

Parameterised Approximation of the Fixation Probability of the Dominant Mutation in the Multi-Type Moran Process*

Leslie Ann Goldberg
Department of Computer Science
University of Oxford
United Kingdom

Marc Roth
School of Electronic Engineering and Computer Science
Queen Mary University of London
United Kingdom

Tassilo Schwarz
Mathematical Institute
University of Oxford
United Kingdom

02 October, 2023

Abstract

The multi-type Moran process is an evolutionary process on a connected graph G in which each vertex has one of k types and, in each step, a vertex v is chosen to reproduce its type to one of its neighbours. The probability of a vertex v being chosen for reproduction is proportional to the fitness of the type of v . So far, the literature was almost solely concerned with the 2-type Moran process in which each vertex is either healthy (type 0) or a mutant (type 1), and the main problem of interest has been the (approximate) computation of the so-called *fixation probability*, i.e., the probability that eventually all vertices are mutants.

In this work we initiate the study of approximating fixation probabilities in the multi-type Moran process on general graphs. Our main result is an FPTRAS (fixed-parameter tractable randomised approximation scheme) for computing the fixation probability of the dominant mutation; the parameter is the number of types and their fitnesses. In the course of our studies we also provide novel upper bounds on the expected *absorption time*, i.e., the time that it takes the multi-type Moran process to reach a state in which each vertex has the same type.

1 Introduction

The study of the (2-type) Moran process dates back to the late 1950s [16] with the goal of modelling and understanding how an advantageous mutation spreads through a finite community: In its initial form, the process starts with n individuals, $n - 1$ of which are healthy (fitness 1), and one of which is a “mutant” with an advantage (fitness $r > 1$). In each step of the process, an individual is chosen with probability proportional to its fitness and replaces the type of another individual, which is chosen uniformly at random, with its own type (healthy or mutant). This process becomes stable if either all individuals are mutants, or if the mutation is extinct. It is well-known that the probability that the mutation takes over, which is called the *fixation probability*, is $(r - 1)/(r^n - 1)$ (see e.g. Lieberman, Hauert, and Nowak [14]).

The Moran process was later generalised to communities with a spatial structure [14, 17]. In the generalisation, individuals are vertices of an n -vertex graph G and each vertex can only replace the type of its neighbours. Moreover, the single mutant at the start of the process is chosen uniformly at random from all vertices. For this *spatial* Moran process, the fixation probability crucially depends

*This research was funded in part by the German Academic Scholarship Foundation. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission. All data is provided in full in the results section of this paper. The second author has been affiliated with the Department of Computer Science, University of Oxford, at the time of writing this paper.

on the structure of the graph. This quantity can be calculated explicitly, but the number of equations for doing so would grow exponentially in the graph's size. Hence, Díaz, Goldberg, Mertzi, Richerby, Serna and Spirakis [5] constructed a *fully polynomial randomised approximation scheme (FPRAS)* for the associated computational problem. This means that they gave an algorithm that takes as input a graph G , a mutation fitness $r \geq 1$, and real numbers ε and δ in the interval $(0,1)$ that guide the accuracy of the approximation. With probability at least $1 - \delta$, the algorithm outputs an ε -approximation of the fixation probability, which means a number which is within $1 \pm \varepsilon$ of the fixation probability. According to the definition of FPRAS, the running time of this algorithm is bounded from above by a polynomial in $|G|$, $1/\varepsilon$ and $\log(1/\delta)$. For more related work on the 2-type Moran process including improved FPRASes, see Section 1.4.

Much less is known about the Moran process when there is more than one type of mutation. Although the multi-type Moran process has been considered for the special case of complete graphs by Ferreira and Neves [8], there is, to the best of our knowledge, no work on the multi-type Moran process on general graphs.

In the setting of multiple types, mutated individuals will compete not only against healthy individuals, but different mutation lineages will also fight for dominance against each other. We initiate in this work the study of the multi-type Moran process on general graphs. The goal is to compute the fixation probability of the dominant mutant, i.e., the type with maximum fitness: If this probability is small enough, then we can hope for the healthy individuals/cells to eventually take over.

Our main result is an efficient algorithm for approximating the fixation probability of the dominant mutant. For the formal statement of our results, we will next introduce the process in detail.

1.1 The Multi-Type Moran Process

The graphs that we consider in this work are simple, undirected, and without loops. Given a vertex v of a graph $G = (V, E)$, we write $N(v) := \{u \mid \{u, v\} \in E\}$ for the (open) neighbourhood of v , and we write $d(v) := |N(v)|$ for the degree of v .

Let $G = (V, E)$ be an n -vertex graph. Let τ be a finite set of types and let τ_{ord} be an element of τ . Intuitively, τ_{ord} is the “ordinary” type — all other types in τ are mutations. Let f be a function from τ to $\mathbb{Q}_{\geq 1}$. We refer to f as a “fitness function” because we use it to measure the “fitness” of the types in τ . Let $f^+ = \max_{i \in \tau} f(i)$ and $\tau_{f^+} = \{i \in \tau \mid f(i) = f^+\}$. Similarly, let $f^- = \min_{i \in \tau} f(i)$ and $\tau_{f^-} = \{i \in \tau \mid f(i) = f^-\}$. In this work, we will study the case where mutations are advantageous, so $\tau_{\text{ord}} \in \tau_{f^-}$. We say that f is an *advantageous fitness function* for τ and τ_{ord} in this case.

Let Ω be the set of functions from V to τ . The elements of Ω will be the states that we study — each state assigns a type to every vertex of G . Given a state $M_0 \in \Omega$, the *Moran process* corresponding to G , τ , f and M_0 is a Markov Chain $M = M_0, M_1, \dots$ with state space Ω and start state M_0 .

Before defining the transitions of M , we give some notation. If S is a state in Ω and v and w are vertices in V then $S|_{v \rightarrow w}$ is the state S' such that $S'(w) = S(v)$ and, for all $u \in V \setminus \{w\}$, $S'(u) = S(u)$. Informally, S' is derived from S by reproducing the type of v onto the type of w . For example, if v is a mutation type, then this mutation spreads to its neighbour w in $S|_{v \rightarrow w}$.

We refer to M_t as the state of the Markov Chain M at (discrete) time t . For any type $j \in \tau$, we write $V_j(t)$ for the set of vertices that are assigned type j at time t , that is, $V_j(t) := \{v \in V \mid M_t(v) = j\}$. To simplify the notation, we set $f(v, t) := f(M_t(v))$ for the fitness of the type of v at time t . Given a subset $U \subseteq V$, we write $f(U, t) := \sum_{v \in U} f(v, t)$ to denote the total fitness of the vertices in U at time t . The *total fitness* $F(t)$ at time t is the total fitness of all vertices, that is, $F(t) := f(V, t)$.

The transition from M_t to M_{t+1} is defined as follows. A random vertex $v \in V$ is chosen with probability $f(v, t)/F(t)$. A vertex $w \in N(v)$ is chosen with probability $1/d(v)$. Then $M_{t+1} := M_t|_{v \rightarrow w}$. Intuitively, in each transition, a vertex v is chosen with probability proportional to the fitness of its type at time t and a vertex w is chosen uniformly at random from the neighbours of v . Then the type of w is replaced with the type of v .

If M_t assigns the same type to all vertices then the state of the Markov chain M does not change after time t . We say that type $j \in \tau$ *fixates* by time t if M_t assigns type j to all vertices. If G is connected, then with probability 1, some type in τ fixates. We use π_j to denote the probability that j fixates.

If we want to clarify G , τ , f , and M_0 , then we refer to M as $M(G, \tau, f, M_0)$ and we refer to π_j as $\pi_j(G, \tau, f, M_0)$.

Sometimes it will be useful to consider a distribution on the start state M_0 . If D is a distribution on Ω then $\pi_j(G, \tau, f, D) = \sum_{M_0 \in \Omega} \Pr_D(M_0) \pi_j(G, \tau, f, M_0)$.

For our purposes, not every distribution D on Ω will be important. In the well-studied case where $|\tau| = 2$ and the fitness function is advantageous (so $\tau_{\text{ord}} \in \tau_{f-}$) the idea is that the type $\alpha \in \tau \setminus \{\tau_{\text{ord}}\}$ is a “mutant” type with a high fitness that arises as a result of a single mutation. It is therefore appropriate to consider the case where the start state M_0 is drawn from the distribution $D_{\text{mut}} = D_{\text{mut}}(G, \tau, \tau_{\text{ord}})$ which is uniform on states $M_0 \in \Omega$ in which exactly one vertex has the type in $\tau \setminus \{\tau_{\text{ord}}\}$. This is the distribution that has been studied in previous works.

There is a natural generalisation of D_{mut} to the multi-type Moran process where $|\tau| = k > 2$. In this case, D_{mut} is the uniform distribution on states $M_0 \in \Omega$ which have exactly one vertex assigned to every type $i \in \tau \setminus \{\tau_{\text{ord}}\}$.

Our results will apply to D_{mut} and also to several other distributions. Let $V[k]$ be the set containing all k -tuples of distinct vertices from V and let $\tau[k]$ be the set containing all k -tuples of distinct types in τ . For every $\mathbf{u} \in V[k]$ and $\gamma \in \tau[k]$, let $\Omega(\mathbf{u}, \gamma)$ be the set of states in Ω that map vertices in \mathbf{u} respectively to types in γ . When a state $S \in \Omega$ is drawn from D_{mut} , there is a clear lower bound on the probability that $S \in \Omega(\mathbf{u}, \gamma)$. Our results will apply to any distribution with this property. For concreteness, we identify a set of suitable distributions as follows. The set of distributions $\mathcal{D} = \mathcal{D}(G, \tau)$ is the set containing every distribution D on Ω that meets the following criteria.

- (1) For every pair $(\mathbf{u}, \gamma) \in V[k] \times \tau[k]$, there is some distribution $D_{\mathbf{u}, \gamma}$ on $\Omega(\mathbf{u}, \gamma)$, such that (2) holds.
- (2) For any $S \in \Omega$, $\Pr_D(S) = \frac{1}{|V[k] \times \tau[k]|} \sum_{(\mathbf{u}, \gamma) \in V[k] \times \tau[k]} \Pr_{D_{\mathbf{u}, \gamma}}(S)$.

It is easy to see that $D_{\text{mut}}(G, \tau, \tau_{\text{ord}}) \in \mathcal{D}(G, \tau)$ since, in this case, $D_{\mathbf{u}, \gamma}$ can be taken to be the distribution containing the single state in $\Omega(\mathbf{u}, \gamma)$ such that every vertex outside of \mathbf{u} is assigned type τ_{ord} .

Recall that Ω is the set of functions from V to τ . Let $\Omega_0(G, \tau)$ be the set of all states $S \in \Omega$ with range τ . That is, for every element $i \in \tau$, there is some $v \in V$ such that $S(v) = i$. By construction, any sample S drawn from a distribution $D \in \mathcal{D}(G, \tau)$ is in $\Omega_0(G, \tau)$.

1.2 Parameterised and Approximation Algorithms

We will study the following problems concerned with the fixation probability of a type with maximum fitness.

Definition 1 (MAXFIXPROB_{mut}). *The problem MAXFIXPROB_{mut} takes as input a connected graph G , a set of types τ with ordinary type $\tau_{\text{ord}} \in \tau$, an advantageous fitness function $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$, and a type $\alpha \in \tau_{f+}$. The goal is to compute $\pi_\alpha(G, \tau, f, D_{\text{mut}}(G, \tau, \tau_{\text{ord}}))$.*

Definition 2 (MAXFIXPROB). *The problem MAXFIXPROB takes as input a connected graph G , a set of types τ , a fitness function $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$, a type $\alpha \in \tau_{f+}$ and a distribution $D \in \mathcal{D}(G, \tau)$ (via a black-box oracle which provides samples from D). The goal is to compute $\pi_\alpha(G, \tau, f, D)$.*

As explained in the introduction, we are interested in efficient approximation algorithms for MAXFIXPROB and MAXFIXPROB_{mut}.

As is usual in the setting of parameterised algorithms (and is natural in the case of this application), we will assume that the number of types (and the encoding of their fitnesses) is significantly smaller than the size of the input graph G , which we denote by $|G|$. Thus, we are happy to accept a factor in the running time that depends on τ and f , say $g(\tau, f)$ — but we would not be happy with an algorithm whose running time is as high as $|G|^{g(\tau, f)}$ for any function g .

Formally, we let $\|f\|$ be the description length of the fitness function f . The exact details about how f is encoded will not be important here since $|G|$ is assumed to be much bigger than $\|f\|$. We *parameterise* our problems in terms of the parameter $\kappa := \|f\|$ and we aim to construct a *fixed-parameter tractable approximation scheme (FPTRAS)*, which is the standard notion for efficient parameterised approximation algorithms (see Arvind and Raman [1]).

Definition 3 (FPTRAS). A fixed-parameter tractable randomised approximation scheme (FPTRAS) for a problem P is a randomised algorithm \mathbb{A} that takes as input a triple (x, ε, δ) where x is an input to the problem P and ε and δ are real numbers in $(0, 1)$. The output of $\mathbb{A}(x, \varepsilon, \delta)$ is a random variable \hat{P} such that $\Pr((1 - \varepsilon)P(x) \leq \hat{P} \leq (1 + \varepsilon)P(x)) \geq 1 - \delta$. In order for \mathbb{A} to be an FPTRAS, there must be a computable function g such that the running time of \mathbb{A} is bounded from above by $g(\kappa) \cdot \text{poly}(|x|, 1/\varepsilon, \log(1/\delta))$, where κ is the parameter of x .

1.3 Main Results

We are now able to state our main results.

Theorem 4. MAXFIXPROB has an FPTRAS when parameterised by $\kappa := \|f\|$.

Recall that MAXFIXPROB allows for oracle access to the initial distribution. The next theorem asserts that the oracle access is not necessary for an FPTRAS if the initial distribution is the standard distribution.

Theorem 5. MAXFIXPROB_{mut} has an FPTRAS when parameterised by $\kappa := \|f\|$.

In the course of proving Theorems 4 and 5 we also obtain upper bounds on the expected absorption time, i.e. the time that the process runs until a type fixates. For the multi-type Moran process, this upper bound is of the form $g(\|f\|) \cdot \text{poly}(|G|)$ (see Theorem 18).

1.4 Related Work

Approximation algorithms for the 2-type spatial Moran process

Work on approximating the fixation probability in the spatial Moran process has so far focused on the 2-type process, for which Díaz et al. [5] developed the first FPRAS. Denoting the types by $\tau = \{1, 2\}$ with $\tau_{\text{ord}} = 1$ being the ordinary type, the idea to approximate the fixation probability π_2 of the mutation is to perform sufficiently many Markov chain Monte Carlo simulations of the process, where the number of steps and simulations can be bounded by proving (a) an upper bound on the expected absorption time and (b) a lower bound on the fixation probability π_2 . Their work applies to the case where the mutation is advantageous, so $\tau_{f-} = \{1\}$ and $f(2) > f(1)$. Note that when the mutation is disadvantageous, its fixation probability π_2 can be exponentially small as a function of the size of the graph, even when the graph is the complete graph, so this approach to approximation is infeasible.

In a first improvement, Chatterjee, Ibsen-Jensen, and Nowak [4] exploit the observation that a vertex reproducing to a neighbour of the same type does not change the overall state. By simulating only those steps that change the underlying Markov chain's state, they speed the FPRAS up.

The effect of $f(2)/f(1)$ on the fixation probability π_2 in *every* connected, undirected graph was characterised by Goldberg, Lapinskas, and Richerby [10] by demonstrating a phase transition in π_2 as a function of $f(2)/f(1)$. This phase transition also entails an improved lower bound on π_2 in the advantageous case. Combining this bound with early termination of the Monte Carlo simulation once fixation is sufficiently likely, they provide a faster FPRAS.

Multiple types in related models

The idea of modelling multiple mutation types has been explored for both the Wright-Fisher model [2] and, more recently, the non-spatial Moran process [7, 8, 3]. The Wright-Fisher model [9, 18] is a non-spatial model with a coarse-grained time scale where all individuals reproduce simultaneously, resulting in non-overlapping generations. This model served as the predecessor of the more fine-grained Moran process; see Lanchier [13] for a detailed definition and comparison to the non-spatial Moran process.

For the Wright-Fisher model, Beerenwinkel, Antal, Dingli, Traulsen, Kinzler, Velculescu, Vogelstein and Nowak [2] provide a generalisation to account for the evolution of multiple types: They consider a sequentially ordered set τ of infinitely many types, where higher types have higher fitness. Initially, all individuals are of some ordinary type τ_{ord} . Besides reproduction happening generation by generation, individuals can progress to a higher type at some rate which depends on how close the initial type is

to the new one in the ordering of τ . Beerenwinkel et al. perform computer simulations of this process which models the evolution of cancer cells. Their goal is to approximate the expected time until a cancer phenotype appear, which they define as the waiting time until some fixed type occurs as a mutation. The focus of their model and simulations is thus on the progression of types.

Etheridge and Griffiths [7] define for the non-spatial Moran process a finite set of types τ where, in addition to the reproduction dynamics, mutation happens between different types at given rates. Chao and Schweinsberg [3] extend the study to a population on a continuous torus, where the focus is on the progression between types (similar to Beerenwinkel et al. [2]).

For the Moran process without spatial dynamics and without progression between types, Ferreira and Neves [8] introduce in recent work the concept of multiple types. Their model is the same as ours except that they focus on the special case where the graph G is the complete graph. They provide bounds on the fixation probabilities in the 3-type case. In the context of our work on the *spatial* Moran process, these bounds can be derived by applying our coupling in the special case where G is a complete graph and $|\tau| = 3$, as we show in Lemma 22.

2 The Continuous Moran Process

Although our main results are about the discrete-time Moran process, it is useful in proofs to consider a continuous-time version of the Moran process. Concretely, as shown by Díaz et al. [6], in the discrete-time case it is not always possible to construct a coupling between two Moran processes with a natural domination property we require for bounding the fixation probabilities.

We use $\text{Exp}(\lambda)$ to denote an exponential distribution with parameter λ . Let M be the (discrete-time) Moran process corresponding to $G = (V, E)$, τ , f , and M_0 . Following the work of Diaz, Goldberg, Richerby, and Serna [6] in the case where $|\tau| = 2$, we define a continuous-time Moran process $\tilde{M} = \tilde{M}(G, \tau, f, M_0)$ corresponding to M . The process \tilde{M} has the same state space as M , namely the set of function Ω from V to τ . and it starts with $\tilde{M}_0 = M_0$. The evolution of \tilde{M} is guided by a sequence T_0, T_1, \dots of non-negative real numbers with $T_0 = 0$. For any non-negative integer i , the evolution from \tilde{M}_{T_i} proceeds as follows. Each vertex $v \in V$ is equipped with a “clock” c_v^i which is an exponentially distributed random variable with parameter $\tilde{f}(v, T_i) := f(\tilde{M}_{T_i}(v))$. The distributions of the clocks c_v^i are mutually independent. Let $\Delta_i := \min_{w \in V} c_w^i$ and $T_{i+1} = T_i + \Delta_i$. For $t \in [T_i, T_{i+1})$, we define $\tilde{M}_t = \tilde{M}_{T_i}$. We define $\tilde{M}_{T_{i+1}}$ as follows. Let $v := \arg \min \{c_u^i \mid u \in V\}$. Choose $w \in N(v)$ with probability $1/d(v)$. Then $\tilde{M}_{T_{i+1}} := \tilde{M}_{T_i}|_{v \rightarrow w}$.

The reason that Díaz et al. [6] chose the exponential distribution for the clocks is that it is memoryless, so it is convenient for couplings. Moreover, the minimum of a set S of exponential distributions is exponentially distributed with a new parameter that is the sum of the parameters in S . Therefore, for all v and T_i , $\Pr(\Delta_i = c_v^i) = \tilde{f}(v, T_i) / \tilde{F}(T_i)$ where $\tilde{F}(T_i) = \sum_{v \in V} \tilde{f}(v, T_i)$. This is the same as the probability that v is chosen for reproduction in a discrete step starting from \tilde{M}_{T_i} . Thus, the process $\tilde{M}_{T_0}, \tilde{M}_{T_1}, \dots$ is a faithful copy of M . It follows that $\tilde{\pi}_j(G, \tau, f, \tilde{M}_0)$, the probability that j fixates in \tilde{M} , is equal to $\pi_j(G, \tau, f, M_0)$.

3 Coupling Continuous Multi-Type Moran Processes

The following lemma is a generalisation of [6, Lemma 5], which was for the case $k = 2$. On an intuitive level, the lemma establishes the existence of a coupling between two Moran processes that satisfies a domination property which will later be used for bounding the fixation probabilities and the expected absorption times.

Lemma 6. *Let $G = (V, E)$ be a graph, let τ be a set of types, and let $f, f': \tau \rightarrow \mathbb{Q}_{\geq 1}$ be fitness functions. Suppose that there is a type $\alpha \in \tau$ with $f(\alpha) = f'(\alpha)$ and that, for any types $\beta, \beta' \in \tau \setminus \{\alpha\}$, $f(\beta) \leq f'(\beta')$. Let \tilde{M}_0 and \tilde{M}'_0 be states in the set Ω of functions from V to τ . Let $\tilde{M} = \tilde{M}(G, \tau, f, \tilde{M}_0)$ and $\tilde{M}' = \tilde{M}'(G, \tau, f', \tilde{M}'_0)$ be continuous-time Moran processes. For all $j \in \tau$, let $\tilde{V}_j(t) = \{v \in V \mid \tilde{M}_t(v) = j\}$ and $\tilde{V}'_j(t) = \{v \in V \mid \tilde{M}'_t(v) = j\}$. Suppose that $V'_\alpha(0) \subseteq V_\alpha(0)$. Then there is a coupling between \tilde{M} and \tilde{M}' such that, for any time $t \in \mathbb{R}_{\geq 0}$, $V'_\alpha(t) \subseteq V_\alpha(t)$.*

Proof. For convenience, let $\tilde{f}(v, T_i) := f(\tilde{M}_{T_i}(v))$ and $\tilde{f}'(v, T_i) := f(\tilde{M}'_{T_i}(v))$. The joint process is guided by a sequence T_0, T_1, \dots of non-negative real numbers with $T_0 = 0$. At time T_0 , both processes are in their respective initial states \tilde{M}_0 and \tilde{M}'_0 . For any non-negative integer i , the evolution of the joint process from time T_i is defined as follows. First, we define some exponentially-distributed random variables.

$$\begin{aligned} \forall v \in V : \quad C_v^0(i) &\sim \text{Exp}(\min(\tilde{f}(v, T_i), \tilde{f}'(v, T_i))). \\ \forall v \in V \text{ with } \tilde{f}(v, T_i) > \tilde{f}'(v, T_i) : \quad C_v^1(i) &\sim \text{Exp}(\tilde{f}(v, T_i) - \tilde{f}'(v, T_i)). \\ \forall v \in V \text{ with } \tilde{f}(v, T_i) < \tilde{f}'(v, T_i) : \quad C_v^2(i) &\sim \text{Exp}(\tilde{f}'(v, T_i) - \tilde{f}(v, T_i)). \end{aligned}$$

The earliest of these random variables is the random variable $\hat{\Delta}(i)$ defined by

$$\hat{\Delta}(i) := \min \{C_v^0(i) \mid v \in V\} \cup \{C_v^1(i) \mid \tilde{f}(v, T_i) > \tilde{f}'(v, T_i)\} \cup \{C_v^2(i) \mid \tilde{f}(v, T_i) < \tilde{f}'(v, T_i)\}.$$

Let $T_{i+1} = T_i + \hat{\Delta}(i)$. The two chains evolve from T_i as follows. Both \tilde{M} and \tilde{M}' stay in their respective states over the interval $[T_i, T_{i+1})$. At time T_{i+1} the state changes as follows.

Case 0. If $\hat{\Delta}(i) = C_v^0(i)$ for some $v \in V$, select $w \in N(v)$ uniformly at random and set $\tilde{M}_{T_{i+1}} := \tilde{M}_{T_i}|_{v \rightarrow w}$ and $\tilde{M}'_{T_{i+1}} := \tilde{M}'_{T_i}|_{v \rightarrow w}$.

Case 1. If $\hat{\Delta}(i) = C_v^1(i)$ for some $v \in V$, select $w \in N(v)$ u.a.r and set $\tilde{M}_{T_{i+1}} := \tilde{M}_{T_i}|_{v \rightarrow w}$ and $\tilde{M}'_{T_{i+1}} := \tilde{M}'_{T_i}$.

Case 2. If $\hat{\Delta}(i) = C_v^2(i)$ for some $v \in V$, select $w \in N(v)$ u.a.r and set $\tilde{M}_{T_{i+1}} := \tilde{M}_{T_i}$ and $\tilde{M}'_{T_{i+1}} := \tilde{M}'_{T_i}|_{v \rightarrow w}$.

It is easy to see that this joint evolution is a correct evolution for the individual chains \tilde{M} and \tilde{M}' . For example, consider \tilde{M} (the argument for \tilde{M}' is symmetric). We will show that for every $v \in V$, the time until v is chosen for reproduction is exponentially distributed with parameter $\tilde{f}(v, T_i)$. If $\tilde{f}(v, T_i) \leq \tilde{f}'(v, T_i)$ this follows since $C_v^0(i)$ is exponentially distributed with parameter $\tilde{f}(v, T_i)$ and \tilde{M} does not evolve in Case 2. Otherwise, it follows since $\min\{C_v^0(i), C_v^1(i)\}$ is exponentially distributed with parameter $\tilde{f}(v, T_i)$. Now assume $V'_\alpha(T_i) \subseteq V_\alpha(T_i)$. To finish the proof, we will show $V'_\alpha(T_{i+1}) \subseteq V_\alpha(T_{i+1})$. Specifically, we will rule out the possibility that $w \in V'_\alpha(T_{i+1}) \setminus V_\alpha(T_{i+1})$.

In Case 0 the only way that $w \in V'_\alpha(T_{i+1})$ is that $v \in V'_\alpha(T_i)$. This implies $v \in V_\alpha(T_i)$ so $w \in V_\alpha(T_{i+1})$, as required.

In Case 1 the only way that $w \in V'_\alpha(T_{i+1}) \setminus V_\alpha(T_{i+1})$ is that $v \notin V_\alpha(T_i)$ so $v \notin V'_\alpha(T_i)$. But then Case 1 requires $\tilde{f}(v, T_i) > \tilde{f}'(v, T_i)$ and this is not possible since $v \notin V_\alpha(T_i)$ and $v \notin V'_\alpha(T_i)$ (given the condition about types β, β' in the lemma statement).

In Case 2 the only way that $w \in V'_\alpha(T_{i+1})$ is that $v \in V'_\alpha(T_i)$ so $v \in V_\alpha(T_i)$. But then we do not have $\tilde{f}(v, T_i) < \tilde{f}'(v, T_i)$ as required for Case 2. \square

Lemma 6 implies Corollary 9 which gives a lower bound on the probability that a maximally-fit mutation fixates in distributions in \mathcal{D} . In order to prove Corollary 9 we will use Corollary 7, which reduces the problem of lower-bounding the fixation probability in the multi-type Moran process to the problem of lower-bounding the fixation probability in the two-type Moran process. We will also use Lemma 8, which is a consequence of earlier work on the two-type Moran process.

Corollary 7. Let $G = (V, E)$ be a connected graph, let τ be a set of types, and let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function. Let α be a type in τ . Let $\beta = \arg \max\{f(j) \mid j \in \tau \setminus \{\alpha\}\}$. Let $\tau' = \{\alpha, \beta\}$. Let Ω be the set of functions from V to τ and let Ω' be the set of functions from V to τ' . Let $g: \Omega \rightarrow \Omega'$ be defined as follows. For every state $S \in \Omega$, $g(S)$ is the state such that $g(S)(v) = \alpha$ if $S(v) = \alpha$ and $g(S)(v) = \beta$, otherwise. Then, for every state $M_0 \in \Omega$, $\pi_\alpha(G, \tau, f, M_0) \geq \pi_\alpha(G, \tau', f, g(M_0))$.

Proof. There are three steps to the argument. First let $f': \tau \rightarrow \mathbb{Q}_{\geq 1}$ be the fitness function with $f'(\alpha) = \alpha$ and, for any $\beta' \in \tau \setminus \{\alpha\}$, $f'(\beta') = f(\beta)$. Note that for any $M_0 \in \Omega$, $\pi_\alpha(G, \tau', f, g(M_0)) = \pi_\alpha(G, \tau, f', M_0)$, so the goal is to prove $\pi_\alpha(G, \tau, f, M_0) \geq \pi_\alpha(G, \tau, f', M_0)$.

Second, let $\tilde{M} = \tilde{M}(G, \tau, f, M_0)$ and let $\tilde{M}' = \tilde{M}'(G, \tau, f', M_0)$. For all $j \in \tau$, let $\tilde{V}_j(t) = \{v \in V \mid \tilde{M}_t(v) = j\}$ and $\tilde{V}'_j(t) = \{v \in V \mid \tilde{M}'_t(v) = j\}$. Applying Lemma 6 we find that for any time $t \in \mathbb{R}_{\geq 0}$, $\tilde{V}'_\alpha(t) \subseteq \tilde{V}_\alpha(t)$. so $\tilde{\pi}_\alpha(G, \tau, f, M_0) \geq \tilde{\pi}_\alpha(G, \tau, f', M_0)$.

Finally, recall (from the construction of the continuous Moran process) that $\tilde{\pi}_\alpha(G, \tau, f, M_0) = \pi_\alpha(G, \tau, f, M_0)$ and $\tilde{\pi}_\alpha(G, \tau, f', M_0) = \pi_\alpha(G, \tau, f', M_0)$. Putting these together proves the corollary. \square

Lemma 8. *Let $G = (V, E)$ be a connected graph, let τ be a set of types with $|\tau| = 2$, and let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function. Let α be a type in τ_{f+} . Let D be a distribution in $\mathcal{D}(G, \tau)$. Then $\pi_\alpha(G, \tau, f, D) \geq 1/n$.*

Proof. Let τ_{ord} be the type in $\tau \setminus \{\alpha\}$. Recall the distribution $D_{\text{mut}} = D_{\text{mut}}(G, \tau, \tau_{\text{ord}})$ which is uniform on states $M_0 \in \Omega$ in which exactly one vertex has type α .

We will construct a coupling of D and D_{mut} such that, in the random sample (M_0, M_0^{mut}) drawn from the coupling, every state that is assigned type α in M_0^{mut} is also assigned type α in M_0 . We can then apply Theorem 6 of [6] to show that $\pi_\alpha(G, f, \tau, D) \geq \pi_\alpha(G, f, \tau, D_{\text{mut}})$, completing the proof.

To construct the coupling, consider the distribution D . By definition, D is defined as follows: Choose (\mathbf{u}, γ) uniformly at random from $V[2] \times \tau[2]$ and choose M_0 from the distribution $D_{\mathbf{u}, \gamma}$ on $\Omega(\mathbf{u}, \gamma)$. Let $n = |V|$. In the coupling, the outcome in which M_0^{mut} assigns vertex u to type α is coupled with all choices (\mathbf{u}, γ) such that one of the following occurs: (1) u is the first vertex of \mathbf{u} and α is the first type in γ and (2) u is the second vertex of \mathbf{u} and α is the second type in γ . Note that $1/n$ of the pairs (\mathbf{u}, γ) have this property. Also, every outcome in $D_{\mathbf{u}, \gamma}$ assigns α to u , as required. \square

Corollary 9. *Let $G = (V, E)$ be a connected graph, let τ be a set of types, and let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function. Let α be a type in τ_{f+} . Let D be a distribution in $\mathcal{D}(G, \tau)$. Then $\pi_\alpha(G, f, \tau, D) \geq 1/n$.*

Proof. Let $k = |\tau|$. If $k = 2$, the result follows directly from Lemma 8. Otherwise, Let $\beta = \arg \max \{f(j) \mid j \in \tau \setminus \{\alpha\}\}$. Let $\tau' = \{\alpha, \beta\}$. Let Ω be the set of functions from V to τ and let Ω' be the set of functions from V to τ' . Let $g: \Omega \rightarrow \Omega'$ be defined as follows. For every state $S \in \Omega$, $g(S)$ is the state such that $g(S)(v) = \alpha$ if $S(v) = \alpha$ and $g(S)(v) = \beta$, otherwise. Corollary 7 ensures that, for every state $M_0 \in \Omega$, $\pi_\alpha(G, \tau, f, M_0) \geq \pi_\alpha(G, \tau', f, g(M_0))$.

Let $g(D)$ be the distribution on Ω' induced by D and g . Specifically, to sample from $g(D)$, take a sample S from D and then output $g(S)$. Then $\pi_\alpha(G, \tau, f, D) \geq \pi_\alpha(G, \tau', f, g(D))$. To finish, we will show that $g(D) \in \mathcal{D}(G, \tau')$. This implies the corollary since then we can apply Lemma 8 to conclude $\pi_\alpha(G, \tau', f, g(D)) \geq 1/n$.

To prove that $g(D) \in \mathcal{D}(G, \tau')$, recall what it means that $D \in \mathcal{D}(G, \tau)$. This means that for every $\mathbf{u} \in V[k]$ and $\gamma \in \tau[k]$, there is a distribution $D_{\mathbf{u}, \gamma}$ on $\Omega(\mathbf{u}, \gamma)$ such that D can be defined as follows: Choose (\mathbf{u}, γ) uniformly at random from $V[k] \times \tau[k]$ and choose a sample S from $D_{\mathbf{u}, \gamma}$.

For every pair (\mathbf{u}', γ') in $V[2] \times \tau'[2]$ let $W_{\mathbf{u}', \gamma'}$ be the set of pairs $(\mathbf{u}, \gamma) \in V[k] \times \tau[k]$ such that \mathbf{u}' is a sub-tuple of \mathbf{u} and γ' is the same sub-tuple (using the same indices) of γ . Let $D'_{\mathbf{u}', \gamma'}$ be the following distribution: Choose (\mathbf{u}, γ) uniformly at random from $W_{\mathbf{u}', \gamma'}$, Choose S from $D_{\mathbf{u}, \gamma}$ and output $g(S)$. Let D' be the distribution defined as follows: Choose (\mathbf{u}', γ') uniformly at random from $V[2] \times \tau'[2]$ and sample an output from $D'_{\mathbf{u}', \gamma'}$. It is clear from construction that D' is in $\mathcal{D}(G, \tau')$. So to prove $g(D) \in \mathcal{D}(G, \tau')$ we need only show $g(D) = D'$. This follows from that fact that, in the execution of D' , each pair $(\mathbf{u}, \gamma) \in V[k] \times \tau[k]$ is equally likely to be chosen. \square

4 Bounding the Absorption Time

The goal of this section is to bound the absorption time of the multi-type Moran process. We will start from initial states $M_0 \in \Omega_0(G, \tau)$ since all distributions in $\mathcal{D}(G, \tau)$ are distributions on the state space $\Omega_0(G, \tau)$.

Definition 10 (Absorption time). Let $G = (V, E)$ be a connected graph, let τ be a set of types, let $f : \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function and let $M_0 \in \Omega_0(G, \tau)$ be an initial state. For the Moran process $M = M(G, \tau, f, M_0)$ and for any type $j \in \tau$ we define the stopping time

$$A_j(G, \tau, f, M_0) := \min\{t \in \mathbb{Z}_{\geq 0} \mid V_j(t) = V \text{ or } V_j(t) = \emptyset\}.$$

We refer to A_j as the absorption time of type j in M . The total absorption time of M is defined as

$$A(G, \tau, f, M_0) := \min\{t \in \mathbb{Z}_{\geq 0} \mid \text{for some } j \in \tau, V_j(t) = V\}.$$

If G, τ, f , and M_0 are clear from the context, we will just write A_j and A .

We will upper bound the absorption time for each type, which will then easily translate to an upper bound on the total absorption time. To this end, we start with the most fit type.

4.1 Expected absorption time of the fittest mutant

To upper bound the expected absorption time, we adopt the potential function for the 2-type process from the work of Díaz et al. [5, Section 3] adapted for a maximally-fit type in the multi-type process.

Definition 11. Let $G = (V, E)$ be a graph and let τ be a set of types. Let Ω be the set of functions from V to τ . For any type $j \in \tau$, the potential function Ψ_j is the function that maps every state $S \in \Omega$ to the value $\Psi_j(S) = \sum_{v:S(V)=j} 1/d(v)$.

Lemma 12 applies in situations where the most fit type is unique. The proof is almost identical to the proof of [5, Lemma 5], but we include the details for completeness.

Lemma 12. Let $G = (V, E)$ be a connected graph, let τ be a set of types with $|\tau| = k > 1$, and let $f : \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function. Let α be a type in τ_{f^+} . Let $f^* = \max\{f(j) \mid j \in \tau \setminus \alpha\}$, and suppose that $f^* < f(\alpha) = f^+$. Suppose that $V_\alpha(t) \notin \{\emptyset, V\}$ for a non-negative integer t . Then

$$\mathbb{E}(\Psi_\alpha(M_{t+1}) - \Psi_\alpha(M_t)) > (1 - f^*/f^+)/n^3.$$

Proof. First note that $\mathbb{E}(\Psi_\alpha(M_{t+1}) - \Psi_\alpha(M_t))$ is equal to

$$\begin{aligned} & \sum_{\substack{\{u,v\} \in E \\ u \in V_\alpha(t) \\ v \in V_j(t), j \neq \alpha}} \frac{f^+}{F(t)} \frac{1}{d(u)} (\Psi_\alpha(M_t|_{u \rightarrow v}) - \Psi_\alpha(M_t)) + \frac{f(j)}{F(t)} \frac{1}{d(v)} (\Psi_\alpha(M_t|_{v \rightarrow u}) - \Psi_\alpha(M_t)) \\ &= \sum_{\substack{\{u,v\} \in E \\ u \in V_\alpha(t) \\ v \in V_j(t), j \neq \alpha}} \frac{f^+}{F(t)} \frac{1}{d(u)} \frac{1}{d(v)} - \frac{f(j)}{F(t)} \frac{1}{d(v)} \frac{1}{d(u)} \end{aligned}$$

Next observe that $F(t) < nf^+$ and $f(j) \leq f^*$ for all $j \in \tau \setminus \alpha$. Thus

$$\mathbb{E}(\Psi_\alpha(M_{t+1}) - \Psi_\alpha(M_t)) > \frac{f^+ - f^*}{nf^+} \sum_{\substack{\{u,v\} \in E \\ u \in V_\alpha(t) \\ v \in V_j(t), j \neq \alpha}} \frac{1}{d(u)d(v)} > \left(1 - \frac{f^*}{f^+}\right)/n^3,$$

concluding the proof. \square

The technique that we use to upper bound $\mathbb{E}(A_\alpha(G, \tau, f, M_0))$ relies on martingale techniques which are made explicit by Hajek [11]. Díaz et. al [5, Theorem 6] used this technique to study the 2-type process, following the approach of He and Yao [12]. For convenience, we use here a multi-type generalisation of the statement from Goldberg et al. [10].

Lemma 13 ([10] Lemma 45). *Let Y be a Markov Chain with a finite state space Ω_Y . Let $c_1, c_2 > 0$, let $\Psi : \Omega_Y \rightarrow \mathbb{R}_{\geq 0}$ be a function, and let $A \geq 0$ be a stopping time with $A \leq \min\{t \mid \Psi(Y_t) = 0 \text{ or } \Psi(Y_t) \geq c_1\}$. Suppose that*

- (i) *from every state $S_1 \in \Omega_Y$ with $0 < \Psi(S_1) < c_1$, there is a path in Y from S_1 to some state S_2 with $\Psi(S_2) = 0$ or $\Psi(S_2) \geq c_1$;*
- (ii) *for all $t \geq 0$, if $\Psi(Y_t) < c_1$ then $\Psi(Y_{t+1}) \leq c_1 + 1$; and*
- (iii) *for all $t \geq 0$ and all $S \in \Omega_Y$ such that the events $A > t$ and $Y_t = S$ are consistent, $\mathbb{E}(\Psi(Y_{t+1}) - \Psi(Y_t) \mid Y_t = S) \geq c_2$.*

Then we have $\mathbb{E}(A) \leq (c_1 - \Psi(Y_0) + 1)/c_2$.

Corollary 14. *Let $G = (V, E)$ be a connected graph, let τ be a set of types with $|\tau| = k > 1$, and let $f : \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function. Let α be a type in τ_{f^+} . Let $f^* = \max\{f(j) \mid j \in \tau \setminus \alpha\}$, and suppose that $f^* < f(\alpha) = f^+$. Consider a state $M_0 \in \Omega_0(G, \tau)$. Then*

$$\mathbb{E}(A_\alpha(G, \tau, f, M_0)) \leq \frac{f^+}{f^+ - f^*} (n + 1) n^3.$$

Proof. Set $c_1 := \sum_{v \in V} 1/d(v)$ and $c_2 := (1 - f^*/f^+)/n^3$. Let $A_\alpha = A_\alpha(G, \tau, f, M_0)$ and let Ψ_α be the potential function from Definition 11. Let $M = M(G, \tau, f, M_0)$. By construction $A_\alpha = \min\{t \mid \Psi(M_t) = 0 \text{ or } \Psi(M_t) = c_1\}$. If $V_\alpha(0) \in \{V, \emptyset\}$ then $A_\alpha = 0$ and we are done. Otherwise, note that the state space Ω of M is finite and that Ψ_α is non-negative for every state. We now show that all three requirements (i) – (iii) of Lemma 13 are satisfied.

For (i), suppose $0 < \Psi(M_t) < c_1$ for some non-negative integer t . From the definition of A_α , $V_\alpha \notin \{\emptyset, V\}$. Since G is connected, there is a positive probability that M moves from M_t to a state S in which $V_\alpha = V$ so $\Psi(S) = c_1$. For (ii), note that the maximum potential is c_1 . Finally, for (iii), Lemma 12 ensures that $\mathbb{E}(\Psi_\alpha(M_{t+1}) - \Psi_\alpha(M_t)) > c_2 > 0$.

Lemma 13 thus implies $\mathbb{E}(A_\alpha) \leq (c_1 - \Psi_\alpha(M_0) + 1)/c_2 \leq (c_1 + 1)/c_2$. Since $c_1 \leq n$, the corollary follows. \square

4.2 Expected total absorption time

Now that we have an upper bound on the expectation of $A_\alpha(G, \tau, f, M_0)$ for the most-fit type α we use it to derive an upper bound on the expectation of $A_j(G, \tau, f, M_0)$ for any $j \in \tau$. Our strategy is as follows. If $j = \alpha$, we are finished. Otherwise after time $A_\alpha(G, \tau, f, M_0)$ there are two possibilities. If type α has fixated, then $A_j(G, \tau, f, M_0)$ occurs. Otherwise, type α has become extinct, so we repeat with the new set of types $\tau \setminus \alpha$. This inductive approach applies to the setting where the fitnesses of the types are pairwise distinct, which we will assume in the first part of this section. The general result (without this assumption) will be derived afterwards.

When we assume that the fitnesses are distinct we will typically assume that $\tau = \{1, \dots, k\}$ and that $f(1) < f(2) < \dots < f(k)$. This simplifies the notation with no loss of generality. Note that $\tau_{f^+} = \{k\}$ in this case.

Definition 15 ($A_{\geq j}$). *Let G be a connected graph. Let k be a positive integer. Let $\tau = \{1, \dots, k\}$. Let $f : \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function such that $f(1) < f(2) < \dots < f(k)$. Let M_0 be a state in $\Omega_0(G, \tau)$ and let $M = M(G, \tau, f, M_0)$. For any type $j \in \{2, \dots, k\}$, define*

$$A_{\geq j} := \min\{t \in \mathbb{Z}_{\geq 0} \mid \cup_{i \geq j} V_i(t) = \emptyset \text{ or, for some } i \geq j, V_i(t) = V\}.$$

$A_{\geq j}$ is thus the first point in time when either all types j, \dots, k are extinct or one of them has fixated.

Observation 16. *Let G be a connected graph. Let k be a positive integer. Let $\tau = \{1, \dots, k\}$. Let $f : \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function such that $f(1) < f(2) < \dots < f(k)$. Let M_0 be a state in $\Omega_0(G, \tau)$ and let $M = M(G, \tau, f, M_0)$. Then $A_{\geq j}$ is monotonically decreasing in j . That is, for all $j \in \{2, \dots, k\}$, $A_{\geq j} \geq A_{\geq j+1}$.*

Proof. At time $t = A_{\geq j}$, either $\cup_{i \geq j} V_i(t) = \emptyset$ or, for some $i \geq j$, $V_i(t) = V$. The only case that doesn't directly imply $A_{\geq j+1} \leq t$ (according to the definition) is $V_j(t) = V$ but this case implies $\cup_{i \geq j+1} V_i(t) = \emptyset$, hence $A_{\geq j+1} \leq t$. \square

Lemma 17. *Let G be a connected graph with n vertices. Let k be a positive integer. Let $\tau = \{1, \dots, k\}$. Let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function such that $f(1) < f(2) < \dots < f(k)$. Let M_0 be a state in $\Omega_0(G, \tau)$ and let $M = M(G, \tau, f, M_0)$. Let j be a type in $\{2, \dots, k\}$. Then*

$$\mathbb{E}(A_{\geq j}) \leq \sum_{i=j}^k \frac{f(i)}{f(i) - f(i-1)} (n+1)n^3.$$

Proof. The lemma is proved by induction on $k - j$. For the base case $j = k$ note that $A_{\geq k} = A_k$. Thus the claim holds for $j = k$ by Corollary 14.

For the inductive step, assume that the claim holds for $j + 1$; we will show that it holds for j .

To this end, let Ω^j be the set of all states with no vertices assigned to types $\{j + 1, \dots, k\}$. Using Corollary 14 on initial states in Ω^j , we obtain

$$\begin{aligned} \mathbb{E}(A_{\geq j} - A_{\geq j+1}) &\leq \sum_{S \in \Omega^j} \mathbb{E}(A_{\geq j} - A_{\geq j+1} \mid M_{A_{\geq j+1}} = S) \cdot \Pr(M_{A_{\geq j+1}} = S) \\ &\leq \frac{f(j)}{f(j) - f(j-1)} (n+1)n^3 \cdot \underbrace{\sum_{S \in \Omega^j} \Pr(M_{A_{\geq j+1}} = S)}_{\leq 1}. \end{aligned}$$

Consequently, in combination with the induction hypothesis for $j + 1$, we have

$$\begin{aligned} \mathbb{E}(A_{\geq j}) &= \mathbb{E}(A_{\geq j} - A_{\geq j+1}) + \mathbb{E}(A_{\geq j+1}) \\ &\leq \frac{f(j)}{f(j) - f(j-1)} (n+1)n^3 + \mathbb{E}(A_{\geq j+1}) \\ &\leq \frac{f(j)}{f(j) - f(j-1)} (n+1)n^3 + \sum_{i=j+1}^k \frac{f(i)}{f(i) - f(i-1)} (n+1)n^3, \end{aligned}$$

concluding the proof. \square

In the final step, we extend our bounds on the absorption times to processes in which distinct types may have the same fitness.

Theorem 18. *Let $G = (V, E)$ be a connected graph with n vertices, let τ be a set of types with $|\tau| > 1$, let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function, and let M_0 be a state in $\Omega_0(G, \tau)$. Suppose that the types in τ have k distinct fitnesses and let $f^- = f_1 < f_2 < \dots < f_k = f^+$ be these fitnesses. For each $i \in \{1, \dots, k\}$, let $\ell_i = |f^{-1}(f_i)|$ be the number of types with fitness f_i and let $\ell = \max\{\ell_i \mid i \in \{1, \dots, k\}\}$ be the maximum number of types that share the same fitness. Then*

(1) *The total absorption time is bounded by*

$$\mathbb{E}(A(G, \tau, f, M_0)) \leq (\ell - 1)n^6 + \sum_{i=2}^k \frac{f_i}{f_i - f_{i-1}} (n+1)n^3.$$

(2) *For any $j \in \{2, \dots, k\}$ the absorption time of each type β with $f(\beta) = f_j$ is bounded by*

$$\mathbb{E}(A_\beta(G, \tau, f, M_0)) \leq n^6 + \sum_{i=j}^k \frac{f_i}{f_i - f_{i-1}} (n+1)n^3.$$

(3) The absorption time of each type β with $f(\beta) = f_1$ is bounded by

$$\mathbb{E}(A_\beta(G, \tau, f, M_0)) \leq n^6 + \sum_{i=2}^k \frac{f_i}{f_i - f_{i-1}} (n+1)n^3.$$

Proof. We first treat types with equal fitnesses as the same type. Let $\tau' = \{1, \dots, k\}$ and let $f': \tau' \rightarrow \mathbb{Q}_{\geq 1}$ be the fitness function such that, for each $i \in \tau'$, $f'(i) = f_i$. Let Ω be the set of functions from V to τ and let Ω' be the set of functions from V to τ' . As usual, let $\Omega_0(G, \tau)$ be the set of functions in Ω with range τ and let $\Omega'_0(G, \tau')$ be the set of functions in Ω' with range τ' .

We will define a function g from Ω to Ω' . Given $S \in \Omega$, let $g(S)$ be the state in Ω' defined as follows. For each $v \in V$, identify the value $i \in \{1, \dots, k\}$ such that $f(v) = f_i$. Then set $g(S)(v) = i$. Note that if $S \in \Omega_0(G, \tau)$ then $S' \in \Omega'_0(G, \tau')$. Let $M'_0 = g(M_0)$.

Let $M = M(G, \tau, f, M_0)$ and $M' = M'(G, \tau', f', M'_0)$. For any $j \in \tau'$, let $V'_j(t) = \{v \in V \mid M'_t(v) = j\}$. Following Definition 10 and 15, define

$$\begin{aligned} A'_j &= \min\{t \in \mathbb{Z}_{\geq 0} \mid V'_j(t) = V \text{ or } V'_j(t) = \emptyset\}, \\ A' &= \min\{t \in \mathbb{Z}_{\geq 0} \mid \text{for some } j \in \tau, V'_j(t) = V\}, \text{ and} \\ A'_{\geq j} &= \min\{t \in \mathbb{Z}_{\geq 0} \mid \cup_{i \geq j} V'_i(t) = \emptyset \text{ or, for some } i \geq j, V'_i(t) = V\}. \end{aligned}$$

Since $A'_{\geq 2} = A'$, Lemma 17 applied to M' implies

$$\mathbb{E}(A') \leq \sum_{i=2}^k \frac{f_i}{f_i - f_{i-1}} (n+1)n^3. \quad (1)$$

Note that $g(M_0), g(M_1), \dots$ is a faithful copy of the chain M' . If M and M' are coupled using this coupling, then at time A' , all vertices are assigned types with the same fitness in $M_{A'}$.

Consider any state $S \in \Omega$ in which vertices are assigned q different types, all with the same fitness. Theorem 11 of the work of Díaz et al. [5] shows that in the 2-type Moran process, from any state with 2 types, the expected time until there is at most one type is at most n^6 .

Applying this argument to S by distinguishing the first type from the $q - 1$ others, the expected time until there are at most $q - 1$ types is at most n^6 . Since there are at most ℓ distinct types with any given fitness, $\mathbb{E}(A(G, \tau, f, S)) \leq (\ell - 1)n^6$. The upper bound (1) on $\mathbb{E}(A(G, \tau, f, M_0))$ given in the statement of the theorem follows.

In order to prove (2), note that $A'_j \leq A'_{\geq j}$. Lemma 17 applied to M' implies

$$\mathbb{E}(A'_{\geq j}) \leq \sum_{i=j}^k \frac{f_i}{f_i - f_{i-1}} (n+1)n^3.$$

At time $\mathbb{E}(A'_{\geq j})$, there are three possibilities. (i) Possibly there are no remaining types with fitness at least f_j . (ii) Possibly for some $i > j$ all remaining types have fitness f_i . (iii) Alternatively, all remaining types have fitness f_j . In the first two cases, $A_\beta(G, \tau, f, M_0)$ has occurred. In the third case, viewing type β separately to the other remaining types, the expected time until $A_\beta(G, \tau, f, M_0)$ occurs is at most n^6 more steps, proving item (2).

Finally, for the proof of (3) observe that after time A' , either type β is extinct, or all remaining types have fitness $f(\beta) = f_1$. The argument is then the same as Item (2). This concludes the proof. \square

The $\ell - 1$ in Item (1) of Theorem 18 could be replaced with $2\log(\ell)$, which is an improvement for $\ell \geq 7$. The way to do this is to split the q initial types into two groups of at most $q/2 + 1$ and iterating.

Bounding the absorption time from an initial state chosen from a distribution in $\mathcal{D}(G, \tau)$ is an immediate consequence of Theorem 18.

Definition 19 (Absorption time from random initial state). Let $G = (V, E)$ be a connected graph, let τ be a set of types, let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function and let $D \in \mathcal{D}(G, \tau)$. For the Moran process $M = M(G, \tau, f, D)$ and for any type $j \in \tau$ we define the stopping time

$$A_j = A_j(G, \tau, f, D) := \min\{t \in \mathbb{Z}_{\geq 0} \mid V_j(t) = V \text{ or } V_j(t) = \emptyset\}.$$

We refer to A_j as the absorption time of type j in M from D . The total absorption time of M from D is defined as

$$A = A(G, \tau, f, D) := \min\{t \in \mathbb{Z}_{\geq 0} \mid \text{for some } j \in \tau, V_j(t) = V\}.$$

Since the support of any distribution $D \in \mathcal{D}(G, \tau)$ is contained in $\Omega_0(G, \tau)$, the bounds in Theorem 18 also hold if M_0 is chosen from D . We make this explicit for the most fit type, since this is the result that we will use to establish our FPTRAS.

Corollary 20. Let G be a connected graph with n vertices, let τ be a set of types with $|\tau| > 1$, let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function, and let $D \in \mathcal{D}(G, \tau)$. Let $f^* = \max\{f(j) \mid j \in \tau \setminus \tau_{f+}\}$. The total absorption time of any $\alpha \in \tau_{f+}$ is bounded by

$$\mathbb{E}(A_\alpha(G, \tau, f, D)) \leq n^6 + \frac{f(\alpha)}{f(\alpha) - f^*}(n + 1)n^3.$$

Proof. Note that $\mathbb{E}(A_\alpha(G, \tau, f, D)) = \sum_{S \in \Omega} \Pr_D(S) E(A_\alpha(G, \tau, f, S))$. Since the support of D is contained in $\Omega_0(G, \tau)$, the corollary follows from Theorem 18. \square

5 Construction of the FPTRAS

Theorem 4 (restated). MAXFIXPROB has an FPTRAS when parameterised by $\kappa := \|f\|$.

Proof. Consider first the following algorithm \mathbb{A}' that takes as input a connected graph G , a set of types τ with $|\tau| > 1$, a fitness function $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$, a type $\alpha \in \tau_{f+}$, accuracy parameters ε and δ' , and (via oracle access) a distribution $D \in \mathcal{D}(G, \tau)$. By Corollary 9, $\pi_\alpha(G, \tau, f, D) \geq 1/n$.

Set $t = \lceil 3 \ln(2/\delta')n/\varepsilon^2 \rceil$. The algorithm \mathbb{A}' performs t independent simulations of $M(G, \tau, f, D)$, stopping the simulation if either (i) all vertices have type α , or (ii) type α is extinct. In each simulation, the initial state is sampled using the oracle for D .

Let $f^* = \max\{f(j) \mid j \in \tau \setminus \tau_{f+}\}$. Corollary 20 implies that the expected time of each simulation is at most $\text{poly}(n) \cdot f(\alpha)/(f(\alpha) - f^*)$. Let w be the number of simulations in which type α fixates. The output of algorithm \mathbb{A}' is w/t .

For $i \in \{1, \dots, t\}$ let X_i be the indicator for the event that α fixates in the i 'th simulation. Since $\mathbb{E}(X_i) = \pi_\alpha(G, \tau, f, D) \geq 1/n$, it follows that $\mathbb{E}(w/t) = \mathbb{E}(\frac{1}{t} \sum_{i=1}^t X_i) = \pi_\alpha(G, \tau, f, D)$. By a Chernoff bound (see e.g. the standard textbook of Mitzenmacher and Upfal [15, Theorem 11.1]),

$$\Pr(|w/t - \pi_\alpha(G, \tau, f, D)| < \varepsilon \pi_\alpha(G, \tau, f, D)) \geq 1 - \delta'.$$

The expected total running time of \mathbb{A}' is bounded from above by $t \cdot \text{poly}(n) \cdot f(\alpha)/(f(\alpha) - f^*)$. We next obtain an ε -approximation algorithm \mathbb{A} from \mathbb{A}' using the following standard construction. The algorithm \mathbb{A} takes as input $G, \tau, f, \alpha, \varepsilon$, and (via oracle access) the distribution D . The algorithm \mathbb{A} simulates \mathbb{A}' with these inputs and with $\delta' = 1/8$ for at most $8t \cdot \text{poly}(n) \cdot f(\alpha)/(f(\alpha) - f^*)$ steps. If \mathbb{A}' terminates within this number of steps, then algorithm \mathbb{A} returns its output. Otherwise, algorithm \mathbb{A} stops the simulation and returns 0.

By Markov's inequality, the probability that \mathbb{A} stops the simulation is at most $1/8$. Thus, the probability that \mathbb{A} returns an ε -approximation to $\pi_\alpha(G, \tau, f, D)$ is at least $1 - \delta' - 1/8 \geq 3/4$.

The desired FPTRAS takes the median of the outputs after repeating algorithm \mathbb{A} $O(\log(1/\delta))$ times. \square

Theorem 5 (restated). MAXFIXPROB_{mut} has an FPTRAS when parameterised by $\kappa := \|f\|$.

Proof. The proof is the same as the proof of Theorem 4 except that the algorithms can directly sample from $D_{\text{mut}}(G, \tau, \tau_{\text{ord}})$ in polynomial time, without using the oracle. \square

6 Upper-bounding fixation probabilities in terms of the 2-type Moran process

Corollary 7 gives a lower bound on the fixation probability of a type α in the multi-type Moran process, based on the fixation probability of α in a related two-type Moran process. It is just as easy to get upper bounds with this idea, as the following corollary shows.

Corollary 21. *Let $G = (V, E)$ be a connected graph, let τ be a set of types, and let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function. Let α be a type in τ . Let $\beta = \arg \min\{f(j) \mid j \in \tau \setminus \{\alpha\}\}$. Let $\tau' = \{\alpha, \beta\}$. Let Ω be the set of functions from V to τ and let Ω' be the set of functions from V to τ' . Let $g: \Omega \rightarrow \Omega'$ be defined as follows. For every state $S \in \Omega$, $g(S)$ is the state such that $g(S)(v) = \alpha$ if $S(v) = \alpha$ and $g(S)(v) = \beta$, otherwise. Then, for every state $M_0 \in \Omega$, $\pi_\alpha(G, \tau, f, M_0) \leq \pi_\alpha(G, \tau', f, g(M_0))$.*

Proof. Let $f': \tau \rightarrow \mathbb{Q}_{\geq 1}$ be the fitness function with $f'(\alpha) = \alpha$ and, for any $\beta' \in \tau \setminus \{\alpha\}$, $f'(\beta') = f(\beta)$. Note that $\pi_\alpha(G, \tau', f, g(M_0)) = \pi_\alpha(G, \tau, f', M_0)$, so the goal is to prove $\pi_\alpha(G, \tau, f, M_0) \leq \pi_\alpha(G, \tau, f', M_0)$.

Let $\tilde{M} = \tilde{M}(G, \tau, f, M_0)$ and let $\tilde{M}' = \tilde{M}'(G, \tau, f', M_0)$. For all $j \in \tau$, let $\tilde{V}_j(t) = \{v \in V \mid \tilde{M}_t(v) = j\}$ and $\tilde{V}'_j(t) = \{v \in V \mid \tilde{M}'_t(v) = j\}$. Applying Lemma 6 (with the roles of f and f' reversed) we find that for any time $t \in \mathbb{R}_{\geq 0}$, $\tilde{V}'_\alpha(t) \supseteq \tilde{V}_\alpha(t)$. so $\tilde{\pi}_\alpha(G, \tau, f, M_0) \leq \tilde{\pi}_\alpha(G, \tau, f', M_0)$. As in the proof of Corollary 7 the final step is to observe that $\tilde{\pi}_\alpha(G, \tau, f, M_0) = \pi_\alpha(G, \tau, f, M_0)$ and $\tilde{\pi}_\alpha(G, \tau, f', M_0) = \pi_\alpha(G, \tau, f', M_0)$. \square

We conclude the paper by using Corollaries 7 and 21 to re-derive the recent bounds of Ferreira and Neves on fixation probabilities of the 3-type Moran process in the very special case where G is a complete graph. Their main result on the Moran process [8, Theorem 4] corresponds to the special case of Lemma 22 with $\tau = \{A, B, C\}$ and $f(A) > f(B) > f(C)$.

Lemma 22. *Let n be a positive integer and let $G = (V, E)$ be a complete graph with n vertices. Let τ be a set of types and let $f: \tau \rightarrow \mathbb{Q}_{\geq 1}$ be a fitness function. Let α be a type in τ . Let $\beta^+ = \arg \max\{f(\beta) \mid \beta \in \tau \setminus \{\alpha\}\}$ and $\beta^- = \arg \min\{f(\beta) \mid \beta \in \tau \setminus \{\alpha\}\}$. Suppose that $f(\alpha) \neq f(\beta^+)$ and $f(\alpha) \neq f(\beta^-)$. Let M_0 be a state in the set Ω of functions from V to τ . Let i be the number of vertices in V which are mapped to type α by M_0 . Then*

$$\frac{1 - \left(\frac{f(\beta^+)}{f(\alpha)}\right)^i}{1 - \left(\frac{f(\beta^+)}{f(\alpha)}\right)^n} \leq \pi_\alpha(G, \tau, f, M_0) \leq \frac{1 - \left(\frac{f(\beta^-)}{f(\alpha)}\right)^i}{1 - \left(\frac{f(\beta^-)}{f(\alpha)}\right)^n}.$$

Proof. We first give the lower bound. Let $\beta = \beta^+$. Let $\tau' = \{\alpha, \beta\}$. Let Ω' be the set of functions from V to τ' . Let $g: \Omega \rightarrow \Omega'$ be defined as follows. For every state $S \in \Omega$, $g(S)$ is the state such that $g(S)(v) = \alpha$ if $S(v) = \alpha$ and $g(S)(v) = \beta$, otherwise. Corollary 7 shows that $\pi_\alpha(G, \tau, f, M_0) \geq \pi_\alpha(G, \tau', f, g(M_0))$. To conclude the proof, we note that $\pi_\alpha(G, \tau', f, g(M_0))$ is equal to the lower bound given in the lemma statement. This well-known fact follows because $\pi_\alpha(G, \tau', f, g(M_0))$ is the probability of hitting the lower boundary in a biased one-dimensional random walk. The walk is biased since $f(\alpha) \neq f(\beta)$. See Lieberman et al. [14] for details. The upper bound is the same, except $\beta = \beta^-$ and Corollary 21 is used instead of Corollary 7. \square

References

- [1] Vikraman Arvind and Venkatesh Raman. Approximation algorithms for some parameterized counting problems. In Prosenjit Bose and Pat Morin, editors, *Algorithms and Computation, 13th International Symposium, ISAAC 2002 Vancouver, BC, Canada, November 21-23, 2002, Proceedings*, volume 2518 of *Lecture Notes in Computer Science*, pages 453–464. Springer, 2002. doi:10.1007/3-540-36136-7_40.

- [2] Niko Beerenwinkel, Tibor Antal, David Dingli, Arne Traulsen, Kenneth W. Kinzler, Victor E. Velculescu, Bert Vogelstein, and Martin A. Nowak. Genetic Progression and the Waiting Time to Cancer. *PLOS Computational Biology*, 3(11):e225, November 2007. doi:[10.1371/journal.pcbi.0030225](https://doi.org/10.1371/journal.pcbi.0030225).
- [3] Brian Chao and Jason Schweinsberg. A Spatial Mutation Model with Increasing Mutation Rates. *CoRR*, abs/2108.09590, 2021. arXiv:[2108.09590](https://arxiv.org/abs/2108.09590), doi:[10.48550/arXiv.2108.09590](https://doi.org/10.48550/arXiv.2108.09590).
- [4] Krishnendu Chatterjee, Rasmus Ibsen-Jensen, and Martin A. Nowak. Faster monte-carlo algorithms for fixation probability of the moran process on undirected graphs. In *42nd International Symposium on Mathematical Foundations of Computer Science*, volume 83 of *LIPICs*, pages 61:1–61:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:[10.4230/LIPICs.MFCS.2017.61](https://doi.org/10.4230/LIPICs.MFCS.2017.61).
- [5] Josep Díaz, Leslie Ann Goldberg, George B. Mertzios, David Richerby, Maria Serna, and Paul G. Spirakis. Approximating Fixation Probabilities in the Generalized Moran Process. *Algorithmica*, 69(1):78–91, May 2014. doi:[10.1007/s00453-012-9722-7](https://doi.org/10.1007/s00453-012-9722-7).
- [6] Josep Díaz, Leslie Ann Goldberg, David Richerby, and Maria Serna. Absorption time of the Moran process. *Random Structures & Algorithms*, 49(1):137–159, 2016. doi:[10.1002/rsa.20617](https://doi.org/10.1002/rsa.20617).
- [7] A. M. Etheridge and R. C. Griffiths. A coalescent dual process in a Moran model with genic selection. *Theoretical Population Biology*, 75(4):320–330, 2009. doi:[10.1016/j.tpb.2009.03.004](https://doi.org/10.1016/j.tpb.2009.03.004).
- [8] Eliza M. Ferreira and Armando G. M. Neves. Fixation probabilities for the Moran process with three or more strategies: General and coupling results. *Journal of Mathematical Biology*, 81(1):277–314, July 2020. doi:[10.1007/s00285-020-01510-0](https://doi.org/10.1007/s00285-020-01510-0).
- [9] R. A. Fisher. On the Dominance Ratio. *Proceedings of the Royal Society of Edinburgh*, 42:321–341, 1923. doi:[10.1017/S0370164600023993](https://doi.org/10.1017/S0370164600023993).
- [10] Leslie Ann Goldberg, John Lapinskas, and David Richerby. Phase transitions of the Moran process and algorithmic consequences. *Random Structures & Algorithms*, 56(3):597–647, 2020. doi:[10.1002/rsa.20890](https://doi.org/10.1002/rsa.20890).
- [11] Bruce Hajek. Hitting-Time and Occupation-Time Bounds Implied by Drift Analysis with Applications. *Advances in Applied Probability*, 14(3):502–525, 1982. doi:[10.2307/1426671](https://doi.org/10.2307/1426671).
- [12] Jun He and Xin Yao. Drift analysis and average time complexity of evolutionary algorithms. *Artificial Intelligence*, 127(1):57–85, 2001. doi:[10.1016/S0004-3702\(01\)00058-3](https://doi.org/10.1016/S0004-3702(01)00058-3).
- [13] Nicolas Lanchier. Wright–Fisher and Moran models. In Nicolas Lanchier, editor, *Stochastic Modeling*, Universitext, pages 203–218. Springer International Publishing, 2017. doi:[10.1007/978-3-319-50038-6_12](https://doi.org/10.1007/978-3-319-50038-6_12).
- [14] Erez Lieberman, Christoph Hauert, and Martin A. Nowak. Evolutionary dynamics on graphs. *Nature*, 433(7023):312–316, January 2005. doi:[10.1038/nature03204](https://doi.org/10.1038/nature03204).
- [15] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005. doi:[10.1017/CB09780511813603](https://doi.org/10.1017/CB09780511813603).
- [16] P. A. P. Moran. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society*, 54(1):60–71, 1958. doi:[10.1017/S0305004100033193](https://doi.org/10.1017/S0305004100033193).
- [17] Martin A. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, 2006. doi:[10.2307/j.ctvjghw98](https://doi.org/10.2307/j.ctvjghw98).
- [18] Sewall Wright. Evolution in Mendelian Populations. *Genetics*, 16(2):97–159, 1931. doi:[10.1093/genetics/16.2.97](https://doi.org/10.1093/genetics/16.2.97).