

Implementing Minimum Cycle Basis Algorithms

KURT MEHLHORN and DIMITRIOS MICHAÏL
Max-Planck-Institut für Informatik

In this paper, we consider the problem of computing a minimum cycle basis of an undirected graph $G = (V, E)$ with n vertices and m edges. We describe an efficient implementation of an $O(m^3 + mn^2 \log n)$ algorithm. For sparse graphs, this is the currently best-known algorithm. This algorithm's running time can be partitioned into two parts with time $O(m^3)$ and $O(m^2n + mn^2 \log n)$, respectively. Our experimental findings imply that for random graphs the true bottleneck of a sophisticated implementation is the $O(m^2n + mn^2 \log n)$ part. A straightforward implementation would require $\Omega(nm)$ shortest-path computations. Thus, we develop several heuristics in order to get a practical algorithm. Our experiments show that in random graphs our techniques result in a significant speed-up. Based on our experimental observations, we combine the two fundamentally different approaches to compute a minimum cycle basis to obtain a new hybrid algorithm with running time $O(m^2n^2)$. The hybrid algorithm is very efficient, in practice, for random dense unweighted graphs. Finally, we compare these two algorithms with a number of previous implementations for finding a minimum cycle basis of an undirected graph.

Categories and Subject Descriptors: G.2.2 [**Discrete Mathematics**]: Graph Theory—*Graph algorithms*

General Terms: Algorithms, Design, Experimentation

Additional Key Words and Phrases: Cycle basis, graph algorithms

1. INTRODUCTION

1.1 Preliminaries

Let $G = (V, E)$ be an undirected graph. A *cycle* of G is any subgraph in which each vertex has an even degree. Associated with each cycle is an *incidence vector* x , indexed on E , where $x_e = 1$, if e is an edge of C , $x_e = 0$ otherwise. The vector space over $GF(2)$ generated by the incidence vectors of cycles is called the *cycle*

A preliminary version of this paper appeared in the Proc. 4th International Workshop on Efficient and Experimental Algorithms (WEA), 2005 [Mehlhorn and Michail 2005]

Authors' addresses: Kurt Mehlhorn and Dimitrios Michail, Max-Planck-Institut für Informatik, Algorithms and Complexity Group, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany; email: {mehlhorn, michail}@mpi-inf.mpg.de

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., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org. © 2006 ACM 1084-6654/2006/0001-ART2.5 \$5.00 DOI 10.1145/1187436.1216582 http://doi.acm.org/10.1145/1187436.1216582

space of G . It is well known that this vector space has dimension $N = m - n + \kappa$, where m is the number of edges, n is the number of vertices, and κ the number of connected components of G . A maximal set of linearly independent cycles is called a *cycle basis*.

The edges of G have non negative weights. The weight of a cycle is the sum of the weights of its edges. The weight of a cycle basis is the sum of the weights of its cycles. We consider the problem of computing a cycle basis of minimum weight in an undirected graph. (We use the abbreviation MCB to refer to a minimum cycle basis.)

The problem has been extensively studied, both in its general setting and in special classes of graphs. Its importance lies in its use as a preprocessing step in several algorithms. Such algorithms include diverse applications like electrical circuit theory [Chua and Chen 1973], structural engineering [Cassell et al. 1976], and surface reconstruction [Tewari et al. 2005; Gotsman et al. 2005].

The first polynomial time algorithm for the minimum cycle basis problem was given by Horton [1987] with running time $O(m^3n)$. de Pina [1995] gave an $O(m^3 + mn^2 \log n)$ algorithm by using a different approach. Golynski and Horton [2002] improved Horton's algorithm to $O(m^\omega n)$ by using fast matrix multiplication. It is presently known [Coppersmith and Winograd 1990] that $\omega < 2.376$. Recently, Berger et al. [2004] gave another $O(m^3 + mn^2 \log n)$ algorithm by using similar ideas as de Pina. Finally, Kavitha et al. [2004] improved de Pina's algorithm to $O(m^2n + mn^2 \log n)$ again by using fast matrix multiplication. In the same paper, a faster $1 + \epsilon$ approximation algorithm, for any $\epsilon > 0$, is presented.

In this paper, we report our experimental findings from our implementation of the $O(m^3 + mn^2 \log n)$ algorithm presented in de Pina [1995]. Our implementation uses LEDA [Mehlhorn and Naher 1999]. We develop a set of heuristics, which improve the best-case performance of the algorithm while maintaining its asymptotics. Finally, we consider a hybrid algorithm obtained by combining the two different approaches used in de Pina [1995], Kavitha et al. [2004], Horton [1987], and Golynski and Horton [2002], with running time $O(m^2n^2)$. We then compared the implementations. The new algorithm is motivated by our need to reduce the cost of the shortest-path computations. The resulting algorithm seems to be very efficient in practice for random, dense unweighted graphs. Finally, we compare our implementations with previous implementations of minimum cycle basis algorithms [Huber 2002; Kreisbasenbibliothek 2004].

The paper is organized as follows. In Section 2, we briefly describe the algorithms. In Section 2.1, we describe our heuristics and in Section 2.2 we present our new algorithm. In Section 3, we present and discuss our experimental results. We summarize and offer conclusions in Section 4.

2. ALGORITHMS

We briefly describe the algorithms. Let $G(V, E)$ be an undirected graph with m edges and n vertices. Let $l : E \mapsto \mathbb{R}_{\geq 0}$ be a nonnegative length function on

Algorithm 1. Construct an MCB

```

Set  $S_i = \{e_i\}$  for all  $i = 1, \dots, N$ 
for  $i = 1$  to  $N$  do
  Find  $C_i$  as the shortest cycle in  $G$  s.t  $\langle C_i, S_i \rangle = 1$ 
  for  $j = i + 1$  to  $N$  do
    if  $\langle C_i, S_j \rangle = 1$  then
       $S_j = S_j \oplus S_i$ 
    end
  end
end

```

the edges. Let κ be the number of connected components of G and let T be any spanning forest of G . Also, let e_1, \dots, e_N be the edges of $G \setminus T$ in some arbitrary, but fixed, order. Note that $N = m - n + \kappa$ is exactly the dimension of the cycle space.

The algorithm [de Pina 1995] computes the cycles of an MCB and their *witnesses*. A witness S of a cycle C is a subset of $\{e_1, \dots, e_N\}$, which will prove that C belongs to the MCB [de Pina 1995; Kavitha et al. 2004]. We view these subsets in terms of their incidence vectors over $\{e_1, \dots, e_m\}$. Hence, both cycles and witnesses are vectors in the space $\{0, 1\}^m$. $\langle C, S \rangle$ stands for the standard inner product of vectors C and S . Since we are in the field $GF(2)$, observe that $\langle C, S \rangle = 1$ if, and only if, the cardinality of the intersection of the two edge sets is odd. Finally, adding two vectors C and S in $GF(2)$ is the same as the symmetric difference of the two edge sets. Algorithm 1 gives a full description, where \oplus denotes the symmetric difference.

The algorithm in phase i has two parts, one is the computation of the cycle C_i and the second part is the update of the sets S_j for $j > i$. Note that updating the sets S_j for $j > i$ is nothing more than maintaining a basis $\{S_{i+1}, \dots, S_N\}$ of the subspace orthogonal to $\{C_1, \dots, C_i\}$.

2.1 Computing the Cycles

Given S_i , it is easy to compute a shortest-cycle C_i such that $\langle C_i, S_i \rangle = 1$ by reducing it to n shortest-path computations in an appropriate graph G_i . The following construction is well known.

G_i has two copies v^+ and v^- of each vertex $v \in V$. For each edge $e = (u, v) \in E$ do: if $e \notin S_i$, then add edges (u^+, v^+) and (u^-, v^-) to the edge set of G_i and assign their weights to be the same as e . If $e \in S_i$, then add edges (u^+, v^-) and (u^-, v^+) to the edge set of G_i and assign their weights to be the same as e . G_i can be visualized as two levels of G (the $+$ level and the $-$ level). Within each level, we have edges of $E \setminus S_i$. Between the levels, we have the edges of S_i . Call G_i the *signed* graph.

Any v^+ to v^- path p in G_i corresponds to a cycle in G by identifying edges in G_i with their corresponding edges in G . If an edge $e \in G$ occurs multiple times, we include it if the number of occurrences of e modulo 2 is 1. Because we identify v^+ and v^- with v , the path in G resulting from p is a cycle C . Since we start from a positive vertex and end in a negative one, the cycle has to change

sign an odd number of times and, therefore, uses an odd number of edges from S_i . In order to find a shortest cycle, we compute a shortest path from v^+ to v^- for all $v \in V$.

2.2 Running Time

In each phase, we have the shortest-path computations which take time $O(n(m + n \log n))$ and the update of the sets, which take $O(m^2)$ time. We execute $O(m)$ phases and, therefore, the running time is $O(m^3 + m^2n + mn^2 \log n)$.

2.3 Heuristic Improvements

In this section, we present several heuristics that can improve the running time substantially. All heuristics preserve the worst-case time and space bounds.

2.3.1 Compressed Representation (H1). All vectors (sets S and cycles C), which are handled by the algorithm, are in $\{0, 1\}^m$. Moreover, any operations performed are normal set operations. This allows us to use a compressed representation where each entry of these vectors is represented by a bit of an integer. This allows us to save up space and, at the same time, to perform 32 or 64 bitwise operations in parallel.

2.3.2 Upper Bounding the Shortest Path (H2). During phase i , we might perform up to n shortest-path computations in order to compute the shortest cycle C_i with an odd intersection with the set S_i . We can use the shortest path found so far as an upper bound on the shortest path. This is implemented as follows; a node is only added in the priority queue of Dijkstra's implementation if its current distance is not more than our current upper bound.

2.3.3 Reducing the Shortest-Path Computations (H3). We come to the most important heuristic. In each of the N phases we are performing n shortest-path computations. This results in $\Omega(mn)$ shortest-path computations.

Let $S = \{e_1, e_2, \dots, e_k\}$ be a *witness* at some point of the execution. We need to compute the shortest cycle C s.t. $\langle C, S \rangle = 1$. We can reduce the number of shortest-path computations based on the following observation.

Let $C_{\geq i}$ be the shortest cycle in G s.t. $\langle C_{\geq i}, S \rangle = 1$, and $C_{\geq i} \cap \{e_1, \dots, e_{i-1}\} = \emptyset$, and $e_i \in C_{\geq i}$. Then, cycle C can be expressed as

$$C = \min_{i=1, \dots, k} C_{\geq i}.$$

We can compute $C_{\geq i}$ in the following way. We delete edges $\{e_1, \dots, e_i\}$ from the graph G and the corresponding edges from the signed graph G_i . Let $e_i = (v, u) \in G$. We then compute a shortest path in G_i from v^+ to u^+ . The path computed will have an even number of edges from the set S , and together with e_i , an odd number. Since we deleted edges $\{e_1, \dots, e_i\}$, the resulting cycle does not contain any edges from $\{e_1, \dots, e_{i-1}\}$.

Using, the above observation, we can compute each cycle in $O(kSP(n, m))$ time when $|S| = k < n$, and in $O(nSP(n, m))$, when $|S| \geq n$. Thus, the running




Fig. 1. Horton's candidate cycle for the MCB.

time for the cycles computations is equal to

$$SP(n, m) \sum_{i=1, \dots, N} \min\{n, |S_i|\}$$

where $SP(n, m)$ is the time to compute a single-source shortest path on an undirected weighted graph with m edges and n vertices.

2.4 A New Hybrid Algorithm

The first polynomial algorithm [Horton 1987] developed did not compute the cycles one by one, but instead computed a superset of the MCB and then greedily extracted the MCB by Gaussian elimination. This superset contains $O(mn)$ cycles, which are constructed in the following way.

For each vertex $w \in V$ and edge $e = (v, u) \in E$, construct the cycle $C = SP(w, v) + SP(w, u) + (v, u)$, where $SP(a, b)$ is the shortest path from a to b (Figure 1). If these two shortest paths do not contain a vertex other than w in common then keep the cycle, otherwise discard it. Let us call this set of cycles, the *Horton set*. It was shown in Horton [1987] that the Horton set always contains an MCB. However, not every MCB is contained in the Horton set.

Based on the above and motivated by the need to reduce the cost of the shortest-path computations, we develop a new algorithm, which combines the two approaches. That is, compute the Horton set and extract the MCB not by using Gaussian elimination, which would take time $O(m^3n)$, but by using the orthogonal space, of the cycle, space as we did in Algorithm 1. The Horton set contains an MCB, but not necessarily all the cycles that belong to any MCB. We resolve this difficulty by ensuring uniqueness of the MCB. We ensure uniqueness by ensuring uniqueness of the shortest-path distances on the graph (either by perturbation or by lexicographic ordering). After the preprocessing step, every cycle of the MCB will be contained in the Horton set and, therefore, we can query the superset for the cycles instead of the graph G . A succinct description can be found in Algorithm 2.

The above algorithm has worst case running time $O(m^2n^2)$. This is because the Horton set contains, at most, mn cycles. We, therefore, need to search for, at most, m cycles and each cycle contains, at most, n edges. The space requirement is $O(mn^2)$, a factor of n more than the lower bound, which is $O(mn)$ to represent the cycles of the MCB.

The important property of this algorithm is that the time to actually compute the cycles is only $O(n^2m)$, which is, by a factor of $\frac{m}{n}$, better than the $O(m^2n)$

Algorithm 2. Construct an MCB

```

Ensure uniqueness of shortest path distances of  $G$  (lexicographically or by
perturbation)
Construct superset (Horton set)  $S$  of the MCB
Set  $S_i = \{e_i\}$  for all  $i = 1, \dots, N$ 
for  $i = 1$  to  $N$  do
  Find  $C_i$  as the shortest cycle in  $S$  s.t  $\langle C_i, S_i \rangle = 1$ 
  for  $j = i + 1$  to  $N$  do
    if  $\langle C_i, S_j \rangle = 1$  then
       $S_j = S_j \oplus S_i$ 
    end
  end
end

```

time required by Algorithm 1. Together with the experimental observation that in general the linear independence step is not the bottleneck, we actually hope to have developed a very efficient algorithm.

3. EXPERIMENTS

3.1 Setup

We perform several experiments in order to understand the running time of the algorithms using the previously presented heuristics. In order to understand the speedup obtained, especially from the use of the $H3$ heuristic, we study, in more detail the cardinalities of the sets S during the algorithm, as well as how many operations are required in order to update these sets. We also compare the running times of Algorithms 1 and 2 with previous implementations.

All experiments are done using random sparse and dense graphs. All graphs were constructed using the $G(n; p)$ model, for $p = 4/n, 0.3, 0.5$ and 0.9 . Our implementation uses LEDA [Mehlhorn and Naher 1999]. All experiments were performed on a Pentium 1.7-Ghz machine with 1 GB of memory, running GNU/Linux. We used the GNU g++ 3.3 compiler with the $-O$ optimization flag. All other implementations use the boost C++ libraries [Boost 2001].

3.2 Updating S_i 's

In this section, we present experimental results, which suggest that the dominating factor of the running time of Algorithm 1 (at least for random graphs) is not the time needed to update the witnesses S , but the time to compute the cycles.

Note that the time to update the witnesses is $O(m^3)$ and the time to compute the cycles is $O(m^2n + mn^2 \log n)$. Thus, on sparse graphs, this algorithm has the same running time $O(n^3 \log n)$ as the fastest known. The currently fastest algorithm [Kavitha et al. 2004] for the MCB problem has running time $O(m^2n + mn^2 \log n + m^\omega)$; the m^ω factor is dominated by the m^2n , but we present it here in order to understand what type of operations the algorithm performs. This

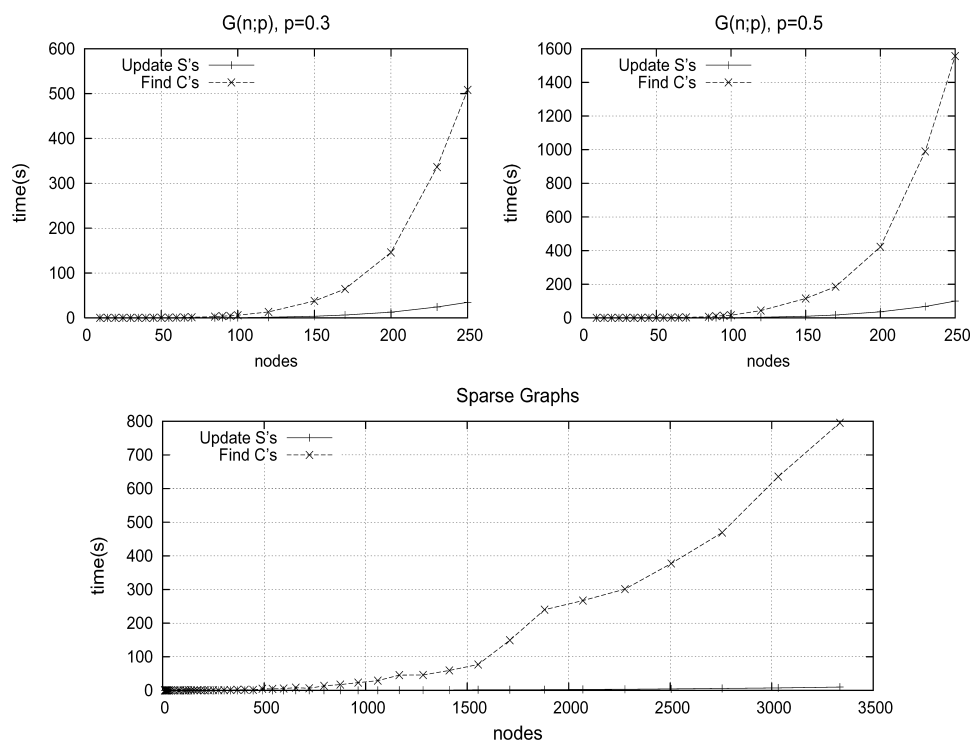


Fig. 2. Comparison of the time taken to update the sets S and the time taken to calculate the cycles on random weighted graphs by Algorithm 1.

algorithm improves upon de Pina [1995] w.r.t the time needed to update the sets S by using fast-matrix multiplication techniques.

Although fast-matrix multiplication can be practical for medium- and large-sized matrices, our experiments show that the time needed to update the sets S is a small fraction of the time needed to compute the cycles. Figure 2 presents a comparison of the required time to update the sets S_i and to calculate the cycles C_i by using the signed graph for random weighted graphs.

In order to get a better understanding of this fact, we performed several experiments. As it turns out, in practice, the average cardinality of the sets S is much less than N and, moreover, the number of times we actually perform set updates (if $\langle C_i, S_j \rangle = 1$) is much less than $N(N - 1)/2$. Furthermore, heuristic $H1$ decreases the constant factor of the running time (for updating S 's) substantially by performing 32 or 64 operations in parallel. This constant factor decrease does not concern the shortest-path computations. Table I summarizes our results.

3.3 Number of Shortest-Path Computations

Heuristic $H3$ improves the best case of the algorithm, while maintaining, at the same time, the worst case. Instead of $\Omega(nm)$ shortest-path computations,

Table I. Statistics about Sets S Sizes on Sparse Random Graphs with $p = 4/n$ and Dense Random Graphs for $p = 0.3$ and 0.5^a

n	m	N	$N(N-1)/2$	$\max(S)$	$\text{avg}(S)$	$\#(S, C) = 1$
sparse ($m \approx 2n$)						
10	19	10	45	4	2	8
51	102	52	1326	25	3	113
104	208	108	5778	44	4	258
206	412	212	22366	108	5	760
491	981	500	124750	226	7	2604
596	1192	609	185136	315	6	2813
963	1925	985	484620	425	7	5469
1060	2120	1084	586986	498	7	5980
1554	3107	1581	1248990	537	8	9540
2070	4139	2105	2214460	1051	13	20645
3032	6064	3092	4778686	1500	13	31356
4441	8882	4525	10235550	2218	17	58186
$p = 0.3$						
10	13	4	6	2	2	2
25	90	66	2145	27	3	137
50	367	318	50403	133	5	1136
75	832	758	286903	370	6	3707
100	1485	1386	959805	613	7	8103
150	3352	3203	5128003	1535	9	22239
200	5970	5771	16649335	2849	10	49066
300	13455	13156	86533590	6398	10	116084
500	37425	36926	681746275	18688	14	455620
$p = 0.5$						
10	22	13	78	7	2	14
25	150	126	7875	57	4	363
50	612	563	158203	298	6	2527
75	1387	1313	861328	654	6	6282
100	2475	2376	2821500	1168	8	15771
150	5587	5438	14783203	2729	9	39292
200	9950	9751	47536125	4769	11	86386
300	22425	22126	244768875	10992	13	227548
500	62375	61876	1914288750	30983	15	837864

^aSets are considered during the whole execution of the algorithm. Column $\#(S, C) = 1$ denotes the number of updates performed on sets S . An upper bound on this is $N(N-1)/2$, which we actually use when bounding the algorithm's running time. Note that the average cardinality of S is very small compared to N , although the maximum cardinality of some S is in $O(N)$.

we hope to perform much less. In Table II, we study the sizes of the sets S_i for $i = 1, \dots, N$ used to calculate the cycles for sparse and dense graphs, respectively.

In both sparse and dense graphs, although the maximum set can have quite large cardinality, the average set size is much less than n . Moreover, in sparse graphs every set used has cardinality less than n . On dense graphs the sets with cardinality less than n are more than 95%. This implies a significant speedup as a result of the $H3$ heuristic.

Figure 3 compares the running times of Algorithm 1 with and without the $H3$ heuristic.

Table II. Statistics about Sets S_i Sizes on Sparse Random Graphs with $p = 4/n$ and Dense Random Graphs for $p = 0.3$ and 0.5^a

n	m	N	$\max(S_i)$	$\lceil \text{avg}(S_i) \rceil$	$ \{S_i : S_i < n\} $
sparse ($m \approx 2n$)					
10	19	10	4	2	10
51	102	52	16	5	52
104	208	108	39	5	108
206	412	212	106	10	212
491	981	498	246	13	498
596	1192	609	220	11	609
963	1925	980	414	11	980
1060	2120	1076	496	17	1076
1554	3107	1573	795	21	1573
2070	4139	2108	1036	27	2108
3032	6064	3092	1468	33	3092
4441	8882	4522	1781	33	4522
$p = 0.3$					
10	13	4	2	2	4
25	90	66	20	4	66
50	367	318	153	10	302
75	832	758	357	15	721
100	1485	1386	638	15	1343
150	3352	3203	1534	18	3133
200	5970	5771	2822	29	5635
300	13455	13156	6607	32	12968
500	37425	36926	15965	39	36580
$p = 0.5$					
10	22	13	7	3	13
25	150	126	66	5	121
50	612	563	222	12	532
75	1387	1313	456	10	1276
100	2475	2376	1094	15	2314
150	5587	5438	2454	19	5338
200	9950	9751	4828	28	9601
300	22425	22126	10803	33	21875
500	62375	61876	30877	38	61483

^aOnly sets used to calculate cycles are considered.

3.4 Running Time

In this section, we compare the various implementations for computing a minimum cycle basis. Except from Algorithms 1 (DP) and 2 (HYB), we include in the comparison two implementations, (HOR) or (HOR.U1) [Kreisbasenbibliothek 2004] and (HOR.U2) [Huber 2002] of Horton's algorithm with running time $O(m^3n)$, and an implementation (FEAS) [Kreisbasenbibliothek 2004] of the $O(m^3 + mn^2 \log n)$ algorithm presented in Berger et al. [2004]. Algorithms 1 and 2 are implemented with compressed integer sets. Fast-matrix multiplication [Kavitha et al. 2004; Golynski and Horton 2002] can nicely improve many parts of these implementations with respect to the worst-case complexity. We did not experiment with these versions of the algorithms.

The comparison of the running times is performed for three different types of undirected graphs: (a) random sparse graphs, where $m \approx 2n$, (b) random graphs

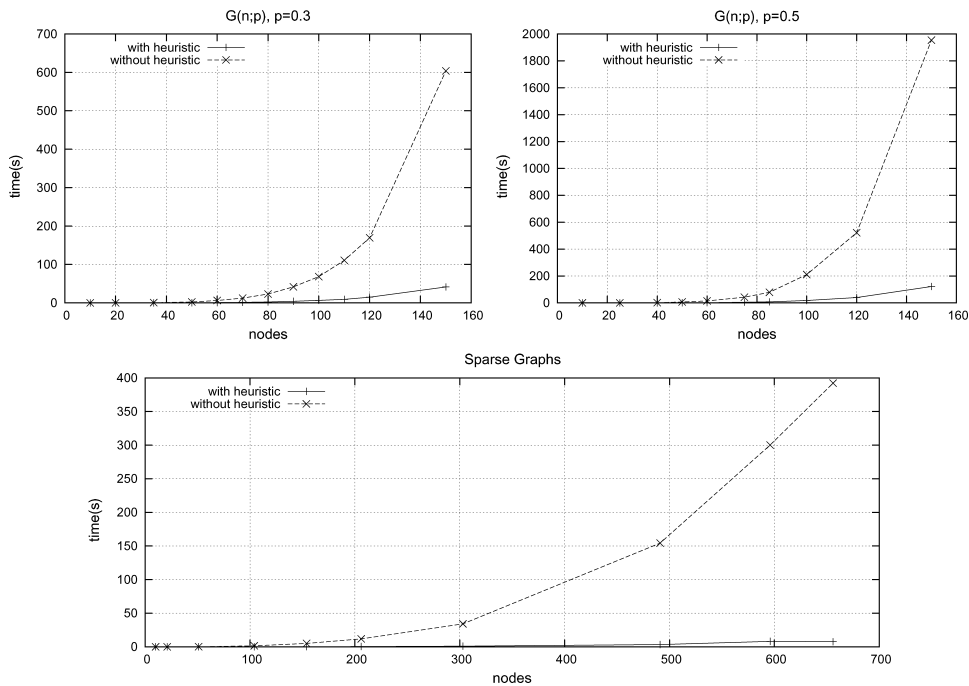


Fig. 3. Running times of Algorithm 1 with and without the $H3$ heuristic. Without the heuristic the algorithm is forced to perform $\Omega(nm)$ shortest path computations.

from $G(n,p)$ with different density $p = 0.3, 0.5$, and 0.9 , and (c) hypercubes. Tests are performed for both weighted and unweighted graphs. In the case of weighted graphs, the weight of an edge is an integer chosen independently at random from the uniform distribution in the range $[0 \dots 2^{16}]$.

Figures 4 and 5 summarize the results of these comparisons. In the case of weighted graphs, Algorithm 1 is definitely the winner. On the other hand, in the case of dense unweighted graphs, Algorithm 2 performs much better. As can be easily observed the differences on the running time of the implementations are rather small for sparse graphs. For dense graphs, however, we observe a substantial difference in performance.

3.4.1 Dense Unweighted Graphs. In the case of dense unweighted graphs, the hybrid algorithm performs better than the other algorithms. However, even on the exact same graph, the addition of weights changes the performance substantially. This change in performance is not because of the difference in size of the produced Horton set between the unweighted and the weighted case, but because of the total number of queries that have to be performed in this set.

In the hybrid algorithm before computing the MCB, we sort the cycles of the Horton set. Then, for each of the N phases, we *query* the Horton set from the least costly cycle to the most, until we find a cycle with an odd intersection with our current witness S . Figure 6 plots the number of cycles in the Horton set and the number of queries required in order to extract the MCB from this set

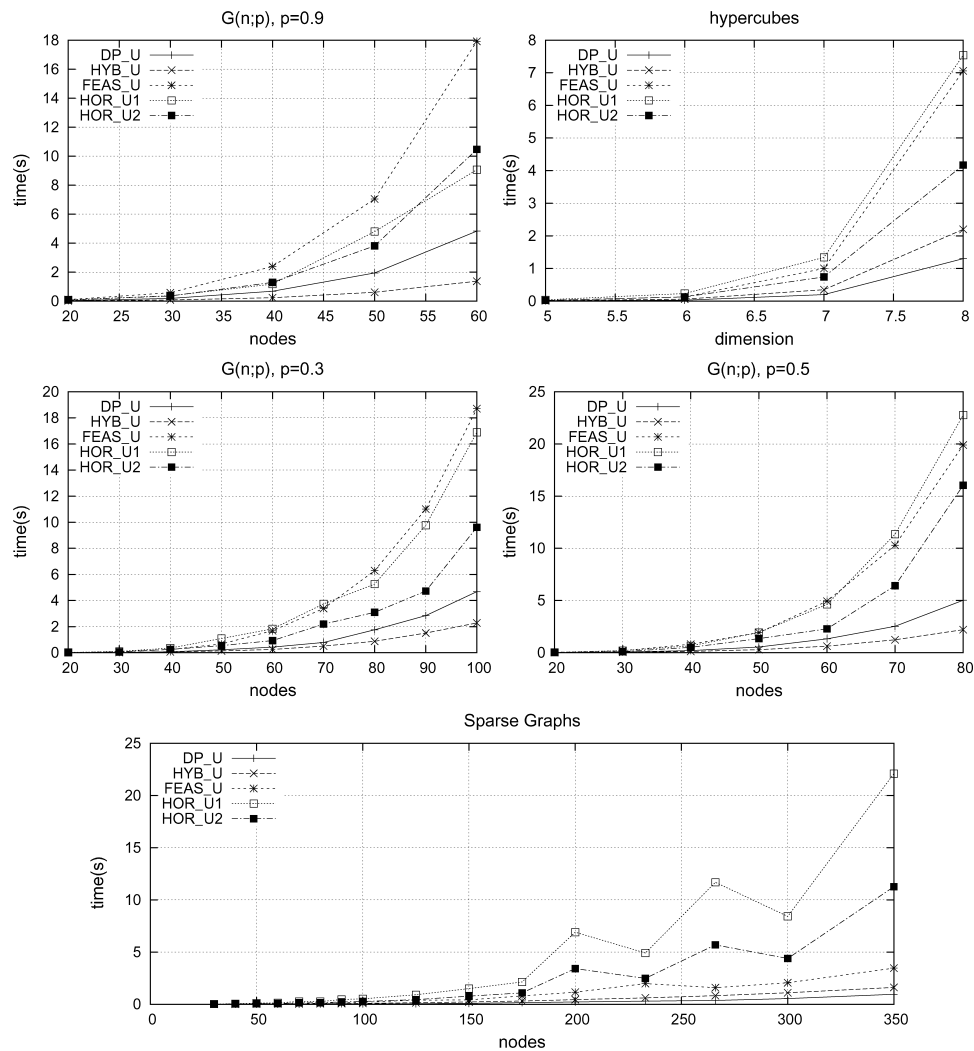


Fig. 4. Comparison of various algorithms for random *unweighted* graphs. Algorithm 1 is denoted as *DP_U* and Algorithm 2 as *HYB_U*. *HOR_U1* and *HOR_U2* are two different implementation of Horton’s algorithm. *FEAS_U* is an implementation of another $O(m^3)$ algorithm. See Section 3.4 for details.

for dense graphs. In the case of unweighted graphs, the number of queries is substantially smaller than in the case of weighted graphs. This is exactly the reason why the hybrid algorithm outperforms the others in unweighted dense graphs.

4. CONCLUSIONS

In this paper, we described an efficient implementation of an $O(m^3 + mn^2 \log n)$ algorithm for calculating a minimum cycle basis of an undirected graph. Using several heuristics, we improved its best case in order to achieve low running

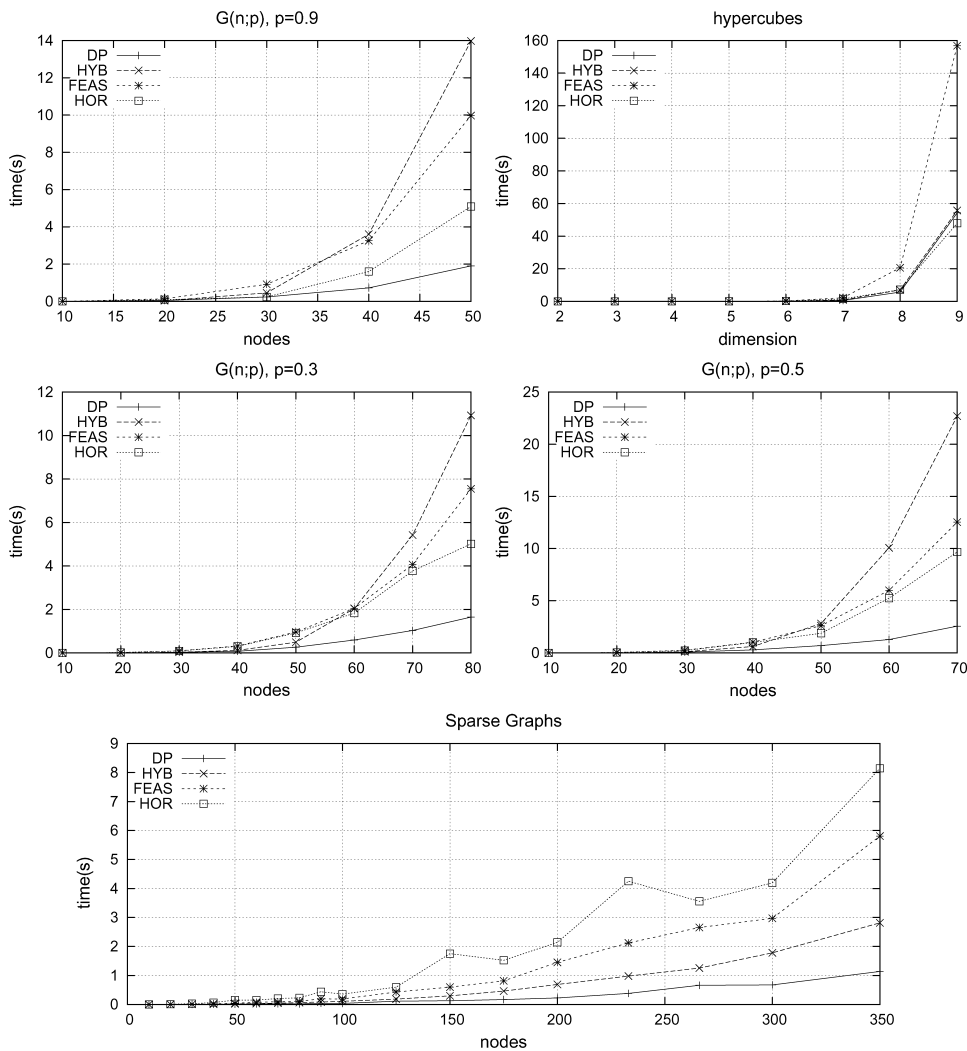


Fig. 5. Comparison of various algorithms for random *weighted* graphs. Algorithm 1 is denoted as *DP* and Algorithm 2 as *HYB*. *HOR* is Horton’s algorithm and *FEAS* is another $O(m^3 + mn^2 \log n)$ algorithm. See Section 3.4 for details.

times. By experimenting on sparse and dense graphs we justified, for random graphs, the usefulness of our heuristics. We also observed that the $O(m^2n + mn^2 \log n)$ factor is the bottleneck of the algorithm in practice.

Moreover, we compared our implementation with a new “hybrid” algorithm, which combines the two different approaches used by the previous algorithms for solving the minimum cycle basis problem. This algorithm has running time $O(m^2n^2)$ and performs very well, in practice, for dense unweighted graphs.

Finally, we performed a comparison between the various implementations available in order to compute a minimum cycle basis and our new implementation. These experimental results suggest that future research should focus

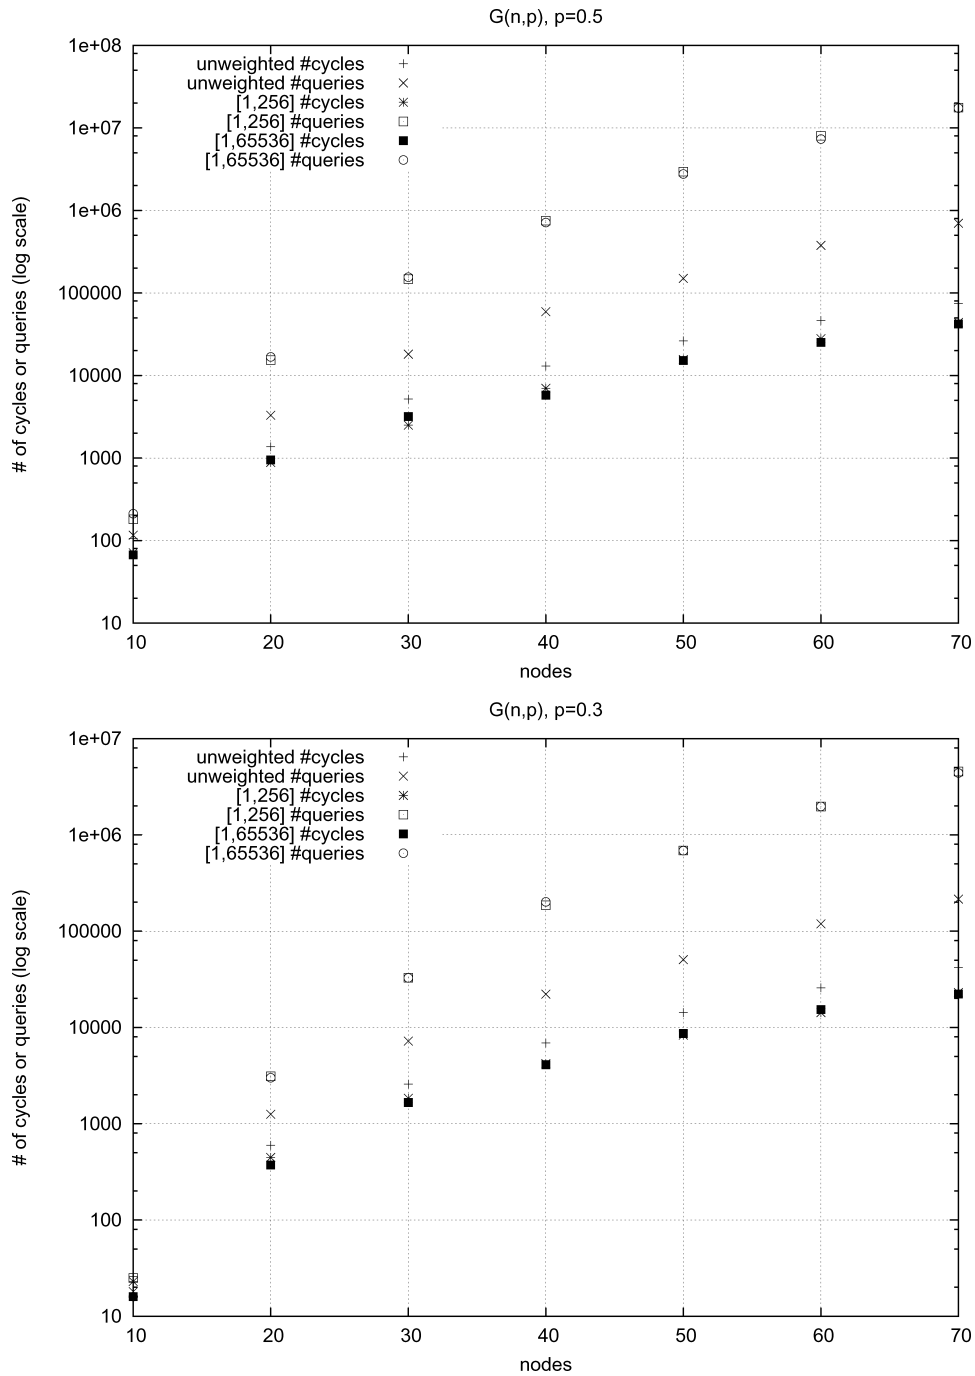


Fig. 6. Number of cycles in the Horton set (set with duplicates) and number of queries required in this set (set sorted by cycle weight) in order to extract the MCB for random dense graphs with random weights of different ranges. Each random graph is considered with three different edge weight ranges: (a) unweighted, (b) weights in $[1, 2^8]$, and (c) weights in $[1, 2^{16}]$.

on computing each cycle using a different approach, which needs less than n shortest-path computations. This could be accomplished as an independent problem or in relation with the structure of the particular sets S , which are produced during the MCB algorithm's execution.

REFERENCES

- BERGER, F., GRITZMANN, P., AND DE VRIES, S. 2004. Minimum cycle basis for network graphs. *Algorithmica* 40, 1, 51–62.
- BOOST. 2001. *C++ Libraries*. <http://www.boost.org>.
- CASELL, A. C., HENDERSON, J. C., AND RAMACHANDRAN, K. 1976. Cycle bases of minimal measure for the structural analysis of skeletal structures by the flexibility method. *Proc. Royal Society of London Series A* 350, 61–70.
- CHUA, L. O. AND CHEN, L. 1973. On optimally sparse cycle and coboundary basis for a linear graph. *IEEE Trans. Circuit Theory CT-20*, 495–503.
- COPPERSMITH, D. AND WINOGRAD, S. 1990. Matrix multiplications via arithmetic progressions. *Journal of Symb. Comput.* 9, 251–280.
- DE PINA, J. 1995. Applications of shortest path methods. Ph.D. thesis, University of Amsterdam, Netherlands.
- GOLYNSKI, A. AND HORTON, J. D. 2002. A polynomial time algorithm to find the minimum cycle basis of a regular matroid. In *8th Scandinavian Workshop on Algorithm Theory*.
- GOTSMAN, C., KALIGOSI, K., MEHLHORN, K., MICHAIL, D., AND PYRGA, E. 2005. Cycle bases of graphs and sampled manifolds. manuscript. <http://www.mpi-inf.mpg.de/~michail/papers/mcb-manifold.pdf>.
- HORTON, J. D. 1987. A polynomial-time algorithm to find a shortest cycle basis of a graph. *SIAM Journal of Computing* 16, 359–366.
- HUBER, M. 2002. Implementation of algorithms for sparse cycle bases of graphs. <http://www-m9.ma.tum.de/dm/cycles/mhuber>.
- KAVITHA, T., MEHLHORN, K., MICHAIL, D., AND PALUCH, K. E. 2004. A faster algorithm for minimum cycle basis of graphs. In *31st International Colloquium on Automata, Languages and Programming*, Finland. 846–857.
- KREISBASENBIBLIOTHEK. 2004. *CyBaL*. <http://www-m9.ma.tum.de/dm/cycles/cybal>.
- MEHLHORN, K. AND MICHAIL, D. 2005. Implementing minimum cycle basis algorithms. In *Experimental and Efficient Algorithms, 4th International Workshop, WEA 2005, Santorini Island, Greece, May 10–13, 2005, Proceedings*, S. E. Nikolettseas, Ed. Lecture Notes in Computer Science, vol. 3503. Springer, New York. 32–43.
- MEHLHORN, K. AND NAHER, S. 1999. *LEDA: A Platform for Combinatorial and Geometric Computing*. Cambridge University Press, Cambridge.
- TEWARI, G., GOTSMAN, C., AND GORTLER, S. J. 2005. Meshing genus-1 point clouds using discrete one-forms. Tech. Rep. TR-14-05, Computer Science Group, Harvard University.

Received September 2005; revised January 2006; accepted January 2006