

1 Parallel Variable-Resolution Bathymetric Estimation 2 with Static Load Balancing

3 B.R. Calder^{a,*}

4 ^a*Center for Coastal and Ocean Mapping and NOAA-UNH Joint Hydrographic Center,*
5 *University of New Hampshire, Durham, NH 03824, USA*

6 Abstract

7 A method for partitioning a large computation task (direct, variable resolution
8 bathymetric grid construction from raw observations) into thread-parallel code
9 is described. Based on the data density estimated for the first pass of the CHRT
10 algorithm, this algorithm statically partitions the estimation task into spatially
11 distinct blocks of approximately equal total data observation count so that each
12 can be executed in parallel and be expected to complete approximately concur-
13 rently. No communication between blocks or further load balancing is therefore
14 required. A branch-and-bound algorithm is used to control the complexity of
15 the partitioning task, but the computation time increases significantly as more
16 partitions are required, leading to a degree of diminishing returns for allocating
17 further computational resources and suggesting alternative approaches for high
18 thread-count systems. Speed-up of the algorithm over a pair of test datasets
19 (using real-world hydrographic survey data) shows that the performance con-
20 sistentlly improves with the number of computational tasks assigned, initially
21 (super-) linearly, although ultimately sub-linearly as other resource sharing lim-
22 itations take over. An overall speedup of 4.1 times is demonstrated with a
23 quad-core single-processor workstation.

24 *Keywords:* Parallel processing, CHRT, CUBE, Data-driven Estimation,
25 Branch and Bound, Bathymetric Data Processing, Surface Estimation

26 1. Introduction

27 Bathymetry is often a base layer in marine spatial modelling, providing a
28 important constraint on the physical environment (e.g., defining the waveguide
29 for acoustic propagation studies) and driving other analyses. The reconstruc-
30 tion of a best-estimate of depth (or, more generally, any scalar field) within a
31 given area based on remote-sensed observations (Krishnan et al., 2010; Hofierka
32 et al., 2017) is therefore an important problem with many practical applications,
33 including ocean mapping, geophysical modelling, coastal zone management, and
34 nautical charting.

35 The problem is computationally challenging. The datasets are often large
36 (order $10^9 - 10^{10}$ observations), and the algorithms can be complex due to

*E-mail: brc@ccom.unh.edu, Tel.: +1 603 862 0526. Declarations of Interest: None.
Preprint submitted to Elsevier *September 6, 2018*

37 dataset features such as observational blunders (Calder and Mayer, 2003; Debese
38 et al., 2012; Isenburg et al., 2006). The datasets may also have a spatial-varying
39 data density, requiring spatial adaptation of reconstruction resolution to avoid
40 spatial aliasing or over-smoothing. (In hydrographic practice, over-smoothing
41 could result in missed navigationally significant objects, which are a primary
42 concern.) In many cases, the data density is approximately a function of the
43 water depth, so deeper areas can only be reconstructed at significantly lower
44 resolution; these changes can happen within very short distances, for example
45 in fjord-like environments.

46 Efficient computation of the reconstruction is therefore essential. In addition
47 to minimising processing time, fast computation allows for more advanced
48 algorithms to be built around the basic estimation task. For example, it can
49 be difficult to compensate for the effects of slope in CHRT (the estimation algorithm
50 considered here) *a priori* because there is no good estimate of slope
51 until the reconstruction is computed, but that reconstruction is biased by lack
52 of slope compensation. Iterating to solution is plausible, but if the algorithm
53 is expensive to compute, the iterations might take sufficiently long as to render
54 the method ineffective.

55 Computing a reconstruction in parallel is therefore advantageous. Many
56 algorithms, however, are either global (i.e., require all data in an area to proceed,
57 for example the surface fitting of Debese et al. (2012)), or non-local (i.e., need to
58 interact with nearby estimation sites to complete the estimation, for example the
59 continuous spline in tension method of Smith and Wessel (1990)), which can
60 make them difficult to segment for parallel implementation. Sending packets
61 of observations to different threads for update of a common data structure
62 generally requires a significant level of complexity in the data structure to allow
63 simultaneous access, while splitting spatially has difficulties in ensuring that
64 the sub-tasks are well balanced. A solution which is well balanced, does not
65 require memory locking, and can be scaled to many computational resources is
66 therefore key.

67 Load-balancing, or the more general case of splitting a non-even workload
68 computational domain over a number of computational resources so as to achieve
69 some metric, is a common problem, and has therefore received much attention
70 in the literature. The problem of finding the optimal general partition in two
71 dimensions is known to be NP-Hard (Khanna et al., 1998), or NP-Complete
72 in some cases such as the Generalized Block Distribution (Grigni and Manne,
73 1996); even achieving a bound on the performance within a factor of two is
74 NP-Hard (Aspvall et al., 2001). Consequently, most of the research on the
75 matter has revolved around finding better approximations to the problem, and
76 lowering the upper bound on performance (see, e.g., (Manne and Sørensen, 1996),
77 (Aspvall et al., 2001), (Berman et al., 2001), (Lorys and Paluch, 2003), (Saule
78 et al., 2012), etc.).

79 The most general case of partitioning would allow for arbitrary segments to
80 be assigned to a computational resource; in keeping with Tobler's Law (Tobler,
81 1970), however, most solutions focus on assigning rectangular areas in order to
82 maintain advantages of spatial locality in the computations. Most work has been

83 done on recursive partitions (Berger and Bokhari, 1987), the rectilinear parti-
84 tion (Nichol, 1991), also known as the Generalized Block Distribution (Grigni
85 and Manne, 1996), and the jagged partition (Manne and Sørenvik, 1996), also
86 known as the Semi-generalized Block Distribution (see (Saule et al., 2012) for a
87 good overview), although there are many more potential partitioning schemes.

88 The complexity and performance of the best known algorithms for the var-
89 ious approximations vary widely. Saule et al. (2012) compare a variety of al-
90 gorithms, with different heuristics designed to achieve better performance, in-
91 cluding hierarchical sub-division, rectilinear division, and jagged partitions, for
92 which polynomial time algorithms are available. They conclude that their jagged
93 partition variant, or a hierarchical subdivision may achieve better load balanc-
94 ing than other algorithms while still being runtime efficient, while a combination
95 of algorithms (a “hybrid”) can improve on achievable balance even further. (In-
96 triguingly, the best case load imbalances reported are of similar magnitude to
97 those reported here.)

98 These approaches, while efficient, are more constrained than is required in
99 the case presented here, primarily due to a basic assumption that communica-
100 tions costs between the segments of the partition are important. Here, however,
101 given the observations, each segment of the partition can compute independ-
102 ently its part of the solution to the overall estimation problem, so there is no
103 limitation to the arrangement of the segments, and a more general solution can
104 be attempted. Furthermore, the focus here is on splitting the overall task over
105 a relatively small number of computational resources, since the primary goal
106 is a multithreaded, single CPU solution, since it remains the most commonly
107 available computational resource for most users in the field (Qin and Zhan,
108 2012). Consequently, the more sophisticated heuristic-based algorithms are not
109 required, and an optimal solution of the (constrained) partitioning problem can
110 be used, simplifying the implementation. (The problem of scaling to higher
111 numbers of computational resources is considered in Section 4.)

112 In this paper, therefore, a spatial partitioning algorithm is proposed for the
113 CHRT (CUBE with Hierarchical Resolution Techniques) algorithm (Calder and
114 Rice, 2017) which takes advantage of the structure of CHRT to ensure that each
115 computational resource can operate independently of the others without com-
116 munication or interlocks so long as they have global access to all of the observa-
117 tions. (Section 2.1 has an outline of the CHRT algorithm.) The algorithm uses
118 the data density estimates computed during the first pass of the CHRT algorithm
119 to drive the partition, which allows for the partition segments to contain ap-
120 proximately the same number of observations, and consequently to have nearly
121 uniform computation time. Our goal here is not to minimise the maximum cost
122 for any segment (c.f. (Saule et al., 2012), or (Muthukrishnan and Suel, 2005))
123 but to have even load across all of the computational resources, hence keeping
124 them all busy for the minimal time with highest efficiency. Operating system file
125 caching assists in delaying IO-limited performance (i.e., where recently used files
126 remain in memory, and therefore are not subject to spinning-disc latency), and a
127 branch-and-bound evaluator allows the partition to be computed efficiently. Use
128 of the partitioning algorithm allows ready extension of the CHRT algorithm to

129 a multi-threaded implementation, with consequent performance improvement.

130 The remainder of the paper outlines the relevant features of the CHRT al-
131 gorithm that support the partitioning algorithm, its implementation, and the
132 performance improvements achieved using commodity single-processor worksta-
133 tion hardware. Finally, some perspectives on the ability of the algorithm to be
134 scaled, and generalised to a distributed (i.e., network-connected) implementa-
135 tion, are offered.

136 2. Methods

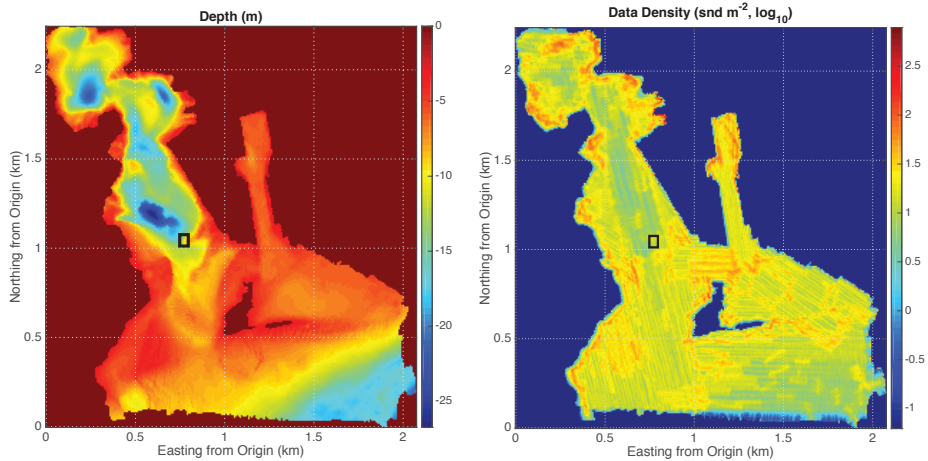
137 2.1. Core Estimator

138 The CHRT (CUBE with Hierarchical Resolution Techniques) algorithm (Calder
139 and Rice, 2017), a development of the CUBE (Combined Uncertainty and Bathy-
140 metry Estimator) algorithm (Calder and Mayer, 2003) was used as the basis for
141 the current work. The CHRT algorithm was developed to estimate variable reso-
142 lution depths from raw observational data based on the premise that in regions
143 where there is higher data density it should be possible to reconstruct with
144 smaller sample spacings, giving higher resolution reconstructions of the surface.
145 The algorithm starts with a low-resolution virtual tile (Yildirim et al., 2015)
146 grid across the area of interest, and at each grid node estimates the data den-
147 sity of the observations. A piecewise constant sample spacing (PCSS) grid is
148 then constructed by replacing each low-resolution grid cell with a regular grid
149 at the sample spacing determined by the data density, after which a variable
150 resolution depth reconstruction can be computed in a second pass. Figure 1
151 shows an example of the first pass of the algorithm applied to a hydrographic
152 survey in Woods Hole, MA.

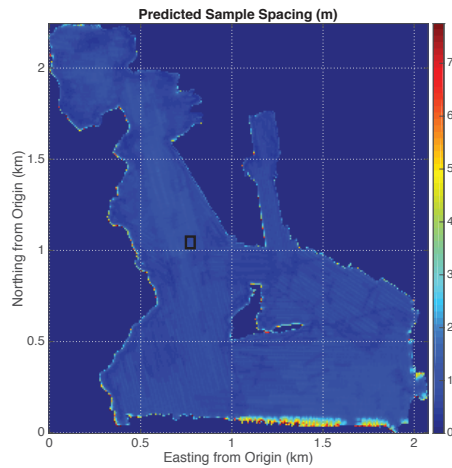
153 A basic problem for any parallel algorithm is how to split the task into man-
154 ageable sub-tasks. The low-resolution grid used in CHRT allows for a relatively
155 simple solution to this problem since the refinement grids established after the
156 first pass of the algorithm are by design constrained to lie entirely within their
157 parent cell, Figure 2. Consequently, given the observations that contribute to
158 the cell (which may include some immediately adjacent in order to avoid edge
159 effects), the computation of each cell is independent of the others, and there-
160 fore can be processed without any communication or interlock, and there is no
161 requirement for “ghost cell” edge buffers (Tesfa et al., 2011). Any sub-group of
162 cells can therefore be assigned to any available computational resource, so long
163 as it has access to all of the observations. This decomposition of the base algo-
164 rithm avoids having to design a variant for parallel implementation, with all of
165 the associated development and maintenance costs (Hofierka et al., 2017). The
166 CHRT Conformance Test Suite (Calder and Plumlee, 2017) ensures equivalence
167 of serial and parallel computation.

168 2.2. Partitioning Scheme

169 For the CHRT algorithm, the processing cost is reasonably approximated by
170 the number of observations that have to be assimilated at a particular recon-
171 struction location. A plausible load balancing partitioning scheme is therefore



(a) Low resolution bathymetry (m); the gross features are clear, but significant objects (e.g., mooring blocks, pilings) require the variable-resolution data to resolve. (b) Data density (snd m^{-2} , log scale). Note the hole in the middle caused by a very shallow area that was not surveyed for safety.



(c) Predicted sample spacing (m) for refinements.

Figure 1: Example of the CHRT algorithm applied to a NOAA survey in Woods Hole, MA, showing (a) first-pass low-resolution depth estimate, (b) data density estimate (note logarithmic scale), and (c) estimated sample spacing for each low-resolution cell. Black rectangle is shown in detail in Figure 2. This figure is reproduced from Calder and Rice (2017).




Figure 2: Example of variable resolution depth reconstruction, and the location of variable resolution reconstruction points derived from data density estimates during the first pass for the data indicated in the black rectangle in Figure 1. Each refinement grid is constrained to be entirely within the parent low-resolution grid cell (here, at 8 m intervals). Colours represent the estimated depths; white dots mark the locations of the variable resolution estimation points. Labels on the geographic axes mark the edges of the low-resolution cells containing the refined (white dot) grids. This figure is reproduced from Calder and Rice (2017).

172 to split the overall area to be processed into sub-areas that contain approxi-
 173 mately the same number of observations. (Alternatives, such as partitioning by
 174 input files and then recombining partial grids, or dynamic partitioning of sub-
 175 groups (Yıldırım et al., 2015) would lead to serialized code or higher communi-
 176 cations costs, respectively.) In order to facilitate this, the partition algorithm
 177 assumes that a spatial observation density estimate is available from the first
 178 pass of the CHRT algorithm (this is a core component of the base algorithm).

179 Let the data density estimates be arranged in a grid, $\rho(u, v), 0 \leq u < U, 0 \leq$
 180 $v < V, (u, v) \in \mathbb{Z}^2$, with N total observations in the dataset spread through
 181 the area. The goal is to find an optimal partition of the overall domain $S_0 =$
 182 $\{(u, v) : (u, v) \in [0, U) \times [0, V)\}$ into C segments, one for each computational
 183 resource, each containing an equal number of observations. The grid could
 184 be partitioned into arbitrarily-shaped groups of cells with equal numbers of
 185 observations, but to contain the complexity consider admitting only north-south
 186 or east-west segment boundaries (Berger and Bokhari, 1987).

187 The algorithm solves this problem recursively over the general segment of
 188 the grid, $S = \{(u, v) : (u, v) \in [u_0, u_1 - 1] \times [v_0, v_1 - 1], 0 \leq u_0 < u_1 < U, 0 \leq v_0 <$
 189 $v_1 < V\}$, which is to be split into $C_S \leq C$ segments each of N_S/C_S observations,
 190 where $N_S \leq N$ is the observation count for S . The recursion root is $S = S_0,$
 191 $C_S = C, N_S = N$.

192 Consider first an algorithm that enumerates all possible partitions. Each par-
 193 tition line is placed to split S into sub-segments of some multiple $cN_S/C_S, 1 \leq$
 194 $c < C_S$ of the observations in the segment, Figure 3. Given an area A for each
 195 cell in the domain, the partition line is placed at $\{u_c, 1 \leq c < C_S\}$ where

$$196 \quad u_c = \max_{u_0 \leq x < u_1} \left\{ x : \sum_{u=u_0}^x \sum_{v=v_0}^{v_1-1} \rho(u, v) A < cN_S/C_S \right\} \quad (1)$$

197 or at $\{v_c, 1 \leq c < C_S\}$ where

$$198 \quad v_c = \max_{v_0 \leq y < v_1} \left\{ y : \sum_{v=v_0}^y \sum_{u=u_0}^{u_1-1} \rho(u, v) A < cN_S/C_S \right\} \quad (2)$$

199 for north-south and east-west partition lines, respectively, giving $2(C_S - 1)$ po-
 200 tential partial segmentations.

201 For north-south partitions, these potential positions split S into

$$202 \quad S_L = \{(u, v) : (u, v) \in [u_0, u_c] \times [v_0, v_1 - 1]\} \quad (3)$$

203 and

$$204 \quad S_R = \{(u, v) : (u, v) \in [u_c + 1, u_1 - 1] \times [v_0, v_1 - 1]\} \quad (4)$$

205 so that $S = S_L \cup S_R$ and $S_L \cap S_R = \emptyset$, and equivalently for east-west partitions.
 206 The algorithm can then be applied recursively to S_L and S_R , each now of N_{S_L}
 207 and N_{S_R} observations, respectively, with a target of $C_S - 1$ computational re-
 208 sources assigned to them. The recursion terminates when $C_S = 1$. Since each

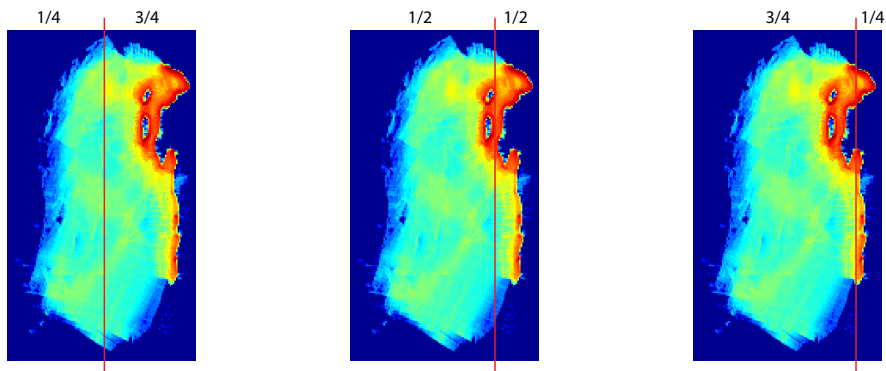


Figure 3: Example of potential position points for the first stage of the partitioning algorithm with four computational resources, where the goal is to split off one, two or three quarters of the observations (with regions of two or three quarters being split in later stages of algorithm); the background images are the data density estimated from the observations in the Ernest Sound, AK test dataset (see Section 3.1). Note the significantly larger area associated with the $1/4$ position (left) image compared to the $3/4$ position (right) image, due to the significantly lower data density to the west of the partition line in the former case. Solving for a different number of computational resources would have different partition points as the algorithm attempts to split off different sized portions of the data.




Figure 4: Example of a partial hierarchical tree for four computational resources, indicating some of the potential combinations of north-south and east-west partitions applied in sequence to split the problem into four segments, each of one quarter of the problem. Note the practical redundancy in many of the logically separate partitions, which can reduce the efficiency of the search if the whole tree is enumerated.

209 partition can be placed either north-south or east-west on each occasion, this al-
 210 gorithm naturally leads to a tree of potential partition schemes, Figure 4. Many
 211 of the partitions generated are logically separate (i.e., the order in which the
 212 partition lines were generated are different) but practically the same (i.e., the
 213 resulting segments are identical). This can lead to implementation efficiencies,
 214 a topic pursued in the following section.

215 In principle, this algorithm can be applied from S_0 to enumerate all leaves
 216 of the partition tree. Due to the granularity of the low-resolution cells, it is
 217 unlikely that any given partition will exactly split the problem in C segments
 218 of N/C observations. The viability of the different leaves of the partition tree
 219 can therefore be assessed according to how closely they achieve this goal, with
 220 the closest match being the preferred solution. An example of an 8-partition
 221 applied to the data in Figure 3 is shown in Figure 5.

222 2.3. Partitioning Algorithm Implementation

223 In theory, the partition that best matches the ideal, even, distribution of
 224 observations could be determined by simply enumerating the tree of potential
 225 partitions. The first stage of splitting has $2(C - 1)$ potential splits; the second
 226 has $2(C - 2)$, and so on, for a total of $2^{C-1}(C - 1)!$ potential solutions. Some
 227 reduction in effort could be obtained by exploiting similarities in the potential
 228 partitions, but for any reasonable target number of segments the size of the tree
 229 rapidly makes a full enumeration intractable: for $C = 8$, for example, a reason-
 230 able choice for quad-core hyper-threaded processors, a total of 645,120 potential




Figure 5: Example of a partition for eight computational resources applied to the Ernest Sound data (Figure 3). The segments selected by the algorithm are shown as white outlines (with semi-transparent colours) over the sounding count (in log-scale). Note that the segments are shown with boundaries slightly separated for clarity; in reality, they completely tile the computational area.

231 solutions would have to be enumerated; at $C = 16$, the total is 4.285×10^{16} .
 232 The goal is still to evaluate the whole tree, however, so the algorithm applies
 233 the “branch and bound” technique (Land and Doig, 1960) to avoid evaluating
 234 inefficient branches of the tree as often as possible, and applies heuristics to
 235 attempt to accelerate the process.

236 Consider the situation at any node in the tree of Figure 4. Assume that
 237 for any individual segment S there is a cost function $P(S, N, C)$ that rep-
 238 represents the penalty for not matching the nominal ideal observation count of
 239 N/C observations per segment. If $C_S = 1$ (i.e., the segment represents the
 240 best approximation to a single quantum of observations), then the cost can
 241 be evaluated directly; otherwise, the potential refinements of the segment are
 242 $R = \{v_1, \dots, v_{C_S-1}, h_1, \dots, h_{C_S-1}\}$ where the v_i represent north-south, and h_i
 243 east-west partitions, respectively, computed according to (1)–(2). Each refine-
 244 ment induces a pair of segments $(S_P(r), S_S(r))$, $r \in R$ according to (3)–(4),
 245 corresponding to the segment prior to, and subsequent to, the partition lo-
 246 cation, respectively. The overall penalty for each potential refinement of the
 247 segment can therefore be computed as

$$\begin{aligned}
 248 \quad P(r) &= P(S_P(r) \cup S_S(r), N, C) \\
 249 \quad &= P(S_P(r), N, C) + P(S_S(r), N, C), \tag{5}
 \end{aligned}$$

251 with the individual penalties being evaluated recursively. Clearly, the optimal
 252 refinement is

$$253 \quad r^* = \arg \min_r P(r), \tag{6}$$

254 and the parent node therefore has a “best known partition” penalty of $P(r^*)$.

255 At each node in the tree, the partitioning decisions made further up the tree
 256 lead to a penalty which the algorithm has already assumed in order to get to the
 257 decision point represented by the node. An allowable penalty can therefore be
 258 passed to each node by its parent, indicating the maximum penalty remaining to
 259 the branch for any refinement to be viable in comparison with the best available
 260 refinement elsewhere; testing against this limit can therefore reduce the number
 261 of evaluations that need to be attempted.

262 Let α_S be the available penalty provided to the node for segment S ; to
 263 seed the recursion, let $\alpha_{S_0} \rightarrow \infty$, or in practice the maximum value available.
 264 Clearly, if $P(S_P(r), N, C) > \alpha_S$, the proposed refinement is not viable, irrespec-
 265 tive of $P(S_S(r), N, C)$, and the evaluation of potential refinements of S_S need
 266 not be computed (and vice versa). (Observe that if $P(S_P(r), N, C) < \alpha_S$, then
 267 the bound for evaluating $S_S(r)$ should be $P(S_P(r), N, C) - \alpha_S$, which can help
 268 to reject more potential refinements, further reducing the computation cost.)
 269 This bound can be incrementally tightened by observing that each refinement
 270 evaluated can provide a better target if $P(r) < \alpha_S$. Therefore, define

$$271 \quad \alpha(0) = \alpha_S \tag{7}$$

$$272 \quad \alpha(r) = \min(\alpha(r-1), P(r)) \tag{8}$$

274 so that the target that candidate refinements have to better tightens as good
 275 refinements are determined. The ordering in which refinements are attempted
 276 is essentially arbitrary, but the choice of which to evaluate first for most effi-
 277 cient evaluation is not. Given that one side or other of the partition might be
 278 eliminated from consideration after the other is evaluated, it is advantageous to
 279 evaluate first the side that is shallowest (i.e., with smallest N_S).

280 The efficiency of the pruning algorithm is maximised if the algorithm can es-
 281 tablish a plausible solution (i.e., one close to the optimal) early in the sequence
 282 of evaluations, since it will lead to many more branches being pruned more
 283 quickly. There is no way to predict where a “good” solution would lie *a priori*,
 284 but a useful heuristic is to observe that splitting off a segment of N/C obser-
 285 vations early in the sequence is unlikely to provide a good solution, since the
 286 split position can only be adjusted by a whole row or column of low-resolution
 287 cells, which can contain many observations. If, on the other hand, the first split
 288 breaks the area approximately in half (c.f. (Berger and Bokhari, 1987)), then
 289 any error in the observation count can be amortised over all of the remaining
 290 splits. The algorithm therefore starts at the $c = \lfloor C/2 \rfloor$ position, and moves
 291 outward towards either extreme, swapping sides after each step (i.e., evaluating
 292 at $\lfloor C/2 \rfloor, \lfloor C/2 \rfloor - 1, \lfloor C/2 \rfloor + 1, \lfloor C/2 \rfloor - 2, \dots$).

293 Evaluating the number of observations within a proposed segment is the
 294 most expensive part of the partition computation. The count of observations
 295 within each low-resolution cell is, however, fixed. Therefore it is possible to
 296 utilise a variant of the summed-area table technique (Crow, 1984), also known
 297 as a prefix sum table (Ladner and Fischer, 1980), to cache cumulative sums and
 298 hence significantly improve the computation time.

299 There are a number of plausible definitions for the cost function. Here, a
 300 simple comparison against the nominal observation count per computational
 301 resource is used,

$$302 \quad P(S, N, C) = |N_S - N/C|. \quad (9)$$

303 In the simplest case, $N_S = \sum_{(u,v) \in S} \rho(u,v)A$. Evaluation of cost functions in
 304 bounded arithmetic (i.e., where there is a maximum representable cost, P_{\max})
 305 requires some care. In particular, saturation addition is required, so that if

$$306 \quad P'(S, N, C) = \min(P_{\max}, P(S, N, C)) \quad (10)$$

307 is the bounded arithmetic representation of the cost function, then

$$308 \quad P'(r) = P'(S_P, N, C) + \min(P'(S_S, N, C), P_{\max} - P'(S_P, N, C)). \quad (11)$$

311 The CHRT algorithm utilises some observations from just outside an assigned
 312 computation domain so as to avoid any edge effects in the estimation. If these
 313 are ignored, then the computational cost of processing a segment will be un-
 314 derestimated, potentially significantly in shallow areas with dense observations.
 315 Let S_E be the annulus, one low-resolution cell wide, around segment S , with
 316 N_E observations. The CHRT algorithm, by default, uses observations only out

317 to $\sqrt{2}W$ from the segment boundary for cells of width (and height) W m. The
 318 annulus therefore provides approximately $\sqrt{2}N_E$ observations (simplifying for
 319 the corner cells), and the total effective number of observations used in (9) is

$$320 \quad N_S = \sqrt{2} \sum_{(u,v) \in \mathcal{S}_E} \rho(u,v)A + \sum_{(u,v) \in \mathcal{S}} \rho(u,v)A. \quad (12)$$

321 *2.4. Parallel Estimator Implementation*

322 Although the core estimation algorithm is the same for serial and parallel
 323 computation, some care in staging is required. The CHRT algorithm supports
 324 virtual tiles, with memory-mapped files that are demand paged with least-
 325 recently-used cache replacement. The memory-mapped structures may cross
 326 segment boundaries, however, and to avoid interlocks it is therefore necessary
 327 for each computational resource to have a separate copy of the results of the
 328 first pass of the algorithm for the tiles associated with the segment assigned.
 329 Knowledge of the segment bounds allows this computation to be done *a priori*,
 330 and the parallel wrapper code can pre-copy the required files along with the
 331 base metadata for the data structure. After the initial configuration, the com-
 332 putational resources are independently scheduled as separate threads within the
 333 main process.

334 A producer-consumer pattern is used with one consumer implementing the
 335 estimator for each segment of the partition. The producer implements the Com-
 336 mand pattern (Gamma et al., 1994) by constructing work packages for each
 337 stage of the computation, derived from an abstract interface, that are queued
 338 for all of the worker threads to execute. A modified barrier synchronisation
 339 pattern (Wilkinson and Allen, 2005) allows for the threads to be marshalled,
 340 indicating to the producer that all required computations have been completed
 341 (e.g., so that the client interface can determine when it is safe to request a
 342 reconstruction take place).

343 *2.5. Partial Result Reassembly*

344 As with the partitioning problem, reassembly of the partial results from each
 345 computational resource is made simpler because the segments are aligned with
 346 low-resolution cell boundaries. Each computational resource can therefore, on
 347 demand, generate its partial result and write them into the shared output data
 348 structure for the overall result without interlocks.

349 In theory, the partial reconstruction computations for each segment could
 350 be overlapped with any remaining primary computation in order to avoid any
 351 serial-code delays. The operational paradigm for reconstruction is user-driven,
 352 however, so it is not necessarily the case that reconstruction immediately fol-
 353 lows primary computation. The test implementation therefore treats these as
 354 separate events.

355 3. Results

356 3.1. Test Datasets

357 Two datasets were used to test the performance of the partitioning algo-
358 rithm, and the parallel version of CHRT; both are hydrographic datasets col-
359 lected by the U.S. National Oceanographic and Atmospheric Administration as
360 part of the U.S. national charting programme.

361 The first, a primary hydrographic survey in the vicinity of Woods Hole,
362 MA, was conducted by the NOAA Ship *Whiting* in 2001 (Barnum, 2001), and
363 consists of a total of 37.7×10^6 observations in depth ranges from 2–30 m. The
364 second dataset is a portion of the survey conducted in Ernest Sound, AK in the
365 vicinity of Union Point by the NOAA Ship *Fairweather* in 2009 (Baird, 2009),
366 and consists of a total of 9.3×10^6 observations in depth ranges from 4–220 m.

367 Both datasets were used previously to demonstrate the development and
368 behaviour of the CHRT algorithm, and are more fully described in Calder and
369 Rice (2017).

370 3.2. Reconstruction Partitioning

371 To assess the performance of the partitioning algorithm itself, the data den-
372 sity estimates from both datasets were used to compute a partition for differing
373 computational resource counts. The run-time efficiency of the algorithm is es-
374 sential to its use: if the algorithm takes longer to compute the partition than
375 having the problem partitioned improves the run-time of the estimation algo-
376 rithm, then the effort is wasted. Figure 6 illustrates the actual and relative
377 computational time observed on a particular computer system, and in particu-
378 lar the rapid increase in run-time engendered by increasing computational
379 resource allocation (c.f. Saule et al. (2012)), as might be expected given the
380 known complexity class of the general case. For the Woods Hole data, slightly
381 higher run-times are observed, corresponding to the larger area being surveyed
382 and the 8 m low-resolution cell size compared to the 32 m size used for Ernest
383 Sound. For moderate (e.g., single workstation) resource counts, however, the
384 actual run-time is significantly smaller than the estimator run-time, making
385 the algorithm a pragmatic solution. (Note that the actual computational time
386 is only illustrative, since it will vary with hardware and compiler selections.)
387 Trade-offs between the estimator and partitioning algorithm runtime, and their
388 implications for how to partition the problem, are considered further in Sec-
389 tion 4.

390 The potential performance of the algorithm depends on the evenness with
391 which the observations can be distributed among the segments of the partition
392 (i.e., the degree of deviation from the nominal allocation of N/C observations).
393 Figure 7 shows the mean deviation per segment as a function of the number of
394 computational resources assigned, clearly showing that the percentage mismatch
395 between actual and ideal workload per computation resource is on the order of
396 a few percent of nominal workload. The mismatch rises with computational re-
397 sources since each segment becomes smaller, making each row or column added

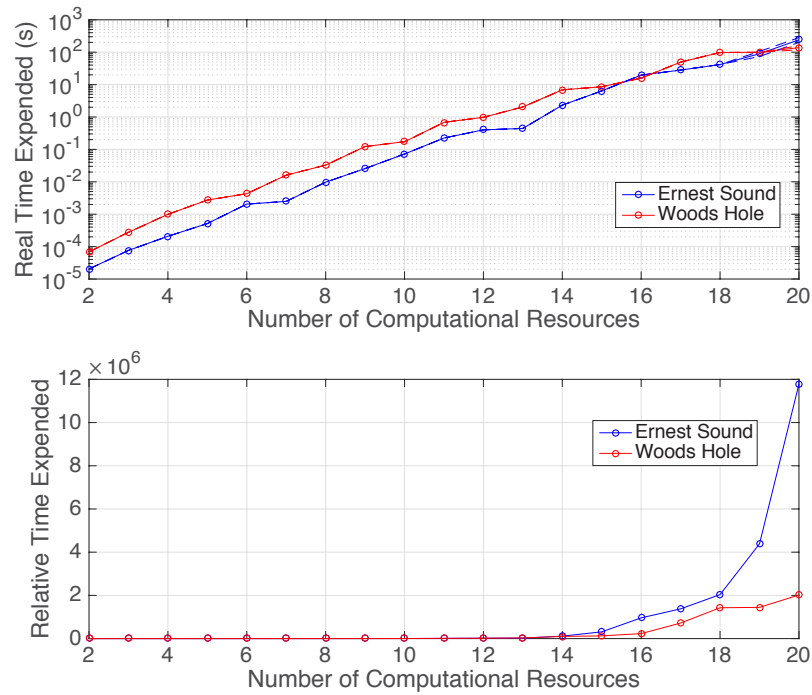


Figure 6: Estimate of absolute and relative run-time to compute a partition as a function of computational resource count. Note logarithmic scale on absolute run-time plot; dashed lines (only visible to the right of the absolute run-time plot due to scale) are 95% CI limits for $N = 100$ runs of the partitioning algorithm. Relative run-times are computed with respect to that for a computational resource count of two units.




Figure 7: Percentage average deviation from ideal observation distribution per segment as a function of computational resource count.

398 or removed a larger (potential) percentage of the nominal workload. The differ-
 399 ence between the two datasets is due to dataset size and geographical extent.

400 3.3. Speedup and Processing Rate

401 The two test datasets were processed using first the serial version of the
 402 algorithm, and then the parallel version, repeating the process 100 times in
 403 each case in order to gather statistics on variability. The test hardware having
 404 a quad-core, hyper-threaded processor, a range of 2–8 computational elements
 405 were considered.

406 The speedup achieved for the multi-threaded version of the algorithm is
 407 shown in Figure 8, and the efficiency (also known as strong scaling (Barnes,
 408 2016)) is shown in Figure 9. The algorithm demonstrates almost perfect (and
 409 very slightly super-linear) speedup for 2–3 computational elements, but then
 410 starts to diverge from ideal speedup as the effects of cache, memory band-
 411 width, and I/O contention start to take effect; the corresponding efficiency
 412 shows the equivalent relatively gentle decline with additional computational re-
 413 source. Note, however, that performance does not decrease as further resources
 414 are added. An overall speedup of about 4.1 times is achieved, which is perhaps
 415 not surprising on a single quad-core (albeit hyper-threaded) processor.

416 The overall computation rate per thread is shown in Figure 10. Clearly,
 417 the theoretical observation processing rate for each thread is constant; the ap-

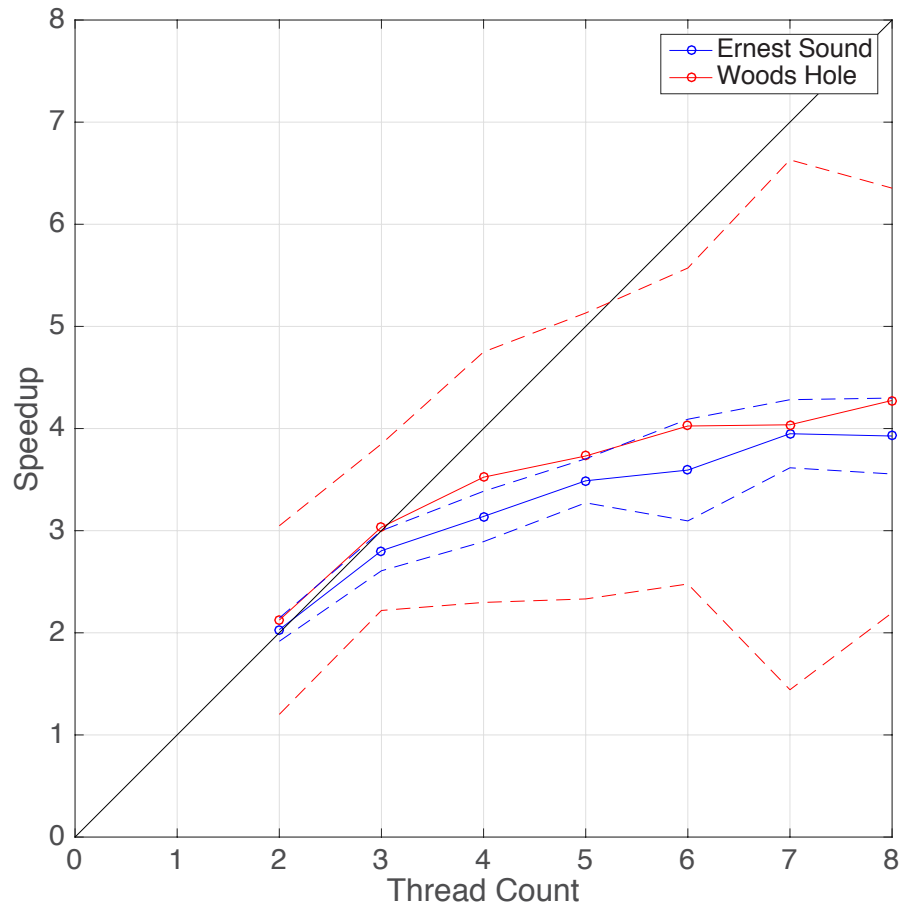


Figure 8: Speedup achieved by the algorithm with 2–8 threads on a single, quad-core, hyper-threaded CPU. Dashed lines indicate 95% CI limits.




Figure 9: Efficiency of computation (i.e., speedup per computational resource committed to the task) corresponding to Figure 8. Dashed lines indicate 95% CI limits.

418 parent processing rate, however, drops as the number of threads increases and
 419 contention for resources takes effect. For small numbers of threads, the addi-
 420 tional threads lead to sufficient improvement to compensate for the reduction in
 421 apparent processing rate; for larger number of threads, the resource contention
 422 overwhelms the benefit of extra threads, leading to reduced speed improvements.

423 The distribution of observations at the threads is given in Figure 11. A sig-
 424 nificant difference is observed between the two datasets due to the differences
 425 in bathymetry in the regions, and the type of echosounder used during the sur-
 426 veys. The more even distribution achieved with the Woods Hole dataset is one
 427 reason for the slightly improved speed-up observed. Analysis of the observation
 428 counts recorded at the threads indicates that the over-computation (i.e., the
 429 observations that need to be redundantly included in the partial computations
 430 so that no edge effects are engendered) average over all of the segments in the
 431 partition to approximately 2% of the total number of observations.

432 4. Discussion

433 The multi-threaded implementation of the CHRT algorithm clearly improves
 434 on the overall run-time for the algorithm, being limited on a single processor
 435 by factors other than the CPU bound of the algorithm. That is, the algo-
 436 rithm could theoretically complete faster if higher memory and disc bandwidth,

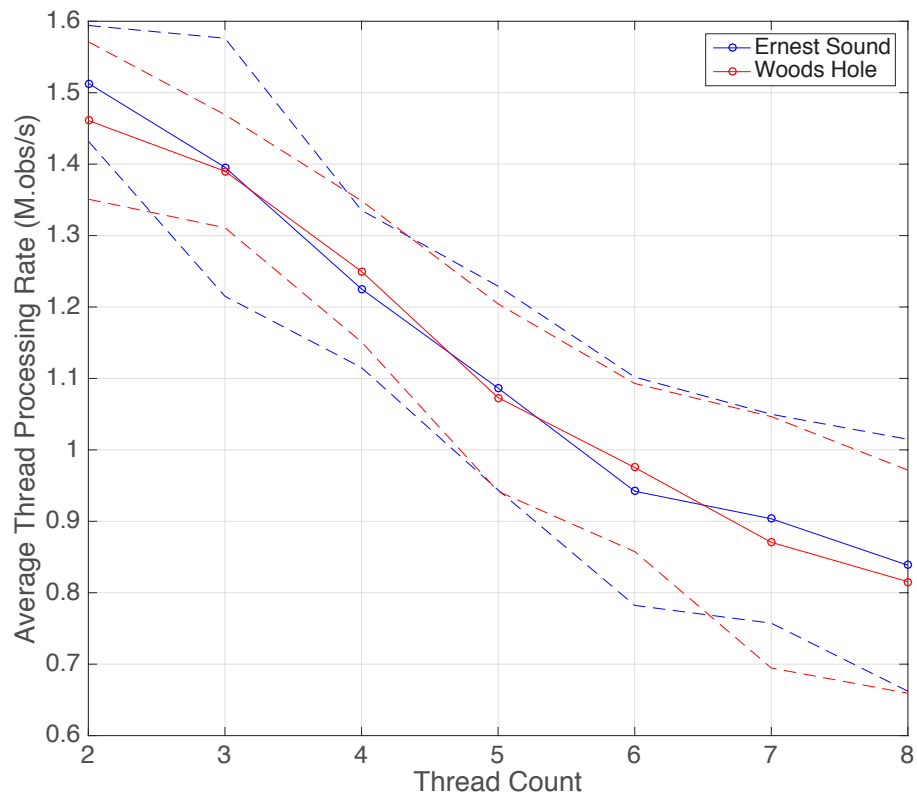


Figure 10: Processing rate per thread (in millions of observations per second processed) for 2–8 threads on a single, quad-core, hyper-threaded CPU. Dashed lines indicate 95% CI limits.

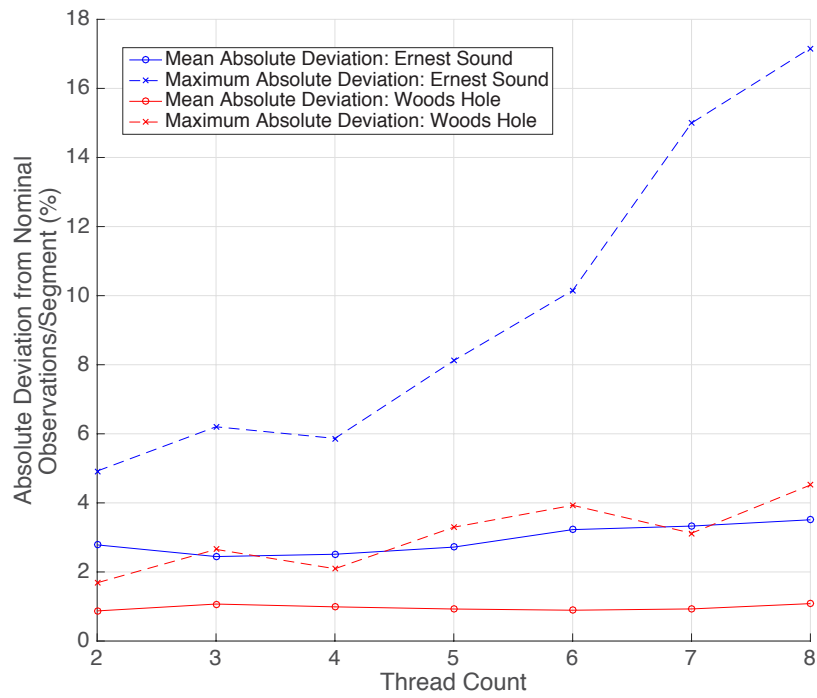


Figure 11: Maximum and mean absolute deviation of processed observations per partition segment from nominal “even” division as a percentage of nominal observation count. Non-zero deviations cause non-uniform thread run-times and hence lower overall efficiency.

437 and/or larger caches were available. Compressing the virtual tiles that act as
438 intermediate results before serialization (Barnes, 2016) or the addition of solid
439 state disc buffers (Barnes, 2017) might also improve the situation. On a sin-
440 gle processor, however, there is a limit to achievable performance improvement,
441 which suggests that it might be advantageous to distribute the algorithm over
442 more nodes in order to achieve greater speedups, a topic of current research.

443 An increase in the number of computational resources committed to the al-
444 gorithm has implications for the overall efficiency of the algorithm, and may not
445 always be advantageous. That is, although increasing numbers of computational
446 resources on distributed nodes will reduce the estimation algorithm’s run-time,
447 it remains an open question whether the gain will be sufficient to offset the
448 partition run-time costs for larger computational resource counts. This in turn
449 suggests that it might make sense to allow for a logical grouping of computa-
450 tional resources (i.e., making sub-clusters), partitioning over the groups at the
451 global level, and then sub-dividing the assigned segment locally within the group
452 either equally, or through an iteration of the partitioning algorithm applied to
453 the assigned segment; this is similar in spirit to the “hybrid” solution of Saule
454 et al. (2012). This would minimise the run-time for the partitioning algorithm
455 (the more so because the local sub-division could be computed in parallel),
456 although it would result in a globally sub-optimal partition. The difference be-
457 tween the performance of a locally optimal but globally sub-optimal solution
458 and the globally optimal partition when all effects are taken into account is not
459 obvious, and would require further investigation.

460 The results demonstrate that the degree of even distribution of observations
461 between segments depends on the problem itself, although the mean perform-
462 ance is within 1-2% of nominal for the two (very different) datasets tested.
463 Absolutely even distribution is likely impossible without further complexity in
464 the partitioning algorithm to allow for non-rectangular segments. This might
465 not be beneficial, however. Due to the CHRT algorithm’s use of observations
466 surrounding each segment to ensure that there are no edge effects, a longer
467 perimeter, such as could be generated with non-rectangular segments, would
468 lead to more observations being drawn into a segment. This extra computation
469 is redundant in the sense that more than one computational resource will have
470 to do the same base computations for the observation. Although current evi-
471 dence is that this is a small effect (approximately 2% of the overall load on each
472 thread), increased numbers of computational resources (resulting in smaller seg-
473 ments with higher perimeter to area ratios) and non-rectangular segments could
474 potentially increase it to a significant degree.

475 Consequently, it seems likely that further improvements to the algorithm
476 do not necessarily pertain to larger numbers of computational resources. Multi-
477 threading of the algorithm as applied to a single segment, for example by having
478 one thread set up the data at each low-resolution cell while one or more threads
479 do the processing within cells, might be a productive line of investigation.

480 5. Conclusions

481 The time taken to compute a bathymetric (or other scalar field) reconstruc-
482 tion from raw observations is critical for practical data processing methods;
483 acceleration of the computation can also be an enabler for more advanced algo-
484 rithms built on the base computation.

485 The results here demonstrate that it is possible to efficiently pre-partition
486 the computational task for a bathymetric reconstruction algorithm (in this case
487 CHRT) into a fixed number of segments, each of which has approximately the
488 same amount of computational effort. This allows the computation to proceed
489 without further communication between computational units, avoiding commu-
490 nication or synchronisation overhead. Partition times of order 10-100 millisec-
491 onds are observed for small numbers of computational resources, along with
492 mean absolute deviations from even distribution of effort on order 1-2%.

493 The resulting multi-threaded demonstration implementation of a parallel
494 CHRT, for use on a single, quad-core CPU, is observed to achieve maximum
495 speed-up of 4.1 on eight threads, with the sub-linear performance being driven
496 by cache, memory, and disc contention between the threads. Nominal processing
497 rates of up to 1.5×10^6 observations per second per thread are observed.

498 Examination of algorithm behaviour (partition computation rate, redundant
499 but necessary computations, and observation count balance) with increasing
500 numbers of computational resources indicate that it might be fruitful to examine
501 either distribution over multiple compute nodes, or multi-threading the core
502 algorithm to further improve the performance of the algorithm.

503 Acknowledgements

504 The support of NOAA grants NA10NOS4000073 and NA15NOS4000200 for this
505 work is gratefully acknowledged.

506 Computer Code Availability

507 An example implementation of the algorithm, written in C++11, is avail-
508 able at <https://github.com/brian-r-calder/density-partition.git>, using the
509 GNU GPL, version 2. The code was written to be portable, and therefore
510 should require only a C++11 compiler for use; it was developed primarily
511 on macOS, but has also been tested on both Windows and Linux platforms.
512 Further details on compilation are provided in the source distribution. Ex-
513 ample input data density files, and expected output, are also provided. An
514 example implementation of a one-dimensional version of the CHRT algorithm
515 was published to accompany Calder and Rice (2017), and can be found at
516 <https://github.com/brian-r-calder/vr-grid-estimator.git>. The correspond-
517 ing author may be contacted for further details.

518 **References**

- 519 Aspvall, B., Halldórsson, M. M., Manner, F., 2001. Approximations for the
520 general block distribution of a matrix. *Theoretical Computer Science* 262,
521 145–160.
- 522 Baird, D. D., 2009. Descriptive report: Hydrographic survey H11825 (Ernest
523 Sound, AK). Tech. rep., National Oceanic and Atmospheric Administration.
- 524 Barnes, R., 2016. Parallel Priority-Flood depression filling for trillion cell digital
525 elevation models on desktops or clusters. *Computers and Geosciences* 96, 56–
526 68.
- 527 Barnes, R., 2017. Parallel non-divergent flow accumulation for trillion cell digi-
528 tal elevation models on desktops and clusters. *Environmental Modelling and*
529 *Software* 92, 202–212.
- 530 Barnum, S. R., 2001. Descriptive report: Hydrographic survey H11077 (Woods
531 Hole, MA). Tech. rep., National Oceanic and Atmospheric Administration,
532 1315 East West Highway, Silver Spring, MD.
- 533 Berger, M. J., Bokhari, S. H., 1987. A partitioning strategy for nonuniform
534 problems on multiprocessors. *IEEE Trans. Comp.* C-36 (5), 570–580.
- 535 Berman, P., DasGupta, B., Muthukrishnan, S., Ramaswami, S., 2001. Efficient
536 approximation algorithms for tiling and packing problems with rectangles. *J.*
537 *Algorithms* 41, 443–470.
- 538 Calder, B. R., Mayer, L. A., 2003. Automatic processing of high-rate, high-
539 density multibeam echosounder data. *Geochem., Geophys. and Geosystems*
540 (G3) DID 10.1029/2002GC000486 4 (6).
- 541 Calder, B. R., Plumlee, M. D., 2017. On testing of complex hydrographic data
542 processing algorithms. In: *Proc. U.S. Hydrographic Conference 2017. The*
543 *Hydrographic Society of America, Galveston, TX*, pp. 1–6.
- 544 Calder, B. R., Rice, G., 2017. Computationally efficient variable resolution depth
545 estimation. *Computers and Geosciences* 106, 49–59.
- 546 Crow, F. C., 1984. Summed-area tables for texture mapping. *Proc. ACM SIG-*
547 *GRAPH* 18 (3), 207–212.
- 548 Debese, N., Moitié, R., Seube, N., 2012. Multibeam echosounder data cleaning
549 through a hierarchic adaptive and robust local surfacing. *Computers and*
550 *Geosciences* 46, 330–339.
- 551 Gamma, E., Helm, R., Johnson, R., Vissides, J., 1994. *Design Patterns.*
552 Addison-Wesley Professional.

- 553 Grigni, M., Manne, F., 1996. On the complexity of the generalized block distri-
554 bution. In: Proc. International Workshop of Parallel Algorithms for Irregu-
555 larly Structured Problems. Springer, Berlin Heidelberg, pp. 319–326.
- 556 Hofierka, J., Lacko, M., Zubal, S., 2017. Parallelization of interpolation, solar
557 radiation and water flow simulation modules in GRASS GIS using OpenMP.
558 Computers and Geosciences 107, 20–27.
- 559 Isenburg, M., Liu, Y., Shewchuck, J., Snoeyink, J., Thirion, T., 2006. Generat-
560 ing raster DEM from mass points via TIN streaming. In: M. Raubal et al.
561 (Ed.), Geographic Information Science (Lecture Notes in Computer Science).
562 Vol. 4197. Springer-Verlag, pp. 186–198.
- 563 Khanna, S., Muthukrishnan, S., Paterson, M., 1998. On approximating rectan-
564 gle tiling and packing. In: Proc. ACM Symp. on Discrete Algorithms. Assoc.
565 Comp. Machinery, pp. 384–393.
- 566 Krishnan, S., Baru, C., Crosby, C., Nov. 2010. Evaluation of MapReduce for
567 gridding lidar data. In: Qiu, J., Zhao, G., Rong, C. (Eds.), Proc. Second IEEE
568 International Conference on Cloud Computing Technology and Science. pp.
569 33–40.
- 570 Ladner, R. E., Fischer, M. J., 1980. Parallel prefix computation. J. Assoc. Comp.
571 Mach. 27 (4), 831–838.
- 572 Land, A. H., Doig, A. G., 1960. An automatic method of solving discrete pro-
573 gramming problems. Econometrica 28 (3), 497–520.
- 574 Lorys, K., Paluch, K. E., 2003. New approximation algorithm for RTILE prob-
575 lem. Theoretical Computer Science 303, 517–537.
- 576 Manne, F., Sørenvik, T., 1996. Partitioning an array onto a mesh of processors.
577 In: Waśniewski, J., Dongarra, J., Madsen, K., Olesen, D. (Eds.), Proc. Inter-
578 national Workshop of Applied Parallel Computing. No. 1184 in Lecture Notes
579 in Computer Science. Springer, Berlin Heidelberg, pp. 467–477.
- 580 Muthukrishnan, S., Suel, T., 2005. Approximation algorithms for array parti-
581 tioning problems. J. Algorithms 54, 85–104.
- 582 Nichol, D. M., 1991. Rectilinear partitioning of irregular data parallel compu-
583 tations. Tech. Rep. NASA-CR-187601, NASA.
- 584 Qin, C.-Z., Zhan, L., 2012. Parallelizing flow-accumulation calculations on graph-
585 ics processing units—from iterative DEM preprocessing to recursive multiple-
586 flow-direction algorithm. Computers and Geosciences 43, 7–16.
- 587 Saule, E., Bas, E. Ö., Çatalyürek, Ü. V., 2012. Load-balancing spatially located
588 computations using rectangular partitions. J. Parallel Distrib. Comput. 72,
589 1201–1214.

- 590 Smith, W. H. F., Wessel, P., 1990. Gridding with continuous curvature splines
591 in tension. *Geophysics* 55 (3), 293–305.
- 592 Tesfa, T. K., Tarboton, D. G., Watson, D. W., Schreuders, K. A. T., Baker,
593 M. E., Wallace, R. M., 2011. Extraction of hydrological proximity measures
594 using DEMs using parallel processing. *Environmental Modelling and Software*
595 26, 1696–1709.
- 596 Tobler, W., 1970. A computer movie simulating urban growth in the Detroit
597 Region. *Economic Geograpy* 46, 234–240.
- 598 Wilkinson, B., Allen, M., 2005. *Parallel Programming*, 2nd Edition. Pearson
599 Education, Upper Sadde River, NJ 07458.
- 600 Yıldırım, A. A., Watson, D., Tarboton, D., Wallace, R. M., 2015. A virtual tile
601 approach to raster-based calculations of large digital elevation models in a
602 shared-memory system. *Computers and Geosciences* 82, 78–88.