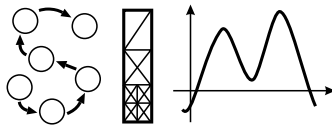


Lehrstuhl für Informatik 10 (Systemsimulation)



Hierarchical Hybrid Grids:

A Framework for Efficient Multigrid
on High Performance Architectures

Benjamin Karl Bergen

Abstract

For many scientific and engineering applications it is often desirable to use unstructured grids to represent complex geometries. Unfortunately the data structures required to represent discretizations on such grids typically result in extremely inefficient performance on current high-performance architectures. Here we introduce a grid framework using patch-wise, regular refinement that retains the flexibility of unstructured grids while achieving performance comparable to that seen with purely structured grids. This approach leads to a grid hierarchy suitable for use with geometric multigrid methods thus combining asymptotically optimal algorithms with extremely efficient data structures to obtain a powerful technique for large scale simulations.

1 Introduction

We begin in section 2 by discussing the underlying causes of the poor performance experienced when using purely unstructured grids. This will be followed in section 3 by a description of the Hierarchical Hybrid Grids(HHG) framework which offers one possible solution to this problem. Also in this section we will consider a caveat of the communication model being used and show that it does not have an adverse effect on the convergence rates obtained from a geometric multigrid algorithm. Finally, we will present performance results obtained on a Hitachi SR8000-F1 super computer for both a purely unstructured implementation and the HHG implementation.

2 Problems With Unstructured Grids

In this section we will look at three causes of poor performance on unstructured grids. In all three cases the underlying culprit is the type of data structure required to represent an unstructured grid, so it is worth our time to talk about the general characteristics of such data structures before we consider why they cause performance problems.

For our purposes when we talk about an unstructured grid we are really referring to a element mesh that might be used to represent a problem domain on which we would like to solve some equation or system of equations using a finite element or finite volume discretization. Such discretizations lead to generally sparse matrices that represent couplings between unknowns in the grid. Because the resulting matrix is sparse it makes sense to only store the non-zero values[1], however, because it is unstructured we must introduce some type of indirect indexing so that we can locate the non-zero entries on a row when we are applying an update. As it turns out this last statement contains the causes of all three of the problems that we will consider: First, indirect indexing means that we do not know which value we will need next; second, it is unlikely that related unknowns will be in close proximity to one another; and finally, the dereferencing needed to locate the positions of the non-zero values will require many integer operations. These causes correspond respectively to the problems discussed in the following subsections.

2.1 Prefetch Optimizations

This first problem we will address is caused by the extra level of indirection introduced by the use of indirect indexing. Because the current memory bandwidth available is not fast enough to keep up with the demands of modern CPUs[5], the ability to prefetch data before it is needed is of critical importance in producing efficient codes. Whether or not data can be prefetched depends largely on the predictability of data access. If the compiler is unable to determine a pattern of data access then it cannot schedule prefetch instructions which results in a substantial loss in performance. For unstructured grids it is not possible for the compiler to determine the data access patterns due to the indirect indexing of the data.

2.2 Cache Effects

The undesirable cache effects that one encounters with purely unstructured grids are largely due to a lack of *spatial locality*[4] of the underlying data structures used to represent the grid. The term *spatial locality* is used to indicate that if a cache line is retrieved because the processor has requested a particular byte of data¹ then the rest of the data in the cache line are likely to

¹The specific byte requested by the CPU is called the **critical word**

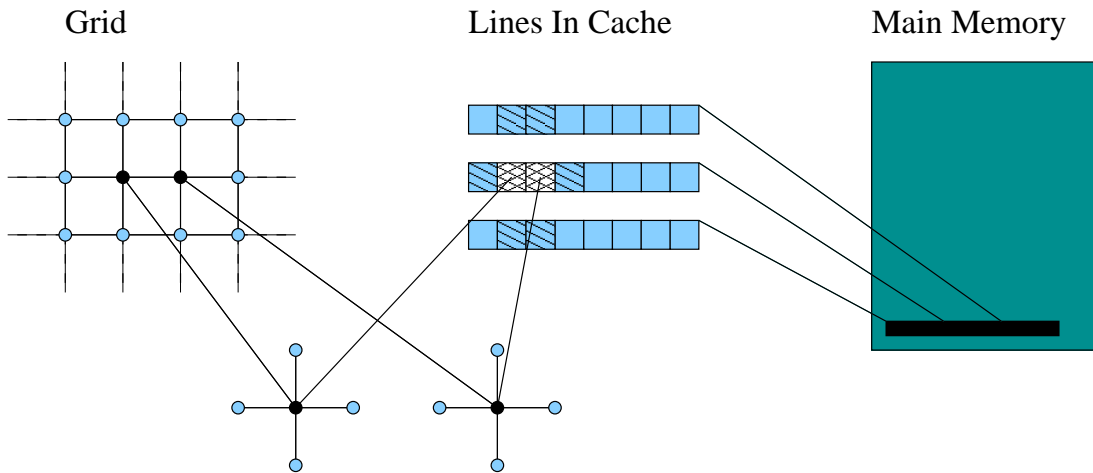


Figure 1: To understand what is meant by spatial locality consider the following scenario: Assume that no data is currently cached and the CPU is to update the left-most black data unknown by taking a weighted average of its nearest neighbors. It begins this process by loading the data unknown above the one being updated. This results in a cache miss which causes the top line in the cache lines diagram to be filled. Because the data is stored in a contiguous array the neighboring value of the right-most data unknown is also now in cache. The CPU now needs the data unknown to the left of the one being updated. This also results in a cache miss which causes the middle line in the cache lines diagram to be filled. Again, because of the contiguous storage scheme, the three data unknowns to the right of the one that caused the cache miss are also now in cache. The same step is repeated for the unknown below the one being updated with the result that when the right-most black data unknown is updated there are no cache misses because all of the required unknowns are already in the cache due to their spatial locality in memory to the unknowns of the previous step. Note that this example ignores the possibility of a cache conflict occurring.

be needed as well. To see how this works in a structured grid consider figure 1. Here we see that not only is the next unknown to be updated likely to be fetched in the case of a cache miss, but also the neighbors of the next unknown, since they are in close proximity to the neighbors of the current unknown in memory. In an unstructured grid it is generally not possible to come up with an ordering of the unknowns that will result in such favorable memory access patterns which consequently leads to poorer performance.

2.3 Instruction Mix

The final problem that we will consider has to do with the types of instructions that must be executed to access the non-zero entries of a matrix that represents an unstructured grid. Again indirect indexing is the underlying cause of the problem, however, in this case, we are not worried about the irregularity of the data access. Instead the concern here is the mix of operations that must be performed in updating an unknown. Most modern CPUs now have multiple *Arithmetic Logic Units* which means that under favorable conditions they are able to perform several operations at each clock cycle. Depending on the hardware capabilities of the processor and the mixture of tasks to be performed, i.e., integer versus floating point operations, many processors can achieve on-chip parallelism that can vastly increase performance. To see how indirect indexing might cause a problem in this context consider the schematic processor execution core in figure 2 which is typical of current (super-scalar) architectures and is able to perform up to four floating point operations per clock cycle. It should be clear that such processors might be at a disadvantage if one is trying to work with purely unstructured grids. For example, imagine that we would like to update an unknown by taking a weighted average of its neighbors. If the unknown in question has four neighbors then we would need five integer operations and seven floating point operations which would take no fewer than six clock cycles to complete². On a structured grid the same computation would take only three clock cycles. In this case the problem is not only that there is more work to be done than if we were working

²This does not include load and store operations and also assumes that the weights are constant.

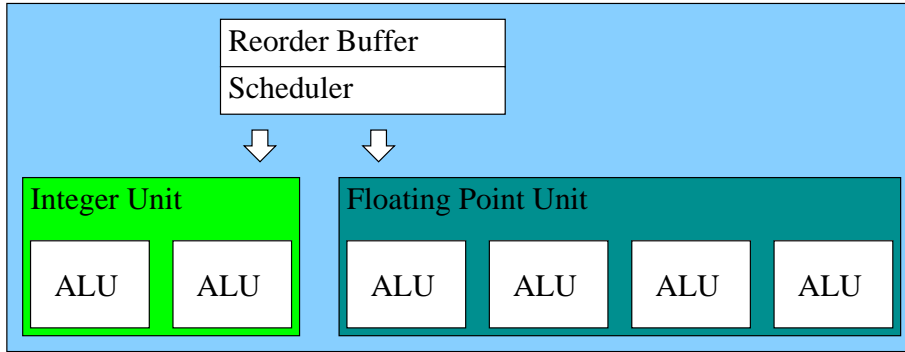


Figure 2: Processor execution core with two integer units and four floating point units. The efficiency of an algorithm often depends on having a favorable mix of instructions. The indirect indexing used in representing unstructured grids requires many integer operations and may cause an imbalance that will have a negative effect on performance.

with the structured grid, there is also the matter of scheduling to consider. Because the CPU only has two integer units the floating point operations cannot begin until the memory locations of at least two of the neighboring points have been resolved. Therefore, using an unstructured grid results in a bad instruction mix for this CPU. If the CPU had an equal number of integer and floating point units we would not see as much discrepancy in the number of clock cycles needed to complete the task so this problem is clearly architecture dependent.

3 Hierarchical Hybrid Grids

So far we have seen that the data structures needed to represent unstructured grids can cause some performance problems. Still, there are many applications for which we would like to use unstructured grids because of the flexibility they provide. In the next few subsections we present a possible solution to these problems that uses patch-wise regular refinement of a purely unstructured input grid to achieve nearly the same performance as that seen on structured grids (some early steps in this direction appear in [3]). We begin our discussion with an outline of the basic principles of HHG.

3.1 Basic Principles

The idea behind HHG is as follows: We begin with a purely unstructured input grid. This grid is assumed to be fairly coarse and is only meant to resolve the geometry of the problem being solved. This means, for example, that the input grid will resolve different material parameters but will not attempt to resolve problems such as shocks or singularities. It is also assumed that the desired resolution of the solution is much higher than that of the input grid so that some type of refinement is necessary both to ensure the proper resolution and to handle mathematical problems like those already mentioned. We then apply regular refinement to each patch of the input grid. Doing this successively, generates a nested grid hierarchy that is suitable for use with geometric multigrid algorithms. For an example of what such a grid hierarchy looks like consider figure 3.

Each grid in the new hierarchy is still logically unstructured, however, by using the right data structures we can now exploit the regularity of the patches. What we would ultimately like is for the neighbors of the interior points to occur at known, regular offsets from those points. This can now be accomplished by allocating separate blocks of memory for each individual patch. Then, patch-wise, we will have memory access patterns similar to a structured grid thus allowing stencil based operations for such things as smoothing, residual calculation, and transfer operations, and thereby avoiding the performance penalties associated with indirect indexing. This is essentially a variant of using block structured grids with the advantage that the structure and the resolution of the block connections are generated automatically. We will see in the next two subsections how the blocks are joined.

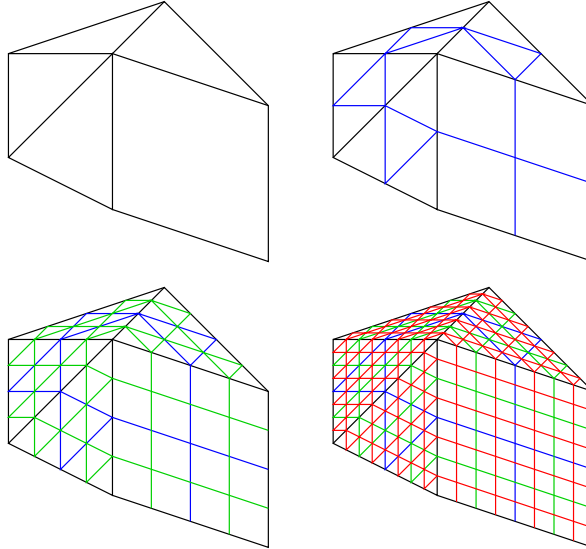


Figure 3: Beginning with the hybrid input grid on the top left, regular refinement is applied once to obtain the grid on the top right, and again to obtain the lower left grid, and again to finally obtain the lower right grid. After two levels of refinement (lower left grid) we can see that each patch has a structured interior. This structure may now be exploited to achieve higher performance.

3.2 Logical Objects

Exploiting the patch-wise regularity of our new grids is such a nice idea that it makes sense to ask if there is some more structure in this approach that may be exploited. The answer is yes, so long as we decompose the grid into the correct logical objects. Figure 4 shows that there is also a regular geometric structure along the edge unknowns. In order to use this regularity we again need to use the right data structures, i.e., each edge must be treated as a separate object in memory so that its neighbors occur at regular offsets. The relationship between the geometric structure of an edge and its memory representation is shown in figure 5 for an edge between two triangles.

Because they are inherently unstructured, it is not possible to treat the vertices of the patches in a structured manner, however, it does make sense to treat them as separate memory objects since they are updated differently from the elements and edges. This implies that we should decompose the input grid into elements, edges, and vertices³, and apply refinement to each object individually.

3.3 Communication Model

As a consequence of our grid decomposition scheme we must now introduce some communication between objects. This is necessary to ensure that every interior point in a structured region may be updated in the same manner and also to resolve the dependencies that exist between objects, e.g., some interior points of an element have neighboring points that are interior to a neighboring edge. We must therefore add *ghost* values around the edges of each logical object which correspond to the interior points of that objects neighbors. Interior and ghost points are shown in figure 6 for an edge and its neighboring elements.

In addition to communicating dependencies we must also rethink what an update of the unknowns on the grid should look like. In a normal treatment we would just loop over all unknowns in the grid applying an operation. However, this would simply be treating the refined grids as global unstructured grids and would then suffer from the same performance problems mentioned before. Exploiting the regularity of our new data structures is slightly more complicated although still straight forward.

The change is that we must now loop over each object type, i.e., elements, edges, and

³In three dimensions we have elements, faces, edges, and vertices where both the vertices and edges exhibit the unstructured nature of the input grid. Therefore, the elements and faces may be regularly refined and the resulting structure exploited.

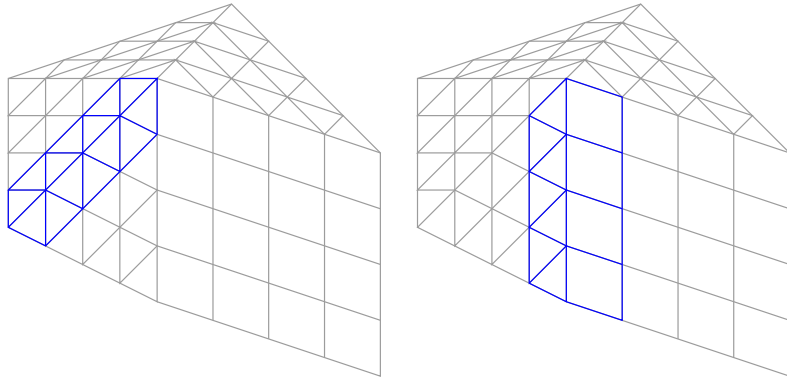


Figure 4: Edges also exhibit a regular structure which may be exploited. The grid on the left shows the structure of an edge between two triangular patches while the grid on the right shows the structure of an edge between a triangle and a quadrilateral. Although not pictured here, edges between two quadrilaterals also have a regular structure that may be exploited.

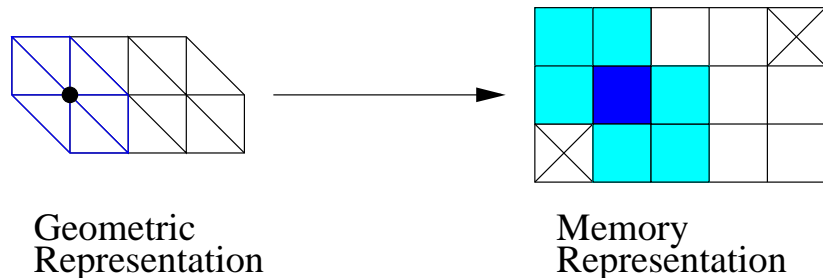


Figure 5: This edge may be expressed as a two-dimensional array in which case the light blue memory squares represent the neighboring points each of which has a regular offset from the dark blue interior point. The X's in the lower left and top right corners of the memory representation denote that these cells are not used. This edge has been regularly refined twice with the result that it has three interior points that may be operated on by the use of stencils.

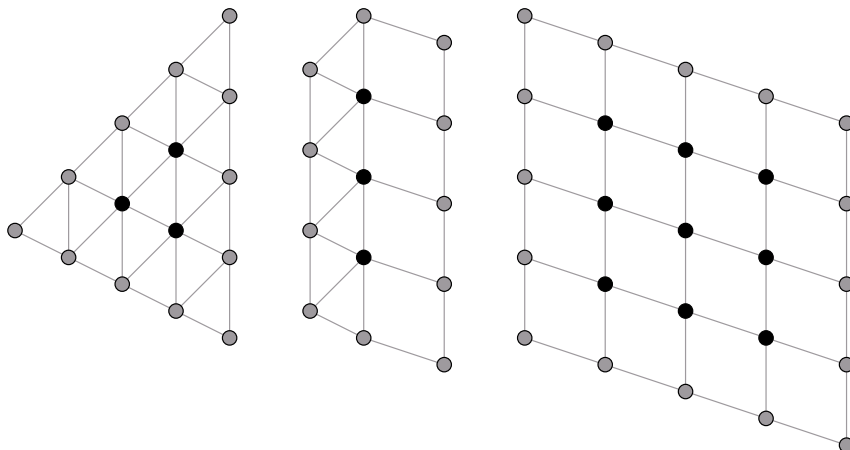


Figure 6: These three objects represent the result of decomposing two patches from the input grid into two elements and an edge. The black interior points may be updated independently provided that the gray ghost values are up to date.

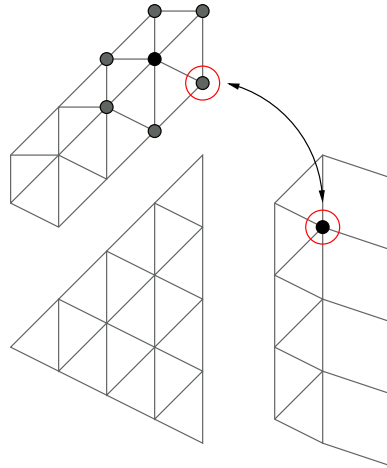


Figure 7: If the edge to the right of the element has already been updated then the edge above the element will not have the current value of the circled point.

vertices. As an example consider performing a GAUSS–SEIDEL step on one of the grids in our new hierarchy,

```

DO I=1,NUM_VERTICES
  ...smooth vertex(I)...
END DO

update dependencies

DO I=1,NUM_EDGES
  ...smooth edge(I)...
END DO

update dependencies

DO I=1,NUM_ELEMENTS
  ...smooth element(I)...
END DO

```

The differences here are that the edges and elements both contain unknowns that are being updated, and that updates of each object type require dependency information to be traded. Otherwise, the operations are the same albeit with a different ordering of the unknowns.

In order to simplify our implementation and to cut down on the amount of communication that must occur during an update, we have used what might be called a *one-way* communication model. What we mean by this is that each object type only communicates with the next *lower/simpler* object type, i.e., elements only communicate with edges, and edges only communicate with vertices. There is also no direct communication between objects of the same type. For the element object type this communication model does not cause any problems because no interior points of an element depend directly on the interior points of another element⁴. However, this does cause a difficulty for the edge objects. The problem here is that the interior points at each end of an edge are directly dependent on interior points from its two neighboring edges. To see this consider figure 7.

In the course of updating all edge objects, some dependencies will be updated but not communicated since there is no direct communication between the edges. For a GAUSS–SEIDEL iteration this will result in a JACOBI-like⁵ iteration for some of the

⁴In this case there is always an edge object in between whose interior unknowns provide a buffer.

⁵For any of these points at most two of its neighbors will have *stale* values from the previous iteration so that this type of update is somewhere between JACOBI and GAUSS–SEIDEL.

points. However, these missing dependencies do get updated before the next iteration by communication through the elements and for a specific input grid there are a fixed number of such points regardless of the level of refinement⁶. This communication model dramatically cuts down on the complexity of the HHG implementation and as we will see in the next section, it seems to have little or no effect on convergence.

3.4 Multigrid Convergence

To determine the effects of our communication model on multigrid convergence we solve the following problem,

$$\begin{aligned} -\Delta u &= f & \vec{x} \in \Omega &= [0, 1]^2 \\ u &= 0 & \vec{x} \in \partial\Omega & \end{aligned}$$

using a standard finite element discretization. If we use as our input grid the unit square subdivided into eight regular, triangular elements and apply five levels of refinement we obtain a similar problem to that for which multigrid convergence rates are given in [2] where $h = \frac{1}{64}$. We apply four-color GAUSS-SEIDEL to the elements and lexicographic GAUSS-SEIDEL/JACOBI to the edges. As reported in [2] we calculate our convergence rates by averaging the first five \mathcal{V} -cycles. The convergence rate for a $\mathcal{V}(2, 1)$ cycle for HHG is 0.09 which is comparable to the rate of 0.04 for red-black GAUSS-SEIDEL as reported in [2] and much better than JACOBI with a reported convergence rate of 0.24. We can therefore conclude that the ignored edge dependencies in our communication model do not have a serious effect on convergence.

4 Performance Results

Figure 8 shows our performance results which were computed on the porous media flow problem in figure 9. The results were generated on a single processor on a Hitachi SR8000-F1 super computer by applying different levels of regular refinement to the thirty eight element, unstructured input grid. The varying levels of performance seen in the different HHG schemes are largely due to the number of unknowns that may be updated in continuous loops without communication or transfer and residual operations being needed. For all of the HHG schemes high levels of refinement mean that many points can be updated without checking loop logic or communicating dependency information, so we see extremely good performance. When smoothing on all levels is added there is a performance hit because we are now working on some grids which are very coarse and require more book keeping operations than real work. The multigrid algorithm adds the extra burden of applying transfer operations and also works on all levels of refinement which gives it slightly lower performance in terms of MFLOP/s rates although it pays for itself many times over in convergence. As a point of reference, the results from an unstructured code are also present and show that HHG can be up to thirty times faster than conventional algorithms on grids that are heavily refined.

5 Conclusion

As scientist and engineers attempt to solve more difficult problems on complex geometries there will be an increasing need to turn to unstructured grid representations of problem domains. Traditional data structures for these grids fail to make efficient use of the resources of many high-performance architectures. Fortunately, as we have shown, it is possible to create efficient data structures that achieve extremely high performance as measured in MFLOP/s that still retain the salient features of purely unstructured grids. The process of developing such data structures requires careful consideration of

⁶Each edge object has two such points, one at each end. This does not depend on how many times an edge has been regularly refined.

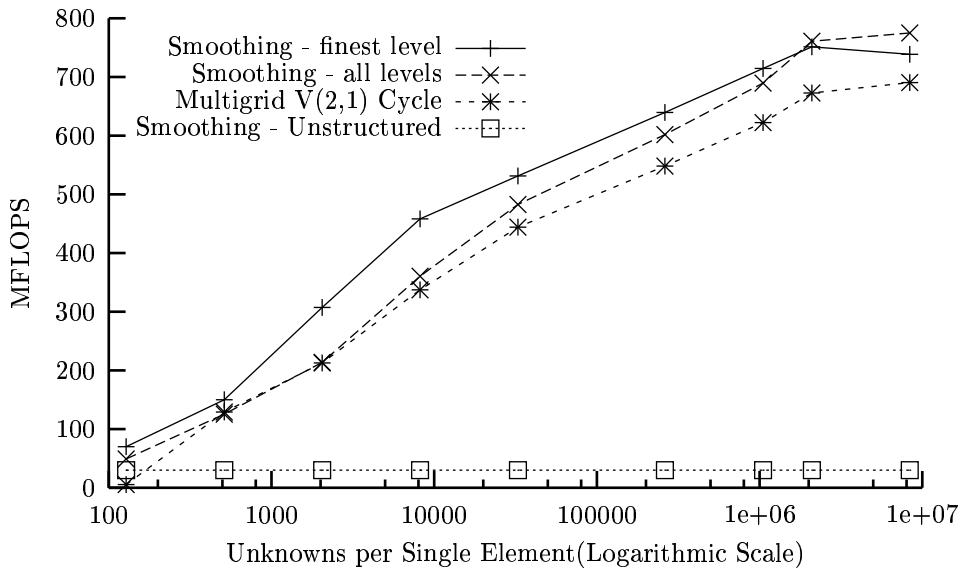


Figure 8: Performance results for the Hitachi SR8000-F1 for a single processor with a theoretical peak performance of 1.5 GFLOP/s. Here we see that smoothing only on the finest level of refinement gets the best performance. The performance decreases seen in smoothing all levels and in performing multigrid are due to the work that must be done on coarser grids which have fewer unknowns. On these levels the ratio of structured to unstructured work being done is not as favorable so we expect to see a slight drop in performance. The unstructured code performs poorly regardless of the number of unknowns which we should expect due to the problems cited in section 2.

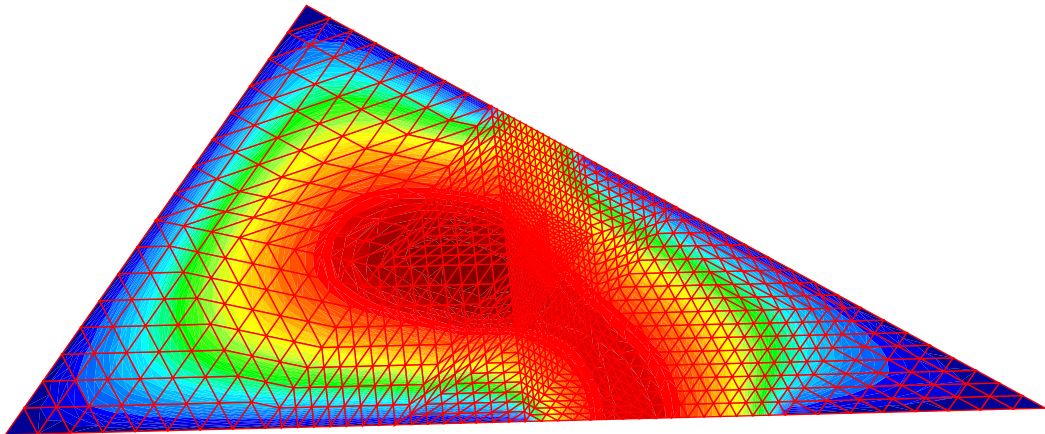


Figure 9: Porous media flow problem with constant forcing. The input grid for this problem has thirty eight elements.

the underlying processes that are being carried out and an increasingly acute understanding of how high-performance architectures work. However, in the end, we see that although HHG requires more elaborate data structures than a conventional approach the gain in performance is worth the effort.

References

- [1] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd Edition*. SIAM, Philadelphia, PA, 1994.
- [2] William L. Briggs, Van Emden Henson, and S. F. McCormick. *A Multigrid Tutorial*. SIAM, Philadelphia, PA, 2000.
- [3] F. Hülsemann, P. Kipfer, U. Rüde, and G. Greiner. gridlib: Flexible and efficient grid management for Simulation and Visualization. In *Computational Science - ICCS 2002*, Lecture Notes in Computer Science, pages 652–661. Springer, 2002.
- [4] L. Stals and U. Rüde. Techniques for Improving The Data Locality of Iterative Methods. Technical Report MRR97-038, School of Mathematical Science, Australian National University, October 1997.
- [5] C. Weiß, W. Karl, M. Kowarschik, and U. Rüde. Memory Characteristics of Iterative Methods. In *Proceedings of the ACM/IEEE SC99 Conference*, Portland, Oregon, November 1999.