A Finite Compactness Notion, and Property Testing

(extended/open abstract for LLC 2011)

Richard J. Lipton   Kenneth W. Regan

June 25, 2011

Compactness in Logic

Say a property Π of $\mathcal{L}$-structures is *compact* if for all $\mathcal{L}$-structures Ω, whenever all finite substructures of Ω have Π, then Ω has Π.

- Usually one says that a collection $\mathcal{P}$ of properties is compact, where $\mathcal{P}$ is defined by expressibility in $\mathcal{L}$. For instance, first-order definable properties are compact.
- Here we wish to focus on single properties.
- We consider mostly hereditary properties, so "whenever" can mean *iff* not just *if*.
- We focus on simple (un)directed graphs as the structures. Example: *Bipartiteness*.

Complexity-Scaled Version

Let $N = 2^n$, where n is the input-size parameter. Consider graphs G of size N nodes that offer query access to (non-)edges $E(i, j)$.

- Size close to N, e.g. $(1 - \epsilon)N$ = "large" ≈ infinite.
- Size $n^{O(1)}$ = "small' ≈ finite.

No longer true for bipartiteness—the smallest odd cycles in a non-bipartite graph of size N can have size $\approx N$.

How to have a compactness notion for already-finite structures?

Solution: make the target notion approximate.

Approximate Compactness

Idea: Π is *approximately compact* if for all large G, whenever all small subgraphs have Π, then some large subgraph has Π.

Definition

A graph property Π is $(f(N), g(N))$-*compact* if for all N and all graphs G of size N, if all $f(N)$-node subgraphs H of G have Π, then there is a $g(N)$-node subgraph G' that has Π.

Subgraphs are *vertex*-induced; original thought was an *edge*-induced concept.

Example: Bipartiteness

If all k-node subgraphs are bipartite, does it follow that some $(1-\epsilon)N$-node subgraph is bipartite?

Depends on k. Here k may depend on N and ϵ, but ϵ is fixed for all N, so $k = f(N)$. With $g(N)$ fixed as $(1-\epsilon)N$, only k varies.

The following blog-procured results by Noga Alon and Luca Trevisan act as asymptotic bookends on k.

Contributed Results

Theorem (Alon)

For every $\epsilon > 0$ *there exists* $C > 0$ *depending only on* ϵ *such that bipartiteness is* $(C \log N, (1-\epsilon)N)$*-compact. Moreover, we can construct the subgraph of size* $(1-\epsilon)N$ *in* $N^{O(1)}$ *time.*

Theorem (Trevisan)

If $f(N)$ *is such that bipartiteness is* $(f(N), (1-\epsilon)N)$*-compact, then* $f(N) = \Omega(\log N)$*.*

Proof of Theorem 3

Proof.

There exist size-N expanders G of girth $k = \Omega(\log N)$, so certainly any odd cycle has that size, thus all subgraphs of size $k-1$ are bipartite. By the *expander mixing lemma* G cannot have even a bipartite subgraph of size $> N/2$, for large enough N. □

Proof of Theorem 2

Fix $\epsilon > 0$; fix $C > 0$ later, and let $k = \frac{1}{2} C \log_2 N$. A graph is non-bipartite if and only if it has an odd-length cycle. Hence if all k-node subgraphs of G are bipartite, then G has no odd cycles of length k or less.

To construct a bipartite subgraph H, start with any vertex v. For each $i \geq 1$ define $N_i(v)$ to be the neighborhood of vertices within i steps of v, and $S_i = N_i \setminus N_{i-1}$ to be the "shell" of those at exactly distance i. Now consider the least j such that

$$|S_j| < \epsilon |N_j|.$$

Proof, how to avoid $j > k$

If $j > k$, then we have for all $i \leq k$, $|S_i| \geq \epsilon |N_i|$, so $|S_i| \geq \epsilon(|S_i| + |N_{i-1}|)$, so $|S_i| \geq (1/(1-\epsilon))|N_{i-1}|$, which in turn trivially implies $|S_i| \geq (1/(1-\epsilon))|N_{i-1}|$. This implies

$$|S_i| \geq (\frac{1}{1-\epsilon})^k,$$

which is $> N$ when $\frac{C}{2} \log_2 N \log_2(\frac{1}{1-\epsilon}) > \log_2 N$, so when

$$C > \frac{2}{\log_2(\frac{1}{1-\epsilon})}.$$

Fixing C to be otherwise forces $j \leq k$.

Proof, case $j \leq k$

Now we observe that every S_i for $i < j$ is an independent set. If it has an edge (s_i, t_i), then the paths of length i from s_i and t_i back to v come together at v or some earlier node in a way that forms an odd cycle of length at most $2i + 1 \leq C \log_2 n$, contradicting the assumption.

It follows that the subgraph H_j induced by $N_j \setminus S_j$ is bipartite, since the S_i give the 2-coloring. Putting $n = |N_j|$, note that $|S_j| \leq \epsilon n$.

By induction we have that the leftover graph induced by $V(G) \setminus N_j$ has a bipartite subgraph H' of size at least $(1 - \epsilon)(N - n)$. Since $H_j \cup H'$ is separated, it is bipartite as well, and has size at least $(1 - \epsilon)N$. Clearly this induction yields a polynomial-time algorithm. $\square$

Property Testing

A *tester* is a randomized algorithm A that probes edges of the graph, such that:

1. If G has the property Π, then after the probing the algorithm A says 'yes' with probability at least $2/3$. In the *one-sided* model, the algorithm always says 'yes' in this case.
2. If the graph does not have Π, and is not "near" a graph in Π, then A returns a 'no' with probability at least $2/3$.

Here "near" is defined via a parameter $\epsilon > 0$ that is also given to the algorithm. Two graphs G, G' are ϵ-*close* if they have the same size N and for all but $\epsilon\binom{N}{2}$ pairs $1 \leq i < j \leq N$, $E(i,j) \leftrightarrow E'(i,j)$. **Relates to edge-induced subgraphs.**

Query complexity q of A is the maximum number of edge probes. **Poly-testable** if $q = q(\epsilon, N) = O(n^c)$ where the "O" depends on ϵ.

Poly-Testable $\implies$ Finitely Compact? (No.)

Idea: If all k-node subgraphs have Π then the tester A should accept, which means G is ϵ-close to some G' that has property Π. From G' we aim to produce a large vertex-induced subgraph H with Π, showing that Π is $(k, (1-\epsilon)N)$-compact.

Counterexample: Bipartiteness is testable for fixed k, but not finite-compact with fixed k. What could have gone wrong?

- The tester A need not work by probing small subgraphs for the property Π itself.
- The closeness condition counts edges, and might not carry over to *vertex*-induced subgraphs.

Cognizance and Edge-Induced Subgraphs

We address the former issue by strengthening a notion called "canonical" by Goldreich and Trevisan [2003].

Definition

A property testing algorithm A for Π is **cognizant** if it generates one or more vertex-induced subgraphs H of G, probing only the edges in H, and accepts if and only if the majority of the probed graphs have property Π.

The following theorem is credited to Alon in "Appendix D" of that paper.

Theorem (Alon in Goldreich-Trevisan, 2003)

Every testable property that is closed under edge-induced subgraphs has a cognizant tester.

“Edge-Induced” Finite Compactness

Definition

Π is $(f(N), \epsilon)$-*edge-compact* if for all N and all graphs G of size N, if all $f(N)$-node subgraphs H of G have Π, then by changing at most ϵN^2 edges we can get a graph G' that has Π.

An N-vertex graph is *dense* if it has δN^2 edges, where we intend $\delta > \epsilon$ and fixed.

Theorem

For sufficiently large k, bipartiteness of dense graphs is $(k, 1/k^3)$ edge-compact (with edge-removals only).

Proof of Theorem 7

Given $\epsilon > 0$, the known cognizant tester for bipartiteness [Goldreich survey, 2010] selects $k = O(\log(1/\epsilon)/\epsilon^2)$ vertices uniformly at random, and accepts iff the subgraph R they induce is bipartite. This makes $1/k^3 < \epsilon$. Suppose G is a graph for which all k-node subgraphs are bipartite. Then the tester accepts with certainty. By definition of being a tester, G is near a graph G' in Π, which here entails that ϵN^2 edges can be deleted from G to yield H. $\square$

- Moreover, a suitable H can be described succinctly in terms of choices that accompany R.
- Stronger bounds than Theorem 2, since $k = O(1)$, but for weaker notion—H is not a vertex-induced subgraph, and Theorem 3 shows it cannot be made so.

Some Other Properties

Theorem

If 3-colorability is $(f(N), \Theta(N))$-compact, any f, with a (random) polynomial-time algorithm for finding a $\Theta(N)$-sized subgraph and 3-coloring it, then 3-colorable graphs can be colored with $O(\log N)$ colors in (random) polynomial time.

Proof.

Removing the 3-colored subgraph always shrinks the graph by a constant factor, and since we can use fresh colors for the rest, the iteration uses $O(\log N)$ colors overall. □

Relation to Approximate Coloring

- Chlamtac [2007] colors any given 3-colorable graph G in $O(N^{0.2072})$ colors.
- Meanwhile, Guruswami and Khanna [2004] showed that it is NP-hard to find a 4-coloring.
- Still best known upper and lower bounds on the number of colors?
- Zuckerman [2007] showed that for all $\epsilon > 0$, approximating the chromatic number of a graph to wihin a factor of $N^{1-\epsilon}$ is NP-hard.
- This seems to be reason to suspect that $O(N^c)$ for some fixed $c < 0.2072$ should be a lower bound, but the consequence doesn't immediately apply to 3-colorable graphs.

Local-Global / Almost-Global

The following are considered "Local-Global" *pairs of* properties—see paper for references and links to more:

(a) If $(\forall u \neq v \in V)(\exists! w)[E(u,w) \wedge E(v,w)]$, then $(\exists u)(\forall v \neq u)E(u,v)$.

(b) If every k-node subgraph is bipartite, then G can be colored with $N^{O(1/k)}$ colors.

(c) If every k-node subgraph is 3-colorable, then G can be colored with $n^{1/2+r(k)}$ colors, where $r(k) \longrightarrow 0$ as k grows.

(d) If $h : V \longrightarrow \mathbf{R}^+$ has average value at least μ on $N_t(v)$ for all $v \in V$ and $t \leq r$, $t \geq 1$, then its average on V is at least $\mu / n^{O(1/\log r)}$.

(e) If G is s-connected and has no independent sets of size $s+1$, then G has a Hamiltonian circuit.

Does relaxing to *almost*-global enable more properties, keeping the same property?

Things To Do

1. Prove more finite-compact properties, for interesting bounds $f(N)$ for "small," $g(N)$ for "large."
2. Find a better relationship to property testing?
3. Find relations to classes of logical formulas defining the properties.

On the last, if Π is defined by a first-order sentence

$$\phi = (\forall x_1, \ldots, x_k)(\exists \ldots)M$$

(without constants, and with M quantifier-free), and ϕ holds for all k-node subgraphs, then it is true for the whole graph. However, it seems hard to say more than this in short order, or to make direct use of the weaker goal of needing ϕ to be true only of a large subgraph, in relation to the formula's structure.

Conclusions

- "With a Little Help From Our Friends," we have shown some non-trivial results and differences for a fairly natural poly-versus-exp finitary analogue of compactness.
- Motivated by, and perhaps can inform, the important field of Property Testing.
- "Open Paper" connected to our blog—anyone can pitch in.
- Thanks to the organizers for giving us this opportunity.