Towards Full Completeness for the Linear Logic of Chu Spaces

Vaughan Pratt*
Dept. of Computer Science
Stanford University
Stanford, CA 94305-2140
pratt@cs.stanford.edu

April 27, 1997

Abstract

We prove full completeness for a fragment of the linear logic of the self-dual monoidal category of Chu spaces over 2, namely that the proofs between semisimple (conjunctive normal form) formulas of multiplicative linear logic without constants having two occurrences of each variable are in bijection with the dinatural transformations between the corresponding functors. The proof assigns to variables domains having at most four elements, demonstrating a uniform finite model property for this fragment. We define a notion of proof function analogous to the notion of truth function, determining a transformation between functors, and show that the transformation denoted by a proof net is dinatural if and only if the proof net is sound, namely acyclic and connected. Proof functions are of independent interest as a 2-valued model of MLL with MIX.

1 Introduction

Whereas ordinary logic axiomatizes theorems, linear logic axiomatizes proofs. The semantic criterion for theoremhood is *validity*: the truth function denoted by a formula is required to be universally true. The semantic criterion for proofhood is *naturality*: the transformation denoted by a proof is required to commute with all morphisms of all objects denotable by variables. Given a proof of the sequent $\varphi \vdash \psi$, the transformation it denotes is between the functors denoted by φ and ψ. When variables may appear both positively and negatively, naturality must be generalized to dinaturality to accommodate the resulting reversal of some arrows in the commuting diagrams witnessing naturality.

In ordinary logic, an axiom system that generates all valid formulas is called complete. In linear logic, an axiom system that generates all proofs denoting a dinatural transformation has been called *fully complete* by Abramsky and Jagadeesan [1], who interpret multiplicative linear logic (MLL) over a category of games as objects and history-free uniform strategies as morphisms and show full completeness of MLL with MIX. But the MIX rule, from $\vdash \Gamma$ and $\vdash \Delta$ infer $\vdash \Gamma, \Delta$, although validated by most naturally arising models of linear logic, is nevertheless omitted from "pure" linear logic. Seeking the real McCoy, Hyland and Ong [6] modify AJ's model to refute MIX by disallowing certain plays they characterize as "unfair."

*This work was supported by ONR under grant number N00014-92-J-1974

This paper proves full completeness of a modest fragment of MLL without MIX interpreted over Chu spaces. As a static object a Chu space (A, R, X) is merely a binary relation $R \subseteq A \times X$. The interest in Chu spaces is in their dynamics: Chu spaces transform via a pair of functions acting covariantly on A but contravariantly on X, as though they were topological spaces having A as the set of points, X as the set of open subsets of A, and R as the membership relation between points and open sets. However there is no requirement that the open sets be closed under arbitrary union and finite intersection, making the category **Chu** of Chu spaces and their continuous functions a full supercategory of the category **Top** of topological spaces. Unlike **Top** however, **Chu** is a self-dual closed category, admitting functors *perp*, $-^{\perp} : \mathbf{Chu}^{\mathrm{op}} \to \mathbf{Chu}$, and *tensor*, $\otimes : \mathbf{Chu}^2 \to \mathbf{Chu}$, which interpret the corresponding basic operations of MLL.

Our interest in Chu spaces as a model of linear logic is two-fold: as a model of concurrency, with linear logic serving as a process algebra [5, 4, 8], and as "universal topology," a generalization and simplification of universal algebra in which the operations of linear logic constitute "pure" versions of their counterparts in more application-specific categories, e.g. direct sum $U \oplus V$, tensor product $U \otimes V$, and dual U^* of vector spaces [9, 10].

All proofs between MLL terms, such as of $p \vdash p$ and $(p \multimap q) \otimes r \vdash p \multimap (q \otimes r)$, have interpretations as dinatural transformations, showing the soundness of MLL for Chu spaces. Here we prove the converse for a fragment of MLL: every dinatural transformation between the interpretations of terms of this fragment interprets some proof.

Our "modest fragment" of MLL restricts to conjunctive normal form (CNF) formulas (called semisimple by Hyland and Ong [6]), those expressible as the tensor product $c_1 \otimes \ldots \otimes c_n$ of *clauses* each formed as a par $c_i = P_1 \parr \ldots \parr P_{m_i}$ of *literals*, namely atoms $P_j = p$ or negated atoms $P_j = p^{\perp}$. A more serious restriction is that we permit formulas to contain at most one occurrence of each sign of each variable.

On the syntactic side these restrictions trivialize linear logic to the point where a proof is just an undirected tree whose vertices are clauses and whose edges connect clauses sharing a variable p whose occurrences in those respective clauses have opposite signs. The edges constitute *axiom links* $p \vdash p$ and can be thought of as specifying steps in a resolution proof. Each linear logic theorem in this fragment has a unique such proof, which can be found simply by locating each complementary pair of variables, and which can be verified to be a proof by checking that the resulting graph is connected and acyclic (the Danos-Regnier condition for this special case).

The semantic side however is much more interesting, even for this limited fragment. We show that the dinatural transformations between the terms of this fragment are in bijection with the proof nets of the corresponding linear logic theorems. Two aspects of this proof are of particular interest.

First, we are able to show completeness by assigning to variables only Chu spaces with at most 4 points. Although there are other completeness results where finite objects are assigned to variables, e.g. [7], this is the first for which the objects assigned to variables are of *uniformly* bounded finite size. We expect that this property survives lifting the CNF restriction, but have no expectations either way for when the restriction to two occurrences per variable is lifted.

Second, we introduce the notion of *proof function* analogous to truth function. Semantic theoremhood of a formula is a property of the truth function it denotes, which must be constantly true or *valid*. Semantic soundness of a proof is a property of the proof function it denotes, which must represent a dinatural transformation.

For Boolean theoremhood the set $\{0, 1\}$ is central. Its counterpart for proofs will in this paper be played by the *quad* $\mathcal{Q} = \{0,1,2,3\}$, from which we extract $\{0,1\}$ for our notion of proof function. Both truth and proof functions will be n-ary operations on $\{0,1\}$, i.e. Boolean operations. But whereas the value of a truth function is determined by independent assignments to *variables*, that of our proof function is determined by independent assignments to variable *occurrences*. And whereas the evaluation of a truth function proceeds up the formula from leaves to root, that of a proof function propagates in both directions along axiom links of the proof net. And with regard to the meaning of 0 and 1, for truth functions it is absolute but for proof functions it is relative to the truth of neighboring clauses, calling for a relativization or correction phase in the evaluation process.

We do not currently see how to extend our result beyond the binary case. For theorems with four or more occurrences of some variable, elements that are invariant with respect to automorphisms of the quad $\mathcal{Q}$ proliferate exponentially in the number of occurrences. We do not presently know whether this is endemic to the whole of **Chu** or just a limitation of $\mathcal{Q}$.

2 Proof Functions

The path from the definition of Chu spaces and their MLL operations to the characterization of the dinatural transformations between MLL functors is quite tedious. At the start of this path we have the simple idea of a Chu space as just a binary relation transforming like a topological space, with the main Chu space we need, the quad $\mathcal{Q}$, consisting of the integers mod 4 with the consecutive pair topology. At the other end, the dinatural transformations we arrive at can be described in elementary terms as "proof functions" without reference to Chu spaces, functors, or dinatural transformations.

We therefore reverse the natural order of development by treating proof functions now. This way, even if dinaturality proves heavy going (as much the fault of our bad exposition as of the complexity of the material), the relatively straightforward notion of a proof function will have been of interest as well as giving at least some insight into the structure of dinaturals in Chu.

Proof functions constitute a 2-valued model of MLL with MIX, or at least the CNF or semisimple fragment of MLL we treat below, which makes them of independent interest in their own right. Proof functions being just a syntactic formulation of the dinatural transformations which this paper is about, they have all the same limitations, including restriction to two occurrences of each variable. The alternative is to drop this restriction but replace it with the idea of a proof function as a semantic object denoted by a formula together with a proof in the form of a linking of opposite variables. Thus whereas a truth function is denoted by a formula, a proof function is denoted by a proof. The proof is sound just when the proof function exists.

Like truth functions, a proof function is an m-ary Boolean operation $\pi : 2^m \to 2$ denoted by a formula. For this paper we restrict attention to CNF formulas, the kind that resolution theorem provers test for unsatisfiability to decide whether their negation is a theorem. Whereas truth functions can only predict the outcome of a resolution proof, proof functions can be viewed as a two-valued semantics for the resolution process itself that recognizes more explicitly than truth functions such concerns as whether a locally sound argument is globally unsound due to a cycle in the proof structure.

The differences from truth functions are as follows. First, in evaluating a proof function, truth assignments are made not to variables but to *occurrences* of variables. Thus two occurrences of p

may be assigned contradictory values, and the literals p and $p^\perp$ may have the same value.

Second, singleton clauses are evaluated specially. The single-literal clause p evaluates as though it were a doubleton clause $p \parr q$ where p is the same literal while q is a new variable (and hence a positive literal) used nowhere else, such that when assigning values to occurrences, q receives the value intended for p while p receives 0.

Third, evaluation of a proof function entails "systematic lying," formalized by spuriously negating (i.e. changing the sign of) certain literals and clauses. Which ones lie in this way depends on the assignment, subject to the following constraints.

1. Conjunction is undefined off the diagonal, that is, consensus is required. We have $0 \otimes 0 = 0$ and $1 \otimes 1 = 1$ as usual, but $0 \otimes 1$ and $1 \otimes 0$ are undefined. It follows that for the proof function to exist, some clauses may need to lie in order to reach consensus. The consensus is then taken to be the value of the proof function at that assignment.

2. $\pi(0, 0, \ldots, 0) = 0$, an arbitrary choice. Condition 2 then propagates the value at the all-zeros assignment to the other assignments.

Now these constraints do not serve to restrict π very much, since when changing the assignment at occurrences within a clause, literals outside that clause can be changed at will. The following further constraint on proof functions takes into account not only the formula but its accompanying proof. It is possible for a proof function to meet these conditions if and only if the CNF formula in question is a theorem of MLL.

Condition 3 depends on the following notions. Proof functions take an additional parameter besides the formula itself, namely a *mating*, defined as a perfect matching of complementary pairs or *links* $p, p^\perp$ of variable occurrences. It is a folk theorem that the theorems of MLL are exactly those MLL sequents derivable from $(p \multimap p) \otimes (q \multimap q) \otimes \ldots$ by weak distributivity, which is the rule, replace a positive occurrence of $(p \multimap q) \otimes r$ by $p \multimap (q \otimes r)$. The links originate in the $p \multimap p$'s and can understood as being preserved by weak distributivity. For the two-variable fragment we restrict to for dinaturality and the quad, this mating is uniquely determined by the formula and hence redundant, but it is required for proof functions in their own right when this restriction is dropped.

The mate relation determines the inlaw relation. A literal (possibly negated variable) occurring in the same clause as the mate $p^\perp$ of p is called an *inlaw* of p. The mate itself of p is not an inlaw of p. p will be its own inlaw only when its mate is in the same clause. For a singleton clause that has been split into $p \parr q$, p keeps its mate, while q has no mate but is the only inlaw of p's mate. Note that p is no one's inlaw since q is no one's mate.

3. Literals lie just when some inlaw (including its spurious negation if any) evaluates to true (under the present assignment).

When the proof is acyclic, that is, when the transitive closure of the inlaw relation is irreflexive, the determination of which literals lie can then be propagated from inlaw to inlaw, starting with those new literals in singleton clauses that are inlaws (since they cannot have inlaws).

Conversely any failure of reflexivity, i.e. a cycle, makes condition 3 impossible, in particular at any assignment that sets one literal on the cycle to 1 and the rest to 0. If we make 1 honest, then by induction, each 0 around the cycle must lie, forcing 1 to lie, a contradiction. Conversely if we make 1 lie, then by induction, each 0 around the cycle must be honest, forcing 1 to be honest, also a contradiction. Hence in the presence of an inlaw cycle, no choice of lying literals can meet condition 1. The significance of this is that a cylic proof cannot have a proof function. The nontheorems of

CNF MLL are those for which every mating has a cycle (which includes those with no mating at all, i.e. no perfect matching of opposite literals).

We add one more condition.

4. Changing the value assigned to an occurrence cannot change whether or not the clause containing it is lying.

The same invariance applied to the literals in a clause, that whether they lie does not depend on changes made to assignments to variables in the same clause, is a consequence of Condition 3. This is because if there is a proof function at all, the transitive closure of the inlaw relation cannot contain cycles, whence changes to occurrences in a clause must propagate out of the clause immediately and be unable to return.

With these conditions there can be at most one proof function denoted by a proof as a mating of variables.

We illustrate these conditions with the proof function of the theorem $p \vdash p$, with the evident mating. Moving the consequent to the left hand side yields $p \otimes p^{\perp} \vdash$, a CNF formula consisting of two singleton clauses p, $p^{\perp}$. Expanding each clause with a dummy variable yields $(p \parr q) \otimes (p^{\perp} \parr r)$.

We now consider the four assignments 00, 10, 11, 01 in turn. For each, q and r receive the respective values of the assignment while the two occurrences of p both receive 0 independently of the assignment.

For 00, both q and r receive 0 and so do not affect the occurrences of those literals they are the inlaws of. The clauses evaluate to $0 \parr 0 = 0$ and $0^{\perp} \parr 0 = 1$. Taking $\pi(0,0)$ to be 0, we force consensus by complementing the second clause to give $(p \parr q) \otimes (p^{\perp} \parr r)^{\perp}$.

For 10, change the first occurrence of p to 1, actually q. But q is $p^{\perp}$'s inlaw, so complement $p^{\perp}$. The clauses then evaluate to $0 \parr 1 = 1$ and $(0 \parr 0)^{\perp} = 1$. Since the changed variable was in the first clause, we take $\pi(1,0) = 1$. We already have consensus so no further change to the clause signs is needed, finalizing the formula as $(p \parr q) \otimes (p \parr r)^{\perp}$ for this assignment.

For 11, change the second occurrence of p to 1, actually r. This being an inlaw of the first p, we complement that p and evaluate: $0^{\perp} \parr 1 = 1$ and $(0 \parr 1)^{\perp} = 0$. We are now tracking the second clause so we take $\pi(1,1) = 0$. For consensus we move to $(p^{\perp} \parr q)^{\perp} \otimes (p \parr r)^{\perp}$.

For 01, change the first p, or rather q, back to 0. We therefore complement the second p. Evaluation yields $(0^{\perp} \parr 0)^{\perp} = 0$, $(0 \parr 1)^{\perp} = 0$. So $\pi(0,1) = 0$ and the formula giving it is $(p^{\perp} \parr q)^{\perp} \otimes (p^{\perp} \parr r)^{\perp}$.

So in summary the proof function, evaluated in the order 00, 01, 10, 11, is 0010.

Note that the final formula arrived at for each assignment has the property that it agrees with π at that assignment. Also note that if we had moved directly from 00 to 01 we would have arrived at the same formula. This is an instance of the central consistency property of this model, that all ways of starting from an arbitrarily chosen value for π at one assignment and arriving at another assignment yield the same modified formula at that assignment.

For a bigger example take the MLL theorem $p \otimes (p \multimap q) \vdash q$, again with the evident mating. The corresponding CNF formula to be refuted is $q^{\perp} \otimes p \otimes (q \parr p^{\perp})$ (with variables reordered for convenience). Expanding singletons gives $(q^{\perp} \parr r) \otimes (p \parr s) \otimes (q \parr p^{\perp})$. Here we have 16 assignments, 0000,0001,...,1111. For the first assignment, the first q's inlaw, $p^{\perp}$, is 1, so we complement q. The clauses then evaluate as $0 \parr 0 = 0$, $0 \parr 0 = 0$, and $0 \parr 1 = 1$. We pick $\pi(0,0,0,0) = 0$ and bring the third clause into line by complementing it to give $(q \parr r) \otimes (p \parr s) \otimes (q \parr p^{\perp})^{\perp}$.

At 0001, the first q's inlaw, $p^\perp$, becomes false, so we restore that q to its old sign to give $(q^\perp \parr r) \otimes (p \parr s) \otimes (q \parr p^\perp)^\perp$. The clauses evaluate to $1 \parr 0 = 1$, $0 \parr 0 = 0$, $(0 \parr 0)^\perp = 1$ or 1,0,1. Since we changed a variable in the third clause, we use its value to get $\pi(0,0,0,1) = 1$, and complement the middle clause to bring it into line giving $(q^\perp \parr r) \otimes (p \parr s)^\perp \otimes (q \parr p^\perp)^\perp$.

Jumping ahead to 0100, the second p's inlaw, s, becomes true so we complement that p. The clauses are $1 \parr 0 = 1$, $0 \parr 1 = 1$, $0 \parr 1 = 1$. Since 0100 changes the first p in 0000, we go back to that assignment, take the middle clause $(p \parr s)$, evaluate it here to get 1, and take $\pi(0,1,0,0) = 1$ so that π behaves like the middle clause in the passage from 0000 to 0100, since the occurrence that changed is in the middle clause. For the details of the remaining assignments the reader may ftp, compile and run the C program ftp://boole.stanford.edu/pub/proofun.c.

Continuing in this way, we obtain the proof function 0100101101000100 whose respective bits are the values of π at 0000 through 1111. We will return to this example at the end of the paper.

Note that when the proof net is acyclic, removing any axiom link partitions the net. The two sign changes applied to each end of any link can then be seen to be computed from the respective nets obtained by removing that link, with each end getting its information from the net it is *not* attached to. In this sense data flows in both directions along each link.

3 Chu spaces

Proof functions express the operational essence of dinatural transformations between MLL functors in the category of Chu spaces, which we now treat.

A *Chu space* $\mathcal{A} = (A, R, X)$ *over* the set $2 = \{0, 1\}$ consists of sets A and X and a binary relation $R \subseteq A \times X$. We interpret A and X as the carrier and cocarrier of $\mathcal{A}$, consisting of respectively *points* $a, b, \ldots$ and *states* $x, y, \ldots$.

We write R_a for $\{x \in X \mid R(a, x)\}$ and R^x for $\{a \in A \mid R(a, x)\}$, called respectively *row* a and *column* x. Rows and columns are to be understood as representations of points and states respectively. A Chu space no two states (respectively points) of which have the same representation is called *extensional* (respectively *separated* or T_0); with both properties it is called *biextensional*. (So an extensional Chu space whose columns are closed under arbitrary union and finite intersection constitutes a topological space whose open sets are the columns.) In this paper we shall work mainly with biextensional spaces.

A Chu space is called *discrete* when it is extensional and every subset of its carrier appears as a column, and *coherent* when it is separated and every subset of its cocarrier appears as a row.

As examples of Chu spaces we give the two spaces we shall be assigning as values of variables in the proof of the main theorem. The smaller is the discrete space $1 = (1, \in, 2)$ where 1 denotes $\{0\}$, 2 its power set, and $\in$ the membership relation between them. The larger is $\mathcal{Q} = (4, Q, 4)$ where $4 = \{0, 1, 2, 3\}$ and $Q(a, x)$ is the relation $a \in \{x, x+1\}$. $\mathcal{Q}$ is neither discrete, lacking 12 states, nor coherent, lacking 12 points.

The *dual* of $\mathcal{A} = (A, R, X)$, denoted $\mathcal{A}^\perp$, is $(X, R^\smile, A)$ where the converse $R^\smile$ is defined as $R^\smile(x, a) = R(a, x)$. Duality interchanges points and states, rows and columns, separation and extensionality, and discreteness and coherence. Clearly $\mathcal{A}^{\perp\perp} = \mathcal{A}$.

Statically a Chu space is just a binary relation. The intrinsic interest in Chu spaces only emerges when its transformational structure, or some adequate reflection thereof, is included in the picture, as follows.

A *Chu transform*, or just *map*, from (A, R, X) to (B, S, Y) is an *adjoint pair* $(f, f^\perp)$ of functions $f{:}A{\to}B$, $f^\perp{:}Y{\to}X$, defined as satisfying $S(f(a), y) = R(a, f^\perp(y))$ for all $a \in A$ and $y \in Y$. The composite

$$(A, R, X) \stackrel{(f,f^\perp)}{\longrightarrow} (B, S, Y) \stackrel{(g,g^\perp)}{\longrightarrow} (C, T, Z)$$

defined by $(g, g^\perp)(f, f^\perp) = (gf, f^\perp g^\perp)$ is a Chu transform because

$$T(gf(a), z) = S(f(a), g^\perp(z)) = R(a, f^\perp g^\perp(z)).$$

Hence Chu spaces and their maps so composed form the category **Chu**.[1]

When the domain of a map is extensional, it may be seen that $f^\perp{:}Y{\to}X$ is determined by $f{:}A{\to}B$, in which case $(f, f^\perp)$ may be abbreviated to f without ambiguity as to $f^\perp$. Dually when its target is separated, $f^\perp$ determines f. When (A, R, X) is discrete, the maps to (B, S, Y) are exactly the functions $f{:}A{\to}B$, with the evident dual when (B, S, Y) is coherent.

Duality defines a contravariant functor on **Chu** in the evident way, turning each map $(f, f^\perp) : \mathcal{A} \to \mathcal{B}$ around to become $(f^\perp, f) : \mathcal{B}^\perp \to \mathcal{A}^\perp$.

For the case of topological spaces represented as above, the adjointness condition for $(f, f^\perp)$ is readily seen to be equivalent to continuity of f (and f determines $f^\perp$ by extensionality). In this sense a Chu space is a generalized topological space.

As other examples of Chu transforms we enumerate all morphisms involving the spaces 1 and $\mathcal{Q}$. Since 1 is discrete, maps from 1 to (A, R, X) may be identified as usual with the elements of A; in particular 1 has one endomorphism, and four maps from it to $\mathcal{Q}$. However there are no maps from $\mathcal{Q}$ to 1 because the latter has an empty column, whose inverse image (reasoning topologically) must be empty, but $\mathcal{Q}$ has no empty columns.

This leaves just the endomorphisms of $\mathcal{Q}$ to consider. Now every state x of $\mathcal{Q} = (4, Q, 4)$ contains exactly one of x and its "complement" $x + 2 \bmod 4$. Hence any endomorphism $f{:}\mathcal{Q}{\to}\mathcal{Q}$ must preserve complements (since if not then $f(a)$ and $f(a+2)$ will appear together in some state x, but then the state $f^\perp(x)$ must contain both x and $x+2$, but $\mathcal{Q}$ has no such state). That is, the two diameters of $\mathcal{Q}$, $\{0, 2\}$ and $\{1, 3\}$, are sent to diameters. Each of the two diameters can be mapped to a diameter in four ways, for a total of 16 ways, these being all and only the endomorphisms of $\mathcal{Q}$. Half of these are bijections, constituting the automorphism group of the quad, which can be seen to form the *dihedral group* D_4 of rotations and inversions of the square.

In summary we have the following cardinalities of homsets.

Hom	$bf1$	$\mathcal{Q}$
$bf1$	1	4
$\mathcal{Q}$	0	16

Tensor Product. The *tensor product* of Chu spaces $\mathcal{A} = (A, R, X)$ and $\mathcal{B} = (B, S, Y)$, denoted $\mathcal{A} \otimes \mathcal{B}$, is defined as $(A \times B, T, |\mathcal{A} {\multimap} \mathcal{B}|)$ where $|\mathcal{A} {\multimap} \mathcal{B}|$ denotes the set of maps $(f, f^\perp) : \mathcal{A} \to \mathcal{B}^\perp$ and T is the relation defined by $T((a, b), (f, f^\perp) = S(b, f(a)) = R(a, f^\perp(b))$.

[1]This is the case $V = \mathbf{Set}$, $k = 2$ of the category $\mathbf{Chu}(V, k)$ of V-enriched Chu spaces over an object k of symmetric closed category V, defined by M. Barr and studied by P. Chu [2]. In this more general setting A and X are objects of V and $R : A \otimes X \to k$ is a morphism of V.

Now a column $T^{(f,f^\perp)}$ of $\mathcal{A} \otimes \mathcal{B}$, as a subset of $A \times B$, is a binary relation between A and B, equivalently an $A \times B$ matrix over 2. Define the a-th row of this matrix, denoted $T_a^{(f,f^\perp)}$, to be $\{b \mid T((f,f^\perp),(a,b))\}$, a subset of B, and similarly the b-th column, $T_b^{(f,f^\perp)}$, as $\{a \mid T((f,f^\perp),(a,b))\}$, a subset of A.

The condition $T((a,b),(f,f^\perp) = R(a,f^\perp(b))$ can then be rephrased as

$T_b^{(f,f^\perp)} = R_{f^\perp(b)}$, while the condition $T((a,b),(f,f^\perp) = S(b,f(a))$ can similarly be expressed as $T_a^{(f,f^\perp)} = S_{f(a)}$. But $R_{f^\perp(b)}$ and $S_{f(a)}$ are columns of respectively $\mathcal{A}$ and $\mathcal{B}$. This yields a relatively uncluttered reformulation of the condition, namely that every column of $\mathcal{A} \otimes \mathcal{B}$, when organized as an $A \times B$ matrix, has for its columns columns of $\mathcal{A}$, and for its rows columns of $\mathcal{B}$.

The roles of f and $f^\perp$ here are to specify respectively which column of B provides row a of $T^{(f,f^\perp)}$, and which column of A provides column b.

Each column of $\mathcal{A} \otimes \mathcal{B}$ therefore uniquely determines f when $\mathcal{B}$ is extensional and $f^\perp$ when $\mathcal{A}$ is extensional. Hence *tensor product preserves extensionality.*

Functoriality. Thus far we have only defined the object part of the functors perp and tensor. We now describe the maps of $\mathcal{A}^\perp$ and $\mathcal{A} \otimes \mathcal{B}$ induced by maps of $\mathcal{A}$ and $\mathcal{B}$.

Perp is easy: the map $(f,f^\perp) : \mathcal{A} \to \mathcal{B}$ induces the dual map $(f^\perp, f) : \mathcal{B}^\perp \to \mathcal{A}^\perp$, evidently a Chu transform by the symmetry of the definition of Chu transform.

For tensor, maps $(f,f^\perp) : (A,R,X) \to (A',R',X')$ and $(g,g^\perp) : (B,S,Y) \to (B',S',Y')$ induce the map $f \otimes g : \mathcal{A} \otimes \mathcal{B} \to \mathcal{A}' \otimes \mathcal{B}'$ whose covariant function is simply the function sending (a,b) to $(f(a),g(b))$. The contravariant function is the function from $|\mathcal{A}' \multimap {B'}^\perp|$ to $|\mathcal{A} \multimap \mathcal{B}^\perp|$ sending $h{:}A' \to Y'$ to $A \xrightarrow{f} A' \xrightarrow{h} Y' \xrightarrow{g^\perp} Y$, i.e. $g^\perp h f$. We need to verify that this composite belongs to $|\mathcal{A} \multimap \mathcal{B}^\perp|$, i.e. is a Chu transform from $\mathcal{A}$ to $\mathcal{B}^\perp$. For this it suffices to produce an adjoint; we claim that $X \xleftarrow{f^\perp} X' \xleftarrow{h^\perp} B' \xleftarrow{g} B$, i.e. $f^\perp h^\perp g$, is an adjoint, which follows from

$$S(g^\perp hf(a),b) = S'(hf(a),g(b)) = R'(f(a),h^\perp g(b)) = R(a,f^\perp h^\perp g(b)).$$

4 Dinaturality

Dinaturality extends naturality to accommodate variables of both variances in the same pair of functors. This is achieved straightforwardly by treating the positive (covariant) and negative (contravariant) variables as independent. When variables $a_1, a_2, \ldots$ only appear covariantly, we write $F(a_1, a_2, \ldots)$, but when they can appear with both variances we write $F(a_1, a_2, \ldots; a_1, a_2, \ldots)$ to separate them into respectively their negative and positive occurrences.

Following this convention for accommodating both variances, arrows of the form $F(a) \xrightarrow{F(f)} F(b)$ in a naturality diagram of the kind shown on the left below are split as $F(a,a) \xleftarrow{F(f,a)} F(b,a) \xrightarrow{F(b,f)} F(b,b)$ to yield the hexagonal diagram on the right.

$$
\begin{tikzcd}
F(a) \arrow[r, "\tau_a"] \arrow[d, "F(f)"'] & G(a) \arrow[d, "G(f)"] \\
F(b) \arrow[r, "\tau_b"] & G(b)
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
F(a,a) \arrow[r, "\tau_a"] & G(a,a) \arrow[d, "G(a,f)"] \\
F(b,a) \arrow[u, "F(f,a)"] \arrow[d, "F(b,f)"'] & G(a,b) \\
F(b,b) \arrow[r, "\tau_b"] & G(b,b) \arrow[u, "G(f,b)"']
\end{tikzcd}
$$

Linear logic permits dinaturality to be understood in more elementary terms, in that it permits trivializing either of F or G. This is because the antecedent can be assumed vacuous (the sequent is $\vdash \psi$, equivalently $\psi^\perp \vdash$ which we shall sometimes find more convenient). Secondly, it will simplify language to assume (as is the case for **Chu**) that the functors F and G are valued in a concrete category.

Combining these two, so that we have just the one functor F on the right say, a dinatural transformation then becomes just a distinguished element of every concrete object $F(a_1, a_2, \ldots; b_1, b_2, \ldots)$ in the image of F, with every tuple of morphisms $f_i{:}a_i{\to}a'_i$ and $g_i{:}b_i{\to}b'_i$ determining a function $F(f_1, f_2, \ldots; g_1, g_2, \ldots) : F(a'_1, a'_2, \ldots; b_1, b_2, \ldots) \to F(a_1, a_2, \ldots; b'_1, b'_2, \ldots)$ which is required to preserve the distinguished point.

As an alternative simplification it will sometimes be possible to find morphisms $f_i : a_i \to b_i$ witnessing nondinaturality that are isomorphisms. Isomorphisms have the advantage that they can be understood as running in either direction; the up arrows in the hexagon above can then be reversed to make it an ordinary square. In such situations it is often possible to restrict to automorphisms $f_i : a_i \to b_i$, there being no reason to keep multiple copies of isomorphic objects..

With these simplifications, the difference between validity in order logic and dinaturality in categorical logic can then be described as follows. With validity one assigns a single point of a space to each variable, evaluates the formula and tests whether the single output value is *true*, and then varies the input values to ensure that the output value never changes.

With dinaturality, one instead assigns a whole space $a_1, a_2, \ldots$ at once to each variable $x_1, x_2, \ldots$, and the formula $\psi(x_1, x_2, \ldots)$ then evaluates to a whole space $F(a_1, a_2, \ldots)$ where F is the functor interpreting the formula ψ. Any collection of automorphisms $f_1{:}a_1{\to}a_1,\ f_2{:}a_2{\to}a_2, \ldots$ applied to the input spaces induces an automorphism of $F(a_1, a_2, \ldots)$. (Contravariant variables in F are accommodated by using the inverse of f when a contravariant variable has value f.) Whereas validity requires *the* single output value to be invariant, dinaturality requires that the output space contain *an* invariant point, namely a common fixpoint of all such induced transformations.

5 CNF Evaluation for the Quad

We illustrated the two basic operations of MLL, $\mathcal{A}^\perp$ and $\mathcal{A} \multimap \mathcal{B}$, for $\mathcal{A} = \mathcal{B} = \mathcal{Q}$ only. We now describe the values of larger formulas when their variables are all set to $\mathcal{Q}$, confining ourselves to semisimple (DNF) formulas in the sense of Hyland and Ong [6]. We assume that all negations $\mathcal{A}^\perp$ have been "pushed down" to the literals, leaving just $\otimes$ and $⅋$ as the two operators higher up in the formula. Since $\mathcal{Q}$ is self-dual we can assume all literals to be evaluated at the quad. (Positive and negative literals are therefore not distinguishable by the spaces they evaluate to but only by how they transform.) We write $\mathcal{Q}^n$ for the n-fold tensor product $\mathcal{Q} \otimes \ldots \otimes \mathcal{Q}$.

Proposition 1 *$\mathcal{Q}^n = (2.D, X)$ where $D = 2^n$, $2.D$ denotes two copies of D, the first copy of which contains elements $(d,1)$ and the second $(d,0)$, and X consists of the 2^{2^n} graphs of characteristic functions of subsets of D.*

Hence for any $x \in X$ representing a subset of D, $|x| = |D|$, and for all $d \in D$, exactly one of $(d,0)$ and $(d,1)$ is in x, namely $(d,1)$ if d is in that subset, and $(d,0)$ otherwise. For $n = 1$, $D = \{0,1\}$ (taking 2^n to be either $\{0,1,\ldots,2^n-1\}$ or the power set of $\{0,1,\ldots,n-1\}$ as convenient), whence $2.D$ consists of $(0,1)$, $(1,1)$, $(0,0)$, and $(1,0)$, corresponding to 0,1,2,3 respectively in $\mathcal{Q}$, with the four subsets of D corresponding to the states of $\mathcal{Q}$. For $n = 2$, $D = \{0,1,2,3\}$, whence $2.D$ consists of $(0,1)$, $(1,1)$, $(2,1)$, $(3,1)$, $(0,0)$, $(1,0)$, $(2,0)$, $(3,0)$, corresponding to the 8 points of $\mathcal{Q}^2$, and there are 16 ways of choosing four complementary pairs from these, corresponding to the 16 states of $\mathcal{Q}^2$.

Proof: We proceed by induction on n. The case $n = 1$ appears in the example above. Now consider $(\mathcal{Q}^{n+1})^\perp = \mathcal{Q} \multimap (\mathcal{Q}^n)^\perp$. This has $(2^{2^n})^2 = 2^{2^{n+1}}$ points, namely all functions sending each of 0 and 1 to one of the 2^{2^n} states of $\mathcal{Q}^n$ (and hence 2 and 3 to the respective complements). There are $4.2.2^n$ pairs (a,y) for $a \in Q$ and $y \in \mathcal{Q}^n$, but $[a,y] = [a+2, \bar{y}]$ (where $\bar{y} = \{(a, 1-i) \mid (a,i) \in y\}$), so $(\mathcal{Q}^{n+1})^\perp$ only has $2.2.2^n = 2.2^{n+1}$ states. Transposing yields $\mathcal{Q}^{n+1}$. ∎

We now wish to characterize the Chu spaces formed as a par of k such tensor powers, namely $Q^{m_1} \parr Q^{m_2} \parr \ldots \parr Q^{m_k}$. It is more natural to work with its dual $(Q^{m_1})^\perp \otimes (Q^{m_2})^\perp \otimes \ldots \otimes (Q^{m_k})^\perp$, a Chu space (A, X) for which A is a quotient of the cartesian product $\prod_{i=1}^k 2^{2^{m_i}}$ and X is a set of subsets of A, which we will identify with their characteristic functions $f(x_1, \ldots, x_k)$ where each x_i ranges over the points of $(Q^{m_i})^\perp$. This would appear to make the domain of f $\prod_{i=1}^k 2^{2^{m_i}}$ itself rather than a quotient thereof. However our treatment of these f's commutes with this quotienting, which we can therefore assume has been postponed as long as we like or even not done at all.

As a consequence of the definition of $\otimes$, f is a state of X if and only if, for all $i = 1, \ldots, k$, and for all values of x_j, $j \neq i$, the value of f as a function of x_i alone is the characteristic function of a state of $(\mathcal{Q}^{m_i})^\perp$, or point of $\mathcal{Q}^{m_i}$. But this is either $(d,1)$, and hence an ultrafilter of $2^{2^{m_i}}$ as a Boolean algebra, or $(d,0)$ and hence a maximal ideal ("ultraideal" or complement of an ultrafilter) of the same Boolean algebra. In the former case it will be the set of all subsets of 2^{m_i} containing d, in the latter the complement of that set. (Everything is finite here.)

We may visualize A as a k-dimensional brick, the i-th side of which is of length $2^{2^{m_i}}$, and each state f as a two-color "painting" of this brick in which every line parallel to the i-th axis is painted as either an ultrafilter or maximal ideal of $2^{2^{m_i}}$.

Now sides of length 2^{2^m} make A a pretty big brick. The following proposition shows how to shrink it to an exponentially smaller brick in each of its k dimensions. The crucial observation is that when each $(Q^{m_i})^\perp$ is replaced by that subspace $P_i \subseteq (Q^{m_i})^\perp$ consisting of those subsets of D having at most one element, with the usual subspace topology (intersect the original open sets with the subspace), the states of $P_1 \otimes \ldots \otimes P_k$ are in bijection with those of $(Q^{m_1})^\perp \otimes \ldots \otimes (Q^{m_k})^\perp$, via a correspondence making each of the latter states the unique extension of the former.

For example when $k = 2$ and $m_1 = m_2 = 2$, $(Q^2)^\perp$ has 8 states ($D = 2^2$) and $2^D = 16$ points, while $(Q^2)^\perp \otimes (Q^2)^\perp$ has $16.16/2 = 128$ points and 320 states. But if we take P to be the 5-point subspace of $(Q^2)^\perp$ consisting of the empty and singleton subsets of D, still having 8 states (each now intersected with the carrier of P), then $P \otimes P$ has only 25 points but it still has 320 states.

The states of the little brick are characterized as having every line parallel to the i-th axis painted as a singleton if that line had been painted as an ultrafilter in the big brick, or a cosingleton if previously a maximal ideal. Furthermore the one element of the singleton, or missing element of the cosingleton, is an atom of $2^{2^{m_i}}$, namely a singleton subset of 2^{m_i}, by the construction, whose one element is therefore a subset of m_i.

The following lemma has two applications, the first of which is in the verification of the above.

Lemma 1 *A two-dimensional little brick $P_1 \otimes P_2$ has two types of states, st and sutv, where $0 \leq s, u < 2^{m_1}$, $u \neq s$, and $0 \leq t, v < 2^{m_2}$, $v \neq t$. In the first type the 1's are on the s-th row and the t-th column, except at their intersection $f(s,t)$ which is 0. In the second type the 1's are on the s-th row and the t-th column (including $f(s,t)$), with three exceptions: $f(s,v) = f(u,t) = 0$, and $f(u,v) = 1$.*

Some discussion may clarify the situation. We have a $(1+2^{m_1}) \times (1+2^{m_2})$ rectangle whose states are matrices $f(i,j)$ of 0's and 1's such that every row and column, each indexed by $\{*, 0, 1, \ldots, 2^{m_i}-1\}$, where * is for the "1+", is constant-but-one (denoting either a singleton or a cosingleton). Moreover the exceptional entry cannot be in the *-th position of that row or column, whence the *-th position can be relied on to specify whether the row or column is a singleton (when 0) or a cosingleton; we shall view the *-th bit as the "sign bit" of the row or column. Entry $f(*,*)$ serves as the sign bit of the whole of f. It is easily seen (by counting 1's on the rows and noting that exactly one row must be almost all 1's, namely the row for which $f(i,*) \neq f(*,*)$) that exactly $2^m + 2^n$ entries differ from $f(*,*)$.

Note that rows s and u are complementary, as are columns t and v.

Proof: Assume without loss of generality that $f(*,*) = 0$. Now s and t are determined by the position of the exception in respectively the *-th column and *-th row. Hence $f(s,j) = 1$ for all but one column j. If that $j = t$, i.e. $f(s,t) = 0$, then f is now fully determined and we have an *st* state. If $j \neq t$ then $f(s,t) = 1$, and we have an *sutv* state for which v is that j, i.e. $f(s,v) = 0$. So now we must have exactly one row $i \neq s$ such that $f(i,v) = 1$; take u to be that i, i.e. $f(u,v) = 1$. Since $f(i,*) = 0$, $f(i,t) = 0$ as well, and now f is again fully determined. ∎

For $m_1 = m_2 = 2$, s and t can each be any of four values, giving 16 combinations. For each combination there is just one state of the first type. For the second type, u and v can each be any of the three remaining values, giving 9 combinations. These plus the single first-type state add up to 10, giving a total of 160 states with $f(*,*) = 0$, and hence 320 states of both signs, as promised earlier.

We now verify the correspondence between states of the big and little bricks.

Proposition 2 *Let $P_i \subseteq (Q^{m_i})^\perp$, $i = 1, \ldots, k$, be constructed as above, whence $\prod_i P_i \subseteq \prod_i (Q^{m_1})^\perp$ (the $\prod_i$'s being tensor products). Then the set of states f of $\prod_i P_i$ (as functions $f : \prod_i |P_i| \to 2$) and the set of states f^+ of $\prod_i (Q^{m_1})^\perp$ (as functions $f+ : \prod_i |(Q^{m_i})^\perp| \to 2$) are in a bijective correspondence such that each f is the restriction of its corresponding f^+ to $\prod_i |P_i|$, and extends uniquely to that f^+.*

Proof: The restriction of state f^+ to f is a state of $\prod_i P_i$ because the P_i's have the subspace topology. It remains to show that this restriction is a bijection. Given any function f representing a state of $\prod_i P_i$, extend it to $\prod_i (Q^{m_1})^\perp$ by forming $f^+(x_1, \ldots, x_k)$, where each x_i is a subset of $D_i = 2^{m_i}$, inductively as follows. Initially $f^+(s_1, \ldots, s_k) = f(s_1, \ldots, s_k)$ where each s_i is either a singleton or empty (the domain of f). Now suppose we have extended f to f^+ at the first j

arguments, $j \geq 0$, that is, we know $f^+(x_1, \ldots, x_j, s_{j+1}, \ldots, s_k)$ for all subsets x_i of 2^{m_i}, $i \leq j$ and for all s_i's, $j < i \leq n$.

Let $t = f^+(x_1, \ldots, x_j, \emptyset, s_{j+2}, \ldots, s_k)$ ($t = 0$ or 1).

Claim: $f^+(x_1, \ldots, x_j, s_{j+1}, \ldots, s_k) = t$ for all but exactly one singleton (hence nonempty) choice of s_{j+1}, all other arguments in positions $i \neq j+1$ being held fixed.

Before proving the claim we apply it to the construction of f^+. Let $d \in D_{j+1} = 2^{m_{j+1}}$ be the element of the singleton in the claim. For each $x_{j+1} \subseteq D_{j+1}$ define $f^+(x_1, \ldots, xj+1, \ldots, s_k)$ as $\bar{t}$ if $d \in x_{j+1}$ and t otherwise. This construction evidently makes the corresponding line parallel to the j-th axis an ultrafilter if $t = 0$, otherwise a maximal ideal. Furthermore the choice of ultrafilter or maximal ideal is uniquely determined by t and the one singleton s_{j+1} at which f became $\bar{t}$. Moreover the choice agrees with f^+ where already defined at empty-or-singleton values of s_{j+1}, whence this remains an extension of f. We continue in this way until $j = k$ to produce f^+.

It remains to prove the claim. We proceed by induction on j. For $j = 0$ we observed earlier that the lines in the reduced brick are singletons or cosingletons, and moreover the one element of that singleton or missing element of that cosingleton is itself a singleton (and not the empty subset, which is why we chose t at $s_{j+1} = \emptyset$). Now let $j \geq 1$ and assume we have verified the claim up to $j-1$. We extend the claim to j, i.e. $f(x_1, \ldots, x_{j-1}, x_j, s_{j+1}, \ldots, s_k) = t$ for all but one value of s_{j+}, as follows.

Consider the two-dimensional little rectangle obtained by varying s_j and s_{j+1} over the empty-or-singleton and their complements, oriented for convenience of reference so that s_j ranges over rows and s_{j+1} over columns. By the lemma the exceptional element of each column is on row s, unless this is a type *sutv* state in which case in column v it is on row u.

Now consider $f(x_1, \ldots, x_{j-1}, x_j, s_{j+1}, \ldots, s_k)$, corresponding to an expansion of the above two-dimensional little rectangle so that there are rows x_j for all subsets of $D_j = 2^{mj}$ and not just the empty-or-singleton rows s_j. For an *st* state there are two cases, depending on whether $s \in x_j$. For an *sutv* state there are four cases, depending on whether $s \in x_j$ and whether $u \in x_j$. In every case the state at row x_j equals one of four rows of the little rectangle. If u does not exist or x_j does not contain u, then row x_j equals row s or row * depending respectively on whether or not $s \in x_j$. Otherwise row x_j equals the complement of either row s or row * depending respectively on whether or not $s \in x_j$. In all cases this is a row meeting the condition in the claim. ∎

(The role of the case $s_{j+1} = \emptyset$ in producing t is essential in the proof. The proposition is falsified by replacing "singleton-or-empty" by "singleton" in the construction of P_i, e.g. $P \otimes P$ where $P \subseteq \mathcal{Q}^2$ acquires 192 spurious states, including the 4×4 identity matrix. $P \otimes P$ must be a 5×5 brick to get the correct count of 320 states.)

The foregoing describes all Chu spaces denoted by all semisimple MLL$^-$ formulas when their variables are assigned the quad. Both points and states of these Chu spaces come in complementary pairs.

We now focus on a single axiom link $p \vdash p$ of a proof net. This link is associated with two clauses, and hence two dimensions of the brick. An automorphism of the object A assigned to variable p induces an automorphism of the whole brick which moves points around within each of the parallel rectangles whose two axes correspond to the respective clauses containing that variable.

The complement of a state is obtained by complementing all bits of f. Without loss of generality we can assume $f(*, *) = 0$, giving one representative of each complementary pair of states. Evidently a state is invariant if and only if its complement is invariant, whence invariants of functors evaluated at the quad must come in complementary pairs. We may eliminate one member of such a pair by permitting the discrete singleton $bf1 = (1, \in, 2)$, where 2 is the power set of 1, to be assigned to variables, along with an arbitrarily chosen morphism $f : bf1 \to \mathcal{Q}$. This setup is similar to the subobject classifier Ω in a topos, which is accompanied by a final object and a morphism from it to Ω singling out an element.

We indicate only briefly and by example how this method works. Taking as our example the theorem $p \otimes (p \multimap q) \multimap q$, the relevant diagram is $1 \otimes (1 \multimap 1) \multimap 1 \longrightarrow 1 \otimes (1 \multimap 1) \multimap \mathcal{Q} \longleftarrow 1 \otimes (1 \multimap \mathcal{Q}) \multimap \mathcal{Q} \longrightarrow 1 \otimes (\mathcal{Q} \multimap \mathcal{Q}) \multimap \mathcal{Q} \longleftarrow \mathcal{Q} \otimes (\mathcal{Q} \multimap \mathcal{Q}) \multimap \mathcal{Q}$. The arrows are the evident maps, e.g. the first is $1 \otimes (1 \multimap 1) \multimap f$, the next $1 \otimes (1 \multimap f) \multimap \mathcal{Q}$, etc. Evaluating the five functors in the diagram yields respectively $bf1$, $\mathcal{Q}$, $\mathcal{Q} \multimap \mathcal{Q}$, $(\mathcal{Q} \multimap \mathcal{Q}) \multimap \mathcal{Q}$, and $\mathcal{Q} \otimes (\mathcal{Q} \multimap \mathcal{Q}) \multimap \mathcal{Q}$. The reader may verify that the dinaturality condition, namely that each of the four maps in this diagram must preserve the invariant element of its domain, constrains this element to be unique in all five objects. Starting with $bf1$, for which this is trivial, and work through to the other end of this zig-zag. Following arrows in the forward direction, this depends on $f : bf1 \to \mathcal{Q}$, and its images under the functors, being injective. Following the two running backwards, the idea is that each functor composes elements of $\mathcal{Q} \multimap \mathcal{Q}$ with f, whose effect is to mark an element of $\mathcal{Q}$, with dinaturality then holding that element fixed. This prevents the 180° rotation that as an element of $\mathcal{Q} \multimap \mathcal{Q}$ on its own is invariant, i.e. dinatural.

The above uses $bf1$ to show that one of the two dinatural transformations denoted by each proof is peculiar to the quad, and is rejected when $bf1$ and an arrow from $bf1$ to the quad is added. But $bf1$ also serves the purpose of rejecting MIX. Lafont and Streicher observed that $bf1^\perp \otimes bf1^\perp \vdash bf1^\perp \parr bf1^\perp$ has no morphism in **Chu** because the left hand side has no states while the right hand side has one state.

6 Induced Automorphisms

We now describe in outline the automorphisms of the big brick induced by an automorphism of variable j in clause i, deferring details to the full paper. The orbits of the induced automorphism are parallel to the axis associated with clause i, and hence remain within each line along that axis, of length $2^{2^{m_i}}$. For a given state of the big brick these lines are painted as either an ultrafilter or a maximal ideal, and as such represent elements (d, s) of $2.D_i$ where $s = 0$ or 1 indicates which copy of D_i this element belongs to.

We associate (d, s) with the point $2s + (j \in d)$ where $j \in d$ is 1 or 0 according to whether or not the subset $d \subseteq m_i$ contains j. (The idea is that d associates one of the two generators 0 and 1 of $\mathcal{Q}$ with each variable in the clause. The m_i bits that distinguish 0 from 2 and 1 from 3 for each variable are all xor'd together into one global sign bit for the whole clause, otherwise we would have $2m_i$ bits of information associated with each line instead of just $1 + m_i$. This behavior is a consequence of the definition of $\mathcal{A} \otimes \mathcal{B}$, resulting from repeated applications of the identification $[a, y] = [a + 2, \bar{y}]$.)

An automorphism of variable j in clause i maps the line painted (d, s) as follows. Write $2s + (j \in d)$ in binary as sb, and let the automorphism carry it to $s'b'$. The induced automorphism of the big brick repaints that line as (d', s') where $d' = d$ if $b' = b$ and otherwise d' and d differ in one element, namely j. For example if 1 ($s = 0$, $b = 1$) is rotated to 2 ($s = 1$, $b = 0$) then s changes so the line is complemented. (Understood as an automorphism of $2^{2^{m_i}}$, this is just Boolean complement of this Boolean algebra.) But b also changes, which replaces ultrafilter-or-ideal d by d'. (This too can be described as a permutation of $2^{2^{m_i}}$, namely as a transposition of its axes d and d'.)

This behavior is an automorphism of the big brick but not of the little brick, since orbits need not be confined to the latter. We could deal with this by extending a state of the little brick to the big one, but this is not necessary because the effect on the state of the little brick can be described straightforwardly as follows.

For our purposes it is convenient to use a brick even smaller than the little brick. Although the *-th row and column are required to avoid spurious states, for $m \geq 2$ they are redundant with regard to *specifying* a state, since the $*$-th row can be reconstructed column by column by taking a vote in each column. For $m = 1$ this is not the case since the remaining two rows or columns are not enough for a vote. On the other hand, in that case rows 0 and 1 necessarily complement each other, whence one of *those* rows is redundant with regard to specifying states. Hence for all $m \geq 1$ a state can be specified on a k-dimensional brick whose i-th side is of length 2^{m_i}, albeit with different codings for $m_i = 1$ and $m_i \geq 2$. We distinguish this *-free brick from the little brick by the name *tiny brick*.

For the tiny brick, a line painted with the singleton $\{d\}$ represents (d, s) with $s = 0$, while the complementary line has $s = 1$. To change s, complement the line, thereby interchanging singletons and cosingletons. To change bit b, change the membership of j in d to produce d': this replaces $\{d\}$ by $\{d'\}$ and likewise for cosingletons.

For example if $m = 2$ and $d = 1$, the singleton $\{1\}$ is painted as the length-4 line 0010, and becomes either 0001 or 1000 depending on whether j is variable 0 (representing the units position) or variable 1 (two's position) in the clause. This repainting is how the induced automorphism of the big brick appears as seen through the limited aperture of the little brick.

The case $m = 1$ requires special treatment. Here lines are of length 2 and so consist of two bits. We define these as membership in the respective open sets $\{1, 2\}$ and $\{2, 3\}$; thus 0,1,2,3 are respectively 00, 10, 11, 01 (picture the bit pattern 00110011... as seen through a two-bit aperture and moving to the right). The low order bit is the sign bit s, in agreement with the sign bit for $m \geq 2$, but the high order bit is not $b (= j \in d)$ but rather $e = b + s \bmod 2$ (because we need the sign bit in order to tell whether the high order bit is the exceptional bit).

7 Invariants

We now combine the previous two sections to deduce which states are invariant under automorphisms.

A state f of the big brick (A, X) is invariant when $fg = f$ for all automorphisms $g : A \to A$ induced by automorphisms of the $\sum_i m_i$ copies of $\mathcal{Q}$ (where the automorphisms of $\mathcal{Q}$'s instantiating the same variable must be the same). With the states-as-paintings visualization, an invariant state is a painting of A every point of which remains the same color under the $2v$ induced automorphisms of A. Since states of the tiny brick uniquely determine states of the big brick, it follows that a state of the tiny brick is invariant under an automorphism of a variable if and only if the corresponding state of the big brick is invariant. Now we cannot speak of an automorphism of the tiny brick, but we can speak of the restriction of fg and f (as predicates on the big brick, where g is an automorphism of the big brick) to the tiny brick. Invariance of a state of the tiny brick under an automorphism of a variable can then be stated as equality of these two restrictions. The visual effect of invariance, that the "painting does not change," is then the same for the big and tiny bricks.

The dihedral group D_4 can be generated by two of its elements, constituting a sufficient "test set" of automorphisms of $\mathcal{Q}$. As generators we select *rotation*, defined as successor mod 4, and *flip*, defined as interchanging 1 and 3 leaving 0 and 2 fixed. (If we identify $q \in \mathcal{Q}$ with the fourth root of unity i^q, rotation is multiplication by i and flip is complex conjugate.)

A state is invariant under automorphisms applied to two or more variables simultaneously if and only if for every choice of a single variable, that state is invariant under all automorphisms of that variable. Hence for v distinct variables we need consider only $2v$ induced automorphisms of the big brick.

Automorphisms of a variable that occurring in clauses i and j induces an automorphism of the big brick whose orbits are confined to the two-dimensional planes containing the lines parallel to the i-th and j-th axes.

With all formulas in DNF, negations occur only at literals. As far as rotation is concerned, $\mathcal{Q}$ and $\mathcal{Q}^\perp$ are indistinguishable, since as objects they are isomorphic (in eight possible ways, namely via the eight automorphisms of $\mathcal{Q}$), and the inverse image map f^{-1}, which one might think at first rotates $\mathcal{Q}^\perp$ backwards, only does so as a map from target to source. As a map from source to target, namely as $(f^{-1})^{-1}$ where the first $-^{-1}$ means inverse image but the second $-^{-1}$ means inverse of a bijection, it rotates $\mathcal{Q}^\perp$ in the same direction as $\mathcal{Q}$ when they are related to each other by one of the above isomorphisms. For definiteness we take the isomorphism that pairs state $\{i, i+1\}$ with point $i+1 \pmod 4$.

For $m = 1$, rotation of either $\mathcal{Q}$ or $\mathcal{Q}^\perp$ cycles through $00 \to 10 \to 11 \to 01 \to 00$. This can be described uniformly as replacing es by $\bar{s}e$. For $m \geq 2$ we switch to binary, i.e. rotation cycles through $00 \to 01 \to 10 \to 11 \to 00$. At $m = 2$, $j = 0$, this translates to cycling through $0001 \to 0010 \to 1110 \to 1101 \to 0001$, or $0100 \to 1000 \to 1011 \to 0111 \to 0100$. For $m = 2$, $j = 1$, the corresponding rotations are $0001 \to 0100 \to 1110 \to 1011 \to 0001$, or $0010 \to 1000 \to 1101 \to 0111 \to 0010$. For $m = 3$, j can be 0, 1, or 2, and so on. Note that in all cases two rotations of a given variable (one 180° rotation) complements the state of the whole brick.

$\mathcal{Q}$ and $\mathcal{Q}^\perp$ are distinguished by *flip*, which leaves points 0 and 2 of $\mathcal{Q}$ invariant, but leaves no states of $\mathcal{Q}$ invariant since it swaps open set $\{0,1\}$ with $\{3,0\}$, and $\{1,2\}$ with $\{2,3\}$.

For $m = 1$, flipping $\mathcal{Q}$ interchanges 10 and 01 leaving 00 and 11 fixed. This can be described as replacing es by se. For $m = 2$, $j = 0$, the basic interchange is 01 (binary for 1) with 11 (binary for 3), i.e. replace ab with $(a+b)b \pmod 2$. This translates to interchanging 0010 with 1101, and 1000 with 0111, all else remaining invariant. For $m = 2$, $j = 1$, we interchange 0100 with 1011 and 1000 with 0111, all else being fixed. Etc.

Flipping $\mathcal{Q}^\perp$ switches 0 with 1, and 2 with 3. For $m = 1$ this replaces es by $\bar{e}s$. For $m = 2$, $j = 0$ this translates to switching 0001 with 0010 and 0100 with 1000, and likewise for their complements.

We now apply these rules to the determination of the invariant points of the formula $p \multimap p$, or $p^\perp \invamp p$, equivalently the invariant states of $p \otimes p^\perp$, when $p = \mathcal{Q}$. Here $m_1 = m_2 = 1$, whence the big brick is 4×4 and the tiny brick is 2×2. All 16 states of the latter are legal; write these as binary digits $abcd$, with ab and cd being lines parallel to the p (first occurrence) axis (with b and d the respective low order or sign bits) and ac and bd parallel to the $p^\perp$ axis (with c and d as the respective sign bits).

Rotation of p sends $abcd$ to $b\bar{a}d\bar{c}$ while rotation of $p^\perp$ sends $abcd$ to $cd\bar{a}\bar{b}$. Composing these sends

$abcd$ to $d\bar{c}\bar{b}a$. Rotational invariance therefore requires $a = d$ and $b = \bar{c}$.

Flipping the first (positive) occurrence sends $abcd$ to $badc$. Flipping the second (negative) occurrence sends $abcd$ to $\bar{a}\bar{b}cd$. Composing these sends $abcd$ to $\bar{b}\bar{a}dc$. Invariance under flipping therefore requires $a = \bar{b}$, $c = d$.

Invariance under rotation and flipping (and hence under all automorphisms of $\mathcal{Q}$) therefore requires $a = \bar{b} = c = d$. This is satisfied when $abcd$ is either 0100 or its complement 1011. (The 0010 we obtained for $p \multimap p$ via proof nets was for the order $acbd$, with $\bar{p}$ changing faster.) Together these constitute the single invariant pair of $p \multimap p$. As elements of $\mathcal{Q} \multimap \mathcal{Q}$ these correspond to the identity map and 180° rotation.

The formula $q ⅋ q^\perp ⅋ p ⅋ p^\perp$ at $p = q = \mathcal{Q}$ denotes a 2^4-point brick. Associating the four variables with axes that stride respectively 8, 4, 2, and 1 bits, its invariant pair may be written 0100101101000100. Here each group of four bits is an invariant state of $p \otimes p^\perp$, while each group can be seen to correspond to one bit of the invariant state 0100 or 1011 of $q \otimes q^\perp$. This example is consistent with the MIX rule. Note how the complementary pair of invariants for the first half provided the 0 and 1 bits for a complementary pair of the second half. Also note how the two constrained each other to make this the only invariant pair.

We now consider one axiom link in a general semisimple formula, between variables k and l in clauses i and j respectively; call this an ij link.

Proposition 3 *For every ij link of a formula, an invariant state must have every ij plane of the tiny brick be of the second type, $sutv$. Furthermore s and u must correspond to subsets d, d' of $D_i = 2^{m_i}$ differing just in k (i.e. they must be aligned along the k axis of the i dimension viewed as a 2^{m_i} hypercube), and similarly t and v must correspond to subsets d, d' of $D_j = 2^{m_j}$ differing just in l (alignment along the l axis within the j axis).*

Proof: (Sketch) We have already shown this for $m_i = 1$, whose unique pair of invariants satisfy the proposition. For larger m_i's the same invariant structure must obtain, but now with $2^{m_i - 1}$ room per axis for d and d' to rattle around in. ∎

This narrows the choices of s and u together down to 2.2^{m_i-1} possibilities, namely 2^{m_i-1} complementary pairs, determined by the choice of $d - \{k\}$ for $d \in D_i = 2^{m_i}$, and likewise t and v together have 2.2^{m_j-1} choices.

But each further axiom link incident on clause i further narrows the choices. So after w axiom links have constrained clause i we are left with 2.2^{m_i-w} choices associated with that clause. In particular when all axiom links have been accounted for, only 2 choices remain, forming a complementary pair, namely the unique invariant pair. But this is 2 choices altogether, not 2 choices per clause, since the clauses are interlocked with each other, similar to the interlocking we saw with the $p^\perp ⅋ p ⅋ q^\perp ⅋ q$ example earlier.

We summarize all this as the main theorem of this paper.

Theorem 1 *For any semisimple formula φ of binary MLL$^-$, proofs of*

$\varphi(p_1, \ldots, p_n)$ are in 1-1 correspondence with invariant pairs of $\varphi(\mathcal{Q}, \ldots, \mathcal{Q})$.

We illustrate this with the example used to illustrate proof functions, namely $p \otimes (p \multimap q) \vdash q$. The proof function we obtained, 0100100011101101, can be laid out more visibly for our purposes as

$\genfrac{}{}{0pt}{}{0100}{1000}\ \genfrac{}{}{0pt}{}{1110}{1101}$, with the vertical direction indexed by the first occurrence of p and the major horizontal direction by $q^\perp$. We consider each of the axiom links $p \vdash p$ and $q \vdash q$ in turn.

For the link $p \vdash p$, since $p^\perp$ is variable 0 in the $q \parr p^\perp$ clause, we can test for invariance under rotation of p by first rotating the $p^\perp$ occurrence using the rotation rules for $m = 2$, $j = 0$, which transform the four rows to give $\genfrac{}{}{0pt}{}{1000}{1011}$ and $\genfrac{}{}{0pt}{}{1101}{0001}$. Now rotate each of the eight columns down to recover what we had, demonstrating invariance under rotation.

For invariance under flipping, first flip $p^\perp$, which by the rules for negative flipping with $m = 2$ and $j = 0$ gives $\genfrac{}{}{0pt}{}{1000}{0100}$ and $\genfrac{}{}{0pt}{}{1101}{1110}$. Now flip p, which simply exchanges the top and bottom rows, restoring what we had.

That completes the link for $p \vdash p$. For the link $q \vdash q$, transpose the 2×2 matrix of rows as $\genfrac{}{}{0pt}{}{0100}{1110}$ and $\genfrac{}{}{0pt}{}{1000}{1101}$, indexed by $q^\perp$ vertically and $q \parr p^\perp$ within groups of four bits.

Since q is variable 1 in the $q \parr p^\perp$ clause we can test for invariance under rotation of q by transforming the four rows to $\genfrac{}{}{0pt}{}{1110}{1011}$ and $\genfrac{}{}{0pt}{}{1101}{0111}$ via the rules for positive flip with $m = 2$, $j = 1$. Now rotate down to recover what we had.

Since $p^\perp$ is variable 0 in the $q \parr p^\perp$, we can test for invariance under rotation of p by first rotating the $p^\perp$ occurrence using the rotation rules for $m = 2$, $j = 0$, which transform the four rows to give $\genfrac{}{}{0pt}{}{1000}{1011}$ and $\genfrac{}{}{0pt}{}{1101}{0001}$. Now rotate each of the eight columns down to recover what we had, demonstrating invariance under rotation.

For invariance under flipping, first flip $p^\perp$, which by the rules for negative flipping with $m = 2$ and $j = 0$ gives $\genfrac{}{}{0pt}{}{1000}{0100}$ and $\genfrac{}{}{0pt}{}{1101}{1110}$. Now flip p, which simply exchanges the top and bottom rows, restoring what we had.

For the link $q \vdash q$, reorganize the four rows as $\genfrac{}{}{0pt}{}{0100}{1110}$ and $\genfrac{}{}{0pt}{}{1000}{1101}$, indexed by $q^\perp$ and $q^\perp \parr q$. Since q is variable 1 in the $q \parr p^\perp$ clause we can test for invariance under rotation of q by transforming the four rows to $\genfrac{}{}{0pt}{}{1110}{1011}$ and $\genfrac{}{}{0pt}{}{1101}{0111}$ via the rules for positive flip with $m = 2$, $j = 1$. Now rotate down to recover what we had.

Flipping positive q with $m = 2$ and $j = 1$ gives $\genfrac{}{}{0pt}{}{1011}{1110}$ and $\genfrac{}{}{0pt}{}{0111}{1101}$. To flip $q^\perp$, simply complement the top row, which again restores what we had.

8 Remarks and Questions

Validity and naturality are dual in certain respects, one of which is complexity. In Boolean logic the naive test of validity of a formula entails evaluating it at exponentially many combinations of inputs. But for naturality it is a consequence of functoriality that an output point is invariant under all possible combinations of transformations of input variables when it is invariant under transformations of individual variables. This is because a transformation of several variables at once is expressible as the composition of transformations applied to individual variables, and if each of those leave the same point fixed then so do their composites. Hence to look for invariants, it suffices to try a generating set of endomorphisms for each space. For the spaces of this paper we need just two generators per variable, calling for a linear number of transformations of the input spaces to be performed.

On the other hand the output space for ordinary Boolean evaluation is a single bit, which is easily monitored. But the semantics of Chu spaces, the domain of interest for this paper, are such that in searching for invariants in the output space, the size of this search space is exponential in the number of variables. What we gained on the swings we lost on the merry-go-round.

We conclude with the following questions.

1. Explain the exponentially many dinaturals that arise for multiple occurrences of a given literal. Are they peculiar to the quad or are they endemic to the linear logic of **Chu**?

2. Is the small brick special to the quad, or does an analogous exponential reduction obtain when evaluating a CNF formula at arbitrary Chu spaces.

3. $\mathbf{Chu}(\mathbf{Set}, K)$ inherits the bicompleteness of **Set**, whence the additive connectives of linear logic are defined, namely as coproduct $A \oplus B$ and product $A\&B$. Full completeness of MALL, MLL together with the additives, is an obvious problem to pursue once full completeness for the multiplicatives alone is understood.

Acknowledgments. This work grew out of a summer collaboration on this problem with Gordon Plotkin, who suggested the quad, and Richard Blute.

References

[1] Samson Abramsky and Radha Jagadeesan. Games and full completeness for multiplicative linear logic. *Journal of Symbolic Logic*, 59(2):543–574, 1994.

[2] M. Barr. *$*$-Autonomous categories*, volume 752 of *Lecture Notes in Mathematics*. Springer-Verlag, 1979.

[3] R. F. Blute. Hopf algebras and linear logic (with appendix by M. Barr). *Mathematical Structures in Computer Science*, 6(2):189–217, April 1996.

[4] V. Gupta. *Chu Spaces: A Model of Concurrency*. PhD thesis, Stanford University, September 1994. Tech. Report, available as ftp://boole.stanford.edu/pub/gupthes.ps.Z.

[5] V. Gupta and V.R. Pratt. Gates accept concurrent behavior. In *Proc. 34th Ann. IEEE Symp. on Foundations of Comp. Sci.*, pages 62–71, November 1993.

[6] J.M.E. Hyland and C.-H.L. Ong. Fair games and full completeness for multiplicative linear logic without the mix-rule. Available by ftp from ftp.comlab.ox.ac.uk as fcomplete.ps.gz in /pub/Documents/techpapers/Luke.Ong, 1993.

[7] Y. Lafont. The finite model property for various fragments of linear logic. Available as http://hypatia.dcs.qmw.ac.uk/authors/L/LafontYGA/papers/model.ps.Z, 1996.

[8] V.R. Pratt. Chu spaces and their interpretation as concurrent objects. In J. van Leeuwen, editor, *Computer Science Today: Recent Trends and Developments*, volume 1000 of *Lecture Notes in Computer Science*, pages 392–405. Springer-Verlag, 1995.

[9] V.R. Pratt. The Stone gamut: A coordinatization of mathematics. In *Logic in Computer Science*, pages 444–454. IEEE Computer Society, June 1995.

[10] V.R. Pratt. Broadening the denotational semantics of linear logic. In *Special Issue on the Linear Logic 96*, volume 3 of *ENTCS (Electronic Notes of Theoretical Computer Science)*, Tokyo, 1996.