

PARALLEL MULTILEVEL TETRAHEDRAL GRID REFINEMENT

SVEN GROSS AND ARNOLD REUSKEN*

Abstract. In this paper we introduce and analyze a parallel version of a multilevel red/green local refinement algorithm for tetrahedral meshes. The serial version of this algorithm is presented in [6, 7]. We introduce a new data distribution format that is very suitable for parallelization of the multilevel refinement algorithm. This format is called an admissible hierarchical decomposition. We will prove that the application of the parallel refinement algorithm to an input admissible hierarchical decomposition yields an admissible hierarchical decomposition. We will show that due to this property one has good data locality (hence little communication) and at the same time only a small amount of copies (small storage overhead).

AMS subject classifications. 65N50, 65N55

Key words. tetrahedral grid refinement, parallelization, stable refinement, consistent triangulations

1. Introduction. In the field of numerical solution methods for partial differential equations adaptive discretization methods gain growing acceptance. In such adaptive methods both simplicial and hexahedral meshes are used. In this paper we only consider the former class and we restrict ourselves to the three dimensional situation, i.e., we treat tetrahedral meshes (which we also call triangulations). For well-known reasons one usually requires these meshes to fulfill a stability and consistency condition. The stability condition means that in a repeated refinement process, the interior angles of all tetrahedra that are generated must be uniformly bounded away from zero. A triangulation is called consistent if the intersection of any two tetrahedra from the triangulation is either empty, a common face, a common edge or a common vertex. Several refinement algorithms are known that satisfy both conditions for consistent input triangulations. These algorithms can be divided in two classes: *red/green refinement methods* (for example, [6, 8, 19, 26]) and *bisection methods* (for example, [1, 18, 20, 22, 23, 24]). For a discussion of the main differences between these classes of methods we refer to the literature (e.g., [6, 8, 24]). In this paper we only consider the three dimensional red/green refinement method as presented in [6, 8, 26]. This method is a 3D generalization of the well-known 2D red/green refinement technique due to Bank [2, 3]. Based on an idea from [4] a 3D *multilevel* red/green refinement algorithm is presented in [6, 7]. Below we will explain the term “multilevel”. Variants of this algorithm are used in several packages, for example in UG [5], KASKADE [17], DROPS [12], PML/MG [21]. The main topic of this paper is the introduction and analysis of a *parallel* version of this algorithm. Although parallel versions of this algorithm have already been implemented (e.g., in UG, PML/MG) we do not know of any paper in which a clear formulation of such a parallel version, which turns out to be rather complicated (cf. section 4.3), and an analysis of important properties are given.

To be able to summarize some interesting properties of the (parallel) refinement algorithm we first have to introduce a few notions. A sequence of triangulations $\mathcal{M} = (\mathcal{T}_0, \dots, \mathcal{T}_J)$ of some domain Ω is called a *multilevel triangulation* if: 1. each tetrahedron in \mathcal{T}_k ($k \geq 1$) is either in \mathcal{T}_{k-1} or it is obtained by a (red or green) refinement of its parent in \mathcal{T}_{k-1} ; 2. if a tetrahedron in \mathcal{T}_k is not refined when going

*Institut für Geometrie und Praktische Mathematik, RWTH Aachen, D-52056 Aachen, reusken@igpm.rwth-aachen.de, gross@igpm.rwth-aachen.de

to \mathcal{T}_{k+1} , then it remains unrefined in all \mathcal{T}_ℓ , $\ell > k$. Such a multilevel triangulation has nice structural properties. For example, the whole sequence \mathcal{M} can be uniquely reconstructed from \mathcal{T}_0 and \mathcal{T}_J . A related property is that all the information on \mathcal{M} can be represented in a natural way using a *hierarchical decomposition* $\mathcal{H} = (\mathcal{G}_0, \dots, \mathcal{G}_J)$, $\mathcal{G}_0 := \mathcal{T}_0$, $\mathcal{G}_k := \mathcal{T}_k \setminus \mathcal{T}_{k-1}$, $k \geq 1$. The hierarchical surplus on level k , \mathcal{G}_k , consists of all tetrahedra in \mathcal{M} that are on the same level k . In the implementation of the refinement algorithm this hierarchical decomposition plays a key role.

Now assume that based on some error indicator certain tetrahedra in the finest triangulation \mathcal{T}_J are marked for refinement. In many refinement algorithms one then modifies the finest triangulation \mathcal{T}_J resulting in a new one, \mathcal{T}_{J+1} . Using such a strategy (which we call a *one-level method*) the new sequence $(\mathcal{T}_0, \dots, \mathcal{T}_{J+1})$ is in general not a multilevel triangulation because the nestedness property 1 does not hold. We also note that when using such a method it is difficult to implement a reasonable coarsening strategy. In *multilevel* refinement algorithms as in [12, 6] the whole sequence \mathcal{M} is used and as output one obtains a sequence $\mathcal{M}' = (\mathcal{T}'_0, \dots, \mathcal{T}'_{J'})$, with $\mathcal{T}'_0 = \mathcal{T}_0$ and $J' \in \{J-1, J, J+1\}$. In general one has $\mathcal{T}'_k \neq \mathcal{T}_k$ for $k > 0$. We list a few important properties of this method:

- Both the input and output are *multilevel triangulations*.
- The method is *stable* and *consistent*.
- Local refinement and *coarsening* are treated in a similar way.
- The implementation uses only the hierarchical decomposition of \mathcal{M} . This allows *relatively simple data structures* without storage overhead.
- The costs are proportional to the number of tetrahedra in \mathcal{T}_J .

For a detailed discussion of these and other properties we refer to [6, 7]. In this paper we consider the parallelization of such a multilevel refinement algorithm. To make the paper more self-contained and in view of a better understanding of the parallel algorithm we decided to give a rather detailed description of the serial method (section 3).

The parallel method that we propose produces, for a given input multilevel triangulation, the same output multilevel triangulation as the serial method. In this sense the “computational part” of the algorithm is not changed. The main problem related to the parallelization is how to store the distributed data on the different processors¹ of the distributed memory machine in a useful way: One has to find a good compromise between storage overhead on the different processors and the costs of communication between the processors. In the case of the multilevel refinement method, a crucial point is that for the refinement (or unrefinement) of a tetrahedron T one needs information about the children of T . If the generic case is such that T and its children are not stored on the same processor one needs a large amount of small messages and the high communication costs will result in low parallel efficiency. Hence, in view of a reduction of communication costs it is advantageous to have T (or a copy of T) and its children on the same processor. If this condition has to be fulfilled, one must allow copies (i.e., overlap) of tetrahedra on different processors since otherwise a domino effect can occur: all ancestors and descendants of a tetrahedron have to be on the same processor. This causes very coarse data granularity and hence low parallel efficiency.

We will introduce and analyze a new data distribution format that is very suitable for parallelization of the multilevel refinement algorithm. This format is called an *admissible hierarchical decomposition*. We will prove that the application of the

¹here ‘processor’ means CPU with associated memory

parallel refinement algorithm to an input admissible hierarchical decomposition yields an admissible hierarchical decomposition. We will show that due to this property one has good data locality (hence little communication) and at the same time only a small amount of copies (small storage overhead).

This paper is organized as follows. In section 2 we collect some definitions. In section 3 a rather detailed description of the serial multilevel refinement algorithm is given. In section 4 we introduce and analyze the admissible hierarchical decomposition. Furthermore, the parallel multilevel refinement algorithm is presented. In section 4.1 we summarize the main new results of this paper. In section 5 a theoretical analysis is given which proves that the application of the parallel refinement algorithm to an input admissible hierarchical decomposition yields an admissible hierarchical decomposition. Finally, in section 6 we give results of a few numerical experiments.

2. Definitions and notation. In this section we collect notations and definitions that will be used in the remainder of this paper. Let Ω be a polyhedral domain in \mathbb{R}^3 .

DEFINITION 1 (Triangulation). A finite collection \mathcal{T} of tetrahedra $T \subset \bar{\Omega}$ is called a *triangulation* of Ω (or $\bar{\Omega}$) if the following holds:

1. $\text{vol}(T) > 0$ for all $T \in \mathcal{T}$,
2. $\bigcup_{T \in \mathcal{T}} T = \bar{\Omega}$,
3. $\text{int}(S) \cap \text{int}(T) = \emptyset$ for all $S, T \in \mathcal{T}$ with $S \neq T$.

DEFINITION 2 (Consistency). A triangulation \mathcal{T} is called *consistent* if the intersection of any two tetrahedra in \mathcal{T} is either empty, a common face, a common edge or a common vertex.

DEFINITION 3 (Stability). A sequence of triangulations $\mathcal{T}_0, \mathcal{T}_1, \mathcal{T}_2, \dots$ is called *stable* if all angles of all tetrahedra in this sequence are uniformly bounded away from zero.

DEFINITION 4 (Refinement). For a given tetrahedron T a triangulation $\mathcal{K}(T)$ of T is called a *refinement* of T if $|\mathcal{K}(T)| \geq 2$ and any vertex of any tetrahedron $T' \in \mathcal{K}(T)$ is either a vertex or an edge midpoint of T . In this case T' is called a child of T and T is called the parent of T' . A refinement $\mathcal{K}(T)$ of T is called *regular* if $|\mathcal{K}(T)| = 8$, otherwise it is called *irregular*. A triangulation \mathcal{T}_{k+1} is called *refinement* of a triangulation $\mathcal{T}_k \neq \mathcal{T}_{k+1}$ if for every $T \in \mathcal{T}_k$ either $T \in \mathcal{T}_{k+1}$ or $\mathcal{K}(T) \subset \mathcal{T}_{k+1}$ for some refinement $\mathcal{K}(T)$ of T .

DEFINITION 5 (Multilevel triangulation). A sequence of consistent triangulations $\mathcal{M} = (\mathcal{T}_0, \dots, \mathcal{T}_J)$ is called a *multilevel triangulation* of Ω if the following holds:

1. For $0 \leq k < J$: \mathcal{T}_{k+1} is a refinement of \mathcal{T}_k .
2. For $0 \leq k < J$: $T \in \mathcal{T}_k \cap \mathcal{T}_{k+1} \Rightarrow T \in \mathcal{T}_J$.

The tetrahedra $T \in \mathcal{T}_J$ are called the leaves of \mathcal{M} . Note that T is a leaf iff T has no children in \mathcal{M} . A tetrahedron $T \in \mathcal{M}$ is called *regular* if $T \in \mathcal{T}_0$ or T resulted from a regular refinement of its parent. Otherwise T is called *irregular*. A multilevel triangulation \mathcal{M} is called regular if all irregular $T \in \mathcal{M}$ are leaves (i.e., have no children in \mathcal{M}).

REMARK 1. Let \mathcal{M} be a multilevel triangulation and V_k ($0 \leq k \leq J$) be the corresponding finite element spaces of continuous functions $p \in C(\bar{\Omega})$ such that $p|_T \in \mathcal{P}_q$ for all $T \in \mathcal{T}_k$ ($q \geq 1$). The refinement property 1 in definition 5 implies nestedness of these finite element spaces: $V_k \subset V_{k+1}$.

DEFINITION 6 (Hierarchical decomposition of \mathcal{M}). Let $\mathcal{M} = (\mathcal{T}_0, \dots, \mathcal{T}_J)$ be a multilevel triangulation of Ω . For every tetrahedron $T \in \mathcal{M}$ a unique level number

$\ell(T)$ is defined by $\ell(T) := \min\{k \mid T \in \mathcal{T}_k\}$. The set $\mathcal{G}_k \subset \mathcal{T}_k$

$$\mathcal{G}_k := \{T \in \mathcal{T}_k \mid \ell(T) = k\}$$

is called the *hierarchical surplus* on level k ($k = 0, 1, \dots, J$). Note that $\mathcal{G}_0 = \mathcal{T}_0$, $\mathcal{G}_k = \mathcal{T}_k \setminus \mathcal{T}_{k-1}$ for $k = 1, \dots, J$. The sequence $\mathcal{H} = (\mathcal{G}_0, \dots, \mathcal{G}_J)$ is called the *hierarchical decomposition* of \mathcal{M} . Note that the multilevel triangulation \mathcal{M} can be reconstructed from its hierarchical decomposition.

REMARK 2. The hierarchical decomposition induces simple data structures in a canonical way. The tetrahedra of each hierarchical surplus \mathcal{G}_k are stored in a separate list. Thus every tetrahedron $T \in \mathcal{M}$ is stored exactly once since T has a unique level number $\ell(T)$. By introducing unique level numbers also for vertices, edges and faces, these subsimplices can be stored in the same manner: For a subsimplex S the level number $\ell(S)$ is defined as the level of its first appearance. Additionally, the objects are linked to certain corresponding objects by pointers (e.g., a tetrahedron is linked to its vertices, edges, faces, children and parent).

3. A serial multilevel refinement algorithm. In this section we describe a refinement algorithm which is, apart from some minor modifications, the algorithm presented in [6, 7]. This method is based on similar ideas as the refinement algorithms in [3, 4, 5]. At the end of this section (remark 4) we discuss the differences between the method presented here and the one from [6, 7]. We give a rather detailed description of the algorithm to facilitate the presentation of the *parallel* version of this method that is given in section 4.3. The refinement strategy is based on a set of regular and irregular refinement rules (also called red and green rules, due to [2, 3]). The regular and irregular rules are local in the sense that they are applied to a single tetrahedron. These rules are applied in a (global) refinement algorithm (section 3.3) that describes how the local rules can be combined to ensure consistency and stability.

3.1. The regular refinement rule. Let T be a given tetrahedron. For the construction of a regular refinement of T it is natural to connect midpoints of the edges of T by subdividing each of the faces into four congruent triangles. This yields four subtetrahedra at the corners of T (all congruent with T) and an octahedron in the middle. For the subdivision of this octahedron into four subtetrahedra with equal volume there are three possibilities corresponding to the three diagonals that connect opposite vertices of the octahedron. One has to be careful in choosing an appropriate subdivision since in [26] it has been shown that the wrong choice may lead to a sequence of triangulations that is not stable. A stable tetrahedral regular refinement strategy, based on an idea from [11], is presented in [6, 8]. We recall this method.

Let $T = [x^{(1)}, x^{(2)}, x^{(3)}, x^{(4)}]$ be a tetrahedron with *ordered* vertices $x^{(1)}$, $x^{(2)}$, $x^{(3)}$, $x^{(4)}$ and

$$x^{(ij)} := \frac{1}{2}(x^{(i)} + x^{(j)}), \quad 1 \leq i < j \leq 4,$$

the midpoint of the edge between $x^{(i)}$ and $x^{(j)}$. The regular refinement $\mathcal{K}(T) := \{T_1, \dots, T_8\}$ of T is constructed by the (red) rule

$$\begin{aligned} T_1 &:= [x^{(1)}, x^{(12)}, x^{(13)}, x^{(14)}], & T_5 &:= [x^{(12)}, x^{(13)}, x^{(14)}, x^{(24)}], \\ T_2 &:= [x^{(12)}, x^{(2)}, x^{(23)}, x^{(24)}], & T_6 &:= [x^{(12)}, x^{(13)}, x^{(23)}, x^{(24)}], \\ T_3 &:= [x^{(13)}, x^{(23)}, x^{(3)}, x^{(34)}], & T_7 &:= [x^{(13)}, x^{(14)}, x^{(24)}, x^{(34)}], \\ T_4 &:= [x^{(14)}, x^{(24)}, x^{(34)}, x^{(4)}], & T_8 &:= [x^{(13)}, x^{(23)}, x^{(24)}, x^{(34)}]. \end{aligned} \quad (3.1)$$

In [8] it is shown that for any T the repeated application of this rule produces a sequence of consistent triangulations of T which is stable. For a given T all tetrahedra that are generated in such a recursive refinement process form at most three congruence classes.

3.2. Irregular refinement rules. Let \mathcal{T} be a given consistent triangulation. We select a subset \mathcal{S} of tetrahedra from \mathcal{T} and assume that the regular refinement rule is applied to each of the tetrahedra from \mathcal{S} . In general the resulting triangulation \mathcal{T}' will not be consistent. The irregular (or green) rules are used to make this new triangulation consistent. For this we introduce the notions of an edge counter and edge refinement pattern. The *edge counter* $C(E)$, that depends on \mathcal{T} and \mathcal{S} , assigns an integer value to each of the edges E of \mathcal{T} as follows: $C(E) = m$ if the edge E is an edge of precisely m elements from \mathcal{S} . Hence, the edge E has been refined when going from \mathcal{T} to \mathcal{T}' iff $C(E) > 0$ holds. Related to this, for $T \in \mathcal{T}$ we define the *edge refinement pattern* $R(T)$ as follows. Let E_1, \dots, E_6 be the ordered edges of T . We define the 6-tuple

$$R = (r_1, \dots, r_6) \in \{0, 1\}^6$$

by: $r_i = 0$ if $C(E_i) = 0$ and $r_i = 1$ if $C(E_i) > 0$. This edge refinement pattern describes a local property in the sense that it is related to a single tetrahedron and its value directly follows from the values of the edge counter. For $T \in \mathcal{S}$ we have $R(T) = (1, \dots, 1)$. For $T \in \mathcal{T} \setminus \mathcal{S}$ the case $R(T) = (0, \dots, 0)$ corresponds to the situation that the tetrahedron T does not contain any vertices from \mathcal{T}' at the midpoints of its edges. For each of the $2^6 - 1$ possible patterns $R \neq (0, \dots, 0)$ there exists a corresponding refinement $\mathcal{K}(T)$ of T (in the fashion of (3.1)) for which the vertices of the children coincide with vertices of T or with the vertices at the midpoints on the edges E_i with $r_i = 1$. This refinement, however, is not always unique. This is illustrated in figure 3.1.

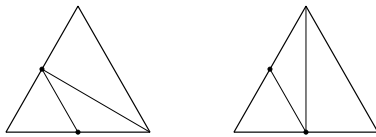


FIG. 3.1. *Non-unique face refinement*

To obtain a consistent triangulation in which the subdivision of adjacent faces matches special care is needed. One way to ensure consistency is by introducing a so-called consistent vertex numbering:

DEFINITION 7 (Consistent vertex numbering). Let T_1 and T_2 be two adjacent tetrahedra with a common face $F = T_1 \cap T_2$ and local vertex ordering

$$T_l = [x_l^{(1)}, x_l^{(2)}, x_l^{(3)}, x_l^{(4)}], \quad l = 1, 2.$$

Let the vertex set of F be given by

$$\{x_1^{(i_1)}, x_1^{(i_2)}, x_1^{(i_3)}\} = \{x_2^{(j_1)}, x_2^{(j_2)}, x_2^{(j_3)}\}.$$

Without loss of generality we assume $i_1 < i_2 < i_3$ and $j_1 < j_2 < j_3$. The pair (T_1, T_2) has a *consistent vertex numbering* if

$$x_1^{(i_k)} = x_2^{(j_k)}, \quad k = 1, 2, 3$$

holds, i.e., if the ordering of the vertices of F induced by the vertex numbering of T_1 coincides with the one induced by the vertex numbering of T_2 . A consistent triangulation \mathcal{T} has a *consistent vertex numbering* if every two neighbouring tetrahedra have this property.

We note that a consistent vertex numbering can be constructed in a rather simple way. Consider an (initial) triangulation $\tilde{\mathcal{T}}$ with an arbitrary numbering of its vertices. This global numbering induces a canonical local vertex ordering which is a consistent vertex numbering of $\tilde{\mathcal{T}}$. Furthermore, each refinement rule can be defined such that the consistent vertex numbering property of the parent is inherited by its children by prescribing suitable local vertex orderings of the children. (3.1) is an example of such a rule. Using such a strategy a consistent triangulation $\tilde{\mathcal{T}}'$ that is obtained by refinement of $\tilde{\mathcal{T}}$ according to these rules also has a consistent vertex numbering.

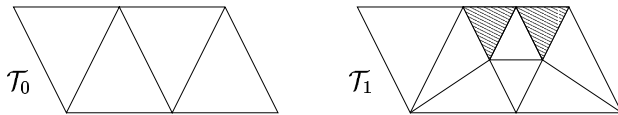
ASSUMPTION 1. *In the remainder of this paper we always assume that the initial triangulation \mathcal{T}_0 is consistent and has a consistent vertex numbering.*

Assume that the given triangulation \mathcal{T} has a consistent vertex numbering. For a face with a pattern as in figure 3.1 one can then define a unique face refinement by connecting the vertex with the smallest number with the midpoint of the opposite edge. *For each edge refinement pattern $R \in \{0, 1\}^6$ we then have a unique rule.* We emphasize that if the edge refinement pattern is known *for the application of the regular or irregular rules to a given tetrahedron no information from neighbouring tetrahedra is needed.* Clearly, for parallelization this is a very nice property.

Up to now we discussed how the *consistency* of a triangulation can be achieved by the choice of suitable irregular refinement rules based on the consistent vertex numbering property. We will now explain how the regular and irregular rules can be combined in a repeated refinement procedure to obtain a *stable* sequence of consistent triangulations. The crucial point is to *allow only the refinement of regular tetrahedra*, i.e. children of irregularly refined tetrahedra, also called *green children*, are never refined. If such a green child T is marked for refinement, instead of refining T the irregular refinement of the parent will be replaced by a regular one. As the application of the regular rule (3.1) creates tetrahedra of at most 3 congruence classes (cf. [11, 8]), the tetrahedra created by a refinement procedure according to this strategy belong to an a-priori bounded number of congruence classes. Hence the obtained sequence of triangulations is stable.

3.3. The multilevel refinement algorithm. In this section we describe how the regular and irregular refinement rules are used in a multilevel refinement algorithm (as in [12, 6, 7]).

We first illustrate the difference to well-known one-level refinement methods as in, e.g., [2, 3]. For ease of presentation, we use triangles instead of tetrahedra in our examples. Consider the following multilevel triangulation $(\mathcal{T}_0, \mathcal{T}_1)$:



In \mathcal{T}_1 two triangles are marked (by shading) for further refinement. The one-level method from [3] uses the finest triangulation \mathcal{T}_1 as input and applies a regular refinement rule followed by irregular refinement rules:

As output one obtains a consistent triangulation \mathcal{T}_2 which is *not* a refinement of \mathcal{T}_1 (in the sense of definition 4). Related to this we note that the finite element

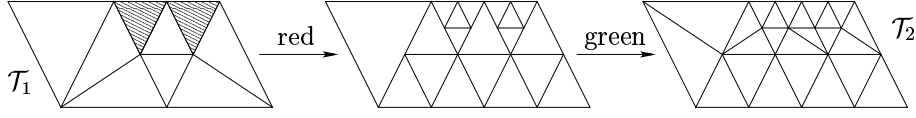


FIG. 3.2. *One-level red/green refinement*

spaces corresponding to $\mathcal{T}_1, \mathcal{T}_2$ are not nested (cf. remark 1) and that it is not obvious how to construct a hierarchical decomposition of the sequence $\mathcal{T}_0, \mathcal{T}_1, \mathcal{T}_2$. These are disadvantages if one uses such grid refinement techniques in combination with multigrid solvers for the numerical solution of partial differential equations. An advantage of the one-level method compared to the multilevel strategy discussed below is its simplicity. In multilevel refinement algorithms both the input and the output are multilevel triangulations (definition 5). In the implementation only the hierarchical decompositions of these multilevel triangulations are used. The algorithmic structure of the method presented in [6, 7] is shown in figure 3.3.

Algorithm <i>SerRefinement</i> ($\mathcal{G}_0, \dots, \mathcal{G}_J$)	
for $k = J, \dots, 0$ do	// phase I
<i>DetermineMarks</i> (\mathcal{G}_k);	(1)
<i>MarksForClosure</i> (\mathcal{G}_k);	(2)
for $k = 0, \dots, J$ do if $\mathcal{G}_k \neq \emptyset$ then	// phase II
if $k > 0$ then <i>MarksForClosure</i> (\mathcal{G}_k);	(3)
if $k < J$ then <i>Unrefine</i> (\mathcal{G}_k);	(4)
<i>Refine</i> (\mathcal{G}_k);	(5)
if $\mathcal{G}_J = \emptyset$ then $J := J - 1$;	(6)
else if $\mathcal{G}_{J+1} \neq \emptyset$ then $J := J + 1$;	(7)

FIG. 3.3. *Serial multilevel refinement algorithm.*

We now introduce the notions of **status** and **mark** of a tetrahedron that will be used in the subroutines of this algorithm. Let $\mathcal{M} = (\mathcal{T}_0, \dots, \mathcal{T}_J)$ be a multilevel triangulation that has been constructed by applying the regular and irregular refinement rules and let $\mathcal{H} = (\mathcal{G}_0, \dots, \mathcal{G}_J)$ be the corresponding hierarchical decomposition. Every tetrahedron $T \in \mathcal{H}$ is either a leaf of \mathcal{M} (i.e., $T \in \mathcal{T}_J$) or it has been refined. The label **status** is used to describe this property of T :

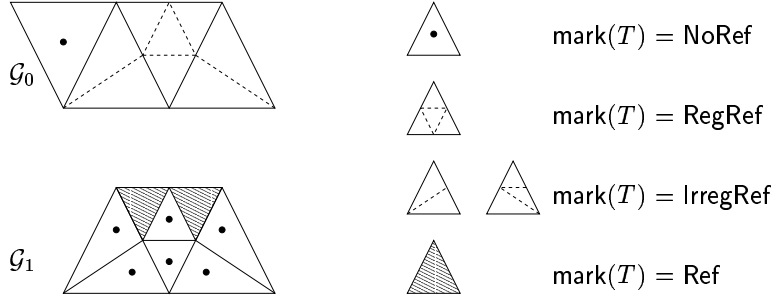
$$\text{For } T \in \mathcal{H}: \quad \text{status}(T) = \begin{cases} \text{NoRef} & \text{if } T \text{ is a leaf of } \mathcal{M} \\ \text{RegRef} & \text{if } T \text{ is regularly refined in } \mathcal{M} \\ \text{IrregRef} & \text{if } T \text{ is irregularly refined in } \mathcal{M} \end{cases}$$

We note that the label **IrregRef** contains the number of the irregular refinement rule (one out of 63) that has been used to refine T , i.e., **status**(T) coincides with the edge refinement pattern of T .

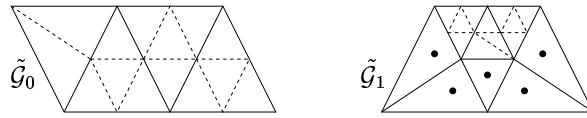
In adaptive refinement an error estimator (or indicator) is used to mark certain elements of \mathcal{T}_J for further refinement or for deletion. For this the label **mark** is used:

$$\text{For } T \in \mathcal{H}: \quad \text{mark}(T) = \begin{cases} \text{Ref} & \text{if } T \in \mathcal{T}_J \text{ is marked for refinement} \\ \text{Del} & \text{if } T \in \mathcal{T}_J \text{ is marked for deletion} \\ \text{status}(T) & \text{otherwise} \end{cases}$$

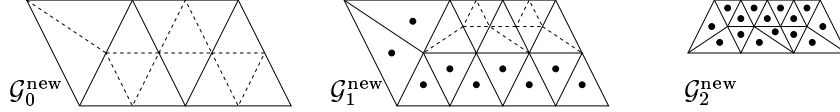
The main idea underlying the algorithm *SerRefinement* is illustrated using the multilevel triangulation $(\mathcal{T}_0, \mathcal{T}_1)$ from above. The hierarchical decomposition and the corresponding marks are as follows:



Note that for the two shaded triangles in \mathcal{G}_1 we have $\text{status}(T) \neq \text{mark}(T)$. For all other triangles $\text{status}(T) = \text{mark}(T)$ holds. In phase I of the algorithm (top-down: (1),(2)) only marks are changed:



Once phase I has been completed the marks have been changed such that $\text{mark}(T) \in \{\text{NoRef}, \text{RegRef}, \text{IrregRef}\}$ holds for all $T \in \mathcal{H}$. We emphasize that all green children in $\tilde{\mathcal{G}}_1$ have $\text{mark}(T) = \text{NoRef}$, as they are not refined because of stability reasons. Instead the corresponding irregularly refined parents in $\tilde{\mathcal{G}}_0$ are labeled by $\text{mark}(T) = \text{RegRef}$. In the second phase (bottom-up: (3)-(5)) the actual refinement (unrefinement is not needed in our example) is constructed:



In the output hierarchical decomposition $\mathcal{H} = (\mathcal{G}_0^{\text{new}}, \mathcal{G}_1^{\text{new}}, \mathcal{G}_2^{\text{new}})$ we have $\text{mark}(T) = \text{status}(T)$ for all $T \in \mathcal{H}$. The output multilevel triangulation $\mathcal{M} = (\mathcal{T}_0^{\text{new}}, \mathcal{T}_1^{\text{new}}, \mathcal{T}_2^{\text{new}})$ is *regular* (cf. definition 5) and is given by

$$\mathcal{T}_0^{\text{new}} = \mathcal{G}_0^{\text{new}}, \quad \mathcal{T}_1^{\text{new}} = \mathcal{G}_1^{\text{new}}, \quad \mathcal{T}_2^{\text{new}} = \mathcal{G}_2^{\text{new}} \cup \{T \in \mathcal{G}_1^{\text{new}} \mid \text{mark}(T) = \text{NoRef}\}.$$

Note that $\mathcal{T}_0^{\text{new}} = \mathcal{T}_0$, $\mathcal{T}_1^{\text{new}} \neq \mathcal{T}_1$ (!) and that the new finest triangulation $\mathcal{T}_2^{\text{new}}$ is the same as the triangulation \mathcal{T}_2 in figure 3.2 resulting from the one-level algorithm.

Below we describe the subroutines used in algorithm *SerRefinement*. A detailed discussion of these subroutines is given in [6, 7].

DetermineMarks. In this subroutine only marks are changed. The new values of the marks that are changed are of the type *RegRef* or *NoRef*. The value $\text{mark}(T) = \text{RegRef}$ is assigned if

- T is a regular leaf with $\text{mark}(T) = \text{Ref}$
- T has been irregularly refined ($\text{status}(T) = \text{IrregRef}$) and at least one of its children is marked for refinement (either by $\text{mark}(\text{child}) = \text{Ref}$ or by a certain edge refinement pattern)

The value $\text{mark}(T) = \text{NoRef}$ is assigned if

- $\text{status}(T) = \text{RegRef}$ and all children of T are marked for deletion ($\text{mark}(T') = \text{Del}$ for all $T' \in \mathcal{K}(T)$)
- T has been irregularly refined ($\text{status}(T) = \text{lrregRef}$) and none of its children is marked for refinement

The subroutine is described in detail in figure 3.4. Another important task of *De-*

```

Function DetermineMarks( $\mathcal{G}_k$ )
  for  $T \in \mathcal{G}_k$  do
    if  $\text{status}(T) = \text{NoRef}$  then
      if  $T$  is regular and  $\text{mark}(T) = \text{Ref}$  then
         $\text{mark}(T) := \text{RegRef};$  (1)
        increase edge counters; (2)
      if  $k = 0$  and  $\text{mark}(T) = \text{Del}$  then  $\text{mark}(T) := \text{NoRef};$  (3)
    else if  $\text{status}(T) = \text{RegRef}$  then
      if  $\forall T' \in \mathcal{K}(T) : \text{mark}(T') = \text{Del}$  then
         $\text{mark}(T) := \text{NoRef};$  (4)
        decrease edge counters; (5)
      for  $T' \in \mathcal{K}(T)$  do
        if  $\text{mark}(T') = \text{Del}$  then  $\text{mark}(T') := \text{NoRef};$  (6)
    else //  $\text{status}(T) = \text{lrregRef}$ 
      if  $\exists T' \in \mathcal{K}(T) : \text{mark}(T') = \text{Ref}$ 
      or an edge of a child, which is not an edge of  $T$ ,
        is marked for refinement then
           $\text{mark}(T) := \text{RegRef};$  (7)
          increase edge counters; (8)
        else  $\text{mark}(T) := \text{NoRef};$  (9)
        for  $T' \in \mathcal{K}(T)$  do  $\text{mark}(T') := \text{NoRef};$  (10)

```

FIG. 3.4. Subroutine *DetermineMarks*.

termineMarks is the book-keeping of the edge counters. If a tetrahedron is to be regularly refined the counters of its edges are increased. Similarly, if all children of a regularly refined tetrahedron are removed the edge counters of its edges are decreased.

MarksForClosure. In this subroutine an appropriate irregular refinement rule is determined for an element $T \in \mathcal{G}_k$ to avoid hanging nodes. The tetrahedron T must be regular (irregular elements are never refined) and it should not be marked for regular refinement. The subroutine is described in figure 3.5.

```

Function MarksForClosure( $\mathcal{G}_k$ )
  for  $T \in \mathcal{G}_k$  do
    if  $T$  is regular and  $\text{mark}(T) \neq \text{RegRef}$  then
      Determine the edge refinement pattern  $R$ 
      of  $T$  (using the edge counters); (1)
    if  $R = (0, \dots, 0)$  and  $\text{mark}(T) = \text{Del}$  then
      do nothing;
      // NoRef-mark is set in DetermineMarks( $\mathcal{G}_{k-1}$ )
    else
       $\text{mark}(T) := R;$  (2)

```

FIG. 3.5. Subroutine *MarksForClosure*.

Unrefine. In the call of this subroutine on level k all the tetrahedra, vertices,

edges and faces on level $k + 1$ that are not needed anymore (due to changed marks) are removed. More details are given in figure 3.6. We note that for an efficient implementation one could check also for the case $\text{status}(T) \neq \text{NoRef}$ and $\text{mark}(T) \neq \text{status}(T)$ whether certain already known objects (tetrahedra, edges, etc.) on level $k + 1$ can be reused in the refinement.

Function $\text{Unrefine}(\mathcal{G}_k)$
 Label all tetrahedra, vertices, edges and faces on level $k + 1$ for deletion; (1)
for $T \in \mathcal{G}_k$ **do** **if** $(\text{status}(T) \neq \text{NoRef} \text{ and } \text{mark}(T) = \text{status}(T))$ **then**
 Remove all deletion labels of the
 children of T and of their vertices, edges, faces; (2)
 Remove all tetrahedra, vertices, edges and faces on level $k + 1$
 that are labeled for deletion; (3)

FIG. 3.6. Subroutine *Unrefine*.

Refine. In the subroutine *Refine*, if $\text{mark}(T) \neq \text{status}(T)$, an appropriate refinement of T , based on $\text{mark}(T)$, is made and new objects (tetrahedra, vertices, edges and faces) are created on level $k + 1$. For $k = J$, a new hierarchical surplus \mathcal{G}_{J+1} is constructed. The refined tetrahedron T is labeled with its new **status**. After application of the subroutine *Refine* on level k all tetrahedra on this level have a **status** which corresponds to their **mark**. Leaves are given the **status** **NoRef**. Further details are presented in figure 3.7.

Function $\text{Refine}(\mathcal{G}_k)$
if $k = J$ **then** $\mathcal{G}_{k+1} := \emptyset$; (1)
for $T \in \mathcal{G}_k$ **do**
if $\text{mark}(T) \neq \text{status}(T)$ **then**
 Refine T according to $\text{mark}(T)$; (2)
 $\text{status}(T) := \text{mark}(T)$; (3)
 for $T' \in \mathcal{K}(T)$ **do**
 Find existing vertices, edges and faces of T' ; (4)
 create missing vertices, edges and faces of T' ; (5)
 $\text{status}(T') := \text{NoRef}$; (6)

FIG. 3.7. Subroutine *Refine*.

REMARK 3. Applying the subroutines to \mathcal{G}_k , for any $T \in \mathcal{G}_k$ one only needs information that is directly connected with T (e.g., $\text{mark}(T)$), information about a possible parent (to decide, for example, whether T is regular or not) and if T is not a leaf, information about the children of T .

REMARK 4. As already noted above, algorithm *SerRefinement* is very similar to the method of BEY ([6, 7]). There are, however, some differences. We use a complete set of 64 refinement rules, whereas BEY uses a much smaller number of green rules, which may cause a domino effect. In view of data locality we wanted to avoid this effect. Concerning the implementation there are two major differences with the approach of BEY. Firstly, in our implementation faces are explicitly represented in the data structure, which simplifies the access from one tetrahedron to its four neighbour tetrahedra. Secondly, each vertex, edge and face is stored only once (and not on every level where corresponding tetrahedra occur).

4. Parallelization of the multilevel refinement method.

4.1. Introduction. In many applications (for example, CFD) the computational complexity is very high and one wants to use parallel machines. Here we only consider distributed memory machines which use message passing (MPI). In this section we reconsider the local refinement algorithm *SerRefinement* and present a version, which is called *ParRefinement*, that is suitable for such a machine. For a given input-multilevel triangulation the parallel method produces the same output-multilevel triangulation as the serial method. In this sense the “computational part” of the algorithm is not changed. In the parallel case load has to be distributed uniformly among the processors, so in practice an adaptive parallel refinement algorithm is combined with dynamic load balancing and data migration between the processors. One of the main decisions related to the parallelization of the algorithm *SerRefinement* is to find an appropriate distributed storage of data with regard to data locality. Therefore one has to find a good compromise between storage overhead on the different processors and the costs of communication between the processors. In the case of the multilevel refinement method, a crucial point is that for a tetrahedron T one needs information about the children of T . If T and its children are not necessarily stored on the same processor, then in general one needs a large amount of small messages and the high communication costs will result in low parallel efficiency. Hence, in view of a reduction of communication costs it is advantageous to *have T (or a copy of T) and its children on the same processor*. If this condition has to be fulfilled, *one must allow (ghost) copies* (i.e., overlap) of tetrahedra on different processors since otherwise a domino effect can occur: all ancestors and descendants of a tetrahedron must be on the same processor. This would cause very coarse data granularity, poor load balancing and hence low parallel efficiency.

One main topic of this paper is the introduction and analysis of a new data distribution format that is very suitable for parallelization of the multilevel refinement algorithm. This data distribution format is such that the following holds:

1. Let $T \in \mathcal{G}_k$ be an element from the hierarchical surplus on level k . Then T is stored on one processor, say p . In certain cases (explained below) a ghost copy of T is stored on *one* other processor, say q .
2. The children of T (if they exist) are all stored either on processor p or, if T has a ghost copy, on processor q . For $T \in \mathcal{G}_k$, $k > 0$, the parent of T or a copy of it is stored on processor p .

Due to property 2 the refinement algorithm can be applied on level k (i.e., to the hierarchical surplus \mathcal{G}_k) with low communication costs. Communication in the form of relatively few large messages will be needed to transfer information from elements on one processor to copies (if they exist) on other processors and to exchange information on edge counters between the processors. The first property shows that in a certain sense the overlap is small.

The new data distribution format will be made mathematically precise by a formal specification of a so-called *admissible hierarchical decomposition*. This is presented in section 4.2. In section 4.3 we introduce the *parallel* version of the multilevel refinement algorithm (*ParRefinement*).

The main results concerning the admissible hierarchical decomposition and the parallel multilevel refinement method can be summarized as follows:

- a. An admissible hierarchical decomposition has the desirable properties 1 (small storage overhead) and 2 (data locality) from above. This is proved in section 4.2.

- b. The application of the algorithm *ParRefinement* to an admissible hierarchical decomposition results in an admissible hierarchical decomposition. This is proved in section 5.
- c. Given an admissible hierarchical decomposition one can apply a suitable load balancing and data migration algorithm such that after data migration one still has an admissible hierarchical decomposition. We comment on this in remark 7 below.

4.2. Admissible hierarchical decomposition. Let the sequence $\mathcal{M} = (\mathcal{T}_0, \dots, \mathcal{T}_J)$ of triangulations be a multilevel triangulation and $\mathcal{H} = (\mathcal{G}_0, \dots, \mathcal{G}_J)$ the corresponding hierarchical decomposition. In this section we introduce a particular format for the distribution of the tetrahedra in \mathcal{H} among processors on a parallel machine. We assume that the processors are numbered $1, \dots, P$.

For the set of elements in the hierarchical surplus on level k that are stored on processor p we introduce the notation

$$\mathcal{G}_k(p) := \{ T \in \mathcal{G}_k \mid T \text{ is stored on processor } p \}$$

and we define

$$\mathcal{H}(p) := (\mathcal{G}_0(p), \dots, \mathcal{G}_J(p)) .$$

In general the intersection $\mathcal{G}_k(p) \cap \mathcal{G}_k(q)$, $p \neq q$, may be nonempty. Also note that in general $\mathcal{H}(p)$ is not a hierarchical decomposition (in the sense of definition 6). The sequence

$$\tilde{\mathcal{H}} = (\mathcal{H}(1), \dots, \mathcal{H}(P)) \tag{4.1}$$

is called a *distributed hierarchical decomposition* (corresponding to \mathcal{H}).

For each level k and processor p we introduce a set of *master elements*, $\text{Ma}_k(p) \subset \mathcal{G}_k(p)$, and a set of *ghost elements*, $\text{Gh}_k(p) \subset \mathcal{G}_k(p)$. In the formulation of the conditions below we use: $\mathcal{K}(T) := \emptyset$ if $\text{status}(T) = \text{NoRef}$, and $\text{Ma}_{J+1}(p) := \emptyset$.

We now formalize the conditions on data distribution as follows.

DEFINITION 8 (Admissible hierarchical decomposition). The distributed hierarchical decomposition $\tilde{\mathcal{H}}$ is called an *admissible hierarchical decomposition* if for all $k = 1, \dots, J$ the following conditions are fulfilled:

- (A1) **Partitioning of $\mathcal{G}_k(p)$:** The sets of masters and ghosts form a disjoint partitioning of $\mathcal{G}_k(p)$:

$$\forall p \quad \text{Ma}_k(p) \cup \text{Gh}_k(p) = \mathcal{G}_k(p) \quad \text{and} \quad \text{Ma}_k(p) \cap \text{Gh}_k(p) = \emptyset$$

- (A2) **Existence:** Every element from \mathcal{G}_k is represented as a master element on level k :

$$\mathcal{G}_k = \bigcup_{p=1}^P \text{Ma}_k(p)$$

- (A3) **Uniqueness:** Every element from \mathcal{G}_k is represented by at most one master element on level k :

$$\forall p_1, p_2 : \quad \text{Ma}_k(p_1) \cap \text{Ma}_k(p_2) \neq \emptyset \Rightarrow p_1 = p_2$$

(A4) **Child–parent locality:** A child master element and its parent (as master or ghost) are stored on the same processor:

$$\forall p \quad \forall T \in \mathcal{G}_k \quad \forall T' \in \mathcal{K}(T) : \quad T' \in \mathbf{Ma}_{k+1}(p) \Rightarrow T \in \mathcal{G}_k(p)$$

(A5) **Ghosts are parents:** Ghost elements always have children:

$$\forall p \quad \forall T \in \mathbf{Gh}_k(p) : \quad \mathcal{K}(T) \neq \emptyset$$

(A6) **Ghost–children locality:** A ghost element and its children are stored on the same processor:

$$\forall p \quad \forall T \in \mathbf{Gh}_k(p) : \quad \mathcal{K}(T) \subset \mathbf{Ma}_{k+1}(p)$$

REMARK 5. Consider a consistent initial triangulation $\mathcal{T}_0 = \mathcal{G}_0$ with a nonoverlapping distribution of the tetrahedra: $\mathcal{G}_0(p) \cap \mathcal{G}_0(q) = \emptyset$ for all $p \neq q$. In this case all tetrahedra can be stored as masters and there are no ghosts. Then the distributed hierarchical decomposition $\tilde{\mathcal{H}} = ((\mathcal{G}_0(1)), \dots, (\mathcal{G}_0(P)))$ is admissible.

Two elementary results are given in:

LEMMA 4.1. *Let $\tilde{\mathcal{H}}$ as in (4.1) be a distributed hierarchical decomposition. The following holds:*

1. *If the conditions (A3), (A5) and (A6) are satisfied then for any element from \mathcal{G}_k there is at most one corresponding ghost element:*

$$\forall T \in \mathcal{G}_k \quad \forall p, q : \quad T \in \mathbf{Gh}_k(p) \cap \mathbf{Gh}_k(q) \Rightarrow p = q$$

2. *If the conditions (A1), (A2), (A3), (A4) and (A6) are satisfied then all children of a parent are stored as master elements on one processor:*

$$\forall T \in \mathcal{G}_k \quad \exists p : \quad \mathcal{K}(T) \subset \mathbf{Ma}_{k+1}(p)$$

Proof.

1. Take $T \in \mathcal{G}_k$ and p, q such that $T \in \mathbf{Gh}_k(p) \cap \mathbf{Gh}_k(q)$. Then because of (A5) and (A6) we obtain $\emptyset \neq \mathcal{K}(T) \subset \mathbf{Ma}_{k+1}(p) \cap \mathbf{Ma}_{k+1}(q)$. From (A3) we conclude $p = q$.
2. Take $T \in \mathcal{G}_k$. For $\mathcal{K}(T) = \emptyset$ nothing has to be proven, so we consider $\mathcal{K}(T) \neq \emptyset$. Choose $T' \in \mathcal{K}(T) \subset \mathcal{G}_{k+1}$. Then because of (A2)

$$\exists p : \quad T' \in \mathbf{Ma}_{k+1}(p) ,$$

and due to (A4) T is also stored on processor p : $T \in \mathcal{G}_k(p)$. From (A1) it follows that T is stored either as a ghost element or a master element on processor p . First consider $T \in \mathbf{Gh}_k(p)$. Then (A6) implies $\mathcal{K}(T) \subset \mathbf{Ma}_{k+1}(p)$, *qed*.

Now consider the case $T \in \mathbf{Ma}_k(p)$. Suppose another child $T'' \in \mathcal{K}(T)$, $T'' \neq T'$, is *not* stored as master element on processor p , i.e. $T'' \notin \mathbf{Ma}_{k+1}(p)$. From (A2) it follows that there is a processor q , $q \neq p$, with $T'' \in \mathbf{Ma}_{k+1}(q)$, and (A4) yields $T \in \mathcal{G}_k(q)$. From $T \in \mathbf{Ma}_k(p)$, (A3) and (A1) we obtain $T \in \mathbf{Gh}_k(q)$ and hence (A6) yields $\mathcal{K}(T) \subset \mathbf{Ma}_{k+1}(q)$. In particular we have $T' \in \mathbf{Ma}_{k+1}(q)$ and thus

$$T' \in \mathbf{Ma}_{k+1}(p) \cap \mathbf{Ma}_{k+1}(q) ,$$

with $p \neq q$. This is a contradiction to (A3). Hence for all children $T'' \in \mathcal{K}(T)$ $T'' \in \mathbf{Ma}_{k+1}(p)$ holds. □

REMARK 6. Due to the conditions (A2) and (A3) every tetrahedron $T \in \mathcal{H}$ can be assigned a unique processor on which T is stored as a master element. In other words, we have a well-defined function $\text{master} : \mathcal{H} \rightarrow \{1, \dots, P\}$ that is given by

$$\text{master}(T) = p \Leftrightarrow T \in \text{Ma}_{\ell(T)}(p) .$$

Let $T \in \text{Ma}_k(p)$ be a parent master element. From the second result in lemma 4.1 and (A4) it follows that either all children are masters on the same processor p as T , or they are masters on some other processor q . In the latter case, the element T has a corresponding ghost element on processor q . Due to this property, in the parallel refinement algorithm below we use the strategy:

If a parent tetrahedron T has a ghost copy then operations that involve children of T are performed on the processor on which the ghost and the children are stored. (4.2)

From condition (A4) it follows that a child master element has its parent (as ghost or as master) on the same processor. Therefore we use the strategy:

If a child tetrahedron T has a ghost copy then operations that involve the parent of T are performed on the processor on which the master element of T and its parent are stored. (4.3)

The first result in lemma 4.1 shows that every $T \in \mathcal{H}$ has at most one ghost copy. Moreover, due to (A5) all leaves ($T \in \mathcal{T}_J$) have no ghost copies. In this sense the overlap of tetrahedra between the processors is small.

4.3. Parallel multilevel refinement algorithm. In figure 4.1 the parallel version of algorithm *SerRefinement* is given. For the subroutines that are new compared to the serial version of the refinement algorithm we use a different font.

Algorithm <i>ParRefinement</i> ($\mathcal{G}_0(p), \dots, \mathcal{G}_J(p)$)	
for $k = J, \dots, 0$ do	// phase I
<i>DetermineMarks</i> ($\mathcal{G}_k(p)$);	(1)
<i>CommunicateMarks</i> ($\mathcal{G}_k(p)$);	(2)
<i>AccumulateEdgeCounter</i> ($\mathcal{G}_k(p)$);	(3)
<i>MarksForClosure</i> ($\mathcal{G}_k(p)$);	(4)
for $k = 0, \dots, J$ do if $\mathcal{G}_k(p) \neq \emptyset$ then	// phase II
if $k > 0$ then <i>MarksForClosure</i> ($\mathcal{G}_k(p)$);	(5)
if $k < J$ then <i>Unrefine</i> ($\mathcal{G}_k(p)$);	(6)
<i>Refine</i> ($\mathcal{G}_k(p)$);	(7)
<i>DetermineNewestLevel</i> ();	(8)

FIG. 4.1. *Parallel multilevel refinement algorithm.*

The structure of the algorithm is the same as in the serial case. Below we discuss the subroutines used in the algorithm *ParRefinement*. Most of these subroutines are modifications of the subroutines (with the same name) that are used in the serial algorithm in section 3.3. For these we only explain the main differences between the parallel and the serial versions.

Let us first explain some notions that are used in the following to describe certain parallelization aspects: An object is said to be stored on a *processor boundary*, if it

is stored on several processors at the same time. For these objects one often needs communication: The exchange of messages between corresponding local copies of objects on processor boundaries is called *interface communication*. The correlation of *local copies* and the corresponding *global objects* is realized by global identification numbers: Local copies (on different processors) with the same global ID represent the same global object.

The set of master elements on processor p which have ghost elements on some other processor is defined by

$$\text{HasGh}_k(p) := \text{Ma}_k(p) \cap \bigcup_{q=1, \dots, P} \text{Gh}_k(q) .$$

This set plays an important role in the description of the subroutines for the parallel case. We note that if a tetrahedron T has a corresponding ghost element on some other processor then for certain operations which involve T , its children or its parent, one has to decide whether these are performed using T or the ghost copy of T . This decision is based on the two strategies (4.2) and (4.3) formulated in remark 6.

DetermineMarks. This subroutine is presented in figure 4.2 and is almost the same as for the serial case. Only $T \in \mathcal{G}_k(p) \setminus \text{HasGh}_k(p)$ are involved in the block (4)–(10) since access to the children is needed (cf. (4.2)). In (5) and (8) edge counters

```

Function DetermineMarks( $\mathcal{G}_k(p)$ )
  for  $T \in \mathcal{G}_k(p)$  do
    if status( $T$ ) = NoRef then
      if  $T$  is regular and mark( $T$ ) = Ref then
        mark( $T$ ) := RegRef; (1)
        increase edge counters; (2)
      if  $k = 0$  and mark( $T$ ) = Del then
        mark( $T$ ) := NoRef; (3)
    else if  $T \notin \text{HasGh}_k(p)$  then
      if status( $T$ ) = RegRef then
        if  $\forall T' \in \mathcal{K}(T) : \text{mark}(T') = \text{Del}$  then
          mark( $T$ ) := NoRef; (4)
          if  $T \notin \text{Gh}_k(p)$  then decrease edge counters; (5)
        for  $T' \in \mathcal{K}(T)$  do
          if mark( $T'$ ) = Del then
            mark( $T'$ ) := NoRef; (6)
        else // status( $T$ ) = IrregRef
          if  $\exists T' \in \mathcal{K}(T) : \text{mark}(T') = \text{Ref}$ 
          or an edge of a child, which is not an edge of  $T$ ,
          is marked for refinement then
            mark( $T$ ) := RegRef; (7)
            if  $T \notin \text{Gh}_k(p)$  then increase edge counters; (8)
          else
            mark( $T$ ) := NoRef; (9)
          for  $T' \in \mathcal{K}(T)$  do
            mark( $T'$ ) := NoRef; (10)

```

FIG. 4.2. Subroutine DetermineMarks; parallel case.

are updated only by master tetrahedra to avoid multiple updates. The condition **if** T is regular can be checked without communication. For this it is important to note

that *ghost elements are always regular*: if a ghost element T was irregular, it would not have any children, which contradicts (A5). For a master element its parent is stored on the same processor (cf. (A4)). Hence no communication is needed in this subroutine.

CommunicateMarks. This subroutine is given in figure 4.3. The marks from the ghost copies are communicated to the corresponding master tetrahedra (interface communication). Once this information is available, the marks and edge counters of the tetrahedra $T \in \text{HasGh}_k(p)$ which were not treated in *DetermineMarks* are modified, if necessary. Note that, due to (A5), for $T \in \text{HasGh}_k(p)$ always $\text{status}(T) \neq \text{NoRef}$ holds.

```

Function CommunicateMarks( $\mathcal{G}_k(p)$ )
  for  $T \in \text{Gh}_k(p)$  do
     $m_T := \text{mark}(T);$  (1)
  for ( $q = 1, \dots, P, q \neq p$ ) do
     $\text{send } \{m_T : T \in \text{Gh}_k(p) \cap \text{Ma}_k(q)\}$ 
     $\text{to processor } q;$  (2)
  Receive corresponding messages
  from other processors; (3)
  for  $T \in \text{HasGh}_k(p)$  do
     $\text{mark}(T) := m_T;$  (4)
    if  $\text{status}(T) = \text{RegRef}$  then
      if  $m_T \neq \text{RegRef}$  then
        decrease edge counters; (5)
      else //  $\text{status}(T) = \text{IrregRef}$ 
        if  $m_T = \text{RegRef}$  then
          increase edge counters; (6)

```

FIG. 4.3. Subroutine CommunicateMarks.

AccumulateEdgecounter. For an edge E that is stored on several processors the value of its edge counter $C(E)$ is stored in a distributed manner, i.e. each local copy stores a contribution of the global value. Thus in order to obtain the global counter value $C(E)$ on each local copy of E , the local values have to be added up among the local copies. For this purpose interface communication involving the edges is needed.

MarksForClosure. This subroutine is the same as the one in the serial case in figure 3.5 (one only has to replace \mathcal{G}_k by $\mathcal{G}_k(p)$). Since after the call of *AccumulateEdgecounter*($\mathcal{G}_k(p)$) the edge refinement pattern is available on each tetrahedron $T \in \mathcal{G}_k(p)$ the marks can be set without communication.

Unrefine. This subroutine is shown in figure 4.4. In the first for-loop $T \in \text{HasGh}_k(p)$ are skipped since in (2) access to the children is needed (see (4.2)). Ghost elements that are marked for no refinement are deleted in (3) in order to ensure (A5). The second for-loop (4) is needed because it may happen that for $T' \in \text{Gh}_{k+1}(p)$ there is no parent on the same processor that removed the deletion labels from T' in (2). Note that the deletion of local objects that have copies on other processors causes communication since all copies have to be informed about the local deletion.

Refine. This subroutine is given in figure 4.5. It is almost the same as for the serial case except that the for-loop only iterates over $\mathcal{G}_k(p) \setminus \text{HasGh}_k(p)$, cf. (4.2). We emphasize that local objects created on processor boundaries have to be identified with each other. Consider, for example, an unrefined edge E which is stored on two

Function $Unrefine(\mathcal{G}_k(p))$
 Label all tetrahedra, vertices, edges and faces on level $k + 1$ for deletion; (1)
for $T \in \mathcal{G}_k(p) \setminus \text{HasGh}_k(p)$ **do if** $\text{status}(T) \neq \text{NoRef}$ **then**
 if $\text{mark}(T) = \text{status}(T)$ **then**
 Remove all deletion labels of the
 children of T and of their vertices, edges, faces; (2)
 else if $T \in \text{Gh}_k(p)$ **and** $\text{mark}(T) = \text{NoRef}$ **then**
 Remove T and its vertices, edges, faces,
 that are not needed anymore by other tetrahedra; (3)
for $T \in \text{Gh}_{k+1}(p)$ **do**
 Remove all deletion labels of T and of its vertices, edges, faces; (4)
 Remove all tetrahedra, vertices, edges and faces on level $k + 1$,
 that are labeled for deletion; (5)

FIG. 4.4. Subroutine *Unrefine*; parallel case.

different processors, say p and q , with a counter $C(E)$ that has changed and is now positive. Then for both local copies of E new midvertices V_p and V_q are created independently on processor p and q , respectively. The information that V_p and V_q represent the same global object has to be communicated between the processors p and q by identifying both midvertices with each other. The identification of edges and faces can be done in a similar way. In our implementation we use the DDD library to take care of this identification (cf. remark 8).

Function $Refine(\mathcal{G}_k(p))$
if $k = J$ **then** $\mathcal{G}_{k+1}(p) := \emptyset$; (1)
for $T \in \mathcal{G}_k(p) \setminus \text{HasGh}_k(p)$ **do**
 if $\text{mark}(T) \neq \text{status}(T)$ **then**
 Refine T according to $\text{mark}(T)$, i.e.
 create children of T as master elements; (2)
 $\text{status}(T) := \text{mark}(T)$; (3)
 for $T' \in \mathcal{K}(T)$ **do**
 Find existing vertices, edges and faces of T' ; (4)
 create missing vertices, edges and faces of T' ; (5)
 $\text{status}(T') := \text{NoRef}$; (6)

FIG. 4.5. Subroutine *Refine*; parallel case.

DetermineNewestLevel. In this subroutine, which is described in figure 4.6, the new global finest level index $\in \{J - 1, J, J + 1\}$ is determined. For this an MPI – Allreduce operation is used in (4).

REMARK 7. We note that *ghost copies are not created within the parallel refinement algorithm*. These are generated only in a data migration procedure that is used for load balancing. In this paper we do not discuss the load balancing and data migration algorithm. This will be presented in a separate paper. The main idea is as follows. The load balancing is performed by means of a graph partitioning method which assumes the dual graph G of a triangulation as input. In order to ensure that the children of a common parent are all assigned to the same processor, they are combined to one multi-node, inducing a reduced dual graph G' which is then passed to the graph partitioner. Based on the output of the graph partitioning method certain tetrahedra are migrated and certain ghost copies are created. It can be shown that

```

Function DetermineNewestLevel()
  if  $\mathcal{G}_J(p) = \emptyset$  then  $J_p := J - 1;$  (1)
  else if  $\mathcal{G}_{J+1}(p) \neq \emptyset$  then  $J_p := J + 1;$  (2)
  else  $J_p := J;$  (3)
  Determine  $J' := \max_{1 \leq p \leq P} J_p ;$  (4)
  if  $J_p < J'$ 
    for  $k = J_p + 1, \dots, J'$  do
       $\mathcal{G}_k(p) := \emptyset;$  (5)

```

FIG. 4.6. *Subroutine DetermineNewestLevel.*

given an input admissible hierarchical decomposition the data migration algorithm is such that its output is again an admissible hierarchical decomposition.

REMARK 8. We briefly comment on some implementation issues. The management of the distributed data is handled by the library DDD (which stands for *dynamic distributed data*, cf. [9]) also used in the package UG [5] for the same task. The behaviour of DDD can be adjusted to the user's requirements by means of several user-defined handler functions. The functionality of DDD includes the deletion of objects and the transfer of local objects from one processor to another (transfer module), interface communication across processor boundaries (interface module) and the identification of local objects from different processors (identification module). All communication is bundled such that the number of exchanged messages is kept small. All parallel communication and administration actions occurring in *ParRefinement* (e.g., the accumulation of edge counters by interface communication, the consistent deletion of local objects or the identification of local objects with each other) can be handled by DDD.

5. Analysis of the parallel refinement algorithm. In this section we analyze the algorithm *ParRefinement* applied to an input *admissible* hierarchical decomposition

$$\tilde{\mathcal{H}}^{\text{old}} = (\mathcal{H}^{\text{old}}(1), \dots, \mathcal{H}^{\text{old}}(P)) \quad (5.1)$$

as in (4.1) with $\mathcal{H}^{\text{old}}(p) = (\mathcal{G}_0^{\text{old}}(p), \dots, \mathcal{G}_J^{\text{old}}(p))$. The main result we prove (theorem 5.7) is that the output of this algorithm

$$\tilde{\mathcal{H}}^{\text{new}} = (\mathcal{H}^{\text{new}}(1), \dots, \mathcal{H}^{\text{new}}(P)) \quad (5.2)$$

is again an admissible hierarchical decomposition.

First we prove that the algorithm *ParRefinement* is well-defined:

THEOREM 5.1. *The algorithm ParRefinement applied to an admissible hierarchical decomposition is well-defined.*

Proof. We have to show that all data that are needed in the subroutines of *ParRefinement* applied to $\mathcal{G}_k(p)$ are available on processor p . By definition a tetrahedron $T \in \mathcal{G}_k(p)$ is stored (as master or ghost) on processor p . A tetrahedron is always stored together with its vertices, edges and faces. Hence information on these subsimplices is available on processor p if the tetrahedron is available on processor p .

For $T \in \mathcal{G}_k(p)$ the two critical points are references to its parent and references to its children. The former is needed only to check the condition *if T is regular*, which occurs at several places.

We first consider this condition. If T is a ghost element then it must be regular (hence the condition must be fulfilled), since if we assume T to be irregular then T has

no children and a contradiction follows from (A5). If T is a master element, then due to (A4) its parent is stored (as ghost or master) on processor p . Hence the reference to its parent is available on processor p and the condition can be checked.

We now consider the references to the children of T that occur at several places in the subroutines. A look at the subroutines shows that such a reference is needed only for $T \in \mathcal{G}_k(p) \setminus \text{HasGh}_k(p)$. Such a T is either a master element (on processor p) which does not have a ghost copy on another processor or a ghost element. In the former case it follows using (A4) that children of T must be master elements on processor p . If T is a ghost element then it follows from (A6) that the children are stored on processor p . We conclude that in all cases the references to children are available on processor p . \square

It is straightforward to check that for a given hierarchical decomposition of an input multilevel triangulation the serial algorithm *SerRefinement* and its parallel version *ParRefinement* yield the same output triangulation. In [6, 7] it is proved that this output is again a multilevel triangulation. Hence the distributed hierarchical decomposition in (5.2), which corresponds to the output multilevel triangulation is well-defined. In the remainder of this section we will prove that this output distributed hierarchical decomposition is admissible. We first show that the condition (A1) is satisfied. Here and in the remainder we always assume that *the input distributed hierarchical decomposition (5.1) is admissible*.

LEMMA 5.2. *The output distributed hierarchical decomposition (5.2) of algorithm ParRefinement satisfies condition (A1).*

Proof. The new set $\mathcal{G}_k^{\text{new}}(p)$ is obtained from the old one $\mathcal{G}_k^{\text{old}}(p)$ by removing elements (in *Unrefine*) and by creating new *master* elements (in *Refine*). The removal of elements does not destroy property (A1). If a new tetrahedron T is added to $\mathcal{G}_k^{\text{old}}(p)$ we obtain $\mathcal{G}'_k(p) := \mathcal{G}_k(p) \cup \{T\}$ and the partitioning

$$\text{Ma}'_k(p) := \text{Ma}_k(p) \cup \{T\}, \quad \text{Gh}'_k(p) := \text{Gh}_k(p),$$

which still has property (A1). \square

In order to prove that the other conditions (A2)-(A6) also hold we need a few lemmas. We use the notation $F(T)$ for the parent of T .

LEMMA 5.3. *Let $k \geq 1$ and let $T \in \mathcal{G}_k^{\text{new}}$ be a tetrahedron such that $T \notin \mathcal{G}_k^{\text{old}}$. Then the following holds:*

$$\exists p : T \in \text{Ma}_k^{\text{new}}(p) \tag{5.3}$$

$$T \in \text{Ma}_k^{\text{new}}(p) \Rightarrow F(T) \in \mathcal{G}_{k-1}^{\text{new}}(p) \tag{5.4}$$

Proof. A new tetrahedron $T \in \mathcal{G}_k^{\text{new}} \setminus \mathcal{G}_k^{\text{old}}$ can only be constructed in the subroutine *Refine*. There it is created as a master element on some processor p . Hence, the result (5.3) holds.

Let $T \in \text{Ma}_k^{\text{new}}(p)$ be a new master element which has a parent, i.e., it is obtained by applying a refinement rule to its parent $F(T) \in \mathcal{G}_{k-1}^{\text{old}}$. Since T is on processor p , this refinement must be performed in the call of the subroutine *Refine*($\mathcal{G}_{k-1}^{\text{old}}(p)$), and thus $F(T) \in \mathcal{G}_{k-1}^{\text{old}}(p)$. The element $F(T)$ can only be removed in the subroutine *Unrefine* on level $\ell < k$. This can not happen because the second phase of the refinement algorithm ((5)-(7)) is a bottom-up procedure. We conclude that $F(T) \in \mathcal{G}_{k-1}^{\text{new}}(p)$ holds. \square

LEMMA 5.4. Let $k \geq 0$. For $T \in \mathcal{G}_k^{\text{new}} \cap \mathcal{G}_k^{\text{old}}$ the following holds:

$$T \in \text{Ma}_k^{\text{old}}(p) \Leftrightarrow T \in \text{Ma}_k^{\text{new}}(p) \quad (5.5)$$

$$T \in \mathcal{G}_k^{\text{old}}(p) \text{ and } \mathcal{K}^{\text{new}}(T) \neq \emptyset \Rightarrow T \in \mathcal{G}_k^{\text{new}}(p) \quad (5.6)$$

Proof. Consider a tetrahedron $T \in \mathcal{G}_k^{\text{old}} \cap \mathcal{G}_k^{\text{new}}$. The result in (5.5) follows from the fact that nowhere in the algorithm *ParRefinement* a tetrahedron is moved from one processor to another or a master element is changed to a ghost element.

We now consider (5.6). If $T \in \text{Ma}_k^{\text{old}}(p)$ then the result directly follows from (5.5). Hence, we only have to consider a ghost element $T \in \text{Gh}_k^{\text{old}}(p)$ with $\mathcal{K}^{\text{new}}(T) \neq \emptyset$. Because the tetrahedron T has children in $\tilde{\mathcal{H}}^{\text{new}}$ we have $\text{mark}^{\text{new}}(T) \neq \text{NoRef}$. Assume that $T \notin \mathcal{G}_k^{\text{new}}(p)$ holds. Since tetrahedra are not moved between processors it follows that T has been removed from $\mathcal{G}_k^{\text{old}}(p)$. This can be caused only by the assignments (5) in $\text{Unrefine}(\mathcal{G}_{k-1}^{\text{old}}(p))$ or (3) in $\text{Unrefine}(\mathcal{G}_k^{\text{old}}(p))$. We first consider the latter case. From the bottom-up structure of the second phase of the refinement algorithm it follows that $\text{mark}(T)$ is not changed after the call of $\text{Unrefine}(\mathcal{G}_k^{\text{old}}(p))$. Thus, in this call we have $\text{mark}(T) \neq \text{NoRef}$ and the assignment (3) is not applied.

We now treat the other case. As outlined above we only have to consider a ghost element $T \in \text{Gh}_k^{\text{old}}(p)$. Thus assignment (4) in $\text{Unrefine}(\mathcal{G}_{k-1}^{\text{old}}(p))$ is applied to T and so T won't be removed in assignment (5) of this routine. Since there are no other cases left, we conclude that T is not removed from $\mathcal{G}_k^{\text{old}}(p)$, which yields a contradiction. \square

The result in (5.5) shows that master elements from $T \in \mathcal{G}_k^{\text{new}} \cap \mathcal{G}_k^{\text{old}}$ remain on the same processor. The result in (5.6) yields that elements from $T \in \mathcal{G}_k^{\text{new}} \cap \mathcal{G}_k^{\text{old}}$ that are refined in $\tilde{\mathcal{H}}^{\text{new}}$ remain on the same processor. The refinement condition $\mathcal{K}^{\text{new}}(T) \neq \emptyset$ is necessary (there may be ghost elements in $\mathcal{G}_k^{\text{old}}$ that are removed).

The next lemma shows that ghost elements in $\mathcal{G}_k^{\text{new}}$ on processor p must be ghost elements in $\mathcal{G}_k^{\text{old}}$ on the same processor.

LEMMA 5.5. The following holds:

$$T \in \text{Gh}_k^{\text{new}}(p) \Rightarrow T \in \text{Gh}_k^{\text{old}}(p)$$

Proof. This can be concluded from the following two properties. Firstly, tetrahedra are not moved from one processor to another. Secondly, new tetrahedra (created in the subroutine *Refine*) always are masters. \square

We now formulate a main result of our analysis.

THEOREM 5.6. *Let the input hierarchical decomposition $\tilde{\mathcal{H}}^{\text{old}}$ be admissible, i.e., it satisfies the conditions (A1)-(A6) for the levels $k = 0, \dots, J$. Then for the output hierarchical decomposition $\tilde{\mathcal{H}}^{\text{new}}$ the conditions (A2)-(A6) are satisfied for the levels $k = 0, \dots, J - 1$.*

Proof. Note that the levels $k = 0, \dots, J - 1$ exist in $\tilde{\mathcal{H}}^{\text{new}}$. Since the distributed input hierarchical decomposition $\tilde{\mathcal{H}}^{\text{old}}$ is admissible we can use that on the levels $k = 0, \dots, J$ the conditions (A1)-(A6) hold for $\tilde{\mathcal{H}}^{\text{old}}$. These properties are denoted by $(\text{A1})^{\text{old}} - (\text{A6})^{\text{old}}$.

(A2). By definition we have $\mathcal{G}_k^{\text{new}} = \bigcup_{p=1}^P \mathcal{G}_k^{\text{new}}(p)$ and $\text{Ma}_k^{\text{new}}(p) \subset \mathcal{G}_k^{\text{new}}(p)$. From this we get $\bigcup_{p=1}^P \text{Ma}_k^{\text{new}}(p) \subseteq \mathcal{G}_k^{\text{new}}$. For the other inclusion \supseteq we consider $T \in \mathcal{G}_k^{\text{new}}$. If $T \in \mathcal{G}_k^{\text{old}}$ then using (A2)^{old} we have:

$$\exists p \quad T \in \text{Ma}_k^{\text{old}}(p)$$

and thus, from lemma 5.4 it follows that $T \in \text{Ma}_k^{\text{new}}(p)$ holds. Now consider the case $T \notin \mathcal{G}_k^{\text{old}}$. On level 0 no new tetrahedra are created and thus $k \geq 1$ must hold. From lemma 5.3 we obtain:

$$\exists p \quad T \in \text{Ma}_k^{\text{new}}(p) .$$

We conclude that the inclusion $\bigcup_{p=1}^P \text{Ma}_k^{\text{new}}(p) \supseteq \mathcal{G}_k^{\text{new}}$ holds, too.

(A3). We have to prove:

$$\forall p, q : \quad \text{Ma}_k^{\text{new}}(p) \cap \text{Ma}_k^{\text{new}}(q) \neq \emptyset \Rightarrow p = q \quad (5.7)$$

This will be proved using an induction argument. For $k = 0$ we have $\text{Ma}_0^{\text{new}}(r) = \text{Ma}_0^{\text{old}}(r)$ for $r = 1, \dots, P$. Thus the result (5.7) for $k = 0$ follows from (A3)^{old}.

Now we take $k \geq 1$ and assume that the result in (5.7) holds for the lower levels ℓ , $0 \leq \ell < k$. Consider $T \in \text{Ma}_k^{\text{new}}(p) \cap \text{Ma}_k^{\text{new}}(q)$.

If $T \in \mathcal{G}_k^{\text{old}}$, then from lemma 5.4 we obtain

$$T \in \text{Ma}_k^{\text{old}}(p) \quad \text{and} \quad T \in \text{Ma}_k^{\text{old}}(q) .$$

From (A3)^{old} it follows that $p = q$ holds.

We now treat the other, more involved case $T \notin \mathcal{G}_k^{\text{old}}$. Then T must have a parent $F(T)$. From lemma 5.3 we get

$$F(T) \in \mathcal{G}_{k-1}^{\text{new}}(p) \cap \mathcal{G}_{k-1}^{\text{new}}(q) .$$

This parent can be stored as master or as ghost. We analyze the three possible situations:

1. Both on processor p and processor q the parent $F(T)$ is stored as a master element: $F(T) \in \text{Ma}_{k-1}^{\text{new}}(p) \cap \text{Ma}_{k-1}^{\text{new}}(q)$. From (5.7) on level $k - 1$ it follows that $p = q$ holds.
2. Both on processor p and processor q the parent $F(T)$ is stored as a ghost element: $F(T) \in \text{Gh}_{k-1}^{\text{new}}(p) \cap \text{Gh}_{k-1}^{\text{new}}(q)$. From lemma 5.5 it follows that

$$F(T) \in \text{Gh}_{k-1}^{\text{old}}(p) \cap \text{Gh}_{k-1}^{\text{old}}(q) .$$

From the first result in lemma 4.1 we conclude that $p = q$ holds.

3. Assume that $F(T)$ is stored as ghost on one processor and as a master on another one:

$$F(T) \in \text{Gh}_{k-1}^{\text{new}}(p) \quad \text{and} \quad F(T) \in \text{Ma}_{k-1}^{\text{new}}(q) , \quad p \neq q .$$

From lemma 5.5 we get $F(T) \in \text{Gh}_{k-1}^{\text{old}}(p)$ and thus also $F(T) \in \mathcal{G}_{k-1}^{\text{old}}$. Application of lemma 5.4 yields $F(T) \in \text{Ma}_{k-1}^{\text{old}}(q)$ and we conclude:

$$F(T) \in \text{HasGh}_{k-1}^{\text{old}}(q) . \quad (5.8)$$

The new tetrahedron $T \in \text{Ma}_k^{\text{new}}(q)$ can only be created as refinement of $F(T)$ in instruction (2) in the call of the subroutine $\text{Refine}(\mathcal{G}_{k-1}^{\text{old}}(q))$. However, this instruction is only performed if $F(T) \notin \text{HasGh}_{k-1}^{\text{old}}(q)$. Thus we obtain a contradiction with (5.8). We conclude that this third case can not occur.

(A4). Take $T \in \mathcal{G}_k^{\text{new}}$. If $\mathcal{K}^{\text{new}}(T) = \emptyset$ there is nothing to prove. We consider the situation $\mathcal{K}^{\text{new}}(T) \neq \emptyset$, $T' \in \mathcal{K}^{\text{new}}(T)$ with $T' \in \text{Ma}_{k+1}^{\text{new}}(p)$.

We first treat the case $T' \in \mathcal{G}_{k+1}^{\text{old}}$. Then for its parent T we have $T \in \mathcal{G}_k^{\text{old}}$ and $T' \in \mathcal{K}^{\text{old}}(T)$. Take q such that $T' \in \text{Ma}_{k+1}^{\text{old}}(q)$. Using lemma 5.4 we get $T' \in \text{Ma}_{k+1}^{\text{new}}(q)$. From the uniqueness property (A3) we conclude $q = p$. We thus have $T' \in \text{Ma}_{k+1}^{\text{old}}(p)$, and together with (A4)^{old} this yields $T \in \mathcal{G}_k^{\text{old}}(p)$. Now note that $\mathcal{K}^{\text{new}}(T) \neq \emptyset$ and use the second result in lemma 5.4. This implies

$$T \in \mathcal{G}_k^{\text{new}}(p) .$$

We now consider the case $T' \notin \mathcal{G}_{k+1}^{\text{old}}$. From lemma 5.3 we obtain:

$$\exists q \quad T' \in \text{Ma}_{k+1}^{\text{new}}(q) .$$

Due to the uniqueness property (A3) we have $q = p$. Application of the second result in lemma 5.3 yields

$$T = F(T') \in \mathcal{G}_k^{\text{new}}(p) .$$

So in both cases we have that the parent T is stored on the same processor p as T' .

(A5). Take $T \in \text{Gh}_k^{\text{new}}(p)$. Lemma 5.5 implies

$$T \in \text{Gh}_k^{\text{old}}(p) ,$$

and $\mathcal{K}^{\text{old}}(T) \neq \emptyset$ because of (A5)^{old}. Assume that $\mathcal{K}^{\text{new}}(T) = \emptyset$ holds. This implies $\text{mark}^{\text{new}}(T) = \text{NoRef}$. In the second phase of the refinement algorithm the mark of T is not changed after the call of $\text{Unrefine}(\mathcal{G}_k(p))$, and thus $\text{mark}(T) = \text{NoRef}$ during this call. In the instruction (2) of this call the element T is then removed, which implies $T \notin \text{Gh}_k^{\text{new}}(p)$. We obtain a contradiction and conclude that $\mathcal{K}^{\text{new}}(T) \neq \emptyset$ holds.

(A6). Take $T \in \text{Gh}_k^{\text{new}}(p)$. Lemma 5.5 implies $T \in \text{Gh}_k^{\text{old}}(p)$. From (A5)^{old} it follows that $\mathcal{K}^{\text{old}}(T) \neq \emptyset$. From (A5) we have $\mathcal{K}^{\text{new}}(T) \neq \emptyset$. Let

$$T' \in \mathcal{K}^{\text{new}}(T) \subset \mathcal{G}_{k+1}^{\text{new}}$$

be one of the children of T in $\tilde{\mathcal{H}}^{\text{new}}$.

First consider the case $T' \in \mathcal{K}^{\text{old}}(T)$. Due to $T \in \text{Gh}_k^{\text{old}}(p)$ and (A6)^{old} we get $T' \in \text{Ma}_{k+1}^{\text{old}}(p)$ and lemma 5.4 yields $T' \in \text{Ma}_{k+1}^{\text{new}}(p)$.

We now treat the other case, $T' \notin \mathcal{K}^{\text{old}}(T)$. Since $F(T') = T \in \text{Gh}_k^{\text{old}}(p)$ the new child tetrahedron T' must have been created in $\text{Refine}(\mathcal{G}_k(p))$, instruction (2). Thus $T' \in \text{Ma}_{k+1}^{\text{new}}(p)$ holds.

We conclude that $\mathcal{K}^{\text{new}}(T) \subset \text{Ma}_{k+1}^{\text{new}}(p)$. \square

Finally, we prove the main result:

THEOREM 5.7. *The output distributed hierarchical decomposition (5.2) of algorithm ParRefinement satisfies the conditions (A1)-(A6).*

Proof. The result concerning (A1) follows from lemma 5.3. It remains to prove that the conditions (A2) – (A6) are satisfied for the sets $\mathcal{G}_k^{\text{new}}(p)$, $k = 0, \dots, J^{\text{new}}$. Note that $J^{\text{new}} \in \{J-1, J, J+1\}$. We consider these three cases.

$J^{\text{new}} = J-1$. For this situation the desired result is proved in theorem 5.6.

$J^{\text{new}} = J$. From theorem 5.6 it follows that the sets $\mathcal{G}_k^{\text{new}}(p)$, $k = 0, \dots, J-1$ satisfy the conditions (A2) –(A6), hence we only have to consider $\mathcal{G}_J^{\text{new}}(p)$. First note

that, due to (A5)^{old} we have $\mathcal{G}_J^{\text{old}}(p) = \text{Ma}_J^{\text{old}}(p)$, i.e., the input finest level contains only master elements. Now consider an element $T \in \mathcal{G}_J^{\text{new}}(p)$. Then either $T \in \mathcal{G}_J^{\text{old}}(p)$, in which case T is a master element or T is a child of some element $F(T) \in \mathcal{G}_{J-1}^{\text{new}}$. In the latter case T is created in *Refine* (2) and must be a master element. We conclude that $\mathcal{G}_J^{\text{new}}(p) = \text{Ma}_J^{\text{new}}(p)$ and $\text{Gh}_J^{\text{new}}(p) = \emptyset$, i.e., there are no ghost elements in the output on level J . One easily verifies that for this situation the conditions (A1)-(A6) are satisfied for level J (as for \mathcal{G}_0 in remark 5).

$\mathbf{J}^{\text{new}} = \mathbf{J} + 1$. The arguments used in the case $J^{\text{new}} = J$ still hold and thus it follows that the conditions (A1)-(A6) hold for the levels $k = 0, \dots, J$. We only have to consider the set $\mathcal{G}_{J+1}^{\text{new}}(p)$. This set contains only new tetrahedra, i.e. elements which are created in *Refine* (2) by refinement of elements from $\mathcal{G}_J^{\text{new}}$. Thus all these elements are masters: $\mathcal{G}_{J+1}^{\text{new}}(p) = \text{Ma}_{J+1}^{\text{new}}(p)$ and $\text{Gh}_{J+1}^{\text{new}}(p) = \emptyset$. The same argument as used in the case $J^{\text{new}} = J$ yields that the conditions (A1)-(A6) are satisfied for $\mathcal{G}_{J+1}^{\text{new}}(p)$. \square

6. Numerical experiments. In this section we present results of a few numerical experiments. The computations were executed on the RWTH SunFire SMP cluster, using up to 64 900-MHz-processors. 1 GB RAM per processor was reserved by the batch system when starting a parallel job. The SMP nodes are connected to each other by Gigabit Ethernet, messages between processors of the same node are exchanged through the shared memory. Because of the multi-user mode one can not expect optimal runtimes. All runs were performed twice, then for each number of processors P the average of the measured runtimes was taken. The parallel runtimes are compared to the serial runtime in terms of the following definition to obtain a rough measure for the scalability of the parallel algorithm.

DEFINITION 9 (Scaled efficiency). Consider an algorithm of complexity $\mathcal{O}(n)$ where n denotes the characteristic problem size. For a given number of processor P let $T_n(P)$ denote the runtime of the parallel algorithm on P processors for a problem of total size n . Choose a *fixed local problem size* n' (based on the memory available per processor). The ratio

$$E_{n'}^{\text{sc}}(P) := \frac{T_{n'}(1)}{T_{P \cdot n'}(P)}$$

of serial runtime to parallel runtime on P processors for a fixed local problem size n' is called the *scaled efficiency*.

Note that the refinement algorithm is of complexity $\mathcal{O}(n)$ where n is the number of tetrahedra.

6.1. Experiment 1: global refinement. First, we consider the global refinement of the unit cube $\Omega = [0, 1]^3$. For P processors the initial triangulation \mathcal{T}_0 consists of $P \cdot 384$ tetrahedra. Thus, the problem size grows proportional to the processor number and we have a fixed local problem size. This triangulation is refined globally four times. The resulting finest triangulation \mathcal{T}_4 then consists of $P \cdot 1\,572\,864$ tetrahedra, the whole final multilevel triangulation \mathcal{M} contains $P \cdot 1\,797\,504$ tetrahedra.

The parallel runtimes (in seconds) for the last refinement step and the scaled efficiency for several numbers of processors P are shown in table 6.1. Satisfactory efficiencies are obtained and thus the parallel refinement algorithm turns out to be scalable for uniform refinement (at least, up to $P = 64$ processors).

6.2. Experiment 2: local refinement. In the second and third experiment we treat a problem with local refinement. In the second experiment we do not use load balancing, whereas in the third experiment a load distribution is used. We again

P	1	2	4	8	16	32	64
$T(P)$	69.45	70.96	71.31	73.42	79.48	86.48	94.79
$E^{sc}(P)$	100%	98%	97%	95%	87%	84%	80%

TABLE 6.1

Experiment 1: global refinement, last refinement step.

consider the unit cube $\Omega = [0, 1]^3$ with the same initial triangulation as in experiment 1. Let $B \subset \Omega$ be the ball with radius $r = 0.3$ centered at $(0.4, 0.4, 0.4)$. Before applying the refinement algorithm only those tetrahedra are marked for refinement, whose barycenter is located within B . This is repeated four times. The parallel runtimes (in seconds) for the last refinement step are presented in table 6.2. After the last refinement the number of tetrahedra per processor ranges from 384 up to 927 832. Due to this heavy load imbalance between the processors two effects can be observed. On the one hand efficiency is very poor even for fairly low processor numbers due to the fact that for some processors the local grid is refined almost uniformly whereas on other processors only few tetrahedra are refined. Thus some of the processors are idle waiting for others with much more computational load. On the other hand a large amount of the available memory will not be used, so the theoretical maximum problem size n_{\max} , which is determined by the overall memory supply of the parallel machine, cannot be reached without a suitable load balancing strategy.

P	1	2	4	8	16	32	64
$T(P)$	7.44	11.11	17.03	29.42	42.52	57.06	112.23
$E^{sc}(P)$	100%	67%	44%	25%	17%	13%	7%

TABLE 6.2

Experiment 2: local refinement without load balancing, last refinement step.

6.3. Experiment 3: local refinement with load balancing. Because of the disadvantages mentioned above it is clear, that one has to use load balancing to obtain better efficiency and better memory utilization. In the third experiment we consider the situation as in experiment 2 but now additionally apply load balancing before the last refinement step. In contrast to the previous two experiments ghost copies appear which are created during the load balancing phase. The parallel runtimes (in seconds) for the last refinement step are shown in table 6.3. One of course obtains much better results as for the imbalanced case in table 6.2. Nevertheless we get worse efficiency with local refinement as compared to the results for uniform refinement in table 6.1. This is not surprising as there is more communication (because of the ghost copies) and at the same time much less computational work: The average number of tetrahedra per processor after the last refinement step is about 200 000 whereas in the first experiment this number is about $1.8 \cdot 10^6$.

P	1	2	4	8	16	32	64
$T(P)$	8.07	8.63	9.37	10.47	12.14	14.73	23.97
$E^{sc}(P)$	100%	94%	86%	77%	67%	55%	34%

TABLE 6.3

Experiment 3: local refinement with load balancing, last refinement step.

Figures 6.1–6.3 show the distributed triangulation on $P = 4$ processors after the last




FIG. 6.1. *Distributed triangulation.*




FIG. 6.2. *Interior.*




FIG. 6.3. *Detail.*

refinement.

REMARK 9. In an adaptive parallel CFD code the parallel multilevel refinement algorithm is used in combination with discretization routines (creation of the stiffness matrix) and solution routines (iterative solvers). Usually the bulk of the arithmetic work is needed in the solution routine. In such a setting a deterioration of the parallel efficiency as in table 6.3 will probably hardly influence the overall parallel efficiency. The latter is mainly determined by the parallel efficiency of the solution method.

REFERENCES

- [1] E. Bänsch: *Local mesh refinement in 2 and 3 dimensions*, Impact of Computing in Science and Engineering 3, 1991, 181–191.
- [2] R.E. Bank: *PLTMG: A software package for solving elliptic partial differential equations, users' guide 6.0*, SIAM, Philadelphia, 1990.
- [3] R.E. Bank, A.H. Sherman, A. Weiser: *Refinement algorithms and data structures for regular local mesh refinement*, in Scientific computing (R. Stepleman, ed.), North-Holland, Amsterdam (1983), pp. 3–17.
- [4] P. Bastian: *Parallele adaptive Mehrgitterverfahren*. Teubner, Stuttgart 1996.
- [5] P. Bastian, K. Birken, K. Johannsen, S. Lang, N. Neuß, H. Rentz-Reichert and C. Wieners, *UG - A flexible software toolbox for solving partial differential equations*, Computing and Visualization in Science 1, 1997, 27–40.
- [6] J. Bey: *Tetrahedral grid refinement*, Computing 55, 1995, 355–378.
- [7] J. Bey: *Finite-Volumen- und Mehrgitterverfahren für elliptische Randwertprobleme*. Dissertation, Universität Tübingen, 1997.
- [8] J. Bey: *Simplicial grid refinement: on Freudenthal's algorithm and the optimal number of congruence classes*, Numer. Math. 85, 2000, pp. 1–29.
- [9] K. Birken: *Dynamic Distributed Data (DDD) — a software tool for distributed memory parallelization*. Dokumentation, 1997.
<http://cox.iwr.uni-heidelberg.de/~ddd/>
- [10] B. L. Chamberlain: *Graph Partitioning Algorithms for Distributing Workloads of Parallel Computations*. Technical report, University of Washington, 1998.
<http://www.cs.washington.edu/homes/brad/cv/pubs/degree/generals.html>
- [11] H. Freudenthal: *Simplizialzerlegungen von beschränkter Flachheit*, Annals of Mathematics 43, 1942, pp. 580–582.
- [12] S. Groß, J. Peters, V. Reichelt, A. Reusken: *The DROPS Package for Numerical Simulations of Incompressible Flows Using Parallel Adaptive Multigrid Techniques*. IGPM Report 211, 2002.
http://www.igpm.rwth-aachen.de/reports/html/rep_2002.html
- [13] S. Groß: *Parallelisierung eines adaptiven Verfahrens zur numerischen Lösung partieller Differentialgleichungen*, Master thesis, in German, RWTH Aachen, 2002.
<http://www-igpm.rwth-aachen.de/www/diplomarbeiten.html>
- [14] G. Haase: *Parallelisierung numerischer Algorithmen für partielle Differentialgleichungen*.

- Teubner, Stuttgart; Leipzig 1999.
- [15] G. Karypis, K. Schloegel, V. Kumar: *ParMETIS. Parallel Graph Partitioning and Sparse Matrix Ordering Library*. Dokumentation, University of Minnesota, Army HPC Research Center Minneapolis, 1998.
<http://www-users.cs.umn.edu/~karypis/metis/parmetis/>
 - [16] G. Karypis, K. Schloegel, V. Kumar: *Parallel Static and Dynamic Multi-Constraint Graph Partitioning*. University of Minnesota, Supercomputing Institute, 2001.
 - [17] KASKADE. A toolbox for adaptive multilevel codes.
<http://www.zib.de/Scisoft/kaskade2/>
 - [18] A. Liu, B. Joe: *Quality local refinement of tetrahedral meshes based on bisection*, SIAM J. Sci. Comput. 16, 1995, 1269–1291.
 - [19] A. Liu, B. Joe: *Quality local refinement of tetrahedral meshes based on 8-subtetrahedron subdivision*, Math. Comp. 65, 1996, 1183–1200.
 - [20] J.M.L. Maubach: *Local bisection refinement for N-simplicial grids generated by reflection*, SIAM J. Sci. Comput. 16, 1995, 210–227.
 - [21] MG. A parallel multilevel platform for unstructured grids.
<http://rcswww.urz.tu-dresden.de/~jstiller/projects/mg/>
 - [22] W.F. Mitchell: *Adaptive refinement for arbitrary finite-element spaces with hierarchical basis*, J. Comput. Appl. Math. 36, 1991, 65–78.
 - [23] M.C. Rivara: *Algorithms for refining triangular grids suitable for adaptive and multigrid techniques*, International Journal of Numerical Methods in Engineering 20, 1984, 745–756.
 - [24] C.T. Traxler: *An algorithm for adaptive mesh refinement in n dimensions*, Computing 59, 1997, 115–137.
 - [25] C. Walshaw, M. Cross, M. Everett: *Parallel Dynamic Graph Partitioning for Adaptive Unstructured Meshes*. J. Parallel Distrib. Comput. 47 (1997), 102–108.
 - [26] S. Zhang: *Successive subdivisions of tetrahedra and multigrid methods on tetrahedral meshes*, Houston J. Math. 21, 1995, 541–556.