

SYNOPSIS OF PROOF

The proof can be described in the following steps.

1. We show that for a system of n interacting random variables that follow a *directed* graphical model, we can always obtain a parametrization of the distribution in the form of Gibbs potentials due to the recursive factorization that directed models admit.
2. We use the fact that FO(LFP) captures polynomial time on successor structures. The crucial observation is that the flow of information in FO(LFP) is *directed*, since the operator that is being iterated is a monotone operator. Using this, we can embed this flow in a larger graphical model that is directed and has $\text{poly}(n)$ variables.
3. We encode k -SAT (for $k \geq 9$) as a FO(LFP) query on a successor structure, and run this query multiple times with exponentially many initial partial assignments. In each run, the final assignment is computed using LFP (of some fixed arity).
4. We show therefore that the distribution of solutions would be represented as a mixture of directed graphical models each representing a LFP computation.
5. Each model in this mixture is directed (due to the monotonicity of the LFP operator) and therefore parametrizable with Gibbs potentials.
6. We show that Gibbs potentials in exponentially numerous mixtures of such models “patch together” through a sheaf-like property in an efficient manner. This is because each component in the mixture represents a computation by the *same monotone operator which comes from a positive first order formula* that is being iterated. This allows us to patch together different runs of the computation. In this way, we obtain a characterization of the

number of independent parameters it takes to specify the mixture¹. We show that the joint distribution over all mixtures is parameterizable with only $2^{\text{poly}(\log n)}$ independent parameters.

7. The above should not be surprising given that first order logic uses limited amount of information at each stage of the computation, and so the number of distinct independent joint behaviors possible is limited.
8. Finally, to demonstrate the contradiction, we show that in the known hard phase of random k -SAT (for $k \geq 9$), the number of independent parameters required to specify the joint distribution is $O(c^n)$, $c > 1$. In order to see this, we prove that each set of values the core takes in the exponentially many clusters must be specified as an independent parameter of the joint distribution. This completes the proof that $\mathbf{P} \neq \mathbf{NP}$.
9. We show that the above does not hold for linear systems such as XORSAT that also have exponentially many clusters since the linearity limits the number of independent parameters that can be used to specify the potentials. Therefore, in such systems, each cluster does not require an additional independent parameter to specify. The case of 2-SAT is also dissimilar since it does not enter a d1RSB phase (such a phase being proven to exist only for $k \geq 9$, and strong empirical evidence that it does not exist even for $k = 3$). In both the cases of XORSAT and 2-SAT, we have a parametrization of the joint distribution in terms of potentials with $O(c^{\text{poly}(\log n)})$ independent parameters, which corresponds to their being in \mathbf{P} . In other words, it is not only the geometry of the solution space, but

¹In earlier versions of the paper, we used locality. We later realized that locality is not the fundamental property of first order logic that we require. What we require is that the first order formula bases its decision at each stage only on the presence of fixed formulae in a fixed number of element types, as described by the Gaifman and Schwentick-Barthelmann Normal forms. Both monadic and complex LFP have this property at each stage, since it characterizes all first order formulae. It simplifies our proof and resolves the issues raised about the finite model theory portion of the proof.

the effect that it has on the number of independent parameters required to specify the distribution, that has to be taken into account.

10. In summary, the complexity of the joint distribution generated by LFP would be insufficient to explain the complexity of the joint distribution of variables in the hard phases of k -SAT for $k \geq 9$. In other words, the LFP interaction model is not strong enough to capture the complex interactions between variables that exist in the d1RSB phase of k -SAT for $k \geq 9$. Its potentials either have factors that are too small (the case of range limited interactions, such as monadic LFP), or when they do have factors of scope $O(n)$, these factors are parametrized over a small number of independent behaviors (the case of value limited interactions, such as complex LFP). In both cases, $2^{\text{poly}(\log n)}$ independent parameters suffice to specify the potentials for the distribution.