



Computing water flow through complex landscapes, Part 3: Fill-Spill-Merge: Flow routing in depression hierarchies

Richard Barnes^{1,2,3}, Kerry L. Callaghan^{4,5}, and Andrew D. Wickert^{4,5}

¹Energy & Resources Group (ERG), University of California, Berkeley, USA

²Electrical Engineering & Computer Science, University of California, Berkeley, USA

³Berkeley Institute for Data Science (BIDS), University of California, Berkeley, USA

⁴Department of Earth & Environmental Sciences, University of Minnesota, Minneapolis, USA

⁵Saint Anthony Falls Laboratory, University of Minnesota, Minneapolis, USA

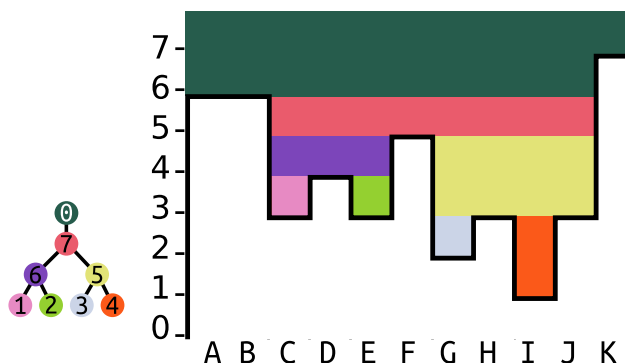
Correspondence: Richard Barnes (richard.barnes@berkeley.edu)

1 **Abstract.** Depressions—inwardly-draining regions—are common to many landscapes. When there is sufficient moisture, de-
2 pressions take the form of lakes and wetlands; otherwise, they may be dry. Hydrological flow models used in geomorphology,
3 hydrology, planetary science, soil and water conservation, and other fields often eliminate depressions through filling or breach-
4 ing; however, this can produce unrealistic results. Models that retain depressions, on the other hand, are often undesirably
5 expensive to run. In previous work we began to address this by developing a depression hierarchy data structure to capture the
6 full topographic complexity of depressions in a region. Here, we extend this work by presenting a Fill-Spill-Merge algorithm
7 that utilizes our depression hierarchy to rapidly process and distribute runoff. Runoff fills depressions, which then overflow
8 and spill into their neighbors. If both a depression and its neighbor fill, they merge. We provide a detailed explanation of the
9 algorithm as well as results from two sample study areas. In these case studies, the algorithm runs 90–2,600× faster (with a
10 2,000–63,000× reduction in compute time) than the commonly-used Jacobi iteration and produces a more accurate output.
11 Complete, well-commented, open-source code is available on Github and Zenodo.

12 1 Introduction

13 Depressions (see Lindsay (2015) for a typology) are inwardly-draining regions of a DEM that lack any outlet to an ocean or
14 other designated base elevation. Depressions occur naturally, and can be formed by glacial erosion and/or deposition (Brecken-
15 ridge and Johnson, 2009), compressional and/or extensional tectonics (Reheis, 1999; Hilley and Strecker, 2005), and cratering
16 Cabrol and Grin (1999). They often host lakes and wetlands by retaining water locally. Depressions may themselves contain
17 depressions. Such regions confound algorithms for geomorphological and terrain analysis, as well as those for hydrological
18 modeling, because many such algorithms simply route water down topographic slope following the local gradient: depressions
19 neither fill with water, nor drain.

20 Many hydrological models deal with the complexity of depressions by removing them. This can be done either by filling the
21 depressions with earth so that they form a flat region of landscape (e.g. Jenson and Domingue (1988); Martz and Jong (1988));
22 breaching (Martz and Garbrecht, 1998) or carving them (Soille et al., 2003) so that water flows from their lowest point through



42

43 **Figure 1. A single subtree of a depression hierarchy and the depression it represents.** Depressions 1–4 are leaf depressions. Depression
44 6 is a parent depression (also termed a meta-depression) that contains depressions 1 and 2. Water from the plateau on the left above cells *A*
45 and *B* might *fill* Depression 1 (cell *C*), causing it to *spill* into Depression 2 (cell *E*). Only when both depressions are full do they *merge* and
46 begin filling Depression 6 (cells *C*, *D*, and *E*). Modified from Barnes et al. (2020).

23 the carved channel and onward to downstream regions; or some combination of these (Lindsay and Creed, 2005b; Schwanghart
24 and Scherler, 2017; Soille, 2004; Lindsay, 2016). This approach is justified for situations in which spatiotemporal aspects of
25 the analysis allow depressions to be ignored or for cases in which all depressions can be considered to be data errors (Lindsay
26 and Creed, 2005a). Historically, many DEMs were constructed from sparse data, and small data errors produced depressions,
27 especially in flat areas (O’Callaghan and Mark, 1984). Such an assumption is no longer justified, as improved and increasingly
28 high-resolution data have become available (Li et al., 2011). Even coarse-resolution data are capable of resolving real-world
29 depressions (e.g. Riddick et al., 2018; Wickert, 2016). With this in mind, new approaches are beginning to be examined,
30 particularly in post-glacial landscapes where depressions have a significant impact on local hydrology (e.g. Lai and Anders
31 (2018)) and therefore cannot be ignored during modeling.

32 FlowFill (Callaghan and Wickert, 2019) began to combat this problem by routing water across landscapes in a way that
33 conserved water volume, creating flow-routing surfaces that could still contain real depressions. Under reasonable runoff con-
34 ditions, their results show landscapes that still contain depressions and disrupted flow routes. The FlowFill method iteratively
35 routes water from higher to lower terrain. As depressions fill, they pose an extreme challenge to such a method: since water
36 seeks a level surface, the surface of a filled depression must eventually become flat and any fluid flowing onto the surface
37 diffuses across it. Even for moderately-sized surfaces it can take many iterations for a solver to reach steady state; we provide
38 a theoretical analysis of this in Section 4. Runtimes for FlowFill ranged from seconds to days: large datasets quickly became
39 unwieldy. Of those examples tested by Callaghan and Wickert (2019), the slowest was a dataset of 4,176,000 cells which took
40 approximately 33 hours for FlowFill to process. In contrast, the Fill-Spill-Merge algorithm presented here fills a similarly-sized
41 dataset in 8.7 s.

47 To achieve this, we developed a data structure—the *depression hierarchy*—which represents the topologic and geographic
48 structure of depressions. In an accompanying paper, we provide details concerning the depression hierarchy and its construc-



49 tion (Barnes et al., 2020). In this paper, we explain how the depression hierarchy can be leveraged to accelerate hydrological
50 models using a paradigm we call *Fill-Spill-Merge*.

66 2 Using The Depression Hierarchy

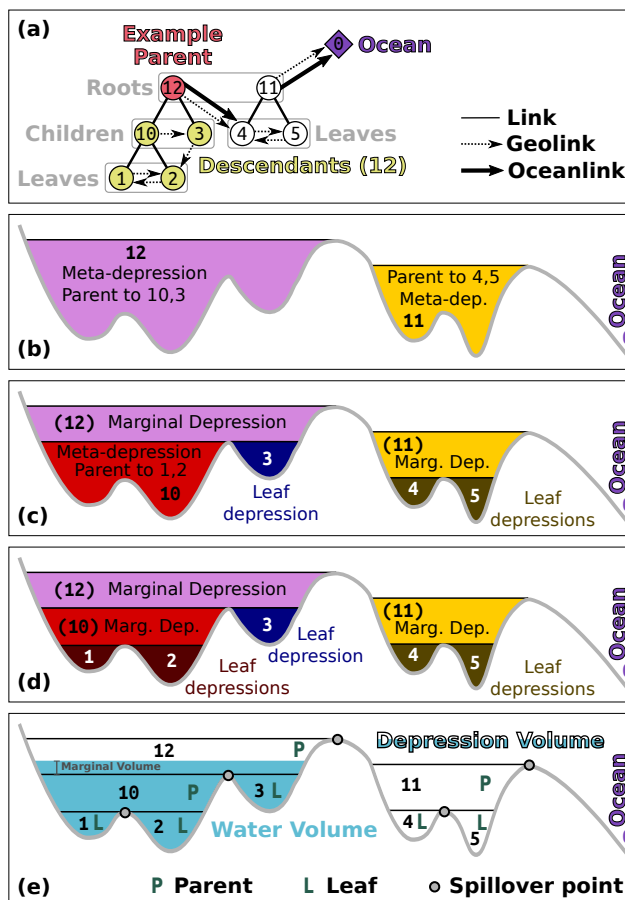
67 Depressions can themselves contain depressions, as shown in Figure 1. A depression hierarchy (DH) is a forest of binary trees,
68 as shown in Figure 2a, that represents the relationships between depressions (Figure 2a–d). Each node in the DH represents a
69 depression. Nodes higher in the DH are depressions that themselves contain depressions; we term these *meta-depressions*. A
70 node in the DH can have several classifications:

- 71 – **Parent:** A node, such as #10 in Figure 2, that represents a meta-depression, and whose topological descendants therefore
72 also form depressions.
- 73 – **Child:** A depression, such as both #10 and #1 in Figure 2, that geographically and topologically exists within the meta-
74 depression formed by its parent.
- 75 – **Leaf:** A depression, such as #1 and #2 in Figure 2, that has no children. The leaves of the binary trees represent the
76 smallest, most deeply-nested depressions. If a landscape were initially devoid of water, then water flowing down slopes
77 would begin to collect in some subset of these leaf depressions before it would begin to fill their parent depressions.
- 78 – **Root:** A depression, such as #0 in Figure 2, that has no parent. This term may also refer to any node that is used as the
79 starting point for a traversal that only considers the node and its descendants.
- 80 – **Descendant:** A child of a given parent, or the child of a child of that parent, and so on. In Figure 2, #1, #2, #3, and #10
81 are all descendants of #12.
- 82 – **Sibling:** Every node has either no children (leaf nodes) or two children. Nodes which share a parent are siblings. In
83 Figure 2, #1 and #2 are siblings, as are #4 and #5.

84 As depressions fill, their water surfaces eventually reach a *spill elevation* (Figure 2e) at which they overflow into neigh-
85 boring depressions. During this spilling, water flows from a depression into a geographically neighboring leaf depression,
86 topologically connected by a *geolink*. The spill elevations in Figure 1 are the highest points of each band of color.

87 Each node in the DH is associated with several properties:

- 88 – **Depression volume:** This is the *total* volume of water that the depression, including all of its descendants, can contain
89 before spilling over.
- 90 – **Water volume:** This is the *total* volume of water *actually being stored* in the depression. A parent depression will have
91 a non-zero water volume only if all of its children are completely full and the parent itself contains some additional
92 volume of water. In this case, the water volume will be the sum of the water volumes of the children and the additional



51

52 **Figure 2. Terminology for the depression hierarchy and water flow through it.** The depression hierarchy shown here is drawn from the
 53 left hand side of Figure 1 from the companion paper by Barnes et al. (2020). (a) Topology. A *parent* and its *descendants* are associated with
 54 depressions (b–d). Direct descendants are called *children*. *Leaves* are the terminal members of the depression hierarchy; they have no children
 55 and represent simple depressions (i.e., those that are not meta-depressions). Members of a single *binary tree* are joined in their hierarchy
 56 through *links*; directional links that represent water-spillover directions between geospatially adjacent depressions are called *geolinks*. Flow
 57 from one binary tree into another and towards the ocean follows the *oceanlinks*. Though only one binary tree is shown, the ocean may be the
 58 parent to an arbitrarily large *forest* of binary trees. (b) Parents in the hierarchy form *meta-depressions* — depressions that encompass other
 59 depressions. (c) These meta-depressions contain *leaf depressions* — depressions that themselves contain no depressions. These are associated
 60 with leaves in the depression hierarchy. Meta-depression 12 also contains another meta-depression, 10. The regions of Depressions 11 and
 61 12 that lie above their child depressions are termed “marginal depressions”. (d) Meta-depression 10 contains leaf depressions 1 and 2. (e)
 62 Water flow in the depression hierarchy. Water first fills *leaf depressions* before flooding into neighboring *depressions*. Once a depression and
 63 its neighbor are completely filled, their *parent* begins to flood. The *depression volume* is the full geometric volume of the depression. The
 64 *water volume*, naturally, is the volume of water within a given depression. The *marginal volume* is the volume of water partially filling the
 65 top-level meta-depression; appropriately spreading this water across the landscape is the topic of Section 3.3.



93 margin of water contained within the parent (i.e., the “marginal volume” indicated on Figure 2). Parents whose children
94 are not all filled with water will have a water volume equal to zero. In this way, we can use this property to determine
95 which portions of the DH are fully or partially filled, and which are the highest water-containing nodes in any of the
96 binary trees.

97 – **Geolink:** When a depression spills, its water passes into the subtree rooted by its sibling. However, in a full model of
98 flow, the water would move downslope from the spill cell into whichever leaf depression of the sibling is geographically
99 proximal to the spill cell. *Geolinks* are pointers from depressions higher in the DH to the leaf depressions that receive
100 their water if they overflow. These are the dashed lines shown in Figure 2. Geolinks are similar to the connections used
101 in a threaded binary tree (Fenner and Loizou, 1984).

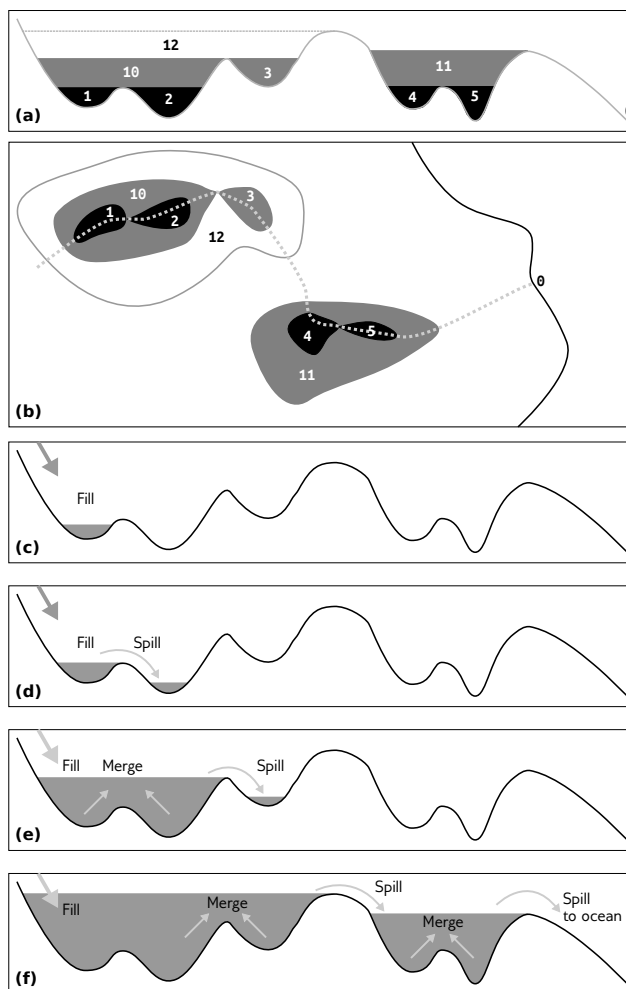
102 – **Ocean link:** Depressions high in the mountains may overflow down escarpments to depressions far below. In this case,
103 the depressions do not overflow into each other: the relationship is one-way. There can be multiple such escarpments, so
104 this can happen multiple times. In such cases, each group of depressions forms a proper binary tree. However, the root
105 of one of the trees has both an *ocean link* and a geolink to a leaf node of the downstream binary tree. In Figure 2, both
106 #11 and #12 are the root nodes of a set of nested depressions. #12 has an ocean link (heavy arrow) to #4, one of the leaf
107 depressions of #11. #12 also has a geolink (dotted arrow) to #4. #11 itself has an ocean link and a geolink to the ocean. In
108 many of the algorithms discussed below, ocean-linked nodes are processed similarly to children; however, information
109 is usually not passed across ocean links. Oceanlinks are used solely for guiding traversals of the depression hierarchy
110 whereas water is passed through geolinks.

111 3 The Algorithm

130 The Fill-Spill-Merge algorithm consists of several steps, outlined here, depicted in Figures 3 and 4, and shown in flowchart
131 form in Figure 5. First (Section 3.1), surface water needs to move downhill, either to the ocean (i.e., a designated sink region
132 or the map edge) or to collect in pit cells – the deepest points within leaf depressions. This operation takes place across all the
133 cells of the DEM. Second (Section 3.2), water is redistributed across the depression hierarchy such that any depressions that
134 have filled sufficiently must spill over into neighboring depressions and, if both depressions are full, flood their parent to merge
135 into a single, larger body of water within a meta-depression. This operation is done without explicitly considering the cells of
136 the DEM, which makes it very fast. Third and finally (Section 3.3), the water within the depression hierarchy is translated into
137 an extent and depth of flooding across the topographic surface (DEM).

140 Computing a depression hierarchy (Barnes et al., 2020) is a necessary precursor to running Fill-Spill-Merge. The specific
141 outputs from the depression hierarchy that are used in the Fill-Spill-Merge algorithm are:

- 142 – *DH*: the depression hierarchy itself.
- 143 – *Flowdirs*: a matrix of flow directions, indicating which of a cell’s neighbors receives its flow. Because Priority-Flood
144 (Barnes et al., 2014) is used to generate the depression hierarchy, flat areas are automatically resolved.




112


113 **Figure 3. Fill-Spill-Merge process.** Water moves through topographic depressions by filling them, spilling into
114 meta-depressions. (a) Topographic cross section with labeled leaf depressions and their parents, following the left-hand side of the
115 depression hierarchy in Figure 2. “0” represents the ocean; other numbers represent leaves and parents that together form depressions and
116 meta-depressions. (b) Map showing this depression structure; the cross-section in (a) follows the dotted gray line. (c) A water source to the
117 left begins to fill Depression 1. (d) Continued water input causes Depression 1 to overflow and spill into Depression 2. (e) Depression 2 fills,
118 causing Depressions 1 and 2 to fill their parent (10) and merge to form a metadepression. This metadepression overflows into Depression 3.
119 (f) Depression 3 fills and merges with Meta-Depression 10 (1 and 2 being implied members based on their position in the hierarchy) to flood
120 their parent, 12. After Meta-Depression 12 overflows, it enters Depression 4, which then fills and spills into Depression 5. After Depression
121 5 floods, its waters join with those from Depression 4 to fill Meta-Depression 11, which then spills to the ocean. Figures 4 and 5 describe the
122 algorithm in more specific detail.



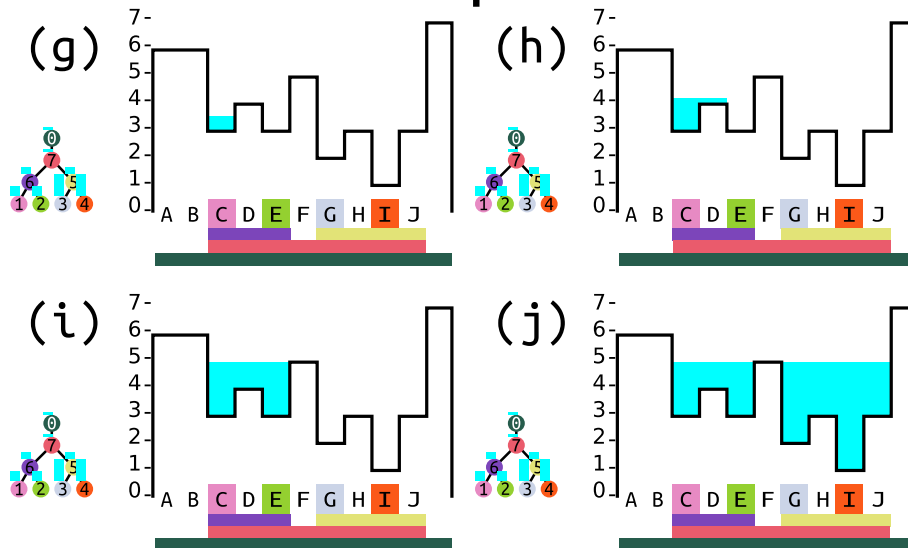
Move water downhill to pits



Overflow and merge depressions

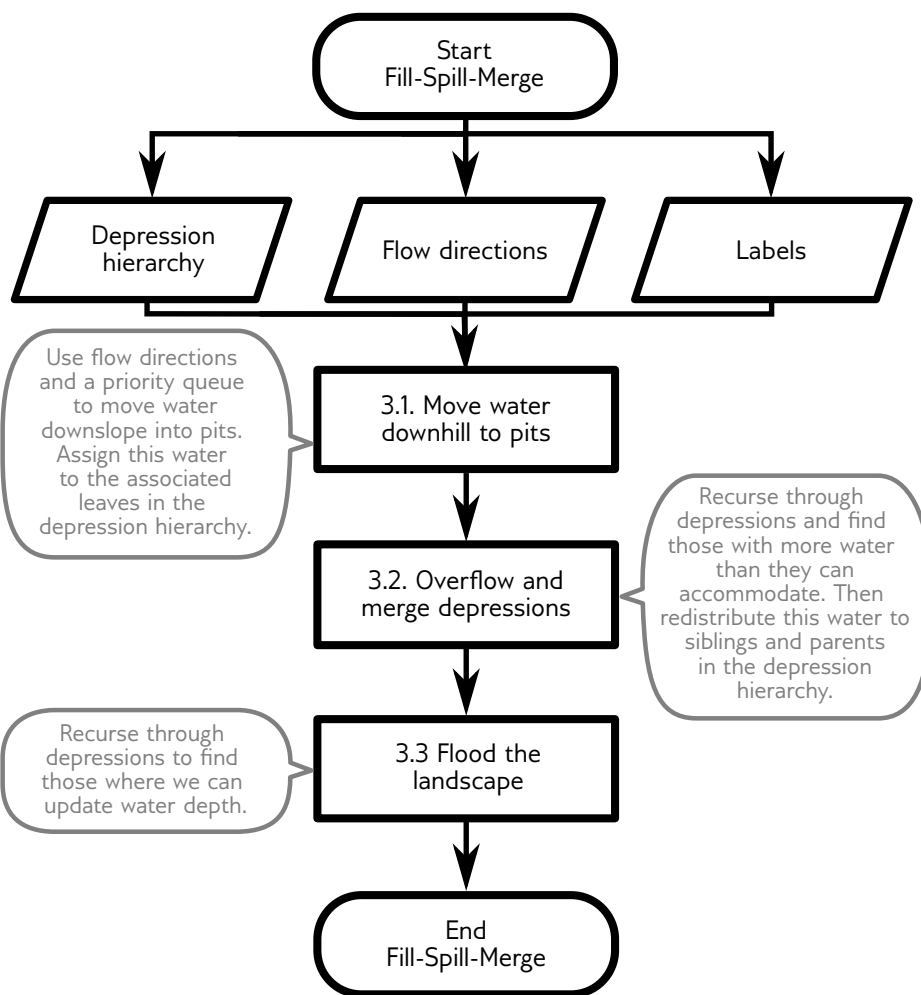


Flood the landscape



123

124 **Figure 4. Visual Overview of the Algorithm.** In this figure the heights of the water bars are non-additive: only the changes between panels
 125 are important. The algorithm consists of three major stages (Figure 5). From its initial distribution (A), water is moved downhill into pit cells
 126 (B, §3.1). Water is then moved within the depression hierarchy (C–F, §3.2): water in depressions with insufficient volume overflows first into
 127 their sibling depressions (D) and then – if the sibling depression becomes filled – passes to their parents (E, F). Any leftover water overflows
 128 into the ocean (F) and is forgotten. Depressions to be flooded are then identified and flooded (§3.3) starting from an arbitrarily-chosen pit
 129 cell (G–J).



138

139 **Figure 5. Flowchart showing the main steps taken by the algorithm.** These steps are described in more detail in §3.1 to §3.3.

145 – *Labels*: a matrix indicating the leaf depression to which each cell belongs.

146 By routing water according to the DH, we significantly accelerate the compute speed and ensure that the full network of
147 depressions is a topologically correct directed tree. Each of the following subsections details one of the numbered steps along
148 the central path of the flowchart shown in Figure 5.

149 3.1 Move Water Downhill to Pits

150 We route water in a similar way to standard flow-accumulation algorithms (Mark, 1988; Wallis et al., 2009; Barnes, 2017), but
151 for completeness summarize our approach here. Flow directions for each cell have already been identified by the DH. Each
152 cell calculates how many of its neighbors flow into it. We call this value the cell's dependency count, as it describes the number



153 of immediate upstream cells whose flow accumulation must be resolved before flow accumulation at the given cell can be
154 computed. Local maxima in the DEM are identified as those cells that receive no flow from any neighbor. These local maxima
155 are placed in a queue. Cells are then popped (i.e., noted while being removed) from this queue. The cells determine how
156 much flow they generate locally (perhaps referring to matrix of rainfall values) and add this to their flow accumulation value.
157 They then add their flow accumulation to their downstream neighbor's and set their own flow accumulation value to zero. The
158 neighbor's dependency count is then decremented. If the neighbor's dependency count has reached zero during this step, it is
159 added to the end of the queue. This process of accumulating flow, passing it downstream, decrementing the dependency count,
160 and adding cells to the queue continues until the queue is empty, at which point every cell on the map has been visited and any
161 water has been moved downslope. Braun and Willett (2013) present an alternative formulation based on a depth-first traversal,
162 but Barnes (2019) demonstrates that a breadth-first ordering, such as that presented here, is better suited to parallelism.

163 When the accumulated flow reaches the pit cell of a depression, the downhill-directed flow routing stops because there is no
164 downhill neighbor to receive the flow. At this point, all of the flow-accumulated water in the pit cell is moved into the pit cell's
165 associated leaf depression in the DH. That is, the water is moved out of the geographic space and into the topologic space. This
166 then enables mass-conserving depression flooding via rapid Fill-Spill-Merge calculations, as detailed below.

167 3.2 Overflow and Merge Depressions

168 At this point, the Fill-Spill-Merge algorithm has routed all of the surface water into either the ocean or into the leaf nodes of the
169 DH. The next step is to redistribute this water through the DH to nodes with enough volume to contain the water, and to send
170 any excess water to the ocean. This set of operations can be performed entirely in the depression hierarchy without reference
171 to the digital elevation model.

172 Intuitively, the process of filling, spilling, and merging can be visualized as occurring from leaf nodes to their parents
173 (Figure 3). Water must be redistributed such that leaf depressions containing more water than they can hold spill over into their
174 neighboring depression. If this neighboring depression is already full, then the excess water must pass to the parent of both the
175 depression and its neighbor. This process continues recursively until either the supplied water is exhausted or this water reaches
176 the ultimate parent, the ocean. In this latter case, all excess water is dropped from the model and the ocean is unaffected.

177 To efficiently redistribute water, the Fill-Spill-Merge algorithm performs nested depth-first traversals of the DH. The outer
178 traversal is post-order and considers each meta-depression in turn, from the most deeply nested to the least. For each meta-
179 depression, an inner traversal handles its overflows by moving water to its sibling (starting by filling the sibling's descendants)
180 and, if there's any left, passing it to the depression's parent. In this way, the outer traversal maintains an invariant: any meta-
181 depression it has processed does not contain an overflow.

182 The outer traversal of the DH (which is, after all, a forest of binary trees) begins with the ocean. For each depression, the
183 algorithm first recurses into the depression's left child and then into its right child. If any oceanlinks are found, the algorithm
184 also recurses into them. In the post-order portion of the traversal (which starts from the leaves and moves back up through
185 the depression hierarchy), the algorithm identifies any depressions containing more water than they can accommodate. This



186 process continues until the recursion returns to the ocean, at which point any additional water is assumed to be added to the
187 ocean without impacting sea level.

188 When an overflowed depression is located, the inner traversal redistributes this water. Let us call this overflowed depression *A*
189 and note that it contains some amount of excess water — that is, water beyond its depression capacity. Our goal is to distribute
190 this fixed amount of excess into neighbouring depressions. At each step below, the amount of this excess water remaining to be
191 distributed will either remain the same or decrease. When we pass water into a depression, it can go to one of three places: the
192 depression itself, its sibling, or its parent. Distributing the water to any of these places may itself cause an overflow. Therefore,
193 the inner (pre-order) traversal comprises the following steps:

194 1. Call the depression that we are currently considering *B*. This may be the depression we originally considered, depression
195 *A*, or it may be some other depression reached during the steps detailed below. We add water to *B* until either it fills or
196 all of the water is used. At this point, this part of the algorithm can terminate if: (i) there is no water left, (ii) *B* is the
197 parent of *A*, (iii) *B* acts as a parent of *A* by receiving its overflow via an oceanlink–geolink pair while not being a sibling
198 or descendant, or (iv) *B* is the ocean.

199 2. Otherwise, if *B* has a sibling and the sibling's water volume is less than its depression volume, then start from Step 1
200 with the new *B* set as the depression pointed to by the current *B*'s geolink.

201 3. Otherwise, if *B* has no sibling or the sibling's water volume is equal to its depression volume, then start from Step 1 with
202 the new *B* set as the parent of the current *B*. (Note that the parent may be the ocean or a node reached via an oceanlink).

203 During each such pass through the inner traversal, water moves at most one step in the DH towards the ocean.

204 The next step of the outer traversal, which begins one level in the DH closer to the ocean, identifies a less nested metade-
205 pression for which the inner traversal might need to be run. If this step were not supplied with information about prior water
206 redistribution, it could cause multiple traversals of a subtree of the DH, which would be computationally wasteful. To prevent
207 this, the inner traversal returns the ID of the final node in which it placed water: this node is the only node in the traversal with
208 spare capacity so future traversals can begin there. Therefore, on subsequent overflows, if such a cached value is available, then
209 the recursion skips directly to that node. This ensures that all the calls to this part of the algorithm take no more than $O(N)$
210 time collectively.

211 The following examples uses the geometry from Figure 2 to describe a set of inner traversals, starting with an overflowing
212 Depression #12. Step numbers mirror those above; numbers in parentheses indicate the number of recursions – that is, the
213 number of times that the inner-traversal algorithm has returned to Step 1:

214 1 Depression #12 fills and overflows.

215 2 Depression #12's water overflows into Depression #4, which is not full, following its geolink.

216 1(r1) Depression #4 acts as Depression #12's parent via a geolink–oceanlink pair. The inner traversal terminates.



217 At this point, the outer traversal moves one level closer to the ocean, and the inner traversal repeats, this time starting at
218 Depression #4.

219 1 Depression #4 fills and overflows.

220 2 Depression #4's water overflows into its sibling, Depression #5, which is not full and is a leaf depression. If Depression
221 #5 had descendants, water overflowing from Depression #4 would have followed a geolink to one of these.

222 1(r1) Depression #5s fills and overflows.

223 2(r1) Depression #4 is full.

224 3(r1) Depression #5 overflows into its parent, Depression #11.

225 1(r2) Depression #11 overflows into the ocean; the inner traversal terminates.

226 Now the outer traversal moves yet another level closer to the ocean, and the new inner traversal starts at Depression #11.

227 1 Depression #11 fills and overflows.

228 2 Depression #11 has no sibling.

229 3 Depression #11 overflows into its parent, the ocean; all remaining excess water is absorbed into an infinite sink.

230 1(r1) The now-selected node is the ocean; the inner traversal terminates.

231 At this point, the outer traversal moves one level closer to the ocean, and arrives at the ocean. The outer traversal also terminates.

232 3.3 Flood the landscape

233 After water moves through the DH (Section 3.2, above), each node in the DH exists in one of the three following states:

234 1. **Empty:** The depression's water volume is equal to zero. In this case, nothing needs to be done. The depression's descen-
235 dants might contain water, but the water never propagates to this level of the DH.

236 2. **Full:** The depression's water volume is equal to the volume of the depression itself. In this case, the depression is entirely
237 full. If the depression's parent contains water, then the calculation of water depth is dealt with at a higher stage in the
238 DH. If the depression's parent is empty, or if the depression's parent is the ocean, then the calculation is performed at
239 this level as described below.

240 3. **Partially filled:** The depression's water volume is less than its depression volume. In this case, the depth of water across
241 the depression and all its descendants' cells must be calculated at this level so that the depression fills to an appropriate
242 level. This is described below and indicated as the *marginal volume* on Figure 2e.



243 The next step is to distribute this water across the DEM, appropriately flooding geographic depressions.

244 Given the three states described above, the algorithm locates the highest-level node within each binary tree that contains
245 water. It does so by first traversing from the ocean to each leaf depression by recursively traveling to each node's children in
246 turn. Each time it reaches a leaf, the algorithm notes its label and pit cell. After identifying each of these, the algorithm reverses
247 direction, moving from child to parent so long as the parent node contains water. Therefore, this traversal towards the ocean
248 ends at the highest-level node whose parent does not contain water. Call this node L . The water volume contained within the
249 depression will only very rarely be exactly enough to perfectly flood it; therefore, we must spread water across the depression
250 to create a flat water surface.

251 To calculate water level within a depression, the algorithm begins by picking an arbitrary pit cell within it, and then uses
252 this as a seed from which to start building a priority queue through the depression. The priority queue returns cells ordered
253 from lowest to highest elevation. At each step through the priority queue, the algorithm checks whether a depression whose
254 outlet is at this elevation would have enough volume to hold the water. If so, the algorithm exits, having successfully defined
255 the flooded area. If not, it continues to build the priority queue.

256 To expand this brief conceptual discussion into a more formal set of steps, let us begin by calling the active cell – that is,
257 the one that is currently being considered by the algorithm – c_p . This cell is initially the arbitrary pit mentioned above, and is
258 added to the priority queue. The algorithm marks c_p , which stands for “cell of current highest priority”, as *visited*; all other
259 cells remain unvisited. The algorithm then follows these steps:

- 260 1. Pop c_p from the priority queue and use its elevation to calculate the volume of water that can be accommodated in the set
261 of cells processed so far (Equation 3, below). If this volume is enough to accommodate the volume of water available,
262 exit the loop and compute the final water level (Equation 4, below). Otherwise, proceed to Step 2.
- 263 2. Add the former c_p (which was popped in Step 1) to a plain queue, which records all of the cells scanned so far; these
264 cells will later be inundated.
- 265 3. Add the cells neighboring the former c_p that are not marked as *visited* to one of two lists. If an unvisited neighboring cell
266 shares a label with L or any of its descendants, then this neighboring cell is added to the priority queue. Each of these
267 neighboring cells is then marked as *visited*.
- 268 4. Choose the lowest-elevation cell in the priority queue and label it as the new c_p and return to Step 1. If the priority queue
269 is empty, then all cells in the same meta-depression as c_p or its descendants have been visited and we are now guaranteed
270 to have sufficient depression volume to hold all of the water.

271 Step 1 in this approach requires an efficient way to determine the volume of a depression below any given elevation. To do
272 so, we imagine a hypothetical outlet that drains the depression. If the depression is full enough that all of its cells receive water,
273 then the elevation of this hypothetical outlet is simply that of the topographic outlet from the depression. If the depression is not
274 yet completely filled, it can be visualized as a pipe in the side of the depression that is an infinite sink for any water entering it,
275 thereby acting analogously to an overflow drain below the edge of a sink or bathtub. If we call the elevation of this hypothetical



276 outlet is o and a depression contains cells of elevations $\{a, b, c, d, \dots\}$, then the capacity of the depression is

$$277 (o - a) + (o - b) + (o - c) + (o - d) + \dots = No - a - b - c - d - \dots \quad (1)$$

$$278 = No - \sum_{i=1}^N (\text{elevations}) \quad (2)$$

279 Now, consider cells $c_i = c_1, \dots, c_N$ in the plain queue (i.e., those that have been visited and popped from the priority queue),
280 we can calculate the volume of the depression below that of the last cell popped from the priority queue, the sill z_s , as:

$$281 V_{\text{dep}, z_s} = z_s N - \sum_{i=1}^N z_i \quad (3)$$

282 Here, V_{dep, z_s} is the volume of the depression below z_s , and z_i is the elevation of cell c_i . Thus, if we keep track of the number
283 of cells in a depression and their total elevation, it is possible to calculate the volume of a depression at any hypothetical outlet
284 level.

285 Once V_{dep, z_s} is greater than or equal to the volume of water in the depression, V_w , the plain queue contains all cells to be
286 flooded. At this point, the algorithm updates z_w , which is the water level within this depression. If $V_w = V_{\text{dep}, z_s}$, the algorithm
287 sets $z_w = z_N$. If instead $V_w < V_{\text{dep}, z_s}$, the available volume is greater than the water volume, and the algorithm calculates z_w
288 in the depression as follows:

$$289 z_w = \frac{1}{N} \left(V_w + \sum_{i=1}^N z_i \right). \quad (4)$$

290 We call this the Lake-Level Equation (LLE). The conditional usage of the LLE described above is purely for computational
291 efficiency: if $V_w = V_{\text{dep}, z_s}$, its solution is that $z_w = z_N$.

292 After solving for the water-surface elevation, the algorithm pops each cell in the plain queue ($c_i = c_1, \dots, c_N$), corresponding
293 to the flooded region, and sets its water elevation to the computed z_w . This is the final step of the Fill-Spill-Merge algorithm. At
294 this point, it outputs a file representing the topography plus water thickness across the domain (i.e., topography with depressions
295 filled or partially filled with water).

296 4 Theoretical Analysis

297 Here we use computational complexity as a means of contrasting the expected run-time of our algorithm against previous
298 algorithms such as FlowFill. To do so, we describe a simple iterative solver similar to FlowFill whose goal is to determine
299 an appropriate water level for a depression. The solver operates on a one-dimensional domain of cells bounded by high cliffs
300 on either side in which each cell may have a column of water. At each step, if the solver finds a discontinuity in water levels
301 between two cells, it responds by averaging the heights of the cells' water columns. (The solver we describe is known as
302 Jacobi's method.) The challenge we present to this solver is a direct analogue of routing flow along a stretch of river with
303 negligible gradient and is very similar to routing flow across the surface of a lake or ocean.



304 For our analysis, we imagine that the system is initialized with a high column of water on the left and no water anywhere
305 else. We call the cell with the water Cell 1. We call the cells to its right 2, 3, 4, and so on. During the solver's first step, Cell 1
306 is initialized. On its second step, Cell 1 averages its height with Cell 2. On the third step, Cell 2 averages with Cell 3 and Cell
307 1 then averages with Cell 2. On the fourth step, Cell 3 averages to 4, 2 averages to 3, and 1 averages with 2. Thus, the number
308 of cells affected at each step are: 1, 2, 3, 4, and so on. Since there must be *at least* as many steps as there are cells, we can say
309 that there are N steps. The total time, t_{compute} , is then

$$310 \quad t_{\text{compute}} = \sum_{i=1}^N i = \frac{N(N+1)}{2} \quad (5)$$

311 Thus, for any model (Callaghan and Wickert, 2019; Fan et al., 2013) that uses a scheme similar to our simple solver, the time
312 required to solve the model is in $O(N^2)$.

313 In contrast, the new algorithm runs in $O(N \log N)$ time in the worst case. Moving water downhill (Section 3.1) is a flow-
314 accumulation algorithm. This is known to take $O(N)$ time (Mark, 1988) and efficient variants exist for performing flow
315 accumulation in parallel on large datasets (Barnes, 2017) and on GPUs (Barnes, 2019), though for simplicity we do not use
316 these techniques here. Moving water within the depression hierarchy (Section 3.2) requires a depth-first post-order traversal of
317 the entire hierarchy. This type of traversal is a foundational algorithm in computer science and takes $O(N)$ time. Each node
318 in this traversal has the potential to overflow, which also results in a depth-first traversal, thereby requiring up to $O(N)$ time.
319 However, by using a jump table that persists between calls to the overflow function, we ensure that it is able to identify the
320 target of the overflow in amortized constant time; that is, the function is able to skip over fully-filled depressions. Finally, the
321 algorithm floods the digital elevation model from the pit cells up. This requires a depth-first post-order traversal, which calls
322 a flooding function (Section 3.3) on select subtrees of the DH. The depth-first traversal takes $O(N)$ time, as described above.
323 The priority queue used for flooding nominally takes $O(N \log N)$ time in the worst case for floating-point data and $O(N)$
324 time in the worst case for integer data (Barnes et al., 2014). However, with specialized data structures the time can be reduced
325 to $O(N)$ for both floating-point and integer data (Barnes et al., 2014). Most real datasets consist of many small depressions
326 whose cell counts $N_{\text{cells-in-dep}}$ are much smaller than the total number of cells in the digital elevation model. Therefore, the
327 actual time is for this step is $O(N_{\text{dep}} N_{\text{cells-in-dep}})$, where N_{dep} is the total number of depressions and $N_{\text{dep}} N_{\text{cells-in-dep}}$ can
328 be much less than N . Because the worst-case time complexity of any operation is $O(N)$, this bounds the time of the algorithm
329 as a whole. However, to reduce the potential for bugs, we use the C++ standard library's $O(N \log N)$ priority queue in our
330 implementation, at the cost of reduced performance.

331 To put this in more concrete terms, consider a long stretch of low-gradient river. Such a feature poses a lower bound on the
332 time of our simple solver. North America's Red River of the North runs for 885 km with a gradient that is often on the order of
333 0.03 m km^{-1} . On a 90 m grid of floating-point data, the river would be 9,833 cells long. Our simple (Jacobi) solver would then
334 take an estimated 97 million time units to reach a solution, whereas the new solver that we describe in this paper would take
335 9,833 time units, a $10,000\times$ speed-up. Our empirical results, below, support both the theory and this expected value.



Dataset	Dimensions	Cells	FSM Time (s)	Total Time (s)
Madagascar	2000×1000	$2.0 \cdot 10^6$	0.1	0.4
U.S. Great Basin	1920×2400	$4.6 \cdot 10^6$	0.2	8.7
Australia	5640×4200	$2.3 \cdot 10^7$	9.1	15.6
Africa	9480×9000	$8.5 \cdot 10^7$	65.3	118.0
N&S America	18720×17400	$3.2 \cdot 10^8$	53.2	231.6
Minnesota 30m topobathy	34742×23831	$8.2 \cdot 10^8$	307.8	792.6

337

338 **Table 1. Datasets used, their dimensions, and algorithm wall-times.** Tests were performed on the Comet cluster run by XSEDE (see
339 main text for full specifications). Times for Fill-Spill-Merge (“FSM Time”) alone and this time plus the depression hierarchy construction
340 time (“Total Time”) are shown. Topographic data for Madagascar, the U.S. Great Basin, Australia, Africa, and North & South America,
341 were clipped from the global GEBCO_08 30-arcsecond global combined topographic and bathymetric elevation data set (GEBCO, 2010).
342 The Minnesota 30m topobathy data is the merged result of two data sources. The topography is resampled from the Minnesota Geospatial
343 Information Office’s 1m LiDAR Elevation Dataset (MNGEO - Minnesota Geospatial Information Office, 2019). Bathymetric data were
344 provided by the Minnesota Department of Natural Resources (MNDNR - Minnesota Department of Natural Resources, 2014). Richard
345 Lively of the Minnesota Geological Survey merged and combined these data sets.

336 5 Empirical Tests

346 We have implemented the algorithm described above in C++11 using the Geospatial Data Abstraction Library (GDAL) (GDAL
347 Development Team, 2016) to read and write data. There are 981 lines of code of which 50% are or contain comments. The
348 code can be acquired from <https://github.com/r-barnes/Barnes2020-FillSpillMerge> and Zenodo (Barnes and Callaghan, 2020).

349 Tests were run on the Comet machine of the Extreme Science and Engineering Discovery Environment (XSEDE) (Towns
350 et al., 2014). Each node of the machine has 2.5 GHz Intel Xeon E5-2680v3 processors with 24 cores per node and 128 GB of
351 DDR4 DRAM. Code was compiled using GNU g++ 7.2.0 with full optimizations enabled. Scaling tests on datasets spanning
352 three orders of magnitude in terms of their number of cells are shown in Table 1. The GuessComp package written in the R
353 programming language by Agenis-Nevers et al. (2019) shows that an $O(N \log N)$ scaling relationship gives the best fit to the
354 data, which agrees with the theory. Further tests are described in our Applications section (§6), below.

355 6 Applications

356 Given a depression hierarchy, Fill-Spill-Merge provides an efficient method to route water across any surface while taking
357 depressions into account. Furthermore, Fill-Spill-Merge can be used to assess which depressions are most important in day-
358 to-day or seasonal changes to the hydrologic system. For example, small depressions will become flooded and spill over even
359 with relatively small amounts of water reaching them, while larger depressions may not be completely filled. These depressions
360 impact the hydrologic connectivity of the landscape (Callaghan and Wickert, 2019).




361 6.1 Field applications

373 We have compared Fill-Spill-Merge with a prior algorithm, FlowFill, at the same two sites used by Callaghan and Wickert
374 (2019): a reach of the Sangamon River in Illinois (Figure 6) and the Río Toro basin in Argentina (Figure 7). Like Fill-Spill-
375 Merge, FlowFill can be used to route water across a landscape while preserving real depressions, but the algorithm is sig-
376 nificantly slower (Table 2). The two selected study sites provide very different landscapes for testing the performance of the
377 algorithm. The Sangamon River site is located at 39.97°N, 88.72°W, in Illinois, USA. It is a low-relief, post-glacial land-
378 scape containing many closed depressions, which may impact hydrologic connectivity as a function of runoff (Lai and Anders,
379 2018). It furthermore contains a grid of roads and associated embankments whose elevations are significant when compared to
380 regional relief and impact water flow paths and storage. Callaghan and Wickert (2019) resampled the 2.5 ft (0.76 m) resolution
381 LiDAR DEM Illinois Geospatial Data Clearinghouse (2020) to 15 m resolution for analysis and manually removed several road
382 bridges using GRASS GIS to prevent artificial pooling behind these; here we use the same modified DEM to enable a direct
383 comparison between the algorithms. The Río Toro site is located mainly in Salta Province, Argentina, around 24.5°S, 65.8°W.
384 This site exhibits more rugged fluvially sculpted topography (Hilley and Strecker, 2005). Callaghan and Wickert (2019) re-
385 sampled the 12-m TanDEM-X DEM of this region (Krieger et al., 2013; Rizzoli et al., 2017) to 120 m resolution. Here we use
386 this same resampled DEM for comparison.

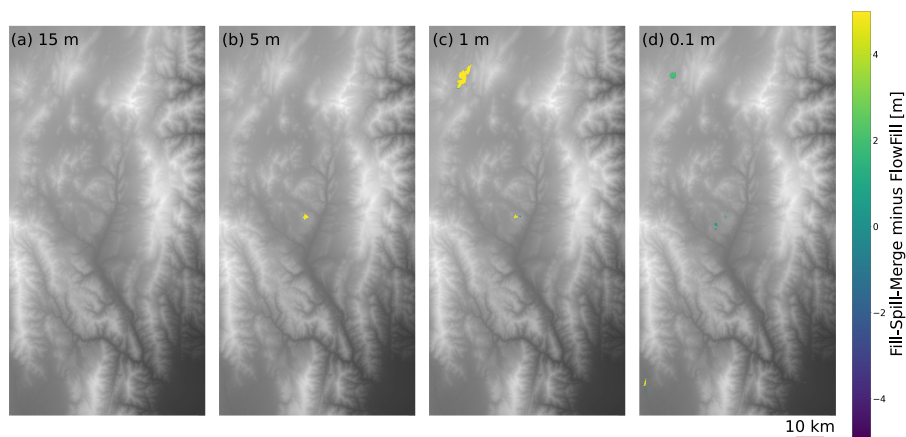
387 As shown in Table 2, wall-times using Fill-Spill-Merge ranged from 0.227–0.243 s for the Sangamon River site and 0.300–
388 0.319 s for the Río Toro site. This compares with times ranging from 20–643 s and 31–155 s, respectively, for FlowFill. These
389 times for both sites correspond to a 86–2,645× reduction in wall-time. Since FlowFill was run with 24 processors, this translates
390 to a 2,064–63,480× reduction in compute time. Considering that each of these example DEMs is quite small relative to
391 modern full-resolution LiDAR-derived elevation data sets or continental-scale 30-meter DEMs (Table 1), this speed-up and its
392 associated $O(N \log N)$ scaling provides a significant advantage for topographic analysis and solving associated problems in
393 hydrology and geomorphology.

403 Although both FlowFill and Fill-Spill-Merge route water downslope, flooding depressions based on the quantity of available
404 water, our results differ in some ways from those of FlowFill (Callaghan and Wickert, 2019). In both Figures 6 and
405 7, Fill-Spill-Merge flooded some depressions more deeply than FlowFill did, and, to a lesser extent, flooded a few depressions
406 with less water. One possible cause for this discrepancy is FlowFill’s asymptotic approach to an equilibrium water level, which
407 may prevent small volumes of water from reaching the depression to which they belong. On the other hand, depressions with
408 a narrow outlet could be especially prone to being overfilled by FlowFill because its cell-by-cell algorithm could dynamically
409 dam this outlet, routing additional water into the depression. Both of these possibilities are further linked to the fact that
410 FlowFill dynamically evolves a land-plus-water flow-routing surface, whereas Fill-Spill-Merge routes flow just over the land
411 surface. These differences make FlowFill more useful for understanding temporal changes in surface water distribution, while
412 Fill-Spill-Merge provides a more accurate snapshot of surface hydrology under equilibrium conditions.



362

363 **Figure 6.** The difference between results of Fill-Spill-Merge and FlowFill at the Sangamon River site. The values for panels (a) to (e)
364 indicate the depth of uniform runoff applied across the landscape for both algorithms. For example, in (a), each cell across the domain starts
365 with 0.2 m of surface water. Green to yellow colors indicate locations where Fill-Spill-Merge had more water, and blue to purple colors
366 indicate locations where FlowFill had more water. Differences of less than 3 mm have been masked out. Commonly, Fill-Spill-Merge retains
367 slightly more water in depressions than FlowFill does. This could be due to the iterative nature of the FlowFill algorithm, which causes it to
368 asymptotically approach the correct values. In some locations, FlowFill has retained more water. One possible reason for this is that some
369 depressions have a narrow outlet, through which Fill-Spill-Merge is able to move all water as appropriate but the cell-by-cell movement
370 of water with FlowFill can produce transient dams that reroute additional water towards these subcatchments. This DEM was prepared by
371 Lai and Anders (2018) and Callaghan and Wickert (2019) from LiDAR topographic data provided by the Illinois State Geological Survey
372 (Illinois Geospatial Data Clearinghouse, 2020).



394

395 **Figure 7.** The difference between results of Fill-Spill-Merge and FlowFill at the Río Toro site. The values for panels (a) to (d) indicate the
 396 depth of uniform runoff applied across the landscape for both algorithms. For example, in (a), each cell across the domain starts with 15 m of
 397 surface water. Green to yellow colors indicate locations where Fill-Spill-Merge had more water, and blue to purple colors indicate locations
 398 where FlowFill had more water. Differences of less than 3 mm have been masked out. In panel (a), 15 m of water was enough to fill all
 399 depressions with both algorithms, so there are no differences between the two. The most significant difference is seen in panel (c), where
 400 Fill-Spill-Merge retained additional water in a large depression. This is likely due to transient damming of its narrow inlet in FlowFill’s
 401 cell-by-cell method of moving water, which may have prevented the full volume of water from flowing into the depression. This DEM was
 402 generated with data acquired from the TanDEM-X mission (Krieger et al., 2013; Rizzoli et al., 2017).

Runoff depth [m]	Sangamon			Río Toro		
	FlowFill	FSM	Speed-up	FlowFill	FSM	Speed-Up
15	642.65	0.243	2645	154.70	0.317	488
10	626.59	0.241	2600	124.37	0.309	402
5	570.02	0.241	2365	93.56	0.300	312
1	472.33	0.241	1960	53.09	0.316	168
0.2	508.87	0.235	2165	38.30	0.316	121
0.1	464.15	0.230	2018	35.75	0.301	119
0.05	418.71	0.243	1723	33.62	0.316	106
0.01	200.81	0.227	885	32.06	0.315	102
0.001	20.12	0.235	86	30.99	0.319	97

413

414 **Table 2. Time comparison of Fill-Spill-Merge vs FlowFill.** Wall-times are in seconds comparing FlowFill (Callaghan and Wickert, 2019)
 415 parallelized across 24 cores versus Fill-Spill-Merge on a single core. Using FlowFill, wall-times increased with the depth of applied runoff
 416 and on flatter landscapes. Using FSM, wall-time is independent of depth of applied runoff and ruggedness of landscape, but increases for
 417 larger domains. FSM’s wall-times were 86–2,645 times faster than FlowFill for these examples; compute times were 2,064–63,480 times
 418 faster.



419 7 Conclusions

420 Here we leverage the depression hierarchy data structure (Barnes et al., 2020) to route flow through surface depressions in
421 a realistic, yet efficient, manner. In comparison to previous approaches, such as Jacobi iteration, the new algorithm runs in
422 log-linear time in the input size and is accompanied by extensively-commented source code. This computationally efficient
423 algorithm may help us to better understand hydrologic connectivity and water storage across hummocky land surfaces, and is
424 an important step forwards in recognising the importance of depressions as real-world features in digital elevation models.

425 *Code availability.* Complete, well-commented source code, an associated makefile, and correctness tests are available from <https://github.com/r-barnes/Barnes2020-FillSpillMerge> and Zenodo (Barnes and Callaghan, 2020).

427 *Author contributions.* KC and AW conceived the problem. RB conceived the algorithm and developed initial implementations. KC and RB
428 completed, debugged and tested the algorithm. All authors contributed to the preparation of the manuscript.

429 *Competing interests.* The authors declare that they have no conflict of interest.

430 *Acknowledgements.* RB was supported by the Department of Energy's Computational Science Graduate Fellowship (Grant No. DE-FG02-
431 97ER25308) and, through the Berkeley Institute for Data Science's PhD Fellowship, by the Gordon and Betty Moore Foundation (Grant
432 GBMF3834) and by the Alfred P. Sloan Foundation (Grant 2013-10-27).

433 KLC was supported by the National Science Foundation under grant no. EAR-1903606, the University of Minnesota Department of Earth
434 Sciences Junior F Hayden Fellowship, the University of Minnesota Department of Earth Sciences H.E. Wright Footsteps Award, and start-up
435 funds awarded to AW by the University of Minnesota.

436 Empirical tests and results were performed on XSEDE's Comet supercomputer (Townes et al., 2014), which is supported by the National
437 Science Foundation (Grant No. ACI-1053575). Portability and debugging tests were performed on the Mesabi machine at the Minnesota
438 Supercomputing Institute (MSI) at the University of Minnesota (<http://www.msi.umn.edu>).

439 The Deutsches Zentrum für Luft- und Raumfahrt (DLR) provided 12 m TanDEM-X DEM coverage of the Río Toro catchment via proposal
440 DEM_GEOL1915 awarded to Taylor Schildgen, Andrew Wickert, Stefanie Tofelde, and Mitch D'Arcy. Jingtao Lai and Alison Anders
441 provided a copy of their Sangamon River DEM.

442 This collaboration resulted from a serendipitous meeting at the Community Surface Dynamics Modeling System (CSDMS) annual meet-
443 ing, which RB had attended on a CSDMS travel grant.



444 References

- 445 Agenis-Nevers, M., Bokde, N. D., Yaseen, Z. M., and Shende, M.: GuessComp: An empirical complexity estimation in R,
446 arXiv:1911.01420v1, 2019.
- 447 Barnes, R.: Parallel non-divergent flow accumulation for trillion cell digital elevation models on desktops or clusters, *Environmental Mod-*
448 *elling & Software*, 92, 202–212, <https://doi.org/10.1016/j.envsoft.2017.02.022>, 2017.
- 449 Barnes, R.: Accelerating a fluvial incision and landscape evolution model with parallelism, *Geomorphology*, 330, 28–39,
450 <https://doi.org/10.1016/j.geomorph.2019.01.002>, 2019.
- 451 Barnes, R. and Callaghan, K.: Fill-Spill-Merge Source Code, <https://doi.org/10.5281/zenodo.3755142>, 2020.
- 452 Barnes, R., Lehman, C., and Mulla, D.: Priority-flood: An optimal depression-filling and watershed-labeling algorithm for digital elevation
453 models, *Computers & Geosciences*, 62, 117 – 127, <https://doi.org/10.1016/j.cageo.2013.04.024>, 2014.
- 454 Barnes, R., Callaghan, K., and Wickert, A.: Computing water flow through complex landscapes, Part 2: Finding hierarchies in depressions
455 and morphological segmentations, *Earth Surface Dynamics*, <https://doi.org/10.5194/esurf-2019-34>, 2020.
- 456 Braun, J. and Willett, S. D.: A very efficient O(n), implicit and parallel method to solve the stream power equation governing fluvial incision
457 and landscape evolution, *Geomorphology*, 180–181, 170–179, <https://doi.org/10.1016/j.geomorph.2012.10.008>, 2013.
- 458 Breckenridge, A. and Johnson, T. C.: Paleohydrology of the upper Laurentian Great Lakes from the late glacial to early Holocene, *Quaternary*
459 *Research*, 71, 397–408, <https://doi.org/10.1016/j.yqres.2009.01.003>, 2009.
- 460 Cabrol, N. A. and Grin, E. A.: Distribution, classification, and ages of Martian impact crater lakes, *Icarus*, 142, 160–172, 1999.
- 461 Callaghan, K. L. and Wickert, A. D.: Computing water flow through complex landscapes – Part 1: Incorporating depressions in flow routing
462 using FlowFill, *Earth Surface Dynamics*, 7, 737–753, <https://doi.org/10.5194/esurf-7-737-2019>, 2019.
- 463 Fan, Y., Li, H., and Miguez-Macho, G.: Global Patterns of Groundwater Table Depth, *Science*, 339, 940–943,
464 <https://doi.org/10.1126/science.1229881>, 2013.
- 465 Fenner, T. I. and Loizou, G.: Loop-free Algorithms for Traversing Binary Trees, *BIT*, 24, 33–44, <https://doi.org/10.1007/BF01934513>, 1984.
- 466 GDAL Development Team: GDAL - Geospatial Data Abstraction Library, Open Source Geospatial Foundation, available at <http://www.gdal.org>, 2016.
- 467
- 468 GEBCO: General Bathymetric Chart of the Oceans (GEBCO), GEBCO_08 grid, version 20100927, <http://www.gebco.net>, 2010.
- 469 Hilley, G. E. and Strecker, M. R.: Processes of oscillatory basin filling and excavation in a tectonically active orogen: Quebrada del Toro
470 Basin, NW Argentina, *GSA Bulletin*, 117, 887–901, <https://doi.org/10.1130/B25602.1>, 2005.
- 471 Illinois Geospatial Data Clearinghouse: Illinois Height Modernization (ILHMP), <https://clearinghouse.isgs.illinois.edu/data/elevation/illinois-height-modernization-ilhmp-lidar-data>, 2020.
- 472
- 473 Jenson, S. and Domingue, J.: Extracting Topographic Structure from Digital Elevation Data for Geographic Information System Analysis,
474 *Photogrammetric Engineering and Remote Sensing*, 54, 1–5, [https://doi.org/0099-1112/88/5411-1593\\$02.25/0](https://doi.org/0099-1112/88/5411-1593$02.25/0), 1988.
- 475 Krieger, G., Zink, M., Bachmann, M., Bräutigam, B., Schulze, D., Martone, M., Rizzoli, P., Steinbrecher, U., Antony, J. W., De Zan, F., et al.:
476 TanDEM-X: A radar interferometer with two formation-flying satellites, *Acta Astronautica*, 89, 83–98, 2013.
- 477 Lai, J. and Anders, A. M.: Modeled Postglacial Landscape Evolution at the Southern Margin of the Laurentide Ice Sheet: Hydrological
478 Connection of Uplands Controls the Pace and Style of Fluvial Network Expansion, *Journal of Geophysical Research: Earth Surface*, 123,
479 967–984, <https://doi.org/10.1029/2017JF004509>, 2018.



- 480 Li, S., MacMillan, R., Lobb, D. A., McConkey, B. G., Moulin, A., and Fraser, W. R.: Lidar DEM error analyses and topo-
481 graphic depression identification in a hummocky landscape in the prairie region of Canada, *Geomorphology*, 129, 263–275,
482 <https://doi.org/10.1016/j.geomorph.2011.02.020>, 2011.
- 483 Lindsay, J. and Creed, I.: Removal of artifact depressions from digital elevation models: towards a minimum impact approach, *Hydrological*
484 *processes*, 19, 3113–3126, <https://doi.org/10.1002/hyp.5835>, 2005a.
- 485 Lindsay, J. B.: Efficient hybrid breaching-filling sink removal methods for flow path enforcement in digital elevation models: Efficient Hybrid
486 Sink Removal Methods for Flow Path Enforcement, *Hydrological Processes*, <https://doi.org/10.1002/hyp.10648>, 2015.
- 487 Lindsay, J. B.: Efficient hybrid breaching-filling sink removal methods for flow path enforcement in digital elevation models: Efficient
488 Hybrid Sink Removal Methods for Flow Path Enforcement, *Hydrological Processes*, 30, 846–857, <https://doi.org/10.1002/hyp.10648>,
489 <http://doi.wiley.com/10.1002/hyp.10648>, 2016.
- 490 Lindsay, J. B. and Creed, I. F.: Removal of artifact depressions from digital elevation models: Towards a minimum impact approach, *Hydro-*
491 *logical Processes*, 19, 3113–3126, <https://doi.org/10.1002/hyp.5835>, 2005b.
- 492 Mark, D.: *Modelling in Geomorphological Systems*, chap. Network models in geomorphology, pp. 73–97, John Wiley & Sons, 1988.
- 493 Martz, L. W. and Garbrecht, J.: The treatment of flat areas and depressions in automated drainage analysis of raster digital elevation models,
494 *Hydrological Processes*, 12, 843–855, [https://doi.org/10.1002/\(SICI\)1099-1085\(199805\)12:6<843::AID-HYP658>3.0.CO;2-R](https://doi.org/10.1002/(SICI)1099-1085(199805)12:6<843::AID-HYP658>3.0.CO;2-R), 1998.
- 495 Martz, L. W. and Jong, E. d.: CATCH: A FORTRAN program for measuring catchment area from digital elevation models, *Computers and*
496 *Geosciences*, 14, 627–640, [https://doi.org/10.1016/0098-3004\(88\)90018-0](https://doi.org/10.1016/0098-3004(88)90018-0), 1988.
- 497 MNDNR - Minnesota Department of Natural Resources: Lake Bathymetric Outlines, Contours, Vegetation, and DEM, [https://gisdata.mn.](https://gisdata.mn.gov/dataset/water-lake-bathymetry)
498 [gov/dataset/water-lake-bathymetry](https://gisdata.mn.gov/dataset/water-lake-bathymetry), 2014.
- 499 MNGEO - Minnesota Geospatial Information Office: LiDAR Elevation Data for Minnesota, [http://www.mngeo.state.mn.us/choose/elevation/](http://www.mngeo.state.mn.us/choose/elevation/lidar.html)
500 [lidar.html](http://www.mngeo.state.mn.us/choose/elevation/lidar.html), 2019.
- 501 O’Callaghan, J. and Mark, D.: The extraction of drainage networks from digital elevation data, *Computer Vision, Graphics, and Image*
502 *Processing*, 28, 323–344, [https://doi.org/10.1016/S0734-189X\(84\)80011-0](https://doi.org/10.1016/S0734-189X(84)80011-0), 1984.
- 503 Reheis, M.: Highest Pluvial-Lake Shorelines and Pleistocene Climate of the Western Great Basin, *Quaternary Research*, 52, 196–205,
504 <https://doi.org/10.1006/qres.1999.2064>, 1999.
- 505 Riddick, T., Brovkin, V., Hagemann, S., and Mikolajewicz, U.: Dynamic hydrological discharge modelling for coupled climate model
506 simulations of the last glacial cycle: the MPI-DynamicHD model version 3.0, *Geoscientific Model Development*, 11, 4291–4316,
507 <https://doi.org/10.5194/gmd-11-4291-2018>, 2018.
- 508 Rizzoli, P., Martone, M., Gonzalez, C., Wecklich, C., Tridon, D. B., Bräutigam, B., Bachmann, M., Schulze, D., Fritz, T., Huber, M., et al.:
509 Generation and performance assessment of the global TanDEM-X digital elevation model, *ISPRS Journal of Photogrammetry and Remote*
510 *Sensing*, 132, 119–139, 2017.
- 511 Schwanghart, W. and Scherler, D.: Bumps in river profiles: uncertainty assessment and smoothing using quantile regression techniques, *Earth*
512 *Surface Dynamics*, 5, 821–839, <https://doi.org/10.5194/esurf-5-821-2017>, 2017.
- 513 Soille, P.: Optimal removal of spurious pits in grid digital elevation models, *Water Resources Research*, 40, 1–9,
514 <https://doi.org/10.1029/2004WR003060>, 2004.
- 515 Soille, P., Vogt, J., and Colombo, R.: Carving and adaptive drainage enforcement of grid digital elevation models, *Water Resources Research*,
516 39, 1366, <https://doi.org/10.1029/2002WR001879>, 2003.



- 517 Towns, J., Cockerill, T., Dahan, M., Foster, I., Gaither, K., Grimshaw, A., Hazlewood, V., Lathrop, S., Lifka, D., Peterson, G. D., et al.:
518 XSEDE: accelerating scientific discovery, *Computing in Science & Engineering*, 16, 62–74, 2014.
- 519 Wallis, C., Watson, D., Tarboton, D., and Wallace, R.: Parallel flow-direction and contributing area calculation for hydrology analysis in
520 digital elevation models, in: *Proceedings of the 2009 International Conference on Parallel and Distributed Processing Techniques and*
521 *Applications*, Las Vegas, Nevada, <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.158.2864>, 2009.
- 522 Wickert, A. D.: Reconstruction of North American drainage basins and river discharge since the Last Glacial Maximum, *Earth Surface*
523 *Dynamics*, 4, 831–869, <https://doi.org/10.5194/esurf-4-831-2016>, 2016.