Formalization of Lexicographic Breadth-First Search

Thomas Holme Surlykke and Dmitriy Traytel

November 26, 2026

Abstract

Despite widespread use of the lexicographic breadth-first search graph ordering algorithm, the literature lacks a complete proof characterizing its output. We verify that the standard characterization provides a sound and complete specification. Our approach introduces a binary reformulation of the label function, exposing intermediate states and enabling proof through monadic refinement.

Contents

1 Introduction **2**

2 Preliminaries **3**
2.1 Triples . . . 3
2.2 Lists . . . 3

3 Implementation and Characterization **4**

4 Soundness **4**
4.1 The binary algorithm and its soundness . . . 5
4.1.1 Preservation of invariants . . . 6
4.2 Soundness of *lexbfs-bin* . . . 7
4.3 *lexbfs* refines *lexbfs-bin* . . . 7
4.4 Soundness of *lexbfs* . . . 9

5 Completeness **9**
5.1 Completeness of *lexbfs-bin* . . . 9
5.1.1 Reconstructing the state of *lexbfs-bin* . . . 9
5.1.2 Preservation of the constructed state . . . 11
5.2 *lexbfs-bin* refines *lexbfs* . . . 12
5.3 Completeness of *lexbfs* . . . 12

6 Correctness **12**

1 Introduction

Lexicographic breadth-first search (LexBFS) is a graph vertex ordering algorithm implementable in linear time [5]. Rose, Tarjan, and Lueker [6] developed it to decompose graphs into slices according to vertex distance from an initial vertex. They use it to find *minimal orderings* and *perfect orderings* of undirected graphs. The algorithm has since found applications in *Chordal Graphs* and *Graph Coloring* [5].

We call any complete list of vertices in a graph a *LexBFS ordering* if it is a possible output of the LexBFS algorithm. As the algorithm selects an arbitrary maximal vertex when multiple exist, different implementations may produce different orderings. Therefore, it is a non-trivial task to characterize the structure and behavior of LexBFS orderings. The Wikipedia page *Lexicographic breadth-first search* [5] provides a mathematical characterization of all LexBFS orderings:

> **Theorem** (LexBFS Property [5])**.** *Let $G = (V, E)$ be a graph with n vertices. Recall that $N(v)$ is the set of neighbors of v. Let $\sigma = (v_1, \ldots, v_n)$ be an enumeration of the vertices of V. The enumeration σ is a LexBFS ordering (with source v_1) if, for all $1 \leq i < j < k \leq n$ with $v_i \in N(v_k) - N(v_j)$, there exists $m < i$ such that $v_m \in N(v_j) - N(v_k)$.*

Wikipedia does not reference a proof, but references Corneil [2], which also describes the LexBFS algorithm and states a variant of the property:

> **Theorem** (LexBFS Property [2])**.** *An ordering $\prec$ of the vertices of an arbitrary graph $G(V, E)$ is an LexBFS ordering if and only if for all vertices a, b, c of G such that $ac \in E$ and $bc \notin E$, $c \prec b \prec a$ implies the existence of a vertex d in G, adjacent to b but not to a and such that $d \prec c$.*

Corneil [2] also provides no proof, but instead references [1], which defines the property as:

(P1) If $a < b < c$ and $ac \in E$ and $ab \notin E$ then there exists a vertex d such that $c < d$, $db \in E$ and $da \notin E$.

and states:

> **Theorem** (LexBFS Property [1])**.**
>
> *(1) Any LexBFS-ordering has property (P1).*
>
> *(2) Any ordering fulfilling (P1) can be generated by LexBFS.*

This paper does prove (2), but for (1) they merely state:

> *Proof.* (1) We refer to the well known proof in [21]. □

Their reference [21] is Golumbic [4], who lists yet another variant of the property as property (L3):

> **Theorem** (LexBFS Property (L3) [4])**.** *if* $\sigma^{-1}(a) < \sigma^{-1}(b) < \sigma^{-1}(c)$ *and* $c \in \mathrm{Adj}(a) - \mathrm{Adj}(b)$*, then there exists a vertex* $d \in \mathrm{Adj}(b) - \mathrm{Adj}(a)$ *with* $\sigma^{-1}(c) < \sigma^{-1}(d)$*.*

Again, despite what was written in [1], no proof is listed. The "well-known" proof is merely an exercise [4]:

> Exercises
> ...
> 3. Prove properties (L1)–(L3) of lexicographic breadth-first search (Section 4.3).

This formalization is motivated by the absence of a published proof, combined with the inconsistent notation across sources. We verify that the LexBFS Property is both sound and complete as a specification of the LexBFS algorithm: Any LexBFS ordering satisfies the property (soundness), and any ordering satisfying the property can be produced by LexBFS (completeness).

2 Preliminaries

2.1 Triples

Three indices meaning c picked before b, picked before a.

abbreviation *triple* σ *ci bi ai* $\equiv$ *ci* $<$ *bi* $\wedge$ *bi* $<$ *ai* $\wedge$ *ai* $<$ *length* σ

2.2 Lists

lemma *nth-not-mem-take-if-ge*: $[\![i \leq j; j < \mathit{length}\ xs; \mathit{distinct}\ xs]\!] \Longrightarrow xs\ !\ j \notin \mathit{set}\ (\mathit{take}\ i\ xs)$
⟨*proof*⟩

We compare lists lexicographically.

lemma *list-lt-iff-ind-lt*:
fixes *xs*::('*c*::*linorder list*)
assumes *len-eq*: *length xs* = *length ys*
shows *xs* < *ys* $\longleftrightarrow$ ($\exists$ *i* < *length xs*. *xs* ! *i* < *ys* ! *i* $\wedge$ ($\forall$ *j*<*i* . *xs* ! *j* = *ys* ! *j*))
⟨*proof*⟩

lemma *prev-lt-if-ind-gt*:
  fixes *xs* :: *'c::linorder list*
  assumes *length xs* = *length ys*
    and *i* < *length xs*
    and *xs* ≤ *ys*
    and *xs*!*i* > *ys*!*i*
  shows ∃ *j* < *i* . *xs*!*j* < *ys*!*j*
  ⟨*proof*⟩

context *Graph*
begin

abbreviation *is-arc* :: *'a* ⇒ *'a* ⇒ *bool* (**infixl** ∼ *60*) **where** *v*∼*w* ≡ (*v*,*w*) ∈ *E*
abbreviation *is-not-arc* :: *'a* ⇒ *'a* ⇒ *bool* (**infixl** ∼∗ *60*) **where** *v* ∼∗ *w* ≡ ¬ *v*∼*w*

3 Implementation and Characterization

The LexBFS algorithm, adapted as closely as possible from Algorithm 1 of Corneil et al. [3]:

definition *lexbfs* **where** *lexbfs* = *do* {
  (*i*, *U*, *L*, *σ*) ← *WHILE*$_T$ (λ(*i*, *U*, *L*, *σ*). *U* ≠ {}) (λ(*i*, *U*, *L*, *σ*). *do* {
    x <− *RES* {*x* ∈ *U*. *L x* = *Max* (*L* ‘ *U*)};
    let U = *U* − {*x*};
    let L = (λ*y*. *if y* ∈ *U* ∧ *x* ∼ *y then L y* @ [*i*] *else L y*);
    RETURN (*i* − *1*, *U*, *L*, *σ* @ [*x*])
}) (*card V*, *V*, λ - . [], []);
  RETURN σ
}

The LexBFS property.

definition *lexbfs-prop σ* ≡
  ∀ *ic ib ia*. *triple σ ic ib ia* ∧ *σ*!*ic* ∼ *σ*!*ia* ∧ *σ*!*ic* ∼∗ *σ*!*ib*
    ⟶ (∃ *id*<*ic*. *σ*!*id* ∼∗ *σ*!*ia* ∧ *σ*!*id* ∼ *σ*!*ib*)

For *lexbfs-prop* to be an actual characterization, we need the additional assumption that the elements of σ correspond precisely to the vertices of V.

definition *is-lexbfs σ* ⟷ *distinct σ* ∧ *set σ* = *V* ∧ *lexbfs-prop σ*

4 Soundness

In this section we prove soundness of the LexBFS algorithm with respect to the specification *is-lexbfs*. In other words, we show *lexbfs* ≤ *SPEC is-lexbfs*.

4.1 The binary algorithm and its soundness

The algorithm with $\mathcal{L}$ replaced by the binary label function $\mathcal{B}$. We will, however, define the algorithm's invariants before defining the algorithm.

Convention: We use $\mathcal{B}$ for the binary label function, and $\mathcal{L}$ for the ordinary label function.

definition *partition-V U σ* ≡ *distinct σ* ∧ *U* ∪ *set σ* = *V* ∧ *U* ∩ *set σ* = {}
lemma *partition-V-empty*[*simp*]: *partition-V V* []
⟨*proof*⟩

lemma *U-finite-if-partition-V*: *partition-V U σ* ⟹ *finite U*
⟨*proof*⟩

lemma *len-sigma-if-partition-V*: *partition-V U σ* ⟹ *length σ* = *card* (*set σ*)
⟨*proof*⟩

lemma *card-V-split-if-partition-V*: *partition-V U σ* ⟹ *card V* = *card U* + *length σ*
⟨*proof*⟩

This states that, if we look at any element $y \in \sigma$ and x later than y (either to the right of y in σ, or not in σ at all), then x must have a less-or-equal $\mathcal{B}$-label value, if we truncate $\mathcal{B}$ to the point where y was added.

definition *label-le-earlier* $\mathcal{B}$ *σ* ≡
∀ *iy*<*length σ*. ∀ *x*∈*V* − *set* (*take iy σ*). *take iy* ($\mathcal{B}$ (*σ* ! *iy*)) ≥ *take iy* ($\mathcal{B}$ *x*)

lemma *label-le-earlierD*[*simp*]:
⟦*label-le-earlier* $\mathcal{B}$ *σ*; *iy* < *length σ*; *x* ∈ *V* − *set* (*take iy σ*)⟧
⟹ *take iy* ($\mathcal{B}$ (*σ*!*iy*)) ≥ *take iy* ($\mathcal{B}$ *x*)
⟨*proof*⟩

States that each entry of $\mathcal{B}$ x determines whether a particular vertex in σ is connected to x or not.

definition *label-iff-connected* $\mathcal{B}$ *σ* ≡ ∀ *ix*<*length σ*. ∀ *y*. ($\mathcal{B}$ *y*)!*ix* ⟷ (*σ*!*ix*)∼*y*
lemma *label-iff-connectedD*[*simp*]:
⟦*label-iff-connected* $\mathcal{B}$ *σ*; *iy* < *length σ*⟧ ⟹ $\mathcal{B}$ *x*!*iy* ⟷ (*σ*!*iy*)∼*x*
⟨*proof*⟩

definition *label-len-eq* $\mathcal{B}$ *σ* ≡ ∀ *x* . *length* ($\mathcal{B}$ *x*) = *length σ*

This simp rule chooses *length σ* as simplest representation, and helps a lot in automating the proofs.

lemma *len-eq-sigma-if-label-len-eq*[*simp*]: *label-len-eq* $\mathcal{B}$ *σ* ⟹ *length* ($\mathcal{B}$ *x*) = *length σ*
⟨*proof*⟩

lemma *len-take-eq-if-label-len-eq*:

label-len-eq $\mathcal{B}$ $\sigma \Longrightarrow$ *length* (*take* i ($\mathcal{B}$ x)) = *length* (*take* i ($\mathcal{B}$ y))
⟨*proof*⟩

Combining the invariants.

definition *lexbfs-bin-INV* ≡ λ(U, $\mathcal{B}$, σ) .
partition-V U σ
∧ *label-le-earlier* $\mathcal{B}$ σ
∧ *label-iff-connected* $\mathcal{B}$ σ
∧ *label-len-eq* $\mathcal{B}$ σ

The refine WHILE procedure requires writing *lexbfs-bin-INV* as a function on a tuple. This unfortunately turns unfolding *lexbfs-bin-INV* into a case statement with a single case. This lemma automatically unfolds *lexbfs-bin-INV* into the auxiliary invariants.

lemma *INV-iff*[*iff*]:
lexbfs-bin-INV (U, $\mathcal{B}$, σ)
⟷ *partition-V* U σ ∧ *label-le-earlier* $\mathcal{B}$ σ ∧ *label-iff-connected* $\mathcal{B}$ σ ∧ *label-len-eq* $\mathcal{B}$ σ
⟨*proof*⟩

Finally, the binary version of the LexBFS algorithm.

definition *lexbfs-bin* **where** *lexbfs-bin* = *do* {
(U, $\mathcal{B}$, σ) ← *WHILE*$_T$ *lexbfs-bin-INV* (λ(U, $\mathcal{B}$, σ). $U \neq \{\}$) (λ(U, $\mathcal{B}$, σ). *do* {
x <− *RES* {$x \in U$. $\mathcal{B}$ x = *Max* ($\mathcal{B}$ ‘ U)};
let U = U − {x};
let $\mathcal{B}$ = (λy. $\mathcal{B}$ y @ [x∼y]);
RETURN (U, $\mathcal{B}$, σ @ [x])
}) (V, λ - . [], []);
RETURN σ
}

4.1.1 Preservation of invariants

We start with the main point of the invariants – if the invariants hold for the terminal case (where $U = \{\}$), then the output satisfies *is-lexbfs*. The remaining work will then be showing that the invariants are preserved.

lemma *lexbfs-prop-if-INV*:
assumes *eq-V*: *partition-V* {} σ
and *aux1*: *label-le-earlier* $\mathcal{B}$ σ
and *aux2*: *label-iff-connected* $\mathcal{B}$ σ
and *aux3*: *label-len-eq* $\mathcal{B}$ σ
shows *lexbfs-prop* σ
⟨*proof*⟩

lemma *is-lexbfs-if-INV*:
lexbfs-bin-INV ({}, $\mathcal{B}$, σ) ⟹ *is-lexbfs* σ

⟨*proof*⟩

The following lemmas show that if an auxiliary invariant holds, then it also holds during the next step of the while loop.

lemma *partition-V-next*:
⟦*partition-V U σ*; *x* ∈ *U*⟧ ⟹ *partition-V* (*U* − {*x*}) (*σ*@[*x*])
⟨*proof*⟩

lemma *label-le-earlier-next*:
assumes *partition*: *partition-V U σ*
and *aux1*: *label-le-earlier* ℬ *σ*
and *aux3*: *label-len-eq* ℬ *σ*
and ℬ *a* = *Max* (ℬ ‘ *U*)
shows *label-le-earlier* (λ*y*. ℬ *y* @[*a*∼*y*]) (*σ*@[*a*])
⟨*proof*⟩

lemma *label-iff-connected-next*:
assumes *aux2*: *label-iff-connected* ℬ *σ* **and** *aux3*: *label-len-eq* ℬ *σ*
shows *label-iff-connected* (λ*y*. ℬ *y* @[*a*∼*y*]) (*σ*@[*a*])
⟨*proof*⟩

lemma *label-len-eq-next*:
label-len-eq ℬ *σ* ⟹ *label-len-eq* (λ*y*. ℬ *y* @ [*x*∼*y*]) (*σ*@[*x*])
⟨*proof*⟩

4.2 Soundness of *lexbfs-bin*

proposition *lexbfs-bin-sound*: *lexbfs-bin* ≤ *SPEC is-lexbfs*
⟨*proof*⟩

4.3 *lexbfs* refines *lexbfs-bin*

For ease, we use an intermediate algorithm, where each label is also updated outside of *U*.

definition *lexbfs-helper* **where** *lexbfs-helper* = *do* {
(*i*, *U*, ℒ, *σ*) ← *WHILE*$_T$ (λ(*i*, *U*, ℒ, *σ*). *U* ≠ {}) (λ(*i*, *U*, ℒ, *σ*). *do* {
x <− *RES* {*x* ∈ *U*. ℒ *x* = *Max* (ℒ ‘ *U*)};
let U = *U* − {*x*};
— Also updates vertices already checked
let ℒ = (λ*y*. *if x* ∼ *y then* ℒ *y* @ [*i*] *else* ℒ *y*);
RETURN (*i* − *1*, *U*, ℒ, *σ* @ [*x*])
}) (*card V*, *V*, λ - . [], []);
RETURN σ
}

Transforms the list of Booleans by placing *V*-index on *True*, and removing *False*.

fun *transform-bin* :: *nat* ⇒ *bool list* ⇒ *nat list* **where**
transform-bin - [] = [] |
transform-bin n (*False*#*xs*) = *transform-bin* (*n*−*1*) *xs*|
transform-bin n (*True*#*xs*) = *n* # *transform-bin* (*n*−*1*) *xs*

lemma *transform-bin-split*[*simp*]: *transform-bin n* (*xs* @ *ys*) = *transform-bin n xs* @ *transform-bin* (*n* − *length xs*) *ys*
⟨*proof*⟩

lemma *hd-transform-bin-le*[*dest*]: *transform-bin n xs* = *a* # *as* ⟹ *a* ≤ *n*
⟨*proof*⟩

A rephrasing of *transform-bin* (*Suc n*) (*x* # *xs*) and *transform-bin* (*Suc n*) (*y* # *ys*), when *x* < *y*, i.e. *x* = *False* and *y* = *True*.

lemma *transform-bin-decreasing*[*simp*]: *transform-bin n xs* < *Suc n* # *transform-bin n ys*
⟨*proof*⟩

proposition *transform-bin-le-iff*:
⟦*length xs* = *length ys*; *length ys* ≤ *n*⟧ ⟹ *xs* ≤ *ys* ⟷ *transform-bin n xs* ≤ *transform-bin n ys*
⟨*proof*⟩

Maps the binary label function $\mathcal{B}$ to the ordinary label function $\mathcal{L}$, except this $\mathcal{L}$ also updates value of vertices already processed.

abbreviation *transform-B* $\mathcal{B}$ *x* ≡ *transform-bin* (*card V*) ($\mathcal{B}$ *x*)

lemma *transform-B-refl-Max*: ⟦*partition-V U* σ; *label-len-eq* $\mathcal{B}$ σ; *x* ∈ *U*⟧
⟹ $\mathcal{B}$ *x* = *Max* ($\mathcal{B}$ ‘ *U*) ⟷ *transform-B* $\mathcal{B}$ *x* = *Max* (*transform-B* $\mathcal{B}$ ‘ *U*)
⟨*proof*⟩

Note: The *Id* allows us to specify what relation these algorithms are equal through. So the outer monad (that returns σ after while loop exits) is completely equal, while the WHILE loops are refined through the relation below.

abbreviation *R-helper-bin* ≡ {(*b*, (*U*, $\mathcal{B}$, σ)). *b* = (*card U*, *U*, *transform-B* $\mathcal{B}$, σ)}

lemma *lexbfs-helper-refines-lexbfs-bin*: *lexbfs-helper* ≤ ⇓ *Id lexbfs-bin*
⟨*proof*⟩

The relation that identifies *lexbfs-helper* and *lexbfs*. Note that this relation is symmetric, so we can reuse it later when we show *lexbfs-helper* ≤ *lexbfs*.

abbreviation *R-lexbfs-helper* ≡ {((*i*, *U*, $\mathcal{L}$, σ), (*i*′, *U*′, $\mathcal{L}$′, σ′)) .
i = *i*′ ∧ *U* = *U*′ ∧ σ = σ′ ∧ (∀ *x* ∈ *U*′ . $\mathcal{L}$ *x* = $\mathcal{L}$′ *x*)}

lemma *lexbfs-refines-lexbfs-helper*: *lexbfs* ≤ ⇓ *Id lexbfs-helper*

⟨*proof*⟩

4.4 Soundness of *lexbfs*

Combining the previous results, and using transitivity of refinement, we can now show that LexBFS is sound with regard to the specification *is-lexbfs*. In other words: Any list that can be produced by the *lexbfs* algorithm satisfies the *lexbfs-prop*.

theorem *lexbfs-sound*: *lexbfs* ≤ *SPEC is-lexbfs*
⟨*proof*⟩

5 Completeness

5.1 Completeness of *lexbfs-bin*

5.1.1 Reconstructing the state of *lexbfs-bin*

By definition of the binary label-function $\mathcal{B}$, and as indicated by the invariant *label-iff-connected*, it is possible and easy to obtain $\mathcal{B}$ from the currently picked vertices σ.

definition *construct-B* σ $y \equiv$ *map* $(\lambda x.\ x{\sim}y)$ σ

lemma *construct-B-empty*[*simp*]: *construct-B* [] = $(\lambda$ - . [])
⟨*proof*⟩

lemma *construct-B-tail*[*simp*]: *construct-B* $(\sigma@[x])$ y = *construct-B* σ y @ $[x{\sim}y]$
⟨*proof*⟩

The constructed $\mathcal{B}$ automatically satisfies two of our auxiliary invariants: *label-len-eq* and *label-iff-connected*.

lemma *construct-B-iff-connected*:
$i <$ *length* $\sigma \Longrightarrow$ *construct-B* σ a ! $i <$ *construct-B* σ b ! $i \longleftrightarrow \sigma!i \sim* a \wedge \sigma!i \sim b$
⟨*proof*⟩

lemma *construct-B-label-iff-connected*: *label-iff-connected* (*construct-B* σ) σ
⟨*proof*⟩

By making this a simp rule, we also easily obtain that *length* (*construct-B* σ x) = *length* (*construct-B* σ y), for any x, y.

lemma *construct-B-len-eq-sigma*[*simp*]: *length* (*construct-B* σ x) = *length* σ
⟨*proof*⟩

lemma *construct-B-len-eq*[*simp*]: *label-len-eq* (*construct-B* σ) σ
⟨*proof*⟩

Reconstruction of U given the currently picked vertices σ. We hide this behind a definition, as being able to control what is expanded through the simp rules below is very useful in the *refine-rcg* proofs, as those have a lot of assumptions and statements in them.

definition *U-at* $\sigma \equiv V - set\ \sigma$

lemma *U-at-finite*[*simp*]: *finite* (*U-at* σ)
⟨*proof*⟩

lemma *U-at-eq-rest*: *is-lexbfs* (σ @ σ-*rest*) $\Longrightarrow$ *U-at* σ = *set* σ-*rest*
⟨*proof*⟩

lemma *U-at-not-empty*: *is-lexbfs* (σ @ *a* # σ-*rest*) $\Longrightarrow$ *U-at* $\sigma \neq \{\}$
⟨*proof*⟩

lemma *U-at-empty*[*simp*]: *is-lexbfs* $\sigma \Longrightarrow$ *U-at* $\sigma = \{\}$
⟨*proof*⟩

lemma *U-at-empty-eq-V*[*simp*]: *U-at* [] = *V*
⟨*proof*⟩

lemma *U-at-next*: $x \in$ *U-at* $\sigma \Longrightarrow$ *U-at* (σ @ [x]) = *U-at* $\sigma - \{x\}$
⟨*proof*⟩

We also add an abbreviation for the reconstructed state.

abbreviation *bin-state-at* $\sigma \equiv$ (*U-at* σ, *construct-B* σ, σ)

By *candidate* we mean that it can be picked in a given step of the algorithm.

abbreviation *candidates* $\sigma \equiv \{x \in$ *U-at* σ . *construct-B* σ x = *Max* (*construct-B* σ ‘ *U-at* σ)$\}$

This shows that it is always possible to pick a next element to add to σ, i.e. that the algorithm can't fail.

lemma *obtains-candidate*: **assumes** *is-lexbfs* (σ @ *a* # σ-*rest*)
obtains x **where** $x \in$ *candidates* σ
⟨*proof*⟩

The main lemma of this section – if a full list satisfies the LexBFS property, then any element inside this list is a possible pick by the algorithm, if we reconstruct the state from the elements preceding it. This is the major step in showing completeness of the Binary LexBFS, i.e. that any list satisfying *is-lexbfs* can be produced by *lexbfs-bin*. The proof follows the mathematical proof from Brandstädt et al. [1].

lemma *candidate-if-is-lexbfs*:
assumes *is-lexbfs*: *is-lexbfs* (σ @ *a* # σ-*rest*) (**is** *is-lexbfs* ?σ-*all*)
shows *a* ∈ *candidates* σ
⟨*proof*⟩

5.1.2 Preservation of the constructed state

For this step, we want to talk about specific steps of the *binary* algorithm. In the following, we need to show that the reconstructed state is preserved at each step, i.e. that what we reconstruct corresponds to what the algorithm outputs. For this, it will be useful to be able to keep parts of the loop from unfolding, and enabling starting the algorithm at a separate step, hence the three sub-definitions in the following.

definition *bin-step* $\equiv \lambda\ (U,\ \mathcal{B},\ \sigma)\ .\ do\ \{$
  $x <-\ RES\ \{x \in U\ .\ \mathcal{B}\ x = Max\ (\mathcal{B}\ `\ U)\};$
  $let\ U = U - \{x\};$
  $let\ \mathcal{B} = (\lambda y\ .\ \mathcal{B}\ y\ @\ [x \sim y]);$
  $RETURN\ (U,\ \mathcal{B},\ \sigma\ @\ [x])$
$\}$

definition *bin-loop* $\equiv$ $WHILE_T\ (\lambda(U,\ \mathcal{B},\ \sigma)\ .\ U \neq \{\})$ *bin-step*

Note that, disregarding the added invariant, *lexbfs-bin-in-steps* = *lexbfs-bin*. We therefore name any theorem relating to *lexbfs-bin-in-steps* with the name *lexbfs-bin*. At the end, when we have shown LexBFS to be correct (i.e. sound and complete), we will use this to prove that *lexbfs-bin-in-steps* = *lexbfs-bin*.

definition *lexbfs-bin-in-steps* $\equiv$ *do* $\{$
  $(U,\ \mathcal{B},\ \sigma) <-$ *bin-loop* (*bin-state-at* []);
  $RETURN\ \sigma$
$\}$

We also need that *bin-loop* cannot fail. This amounts to just carefully unfolding the definition the right amount, as it follows from U being finite and not empty.

lemma *nofail-bin-loop*: *finite* $U \Longrightarrow$ *nofail* (*bin-loop* $(U,\ \mathcal{B},\ \sigma)$)
  ⟨*proof*⟩

The following are many lemmas that show that the possible results starting at the constructed state at $\sigma\ @\ [a]$ are contained in the results starting at σ.

Abbreviation for *inres*, that flips order and adds infix notation, to make results more readable. Note that the $\in$ is bold, unlike normal set-membership.

abbreviation *inres-rev* (**infix** $\in$ *50*) **where** *inres-rev x y* $\equiv$ *inres y x*

Next element in list can be picked by *bin-step*.

lemma *inres-bin-step-if-is-lexbfs*:
  is-lexbfs $(\sigma\ @\ a\#\sigma') \Longrightarrow$ *bin-state-at* $(\sigma\ @\ [a]) \in$ *bin-step* (*bin-state-at* σ)
  ⟨*proof*⟩

Full list can be output, starting at any split.

lemma *inres-bin-loop-if-is-lexbfs*:
is-lexbfs (σ @ σ') $\Longrightarrow$ *bin-state-at* (σ @ σ') $\in$ *bin-loop* (*bin-state-at* σ)
⟨*proof*⟩

Full list is valid output.

lemma *inres-lexbfs-bin-if-is-lexbfs*: *is-lexbfs* $\sigma \Longrightarrow \sigma \in$ *lexbfs-bin-in-steps*
⟨*proof*⟩

lemma *lexbfs-bin-complete*: *SPEC is-lexbfs* $\leq$ *lexbfs-bin-in-steps*
⟨*proof*⟩

5.2 *lexbfs-bin* refines *lexbfs*

The inverse sequence of refinements. We construct the ordinary label function $\mathcal{L}$ by first constructing $\mathcal{B}$, and then transforming it.

abbreviation *R-bin-helper* $\equiv \{((U, \mathcal{B}, \sigma), b).$
$b =$ (*card* U, U, *transform-B* (*construct-B* σ), σ) $\wedge$
$\mathcal{B} =$ *construct-B* $\sigma \wedge$
$U =$ *U-at* $\sigma \wedge$
partition-V U $\sigma\}$

lemma *lexbfs-bin-refines-lexbfs-helper*: *lexbfs-bin-in-steps* $\leq \Downarrow$ *Id lexbfs-helper*
⟨*proof*⟩

lemma *lexbfs-helper-refines-lexbfs*: *lexbfs-helper* $\leq \Downarrow$ *Id lexbfs*
⟨*proof*⟩

5.3 Completeness of *lexbfs*

theorem *lexbfs-complete*: *SPEC is-lexbfs* $\leq$ *lexbfs*
⟨*proof*⟩

6 Correctness

corollary *lexbfs-correct*: *lexbfs* $=$ *SPEC is-lexbfs*
⟨*proof*⟩

By squeezing the different algorithms in between *lexbfs* and *SPEC is-lexbfs*, we can show they are all equivalent. In particular, none of the algorithms can fail.

corollary *lexbfs-equiv*:
lexbfs $=$ *lexbfs-bin*
lexbfs $=$ *lexbfs-helper*
lexbfs $=$ *lexbfs-bin-in-steps*
⟨*proof*⟩

corollary *lexbfs-nofail*:
nofail lexbfs
nofail lexbfs-helper
nofail lexbfs-bin
nofail lexbfs-bin-in-steps
⟨*proof*⟩

end — Graph locale end

References

[1] A. Brandstädt, F. F. Dragan, and F. Nicolai. LexBFS-orderings and powers of chordal graphs. *Discret. Math.*, 171(1–3):27–42, 1997.

[2] D. G. Corneil. Lexicographic breadth first search—a survey. In J. Hromkovič, M. Nagl, and B. Westfechtel, editors, *WG 2004*, volume 3353 of *LNCS*, pages 1–19. Springer, 2004.

[3] D. G. Corneil, M. Habib, C. Paul, and M. Tedder. A recursive linear time modular decomposition algorithm via LexBFS. *CoRR*, abs/0710.3901, 2024.

[4] M. C. Golumbic. Triangulated graphs. In *Algorithmic Graph Theory and Perfect Graphs*, chapter 4, pages 81–104. Academic Press, 1980.

[5] Lexicographic breadth-first search. Wikipedia, https://en.wikipedia.org/w/index.php?title=Lexicographic_breadth-first_search&oldid=1314187570, Sept. 2025. Accessed: Oct. 25, 2025.

[6] D. J. Rose, R. E. Tarjan, and G. S. Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM J. Comput.*, 5(2):266–283, 1976.