Finding Lean Induced Cycles
in Binary Hypercubes[*]

Yury Chebiryak[1], Thomas Wahl[1,2], Daniel Kroening[1,2], Leopold Haller[2]

[1] Computer Systems Institute, ETH Zurich, Switzerland
[2] Computing Laboratory, Oxford University, United Kingdom

Abstract. Induced (chord-free) cycles in binary hypercubes have many applications in computer science. The state of the art for computing such cycles relies on genetic algorithms, which are, however, unable to perform a complete search. In this paper, we propose an approach to finding a special class of induced cycles we call *lean*, based on an efficient propositional SAT encoding. Lean induced cycles dominate a minimum number of hypercube nodes. Such cycles have been identified in Systems Biology as candidates for stable trajectories of gene regulatory networks. The encoding enabled us to compute lean induced cycles for hypercubes up to dimension 7. We also classify the induced cycles by the number of nodes they fail to dominate, using a custom-built All-SAT solver. We demonstrate how *clause filtering* can reduce the number of blocking clauses by two orders of magnitude.

1 Introduction

Cycles through binary hypercubes have applications in numerous fields in computing. The design of algorithms that reason about them is an active area of research. This paper is concerned with obtaining a subclass of these cycles with applications in Systems Biology.

Biochemical reactions in gene networks are frequently modeled using a system of piece-wise linear ordinary differential equations (PLDE), whose number corresponds to the number of genes in the network [4]. It is of critical importance to obtain *stable* solutions, because only stable orbits describe biologically relevant dynamics of the genes. We focus on Glass PLDE, a specific type of PLDE that simulates neural and gene regulatory networks [7].

The phase flow of Glass networks spans a sequence of coordinate orthants, which can be represented by the nodes of a binary hypercube. The orientation of the edges of the hypercube is determined by the choice of focal points of the PLDE. The orientation of the edge shows the direction of the phase flow

---

at the coordinate plane separating the orthants. Thus, the paths in oriented binary hypercubes serve as a discrete representation of the continuous dynamics of Glass gene regulatory networks. A special kind of such paths, *coil-in-the-box* codes, is used for the identification of stable periodic orbits in the Glass PLDE. Coil-in-the-box codes with maximum length represent the networks with longest sequence of gene states for a given number of genes [10].

If a cycle in the hypercube is defined by a coil-in-the-box code, the orientation of all edges adjacent to the cycle can be chosen to direct the flow towards it (the cycle is then called a *cyclic attractor*). Such orientation ensures the convergence of the flow to a periodic attractor that lies in the orthants included in the path. If a node of the hypercube is not adjacent to the cycle, the node does not have edges adjacent to the cycle, and the orientation of the edges at this node does not affect the stability of the flow along the orthants that are defined by the coil-in-the-box code. The choice of edge orientation in turn is linked to the specification of focal points of the PLDE. Therefore, the presence of nodes that are not dominated indicates that the phase flow along the attractor is robust to any variations of the coefficients that define the equations in the orthant corresponding to that node [20]. We say that a node that is not dominated by the cycle is *shunned* by the cycle.

The computation of (preferably long) induced (i.e., chord-free) cycles that dominate as few nodes as possible is therefore highly desirable in this context. We call such cycles *lean induced cycles*.

The state-of-the art in computing longest induced cycles and paths relies on genetic algorithms [5]. However, while this technique is able to identify individual cycles with desired properties, it cannot guarantee completeness, i.e., it may miss specific cycles. Many applications, including those in Systems Biology, rely on a classification of *all* solutions, which precludes the use of any incomplete random search technique.

Recent research suggests that SAT-based algorithms can solve many combinatorial problems efficiently: applications include oriented matroids [18], the coverability problem for unbounded Petri nets [1], bounds on van der Waerden numbers [12, 6], and many more. Solving a propositional formula that encodes a desired combinatorial object with a state-of-the-art SAT solver can be more efficient than the alternatives.

Contribution We encode the problem of identifying lean induced cycles in binary hypercubes as a propositional SAT formula and compute solutions using a state-of-the-art solver. As we aim at the complete set of cycles, we modify the solver to solve the All-SAT problem, and present three orthogonal optimizations that reduce the number of required blocking clauses by two orders of magnitude.

Our implementation enabled us to obtain a broad range of new results on cycles of this kind. L. Glass presented a coil-in-the-box code with one shunned node in the 4-cube [10]. We show that this is the maximum number of shunned

nodes that any *lean* induced cycle may have for that dimension. Then, we show that the longest induced cycles in the next two dimensions are *cube-dominating*: these cycles dominate every node of the cube. In dimension 7, where an induced cycle can be almost twice as long as the shortest cube-dominating cycles, there are lean induced cycles shunning at least three nodes.

2 Preliminaries

We define basic concepts used frequently throughout the paper. The *Hamming distance* between two bit-strings $u = u_1 \ldots u_n$, $v = v_1 \ldots v_n \in \{0,1\}^n$ of length n is the number of bit positions in which u and v differ:

$$d_H^n(u,v) = |\{ i \in \{1,\ldots,n\} : u_i \neq v_i \}|.$$

The *n-dimensional Hypercube*, or *n-cube* for short, is the graph (V, E) with $V = \{0,1\}^n$ and $(u,v) \in E$ exactly if $d_H^n(u,v) = 1$ (see also [14]). The *n-cube* has $n \cdot 2^{n-1}$ edges. We use the standard definitions of *path* and *cycle* through the hypercube graph. The length of a path is the number of its vertices. A Hamiltonian path (cycle) through the *n*-cube is called a *(cyclic) Gray code*. The *cyclic distance* of two nodes W_j and W_k along a cycle of length L in the *n*-cube is

$$d_C^n(W_j, W_k) = \min\{|k - j|, \ L - |k - j|\}.$$

In this paper, we are concerned with particular cycles through the *n*-cube.

Definition 1. *An **induced cycle** $I_0 \ldots I_{L-1}$ in the n-cube is a cycle such that any two nodes on the cycle that are neighbors in the n-cube are also neighbors in the cycle:*

$$\forall j, k \in \{0, \ldots, L-1\} \quad (d_H^n(I_j, I_k) = 1 \ \Rightarrow \ d_C^n(I_j, I_k) = 1). \tag{1}$$

Fig. 1 shows an induced cycle (bold edges) in the 4-cube. In this paper, we are also interested in the immediate neighborhood of the cycle:

Definition 2. *The cycle $I_0 \ldots I_{L-1}$ **dominates** node W of the n-cube if W is adjacent to some node of the cycle:*

$$\exists j \in \{0, \ldots, L-1\} \quad d_H^n(I_j, W) = 1. \tag{2}$$

We say the cycle *shuns* the nodes it does not dominate. A cycle is called *cube-dominating* if it dominates every node of the *n*-cube; such cycles can be thought of as "fat". In contrast, in this paper we are interested in "lean" induced cycles, which dominate as few nodes as possible:

Definition 3. *A **lean** induced cycle is an induced cycle through the n-cube that dominates a minimum number of cube nodes, among all induced n-cube cycles **of the same length**.*

Especially significant are induced cycles of maximum length. The induced cycle in Fig. 1 is longest (length 8) in dimension 4. It is also lean, as it dominates 15 of the 16 cube nodes, and there is no induced cycle of length 8 dominating less than 15 nodes.

3

Fig. 1. A lean induced cycle in the 4-cube. The cycle shuns node 1101

Lean induced cycles in cell biology. Hypercubes with lean induced cycles can aid the synthesis of Glass Boolean networks with stable periodic orbits and stable equilibrium states. For example, *C. elegans* vulval development is known to exhibit a series of cell divisions with 22 nuclei formed in the end of the development. The cell division represents a complex reactive system and includes at least four different molecular signaling pathways [15]. If the state of every signaling pathway is represented by a valuation of a Boolean variable, the 4-cube in Fig. 1 is useful for synthesizing a Glass Boolean network with a stable periodic orbit describing the cell division and an equilibrium depicting the finale state (at node 1101) of the gene regulatory system.

Co-existence of an induced cycle of maximum length and a shunned node in a hypercube indicates that during cell division, the gene network may traverse the maximum possible number of the different states before switching to the final equilibrium.

3 Computing Lean Induced Cycles

In this section, we describe an encoding of induced cycles of a given length into a propositional-logic formula. We then strengthen the encoding to assert the existence of a certain number of shunned nodes. We finally illustrate how we used the MiniSat solver to determine lean induced cycles where this number of shunned nodes is maximized.

3.1 A SAT-Encoding of Induced Cycles with Shunned Nodes

Our encoding relies heavily on comparing the Hamming distance between two hypercube nodes against some constant. We implement such comparisons

4

efficiently using *once-twice* chains, as described in [3]. In brief, a once-twice chain identifies differences between two strings up to some position j based on (i) comparing them at position j, and (ii) recursively comparing their prefixes up to position $j - 1$.

Induced Cycles. We use $n \cdot L$ Boolean variables $I_j[k]$, where $0 \leq j < L$ and $0 \leq k < n$, to encode the coordinates of an induced cycle of length L in the n-cube. The variable $I_j[k]$ denotes the k-th coordinate of the j-th node. In order to form a cycle in an n-cube, consecutive nodes of the sequence must have Hamming distance 1, including the last and the first:

$$\varphi_{cycle} \quad := \quad \left(\bigwedge_{i=0}^{L-2} d_H^n(I_i, I_{i+1}) = 1 \right) \quad \wedge \quad d_H^n(I_{L-1}, I_0) = 1 \,.$$

To make the cycle induced, we eliminate chords as follows:

$$\varphi_{chord\text{-}free} \quad := \quad \bigwedge_{\substack{0 \leq i < j < L, \\ d_C^n(I_i, I_j) \geq 2}} d_H^n(I_i, I_j) \geq 2 \,.$$

This also ensures that the nodes along the cycle are pairwise distinct. In practice, the formula $\varphi_{chord\text{-}free}$ can be optimized by eliminating half of its clauses, using an argument presented in [2].

The conjunction of these constraints is an encoding of induced cycles:

$$\varphi_{IC} \quad := \quad \varphi_{cycle} \wedge \varphi_{chord\text{-}free} \,.$$

Shunned Nodes. We encode the property that a cycle $I_0 \ldots I_{L-1}$ shuns nodes $u_0, \ldots, u_{S-1}$, by requiring the distance of the nodes to the cycle to be at least 2:

$$\varphi_{shunned} \quad := \quad \bigwedge_{i=0}^{S-1} \bigwedge_{j=0}^{L-1} d_H^n(u_i, I_j) \geq 2 \,.$$

We combine this with the condition that the nodes are distinct,

$$\varphi_{distinct} \quad := \quad \bigwedge_{0 \leq i < j < S} d_H^n(u_i, u_j) \geq 1 \,,$$

to obtain an encoding of induced cycles with at least S shunned nodes:

$$\varphi_{ICS} \quad := \quad \varphi_{IC} \wedge \varphi_{shunned} \wedge \varphi_{distinct} \,. \tag{3}$$

We point out some basic monotonicity properties of formula φ_{ICS}. Let $IC(n, L, S^+)$ be the number of induced cycles of length L in the n-cube with at least S shunned hypercube nodes. It is easy to see that

$$\begin{aligned} n_1 \leq n_2 \quad &\Rightarrow \quad IC(n_1, L, S^+) \ \leq \ IC(n_2, L, S^+), \quad \text{and} \\ S_1 \leq S_2 \quad &\Rightarrow \quad IC(n, L, S_1^+) \ \geq \ IC(n, L, S_2^+) \,. \end{aligned}$$

There is no analogous monotonicity law for the length parameter L of an induced cycle. Intuitively, a medium value for L provides the greatest degree of freedom for a cycle.

5

dim. n	length L	max. # shunned nodes
3	6	0
4	8	1
5	14	0
6	26	0
7	48	≥ 3

Table 1. Length of longest induced cycles, and number of shunned nodes

3.2 Computing Lean Induced Cycles using a SAT Solver

Every solution to equation (3) corresponds to an induced cycle of length L in the n-cube with at least S shunned nodes. In order to make the cycle *lean*, we need to maximize S. We achieve this by starting with *cube-dominating* induced cycles, i.e., with $S = 0$, and increasing S in equation (3) until the SAT solver reports unsatisfiability.[3]

Table 1 shows our findings for hypercubes up to dimension 7. For the classical cube of dimension 3, the longest induced cycles have length 6. All of those are cube-dominating. In dimension 4, the longest induced cycles have length 8; an example is shown in Fig. 1. Some of these cycles shun 1 of the 16 cube nodes; the others are cube-dominating. Interestingly, in dimensions 5 and 6, all longest induced cycles are again cube-dominating.

In dimension 7, we found longest (length 48) induced cycles shunning 3 nodes. For larger values of S, our search timed out after 24h. In our experiments, we used the MINISAT solver by Eén and Sörensson [9]. MINISAT provides interfaces for incremental solving and All-SAT; the current version uses preprocessing techniques [8] that simplify the original formula. All experiments were carried out on an Intel Xeon 3.0 GHz, 4-GB RAM PC running Linux.

4 Classification of Induced Cycles

The goal of this section is to determine how many distinct induced cycles of length L and with S shunned nodes exist in the n-cube, for a given triple (n, L, S). By *distinct*, we mean that the cycles cannot be transformed into each other by applying a symmetry permutation of the n-cube. That is, for each tuple (n, L, S), we *classify* the induced cycles into equivalence classes.

The classification of induced cycles with respect to symmetries of a hypercube is of interest in Glass models for neural and gene regulatory networks, because the number of the equivalence classes of the codes indicates how many different types of cells can be regulated by a set of genes [21, 10].

---

[3] Since the range of values for S for which (3) is satisfiable is contiguous, a binary search strategy is also possible, using a heuristically determined initial value for S.

The enumeration of the equivalence classes is achieved using a custom-made All-SAT solver derived from MiniSat. We introduce blocking clauses that suppress solutions symmetric to one encountered before. We observe that cycles identical up to cube symmetries belong to the same class (n, L, S). This ensures that the symmetry breaking does not eliminate solutions with a different set of parameters. In the rest of this section, we describe the classification and the symmetry breaking in more detail.

4.1 Identifying Equivalence Classes using Coordinate Sequences

In order to identify symmetry equivalence classes of cycles, it proved efficient to encode cycles in a slightly different way.

Definition 4 ([10]). *The **coordinate sequence** of a cycle $I_0 \ldots I_{L-1}$ in the n-cube is the sequence $(c_0, \ldots, c_{L-1}) \in \{0, \ldots, n-1\}^L$ such that c_i is the unique coordinate that distinguishes I_i and $I_{i+1 \bmod L}$.*

For example, the coordinate sequence of the cycle in Fig. 1 is the sequence $cs := (0, 1, 2, 0, 3, 2, 1, 3)$, assuming $I_0 = 0000$ and $I_1 = 0001$. In the figure, the dimensions are listed in the order 3210.
Given coordinate sequences, we can define cube symmetries.

Definition 5. *Two cycles C_1 and C_2 in the n-cube are **equivalent**, $C_1 \sim C_2$, if their coordinate sequences are identical up to axis permutations, reflections about the center position, and rotations by an arbitrary number of coordinates.*

Given n and L, let CS denote the set of coordinate sequences of cycles of length L in the n-cube. A reflection or rotation on CS is a permutation π on the set $\{0, \ldots, L-1\}$ that maps a coordinate sequence $(c_i)_{i=0}^{L-1}$ to the sequence $(c_{\pi(i)})_{i=0}^{L-1}$, that is, by acting on the position indices of the sequence. In contrast, an axis permutation on CS is a permutation π on the set $\{0, \ldots, n-1\}$ that maps a coordinate sequence $(c_i)_{i=0}^{L-1}$ to the sequence $(\pi(c_i))_{i=0}^{L-1}$, that is, by acting on the coordinate values of the sequence. For example, the coordinate sequence $cs' := (1, 0, 2, 3, 0, 1, 3, 2)$ is equivalent to sequence cs above, since cs' can be obtained from cs by a left-rotation by one position, followed by a reflection and an axis permutation $(1\ 2\ 3\ 0)$, mapping 1 to 2, 2 to 3, etc.
Our goal is to classify induced cycles based on cube symmetries, for a given parameter tuple (n, L, S). In order for this classification to be sound, the symmetry permutations must not alter the (n, L, S) parameters of a cycle.

Lemma 1. *Let C_1 and C_2 be two equivalent cycles. Then C_1 and C_2 have the same length and shun the same number of cube nodes.*

Proof (sketch). Since C_1 and C_2 are equivalent, there is a sequence Π of permutations, of the type mentioned in definition 5, such that $\Pi(C_1) = C_2$. Reflections and rotations of the coordinate sequence of C_1 translate to reversals of

C_1's orientation, and to rotations of C_1, respectively. These operations change neither the length of the cycle, nor the distance of cube nodes to it.

For an axis permutation π, we have to show that definition 2, *dominates*, is invariant under π. We omit the technical derivation of this property. □

As an example, the unique cycle of the 4-cube corresponding to the above-mentioned coordinate sequence cs', after fixing $I_0 := 0000$ and $I_1 := 0010$, is lean and induced, as is the cycle in Fig. 1. Both cycles shun one cube node. Conversely, cycles with the same parameters (n, L, S) may not be equivalent: Table 2 (see Appendix) lists two distinct – in the above sense – cycles with $(n, L, S) = (4, 8, 0)$.

We determine the number $IC(n, L, S)$ of $\sim$ equivalence classes of induced cycles of length L with *exactly* S shunned nodes as the difference between the number of classes of cycles shunning at least S and $S + 1$ nodes, respectively:

$$IC(n, L, S) \quad = \quad IC(n, L, S^+) \; - \; IC(n, L, (S + 1)^+). \qquad (4)$$

The quantities on the right are computed, separately for S and $S + 1$, by enumerating satisfying assignments to Eq. (3), using an All-solutions SAT solver, implemented on top of MiniSat (see Algorithm 1 on the next page).

As proposed in [3], we encode coordinate sequences using XOR gates on Boolean variables denoting coordinates of a cycle. We write $xor^k[m]$ to refer to the m-th bit in bitwise xor-operation over coordinates of nodes I_k and $I_{k+1 \mod L}$. For example, if $xor^3[2]$ evaluates to true, dimension 2 is traversed while going from I_3 to I_4. We call the variables $xor^k[m]$ the "xor-variables".

To ensure a single representative for each $\sim$ equivalence class, we add blocking clauses for each solution found that prevent permutations of axes, rotations and reflections of the coordinate sequence of the solution. The number of blocking clauses to add per solution is $(2L \cdot n!)$. This is clearly a computational burden for the SAT solver, especially when the solution space is nearly exhausted, and the All-SAT procedure is about to find the formula to be unsatisfiable. In the rest of this section, we present techniques that reduce both the number and the length of the blocking clauses.

4.2 Optimizations

Compressing blocking clauses. A blocking clause for a given induced cycle, barring permutations of axes and rotations/reflections of a coordinate sequence, is expressed in terms of the variables encoding the sequence. For instance, to

---

Algorithm 1: COMPUTE-EQUIVALENCE-CLASSES

Input: the SAT instance I with fixed n, L, S;
         the equivalence relation $\sim$
Output: The set of equivalence classes EC

1:   $EC := \{\}$
2:   SAT_solver.solve(I)
3:   **while** SAT
4:     **do** IC = SAT_solver.decode()
5:       $EC \leftarrow EC \cup \{IC\}$
6:       $\forall IC_j \sim IC.$ I.add_blocking_clause(IC_j)
7:       SAT_solver.solve(I)

---

block permutations of the cycle in Fig. 1, we add the following clause:

$$(\neg xor^0[0] \lor xor^0[1] \lor xor^0[2] \lor xor^0[3]$$
$$\lor xor^1[0] \lor \neg xor^1[1] \lor xor^1[2] \lor \ xor^1[3]$$
$$\vdots$$
$$\lor xor^7[0] \lor xor^7[1] \lor xor^7[2] \lor \neg xor^7[3] \ ) \ .$$

The length of this blocking clause is $(n \cdot L)$. Our first, and simplest, optimization is to omit literals that evaluate to *false*, since we know that these variables encode unit Hamming distance:

$$(\neg xor^0[0] \lor \neg xor^1[1] \lor \neg xor^2[2] \lor \neg xor^3[0] \lor \dots) \ .$$

This reduces the length of a clause to L.

Symmetric Cycles. The following optimization applies to specific cycles, called *symmetric induced cycles*. A Gray code is *symmetric*[4] if elements of its coordinate sequence that are $L/2$ apart are identical [19]. For a symmetric induced cycle, the number of blocking clauses to be added can be reduced by one-half: rotations by more than $L/2$ positions result in cycles that were already blocked.

Prefix Filtering. Without loss of generality, we fix the first two elements of the coordinate sequence to $(0, 1)$. For the next coordinate, dimension 0 cannot be traversed because this would form a chord. Neither can dimension 1, since

---

[4] This definition is not to be confused with the definition in [13], where this term refers to a code for which the number of bit changes is uniformly distributed among the bit positions, hence called a *balanced* Gray code in [17, p. 7].

the cycle must be simple. Out of higher dimensions, we can restrict the search to the *canonical class*[5] with prefix $(0, 1, 2)$. We enforce this prefix by fixing the values of the corresponding xor-variables using the following three clauses:

$$xor^0[0] \wedge xor^1[1] \wedge xor^2[2] \ . \tag{5}$$

This drastically reduces the number of solutions in each equivalence class, and eliminates a large number of blocking clauses. For example, it becomes unnecessary to add a blocking clause for the coordinate sequence cs' on page 7, as cs' is blocked by Eq. (5).

Phase Saving. In an attempt to speed up the enumeration of solutions, we added phase-saving [16] to MiniSat. By default, MiniSat assigns *false* to all decision variables. With phase saving, they are assigned their most recent values in the search. Phase saving combines well with aggressive restarting schemes, since it retains more information between restarts. Our intuition was that after finding a solution, the solver might be able to quickly identify neighboring solutions. Phase-saving alone, however, did not result in any speedups.

Ordering decision variables. Upon closer inspection of All-SAT runs, we found that the activity-based variable selection heuristic mainly chooses from a small set of branching variables. These variables correspond closely to the encoding of solutions in the input CNF. In order to make use of this insight, we extended the solver to allow for prioritization of important variables in the decision heuristic: In this modification, unprioritized variables are only considered for branching after all prioritized variables are assigned a value. We tested a number of possible restrictions, and found that prioritizing the variables that encode the induced-cycle nodes $I_0, \ldots, I_{L-1}$ works well for some instances, but yields bad results in general.

Combined Restart Policy. We found that the enumeration of solutions could be sped up by disabling the geometric restart scheme, but this led to bad performance on the final hard instances. By combining an initial high restart limit (100000 conflicts) with a subsequent switch to MiniSat's original geometric policy, starting again from a very low limit (100 conflicts), we were able to gain a 20% overall speed-up. Easier SAT instances can then be solved before the first restart, while hard instances still profit from aggressive restarts.

Further experiments with different combinations of the discussed strategies revealed that a combination of a high-restart limit, variable prioritization, and phase saving also led to a performance increase of about 20%.

---

[5] A canonical coordinate sequence is the one in which each coordinate k appears before the first appearance of $k + 1$ [11].

4.3 Evaluation

Using prefix filtering and the optimizations for symmetric cycles, we are able to reduce the number of clauses drastically. As an example, consider an instance encoding induced cycles of length 26 in a 6-dimensional hypercube. In order to block a solution, we need to add only 312 blocking clauses in the non-symmetric case and 156 clauses for a symmetric cycle, instead of originally 37440. Our findings are presented in Fig. 2 and extend the classification presented in [22].

For some circuit length values L, the time required by the All-SAT solver increases with the number of shunned nodes. For such values of L, it is faster to perform the classification for a small value of S and then check how many nodes the cycles dominate.

In general, the time required to find the first induced cycle is a few orders of magnitude less than that to perform the classification, even in the case of one class only, as the run-time is dominated by the final unsatisfiable instance.

5 Conclusion

In this paper we have formalized a combinatorial problem relevant in Systems Biology: finding lean induced cycles in a hypercube, i.e., induced cycles that dominate a minimum number of hypercube nodes. We have presented a solution to this problem based on an efficient SAT encoding, and used this encoding to find lean induced cycles using a SAT solver. When compared to genetic algorithms, our method can provide guarantees for finding solutions, or prove the absence thereof.

Our method is suitable for classifying large sets of solutions into symmetry equivalence classes. As suggested by Fig. 2, this allows insights into the distribution of distinct solutions across the parameters n, L, and S. The SAT solver's performance is improved by filtering blocking clauses based on combinatorial properties of induced cycles, and by applying All-SAT specific internal tunings.

Fig. 2. Classification of induced cycles by cube symmetries, for select triples (n, L, S)

References

1. P. A. Abdulla, S. P. Iyer, and A. Nylén. SAT-solving the coverability problem for Petri nets. *Formal Methods in System Design*, 24(1):25–43, 2004.
2. Y. Chebiryak and D. Kroening. An efficient SAT encoding of circuit codes. In *Procs. IEEE International Symposium on Information Theory and its Applications*, pages 1235–1238, Auckland, New Zealand, December 2008.
3. Y. Chebiryak and D. Kroening. Towards a classification of Hamiltonian cycles in the 6-cube. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)*, 4:57–74, 2008.
4. H. de Jong and M. Page. Search for steady states of piecewise-linear differential equation models of genetic regulatory networks. *IEEE/ACM Trans. Comput. Biology Bioinform.*, 5(2):208–222, 2008.
5. P. A. Diaz-Gomez and D. F. Hougen. Genetic algorithms for hunting snakes in hypercubes: Fitness function analysis and open questions. In *SNPD-SAWN '06: Proceedings of the Seventh ACIS International Conference on Software Engineering, Artificial*

Intelligence, Networking, and Parallel/Distributed Computing, pages 389–394, Washington, DC, USA, 2006. IEEE Computer Society.

6. M. R. Dransfield, V. W. Marek, and M. Truszczynski. Satisfiability and computing van der Waerden numbers. In E. Giunchiglia and A. Tacchella, editors, *SAT*, volume 2919 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 2003.

7. R. Edwards. Symbolic dynamics and computation in model gene networks. *Chaos*, 11(1):160–169, 2001.

8. N. Eén and A. Biere. Effective preprocessing in SAT through variable and clause elimination. In *Theory and Applications of Satisfiability Testing (SAT)*, volume 3569 of *LNCS*, pages 61–75. Springer, 2005.

9. N. Eén and N. Sörensson. An extensible SAT-solver. In *Theory and Applications of Satisfiability Testing (SAT)*, volume 2919 of *LNCS*, pages 502–518. Springer, 2004.

10. L. Glass. Combinatorial aspects of dynamics in biological systems. In U. Landman, editor, *Statistical mechanics and statistical methods in theory and applications*, pages 585–611. Plenum Press, 1977.

11. D. E. Knuth. *The Art of Computer Programming*, volume 4, fascicle 2: Generating All Tuples and Permutations. Addison-Wesley Professional, 2005.

12. M. Kouril and J. V. Franco. Resolution tunnels for improved SAT solver performance. In F. Bacchus and T. Walsh, editors, *Theory and Applications of Satisfiability Testing (SAT)*, volume 3569 of *Lecture Notes in Computer Science*, pages 143–157. Springer, 2005.

13. X. Liu and G. F. Schrack. A heuristic approach for constructing symmetric Gray codes. *Appl. Math. Comput.*, 155(1):55–63, July 2004.

14. M. Livingston and Q. Stout. Perfect dominating sets. *Congressus Numerantium*, 79(187–203):187–203, 1990.

15. D. Na'aman Kam, H. Kugler, A. Rami Marelly, J. Hubbard, and M. Stern. Formal modelling of C. elegans development. A scenario-based approach. *Modelling in Molecular Biology*, pages 151–174, 2004.

16. K. Pipatsrisawat and A. Darwiche. A lightweight component caching scheme for satisfiability solvers. In J. Marques-Silva and K. A. Sakallah, editors, *SAT*, volume 4501 of *Lecture Notes in Computer Science*, pages 294–299. Springer, 2007.

17. C. Savage. A survey of combinatorial Gray codes. *SIAM Review*, 39(4):605–629, 1997.

18. L. Schewe. Generation of oriented matroids using satisfiability solvers. In A. Iglesias and N. Takayama, editors, *ICMS*, volume 4151 of *Lecture Notes in Computer Science*, pages 216–218. Springer, 2006.

19. R. C. Singleton. Generalized snake-in-the-box codes. *IEEE Transactions on Electronic Computers*, EC-15(4):596–602, 1966.

20. I. Zinovik, Y. Chebiryak, and D. Kroening. Cyclic attractors in Glass models for gene regulatory networks. *IEEE Trans Inf Theory: Special Issue on Molecular Biology and Neuroscience*, December 2009. Accepted.

21. I. Zinovik, D. Kroening, and Y. Chebiryak. An algebraic algorithm for the identification of Glass networks with periodic orbits along cyclic attractors. In H. Anai, K. Horimoto, and T. Kutsia, editors, *Algebraic Biology (AB)*, volume 4545 of *LNCS*, pages 140–154. Springer, 2007.

22. I. Zinovik, D. Kroening, and Y. Chebiryak. Computing binary combinatorial Gray codes via exhaustive search with SAT-solvers. *IEEE Transactions on Information Theory*, 54(4):1819–1823, April 2008.

6 Appendix

Table 2 shows runtimes, and number of equivalence classes of induced cycles found, for various values of (n, L, S).

Table 2: Classification of induced cycles, with runtimes

n	L	S	Time (sec)		#classes	
			first cycle	All-SAT	$IC(n, L, S^+)$	$IC(n, L, S)$
4	6	0	0.003	0.010	1	0
		1	0.006	0.017	1	0
		2	0.010	0.028	1	1
		3		0.031		0
	8	0	0.007	0.305	3	2
		1	0.009	0.048	1	1
		2		0.015		0
5	10	0	0.016	400.378	10	0
		1	0.017	419.881	10	0
		2	0.027	392.274	10	3
		3	0.031	370.277	7	3
		4	0.047	356.335	4	3
		5	0.043	210.137	1	0
		6	0.095	183.403	1	1
		7		167.397		0
	14	0	0.033	535.027	3	3
		1		3.012		0
6	16	0	0.02	486.37	563	1
		1	0.01	534.55	562	0
		2	0.03	481.12	562	1
		3	0.02	514.36	561	1
		4	0.04	501.77	560	13
		5	0.04	768.08	547	14
		6	0.04	3252.77	533	44
6	24	0	0.08	1183.50	110	76
		1	0.07	695.37	34	14
		2	0.30	689.22	20	15
		3	1.76	592.51	5	3
		4	5.65	1062.92	2	1
		5	26.56	1364.86	1	1
		6	-	1014.34		0
6	26	0	0.42	583.43	4	4
		1	-	750.39	0	0