

Greed is good*: Approximating independent sets in sparse and bounded-degree graphs †

Magnús M. Halldórsson‡

Jaikumar Radhakrishnan§

Abstract

The *minimum-degree greedy* algorithm, or **Greedy** for short, is a simple and well studied method for finding independent sets in graphs. We show that it achieves a performance ratio of $(\Delta + 2)/3$ for approximating independent sets in graphs with degree bounded by Δ . The analysis yields a precise characterization of the size of the independent sets found by the algorithm as a function of the independence number, as well as a generalization of Turán’s bound. We also analyze the algorithm when run in combination with a known preprocessing technique, and obtain an improved $(2\bar{d} + 3)/5$ performance ratio on graphs with average degree \bar{d} , improving on the previous best $(\bar{d} + 1)/2$ of Hochbaum. Finally, we present an efficient parallel and distributed algorithm attaining the performance guarantees of **Greedy**.

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. 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.

One of the most ubiquitous heuristic methods for this problem is the *greedy algorithm* which selects a vertex of minimum degree, deletes that vertex and all of its neighbors from the graph, and repeats this process 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 qualities: attaining the Turán bound, and its generalization in terms of degree sequences [30, 8]; almost always obtaining a solution at least half the size of an optimal solution in a random graph [23]; yielding a non-trivial graph coloring approximation [16], and a “light” coloring with a small chromatic sum [19], when applied iteratively as a coloring method; and finding optimal independent sets in trees, series-parallel, cographs, and graphs of degree at most 2.

While the performance ratio of **Greedy** has been analyzed before to some extent, the true extent of its performance has apparently not been determined before. The best ratio previously

*Gordon Gekko [29].

†A preliminary version of this paper appeared in the 26th ACM Symposium on Theory of Computing, 1994.

‡Science Institute, University of Iceland, IS-107 Reykjavik, Iceland. email: mmh@rhi.hi.is. Work done at Japan Advanced Institute of Science and Technology – Hokuriku.

§Theoretical Computer Science Group, Tata Institute of Fundamental Research, Bombay, India 400 005. email: jaikumar@tcs.tifr.res.in. Work done while visiting JAIST.

claimed for **Greedy** was $\Delta - 1$ on graphs with maximum degree Δ [27] and $\bar{d} + 1$ on graphs of average degree \bar{d} [14].

Our main result is that **Greedy** is much better than previously claimed. We obtain a tight performance ratio of $(\Delta + 2)/3$ in terms of maximum degree, and an asymptotically optimal bound of $(\bar{d} + 2)/2$ in terms of average degree. 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.

We further analyze **Greedy** extended with a preprocessing method of Hochbaum [14]. We use it to improve the best performance ratio known in terms of average degree to $(2\bar{d} + 3)/5$, but show it to be of limited use in terms of maximum degree.

It follows from our analysis that globally minimum degree is not required for **Greedy** to achieve the performance guarantees claimed above; in fact, it holds for any vertex whose degree is at most the average of the degrees of it and its neighbors. This is a locally evaluated property that naturally leads to a parallel and distributed algorithm inheriting the approximative properties of **Greedy**, for the first non-trivial such approximations known to us.

The remainder of the paper is organized as follows. In Section 2 we present the **Greedy** algorithm and some of its properties including the Turán bound, and review other results on approximating this problem. We analyze **Greedy** in detail in Section 3, starting with a generalization of the Turán bound in Section 3.1, followed by tight performance ratios in Section 3.2, and limitations on its performance in Section 3.3. In Section 4, we consider improvements obtained by additionally applying a preprocessing technique, and describe in Section 5 a parallel algorithm attaining the bounds proved for **Greedy**.

1.1 Notation

We use fairly standard graph notation and terminology. For the graph in question, usually denoted $G = (V, E)$, n denotes the number of vertices, Δ the maximum degree, \bar{d} the average degree, α the independence number (the size of the largest independent set), and τ the independence fraction (that is, 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 , $A(G)$ is the size of the solution obtained by A on graph $G = (V, E)$. The performance ratio ρ_A of A is defined by

$$\rho_A = \rho_A(n) = \max_{G, |G|=n} \frac{\alpha(G)}{A(G)}.$$

We focus on **Greedy**, denoted by Gr .

2 The Greedy Algorithm and Related Results

2.1 The Greedy Algorithm

The *minimum-degree greedy* algorithm, or **Greedy** for short, incrementally constructs an independent set by selecting some vertex of minimum degree, removing it and its neighbors from the graph, and iterating on the remaining graph until empty.

```

Greedy(G)
  I ← ∅
  while G ≠ ∅ do

```

```

Choose  $v$  such that  $d(v) = \min_{w \in V(G)} d(w)$ 
 $I \leftarrow I \cup \{v\}$ 
 $G \leftarrow G - \{v\} \cup N(v)$ 
od
Output  $I$ 
end

```

The algorithm can be implemented in time linear in the number of edges and vertices, independent of the degree.

We call a node *critical* if its degree is at most the average of the degrees of it and its neighbors. That is, v is critical if it satisfies:

$$d^2(v) \leq \sum_{w \in N(v)} d(w). \quad (1)$$

A vertex of minimum degree is critical, hence such a node always exists. Although we state **Greedy** with this minimum-degree pivoting rule, the only property that we shall use is that the selected vertex is critical.

Consider an execution of the algorithm to be a sequence of reductions, each corresponding to an iteration. In a reduction, a vertex is selected, added to the solution, and then removed along with its neighborhood from the graph. Let t denote the number of reductions and d_i the degree in the remaining graph of the i -th vertex selected, $i = 1, 2, \dots, t$. The number of vertices removed in the i -th reduction is thus $d_i + 1$.

The main property of the algorithm that we shall be using in our analysis is that the sum of the degrees of the $d_i + 1$ vertices removed in the i -th reduction must be at least $d_i(d_i + 1)$. This allows us to bound from below the number of edges removed in each step.

2.2 Previous results on Greedy

A classical theorem in graph theory, due to Turán [28], states that for any graph G ,

$$\alpha(G) \geq n/(\bar{d} + 1),$$

and that the inequality is tight only for the graph consisting of $\alpha(G)$ disjoint cliques of size as equal as possible. Wei [30] (see [10]) proved an extension,

$$\alpha(G) \geq \sum_{v \in V} 1/(d(v) + 1),$$

which he showed was attained by **Greedy**. This result actually follows from an earlier theorem of Erdős[8] (see [10]), which states that if a graph contains no independent set of size $k + 1$, then there is graph consisting of k disjoint cliques whose degree sequence dominates that of the original graph. The proof of his theorem can also be seen to indirectly refer to **Greedy**.

We include a proof here of the fact that **Greedy** attains the Turán bound, since the proofs of our central results build directly upon it.

Theorem 1 $\alpha(G) \geq \frac{n}{\bar{d} + 1}$

Proof. The proof is a variation of the one given by Hochbaum [14]. The **Greedy** algorithm performs a sequence of t reductions, each time picking a vertex and deleting it and its neighbors from the graph. We count the number of vertices and edges deleted in each reduction.

The removal of vertices in each reduction partitions the vertex set; thus:

$$\sum_{i=1}^t (d_i + 1) = n \tag{2}$$

Since **Greedy** always selects a critical vertex, the sum of the degrees of the vertices deleted in step i is at least $d_i(d_i + 1)$, and thus the number of edges deleted is at least half that amount. Summing over all the reductions:

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

We now add (2) and twice (3), and obtain:

$$(\bar{d} + 1)n \geq \sum_{i=1}^t (d_i + 1)^2.$$

Using the Cauchy-Schwarz inequality and (2), we get:

$$(\bar{d} + 1)n \geq n^2/t.$$

Rearranging the inequality, we obtain the desired bound on t , which is precisely the number of vertices found by **Greedy**. ■

Another simple algorithm operates by a rule that is the inverse of **Greedy**: it deletes a vertex of maximum degree, until no edge remains. Surprisingly, this algorithm also attains the Turán bound, as proved independently by Griggs [10] and Chvátal and McDiarmid [7]. Its approximative properties are however weaker. On the graph with $2s$ vertices that is complete bipartite less a single perfect matching, the algorithm may find only a two vertex independent set for an approximation ratio of $s/2 = (\Delta + 1)/2$. As a result, we do not consider this method further.

An upper bound of $\Delta - 1$ of the performance ratio of **Greedy** (with the minimum-degree rule) can be obtained by rudimentary arguments. Observe that in a graph with minimum degree δ , the independence number is bounded from above by $n\Delta/(\Delta + \delta)$. This is because at least $\delta|I|$ edges must exit an independent set I while at most $\Delta(n - |I|)$ edges can be incident on the remaining vertices. Thus in a regular connected component, the independence number is at most $n/2$, for a ratio of $(\Delta + 1)/2$. This is at most $\Delta - 1$ for $\Delta \geq 3$, and we also know that **Greedy** is optimal when $\Delta = 2$.

Consider now the case of a non-regular component. In each step there is a vertex of degree at most $\Delta - 1$, so **Greedy** finds at least n/Δ independent vertices. As long as **Greedy** selects a vertex of degree one it proceeds optimally, since the optimal solution can contain at most one of the two vertices deleted in each step. Thus, we may assume without loss of generality that the minimum degree is at least two. Hence, by the previous argument, the independence number is at most $n\Delta/(\Delta + 2)$. Combined, this yields a ratio of $\Delta^2/(\Delta + 2)$, which is also at most $\Delta - 1$. The above argument may be what is alluded to in [27, p.306].

2.3 Related results

Any maximal independent set is of size at least $n/(\Delta + 1)$, which results in a trivial performance ratio of $\Delta + 1$. In fact, a ratio of Δ holds, since an optimal solution can contain at most $\Delta|I|$ vertices not already in a maximal solution I .

The theorem of Brooks [5] is an early result that states that any connected graph can be colored with Δ colors unless it consists of $\Delta + 1$ -clique or an odd cycle. Since we can dispose of the exceptions optimally, this yields a stronger bound of n/Δ on the size of the independent set we find. Lovász [22] gave an elegant proof that can be turned into an efficient algorithm.

Hochbaum [14] introduced a preprocessing technique whose effect was to obtain stronger upper bounds on the size of the optimal solution. The technique was based on results of Nemhauser and Trotter [24] on solutions of the linear programming relaxation of the independent set problem. We describe this approach in more detail in Section 4. She obtained approximations of weighted independent sets by applying coloring heuristics and selecting the heaviest color as a solution. In particular, she obtained ratios of $\Delta/2$ and $(d(G) + 1)/2$, where $d(G)$ is the largest minimum degree of any induced subgraph of G , as well as a ratio of $(\bar{d} + 1)/2$ in the unweighted case.

A decade passed with little effect. Independent of the current work, Berman and Fürer [4] followed by Berman and Fujito [3] gave significantly improved ratios of $(\Delta + 3)/5$. The drawback of their methods is the exorbitant time complexity of more than $\exp(32\Delta^{10} \log n)$. Their approach is a local search method, that seeks a larger solution primarily by deleting a moderate number of vertices while adding a greater number.

Khanna et al. [18] also considered a simpler version of this local search strategy, that merely deletes one while adding two or more. They obtained ratios that improved on the $\Delta/2$ bound of [14].

Alongside with the current work, we studied subgraph removal methods that are motivated by results in graph theory that state that graphs without small cliques contain provably larger independent sets. This allows us to obtain performance ratios of $\Delta/6(1 + o(1))$ for graphs of small to fairly large degree, as well as an asymptotic ratio of $O(\Delta/\log \log \Delta)$ [11, 12]. We have also analyzed further the local search method of the previous two papers, decreasing the time complexity needed for the result of [4], and obtained a performance ratio of $(\Delta + 3)/4$ [11] with an efficient version, illustrating further time/approximation tradeoff possibilities.

Halldórsson and Yoshihara [13] have used ideas from the current paper to analyze modified versions of the greedy algorithm schema. In particular, for graphs of maximum degree three, they give a linear-time algorithm that attains a ratio of $9/7 \approx 1.286$.

The preceding has focused on possibility results, but non-approximability results have made great strides in recent years. The independent set problem in bounded-degree graphs (for each fixed Δ) belongs to the class of MAX SNP-complete problems introduced by Papadimitriou and Yannakakis [25]. The groundbreaking results on the theory of interactive proofs, resulting in the landmark paper of Arora et al. [2], show that there is a fixed $\epsilon > 0$ such that approximating the independent set problem in bounded degree graphs within a factor of $1 + \epsilon$ is \mathcal{NP} -hard. Even stronger results hold when degree restrictions are lifted: the independent set problem cannot be approximated with n^ϵ , for some ϵ , unless $\mathcal{P} = \mathcal{NP}$ [2]. Alon et al. [1] have recently been able to scale the latter results down to show that Δ^ϵ approximation on bounded-degree graphs is similarly \mathcal{NP} -hard.

3 Analysis of Greedy

3.1 Relative size of Greedy solutions

We start by strengthening the constructive version of Turán’s theorem, by expressing the size of the obtainable independent set as a function of the independence number. Observe that our bound dominates Turán’s, yielding strict improvements whenever the independence number exceeds the promise of Turán’s bound. Recall that τ is the independence fraction, α/n .

Theorem 2 $Gr \geq \frac{1 + \tau^2}{\bar{d} + 1 + \tau} n$

Proof. Our proof follows that of Theorem 1, while we now additionally keep count the number of vertices deleted that belong to some maximum cardinality independent set.

Fix an independent set of maximum cardinality α , 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 maximum independent set. Then:

$$\sum_{i=1}^t k_i = \alpha. \quad (4)$$

Recall that 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 endpoints in the maximum independent set. Thus, the number of edges deleted is at least $\binom{d_i+1}{2} + \binom{k_i}{2}$. Hence we obtain the following strengthening of ineq. 3:

$$\frac{\bar{d}}{2} n = |E| \geq \sum_{i=1}^t \left(\binom{d_i+1}{2} + \binom{k_i}{2} \right). \quad (5)$$

We now add (2), (4) and twice (5), and apply the Cauchy-Schwarz inequality to obtain:

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

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 3 $Gr \geq \frac{1 - \tau(1 - \tau)}{(1 - \tau)\Delta + 1} n$

Proof. We extend the proof of Theorem 2. 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 have one end in the independent set and one outside, and z_i have both ends outside. Then we have:

$$x_i + 2(y_i + z_i) \geq d_i(d_i + 1), \quad (6)$$

$$y_i \leq k_i(d_i + 1 - k_i), \quad (7)$$

Multiply (7) by -1 (reversing the inequality) and add it to (6) to obtain:

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

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 (5):

$$|E| \geq \sum_{i=1}^t \left(\binom{d_i+1}{2} + \binom{k_i}{2} + \binom{d_i+1-k_i}{2} \right) - z_i \quad (8)$$

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.

$$(n - \alpha)\Delta \geq \sum_{i=1}^t z_i + |E|. \quad (9)$$

Now add twice (8) and twice (9) 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 $\sum_{i=1}^t (d_i + 1) + k_i + (d_i + 1 - k_i)$ and compensate for this by adding $2n$ to the left hand side (invoking (2)). Then:

$$2(n - \alpha)\Delta + 2n \geq \sum_{i=1}^t (d_i + 1)^2 + k_i^2 + (d_i + 1 - k_i)^2. \quad (10)$$

Using (2), (4) and the Cauchy-Schwarz inequality, we obtain:

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

The claim follows from this. ■

3.2 Performance guarantees

The following bound on the performance of **Greedy** on sparse graphs follows easily. We use the bound of Theorem 2 in the denominator of the performance ratio function, use the identity $\alpha = \tau n$ in the numerator, and observe that the ratio is maximized when $\tau = 1$.

Corollary 4 $\rho_{Gr} \leq \frac{\bar{d} + 2}{2}$

For bounded-degree graphs, the general expression obtained in Theorem 3 almost – but not quite – yields our main claim about the performance ratio of **Greedy**. We now proceed to analyze the performance of **Greedy** using a finer scalpel.

Let $t_{d,k}$ be the number of reductions performed by **Greedy** where a vertex of degree d was chosen and exactly k vertices of the independent set were removed. More precisely, for $d = 0, 1, \dots, \Delta$ and $k = 0, 1, \dots, \max(d, 1)$, we define $t_{d,k} = |\{i : d_i = d \text{ and } k_i = k\}|$. With this notation, we may rewrite the constraints (2), (4) and (10) as:

$$\sum_{d,k} (d + 1)t_{d,k} = n, \quad (11)$$

$$\sum_{d,k} kt_{d,k} = \alpha, \quad (12)$$

$$\sum_{d,k} [(d + 1)^2 + k^2 + (d + 1 - k)^2]t_{d,k} \leq 2n(\Delta + 1) - 2\Delta\alpha. \quad (13)$$

We wish to extract from these constraints the best possible lower bound for $t = \sum_{d,k} t_{d,k}$. For this we use the method of multipliers described in Chvátal [6, page 54].

Theorem 5 $\rho_{Gr} \leq (\Delta + 2)/3$

Proof. We need to consider two cases based on the value of Δ .

Case $\Delta \equiv 0, 1 \pmod{3}$ [i.e. $\Delta + 1 \equiv \pm 1 \pmod{3}$]. We construct the linear combination of the constraints (11) and (13), with multipliers $2(\Delta + 1)$ and -1 respectively, and obtain:

$$\sum_{d,k} [2(d+1)(\Delta+1) - ((d+1)^2 + k^2 + (d+1-k)^2)] t_{d,k} \geq 2\Delta\alpha. \quad (14)$$

Let

$$\beta(d) = \min_k [(d+1)^2 + k^2 + (d+1-k)^2]$$

and

$$C(d) = 2(d+1)(\Delta+1) - \beta(d).$$

Set $t_d = \sum_k t_{d,k}$ and conclude from (14) that:

$$\sum_d C(d)t_d \geq 2\Delta\alpha, \quad (15)$$

We show below that, for $d = 0, 1, \dots, \Delta$, $C(d) \leq 2\Delta(\Delta+2)/3$. That with (15) then gives

$$\frac{\Delta+2}{3}t \geq \alpha,$$

as required. It remains only to establish the following claim.

Claim: For $d = 0, 1, \dots, \Delta$, $C(d) \leq \frac{2}{3}\Delta(\Delta+2)$.

Proof of claim: It can easily be verified that:

$$\beta(d) = \begin{cases} \frac{3}{2}(d+1)^2 & \text{if } d \text{ is odd} \\ \frac{3}{2}(d+1)^2 + \frac{1}{2} & \text{if } d \text{ is even} \end{cases}$$

Let $f_0 : \Re \rightarrow \Re$ and $f_1 : \Re \rightarrow \Re$ be defined by:

$$f_0(x) = 2x(\Delta+1) - \frac{3}{2}x^2, \quad f_1(x) = f_0(x) - \frac{1}{2}.$$

Note:

$$C(d) = \begin{cases} f_0(d+1) & \text{if } d \text{ is odd} \\ f_1(d+1) & \text{if } d \text{ is even} \end{cases}$$

Now, $f_0'(x), f_1'(x) = 0$ iff $x = 2(\Delta+1)/3$, and $f_0''(x), f_1''(x) = -3$. Thus, both f_0 and f_1 are concave functions that achieve their unique maximum at $\hat{x} = 2(\Delta+1)/3$. Since f_0 and f_1 are polynomials of degree 2 in x , $f_0(\hat{x} + \epsilon) = f_0(\hat{x} - \epsilon)$ and $f_1(\hat{x} + \epsilon) = f_1(\hat{x} - \epsilon)$, for all ϵ . Thus, to establish the claim, it is enough to verify that f_0 at the nearest even integer to \hat{x} and f_1 at the nearest odd integer to \hat{x} are at most $2\Delta(\Delta+2)/3$.

Let $\Delta + 1 = 3m + r$, where $r = \pm 1$ (since $\Delta + 1 \equiv \pm 1 \pmod{3}$). The nearest even integer to \hat{x} is $2m$ and the nearest odd integer to \hat{x} is $2m + r$. Plugging in, we find that

$$f_0(2m) = f_1(2m+r) = 6m^2 + 4mr = \frac{2}{3}\Delta(\Delta+2),$$

establishing the claim.

Case $\Delta \equiv -1 \pmod{3}$. This time, we construct a linear combination of the constraints (11), (12) and (13), with multipliers $2(\Delta + 1)$, 2 and -1 respectively, and obtain:

$$\sum_{d,k} C(d,k)t_{d,k} \geq 2(\Delta + 1)\alpha, \quad (16)$$

where

$$\begin{aligned} C(d,k) &= 2(d+1)(\Delta+1) + 2k - (d+1)^2 - k^2 - (d+1-k)^2 \\ &= 2(d+1)(\Delta-d) + 2k(d+2-k). \end{aligned}$$

Let $C(d) = \max_k C(d,k)$. To obtain the an upper bound on $C(d)$ in terms of Δ , we consider the functions $f_0 : \mathfrak{R} \rightarrow \mathfrak{R}$ and $f_1 : \mathfrak{R} \rightarrow \mathfrak{R}$ defined as follows.

$$\begin{aligned} f_0(x) &\doteq 2(x+1)(\Delta-x) + \frac{1}{2}(x+2)^2 - \frac{1}{2} = (x+1)(2\Delta - \frac{3}{2}x + \frac{3}{2}) \\ f_1(x) &\doteq 2(x+1)(\Delta-x) + \frac{1}{2}(x+2)^2 = x(2\Delta - \frac{3}{2}x) + 2(\Delta+1) \end{aligned}$$

It can be verified that

$$C(d) = \begin{cases} f_0(d) & \text{if } d \text{ is odd} \\ f_1(d) & \text{if } d \text{ is even} \end{cases}$$

We wish to bound $f_0(d)$ for odd values of d , and $f_1(d)$ for even values of d . Now, $f_0'(x), f_1'(x) = 0$ iff $x = 2\Delta/3$ and $f_0''(x), f_1''(x) = -3$; thus f_0 and f_1 attain their unique maximum at $\hat{x} = 2\Delta/3$. Since f_0 and f_1 are polynomials of degree 2 in x , it is enough to bound f_0 at the nearest odd integer to \hat{x} (i.e. $(2\Delta - 1)/3$), and f_1 at the nearest even integer to \hat{x} (i.e. $2(\Delta + 1)/3$). Plugging in, we find that

$$\begin{aligned} f_0\left(\frac{2\Delta-1}{3}\right) &= \left(\frac{2\Delta-1}{3} + 1\right)\left(2\Delta - \frac{2\Delta-1}{2} + \frac{3}{2}\right) \\ &= \frac{2}{3}(\Delta+1)(\Delta+2); \\ f_1\left(\frac{2(\Delta+1)}{3}\right) &= \frac{2(\Delta+1)}{3}(2\Delta - (\Delta+1)) + 2(\Delta+1) \\ &= \frac{2}{3}(\Delta+1)(\Delta-1+3) \\ &= \frac{2}{3}(\Delta+1)(\Delta+2). \end{aligned}$$

Thus, $C(d) \leq \frac{2}{3}(\Delta+1)(\Delta+2)$, and using (16) we obtain $\frac{\Delta+2}{3}t \geq \alpha$. ■

3.3 Limitations

The performance ratios proved above cannot be improved.

Theorem 6 $\rho_{Gr} \geq \frac{\Delta+2}{3} - O(\Delta^2/n)$, for every $\Delta \geq 3$.

Proof. We give a detailed construction for $\Delta \equiv 1 \pmod{3}$. Consider the following family of graphs H_ℓ , $\ell \geq 2$. We have 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




Figure 1: Initial portion of a hard graph for **Greedy**, $\Delta = 7$

a single perfect matching (i.e. each vertex in the independent set is adjacent to $\ell - 1$ vertices in the following clique). The chain ends with one additional clique.

An instance of this graph with $\ell = 3$ is shown in Figure 1, with the vertices picked by **Greedy** shown in black and the maximum independent set vertices in grey.

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 $(n - \ell)/2\ell + 1$. The optimal solution contains all the independent set vertices for a total of $(n - \ell)/2$. This yields a ratio of

$$\rho_{Gr}(H_\ell) \geq \ell - 2\ell^2/n.$$

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

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

Thus,

$$\rho_{Gr} \geq \frac{\Delta + 2}{3} - O(\Delta^2/n) \tag{17}$$

and

$$\rho_{Gr} \geq \frac{2\bar{d} + 3}{5} - O(\bar{d}^2/n) \tag{18}$$

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

For the case of $\Delta \equiv 0, 2 \pmod{3}$ we need more complicated chains of groups of six subgraphs. For $0 \pmod{3}$, the elements are of the form:

$$K_{\ell-1} - \overline{K_\ell} - K_{\ell-1} - \overline{K_\ell} - K_\ell - \overline{K_{\ell-1}}$$

where K_s ($\overline{K_s}$) denotes a clique (independent set) on s vertices, respectively. In all cases, a clique is completely connected to the following independent set, while connections from the independent set to the next clique miss a single perfect matching. In addition, $\ell - 1$ edges from the first independent set (the second subgraph) go towards the first clique in the next group rather than to the one immediately following. The chain is finished with an additional $2\ell - 1$ clique.




Figure 2: Initial portion of a hard graph for **Greedy**, $\Delta = 6$

An example for $\ell = 3$ is given in Figure 2.

The graphs are designed so that minimum degree will stay as $2\ell - 2$ and equal the degree of the left-most remaining clique. Hence, we may assume that **Greedy** will pick exactly one vertex per clique, or 3 nodes per group, while the number of vertices in the maximum independent set is $3\ell - 1$ per group. Hence, ignoring the end of the chain, the approximation ratio is $\ell - 1/3 = (\Delta + 2)/3$, since maximum degree is $3\ell - 3$.

The case of $\Delta \equiv 2 \pmod{3}$ is similar, with each element of the form:

$$K_{\ell-1} - \overline{K_{\ell+1}} - K_{\ell} - \overline{K_{\ell}} - K_{\ell} - \overline{K_{\ell}}$$

and edges going from the second to the fifth subgraph. We leave the details to the curious reader. ■

We now show that the bound on the performance ratio in terms of average degree (Corollary 4) is optimal in an asymptotic sense.

Theorem 7 $\rho_{Gr} \geq \frac{\bar{d} + 2}{2} - O(1/\bar{d})$

Proof. For each value of Δ , we describe an infinite family of graphs for which the claimed ratio holds. The graphs consists of chains of pairs of subgraphs as in the previous example with the cliques reduced to single vertices. Each pair consists of a vertex adjacent to a ℓ -element independent set each of whose nodes are adjacent to the $\ell - 1$ following single vertex cliques. The chain then ends with a single ℓ -clique that the vertices of the last $\ell - 1$ independent sets are adjacent to. An example is given in Figure 3.




Figure 3: Hard graph for **Greedy** in terms of average degree

There are ℓ^2 edges for each pair consisting of $\ell + 1$ vertices. Thus, if we ignore the end of the chain, the average degree is $2\ell^2/(\ell + 1) = 2\ell - 2 + 2/(\ell + 1)$ and the ratio obtained is:

$$\ell \geq \frac{\bar{d} + 2}{2} - \frac{2}{\bar{d} + 2}.$$

The end of the chain increases the average degree by roughly $O(\ell^2/n)$, which disappears into the lower order term. ■

Variations of the above constructions show that the bounds of Theorems 2 and 3 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.

4 Greedy with preprocessing

Finding a maximum independent set of a graph $G = (V, E)$ can be formulated as an integer programming problem which maximizes $\sum x_i$ over the 0-1 solutions of the system of linear inequalities:

$$x_i \geq 0 \quad \text{for each } v_i \in V$$

and

$$x_i + x_j \leq 1 \quad \text{for each } (v_i, v_j) \in E.$$

If the integrality condition on the solutions is dropped, we obtain the *fractional independent set* problem, or the *linear programming relaxation* of the independent set problem. As shown by Edmonds and Pulleyblank (see [24]), this can be solved efficiently via a bipartite graph constructed as follows: Form two copies of the vertex set V , and let vertices in different copies be adjacent iff they correspond to adjacent vertices in G . Given a characteristic vector $(y_1, y_2, \dots, y_n, y'_1, y'_2, \dots, y'_n)$ of a maximum independent set of this bipartite graph, an optimal LP solution of G is given by:

$$x_i = \frac{1}{2}(y_i + y'_i).$$

By the König-Egervary theorem [20, p.90], the independence number of a bipartite graph equals the number of vertices less the number of edges of a maximum matching. The bipartite matching problem can be computed in time $O(\sqrt{n}|E|)$ by an algorithm of Hopcroft and Karp [15], and hence the LP relaxation of the independent set problem.

We can partition V into the sets ONE, HALF, and ZERO, corresponding to the values of the vertices in a given (arbitrary but fixed) LP solution. Nemhauser and Trotter [24] showed that the set ONE is contained in some optimal independent set O of G . Furthermore, ZERO cannot be contained in O , as otherwise the value of those vertices could be unilaterally increased.

Based on these results, Hochbaum [14] proposed a preprocessing method to improve the performance of approximate independent set algorithms. Compute an LP solution, and apply the approximation algorithm only on the graph H induced by HALF. To the solution found by the algorithm, we can now add the vertices of ONE, possibly resulting in a considerable improvement. The independence number of H can be at most the sum of the values of the LP solution, or $n/2$, so this approach is particularly valuable when the independence number of G is greater than that.

Theorem 8 *Greedy with preprocessing attains a ratio of $(\Delta + 2)/3$, which is tight when $\Delta \equiv 1 \pmod{3}$.*

Proof. The upper bound of $(\Delta + 2)/3$ already follows from Theorem 5, but here it can also be obtained directly from Theorem 3 by arguing that $\tau \leq 1/2$. This is because both the approximate and the optimal solution consist of ONE along with the respective solutions on the graph H , and in H , the independence fraction is at most $1/2$.

For $\Delta \equiv 1 \pmod{3}$, the tightness of this ratio is demonstrated by the graphs given in the proof of Theorem 6. They have the property that, for any independent set I , the number of neighbors of the vertices in I exceeds the size of I . It then follows from Theorem 4 of [24] that any LP solution has all values equal to HALF, i.e. preprocessing is of no help, so $H = G$. ■

We can improve on this bound slightly when $\Delta \equiv 0, 2 \pmod{3}$ to $(\Delta + 2)/3 - 1/(3\Delta + 2)$. 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 the $2.3\bar{3}$ promised by Theorem 5, while there is a graph that forces a ratio of 2.27. We omit the details.

4.1 Average degree

A greater care is needed for a result in terms of average degree. Once preprocessing 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 [14] showed that the Turán bound on **Greedy** can be complemented with the $\tau \leq 1/2$ promise to yield a performance ratio bound of $(\bar{d} + 1)/2$. We follow her lead to obtain a similar result as expected from Theorem 2.

Theorem 9 $\rho_{Gr+Pre} \leq \frac{2\bar{d} + 3}{5}$

Proof. Let d', τ', n' denote the average degree, independence fraction, and number of vertices of H respectively. Recall that H is the subgraph induced by **HALF**, and note that $\tau' \leq 1/2$.

Following [14], we use two properties. The first is that $|\mathbf{ONE}| \geq |\mathbf{ZERO}|$, as otherwise the LP solution value is smaller than if **HALF** = V . The second property is that the number of edges in G is at least $\frac{d'}{2}n' + |\mathbf{ONE}| + |\mathbf{ZERO}|$, since we may assume that G is a connected graph and that n' is positive. Thus, the average degree of G is bounded by:

$$\bar{d} \geq \frac{d'n' + 2|\mathbf{ONE}|}{n' + 2|\mathbf{ONE}|}.$$

The size of the greedy solution will be at least $\frac{1+\tau'^2}{d'+1+\tau'}n' + |\mathbf{ONE}|$, invoking Theorem 2, and the size of the optimal solution $\tau'n' + |\mathbf{ONE}|$. Our goal is now to argue that:

$$\frac{\tau'n' + |\mathbf{ONE}|}{\frac{1+\tau'^2}{d'+1+\tau'}n' + |\mathbf{ONE}|} \leq \frac{(\frac{2}{5}d' + \frac{3}{5})n' + 2|\mathbf{ONE}|}{n' + 2|\mathbf{ONE}|}$$

from which the theorem follows. Cross-multiplying, we find that this entails establishing:

$$\frac{\tau'}{1+\tau'^2}(d'+1+\tau') \leq (2d'+3)/5 \tag{19}$$

and, if $|\mathbf{ONE}| > 0$,

$$2\tau' + 1 \leq 2\frac{1+\tau'^2}{d'+1+\tau'} + \frac{2d'+3}{5}. \tag{20}$$

Ineq. (19) holds because the left hand side is monotone increasing with τ' for $\tau' \leq 1$. To establish (20), we may note that the function $f(x) = 2x + 1 - 2(1+x^2)/(d'+1+x)$ is monotone increasing for $0 \leq x \leq 1/2$. Thus it suffices to replace τ' with $1/2$ in (20) to obtain:

$$2 \leq \frac{5}{2d'+3} + \frac{2d'+3}{5}$$

which is true, since $a/b + b/a$ is always at least 2 when a, b are positive. ■

This again is tight for $\Delta \equiv 1 \pmod{3}$ by Theorem 6, Ineq. (18).

5 Parallel algorithm

The minimum-degree 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, and offers little freedom for heuristic improvements. Fortunately, this is a case where the analysis guides us towards the design of better and/or more general algorithms.

As observed in Section 2, it suffices to select a *critical vertex*, i.e. one satisfying (1). This has some interesting implications. For one, it opens up the possibility of the design of heuristics using secondary selection rules or ordering heuristics while retaining the performance ratios of Greedy.

Another application is a straightforward derivation of a parallel as well as a distributed approximation algorithm attaining the same performance ratios. 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 four or greater) can be selected and processed concurrently.

This suggests a natural approach to a parallel algorithm. Let G^3 denote the graph obtained by taking the adjacency matrix of G to the third power – two vertices are adjacent in G^3 if they are within distance three in G . The adjacency matrices used here are assumed to contain 1's on the diagonal. For a vertex subset S , let $N(S)$ denote the set of vertices adjacent to some node in S .

```

ParallelGreedy( $G$ )
   $I \leftarrow \emptyset$ 
  while ( $V(G) \neq \emptyset$ )
     $W \leftarrow \{v \in V(G) : v \text{ is critical}\}$ 
     $H \leftarrow$  subgraph of  $G^3$  induced by  $W$ 
     $MIS \leftarrow$  maximal independent set of  $H$ 
     $I \leftarrow I \cup MIS$ 
     $G \leftarrow G - (MIS \cup N(MIS))$ 
  od
  return  $I$ 
end

```

We assume the PRAM model, with a finer distinction depending on the MIS algorithm used. We remark that this algorithm can also be implemented in the distributed model of computation [21].

The following lemma due to Alon and Szegedy (private communication) shows that a significant fraction of the vertices must simultaneously satisfy property (1).

Lemma 10 *At least $\frac{4}{\Delta^2+4}n$ vertices are critical.*

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

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

which implies the lemma. As observed by Shearer [26], $E_v[D_v] = 0$. 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 11 *ParallelGreedy finds an independent set of size and performance satisfying Theorems 3 and 5 in time $O(\log^* n \min(\text{poly}(\Delta) \log n, (\log \Delta)\Delta!))$ using n processors in the EREW model.*

Proof. Each vertex added to the solution will satisfy property (1) 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. From Lemma 10, the size of W is at least n/Δ^2 . Every vertex in W is reachable (in G) from $MIS \cup N(MIS)$ by a (not necessarily simple) path of length 2. Since the number of vertices reachable from a fixed vertex in two steps is at most Δ^2 , we have:

$$|MIS \cup N(MIS)|\Delta^2 \geq |W| \geq n/\Delta^2.$$

Thus, the number of deleted vertices, that is $|MIS \cup N(MIS)|$, is at least n/Δ^4 . Hence, the number of iterations is at most $\Delta^4 \log n$.

Also notice that for any vertex v in the graph, some vertex u of distance at most $\Delta - 1$ is critical. Either u is selected or some vertex within distance 2 of u , and thus within distance at most $\Delta + 1$ of v . There are at most $\Delta^{\Delta+1}$ such nodes, and hence, number of rounds. A more careful counting shows that the number of rounds is at most $\Delta!$.

The only non-trivial step in each round is the computation of a maximal independent set of the graph H . An algorithm of Goldberg et al. [9] finds a maximal independent set in time $O(\log \Delta(H)(\Delta^2(H) + \log^* n))$ using a linear number of processors.

The combined time complexity is therefore bounded by $O((\Delta^6 + \log^* n)(\log \Delta) \min(\Delta^4 \log n, \Delta!))$. Note that this is $O(\log^* n)$ on constant degree graphs. ■

While the above algorithm yields a solution satisfying Theorem 2, its time complexity is not well bounded in terms of the average degree \bar{d} . Fortunately, this can be attained by first deleting all vertices of high degree.

Lemma 12 *Theorem 2 and Corollary 4 hold even if we first eliminate (simultaneously) all vertices of degree at least $3\bar{d} + 4$.*

Proof. Consider the subgraph induced by vertices of degree less than $3\bar{d} + 4$. Let $(1 - \gamma)n$, d' , τ' denote the number of vertices, average degree, and independence fraction, respectively. Thus, γ represents the fraction of the vertices that were of high degree. The proposition is that

$$LHS = \frac{1 + \tau'^2}{d' + 1 + \tau'}(1 - \gamma)n \geq \frac{1 + \tau^2}{\bar{d} + 1 + \tau}n = RHS. \quad (21)$$

At least $((3\bar{d} + 3)/2)\gamma n$ edges are deleted, so

$$d' \leq \frac{2|E| - (6|E|/n + 4)\gamma n}{n - \gamma n} \leq \frac{\bar{d}(1 - 3\gamma) - 4\gamma}{1 - \gamma}.$$

The independence number of the remaining graph is at least the original value less the number of deleted vertices, and thus $\tau' \geq \tau - \gamma$. We shall consider two cases, depending on the value of $\tau - \gamma$.

Case $\tau - \gamma \geq 1/(\bar{d} + 1)$. *LHS* is minimized when τ' is at its minimum, which by Turán's theorem is $1/(\bar{d} + 1)$. Thus, we plug in $\tau - \gamma$ for τ' , for:

$$\frac{(1 + (\tau - \gamma)^2)(1 - \gamma)}{\bar{d}(1 - 3\gamma) - 4\gamma + (1 + \tau)(1 - \gamma)} n. \quad (22)$$

The numerator is at least $(1 + \tau^2)(1 - 2\gamma)(1 - \gamma)$, which is at least $(1 + \tau^2)(1 - 3\gamma)$. Also, $(1 + \tau)(1 - \gamma) - 4\gamma$ is at most $(1 + \tau)(1 - 3\gamma)$, since $\tau \leq 1$. Hence, (22) is at least *RHS*.

Case $\tau - \gamma \leq 1/(\bar{d} + 1)$. *LHS* is at least

$$\frac{1}{d' + 1}(1 - \gamma)n \geq \frac{(1 - \gamma)^2}{(1 - 3\gamma)(\bar{d} + 1)} n = \frac{1 + \gamma + (4\gamma^2)/(1 - 3\gamma)}{\bar{d} + 1} n. \quad (23)$$

On the other hand, *RHS* is monotone increasing with τ , so when $\tau \leq 1/(\bar{d} + 1) + \gamma$, its value is at most

$$\frac{1 + (1/(\bar{d} + 1) + \gamma)^2}{\bar{d} + 1 + 1/(\bar{d} + 1) + \gamma} n \leq \frac{1 + \gamma + \gamma^2}{\bar{d} + 1}. \quad (24)$$

Since (23) is at least (24), (21) holds. ■

By deleting first all high-degree vertices (in parallel) before applying `ParallelGreedy`, we obtain an efficient approximation in terms of \bar{d} .

Corollary 13 *There is a EREW parallel algorithm that finds an independent set satisfying Theorem 2 and Corollary 4 in time $O(\log^* n \min(\text{poly}(\bar{d}) \log n, \bar{d}!))$ using n processors.*

The preprocessing method of last section requires the solution of bipartite matching, for which no deterministic parallel algorithms are known. However, randomized parallel algorithms are known [17], and thus the same applies to Theorem 9. An efficient deterministic approach that nearly matches the bound of Theorem 9 is to use either the complement of a maximal matching or the output of the above algorithm, whichever independent set is larger. As we argue in [12], this approach yields a performance ratio of $(2\bar{d} + 4.5)/5$.

Acknowledgments We thank Noga Alon and Mario Szegedy for kindly proving Lemma 10, Vasék Chvátal for comments and helpful advice, Dorit Hochbaum for comments and corrections, and the anonymous referees for thorough important corrections.

References

- [1] N. Alon, U. Feige, A. Wigderson, and D. Zuckerman. Derandomized graph products. *Computational Complexity*, 1995.
- [2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and hardness of approximation problems. In *Proc. 33rd Ann. IEEE Symp. on Found. of Comp. Sci.*, pages 14–23, Oct. 1992.
- [3] P. Berman and T. Fujito. On the approximation properties of independent set problem in degree 3 graphs. In *Fourth Workshop on Algorithms and Data Structures*, Aug. 1995.
- [4] P. Berman and M. Fürer. Approximating maximum independent set in bounded degree graphs. In *Proc. Fifth Ann. ACM-SIAM Symp. on Discrete Algorithms*, pages 365–371, Jan. 1994.

- [5] R. L. Brooks. On coloring the nodes of a network. *Proc. Cambridge Philos. Soc.*, 37:194–197, 1941.
- [6] V. Chvátal. *Linear Programming*. Freeman, New York, 1983.
- [7] V. Chvátal and C. McDiarmid. Small transversals in hypergraphs. *Combinatorica*, 12(1):19–26, 1992.
- [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. *SIAM J. Disc. Math.*, 1(4):434–446, Nov. 1988.
- [10] J. R. Griggs. Lower bounds on the independence number in terms of the degrees. *J. Combin. Theory Ser. B*, 34:22–39, 1983.
- [11] M. M. Halldórsson and J. Radhakrishnan. Improved approximations of independent sets in bounded-degree graphs. In *Proc. Fourth Scand. Workshop on Algorithm Theory*, pages 195–206. Springer-Verlag LNCS #824, 1994.
- [12] M. M. Halldórsson and J. Radhakrishnan. Improved approximations of independent sets in bounded-degree via subgraph removal. *Nordic J. Computing*, 1(4):475–492, 1994.
- [13] M. M. Halldórsson and K. Yoshihara. Greedy approximations of independent sets in low degree graphs. In *Proc. Sixth Intl. Symp. on Algorithms and Computation*. Springer LNCS, Dec. 1995.
- [14] D. S. Hochbaum. Efficient bounds for the stable set, vertex cover, and set packing problems. *Disc. Applied Math.*, 6:243–254, 1983.
- [15] J. Hopcroft and R. Karp. An $n^{5/2}$ algorithm for maximal matchings in bipartite graphs. *SIAM J. Comput.*, 4:225–231, 1973.
- [16] 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.
- [17] R. Karp and V. Ramachandran. A survey of parallel algorithms for shared-memory machines. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 17, pages 869–941. Elsevier Science Publishers B.V., 1990.
- [18] S. Khanna, R. Motwani, M. Sudan, and U. Vazirani. On syntactic versus computational views of approximability. In *Proc. 35th Ann. IEEE Symp. on Found. of Comp. Sci.*, pages 819–830, 1994.
- [19] E. Kubicka, G. Kubicki, and D. Kountanis. Approximation algorithms for the chromatic sum. In *Proc. 1st Great Lakes Computer Science Conf.* Springer LNCS # 507., Oct. 1989.
- [20] E. Lawler. *Combinatorial Optimization: Networks and Matroids*. Holt, Rinehart and Winston, 1976.
- [21] N. Linial. Locality in distributive algorithms. *SIAM J. Comput.*, 21:193–201, 1992.
- [22] L. Lovász. Three short proofs in graph theory. *J. Combin. Theory Ser. B*, 19:269–271, 1975.

- [23] C. McDiarmid. Colouring random graphs. *Annals of Operations Research*, 1:183–200, 1984.
- [24] G. L. Nemhauser and L. Trotter. Vertex packings: Structural properties and algorithms. *Math. Programming*, 8:232–248, 1975.
- [25] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. Syst. Sci.*, 43:425–440, 1991.
- [26] J. B. Shearer. A note on the independence number of triangle-free graphs. *Discrete Math.*, 46:83–87, 1983.
- [27] H. U. Simon. On approximate solutions for combinatorial optimization problems. *SIAM J. Disc. Math.*, 3(2):294–310, May 1990.
- [28] P. Turán. On an extremal problem in graph theory (in Hungarian). *Mat. Fiz. Lapok*, 48:436–452, 1941.
- [29] Twentieth Century Fox. *Wall Street*. Motion picture, 1987.
- [30] V. K. Wei. A lower bound on the stability number of a simple graph. Technical Memorandum No. 81-11217-9, Bell Laboratories, 1981.