

1 A PROOF OF INVARIANTS

2 A.1 At least one of S or O is empty

3 When $BUILDINCA''$ is first called, O has no splits, so the invariant holds.

4 The $REMOVEIRRELEVANTSPLITS$ step can alter S and O in two ways. First, if O contains a solution S'
5 with the same taxon set as S , then we can replace S with S' and clear O . Second, if one of the sub-solutions
6 in O is punctured, we remove it from O and add its child solutions to O , but S remains empty. In both of
7 these cases, the invariant is preserved. Further steps inside $BUILDINCA''$ do not modify S or O .

8 When $BUILDINCA''$ is recursively called on an unmodified component, O will be empty. When
9 $BUILDINCA''$ is called on a modified component or a new component, S will be empty. So, this invariant
10 is preserved in recursive calls to $BUILDINCA''$.

11 Therefore, since the invariant is initially true, and is preserved within $BUILDINCA''$ and at recursive
12 calls, it holds by induction.

13 A.2 Original solutions in O have a taxon set inside $S.T$

14 O is empty when $BUILDINCA''$ is first called. When $C.O$ is initially filled by merging two components, all
15 the original solutions in $C.O$ are for original components that are subsets of the new component. Therefore,
16 when calling $BUILDINCA''(C.S, C.\Delta\Sigma, C.O)$, all the solutions in $O=C.O$ will have a taxon set inside $S.T$.

17 When $REMOVEIRRELEVANTSPLITS$ replaces a punctured solution with its child solutions, the child
18 solutions remain inside the taxon set of the current problem, since they are contained within their parent
19 taxon set, which was inside the taxon set of the current problem. QED.

20 A.3 Original solutions in O have non-overlapping taxon sets

21 O is empty when $BUILDINCA''$ is first called. When O is initially filled by merging a component, all the
22 solutions in O must have non-overlapping taxon sets, since they are solutions for different equivalence
23 classes on the same level.

24 When $REMOVEIRRELEVANTSPLITS$ replaces a punctured solution by its child solutions, it preserves
25 this property, since the taxon sets of the child solutions are non-overlapping, and are contained within the
26 taxon set of their removed parent. Thus the solutions in O will always have non-overlapping taxon sets.

27 B PROPERTIES OF SOLUTION S

28 B.1 Extracting the tree from a Solution

29 If the full $BUILD$ algorithm succeeds, a tree can be extracted from the $Solution$ object. A $Solution$
30 object S maps to a node n in the tree. The node's children consist of either (i) internal nodes (one for each
31 component in $S.\mathcal{C}$) or (ii) leaves (one for each trivial component; *i.e.* taxon identifier for which $S.M$ is a
32 $NULL$ reference). The original $Solution$ object created in $BUILD$ corresponds to the root of the tree.

33 B.2 Extracting the splits Σ from a Solution

34 Each $Solution$ object S is created by performing $BUILDA$ on a set of splits Σ . If $BUILDA(S, \Sigma)$
35 succeeds, then after it is complete S will contain the splits Σ . However, only the splits $S.I$ are *directly*
36 contained in S . The remaining splits of Σ that are not stored in $S.I$ are contained in the child solutions of
37 S (and their descendants).

We can recover the set of splits in S as follows:

$$\Sigma(S) = S.I \cup \left(\bigcup_{C \in S.\mathcal{C}} \Sigma(C.S) \right). \quad (1)$$

38 A solution is considered "empty" if it contains no splits. This can only happen if both $S.I$ and $S.\mathcal{C}$ are
39 empty. It is possible for $S.\mathcal{C}$ to be empty while $S.I$ is non-empty, so both fields must be used.

40 B.3 The relationship between $S.T$ and $\Sigma(S)$

Additionally, if S is a solution to some component C , then the taxon set $S.T$ is determined by the splits
 $\Sigma(S)$:

$$S.T = \bigcup_{\sigma \in \Sigma(S)} \sigma_1 \quad (2)$$

41 Additionally the taxa in $S.T$ form a connected component that is connected *only* through edges correspond-
42 ing to splits in $\Sigma(S)$.

43 The only exception to equation (2) is the top level `Solution`. The top level taxon set is directly
44 initialized by the user and does not correspond to a component that was identified by running `BUILDA`, so
45 there may be taxa in T that are not mentioned in Σ . These taxa will connect directly to the root since they
46 are not in the include group of any split. However, for `Solution` objects below the top level, equation
47 (2) holds.