Publications of Andreas Göbel

This document lists all peer-reviewed publications of Andreas Göbel, Chair for Algorithm Engineering, Hasso Plattner Institute, Potsdam, Germany. This listing was automatically generated on December 7, 2021. An up-to-date version is available online at hpi.de/friedrich/docs/publist/goebel.pdf.

Journal articles

[1] Bampas, E., Göbel, A.-N., Pagourtzis, A., Tentes, A., On the connection between interval size functions and path counting. In: *Computational Complexity* 26, pp. 421–467, 2017.

We investigate the complexity of hard (#P-complete) counting problems that have easy decision version. By 'easy decision,' we mean that deciding whether the result of counting is nonzero is in P. This property is shared by several well-known problems, such as counting the number of perfect matchings in a given graph or counting the number of satisfying assignments of a given DNF formula. We focus on classes of such hard-to-count easy-to-decide problems which emerged through two seemingly disparate approaches: one taken by Hemaspaandra et al. (SIAM J Comput 36(5):1264–1300, 2007), who defined classes of functions that count the size of intervals of ordered strings, and one followed by Kiayias et al. (Lect Notes Comput Sci 2563:453–463, 2001), who defined the classTotP, consisting of functions that count the total number of paths of NP computations. We provide inclusion and separation relations between TotP and interval size counting classes, by means of new classes that we define in this work. Our results imply that many known #P-complete problems with easy decision are contained in the classes defined by Hemaspaandra et al., but are unlikely to be complete for these classes under reductions under which these classes are downward closed, e.g., parsimonious reductions. This, applied to the #MONSAT problem, partially answers an open question of Hemaspaandra et al. We also define a new class of interval size functions which strictly contains FP and is strictly contained in TotP under reasonable complexity-theoretic assumptions. We show that this new class contains hard counting problems.

[2] Galanis, A., Göbel, A., Goldberg, L. A., Lapinskas, J., Richerby, D., Amplifiers for the Moran Process. In: *Journal of the ACM* 64, pp. 5:1–5:90, 2017.

The Moran process, as studied by Lieberman, Hauert, and Nowak, is a randomised algorithm modelling the spread of genetic mutations in populations. The algorithm runs on an underlying graph where individuals correspond to vertices. Initially, one vertex (chosen uniformly at random) possesses a mutation, with fitness $r > 1$. All other individuals have fitness 1. During each step of the algorithm, an individual is chosen with probability proportional to its fitness, and its state (mutant or nonmutant) is passed on to an out-neighbour which is chosen uniformly at random. If the underlying graph is strongly connected, then the algorithm will eventually reach fixation, in which all individuals are mutants, or extinction, in which no individuals are mutants. An infinite family of directed graphs is said to be strongly amplifying if, for every $r > 1$, the extinction probability tends to 0 as the number of vertices increases. A formal definition is provided in the article. Strong amplification is a rather surprising property-it means that in such graphs, the fixation probability of a uniformly placed initial mutant tends to 1 even though the initial mutant only has a fixed selective advantage of $r > 1$ (independently of n). The name "strongly amplifying" comes from the fact that this selective advantage is "amplified." Strong amplifiers have received quite a bit of attention, and Lieberman et al. proposed two potentially strongly amplifying families-superstars and metafunnels. Heuristic arguments have been published, arguing that there are infinite families of superstars that are strongly amplifying. The same has been claimed for metafunnels. In this article, we give the first rigorous proof that there is an infinite family of directed graphs that is strongly amplifying. We call the graphs in the family "megastars." When the algorithm is run on an n-vertex graph in this family, starting with a uniformly chosen mutant, the extinction probability is roughly $n - 1/2$ (up to logarithmic factors). We prove that all infinite families of superstars and metafunnels have larger extinction probabilities (as a function of n). Finally, we prove that our analysis of megastars is fairly tight - there is no infinite family of megastars such that the Moran algorithm gives a smaller extinction probability (up to logarithmic factors). Also, we provide a counterexample which clarifies the literature concerning the isothermal theorem of Lieberman et al.

[3] Göbel, A., Goldberg, L. A., Richerby, D., Counting Homomorphisms to Square-Free Graphs, Modulo 2. In: *Transactions on Computation Theory* 8, pp. 12:1–12:29, 2016.

We study the problem ⊕HomsToH of counting, modulo 2, the homomorphisms from an input graph to a fixed undirected graph H. A characteristic feature of modular counting is that cancellations make wider classes of instances tractable than is the case for exact (nonmodular) counting; thus, subtle dichotomy theorems can arise. We show the following dichotomy: for any H that contains no 4-cycles, ⊕HomsToH is either in polynomial time or is ⊕P-complete. This partially confirms a conjecture of Faben and Jerrum that was previously only known to hold for trees and for a restricted class of tree-width-2 graphs called cactus graphs. We confirm the conjecture for a rich class of graphs, including graphs of unbounded tree-width. In particular, we focus on square-free graphs, which are graphs without 4-cycles. These graphs arise frequently in combinatorics, for example, in connection with the strong perfect graph theorem and in certain graph algorithms. Previous dichotomy theorems required the graph to be tree-like so that tree-like decompositions could be exploited in the proof. We prove the conjecture for a much richer class of graphs by adopting a much more general approach.

[4] Göbel, A., Goldberg, L. A., McQuillan, C., Richerby, D., Yamakami, T., Counting List Matrix Partitions of Graphs. In: *SIAM Journal on Computing* 44, pp. 1089–1118, 2015.

Given a symmetric $D \times D$ matrix M over $0, 1, *$, a list M-partition of a graph G is a partition of the vertices of G into D parts which are associated with the rows of M. The part of each vertex is chosen from a given list in such a way that no edge of G is mapped to a 0 in M and no nonedge of G is mapped to a 1 in M. Many important graph-theoretic structures can be represented as list M-partitions including graph colorings, split graphs, and homogeneous sets and pairs, which arise in the proofs of the weak and strong perfect graph conjectures. Thus, there has been quite a bit of work on determining for which matrices M computations involving list M-partitions are tractable. This paper focuses on the problem of counting list M-partitions, given a graph G and given a list for each vertex of G. We identify a certain set of "tractable" matrices M. We give an algorithm that counts list M-partitions in polynomial time for every (fixed) matrix M in this set. The algorithm relies on data structures such as sparse-dense partitions and subcube decompositions to reduce each problem instance to a sequence of problem instances in which the lists have a certain useful structure that restricts access to portions of M in which the interactions of 0s and 1s are controlled. We show how to solve the resulting restricted instances by converting them into particular counting constraint satisfaction problems ($\#CSP$s), which we show how to solve using a constraint satisfaction technique known as arc-consistency. For every matrix M for which our algorithm fails, we show that the problem of counting list M-partitions is $\#P$-complete. Furthermore, we give an explicit characterization of the dichotomy theorem: counting list M-partitions is tractable (in FP) if the matrix M has a structure called a derectangularizing sequence. If M has no derectangularizing sequence, we show that counting list M-partitions is $\#P$-hard. We show that the metaproblem of determining whether a given matrix has a derectangularizing sequence is NP-complete. Finally, we show that list M-partitions can be used to encode cardinality restrictions in M-partitions problems, and we use this to give a polynomial-time algorithm for counting homogeneous pairs in graphs.

[5] Göbel, A., Goldberg, L. A., Richerby, D., The complexity of counting homomorphisms to cactus graphs modulo 2. In: *Transactions on Computation Theory* 6, pp. 17:1–17:29, 2014.

A homomorphism from a graph G to a graph H is a function from $V(G)$ to $V(H)$ that preserves edges. Many combinatorial structures that arise in mathematics and in computer science can be represented naturally as graph homomorphisms and as weighted sums of graph homomorphisms. In this article, we study the complexity of counting homomorphisms modulo 2. The complexity of modular counting was introduced by Papadimitriou and Zachos and it has been pioneered by Valiant who famously introduced a problem for which counting modulo 7 is easy but counting modulo 2 is intractable. Modular counting provides a rich setting in which to study the structure of homomorphism problems. In this case, the structure of the graph H has a big influence on the complexity of the problem. Thus, our approach is graph-theoretic. We give a complete solution for the class of cactus graphs, which are connected graphs in which every edge belongs to at most one cycle. Cactus graphs arise in many applications such as the modelling of wireless sensor networks and the comparison of genomes. We show that, for some cactus graphs H, counting homomorphisms to H modulo 2 can be done in polynomial time. For every other fixed cactus graph H, the problem is complete in the complexity class $\oplus P$, which is a wide complexity class to which every problem in the polynomial hierarchy can be reduced (using randomised reductions). Determining which H lead to tractable problems can be done in polynomial time. Our result builds upon the work of Faben and Jerrum, who gave a dichotomy for the case in which H is a tree.

Conference papers

[6] Quinzan, F., Doskoč, V., Göbel, A., Friedrich, T., Adaptive Sampling for Fast Constrained Maximization of Submodular Functions. In: *Artificial Intelligence and Statistics (AISTATS)*, pp. 964–972, 2021.

Several large-scale machine learning tasks, such as data summarization, can be approached by maximizing functions that satisfy submodularity. These optimization problems often involve complex side constraints, imposed by the underlying application. In this paper, we develop an algorithm with poly-logarithmic adaptivity for non-monotone submodular maximization under general side constraints. The adaptive complexity of a problem is the minimal number of sequential rounds required to achieve the objective. Our algorithm is suitable to maximize a non-monotone submodular function under a p-system side constraint, and it achieves a $(p+O(\sqrt{p}))$-approximation for this problem, after only poly-logarithmic adaptive rounds and polynomial queries to the valuation oracle function. Furthermore, our algorithm achieves a $p+O(1))$-approximation when the given side constraint is a p-extendible system. This algorithm yields an exponential speed-up, with respect to the adaptivity, over any other known constant-factor approximation algorithm for this problem. It also competes with previous known results in terms of the query complexity. We perform various experiments on various real-world applications. We find that, in comparison with commonly used heuristics, our algorithm performs better on these instances.

[7] Lagodzinski, J. A. G., Göbel, A., Casel, K., Friedrich, T., On Counting (Quantum-)Graph Homomorphisms in Finite Fields of Prime Order. In: *International Colloquium on Automata, Languages and Programming (ICALP)*. 48th International Colloquium on Automata, Languages, and Programming (ICALP 2021), pp. 91:1–91:15, 2021.

We study the problem of counting the number of homomorphisms from an input graph G to a fixed (quantum) graph $\bar{H}$ in any finite field of prime order $\mathbb{Z}_p$. The subproblem with graph H was introduced by Faben and Jerrum [ToC'15] and its complexity is still uncharacterised despite active research, e.g. the very recent work of Focke, Goldberg, Roth, and Zivný [SODA'21]. Our contribution is threefold. First, we introduce the study of quantum graphs to the study of modular counting homomorphisms. We show that the complexity for a quantum graph $\bar{H}$ collapses to the complexity criteria found at dimension 1: graphs. Second, in order to prove cases of intractability we establish a further reduction to the study of bipartite graphs. Lastly, we establish a dichotomy for all bipartite $(K_{3,3}\backslash\{e\}$, *domino*)-free graphs by a thorough structural study incorporating both local and global arguments. This result subsumes all results on bipartite graphs known for all prime moduli and extends them significantly. Even for the subproblem with $p = 2$ this establishes new results.

[8] Friedrich, T., Göbel, A., Krejca, M., Pappik, M., A spectral independence view on hard spheres via block dynamics. In: *International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 66:1–66:15, 2021.

The hard-sphere model is one of the most extensively studied models in statistical physics. It describes the continuous distribution of spherical particles, governed by hard-core interactions. An important quantity of this model is the normalizing factor of this distribution, called the partition function. We propose a Markov chain Monte Carlo algorithm for approximating the grand-canonical partition function of the hard-sphere model in d dimensions. Up to a fugacity of $\lambda < \mathrm{e}/2^d$, the runtime of our algorithm is polynomial in the volume of the system. This covers the entire known real-valued regime for the uniqueness of the Gibbs measure. Key to our approach is to define a discretization that closely approximates the partition function of the continuous model. This results in a discrete hard-core instance that is exponential in the size of the initial hard-sphere model. Our approximation bound follows directly from the correlation decay threshold of an infinite regular tree with degree equal to the maximum degree of our discretization. To cope with the exponential blow-up of the discrete instance we use clique dynamics, a Markov chain that was recently introduced in the setting of abstract polymer models. We prove rapid mixing of clique dynamics up to the tree threshold of the univariate hard-core model. This is achieved by relating clique dynamics to block dynamics and adapting the spectral expansion method, which was recently used to bound the mixing time of Glauber dynamics within the same parameter regime.

[9] Bläsius, T., Friedrich, T., Göbel, A., Levy, J., Rothenberger, R., The Impact of Heterogeneity and Geometry on the Proof Complexity of Random Satisfiability. In: *Symposium on Discrete Algorithms (SODA)* (Symposium of discrete algorithms), pp. 42–53, 2021.

Satisfiability is considered the canonical NP-complete problem and is used as a starting point for hardness reductions in theory, while in practice heuristic SAT solving algorithms can solve large-scale industrial SAT instances very efficiently. This disparity between theory and practice is believed to be a result of inherent properties of industrial SAT instances that make them tractable. Two characteristic properties seem to be prevalent in the majority of real-world SAT instances, heterogeneous degree distribution and locality. To understand the impact of these two properties on SAT, we study the proof complexity of random k-SAT models that allow to control heterogeneity and locality. Our findings show that heterogeneity alone does not make SAT easy as heterogeneous random k-SAT instances have superpolynomial resolution size. This implies intractability of these instances for modern SAT-solvers. On the other hand, modeling locality with an underlying geometry leads to small unsatisfiable subformulas, which can be found within polynomial time. A key ingredient for the result on geometric random k-SAT can be found in the complexity of higher-order Voronoi diagrams. As an additional technical contribution, we show an upper bound on the number of non-empty Voronoi regions, that holds for points with random positions in a very general setting. In particular, it covers arbitrary p-norms, higher dimensions, and weights affecting the area of influence of each point multiplicatively. Our bound is linear in the total weight. This is in stark contrast to quadratic lower bounds for the worst case.

[10] Doskoč, V., Friedrich, T., Göbel, A., Neumann, A., Neumann, F., Quinzan, F., Non-Monotone Submodular Maximization with Multiple Knapsacks in Static and Dynamic Settings. In: *European Conference on Artificial Intelligence (ECAI)*, pp. 435–442, 2020.

We study the problem of maximizing a non-monotone submodular function under multiple knapsack constraints. We propose a simple discrete greedy algorithm to approach this problem, and prove that it yields strong approximation guarantees for functions with bounded curvature. In contrast to other heuristics, this does not require problem relaxation to continuous domains and it maintains a constant-factor approximation guarantee in the problem size. In the case of a single knapsack, our analysis suggests that the standard greedy can be used in non-monotone settings. Additionally, we study this problem in a dynamic setting, in which knapsacks change during the optimization process. We modify our greedy algorithm to avoid a complete restart at each constraint update. This modification retains the approximation guarantees of the static case. We evaluate our results experimentally on a video summarization and sensor placement task. We show that our proposed algorithm competes with the state-of-the-art in static settings. Furthermore, we show that in dynamic settings with tight computational time budget, our modified greedy yields significant improvements over starting the greedy from scratch, in terms of the solution quality achieved.

[11] Friedrich, T., Göbel, A., Neumann, F., Quinzan, F., Rothenberger, R., Greedy Maximization of Functions with Bounded Curvature Under Partition Matroid Constraints. In: *Conference on Artificial Intelligence (AAAI)*, pp. 2272–2279, 2019.

We investigate the performance of a deterministic GREEDY algorithm for the problem of maximizing functions under a partition matroid constraint. We consider non-monotone submodular functions and monotone subadditive functions. Even though constrained maximization problems of monotone submodular functions have been extensively studied, little is known about greedy maximization of non-monotone submodular functions or monotone subadditive functions. We give approximation guarantees for GREEDY on these problems, in terms of the curvature. We find that this simple heuristic yields a strong approximation guarantee on a broad class of functions. We discuss the applicability of our results to three real-world problems: Maximizing the determinant function of a positive semidefinite matrix, and related problems such as the maximum entropy sampling problem, the constrained maximum cut problem on directed graphs, and combinatorial auction games. We conclude that GREEDY is well-suited to approach these problems. Overall, we present evidence to support the idea that, when dealing with constrained maximization problems with bounded curvature, one needs not search for (approximate) monotonicity to get good approximate solutions.

[12] Göbel, A., Lagodzinski, J. A. G., Seidel, K., Counting Homomorphisms to Trees Modulo a Prime. In: *Mathematical Foundations of Computer Science (MFCS)*, pp. 49:1–49:13, 2018.

Many important graph theoretic notions can be encoded as counting graph homomorphism problems, such as partition functions in statistical physics, in particular independent sets and colourings. In this article we study the complexity of $\#_p\mathrm{HomsTo}H$, the problem of counting graph homomorphisms from an input graph to a graph H modulo a prime number p. Dyer and Greenhill proved a dichotomy stating that the tractability of non-modular counting graph homomorphisms depends on the structure of the target graph. Many intractable cases in non-modular counting become tractable in modular counting due to the common phenomenon of cancellation. In subsequent studies on counting modulo 2, however, the influence of the structure of H on the tractability was shown to persist, which yields similar dichotomies. Our main result states that for every tree H and every prime p the problem $\#_p\mathrm{HomsTo}H$ is either polynomial time computable or $\#_p$P-complete. This relates to the conjecture of Faben and Jerrum stating that this dichotomy holds for every graph H when counting modulo 2. In contrast to previous results on modular counting, the tractable cases of $\#_p\mathrm{HomsTo}H$ are essentially the same for all values of the modulo when H is a tree. To prove this result, we study the structural properties of a homomorphism. As an important interim result, our study yields a dichotomy for the problem of counting weighted independent sets in a bipartite graph modulo some prime p. These results are the first suggesting that such dichotomies hold not only for the one-bit functions of the modulo 2 case but also for the modular counting functions of all primes p.

[13] Friedrich, T., Göbel, A., Quinzan, F., Wagner, M., Heavy-tailed Mutation Operators in Single-Objective Combinatorial Optimization. In: *Parallel Problem Solving From Nature (PPSN)*, pp. 134–145, 2018.

A core feature of evolutionary algorithms is their mutation operator. Recently, much attention has been devoted to the study of mutation operators with dynamic and non-uniform mutation rates. Following up on this line of work, we propose a new mutation operator and analyze its performance on the (1+1) Evolutionary Algorithm (EA). Our analyses show that this mutation operator competes with pre-existing ones, when used by the (1+1)EA on classes of problems for which results on the other mutation operators are available. We present a jump function for which the performance of the (1+1)EA using any static uniform mutation and any restart strategy can be worse than the performance of the (1+1)EA using our mutation operator with no restarts. We show that the (1+1)EA using our mutation operator finds a (1/3)-approximation ratio on any non-negative submodular function in polynomial time. This performance matches that of combinatorial local search algorithms specifically designed to solve this problem. Finally, we evaluate experimentally the performance of the (1+1)EA using our operator, on real-world graphs of different origins with up to ∼37 000 vertices and ∼1.6 million edges. In comparison with uniform mutation and a recently proposed dynamic scheme our operator comes out on top on these instances.

[14] Galanis, A., Göbel, A., Goldberg, L.-A., Lapinskas, J., Richerby, D., Amplifiers for the Moran Process. In: *International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 62:1–62:13, 2016.

The Moran process, as studied by Lieberman, Hauert and Nowak, is a randomised algorithm modelling the spread of genetic mutations in populations. The algorithm runs on an underlying graph where individuals correspond to vertices. Initially, one vertex (chosen uniformly at random) possesses a mutation, with fitness $r > 1$. All other individuals have fitness 1. During each step of the algorithm, an individual is chosen with probability proportional to its fitness, and its state (mutant or non-mutant) is passed on to an out-neighbour which is chosen uniformly at random. If the underlying graph is strongly connected then the algorithm will eventually reach fixation, in which all individuals are mutants, or extinction, in which no individuals are mutants. An infinite family of directed graphs is said to be strongly amplifying if, for every $r > 1$, the extinction probability tends to 0 as the number of vertices increases. Strong amplification is a rather surprising property - it means that in such graphs, the fixation probability of a uniformly-placed initial mutant tends to 1 even though the initial mutant only has a fixed selective advantage of $r > 1$ (independently of n). The name "strongly amplifying" comes from the fact that this selective advantage is "amplified". Strong amplifiers have received quite a bit of attention, and Lieberman et al. proposed two potentially strongly-amplifying families - superstars and metafunnels. Heuristic arguments have been published, arguing that there are infinite families of superstars that are strongly amplifying. The same has been claimed for metafunnels. We give the first rigorous proof that there is an infinite family of directed graphs that is strongly amplifying. We call the graphs in the family "megastars". When the algorithm is run on an n-vertex graph in this family, starting with a uniformly-chosen mutant, the extinction probability is roughly $n^{-1/2}$ (up to logarithmic factors). We prove that all infinite families of superstars and metafunnels have larger extinction probabilities (as a function of n). Finally, we prove that our analysis of megastars is fairly tight - there is no infinite family of megastars such that the Moran algorithm gives a smaller extinction probability (up to logarithmic factors). Also, we provide a counterexample which clarifies the literature concerning the isothermal theorem of Lieberman et al. A full version [Galanis/Göbel/Goldberg/Lapinskas/Richerby, Preprint] containing detailed proofs is available at http://arxiv.org/abs/1512.05632. Theorem-numbering here matches the full version.

[15] Göbel, A., Goldberg, L. A., Richerby, D., Counting Homomorphisms to Square-Free Graphs, Modulo 2. In: *International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 642–653, 2015.

We study the problem ⊕HomsToH of counting, modulo 2, the homomorphisms from an input graph to a fixed undirected graph H. A characteristic feature of modular counting is that cancellations make wider classes of instances tractable than is the case for exact (nonmodular) counting; thus, subtle dichotomy theorems can arise. We show the following dichotomy: for any H that contains no 4-cycles, ⊕HomsToH is either in polynomial time or is ⊕P-complete. This partially confirms a conjecture of Faben and Jerrum that was previously only known to hold for trees and for a restricted class of tree-width-2 graphs called cactus graphs. We confirm the conjecture for a rich class of graphs, including graphs of unbounded tree-width. In particular, we focus on square-free graphs, which are graphs without 4-cycles. These graphs arise frequently in combinatorics, for example, in connection with the strong perfect graph theorem and in certain graph algorithms. Previous dichotomy theorems required the graph to be tree-like so that tree-like decompositions could be exploited in the proof. We prove the conjecture for a much richer class of graphs by adopting a much more general approach.

[16] Göbel, A., Goldberg, L. A., McQuillan, C., Richerby, D., Yamakami, T., Counting List Matrix Partitions of Graphs. In: *Conference on Computational Complexity (CCC)*, pp. 56–65, 2014.

Given a symmetric $D \times D$ matrix M over $0, 1, *$, a list M-partition of a graph G is a partition of the vertices of G into D parts which are associated with the rows of M. The part of each vertex is chosen from a given list in such a way that no edge of G is mapped to a 0 in M and no non-edge of G is mapped to a 1 in M. Many important graph-theoretic structures can be represented as list M-partitions including graph colourings, split graphs and homogeneous sets, which arise in the proofs of the weak and strong perfect graph conjectures. Thus, there has been quite a bit of work on determining for which matrices M computations involving list M-partitions are tractable. This paper focuses on the problem of counting list M-partitions, given a graph G and given lists for each vertex of G. We give an algorithm that solves this problem in polynomial time for every (fixed) matrix M for which the problem is tractable. The algorithm relies on data structures such as sparse-dense partitions and sub cube decompositions to reduce each problem instance to a sequence of problem instances in which the lists have a certain useful structure that restricts access to portions of M in which the interactions of 0s and 1s is controlled. We show how to solve the resulting restricted instances by converting them into particular counting constraint satisfaction problems (#CSPs) which we show how to solve using a constraint satisfaction technique known as "arc-consistency". For every matrix M for which our algorithm fails, we show that the problem of counting list M-partitions is #P-complete. Furthermore, we give an explicit characterisation of the dichotomy theorem - counting list M-partitions is tractable (in FP) if and only if the matrix M has a structure called a derectangularising sequence. Finally, we show that the meta-problem of determining whether a given matrix has a derectangularising sequence is NP-complete.

[17] Göbel, A., Goldberg, L. A., Richerby, D., Counting Homomorphisms to Cactus Graphs Modulo 2. In: *Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 350–361, 2014.

A homomorphism from a graph G to a graph H is a function from $V(G)$ to $V(H)$ that preserves edges. Many combinatorial structures that arise in mathematics and in computer science can be represented naturally as graph homomorphisms and as weighted sums of graph homomorphisms. In this article, we study the complexity of counting homomorphisms modulo 2. The complexity of modular counting was introduced by Papadimitriou and Zachos and it has been pioneered by Valiant who famously introduced a problem for which counting modulo 7 is easy but counting modulo 2 is intractable. Modular counting provides a rich setting in which to study the structure of homomorphism problems. In this case, the structure of the graph H has a big influence on the complexity of the problem. Thus, our approach is graph-theoretic. We give a complete solution for the class of cactus graphs, which are connected graphs in which every edge belongs to at most one cycle. Cactus graphs arise in many applications such as the modelling of wireless sensor networks and the comparison of genomes. We show that, for some cactus graphs H, counting homomorphisms to H modulo 2 can be done in polynomial time. For every other fixed cactus graph H, the problem is complete in the complexity class $\oplus P$, which is a wide complexity class to which every problem in the polynomial hierarchy can be reduced (using randomised reductions). Determining which H lead to tractable problems can be done in polynomial time. Our result builds upon the work of Faben and Jerrum, who gave a dichotomy for the case in which H is a tree.

[18] Bampas, E., Göbel, A.-N., Pagourtzis, A., Tentes, A., On the Connection between Interval Size Functions and Path Counting. In: *Theory and Applications of Models of Computation (TAMC)*, pp. 108–117, 2009.

We investigate the complexity of hard counting problems that belong to the class #P but have easy decision version; several well-known problems such as #Perfect Matchings, #DNFSat share this property. We focus on classes of such problems which emerged through two disparate approaches: one taken by Hemaspaandra et al. [1] who defined classes of functions that count the size of intervals of ordered strings, and one followed by Kiayias et al. [2] who defined the class TotP, consisting of functions that count the total number of paths of NP computations. We provide inclusion and separation relations between TotP and interval size counting classes, by means of new classes that we define in this work. Our results imply that many known #P-complete problems with easy decision are contained in the classes defined in [1]-but are unlikely to be complete for these classes under certain types of reductions. We also define a new class of interval size functions which strictly contains FP and is strictly contained in TotP under reasonable complexity-theoretic assumptions. We show that this new class contains some hard counting problems.