

Article

Not peer-reviewed version

---

# An Approximate Solution to the Minimum Dominating Set Problem: The Furones Algorithm

---

[Frank Vega](#) \*

Posted Date: 26 May 2026

doi: 10.20944/preprints202504.0522.v9

Keywords: dominating set; approximation algorithm; planar graphs; Baker's PTAS; P vs NP



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# An Approximate Solution to the Minimum Dominating Set Problem: The Furones Algorithm

Frank Vega 

Information Physics Institute, 840 W 67th St, Hialeah, FL 33012, USA; vega.frank@gmail.com

## Abstract

The Minimum Dominating Set problem is NP-hard, and the best known polynomial-time approximation factor is  $O(\ln n)$ , which is provably tight unless  $P = NP$ . We present a polynomial-time algorithm that reduces an arbitrary input graph to a planar kernel through forced-vertex extraction, pendant elimination, and a linear-time forest projection, and then applies Baker's PTAS to that kernel. The reduction outside the planar PTAS runs in  $O(n + m)$  worst-case time; with a linear-time fixed- $\epsilon$  implementation of Baker's planar subroutine, the full algorithm is linear. The algorithm is provably within twice the optimum whenever the reduction is tight, and the implementation provides a linear-time sufficient --consistency certificate based on the proved forced-boundary inequality. We give a structural *witness mapping* that injects the post-pruning forced-boundary set into the rest of the planar kernel, explaining why pruning often improves the observed bound without turning the remaining gap into an established universal guarantee. We complement the theory with two experiments using Furones v0.2.8 with the --consistency flag: a certified run on 68 NPbench DIMACS clique-complement instances, and a comparison against NetworkX's logarithmic approximation on a 64-graph compendium from VC-Bench. The NetworkX comparison certifies 63 of 64 instances; the lone uncertified instance still has logged upper-bound ratio  $1.603 < 2$ , illustrating incompleteness of the current sufficient certificate rather than a poor empirical outcome.

**Keywords:** dominating set; approximation algorithm; planar graphs; Baker's PTAS; P vs NP

**MSC:** 05C69; 68Q25; 90C27

## 1. Introduction

The **Minimum Dominating Set** (DS) problem is one of the canonical NP-hard problems in graph theory [1]. Given an undirected graph  $G = (V, E)$ , a set  $D \subseteq V$  is a *dominating set* if every vertex of  $V$  is either in  $D$  or has a neighbour in  $D$ ; the goal is to find such a set of minimum cardinality, denoted  $\gamma(G)$ . Dominating sets arise in wireless backbone selection, sensor placement, monitoring infrastructure, and influence maximisation in social networks.

The problem is notoriously hard to approximate. The greedy set-cover reduction yields an  $O(\ln n)$  approximation, and this is essentially the best one can do: Feige [2] showed that set cover (and hence dominating set) cannot be approximated within  $(1 - \epsilon) \ln n$  for any fixed  $\epsilon > 0$  unless  $P = NP$ , and Raz and Safra [3] sharpened this barrier. A direct consequence is that any polynomial-time algorithm achieving a constant approximation ratio — such as 2 — would imply  $P = NP$ .

Despite this barrier, we present an algorithm that is provably a 2-approximation in the tight case and, in the experiments below, often satisfies the linear-time sufficient --consistency certificate. The algorithm, called `find_dominating_set` and shipped in the Furones package (v0.2.8), proceeds in five phases:

1. **Preprocessing:** self-loops are removed and isolated vertices, which must lie in any dominating set, are extracted.

2. **Reduction to a planar kernel:** two reduction rules (forced isolated vertices and pendant elimination with selective neighbour removal) are applied until the working graph  $H$  has minimum degree at least 2. If  $H$  is non-planar, a planar spanning forest is built in linear time, and the reduction is applied a second time. The result is a planar graph  $G_{\text{red}}$  of minimum degree  $\geq 2$ .
3. **Approximate solution on the kernel:** Baker's PTAS [4] is run on  $G_{\text{red}}$  with parameter  $\varepsilon = 1$  (i.e.  $k = \lceil 1/\varepsilon \rceil = 1$ ), yielding a 2-approximation on the planar kernel.
4. **Lifting:** the kernel solution is combined with the forced vertices and the isolated vertices to produce a dominating set of  $G$ .
5. **Pruning:** redundant vertices are greedily removed from the lifted solution while preserving feasibility.

The dominant step in earlier versions was greedy planarisation. The present implementation replaces it with a linear-time forest projection, so the reduction outside the planar PTAS is  $O(n + m)$ . The full Python implementation is available at the link in Table 1.

### Contributions

The main contributions of this paper are the following.

- A practical reduction whose non-PTAS part runs in  $O(n + m)$  time, with a provable 2-approximation guarantee in the *tight* case (the residual planar kernel does not touch any forced vertex).
- A new structural lemma — the *witness mapping* — which injects the post-pruning forced-boundary set  $F_R^{\text{pruned}}$  into the complement  $R \setminus F_R$  of the forced-boundary inside the kernel. This identifies the part of the non-tight case controlled by pruning, while the implemented certificate remains the conservative proved condition on the full forced-boundary set  $F_R$ .
- A theoretical consequence: if the algorithm achieves a 2-approximation universally, then  $P = NP$ , by the  $(1 - \varepsilon) \ln n$  inapproximability lower bound.
- Experimental studies on 68 NP-Bench DIMACS clique-complement instances and a 64-graph compendium from VC-Bench, all run with `--consistency`; the second study also compares against NetworkX's approximation implementation.
- An open-source implementation, Furonos v0.2.8.

The paper is organised as follows. Section 3 describes the algorithm. Section 4 proves that it always outputs a valid dominating set. Section 5 analyses the approximation ratio, develops the witness mapping, and discusses the consequences for the P vs NP question. Section 6 establishes the time complexity. Section 7 reports the experimental results. Section 8 concludes.

## 2. Research Data

The algorithm is implemented in the Python package Furonos, freely available on the Python Package Index [5]. The current version is v0.2.8; code metadata is summarised in Table 1.

**Table 1.** Code metadata for the Furonos package (v0.2.8).

Nr.	Code metadata description	Metadata
C1	Current code version	v0.2.8
C2	Permanent link to code repository	<a href="https://github.com/frankvegadelgado/furonos">https://github.com/frankvegadelgado/furonos</a>
C3	Permanent link to reproducible capsule	<a href="https://pypi.org/project/furonos/">https://pypi.org/project/furonos/</a>
C4	Legal Code License	MIT License
C5	Code versioning system used	git
C6	Languages, tools, and services used	Python $\geq 3.12$ , NetworkX $\geq 3.0$
C7	Compilation requirements and dependencies	Python, NetworkX, itertools

### 3. Description of the Algorithm

We now present the pseudo-code of `find_dominating_set`. The full Python implementation is shipped with the `Furones` package. The implementation also supports a linear-time sufficient `--consistency` flag: when enabled, the routine returns only if the current proof certifies the 2-approximation bound; otherwise it raises a certification error rather than claiming an unproved universal guarantee.

---

#### Algorithm 1 FINDDOMINATINGSET( $G, \varepsilon, consistency$ )

---

**Require:** Undirected graph  $G = (V, E)$ , approximation parameter  $\varepsilon \in (0, 1]$  (default  $\varepsilon = 1$ ), Boolean *consistency* (default false)

**Ensure:** A dominating set  $D \subseteq V$ ; if *consistency* is true, the returned set is certified by the linear-time sufficient conditions below

```

1: Remove all self-loops from G
2:  $I_{iso} \leftarrow \{v \in V : \deg(v) = 0\}$  ▷ isolated vertices
3:  $G \leftarrow G \setminus I_{iso}$ 
4: if G has no vertices then return  $I_{iso}$ 
5: end if
6:  $H \leftarrow G.copy()$ 
7:  $D_{forced} \leftarrow \emptyset$ 
8: RUNCASCADE( $H, D_{forced}, G$ ) ▷ DS reduction rules
9: if H is non-planar then
10:    $P \leftarrow SPANNINGFOREST(H)$ 
11:    $H \leftarrow P$ 
12:   RUNCASCADE( $H, D_{forced}, G$ )
13: end if
14: if H has vertices then
15:    $D_{red} \leftarrow BAKERPTAS(H, \varepsilon)$  ▷  $(1 + \varepsilon)$ -approximation on planar graph
16: else
17:    $D_{red} \leftarrow \emptyset$ 
18: end if
19:  $D \leftarrow PRUNEREDUNDANT(G, D_{forced} \cup D_{red} \cup I_{iso})$ 
20: if D is not a dominating set of G then raise RUNTIMEERROR
21: end if
22: if consistency and not CERTIFIEDTWOAPPROX( $G, H, D_{forced}, D \setminus I_{iso}$ ) then
23:   raise APPROXIMATIONNOTCERTIFIEDERROR
24: end if
25: return D

```

---

The consistency check is deliberately conservative and linear in the size of the graph after the dominating-set verification. It computes

$$F_R = \{v \in V(H) : N_G(v) \cap D_{forced} \neq \emptyset\}$$

and accepts when either  $F_R = \emptyset$  (the tight case of Theorem 2) or  $|D_{forced}| \geq 2|F_R|$  (the sufficient condition in Corollary 1). This check intentionally uses the full forced-boundary set, not only the post-

pruning subset  $F_R^{\text{pruned}}$ . Thus consistency=True does not solve the open universal case; it prevents the implementation from reporting a 2-approximation certificate on instances outside the currently proved conditions.

The reduction cascade (Algorithm 2) implements two rules.

- **Rule 0 (isolated vertex).** If  $\deg_H(v) = 0$  and  $v$  is not already dominated through an edge of  $G$  by some vertex in  $D_{\text{forced}}$ , add  $v$  to  $D_{\text{forced}}$ . In either case, remove  $v$  from  $H$ .
- **Rule 1 (pendant vertex).** If  $\deg_H(v) = 1$  with unique  $H$ -neighbour  $u$  and  $v$  is not already dominated, add  $u$  to  $D_{\text{forced}}$ . Remove both  $u$  and  $v$  from  $H$ , and additionally remove every neighbour  $w$  of  $u$  whose  $H$ -neighbours all lie inside  $\{u\} \cup N_H(u)$ . Neighbours of  $u$  that have at least one edge leaving  $\{u\} \cup N_H(u)$  are kept — these are the only vertices that may end up in the boundary set  $F_R$ .

The rules are applied repeatedly until no vertex of degree at most 1 remains.

---

### Algorithm 2 RUNCASCADE( $H, D_{\text{forced}}, G$ )

---

**Require:** Graph  $H$  (modifiable), forced set  $D_{\text{forced}}$ , original graph  $G$

**Ensure:**  $H$  has minimum degree  $\geq 2$  after removal

```

1:  $Q \leftarrow \{v \in V(H) : \deg_H(v) \leq 1\}$ 
2: while  $Q \neq \emptyset$  do
3:    $v \leftarrow Q.\text{pop}()$ 
4:   if  $v \notin V(H)$  then continue
5:   end if
6:   if  $\deg_H(v) = 0$  then
7:     if  $\exists u \in G.\text{neighbors}(v) \cap D_{\text{forced}}$  then  $H.\text{remove}(v)$ 
8:     else  $D_{\text{forced}}.\text{add}(v)$ ;  $H.\text{remove}(v)$ 
9:     end if
10:  else if  $\deg_H(v) = 1$  then
11:     $u \leftarrow$  unique neighbour of  $v$  in  $H$ 
12:    if  $\exists w \in G.\text{neighbors}(v) \cap D_{\text{forced}} \setminus \{u\}$  then
13:       $H.\text{remove}(v)$ 
14:      if  $\deg_H(u) \leq 1$  then  $Q.\text{add}(u)$ 
15:      end if
16:    else
17:       $D_{\text{forced}}.\text{add}(u)$ 
18:       $N_u \leftarrow H.\text{neighbors}(u)$ 
19:       $R \leftarrow \{u, v\} \cup \{w \in N_u : \forall x \in H.\text{neighbors}(w), x \in N_u \cup \{u\}\}$ 
20:       $H.\text{remove\_nodes}(R)$ 
21:      for  $w \in N_u \setminus R$  do
22:        if  $\deg_H(w) \leq 1$  then  $Q.\text{add}(w)$ 
23:        end if
24:      end for
25:    end if
26:  end if
27: end while

```

---

The planar subgraph construction is given in Algorithm 3. The implementation uses the fastest safe projection: if the residual graph  $H$  is already planar, it is kept unchanged; otherwise a depth-first search keeps only discovery edges. The result is a spanning forest, hence planar by construction. This removes both the  $O(m \log m)$  edge sort and the repeated planarity checks used in earlier greedy versions. The trade-off is that dense non-planar residuals may lose cycle-closing edges that a maximal planar subgraph could have kept.

---

**Algorithm 3** SPANNINGFOREST( $H$ )
 

---

**Require:** Graph  $H$  (possibly non-planar)

**Ensure:** Planar forest  $P$  with  $V(P) = V(H)$  and  $E(P) \subseteq E(H)$

```

1: if ISPLANAR( $H$ ) then return  $H$ .copy()
2: end if
3:  $P \leftarrow$  Graph( $V(H)$ ) with no edges
4:  $S \leftarrow \emptyset$  ▷ visited vertices
5: for  $r \in V(H)$  do
6:   if  $r \in S$  then continue
7:   end if
8:   Push  $r$  onto a stack and add  $r$  to  $S$ 
9:   while the stack is non-empty do
10:     $u \leftarrow$  POP()
11:    for  $v \in N_H(u)$  do
12:      if  $v \notin S$  then
13:        add  $v$  to  $S$ ; push  $v$ ;  $P$ .add_edge( $u, v$ )
14:      end if
15:    end for
16:  end while
17: end for
18: return  $P$ 

```

---

Finally, Baker's PTAS for planar dominating set is invoked with  $\varepsilon = 1$ , which yields a 2-approximation on the planar kernel  $G_{\text{red}}$ . With  $k = \lceil 1/\varepsilon \rceil = 1$ , the underlying treewidth dynamic programme runs in  $O(3^k n) = O(n)$  time on planar graphs [4,6].

#### 4. Correctness of the Algorithm

**Theorem 1.** *Algorithm 1 always returns a valid dominating set of the input graph  $G$ .*

**Proof.** We argue that every vertex of  $G$  is either contained in the output set  $D$  or adjacent to a vertex of  $D$ . Three classes of vertices arise.

**Isolated vertices.** Any vertex of degree 0 in  $G$  is moved to  $I_{\text{iso}}$  and added to  $D$ , hence trivially dominated.

**Vertices removed during the cascade.** The cascade modifies only the working copy  $H$ . Each removal is justified as follows.

- If  $v$  is isolated in  $H$  but already has a neighbour in  $D_{\text{forced}}$  via an edge of  $G$ , then  $v$  is already dominated and can be safely removed.

- Otherwise, an isolated  $v$  has no surviving neighbour and can only be dominated by itself, so it is forced into  $D_{\text{forced}}$ .
- For a pendant  $v$  with unique  $H$ -neighbour  $u$ : if  $v$  is already dominated through some other neighbour,  $v$  is removed safely; otherwise the exchange argument (Proposition 1) certifies the existence of an optimal DS that contains  $u$ , so forcing  $u$  is correct. The vertices removed jointly with  $u$  and  $v$  are either  $u$ ,  $v$ , or other neighbours of  $u$  that have no external edges — once  $u \in D_{\text{forced}}$  all such vertices are dominated by  $u$ .

After the cascade, every removed vertex is either in  $D_{\text{forced}}$  or adjacent to a vertex of  $D_{\text{forced}}$ .

**Vertices that survive in the planar kernel  $G_{\text{red}}$ .** Greedy planarisation (Algorithm 3) only deletes edges, so every vertex of  $H$  remains. Each edge of  $G_{\text{red}}$  is also an edge of  $G$ , so any domination relation witnessed in  $G_{\text{red}}$  is equally valid in  $G$ . Baker's PTAS returns a feasible dominating set  $D_{\text{red}}$  of  $G_{\text{red}}$ , which therefore dominates  $V(G_{\text{red}})$  in  $G$  as well.

The union  $D_{\text{forced}} \cup D_{\text{red}} \cup I_{\text{iso}}$  dominates every vertex of  $G$ . Pruning only removes a vertex  $v$  if (i)  $v$  has another neighbour in the current  $D$  and (ii) every neighbour of  $v$  either lies in  $D$  or has at least two  $D$ -neighbours; both conditions guarantee that the remaining set is still a dominating set.  $\square$

## 5. Approximation Ratio and Consequences for P vs NP

### 5.1. Structural Setup

Throughout this section we fix the following notation. Let  $G$  be the input graph,  $F = D_{\text{forced}}$  the set of forced vertices produced by the cascade,  $R = V(G_{\text{red}})$  the vertex set of the residual planar kernel, and  $D$  the algorithm's output. The sets  $F$  and  $R$  are disjoint by construction, and we write  $\gamma(\cdot)$  for the size of a minimum dominating set.

The *forced-boundary set* is

$$F_R = \{v \in R : N_G(v) \cap F \neq \emptyset\},$$

the vertices of  $G_{\text{red}}$  that have at least one forced neighbour in  $G$ . When  $F_R = \emptyset$  we call the reduction *tight*; otherwise *non-tight*. The implementation exposes this flag as `no_bridge_forced` in `verify_reduction_ds`.

### 5.2. Exchange Argument and Optimality of Forced Vertices

**Proposition 1** (Optimality of forced vertices). *There exists a minimum dominating set  $D^*$  of  $G$  with  $F \subseteq D^*$ .*

**Proof.** We treat each forcing rule separately.

**Rule 0 (isolated  $v$ ).** A vertex with no neighbours can only be dominated by itself, hence belongs to every dominating set of  $G$ .

**Rule 1 (pendant  $v$  with unique  $H$ -neighbour  $u$ ).** Any DS must dominate  $v$ , so  $v \in D$  or  $u \in D$ . If  $D$  contains  $v$  but not  $u$ , then  $D' = (D \setminus \{v\}) \cup \{u\}$  has the same size and satisfies  $N_G[u] \supseteq \{u, v\} = N_G[v]$ , so  $D'$  also dominates everything  $D$  dominated. Iterating this exchange over all pendant reductions yields an optimal DS containing every vertex in  $F$ .  $\square$

### 5.3. Key Inequalities

**Lemma 1** (Reduction bounds). *The following inequalities hold:*

- (i)  $\gamma(G) \leq |F| + \gamma(G_{\text{red}})$ .
- (ii)  $\gamma(G_{\text{red}}) \leq \gamma(G) - |F| + |F_R|$ .
- (iii)  $|D| \leq 2\gamma(G) - |F| + 2|F_R|$ .

**Proof.** (i) Let  $D_r$  be any dominating set of  $G_{\text{red}}$ . Then  $F \cup D_r$  dominates  $G$ : forced vertices are covered by themselves; vertices removed during the cascade are covered by their forced neighbour in  $F$ ; and vertices in  $R$  are covered by  $D_r$  via edges of  $G_{\text{red}} \subseteq G$ . Minimising over  $D_r$  gives  $\gamma(G) \leq |F| + \gamma(G_{\text{red}})$ .

(ii) Fix  $D^*$  with  $F \subseteq D^*$  and  $|D^*| = \gamma(G)$  (Proposition 1). Set  $D_R^* = D^* \cap R$ , so  $|D_R^*| = \gamma(G) - |F|$ . The set  $D_R^*$  may fail to dominate some vertex  $v \in R$ : this happens precisely when  $v$  is dominated in  $D^*$  exclusively by a forced neighbour, in which case  $v \in F_R$ . For every such  $v$  pick a neighbour  $s_v \in N_{G_{\text{red}}}(v)$ , which exists because  $G_{\text{red}}$  has minimum degree  $\geq 2$ . Then  $D_R^* \cup \{s_v\}$  is a dominating set of  $G_{\text{red}}$  of size at most  $|D_R^*| + |F_R| = \gamma(G) - |F| + |F_R|$ .

(iii) Baker's PTAS with  $k = 1$  returns  $D_{\text{red}}$  with

$$|D_{\text{red}}| \leq \frac{k+1}{k} \gamma(G_{\text{red}}) = 2 \gamma(G_{\text{red}}).$$

Before pruning, the lifted solution has size  $|F| + |D_{\text{red}}| \leq |F| + 2 \gamma(G_{\text{red}})$ , and pruning can only decrease this. Substituting (ii) gives

$$|D| \leq |F| + 2(\gamma(G) - |F| + |F_R|) = 2\gamma(G) - |F| + 2|F_R|.$$

□

#### 5.4. Tight Case: Guaranteed 2-Approximation

**Theorem 2** (2-approximation for tight reductions). *If the reduction is tight ( $F_R = \emptyset$ ), then  $|D| \leq 2\gamma(G)$ .*

**Proof.** When  $F_R = \emptyset$ , every  $v \in R$  has no forced neighbour, so in any optimal DS  $D^*$  with  $F \subseteq D^*$  each  $v \in R$  must be dominated by some vertex in  $D^* \cap R$ . Hence  $D_R^* = D^* \cap R$  already dominates  $G_{\text{red}}$ , giving  $\gamma(G_{\text{red}}) \leq \gamma(G) - |F|$ , i.e.  $\gamma(G) = |F| + \gamma(G_{\text{red}})$ . Applying Lemma 1(iii) with  $F_R = \emptyset$  yields  $|D| \leq 2\gamma(G) - |F| \leq 2\gamma(G)$ . □

#### 5.5. Non-Tight Case: The Witness Mapping and a Refined Bound

When  $F_R \neq \emptyset$ , Lemma 1(iii) gives  $|D| \leq 2\gamma(G) - |F| + 2|F_R|$ . By itself this implies a 2-approximation only when  $|F| \geq 2|F_R|$ . The pruning step is designed to recover the slack: many vertices of  $F_R$  that Baker's PTAS chooses are then deleted as redundant, because their forced neighbour in  $F$  already dominates them in  $G$ .

Vertices in  $F_R$  arise exclusively from the selective-removal step of Rule 1: when a pendant  $v$  triggers the forcing of its neighbour  $u$ , every neighbour of  $u$  that has at least one edge leaving  $\{u\} \cup N_H(u)$  is kept in  $H$ , and these kept vertices are precisely the candidates for  $F_R$ . By construction, every  $w \in F_R$  has a forced  $G$ -neighbour, namely  $u$ , and minimum degree at least 2 in  $G_{\text{red}}$ .

Let

$$F_R^{\text{pruned}} = F_R \cap D$$

denote the subset of forced-boundary vertices that survive pruning. We give two complementary results: a structural witness mapping for  $F_R^{\text{pruned}}$  (Theorem 3), and a quantitative consequence (Corollary 1) which formalises the conservative full-boundary certificate currently used by the implementation.

**Theorem 3** (Witness mapping). *For every  $w \in F_R^{\text{pruned}}$  there exists a vertex  $x_w \in R \setminus F_R$  with*

- $x_w \notin D$ ,
- $\{w, x_w\} \in E(G)$ , and
- $w$  is the unique  $D$ -neighbour of  $x_w$  in  $G$ .

Moreover the assignment  $w \mapsto x_w$  is injective. Consequently

$$|F_R^{\text{pruned}}| \leq |R \setminus F_R|.$$

**Proof.** Let  $w \in F_R^{\text{pruned}}$ . Because  $w \in F_R$ , it has a forced  $G$ -neighbour  $u \in F \subseteq D$ , so condition (a) of the pruning rule ("w has a  $D$ -neighbour") is satisfied. The fact that  $w$  was nevertheless retained means condition (b) of the pruning rule must fail for  $w$ : there exists a  $G$ -neighbour  $x_w$  of  $w$  with  $x_w \notin D$  and

$|N_G(x_w) \cap D| < 2$ . Since  $w \in D$  contributes one to that count,  $w$  is the unique  $D$ -neighbour of  $x_w$ . This proves (a), (b) and (c).

We now locate  $x_w$ . First,  $x_w \notin F$ , because  $F \subseteq D$  but  $x_w \notin D$ . Second,  $x_w$  cannot have been removed during the cascade: every cascade-removed vertex either is itself forced or has a forced  $G$ -neighbour, both of which would contribute a  $D$ -neighbour distinct from  $w$  (recall  $w \in R$  and  $F \cap R = \emptyset$ ), contradicting (c). Hence  $x_w \in R$ . Finally, if  $x_w$  had a forced neighbour, the forced neighbour would be in  $D$  and distinct from  $w$ , again contradicting (c); so  $x_w \notin F_R$ , i.e.  $x_w \in R \setminus F_R$ .

For injectivity, suppose  $x_w = x_{w'}$  for some  $w \neq w'$  in  $F_R^{\text{pruned}}$ . Then  $x_w$  would have both  $w$  and  $w'$  as  $D$ -neighbours, contradicting condition (c).  $\square$

Theorem 3 gives the cleanest precise statement available about the post-pruning forced-boundary: every surviving forced-boundary vertex is “earning its keep” by being the sole dominator of a private witness in  $R \setminus F_R$ . The map is into  $R \setminus F_R$  rather than into  $V \setminus F$ , which is the strongest possible target given that vertices in  $F_R$  already have a forced neighbour and therefore cannot serve as witnesses.

**Proposition 2** (Pruning removes forced-boundary redundancy). *A vertex  $w \in F_R$  that lies in  $D_{\text{red}}$  but is not the unique  $D$ -neighbour of any vertex outside  $D$  is removed from  $D$  by PRUNEREDUNDANT.*

**Proof.** Such a  $w$  has its forced neighbour  $u \in F \subseteq D$  available, so condition (a) of the pruning rule holds. Condition (b) of the rule requires that every neighbour  $u'$  of  $w$  either lies in  $D$  or has at least two  $D$ -neighbours; the hypothesis says exactly this. Both conditions hold, so  $w$  is removed.  $\square$

**Corollary 1** (Refined approximation bound). *The output  $D$  of Algorithm 1 satisfies*

$$|D| \leq 2\gamma(G) - |F| + 2|F_R|,$$

and equality is attained only when no vertex of  $F_R$  is pruned. In particular,  $|D| \leq 2\gamma(G)$  whenever  $|F| \geq 2|F_R|$ .

**Proof.** The bound is Lemma 1(iii). Equality requires that every  $F_R$ -vertex chosen by Baker’s PTAS be retained — but Proposition 2 shows that any such vertex with no exclusive role is removed.  $\square$

### 5.6. What the Certificate Does and Does Not Prove

The proved bound from Corollary 1 is universal but loose. The certified NPbench experiment in Section 7 shows that every clique-complement instance in that run satisfies the implemented --consistency check and therefore receives the bound  $|D| \leq 2\gamma(G)$ . The VC-Bench comparison, however, contains one graph, `scc_rt_justinbieber`, for which the sufficient certificate does not fire even though the logged upper-bound ratio is below 2. This is the correct interpretation of a failed consistency run: the implementation has declined to certify the instance, not disproved the usefulness of the returned dominating set.

The post-pruning inequality

$$|F| \geq 2|F_R^{\text{pruned}}|.$$

is therefore not used as an implemented universal certificate. It is a natural strengthening suggested by the witness mapping, because only vertices of  $F_R$  that survive pruning can contribute directly to the final lifted solution. At present, however, the proved and enforced sufficient condition remains the stronger full-boundary inequality  $|F| \geq 2|F_R|$  from Corollary 1. Establishing a sound bound in terms of  $F_R^{\text{pruned}}$ , or finding counterexamples to such a strengthening, is left as future work.

Theorem 3 establishes the strongest one-sided structural fact we have so far: each  $w \in F_R^{\text{pruned}}$  pays for at least one private vertex outside  $F_R$ . What remains open is whether this post-pruning structure can be converted into a sound universal approximation proof.

### 5.7. Consequences for P vs NP

**Theorem 4** (Implication for P vs NP). *If Algorithm 1 with  $\varepsilon = 1$  returns  $|D| \leq 2\gamma(G)$  on every undirected graph  $G$ , then  $P = NP$ .*

**Proof.** A polynomial-time 2-approximation for Minimum Dominating Set would, for every  $\varepsilon > 0$  and every sufficiently large  $n$ , beat the  $(1 - \varepsilon) \ln n$  inapproximability barrier of Feige [2], strengthened by Raz and Safra [3]. Hence such an algorithm forces  $P = NP$ .  $\square$

Whether Algorithm 1 actually attains a universal 2-approximation remains open. Theorem 2 establishes the guarantee unconditionally for tight reductions, while Theorem 3 explains the residual post-pruning structure in the non-tight case. For consistency with this state of knowledge, the implementation's `--consistency` mode is a linear-time sufficient certification mode rather than an exponential verifier or an unconditional guarantee.

## 6. Runtime Analysis

**Theorem 5.** *The reduction and lifting parts of Algorithm 1 run in time  $O(n + m)$ , where  $n = |V|$  and  $m = |E|$ . The full running time is  $O(n + m + T_{\text{BAKER}}(n, \varepsilon))$ , where  $T_{\text{BAKER}}$  is the time used by the planar dominating-set subroutine on the reduced kernel.*

**Proof.** We bound each step.

- **Preprocessing** (self-loops and isolates):  $O(n + m)$ .
- **Cascade** (Rules 0 and 1): each vertex is removed at most once, and each removal scans its neighbour list, giving total work  $O(n + m)$ .
- **Forest projection:** Algorithm 3 first performs one LR planarity test in  $O(n + m)$ . If the residual is already planar, it is returned unchanged. Otherwise, a DFS scan keeps a spanning forest in  $O(n + m)$ . No edge sorting and no repeated planarity checks are performed.
- **Baker's PTAS on the planar kernel:** the kernel has  $n' \leq n$  vertices and  $m' = O(n')$  edges. In the theoretical fixed- $\varepsilon$  Baker framework this step is linear in  $n'$ ; the reference Python implementation uses a practical decomposition/repair routine, so its observed cost is represented by  $T_{\text{BAKER}}(n, \varepsilon)$ .
- **Lifting, pruning, and verification:**  $O(n + m)$ .

Thus the implemented reduction has worst-case linear time, and the full algorithm is linear whenever the planar PTAS backend is implemented in linear time for fixed  $\varepsilon$ .  $\square$

## 7. Experimental Study

### 7.1. Setup

We evaluated Furones v0.2.8 on the NPbench collection of DIMACS clique-complement graphs [7]. These are complements of classical DIMACS maximum-clique instances and are hard benchmark instances for dominating set. The exact command was

```
python -m furones.batch -i .\experiments\ -c -l --consistency,
```

equivalently `batch_asia -i .\experiments\ -c -l --consistency`. The generated log was copied to `experiments/np_bench.txt`; Table 2 is derived only from that log. The time column is the algorithmic phase reported by the line "Our Algorithm ... done in:" and does not include parsing. Because every run completed with `--consistency`, each returned solution is certified by the implemented linear-time sufficient conditions and receives the certified approximation bound  $|D|/\gamma(G) \leq 2$ .

**Table 2.** Certified Furones v0.2.8 results on NPbench DIMACS clique-complement instances.  $T_F$  is the algorithmic time in milliseconds reported in `experiments/np_bench.txt`. The `--consistency` column records that the certification flag was active and did not reject the instance.

Instance	$n$	$m$	$ D $	$T_F$ ms	--consistency	Bound
brock200_1	200	5066	12	79.104	yes	$\leq 2$
brock200_2	200	10024	7	69.520	yes	$\leq 2$
brock200_3	200	7852	7	45.499	yes	$\leq 2$
brock200_4	200	6811	9	41.027	yes	$\leq 2$
brock400_1	400	20077	13	127.964	yes	$\leq 2$
brock400_2	400	20014	13	141.828	yes	$\leq 2$
brock400_3	400	20119	12	127.718	yes	$\leq 2$
brock400_4	400	20035	13	96.426	yes	$\leq 2$
brock800_1	800	112095	13	859.918	yes	$\leq 2$
brock800_2	800	111434	12	653.505	yes	$\leq 2$
brock800_3	800	112267	12	588.356	yes	$\leq 2$
brock800_4	800	111957	12	1035.418	yes	$\leq 2$
c-fat200-1	200	18366	2	76.947	yes	$\leq 2$
c-fat200-2	200	16665	2	81.690	yes	$\leq 2$
c-fat200-5	200	11427	2	60.935	yes	$\leq 2$
c-fat500-1	500	120291	2	533.267	yes	$\leq 2$
c-fat500-10	500	78123	2	379.581	yes	$\leq 2$
c-fat500-2	500	115611	2	509.301	yes	$\leq 2$
c-fat500-5	500	101559	2	706.302	yes	$\leq 2$
C1000.9	1000	49421	36	275.494	yes	$\leq 2$
C125.9	125	787	22	0.000	yes	$\leq 2$
C250.9	250	3141	25	10.970	yes	$\leq 2$
C500.9	500	12418	29	66.273	yes	$\leq 2$
gen200_p0.9_44	200	1990	24	22.267	yes	$\leq 2$
gen200_p0.9_55	200	1990	23	17.404	yes	$\leq 2$
gen400_p0.9_55	400	7980	31	143.843	yes	$\leq 2$
gen400_p0.9_65	400	7980	31	47.538	yes	$\leq 2$
gen400_p0.9_75	400	7980	30	43.435	yes	$\leq 2$
hamming10-2	1024	5120	172	38.668	yes	$\leq 2$
hamming10-4	1024	89600	18	535.905	yes	$\leq 2$
hamming6-2	64	192	15	0.000	yes	$\leq 2$
hamming6-4	64	1312	2	15.273	yes	$\leq 2$
hamming8-2	256	1024	48	9.930	yes	$\leq 2$
hamming8-4	256	11776	6	70.361	yes	$\leq 2$
johnson16-2-4	120	1680	8	7.148	yes	$\leq 2$
johnson32-2-4	496	14880	16	64.438	yes	$\leq 2$
johnson8-2-4	28	168	4	0.000	yes	$\leq 2$
johnson8-4-4	70	560	8	0.000	yes	$\leq 2$
keller4	171	5100	7	27.093	yes	$\leq 2$
keller5	776	74710	16	413.935	yes	$\leq 2$
MANN_a27	378	702	39	10.770	yes	$\leq 2$
MANN_a45	1035	1980	66	47.806	yes	$\leq 2$
MANN_a81	3321	6480	120	173.083	yes	$\leq 2$
MANN_a9	45	72	14	12.538	yes	$\leq 2$
p_hat1000-3	1000	127754	20	743.294	yes	$\leq 2$
p_hat300-1	300	33917	4	147.821	yes	$\leq 2$
p_hat300-2	300	22922	9	89.907	yes	$\leq 2$
p_hat300-3	300	11460	16	55.926	yes	$\leq 2$
p_hat500-1	500	93181	5	441.127	yes	$\leq 2$
p_hat500-2	500	61804	12	300.894	yes	$\leq 2$

Instance	$n$	$m$	$ D $	$T_F$ ms	--consistency	Bound
p_hat500-3	500	30950	26	344.563	yes	$\leq 2$
p_hat700-1	700	183651	4	906.348	yes	$\leq 2$
p_hat700-2	700	122922	9	615.148	yes	$\leq 2$
p_hat700-3	700	61640	22	320.006	yes	$\leq 2$
san200_0.7_1	200	5970	9	30.371	yes	$\leq 2$
san200_0.7_2	200	5970	8	255.688	yes	$\leq 2$
san200_0.9_1	200	1990	20	13.806	yes	$\leq 2$
san200_0.9_2	200	1990	21	10.775	yes	$\leq 2$
san200_0.9_3	200	1990	22	15.211	yes	$\leq 2$
san400_0.5_1	400	39900	6	190.746	yes	$\leq 2$
san400_0.7_1	400	23940	10	109.924	yes	$\leq 2$
san400_0.7_2	400	23940	11	106.871	yes	$\leq 2$
san400_0.7_3	400	23940	10	114.503	yes	$\leq 2$
san400_0.9_1	400	7980	30	44.115	yes	$\leq 2$
sanr200_0.7	200	6032	11	98.955	yes	$\leq 2$
sanr200_0.9	200	2037	23	9.569	yes	$\leq 2$
sanr400_0.5	400	39816	7	191.481	yes	$\leq 2$
sanr400_0.7	400	23931	13	108.915	yes	$\leq 2$
<b>Summary:</b> 68 certified runs; total algorithm time 13.534 s; median 85.799 ms; maximum 1035.418 ms.						

### 7.2. Certified Benchmark Results

The largest logged dominating set has size 172 on hamming10-2; the smallest has size 2 on the c-fat and small Hamming instances. The slowest algorithmic run is brock800\_4 at 1035.418 ms, and the total algorithmic time across all 68 instances is 13.534 s. Since every instance was run with --consistency, Table 2 reports certified approximation bounds rather than empirical ratios against an ILP optimum.

### 7.3. NetworkX Comparison on VC-Bench Real-World Graphs

We also evaluated Furones v0.2.8 on a 64-graph compendium selected from Cai's VC-Bench collection of real-world undirected graphs [8]. Cai's source collection is a large DIMACS-format graph repository; our experiment uses this selected compendium rather than the full repository. The command was

```
python -m furones.batch -i .\network\ -c -a -l --consistency,
```

equivalently batch\_asia -i .\network\ -c -a -l --consistency. The -a flag runs NetworkX's default dominating-set approximation implementation in NetworkX 3.4.2 before Furones; this is the logarithmic-ratio approximation based on the standard set-cover analysis [9]. The resulting log is network/network.txt. Table 3 reports the per-instance algorithmic times, solution sizes, logged upper-bound ratio, and certification outcome. In the current command-line implementation, this upper bound is computed as

$$\log_2(n) \cdot |D_{\text{FURONES}}| / |D_{\text{NETWORKX}}|,$$

using NetworkX's logarithmic approximation as the comparison baseline.

**Table 3.** Detailed VC-Bench comparison from network/network.txt.  $T_F$  and  $T_{NX}$  are the algorithmic times in milliseconds reported by the corresponding "done in" log lines. UB is the logged upper bound  $\log_2(n)|D_F|/|D_{NX}|$ . The final column records whether the linear-time sufficient --consistency check certified the Furones output.

Instance	$ D_F $	$T_F$ ms	$ D_{NX} $	$T_{NX}$ ms	UB	Cert.
bio-celegans	35	21.829	241	79.567	1.281	yes
bio-diseasome	97	13.082	276	85.794	3.167	yes
bio-dmela	1481	435.197	2664	12764.554	7.145	yes
bio-yeast	356	31.494	463	333.162	8.081	yes
ca-CSphd	523	28.987	554	476.615	10.269	yes

Continued on next page

**Table 3.** Detailed VC-Bench comparison from network/network.txt (continued).

Instance	$ D_F $	$T_F$ ms	$ D_{NX} $	$T_{NX}$ ms	UB	Cert.
ca-Erdos992	440	79.386	461	1363.359	11.754	yes
ca-GrQc	788	159.655	2218	4862.762	4.271	yes
ca-netscience	57	34.949	138	22.417	3.538	yes
ia-email-EU	755	670.716	820	16967.082	13.797	yes
ia-email-univ	233	116.958	604	457.616	3.914	yes
ia-enron-only	25	10.048	73	8.127	2.452	yes
ia-fb-messages	263	68.861	594	510.342	4.563	yes
ia-infect-dublin	64	17.877	241	74.469	2.305	yes
ia-infect-hyper	7	26.160	5	0.000	9.548	yes
ia-reality	81	79.358	81	286.423	12.733	yes
inf-power	1500	180.161	2263	5326.914	8.133	yes
rt-retweet	32	7.856	33	0.000	6.385	yes
rt-twitter-copen	199	0.505	236	100.225	8.071	yes
scc_enron-only	2	49.686	1	0.000	14.380	yes
scc_fb-forum	28	369.487	322	963.753	0.777	yes
scc_infect-hyper	1	46.864	1	0.000	6.820	yes
scc_retweet	156	469.269	562	1849.766	2.841	yes
scc_retweet-crawl	6079	508.851	8447	73413.115	10.123	yes
scc_rt_alwefaq	17	9.646	35	0.000	2.997	yes
scc_rt_assad	9	0.000	16	0.000	2.862	yes
scc_rt_bahrain	27	5.158	37	6.804	4.502	yes
scc_rt_barackobama	15	0.000	29	0.000	3.270	yes
scc_rt_damascus	11	3.064	15	0.842	3.731	yes
scc_rt_dash	12	2.242	15	0.534	3.963	yes
scc_rt_gmanews	12	4.640	45	6.426	1.887	yes
scc_rt_gop	6	0.000	6	0.000	3.700	yes
scc_rt_http	1	0.000	1	0.000	2.322	yes
scc_rt_israel	11	0.000	11	0.000	4.459	yes
scc_rt_justinbieber	7	0.000	26	0.000	1.603	no
scc_rt_ksa	8	0.000	12	0.000	2.928	yes
scc_rt_lebanon	5	0.000	5	0.000	3.322	yes
scc_rt_libya	10	0.000	12	0.000	3.962	yes
scc_rt_lolgop	14	25.584	104	37.097	1.089	yes
scc_rt_mittromney	36	0.000	42	0.000	5.719	yes
scc_rt_obama	4	1.006	4	0.000	3.000	yes
scc_rt_occupy	19	1.662	22	2.577	4.993	yes
scc_rt_occupywallstnyc	10	6.294	45	3.365	1.553	yes
scc_rt_oman	5	1.049	6	0.000	3.333	yes
scc_rt_onedirection	3	2.218	27	1.122	0.570	yes
scc_rt_p2	12	0.000	12	0.000	4.700	yes
scc_rt_qatif	4	0.000	5	0.000	3.046	yes
scc_rt_saudi	7	2.041	17	1.019	1.979	yes
scc_rt_tcot	11	1.112	12	1.032	4.309	yes
scc_rt_tlot	6	0.000	6	0.000	3.700	yes
scc_rt_uae	7	1.039	8	1.021	3.649	yes
scc_rt_voteonedirection	3	0.000	3	0.000	2.807	yes
soc-dolphins	14	4.693	33	1.121	2.526	yes
soc-karate	4	2.231	8	0.000	2.544	yes
soc-wiki-Vote	214	127.561	415	204.852	5.051	yes
tech-as-caida2007	2401	687.418	3691	61090.660	9.557	yes
tech-routers-rf	490	127.025	805	875.531	6.723	yes
tech-WHOIS	682	671.057	2285	13022.264	3.841	yes
web-BerkStan	3067	588.120	5472	34108.009	7.615	yes
web-edu	249	267.021	1451	2493.434	1.985	yes
web-google	205	115.451	497	1262.731	4.266	yes
web-indochina-2004	1493	1411.469	7303	128617.267	2.754	yes
web-polblogs	108	83.663	246	297.053	4.096	yes
web-spam	859	1066.798	2346	25313.327	4.474	yes
web-webbase-2001	1277	1180.829	2682	74350.114	6.652	yes
<b>Summary:</b> 64 instances; Furones total 9.827 s; NetworkX total 461.644 s; 63 certified; 1 uncertified with UB 1.603 < 2. Furones total size 24557; NetworkX total size 49110; Furones returned a no-larger set on 62 of 64 instances.						

The comparison shows a strong aggregate size advantage for Furones: its returned dominating set is no larger than NetworkX's on 62 of 64 instances and the total returned size is roughly half of NetworkX's. NetworkX has a slightly smaller median runtime because many tiny SCC instances finish in effectively zero milliseconds, but Furones is much faster in total time and on the largest expensive cases; for example, on web-indochina-2004, Furones takes 1411.469 ms versus NetworkX's 128617.267 ms while returning a set of size 1493 instead of 7303. The sole certification failure, scc\_rt\_justinbieber,

therefore appears as a limitation of the current sufficient certificate rather than as a poor empirical outcome: the logged upper-bound ratio is  $1.603 < 2$ .

## 8. Conclusions

We have presented a polynomial-time algorithm for the Minimum Dominating Set problem that combines a two-phase reduction to a planar kernel with Baker's PTAS and a redundancy-pruning step. The algorithm always returns a valid dominating set, its non-PTAS reduction runs in  $O(n + m)$  time, and it provably achieves a 2-approximation whenever the reduction is tight; the implementation can enforce this claim through a linear-time sufficient --consistency check that refuses uncertified instances. In the non-tight case our new *witness mapping* (Theorem 3) provides an injection from the post-pruning forced-boundary  $F_R^{\text{pruned}}$  into  $R \setminus F_R$ , clarifying why pruning can improve the observed result without converting the current proof into a universal 2-approximation guarantee.

Across 68 NPBench DIMACS clique-complement graphs, the --consistency run certifies the 2-approximation bound on every instance and finishes the algorithmic phase in 13.534 seconds total. On the 64-graph VC-Bench compendium, Furones is compared against NetworkX's approximation implementation; it returns no larger a set on 62 instances, is much faster in total algorithmic time, and has only one consistency failure, whose logged upper-bound ratio is nevertheless below 2.

Future work includes strengthening the certificate beyond the current full-boundary sufficient condition, searching for counterexamples to post-pruning universal bounds, extending the reduction framework to other hard problems such as set cover and vertex cover, and studying hybrid projections that recover some dense planar edges while preserving near-linear time. The implementation is freely available as the Furones package (v0.2.8) on PyPI.

## References

1. Karp, R.M. Reducibility Among Combinatorial Problems. In *50 Years of Integer Programming 1958–2008: From the Early Years to the State-of-the-Art*; Springer: Berlin, Germany, 2010; pp. 219–241. [https://doi.org/10.1007/978-3-540-68279-0\\_8](https://doi.org/10.1007/978-3-540-68279-0_8).
2. Feige, U. A threshold of  $\ln n$  for approximating set cover. *Journal of the ACM* **1998**, *45*, 634–652. <https://doi.org/10.1145/285055.285059>.
3. Raz, R.; Safra, S. The distance of the minimum dominating set problem from  $(\ln n)$ -approximation. In *Proceedings of the Proceedings 43rd Annual IEEE Symposium on Foundations of Computer Science*, Washington, DC, USA, 2002; pp. 311–320. <https://doi.org/10.1109/SFCS.2002.1181958>.
4. Baker, B.S. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM* **1994**, *41*, 153–180. <https://doi.org/10.1145/174644.174650>.
5. Vega, F. Furones: Approximate Dominating Set Solver. <https://pypi.org/project/furones>, 2026. Version 0.2.8, Accessed May 23, 2026.
6. Storer, J.A. Exact and approximation algorithms for the minimum dominating set problem on planar graphs. *SIAM Journal on Computing* **1983**, *12*, 679–696. <https://doi.org/10.1137/0212046>.
7. Nguyen, T.; Bui, T. NP-Complete Benchmark Instances. <https://roars.dev/npbench/>. Vertex cover DIMACS clique complements.
8. Cai, S. A Collection of Large Graphs for Vertex Cover Benchmarking. <https://lcs.ios.ac.cn/~caisw/graphs.html>, 2017. Accessed: 2026-05-22.
9. Vazirani, V.V. *Approximation Algorithms*; Springer Science & Business Media, 2001.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.