

ON THE RELATIVE PROOF COMPLEXITY OF DEEP INFERENCE VIA ATOMIC FLOWS

ANUPAM DAS

ABSTRACT. We consider the proof complexity of the minimal complete fragment, KS , of standard deep inference systems for propositional logic. To examine the size of proofs we employ *atomic flows*, diagrams that trace structural changes through a proof but ignore logical information. As results we obtain a polynomial simulation of versions of Resolution, along with some extensions. We also show that these systems, as well as bounded-depth Frege systems, cannot polynomially simulate KS , by giving polynomial-size proofs of certain variants of the propositional pigeonhole principle in KS .

1. INTRODUCTION

Deep inference is a relatively recent proof methodology whose systems differ from other formalisms by allowing derivations themselves to be composed by logical connectives. One of its main features is *locality*, i.e. inference steps can be checked in constant time, a property that is impossible to achieve in Gentzen systems [Brü03]. In recent years there has been an increasing interest in the proof complexity of deep inference [BG09] [Jeř09] [Str12] [BGGP10] [Das11] [Das14], in particular the weaker systems initially introduced by Brünnler and Tiu [BT01]. Perhaps the most notable result is that a certain system, denoted $\text{KS} \cup \{\text{c}\uparrow\}$, quasipolynomially¹ simulates Frege systems [Jeř09] [BGGP10]. It is conjectured that this can be improved to a polynomial simulation, and so proving nontrivial lower bounds for $\text{KS} \cup \{\text{c}\uparrow\}$ is likely equivalent to proving them for Frege systems, a task which has escaped proof complexity theorists for years.

However this quasipolynomial simulation relies crucially on the presence of “dag-like behaviour”, manifested in deep inference by a particular rule, *cocontraction*: $\text{c}\uparrow \frac{A}{A \wedge A}$. Without it we have a minimal complete system closed under deep inference, KS . This system is free of compression mechanisms, in the sense that a proof of a conjunction can be ‘partitioned’ into proofs of each conjunct, unlike proofs in systems that are dag-like or contain cut. This is explained further in [Das12a].

It is conjectured that KS is unable to polynomially simulate $\text{KS} \cup \{\text{c}\uparrow\}$ [BG09] [BGGP10] [Das11] [Str12], raising the question of where exactly it fits in the hierarchy of proof systems.

Atomic flows are diagrams that track structural changes in a proof (duplication, creation and destruction of atoms) but ignore logical information. In this work they serve as

Received by the editors February 20, 2015.

¹A quasipolynomial in n is a function $2^{\log^k n}$.

a useful abstraction because of certain rewriting procedures on them which can be used to manipulate derivations soundly without any mention of logical syntax. Atomic flows are introduced formally in [GG08] and a comprehensive account can be found in [Gun09].

In this paper we focus on upper bounds and simulations to demonstrate the relative strength of KS. The starting points in our arguments are proofs in a system obtained by extending $\text{KS} \cup \{c\uparrow\}$ by the *coweakening* rule: $w\uparrow \frac{A}{\top}$; the resulting system $\text{KS} \cup \{c\uparrow, w\uparrow\}$ is

denoted KS^+ in this paper. We then appeal to sound rewriting rules on the atomic flows of these proofs to show that cocontraction and coweakening steps can, in certain cases, be eliminated from a proof in polynomial time.

It is worth mentioning here that the addition of coweakening makes little difference to proof complexity, indeed it is not difficult to see from the rewriting rules in Fig. 1 that coweakening steps can be eliminated in time linear in the size of the proof. Rather the real generators of complexity in our proofs are the interactions between contraction and cocontraction nodes in the atomic flows, as we show in Prop. 29 and Lemma 32.

In Sect. 4 we give a simple example of how atomic flows can be used to normalise a naïve encoding of truth table proofs in KS^+ to produce a polynomial simulation in KS. As a corollary we obtain a superpolynomial separation of KS from tree-like cut-free Gentzen systems, since they are unable to simulate truth tables [D'A92], a new proof of a result appearing in [BG09].

In Sect. 5 we consider stronger systems; we improve a result of Jeřábek's that KS^+ has polynomial-size proofs of the functional and onto versions of the propositional pigeonhole principle by showing that they can be polynomially transformed into KS proofs of the same conclusions. This immediately entails that cut-free Gentzen systems, Resolution and even bounded-depth Frege systems are exponentially separated from KS by the results of [PBI93] and [KPW95].

In Sect. 6 we consider simulations in KS of other proof systems. There is already a naïve simulation of tree-like cut-free Gentzen sequent calculi, appearing in [BG09], and what amounts to a polynomial simulation of 'Resolution with multisets' is outlined in [Gug03]. Here we formalise the latter result and also give a polynomial simulation of tree-like Resolution systems, even when sets are the basic data structure. We show that both simulations extend to so-called Resolution(f) systems, introduced by Krajíček in [Kra01] and known to be strictly stronger than usual versions of Resolution [SBI02] [EGM04].

This paper is a full version of [Das12b], and differs from that work as follows:

- (1) Full proofs are given where they were brief or omitted previously.
- (2) We expand on some of the preliminary work on the complexity of normalisation induced by flow rewriting in Sect. 2. In particular we provide full proofs of termination and confluence for the rewriting system `norm` and give explicit reduction strategies that achieve the complexity bounds given in the previous work.
- (3) In Sect. 2 we use the definitions and notations found in [GG08] and [Gun09] for atomic flows, to maintain consistency with the existing literature, whereas in [Das12b] self-contained definitions were preferred for brevity.
- (4) There were some errors in the statements of results in the previous work, which have been corrected here. In particular the previously stated simulations of dag-like cut-free Gentzen systems and dag-like Resolution systems are incorrect as presented and we have not been able to amend them. Here we only obtain such simulations for certain cases, namely the versions of Resolution in Sect. 6.

2. DEEP INFERENCE, ATOMIC FLOWS AND NORMALISATION

In this section we introduce deep inference systems for propositional logic and *atomic flows*, diagrams that trace the structural changes in a proof. We consider a graph rewriting system norm on flows, corresponding to sound manipulations on proofs, and analyse the complexity of termination in this system.

2.1. Deep inference. We consider propositional logic with formulae constructed from literals (propositional variables and their duals), also called *atoms*, over the basis $\{\top, \perp, \wedge, \vee\}$, and use the infix symbol \equiv to denote equivalence of expressions. The variables a, b, c, d range over literals, with \bar{a}, \bar{b}, \dots denoting their duals, and A, B, C, D range over formulae; both sets of variables may include subscripts or superscripts as necessary.

For clarity we use square brackets $[,]$ for disjunctions and round ones $(,)$ for conjunctions. We generally omit external brackets of an expression, and also internal ones under associativity. This does not cause any confusion when it comes to proofs, since any valid bracketing can be reduced to any other by the $=$ rule in Dfn. 1.

Note that we do not have a symbol for negation in our language, formulae are always in negation normal form. We may however write \bar{A} to denote the De Morgan dual of a formula A , obtained by the following rules:

$$\bar{\perp} \equiv \top, \quad \bar{\top} \equiv \perp, \quad \bar{\bar{a}} \equiv a, \quad \overline{A \wedge B} \equiv \bar{A} \vee \bar{B}, \quad \overline{A \vee B} \equiv \bar{A} \wedge \bar{B}$$

Definition 1 (Rules and systems). An *inference rule* is a sound binary relation on formulae decidable in polynomial time, and a *system* is a set of rules. We define the rules we use below, and the systems $\text{KS} = \{\text{ai}\downarrow, \text{aw}\downarrow, \text{ac}\downarrow, \text{s}, \text{m}\}$, $\text{KS}^+ = \text{KS} \cup \{\text{aw}\uparrow, \text{ac}\uparrow\}$ and $\text{SKS} = \text{KS}^+ \cup \{\text{ai}\uparrow\}$.

Atomic structural rules			Linear logical rules
$\text{ai}\downarrow \frac{\top}{a \vee \bar{a}}$ <i>identity</i>	$\text{aw}\downarrow \frac{\perp}{a}$ <i>weakening</i>	$\text{ac}\downarrow \frac{a \vee a}{a}$ <i>contraction</i>	$\text{s} \frac{A \wedge [B \vee C]}{(A \wedge B) \vee C}$ <i>switch</i>
$\text{ai}\uparrow \frac{a \wedge \bar{a}}{\perp}$ <i>cut</i>	$\text{aw}\uparrow \frac{a}{\top}$ <i>coweakening</i>	$\text{ac}\uparrow \frac{a}{a \wedge a}$ <i>cocontraction</i>	$\text{m} \frac{(A \wedge B) \vee (C \wedge D)}{[A \vee C] \wedge [B \vee D]}$ <i>medial</i>

Note in particular our distinction between variables for literals and formulae in the above rules, and between ‘structural’ and ‘logical’ rules.

We also have the logical rule $=$ which is obtained by closing the equations below under reflexivity, symmetry, transitivity and by applying context closure. We implicitly assume that it is contained in every system.

Commutativity

$$A \vee B = B \vee A$$

$$A \wedge B = B \wedge A$$

Associativity

$$[A \vee B] \vee C = A \vee [B \vee C]$$

$$(A \wedge B) \wedge C = A \wedge (B \wedge C)$$

Units

$$A \vee \perp = A$$

$$A \wedge \top = A$$

$$\top \vee \top = \top$$

$$\perp \wedge \perp = \perp$$

Context closure: if $A = B$, $\star \in \{\wedge, \vee\}$, then $C \star A = C \star B$

A proof that $=$ is decidable in polynomial time can be found in [BG09]. Essentially both formulae are just reduced to some canonical form and then compared. Consequently we often omit occurrences of the $=$ rule in proofs and derivations.

Definition 2 (Proofs and derivations). We define derivations and premiss and conclusion functions (pr and cn respectively) inductively.

- (1) Each formula A is a derivation with premiss and conclusion A .
- (2) If Φ and Ψ are derivations and $\star \in \{\wedge, \vee\}$ then $(\Phi \star \Psi)$ is a derivation with premiss $\text{pr}(\Phi) \star \text{pr}(\Psi)$ and conclusion $\text{cn}(\Phi) \star \text{cn}(\Psi)$.
- (3) If Φ and Ψ are derivations and $\rho \frac{\text{cn}(\Phi)}{\text{pr}(\Psi)}$ is an instance of some inference rule ρ then $\rho \frac{\Phi}{\Psi}$ is a derivation with premiss $\text{pr}(\Phi)$ and conclusion $\text{cn}(\Psi)$.

If $\text{pr}(\Phi) \equiv \top$ then we call Φ a *proof*. If Φ is a derivation where all inference steps are instances of rules in a system \mathcal{S} with premiss A , conclusion B , we write $\Phi \parallel_{\mathcal{S}} \frac{A}{B}$. Furthermore, if $A \equiv \top$, i.e. Φ is a proof in a system \mathcal{S} , we write $\Phi \parallel_{\mathcal{S}} \frac{}{B}$.

While our structural rules only have atoms in their premisses and conclusions, the notion of derivation above allows us to extend these to arbitrary formulae, as stated in the proposition below. We often use these ‘generic rules’ rather than their full derivations for convenience.

Proposition 3 (Generic rules). Each rule below is derivable from s , m , and its respective atomic structural rule in polynomial time.

$$\text{i}\downarrow \frac{\top}{A \vee \bar{A}} \quad \text{w}\downarrow \frac{\perp}{A} \quad \text{c}\downarrow \frac{A \vee A}{A} \quad \text{i}\uparrow \frac{A \wedge \bar{A}}{\perp} \quad \text{w}\uparrow \frac{A}{\top} \quad \text{c}\uparrow \frac{A}{A \wedge A}$$

Proof. See [BT01] for full proofs. Here we just give an example of the case for contraction, since that is the only structural rule of the sequent calculus that cannot be reduced to atomic form [Brü03]. The proof is by induction on the depth of the conclusion of a $\text{c}\downarrow$ step.


$$\begin{aligned} \text{c}\downarrow \frac{[A \vee B] \vee [A \vee B]}{A \vee B} &\rightarrow = \frac{[A \vee B] \vee [A \vee B]}{\text{c}\downarrow \frac{A \vee A}{A} \vee \text{c}\downarrow \frac{B \vee B}{B}} \\ \text{c}\downarrow \frac{(A \wedge B) \vee (A \wedge B)}{A \wedge B} &\rightarrow \text{m} \frac{(A \wedge B) \vee (A \wedge B)}{\text{c}\downarrow \frac{A \vee A}{A} \wedge \text{c}\downarrow \frac{B \vee B}{B}} \end{aligned}$$

Note that the case for cocontraction is dual to this: one can just flip the derivations upside down and replace every formula with its De Morgan dual. \square

Definition 4 (Complexity). We define the size $|\Phi|$ of a derivation Φ to be the number of atom occurrences in Φ . A system \mathcal{S} *polynomially simulates* a system \mathcal{T} if each proof in \mathcal{T} can be polynomially transformed into a proof in \mathcal{S} of the same conclusion.

2.2. **Atomic flows.** We give only an informal definition of atomic flows here, but refer the reader to [GG08], [Gun09] for a formal account of atomic flows.

Definition 5 (Atomic flows). For an SKS derivation Φ we define its *atomic flow*, $fl(\Phi)$, to be the diagram obtained by tracing the path of each atom through the derivation, designating the creation, duplication and destruction of atoms by the following corresponding nodes:



We do not have nodes for s , m or $=$ since they do not create, destroy or duplicate any atom occurrences.

More generally an atomic flow, not necessarily of a derivation, is a (vertically) directed graph embedded in the plane generated from the six types of node above.

Atomic flows are considered equivalent up to continuous deformation preserving the (vertical) ordering of connected edges. Note that edges may be *pending* at either end.

We define the size of a flow ϕ , denoted $|\phi|$, to be its number of edges.

In previous works atomic flows have been equipped with a labeling of the edges, or a polarity assignment, for example to avoid the following impossible situation:



Since we are only concerned with the complexity of flows and their transformations we do not include this extra structure; this does not affect the soundness or termination of our rewriting systems, and in fact is crucial in order to obtain confluence, for which labellings of edges can cause problems.

We do, however, insist that edges are vertically directed and so we are often able to talk about one node being ‘above’ another node. Notice that this order is not generally preserved under deformation, e.g. if two nodes are in disconnected components. Whenever we use this notion in arguments it should be clear that it is being used correctly.

Definition 6. A *flow rewriting rule* is an ordered pair of flows, written $\phi \rightarrow \psi$. A *flow rewriting system* (FRS) is a set of flow rewriting rules. A *one-step reduction* of a flow ϕ in an FRS r is a flow ψ that is obtainable from ϕ by replacing some induced subgraph that is the left hand side of some rule in r with its right hand side.

Definition 7. We define a graph rewriting system *norm* on flows in Fig. 1.

The system *norm* is essentially the system $c \cup w$ in [GG08], without the rules for $i\uparrow$. The proof of termination that follows is similar to that for cycle-free flows in [GG08], and the proof of confluence is similar to that in [GGS10].

Notation 8. Let r be a flow rewriting system. We use the following notation:

- We write $\phi \xrightarrow[r]{} \psi$ if there is a one-step reduction from a flow ϕ to a flow ψ using a rule in r .
- We denote by $\xrightarrow[r]{*}$ the reflexive transitive closure of $\xrightarrow[r]{}.$

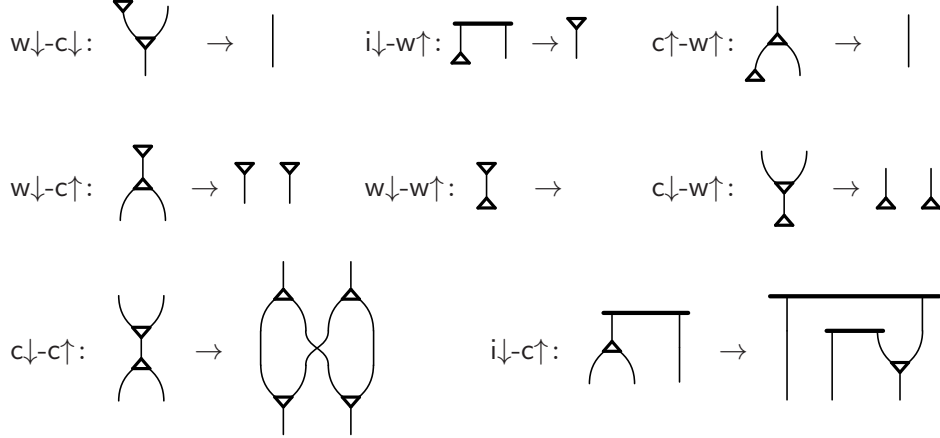


Figure 1: Local rewriting rules for the system norm.

- If a flow ϕ has a unique normal form under r we denote it by $\phi \downarrow_r$.²

In all cases we might omit the subscript r if it is clear from context.

Example 9. We give an example of a flow associated with a derivation in Fig. 2, as well as a reduction under norm , as defined in Dfn. 7, applying $w\uparrow$ rules first.

The first equality follows by the definition of a flow, the second by deformation and the final by definition again. The intermediate steps are as follows:

- (1) Apply $c\downarrow-w\uparrow$ on the left and $c\downarrow-c\uparrow$ on the right.
- (2) Apply $w\downarrow-w\uparrow$ on the left, $i\downarrow-w\uparrow$ in the middle and $i\downarrow-c\uparrow$ on the right.
- (3) Apply $w\downarrow-c\uparrow$ on the left.
- (4) Apply $w\downarrow-c\downarrow$ twice on the left.

The use of colours in the initial and final flow identifies which edges corresponds to which atoms; in the intermediate flows the colours should aid the reader in reconstructing the corresponding transformations on the derivation.

We now proceed to prove that reducing under norm is terminating and confluent. For this we use the usual critical pair lemma, that one only needs to check that the one-step contracta of every overlapping pair of redexes are joinable in order to conclude local confluence. It is simple to see that this is still valid in the flow rewriting setting, since the only other case to check for flows is when the redexes are disjoint, which is trivial. We point out that Newman's lemma, that any locally confluent terminating system is confluent, is true more generally for any binary relation on any set, and so indeed holds in this setting.

Theorem 10. $\xrightarrow{\text{norm}}$ is terminating and confluent.

Proof. For a node ν in a flow ϕ , let $d(\nu, \phi)$ be the distance of ν from the top of ϕ , i.e. the minimum length of a (vertically) directed path from ν to an $\text{ai}\downarrow$ node, an $\text{aw}\downarrow$ node or an edge with upper end pending. For an atomic structural rule ρ , let $D(\rho, \phi)$ be the sequence of natural numbers that counts how many ρ nodes in ϕ have each d -value, i.e. the sequence (n_i) such that, for each i , n_i is the number of ρ nodes ν in ϕ with $d(\nu, \phi) = i$, and consider

²Note that we are using downward arrows in both the names of deep inference rules and to denote normal forms under rewriting systems. Unfortunately both notations are standard in their respective literature, however there should be no ambiguity in their usage so hopefully this will cause little confusion.

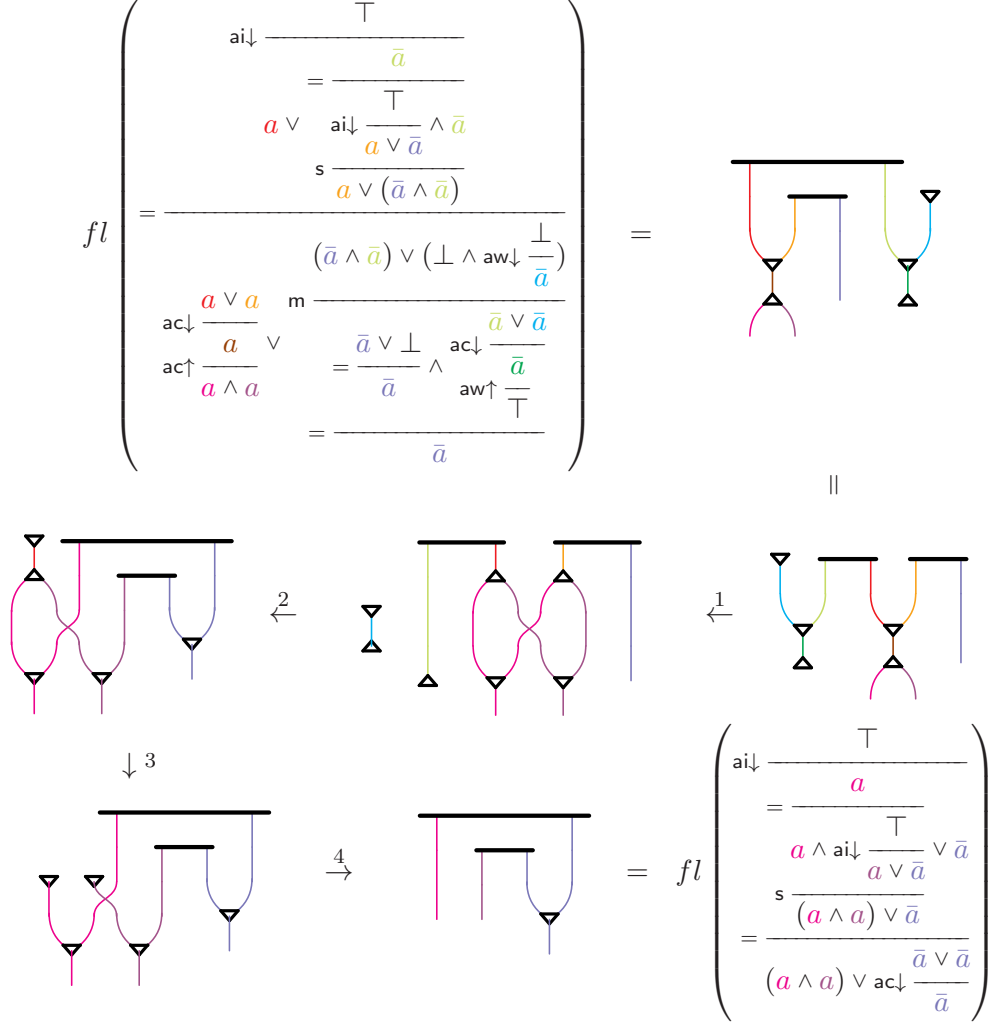



Figure 2: A proof, its flow and a reduction under norm.





the lexicographical ordering $<$ on such sequences, with $(m_i) < (n_i)$ just if $m_k < n_k$ for some k and $m_i \leq n_i$ for $i > k$.

Clearly the rules $c\downarrow-c\uparrow$, $i\downarrow-c\uparrow$ and $w\downarrow-c\uparrow$ strictly reduce the $D(\text{ac}\uparrow, \cdot)$ value of a flow, while the other rules of **norm** (as well as $w\downarrow-c\uparrow$) strictly reduce a flow's size, while not increasing the $D(\text{ac}\uparrow, \cdot)$ value. Therefore each application of a **norm** rule strictly reduces the lexicographical product $D(\text{ac}\uparrow, \cdot) \times |\cdot|$.

Since `norm` is terminating, it suffices to check local confluence for the critical pairs, which are the following by inspection:




Note that every other overlapping pair can be deformed so that each rule application trivially commutes. We consider each case below.



(6) This case is dual to (1).

(7) This case is dual to (2).

(8) This case is dual to (3).



The cases where we appeal to duality follow by simply flipping the indicated reductions upside down and relabeling nodes and reduction steps appropriately. \square

The significance of the rewriting system **norm** is evident from the following results, where we show that reduction steps correspond to sound manipulations of SKS derivations.

Definition 11. If \mathcal{R} is a relation on flows we say that \mathcal{R} *lifts polynomially* to SKS if, whenever $(fl(\Phi), \psi) \in \mathcal{R}$, we can construct a derivation Ψ in time polynomial in $|\Phi| + |\psi|$ with the same premiss and conclusion as Φ and flow ψ . If f is a function on flows then we say that f lifts polynomially to SKS just if the relation $f(\cdot) = \cdot$ lifts polynomially to SKS.

Sometimes we simply say that an individual flow rewrite rule is *sound* rather than saying that it lifts polynomially to SKS.

Remark 12. In fact a derivation can always be manipulated (preserving premiss, conclusion and flow) so that it has size at most polynomial in the size of its flow, as (essentially) shown in [Das13], so the dependence on size of derivation in Dfn. 11 is somewhat redundant. However this is beyond the scope of this work, and it does no harm for us to include this dependence.

Theorem 13. $\xrightarrow{\text{norm}}$ lifts polynomially to SKS.

Proof. See [GG08]. Essentially the proof shows that each local rewrite step on a flow of a derivation induces a sound manipulation of polynomial size on that derivation. We give as an example the case for $c\downarrow$ - $c\uparrow$, since that rule is of particular interest in this work.

Let $\xi\{ \}$ denote an arbitrary context, i.e. a formula with a single hole occurring in place of a subformula, and let $\xi\{A\}$ denote the result of substituting a formula A for the hole in $\xi\{ \}$. The manipulation is as follows:

$$c\downarrow$$
- $c\uparrow$:
$$\xi \left\{ \begin{array}{c} a \vee a \\ \text{ac}\downarrow \frac{\quad}{a} \end{array} \right\} \quad \Phi \parallel \quad \zeta \left\{ \begin{array}{c} a \\ \text{ac}\uparrow \frac{\quad}{a \wedge a} \end{array} \right\} \quad \rightarrow \quad \xi \left\{ \begin{array}{c} \text{ac}\uparrow \frac{a}{a \wedge a} \vee \text{ac}\uparrow \frac{a}{a \wedge a} \\ \text{m} \frac{\quad}{a \wedge a} \\ \text{ac}\downarrow \frac{a \vee a}{a} \wedge \text{ac}\downarrow \frac{a \vee a}{a} \end{array} \right\} \quad \Phi\{(a/(a \wedge a))\} \parallel \quad \zeta \{a \wedge a\}$$

where $\Phi\{a/(a \wedge a)\}$ is obtained by replacing a by $(a \wedge a)$ everywhere in Φ . □

Corollary 14. Given an SKS derivation Φ and a reduction $fl(\Phi) = \phi_1 \xrightarrow{\text{norm}} \cdots \xrightarrow{\text{norm}} \phi_n = fl(\Phi)\downarrow_{\text{norm}}$ we can construct an SKS derivation with the same premiss and conclusion as Φ and with flow $fl(\Phi)\downarrow_{\text{norm}}$ in time polynomial in $|\Phi| + \sum_{i=1}^n |\phi_i|$.


Proof. By induction on the length n of a **norm** derivation. □

2.3. Reduction strategies. We analyse the complexity of normalising a flow under **norm**, presenting a class of reduction strategies that optimise the size of a $\xrightarrow{\text{norm}}$ derivation from a flow to its normal form, up to a polynomial. Our aim is to prove the following theorem:

Theorem 15. The function $\cdot \downarrow_{\text{norm}}$ mapping a flow to its unique normal form under $\xrightarrow{\text{norm}}$ lifts polynomially to SKS.

The result follows from Cor. 14 if we can find appropriate reductions with size only polynomially dependent on the initial derivation and normal form of its flow.

Example 16. Consider the following flow,



where there are n $\text{ac}\uparrow$ nodes. Notice that this flow has normal form Υ , but the complexity of a derivation witnessing this can vary significantly. If we apply $\text{c}\downarrow\text{-c}\uparrow$ steps first then the length of a derivation to normal form is $\Theta(2^n)$, whereas applying $\{\text{w}\downarrow\text{-c}\downarrow, \text{w}\downarrow\text{-c}\uparrow\}$ steps first results in length $\Theta(n)$. These bounds follow from later results in this section, but are not difficult to prove directly.

In fact, this sort of unnecessary exponential blowup can always be avoided by applying ‘weakening’ rules first.

Definition 17. We define the following subsystems of norm:

- (1) $\text{wk} = \{\text{w}\downarrow\text{-c}\downarrow, \text{i}\downarrow\text{-w}\uparrow, \text{c}\uparrow\text{-w}\uparrow, \text{w}\downarrow\text{-c}\uparrow, \text{w}\downarrow\text{-w}\uparrow, \text{c}\downarrow\text{-w}\uparrow\}$.
- (2) $\text{cont} = \text{norm} \setminus \text{wk} = \{\text{c}\downarrow\text{-c}\uparrow, \text{i}\downarrow\text{-c}\uparrow\}$.

Proposition 18. Given a flow ϕ we have:

- (1) If $\phi \xrightarrow{\text{wk}} \psi$ then $|\phi| > |\psi|$.
- (2) If $\phi \xrightarrow{\text{cont}} \psi$ then $|\phi| < |\psi|$.

Lemma 19. Given a flow ϕ we have $(\phi \downarrow_{\text{wk}}) \downarrow_{\text{cont}} = \phi \downarrow_{\text{norm}}$.

Proof. By inspecting the rules of cont , we observe that if there are no wk -redexes in a flow ψ and $\psi \xrightarrow{\text{cont}} \theta$ then there are no wk -redexes in θ . By induction on the length of a $\xrightarrow{\text{cont}}$ derivation we then have that there are no wk -redexes in $(\phi \downarrow_{\text{wk}}) \downarrow_{\text{cont}}$. But since there are also no cont -redexes, by definition of normal form, and $\text{wk} \cup \text{cont} = \text{norm}$ we can conclude that $(\phi \downarrow_{\text{wk}}) \downarrow_{\text{cont}}$ is already in normal form for norm . \square

We can now give a proof of the main theorem of this section.

Proof of Thm. 15. Let Φ be an SKS derivation with flow ϕ whose normal form under norm is ψ , and fix a derivation $\phi = \phi_1 \xrightarrow{\text{wk}} \cdots \xrightarrow{\text{wk}} \phi_m = \psi_1 \xrightarrow{\text{cont}} \cdots \xrightarrow{\text{cont}} \psi_n = \psi$, which must exist by Lemma 19.

By Cor. 14 we can construct an SKS derivation Ψ with same premiss and conclusion as Φ and flow ψ in time polynomial in $|\Phi| + \sum_{i=1}^m |\phi_i| + \sum_{j=1}^n |\psi_j|$. However, by Prop. 18 we have that $|\phi_i| < |\phi_1| = |\phi| \leq |\Phi|$ for all i and $|\psi_j| < |\psi_n| = |\psi|$ for all j , and so this construction can be done in time polynomial in $(m+n)(|\Phi| + |\psi|)$.

Finally, notice also by Prop. 18 that each $\xrightarrow{\text{wk}}$ step decreases the size of a flow, so m is bounded above by $|\phi|$ (and so also $|\Phi|$), and each **cont** step increases the size of a flow, so n is bounded above by $|\psi|$, whence the result follows. \square

3. COMPLEXITY OF NORMAL FORMS

In this section we specialise previous results to KS^+ flows, reducing the complexity of **norm** reduction to counting the number of certain paths in the initial flow. We then give two simple ways of estimating this number, which we use in later sections to obtain complexity results.

Proposition 20. If ϕ is the flow of a KS^+ proof, with normal form ψ under **norm**, then ψ is free of **aw \downarrow** and **ac \uparrow** nodes, i.e. contains just **KS** nodes.


Proof. We argue by contradiction. Notice that by **c \downarrow -c \uparrow** we can assume that all **ac \uparrow** nodes are above all **ac \downarrow** nodes in ψ , by deformation, and so must have upper end directly connected to an **aw \downarrow** or **ai \downarrow** node, since ϕ is associated with a proof. However then ψ can be reduced by either **i \downarrow -c \uparrow** or **w \downarrow -c \uparrow** , contradicting normality. An **aw \uparrow** node, similarly, must be above all **ac \downarrow** and **ac \uparrow** nodes by **c \downarrow -w \uparrow** and **c \uparrow -w \uparrow** , and so must have upper end directly connected to a **aw \downarrow** or **ai \downarrow** node, and again ψ can be reduced by **w \downarrow -w \uparrow** or **i \downarrow -w \uparrow** , contradicting normality. \square

Notice that the above proposition, along with previous results in this section, allows us to transform KS^+ proofs to **KS** proofs of the same conclusion. We state this formally in Thm. 27 once we have determined more about the complexity of this transformation, which we address in the proceeding results.

Definition 21. An **ai-path** is a (simple) path that changes (vertical) direction only at **ai \downarrow** and **ai \uparrow** nodes. We say that an **ai-path** is *maximal* if it cannot be extended, and that it is *open* if it begins and ends at a pending end of an edge.

The *inversion* of a path is just the same path in the reverse direction.

Example 22. The paths on the right, and their inversions, are exactly all the maximal **ai-paths** of the flow below. All of these are open except 0, since one of its ends is a **aw \downarrow** node. More generally all open paths are maximal but not vice-versa.



The following results allow us to estimate the size of the normal form of a flow, under **norm**, without actually constructing it.

Observation 23. $\xrightarrow{\text{norm}}$ preserves the number of open **ai-paths** in a flow.

Notation 24. We write $\lceil \phi \rceil$ to denote the number of open **ai-paths** in a flow ϕ , modulo inversions, and $\#(\rho, \phi)$ to denote the number of ρ nodes in ϕ for an atomic structural rule ρ .

Lemma 25. If ϕ is the flow of a **KS** proof then $\lceil \phi \rceil = \#(\text{ai}\downarrow, \phi)$.

Proof. Since the flow of a proof can have no edge with upper end pending, every edge must be path-connected to a $\text{aw}\downarrow$ or $\text{ai}\downarrow$ node. Since no open path goes through an $\text{aw}\downarrow$ node, and there are no $\text{ai}\uparrow$ nodes, every open ai -path goes through a unique $\text{ai}\downarrow$ node, and every $\text{ai}\downarrow$ node accommodates at least one such path since there are no $\text{aw}\uparrow$ nodes.

Now, the only other node a path can go through is an $\text{ac}\downarrow$ node. Consequently every node in an open ai -path has in-degree 1 before passing its unique $\text{ai}\downarrow$ node, and out-degree 1 after. Hence each $\text{ai}\downarrow$ node accommodates exactly one open ai -path. \square

Lemma 26. If ϕ is the flow of a KS^+ proof, with normal form ψ under norm , then $|\psi| = O(|\phi| + \lceil \phi \rceil)$.

Proof. Let $\chi = \phi \downarrow_{\text{wk}}$ so that $\phi \xrightarrow[\text{wk}]{*} \chi \xrightarrow[\text{cont}]{*} \psi$ by Lemma 19. By inspection of the rules of wk we must be able to decompose χ into two disjoint components: χ_1 , consisting of just $\text{ai}\downarrow$, $\text{ac}\downarrow$ and $\text{ac}\uparrow$ nodes, and χ_2 consisting of just $\text{aw}\downarrow$ and $\text{aw}\uparrow$ nodes. In fact, since ϕ was the flow of a proof, χ_2 consists of just $\text{aw}\downarrow$ nodes.

Now notice that any $\xrightarrow[\text{cont}]{} \psi$ derivation from χ to ψ acts only on χ_1 , and so ψ can be decomposed into two disjoint components: ψ_1 , which is the normal form of χ_1 under norm , consisting of just $\text{ai}\downarrow$ and $\text{ac}\downarrow$ nodes by Prop. 20, and $\psi_2 = \chi_2$, consisting of just $\text{aw}\downarrow$ nodes.

We have that $|\psi_2| = |\chi_2| \leq |\chi| \leq |\phi|$ by Prop. 18.

Notice that $|\psi_1| = 2 \cdot \#(\text{ai}\downarrow, \psi_1) + \#(\text{ac}\downarrow, \psi_1)$, since each $\text{ai}\downarrow$ node has two edges and each $\text{ac}\downarrow$ node adds a single extra edge. Since an $\text{ac}\downarrow$ node has in-degree 2 and out-degree 1, the number of $\text{ac}\downarrow$ nodes cannot outnumber the number of edges emanating from $\text{ai}\downarrow$ nodes in ψ_1 , i.e. $\#(\text{ac}\downarrow, \psi_1) \leq 2 \cdot \#(\text{ai}\downarrow, \psi_1)$, so we have $|\psi_1| \leq 4 \cdot \#(\text{ai}\downarrow, \psi_1)$. By Lemma 25 we then have that $|\psi_1| \leq 4 \cdot \lceil \psi \rceil$, and by Obs. 23 that $|\psi_1| \leq 4 \cdot \lceil \phi \rceil$.

Putting these together we obtain $|\psi| = |\psi_1| + |\psi_2| \leq |\phi| + 4 \cdot \lceil \phi \rceil$ as required. \square

Theorem 27. If Φ is a KS^+ proof with flow ϕ then we can transform it to a KS proof of the same conclusion in time polynomial in $|\Phi| + \lceil \phi \rceil$.

Proof. Let ψ be the normal form of ϕ under norm . By Prop. 20 we have that ψ contains just KS nodes, and so by Thm. 15 we have that a KS proof with same conclusion as Φ can be constructed in time polynomial in $|\Phi| + |\psi|$. The statement follows by the bound $|\psi| = O(|\phi| + \lceil \phi \rceil) = O(|\Phi| + \lceil \phi \rceil)$ given by Lemma 26. \square

3.1. Length of atomic flows. It is not difficult to see that the main contributor to an increase of flow size reducing under norm is the rule $\text{c}\downarrow\text{-c}\uparrow$. It can sometimes cause an exponential blowup, as evident in Ex. 16 if there were no $\text{aw}\downarrow$ node at the top.

The following result provides a simple estimate of the number of paths in a flow, and also the complexity of flow normalisation under norm , and is just an adaptation of well-known techniques to estimate the number of paths in any directed acyclic graph.

Definition 28 (Dimensions of a flow). The *length* of a flow is the maximum number of times the type of node changes in an ai -path. The *width* of a flow is the maximum size of a connected subflow containing just one type of node. The *breadth* of a flow is the number of connected components it has.

The above definition is perhaps most easily understood by allowing $\text{ac}\downarrow$ and $\text{ac}\uparrow$ nodes to have unbounded in-degree and out-degree respectively. For example, by just collapsing

any configuration of $n - 1$ connected $\text{ac}\downarrow$ nodes to a single node $\begin{array}{c} \cdots \\ \cup \\ n \\ \cap \\ \cdots \end{array}$, and similarly for configurations of $\text{ac}\uparrow$ nodes. Then we can consider the length of the flow to be essentially just the maximum length, in the usual sense, of an ai-path, and the width is the maximum out-degree or in-degree.

It is worth mentioning here that replacing $\text{ac}\downarrow$ and $\text{ac}\uparrow$ nodes with these ‘super’ nodes can be lifted polynomially to proofs, in the sense that $\text{ac}\downarrow$ and $\text{ac}\uparrow$ steps in a proof can be soundly replaced by similar ‘super’ steps with only linear change in size.

Proposition 29. If ϕ is a flow consisting of just KS^+ nodes with width w , length l and breadth b , then $\lceil \phi \rceil = b \cdot w^{\frac{l}{2} + O(1)}$.

Proof. For simplicity we write $\begin{array}{c} \cdots \\ \cup \\ n \\ \cap \\ \cdots \end{array}$ for some configuration of $n - 1$ connected $\text{c}\downarrow$ nodes, and similarly for a $\text{c}\uparrow$ configuration. Notice that it suffices to consider the case when $b = 1$, since paths in different connected components are disjoint.

In the worst case scenario we just have a sequence of $\begin{array}{c} \cdots \\ \cup \\ w \\ \cap \\ \cdots \end{array}$ configurations in series vertically, and each configuration multiplies the number of paths by w . It does not make much difference if there is a $\begin{array}{c} \text{---} \\ | \quad | \\ \cap \quad \cap \\ w \quad w \\ \cap \quad \cap \\ \cdots \quad \cdots \end{array}$ configuration at the top of the flow because of the way we have defined an ai-path; the exponent differs by just the addition of a constant. \square


Remark 30. We generally use the trivial upper bound of size of flow for width and breadth, yielding the estimate $\lceil \phi \rceil = |\phi|^{\frac{l}{2} + O(1)}$ for a KS^+ flow ϕ .

3.2. Contraction loops in atomic flows. Sometimes the estimate of number of open ai-paths given via length of a flow in the previous section is not sufficiently accurate. An example of this is given in Sect. 6.2 where the length of flows in the translation R is polynomial in their size, and so the estimate is exponential, yet the actual number of open ai-paths is bounded above by a polynomial.

Unsurprisingly, it is only certain interactions between $\text{ac}\downarrow$ and $\text{ac}\uparrow$ nodes that generate significant complexity in a KS^+ flow; in this section we focus on one type of interaction we call a ‘contraction loop’. In the absence of these the number of open ai-paths in a flow is still polynomial in its size, regardless of its length.

Definition 31. A *contraction loop* in a flow is a pair of $(\text{ac}\uparrow, \text{ac}\downarrow)$ nodes (ν_1, ν_2) such that there are at least two disjoint (directed) paths between ν_1 and ν_2

For example we give the following flow and all its contraction loops,



whereas every other pair has only one path between them. If the edge \star were broken and there was no path from w to x then there would be no contraction loops at all.

Lemma 32. If there are no contraction loops in a KS^+ flow ϕ then $\lceil \phi \rceil \leq |\phi|^3$.

Proof. For an edge ϵ in ϕ consider the following two notions:

- The *weight* of ϵ , denoted $w(\epsilon)$, is the number of directed paths from ϵ to the bottom of ϕ , i.e. to a $\text{aw}\uparrow$ node or an edge with lower end pending.
- For an atomic structural rule ρ let $N(\rho, \epsilon)$ denote the number of ρ nodes below ϵ that are connected to ϵ by a directed path.

We show that $w(\epsilon) \leq N(\text{aw}\uparrow, \epsilon) + N(\text{ac}\uparrow, \epsilon) + 1$ by induction on the distance of ϵ from the bottom of ϕ . The inequality is clear for the base cases when ϵ is directly connected to a $\text{aw}\uparrow$ node, $\epsilon \downarrow \triangle$, and when ϵ has lower end pending, $\epsilon \downarrow$. We have two inductive steps:


- (1) ϵ is an upper edge of a $\text{ac}\downarrow$ node, $\epsilon \downarrow \delta \downarrow$. In this case we clearly have that $w(\epsilon) = w(\delta)$ and so the inequality follows by the inductive hypothesis.
- (2) ϵ is the upper edge of a $\text{ac}\uparrow$ node, $\epsilon \downarrow \gamma \downarrow \delta$. Observe that, since there are no contraction loops in ϕ and $\text{ac}\downarrow$ is the only node type with in-degree greater than 1, any node below this $\text{ac}\uparrow$ node can be directed-path-connected to at most one of γ or δ . Consequently we have that $N(\text{aw}\uparrow, \epsilon) = N(\text{aw}\uparrow, \gamma) + N(\text{aw}\uparrow, \delta)$ and $N(\text{ac}\uparrow, \epsilon) = N(\text{ac}\uparrow, \delta) + N(\text{ac}\uparrow, \gamma) + 1$. Therefore,

$$\begin{aligned} w(\epsilon) &= w(\delta) + w(\gamma) \\ &\leq (N(\text{aw}\uparrow, \delta) + N(\text{ac}\uparrow, \delta) + 1) + (N(\text{aw}\uparrow, \gamma) + N(\text{ac}\uparrow, \gamma) + 1) \\ &\leq (N(\text{aw}\uparrow, \delta) + N(\text{aw}\uparrow, \gamma) + (N(\text{ac}\uparrow, \delta) + N(\text{ac}\uparrow, \gamma) + 1) + 1) \\ &\leq N(\text{aw}\uparrow, \epsilon) + N(\text{ac}\uparrow, \epsilon) + 1 \end{aligned}$$

Finally, notice that $|\phi| \geq N(\text{aw}\uparrow, \phi) + N(\text{ac}\uparrow, \phi) + 1$ and so we have shown that $w(\epsilon) \leq |\phi|$.

Clearly the number of open ai -paths going through an edge with upper end pending is bounded above by its weight, and so by $|\phi|$ by the bound above, while the number of open ai -paths going through any $\text{ai}\downarrow$ node is bounded above by the product of the weights of each of its edges, and so by $|\phi|^2$. In particular we have that the number of open ai -paths going through any edge at the top of a flow is bounded above by $|\phi|^2$, and there are at most $|\phi|$ many such edges, whence the bound follows. □

Example 33. In fact the bound given above is optimal, up to multiplication by a constant. Consider the flow ϕ below, where there are n ai \downarrow nodes:



Clearly the flow has size linear in n and notice that the weights (as defined in the above proof) of the topmost edges, starting from the left, are $1, 2, \dots, n, n, \dots, 1$ respectively. Consequently the number of open ai-paths going through the ai \downarrow nodes, starting from the outside, is $1^2, 2^2, \dots, n^2$ respectively, and so taking the sum we obtain $\lceil \phi \rceil = \Omega(n^3)$.

Notice also that ϕ has length linear in n and width 2, and so the upper bound on $\lceil \phi \rceil$ given by Prop. 29 is exponential, considerably worse than the bound given in Lemma 32.

4. TRUTH TABLES AND TREE-LIKE GENTZEN SYSTEMS

KS polynomially simulates tree-like cut-free Gentzen systems since its rules are generalisations of Gentzen rules [BT01] [BG09]. In the other direction Bruscoli and Guglielmi have proved in [BG09] that the converse does not hold, by way of the so-called ‘Statman tautologies’. We offer a new proof here, via truth tables, as an exercise prior to the main results.

Let LK^- denote the cut-free sequent calculus, and $tree-LK^-$ denote the system of tree-like proofs in this calculus.³

Proposition 34 (D’Agostino). *Tree- LK^- cannot polynomially simulate truth tables.*

Proof. See [D’A92]. □

To expand slightly on the above proposition, truth tables are efficient when there are exponentially many occurrences of each atom, and some such tautologies are hard for $tree-LK^-$. One such example, used by D’Agostino, is simply the disjunction of every assignment on k propositional variables, what we call $\bigvee_{\mathcal{A}} \gamma_{\mathcal{A}}$ below. Such a formula has size $k \cdot 2^k$, although any tree-like cut-free sequent proof must contain at least $k!$ branches, while a truth table contains 2^k rows and $k \cdot 2^k$ columns.

Lemma 35. KS^+ polynomially simulates truth tables.

³It does not matter too much which version of the sequent calculus we choose as complexity properties are typically preserved across common variations.

Proof. Let τ be a tautology. For each partial assignment \mathcal{A} , defined on just those variables appearing in τ , and each formula A satisfied by \mathcal{A} construct a derivation $\Phi_{\mathcal{A}}(A)$ by structural induction on A as follows:

$$\Phi_{\mathcal{A}}(a) \equiv a \quad , \quad \Phi_{\mathcal{A}}(B \wedge C) \equiv \Phi_{\mathcal{A}}(B) \wedge \Phi_{\mathcal{A}}(C) \quad , \quad \Phi_{\mathcal{A}}(B \vee C) \equiv \frac{\Phi_{\mathcal{A}}(B)}{B \vee w \downarrow \frac{\perp}{C}}$$

where, in the last case, when A is a disjunction, the disjunct B was chosen such that B is satisfied by \mathcal{A} . It is clear that each $\Phi_{\mathcal{A}}(\tau)$ has conclusion τ and premiss a conjunction of literals; moreover this conjunction of literals is satisfied by \mathcal{A} .

Let $\gamma_{\mathcal{A}}$ be the conjunction of all literals satisfied by \mathcal{A} , so that each variable appears exactly once. Then we can easily construct derivations $\frac{\gamma_{\mathcal{A}}}{\text{pr}(\Phi_{\mathcal{A}}(\tau))} \parallel \{\text{aw}\uparrow, \text{ac}\uparrow\}$.

Now construct a proof Ψ of $\bigvee_{\mathcal{A}} \gamma_{\mathcal{A}}$ in $\{\text{ai}\downarrow, \text{ac}\uparrow, \text{s}, \text{m}\}$ by induction on the number of distinct variables, as shown below:

$$\text{Base case: } \text{ai}\downarrow \frac{\top}{a \vee \bar{a}} \quad , \quad \text{Inductive step: } \bigvee_{\mathcal{A}} \left(\frac{\Psi \parallel \{\text{ai}\downarrow, \text{ac}\uparrow, \text{s}, \text{m}\}}{\frac{\gamma_{\mathcal{A}}}{\text{ai}\downarrow \frac{\top}{a \vee \bar{a}} \wedge \text{c}\uparrow \frac{\gamma_{\mathcal{A}}}{\gamma_{\mathcal{A}} \wedge \gamma_{\mathcal{A}}}}} \right) \stackrel{2\text{-s}}{\bigvee_{\mathcal{A}} (a \wedge \gamma_{\mathcal{A}}) \vee \bigvee_{\mathcal{A}} (\bar{a} \wedge \gamma_{\mathcal{A}})}$$

Finally we put these together and apply contractions to obtain a KS^+ proof of τ :

$$\bigvee_{\mathcal{A}} \left(\begin{array}{c} \Psi \parallel \{\text{ai}\downarrow, \text{ac}\uparrow, \text{s}, \text{m}\} \\ \gamma_{\mathcal{A}} \\ \parallel \{\text{aw}\uparrow, \text{ac}\uparrow\} \\ \text{pr}(\Phi_{\mathcal{A}}(\tau)) \\ \Phi_{\mathcal{A}} \parallel \{\text{w}\downarrow\} \\ \tau \\ \parallel \{\text{c}\downarrow\} \\ \tau \end{array} \right)$$

It is clear that the derivations inside the large parentheses have size polynomial in $|\tau|$, which is the number of columns in a truth table, and the number of these derivations appearing in disjunction is just the number of assignments, which is the number of rows in a truth table. \square

Theorem 36. KS polynomially simulates truth tables.

Proof. Notice that, in the above simulation, all $\text{ac}\uparrow$ steps are above all $\text{ac}\downarrow$ steps, and so the associated flow will have bounded length. The result follows by Cor. 14, Lemma 26 and Prop. 29. \square

Corollary 37. Tree-LK^- cannot polynomially simulate KS .

Proof. Immediate from Prop. 34 and Thm. 36. \square

Remark 38. It should be noted that D’Agostino’s separation is only quasipolynomial,⁴ and so our separation is also only quasipolynomial, while the proof using the Statman tautologies in [BG09] yields an exponential separation. Nonetheless, an exponential separation follows from the results in the next section.

5. SEPARATIONS VIA VARIANTS OF THE PIGEONHOLE PRINCIPLE

Jeřábek has shown that KS^+ has polynomial-size proofs of the functional and onto variants of the pigeonhole principle [Jeř09]. We show that these proofs reduce under norm to KS proofs with only polynomial increase in size, and so cut-free sequent calculi, Resolution and bounded-depth Frege systems are unable to polynomially simulate KS .

The pigeonhole principle states that, if n pigeons sit in $n - 1$ holes, some hole must contain more than one pigeon. In its unrestricted formulation, the mapping from pigeons to holes can be many-many, while in the functional variant it must be many-one, i.e. a function, and the onto variant insists that each hole must be occupied. It is clear that the two variants are weaker than the unrestricted version, and the variant containing both criteria, the onto functional pigeonhole principle, is weaker still. We will see this more clearly in the following definition.

The propositional encodings of pigeonhole principle variants below are most easily understood by interpreting the atoms a_{ij} as “pigeon i sits in hole j ”. Recall that it does not matter how large disjunctions and conjunctions are bracketed, since any valid bracketing can be easily reduced to any other by the $=$ rule.

Definition 39 (Pigeonhole principles). We define the following formulae,

$$\text{PHP}_n \equiv \bigvee_{i=0}^n \bigwedge_{j=1}^n \bar{a}_{ij} \vee \bigvee_{i=0}^n \bigvee_{j < j'} (a_{ij} \wedge a_{ij'}), \quad \text{F}_n \equiv \bigvee_{j=1}^n \bigvee_{i < i'} (a_{ij} \wedge a_{i'j}), \quad \text{O}_n \equiv \bigvee_{j=1}^n \bigwedge_{i=0}^n \bar{a}_{ij}$$

and denote by FPHP_n , OPHP_n and OFPHP_n the formulae obtained by putting in disjunction the associated formulae, i.e. $\text{FPHP}_n \equiv \text{F}_n \vee \text{PHP}_n$, $\text{OPHP}_n \equiv \text{O}_n \vee \text{PHP}_n$ and $\text{OFPHP}_n \equiv \text{O}_n \vee \text{F}_n \vee \text{PHP}_n$.

We can see in the above definition that any variant can be obtained from a stronger variant, i.e. one with a subset of disjuncts, by a simple application of generic weakening $w\downarrow$. Consequently upper bounds on the size of proofs of one variant yield upper bounds for all weaker variants, and lower bounds vice-versa.

The following result was proved by Beame, Impagliazzo and Pitassi, and independently by Krajíček, Pudlák and Woods.

Theorem 40. Bounded-depth Frege systems have only exponential-size proofs of OFPHP_n .

Proof. See [KPW95], [PBI93]. □

Corollary 41. Bounded-depth Frege systems, Resolution and cut-free sequent calculi have only exponential-size proofs of all variants of the pigeonhole principle.

Proof. All the systems are just special cases of bounded-depth Frege, and a proof of any variant can be extended to one of OFPHP_n by an application of generic weakening $w\downarrow$. □

⁴A quasipolynomial is a function of size $n^{\log^{\Theta(1)} n}$. By taking logarithms and using the integral bound it is not difficult to see that $k!$ grows quasipolynomially in 2^k .

On the other hand we have the following:

Theorem 42 (Buss). There are polynomial-size Frege proofs of PHP_n , and so all variants of the pigeonhole principle.

Proof. See [Bus87]. \square

From here, polynomial-size SKS proofs of PHP_n follow from the following observation:

Proposition 43 (Bruscoli and Guglielmi). SKS is polynomially equivalent to Frege systems.

Proof. See [BG09]. \square

Notice that one direction of the above proposition, indeed the direction that we require, that SKS polynomially simulates Frege systems, can be obtained by recognising that the rules of SKS are just generalisations of the rules of Gentzen systems with cut, which are well-known to be polynomially equivalent to Frege systems.

The following trick, now standard in the deep inference literature, is very useful for proving certain tautologies in KS, as we will see. The idea is that if we know there is an SKS proof of a tautology, then we can transform it into a KS proof of that tautology in disjunction with trivial contradictions. If we can find derivations from each of these contradictions to the tautology we want to prove, then the composition mechanisms for derivations in deep inference guarantee that we can build a proof of the tautology.

Lemma 44. Let A be a formula over the atoms a_1, \dots, a_n . Every SKS proof Φ of A can be polynomially transformed to a KS proof of $A \vee \bigvee_i (a_i \wedge \bar{a}_i)$.

Proof. See e.g. [Jeř09], [BGGP10]. \square

Lemma 45 (Jeřábek). There are polynomial-size proofs of FPHP_n and OPHP_n in KS^+ .

Proof. By Thm. 42, Prop. 43 and Lemma 44 there are polynomial-size KS proofs with conclusion $\text{PHP}_n \vee \bigvee_{i,j} (a_{ij} \wedge \bar{a}_{ij})$. For each atom a_{st} we construct a derivations $\Phi_n^{a_{st}}$ in $\text{KS}^+ \setminus \{\text{ac}\downarrow\}$ with premiss $a_{st} \wedge \bar{a}_{st}$ and conclusion FPHP_n as follows:

$$\begin{array}{c}
 = \frac{a_{st} \wedge \bar{a}_{st}}{\top} \\
 \frac{a_{st} \wedge \bar{a}_{st} \wedge \text{i}\downarrow \frac{\bigwedge_{j \neq t} \bar{a}_{sj} \vee \bigvee_{j \neq t} a_{sj}}{\top}}{2\text{-s}} \\
 \frac{\bigwedge_j \bar{a}_{sj} \vee \left(\begin{array}{c} (n-2)\cdot\text{ac}\uparrow \frac{a_{st}}{a_{st} \wedge \dots \wedge a_{st}} \wedge \bigvee_{j \neq t} a_{sj} \\ \parallel \{s\} \\ \bigvee_{j \neq t} a_{st} \wedge a_{sj} \end{array} \right)}{\text{w}\downarrow} \\
 \text{FPHP}_n
 \end{array}$$

We then put these together and apply contractions to obtain proofs of FPHP_n :

$$\left[\begin{array}{c} \text{KS} \\ \text{PHP}_n \vee \bigvee_{i,j} \left(\begin{array}{c} a \wedge \bar{a} \\ \Phi_n^{a_{ij}} \text{KS}^+ \setminus \{\text{ac}\downarrow\} \\ \text{FPHP}_n \end{array} \right) \\ \text{FPHP}_n \end{array} \right] \begin{array}{c} \text{KS} \\ \{\text{c}\downarrow\} \\ \text{FPHP}_n \end{array}$$

We can construct similar derivations from $a_{st} \wedge \bar{a}_{st}$ to OPHP_n , given below, and put them together in a similar way to obtain proofs of OPHP_n in KS^+ .

$$\begin{array}{c} \frac{a_{st} \wedge \bar{a}_{st}}{\top} \\ \frac{a_{st} \wedge \bar{a}_{st} \wedge \text{ai}\downarrow \frac{\bigwedge_{i \neq s} \bar{a}_{it} \vee \bigvee_{i \neq s} a_{it}}{\text{2-s}}}{\bigwedge_i \bar{a}_{it} \vee \left(\begin{array}{c} (n-2) \cdot \text{ac}\uparrow \frac{a_{st}}{a_{st} \wedge \dots \wedge a_{st}} \wedge \bigvee_{i \neq s} a_{it} \\ \{\text{s}\} \\ \bigvee_{i \neq s} a_{st} \wedge a_{it} \end{array} \right)} \\ \frac{\text{w}\downarrow}{\text{OPHP}_n} \end{array}$$

□

Theorem 46. There are polynomial-size KS proofs of FPHP_n , OPHP_n and so also OFPHP_n .

Proof. In the proofs of FPHP_n constructed above, Lemma 45, notice that the only $\text{ac}\uparrow$ steps occur in $\Phi_n^{a_{st}}$ where there are also no $\text{ac}\downarrow$ steps, and similarly for OPHP_n . It follows that there are only two alternations between $\text{ac}\downarrow$ and $\text{ac}\uparrow$ steps in the path of any atom from an $\text{ai}\downarrow$ step, and so the atomic flows of these proofs will have bounded length. The result follows by Thm. 27 and Prop. 29. □

Corollary 47. Cut-free sequent calculi, Resolution and bounded-depth Frege systems are exponentially separated from KS.

Proof. Immediate from Cor. 41 and Thm. 46. □

6. POLYNOMIAL SIMULATIONS OF VERSIONS OF RESOLUTION

In this section we present a polynomial simulation in KS of tree-like and multiset Resolution systems. We point out that, in previous work [Das12b], a more general result was claimed, namely a simulation of unrestricted Resolution operating with sets, however that proof contained errors that we could not amend for this version of the article. The status of those results is unresolved.

We define the Resolution system below in both its set and multiset formulations.

Definition 48 (Resolution). We use symbols Γ, Δ etc. to range over (multi)sets of literals and write ‘ \cup ’ to denote (multi)set union. We define the system (*multiset-*)*Resolution* by the following rules:

$$\text{wk} \frac{\Gamma}{\Gamma \cup \Delta}, \quad \text{res} \frac{\Gamma \cup \{a} \quad \Delta \cup \{\bar{a}\}}{\Gamma \cup \Delta}$$

A *derivation* from (multi)sets $\Gamma_1, \dots, \Gamma_s$ is a list⁵ $\pi = (\Delta_1, \dots, \Delta_n)$ where each Δ_i is some Γ_j or follows by one of the rules above whose premisses have occurred previously in the list. If we further have that $\Delta_n = \emptyset$ then we call π a *refutation* of $\Gamma_1, \dots, \Gamma_s$.

We call a derivation $(\Delta_1, \dots, \Delta_n)$ *tree-like* if each Δ_i is used at most once as the premiss of any inference step concluding some Δ_j .

To simplify the treatment of (multiset-)Resolution derivations we address certain rather pathological situations below. These are also the reason why we opt to include weakening in our formulation.⁶

Remark 49 (Assumptions on derivation format). In a (multiset-)Resolution derivation, we can assume that neither premiss of a *res* step contains both a resolved atom and its dual. Any such step would have the following format,

$$\text{res} \frac{\Gamma \cup \{a} \quad \Delta \cup \{a, \bar{a}\}}{\Gamma \cup \Delta \cup \{a\}}$$

which can be simulated by an application of *wk* to the left premiss, if necessary at all.

We consequently have that no *res* step has identical premisses, and so there is no need to consider this case in the translations that follow.

Finally, in the case of sets, we further assume that neither the resolved atom nor its dual appear in the conclusion of a *res* step. I.e., in the definition of the *res* above, we assume that Γ and Δ contain neither a nor \bar{a} . Aside from the above situation this phenomenon might also occur if, say, $a \in \Gamma$ and so $\Gamma = \Gamma \cup \{a\}$, whence the step may be simulated by an application of *wk* in a similar way.

Before presenting our simulations in KS of versions of Resolution, we set some notational conventions below in order to easily switch between the two settings.

Notation 50 (Set symbols in deep inference). To reduce the amount of syntax in our deep inference derivations, we will simply write Γ instead of $\bigvee \Gamma$ for the disjunction of the members of Γ , as an abuse of notation.

We similarly use other (multi)set-theoretic notation. In particular, in light of the above remark, we will always have that $\Gamma \cup \{a\} = \Gamma \vee a$ in both the set and multiset settings whenever it occurs. In the multiset setting we further have that $\Gamma \cup \Delta = \Gamma \vee \Delta$ for all Γ, Δ .

Remark 51. Throughout this section, when we say that a proof system polynomially simulates a refutation system, we mean that every refutation of A in the latter can be polynomially transformed to a proof of \bar{A} in the former.

⁵We do not mind how the list is delimited. In the translations below they are written vertically for presentation reasons.

⁶It is straightforward to show that the formulation without weakening is polynomially equivalent by a rule permutation argument.

Definition 52 (Dual systems). The dual of a deep inference rule $\rho \frac{A}{B}$ is $\bar{\rho} \frac{\bar{B}}{\bar{A}}$. E.g. in SKS a rule $x \downarrow$ is dual to $x \uparrow$, while \mathfrak{s} and \mathfrak{m} are self-dual. The set of duals of a system \mathcal{S} is denoted $\bar{\mathcal{S}}$. Notice that a rule is sound if and only if its dual is, by the law of contraposition.

The dual of a derivation $\Phi \parallel \mathcal{S}$ is the derivation $\bar{\Phi} \parallel \bar{\mathcal{S}}$ obtained by flipping Φ upside down

and replacing each inference step with its dual.

Similarly we define the dual of a flow-rewriting rule to be the rule flipped upside-down, replacing each node with its dual, and the dual of a rewriting system is just the set of its duals. A rule/system is sound, terminating and/or confluent if and only if its dual is.

We are now ready to define our basic translation from Resolution derivations to deep inference, on which our simulation results in later sections will be based.

Definition 53 (Translation of Resolution derivations). We give the following translation R of a Resolution step to a $\overline{\text{KS}}^+$ derivation,

$$\begin{aligned}
 R : \quad \text{wk} \frac{\Gamma}{\Gamma \cup \Delta} &\mapsto \frac{\Gamma}{\Gamma \vee \text{wk} \downarrow \frac{\perp}{\Delta \setminus \Gamma}} \\
 &= \frac{\Gamma}{\Gamma \cup \Delta} \\
 R : \quad \text{res} \frac{\Gamma \cup \{a\} \quad \Delta \cup \{\bar{a}\}}{\Gamma \cup \Delta} &\mapsto \frac{\Gamma \cup \{a\} \quad \Delta \cup \{\bar{a}\}}{\Gamma \vee a \quad \Delta \vee \bar{a}} \\
 &\stackrel{2\text{-s}}{=} \frac{\Gamma \vee \Delta \vee \text{ait} \uparrow \frac{a \wedge \bar{a}}{\perp}}{\Gamma \vee \Delta} \\
 &= \frac{\Gamma \vee \Delta}{(\text{c} \downarrow \frac{\Gamma \cup \Delta}{\Gamma \cup \Delta})}
 \end{aligned}$$

where the parenthesised $\text{c} \downarrow$ steps are present only when Γ and Δ are sets (not multisets) that have nonempty intersection.

We extend the definition of R to any (multiset-)Resolution derivation $\pi = (\Delta_1, \dots, \Delta_n)$ from $\Gamma_1, \dots, \Gamma_s$, such that R has the following format:

$$R : \quad \begin{array}{c} \Delta_1 \\ \vdots \\ \Delta_n \end{array} \mapsto \frac{\bigwedge_{r=1}^s \Gamma_r}{\bigwedge_{m=1}^n \Delta_m} \parallel \overline{\text{KS}}^+$$

The definition is by induction on the length n of the Resolution derivation π .

If $n = 0$, i.e. π is an empty list, then $\bigwedge_m \Delta_m = \top$ and we define:⁷

$$R : \quad \pi \mapsto \text{wk} \uparrow \frac{\bigwedge_{r=1}^s \Gamma_r}{\top}$$

⁷Throughout this translation we shall omit certain $=$ rules required to handle units, e.g. we associate the empty disjunction with \perp and the empty conjunction with \top . By now this is a routine consideration.

If π is extended by a (multi)set Γ_i then we define,

$$R : \begin{array}{c} \Delta_1 \\ \vdots \\ \Delta_n \\ \Gamma_i \end{array} \mapsto \frac{\frac{\bigwedge_{r=1}^s \Gamma_r}{\bigwedge_{r \neq i} \Gamma_r \wedge \text{c}\uparrow \frac{\Gamma_i}{\Gamma_i \wedge \Gamma_i}}}{\frac{\bigwedge_{r=1}^s \Gamma_r}{ID \parallel} \wedge \Gamma_i \quad \bigwedge_{m=1}^n \Delta_m}$$

where the derivation marked ID is already defined by induction.

If π is extended by a wk step then we define,

$$R : \begin{array}{c} \Delta_1 \\ \vdots \\ \Delta_i \\ \vdots \\ \Delta_n \\ \Delta_i \cup \Sigma \end{array} \mapsto \frac{\frac{\bigwedge_{r=1}^s \Gamma_r}{ID \parallel} \bigwedge_{m=1}^n \Delta_m}{\text{c}\uparrow \frac{\Delta_i}{\Delta_i}} \wedge \bigwedge_{m \neq i} \Delta_m \wedge \frac{\Delta_i}{\Delta_i \wedge R(wk) \parallel} \parallel \frac{\Delta_i}{\Delta_i \cup \Sigma}$$

$$= \frac{\bigwedge_{m=1}^{n+1} \Delta_m}{\Delta_i \cup \Sigma}$$

where $\Delta_{n+1} = \Delta_i \cup \Sigma$ and the derivation marked ID is already defined by induction.

If π is extended by a res step then we define,

$$R : \begin{array}{c} \Delta_1 \\ \vdots \\ \Delta'_i \cup \{a\} \\ \vdots \\ \Delta'_j \cup \{\bar{a}\} \\ \vdots \\ \Delta_n \\ \Delta'_i \cup \Delta'_j \end{array} \mapsto \frac{\frac{\bigwedge_{r=1}^s \Gamma_r}{ID \parallel} \bigwedge_{m=1}^n \Delta_m}{\text{c}\uparrow \frac{\Delta_i}{\Delta_i \wedge \Delta_i} \wedge \text{c}\uparrow \frac{\Delta_j}{\Delta_j \wedge \Delta_j}} \wedge \bigwedge_{m \neq i, j} \Delta_k \wedge \frac{\Delta_i \wedge \Delta_j}{\Delta_i \wedge \Delta_j \wedge R(res) \parallel} \parallel \frac{\Delta'_i \cup \Delta'_j}{\Delta'_i \cup \Delta'_j}$$

$$= \frac{\bigwedge_{m=1}^{n+1} \Delta_m}{\Delta'_i \cup \Delta'_j}$$

where $\Delta_i = \Delta'_i \cup \{a\}$, $\Delta_j = \Delta'_j \cup \{\bar{a}\}$, $\Delta_{n+1} = \Delta'_i \cup \Delta'_j$ and the derivation marked ID is already defined by induction.

6.1. Multiset Resolution. We consider the multiset setting, for arbitrary derivations (i.e. not necessarily tree-like). Here we achieve a polynomial simulation in KS rather easily by using the identities pointed out in Not. 50.

Observation 54. The image of a multiset-Resolution derivation under R has no $\text{ac}\downarrow$ steps.

Proof. The only $\text{ac}\downarrow$ steps in the definition of R occur in the translation of individual res steps, where the format of the premiss is $\Gamma \vee \Delta$ and the conclusion $\Gamma \cup \Delta$. However, as already pointed out in Not. 50, these are equivalent in the multiset setting, and so no $\text{ac}\downarrow$ steps are necessary at all. \square

Theorem 55. KS polynomially simulates multiset resolution.

Proof. For a multiset resolution refutation π of $\Gamma_1, \dots, \Gamma_s$ we have that $R\pi$ is a $\overline{\text{KS}}^+$ derivation from $\bigwedge_{r=1}^s \Gamma_r$ to \perp , by definition, not containing $\text{ac}\downarrow$ steps, by the observation above.

Consequently we have that $\overline{R\pi}$ has the following format,

$$\begin{array}{c} \top \\ \overline{R\pi} \parallel \text{KS} \cup \{\text{aw}\uparrow\} \\ \bigvee_{r=1}^s \bar{\Gamma}_r \end{array}$$

whose flow has bounded length (due to the absence of $\text{ac}\uparrow$ nodes) and so can be transformed to a KS proof of the same conclusion in polynomial time by Prop. 29 and Thm. 27. \square

6.2. Tree-like Resolution. We now consider Resolution derivations over sets and show a polynomial simulation of tree-like refutations in KS. Here the argument is only slightly more involved: there are both $\text{ac}\downarrow$ and $\text{ac}\uparrow$ steps occurring, but they do not interact in any complex way.

We point out that the proof of this result could perhaps be carried out more simply by writing tree-like derivations as trees and simulating steps locally, discarding premisses once they are used. However the current presentation allows us to use the same translation R for this and the multiset setting, as well as the extensions introduced in the next section. This uniform treatment is possible due to Lemma 32 on the number of open ai-paths in the absence of contraction loops.

Proposition 56. If π is a tree-like Resolution derivation then every $\text{ac}\uparrow$ node in $fl(R\pi)$ has one edge pending.

Proof. By induction on the length of π . On inspection of the atomic flows, notice that each translation step introduces only $\text{c}\uparrow$ nodes whose lower left edge is pending (with respect to the horizontal order of atoms given in the definition of R in Dfn. 53), so it suffices to show that new nodes are not attached to the lower edge of existing $\text{c}\uparrow$ nodes.

If any such situation existed then, by construction, each $\text{c}\uparrow$ node must be associated with the same set Δ_i that is the premiss of distinct steps in π , contradicting the fact that π is tree-like. \square

Recall the notion of contraction loop from Dfn. 31.

Corollary 57. If π is a tree-like Resolution derivation then $fl(R\pi)$ has no contraction loops.

Proof. Any contraction loop would violate Prop. 56 above. \square

Theorem 58. KS polynomially simulates tree-like Resolution.

Proof. For a tree-like Resolution refutation π of $\Gamma_1, \dots, \Gamma_s$, as in the proof of Thm. 55, we have that $\overline{R\pi}$ has the format:

$$\begin{array}{c} \top \\ \overline{R\pi} \parallel \text{KS}^+ \\ \bigvee_{r=1}^s \bar{\Gamma}_r \end{array}$$

We also have that $fl(\overline{R\pi})$ contains no contraction loops, since otherwise there would also be a contraction loop in $fl(R\pi)$, violating Cor. 57. The result now follows from Lemma 32 and Thm. 27. \square

6.3. Extensions of Resolution. Finally, we notice that this simulation extends to some extensions of Resolution, operating on (multi)sets of conjunctions of literals, introduced by Krajíček in [Kra01]. These systems are known to be strictly stronger than their counterparts we have so far dealt with [SBI02] [EGM04].

Definition 59. Let s, t , etc. denote terms, i.e. conjunctions of literals, and Γ, Δ , etc. now denote sets of terms. We define the following rules,

$$\text{wk} \frac{\Gamma}{\Gamma \cup \Delta} \quad , \quad \wedge \frac{\Gamma \cup \{s\} \quad \Delta \cup \{t\}}{\Gamma \cup \Delta \cup \{s \wedge t\}} \quad , \quad \text{res} \frac{\Gamma \cup \{s\} \quad \Delta \cup \{\bar{a}_1, \dots, \bar{a}_k\}}{\Gamma \cup \Delta}$$

where in the **res** rule we require that $s = \bigwedge_{i=1}^k a_i$.

For a function $f : \mathbb{N} \rightarrow \mathbb{N}$, a (tree-like) (multiset-)Resolution(f) derivation/refutation is defined analogously to Dfn. 48, using the rules above, with the additional proviso that no term has size bigger than $f(N)$, where N is the number of inference steps in the refutation.

Theorem 60. For any $f : \mathbb{N} \rightarrow \mathbb{N}$ we have the following:

- (1) KS polynomially simulates multiset-Resolution(f).
- (2) KS polynomially simulates tree-like Resolution(f).

Proof sketch. The proofs are analogous to those appearing earlier for the tree-like and multiset variants of Resolution respectively, interpreting Γ, Δ etc. as (multi)sets of terms and replacing variables a, \bar{a} etc. by term variables and their duals t, \bar{t} etc.⁸

⁸In light of the aforementioned abuse of notation identifying a (multi)set with the disjunction of its elements, we associate the dual of term $\bigwedge_{i=1}^k a_i$ with the (multi)set $\{\bar{a}_1, \dots, \bar{a}_k\}$.

We extend the definition of R on individual steps to \wedge steps as follows:

$$R : \frac{\Gamma \cup \{s\} \quad \Delta \cup \{t\}}{\Gamma \cup \Delta \cup \{s \wedge t\}} \mapsto \frac{\frac{\Gamma \cup \{s\}}{\Gamma \vee s} \wedge \frac{\Delta \cup \{t\}}{\Delta \vee t}}{\frac{\Gamma \vee \Delta \vee (s \wedge t)}{\Gamma \cup \Delta \cup \{s \wedge t\}} \text{ (c}\downarrow\text{)}} \text{ (c}\downarrow\text{)}$$

R is extended to Resolution(f)-derivations similarly to Dfn. 53, dealing with \wedge steps in the same way res steps.

For 1 the argument is identical to that for Thm. 55 since, again, $\text{c}\downarrow$ steps are not introduced for the same reason, cf. Not. 50.

For 2 the argument is identical to that for Thm. 58 since the argument of Prop. 56 remains valid and so there are still no contraction loops in the flow of a derivation in the image of R . \square

7. CONCLUSIONS

We have presented a series of upper bound results for the deep inference system KS . Polynomial simulations were given for truth tables and versions of Resolution, while polynomial-size proofs of functional and onto variants of the propositional pigeonhole principle were presented, yielding exponential separations from all these systems, as well as bounded-depth Frege systems.

We have seen that atomic flows can act as a powerful tool to analyse and manipulate derivations, and that often we can avoid the possibly exponential blowup arising from the $\text{c}\downarrow\text{-c}\uparrow$ rule. A relevant pursuit would be to investigate whether we can *always* avoid this blowup via, perhaps, a polynomial-time local or global flow reduction; this would imply that KS polynomially simulates KS^+ , and thus quasipolynomially simulates SKS and Frege systems. Despite much work in this direction a result, positive or negative, seems difficult to obtain; this remains arguably the most important question in the proof complexity of deep inference.

Question 61. Does KS (quasi)polynomially simulate KS^+ ?

This question has already been asked in previous works, e.g. [BG09], [Jeř09], [Str12] and [Das11], with both positive and negative answers conjectured. We point out that our contribution might be helpful to any work towards a positive answer.

We also point out that it could be that atomic flows do not themselves include sufficient information to carry out an efficient normalisation procedure from KS^+ to KS , if one exists at all. To this end there is work ongoing by various researchers to augment atomic flows with certain *logical* information, in the hope of accessing further normalisation procedures and the ability to ‘rebuild’ deep inference proofs efficiently from flow objects.⁹

One might argue that our proofs in Sects. 4 and 5 eventually reduced to flows of bounded length, and so the complexity of normalisation was trivially polynomial. A more sophisticated situation might involve flows of bounded width and logarithmic length, again resulting in a polynomial blowup, or quasipolynomial width and polylogarithmic length, giving a quasipolynomial blowup. We refer the reader to the recent article [Das14] for an

⁹This is known to be impossible in polynomial time, by [Das13], under certain hardness assumptions.

example of this, where quasipolynomial-size proofs of the unrestricted pigeonhole principle are given in KS, utilising the techniques of this paper.

All the results that appear in this work, and indeed all other works on proof complexity of deep inference, e.g. [BG09] [Jeř09] [Das14], present only simulations and upper bounds. We currently know of no technique for proving lower bounds for deep inference systems and, in light of the question above and results in [Jeř09] and [Das14], it may be that such an endeavour, even for KS, could be as difficult as that of proving lower bounds for Frege systems.

ACKNOWLEDGEMENTS

I am grateful to my advisor, Alessio Guglielmi, for many fruitful discussions on this and related work, as well as Tom Gundersen, Lutz Straßburger and the anonymous reviewers for this paper and previous versions of it. I owe particular thanks to Arnold Beckmann who spotted crucial errors in previous versions and gave detailed feedback.

REFERENCES

- [BG09] 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/PrComp1DI.pdf>.
- [BGGP10] Paola Bruscoli, Alessio Guglielmi, Tom Gundersen, and Michel Parigot. A quasipolynomial cut-elimination procedure in deep inference via atomic flows and threshold formulae. In Edmund M. Clarke and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR-16)*, volume 6355 of *Lecture Notes in Computer Science*, pages 136–153. Springer-Verlag, 2010. <http://cs.bath.ac.uk/ag/p/QPNDI.pdf>.
- [Brü03] Kai Brünnler. Two restrictions on contraction. *Logic Journal of the IGPL*, 11(5):525–529, 2003. <http://www.iam.unibe.ch/~kai/Papers/RestContr.pdf>.
- [BT01] Kai Brünnler and Alwen Fernanto Tiu. A local system for classical logic. Technical Report WV-01-02, Technische Universität Dresden, 2001. <http://iccl.tu-dresden.de/~kai/LocalClassicalLogic-tr.pdf>.
- [Bus87] Samuel R. Buss. Polynomial size proofs of the propositional pigeonhole principle. *Journal of Symbolic Logic*, 52(4):916–927, 1987.
- [D’A92] Marcello D’Agostino. Are tableaux an improvement on truth-tables? *Journal of Logic, Language and Information*, 1:235–252, 1992. 10.1007/BF00156916.
- [Das11] Anupam Das. On the proof complexity of cut-free bounded deep inference. In Kai Brünnler and George Metcalfe, editors, *Tableaux 2011*, volume 6793 of *Lecture Notes in Artificial Intelligence*, pages 134–148. Springer-Verlag, 2011. <http://www.anupamdas.com/items/PrCompII/ProofComplexityBoundedDI.pdf>.
- [Das12a] Anupam Das. Characterising aspects of proof compression. <http://www.anupamdas.com/items/CompMech/CompMech.pdf>, 2012.
- [Das12b] Anupam Das. Complexity of deep inference via atomic flows. In S. Barry Cooper, Anuj Dawar, and Benedikt Löwe, editors, *Computability in Europe*, volume 7318 of *Lecture Notes in Computer Science*, pages 139–150. Springer-Verlag, 2012. <http://www.anupamdas.com/items/RelComp/RelComp.pdf>.
- [Das13] Anupam Das. Rewriting with linear inferences in propositional logic. In Femke van Raamsdonk, editor, *RTA*, volume 21 of *LIPICs*, pages 158–173. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- [Das14] Anupam Das. On the pigeonhole and related principles in deep inference and monotone systems. In *Joint Meeting of the 23rd EACSL Annual Conference on Computer Science Logic (CSL) and the 29th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2014. To appear.

- [EGM04] Juan Luis Esteban, Nicola Galesi, and Jochen Messner. On the complexity of resolution with bounded conjunctions. In *IN 29TH INTERNATIONAL COLLOQUIUM ON AUTOMATA, LANGUAGES AND PROGRAMMING*, pages 220–231. Springer-Verlag, 2004.
- [GG08] 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://www.lmcs-online.org/ojs/viewarticle.php?id=341>.
- [GGS10] Alessio Guglielmi, Tom Gundersen, and Lutz Straßburger. Breaking paths in atomic flows for classical logic. In Jean-Pierre Jouannaud, editor, *25th Annual IEEE Symposium on Logic in Computer Science*, pages 284–293. IEEE, 2010. <http://www.lix.polytechnique.fr/~lutz/papers/AFII.pdf>.
- [Gug03] Alessio Guglielmi. Resolution in the calculus of structures. <http://cs.bath.ac.uk/ag/p/AG10.pdf>, 2003.
- [Gun09] Tom Gundersen. *A General View of Normalisation Through Atomic Flows*. PhD thesis, University of Bath, 2009.
- [Jeř09] Emil Jeřábek. Proof complexity of the cut-free calculus of structures. *Journal of Logic and Computation*, 19(2):323–339, 2009. <http://www.math.cas.cz/~jerabek/papers/cos.pdf>.
- [KPW95] Jan Krajíček, Pavel Pudlák, and Alan Woods. An exponential lower bound to the size of bounded depth frege proofs of the pigeonhole principle. *Random Structures & Algorithms*, 7(1):15–39, 1995.
- [Kra01] Jan Krajíček. On the weak pigeonhole principle. *Fundamenta Mathematicae*, 170(1-3):123–140, 2001.
- [PBI93] Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3:97–140, 1993. 10.1007/BF01200117.
- [SBI02] Nathan Segerlind, Sam Buss, and Russell Impagliazzo. A switching lemma for small restrictions and lower bounds for k-dnf resolution (extended abstract). In *SIAM J. Comput*, pages 604–613. IEEE, 2002.
- [Str12] Lutz Straßburger. Extension without cut. *Annals of Pure and Applied Logic*, 163(12):1995–2007, 2012.