BOTTOM-UP CONSTRUCTION AND 2:1 BALANCE REFINEMENT OF LINEAR OCTREES IN PARALLEL*

HARI SUNDAR†, RAHUL S. SAMPATH‡, AND GEORGE BIROS§

Abstract. In this article, we propose new parallel algorithms for the construction and 2:1 balance refinement of large linear octrees on distributed memory machines. Such octrees are used in many problems in computational science and engineering, e.g., object representation, image analysis, unstructured meshing, finite elements, adaptive mesh refinement and N-body simulations. Fixed-size scalability and isogranular analysis of the algorithms, using an MPI-based parallel implementation, was performed on a variety of input data and demonstrated good scalability for different processor counts (1 to 1024 processors) at the Pittsburgh Supercomputing Center's TCS-1 AlphaServer. The results are consistent for different data distributions. Octrees with over a billion octants were constructed and balanced in less than a minute on 1024 processors. Like other existing algorithms for constructing and balancing octrees, our algorithms have $\mathcal{O}(n \log n)$ work and $\mathcal{O}(n)$ storage complexity. Under reasonable assumptions on the distribution of octants and the work per octant, the parallel time complexity is $\mathcal{O}(n/n_p \log(n/n_p) + n_p \log n_p)$, were n is the final number of leaves and n_p is the number of processors.

Key words. Linear octrees, Balance refinement, Morton encoding, large scale parallel computing, space filling curves

AMS subject classifications. 65N50, 65Y05, 68W10, 68W15

1. Introduction.

Spatial decompositions of the d-dimensional cube have important applications in scientific computing: they can be used as algorithmic foundations for adaptive finite element methods [3, 19], adaptive mesh refinement methods [16, 23], and many-body algorithms [15, 17, 31, 36, 38, 39]. The earliest examples of tree-based spatial decompositions of $\mathcal{R}^n$ can be traced to the use of *binary space partitions* [11, 12, 26]. Binary space partitioning (BSP) is a method for recursively subdividing a space into convex sets by hyperplanes. This subdivision gives rise to a representation of the space by means of a tree data structure known as a BSP tree. The BSP partitions its domain into two subregions and therefore the BSP tree is a binary tree. BSP trees can be used in spaces with any number of dimensions, but *quadtrees* [9] and *octrees* [21] are most useful in partitioning 2D and 3D domains, respectively. These use axis aligned lines and planes, respectively, instead of arbitrary hyperplanes and are efficient for the specific domains they are intended to work on.

Octrees and quadtrees are usually employed while solving two types of problems: searching and partitioning. Searches within a domain using d-trees (d-dimensional trees with a maximum of 2^d children per node), benefit from the reduction of the complexity of the search from $\mathcal{O}(n)$ to $\mathcal{O}(\log n)$. Similar benefits can be obtained by the use of *space-filling curves* [32]. This equivalence has been used in linear quadtree and

* This work was supported by the U.S. Department of Energy under grant DE-FG02-04ER25646, and the U.S. National Science Foundation grants CCF-0427985, CNS-0540372, and DMS-0612578. Computing resources on the TeraGrid's HP AlphaCluster system at the Pittsburgh Supercomputing Center were provided under the award MCA04T026.

†Department of Bioengineering, University of Pennsylvania, 120 Hayden Hall, 3320 Smith Walk, Philadelphia, PA, 19104-2688 (hsundar@seas.upenn.edu).

‡Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, 220 S, 33rd Street, Philadelphia, PA, 19104-6315 (rahulss@seas.upenn.edu).

§Departments of Mechanical Engineering and Applied Mechanics, and Computer and Information Science, University of Pennsylvania, 220 S, 33rd Street, Philadelphia, PA, 19104-6315 (biros@seas.upenn.edu).

octree representations [4, 35], including the current work. Unstructured meshes are often preferred over uniform discretizations because they can be used with complicated domains and permit rapid grading from small to large elements. However, generating large unstructured meshes is a challenging task [28]. On the contrary, octree-based unstructured hexahedral meshes can be constructed efficiently [5, 14, 24, 25, 27, 34]. Although they are not suitable for highly complicated geometries, they provide a good compromise between adaptivity and simplicity for numerous applications like solid modeling [21], object representation [1, 6], visualization [10], image segmentation [30], adaptive mesh refinement [16, 23] and N-body simulations [15, 17, 31, 36, 38, 39].

Octree data structures used in discretizations of partial differential equations should satisfy certain spatial distribution of octant size [4, 35]. That is, adjacent octants should not differ greatly in size[1]. Furthermore, conforming discretizations require a so-called 'balance condition' that is necessary to construct appropriate function spaces. In particular, when the 2:1 balance constraint is imposed on octree-based hexahedral meshes, it ensures that there is at most one *dangling* node on any edge or face. What makes the balance-refinement problem difficult and interesting is a property known as the *ripple effect*: An octant can trigger a sequence of splits whereby it can force an octant to split, even if it is not in its immediate neighborhood. Hence, balance-refinement is an inherently iterative process. Solving the balance-refinement problem in parallel, introduces further challenges in terms of synchronization and communication since the ripple can propagate across multiple processors.

Related Work. Limited work has been done on large scale parallel construction [17, 37, 39] and balance refinement [19, 35] of octrees, and the best known algorithms exhibit suboptimal isogranular scalability. The key component in constructing octrees is the partitioning of the input in order to achieve good load balancing. The use of space-filling curves for partitioning data has been quite popular [17, 35, 37, 39]. The proximity preserving property of space-filling curves makes them attractive for data partitioning. All the existing algorithms for constructing octrees use a top-down approach after the initial partition. The major hurdle in using a parallel top-down approach is avoiding overlaps. This typically requires some synchronization after constructing a portion of the tree [35, 37, 39]. Section 3.1 describes the issues that arise in using a parallel top-down approach.

Bern et al. [4] proposed an algorithm for constructing and balancing quadtrees for EREW PRAM architectures. However, it cannot be easily adapted for distributed architectures. In addition, the balanced quadtree produced is suboptimal and can have up to 4 times as many cells as the optimal balanced quadtree. Tu et al. [35] propose a more promising approach, which was evaluated on large octrees. They construct and balance 1.22B octants for the Greater Los Angeles basin dataset [20] on 2000 processors in about 300 seconds. This experiment was performed on the TCS-1 terascale computing HP AlphaServer Cluster at the Pittsburgh Supercomputing Center. In contrast, we construct and balance[2] 1B octants (approximately) for three different point distributions (Gaussian, Log-Normal and Uniform) on 1024 processors on the same cluster in about 60 seconds.

Synopsis and Contributions. In this paper we present two parallel algorithms:

[1] This is referred to as the 2:1 balance constraint. A formal definition of this constraint is given in section 2.2.

[2] While we enforce the 0-*balance* constraint, [35] only enforce the 1-*balance* constraint. Note that it is harder to 0-*balance* a given octree. See section 2.2 for more details on the different balance constraints.

one to construct complete linear octrees from a list of points; and one to enforce an optimal 2:1 balance constraint[3] on complete linear octrees. We use a linear octree Morton-encoding-based representation. Given a set of points, partitioned across processors, we create a set of octants that we sort and repartition using the Morton ordering. A complete linear octree is constructed using the seed octants. Then, we built an additional auxiliary list of a small number of coarse octants or *blocks*. This auxiliary octant set encapsulates and compresses the local spatial distribution of octants; it is used to accelerate the 2:1-balance refinement, which we implement using a hybrid strategy: intra-block balancing is performed by a classical level-by-level balancing/duplicate-removal scheme; and inter-block balancing is performed by a variant of the ripple-propagation algorithm proposed in [35]. The main parallel tools used are sample sorts (accelerated by biotonic sorts), and standard point-to-point/collective communication calls.[4]

In a nutshell, the major contributions of this work are:

- A parallel bottom-up algorithm for coarsening octrees, which is also used for partitioning the input in our other algorithms.
- A parallel bottom-up algorithm for constructing linear octrees. We avoid the synchronization issues that are usually associated with parallel top-down approaches to the problem.
- An algorithm for enforcing 2:1 balance refinement in parallel. The algorithm constructs the minimum number of nodes to satisfy the 2:1 constraint. Its key feature is that it avoids parallel searches, which as we show in sections 3.5.6 and 3.5.7, are the main hurdles in achieving good isogranular scalability.

Remark: The main parallel cost of the algorithm is that related to the parallel sorts that run in $\mathcal{O}(n \log n)$ work and $\mathcal{O}(n/n_p \log(n/n_p) + n_p \log(n_p))$ time, assuming uniformly distributed points [13]. In the following sections we present several algorithms for which we give precise `work` and `storage` complexity. For some of the parallel algorithms we also give `time` complexity estimates; this corresponds to wall-clock time and includes work/per processor and communication costs. The precise number depends on on the initial distribution and the effectiveness of the partitioning. Thus the numbers for `time` are only an estimate under uniform distribution assumptions. If the `time` complexity is not specifically mentioned then it is comparable to that of a sample-sort.

Organization of the paper. In section 2 we introduce some terminology that will be used in the rest of the paper. Section 3 describes the various components of our construction and balance refinement algorithms. In Section 4 we present numerical experiments, including fixed size and isogranular scalability tests on different data distributions. Finally, in Section 5 shortcomings of the proposed approach are discussed and some suggestions for future work are also offered. Table 1.1 summarizes the notation that is used in the subsequent sections.

2. Background. Octrees are trees in which every node has a maximum of eight children. They are analogous to binary trees (maximum of 2 children per node) in 1-D and quadtrees (maximum of 4 children per node) in 2-D. A node with no children is called a *leaf*. The only node with no parent is the *root*. Nodes that have the same parent are called *siblings*. A node's children, grandchildren and so on and so forth are collectively referred to as the node's *descendants* and this node will be an *ancestor* of

[3] There exists a unique least common balance refinement for a given octree [22].

[4] When we discuss communication costs we assume a Hypercube network topology with $\Theta(n_p)$ Bisection Width.

TABLE 1.1
Notation

D_{max}	Maximum depth of the tree.
$\mathcal{L}(N)$	Level of octant N.
$\mathcal{P}(N)$	Parent of octant N.
$\mathcal{S}(N)$	Siblings (sorted) of octant N.
$\mathcal{C}(N)$	Children (sorted) of octant N.
$\mathcal{D}(N)$	Descendant of octant N.
$\mathcal{FC}(N)$	First child of octant N.
$\mathcal{LC}(N)$	Last child of octant N.
$\mathcal{FD}(N,l)$	First descendant of octant N at level l.
$\mathcal{LD}(N,l)$	Last descendant of octant N at level l.
$\mathcal{DFD}(N)$	Deepest first descendant of octant N.
$\mathcal{DLD}(N)$	Deepest last descendant of octant N.
$\mathcal{A}(N)$	Ancestor of octant N.
$\mathcal{A}_{finest}(N,K)$	Nearest Common Ancestor of octants N and K.
$\mathcal{N}(N,l)$	List of all potential neighbors of octant N at level l.
$\mathcal{N}^s(N,l)$	A subset of $\mathcal{N}(N,l)$, with the property that all of these share the same common corner with N. This is also the corner that N shares with its parent.
$\mathcal{N}(N)$	Neighbor of N at any level.
$\mathcal{I}(N)$	Insulation layer around octant N.
$\{\ldots\}$	A set of elements.
$\emptyset$	The empty set.
$A \leftarrow B$	Assignment operation.
$A \oplus B$	Bitwise A `XOR` B.
$\{A\} \cup \{B\}$	Union of the sets A and B. The order is preserved, if possible.
$\{A\} \cap \{B\}$	Intersection of the sets A and B.
$A + B$	The list formed by concatenating the lists A and B.
$A - B$	Remove the contents of B from A.
$A[i]$	i^{th} element in list A.
`len`(A)	Number of elements in list A.
`Sort`(A)	Sort A in the ascending Morton order.
A`.push_front`(B)	Insert B to the beginning of A.
A`.push_back`(B)	Append B to the end of A.
A_{global}	Union of the list A from all the processors.
n_p	Total number of processors.
`Send`(A, r)	Send A to processor with rank $= r$.
`Receive()`	Receive from any processor.
N^p_{max}	Maximum number of points per octant.

its descendants. A node along with all its descendants can be viewed as a separate tree in itself with this node as its root. Hence, this set is also referred to as a *subtree* of the original tree. The depth of a node from the root is referred to as its *level*. As shown in Fig. 2.1(a), the root of the tree is at level 0 and the children of any node are one level higher than the parent.

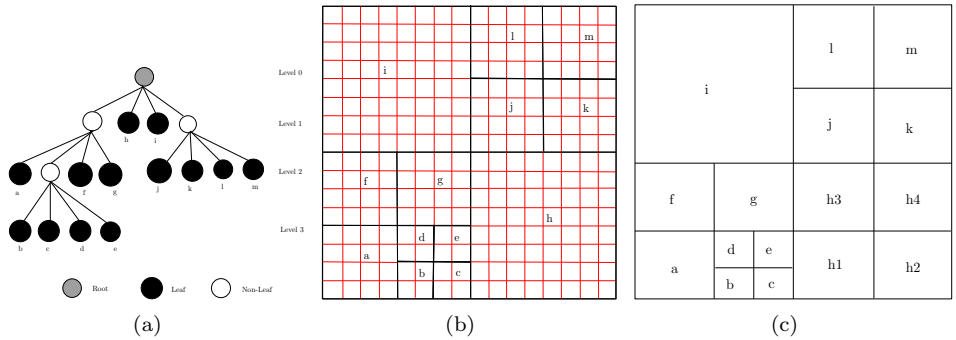

FIG. 2.1. (a) *Tree representation of a quadtree and* (b) *decomposition of a square domain using the quadtree, superimposed over an uniform grid, and* (c) *a balanced linear quadtree: result of balancing the quadtree.*

Octrees and quadtrees[5] can be used to partition cuboidal and rectangular regions, respectively (Fig. 2.1(b)). These regions are referred to as the domain of the tree. A set of octants is said to be complete if the union of the regions spanned by them covers the entire domain. To reduce storage costs, only the complete list of leaf nodes is stored, i.e., as a linear octree. To use a linear representation, a *locational code* is needed to identify the octants. A locational code is a code that contains information about the position and level of the octant in the tree. The following section describes one such locational code known as the *Morton encoding*[6].

2.1. Morton encoding. In order to construct a Morton encoding, the maximum possible depth, D_{max}, of the tree is specified *a priori*. The domain is represented by an uniform grid of $2^{D_{max}}$ indivisible cells in each dimension (Fig. 2.1(b)). Each cell is identified by an integer triplet representing its x, y and z coordinates, respectively. Any octant in the domain can be uniquely identified by specifying one of its corners, also known as its anchor, and its level in the tree (Fig. 2.2).

The Morton encoding for any octant is derived by interleaving[7] the binary representations (D_{max} bits each) of the three coordinates of the octant's anchor, and then appending the binary representation ($(\lfloor(\log_2 D_{max})\rfloor + 1)$ bits) of the octant's level to this sequence of bits. Interesting properties of the Morton encoding scheme are listed in Appendix A. In the rest of the paper the terms *lesser* and *greater* and the symbols $<$ and $>$ are used to compare octants based on their Morton ids, and *coarser* and *finer* to compare them based on their relative sizes, i.e., their levels in the octree.

2.2. Balance Constraint. In many applications involving octrees, it is desirable that adjacent elements do not differ greatly in size [18, 19, 35]. Generalizing Moore's [22] categorization of the general balance conditions, we have the following definition:

[5] All the algorithms described in this paper are applicable to both octrees and quadtrees. For simplicity, we will use quadtrees to illustrate the concepts in this paper and use the terms 'octrees' and 'octants', consistently, in the rest of the paper.

[6] Morton encoding is one of many space-filling curves [7]. Our algorithms are generic enough to work with other space-filling curves as well. However, Morton encoding is relatively simpler to implement since, unlike other space-filling curves, no rotations or reflections are performed.

[7] Instead of bit-interleaving as described here, we use a multicomponent version (Appendix B) of the Morton encoding scheme.

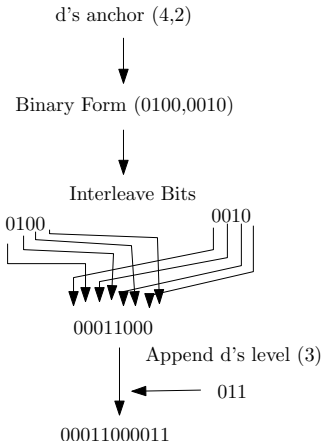

FIG. 2.2. *Computing the Morton id of quadrant 'd' in the quadtree shown in Fig. 2.1(b). The anchor for any quadrant is it's lower left corner.*

DEFINITION 1. *A linear d-tree is k-balanced if and only if, for any* $l \in [1, D_{max})$*, no leaf at level* l *shares an* m*-dimensional face*[8] $(m \in [k, d))$ *with another leaf, at level greater than* $l + 1$*.*

For the specific case of octrees we use 2-*balanced* to refer to octrees that are balanced across faces, 1-*balanced* to refer to octrees that are balanced across edges and faces, and 0-*balanced* to refer to octrees that are balanced across corners, edges and faces. An example of a 0-*balanced* quadtree is shown in Figure 2.1(c). The balance algorithm proposed in this work is capable of *k-balancing* a given complete linear octree, and since it is hardest to 0-*balance* a given octree we report all results for the 0-*balance* case.

3. Algorithms.

3.1. Constructing large linear octrees in parallel. Octrees are usually constructed by using a top-down approach: starting with the root octant, cells are split iteratively based on some criteria, until no further splits are required. This is a simple and efficient sequential algorithm. However, it's parallel analogue is not so. We use the case of point datasets to discuss some shortcomings of a parallel top-down tree construction. Formally, the problem might be stated as: Construct a complete linear octree in parallel from a distributed set of points in a domain with the constraint that no octant should contain more than (N^p_{max}) number of points. Each processor can independently construct a tree using a top-down approach on its local set of points. Constructing a global linear octree requires a parallel merge. Merging however, is not straightforward.

1. Consider the case where the local number of points in some region on every processor was less than (N^p_{max}), and hence all the processors end up having the same level of coarseness in the region. However, the total number of points in that region could be more than (N^p_{max}) and hence the corresponding octant should be refined further.

[8] A corner is a 0-dimensional face, an edge is a 1-dimensional face and a face is a 2-dimensional face.

Algorithm 1. CONSTRUCTING A COMPLETE LINEAR OCTREE FROM A DISTRIBUTED LIST OF POINTS (PARALLEL) - `Points2Octree`

Input: A distributed list of points, L and a parameter, (N^p_{max}), which specifies the maximum number of points per octant.
Output: Complete linear Octree, B.
Work: $\mathcal{O}(n \log n)$, where $n =$ len(L).
Storage: $\mathcal{O}(n)$, where $n =$ len(L).

1. $F \leftarrow$ [Octant$(p, D_{max}), \forall p \in L$]
2. Sort(F)
3. $B \leftarrow$ BlockPartition(F) (Algorithm 4)
4. **for each** $b \in B$
5. **if** NumberOfPoints$(b) > N^p_{max}$
6. $B \leftarrow B - b + \mathcal{C}(b)$
7. **end if**
8. **end for**

2. In most applications, we would also like to associate a unique processor to each octant. Thus, duplicates across processors must be removed.
3. For linear octrees overlaps across processors must be resolved.
4. Since there might be overlaps and duplicates, not all the work done by the processors can be accounted as useful work. This is a subtle yet important point to consider while analyzing the algorithm for load-balancing.

Previous work [17, 35, 37, 39] on this problem has addressed these issues; However, all the existing algorithms involve many synchronization steps and thus suffer from a sizable overhead, resulting in suboptimal isogranular scalability. Instead, we propose a bottom-up approach for constructing octrees from points. The crux of the algorithm is to distribute the data across the processors in such a way that there is uniform load distribution across processors and the subsequent operations to build the octree can be performed by the processors independently, i.e., requiring no additional communication.

First, all points are converted into octants at the maximum depth and then partitioned across the processors using the algorithm described in Section 3.4.1. This produces a contiguous set of coarse blocks (with their corresponding points) on each processor. The complete linear octree is generated by iterating through the blocks and by splitting them based on number of points per block[9]. This process is continued until no further splits are required. This procedure is summarized in Algorithm 1.

Next we describe algorithmic components that are fundamental to our construction and balance refinement algorithms: 1) the generation of a coarse linear octree between two given octants; 2) generation of a complete linear octree from a partial set of octants; and 3) coarsening of octrees.

3.2. Constructing a minimal linear octree between two octants. Given two octants, a and $b > a$, we wish to generate the minimal number of octants that

[9]Refer to the Appendix D on how to sample the points in order to construct the coarsest possible octree.

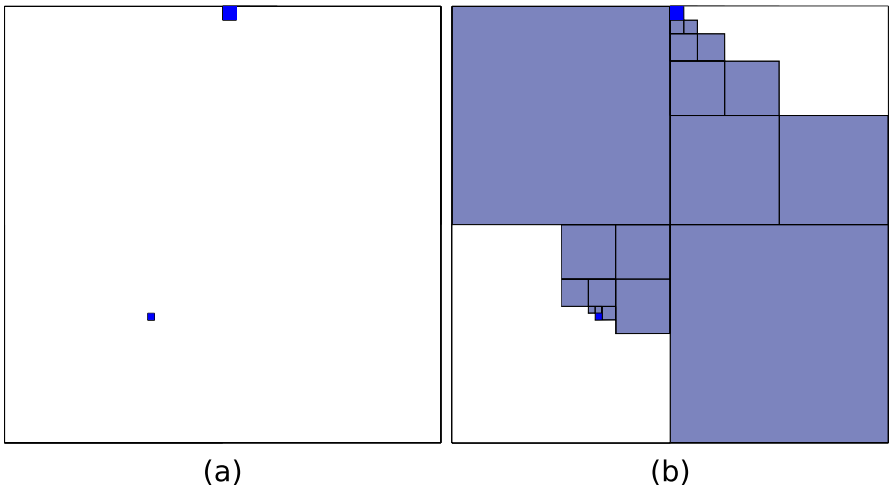

FIG. 3.1. (a) *Two cells: Input to* Algorithm 2. (b) *The minimal number of octants between the cells given in* (a). *This is produced by using* (a) *as input to* Algorithm 2.

Algorithm 2. CONSTRUCTING A MINIMAL LINEAR OCTREE BETWEEN TWO OCTANTS (SEQUENTIAL) - CompleteRegion

---

Input: Two octants, a and $b > a$.
Output: R, the minimal linear octree between a and b.
Work: $\mathcal{O}(n \log n)$, where $n =$ len(R).
Storage: $\mathcal{O}(n)$, where $n =$ len(R).

1. $W \leftarrow \mathcal{C}(\mathcal{A}_{finest}(a, b))$
2. **for each** $w \in W$
3. **if** $(a < w < b)$ AND $(w \notin \{\mathcal{A}(b)\})$
4. $R \leftarrow R + w$
5. **else if** $(w \in \{\{\mathcal{A}(a)\}, \{\mathcal{A}(b)\}\})$
6. $W \leftarrow W - w + \mathcal{C}(w)$
7. **end if**
8. **end for**
9. Sort(R)

---

span the region between a and b according to the Morton ordering. The algorithm (Algorithm 2) first calculates the nearest common ancestor of the octants a and b. This octant is split into its eight children. Out of these, only the octants that are either greater than a and lesser than b or ancestors of a are retained and the rest are discarded. The ancestors of either a or b are split again and we iterate until no further splits are necessary. This produces the minimal coarse complete linear octree between the two octants a and b. This is illustrated in Figure 3.1. This algorithm is based on the Properties 3 and 4 of the Morton ordering, which are listed in Appendix A.

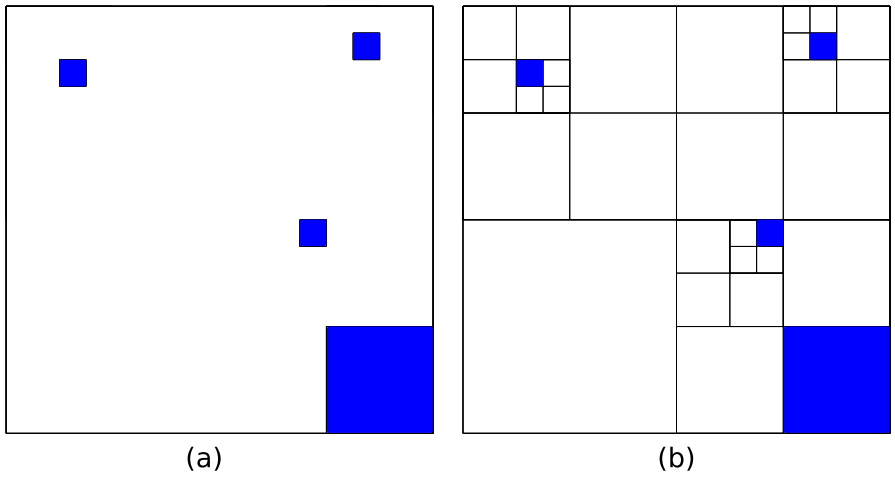

FIG. 3.2. `(a)` *A partial set of quadrants: Input to* `Algorithm 3`. `(b)` *A complete linear quadtree containing the cells in* `(a)`. *This is produced by using* `(a)` *as input to* `Algorithm 3`.

3.3. Constructing complete linear octrees from a partial set of octants. In order to construct a complete linear octree from a partial set of octants (e.g. Figure 3.2(a)), the octants are initially sorted based on the Morton ordering. Algorithm 8 is subsequently used to remove overlaps, if any. Two additional octants are added to complete the domain; the first one is the coarsest ancestor of the least possible octant (the deepest first descendant of the root octant, Property 7), which does not overlap the first given octant, and the second is the coarsest ancestor of the greatest possible octant (the deepest last descendant of the root octant, Property 9), which does not overlap the last given octant. The octants are distributed across the processors to get a weight-based uniform load distribution, and such that the last element on any processor is the same as the first element on the next processor. The local complete linear octree is subsequently generated by completing the region between every consecutive pair of octants as described in Section 3.2. The overlap guarantees that the union of these local complete octrees gives us a global complete and linear octree. We ignore the last octant on each processor, except the last, since these were replicated. This is illustrated in Figure 3.2.

3.4. Parallel bottom-up coarsening of octrees. Given a distributed list of leaves, we want to construct a complete linear coarse octree. We first sort the leaves according to their Morton ordering and then distribute them across the processors so that every processor has the same number of leaves. We select the least and the greatest octant at each processor (e.g., octants a and h from Figure 3.3(a)) and complete the region between them, as described in Section 3.2, to obtain a list of coarse octants. We then select the coarsest cell(s) out of this list of coarse octants (octant e in Figure 3.3(a)). We use the selected octants at each processor and construct a complete linear octree as described in Section 3.3. This gives us a global coarse complete linear octree that is based on the underlying data distribution[10].

3.4.1. Using the parallel coarsening algorithm for partitioning octants. A simple way to partition the domain into an union of blocks would be to take a

[10]Refer to the Appendix C for an estimate of the number of blocks produced.

Algorithm 3. Constructing a complete linear octree from a partial (incomplete) set of octants (parallel) - `CompleteOctree`

Input: A distributed sorted list of octants, L.
Output: R, the complete linear octree.
Work: $\mathcal{O}(n \log n)$, where $n =$ `len`(R).
Storage: $\mathcal{O}(n)$, where $n =$ `len`(R).

```
1.  RemoveDuplicates(L)
2.  L ← Linearise(L) ( Algorithm 8 )
3.  Partition(L) ( Algorithm 5 )
4.  if rank = 0
5.     L.push_front(FC (A_finest (DFD( root ), L[1])))
6.  end if
7.  if rank = (n_p − 1)
8.     L.push_back(LC (A_finest (DLD( root), L[len (L)])))
9.  end if
10. if rank > 0
11.    Send(L[1] ,(rank−1) )
12. end if
13. if rank < (n_p − 1)
14.    L.push_back(Recieve() )
15. end if
16. for i ← 1 to (len(L) − 1)
17.       A ← CompleteRegion (L[i], L[i + 1]) ( Algorithm 2 )
18.       R ← R + L[i] + A
19. end for
20. if rank = (n_p − 1)
21.    R ← R + L[len(L)]
22. end if
```

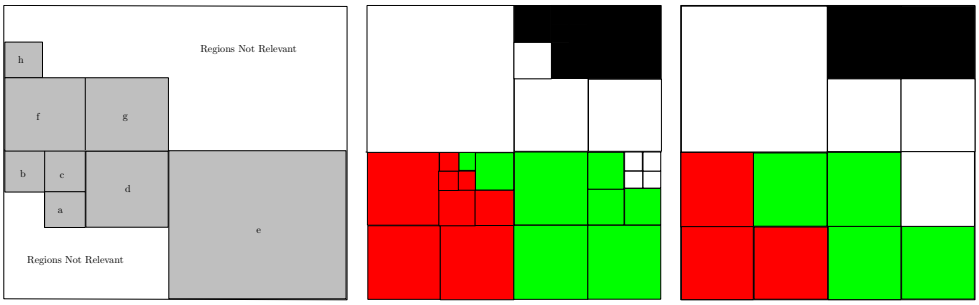

(a) (b) (c)

Fig. 3.3. (a) *A minimal list of quadrants covering the local domain on some processor, and* (b) *A Morton ordering based partition of a quadtree across* 4 *processors, and* (c) *Coarse quadrants and partition produced by using the quadtree shown in* (b) *as input to* `Algorithm 4`.

Algorithm 4. PARTITIONING OCTANTS INTO LARGE CONTIGUOUS BLOCKS (PARALLEL) - `BlockPartition`

Input: A distributed sorted list of octants, F.
Output: A list of the blocks, G. F is re-distributed, but the relative order of the octants is preserved.
Work: $\mathcal{O}(n)$, where $n = \texttt{len}(F)$.
Storage: $\mathcal{O}(n)$, where $n = \texttt{len}(F)$.
Time: Refer to the Appendix C.

1. $T \leftarrow$ `CompleteRegion`$(F[1], F[\texttt{len}(F)])$ (Algorithm 2)
2. $C \leftarrow \{x \in T \mid \forall y \in T,\ \mathcal{L}(x) \leq \mathcal{L}(y)\}$
3. $G \leftarrow$ `CompleteOctree`(C) (Algorithm 3)
4. **for each** $g \in G$
5. `weight`$(g) \leftarrow \texttt{len}(F_{global} \cap \{g, \{\mathcal{D}(g)\}\})$
6. **end for**
7. `Partition`(G) (Algorithm 5)
8. $F \leftarrow F_{global} \cap \{\{g, \{\mathcal{D}(g)\}\},\ \forall\ g \in G\}$

top-down approach and create a coarse regular grid, which can be divided[11] amongst the processors. However, this approach does not take load balancing into account since it does not use the underlying data distribution. Alternatively, one could use a space-filling curve to sort the octants and then partition them so that every processor gets an almost equal sized chunk of octants, contiguous in this order. Two desirable qualities of any partitioning strategy are load balancing, and minimization of overlap between the processor domains. We use our parallel bottom-up coarsening strategy to achieve these. First, we construct a global coarse linear octree based on the underlying distribution as described in section 3.4. In order to assign these coarse blocks to the processors, we first compute the load of each block by computing the number of original octants that lie within each of these blocks. The blocks are then distributed across the processors such that the total weight on each processor is roughly the same[12]. Note that the domain occupied by the blocks and the original octants on any given processor is not the same, but it does overlap to a large extent. The overlap is guaranteed by the fact that both are sorted according to the Morton ordering and that the partitioning was based on the same weighting function (i.e., the number of original octants). The original octants are then partitioned to align with the coarse block boundaries. Algorithm 4 lists all the steps described above and Figures 3.3(b) and 3.3(c) illustrate a sample input to Algorithm 4 and the corresponding output, respectively.

3.5. Balancing large linear octrees in parallel. Balance refinement is the process of refining (subdividing) nodes in a complete linear octree, which fail to satisfy the balance constraint described in Section 2.2. The nodes are refined until all

[11] If we create a regular grid at level l then the number of cells will be $n = 2^{dl}$, where d is the dimension. l is chosen in such a way that $n > p$.

[12] Some of the coarse blocks could be split if it facilitates achieving better load balance across the processors.

Algorithm 5. Partitioning a distributed list of octants (parallel) - `Partition`

Input: A distributed list of octants, W.
Output: The octants re-distributed across processors so that the total weight on each processor is roughly the same. The relative order of the octants is preserved.
Work: $\mathcal{O}(n)$, where $n = \texttt{len}(W)$.
Storage: $\mathcal{O}(n)$, where $n = \texttt{len}(W)$.

1. $S \leftarrow$ **Scan**(weight(W))
2. **if** rank = $(n_p - 1)$
3. TotalWeight $\leftarrow \max(S)$
4. **Broadcast**(TotalWeight)
5. **end if**
6. $\bar{w} \leftarrow \frac{\texttt{TotalWeight}}{n_p}$
7. $k \leftarrow$ (TotalWeight) $\mod n_p$
8. **for** $p \leftarrow 1$ **to** n_p
9. **if** $p \leq k$
10. $Q \leftarrow \{x \in W \mid (p-1).(\bar{w}+1) \leq S(x) < p.(\bar{w}+1)\}$
11. **else**
12. $Q \leftarrow \{x \in W \mid (p-1).\bar{w} + k \leq S(x) < p.\bar{w} + k\}$
13. **end if**
14. $Q_{tot} \leftarrow Q_{tot} + Q$
15. **Send**(Q, $(p-1)$)
16. **end for**
17. $R \leftarrow$ **Receive**()
18. $W \leftarrow W - Q_{tot} + R$

their descendants, which are created in the process of subdivision, satisfy the balance constraint. These subdivisions could in turn introduce new imbalances and so the process has to be repeated iteratively. The fact that an octant can affect octants not immediately adjacent to is known as the *ripple effect*.

We use a two-stage balancing scheme: first we perform local balancing on each processor, and follow this up by balancing across the inter-processor boundaries. One of the goals is to get a union of blocks (select non-leaf nodes of the octree) to reside on each processor so that the surface area and thereby the corresponding inter-processor boundaries are minimized. Determining whether a given partition provides the minimal surface area[13] is NP complete and determining the optimal partition is NP hard, since the problem is equivalent to the set-covering problem [8].

We use the parallel coarsening and partitioning algorithm (described in sections 3.4 and 3.4.1) to construct coarse blocks on each processor and to distribute the underlying octants. By construction, the domains covered by these blocks are disjoint and the union of these blocks covers the entire domain. We use the blocks as a

[13]The number of cells at the boundary depends on the underlying distribution and cannot be known *a priori*. This further complicates the balancing algorithm.

means to minimize the number of octants that need to be split due to inter-processor violations of the 2:1 balancing rule.

3.5.1. Local balancing. There are two approaches for balancing a complete octree. In the first approach, every node constructs the coarsest possible neighbors satisfying the balance constraint, and subsequently duplicates and overlaps are removed [4]. We describe this approach in Algorithm 6. In an alternative approach, the nodes search for neighbors and resolve any violations of the balance constraint [33, 35]. The main advantage of the former approach is that constructing nodes is inexpensive, since it does not involve any searches. However, this could produce a lot of duplicates and overlaps making the linearizing operations expensive. Another disadvantage of this approach is that it cannot handle incomplete domains, and can only operate on subtrees. The advantage of the second approach is that the list of nodes is complete and linear at any stage in the algorithm. The drawback, however, is that searching for neighbors is an expensive operation. Our algorithm uses a hybrid approach: it keeps the number of duplicates and overlaps to a minimum and also reduces the search space thereby reducing the cost of the searching operation. The complete linear octree is first partitioned into coarse blocks using the algorithm described in Section 3.4.1. The descendants of any block, which are present in the fine octree, form a linear subtree with this block as its root. This block-subtree is first balanced using the approach described in Section 3.5.2; the size of this tree will be relatively small, and hence the number of duplicates and overlaps will be small too. After balancing all the blocks, the inter-block boundaries in each processor are balanced using a variant of the *ripple propagation* algorithm [35] described in Section 3.5.4. The performance improvements from using the combined approach are presented in Section 4.2.

3.5.2. Balancing a local block. In principle, Algorithm 6 can be used to construct a complete balanced subtree of this block for each octant in the initial unbalanced linear subtree. Note that these balanced subtrees may have substantial overlap. Hence, Algorithm 8 is used to remove these overlaps. Lemma 3.1 shows that this process of merging these different balanced subtrees results in a complete linear balanced subtree. However, this implementation would be inefficient due to the number of overlaps, which would in turn increase the storage costs and also make the subsequent operations of sorting and removing duplicates and overlaps more expensive. Instead, we interleave the two operations: constructing the different complete balanced subtrees and merging them. The overall scheme is described in Algorithm 7.

We note that a list of octants form a balanced complete octree, if and only if for every octant all its neighbors are at the same level as this octant or one level finer or one level coarser. Hence, the coarsest possible octants in a complete octree that will be balanced against this octant are the siblings and the neighbors at the level of this octant's parent. Starting with the finest level and iterating over the levels up to but not including the level of the block, the coarsest possible (without violating the balance constraint) neighbors (Figure 3.4) of every octant at this level in the current tree (union of the initial unbalanced linear subtree and newly generated octants) are generated. After processing all the octants at any given level, the list of newly introduced coarse octants is merged with the previous list of octants at this level and duplicate octants are removed. The newly created octants are included while working on subsequent levels. Algorithm 8 still needs to be used in the end to remove overlaps, but the working size is much smaller now compared to the earlier case (Algorithm 6). To avoid redundant work and to reduce the number of duplicates to be removed

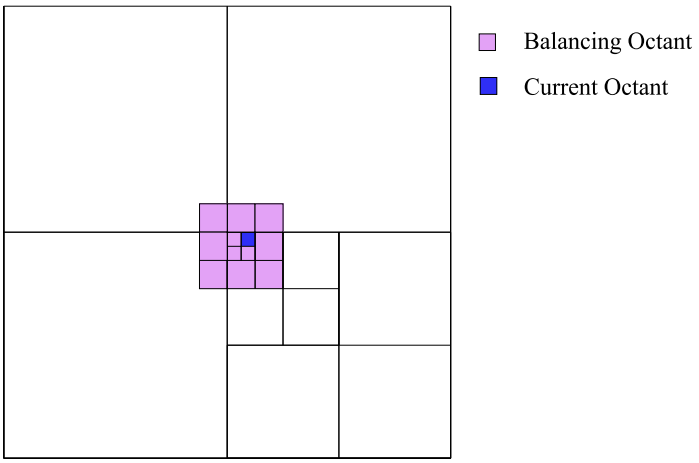

FIG. 3.4. *The minimal list of balancing quadrants for the current quadrant is shown. This list of quadrants is generated in one iteration of* `Algorithm 6`.

in the end, we ensure that no two elements in the working list at any given level are siblings of one another. This can be done in a linear pass on the working list for that level as shown in Algorithm 7.

LEMMA 3.1. *Let T_1 and T_2 be two complete balanced linear octrees with n_1 and n_2 number of potential ancestors respectively, then*

$$T_3 = (T_1 \cup T_2) - \left(\sum_{i=1}^{n_1} \{\mathcal{A}(T_1[i])\}\right) - \left(\sum_{j=1}^{n_2} \{\mathcal{A}(T_2[j])\}\right)$$

is a complete linear balanced octree.

Proof. $T_4 = (T_1 \cup T_2)$ is a complete octree. Now,

$$\left(\left(\sum_{i=1}^{n_1} \{\mathcal{A}(T_1[i])\}\right) + \left(\sum_{j=1}^{n_2} \{\mathcal{A}(T_2[j])\}\right)\right) = \left(\sum_{k=1}^{n_3} \{\mathcal{A}(T_4[k])\}\right)$$

So, $T_3 = \left(T_4 - \left(\sum_{k=1}^{n_3} \{\mathcal{A}(T_4[k])\}\right)\right)$ is a complete linear octree.
Now, suppose that a node $N \in T_3$ has a neighbor $K \in T_3$ such that $\mathcal{L}(K) \geq (\mathcal{L}(N) + 2)$. It is obvious that exactly one of N and K must be present in T_1 and the other must be present in T_2. Without loss of generality, assume that $N \in T_1$ and $K \in T_2$. Since T_2 is complete, there exists at least one neighbor of K,$L \in T_2$, which overlaps N. Also, since T_2 is balanced $\mathcal{L}(L) = \mathcal{L}(K)$ or $\mathcal{L}(L) = (\mathcal{L}(K) - 1)$ or $\mathcal{L}(L) = (\mathcal{L}(K) + 1)$. So, $\mathcal{L}(L) \geq (\mathcal{L}(N) + 1)$. Since L overlaps N and since $\mathcal{L}(L) \geq (\mathcal{L}(N) + 1)$, $L \in \{\mathcal{D}(N)\}$. Hence, $N \notin T_3$. This contradicts the initial assumption. Therefore, T_3 is also balanced. ∎

3.5.3. Searching for neighbors. A leaf needs to be refined if and only if the level of one of its neighbors is at least 2 levels finer than its own. In terms of a search this presents us two options: search for coarser neighbors or search for finer neighbors. It is much easier to search for coarser neighbors than it is to search for finer

Algorithm 6. CONSTRUCTING A COMPLETE BALANCED SUBTREE OF AN OCTANT, GIVEN ONE OF ITS DESCENDANTS (SEQUENTIAL)

Input: An octant, N, and one of its descendants, L.
Output: Complete balanced subtree, R.
Work: $\mathcal{O}(n \log n)$, where $n =$ len(R).
Storage: $\mathcal{O}(n)$, where $n =$ len(R).

```
1.  W ← L, T ← ∅, R ← ∅
2.  for l ← D_max to (L(N) + 1)
3.      for each w ∈ W
4.          R ← R + w + S(w)
5.          T ← T + N(P(w), l − 1)
6.      end for
7.      W ← T, T ← ∅
8.  end for
9.  Sort(R)
10. RemoveDuplicates(R)
11. R ← Linearise(R) ( Algorithm 8 )
```

Algorithm 7. BALANCING A LOCAL BLOCK (SEQUENTIAL) - BalanceSubtree

Input: An octant, N, and a partial list of its descendants, L.
Output: Complete balanced subtree, R.
Work: $\mathcal{O}(n \log n)$, where $n =$ len(R).
Storage: $\mathcal{O}(n)$, where $n =$ len(R).

```
1.  W ← L, P ← ∅, R ← ∅
2.  for l ← D_max to (L(N) + 1)
3.      Q ← {x ∈ W | L(x) = l}
4.      Sort(Q)
5.      T ← {x ∈ Q | S(x) ∉ T}
6.      for each t ∈ T
7.          R ← R + t + S(t)
8.          P ← P + N(P(t), l − 1)
9.      end for
10.     P ← P + {x ∈ W | L(x) = l − 1}
11.     W ← {x ∈ W|L(x) ≠ l − 1}
12.     RemoveDuplicates(P)
13.     W ← W + P, P ← ∅
14. end for
15. Sort(R)
16. R ← Linearise(R) ( Algorithm 8 )
```

Algorithm 8. Removing overlaps from a sorted list of octants (sequential) - `Linearise`

Input: `A sorted list of octants,` W.
Output: R, `an octree with no overlaps.`
Work: $\mathcal{O}(n)$, `where` $n =$ `len`(W).
Storage: $\mathcal{O}(n)$, `where` $n =$ `len`(W).

```
1.  for i ← 1 to (len(W) − 1)
2.      if (W[i] ∉ {A(W[i + 1])})
3.          R ← R + W[i]
4.      end if
5.  end for
6.  R ← R + W[len(W)]
```

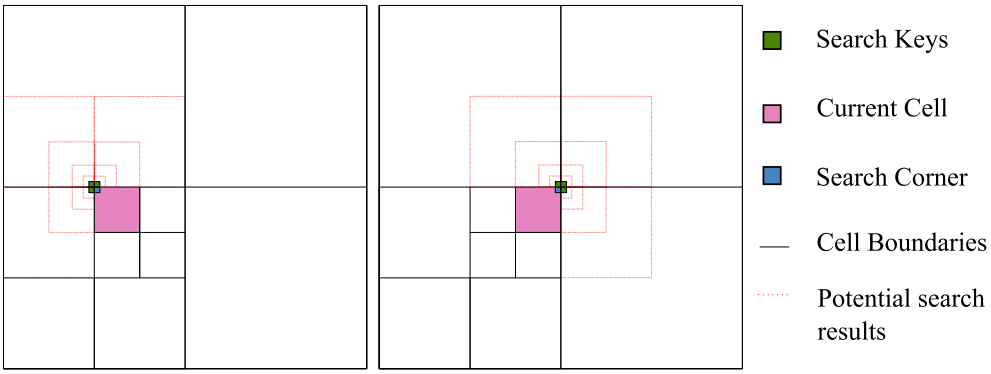

Fig. 3.5. *To find neighbors coarser than the current cell, we first select the finest cell at the far corner. The far corner is the one that is not shared with any of the current cell's siblings. The neighbors of this corner cell are determined and used as the search keys. The search returns the greatest cell lesser than or equal to the search key. The possible candidates in a complete linear quadtree, as shown, are ancestors of the search key.*

neighbors. If we consider the 2D case, only 3 neighbors coarser than the current cell need to be searched for. However, the number of potential neighbors finer than the cell is extremely large, (in 2D it is $2 \cdot 2^{D_{max}-l} + 3$, where l is the level of the current quadrant), and therefore not practical to search. In addition the search strategy depends on the way the octree is stored; the pointer based approach being more popular [4, 33], but has the overhead that it has to be rebuilt every time octants are communicated across processors. In the proposed approach the octree is stored as a linear octree in which the octants are sorted globally in the ascending Morton order, allowing us to search in $\mathcal{O}(\log n)$.

In order to find neighbors coarser than the current cell, we use the approach illustrated in Figure 3.5. First, the finest cell at the far corner (marked as 'Search Corner' in Figure 3.5) is determined. This is the corner that this octant shares with its parent. This is also the corner diagonally opposite to the corner common to all the

siblings of the current cell[14]. The neighbors (at the finest level) of this cell (N) are then selected and used as the search keys. These are denoted by $\mathcal{N}^s(N, D_{max})$. The maximum lower bound[15] for the given search key is determined by searching within the complete linear octree. In a complete linear octree, the maximum lower bound of a search key returns its finest ancestor. If the search result is at a level finer than or equal to the current cell then it is guaranteed that no coarser neighbor can exist in that direction. This idea can be extended to incomplete linear octrees (including multiply connected domains). In this case, the result of a search is ignored if it is not an ancestor of the search key.

3.5.4. Ripple propagation. A variant (Algorithm 9) of the *prioritized ripple propagation* algorithm first proposed by Tu et al. [33], modified to work with linear octrees, is used to balance the boundary leaves. The algorithm selects all leaves at a given level (successively decreasing levels starting with the finest), and searches for neighbors coarser than itself. A list of balancing descendants[16] for neighbors that violate the balance condition are stored. At the end of each level, any octant that violated the balance condition is replaced by a complete linear subtree. This subtree can be obtained either by using the sequential version of Algorithm 3 or by using Algorithm 10, which is a variant of Algorithm 7. Both the algorithms perform equally well.[17]

One difference with earlier versions of the ripple propagation algorithm is that our version works with incomplete domains. In addition, earlier approaches [4, 33, 35] have used pointer-based representations of the local octree, which incurs the additional cost of constructing the pointer-based tree from the linear representation and also increases the memory footprint of the octree as 9 additional pointers[18] are required per octant. The work and storage costs incurred for balancing using the proposed algorithm to construct n balanced octants are $\mathcal{O}(n \log n)$ and $\mathcal{O}(n)$, respectively. This is true irrespective of the domain, including domains that are not simply connected.

3.5.5. Insulation against the ripple-effect. An interesting property of complete linear octrees is that a boundary octant cannot be finer than its internal neighbors[19] (Figure 3.6(a)) [33]. So, if a node (at any level) is internally balanced then to balance it with all its neighboring domains, it is sufficient to appropriately refine the internal boundary leaves[20]. The interior leaves need not be refined any further. Since the interior leaves are also balanced against all their neighbors, they will not force any other octant to split. Hence, interior octants do not participate in the remaining stages of balancing.

Observe that the phenomenon with interior octants described above is only an

[14] We do not need to search in the direction of the siblings.

[15] The greatest cell lesser than or equal to the search key is referred to as its maximum lower bound.

[16] Balancing descendants are the minimum number of descendants that will balance against the octant that performed the search.

[17] We indicate which algorithms are parallel and which are sequential. In our notation the sequential algorithms are sometimes invoked with a distributed object: it is implied with the input is the local instance of the distributed object.

[18] One pointer to the parent and eight pointers to its children.

[19] A neighbor of a boundary octant that does not touch the boundary is referred to as an internal neighbor of the boundary octant.

[20] We refer to the descendants of a node that touch its boundary from the inside as its internal boundary leaves.

Algorithm 9. Ripple propagation on incomplete domains (sequential) - `Ripple`

Input: L, a sorted incomplete linear octree.
Output: W, a balanced incomplete linear octree.
Work: $\mathcal{O}(n \log n)$, where $n =$ len(L).
Storage: $\mathcal{O}(n)$, where $n =$ len(L).

```
1.  W ← L
2.  for l ← D_max to (L(N) + 1)
3.      for each w ∈ W
4.          if L(w) = l
5.              K ← search_keys(w) ( Section 3.5.3 )
6.              (B, J) ← maximum_lower_bound (K, W)
                          (J is the index of B in W)
7.              for each (b, j) ∈ (B, J) | l > (L(b) + 1) & b ∈ A(K)
8.                  T[j] ← T[j] + ({N^s (w, (l − 1))} ∩ {A(K)})
9.              end for
10.         end if
11.     end for
12.     for i ← 1 to len(W)
13.         if T[i] ≠ ∅
14.             R ← R+ CompleteSubtree(W[i], T[i]) ( Algorithm 10 )
15.         else
16.             R ← R + W[i]
17.         end if
18.     end for
19.     W ← R, T, R ← ∅
20. end for
```

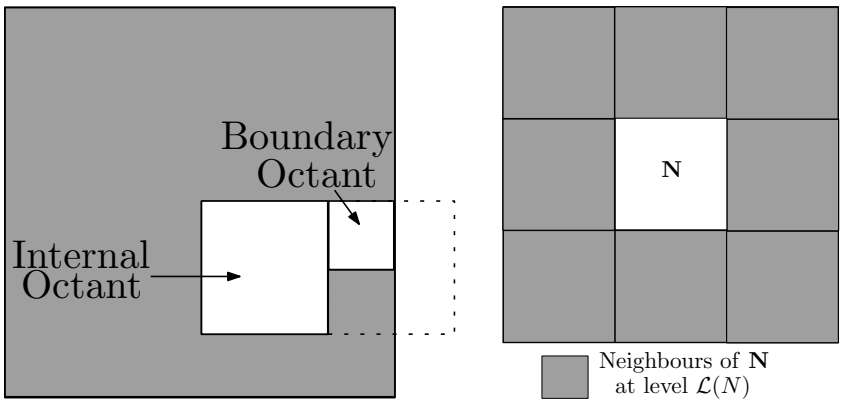

FIG. 3.6. (a) *A boundary octant cannot be finer than its internal neighbors, and* (b) *an illustration of an insulation layer around octant N. No octant outside this layer of insulation can force a split on N.*

Algorithm 10. COMPLETING A LOCAL BLOCK (SEQUENTIAL) - `CompleteSubtree`

Input: `An octant,` N, `and a partial list of its descendants,` L.
Output: `Complete subtree,` R.
Work: $\mathcal{O}(n \log n)$, `where` $n =$ `len`(R).
Storage: $\mathcal{O}(n)$, `where` $n =$ `len`(R).

```
1.  W ← L
2.  for l ← D_max to L(N) + 1
3.      Q ← {x ∈ W | L(x) = l}
4.      Sort(Q)
5.      T ← {x ∈ Q | S(x) ∉ T}
6.      for each t ∈ T
7.          R ← R + t + S(t)
8.          P ← P + S(P(t))
9.      end for
10.     P ← P + {x ∈ W | L(x) = l − 1}
11.     W ← {x ∈ W | L(x) ≠ l − 1}
12.     RemoveDuplicates(P)
13.     W ← W + P, P ← ∅
14. end for
15. Sort(R)
16. R ← Linearise(R) ( Algorithm 8 )
```

example of a more general property:

DEFINITION 2. *For any octant,* N, *in the octree, we refer to the union of the domains occupied by its potential neighbor's at the same level as* N $(\mathcal{N}(N, \mathcal{L}(N)))$ *as the insulation layer around octant* N. *This will be denoted by* $\mathcal{I}(N)$.

PROPERTY 1. *No octant outside the insulation layer around octant* N *can force* N *to split (Figure 3.6(b)).*

This property allows us to decouple the problem of balancing and allows us to work on only a subset of nodes in the octree and yet ensure that the entire octree is balanced.

3.5.6. Balancing inter-processor boundaries. After the intra-processor, and inter-block boundaries are balanced, the inter-processor boundaries need to be balanced. Unlike the internal leaves (Section 3.5.5), the octants on the boundary do not have any insulation against the ripple-effect. Moreover, a ripple can propagate across multiple processors. Most approaches to perform this balance have been based on extensions of the sequential ripple algorithm to a parallel case by performing parallel searches. In an earlier attempt we developed efficient parallel search strategies allowing us to extend our sequential balancing algorithms to the parallel case. Although this approach works well for small problems on a small number of processors, it shows suboptimal isogranular scalability as has been seen with other similar approaches to the problem [35]. The main reason is iterative communication. Although there are many examples of scalable parallel algorithms that involve iterative communication,

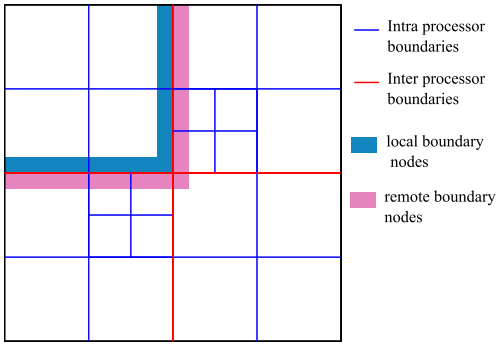

FIG. 3.7. *A coarse quadtree illustrating inter and intra processor boundaries. First, every processor balances each of its local blocks. Then, each processor balances the cells on its intra-processor boundaries. The octants that lie on inter-processor boundaries are then communicated to the respective processors and each processor balances the combined list of local and remote octants.*

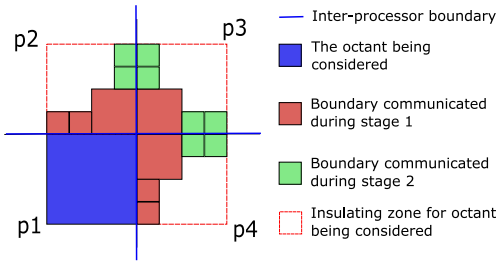

FIG. 3.8. *Communication for inter-processor balancing is done in two stages: First, every octant on the inter-processor boundary is communicated to processors that overlap with its insulation layer. Next, all the local inter-processor boundary octants that lie in the insulation layer of a remote octant received from another processor are communicated to that processor.*

they overlap communication with computation to reduce the overhead associated with communication [13, 29]. Currently, there is no method that overlaps communication with computation for the balancing problem. Thus, any algorithm that uses iterative parallel searches for balancing octrees will be have high communication costs.

In order to avoid parallel searches, the problem of balancing is decoupled. In other words, each processor works independently without iterative communication. To achieve this, two properties are used: `(1)` the only octants that need to be refined after the local balancing stage are the ones that lie on inter-processor boundaries and `(2)` an artificial insulation layer (Property 1) for the boundary octants can be constructed with little communication overhead (Section 3.5.7). The construction of this insulation layer is done in two stages (Figure 3.8): First, every local octant on the inter-processor boundary (Figure 3.7) is communicated to processors that overlap with its insulation layer. These processors can be determined by comparing the local boundary octants against the global coarse blocks. In the second stage of communication, all the local inter-processor boundary octants that overlap with the insulation layer of a remote octant received from another processor are communicated to that processor. Octants that were communicated in the first stage are not communicated to the same processor again. For simplicity, Algorithm 11 only describes a naïve implementation for determining teh octants that need to be communicated at this stage. This, however, can be performed much more efficiently using the results of Lemma

3.2 and Lemma 3.3. After this two-stage communication, each processor balances the union of the local and remote boundary octants using the ripple propagation based method (Section 3.5.4). At the end only the octants spanning the original domain spanned by the processors are retained. Although there is some redundancy in the work, it is compensated by the fact that we avoid iterative communications. Section 3.5.7 gives a detailed analysis of the communication cost involved.

LEMMA 3.2. *If octants a and $b > a$ do not overlap, then there can be no octant $c > b$ that overlaps a.*

Proof. If a and c overlap, then either $a \in \{\mathcal{A}(c)\}$ or $a \in \{\mathcal{D}(c)\}$. Since $c > a$, the latter is a direct violation of Property 4 and hence is impossible. Hence, assume that $c \in \{\mathcal{D}(a)\}$. By Property 9, $c \leq \mathcal{DLD}(a)$. Property 10 would then imply that $b \in \{\mathcal{D}(a)\}$. Property 5 would then imply that a and b must overlap. Since, this is not true our initial assumption must be wrong. Hence, a and c can not overlap. ∎

LEMMA 3.3. *Let N be an inter-processor boundary octant belonging to processor q and let it be sent to processor p during the first stage of communication. If all elements in $\mathcal{I}(N)$ overlap some octant in q or p, then the inter-processor boundary octants on p that overlap with some element in $\mathcal{I}(N)$ and that were not communicated to q in the first stage, will not force a split on N.*

Proof. Note that at this stage both p and q are internally balanced. Thus, N will be forced to split if and only if there is an inter-processor boundary octant, a, on the p touching an octant, b, on q such that $\mathcal{L}(a) > (\mathcal{L}(b) + 1)$ and when b is split it starts a cascade of splits on octants in q that in turn force N to split. Since every inter-processor boundary octant is sent to all its adjacent processors, a must have been sent to q during the first stage of communication. ∎

The keys steps involved in parallel balancing are summarized in Algorithm 11.

3.5.7. Communication costs for parallel balancing. Here, we compare the communication costs associated with the two approaches (upfront communication versus iterative communication). Let us assume that prior to parallel balancing there are a total of N octants in the global octree. The octants that lie on the inter-processor boundary can be classified based on the *degree of the face*[21] that they share with the inter-processor boundary. We use N_k to represent the number of octants that touch any m-dimensional face ($m \in [0, k]$) of the inter-processor boundary.

Note that all vertex boundary octants are also edge and face boundaries and that all edge boundary octants are also face boundary octants. Therefore we have, $N \geq N_2 \geq N_1 \geq N_0$, and for $N \gg n_p$, we have $N \gg N_2 \gg N_1 \gg N_0$.

Although it is theoretically possible that an insulation layer of some octant encloses the domains controlled by multiple processors, it is unlikely for dense octrees. Hence, it is reasonable to assume that in the first stage the octants are only communicated via near-neighbor point-to-point operations. Under this assumption, the total number of octants of a d-tree that need to be communicated in the first stage of the proposed approach is given by

$$N_u = \sum_{k=1}^{d} 2^{d-k} N_{k-1}. \tag{3.1}$$

Consider the example shown in Figure 3.9. The domain on the left is partitioned into two regions, and in this case all boundary octants need to be transmitted to

[21] A corner is a 0-degree face, an edge is a 1-degree face and a face is a 2-degree face.

Algorithm 11. Balancing complete linear octrees (parallel)

Input: A distributed sorted complete linear octree, L.
Output: A distributed complete balanced linear octree, R.
Work: $\mathcal{O}(n \log n)$, where $n = \texttt{len}(L)$.
Storage: $\mathcal{O}(n)$, where $n = \texttt{len}(L)$.
Time: Refer to Section 3.5.7.

1. $B \leftarrow$ BlockPartition(L) (Algorithm 4)
2. $C \leftarrow$ BalanceSubtree (B, L) (Algorithm 7.)
3. $D \leftarrow \{x \in C \mid \exists\ z \in \{\mathcal{N}(x)\} \mid \{\{\{z, \{\mathcal{A}(z)\}\} - \{\mathcal{A}(x)\}\} \cap B\} \neq \emptyset\}$
   (intra-processor boundary octants)
4. $S \leftarrow$ Ripple(D) (Algorithm 9)
5. $F \leftarrow$ Linearise($C \cup S$)
6. $G \leftarrow \{x \in F \mid \exists\ z \in \{\mathcal{N}(x)\} \mid \{\{z, \{\mathcal{A}(z)\}\} \cap B\} = \emptyset\}$
   (inter-processor boundary octants)
7. **for each** $g \in G$
8. **for each** $b \in B_{global} - B$
9. **if** $\{b \cap \mathcal{I}(g)\} \neq \emptyset$
10. **Send**(g, rank(b))
11. **end if**
12. **end for**
13. **end for**
14. $T \leftarrow$ Receive()
15. **for each** $g \in G$
16. **for each** $t \in T$
17. **if** $\{g \cap \mathcal{I}(t)\} \neq \emptyset$
18. **if** g was not sent to rank(t) in Step 10
19. **Send**(g, rank(t))
20. **end if**
21. **end if**
22. **end for**
23. **end for**
24. $K \leftarrow$ Receive()
25. $H \leftarrow$ Ripple($G \cup T \cup K$)
26. $R \leftarrow \{x \in \{H \cup F\} \mid \{B \cap \{x, \{\mathcal{A}(x)\}\}\} \neq \emptyset\}$
27. $R \leftarrow$ Linearise(R) (Algorithm 8)

exactly one other processor. The addition of the additional boundary, in the figure on the right, does not affect most boundary nodes, except for the boundary octants that share a corner, i.e., a 0-dimensional face with the inter processor boundaries. These octants need to be sent to an additional 2 processors, and that is the reason we have a factor of 2^{d-k} in Equation 3.1. For the case of octrees, additional communication is incurred because of edge boundaries as well as vertex boundaries. Edge boundary octants need to be communicated to 2 additional processors whereas the vertex boundary octants need to be communicated to 4 additional processors (7 processors in all).

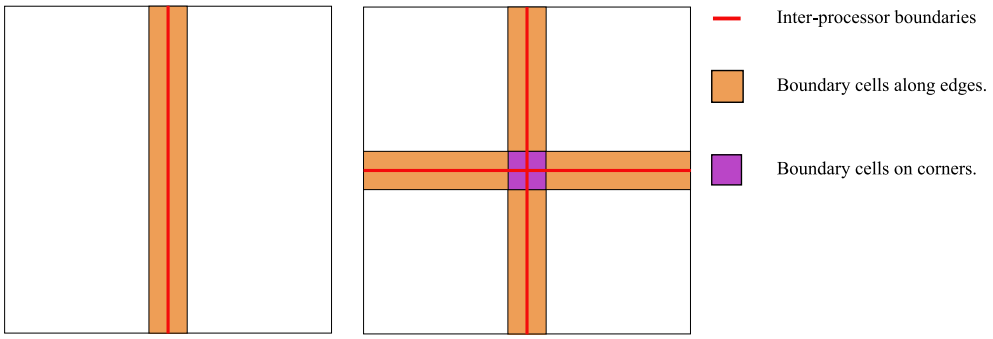

FIG. 3.9. *Cells that lie on the inter-processor boundaries. The figure on the left shows an inter-processor boundary involving 2 processors and the figure on the right shows an inter-processor boundary involving 4 processors.*

Now, we analyze the cost associated with the second communication step in our algorithm. Consider the example shown in Figure 3.8. Note that all the immediate neighbors of the octant under consideration (Octant on processor 1 in the figure), were communicated during the first stage. The octants that lie in the insulation zone of this octant and that were not communicated in the first stage are those that lie in a direction normal to the inter-processor boundary. However, most octants that lie in a direction normal to the inter-processor boundary are internal octants on other processors. As shown in Figure 3.8, the only octants that lie in a direction normal to one inter-processor boundary and are also tangential to another inter-processor boundary are the ones that lie in the shadow of some edge or corner boundary octant. Therefore, we only communicate $\mathcal{O}(N_1+N_0)$ octants during this stage. Since $N \gg n_p$ and $N_2 \gg N_1 \gg N_0$ for most practical applications, the cost for this communication step can be ignored.

The minimum number of search keys that need to be communicated in a search based approach is given by

$$N_s = \sum_{k=1}^{d} 2^{k-1} N_{k-1}. \tag{3.2}$$

Again considering the example shown in Figure 3.9, each boundary octant in the figure shown on the left, generates 3 search keys, out of which one lies on the same processor. The other two need to be communicated to the other processor. The addition of the extra boundary, in the figure on the right, does not affect most boundary nodes, except for the boundary octants that share a corner, i.e., a 0-dimensional face with the inter processor boundaries. These octants need to be sent to an additional processor, and that is the reason we have a factor of 2^{k-1} in Equation 3.2. It is important to observe the difference between the communication estimates for upfront communication, 3.1, with that of the search based approach, 3.2. For large octrees,

$$N_u \approx N_2,$$

while,

$$N_s \approx 4N_2.$$

Note, that in arriving at the communication estimate for the search based approaches, we have not accounted for the additional octants created during the inter-processor balancing. In addition, iterative search based approaches are further affected by communication lag and synchronization. Our approach in contrast requires no subsequent communication.

In conclusion, the communication cost involved in the proposed approach is lower than that of search based approaches[22].

4. Results. The performance of the proposed algorithms is evaluated by a number of numerical experiments, including fixed-size and isogranular scalabilty analysis. The algorithms were implemented in C++ using the MPI library. A variant of the sample sort algorithm was used to sort the points and the octants, which incorporates a parallel bitonic sort to sort the sample elements as suggested in [13]. PETSc [2] was used for profiling the code. All tests were performed on the Pittsburgh Supercomputing Center's TCS-1 terascale computing HP AlphaServer Cluster comprising of 750 SMP ES45 nodes. Each node is equipped with four Alpha EV-68 processors at 1 GHz and 4 GB of memory. The peak performance is approximately 6 Tflops, and the peak performance for the top-500 LINPACK benchmark is approximately 4 Tflops. The nodes are connected by the Quadrics interconnect, which delivers over 500 MB/s of message-passing bandwidth per node and has a bisection bandwidth of 187 GB/s. In our tests we have used 4 processors per node, wherever possible.

First, results from an experiment to compare different strategies for the local balancing stage is presented. This highlights the advantages of using the two-stage approach over existing approaches. Following this, results from the fixed-size and isogranular scalability analysis are presented.

4.1. Test Data. Point data of different sizes were generated for three different distributions; Gaussian, Log-normal and Regular. The regular distribution corresponds to a set of points uniformly distributed such that they are regularly spaced. Datasets were generated for all three distributions of increasing sizes that result in a balanced octree with octants ranging from 10^6(1M) to 10^9(1B). The fixed size scalability analysis was performed by selecting the 1M, 32M and 128M Gaussian point distributions to represent small, medium and large problems. Since the input sizes for different stages of the algorithm vary, we provide the input and output sizes for different stages in table 4.1. Since the regularly spaced points produce octrees that are inherently balanced, these datasets serve to estimate the communication costs associated with the parallel balancing algorithm.

4.2. Comparison between different strategies for the local balancing stage. In order to assess the advantages of using a two-stage approach for local balancing over existing methods, we compared the runtimes on different problem sizes. Since the comparison was for different strategies for local balancing, it does not involve any communication and hence was evaluated on a quad dual-core shared memory machine. We compared our two-stage approach, discussed in Section 3.5.1, with two other approaches; the first approach is the prioritized ripple propagation idea applied on the entire local domain [35], and the second approach is to use ripple propagation in 2 stages, where the local domain is first split into coarser blocks[23] and ripple propagation is applied first to each local block and then repeated on the

[22] We are assuming that both the approaches use the same partitioning of octants.

[23] The same partitioning strategy as used in our two-stage algorithm was used to obtain the coarser blocks.

Problem size	Gaussian			log-Normal			Regular	
	Points	Unbalanced Octants	Balanced Octants	Points	Unbalanced Octants	Balanced Octants	Points	Octants
1M	180K	607K	995.4K	180K	607K	987.5K	405.2K	994.9K
2M	361K	1211K	2000.5K	361K	1214K	2016.9K	2097.2K	2097.2K
4M	720K	2434K	3973.9K	720K	2434K	3958.5K	2.4M	4.06M
8M	1466K	4912K	8.0M	1466K	4920K	8.1M	3.24M	7.96M
16M	2886K	9686K	16M	2886K	9712K	16M	16.78M	16.78M
32M	5.8M	19.6M	31.9M	5.8M	19.6M	31.8M	19.25M	32.53M
64M	11.7M	39.3M	64.4M	11.7M	39.3M	64.7M	25.93M	63.67M
128M	23.5M	79.3M	130.6M	23.5M	79.4M	130.1M	134.22M	134.22M
256M	47M	158.6M	256.8M	47M	158.3M	256.4M	153.99M	260.24M
512M	94M	315.1M	516.9M	94M	315.5M	520.8M	168.20M	337.66M
1B	164.5M	553.6M	914.9M	164.5M	555.2M	912.2M	1.07B	1.07B

TABLE 4.1

Input and output sizes for the construction and balancing algorithms for the scalability experiments on Gaussian, Log-Normal, and Regular point distributions. Note that Regular point distributions are inherently balanced, and therefore the number of octants is only reported once.

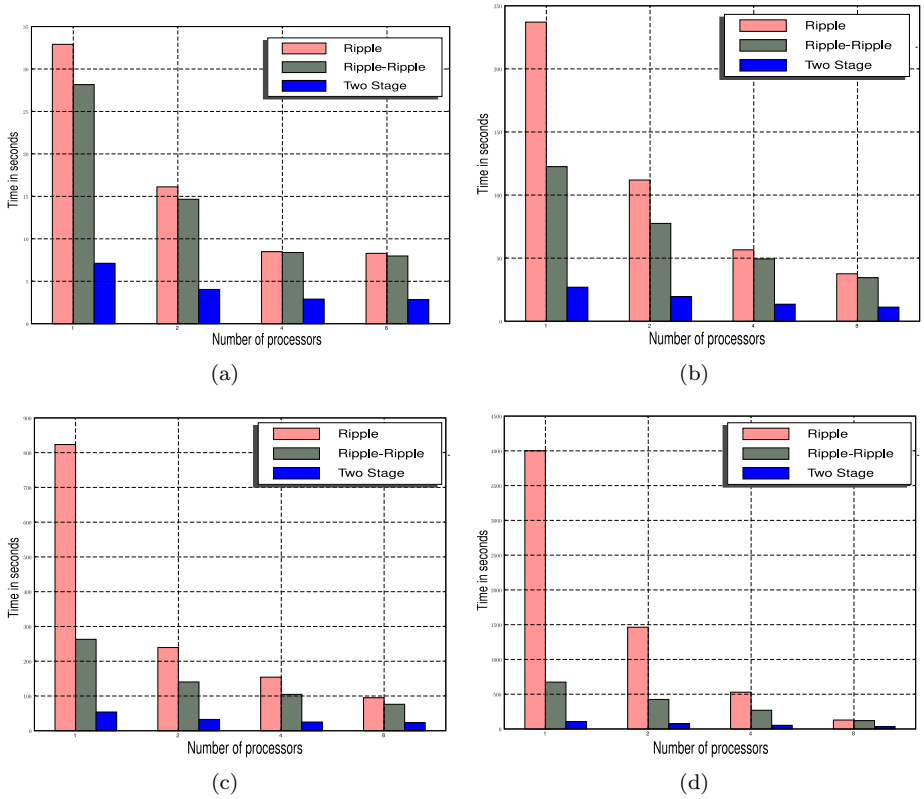

FIG. 4.1. *Comparison of three different approaches for balancing linear octrees* `(a)` *for a Gaussian distribution of* `1M` *octants,* `(b)` *for a Gaussian distribution of* `4M` *octants,* `(c)` *for a Gaussian distribution of* `8M` *octants, and* `(d)` *for a Gaussian distribution of* `16M` *octants.*

boundaries of all local blocks. Fixed size scalability analysis was performed to compare the above mentioned three approaches with problem sizes of 1, 4, 8, and 16 million points. The results are shown in in Figure 4.1. All three approaches demonstrate good fixed size scalability, but the two-stage approach demonstrate better absolute runtime.

4.3. Scalability analysis. In this section we provide experimental evidence on the good scalability of our algorithms. We present both fixed-size and isogranular

scalability analysis. Fixed size scalability was performed for different problem sizes to analyze the improvement in performance when the problem size is kept constant and the number of processors are increased. Fixed size scalability allows us to determine the problem sizes for which speedup can be expected by increasing the processor count. Isogranular scalability analysis is performed by tracking the execution time while increasing the problem size and the number of processors proportionately. By maintaining the problem size per processor (relatively) constant as the number of processors is increased, we can identify communication problems related to the size and frequency of the messages as well as global reductions and problems with algorithmic scalability.

One of the important components in our algorithms is the sample sort routine, which has a complexity of $\mathcal{O}(\frac{N}{n_p}\log\frac{N}{n_p} + n_p^2 \log n_p)$ if the samples are sorted using a serial sort. This causes problems when $N = \mathcal{O}(n_p^2)$ as the serial sort begins to dominate and results in poor scalability. For example, at $n_p = 1024$ we would require $N > 10^6$ to obtain good scalability. This presents some problems as it becomes difficult to fit arbitrarily large problems on a given processor. Using the parallel bitonic sort to sort the samples [13] reduces the complexity to $\mathcal{O}(\frac{N}{n_p}\log\frac{N}{n_p} + n_p \log n_p)$.

Isogranular scalability analysis was performed for all three distributions with an output size of roughly 1M octants per processor, for processor counts ranging from 1 to 1024. In all cases the tests were performed with a maximum of one point per octant. This resulted in octrees with over 1 billion octants in some cases. Since the regularly spaced distribution is inherently balanced, the input point sizes were much greater for this case than those for Gaussian and Log-normal distributions. Both the Gaussian and Log-normal distributions are imbalanced and it can be seen in Table 4.1 that on an average the number of unbalanced octants is 3 times the number of input points, and that the number of octants doubles as a result of balancing. For the regularly spaced distribution, we observe that in some cases the number of octants is the same as the number of input points (2M, 16M, 128M and 1B). These are special cases where the resulting grid is a regular grid .

Wall-clock timings, speedup and efficiency for the isogranular analysis for the three distributions are shown in Figures 4.2, 4.3, and 4.4. The isogranular analysis for the regularly spaced distribution allows us to estimate the overhead involved in communication, partitioning and in balancing the local blocks using Algorithm 7. The ripple propagation algorithm does not do any work in this experiment. The plots demonstrate the good isogranular scalability of the algorithm. We achieve near optimal isogranular scalability for all three distributions (50s per 10^6 octants per processor for the Gaussian and Log-normal distributions and 25s for the regularly spaced distribution.).

Fixed size scalability tests were also performed for three problem set sizes, small (1 million points), medium (32 million points) and large (128 million points), for the Gaussian distribution. These results are plotted in Figures 4.5, 4.6 and 4.7.

5. Conclusions. We have presented two new parallel algorithms for constructing and balancing large linear octrees on distributed memory machines. We have also tested MPI-based scalable parallel implementations for both the algorithms. Our algorithms have several important features:

- Experiments on three different types of input distributions demonstrate that the algorithms are insensitive to the underlying data distribution.
- Our algorithms avoid iterative communications and thus are able to achieve low absolute runtime and good scalability.

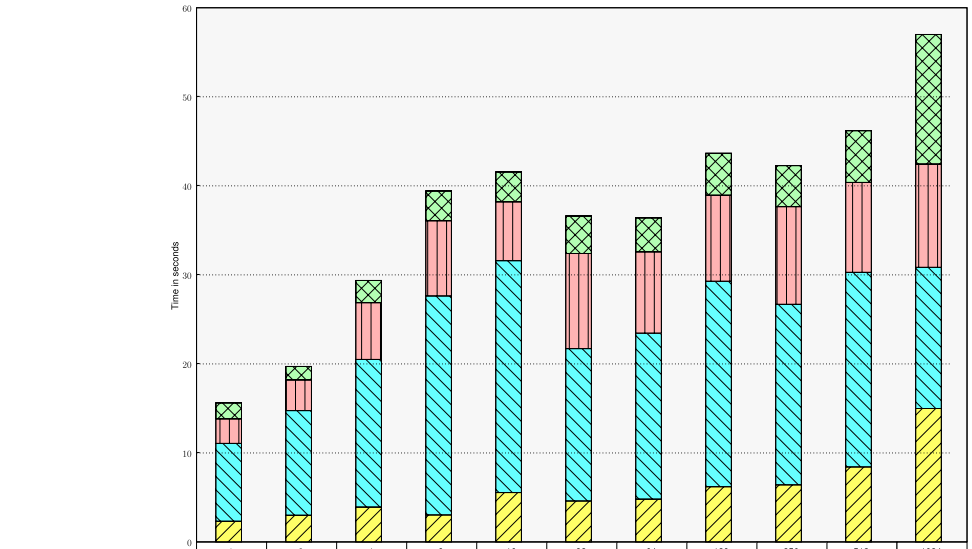

	1	2	4	8	16	32	64	128	256	512	1024
Construction	2.3	3.0	3.9	3.0	5.6	4.6	4.8	6.2	6.4	8.4	15
Internal Balance	8.7	11.8	16.6	24.6	26.0	17.1	18.6	23.1	20.3	21.8	15
Boundary Balance	2.8	3.4	6.4	8.4	6.6	10.7	9.1	9.7	11.0	10.1	11
Communication for Balance	1.8	1.5	2.5	3.3	3.4	4.2	3.8	4.7	4.6	5.8	14

FIG. 4.2. *Isogranular scalability for Gaussian distribution of* `1M` *octants per processor. From left to right, the bars indicate the time taken for the different components of our algorithms for increasing processor counts. The bar for each processor is partitioned into* `4` *sections. From top to bottom, the sections represent the time taken for* `(1)` *communication (including related pre-processing and post-processing) during balance refinement* `(Algorithm 11)`*,* `(2)` *balancing across intra and inter processor boundaries* `(Algorithm 9)`*,* `(3)` *balancing the blocks* `(Algorithm 7)` *and* `(4)` *construction from points* `(Algorithm 1)`*.*

- The experiments for comparing the performance of different algorithms for the local balancing stage demonstrated that the one proposed in this paper has a significantly lower running time than the others.
- We demonstrated scalability up to 1024 processors: we were able to construct and balance octrees with over 1 billion octants in less than a minute.

We need to consider the following factors to improve the performance of the proposed algorithms. In order to minimize communication costs, it is desirable to have as large coarse blocks as possible since the communication cost is proportional to the area of the inter-processor boundaries. However, too coarse blocks will increase the work for the local block balancing stage (Section 3.5.2). If additional local splits are introduced, then the intra-block boundaries increase causing the work load for the first ripple balance to increase. The local balancing step of the algorithm can be made more efficient by performing the local balancing recursively by estimating the correct size of the block that can be balanced by the search-free approach. Such an approach should be based on low-level architecture details, like the cache size.

Acknowledgments. The authors thank Santi Swaroop Adavani and Shravan Veerapaneni of University of Pennsylvania for providing us with an implementation of the parallel `Sample Sort` algorithm.

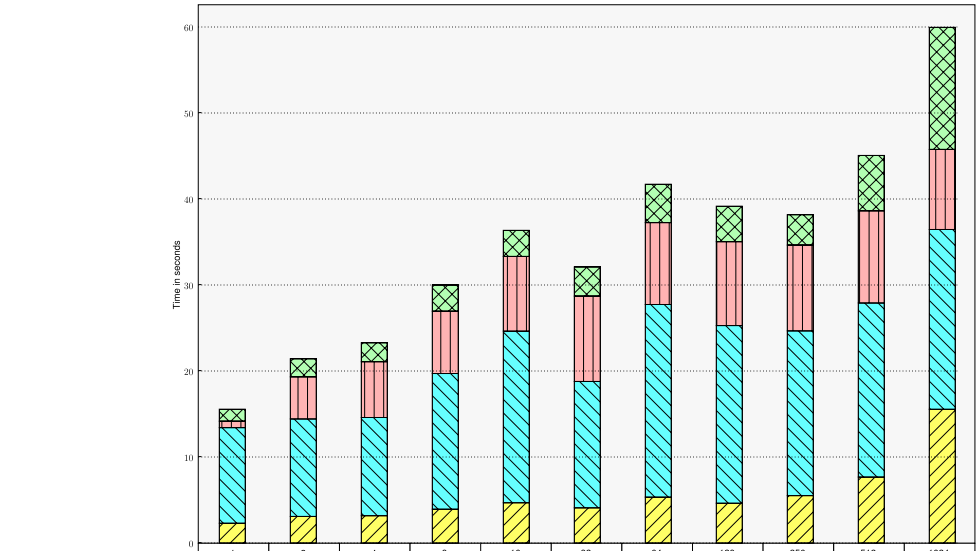

	1	2	4	8	16	32	64	128	256	512	1024
Construction	2.3	3.1	3.2	3.9	4.7	4.1	5.3	4.6	5.5	7.7	15
Internal Balance	11.1	11.3	11.4	15.8	20.0	14.7	22.4	20.7	19.2	20.2	20
Boundary Balance	0.7	4.9	6.5	7.2	8.7	9.9	9.5	9.8	10.0	10.7	9
Communication for Balance	1.4	2.1	2.2	3.1	3.0	3.4	4.5	4.1	3.5	6.4	14

FIG. 4.3. *Isogranular scalability for Log-normal distribution of* `1M` *octants per processor. From left to right, the bars indicate the time taken for the different components of our algorithms for increasing processor counts. The bar for each processor is partitioned into* `4` *sections. From top to bottom, the sections represent the time taken for* `(1)` *communication (including related pre-processing and post-processing) during balance refinement* `(Algorithm 11)`*,* `(2)` *balancing across intra and inter processor boundaries* `(Algorithm 9)`*,* `(3)` *balancing the blocks* `(Algorithm 7)` *and* `(4)` *construction from points* `(Algorithm 1)`*.*

REFERENCES

[1] D. AYALA, P. BRUNET, R. JUAN, AND I. NAVAZO, *Object representation by means of nonminimal division quadtrees and octrees*, ACM Trans. Graph., 4 (1985), pp. 41–59.

[2] S. BALAY, K. BUSCHELMAN, W. D. GROPP, D. KAUSHIK, M. G. KNEPLEY, L. C. MCINNES, B. F. SMITH, AND H. ZHANG, *PETSc Web page*, 2001. http://www.mcs.anl.gov/petsc.

[3] R. BECKER AND M. BRAACK, *Multigrid techniques for finite elements on locally refined meshes*, Numerical Linear Algebra with applications, 7 (2000), pp. 363–379.

[4] M. W. BERN, D. EPPSTEIN, AND S.-H. TENG, *Parallel construction of quadtrees and quality triangulations*, International Journal of Computational Geometry and Applications, 9 (1999), pp. 517–532.

[5] G. BERTI, *Image-based unstructured 3-d mesh generation for medical application*, in European Congress on Computational Methods in Applied Sciences and Engineering, Jyväskylä, Finland, 2004.

[6] P. BRUNET AND I. NAVAZO, *Solid representation and operation using extended octrees*, ACM Trans. Graph., 9 (1990), pp. 170–197.

[7] P. M. CAMPBELL, K. D. DEVINE, J. E. FLAHERTY, L. G. GERVASIO, AND J. D. TERESCO, *Dynamic octree load balancing using space-filling curves*, Tech. Report CS-03-01, Williams College Department of Computer Science, 2003.

[8] T. CORMAN, C. LEISERSON, AND R. RIVEST, *Introduction to Algorithms*, MIT Press, 1990.

[9] R. A. FINKEL AND J. L. BENTLEY, *Quad trees: A data structure for retrieval on composite keys*, Acta Inf., 4 (1974), pp. 1–9.

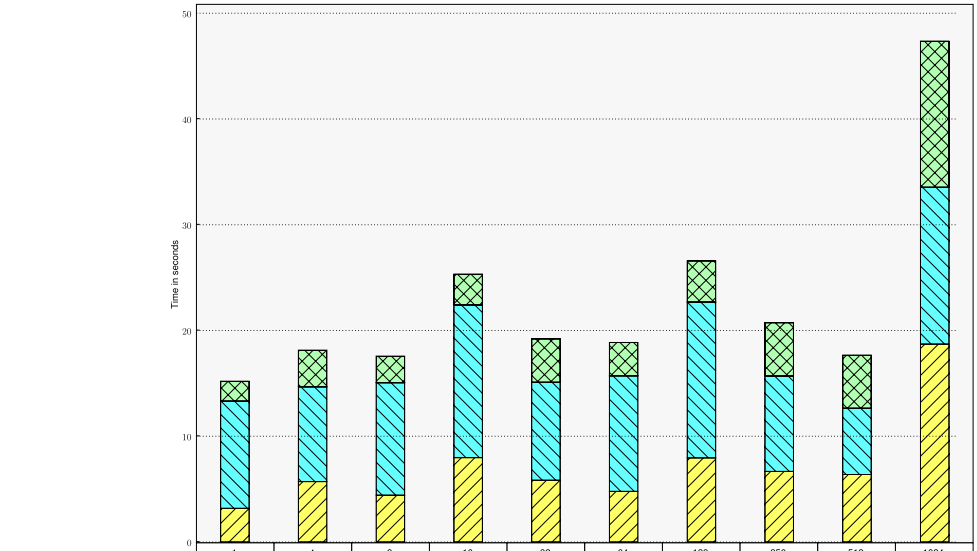

	1	4	8	16	32	64	128	256	512	1024
Construction	3.2	5.7	4.4	8.0	5.8	4.8	8.0	6.7	6.4	18
Internal Balance	10.1	9.0	10.6	14.4	9.3	10.9	14.7	9.0	6.3	14
Boundary Balance	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Communication for Balance	1.9	3.5	2.5	2.9	4.1	3.1	3.9	5.0	5.0	19

FIG. 4.4. *Isogranular scalability for uniformly spaced points with* `1M` *octants per processor. From left to right, the bars indicate the time taken for the different components of our algorithms for increasing processor counts. The bar for each processor is partitioned into* `4` *sections. From top to bottom, the sections represent the time taken for* `(1)` *communication (including related pre-processing and post-processing) during balance refinement* `(Algorithm 11)`*,* `(2)` *balancing across intra and inter processor boundaries* `(Algorithm 9)`*,* `(3)` *balancing the blocks* `(Algorithm 7)` *and* `(4)` *construction from points* `(Algorithm 1)`*. While both the input and output grain sizes remain almost constant for the Gaussian and LogNormal distributions, only the output grain size remains constant for the Uniform distribution. Hence, the trend seen in this study is a little different from those for the Gaussian and LogNormal distributions.*

[10] L. FREITAG AND R. LOY, *Adaptive, multiresolution visualization of large data sets using a distributed memory octree*, in ACM/IEEE 1999 Conference on Supercomputing, 13-18 Nov. 1999.

[11] H. FUCHS, G. D. ABRAM, AND E. D. GRANT, *Near real-time shaded display of rigid objects*, in SIGGRAPH '83: Proceedings of the 10th annual conference on Computer graphics and interactive techniques, New York, NY, USA, 1983, ACM Press, pp. 65–72.

[12] H. FUCHS, Z. M. KEDEM, AND B. F. NAYLOR, *On visible surface generation by a priori tree structures*, in SIGGRAPH '80: Proceedings of the 7th annual conference on Computer graphics and interactive techniques, New York, NY, USA, 1980, ACM Press, pp. 124–133.

[13] A. GRAMA, A. GUPTA, G. KARYPIS, AND V. KUMAR, *An Introduction to Parallel Computing: Design and Analysis of Algorithms*, Addison Wesley, second ed., 2003.

[14] D. M. GREAVES AND A. G. L. BORTHWICK, *Hierarchical tree-based finite element mesh generation*, International Journal for Numerical Methods in Engineering, 45 (1999), pp. 447–471.

[15] L. GREENGARD AND W. GROPP, *A parallel version of the fast multipole method-invited talk*, in Proceedings of the Third SIAM Conference on Parallel Processing for Scientific Computing, Philadelphia, PA, USA, 1989, Society for Industrial and Applied Mathematics, pp. 213–222.

[16] M. GRIEBEL AND G. ZUMBUSCH, *Parallel multigrid in an adaptive PDE solver based on hashing*, in Parallel Computing: Fundamentals, Applications and New Directions, Proceedings of the Conference ParCo'97, 19-22 September 1997, Bonn, Germany, E. H. D'Hollander, G. R. Joubert, F. J. Peters, and U. Trottenberg, eds., vol. 12, Amsterdam, 1998, Elsevier, North-Holland, pp. 589–600.

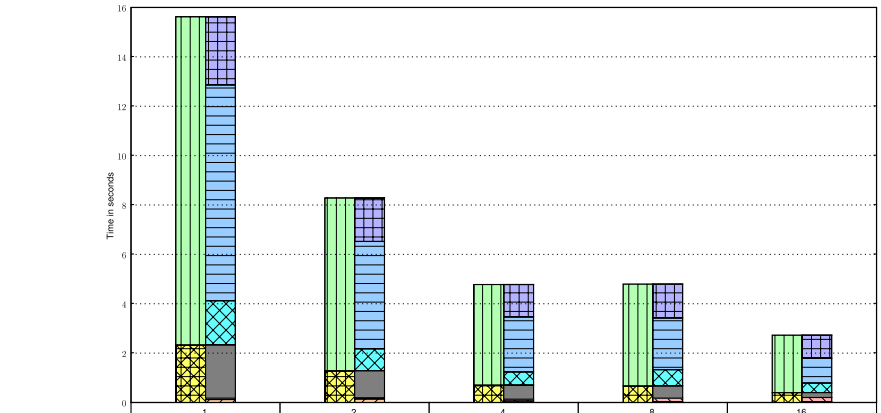

		1	2	4	8	16
Construction		2.3	1.3	0.7	0.7	0.4
Sort - Construction		0.1	0.1	0.1	0.0	0.0
BlockPartition		0.1	0.0	0.1	0.1	0.2
Balance		13.3	7.0	4.1	4.1	2.3
Commn - Balance		1.8	0.9	0.5	0.7	0.4
Internal Balance		8.7	4.3	2.2	2.1	1.0
Boundary Balance		2.8	1.8	1.3	1.4	0.9

FIG. 4.5. *Fixed size scalability for Gaussian distribution of* `1M` *octants. From left to right, the bars indicate the time taken for the different components of our algorithms for increasing processor counts. The bar for each processor is partitioned into* `2` *columns, which are further subdivided. The left column is subdivided into* `2` *sections and the right column is subdivided into* `6` *sections. The top and bottom sections of the left column represent the total time taken for* `(1)` *balance refinement* `(Algorithm 11)` *and* `(2)` *construction* `(Algorithm 1)`*, respectively. From top to bottom, the sections of the right column represent the time taken for* `(1)` *balancing across intra and inter processor boundaries* `(Algorithm 9)`*,* `(2)` *balancing the blocks* `(Algorithm 7)`*,* `(3)` *communication (including related pre-processing and post-processing) during balance refinement,* `(4)` *local processing during construction,* `(5) BlockPartition` *and* `(6) Sample Sort`*.*

[17] B. HARIHARAN, S. ALURU, AND B. SHANKER, *A scalable parallel fast multipole method for analysis of scattering from perfect electrically conducting surfaces*, in Supercomputing '02: Proceedings of the 2002 ACM/IEEE conference on Supercomputing, Los Alamitos, CA, USA, 2002, IEEE Computer Society Press, pp. 1–17.

[18] B. V. HERZEN AND A. H. BARR, *Accurate triangulations of deformed, intersecting surfaces*, in SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques, New York, NY, USA, 1987, ACM Press, pp. 103–110.

[19] E. KIM, J. BIELAK, O. GHATTAS, AND J. WANG, *Octree-based finite element method for large-scale earthquake ground motion modeling in heterogeneous basins*, AGU Fall Meeting Abstracts, (2002), pp. B1221+.

[20] H. MAGISTRALE, S. DAY, R. CLAYTON, AND R. GRAVES, *The scec southern california reference three-dimensional seismic velocity model version 2*, Bulletin of the Seismological Soceity of America, (2000).

[21] D. MEAGHER, *Geometric modeling using octree encoding*, Computer Graphics and Image Processing, 19 (1982), pp. 129–147.

[22] D. MOORE, *The cost of balancing generalized quadtrees.*, in Symposium on Solid Modeling and Applications, 1995, pp. 305–312.

[23] S. POPINET, *Gerris: a tree-based adaptive solver for the incompressible euler equations in complex geometries*, Journal of Computational Physics, 190 (20 September 2003), pp. 572–600(29).

[24] R. SCHNEIDERS, *An algorithm for the generation of hexahedral element meshes based on an octree technique*, 1997.

[25] R. SCHNEIDERS, R. SCHINDLER, AND F. WEILER, *Octree-based generation of hexahedral element meshes*, 1996.

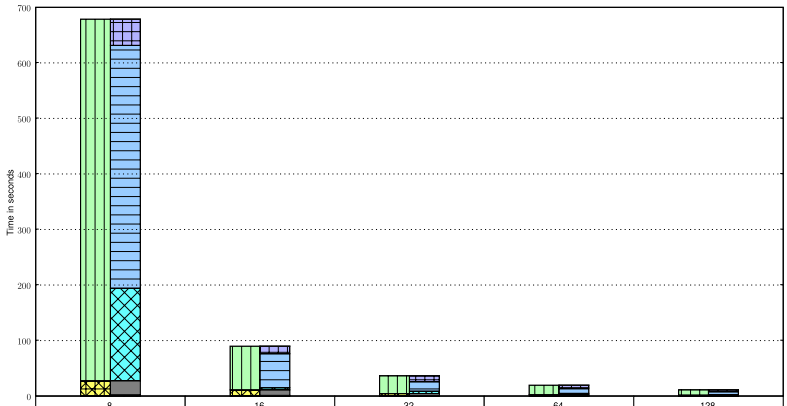

	8	16	32	64	128
Construction	27.7	11.3	4.6	2.5	1.7
Sort - Construction	1.7	0.8	0.4	0.2	0.2
BlockPartition	0.8	0.5	0.4	0.4	0.6
Balance	651.2	78.7	32.0	17.2	9.6
Commn - Balance	166.2	3.9	4.2	2.3	1.1
Internal Balance	437.5	61.2	17.1	9.5	5.0
Boundary Balance	47.4	13.6	10.7	5.3	3.4

FIG. 4.6. *Fixed size scalability for Gaussian distribution of* `32M` *octants. From left to right, the bars indicate the time taken for the different components of our algorithms for increasing processor counts. The bar for each processor is partitioned into* `2` *columns, which are further subdivided. The left column is subdivided into* `2` *sections and the right column is subdivided into* `6` *sections. The top and bottom sections of the left column represent the total time taken for* `(1)` *balance refinement* `(Algorithm 11)` *and* `(2)` *construction* `(Algorithm 1)`*, respectively. From top to bottom, the sections of the right column represent the time taken for* `(1)` *balancing across intra and inter processor boundaries* `(Algorithm 9)`*,* `(2)` *balancing the blocks* `(Algorithm 7)`*,* `(3)` *communication (including related pre-processing and post-processing) during balance refinement,* `(4)` *local processing during construction,* `(5) BlockPartition` *and* `(6) Sample Sort`*.*

[26] R. SCHUMAKER, P. BRAND, M. GILLILAND, AND W. SHARP, *Study for applying computer-generated images to visual simulation*, U.S. Air Force Human Resources Laboratory, Technical Report, 69 (1969).

[27] M. SHEPHARD AND M. GEORGES, *Automatic three-dimensional mesh generation by the finite octree technique*, International Journal for Numerical Methods in Engineering, 26 (1991), pp. 709–749.

[28] J. R. SHEWCHUK, *Tetrahedral mesh generation by delaunay refinement*, in Proceedings of the Fourteenth Annual Symposium on Computational Geometry, Minneapolis, Minnesota, USA, June 1998, Association for Computing Machinery, pp. 86–95.

[29] A. K. SOMANI AND A. M. SANSANO, *Minimizing overhead in parallel algorithms through overlapping communication/computation*, tech. report, Institute for Computer Applications in Science and Engineering (ICASE), 1997.

[30] K. STRASTERS AND J. GERBRANDS, *3-dimensional image segmentation using a split, merge and group-approach*, Pattern Recognition Letters, 12 (1991), pp. 307–325.

[31] S.-H. TENG, *Provably good partitioning and load balancing algorithms for parallel adaptive n-body simulation*, SIAM Journal on Scientific Computing, 19 (1998).

[32] H. TROPF AND H. HERZOG, *Multidimensional range search in dynamically balanced trees*, Angewandte Informatik, 2 (1981), pp. 71–77.

[33] T. TU AND D. R. O'HALLARON, *Balance refinement of massive linear octree datasets*, CMU Technical Report, CMU-CS-04 (2004).

[34] ———, *Extracting hexahedral mesh structures from balanced linear octrees*, in 13th International Meshing Roundtable, Williamsburg, VA, Sandia National Laboratories, September 19-22 2004, pp. 191–200.

[35] T. TU, D. R. O'HALLARON, AND O. GHATTAS, *Scalable parallel octree meshing for terascale applications*, in SC '05: Proceedings of the 2005 ACM/IEEE conference on Supercomputing, Washington, DC, USA, 2005, IEEE Computer Society, p. 4.

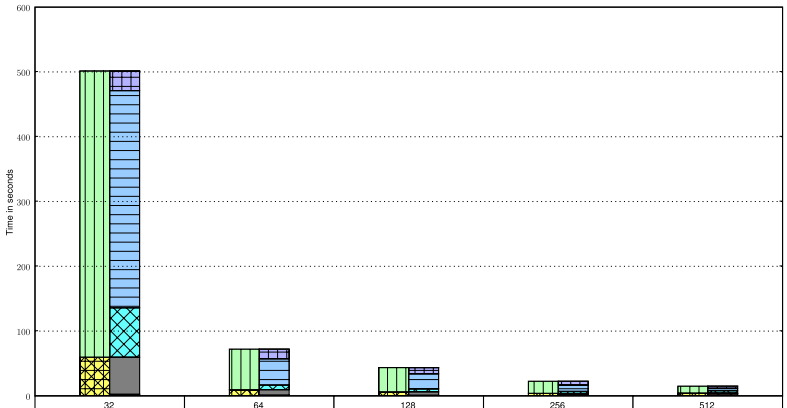

	32	64	128	256	512
Construction	59.7	9.5	6.2	3.9	4.3
Sort - Construction	2.0	0.8	0.4	0.6	1.4
BlockPartition	0.8	0.7	0.9	1.4	2.2
Balance	441.8	62.8	37.5	18.8	10.8
Commn - Balance	76.4	7.5	4.7	3.1	3.8
Internal Balance	335.1	40.3	23.1	9.7	3.7
Boundary Balance	30.3	15.0	9.7	5.9	3.3

FIG. 4.7. *Fixed size scalability for Gaussian distribution of* `128M` *octants. From left to right, the bars indicate the time taken for the different components of our algorithms for increasing processor counts. The bar for each processor is partitioned into* `2` *columns, which are further subdivided. The left column is subdivided into* `2` *sections and the right column is subdivided into* `6` *sections. The top and bottom sections of the left column represent the total time taken for* `(1)` *balance refinement* `(Algorithm 11)` *and* `(2)` *construction* `(Algorithm 1)`*, respectively. From top to bottom, the sections of the right column represent the time taken for* `(1)` *balancing across intra and inter processor boundaries* `(Algorithm 9)`*,* `(2)` *balancing the blocks* `(Algorithm 7)`*,* `(3)` *communication (including related pre-processing and post-processing) during balance refinement,* `(4)` *local processing during construction,* `(5) BlockPartition` *and* `(6) Sample Sort`*.*

[36] M. S. WARREN AND J. K. SALMON, *Astrophysical n-body simulations using hierarchical tree data structures*, in Supercomputing '92: Proceedings of the 1992 ACM/IEEE conference on Supercomputing, Los Alamitos, CA, USA, 1992, IEEE Computer Society Press, pp. 570–576.

[37] M. S. WARREN AND J. K. SALMON, *A parallel hashed oct-tree n-body algorithm*, in Proceedings of Supercomputing '93, 31 March 1993.

[38] L. YING, G. BIROS, AND D. ZORIN, *A kernel-independent adaptive fast multipole algorithm in two and three dimensions*, J. Comput. Phys., 196 (2004), pp. 591–626.

[39] L. YING, G. BIROS, D. ZORIN, AND H. LANGSTON, *A new parallel kernel-independent fast multipole method*, in SC '03: Proceedings of the 2003 ACM/IEEE conference on Supercomputing, Washington, DC, USA, 2003, IEEE Computer Society, p. 14.

Appendix A. Properties of Morton encoding.

PROPERTY 2. *Sorting all the leaves in the ascending order of their Morton ids is identical to a preorder traversal of the leaves of the octree. If one connects the centers of the leaves in this order, one can observe a Z-pattern in the Cartesian space. The space-filling Z-order curve has the property that spatially nearby octants tend to be clustered together. The octants in Figures 2.1(b) and 2.1(c) are all labeled according to this order. Depending on the order of interleaving the coordinates, different Z-order curves are obtained. The two possible Z-curves in 2-D are shown in the Figure A.1. Similarly, in 3-D six different types of Morton ordering are possible.*

PROPERTY 3. *Given three octants,* $a < b < c$ *and* $c \notin \{\mathcal{D}(b)\}$*:*

$$a < d < c, \quad \forall d \in \{\mathcal{D}(b)\}.$$

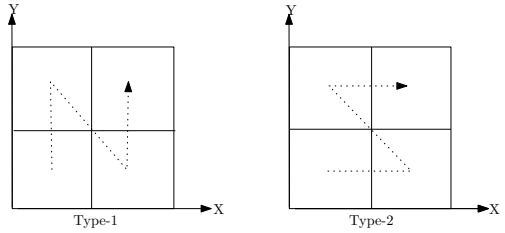

FIG. A.1. *Two types of z-ordering in quadtrees.*

PROPERTY 4. *The Morton id of any node is less than those of its descendants.*

PROPERTY 5. *Two distinct octants overlap if and only if one is an ancestor of the other.*

PROPERTY 6. *The Morton id of any node and of its first child*[24] *are consecutive. It follows from Property 4 that the first child is also the child with the least Morton id.*

PROPERTY 7. *The first descendant at level l, denoted by $\mathcal{FD}(N,l)$, of any node N is the descendant at level l with the least Morton id. This can be arrived at by following the first child at every level starting from N. $\mathcal{FD}(N,D_{max})$ is also the anchor of N and is also referred to as the* deepest first descendant*, denoted by $\mathcal{DFD}(N)$, of node N.*

PROPERTY 8. *The range $(N,\mathcal{DFD}(N)]$ only contains the first descendants of N at different levels and hence there can be no more than one leaf in this range in the entire linear octree.*

PROPERTY 9. *The* last descendant *at level l, denoted by $\mathcal{LD}(N,l)$, of any node N is the descendant at level l with the greatest Morton id. This can be arrived at by following the* last child[25] *at every level starting from N. $\mathcal{LD}(N,D_{max})$ is also referred to as the* deepest last descendant*, denoted by $\mathcal{DLD}(N)$, of node N.*

PROPERTY 10. *Every octant in the range $(N,\mathcal{DLD}(N)]$ is a descendant of N.*

Appendix B. Multicomponent Morton Representation. Every Morton id is a set of 4 entities: The three co-ordinates of the anchor of the octant and the level of the octant. We have implemented the node as a C++ class, which contains these 4 entities as its member data. To use this set as a locational code for octants, we define two primary binary logical operations on it: a) Comparing if 2 ids are equal and b) Comparing if one id is lesser than the other.

Two ids are equal if and only if all the 4 entities are respectively equal. If two ids have the same anchor then the one at a coarser level has a lesser Morton id. If the anchors are different, then we can use Algorithm 12 to determine the lesser id. The Z-ordering produced by this operator is identical to that produced by the scalar Morton ids described in section 2.1. The other logical operations can be readily derived from these two operations.

Appendix C. Analysis of the Block Partitioning Algorithm. Assume that the input to the partitioning algorithm is a sorted distributed list of N octants. Then, we can guarantee coarsening of the input if there are more than 8 octants[26] per

[24] the child that has the same anchor as the parent

[25] child with the greatest Morton id

[26] 2^d cells for a d-tree.

Algorithm 12. FINDING THE LESSER OF TWO MORTON IDS (SEQUENTIAL)

Input: Two Morton ids, A and B with different anchors.
Output: R, the lesser of the two Morton ids.

```
1. X_i ← (A_i ⊕ B_i),  i ∈ {x, y, z}
2. e ← arg max_i (log_2(X_i))
3. if A_e < B_e
       R ← A
4. else
       R ← B
5. end if
```

processor. The minimum number of octants on any processor, n_{min}, can be expressed in terms of N and the imbalance factor[27], c, as follows:

$$n_{min} = \frac{N}{1 + c(n_p - 1)}.$$

This implies that the coarsening algorithm will coarsen the octree if,

$$\begin{aligned} n_{min} &= \frac{N}{1 + c(n_p - 1)} > 2^d, \\ \implies N &> 2^d(1 + c(n_p - 1)). \end{aligned}$$

The total number of blocks created by our coarsening algorithm is $\mathcal{O}(p)$. Specifically, the total number of blocks produced by the coarsening algorithm, N_{blocks}, satisfies:

$$p \leq N_{blocks} < 2^d p.$$

If the input is sorted and if $c \approx 1$, then the communication cost for this partition is $\mathcal{O}(\frac{N}{n_p})$.

Appendix D. Special case during construction. We can not always guarantee the coarsest possible octree for an arbitrary distribution of N points and arbitrary values of N^p_{max}, especially when $N^p_{max} \approx \frac{N}{n_p}$. However, if every processor has at least 2 *well-separated* [28] points and if $N^p_{max} = 1$, then the algorithm will produce the coarsest possible octree under these constraints. However, this is not too restrictive because the input points can always be sampled in such a way that the algorithm produces the desired octree. Besides, the maximum depth of the octree can also be used to control the coarseness of the resulting octree. In all our experiments, we used $N^p_{max} = 1$ and we always got the same octree for different number of processor counts (Table 4.1).

[27] The imbalance factor is the ratio between the maximum and minimum number of octants on any processor.

[28] Convert the points into octants at D_{max} level. If there exists at least one coarse octant between these two octants, then the points are considered to be well-separated.