Domination without independence, a tale of parameterized inapproximability

Yijia Chen

Fudan University

Rodfest, January 6th, 2017

Joint work with Bingkai Lin (NII)

Domination with independence

Definition

Let G be a graph. Then a subset $D \subseteq V(G)$ is an independent dominating set if

(i) for every $u \in V(G)$ there is a vertex $v \in D$ with $u = v$ or $\{u, v\} \in E(G)$,

(ii) and for every $u, v \in D$ we have $\{u, v\} \notin E(G)$.

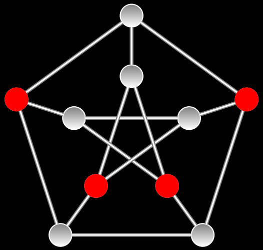

The peterson graph with an independent dominating set of size 4.

Domination without independence

Definition

Let G be a graph. Then a subset $D \subseteq V(G)$ is a dominating set, if for every $u \in V(G)$ there is a vertex $v \in D$ with $u = v$ or $\{u, v\} \in E(G)$.

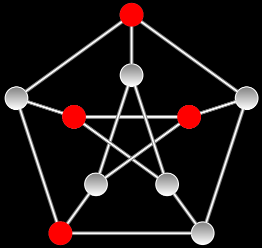

The peterson graph with a dominating set of size 4 which is not an independent set.

Compute a dominating set

- as small as possible – approximation algorithms
- and as efficiently as possible – parameterized algorithms

Parameterized Algorithms and Complexity

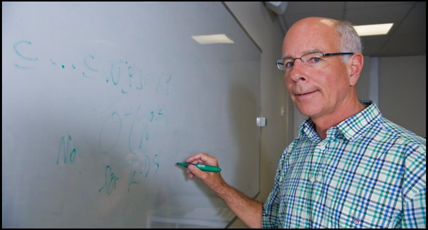

Its vastness and influence

- Five monographs, including
  1. R.G. Downey and M.R. Fellows. Parameterized Complexity, 1999.
  2. R.G. Downey and M.R. Fellows. Fundamentals of Parameterized Complexity, 2013. More than 700 pages!
- One dedicated conference – The International Symposium on Parameterized and Exact Computation (IPEC), with one dedicated Nerode Prize.
- Frequent appearances in top computer science theory conferences – FOCS, STOC, SODA, with two recent best SODA papers.
- When visiting ftp servers, I often type

  $$\texttt{fpt}://***.***.***$$

  fixed-parameter tractable – the central notion of parameterized complexity.

This talk

Using the dominating set problem as an example:

- What is fixed-parameter tractability? And what is its connection to classical complexity?
- What is parameterized approximation - the combination of classical approximation and parameterized algorithms?
- Do we have any parameterized approximation of the (independent) dominating set problem?

Rod has played major roles in all the above developments.

The dominating set problem

min-DS	
Input:	A graph G.
Solution:	A dominating set $D \subseteq V(G)$ of G.
Cost:	$\lvert D \rvert$.
Goal:	min.

The decision version:

DS	
Input:	A graph G and $k \in \mathbb{N}$.
Problem:	Does G contain a dominating set of size at most k?

min-IDS and IDS are similarly defined for the independent dominating set problem.

Theorem (Karp, 1972)
DS *is complete for* NP, *hence* min-DS *is hard for* NP.

Two approaches for dealing with NP-hard problems:

1. Approximation algorithms.
2. Parameterized algorithms.

Parameterized Complexity of DS

p-DS

Input: A graph $G = (V, E)$ and $k \in \mathbb{N}$.
Parameter: k.
Problem: Does G contain a dominating set of size at most k?

Can we decide p-DS in polynomial time provided k is sufficiently small, e.g.,

$$k \leq \log\log|V| \quad \text{or even} \quad k \leq \log\log\log\log|V|?$$

Fixed-parameter tractability

Definition

p-DS is fixed-parameter tractable if there is an algorithm deciding p-DS in time

$$f(k) \cdot |G|^{O(1)},$$

where $f : \mathbb{N} \to \mathbb{N}$ is an arbitrary computable function.

Lemma

p-DS is fixed-parameter tractable if and only if there is an algorithm computing min-DS *in time*

$$f(\mathrm{ds}(G)) \cdot |G|^{O(1)},$$

where $f : \mathbb{N} \to \mathbb{N}$ is an arbitrary computable function.

Is p-DS fixed-parameter tractable?

Theorem (Downey and Fellows, 1999)

p-DS is not fixed-parameter tractable, unless ETH [Impagliazzo, Paturi, and Zane, 2001] *fails.*

$$\begin{aligned}\text{ETH} &= \text{Exponential Time Hypothesis} \\ &= \text{ there is no algorithm deciding } 3\textsc{Sat} \text{ in time } 2^{o(n)}.\end{aligned}$$

Perhaps one of the most important paper concerning lower bounds here is the somewhat neglected paper of Impagliazzo, Paturi and Zane.

– R.G. Downey. Parameterized Complexity for the Skeptic, 2003.

The greedy approximation algorithm

1. $D \leftarrow \emptyset$
2. `while` D `is not a dominating set do`
3. $\quad v \leftarrow$ `a vertex adjacent to the largest number of un-dominated vertices`
4. $\quad D \leftarrow D \cup \{v\}$
5. `output` D`.`

The approximation ratio of the greedy algorithm

Theorem (Johnson, 1974; Stein, 1974; Lovász, 1975; etc)

On an input graph G with n vertices, the greedy algorithm always computes a dominating set of size at most

$$\ln n \cdot \mathrm{ds}(G),$$

where $\mathrm{ds}(G)$ *is the size of a minimum dominating set in* G*.*

Moreover, there are graphs showing that the above bound is almost tight.

Lower bound

Building on the PCP theorem and Raz's Parallel Repetition Theorem, after a long sequence of papers:

Theorem (Dinur and Steurer, 2014)

Assume P $\neq$ NP*. Then for every* $\varepsilon > 0$ *there is no polynomial time algorithm that always computes a dominating set of size at most*

$$(1-\varepsilon)\cdot \ln n \cdot \mathrm{ds}(G).$$

A happy ending for classical complexity?

Do we have a polynomial time approximation algorithm that always computes a dominating set of size at most

$$2^{2^{2^{2^{\mathrm{ds}(G)}}}}?$$

1. The examples for the greedy lower bound have constant-size $\mathrm{ds}(G)$. Thus, $\ln \mathrm{ds}(G) \cdot \mathrm{ds}(G)$ is not a greedy upper bound.
2. The existing proofs only rule out algorithms computing a dominating set of size bounded by
   - $\ln \mathrm{ds}(G) \cdot \mathrm{ds}(G)$ under $\mathrm{P} \neq \mathrm{NP}$.
   - $\mathrm{polylog}(\mathrm{ds}(G)) \cdot \mathrm{ds}(G)$ assuming SAT can not be solved in subexponential time [Nelson, 2007].

Combine fixed-parameter and approximation algorithms

Three papers

- L. Cai and X. Huang. Fixed-parameter approximation: Conceptual framework and approximability results. 2006.
- Y. Chen, M. Grohe, and M. Grüber. On parameterized approximability, 2006.
- R. G. Downey, M. R. Fellows, and C. McCartin. Parameterized approximation problems, 2006.

Question

Is there an algorithm $\mathbb{A}$ such that for every input graph G the following two conditions are satisfied?

(A1) *$\mathbb{A}$ outputs an (independent) dominating set of size*

$$\rho(\mathrm{ds}(G)) \cdot \mathrm{ds}(G),$$

i.e., the approximation ratio $\rho : \mathbb{N} \to \mathbb{N}$ is measured in terms of the optimum.

(A2) *The running time of $\mathbb{A}$ is bounded by*

$$f(\mathrm{ds}(G)) \cdot |G|^{O(1)},$$

for a computable function $f : \mathbb{N} \to \mathbb{N}$.

$\mathbb{A}$ *is an* fpt *approximation of* min-DS*(*min-IDS*) with ratio ρ.*

Even if you are only interested in polynomial time

Lemma

p-DS has an fpt *approximation for some ratio ρ if and only if* min-DS *has a polynomial time approximation for some ratio $\rho'(\mathrm{ds}(G))$.*

Rod's results

Theorem (Downey, Fellows, McCartin, and Rosamond, 2008)
Assume that p-DS is not fixed-parameter tractable.

p-DS has no fpt *approximation with ratio $\rho(k) = k + c$ for any constant $c \in \mathbb{N}$.*

p-IDS has no fpt *approximation for any ratio $\rho : \mathbb{N} \to \mathbb{N}$.*

p-IDS is one of the first natural problems proved to have no fpt approximation at all.

It is hideous.

– M.R. Fellows.

The difference between IDS and DS

If D is a dominating set of G, then any superset $D' \supset D$ is also a dominating set. That is, DS is a monotone problem.

IDS is not monotone.

Our results

Theorem

1. *Assume* FPT $\neq$ W[1]*. Then there is no* fpt *approximation of* min-DS *with any constant approximation ratio. That is, for every $c \in \mathbb{N}$ there is no* fpt*-algorithm $\mathbb{A}$ which always outputs a dominating set of size at most*

$$c \cdot \mathrm{ds}(G).$$

2. *Assume* ETH*. Then for every $\varepsilon > 0$ there is no* fpt *approximation of* min-DS *with approximation ratio*

$$\sqrt[4+\varepsilon]{\log\left(\mathrm{ds}(G)\right)}.$$

What is FPT $\neq$ W[1]?

Theorem

FPT $\neq$ W[1] *if and only if there is no* fpt-*algorithm that decides*

p-CLIQUE	
Input:	A graph G and $k \in \mathbb{N}$.
Parameter:	k.
Problem:	Does G contain a clique of size k?

where a clique is a subset of V with all possible edges present in G.

Theorem (Downey and Fellows, 1999)

$$\text{ETH} \implies \text{FPT} \neq \text{W[1]}$$
$$\iff p\text{-}\textsc{Clique} \textit{ is not fixed-parameter tractable.}$$

Some consequence

Corollary

Let $\beta : \mathbb{N} \to \mathbb{N}$ be a nondecreasing and unbounded computable function. Consider the following promise problem

min-DS_β	
Input:	A graph G with $\mathrm{ds}(G) \le \beta(\|V(G)\|)$.
Solution:	A dominating set D of G.
Cost:	$\|D\|$.
Goal:	min.

Then there is no polynomial time constant approximation algorithm for min-DS_β*, unless* $\mathrm{FPT} = \mathrm{W}[1]$*.*

Remark

It would be difficult to prove the above corollary using only the classical PCP*-based approach.*

The starting point of the proof

Theorem (Lin, 2015)

p-BiClique	
Input:	A graph $G = (V, E)$ and $k \in \mathbb{N}$.
Parameter:	k.
Problem:	Does G contain a biclique whose both side has size k?

is W[1]*-hard. Here, a biclique* $(C_1, C_2) \subseteq V \times V$ *in* G *satisfies that* $\{u, v\} \in E$ *for every* $u \in C_1$ *and* $v \in C_2$*.*

The most infamous ... It is rather an embarrassment to the field that the question remains open after all these years!

– R.G. Downey and M.R. Fellows. Fundamentals of Parameterized Complexity, 2013.

Lin's construction

Theorem

There is an fpt-*algorithm* $\mathbb{A}$ *such that for every graph* G *with* n *vertices and* $k \in \mathbb{N}$ *the algorithm* $\mathbb{A}$ *constructs a bipartite graph* $H = (A \dot\cup B, E)$ *satisfying:*

(i) *if* G *contains a* k*-clique, then there are* $\binom{k}{2}$ *vertices in* A *with at least* $\left\lceil n^{\frac{6}{k+1}} \right\rceil$ *common neighbors in* B*;*

(ii) *otherwise, every* $\binom{k}{2}$ *vertices in* A *have at most* $(k+1)!$ *common neighbors in* B*.*

The gap

(i) If G has no k-clique, then any biclique (C_1, C_2) in H with $|C_1| = \binom{k}{2}$ must have $|C_2| \leq (k+1)!$.

(ii) If G has a k-clique, then there is a biclique (C_1, C_2) in H with $|C_1| = \binom{k}{2}$ and

$$|C_2| \geq n^{\frac{6}{k+1}} \geq h(k)$$

for any $h : \mathbb{N} \to \mathbb{N}$ if n is sufficiently large.

We might choose $h(k) = (k+1)! + 1$, $h(k) = k^k$, or $h(k) = 2^{2^{2^{2^k}}}$.

The difficulty

In the beginning we tried to use the gap to prove p-$\textsc{Clique}$ is hard to approximate, but

(iii) even if G has no k-clique, H might still have a biclique (C_1, C_2) in H with $|C_1| = \binom{k}{2} - 1$ and

$$|C_2| \geq n^{\frac{6}{k+1}} \geq h(k)$$

for any $h : \mathbb{N} \to \mathbb{N}$ if n is sufficiently large.

The Color-Coding

Lemma (Alon, Yuster, and Zwick, 1995)

For every $n, k \in \mathbb{N}$ there is a family $\Lambda_{n,k}$ of polynomial time computable functions from $[n]$ to $[k]$ such that for every k-element subset X of $[n]$, there is an $h \in \Lambda_{n,k}$ such that h is injective on X. Moreover, $\Lambda_{n,k}$ can be computed in time $2^{O(k)} \cdot n^{O(1)}$.

We can use the color-coding technique to encode in a DS instance the condition that $|C_1| = \binom{k}{2}$.

The case $\rho < 3/2$

Theorem

Let $\rho < 3/2$. Then there is no fpt *approximation of the parameterized dominating set problem achieving ratio ρ unless* $\mathrm{FPT} = \mathrm{W}[1]$.

Lemma

Let $\rho < 3/2$. Then there is an fpt-*algorithm $\mathbb{A}$ such that for every graph G and $k \in \mathbb{N}$ the algorithm $\mathbb{A}$ outputs a graph H satisfying the following conditions.*

(i) *If G has a k-clique, then $\mathrm{ds}(H) \leq h(k)$.*

(ii) *If G has no k-clique, then $\mathrm{ds}(H) > \rho \cdot h(k)$.*

Here, $h : \mathbb{N} \to \mathbb{N}$ is an appropriate computable function.

Amplifying the gap from 3/2?

The standard approach would be to define a certain graph product, e.g., G^2, such that

$$\mathrm{ds}(G^2) = \mathrm{ds}(G)^2 \tag{1}$$

for every graph G.

Lemma

Assume $\mathsf{P} \neq \mathsf{NP}$*. Then there is no polynomial time algorithm which computes a graph* G^2 *satisfying* (1)*.*

Amplifying the gap from 3/2

We have to define a tailor-made graph product on the graphs produced by the reduction in the case of $\rho < 3/2$.

Theorem

There is an algorithm $\mathbb{A}$ *such that on input a graph* G, $k \geq 3$, *and* $c \in \mathbb{N}$ *the algorithm* $\mathbb{A}$ *computes a graph* G_c *such that*

(i) *if* G *contains a* k*-clique, then* $\mathrm{ds}(G_c) < 1.1 \cdot d^c$;

(ii) *if* G *does not contain a* k*-clique, then* $\mathrm{ds}(G_c) > c \cdot d^c/3$,

where $d = \left(30 \cdot c^2 \cdot (k+1)^2\right)^{4 \cdot k^3 + 3c}$. *Moreover the running time of* $\mathbb{A}$ *is bounded by* $f(k, c) \cdot |G|^{O(c)}$ *for a computable function* $f : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$.

Main results (again)

Theorem

1. *Assume* FPT $\neq$ W[1]. *Then there is no* fpt *approximation of* min-DS *with any constant approximation ratio. That is, for every* $c \in \mathbb{N}$ *there is no* fpt-*algorithm* $\mathbb{A}$ *which always outputs a dominating set of size at most*
$$c \cdot \mathrm{ds}(G).$$

2. *Assume* ETH. *Then for every* $\varepsilon > 0$ *there is no* fpt *approximation of* min-DS *with ratio*
$$\sqrt[4+\varepsilon]{\log\big(\mathrm{ds}(G)\big)}.$$

An AC^0 version

Building on a strong AC^0 version [Chen and Flum, 2016] of the the planted clique conjecture:

Theorem

There are no AC^0*-circuits which approximate* min-DS *with ratio*

$$\frac{\log \mathrm{ds}(G)}{\omega(\log \log \mathrm{ds}(G))}.$$

AC^0

A sequence $(C_n)_{n \in \mathbb{N}}$ of Boolean circuits are AC^0-circuits if for every $n \in \mathbb{N}$

(i) C_n has n inputs;

(ii) the depth of C_n is bounded by a fixed constant d;

(iii) the size of C_n is polynomially bounded in n.

$(C_n)_{n \in \mathbb{N}}$ is dlogtime-uniform if

$$1^n \mapsto C_n$$

can be computed by a deterministic logarithmic time Turing machine.

AC^0 and first-order logic

We assume every graph G has $V(G) = \{0, \ldots, n-1\}$, thus we can extend G to $(G, +, \times)$.

Theorem (Barrington, Immerman, and Straubing, 1990)

Let K be a class of graphs. Then the following are equivalent:

- *There is a family of dlogtime-uniform* AC^0*-circuits* $\left(\mathrm{C}_n\right)_{n \in \mathbb{N}}$ *such that*

  $$G \in K \iff \mathrm{C}(G) = 1$$

  for every graph G.

- *There is an* FO*-sentence φ such that*

  $$G \in K \iff (G, +, \times) \models \varphi$$

  for every graph G.

Inapproximability of min-DS by FO

Corollary

Let $\iota : \mathbb{N} \to \mathbb{R}^+$ *be a function with* $\iota(k) = (\log k)/\omega(\log\log k)$*. Then there exist two functions* $\delta_0, \delta_1 : \mathbb{N} \to \mathbb{N}$ *with*

$$\delta_1(n) \geq \iota\big(\delta_0(n)\big) \cdot \delta_0(n)$$

for all $n \in \mathbb{N}$ *such that the following promise problem is not definable in* FO*.*

$\mathrm{GAP}_{\delta_0,\delta_1}$-DS

Input:	A graph G with n vertices.
Parameter:	$\mathrm{ds}(G) \leq \delta_0(n)$ or $\mathrm{ds}(G) > \delta_1(n)$.
Problem:	Is $\mathrm{ds}(G) > \delta_1(n)$?

Approximability of the vertex cover problem by FO

Theorem

There is a computable function $\rho : \mathbb{N} \to \mathbb{N}$ *and an* FO*-formula* $\varphi(x)$ *such that for every graph* G *the set*

$$C = \{v \in V(G) \mid (G, +, \times) \models \varphi(v)\}$$

is a vertex cover of G *with*

$$|C| \leq \rho(\mathrm{vc}(G)) \cdot \mathrm{vc}(G),$$

where $\mathrm{vc}(G)$ *is the size of a minimum vertex cover of* G*.*

Conclusions

1. Parameterized approximation is still a very young subarea of parameterized algorithms and complexity, with many deep and challenging open problems, e.g., the clique problem.

2. The existing results already show some great potentials, e.g., under ETH we might bypass the daunting PCP machinery.

   - Our lower bound holds for AC^0. On the other hand, it is known that PCP cannot be done in AC^0.
   - Can we reverse-engineer a parameterized PCP theorem?

An email exchange with Rod

Dear Rod,

I enclose a paper I just finished with Bingkai Lin proving that the dominating set problem has no fpt approximation with any constant ratio.

Best wishes,

Yijia

cool!

rod

Happy 60, Rod!