A Quasipolynomial Cut-Elimination Procedure in Deep Inference via Atomic Flows and Threshold Formulae

Paola Bruscoli[1⋆], Alessio Guglielmi[1⋆⋆], Tom Gundersen[1⋆⋆⋆], and Michel Parigot[2†]

[1] University of Bath (UK) and LORIA & INRIA Nancy-Grand Est (France)
[2] Laboratoire PPS, UMR 7126, CNRS & Université Paris 7 (France)

Abstract. Jeřábek showed in 2008 that cuts in propositional-logic deep-inference proofs can be eliminated in quasipolynomial time. The proof is an indirect one relying on a result of Atserias, Galesi and Pudlák about monotone sequent calculus and a correspondence between this system and cut-free deep-inference proofs. In this paper we give a direct proof of Jeřábek's result: we give a quasipolynomial-time cut-elimination procedure in propositional-logic deep inference. The main new ingredient is the use of a computational trace of deep-inference proofs called atomic flows, which are both very simple (they trace only structural rules and forget logical rules) and strong enough to faithfully represent the cut-elimination procedure.

1 Introduction

Deep inference is a deduction framework (see [Gug07,BT01,Brü04]), where deduction rules apply arbitrarily deep inside formulae, contrary to traditional proof systems like natural deduction and sequent calculus, where deduction rules deal only with their outermost structure. This greater freedom is both a source of immediate technical difficulty and the promise, in the long run, of new powerful proof-theoretic methods. A general methodology allows to design deep-inference deduction systems having more symmetries and finer structural properties than sequent calculus. For instance, cut and identity become really dual of each other,

⋆ Supported by EPSRC grant EP/E042805/1 *Complexity and Non-determinism in Deep Inference* and by an ANR *Senior Chaire d'Excellence* titled *Identity and Geometric Essence of Proofs*.

⋆⋆ Supported by EPSRC grant EP/E042805/1 *Complexity and Non-determinism in Deep Inference* and by an ANR *Senior Chaire d'Excellence* titled *Identity and Geometric Essence of Proofs*.

⋆⋆⋆ Supported by an *Overseas Research Studentship* and a *Research Studentship*, both of the University of Bath, and by an ANR *Senior Chaire d'Excellence* titled *Identity and Geometric Essence of Proofs*.

† Supported by *Project INFER—Theory and Application of Deep Inference* of the *Agence Nationale de la Recherche*.

whereas they are only morally so in sequent calculus, and all structural rules can be reduced to their atomic form, whereas this is false for contraction in sequent calculus.

All usual logics have deep-inference deduction systems enjoying cut elimination (see [Gug] for a complete overview). The traditional methods of cut elimination of sequent calculus can be adapted to a large extent to deep inference, despite having to cope with a higher generality. New methods are also achievable. The standard proof system for propositional classical logic in deep inference is system SKS [BT01,Brü04]. Its cut elimination has been proved in several different ways [BT01,Brü04,GG08].

Recently, Jeřábek showed that cut elimination in SKS proofs can be done in quasipolynomial time [Jeř08], i.e., in time $n^{O(log(n))}$. The result is surprising because all known cut-elimination methods for classical-logic proof systems require exponential time, in particular for Gentzen's sequent calculus. Jeřábek obtained his result by relying on a construction over threshold functions by Atserias, Galesi and Pudlák, in the monotone sequent calculus [AGP02]. Note that, contrary to SKS, the monotone sequent calculus specifies a weaker logic than propositional logic because negation is not freely applicable.

The technique that Jeřábek adopts is indirect because cut elimination is performed over proofs in the sequent calculus, which are, in turn, related to deep-inference ones by polynomial simulations, originally studied in [Brü06] and [BG09].

In this paper we give a direct proof of Jeřábek's result: that is, we give a quasipolynomial-time cut-elimination procedure in propositional-logic deep inference, which, in addition to being internal, has a strong computational flavour. This proof uses two ingredients:

1. an adaptation of Atserias-Galesi-Pudlák technique to deep inference, which slightly simplifies the technicalities associated with the use of threshold functions; in particular, the formulae and derivations that we adopt are simpler than those in [AGP02];
2. a computational trace of deep-inference proofs called atomic flows, which are both very simple (they trace only structural rules and forget logical rules) and strong enough to faithfully represent cut elimination.

Atomic flows, which can be considered as specialised Buss flow graphs [Bus91], play a major role in designing and controlling the cut elimination procedure presented in this paper. They contribute to the overall clarification of the procedure, by reducing our dependency on syntax. The techniques developed via atomic flows tolerate variations in the proof system specification. In fact, their geometric nature makes them largely independent of syntax, provided that certain linearity conditions are respected (and this is usually achievable in deep inference).

The paper is self-contained. Sections 2 and 3 are devoted, respectively, to the necessary background on deep inference and atomic flows. Threshold functions and formulae are introduced in Sect. 5.

We normalise proofs in two steps, each of which has a dedicated section in the paper:

1. We transform any given proof into what we call its 'simple form'. No use is made of threshold formulae and no significant proof complexity is introduced. This is presented in Sect. 4, which constitutes a good exercise on deep inference and atomic flows.
2. In Sect. 6, we show the cut elimination step, starting from proofs in simple form. Here, threshold formulae play a major role.

Section 7 concludes the paper with some comments on future research directions.

2 Propositional Logic in Deep Inference: The SKS System

Formulae and Contexts

Two *logical constants*, $\mathtt{f}$ (false) and $\mathtt{t}$ (true) and a countable set of *propositional letters*, denoted by p and q, are given. A primitive *negation* $\bar{\cdot}$ is defined on propositional letters: to each propositional letter p is associated its negation $\bar{p}$. *Atoms* are propositional letters and their negation; they are denoted a, b, c, d and e. Negation is extended to the set of atoms by defining $\bar{\bar{p}} = p$, for each negated propositional letter $\bar{p}$. Being in classical logic, one can always exchange an atom with its negation: at the level of atoms, it doesn't matter which one is the propositional letter or its negation.
Formulae, denoted by A, B, C and D, are freely built from logical constants and atoms using *disjunction* and *conjunction*. The disjunction and conjunction of two formulae A and B are denoted respectively $[A \vee B]$ and $(A \wedge B)$: the different brackets have the only purpose of improving legibility. We usually omit external brackets of formulae and sometimes we omit superfluous brackets under associativity. Negation can be extended to arbitrary formulae in an obvious way using De Morgan's laws, but we do not need it in this paper. We write $A \equiv B$ for literal equality of formulae.
We denote (*formula*) *contexts*, i.e., formulae with a hole, by $K\{\ \}$; for example, if $K\{a\}$ is $b \wedge [a \vee c]$, then $K\{\ \}$ is $b \wedge [\{\ \} \vee c]$, $K\{b\}$ is $b \wedge [b \vee c]$ and $K\{a \wedge d\}$ is $b \wedge [(a \wedge d) \vee c]$.

Derivations and Proofs

An (*inference*) *rule* ρ is an expression $\rho\,\dfrac{A}{B}$, where the formulae A and B are called *premiss* and *conclusion*, respectively. In deep inference, rules are applied in arbitrary contexts: an (*inference*) *step* corresponding to rule ρ is an expression $\rho\,\dfrac{K\{C\}}{K\{D\}}$, where $K\{\ \}$ is a context and $\rho\,\dfrac{C}{D}$ is an instance of $\rho\,\dfrac{A}{B}$.
A *derivation*, Φ, *from* A (*premiss*) *to* B (*conclusion*) is a chain of inference steps

with A at the top and B at the bottom, and is usually indicated by $\Phi \Vdash_{\mathcal{S}}^{A}_{B}$, where $\mathcal{S}$ is the name of the deduction system or a set of inference rules (we might omit Φ and $\mathcal{S}$); we also use the notation $\Phi : A \to B$. Sometimes we group $n \geq 0$ inference steps of the same rule ρ together into one step, and we label the step with $n \times \rho$. Besides Φ, we denote derivations with Ψ.
A *proof*, often denoted by Π, is a derivation with premiss t.
The *size* $|A|$ of a formula A is the number of unit and atom occurrences appearing in it. The *size* $|\Phi|$ of a derivation Φ is the sum of the sizes of the formulae occurring in it. The *length* of a derivation is the number of inference steps applied in the derivation. The *width* of a derivation is the maximal size of the formulae occurring in it.

Substitution

By $A\{a_1/B_1, \ldots, a_h/B_h\}$, we denote the operation of simultaneously substituting formulae B_1, ..., B_h into all the occurrences of the atoms a_1, ..., a_h in the formula A, respectively. By defining the substitution at the level of atoms, where atoms and their negation are equal citizen, we mean that the substitution to the occurrences of an atom doesn't touch the occurrences of its negation. Often, we only substitute certain occurrences of atoms: there will be no ambiguity because this is done in the context of atomic flows, where occurrences are distinguished with superscripts. The notion of substitution is extended to derivations in the natural way.

Inference Rules of SKS

Structural inference rules:

$$\mathsf{ai}{\downarrow}\,\frac{\mathsf{t}}{a \vee \bar{a}} \qquad \mathsf{aw}{\downarrow}\,\frac{\mathsf{f}}{a} \qquad \mathsf{ac}{\downarrow}\,\frac{a \vee a}{a}$$

identity (interaction) *weakening* *contraction*

$$\mathsf{ai}{\uparrow}\,\frac{a \wedge \bar{a}}{\mathsf{f}} \qquad \mathsf{aw}{\uparrow}\,\frac{a}{\mathsf{t}} \qquad \mathsf{ac}{\uparrow}\,\frac{a}{a \wedge a}$$

cut (cointeraction) *coweakening* *cocontraction*

Logical inference rules:

$$\mathsf{s}\,\frac{A \wedge [B \vee C]}{(A \wedge B) \vee C} \qquad \mathsf{m}\,\frac{(A \wedge B) \vee (C \wedge D)}{[A \vee C] \wedge [B \vee D]}$$

switch *medial*

There are also *equality* rules $=\dfrac{C}{D}$, for C and D on opposite sides in one of the following equations:

$$\begin{array}{cc} A \vee B = B \vee A & A \vee \mathsf{f} = A \\ A \wedge B = B \wedge A & A \wedge \mathsf{t} = A \\ [A \vee B] \vee C = A \vee [B \vee C] & \mathsf{t} \vee \mathsf{t} = \mathsf{t} \\ (A \wedge B) \wedge C = A \wedge (B \wedge C) & \mathsf{f} \wedge \mathsf{f} = \mathsf{f} \end{array} \quad . \tag{1}$$

Conventions

(a) In derivations we freely use equality rules without mentioning them. For instance

$$\mathsf{ai}{\downarrow}\dfrac{\dfrac{\dfrac{A}{A \wedge \mathsf{t}}=}{\mathsf{t} \wedge A}=}{[p \vee \bar{p}] \wedge A} \quad \text{is written} \quad \mathsf{ai}{\downarrow}\dfrac{A}{[p \vee \bar{p}] \wedge A} \ .$$

(b) The structural rules have been given in atomic form in SKS. This is possible because in deep inference the general form of the structural rules, given below, is derivable from their atomic form, moreover it is derivable with a polynomial cost.

$$\begin{array}{ccc} \mathsf{i}{\downarrow}\dfrac{\mathsf{t}}{A \vee \bar{A}} & \mathsf{w}{\downarrow}\dfrac{\mathsf{f}}{A} & \mathsf{c}{\downarrow}\dfrac{A \vee A}{A} \\ \mathsf{i}{\uparrow}\dfrac{A \wedge \bar{A}}{\mathsf{f}} & \mathsf{w}{\uparrow}\dfrac{A}{\mathsf{t}} & \mathsf{c}{\uparrow}\dfrac{A}{A \wedge A} \end{array} \quad .$$

We will freely use a nonatomic rule instance to stand for some derivation in SKS that derives that instance.

Operations on Derivations

Inference rules being applicable in any context, given a context $K\{\ \}$ and a derivation $\Phi : A \to B$, one can form a derivation $K\{\Phi\} : K\{A\} \to K\{B\}$ by adding the context $K\{\ \}$ at each inference step of the derivation. Given two derivations $\Phi : A \to B$ and $\Psi : C \to D$, one can form in this way the derivations $\Phi \wedge C : A \wedge C \to B \wedge C$ and $B \wedge \Psi : B \wedge C \to B \wedge D$. Then one can put one after the other to get a derivation $\begin{array}{c} A \wedge C \\ \| \\ B \wedge D \end{array}$ of $B \wedge D$ from $A \wedge C$; we denote by $\Phi \wedge \Psi : A \wedge C \to B \wedge D$ this derivation which consists in making Φ and then Ψ. In the same way, one can get a derivation $\Phi \vee \Psi : A \vee C \to B \vee D$ of $B \vee D$ from $A \vee C$. We will freely use these constructions throughout the paper.

3 Atomic Flows

Atomic flows, which have been introduced in [GG08], are, essentially, specialised Buss flow graphs [Bus91]. They are particular directed graphs associated with SKS derivations: every derivation yields one atomic flow obtained by tracing the atoms (propositional letters and their negation) occurrences in the derivation. More precisely, one traces the behaviour of these occurrences through the structural rules: creation / destruction / duplication. No information about instances of logical rules is kept, only structural rules play a role and, as a consequence, infinitely many derivations correspond to each atomic flow. As shown in [GG08], it turns out that atomic flows contain sufficient structure to control cut elimination procedures, providing in particular induction measures that can be used to ensure termination. Such cut-elimination procedures require exponential time on the size of the derivation to be normalised. In the present work, we improve the complexity of cut elimination to quasipolynomial time, using in addition threshold formulae, which are independent from the given proof.

Atomic Flow Associated to a Derivation

We first index occurrences of atoms in derivations with natural numbers in such a way that:

- different occurrences of atoms in formulae have different indexes;
- indexes are preserved by logical rules and the context part of structural rules;
- in each instance of a structural rule, active occurrences of atoms have different indexes; for example an instance of the contraction rule becomes $\mathsf{ac}{\downarrow}\dfrac{a^{\mathbf{1}} \vee a^{\mathbf{2}}}{a^{\mathbf{3}}}$.

We associate inductively (say, in a top-down manner) to each derivation with indexed occurrences of atoms an atomic flow as follows:

- to a formula $A(a^{\mathbf{1}}, \ldots, a^{n})$ with exactly n occurrences of atoms, the following flow, consisting of n edges, is associated:

$$\mathbf{1}\,|\cdots|\,\mathbf{n}\ ;$$

- the logical rules and the context part of structural rule do not change the flow;
- each instance of a structural rule adds a *vertex*, whose incident edges correspond to active occurrences of atoms in the rule; the association of vertices to structural rules is illustrated below:

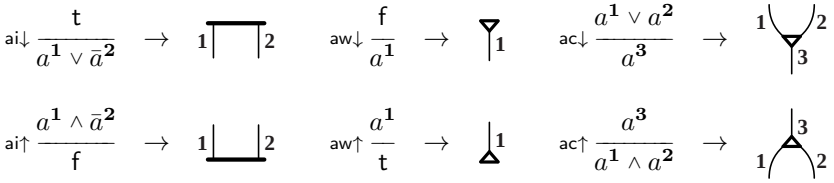

The left-hand side of each arrow shows an instance of a structural rule, whose atom occurrences are labelled by small numerals. Correspondingly, the right-hand side of the same arrow, shows the vertex associated to the given rule instance: the labelling of incident edges respects the labelling of atom occurrences. In a top-down inductive reading of the proof, the upper edges of the vertices are meant to be associated to the already defined flow and the lower edges are new ones. Moreover, we qualify each vertex according to the rule it corresponds to: for example, in a given atomic flow, we might talk about a *contraction vertex*, or a *cut vertex*, and so on. Instead of small numerals, sometimes we use ϵ or ι to label edges (as well as atom occurrences), but we do not always use labels.

All edges are directed, but we do not explicitly show the orientation. Instead, we consider it as implicitly given by the way we draw them, i.e., edges are oriented along the vertical direction. So, the vertices corresponding to dual rules, are mutually distinct: for example, an identity vertex and a cut vertex are different because the orientation of their edges is different. On the other hand, the horizontal direction plays no role in distinguishing atomic flows; this corresponds to commutativity of logical relations. Here are for instance three representations of the same flow:

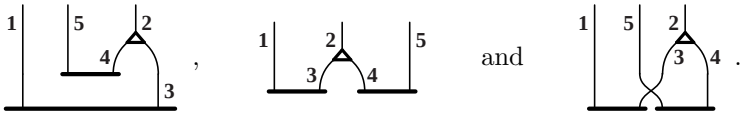

It should be noted that atomic flows built from derivations have no directed cycles and bear a natural polarity assignment (corresponding to atoms versus negated atoms in the derivation), that is a mapping of each edge to an element of $\{-, +\}$, such that the two edges of each identity or cut vertex map to different values and the three edges of each contraction or cocontraction vertex map to the same value. We denote atomic flows by ϕ and ψ.

Examples of Atomic Flows Associated to Derivations

$$\mathsf{ac}{\downarrow}\,\frac{a^{\mathbf{1}} \wedge \left[b^{\mathbf{2}} \vee \left[a^{\mathbf{3}} \vee a^{\mathbf{4}}\right]\right]}{a^{\mathbf{1}} \wedge \left[b^{\mathbf{2}} \vee a^{\mathbf{5}}\right]} \quad \rightarrow \quad$$

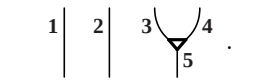

$$
\begin{array}{l}
\mathsf{ai}{\downarrow}\ \dfrac{\left(a^{\mathbf{1}} \wedge \left[\bar{a}^{\mathbf{3}} \vee \mathsf{t}\right]\right) \wedge \bar{a}^{\mathbf{8}}}{\left(a^{\mathbf{1}} \wedge \left[\bar{a}^{\mathbf{3}} \vee \left[\bar{a}^{\mathbf{4}} \vee a^{\mathbf{5}}\right]\right]\right) \wedge \bar{a}^{\mathbf{8}}} \\
=\ \dfrac{\left(a^{\mathbf{1}} \wedge \left[\left[\bar{a}^{\mathbf{3}} \vee \bar{a}^{\mathbf{4}}\right] \vee a^{\mathbf{5}}\right]\right) \wedge \bar{a}^{\mathbf{8}}}{\left[\left(a^{\mathbf{1}} \wedge \left[\bar{a}^{\mathbf{3}} \vee \bar{a}^{\mathbf{4}}\right]\right) \vee a^{\mathbf{5}}\right] \wedge \bar{a}^{\mathbf{8}}} \\
\mathsf{s}\ \dfrac{\left[\left(a^{\mathbf{1}} \wedge \left[\bar{a}^{\mathbf{3}} \vee \bar{a}^{\mathbf{4}}\right]\right) \vee a^{\mathbf{5}}\right] \wedge \bar{a}^{\mathbf{8}}}{\left[\left(a^{\mathbf{1}} \wedge \bar{a}^{\mathbf{2}}\right) \vee a^{\mathbf{5}}\right] \wedge \bar{a}^{\mathbf{8}}} \\
\mathsf{ac}{\downarrow}\ \dfrac{\left[\left(a^{\mathbf{1}} \wedge \bar{a}^{\mathbf{2}}\right) \vee a^{\mathbf{5}}\right] \wedge \bar{a}^{\mathbf{8}}}{\left[\mathsf{f} \vee a^{\mathbf{5}}\right] \wedge \bar{a}^{\mathbf{8}}} \\
\mathsf{ai}{\uparrow}\ \dfrac{\left[\mathsf{f} \vee a^{\mathbf{5}}\right] \wedge \bar{a}^{\mathbf{8}}}{a^{\mathbf{5}} \wedge \bar{a}^{\mathbf{8}}} \\
=\ \dfrac{a^{\mathbf{5}} \wedge \bar{a}^{\mathbf{8}}}{\left(a^{\mathbf{6}} \wedge a^{\mathbf{7}}\right) \wedge \bar{a}^{\mathbf{8}}} \\
\mathsf{ac}{\uparrow}\ \dfrac{\left(a^{\mathbf{6}} \wedge a^{\mathbf{7}}\right) \wedge \bar{a}^{\mathbf{8}}}{a^{\mathbf{6}} \wedge \left(a^{\mathbf{7}} \wedge \bar{a}^{\mathbf{8}}\right)} \\
=\ \dfrac{a^{\mathbf{6}} \wedge \left(a^{\mathbf{7}} \wedge \bar{a}^{\mathbf{8}}\right)}{a^{\mathbf{6}} \wedge \mathsf{f}} \\
\mathsf{ai}{\uparrow}
\end{array}
\quad \rightarrow \quad
$$

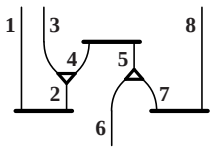

.

Abstract Notation of Atomic Flows

When certain details of a flow are not important, but only the vertex kinds and its upper and lower edges are, we can use boxes, labelled with all the vertex kinds that can appear in the flow they represent. For example:

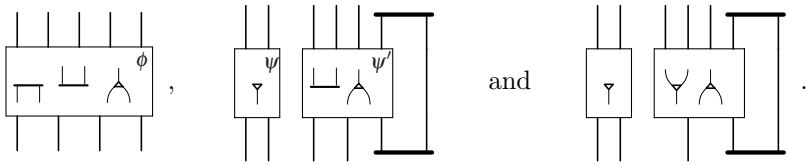

.

When no vertex labels appear on a box, we assume that the vertices in the corresponding flow can be any (so, it does not mean that there are no vertices in the flow).

We sometimes use a double line notation for representing multiple edges. For example, the following diagrams represent the same flow:

,

where $l, m \geq 0$; note that we use $\boldsymbol{\epsilon}_1^l$ to denote the vector $(\epsilon_1, \ldots, \epsilon_l)$. We might label multiple edges with either a vector of the associated atom occurrences in a derivation or one of the formulae the associated atom occurrences belong to in a derivation.

We extend the double line notation to collections of isomorphic flows. For example, for $l \geq 0$, the following diagrams represent the same flow:

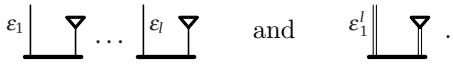

.

We observe that the flow of every SKS derivation can always be represented as follows:

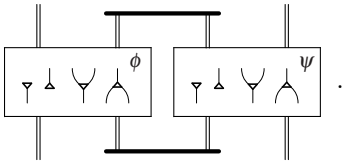

.

4 Normalisation Step 1: Simple Form

The first step in our normalisation procedure, defined here, consists in routine deep-inference manipulations, which are best understood in conjunction with atomic flows. For this reason, this section is a useful exercise for a reader who is not familiar with deep inference and atomic flows.

In Theorem 5 of this section, we show that every proof can be transformed into 'simple form'. Proofs in simple form are such that we can substitute formulae for all the atom occurrences that appear in cut instances, without substituting for atom occurrences that appear in the conclusion of the derivation. Of course, doing this would invalidate identity and cut instances, but in Sect. 6 we see how we can build a valid cut-free proof from the broken derivations obtained by substituing formulae into derivations in simple form.

We first show some standard deep-inference results. We will see how we can permute all the identity (resp., cut) rule instances to the top (resp., bottom) of a proof, without changing the atomic flow of the proof, and without significantly changing the size of the proof.

Lemma 1. *Given a context $K\{\ \}$ and a formula A, there exist derivations*

$$\begin{array}{c} A \wedge K\{\mathsf{t}\} \\ \| \{\mathsf{s}\} \\ K\{A\} \end{array} \quad \textit{and} \quad \begin{array}{c} K\{A\} \\ \| \{\mathsf{s}\} \\ K\{\mathsf{f}\} \vee A \end{array} \ ,$$

each of whose width is the size of $K\{A\}$ plus one and length is bounded by a polynomial in the size of $K\{\ \}$.

Proof. The result follows by structural induction on $K\{\ \}$: The base cases are:

$$= \frac{A \wedge \{\mathsf{t}\}}{\{A\}} \quad \text{and} \qquad = \frac{\{A\}}{\{\mathsf{f}\} \vee A} \ .$$

The inductive cases are

$$= \frac{\begin{array}{c} A \wedge (B \wedge K\{\mathsf{t}\}) \\ \| \{\mathsf{s}\} \\ B \wedge K\{A\} \end{array}}{B \wedge (A \wedge K\{\mathsf{t}\})} \ , \qquad \mathsf{s} \frac{\begin{array}{c} A \wedge [B \vee K\{\mathsf{t}\}] \\ \| \{\mathsf{s}\} \\ B \vee K\{A\} \end{array}}{B \vee (A \wedge K\{\mathsf{t}\})} \quad \text{and}$$

$$\mathsf{s}\,\dfrac{\begin{array}{c}B \wedge K\{A\}\\ \|\{\mathsf{s}\}\\ B \wedge [K\{\mathsf{f}\} \vee A]\end{array}}{(B \wedge K\{\mathsf{f}\}) \vee A}\ , \qquad =\dfrac{\begin{array}{c}B \vee K\{A\}\\ \|\{\mathsf{s}\}\\ B \vee [K\{\mathsf{f}\} \vee A]\end{array}}{[B \vee K\{\mathsf{f}\}] \vee A}\ .$$

□

Lemma 2. *Given a derivation $\Phi : A \to B$, with flow*

$$\phi \;=\; \begin{array}{c} A \;\; a_1^n \;\; \bar{a}_1^n \\ \boxed{\phi'} \\ b_1^m \;\; \bar{b}_1^m \;\; B \end{array}\ ,$$

there exists a derivation

$$\Psi \;=\; \dfrac{\begin{array}{c}\dfrac{A}{A \wedge [a_1 \vee \bar{a}_1] \wedge \cdots \wedge [a_n \vee \bar{a}_n]}\,n\times\mathsf{ai}{\downarrow}\\ \Psi' \|\\ \dfrac{(b_1 \wedge \bar{b}_1) \vee \cdots \vee (b_m \wedge \bar{b}_m) \vee B}{B}\,m\times\mathsf{ai}{\uparrow}\end{array}}{B}\ ,$$

for some atoms $a_1, \ldots, a_n, b_1, \ldots, b_m$ and some derivation Ψ', such that the flow of Ψ is ϕ, the flow of Ψ' is ϕ' and the size of Ψ is bounded by a polynomial in the size of Φ.

Proof. For every relevant interaction we perform the following transformation:

$$\mathsf{ai}{\downarrow}\,\dfrac{\begin{array}{c}A\\ \Phi' \|\\ K\{\mathsf{t}\}\end{array}}{\begin{array}{c}K\{a \vee \bar{a}\}\\ \Phi'' \|\\ B\end{array}} \quad \text{is transformed into} \quad \begin{array}{c}\dfrac{A}{[a \vee \bar{a}] \wedge A}\,\mathsf{ai}{\downarrow}\\ [a \vee \bar{a}] \wedge \Phi' \|\\ [a \vee \bar{a}] \wedge K\{\mathsf{t}\}\\ \|\{\mathsf{s}\}\\ K\{a \vee \bar{a}\}\\ \Phi'' \|\\ B\end{array}\ ,$$

which is possible by Lemma 1. Instances of cut rules can be dealt with in a symmetric way.

Each transformation increases the width of the derivation by a constant and increases the length by at most a polynomial in the width of the derivation. Hence, the size of Ψ is bounded by a polynomial in the size of Φ. □

We now show how to extend substitutions from formulae to derivations. Using atomic flows, we single out some atom occurrences that we substitute for. Substitutions play a crucial role in Theorem 5 and in Theorem 11. It is important to notice that a substitution only copies atomic flows, it does not introduce new vertices; and that the cost of substitution is polynomial.

Lemma 3. *Given a derivation $\Phi : A \to B$, let its associated flow have shape*

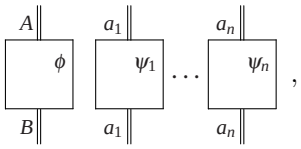

such that, for $1 \leq i \leq n$, all the edges of ψ_i are mapped to from occurrences of a_i, then, for any formulae $C_1, \ldots, C_n$ there exists a derivation

$$\Psi \;=\; \begin{array}{c} A\{a_1^{\psi_1}/C_1, \ldots, a_n^{\psi_n}/C_n\} \\ \| \\ B\{a_1^{\psi_1}/C_1, \ldots, a_n^{\psi_n}/C_n\} \end{array} ,$$

whose flow is

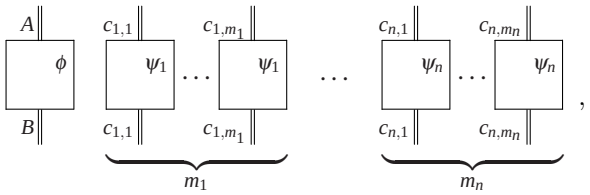

where, for every $1 \leq i \leq n$, the atom occurrences of C_i are $c_{i,1}, \ldots, c_{i,m_i}$; moreover, the size of Ψ is bounded by a polynomial in the size of Φ and, for each $1 \leq i \leq n$, the size of C_i.

Proof. We sketch the proof: For each $1 \leq i \leq n$, we can proceed by structural induction on C_i and then on ψ_i. For the two cases of $C_i \equiv D \vee E$ and $C_i \equiv D \wedge E$ we have to consider, for each vertex of ψ_i, one of the following situations (notice that ψ_i can not contain interaction or cut vertices):

$$\mathsf{w}\!\downarrow \frac{\mathsf{w}\!\downarrow \dfrac{\mathsf{f}}{E}}{D \vee E} \quad , \quad \mathsf{w}\!\downarrow \frac{\mathsf{w}\!\downarrow \dfrac{\mathsf{f}}{\mathsf{f} \wedge E}}{D \wedge E} \quad , \quad \mathsf{c}\!\downarrow \frac{\mathsf{c}\!\downarrow \dfrac{D \vee D \vee E \vee E}{D \vee D \vee E}}{D \vee E} \quad , \quad \mathsf{c}\!\downarrow \frac{\mathsf{m}\,\dfrac{(D \wedge E) \vee (D \wedge E)}{[D \vee D] \wedge [E \vee E]}}{\mathsf{c}\!\downarrow \dfrac{[D \vee D] \wedge E}{D \wedge E}} \quad ,$$

and their dual ones. □

Notation 4. When we write $\Phi\{a_1^{\psi_1}/C_1, \ldots, a_n^{\psi_n}/C_n\}$, we mean the derivation Ψ obtained in the proof of Lemma 3.

We now present the main result of this section. We show that any derivation can be transformed into a derivation whose atomic flow is on what we call 'simple form'. Referring to the second flow in Theorem 5, we observe that we could substitue for the atom occurrences corresponding to the rightmost copy of ϕ without substituting for any atom occurrence appearing in the conclusion of the proof. This is one of the two main ingredients in our normalisation procedure.

Theorem 5. *Given a proof Φ of A, with flow*

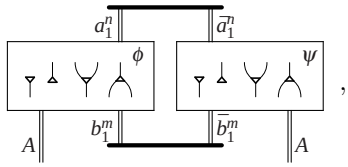

there exists a proof Ψ of A, with flow

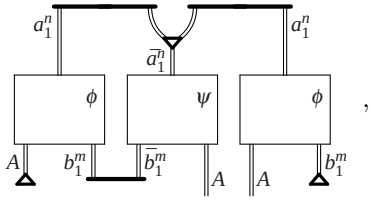

such that the size of Ψ is bounded by a polynomial in the size of Φ.

Proof. Consider the derivation

$$\begin{array}{c} \left[a_1^\phi \vee \bar{a}_1^\psi\right] \wedge \cdots \wedge \left[a_n^\phi \vee \bar{a}_n^\psi\right] \\ \Phi' \| \\ \left(b_1^\phi \wedge \bar{b}_1^\psi\right) \vee \cdots \vee \left(b_m^\phi \wedge \bar{b}_m^\psi\right) \vee A \end{array} ,$$

with atomic flow

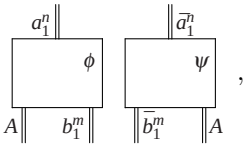

which exists and whose size is bounded by a polynomial in the size of Φ by Lemma 2. Let $a_1, \ldots, a_n, b_1, \ldots, b_m, c_1, \ldots, c_l$ be all the atoms whose occurrences are mapped to edges in ϕ and let

$$\begin{aligned} \sigma = \{&a_1^\phi / (a_1 \wedge a_1), \ldots, a_n^\phi / (a_n \wedge a_n), b_1^\phi / (b_1 \wedge b_1), \ldots, \\ &b_m^\phi / (b_m \wedge b_m), c_1^\phi / (c_1 \wedge c_1), \ldots, c_l^\phi / (c_l \wedge c_l)\} \ . \end{aligned}$$

We then construct Ψ:

$$\begin{array}{c} \mathsf{t} \\ 2n\times\mathsf{ai}\downarrow \dfrac{}{\left[a_1 \vee \bar{a}_1\right] \wedge \left[a_1 \vee \bar{a}_1\right] \wedge \cdots \wedge \left[a_n \vee \bar{a}_n\right] \wedge \left[a_n \vee \bar{a}_n\right]} \\ 2n\times\mathsf{s} \dfrac{}{\left[(a_1 \wedge a_1) \vee \bar{a}_1 \vee \bar{a}_1\right] \wedge \cdots \wedge \left[(a_n \wedge a_n) \vee \bar{a}_n \vee \bar{a}_n\right]} \\ n\times\mathsf{ac}\downarrow \dfrac{}{\left[(a_1 \wedge a_1) \vee \bar{a}_1\right] \wedge \cdots \wedge \left[(a_n \wedge a_n) \vee \bar{a}_n\right]} \\ \Phi'\sigma \| \\ \left((b_1 \wedge b_1) \wedge \bar{b}_1\right) \vee \cdots \vee \left((b_m \wedge b_m) \wedge \bar{b}_m\right) \vee A\sigma \\ \|\{\mathsf{aw}\uparrow\} \\ m\times\mathsf{ai}\uparrow \dfrac{\left(b_1 \wedge \bar{b}_1\right) \vee \cdots \vee \left(b_m \wedge \bar{b}_m\right) \vee A}{A} \end{array} ,$$

with the required atomic flow, where, by Lemma 3, the derivation $\Phi'\sigma$ exists and its size is bounded by a polynomial in the size of Φ'. □

5 Threshold Formulae

Threshold formulae realise boolean threshold functions, which are defined as boolean functions that are true if and only if at least k of n inputs are true (see [Weg87] for a thorough reference on threshold functions).

There are several ways of encoding threshold functions into formulae, and the problem is to find, among them, an encoding that allows us to obtain Theorem 10. Efficiently obtaining the property stated in Theorem 10 crucially depends also on the proof system we adopt.

The following class of threshold formulae, which we found to work for system SKS, is a simplification of the one adopted in [AGP02].

In the rest of this paper, whenever we have a sequence of atoms $a_1, \ldots, a_n$, we will assume, without loss of generality, that n is a power of two.

Definition 6. *For every $n = 2^m$, with $m \geq 0$, and $k \geq 0$, we define the operator θ_k^n inductively as follows:*

$$\theta_k^n(a_1, \ldots, a_n) = \begin{cases} \mathsf{t} & \text{if } k = 0 \\ \mathsf{f} & \text{if } k > n \\ a_1 & \text{if } n = k = 1 \\ \bigvee_{\substack{i+j=k \\ 0 \leq i,j \leq n/2}} \left(\theta_i^{n/2}(a_1, \ldots, a_{n/2}) \wedge \theta_j^{n/2}(a_{n/2+1}, \ldots, a_n) \right) & \text{otherwise.} \end{cases}$$

For any n atoms $a_1, \ldots, a_n$, we call $\theta_k^n(a_1, \ldots, a_n)$ the threshold formula at level k *(*with respect to $a_1, \ldots, a_n$*).*

The size of the threshold formulae dominates the cost of the normalisation procedure, so, we evaluate their size.

Lemma 7. *For any $n = 2^m$, with $m \geq 0$, and $k \geq 0$ the size of $\theta_k^n(a_1, \ldots, a_n)$ has a quasipolynomial bound in n.*

Proof. We show that the size of $\theta_k^n(a_1, \ldots, a_n)$ is bounded by $n^{2 \log n}$. We reason by induction on n; the case $n = 1$ trivially holds. For $n > 1$, we consider that the size of $\theta_k^n(a_1, \ldots, a_n)$ is bounded by

$$\sum_{\substack{i+j=k \\ 0 \leq i \leq n \\ 0 \leq j \leq n}} 2n/2^{2 \log n/2} \quad .$$

We then have

$$\sum_{\substack{i+j=k \\ 0 \leq i \leq n \\ 0 \leq j \leq n}} 2n/2^{2 \log n/2} \leq \sum_{\substack{i+j=n/2 \\ 0 \leq i \leq n \\ 0 \leq j \leq n}} 2n/2^{2 \log n/2} \leq (n+2)n/2^{2 \log n/2} \quad ,$$

and since $n+2 \leq n^2$ and $n/2 < n$, we have

$$(n+2)n/2^{2\log n/2} \leq n^2 n^{2\log n/2} = n^{2\log n - 2\log 2 + 2} = n^{2\log n} \ ,$$

as required. □

Lemma 8. *For any $n = 2^m$, with $m \geq 0$, $k \geq 0$ and $1 \leq i \leq n$, there exists a derivation*

$$\Gamma_k^i \ = \ \begin{array}{c} \theta_k^n(a_1, \ldots, a_n)\{a_i/\mathsf{f}\} \\ \| \{\mathsf{aw}\downarrow, \mathsf{aw}\uparrow\} \\ \theta_{k+1}^n(a_1, \ldots, a_n)\{a_i/\mathsf{t}\} \end{array} \ ,$$

whose size has a quasipolynomial bound in n.

Proof. The result follows by Lemma 7 and structural induction on Definition 6. It is worth noting that both the premiss and the conclusion of Γ_k^i are logically equivalent to $\theta_k^{n-1}(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$. □

Lemma 9. *Given a formula A and an atom a that occurs in A, there exist derivations* $\begin{array}{c} a \wedge A\{a/\mathsf{t}\} \\ \| \{\mathsf{ac}\uparrow, \mathsf{s}\} \\ A \end{array}$ *and* $\begin{array}{c} A \\ \| \{\mathsf{ac}\downarrow, \mathsf{s}\} \\ A\{a/\mathsf{f}\} \vee a \end{array}$ *such that their sizes are both bounded by a polynomial in the size of A.*

Proof. The result follows by induction on the number of occurrences of a in A, and Lemma 1. □

We now present the main result of this section. We show that, using threshold functions, we are able to deduce a conjunction of disjunctions from a disjunction of (slightly different) conjunctions. This construction is based on (seen top-down) contractions meeting cocontrations, and can be thought of as a generalisation of the simple sharing mechanism that allows us to deduce $a \wedge \cdots \wedge a$ from $a \vee \cdots \vee a$.

In Theorem 11 we will see how using this sharing mechanism allows us to glue together several 'broken' derivations in order to build a cut-free proof.

Theorem 10. *Let, for some $n = 2^m$ with $m \geq 0$, a_1, …, a_n be distinct atoms. Then, for every $1 \leq k \leq n+1$, there exists a derivation*

$$\Gamma_k \ = \ \begin{array}{c} \big(a_1 \wedge \theta_{k-1}^n(a_1, \ldots, a_n)\{a_1/\mathsf{f}\}\big) \vee \cdots \vee \big(a_n \wedge \theta_{k-1}^n(a_1, \ldots, a_n)\{a_n/\mathsf{f}\}\big) \\ \Big\| \mathsf{SKS} \setminus \{\mathsf{ai}\downarrow, \mathsf{ai}\uparrow\} \\ [a_1 \vee \theta_k^n(a_1, \ldots, a_n)\{a_1/\mathsf{f}\}] \wedge \cdots \wedge [a_n \vee \theta_k^n(a_1, \ldots, a_n)\{a_n/\mathsf{f}\}] \end{array} \ ,$$

such that the size of Γ_k has a quasipolynomial bound in n.

Proof. For $1 \leq k \leq n+1$, we construct:

$$\Gamma_k \quad = \quad \begin{array}{c} \big(a_1 \wedge \theta^n_{k-1}(a_1,\ldots,a_n)\{a_1/\mathsf{f}\}\big) \vee \cdots \vee \big(a_n \wedge \theta^n_{k-1}(a_1,\ldots,a_n)\{a_n/\mathsf{f}\}\big) \\ \Gamma^1_k \vee \cdots \vee \Gamma^n_k \,\|\{\mathsf{aw}\downarrow,\mathsf{aw}\uparrow\} \\ (a_1 \wedge \theta^n_k(a_1,\ldots,a_n)\{a_1/\mathsf{t}\}) \vee \cdots \vee (a_n \wedge \theta^n_k(a_1,\ldots,a_n)\{a_n/\mathsf{t}\}) \\ \Phi_1 \,\|\{\mathsf{ac}\uparrow,\mathsf{s}\} \\ \bigvee_{1\leq i\leq n} \theta^n_k(a_1,\ldots,a_n) \\ \|\{\mathsf{c}\downarrow\} \\ \theta^n_k(a_1,\ldots,a_n) \\ \|\{\mathsf{c}\uparrow\} \\ \bigwedge_{1\leq i\leq n} \theta^n_k(a_1,\ldots,a_n) \\ \Phi_2 \,\|\{\mathsf{ac}\downarrow,\mathsf{s}\} \\ [a_1 \vee \theta^n_k(a_1,\ldots,a_n)\{a_1/\mathsf{f}\}] \wedge \cdots \wedge [a_n \vee \theta^n_k(a_1,\ldots,a_n)\{a_n/\mathsf{f}\}] \end{array} ,$$

where Φ_1 and Φ_2 exist by Lemma 9 and, for $1 \leq i \leq n$, Γ^i_k exists by Lemma 8. The size of Γ_k is quasipolynomial in n, by Lemma 8 and Lemma 9. □

6 Normalisation Step 2: Cut Elimination

We now show the main construction of this paper. A cut-elimination result for derivations in simple form. The procedure uses a class of external and independent derivations in order to glue together pieces of the original proof. One valid class of such derivations are the ones shown in Sect. 5.

Theorem 11. *Let*

1. *$N > 0$ be an integer;*
2. *$a_1, \ldots, a_n$ be distinct atoms, where $n = 2^m$ for some $m \geq 0$;*
3. *there be, for every $0 < k < N$ and $1 \leq i \leq n$, a formula $C^{a_i}_k$;*
4. *there be, for every $1 \leq k \leq N$, a derivation*

$$\Gamma_k \quad = \quad \begin{array}{c} \big(a_1 \wedge C^{a_1}_{k-1}\big) \vee \cdots \vee \big(a_n \wedge C^{a_n}_{k-1}\big) \\ \|\mathsf{SKS}\backslash\{\mathsf{ai}\downarrow,\mathsf{ai}\uparrow\} \\ [a_1 \vee C^{a_1}_k] \wedge \cdots \wedge [a_n \vee C^{a_n}_k] \end{array} ,$$

where $C^{a_1}_0 \equiv \cdots \equiv C^{a_n}_0 \equiv \mathsf{t}$ and $C^{a_1}_N \equiv \cdots \equiv C^{a_n}_N \equiv \mathsf{f}$.

For every proof Φ of A, whose flow is

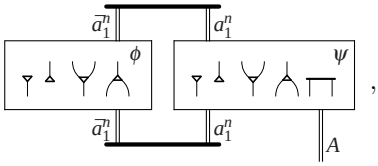

,

where only occurrences of the atoms $\bar{a}_1, \ldots, \bar{a}_n$ are mapped to edges in ϕ, there exists a cut-free proof Ψ of A whose size is bounded by a polynomial in N, the size of Φ and, for $1 \leq k \leq N$, the size of Γ_k.

Proof. For every $1 \leq i \leq n$, let m_i (resp., m'_i) be the number of interactions (resp., cuts) where a_i^ψ and $\bar{a}_i^\phi$ appears in Φ, and consider the derivation

$$\Phi' \;=\; \begin{array}{c} \bigwedge_{1\leq j\leq m_1}\left[a_1^\psi \vee \bar{a}_1^\phi\right] \wedge \cdots \wedge \bigvee_{1\leq j\leq m_n}\left[a_n^\psi \vee \bar{a}_n^\phi\right] \\ \| \\ A \vee \bigvee_{1\leq j\leq m'_1}\left(a_1^\psi \wedge \bar{a}_1^\phi\right) \vee \cdots \vee \bigvee_{1\leq j\leq m'_n}\left(a_n^\psi \wedge \bar{a}_n^\phi\right) \end{array},$$

with atomic flow

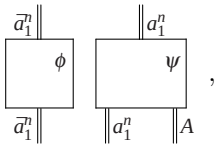

which exists by Lemma 2. Then, for $0 \leq k \leq N$, construct the following derivation:

$$\Phi_k \;=\; \begin{array}{c} [a_1 \vee C_k^{a_1}] \wedge \cdots \wedge [a_n \vee C_k^{a_n}] \\ \|\{\mathsf{c}\uparrow,\mathsf{w}\uparrow\} \\ \bigwedge_{1\leq j\leq m_1}[a_1 \vee C_k^{a_1}] \wedge \cdots \wedge \bigvee_{1\leq j\leq m_n}[a_n \vee C_k^{a_n}] \\ \Phi'\{\bar{a}_1^\phi/C_k^{a_1},\ldots,\bar{a}_n^\phi/C_k^{a_n}\}\|\mathsf{SKS}\setminus\{\mathsf{ai}\uparrow\} \\ A \vee \bigvee_{1\leq j\leq m'_1}(a_1 \wedge C_k^{a_1}) \vee \cdots \vee \bigvee_{1\leq j\leq m'_n}(a_n \wedge C_k^{a_n}) \\ \|\{\mathsf{c}\downarrow,\mathsf{w}\downarrow\} \\ A \vee (a_1 \wedge C_k^{a_1}) \vee \cdots \vee (a_n \wedge C_k^{a_n}) \end{array},$$

which exists, and whose size is bounded by a polynomial in the size of Φ and the size of Γ_k, by Lemma 3. We then construct the cut-free derivation Ψ as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\mathsf{t}}{([a_1 \vee \mathsf{t}] \wedge \cdots \wedge [a_n \vee \mathsf{t}])}\, n\times\mathsf{aw}\downarrow}}{A \vee (a_1 \wedge C_0^{a_1}) \vee \cdots \vee (a_n \wedge C_0^{a_n})}\,\Phi_0\|\mathsf{SKS}\setminus\{\mathsf{ai}\uparrow\}}}{A \vee ([a_1 \vee C_1^{a_1}] \wedge \cdots \wedge [a_n \vee C_1^{a_n}])}\,A\vee\Gamma_1\|\mathsf{SKS}\setminus\{\mathsf{ai}\downarrow,\mathsf{ai}\uparrow\}}}{A \vee [A \vee (a_1 \wedge C_1^{a_1}) \vee \cdots \vee (a_n \wedge C_1^{a_n})]}\,A\vee\Phi_1\|\mathsf{SKS}\setminus\{\mathsf{ai}\uparrow\}}}{\vdots}}{A \vee A \vee \cdots \vee A \vee \left([a_1 \vee C_{N-1}^{a_1}] \wedge \cdots \wedge [a_n \vee C_{N-1}^{a_n}]\right)}\,A\vee A\vee\cdots\vee A\vee\Phi_{N-1}\|\mathsf{SKS}\setminus\{\mathsf{ai}\uparrow\}}}{A \vee A \vee \cdots \vee A \vee \left[A \vee \left(a_1 \wedge C_{N-1}^{a_1}\right) \vee \cdots \vee \left(a_n \wedge C_{N-1}^{a_n}\right)\right]}\,A\vee A\vee\cdots\vee A\vee A\vee\Gamma_N\|\mathsf{SKS}\setminus\{\mathsf{ai}\downarrow,\mathsf{ai}\uparrow\}}}{A \vee A \vee \cdots \vee A \vee A \vee ([a_1 \vee C_N^{a_1}] \wedge \cdots \wedge [a_n \vee C_N^{a_n}])}\,A\vee A\vee\cdots\vee A\vee A\vee\Phi_N\|\mathsf{SKS}\setminus\{\mathsf{ai}\uparrow\}}}{A \vee A \vee \cdots \vee A \vee A \vee [A \vee (a_1 \wedge \mathsf{f}) \vee \cdots \vee (a_n \wedge \mathsf{f})]}\, n\times\mathsf{aw}\uparrow}}{A}\,(N-1)\times\mathsf{c}\downarrow}$$

□

It is worth noting that if we fix $N = n + 1$ in Theorem 11, the formulae $C_k^{a_i}$ are bound to be threshold formulae.

Corollary 12. *Given a proof Φ of A, there exists a cut-free proof Ψ of A, whose size is bounded by a quasipolynomial in the size of Φ.*

Proof. The result follows by Theorem 5, Theorem 10 and Theorem 11. □

7 Final Comments

The quasipolynomial cut-elimination procedure makes use of the cocontraction rule. But the cocontraction rule can also be eliminated. A natural question is whether one can extend the quasipolynomial cut elimination to a cocontraction elimination or to say it in another way, whether one can eliminate cuts in quasipolynomial time without the help of cocontractions. This is probably an important question because all indications we have point to an essential role being played by cocontraction in keeping the complexity low. Cocontraction has something to do with sharing, it seems to provide a typical 'dag-like' speed-up over the corresponding 'tree-like' expansion.

The role played by cocontractions is the most immediate explanation of why quasipolynomial cut elimination works in Deep Inference and not, at the present stage, in the sequent calculus (for full propositional logic). The reason seems to be that exploiting cocontraction in the absence of cut is an intrinsic feature of deep inference, not achievable in Gentzen theory because of the lack of a top-down symmetry therein.

Another natural question is whether quasipolynomial time is the best we can do: there is no obvious objection to the existence of a polynomial cut-elimination procedure. It is possible to express threshold functions with polynomial formulae, but the hardest problem seems to be to obtain corresponding derivations of polynomial length. Deep inference flexibility in constructing derivations might help here.

The cut-elimination procedure presented here is peculiar because it achieves its result by using an external scheme, constituted by the threshold functions and the corresponding derivations, which does not depend on the particular derivation we are working on. It is as if the threshold construction was a catalyzer that shorten the cut elimination. It would be interesting to interpret this phenomenon computationally, in some sort of Curry-Howard correspondence, where the threshold construction implements a clever operator. We intend to explore this path in the near future.

This leads to the wider question of a computational interpretation of deep inference. Atomic flows are a weak computational trace, which takes only the structural rules into account. It is surprising that such a trace, which forgets all the information given by the logical rules, is powerful enough to drive the cut-elimination procedure. We intend to carefully study its computational power and to see whether one can construct on this ground an original computational interpretation of proofs.

References

1. Albert Atserias, Nicola Galesi, and Pavel Pudlák. Monotone simulations of non-monotone proofs. *Journal of Computer and System Sciences*, 65(4):626–638, 2002.
2. Kai Brünnler. *Deep Inference and Symmetry in Classical Proofs*. Logos Verlag, Berlin, 2004. http://www.iam.unibe.ch/ kai/Papers/phd.pdf.
3. Kai Brünnler. Deep inference and its normal form of derivations. In Arnold Beckmann, Ulrich Berger, Benedikt Löwe, and John V. Tucker, editors, *Computability in Europe 2006*, volume 3988 of *Lecture Notes in Computer Science*, pages 65–74. Springer-Verlag, July 2006. http://www.iam.unibe.ch/ kai/Papers/n.pdf.
4. Kai Brünnler and Alwen Fernanto Tiu. A local system for classical logic. In R. Nieuwenhuis and A. Voronkov, editors, *LPAR 2001*, volume 2250 of *Lecture Notes in Artificial Intelligence*, pages 347–361. Springer-Verlag, 2001. http://www.iam.unibe.ch/ kai/Papers/lcl-lpar.pdf.
5. Paola Bruscoli and Alessio Guglielmi. On the proof complexity of deep inference. *ACM Transactions on Computational Logic*, 10(2):1–34, 2009. Article 14. http://cs.bath.ac.uk/ag/p/PrComplDI.pdf.
6. Samuel R. Buss. The undecidability of k-provability. *Annals of Pure and Applied Logic*, 53(1):75–102, 1991.
7. Alessio Guglielmi. Deep inference and the calculus of structures. Web site at http://alessio.guglielmi.name/res/cos.
8. Alessio Guglielmi. A system of interaction and structure. *ACM Transactions on Computational Logic*, 8(1):1–64, 2007. http://cs.bath.ac.uk/ag/p/SystIntStr.pdf.
9. Alessio Guglielmi and Tom Gundersen. Normalisation control in deep inference via atomic flows. *Logical Methods in Computer Science*, 4(1:9):1–36, 2008. http://arxiv.org/pdf/0709.1205.
10. Emil Jeřábek. Proof complexity of the cut-free calculus of structures. *Journal of Logic and Computation*, 2008. In press. http://www.math.cas.cz/ jerabek/papers/cos.pdf.
11. Ingo Wegener. *The Complexity of Boolean Functions*. John Wiley & Sons Ltd and B. G. Teubner, Stuttgart, 1987.