This paper appeared in: Discrete Applied Mathematics 155 (2007), 2539-2545

Computation of Best Possible Low Degree Expanders

Stefan Hougardy

Forschungsinstitut für Diskrete Mathematik
Universität Bonn
Lennéstr. 2, 53113 Bonn, GERMANY
hougardy@or.uni-bonn.de

Ivo Köthnig

Institut für Informatik
Humboldt-Universität zu Berlin
10099 Berlin, GERMANY
koethnig@informatik.hu-berlin.de

Abstract. We present an algorithm for computing a best possible bipartite cubic expander for a given number of vertices. Such graphs are needed in many applications and are also the basis for many results in theoretical computer science. Known construction methods for expander graphs yield expanders that have a fairly poor expansion compared to the best possible expansion. Our algorithm is based on a lemma which allows to calculate an upper bound for the expansion of cubic bipartite graphs.

Keywords. expander graph, orderly algorithm

1 Introduction

An expander is a graph that is sparse (i.e. it has only few edges) and highly connected, in the sense that the neigborhood of each set of vertices has large cardinality. Expanders are needed for many applications and are also used for many results in theoretical computer science, for example network design [Pip87, PY82], complexity theory [Val77, Sip88], derandomization [IW97], coding theory [SS96], cryptography [GIL+90], and non-approximability results [Pap94].

In this paper we consider regular bipartite expanders, i.e. bipartite graphs in which all vertices have the same degree. A bipartite graph $G = (A \cup B, E)$ is said to have expansion c for A if for every subset $S \subseteq A$ with $|S| \leq |A|/2$ the neighborhood $N(S)$ of S has size at least $c \cdot |S|$. The goal is to find k-regular bipartite expanders on n vertices with highest possible value for the expansion c.

It is rather easy to prove by a probabilistic argument the existence of k-regular expanders with an expansion strictly larger than 1 [Pin73]. For many applications explicit constructions of such expanders are needed. Several such constructions are known [Mar73, GG81, RVW01] but they only produce expanders with a fairly small expansion compared to the best possible expansion. This also holds for the expander graphs constructed in [LPS88] that meet the eigenvalue bound and are in this sense best possible. In this paper we present an algorithm which computes either one or all best possible k-regular bipartite expanders for a given number of vertices.

This paper appeared in: Discrete Applied Mathematics 155 (2007), 2539-2545

The simple brute force approach is to enumerate all unlabeled k-regular bipartite graphs and afterwards check their expansion. In our algorithm the expansion of a graph is already checked when only some part of it has been generated. This allows us to skip in a very early stage the enumeration of graphs which can not beat the best graph which is already computed. The heart of our algorithm is a lemma that allows the prediction of the best possible expansion for subgraphs of cubic graphs.

This way the algorithm generates much fewer graphs than a brute force approach would require. Our approach allows to generate a best possible cubic expander on up to 48 vertices, while a brute force approach is limited to 34 vertices.

2 The Algorithm

The main idea of our algorithm is to enumerate for a given number n of vertices and degree k only those k-regular bipartite graphs on $2n$ vertices that are candidates for having the best possible expansion. There exist very fast algorithms for enumerating all k-regular bipartite graphs on $2n$ vertices [Mer96]. As the number of such graphs grows extremly fast, it is not feasable to first generate all bipartite k-regular graphs on $2n$ vertices and then to choose one having the largest expansion. This approach is especially not feasable as computing the expansion of a graph is a co-NP-hard problem [BKV+81].

FIND-BEST-EXPANDER

```
VALID(G)
  RETURN
    regularity(G) AND
    maximality(G) AND
    expansion(G)

DFS(i)
  IF (i = n)
    expansion := expansion of G
  ELSE
    FOR EACH v_1 < ... < v_k ∈ B
      G := G + {{i, v_1}, ..., {i, v_k}}
      IF (VALID(G))
        DFS(i + 1)
      G := G − {{i, v_1}, ..., {i, v_k}}

MAIN(n)
  A := {0, ..., n − 1}
  B := {n, ..., 2n − 1}
  G := (A ∪ B, ∅)
  expansion := 1
  DFS(0)
```

Figure 2.1: Outline of an algorithm for computing expanders.

This paper appeared in: Discrete Applied Mathematics 155 (2007), 2539-2545

Our approach uses an orderly algorithm that starts from the empty graph and extends the set of currently produced subgraphs. While generating these subgraphs we compute the expansion of all their completed subsets (i.e. those in which all vertices already have degree k) and remove all subgraphs that cannot yield expanders that are better than the best expander currently found.

To be more precise the algorithm works as shown in Figure 2.1. It starts with the procedure MAIN which generates an empty graph on $2n$ vertices. The vertices 0 to $n-1$ belong to the bipartition set A and the vertices n to $2n-1$ belong to the bipartition set B of the graph. Initially we assume that the expansion is at least 1, which is justified by Hall's theorem.

Next we call a recursive procedure DFS. The parameter of DFS determines for which vertex of A all neighbors have to be computed as next. All possible supergraphs will be checked by the function VALID, which tests if the resulting graph is

- a subgraph of a regular graph,
- in maximal representation and
- has the expansion property.

Since expansion is invariant up to isomorphism we only generate graphs in lexicographic maximal representation, i.e. for all permutations of $\{0,\ldots,n-1\}$ and $\{n,\ldots,2n-1\}$ we check if the corresponding adjacency matrix is lexicographically smaller than the adjacency matrix of the current graph. For this we use some techniques similar to those in [Mer96]. Details of the algorithm for this test are described in [KO02].

To check if the graph has the expansion property we test for all subsets of $\{0,\ldots,i\}$ containing i, that the size of the neighborhood is not to small. Since we have already tested all subsets of $\{0,\ldots,i\}$ not containing i in the previous iterations we do not need to check them again.

It follows directly from the definition, that it is not necessary to check the subsets of B. But there are good reasons doing this. If we check the size of the neighborhood of subsets of B, we can bound the size of the neighborhood of subsets of A even if we do not know their neighborhoods since it may not have been generated so far. For this we use the following lemma, which is easily seen to be true.

Lemma 1 *Let $G=(A\cup B,E)$ be a regular bipartite graph, R a subset of B and S a subset of A that contains the neighborhood of R in G. Furthermore let $P=B\setminus R$ and $Q=A\setminus S$. Then the neighborhood of Q is included in P.* □

This lemma is applied as follows. Assume we know that a subset R of B exists with size x and there is a subset S of A with size y containing the neighborhood of R. Then we know, that a subset Q of A with size $n-y$ exists and the size of its neighborhood is bounded by $n-x$ (n is the size of A and B).

The test of the expansion of the subsets of B works like the test of the expansion of the subsets of A. It is just a little bit more difficult to implement since the vertices of B already having degree k and thus which subsets are allowed to be tested must be known.

This paper appeared in: Discrete Applied Mathematics 155 (2007), 2539-2545

3 Calculating an Upper Bound for the Expansion of Supersets

In this section we present a theorem which shows, that for a given subset one can estimate an upper bound for the size of the neighborhood of some of its small supersets (i.e supersets with just a few more vertices). Notice, that we do not know for which superset we can bound the size of the neighborhood (we just know the size of the superset). Furthermore it is not necessary to explictly know the subset from which we derive the existence of a superset with small neighborhood. We just need its size and a bound of its neighborhood. Thus we can recursivly apply the theorem to bound the size of the neigborhood of some arbitrary large supersets. To prove the theorem we first present some necessary lemmas.

Lemma 2 *Let $G = (A \cup B, E)$ a connected bipartite graph, $n = |A|$, such that each vertex of A has at most degree k. Then for each $x \in \{1, \ldots, n\}$ there exists an $X \subseteq A$ with*

$$|N(X)| \leq x(k-1) + 1.$$

Proof. We proof this by induction on x. Let $x = 1$. We choose an arbitrary vertex $v \in A$. Vertex v has degree at most k and thus the lemma holds. Now let $n \geq x > 1$. By induction hypothesis there is a set $Y \subseteq A$ of size $x-1$ with $|N(Y)| \leq (x-1)(k-1)+1$. Since the graph is connected there are vertices in $A \setminus Y$, which belong to the neighborhood of $N(Y)$. Thus, we can choose an arbitrary vertex in $A \setminus Y$, which is adjacent to $N(Y)$. This vertex has at most $k-1$ neighbors in $B \setminus N(Y)$. For $X = Y \cup \{v\}$ it follows $|X| = x$ and

$$\begin{aligned} N(X) &\leq (x-1)(k-1) + 1 + k - 1 \\ &= x(k-1) + 1. \end{aligned}$$

□

Corollary 3 *Let $G = (A \cup B, E)$ be a connected bipartite graph, $n = |A|$, such that each vertex of A has at most degree 2. Then there is for each $x \in \{1, \ldots, n\}$ an $X \subseteq A$ with $|N(X)| \leq x + 1$.*

Lemma 4 *Let $G = (A \cup B, E)$ be a connected bipartite graph, $n = |A|$, such that each vertex of A has at most degree 2. Then the following hold:*

$$\begin{aligned} |B| &= |A| + 1, \textit{ if } d(v) = 2 \textit{ for each } v \in A \textit{ and } G \textit{ is a tree, and} \\ |B| &< |A| + 1, \textit{ else.} \end{aligned}$$

Proof. Since G is connected it follows $|A|+|B|-1 = |E|$, if G is a tree and $|A|+|B|-1 < E$ otherwise. Since each edge is incident with a vertex in A and a vertex in B and each vertex in $|A|$ is incident with at most 2 edges it follows $|E| = 2|A|$, if $d(v) = 2$ for each $v \in A$ and $|E| < 2|A|$ otherwise. Together:

$$\begin{aligned} |A| + |B| - 1 &= 2|A|, \text{ if } d(v) = 2 \text{ for each } v \in A \text{ and } G \text{ is a tree, and} \\ |A| + |B| - 1 &< 2|A|, \text{ else.} \end{aligned}$$

Adding $1 - |A|$ on both sides gives the result. □

This paper appeared in: Discrete Applied Mathematics 155 (2007), 2539-2545

Lemma 5 *Let a, b, x and y be natural numbers, $y \geq x$ and $a > x$. If $b \cdot y < a\,(y+1)$ then:*

$$(b-x)\,(y-x) < (a-x)\,(y-x+1)\,.$$

Proof. We concider two cases.
Case 1: Let $b \geq a+1$. Then it follows:

$$\begin{aligned}(b-x)\,(y-x) &= by - bx - xy + x^2 \\ &< a\,(y+1) - (a+1)\,x - xy - x^2 \\ &= ay - ax + a - xy + x^2 - x \\ &= (a-x)\,(y-x+1)\,.\end{aligned}$$

Case 2: Let $b \leq a$. Then also $b - x \leq a - x$ holds. Since $y \geq x$ the term $y - x$ is not negative and it follows:

$$(b-x)\,(y-x) \quad \leq \quad (a-x)\,(y-x)\,.$$

Since $a > x$ the term $a - x$ is positive and it follows:

$$\begin{aligned}(b-x)\,(y-x) &\leq (a-x)\,(y-x) \\ &< (a-x)\,(y-x) + a - x \\ &= (a-x)\,(y-x+1)\,.\end{aligned}$$

□

We now can proof the following theorem.

Theorem 6 *Let $G = (A \cup B, E)$ be a connected 3-regular bipartite graph, $R \subseteq A$, $S \subseteq B$ with $N(R) \subseteq S$. Furthermore let $P = A \setminus R$ and $Q = B \setminus S$. If $y+1 \leq |P|$ and*

$$|Q|\,y < 3\,(|S| - |R|)\,(y+1)\,,$$

then there exists a subset $X \subseteq P$ with $|X| = x$ and $|N(X) \cap Q| \leq x+1$ for each $x \in \{1, \ldots, y+1\}$.

Proof. We consider the set $P' = N(S) \cap P$ and the induced subgraph $G' = G[P' \cup Q]$ of G. The proof works by induction on y. Since a vertex in P' has at most two neighbors in Q, the result can be seen easily for the case $y = 0$.

Let now $y > 0$. We first assume, that the vertices of P' have degree exactly two in G' and all components of G' are trees. Using Corollary 3 we only need to show, that there exists a component in G' containing $y+1$ vertices of P'. By our assumption each component K of G' fulfills the premises of Proposition 4, such that $|B_K| = |A_K| + 1$, where A_K is the set of vertices in P' of component K and B_K ist the set of vertices in Q of component K. This implies that $k = |Q| - |P'|$ is the number of components in G'. Since S is adjacent with exactly $3\,|S|$ edges and R is adjacent with exactly $3\,|R|$ edges, which all connect to vertices in S, there exist $3\,(|S| - |R|)$ edges between S and P'. Since all vertices in P' have degree 2 in G' two different edges between S and P' can not be incident with the same vertex of P'.

This paper appeared in: Discrete Applied Mathematics 155 (2007), 2539-2545

Thus it follows $|P'| = 3(|S| - |R|)$. Assume now that each component of G' contains at most y vertices in P'. Then we get the following contradiction:

$$\begin{aligned} |P'| &\leq ky \\ &= (|Q| - |P'|)\, y \\ &= |Q|\, y - |P'|\, y \\ &< 3(|S| - |R|)(y+1) - |P'|\, y \\ &= |P'|(y+1) - |P'|\, y \\ &= |P'|. \end{aligned}$$

We now study the case, that there either exists a vertex of P' which has degree at most one in G' or there exist components of G' which are not trees. We choose an arbitrary component K of G' containing either a vertex of P' that has degree at most one in G' or is not a tree. If $|A_K| \geq y+1$ the result follows from Corollary 3. Otherwise we have to show, the result for each x with $|A_K| < x \leq y+1$.

From Proposition 4 it follows $|B_K| \leq |A_K|$. Let $B'_K \subseteq Q$, such that $|B'_K| = |A_K|$ and $B_K \subseteq B'_K$. We now consider the sets $R^+ = R \cup A_K$, $S^+ = S \cup B'_K$, $P^+ = A \setminus A_K = P \setminus A_K$ and $Q^+ = Q \setminus B'_K = Q \setminus B'_K$. Let $y^+ = y - |A_k|$. We just need to show now, that for each $x \in \{1, \ldots, y^+\}$ there exists a $X \subseteq P^+$ with $|X| = x$ and at most $x+1$ neighbors in Q^+. Since $y^+ < y$, $|P^+| = |P| - |A_K| \geq y + 1 - |A_K| = y^+ + 1$ and

$$\begin{aligned} |Q^+|\, y^+ &= (|Q| - x^+)(y - x^+) \\ &< (3(|S| - |R|) - x^+)(y - x^+ + 1) \\ &< 3(|S| - |R|)(y^+ + 1) \\ &= 3(|S^+| - |R^+|)(y^+ + 1) \end{aligned}$$

(use Lemma 5 for step 2), this follows by induction. □

We can now deduce from the existence of a small set R with a small neighborhood S the existence of a superset of R with small neighborhood. We just calculate the maximal y with $y + 1 \leq |P|$ and $|Q|\, y < 3(|S| - |R|)$. It follows by Theorem 6, that there must be a subset R' of A with $|R| + y + 1$ vertices and a subset S' of B with size $|S| + y + 2$ that contains all neighbors of R. This step can be repeated until the claimed subset of A or the claimed neigborhood of the claimed subset of B has size $|A|/2$ depending if we consider supersets from a subset of A or supersets from a subset of B.

4 Results

Figure 1 shows the best possible expansion for 3-regular bipartite graphs with $2n$ vertices graphically. Interesting is to observe the pattern depending on the divisibility by 6 and the seemingly convergence to 4/3. Monien and Preis [MP01] have shown that for large n a cubic bipartite expander can have an expansion of at most $4/3 + \epsilon$.

Table 1 shows the best possible expansion. It also shows the runtime our algorithm needed on an Athlon-CPU with 3GHz and the number of tested subgraphs.

Table 2 shows the number of 3-regular bipartite graphs with $2n$ vertices, which have at least for one bipartion set the best possible expansion (last column). Column $A \vee B$ contains

This paper appeared in: Discrete Applied Mathematics 155 (2007), 2539-2545

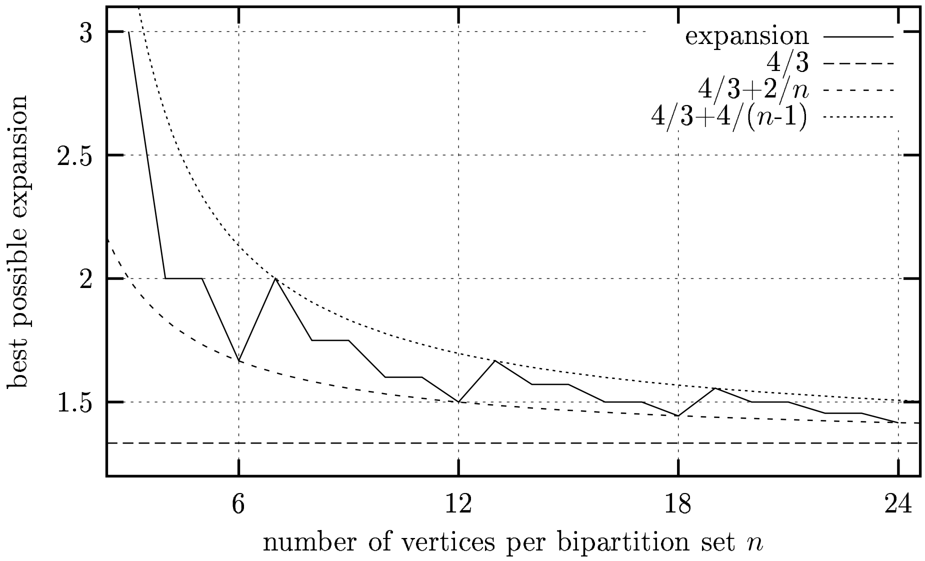

Figure 4.1: best possible expansion of 3-regular bipartite graphs

n	subgraphs	time	expansion
3	2	0 sec	3/1
4	4	0 sec	4/2
5	10	0 sec	4/2
6	18	0 sec	5/3
7	19	0 sec	6/3
8	42	0 sec	7/4
9	51	0 sec	7/4
10	480	0 sec	8/5
11	1467	0 sec	8/5
12	272	0 sec	9/6
13	552	0 sec	10/6
14	677	0 sec	11/7
15	5727	0.1 sec	11/7
16	32935	1.3 sec	12/8
17	6171	0.3 sec	12/8
18	61456	5.2 sec	13/9
19	96423	7.1 sec	14/9
20	8151	0.5 sec	15/10
21	1879924	1 min 11 sec	15/10
22	21052389	1 h 29 min	16/11
23	31588879	2 h 43 min	16/11
24	81502779	4 h 30 min	17/12

Table 1: best possible expansion of 3-regular bipartite graphs

n	subgraphs	time	$A \vee B$	not sym.	sym.	$A = B$	sum
3	3	0 sec	0	0	1	1	1
4	7	0 sec	0	0	1	1	1
5	20	0 sec	0	0	1	1	1
6	50	0 sec	0	0	3	3	3
7	68	0 sec	0	0	1	1	1
8	113	0 sec	0	0	1	1	1
9	480	0 sec	0	0	7	7	7
10	1885	0 sec	11	4	25	29	40
11	15166	0.2 sec	19	244	202	446	465
12	148230	3.6 sec	2292	2931	752	3683	5975
13	3748	0 sec	0	0	1	1	1
14	15831	0.2 sec	0	0	4	4	4
15	1601611	43.3 sec	2945	793	423	1216	4161
16	36034623	34 min 44 sec	177411	25924	2979	28903	206314
17	1318065325	68 h 46 min	15354167	8708524	75431	8783955	24138122
18							
19	216907	16.3 sec	1	0	0	0	1
20	917948423	17 h 52 min	1	0	4	4	5

Table 2: number of best expanders

the graphs with best possible expansion for exactly one bipartition set. Graphs with best possible expansion for both bipartition sets are divided by those that are symmetrical and those that are not symmetrical, i.e there is an/no automorphism that maps vertices from B to A and vice versa. Column $A = B$ shows the sum of these columns.

References

[BKV+81] Blum, Karp, Vornberger, Papadimitriou, Yannakakis, The complexity of testing whether a graph is a superconcentrator, Inform. Process. Letters 13, 164-167, 1981.

[GG81] O. Gabber, Z. Galil, Explicit constructions of linear-sized superconcentrators. J. Comput.System Sci., 22(3), 407-420, 1981.

[GIL+90] O. Goldreich, R. Impagliazzo, L. Levin, R. Venkatesan, and D. Zuckermann. Security preserving amplification of hardness. FOCS 1990, 318-326.

[IW97] R. Impagliazzo and A. Wigderson. $P = BPP$ if E has sub-exponential circuits: Derandomizing the XOR lemma. STOC 1997, 220-229.

[KO02] I. Köthnig. Expansion 3-regulärer bipartiter Graphen. Studienarbeit. Lehrstuhl für Algorithmen und Komplexität, Institut für Informatik, Humboldt-Universität zu Berlin, 2002.

[LPS88] A. Lubotzky, R. Phillips, P. Sarnak, Ramanujan graphs, Combinat. 8 (1988), 261–277.

This paper appeared in: Discrete Applied Mathematics 155 (2007), 2539-2545

[Mar73] G.A.Margulis, Explicit constructions of expanders, Problemy Peredaci Informacii, 9(4), 71-80, 1973.

[Mer96] M. Meringer, Erzeugung regulärer Graphen, Diplomarbeit,Bayreuth, 1996.

[MP01] B. Monien, R. Preis, Upper Bounds on the Bisection Width of 3- and 4-regular Graphs, In: 26th International Symposium on Mathematical Foundations of Computer Science (MFCS) 2001, LNCS 2136, Springer 2001, 524–536.

[Pap94] C. H. Papadimitriou, Computational Complexity, Addison Wesley, 1994

[Pin73] M. Pinsker, On the complexity of a concentrator. In 7th Annual Teletraffic Conference 318/1-318/4, Stockholm, 1973.

[Pip87] N. Pippenger. Sorting and selecting in rounds. SIAM Journal on Computing, 16(6):1032-1038, December 1987.

[PY82] N. Pippenger and A.C. Yao. Rearrangeable networks with limited depth. SIAM J. Algebraic and Discrete Methods, 3:411-417, 1982.

[RVW01] Reingold, Vadhan, Wigderson, Entropy Waves, the Zig-Zag Graph Product and New Constant-Degree Expanders and Extractors, ECCC Report No. 18, 2001.

[Sip88] Michael Sipser. Expanders, randomness, or time versus space. J. Comput. System Sci., 36(3):379-383, 1988.

[SS96] M. Sipser and D. A. Spielman. Expander Codes. IEEE Trans. Inform. Theory, 42(6, part1):1710-1722, 1996. Codes and complexity.

[Val77] L. G. Valiant. Graph-theoretic arguments in low-level complexity. LNCS 53, 1977, 162-176