THIS DOCUMENT IS THE ONLINE-ONLY APPENDIX TO:

Disjunction and Modular Goal-directed Proof Search

MATTHEW STONE
Rutgers, the State University of New Jersey

ACM Transactions on Computational Logic, Vol. V, No. N, February 2004, Pages 1–36.

A. PROOF OF THEOREM 2.23

We wish to eliminate delayed inferences from SCL proofs. This transformation depends on a generalization of delayed inferences, which we can term *misplaced* inferences since we intend to eliminate them. We assume an overall derivation $\mathcal{D}$, and consider a right inference R that applies to principal E within some subderivation $\mathcal{D}'$ of $\mathcal{D}$.

Definition A.1. We say a right inference R is *right-based* on an inference R' in $\mathcal{D}$ if $R = R'$ or R is based on R' and every inference on which R is based above and including R' is a right inference. Then R is *misplaced* in $\mathcal{D}'$ exactly when there are inferences M and R' in $\mathcal{D}'$ such that, in $\mathcal{D}$, M is based on an inference L, R is right-based on R', and R' is delayed with respect to L.

In this case we will also say R is misplaced *with respect to M*. We can abstract a key case of misplaced inferences by the following schematic derivation:

$$\begin{array}{lcl}
 & \overline{\qquad\qquad\qquad} & R \\
\text{Right inferences and inferences } R \text{ not based in} & \overline{\qquad\vdots\qquad} & M \\
R' \text{ delayed wrt } L & \overline{\quad \ldots E \ldots \quad} & R' \\
(M \text{ based in } L) & \overline{\qquad\downarrow\qquad} & L \\
 & \ldots E \ldots &
\end{array}$$

This schematic derivation shows informally how *misplaced inferences* help provide an inductive characterization of the inferences that stand in the way of obtaining an eager derivation. In an eager derivation, it will be impossible for R to appear above L. For R' cannot be delayed with respect to L, but once R' and L are interchanged, we will obtain a new delayed inference that R is based in, until finally we must interchange L and R. Of course, to do this, we must first interchange R with the *misplaced* inferences, such as M, which stand between R and L and cannot themselves be interchanged with L because they are based in L.

Observe that the relation R is misplaced with respect to M is asymmetrical. To see this, suppose R is misplaced with respect to M. By definition, R is right-based on R' which is delayed with respect to a left inference L on which M is based. Meanwhile, for M to be misplaced with respect to R, by definition, we must have M right-based on M' and R based in some left rule L_R. Any such M' would have to be based in L since no left inferences intervene between M and M'; M' must thus appear *inside* a schematic like that above. At the same time, since no left inferences intervene between R and R', R' would have to be based in any such L_R, which must thus appear *outside* such a schematic, closer to the root of the overall derivation. Accordingly, any such L_R must occur closer to the root of $\mathcal{D}$ than L; meanwhile the principal of M' is introduced further from the root than L. So we will not have M' delayed with respect to L_R.

Call R *badly misplaced* in $\mathcal{D}'$ if R is misplaced with respect to M and M occurs closer to the root than R. A subderivation $\mathcal{D}'$ with no badly misplaced inferences will be called *good*. An overall good derivation is also eager, since any delayed inference is badly misplaced.

LEMMA A.2. *Consider a subderivation $\mathcal{D}'$ of an overall derivation $\mathcal{D}$, with the property that $\mathcal{D}'$ has good immediate subderivations and that $\mathcal{D}'$ ends in inference M. From $\mathcal{D}'$ we can construct a derivation with the same end-sequent that is good.*

PROOF. The assumption that the immediate subderivations of $\mathcal{D}'$ are good is a very powerful one. For suppose that some inference is badly misplaced with respect to some other in $\mathcal{D}'$. Then we can only have some rule R badly misplaced with respect to M—anything else would contradict that assumption.

In fact, we can show that some such R must be adjacent to M. Consider an inference S that intervenes between R and M: we will show that S must be badly misplaced with respect to M too. By the definition of misplaced, M is based on some left rule L in $\mathcal{D}$, R is right-based on R', and R' is delayed with respect to L. Now consider the inferences that S is based on above L. If any of these is a left inference L', or S is itself a left inference, then R is also misplaced with respect to S—indeed, badly misplaced. This contradicts the assumption that the subderivations of $\mathcal{D}'$ are good. So none of these inferences can be a left inference, which means S is a right inference that is right-based on some inference S' above L. S' must be delayed with respect to L. Hence S is badly misplaced with respect to M.

Now we can proceed after [Kleene 1951, Lemma 10]. Define the *grade* of $\mathcal{D}'$ as the number of badly misplaced inferences in $\mathcal{D}'$. We show by induction on the grade that $\mathcal{D}'$ can be transformed to a good one.

The base case is a derivation of grade 0. This case has $\mathcal{D}'$ itself good. Thus, suppose the lemma holds for derivations of grade g, and consider $\mathcal{D}'$ of grade $g+1$. By the argument just given, one immediate subderivation—call it $\mathcal{D}''$—must end with an inference R which is badly misplaced with respect to M. Such an R of course cannot be based in M, so we interchange inferences R and M. In the result, the subderivation(s) ending in M satisfy the condition of the lemma with grade g or less. By applying the induction hypothesis, we can replace these subderivations with good ones. By asymmetry, M is not now badly misplaced with respect to R, nor can any of the other inferences be badly misplaced with respect to R, since they were not so in the original derivation. It follows that the result is a good derivation. □

Using this lemma, we can now present the proof of Theorem 2.23 in full.

THEOREM A.3 (THEOREM 2.23). Any SCL(I) derivation $\mathcal{D}$ is equal to an eager derivation $\mathcal{D}'$ up to permutations of inferences.

PROOF. Define the *reluctance* of $\mathcal{D}$ to be the number of rule applications R such that the subderivation $\mathcal{D}_R$ of $\mathcal{D}$ rooted in R is not good. We proceed by induction on reluctance. If reluctance is zero, $\mathcal{D}$ is itself good.

Now suppose the theorem holds for derivations of reluctance d, and consider $\mathcal{D}$ of reluctance $d+1$. Since $\mathcal{D}$ is finite, there must be a highest inference R such that some inference is badly misplaced with respect to R in the subderivation $\mathcal{D}_R$ rooted at R. This $\mathcal{D}_R$ satisfies the condition of Lemma A.2. Therefore this $\mathcal{D}_R$ can be replaced with a corresponding eager derivation, giving a new derivation of smaller reluctance. The induction hypothesis then shows that the resulting derivation can be made eager. □

B. PROOF OF THEOREM 3.5

THEOREM B.1 (THEOREM 3.5). *Let Γ and Δ be multisets of tracked prefixed expressions in which each formula is tracked by the empty set and prefixed by the empty prefix. There is a proof of $\Gamma \longrightarrow \Delta$ in SCL exactly when there is a proof of $\Gamma; \longrightarrow; \Delta$ in SCLP in which every block is canceled.*

PROOF. The argument for Theorem 3.5 depends on three lemmas: Lemma 3.8, proved in Section B.1; Lemma 3.16, proved in Section B.2; and Lemma 3.17, proved in Section B.3.

As observed already in Section 2.4, there is an SCL proof of $\Gamma \longrightarrow \Delta$ exactly when there is an SCLI proof of $\Gamma \longrightarrow \Delta$. By Theorem 2.23 of Section 2.4, there is an SCLI proof of $\Gamma \longrightarrow \Delta$ exactly when there is an *eager* SCLI proof of $\Gamma \longrightarrow \Delta$. By Lemma 3.6, there is an eager SCLI proof of $\Gamma \longrightarrow \Delta$ exactly when there is an eager articulated SCLI proof of $\Gamma; \longrightarrow; \Delta$. And by Lemma 3.8, there is an eager articulated SCLI proof of $\Gamma; \longrightarrow; \Delta$ exactly when there is an eager SCLS proof of $\Gamma; \longrightarrow; \Delta$.

Continuing through the argument, by the Contraction Lemma, we may assume without loss of generality that $\Gamma; \longrightarrow; \Delta$ is a simple sequent. We know from its lack of prefixes that the sequent $\Gamma; \longrightarrow; \Delta$ is also spanned and balanced. By Lemma 3.16 of Section B.2.3, then, there is an eager SCLS proof of $\Gamma; \longrightarrow; \Delta$ exactly when there is a blockwise eager SCLB derivation of $\Gamma; \longrightarrow; \Delta$ in which every block is canceled, linked, isolated, simple, balanced and spanned. And by Lemma 3.17, there is a blockwise eager SCLB derivation of $\Gamma; \longrightarrow; \Delta$ in which every block is canceled, linked, isolated, simple, balanced and spanned exactly when there is an SCLP derivation of $\Gamma; \longrightarrow; \Delta$ in which every inference is linked. And if every inference is linked, every block is canceled. □

B.1 Proof of Lemma 3.8

We show in this section that an articulated SCLI proof with end-sequent $\Pi; \longrightarrow; \Theta$ corresponds to an SCLS proof with end-sequent $\Pi; \longrightarrow; \Theta$, and vice versa. In fact, to transform SCLS to articulated SCLI we have a simple structural induction which replaces $(\supset\rightarrow^S)$ with $(\supset\rightarrow)$ using the weakening lemma; the soundness of SCLS over SCLI then follows by Lemma 3.6. Thus, here we are primarily concerned with completeness of a new sequent inference figure.

The use of $(\supset\rightarrow^S)$ in eager derivations ensures that the processing of each new goal refers directly to global program statements. To formalize this idea, we introduce the notion of a *fresh* inference.

Definition B.2 *Fresh*. Let $\mathcal{D}$ be an SCLV derivation. An inference R in $\mathcal{D}$ is *fresh* exactly when R is a right inference and the path from R to the root never follows the left spur of any $(\supset\rightarrow)$ inference.

LEMMA B.3. *Let $\mathcal{D}$ be an eager SCLV derivation with an end-sequent of the form*

$$\Pi; \rightarrow \Delta; \Theta$$

and consider a subderivation $\mathcal{D}'$ of $\mathcal{D}$ rooted in a fresh inference R. Then the end-sequent of $\mathcal{D}'$ also has the form

$$\Pi'; \rightarrow \Delta'; \Theta'$$

for some Π', Δ' and Θ'.

PROOF. Suppose otherwise, and consider a maximal $\mathcal{D}'$ whose end-sequent contains a non-empty multiset of local statements Γ. We can describe $\mathcal{D}'$ equivalently as the subderivation of $\mathcal{D}$ that is rooted in a lowest fresh inference R when the end-sequent of $\mathcal{D}$ contains some local statements. R cannot be the first inference of $\mathcal{D}$, so there must be an inference S in $\mathcal{D}$ immediately below R. If S is a left rule, then the fact that $\mathcal{D}$ is eager leads to a contradiction. R must be based in S, or else R will be delayed. This means S is an implication inference; but given that R is fresh, R must appear along the branch of $(\supset\rightarrow^S)$ without local statements. Meanwhile, if S is a right rule, it follows from the formulation of the rules that if the end-sequent of $\mathcal{D}_R$ has non-empty local statements then the end-sequent of $\mathcal{D}_L$ must also. This contradicts the assumption that R is first. □

Now we proceed with the proof of Lemma 3.8.

LEMMA B.4 (LEMMA 3.8). *An eager articulated SCLI derivation whose end-sequent is of the form*

$$\Pi; \rightarrow \Delta; \Theta$$

can be transformed to an eager SCLS derivation of the same end-sequent.

PROOF. We assume an eager SCLV derivation $\mathcal{D}$ with such an end-sequent; we show that we can transform it into an eager SCLS derivation $\mathcal{D}'$ with the same end-sequent. The proof is by induction on the number of occurrences of $(\supset\rightarrow)$ inferences in $\mathcal{D}$.

In the base case, there are no $(\supset\rightarrow)$ inferences and $\mathcal{D}'$ is just $\mathcal{D}$.

Suppose the claim holds for derivations where $(\supset\rightarrow)$ is used fewer than n times, and suppose $\mathcal{D}$ is a derivation in which $(\supset\rightarrow)$ is used n times. Choose an inference L of $(\supset\rightarrow)$ with no other $(\supset\rightarrow)$ inference closer to the root of $\mathcal{D}$; we must rewrite the left subderivation at L to match the $(\supset\rightarrow^S)$ inference figure. To do this we will draw on additional inferences from $\mathcal{D}$. We find these inferences in a subderivation $\mathcal{D}'$ of $\mathcal{D}$ distinguished as a function of L—in particular, we identify $\mathcal{D}'$ as the largest subderivation of $\mathcal{D}$ containing L but no right inferences or segment boundaries below L.

Using Lemma B.3, we develop a schema of $\mathcal{D}'$ thus:

$$\dfrac{\dfrac{\begin{array}{c}\mathcal{D}^A\\ \Pi; \Gamma, A \supset B^{\mu}_X \rightarrow A^{\mu}_X, \Delta; \Theta\end{array} \qquad \begin{array}{c}\mathcal{D}^B\\ \Pi; \Gamma, A \supset B^{\mu}_X, B^{\mu}_X \rightarrow \Delta; \Theta\end{array}}{\mathcal{D}^L \left\{ \begin{array}{c} \Pi; \Gamma, A \supset B^{\mu}_X \longrightarrow \Delta; \Theta \\ \vdots \\ \Pi; \longrightarrow \Delta; \Theta \end{array}\right.}\,L}{\text{(Segment boundary or right rule)}}$$

We suppose L applies to an expression $A \supset B^{\mu}_{X}$; the left subderivation of L, $\mathcal{D}^A$ adds the goal A; the right, $\mathcal{D}^B$, uses the assumption B. The subderivation of $\mathcal{D}'$ from the end-sequent of L abstracts the left inferences performed elsewhere in this segment (and any subgoals that these inferences trigger). We notate this tree of inferences $\mathcal{D}^L$. By Lemma B.3, $\mathcal{D}'$ ends with a sequent of the form $\Pi; \longrightarrow \Delta; \Theta$. Because of the form of the intervening rules, we have the same succedent $\Delta; \Theta$ at L, as well as the same global statements Π.

We use $\mathcal{D}^L$ to construct an eager SCLS derivation $\mathcal{A}$ corresponding to $\mathcal{D}^A$; we will substitute the result for the left subtree at L to revise L to fit the $(\supset\rightarrow^S)$ figure. In outline, the derivation we aim for is an eager SCLS version of:

$$\frac{\mathcal{D}^A}{\mathcal{D}^L + A^{\mu}_{X}}$$

The problem is that if $\mathcal{D}^A$ is rooted in a right inference to A, we will not obtain an eager derivation when we reassemble L. The SCLS derivation $\mathcal{A}$ we use is actually constructed by recursion on the structure of $\mathcal{D}^A$, applying this kind of transformation at appropriate junctures. At each stage, we call the subderivation of $\mathcal{D}^A$ we are considering $\mathcal{D}'^A$.

For the base case, this subderivation is an axiom, and we construct this subderivation as a result. If $\mathcal{D}'^A$ ends in a right rule, the construction proceeds inductively by constructing corresponding subderivations and recombining them by the same right rule. With a right inference here, the resulting derivation must be eager since the subderivations are eager.

If $\mathcal{D}'^A$ ends in a left inference, the construction is immediate. We observe that $\mathcal{D}'^A$ has an end-sequent of the form

$$\Pi, \Pi'; \longrightarrow \Delta, \Delta'; \Theta, \Theta'$$

(The inventory of expressions can only be expanded, and that only in certain places, as we follow right inferences to reach $\mathcal{D}'^A$.) So we first weaken $\mathcal{D}^L$ by the needed additional expressions—Π' on the left and Δ' (locally) and Θ' (globally) on the right; then we identify the open leaf in $\mathcal{D}^L$ with $\mathcal{D}'^A$, obtaining a larger derivation $\mathcal{D}_I$ defined as:

$$\frac{\mathcal{D}'^A}{\Pi' + \mathcal{D}'^L + A^{\mu}_{X} + \Delta'; \Theta'}$$

Any delayed inference in $\mathcal{D}_I$ would in fact be delayed in $\mathcal{D}'^A$, so this is an eager derivation. The result has, moreover, fewer than n $(\supset\rightarrow)$ inferences, since it omits at least L from $\mathcal{D}'$. Then the induction hypothesis applies to give the needed SCLS derivation $\mathcal{A}$.

Given the derivation $\mathcal{A}$ so constructed, we substitute $\mathcal{A}$ for $\mathcal{D}^A$ in $\mathcal{D}$. The result $\mathcal{D}^*$ is an eager derivation; $\mathcal{D}^*$ contains an $(\supset\rightarrow^S)$ inference corresponding to L and therefore contains fewer than n uses of $(\supset\rightarrow)$. The induction hypothesis applies to transform $\mathcal{D}^*$ to the needed overall derivation. □

B.2 Proof of Lemma 3.16

B.2.1 *Replacing Herbrand terms.* To begin, it is convenient to observe that the use of indexed Herbrand terms allows us to rename Herbrand terms in a proof under certain conditions.

LEMMA B.5 SUBSTITUTION. *Let $\mathcal{D}$ be an SCLU derivation with end-sequent*

$$\Pi; \longrightarrow; \Theta$$

in which no Herbrand terms or Herbrand prefixes appear; consider a spanned simple subderivation $\mathcal{D}'$ in which a modal Herbrand function η^u_A occurs in some sequent, but does not occur in the end-sequent. Let η^v_A be a Herbrand function that does not occur in $\mathcal{D}$. Then we can construct a proof $\mathcal{D}^*$ containing corresponding inferences in a corresponding order to $\mathcal{D}$ but in which Herbrand terms and Herbrand prefixes are adjusted so that η^v_A is used in place of η^u_A precisely in the subderivation corresponding to $\mathcal{D}'$.

PROOF. The argument proceeds by induction on the structure of derivations. A complex substitution may be required, because the Herbrand calculus may require not only the replacement of η^u_A itself but also the replacement of Herbrand terms that depend indirectly on η^u_A. It is convenient to begin by replacing any first-order Herbrand term not introduced by a $(\exists \rightarrow)$ or $(\rightarrow \forall)$ inference by a distinguished constant c_0—starting with leaves of the derivation and working downward. This replacement is to ensure that each first-order and modal Herbrand term in $\mathcal{D}$ is determined from an expression in the end-sequent of $\mathcal{D}$ by a finite number of steps of inference. We continue with the systematic replacement of η^u_A and its dependents. In both cases, the form of $\mathcal{D}$ ensures that a finite substitution can systematically rename all these Herbrand terms as required. We use the fact that each sequent is simple and spanned to extend this substitution inductively upward. Because each sequent is spanned the substitution does not need to be extended at $(\Box \rightarrow)$ inferences; because each sequent is simple the substitution can be extended freshly at $(\rightarrow \Box)$ and $(\rightarrow >)$ inferences. Finally, the form of first-order Herbrand terms ensures that a finite extension of the substitution suffices for $(\rightarrow \exists)$ and $(\forall \rightarrow)$ inferences. $\Box$

B.2.2 *Rectifying blocks.* The transformation of individual blocks appeals to the following definition of *required* elements of proofs.

Definition B.6 *Required*. Given a derivation $\mathcal{D}$ with end-sequent

$$\Pi;\Gamma \longrightarrow \Delta;\Theta$$

we say that an expression occurrence E in Θ or Π is *required* iff either it is linked or some block in $\mathcal{D}$ is adjacent to the root block and has an end-sequent

$$\Pi'; \longrightarrow ;\Theta'$$

in which Π' or Θ' contains an expression occurrence based in E.

LEMMA B.7 RECTIFICATION. *We are given a blockwise eager SCLU derivation $\mathcal{D}$ such that: every block in $\mathcal{D}$ is canceled and isolated; every block in $\mathcal{D}$ other than the root is spanned, linked, balanced and simple; and the end-sequent of $\mathcal{D}$ is balanced. We transform $\mathcal{D}$ to an SCLU derivation $\mathcal{D}'$ in which every block is canceled, linked, isolated, balanced and simple and every block other than the root is spanned. Every block in $\mathcal{D}'$ other than the root block is identical to a block of $\mathcal{D}$; and the inferences in the root block of $\mathcal{D}$ correspond to inferences in the same order in $\mathcal{D}$ (and so $\mathcal{D}'$ is blockwise eager). If the end-sequent of $\mathcal{D}$ is spanned then $\mathcal{D}'$ is spanned and isolated.*

PROOF. We describe a transformation that establishes the following inductive property given $\mathcal{D}$. There are simple multisets $\Pi_M \subseteq \Pi$ and $\Theta_M \subseteq \Theta$, together with multisets $\Gamma' \subseteq \Gamma$ and $\Delta' \subseteq \Delta$ such that: any Θ' that spans Π_M includes Θ_M; and for any simple Π' with $\Pi_M \subseteq \Pi' \subseteq \Pi$ and any simple Θ' with $\Theta' \subseteq \Theta$ such that Π' and Θ' are spanned by Θ' and the pair Π', Θ' is balanced, there is a $\mathcal{D}'$ in which every block is canceled, linked, balanced,

balanced and simple, with end-sequent:

$$\Pi';\Gamma' \longrightarrow \Delta';\Theta'$$

In this $\mathcal{D}'$, each expression in Γ' is linked; each expression in Δ' is linked; each Π_M expression that occurs in Π' is required and each Θ_M expression that occurs in Θ' is linked. Every block in $\mathcal{D}'$ other than the root block is identical to a block of $\mathcal{D}$; and the inferences in the root block of $\mathcal{D}$ correspond to inferences in the same order in $\mathcal{D}$. Finally, if Γ' and Δ' are spanned by Θ' then $\mathcal{D}'$ is spanned; if $\mathcal{D}$ is linked then $\mathcal{D}'$ contains all the axioms of $\mathcal{D}$.

At axioms, for $\mathcal{D}$ of

$$\Pi;\Gamma,A^{\mu}_{X} \longrightarrow A^{\mu}_{Y},\Delta;\Theta$$

Π_M and Θ_M are empty, while $\Gamma' = A^{\mu}_{X}$ and $\Delta' = A^{\mu}_{X}$. Assume we are given simple Π' from Π and simple Θ' from Θ with Π' and Θ' spanned by Θ'. We construct $\mathcal{D}'$ of

$$\Pi';A^{\mu}_{X} \longrightarrow A^{\mu}_{Y};\Theta'$$

If A^{μ}_{X} is spanned by Θ', this axiom is spanned too; the remaining conditions are immediate.

At inferences, consider as a representative case $(\vee \rightarrow)$. $\mathcal{D}$ ends:

$$\dfrac{\stackrel{\mathcal{D}_1}{\Pi;\Gamma,A\vee B^{\mu}_{X},A^{\mu}_{X} \longrightarrow \Delta;\Theta} \qquad \stackrel{\mathcal{D}_2}{\Pi;\Gamma,A\vee B^{\mu}_{X},B^{\mu}_{X} \longrightarrow \Delta;\Theta}}{\Pi;\Gamma,A\vee B^{\mu}_{X} \longrightarrow \Delta;\Theta}$$

The blocks of $\mathcal{D}_1$ and $\mathcal{D}_2$ either contain the root or are blocks from $\mathcal{D}$; the Herbrand prefixes in the end-sequents of $\mathcal{D}_1$ and $\mathcal{D}_2$ occur with the same distribution as in $\mathcal{D}$. Therefore we can apply the induction hypothesis to get Π_{M1}, Θ_{M1}, Γ'_1 and Δ'_1 for $\mathcal{D}_1$; we can apply it to get Π_{M2}, Θ_{M2}, Γ'_2 and Δ'_2 for $\mathcal{D}_2$. To transform $\mathcal{D}$ itself, we perform case analysis on Γ'_1 and Γ'_2.

If Γ'_1 does not contain an occurrence of A^{μ}_{X}, then $\Pi_M = \Pi_{M1}$, $\Theta_M = \Theta_{M1}$, $\Gamma' = \Gamma'_1$ and $\Delta' = \Delta'_1$; $\mathcal{D}'_1$ suffices to carry through the induction hypothesis.

Similarly, if Γ'_2 does not contain an occurrence of B^{μ}_{X}, then $\Pi_M = \Pi_{M2}$, $\Theta_M = \Theta_{M2}$, $\Gamma' = \Gamma'_2$ and $\Delta' = \Delta'_2$; $\mathcal{D}'_2$ suffices to carry through the induction hypothesis.

Otherwise, we will set up $\Pi_M = \Pi_{M1} \cup \Pi_{M2}$ and $\Theta_M = \Theta_{M1} \cup \Theta_{M2}$ (as sets); by the inductive characterization of Π_{M1}, Π_{M2}, Θ_{M1} and Θ_{M2}, any Θ' that spans both Π_{M2} and Π_{M2} includes both Θ_{M1} and Θ_{M2}. We also set up Γ' as the multiset containing at least one occurrence of $A\vee B^{\mu}_{X}$ and as many expression occurrences of any expression as either are found in $\Gamma'_1 \backslash A^{\mu}_{X}$ or are found in $\Gamma'_2 \backslash B^{\mu}_{X}$; we set up Δ' as the multiset containing as many expression occurrences of any expression as are found in either Δ'_1 or Δ'_2.

To continue, we now consider simple Π' from Π and simple Θ' from Θ such that $\Pi_{M1} \subseteq \Pi'$, $\Pi_{M2} \subseteq \Pi'$, Π' and Θ' are spanned by Θ', and the pair Π',Θ' is balanced. We know that Θ' includes Θ_M. We can apply the inductive property to transform $\mathcal{D}_1$ and $\mathcal{D}_2$ into derivations with the inductive property:

$$\stackrel{\mathcal{D}'_1}{\Pi';\Gamma'_1 \longrightarrow \Delta'_1;\Theta'} \qquad \stackrel{\mathcal{D}'_2}{\Pi';\Gamma'_2 \longrightarrow \Delta'_2;\Theta'}$$

We weaken *the lowest block* of $\mathcal{D}'_1$ on the left by the expressions in Γ^{+} and not already in Γ' and on the right by the expressions in Δ^{+} and not already in Δ', giving $\mathcal{D}^{+}_1$. We similarly weaken the lowest block of $\mathcal{D}'_2$ on the left by the expressions in Γ^{+} and not already in Γ'_2

and on the right by the expressions in Δ^+ and not already in Δ_2', giving $\mathcal{D}_2^+$. Only the lowest blocks are affected by the weakening transformations, so other blocks remain canceled, linked, spanned, isolated and simple; the lowest block in each case remains canceled. The lowest blocks also remain linked since no inferences are added; and they remain simple (and balanced) because no weakening occurs in the global areas. Construct $\mathcal{D}'$ as

$$\dfrac{\dfrac{\mathcal{D}_1^+}{\Pi';\Gamma^+,A_X^\mu \longrightarrow \Delta^+;\Theta'} \qquad \dfrac{\mathcal{D}_2^+}{\Pi';\Gamma^+,B_X^\mu \longrightarrow \Delta^+;\Theta'}}{\Pi';\Gamma^+ \longrightarrow \Delta^+;\Theta'}$$

The end-sequent is simple and balanced so the root block is simple and balanced; the inference is linked since A_X^μ and B_X^μ are linked in the subderivations, so the root block is linked. The root block remains canceled as always.

Any Π_M expression is required here because it is required either in $\mathcal{D}_1^+$ in virtue of its membership in Π_{M1} or in $\mathcal{D}_2^+$ in virtue of its membership in Π_{M2}; likewise any Θ_M expression is linked here because it is linked either in $\mathcal{D}_1^+$ in virtue of its membership in Θ_{M1} or in $\mathcal{D}_2^+$ in virtue of its membership in Θ_{M2}. Thus, except for the spanning conditional, we have shown everything we need of this $\mathcal{D}'$.

Finally, then, if Γ' and Δ' is spanned by Θ', Δ_1' and Δ_2' are spanned by Θ' and Γ_1' and Γ_2' are spanned by Θ' in the resulting (spanned) subderivations $\mathcal{D}_1'$ and $\mathcal{D}_2'$. This shows that the end-sequent of $\mathcal{D}'$ is also spanned, so $\mathcal{D}'$ itself is spanned.

This reasoning is representative of the construction required also for $(\wedge \rightarrow)$, $(\exists \rightarrow)$, $(\forall \rightarrow)$, $(\rightarrow \wedge)$, $(\rightarrow \vee)$, $(\rightarrow \exists)$, $(\rightarrow \forall)$, (decide) and (restart). It applies also for $(\supset\rightarrow^S)$, with the obvious caveat that we do not weaken the left subderivation to match local left expressions, since the form of the $(\supset\rightarrow^S)$ inference requires there to be none.

Next we have $(\vee \rightarrow^B)$; we consider the representative case of $(\vee \rightarrow_L^B)$. $\mathcal{D}$ ends:

$$\dfrac{\dfrac{\mathcal{D}_1}{\Pi_0,\Pi;\Gamma,A\vee B_X^\mu,A_X^\mu \longrightarrow \Delta;\Theta_0,\Theta} \qquad \dfrac{\mathcal{D}_2}{\Pi_0,B_X^\mu; \longrightarrow \Theta_0}}{\Pi_0,\Pi;\Gamma,A\vee B_X^\mu \longrightarrow \Delta;\Theta_0,\Theta}$$

We treat this specially to respect the block boundary before $\mathcal{D}_2$. In particular, we apply the induction hypothesis to $\mathcal{D}_1$ (as we may since its end-sequent has the same distribution of Herbrand prefixes as does that of $\mathcal{D}$), to get Π_{M1}, Θ_{M1}, Γ_1' and Δ_1'. If A_X^μ does not occur in Γ_1', we let $\Pi_M = \Pi_{M1}$, $\Theta_M = \Theta_{M1}$, $\Gamma' = \Gamma_1'$ and $\Delta' = \Delta_1'$; any derivation $\mathcal{D}_1'$ constructed from appropriate Π' and Θ' suffices to carry through the induction hypothesis.

Otherwise, we get $\Pi_M = \Pi_{M1} \cup \Pi_{e0}$ (as a set), $\Theta_M = \Theta_{M1}$; any Θ' that spans Π_M also spans Π_{M1} and so includes Θ_M. $\Delta' = \Delta_1'$ and Γ' contains Γ_1' with the occurrence of A_X^μ removed, together with an occurrence of $A\vee B_X^\mu$ if Γ_1' does not already contain such an expression.

Assume simple Π' with $\Pi_M \subseteq \Pi' \subseteq \Pi$ and simple Θ' with $\Theta' \subseteq \Theta$ with Π' and Θ' spanned by Θ' and the pair Π',Θ' balanced. As before, we must have Θ_M included in Θ'. We therefore obtain $\mathcal{D}_1'$ by the inductive property; we then weaken $\mathcal{D}_1'$ locally within the lowest block by $A\vee B_X^\mu$ on the left if necessary, to obtain a good derivation $\mathcal{D}_1^*$.

The needed $\mathcal{D}'$ is now constructed as:

$$\dfrac{\dfrac{\mathcal{D}_1^*}{\Pi';\Gamma',A_X^\mu \longrightarrow \Delta';\Theta'} \qquad \dfrac{\mathcal{D}_2}{\Pi_0,B_X^\mu; \longrightarrow \Theta_0}}{\Pi';\Gamma' \longrightarrow \Delta';\Theta'}$$

We first argue that the construction instantiates the $(\vee \rightarrow_L^B)$ inference rule. Every Herbrand prefix in Π_{0e} and B_X^μ occurs in Π' or Γ', so Π_{0e} and B_X^μ are spanned by Θ'. But because the root block in $\mathcal{D}$ is isolated, Π_{0e} and B_X^μ are spanned minimally by Θ_0. Thus $\Theta_0 \subseteq \Theta'$. $\Pi_{0e} \subseteq \Pi_M$ by construction; by isolation Π_0 is the smallest set such that the pair of Π_0, Θ_0 is balanced. But since Π', Θ' is balanced, $\Pi_0 \subseteq \Pi'$.

Now we show that $\mathcal{D}'$ so constructed has the needed properties. The end-sequent is simple and balanced so the root block is simple and balanced. The inference is linked: A_X^μ is linked in $\mathcal{D}_1'$ by the induction hypothesis; B_X^μ is linked in $\mathcal{D}_2$ because $\mathcal{D}_2$ begins a new block which by assumption is canceled. The root block remains canceled as always. Any Π_M expression is required here because either a corresponding expression Π_{0e} in the new block at the left subderivation is based on it, or because it is required in $\mathcal{D}_1'$. Every Θ_M is linked because it is linked in $\mathcal{D}_1^*$.

Finally, if Γ' and Δ' are spanned by Θ', then Δ_1' and Γ_1' are spanned by Θ_1'. The new subderivation $\mathcal{D}_1'$ is therefore spanned by the inductive property; this ensures that the overall derivation is spanned.

Next consider $(\Box \rightarrow)$. $\mathcal{D}$ ends:

$$\dfrac{\begin{array}{c}\mathcal{D}_1\\ \Pi;\Gamma,\Box_i A_X^\mu, A_{X,\mu\nu}^{\mu\nu} \longrightarrow \Delta;\Theta\end{array}}{\Pi;\Gamma,\Box_i A_X^\mu \longrightarrow \Delta;\Theta}$$

As always, we apply the induction hypothesis to $\mathcal{D}_1$ (as we may since the Herbrand prefixes on Π and Θ formulas remain the same) to obtain Π_{M1}, Θ_{M1}, Γ_1' and Δ_1'. If $A_{X,\mu\nu}^{\mu\nu}$ does not occur in Γ_1', we let $\Pi_M = \Pi_{M1}$, $\Theta_M = \Theta_{M1}$, $\Gamma' = \Gamma_1'$ and $\Delta' = \Delta_1'$; any subderivation $\mathcal{D}_1'$ obtained by the inductive property suffices to witness the inductive property for $\mathcal{D}$.

Otherwise we obtain Γ' by extending Γ_1' by the principal expression $\Box_i A_X^\mu$ if necessary and eliminating the side expression $A_{X,\mu\nu}^{\mu\nu}$; $\Pi_M = \Pi_{M1}$, $\Theta_M = \Theta_{M1}$ and $\Delta' = \Delta_1'$. (Since these are common to the subderivation, any Π' that spans Π_M includes Θ_M.) Now we consider Π' with $\Pi_M \subseteq \Pi' \subseteq \Pi$ and Θ' with $\Theta' \subseteq \Theta$, Π' and Θ' spanned by Θ' and the pair Π', Θ' balanced. As always, we have $\Theta_M \subseteq \Theta'$. We obtain $\mathcal{D}_1'$ using Π' and Θ', and weaken the lowest block by local formulas; calling the result $\mathcal{D}_1^+$, we can produce $\mathcal{D}'$ by the following construction:

$$\dfrac{\begin{array}{c}\mathcal{D}_1^+\\ \Pi';\Gamma', A_{X,\mu\nu}^{\mu\nu} \longrightarrow \Delta';\Theta'\end{array}}{\Pi';\Gamma' \longrightarrow \Delta';\Theta'}$$

Everything is largely as before. The key new reasoning comes when we assume that Γ' and Δ' are spanned by Θ'. We must argue that $\Gamma', A_{X,\mu\nu}^{\mu\nu}$ is in fact spanned by Θ'. Since $A_{X,\mu\nu}^{\mu\nu}$ is linked in $\mathcal{D}_1^+$, there must be an axiom in this block which is based in $A_{X,\mu\nu}^{\mu\nu}$; indeed, since the expression occurs as a local antecedent, this axiom must occur within the segment. This axiom must pair expressions prefixed by a path μ' where $\mu\nu$ is a prefix of μ'. But because $\mathcal{D}'$ remains blockwise eager, no inferences apply to Δ' or Θ' formulas within the segment (nor can they in this fragment augment the Δ' or Θ' formulas within the segment); therefore some Δ' expression is associated with Herbrand prefix μ'. But since Δ' is spanned by Θ', we have that every prefix of μ' is associated with some Θ' expression; so every prefix of $\mu\nu$ is associated with some Θ' expression. Thus $\mathcal{D}_1^+$ is spanned and in turn $\mathcal{D}'$ is spanned.

We have one last representative class of inferences in $\mathcal{D}$: $(\rightarrow \Box)$ and $(\rightarrow >)$. We illustrate with the case where $\mathcal{D}$ ends in $(\rightarrow>)$:

$$\dfrac{\begin{array}{c}\mathcal{D}_1\\ \Pi, A^{\mu\eta}_{X,\mu\eta}; \Gamma \longrightarrow \Delta, A >_i B^{\mu}_{X}; \Theta, B^{\mu\eta}_{X,\mu\eta}\end{array}}{\Pi; \Gamma \longrightarrow \Delta, A >_i B^{\mu}_{X}; \Theta}$$

We begin by applying the induction hypothesis to $\mathcal{D}_1$ (as we can, given the symmetric extension of Π and Θ by labeled expressions). We obtain Θ_{M1}, Π_{M1}, Γ'_1 and Δ'_1; we consider alternative cases in response to Θ and Θ_{M1}. First we suppose $B^{\mu\eta}_{X,\mu\eta} \notin \Theta$. It follows by our assumption about $\mathcal{D}$ that $A^{\mu\eta}_{X,\mu\eta} \notin \Pi$ either, nor does η occur in Θ. For this case, we start by defining an overall Π_M and Θ_M: Θ_M is Θ_{M1} with any occurrence of $B^{\mu\eta}_{X,\mu\eta}$ eliminated; Π_M is Π_{M1} with any occurrence of $A^{\mu\eta}_{X,\mu\eta}$ eliminated. Π_M contains no occurrences of $\mu\eta$, since Π does not; thus given the inductive property of Θ_{M1} and Π_{M1}, any Θ' that spans Π_M spans Θ_M. We define Γ' and Δ' so that $\Gamma' = \Gamma'_1$ and Δ' contains Δ'_1 together with an occurrence of $A >_i B^{\mu}_{X}$, provided Δ'_1 does not already contain one and $B^{\mu\eta}_{X,\mu\eta} \in \Theta_{M1}$ or $A^{\mu\eta}_{X,\mu\eta} \in \Pi_{M1}$. So, assume we are given simple Π' with $\Pi_M \subseteq \Pi' \subseteq \Pi$ and simple Θ' with $\Theta' \subseteq \Theta$ (and so $\Theta_M \subseteq \Theta'$) such that Π' and Θ' are spanned by Θ' and the pair Π', Θ' is balanced.

We consider whether $B^{\mu\eta}_{X,\mu\eta} \in \Theta_{M1}$ or $A^{\mu\eta}_{X,\mu\eta} \in \Pi_{M1}$. If neither, we apply the induction hypothesis to $\mathcal{D}_1$ for the case that Θ'_1 is Θ' and Π'_1 is Π'. The resulting derivation $\mathcal{D}'_1$ serves as $\mathcal{D}'$.

Otherwise, $B^{\mu\eta}_{X,\mu\eta} \in \Theta_{M1}$ or $A^{\mu\eta}_{X,\mu\eta} \in \Pi_{M1}$; we apply the inductive property of $\mathcal{D}_1$ for the case that Θ'_1 is $\Theta', B^{\mu\eta}_{X,\mu\eta}$ and Π'_1 is $\Pi', A^{\mu\eta}_{X,\mu\eta}$ (clearly Π'_1 and Θ'_1 are spanned by Θ'_1 assuming Π' and Θ' are spanned by Θ'; the pair Π'_1, Θ'_1 is also balanced given its symmetric extension). If $B^{\mu\eta}_{X,\mu\eta} \in \Theta_{M1}$, by the inductive property it is linked. If $A^{\mu\eta}_{X,\mu\eta} \in \Pi_{M1}$, it is required, but we shall show that it is in fact linked. By the definition of being required, the other possibility is that there is a block adjacent to the root block of $\mathcal{D}'_1$ with end-sequent

$$\Pi'', E; \longrightarrow \Theta''$$

in which the $(\vee \rightarrow^{B})$ inference R that bounds the block is based in E and Π'', E or Θ'' contains an expression occurrence based in $A^{\mu\eta}_{X,\mu\eta}$. But since the original block is isolated in the original $\mathcal{D}$, it is E that must be based in $A^{\mu\eta}_{X,\mu\eta}$. But then R is based in $A^{\mu\eta}_{X,\mu\eta}$ and R is linked: in particular its side expression in the left spur) must be linked; so $A^{\mu\eta}_{X,\mu\eta}$ is linked too.

Thus we can weaken $\mathcal{D}'_1$ in its lowest block if necessary by $A >_i B^{\mu}_{X}$ as a local right formula (in Γ), producing $\mathcal{D}^{+}_1$; $\mathcal{D}^{+}_1$ remains good by the same argument as the earlier cases. Thus we can construct $\mathcal{D}'$ as:

$$\dfrac{\begin{array}{c}\mathcal{D}^{+}_1\\ \Pi', A^{\mu\eta}_{X,\mu\eta}; \Gamma' \longrightarrow \Delta', A >_i B^{\mu}_{X}; \Theta', B^{\mu\eta}_{X,\mu\eta}\end{array}}{\Pi'; \Gamma' \longrightarrow \Delta'; \Theta'}$$

The end-sequent here is simple and balanced, so the whole root block is simple and balanced. The new inference is linked (in virtue of the linked occurrence of one side expression—$A^{\mu\eta}_{X,\mu\eta}$ or $B^{\mu\eta}_{X,\mu\eta}$) so the whole root block is linked. The root block is of course canceled. Each element of Π_M is required because it is an element of Π_{M1} and required in the immediate subderivation; each element of Θ_M is linked, because it is an element of Θ_{M1} and therefore linked in the immediate subderivation.

To conclude the case, suppose the end-sequent of $\mathcal{D}$ is spanned and that Γ' and Δ' are spanned by Θ'; it follows that same property applies to $\mathcal{D}_1$ so the subderivation is spanned. Then the end-sequent must also be spanned.

The alternative case has $B^{\mu\eta}_{X,\mu\eta} \in \Theta$. By assumption, it also has $A^{\mu\eta}_{X,\mu\eta} \in \Pi$. We therefore define an overall Π_M and Θ_M directly as Π_{M1} and Θ_{M1}, respectively; similarly $\Gamma' = \Gamma'_1$ and $\Delta' = \Delta'_1$. To construct the needed $\mathcal{D}'$ for appropriate Π' and Θ', we simply apply the induction hypothesis to $\mathcal{D}_1$ for the case that Θ'_1 is Θ' and Π'_1 is Π'. The resulting derivation $\mathcal{D}'_1$ suffices.

Having completed the induction, we argue that we can obtain an overall $\mathcal{D}'$ that is isolated, assuming the original $\mathcal{D}$ is not only isolated but spanned. Apply the inductive result to $\mathcal{D}$ for the case $\Pi' = \Pi$ and $\Theta' = \Theta$; since $\Gamma' \subseteq \Gamma$ and $\Delta' \subseteq \Delta$ we obtain a spanned derivation $\mathcal{D}'$ ending

$$\Pi;\Gamma' \longrightarrow \Delta';\Theta$$

Consider the end-sequent of any block other than the root in $\mathcal{D}'$; it is

$$\Pi_0, E; \longrightarrow ; \Theta_0$$

where a corresponding block occurs in $\mathcal{D}$. I argue by contradiction that for any $F \in \Pi_0$ either $F \in \Pi$ or F is based in an occurrence of F as the side expression of an inference in $\mathcal{D}'$ in which E is also based. (This will show that $\mathcal{D}'$ is isolated.) So consider an exceptional F. Since $\mathcal{D}$ is isolated, if $F \notin \Pi$, F is based in an occurrence of F as the side expression of an inference in $\mathcal{D}$ in which E is also based; this inference introduces some path symbol η which occurs in the label of F and E. In $\mathcal{D}'$, E can not be based in such an inference; otherwise F would also be based in that inference, since $\mathcal{D}'$ is simple. (We have assumed that F is not based in such an inference.) But in this case the expression in the end-sequent of $\mathcal{D}'$ on which E is based must contain η. Because the end-sequent of $\mathcal{D}'$ is spanned the form of Π and Θ is constrained in $\mathcal{D}$, F must occur in Π. This is absurd. □

We conclude Section B.2.2 by observing some facts about this construction. First, let $\mathcal{D}'$ be a derivation obtained by the construction of Lemma B.7, and suppose $\mathcal{D}'$ is weakened (in a spanned and balanced way) to $\mathcal{D}''$ by adding occurrences of global expressions that either already occur in the end-sequent of $\mathcal{D}'$ or never occur as global expressions in $\mathcal{D}'$. Then a straightforward induction shows that $\mathcal{D}'$ is obtained again from $\mathcal{D}''$ by the construction of Lemma B.7.

Second, observe that if $\mathcal{D}'$ is a derivation obtained by the construction of Lemma B.7, and $\mathcal{D}''$ is obtained from $\mathcal{D}''$ by the renaming of Herbrand prefixes (as in Lemma B.5), then straightforward induction shows that $\mathcal{D}''$ is obtained again from $\mathcal{D}''$ by the construction of Lemma B.7.

Third, let $\mathcal{D}'$ be a derivation for which the construction of Lemma B.7 yields itself. Let ν be a prefix and let the $\Pi;\Theta$ be the smallest balanced pair where Θ contains all the carriers of prefixes of ν introduced in $\mathcal{D}'$. Suppose each expression in Π and Θ has the property that at most one inference of $\mathcal{D}'$ has an occurrence of that expression as a side expression. Consider a derivation $\mathcal{D}''$ obtained from $\mathcal{D}'$ by weakening globally by Π (on the left) and by Θ (on the right). Let $\mathcal{D}^*$ be the result of applying the construction of Lemma B.7 to $\mathcal{D}''$. Then $\mathcal{D}^*$ contains any subderivation of $\mathcal{D}'$ whose end-sequent contains Π and Θ as global formulas. Again this is a straightforward induction; the base case considers a subderivation of $\mathcal{D}'$ whose end-sequent contains Π and Θ as global formulas; in this case we apply the first observation. Unary inferences extend the claim immediately. At binary inferences, one subderivation must be unchanged, by the first observation: since Π and Θ are introduced on a unique path, each Π and Θ formula never occurs or already occurs in the end-sequent in that subderivation. Thus the other subderivation necessarily appears in the derivation obtained by the construction of Lemma B.7.

B.2.3 *Block conversion.* We now have the background required to perform the conversion to block structure, and complete the proof of Lemma 3.16.

LEMMA B.8 (LEMMA 3.16). *We are given a blockwise eager SCLS derivation $\mathcal{D}$ whose end-sequent is spanned and balanced and takes the form:*

$$\Pi; \longrightarrow; \Theta$$

We transform $\mathcal{D}$ into a blockwise eager SCLB derivation in which every block is canceled, linked, isolated, simple, balanced and spanned.

PROOF. Our induction hypothesis is stronger than the lemma. We assume a blockwise eager SCLU derivation $\mathcal{D}$ with end-sequent of the form

$$\Pi; \longrightarrow; \Theta$$

in which every block is canceled, linked, isolated, simple, balanced and spanned, such that that the subproof rooted at any $(\vee \rightarrow)$ inference in $\mathcal{D}$ is an SCLS derivation. And we identify a distinguished expression occurrence E in the end-sequent of $\mathcal{D}$ which is linked. By Lemma B.7, it is straightforward to obtain such a derivation from the SCLS derivation (containing only a single block) that we have assumed. We transform $\mathcal{D}$ into a blockwise eager SCLB derivation in which every block is canceled, linked, isolated, simple, balanced and spanned and in which E is also linked; we perform induction on the number of $(\vee \rightarrow)$ inferences in $\mathcal{D}$.

In the base case there are no $(\vee \rightarrow)$ inferences, so $\mathcal{D}$ itself is an SCLB derivation.

In the inductive case, we assume $\mathcal{D}$ with n $(\vee \rightarrow)$ inferences, and assume the hypothesis true for derivations with fewer. We find an application L of $(\vee \rightarrow)$ with no other closer to the root of $\mathcal{D}$. We will transform $\mathcal{D}$ to eliminate L.

Let $\mathcal{D}'$ denote the smallest subderivation of $\mathcal{D}$ containing the full block of $\mathcal{D}$ in which L occurs. Explicitly, $\mathcal{D}'$ may be $\mathcal{D}$ itself; otherwise, $\mathcal{D}'$ is rooted at the right subderivation of the highest $(\vee \rightarrow^B)$ inference below L—an inference we will refer to as H. In either case, our assumptions allow us to identify a distinguished linked expression F in the end-sequent of $\mathcal{D}'$: either the assumed E from $\mathcal{D}$, or the side expression of the inference H (assumed canceled). Suppose $A \vee B_Y^{\nu}$ is the principal of L. We can apply Lemma B.5 to rename $A \vee B_Y^{\nu}$ to $A \vee B_X^{\mu}$ in such a way that each symbol in μ that is introduced in $\mathcal{D}'$ is introduced by a unique inference there. Now we can infer the following schema for $\mathcal{D}'$:

$$\left[\dfrac{\begin{array}{c}\mathcal{D}^A\\ \Pi_0, F, \Pi; \Gamma, A \vee B_X^{\mu}, A_X^{\mu} \longrightarrow \Delta; \Theta_0, \Theta\end{array} \qquad \begin{array}{c}\mathcal{D}^B\\ \Pi_0, F, \Pi; \Gamma, A \vee B_X^{\mu}, B_X^{\mu} \longrightarrow \Delta; \Theta_0, \Theta\end{array}}{\Pi_0, F, \Pi; \Gamma, A \vee B_X^{\mu} \longrightarrow \Delta; \Theta_0, \Theta}\, L\right]$$
$$\begin{array}{c}\mathcal{D}^L\\ \Pi_0, F; \longrightarrow; \Theta_0\end{array}$$

That is, the subderivation of $\mathcal{D}'$ below L is$\mathcal{D}^L$; the right subderivation above L (in which B is assumed) is $\mathcal{D}^B$; the left is $\mathcal{D}^A$.

We will use the inferences from $\mathcal{D}^L$ to construct alternative smaller derivations in place of $\mathcal{D}^A$ and $\mathcal{D}^B$. By Θ', indicate the minimal set of formulas required in addition to Θ_0 to span A^μ_X; by Π' indicate the minimal set of formulas required in addition to Π_0, F and A^μ_X to ensure that the pair given by Π_0, Π', F, A^μ_X and Θ_0, Θ' is balanced. (This is well-defined because the sequent $\Pi_0, F \longrightarrow \Theta_0$ is already spanned and balanced.) Now we can construct two new subderivations $\mathcal{D}'^A$ and $\mathcal{D}'^B$ given respectively as follows:

$$\begin{array}{c} \left[\dfrac{\begin{array}{c}\Pi' + A^\mu_X + \mathcal{D}^A + \Theta' \\ \Pi_0, F, \Pi, \Pi', A^\mu_X; \Gamma, A \vee B^\mu_X, A^\mu_X \longrightarrow \Delta; \Theta_0, \Theta, \Theta'\end{array}}{\Pi_0, F, \Pi, \Pi', A^\mu_X; \Gamma, A \vee B^\mu_X \longrightarrow \Delta; \Theta_0, \Theta, \Theta'}\text{decide} \right] \\ \Pi' + A^\mu_X + \mathcal{D}^L + \Theta' \\ \Pi_0, F, \Pi', A^\mu_X; \longrightarrow; \Theta_0, \Theta' \end{array}$$

$$\begin{array}{c} \left[\dfrac{\begin{array}{c}[\Pi' + B^\mu_X + \mathcal{D}^B + \Theta'] \\ \Pi_0, F, \Pi, \Pi', B^\mu_X; \Gamma, B \vee B^\mu_X, B^\mu_X \longrightarrow \Delta; \Theta_0, \Theta, \Theta'\end{array}}{\Pi_0, F, \Pi, \Pi', B^\mu_X; \Gamma, B \vee B^\mu_X \longrightarrow \Delta; \Theta_0, \Theta, \Theta'}\text{decide} \right] \\ \Pi' + B^\mu_X + \mathcal{D}^L + \Theta' \\ \Pi_0, F, \Pi', B^\mu_X; \longrightarrow; \Theta_0, \Theta' \end{array}$$

That is, we weaken $\mathcal{D}^A$ and $\mathcal{D}^B$ by global versions of the side expression of inference L throughout their *lowest blocks*; we apply a (decide) inference to obtain a new subderivation to substitute for the subderivation rooted at L in $\mathcal{D}^L$. We weaken by sufficient additional formulas globally in the *lowest blocks* to ensure that the end-sequents of these derivations are balanced and spanned.

Since we have changed only the lowest block here, and have ensured that this block remains isolated and canceled, we can now apply Lemma B.7 to obtain corresponding derivations $\mathcal{D}^A_I$ and $\mathcal{D}^B_I$ in which every block is canceled, linked, isolated, simple, balanced and spanned. In light of our first observation about the construction of Lemma B.7, we can see that the inferences of $\mathcal{D}^A$ are preserved up to the new (decide) inference. And in light of our third observation about the construction of Lemma B.7, given the unique inferences introducing Θ_0 and Π_0, this (decide) inference must be preserved in $\mathcal{D}^A_I$. Thus A^μ_X is linked in $\mathcal{D}^A_I$ and for analogous reasons B^μ_X is linked in $\mathcal{D}^B_I$. These derivations satisfy the induction hypothesis as deductions with fewer than n $(\vee \rightarrow)$ inferences; we can apply the induction hypothesis with A^μ_X and B^μ_X as the distinguished linked formulas to preserve. This results in SCLB derivations $\mathcal{A}$ and $\mathcal{B}$ with the same end-sequents as $\mathcal{D}'^A$ and $\mathcal{D}'^B$, in which every block is canceled, linked, isolated, simple and spanned, and in which respectively A^μ_X and B^μ_X are linked.

We need only one of $\mathcal{A}$ and $\mathcal{B}$ to reconstruct $\mathcal{D}'$ using blocking inferences. For example, we obtain a proof using $(\vee \rightarrow^B_L)$ by using $\mathcal{B}$ in place of $\mathcal{D}^B$ as schematized below:

$$\begin{array}{c} \left[\dfrac{\begin{array}{c}\mathcal{D}^A \\ \Pi_0, F, \Pi; \Gamma, A \vee B^\mu_X, A^\mu_X \longrightarrow \Delta; \Theta_0, \Theta\end{array} \qquad \begin{array}{c}\mathcal{B} \\ \Pi_0, F, \Pi', B^\mu_X \longrightarrow \Theta_0, \Theta'\end{array}}{\Pi_0, F, \Pi; \Gamma, A \vee B^\mu_X \longrightarrow \Delta; \Theta_0, \Theta} \vee \rightarrow^B_L \right] \\ \mathcal{D}^L \\ \Pi_0, F; \longrightarrow; \Theta_0 \end{array}$$

In a complementary way, we obtain a proof using $(\vee \rightarrow^B_R)$ by using $\mathcal{A}$ in place of $\mathcal{D}^A$ as schematized below:

$$\left[\dfrac{\dfrac{\mathcal{D}^B}{\Pi_0, F, \Pi; \Gamma, A \vee B_X^\mu, B_X^\mu \longrightarrow \Delta; \Theta_0, \Theta} \qquad \dfrac{\mathcal{A}}{\Pi_0, F, \Pi', A_X^\mu \longrightarrow \Theta_0, \Theta'}}{\Pi_0, F, \Pi; \Gamma, A \vee B_X^\mu \longrightarrow \Delta; \Theta_0, \Theta} \vee \to_L^B\right]$$
$$\mathcal{D}^L$$
$$\Pi_0, F; \longrightarrow; \Theta_0$$

Note that the root block is isolated in both cases, because we have added only as many formulas to Π' and Θ' as are necessary to obtain a balanced, spanned sequent; the remaining expressions originate in the end-sequent of the previous block, which we know was isolated. Thus, in both cases, we have blockwise eager derivations in which every block is canceled, isolated, simple, balanced and spanned, in which fewer than n $(\vee \to)$ inferences are used, and in which only the root block may fail to be linked. We thus need to apply the construction of Lemma B.7 again to ensure that the root block is linked. It is possible for the distinguished occurrence of F not to be linked in one of the resulting derivations, but not both. To see this, consider applying the construction of Lemma B.7 to $\mathcal{D}'$ itself, as a test: the result will be $\mathcal{D}'$ since $\mathcal{D}'$ is linked. Starting from $\mathcal{D}^A$ and $\mathcal{D}^B$ and axioms elsewhere, each inference in $\mathcal{D}'$ corresponds to an inference in the alternative derivations schematized above. We can argue by straightforward induction that no formula is linked in the reconstructed $\mathcal{D}'$ unless it is also linked in the one of the corresponding reconstructed alternative derivations. And F is linked in $\mathcal{D}'$.

Call the derivation in which F is linked $\mathcal{D}''$; we substitute $\mathcal{D}''$ for $\mathcal{D}'$ in $\mathcal{D}$. Since F remains linked in $\mathcal{D}''$, when we do so, we obtain a blockwise eager SCLU derivation with an appropriate end-sequent, with fewer original $(\vee \to)$ inferences, and in which every block remains canceled, linked, isolated, simple, balanced and spanned, and in which $(\vee \to)$ inferences lie at the root of SCLS derivations. Applying the induction hypothesis to the result gives the required SCLB derivation. □

B.3 Proof of Lemma 3.17

LEMMA B.9 (LEMMA 3.17). *Given a blockwise eager SCLB derivation $\mathcal{D}$, with end-sequent*

$$\Pi; \longrightarrow; \Theta$$

in which every block is linked, simple and spanned, we can construct a corresponding SCLP derivation of the same end-sequent in which every block remains linked.

PROOF. We again use an induction hypothesis stronger than the lemma. Given the conditions of the lemma, we construct an SCLP derivation $\mathcal{D}'$ in which four additional properties hold:

—the end-sequent of $\mathcal{D}'$ takes the form

$$\Pi; \Gamma' \longrightarrow \Delta'; \Theta$$

with $\Gamma' \subseteq \Gamma$ and $\Delta' \subseteq \Delta$;

—$\mathcal{D}'$ contains in each segment or block all and only the axioms of the corresponding segment or block of $\mathcal{D}$;

—whenever $\mathcal{D}'$ contains a sequent of the form

$$\Pi^*; \Gamma^* \to F; \Theta^*$$

F is the only right expression on which an axiom in that block is based; and

—whenever $\mathcal{D}'$ contains a sequent of the form

$$\Pi^*; F \to \Delta^*; \Theta^*$$

then F is the only left expression on which an axiom in that segment is based.

In the base case, $\mathcal{D}$ is

$$\Pi; \Gamma, A^{\mu}_{X} \longrightarrow B^{\nu}_{X}, \Delta; \Theta$$

and $\mathcal{D}'$ is

$$\Pi; A^{\mu}_{X} \longrightarrow B^{\nu}; \Theta$$

Supposing the claim true for proofs of height h, consider a proof $\mathcal{D}$ with height h+1. We consider cases for the different rules with which $\mathcal{D}$ could end.

The treatment of $(\to\wedge)$ is representative of the case analysis for the right rules other than $(\to>)$. $\mathcal{D}$ ends

$$\frac{\Pi; \longrightarrow A^{\mu}_{X}, A\wedge B^{\mu}_{X}, \Delta; \Theta \qquad \Pi; \longrightarrow B^{\mu}_{X}, A\wedge B^{\mu}_{X}, \Delta; \Theta}{\Pi; \longrightarrow A\wedge B^{\mu}_{X}, \Delta; \Theta} \to\wedge$$

(It is a consequence of Lemma B.3 that in the initial derivation there is an empty local area.) We simply apply the induction hypotheses to the immediate subderivations. If the resulting derivations end with (restart), consider the immediate subderivation of the results, otherwise consider the results themselves. These derivations end

$$\Pi; \longrightarrow C; \Theta$$
$$\Pi; \longrightarrow D; \Theta$$

We must have $C = A^{\mu}_{X}$; we know from the structure of $\mathcal{D}$ that A^{μ}_{X} is linked, and A^{μ}_{X} could not be linked in $\mathcal{D}$ unless $C = A^{\mu}_{X}$ since $\mathcal{D}'$ shows that all of the axioms in $\mathcal{D}$ derive from C. For the same reason $D = B^{\mu}_{X}$. So we can combine the resulting proofs by an $(\to\wedge)$ inference to give the needed $\mathcal{D}'$.

The case of $(\to>)$ proceeds similarly, but relies on an additional observation. $\mathcal{D}$ ends

$$\frac{\begin{array}{c}\mathcal{D}_1\\ \Pi, A^{\mu\eta}_{X,\mu\eta}; \longrightarrow \Delta, A >_i B^{\mu}_{X}; B^{\mu\eta}_{X,\mu\eta}, \Theta\end{array}}{\Pi; \longrightarrow \Delta, A >_i B^{\mu}_{X}; \Theta} \to>$$

We apply the induction hypothesis to $\mathcal{D}_1$ and eliminate any final (restart) inference. This gives us a derivation $\mathcal{D}'_1$ of

$$\Pi, A^{\mu\eta}_{X,\mu\eta}; \longrightarrow E; B^{\mu\eta}_{X,\mu\eta}, \Theta$$

If we know that the B-side expression of this inference is linked in this block, then we can conclude, as before, that E is an occurrence of the expression $B^{\mu\eta}_{X,\mu\eta}$. We show this as follows. We know from the structure of $\mathcal{D}$ only that *one* of the A-expression and the B-expression must be linked. However, it is straightforward to show that no left expression $A^{\mu\eta}_{X,\mu\eta}$ is linked in an SCLP derivation with a local goal C^{ν}_{Y} unless $\mu\eta$ is a prefix of ν. (The argument is a straightforward variant for example of [Stone 1999, Lemma 2].) Since $\mathcal{D}$ is simple and spanned, η must be new; $B^{\mu\eta}_{X,\mu\eta}$ is the only expression whose associated path term has $\mu\eta$ as a prefix.

Thus, we construct $\mathcal{D}'$ using an SCLP inference as

$$\dfrac{\begin{array}{c}\mathcal{D}_1'\\ \Pi, A^{\mu\eta}_{X,\mu\eta}; \longrightarrow B^{\mu\eta}_{X,\mu\eta}; B^{\mu\eta}_{X,\mu\eta}, \Theta\end{array}}{\Pi; \longrightarrow A >_i B^{\mu}_{X}; \Theta} \to>$$

Now suppose $\mathcal{D}$ ends in a left rule other than $(\supset\to^S)$ or $(\vee \to^B)$. We take $(\wedge \to)$ as a representative case; then $\mathcal{D}$ is:

$$\dfrac{\begin{array}{c}\mathcal{D}_1\\ \Pi; \Gamma, A \wedge B^{\mu}_{X}, A^{\mu}_{X}, B^{\mu}_{X} \longrightarrow \Delta; \Theta\end{array}}{\Pi; \Gamma, A \wedge B^{\mu}_{X} \longrightarrow \Delta; \Theta} \wedge \to$$

Apply the induction hypothesis to $\mathcal{D}_1$. If the result ends in a (decide) inference, let $\mathcal{D}_1'$ be the immediate subderivation of the result; otherwise let $\mathcal{D}_1'$ be the result itself. $\mathcal{D}_1'$ is an SCLP derivation with an end-sequent of the form:

$$\Pi; E \longrightarrow F; \Theta$$

E must be a side expression of the inference in question, here $(\wedge \to)$; otherwise the corresponding inference could not have been linked in $\mathcal{D}$. One of the inference figures $(\wedge \to_L)$ and $(\wedge \to_R)$ must apply depending on which side expression E is. For concrete illustration, we suppose E is A^{μ}_{X}; then we construct $\mathcal{D}'$ as:

$$\dfrac{\begin{array}{c}\mathcal{D}_1'\\ \Pi; A^{\mu}_{X} \longrightarrow F; \Theta\end{array}}{\Pi; A \wedge B^{\mu}_{X} \longrightarrow F; \Theta} \wedge \to_L$$

Next, we suppose $\mathcal{D}$ ends in $(\supset\to^S)$, as follows:

$$\dfrac{\begin{array}{c}\mathcal{D}_1\\ \Pi; \longrightarrow A^{\mu}_{X}, \Delta; \Theta\end{array} \qquad \begin{array}{c}\mathcal{D}_2\\ \Pi; \Gamma, A \supset B^{\mu}_{X}, B^{\mu}_{X} \longrightarrow \Delta; \Theta\end{array}}{\Pi; \Gamma, A \supset B^{\mu}_{X} \longrightarrow \Delta; \Theta} \supset\to^S$$

We begin by applying the induction hypothesis to the subderivation $\mathcal{D}_1$. After stripping off any (restart), we obtain an SCLP derivation $\mathcal{D}_1$ with end-sequent

$$\Pi; \longrightarrow C; \Theta$$

By the usual linking argument, the expression C must be identical to A^{μ}_{X}. We then apply the induction hypothesis also to the right subderivation. Again, after stripping off any (decide), we get an SCLP derivation $\mathcal{D}_2$ with end-sequent

$$\Pi; D \longrightarrow E; \Theta$$

By the usual linking argument, D must in fact be identical to B^{μ}_{X}. Thus we obtain the needed $\mathcal{D}'$ by combining the two derivations by the SCLP $(\supset\to)$ rule:

$$\dfrac{\begin{array}{c}\mathcal{D}_1'\\ \Pi; \longrightarrow A^{\mu}_{X}; \Theta\end{array} \qquad \begin{array}{c}\mathcal{D}_2'\\ \Pi; B^{\mu}_{X} \longrightarrow E; \Theta\end{array}}{\Pi; A \supset B^{\mu}_{X} \longrightarrow E; \Theta} \supset\to$$

Finally, for $(\vee \rightarrow^B)$, we consider the representative case of $\mathcal{D}$ as schematized below:

$$\dfrac{\begin{array}{c}\mathcal{D}_1\\ \Pi;\Gamma,A^\mu_X \longrightarrow \Delta;\Theta\end{array} \qquad \begin{array}{c}\mathcal{D}_2\\ \Pi',B^\mu_X; \longrightarrow ;\Theta'\end{array}}{\Pi;\Gamma,A\vee B^\mu_X \longrightarrow \Delta;\Theta} \vee \rightarrow^B_L$$

We begin by applying the induction hypothesis to $\mathcal{D}_1$, the subderivation in the current block; if necessary, we strip off any initial (decide) inference, obtaining $\mathcal{D}'_1$ with an end-sequent that by linking takes the form:

$$\Pi;A^\mu_X \longrightarrow E;\Theta$$

Next, we apply the induction hypothesis to the other subderivation. Since both local areas are empty in the input subderivation, they remain empty in the result subderivation: this gives $\mathcal{D}'_2$ with end-sequent:

$$\Pi',B^\mu_X; \longrightarrow ;\Theta'$$

The two subderivations can be recombined by the SCLP $(\vee \rightarrow_L)$ inference to obtain the needed $\mathcal{D}'$:

$$\dfrac{\begin{array}{c}\mathcal{D}'_1\\ \Pi;A^\mu_X \longrightarrow E;\Theta\end{array} \qquad \begin{array}{c}\mathcal{D}'_2\\ \Pi',B^\mu_X; \longrightarrow ;\Theta'\end{array}}{\Pi;A\vee B^\mu_X; \longrightarrow E;\Theta} \vee \rightarrow_L$$

□