On the pigeonhole and related principles in deep inference and monotone systems

Anupam Das

► To cite this version:

Anupam Das. On the pigeonhole and related principles in deep inference and monotone systems. CSL-LICS '14 Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), Jul 2014, Vienna, Austria. pp.1 - 10, 10.1145/2603088.2603164 . hal-01096154

HAL Id: hal-01096154

<https://hal.inria.fr/hal-01096154>

Submitted on 16 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the Pigeonhole and Related Principles in Deep Inference and Monotone Systems

Anupam Das

INRIA & University of Bath

anupam.das@inria.fr

Abstract

We construct quasipolynomial-size proofs of the propositional pigeonhole principle in the deep inference system KS, addressing an open problem raised in previous works and matching the best known upper bound for the more general class of monotone proofs.

We make significant use of monotone formulae computing boolean threshold functions, an idea previously considered in works of Atserias et al. The main construction, monotone proofs witnessing the symmetry of such functions, involves an implementation of merge-sort in the design of proofs in order to tame the structural behaviour of atoms, and so the complexity of normalization. Proof transformations from previous work on atomic flows are then employed to yield appropriate KS proofs.

As further results we show that our constructions can be applied to provide quasipolynomial-size KS proofs of the parity principle and the generalized pigeonhole principle. These bounds are inherited for the class of monotone proofs, and we are further able to construct $n^{O(\log \log n)}$ -size monotone proofs of the weak pigeonhole principle with $(1 + \varepsilon)n$ pigeons and n holes for $\varepsilon = 1/\log^k n$, thereby also improving the best known bounds for monotone proofs.

1. Introduction

The *pigeonhole principle* states that if m pigeons are sitting in n holes, and $m > n$, then two pigeons must be in the same hole. It can be expressed in propositional logic as follows,

$$\text{PHP}_n^m : \bigwedge_{i=1}^m \bigvee_{j=1}^n p_{ij} \rightarrow \bigvee_{j=1}^n \bigvee_{i=1}^{m-1} \bigvee_{i'=i+1}^m p_{ij} \wedge p_{i'j}$$

where p_{ij} should be interpreted as “pigeon i sits in hole j ”.¹ This encoding forms a class of propositional tautologies, for $m > n$, that has become a benchmark in proof complexity [24]. For the case of

¹ Notice that the above formula allows the mapping from pigeons to holes to be many-many. Additional restrictions can be placed on the mapping, demanding that it is a function or that it is onto, resulting in a logically weaker formula, but here we consider only the version above.

$m = n + 1$ many propositional proof systems, such as the cut-free sequent calculus, Resolution and bounded-depth Frege, only have proofs of size exponential in n [21] [23], whereas small proofs (of size polynomial in n) have been constructed for Frege systems, and so also sequent calculi with cut [11].

This paper presents a novel proof structure for PHP_n^m , inspired by previous works of Atserias et al. [2] [3], implemented in a representation of monotone proofs² as rewriting derivations [18]. Consequently, we obtain quasipolynomial³-size proofs of PHP_n^{n+1} in the minimal deep inference system for propositional logic, KS. This answers questions previously raised in [9] [18] [25] [13] on the complexity of KS proofs of PHP_n^m by matching the best known bound for the more general class of monotone proofs [2].

By making certain generalizations we are able to apply our methods to obtain quasipolynomial-size KS proofs of the parity principle and the generalized pigeonhole principle, bounds that are inherited by the class of monotone proofs. Finally we show that our proof structure can be applied to yield $n^{O(\log \log n)}$ -size monotone proofs of $\text{PHP}_n^{(1+\varepsilon)n}$ where $\varepsilon = 1/\log^k n$ for $k > 1$, significantly improving the best known bound of $n^{O(\log n)}$ inherited from proofs of PHP_n^{n+1} in [2].⁴ We point out that this is the first example where considerations in the complexity of deep inference have yielded improved results for more mainstream systems in proof complexity.

Deep inference systems for classical propositional logic were introduced by Guglielmi et al. [15] [8] and, despite significant progress in recent years on the complexity of deep inference, the classification of the system KS remains open.

In [9] it was shown that KS polynomially simulates (tree-like) cut-free sequent calculi but not vice-versa. This result was strengthened in [13] where it was shown that KS polynomially simulates certain fragments of Resolution and dag-like cut-free sequent calculi, and it was also shown that these systems, as well as bounded-depth Frege systems, cannot polynomially simulate KS. This work made significant use of proof transformations induced by certain graph rewriting techniques from [16]. In this way the complexity of normalizing a monotone proof to a KS proof was reduced to counting the number of paths in the associated *atomic flow*, the graph obtained by tracing the journey of each atom through the proof.

It was asked in [25] and [9] whether polynomial-size proofs of PHP_n^m exist in KS, and in [25] it was conjectured that no polynomial-size proofs exist. On the other hand, in [18] Jeřábek gives proofs in an extended system of weaker variants of the pi-

² A monotone proof is a proof in the sequent calculus free of negation-steps.

³ A quasipolynomial in n is a function of size $n^{\log^{\Theta(1)} n}$.

⁴ This result also supports the more general conjecture in the community that the class of monotone proofs polynomially simulates Frege systems [3] [19] [20].

geonhole principle, where the mapping from pigeons to holes is required to be functional or onto, which normalize to KS proofs of polynomial size [13]. He uses an elegant black-box construction relying on the existence of the aforementioned Frege proofs, although he notes that this method does not seem to generalize to PHP_n^m .

In this work we rely heavily on a propositional encoding of *threshold functions*, yielding formulae that count how many of their arguments are true, and our construction is inspired by the monotone proofs of PHP_n^m given by Atserias et al. [2]. They use the same threshold formulae as us but our main construction, short proofs that permute the arguments of a threshold formula, is considerably more involved than the analogous construction in their paper due to technicalities of the weaker system KS. The tradeoff is that this more sophisticated proof structure enables us to later achieve the aforementioned improvement in upper bounds on the size of monotone proofs for the weak pigeonhole principle.

In [2] simple proofs are provided for each transposition, whence the result follows since each permutation can be expressed as a product of at most polynomially many transpositions, resulting in monotone proofs whose atomic flows have polynomial length. However due to this length bound such proofs normalize to exponential-size KS proofs under the aforementioned transformations. Instead we notice in Sect. 3 that the specific permutation required, corresponding to the transposition of a matrix, has a particularly simple decomposition into logarithmically many *interleavings*, which we implement as monotone proofs whose atomic flows have polylogarithmic length and hence normalize to KS proofs in quasipolynomial time.

In Sect. 4 we generalize this construction by noticing that any permutation can be decomposed into a product of logarithmically many *riffle shuffles*; this is equivalent to the action of applying merge-sort to the inverse of a permutation. In Sect. 5, we show that these techniques can be applied to yield the aforementioned proofs of the parity principle, generalized pigeonhole principle and the weak pigeonhole principle. In the final result the polylogarithmic length of the atomic flows of our proof structure is crucial since it allows us to use smaller monotone formulae that only *approximate* threshold functions, and to maintain a sufficiently accurate approximation throughout the various permutations of their arguments.

2. Preliminaries

Deep inference systems for classical logic were introduced in [8] and studied in detail in [6] and [9]. The representation of proofs we use here was introduced in [17].

2.1 Propositional Logic

Propositional formulae are constructed freely from atoms (propositional variables and their duals), also known as literals, over the basis $\{\top, \perp, \wedge, \vee\}$, with their usual interpretations. The variables a, b, c, d range over atoms, with $\bar{a}, \bar{b}, \dots$ denoting their duals, and A, B, C, D range over formulae. There is no object-level symbol for negation; instead we may write $\bar{A}$ to denote the De Morgan dual of A , obtained by the following rules:

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

For convenience, we consider formulae equivalent under the smallest equivalence relation generated by the equations below.

$$\begin{array}{ll} [A \vee B] \vee C = A \vee [B \vee C] & A \vee \perp = A \\ (A \wedge B) \wedge C = A \wedge (B \wedge C) & A \wedge \top = A \\ A \vee B = B \vee A & \top \vee \top = \top \\ A \wedge B = B \wedge A & \perp \wedge \perp = \perp \end{array}$$

$$\text{If } A = B, \star \in \{\wedge, \vee\}, \text{ then } C \star A = C \star B$$

For this reason we generally omit internal brackets of a formula, under associativity, as well as external brackets. For clarity we also use square brackets $[,]$ for disjunctions and round ones $(,)$ for conjunctions.

Remark 1 (Equality). Equality of formulae $=$, as defined above, is usually implemented as an inference rule in deep inference. It is decidable in polynomial time [9], and whether it is implemented as an inference rule or equivalence relation is purely a matter of convention. Nonetheless we sometimes use it as a ‘fake’ inference rule, to aid the reader.

It will sometimes be convenient to represent the arguments of a boolean function as a vector or matrix of atoms. However the order in which the atoms are to be read is sensitive, and so we introduce the following notation.

Notation 2 (Vectors and Matrices of Variables). We use bold lowercase letters $\mathbf{a}, \mathbf{b}, \dots$ to denote (row-)vectors of atoms and bold uppercase letters $\mathbf{A}, \mathbf{B}, \dots$ to denote matrices of atoms. Vectors are read in their natural order, and we associate a matrix with the vector obtained by reading it rows-first. In this way the transpose of a matrix is equivalent to the vector obtained by reading it columns-first.

The notation $(\mathbf{a}, \mathbf{b})$ denotes the horizontal concatenation of vectors $\mathbf{a}$ and $\mathbf{b}$, and compound matrices are similarly written in the usual way. The notation $(a_i)_{i=1}^n$ denotes the vector $(a_1, \dots, a_n)$.

Definition 3 (Rules and Systems). An *inference rule* is a binary relation on formulae decidable in polynomial time, and a *system* is a set of rules. We define the deep inference system SKS as the set of all inference rules in Fig. 1, and also the subsystem $\text{KS} = \{\text{ai}\downarrow, \text{aw}\downarrow, \text{ac}\downarrow, \text{s}, \text{m}\}$. Note in particular the distinction between variables for atoms and formulae.

Remark 4. It is worth pointing out that the formulation of deep inference with units is very convenient for proof-theoretic manipulation of derivations, and we exploit this throughout. However we could equally formulate our systems without units with no significant change in complexity; this approach is taken in [25] and the equivalence of these two formulations is shown in [12].

Definition 5 (Proofs and Derivations). We define derivations, and premiss and conclusion functions, pr , cn resp., inductively:

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

If $\text{pr}(\Phi) = \top$ 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 $\Phi \Vdash_{\mathcal{S}} B$. Furthermore, if $A = \top$, i.e. Φ is a proof in a system $\mathcal{S}$, we write $\Phi \Vdash_{\mathcal{S}}$.

We extend our structural rules beyond atoms, to general formulae, below.

Proposition 6 (Generic Rules). *Each rule below has polynomial-size derivations in the system containing s , m , and its respective*

Atomic structural rules

$$\begin{array}{ccc}
 \text{ai}\downarrow \frac{\top}{a \vee \bar{a}} & \text{aw}\downarrow \frac{\perp}{a} & \text{ac}\downarrow \frac{a \vee a}{a} \\
 \text{identity} & \text{weakening} & \text{contraction} \\
 \\
 \text{ai}\uparrow \frac{a \wedge \bar{a}}{\perp} & \text{aw}\uparrow \frac{a}{\top} & \text{ac}\uparrow \frac{a}{a \wedge a} \\
 \text{cut} & \text{coweakening} & \text{cocontraction}
 \end{array}$$

Logical rules

$$\begin{array}{ccc}
 \text{s} \frac{A \wedge [B \vee C]}{(A \wedge B) \vee C} & & \\
 \text{switch} & & \\
 \\
 \text{m} \frac{(A \wedge B) \vee (C \wedge D)}{[A \vee C] \wedge [B \vee D]} & & \\
 \text{medial} & &
 \end{array}$$

Figure 1. Rules of the deep inference system SKS.

atomic structural rule.

$$\begin{array}{ccc}
 \text{i}\downarrow \frac{\top}{A \vee \bar{A}} & \text{w}\downarrow \frac{\perp}{A} & \text{c}\downarrow \frac{A \vee A}{A} \\
 \text{i}\uparrow \frac{A \wedge \bar{A}}{\perp} & \text{w}\uparrow \frac{A}{\top} & \text{c}\uparrow \frac{A}{A \wedge A}
 \end{array}$$

Proof Sketch. See [8] for full proofs. We just consider the case for contraction, since that is the only structural rule Gentzen calculi cannot reduce to atomic form [5]. The proof is by induction on the depth of the conclusion of a $\text{c}\downarrow$ -step.

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

Note that the case for $\text{c}\uparrow$ is dual to this: one can just flip the derivations upside down and replace every formula with its De Morgan dual. $\text{c}\downarrow$ -steps become $\text{c}\uparrow$ -steps and s and m steps remain valid. $\square$

We often use these ‘generic’ rules in proof constructions, which should be understood as abbreviations for the derivations mentioned above.

Definition 7 (Complexity). The *size* of a derivation Φ , denoted $|\Phi|$, is the number of atom occurrences in it. For a vector $\mathbf{a}$ or matrix $\mathbf{A}$, let $|\mathbf{a}|, |\mathbf{A}|$ denote its number of elements, respectively.

We will generally omit complexity arguments when they are routine, for convenience. However we outline the main techniques used to control complexity in the following sections.

2.2 Monotone and Normal Derivations

We define monotone and normal derivations and relate them to proof systems in deep inference. We point out that the notion of monotone derivation given here is polynomially equivalent to the tree-like monotone sequent calculus [18], and so is consistent with the usual terminology from the point of view of proof complexity.

Definition 8. A derivation is *monotone* if it does not contain the rules $\text{ai}\downarrow, \text{ai}\uparrow$. A monotone derivation is said to be *normal* if it has the following shape:

$$\begin{array}{c}
 A \\
 \parallel \{\text{aw}\uparrow, \text{ac}\uparrow, \text{s}, \text{m}\} \\
 B \\
 \parallel \{\text{aw}\downarrow, \text{ac}\downarrow, \text{s}, \text{m}\} \\
 C
 \end{array}$$

The significance of normal derivations is that they can be efficiently transformed into KS-proofs of the implication they derive, as demonstrated in the following proposition.

Proposition 9. A normal derivation $\begin{array}{c} A \\ \Phi \parallel \{\text{aw}\uparrow, \text{ac}\uparrow, \text{s}, \text{m}\} \\ B \\ \Psi \parallel \{\text{aw}\downarrow, \text{ac}\downarrow, \text{s}, \text{m}\} \\ C \end{array}$ can be transformed in linear time to a KS-proof of $\bar{A} \vee C$.

Proof Sketch. Define the derivation $\bar{\Phi} \parallel \{\text{aw}\downarrow, \text{ac}\downarrow, \text{s}, \text{m}\}$ by flipping Φ upside-down, replacing every atom with its dual, $\wedge$ for $\vee$ and vice-versa. $\text{aw}\uparrow$ -steps become $\text{aw}\downarrow$ -steps, $\text{ac}\uparrow$ -steps become $\text{ac}\downarrow$ -steps and s and m steps remain valid. Now construct the required derivation: $\begin{array}{c} \top \\ \text{i}\downarrow \frac{\bar{B}}{B} \\ \bar{\Phi} \parallel \vee \Psi \parallel \\ A \quad C \end{array}$. $\square$

We emphasize that it is the existence of an ‘intermediate’ formula in normal derivations, e.g. B in the proof above, that allows us to isolate all the $\uparrow$ steps and flip them into $\downarrow$ steps, resulting in a KS proof. If we started with an arbitrary monotone derivation there may be no such formula, and so any choice of an intermediate formula would also flip some $\downarrow$ steps into $\uparrow$ steps.

2.3 Atomic Flows and Normalization

We are particularly interested in those monotone derivations that can be efficiently transformed to normal ones. A thorough analysis of the complexity of such transformations is carried out in [13] in the setting of graph rewriting. We state informally the main concepts and results here.

Atomic flows and various rewriting systems on them were introduced formally in [16].

Definition 10 (Atomic Flows and Dimensions). The *atomic flow*, or simply *flow*, of a monotone derivation is the vertically directed graph obtained by tracing the paths of each atom through the derivation, designating the creation, destruction and duplication of atom occurrences by the following nodes:

$$\begin{array}{ccc}
 \text{aw}\downarrow \frac{\perp}{a} & \rightarrow & \text{Y-shape} \\
 \text{aw}\uparrow \frac{a}{\top} & \rightarrow & \text{X-shape} \\
 \text{ac}\downarrow \frac{a \vee a}{a} & \rightarrow & \text{V-shape} \\
 \text{ac}\uparrow \frac{a}{a \wedge a} & \rightarrow & \text{A-shape}
 \end{array}$$

Figure 2. Graph rewriting rules for atomic flows.

We do not have nodes for s and m steps since they do not create, destroy or duplicate any atom occurrences, and we generally consider flows equivalent up to continuous deformation preserving the vertical order of edges.

The *size* of a flow is its number of edges. The *length* of a flow is the maximum number of times the node type changes in a (vertically directed) path. The *width* of a flow is the maximum number of input or output edges in a subgraph of a connected component.

For intuition, the width of a flow can be thought of as a measure of how much a configuration of $ac\uparrow$ nodes increases the number of edges in a connected component before a configuration of $ac\downarrow$ nodes decreases it.

Example 11. We give an example of a monotone derivation and its flow below:

The flow has length 3, measured from the top-right $aw\downarrow$ node to the bottom-left $ac\uparrow$ node, and width 4, measured either as the outputs of the two top $ac\uparrow$ nodes or the inputs of the two bottom $ac\downarrow$ nodes.

Observation 12. A normal derivation has flow length 1.

Theorem 13 (Normalization). A monotone derivation Φ whose flow has width w and length l can be transformed into a normal derivation of size $|\Phi| \cdot w^{l+O(1)}$, preserving premiss and conclusion.

While the proof of the above theorem can be found in [13], we outline the main ideas to give the reader an intuition of the argument.

Proof Sketch. The graph rewriting rules in Fig. 2 induce transformations on monotone derivations by consideration of the corresponding rule permutations; note that, due to atomicity of the structural rules, permutations with logical steps are trivial. The system is terminating and the flows of normal derivations are all normal forms of this system. Each rewrite step preserves the number of maximal paths between pending edges, and a normal derivation has size polynomial in this measure.

Consequently the complexity of normalizing a monotone derivation is polynomial in its size and the number of maximal paths in its flow, and this is estimated by the given bound. $\square$

Notice, in particular, that any rewrite derivation on atomic flows acts independently on different connected components. Therefore the complexity of normalization is determined by the structural behaviour of individual atoms - there is no interaction between distinct atoms during normalization.

Finally, most of the proofs in this work are inductions, and for the base cases it will typically suffice to build any monotone proof of a single formula or simple class of formulae, since we are interested in how the size (or width, length) of the proofs grow and not their initial values. For this reason, the following result will be useful, and we implicitly assume it when omitting base cases of inductions.

Proposition 14 (Monotone Implicational Completeness). Let A, B be negation-free formulae such that $A \rightarrow B$ is valid. Then there is a monotone derivation $\frac{A}{B}$.

Proof Sketch. Construct a disjunctive normal form A' of A and conjunctive normal form B' of B by distributivity. Note that all distributivity laws are derivable by Dfn. 19 and duality, so there are monotone derivations $\frac{B'}{B}$ and $\frac{A}{A'}$. Clearly each conjunction

of A' logically implies each disjunction of B' and so there must be derivations in $\{aw\downarrow, aw\uparrow\}$ witnessing this fact. Using these derivations and applying $c\downarrow, c\uparrow$ appropriately we can construct a monotone derivation $\frac{A'}{B'}$, whence the result follows by sequential composition of these derivations. $\square$

By appealing to Thm. 13 we then obtain the following result.

Corollary 15. Normal derivations are monotone implicationally complete.

3. Short Proofs of the Pigeonhole Principle

Throughout this section the variables m and n are powers of 2 and $m \leq n$. All proofs in this section are monotone unless otherwise mentioned.

3.1 Threshold Formulae and Permutations

Threshold functions are a class of boolean functions $TH_k^n : \{0, 1\}^n \rightarrow \{0, 1\}$ by $TH_k^n(\sigma_1 \cdots \sigma_n) = 1$ just if $\sum_{i=1}^n \sigma_i \geq k$.

In this section we define quasipolynomial-size monotone formulae computing such functions and construct derivations whose

flows have length $\log^{O(1)} n$ and width $O(n)$ that conduct certain permutations on the arguments of such formulae.

Definition 16 (Threshold Formulae). We define the formulae,

$$\text{th}_k^1(a) := \begin{cases} \top & k = 0 \\ a & k = 1 \\ \perp & k > 1 \end{cases}$$

$$\text{th}_k^{2n}(a, b) := \bigvee_{i+j=k} \text{th}_i^n(a) \wedge \text{th}_j^n(b)$$

for vectors a, b of length n .

Observation 17. th_k^n computes the threshold function TH_k^n , and has size $n^{O(\log n)}$ and depth $O(\log n)$.

Definition 18 (Interleaving). For $a = (a_1, \dots, a_n), b = (b_1, \dots, b_n)$ let $a \parallel b$ denote the interleaving of a with b : $(a_1, b_1, \dots, a_n, b_n)$. More generally, we denote by $a \parallel_m b$ the m -interleaving:

$$(a_1, \dots, a_m, b_1, \dots, b_m, \dots, a_{n-m+1}, \dots, a_n, b_{n-m+1}, \dots, b_n)$$

Definition 19 (Distributivity). We define distributivity rules as abbreviations for the following derivations:

$$\text{dist } \uparrow: \frac{\text{c}\uparrow \frac{A}{A \wedge A} \wedge [B \vee C]}{2\text{-s} \frac{(A \wedge B) \vee (A \wedge C)}{A \wedge A} \wedge [B \vee C]}$$

$$\text{dist } \downarrow: \frac{m \frac{(A \wedge B) \vee (A \wedge C)}{A \vee A}}{\text{c}\downarrow \frac{A \vee A}{A} \wedge [B \vee C]}$$

Lemma 20. There are monotone derivations,

$$\text{th}_k^{2n}(a, b) \parallel \text{th}_k^{2n}(a \parallel_m b)$$

whose flows have length $O(\log n)$ and width $O(n)$.

Proof. We use the following identity:

$$(a, b) \parallel_m (c, d) = (a \parallel_m c, b \parallel_m d)$$

We give an inductive step from n to $2n$ in Fig. 3 where derivations marked IH are obtained by the inductive hypothesis.

The $\text{dist } \uparrow$ steps duplicate each atom at most r times and so, analyzing the associated flow, each inductive step adds $O(r)$ configurations of $\text{ac}\uparrow$ and $\text{ac}\downarrow$ nodes of width $O(r)$ on top of $O(r)$ copies of the inductive hypothesis in parallel. The induction terminates in $\log \frac{n}{m}$ steps, whence the bound on length is obtained. $\square$

Observation 21. For matrices B and C of equal dimensions we have:⁵

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}^\top = \begin{pmatrix} A^\top & C^\top \\ B^\top & D^\top \end{pmatrix}$$

Recall that a matrix of atoms is equivalent to the vector obtained from a rows-first reading of it.

Theorem 22 (Transposition). There are monotone derivations,

$$\text{th}_k^n(X) \parallel \text{th}_k^n(X^\top)$$

⁵Of course, in any such situation, A and D will also have equal dimensions.

whose flows have length $O(\log^2 n)$ and width $O(n)$.

Proof. Let A, B, C, D be the four quadrants of X . We give an inductive step from n to $2n$,

$$\frac{\text{th}_k^{2n} \begin{pmatrix} A & B \\ C & D \end{pmatrix}}{\bigvee_{i+j=k} \left(\text{th}_i^n(A \parallel B) \parallel \text{th}_j^n(C \parallel D) \right)} \wedge \frac{\text{th}_i^n(A^\top) \parallel \text{th}_j^n(C^\top)}{\text{th}_k^{2n} \left(\begin{pmatrix} A^\top \\ B^\top \end{pmatrix}, \begin{pmatrix} C^\top \\ D^\top \end{pmatrix} \right)}$$

interleave $\parallel$

$$\text{th}_k^{2n} \begin{pmatrix} A^\top & C^\top \\ B^\top & D^\top \end{pmatrix}$$

where the derivations marked IH are obtained by the inductive hypothesis and Obs. 21, and the derivation marked ‘interleave’ is obtained by applying Lemma 20 to interleave the rows of the two matrices.

Analyzing the associated flow, each inductive step adds an interleaving below $O(k)$ copies of the inductive hypothesis in parallel, thereby adding $O(\log n)$ to the length and maintaining a width of $O(n)$, by Lemma 20. The induction terminates in $O(\log n)$ steps, whence the upper bound on length is obtained. $\square$

3.2 From Threshold Formulae to the Pigeonhole Principle

The previous section showed that there are ‘short’ derivations transposing a matrix of arguments of a threshold formula. We show here how such derivations are used to obtain quasipolynomial-size proofs of the pigeonhole principle.

In this section almost all derivations are normal, so we omit their flows and complexity analysis.

Definition 23 (Pigeonhole Principle). We define the following:

$$\text{LPHP}_n := \bigwedge_{i=1}^n \bigvee_{j=1}^{n-1} p_{ij} \quad \text{RPHP}_n := \bigvee_{j=1}^{n-1} \bigvee_{i=1}^n \bigvee_{i'=i+1}^n (p_{i'j} \wedge p_{ij})$$

$$\text{PHP}_n := \text{LPHP}_n \rightarrow \text{RPHP}_n$$

Definition 24. Let $\perp_{mn}$ be the $(m \times n)$ matrix with the constant $\perp$ at every entry. Define $P_n = ((p_{ij} \quad \perp_{n1}))$, with i, j ranging as in Dfn. 23. I.e. P_n is obtained by extending (p_{ij}) with an extra column of $\perp$ -entries, so that it is a square matrix.

Our aim in this section is to prove the following theorem, from which we can extract proofs of PHP_n in KS by the results in earlier sections.

Theorem 25. There are normal derivations,

$$\text{LPHP}_n \parallel \text{th}_n^{n^2}(P_n) \quad , \quad \text{th}_n^{n^2}(P_n^\top) \parallel \text{RPHP}_n$$

of size $n^{O(\log n)}$.

Before we can give a proof, we need some intermediate results. It should be pointed out that similar results were provided in [2] for the monotone sequent calculus, which could be translated to deep inference by [7] [18], but we include them for completeness. Indeed, similar results appeared in [10]. These intermediate results are fairly routine, and there is nothing intricate from the point of view of complexity.

Figure 3. Interleaving the arguments of a threshold formula.

Proposition 26. For $l \geq k$ there are normal derivations,

$$\begin{array}{c} \text{th}_l^n(a) \\ \parallel \\ \text{th}_k^n(a) \end{array}$$

of size $n^{O(\log n)}$.

Proof. We give an inductive step from n to $2n$,

$$\begin{array}{c} \text{th}_l^{2n}(a, b) \\ = \frac{\bigvee_{i+j=l} \left(\begin{array}{c} \text{th}_i^n(a) \quad \text{th}_j^n(b) \\ IH \parallel \quad \wedge \quad IH \parallel \\ \text{th}_{i'}^n(a) \quad \text{th}_{j'}^n(b) \end{array} \right)}{\text{th}_k^{2n}(a, b)} \end{array}$$

where i' and j' are chosen such that $i' \leq i$, $j' \leq j$ and $i' + j' = k$, and derivations marked IH are obtained by the inductive hypothesis. $\square$

Lemma 27 (Evaluation). There are normal derivations,

$$\begin{array}{c} \text{th}_{r+s}^{2n}(a, b) \\ \parallel \\ \text{th}_{r+1}^n(a) \vee \text{th}_s^n(b) \end{array}$$

of size $n^{O(\log n)}$.

Proof. Notice that if $i + j = r + s$ then $i > r$ or $j \geq s$. We give a construction in Fig. 4, where Φ and Ψ denote derivations obtained by Prop. 26. $\square$

Lemma 28. For vectors $\mathbf{a}^1, \dots, \mathbf{a}^m$ of atoms there are normal derivations,

$$\begin{array}{c} \bigvee_{r=1}^m \text{th}_k^n(\mathbf{a}^r) \\ \parallel \\ \text{th}_k^{mn}(\mathbf{a}^r)_{r=1}^n \end{array}, \quad \begin{array}{c} \bigwedge_{r=1}^m \text{th}_k^n(\mathbf{a}^r) \\ \parallel \\ \text{th}_{mk}^{mn}(\mathbf{a}^r)_{r=1}^n \end{array}$$

of size $n^{O(\log n)}$.

Proof. By induction on m . Simply apply $=$ and $w\downarrow$ to fill out the formula. $\square$

We are now in a position to prove Thm. 25.

Proof Sketch of Thm. 25. Repeatedly apply Lemma 27 to $\text{th}_n^{n^2}(\mathbf{P}_n^1)$, always setting $r = s$ or $r = s + 1$, until a disjunction of threshold formulae with n arguments each is obtained. By the ordering of the atoms in $\mathbf{P}_n^1$ these threshold formulae will have as arguments $(p_{ij})_{i=1}^n$ for some j , or all $\perp$; in the latter case any such formula is equivalent to $\perp$, since the threshold will be at least 1.

In the former case, by the choice of r and s at each stage, we have that the threshold of each of these formulae is at least 2. From here Prop. 26 can be applied so that all thresholds are 2, whence RPHP_n can be easily derived.

For the other direction, construe each variable p_{ij} as a threshold formula $\text{th}_1^1(p_{ij})$ and apply Lemma 28 to obtain a derivation from LPHP_n to $\text{th}_n^{n^2}(\mathbf{P})$.

In both cases normality is established by the normalization procedure of Thm. 13. We have chained together finitely many normal derivations so the length of the associated flows is bounded by a constant, whence the upper bound on size is obtained. $\square$

Theorem 29. There are normal derivations,

$$\begin{array}{c} \text{LPHP}_n \\ \parallel \\ \text{RPHP}_n \end{array}$$

of size $n^{O(\log^2 n)}$.

Proof. By Thms. 22 and 25 there are monotone derivations with same premiss and conclusion of length $O(\log^2 n)$ and width $O(n)$. The result then follows by Thm. 13. $\square$

Corollary 30. There are KS proofs of PHP_n of size $n^{O(\log^2 n)}$.

Proof. By Prop. 9. $\square$

$$\begin{aligned}
&= \frac{\text{th}_{r+s}^{2n}(\mathbf{a}, \mathbf{b})}{\bigvee_{i+j=r+s} \text{th}_i^n(\mathbf{a}) \wedge \text{th}_j^n(\mathbf{b})} \\
&= \frac{\bigvee_{\substack{i > r \\ j=r+s-i}} \left(\frac{\text{th}_i^n(\mathbf{a}) \wedge \text{w}\uparrow \frac{\text{th}_j^n(\mathbf{b})}{\top}}{\text{th}_i^n(\mathbf{a})} \right) \vee \bigvee_{\substack{j \geq s \\ i=r+s-j}} \left(\frac{\text{w}\uparrow \frac{\text{th}_i^n(\mathbf{a})}{\top} \wedge \text{th}_j^n(\mathbf{b})}{\text{th}_j^n(\mathbf{b})} \right)}{\frac{\text{th}_{r+1}^n(\mathbf{a})}{\text{th}_{r+1}^n(\mathbf{a})} \vee \frac{\text{th}_s^n(\mathbf{b})}{\text{th}_s^n(\mathbf{b})}} \\
&\quad \text{with } s \cdot c \downarrow \text{ on the left and } r \cdot c \downarrow \text{ on the right.}
\end{aligned}$$

Figure 4. Proof of Lemma 27.

3.3 The Case when n is not a Power of 2

Though we have assumed that n is a power of 2 throughout this section, the proof is actually sufficient for all n , as pointed out in [2].

Definition 31. For $r \leq s$ given, define $\text{LPHP}_s(r)$ by substituting $\perp$ for every atom p_{ij} where $i > r$ or $j \geq r$. Define $\text{RPHP}_s(r)$ analogously.

Observation 32. For all $r \leq s$ we have that $\text{LPHP}_s(r) = \text{LPHP}_r$ and $\text{RPHP}_s(r) = \text{RPHP}_r$.⁶ Consequently a proof of PHP_r is just a proof of PHP_n , where n is the power of 2 such that $r \leq n < 2r$.

4. Arbitrary Permutations

Interleavings by themselves do not form a generating set for the symmetric group, and so cannot be used to generate derivations for arbitrary permutations of arguments of threshold formulae. However a generalization of them, corresponding to the set of riffle shuffles on a deck of cards, do form such a set. In this section we show how they may be used to generate arbitrary permutations on the arguments of threshold formulae.

The proofs in this section are similar to those in Sect. 3, and so we omit them for brevity, instead providing the general proof structure as intermediate results.

Recall that our original definition of threshold formulae used a symmetric divide-and-conquer strategy, generated from a complete binary tree in the natural way. In this section it will be useful to have a more general definition of threshold formulae, based on any tree decomposition of the divide-and-conquer strategy.

Throughout this section we assume all trees are binary.

Definition 33. For a tree T , let $d(T)$ denote its depth, $l(T)$ its number of leaves and $|T|$ denote its number of nodes. For a binary tree T , let T_0 denote its left subtree (from the root) and T_1 its right. Thus any string $\sigma \in \{0, 1\}^k$ determines a unique subtree T_σ of T , for $k \leq d(T)$.

Definition 34 (General Threshold Formulae). For a binary tree T and vectors $\mathbf{a}, \mathbf{b}$ with $|\mathbf{a}| = l(T_0)$, $|\mathbf{b}| = l(T_1)$, define

$$\text{th}_k^T(\mathbf{a}, \mathbf{b}) = \bigvee_{i+j=k} \text{th}_i^{T_0}(\mathbf{a}) \wedge \text{th}_j^{T_1}(\mathbf{b})$$

with the base case the same as in Dfn. 16.

The following proposition gives an estimate of the size of these threshold formulae.

Proposition 35. For a binary tree T , $|\text{th}_k^T(\mathbf{a})| = l(T)^{O(d(T))}$.

⁶ Recall that formulae are equivalent up to $\equiv$.

Proof. In the worst case, every level of the binary tree is full, whence the bound is obtained by Obs. 17. $\square$

What we define as a shuffle below corresponds to the common riffle method of shuffling a deck of cards: cut the deck anywhere, partitioning it into a left and right part, and then interleave these in any way, maintaining the relative order of cards in either partition. Under this analogy each card of the deck will correspond to a leaf of the tree determining a threshold formula.

Definition 36 (Cuts and Shuffles). A cut of a vector $(a_1, \dots, a_n)$ is a pair $\{(a_1, \dots, a_m), (a_{m+1}, \dots, a_n)\}$. A riffle shuffle, or simply shuffle, of length n is a string $\sigma \in \{0, 1\}^n$.

For a vector $\mathbf{a}$ and shuffle σ of length $n = |\mathbf{a}|$ we write $\sigma(\mathbf{a})$ to denote the following action of σ on $\mathbf{a}$: let Σ_i denote the number of 1s in $\sigma_1 \dots \sigma_i$, so that $i - \Sigma_i$ is the number of 0s in $\sigma_1 \dots \sigma_i$ and $k = \Sigma_n$ is the number of 1s in σ ; we give a componentwise definition of $\sigma(\mathbf{a})$:

$$(\sigma(\mathbf{a}))_i = \begin{cases} a_{i-\Sigma_i} & \sigma_i = 0 \\ a_{n-k+\Sigma_i} & \sigma_i = 1 \end{cases}$$

In the above definition, one should think of σ determining a cut $\{(a_1, \dots, a_{n-k}), (a_{n-k+1}, \dots, a_n)\}$, and where each bit indicates which side of the cut the next element of $\sigma(\mathbf{a})$ comes from.

Lemma 37 (Cutting). For any tree T and cut $\{\mathbf{a}, \mathbf{b}\}$ there are trees S_0, S_1 with $d(S_0), d(S_1) \leq d(T)$ such that there are monotone derivations,

$$\begin{aligned} &\text{th}_k^T(\mathbf{a}, \mathbf{b}) \\ &\quad \parallel \\ &\bigvee_{i+j=k} \text{th}_i^{S_0}(\mathbf{a}) \wedge \text{th}_j^{S_1}(\mathbf{b}) \end{aligned}$$

whose flows have length $O(d(T))$ and width $O(l(T))$.

Proof of Lemma 37. By induction on $l(T)$. Without loss of generality, suppose $\mathbf{b}$ is contained entirely in T_1 (otherwise, $\mathbf{a}$ is contained entirely in T_0 and the argument is symmetric). We construct the

following derivation,

$$\begin{aligned}
& \frac{\text{th}_r^T(\mathbf{a}, \mathbf{b})}{\bigvee_{s+t=r} \left(\left(\begin{array}{c} \text{th}_s^{T_0}(\mathbf{a}^1) \wedge \text{th}_t^{T_1}(\mathbf{a}^2, \mathbf{b}) \\ \text{IH} \parallel \\ \bigvee_{i+j=t} \text{th}_i^{S'_0}(\mathbf{a}^2) \wedge \text{th}_j^{S_1}(\mathbf{b}) \end{array} \right) \right)} \\
&= \frac{\bigvee_{i+j=t} \text{th}_s^{T_0}(\mathbf{a}^1) \wedge \text{th}_i^{S'_0}(\mathbf{a}^2) \wedge \text{th}_j^{S_1}(\mathbf{b})}{\bigvee_{s'+t'=r} \left(\begin{array}{c} \bigvee_{k+l=s'} \text{th}_k^{T_0}(\mathbf{a}^1) \wedge \text{th}_l^{S'_0}(\mathbf{a}^2) \wedge \text{th}_{t'}^{S_1}(\mathbf{b}) \\ \text{dist} \downarrow \\ \left(\begin{array}{c} \bigvee_{k+l=s'} \text{th}_k^{T_0}(\mathbf{a}^1) \wedge \text{th}_l^{S'_0}(\mathbf{a}^2) \\ \text{th}_{s'}^{S_0}(\mathbf{a}) \end{array} \right) \wedge \text{th}_{t'}^{S_1}(\mathbf{b}) \end{array} \right)}
\end{aligned}$$

where the derivation marked *IH* is obtained by the inductive hypothesis. $\square$

Lemma 38 (Shuffling). *Let S be a tree and σ a shuffle of length $l(S)$. There is a tree T with $d(T) = O(d(S))$ and monotone derivations,*

$$\begin{array}{c} \text{th}_k^S(\mathbf{v}) \\ \parallel \\ \text{th}_k^T(\sigma(\mathbf{v})) \end{array}$$

whose flows have length $O(d(S)^2)$ and width $O(l(S))$.

Proof of Lemma 38. By induction on $l(S)$, we give the inductive step in Fig. 5. In the construction we set $\mathbf{v} = (\mathbf{w}, \mathbf{x})$, defined by $l(S_0)$ and $l(S_1)$, and $\sigma(\mathbf{v}) = (\mathbf{y}, \mathbf{z})$. The argument is analogous to the one in Lemma 3, with derivations marked ‘cut’ obtained from Lemma 37 and derivations marked *IH* obtained from the inductive hypothesis.

In particular the cuts are chosen⁷ such that $|\mathbf{b}^2| = |\mathbf{c}^1|$ and so that there are shuffles σ_1, σ_2 with

$$\sigma(\mathbf{v}) = (\sigma_1(\mathbf{a}, \mathbf{b}^1, \mathbf{c}^1), \sigma_2(\mathbf{b}^2, \mathbf{c}^2, \mathbf{d}))$$

$\square$

Theorem 39 (Merge Sort). *For any tree S and permutation π on $\{1, \dots, l(S)\}$ there is a tree T with $d(T) = O(d(S))$ and monotone derivations,*

$$\begin{array}{c} \text{th}_k^S(a_{i\pi})_{i=1}^n \\ \parallel \\ \text{th}_k^T(a_i)_{i=1}^n \end{array}$$

whose flows have length $O(d(S)^3)$ and width $O(l(S))$.

Proof of Thm. 39. By induction on $l(S) = l(T)$. We construct the following derivation,

$$\begin{array}{c} \text{th}_r^S(a_{i\pi})_{i=1}^n \\ \bigvee_{s+t=r} \left(\begin{array}{c} \text{th}_s^{S_0}(a_{i\pi})_{i=1}^m \quad \text{th}_t^{S_1}(a_{i\pi})_{i=m+1}^n \\ \text{IH} \parallel \quad \wedge \quad \text{IH} \parallel \\ \text{th}_s^{T'_0}(\mathbf{a}^1) \quad \text{th}_t^{T'_1}(\mathbf{a}^2) \end{array} \right) \\ \text{shuffle} \parallel \\ \text{th}_r^T(a_i)_{i=1}^n \end{array}$$

⁷Such a choice exists (and is unique) by the discrete intermediate value theorem.

where the derivations marked *IH* are obtained from the inductive hypothesis, sorting the inputs of the left and right subtrees of S to vectors $\mathbf{a}^1$ and $\mathbf{a}^2$ resp., and the derivation marked ‘shuffle’, obtained from Lemma 38, carries out the unique shuffle on $(\mathbf{a}^1, \mathbf{a}^2)$ resulting in a completely sorted vector. $\square$

Proposition 40 (Repartitionings). *For trees S, T with the same number of leaves there are monotone derivations,*

$$\begin{array}{c} \text{th}_k^S(\mathbf{a}) \\ \parallel \\ \text{th}_k^T(\mathbf{a}) \end{array}$$

whose flows have length $O(d(S)^2)$ and width $O(l(S))$.

Proof of Prop. 40. By induction on $l(S) = l(T)$. Let $\{\mathbf{b}, \mathbf{c}\}$ be the cut of $\mathbf{a}$ such that $|\mathbf{b}| = l(T_0)$ and $|\mathbf{c}| = l(T_1)$. We construct the following derivation,

$$\begin{array}{c} \text{th}_k^S(\mathbf{a}) \\ \text{cut} \parallel \\ \bigvee_{i+j=k} \left(\begin{array}{c} \text{th}_i^{R_0}(\mathbf{b}) \quad \text{th}_j^{R_1}(\mathbf{c}) \\ \text{IH} \parallel \quad \wedge \quad \text{IH} \parallel \\ \text{th}_i^{T_0}(\mathbf{b}) \quad \text{th}_j^{T_1}(\mathbf{c}) \end{array} \right) \\ \parallel \\ \text{th}_k^T(\mathbf{b}, \mathbf{c}) \end{array}$$

where the derivation marked ‘cut’ is obtained from Lemma 37 and the derivations marked *IH* are obtained from the inductive hypothesis. $\square$

Corollary 41. *For any tree T and permutation π on $\{1, \dots, l(T)\}$ there are normal derivations,*

$$\begin{array}{c} \text{th}_k^T(a_i)_{i=1}^n \\ \parallel \\ \text{th}_k^T(a_{i\pi})_{i=1}^n \end{array}$$

of size $l(T)^{O(d(T)^3)}$.

Proof. By Thm. 39, Prop. 40 and Thm. 13. $\square$

5. Further Results and Applications

We give some examples of how the techniques developed in previous sections can be applied to yield further results, namely quasipolynomial-size normal proofs of the Generalized Pigeonhole principle and the Parity principle. Both bounds are also inherited for monotone proofs although, while these have not appeared in the literature, we point out that such monotone proofs could also have been constructed using the permutation arguments of Atserias et al. in [2].

More interestingly we provide $n^{O(\log \log n)}$ -size monotone proofs for the *weak* pigeonhole principle, with $(1+\varepsilon)n$ pigeons and n holes for every $\varepsilon = 1/\log^{\Omega(1)} n$, improving the previous best known bound of $n^{O(\log n)}$ inherited from the proofs of PHP $_n^{n+1}$ given in [2].

5.1 Generalized Pigeonhole Principle

If there are 45 hats that are either red or green, then there must be 23 of the same colour. This exemplifies a generalization of the pigeonhole principle where sufficiently many pigeons may guarantee more than two in some hole [14]. If $k+1$ pigeons in some hole are required then $nk+1$ pigeons are necessary, so this principle can be encoded as follows:

$$\bigwedge_{i=0}^{nk+1} \bigvee_{j=1}^n a_{ij} \rightarrow \bigvee_{i_r < i_{r+1}} \bigwedge_{r=1}^{k+1} a_{i_r, j}$$

$$\begin{aligned}
& \text{th}_r^S(w, x) \\
&= \frac{\bigvee_{s+t=r} \left(\left(\text{th}_i^{S'_0}(a, b^1) \wedge \text{th}_j^{R_0}(b^2) \right) \wedge \left(\text{th}_k^{R_1}(c^1) \wedge \text{th}_l^{S'_1}(c^2, d) \right) \right)}{\left(\text{th}_i^{S'_0}(a, b^1) \wedge \text{th}_j^{R_0}(b^2) \right) \wedge \left(\text{th}_k^{R_1}(c^1) \wedge \text{th}_l^{S'_1}(c^2, d) \right)} \\
&= \frac{\bigvee_{s+t=r} \left(\left(\text{th}_i^{S'_0}(a, b^1) \wedge \text{th}_k^{R_1}(c^1) \right) \wedge \left(\text{th}_j^{R_0}(b^2) \wedge \text{th}_l^{S'_1}(c^2, d) \right) \right)}{\left(\text{th}_i^{S'_0}(a, b^1) \wedge \text{th}_k^{R_1}(c^1) \right) \wedge \left(\text{th}_j^{R_0}(b^2) \wedge \text{th}_l^{S'_1}(c^2, d) \right)} \\
&= \frac{\bigvee_{s+t=r} \left(\text{th}_s^{T'_0}(a, b^1, c^1) \wedge \text{th}_t^{T'_1}(b^2, c^2, d) \right)}{\left(\text{th}_s^{T'_0}(a, b^1, c^1) \wedge \text{th}_t^{T'_1}(b^2, c^2, d) \right)} \\
&= \text{th}_r^T(y, z)
\end{aligned}$$

Figure 5. Riffle shuffling the inputs of a threshold formula.

This formula has size $O(n^{k+1})$, polynomial for fixed k . If, however k is large relative to n , e.g. $n/2$ or $\sqrt{n}$, then one can always express the right hand side using threshold formulae to obtain an encoding of quasipolynomial-size.

It is simple to see that our proofs of PHP_n can be generalized to this class of tautologies, by the same arguments as in Sect. 3.

5.2 Parity Principle

The *parity principle* states that one cannot partition an odd-size set into pairs, and is usually encoded by the following tautologies,

$$\text{PAR}_n : \bigwedge_{i=0}^{2n} \bigvee_{j \neq i} a_{\{i,j\}} \rightarrow \bigvee_{j \neq i > i' \neq j} a_{\{i,j\}} \wedge a_{\{i',j\}}$$

where $a_{\{i,j\}}$ should be interpreted as “element i is paired with element j ”.

These tautologies have similar structure to PHP_n , but in many proof systems these tautologies are in fact *harder* to prove. For example, in bounded-depth Frege systems PHP_n can be efficiently derived from PAR_n but not vice-versa [1] [4].

However, in KS, we can construct quasipolynomial-size proofs of PAR_n using similar methods to those for PHP_n , and we give an outline of these constructions in this subsection.

We omit proofs corresponding to basic properties of threshold functions, since they are fairly routine inductions of which Sect. 3 has given many examples, and also often do not specify precise orderings of variables or tree-structures of a threshold formulae, since these can all be reduced to any other in quasipolynomial time, by the results of Sect. 4.

Let LPAR_n and RPAR_n denote the left and right hand sides of PAR_n respectively. By a similar argument to Prop. 25 we obtain normal derivations of the following form,

$$\begin{aligned}
& \text{LPAR}_n \\
& \parallel \\
& \text{th}_{2n+1}^{2n(2n+1)}(a^2)
\end{aligned}$$

where a^2 is an appropriate sequence of the variables $a_{\{i,j\}}$ in which each variable occurs exactly twice, as in LPAR_n .

Let (a, a) be a permutation of a^2 so that each variable occurs exactly once in a . Now we can construct the following derivation,

$$\begin{aligned}
& \text{th}_{2n+1}^{2n(2n+1)}(a^2) \\
& \parallel \text{permute} \\
& \text{th}_{2n+1}^{2n(2n+1)}(a, a) \\
& \parallel \text{evaluate} \\
& \text{th}_{n+1}^{n(2n+1)}(a) \vee \text{th}_{n+1}^{n(2n+1)}(a) \\
& \text{c}\downarrow \\
& \text{th}_{n+1}^{n(2n+1)}(a)
\end{aligned}$$

where the derivation marked ‘permute’ applies the results of Sect. 4, namely Cor. 41, to permute the arguments of a threshold formula, and the derivation marked ‘evaluate’ is obtained by Lemma 27, setting $r = n$ and $s = n + 1$.

Now notice that, if $n + 1$ of the variables $a_{\{i,j\}}$ are true, i.e. we have $n + 1$ pairs out of $2n + 1$ variables, we must have some j which is paired with two distinct variables, and this can be realized as derivations,

$$\begin{aligned}
& \text{th}_{n+1}^{n(2n+1)}(a) \\
& \parallel \\
& \text{RPAR}_n
\end{aligned}$$

in a similar way to Prop. 25.

Chaining all these normal derivations together gives us monotone derivations $\parallel$ of quasipolynomial size and with flows of RPAR_n bounded length, and from here we can construct quasipolynomial-size KS-proofs of PAR_n in the usual way.

5.3 Monotone Proofs of the Weak Pigeonhole Principle

The results of this section provide the first example of considerations in the complexity of deep inference yielding new results for more mainstream systems in proof complexity. Unlike the previous two results our proofs of the weak pigeonhole principle rely crucially on the fact that the proofs permuting threshold arguments we constructed have flows of polylogarithmic length. The basic idea is to begin with formulae *approximating* threshold functions and

bound how much worse the approximation develops as the interleaving and transposition arguments of Sect. 3.1 are applied.

5.3.1 Approximating Threshold Functions

It is not quite correct to call the formulae we define below as ε -approximators of threshold functions, since in fact they output incorrectly on a large proportion of inputs. Rather they output 1 just if the actual threshold is within some predetermined factor of the threshold being measured. The tradeoff is that we are able to define monotone formulae that are much smaller than the usual threshold formulae we have used until now.

Definition 42 (Threshold Approximators). Let $|a| = |b| = n$. We define the (p, q) -approximator $T_k^n[p, q]$ of TH_k^n as follows,

$$T_k^n[p, q](a, b) = \bigvee_{i+j=p} T_{\frac{ik}{q}}^n[p, q](a) \wedge T_{\frac{jk}{q}}^n[p, q](b)$$

where we assume that k is some power of q and n is a power of 2, for example by adding a string of $\top$ s and $\perp$ s of appropriate format to the arguments.⁸

It is not easy to understand the semantics of these approximators, and in the next section we provide solely proof-theoretic arguments rather than semantic intuition. We do, however, make the following observations, provable by straightforward inductions.

Observation 43. *We have the following properties,*

1. $TH_k^n \Rightarrow T_k^n[p, q]$ for all $p < q$.
2. $T_k^n[p, q] \Rightarrow TH_{k(p/q)^{\log n}}^n$.
3. $|T_k^n[p, q](a)| = O(p^{\log n})$

where $\Rightarrow$ denotes logical implication.

5.3.2 Manipulating Arguments in Threshold Approximators

In this section we return to the derivations proved in Sect. 3.1 on interleaving and transposing arguments of a formula. Since the approximators we now consider do not exactly compute threshold functions, they are no longer symmetric and so similar derivations cannot be constructed. Rather we show that witnessing certain permutations requires a bounded deterioration in the accuracy of the approximation. Ultimately we will choose an initial approximation that is accurate enough to ensure that this deterioration does not become too excessive.

We first state a basic fact allowing us to infer weaker approximations from stronger ones. The proof is straightforward and follows the same kind of induction argument as previous proofs.

Lemma 44. *For $p \geq p', k \geq k'$ there are normal derivations,*

$$\begin{array}{c} T_k^n[p, q](a) \\ \parallel \\ T_{k'}^n[p', q](a) \end{array}$$

of size $p^{O(\log n)}$.

The following is the main result that will control the deterioration in approximation throughout our overall argument.

Lemma 45. *There are normal derivations,*

$$\begin{array}{c} T_k^n[p, q](a, b, c, d) \\ \parallel \\ T_k^n[p-1, q](a, c, b, d) \end{array}$$

of size $p^{O(\log n)}$.

We first need the following result, whose proof is straightforward.

⁸ This increases the number of arguments by at most multiplication by $2q$.

Proposition 46. *For $x, y \geq 0$, if $(x + y) \in \mathbb{N}$ then $(x + y) - (\lfloor x \rfloor + \lfloor y \rfloor) \leq 1$.*

Proof of Lemma 45. We construct appropriate derivations in Fig. 6 where,

$$\begin{array}{lcl} s' & = & \lfloor \frac{is+kt}{p} \rfloor \\ t' & = & \lfloor \frac{js+lt}{p} \rfloor \end{array}$$

and,

$$\begin{array}{lcl} i' & = & \lfloor \frac{isp}{is+kt} \rfloor \\ j' & = & \lfloor \frac{jsp}{js+lt} \rfloor \\ k' & = & \lfloor \frac{ktp}{is+kt} \rfloor \\ l' & = & \lfloor \frac{ltp}{js+lt} \rfloor \end{array}$$

so we have $s' + t', i' + j', k' + l' \geq p - 1$ by Prop. 46 and also $i's' \leq is, j't' \leq js, k's' \leq kt, l't' \leq lt$, so that the derivations marked 'decrease' are obtained by Lemma 44. $\square$

The following result has proof similar to that of Lemma 20, using the above lemma in the inductive steps to measure the deterioration of the approximator.

Proposition 47. *There are monotone derivations,*

$$\begin{array}{c} T_k^n[p, q](a, b) \\ \parallel \\ T_k^n[p - \log n, q](a \parallel b) \end{array}$$

of size $p^{O(\log n)}$ whose flows have length $O(\log n)$ and width $O(p)$.

The following result has proof similar to that of Thm. 22.

Theorem 48. *There are monotone derivations,*

$$\begin{array}{c} T_k^n[p, q](X) \\ \parallel \\ T_k^n[p - \log^2 n, q](X^\top) \end{array}$$

of size $p^{O(\log n)}$ whose flows have length $O(\log^2 n)$ and width $O(p)$.

5.3.3 From Approximators to the Weak Pigeonhole Principle

Recall the definition of PHP_n^m , where m denotes an arbitrary number of pigeons greater than the number of holes n , and define $LPHP_n^m$ and $RPHP_n^m$ analogously to Dfn. 23. In this section we essentially mimic the results of Sect. 3.2 to complete our proofs of the weak pigeonhole principle.

First we will need the following well known result whose proof follows, for example, by consideration of the inclusion-exclusion principle in the binomial expansion.

Proposition 49. *For $\varepsilon \leq 1$ we have that $(1 - \varepsilon)^k \geq 1 - \varepsilon k$.*

Proof. Let $X_0, \dots, X_{k-1}$ be independent identically distributed Bernoulli random variables with $\Pr[X_i = 1] = \varepsilon$ and $\Pr[X_i = 0] = (1 - \varepsilon)$ for each i . Then:

$$\begin{array}{lll} (1 - \varepsilon)^k & = & \Pr \left[\bigcap_{i < k} X_i = 0 \right] \quad \text{by independence.} \\ & = & \Pr \left[\bigcup_{i < k} \overline{X_i} = 1 \right] \quad \text{by De Morgan laws.} \\ & = & 1 - \Pr \left[\bigcup_{i < k} X_i = 1 \right] \quad \text{by complements.} \\ & \geq & 1 - \sum_{i < k} \Pr[X_i = 1] \quad \text{by the union bound.} \\ & \geq & 1 - \varepsilon k \end{array}$$

$\square$

The following result has proof similar to that of Thm. 25.

$$\begin{aligned}
&= \frac{\tau_r^{4n}[p, q](a, b, c, d)}{\left(\begin{array}{c} \tau_r^{2n}[p, q](a, b) \quad \tau_r^{2n}[p, q](c, d) \\ \bigvee_{i+j=p} \tau_{\frac{isr}{q^2}}^n[p, q](a) \wedge \tau_{\frac{jstr}{q^2}}^n[p, q](b) \quad \wedge \quad \bigvee_{k+l=p} \tau_{\frac{ktr}{q^2}}^n[p, q](c) \wedge \tau_{\frac{ltr}{q^2}}^n[p, q](d) \\ \parallel^{\text{dist}\uparrow} \\ \left(\tau_{\frac{isr}{q^2}}^n[p, q](a) \wedge \tau_{\frac{jstr}{q^2}}^n[p, q](b) \right) \wedge \left(\tau_{\frac{ktr}{q^2}}^n[p, q](c) \wedge \tau_{\frac{ltr}{q^2}}^n[p, q](d) \right) \\ \bigvee_{\substack{i+j=p \\ k+l=p}} \left(\begin{array}{cc} \tau_{\frac{isr}{q^2}}^n[p, q](a) & \tau_{\frac{ktr}{q^2}}^n[p, q](c) \\ \parallel^{\text{decrease}} \wedge \parallel^{\text{decrease}} \\ \tau_{\frac{i's'r}{q^2}}^n[p-1, q](a) & \tau_{\frac{k't'r}{q^2}}^n[p-1, q](c) \end{array} \right) \wedge \left(\begin{array}{cc} \tau_{\frac{jstr}{q^2}}^n[p, q](b) & \tau_{\frac{ltr}{q^2}}^n[p, q](d) \\ \parallel^{\text{decrease}} \wedge \parallel^{\text{decrease}} \\ \tau_{\frac{j't'r}{q^2}}^n[p-1, q](b) & \tau_{\frac{l't'r}{q^2}}^n[p-1, q](d) \end{array} \right) \end{array} \right)} \\
&= \frac{\left(\begin{array}{c} \bigvee_{\substack{i+k=p-1 \\ j+l=p-1}} \left(\tau_{\frac{isr}{q^2}}^n[p-1, q](a) \wedge \tau_{\frac{kstr}{q^2}}^n[p-1, q](c) \right) \wedge \left(\tau_{\frac{jstr}{q^2}}^n[p-1, q](b) \wedge \tau_{\frac{ltr}{q^2}}^n[p-1, q](d) \right) \\ \parallel^{\text{dist}\downarrow} \\ \bigvee_{i+k=p-1} \tau_{\frac{isr}{q^2}}^n[p-1, q](a) \wedge \tau_{\frac{kstr}{q^2}}^n[p-1, q](c) \quad \wedge \quad \bigvee_{i+k=p-1} \tau_{\frac{jstr}{q^2}}^n[p-1, q](b) \wedge \tau_{\frac{ltr}{q^2}}^n[p-1, q](d) \\ \tau_r^{2n}[p-1, q](a, c) \quad \wedge \quad \tau_r^{2n}[p-1, q](b, d) \end{array} \right)}{\tau_r^{4n}[p-1, q](a, c, b, d)}
\end{aligned}$$

Figure 6. Single interleaving step for threshold approximators.

Lemma 50. For $q > p$ and $k > \frac{n}{(p/q)^{\log n}}$ there are normal derivations,

$$\begin{array}{ccc}
\text{LPHP}_n^m & & \tau_k^{mn}[p, q](p_{ij})^\top \\
\parallel & , & \parallel \\
\tau_m^{mn}[p, q](p_{ij}) & & \text{RPHP}_n^m
\end{array}$$

of size $p^{O(\log n)}$.

Theorem 51. For $\varepsilon = 1/\log^{\Omega(1)} n$ there are monotone derivations,

$$\begin{array}{ccc}
\text{LPHP}_{(1-\varepsilon)n}^n & & \\
\parallel & & \\
\text{RPHP}_{(1-\varepsilon)n}^n & &
\end{array}$$

of size $n^{O(\log \log n)}$, width $O(\log n)$ and length $O(\log^2 n)$.

Proof. For $\varepsilon = 1/\log^d n$, choose $q = 3 \log^{d+3} n$ and $p = q - 1$. Since $\varepsilon > \frac{1}{q}$, there is a trivial derivation from $\text{LPHP}_{(1-\varepsilon)n}^n$ to $\text{LPHP}_{(1-\frac{1}{q})n}^n$ in $w\downarrow$, and by chaining this to the derivations from Lemmata 50 and 48 we obtain monotone derivations from $\text{LPHP}_{(1-\varepsilon)n}^n$ to $\tau_n^{(1-\frac{1}{q})n^2}[p - \log^2 n, q](p_{ij})^\top$.

We now need to check that $n > \frac{(1-\varepsilon)n}{((p-\log^2 n)/q)^{\log n}}$ before applying Lemma 50. Now we have that,

$$\begin{aligned}
\left(\frac{p - \log^2 n}{q} \right)^{\log n} &= \left(\frac{3 \log^{d+3} n - \log^2 n - 1}{3 \log^{d+3} n} \right)^{\log n} \\
&= \left(1 - \frac{\log^2 n + 1}{3 \log^{d+3} n} \right)^{\log n} \\
&\geq \left(1 - \frac{2}{3 \log^{d+1} n} \right)^{\log n} \\
&\geq 1 - \frac{2 \log n}{3 \log^{d+1} n} \geq 1 - \frac{2}{3 \log^d n}
\end{aligned}$$

by Prop. 49. Consequently we have that,

$$\frac{(1-\varepsilon)}{((p - \log^2 n)/q)^{\log n}} \leq \frac{1 - \frac{1}{\log^d n}}{1 - \frac{2}{3 \log^d n}} < 1$$

giving monotone derivations from $\text{LPHP}_{(1-\varepsilon)n}^n$ to $\text{RPHP}_{(1-\varepsilon)n}^n$ by Lemma 50. From previous bounds, the size of these derivations is $p^{O(\log n)} = (\log n)^{O(\log n)} = n^{O(\log \log n)}$ as required. $\square$

Since the width of these derivations is $O(\log n)$ we also gain a minor improvement in the complexity of KS proofs of $\text{PHP}_{(1-\varepsilon)n}^n$ over those appearing in Sect. 3.

Corollary 52. There are KS proofs of $\text{PHP}_{(1-\varepsilon)n}^n$, for $\varepsilon = 1/\log^{\Omega(1)} n$, of size $n^{O(\log n \log \log n)}$.

6. Final Comments

We constructed explicit quasipolynomial-size proofs of the pigeonhole principle in KS, and generalized our techniques to further yield quasipolynomial-size proofs of the parity principle and quite strong variants of the weak pigeonhole principle. In particular the existence of $n^{O(\log \log n)}$ -size monotone proofs of the most common variant, with $2n$ pigeons and n holes, are implied by our construction. We repeat that this is the first time when considerations in the complexity of deep inference proofs have led to improvements for systems in mainstream proof complexity.

The various proof structures used throughout this work are similar in concept and fairly uniform, and so it might be pertinent to design high-level tools to more easily manipulate deep inference proofs, respecting certain complexity properties. One such approach might be to design an associated theory of *bounded arithmetic*, as done for other propositional proof systems, e.g. the theory IA_0 for bounded-depth Frege systems [22]. Work in this direction is ongoing.

A natural question is whether the methods used here could be generalized to yield a simulation of Frege proofs, as done for KS^+ in [10] [18]. That construction is also relies heavily on threshold formulae, however it is not clear how to restrict the length of flows in the same way as we did here.

Note that the proofs we have given, albeit $n^{O(\log^2 n)}$ so quasipolynomial in size, are not in polynomial correspondence with those

constructed in [2] for the monotone sequent calculus, and so KS^+ , which have smaller quasipolynomial size $n^{O(\log n)}$. In fact it is conjectured that there are polynomial-size proofs in KS^+ , due to the more general conjecture that the monotone sequent calculus polynomially simulates the full sequent calculus over monotone sequents. Consequently we cannot rule out the possibility that proofs of PHP_n witness a superpolynomial separation between KS and KS^+ .

References

- [1] M. Ajtai. Parity and the pigeonhole principle. In S. Buss and P. Scott, editors, *Feasible Mathematics*, volume 9 of *Progress in Computer Science and Applied Logic*, pages 1–24. Birkhäuser Boston, 1990. ISBN 978-0-8176-3483-4. URL http://dx.doi.org/10.1007/978-1-4612-3466-1_1.
- [2] A. Atserias, N. Galesi, and R. Gavalda. Monotone proofs of the pigeonhole principle. *Mathematical Logic Quarterly*, 47(4):461–474, 2001. ISSN 1521-3870. URL <http://www.lsi.upc.edu/~atserias/papers/php/proof9.pdf>.
- [3] A. Atserias, N. Galesi, and P. Pudlák. Monotone simulations of non-monotone proofs. *Journal of Computer and System Sciences*, 65(4):626–638, 2002.
- [4] P. Beame and T. Pitassi. An exponential separation between the parity principle and the pigeonhole principle. pages 195–228, 1996.
- [5] K. Brännler. Two restrictions on contraction. *Logic Journal of the IGPL*, 11(5):525–529, 2003. <http://www.iam.unibe.ch/~kai/Papers/RestContr.pdf>.
- [6] K. Brännler. *Deep Inference and Symmetry in Classical Proofs*. Logos Verlag, Berlin, 2004. <http://www.iam.unibe.ch/~kai/Papers/phd.pdf>.
- [7] K. Brännler and R. McKinley. An algorithmic interpretation of a deep inference system. In I. Cervesato, H. Veith, and A. Voronkov, editors, *LPAR 2008*, volume 5330 of *Lecture Notes in Computer Science*, pages 482–496. Springer-Verlag, 2008. <http://www.iam.unibe.ch/~kai/Papers/2008aidis.pdf>.
- [8] K. Brännler and A. F. Tiu. A local system for classical logic. Technical report, 2001. <http://www.iam.unibe.ch/~kai/Papers/1cl1-lpar.pdf>.
- [9] P. Bruscoli and A. 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>.
- [10] P. Bruscoli, A. Guglielmi, T. Gundersen, and M. Parigot. A quasipolynomial cut-elimination procedure in deep inference via atomic flows and threshold formulae. In E. M. Clarke and A. 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. URL <http://cs.bath.ac.uk/ag/p/QPNDI.pdf>.
- [11] S. R. Buss. Polynomial size proofs of the propositional pigeonhole principle. *Journal of Symbolic Logic*, 52(4):916–927, 1987.
- [12] A. Das. On the proof complexity of cut-free bounded deep inference. In K. Brännler and G. Metcalfe, editors, *Tableaux 2011*, volume 6793 of *Lecture Notes in Artificial Intelligence*, pages 134–148. Springer-Verlag, 2011. URL <http://www.anupamdas.com/items/PrCompII/ProofComplexityBoundedDI.pdf>.
- [13] A. Das. Complexity of deep inference via atomic flows. In S. B. Cooper, A. Dawar, and B. 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>.
- [14] E. W. Dijkstra. The undeserved status of the pigeon-hole principle. Mar. 1991. URL <http://www.cs.utexas.edu/users/EWD/ewd10xx/EWD1094.PDF>.
- [15] A. 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>.
- [16] A. Guglielmi and T. 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>.
- [17] A. Guglielmi, T. Gundersen, and M. Parigot. A proof calculus which reduces syntactic bureaucracy. In C. Lynch, editor, *RTA 2010*, volume 6 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 135–150. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2010. <http://drops.dagstuhl.de/opus/volltexte/2010/2649>.
- [18] E. 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>.
- [19] E. Jeřábek. A sorting network in bounded arithmetic. *Annals of Pure and Applied Logic*, 162(4):341–355, 2011.
- [20] E. Jeřábek. Proofs with monotone cuts. *Mathematical Logic Quarterly*, 58(3):177–187, 2012.
- [21] J. Krajčůček, P. Pudlák, and A. 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. ISSN 1098-2418. URL <http://dx.doi.org/10.1002/rsa.3240070103>.
- [22] J. Paris and A. Wilkie. δ_0 sets and induction. *Open Days in Model Theory and Set Theory*, W. Guzicki, W. Marek, A. Pelc, and C. Rauszer, eds, pages 237–248, 1981.
- [23] T. Pitassi, P. Beame, and R. Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3:97–140, 1993. ISSN 1016-3328. URL <http://dx.doi.org/10.1007/BF01200117>.
- [24] A. Razborov. Proof complexity of pigeonhole principles. In *Developments in Language Theory*, pages 203–206. Springer, 2002.
- [25] L. Straßburger. Extension without cut. *Annals of Pure and Applied Logic*, 163(12):1995–2007, 2012. URL <http://www.lix.polytechnique.fr/~lutz/papers/psppp.pdf>.