A Performance Study of Data Layout Techniques for Improving Data Locality in Refinement-Based Pathfinding

ROBERT NIEWIADOMSKI, JOSÉ NELSON AMARAL, and ROBERT C. HOLTE University of Alberta

The widening gap between processor speed and memory latency increases the importance of crafting data structures and algorithms to exploit temporal and spatial locality. Refinement-based pathfinding algorithms, such as Classic Refinement (CR), find quality paths in very large sparse graphs where traditional search techniques fail to generate paths in acceptable time. In this paper, we present a performance evaluation study of three simple data structure transformations aimed at improving the data reference locality of CR. These transformations are robust to changes in computer architecture and the degree of compiler optimization. We test our alternative designs on four contemporary architectures, using two compilers for each machine. In our experiments, the application of these techniques results in performance improvements of up to 67% with consistent improvements above 15%. Analysis reveals that these improvements stem from improved data reference locality at the page level and to a lesser extent at the cache line level.

Categories and Subject Descriptors: E.1 [**Data Structures**]: Lists, Stacks, and Queues; I.2.8 [**Computing Methodologies**]: Artificial Intelligence—*Problem solving, control methods, and search:* Graph and tree search strategies

General Terms: Algorithms, Experimentation, Performance

Additional Key Words and Phrases: Cache-conscious algorithms, classical refinement, pathfinding

1. INTRODUCTION

The problem of finding a path linking two vertices in a graph is generally referred to as *pathfinding*. Pathfinding has applications in many industries such as computer games, freight transport, travel planning, circuit routing, network packet routing, and so on. For instance, in the real-time strategy (RTS) video-game genre pathfinding is used for conducting troop movement on the game map [DeLoura 2000]. In these games, pathfinding consumes up

Authors' address: R. Niewiadomski, J. N. Amaral, and R. C. Holte, Department of Computing Science, University of Alberta, Edmonton, AB, T6G-2E8, Canada; email: {niewiado,amaral,holte}@ cs.ualberta.ca.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org. © 2004 ACM 1084-6654/04/0001-ART01 \$5.00

to 50% of total computation time [Correspondence with David C. Pottinger of Ensemble Studios; Pottinger 2000]. Pathfinding typically entails finding a minimum length path. We can find a minimum length path linking two vertices in O(V+E) using the Breadth-First Search (BFS) algorithm [Cormen et al. 1991]. However, in a time-sensitive application, such as RTS video games, hundreds of paths may need to be generated in the span of a single second in a graph featuring millions of vertices. Under such circumstances the use of BFS results in unacceptable pathfinding performance. Applications can substantially improve pathfinding performance by settling for quality paths, that is, paths with lengths slightly longer than the minimum. Refinement-based search (RBS) algorithms generate a quality path linking two vertices by way of restricting the portion of the graph that is searched [Holte et al. 1996]. RBS algorithms have been shown to generate paths that are only 10% longer than the minimum while doing a fraction of the work done by BFS. Unsurprisingly, variants of RBS are employed in time-sensitive applications to achieve acceptable pathfinding performance. Classic Refinement (CR), an RBS algorithm, partitions a large graph into many subgraphs, and generates an *abstract graph* that describes the interconnections among the subgraphs. A path between two vertices, u and v in the original graph is found by: (1) identifying the vertices in the abstract graph that correspond to the partitions containing u and v; (2) finding a path, in the abstract graph, between the identified vertices; (3) using this abstract path to find a path in the original graph.

This paper presents a performance evaluation study of three techniques for improving the data reference locality of CR: (1) data reordering; (2) data duplication; (3) and merging of independent data structures into a common memory area. We demonstrate that combining these techniques can result in performance improvements of up to 67% with consistent improvements above 15%. Through analysis of hardware counter profiles, as well as memory access trace data, we demonstrate that these results stem from improved data reference locality at the page level and, to a lesser extent, at the cache line level. By testing on four different architectures with GCC and vendor compilers, we also demonstrate that our techniques are robust to changes in hardware as well as to the level of compiler optimization.

Section 2 presents a generic method for graph abstraction for RBS as well as the CR algorithm. Section 3 describes the baseline implementation and our three techniques. Section 4 presents our experimental framework while Section 5 contains the results and subsequent analysis of our experiments. Finally, Section 6 discusses related work.

2. ABSTRACTION AND SEARCH

For completeness, we describe the process used to create abstraction hierarchies and present the CR search algorithm.

2.1 Graph Abstraction

Let $G_0(V_0, E_0)$ be the input graph, also called the *source graph*, to a RBS algorithm. Let $G_1(V_1, E_1)$ be an *abstraction* of G_0 as defined below. Both G_0 and

A Performance Study of Data Layout Techniques 3



Please check figure closely as supplied art does not match art in PDF.

Fig. 1. A three-level abstraction hierarchy.

 G_1 are undirected and unweighted graphs. G_0 is partitioned into connected subgraphs. The abstract graph G_1 must have one vertex for each subgraph of G_0 . If a vertex v_i^0 in G_0 maps to a vertex v_p^1 in G_1 , we say that v_p^1 is the *image* of v_i^0 at abstraction level 1. (Note: v_p^1 should be read as "vertex p at abstraction level 1".) We also say that the set of vertices in G_0 that map to vertex v_n^1 in G_1 is the *preimage* of v_p^1 . The abstract graph G_1 has an edge (v_p^1, v_q^1) if and only if there is an edge (v_i^0, v_j^0) in G_0 such that v_i^0 belongs to the preimage of v_p^1 and v_i^0 belongs to the preimage of v_a^1 . This transformation ensures that paths in G_0 can be mapped to corresponding paths in G_1 .

2.2 Abstraction Hierarchies

Because we can create an abstract graph for any undirected graph, we can create an abstraction of an abstraction to generate an abstraction hierarchy. A sequence of graphs $\{G_0, G_1, \ldots, G_{n-1}\}$ is an abstraction hierarchy for source graph G_0 if G_{i+1} is an abstraction of G_i for $0 \le i < n-1$. An example of a three-level abstraction hierarchy is shown in Figure 1.

2.3 Abstraction Generation

To generate an abstraction hierarchy, we use the "max-degree" STAR algorithm [Holte et al. 1996]. The input is the source graph G_0 and a radius r. Given a graph G_a , the first step is to generate an abstraction G_{a+1} by partitioning G_a into connected subgraphs. The STAR algorithm partitions G_a into subgraphs whose maximum diameter is at most 2r. The algorithm selects a vertex v_i^a in G_a . If v_i^a is not in a subgraph, a new subgraph is created. Next a breadth-first traversal of G_a starting in v_i^a is used to find the other vertices of

this subgraph. This traversal ignores vertices that already belong to another subgraph. The traversal stops at vertices that are r edges away from v_i^a . Once all vertices in G_a are assigned to subgraphs, we proceed to generate G_{a+1} .

2.4 Refinement-Based Search

RBS algorithms use an abstraction hierarchy to find short paths between a start vertex and a goal vertex. Examples of RBS algorithms are CR, Path Marking, and Alternating Opportunism [Holte et al. 1996]. These algorithms vary in terms of efficiency and the quality of the paths generated. This paper focuses on $\rm CR.^1$

2.4.1 Definitions and Notations.

Definition 2.1. An ordered list of G_a vertices, $P = \{v_0^a, v_0^a, v_1^a, \dots, v_{k-1}^a\}$, is a *path* in G_a if and only if G_a contains the edges $(v_0^a, v_1^a), (v_1^a, v_2^a), \dots, (v_{k-2}^a, v_{k-1}^a)$. We use the notation P[j] to refer to the *j*th element in path P.

Definition 2.2. A path $P = \{v_0^a, v_1^a, \dots, v_{k-1}^a\}$ in G_a is a constrained path if and only if it is the shortest path between v_0^a and v_{k-1}^a , such that vertices $v_0^a, v_1^a, \dots, v_{k-1}^a$ belong to the preimage of the same vertex v_p^{a+1} .

Because the preimage of v_p^{a+1} is a connected subgraph of G_a , when computing a constrained path, a search algorithm can restrict its search space to the vertices in the preimage of v_p^{a+1} . In a constrained path all vertices are in the same preimage. G_0 may contain a shorter path between v_0^a and v_{k-1}^a than the constrained path P, but any such path will contain at least one vertex outside the preimage of v_p^{a+1} , and therefore will not be a constrained path.

Definition 2.3. Let v_p^{a+1} and v_q^{a+1} be two vertices in G_{a+1} such that (v_p^{a+1}, v_q^{a+1}) is an edge in G_{a+1} . Let v_i^a be a vertex in the preimage of v_p^{a+1} . Then there exists a path from v_i^a to any vertex in the preimage of v_q^{a+1} . A constrained jump path J from v_i^a to the preimage of v_q^{a+1} is the shortest path between v_i^a and any vertex in the preimage of v_q^{a+1} , such that any edge traversed by J connects vertices that belong either to the preimage of v_p^{a+1} or to the preimage of v_q^{a+1} .

Again, a shorter path from v_i^a to the preimage of v_p^{a+1} may exist in G_0 , but it would have to include at least one vertex outside the preimage of v_p^{a+1} or v_q^{a+1} and thus not be constrained.

2.4.2 *Classic Refinement.* Figure 2 presents pseudocode for the CR algorithm. Given a source graph G_0 and an abstraction hierarchy $A = \{G_0, G_1, \ldots, G_{n-1}\}$, we are interested in finding a path in G_0 between a source vertex s^0 and a goal vertex g^0 . The CR algorithm starts by finding a path, P_{n-1} , between s^{n-1} and g^{n-1} , the images of the source and goal vertices in the highest level of the hierarchy, G_{n-1} . If no such path exists then the algorithm returns a

¹Although not the focus of this paper, the optimizations presented in Section 3 are also applicable for improving the performance of Path Marking and Alternating Opportunism.

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.

```
CLASSIC REFINEMENT (A, s^0, g^0, n)
 1: s^{n-1} \leftarrow \text{LOOKUPVERTEXIMAGE}(s^0, n-1)
 2: g^{n-1} \leftarrow \text{LookupVertexImage}(g^0, n-1)
 3: P_{n-1} \leftarrow \text{FindPath}(s^{n-1}, g^{n-1}, n-1)
 4: if |P_{n-1}| = 0 then
          return NULL
 5:
 6: for i = n - 2 to i = 0
           P_i \leftarrow \{\}
 7:
           b \leftarrow \text{LOOKUPVERTEXIMAGE}(s^0, i)
 8:
           for j \leftarrow 0 to j = |P_{i+1}| - 1
 9:
                J \leftarrow \text{FindConstrainedJumpPath}(G_i, b, P_{i+1}[j+1])
10:
                P_i \leftarrow \operatorname{Append}(P_i, J)
11:
                b \leftarrow \text{LastVertex}(J)
12:
13:
           endfor
           g_i \leftarrow \text{LOOKUPVERTEXIMAGE}(g^0, i)
14:
15:
           C \leftarrow \text{FindConstrainedPath}(b, g_i, i)
16:
           P_i \leftarrow \text{Append}(P_i, C)
17: endfor
18: return P_0
```

Fig. 2. CR algorithm.

NULL path indicating that no path exists between s^0 and g^0 in G_0 (steps 1–5). If a path was found in G_{n-1} , CR iterates through each remaining level of abstraction, starting at G_{n-2} and going down the abstraction hierarchy to G_0 (for loop at step 6).

Let $P_{i+1} = \{s^{i+1}, v_1^{i+1}, \dots, v_{k-2}^{i+1}, g^{i+1}\}$ be the path found in G_{i+1} . In order to compute the path P_i , CR initializes b to the image of s^0 at abstraction level i. CR then computes the constrained jump path J from b to a vertex in the preimage of the next P_{i+1} vertex, $P_{i+1}[j+1]$ (step 10). By definition the last vertex in Jis the first vertex in the preimage of $P_{i+1}[j+1]$ visited by J. CR appends the constrained jump path J to P_i and updates b so that it is now the first vertex visited in $P_{i+1}[j+1]$ and iterates until the preimage of g^{i+1} is reached.

Finally, when b is the initial vertex in the preimage of g^{i+1} , CR computes a constrained path C between b and g^i , the image of g^0 in G_i (step 15) and appends C to P_i .

Once the algorithm reaches the bottom-most level of the abstraction and P_0 has been generated, the algorithm is finished and path P_0 is returned.

3. DATA STRUCTURE TRANSFORMATIONS

This section describes the three data structure transformations that we propose for improving the performance of CR. The techniques are characterized as follows: We use the graph shown in Figure 3 as an example in our presentation. We begin with a description of our baseline implementation.

3.1 Baseline

Our baseline implementation of CR is based on sound implementation techniques for sparse graph traversal algorithms. Because we are interested in graphs with a low degree of connectivity, we utilize adjacency lists to represent

6 R. Niewiadomski et al.



Fig. 4. Fields in the data structure of a vertex.

graphs. Figure 4 shows the 32-bit fields in the data structure used to represent a vertex v_a^i in the baseline implementation. ID is the vertex unique identification, the traversal visit marker (TVM) indicates if the vertex has been visited, BP is a back pointer used to store the previous node visited. BP is used to reconstruct the path once the goal node is encountered. IMG is a pointer to the vertex's image, and DEG is the degree of v_a^i . Figure 4 illustrates the data structure for v_0^0 in Figure 3.

The TVM is used to determine if the vertex has already been visited in the current search. The TVM of every vertex is initialized to zero before the first search takes place. A global search counter (GSC) is maintained. Before the first search takes place GSC is set to 1. Whenever vertex v_a^i is visited during the *z*th search, its TVM is set to GSC. When the *z*th search completes, GSC is incremented. Therefore, any vertex that has a TVM smaller than GSC during search z, has not been visited yet. An alternative design that would be very conservative in space usage would use a bit vector to represent the visited state, with one bit per vertex. Such a design has three disadvantages when compared with our TVM approach: (1) it requires bit vector manipulation, thus making the code more difficult to maintain; (2) it requires the reinitialization of the visited field before the start of each search; (3) the memory locations that contain the bits may be distant from the data structure that contains the vertex information in memory, and thus may not benefit from the spatial locality encountered when the TVM is embedded within the vertex structure. Because the TVM is a 32-bit integer, we do not have to reinitialize it unless we search more than four-billion times.

Through manual data placement, we ensure that all vertices in a given graph are stored in a single contiguous region of memory with one vertex data

A Performance Study of Data Layout Techniques



Fig. 5. Memory layout without vertex clustering.

structure immediately followed by another. Such an approach is advantageous compared to dynamically allocating memory on a per node basis. In addition to being faster, our manual data placement method ensures that the in-memory distance between vertices is consistent between different trials of the same experiment.

3.1.1 Constrained Path Finding Algorithm and Queue Representation. We use BFS to search for constrained paths and constrained jump paths. BFS stores vertices to be visited in a working queue. This queue is sometimes implemented as a circular buffer to save memory [Cormen et al. 1991]. However, we found that the overhead of checking for wraparound and overflow is high. Our approach eliminates this bookkeeping by allocating two buffers, each of which is large enough to accommodate a pointer to every vertex in the graph. During search we alternate between buffers, using one buffer to dequeue nodes discovered in the previous round of node expansion and the other buffer to queue up newly discovered nodes. Even though we could just as easily use only a single buffer to implement the queue for BFS, we found the two queues superior to a single queue. Using buffers instead of circular queues is a practice found in pathfinding engine implementations in video games [Freecraft; DeLoura 2000], as well as BFS-based implicit graph search [Korf 2003].

3.2 Vertex Clustering

Vertex Clustering is our data reordering technique. Consider the input graph shown in Figure 3. A naive implementation of the abstraction generation algorithm stores the vertices in memory in the order in which they appear in the original graph. Figure 5 shows the resulting layout of the vertex data structures in memory. In this figure, each small box represents a 32-bit memory field. For convenience of drawing, we present eight 32-bit fields per line. We

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.

7



Fig. 6. Memory layout with vertex clustering.

identify the 32-bit field where the data structure corresponding to each vertex of Figure 3 starts. Consider the search for a constrained jump path from v_0^1 to v_2^1 starting at v_0^0 and ending at v_3^0 in the example of Figure 3. The shaded areas in Figure 5 show the memory locations that are accessed in this search. The lighter shade denotes a single access while the darker shade marks a location accessed twice during the search. Besides the irregular memory access pattern shown in Figure 5, the baseline implementation also keeps a separate work queue, and therefore performs accesses to a separate region of memory that are interleaved with the accesses shown in Figure 5. These accesses not only do not benefit from spatial locality, they also are potential source of conflict misses in the data caches.

We use the technique of *vertex clustering* that consists of rearranging the vertex data structures in memory, such that vertices that map to the same image are located in close proximity of each other in memory. Clustering takes place after the abstraction hierarchy formation. To cluster vertices we simply place all vertices that have a common image into a contiguous memory region. Figure 6 shows a memory layout after one possible application of vertex clustering. The shaded areas in this figure are the memory locations that are accessed for the same constrained jump path search from v_0^1 to v_2^1 starting at v_0^0 and ending at v_3^0 . Notice how the memory accesses are much closer to each other in memory. Though not addressed in this paper, we expect the benefits of vertex clustering to be more pronounced in abstractions generated with a larger radius.

3.3 Image Mapping

Image mapping replicates data to improve spatial locality. With the data structure shown in Figure 4, pathfinding exhibits poor spatial locality even after vertex clustering has been applied. For instance, consider the constrained jump

A Performance Study of Data Layout Techniques •

ID	TVM	BP	IMG	DEG		Adjacency List						
\mathbf{v}_0^0	0	NULL	\mathbf{v}_0^{1}	3	\mathbf{v}_1^0	\mathbf{v}_1^1	\mathbf{v}_4^0	\mathbf{v}_0^1	\mathbf{v}_7^0	\mathbf{v}_{0}^{1}		

Fig. 7. Vertex data structure for abstract map.

ID	TVM	BP	EQP	IMG	DEG	Adjacency List					
\mathbf{v}_0^0	0	NULL	NULL	\mathbf{v}_0^1	3	\mathbf{v}_1^0	\mathbf{v}_1^1	v ⁰ ₄	\mathbf{v}_0^1	v ⁰ ₇	\mathbf{v}_0^1

Fig. 8. Vertex data structure for abstract map with embedded queue.

path example in Section 3.2. In order to constrain search to vertices that map to v_0^1 we need to access the IMG field of each vertex encountered during search. As a result, in Figure 6, we see that the baseline implementation accesses memory locations that are far from the locations where we find data structures for vertices in the preimage of v_0^1 .

The *image mapping* technique eliminates this unfavorable memory access pattern through an augmentation of the vertex data structure as shown in Figure 7. The change from Figure 4 is that the adjacency list contains not only a pointer to the neighboring vertices, but also the IMG field of each neighbor. Thus when finding paths we do not need to access remote memory locations to determine the image of a given vertex.

3.4 Embedded Queue

The next source of poor memory reference pattern is the working queue of BFS. This queue is likely to reside in a remote memory region. Thus, a constrained path search exhibits interleaved accesses to two separate regions of memory: (1) the graph vertex region and (2) the BFS working queue region. Frequent switching between two potentially distant memory regions has two adverse effects. First, it may cause *cache thrashing*, that is, entries that will be used later are discarded because of memory conflicts. Second, it may prevent the memory accesses from benefiting from the free prefetching that most cache memories offer in the form of large cache lines. In other words, because of poor spatial locality in the memory accesses, the algorithm may not benefit from implicit prefetching of the vertex pointers in the working queue.

The *embedded queue* technique keeps the information about vertices yet to be visited by BFS within the vertex's data structures. Embedding the queue merges two independent data structures, the graph and the queue, into a common memory area. To implement a BFS embedded queue, we augment the vertex data structure with an additional field, the embedded queue pointer (EQP), as shown in Figure 8. The EQP field of a vertex contains a pointer to the last vertex that was added to the working queue. We now present the modified constrained jump path search algorithm with these three modifications.

3.5 The Embedded Queue Constrained Jump Path Algorithm

An important component of CR is the algorithm that finds a constrained jump path from a start vertex *s* to any vertex in an image *I*. This algorithm assumes that if vertex *s* is in abstraction level *a*, then G_{a+1} has an edge between the image

9

```
EMBEDDEDQUEUECONSTRAINEDJUMPPATH(G(V, E), s, I)
 1: EQP(s) \leftarrow \text{NULL}
 2: w \leftarrow s
 3: w' \leftarrow \text{NULL}
 4: while TRUE
          while w \neq \text{NULL}
 5:
 6:
               for v \in V such that (w, v) \in E
 7:
                    if Image(v) = I
 8:
                        BP(v) \leftarrow w
 9:
                        return v
10:
                    if Image(v) \neq Image(s)
11:
                        continue
12:
                    if TVM(v) = GSC
13:
                        continue
14:
                    BP(v) \leftarrow w
                    EQP(v) \leftarrow w'
15:
16:
                    w' \leftarrow v
17:
                    TVM(v) \leftarrow GSC
18:
               endfor
19:
               w \leftarrow EQP(w)
20:
          endwhile
21:
          w \leftarrow w'
22:
          w' \leftarrow \text{NULL}
23: endwhile
```

Fig. 9. Embedded queue constrained path algorithm with abstract map.

of s and the vertex representing I. Figure 9 presents the modified constrained jump path algorithm that uses embedded queues to store the work queue.

The pattern of vertex visitation in BFS can be viewed as an expanding wave that starts at the initial vertex. If we divide this expansion into phases, in phase 0 we visit the starting vertex s, in phase 1 we visit all the immediate neighbors of s. In phase 2 we visit all the vertices that are two hops away from the starting vertex, and so on. The embedded queue algorithm uses w to access the linked list formed by the EQPs of the vertices that are being visited in the current phase. It uses w' to build the linked list of the vertices to be visited in the next phase.

When traversing a list in a given phase of the BFS algorithm, we use EQP to find the next vertex to be visited. In the initialization (steps 1–3), the EQP of the starting vertex s is assigned NULL to ensure that the phase 0 will terminate, and NULL is also assigned to w' to ensure that the next phase will also terminate. The first vertex of phase 0 is s. The algorithm terminates when a vertex whose image is I is encountered (step 7). The accesses to the neighbors of v in the for loop (step 6) benefit from spatial locality because they are stored in an adjacency list. Vertices that are not in the same image as the starting vertex (step 10) or that have already being visited (step 12) are not included in the working list for the next phase.

The image mapping technique ensures that the comparison between the image of v and the image of the starting vertex s (step 10) accesses data internal to the data structure of v and thus benefits from spatial locality. With the use

	Embedded	Vertex	Image
Implementation	Queues	Clustering	Mapping
Baseline			
Q	\checkmark		
-V-		\checkmark	
I			\checkmark
QV-	\checkmark	\checkmark	
Q-I	\checkmark		\checkmark
-VI		\checkmark	\checkmark
QVI	\checkmark	\checkmark	\checkmark

Table I. Techniques Featured in Each Implementation

of the embedded queue technique, the accesses to EQP (steps 15 and 19) are within the data structures of vertices v and w and thus also benefit from spatial locality.

The direction in which the embedded queue is constructed and traversed in the algorithm shown in Figure 9 matters. We build a backward queue in the sense that the newly discovered vertex v is placed at the front of w', not the rear. The advantage of this traversal direction is that when we finish building the queue, we start to visit vertices in the reverse order in which they were added to the queue. Thus, we are likely to visit vertices that we have recently visited and benefit from temporal locality in the cache memories.

4. EXPERIMENTAL FRAMEWORK

In order to test the effect of the techniques described in Section 3, we implemented eight versions of the constrained path finding algorithm.² We extensively tested the performance of these implementations on four different machines using two compilers on each system and four regular graphs as input. This section describes the algorithm implementations, the machines and compilers used, the input graphs, and the conditions under which the various experiments are performed.

4.1 Our Implementations

Using ANSI C we wrote eight CR implementations. Each implementation features a different mix of the techniques of queue embedding (Q), vertex clustering (V), and image mapping (I). The Baseline implementation features none of these techniques. Implementations are referred to by three character names, where each character indicates either the presence or absence of a given technique. For example, Q-I features embedded queues and image mapping, but not vertex clustering. Table I shows how the implementations are named.

4.2 Experiments and Hardware

We conducted experiments on four systems based on different processors: SGI (MIPS R12K), IBM (IBM Power3), AMD (AMD Athlon XP), and INTEL (Intel

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.

11

²The source code of our implementations along with sample input data and documentation can be found at the following URL: http://www.cs.ualberta.ca/~amaral/sources/JEA/.

]	Feature	SGI	IBM	AMD	INTEL
Processor	Туре	MIPS R12K	POWER3	2000+XP	Pentium 4
	Clock frequency	$350 \mathrm{~MHz}$	$450 \mathrm{~MHz}$	$1667 \mathrm{~MHz}$	$2260 \mathrm{~MHz}$
	Capacity	32 KB	64 KB	64 KB	8 KB
Data Cache L1	Associativity	2-way	128-way	2-way	4-way
	Line size	32 bytes	128 bytes	64 bytes	64 bytes
	Capacity	4 MB	8 MB	256 KB	512 KB
Data Cache L2	Associativity	2-way	4-way	16-way	8-way
	Line size	128 bytes	128 bytes	64 bytes	64 bytes
Data TLB L1	Capacity	56 entries	256 entries	32 entries	128 entries
	Associativity	Fully	2-way	Fully	Fully
Data TLB L2	Capacity	None	None	256 entries	None
	Associativity			4-way	
Virtual memory	y page size	16 KB	4 KB	4 KB	4 KB
Data TLB cover	rage	896 KB	1,024 KB	$1,152~\mathrm{KB}$	512 KB
TLB miss hand	ler	Software	Hardware	Hardware	Hardware
Main memory c	apacity	1 GB	1 GB	1 GB	1 GB
Compilors		MIPSpro (7.2.1)	IBM XLC (6.0)	Intel (6.0)	Intel (6.0)
Computers		GCC (2.7.2)	GCC (2.9)	$GCC\left(2.96 ight)$	GCC(2.96)
Hardware even	t measurement library	Perfex	PMAPI	PMC	PAPI

Table II. Systems and Compilers Used in the Experiments

Pentium 4). On each system we used GCC and the processor vendor's compiler to build our implementations with -00 and $-03.^3$ Unless specified otherwise, all results presented in this paper were obtained with -03 and a vendor-based compiler, except in the case of the SGI system where GCC was used. Detailed specifications of each system as well as information about compilers can be found in Table II.

All measurements encompassed the computation of 10,000 paths between random pairs of vertices.⁴ Computing each path includes: (1) searching for the path and (2) reconstructing the path into a linked list via a BP pointer traversal.⁵ Our main performance metric is execution time (measured as wallclock time). We also used hardware event counters to measure various processor and memory events.⁶ All execution times and hardware event totals were obtained on an unloaded machine. We always report the best result from several runs of each experiment. We always run our experiments as the single user in the machine and without other concurrent applications. In our experiments, we observed a variance between different runs of the same experiment that was consistently well below 1%. Thus choosing the best of several runs for each experiment eliminates spurious interactions with the operating system operation. We also collected traces of memory accesses made to heap memory associated with the abstraction hierarchy and, when applicable, with the queue vectors. This additional data was collected in separate runs to prevent interference with the time and hardware event measurements. This data enables us to construct

³On the AMD, we used the Intel C compiler as the processor vendor's compiler.

⁴We used a low-overhead portable and deterministic random number generator described in Press et al. [1992].

 $^{^5 \}rm Reconstruction$ of paths took between 1% and 17% of the execution time.

⁶We used hardware event measurement libraries (see Table II).

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.

various metrics. For instance, we computed the average number of distinct memory blocks accessed during search. By using block sizes corresponding to the cache line and page capacities found in our systems, and assuming minimal amounts of data reuse between successive searches, we can approximate the amount of traffic at each level of the memory hierarchy.

We tried to minimize code discrepancies between implementations. The main changes from the Baseline are as follows: (1) for vertex clustering the code stays the same; (2) for image mapping we slightly alter the lines of code that determine the image of a neighbor vertex; (3) for embedded queues, however, we had to replace the queue vector code with a pointer manipulating code that traverses the embedded queues.

4.3 Input Graphs

The core of our experiments was performed using four types of input graphs. We considered several instances of each graph type:

- 2D-Grid. A two-dimensional grid with a height of h and a width of w. A 2D-Grid has $h \times w$ vertices. Except for border vertices, a vertex located at (x, y) has edges to all vertices located at (x + a, y) and (x, y + b), where $a, b \in (-1, 1)$, thereby yielding an average vertex degree of approximately 4. We made h = w and varied its value from h = 256 to h = 1024 by increments of 128.
- 2DD-Grid. Same as 2D-Grid, but the average vertex degree is approximately 8 due to the introduction of diagonal edges to all vertices located at (x+a, y+b), where $a, b \in (-1, 1)$.
- 3D-Grid. A three-dimensional grid with height h, width w, and depth d. A 3D-grid has $h \times w \times d$ vertices. Except for border vertices, a vertex located at (x, y, z) has edges to all vertices located at (x + a, y, z), (x, y + b, z), and (x, y, z + c), where $a, b, c \in (-1, 1)$, thereby yielding an average vertex degree of approximately 6. We made h = w = d and varied h from s = 48 to s = 96 by increments of 8.
- 3DD-Grid. Same as 3D-Grid, but the average vertex degree is approximately 26 due to the introduction of diagonal edges. A vertex (x, y) has an edge to all vertices, except for itself, located at (x + a, y + b, z + c), where $a, b, c \in (-1, 0, 1)$.

Collectively, these graphs are representative of graphs with two-dimensional and three-dimensional topologies. For example, 2D-Grid and 2DD-Grid graphs are similar to graphs representing RTS game-worlds and road-maps. 3D-Grid and 3DD-Grid graphs on the other hand are similar to graphs used to represent three-dimensional objects, such as buildings and bridges. These grids are *empty* in the sense that they are not populated with obstacles as a grid in a computer game would be. In an empty grid the shortest path between any two points is always a straight line. However, our pathfinding algorithm does not use this knowledge. By using empty graphs to study our techniques, we avoid side effects caused by irregularities in the graphs. For completeness, in Section 5.6 we present a performance-oriented case study using nonempty graphs of realistic terrains.

Graph Type	2D-0	Grid	2DD-	-Grid	3D-	Grid	3DD-Grid		
Size	10	24	10	24		96	96		
Measure	V	E	V $ E $		V	E	V	E	
G_0	1,048,576	2,096,104	1,048,576	4,188,162	884,736	$2,\!626,\!560$	884,736	11,254,460	
G_1	196,779	559,645	116,964	465,806	116,784	644,850	32,768	398,908	
G_2	24,480	71,371	12,996	51,302	9,369	63,754	1331	14,230	
G_3	2973	8578	1444	5550	699	3375	64	468	
G_4	392	1,039	169	600	154	281	8	28	
G_5	68	137	25	72	68	67	1	0	
G_6	9	14	4	6	1	0	-	-	
G_7	3	2	1	0	-	-	-	-	
G_8	1	0	-	-	-	-	-	-	

Table III. Number of Vertices and Edges at Each Level of the Abstraction Hierarchy for the Largest Instance of Each Graph (A dash indicated the absence of a level.)

Table IV. Average Vertex Degree at Each Level of the Abstraction Hierarchy for the Largest Instance of Each Graph (A dash indicated the absence of a level.)

Graph Type	Size	G_0	G_1	G_2	G_3	G_4	G_5	G_6	G_7	G_8
2D-Grid	1024	4.0	5.7	5.8	5.8	5.3	4.0	3.1	1.3	0
2DD-Grid	1024	8.0	8.0	7.9	7.7	7.1	5.8	3.0	0	-
3D-Grid	96	5.9	11.0	13.6	9.7	3.6	2.0	0	-	-
3DD-Grid	96	25.4	24.3	21.4	14.6	7.0	0	1	-	-

All graphs are connected thereby enabling the search for a path between two randomly selected vertices. Abstraction hierarchies are generated with the STAR method with a radius of 2. The vertices of G_0 are placed in memory in the order in which they are read-in (see Section 5.1). The generation of G_{i+1} iterates through vertices in G_i in the order in which they appear in memory, selecting yet to be classified vertices as starting points for new subgraphs. Newly generated G_{i+1} vertices are placed in memory in the order in which they are created. The abstraction generation algorithm iterates until the resulting abstraction level has no edges. For vertex clustering, after the abstraction hierarchy is generated, vertices that belong to the same image are stored, one after another, in a contiguous memory location.⁷

Table III contains a summary of the abstraction hierarchy of the largest instance of each graph type in terms of vertex and edge totals present at each level of abstraction. Table IV does the same for the average vertex degree.

5. RESULTS AND ANALYSIS

Our experimental study findings can be summarized as follows:

-The combination of vertex clustering and image mapping (-VI) can reduce baseline execution times by upward of 43%. If vertices in the input graph are not ordered according to their vicinity, -VI can reduce baseline execution time by as much as 67%.

 $^{^7}$ The order in which vertices appear in the contiguous memory location reflects the order in which they were created during the abstraction generation process.

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.

- —The reduction in execution time is correlated with the reduction in TLB misses, and to a lesser degree to cache miss reductions. This observation supports the assertion that our techniques improve search performance through improved data reference locality.
- -The execution time improvements are robust to changes in the hardware architecture, compilers, and in the level of compiler optimization.
- —Although the addition of queue embedding to -VI produces performance improvements in some system and input graph combinations, indiscriminate use of the queue embedding technique is not advisable.
- --While some of our techniques increase the memory space required to store the abstraction hierarchy, they can also reduce the dynamic memory footprint (DMF) of search, thereby decreasing memory subsystem traffic as a whole, especially at the page level.
- —The results of a case study involving graphs representing realistic terrains suggest that our techniques are applicable to improving pathfinding performance in a variety of graphs.

5.1 Initial Vertex Order

The order in which vertices appear in the input graph can have a significant effect on the performance gains produced by our techniques. In our experiments, with the absence of vertex clustering, the order in which G_0 vertices appear in memory reflects the order in which they appear in the input graph. In addition, due to the nature of our abstraction generation process, the vertex order present in the input graph may also influence the order of vertices at higher levels of abstraction. Thus, the order in which vertices appear in the input graph can have an effect on data reference locality by affecting the in-memory proximity of vertices belonging to the same preimage. To study this effect, we define two types of input graph vertex ordering. In the *default vertex order* (DVO), vertices are ordered to capture the natural vertex order for the given graph type. The DVO of each graph type is as follows: 2D-Grid and 2DD-Grid vertices are ordered as left-to-right rows stacked on top of each other in a top-to-bottom manner, 3D-Grid vertices are organized as front-to-back ordered instances of 2D-Grid graphs, and 3DD-Grid graph vertices are organized as front-to-back instances of 2DD-Grid graphs. The randomized vertex order (RVO) of each graph type is generated by randomly scrambling the vertex order of DVO.

Performance gains achieved for DVO are indicative of the gains one can expect from the techniques presented in this paper. We use RVO to explore the upper bounds of the performance gains achievable with our techniques. Arbitrary vertex orderings such as RVO may seem artificial. After all, programmers do not tend to intentionally scramble the vertex order of their two-dimensional grids. However, arbitrary vertex orders can occur in practice. For example, consider a graph representing a national road system where vertices correspond to cities and junctions. If the vertices appear in the input graph in an order that reflects an alphabetical sort of their labels (that is, city names), their ordering might have little correlation with their geographic proximity. Previous work has shown that it is beneficial to place vertices in memory in an order based

			DV	0		RVO			
System	Version	2D	2DD	3D	3DD	2D	2DD	3D	3DD
	Q	3.0	-1.2	2.3	-16.4	-1.0	-1.8	2.0	0.6
	-V-	0.7	10.7	4.4	2.3	54.9	56.5	51.6	56.6
	I	-2.8	0.7	14.7	24.0	18.9	19.5	33.6	40.0
SGI	QV-	6.2	8.3	6.9	-13.3	52.9	54.0	49.8	51.3
	Q-I	-2.4	-6.0	11.3	0.7	17.1	17.5	31.2	38.0
	-VI	15.5	22.1	34.3	33.8	65.8	62.1	67.8	53.9
	QVI	12.2	8.8	30.2	5.1	62.6	58.5	64.9	57.2
	Q	3.4	3.8	-6.0	-9.2	-2.6	-2.3	-1.1	-1.0
	-V-	0.8	14.9	9.3	3.5	52.2	51.3	45.3	45.9
	I	-9.5	-9.4	14.8	27.5	9.1	9.8	25.4	30.9
IBM	QV-	6.3	13.8	10.2	-7.0	50.5	55.2	44.1	43.5
	Q-I	-4.9	0.1	17.1	14.8	7.3	11.0	23.9	30.5
	-VI	5.1	18.7	41.0	39.5	61.7	58.5	63.9	56.0
	QVI	16.3	22.7	43.8	26.9	62.6	61.3	65.2	56.7
	Q	1.2	-1.8	-1.0	-20.3	-2.9	-5.1	-6.3	-6.0
	-V-	4.4	10.0	11.4	5.0	42.8	46.3	34.9	48.2
	I	-4.1	-3.7	18.7	17.7	8.4	1.2	19.9	21.9
AMD	QV-	7.3	3.7	9.4	-14.0	40.4	41.5	31.0	30.7
	Q-I	-3.5	-7.2	17.0	-8.4	5.7	-3.2	16.3	16.8
	-VI	13.2	13.5	39.0	28.6	54.2	51.8	55.5	54.0
	QVI	16.2	7.6	37.2	3.3	52.5	47.8	52.6	47.1
	Q	-0.7	2.7	-3.4	-9.1	-4.9	-5.9	-5.5	-8.8
	-V-	2.0	17.3	9.9	7.9	48.7	48.6	41.9	40.9
	I	-7.2	-8.2	15.2	24.6	15.5	10.5	26.7	20.6
INTEL	QV-	4.9	15.0	7.6	-1.4	45.8	44.4	38.9	-8.8
	Q-I	-7.2	-4.1	12.9	11.6	12.9	4.2	23.0	13.9
	-VI	10.8	16.1	40.3	38.0	58.9	52.9	63.1	55.8
	QVI	12.5	17.3	37.6	24.6	57.0	50.6	61.3	47.3

Table V. Percentage Improvement Over Baseline for Seven Combinations of the Data Layout Techniques on Four Machines, Using the Largest Instance of Each Graph Type

The best improvement for each machine/graph combination is displayed in bold text.

on a sorting of their geographical locations [Edelkamp and Schrödl 2000]. A comparison of the RVO with the DVO performance gains in Table V confirms this observation. For the remainder of the experiments and analysis in this paper, we use DVO exclusively.

5.2 Execution Time

We begin by exploring how the performance of our techniques compares with the Baseline. Which one is the best performing implementation? From the results shown in Table V, with the exception of 2D-Grid graphs, -VI generally produces the best performance improvements over Baseline. In the case of 2D-Grid graphs however, QVI outperforms -VI.

Figures 10(a)-10(b) illustrate the performance gains achieved by QVI for 2D-Grid graphs and by -VI for 2DD-Grid, 3D-Grid, and 3DD-Grid graphs. These figures showcase the percentage reductions of execution time produced by the best performing implementation for each graph type in comparison with the Baseline implementation. Each figure presents results for all sizes of the given graph type using all four systems. The bars are composed of two segments



Fig. 10. Execution time for the best performing implementation for each graph type compared with the Baseline implementation. The top of each bar is annotated with the corresponding percentage reduction in execution time.

stacked on top of each other. The lighter segment denotes the execution time registered by QVI, in the case of Figure 10(a), and -VI in the case of Figures 10(b)–10(d), while the darker segment corresponds to the additional execution time required by Baseline. The top of each bar is annotated with the magnitude of the performance improvement. Performance gains vary from graph type to graph type. For 2D-Grid and 2DD-Grid graphs the performance gains range from 0.5% to 22.1%, while for 3D-Grid and 3DD-Grid graphs the performance gains are positive and generally consistent across hardware platforms for each graph type. This indicates that QVI and -VI are robust with respect to changes in the hardware architecture.

Are QVI and -VI robust to compiler changes as well? In short the answer is *Yes*. Figure 11 is a study of the percentage reduction in execution time produced by QVI or -VI over Baseline for the largest instance of each graph type. Results are presented for all systems and compilers, comparing two levels of optimization, -00 and -03. In all instances QVI and -VI produce nontrivial performance gains over Baseline. Because similar performance gains are obtained with -00 and -03, with both GCC and vendor compilers, QVI and -VI improve performance improvements seems to be orthogonal to compiler optimizations.

5.3 Data Reference Locality Improvement

Using hardware event counters for runs presented in Figures 10(a)-10(d), we examined the effectiveness of our best performing implementations in terms of their effect on data reference locality during search.



Fig. 11. A study of the percentage reduction in execution time produced by QVI or -VI over Baseline for the largest instance of each graph type. Results are presented for all systems and compilers, comparing two levels of optimization, -00 and -03.

Figures 12(a)-12(d) compare Baseline and the best performing implementation for some graph/system combinations in terms of the total number of L2 data cache and TLB misses. In each figure, we consider all graph sizes and we order the measurements, along the horizontal axis, according to the percentage reduction in execution time. All figures show a correlation between the reduction in TLB misses and the reduction in execution time. This result indicates that page-level data reference locality improves with our techniques and is responsible for some of the performance improvement.

Figures 12(a) and 12(b) show that data reference locality at the cache line level generally did not improve in the case of 2D-Grid and 2DD-Grid graphs.⁸ We suspect that because the abstraction hierarchy levels of two-dimensional grids have lower average vertex degrees (see Table IV) than three-dimensional grid abstraction hierarchy levels, there is less opportunity to eliminate secondary cache misses with respect to the baseline implementation. On the other hand, Figures 12(c) and 12(d) show significant improvements in data reference locality at the cache line level for 3D-Grid and 3DD-Grid graphs. Collectively, these figures present a nontrivial link between performance gains and data reference locality improvements, especially at the page level.

5.4 Should Queue Embedding be Used?

Table V shows that the combination -VI produces definite improvements on all the machines and graphs studied. The advantage of queue embedding, however, is questionable at best. Queue embedding is motivated by the following intuition: placing the information stored in the queue in the same memory region as the vertex information should prevent frequent switches between two distant memory regions and thus improve performance. However, the execution time improvements presented in Table V indicate that queue embedding is not advantageous for several graphs and systems. Figure 13 underscores this point. This graph shows the percentage change, in the execution time, when queue

 $^{^8 \}rm Occasionally,$ the number of L2 cache misses increases in experiments involving 2D-Grid and 2DD-Grid graphs.

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.



Fig. 12. Study of the correlation between execution time reduction, reduction in TLB misses, and changes in L2 cache misses for select implementations.

embedding is added to a -VI implementation. In many instances adding queue embedding to -VI significantly hurts performance. As this result goes against our original intuition, we want to investigate why this is the case. Some immediate observations are: (1) Q is detrimental to the 3DD-Grid on all systems; (2) for the remaining graphs, the influence of Q on performance is very different on the IBM than on the other systems.

• R. Niewiadomski et al.



Fig. 13. Percentage reduction in execution time when Q is added to -VI.

Table VI. Misses in Data Cache (DC), Secondary Cache (DSC), and TLB, and Graduated Instructions (all Measured in Millions) for -VI and the Percentile Change for QVI

		2D-Grid		2DD-Grid		3D-Grid		3DD-Grid	
			QVI		QVI		QVI		QVI
System	Metric	-VI	Change	-VI	Change	-VI	Change	-VI	Change
aat	DC Misses	80.3	10.0%	66.0	23.9%	22.3	14.7%	35.9	39.3%
	DSC Misses	12.1	40.2%	10.9	68.6%	4.6	28.3%	9.6	58.2%
DGI	TLB Misses	7.2	-6.6%	5.6	-15.0%	1.7	-5.5%	0.87	-10.9%
	Instructions	2,782	-16.2%	$2,\!480$	1.9%	908	-8.4%	1,667	23.9%
	DC Misses	48.6	-44.0%	36.7	-32.1%	7.0	8.0%	13.3	9.8%
TDM	DSC Misses	7.3	9.2%	7.1	9.8%	3.3	4.2%	3.3	4.2%
TBW	TLB Misses	10.3	-38.4%	6.8	-42.5%	1.3	1.7%	1.0	0.0%
	Instructions	$2,\!257$	-7.3%	1,775	9.9%	639	-1.1%	1,012	28.9%

Table VI shows the variation in primary (DC) and secondary (DSC) cache misses, TLB misses, and number of instructions graduated for the IBM and SGI systems when Q is added to -VI. The number of TLB misses drops significantly when queue embedding is used, confirming our speculation of higher locality at the page level when the data accesses are not switched frequently from the area where the vertex's data are stored to the separate area where the queue is stored. However, embedding the queue results in a much higher rate of cache misses both for the primary and secondary caches. Our best explanation for this effect is that queue embedding accesses suffer from the "pointer chasing" problems, that is, a field in the current vertex must be accessed before the next element of the queue is known. In general, compilers have a hard time optimizing code with this kind of data access pattern [Chilimbi and Hirzel 2002; Stoutchinin et al. 2001]. On the other hand, a vector-based queue lends itself well to optimizations both by the compiler and by the underlying hardware. For example, nonbinding prefetch instructions can be issued by the hardware and/or inserted by the compiler. Also vector-based queue benefits from wide cache lines because a single cache line contains several queue entries. In contrast, with embedded queues a single cache line is not likely to contain more than



Fig. 14. Abstraction hierarchy memory footprints of all implementations for all instances of each graph type.

one or two queue entries. Queue embedding also results in a larger number of instructions graduated for graphs with a high degree of connectivity (high ratio of edges per vertex) such as 2DD-Grid and 3DD-Grid. In such graphs, the working queue is much larger than in 2D-Grid and 3D-Grid. This empirical evidence is indicative of the code generator having less success with the linked-list nature of the queue embedded code.

We also experimented with comparisons between compilation at levels -00 and -03 on the SGI machine and discovered that the compiler eliminates approximately 10% more stores when operating at a higher level of optimization for -VI than for QVI. This is further evidence that the code generated for QVI is more difficult to optimize.

In summary, we do not recommend the use of queue embedding for graphs that have high connectivity because of the difficulty that compilers have optimizing the code. This lack of efficient optimization manifests itself in higher primary and secondary cache misses, and a larger number of instructions graduated.

5.5 Memory Space Requirements and Memory Footprints

With the exception of vertex clustering, the techniques described in this paper increase the memory space required to store the data for a given abstraction hierarchy because the data structure for each individual vertex is larger.

Figure 14 shows the growth in the memory footprint of the abstraction hierarchy for all combinations of graph instance and implementation used in our experiments.⁹ Each bar is composed of four segments stacked on top of one

⁹We refer to the amount of space required to store the abstraction hierarchy as its memory footprint.

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.



Fig. 15. The percentage reduction produced by each implementation over Baseline in the dynamic memory footprint. Results are presented for all graph types using two block sizes, capturing typical cache line and page capacities.

another showing the abstraction hierarchy memory footprint of all implementations for a given graph instance. On average, queue embedding increased the memory footprint by 10.7% for 2D-Grid graphs, 7.7% for 2DD-Grid graphs, 8.7% for 3D-Grid graphs, and 3.2% for 3DD-Grid graphs. For image mapping, average increases were 46.2% for 2D-Grid graphs, 61.4% for 2DD-Grid graphs, 56.7% for 3D-Grid graphs, and 83.4% for 3DD-Grid graphs. In all instances, the abstraction hierarchy memory footprint is well below the main memory capacity of all of our test systems. Thus, we observed virtually no page faults.

In contemporary architectures with virtual memory management, the memory space allocated for data storage is not necessarily a major concern. Instead, an algorithm designer seeking to improve performance should primarily be concerned with the amount of memory accessed during program execution, that is, the DMF. For example, relational database systems have long used data redundancy to improve query performance. Although the memory footprint of the database system may increase, the amount of memory accessed to execute a query decreases. With that in mind, we measured the DMF as the average number of distinct memory blocks referenced during a single search.¹⁰ The DMF measure encompasses stores and loads made to the memory region containing the abstraction hierarchy, and, when applicable, the queue vectors. In measuring the DMF of each implementation we used two block sizes, 64 and 4096 bytes, capturing typical cache line and page capacities. Assuming a negligible amount of data reuse from one search to the next, we can use DMF to gage the amount of memory subsystem traffic generated by each implementation, both at the cache and at the page level.

Figures 15(a) and 15(b) present a study of the effect of our techniques on reducing the DMF using block sizes of 64 and 4096 bytes, respectively. The figures compare the DMF of each implementation to the DMF of Baseline for the largest instance of each graph type. Each bar represents the percentage reduction in the DMF. For example, -11.1% for the 2D-Grid in Figure 15(a)

 $^{^{10} \}rm DMF \ is \ somewhat \ akin \ to \ the \ I/O \ measure \ commonly \ used \ in \ external \ memory \ algorithm \ analysis.$

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.



Fig. 16. Overhead snapshots of 5 of the 16 terrain maps used in our nonempty graph performance study. Each snapshot is annotated with the grade of obstacle coarseness in the corresponding map.

means that QV- references 11.1% more distinct 64-byte memory blocks than Baseline.

At the cache line level, Figure 15(a), image mapping is the most effective technique in reducing the DMF. When used in isolation, techniques other than image mapping cause the DMF to increase. When image mapping is combined with the other techniques, the DMF improvement is greater than the sum of the DMF improvements of each technique. This result indicates that our techniques complement each other.

At the page level, Figure 15(b), vertex clustering produces the largest reductions in DMF. Image mapping also yields DMF improvements, although only for three-dimensional grids. The use of queue embedding makes very little difference on the DMF at the page level, although it appears to be somewhat more effective than at the cache line level. Combining techniques appears to have a similar effect at the page level as it does at the cache line level.

Overall, our techniques yield significant decreases in the DMF of search. The largest improvements result from combining techniques and are generally better at the page level than at the cache line level. In addition, two-dimensional grids tend to benefit more than three-dimensional grids, especially at the cache line level.

5.6 Case Study

So far our investigation used only empty graphs. What about nonempty graphs? Do our techniques improve search performance for nonempty graphs? To shed light on this issue, we present a case study involving a class of nonempty graphs. In particular, we examine the performance of our implementations using fractal-generated two-dimensional terrain maps.

Researchers in the field of *mobile robotics* often utilize fractal-generated terrains to model both terrestrial and extraterrestrial terrains [Singh et al. 2000; Yahja et al. 1998]. We generated sixteen 1024×1024 fractal-based terrain *maps*. Approximately 60% of the points in each map correspond to free space, with the remaining portion representing obstacles. The maps vary in terms of the coarseness of terrain obstacles, ranging from a grade of 4 to a grade of 64 in increments of four.¹¹ Figure 16 presents overhead snapshots of five of our maps with increasing grades of obstacle coarseness. Each map

¹¹Coarseness grade is roughly equivalent to the maximum diameter of an obstacle.

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.

was converted into two explicit graph instances, a 2D form and a 2DD form. A 2D instance is akin to a 2D-Grid in that it features only perpendicular edges, while a 2DD instance is similar to a 2DD-Grid because it has both perpendicular and diagonal edges. Because the maps are essentially two-dimensional grids, we used an input graph vertex order similar to the DVO of our 2D-Grid and 2DD-Grid graphs. We note, however, that unlike our empty graphs, our fractal-generated graphs are not connected. As a result, a path between two randomly selected vertices in the graph is not always possible. The manner in which we generate the abstraction hierarchy lends itself to detecting whether or not two vertices are connected (see Section 4.3). In particular, to check whether or not s^0 is connected to g^0 we merely check if $s^{n-1} = g^{n-1}$; if $s^{n-1} \neq g^{n-1}$ then s^0 is not connected to g^{0} .¹²

The experimental results show our techniques yielding nontrivial performance gains for both the 2D and 2DD graph instances of the fractal-generated terrain maps. In general, QVI was the best performing implementation on the IBM while -VI was superior to other implementations on the remaining systems. Figure 17 shows the percentage reduction in execution time achieved over Baseline by QVI, in the case of the IBM, and -VI, in the case of the remaining systems. Results are presented for all 2D instances in Figure 17(a), and for all 2DD instances in Figure 17(b). Both figures show performance gains on each system being relatively consistent across all grades of obstacle coarseness. The average performance gain on each system varies between 10.6% and 19.6% for 2D instances, and between 5.8% and 15.2% for 2DD instances. How do these gains compare to those achieved for the empty 2D-Grid and 2DD-Grid graphs? As a basis of comparison, we compare the results from our experiments involving the 768 two-dimensional grids, since the 768 grids feature approximately the same number of vertices as our maps. On average, the performance gains achieved for fractal instances were 11.9% larger than the ones attained for the 768 2D-Grid graph. In the case of 2DD instances, however, performance gains were typically 32.2% smaller than in the case of the 768 2DD-Grid graph. These experimental results with nonempty graphs indicate that our techniques improve path finding in irregular graphs.

6. RELATED WORK

This paper presents an extensive performance evaluation of the techniques presented in Section 3. These techniques were first described in Niewiadomski et al. [2003]. To our knowledge there is no previous work specifically addressing the locality of abstraction search algorithms such as CR. Nonetheless, we find research addressing graph search locality in general. For instance, Edelkamp and Schrödl address the problem of thrashing of pages at the virtual memory level [Edelkamp and Meyer 2001; Edelkamp and Schrödl 2000]. They apply their localized A^{*} to a route planning system. In a nutshell, their strategy

 $^{^{12}}$ In our code we generate a pair of randomly selected vertices. If the two vertices are not connected we do not try to find a path between them and move onto generating another pair of randomly selected vertices. We do this a total of 10,000 times. As such, in our experiments involving the fractal-generated graphs the total number of paths actually found is less than 10,000.

ACM Journal of Experimental Algorithmics, Vol. 9, Article No. 1.2, 2004.



% Reduction in Execution Time



Fig. 17. Summary of the performance gains attained by the best performing implementations on each system for both the 2D and 2DD graph instances of our fractal-generated terrain maps. For each machine and implementation the bars correspond to the percentage reductions of Baseline execution times. Each bar corresponds to a given graph instance, with the number at the bottom of the bar denoting the grade of obstacle coarseness in the corresponding map.

involves improving reference locality by sorting vertices based on their relative geographic locations and altering the order in which states are expanded during search.

In the field of *external memory* algorithms, we find various techniques for improving the I/O efficiency of external memory graph search [Agarwal et al. 1998; Arge et al. 2000; Chiang et al. 1995; Meyer et al. 2003; Nodine et al. 1993; Vitter 2001]. Typically, external memory algorithm techniques are akin to vertex clustering (grouping) and image mapping (data redundancy). For instance, blocking of data is used to minimize the number of page faults incurred during the traversal of paths in planar graphs. Variants of vertex grouping are also used to increase the performance of sparse matrix multiplication [Gibbs et al. 1976; Pinar and Heath 1999].

Graph partitioning, needed for abstraction generation, is a well-studied problem [Daz et al. 2002; Karypis et al. 1997; Patel and Cote 1981]. For example, consider the problem of partitioning a graph into k subgraphs such that each subgraph has roughly the same number of vertices and the number of intersubgraph edges is minimized. The ability to partition graphs in this manner could enable the targeting of specific page and/or cache-line capacities during the abstraction generation process. Although, such an approach could lead to enhanced data reference locality, its effect on path quality is unclear to us.

Also of interest are the ideas explored in the realm of *cache oblivious* algorithms [Arge et al. 2002; Brodal and Fagerberg 2003; Frigo et al. 1999]. In general, the aim of the cache oblivious paradigm is to improve the data access locality of algorithms independent of memory hierarchy parameters. However, we are yet to come across a cache oblivious approach that enhances the locality of general sparse graph search. Korf finds that when searching an implicit graph with BFS, compared to checking for duplicates as you go, sorting the working queue before its expansion permits the elimination of duplicates, leading to better performance and possibly fewer data cache misses [Korf 2003]. It remains to be seen if this approach can be effective for explicit graph BFS involving graphs such as two- and three-dimensional grids.

The Artificial Intelligence community focuses on reducing the search space [Russell 1992], to produce improvements of orders of magnitude. The gains obtained with the data structure transformation oriented techniques presented in this paper are orthogonal to the search space reduction, and the two techniques can be easily combined. These techniques are also orthogonal to performance improvements obtained through compiler transformations that improve data placement [Calder et al. 1998; Chilimbi et al. 1999]. However, the automated techniques found in contemporary compilers are quite inept at improving data locality with respect to graph search in general. Even with the ongoing development of profile-oriented compilation, we expect this to continue to be the case because techniques such as our embedded queue and image mapping methods not only require a change in the manner data is laid out in memory but also require changes to the search algorithms themselves.

Our techniques are particularly efficient at improving page-level locality as indicated by the large TLB miss reductions. Increasing the page size is a general-purpose approach for improving page-level locality [Winwood et al. 2002]. However, there are numerous drawbacks and implementation details that still need to be addressed. For instance, page swap time and fragmentation become bigger issues as page size grows. In terms of use of large page sizes for RBS pathfinding codes, we believe that it may be warranted in scientific computations but generally undesirable for interactive desktop applications such as commercial video games.

7. CONCLUSION

Research in the AI and computer game communities has produced algorithms to quickly find short paths in very large sparse graphs. However, the effects of temporal and spatial locality in the implementation of these algorithms have been mostly overlooked. This paper demonstrates that three simple data structure transformation techniques can consistently improve the performance of CR pathfinding for sparse graphs. In our experiments, these techniques improved data reference locality resulting in performance improvements of up to 67% with consistent improvements above 15%. In addition, these techniques appear to be orthogonal to compiler optimizations and robust with respect to hardware architecture.

ACKNOWLEDGMENTS

This research is partially funded by grants from the Natural Sciences and Engineering Research Council (NSERC) of Canada through its Collaborative Research and Development (CRD) Grants program, by IBM Canada through the Center for Advanced Studies (CAS) and the IBM Toronto Laboratory, by the Canadian Foundation for Innovation (CFI), and by the Alberta Ingenuity Fund.

REFERENCES

- AGARWAL, P. K., ARGE, L., MURALI, T. M., VARADARAJAN, K. R., AND VITTER, J. S. 1998. I/O-efficient algorithms for contour-line extraction and planar graph blocking. In *Ninth Symposium on Discrete Algorithms (SODA)*. 117–126.
- ARGE, L., BENDER, M. A., DEMAINE, E. D., HOLLAND-MINKLEY, B., AND MUNRO, J. I. 2002. Cacheoblivious priority queue and graph algorithm applications. In *Annual ACM Symposium on Theory* of Computing. 268–276.
- ARGE, L., BRODAL, G. S., AND TOMA, L. 2000. On external memory MST, SSSP and multi-way planar graph separation. Available at citeseer.nj.nec.com/293177.html.
- BRODAL, G. S. AND FAGERBERG, R. 2003. On the limits of cache-obliviousness. In ACM Symposium on Theory of Computing. 307–315.
- CALDER, B., CHANDRA, K., JOHN, S., AND AUSTIN, T. 1998. Cache-conscious data placement. In Architectural Support for Programming Languages and Operating Systems (ASPLOS-VIII). 139–149.
- CHIANG, Y.-J., GOODRICH, M. T., GROVE, E. F., TAMASSIA, R., VENGROFF, D. E., AND VITTER, J. S. 1995. External-memory graph algorithms. In *Symposium on Discrete Algorithms*. 139–149.
- CHILIMBI, T. M. AND HIRZEL, M. 2002. Dynamic hot data stream prefetching for general-purpose programs. In Programming Language Design and Implementation (PLDI), Berlin, Germany. 199–209.
- CHILIMBI, T. M., HILL, M. D., AND LARUS, J. R. 1999. Cache-conscious structure layout. In Programming Language Design and Implementation (PLDI), Atlanta, GA. 1–12.
- CORMEN, T. H., LEISERSON, C. E., AND RIVEST, R. L. 1991. Introduction to Algorithms. The MIT Press.
- CORRESPONDENCE WITH DAVID C. POTTINGER OF ENSEMBLE STUDIOS.
- DELOURA, M., ED. 2000. Game Programming Gems, vol. 1. Charles River Media.
- DAZ, J., PETIT, J., AND SERNA, M. 2002. A survey of graph layout problems. ACM Computing Surveys (CSUR) 34, 3, 313–356.
- EDELKAMP, S. AND MEYER, U. 2001. Theory and practice of time-space trade-offs in memory limited search. In *Lecture Notes in Computer Science*, vol. 2174. 169–184.
- EDELKAMP, S. AND SCHRÖDL, S. 2000. Localizing A*. In Seventeenth National Conference on Artificial Intelligence (AAAI-2000). 885–890.
- Freecraft real-time strategy gaming engine. Available at http://www.freecraft.net.
- FRIGO, M., LEISERSON, C. E., PROKOP, H., AND RAMACHANDRAN, S. 1999. Cache-oblivious algorithms. In 40th Annual Symposium on Foundations of Computer Science. IEEE, 285–298.
- GIBBS, N. E., W. G. POOLE, J., AND STOCKMEYER, P. K. 1976. A comparison of several bandwidth and profile reduction algorithms. ACM Trans. Math. Softw. 2, 4, 322–330.
- HOLTE, R. C., MKADMI, T., ZIMMER, R. M., AND MACDONALD, A. J. 1996. Speeding up problem solving by abstraction: A graph oriented approach. *Artif. Intell.* 85, 1–2, 321–361.
- KARYPIS, G., AGGARWAL, R., KUMAR, V., AND SHEKHAR, S. 1997. Multilevel hypergraph partitioning: application in VLSI domain. In *Thirty-Fourth Annual Conference on Design Automation Confer*ence. 526–529.
- KORF, R. 2003. Delayed duplicate detection: Extended abstract. In International Joint Conference on Artificial Intelligence. 1539–1541.
- MEYER, U., SANDERS, P., AND SIBEYN, J. F., EDS. 2003. Algorithms for Memory Hierarchies, Advanced Lectures [Dagstuhl Research Seminar, March 10–14, 2002]. Lecture Notes in Computer Science, vol. 2625. Springer, Berlin.

- NIEWIADOMSKI, R., AMARAL, J. N., AND HOLTE, R. C. 2003. Crafting data structures: A study of reference locality in refinement-based path finding. In *International Conference on High Performance Computing (HiPC)*, Hyderabad, India.
- NODINE, M. H., GOODRICH, M. T., AND VITTER, J. S. 1993. Blocking for external graph searching. In Symposium on Principles of Database Systems (PODS 93). 222–232.
- PATEL, A. M. AND COTE, L. C. 1981. Partitioning for VLSI placement problems. In Eighteenth Design Automation Conference. 411–418.
- PINAR, A. AND HEATH, M. T. 1999. Improving performance of sparse matrix-vector multiplication. In ACM/IEEE Conference on Supercomputing.
- POTTINGER, D. C. 2000. Terrain analysis in realtime strategy games. In *Game Developers Conference Proceedings*.
- PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., AND FLANNERY, B. P. 1992. Numerical Recipes in C: The Art of Scientific Computing, 2nd ed. Cambridge University Press. 284.
- RUSSELL, S. J. 1992. Efficient memory-bounded search methods. In 10th European Conference on Artificial Intelligence (ECAI 92), Vienna, Austria. 1–5.
- SINGH, S., SIMMONS, R., SMITH, T., STENTZ, A., VERMA, V., YAHJA, A., AND SCHWEHR, K. 2000. Recent progress in local and global traversability for planetary rovers. In *IEEE International Conference on Robotics and Automation*.
- STOUTCHININ, A., AMARAL, J. N., GAO, G. R., DEHNERT, J., JAIN, S., AND DOUILLET, A. 2001. Speculative pointer prefetching of induction pointers. In Compiler Construction 2001—European Joint Conferences on Theory and Practice of Software, Genova, Italy. 289–303.
- VITTER, J. S. 2001. External memory algorithms and data structures: dealing with massive data. *ACM Comput. Surv.* 33, 2, 209–271.
- WINWOOD, S., SHUF, Y., AND FRANKE, H. 2002. Multiple page size support in the Linux kernel. In *Proceedings of the Ottawa Linux Symposium*. 573–593.
- YAHJA, A., STENTZ, A., SINGH, S., AND BRUMITT, B. L. 1998. Framed-quadtree path planning for mobile robots operating in sparse environments. In *IEEE International Conference on Robotics* and Automation.

Received November 2003; accepted June 2004