Branching Processes for QuickCheck Generators

Agustín Mista
Universidad Nacional de Rosario
Rosario, Argentina
amista@dcc.fceia.unr.edu.ar

Alejandro Russo
Chalmers University of Technology
Gothenburg, Sweden
russo@chalmers.se

John Hughes
Chalmers University of Technology
Gothenburg, Sweden
rjmh@chalmers.se

Abstract

In *QuickCheck* (or, more generally, random testing), it is challenging to control random data generators' distributions—specially when it comes to *user-defined algebraic data types* (ADT). In this paper, we adapt results from an area of mathematics known as *branching processes*, and show how they help to analytically predict (at compile-time) the expected number of generated constructors, even in the presence of mutually recursive or composite ADTs. Using our probabilistic formulas, we design heuristics capable of automatically adjusting probabilities in order to synthesize generators which distributions are aligned with users' demands. We provide a Haskell implementation of our mechanism in a tool called DRAGEN and perform case studies with real-world applications. When generating random values, our synthesized *QuickCheck* generators show improvements in code coverage when compared with those automatically derived by state-of-the-art tools.

CCS Concepts • **Software and its engineering → Software testing and debugging**;

Keywords Branching process, QuickCheck, Testing, Haskell

1 Introduction

Random property-based testing is an increasingly popular approach to finding bugs [3, 16, 17]. In the Haskell community, *QuickCheck* [9] is the dominant tool of this sort. *QuickCheck* requires developers to specify *testing properties* describing the expected software behavior. Then, it generates a large number

of random *test cases* and reports those violating the testing properties. *QuickCheck* generates random data by employing *random test data generators* or *QuickCheck* generators for short. The generation of test cases is guided by the *types* involved in the testing properties. It defines default generators for many built-in types like booleans, integers, and lists. However, when it comes to user-defined ADTs, developers are usually required to specify the generation process. The difficulty is, however, that it might become intricate to define generators so that they result in a suitable distribution or enforce data invariants.

The state-of-the-art tools to derive generators for user-defined ADTs can be classified based on the automation level as well as the sort of invariants enforced at the data generation phase. *QuickCheck* and *SmallCheck* [27] (a tool for writing generators that synthesize small test cases) use type-driven generators written by developers. As a result, generated random values are well-typed and preserve the structure described by the ADT. Rather than manually writing generators, libraries *derive* [24] and *MegaDeTH* [13, 14] automatically synthesize generators for a given user-defined ADT. The library *derive* provides no guarantees that the generation process terminates, while *MegaDeTH* pays almost no attention to the distribution of values. In contrast, *Feat* [11] provides a mechanism to uniformly sample values from a given ADT. It enumerates all the possible values of a given ADT so that sampling uniformly from ADTs becomes sampling uniformly from the set of natural numbers. *Feat*'s authors subsequently extend their approach to *uniformly* generate values constrained by user-defined predicates [8]. Lastly, *Luck* is a domain specific language for manually writing *QuickCheck* properties in tandem with generators so that it becomes possible to finely control the distribution of generated values [18].

In this work, we consider the scenario where developers are not fully aware of the properties and invariants that input data must fulfill. This constitutes a valid assumption for *penetration testing* [2], where testers often apply fuzzers in an attempt to make programs crash—an anomaly which might lead to a vulnerability. We believe that, in contrast, if users can recognize specific properties of their systems then it is preferable to spend time writing specialized generators for that purpose (e.g., by using *Luck*) instead of considering automatically derived ones.

Our realization is that *branching processes* [29], a relatively simple stochastic model conceived to study the evolution of populations, can be applied to predict the generation distribution of ADTs' constructors in a simple and automatable

manner. To the best of our knowledge, this stochastic model has not yet been applied to this field, and we believe it may be a promising foundation to develop future extensions. The contributions of this paper can be outlined as follows:

▶ We provide a mathematical foundation which helps to analytically characterize the distribution of constructors in derived *QuickCheck* generators for ADTs.

▶ We show how to use type reification to simplify our prediction process and extend our model to mutually recursive and composite types.

▶ We design (compile-time) heuristics that automatically search for probability parameters so that distributions of constructors can be adjusted to what developers might want.

▶ We provide an implementation of our ideas in the form of a Haskell library[1] called **DRAGEN** ⸙ (the Danish word for *dragon*, here standing for *Derivation of RAndom GENerators*).

▶ We evaluate our tool by generating inputs for real-world programs, where it manages to obtain significantly more code coverage than those random inputs generated by *MegaDeTH*'s generators.

Overall, our work addresses a timely problem with a neat mathematical insight that is backed by a complete implementation and experience on third-party examples.

2   Background

In this section, we briefly illustrate how *QuickCheck* random generators work. We consider the following implementation of binary trees:

data *Tree* = *Leaf$_A$* | *Leaf$_B$* | *Leaf$_C$* | *Node Tree Tree*

In order to help developers write generators, *QuickCheck* defines the *Arbitrary* type-class with the overloaded symbol *arbitrary* :: *Gen a*, which denotes a monadic generator for values of type *a*. Then, to generate random trees, we need to provide an instance of the *Arbitrary* type-class for the type *Tree*. Figure 1 shows a possible implementation. At the top level, this generator simply uses *QuickCheck*'s primitive *oneof* :: [*Gen a*] → *Gen a* to pick a generator from a list of generators with uniform probability. This list consists of a random generator for each possible choice of data constructor of *Tree*. We use *applicative style* [21] to describe each one of them idiomatically. So, *pure Leaf$_A$* is a generator that always generates *Leaf$_A$*s, while *Node* ⟨\$⟩ *arbitrary* ⟨∗⟩ *arbitrary* is a generator that always generates *Node* constructors, "filling" its arguments by calling *arbitrary* recursively on each of them.

---
[1]Available at https://bitbucket.org/agustinmista/dragen

Although it might seem easy, writing random generators becomes cumbersome very quickly. Particularly, if we want to write a random generator for a user-defined ADT *T*, it is also necessary to provide random generators for every user-defined ADT inside of *T* as well! What remains of this section is focused on explaining the state-of-the-art techniques used to *automatically* derive generators for user-defined ADTs via type-driven approaches.

2.1   Library *derive*

The simplest way to automatically derive a generator for a given ADT is the one implemented by the Haskell library *derive* [24]. This library uses Template Haskell [28] to automatically synthesize a generator for the data type *Tree* semantically equivalent to the one presented in Figure 1.

While the library *derive* is a big improvement for the testing process, its implementation has a serious shortcoming when dealing with recursively defined data types: in many cases, there is a non-zero probability of generating a recursive type constructor every time a recursive type constructor gets generated, which can lead to infinite generation loops. A detailed example of this phenomenon is presented in the supplementary material [23]. In this work, we only focus on derivation tools which accomplish terminating behavior, since we consider this an essential component of well-behaved generators.

2.2   *MegaDeTH*

The second approach we will discuss is the one taken by *MegaDeTH*, a meta-programming tool used intensively by *QuickFuzz* [13, 14]. Firstly, *MegaDeTH* derives random generators for ADTs as well as all of its nested types—a useful feature not supported by *derive*. Secondly, *MegaDeTH* avoids potentially infinite generation loops by setting an upper bound to the random generation recursive depth.

Figure 2 shows a simplified (but semantically equivalent) version of the random generator for *Tree* derived by *MegaDeTH*. This generator uses *QuickCheck*'s function *sized* :: (*Int* → *Gen a*) → *Gen a* to build a random generator based on a function (of type *Int* → *Gen a*) that limits the possible recursive calls performed when creating random values. The integer passed to *sized*'s argument is called the *generation size*. When the generation size is zero (see definition *gen* 0), the generator only chooses between the *Tree*'s terminal constructors—thus

---

instance *Arbitrary Tree* **where**
  arbitrary = *oneof* [*pure Leaf$_A$*, *pure Leaf$_B$*, *pure Leaf$_C$*
                   , *Node* ⟨\$⟩ *arbitrary* ⟨∗⟩ *arbitrary*]

Figure 1. Random generator for *Tree*.

---

instance *Arbitrary Tree* **where**
  arbitrary = *sized gen* **where**
    gen 0 = *oneof*
      [*pure Leaf$_A$*, *pure Leaf$_B$*, *pure Leaf$_C$*]
    gen n = *oneof*
      [*pure Leaf$_A$*, *pure Leaf$_B$*, *pure Leaf$_C$*
      , *Node* ⟨\$⟩ *gen* (*div n* 2) ⟨∗⟩ *gen* (*div n* 2)]

Figure 2. *MegaDeTH* generator for *Tree*.

ending the generation process. If the generation size is strictly positive, it is free to randomly generate any *Tree* constructor (see definition *gen n*). When it chooses to generate a recursive constructor, it reduces the generation size for its subsequent recursive calls by a factor that depends on the number of recursive arguments this constructor has (*div n* 2). In this way, *MegaDeTH* ensures that all generated values are finite.

Although *MegaDeTH* generators always terminate, they have a major practical drawback: in our example, the use of *oneof* to uniformly decide the next constructor to be generated produces a generator that generates leaves approximately three quarters of the time (note this also applies to the generator obtained with *derive* from Figure 1). This entails a distribution of constructors heavily concentrated on leaves, with a very small number of complex values with nested nodes, regardless how large the chosen generation size is—see Figure 3 (left).
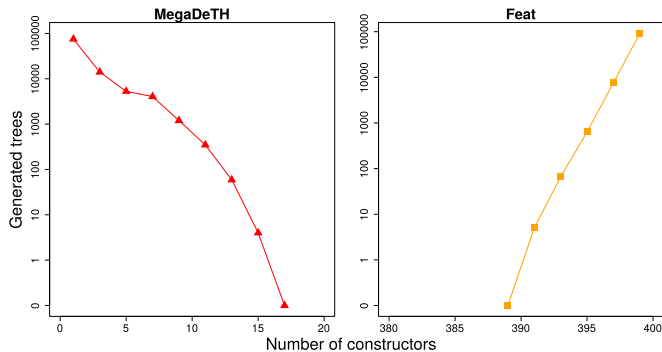
2.3 Feat

The last approach we discuss is *Feat* [11]. This tool determines the distribution of generated values in a completely different way: it uses uniform generation based on an *exhaustive enumeration of all the possible values of the ADTs being considered*. *Feat* automatically establishes a bijection between all the possible values of a given type *T*, and a finite prefix of the natural numbers. Then, it guarantees a *uniform generation over the complete space of values of a given data type T* up to a certain size.[2] However, the distribution of size, given by the number of constructors in the generated values, is highly dependent on the structure of the data type being considered.

Figure 3 (right) shows the overall distribution shape of a *QuickCheck* generator derived using *Feat* for *Tree* using a generation size of 400, i.e., generating values of up to 400 constructors.[3] Notice that all the generated values are close to the maximum size! This phenomenon follows from the exponential growth in the number of possible *Tree*s of *n* constructors as we increase *n*. In other words, the space of *Tree*s up to 400 constructors is composed to a large extent of values with around 400 constructors, and (proportionally) very few with a smaller number of constructors. Hence, a generation process based on uniform generation of a natural number (which thus ignores the structure of the type being generated) is biased very strongly towards values made up of a large number of constructors. In our tests, no *Tree* with less than 390 constructors was ever generated. In practice, this problem can be partially solved by using a variety of generation sizes in order to get more diversity in the generated values. However, to decide which generation sizes are the best choices is not a trivial task either. As consequence, in this work we consider only the case of fixed-size random generation.

---

[2]We avoid including any source code generated by *Feat*, since it works by synthetizing *Enumerable* type-class instances instead of *Arbitrary* ones. Such instances give no insight into how the derived random generators work.

[3] We choose to use this generation size here since it helps us to compare *MegaDeTH* and *Feat* with the results of our tool in Section 8.

Figure 3. Size distribution of 100000 randomly generated *Tree* values using *MegaDeTH* (▲) with generation size 10, and *Feat* (■) with generation size 400.

As we have shown, by using both *MegaDeTH* and *Feat*, the user is tied to the fixed generation distribution that each tool produces, which tends to be highly dependent on the particular data type under consideration on each case. Instead, this work aims to provide a *theoretical framework able to predict and later tune the distributions of automatically derived generators*, giving the user a more flexible testing environment, while keeping it as automated as possible.

3 Simple-Type Branching Processes

Galton-Watson Branching processes (or branching processes for short) are a particular case of Markov processes that model the growth and extinction of populations. Originally conceived to study the extinction of family names in the Victorian era, this formalism has been successfully applied to a wide range of research areas in biology and physics—see the textbook by Haccou et al. [15] for an excellent introduction. In this section, we show how to use this theory to model *QuickCheck*'s distribution of constructors.

We start by analyzing the generation process for the *Node* constructors in the data type *Tree* as described by the generators in Figure 1 and 2. From the code, we can observe that the stochastic process they encode satisfies the following assumptions (which coincide with the assumptions of Galton-Watson branching processes): i) With a certain probability, it starts with some initial *Node* constructor. ii) At any step, the probability of generating a *Node* is not affected by the *Node*s generated before or after. iii) The probability of generating a *Node* is independent of where in the tree that constructor is about to be placed.

The original Galton-Watson process is a simple stochastic process that counts the population sizes at different points in time called *generations*. For our purposes, populations consist of *Node* constructors, and generations are obtained by selecting tree levels.

Figure 4 illustrates a possible generated value. It starts by generating a *Node* constructor at generation (i.e., depth) zero (G_0), then another two *Node* constructors as left and right

subtrees in generation one (G_1), etc. (Dotted edges denote further constructors which are not drawn, as they are not essential for the point being made.) This process repeats until the population of *Node* constructors becomes extinct or stable, or alternatively grows forever.

The mathematics behind the Galton-Watson process allows us to predict the expected number of offspring at the *n*th-generation, i.e., the number of *Node* constructors at depth *n* in the generated tree. Formally, we start by introducing the random variable R to denote the number of *Node* con-

G_0	*Node*	
G_1	*Node*	*Node*
G_2	*Leaf$_A$* *Node*	*Leaf$_B$* *Node*

Figure 4. Generation of *Node* constructors.

structors in the next generation generated by a *Node* constructor in this generation—the R comes from "reproduction" and the reader can think it as a *Node* constructor reproducing *Node* constructors. To be a bit more general, let us consider the *Tree* random generator automatically generated using *derive* (Figure 1), but where the probability of choosing between any constructor is no longer uniform. Instead, we have a p_C probability of choosing the constructor C. These probabilities are external parameters of the prediction mechanism, and Section 7 explains how they can later be instantiated with actual values found by optimization, enabling the user to tune the generated distribution.

We note p_{Leaf} as the probability of generating a leaf of any kind, i.e., $p_{Leaf} = p_{LeafA} + p_{LeafB} + p_{LeafC}$. In this setting, and assuming a parent constructor *Node*, the probabilities of generating R numbers of *Node* offspring in the next generation (i.e., in the recursive calls of *arbitrary*) are as follows:

$$P(R = 0) = p_{Leaf} \cdot p_{Leaf}$$
$$P(R = 1) = p_{Node} \cdot p_{Leaf} + p_{Leaf} \cdot p_{Node} = 2 \cdot p_{Node} \cdot p_{Leaf}$$
$$P(R = 2) = p_{Node} \cdot p_{Node}$$

One manner to understand the equations above is by considering what *QuickCheck* does when generating the subtrees of a given node. For instance, the cases when generating exactly one *Node* as descendant ($P(R = 1)$) occurs in two situations: when the left subtree is a *Node* and the right one is a *Leaf*; and viceversa. The probability for those events to occur is $p_{Node} * p_{Leaf}$ and $p_{Leaf} * p_{Node}$, respectively. Then, the probablity of having exactly one *Node* as a descendant is given by the sum of the probability of both events—the other cases follow a similar reasoning.

Now that we have determined the distribution of R, we proceed to introduce the random variables G_n to denote the population of *Node* constructors in the *n*th generation. We write ξ_i^n for the random variable which captures the number of (offspring) *Node* constructors at the *n*th generation produced by the *i*th *Node* constructor at the *(n-1)*th generation. It is easy

to see that it must be the case that $G_n = \xi_1^n + \xi_2^n + \cdots + \xi_{G_{n-1}}^n$. To deduce $E[G_n]$, i.e. the expected number of *Node*s in the *n*th generation, we apply the (standard) Law of Total Expectation $E[X] = E[E[X|Y]]$ [4] with $X = G_n$ and $Y = G_{n-1}$ to obtain:

$$E[G_n] = E[E[G_n|G_{n-1}]]. \tag{1}$$

By expanding G_n, we deduce that:

$$E[G_n|G_{n-1}] = E[\xi_1^n + \xi_2^n + \cdots + \xi_{G_{n-1}}^n|G_{n-1}]$$
$$= E[\xi_1^n|G_{n-1}] + E[\xi_2^n|G_{n-1}] + \cdots + E[\xi_{G_{n-1}}^n|G_{n-1}]$$

Since ξ_1^n, ξ_2^n, ..., and $\xi_{G_{n-1}}^n$ are *all governed by the distribution captured by the random variable R* (recall the assumptions at the beginning of the section), we have that:

$$E[G_n|G_{n-1}] = E[R|G_{n-1}] + E[R|G_{n-1}] + \cdots + E[R|G_{n-1}]$$

Since R is *independent of the generation where Node constructors decide to generate other Node constructors*, we have that

$$E[G_n|G_{n-1}] = \underbrace{E[R] + E[R] + \cdots + E[R]}_{G_{n-1} \text{ times}} = E[R] \cdot G_{n-1} \tag{2}$$

From now on, we introduce m to denote the mean of R, i.e., *the mean of reproduction*. Then, by rewriting $m = E[R]$, we obtain:

$$E[G_n] \overset{(1)}{=} E[E[G_n|G_{n-1}]] \overset{(2)}{=} E[m \cdot G_{n-1}] \overset{m \text{ is constant}}{=} E[G_{n-1}] \cdot m$$

By unfolding this recursive equation many times, we obtain:

$$E[G_n] = E[G_0] \cdot m^n \tag{3}$$

As the equation indicates, the expected number of *Node* constructors at the *n*th generation is affected by the mean of reproduction. Although we obtained this intuitive result using a formalism that may look overly complex, it is useful to understand the methodology used here. In the next section, we will derive the main result of this work following the same reasoning line under a more general scenario.
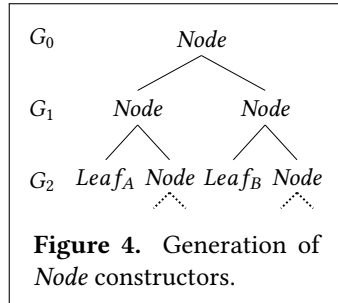
We can now also predict the total expected number of individuals *up to* the *n*th generation. For that purpose, we introduce the random variable P_n to denote the population of *Node* constructors up to the *n*th generation. It is then easy to see that $P_n = \sum_{i=0}^n G_i$ and consequently:

$$E[P_n] = \sum_{i=0}^n E[G_i] \overset{(3)}{=} \sum_{i=0}^n E[G_0] \cdot m^i = E[G_0] \cdot \left(\frac{1 - m^{n+1}}{1 - m} \right) \tag{4}$$

where the last equality holds by the geometric series definition. This is the general formula provided by the Galton-Watson process. In this case, the mean of reproduction for *Node* is given by:

$$m = E[R] = \sum_{k=0}^2 k \cdot P(R = k) = 2 \cdot p_{Node} \tag{5}$$

---

[4] $E[X|Y]$ is a function on the random variable Y, i.e., $E[X|Y]y = E[X|Y = y]$ and therefore it is a random variable itself. In this light, the law says that if we observe the expectations of X given the different y_s, and then we do the expectation of all those values, then we have the expectation of X.

By (4) and (5), the expected number of *Node* constructors up to generation n is given by the following formula:

$$E[P_n] = E[G_0] \cdot \left(\frac{1 - m^{n+1}}{1-m} \right) = p_{Node} \cdot \left(\frac{1 - (2 \cdot p_{Node})^{n+1}}{1 - 2 \cdot p_{Node}} \right)$$

If we apply the previous formula to predict the distribution of constructors induced by *MegaDeTH* in Figure 2, where $p_{LeafA} = p_{LeafB} = p_{LeafC} = p_{Node} = 0.25$, we obtain an expected number of *Node* constructors up to level 10 of 0.4997, which denotes a distribution highly biased towards small values, since we can only produce further subterms by producing *Node*s. However, if we set $p_{LeafA} = p_{LeafB} = p_{LeafC} = 0.1$ and $p_{Node} = 0.7$, we can predict that, as expected, our general random generator will generate much bigger trees, containing an average number of 69.1173 *Node*s up to level 10! Unfortunately, we cannot apply this reasoning to predict the distribution of constructors for derived generators for ADTs with more than one non-terminal constructor. For instance, let us consider the following data type definition:

data *Tree′* = *Leaf* | *Node_A Tree′ Tree′* | *Node_B Tree′*

In this case, we need to separately consider that a $Node_A$ can generate not only $Node_A$ but also $Node_B$ offspring (similarly with $Node_B$). A stronger mathematical formalism is needed. The next section explains how to predict the generation of this kind of data types by using an extension of Galton-Waston processes known as *multi-type branching processes*.

4 Multi-Type Branching Processes

In this section, we present the basis for our main contribution: *the application of multi-type branching processes to predict the distribution of constructors*. We will illustrate the technique by considering the *Tree′* ADT that we concluded with in the previous section.

Before we dive into technicalities, Figure 5 shows the automatically derived generator for *Tree′* that our tool produces. Our generators depend on the (possibly) different probabilities that constructors have to be generated—variables p_{Leaf}, p_{NodeA}, and p_{NodeB}. These probabilities are used by the function *chooseWith* :: [(*Double*, *Gen a*)] → *Gen a*, which picks a random generator of type *a* with an explicitly given probability from a list. This function can be easily expressed by using *QuickCheck*'s primitive operations and therefore we omit its

---

instance *Arbitrary Tree′* **where**
  arbitrary = *sized gen* **where**
    gen 0 = *pure Leaf*
    gen n = *chooseWith*
      [(p_{Leaf} , *pure Leaf*)
      , (p_{NodeA}, *Node_A* ⟨$⟩ *gen* (*n*−1) ⟨∗⟩ *gen* (*n*−1))
      , (p_{NodeB}, *Node_B* ⟨$⟩ *gen* (*n*−1))]

Figure 5. ᗡᖇᴀGᴇN generator for *Tree′*

---

implementation. Additionally note that, like *MegaDeTH*, our generators use *sized* to limit the number of recursive calls to ensure termination. We note that the theory behind branching processes is able to predict the termination behavior of our generators and we could have used this ability to ensure their termination without the need of a depth limiting mechanism like *sized*. However, using *sized* provides more control over the obtained generator distributions.

To predict the distribution of constructors provided by ᗡᖇᴀGᴇN generators, we introduce a generalization of the previous Galton-Watson branching process called multi-type Galton-Watson branching process. This generalization allows us to consider several *kinds of individuals*, i.e., constructors in our setting, to procreate (generate) different *kinds of offspring* (constructors). Additionally, this approach allows us to consider not just one constructor, as we did in the previous section, but rather to consider all of them at the same time.

Before we present the mathematical foundations, which follow a similar line of reasoning as that in Section 3, Figure 6 illustrates a possible generated value of type *Tree′*.

In the generation process, it is assumed that *the kind (i.e., the constructor) of the parent might affect the probabilities of reproducing (generating) offspring of a certain kind.* Observe that this is the case for a wide range of derived ADT generators, e.g., choosing a terminal constructor (e.g., *Leaf*) affects the probabilities of generat-
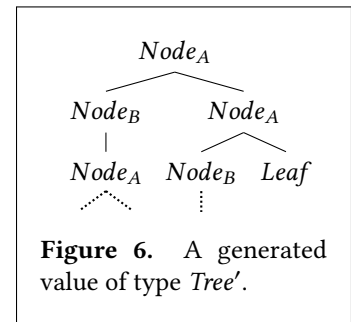

Figure 6. A generated value of type *Tree′*.

ing non-terminal ones (by setting them to zero). The population at the nth generation is then characterized as a vector of random variables $G_n = (G_n^1, G_n^2, \cdots, G_n^d)$, where d is the number of different kinds of constructors. Each random variable G_n^i captures the number of occurrences of the ith-constructor of the ADT at the nth generation. Essentially, G_n "groups" the population at level n by the constructors of the ADT. By estimating the expected shape of the vector G_n, it is possible to obtain the expected number of constructors at the nth generation. Specifically, we have that $E[G_n] = (E[G_n^1], E[G_n^2], \cdots, E[G_n^d])$. To deduce $E[G_n]$, we focus on deducing each component of the vector.

As explained above, the reproduction behavior is determined by the kind of the individual. In this light, we introduce random variable R_{ij} to denote a parent ith constructor reproducing a jth constructor. As we did before, we apply the equation $E[X] = E[E[X|Y]]$ with $X = G_n^j$ and $Y = G_{n-1}$ to obtain $E[G_n^j] = E[E[G_n^j|G_{n-1}]]$. To calculate the expected number of jth constructors at the level n produced by the constructors present at level $(n-1)$, i.e., $E[G_n^j|G_{(n-1)}]$, it is enough to count the expected number of children of kind j produced by the

different parents of kind i, i.e., $E[R_{ij}]$, times the amount of parents of kind i found in the level $(n-1)$, i.e., $G^i_{(n-1)}$. This result is expressed by the following equation marked as $(\star)$, and is formally verified in the supplementary material.

$$E[G^j_n|G_{n-1}] \overset{(\star)}{=} \sum_{i=1}^{d} G^i_{(n-1)} \cdot E[R_{ij}] = \sum_{i=1}^{d} G^i_{(n-1)} \cdot m_{ij} \quad (6)$$

Similarly as before, we rewrite $E[R_{ij}]$ as m_{ij}, which now represents a single expectation of reproduction indexed by the kind of both the parent and child constructor.

Mean matrix of constructors In the previous section, m was the expectation of reproduction of a single constructor. Now we have m_{ij} as the expectation of reproduction indexed by the parent and child constructor. In this light, we define M_C, the *mean matrix of constructors* (or mean matrix for simplicity) such that each m_{ij} stores the expected number of jth constructors generated by the ith constructor. M_C is a parameter of the Galton-Watson multi-type process and can be built at compile-time using statically known type information. We are now able to deduce $E[G^j_n]$.

$$E[G^j_n] = E[E[G^j_n|G_{n-1}]] \overset{(6)}{=} E\left[\sum_{i=1}^{d} G^i_{(n-1)} \cdot m_{ij}\right]$$

$$= \sum_{i=1}^{d} E[G^i_{(n-1)} \cdot m_{ij}] = \sum_{i=1}^{d} E[G^i_{(n-1)}] \cdot m_{ij}$$

Using this last equation, we can rewrite $E[G_n]$ as follows.

$$E[G_n] = \left(\sum_{i=1}^{d} E[G^1_{(n-1)}] \cdot m_{i1}, \cdots, \sum_{i=1}^{d} E[G^d_{(n-1)}] \cdot m_{id}\right)$$

By linear algebra, we can rewrite the vector above as the matrix multiplication $E[G_n]^T = E[G_{n-1}]^T \cdot M_C$. By repeatedly unfolding this definition, we obtain that:

$$E[G_n]^T = E[G_0]^T \cdot (M_C)^n \quad (7)$$

This equation is a generalization of (3) when considering many constructors. As we did before, we introduce a random variable $P_n = \sum_{i=0}^{n} G_i$ to denote the population up to the nth generation. It is now possible to obtain the expected population of all the constructors but in a clustered manner:

$$E[P_n]^T = E\left[\sum_{i=0}^{n} G_i\right]^T = \sum_{i=0}^{n} E[G_i]^T \overset{(7)}{=} \sum_{i=0}^{n} E[G_0]^T \cdot (M_C)^n \quad (8)$$

It is possible to write the resulting sum as the closed formula:

$$E[P_n]^T = E[G_0]^T \cdot \left(\frac{I - (M_C)^{n+1}}{I - M_C}\right) \quad (9)$$

where I represents the identity matrix of the appropriate size. Note that equation (9) only holds when $(I - M_C)$ is nonsingular, however, this is the usual case. When $(I - M_C)$ is singular, we resort to using equation (8) instead. Without losing generality, and for simplicity, we consider equations (8)

and (9) as interchangeable. They are the general formulas for the Galton-Watson multi-type branching processes.

Then, to predict the distribution of our *Tree'* data type example, we proceed to build its mean matrix M_C. For instance, the mean number of *Leaf*s generated by a $Node_A$ is:

$$
\begin{aligned}
m_{Node_A, Leaf} \quad = \quad & \underbrace{1 \cdot p_{Leaf} \cdot p_{Node_A} + 1 \cdot p_{Leaf} \cdot p_{Node_B}}_{\textit{One Leaf as left-subtree}} \\
+ \quad & \underbrace{1 \cdot p_{Node_A} \cdot p_{Leaf} + 1 \cdot p_{Node_B} \cdot p_{Leaf}}_{\textit{One Leaf as right-subtree}} \\
+ \quad & \underbrace{2 \cdot p_{Leaf} \cdot p_{Leaf}}_{\textit{Leaf as left- and right-subtree}} \\
= \quad & 2 \cdot p_{Leaf} \quad (10)
\end{aligned}
$$

The rest of M_C can be similarly computed, obtaining:

$$M_C = \begin{array}{c} \\ Leaf \\ Node_A \\ Node_B \end{array} \begin{array}{c} Leaf \qquad Node_A \qquad Node_B \\ \left[\begin{array}{ccc} 0 & 0 & 0 \\ 2 \cdot p_{Leaf} & 2 \cdot p_{Node_A} & 2 \cdot p_{Node_B} \\ p_{Leaf} & p_{Node_A} & p_{Node_B} \end{array}\right] \end{array} \quad (11)$$

Note that the first row, corresponding to the *Leaf* constructor, is filled with zeros. This is because *Leaf* is a terminal constructor, i.e., it cannot generate further subterms of any kind.[5]

With the mean matrix in place, we define $E[G_0]$ (the initial vector of mean probabilities) as $(p_{Leaf}, p_{Node_A}, p_{Node_B})$. By applying (9) with $E[G_0]$ and M_C, we can predict the expected number of generated *non-terminal* $Node_A$ constructors (and analogously $Node_B$) with a size parameter n as follows:

$$E[\text{Node}_A] = \left(E[P_{n-1}]^T\right).\text{Node}_A = \left(E[G_0]^T \cdot \left(\frac{I - (M_C)^n}{I - M_C}\right)\right).\text{Node}_A$$

Function $(\_).C$ simply projects the value corresponding to constructor C from the population vector. It is very important to note that the sum only includes the population up to level $(n - 1)$. This choice comes from the fact that our *QuickCheck* generator can choose between only terminal constructors at the last generation level (recall that *gen 0* generates only *Leaf*s in Figure 5). As an example, if we assign our generation probabilities for *Tree'* as $p_{Leaf} \mapsto 0.2$, $p_{Node_A} \mapsto 0.5$ and $p_{Node_B} \mapsto 0.3$, then the formula predicts that our *QuickCheck* generator with a size parameter of 10 will generate on average 21.322 $Node_A$s and 12.813 $Node_B$s. This result can easily be verified by sampling a large number of values with a generation size of 10, and then averaging the number of generated $Node_A$s and $Node_B$s across the generated values.

In this section, we obtain a prediction of the expected number of non-terminal constructors generated by DRAGEN generators. To predict terminal constructors, however, requires a special treatment as discussed in the next section.

---

[5]The careful reader may notice that there is a pattern in the mean matrix if inspected together with the definition of *Tree'*. We prove in Section 6 that each m_{ij} can be automatically calculated by simply exploiting type information.

5   Terminal constructors

In this section we introduce the special treatment required to predict the generated distribution of terminal constructors, i.e. constructors with no recursive arguments.

Consider the generator in Figure 5. It generates terminal constructors in two situations, i.e., in the definition of *gen* 0 and *gen n*. In other words, the random process introduced by our generators can be considered to be composed of two independent parts when it comes to terminal constructors—refer to the supplementary material for a graphical interpretation. In principle, the number of terminal constructors generated by the stochastic process described in *gen n* is captured by the multi-type branching process formulas. However, to predict the expected number of terminal constructors generated by exercising *gen* 0, we need to separately consider a random process that *only generates terminal constructors* in order to terminate. For this purpose, and assuming a maximum generation depth n, we need to calculate the number of terminal constructors required to stop the generation process at the recursive arguments of each non-terminal constructor at level $(n-1)$. In our *Tree'* example, this corresponds to two *Leaf*s for every $Node_A$ and one *Leaf* for every $Node_B$ constructor at level $(n - 1)$.

Since both random processes are independent, to predict the overall expected number of terminal constructors, we can simply add the expected number of terminal constructors generated in each one of them. Recalling our previous example, we obtain the following formula for *Tree'* terminals as follows:

$$E[Leaf] = \underbrace{\left(E[P_{n-1}]^T\right).\text{Leaf}}_{\text{branching process}} + \underbrace{2\cdot\left(E[G_{n-1}]^T\right).\text{Node}_A}_{\text{case }(Node_A\ Leaf\ Leaf)}$$
$$+ \underbrace{1\cdot\left(E[G_{n-1}]^T\right).\text{Node}_B}_{\text{case }(Node_B\ Leaf)}$$

The formula counts the *Leaf*s generated by the multi-type branching process up to level $(n - 1)$ and adds the expected number of *Leaf*s generated at the last level.

Although we can now predict the expected number of generated *Tree'* constructors regardless of whether they are terminal or not, this approach only works for data types with a single terminal constructor. If we have a data type with multiple terminal constructors, we have to consider the probabilities

---

instance *Arbitrary Tree''* **where**
  arbitrary = *sized gen* **where**
    gen 0 = *chooseWith*
      $[(p^*_{LeafA}, pure\ Leaf_A), (p^*_{LeafB}, pure\ Leaf_B)]$
    gen n = *chooseWith*
      $[(p_{LeafA},\ pure\ Leaf_A), (p_{LeafB}, pure\ Leaf_B)$
      $,(p_{NodeA}, Node_A\ \langle\$\rangle\ gen\ (n{-}1)\ \langle*\rangle\ gen\ (n{-}1))$
      $,(p_{NodeB}, Node_B\ \langle\$\rangle\ gen\ (n{-}1))]$

Figure 7. Derived generator for *Tree''*

---

of choosing each one of them when filling the recursive arguments of non-terminal constructors at the previous level. For instance, consider the following ADT:

data $Tree'' = Leaf_A\,|\,Leaf_B\,|\,Node_A\ Tree''\ Tree''\,|\,Node_B\ Tree''$

Figure 7 shows the corresponding DRAGEN generator for *Tree''*. Note there are two sets of probabilities to choose terminal nodes, one for each random process. The p^*_{LeafA} and p^*_{LeafB} probabilities are used to choose between terminal constructors at the last generation level. These probabilities preserve the same proportion as their non-starred versions, i.e., they are normalized to form a probability distribution:

$$p^*_{LeafA} = \frac{p_{LeafA}}{p_{LeafA} + p_{LeafB}} \qquad p^*_{LeafB} = \frac{p_{LeafB}}{p_{LeafA} + p_{LeafB}}$$

In this manner, we can use the same generation probabilities for terminal constructors in both random processes—therefore reducing the complexity of our prediction engine implementation (described in Section 7).

To compute the overall expected number of terminals, we need to predict the expected number of terminal constructors at the last generation level which could be descendants of non-terminal constructors at level $(n - 1)$. More precisely:

$$E[Leaf_A] = \underbrace{\left(E[P_{n-1}]^T\right).\text{Leaf}_A}_{\text{branching process}} + \underbrace{2\cdot p^*_{LeafA}\cdot\left(E[G_{n-1}]^T\right).\text{Node}_A}_{\text{expected leaves to fill }Node_A s}$$
$$+ \underbrace{1\cdot p^*_{LeafA}\cdot\left(E[G_{n-1}]^T\right).\text{Node}_B}_{\text{expected leaves to fill }Node_B s}$$

where the case of $E[Leaf_B]$ follows analogously.

6   Mutually-recursive and composite ADTs

In this section, we introduce some extensions to our model that allow us to derive DRAGEN generators for data types found in existing off-the-shelf Haskell libraries. We start by showing how multi-type branching processes naturally extend to mutually-recursive ADTs. Consider the mutually recursive ADTs T_1 and T_2 with their automatically derived generators shown in Figure 8. Note the use of the *QuickCheck*'s function *resize*::$Int \rightarrow Gen\ a \rightarrow Gen\ a$, which resets the generation size of a given generator to a new value. We use it to decrement the generation size at the recursive calls of *arbitrary* that generate subterms of a mutually recursive data type.

The key observation is that we can ignore that A, B, C and D are *constructors belonging to different data types* and just consider each of them as a kind of offspring on its own. Figure 9 visualizes the possible offspring generated by the non-terminal constructor B (belonging to T_1) with the corresponding probabilities as labeled edges. Following the figure, we obtain the expected number of Ds generated by B constructors as follows:

$$m_{BD} = 1 \cdot p_A \cdot p_D + 1 \cdot p_B \cdot p_D = p_D \cdot (p_A + p_B) = p_D$$

```
    data T₁ = A | B T₁ T₂
    data T₂ = C | D T₁

    instance Arbitrary T₁ where
       arbitrary = sized gen where
          gen 0 = pure A
          gen n = chooseWith
             [ (p_A, pure A)
             , (p_B, B ⟨$⟩ gen (n−1) ⟨∗⟩ resize (n−1) arbitrary) ]

    instance Arbitrary T₂ where
       arbitrary = sized gen where
          gen 0 = pure C
          gen n = chooseWith
             [ (p_C, pure C), (p_D, D ⟨$⟩ resize (n−1) arbitrary) ]
```

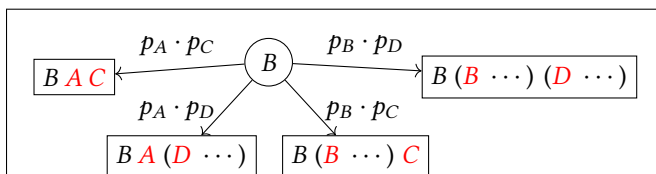**Figure 8.** Mutually recursive types T_1 and T_2 and their DRAGEN generators.

Doing similar calculations, we obtain the mean matrix M_C for A, B, C, and D as follows:

$$
M_C = \begin{array}{c} \\ A \\ B \\ C \\ D \end{array}
\begin{array}{c}
\begin{array}{cccc} A & B & C & D \end{array} \\
\left[\begin{array}{cc|cc}
0 & 0 & 0 & 0 \\
p_A & p_B & p_C & p_D \\
\hline
0 & 0 & 0 & 0 \\
p_A & p_B & 0 & 0
\end{array} \right]
\end{array}
\tag{12}
$$

We define the mean of the initial generation as $E[G_0] = (p_A, p_B, 0, 0)$—we assing $p_C = p_D = 0$ since we choose to start by generating a value of type T_1. With M_C and $E[G_0]$ in place, we can apply the equations explained through Section 4 to predict the expected number of A, B, C and D constructors.

While this approach works, it completely ignores the types T_1 and T_2 when calculating M_C! For a large set of mutually-recursive data types involving a large number of constructors, handling M_C like this results in a high computational cost. We show next how we cannot only shrink this mean matrix of constructors but also compute it automatically by making use of data type definitions.

Mean matrix of types   If we analyze the mean matrices of *Tree'* (11) and the mutually-recursive types T_1 and T_2 (12), it seems that determining the expected number of offspring generated by a non-terminal constructor requires us to *count*

Figure 9. Possible offspring of constructor B.

the number of occurrences in the ADT which the offspring belongs to. For instance, $m_{Node_A, Leaf}$ is $2 \cdot p_{Leaf}$ (10), where 2 is the number of occurrences of *Tree'* in the declaration of $Node_A$. Similarly, m_{BD} is $1 \cdot p_D$, where 1 is the number of occurrences of T_2 in the declaration of B. This observation means that instead of dealing with constructors, we could directly deal with types!

We can think about a branching process as generating "place holders" for constructors, where place holders can only be populated by constructors of a certain type.

Figure 10 illustrates offspring as types for the definitions T_1, T_2, and *Tree'*. A place holder of type T_1 can generate a place holder for type T_1 and a place holder for type T_2. A place holder of type T_2 can generate a place holder of type T_1. A place holder of type *Tree'* can generate two place holders of type *Tree'* when generating $Node_A$, one place holder when generating $Node_B$, or zero place holders when generating a *Leaf* (this last case is not shown in the figure since it is void). With these considerations, the mean matrices of types for *Tree'*, written $M_{Tree'}$; and types T_1 and T_2, written $M_{T_1T_2}$ are defined as follows:
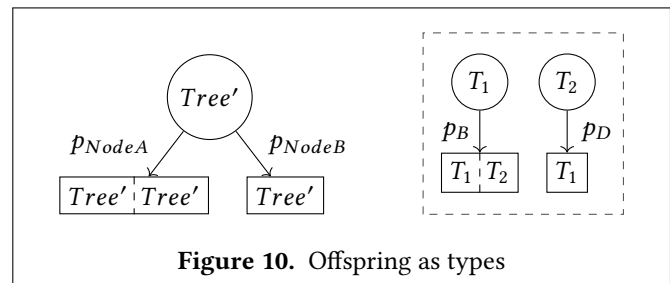
$$
M_{Tree'} = \begin{array}{c} \\ Tree' \end{array}
\begin{array}{c}
\begin{array}{c} Tree' \end{array} \\
\left[2 \cdot p_{Node_A} + p_{Node_B} \right]
\end{array}
\qquad
M_{T_1T_2} = \begin{array}{c} \\ T_1 \\ T_2 \end{array}
\begin{array}{c}
\begin{array}{cc} T_1 & T_2 \end{array} \\
\left[\begin{array}{cc} p_B & p_B \\ p_D & 0 \end{array} \right]
\end{array}
$$

Note how $M_{Tree'}$ shows that the mean matrices of types might reduce a multi-type branching process to a simple-type one.

Having the type matrix in place, we can use the following equation (formally stated and proved in the supplementary material) to soundly predict the expected number of constructors of a given set of (possibly) mutually recursive types:

$$
(E[G_n^C]).C_i^t = (E[G_n^T]).T_t \cdot p_{C_i^t} \qquad\qquad (\forall n \geq 0)
$$

Where G_n^C and G_n^T denotes the nth-generations of constructors and type place holders respectively. C_i^t represents the ith-constructor of the type T_t. The equation establishes that, the expected number of constructors C_i^t at generation n consists of the expected number of type place holders of its type (i.e., T_t) at generation n times the probability of generating that constructor. This equation allows us to simplify many of our calculations above by simply using the mean matrix for types instead of the mean matrix for constructors.

Figure 10. Offspring as types

6.1 Composite types

In this subsection, we extend our approach in a *modular* manner to deal with composite ADTs, i.e., ADTs which use already defined types in their constructors' arguments and which are not involved in the branching process. We start by considering the ADT *Tree* modified to carry booleans at the leaves:

data *Tree = Leaf$_A$ Bool | Leaf$_B$ Bool Bool | · · ·*

Where · · · denotes the constructors that remain unmodified. To predict the expected number of *True* (and analogously of *False*) constructors, we calculate the multi-type branching process for *Tree* and multiply each expected number of leaves by the number of arguments of type *Bool* present in each one:

$$E[True] = p_{True} \cdot (\underbrace{1 \cdot E[Leaf_A]}_{\text{case } Leaf_A} + \underbrace{2 \cdot E[Leaf_B]}_{\text{case } Leaf_B})$$
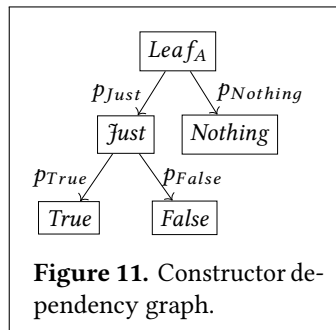
In this case, *Bool* is a ground type like *Int*, *Float*, etc. Predictions become more interesting when considering richer composite types involving, for instance, instantiations of polymorphic types. To illustrate this point, consider a modified version of *Tree* where *Leaf$_A$* now carries a value of type *Maybe Bool*:

data *Tree = Leaf$_A$ (Maybe Bool) | Leaf$_B$ Bool Bool | · · ·*

In order to calculate the expected number of *True*s, now we need to consider the cases that a value of type *Maybe Bool actually carries a boolean value*, i.e., when a *Just* constructor gets generated:

$$E[True] = p_{True} \cdot (1 \cdot E[Leaf_A] \cdot p_{Just} + 2 \cdot E[Leaf_B])$$

In the general case, for constructor arguments utilizing other ADTs, it is necessary to know the chain of constructors required to generate "foreign" values—in our example, a *True* value gets generated if a *Leaf$_A$* gets generated with a *Just* constructor "in between." To obtain such information, we

Figure 11. Constructor dependency graph.

create of a *constructor dependency graph* (CDG), that is, a directed graph where each node represents a constructor and each edge represents its dependency. Each edge is labeled with its corresponding generation probability. Figure 11 shows the CDG for *Tree* starting from the *Leaf$_A$* constructor. Having this graph together with the application of the multi-type branching process, we can predict the expected number of constructors belonging to external ADTs. It is enough to multiply the probabilities at each edge of the path between every constructor involved in the branching process and the desired external constructor.

The extensions described so far enable our tool (presented in the next section) to make predictions about *QuickCheck* generators for ADTs defined in many existing Haskell libraries.

7 Implementation

DRAGEN is a tool chain written in Haskell that implements the multi-type branching processes (Section 4 and 5) and its extensions (Section 6) together with a distribution optimizer, which calibrates the probabilities involved in generators to fit developers' demands. **DRAGEN** synthesizes generators by calling the Template Haskell function *dragenArbitrary :: Name → Size → CostFunction → Q [Dec]*, where developers indicate the target ADT for which they want to obtain a *QuickCheck* generator; the desired generation size, needed by our prediction mechanism in order to calculate the distribution at the last generation level; and a *cost function* encoding the desired generation distribution.

The design decision to use a probability optimizer rather than search for an analytical solution is driven by two important aspects of the problem we aim to solve. Firstly, the computational cost of exactly solving a non-linear system of equations (such as those arising from branching processes) can be prohibitively high when dealing with a large number of constructors, thus a large number of unknowns to be solved for. Secondly, the existence of such exact solutions is not guaranteed due to the implicit invariants the data types under consideration might have. In such cases, we believe it is much more useful to construct a distribution that approximates the user's goal, than to abort the entire compilation process. We give an example of this approximate solution finding behavior later in this section.

7.1 Cost functions

The optimization process is guided by a user-provided cost function. In our setting, a cost function assigns a real number (a cost) to the combination of a generation size (chosen by the user) and a mapping from constructors to probabilities:

type *CostFunction = Size → ProbMap → Double*

Type *ProbMap* encodes the mapping from constructor names to real numbers. Our optimization algorithm works by generating several *ProbMap* candidates that are evaluated through the provided cost function in order to choose the most suitable one. Cost functions are expected to return a smaller positive number as the predicted distribution obtained from its parameters gets closer to a certain *target distribution*, which depends on what property that particular cost function is intended to encode. Then, the optimizator simply finds the best *ProbMap* by minimizing the provided cost function.

Currently, our tool provides a basic set of cost functions to easily describe the expected distribution of the derived generator. For instance, *uniform :: CostFunction* encodes constuctor-wise uniform generation, an interesting property that naturally arises from our generation process formalization. It guides the optimization process to a generation distribution that minimizes the difference between the expected number of each generated constructor and the generation size.

Moreover, the user can restrict the generation distribution to a certain subset of constructors using the cost functions *only* :: [*Name*] → *CostFunction* and *without* :: [*Name*] → *CostFunction* to describe these restrictions. In this case, the whitelisted constructors are then generated following the *uniform* behavior. Similarly, if the branching process involves mutually recursive data types, the user could restrict the generation to a certain subset of data types by using the functions *onlyTypes* and *withoutTypes*. Additionally, when the user wants to generate constructors according to certain proportions, *weighted* :: [(*Name*, *Int*)] → *CostFunction* allows to encode this property, e.g. three times more *Leaf$_A$*'s than *Leaf$_B$*'s.

Table 1 shows the number of expected and observed constructors of different *Tree* generators obtained by using different cost functions. The observed expectations were calculated averaging the number of constructors across 100000 generated values. Firstly, note how the generated distributions are soundly predicted by our tool. In our tests, the small differences between predictions and actual values dissapear as we increase the number of generated values. As for the cost functions' behavior, there are some interesting aspects to note. For instance, in the *uniform* case the optimizer cannot do anything to break the implicit invariant of the data type: every binary tree with *n* nodes has *n* + 1 leaves. Instead, it converges to a solution that "approximates" a uniform distribution around the generation size parameter. We believe this is desirable behavior, to find an approximate solution when certain invariants prevent the optimization process from finding an exact solution. This way the user does not have to be aware of the possible invariants that the target data type may have, obtaining a solution that is good enough for most purposes. On the other hand, notice that in the *weighted* case at the second row of Table 1, the expected number of generated *Node*s is considerably large. This constructor is not listed in the proportions list, hence the optimizer can freely adjust its probability to satisfy the proportions specified for the leaves.

7.2 Derivation Process

DRAGEN's derivation process starts at compile-time with a type reification stage that extracts information about the structure of the types under consideration. It follows an intermediate stage composed of the optimizer for probabilities used in generators, which is guided by our multi-type branching process model, parametrized on the cost function provided. This optimizer is based on a standard local-search optimization algorithm that recursively chooses the best mapping from constructors to probabilities in the current neighborhood. Neighbors are *ProbMap*s, determined by individually varying the probabilities for *each constructor* with a predetermined Δ. Then, to determine the "best" probabilities, the local-search applies our prediction mechanishm to the immediate neighbors that have not yet been visited by evaluating the cost function to select the most suitable next candidate. This process continues until a local minimum is reached when there are no new

neighbors to evaluate, or if each step improvement is lower than a minimum predetermined ε.

The final stage synthesizes a *Arbitrary* type-class instance for the target data types using the optimized generation probabilities. For this stage, we extend some functionality present in *MegaDeTH* in order to derive generators parametrized by our previously optimized probabilities. Refer to the supplementary material for further details on the cost functions and algorithms addressed by this section.

8 Case Studies

We start by comparing the generators for the ADT *Tree* derived by *MegaDeTH* and *Feat*, presented in Section 2, with the corresponding generator derived by **DRAGEN** using a *uniform* cost function. We used a generation size of 10 both for *MegaDeTH* and **DRAGEN**, and a generation size of 400 for *Feat*—that is, *Feat* will generate test cases of maximum 400 constructors, since this is the maximum number of constructors generated by our tool using the generation size cited above. Figure 12 shows the differences between the complexity of the generated values in terms of the number of constructors. As shown in Figure 3, generators derived by *MegaDeTH* and *Feat* produce very narrow distributions, being unable to generate a diverse variety of values of different sizes. In contrast, the **DRAGEN** optimized generator provides a much wider distribution, i.e., from smaller to bigger values.

It is likely that the richer the values generated, the better the chances of covering more code, and thus of finding more bugs. The next case studies provide evidence in that direction.

Although **DRAGEN** can be used to test Haskell code, we follow the same philosophy as *QuickFuzz*, targeting three complex and widely used external programs to evaluate how well our derived generators behave. These applications are *GNU bash 4.4*—a widely used Unix shell, *GNU CLISP 2.49*—the GNU Common Lisp compiler, and *giffix*—a small test utility from the *GIFLIB 5.1* library focused on reading and writing Gif images. It is worth noticing that these applications are not written in Haskell. Nevertheless, there are Haskell libraries designed to inter-operate with them: *language-bash*, *atto-lisp*, and *JuicyPixels*, respectively. These libraries provide ADT definitions
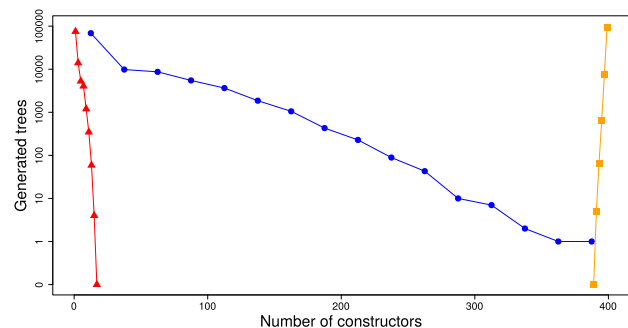

Figure 12. *MegaDeTH* (▲) vs. *Feat* (■) vs. **DRAGEN** (●) generated distributions for type *Tree*.

Table 1. Predicted and actual distributions for *Tree* generators using different cost functions.

Cost Function	Predicted Expectation				Observed Expectation			
	$Leaf_A$	$Leaf_B$	$Leaf_C$	$Node$	$Leaf_A$	$Leaf_B$	$Leaf_C$	$Node$
uniform	5.26	5.26	5.21	14.73	5.27	5.26	5.21	14.74
weighted $[('Leaf_A, 3), ('Leaf_B, 1), ('Leaf_C, 1)]$	30.07	9.76	10.15	48.96	30.06	9.75	10.16	48.98
weighted $[('Leaf_A, 1), ('Node, 3)]$	10.07	3.15	17.57	29.80	10.08	3.15	17.58	29.82
only $['Leaf_A, 'Node]$	10.41	0	0	9.41	10.43	0	0	9.43
without $['Leaf_C]$	6.95	6.95	0	12.91	6.93	6.92	0	12.86

which we used to synthesize **DRAGEN** generators for the inputs of the aforementioned applications. Moreover, they also come with serialization functions that allow us to transform the randomly generated Haskell values into the actual test files that we used to test each external program. The case studies contain mutually recursive and composite ADTs with a wide number of constructors (e.g., GNU bash spans 31 different ADTs and 136 different constuctors)—refer to the supplementary material for a rough estimation of the scale of such data types and the data types involved with them.

For our experiments, we use the coverage measure known as *execution path* employed by American Fuzzy Lop (AFL) [20]—a well known fuzzer. It was chosen in this work since it is also used in the work by Grieco et al. [14] to compare *MegaDeTH* with other techniques. The process consists of the *instrumentation* of the binaries under test, making them able to return the path in the code taken by each execution. Then, we use AFL to count how many different executions are triggered by a set of randomly generated files—also known as a corpus. In this evaluation, we compare how different *QuickCheck* generators, derived using *MegaDeTH* and using our approach, result in different code coverage when testing external programs, as a function of the size of a set of independently, randomly generated corpora. We have not been able to automatically derive such generators using *Feat*, since it does not work with some Haskell extensions used in the bridging libraries.

We generated each corpus using the same ADTs and generation sizes for each derivation mechanism. We used a generation size of 10 for CLISP and bash files, and a size of 5 for Gif files. For **DRAGEN**, we used *uniform* cost functions to reduce any external bias. In this manner, any observed difference in the code coverage triggered by the corpora generated using each derivation mechanism is entirely caused by the optimization stage that our predictive approach performs, which does not represent an extra effort for the programmer. Moreover, we repeat each experiment 30 times using independently generated corpora for each combination of derivation mechanism and corpus size.

Figure 13 compares the mean number of different execution paths triggered by each pair of generators and corpus sizes, with error bars indicating 95% confidence intervals of the mean. It is easy to see how the **DRAGEN** generators synthesize test cases capable of triggering a much larger number of different

execution paths in comparison to *MegaDeTH* ones. Our results indicate average increases approximately between 35% and 41% with an standard error close to 0.35% in the number of different execution paths triggered in the programs under test.

An attentive reader might remember that *MegaDeTH* tends to derive generators which produce very small test cases. If we consider that small test cases should take less time (on average) to be tested, is fair to think there is a trade-off between being able to test a bigger number of smaller test cases or a smaller number of bigger ones having the same time available. However, when testing external software like in our experiments, it is important to consider the time overhead introduced by the operating system. In this scenario, it is much more preferable to test interesting values over smaller ones. In our tests, size differences between the generated values of each tool does do not result in significant differences in the runtimes required to test each corpora—refer to the supplementary material for further details. A user is most likely to get better results by using our tool instead of *MegaDeTH*, with virtually *the same effort.*

We also remark that, if we run sufficiently many tests, then the expected code coverage will tend towards 100% of the reachable code in both cases. However, in practice, our approach is more likely to achieve higher code coverage for the same number of test cases.

9 Related Work

Fuzzers are tools to tests programs against randomly generated unexpected inputs. *QuickFuzz* [13, 14] is a tool that synthesizes data with rich structure, that is, well-typed files which can be used as initial "seeds" for state-of-the-art fuzzers—a work flow which discovered many unknown vulnerabilities. Our work could help to improve the variation of the generated initial seeds, by varying the distribution of *QuickFuzz* generators—an interesting direction for future work.

SmallCheck [27] provides a framework to exhaustively test data sets up to a certain (small) size. The authors also propose a variation called *Lazy SmallCheck*, which avoids the generation of multiple variants which are passed to the test, but not actually used.

QuickCheck has been used to generate well-typed lambda terms in order to test compilers [25]. Recently, Midtgaard et al. extend such a technique to test compilers for impure programming languages [22].
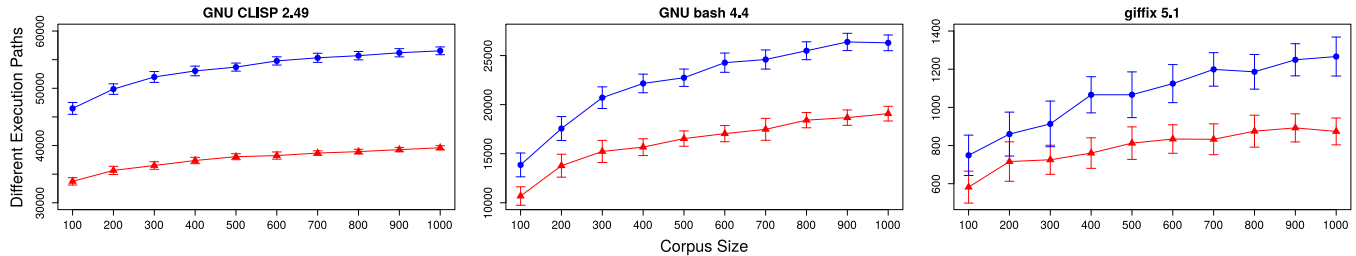
Figure 13. Path coverage comparison between *MegaDeTH* (▲) and **DRAGEN** (●).

Luck [18] is a domain specific language for writing testing properties and *QuickCheck* generators at the same time. We see *Luck*'s approach as orthogonal to ours, which is mostly intended to be used when we do not know any specific property of the system under test, although we consider that borrowing some functionality from *Luck* into **DRAGEN** is an interesting path for future work.

Recently, Lampropoulos et al. propose a framework to automatically derive random generators for a large subclass of Coqs' inductively defined relations [19]. This derivation process also provides proof terms certifying that each derived generator is sound and complete with respect to the inductive relation it was derived from.

Boltzmann models [10] are a general approach to randomly generating combinatorial structures such as trees and graphs—also extended to work with closed simply-typed lambda terms [4]. By implementing a *Boltzmann sampler*, it is possible to obtain a random generator built around such models which uniformly generates values of a target size with a certain size tolerance. However, this approach has practical limitations. Firstly, the framework is not expressive enough to represent complex constrained data structures, e.g red-black trees. Secondly, Boltzmann samplers give the user no control over the distribution of generated values besides ensuring size-uniform generation. They work well in theory but further work is required to apply them to complex structures [26]. Conversely, **DRAGEN** provides a simple mechanism to predict and tune the overall distribution of constructors *analytically at compile-time*, using statically known type information, and requiring no runtime reinforcements to ensure the predicted distributions. Future work will explore the connections between branching processes and Boltzmann models.

Similarly to our work, Feldt and Poulding propose *GödelTest* [12], a search-based framework for generating biased data. It relies on non-determinism to generate a wide range of data structures, along with metaheuristic search to optimize the parameters governing the desired biases in the generated data. Rather than using metaheuristic search, our approach employs a completely analytical process to predict the generation distribution at each optimization step. A strength of the *GödelTest* approach is that it can optimize the probability parameters even when there is no specific target distribution over the

constructors—this allows exploiting software behavior under test to guide the parameter optimization.

The efficiency of random testing is improved if the generated inputs are evenly spread across the input domain [5]. This is the main idea of *Adaptive Random Testing* (ART) [6]. However, this work only covers the particular case of testing programs with numerical inputs and it has also been argued that adaptive random testing has inherent inefficiencies compared to random testing [1]. This strategy is later extended in [7] for object-oriented programs. These approaches present no analysis of the distribution obtained by the heuristics used, therefore we see them as orthogonal work to ours.

10   Final Remarks

We discover an interplay between the stochastic theory of branching processes and algebraic data types structures. This connection enables us to describe a solid mathematical foundation to capture the behavior of our derived *QuickCheck* generators. Based on our formulas, we implement a heuristic to automatically adjust the expected number of constructors being generated as a way to control generation distributions.

One holy grail in testing is the generation of structured data which fulfills certain invariants. We believe that our work could be used to enforce some invariants on data "up to some degree." For instance, by inspecting programs' source code, we could extract the pattern-matching patterns from programs (e.g., (*Cons* (*Cons* *x*))) and derive generators which ensure that such patterns get exercised a certain amount of times (on average)—intriguing thoughts to drive our future work.

Acknowledgments

References

[1] A. Arcuri and L. Briand. 2011. Adaptive Random Testing: An Illusion of Effectiveness?. In *Proc. of the International Symposium on Software Testing and Analysis (ISSTA '11)*. ACM.

[2] B. Arkin, S. Stender, and G. McGraw. 2005. Software penetration testing. *IEEE Security Privacy* (2005).

[3] T. Arts, J. Hughes, U. Norell, and H. Svensson. 2015. Testing AUTOSAR software with QuickCheck. In *In Proc. of IEEE International Conference on Software Testing, Verification and Validation, ICST Workshops.*

[4] M. Bendkowski, K. Grygiel, and P. Tarau. 2017. Boltzmann Samplers for Closed Simply-Typed Lambda Terms. In *In Proc. of International Symposium on Practical Aspects of Declarative Languages.* ACM.

[5] F.T. Chan, T.Y. Chen, I.K. Mak, and Y.T. Yu. 1996. Proportional sampling strategy: guidelines for software testing practitioners. *Information and Software Technology* 38, 12 (1996), 775 – 782.

[6] T. Y. Chen, H. Leung, and I. K. Mak. 2005. Adaptive Random Testing. In *Advances in Computer Science - ASIAN 2004. Higher-Level Decision Making*, Michael J. Maher (Ed.). Springer Berlin Heidelberg.

[7] I. Ciupa, A. Leitner, M. Oriol, and B. Meyer. 2008. ARTOO: adaptive random testing for object-oriented software. In *Proc. of International Conference on Software Engineering.* ACM/IEEE.

[8] K. Claessen, J. Duregård, and M. H. Palka. 2014. Generating Constrained Random Data with Uniform Distribution. In *Proc. of the Functional and Logic Programming FLOPS.*

[9] K. Claessen and J. Hughes. 2000. QuickCheck: A Lightweight Tool for Random Testing of Haskell Programs. In *Proc. of the ACM SIGPLAN International Conference on Functional Programming (ICFP).*

[10] P. Duchon, P. Flajolet, G. Louchard, and G. Schaeffer. 2004. Boltzmann Samplers for the Random Generation of Combinatorial Structures. *Combinatorics, Probability and Computing.* 13 (2004).

[11] J. Duregård, P. Jansson, and M. Wang. 2012. Feat: Functional enumeration of algebraic types. In *Proc. of the ACM SIGPLAN Symposium on Haskell.*

[12] R. Feldt and S. Poulding. 2013. Finding test data with specific properties via metaheuristic search. In *Proc. of International Symp. on Software Reliability Engineering (ISSRE).* IEEE.

[13] G. Grieco, M. Ceresa, and P. Buiras. 2016. QuickFuzz: An automatic random fuzzer for common file formats. In *Proc. of the International Symposium on Haskell.* ACM.

[14] G. Grieco, M. Ceresa, A. Mista, and P. Buiras. 2017. QuickFuzz testing for fun and profit. *Journal of Systems and Software* 134, Supp. C (2017).

[15] P. Haccou, P. Jagers, and V. Vatutin. 2005. *Branching processes. Variation, growth, and extinction of populations.* Cambridge University Press.

[16] J. Hughes, C. Pierce B, T. Arts, and U. Norell. 2016. Mysteries of DropBox: Property-Based Testing of a Distributed Synchronization Service. In *Proc.*

[17] J. Hughes, U. Norell, N. Smallbone, and T. Arts. 2016. Find more bugs with QuickCheck!. In *Proc. of the International Workshop on Automation of Software Test, AST@ICSE.*

[18] L. Lampropoulos, D. Gallois-Wong, C. Hritcu, J. Hughes, B. C. Pierce, and L. Xia. 2017. Beginner's luck: a language for property-based generators. In *Proc. of the ACM SIGPLAN Symposium on Principles of Programming Languages, POPL.*

[19] L. Lampropoulos, Z. Paraskevopoulou, and B. C. Pierce. 2017. Generating Good Generators for Inductive Relations. *In Proc. ACM on Programming Languages* 2, POPL, Article 45 (2017).

[20] M. Zalewski. 2010. American Fuzzy Lop: a security-oriented fuzzer. http://lcamtuf.coredump.cx/afl/. (2010).

[21] C. McBride and R. Paterson. 2008. Applicative Programming with Effects. *Journal of Functional Programming* 18, 1 (Jan. 2008).

[22] J. Midtgaard, M. N. Justesen, P. Kasting, F. Nielson, and H. R. Nielson. 2017. Effect-driven QuickChecking of compilers. *In Proceedings of the ACM on Programming Languages, Volume 1* ICFP (2017).

[23] A. Mista, A. Russo, and J. Hughes. 2018. Branching Processes for QuickCheck Generators (extended version). https://bitbucket.org/agustinmista/dragen/downloads/full-paper.pdf. (2018).

[24] N. Mitchell. 2007. Deriving Generic Functions by Example. In *Proc. of the 1st York Doctoral Syposium.* Tech. Report YCS-2007-421, Department of Computer Science, University of York, UK, 55–62.

[25] M. Palka, K. Claessen, A. Russo, and J. Hughes. 2011. Testing and Optimising Compiler by Generating Random Lambda Terms. In *The IEEE/ACM International Workshop on Automation of Software Test (AST 2011).*

[26] S. M. Poulding and R. Feldt. 2017. Automated Random Testing in Multiple Dispatch Languages. *IEEE International Conference on Software Testing, Verification and Validation (ICST)* (2017).

[27] C. Runciman, M. Naylor, and F. Lindblad. 2008. Smallcheck and Lazy Smallcheck: automatic exhaustive testing for small values. In *Proc. of the ACM SIGPLAN Symposium on Haskell.*

[28] T. Sheard and Simon L. Peyton Jones. 2002. Template meta-programming for Haskell. *SIGPLAN Notices* 37, 12 (2002), 60–75.

[29] H. W. Watson and F. Galton. 1875. On the probability of the extinction of families. *The Journal of the Anthropological Institute of Great Britain and Ireland* (1875).