

Dynamic State-Space Partitioning in External-Memory Graph Search

Rong Zhou

Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
rzhou@parc.com

Eric A. Hansen

Dept. of Computer Science and Eng.
Mississippi State University
Mississippi State, MS 39762
hansen@cse.msstate.edu

Abstract

State-of-the-art external-memory graph search algorithms rely on a hash function, or equivalently, a state-space projection function, that partitions the stored nodes of the state-space search graph into groups of nodes that are stored as separate files on disk. The scalability and efficiency of the search depends on properties of the partition: whether the number of unique nodes in a file always fits in RAM, the number of files into which the nodes of the state-space graph are partitioned, and how well the partitioning of the state space captures local structure in the graph. All previous work relies on a static partitioning of the state space. In this paper, we introduce a method for dynamic partitioning of the state-space search graph and show that it leads to substantial improvement of search performance.

1 Introduction

Recently-developed algorithms for external-memory graph search – including hash-based delayed duplicate detection [Korf and Schultze, 2005; Korf, 2008] and structured duplicate detection [Zhou and Hansen, 2004; 2006b; 2007] – rely on a hash function, or equivalently, a state-space projection function, that partitions the nodes of the state-space search graph into *buckets* of nodes that are stored as separate files on disk. For both hash-based delayed duplicate detection and structured duplicate detection, the state-space projection function must satisfy the same criteria. The set of unique nodes in each bucket must fit in RAM. For best performance, there should not be many more buckets than necessary, which means nodes should be relatively evenly distributed among buckets and buckets should be relatively full. Finally, search efficiency depends on how well the state-space projection function captures local structure in the graph, where locality takes this form; for any bucket of nodes, their successor nodes are found in only a small number of other buckets. Hash-based delayed duplicate detection leverages this form of local structure in order to *interleave expansion and merging* of nodes and structured duplicate detection uses it in its concept of *duplicate detection scope*.

Finding a projection function that satisfies all of these criteria presents a challenge. Korf relies on handcrafted projection functions (which he calls hash functions) that are tailored to specific search domains with well-understood structure. Zhou and Hansen [2006b] describe how to automatically generate an appropriate projection function by heuristic-guided greedy search through the space of possibilities. All previous work uses a static projection function that does not change during the progress of the search, but this has drawbacks. Static partitioning can capture local graph structure. But for a static partition that captures local structure, it is difficult to predict in advance the number of nodes that will map to each bucket; in practice, the distribution of nodes to buckets can be very uneven. A randomized hash function creates a static partition that evenly distributes nodes among buckets. But it does not capture any local structure, since it allows nodes in one bucket to have successor nodes in any other random bucket. In short, it is difficult to design a projection function that both captures local structure and evenly distributes nodes among buckets.

In this paper, we describe an approach to improving the performance of structured duplicate detection – and, by implication, also hash-based delayed duplicate detection – by dynamically adjusting the projection function in the course of the search. This allows the search algorithm to monitor the distribution of nodes in buckets at runtime, and modify the projection function to improve search performance. In case the set of nodes in a bucket does not fit in RAM, for example, the algorithm changes the partition so that the search can continue. We show that the overhead for dynamically re-partitioning the state space is modest, and, in practice, it is more than compensated for by an improvement in the space and time complexity of the search.

2 Background and motivation

We begin with a review and comparison of hash-based delayed duplicate detection and structured duplicate detection, and focus on the role played by the state-space projection function in each approach. A projection function is used to partition the stored nodes of a state-space search graph into *buckets*, where each bucket of nodes is stored in a separate file. A common way to specify a projection function is by specifying a subset of state variables; states that have the same values for this subset of variables map to the same

bucket. A projection function captures local structure in a graph if for any bucket of nodes, successor nodes are found in only a small number of other buckets, called its *neighbors*. In other words, it captures local structure when the largest number of neighbors of any bucket is very small relative to the total number of buckets.

2.1 Hash-based delayed duplicate detection

Korf introduced hash-based delayed duplicate detection (DDD) to avoid the overhead of sorting-based DDD, which relies on external sorting of files to detect and remove duplicate nodes.

Hash-based DDD uses two hash functions, where the first corresponds to what we call a state-space projection function. As new states are generated they are placed into separate files based on this first hash function. This guarantees that all duplicate nodes end up in the same file. To remove duplicates from a file, the file is read into a hash table that fits in RAM and is associated with a second hash function. Duplicate nodes that map to the same slot of the hash table are “merged.” Finally, the contents of this hash table are written back to disk as a file of unique nodes.

Consider a breadth-first search algorithm that expands one level of the search space at a time. The files at the current level of the search space contain no duplicate nodes, and are called *expansion files*. As the nodes in each of these files are expanded, their successor nodes are written to files at the next depth of the search space, with the value of the first hash function determining which file they are written to. Since these files contain duplicate nodes that will be removed in the merge phase of the algorithm, they are called *merge files*. Because expanding all files at the current level before merging any files at the next level could require a large amount of extra disk space to store all of the duplicate nodes, Korf [2005] proposes to *interleave expansion and merging*. This is possible only if the projection function captures local structure in the state-space search graph; that is, it depends on the nodes in each expansion file having successor nodes in only a small number of merge files. As soon as all the expansion files that have successors in a particular merge file have been expanded, duplicates can be removed from the merge file (by copying the file into a hash table in RAM) even if all expansion files in the current level of the search space have not yet been expanded. To save disk space, merging duplicates in a file is given priority over expanding another file.

2.2 Structured duplicate detection

The idea of leveraging local graph structure to interleave the expansion and merging steps of hash-based DDD was adapted from an approach to external-memory graph search called structured duplicate detection (SDD) [Zhou and Hansen, 2004]. Here we describe a form of SDD that uses a technique called *edge partitioning* [Zhou and Hansen, 2007]. Comparing hash-based DDD to this form of SDD will help to clarify how the two approaches are related. It will also show why the state-space projection function used by both approaches must satisfy the same criteria.

Like hash-based DDD, SDD uses two “hash functions.” It partitions the nodes in each level of the search space into sep-

arate files based on the first hash function, that is, the projection function. For each file, it expands the nodes contained in the file and writes the successor nodes to files in the next level of the search space. The key difference between SDD and hash-based DDD is that SDD does not delay duplicate detection; all duplicate nodes are detected and eliminated as soon as they are generated, without writing any duplicate node to disk. In its original form, SDD accomplishes this by copying the *duplicate-detection scope* of the currently-expanding file into a hash table in RAM associated with the second hash function. The duplicate-detection scope consists of all nodes in any of the neighbor files of the expanding file. Since this requires the nodes in multiple files to fit in RAM at once, this form of SDD has a larger internal-memory requirement than hash-based DDD, *if they both use the same projection function*. (Recall that hash-based DDD never requires more than one file of unique nodes to fit in RAM at once.)

Edge partitioning reduces the internal-memory requirement of SDD in the following way. If the duplicate-detection scope of an expansion file does not fit in RAM, then one or more of its neighbor files are not copied into RAM. Instead, when expanding the nodes in the expansion file, the successor nodes that map to one of these neighbor files are simply not generated. After every node in the expansion file is expanded, the ignored neighbor file(s) are copied into RAM, replacing the neighbor files that were previously in RAM, and the nodes in the same parent file are expanded again. This time, the successor nodes that mapped to one of the neighbor files that is no longer in RAM are not generated, and only the successor nodes that map to the neighbor file(s) in RAM are generated and saved, as long as they are not duplicates. Thus, by *incrementally* expanding the nodes in a file, the internal-memory requirements of SDD can be reduced to the point where no more than one (neighbor) file needs to be stored in RAM at once, the same as for hash-based DDD. The time and space overhead of writing all duplicate nodes to disk and removing them in a later merge step is avoided. But whenever a duplicate-detection scope does not fit in RAM, SDD with edge partitioning incurs the different time overhead of having to read the same expansion file multiple times.

2.3 Comparison of approaches

Assuming that hash-based DDD and SDD with edge partitioning use the same projection function, it is clear that they have the same peak RAM requirement: it is the amount of RAM required to store all the unique nodes in the largest file.

As for disk storage, hash-based DDD always needs at least as much disk space as SDD, since both approaches store all unique nodes. In addition, hash-based DDD needs disk space to store duplicate nodes. How much additional disk storage it needs depends on how well the projection function captures local structure in the state-space search graph. If local structure is leveraged to allow interleaving of expansion and merging, it may need very little additional disk space. In the worst case, when merging of files must be postponed until all files in the current layer are expanded, it could require much more disk storage than SDD, by a factor equal to the ratio of duplicate nodes to unique nodes. Note that if the projection function does not capture any local structure, both the inter-

nal memory requirement and the disk storage requirement of SDD with edge partitioning remain the same; only its time complexity increases due to incremental node expansions and multiple reads of the same expansion file.

Comparing the time complexity of hash-based DDD and SDD is more challenging, but we can make some general remarks about their relative advantages and disadvantages. Both approaches perform extra work that is not performed by the other approach, and that is what we compare. For hash-based DDD, the extra work is writing all duplicate nodes to disk and then eliminating them in a later merge step that copies the nodes in each merge file back to RAM, eliminates duplicates, and writes an expansion file of unique nodes. For SDD with edge partitioning, the extra work consists of incremental node expansions and reading the same file from disk multiple times. The extra work performed by hash-based DDD is proportional to the ratio of duplicate nodes to unique nodes in the search space, which is problem-dependent. The extra work performed by SDD depends on how much local structure is captured by the projection function, since the number of times a file may need to be read from disk is bounded above by the number of its neighbor files.

Hash-based DDD may have an advantage in time complexity when there are few duplicates relative to unique nodes in the search space. (An example of such a problem would be the Rubik's Cube search problem used a test case by Korf [2008].) SDD may have an advantage when the ratio of duplicates to unique nodes is large, and the projection function captures local structure in the state-space search graph. The best approach is likely to be problem-dependent. The point of our comparison is to show that both approaches rely on a state-space projection function that must satisfy the same criteria. It follows that the method for dynamic state-space partitioning introduced in this paper can be effective for both approaches. Our experimental results will demonstrate the effectiveness of dynamic partitioning for SDD.

2.4 Criteria of a good projection function

As we have seen, one criterion the projection function should satisfy is that it should capture local structure in the state-space search graph. For hash-based DDD, this allows interleaving of expansion and merging. For SDD with edge partitioning, it limits the time overhead of incremental expansions.

A more important criterion the projection function should satisfy is that the set of unique nodes in each file must fit in a hash table in RAM. One way to ensure this is to use a high-resolution projection function that partitions the nodes of the state-space search graph into so many files that each is guaranteed to fit in RAM. But partitioning the state space into too many files can degrade search performance. Besides decreasing the average size of a file, the typically uneven distribution of nodes among files means that many files could be empty or nearly empty, making access to disk less sequential. In addition, the search algorithm needs to maintain a table in RAM that keeps track of all files, whether they are open or not, whether they have a write or read buffer, their status as an expansion or merge file (in hash-based DDD), a list of their neighbor files, etc., and the size of this table can grow exponentially with the resolution of the projection function.

Refining the partition also tends to increase the number of neighbors of a file. In hash-based DDD, this means that more file buffers must be maintained to allow generated nodes to be written to their corresponding file; since the operating system limits the number of open files, this could be a problem. For SDD, it could lead to more incremental expansions. In general, increasing the resolution of the projection function reduces the peak RAM consumption of the search algorithm in exchange for an increase in its running time, for all of the reasons we have mentioned. It is a classic space-time trade-off.

Thus we have two competing criteria. The set of unique nodes in each file must fit in RAM but this criterion should be satisfied while also allowing as coarse a partition as possible, in order not to degrade search performance too much. Given an uneven distribution of nodes among files, however, a coarse partition increases the risk that the set of unique nodes in a particular file may not fit in RAM. As we will see, dynamic state-space partitioning allows us to manage this trade-off more effectively. It will let us find the coarsest partition that still allows the largest file to fit in RAM.

Although this issue arises in general for hash-based DDD, it does not arise for two test domains for which Korf reports most of his experimental results: sliding-tile puzzles and the four-peg Towers of Hanoi. For these domains, Korf uses handcrafted hash functions that are perfect and invertible, allowing the use of direct-address tables in memory that just need to store a few bits of information for each entry, instead of the entire state description. This allows him to partition the nodes of the state-space search graph into sufficiently many files that the maximum number of unique nodes in each file is guaranteed to fit in the direct-address table in RAM.

But in general, perfect *and* invertible hash functions are not possible; for example, they are not feasible for either domain-independent planning or model checking. (Edelkamp and Sulewski [2008] make this point about model checking. An application of hash-based DDD to model checking is described by Evangelista [2008].) Instead, the hash table must store the complete state description with each entry. Since it is usually unrealistic for the projection function to partition the nodes of the search graph into so many files that it is possible to guarantee that the set of all possible nodes that map to a bucket can fit in RAM at once, an open-address hash table is used that allows collisions. The hope is that the actual nodes in a bucket at any point during the search will fit in RAM. The fact that the number of nodes that will be generated and stored in any file is not known until run time motivates a dynamic approach to state-space partitioning.

A final remark about Korf's [2008] examples is relevant to the question of the distribution of nodes among buckets. Most of his examples involve exhaustive breadth-first search of a graph. By contrast, the test examples reported in this paper involve heuristic search for the shortest path from a start state to a goal state, which means that most of the graph is not explored. In our experience, the distribution of nodes among buckets is more uneven for a heuristic search than an exhaustive search due to the uneven effect of reachability on the contents of the various buckets.

2.5 Pathological state-space projection functions

To highlight the influence of heuristic search on the distribution of nodes among buckets, we use the 15-Puzzle example shown in Figure 1. Suppose the state-space projection function hashes a node to a bucket based on the position of tiles 3, 7, 11, 12, 13, 14, and 15 (shown as gray tiles). Since there are altogether $16!/9! = 57,657,600$ different combinations for the positions of these 7 gray tiles, each bucket should get a fraction of $57,657,600^{-1} = 1.73 \times 10^{-8}$ of the total number of nodes generated, if the distribution of nodes are perfectly balanced among all buckets. However, for the start state shown in Figure 1, all the gray tiles are already at their goal positions. Thus an optimal solution does not need to move these gray tiles far away from their current positions. In fact, since the white tiles in Figure 1 form a solvable instance of (a variant of) the 8-Puzzle, we know that there is at least one optimal solution that does not require moving a single gray tile.

Since all states that share the same positions of these gray tiles are mapped to the same bucket, the bucket with all the gray tiles located at their goal positions would get the majority of nodes. This is because the heuristic biases the search to move the white but not the gray tiles. Unfortunately, this means that almost every node generated is mapped to the same bucket. We refer to projection functions that create highly imbalanced buckets as *pathological state-space projection functions*. In this example, an external-memory search algorithm that uses either hash-based delayed or structured duplicate detection would not save any RAM, and could potentially use *more* RAM, because the abstract state-space graph could have orders of magnitude more abstract nodes than the number of search nodes expanded in solving the problem. This example also shows that increasing the resolution of the state-space projection function is not guaranteed to work in heuristic search. If the projection function used is pathological, then creating more buckets does not necessarily reduce the size of the largest bucket, which determines the peak RAM requirements in both delayed and structured duplicate detection.

Note that the same state-space projection function would have been perfectly fine if it were used inside a brute-force breadth-first search algorithm, because, in the absence of any search bias, the search is just as likely to move the gray tiles as it is to move the white tiles, resulting in a more balanced distribution of nodes among all the buckets.

3 Dynamic state-space partitioning

Zhou and Hansen [2006b] describe an automatic state-space partitioning algorithm for a domain-independent STRIPS planner that uses external-memory graph search with SDD. The projection function that partitions the state space is defined by selecting a subset of state variables. Beginning with the null set of variables, the algorithm performs greedy search in the space of projection functions by selecting at each step a multi-valued variable (or a related group of Boolean variables) that maximizes the locality of the partition, where locality is defined as the largest number of neighbors of any bucket divided by the total number of buckets. As they point




Figure 1: An example of a pathological state-space projection function for an instance of the 15-Puzzle. The gray tiles are the tiles whose positions are considered in the state-space projection function.

out, this measure of locality is a good objective function for the greedy search under the assumption that the projection function evenly partitions the stored nodes of the graph. But as we have argued (and experimental results will show) this assumption is an over-idealization; in practice, the distribution of nodes among buckets can be very uneven.

In this paper, we improve on this static approach by introducing a dynamic partitioning algorithm that monitors the distribution of nodes among buckets in the course of the search and modifies the projection function to adapt to the distribution. The dynamic partitioning algorithm searches for a projection function that both captures local structure *and* keeps the size of the largest bucket of nodes as small as possible – in particular, small enough to fit in RAM. The algorithm we describe is simple and could be improved in obvious ways, but it is sufficient to show the effectiveness of the approach.

Like the static partitioning algorithm, the dynamic partitioning algorithm is greedy and adds a new state variable to the projection function each iteration. In the initial iteration, no state variables have been selected and the partition consists of a single bucket for all nodes. For each bucket in the partition, it keeps a vector of counters, one for each state variable that has not yet been selected. It scans all nodes on disk and computes values for the counters, as follows. As it reads each node, it maps it to one of the buckets of the partition created in the previous iteration. Then, for each state variable that has not yet been selected and the corresponding bucket in a refined partition, it determines whether the node maps to this potential bucket. If so, it increments the corresponding counter. At the end of the iteration, the algorithm selects the state variable which results in the greatest reduction in the size of the largest bucket (and also captures locality in the state-space graph) and adds it to the projection function. This refines the partition. Then the process repeats, with new counters. The algorithm terminates when either the size of the largest bucket is below some threshold or the maximum number of buckets has been reached. It checks whether the partition found by the dynamic algorithm is significantly better than the partition used to organize the current set of files. If so, it creates a new set of files based on the new partition and copies the nodes on disk to the new files. To save disk space, a file that corresponds to an old bucket is deleted immediately after all of its nodes are moved to their new buckets. Thus, dynamically changing the partition does not change the peak

Vars	Values	Nodes
$\{X\}$	$\{X = 1\}$	$\{a, b\}$
	$\{X = 2\}$	$\{c, d\}$
	$\{X = 3\}$	$\{e, f\}$
$\{Y\}$	$\{Y = 4\}$	$\{a, c, e\}$
	$\{Y = 5\}$	$\{b, d, f\}$
$\{Z\}$	$\{Z = 6\}$	$\{a, b, c, d\}$
	$\{Z = 7\}$	$\{e, f\}$

Table 1: First iteration of dynamic partitioning.

Vars	Values	Nodes
$\{X, Y\}$	$\{X = 1, Y = 4\}$	$\{a\}$
	$\{X = 1, Y = 5\}$	$\{b\}$
	$\{X = 2, Y = 4\}$	$\{c\}$
	$\{X = 2, Y = 5\}$	$\{d\}$
	$\{X = 3, Y = 4\}$	$\{e\}$
	$\{X = 3, Y = 5\}$	$\{f\}$
$\{X, Z\}$	$\{X = 1, Z = 6\}$	$\{a, b\}$
	$\{X = 1, Z = 7\}$	\emptyset
	$\{X = 2, Z = 6\}$	$\{c, d\}$
	$\{X = 2, Z = 7\}$	\emptyset
	$\{X = 3, Z = 6\}$	\emptyset
	$\{X = 3, Z = 7\}$	$\{e, f\}$

Table 2: Second iteration of dynamic partitioning.

disk space requirements of the search algorithm; its only effect is to reduce the peak RAM requirements.

Example An example illustrates how the greedy dynamic partitioning algorithm works. Suppose a search problem has three state variables: $X \in \{1, 2, 3\}$, $Y \in \{4, 5\}$, and $Z \in \{6, 7\}$. The algorithm has generated and stored 6 states (encoded in $\langle X, Y, Z \rangle$ format): $a = \langle 1, 4, 6 \rangle$, $b = \langle 1, 5, 6 \rangle$, $c = \langle 2, 4, 6 \rangle$, $d = \langle 2, 5, 6 \rangle$, $e = \langle 3, 4, 7 \rangle$, and $f = \langle 3, 5, 7 \rangle$. Tables 1 and 2 show the first two iterations of the algorithm. In both tables, the “Vars” column shows the set of state variables being considered for the projection function, the “Values” column shows the assignment of values to the state variables, and the “Nodes” column shows the set of stored nodes whose state encoding matches the corresponding “Values” column. In the first iteration, only three singleton state-variable sets, $\{X\}$, $\{Y\}$, and $\{Z\}$, are considered. The largest bucket size, as a result of using a single state variable, is 2 for X , 3 for Y , and 4 for Z . Thus, at the end of the first iteration, the state variable X is chosen for the projection function. Since there are only two variables Y and Z left, the second iteration only has two candidates – the variable sets $\{X, Y\}$ and $\{X, Z\}$ – for the refined projection function. Clearly, the variable set $\{X, Y\}$ should be chosen, because it reduces the largest bucket size to one, achieving a perfect balance across all buckets.

In our initial description of the algorithm, we assumed that all nodes are stored on disk when the dynamic partitioning algorithm is invoked. In fact, some could be stored only in RAM and we need to handle the case where changing the partition requires moving nodes from RAM to disk, or vice

versa. The following four cases need to be handled correctly: (1) moving a RAM node to a RAM bucket, (2) moving a disk node to a disk bucket, (3) moving a RAM node to a disk bucket, and (4) moving a disk node to a RAM bucket. Furthermore, one or more RAM buckets may need to be flushed to disk, if internal memory is exhausted in the middle of moving nodes to their new buckets. Thus, in the fourth case above, the algorithm needs to make sure there is space in RAM to hold a node read from disk; if not, all the nodes in the new bucket need to be written to disk. The procedure for handling the second case is then invoked, because the new bucket is no longer in RAM. To save RAM, once a bucket is flushed to disk, it is never read back into RAM until the search resumes.

Excessive overhead for dynamic partitioning is avoided in a couple of ways. First, since repartitioning the state space involves the time-consuming process of moving nodes on disk and creating new files, it is done only when the dynamic partitioning algorithm finds a significantly better partition. (In our implementation, the reduction in the largest bucket size must be greater than 10%.) Second, dynamic partitioning is invoked only if there is a significant imbalance in bucket sizes *and* the largest bucket consumes a substantial fraction of available RAM. (In our implementation, the ratio of largest bucket size to average bucket size must be greater than three *and* the largest bucket size must be greater than half of available RAM; these choices could be tuned to improve performance.) How often dynamic partitioning is invoked is problem-dependent. If a good partition is found early in the search, it may not need to be changed. This can be viewed as choosing a partition based on a sampling of the search space.

4 Experimental results

We implemented our dynamic state-space partitioning algorithm inside an external-memory domain-independent STRIPS planner that uses as its underlying graph-search algorithm *breadth-first heuristic search* [Zhou and Hansen,

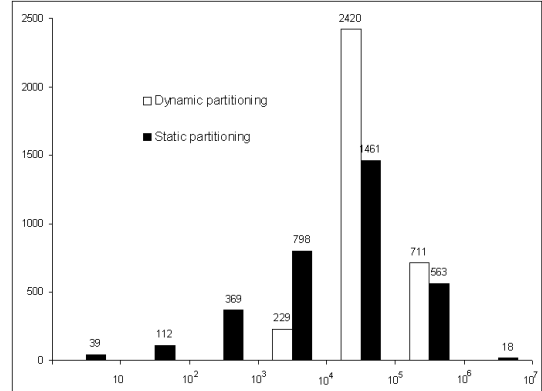


Figure 2: Distribution of nodes among buckets using static and dynamic partitioning for Korf’s 15-puzzle problem instance #88. The x-axis is bucket size in number of nodes and the y-axis is count of buckets that have a size that falls in one of the following 7 ranges; $[0, 10]$, $(10, 10^2]$, $(10^2, 10^3]$, $(10^3, 10^4]$, $(10^4, 10^5]$, $(10^5, 10^6]$, $(10^6, 10^7]$.

#	Len	Static partitioning				Dynamic partitioning			
		RAM	Disk	Increm Exp	Secs	RAM	Disk	Increm Exp	Secs
17	66	908,902	50,871,643	711,180,658	589	149,054	51,443,638	718,502,872	854
49	59	927,906	80,987,861	812,948,341	697	105,021	81,209,329	817,962,744	1,021
53	64	244,889	48,518,100	592,797,672	511	102,365	48,650,054	593,080,899	690
56	55	498,854	49,436,882	477,575,355	424	95,883	49,570,631	480,107,665	522
59	57	957,496	52,834,528	531,743,811	484	169,941	53,504,361	539,097,048	543
60	66	867,509	218,185,611	2,582,825,054	2,184	680,840	218,181,871	2,566,169,284	2,568
66	61	309,651	81,919,509	920,508,447	793	124,377	82,084,656	922,526,667	1,014
82	62	385,486	177,927,698	1,865,565,899	1,582	245,695	177,963,645	1,859,645,376	2,115
88	65	1,776,317	295,406,768	3,357,109,415	2,923	349,901	296,507,472	3,354,622,475	3,234
92	57	324,196	48,085,400	512,701,523	450	71,998	48,130,370	511,378,919	702

Table 3: Comparison of edge partitioning with and without dynamic state-space partitioning on the 10 hardest of Korf’s 100 15-Puzzle instances encoded as STRIPS planning problems. The number of buckets in the partition is the same for both static and dynamic partitioning. Columns show solution length (Len), peak number of nodes stored in RAM (RAM), peak number of nodes stored on disk (Disk), number of incremental node expansions (Exp), and running time in CPU seconds (Secs).

Problem	Len	Static partitioning					Dynamic partitioning			
		Disk	RAM	Increm Exp	Secs	Buckets	RAM	Increm Exp	Secs	Buckets
blocks-14	38	381,319	37,129	10,763,944	21	2,660	8,637	13,738,732	40	2,644
gripper-7	47	2,792,790	13,000	177,532,311	506	3,726	14,999	169,318,244	298	2,568
freecell-3	18	4,279,315	151,546	107,699,115	284	1,764	35,247	101,070,199	327	1,764
depots-7	21	12,877,783	410,815	184,606,201	300	4,240	32,000	255,291,983	584	4,240
driverlog-11	19	15,780,803	89,999	233,976,409	305	3,848	75,000	299,947,271	414	2,752
gripper-8	53	14,099,800	59,999	894,274,064	1,427	4,212	50,000	857,433,260	1,210	3,210
depots-13	25	1,110,708	81,003	27,711,837	25	625	14,653	25,411,325	38	700
driverlog-14	28	26,356,967	911,288	664,087,448	775	784	472,011	544,381,999	810	616
logistics-10	42	81,728,366	8,505,120	1,434,271,308	3,267	2,744	448,343	1,961,380,522	3,959	2,940

Table 4: Comparison of edge partitioning with and without dynamic state-space partitioning on STRIPS planning domains. Columns show solution length (Len), peak number of nodes stored on disk (Disk), peak number of nodes stored in RAM (RAM), number of incremental node expansions (Exp), running time in CPU seconds (Secs), and the peak number of buckets. For the first six problems, the planner used the max-pair heuristic; for the last three, it used a more accurate pattern database.

2006a]. Experiments were performed on a machine with Intel Xeon 3.0 GHz processors, 8 GB of RAM and 6 MB of L2 cache. No parallel processing was used in the experiments.

We first tested the performance of the external-memory STRIPS planner on the 15-Puzzle with results shown in Table 3. Although state-of-the-art domain-specific solvers can find optimal solutions to randomly generated instances of the 24-Puzzle and some easy instances of the 35-Puzzle, the 15-Puzzle remains a challenge for domain-independent planners. Using an advanced disjoint pattern database heuristic, the best domain-independent solver [Haslum *et al.*, 2007] can solve only 93 of Korf’s 100 15-Puzzle instances. Here we show that with a very basic pattern database heuristic (equivalent to the Manhattan-Distance heuristic for domain-specific solvers), our external-memory STRIPS planner is able to solve the entire set with ease. For example, the most difficult instance (#88) can be solved by our planner in less than an hour, storing less than 350 thousand nodes in RAM and less than 300 million nodes on disk. At the size of 36 bytes per RAM node and 14 bytes per disk node, it turns out the most difficult instances of the 15-Puzzle can be solved using less than 20 megabytes of RAM. Without external-memory search, it would take roughly 10.8 gigabytes of RAM just to store the nodes, even though the underlying breadth-first heuristic search algorithm only stores 30% of the nodes ex-

panded by A* [Zhou and Hansen, 2006a].

For all instances in Table 3 and for both static and dynamic partitioning, the number of buckets in the partition is 3360. Dynamic partitioning uses a slight amount of extra disk space for moving nodes on disk when the partition is changed dynamically. Some of the extra time overhead for dynamic partitioning is due to the partitioning algorithm itself. Most is due to extra file I/O as a result of using less RAM. Finally, Figure 2 provides additional insight into why the approach is effective. With dynamic partitioning, the distribution of stored nodes among files is more concentrated around an average file size than with static partitioning, which has both larger files and empty files.

We also tested the external-memory planner on domains from the biennial planning competition. The results are shown in Table 4. A couple interesting observations can be made. First, peak RAM consumption for problems such as *depots-7* and *logistics-10* is substantially reduced. Previously, the only way to reduce peak RAM consumption was to use a more fine-grained projection function that creates more buckets. But note that in some cases, the reductions we achieved in peak RAM consumption is a result of using a *coarser* partition with fewer buckets. This shows that the resolution of the projection function, while important, is not the only factor that determines the amount of RAM saved in

external-memory search. Our results show that there can be a large difference in peak RAM consumption among projection functions that have the same resolution. This illustrates the benefit of dynamic partitioning based on monitoring the actual distribution of nodes among buckets. The only problem instance for which dynamic partitioning uses more RAM than static partitioning is `gripper-7`. In this case, we intentionally forced the dynamic partitioning algorithm to use a coarser projection function in order to see if it could still find a good partition. Note that with fewer buckets for `gripper-7`, the external-memory search algorithm runs 70% faster in return for a 15% increase in peak RAM consumption.

It is also worth noting that the columns labeled “RAM” in Tables 3 and 4, which show the peak RAM nodes for structured duplicate detection, also show the peak RAM nodes for hash-based delayed duplicate detection if it uses the same underlying search algorithm (e.g., breadth-first heuristic search) and the same static or dynamic partitioning method. Thus, from the viewpoint of reducing the peak RAM requirements, dynamic partitioning can improve SDD and DDD *equally*.

5 Conclusion

We have introduced an approach to dynamic state-space partitioning in external-memory graph search that substantially reduces peak RAM consumption in exchange for a modest increase in running time. It can also reduce running time while leaving peak RAM consumption roughly the same. It achieves improved performance and a more favorable time-memory tradeoff than static partitioning because the partition is adapted to the actual distribution of stored nodes. Although the timing results of our experiments are for structured duplicate detection only, the approach could be used to reduce the peak RAM requirements of both SDD and hash-based DDD.

References

- [Edelkamp and Sulewski, 2008] S. Edelkamp and D. Sulewski. Model checking via delayed duplicate detection on the GPU. Technical report, University of Dortmund, 2008.
- [Evangelista, 2008] S. Evangelista. Dynamic delayed duplicate detection for external memory model checking. In *Proc. of the 15th Int. SPIN workshop*, pages 77–94, 2008.
- [Haslum *et al.*, 2007] P. Haslum, M. Helmert, B. Bonet, A. Botea, and S. Koenig. Domain-independent construction of pattern database heuristics for cost-optimal planning. In *Proceedings of the 22nd Conference on Artificial Intelligence (AAAI-07)*, pages 1007–1012, 2007.
- [Korf and Schultze, 2005] R. Korf and P. Schultze. Large-scale parallel breadth-first search. In *Proc. of the 20th National Conference on Artificial Intelligence (AAAI-05)*, pages 1380–1385, 2005.
- [Korf, 2008] R. Korf. Linear-time disk-based implicit graph search. *Journal of the ACM*, 35(6), 2008.
- [Zhou and Hansen, 2004] R. Zhou and Eric Hansen. Structured duplicate detection in external-memory graph search. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI-04)*, pages 683–688, 2004.
- [Zhou and Hansen, 2006a] R. Zhou and E. Hansen. Breadth-first heuristic search. *Artificial Intelligence*, 170(4-5):385–408, 2006.
- [Zhou and Hansen, 2006b] R. Zhou and E. Hansen. Domain-independent structured duplicate detection. In *Proc. of the 21st National Conference on Artificial Intelligence (AAAI-06)*, pages 1082–1087, 2006.
- [Zhou and Hansen, 2007] R. Zhou and Eric Hansen. Edge partitioning in external-memory graph search. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 2410–2416, 2007.