アルゴリズム 39－9
（1994. 5. 13）

限定次数最大独立集合問題の近似について

マグナス ハルダースソン
北陸先端科学技術大学院大学 (情)

ジャイクマル ラダクリシュナン
タタ インスチチュト

最大独立集合問題は、組合せ最適化の基本的な問題である。この問題の最も重要な場合は、最大頂点次数が定数によって限定された場合であるが、その場合でさえ、この問題は NP 完全であることが知られている。一方、この問題に対する自然な近似アルゴリズムは Greedy と呼ばれ、線形時間で実行できる。本論ではまず、このアルゴリズムの近似率を解析し、$(\Delta+2)/3$ 近似率を示す。次に、この解析に基づいて、$O(\log^* n)$ 時間の並列および分散アルゴリズムを設計する。最後に, 独立集合アルゴリズムの近似率を改良するための一般的手法を紹介し、その手法を用いて $O(\Delta/\log\log\Delta)$ 近似率を示す。これは始めての $o(\Delta)$ 近似率である。

グラフ理論、近似アルゴリズム、限定次数最大独立集合問題

Recent Results on Approximating Independent Sets in Sparse and Bounded-degree Graphs

Magnús M. Halldórsson
School of Information Science
Japan Adv. Institute of Science and Tech.
Ishikawa 923-12, JAPAN
magnus@jaist.ac.jp

Jaikumar Radhakrishnan
Theoretical Computer Science Group
Tata Institute of Fundamental Research
Bombay, India
jaikumar@tifrvax.bitnet

The *maximum independent set problem* is one of the fundamental problems of combinatorial optimization. The most important case for which the problem remains NP-complete is when the maximum vertex degree is bounded by a constant. We analyze the simple and natural *minimum-degree Greedy* algorithm and obtain several improved bounds: $(\Delta+2)/3$ performance ratio on graphs with maximum degree Δ, $(2\bar{d}+3)/5$ ratio on graphs with average degree $\bar{d}$, and a $O(\log^* n)$ parallel algorithm attaining these bounds on constant degree graphs. Finally, we introduce a generally applicable technique for improving the approximation ratios of independent set algorithms. We use this technique to obtain a $O(\Delta/\log\log\Delta)$ ratio for large values of Δ, for the first $o(\Delta)$ ratio.

graph theory, approximation algorithms, bounded degree independent set problem

1 Introduction

An *independent set* in a graph is a collection of vertices that are mutually non-adjacent. The problem of finding an independent set of maximum cardinality is one of the fundamental combinatorial problems. Unfortunately, it is known to be $\mathcal{NP}$-complete, even for bounded-degree graphs, and therefore no efficient algorithms are in sight.

Given the hardness of exact computation, we are interested in approximation algorithms for the independent set problem in bounded-degree graphs. In particular, we seek an algorithm with a good *performance ratio*, which is a bound on the maximum ratio between the optimal solution size (i.e. the independence number) and the size of the solution found by the heuristic. The study of such algorithm has become increasingly more prevalent.

One of the most ubiquitous heuristic methods for this problem is the *greedy algorithm*. It iteratively selects a vertex of minimum degree and deletes that vertex and all of its neighbors from the graph, until the graph becomes empty. As a delightfully simple and efficient algorithm, the Greedy method deserves a particularly detailed analysis. It is already known to possess several important properties: attaining the Turán bound, and its generalization in terms of degree sequences [8]; almost always obtaining a solution at least half the size of an optimal solution in a general random graph [3]; yielding a non-trivial graph coloring approximation [14] when applied iteratively; and finding optimal independent sets in forests, series-parallel graphs, and cographs. While Greedy has been frequently studied before, the true extent of its performance ratio has apparently not been determined previously. The best ratio previously claimed was $\Delta - 1$ on graphs with maximum degree Δ [22] and $\overline{d} + 1$ on graphs of average degree $\overline{d}$.

Our main result is that Greedy is surprisingly much better than previously expected. We obtain a tight performance ratio of $(\Delta + 2)/3$ in terms of maximum degree, and an asymptotically optimal bound of $(\overline{d}+2)/2$ in terms of average degree. It comes as a considerable surprise that this simple, linear time method performs as well as it does. In the process, we give a natural extension of Turán's bound that incorporates the actual independence number of the graph, and give a general, tight expression of the size of the solution found as a function of the independence number and the number of vertices. Section 2 contains the detailed analysis of Greedy.

We further analyze Greedy in combination with a fractional relaxation technique of Nemhauser and Trotter [17, 13], in subsections 2.5 and 2.6. We use it to improve the best performance ratio known in terms of average degree to $(2\overline{d}+3)/5$, but show it to be of limited use in terms of maximum degree.

We introduce in section 3 a new technique that involves removing small dense subgraphs. This technique is a general schema that can potentially be applied to any approximation algorithm for this and related problems. By removing all cliques of fixed size from the graph, we can either find a larger solution or obtain a better upper bound on the size of the optimal solution. By implementing a graph theorem of Ajtai et al [1], we can obtain by this scheme a performance ratio of $O(\Delta/\log\log\Delta)$, which is the first $o(\Delta)$ ratio. Using different component algorithms, we can also obtain a practical method that improves the best performance ratio for graphs of intermediate maximum degree.

Finally, in section 4 we show that the performance ratios proved for Greedy in bounded-degree graphs can also be obtained by a simple parallel algorithm. Our analysis of Greedy suggests that globally minimum degree is not required, in fact, any vertex satisfying the locally evaluated property *"degree of vertex is at most the average of its neighbors' degrees"* can be selected. This is a simple *local rule* that can be implemented efficiently both in parallel and distributed.

1.1 Related results

Until very recently, the best ratio claimed for any approximation algorithm for independent sets in bounded degree graphs was $\Delta/2$ [13]. Suddenly, several developments took place.

Berman and Fürer [4] obtained significantly improved performance ratios of $(\Delta+3)/5+\epsilon$ and $(\Delta+3.25)/5+\epsilon$ for even and odd degrees respectively. Their method is a type of a local improvement method. It's drawback is the astoundingly huge complexity of $n^{32\Delta^{4/\epsilon}}$, which for typical solution quality means on the order of $n^{2^{100}}$. While some improvements in the complexity are possible [12], even a n^{50} complexity appears out of reach. Khanna et al. [15] considered a very simple (and fast) local improvement method. Using their analysis, it is easy to show that the performance of that algorithm when complemented with Nemhauser-Trotter is $(\Delta + 2)/3$, same as Greedy.

The results covered here are combined from [11] and [12]. Also reported in [12] is further analysis of local improvement methods. We can obtain a $(\Delta+3)/4$ using time that is linear in n if Δ is a constant. It can be improved to $(\Delta + 2)/4 + \epsilon$ using $O(\Delta^{\Delta/\epsilon}n)$ time.

In spite of these related results, we believe the results on Greedy reported here hold their own, given the algorithm's simplicity, superior complexity, and general applicability.

In sparse graphs, the only result we are aware of is a $(\overline{d}+1)/2$ performance ratio of Greedy with Nemhauser and Trotter's method [13]. No previous parallel approximation algorithms were known to us, except the Δ and $\delta + 1$ ratios that can be obtained from parallel implementations of Brooks' (see [18] for references) and Turán's theorems [10], respectively.

1.2 Notation

We use fairly standard graph notation and terminology. For the graph in question, usually denoted by G, n denotes the number of vertices, Δ the maximum degree, $\overline{d}$ the average degree, α the independence number (the size of the largest independent set), and τ the independence fraction (i.e. α/n). For a vertex v, $d(v)$ denotes the degree of v, and $N(v)$ the set of neighbors of v.

For an independent set algorithm A, the performance ratio of A is defined by

$$\rho_A = \max_G \frac{\alpha(G)}{A(G)}$$

where $A(G)$ is the size of the solution obtained by algorithm A on graph G. We particularly consider two algorithms: Greedy, for short Gr, and the combination of Greedy and Nemhauser-Trotter, denoted by $Gr + NT$. Gr is also a shorthand for $Greedy(G)$.

2 Analysis of Greedy

The Minimum-Degree Greedy algorithm, or Greedy for short, operates as follows. It executes a sequence of *reductions*, each of which corresponds to an iteration, where a vertex is selected and added to the solution, and then it and its neighborhood are removed from the graph. It stops when the graph has been exhausted and outputs the set of selected vertices.

```
Greedy(G)
  I ← ∅
  while G ≠ ∅ do
     Choose v such that d(v) = min_{w∈V(G)} d(w)
     I ← I ∪ {v}
     G ← G − {v} ∪ N(v)
  od
  Output I
end
```

The algorithm can be implemented in time linear in the number of edges and vertices, independent of degree. This involves maintaining a multiset of small nonnegative integers (the degrees of the vertices), along with the associated vertex number, under the operations of unit decrease, deletion, and finding the minimum. This could be implemented by an array of linked lists, one for each degree value, along with an appropriate array structure to provide for direct referencing.

We use the following notation for the operation of Greedy. Let t be the number of reductions performed by Greedy and let $d_1, d_2, \ldots, d_t$ be the degrees of the vertices selected. The number of vertices removed in the i-th reduction is thus $d_i + 1$. The main property of the algorithm that we use in our analysis is that the sum of the degrees of the vertices removed in the i-th reduction must be at least $d_i(d_i + 1)$. This allows us to lower bound the number of edges removed in each step.

2.1 Relative size of Greedy solutions

We start by generalizing the constructive version of Turán's theorem by pushing the independence ratio into the expression.

Theorem 1 $Gr \geq \dfrac{1+\tau^2}{\overline{d}+1+\tau}n.$

Proof. Counting the number of vertices and edges deleted in each reduction gives us the necessary inequalities to prove the claim.. The removal of vertices in each reduction partitions the vertex set, yielding

$$\sum_{i=1}^{t}(d_i+1) \;=\; n. \tag{1}$$

Fix a maximum independent set and let k_i be the number of vertices among the $d_i + 1$ vertices deleted in reduction i that are also contained in that independent set. Then,

$$\sum_{i=1}^{t} k_i \;=\; \alpha. \tag{2}$$

Since Greedy always picks a vertex of minimum degree, the sum of the degrees of the vertices deleted in step i is at least $d_i(d_i + 1)$. Note that no edge can have both its end points in the maximum independent set. Then, it can be shown that the number of edges deleted in step i is at least $\binom{d_i+1}{2} + \binom{k_i}{2}$. Hence,

$$\frac{\overline{d}}{2}n \;=\; |E| \;\geq\; \sum_{i=1}^{t}\binom{d_i+1}{2} + \binom{k_i}{2}. \tag{3}$$

We now add (1), (2) and twice (3), and apply the Cauchy-Schwarz inequality to obtain

$$(\overline{d}+1+\tau)n \;\geq\; \sum_{i=1}^{t}(d_i+1)^2 + k_i^2 \;\geq\; (1+\tau^2)n^2/t.$$

Rearranging the inequality, we obtain the desired bound on t. ∎

We now turn our attention to bounded-degree graphs, using techniques similar to the preceding proof to obtain bounds parametrized by the maximum degree Δ.

Theorem 2 $Gr \;\geq\; \dfrac{1-\tau(1-\tau)}{(1-\tau)\Delta+1}n.$

Proof. The proof follows the proof of the preceding theorem with some extensions. In the i-th step $d_i + 1$ vertices and all edges incident on them are deleted. Of these edges, let x_i have only one end in these $d_i + 1$ vertices; the remaining edges have both ends among the $d_i + 1$ vertices: of these, let y_i of these have one end in the independent set and one outside, and z_i have both ends outside. Then we have

$$x_i + 2(y_i + z_i) \;\geq\; d_i(d_i+1), \tag{4}$$

$$y_i \;\leq\; k_i(d_i+1-k_i), \tag{5}$$

Multiply (5) by -1 (reversing the inequality) and add it to (4) to obtain

$$\begin{aligned} x_i + y_i + 2z_i \;&\geq\; d_i(d_i+1) - k_i(d_i+1-k_i) \\ &\geq\; \binom{d_i+1}{2} + \binom{d_i+1}{2} - k_i(d_i+1-k_i) \\ &=\; \binom{d_i+1}{2} + \binom{k_i}{2} + \binom{d_i+1-k_i}{2}. \end{aligned}$$

Since the number of edges deleted in the i-th step is precisely $x_i + y_i + z_i$, we have the following extension of (3):

$$|E| \;\geq\; \sum_{i=1}^{t}\binom{d_i+1}{2} + \binom{k_i}{2} + \binom{d_i+1-k_i}{2} - z_i. \tag{6}$$

We also count the total degree of vertices outside the maximum independent set, which entails counting edges incident on the independent set vertices once but those fully outside the independent set twice.

$$(n-\alpha)\Delta \;\geq\; \sum_{i=1}^{t} z_i + |E|. \qquad (7)$$

Now add twice (6) and twice (7) to obtain

$$2(n-\alpha)\Delta \;\geq\; \sum_{i=1}^{t} d_i(d_i+1)+k_i(k_i-1)+(d_i+1-k_i)(d_i-k_i).$$

To simplify the right hand side, we add $\sum_{i=1}^{t} d_i + k_i + (d_i - k_i)$ and compensate for this by adding $2n$ to the left hand side (invoking (1)). Then,

$$2\Delta(n-\alpha)+2n \;\geq\; \sum_{i=1}^{t}(d_i+1)^2+k_i^2+(d_i+1-k_i)^2. \quad (8)$$

Using (1), (2) and the Cauchy-Schwarz inequality, we obtain

$$(2\Delta(1-\tau)+2\tau)n \;\geq\; [1+\tau^2+(1-\tau)^2]n^2/t.$$

The claim follows from this. ∎

2.2 Performance guarantee

The following bound on the performance of Greedy on sparse graphs follows immediately since the ratio of the independence number τn to the bound in thm. 1 is maximized when $\tau = 1$.

Corollary 3 $\rho_{G_r} \;\leq\; \dfrac{\overline{d}+2}{2}$.

For bounded-degree graphs, the general expression obtained in thm. 2 almost – but not quite – yields our main claim about the performance ratio of Greedy. We omit the proof.

Theorem 4 $\rho_{G_r} \;\leq\; (\Delta+2)/3$.

2.3 Limitations

The performance ratios proved above cannot be improved.

Theorem 5 $\rho_{G_r} \;\geq\; \dfrac{\Delta+2}{3} - O(\Delta^2/n)$, *for* $\Delta \geq 3$.

Proof. We give a detailed construction for $\Delta \equiv 1 \pmod 3$. We construct a graph, parametrized by integer $l \geq 2$, that consists of a chain of repetitions of a pair of subgraphs: a clique on l vertices followed by an independent set on l vertices. The two subgraphs are completely connected, while the connections between the independent set and the clique of the following pair miss only a single matching (i.e. each vertex is of degree $l-1$). The chain ends with one additional clique.

The essential property of the graph is that the degree of the independent set vertices equals the degree of the vertices of the first clique of the chain. We can therefore assume that Greedy will pick one of the vertices from the first clique and remove the remaining vertices from the pair, reducing the graph to an identical chain with one fewer pairs. Thus, Greedy selects one vertex from each pair, plus one from the final clique, for a total of $(n-l)/2l+1$. The optimal solution contains all the independent set vertices for a total of $(n-l)/2$. This yields a ratio of

$$\rho_{G_r} \;\geq\; l - 2l^2/n.$$

To relate that to the degree measures, we have that $\Delta = 3l-2$, and

$$\overline{d} \;\leq\; 2\left(\binom{l}{2} + l^2 + l(l-1)\right)/2l$$
$$= \frac{5l-3}{2}.$$

Thus,

$$\rho_{G_r} \;\geq\; \frac{\Delta+2}{3} - O(\Delta^2/n), \qquad (9)$$

and

$$\rho_{G_r} \;\geq\; \frac{2\overline{d}+3}{5} - O(\overline{d}^2/n), \qquad (10)$$

even when $\tau \leq 1/2$.

The construction for $\Delta \equiv 0, 2 \pmod 3$ is similar, but less regular, and is omitted. ∎

We can also show that the bound on the performance ratio in terms of average degree (Corollary 3) approaches optimality as $\overline{d}$ gets larger. The proof is omitted.

Theorem 6 $\rho_{G_r} \;\geq\; \dfrac{\overline{d}+2}{2} - O(1/\overline{d})$.

The graph for which this ratio is attained consists of a chain of pairs of subgraphs as in the previous example, with each clique reduced to a single vertex adjacent to several of the following single-vertex cliques. We omit the detailed proof.

Variations of the above constructions show that the bounds of thms. 1 and 2 are tight for a range of values of τ. This involves varying the size of the cliques relative to the size of the independent sets, possibly by connecting each independent set to several subsequent cliques. We omit the details.

2.4 Nemhauser-Trotter

A method of Nemhauser and Trotter [17] for the fractional vertex cover problem has been used successfully to obtain better approximations for both the vertex cover and the independent set problems on bounded-degree graphs. The time complexity of this method is $O(n^{3/2} + |E|)$.

Their method yields an optimal solution to the linear relaxation of the integer problem where each variable, representing whether a node is in the independent set or the vertex cover, has value from $\{0, 1, 1/2\}$. This corresponds to partitioning the vertices into three sets R, P, and Q with the following important properties:

1. Some maximum independent set of G contains all the vertices of R and none of the vertices in P.

2. The independence fraction of the subgraph H induced by Q is at most $1/2$.

This in effect means that the method can be used as a preprocessor for approximation algorithms; since the portions R and P are solved optimally, it suffices to focus our attention on the graph H where the independence number is guaranteed to be small.

This immediately implies, for instance, that the $n/(\Delta+1)$ lower bound on the size of the Greedy solution along with the above upper bound on the independence number yields a $(\Delta+1)/2$ performance ratio for the combination of Greedy and Nemhauser-Trotter methods. A stronger bound follows immediately from thm. 2.

Corollary 7 $\rho_{Gr+NT} \leq \dfrac{\Delta+2}{3}$.

It follows from (9) that this ratio is essentially tight for $\Delta \equiv 1 \pmod 3$. That is, if the only property used about the Nemhauser-Trotter method is that it allows us to assume that the independence fraction is at most half, then we can do no better than we have shown. As we have not analyzed the method in detail, it is possible however that it will have other properties that inhibit examples where Greedy has poor performance.

We can improve on this bound slightly to $(\Delta+2)/3 - 1/(3\Delta+2)$ when $\Delta \equiv 0, 2 \pmod 3$. For instance, we have a tight ratio of $3/2$ when $\Delta = 3$. Also, for $\Delta = 5$ we get a ratio of $16/7 \approx 2.286$, down from $2.3\overline{3}$, while there is a graph that forces a ratio of 2.27. We omit the details.

2.5 Average degree

Once Nemhauser-Trotter has been applied, the average degree may have changed for the worse and we cannot immediately apply the bounds proved on Greedy. Nevertheless, a closer look shows that the bounds will complement each other as hoped for. Hochbaum [13], showed that the Turán bound on Greedy can be complemented with the $\tau \leq 1/2$ promise of Nemhauser-Trotter to yield a performance ratio of $(\overline{d}+1)/2$.

Our bound on Greedy, when complemented by Nemhauser-Trotter, yields a considerably performance ratio. The proof is similar to that of Hochbaum and is omitted.

Theorem 8 $\rho_{Gr+NT} \leq \dfrac{2\overline{d}+3}{5}$.

This again is tight for $\Delta \equiv 1 \pmod 3$, from (10).

3 Subgraph Removal

We present a generally applicable method for improving the performance ratio of independent set approximation algorithms. Any algorithm whose performance ratio decreases as the independence fraction of the graph decreases, can be enhanced using this approach, with greater improvements as the maximum degree gets larger.

The idea comes from the observation that graphs without dense subgraphs, particularly cliques, contain provably larger independent sets than graphs do in general, and moreover these larger solutions can be found effectively. We remove all cliques of certain size from the input graph and apply the improved algorithms on the remaining graph. This will be advantageous as long as the input graph contains few disjoint cliques; if it contains many disjoint cliques, the independence number must be low and our performance ratio will improve in either case. This idea was previously used to approximate the Independent Set problem in general graphs [5].

This schema uses as subroutine two types of algorithms: an approximation algorithm for general graphs, and algorithms that find large independent sets in k-clique free graphs with possibly different algorithms for different values of k. Except for the case of triangle-free graphs, we consider here only the case of Greedy. Better algorithms of either type translate immediately to better performance ratios. In fact, we have recently improved the best performance ratios known for moderate to large values of Δ [12] using better algorithms for clique-free graphs.

We first illustrate this technique in its simplest form, when removing disjoint 2-cliques (i.e. a matching). The next step is to consider removing triangles, where we can use an effective algorithm of Shearer. We then present the general schema, and prove improved ratios for Greedy in the context of this scheme.

Removing edges The following simple idea has also been considered in [19] and [4]. We use it to get a linear time algorithm with a good performance ratio in terms of average degree.

Find an independent in two ways, and retain the larger one. One way is to find a maximal matching with m edges, and use the complement – the vertices not appearing in the matching – as an independent set approximation. The size of this set is $n-2m$ which is at least $2\alpha - n$, since the independence number α is at most $n-m$. Hence, the performance ratio is at most $\frac{\tau}{2\tau-1}$. The other way is to use the Greedy algorithm, with a performance guarantee of $\frac{(\overline{d}+1+\tau)\tau}{1+\tau^2}$ (see thm. 3).

Observe that the former ratio is monotone decreasing with τ ($\tau > 1/2$), while the latter one is monotone increasing. A close study shows that the value of τ for which the ratios agree is at most $\frac{1}{2}+\frac{5/4}{2\overline{d}+2}$. If we plug that into the higher ratio, that of the maximal matching, this yields a performance ratio of $\frac{2\overline{d}+4.5}{5}$. Hence, in fully linear time we come within an additive 0.3 of the $(2\overline{d}+3)/5$ bound in sec. 2.5 that requires $\Omega(n^{3/2})$ time.

3.1 AEKS

Ajtai, Erdős, Komlós and Szemerédi [2] proved the following result about K_ℓ free graphs.

Theorem 9 (AEKS) *There exists an absolute constant* c_1 *such that for any* K_ℓ*-free graph* G,

$$\alpha(G) \geq c_1 \frac{\log((\log \overline{d})/\ell)}{\overline{d}} n.$$

We have obtained an algorithm AEKS that constructs such an independent set in polynomial time, by derandomizing the parts of the proof where probabilistic existence arguments are used. For lack of space, we omit its description.

We then use the following simple algorithm to approximate independent sets. An independent set is *maximal* (MIS) if adding any further vertices to the set violates independence. An MIS is easy to find and provides a sufficient general upper bound of $n/(\Delta+1)$.

AEKS-SR(G)
  $G' \leftarrow G -$ CliqueCollection$(G,\ c_1 \log\log\Delta)$
  return $\max$(AEKS(G'), MIS(G))
end

Theorem 10 *The performance ratio of* AEKS-SR *is* $O(\Delta/\log\log\Delta)$.

Proof. Let k denote $c_1 \log\log\Delta$, and let n' denote the size of $V(G')$. The independence number collects at most one from each k-clique, for at most

$$\alpha \ \leq\ n/k + n' \ \leq\ 2\max(n/k, n')$$

while the size of the solution found by AEKS-SR is at least

$$\text{AEKS-SR}(G) \ \geq\ \max(\frac{1}{\Delta+1}n, \frac{k}{\Delta}n') \ \geq\ \frac{k}{\Delta+1}\max(n/k, n').$$

The ratio between the two clearly satisfies the claim. ∎

Observe that the combined method runs in polynomial time for Δ as large as $n^{1/\log\log n}$.

3.2 Generic Clique Removal Schema

We now give the general algorithm, indexed by k. The algorithm is really a schema that depends on what subordinate methods are being used. One is algorithm General-BDIS-Algorithm used to find independent sets in general (bounded-degree) graphs. Also we need methods to find independent sets in ℓ-clique free graphs, possibly one for each value of $\ell \geq 3$. Finally, observe that S defined below consists of a collection of vertex subsets that are mutually disjoint and each induce a clique of size ℓ.

CliqueRemoval$_k(G)$
  $A_0 \leftarrow$ General-BDIS-Algorithm(G)
  for $\ell = k$ downto 2 do
    $S \leftarrow$ CliqueCollection(G,ℓ)
    $G \leftarrow G - S$
    $A_\ell \leftarrow K_\ell$-free-BDIS-Algorithm(G)
  od
  Output A_i of maximum cardinality
end

3.3 Effective method for moderately large maximum degree

While the clique removal method in combination with AEKS yields good asymptotic performance ratios, Δ must be quite high for the gained $\log\log\Delta$ factor to overcome the large constants involved.

We now turn our attention to practical methods that can benefit from the clique removal schema. This involves an algorithm of Shearer [20] for 3-clique-free graphs, and a simple local search algorithm for other k-clique-free graphs as well as for use as the general BDIS algorithm.

2-opt Khanna et al. [15] studied a simple local search algorithm that we have named 2-opt. Starting with an initial maximal independent sets, it tries all possible ways of adding two vertices and removing only one while retaining the independence property. We say that a triple $\langle v_1, v_2, u\rangle$ is a *2-improvement* of an independent set I iff vertices v_1, v_2 are outside of I, u is in I, and adding the former two to I while removing the latter retains the independence property. Since I can be assumed to be a maximal independent set, it suffices to look at pairs adjacent to a common vertex in I.

2-opt(G)
  $I \leftarrow$ MIS(G)
  while ($\exists$ 2-improvement $\langle v_1,\ v_2,\ u\rangle$ of I)
    $I \leftarrow I \cup \{v_1, v_2\} - \{u\}$
  return I
end

Using proper data structures, the algorithm can be implemented in time $O(poly(\Delta)n)$ time.

The following was shown by Khanna et al [15].

Lemma 11 2-opt $\geq \dfrac{1+\tau}{\Delta+2}n$

We can get improved bounds for k-clique free graphs.

Lemma 12 *On a k-clique free graph G,*

$$\text{2-opt}(G) \ \geq\ \frac{2}{\Delta+k}n.$$

Shearer Shearer [20] proved the following theorem, improving a previous result of Ajtai, Komlós and Szemerédi [2].

Theorem 13 (Shearer [20]) *Let* $f_s(d) = (d\log_e d - d + 1)/(d-1)^2$, $f_s(0) = 1$, $f_s(1) = \frac{1}{2}$. *For a triangle-free graph* G, $\alpha(G) \geq f_s(\overline{d})n$.

Moreover, he gave the following algorithm that attains the claimed bound.

Shearer(G)
  $A \leftarrow \emptyset$.
  while $G \neq \emptyset$ do
    Pick a vertex v of degree d_v such that
      $(d_v+1)f_s(\overline{d}) \ \leq\ 1+(\overline{d}d_v+\overline{d}-2\sum_{w\in N(v)} d(w))f_s'(\overline{d})$
    $A \leftarrow A \cup \{v\}$

```
    G ← G − (N(v) ∪ {v}).
  od
  return A.
end
```

We shall only need the obvious corollary that Shearer$(G) \geq f_s(\Delta)n \approx n(\log\Delta)/\Delta$.

Using an appropriate data structure to maintain the f-values of the vertices, the algorithm can be implemented in time $O(poly(\Delta)n)$. In fact, the claim is also obtained in fully linear time by a simple randomized greedy algorithm, that chooses a random non-adjacent vertex in each step.

To improve the approximation further, we apply the method of Nemhauser and Trotter on each incarnation of G. That will allow us to assume that nothing will be left after the edges (2-clique) are removed.

We obtain the following explicit, if less than compact, bound on the performance ratio. H_k is the k-th Harmonic number.

Theorem 14 CliqueRemoval$_k$, *using* 2-opt *and* Shearer *attains a performance ratio of at most*

$$\left[\frac{\Delta}{2}+2+\frac{k}{2}(H_{k-1}+\frac{1}{3f_s(\Delta)}-\frac{3}{2}+\frac{\Delta}{3})\right]/(k+1)$$

for graphs of maximum degree $\Delta \geq 5$.

4 Parallel and Distributed Algorithm

The Greedy algorithm stipulates that in each step a vertex of globally minimum degree be selected, added to the solution, and removed from the graph along with its neighbors. As such, it looks impossible to parallelize, as well as offering little freedom for heuristic improvements. Fortunately, this is one of the delightful cases when the analysis guides us towards the design of better and/or more general algorithms.

We observe that the selection of a vertex is prescribed by a simple local rule and that a constant fraction of the vertices in a bounded-degree graph satisfies this rule. This has some interesting implications. For one, it opens up the possibility of the design of heuristics using secondary selection rules that retain the performance ratios of thms. 1 and 2. Another is a straightforward derivation of a parallel as well as a distributed approximation algorithm attaining these performance ratios.

From the proofs of thms. 1 and 2, we find that a sufficient criteria for the selected vertex v is that its degree be less than the average of its neigbors' degrees. That is,

$$d(v) \leq \frac{\sum_{w\in N(v)} d(w)}{d(v)}. \tag{11}$$

A heuristic may choose any ordering that obeys the above property. Vertices can be selected in parallel as long as the selection of one doesn't affect the above criteria for the other. In particular, vertices with disjoint and non-adjacent neighborhoods (i.e. of distance three or greater) can be selected and processed concurrently.

This suggests a natural approach to a parallel algorithm:

1. Find a set W of vertices satisfying (11).
2. Form a graph H on vertex set W, with edge between vertices that may interact with respect to the Greedy reduction (i.e. of distance three or less).
3. Find a maximal independent set MIS in H.
4. Perform the Greedy reduction on these vertices in parallel. Namely, add the vertices MIS to the solution, and delete them and their neighbors from the graph.
5. Repeat from step 1 until the graph is empty.

The following lemma due to Alon and Szegedy (private communication) shows that a significant fraction of the vertices must have the above property simultaneously.

Lemma 15 *In a graph on n vertices with maximum degree Δ, at least $\frac{4}{\Delta^2+4}n$ vertices satisfy property (11).*

Proof. Let D_v denote $d(v)^2 - \sum_{w\in N(v)} d(w)$. We shall show that

$$\Pr_v[D_v \geq 0] \;\geq\; \frac{4}{\Delta^2+4},$$

which implies the lemma. As observed by Shearer [20], $E_v[D_v] = 0$. The value of D_v is bounded above by $d(v)(\Delta - d(v)) \leq \Delta^2/4$, and since it is integral, it must differ from zero by at least one when negative. Thus,

$$-1(1-\Pr[D_v \geq 0]) + \Delta^2/4 \cdot \Pr[D_v \geq 0] \geq 0.$$

The claim now follows. ∎

Theorem 16 *There is an EREW parallel algorithm that finds an independent set of size and performance satisfying theorems 2 and 4 in time* $\log^* n \min(poly(\Delta)\log n, \Delta^\Delta)$ *using n processors.*

Proof. Each vertex added to the solution will satisfy property (11) regardless of the order of removal of the simultaneously chosen vertices. Hence, the results of the theorems apply to this algorithm.

Let us now estimate the time complexity, starting with the number of iterations. The first step reduces the number of vertices by a factor of at most $O(\Delta^2)$, as per the lemma above. The number of vertices deleted in the fourth step, which are the selected vertices and their neighbors, is at most another Δ^2 factor smaller. Thus at least n/Δ^4 vertices are removed from the graph in each step, so the number of rounds is bounded by $\Delta^4\log n$. Also notice that for any vertex in the graph, some vertex of distance at most Δ gets eliminated in each round. Hence the number of rounds is also bounded by Δ!

The only non-trivial step in each round is the computation of a maximal independent set (MIS) of the graph H. An algorithm of Goldberg et al. [9] finds an MIS in time $O(\Delta(H)\log\Delta(H)(\Delta(H)+\log^* n))$ using linear number of processors. The combined time complexity is therefore bounded by $O(\Delta^7\log\Delta(\Delta^3+\log^* n)\log n)$. The processor count is linear in n, and considering the total amount of work can probably be made some polynomial of Δ smaller. ∎

The algorithm given above also satisfies the criteria of a distributed algorithm.

5 Comparison

We compare the ratios obtained by the practical version of the clique removal method of sec. 3, the expensive local improvement method of Berman and Furer [4], the more efficient version of [12], and the Greedy algorithm of sec. 2. The methods offer a wide range of performance/complexity tradeoffs, with strengths in different regions

Δ	Clique rem.	[4]	[12]	Greedy
10	3.54	2.60	2.75	4.00
33	8.92	7.25	8.50	11.66
100	23.01	20.60	25.25	34.00
1024	201.57	205.40	256.25	342.00
8192	1535.20	1639.00	2048.25	2731.33

Acknowledgments We thank Noga Alon and Mario Szegedy for kindly proving lemma 15.

References

[1] M. Ajtai, P. Erdős, J. Komlós, and E. Szemerédi. On Turán's theorem for sparse graphs. *Combinatorica*, 1(4):313–317, 1981.

[2] M. Ajtai, J. Komlós, and E. Szemerédi. A note on Ramsey numbers. *J. Combin. Theory Ser. A*, 29:354–360, 1980.

[3] N. Alon and J. Spencer. *The Probabilistic Method.* Wiley, 1992.

[4] P. Berman and M. Fürer. Approximating maximum independent set in bounded degree graphs. In *Proc. Fifth ACM-SIAM Symp. on Discrete Algorithms*, Jan. 1994.

[5] R. B. Boppana and M. M. Halldórsson. Approximating maximum independent sets by excluding subgraphs. *BIT*, 32(2):180–196, June 1992.

[6] V. Chvátal. A greedy heuristic for the set covering problem. *Math. Oper. Res.*, 4:233–235, 1979.

[7] V. Chvátal. *Linear Programming.* Freeman, New York, 1983.

[8] P. Erdős. On the graph theorem of Turán (in Hungarian). *Mat. Lapok*, 21:249–251, 1970.

[9] A. V. Goldberg, S. A. Plotkin, and G. E. Shannon. Parallel symmetry-breaking in sparse graphs. In *Proc. 19th ACM Symp. on Theory of Computing*, pages 315–324, May 1987.

[10] M. Goldberg and T. Spencer. An efficient parallel algorithm that finds independent sets of guaranteed size. In *Proc. First ACM-SIAM Symp. on Discrete Algorithms*, pages 219–225, Jan. 1990.

[11] M. M. Halldórsson and J. Radhakrishnan. Greed is good: Approximating independent sets in sparse and bounded-degree graphs. To appear in *Proc. 26th ACM Symp. on Theory of Computing*, Montreal, May 1994.

[12] M. M. Halldórsson and J. Radhakrishnan. Improved approximations of independent sets in bounded-degree graphs. To appear in *4th Scandinavian Workshop on Algorithm Theory*, Denmark, July 1994.

[13] D. S. Hochbaum. Efficient bounds for the stable set, vertex cover, and set packing problems. *Disc. Applied Math.*, 6:243–254, 1983.

[14] D. S. Johnson. Worst case behavior of graph coloring algorithms. In *Proc. 5th Southeastern Conf. on Combinatorics, Graph Theory, and Computing. Congressus Numerantium X*, pages 513–527, 1974.

[15] S. Khanna, R. Motwani, M. Sudan, and U. Vazirani. On syntactic versus computation views of approximability. Manuscript, Dec. 1993.

[16] B. Monien and E. Speckenmeyer. Some further approximation algorithms for the vertex cover problem. In *Lecture Notes in Computer Science #159*, pages 341–349. Springer Verlag, 1983.

[17] G. L. Nemhauser and W. T. Trotter. Vertex packing: Structural properties and algorithms. *Math. Programming*, 8:232–248, 1975.

[18] A. Panconesi and A. Srinivasan. Improved distributed algorithms for coloring and network decomposition problems. In *Proc. 24nd Ann. ACM Symp. on Theory of Computing*, pages 581–592, May 1992.

[19] V. T. Paschos. A $(\delta/2)$-approximation algorithm for the maximum independent set problem. *Inf. Process. Lett.*, 44:11–13, Nov. 1992.

[20] J. B. Shearer. A note on the independence number of triangle-free graphs. *Discrete Math.*, 46:83–87, 1983.

[21] J. B. Shearer. A note on the independence number of triangle-free graphs, II. *J. Combin. Theory Ser. B*, 53:300–307, 1991.

[22] H. U. Simon. On approximate solutions for combinatorial optimization problems. *SIAM J. Disc. Math.*, 3(2):294–310, May 1990.

[23] P. Turán. On an extremal problem in graph theory (in Hungarian). *Mat. Fiz. Lapok*, 48:436–452, 1941.