Test Generation for Model-Based Diagnosis

Gregory Provan[1]
Department of Computer Science, University College Cork, Cork, Ireland
g.provan@cs.ucc.ie

ABSTRACT
This article formalises the dual problem to model-based diagnosis (MBD), i.e., generating *tests* to isolate multiple simultaneous faults. Using a standard propositional MBD framework, we first define a test of minimal size that can isolate multiple simultaneous faults of an arbitrary nature. Second, we prove complexity results for multiple-fault tests of minimal size in propositional system models, showing such problems have complexity similar to those of MBD problems, i.e., complexity at the second level of the polynomial hierarchy.

1 Introduction

Many real-world applications have focused on identifying minimal sensor observations, or *tests*, that will accurately isolate faults. For example, in the aerospace industry, it is common practice to pre-compute minimal test sets for systems, and also to build special equipment, called Built-in Test Equipment (or BITE), to conduct such fault-isolation tests (called Built-In Tests) [11]. The BITE will then conduct pre-defined sensor analysis upon system start-up or malfunction. This area of diagnosis, called Automated Test Pattern Generation (ATPG) [3], has significant practical importance. ATPG, in more precise terms, concerns pre-computing a set of tests that can validate circuits or diagnose hardware in embedded applications [3].

In contrast, model-based diagnosis (MBD) [14] addresses the task of isolating multiple simultaneous faults, given an *arbitrary observation* α. Although the objectives of efficient fault isolation are similar, the approaches, and the entailed algorithms, are quite different.

This article formalises ATPG as a dual problem to MBD. Using a standard propositional MBD framework, we first define a test that can isolate multiple simultaneous faults, of an arbitrary nature. This formalisation extends standard notions of ATPG, which are restricted to generating tests that can isolate *single, stuck-at* faults, namely, faults in discrete circuits that occur when wires or gates are stuck at either 0 or 1 [3, 15]. In this article we address models framed as Boolean functions, since all existing ATPG is based on Boolean functions. However, our framework applies to more general models; for example, we can generalise our results to multi-valued propositional functions in a straightforward fashion, thus allowing our results to be applicable to qualitative MBD/ATPG models with finite, discrete-valued variables.

Our contributions are as follows. First, we define a general ATPG model, cast within a satisfiability (SAT) framework, which adopts the multiple-fault definitions of MBD [14]. This general ATPG problem addresses multiple simultaneous faults of an arbitrary nature. Second, we prove complexity results for these more general ATPG models, demonstrating that these problems lie at the second level of the polynomial hierarchy. Given the intractability of MBD inference on real-world circuits [13], it is likely that multiple-fault ATPG will also be intractable, since both the MBD and multiple-fault ATPG are at the second level of the polynomial hierarchy.

[1] Supported by SFI grant 04/IN3/I524.

2 Related Work

This article applies a standard MBD framework [14] to test generation. To our knowledge, no article has formalised a similar approach for multiple-fault ATPG using the MBD consistency-based framework. Further, the ATPG literature lacks any methodology that directly generates tests to isolate faults of arbitrary size, often leading to inefficient and inaccurate fault isolation; ATPG focuses instead on iterated single-fault approaches for this task, e.g., [18, 19]. ATPG also differs from *sequential diagnosis* [17], in that it computes all tests *a priori*; in contrast, test sequencing computes the next best test *dynamically*.

The complexity of test generation has been carefully studied, and we review the complexity of ATPG and MBD.

ATPG Complexity: The standard single-fault ATPG model has been proven to be NP-complete [8], as it is an instance of the well-known SAT and CIRCUIT-SAT problems. Although this worst-case result indicates intractability, in practice test-generation is tractable because of the structural properties of circuits: the complexity is exponential in the undirected circuit cut-width [12], and the average-case complexity is poly-time in circuit size for circuits with bounded cut-width. However, in the real world systems, multiple faults can occur simultaneously, necessitating multiple-fault ATPG methods.

MBD Complexity: Computing a diagnosis which is minimal with respect to subset-inclusion or fault cardinality is NP-complete for propositional Horn models [1, 6]. The complexity for arbitrary propositional models, for several minimality conditions, has been shown to be at the second level of the polynomial hierarchy [4]. The complexity becomes less tractable if we instead consider the problem of computing the *set of all minimal diagnoses*. This problem is at least as difficult as counting the number of diagnoses, which has been shown to be #co-NP-Complete [7]. These results indicate that the diagnosis problems in which we are interested, i.e., computing the set of minimal-cardinality diagnoses over propositional models, are intractable.

3 Definitions and Notation

We now introduce notation for Boolean functions, MBD and ATPG.

3.1 Boolean Functions

Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be a Boolean function over a set $X = \{x_1, x_2, ..., x_n\}$ of n variables. A conjunctive normal form (CNF) Boolean formula f on Boolean variables X is a conjunction of m clauses $\{C_1, C_2, ...C_m\}$. Each clause C_i is a disjunction of k_i literals $l_1, ...l_{k_i}$. A literal is an instance of a variable or its complement.

We call λ_f an *assignment* for f if we set each of the variables in f to either 0 or 1; such an assignment may be represented by an n-bit vector in $\{0,1\}^n$ in the natural way.

A *satisfying assignment* λ, or model, for f is one for which $f(\lambda) = 1$. The set of satisfying assignments for f is denoted by Λ_f. A *partial assignment* is obtained when only a subset of variables in X is assigned values. A partial assignment may be represented by a vector of length n, each of whose elements is either 0, 1, or $*$.

The generic satisfiability problem can be defined as follows.

Definition 1 (SAT) *Given a CNF formula, $f(x_1, x_2, ..., x_n)$, the Boolean satisfiability problem SAT (f) has an answer YES iff there exists an assignment λ of Boolean values to the variables $x_1, x_2, ..., x_n$, i.e., $\exists\lambda$ such that f evaluates to 1.*

3.2 Model-Based Diagnosis Problem

This section introduces the notion of system model and diagnosis that we use to generalise ATPG to multiple-fault scenarios. MBD models are applicable to arbitrary systems, but in the following we assume that we are defining the MBD problem for circuits, to ensure consistency between the MBD and ATPG formalisms. In future work we will define ATPG for arbitrary systems.

Central to MBD, a *model* of an artifact is represented as a Boolean propositional function f over X. Distinguishing two subsets of these variables as *assumable* and *observable*[2] variables gives us a diagnostic system. We use a standard specification of MBD [14], except that we use the notion of health assignment (defined below) rather than sets of assumables to denote system health status.

Definition 2 (Diagnostic System) *A diagnostic system Ψ is defined as the triple $\Psi = \langle f, \mathcal{H}, \mathcal{O}\rangle$, where f is a propositional theory over a set of variables X, $\mathcal{H} \subseteq X$ is the set of assumables, and $\mathcal{O} \subseteq X$ is the set of observables.*

Throughout this paper we will assume that $\mathcal{O} \cap \mathcal{H} = \emptyset$ and $f \not\models \bot$.

The traditional query in MBD computes terms of assumable variables, which are explanations for the system description and an observation. By convention, if the set of health variables $\mathcal{H} = \{h_1, ..., h_m\}$, then $h_i = 0$ denotes normal functionality for component i, and $h_i = 1$ denotes a fault.

Definition 3 (Health Assignment) *Given a diagnostic system $\Psi = \langle f, \mathcal{H}, \mathcal{O}\rangle$, an assignment $\lambda_\mathcal{H}$ to all variables in $\mathcal{H}$ is defined as a health assignment.*

The MBD notion of diagnosis covers multiple simultaneous faults, as denoted below.

Definition 4 (Diagnosis) *Given a diagnostic system $\Psi = \langle f, \mathcal{H}, \mathcal{O}\rangle$, and an observation α over some variables in $\mathcal{O}$, a health assignment ω is a diagnosis iff $f \wedge \alpha \wedge \omega \not\models \bot$.*

[2] In the MBD literature the assumable variables are also referred to as "component", "failure-mode", or "health" variables. Observable variables are also called "measurable" or "control" variables.

In the MBD literature, a range of types of "preferred" diagnosis has been proposed. This turns the MBD problem into an optimization problem. In the following definition we consider the common subset- and cardinality-ordering.

Definition 5 (Minimal Diagnosis) *A diagnosis ω is defined as minimal, if no diagnosis ω' exists such that the set of negative literals in ω' form a proper subset of the set of negative literals in ω.*

Definition 6 (Diagnosis Cardinality) *The cardinality of a diagnosis, denoted as $|\omega|$, is defined as the number of negative literals in ω. A diagnosis is defined as cardinality-minimal, denoted $MinCard(\omega)$, if it minimizes the number of negative literals.*

3.3 ATPG Definitions

We now define the tasks performed by ATPG in terms of computing tests (satisfying assignments) that can isolate faults in a system specified in terms of a Boolean function f. ATPG traditionally uses a particular type of Boolean function, a Boolean circuit.

Definition 7 (Boolean circuit) *A Boolean circuit C is a directed acyclic graph (DAG) with distinguished observables* IN, OUT, *such that* IN$\cup$OUT$= \mathcal{O}_C$ *and* IN$\cap$ OUT$= \emptyset$*, where the vertices are labeled as follows:*

- *The input vertices,* IN*, labeled with a variable x_i or a constant (0 or 1), have fan-in 0.*
- *The output vertices,* OUT*, labeled "output", have fan-out 0.*
- *The gate vertices, $\mathcal{H} = \{h_1, ..., h_m\}$, with fan-in $k > 0$, are labeled with a Boolean function $h_i \in \mathcal{H}$ on k inputs $(\vee, \wedge, \neg)$, where the $\neg$ gate has fan-in of 1.*

We now define notions of faults on Boolean circuits, using notation introduced in [12]. The standard ATPG notion of a fault and faulted-circuit is restricted to a single faulty gate with a *stuck-at* fault.

Definition 8 (Single stuck-at-fault) *Given a Boolean circuit C, a single stuck-at fault $\phi(h_i, \nu)$ causes a component $h_i \in \mathcal{H}$ to be permanently stuck at logic value $\nu \in \{0,1\}$.*

Definition 9 (Faulted circuit) *Given a circuit C and a single stuck-at fault $\phi(h_i, \nu)$, a faulted circuit, denoted by C_ϕ, is the circuit C with the output for h_i set to value ν.*

We can translate a circuit into a propositional formula which corresponds to a diagnostic system, $\Psi_C = \langle f, \mathcal{H}, \mathcal{O}_C\rangle$, in which we restrict observables to $\mathcal{O}_C$ (which covers only the inputs and outputs of the system) and translate each gate into an equivalent logical formula. Based on this translation, the notion of faulted circuit corresponds to a single-fault health assignment ω^i in which only one health variable is set to 1, i.e., for one $i \in \{1, ..., m\}$, $h_i = 1$ and all other $h_{j:j\neq i} = 0$. We denote a "nominal" health assignment, where all $h_i = 0$, using $\omega^\emptyset$.

Definition 10 (Test) *Given a circuit diagnostic system $\Psi_C = \langle f, \mathcal{H}, \mathcal{O}_C\rangle$, and a health assignment ω^i, a test is an instantiation α over some variables in $\mathcal{O}$ such that $f \wedge \alpha \wedge \omega^i \not\models \bot$.*

Just as in MBD, we are interested in computing tests with respect to some completeness and minimality criteria. A test set is *complete* if it detects all single stuck-at faults; a complete test set of minimal size is a *minimal* test set.

We now define the (single-fault) ATPG problem using consistency-based terminology from MBD:

Definition 11 (ATPG testability problem) *Given a Boolean circuit diagnostic system, Ψ_C, a single stuck-at fault $\phi(h_i, \nu)$, corresponding to health assignment ω^i, is testable if and only if there exists a test α such that $f \wedge \alpha \wedge \omega^i \not\models \bot$ and $f \wedge \alpha \wedge \omega^\emptyset \models \bot$. Otherwise the fault is said to be untestable.*

The ATPG problem is to compute a minimal set of tests such that every single stuck-at fault is testable.

3.4 Review of Complexity Classes

In this article we will be defining the complexity of several problems, and we introduce the polynomial hierarchy as the standard means for classifying different complexity classes. The best-known complexity classes are P and NP: P is the set of languages possessing algorithms that run in time that is a polynomial in the length of the input; NP is the set of languages possessing algorithms that run in nondeterministic polynomial time. Since we will be considering problems that are harder than those in P and NP, we define the polynomial hierarchy in terms of languages as follows. A language $\mathcal{L}$ is in the class Σ_i^P iff there is another language $\mathcal{L}'$ in the class P and an integer k for which $\mathcal{L} = \{x : (\exists y_1)(\forall y_2)(\exists y_3)\cdots(Qy_i), |y_i| = |x|^k$ for all $i, [(x, y_1, y_2, \cdots, y_i) \in \mathcal{L}']\}$, where the sequence of quantifiers alternates, ending with $Q = \exists$ if i is odd or $Q = \forall$ if i is even. According to this definition, $P = \Sigma_0^P$ and $P = \Sigma_1^P$. We can also define the complementary hierarchy Π_i^P of problems, which denote the problems defined by co$\mathcal{L} = \{L : \bar{L} \in \mathcal{L}\}$, or in other words, $co\Sigma_i^P = \Pi_i^P$, for $i = 0, ..., \infty$. In a manner analogous to NP-hard problems being computationally more difficult than P-hard problems, Σ_2^P-hard problems are computationally more difficult than NP-hard (or Σ_1^P-hard) problems.

4 Multiple-Fault ATPG

We now define more general notions of ATPG problems, using the multiple-fault specifications defined in the previous section.

Definition 12 (Multiple-fault Health Assignment) *Given a diagnostic system $\Psi = \langle f, \mathcal{H}, \mathcal{O}\rangle$, a multiple-fault health assignment ω^I is an instantiation of $\mathcal{H}$ where at least 2 elements of ω^I are set to 1.*

We can now define a multiple-fault ATPG test:

Definition 13 (Multiple-fault Test) *Given a circuit diagnostic system $\Psi_C = \langle f, \mathcal{H}, \mathcal{O}_C\rangle$, and a multiple-fault health assignment ω^I, a multiple-fault ATPG test is an instantiation α over some variables in $\mathcal{O}$ such that $f \wedge \alpha \wedge \omega^I \not\models \bot$ and $f \wedge \alpha \wedge \omega^\emptyset \models \bot$.*

The multiple-fault ATPG problem is hence defined as the problem of computing a test-set that can isolate every multiple fault combination in $2^{\mathcal{H}}$. We now show the complexity of a decision version of Definition 13, MF-TEST.

Problem 1 (MF-TEST) *Given a circuit diagnostic system $\Psi_C = \langle f, \mathcal{H}, \mathcal{O}_C\rangle$, and a multiple-fault health assignment ω^I, does there exist a multiple-fault ATPG test?*

Theorem 1 *The Multiple-fault ATPG problem* MF-TEST *is Σ_2^P-complete.*

Proof: We prove this result using a reduction from a propositional abduction problem (PAP) $\mathcal{P}$ [16], for which the problem of solution existence (as defined below) has been shown to be Σ_2^P-complete [4]. A propositional abduction problem (PAP) can be defined using a tuple $\langle V, \Xi, \mu, \mathcal{T}\rangle$, where V is a set of variables of which Ξ and μ are disjoint subsets, while $\mathcal{T}$ is a (consistent) propositional formula. Ξ is typically referred to as the hypotheses, and μ as the manifestations. A solution, $\delta(\mathcal{T}, \mu)$, given manifestations μ, exists if $\Xi \cup \mathcal{T}$ is consistent and $\Xi \cup \mathcal{T} \models \mu$.

We can reduce a PAP into a generalised ATPG problem using the following procedure:

- for each hypothesis $\Xi_i \in \Xi$ create an observable variable $\mathcal{O}_i \in \mathcal{O}$;
- for each manifestation $\mu_i \in \mu$ create a component variable $h \in \mathcal{H}$;
- for each variable $v_i \in V \setminus (\Xi \cup \mu)$ create a variable x_i;
- for the propositional formula $\mathcal{T}$ create a Boolean function f, in which we have the variable correspondence as defined.

We assume that in both PAP and ATPG, a variable in (Ξ, μ), and its corresponding pair $(\mathcal{O}, h)$, has value 0 denoting "normal" and 1 denoting "abnormal" (for $\Xi, \mathcal{O}$) or "faulty" (μ, h).

Clearly this reduction can be performed in polynomial time. We now show that the PAP has a solution iff the multiple-fault ATPG problem is testable.

$\Rightarrow$: Since $\mathcal{T}$ is a (consistent) propositional formula, then it must be the case that $\Xi \cup \mathcal{T} \models \mu$ when $\mu = \{0, ..., 0\}$, together with $\Xi = \{0, ..., 0\}$, denote a "normal" state.

Assume a solution exists in the PAP with some μ that is abnormal; by our mapping, there exists a solution in MF-TEST such that some faulty h corresponding to μ has a consistent assignment. Consequently, we must have a testable multiple-fault setting in MF-TEST, since there exists an assignment of Boolean values to the primary inputs and output of C (and also C_ϕ) such that the output from components $h_i \in \mathcal{H}$ have complementary logic values in C and C_ϕ.

$\Leftarrow$: Assume that a testable multiple-fault setting exists in our ATPG problem. By our reduction, this means that there must be a solution to our PAP with abnormal setting for μ. □

5 Minimal Multiple-Fault ATPG

When we generate tests for multiple faults, there are two types of minimality that must be considered: (1) minimality of the multiple-fault; and (2) minimality of the size of the test set.

Multiple-fault minimality has largely been ignored within ATPG, since ATPG computes multiple simultaneous faults by isolating single faults in a sequential manner; however, this issue has received considerable attention in MBD, and special-purpose algorithms have been defined for minimal multiple-fault isolation, e.g., [14, 2]. Typically, given an anomalous observation in MBD, one wants to isolate the fault-set of minimal cardinality, since (1) it is most likely that fewer components have failed, and (2) this leads to replacing as few components as possible.

Computing a test set of minimal size is a key underlying ATPG task, since running the fewest tests leads to the most efficient (cheapest) fault isolation procedure. Test-set minimality has been largely ignored in MBD, which assumes arbitrary, rather than optimised, observations will be input to the inference engine. It has been shown that computing the minimum number $T^*(C)$ of tests to identify all *single stuck-at faults* in a circuit C is NP-hard [10]. Further, approximating this minimum test-set size is NP-hard as well [10], i.e., it is NP-hard to define some $\beta > 1$ and number of tests t such that

$T^*(C) \leq t \leq T^*(C) \cdot \beta$. As a consequence, we anticipate that the minimal test-set problem for multiple faults will be intractable, given the intractability of the simpler, single-fault problem.

In the following, we address the issue of generating a test-set for minimal faults which are not restricted to the single-fault case.

5.1 Tests for Minimal Multiple-Fault Diagnoses

Almost all real-world systems suffer from being under-sensed, i.e., having too few sensors to uniquely isolate every multiple-fault scenario, and from the model being over-constrained [15]. A consequence of the diagnostic system being over-constrained and under-sensed is that, for different α, there may be minimal-cardinality diagnoses of differing cardinalities. Multiple-fault ATPG is thus faced with the task of identifying the α leading to the max-cardinality diagnosis amongst the different α. This fault isolation problem has been defined as a Max/Fault Min/Cardinality (MFMC) problem [5]. Here, we show the complexity of generating tests for MFMC diagnoses.

The problem of computing a test for the Max-Fault Min-Cardinality (MFMC) problem [5] can be defined as follows.

Definition 14 (MFMC-Test) *Given a diagnostic system Ψ, a Max-Fault Min-Cardinality (MFMC) test (observation) is an instantiation α over the variables in $\mathcal{O}$ such that $MinCard(f \wedge \alpha)$ is maximized.*

5.1.1 Worst-Case Complexity

We will show that a restricted, decision version of the MFMC-Test problem, $MFMC_D$, is Π_2^P-complete, by showing a reduction from MINMAX-CLIQUE to $MFMC_D$. $MFMC_D$ frames the problem of finding a test α that isolates a cardinality-minimal diagnosis ω of cardinality at least κ.

Problem 2 ($MFMC_D$) *: Given a CNF boolean function f over a set X of variables, for which the variable set is partitioned into disjoint subsets denoting observables $\mathcal{O}$, health variables $\mathcal{H}$ and unobservable variables U, an integer κ, does there exist a test α and a health assignment ω such that $\min_{\omega \models f \wedge \alpha} |\omega| \geq \kappa$?*

Theorem 2 *$MFMC_D$ is Π_2^P-complete.*

Proof: We need to show two things: (1) that $MFMC_D$ is in Π_2^P; and (2) a reduction from a Π_2^P-complete problem.

1. It is easy to see that $MFMC_D$ is in Π_2^P, since it is solvable by a polynomial-time nondeterministic machine with the use of an NP-complete set of an oracle.
2. We now show a reduction from the Π_2^P-complete problem MINMAX-CLIQUE [9].
   MINMAX-CLIQUE is denoted by a graph $G = (V, E)$, a partition $(V_{i,j})_{i \in I, j \in J}$ of V, an integer κ, a function $t : I \rightarrow J$, and χ_t is the size of the largest clique in G restricted to $\bigcup_{i \in I}(V_{i,t(i)})$. In other words, we can denote $\chi_t(G) = \min_t \max_\chi\{|\chi| : \chi \subseteq V$ is a clique in $G_t\}$. G_t denotes the induced subgraph of G on the vertex set $V_i = \bigcup_{i=1}^{I} V_{i,t(i)}$.
   Given the partition, $|V| = I \cdot J = n$; hence $V = \{v_{1,1}, v_{1,2}, ..., v_{1,J}, v_{2,1}, ..., v_{2,J}, ..., v_{I,J}\}$. Partition i is given by $\{v_{i,1}, v_{i,2}, ..., v_{i,J}\}$, for $i = 1, ..., I$.
   We assume that in $MFMC_D$, there are I health variables, each of which has J possible (abnormal) health values, i.e., health variable k has values given by $\{h_{k,1}, h_{k,2}, ..., h_{k,J}\}$, $k = 1, ..., I$. Hence our space of possible health values is given by $\mathcal{H} = \{h_{1,1}, h_{1,2}, ..., h_{1,J}, h_{2,1}, ..., h_{2,J}, ..., h_{I,J}\}$.
   Let $f = C_1 \wedge C_2 \wedge \cdots \wedge C_m$ be a CNF formula with m clauses. We construct a CNF formula f from G as follows.

   Variables The variables in f are as follows:

   - For each vertex $v_{i,j}$, introduce a variable $h_{v_{i,j}}$ (the semantics is that $h_{v_{i,j}}$ is true if $v_{i,j}$ is in the clique, and the only literals occurring in a clique χ are health literals).
   - For each $i = 1, 2, ..., I$, $j = 1, 2, ..., J$ and $k = 0, 1, 2, ..., \kappa$, introduce a variable Y_{ijk}.

   The second set of variables ensures that Y_{ijk} is true if there are at least k true variables in the set $\{h_{v_{1,1}}, ..., h_{v_{1,J}}, h_{v_{2,1}}, ..., h_{v_{i,J}}\}$.

   Clauses The clauses in f are as follows:

   - For each pair of vertices $u, v \in G$ having no edge (u, v), add a clause $\neg h_v \vee \neg h_u$ to f.[3]
   - For each $i = 1, 2, ..., I$, $j = 1, 2, ..., J$ and $k = 1, 2, ..., \kappa$, add clause $Y_{i-1,j,k} \vee (Y_{i-1,j,k-1} \wedge h_{v_{i,j}})$.
   - For each $i = 1, 2, ..., I$, $j = 1, 2, ..., J$ add clauses enforcing $Y_{i,j,0}$ = true
   - For each $k = 1, ..., \kappa$, add clauses enforcing $Y_{0,0,k}$ = false.
   - Finally, add a clause containing just the variable $Y_{I,J,\kappa}$.

   The above reduction clearly can be computed in polynomial time. Next we verify that the reduction is correct.
   $\Rightarrow$: Suppose G has a clique χ_t of size κ. If this is true, then in f we must assign true to the variables h_v such that $v \in \chi_t$, and false to the other variables. We also assign true to those Y_{ijk}'s such that there are at least k true variables in $\{h_{v_1}, h_{v_2}, ..., h_{v_i}\}$, and we assign false to the other Y_{ijk}'s. This gives a satisfying assignment for f.
   Further, under the partition function $t : I \rightarrow J$, $\chi_t(G) = \min_t \max_\chi\{|\chi| : \chi \subseteq V$. This corresponds in our satisfying assignment in f to there existing $\min_{\omega \models f \wedge \alpha} |\omega| = \kappa$, since every node in $\chi_t(G)$ consists of a health literal from a different partition, i.e., it consists of a minimal diagnosis ω such that $|\omega| = \kappa$.
   $\Leftarrow$: Suppose f has a satisfying assignment. Let $\chi = \{v : h_v$ is assigned "true" $\}$. Then χ is a clique because for any pair of vertices $u, v \in \chi$, there cannot be a clause $\neg h_v \vee \neg h_u$ in f (because such a clause would not be satisfied), so there must be an edge $(u, v) \in G$. χ has size κ or more because $Y_{I,J,\kappa}$ must be true.
   Further, a satisfying assignment in f must correspond to there existing $\chi_t(G) = \min_t \max_\chi\{|\chi| : \chi \subseteq V$, since under the satisfying assignment we have $\min_{\omega \models f \wedge \alpha} |\omega| = \kappa$, and every node in $\chi_t(G)$ consists of a health literal from a different partition, i.e., it consists of a minimal diagnosis ω such that $|\omega| = \kappa$. □

5.1.2 Approximation Complexity

We can also prove the complexity of approximating this problem. As based on the approach presented in [9], we will use the method of proving the Π_2^P-completeness results for c-approximation problems MINMAX-A in terms of G-reductions. Ko and Lin [9] describe a G-reduction from MINMAX-SAT (Γ_{SAT}) to MINMAX-A, where MINMAX-A is an arbitrary optimisation problem. In particular, Ko and Lin constructed G-reductions from $\langle \Gamma_{SAT} : |f|; (1-\epsilon)|f| \rangle$ to a pair $\langle \Gamma_{MINMAX-A} : (1-\epsilon_2)size(x); (1-\epsilon_1)size(x) \rangle$, where $\epsilon_1 > \epsilon_2 \geq 0$. Based on this approach, they proved the following:

[3] At this point, the satisfying assignments to f correspond to cliques in G.

Theorem 3 *[9] There exists a constant $c > 1$ such that the c-approximation problem of* MINMAX-CLIQUE *is Π_2^P-complete.*

We can use this result to prove the following:

Theorem 4 ($MFMC_D$-approximation) *There exists a constant $c > 1$ such that the c-approximation problem of $MFMC_D$ is Π_2^P-complete.*

Proof: Let μ be the reduction of theorem 2. In the above proof, we showed that for any graph G, MINMAX-CLIQUE $= MFMC_D(\mu(G))$ For any graph G whose vertex set V is partitioned into $V_{i,j}$, $1 \leq i \leq I$, $1 \leq j \leq J$, we define $size(G) = I$. Similarly, for any CNF formula f whose variables are partitioned into (O, H, U) such that the partition for H follows $H_{i,j}$, $1 \leq i \leq I$, $1 \leq j \leq J$, we define $size(f) = |\omega|$.

Then the above observation implies that μ is a G-reduction from $\langle$ MINMAX-CLIQUE : $size(G); (1-\epsilon)size(G)\rangle$ to $\langle MFMC_D : size(f); (1-\epsilon)size(f)\rangle$. □

5.2 Test-Set Minimality

We now address the problem of computing a minimal test to identify a *multiple fault diagnosis* in a circuit.

Given a test set T, we define the size of T as $|T|$. A given test $\alpha \in T$ may identify a set of diagnoses as being consistent with α. If test $\alpha \in T$ is consistent with fault ω, then we assign $t = 1$; we assign $t = 0$ otherwise. A test set T unambiguously isolates ω if, when applied to $\omega \wedge f$, ω is the only fault consistent with T. A test set T is of minimal size for isolating a fault ω if T unambiguously isolates ω and there is no other test set T' such that $|T'| < |T|$ and T' unambiguously isolates ω.

We define the problem of computing a minimal test-set for a multiple-fault, MIN-MF-TEST, as follows:

Problem 3 (MIN-MF-TEST**)** *: Given a CNF boolean function f over a set X of variables, for which the variable set is partitioned into disjoint subsets denoting observables $\mathcal{O}$, health variables $\mathcal{H}$ and unobservable variables U, a test set T, and an integer κ, does there exist a test α and a health assignment ω such that $\max_{\alpha \models f \wedge \omega} |\alpha| \leq \kappa$?*

Theorem 5 MIN-MF-TEST *is Π_2^P-complete.*

Proof Sketch:[4] We prove this theorem analogously to the proof of Theorem 2, except that we reduce Maxmin-Vertex Cover (MMVC), which is Π_2^P-complete [9], to MIN-MF-TEST. We define MMVC as:

Problem 4 (Maxmin-Vertex Cover (MMVC**))** *: Given a graph $G = (V, E)$, a partition $(V_{i,j})_{i \in I, j \in J}$ of V, an integer κ, a function $t : I \rightarrow J$, and χ_t is the size of the smallest vertex cover in G restricted to $\bigcup_{i \in I}(V_{i,t(i)})$,is $\max_{t \in J^I} \chi_t(G) \leq \kappa$?*

We assume that in MMVC, there are I test variables $t_i \in T$, each of which has J possible (abnormal) values, i.e., test variable k has values given by $\{t_{k,1}, t_{k,2}, ..., t_{k,J}\}$, $k = 1, ..., I$. Hence our space of possible test values is given by $\mathcal{T} = \{t_{1,1}, t_{1,2}, ..., t_{1,J}, t_{2,1}, ..., t_{2,J}, ..., t_{I,J}\}$.

Let $f = C_1 \wedge C_2 \wedge \cdots \wedge C_m$ be a CNF formula with m clauses. We construct a CNF formula f from G similar to the construction of the proof of Theorem 2, except that we introduce test variables for health variables. Given the reduction from MMVC to MIN-MF-TEST, the theorem follows. □

[4] We sketch this proof due to space limitations.

6 Summary and Discussion

We have described a generalised framework for ATPG that allows this important technique to be applied to a wider range of tasks than the current single-fault tasks; such new tasks include multiple-fault inference for a range of embedded and other applications. In addition, we have described the complexity of these generalised ATPG problems, showing that these problems lie on the second level of the polynomial hierarchy.

This multiple-fault ATPG framework offers significant advances to test generation, both in the ability to isolate multiple faults, and in the ability to address systems more complex than digital circuits. The drawback is the intractability of these generalised ATPG problems. Since it is likely that inference will be intractable for real-world artifacts (as is true for MBD [13]), further work needs to be done to develop efficient, approximation algorithms. One promising MBD approach, that of stochastic approximation algorithms for multiple-fault test generation, has been explored in [5].

REFERENCES

[1] T. Bylander, D. Allemang, M.C. Tanner, and J. Josephson, 'The computational complexity of abduction', *Artificial Intelligence*, **49**, 25–60, (1991).

[2] J. de Kleer, 'An Assumption-based TMS', *AI Journal*, **28**, 127–162, (1986).

[3] R. Drechsler and G. Fey, 'Automatic Test Pattern Generation', *Formal Methods for Hardware Verification, LNCS*, 30–55, (2006).

[4] T. Eiter and G. Gottlob, 'The Complexity of Logic-Based Abduction', *J. ACM*, **42**(1), 3–42, (1995).

[5] A. Feldman, G. Provan, and A. van Gemund, 'Computing observation vectors for max-fault min-cardinality diagnoses', in *Proc. AAAI'08*, (July 2008).

[6] G. Friedrich, G. Gottlob, and W. Nejdl, 'Physical impossibility instead of fault models', in *Proc. AAAI*, (1990).

[7] M. Hermann and R. Pichler, 'Counting complexity of propositional abduction', in *IJCAI*, pp. 417–422, (2007).

[8] O.H. Ibarra and S.K. Sahni, 'Polynomially Complete Fault Detection Problems', *IEEE Transactions on Computers*, **C-24**(3), 242–249, (1975).

[9] K-I. Ko and C.L. Lin, 'On the Complexity of MIN-MAX Optimization Problems and Their Approximation', in *Minimax and Applications*, eds., Panos M. Pardalos and Dingzhu Du, Kluwer Academic Publishers, (1995).

[10] B. Krishnamurthy and B. Akers, 'Complexity of estimating the size of a test set', *IEEE Trans. Comp.*, **33**(8), 750–752, (1984).

[11] E J McCluskey, 'Built-in self-test techniques', *IEEE Design Test Comp.*, **2**(2), 21–28, (1985).

[12] M.R. Prasad, P. Chong, and K. Keutzer, 'Why is Combinational ATPG Efficiently Solvable for Practical VLSI Circuits?', *Journal of Electronic Testing*, **17**(6), 509–527, (2001).

[13] G.M. Provan, 'An Empirical Analysis of the Complexity of Model-Based Diagnosis', in *ECAI*, pp. 783–784, (2006).

[14] R. Reiter, 'A Theory of Diagnosis from First Principles', *Artificial Intelligence*, **32**, 57–96, (1987).

[15] W. R. Simpson and J. W. Sheppard, *System Test and Diagnosis*, Kluwer Academic Publishers, 1994.

[16] P. Torasso, L. Console, L. Portinale, and D. Theseider Dupre, 'On the role of abduction', *ACM Comput. Surv.*, **27**(3), 353–355, (1995).

[17] T. Ünlüyurt, 'Sequential Testing of Complex Systems: a Review', *Discrete Applied Mathematics*, **142**(1-3), 189–205, (2004).

[18] Z. Wang, M. Marek-Sadowska, K.H. Tsai, and J. Rajski, 'Multiple fault diagnosis using n-detection tests', *Computer Design, 2003. Proceedings. 21st International Conference on*, 198–201, (2003).

[19] Z. Wang, M. Marek-Sadowska, K.H. Tsai, and J. Rajski, 'Analysis and Methodology for Multiple-Fault Diagnosis', *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, **25**(3), 558–575, (2006).