Genetic Algorithms for Redundancy in Interaction Testing

Ryan E. Dougherty
Colgate University
Hamilton, New York
rdougherty@colgate.edu

ABSTRACT

It is imperative for testing to determine if the components within large-scale software systems operate functionally. Interaction testing involves designing a suite of tests, which guarantees to detect a fault if one exists among a small number of components interacting together. The cost of this testing is typically modeled by the number of tests, and thus much effort has been taken in reducing this number. Here, we incorporate redundancy into the model, which allows for testing in non-deterministic environments. Existing algorithms for constructing these test suites usually involve one “fast” algorithm for generating most of the tests, and another “slower” algorithm to “complete” the test suite. We employ a genetic algorithm that generalizes these approaches that also incorporates redundancy by increasing the number of algorithms chosen, which we call “stages.” By increasing the number of stages, we show that not only can the number of tests be reduced compared to existing techniques, but the computational time in generating them is also greatly reduced.

CCS CONCEPTS

• **Software and its engineering** → **Search-based software engineering**;

KEYWORDS

covering array, genetic algorithm, redundancy, search-based software engineering

ACM Reference Format:

Ryan E. Dougherty. 2020. Genetic Algorithms for Redundancy in Interaction Testing. In *Genetic and Evolutionary Computation Conference (GECCO '20)*, July 8–12, 2020, Cancun, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3377930.3390230>

1 INTRODUCTION

The task of testing software systems has always remained a challenging problem, in that a tester is to find and eliminate any defects that were not found during earlier stages of development. One systematic approach has been interaction testing [6, 12, 20], wherein a system configuration is presented, and the task is to design a series of tests such that any interaction between components of at most a certain size is tested. It is empirically shown in [11] that a very

Browser	OS	Connection	Printer
Safari	Windows	LAN	Local
Safari	Linux	ISDN	Networked
Safari	macOS	PPP	Screen
IE	Windows	ISDN	Screen
IE	macOS	LAN	Networked
IE	Linux	PPP	Local
Chrome	Windows	PPP	Networked
Chrome	Linux	LAN	Screen
Chrome	macOS	ISDN	Local

Figure 1: A Covering Array with 9 tests ($N = 9$), 4 factors ($k = 4$), each with 3 levels ($v = 3$), and strength 2 ($t = 2$).

large percentage of errors within a software system can be detected with interactions of small size.

Suppose that a system to be tested has *factors* $f_1, \dots, f_k$, and each factor f_i has a set of allowed *levels* (i.e., valid inputs to f_i). Let t be a positive integer at most k , the number of factors. We design a test suite (i.e., a set of tests) such that each one is a valuation of each factor to one of its valid inputs, with the property that any set of t or fewer factors is exhaustively tested. By this last term, we mean that no matter the set of t factors chosen, and any level selected for each factor, at least one test has that choice. The parameter t is the *strength* of the test suite.

We present an equivalent formulation. Define an *interaction* of size t (or a t -way interaction) to be a set of the form

$$\{(f_i, \ell_i) : \ell_i \text{ is a level for factor } f_i, 1 \leq i \leq t\}.$$

Let $\mathcal{I}_{k,t}$ be the set of all interactions of size at most t with k factors. Then we say an array A with N rows and k columns is a *mixed-level covering array* $\text{MCA}(N; t, (v_1, \dots, v_k))$ if (1) column i corresponds to factor f_i , (2) v_i is the number of levels for f_i , and (3) all interactions in $\mathcal{I}_{k,t}$ appear at least once in the array. When $v_1 = \dots = v_k = v$, we say that the array A is a (uniform) *covering array* $\text{CA}(N; t, k, v)$. This paper only concerns uniform covering arrays.

To test a system using a covering array, a tester observes each test in turn, and chooses the corresponding level for each factor as dictated by the test. Observe the test suite in Figure 1, reproduced from [4]; we claim it is a $\text{CA}(9; 2, 4, 3)$. There are 9 rows (representing the tests), 4 factors (Browser, OS, Connection, and Printer), 3 levels for each factor (e.g., Safari, IE, and Chrome for the factor Browser), and has strength 2. The tester then records whether the output of the system, after the test is performed, is expected; if this property is true for all tests, then the coverage property guarantees no fault exists in the system that involves at most t factors. If a fault exists within this system, it must be due to an interaction of 3 or more factors, since all 2-way interactions appear at least once.

Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
GECCO '20, July 8–12, 2020, Cancun, Mexico
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7128-5/20/07...\$15.00
<https://doi.org/10.1145/3377930.3390230>

With covering arrays we develop in this paper, we substitute the levels in each factor for 0, 1, 2, $\dots$ for ease of presentation.

With the standard definition, we assume that the system is *deterministic*, in that running a given test multiple times always produces the same result. For example, in the CA from Figure 1, the first test of Safari/Windows/LAN/Local, if repeated, would produce the same output. However, the assumption of the system's being deterministic is often not realistic, particularly when there is *noise* or *randomness* within the testing environment. If we amend our definition to say that each interaction appears *at least a given number of times* λ , then by increasing λ , we gain further confidence in the correctness of the system, even if it is not deterministic. For example, if $\lambda = 5$, then if the resulting output matches what was expected 5 times in a row, it is much less likely that a fault still exists than if, say, $\lambda = 1$.

We update our notation with *index* λ as follows: $CA_\lambda(N; t, k, v)$, in that every interaction appears at least λ times. It is possible to generalize our model by having each interaction I have its own corresponding index λ_I , but we content ourselves with the generality developed here, since we can take the maximum index over all interactions to be the λ we seek.

To minimize total cost of the tester, the most often chosen metric for a covering array of “smallest cost” is finding the smallest number of tests N for which such an array exists. The *covering array number*, $CAN_\lambda(t, k, v)$, is this quantity. Much work has been to determine these values for $\lambda = 1$ [4, 5, 15], but none for when $\lambda > 1$, as far as we are aware. When $\lambda > 1$, we say that the array is of *higher-index*.

Sarkar and Colbourn [16] introduced the *two-stage* framework for constructing covering arrays, which is as follows. The first stage generates an array randomly which covers all but (at most) a certain number of interactions, and the second stage deterministically covers the remaining interactions with more rows until all interactions are covered. Their methods determine the smallest number of rows in the first stage such that the expected number of rows produced after the second stage is minimized. They determined that in the second stage there is a trade-off between computation time and bound on CAN_λ , and no one method is uniformly better than the others. Further, the first stage is comparatively faster than the second, but suffers from not giving a guarantee on the exact number of interactions left for the second stage, or their structure. And finally, they only studied the $\lambda = 1$ case.

There have been much previous work involving genetic algorithms and other metaheuristic techniques for covering arrays and related objects [9, 13, 14, 18, 19]. However, our approach is different because our genetic algorithm is focused on determining an optimal number of *constructive* stages (as well as which method to use for each), whereas all existing techniques focus on mutations to the array itself in the hope of forming a CA. A main advantage of our approach is that the algorithms chosen are fully deterministic and replicable.

The contributions of this paper are that

- (1) we extend the methods of Sarkar and Colbourn by introducing a *multi-stage* framework for higher-index covering arrays using a genetic algorithm;

- (2) our framework is general enough to allow for any choice of algorithm at any stage, with any number of stages, and any index;
- (3) our genetic algorithm's choices of which stages to select, and how many of them, yield a dramatic decrease in computational time in the creation of these arrays, sometimes by two orders of magnitude; and
- (4) we show that there is a pattern to which algorithms are best to choose at each stage by analyzing the Pareto fronts resulting from our genetic algorithm.

2 OUR MULTI-STAGE FRAMEWORK

Denote $MS(A_1(\lambda_1), \dots, A_M(\lambda_m))$ to be a *multi-stage selection with algorithms* $A_1, \dots, A_m$, such that after algorithms $A_1, \dots, A_i$ are applied, then a covering array is produced with index $\lambda_1 + \dots + \lambda_i$. If the desired index is λ , then it is the case that $\lambda = \lambda_1 + \dots + \lambda_m$. We call the application of A_i to be the *i th stage*. Once the i th stage is completed, the resulting array is fed as input to the $(i + 1)$ -st stage. The goal of each algorithm A_i is to be computationally efficient, require few rows to complete, and cover many interactions to reduce the cost of future stages.

2.1 Why Multiple Stages?

It is not immediately clear why multiple stages may be advantageous over a single stage. For multiple stages, a software tester would have to choose which algorithms to use in each stage, what order to choose them, what index to choose for each, and above all, there is not necessarily any communication between the stages on how to optimize the parameters for each. A single stage also is easier to implement and maintain, often has performance guarantees (in terms of run time to produce the covering array and its final size), and the produced arrays often have a structure that can be easily analyzed, since the method to construct it is fixed.

We give a simple, but representative example of why multiple stages make sense, and an intuitive understanding why; a more detailed explanation is in Section 5. The Unix `sort` command has 18 binary flags; an exhaustive test suite to test the correctness of `sort` would involve 2^{18} tests! Suppose that we want to build a test suite with $t = 2$; the best known (and proven optimal) covering array with $\lambda = 1$ has 8 rows [10]. Now suppose we desire to have a redundancy of $\lambda = 5$; in other words, we desire to build a $CA_5(N; 2, 18, 2)$, with N as small as possible. A naïve approach would say that 40 rows is possible, simply by juxtaposing the 8-row array 5 times.

An implementation of an extension to a well-known greedy algorithm (described in Section 3.3) produces an array with 29 rows in a single stage. However, a simple choice of algorithm selection found that there exist 3-stage and 4-stage solutions yielding 27 rows; furthermore, all of these multiple-stage solutions complete in 50% or less time compared to the single-stage one. In fact, all solutions found between 27 and 28 rows had 2 or more stages, among all algorithms tested.

An intuitive understanding of why multiple stages are better is that the algorithm produces the array one row at a time, and is “greedy” in the sense that many interactions are covered early, and only a small number are left uncovered in the last rows. Consider the two CA_{5s} in Figure 2; the one with $N = 29$ rows (right) was

000000000000000000	000000000000000000
111111111111111110	111111111111111110
010101010101010101	010101010101010101
101010101010101011	101010101010101011
001100110011001100	001100110011001100
110011001100110000	110011001100110010
000011110000111110	011010010110100101
111100001111000010	100101101001011011
000000001111111100	000011111111000000
111111110000000000	111100000000111110
100101100110100101	010110100101101011
011010011001011011	101001011010010101
000000000000000010	001111000011110000
111001111101101111	110000111100001110
001110000110110111	011001101001100111
110110111011010001	100110010110011001
010101001010011011	000000001111111100
101011010111001001	111111110000000010
011100100001110101	010101011010101011
100011111100001100	101010100101010101
010011100011101010	001100111100110010
101100011000110110	110011000011001100
001001101101010011	000011110000111100
110110010110101001	111100001111000010
010010101110000100	000000000000000010
100101000001011110	111111111111111100
001000010000100000	000000000000000000
	000000000010000110
	010000001000000000

Figure 2: A $CA_5(27; 2, 18, 2)$ (left) and a $CA_5(29; 2, 18, 2)$ (right).

produced via this greedy method, and the one with $N = 27$ rows (left) was produced with 3 stages, the first and third being this same greedy method, and the second via graph coloring. For the latter, the first and second stage has an index of 1, and the third stage has an index of 3. Therefore, the selection was $MS\langle G(1), C(1), G(1) \rangle$, where G represents greedy, and C represents coloring. The horizontal lines in the left CA_3 indicate where one stage ends, and another begins. Not only is there a reduction in the rows, but the left CA was produced in 75% less time than the right CA.

We investigate these two CAs further. In Figure 3, we give a scatter plot corresponding to each of the CAs as follows. The horizontal axis indicates row numbers, and for a particular row i , we calculate the number of interactions N_i such that they become λ -covered in row i . For example, for the left CA in Figure 2, in the first two columns, the two values 00 appear in rows 0, 4, 6, 8, and 12, among others. Therefore, we have that the interaction $\{(0, 0), (1, 0)\}$ becomes 5-covered in row 12. The vertical axis in the plot indicates a cumulative total of interactions that are 5-covered. The first CA's plot achieves a vertical value of $\binom{k}{t} v^t = 612$ at $N = 27$, and the second at $N = 29$. We can see that when the first stage ends after row 10, the plot for the left CA increases more quickly than for the right CA. What is less clear is why, even though the right CA overtakes the left, it still requires more rows.

Figure 3: Number of interactions first λ -covered in each row of the $CA_5(27; 2, 18, 2)$ (circles) and $CA_5(29; 2, 18, 2)$ (crosses) in Figure 2.

A simple analysis of the interactions left to be covered in the final rows (for the right CA) shows that there are many *conflicts* between the interactions that force them to appear in separate rows. However, the algorithm used for this CA inherently is a one-row-at-a-time method and cannot “backtrack” to make any necessary changes to earlier rows. Even though the left CA uses the same algorithm, it has multiple stages, and thus the stages can “communicate” information about the remaining interactions, whereas the right CA cannot since it uses only one stage.

We ran a systematic experiment to further illustrate why multiple stages are advantageous. Our experiment involved all algorithms explained in Section 3, and all decompositions of a target index $\lambda = 5$ into at most 5 choices of algorithms (possibly repeating). We were able to speed up the computation substantially by noting that all choices of algorithms described in Section 3 are deterministic. Therefore, for any two multi-stage strategies starting with the same algorithm and index can save the result of the first stage without having to repeat the computation ourselves. We saved the time it took to generate that initial array, and added its creation time to the total computation time when the final array is constructed.

The results of this experiment appear in Table 1. One can easily see that the number of rows varies substantially depending on the number of stages taken. As the number of stages increases, the maximum, average, and standard deviation for the number of produced rows all decrease. If one were to pick a number of stages a priori and then choose a set of methods uniformly at random, the obvious best choice would be 5 stages in this case, since the average is the smallest. As an added benefit, the minimum, maximum, average, and median time taken to construct these arrays also decrease, some by over an order of magnitude.

3 DETERMINISTIC ALGORITHMS FOR CAs

In this section, we describe each of the constructive, deterministic algorithms for CAs that were used in our genetic algorithm. Each

NS	Min N	Max N	Avg N	Median N	StdDev N	Min T	Max T	Avg T	Median T	StdDev T
1	29	3060	826	138	1481	0.361	17.129	6.943	4.870	6.460
2	27	2448	541	119	738	0.014	11.614	2.692	1.785	3.141
3	27	1836	388	130	473	0.013	7.076	1.318	0.802	1.623
4	27	1224	297	134	323	0.013	4.006	0.742	0.393	0.886
5	28	627	239	130	228	0.013	1.815	0.455	0.264	0.457

Table 1: Statistics of the number of stages NS when $t = 2, k = 18, v = 2, \lambda = 5$, where N is the number of rows produced, and T is the total computational time (in seconds). The minimum, maximum, average, median, and standard deviation for both N and T are given. The “Avg N ,” “Med N ,” and “StdDev N ” columns are rounded to the nearest integer.

stage starts with every interaction having been covered at least some number of times α , and the “goal” of the stage is to finish with every interaction covered at least another number of times β , where $\alpha < \beta$. We call α the *current index*, and β the *desired index*.

3.1 Basic

This stage involves adding one row for each uncovered interaction with index between α and β . For all entries in the rows that do not involve a specific interaction, we put a fixed value f . The advantage of this method is clear, in that virtually no additional computation is needed; however, this method suffers by having value f appear in many places; further, interactions that involve many occurrences of f are covered many times, whereas for other interactions, not nearly as much.

For these reasons, we make the basic algorithm more adaptive. Instead of inserting the fixed value f , we insert into column c any entry that occurs least frequently among rows already constructed. By doing this, we are attempting to have columns contain all entries as equally often as possible. Because this more adaptive method produced either equal or better results than the original one in practice, and has very small computational difference for cost, we refer to the adaptive method as the basic one from here on.

3.2 Graph Coloring

Sarkar and Colbourn [16] define an *incompatibility graph* for a set of interactions as follows. The vertices are the interactions that are not covered, and an edge is formed between any two interactions I_1, I_2 if in some column shared between I_1, I_2 , the corresponding values are different. The smallest number of colors in a proper coloring of such a graph is the minimum number of rows needed to cover all of these remaining interactions.

We could adapt their strategy for higher-index covering arrays by simply iterating their method multiple times to achieve the desired index. However, we can build a single graph and solve the coloring problem in one stage for higher index, as follows.

The vertices of the graph are all interactions I that are covered at least α times, and fewer than β times, paired with an integer s_I with $\alpha \leq s_I < \beta$. To form edges, let I_1 and I_2 be two such interactions, with values s_{I_1}, s_{I_2} . If I_1, I_2 share a column with different symbols (regardless of the values of s_{I_1}, s_{I_2}), form an edge. Otherwise, if I_1, I_2 involve the *same* set of columns, but $s_{I_1} \neq s_{I_2}$, we form an edge also. This graph has a proper coloring with N colors if and only if N rows can be formed to cover all such interactions.

Since there is no known polynomial-time algorithm for the graph coloring problem, we use the following two “greedy” graph coloring algorithms: “largest first,” and “smallest last.” For both, the vertices are given an order. The former sorts the vertices in non-increasing order of degree. The latter sorts the vertices as $v_1, \dots, v_n$ whenever v_i has the minimum degree in the maximal subgraph among the vertices $v_1, \dots, v_i$ for every i . In all cases, the graph coloring algorithm works as follows. We process the vertices in the given order according to the method. Additionally, an order on the available colors is made; initially, only one color is available. For an uncolored vertex v , we observe the neighbors of v . If any color available is possible to be assigned to v such that no neighbor of v has the same color, we choose the “smallest” color. Otherwise, we allocate a new color, and assign v to that color. At the end, the number of colors allocated is the number of corresponding rows in the covering array.

3.3 Density

The probabilistic method [1] for covering arrays [17] says that for any t, k, v where v, t are fixed, there must exist some $N = \Theta(\log k)$ for which a covering array on that many rows exists. One can show that the optimal number of rows for a higher-index CA asymptotically is competitively small compared to $O(\lambda \cdot \log k)$, involving juxtaposing the array λ times [8]. However, the method does not produce the array, but merely shows that such an array exists.

One can turn this idea into an easy randomized algorithm that generates the array one row at a time, as follows. Initially, have A be an empty array. Suppose that ρ interactions remain to be covered. The probability that any one of these interactions is covered in a row with entries chosen uniformly at random is $1/v^t$. Repeatedly generate rows $R_1, R_2, \dots$ uniformly at random until some row R_i covers a number of interactions for the first time c , where $c \geq \rho/v^t$. When such a row is found, append R_i to A , and update the list of interactions not covered. An easy analysis shows that when all interactions are covered, A will have size at most a constant times $\log k$ (where the constant depends on v, t).

It is desired to have a deterministic algorithm that has this property, so that a guarantee on the number of rows produced can be made. Bryce and Colbourn [3] designed the “conditional expectation” (or *density*) algorithm, which is deterministic, and a covering array is produced that meets the same logarithmic bound. However, their methods do not immediately translate to covering arrays of higher index.

Our method, which generalizes their work, is as follows: let $p = 1/v^t$ be the probability of coverage in a random row as before. The probability that an interaction is not λ -covered in N rows is $\sum_{i=0}^{\lambda-1} \binom{N}{i} p^i (1-p)^{N-i}$. The expected number of interactions not λ -covered after N rows, if chosen uniformly at random, is $\binom{k}{t}$ times this probability (or, in general, it is the number of uncovered interactions times this probability). When this expectation $E(N)$ is strictly less than 1, a higher-index CA exists.

We now state our constructive algorithm. First, find the smallest N such that $E(N) < 1$. Let A be a covering array to-be-built (empty or not). Let T be a t -way interaction, and r a to-be-completed row that has some column c not fixed to a value. For c , choose a factor c_f . Now, determine the expected number of remaining interactions left not λ -covered in the remaining $N - |A|$ incomplete rows, if these rows are chosen at random and column c of row r is fixed to c_f . Then, choose any factor for column c that *minimizes* this expectation. Once all entries of row r are fixed, we decrease the choice for N according to how much the expectation decreased after generating r . At a high level, N is the estimate on how many rows are needed, and when row is completed, the estimate is updated accordingly (for details, see [8]).

4 EXPERIMENTS

We designed our experiments to answer the following research questions:

- RQ1: Do multiple stages improve the sizes of CAs of higher-index compared to a single stage?
- RQ2: Do multiple stages improve the computational cost in generating such CAs compared to a single stage?
- RQ3: For “optimal” multiple-stage selections, is there a pattern in regard to algorithms chosen or index values, and how many stages are best for a given parameter situation?

To address these questions, we implemented a genetic algorithm with individuals encoding the multiple stages of algorithms, as well as the corresponding index for each. The representation of an individual is a multi-stage selection $MS\langle A_1(\lambda_1), \dots, A_m(\lambda_m) \rangle$.

4.1 Mutation

Here we describe the mutation operators used in our genetic algorithm. Suppose that the individual is $MS\langle A_1(\lambda_1), \dots, A_m(\lambda_m) \rangle$. For mutation, there are several operators that we have employed:

- Append: choose an algorithm B uniformly at random from the ones listed in Section 3. Choose any A_i uniformly at random with $\lambda_i > 1$. Then, append B to the individual (at the end) with index λ_{m+1} to be a randomly chosen integer between 1 and $\lambda_i - 1$ (inclusive). We also reduce the index of A_i by λ_{m+1} . If all of the $\lambda_i = 1$, then we remove a randomly chosen A_i before adding B with its index being 1.
- Swap: swap the order of two randomly chosen A_i and A_j and their indexes ($i \neq j$).
- Index Transfer: take two randomly chosen A_i and A_j ($i \neq j$), subtract a randomly chosen integer from one of their indexes, and add the same integer to the other index.
- Modify: substitute a randomly chosen A_i for any other algorithm in its place, with the same index.

- Join: randomly choose A_i and A_j ($i \neq j$), remove them from the individual, and re-insert a new algorithm A_k , chosen uniformly at random from $\{A_i, A_j\}$; furthermore, the new corresponding index is $\lambda_i + \lambda_j$.

When an individual I is to be mutated, one algorithm from these will be selected uniformly at random to apply to I . The selection of these operations was to allow for individuals of all types, including the first stage being of high index, many stages with low indexes each, any combination of algorithms including adjacent repetitions, etc.

4.2 Crossover

Crossover works as follows. Suppose that the two individuals are of the form $MS\langle A_1(\lambda_1), \dots, A_m(\lambda_m) \rangle$, and $MS\langle B_1(\gamma_1), \dots, B_n(\gamma_n) \rangle$. We pick two positive integers a, b at random such that $a + b \leq \lambda$. We form a new individual by choosing a random nonempty subset of size a from $\{A_1, \dots, A_m\}$, another of size b from $\{B_1, \dots, B_n\}$, and composing them together, ordering the chosen algorithms at random.

Suppose our formed individual is $MS\langle C_1(\beta_1), \dots, C_{a+b}(\beta_{a+b}) \rangle$, where the β_i come from the corresponding individual and algorithm. Note that the sum of the β_i in our new individual may exceed λ . If this is the case, we choose a $\beta_i > 1$ at random and decrease β_i by 1; repeat this procedure until the sum of the β_i is equal to λ . If the sum originally is less than λ , then we perform nearly the opposite operation: choose any β_i (may be equal to 1) at random, and increase β_i by 1; repeat until the sum is λ .

The fitness of an individual, naturally, is a tuple (N, T) , where N is the number of rows in the produced covering array, and T is the total computational time. When we sort the population by fitness, we sort by N first, and then if two individuals have the same number of rows, we sort them by their T values. We used the NSGA-II algorithm [7] to observe the Pareto front of the population.

4.3 Results

We evaluated our genetic algorithm on various parameter situations for $t \in \{2, 3, 4\}$. The results are reported in Table 2. Here is a short-hand list of all the algorithms reported there:

- Basic: B
- Coloring (largest first): L
- Coloring (smallest last): S
- Density: D

For ease of presentation in the table, we shorten

$$MS\langle A_1(\lambda_1), \dots, A_m(\lambda_m) \rangle$$

to be a list of the algorithm names, with a subscript indicating the index. For example, the individual representing

$$MS\langle B(1), S(1), B(1), L(2) \rangle$$

is instead written as B_1, S_1, B_1, L_2 . In all our experiments, our population had 300 individuals, and we report results for up to 100 generations. Table 3 contains the results for a single stage. Further, in Figures 4 to 6, we provide figures of scatter plots for the situations:

- $t = 2, k = 10, v = 2, \lambda = 10$;
- $t = 3, k = 10, v = 2, \lambda = 5$; and

Figure 4: Pareto Fronts of $CA_\lambda(N; t, k, v)$ s formed, where $\lambda = 10, t = 2, k = 10, v = 2$, for Generations 1, 10, 50, and 100. Circles are Generation 1, Squares are Generation 10, Diamonds are Generation 50, and Crosses are Generation 100.

Figure 6: Pareto Front of $CA_\lambda(N; t, k, v)$ s formed, where $\lambda = 5, t = 4, k = 10, v = 2$. Circles are Generation 1, Squares are Generation 10, Diamonds are Generation 50, and Crosses are Generation 100.

Figure 5: Pareto Front of $CA_\lambda(N; t, k, v)$ s formed, where $\lambda = 5, t = 3, k = 10, v = 2$. Circles are Generation 1, Squares are Generation 10, Diamonds are Generation 50, and Crosses are Generation 100.

- $t = 4, k = 10, v = 2, \lambda = 5$.

The horizontal axis (time) is given in logarithmic scale. The circles indicate generation 1, squares are generation 10, diamonds are generation 50, and crosses are generation 100. These figures show the Pareto fronts of the genetic algorithm for each of the same generations examined, and are representative over all parameter situations given in Table 2.

5 DISCUSSION

We discuss each of the research questions in turn. For RQ1, it is quite evident that multiple stages do improve the sizes of covering arrays. For example, observe the situation of $t = 4, k = 10, v = 2, \lambda = 5$: the most fit individual produced had 112 rows, whereas even a two-stage solution had 114 for its fittest individual. Even though the individuals with smallest N did not improve the number of rows significantly, what is more striking is when one observes Figures 4 to 6, especially Figure 4: the distribution across the vertical axis decreases substantially as the number of generations increases.

For RQ2, it is even more evident that multiple stages improve the computational time to construct covering arrays. A striking example is when $t = 2, k = 10, v = 2, \lambda = 10$: the individual with 44 rows in Generation 1 had a run-time of 9.559 seconds, and by Generation 100, an individual exists with the same number of rows but only takes 0.083 seconds, a improvement by over two orders of magnitude. In fact, the construction time difference between this final individual and the one with 206 rows is so small that the latter can be completely ignored. Further proof of the strong improvement is in Figures 4 and 5, in that the Pareto fronts shift left considerably as the number of generations increases.

For RQ3, we note that for all entries in the “Lowest N ” columns, the first algorithm chosen is either density (D) or “smallest last” (S). Also, in the “Lowest Time” columns, the first algorithm chosen is always “Basic” (B). Interestingly enough, “Largest First” (L) never was the first algorithm. An analysis of the output data reveals that L does in fact produce covering arrays with a small N relative to the entire distribution, but it is not the smallest N found. An explanation of this may be that this algorithm has to work with a set of interactions that has many symmetries, and one choice among many vertices with high degree may end up with more rows than another choice does.

t	k	v	λ	Generation 1		Generation 10		Generation 50		Generation 100	
				Lowest N	Lowest Time	Lowest N	Lowest Time	Lowest N	Lowest Time	Lowest N	Lowest Time
2	10	2	5	24, 0.482 S_4, L_1	189, 0.003 B_1, S_1, B_1, L_2	24, 0.157 S_2, S_1, D_1, L_1	189, 0.003 B_1, B_2, L_2	24, 0.071 S_1, D_1, L_1, L_1, S_1	189, 0.003 B_1, L_1, L_1, L_2	24, 0.071 S_1, D_1, L_1, L_1, S_1	189, 0.003 B_1, L_1, L_1, L_2
2	10	2	10	44, 9.559 D_9, S_1	360, 0.004 B_2, L_2, B_6	44, 1.685 D_1, S_1, D_1, D_3, D_4	360, 0.004 B_2, B_8	44, 0.272 $S_1, L_1, S_1, D_1, L_1, D_1, D_2, L_2$	206, 0.004 $B_1, L_1, S_1, S_1, S_1, S_1, S_1, L_1, L_1, L_1, L_1$	44, 0.083 $S_1, S_1, D_1, L_1, L_1, L_1, L_1, S_1, L_1, L_1, L_1$	206, 0.004 $B_1, L_1, S_1, S_1, S_1, S_1, S_1, L_1, L_1, L_1, L_1$
2	10	3	5	56, 27.628 D_5	816, 0.009 B_2, L_1, L_2	55, 4.838 S_1, D_1, D_3	442, 0.007 B_1, L_1, B_1, L_1, B_1	55, 4.838 S_1, D_1, D_3	429, 0.007 B_1, B_1, B_1, B_1, L_1	55, 4.838 S_1, D_1, D_3	429, 0.007 B_1, B_1, B_1, B_1, L_1
2	20	2	5	28, 2.179 S_1, S_3	770, 0.019 B_1, S_4	27, 0.369 S_1, S_1, L_1, L_1, L_1	769, 0.017 B_1, L_1, L_3	27, 0.369 S_1, S_1, L_1, L_1, L_1	769, 0.017 B_1, L_1, L_3	27, 0.369 S_1, S_1, L_1, L_1, L_1	769, 0.017 B_1, L_1, L_3
2	20	2	10	48, 26.623 D_1, D_1, S_1, D_7	1520, 0.042 B_2, D_1, L_7	48, 12.732 $D_1, S_1, L_1, D_1, D_1, D_5$	787, 0.024 B_1, L_2, L_6, S_1	48, 7.772 $D_1, S_1, L_1, D_1, D_1, D_4, S_1$	787, 0.024 B_1, L_2, L_6, S_1	48, 7.772 $D_1, S_1, L_1, D_1, D_1, D_4, S_1$	787, 0.024 B_1, L_2, L_6, S_1
3	10	2	5	48, 33.537 D_5	1050, 0.023 B_1, S_1, B_3	48, 3.381 D_1, D_1, D_1, L_1, D_1	976, 0.015 B_1, B_1, B_1, L_1, B_1	48, 3.381 D_1, D_1, D_1, L_1, D_1	976, 0.015 B_1, B_1, B_1, L_1, B_1	48, 3.381 D_1, D_1, D_1, L_1, D_1	976, 0.015 B_1, B_1, B_1, L_1, B_1
4	10	2	5	114, 169.255 D_1, D_4	6720, 0.093 B_2, B_3	112, 53.795 S_1, L_1, D_3	3418, 0.067 B_1, S_1, B_1, S_2	112, 53.795 S_1, L_1, D_3	3416, 0.070 B_1, B_1, L_1, L_1, L_1	112, 53.795 S_1, L_1, D_3	3416, 0.070 B_1, B_1, L_1, L_1, L_1

Table 2: Genetic Algorithm Results, with two individuals reported for each parameter situation: the one with the lowest N , and the one with the lowest time, at generations 1, 10, 50, and 100. The two numbers in each entry signify that individual's N and time, as well as the algorithms with indexes chosen as subscripts.

t	k	v	λ	Lowest N	Lowest Time
2	10	2	5	26, 2.976, D_5	900, 0.036, B_5
2	10	2	10	44, 10.855, D_{10}	1800, 0.069, B_{10}
2	10	3	5	56, 27.120, D_5	2025, 0.083, B_5
2	20	2	5	31, 23.994, D_5	3800, 0.562, B_5
2	20	2	10	49, 97.389, D_{10}	7600, 1.107, B_{10}
3	10	2	5	48, 34.623, D_5	4800, 0.473, B_5
4	10	2	5	121, 262.64, D_5	16800, 3.077, B_5

Table 3: Individuals composed of a single stage with the smallest number of rows N and smallest time T for the same parameter situations as in Table 2.

Furthermore, at Generation 10 and after (apart from two exceptions), the index for the first algorithm is always 1. The index for most of the later stages is 1, or small relative to the index, apart from a few exceptions. This confirms our intuition, in that a low-index first stage is computationally efficient, produces a small number of rows, and simultaneously covers many interactions.

Note that the genetic algorithm has very much the same effect no matter if t increases, k increases, v increases, or λ increases. However, the extent to which the computational time decreases appears to be smaller as t increases, such as for the $t = 4$ situation, where the time of the “Lowest N ” individual only decreases by 68.2%.

An interesting point to discuss relates not to the genetic algorithm, but rather to the experiment with a single stage, in Table 3. Observe the last row, with parameters $t = 4, k = 10, v = 2, \lambda = 5$: the “Lowest N ” individual had 121 rows, completed in 262.640 seconds, and the selection was $MS(D(5))$. The density algorithm is computationally intensive, but here we actually found that this algorithm did not take the longest time. In fact, $MS(L(5))$ produced 740 rows in over 430 seconds, and $MS(S(5))$ produced 137 rows in over 1000 seconds! An explanation is that as the strength t increases, then the graphs being constructed for L and S are very large. Since there is only one stage, the entire graph on $\binom{k}{t}v^t$ vertices is created, one for each interaction. Note that D also maintains a structure for all of these interactions. But the higher-index formulation of the edges imply that the graph is very highly connected, whereas D only has to maintain a constant times the number of interactions. Much of the computation for L and S was dedicated to managing the (large) graph all at once, since the size of the graph was dominated by the number of edges, not vertices.

5.1 Genetic Operators Discussion

It is clear from the results in Table 2 that the two most important operations towards achieving a smaller covering array faster are Append and Modify, since most of the most fit individuals found have many stages with small indexes.

There appears to be a correlation between these two operators, for the sole reason that there are many instances of the same algorithm with the same index being repeated. For example, for $t = 2, k = 10, v = 2, \lambda = 10$, in Generation 50, the “Lowest Time” individual has six occurrences of S with index 1 repeated. It would be far simpler to use one instance of S with index 6.

We explain why this is the case for each of the algorithms in turn. For B, there is no heuristic being calculated as rows are generated, since each interaction currently uncovered is put into its own row(s). Because there is no additional calculation occurring, B is fast but produces far too many rows. By having multiple adjacent occurrences of B with index 1, there is a small “heuristic” being calculated, which happens between the stages in determining which interactions are left. Since this cost is far less after the first stage is complete, we can now see why this occurs for B.

The explanation for L and S is nearly identical. The advantage of coloring algorithms occurs most when the graph is “sparsely” connected; if there are only a few edges comparatively, then the choices made for each of the algorithms will produce a coloring that is closer to the optimal number. Suppose we are observing a coloring stage with a relatively large index, and suppose that our current index is α , with the target index of β . Then a *clique* (i.e., a set of vertices with every pair connected via an edge) is formed among *all* interactions with the same column set and value set, but different index. Therefore, when the index of a stage increases, the graph becomes much more connected than with one stage. Like the case with B before, a stage with high target index has no heuristic other than the choice of vertices, but by having multiple repeated stages with index 1 of the same algorithm, there is a “heuristic” created between the stages.

For D, the algorithm is greedy in that it chooses the “best” symbol in each position of a row being generated. As is shown in [2], the rate of coverage decreases as the number of generated rows increase. If we have a stage for density with high target index, then this phenomenon certainly occurs: we saw this in Figure 2. By having multiple stages of D with index 1, and recalculating the interactions left to cover after a stage is complete, we can now see why multiple sequential stages of D improve over a single stage with high index.

The crossover operator did not prove to be nearly as powerful as the mentioned mutation operators. We can confirm this by a similar experiment that we performed that was mutation-only (i.e., no crossover), and it eventually produced individuals that were either the same, or nearly equivalent, to those shown in Table 2. However, we had to extend the algorithm to nearly 200 generations before these individuals were found. Therefore, we can conclude that the crossover operator was helpful only in improving the number of generations in the genetic algorithm to find “very fit” individuals.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we developed a genetic algorithm that constructs covering arrays of higher index using a sequence of deterministic algorithms. This algorithm is a generalization of existing methods but simultaneously addresses the question of redundancy in interaction testing, which has not been examined in any publication as far as we are aware. We believe that an exploration of redundancy with

covering arrays and similar objects will lead to improved testing practices in systems that are inherently not deterministic.

As a result of our experiments, we can conclude the following different avenues for the construction of covering arrays of higher index. A very good solution can be obtained with all indexes set to 1, with the first algorithm dependent on the desired goal.

- If the goal is mostly computational efficiency and very little about N , have the first algorithm be B, followed by any choice of greedy graph coloring algorithms. Any index can be chosen here, but a lower N is reached with very little additional time if the first index is 1.
- If the goal is a balance between computational efficiency and N , then have most algorithms be greedy graph coloring ones, either L or S. Whichever algorithms are chosen, have B not be in any stage.
- If the goal is to minimize N , choose either S or D as the first stage, as well as several choices of D in subsequent stages. Whichever algorithms are chosen, have B not be in any stage.

Our future work involves introducing randomness into the framework. Although random methods do not have guarantees about the size of the produced array, they are often much faster, and the number of rows in practice is often close or better than deterministic algorithms. We plan to investigate incorporating randomness into the fitness function while also keeping the simulation time low, which was a distinct advantage in our approach of only having deterministic algorithms. One direction that we plan to investigate is to have the framework start with a randomly generated array, have it remain fixed, and then find a sequence of stages starting from this array (our algorithm developed here starts without any rows built).

One aspect of our algorithm worth observing is that it is absolute, in that if a number of rows N and time T are observed for an individual I such that T is the minimum time in the distribution, then it is reported in Table 2. However Figures 4 to 6 show examples of two points (N_1, T_1) and (N_2, T_2) with $T_1 < T_2$, but N_1 is much greater than N_2 . In other words, the time decrease is very small, but the number of rows increase is very large. It would be worth exploring an alteration to the Pareto front calculation that picks the point with minimum N and “small” T such that the time distance between it and the global minimum time is within a small ϵ factor.

Another future work item considers exact methods, such as satisfiability and constraint programming. While they guarantee to produce the smallest number of rows possible in a given stage and are deterministic, they are far slower than any of the algorithms we have tested. Furthermore, since most of the individuals we have found here have each stage representing a small index, we believe that each algorithm (other than B) produced an array very close to the optimal size for a given stage. It would be interesting, however, to see if such exact methods would be useful at a very late stage. This is mainly because most deterministic methods often cover many interactions early, leaving a small number of them remaining, and often having trouble covering them in a small number of rows. Further, since the number of interactions left is small, the time needed to run these exact methods is far reduced.

REFERENCES

- [1] Noga Alon and Joel H. Spencer. 2004. *The probabilistic method*. John Wiley & Sons.
- [2] Renée C. Bryce and Charles J. Colbourn. 2007. One-test-at-a-time heuristic search for interaction test suites. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation*. 1082–1089.
- [3] Renée C. Bryce and Charles J. Colbourn. 2009. A density-based greedy algorithm for higher strength covering arrays. *Software Testing, Verification and Reliability* 19, 1 (2009), 37–53.
- [4] Charles J. Colbourn. 2004. Combinatorial aspects of covering arrays. *Le Matematiche (Catania)* 58 (2004), 121–167.
- [5] Charles J. Colbourn. 2011. Covering arrays and hash families. *NATO Science for Peace and Security Series, D: Information and Communication Security* 29, Information Security, Coding Theory and Related Combinatorics (2011), 99–135. <https://doi.org/10.3233/978-1-60750-663-8-99>
- [6] Siddhartha R. Dalal, Ashish Jain, Nachimuthu Karunanithi, J.M. Leaton, Christopher M. Lott, Gardner C. Patton, and Bruce M. Horowitz. 1999. Model-based testing in practice. In *Proceedings of the 21st international conference on Software engineering*. 285–294.
- [7] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.
- [8] Ryan Dougherty. 2019. *Hash Families and Applications to t -Restrictions*. Ph.D. Dissertation. Arizona State University.
- [9] Ryan E. Dougherty, Erin Lanus, Charles J. Colbourn, and Stephanie Forrest. 2019. Genetic Algorithms for Affine Transformations to Existential t -restrictions. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '19)*. ACM, New York, NY, USA, 1707–1708. <https://doi.org/10.1145/3319619.3326823>
- [10] Daniel J Kleitman and Joel Spencer. 1973. Families of k -independent sets. *Discrete Mathematics* 6, 3 (1973), 255–262.
- [11] D. Richard Kuhn and Michael J. Reilly. 2002. An investigation of the applicability of design of experiments to software testing. In *27th Annual NASA Goddard/IEEE Software Engineering Workshop Proceedings*. IEEE, 91–95.
- [12] D Richard Kuhn, Dolores R Wallace, and Albert M Gallo. 2004. Software fault interactions and implications for software testing. *IEEE transactions on software engineering* 30, 6 (2004), 418–421.
- [13] Kevin Leach, Ryan Dougherty, Chad Spensky, Stephanie Forrest, and Westley Weimer. 2019. Evolutionary computation for improving malware analysis. In *2019 IEEE/ACM International Workshop on Genetic Improvement (GI)*. IEEE, 18–19.
- [14] Kari J Nurmela. 2004. Upper bounds for covering arrays by tabu search. *Discrete applied mathematics* 138, 1-2 (2004), 143–152.
- [15] Kaushik Sarkar and Charles J Colbourn. 2017. Upper bounds on the size of covering arrays. *SIAM Journal on Discrete Mathematics* 31, 2 (2017), 1277–1293.
- [16] Kaushik Sarkar and Charles J. Colbourn. 2019. Two-stage algorithms for covering array construction. *Journal of Combinatorial Designs* 27, 8 (2019), 475–505. <https://doi.org/10.1002/jcd.21657>
- [17] Kaushik Sarkar, Charles J Colbourn, Annalisa De Bonis, and Ugo Vaccaro. 2016. Partial covering arrays: algorithms and asymptotics. In *International Workshop on Combinatorial Algorithms*. Springer, 437–448.
- [18] John Stardom. 2001. *Metaheuristics and the search for covering and packing arrays*. Master's thesis. Simon Fraser University.
- [19] Jimena Adriana Timaná-Peña, Carlos Alberto Cobos-Lozada, and Jose Torres-Jimenez. 2016. Metaheuristic algorithms for building Covering Arrays: A review. *Revista Facultad de Ingeniería* 25, 43 (2016), 31–45.
- [20] Cemal Yilmaz, Myra B Cohen, and Adam A Porter. 2006. Covering arrays for efficient fault characterization in complex configuration spaces. *IEEE Transactions on Software Engineering* 32, 1 (2006), 20–34.