

The effect of parallelization on a tetrahedral mesh optimization method

D. Benitez, E. Rodríguez, J.M. Escobar, R. Montenegro

SIANI Research Institute

University of Las Palmas de Gran Canaria, SPAIN



Motivation

- **Untangling and smoothing of tetrahedral meshes is an important step in the optimization of meshes in problems with moving boundaries because it provides simulations with elements of good quality**
- **Improving the speed of mesh generation helps users iterate problem setup faster**
- **There are currently no **parallel** mesh untangling algorithms in existence**
- **There are also no **parallel** simultaneous mesh untangling and smoothing techniques in the literature**

Outline


- We propose a **new parallel algorithm** for simultaneous untangling and smoothing of **tetrahedral meshes** (it is a tetrahedral mesh optimization method)
- We also provide a detailed analysis of its parallelization on a **many-core computer**:
 - parallel scalability
 - load balancing
 - parallelism bottlenecks
 - influence of 3 graph coloring algorithms

Summary


- **Our mathematical approach to tetrahedral mesh optimization**
- **The novel parallel algorithm**
- **Experimental methodology**
- **Performance scalability**
- **Load balancing**
- **Parallelism bottlenecks**
- **Influence of coloring algorithms on parallel performance**
- **Conclusions and future work**

Our mathematical approach to tetrahedral mesh optimization


Equilateral tetrahedron




Tetrahedral mesh with non-valid tetrahedra (artificially tangled for our experimental tests)



Example of a valid tetrahedral mesh for a cube




Optimized tetrahedral mesh without non-valid tetrahedra




Our mathematical approach to tetrahedral mesh optimization

TANGLED MESHES


Bunny




Tube




Bone







Screwdriver



Toroid




HR toroid



UNTANGLED MESHES

Our mathematical approach to tetrahedral mesh optimization


Equilateral tetrahedron



Its quality (q) specifies the degree to which regularity is achieved:

$q=1$: equilateral tetrahedron

$q<1$: non-equilateral tetrahedron






(a) A small triangle on the surface makes a tetrahedron too thin

(b) A large triangle on the surface makes a tetrahedron too flat

(c) An appropriately sized surface triangle makes a tetrahedron of desired shape

Our mathematical approach to tetrahedral mesh optimization

- M : a tetrahedral mesh
- \mathbf{v} : inner mesh node
- $\mathbf{x}_{\mathbf{v}}$: node position
- $N_{\mathbf{v}}$: the local submesh (set of tetrahedra connected to the node \mathbf{v})
- $K(\mathbf{x}_{\mathbf{v}})$: objective function that measures the quality of the local submesh



Our mathematical approach to tetrahedral mesh optimization

- **Our untangling and smoothing technique finds the new position x_v that each inner mesh node v must hold, in such a way that $K(x_v)$ is optimized**
- **This process repeats several times for all the nodes of the mesh M**
- **Mathematical details:** J. M. Escobar , E. Rodriguez , R. Montenegro , G. Montero , J. M. Gonzalez-Yuste (2003) Simultaneous untangling and smoothing of tetrahedral meshes. Computer Methods in Applied Mechanics and Engineering, 192: 2775-2787

The novel parallel algorithm

- **Sequential algorithm (SUS) for the simultaneous untangling and smoothing of a tetrahedral mesh M**

1. **function** OptimizeNode($x_{\mathbf{v}}$, $N_{\mathbf{v}}$)
2. Optimize objective function $K(x_{\mathbf{v}})$
3. **end function**
4. **procedure** SUS
5. $Q \leftarrow 0$
6. $k \leftarrow 0$
7. **while** $Q < \lambda$ **and** $k < \text{maxIter}$ **do**
8. **for** each vertex $\mathbf{v} \in M$ **do**
9. $x'_{\mathbf{v}} \leftarrow \text{OptimizeNode}(x_{\mathbf{v}}, N_{\mathbf{v}})$
10. **end do**
11. $Q \leftarrow \text{quality}(M)$
12. $k \leftarrow k+1$
13. **end do**
14. **end procedure**

The novel parallel algorithm

- **Sequential algorithm (SUS) for the simultaneous untangling and smoothing of a tetrahedral mesh M**

1. **function** OptimizeNode(x_v, N_v)
2. Optimize objective function $K(x_v)$
3. **end function**
4. **procedure** SUS
5. $Q \leftarrow 0$
6. $k \leftarrow 0$
7. **while** $Q < \lambda$ **and** $k < \text{maxIter}$ **do**
8. **for** each vertex $v \in M$ **do**
9. $x'_v \leftarrow \text{OptimizeNode}(x_v, N_v)$
10. **end do**
11. $Q \leftarrow \text{quality}(M)$
12. $k \leftarrow k+1$
13. **end do**
14. **end procedure**

INPUTS:

- M : tangled tetrahedral mesh,
- maxIter : maximum number of untangling and smoothing iterations
- N_v : set of tetrahedra connected to the free *node* v
- x_v : is the initial position of the free node
- x'_v : its position after optimization, which is implemented with procedure *OptimizeNode()*
- Q measures the lowest quality of a tetrahedron of M
- $\text{quality}()$: function that provides the minimum quality of mesh M

The novel parallel algorithm

- **Sequential algorithm (SUS) for the simultaneous untangling and smoothing of a tetrahedral mesh M**

1. **function** OptimizeNode($x_{\mathbf{v}}$, $N_{\mathbf{v}}$)
2. Optimize objective function $K(x_{\mathbf{v}})$
3. **end function**
4. **procedure** SUS
5. $Q \leftarrow 0$
6. $k \leftarrow 0$
7. **while** $Q < \lambda$ **and** $k < \text{maxIter}$ **do**
8. **for** each vertex $\mathbf{v} \in M$ **do**
9. $x'_{\mathbf{v}} \leftarrow \text{OptimizeNode}(x_{\mathbf{v}}, N_{\mathbf{v}})$
10. **end do**
11. $Q \leftarrow \text{quality}(M)$
12. $k \leftarrow k+1$
13. **end do**
14. **end procedure**

OUTPUT:

- an untangled and smoothed mesh M , whose minimum quality must be larger than an user-specified threshold λ

The novel parallel algorithm

- **Sequential algorithm (SUS) for the simultaneous untangling and smoothing of a tetrahedral mesh M**

1. function OptimizeNode(x_v, N_v)
2. Optimize objective function $K(x_v)$
3. end function
4. procedure SUS
5. $Q \leftarrow 0$
6. $k \leftarrow 0$
7. while $Q < \lambda$ and $k < \text{maxIter}$ do
8. for each vertex $v \in M$ do
9. $x'_v \leftarrow \text{OptimizeNode}(x_v, N_v)$
10. end do
11. $Q \leftarrow \text{quality}(M)$
12. $k \leftarrow k+1$
13. end do
14. end procedure

This algorithm iterates sequentially over all the mesh vertices in some order, at each step adjusting the spatial coordinates x'_v of a free node v

The novel parallel algorithm

- **Sequential algorithm (SUS) for the simultaneous untangling and smoothing of a tetrahedral mesh M**

1. function OptimizeNode(x_v , N_v)
2. Optimize objective function $K(x_v)$
3. end function
4. procedure SUS
5. $Q \leftarrow 0$
6. $k \leftarrow 0$
7. while $Q < \lambda$ and $k < \text{maxIter}$ do
8. for each vertex $v \in M$ do
9. $x'_v \leftarrow \text{OptimizeNode}(x_v, N_v)$
10. end do
11. $Q \leftarrow \text{quality}(M)$
12. $k \leftarrow k+1$
13. end do
14. end procedure

This process repeats several times for all the nodes of the mesh M

The novel parallel algorithm

- **Parallel algorithm (pSUS) for the simultaneous untangling and smoothing of a tetrahedral mesh M**

```
1. function OptimizeNode( $x_v, N_v$ )
2.   Optimize objective function  $K(x_v)$ 
3. end function
4. procedure Coloring( $G=(V,E)$ )
5.    $G$  is partitioned into independent
   sets  $I=\{I_1, I_2, \dots\}$  using  $C_1, C_2$  or  $C_3$ 
   coloring algorithm
6. end procedure
```

Its inputs are the same as described for sequential algorithm

```
7. procedure pSUS
8.    $I \leftarrow \text{Coloring}(G=(V,E))$ 
9.    $k \leftarrow 0$ 
10.   $Q \leftarrow 0$ 
11.  while  $Q < \lambda$  and  $k < \text{maxIter}$  do
12.    for each independent set  $I_i \in I$  do
13.      for each vertex  $v \in I_i$  in parallel do
14.         $x'_v \leftarrow \text{OptimizeNode}(x_v, N_v)$ 
15.      end do
16.    end do
17.     $Q \leftarrow \text{quality}(M)$ 
18.     $k \leftarrow k+1$ 
19.  end do
20. end procedure
```

The novel parallel algorithm

- **Parallel algorithm (pSUS) for the simultaneous untangling and smoothing of a tetrahedral mesh M**

```
1. function OptimizeNode( $x_v, N_v$ )
2.   Optimize objective function  $K(x_v)$ 
3. end function
4. procedure Coloring( $G=(V,E)$ )
5.    $G$  is partitioned into independent
   sets  $I=\{I_1, I_2, \dots\}$  using  $C_1, C_2$  or  $C_3$ 
   coloring algorithm
6. end procedure
```

We implemented graph coloring with procedure `Coloring()`, which partitions the mesh in a disjoint sequence of independent sets: I_1, I_2, \dots

```
7. procedure pSUS
8.    $I \leftarrow \text{Coloring}(G=(V,E))$ 
9.    $k \leftarrow 0$ 
10.   $Q \leftarrow 0$ 
11.  while  $Q < \lambda$  and  $k < \text{maxIter}$  do
12.    for each independent set  $I_i \in I$  do
13.      for each vertex  $v \in I_i$  in parallel do
14.         $x'_v \leftarrow \text{OptimizeNode}(x_v, N_v)$ 
15.      end do
16.    end do
17.     $Q \leftarrow \text{quality}(M)$ 
18.     $k \leftarrow k+1$ 
19.  end do
20. end procedure
```


The novel parallel algorithm

- **Parallel algorithm (pSUS) for the simultaneous untangling and smoothing of a tetrahedral mesh M**

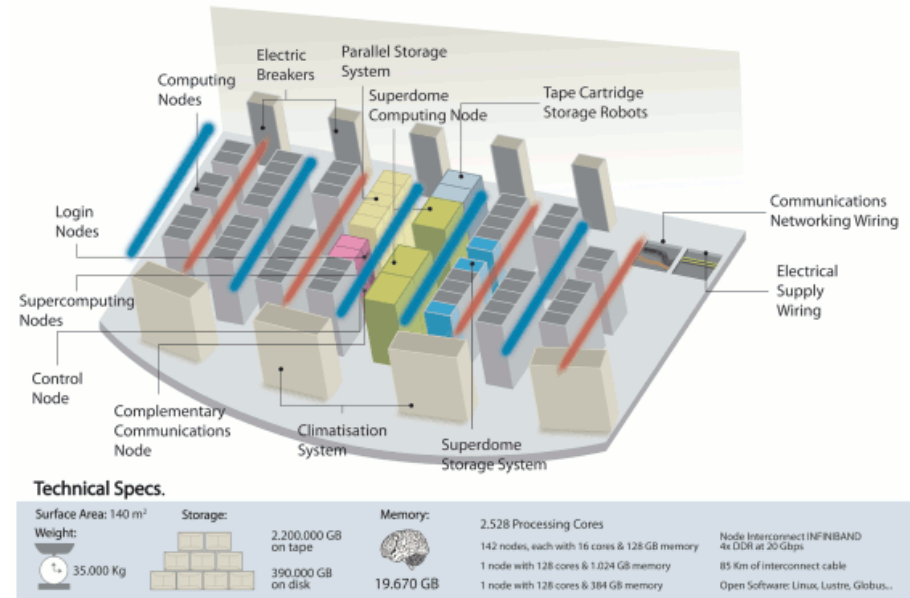
```
1. function OptimizeNode( $x_v, N_v$ )
2.   Optimize objective function  $K(x_v)$ 
3. end function
4. procedure Coloring( $G=(V,E)$ )
5.    $G$  is partitioned into independent
   sets  $I=\{I_1, I_2, \dots\}$  using  $C_1, C_2$  or  $C_3$ 
   coloring algorithm
6. end procedure
```

This parallel algorithm optimize
in parallel the nodes of each
independent set

```
7. procedure pSUS
8.    $I \leftarrow \text{Coloring}(G=(V,E))$ 
9.    $k \leftarrow 0$ 
10.   $Q \leftarrow 0$ 
11.  while  $Q < \lambda$  and  $k < \text{maxIter}$  do
12.    for each independent set  $I_i \in I$  do
13.      for each vertex  $v \in I_i$  in parallel do
14.         $x'_v \leftarrow \text{OptimizeNode}(x_v, N_v)$ 
15.      end do
16.    end do
17.     $Q \leftarrow \text{quality}(M)$ 
18.     $k \leftarrow k+1$ 
19.  end do
20. end procedure
```

Experimental methodology


- **Finis Terrae supercomputer (www.cesga.es), the third largest supercomputer in Spain**



- **1 many-core node HP Integrity Superdome, with 128 cores Itanium Montvale and 1.024 GB (NUMA): shared memory architecture**

Experimental methodology

- **Six different tangled benchmark meshes**



Experimental methodology

- **Description of the tangled benchmark meshes. The quality of non-valid tetrahedra is considered zero. So, the minimum quality is zero for all meshes**

Name	Number of vertices (m)	Number of tetrahedra	Average mesh quality	Number of inverted tetrahedra	Maximum vertex degree	Object
“m=6358”	6358	26446	0.2618	2215	26	Bunny
“m=9176”	9176	35920	0.1707	13706	26	Tube
“m=11525”	11525	47824	0.2660	1924	26	Bone
“m=39617”	39617	168834	0.1302	83417	26	Screwdriver
“m=201530”	201530	840800	0.2409	322255	26	Toroid
“m=520128”	520128	2201104	0.0657	1147390	26	HR toroid

Experimental methodology

- **Intel C++ compiler 11.1 with “O2” flag**
- **Linux system kernel “2.6.16.53-0.8-smp”.**
- **The source code of the parallel version included OpenMP directives, which were disabled when the sequential version was compiled**
- **Both software versions were profiled with PAPI API, which uses performance counter hardware of Itanium 2 processors**
- **Hardware binding: processor and memory**

Experimental methodology


- For each benchmark mesh we run the parallel version multiple times using a given maximum number of active threads between 1 and 128
- Each run is divided into two phases
 - The first of them completely untangles a mesh. This phase loops over all mesh vertices repetitively
 - The second phase smoothes the mesh until successive iterations increases the minimum mesh quality less than 5%

Performance scalability


- True speed-up and parallel efficiency of the **body of the main loop**

$$x'_v \leftarrow \text{OptimizeNode}(x_v, N_v)$$

$$\text{Speed-Up} = \frac{t_s}{t_{N_c}} \quad \text{Parallel Efficiency} = 100\% \times \frac{S_{N_c}}{N_c}$$



m=6358



Performance scalability


- True speed-up and parallel efficiency of the body of the **complete parallel Algorithm**

procedure pSUS

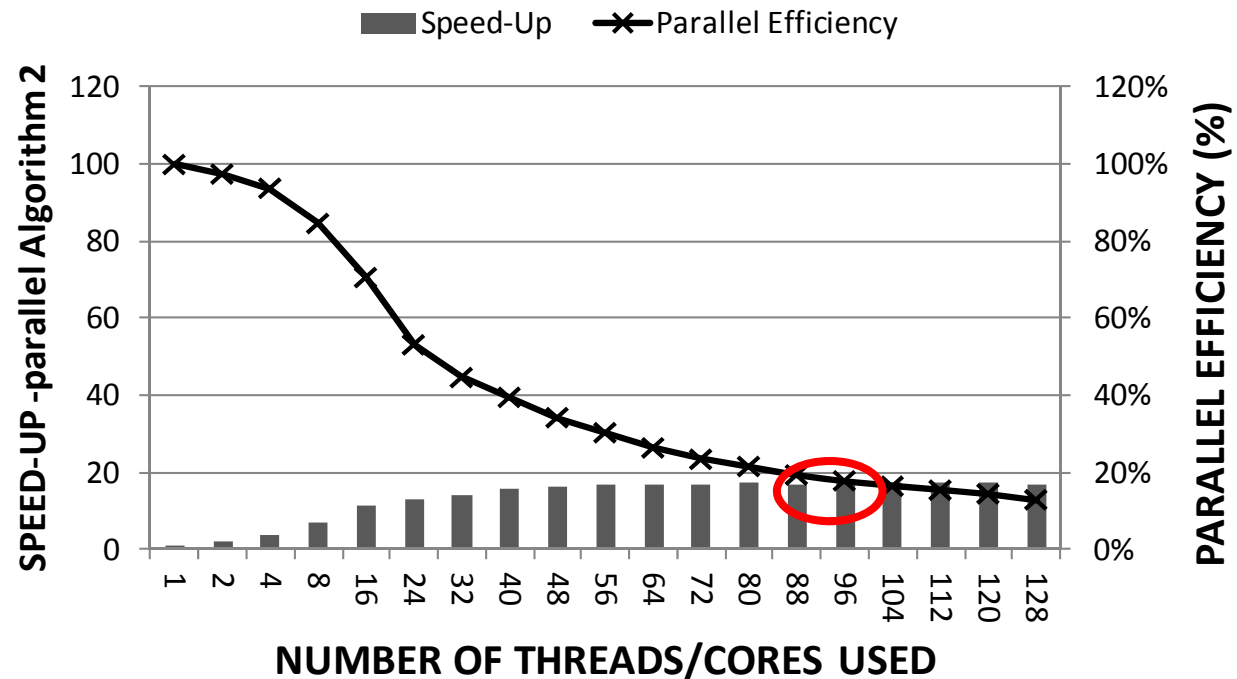
...

```

while Q < λ and k < maxIter do
  for each independent set Ii ∈ I do
    for each vertex v ∈ Ii in parallel do
      x'v ← OptimizeNode(xv, Nv)
  
```



m=6358



Performance scalability

- **Best runtime for the complete parallel algorithm (procedure pSUS)**

Name of tetrahedral mesh	Serial runtime (seconds)	Best parallel runtime (seconds)	Best number of cores	Best Speed-Up	Best parallel efficiency	Best coloring algorithm	Number of colors	Number of iterations (U&S)	Minimum mesh quality	Average mesh quality
m=6358	17.33	1.49	72	11.7X	16.2%	C ₁	29	25	0.1319	0.6564
m=9176	37.25	1.17	88	31.9X	36.3%	C ₃	29	26	0.2580	0.6823
m=11525	33.69	1.13	120	29.7X	24.8%	C ₃	10	38	0.1109	0.6474
m=39617	87.40	1.59	128	54.9X	42.9%	C ₁	31	11	0.1698	0.7329
m=201530	2505.37	81.28	128	30.8X	24.1%	C ₂	21	143	0.2275	0.6687
m=520128	2259.72	41.86	120	54.0X	45.0%	C ₃	34	36	0.2233	0.6750

Performance scalability

- Performance model for our parallel algorithm based on Amdahl's law

t_S	Sequential time	}	Speed-up
$t_{N_C}^P$	Parallel time <u>without</u> overhead		
$t_{N_C} = t_{N_C}^P + t_{N_C}^O$	Parallel time <u>with</u> overhead		

$$S_{N_C} = \frac{t_S}{t_{N_C}^P + t_{N_C}^O}$$

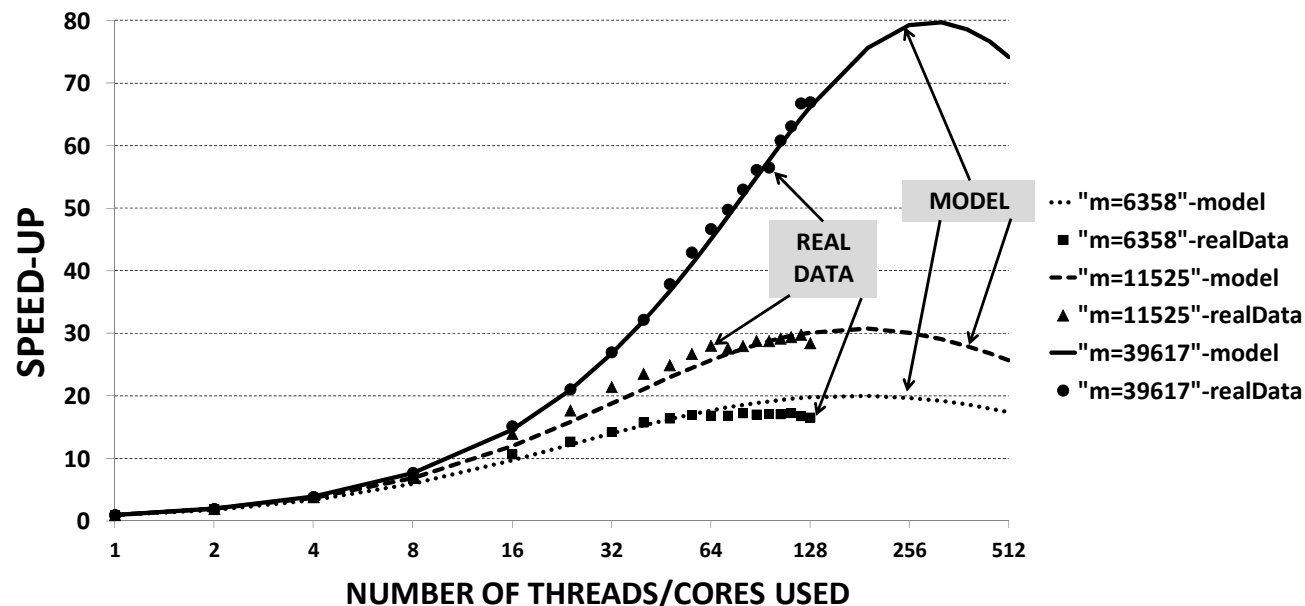
$$t = \frac{N_I}{IPC \times f} \quad \text{CPU time}$$

Performance scalability

- Performance model for our parallel algorithm based on Amdahl's law

$$S_{N_c} = \frac{t_s}{t_{N_c}^P + t_{N_c}^O} = \frac{\frac{t_s}{t_{N_c}^P}}{1 + \frac{t_{N_c}^O}{t_{N_c}^P}} = \frac{\frac{t_s}{t_{N_c}^P}}{1 + \theta_{N_c}} \rightarrow t_{N_c}^P \approx \frac{t_s}{N_c}$$

$$\theta_{N_c} = \frac{t_{N_c}^O}{t_{N_c}^P} \stackrel{f=cons}{=} \frac{IPC_{N_c}^P \times N_I^O}{IPC_{N_c}^O \times N_I^P}$$

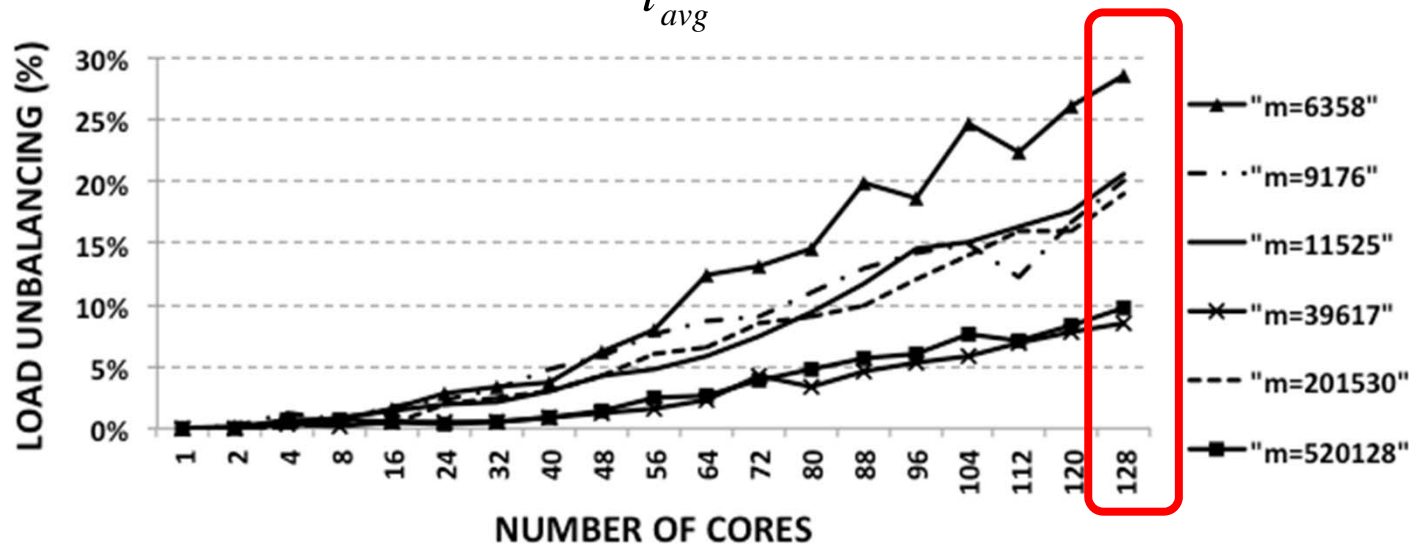


Conclusion:
parallel efficiency deteriorate as the number of threads increases because they tend to be dominated by the thread scheduling overhead

Load balancing

- Load unbalancing when meshes with different number of vertices (m) and up to 128 Itanium2 cores are used. C_3 coloring algorithm and dynamic OpenMP thread scheduling are used


$$L_{N_c} = 100\% \times \frac{t_{\max} - t_{\min}}{t_{\text{avg}}} \quad t_{\max} > t_{\text{avg}} > t_{\min} > 0$$




Parallelism bottlenecks

- During runtime of the main mesh optimization procedure, stall cycles of each parallel thread are in the range from 29%(1 core) to 58%(128 cores)
- These computation bottlenecks are located in:
 - **double-precision floating-point units**: from 70%(1c) to 27%(128c) of stall cycles
 - **data loads**: from 16%(1c) to 55%(128c)
 - cache memories: main source of data load stall cycles
 - NUMA (Non-Uniform Memory Access) memory: less than 1% of data load stall cycles
 - **branch instructions**: from 5%(1c) to 14%(128c)
 - **“no-operation” instructions (40%)**: caused by the long instruction format and compiler inefficiency

Influence of coloring algorithms on parallel performance




Distance-1 coloring :
adjacent nodes do
not have the same
color



An **independent set** of
a graph is a set of not
adjacent vertices

Influence of coloring algorithms on parallel performance



- Percentage of total parallel runtime that is required by the three graph coloring algorithms C_1 , C_2 , and C_3 when the six benchmark meshes are untangled and smoothed and 128 Itanium2 processors are used



- This means that the computational load required by our parallel algorithm is much heavier than required by graph coloring algorithms.

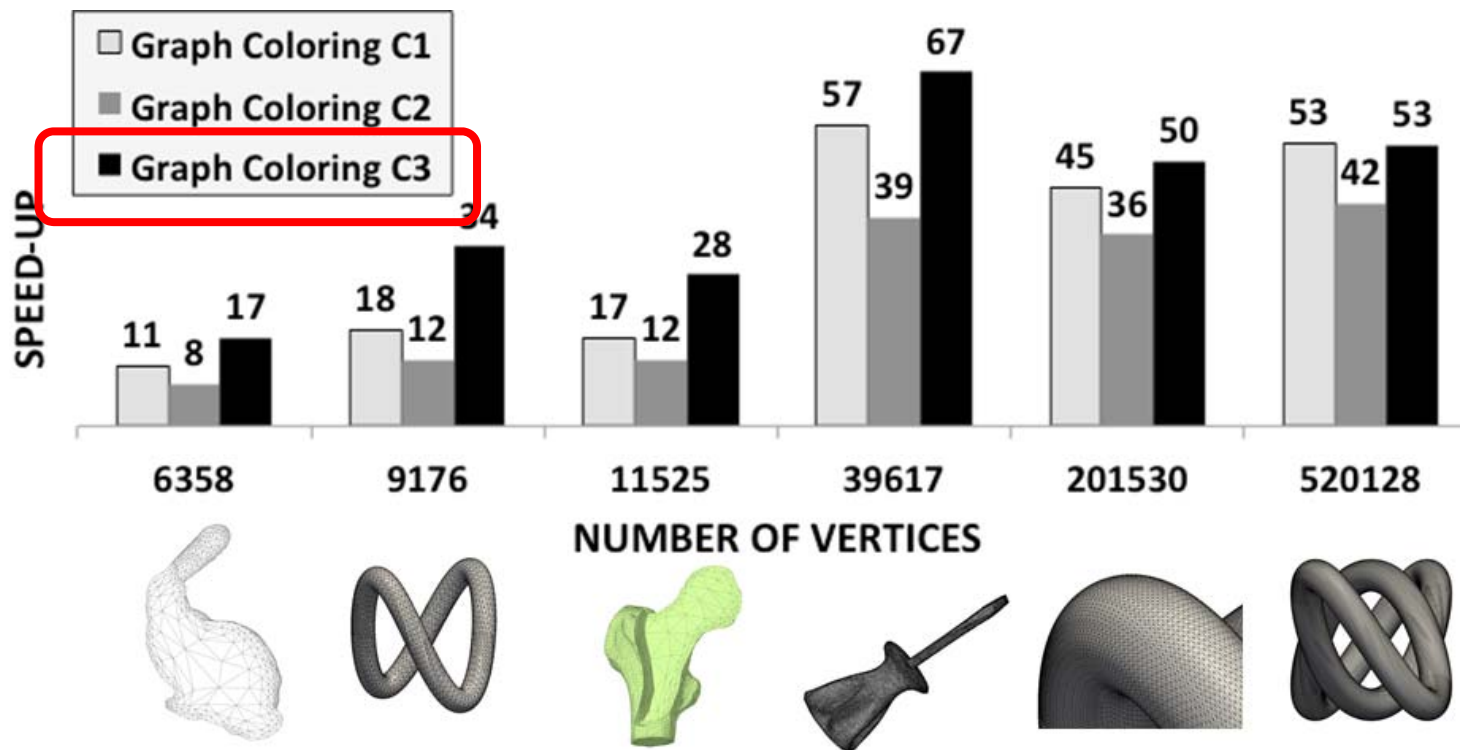
Influence of coloring algorithms on parallel performance

- Speed-up achieved by the complete parallel algorithm (pSUS) depends on the mesh coloring algorithm



Influence of coloring algorithms on parallel performance

- Speed-up achieved by the complete parallel algorithm (pSUS) for all six benchmark meshes when three graph coloring algorithms (C_1 , C_2 , and C_3) are used and all 128 shared-memory Itanium2 processors are active



Conclusions

- **We demonstrate that this algorithm is highly scalable when run on a high-performance shared-memory many-core computer with up to 128 Itanium 2 processors.**
- **It is due to the graph coloring algorithm that is used to identify independent sets of vertices without computational dependency**

Conclusions

- **We have analyzed the causes of its parallel deterioration on a 128-core shared-memory high performance computer using six benchmark meshes.**
- **It is mainly due to loop-scheduling overhead of the OpenMP programming methodology.**
- **The graph coloring algorithm has low impact on the total execution time. However, the total execution time of our parallel algorithm depends on the selected coloring algorithm.**

Future work

- **Our parallel algorithm is CPU bound and its demonstrated scalability potential for many-core architectures encourages us to extend our work to achieve higher performance improvements from massively parallel GPUs.**
- **The main problem will be to reduce the negative impact of global memory random accesses when the non-consecutive mesh vertices are optimized by the same streaming multiprocessor.**