



A new algebraic approach to genome rearrangement models

Venta Terauds¹ · Jeremy Sumner¹

Received: 21 December 2020 / Revised: 24 January 2022 / Accepted: 31 March 2022 /

Published online: 5 May 2022

© The Author(s) 2022

Abstract

We present a unified framework for modelling genomes and their rearrangements in a genome algebra, as elements that simultaneously incorporate all physical symmetries. Building on previous work utilising the group algebra of the symmetric group, we explicitly construct the genome algebra for the case of unsigned circular genomes with dihedral symmetry and show that the maximum likelihood estimate (MLE) of genome rearrangement distance can be validly and more efficiently performed in this setting. We then construct the genome algebra for a more general case, that is, for genomes that may be represented by elements of an arbitrary group and symmetry group, and show that the MLE computations can be performed entirely within this framework. There is no prescribed model in this framework; that is, it allows any choice of rearrangements that preserve the set of regions, along with arbitrary weights. Further, since the likelihood function is built from path probabilities—a generalisation of path counts—the framework may be utilised for any distance measure that is based on path probabilities.

Mathematics Subject Classification 92D15 · 20C05

This work was supported by Australian Research Council Discovery Grant DP180102215. The authors would like to thank Andrew Francis and Joshua Stevenson, for many interesting discussions relating to this work, and the anonymous reviewers, for their constructive comments.

✉ Venta Terauds
venta.terauds@utas.edu.au

Jeremy Sumner
jeremy.sumner@utas.edu.au

¹ Discipline of Mathematics, School of Natural Sciences, University of Tasmania, Private Bag 37, Sandy Bay, TAS 7001, Australia

1 Introduction

In the eight decades since Dobzhansky and Sturtevant observed that differences in fruit fly genomes could be explained by a sequence of reversals of genome segments (Dobzhansky and Sturtevant 1938), the study of evolution via genome rearrangement has developed into a rich and active field, with diverse applications (Chen et al. 2018; Darmon and Leach 2014; Oesper et al. 2017). Much work focuses on the calculation of evolutionary distances under rearrangement models, with the distances subsequently used to reconstruct phylogenetic trees. For example, minimal rearrangement distances between genomes—and other distance estimates based on these—have been studied extensively and, under various model restrictions, can be calculated efficiently (Bader et al. 2001; Wang et al. 2006; Bader and Ohlebusch 2006; Oliveira et al. 2019). There are, however, good arguments for applying stochastic methods that estimate genomic distance, via rearrangement, as evolutionary time elapsed (Serdoz et al. 2017), particularly when such an approach allows various rearrangement models to be considered (Terauds and Sumner 2019).

The maximum likelihood approach detailed in Serdoz et al. (2017) utilised the theory of the symmetric and dihedral groups to model circular genomes and region set-conserving rearrangements, motivated by earlier group-theoretical approaches to rearrangement models (Francis 2014; Egri-Nagy et al. 2014). The combinatorial problem of calculating the maximum likelihood estimate (MLE) of evolutionary distance was then converted into a numerical one in Sumner et al. (2017) via the representation theory of the symmetric group algebra. In Terauds and Sumner (2019), the consideration of symmetry was extended to include symmetry of rearrangement models, the role of this in simplifying calculations was explored, and the concrete implementation of the technique for a general model was described. Whilst the representation theory approach reduces the complexity of the MLE computations, the complexity is still of factorial order, meaning that computations for large number of regions remain, for the moment, out of reach.

In this work, we suggest that the appropriate theoretical setting for such MLE computations is in fact not the symmetric group algebra but a lower-dimensional algebra. In the symmetric group algebra, the basis elements for computations are individual permutations, each representing a rearrangement or a genome in a fixed orientation, and symmetry is incorporated as an extra step in the calculations. To simplify this, we construct an algebra that incorporates the inherent symmetry into each element. Here, the basis elements are *permutation clouds*. These correspond to genomes, by simultaneously including all physical orientations; due to the corresponding symmetries of the rearrangement model, they also represent rearrangements in a natural way.

This approach explains and removes the redundancy in the MLE computations that was observed in Terauds and Sumner (2019). In developing the approach, we firstly focus on the simple concrete case of uni-chromosomal circular genomes modelled with unoriented regions and no distinguished positions, building on previous work (Serdoz et al. 2017; Sumner et al. 2017; Terauds and Sumner 2019). Subsequently, we demonstrate that our results may be applied more generally, for example to genomic models that include region orientation and/or origin and terminus of replication. Further, although our focus is on calculation of MLEs, our approach can be applied to

calculate other measures of genomic distance under rearrangement; in particular, any that utilise path counts or weighted path counts, such as minimum distance. Whilst the framework does not specify a rearrangement model—indeed, one may choose the allowed rearrangements and their weights—we note that the group-based approach limits us to rearrangements that conserve the set of genomic regions, and thus cannot accommodate insertions, deletions or duplications. We are currently working on expanding the framework to a semigroup-based approach that could incorporate at least some of these rearrangement types. Some of the algebra easily extends to the semigroup case (see Remark 4.7, for example), however there is much yet to be done and this is outside the scope of the present paper.

In the next section, we outline the details of the symmetric group algebra approach to calculating MLEs (Sumner et al. 2017; Terauds and Sumner 2019) for pairs of unsigned, circular genomes, which forms the foundation for the current work. Following this, in Sect. 3, we construct the genome algebra, based on permutation clouds, for this case and show that it provides a coherent framework for modelling such genomes and region set-conserving rearrangements and for calculating MLEs. Section 4 outlines the extension of our results and techniques from permutations with dihedral symmetry to an arbitrary group and symmetry group. This verifies that, as well as incorporating flexibility in the rearrangement model, the framework is not specific to one particular genomic model. The paper concludes with a brief discussion section.

2 Background: the permutation approach

In this section we set out the theoretical framework for rearrangement models based on permutations, and recall the key elements of the technique for calculating the maximum likelihood estimate of evolutionary distance. Full derivations and details may be found in Terauds and Sumner (2019) and the earlier papers (Sumner et al. 2017; Serdoz et al. 2017). For the specific case study in this and the next section, we model the evolution of single-strand, circular genomes; we do not consider the regions to be oriented and do not distinguish any positions.¹ Genomes that are to be compared share N identified regions² of interest and we consider only rearrangements that conserve the set of regions. Accordingly, we use unsigned permutations, that is, elements of the symmetric group, \mathcal{S}_N , to represent both genomes and rearrangements. Explicitly: the regions and positions are each labelled by the integers $\{1, 2, \dots, N\}$, and a given genome is represented by a permutation $\sigma \in \mathcal{S}_N$, where

$$\sigma(i) = j \iff \text{region } i \text{ is in position } j.$$

Note that while the region labels are chosen once and are immutable, the position labelling reflects a choice of reference frame (starting position and direction of numbering) that changes when we move the genome in space. Since we do not distinguish any positions, there are $2N$ possible choices of reference frame and thus $2N$ distinct permutations that represent any given genome; we denote these by

¹ These latter two may be considered simplifying assumptions, since they reduce the size of the state space. However, as will be shown later, the framework just as easily accommodates the alternate cases.

² Here, a *region* is a contiguous section of the genome such as a sequence of genes (for an example of how such genomic simplification may be enacted in practice, see Belda et al. 2005).

$$[\sigma] := \{d\sigma : d \in \mathcal{D}_N\}. \quad (1)$$

Here, \mathcal{D}_N is the dihedral group, and the genome has *dihedral symmetry*.

Since \mathcal{D}_N is a subgroup of \mathcal{S}_N , the sets $[\sigma]$ are cosets, that is, each $[\sigma] = \mathcal{D}_N\sigma$ is an equivalence class of \mathcal{S}_N . Since a given genome exists independently of its orientation in space, we may identify it with the entire coset (Serdoz et al. 2017; Egri-Nagy et al. 2014). However, in this initial formulation, we choose any one of the permutations from the coset (1), say σ , to specify the genome, work with this single permutation at first, and incorporate all permutations in the set $[\sigma]$ (all symmetries of the genome) into the likelihood calculations in due course.

We model evolution as a sequence of discrete rearrangement events occurring in continuous time. In this section, as in previous work, we consider a rearrangement to be a single permutation acting on a single permutation; in the next section we shall develop this into the notion of permutation clouds acting on permutation clouds. For now, however, for a genome represented by $\sigma \in \mathcal{S}_N$, a rearrangement event is represented by a permutation, $a \in \mathcal{S}_N$, acting on σ (on the left): $\sigma \mapsto a\sigma$. We refer to permutations a acting in this way as “rearrangements”.

The full biological model for evolution is given by $(\mathcal{M}, w, dist)$, where $\mathcal{M} \subseteq \mathcal{S}_N$ is the set of allowed rearrangements, $w : \mathcal{M} \rightarrow (0, 1]$ is the probability distribution on this set, and $dist$ is the probability distribution of the independent rearrangement events in time. One may have biological evidence for including particular types and sizes of rearrangements in the model with differing relative probabilities, or may wish to compare distances computed under differing models (see Terauds et al. 2021 for some specific examples of models and distance comparisons). The distribution $dist$ may similarly be chosen according to evidence or preference; in this treatment, we use the Poisson distribution.

We shall emphasise at this stage that we make minimal further restrictions on the set of allowed rearrangements \mathcal{M} . Without loss of generality, we assume that \mathcal{M} generates the group \mathcal{S}_N ; this means that any permutation in \mathcal{S}_N may be obtained from any other by applying a sequence of elements from \mathcal{M} . (Note that the case of \mathcal{M} not generating \mathcal{S}_N is simpler: in this case \mathcal{M} generates a subgroup, $H \subseteq \mathcal{S}_N$, and the problem reduces to considering this smaller group, since any pair of elements would simply be unrelated under the model or both be elements of a coset $H\sigma$.) Elements of \mathcal{D}_N are, formally, allowed in the set of rearrangements, although their action does not actually alter the genome. This allows for full generality; for example, if one wishes to include ‘all inversions’ in the model, then the inversion of a region of size $N - 1$ is the same as flipping the genome over in space.

The first model condition simply states that the model should naturally possess the same symmetry as the genome (in the current case, dihedral symmetry). Suppose, for example, that $(1, 2) \in \mathcal{M}$, meaning that the regions in positions 1 and 2 may swap places. Then, since the position labelling is arbitrary, we should have $(\ell, \ell + 1) \in \mathcal{M}$ for all $\ell = 1, 2, \dots, N - 1$, and $(N, 1) \in \mathcal{M}$, meaning that *any* two regions in adjacent positions may swap places; further, these rearrangements should all be equally probable. We refer to this property as *dihedral symmetry of the model* and, mathematically, express the condition as

$$\text{for each } a \in \mathcal{M} \text{ and } d \in \mathcal{D}_N, \quad dad^{-1} \in \mathcal{M} \text{ and } w(dad^{-1}) = w(a). \quad (\text{M1})$$

The second model condition ensures that the modelling is agnostic to the temporal direction of evolution. More precisely, the condition states that for any rearrangement that is allowed, its inverse is also allowed, with the same probability. That is,

$$\text{for each } a \in \mathcal{M}, \quad a^{-1} \in \mathcal{M} \text{ and } w(a^{-1}) = w(a). \quad (\text{M2})$$

We refer to this property as *rearrangement reversibility of the model*, or simply model reversibility. This condition is natural in the current group-based setting, where the typical rearrangements are reversals (which are self-inverse) and translocations (whose inverses are translocations). It is not essential for most of the construction, however does have some nice implications. For example, when we interpret our model of evolution as a Markov process, in Sect. 3.1, we will show that (M2) is equivalent to the time reversibility of the Markov process.

The evolutionary distance measure we consider in this paper is the maximum likelihood estimate of time elapsed (MLE). This is the maximum value of the likelihood function, which gives the probability, for any given time T , that the reference genome has evolved into the target genome in this amount of time. To be precise, for reference genome represented by the identity permutation $e \in \mathcal{S}_N$ and target genome represented by $\sigma \in \mathcal{S}_N$, the MLE is the maximum of the function $L(T|\sigma)$, where

$$L(T|\sigma) := P(\sigma|T) = \sum_{k=0}^{\infty} P(e \mapsto [\sigma] \text{ via } k \text{ events})P(k \text{ events in time } T). \quad (2)$$

Of course, the likelihood function need not have a maximum; this simply means that no evidence of an evolutionary relationship between the reference and the target under the given model can be discerned. This scenario, familiar from DNA sequence alignment paradigms such as the Jukes–Cantor correction (Felsenstein 2004), was discussed in the context of genome rearrangement models in Serdoz et al. (2017) and Terauds and Sumner (2019), where it was observed to occur in a substantial proportion of cases (independently of the chosen biological model).

For each k , the factor $P(k \text{ events in time } T)$ in the likelihood expression is determined by the distribution $dist$. The first factor, $P(e \mapsto [\sigma] \text{ via } k \text{ events})$ is the *genome path probability*, which we shall denote by $\alpha_k(\sigma)$. Since the target genome may be represented by any permutation from the set $[\sigma]$, ‘ $e \mapsto [\sigma]$ ’ is shorthand for “the permutation e is transformed into any permutation from the set $[\sigma]$ ” and we calculate the genome path probability as a sum of *permutation path probabilities*, denoted by $\beta_k(\sigma)$. That is,

$$\alpha_k(\sigma) := P(e \mapsto [\sigma] \text{ via } k \text{ events}) = \sum_{d \in \mathcal{D}_N} P(e \mapsto d\sigma \text{ via } k \text{ events}) = \sum_{d \in \mathcal{D}_N} \beta_k(d\sigma).$$

Given a permutation $\sigma \in \mathcal{S}_N$ and a model $(\mathcal{M}, w, dist)$, we specify each permutation path probability $\beta_k(\sigma)$ by considering the set $\mathcal{P}_k(\sigma)$ of all k -length sequences of

permutations, chosen from \mathcal{M} , that transform e into σ . Since we assume rearrangement events to be independent, the permutation path probability is then the sum of the probabilities of all such sequences, that is,

$$\beta_k(\sigma) = \sum_{(a_1, a_2, \dots, a_k) \in \mathcal{P}_k(\sigma)} w(a_1)w(a_2) \dots w(a_k).$$

We note that permutation path probabilities vary for different elements of $[\sigma]$ (that is, in general, $\beta_k(\sigma) \neq \beta_k(d\sigma)$ for $\sigma \in \mathcal{S}_N, d \in \mathcal{D}_N$). However, *genome* path probabilities are of course constant on the cosets $[\sigma]$. In fact, the model symmetry conditions (M1) and (M2) ensure that there are bigger classes of permutations that all have the same path probabilities (and thus likelihoods). The following results were established in Serdoz et al. (2017) ((i)) and Terauds and Sumner (2019) ((ii) and (iii)).

Theorem 2.1 *Let $(\mathcal{M}, w, dist)$ be a full biological model for evolution. For all $k \in \mathbb{N}_0$ and $\sigma \in \mathcal{S}_N$, the following hold.*

- (i) $\alpha_k(\sigma_1) = \alpha_k(\sigma_2)$ for all $\sigma_1, \sigma_2 \in [\sigma] = \{d\sigma : d \in \mathcal{D}_N\}$.
- (ii) *If the model has the dihedral symmetry property (M1), then $\alpha_k(\sigma_1) = \alpha_k(\sigma_2)$ for all σ_1, σ_2 in the set*

$$[\sigma]_D := \{d_1\sigma d_2 : d_1, d_2 \in \mathcal{D}_N\}.$$

- (iii) *If the model has the dihedral symmetry property (M1) and the reversibility property (M2), then $\alpha_k(\sigma_1) = \alpha_k(\sigma_2)$ for all σ_1, σ_2 in the set*

$$[\sigma]_{DR} := \{d_1\sigma d_2, d_1\sigma^{-1}d_2 : d_1, d_2 \in \mathcal{D}_N\}.$$

It was shown in Sumner et al. (2017) that the combinatorial problem of calculating path probabilities may be converted into a linear algebra problem via the representation theory of the symmetric group algebra, $\mathbb{C}[\mathcal{S}_N]$. For full details in the general model setting, we refer the reader to Terauds and Sumner (2019). We recall the essential steps in the derivation here, since we shall undertake a similar procedure in a lower dimensional algebra in the next section.

Here, we use the term *algebra* to mean a vector space equipped with a bilinear product. In particular, we require the group algebra $\mathbb{C}[\mathcal{S}_N]$ consisting of all formal linear combinations of elements of the group \mathcal{S}_N ; this algebra has natural basis \mathcal{S}_N , and thus dimension $N!$. For detailed background on the symmetric group algebra, and algebras more generally, we refer the reader to Sagan (2001) and Etingof et al. (2011) respectively. The following group algebra elements are key to our calculations.

Definition 2.2 *Let $(\mathcal{M}, w, dist)$ be a biological model for evolution of genomes with N regions. We define the *model element*, \mathbf{s} , and the *symmetry element*, \mathbf{z} , of the group algebra $\mathbb{C}[\mathcal{S}_N]$ by*

$$\mathbf{s} := \sum_{a \in \mathcal{M}} w(a) a \quad \text{and} \quad \mathbf{z} := \frac{1}{2N} \sum_{d \in \mathcal{D}_N} d.$$

To reformulate the path probabilities, we firstly observe that

$$\mathbf{s}^k = \sum_{\tau \in \mathcal{S}_N} \beta_k(\tau) \tau. \tag{3}$$

Then, for $\sigma \in \mathcal{S}_N$, we multiply (3) on the left by σ^{-1} to see that $\beta_k(\sigma)$ is the coefficient of e in the expansion of $\sigma^{-1}\mathbf{s}^k$. The representation theory of the symmetric group algebra tells us that this is exactly $(\frac{1}{N!}$ times) the trace of the regular representation of $\sigma^{-1}\mathbf{s}^k$. That is,

$$\beta_k(\sigma) = \frac{1}{N!} \chi_{\text{reg}}(\sigma^{-1}\mathbf{s}^k). \tag{4}$$

Thus, for $\sigma \in \mathcal{S}_N$, the k th genome path probability is

$$\begin{aligned} \alpha_k(\sigma) &= \sum_{d \in \mathcal{D}_N} \beta_k(d\sigma) = \frac{1}{N!} \sum_{d \in \mathcal{D}_N} \chi_{\text{reg}}(\sigma^{-1}d\mathbf{s}^k) = \frac{2N}{N!} \chi_{\text{reg}}(\sigma^{-1}\mathbf{z}\mathbf{s}^k) \\ &= \frac{2N}{N!} \sum_{p \vdash N} D_p \chi_p(\sigma^{-1}\mathbf{z}\mathbf{s}^k), \end{aligned} \tag{5}$$

where we have used the linearity of the characters to incorporate the symmetry element \mathbf{z} . The final equality is gained by decomposing the regular representation of $\mathbb{C}[\mathcal{S}_N]$ into irreducible representations. Recall that the irreducible representations of $\mathbb{C}[\mathcal{S}_N]$ correspond to the integer partitions of N (Sagan 2001, Prop. 1.10.1); here, we denote a partition of N by $p \vdash N$ and index the representations and related objects accordingly. Specifically, for each partition $p \vdash N$, ρ_p is the irreducible representation corresponding to p , D_p is its multiplicity (and dimension), and χ_p is the character of this representation.

The above derivation of the permutation path probabilities follows that of Sumner et al. (2017); in that paper it was also noted that an alternative derivation is possible via the theory of the Fourier transform on \mathcal{S}_N . That is, one may extend the probability distribution w on \mathcal{M} to w' on the whole of \mathcal{S}_N , notice that the Fourier transform of w' with respect to an irreducible representation ρ_p is equal to $\rho_p(\mathbf{s})$ and that w' convolved with itself k times is exactly the function β_k on \mathcal{S}_N , and then apply the Fourier inversion formula to obtain (4).

Now, for a model with rearrangement reversibility (M2), the irreducible representations of the model element \mathbf{s} are diagonalisable (Terauds and Sumner 2019) and we obtain

$$\alpha_k(\sigma) = \frac{2N}{N!} \sum_{p \vdash N} D_p \sum_{i=1}^{r_p} (\lambda_{p,i})^k \text{tr}(\rho_p(\sigma^{-1}\mathbf{z})E_{p,i}),$$

where for each p , the eigenvalues of $\rho_p(\mathbf{s})$ are $\{\lambda_{p,i} : i = 1, \dots, r_p\}$ and, for each p and i , $E_{p,i}$ is the projection onto the eigenspace of $\lambda_{p,i}$. Substituting this into the likelihood expression (2) and setting the distribution of events in time to be $\text{dist} = \text{Poisson}(1)$, we obtain

$$L(T|\sigma) = e^{-T} \frac{2N}{N!} \sum_{p \vdash N} D_p \sum_{i=1}^{r_p} \text{tr}(\rho_p(\sigma^{-1}\mathbf{z})E_{p,i})e^{\lambda_{p,i}T}, \tag{6}$$

where we have observed that the expression is in fact a power series and, accordingly, have been able to eliminate the infinite sum from the expression.

We note that, for a given model, one need only calculate the eigenvalues of each $\rho_p(\mathbf{s})$ once. Thus the bulk of the calculation burden is now in calculating the *partial traces*, that is, for any given genome, the set $\{\text{tr}(\rho_p(\sigma^{-1}\mathbf{z})E_{p,i}) : p \vdash N, i = 1 \dots r_p\}$ of coefficients that correspond to the distinct eigenvalues in the likelihood equation.

In implementing the likelihood calculations using the expression (6), we observed that for all genomes, *most* of these partial trace coefficients were zero (Terauds and Sumner 2019). That is, most of our calculations ended up not contributing to the final likelihood function. In the next section, we explain the occurrence of these zeroes and show that the redundancy can be removed from the computations.

3 The circular genome algebra

The calculations outlined above are performed in the group algebra $\mathbb{C}[S_N]$, where each permutation in S_N is a distinct basis element. However, we (and, indeed, the computations) do not distinguish between different permutations that represent the same genome—that is, between elements of each equivalence class

$$[\sigma] = \{d\sigma : d \in \mathcal{D}_N\},$$

for $\sigma \in S_N$. We now construct a lower-dimensional algebra by combining these equivalent permutations together to form basis elements—*permutation clouds*—that correspond to circular genomes. Until otherwise stated, assume that we have fixed a number of regions N and a biological model for evolution $(\mathcal{M}, w, \text{dist})$ that has dihedral symmetry and is reversible, that is, satisfies (M1) and (M2).

Definition 3.1 For symmetry element $\mathbf{z} \in \mathbb{C}[S_N]$, the *circular genome algebra* for N regions is

$$\mathcal{A} := \mathbf{z}\mathbb{C}[S_N] = \{\mathbf{z}\boldsymbol{\tau} : \boldsymbol{\tau} \in \mathbb{C}[S_N]\}.$$

Any element of \mathcal{A} of the form $\mathbf{z}\sigma$, where $\sigma \in S_N$, is called a *permutation cloud*.

One easily verifies that \mathcal{A} is a subalgebra of $\mathbb{C}[S_N]$. To see that \mathcal{A} has a natural basis that is in correspondence with the set of genomes, firstly observe that any element of \mathcal{A} can be written as a linear combination of permutation clouds $\mathbf{z}\sigma$, for $\sigma \in S_N$. Thus there exists a basis for \mathcal{A} of the form $\{\mathbf{z}\sigma_1, \dots, \mathbf{z}\sigma_K\}$, for $\sigma_i \in S_N$. Now,

$$\mathbf{z}\sigma_i = \frac{1}{2N} \sum_{d \in \mathcal{D}_N} d\sigma_i,$$

so that each basis element is a weighted sum of elements from a set $[\sigma_i]$, representing a particular genome. Since the sets are equivalence classes, for any $\sigma_1, \sigma_2 \in \mathcal{S}_N$ we have

$$\mathbf{z}\sigma_1 = \mathbf{z}\sigma_2 \iff \sigma_1, \sigma_2 \in [\sigma] \text{ for some } \sigma \in \mathcal{S}_N.$$

This means that the set of distinct permutation clouds corresponds to the set of distinct genomes, and these form a basis for \mathcal{A} . Finally, noting that for all $\sigma \in \mathcal{S}_N, |[\sigma]| = 2N$, we see that

$$\dim(\mathcal{A}) = |{[\sigma] : \sigma \in \mathcal{S}_N}| = \frac{N!}{2N} =: K. \tag{7}$$

For the remainder of this section, we fix a basis for \mathcal{A} ,

$$B := \{\mathbf{z}\sigma : \sigma \in \mathcal{S}_N\} = \{\mathbf{z}\sigma_1, \dots, \mathbf{z}\sigma_K\}, \tag{8}$$

where we have chosen a representative $\sigma_i \in \mathcal{S}_N$ of each equivalence class $[\sigma_i]$ for notational convenience. We set the first basis element to correspond to $[e] = \mathcal{D}_N$, so that $\mathbf{z}\sigma_1 = \mathbf{z}$. Having used the symmetry of the genomes to construct the algebra, we now incorporate the symmetry of the model to extract some useful properties.

Proposition 3.2 *The model and symmetry elements, $\mathbf{s}, \mathbf{z} \in \mathbb{C}[\mathcal{S}_N]$, have the following properties.*

- (i) \mathbf{z} is idempotent;
- (ii) \mathbf{s} and \mathbf{z} commute.

Proof (i) Since \mathcal{D}_N is a group, we have

$$\mathbf{z}^2 = \frac{1}{(2N)^2} \sum_{d \in \mathcal{D}_N} \sum_{f \in \mathcal{D}_N} d f = \frac{1}{(2N)^2} \sum_{d \in \mathcal{D}_N} \sum_{f \in \mathcal{D}_N} f = \frac{1}{2N} \sum_{f \in \mathcal{D}_N} f = \mathbf{z}.$$

(ii) Now we use the dihedral symmetry (M1) of the model to rewrite the model as *m* base rearrangements, $a_1, \dots, a_m \in \mathcal{S}_N$, along with their symmetries. That is,

$$\mathcal{M} = \{da_1d^{-1}, da_2d^{-1}, \dots, da_md^{-1} : d \in \mathcal{D}_N\}. \tag{9}$$

Then, using the same idea as in (i),

$$\mathbf{z}\mathbf{s} = \frac{1}{2N} \sum_{f \in \mathcal{D}_N} \sum_{i=1}^m \sum_{d \in \mathcal{D}_N} w(a_i) f da_i d^{-1} = \frac{1}{2N} \sum_{f \in \mathcal{D}_N} \sum_{i=1}^m \sum_{d \in \mathcal{D}_N} w(a_i) da_i d^{-1} f = \mathbf{s}\mathbf{z}.$$

□

The above properties translate immediately into properties of the representations of \mathbf{z} and \mathbf{s} .

Corollary 3.3 *Let $p \vdash N$ and $\rho_p : \mathbb{C}[\mathcal{S}_N] \rightarrow M_{D_p}(\mathbb{C})$ denote the corresponding irreducible representation of the symmetric group algebra. Then*

- (i) *the only eigenvalues of $\rho_p(\mathbf{z})$ are 0 and 1;*
- (ii) *$\rho_p(\mathbf{z})$ and $\rho_p(\mathbf{s})$ are simultaneously diagonalisable, with real eigenvectors.*

Proof Claim (i) is immediate since \mathbf{z} , and thus $\rho_p(\mathbf{z})$, is idempotent. To show (ii), we firstly choose the representation ρ_p to be orthogonal on \mathcal{S}_N (Sagan 2001). Then the rearrangement reversibility of the model ensures that $\rho_p(\mathbf{s})$ is symmetric (Terauds and Sumner 2019) and, similarly, one may verify directly that $\rho_p(\mathbf{z})^\top = \rho_p(\mathbf{z})$. Thus, since \mathbf{z} and \mathbf{s} commute, the representation matrices commute and are simultaneously diagonalisable. In particular, since these matrices are real symmetric, the orthonormal set of simultaneous eigenvectors may be chosen to be real. \square

Now fix $p \vdash N$ and choose a set of orthonormal vectors $\{v_1, v_2, \dots, v_{D_p}\} \subseteq \mathbb{R}^{D_p}$ that are eigenvectors for both $\rho_p(\mathbf{z})$ and $\rho_p(\mathbf{s})$, ordered so that the first k_p of them are eigenvectors for the eigenvalue 1 of $\rho_p(\mathbf{z})$. Take an eigenvalue, $\lambda_{p,i}$ of $\rho_p(\mathbf{s})$ and let $J_i \subseteq \{1, 2, \dots, D_p\}$ such that $\{v_j : j \in J_i\}$ are the eigenvectors for $\lambda_{p,i}$. Then, for $\sigma \in \mathcal{S}_N$, the partial trace for $\lambda_{p,i}$ may be written as

$$\text{tr}(\rho_p(\sigma^{-1}\mathbf{z})E_{p,i}) = \sum_{j \in J_i} v_j^T \rho_p(\sigma^{-1}\mathbf{z})v_j, \tag{10}$$

where for each j ,

$$v_j^T \rho_p(\sigma^{-1}\mathbf{z})v_j = v_j^T \rho_p(\sigma^{-1})\rho_p(\mathbf{z})v_j = \begin{cases} v_j^T \rho_p(\sigma^{-1})v_j, & \text{if } j \leq k_p; \\ 0, & \text{if } j > k_p. \end{cases} \tag{11}$$

We see then that $\rho_p(\mathbf{z})$ ‘‘knocks out’’ parts of the partial traces; in particular, it does this independently of the genome. We shall establish shortly that, in total, $\frac{2N-1}{2N}$ of the partial traces are knocked out in this way, thus explaining the observation in Terauds and Sumner (2019) that most of the calculated partial traces were zero.

The key to performing MLE computations in the symmetric group algebra is the relationship between the character of the regular representation and the identity element, $e \in \mathbb{C}[\mathcal{S}_N]$: the character $\chi_{\text{reg}}(\boldsymbol{\tau})$ counts occurrences of the identity in a generic element $\boldsymbol{\tau} \in \mathbb{C}[\mathcal{S}_N]$. Since \mathbf{z} is idempotent, it is a left identity (but not a right identity) in the algebra \mathcal{A} . We now construct the regular representation, $\rho_{\text{reg}}^{\mathcal{A}}$, of \mathcal{A} and show that its character, the *regular character* $\chi_{\text{reg}}^{\mathcal{A}}$, functions in exactly this way for the left identity $\mathbf{z} \in \mathcal{A}$.

We construct the regular representation of \mathcal{A} via the left action of elements of \mathcal{A} on the basis $B = \{\mathbf{z}\sigma_1, \dots, \mathbf{z}\sigma_K\}$ fixed above (8). We need only consider the representation of a generic basis element, $\mathbf{z}\sigma$ for $\sigma \in \mathcal{S}_N$, since one may extend linearly to all of \mathcal{A} . For arbitrary $\sigma \in \mathcal{S}_N$, the ij th entry of the matrix $\rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\sigma)$ is the coefficient of $\mathbf{z}\sigma_i$ in the expansion of $(\mathbf{z}\sigma)(\mathbf{z}\sigma_j)$, that is,

$$(\rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\sigma))_{ij} = \frac{1}{2N} |\{d \in \mathcal{D}_N : \sigma d \sigma_j \in [\sigma_i]\}|. \tag{12}$$

One readily verifies that $\rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z})$ is the $K \times K$ identity matrix and that $\rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\sigma)^T = \rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\sigma^{-1})$. The regular character $\chi_{\text{reg}}^{\mathcal{A}}$ is the trace of the regular representation matrix. For a generic basis element $\mathbf{z}\sigma \in \mathcal{A}$,

$$\begin{aligned} \chi_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\sigma) &= \sum_{i=1}^K (\rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\sigma))_{ii} = \frac{1}{2N} \sum_{i=1}^K |\{d_1 \in \mathcal{D}_N : \sigma d_1 \sigma_i \in [\sigma_i]\}| \\ &= \frac{1}{2N} \sum_{i=1}^K \sum_{d_2 \in \mathcal{D}_N} |\{d_1 \in \mathcal{D}_N : \sigma d_1 \sigma_i = d_2 \sigma_i\}| \\ &= \frac{1}{2N} \sum_{i=1}^K \sum_{d_2 \in \mathcal{D}_N} |\{d_1 \in \mathcal{D}_N : \sigma d_1 = d_2\}| \\ &= \frac{1}{2N} K \sum_{d_2 \in \mathcal{D}_N} |\{d_1 \in \mathcal{D}_N : d_1 = \sigma^{-1} d_2\}| \\ &= \begin{cases} K & \text{if } \sigma \in \mathcal{D}_N; \\ 0 & \text{if } \sigma \notin \mathcal{D}_N. \end{cases} \end{aligned} \tag{13}$$

Since $\mathbf{z}\sigma = \mathbf{z}$ if and only if $\sigma \in \mathcal{D}_N$, this shows that we can use the character of the regular representation of the algebra \mathcal{A} to track coefficients of the left identity \mathbf{z} , just as we do for the identity e in $\mathbb{C}[S_N]$. Further, we can express the regular character of \mathcal{A} as a sum over the irreducible characters of $\mathbb{C}[S_N]$, and thus see that the regular characters of \mathcal{A} and $\mathbb{C}[S_N]$ coincide on \mathcal{A} .

Proposition 3.4 For arbitrary $\tau \in \mathcal{A}$,

- (i) $\frac{1}{K} \chi_{\text{reg}}^{\mathcal{A}}(\tau)$ is the coefficient of \mathbf{z} in τ ;
- (ii) $\chi_{\text{reg}}^{\mathcal{A}}(\tau) = \sum_{p \vdash N} \chi_p(e) \chi_p(\tau) = \sum_{p \vdash N} D_p \chi_p(\tau) = \chi_{\text{reg}}(\tau)$.

Proof (i) Given $\tau \in \mathcal{A}$ and the basis B from (8), there exist $c_1, \dots, c_K \in \mathbb{C}$ such that $\tau = c_1 \mathbf{z} + c_2 \mathbf{z}\sigma_2 + \dots + c_K \mathbf{z}\sigma_K$. Then

$$\chi_{\text{reg}}^{\mathcal{A}}(\tau) = c_1 \chi_{\text{reg}}^{\mathcal{A}}(\mathbf{z}) + \sum_{i=2}^K c_i \chi_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\sigma_i) = c_1 K,$$

from (13).

(ii) It suffices to consider a generic basis element $\mathbf{z}\sigma \in \mathcal{A}$, since the characters are linear. We shall apply the dual orthogonality relations on the irreducible characters of S_N (see for example (James and Liebeck 2001, Thm. 16.4)), given for $\sigma, \tau \in S_N$ by

$$\sum_{p \vdash N} \chi_p(\sigma) \overline{\chi_p(\tau)} = \delta(\sigma, \tau) |\text{cent}_{S_N}(\sigma)|, \tag{14}$$

where $\text{cent}_{\mathcal{S}_N}(\sigma) := \{\gamma \in \mathcal{S}_N : \gamma\sigma = \sigma\gamma\}$ is the centraliser of σ and the map $\delta : \mathcal{S}_N \rightarrow \{0, 1\}$ is defined by

$$\delta(\sigma, \tau) = \begin{cases} 1, & \text{if } \sigma, \tau \text{ are in the same conjugacy class in } \mathcal{S}_N; \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

Recall that for $p \nmid N$ and $\tau \in \mathcal{S}_N$, $\overline{\chi_p(\tau)} = \chi_p(\tau^{-1})$.³ Then for $\sigma \in \mathcal{S}_N$, we have

$$\begin{aligned} \sum_{p \nmid N} \chi_p(e)\chi_p(\mathbf{z}\sigma) &= \frac{1}{2^N} \sum_{d \in \mathcal{D}_N} \sum_{p \nmid N} \chi_p(e)\overline{\chi_p(\sigma^{-1}d)} \\ &= \frac{1}{2^N} \sum_{d \in \mathcal{D}_N} \delta(e, \sigma^{-1}d) |\text{cent}_{\mathcal{S}_N}(e)| \\ &= \begin{cases} \frac{N!}{2^N} & \text{if } \sigma \in \mathcal{D}_N; \\ 0 & \text{if } \sigma \notin \mathcal{D}_N. \end{cases} \\ &= \chi_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\sigma), \end{aligned}$$

by (13), recalling that $K = \frac{N!}{2^N}$. □

An immediate consequence of the above is an expression for the dimension of \mathcal{A} in terms of the characters χ_p of $\mathbb{C}[\mathcal{S}_N]$:

$$\dim(\mathcal{A}) = K = \frac{N!}{2^N} = \sum_{p \nmid N} \chi_p(e)\chi_p(\mathbf{z}) = \sum_{p \nmid N} D_p k_p, \tag{16}$$

where, for each $p \nmid N$, k_p is the multiplicity of the eigenvalue 1 of $\rho_p(\mathbf{z})$. The first part of this can also be seen directly from the dual orthogonality relations.

To perform the MLE calculations in the algebra \mathcal{A} efficiently, we will need a decomposition of the regular character in terms of irreducible characters of \mathcal{A} . Firstly, we'll observe that irreducible submodules of $\mathcal{A} = \mathbf{z}\mathbb{C}[\mathcal{S}_N]$ can be produced by acting with \mathbf{z} on the irreducible submodules of $\mathbb{C}[\mathcal{S}_N]$. This is straightforward and, in fact, true in a more general context (see, for example, Steinberg 2016, Lemma 4.15), however we include the details here since we'll use them in our subsequent constructions. We denote the irreducible submodules of $\mathbb{C}[\mathcal{S}_N]$ by $V_p = \mathbb{C}^{D_p}$ for $p \nmid N$, that is, we write

$$\mathbb{C}[\mathcal{S}_N] \cong \bigoplus_{p \nmid N} D_p V_p. \tag{17}$$

Theorem 3.5 *The non-trivial modules gained by acting with the symmetry element \mathbf{z} on the irreducible submodules of $\mathbb{C}[\mathcal{S}_N]$ are irreducible modules of the genome algebra \mathcal{A} .*

³ Recall that the irreducible characters of \mathcal{S}_N are all real valued, however we use the general form of the result here, since we apply the same argument to an arbitrary group in Sect. 4.

Proof Let $p \dashv N$. As above, we may take a set $\{v_1, v_2, \dots, v_{D_p}\}$ of (real, orthonormal, linearly independent) eigenvectors for both $\rho_p(\mathbf{s})$ and $\rho_p(\mathbf{z})$, ordered such that the first k_p of them correspond to the eigenvalue 1 of $\rho_p(\mathbf{z})$. Then $V_p = \text{span}_{\mathbb{C}}\{v_i : i = 1, \dots, D_p\}$ and

$$W_p := \mathbf{z} \cdot V_p = \text{span}_{\mathbb{C}}\{\rho(\mathbf{z})v_i : i = 1, \dots, D_p\} = \text{span}_{\mathbb{C}}\{v_i : i = 1, \dots, k_p\}. \tag{18}$$

It is clear that W_p is an \mathcal{A} -module; we need show that it is either $\{0\}$ or irreducible. Suppose that there exists $U_p \subseteq W_p$ such that U_p is an \mathcal{A} -module, and $0 \neq u \in U_p$. Then

$$U_p \supseteq \text{span}_{\mathbb{C}}\{\rho_p(\mathbf{z}\sigma)u : \sigma \in \mathcal{S}_N\} = \mathbf{z} \cdot \text{span}_{\mathbb{C}}\{\rho_p(\sigma)u : \sigma \in \mathcal{S}_N\} = \mathbf{z} \cdot V_p = W_p,$$

so that $U_p = W_p$. □

It was shown in Terauds and Sumner (2019) that there exist $p \dashv N$ for all $N > 3$ such that $\chi_p(\mathbf{z}) = k_p = 0$; that is, there are always some $\mathbb{C}[\mathcal{S}_N]$ -modules that are projected down to zero in \mathcal{A} . We note that this is not true in the more general case considered in Sect. 4 (for example if \mathbf{z} is constructed from a different symmetry group).

Now, the dimension expression (16) suggests that we will not be able to decompose the algebra \mathcal{A} into a direct sum of irreducible submodules as we can for $\mathbb{C}[\mathcal{S}_N]$ (17). By (Etingof et al. 2011, Thm. 3.5.8), if this were possible with the irreducible submodules W_p from above, then the dimension of \mathcal{A} would be $\sum_{p \dashv N} k_p^2$. We have not yet verified here that these W_p comprise *all* irreducible submodules of \mathcal{A} , nor that they are all distinct (not isomorphic to one another), but this is indeed the case (Steinberg 2016, Thm. 4.23). The difference between the dimension of \mathcal{A} and that gained from the irreducible modules here is signalling that not all of the information about \mathcal{A} can be represented by the action of \mathcal{A} —in this case, the left action of \mathcal{A} on its irreducible modules, and on itself, is not injective.

To see this, let W be an irreducible module of \mathcal{A} . Then, since \mathbf{z} is a left identity in \mathcal{A} , \mathbf{z} must act as the identity on W . But then, for any $\mathbf{z}\sigma, \mathbf{z}\sigma' \in \mathcal{A}$ such that $\mathbf{z}\sigma\mathbf{z} = \mathbf{z}\sigma'\mathbf{z}$,

$$(\mathbf{z}\sigma) \cdot w = (\mathbf{z}\sigma) \cdot (\mathbf{z} \cdot w) = (\mathbf{z}\sigma\mathbf{z}) \cdot w = (\mathbf{z}\sigma'\mathbf{z}) \cdot w = (\mathbf{z}\sigma') \cdot w, \tag{19}$$

for all $w \in W$.

From Theorem 2.1(ii), such $\mathbf{z}\sigma \neq \mathbf{z}\sigma' \in \mathcal{A}$ correspond to physically distinct genomes that share the same path probabilities and likelihood functions: we have $\mathbf{z}\sigma\mathbf{z} = \mathbf{z}\sigma'\mathbf{z}$ if and only if $\sigma' \in [\sigma]_D = \{d\sigma d' : d, d' \in \mathcal{D}_N\}$.

In the language of algebras, \mathcal{A} has a *non-trivial radical* (Etingof et al. 2011, Def. 3.5.1), since (for $N > 3$) there are non-zero elements $\mathbf{z}\sigma - \mathbf{z}\sigma' \in \mathcal{A}$ that annihilate all irreducible modules of \mathcal{A} . As a concrete example, consider the following.

Example 3.6 Let $\sigma = (1, 2), \sigma' = (2, 3) \in \mathcal{S}_N$. Setting $r = (1, 2, \dots, N) \in \mathcal{D}_N$, we observe that $\sigma' = r\sigma r^{-1}$, so that $\mathbf{z}\sigma\mathbf{z} = \mathbf{z}\sigma'\mathbf{z}$. But there exists no $d \in \mathcal{D}_N$ such that $\sigma = d\sigma'$, and thus $\mathbf{z}\sigma \neq \mathbf{z}\sigma'$. ◇

For our practical purposes, this is perfect: the algebra sees genomes as distinct entities, but their representations do not distinguish between genomes corresponding to an equivalence class $[\sigma]_D$, whose likelihood functions are the same. Further, whilst $\mathbf{z}\sigma$ and $\mathbf{z}\sigma'$ correspond to distinct genomes, if we consider them as rearrangements, they are not distinct, since they have the same action. We shall return to this presently, when we define models in the genome algebra.

Proposition 3.7 *Let $\mathbf{z}\sigma, \mathbf{z}\sigma' \in \mathcal{A}$. Then $\rho_{reg}^A(\mathbf{z}\sigma) = \rho_{reg}^A(\mathbf{z}\sigma')$ if and only if $\sigma' \in [\sigma]_D$.*

Proof For the reverse implication, we argue as above, replacing w in (19) by each basis element $\mathbf{z}\sigma_i$ to verify that the matrices are the same. Conversely, if the regular representations coincide, then we immediately have that $(\mathbf{z}\sigma)\mathbf{z} = (\mathbf{z}\sigma')\mathbf{z}$. \square

We now explicitly consider the irreducible representations of \mathcal{A} on the irreducible submodules and use these to rewrite the regular character of \mathcal{A} in terms of irreducible characters of \mathcal{A} . Let $p \dashv N$ such that $k_p > 0$ and consider the module W_p of \mathcal{A} which, as in the proof of Theorem 3.5, has a basis $\{v_1, \dots, v_{k_p}\} \subseteq \mathbb{R}^d$ of orthonormal eigenvectors. Now, the action of \mathcal{A} on $W_p = \mathbf{z} \cdot V_p$ is inherited from the action of $\mathbb{C}[\mathcal{S}_N]$ on V_p , so for arbitrary $\tau \in \mathcal{A}$, we define the $(k_p \times k_p)$ representation matrix $\rho_p^A(\tau)$ on W_p via the action of $\rho_p(\tau)$ on the basis vectors v_j :

$$\left(\rho_p^A(\tau)\right)_{ij} := v_i^T \rho_p(\tau) v_j.$$

More concisely, setting Q_p to be the $(D_p \times k_p)$ matrix with $\{v_1, \dots, v_{k_p}\}$ as columns, we have

$$\rho_p^A(\tau) = Q_p^T \rho_p(\tau) Q_p. \tag{20}$$

Clearly (also c.f. (19)), $\rho_p^A(\mathbf{z})$ is the $(k_p \times k_p)$ identity matrix for each such $p \dashv N$. For each $p \dashv N$ such that $k_p = 0$, we formally define ρ_p^A to be the zero representation. Now we may calculate the irreducible characters χ_p^A of \mathcal{A} and see that they coincide with the irreducible characters χ_p of $\mathbb{C}[\mathcal{S}_N]$ restricted to \mathcal{A} .

Proposition 3.8 *For each $p \vdash N$ and $\tau \in \mathcal{A}$, $\chi_p^A(\tau) = \chi_p(\tau)$.*

Proof Let $p \dashv N$. By linearity, we need only verify the claim on a generic basis element, $\mathbf{z}\sigma \in \mathcal{A}$. Again utilising the orthonormal eigenvectors $\{v_1, \dots, v_{D_p}\}$ of $\rho_p(\mathbf{z})$, where those for $i \leq k_p$ correspond to the eigenvalue 1 and the remainder to the eigenvalue 0, we have

$$\begin{aligned} \chi_p^A(\mathbf{z}\sigma) &= \sum_{i=1}^{k_p} \left(\rho_p^A(\mathbf{z}\sigma)\right)_{ii} = \sum_{i=1}^{k_p} v_i^T \rho_p(\mathbf{z}\sigma) v_i = \sum_{i=1}^{D_p} v_i^T \rho_p(\mathbf{z}\sigma\mathbf{z}) v_i \\ &= \chi_p(\mathbf{z}\sigma\mathbf{z}) = \chi_p(\mathbf{z}\sigma), \end{aligned}$$

where in the final step we have used the cyclicity of the trace and the idempotency of \mathbf{z} (Proposition 3.2). \square

Combining Propositions 3.4 and 3.8 gives the desired character decomposition.

Corollary 3.9 For arbitrary $\tau \in \mathcal{A}$,

$$\chi_{reg}^{\mathcal{A}}(\tau) = \sum_{p \vdash N} D_p \chi_p^{\mathcal{A}}(\tau).$$

□

Having defined and decomposed the regular character of \mathcal{A} , we are ready to return to the likelihood calculations. Using the equivalence of the characters of \mathcal{A} and $\mathbb{C}[S_N]$ on the algebra \mathcal{A} , along with the the interplay between the genome and model symmetry, we now verify that we may work entirely in \mathcal{A} to calculate the genome path probabilities and thus the likelihood functions, as defined in the previous section (2).

Theorem 3.10 Let $\sigma \in S_N$ and $k \in \mathbb{N}_0$. Then

$$\alpha_k(\sigma) = \frac{2N}{N!} \chi_{reg}^{\mathcal{A}}(\mathbf{z}\sigma^{-1}\mathbf{z}\mathbf{s}^k) = \frac{2N}{N!} \sum_{p \vdash N} D_p \chi_p^{\mathcal{A}}(\mathbf{z}\sigma^{-1}\mathbf{z}\mathbf{s}^k). \tag{49}$$

Proof From (5), $\alpha_k(\sigma) = \frac{2N}{N!} \chi_{reg}(\sigma^{-1}\mathbf{z}\mathbf{s}^k) = \frac{2N}{N!} \chi_{reg}(\mathbf{z}\sigma^{-1}\mathbf{z}\mathbf{s}^k)$, since \mathbf{z} and \mathbf{s} commute, \mathbf{z} is idempotent and the trace is cyclic. The first equality is then clear from Proposition 3.4 and the second from Corollary 3.9. □

We have mentioned the importance of the ‘identity counting’ property of the regular character, that is, Proposition 3.4 (i), but this combinatorial component is somewhat hidden in the proof of Theorem 3.10. To highlight it, one may begin with the identity (3) stated in the previous section and, for any given genome $\mathbf{z}\sigma$ ($\sigma \in S_N$), multiply by $\mathbf{z}\sigma^{-1}\mathbf{z}$ to obtain

$$\mathbf{z}\sigma^{-1}\mathbf{z}\mathbf{s}^k = \frac{1}{2N} \sum_{\tau \in S_N} \alpha_k(\tau) \mathbf{z}\sigma^{-1}\tau.$$

By observing that there are exactly $2N$ values of $\tau \in S_N$ for which $\mathbf{z}\sigma^{-1}\tau = \mathbf{z}$, one thus sees directly that the coefficient of \mathbf{z} in the expansion of $\mathbf{z}\sigma^{-1}\mathbf{z}\mathbf{s}^k$ is $\alpha_k(\sigma)$.

Note that we could have simplified the above character expression (21) a little, that is,

$$\chi_p^{\mathcal{A}}(\mathbf{z}\sigma^{-1}\mathbf{z}\mathbf{s}^k) = \chi_p^{\mathcal{A}}(\mathbf{z}\sigma^{-1}\mathbf{s}^k\mathbf{z}) = \chi_p^{\mathcal{A}}(\mathbf{z}\sigma^{-1}\mathbf{s}^k).$$

However, as we did in the algebra $\mathbb{C}[S_N]$, we want to diagonalise the matrices representing the model element, namely the matrices $\rho_p^{\mathcal{A}}(\mathbf{z}\mathbf{s})$. So we keep the middle \mathbf{z} and write

$$\chi_p^{\mathcal{A}}(\mathbf{z}\sigma^{-1}\mathbf{z}\mathbf{s}^k) = \text{tr} \left(\rho_p^{\mathcal{A}}(\mathbf{z}\sigma^{-1}(\mathbf{z}\mathbf{s})^k) \right) = \text{tr} \left(\rho_p^{\mathcal{A}}(\mathbf{z}\sigma^{-1}) \rho_p^{\mathcal{A}}(\mathbf{z}\mathbf{s})^k \right).$$

For each $p \vdash N$, as in the proof of Corollary 3.3, we can choose $\rho_p(\mathbf{s})$ to be symmetric; thus by the definition (20) each matrix $\rho_p^{\mathcal{A}}(\mathbf{z}\mathbf{s})$ is symmetric and thus diagonalisable.

Then we obtain

$$\alpha_k(\sigma) = \frac{2N}{N!} \sum_{p \vdash N} D_p \sum_{i=1}^{R_p} \lambda_{p,i}^k \text{tr}(\rho_p^A(\mathbf{z}\sigma^{-1})E_{p,i}^A), \tag{22}$$

where $E_{p,i}^A$ is the projection onto the eigenspace of the i th eigenvalue, $\lambda_{p,i}$, of $\rho_p^A(\mathbf{z}\mathbf{s})$.

Now, finally substituting the path probabilities (22) into the theoretical likelihood expression (2), we obtain

$$L(T|\sigma) = e^{-T \frac{2N}{N!} \sum_{p \vdash N} D_p \sum_{i=1}^{R_p} \text{tr}(\rho_p^A(\mathbf{z}\sigma^{-1})E_{p,i}^A) e^{\lambda_{p,i} T}}. \tag{23}$$

It is clear from Theorem 3.10 that the likelihood expression (23), involving only elements of the genome algebra \mathcal{A} , is equal to that (6) gained via the group algebra $\mathbb{C}[\mathcal{S}_N]$. We now show that the above is, really, a simplified version of (6): that is, by working in the smaller algebra we have eliminated the many eigenvalue terms that occur with zero coefficients.

Proposition 3.11 *For each $p \vdash N$ such that $W_p \neq \{0\}$, the eigenvalues of the matrix $\rho_p^A(\mathbf{z}\mathbf{s})$ are exactly the eigenvalues of $\rho_p(\mathbf{s})$ that occur with non-zero coefficient in the likelihood expression (6).*

Proof Let $p \vdash N$ such that $W_p \neq \{0\}$. As above, take the set $\{v_1, v_2, \dots, v_{D_p}\} \subseteq \mathbb{R}^{D_p}$ of orthonormal eigenvectors for both $\rho_p(\mathbf{z})$ and $\rho_p(\mathbf{s})$, with the first k_p corresponding to the eigenvalue 1 of $\rho_p(\mathbf{z})$, and form the matrix Q_p with the first k_p vectors as columns. Then, as in (20),

$$\rho_p^A(\mathbf{z}\mathbf{s}) = Q_p^T \rho_p(\mathbf{z}\mathbf{s}) Q_p = Q_p^T \rho_p(\mathbf{s}\mathbf{z}) Q_p = Q_p^T \rho_p(\mathbf{s}) Q_p = \begin{pmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ & & \ddots & \\ 0 & 0 & \dots & \lambda_{k_p} \end{pmatrix}, \tag{24}$$

where each λ_i is clearly an eigenvalue of both $\rho_p^A(\mathbf{z}\mathbf{s})$ and $\rho_p(\mathbf{s})$ (and the λ_i are not necessarily distinct). Suppose λ' is an eigenvalue of $\rho_p(\mathbf{s})$ that does not appear in the matrix of (24). Then λ' has corresponding eigenvector(s) $\{v_j : j \in J'\}$, where $J' \subseteq \{k_{p+1}, k_{p+2}, \dots, D_p\}$. But then, for any $\sigma \in \mathcal{S}_N$, the coefficient of the λ' term in the likelihood expression is the partial trace

$$\sum_{j \in J'} v_j^T \rho_p(\sigma^{-1}\mathbf{z}) v_j = 0,$$

by (10) and (11). □

Note that, although in the proof of Proposition 3.11 we construct each $\rho_p^A(\mathbf{z}\mathbf{s})$ as a diagonal matrix (in which case the projections onto the eigenspaces would be diagonal matrices of 1s and 0s), we do this only to verify that the representation has the required properties, and we utilise the eigenvectors of the representation $\rho_p(\mathbf{s})$. In practice, the whole point is to *not* calculate the much bigger representations $\rho_p(\mathbf{s})$. That is, when implementing calculations, we would expect to construct a basis for each irreducible module W_p directly, hence the general form of the projections in (23).

We note that the equivalence of the path probabilities and thus likelihoods on the classes $[\sigma]_D$ and $[\sigma]_{DR}$ stated in Theorem 2.1 can alternatively be obtained by working directly in the genome algebra \mathcal{A} . We omit the proof here since we shall prove a more general version of the result in Sect. 4.

Since $\dim(\mathbb{C}[\mathcal{S}_N]) = N! = \sum_{p \vdash N} D_p^2$ and $\dim(\mathcal{A}) = \frac{N!}{2N} = \sum_{p \vdash N} D_p k_p$, we have

$$\sum_{p \vdash N} D_p k_p = \sum_{p \vdash N} D_p \frac{D_p}{2N}. \tag{49}$$

Note that this does not imply that $k_p = \frac{1}{2N} D_p$ for each (or any) $p \vdash N$, rather that *on average*, and asymptotically, the dimension of each irreducible submodule W_p of \mathcal{A} is $\frac{1}{2N}$ th of the dimension of the irreducible submodule V_p of $\mathbb{C}[\mathcal{S}_N]$. To put this another way, on average, $\frac{2N-1}{2N}$ ths of the computations in the group algebra, as documented in Terauds and Sumner (2019), resulted in zeroes.

Given that the dimension of the algebra \mathcal{A} is still $\frac{N!}{2N}$, this does not significantly reduce the computational complexity. However, since the multiplicity of the irreducible submodules in \mathcal{A} is the same as in the group algebra (25), the reduction in the dimension of the irreducible submodules is (relatively) much larger than the reduction in total dimension.

Example 3.12 Consider $N = 6$. There are $N! = 720$ permutations in S_6 , so the dimension of the regular representation of $\mathbb{C}[S_6]$ is 720. The dimensions of the irreducible modules V_p of $\mathbb{C}[S_6]$ (given as a list rather than a set as they are not all distinct) are

$$[D_p : p \vdash 6] = [1, 5, 9, 10, 5, 16, 10, 5, 9, 5, 1].$$

Moving to the genome algebra, there are $\frac{N!}{2N} = 60$ distinct genomes, so the dimension of the regular representation of \mathcal{A} is 60. The dimensions of the corresponding irreducible modules $W_p = \mathbf{z} \cdot V_p$ of \mathcal{A} are

$$[k_p : p \vdash 6] = [1, 0, 2, 0, 0, 1, 1, 2, 0, 1, 0].$$

Thus, for any rearrangement model, each likelihood expression will be a sum of at most eight terms, corresponding to at most eight distinct eigenvalues. \diamond

We note that such dimension reductions are less striking for larger N . In any case, we see a significant theoretical gain here: the genome algebra \mathcal{A} incorporates the symmetry of the genomes and models into a unified framework, within which the

problem can be formulated and the computations performed. To highlight this, we next consider the regular representations of \mathbf{s} in the algebra $\mathbb{C}[\mathcal{S}_N]$ and of \mathbf{zs} in the algebra \mathcal{A} as Markov matrices; then we conclude this section by re-formulating the model in the genome algebra framework.

3.1 The Markov interpretation

In the group algebra, $\mathbb{C}[\mathcal{S}_N]$, the rows and columns of the regular representation are determined by the $N!$ permutations $\sigma_i \in \mathcal{S}_N$. In particular,

$$\rho_{\text{reg}}(\mathbf{s}) = \sum_{a \in \mathcal{M}} w(a) \rho_{\text{reg}}(a),$$

where for each rearrangement permutation $a \in \mathcal{M}$, the ij th entry of $\rho_{\text{reg}}(a)$ is 1 if $a\sigma_j = \sigma_i$ and 0 otherwise, so that the $\rho_{\text{reg}}(a)$ matrices have exactly one ‘1’ in each row and column. Then $\rho_{\text{reg}}(\mathbf{s})$, as a convex sum of Markov matrices, is itself a Markov matrix. The j th column of $\rho_{\text{reg}}(\mathbf{s})$ contains $|\mathcal{M}|$ non-zero entries, each equal to a unique $w(a)$, since for the distinct permutations $a \in \mathcal{M}$, the permutations $a\sigma_j$ are all distinct.

Thus $\rho_{\text{reg}}(\mathbf{s})$ is the transition matrix of a discrete Markov chain where the states are the $N!$ permutations in \mathcal{S}_N and the ij th entry is the probability of permutation σ_j transitioning into permutation σ_i via one rearrangement chosen from the model \mathcal{M} . That is,

$$\rho_{\text{reg}}(\mathbf{s})_{ij} = \begin{cases} w(a) & \text{if } a = \sigma_i \sigma_j^{-1} \in \mathcal{M}, \\ 0 & \text{otherwise.} \end{cases}$$

It is clear from this formulation that the matrix $\rho_{\text{reg}}(\mathbf{s})$ is symmetric if and only if the model has the rearrangement reversibility property (M2). Thus, since the stationary distribution on the Markov chain is the uniform distribution on the states, the reversibility property (M2) of the model is equivalent to reversibility of the Markov model.

Now, in the algebra \mathcal{A} , the corresponding matrix representing the model element is

$$\rho_{\text{reg}}^A(\mathbf{zs}) = \sum_{a \in \mathcal{M}} w(a) \rho_{\text{reg}}^A(\mathbf{za}). \tag{26}$$

As above, the matrices on the right hand side represent basis elements of the algebra, here \mathbf{za} for $a \in \mathcal{M}$. Although the basis elements here do not form a group (so their regular representations are not, in general, zero-one matrices), each of the $\rho_{\text{reg}}^A(\mathbf{za})$ is again a Markov matrix: for a given $a \in \mathcal{M}$, the ij th entry of $\rho_{\text{reg}}^A(\mathbf{za})$ is the coefficient of $\mathbf{z}\sigma_i$ in the expansion of $(\mathbf{za})(\mathbf{z}\sigma_j)$ and, since $\mathbf{z} = \frac{1}{2N} \sum_{d \in \mathcal{D}_N} d$, the expansion is a convex sum. Thus the entries in each column of each $\rho_{\text{reg}}^A(\mathbf{za})$ sum to one and $\rho_{\text{reg}}^A(\mathbf{zs})$, as a convex sum of Markov matrices, is indeed a Markov matrix.

Each basis element $\mathbf{z}\sigma_i$ corresponds to a genome, so $\rho_{\text{reg}}^A(\mathbf{z}\mathbf{s})$ is the transition matrix of a Markov chain where the states are genomes. The ij th entry, which we calculated as the proportion of the expansion of $(\mathbf{z}\mathbf{s})(\mathbf{z}\sigma_j)$ that is equal to $\mathbf{z}\sigma_i$, is of course the probability of the genome $(\mathbf{z}\sigma_j)$ transitioning into the genome $\mathbf{z}\sigma_i$ in one step, via the model.

3.2 Permutation clouds: a unifying concept

In the permutation approach detailed in Sect. 2, we considered rearrangement events to be individual permutations acting on individual permutations. In the setting of the genome algebra, we represent both genomes and rearrangements by permutation clouds, each of which is a sum of permutations weighted by their probabilities ($\mathbf{z}\sigma = \frac{1}{2N} \sum_{d \in \mathcal{D}_N} d\sigma$). A single rearrangement event is here modelled by a permutation cloud $\mathbf{z}a$ acting on a permutation cloud $\mathbf{z}\sigma$. Mathematically, this event results in a convex combination of permutation clouds $\sum c_i \mathbf{z}\sigma_i$; biologically, it results in one of the genomes $\mathbf{z}\sigma_i$, according to the probability distribution given by the coefficients c_i .

The permutation cloud view of circular genomes seems to us quite natural. To observe a genome, we fix an orientation and a reference frame, and assign to it a single permutation (any one, from the appropriate equivalence class $[\sigma]$, with probability $\frac{1}{2N}$). We refer to this as an *instance* of the genome. Theoretically, however, the genome exists simultaneously as all of its possible physical orientations in space; it *is* the cloud, $\mathbf{z}\sigma$.

What about rearrangements? For a rearrangement permutation $a \in \mathcal{S}_N$ and $d \in \mathcal{D}_N$, the result of the action da on $\sigma \in \mathcal{S}_N$ is $d(a\sigma) \in [a\sigma]$, that is, it results in the same genome as a acting on σ . So we can think of $\mathbf{z}a$ acting on $\mathbf{z}\sigma$ as encompassing (all orientations of (a acting on (all orientations of σ))).

Of course, the action of $\mathbf{z}a$ on $\mathbf{z}\sigma$ also incorporates the dihedral symmetries of a as an action, that is, dad^{-1} for $d \in \mathcal{D}_N$. For a biological model $(\mathcal{M}, w, dist)$ for evolution of genomes as permutations, under the assumption of dihedral symmetry, we wrote (9)

$$\mathcal{M} = \{da_1d^{-1}, da_2d^{-1}, \dots, da_md^{-1} : d \in \mathcal{D}_N\} \subseteq \mathcal{S}_N, \tag{27}$$

where for each a_k and all $d \in \mathcal{D}_N$, $w(da_kd^{-1}) = w(a_k)$. Since $da_kd^{-1} \in [a_k]_D$, the action of $\mathbf{z}(da_kd^{-1})$ on $\mathbf{z}\sigma$ is the same as the action of $\mathbf{z}a_k$ on $\mathbf{z}\sigma$ (see Proposition 3.7) and thus each $\rho_{\text{reg}}^A(\mathbf{z}(da_kd^{-1})) = \rho_{\text{reg}}^A(\mathbf{z}a_k)$.

Having shown that the MLE computations can be performed in the genome algebra \mathcal{A} , and discussed the representation of both genomes and rearrangements as permutation clouds in this algebra, it remains to reformulate the model within this framework. Given the model (27) in the permutation framework, the equivalent model for evolution in the genome algebra setting is $(\mathcal{M}^A, w^A, dist)$, where

$$\mathcal{M}^A := \{\mathbf{z}a_1, \dots, \mathbf{z}a_m\} \subseteq \mathcal{A},$$

and $w^A(\mathbf{z}a_i) = 2Nw(a_i)$ for each i .

Since the dihedral symmetry of the genomes is built into the algebra \mathcal{A} , specifying the model to consist of elements of \mathcal{A} in this way makes the dihedral symmetry requirement (M1) redundant. Model reversibility in this setting is formulated as

$$\text{for each } \mathbf{za} \in \mathcal{M}^{\mathcal{A}}, \mathbf{za}^{-1} \in \mathcal{M}^{\mathcal{A}} \text{ and } w^{\mathcal{A}}(\mathbf{za}^{-1}) = w^{\mathcal{A}}(\mathbf{za}). \tag{M2^{\mathcal{A}}}$$

This condition is sufficient to ensure that the irreducible representations of \mathbf{zs} are diagonalisable, which is convenient for computations. Although the algebra \mathcal{A} has a left identity, it does not contain inverses, so \mathbf{za}^{-1} is not (in general) an inverse of \mathbf{za} . However, as we shall see in the next section, model reversibility is, further, equivalent to the reversibility of the Markov model. We conclude this section with an example to illustrate some of these key concepts.

Example 3.13 Suppose we wish to consider a model consisting only of “small inversions”, which we will take to be inversions of two or three regions. In the permutation framework, we would define this model to be

$$\begin{aligned} \mathcal{M} &:= \{d(1, 2)d^{-1}, d(1, 3)d^{-1} : d \in \mathcal{D}_N\} \\ &= \{(1, 2), (2, 3), \dots, (N - 1, N), (N, 1), (1, 3), (2, 4), \dots, (N - 1, 1), (N, 2)\}. \end{aligned}$$

Here there are N , rather than $2N$, distinct instances of each rearrangement type, since for inversions,⁴ each flip coincides with a rotation.

For the rearrangement probabilities, one could choose the uniform distribution, $w(a) = \frac{1}{2N}$ for all $a \in \mathcal{M}$, or one may consider the larger inversions to be less likely and set, for all $a \in \mathcal{M}$,

$$w_r(a) = \begin{cases} \frac{2}{3N}, & \text{if } a = d(1, 2)d^{-1}, \text{ some } d \in \mathcal{D}_N; \\ \frac{1}{3N}, & \text{if } a = d(1, 3)d^{-1}, \text{ some } d \in \mathcal{D}_N. \end{cases}$$

In the genome algebra, the model is simpler to express; we take the rearrangement instances $(1, 2)$ and $(1, 3)$ and the model is

$$\mathcal{M}^{\mathcal{A}} := \{\mathbf{z}(1, 2), \mathbf{z}(1, 3)\}.$$

The weight functions corresponding to the above would then be $w^{\mathcal{A}}(\mathbf{z}(1, 2)) = w^{\mathcal{A}}(\mathbf{z}(1, 3)) = \frac{1}{2}$ or $w_r^{\mathcal{A}}(\mathbf{z}(1, 2)) = \frac{2}{3}$, $w_r^{\mathcal{A}}(\mathbf{z}(1, 3)) = \frac{1}{3}$.

One may recall from Example 3.6 that $\mathbf{z}(1, 2) \neq \mathbf{z}(2, 3)$, however, one need not (and indeed should not) include both of these in the rearrangement model since they have the same action: $\mathbf{z}(1, 2) \cdot \mathbf{z}\sigma = \mathbf{z}(2, 3) \cdot \mathbf{z}\sigma$ for any genome $\mathbf{z}\sigma \in \mathcal{A}$, since $\mathbf{z}(1, 2)\mathbf{z} = \mathbf{z}(2, 3)\mathbf{z}$. Further, one must be aware of *complementary rearrangements*.

For example, in the case $N = 5$, the actions of $\mathbf{z}(1, 2)$ and $\mathbf{z}(1, 3)$ coincide: inverting a two-region segment is, under dihedral symmetry, the same rearrangement as inverting the complementary three region segment (correspondingly, when $N = 5$, $\mathbf{z}(1, 2)\mathbf{z} = \mathbf{z}(1, 3)\mathbf{z}$). ◇

⁴ In this non-oriented region case.

One can easily eliminate the possibility of such ‘rearrangement redundancies’ by reformulating the model in the genome algebra as a set of elements of the form \mathbf{zaz} ; we do this in the next section (30). More on such considerations, along with explicit examples of rearrangement models in the oriented region case, may be found in Terauds et al. (2021); a deeper algebraic consideration of rearrangements is given in Stevenson et al. (2022).

4 More general models of genomes

The construction of the genome algebra $\mathcal{A} = \mathbf{z}\mathbb{C}[\mathcal{S}_N]$ in Sect. 3 was determined by assumptions we made about how to model the genomes. In particular, following on from previous work (Serdoz et al. 2017; Sumner et al. 2017; Terauds and Sumner 2019), we chose to model circular genomes, without considering orientation of regions, which meant an instance of the genome could be represented by a permutation $\sigma \in \mathcal{S}_N$. We modelled the genomes without a distinguished position, which meant the genome symmetries corresponded to the dihedral group \mathcal{D}_N . In this section, we outline how the constructions and techniques presented for this specific case can be generalised to cover different genomic models: any for which genome (and rearrangement) instances can be represented as elements of a group G .

Suppose, for example, that one wanted to vary the above model to include an origin of replication in the circular genomes. We would model this as a distinguished position and the genomes would then have no rotational symmetry, only reflectional. The symmetry group would thus be $Z_N = \{e, f\}$, the symmetry element $\mathbf{z} = \frac{1}{2}(e + f)$, and each genome an element $\mathbf{z}\sigma = \frac{1}{2}(\sigma + f\sigma) \in \mathbf{z}\mathbb{C}[\mathcal{S}_N]$. The model would naturally reflect this symmetry, with rearrangements taking the form \mathbf{za} , $a \in \mathcal{S}_N$. In particular, this case allows for rearrangements at different positions on the genome, relative to the origin of replication, to be assigned different probabilities. With appropriate choices of rearrangements, this framework could also be used to represent linear genomes.

To include orientation of genes, one would use a different underlying group, for example the hyperoctahedral group H_N of signed permutations (as outlined in Egri-Nagy et al. 2014), and a symmetry group of choice (for example, a copy of the dihedral group in the case of a circular genome with no distinguished positions). An explicit consideration of the genome algebra for the signed region case, including some detailed examples, may be found in Terauds et al. (2021).

To construct the general genome algebra, we begin with a group G , whose elements represent instances of the genomes of interest, and a subgroup $Z \subseteq G$ that represents the physical symmetries of these genomes.⁵ We consider rearrangements such that a single rearrangement event for an instance $g \in G$ of a genome can be modelled via the left action of a particular element $a \in G$ on g . The terms in the following definition reflect our applications of the objects, but obviously the subsequent results concerning the algebras hold whether or not one applies them to genomes.

⁵ To model genomes as possessing no symmetries, one takes the symmetry group to be trivial; this case thus fits within the framework, but each genome simply corresponds to a single permutation.

Definition 4.1 Let G be a finite group with subgroup $Z \subseteq G$. Define

$$\mathbf{z} := \frac{1}{|Z|} \sum_{z \in Z} z, \quad \mathcal{A} := \mathbf{z}C[G] \quad \text{and} \quad \mathcal{A}_0 := \mathbf{z}C[G]\mathbf{z}.$$

We call \mathcal{A} the *genome algebra of G with Z* , \mathcal{A}_0 the *class algebra of G with Z* and \mathbf{z} the *symmetry element of \mathcal{A}* and \mathcal{A}_0 .

Rather than proceeding as in Sects. 2 and 3, where we first defined the rearrangement model, path probabilities and likelihoods for genome instances (group elements) and then showed that the calculations could be performed in the genome algebra, we will here formulate these concepts (and then perform the computations) entirely in the genome algebra \mathcal{A} . We include the class algebra \mathcal{A}_0 for completeness. Following the observations in the previous section, it seems a natural next step to consider the algebra formed by combining together the elements of \mathcal{A} that act indistinguishably. However, we shall see that this lower dimensional algebra is not the appropriate setting for our calculations.

Lemma 4.2 Let G be a finite group with subgroup $Z \subseteq G$.

- (i) For each $g \in G$, define $[g] := \{zg : z \in Z\}$. Then the sets $\{[g] : g \in G\}$ are equivalence classes of G . For each $g \in G$, $|[g]| = |Z|$.
- (ii) For each $g \in G$, define $[g]_D := \{zg z' : z, z' \in Z\}$. Then the sets $\{[g]_D : g \in G\}$ are equivalence classes of G .

Proof Since the sets $[g]$ and $[g]_D$ for $g \in G$ are respectively right cosets and double cosets of G with respect to the subgroup Z , it is clear that they are equivalence classes. □

We use the label ‘ D ’ for the classes defined in (ii) above to refer to the double coset structure of the sets $[g]_D$ (noting that this conveniently coincides with the original usage (Terauds and Sumner 2019) of the label, which referred to the dihedral symmetry in the circular genome case). The following statements can be derived directly from the subgroup properties of Z and Lemma 4.2 (c.f. the corresponding results in Sect. 3).

Proposition 4.3 Let G be a finite group with subgroup $Z \subseteq G$. Let \mathcal{A} and \mathcal{A}_0 respectively be the genome algebra and the class algebra of G with Z and \mathbf{z} the symmetry element of \mathcal{A} and \mathcal{A}_0 . Then

- (i) \mathbf{z} is idempotent, \mathbf{z} is a left identity in \mathcal{A} , and \mathbf{z} is the identity in \mathcal{A}_0 ;
- (ii) \mathcal{A} has a basis of the form $\{\mathbf{z}g : g \in G\} = \{\mathbf{z}g_1, \dots, \mathbf{z}g_K\}$ and

$$K := \dim(\mathcal{A}) = |\{[g] : g \in G\}| = \frac{|G|}{|Z|};$$

- (iii) \mathcal{A}_0 has a basis of the form $\{\mathbf{z}g\mathbf{z} : g \in G\} = \{\mathbf{z}g_1\mathbf{z}, \dots, \mathbf{z}g_L\mathbf{z}\}$ and

$$L := \dim(\mathcal{A}_0) = |\{[g]_D : g \in G\}|.$$

□

For the remainder of this section, we fix a basis for each of \mathcal{A} and \mathcal{A}_0 , as defined in (ii) and (iii) above.

Remark 4.4 Each equivalence class $[g]_D$ can be viewed as an orbit of G under an action of the group $Z \times Z$ and thus, from (iii) above, the dimension L of \mathcal{A}_0 may be calculated via Burnside’s lemma (James and Liebeck 2001, Prop. 29.4). By combining this with the dual orthogonality relations on the group G , one can directly obtain the dimension result stated below in Theorem 4.5 (i).

We note that working ‘entirely’ in the genome algebra does not mean that we forget about the group G . In practice, one would observe a genome with a particular orientation and reference frame, thus as an instance $g \in G$, and then identify the genome as the cloud $\mathbf{z}g \in \mathcal{A}$ for the purposes of computation. There are $K = \frac{|G|}{|Z|}$ distinct genomes, corresponding to the distinct basis elements $\mathbf{z}g$ of the genome algebra \mathcal{A} . Similarly, one would conceive a rearrangement initially as an instance $a \in G$ and then lift to $\mathbf{z}a$ in the genome algebra. Considering all orientations of a rearrangement instance on all orientations of a genome corresponds to a left action of the algebra \mathcal{A} on itself. Distinct elements of \mathcal{A} that correspond to the same element of \mathcal{A}_0 act indistinguishably, since

$$(\mathbf{z}a) \cdot (\mathbf{z}g) = (\mathbf{z}a\mathbf{z}) \cdot (\mathbf{z}g). \tag{28}$$

Thus there are $L = \dim(\mathcal{A}_0)$ distinct rearrangement actions.⁶

The (left) regular representations $\rho_{\text{reg}}^{\mathcal{A}}$ of the genome algebra \mathcal{A} and ρ_{reg}^0 of the class algebra \mathcal{A}_0 can be constructed in the usual way (c.f. (12)) via the bases fixed above. As in Sect. 3, one readily verifies that $\rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z})$ is the $K \times K$ identity matrix and that $\rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}g)^T = \rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}g^{-1})$. Since \mathbf{z} is an identity in \mathcal{A}_0 , $\rho_{\text{reg}}^0(\mathbf{z})$ is the $L \times L$ identity matrix. In this case, the equivalence classes $[g]_D$ need not be the same size and thus, in general, $\rho_{\text{reg}}^0(\mathbf{z}g\mathbf{z})^T \neq \rho_{\text{reg}}^0(\mathbf{z}g^{-1}\mathbf{z})$.⁷ We denote the regular characters of \mathcal{A} and \mathcal{A}_0 by $\chi_{\text{reg}}^{\mathcal{A}}$ and χ_{reg}^0 respectively and note that these take real values on any algebra element that is a real linear combination of basis elements.

Recall that, by Maschke’s theorem (Etingof et al. 2011, Thm. 4.1.1), the group algebra $\mathbb{C}[G]$ of any finite group G can be written as a direct sum over its irreducible modules.

Theorem 4.5 *Let G be a finite group with subgroup $Z \subseteq G$. Denote the distinct irreducible submodules of $\mathbb{C}[G]$ by V_i , with $\dim(V_i) = D_i$ for each i so that*

$$\mathbb{C}[G] \cong \bigoplus_{i=1}^M D_i V_i,$$

and denote the corresponding irreducible representations and characters of $\mathbb{C}[G]$ by ρ_i and χ_i respectively. Then the following hold.

⁶ We note that most of these mathematically possible rearrangements would not correspond to biologically plausible ones, so would not appear in rearrangement models in practice. For a deeper consideration of biologically plausible rearrangements from an algebraic perspective, see Stevenson et al. (2022).

⁷ One can verify via a simple counting argument that $(\rho_{\text{reg}}^0(\mathbf{z}g\mathbf{z}))_{ij} = \frac{|[g_i]_D|}{|[g_j]_D|} (\rho_{\text{reg}}^0(\mathbf{z}g^{-1}\mathbf{z}))_{ji}$.

(i) For each i , $W_i := \mathbf{z} \cdot V_i$ is either $\{0\}$ or an irreducible \mathcal{A}_0 -module and

$$\mathcal{A}_0 \cong \bigoplus_{\substack{1 \leq i \leq M \\ W_i \neq \{0\}}} k_i W_i, \tag{29}$$

with $k_i = \dim(W_i) = \chi_i(\mathbf{z})$ for each i . Thus, $\dim(\mathcal{A}_0) = L = \sum_{i=1}^M \chi_i(\mathbf{z})^2$.

(ii) The modules $\{W_i : 1 \leq i \leq M, W_i \neq \{0\}\}$ comprise all irreducible modules of \mathcal{A} . Denoting the corresponding irreducible representations of \mathcal{A} and \mathcal{A}_0 by $\rho_i^{\mathcal{A}}$ and ρ_i^0 respectively,

$$\rho_i^{\mathcal{A}}(\mathbf{z}g) = \rho_i^{\mathcal{A}}(\mathbf{z}g\mathbf{z}) = \rho_i^0(\mathbf{z}g\mathbf{z}),$$

for all $g \in G$ and all (relevant) $1 \leq i \leq M$. Thus for all $\mathbf{g} \in \mathcal{A}_0$, $\rho_i^{\mathcal{A}}(\mathbf{g}) = \rho_i^0(\mathbf{g})$.

(iii) Denoting the corresponding characters of \mathcal{A} and \mathcal{A}_0 by $\chi_i^{\mathcal{A}}$ and χ_i^0 respectively, and defining $\chi_i^{\mathcal{A}} = \chi_i^0 \equiv 0$ for each i such that $W_i = \{0\}$,

$$\chi_i(\mathbf{z}g) = \chi_i^{\mathcal{A}}(\mathbf{z}g) = \chi_i^0(\mathbf{z}g\mathbf{z}),$$

for all $g \in G$ and all $1 \leq i \leq M$. Thus the characters χ_i , $\chi_i^{\mathcal{A}}$ and χ_i^0 coincide on \mathcal{A}_0 .

(iv) For all $\mathbf{g} \in \mathcal{A}$,

$$\chi_{reg}^{\mathcal{A}}(\mathbf{g}) = \sum_{i=1}^M D_i \chi_i^{\mathcal{A}}(\mathbf{g}) = \chi_{reg}(\mathbf{g}),$$

where χ_{reg} and $\chi_{reg}^{\mathcal{A}}$ denote the regular characters of $\mathbb{C}[G]$ and \mathcal{A} respectively.

(v) For any $\mathbf{g} \in \mathcal{A}$, $\left(\frac{1}{K} \cdot \chi_{reg}^{\mathcal{A}}(\mathbf{g})\right)$ is the coefficient of \mathbf{z} in \mathbf{g} .

Proof For (i), we use (Steinberg 2016, Prop. 4.18, Thm. 4.23). For the remaining results, we use the observation (28) and proceed just as for the corresponding results in Sect. 3. Note that in this general setting we cannot assume that the irreducible representations ρ_i are orthogonal on G , but we can choose them to be unitary (Etingof et al. 2011, Thm. 4.6.2). This means that the corresponding irreducible representations of \mathbf{z} in $\mathbb{C}[G]$ are self adjoint, and thus each $\rho_i(\mathbf{z})$ is unitarily diagonalisable, so that its eigenvectors form an orthonormal basis for $\mathbb{C}^{D_i} \cong V_i$. Since the eigenvectors need not be real, the only difference in the proofs is that we need the conjugate transposes, not just transposes, of these vectors. \square

The above results imply that $\mathcal{A}_0 \cong \mathcal{A}/\text{Rad}(\mathcal{A})$ (Etingof et al. 2011, Thm. 3.5.4), which formalises the relationship between the genome algebra and the class algebra: \mathcal{A}_0 is obtained from \mathcal{A} by factoring out the elements of \mathcal{A} that act trivially. We have previously expressed this as \mathcal{A}_0 combining together the elements of \mathcal{A} that act indistinguishably. Another aspect of this is the following.

Corollary 4.6 *Let G be a finite group with subgroup $Z \subseteq G$. For any $g \in G$ and each irreducible representation ρ_i^A of \mathcal{A} , $\rho_i^A(\mathbf{z}g) = \rho_i^A(\mathbf{z}g')$ for all $g' \in [g]_D$. For any $g, g' \in G$, $\rho_{reg}^A(\mathbf{z}g) = \rho_{reg}^A(\mathbf{z}g')$ if and only if $g' \in [g]_D$. \square*

Remark 4.7 We note that (Steinberg 2016, Prop. 7.14) implies part of Theorem 4.5 (iii) (namely, that $\chi_i^0(\mathbf{g}) = \chi_i(\mathbf{g})$ for $\mathbf{g} \in \mathcal{A}_0$) in the more general case of G being a semigroup. Extending the framework to algebras based on semigroups would allow us to model further types of rearrangements, such as insertions and deletions (Francis 2014), and we intend to investigate this possibility in future work.

We are ready to proceed with the formulation of path probabilities and the likelihood function within the genome algebra \mathcal{A} . Firstly, we formally define a biological model for evolution in the genome algebra to be $(\mathcal{M}, w, dist)$, where

$$\mathcal{M} := \{\mathbf{za}_1\mathbf{z}, \mathbf{za}_2\mathbf{z}, \dots, \mathbf{za}_q\mathbf{z}\} \subseteq \mathcal{A}, \tag{39}$$

for some $a_1, \dots, a_q \in G$, $w : \mathcal{M} \rightarrow (0, 1)$ is the probability distribution on \mathcal{M} , and $dist$ is the probability distribution of rearrangement events in time. Note that we have used the form \mathbf{zaz} rather than \mathbf{za} to avoid duplicating rearrangements in the model (that is, elements $\mathbf{za}, \mathbf{za}' \in \mathcal{A}$ that are distinct but have the same left action, c.f. Examples 3.6 and 3.13). Presently, we shall also add the condition that the model be *reversible*, that is, that $\mathbf{za}^{-1}\mathbf{z} \in \mathcal{M}$ for every $\mathbf{zaz} \in \mathcal{M}$, and $w(\mathbf{za}^{-1}\mathbf{z}) = w(\mathbf{zaz})$.

We fix the reference genome to be $\mathbf{z} \in \mathcal{A}$ (whose instances are the elements of Z , in particular $e \in Z$). Then, for any target genome $\mathbf{z}g$ (where we have observed the instance $g \in G$) and each $k \in \mathbb{N}_0$, we define the path probability $\alpha_k(\mathbf{z}g)$ to be

$$\alpha_k(\mathbf{z}g) := P(\mathbf{z} \mapsto \mathbf{z}g \text{ via } k \text{ rearrangements}),$$

where $\mathbf{z} \mapsto \mathbf{z}g$ means ‘‘genome \mathbf{z} is transformed into genome $\mathbf{z}g$ ’’. As usual, to find the path probability for an arbitrary genome $\mathbf{z}h$ to be transformed into target genome $\mathbf{z}g$, we can simply translate to the reference; that is, this is exactly the path probability for \mathbf{z} to be rearranged into $\mathbf{z}gh^{-1}$, which is $\alpha_k(\mathbf{z}gh^{-1})$.

Given a model \mathcal{M} , we define the corresponding *model element* of \mathcal{A} to be

$$\tilde{\mathbf{s}} := \sum_{i=1}^q w(\mathbf{za}_i\mathbf{z})\mathbf{za}_i.$$

We write $\tilde{\mathbf{s}}$ to distinguish the model element here from the previous definition of \mathbf{s} in the group algebra (c.f. Definition 2.2), and choose to sum over rearrangements of the form \mathbf{za}_i rather than $\mathbf{za}_i\mathbf{z}$ for simplicity (recall from (28), these have the same action, so either form may be used).

It remains to connect the path probabilities to the regular character of powers of the model element (c.f (3) in Sect. 2). Recall from Sect. 3.1 that $\mathbf{za} \cdot \mathbf{z}g$ gives a convex combination of genomes, that is,

$$\mathbf{z}a \cdot \mathbf{z}g = \sum_{i=1}^K \mathfrak{p}_i \mathbf{z}g_i,$$

where $\{\mathbf{z}g_1, \dots, \mathbf{z}g_K\}$ is our fixed basis for \mathcal{A} and each \mathfrak{p}_i is the proportion of the expansion of $\mathbf{z}a\mathbf{z}g$ that is equal to $\mathbf{z}g_i$ or, equivalently, the probability that the rearrangement $\mathbf{z}a$ acting on the genome $\mathbf{z}g$ will result in the genome $\mathbf{z}g_i$. Thus

$$\tilde{\mathfrak{s}} \cdot \mathbf{z} = \sum_{j=1}^q w(\mathbf{z}a_j\mathbf{z})\mathbf{z}a_j\mathbf{z} = \sum_{j=1}^q w(\mathbf{z}a_j\mathbf{z}) \sum_{i=1}^K \mathfrak{p}_{j,i} \mathbf{z}g_i = \sum_{i=1}^K \left(\sum_{j=1}^q w(\mathbf{z}a_j\mathbf{z})\mathfrak{p}_{j,i} \right) \mathbf{z}g_i,$$

where we have rearranged and collected terms in the final step so that, for each i , $\sum_{j=1}^q w(\mathbf{z}a_j\mathbf{z})\mathfrak{p}_{j,i}$ is the total probability that the genome \mathbf{z} will be transformed into the genome $\mathbf{z}g_i$ via some (single) rearrangement chosen from the model. Thus

$$\tilde{\mathfrak{s}} \cdot \mathbf{z} = \sum_{i=1}^K \alpha_1(\mathbf{z}g_i)\mathbf{z}g_i$$

and, by repeatedly applying $\tilde{\mathfrak{s}}$, one sees that

$$\tilde{\mathfrak{s}}^k \mathbf{z} = \sum_{i=1}^K \alpha_k(\mathbf{z}g_i)\mathbf{z}g_i. \tag{31}$$

Now, for $g \in G$ an instance of the genome of interest, multiply (31) on the right by g^{-1} to obtain

$$\tilde{\mathfrak{s}}^k \mathbf{z}g^{-1} = \sum_{i=1}^K \alpha_k(\mathbf{z}g_i)\mathbf{z}g_i g^{-1}.$$

Since $\mathbf{z}g_i g^{-1} = \mathbf{z}$ if and only if $\mathbf{z}g_i = \mathbf{z}g$, we see that $\alpha_k(\mathbf{z}g)$ is the coefficient of \mathbf{z} in the expansion of $\tilde{\mathfrak{s}}^k \mathbf{z}g^{-1}$, and thus

$$\alpha_k(\mathbf{z}g) = \frac{1}{K} \cdot \chi_{\text{reg}}^{\mathcal{A}} \left(\tilde{\mathfrak{s}}^k \mathbf{z}g^{-1} \right) \tag{32}$$

by Theorem 4.5 (v).

Theorem 4.5 allows us to decompose the regular character in (32) into irreducible characters, however we also will need to diagonalise the irreducible representation matrices.

Lemma 4.8 *Let G be a finite group with subgroup $Z \subseteq G$. Let $(\mathcal{M}, w, \text{dist})$ be a biological model for evolution of genomes represented by elements $\mathbf{z}g \in \mathcal{A}$ of the genome algebra (where $g \in G$) and let $\tilde{\mathfrak{s}} \in \mathcal{A}$ be the corresponding model element. If the model is reversible, then the following hold.*

- (i) The irreducible representation matrices of the model element $\tilde{\mathfrak{s}}$ in \mathcal{A} are diagonalisable.
- (ii) The regular representation of $\tilde{\mathfrak{s}}$ in \mathcal{A} is symmetric.

Proof (i) Suppose that the model is reversible and let $1 \leq i \leq q$. We have (c.f. (20))

$$\rho_i^{\mathcal{A}}(\tilde{\mathfrak{s}}) = \rho_i^{\mathcal{A}}(\tilde{\mathfrak{s}}\mathbf{z}) := \overline{Q}^T \rho_i(\tilde{\mathfrak{s}}\mathbf{z}) Q,$$

where Q is the $D_i \times k_i$ matrix of orthonormal eigenvectors for $\rho_i(\mathbf{z})$. Since G is a finite group, we may choose the irreducible representation ρ_i on G to be unitary. Then, writing $\rho_i(\tilde{\mathfrak{s}}\mathbf{z})$ as a sum of matrices of the form $w(\mathbf{z}\mathbf{a}\mathbf{z}) (\rho_i(\mathbf{z}\mathbf{a}\mathbf{z}) + \rho_i(\mathbf{z}\mathbf{a}^{-1}\mathbf{z}))$ (omitting the second term if $a = a^{-1}$), each of which is self adjoint, we see that $\rho_i(\tilde{\mathfrak{s}}\mathbf{z})$ is self-adjoint and thus so is $\rho_i^{\mathcal{A}}(\tilde{\mathfrak{s}})$. For claim (ii), we proceed similarly, using the observation that $\rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\mathbf{a})^T = \rho_{\text{reg}}^{\mathcal{A}}(\mathbf{z}\mathbf{a}^{-1})$. \square

Theorem 4.9 *Let G be a finite group with subgroup $Z \subseteq G$. Let $(\mathcal{M}, w, \text{dist})$ be a reversible biological model for evolution of genomes represented by basis elements of the genome algebra \mathcal{A} of G with Z and let $\tilde{\mathfrak{s}} \in \mathcal{A}$ be the corresponding model element. Let $g \in G$ be an observed instance of a genome $\mathbf{z}g \in \mathcal{A}$. Then the following hold.*

- (i) For any $k \in \mathbb{N}_0$, the probability that the reference genome \mathbf{z} is transformed into the genome $\mathbf{z}g$ via k rearrangements chosen from the model is

$$\alpha_k(\mathbf{z}g) = \frac{|Z|}{|G|} \chi_{\text{reg}}^{\mathcal{A}}(\tilde{\mathfrak{s}}^k \mathbf{z}g^{-1}) = \frac{|Z|}{|G|} \sum_{i=1}^M D_i \sum_{j=1}^{R_i} \lambda_{i,j}^k \text{tr}(\rho_i^{\mathcal{A}}(\mathbf{z}g^{-1}) E_{i,j}^{\mathcal{A}}),$$

where for each i , $E_{i,j}^{\mathcal{A}}$ is the projection onto the eigenspace of the j th eigenvalue $\lambda_{i,j}$ of $\rho_i^{\mathcal{A}}(\tilde{\mathfrak{s}})$.

- (ii) If the distribution of rearrangement events in time is $\text{dist} = \text{Poisson}(1)$, then the probability that the reference genome \mathbf{z} is transformed into the genome $\mathbf{z}g$ via the given model in time T is given by the likelihood function

$$L(T|\mathbf{g}) = e^{-T} \frac{|Z|}{|G|} \sum_{i=1}^M D_i \sum_{j=1}^{R_i} \text{tr}(\rho_i^{\mathcal{A}}(\mathbf{z}g^{-1}) E_{i,j}^{\mathcal{A}}) e^{\lambda_{i,j} T}.$$

- (iii) For any genome $\mathbf{z}h \in \mathcal{A}$ with an instance $h \in [g]_D \cup [g^{-1}]_D$, the path probabilities and likelihood functions of $\mathbf{z}g$ and $\mathbf{z}h$ coincide.

Proof The first expression for the path probability $\alpha_k(\mathbf{z}g)$ was gained above (32). To gain the second, we use the decomposition of the regular character from Theorem 4.5 (iv), and then for each i , use the cyclicity of the trace to write

$$\chi_i^{\mathcal{A}}(\tilde{\mathfrak{s}}^k \mathbf{z}g^{-1}) = \text{tr} \left(\rho_i^{\mathcal{A}}(\mathbf{z}g^{-1}) \left(\rho_i^{\mathcal{A}}(\tilde{\mathfrak{s}}) \right)^k \right).$$

Then, from Lemma 4.8, we may diagonalise $\rho_i^A(\tilde{\mathfrak{S}})$ to gain the second expression. Analogously to the definition in Sect. 2, but with genomes instead of elements of G , we define the likelihood function as

$$\begin{aligned} L(T|g) &:= P(\mathbf{z}g|T) = \sum_{k=0}^{\infty} P(\mathbf{z} \mapsto \mathbf{z}g \text{ via } k \text{ events})P(k \text{ events in time } T) \\ &= \sum_{k=0}^{\infty} \alpha_k(\mathbf{z}g) \frac{e^{-T} T^k}{k!}. \end{aligned}$$

Then substituting in the expression from (i) and simplifying the power series gives (ii).

(iii) Let $h \in G$ such that $\mathbf{z}h\mathbf{z} = \mathbf{z}g\mathbf{z}$ or $\mathbf{z}h\mathbf{z} = \mathbf{z}g^{-1}\mathbf{z}$. We show that $\alpha_k(\mathbf{z}g) = \alpha_k(\mathbf{z}h)$ for all $k \in \mathbb{N}_0$, which implies (iii). Let $k \in \mathbb{N}_0$. Since the trace is cyclic, we have

$$\chi_{\text{reg}}^A(\tilde{\mathfrak{S}}^k \mathbf{z}g^{-1}) = \text{tr} \left(\rho_{\text{reg}}^A(\mathbf{z}g^{-1}) \rho_{\text{reg}}^A(\tilde{\mathfrak{S}})^k \right) = \text{tr} \left(\rho_{\text{reg}}^A(\tilde{\mathfrak{S}})^k \rho_{\text{reg}}^A(\mathbf{z}g) \right) = \chi_{\text{reg}}^A(\tilde{\mathfrak{S}}^k \mathbf{z}g),$$

where the second equality was obtained by taking the transpose of the argument and applying Lemma 4.8. Then from (i) it is clear that $\alpha_k(\mathbf{z}g) = \alpha_k(\mathbf{z}g^{-1})$ and these coincide with $\alpha_k(\mathbf{z}h)$ by Corollary 4.6. □

Since the regular representation of the model element is symmetric, by Lemma 4.8, and the equilibrium distribution is the uniform distribution on the set of genomes, reversibility of the model \mathcal{M} is equivalent to time reversibility of the underlying Markov process. As in Sect. 3.1, the regular representation of the model element in \mathcal{A} is the transition matrix for a Markov chain with states being genomes, with the probability that genome $\mathbf{z}g_j$ transitions into genome $\mathbf{z}g_i$ via k rearrangement steps from the model given by

$$\rho_{\text{reg}}^A(\tilde{\mathfrak{S}}^k)_{ij} = \alpha_k(\mathbf{z}g_i g_j^{-1}).$$

Reversibility then means that for any genomes $\mathbf{z}g, \mathbf{z}h \in \mathcal{A}$, the probability of $\mathbf{z}g$ transforming into $\mathbf{z}h$ in k steps via the given model is the same as that of $\mathbf{z}h$ transforming into $\mathbf{z}g$ in k steps. In terms of path probabilities,

$$\alpha_k(\mathbf{z}g h^{-1}) = \alpha_k(\mathbf{z}h g^{-1}),$$

which is just a special case of Theorem 4.9 (iii). Model reversibility thus implies that the MLE distance is ‘directionless’, or symmetric, as is any evolutionary distance measure based on path probabilities calculated in this framework.

To conclude this section, we return briefly to the class algebra \mathcal{A}_0 . By constructing simple examples, one can verify that, in general, the regular representation matrices of non-identity basis elements in \mathcal{A}_0 have non-zero entries on the diagonal, and thus see that we do not have an analogue of Theorem 4.5 (v) for \mathcal{A}_0 . That is, the regular character of \mathcal{A}_0 is not counting occurrences of the identity in elements of this algebra, and thus cannot be used to calculate path probabilities as in Theorem 4.9.

Consider the underlying Markov model here, with transition matrix given by the regular representation of the model element analogue $\rho_{\text{reg}}^0(\tilde{\mathbf{S}}\mathbf{z})$. Now the states are the basis elements $\mathbf{z}g_i\mathbf{z}$, each corresponding to an equivalence class $[g_i]_D$. Since each equivalence class $[g]_D$ is the disjoint union of $\frac{|[g]_D|}{|Z|}$ equivalence classes of the form $[gz]$ for some $z \in Z$, each basis element $\mathbf{z}g\mathbf{z}$ is the average of $\frac{|[g]_D|}{|Z|}$ distinct genomes of the form $\mathbf{z}gz$ for some $z \in Z$. Thus an arbitrary element of the matrix gives us the *average* probability of a genome from a certain class transitioning into a genome from another class (and a diagonal element gives the probability of a transition within a class). This is not refined enough for our purposes, since, given one genome $\mathbf{z}g$ and two more genomes $\mathbf{z}g'$ and $\mathbf{z}g''$ that are in the same class ($g' \in [g'']_D$), the probability of transitioning between $\mathbf{z}g$ and $\mathbf{z}g'$ need not be the same as the probability of transitioning between $\mathbf{z}g$ and $\mathbf{z}g''$.

The information is not entirely lost, however; one can use the first column of this Markov matrix to calculate path probabilities. Given an observed instance $g \in G$ of a genome, we find the appropriate basis element $\mathbf{z}g_\ell\mathbf{z}$ of \mathcal{A}_0 such that $g \in [g_\ell]_D$, and then

$$\alpha_k(\mathbf{z}g) = \frac{|Z|}{|[g_\ell]_D|} \left(\rho_{\text{reg}}^0((\tilde{\mathbf{S}}\mathbf{z})^k) \right)_{\ell 1} .$$

Of course, the fact that this path probability information exists in the regular representation does not mean that it is *easy* to obtain, in particular since the size of the regular representation in \mathcal{A}_0 is likely to be rapidly increasing with the number of genomic regions (for example for $G = \mathcal{S}_N$ and $Z = \mathcal{D}_N$, $\dim(\mathcal{A}_0)$ is proportional to $(N - 2)!$) and, being unable to retain the ‘first column’ information through diagonalisation (as one can for the trace), one would need to calculate the k th power of the matrix for each desired path probability. We further note that calculating the equivalence classes themselves, and checking for membership of an equivalence class, is a non-trivial exercise and simply not feasible for large numbers of regions.

In any case, the class algebra \mathcal{A}_0 is nicer in some ways than the genome algebra \mathcal{A} , in particular in that it is decomposable (that is, isomorphic to a direct sum of its irreducible modules). This property is formally known as *semisimplicity*. Then, since the irreducible modules of \mathcal{A}_0 are identical to those of the algebra \mathcal{A} and the irreducible representations of the two algebras not only have the same dimension but coincide on the objects of interest,

$$\rho_p^0(\mathbf{z}g\mathbf{z}) = \rho_p^{\mathcal{A}}(\mathbf{z}g\mathbf{z}) = \rho_p^{\mathcal{A}}(\mathbf{z}g) ,$$

one may in fact choose to implement the calculations of the irreducible representations in \mathcal{A} or \mathcal{A}_0 and then, either way, combine the results together according to the decomposition given in Theorem 4.9.

5 Conclusion

We have presented a coherent algebraic framework for modelling some classes of genomes and rearrangements in an algebra that incorporates the inherent physical symmetries into each element. Algebraic frameworks for modelling genome rearrangement have been studied previously (Meidanis and Dias 2000; Moulton and Steel 2012; Francis 2014), and the importance of including genome symmetry in rearrangement distance calculations has been recognised (Egri-Nagy et al. 2014; Serdoz et al. 2017), however our unified approach, incorporating symmetry into the position paradigm framework (Bhatia et al. 2018), is new.

Beginning with the specific case of circular genomes modelled with unoriented regions and dihedral symmetry, we explicitly constructed the genome algebra from the symmetric group algebra, and showed that the MLE computations can be performed entirely within this algebra. By identifying genomes and rearrangements with single elements—permutation clouds—in the genome algebra, we have advanced previous work that identified genomes with cosets of permutations (where each element of a given coset represents an instance of a genome in a fixed physical orientation) but used the permutations as the basis elements for computation (Egri-Nagy et al. 2014; Serdoz et al. 2017; Sumner et al. 2017; Terauds and Sumner 2019). We have both explained and removed the redundancy that we identified (Terauds and Sumner 2019) in the implementation of the calculations in the symmetric group algebra.

In Terauds and Sumner (2019), we also signalled a desire to extend our technique for calculating the MLE to other settings, for example to include oriented regions or genomes with non-dihedral symmetry. We have not recorded the results of any explicit computations here, however we have algebraically verified that the technique can indeed be extended to a much more general case. For genomes where a single physical orientation can be represented by elements of a group G , and their physical symmetries by the subgroup $Z \subseteq G$, we defined the genome algebra of G with Z ; here, as in the special case described above, genomes and rearrangements correspond to basis elements (clouds). We showed that the path probabilities and thus the MLE can be formulated and the computations performed entirely in this genome algebra. An application of the framework to modelling signed circular genomes, using the hyperoctahedral group and two possible symmetry groups, is presented in Terauds et al. (2021), along with the results of some sample computations that illustrate how the framework may be applied to compare different models and distance measures.

Although the genome algebra has lower dimension than the group algebra (by a factor of $\frac{1}{2N}$ in the \mathcal{D}_N case and $\frac{1}{|Z|}$ in the general case), this does not significantly reduce the computational complexity of calculating the MLE. We have performed distance calculations, in reasonable time, for genomes with up to twelve unoriented regions (unpublished) and up to six oriented regions (Terauds et al. 2021). Work to extend our initial experimental calculations to implementation of the framework for larger numbers of regions is ongoing. In particular, we are exploring the use of simulations and intend to apply numerical approximations to make distance calculations tractable for genomes with larger numbers of regions.

Whilst the framework does not specify a particular rearrangement model (and indeed, allows choice both in the type of rearrangements allowed and their rela-

tive probabilities of occurring), we cannot currently model insertions, deletions, or duplications, since the underlying group structure means we are restricted to rearrangements that do not alter the set of regions. This is a clear limitation of the current approach. To address this, we are currently working on extending the framework to a semigroup-based approach, with the aim of accommodating insertions and deletions. Furthermore, whilst one can apply different probabilities to different rearrangements (and, depending on the genome's symmetry, rearrangements at different genomic positions), the current approach does not incorporate intergenic regions or explicitly consider breakpoints. Whether a group- or semigroup-based genome algebra approach can be devised that incorporates these biological realities, and others such as multiple chromosomes, is another question for future research.

Finally, we note that the applications of this algebraic framework are not limited to calculating MLEs. The likelihood function is built from path probabilities; since our fundamental results hold for these 'building blocks', other rearrangement distance measures that are based on path probabilities may be calculated via the genome algebra. We have shown that, via the regular representation of the genome algebra, the general genome rearrangement model can be viewed as a discrete (or, with the addition of the stochastic component, continuous time) Markov chain, and thus represented as a connected graph, generalising the Cayley graph approach (Moulton and Steel 2012; Clark et al. 2019). This facilitates the calculation of further distance measures, for example mean first passage time, as demonstrated in Terauds et al. (2021).

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bader David A, Moret Bernard ME, Yan Mi (2001) A linear-time algorithm for computing inversion distance between signed permutations with an experimental study. *J Comput Biol* 8(5):483–491
- Bader M, Ohlebusch E (2006) Sorting by weighted reversals, transpositions, and inverted transpositions. In: Research in computational molecular biology, 10th Annual International Conference, RECOMB 2006, Venice, Italy, April 2–5, 2006, Proceedings, pp 563–577
- Belda E, Moya A, Silva Francisco J (2005) Genome rearrangement distances and gene order phylogeny in *y*-proteobacteria. *Mol Biol Evolut* 22(6):1456–1467
- Bhatia S, Feijão P, Francis AR (2018) Position and content paradigms in genome rearrangements: the wild and crazy world of permutations in genomics. *Bull Math Biol* 80(12):3227–3246
- Chen L, Chen P-Y, Xue X-F, Hua H-Q, Li Y-X, Zhang F, Wei S-J (2018) Extensive gene rearrangements in the mitochondrial genomes of two egg parasitoids, *Trichogramma japonicum* and *Trichogramma ostriniae* (hymenoptera: Chalcidoidea: Trichogrammatidae). *Sci Rep* 8
- Clark C, Egri-Nagy A, Francis A, Gebhardt V (2019) Bacterial phylogeny in the Cayley graph. *Discrete Math Algorithms Appl* 11(05):1950059
- Darmon E, Leach DRF (2014) Bacterial genome instability. *Microbiol Mol Biol Rev* 78(1):1–39

- Dobzhansky T, Sturtevant AH (1938) Inversions in the chromosomes of *Drosophila pseudoobscura*. *Genetics* 23(1):28–64
- Egri-Nagy A, Gebhardt V, Tanaka MM, Francis AR (2014) Group-theoretic models of the inversion process in bacterial genomes. *J Math Biol* 69(1):243–265
- Etingof P, Golberg O, Hensel S, Liu T, Schwendner A, Vaintrob D, Yudovina E (2011) Introduction to representation theory, volume 59 of Student Mathematical Library. American Mathematical Society, Providence, RI. With historical interludes by Slava Gerovitch
- Felsenstein J (2004) Inferring phylogenies. Sinauer Associates, Sunderland
- Francis AR (2014) An algebraic view of bacterial genome evolution. *J Math Biol* 69(6–7):1693–1718
- James G, Liebeck M (2001) Representations and characters of groups, 2nd edn. Cambridge University Press, New York
- Meidanis J, Dias Z (2000) An alternative algebraic formalism for genome rearrangements. In: Sankoff D, Nadeau JH (eds) *Comparative genomics*. *Comput Biol* 1: 213–223
- Moulton V, Steel M (2012) The ‘butterfly effect’ in Cayley graphs with applications to genomics. *J Math Biol* 65(6–7):1267–1284
- Oesper L, Dantas S, Raphael BJ (2018) Identifying simultaneous rearrangements in cancer genomes. *Bioinformatics* 34(2):346–352
- Oliveira AR, Jean G, Fertin G, Dias U, Dias Z (2019) Super short operations on both gene order and intergenic sizes. *Algorithms Mol Biol* 14(1):1–17
- Sagan BE (2001) The symmetric group, volume 203 of Graduate texts in mathematics, 2nd edn. Springer, New York
- Serdoz S, Egri-Nagy A, Sumner J, Holland BR, Jarvis PD, Tanaka MM, Francis AR (2017) Maximum likelihood estimates of pairwise rearrangement distances. *J Theor Biol* 423:31–40
- Steinberg B (2016) Representation theory of finite monoids. Springer, Cham
- Stevenson J, Terauds V, Sumner J (2022) Rearrangement events on circular genomes. arXiv preprint [arXiv:2202.01968](https://arxiv.org/abs/2202.01968)
- Sumner JG, Jarvis PD, Francis AR (2017) A representation-theoretic approach to the calculation of evolutionary distance in bacteria. *J Phys A Math Theor* 50(33):335601
- Terauds V, Sumner J (2019) Maximum likelihood estimates of rearrangement distance: implementing a representation-theoretic approach. *Bull Math Biol* 81(2):535–567
- Terauds V, Stevenson J, Sumner J (2021) A symmetry-inclusive algebraic approach to genome rearrangement. *J Bioinform Comput Biol* 19(06):2140015
- Wang L-S, Warnow T, Moret BME, Jansen RK, Raubeson LA (2006) Distance-based genome rearrangement phylogeny. *J Mol Evol* 63(4):473–483

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.