

## Large-Graph Layout Algorithms at Work: An Experimental Study

*Stefan Hachul   Michael Jünger*

Institut für Informatik  
Universität zu Köln  
Pohligstraße 1, 50969 Köln, Germany  
{hachul,mjuenger}@informatik.uni-koeln.de

### Abstract

In the last decade several algorithms that generate straight-line drawings of general large graphs have been invented. In this paper we investigate some of these methods that are based on force-directed or algebraic approaches in terms of running time and drawing quality on a big variety of artificial and real-world graphs. Our experiments indicate that there exist significant differences in drawing qualities and running times depending on the classes of tested graphs and algorithms.

Article Type	Communicated by	Submitted	Revised
Regular paper	P. Eades and P. Healy	January 2006	January 2007

## 1 Introduction

What do biochemical reactions of proteins in baker’s yeast, the ecosystem of plankton, sea perch, and anchovy, the American electricity network, and the international air traffic have in common? They can be modeled as graphs.

In general a graph  $G = (V, E)$  is used to model information that can be described as objects (the node set  $V$ ) and connections between those objects (the edge set  $E$ ).

One fundamental tool for analyzing such graphs is the automatic generation of layouts that visualize the graphs and are easy to understand. A popular class of algorithms that is used to visualize general graphs are force-directed graph-drawing methods. Those methods generate drawings of a given general graph  $G = (V, E)$  in the plane in which each edge is represented by a straight line connecting its two adjacent nodes. The computation of the drawings is based on associating  $G$  with a physical model. Then, an iterative algorithm tries to find a placement of the nodes so that the total energy in the underlying physical system is minimal. Important esthetic criteria are uniformity of edge length, few edge crossings, non-overlapping nodes and edges, and the display of symmetries if some exist.

Classical force-directed algorithms like [5, 16, 7, 4, 6] are used successfully in practice (see e.g., [2]) for drawing general graphs containing few hundreds of vertices. However, in order to generate drawings of graphs that contain thousands or hundreds of thousands of vertices more efficient force-directed techniques have been developed [21, 20, 9, 8, 13, 23, 12, 11]. Besides fast force-directed algorithms other very fast methods for drawing large graphs (see e.g., [14, 17]) have been invented. These methods are based on techniques of linear algebra instead of physical analogies. But they strive for the same esthetic drawing criteria.

Previous experimental tests of these methods are mainly restricted to regular graphs with *grid-like* structures (see e.g., [14, 17, 9, 23, 13]). Since general graphs share these properties quite seldom, and since the test environments of these experiments are different, a standardized comparison of the methods on a wider range of graphs is needed.

In this study we experimentally compare some of the fastest state-of-the-art algorithms for straight-line drawing of general graphs on a big variety of graph classes. In particular, we investigate the force-directed algorithm **GRIP** of Gajer and Kobourov [9] and Gajer et al. [8], the *Fast Multi-scale Method* (**FMS**) of Harel and Koren [13], and the *Fast Multipole Multilevel Method* (**FM<sup>3</sup>**) of Hachul and Jünger [12, 11]. We did not test other efficient force-directed methods like **FADE** of Quigley and Eades [20] and **JIGGLE** of Tunkelang [21] since the original implementations were not available to us or the implementations did not allow to import test graphs, respectively.

The examined algebraic methods are the algebraic multigrid method **ACE** of Koren et al. [17] and the *high-dimensional embedding* approach (**HDE**) of Harel and Koren [14]. Additionally, one of the fastest classical force-directed algorithms, namely the *grid-variant algorithm* (**GVA**) of Fruchterman and Rein-

gold [7], is tested as a benchmark.

After a short description of the tested algorithms in Section 2 and a presentation of the experimental framework in Section 3, our results will be presented in Section 4.

## 2 The Algorithms

### 2.1 The Grid-Variant Algorithm (GVA)

The grid-variant algorithm (GVA) of Fruchterman and Reingold [7] is based on a model of pairwise repelling charged particles (the nodes) and attracting springs (the edges), similar to the model of the **Spring Embedder** of Eades [5]. Since a naive exact calculation of the repulsive forces acting between all pairs of charges needs  $\Theta(|V|^2)$  time per iteration, GVA only calculates the repulsive forces acting between nodes that are placed relatively near to each other. Therefore, the rectangular drawing area is subdivided into a regular square grid. The repulsive forces that act on a node  $v$  that is contained in a grid box  $B$  are approximated by summing up only the repulsive forces that are induced by the nodes contained in  $B$  and the nodes in the grid boxes that are neighbors of  $B$ . If the number of iterations is assumed to be constant, the best-case running time of GVA is  $\Theta(|V| + |E|)$ . The worst-case running time, however, remains  $\Theta(|V|^2 + |E|)$ .

### 2.2 The Method GRIP

Gajer and Kobourov [9] and Gajer et al. [8] developed the force-directed multi-level algorithm GRIP. In general, multilevel algorithms are based on two phases. A *coarsening phase*, in which a sequence of coarse graphs with decreasing sizes is computed and a *refinement phase* in which successively drawings of finer graphs are computed, using the drawings of the next coarser graphs and a variant of a suitable force-directed single-level algorithm.

The coarsening phase of GRIP is based on the construction of a *maximum independent set filtration* or *MIS filtration* of the node set  $V$ . A MIS filtration is a family of sets  $\{V =: V_0, V_1, \dots, V_k\}$  with  $\emptyset \subset V_k \subset V_{k-1} \dots \subset V_0$  so that each  $V_i$  with  $i \in \{1, \dots, k\}$  is a maximal subset of  $V_{i-1}$  for which the graph-theoretic distance between any pair of its elements is at least  $2^{i-1} + 1$ . Gajer and Kobourov [9] use a **Spring Embedder**-like method as a single-level algorithm at each level. The used force vector is similar to that used in the method of Kamada and Kawai [16], but is restricted to a suitable chosen subset of  $V_i$ .

Other notable specifics of GRIP are that it computes the MIS filtration only and no edge sets of the coarse graphs  $G_0, \dots, G_k$  that are induced by the filtrations. Furthermore, it is designed to place the nodes in an  $n$ -dimensional space ( $n \geq 2$ ), to draw the graph in this space, and to project it into two or three dimensions.

The asymptotic running time of the algorithm, excluding the time that is needed to construct the MIS filtration, is  $\Theta(|V|(\log \text{diam}(G))^2)$  for graphs with

bounded maximum node degree, where  $diam(G)$  denotes the diameter of  $G$ .

### 2.3 The Fast Multi-scale Method (FMS)

In order to create the sequence of coarse graphs in the force-directed multilevel method **FMS**, Harel and Koren [13] use an  $O(k|V|)$  algorithm that finds a 2-approximative solution of the  $\mathcal{NP}$ -hard  $k$ -center problem [10]. The node set  $V_i$  of a graph  $G_i$  in the sequence  $G_0, \dots, G_k$  is determined by the approximative solution of the  $k_i$ -center problem on  $G$  with  $k_i > k_{i+1}$  for all  $i \in \{0, \dots, k-1\}$ .

The authors use a variation of the algorithm of Kamada and Kawai [16] as a force-directed single-level algorithm. In order to speed up the computation of this method, they modify the energy function of Kamada and Kawai [16] that is associated with a graph  $G_i$  with  $i \in \{0, \dots, k-1\}$ . The difference to the original energy of Kamada and Kawai [16] is that only some of the  $|V(G_i)| - 1$  springs that are connected with a node  $v \in V(G_i)$  are considered.

The asymptotic running time of **FMS** is  $\Theta(|V||E|)$ . Additionally,  $\Theta(|V|^2)$  memory is needed to store the distances between all pairs of nodes.

### 2.4 The Fast Multipole Multilevel Method (FM<sup>3</sup>)

The force-directed multilevel algorithm **FM<sup>3</sup>** has been introduced by Hachul and Jünger [12, 11]. It is based on a combination of an efficient multilevel technique with an  $O(|V| \log |V|)$  approximation algorithm to obtain the repulsive forces between all pairs of nodes.

In the coarsening step subgraphs with a small diameter (called *solar systems*) are collapsed to obtain a multilevel representation of the graph. In the used single-level algorithm, the bottleneck of calculating the repulsive forces acting between all pairs of charged particles in the **Spring Embedder**-like force model is overcome by rapidly evaluating potential fields using a novel multipole-based tree code. The worst-case running time of **FM<sup>3</sup>** is  $O(|V| \log |V| + |E|)$  with linear memory requirements.

### 2.5 The Algebraic Multigrid Method ACE

In the description of their method **ACE**, Koren et al. [17] define the quadratic optimization problem

$$(P) \quad \min x^T L x \quad \text{so that} \quad x^T x = 1 \quad \text{in the subspace} \quad x^T \mathbf{1}_n = 0.$$

Here,  $n = |V|$  and  $L$  is the *Laplacian* matrix of  $G$ .

The minimum of (P) is obtained by the eigenvector that corresponds to the smallest positive eigenvalue of  $L$ . The problem of drawing the graph  $G$  in two dimensions is reduced to the problem of finding the two eigenvectors of  $L$  that are associated with the two smallest eigenvalues.

Instead of calculating the eigenvectors directly, an *algebraic multigrid algorithm* is used. Similar to the force-directed multilevel ideas, the idea is to

express the originally high-dimensional problem in lower and lower dimensions, solving the problem at the lowest dimension, and progressively solving a high-dimensional problem by using the solutions of the low-dimensional problems.

The authors do not give an upper bound on the asymptotic running time of ACE in the number of nodes and edges.

## 2.6 High-Dimensional Embedding (HDE)

The method HDE of Harel and Koren [14] is based on a two phase approach that, first, generates an embedding of the graph in a very high-dimensional vector space and, then, projects this drawing into the plane.

The high-dimensional embedding of the graph is computed by, first, using a linear time algorithm for approximatively solving the  $k$ -center problem [10]. A fixed value of  $k = 50$  is chosen, and  $k$  is also the dimension of the high-dimensional vector space. Then, breadth-first search starting from each of the  $k$  center nodes is performed resulting in  $k$   $|V|$ -dimensional vectors that store the graph-theoretic distances of each  $v \in V$  to each of the  $k$  centers. These vectors are interpreted as a  $k$ -dimensional embedding of the graph.

In order to project the high-dimensional embedding of the graph into the plane, the  $k$  vectors are used to define a *covariance* matrix  $S$ . The  $x$ - and  $y$ -coordinates of the two-dimensional drawing are obtained by calculating the two eigenvectors of  $S$  that are associated with its two largest eigenvalues. HDE runs in  $\Theta(|V| + |E|)$  time.

## 3 The Experiments

### 3.1 Test-Environment, Implementations, and Parameter Settings

All experiments were performed on a 2.8 GHz Intel Pentium 4 PC with one gigabyte of memory running Linux.

We tested an implementation of GVA by S. Näher and D. Alberts that is part of the AGD [15] library, an implementation of GRIP by R. Yusufov that is available from [24], and implementations of FMS, ACE, HDE by Y. Koren that are available from [18]. Finally, we tested our own implementation of FM<sup>3</sup>.

The provided executables of ACE, HDE, and FMS are Microsoft Windows executables. They were tested on the same machine but using Windows as operating system, instead.

In order to obtain a fair comparison, we ran each algorithm with the same set of standard-parameter settings (given by the authors) on each tested graph. However, we are aware that in some cases it might be possible to obtain better results by spending a considerable amount of time with trial-and-error searching for an optimal set of parameters for each algorithm and graph.

### 3.2 The Set of Test Graphs

Since only few implementations can handle disconnected and weighted graphs, we restrict our attention to connected unweighted graphs, here.

We generated several classes of artificial graphs to examine the scaling of the algorithms on graphs with predefined structures but different sizes.

These are *random grid* graphs that were obtained by, first, creating regular square grid graphs and, then, randomly deleting 3% of the nodes. The *sierpinski* graphs were created by associating the *Sierpinski Triangles* with graphs. Furthermore, we generated complete *6-nary trees*.

The next two classes of artificial graphs were designed to test how well the algorithms can handle highly non-uniform distributions of the nodes and high node degrees. Therefore, we created these graphs in a way so that one can expect that an energy-minimal configuration of the nodes in a drawing that relies on a **Spring Embedder**-like force model induces a tiny subregion of the drawing area which contains  $\Theta(|V|)$  nodes. In particular, we constructed trees that contain a root node  $r$  with  $|V|/4$  neighbors. The other nodes were subdivided into six subtrees of equal size rooted at  $r$ . We called these graphs *snowflake* graphs.

Additionally we created *spider* graphs by constructing a circle  $C$  containing 25% of the nodes. Each node of  $C$  is also adjacent to 12 other nodes of the circle. The remaining nodes were distributed on 8 paths of equal length that were rooted at one node of  $C$ . In contrast to the snowflake graphs, the spider graphs have bounded maximum degree.

The last kind of artificial graphs are graphs with a relatively high edge density  $|E|/|V| \geq 14$ . We call them *flower* graphs. They are constructed by joining 6 circles of equal length at a single node before replacing each of the nodes by a complete subgraph with 30 nodes ( $K_{30}$ ).

The rest of the test graphs are taken from real-world applications. In particular, we selected graphs from the *AT&T graph library* [1], from C. Walshaw's graph collection [22], and a graph that describes a social network of 2113 people that we obtained from C. Lipp.

We partitioned the artificial and real-world graphs into two sets. The first set are graphs that consist of few biconnected components, have a constant maximum node degree, and have a low edge density. Furthermore, one can expect that an energy-minimal configuration of the nodes in a **Spring Embedder** drawing of such a graph does not contain  $\Theta(|V|)$  nodes in an extremely tiny subregion of the drawing area. Since one can anticipate from previous experiments [14, 17, 9, 13] that graphs contained in this set do not cause problems for many of the tested algorithms, we call the set of these graphs *kind*. The second set is the complement of the first one, and we call the set of these graphs *challenging*.

Table 1 gives an overview of the structures of the tested graphs.

Type	Name	$ V $	$ E $	$ B $	$\frac{ E }{ V }$	max. degree
Kind Artificial	rnd_grid_032	985	1834	2	1.8	4
	rnd_grid_100	9497	17849	6	1.8	4
	rnd_grid_320	97359	184532	2	1.9	4
	sierpinski_06	1095	2187	1	2.0	4
	sierpinski_08	9843	19683	1	2.0	4
	sierpinski_10	88575	177147	1	2.0	4
Kind Real World	crack	10240	30380	1	2.9	9
	fe_pwt	36463	144794	55	3.9	15
	finan_512	74752	261120	1	3.4	54
	fe_ocean	143437	409593	39	2.8	6
Challenging Artificial	tree_06_04	1555	1554	1554	1.0	7
	tree_06_05	9331	9330	9330	1.0	7
	tree_06_06	55987	55986	55986	1.0	7
	snowflake_A	971	970	970	1.0	256
	snowflake_B	9701	9700	9700	1.0	2506
	snowflake_C	97001	97000	97000	1.0	25006
	spider_A	1000	2200	801	2.2	18
	spider_B	10000	22000	8001	2.2	18
	spider_C	100000	220000	80001	2.2	18
	flower_A	930	13521	1	14.5	30
	flower_B	9030	131241	1	14.5	30
	flower_C	90030	1308441	1	14.5	30
Challenging Real World	ug_380	1104	3231	27	2.9	856
	esslingen	2075	5530	867	2.6	97
	add_32	4960	9462	951	1.9	31
	dg_1087	7602	7601	7601	1.0	6566
	bcsstk_33	8738	291583	1	33.3	140
	bcsstk_31	35586	572913	48	16.1	188

Table 1: The test graphs and their structures

### 3.3 The Criteria of Evaluation

The natural criteria to evaluate a graph-drawing algorithm in practice are the needed running times and the quality of the drawings.

In order to evaluate the quality of the drawings, we focus on the aesthetic criteria that all tested algorithms strive for (i.e., uniformity of edge length, few edge crossings, non-overlapping nodes and edges, and the display of symmetries).

Suppose  $G = (V, E)$  is a graph and  $\Gamma$  is a drawing of  $G$ . Let  $l_\Gamma(e)$  denote the length of an edge  $e \in E$  in  $\Gamma$  and  $l_\Gamma^{\text{av}}$  denote the average edge length of an edge in  $\Gamma$ . To measure the uniformity of the edges, we calculated the *normalized standard deviation of the edge length* that is

$$\sigma_{\Gamma} := \sqrt{\sum_{e \in E} \frac{(l_{\Gamma}(e) - l_{\Gamma}^{\text{av}})^2}{|E| \cdot (l_{\Gamma}^{\text{av}})^2}}.$$

We counted the number of edge crossings in  $\Gamma$  (denoted by  $ecn_{\Gamma}$ ) as well as the number of pairs of edges that completely overlap each other (denoted by  $eon_{\Gamma}$ ). To compare these measures on graphs with different sizes, we define the *relative edge-crossing number* ( $recn_{\Gamma}$ ) and the *relative edge-overlapping number* ( $reon_{\Gamma}$ ) of a drawing  $\Gamma$  of  $G$  as follows:

$$recn_{\Gamma} := \frac{|ecn_{\Gamma}|}{|E|} \quad reon_{\Gamma} := \frac{|eon_{\Gamma}|}{|E|}$$

We did not measure the number of (partially) overlapping nodes. The reason is that per default some implementations represent nodes by points (**ACE**, **HDE**), whereas the others draw them as circles that occupy nonempty area. Since symmetry-detection is  $\mathcal{NP}$ -hard [19], we did not explicitly measure the symmetry of a given drawing and printed the drawing, instead. One of the most important implicit goals of a graph-drawing algorithm is that an individual user is satisfied with a drawing. Therefore, we provide printouts of the computed drawings, too.

## 4 The Results

### 4.1 Comparison of the Running Times

The running times of the methods **GVA**, **FM<sup>3</sup>**, **GRIP**, **FMS**, **ACE**, and **HDE** for the tested graphs are presented in Table 2.

As expected, in most cases **GVA** is the slowest method among the force-directed algorithms. The largest graph `fe_ocean` is drawn by **GVA** in 5 hours and 20 minutes.

The method **FM<sup>3</sup>** is significantly faster than **GVA** for all tested graphs. The running times range from less than 2 seconds for the smallest graphs to less than 6 minutes for the largest graph `fe_ocean`. The sub-quadratic scaling of **FM<sup>3</sup>** can be experimentally confirmed for all classes of tested graphs.

Except for the dense graphs `flower_B` and `bcsstk_33` **GRIP** is faster than **FM<sup>3</sup>** (up to a factor of 9). Unfortunately, we could not examine the scaling of **GRIP** for the largest graphs due to an error in the executable.

Since the memory requirement of **FMS** is quadratic in the size of the graph, the implementation of **FMS** is restricted to graphs that contain at most 10,000 nodes. The running times of **FMS** are comparable with those of **FM<sup>3</sup>** for the smallest and the medium sized kind graphs. In contrast to this, the CPU times of **FMS** increase drastically for several challenging graphs, in particular for graphs that either contain nodes with a very high degree or have a high edge density.



Type	Name	CPU Time in Seconds					
		GVA	FM <sup>3</sup>	GRIP	FMS	ACE	HDE
Kind Artificial	rnd_grid_032	<u>12.5</u>	1.9	0.3	1.0	< <b>0.1</b>	< <b>0.1</b>
	rnd_grid_100	<u>203.4</u>	19.1	4.4	32.0	0.5	<b>0.1</b>
	rnd_grid_320	<u>6316.1</u>	215.4	( <i>E</i> )	( <i>M</i> )	4.1	<b>1.3</b>
	sierpinski_06	<u>13.1</u>	1.8	0.3	1.0	< <b>0.1</b>	< <b>0.1</b>
	sierpinski_08	<u>171.7</u>	16.8	4.8	33.0	1.0	<b>0.1</b>
	sierpinski_10	<u>3606.4</u>	162.0	( <i>E</i> )	( <i>M</i> )	23.4	<b>1.0</b>
Kind Real World	crack	<u>317.5</u>	23.0	6.8	( <i>M</i> )	0.4	<b>0.2</b>
	fe_pwt	1869.1	69.0	( <i>E</i> )	( <i>M</i> )	( <i>T</i> )	<b>0.5</b>
	finan_512	<u>6319.8</u>	158.2	( <i>E</i> )	( <i>M</i> )	7.5	<b>1.0</b>
	fe_ocean	<u>19247.0</u>	355.9	( <i>E</i> )	( <i>M</i> )	4.0	<b>3.4</b>
Chal- lenging Artificial	tree_06_04	<u>14.3</u>	2.6	0.3	2.0	< <b>0.1</b>	< <b>0.1</b>
	tree_06_05	<u>130.3</u>	17.7	2.4	43.0	0.5	< <b>0.1</b>
	tree_06_06	<u>1769.2</u>	121.3	( <i>E</i> )	( <i>M</i> )	4.5	<b>0.5</b>
	snowflake_A	8.0	1.6	0.4	<u>73.0</u>	0.4	< <b>0.1</b>
	snowflake_B	143.2	17.4	6.1	3320.0	( <i>T</i> )	< <b>0.1</b>
	snowflake_C	14685.7	166.5	( <i>E</i> )	( <i>M</i> )	( <i>T</i> )	<b>0.8</b>
	spider_A	<u>17.6</u>	1.9	0.4	1.0	1.1	< <b>0.1</b>
	spider_B	<u>189.0</u>	17.7	7.2	47.0	8.9	<b>0.1</b>
	spider_C	<u>4568.3</u>	177.2	( <i>E</i> )	( <i>M</i> )	280.7	<b>1.3</b>
	flower_A	<u>61.7</u>	1.2	0.7	1.0	< <b>0.1</b>	< <b>0.1</b>
	flower_B	<u>595.1</u>	11.9	19.3	46.0	1.4	<b>0.2</b>
	flower_C	11841.5	121.4	( <i>E</i> )	( <i>M</i> )	( <i>T</i> )	<b>1.4</b>
Chal- lenging Real World	ug_380	<u>23.1</u>	2.1	0.4	1.0	< <b>0.1</b>	< <b>0.1</b>
	esslingen	43.8	4.0	0.5	<u>404.0</u>	1.0	< <b>0.1</b>
	add_32	80.6	12.1	1.6	17.0	0.5	< <b>0.1</b>
	dg_1087	624.8	18.1	3.6	<u>5402.0</u>	108.4	< <b>0.1</b>
	bcsstk_33	1494.6	23.8	29.1	<u>6636.0</u>	0.4	<b>0.3</b>
	bcsstk_31	<u>4338.4</u>	83.6	( <i>E</i> )	( <i>M</i> )	1.9	<b>0.7</b>

Table 2: The test graphs and the running times that are needed by the tested algorithms to draw them. Explanations: (*E*) No drawing was computed due to an error in the executable. (*M*) No drawing was computed because the memory is restricted to graphs with  $\leq 10,000$  nodes. (*T*) No drawing was computed within 10 hours of CPU time. *B* denotes the sets of biconnected components of the graphs. Best values are printed bold. Worst values are underlined.

The algorithm ACE is much faster than the force-directed algorithms for nearly all kind graphs. However, the running times grow extremely if ACE is used to draw several of the challenging graphs.

The linear time method HDE is by far the fastest algorithm. It needs less than 3.4 seconds for drawing even the largest tested graph.

## 4.2 Comparison of the Drawings

### 4.2.1 Uniformity of Edge Lengths

Figure 1 shows the normalized standard deviations of the edge lengths  $\sigma_\Gamma$  of the drawings  $\Gamma$  for the tested kind and challenging graphs.




Figure 1: The normalized standard deviations of the edge lengths of the drawings for the kind graphs (left) and challenging graphs (right).

For the kind graphs the  $\sigma_\Gamma$  values of the drawings that are computed by GVA, FM<sup>3</sup>, FMS, and HDE are below 1 while the  $\sigma_\Gamma$  values associated with ACE are still below 2. This indicates that for those graphs the goal of generating drawings with uniform edge lengths has been reached well by all algorithms.

The  $\sigma_\Gamma$  values of some drawings of challenging graphs that are generated by HDE and ACE exceed values of 5 and 25, respectively. In contrast to this, the  $\sigma_\Gamma$  values of GVA and FMS remain below 1 for the majority of the drawings of the challenging graphs. All  $\sigma_\Gamma$  values associated with drawings computed by FM<sup>3</sup> are smaller than 1.

Notice that it is not possible to obtain the measures  $\sigma_\Gamma$ ,  $recn_\Gamma$ , and  $reon_\Gamma$  for the drawings generated by GRIP, since the used implementation of GRIP provides screen output only. Hence, in this case we had to restrict ourselves to printing screenshots for visual comparisons.

### 4.2.2 Edge Crossings and Overlapping Edges

The relative edge-crossing numbers  $recn_\Gamma$  and the relative edge-overlapping numbers  $reon_\Gamma$  of the drawings  $\Gamma$  generated by the algorithms are listed in Tables 3 and 4, respectively.

For the kind graphs, the multilevel and algebraic methods generate comparable and significantly smaller numbers of edge crossings than the classical force-directed method GVA.

With two exceptions, the  $recn_\Gamma$  values of the drawings of the challenging graphs that are generated by FM<sup>3</sup> are smaller than the corresponding values of GVA. The  $recn_\Gamma$  values of ACE are much smaller than those of FM<sup>3</sup> for several challenging artificial graphs, while they are comparable with those of FM<sup>3</sup> for the

Type	Name	$recn_{\Gamma}$				
		GVA	FM <sup>3</sup>	FMS	ACE	HDE
Kind Artificial	rnd_grid_032	<u>3.82</u>	<b>0</b>	<b>0</b>	<b>0</b>	<0.01
	rnd_grid_100	<u>14.75</u>	<b>0</b>	<b>0</b>	<b>0</b>	<0.01
	rnd_grid_320	<u>181.51</u>	<b>0</b>	(N)	<0.01	<0.01
	sierpinski_06	<u>2.00</u>	0.05	<0.01	<b>0</b>	0.02
	sierpinski_08	<u>9.49</u>	0.07	<b>0.01</b>	0.02	0.08
	sierpinski_10	<u>99.97</u>	0.09	(N)	0.27	<b>0.01</b>
Kind Real World	crack	<u>30.82</u>	<0.01	(N)	<b>0</b>	0.07
	fe_pwt	<u>150.70</u>	2.45	(N)	(N)	<b>1.61</b>
	finan_512	<u>301.25</u>	18.81	(N)	<b>12.27</b>	21.27
	fe_ocean	<u>622.48</u>	<b>7.13</b>	(N)	9.07	8.24
Chal- lenging Arti- ficial	tree_06_04	<u>2.21</u>	1.16	7.89	0.01	<b>0</b>
	tree_06_05	<u>9.33</u>	1.89	11.48	<b>0</b>	<u>22.92</u>
	tree_06_06	<u>70.68</u>	<b>3.31</b>	(N)	4.16	<u>128.82</u>
	snowflake_A	<u>0.63</u>	<b>0</b>	0.10	<0.01	0.62
	snowflake_B	<u>1.46</u>	<b>0</b>	<u>8.18</u>	(N)	6.92
	snowflake_C	<u>15.53</u>	<b>0</b>	(N)	(N)	<u>195.87</u>
	spider_A	<u>15.62</u>	<u>16.55</u>	<b>1.17</b>	6.60	1.25
	spider_B	<u>154.70</u>	132.96	1.64	<b>0</b>	<b>0</b>
	spider_C	<u>2522.89</u>	1029.64	(N)	<b>0</b>	<b>0</b>
	flower_A	<u>46.71</u>	<u>49.08</u>	5.63	<b>0.26</b>	0.55
	flower_B	<u>64.90</u>	<u>51.57</u>	1.90	<b>0.06</b>	0.34
	flower_C	<u>578.22</u>	<u>53.39</u>	(N)	(N)	<b>0.30</b>
Chal- lenging Real World	ug_380	<u>22.93</u>	19.55	13.67	20.99	<b>1.35</b>
	esslingen	<u>47.52</u>	23.71	28.42	20.81	<b>3.89</b>
	add_32	<u>8.65</u>	1.69	5.75	<b>0.89</b>	5.80
	dg_1087	<u>1.74</u>	< <b>0.01</b>	<u>37.07</u>	5.92	6.49
	bcsstk_33	<u>720.94</u>	<u>376.18</u>	<u>4171.05</u>	413.56	<b>113.86</b>
bcsstk_31	<u>708.69</u>	<u>94.26</u>	(N)	<b>63.00</b>	611.21	

Table 3: The relative edge-crossing numbers ( $recn_{\Gamma}$ ) of the drawings  $\Gamma$  computed by the tested algorithms. The entry (N) indicates that no drawing was computed. Best values are printed bold. Worst values are underlined.

challenging real-world graphs. Depending on the classes of tested challenging graphs, the  $recn_{\Gamma}$  values of the drawings computed by FMS and HDE vary a lot: FMS and HDE generate many crossings for the 6-nary trees and the snowflake graphs but comparatively few crossings for the spider and flower graphs.

No drawing of a kind graph generated by GVA or FM<sup>3</sup> contains overlapping edges. In contrast to this, few pairs of overlapping edges exist in many drawings of kind graphs computed by FMS and ACE, and in all drawings computed by HDE.

The challenging graphs spider\_A and esslingen are the only graphs in the test set that contain parallel edges, resulting in minimum relative edge-overlapping

Type	Name	$reon_{\Gamma}$				
		GVA	FM <sup>3</sup>	FMS	ACE	HDE
Kind Arti- ficial	rnd_grid_032	<b>0</b>	<b>0</b>	<u>0.03</u>	<b>0</b>	<0.01
	rnd_grid_100	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<u>0.03</u>
	rnd_grid_320	<b>0</b>	<b>0</b>	( <i>N</i> )	<b>0</b>	<u>0.08</u>
	sierpinski_06	<b>0</b>	<b>0</b>	<u>0.03</u>	<b>0</b>	0.01
	sierpinski_08	<b>0</b>	<b>0</b>	<u>0.14</u>	<0.01	0.11
	sierpinski_10	<b>0</b>	<b>0</b>	( <i>N</i> )	0.98	<u>1.43</u>
Kind Real World	crack	<b>0</b>	<b>0</b>	( <i>N</i> )	<b>0</b>	<u>0.12</u>
	fe_pwt	<b>0</b>	<b>0</b>	( <i>N</i> )	( <i>N</i> )	<u>34.25</u>
	finan_512	<b>0</b>	<b>0</b>	( <i>N</i> )	3.68	<u>9.14</u>
	fe_ocean	<b>0</b>	<b>0</b>	( <i>N</i> )	0.22	<u>1.67</u>
Chal- lenging Arti- ficial	tree_06_04	<b>0</b>	<b>0</b>	0.38	59.09	<u>78.08</u>
	tree_06_05	<b>0</b>	<b>0</b>	0.86	<u>504.27</u>	466.83
	tree_06_06	<b>0</b>	<b>0</b>	( <i>N</i> )	<u>2929.74</u>	2030.40
	snowflake_A	<b>0</b>	<b>0</b>	42.02	32.11	<u>85.34</u>
	snowflake_B	<b>0</b>	<b>0</b>	169.77	( <i>N</i> )	<u>398.06</u>
	snowflake_C	<b>0</b>	<b>0</b>	( <i>N</i> )	( <i>N</i> )	<u>16818.82</u>
	spider_A	<b>0.27</b>	<b>0.27</b>	0.67	19.57	<u>63.50</u>
	spider_B	<b>0</b>	<b>0</b>	<u>6903.67</u>	4248.23	297.13
	spider_C	<b>0</b>	<b>0</b>	( <i>N</i> )	<u>44554.09</u>	40163.33
	flower_A	<b>0</b>	<b>0</b>	5.14	149.31	<u>160.36</u>
flower_B	<b>0</b>	<b>0</b>	21.63	167.48	<u>210.83</u>	
flower_C	<b>0</b>	<b>0</b>	( <i>N</i> )	( <i>N</i> )	<u>340.19</u>	
Chal- lenging Real World	ug_380	<b>0</b>	<b>0</b>	1.29	0.79	<u>10.69</u>
	esslingen	<b>0.14</b>	<b>0.14</b>	0.97	1.93	<u>2.54</u>
	add_32	<b>0</b>	<b>0</b>	1.21	0.19	<u>5.73</u>
	dg_1087	<b>0</b>	<b>0</b>	109.23	<u>1882.14</u>	1390.30
	bcsstk_33	<b>0</b>	<b>0</b>	<u>3416.70</u>	0.13	329.17
bcsstk_31	<b>0</b>	<b>0</b>	( <i>N</i> )	2.50	<u>2386.76</u>	

Table 4: The relative edge-overlapping numbers ( $reon_{\Gamma}$ ) of the drawings  $\Gamma$  computed by the tested algorithms. The entry (*N*) indicates that no drawing was computed. Best values are printed bold. Worst values are underlined.

numbers of 0.27 and 0.14, respectively. Hence, GVA and FM<sup>3</sup> are the only algorithms that generate drawings with the minimum number of overlapping edges for all tested graphs. In contrast to this, the  $recn_{\Gamma}$  values of all drawings of challenging graphs generated by FMS, ACE, and HDE are positive and reach extremely high values in many cases.

For FMS an explanation of this behavior might be that the positions have integer values only, and that the underlying grid-resolution is too small. For the algebraic methods an explanation of the many overlapping nodes and edges can be found in [3, 17] and [11].

Since overlapping edges are at least as undesirable as crossing edges we additionally compared the sums of the relative edge-crossing numbers and the dedicated relative edge-overlapping numbers (see Figure 2).




Figure 2: The sums of the relative edge-crossing and dedicated relative edge-overlapping numbers of the drawings for the tested kind graphs (left) and challenging graphs (right).

For the kind graphs, the  $recn_{\Gamma}$  values of the drawings computed by **GVA** are significantly larger than the sum of  $recn_{\Gamma}$  and  $reon_{\Gamma}$  of all other algorithms. For those graphs  $FM^3$  and  $ACE$  have smaller values than  $HDE$ .  $FM^3$  reaches the lowest added values of  $recn_{\Gamma}$  and  $reon_{\Gamma}$  for the majority of the challenging graphs. The other multilevel and algebraic methods have frequently higher added  $recn_{\Gamma}$  and  $reon_{\Gamma}$  values than the classical force-directed algorithm **GVA**.

### 4.2.3 The Overall Picture

In the remainder of this section, we will discuss the quality of the drawings of the tested graphs that are presented in Figures 3 to 10 by keeping the modeled esthetic criteria in mind.

Note that due to the very restricted drawing area, not all details of each large graphs can be displayed, here. However, those details can be examined by creating large printouts or by using simple zoom and pan techniques.

For all kind graphs the classical method **GVA** does not untangle the drawings that were induced by the random initial placements. In contrast to this, nearly all algorithms computed relatively pleasing drawings of the kind graphs (see Figure 3, Figure 4, and Figure 5(a)-(d)).

None of the drawings of the complete 6-nary trees (see Figure 5(e)-(j)) is really convincing, since the algorithms either produce many unnecessary edge crossings or they place many nodes at the same coordinates.

Except  $FM^3$  none of the tested algorithms displays the global structure of the snowflake graphs. Even the drawings of the smallest snowflake graph (see Figure 6(a)-(f)) leave room for improvement. However, **GVA** and **GRIP** visualize parts of its structure in an appropriate way.

The drawings of the spider\_A graph (see Figure 6(g)-(l)) that are generated by GRIP, FMS, and HDE are not as symmetric as the drawing computed by FM<sup>3</sup>. But they display the global structure of the graph. The drawing generated by GVA shows the dense subregion, but GVA does not untangle the 8 paths. The paths in the drawing of ACE are not displayed in the same length. The drawings of the larger spider graphs are of comparable quality.

The drawings of the flower\_B graph (see Figure 7(a)-(f)) that are computed by FMS and HDE display the global structure of the graph but the symmetries are not as clear as in the drawing generated by FM<sup>3</sup>. The drawings of the other flower graphs are of comparable quality.

We concentrate on the challenging real-world graphs now. The graphs ug\_380 and dg\_1087 both contain one node with a very high degree. Furthermore, dg\_1087 has many biconnected components, since it is a tree. Only the drawings that are computed by GVA, FM<sup>3</sup>, and GRIP (see Figure 7(g)-(l) and Figure 8(a)-(f)) clearly display the central regions of these graphs.

The social network esslingen (see Figure 8(g)-(l)) consists of two big well-connected subgraphs. This can be visualized by FM<sup>3</sup>, GRIP, and HDE. But the drawings contain several edge crossings.

Since add\_32 that describes a 32 bit adder contains many biconnected components, we expect that the drawings have a tree-like shape. This structure is visualized by GVA, FM<sup>3</sup>, GRIP, and ACE (see Figure 9(a)-(f)). The drawings of GVA and GRIP contain comparatively many edge crossings, while the drawing of ACE displays the global structure, but hides local details.

Finally, we discuss the drawings of the dense graphs bcsttk\_31 and bcsttk\_33. The drawings of bcsttk\_33 (see Figure 9(g)-(l)) that are generated by FM<sup>3</sup>, GRIP, and ACE are comparable and visualize the regular structure of the graph. The car body that is modeled by the graph bcsttk\_31 (see Figure 10) is visualized by FM<sup>3</sup> and ACE only.

## 5 Conclusion

We can summarize that only GVA, FM<sup>3</sup>, and HDE generate drawings of all tested graphs. The force-directed multilevel methods and the algebraic methods are – except the methods FMS and ACE for some graphs – much faster than the comparatively slow classical algorithm GVA. HDE, FM<sup>3</sup> and GRIP scale well on all tested graphs. FM<sup>3</sup> needs a few minutes to draw the largest graphs. GRIP is up to factor 9 faster than FM<sup>3</sup> but it could not be tested on the largest graphs. All tested methods are much slower than HDE that needs only few seconds to draw even the largest graphs.

As expected, all algorithms, except GVA, generate pleasing drawings of the kind graphs with relatively uniform edge length and few edge crossings. In contrast to this, the quality of the computed drawings varies a lot depending on the structures of the tested challenging graphs. In particular, FMS, HDE, and ACE frequently generate drawings with many overlapping edges. Unlike this, FM<sup>3</sup> generates pleasing drawings for the majority of the challenging graphs. But

there still remain classes of tested graphs (e.g., the complete trees and the social network graph esslingen) for which the drawing quality of all tested algorithms leaves much room for improvement.

A practical advice for developers of graph-drawing systems is as follows: First use HDE followed by ACE, since they are the fastest methods in all or many cases, respectively. If the drawings are not satisfactory or one supposes that important details of the graph's structure are hidden, use FM<sup>3</sup> to obtain comparable or better results in reasonable time. If those drawings are still not nice, one should experiment with other methods or try to improve the existing ones with new concepts.

## Acknowledgments

We would like to thank David Alberts, Steven Kobourov, Yehuda Koren, Stefan Näher, and Roman Yusufov for making the implementations of their algorithms available to us. We thank Ulrik Brandes, Carola Lipp and Chris Walshaw for the access to the real-world test graphs.

## References

- [1] The AT&T graph collection: [www.graphdrawing.org](http://www.graphdrawing.org).
- [2] F. Brandenburg, M. Himsolt, and C. Rohrer. An Experimental Comparison of Force-Directed and Randomized Graph Drawing Methods. In *Graph Drawing 1995*, volume 1027 of *LNCS*, pages 76–87. Springer-Verlag, 1996.
- [3] U. Brandes and D. Wagner. In *Graph Drawing Software*, volume XII of *Mathematics and Visualization*, chapter visone - Analysis and Visualization of Social Networks, pages 321–340. Springer-Verlag, 2004.
- [4] R. Davidson and D. Harel. Drawing Graphs Nicely Using Simulated Annealing. *ACM Transactions on Graphics*, 15(4):301–331, 1996.
- [5] P. Eades. A heuristic for graph drawing. *Congressus Numerantium*, 42:149–160, 1984.
- [6] A. Frick, A. Ludwig, and H. Mehldau. A Fast Adaptive Layout Algorithm for Undirected Graphs. In *Graph Drawing 1994*, volume 894 of *LNCS*, pages 388–403. Springer-Verlag, 1995.
- [7] T. Fruchterman and E. Reingold. Graph Drawing by Force-directed Placement. *Software–Practice and Experience*, 21(11):1129–1164, 1991.
- [8] P. Gajer, M. Goodrich, and S. Kobourov. A Multi-dimensional Approach to Force-Directed Layouts of Large Graphs. In *Graph Drawing 2000*, volume 1984 of *LNCS*, pages 211–221. Springer-Verlag, 2001.
- [9] P. Gajer and S. Kobourov. GRIP: Graph Drawing with Intelligent Placement. In *Graph Drawing 2000*, volume 1984 of *LNCS*, pages 222–228. Springer-Verlag, 2001.
- [10] T. Gonzalez. Clustering to Minimize the Maximum Inter-Cluster Distance. *Theoretical Computer Science*, 38:293–306, 1985.
- [11] S. Hachul. *A Potential-Field-Based Multilevel Algorithm for Drawing Large Graphs*. PhD thesis, Institut für Informatik, Universität zu Köln, Germany, 2005. <http://kups.ub.uni-koeln.de/volltexte/2005/1409>.
- [12] S. Hachul and M. Jünger. Drawing Large Graphs with a Potential-Field-Based Multilevel Algorithm (Extended Abstract). In *Graph Drawing 2004*, volume 3383 of *Lecture Notes in Computer Science*, pages 285–295. Springer-Verlag, 2005.
- [13] D. Harel and Y. Koren. A Fast Multi-scale Method for Drawing Large Graphs. In *Graph Drawing 2000*, volume 1984 of *LNCS*, pages 183–196. Springer-Verlag, 2001.



- [14] D. Harel and Y. Koren. Graph Drawing by High-Dimensional Embedding. In *Graph Drawing 2002*, volume 2528 of *LNCS*, pages 207–219. Springer-Verlag, 2002.
- [15] M. Jünger, G. Klau, P. Mutzel, and R. Weiskircher. In *Graph Drawing Software*, volume XII of *Mathematics and Visualization*, chapter AGD - A Library of Algorithms for Graph Drawing, pages 149–172. Springer-Verlag, 2004.
- [16] T. Kamada and S. Kawai. An Algorithm for Drawing General Undirected Graphs. *Information Processing Letters*, 31:7–15, 1989.
- [17] Y. Koren, L. Carmel, and D. Harel. Drawing Huge Graphs by Algebraic Multigrid Optimization. *Multiscale Modeling and Simulation*, 1(4):645–673, 2003.
- [18] Y. Koren’s algorithms: [research.att.com/~yehuda/index\\_programs.html](http://research.att.com/~yehuda/index_programs.html).
- [19] J. Manning. Computational complexity of geometric symmetry detection in graphs. In *Great Lakes Computer Science Conference*, volume 507 of *Lecture Notes in Computer Science*, pages 1–7. Springer-Verlag, 1990.
- [20] A. Quigley and P. Eades. FADE: Graph Drawing, Clustering, and Visual Abstraction. In *Graph Drawing 2000*, volume 1984 of *LNCS*, pages 197–210. Springer-Verlag, 2001.
- [21] D. Tunkelang. JIGGLE: Java Interactive Graph Layout Environment. In *Graph Drawing 1998*, volume 1547 of *LNCS*, pages 413–422. Springer-Verlag, 1998.
- [22] C. Walshaw’s graph collection: [staffweb.cms.gre.ac.uk/~c.walshaw/partition](http://staffweb.cms.gre.ac.uk/~c.walshaw/partition).
- [23] C. Walshaw. A Multilevel Algorithm for Force-Directed Graph Drawing. In *Graph Drawing 2000*, volume 1984 of *LNCS*, pages 171–182. Springer-Verlag, 2001.
- [24] R. Yusufov’s implementation of GRIP: [www.cs.arizona.edu/~kobourov/GRIP](http://www.cs.arizona.edu/~kobourov/GRIP).

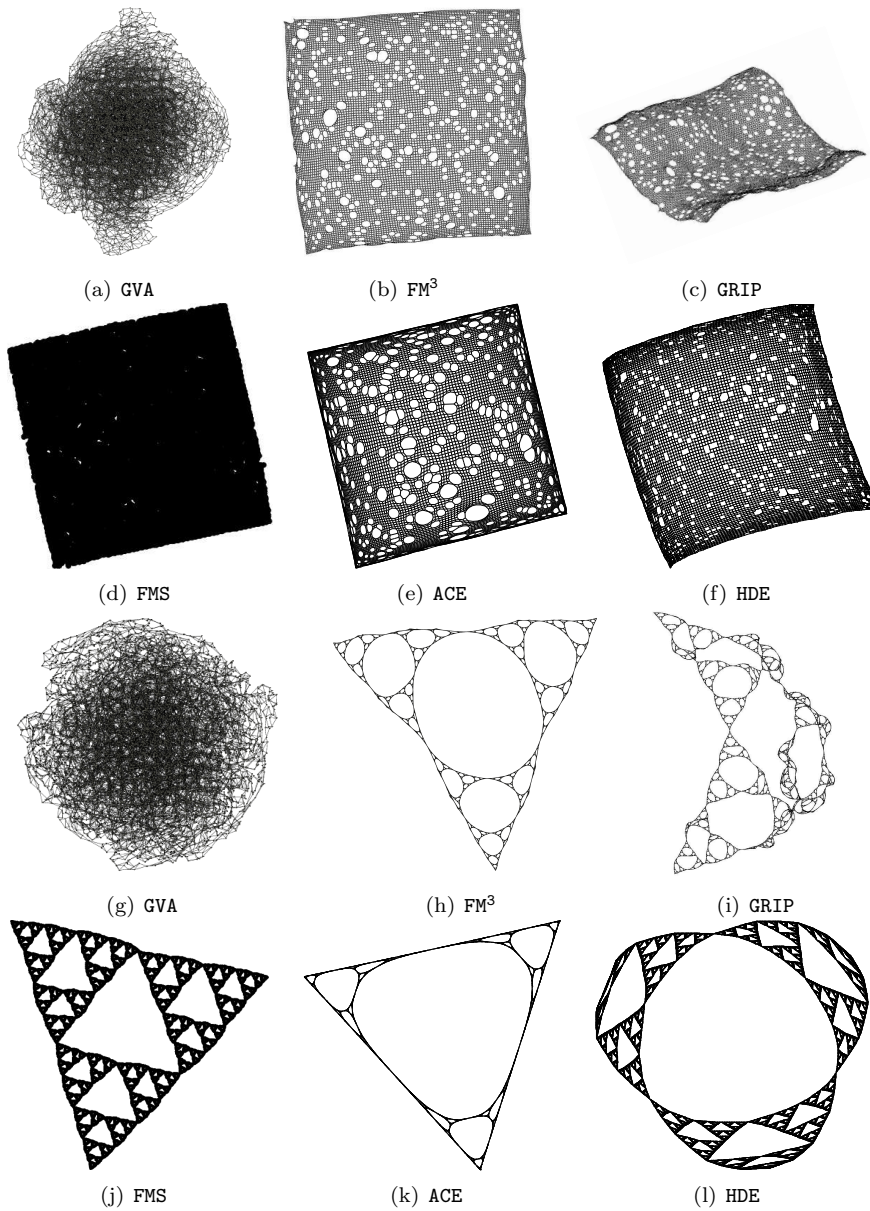


Figure 3: (a)-(f) Drawings of *rnd\_grid\_100* and (g)-(l) *sierpinski\_08* generated by different algorithms.

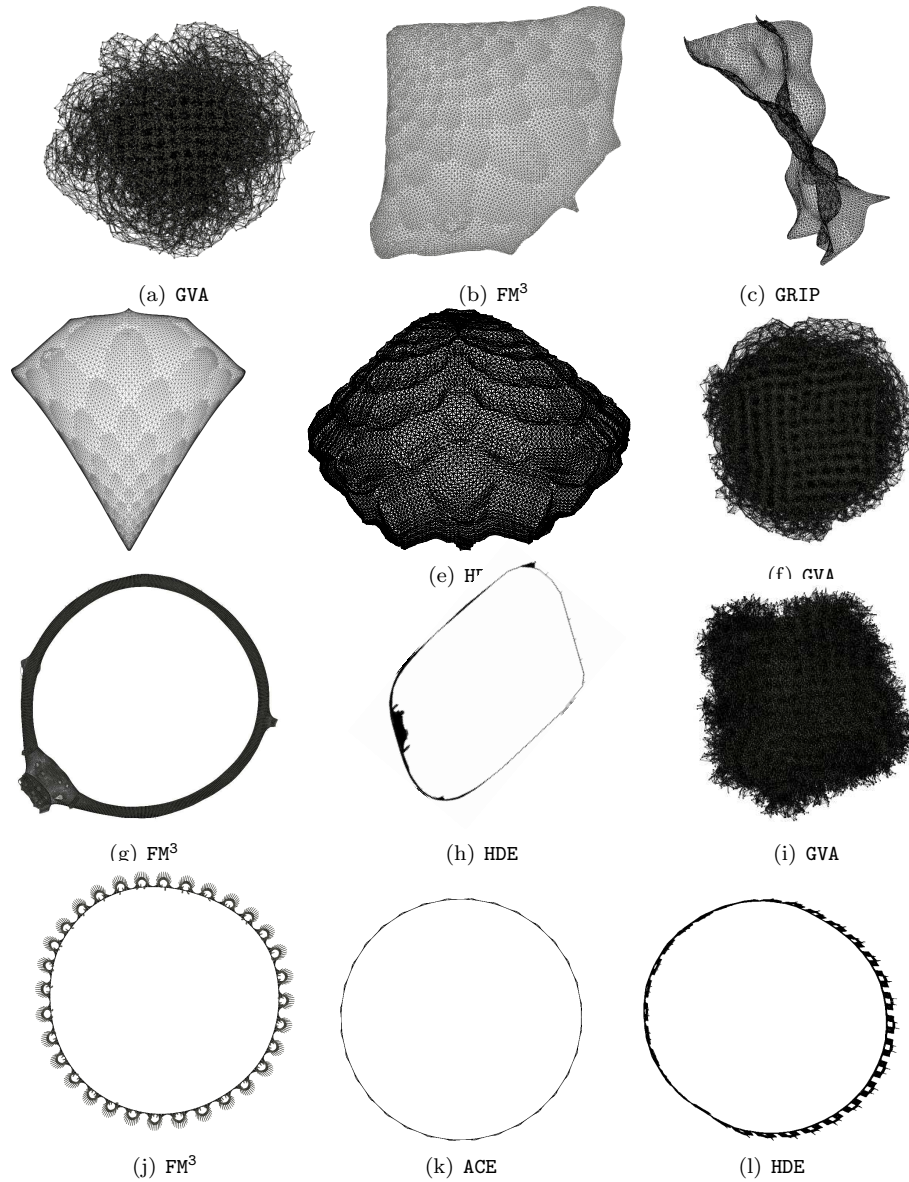


Figure 4: (a)-(e) Drawings of crack, (f)-(h) fe\_pwt, and (i)-(l) finan\_512 generated by different algorithms.

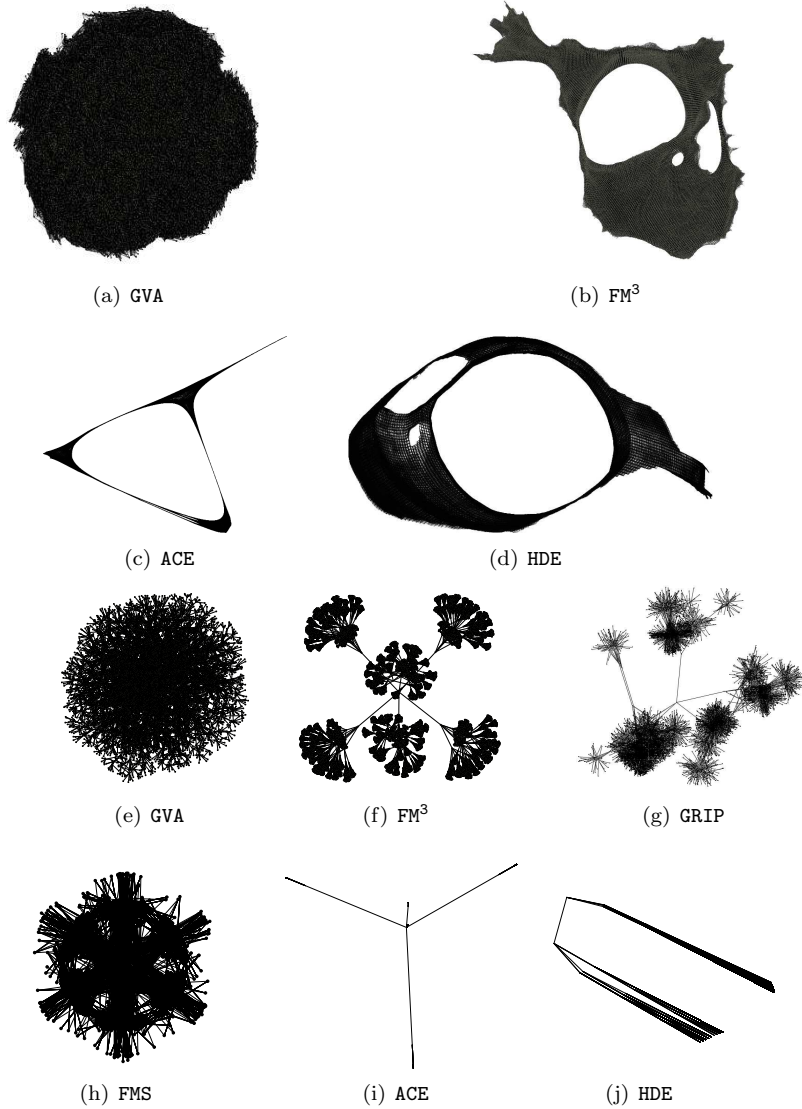


Figure 5: (a)-(d) Drawings of fe.ocean and (e)-(j) tree.06\_05 generated by different algorithms.

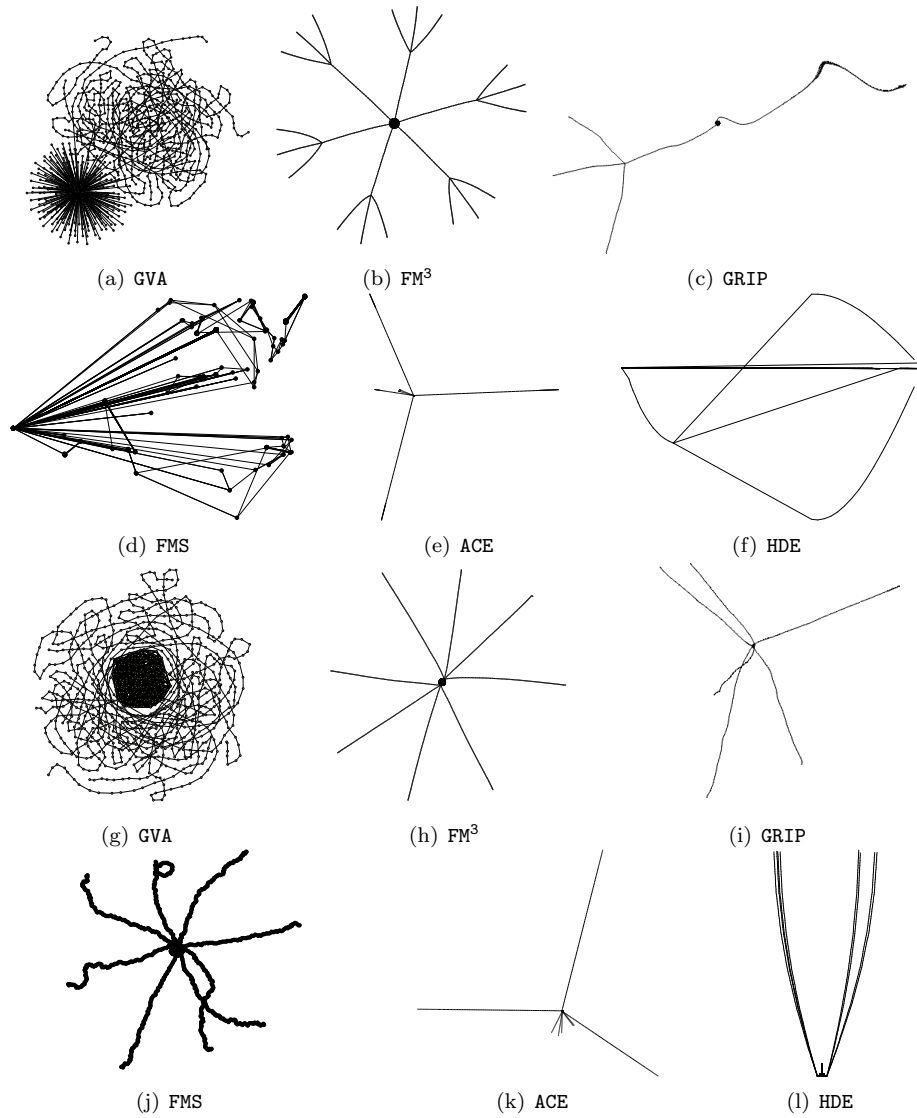


Figure 6: (a)-(f) Drawings of snowflake\_A and (g)-(l) spider\_A generated by different algorithms.

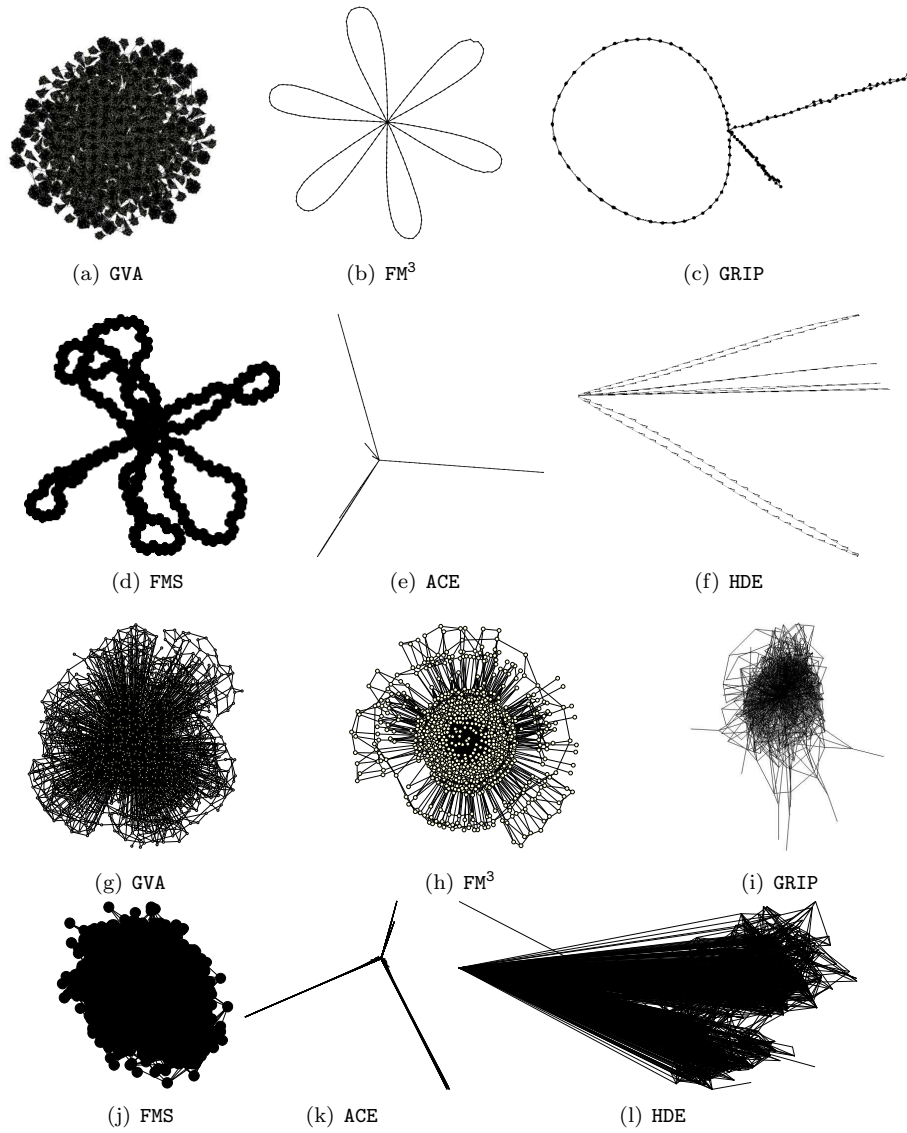


Figure 7: (a)-(f) Drawings of *flower\_B* and (g)-(l) *ug\_380* generated by different algorithms.

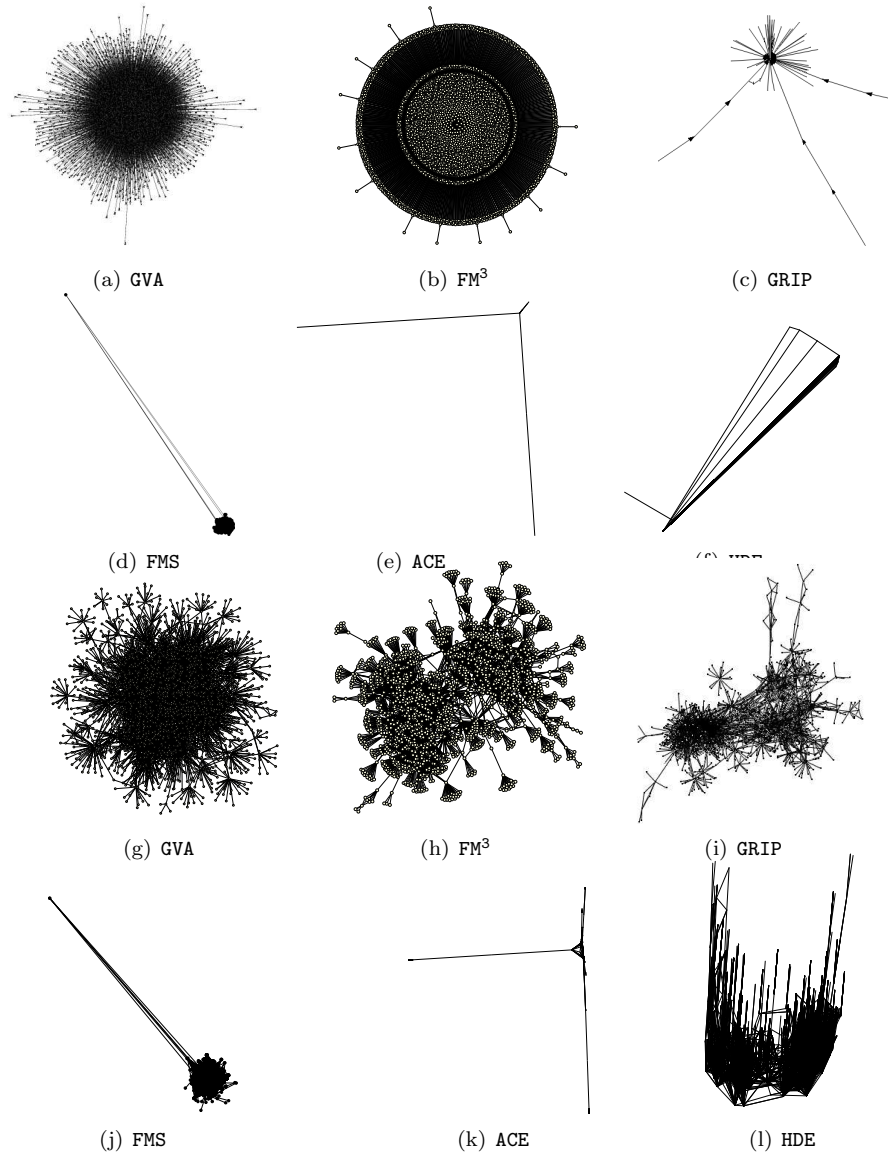


Figure 8: (a)-(f) Drawings of dg\_1087 and (g)-(l) esslingen generated by different algorithms

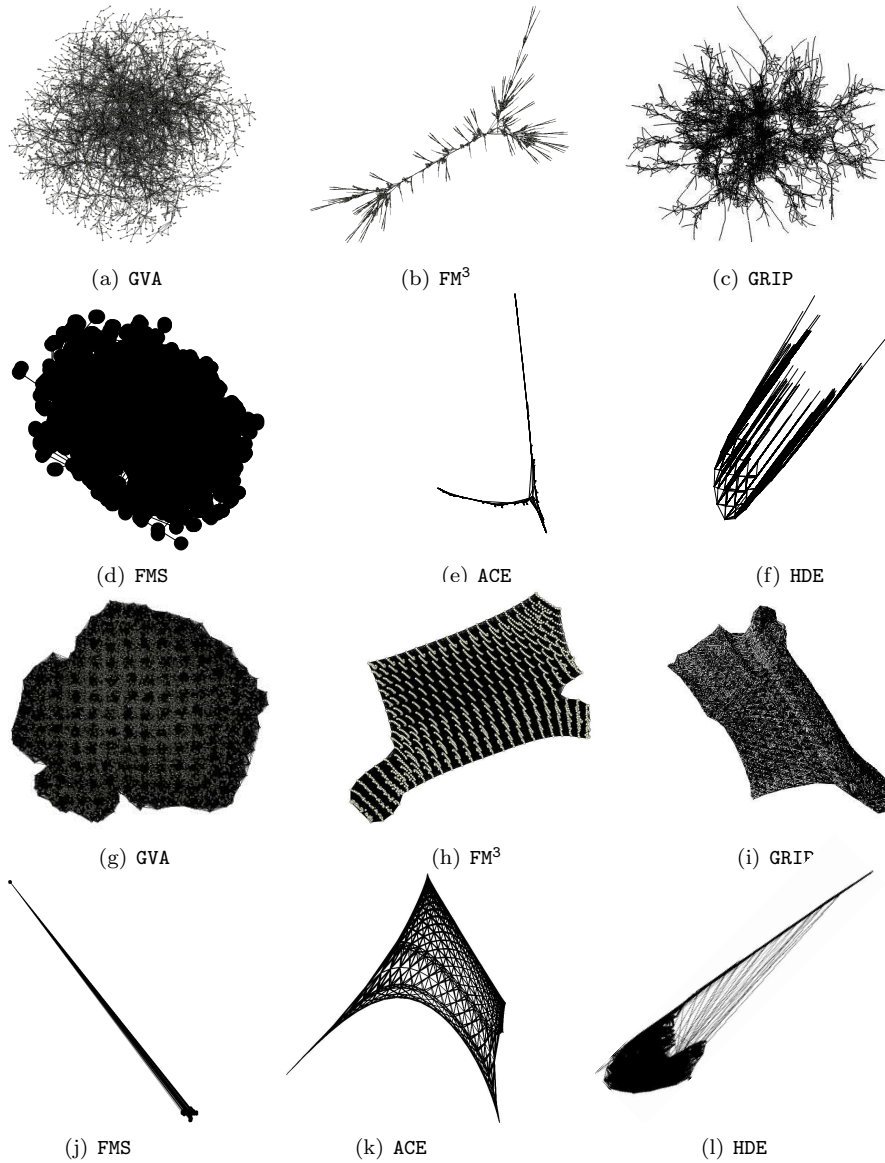


Figure 9: (a)-(f) Drawings of add\_32 and (g)-(l) bcsstk\_33 generated by different algorithms.






Figure 10: (a)-(d) Drawings of *bcsstk\_31* generated by different algorithms.