UCL DEPARTMENT OF
COMPUTER SCIENCE

Research Note
RN/13/17

Learning Combinatorial Interaction Testing Strategies using Hyperheuristic Search

September 17, 2013

Yue Jia

Myra B. Cohen

Mark Harman

Justyna Petke

Abstract

Two decades of bespoke Combinatorial Interaction Testing (CIT) algorithm development have left software engineers with a bewildering choice of configurable system testing techniques. This paper introduces a single hyperheuristic algorithm that earns CIT strategies, providing a single generalist approach. We report experiments that show that our algorithm competes with known best solutions across constrained and unconstrained problems. For all 26 real world subjects and 29 of the 30 constrained benchmark problems studied, it equals or improves upon the best known result. We also present evidence that our algorithm's strong generic performance is caused by its effective unsupervised learning. Hyperheuristic search is thus a promising way to relocate CIT design intelligence from human to machine.

Learning Combinatorial Interaction Testing Strategies using Hyperheuristic Search

Yue Jia
University College London, UK

Myra B. Cohen
University of Nebraska-Lincoln, USA

Mark Harman
University College London, UK

Justyna Petke
University College London, UK

ABSTRACT

Two decades of bespoke Combinatorial Interaction Testing (CIT) algorithm development have left software engineers with a bewildering choice of configurable system testing techniques. This paper introduces a single hyperheuristic algorithm that learns CIT strategies, providing a single *generalist* approach. We report experiments that show that our algorithm competes with known best solutions across constrained and unconstrained problems. For all 26 real world subjects and 29 of the 30 constrained benchmark problems studied, it equals or improves upon the best known result. We also present evidence that our algorithm's strong generic performance is caused by its effective unsupervised learning. Hyperheuristic search is thus a promising way to relocate CIT design intelligence from human to machine.

Categories and Subject Descriptors

D.2.5 [**Software Engineering**]: Testing and Debugging

General Terms

Verification

Keywords

Software Testing, SBSE, Hyperheuristic Search

1. INTRODUCTION

Combinatorial Interaction Testing (CIT) is important because of the increasing reliance on configurable systems. CIT aims to generate samples that cover all possible value combinations between any set of t parameters, where t is fixed (usually between 2 and 6). Software product lines, operating systems, development environments and many other systems are typically governed by large configuration parameter and feature spaces for which CIT has proved useful [30].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ICSE 2014, India
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

Over two decades of research has gone into the development of bespoke CIT test generation techniques, each of which is tailored and tuned to a specific problem. For example, some CIT algorithms have been tuned and evaluated only on unconstrained problems [8, 12, 24, 27], while others have been specifically tuned for constrained interaction testing [4, 19], which prohibits certain configurations. Still other CIT approaches target specific problem structures, such as parameter spaces with few available parameter value choices [28,31], or are tuned to work on a particular set of real world problems [36].

The tester is therefore presented with many different techniques and implementations from which to choose, each of which has its own special properties. Unfortunately, this choice can be a bewildering one, because an algorithm designed for one CIT instance may perform poorly when applied to another (or may even be inapplicable). It may not be reasonable to expect a practicing software tester to perform their own experiments to decide on the best CIT algorithm choice for each and every testing problem. It would be equally unreasonable to expect every organisation to hire an algorithm designer to build bespoke CIT testing implementations for each testing scenario the organisation faces.

To overcome some of the practical difficulties in choosing among available implementations, there has been an attempt to collate algorithms into a common framework with a single unified interface [6]. This framework seeks to ensure that all implementations ought to be at least *easily applicable*. However, it cannot help the tester to choose *which algorithm* to apply to each CIT problem instance. Instead, he or she must run experiments to gather this data.

To address this problem we introduce a simulated annealing hyperheuristic search based algorithm. Hyperheuristics are a new class of Search Based Software Engineering algorithms, the members of which use dynamic adaptive optimisation to learn optimisation strategies without supervision [3, 22]. Our hyperheuristic algorithm learns the CIT strategy to apply dynamically, as it is executed. This single algorithm can be applied to a wide range of CIT problem instances, regardless of their structure and characteristics.

For our new algorithm to be acceptable as a generic solution to the CIT problem, we need to demonstrate that it is effective and efficient across a wide range of CIT problem instances, when compared to other possible algorithm choices. To assess the effectiveness of CIT solutions we use the size of the final sample.

Garvin et al. [19] have demonstrated that this size has the greatest impact on the overall efficacy of CIT. To assess efficiency we report computational time (as is standard in CIT experiments), but we also deploy the algorithms in the cloud. Cloud deployment provides an unequivocal supplemental assessment of monetary cost (as has been done in recent software engineering studies [25]).

We compare our hyperheuristic algorithm, not only against results from other state-of-the-art CIT techniques, but also against the best known results in the literature, garnered over 20 years of analysis of CIT. This is a particularly challenging 'quality comparison' for any algorithm, because some of these best known results are the product of many years of careful analysis by mathematicians, not machines.

We show that our hyperheuristic algorithm performs well on both constrained and unconstrained problems and across a wide range of parameter sizes and data sets that includes both two and three-way coverage. Like the best known results, some of these data sets have been designed using human ingenuity. Human design ensures that these benchmarks capture especially pathological 'corner cases' and problems with specific structures that are known to pose challenges to the CIT algorithms.

Overall, our results provide evidence to support the claim that hyperheuristic search is a promising solution to CIT, potentially replacing the current bewildering range of choices with a single, generic solution that learns to tailor itself to the configuration testing problem to which it is applied.

The primary contributions of this paper are:

1. The formulation of CIT as a hyperheuristic search problem and the introduction of the first hyperheuristic algorithm for solving this problem. This work is also the first use of hyperheuristic learning in the Software Engineering literature.

2. A comprehensive empirical study showing that this approach is both effective and efficient. The study reports results across a wide range of 59 previously studied benchmarks that include known pathological and corner cases. We also study 26 problem instances from two previous studies where each of the 26 CIT problems is drawn from a real world configurable system testing problem. These study subjects cover both constrained and unconstrained CIT problems, and we report results for both 2-way and 3-way interaction coverage for a subset of these.

3. We use the Amazon EC2 cloud to measure the real computational cost (in US dollars) of the algorithms studied. These results indicate that, with default settings, our hyperheuristic algorithm can produce competitive results to state-of-the-art tools at a reasonable cost. For example, our algorithm produces all the pairwise interaction tests reported in the paper for all 26 real world problems and the 44 pairwise benchmarks for a total cost of only $2.09.

4. A further empirical study is used to explore the nature of online learning employed by our algorithm. The results of this study show that the hyperheuristic search productively combines heuristic operators that would have proved to be unproductive in isolation. Our results also demonstrate how the hyperheuristic adapts its choice of operators to the specific problem to which it is applied.

2. PRELIMINARIES

In this section we will give a quick overview of the notation used throughout the paper. CIT produces a Covering Array (CA), which is typically represented as follows in the literature:

$$CA(N; t, v_1^{k_1} v_2^{k_2} ... v_m^{k_m})$$

where N is the size of the array, the sum of $k_1, ..., k_m$ is the number of parameters (or factors), each v_i stands for the number of values for each of the k_i parameters in turn and t is the strength of the array; a t-way interaction test suite aims to cover all possible t-way combinations of values between any t parameters.

Suppose we want to generate a pairwise (aka 2-way) interaction test suite for an instance with 3 parameters, where the first and third parameter can take 4 different values and the second one can only take 3 different values. Then the problem can be formulated as: $CA(N; 2, 4^1 3^1 4^1)$ and the model of the problem is $4^1 3^1 4^1$.

Furthermore, in order to test all combinations one would need $4 * 3 * 4 = 48$ test cases, pairwise coverage reduces this number to 16. Suppose that we have the following constraints: only the first value for the first parameter can ever be combined with values for the other parameters, and the last value for the second parameter can never be combined with values for all the other parameters. Introducing these constraints reduces the size of the test suite further; only 8 test cases are now required to cover all feasible interactions. Since constraints reduce test suite size and naturally occur in real-world problems, constrained CIT is well-fitted for industrial applications [33].

Many different algorithms have been introduced to generate covering arrays. Each of these algorithms is customised for specific problem instances. For example, there have been greedy algorithms, such as the Automatic Efficient Test Case Generator, AETG [8] , the 'In Parameter Order' algorithm (IPO) [26] and PICT [15]. These methods either generate a new test case on-the-fly, seeking to cover the largest number of uncovered t-way interactions, or start with a small number of parameters and iteratively add new columns and rows to fill in the missing coverage.

Other approaches include metaheuristic search algorithms, such as simulated annealing [12, 19, 28] or tabu search [31]. These metaheuristics are usually divided into two phases or stages. In the first stage, binary search, for instance, is used to generate a random test suite, r of fixed size n. In the second stage, metaheuristic search is used to search for a test suite of size n, starting with r, that covers as many interactions as possible. And there are other unique algorithms, such as those that use constraint solving or logic techniques as the core of their approach [6, 24].

3. HYPERHEURISTIC CIT ALGORITHM

Hyperheuristic search is a new class of optimisation algorithms for Search Based Software Engineering [21]. Hyperheuristics have been successfully applied to many operational research problems outside of software engineering [3]. However, though they have been advocated as a possible solution to dynamic adaptive optimisation for software engineering [22], they have not, hitherto, been applied to any software engineering problem [1, 23, 34]. There are two subclasses of hyperheuristic algorithms: generative and selective. Generative hyperheuristics combine low level heuristics to generate new higher level heuristics.

Selective hyperheuristics select from a set of low level heuristics. In this paper we use a selective hyperheuristic algorithm. Selective hyperheuristic algorithms can be further divided into two classes, depending upon whether they are online or offline. We use online selective hyperheuristics, since we want a solution that can learn to select the best CIT heuristic to apply, unsupervised, as it executes.

The hyperheuristic algorithm takes the set of lower level heuristics as input, and layers the heuristic search into two levels that work together to produce the overall solution. The first (or outer) layer uses a normal metaheuristic search to find solutions directly from the solution space of the problem. The inner layer heuristic, searches for the best candidate lower heuristics for the outer layer heuristics in the current problem state. As a result, the inner search adaptively identifies and exploits different strategies according to the characteristics of the problems it faces.

Our algorithm uses Simulated Annealing (SA) as the outer search. We choose SA because it has been successfully applied to CIT problems, yet, even within this class of algorithms, there is a wide choice of available approaches and implementations [11,12,19,20,28]. We use a reinforcement learning agent to perform the inner layer selection to heuristics. Our overall Hyper Heuristic Simulated Annealing algorithm (HHSA) is depicted in Figure 1 and set out more formally as Algorithm 1.

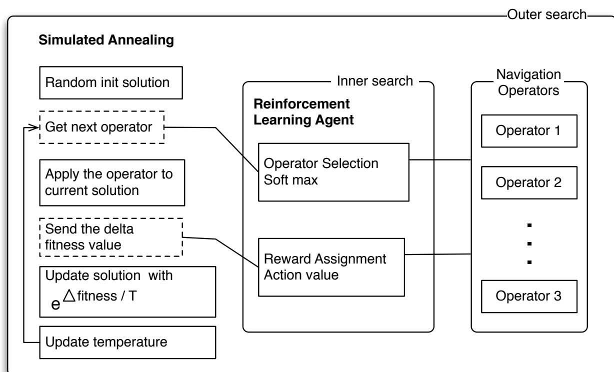

Figure 1: Hyper Heuristic Simulated Annealing

3.1 The Outer Layer: Simulated Annealing

A standard Simulated Annealing (SA) algorithm is used as the outer layer search. The SA algorithm starts with a randomly generated Nxk array as an initial solution. The fitness value for each solution is the number of uncovered t-tuples present in the current array. The fitness value is also used to represent the current state of the problem (i.e. how many tuples remain to be covered). This 'problem state' is used to understand how our algorithm learns throughout different stages of the problem.

In each annealing iteration, the SA algorithm asks the reinforcement learning agent to choose a best operator for the current problem state. It then applies the operator and passes the change in fitness value (delta fitness) back to the agent. The SA algorithm accepts the new solution if its fitness is the same or better than the fitness of the previous solution. Otherwise it uses a probability, $e^{\Delta fitness/T}$, for accepting the current solution based on the current temperature T.

As the SA algorithm proceeds, the temperature, T, is progressively decreased according to a cooling schedule. Decreasing the temperature reduces the probability of SA accepting a move that reduces fitness. The SA algorithm stops when the current array covers all t-tuples or it reaches a preset maximum number of non-improving moves.

To incorporate constraints between parameters, the outer SA first preprocesses the constraints and identifies all invalid tuples which must not be covered. Since previous work [11,19,20] used a SAT-solver, MiniSAT, for constraint solving, we also use it in our implementation. Other constraint solvers could be used, but we wish to be able to compare effectiveness to these existing state-of-the-art CIT systems and the best results reported for them.

The outer SA checks constraint violations after applying each operator and proposes a repair if there are any violations. The constraint fixing algorithm is a simple greedy approach that checks each row of covering array, one at a time. The algorithm is set out formally as Algorithm 2. If the outer SA fails to fix the array, it reapplies the current heuristic operator to generate a new solution.

```
Input  : t, k, v, c, N, MaxNoImprovment
Output: covering array A
A ← initial_array (t,k,v,N)
no_improvement ← 0
curr_missing ← countMissingTuples ( A )
while curr_missing ≠ 0 and MaxNoImprovment ≠ no_improvement do
    op ← rl_agent_choose_action ( curr_missing )
    A' = local_move (op, A )
    while fix_cons_violation (A', c ) do
        A' = local_move (op, A )
    end
    new_missing ← countMissingTuples (A' )
    Δfitness = curr_missing − new_missing
    rl_agent_set_reward ( op, Δfitness )
    if e^(Δfitness/Temp) > rand_0_to_1 () then
        if Δfitness = 0 then
            no_improvement ← no_improvement + 1
        end
        A ← A'
        curr_missing ← new_missing
    end
    Temp ← cool ( Temp )
end
```

Algorithm 1: HHSA

We enclose the SA in a binary search procedure to determine the array size N. This outer binary search procedure is a commonly used solution to iteratively direct an algorithm to CIT problems for different values of N, until a smallest covering array of size N can be found [12]. The outer binary search takes an upper and lower bound on the size of array as input, and returns the covering array with the smallest possible size.

The CASA tool for CIT [19,20] uses a more sophisticated version of the binary search. It first tries the same size multiple times and then does a greedy one sided narrowing to improve the final array size. Our implementation also performs this 'CASA-style' greedy approach to finding the array size, but the use of this approach is tunable (i.e. we use this when we want to search 'harder').

```
Input   : constraints c, row R
Output: has_violation
has_violation ← False
fix_time ← 0
recheck: foreach clause in c do
    foreach term in clause do
        if R has term then
            has_violation ← True
        else
            has_violation ← False
            break
        end
    end
    if has_violation then
        if fix_time = MaxFixTime then
            break
        end
        term =clause_get_random_term (clause )
        R = random_fix_term (R, term )
        fix_time = fix_time +1
        go to recheck
    end
end
```

Algorithm 2: Constraint Violation Fixing

3.2 The Reinforcement Learning Agent

The goal of the inner layer is to select the best operator at the current problem state. This operator selection problem can be considered an n-Armed Bandit Problem, in which the n arms are the n available heuristics and the machine learner needs to determine which of these heuristics offers the best reward at each problem state. We designed a Reinforcement Learning (RL) agent for the the inner search, as RL agents are known to provide generally good solutions to this kind of so-called 'Bandit Problem' [38].

As shown in Figure 1, the RL agent takes a set of operators as input. In each annealing iteration, the RL agent repeatedly chooses the best fit operator, a, based on the expected reward of applying operators at current state of the problem.

After applying the operator a, the RL agent receives a reward value from the outer layer SA algorithm, based on performance. At the end of the iteration, the RL agent updates the expected reward for the chosen operator with the reward value returned.

The goal of the RL agent is to maximise the expected total reward that accrues over the entire running time of the algorithm. Because the reward returned by SA is the improvement of SA's fitness value, the RL agent will thus 'learn' to choose the operators that tend to maximise the SA's fitness improvement, adapting to changes in problem characteristics.

Our RL agent uses an action-value method [38] to estimate the expected rewards for each of the operators available to it at a given problem state. That is, given a set of operators $A = \{a_1, a_2, \ldots, a_i\}$, let $R_i = \{r_{i1}, r_{i2}, \ldots, r_{ik}\}$, be the returned reward values of operator a_i at the k^{th} iteration at which a_i is applied.

The estimated reward a_i is defined as the mean reward value, $\frac{r_{i1}+r_{i2}+\cdots+r_{ik}}{k_a}$, received from SA. We balance the twin learning objectives of exploration and exploitation, the RL agent uses a SOFTMAX selection rule [38].

The SOFTMAX selection rule is a greedy approach that gives the operator with the best estimated reward the highest selection probability. For each operator a_i, the selection probability is defined based on the Gibbs distribution: $\frac{e^{R_{a_i}/T}}{\sum_{j=1}^{n} e^{R_{a_j}/T}}$. A higher value of temperature T makes the selection of all operators more equal while a lower value makes a greater difference in selection probability.

3.3 Search Space Navigation Operators

We have selected a set of six operators to investigate the performance and feasibility of this approach to adaptive learning for CIT. Like any general process, we choose operators that can be widely applicable and which the learner might be able to combine in productive ways. Since we must be general, we cannot exploit specific problem characteristics, leaving it to the learner to find ways to do this through the smart combination of the low level heuristics we define.

We have based our operator selection on the previous algorithms for CIT. None of the operators considers constraints directly, but some have been used for constrained and some for unconstrained problems. Like other machine learning approaches we need a combination of 'smart' heuristic and 'standard' heuristics, since each can act as an enabler for the other. The first three operators are ones we deem to be entirely *standard*; they do not require book keeping or search for particular properties before application. The second set are ones that we deem to be somewhat *smart*; these use information that one might expect could potentially help guide the outer search. The operators are as follows:

1. Single Mutation (Std): Randomly select a row r and a column c, change the value at r, c to a random valid value. This operator matches the neighbourhood transformation in the unconstrained simulated annealing algorithm [12].
2. Add/Del: (Std): Randomly delete a row r and add a new row r' randomly generated. While CASA includes a row replacement operator, that one uses intelligence (i.e. the row is not just randomly generated).
3. Multiple Mutation (Std): Randomly select two rows, r_1 and r_2, and crossover each column of r_1 and r_2 with a probability of 0.2.
4. Single Mutation (Smart): Randomly select a missing tuple, m, which is the combination of columns $c_1, \ldots, c_n$. Go through each row in the covering array, if there exists a duplicated tuple constructed by the same combination of columns $c_1, \ldots, c_n$, find a row containing the duplication randomly and change the row to cover the missing tuple m. Otherwise randomly select a row r and change the row to cover the missing tuple m.
5. Add/Del: (Smart): Randomly delete a row r, and add a new row r' to cover n missing tuples. We define n as the smaller value from $k/2$ (where k is the number of parameters) and the number of missing tuples. This is a simple form of constructing a new row used by AETG [8].
6. Multiple Mutation (Smart): Randomly select two rows, r_1 and r_2, and compare the frequency of a value at each column, f_{c1} and f_{c2}. With probability of 0.2, the column with high frequency will be mutated to a random value.

4. EXPERIMENTS

In order to assess the usefulness of using our hyperheuristic algorithm as a general approach to CIT, we built our version of the hyperheuristic simulated annealing algorithm and posed the following research questions:[1]

RQ1 What is the *quality* of the test suites generated using the hyperheuristic approach?

One of the primary goals of CIT is to find the smallest test suite (defined by the covering array) that achieves the desired strength coverage. It is trivial to generate an arbitrarily large covering test suite - simply include one test case for each interaction to be covered. However, such a naïve approach to test generation would yield exponentially many test cases. All CIT approaches therefore work around problem of finding a minimal size covering array for testing. The goal of CIT is to try to find the smallest test suite that achieves 100% t-way interaction coverage for some chosen strength of interaction t. In our experiment, we compare the size of the test suites generated by the Hyperheuristic Simulated Annealing Algorithm (HSSA) in three different ways. We compare against the

1. Best known results reported in the literature, produced by any approach, including analysis and construction by mathematicians.
2. Best known (i.e. published) results produced by automated tools.
3. A state-of-the-art SA-based tool that was designed to run on unconstrained problems and a state-of-the-art SA-based tool that was designed to handle constrained problems well, CASA.

RQ2 How *efficient* is the hyperheuristic approach and what is the trade off between the quality of the results and the running time?

Another important issue in CIT is the time to find a test suite that is as close to the minimal one as possible given time budgeted for the search. Depending on the application, one might want to sacrifice minimality for efficiency (and vice-versa). This research question therefore investigates whether the hyperheuristic approach can generate small test suites in reasonable time.

If the answers to the first two research questions are favourable to our hyperheuristic algorithm, then we will have evidence that it can be useful. However, usefulness on our set of problems, wide and varied though it is, may not be sufficient for our algorithm to be actually used. We seek to further explore whether its value is merely an artefact of the operators we chose for low level heuristics. We also want to check whether the algorithm is really 'learning'. If not, then it might prove to be insufficiently adaptive to changing problem characteristics. The next two research questions seek to test our algorithms further, by investigating these questions.

RQ3 How efficient and effective is each search navigation operator *in isolation*?

In order to collect baseline results for each of the operators that the hyperheuristic approach can choose, we study the effects of each operator in isolation. That is, we ask how well each operator can perform on its own.

[1]Supplementary data, models and results, can be found on our website (http://cse.unl.edu/~myra/artifacts/HHSA).

Should it turn out that there is a single operator that performs very well, then there would be no need for further study; we could simply use the high performing operator in isolation. Similarly, should one operator prove to perform poorly and to be expensive then we might consider removing it from further study.

RQ4 Do we see evidence that the hyperheuristic approach is learning?

Should it turn out that the hyperheuristic approach performs well, finding competitively sized covering arrays in reasonable time, then we have evidence to suggest that the adaptive learning used by the hyperheuristic approach is able to learn which operator to deploy. However, is it *really learning*? This RQ investigates, in more detail, the nature of the learning taking place as the algorithm searches for interaction test suites. We explore how the problem difficulty varies over time for each of the CIT problems we study, and then ask which operators are chosen at each stage of difficulty; is there evidence that the algorithm is selecting different operators for different types of problems?

4.1 Experimental Setup

In this section we present the experiments conducted.

Subjects Studied. There are five subject sets used in our experiments. The details are summarised below:

[Syn-2] contains 14 different pairwise (2-way) synthetic models without constraints. These are shown in the leftmost column of Table 1. These models are benchmarks that have been used both to compare mathematical constructions as well as search based techniques [12, 19, 27, 37, 39]. We take these from Table 7 from the paper by Garvin et al. [19].

[Syn-3] contains 15 different 3-way synthetic models without constraints. These are shown in the second column of Table 1. These models are benchmarks that have been used for mathematical constructions and search [7, 10, 12]. We take these from Table 7 from the paper by Garvin et al. [19].

[Syn-C2] contains 30 different 2-way synthetic models with constraints (see Table 1, rightmost two columns). These models were designed to simulate configurations with constraints in real world programs, generated by Cohen et al. [9] and adopted in follow-up research by Garvin et al. [19, 20].

[Real-1] contains real world models from a recent benchmark created by Segall et al. [36], shown in Table 2. There are 20 CIT problems in this subject set, generated by or for IBM customers. The 20 problems cover a wide range of applications, including telecommunications, healthcare, storage and banking systems.

[Real-2] contains 6 real world constrained subjects shown in Table 2, which have been widely studied in the literature [9, 11, 19, 20, 35]. The TCAS model was first presented by Kuhn et al. [35]. TCAS is a traffic collision avoidance system from the 'Siemens' suite [16]. The rest of the models in this subject set were introduced by Cohen et al. [9, 11]. SPIN-S and SPIN-V are two components for model simulation and model verification. GCC is a well known compiler system from the GNU Project. Apache is a web server application and Bugzilla is a web-based bug tracking system.

Methodology: All experiments but one are carried out on a desktop computer with a 6 core 3.2GHz Intel CPU and 8GB memory. To answer the last part of RQ2 we used the Amazon EC2 Cloud. All experiments are repeated five times and the results are averaged over these five runs.

Table 1: Synthetic Subjects Syn-2, Syn-3 and Syn-C2. The first subject set contains 2-way synthetic models without constraints from [12, 19, 27, 37, 39]. The second subject set contains 3-way synthetic models without constraints from [7, 10, 12]. The last subject set contains synthetic models designed to simulate real world programs used in [9, 19, 20].

Subject Set: Syn-2		Subject Set: Syn-3		Subject Set: Syn-C2			Subject Set: Syn-C2		
Subjects	Model	Subjects	Model	Subjects	Unconstr. Param.	Constr. Param.	Subjects	Unconstr. Param.	Constr. Param.
S2-1	3^4	S3-1	3^6	C2-S1	$2^{86}3^34^15^56^2$	$2^{20}3^34^1$	C2-S16	$2^{81}3^34^26^1$	$2^{30}3^4$
S2-3	$5^13^82^2$	S3-2	4^6	C2-S2	$2_{86}3^34^35^16^1$	$2^{19}3^3$	C2-S17	$2^{128}3^34^25^16^3$	$22^{25}3^4$
S2-3	3^{13}	S3-3	$3^24^25^2$	C2-S3	$2^{27}4^2$	2^93^1	C2-S18	$2^{127}3^24^25^16^3$	$2^{23}3^44^1$
S2-4	$4^13^{39}2^{35}$	S3-4	5^6	C2-S4	$2^{51}3^44^25^1$	$2^{15}3^2$	C2-S19	$2^{172}3^94^95^36^4$	$2^{38}3^5$
S2-5	$5^14^43^{11}2^5$	S3-5	5^7	C2-S5	$2^{155}3^74^35^56^4$	$2^{32}3^64^1$	C2-S20	$2^{138}3^44^55^46^7$	$2^{42}3^6$
S2-6	$4^{15}3^{17}2^{29}$	S3-6	6^6	C2-S6	$2^{73}4^36^1$	$2^{26}3^4$	C2-S21	$2^{76}3^34^25^16^3$	$2^{40}3^6$
S2-7	$6^15^14^63^82^3$	S3-7	$6^64^22^2$	C2-S7	$2^{29}3^1$	$2^{13}3^2$	C2-S22	$2^{73}3^34^3$	$2^{31}3^4$
S2-8	$7^16^15^14^53^82^3$	S3-8	$10^16^24^33^1$	C2-S8	$2^{109}3^24^25^36^3$	$2^{32}3^44^1$	C2-S23	$2^{25}3^16^1$	$2^{13}3^2$
S2-9	4^{100}	S3-9	8^8	C2-S9	$2^{57}3^14^15^16^1$	$2^{30}3^7$	C2-S24	$2^{110}3^25^36^4$	$2^{25}3^4$
S2-10	6^{16}	S3-10	7^7	C2-S10	$2^{130}3^64^55^26^4$	$2^{40}3^7$	C2-S25	$2^{118}3^64^25^26^6$	$2^{23}3^34^1$
S2-11	7^{16}	S3-11	9^9	C2-S11	$2^{84}3^44^25^26^4$	$2^{28}3^4$	C2-S26	$2^{87}3^14^35^4$	$2^{28}3^4$
S2-12	8^{16}	S3-12	10^6	C2-S12	$2^{136}3^44^35^16^3$	$2^{23}3^4$	C2-S27	$2^{55}3^24^25^16^2$	$2^{17}3^3$
S2-13	8^{17}	S3-13	10^{10}	C2-S13	$2^{124}3^44^15^26^2$	$2^{22}3^4$	C2-S28	$2^{167}3^{16}4^25^36^6$	$2^{31}3^6$
S2-14	10^{20}	S3-14	12^{12}	C2-S14	$2^{81}3^54^36^3$	$2^{13}3^2$	C2-S29	$2^{134}3^75^3$	$2^{19}3^3$
		S3-15	14^{14}	C2-S15	$2^{50}3^44^15^26^1$	$2^{20}3^2$	C2-S30	$2^{72}3^44^16^2$	$2^{20}3^2$

Table 2: Real World Subject Sets. Real-1 (top) contains 20 models from [36]. Real-2 (bottom) contains 6 models with constraints from [9, 11, 19, 20, 35].

Subjects	Unconstrained Parameters	Constrained Param.
Real-1: 2-way		
Concurrency	2^5	$2^43^15^2$
Storage1	$2^13^14^15^1$	4^{95}
Banking1	3^44^1	5^{112}
Storage2	3^46^1	-
CommProtocol	$2^{10}7^1$	$2^{10}3^{10}4^{12}5^{96}$
SystemMgmt	$2^53^45^1$	$2^{13}3^4$
Healthcare1	$2^63^25^16^1$	2^33^{18}
Telecom	$2^53^14^25^16^1$	$2^{11}3^14^9$
Banking2	$2^{14}4^1$	2^3
Healthcare2	$2^53^64^1$	$2^13^65^{18}$
NetworkMgmt	$2^24^15^310^211^1$	2^{20}
Storage3	$2^93^15^36^18^1$	$2^{38}3^{10}$
Proc.Comm1	$2^33^64^6$	2^{13}
Services	$2^33^45^28^210^2$	$3^{386}4^2$
Insurance	$2^63^15^16^211^113^117^131^1$	-
Storage4	$2^53^74^15^26^27^19^113^1$	2^{24}
Healthcare3	$2^{16}3^64^55^16^1$	2^{31}
Proc.Comm2	$2^33^{12}4^85^2$	1^42^{121}
Storage5	$2^53^85^36^28^19^110^211^1$	2^{151}
Healthcare4	$2^{13}3^{12}4^65^26^17^1$	2^{22}
Real-2: 2,3-way		
TCAS	$2^73^24^110^2$	2^3
Spin-S	$2^{13}4^5$	2^{13}
Spin-V	$2^{42}3^24^{11}$	$2^{47}3^2$
GCC	$2^{189}3^{10}$	$2^{37}3^3$
Apache	$2^{158}3^84^45^16^1$	$2^33^14^25^1$
Bugzilla	$2^{49}3^14^2$	2^43^1

Table 3: Settings for the HHSA-L, HHSA-M and HHSA-H configurations.

Config.	Search	InitT	Co-Rate	Co-Step	MaxNo-Imp
HHSA-L	binary	0.3	0.98	2,000	50,000
HHSA-M	binary	0.3	0.998	10,000	50,000
HHSA-H	binary	0.3	0.998	10,000	50,000
	greedy	0.5	0.9998	10,000	100,000

4.2 HHSA Configuration

There are four parameters that impact the computational resources used by our hyperheuristic algorithm, HHSA: the initial temperature, the cooling rate, the cooling step function, and maximum number of non-improvements allowed before termination is forced. A higher initial temperature allows HHSA to spend more effort in exploring the search space. The cooling rate and cooling step function work together to control the cooling schedule for HHSA.

To understand the trade-off between the quality of the results and the efficiency of the hyperheuristic algorithm, we use three different configurations for our hyperheuristic algorithm: HHSA-L (Low), HHSA-M (Medium) and HHSA-H (High). The HHSA-L and HHSA-M configurations only apply the outer binary search to guide HHSA to search for the smallest test suite while the HHSA-H configuration additionally applies the greedy search conducted after the binary search. The settings for these three configurations are shown in Table 3.

We chose these setting after some experimentation because all can be executed in reasonable time for one or more use-cases for CIT. In the low setting, the time taken is low, but the expected result quality is consequently equally low, whereas in the higher settings, we can explore what additional benefits are gained from the allocation of extra computational resource.

5. RESULTS

In this section we provide results aimed at answering each of our research questions.

5.1 RQ1: Quality of Hyperheuristic Search

We begin by looking at the set of unconstrained synthetic problems (Table 4) for 2- (top) and 3-way (bottom) CIT. In this table, we include the best reported solution from the literature followed by the smallest CIT sample and its running time for each of the three settings of the HHSA. The best column follows the format of Table 7 from Garvin et al. [19] and includes results obtained by mathematical or constructive methods as well as search. We also include the size reported in that paper both for the unconstrained SA and CASA tool, which is optimized for constrained problems.

Table 4: Sizes and times (seconds) for Syn-2 (top) and Syn-3 (bottom). The Best column reports the best known results as given in [19]. The SA and CASA columns report the size of the unconstrained simulated annealing algorithm and the CASA algorithm. The Diff-Best column indicates the difference between the smallest HHSA variant and the Best column.

Subject	Best	SA	CASA	HHSA-L		HHSA-M		HHSA-H		Diff-Best	Diff-SA	Diff-CASA
				Size	Time	Size	Time	Size	Time			
S2-1	9	9	9	9	1	9	12	9	44	0	0	0
S2-2	15	15	15	15	1	15	14	15	120	0	0	0
S2-3	15	15	15	15	1	15	14	15	101	0	0	0
S2-4	21	21	22	22	6	21	92	21	1,086	0	0	-1
S2-5	21	21	23	22	1	22	21	21	241	0	0	-2
S2-6	30	30	30	31	4	29	212	29	961	-1	-1	-1
S2-7	30	30	30	30	1	30	41	30	177	0	0	0
S2-8	42	42	46	42	1	42	22	42	175	0	0	-4
S2-9	45	45	46	47	41	46	259	45	2,647	0	0	-1
S2-10	62	62	64	66	2	64	31	63	293	1	1	-1
S2-11	84	87	86	88	3	87	43	86	315	2	-1	0
S2-12	110	112	112	115	6	112	54	111	581	1	-1	-1
S2-13	111	114	114	117	7	115	62	113	644	2	-1	-1
S2-14	162	183	185	195	15	194	98	189	1,201	27	6	4
Overall	757	786	797	814	90	801	975	789	8,586	32	3	-8
S3-1	33	33	33	33	0	33	2	33	5	0	0	0
S2-2	64	64	96	64	0	64	1	64	1	0	0	-32
S3-3	100	100	100	101	1	100	31	100	153	0	0	0
S3-4	125	152	185	176	2	161	21	125	78	0	-27	-60
S3-5	180	201	213	211	3	205	40	202	473	22	1	-11
S3-6	258	300	318	316	4	315	56	308	875	50	8	10
S3-7	272	317	383	345	11	329	123	319	1,893	47	2	-64
S3-8	360	360	360	360	6	360	138	360	498	0	0	0
S3-9	512	918	942	958	39	1,000	187	994	6,966	446	40	16
S3-10	545	552	573	595	14	595	99	575	2,309	30	23	2
S3-11	729	1,426	1,422	1,520	112	1,637	351	1,600	7,206	791	94	98
S3-12	1,100	1,426	1,462	1,440	44	1,530	329	1,496	10,921	340	14	-22
S3-13	1,219	2,163	2,175	2,190	231	2,440	543	2,453	11,138	971	27	15
S3-14	2,190	4,422	4,262	4,760	831	5,080	1,634	5,080	17,679	2,570	338	498
S3-15	3,654	8,092	8,103	9,195	3,684	9,040	5,748	9,039	30,611	5,385	947	936
Overall	11,341	20,526	20,627	22,264	4,982	22,889	9,303	22,748	90,807	10,652	1467	1366

The size and time columns give the smallest size of the CIT sample found by HHSA, and the average running time in seconds over five runs. The Diff-Best column reports the difference between the best known results (first column) and HHSA's best results. We have also reported HHSA vs. SA (Diff-SA) and HHSA vs. CASA (Diff-CASA). A negative value indicates that HHSA found a smaller sample.

The sizes of test suites found by HHSA are very close to the benchmarks for all but one of the 2-way unconstrained synthetic models. In fact, in benchmark S2-6, both the medium and high settings of HHSA find a lower bound. The last subject, S2-14 is interesting because it is pathological and has been studied extensively by mathematicians. The model 10^{20}, has 20 parameters, each with 10 values. The use of customizations for this particular problem, such as symmetry has led to both constructions and post-optimizations.

The discussion of this single model consumes more than half a page in a recent dissertation which is credited with the bound[2] of 162 [29]. The best simulated annealing bound, of 183, is close to the high setting of HHSA (189).

There is a larger gap between the results generated by HHSA and best known results on 3-way synthetic models. On the smaller models, HHSA seems to generate sample sizes between the unconstrained SA technique and CASA. However, on the larger size models HHSA does not fare as well, but we do see improvement as we increase from low to high settings and these are all very large search spaces; we explore this cost-effectiveness trade-off for HHSA in RQ2.

[2]This bound was recently reduced by others to 155.

We now turn to the constrained synthetic models seen in Table 5. In this table the column labelled 'Best' represents the best known results for CASA (the only tool on which these synthetic benchmarks have been reported to date). For the constrained problems HHSA performs as well or better than the best known results (except in one case) despite the fact that CASA is optimized for these subjects. HHSA requires 39 fewer rows overall than the best reported results.

The last comparison we make is with the Real benchmarks. Table 6 shows a comparison for all of our real subjects against a set of existing tools which were reported in the literature (references provided in the table). Again we see that the HHSA algorithm performs as well or better than all of the other tools. For the IBM benchmarks HHSA reduces the overall number of rows in our samples by 52, and for the open source applications HHSA reduces the 2-way by 3 rows, and the 3-way by 54 rows.

Summary of RQ1. *We conclude that the quality of results obtained by using HHSA is high. While we do not produce the best results on every model, we are quite competitive and for all of the real subjects we are as good as, or improve upon the best known results.*

5.2 RQ2: Efficiency of Hyperheuristic HHSA

Table 7 summarizes the average execution time in seconds per subject within each group of benchmarks, using the three configurations of HHSA. The average execution time for the experiments with low configuration is about 1000 seconds (or 17 minutes).

Table 5: Sizes and times (seconds) for Syn-C2. The Best column reports the best results from CASA. The Df column is the difference between the best HHSA setting and the Best.

Sub.	Best	HHSA-L		HHSA-M		HHSA-H		Df	Sub.	Best	HHSA-L		HHSA-M		HHSA-H		Df
		Size	Time	Size	Time	Size	Time				Size	Time	Size	Time	Size	Time	
CS1	38	39	16	37	563	36	3,093	-2	CS16	19	24	27	24	177	24	689	5
CS2	30	30	30	30	391	30	1,074	0	CS17	39	41	16	36	575	36	2,648	-3
CS3	18	18	2	18	24	18	130	0	CS18	43	44	31	41	397	39	5,779	-4
CS4	20	20	7	20	164	20	448	0	CS19	47	50	96	46	1,134	44	10,685	-3
CS5	47	49	59	45	894	44	8,731	-3	CS20	53	55	90	52	1,286	50	12,622	-3
CS6	24	24	16	24	149	24	1,248	0	CS21	36	36	23	36	411	36	2,513	0
CS7	9	9	3	9	74	9	364	0	CS22	36	36	12	36	345	36	2,234	0
CS8	39	41	22	38	875	37	5,362	-2	CS23	12	12	2	12	11	12	188	0
CS9	20	20	27	20	253	20	682	0	CS24	44	46	18	41	283	40	3,909	-4
CS10	43	46	53	43	611	40	8,902	-3	CS25	49	51	37	47	748	46	6,399	-3
CS11	41	43	21	39	222	38	3,096	-3	CS26	30	31	17	28	348	27	1,927	-3
CS12	40	40	32	37	952	36	4,097	-4	CS27	36	36	8	36	151	36	671	0
CS13	36	36	45	36	598	36	3,309	0	CS28	50	53	77	50	902	48	10,709	-2
CS14	36	37	20	36	304	36	1,780	0	CS29	27	30	32	26	528	26	2,995	-1
CS15	30	30	11	30	239	30	628	0	CS30	17	19	12	17	158	16	1,405	-1
									Ov.	1,009	1,046	862	990	13,767	970	108,317	-39

Table 6: Three tables giving the sizes and times (seconds) for Real-1 2-way (top table), Real-2 2-way (middle table) and Real-2 3-way (bottom table). The Best Known column shows the best known results in the literature, and the tools that produced the results. References to the papers where these results are reported are listed.

Sub.	Best Known		HHSA-L		HHSA-M		HHSA-H		Diff
	Size	Tools	Size	Time	Size	Time	Size	Time	
Tools: A-ACTS, F-FoCuS, J-Jenny, P-PICT, C-CASA, T-Ttools									
Subject set: Real-1, 2-way [36]									
Con.	5	A,J	5	0	5	9	5	76	0
Sto.1	17	F	17	2	17	67	17	396	0
Ban.1	14	F	13	1	13	24	13	205	-1
Sto.2	18	F	18	1	18	23	18	100	0
Com.	16	F	16	3	16	86	16	898	0
Sys.	16	F	15	1	15	16	15	103	-1
Hea.1	30	A,F	30	2	30	49	30	193	0
Tel.	30	F	30	2	30	40	30	163	0
Ban.2	10	A	10	1	10	28	10	96	0
Hea.2	18	A,P,F	14	1	14	17	14	143	-4
Net.	115	F	110	2	110	63	110	229	-5
Sto.3	52	A,F	50	5	50	136	50	578	-2
Pro.1	28	J	23	1	22	14	22	123	-6
Ser.	102	F	100	10	100	266	100	1,008	-2
Ins.	527	A,P,F	527	13	527	411	527	1,549	0
Sto.4	130	P,F	117	3	117	80	117	345	-13
Hea.3	35	F	34	2	34	34	34	189	-1
Pro.2	32	A	28	5	27	54	27	66	-5
Sto.5	226	F	215	17	215	415	215	1,501	-11
Hea.4	47	F	46	3	46	45	46	230	-1
Overall	1,468	-	1,418	75	1,416	1,877	1,416	8,191	- 52
Subject set: Real-2, 2-way [5] [20]									
TCAS	100	C,T	100	6	100	166	100	578	0
SPIN-S	19	C	19	1	19	27	19	144	0
SPIN-V	32	C	33	11	31	212	31	1,725	-1
GCC	19	C	19	43	17	578	18	2,552	-2
Apache	30	C,T	31	71	30	656	30	3,676	0
Bugzilla	16	C,T	16	3	16	28	16	119	0
Overall	216	-	218	135	213	1,667	214	8,794	-3
Subject set: Real-2, 3-way [5] [19]									
TCAS	401	T	400	141	400	4,636	400	13,808	-1
SPIN-S	95	C	95	14	80	200	80	680	-15
SPIN-V	232	C	217	818	202	7,942	195	37,309	-37
GCC	94	C	102	7,562	94	83,324	-	-	0
Apache	177	C	193	25,258	176	191,630	-	-	-1
Bugzilla	59	C,T	61	156	59	1,769	60	1,726	0
Overall	1,058	-	1,068	33,949	1,011	289,501	-	-	-54

Despite the overall average of 17 minutes, the majority of the executions require fewer than 5 minutes. The 3-way experiments running GCC and Apache in the Real-2 benchmarks take the longest (1.6 hours on average). The high setting for this subject set was not finished after 3 days so we terminated it (indicated by '-'). HHSA-M is about 12 times slower overall than HHSA-L. However, most of the subjects still run within 10 minutes. The runtime for HHSA-H is about 7 times slower than for HHSA-M and takes at most 1.5 hours for the majority of the subjects.

On the right side of this table we see the 'Time Ratio' between the HHSA-L vs. HHSA-M and HHSA-M and HHSA-H, as well as the 'Size Improvement' which indicates how much smaller the second variant is. As we can see, while it costs us 12 times more to run the HHSA-M variant, it improves our sample sizes by almost 3 percent.

Moving from HHSA-M to HHSA-H improves our results by another 1%, while the cost is 7 times more in algorithm runtime. If we also consider the time to run test suites for this sample (see [19]), then this may make a practical difference. Consider if it takes overnight to run a full test suite for each configuration in our sample. The extra computation time for construction may pay off.

We next examine the practical implications of running the different variants of our algorithm. For this experiment we run all of the 2-way subjects in the Amazon EC2 (Elastic Compute Cloud) with the High-CPU On-Demand Instance (c1.medium) [2], and record not only the time, but the actual cost for obtaining our results.

We run the CASA tool as a baseline and the HHSA-L and HHSA-M settings. The results are shown in Table 8. The times shown represent the total time for all programs in the respective benchmarks (averaged over 5 runs). The HHSA-L setting took about 8 tenths of an hour to run all of the benchmarks, but cost only 13 cents. CASA took more time than the HHSA-L variant (2.9 hours) and cost $0.49. The medium variant of HHSA required the longest runtime (12.7 hours), but still only cost us $2.09.

Summary of RQ2. *We conclude that the HHSA algorithm is efficient when run at the lowest level (HHSA-low). When run at the higher levels we see a cost-quality tradeoff. In practice, however, the monetary cost of running these algorithms is very small.*

Table 7: **Running times (seconds) of the different levels of HHSA. Each time represents the average time for each individual model within the benchmark. Time Ratio and Size Impr. show the ratio and percent (respectively) between the L/M and M/H settings. '-' indicates that no result was obtained after 3 days.**

Subject Sets	HHSA-L Time	HHSA-M Time	HHSA-H Time	HHSA-L vs. HHSA-M		HHSA-M vs. HHSA-H	
				Time Ratio	Size Impr.	Time Ratio	Size Impr.
Syn-2	6	70	613	11	2.6%	9	1.6%
Real-1	4	94	409	25	0.3%	4	0.1%
Real-2	23	278	1,466	12	2.5%	5	0.9%
Syn-C2	29	459	3,611	16	6.1%	8	1.8%
Syn-3	332	620	6,054	2	-0.6%	10	0.4%
Real-2(3way)	5,658	37,007	-	7	5.4%	-	-
Average	1,009	6,421	2,431	12	2.7%	7	1.0%

Table 8: Sizes and times (seconds) and dollar cost for running each of the benchmark sets to completion in the Amazon EC2 Cloud with the High-CPU On-Demand Instance (c1.medium) [2].

Subjects	CASA			HHSA-L			HHSA-M		
	Time (s)	Cost$	Size	Time (s)	Cost $	Size	Time (s)	Cost$	Size
Syn-S2	5,777	0.26	808	220	0.01	820	2,350	0.11	805
Real-2	119	0.01	1,451	185	0.01	1,421	4,660	0.21	1,417
Real-1	265	0.01	233	383	0.02	222	3,971	0.18	216
Syn-C2	4,440	0.20	1,053	2,029	0.09	1,067	34,736	1.59	1,005
Overall	10,601	0.49	3,545	2,817	0.13	3,530	45,717	2.09	3,443

5.3 RQ3: Search Navigation Operator Comparison

We now examine how efficient and effective each of the search navigation operators are in isolation. To answer this question, we built seven versions of the simulated annealing algorithm, all using the HHSA-L settings. Each of the first six versions contains a single operator. For the seventh version, HH-Random, we include all operators, but the operator to use at each stage is chosen at random (with no intelligence).

The overall results for operator comparison are shown in Table 9. Each of the operators is listed in a row (Op1-Op6). The numbers correspond to their earlier descriptions (see Section 3.3). The next row is HH-Random, followed by the HHSA-L variant. The best operators on their own appear to be the "mutation" operators. Operator 4 (multiple mutation) seems to work relatively well on its own as does Operator 1 (single mutation). The HH-Rand variant performs second best which indicates that the combination of operators is helping the search, and it runs relatively fast, however without the guidance from learning it appears not do quite as well as the HHSA-L algorithm.

Summary of RQ3. *We conclude that there is a difference between effectiveness of each of the operators and that combining them contributes to a better quality solution.*

Table 9: "Navigation operator" comparison. Op1 to Op6 uses the standard SA with and individual search operator. HH-Rand makes a random choice at each evaluation. All variants are run using the low configuration. Time is in seconds.

Subjects	Syn-S2		Real-1		Real-2 (2-way)		Syn-C2	
	Size	Time	Size	Time	Size	Time	Size	Time
Op1	841	68	227	117	1,461	35	1,117	862
Op2	1,333	113	263	248	1,500	111	1,376	2,033
Op3	1,235	359	726	868	2,715	90	3,298	5,055
Op4	816	208	227	179	1,420	159	1,070	1,639
Op5	981	254	237	282	1,432	294	1,198	2,884
Op6	880	383	221	562	1,454	97	1,042	3,133
HH-Rand	812	321	218	441	1,419	113	1,024	2,903
HHSA-L	806	975	216	1,666	1,418	1,876	1,003	13,768

5.4 RQ4: Does the Hyperheuristic Algorithm Learn?

To determine if the operators that are selected by the hyperheuristic SA algorithm are *learned* we examine Table 10 and Figure 2. We first look at the graphs. The `x-axis` represents the different *problem states* which is the number of missing tuples that the problem has left to cover. On the left part of the graph, there are many tuples remaining uncovered, and towards the right, very few. We plot the reward scores from our learning algorithm for each operator at each stage (a higher reward score means the operator is more likely to be selected). We show this data for one synthetic and one real subject (due to space limitations), S2-8 (top), and TCAS (bottom). As we can see, early on when the problem is easier, most of the operators are close to the same reward value with one or two standing out (Operator 4 in S2-8 and Operator 5 in TCAS). This changes as we have fewer tuples to cover; most of the operators move towards a negative reward with a few remaining the most useful. Not only do we see different "stages" of operator selection, but we also see two different patterns.

We examine this further by breaking down data from each benchmark set into stages. We evenly split our data by iteration into an early (S1), middle (S2) and late (S3) stage of the search. For each, we select the pairs of operators that are selected most often across each benchmark set. For instance, Op1+Op4 is selected most often at stage S1 for 6 subjects in the set Syn-2 (see Table 10). In stage 1 we see that Op4+Op5 is selected most often overall, while in stage 2 it is Op1+Op4 and in stage 3 it is Op1+Op6. Within each benchmark we see different patterns. For instance, in the first stage Op1+Op5 is selected most often by the Syn-C2 (constrained synthetic) which is different from the others. In stage 2 again we see that the Syn-C2 has a different pattern of operator selection with Op1+Op4 being selected 14 times. In other sets such as the Real 1 we see that the Op4+Op6 combination is chosen most often.

Summary of RQ4. *We see evidence that the Hyperheuristic algorithm is learning both at different stages of search and across different types of subjects.*

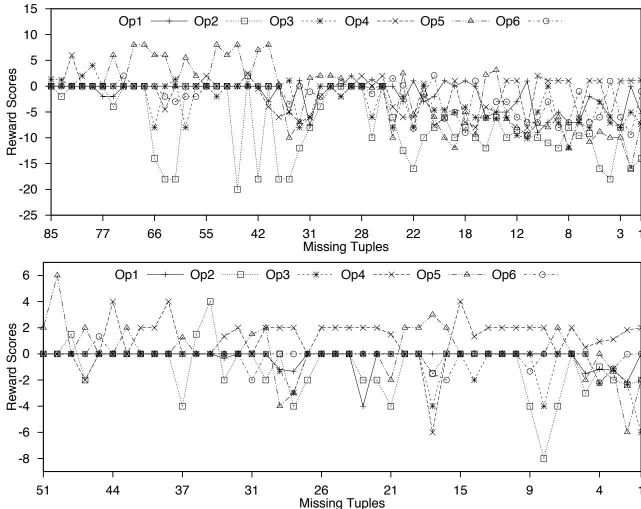

Figure 2: Subject: S2-8 (top) and TCAS (bottom). X-axis shows number of tuples left to cover. Y-axis shows the learning algorithm's reward scores.

Table 10: Learning Strategies. Three stages of the algorithm (S1-early), (S2-middle) and (S3-late) showing the pairs of operators chosen the most often by stage and subject set.

Strategies		Syn-2	Syn-C2	Real-1	Real-2	Ov.
Stage	Operators					
S1	Op1 + Op4	6	2	2	0	10
	Op1 + Op5	0	11	1	1	13
	Op4 + Op5	6	4	13	12	35
	Op4 + Op6	1	1	0	2	4
	Op5 + Op6	1	0	4	3	8
S2	Op1 + Op3	0	1	0	0	1
	Op1 + Op4	0	14	1	2	17
	Op1 + Op5	1	2	2	1	6
	Op1 + Op6	6	6	2	1	15
	Op3 + Op4	0	1	1	0	2
	Op3 + Op5	1	0	3	0	4
	Op3 + Op6	0	1	0	0	1
	Op4 + Op5	0	1	0	0	1
	Op4 + Op6	5	3	7	1	16
	Op5 + Op6	1	1	4	1	7
S3	Op1 + Op3	2	3	2	1	8
	Op1 + Op4	1	2	3	0	6
	Op1 + Op5	0	0	1	1	2
	Op1 + Op6	3	10	6	3	22
	Op2 + Op3	0	0	1	0	1
	Op3 + Op5	0	0	1	0	1
	Op3 + Op6	7	3	3	1	14
	Op4 + Op6	0	10	1	0	11
	Op5 + Op6	1	2	2	0	5

6. RELATED WORK

Research on the generation of covering arrays has a long history [13,30]. Many of the original tools were developed to work on unconstrained problems (or require manual remodeling) [8, 12, 27, 31], while more recent ones are specialized for constrained problems [4, 19, 36], or higher strengths [26].

Another classification is the type of algorithm used. In general the greedy approaches such as AETG, IPO and PICT are fastest, but may create larger sized samples [8, 15, 27] while heuristic search produces smaller array sizes, but requires longer running times [12, 19, 28, 31].

Garvin et al. [19] present a break-even approach to quantify the true cost between algorithms, which considers the time it costs to test the software as well as the time to build the arrays (i.e. this includes the impact of sample size as well as the generation time), and conclude that size tends to be the limiting factor for systems with even a short testing time per element in the covering array.

The mathematics community has developed both mathematical (constructive) and probabilistic proofs of bounds for a wide range of problems (see [14]), but these are synthetic models which may or may not be consistent with practice. We compare against some of these in this paper.

Another trend in CIT has been to specialize the construction for a particular process (incremental or adaptive) [17,18] which either builds covering arrays in stages to map to a particular test process, or iteratively modifies the model as it uncovers unknown constraints. But these techniques use a standard CIT algorithm as a primitive core, and are therefore orthogonal to our work.

Finally there are algorithms that are devised to work on very large models (such as the complete Linux kernel with 6,000+ factors) [32]. Our use of tunable settings for HHSA is consistent with this potential need.

Given the large mix of approaches to date, Calvagna et al. built CITLAB [6], a tool and language that brings together many other different tools for CIT into a single interface and framework. This allows the tester to execute different tools on the same benchmarks where the goal is to determine the "best" choice.

However, this framework does not remove the limitation that we are trying to solve – the tester still has to decide which algorithm works on which problem. Our approach differs from these approaches because our aim is not to be the "best" for any particular problem type, but to provide a generalist tool armed with online learning, that automatically adapts to each different problem model it encounters.

7. CONCLUSIONS

In this paper we have presented a hyperheuristic algorithm for constructing CIT samples. We have shown that the algorithm is general and *learns* as the problem set changes through a large empirical study on a broad set of benchmarks. We have shown that the algorithm is effective when we compare it across the benchmarks and other algorithms and results from the literature.

We have also seen that the use of different tunings for the algorithm (low, medium and high) will provide a quality-cost tradeoff with the higher setting producing better results, but taking longer to run. When we examine the practicality of such an algorithm, we see that the monetary cost for running the algorithm is quite small when using today's cloud ($2.09).

Finally, we have examined the various stages of learning of our algorithm and see that the different heuristic operators are more effective at different stages (early, middle, late) and that they vary across programs and benchmarks. It is this ability to learn and adapt that we believe is the most important aspect of this search.

As future work we will look at alternative tunings for the algorithm so that we can scale to very large problems (a very low setting) and can find even smaller sample sizes (a very high setting). We will also incorporate new operators and will look at alternative algorithms for the outer layer, such as genetic algorithms.

8. REFERENCES

[1] S. Ali, L. C. Briand, H. Hemmati, and R. K. Panesar-Walawege. A systematic review of the application and empirical investigation of search-based test-case generation. *IEEE Transactions on Software Engineering*, pages 742–762, 2010.

[2] Amazon. EC2 (Elastic Compute Cloud). Available at http://aws.amazon.com/ec2/.

[3] E. K. Burke, M. Gendreau, M. Hyde, G. Kendall, G. Ochoa, E. Ozcan, and R. Qu. Hyper-heuristics: a survey of the state of the art. *Journal of the Operational Research Society*, 2013. to appear.

[4] A. Calvagna and A. Gargantini. A formal logic approach to constrained combinatorial testing. *Journal of Automated Reasoning*, 45:331–358, December 2010.

[5] A. Calvagna and A. Gargantini. T-wise combinatorial interaction test suites construction based on coverage inheritance. *Software Testing, Verification and Reliability*, 22(7):507–526, 2012.

[6] A. Calvagna, A. Gargantini, and P. Vavassori. Combinatorial interaction testing with CITLAB. In *ICST*, pages 376–382, 2013.

[7] M. Chateauneuf and D. L. Kreher. On the state of strength-three covering arrays. *Journal of Combinatorial Designs*, 10(4):217–238, 2002.

[8] D. M. Cohen, S. R. Dalal, M. L. Fredman, and G. C. Patton. The AETG system: an approach to testing based on combinatorial design. *IEEE Transactions on Software Engineering*, 23(7):437–444, 1997.

[9] M. Cohen, M. Dwyer, and J. Shi. Constructing interaction test suites for highly-configurable systems in the presence of constraints: A greedy approach. *Software Engineering, IEEE Transactions on*, 34(5):633–650, 2008.

[10] M. B. Cohen, C. J. Colbourn, and A. C. H. Ling. Augmenting simulated annealing to build interaction test suites. In *Proceedings of the 14th International Symposium on Software Reliability Engineering*, ISSRE '03, pages 394–, Washington, DC, USA, 2003. IEEE Computer Society.

[11] M. B. Cohen, M. B. Dwyer, and J. Shi. Interaction testing of highly-configurable systems in the presence of constraints. In *Proceedings of the 2007 international symposium on Software testing and analysis*, ISSTA '07, pages 129–139, New York, NY, USA, 2007. ACM.

[12] M. B. Cohen, P. B. Gibbons, W. B. Mugridge, and C. J. Colbourn. Constructing test suites for interaction testing. In *Proceedings of the 25th International Conference on Software Engineering*, ICSE '03, pages 38–48, Washington, DC, USA, 2003. IEEE Computer Society.

[13] C. J. Colbourn. Combinatorial aspects of covering arrays. *Le Matematiche (Catania)*, 58:121–167, 2004.

[14] C. J. Colbourn. Covering array tables. Available at http://www.public.asu.edu/~ccolbou/src/tabby/catable.html, 2012.

[15] J. Czerwonka. Pairwise testing in real world. In *Pacific Northwest Software Quality Conference*, pages 419–430, October 2006.

[16] H. Do, S. G. Elbaum, and G. Rothermel. Supporting controlled experimentation with testing techniques: An infrastructure and its potential impact. *Empirical Software Engineering: An International Journal*, 10(4):405–435, 2005.

[17] E. Dumlu, C. Yilmaz, M. B. Cohen, and A. Porter. Feedback driven adaptive combinatorial testing. In *Proceedings of the 2011 International Symposium on Software Testing and Analysis*, pages 243–253, 2011.

[18] S. Fouché, M. B. Cohen, and A. Porter. Incremental covering array failure characterization in large configuration spaces. In *International Symposium on Software Testing and Analysis (ISSTA)*, pages 177–187, July 2009.

[19] B. Garvin, M. Cohen, and M. Dwyer. Evaluating improvements to a meta-heuristic search for constrained interaction testing. *Empirical Software Engineering*, 16(1):61–102, 2011.

[20] B. J. Garvin, M. B. Cohen, and M. B. Dwyer. An improved meta-heuristic search for constrained interaction testing. In *Proceedings of the 2009 1st International Symposium on Search Based Software Engineering*, SSBSE '09, pages 13–22, Washington, DC, USA, 2009. IEEE Computer Society.

[21] M. Harman. The current state and future of search based software engineering. In *2007 Future of Software Engineering*, FOSE '07, pages 342–357, Washington, DC, USA, 2007. IEEE Computer Society.

[22] M. Harman, E. Burke, J. Clark, and X. Yao. Dynamic adaptive search based software engineering. In *Proceedings of the ACM-IEEE international symposium on Empirical software engineering and measurement*, ESEM '12, pages 1–8, 2012.

[23] M. Harman, A. Mansouri, and Y. Zhang. Search based software engineering: Trends, techniques and applications. *ACM Computing Surveys*, 45(1):Article 11, November 2012.

[24] B. Hnich, S. Prestwich, E. Selensky, and B. Smith. Constraint models for the covering test problem. *Constraints*, 11:199–219, 2006.

[25] C. Le Goues, M. Dewey-Vogt, S. Forrest, and W. Weimer. A systematic study of automated program repair: Fixing 55 out of 105 bugs for $8 each. In *ICSE*, pages 3–13, 2012.

[26] Y. Lei, R. Kacker, D. R. Kuhn, V. Okun, and J. Lawrence. Ipog-ipog-d: efficient test generation for multi-way combinatorial testing. *Software Testing Verification and Reliability*, 18:125–148, September 2008.

[27] Y. Lei and K. Tai. In-parameter-order: a test generation strategy for pairwise testing. In *High-Assurance Systems Engineering Symposium, 1998. Proceedings. Third IEEE International*, pages 254–261, 1998.

[28] J. Martinez-Pena, J. Torres-Jimenez, N. Rangel-Valdez, and H. Avila-George. A heuristic approach for constructing ternary covering arrays using trinomial coefficients. In *Proceedings of the 12th Ibero-American conference on Advances in artificial intelligence*, IBERAMIA'10, pages 572–581, 2010.

[29] P. Nayeri. *Post-Optimization: Necessity Analysis for Combinatorial Arrays*. Ph.D. thesis, Department of Computer Science and Engineering, Arizona State University, April 2011.

[30] C. Nie and H. Leung. A survey of combinatorial

testing. *ACM Computing Surveys*, 43(2):1–29, February 2011.

[31] K. Nurmela. Upper bounds for covering arrays by tabu search. *Discrete Applied Mathematics*, 138(1-2):143–152, 2004.

[32] G. Perrouin, S. Sen, J. Klein, B. Baudry, and Y. l. Traon. Automated and scalable t-wise test case generation strategies for software product lines. In *Proceedings of the 2010 Third International Conference on Software Testing, Verification and Validation*, pages 459–468, 2010.

[33] J. Petke, S. Yoo, M. B. Cohen, and M. Harman. Efficiency and early fault detection with lower and higher strength combinatorial interaction testing. In *Proceedings of the 9th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE'13)*, pages 26–36, August 2013.

[34] O. Räihä. A survey on search based software design. Technical Report Technical Report D-2009-1, Department of Computer Sciences, University of Tampere, 2009.

[35] D. Richard Kuhn and V. Okum. Pseudo-exhaustive testing for software. In *Proceedings of the 30th Annual IEEE/NASA Software Engineering Workshop*, SEW '06, pages 153–158, Washington, DC, USA, 2006. IEEE Computer Society.

[36] I. Segall, R. Tzoref-Brill, and E. Farchi. Using binary decision diagrams for combinatorial test design. In *Proceedings of the 2011 International Symposium on Software Testing and Analysis*, ISSTA '11, pages 254–264, New York, NY, USA, 2011. ACM.

[37] J. Stardom. *Metaheuristics and the Search for Covering and Packing Arrays.* Thesis (M.Sc.)–Simon Fraser University, 2001.

[38] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction.* MIT Press, 1998.

[39] Y.-W. Tung and W. Aldiwan. Automating test case generation for the new generation mission software system. In *Aerospace Conference Proceedings, 2000 IEEE*, volume 1, pages 431–437 vol.1, 2000.