The Gene Family-Free Median of Three

Daniel Doerr[1(✉)], Pedro Feijão[2], Metin Balaban[1], and Cedric Chauve[3(✉)]

[1] School of Computer and Communication Sciences, EPFL,
1015 Lausanne, Switzerland
daniel.doerr@epfl.ch
[2] Faculty of Technology, Bielefeld University, 33615 Bielefeld, Germany
[3] Department of Mathematics, Simon Fraser University, Burnaby, BC, Canada
cedric.chauve@sfu.ca

Abstract. The gene family-free framework for comparative genomics aims at developing methods for gene order analysis that do not require prior gene family assignment, but work directly on a sequence similarity graph. We present a model for constructing a median of three genomes in this family-free setting, based on maximizing an objective function that generalizes the classical breakpoint distance by integrating sequence similarity in the score of a gene adjacency. We show that the corresponding computational problem is MAX SNP-hard and we present a 0–1 linear program for its exact solution. The result of this program is a median genome with median genes associated to extant genes, in which median adjacencies are assumed to define positional orthologs. We demonstrate through simulations and comparison with the OMA orthology database that the herein presented method is able compute accurate medians and positional orthologs for genomes comparable in size of bacterial genomes.

1 Introduction

The prediction of evolutionary relationships between genomic sequences is a long-standing problem in computational biology. According to Fitch [8], two genomic sequences are called *homologous* if they descended from a common ancestral sequence. Furthermore, Fitch identifies different events that give rise to a branching point in the phylogeny of homologous sequences, leading to the well-established concepts of orthologous genes (who descend from their last common ancestor through a speciation) and paralogous genes (descending from their last common ancestor through a duplication) [9]. Until quite recently, orthology and paralogy relationships were mostly inferred from sequence similarity. However it is now well accepted that the syntenic context can carry valuable evolutionary information, which has lead to the notion of *positional orthologs* [5], which are orthologs whose syntenic context was not changed in a duplication event. In the present work, we describe a method to compute groups of likely orthologous genes for a group of three genomes, through a new problem we introduce, the *gene family-free median of three*.

Most methods for detecting potential orthologous groups require a prior clustering of the genes of the considered genomes into *homologous gene families*,

© Springer International Publishing Switzerland 2016
M. Frith and C.N.S. Pedersen (Eds.): WABI 2016, LNBI 9838, pp. 102–120, 2016.
DOI: 10.1007/978-3-319-43681-4_9

defined as groups of genes assumed to originate from a single ancestral gene. Yet clustering protein sequences into families is already in itself a difficult problem.

Here, we follow the matching-based approach, framed within the gene family-free principle, that embodies the idea to perform gene order analysis without the prerequisite of gene family or homology assignments. Instead, we are given all-against-all *gene similarities* through a symmetric and reflexive *similarity measure* $\sigma : \Sigma \times \Sigma \rightarrow \mathbb{R}_{\geq 0}$ over the universe of genes Σ [3]. We use sequence similarity but other similarity measures can fit the previous definition. This leads to the formalization of the *gene similarity graph* [3], i.e. a graph where each vertex corresponds to a gene of the dataset and where each pair of vertices associated with genes of distinct genomes are connected by a strictly positively weighted edge according to gene similarity measure σ. Gene family or homology assignments represent a particular subgroup of gene similarity functions that require transitivity. Independent of the particular similarity measure σ, relations between genes imposed by σ are considered as candidates for homology assignments. A gene family-free research program was outlined in [3] (see also [7]) and has so far been developed for the pairwise comparison of genomes [6,10,13] and shown to be effective for orthology analysis [11].

In Sect. 2 we introduce a new genome median problem in the family-free framework, that generalizes the traditional breakpoint median problem [16]. For a group of three genomes, the input of the family-free median problem is a tripartite similarity graph of pairwise gene similarities. Informally, a median of three is defined as a genome, and as such is composed of a set of median genes that are associated to the genes of the input genomes and that give rise to one or more linear or circular gene order sequences. The matching of median genes to input genes as well as their ordering in the median genome is subject to an optimization problem. Hereby, our optimization criterion fully integrates both sequence similarity and gene order conservation.

In Sect. 3 we study its the computational complexity and give an exact algorithm for its solution. We show that our method can be used for positional ortholog prediction in simulated and real data sets of bacterial genomes in Sect. 4.

2 The Gene Family-Free Median of Three

Extant genomes, genes and adjacencies. In this work, a genome G is entirely represented by a tuple $G \equiv (\mathcal{C}, \mathcal{A})$, where $\mathcal{C}$ denotes a non-empty set of unique genes, and $\mathcal{A}$ is a set of *adjacencies*. Genes are represented by their *extremities*, i.e., a gene $g \equiv (g^{\mathrm{t}}, g^{\mathrm{h}})$, $g \in \mathcal{C}$, consists of a *head* g^{h} and a *tail* g^{t}. Telomeres are modeled explicitly, as special genes of $\mathcal{C}(G)$ with a single extremity, denoted by "$\circ$". Extremities g_1^a, g_2^b, $a, b \in \{\mathrm{h}, \mathrm{t}\}$ of any two genes g_1, g_2 form an *adjacency* $\{g_1^a, g_2^b\}$ if they are immediate neighbors in their genome sequence. In the following, we will conveniently use the notation $\mathcal{C}(G)$ and $\mathcal{A}(G)$ to denote the set of genes and the set of adjacencies of genome G, respectively. We indicate the presence of an adjacency $\{x_1^a, x_2^b\}$ in an extant genome X by

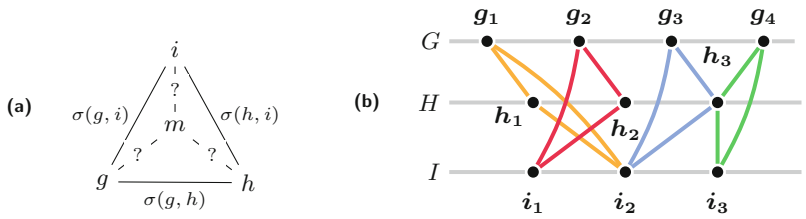

Fig. 1. (a) Illustration of the score of a candidate median gene. (b) Gene similarity graph of three genomes G, H, and I. Colored components indicate candidate median genes $m_1 = (g_1, h_1, i_2)$, $m_2 = (g_2, h_2, i_1)$, $m_3 = (g_3, h_3, i_2)$, and $m_4 = (g_4, h_3, i_3)$. Median gene pairs m_1, m_3 and m_3, m_4 are conflicting.

$$\mathbb{I}_X(x_1^a, x_2^b) = \begin{cases} 1 & \text{if } \{x_1^a, x_2^b\} \in \mathcal{A}(X) \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Given two genomes G and H and gene similarity measure σ, two adjacencies, $\{g_1^a, g_2^b\} \in \mathcal{A}(G)$ and $\{h_1^a, h_2^b\} \in \mathcal{A}(H)$ with $a, b \in \{\text{h, t}\}$ are *conserved* iff $\sigma(g_1, h_1) > 0$ and $\sigma(g_2, h_2) > 0$. We subsequently define the *adjacency score* of any four extremities g^a, h^b, i^c, j^d, where $a, b, c, d \in \{\text{h, t}\}$ and $g, h, i, j \in \Sigma$ as the geometric mean of their corresponding gene similarities [3]:

$$s(g^a, h^b, i^c, j^d) \equiv \sqrt{\sigma(g, h) \cdot \sigma(i, j)} \tag{2}$$

Median genome, genes and adjacencies. Informally, the family-free median problem asks for a fourth genome M that maximizes the sum of pairwise adjacency scores to three given extant genomes G, H, and I. In doing so, the gene content of the requested median M must first be defined: each gene $m \in \mathcal{C}(M)$ must be unambiguously associated with a triple of extant genes (g, h, i), $g \in \mathcal{C}(G)$, $h \in \mathcal{C}(H)$, and $i \in \mathcal{C}(I)$. Moreover, we want to associate to a median gene m a sequence similarity score (g, h, i) relatively to the three extant genes it is related to. As the sequence of the median gene is obviously not available, we define this score as the geometric mean of their pairwise similarities (see Fig. 1 (a)):

$$\sigma(g, m) = \sigma(h, m) = \sigma(i, m) \equiv \sqrt[3]{\sigma(g, h) \cdot \sigma(g, i) \cdot \sigma(h, i)} \tag{3}$$

In the following we make use of mapping $\pi_G(m) \equiv g$, $\pi_H(m) \equiv h$, and $\pi_I(m) \equiv i$ to relate gene m with its extant counterparts. Two candidate median genes or telomeres m_1 and m_2 are *conflicting* if $m_1 \neq m_2$ and the intersection between associated gene sets $\{\pi_G(m_1), \pi_H(m_1), \pi_I(m_1)\}$ and $\{\pi_G(m_2), \pi_H(m_2), \pi_I(m_2)\}$ is non-empty (see Fig. 1 (b) for example). A set of candidate median genes or telomeres $\mathcal{C}$ is called *conflict-free* if no two of its members $m_1, m_2 \in \mathcal{C}$ are conflicting. This definition trivially extends to the notion of a *conflict-free* median.

Problem 1 (FF-Median). Given three genomes G, H, and I, and gene similarity measure σ, find a conflict-free median M, which maximizes the following formula:

$$\mathcal{F}_\lambda(M) = \sum_{\{m_1^a, m_2^b\} \in \mathcal{A}(M)} \sum_{\substack{X \in \{G,H,I\}, \\ \{\pi_X(m_1)^a, \pi_X(m_2)^b\} \in \mathcal{A}(X)}} s(m_1^a, \pi_X(m_1)^a, m_2^b, \pi_X(m_2)^b), \tag{4}$$

where $a, b \in \{\mathrm{h}, \mathrm{t}\}$ and $s(\cdot)$ is the adjacency score as defined by Eq. (2).

Remark 1. The adjacency score for a median adjacency $\{m_1^a, m_2^b\}$ with respect to the corresponding potential extant adjacency $\{\pi_X(m_1)^a, \pi_X(m_2)^b\}$, where $\{m_1^a, m_2^b\} \in \mathcal{A}(M)$ and $X \in \{G, H, I\}$, can be entirely expressed in terms of pairwise similarities between genes of extant genomes using Eq. (3):

$$s(m_1^a, \pi_X(m_1)^a, m_2^b, \pi_X(m_2)^b) = \sqrt[6]{\prod_{\{Y,Z\} \subset \{G,H,I\}} \sigma(\pi_Y(m_1), \pi_Z(m_1)) \cdot \sigma(\pi_Y(m_2), \pi_Z(m_2))}$$

In the following, a median gene m and its extant counterparts (g, h, i) are treated as equivalent. We denote the set of all *candidate median genes* by

$$\Sigma_\lambda = \{(g, h, i) \mid g \in \mathcal{C}(G), h \in \mathcal{C}(H), i \in \mathcal{C}(I) : \sigma(g, h) \cdot \sigma(g, i) \cdot \sigma(h, i) > 0\}. \tag{5}$$

Each pair of median genes $(g_1, h_1, i_1), (g_2, h_2, i_2) \in \Sigma_\lambda$ and extremities $a, b \in \{\mathrm{h}, \mathrm{t}\}$ give rise to a *candidate median adjacency* $\{(g_1^a, h_1^a, i_1^a), (g_2^b, h_2^b, i_2^b)\}$ if $(g_1^a, h_1^a, i_1^a) \neq (g_2^b, h_2^b, i_2^b)$, and (g_1^a, h_1^a, i_1^a) and (g_2^b, h_2^b, i_2^b) are non-conflicting. We denote the set of all candidate median adjacencies and the set of all *conserved* (*i.e.* present in at least one extant genome) candidate median adjacencies by $\mathcal{A}_\lambda = \{\{m_1^a, m_2^b\} \mid m_1, m_2 \in \Sigma_\lambda,\ a, b \in \{\mathrm{h}, \mathrm{t}\}\}$ and $\mathcal{A}_\lambda^C = \{\ \{m_1^a, m_2^b\} \in \mathcal{A}_\lambda \mid \sum_{X \in \{G,H,I\}} \mathbb{I}_X(\pi_X(m_1)^a, \pi_X(m_2)^b) \geq 1\}$, respectively.

Remark 2. A median gene can only belong to a median adjacency with non-zero adjacency score if all pairwise similarities of its corresponding extant genes g, h, i are non-zero. Thus, the search for median genes can be limited to 3-cliques (triangles) in the tripartite similarity graph.

Remark 3. The right-hand side of the above formula for the weight of an adjacency is independent of genome X. From Eq. (4), an adjacency in median M has only an impact in a solution to problem FF-Median if it participates in a gene adjacency in at least one extant genome. So including in a median genome median genes that do not belong to a candidate median adjacency in $\mathcal{A}_\lambda^C$ do not increase the objective function.

Related problems. The FF-median problem relates to previously studied gene order evolution problems. It is a generalization of the tractable mixed multichromosomal median problem introduced in [16], that can indeed be defined as an FF-median problem with a similarity graph composed of disjoint 3-cliques and edges having all the same weight. The FF-median problem also bears similarity with methods aimed at detecting groups of orthologous genes based on gene order evolution, especially the MultiMSOAR [15] algorithm, although other method integrate synteny and sequence conservation for inferring orthogroups, see [5].

Our approach differs first and foremost in its family-free principle (all other methods require a prior gene family assignment). Compared to MultiMSOAR, the only other method that can handle more than two genomes with an optimization criterion that considers gene order evolution, both MultiMSOAR (for three genomes) and FF-median aim at computing a maximum weight tripartite matching. However we differ fundamentally from MultiMSOAR by the full integration of sequence and synteny conservation into the objective function, while MultiMSOAR proceeds first by computing pairwise orthology assignments to define a multipartite graph.

3 Algorithmic and Complexity Results

We now describe our theoretical results: a NP-hardness proof, an exact Integer Linear Program (ILP), and an algorithm to detect local optimal structures.

Theorem 1. *Problem FF-Median is MAX SNP-hard.*

We describe the full hardness proof in Appendix A. It is based on a reduction from the Maximum Independent Set for Graphs of Bounded Degree 3.

An exact ILP algorithm to problem FF-Median. We now present program `FF-Median`, described by Algorithm 1, that exploits the specific properties of problem FF-Median to design an ILP using $\mathcal{O}(n^5)$ variables and statements. Program `FF-Median` makes use of two types of binary variables **a** and **b** as declared in domain specifications (`D.01`) and (`D.02`), that defines the set of median genes Σ_λ and of candidate conserved median adjacencies $\mathcal{A}_\lambda^C$ (Remark 3). The former variable type indicates the presence or absence of candidate genes in an optimal median M. The latter, variable type **b**, specifies if an adjacency between two gene extremities or telomeres is established in M. Constraint (`C.01`) ensures that M is conflict-free, by demanding that each extant gene (or telomere) can be associated with at most one median gene (or telomere). Further, constraint (`C.02`) dictates that a median adjacency can only be established between genes that both are part of the median. Lastly, constraint (`C.03`) guarantees that each gene extremity and telomere of the median participates in at most one adjacency.

Property 1. The size (i.e. number of variables and statements) of any ILP returned by program `FF-Median` is limited by $\mathcal{O}(n^5)$ where $n = \max(|\mathcal{C}(G)|, |\mathcal{C}(H)|, |\mathcal{C}(I)|)$.

Remark 4. The output of the algorithm `FF-Median` is a set of adjacencies between median genes that define a set of linear and/or circular orders, called CARs (Contiguous Ancestral Regions), where linear segments are not capped by telomeres. So formally the computed median might not be a valid genome. However, as adding adjacencies that do not belong to $\mathcal{A}_\lambda^C$ do not modify the score of a given median, a set of median adjacencies can always be completed into

Algorithm 1. Program `FF-Median` for three genomes (G, H, I)

Objective: `Maximize`

$$\sum_{\substack{\{m_1,m_2\}\in\mathcal{A}_\lambda^C,\\ m_1=(g_1,h_1,i_1),\\ m_2=(g_2,h_2,i_2),\\ a,b\in\{\mathrm{h},\,\mathrm{t}\}}} \mathbf{b}(g_1^a, g_2^b, h_1^a, h_2^b, i_1^a, i_2^b) \sum_{\substack{X\in\{G,H,I\},\\ \{\pi_X(m_1)^a,\pi_X(m_2)^b\}\in\mathcal{A}(X)}} s(m_1^a, \pi_X(m_1)^a, m_2^b, \pi_X(m_2)^b)$$

Constraints:

(`C.01`) $\forall\ g' \in \mathcal{C}(G)$: $\sum_{(g',h,i)\in\Sigma_\lambda} \mathbf{a}(g',h,i) \leq 1$

$\forall\ h' \in \mathcal{C}(H)$: $\sum_{(g,h',i)\in\Sigma_\lambda} \mathbf{a}(g,h',i) \leq 1$

$\forall\ i' \in \mathcal{C}(I)$: $\sum_{(g,h,i')\in\Sigma_\lambda} \mathbf{a}(g,h,i') \leq 1$

(`C.02`) $\forall\ \{(g_1,h_1,i_1),(g_2,h_2,i_2)\} \in \mathcal{A}_\lambda^C$ and $\forall\ a,b \in \{\mathrm{h},\,\mathrm{t}\}$:

$$2 \cdot \mathbf{b}(g_1^a, g_2^b, h_1^a, h_2^b, i_1^a, i_2^b) \leq \mathbf{a}(g_1,h_1,i_1) + \mathbf{a}(g_2,h_2,i_2)$$

(`C.03`) $\forall\ (g_1,h_1,i_1) \in \Sigma_\lambda$ and $\forall\ a \in \{\mathrm{h},\,\mathrm{t}\}$:

$$\sum_{(g_2,h_2,i_2)\in\Sigma_\lambda,\ b\in\{\mathrm{h},\,\mathrm{t}\}} \mathbf{b}(g_1^a, g_2^b, h_1^a, h_2^b, i_1^a, i_2^b) \leq 1$$

Domains:

(`D.01`) $\forall\ (g,h,i) \in \Sigma_\lambda$: $\mathbf{a}(g,h,i) \in \{0,1\}$

(`D.02`) $\forall\ \{(g_1,h_1,i_1),(g_2,h_2,i_2)\} \in \mathcal{A}_\lambda^C$ and $\forall\ a,b \in \{\mathrm{h},\,\mathrm{t}\}$:

$$\mathbf{b}(g_1^a, g_2^b, h_1^a, h_2^b, i_1^a, i_2^b) \in \{0,1\}$$

a valid genome by such adjacencies that join the linear segments together and add telomeres. These extra adjacencies would not be supported by any extant genome and thus can be considered as dubious, and in our implementation, we only return the median adjacencies computed by the ILP, *i.e.* a subset of $\mathcal{A}_\lambda^C$.

Remark 5. Following Remark 2, preprocessing the input extant genomes requires to handle the extant genes that do not belong to at least one 3-clique in the similarity graph. Such genes can not be part of any median. So one could decide to leave them in the input, and the ILP can handle them and ensures they are never part of the output solution. However, discarding them from the extant genomes can help recover adjacencies that have been disrupted by the insertion of a mobile element for example, so in our implementation we follow this approach.

As discussed at the end of Sect. 2, the FF-median problem is a generalization of the mixed multichromosomal breakpoint median [16]. However, it was shown in [16] that this breakpoint median problem can be solved in polynomial time by a Maximum-Weight Matching (MWM) algorithm. This motivates the results presented in the next paragraph that use a MWM algorithm to identify optimal median substructures by focusing on conflict-free sets of median genes.

Finding local optimal segments. Tannier *et al.* [16] solve the mixed multichromosomal breakpoint median problem by transforming it into an MWM problem, that we outline now. A graph is defined in which each extremity of a candidate median gene and each telomere gives rise to a vertex. Any two vertices are connected by an edge, weighted according to the number of observed adjacencies between the two gene extremities in extant genomes. Edges corresponding to adjacencies between a gene extremity and telomeres are weighted only by half as much. An MWM in this graph induces a set of adjacencies that defines an optimal median.

We first describe how this approach applies to our problem. We define a graph $\Gamma(\Sigma_\curlywedge)$ constructed from an FF-Median instance (G, H, I, σ) that is similar to that of Tannier *et al.*, deviating by defining vertices as candidate median gene extremities and weighting an edge between two vertices m_1^a, m_2^b, $a, b \in \{\text{h, t}\}$, by

$$w(\{m_1^a, m_2^b\}) = \sum_{X \in \{G,H,I\}} \mathbb{I}_X(\pi_X(m_1)^a, \pi_X(m_2)^b) \cdot s(m_1^a, \pi_X(m_1)^a, m_2^b, \pi_X(m_2)^b). \tag{6}$$

We make first the following observation, where a conflict-free matching is a matching that does not contain two conflicting vertices (candidate median genes):

Observation 2. *Any conflict-free matching in graph $\Gamma(\Sigma_\curlywedge)$ of maximum weight defines an optimal median.*

We show now that we can define notions of sub-instances – of a full FF-median instance – that contains no internal conflicts, for which applying the MWM can allow to detect if the set of median genes defining the sub-instance is part of at least one optimal FF-median. Let $\mathcal{S}$ be a set of candidate median genes. An *internal conflict* is a conflict between two genes from $\mathcal{S}$; an *external conflict* is a conflict between a gene from $\mathcal{S}$ and a candidate median gene not in $\mathcal{S}$. We say that $\mathcal{S}$ is *contiguous* in extant genome X if the set $\pi_X(\mathcal{S})$ forms a unique, contiguous, segment in X. We say that $\mathcal{S}$ is an *internal-conflict free segment* (IC-free segment) if it contains no internal conflict and is contiguous in all three extant genomes; this can be seen as the family-free equivalent of the notion of *common interval in permutations* [2]. An IC-free segment is a *run* if the order of the extant genes is conserved in all three extant genomes, up to a full reversal of the segment.

Intuitively, one can find an optimal solution to the sub-instance defined by an IC-free segment, but it might not be part of an optimal median for the whole instance due to side effects of the rest of the instance. So we need to adapt the graph to which we apply an MWM algorithm to account for such side effects. To do so, we define the *potential* of a candidate median gene m as

$$\Delta(m) = \max_{\{m_1^a, m^b\}, \{m^a, m_2^b\} \in \mathcal{A}_\curlywedge} \big(w(\{m_1^a, m^b\}) + w(\{m^a, m_2^b\})\big).$$

We then extend graph $\Gamma(\mathcal{S}) =: (V, E)$ to graph $\Gamma'(\mathcal{S}) := (V, E')$ by adding edges between the extremities of each candidate median gene of an IC-free segment $\mathcal{S}$,

i.e. $E' = E \cup \{\{m^{\mathrm{h}}, m^{\mathrm{t}}\} \mid m \in \mathcal{S}\}$ (note that when $|\mathcal{S}| > 1$, $w(\{m^{\mathrm{h}}, m^{\mathrm{t}}\}) = 0$ since $\mathcal{S}$ is contiguous in all three extant genomes). In the following we refer to these edges as *conflict edges*. Let $C(m)$ be the set of candidate median genes that are involved in an (external) conflict with a given candidate median gene m of $\mathcal{S}$, then the conflict edge $\{m^{\mathrm{h}}, m^{\mathrm{t}}\} \in E'$ is weighted by the maximum potential of a non-conflicting subset of $C(m)$,

$$w'(\{m^{\mathrm{h}}, m^{\mathrm{t}}\}) = \max(\{\sum_{m' \in C'} \Delta(m') \mid C' \subseteq C(m) : \ C' \text{ is conflict-free}\}) .$$

A conflict-free matching in $\Gamma'(\mathcal{S})$ is a matching without a conflict edge.

Lemma 1. *Given an internal conflict-free segment $\mathcal{S}$, any maximum weight matching in graph $\Gamma'(\mathcal{S})$ that is conflict-free defines a set of median genes and adjacencies that belong to at least one optimal FF-median of the whole instance.*

A proof is presented in Appendix B. Lemma 1 leads to a procedure (Algorithm 2) that iteratively identifies and tests IC-free segments in the FF-Median instance. For each identified IC-free segment S an adjacency graph $\Gamma'(S)$ is constructed and a maximum weight matching is computed (Line 2–3). If the resulting matching is conflict free (Line 4), adjacencies of IC-free segment S are reported and S is removed from an FF-Median instance by masking its internal adjacencies and removing all candidate median genes (and consequently their associated candidate median adjacencies) corresponding to external conflicts (Line 5–6). It then follows immediately from Lemma 1 that the set median genes returned by Algorithm 2 belongs to at least one optimal solution to the FF-median problem.

Algorithm 2. Algorithm `ICF-SEG`

Input: FF-Median instance (G, H, I, σ)
Output: Set of adjacencies $\textsc{Adj}_M$ that is part of a median M of (G, H, I, σ).

1: **while** there exists an unobserved IC-free conserved segment S in (G, H, I, σ) **do**
2: Construct adjacency graph $\Gamma'(S)$ of S
3: Find maximum weight matching $\mathcal{M} \subseteq E(\Gamma'(S))$
4: **if** $A(S) = \mathcal{M}$ **then**
5: Add $A(S)$ to $\textsc{Adj}_{\mathcal{M}}$
6: Remove S including external conflicts from (G, H, I, σ)
7: **end if**
8: **end while**

In the experiments, IC-free runs are used instead of segments. Step 1 is performed efficiently by first identifying maximal IC-free runs, then breaking it down into smaller runs whenever the condition in Step 4 is not satisfied.

4 Experimental Results and Discussion

Our algorithms have been implemented in Python and require CPLEX[1]; they are freely available as part of the family-free genome comparison tool `FFGC` downloadable at http://bibiserv.cebitec.uni-bielefeld.de/ffgc.

[1] http://www.ibm.com/software/integration/optimization/cplex-optimizer/.

In subsequent analyses, gene similarities are based on local alignment hits identified with BLASTP on protein sequences using an e-value threshold of 10^{-5}. In gene similarity graphs, we discard spurious edges by applying a *stringency filter* proposed by Lechner *et al.* [12] that utilizes a local threshold parameter $f \in [0, 1]$ and BLAST bit-scores: a BLAST hit from a gene g to h is only retained if it is has a higher or equal score than f times the best BLAST hit from h to any gene g' that is member of the same genome as g. In all our experiments, we set f to 0.5. Edge weights of the gene similarity graph are then calculated according to the *relative reciprocal BLAST score* (RRBS) [14]. Finally we applied Algorithm ICF-SEG with conserved segments defined as runs.

For solving the FF-Median problem, we granted CPLEX two CPU cores, 4 GB memory, and a time limit of 3 h per dataset.

In our experiments, we compare ourselves against the orthology prediction tool MultiMSOAR [15]. This tool requires precomputed gene families, which we constructed by following the workflow described in [15].

Evaluation on simulated data. We first evaluate our algorithms on simulated data sets obtained by ALF [4]. The ALF simulator covers many aspects of genome evolution from point mutations to global modifications. The latter includes two types of genome rearrangements, as well as various options to customize the process of gene family evolution. In our simulations, we mainly use standard parameters suggested by the authors of ALF and we focus on three parameters that primarily influence the outcome of gene family-free genome analysis: (i) the rate of sequence evolution, (ii) the rate of genome rearrangements, and (iii) the rate of gene duplications and losses. We keep all three rates constant, only varying the evolutionary distance between the generated extant genomes. We confine our simulations to protein coding sequences. A comprehensive list of parameter settings used in our simulations is shown in Table 2 in Appendix C. As root genome in the simulations, we used the genomic sequence of an *E. coli* K-12 strain (Accession no: `NC_000913.2`) which comprises 4,320 protein coding genes. We then generated 7×10 data sets with increasing evolutionary distance ranging from 10 to 130 *percent accepted mutations* (PAM). Details about the generated data sets are shown in Table 1 in Appendix C. Figure 2(a) shows the outcome of our analysis with respect to precision and recall[2] of inferring positional orthologs. In all simulations, FF-Median generated no or very few false positives, leading to perfect or near-perfect precision score, consistently outperforming MultiMSOAR. However, since the objective of FF-Median only takes median genes into account that are conserved by synteny, the increase in mutational changes over evolutionary time causes a growing loss of syntenic context which results in a lower recall. Therefore, MultiMSOAR retains a better recall for larger evolutionary distances, while FF-Median provides better results for more closely related genomes.

Evaluation on real data. We study 15 γ-proteobacterial genomes that span a large taxonomic spectrum and are contained in the OMA database [1]. A complete list

[2] precision: #true positives/(#true positives + #false positives), recall: #true positives/(#true positives + #false negatives).

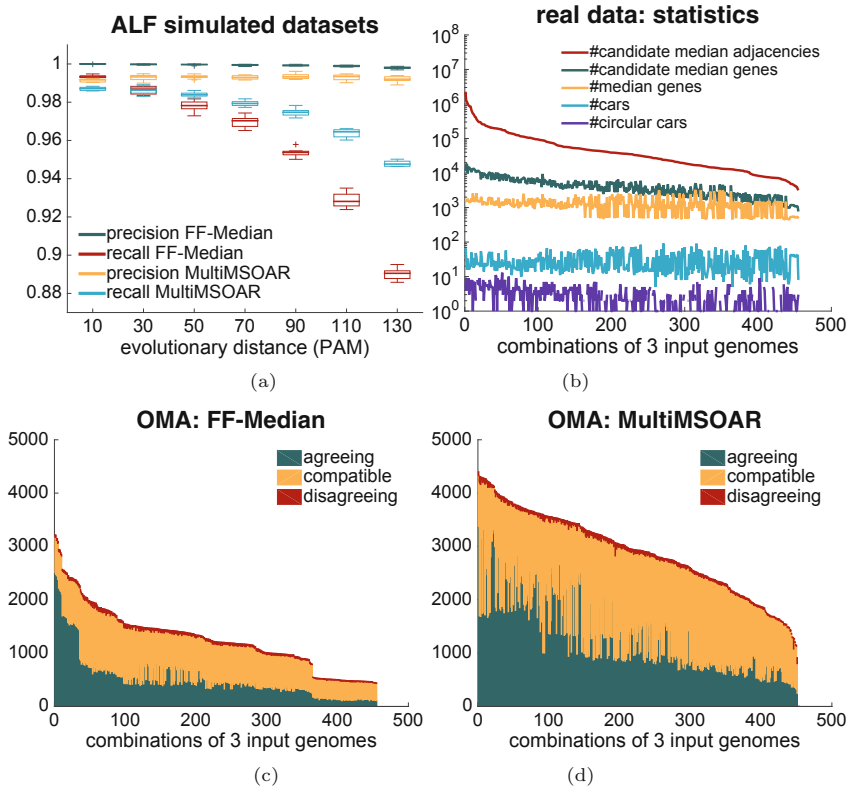

Fig. 2. Top: (a) Precision and recall of FF-Median and MultiMSOAR in simulations; (b) statistical assessment of CARs and median genes on real datasets. Bottom: agreement, compatibility and disagreement of positional orthologs inferred by (c) FF-Median and (d) MultiMSOAR with OMA database. (Color figure online)

of species names is given in Appendix D. We obtained the genomic sequences from the NCBI database and constructed for each combination of three genomes a gene similarity graph following the same procedure as in the simulated dataset. In 9 out of the 455 combinations of genomes the time limit prohibited CPLEX from finding an optimal solution. However, in those cases CPLEX was still able to find integer feasible suboptimal solutions. Figure 2(b) displays statistics of the real dataset. The number of candidate median genes and adjacencies ranges from 442 to 18,043 and 3,164 to 2,261,716, respectively, giving rise to up to 3,227 median genes that are distributed on 5 to 91 CARs per median. Some CARs are circular, indicating dubious conformations mostly arising from tandem duplications, but the number of such cases were low (mean: 2.78, max: 13).

We observed that the gene families in the OMA database are clustered tightly and therefore missing many true orthologies in the considered triples of genomes. As a result, many of the orthologous groups inferred by FF-Median and MultiMSOAR fall into more than one gene family inferred by OMA. We therefore

evaluate our results by classifying the inferred orthologous groups into three categories: An orthologous group *agrees* with OMA if its three genes are in the same OMA group. It *disagrees* with OMA if extant genes x and y (of genomes X and Y respectively) are in different OMA groups but the OMA group of x contains another gene from genome Y. It is *compatible* with OMA if it neither agrees nor disagrees with OMA. We measure the number of median genes as well as the number orthologous groups of MultiMSOAR in each of the three categories. Figure 2(c) and (d) shows the outcome this analysis. MultiMSOAR is generally able to find more orthology relations in the dataset. This comes at no surprise, as it is clear from the objective of problem FF-Median and from the results of the simulated datasets that our method does not retain candidate median genes which have lost their syntenic context, which happens in triples of highly divergent genomes. The number of disagreeing orthologous groups that disagree with OMA is comparably low for both FF-Median (mean: 35.16, var: 348) and MultiMSOAR (mean: 48.61, var: 348).

We then performed another analysis to assess the *robustness* of the positional orthology predictions. To this end, we look at orthologous groups across multiple datasets that share two extant genomes, but vary in the third. Given two genes, x of genome X and y of genome Y, an orthologous group that contains x and y is called *robust* if x and y occur in the same orthologous group, whatever the third extant genome is. We computed the percentage of robust orthologous groups for all gene pairs of randomly-chosen genomes *E. coli K-12 MG 1655* and *S. enterica subsp. enterica serovar Typhimurium str. 14028s* in our dataset. The results indicate that orthologous groups inferred by FF-Median are slightly more robust (95.61 %) than robust those by MultiMSOAR (91.77 %). This is likely due to the strict constraint of defining median adjacencies only from genes that participate in at least one observed adjacency (Remark 4).

Overall, we can observe that FF-Median performed better than MultiMSOAR only for triples of closely related genomes – which is consistent with our observation on simulated data – while being slightly more robust in general. This suggests FF-Median is an interesting alternative to identify higher confidence positional orthologs, at the expense of a higher recall rate.

Future work. We first aim to investigate alternative methods to reduce the computational load of Program FF-Median by identifying further strictly suboptimal and optimal substructures, which might require a better understanding of the impact of internal conflicts within substructures defined by intervals in the extant genomes. Without the need to modify drastically either the FF-median problem definition or the ILP, one can think about more complex weighting schemes for adjacencies that could account for known divergence time between genomes or relaxed notion of adjacencies that would address the high recall rate we observe in FF-Median. Within that regard, it would probably be interesting to combine this with the use of common intervals instead of runs to define conflict-free sub-instances. Finally, ideal family-free analysis should take into account the effects of gene family evolution. However, the presented family-free median model can only resolve certain cases of gene duplication. It is generally susceptible to gene losses that occurred along the evolutionary paths between

the three extant genomes and their common ancestor. The definition of a family-free median model that tolerates events of gene family evolution at a reasonable computational cost is likely an interesting research avenue.

Acknowledgments. The authors are grateful to the anonymous reviewers for their valuable comments. DD and MB wish to thank Bernard M.E. Moret for helpful discussions. CC acknowledges funding from the NSERC Discovery Grant.

References

1. Altenhoff, A.M., Skunca, N., Glover, N., Train, C.-M., Sueki, A., Pilizota, I., Gori, K., Tomiczek, B., Müller, S., Redestig, H., Gonnet, G.H., Dessimoz, C.: The OMA orthology database in 2015: function predictions, better plant support, synteny view and other improvements. Nucleic Acids Res. **43**, 240–249 (2015)
2. Bergeron, A., Chauve, C., Gingras, Y.: Formal models of gene clusters. In: Bioinformatics Algorithms: Techniques and Applications, pp. 177–202. Wiley, New York (2008)
3. Braga, M.D.V., Chauve, C., Doerr, D., Jahn, K., Stoye, J., Thévenin, A., Wittler, R.: The potential of family-free genome comparison. In: Models and Algorithms for Genome Evolution, pp. 287–323. Springer, London (2013)
4. Dalquen, D.A., Anisimova, M., Gonnet, G.H., Dessimoz, C.: Alf - a simulation framework for genome evolution. Mol. Biol. Evol. **29**(4), 1115–1123 (2012)
5. Dewey, C.N.: Positional orthology: putting genomic evolutionary relationships into context. Brief Bioinform. **12**(5), 401–412 (2011)
6. Doerr, D., Thévenin, A., Stoye, J.: Gene family assignment-free comparative genomics. BMC Bioinform. **13**(Suppl. 19), S3 (2012)
7. Dörr, D.: Gene family-free genome comparison. Ph.D. thesis, Universität Bielefeld, Bielefeld, Germany (2016)
8. Fitch, W.M.: Homology a personal view on some of the problems. Trends Genet. **16**, 227–231 (2000)
9. Gabaldón, T., Koonin, E.V.: Functional and evolutionary implications of gene orthology. Nat. Rev. Genet. **14**, 360–366 (2013)
10. Kowada, L.A.B., Doerr, D., Dantas, S., Stoye, J.: New genome similarity measures based on conserved gene adjacencies. In: RECOMB **2016**, pp. 204–224 (2016)
11. Lechner, M., Hernandez-Rosales, M., Doerr, D., Wieseke, N., Thévenin, A., Stoye, J., Hartmann, R.K., Prohaska, S.J., Stadler, P.F.: Orthology detection combining clustering and synteny for very large datasets. PLoS ONE **9**(8), e105015 (2014)
12. Lechner, M., Findeiß, S., Steiner, L., Marz, M., Stadler, P.F., Prohaska, S.J.: Proteinortho: detection of (co-)orthologs in large-scale analysis. BMC Bioinform. **12**, 124 (2011)
13. Martinez, F.V., Feijão, P., Braga, M.D.V., Stoye, J.: On the family-free DCJ distance and similarity. Algorithms Mol. Biol. **10**, 13 (2015)
14. Pesquita, C., Faria, D., Bastos, H., Ferreira, A.E.N., Falcão, A.O., Couto, F.M.: Metrics for GO based protein semantic similarity: a systematic evaluation. BMC Bioinform. **9**(Suppl. 5), S4 (2008)
15. Shi, G., Peng, M.-C., Jiang, T.: Multimsoar 2.0: an accurate tool to identify ortholog groups among multiple genomes. PLoS ONE **6**(6), e20892 (2011)
16. Tannier, E., Zheng, C., Sankoff, D.: Multichromosomal median and halving problems under different genomic distances. BMC Bioinform. **10**, 120 (2009)

A Hardness Proof

A.1 Reduction

The *maximum independent set problem for graphs bounded by node degree* 3, denoted as MAX IS-3 is MAX SNP-hard [3]. The corresponding decision problem can be informally stated as follows: Given a graph Λ bounded by degree 3 and some number $l \geq 1$, does there exists a set of vertices $V' \subseteq V$ of size $|V'| = l$ whose induced subgraph is unconnected? In the following, we present a transformation scheme $\mathbf{R}$ to phrase Λ as FF-median instance $\mathbf{R}(\Lambda) = (G, H, I, \sigma)$ such that the value $\mathcal{F}_{\curlywedge}(M)$ of a median M of $\mathbf{R}(\Lambda)$ is limited by $\mathcal{F}_{\curlywedge}(M) \leq 2 \cdot l + 3$. In doing so, we associate vertices of V with genes of extant genomes G, H and I. In order to keep track of associated genes, we denote by function $\xi(x)$ the list of vertices associated with gene x. We further introduce two types of unassociated genes, "$\emptyset$" and "$*$", whose members are identified by subscript notation.

Transformation $\mathbf{R}$:

1. Construct genome G such that for each vertex $v \in V$ there exists two associated genes $g_v, \bar{g}_v \in \mathcal{C}(G)$, i.e. $\xi(g_v) = \xi(\bar{g}_v) = v$. Further, let each gene pair $g_v, \bar{g}_v$ form a circular chromosome, giving rise to adjacency set $\mathcal{A}(G) = \{\{\bar{g}_v^{\mathrm{h}}, g_v^{\mathrm{t}}\}, \{\bar{g}_v^{\mathrm{h}}, g_v^{\mathrm{t}}\} \mid v \in V,\ g_v, \bar{g}_v \in \mathcal{C}(G)\}$.
2. For each edge $(u, v) \in E$ construct a circular chromosome $\mathcal{X}_{uv}$ hosting two genes $x_{uv}, x_\emptyset \in \mathcal{C}(\mathcal{X}_{uv})$, with gene x_{uv} being associated with both vertices u and v and gene $x_\emptyset$ being unassociated. Further, let both genes form a circular chromosome, giving rise to adjacency set $\mathcal{A}(\mathcal{X}_{uv}) = \{\{x_{uv}^{\mathrm{h}}, x_\emptyset^{\mathrm{t}}\}, \{x_\emptyset^{\mathrm{h}}, x_{uv}^{\mathrm{t}}\}\}$.
3. Assign each chromosome constructed in the previous step either to genome H or to genome I such that each vertex $v \in V$ is associated with at most two genes per genome.
4. Complete genomes H and I with additional circular chromosomes $\mathcal{X}_v$ where $\mathcal{C}(\mathcal{X}_v) = \{x_v, x_\emptyset\}$ and $\mathcal{A}(\mathcal{X}_v) = \{\{x_v^{\mathrm{h}}, x_\emptyset^{\mathrm{t}}\}, \{x_\emptyset^{\mathrm{h}}, x_v^{\mathrm{t}}\}\}$ such that each vertex in V is associated with exactly two genes per genome.
5. For each vertex $v \in V$, let $g, \bar{g} \in \mathcal{C}(G)$, $h, \bar{h} \in \mathcal{C}(H)$, and $i, \bar{i} \in \mathcal{C}(I)$ be the pairs of genes associated with v, i.e. $\xi(g) = \xi(\bar{g}) = \xi(h) \cap \xi(i) = \xi(\bar{h}) \cap \xi(\bar{i}) = v$. Assign gene similarities $\sigma(g, h) = \sigma(g, i) = \sigma(h, i) = 1$ and $\sigma(\bar{g}, \bar{h}) = \sigma(\bar{g}, \bar{i}) = \sigma(\bar{h}, \bar{i}) = 1$.
6. Add a copy of circular chromosome $\mathcal{X}_*$ to each genome G, H, and I, where $\mathcal{C}(\mathcal{X}_*) = \{x_*, \bar{x}_*\}$ and $\mathcal{A}(\mathcal{X}_*) = \{\{x_*^{\mathrm{h}}, \bar{x}_*^{\mathrm{t}}\}, \{\bar{x}_*^{\mathrm{h}}, x_*^{\mathrm{t}}\}\}$. Let $g_*, \bar{g}_* \in \mathcal{C}(G), h_*, \bar{h}_* \in \mathcal{C}(H)$, and $i_*, \bar{i}_* \in \mathcal{C}(I)$, set the gene similarity score between all pairs of genes in $\{g_*, h_*, i_*\}$ and $\{\bar{g}_*, \bar{h}_*, \bar{i}_*\}$ respectively, to 1. Lastly, set the gene similarity score of all pairs of unassociated genes of type "$\emptyset$" including genes $g_*, \bar{g}_*$ to $\frac{1}{4}$.

Except for step 3, none of the instructions of transformation scheme $\mathbf{R}$ are computationally challenging. Note that in step 3 the demanded partitioning of chromosomes into genomes H and I is always possible as consequence of Vizing's Theorem [4], by which every graph with maximum node degree d is

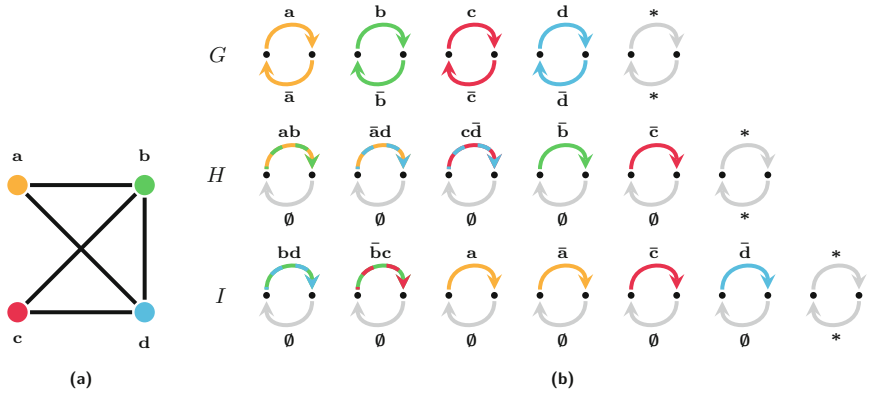

Fig. 3. (a) A simple graph bounded by degree three and (b) a corresponding FF-Median instance constructed with transformation scheme **R**.

edge-colorable using at most d or $d+1$ colors. Hence, using colors $\chi_1, \chi_2, \chi_3, \chi_4$ with $\chi_1 = \chi_2 \equiv I$, $\chi_3 = \chi_4 \equiv H$ and Misra and Gries' algorithm [2], edges of graph $\Lambda = (E, V)$ can be partitioned into two groups in $\mathcal{O}(|E||V|)$ time implying an assignment to genomes H and I.

Example 1. Fig. 3(b) shows a FF-Median instance constructed with transformation scheme **R** from the simple graph depicted in Fig. 3(a). Gene similarities between genes are not shown, but can be derived from the genes' labeling.

We structure our proof that the presented transformation is in fact a valid mapping of an MAX IS-3 instance to an instance of FF-Median into three different lemmas:

Lemma 2. *Given a median M of FF-Median instance $\mathbf{R}(\Lambda) = (G, H, I, \sigma)$, (1) for each median gene $(g, h, i) \in \mathcal{C}(M)$ where g, h, or i are associated with vertices in $V(\Lambda)$ holds $\xi(g) = \xi(h) \cap \xi(i) = v$, $v \in V(\Lambda)$; (2) there exist at most two median genes whose corresponding extant genes are not associated to any vertex in $V(\Lambda)$.*

Proof. Assume for contradiction that claim (1) does not hold. Then either $\xi(g) \neq \xi(h) \cap \xi(i)$, or $\xi(h) \cap \xi(i) = \emptyset$, both of which violate the constraint of establishing gene similarities between associated genes that is given in step 5. For claim (2), observe that the only unassociated genes in genome G are gene g_* and $\bar{g}_*$ introduced in step 6, limiting the overall number of unassociated genes in any median M. □

Lemma 3. *The conserved adjacency set of any median M of FF-Median instance $\mathbf{R}(\Lambda) = (G, H, I, \sigma)$ is of the form $\mathcal{A}(M) \cap \mathcal{A}_\lambda^C = \mathcal{A}_\lambda^G(M) \cup \{\{m_*^{\mathrm{h}}, \overline{m}_*^{\mathrm{t}}\}, \{\overline{m}_*^{\mathrm{h}}, m_*^{\mathrm{t}}\}\}$, where the extant genes corresponding to m_* and $\overline{m}_*$ are all unassociated genes of type "$*$" and $\mathcal{A}(M)_\lambda^G \subseteq \left\{\{m_1^{\mathrm{h}}, m_2^{\mathrm{t}}\} \in A_\lambda^C \mid \xi(\pi_G(m_1)) = \xi(\pi_G(m_2))\right\}$.*

Proof. Observe that both candidate median adjacencies $a_* = \{m_*^{\mathrm{h}}, \overline{m}_*^{\mathrm{t}}\}$ and $\bar{a}_* = \{\overline{m}_*^{\mathrm{h}}, m_*^{\mathrm{t}}\}$ are conserved in all three genomes, whereas all other conserved candidate adjacencies between associated and unassociated genes can be at most conserved in H and I. Establishing adjacencies $a_*, \bar{a}_*$ gives rise to a cumulative adjacency score of 6. Conversely, up to 4 non-conflicting adjacencies between associated and unassociated genes can be established that are conserved in both genomes H and I. However, since such adjacencies are only conserved between unassociated genes of type "$\emptyset$" whose gene similarities are set to $\frac{1}{4}$, the best cumulative adjacency score can not exceed 4. Thus, adjacencies $a_*, \bar{a}_*$ must be contained in any median. Further, because of this and the fact that in both genomes H and I, each gene associated with vertices of $V(\Lambda)$ is only adjacent to an unassociated gene, M cannot contain adjacencies that are conserved in extant genomes other than genome G, which are the adjacencies of each gene pair $(g_v, \bar{g}_v)$ associated with the same vertex $v \in V(\Lambda)$. □

Lemma 4. *Given FF-median instance* $\mathbf{R}(\Lambda) = (G, H, I, \sigma)$*, let* m_u, m_v *be any pair of candidate median adjacencies of* $\mathcal{A}_\curlywedge$ *whose corresponding extant genes are associated to vertices* $u, v \in V(\Lambda)$*, then* m_u, m_v *are conflicting if and only if* $(u, v) \in E$*.*

Proof. By construction in step 5 of transformation scheme $\mathbf{R}$, each vertex $v \in V$ is associated with exactly two candidate median genes $m_v = (g, h, i), \overline{m}_v = (\bar{g}, \bar{h}, \bar{i})$, $m_v, \overline{m}_v \in \Sigma_\curlywedge$, such that $\xi(g) = \xi(h) \cap \xi(i) = v$ and $\xi(\bar{g}) = \xi(\bar{h}) \cap \xi(\bar{i}) = v$. Further, let u be another vertex of $V(\Lambda)$, such that $(u, v) \in E(\Lambda)$, and $m_u, \overline{m}_u$ are its two corresponding candidate median genes. Then, by construction in step 2, there exists exactly one extant gene x with $\xi(x) = uv$ (which, by assignment in step 3, is either contained in genome H or I). Consequently, either m_u is in conflict with m_v, or $\overline{m}_u$ with $\overline{m}_v$, or $\overline{m}_u$ with m_v, or m_u with $\overline{m}_v$. Recall that by construction in step 2 in $\mathbf{R}$ and by Lemma 3, $m_u, \overline{m}_u$ and $m_v, \overline{m}_v$ form conserved candidate adjacencies $\{m_u^{\mathrm{h}}, \overline{m}_u^{\mathrm{t}}\}$, $\{\overline{m}_u^{\mathrm{h}}, m_u^{\mathrm{t}}\}$ and $\{m_v^{\mathrm{h}}, \overline{m}_v^{\mathrm{t}}\}$, $\{\overline{m}_v^{\mathrm{h}}, m_v^{\mathrm{t}}\}$, respectively. Clearly, independent of which of the candidate median gene pairs of u and v are in conflict, both pairs of candidate median adjacencies are in conflict with each other.

Now, let u, v be two vertices of $V(\Lambda)$ such that edge $(u, v) \notin E(\Lambda)$, then there exists no gene x in extant genomes H and I with $\xi(x) = uv$. Even more, due to Lemma 2, there cannot exist a candidate median gene (g, h, i) with $\{u, v\} \subseteq \xi(g) \cup \xi(h) \cup \xi(i)$. Thus, the candidate median genes of u and v are not conflicting and neither are their corresponding candidate median adjacencies. □

We proceed to show that the given transformation scheme gives rise to an approximation preserving reduction known as *L-reduction*. An L-reduction reduces a problem P to a problem Q by means of two polynomial-time computable transformation functions: A function $f : P \to Q' \subseteq Q$ that maps each instance of P onto an instance of Q, herein represented by transformation scheme $\mathbf{R}$, and a function $g : Q' \to P$ to transform any feasible solution of an instance

in Q' to a feasible solution of an instance of P. Here, a *feasible* solution means any – not necessarily *optimal* – solution that obeys the problem's constraints. A feasible solution of FF-Median instance (G, H, I, σ) is an *ancestral genome* X where $\mathcal{C}(X) \subseteq \Sigma_\lambda$ and $\mathcal{A}(X) \subseteq \mathcal{A}_\lambda$ such that $\mathcal{A}(X)$ is conflict-free. We give the following transformation scheme to map ancestral genomes of an FF-Median instance to solutions of an MAX IS-3 instance:

Transformation $\mathbf{S}$: Given any ancestral genome X of $\mathbf{R}(\Lambda)$, return $\{\xi(\pi_G(m_1))|$ $\{m_1^a, m_2^b\} \in \mathcal{A}(X) : \mathbb{I}_G(\pi_G(m_1)^a, \pi_G(m_2)^b) = 1$ and $\xi(\pi_G(m_1)) \neq \emptyset\}$.

We define score function $s_\lambda(X) \equiv \frac{1}{2}\mathcal{F}_\lambda(X) - 3$ of an ancestral genome X. For $(\mathbf{R}, \mathbf{S})$ to be an L-reduction the following two properties must hold for any given MAX IS-3 instance (Λ, l): (1) There is some constant α such that for any median M of the transformed FF-Median instance $\mathbf{R}(\Lambda)$ holds $s_\lambda(M) \leq \alpha \cdot l$; (2) There is some constant β such that for any ancestral genome X of $\mathbf{R}(\Lambda)$ holds $l - |\mathbf{S}(X)| \leq \beta \cdot |s_\lambda(M) - s_\lambda(X)|$. We proceed to proof the following lemma:

Lemma 5. $(\mathbf{R}, \mathbf{S})$ *is an L-reduction of problem MAX IS-3 to problem FF-Median with* $\alpha = \beta = 1$.

Proof. For any median M of FF-Median instance $\mathbf{R}(\Lambda)$, the number of conserved median adjacencies with correspondence to the same vertex of Λ is two, giving rise a cumulative adjacency score of two. From Lemmata 3 and 4 immediately follows that any ancestral genome of $\mathbf{R}(\Lambda)$ that maximizes the number of conserved adjacencies also maximizes the number of independent vertices in Λ. Recall that the two conserved adjacencies between unassociated genes of type "$*$" (which are part of all medians) give rise to a cumulative adjacency score of 6, we conclude that $|\mathcal{A}(M) \cap \mathcal{A}_\lambda^C| - 2 = \frac{1}{2}\mathcal{F}_\lambda(M) - 3 = s_\lambda(M) = l$, thus $\alpha = 1$.

Because $l = s_\lambda(M)$, it remains to show that $l - |S(X)| \leq \beta|l - s_\lambda(X)|$. In a *sub-optimal* ancestral genome of $\mathbf{R}(\Lambda)$, median genes with no association to vertices of Λ can also contain extant genes of type "$\emptyset$". These unassociated median genes can form "mixed" conserved adjacencies with genes that are associated with vertices of Λ. Such mixed conserved adjacencies have no correspondence to vertices in Λ and do not contribute to the transformed solution $\mathbf{S}(X)$ of an ancestral genome X. Yet, as mentioned earlier, the cumulative adjacency score of all mixed conserved adjacencies can not not exceed 4. Therefore it holds that $|S(X)| \geq s_\lambda(X)$ and we conclude $\beta = 1$. □

B Speeding up the Search for a Median

Proof of Lemma 1:

Proof. Given an IC-free segment $\mathcal{S} = \{m_1, \ldots, m_k\}$ of an FF-Median instance (G, H, I, σ). Let M be a conflict-free matching in graph $\Gamma'(\mathcal{S})$. Because M is conflict-free and $\mathcal{S}$ contiguous in all three extant genomes, M must contain all candidate median genes of S. Now, let M' be a median such that $\mathcal{S} \not\subseteq \mathcal{C}(M')$.

Further, let $C(m)$ be the set of candidate median genes that are involved in a conflict with with a given median gene m of $\mathcal{S}$ and $X = \mathcal{C}(M') \cap (\bigcup_{m \in \mathcal{S}} C(m) \cup \mathcal{S})$. Clearly, $X \neq \emptyset$ and for the contribution $\mathcal{F}_\lambda(X)$ must hold $\mathcal{F}_\lambda(X) \geq \mathcal{F}_\lambda(\mathcal{S})$, otherwise M' is not optimal since it is straightforward to construct a median higher score which includes $\mathcal{S}$. Clearly, the contribution $\mathcal{F}(X)$ to the median is bounded by $\max(\{\sum_{m' \in C'} \Delta(m') \mid C' \subseteq C(m) : \ C' \text{is conflict-free}\}) + \mathcal{F}_\lambda(\mathcal{S})$. But since $\mathcal{S}$ gives rise to a conflict-free matching with maximum score, also median M'' with $\mathcal{C}(M'') = (\mathcal{C}(M') \setminus X) \cup \mathcal{C}(\mathcal{S})$ and $\mathcal{A}(M'') = (\mathcal{A}(M') \setminus \mathcal{A}(X)) \cup \mathcal{A}(S))$ must be an (optimal) median. □

C Simulated Sequence Evolution with ALF

Table 1. Average benchmark data of seven evolutionary distances, each comprising ten genomic datasets generated by ALF [1].

PAM	Genome	Inversions	Transpositions	Duplications	Losses
10	*G*	8.7	6.1	7.3	6.9
	H	7.3	4.5	6.3	5.4
	I	8.5	6.6	10.4	5.6
30	*G*	24.5	16.9	21.0	22.7
	H	23.4	19.8	20.6	18.4
	I	25.5	17.2	17.5	20.9
50	*G*	39.9	27.8	32.4	36.7
	H	41.8	31.8	31.0	31.7
	I	43.2	30.0	28.7	39.7
70	*G*	58.6	42.3	41.1	39.2
	H	57.0	43.6	46.3	45.1
	I	60.4	41.4	40.7	39.1
90	*G*	75.0	54.5	53.1	64.2
	H	69.9	50.5	54.1	65.0
	I	75.2	55.5	60.3	58.5
110	*G*	96.3	69.4	67.0	74.6
	H	90.6	64.2	62.5	70.9
	I	90.2	68.5	62.6	61.2
130	*G*	105.7	76.3	74.4	81.0
	H	108.7	78.2	79.6	82.8
	I	110.8	73.6	73.9	77.3

Table 2. Parameter settings for simulations generated by ALF [1].

Parameter name	Value	
Sequence evolution		
Substitution model	WAG (amino acid substitution model)	
Insertion and deletion	Zipfian distribution	exponent $c = 1.8214$
	Insertion rate	0.0003
	Maximum insertion length	50
Rate variation among sites	Γ-distribution	shape parameter $a = 1$
	Number of classes	5
	Rate of invariable sites	0.01
Genome rearrangement		
Inversion	rate	0.0004
	Maximum inversion length	100
Transposition	rate	0.0002
	Maximum transposition length	100
	Rate of inverted transposition	0.1
Gene family evolution		
Gene duplication	Rate	0.0001
	Max. no. of genes involved in one dupl.	5
	Probability of transposition after dupl.	0.5
	Fission/fusion after duplication	0.1
	Probability of rate change	0.2
	Rate change factor	0.9
	Probability of temporary rate change (duplicate)	0.5
	Temporary rate change factor (duplicate)	1.5
	Life of rate change (duplicate)	10 PAM
	Probability of temporary rate change (orig+duplicate)	0.3
	Temporary rate change factor (orig+duplicate)	1.2
	Life of rate change (orig+duplicate)	10 PAM
Gene loss	Rate	0.0001
	Maximum length of gene loss	5
Gene fission/fusion	rate	0.0
	Maximum number of fused genes	–

D Real Genomes Dataset

See Table 3.

Table 3. Dataset of genomes used in comparison with the OMA database.

Genbank ID	Name
U00096.3	Escherichia coli str. K-12 substr. MG1655, complete genome
AE004439.1	Pasteurella multocida subsp. multocida str. Pm70, complete genome
AE016853.1	Pseudomonas syringae pv. tomato str. DC3000, complete genome
AM039952.1	Xanthomonas campestris pv. vesicatoria complete genome
CP000266.1	Shigella flexneri 5 str. 8401, complete genome
CP000305.1	Yersinia pestis Nepal516, complete genome
CP000569.1	Actinobacillus pleuropneumoniae L20 serotype 5b complete genome
CP000744.1	Pseudomonas aeruginosa PA7, complete genome
CP000766.3	Rickettsia rickettsii str. Iowa, complete genome
CP000950.1	Yersinia pseudotuberculosis YPIII, complete genome
CP001120.1	Salmonella enterica subsp. enterica serovar Heidelberg str. SL476, complete genome
CP001172.1	Acinetobacter baumannii AB307-0294, complete genome
CP001363.1	Salmonella enterica subsp. enterica serovar Typhimurium str. 14028S, complete genome
FM180568.1	Escherichia coli 0127:H6 E2348/69 complete genome, strain E2348/69
CP002086.1	Nitrosococcus watsoni C-113, complete genome

References

1. Dalquen, D.A., Anisimova, M., Gonnet, G.H., Dessimoz, C.: Alf - a simulation framework for genome evolution. Mol. Biol. Evol. **29**(4), 1115–1123 (2012)
2. Misra, J., Gries, D.: A constructive proof of Vizing's theorem. Inform. Process. Lett. **41**(3), 131–133 (1992)
3. Papadimitriou, C.H., Yannakakis, M.: Optimization, approximation, and complexity classes. J. Comput. Syst. Sci. **43**(3), 425–440 (1991)
4. Vizing, V.G.: On an estimate of the chromatic class of a p-graph. Diskret. Analiz **3**, 25–30 (1964)