **Figure 11: Simulation results of Parallel SnowModel over CONUS using the WRF projection. (a) Spatial patterns of SWE over the**
 495 **CONUS domain for 12 February 2011, (b) highlighting the SWE distribution (c) and topography with an applied hillshade of a sub-**
 496 **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**
 497 **water year for grid cells (“erode” and “deposit”) identified in panel (b). The “erode” and “deposit” grid cells highlight areas of similar**
 498 **elevation but significant differences in SWE evolution resulting from blowing-snow redistribution processes.**

499 5 Discussion

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

508 Simulation experiments were conducted using Parallel SnowModel to validate the parallel logic, interpret its performance
 509 across different algorithm versions and across different domains sizes, and demonstrate its ability to simulate continental
 510 domains at high-resolution. Code profiling and speedup analyses over the CO Headwaters domain helped identify
 511 bottlenecks in file I/O and processor communication in SnowTran-3D during the development of the parallel algorithm

512 (Sect. 4.1). Corrections to the referred bottlenecks allowed Parallel SnowModel to scale up to regional and continental sized
513 simulations and highlights the value of optimizing scientific code. The scalability analyses showed the effectiveness of
514 running Parallel SnowModel with an increasing number of processes on state, regional, and continental domains that contain
515 different proportions in both size (N_x and N_y) and snow-covered grid cells (Sect 4.2). For Parallel SnowModel scalability is
516 primarily dependent on the number of grid cells per process (N_x and l_{ny}) and is affected by snow redistribution, which is
517 dependent on the proportion of terrestrial grid cells with sufficient winds and available soft snow to be redistributed (Sect.
518 3.1.2.2). For example, a maritime snowpack (e.g. PNW) as compared to a continental snowpack (e.g. CO Headwaters), may
519 be deeper and more spatially extensive but potentially lacks a high frequency of soft snow above tree-line to activate snow
520 redistribution. Furthermore, the similar relationships among efficiency and domain decomposition observed on the simulated
521 domains that vary in size, topography, vegetation, and snow classes (Sturm and Liston, 2021; Liston and Sturm, 2021) (Fig.
522 10), make it reasonable to extrapolate the results from these simulation experiments to other domains within CONUS.
523 Additionally, these experiments emphasize the relationships among speed, memory, and computing resources for Parallel
524 SnowModel. A common laptop (~ 4 processes) has sufficient CPUs to run local sized domains within a reasonable amount
525 of time, but likely does not have sufficient memory for state-sized simulations. Similarly, the minimum memory (1160 GB;
526 Fig. 1) required to run the CONUS domain, could be simulated on a large server (~ 128 processes) with one process per
527 node. However, extrapolating from our scaling results on Cheyenne (Figure 9), we estimate it would take over 2.5 days to
528 run a CONUS simulation for one water year with this configuration. In contrast, it took approximately 5 hours for CONUS
529 to run on the Discover supercomputer using 1800 processes. Therefore, by the time it took the large server to complete a
530 CONUS simulation for one water year, 12 water years could have been simulated on a supercomputer. Lastly, results from
531 the CONUS simulation highlight the ability of Parallel SnowModel to run high-resolution continental simulations, while
532 maintaining fine-scale snow processes that occur at a local level (Sect. 4.3).

533 SnowModel can simulate high-resolution outputs of snow depth, density, SWE, grain size, thermal resistance, snow strength,
534 snow albedo, landscape albedo, meltwater production, snow-water runoff, blowing snow flux, visibility, peak winter SWE,
535 snow-season length, snow onset date, snow-free date, and more, all produced by a physical model that maintains consistency
536 among variables. While several snow data products exist, few capture the suite of snow properties along with the spatio-
537 temporal extents and resolutions that can benefit a wide variety of applications. For example, current snow information
538 products include the NASA daily SWE distributions globally for dry (non-melting) snow on a 25 km grid (Tedesco and
539 Jeyaratnam, 2019), a NASA snow-cover product on a 500 m grid (Hall et al., 2006) that is missing information due to clouds
540 approximately 50% of the time (Moody et al., 2005), and the Snow Data Assimilation System (SNODAS) daily snow
541 information provided by the National Oceanic and Atmospheric Administration (NOAA) and the National Weather Service
542 (NWS) National Operational Hydrologic Remote Sensing Center (NOHRSC) on a 1 km grid (Center, 2004), which is itself
543 model derived and has limited geographic coverage and snow properties. The Airborne Snow Observatory (ASO) provides
544 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
545 (Painter et al., 2016), but has limited spatio-temporal coverage and a high cost of acquisition. Furthermore, there are many

546 fields of study that can benefit from 100 m resolution information of internally consistent snow variables, including wildlife
547 and ecosystem, military, hydrology, weather and climate, cryosphere, recreation, remote sensing, engineering and civil
548 works, and industrial applications. The new Parallel SnowModel described here permits the application of this modeling
549 system to very large domains without sacrificing spatial resolution.

550 **6 Conclusions**

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

564 **Appendix A**

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

568 **Appendix B**

569 Validation SnowModel experiments were run in serial and in parallel over the Tuolumne and CO Headwaters domains (Sect.
570 4.1) using the RMSE statistic (Eq. 3). Important output variables from EnBal, MicroMet, SnowPack, and SnowTran-3D
571 demonstrated similar, if not identical values, when compared to serial results for all timesteps during the simulations; RMSE
572 values were within machine precision ($\sim 10^{-6}$) regardless of the output variable, domain, or number of processes used. The
573 validated output variables include albedo [%], precipitation [m], emitted longwave radiation [$W * m^{-2}$], incoming longwave

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

579 **Code, data availability, and supplement**

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

584 **Author contribution**

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

588 **Competing interests**

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

590 **Disclaimer**

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

593 **Financial support**

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

599 Lab (CISL) and NASA High-End Computing (HEC) Program through the NASA Center for Climate Simulation (NCCS) at
600 Goddard Space Flight Center.
601

602 **Acknowledgements**

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

606 **References**

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

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

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

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

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

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

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

623 Essery, R., Li, L., and Pomeroy, J.: A distributed model of blowing snow over complex terrain, *Hydrological Processes*, 13,
624 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.

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

628 Fang, X. and Pomeroy, J.: Modeling blowing snow redistribution to prairie wetlands, *Hydrological Processes*, 23, 2557-
629 2569, 10.1002/hyp.7348, 2009.

630 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.
631 J., Casey, K. A., and Choudhury, B.: A blended global snow product using visible, passive microwave and scatterometer
632 satellite data, *International Journal of Remote Sensing*, 32, 1371-1395, 10.1080/01431160903548013, 2011.

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

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

637 Hall, D., Riggs, G., and Salomonson, V.: MODIS/Terra Snow Cover 5-Min L2 Swath 500m, Version, 5, 2011167.2011750,
638 2006.

639 Hamman, J. J., Nijssen, B., Bohn, T. J., Gergel, D. R., and Mao, Y.: The Variable Infiltration Capacity model version 5
640 (VIC-5): Infrastructure improvements for new applications and reproducibility, *Geoscientific Model Development*, 11, 3481-
641 3496, 2018.

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

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

649 Hoppinen, Z. M., Oveisgharan, S., Marshall, H.-P., Mower, R., Elder, K., and Vuyovich, C.: Snow Water Equivalent
650 Retrieval Over Idaho, Part B: Using L-band UAVSAR Repeat-Pass Interferometry, *The Cryosphere Discussions*, 2023, 1-24,
651 2023.

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

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

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

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

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

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

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

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

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

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

- 675 Liston, G., Reinking, A. K., and Boleman, N.: Daily SnowModel Outputs Covering the ABoVE Core Domain, 3-km
676 Resolution, 1980-2020, 10.3334/ORNLDAAC/2105, 2022.
- 677 Liston, G. E.: Local advection of momentum, heat, and moisture during the melt of patchy snow covers, *Journal of Applied*
678 *Meteorology and Climatology*, 34, 1705-1715, 1995.
- 679 Liston, G. E.: Representing Subgrid Snow Cover Heterogeneities in Regional and Global Models, *Journal of Climate*, 17,
680 1381-1397, 10.1175/1520-0442(2004)017<1381:Rsschi>2.0.Co;2, 2004.
- 681 Liston, G. E. and Elder, K.: A distributed snow-evolution modeling system (SnowModel), *Journal of Hydrometeorology*, 7,
682 1259-1276, 2006a.
- 683 Liston, G. E. and Elder, K.: A Meteorological Distribution System for High-Resolution Terrestrial Modeling (MicroMet),
684 *Journal of Hydrometeorology*, 7, 217-234, 10.1175/jhm486.1, 2006b.
- 685 Liston, G. E. and Hall, D. K.: An energy-balance model of lake-ice evolution, *Journal of Glaciology*, 41, 373-382, 1995.
- 686 Liston, G. E. and Hiemstra, C. A.: A simple data assimilation system for complex snow distributions (SnowAssim), *Journal*
687 *of Hydrometeorology*, 9, 989-1004, 2008.
- 688 Liston, G. E. and Hiemstra, C. A.: The changing cryosphere: Pan-Arctic snow trends (1979–2009), *Journal of Climate*, 24,
689 5691-5712, 2011.
- 690 Liston, G. E. and Mernild, S. H.: Greenland freshwater runoff. Part I: A runoff routing model for glaciated and nonglaciated
691 landscapes (HydroFlow), *Journal of Climate*, 25, 5997-6014, 2012.
- 692 Liston, G. E. and Sturm, M.: A snow-transport model for complex terrain, *Journal of Glaciology*, 44, 498 - 516, 1998.
- 693 Liston, G. E. and Sturm, M.: Global Seasonal-Snow Classification, Version 1 [dataset], 2021.
- 694 Liston, G. E., Perham, C. J., Shideler, R. T., and Chevront, A. N.: Modeling snowdrift habitat for polar bear dens,
695 *Ecological Modelling*, 320, 114-134, <https://doi.org/10.1016/j.ecolmodel.2015.09.010>, 2016.
- 696 Liston, G. E., Winther, J.-G., Bruland, O., Elvehøy, H., and Sand, K.: Below-surface ice melt on the coastal Antarctic ice
697 sheet, *Journal of Glaciology*, 45, 273-285, 1999.
- 698 Liston, G. E., Haehnel, R. B., Sturm, M., Hiemstra, C. A., Berezovskaya, S., and Tabler, R. D.: Simulating complex snow
699 distributions in windy environments using SnowTran-3D, *Journal of Glaciology*, 53, 241-256, 2007.

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

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

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

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

711 Marsh, C. B., Pomeroy, J. W., Spiteri, R. J., and Wheeler, H. S.: A Finite Volume Blowing Snow Model for Use With
712 Variable Resolution Meshes, *Water Resources Research*, 56, e2019WR025307, <https://doi.org/10.1029/2019WR025307>,
713 2020.

714 Miller, P., Robson, M., El-Masri, B., Barman, R., Zheng, G., Jain, A., and Kalé, L.: Scaling the isam land surface model
715 through parallelization of inter-component data transfer, 2014 43rd International Conference on Parallel Processing, 422-
716 431,

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

719 Moody, E. G., King, M. D., Platnick, S., Schaaf, C. B., and Gao, F.: Spatially complete global spectral surface albedos:
720 Value-added datasets derived from Terra MODIS land products, *IEEE Transactions on Geoscience and Remote Sensing*, 43,
721 144-158, 2005.

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

727 Mortezapour, M., Menounos, B., Jackson, P. L., Erler, A. R., and Pelto, B. M.: The role of meteorological forcing and snow
728 model complexity in winter glacier mass balance estimation, Columbia River basin, Canada, *Hydrological Processes*, 34,
729 5085-5103, <https://doi.org/10.1002/hyp.13929>, 2020.

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

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

734 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)
735 [SnowModel-1.0](https://github.com/NCAR/Parallel-SnowModel-1.0), 2023.

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

738 Nolin, A. W.: Recent advances in remote sensing of seasonal snow, *Journal of Glaciology*, 56, 1141-1150,
739 10.3189/002214311796406077, 2010.

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

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

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

746 Parhami, B.: SIMD machines: do they have a significant future?, *ACM SIGARCH Computer Architecture News*, 23, 19-22,
747 1995.

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

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

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

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

758 Prokop, A. and Schneiderbauer, S.: The atmospheric snow-transport model: SnowDrift3D, *Journal of Glaciology*, 57, 526-
759 542, 10.3189/002214311796905677, 2011.

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

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

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

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

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

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

774 Sharma, V., Swayne, D., Lam, D., MacKay, M., Rouse, W., and Schertzer, W.: Functional Parallelization of a Land Surface
775 Model in Regional Climate Modeling, 2004.

776 Skofronick-Jackson, G. M., Johnson, B. T., and Munchak, S. J.: Detection Thresholds of Falling Snow From Satellite-Borne
777 Active and Passive Sensors, *IEEE Transactions on Geoscience and Remote Sensing*, 51, 4177-4189,
778 10.1109/TGRS.2012.2227763, 2013.

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

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

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

784 Takala, M., Luojus, K., Pulliainen, J., Derksen, C., Lemmetyinen, J., Kärnä, J.-P., Koskinen, J., and Bojkov, B.: Estimating
785 northern hemisphere snow water equivalent for climate research through assimilation of space-borne radiometer data and
786 ground-based measurements, *Remote Sensing of Environment*, 115, 3517-3529, <https://doi.org/10.1016/j.rse.2011.08.014>,
787 2011.

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

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

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

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

797 Vuyovich, C. M., Jacobs, J. M., and Daly, S. F.: Comparison of passive microwave and modeled estimates of total watershed
798 SWE in the continental United States, *Water Resources Research*, 50, 9088-9102, <https://doi.org/10.1002/2013WR014734>,
799 2014.

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

802 Wrzesien, M. L., Durand, M. T., Pavelsky, T. M., Kapnick, S. B., Zhang, Y., Guo, J., and Shum, C. K.: A New Estimate of
803 North American Mountain Snow Accumulation From Regional Climate Model Simulations, *Geophysical Research Letters*,
804 45, 1423-1432, <https://doi.org/10.1002/2017GL076664>, 2018.

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

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

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

812