Constructive Provability Logic

Robert J. Simmons **Bernardo Toninho**

September 12, 2010
CMU-CS-XX-XXX

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

We present a novel formulation of the modal logic **CPL**, a *constructive logic of provability* that is closely connected to the Gödel-Löb logic of provability. Our logical formulation allows modal operators to talk about both *provability* and *non-provability* of propositions at reachable worlds. We are interested in the applications of **CPL** to logic programming; however, this report focuses on the presentation of **CPL** itself and on its formalization in the Agda programming language.

This material is based upon work supported under an X10 Innovation Award from IBM and a National Science Foundation Graduate Research Fellowship for the first author. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of IBM or the National Science Foundation.

Keywords: modal logic, provability logic, judgmental reconstruction, natural deduction, sequent calculus, iterated inductive definitions

1 Motivation

Consider the following propositions (where "$\supset$" represents implication):

$$\forall x.\, \forall y.\, \mathsf{edge}(x, y) \supset \mathsf{edge}(y, x)$$
$$\forall x.\, \forall y.\, \mathsf{edge}(x, y) \supset \mathsf{path}(x, y)$$
$$\forall x.\, \forall y.\, \forall z.\, \mathsf{edge}(x, y) \supset \mathsf{path}(y, z) \supset \mathsf{path}(x, z)$$

One way to think of these propositions is as rules in a *bottom-up logic program*. This gives them an operational meaning: given some known set of facts, a bottom-up logic program uses rules to derive more facts. If we start with the single fact $\mathsf{edge}(\mathsf{a}, \mathsf{b})$, we can derive $\mathsf{edge}(\mathsf{b}, \mathsf{a})$ by using the first rule (taking $x = \mathsf{a}$ and $y = \mathsf{b}$), and then, using this new fact, we can derive $\mathsf{path}(\mathsf{b}, \mathsf{a})$ by using the second rule (taking $x = \mathsf{b}$ and $y = \mathsf{a}$). Finally, from the original $\mathsf{edge}(\mathsf{a}, \mathsf{b})$ fact and the new $\mathsf{path}(\mathsf{b}, \mathsf{a})$ fact, we can derive $\mathsf{path}(\mathsf{a}, \mathsf{a})$ using the third rule (taking $x = \mathsf{a}$ and $y = \mathsf{b}$, and $z = \mathsf{a}$). Once the only new facts we can derive are facts we already know, we say we have reached *saturation* — this will happen in our example when we have derived $\mathsf{edge}(\mathsf{a}, \mathsf{b})$, $\mathsf{edge}(\mathsf{b}, \mathsf{a})$, $\mathsf{path}(\mathsf{a}, \mathsf{b})$, $\mathsf{path}(\mathsf{b}, \mathsf{a})$, $\mathsf{path}(\mathsf{a}, \mathsf{a})$, and $\mathsf{path}(\mathsf{b}, \mathsf{b})$. Bottom-up logic programming is a very simple and intuitive kind of reasoning, but it has shown to be an elegant and powerful way of representing and solving a large number of problems, especially in the field of program analysis (see [18] for a number of references).

Next, consider the following proposition:

$$\forall x.\, \forall y.\, \mathsf{path}(x, y) \supset \neg\mathsf{edge}(x, y) \supset \mathsf{noedge}(x, y)$$

Intuition says that this is meaningful. In our example above, we can derive $\mathsf{path}(\mathsf{a}, \mathsf{a})$, but we can't possibly derive $\mathsf{edge}(\mathsf{a}, \mathsf{a})$, so we should be able to conclude $\mathsf{noedge}(\mathsf{a}, \mathsf{a})$. A bottom-up logic programming semantics based on *stratified negation* verifies this intuition. In a stratified logic program made up of the four previous rules, we can derive all the consequences of the first three rules until saturation is reached. At this point, we know everything there is to know about facts of the form $\mathsf{edge}(X, Y)$ and $\mathsf{path}(X, Y)$. When considering the negated premise $\neg\mathsf{edge}(x, y)$ in the fourth rule, we simply check the saturated database and conclude that the premise holds if the fact does not appear in the database.

Stratified negation would, however, disallow the addition of the following rule as paradoxical or contradictory:

$$\forall x.\, \forall y.\, \mathsf{path}(x, y) \supset \neg\mathsf{edge}(x, y) \supset \mathsf{edge}(x, y)$$

Why is this rule problematic? Operationally, the procedure we used for stratified negation no longer really makes sense: we reached saturation, then concluded that there was no way to prove $\mathsf{edge}(\mathsf{a}, \mathsf{a})$, then used that conclusion to prove $\mathsf{edge}(\mathsf{a}, \mathsf{a})$. But we had just concluded that it wasn't provable! Stratified negation ensures that we never use the fact that there is no proof of A to come up with a proof of A, either directly or indirectly. However, stratified negation is an odd property: the program consisting of the single rule $\neg\mathsf{prop1} \supset \mathsf{prop2}$ is stratified (we consider prop1 first, and then we consider prop2), and the program consisting of the single rule $\neg\mathsf{prop2} \supset \mathsf{prop1}$ is

also stratified (we consider prop2 first, and then we consider prop1), but the two rules cannot be combined as a single stratified logic program. This sort of problem is a large part of the reason why giving a general and proof-theoretic justification for stratified negation has been elusive.

This report considers the proof theory of a logical system, *constructive provability logic*, that we believe can be used to give a complete and satisfying justification for stratified negation in logic programming. However, logic programming will be used in this report only as a motivating example — the precise relationship between this logic and stratified logic programming will be left for a future paper.

1.1 Foundations, formalization, and Agda

It is always the case that the proof theory of a logic needs to be formalized using some metalogic — usually some assumed and largely informal notion of set-theory-based mathematics that admits, at minimum, induction. Any consistency results for the logic are obviously premised upon consistency of the metalogic; worrying about consistency of the metalogic is generally filed under the label "foundational issues." However, as we will see, in constructive provability logic the metalogic is interwoven with the proof theory in a way that is mostly foreign outside of dependent type theories,[1] making these sorts of foundational issues quite a bit more relevant. As a result of this concern, all the systems and theorems in this paper have been formalized in the dependently typed programming language Agda [10]. The code from this formalization is available online at XXX.

2 A judgmental introduction

We will introduce constructive provability logic in a manner consistent with Pfenning and Davies' judgmental reconstruction of modal logic [13] (itself an interpretation of Martin Löf's 1983 Siena Lectures [9]). This section is not intended to be a complete introduction to the judgmental methodology, and we refer readers to the aforementioned papers for a more complete discussion.

The judgmental methodology carefully maintains a separation between propositions (which we write as A, B, etc.) and judgments J. Propositions are syntactic constructs that are built up from some set of *atomic propositions* ($\mathsf{edge}(\mathsf{a}, \mathsf{b})$ is an example of an atomic proposition) using connectives (implication $A \supset B$, conjunction $A \wedge B$, and disjunction $A \vee B$ are examples of connectives), and judgments are things that are proved using rules of inference. In this methodology, given a proposition A we can talk about giving a proof of the judgment $A\ \mathit{true}$ (i.e. "proving that A is true") or perhaps giving a proof of the judgment $A\ \mathit{false}$ (i.e. "proving that A is false"), or even the judgment that A is true at some specific time t. It isn't really meaningful to "prove A" — if we say such a thing, we usually mean it as shorthand for proving that A is true.

A *hypothetical* judgment $J_1, \ldots, J_n \vdash J$ (where the sequence of J_i are called the *antecedents* and J is called the *consequent*) roughly expresses that the judgment that J has a proof if we assume that there are proofs of the assumptions $J_1, \ldots, J_n$. However, a hypothetical judgment does not necessarily have a set meaning; rather, we *define* the meaning of a hypothetical judgment

[1] One exception is Zeilberger et al.'s recent work on higher-order focusing [7, 19].

by defining three things: (1) a *hypothesis rule*, (2) a *weakening principle*, and a (3) a *substitution principle*. These should ideally flow naturally from our preexisting understanding of what the hypothesis rule should mean. A rich family of logics (we call these the *structural logics*) obey a common set of principles (we use Ψ as an abbreviation for $J_1, \ldots, J_n$):

Definition of the hypothetical judgment in structural logics:
- *Hypothesis rule*: If $J \in \Psi$, then $\Psi \vdash J$.
- *Weakening*[2] *principle*: If $\Psi \subseteq \Psi'$ and $\Psi \vdash J$ then $\Psi' \vdash J$.
- *Substitution principle*: If $\Psi \vdash J$ and $\Psi, J \vdash J'$ then $\Psi \vdash J'$.

These weakening and substitution principles have an interesting character. In one sense, they are the last thing we need to consider, as once the logic is fully defined, they are theorems that we have to prove about the logic as a whole. However, the position of the judgmental methodology is that such principles are also the *first* thing that we need to consider. On a philosophical level, this is because these principles should flow from our understanding of the meaning of the hypothetical judgment. On a practical level, the weakening and substitution principles are necessary to have on hand as we work through the two important "sanity checks" on the rules which define the meaning of a new connective (more on this in a moment).

The meaning of a connective is defined by two sets of rules, the *introduction* and *elimination* rules. In the case of implication $A \supset B$, which we will use as an example, there is one introduction rule and one elimination rule. Introduction rules establish how we can obtain proof of a judgment about a certain proposition — they mention the connective in the conclusion.

$$\frac{\Psi, A\ true \vdash B\ true}{\Psi \vdash A \supset B\ true}\ \supset I$$

Elimination rules establish how we can use proof of a judgment about that proposition — they mention the connective in the premise.

$$\frac{\Psi \vdash A \supset B\ true \quad \Psi \vdash A\ true}{\Psi \vdash B\ true}\ \supset E$$

The two sanity checks are usually called *local soundness* and *local completeness*. Local soundness ensures that the elimination rules are not too strong relative to the introduction rules (or, conversely, that the introduction rules are not too weak).[3] Consider a proof $\mathcal{D}$ of the hypothetical judgment $\Psi \vdash C\ true$ where the "last" rule is an elimination rule. In the running example of implication, this means that the elimination rule is $\supset E$ and there are two subproofs: one proves

[2]The weakening principle actually generalizes the principle commonly called weakening (if $\Psi \vdash J$ then $\Psi, J' \vdash J$) as well as the principles commonly called exchange (if $\Psi, J_1, J_2, \Psi' \vdash J$ then $\Psi, J_2, J_1, \Psi' \vdash J$) and contraction (if $\Psi, J', J' \vdash J$ then $\Psi, J' \vdash J$).

[3]Note that the fact that we call this property local **soundness** indicates a bias towards the introduction rules — local soundness proves that the elimination rules are (locally) sound with respect to the introduction rules, but it also proves the introduction rules are (locally) complete with respect to the elimination rules! Dummett labeled this bias towards the introduction rules the *verificationist* perspective and the opposite bias towards the elimination rules the *pragmatist* perspective [4].

$\Psi \vdash A \supset C\ true$ and the other proves $\Psi \vdash A\ true$. Call these two subproofs $\mathcal{D}_1$ and $\mathcal{D}_2$, respectively.

$$\dfrac{\begin{array}{cc}\mathcal{D}_1 & \mathcal{D}_2 \\ \Psi \vdash A \supset C\ true & \Psi \vdash A\ true\end{array}}{\Psi \vdash C\ true}\ \supset E$$

As the "last" rule was an elimination rule, one of the subproofs must mention the relevant connective (in this case $\mathcal{D}_1$). Local soundness is the property that, if the "last" rule in that premise is an *introduction* rule, then both the introduction rule and the elimination rule are unnecessary — we can reconstruct a proof of the ultimate conclusion by using the premises of the introduction rule and any other premises of the elimination rule. In the case of implication, we can get $\mathcal{D}'$ by applying the substitution principle to the subproofs labeled $\mathcal{D}_2$ and $\mathcal{D}'_1$ below.

$$\dfrac{\dfrac{\begin{array}{c}\mathcal{D}'_1 \\ \Psi, A\ true \vdash C\ true\end{array}}{\Psi \vdash A \supset C\ true}\ \supset I \quad \begin{array}{c}\mathcal{D}_2 \\ \Psi \vdash A\ true\end{array}}{\Psi \vdash C\ true}\ \supset E \qquad \Rightarrow_R \qquad \begin{array}{c}\mathcal{D}' \\ \Psi \vdash C\ true\end{array}$$

Local completeness, on the other hand, ensures that the elimination rules are not too weak relative to the introduction rules (or, conversely, that the introduction rules are not too strong). Whereas local soundness has the form of a reduction or simplification, local completeness has the form of an expansion: we show that, given a proof of the connective, we can obtain enough evidence by applying elimination rules to re-apply the introduction rule and reconstruct the proof. In the expansion below, we get $\mathcal{D}'$ by applying the weakening principle to $\mathcal{D}$.

$$\begin{array}{c}\mathcal{D} \\ \Psi \vdash A \supset B\ true\end{array} \qquad \Rightarrow_E \qquad \dfrac{\dfrac{\begin{array}{c}\mathcal{D}' \\ \Psi, A\ true \vdash A \supset B\ true\end{array} \quad \dfrac{}{\Gamma, A\ true \vdash A\ true}\ hyp}{\Psi, A\ true \vdash B\ true}}{\Psi \vdash A \supset B\ true}$$

2.1 Intuitionistic Kripke semantics (a.k.a. "Simpson-style" modal logic)

Modal logic is an extension of regular logic that initially sought to deal with concepts like "possibility" and "necessity." A popular way of understanding and modeling modal logics is through *Kripke semantics*, which explain the meaning of possibility and necessity in terms of some set of *worlds* and some *accessibility relation*. An accessibility relation determines whether you can get from one world from another world. Then, the judgment "A is possibly true at world w" means that, from w, you can get to some world where A is true, and the judgment "A is necessarily true at world w" means that, from w, *everywhere* you can get to is a world where A is true.

The primary contribution of Alex Simpson's Ph.D. thesis was to show that many intuitionistic modal logics could be given proof-theoretic treatment that strongly resembles Kripke semantics by using a structural logic with two judgments. The first judgment is $A[w]$, which expresses that A is true at a specific "world" w. The other judgment is $w \prec w'$, which expresses that, from world w, world w' is accessible. In Simpson's thesis, worlds are just syntactic things (much like

propositions). One complication is that, in order to talk about the definition of the hypothetical judgment in Simpson-style modal logic, we must extend the form of the hypothetical judgment to account for the fact that we have as antecedents not only judgments $A[w]$ and $w \prec w'$, but also *world variables* ω.

A general introduction to hypothetical judgments parametrized by variables would lead us too far astray; Harper has a complete discussion elsewhere [6, Chapter 4]. Specialized to our needs, the form of the hypothetical judgment is $\Phi \vdash_\Sigma J$, where $\Sigma = \omega_1 \ldots \omega_k$, $\Phi = J_1, \ldots, J_n$, and J along with each of the J_i are either $A[w]$ or $w \prec w'$.

Definition of the hypothetical judgment in Simpson-style modal logic:

— *Hypothesis rule*: If $J \in \Phi$, then $\Phi \vdash_\Sigma J$.
— *Weakening principle*: If $\Sigma \subseteq \Sigma'$ and $\Phi \subseteq \Phi'$ and $\Phi \vdash_\Sigma J$ then $\Phi' \vdash_{\Sigma'} J$.
— *Variable substitution principle*: If $\Phi \vdash_{\Sigma,\omega} J$, then $\Phi[w/\omega] \vdash_\Sigma J[w/\omega]$.
— *Substitution principle 1*: If $\Phi \vdash_\Sigma A[w]$ and $\Phi, A[w] \vdash_\Sigma J$ then $\Phi \vdash J$.
— *Substitution principle 2*: If $\Phi \vdash_\Sigma w \prec w'$ and $\Phi, w \prec w' \vdash_\Sigma J$ then $\Phi \vdash J$.

The hypothetical judgment is only well-formed if every world variable mentioned in Φ or J also appears in Σ. This restricts the weakening and variable substitution principles — it isn't possible to weaken the hypothetical judgment $A[\omega_1] \vdash_{\omega_1} C[\omega_1]$ to $A[\omega_1], \omega_1 \prec \omega_3 \vdash_{\omega_1,\omega_2} C[\omega_1]$ because $\omega_3 \notin \{\omega_1, \omega_2\}$, for instance. It is also not possible to use variable substitution to replace $A[\omega_2] \vdash_{\omega_1,\omega_2} C[\omega_2]$ with $A[\omega_3] \vdash_{\omega_1} C[\omega_3]$, and for the same reason.

Now we can define the meaning of connectives and reason about their soundness and completeness in the same way as we did before. The rules for implication are nearly unchanged, and the local soundness and completeness checks behave much as they did before, so we will not repeat them.

$$\frac{\Phi, A[w] \vdash_\Sigma B[w]}{\Phi \vdash_\Sigma A \supset B[w]} \supset I \qquad \frac{\Phi \vdash_\Sigma A \supset B[w] \quad \Phi \vdash_\Sigma A[w]}{\Phi \vdash_\Sigma B[w]} \supset E$$

The point of this new infrastructure is that it allows us to define new connectives, such as *modal possibility*, written $\Diamond A$. The intended meaning of $\Diamond A$ is that it should be true at a given world if there is some accessible world where A is true.

$$\frac{\Phi \vdash_\Sigma w \prec w' \quad \Phi \vdash_\Sigma A[w']}{\Phi \vdash_\Sigma \Diamond A[w]} \Diamond I \qquad \frac{\Phi \vdash_\Sigma \Diamond A[w] \quad \Phi, w \prec \omega'', A[\omega''] \vdash_{\Sigma,\omega''} C[w']}{\Phi \vdash_\Sigma C[w']} \Diamond E$$

The introduction rule just follows our informal definition above: if a world w' is accessible from a world w and A is true at w', then $\Diamond A$ is true at w. The elimination rule is slightly more complicated. If $\Diamond A$ is true at w and we are trying to prove C at some (potentially different) world w', then it suffices to prove $C[w']$ under the additional assumption that A is true at ω'', where ω'' is a newly introduced world variable that represents an arbitrary world accessible from w.

Local soundness for modal possibility is straightforward, though it uses all three substitution principles: from $\mathcal{D}_3$, variable substitution gives us $\Phi, w_1 \prec w_2, A[w_2] \vdash_\Sigma C[w_3]$. Then, from the first substitution principle along with $\mathcal{D}_2'$ (which is $\mathcal{D}_2$ after the weakening principle is used to add

the premise $w_1 \prec w_2$), we get $\Phi, w_1 \prec w_2 \vdash_\Sigma C[w_3]$. Finally, the second substitution principle along with $\mathcal{D}_1$ gives us $\mathcal{D}'$, a proof of $\Phi \vdash_\Sigma C[w_3]$.

$$\dfrac{\dfrac{\begin{array}{c}\mathcal{D}_1\\ \Phi \vdash_\Sigma w_1 \prec w_2\end{array} \quad \begin{array}{c}\mathcal{D}_2\\ \Phi \vdash_\Sigma A[w_2]\end{array}}{\Phi \vdash_\Sigma \Diamond A[w_1]}\Diamond I \quad \begin{array}{c}\mathcal{D}_3\\ \Phi, w_1 \prec \omega, A[\omega] \vdash_{\Sigma,\omega} C[w_3]\end{array}}{\Phi \vdash_\Sigma C[w_3]}\Diamond E \qquad \Rightarrow_R \qquad \begin{array}{c}\mathcal{D}'\\ \Phi \vdash C[w_3]\end{array}$$

Local completeness for modal possibility is straightforward: the two new pieces of information provided by the $\Diamond E$ rule are precisely what we need to apply the $\Diamond I$ rule.

$$\begin{array}{c}\mathcal{D}\\ \Phi \vdash \Diamond A[w]\end{array} \Rightarrow_E$$

$$\dfrac{\begin{array}{c}\mathcal{D}\\ \Phi \vdash \Diamond A[w]\end{array} \quad \dfrac{\dfrac{}{\Phi, w \prec \omega', A[\omega'] \vdash_{\Sigma,\omega'} w \prec \omega'}\,hyp \quad \dfrac{}{\Phi, w \prec \omega', A[\omega'] \vdash_{\Sigma,\omega'} A[\omega']}\,hyp}{\Phi, w \prec \omega', A[\omega'] \vdash_{\Sigma,\omega'} \Diamond A[w]}\Diamond I}{\Phi \vdash \Diamond A[w]}\Diamond E$$

2.2 Reflection over the accessibility relation

Consider a very simple accessibility relation: there are two worlds α and β, and from α, β is accessible (again, we write this $\alpha \prec \beta$). Then assume that $\Diamond A$ is true at α and that $A \supset B$ is true at β. Should we be able to conclude that B is true at β?

We can represent rephrase this question by asking whether the following judgment has a proof:

$$\alpha \prec \beta, \Diamond A[\alpha], A \supset B[\beta] \vdash_{\alpha,\beta} B[\beta]$$

At the level of a word problem, this seems plausible: $\Diamond A$ is true at α, meaning that A is true at some world accessible from α. As β is the only world accessible from α, you could argue that this means A must be true at β, at which point the rest follows by implication elimination:

$$\dfrac{\dfrac{}{\ldots, A[\beta], A \supset B[\beta] \vdash_{\alpha,\beta} A \supset B[\beta]}\,hyp \quad \dfrac{}{\ldots, A[\beta], A \supset B[\beta] \vdash_{\alpha,\beta} A[\beta]}\,hyp}{\ldots, A[\beta], A \supset B[\beta] \vdash_{\alpha,\beta} B[\beta]}\supset E$$

This reasoning, however, is inconsistent with the principles of the logic. If we can prove the judgment $\alpha \prec \beta, \Diamond A[\alpha], A \supset B[\beta] \vdash_{\alpha,\beta} B[\beta]$, the weakening principle says that we must also be able to prove the judgment $\alpha \prec \beta, \alpha \prec \gamma, \Diamond A[\alpha], A \supset B[\beta] \vdash_{\alpha,\beta,\gamma} B[\beta]$. By weakening the previous hypothetical judgment, there are are now two worlds accessible from α, both β and γ. This no longer seems like a hypothetical judgment that should have a proof, as it might be the case that A was true at γ but not at β. In other words, we should hope that the Simpson-style modal logic doesn't allow us to prove this sequent[4] — if it does that would indicate a problem with the weakening principle we started with!

[4] It doesn't.

2.2.1 Modal logic with a pre-defined accessibility relation

One of the reasons that we formalize logic in the first case is to capture and formalize patterns of natural reasoning. Perhaps we want to be able to formalize the informal reasoning above. It is immediately clear that any logic that captures this argument will have different defining principles that differ from the defining principles of Simpson-style logic. In particular, the weakening principle cannot apply in the same way to judgments about accessibility — if we add new worlds or new connections in the accessibility relation, previously provable judgments may become unprovable.

The way we will deal with this is by completely separating reasoning about accessibility and reasoning about truth-at-a-given-world; we will just assume that there is some preexisting set of worlds w and some preexisting accessibility judgment $w \prec w'$ that the logic inherits. The simple accessibility relation that only has $\alpha \prec \beta$ is one possible accessibility relation, and another is represented by the following diagram, where the arrow from α to β indicates that $\alpha \prec \beta$:

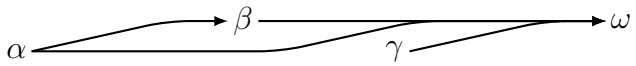

Another accessibility relation might have two worlds and a complete accessibility relation ($\alpha \prec \beta$, $\alpha \prec \alpha$, $\beta \prec \alpha$, $\beta \prec \beta$), and others might have an infinite number of worlds. For instance, the "count-up" accessibility relation has countably infinite worlds arranged like this:

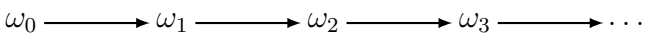

The "count-down" accessibility relation has countably infinite worlds arranged like this:

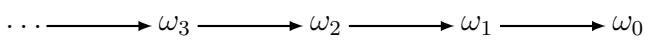

Yet another possibility is this "infinite options" accessibility relation, where α has countably infinite successors and ω countably infinite predecessors:

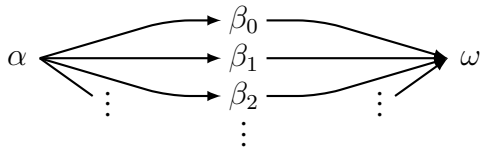

Whatever accessibility relation we use, judgments about accessibility no longer need to appear in the hypothetical judgment, so we can once again define a structural logic. Specializing the original definition to our current logic, we let $\Gamma = A_1[w_1], \ldots, A_n[w_n]$.

Definition of the hypothetical judgment a modal logic with a predefined accessibility relation:
— *Hypothesis rule*: If $A_i[w_i] \in \Gamma$, then $\Gamma \vdash A_i[w_i]$.
— *Weakening principle*: If $\Gamma \subseteq \Gamma'$ and $\Gamma \vdash A[w]$ then $\Gamma' \vdash A[w]$.
— *Substitution principle*: If $\Gamma \vdash A[w]$ and $\Gamma, A[w] \vdash C[w']$ then $\Gamma \vdash C[w']$.

2.2.2 Higher-order formulations of rules

Having set out the principles of the logic, we can talk about connectives. As we will see, this requires us to introduce a significant new idea, rules with premises that are "higher-order." Non-modal connectives (like implication) can be preserved from the Simpson-style modal logic, but the modal operators will, unsurprisingly, need to change. A reasonable introduction rule for modal possibility $\Diamond A$ looks much like it did before, but now the first premise $w \prec w'$ just refers to the predefined accessibility relation:

$$\frac{w \prec w' \quad \Gamma \vdash A[w']}{\Gamma \vdash \Diamond A[w]}\ \Diamond I$$

The elimination rules make things a bit more complicated. Say we're dealing with this accessibility relation:

In a sense, we want our elimination rules to be specific to the world. If we have a proof of $\Diamond A[\alpha]$, we know that at one of the two worlds accessible from α (namely β and ω), A must be true. Therefore, if we can prove $C[w']$ assuming $A[\beta]$, **and** if we can prove $C[w']$ assuming $A[\omega]$, then C is true at w'. This is captured by the following elimination rule:

$$\frac{\Gamma \vdash \Diamond A[\alpha] \quad \Gamma, A[\beta] \vdash C[w'] \quad \Gamma, A[\omega] \vdash C[w']}{\Gamma \vdash C[w']}\ \Diamond E_\alpha$$

Following this strategy, we need an elimination rule for proofs of $\Diamond A$ at each of the other worlds:

$$\frac{\Gamma \vdash \Diamond A[\beta] \quad \Gamma, A[\omega] \vdash C[w']}{\Gamma \vdash C[w']}\ \Diamond E_\beta \qquad \frac{\Gamma \vdash \Diamond A[\gamma] \quad \Gamma, A[\omega] \vdash C[w']}{\Gamma \vdash C[w']}\ \Diamond E_\gamma \qquad \frac{\Gamma \vdash \Diamond A[\omega]}{\Gamma \vdash C[w']}\ \Diamond E_\omega$$

For a logic defined under the accessibility relation above, the introduction rule $\Diamond I$ and the four elimination rules $\Diamond E_\alpha$, $\Diamond E_\beta$, $\Diamond E_\gamma$, and $\Diamond E_\omega$ are, in fact, locally sound and complete. However, it should be obvious that this is not a feasible or scalable way to put together a logic; for instance, we'd need to have an infinite number of rules to handle accessibility relations with an infinite number of objects! (Performing an infinite number of checks for local soundness and completeness is nobody's idea of a good time.)

However, the elimination rules that we wrote for $\Diamond E$ can be generically represented using a *higher-order formulation*. The higher-order formulation of the $\Diamond E$ rule looks like this:

$$\frac{\Gamma \vdash \Diamond A[w] \quad \forall w'.\, w \prec w' \Rightarrow \Gamma, A[w'] \vdash C[w]}{\Gamma \vdash C[w]}\ \Diamond E$$

The second premise quantifies over all worlds w' such that $w \prec w'$ and demands that a proof of $\Gamma, A[w'] \vdash C[w]$ be given for each such w'. We refer to this higher-order formulation as *reflection*

over proofs of another judgment — in this case, we are reflecting over the definition of the accessibility relation. Higher-order formulations are only permissible when we can give a complete definition of the judgment we're reflecting over before we discuss the judgment that uses reflection. In this case, we have already established that we can give a complete definition of the judgment $w \prec w'$ before we say anything about proofs of $\Gamma \vdash A[w]$, so the higher-order formulation is permissible.

If we so desire, we can imagine that the second premise of $\Diamond E$ just takes a given accessibility relation and "macro expands" into as many rules as there are worlds in the accessibility relation. In some cases (the "infinite options" accessibility relation is an example), this means that some rules have an infinite number of premises. In the experience of these authors, that is a difficult concept to wrap one's head around, and it means that even establishing simple properties like local soundness and completeness involve dealing with infinite objects in a way that can be delicate at best. However, it's also not necessary: we can instead treat the arrow "$\Rightarrow$" as "implies" in the sense of "I have to write down a proof of this." We will give an example to show what this means.

2.2.3 Example

In this example, we will use the rules defined above and the four-world accessibility relation given in the previous section. Let $\Gamma_0 = \Diamond A[\alpha], A \supset C[\beta], A \supset B[\omega], A \supset C[\omega]$; we will prove that $\Gamma_0 \vdash \Diamond C[\alpha]$. First, we prove the following theorem:

Theorem 1. *For all w', if $\alpha \prec w'$, then $\Gamma_0, A[w'] \vdash \Diamond C[\alpha]$*

Proof. By case analysis on the accessibility relation, either $w' = \beta$ or $w' = \omega$. If $w' = \beta$, we have the following proof:

$$\dfrac{\dfrac{}{\alpha \prec \beta}\ axiom \quad \dfrac{\dfrac{\overline{\Gamma_0, A[\beta] \vdash A \supset C[\beta]}\ hyp \quad \overline{\Gamma_0, A[\beta] \vdash A[\beta]}\ hyp}{\Gamma_0, A[\beta] \vdash C[\beta]} \supset E}{\Gamma_0, A[\beta] \vdash \Diamond C[\alpha]} \Diamond I$$

If $w' = \omega$, we have the following proof:

$$\dfrac{\dfrac{}{\alpha \prec \omega}\ axiom \quad \dfrac{\dfrac{\overline{\Gamma_0, A[\beta] \vdash B \supset C[\omega]}\ hyp \quad \dfrac{\overline{\Gamma_0, A[\beta] \vdash A \supset B[\omega]}\ hyp \quad \overline{\Gamma_0, A[\beta] \vdash A[\omega]}\ hyp}{\Gamma_0, A[\beta] \vdash B[\omega]} \supset E}{\Gamma_0, A[\omega] \vdash C[\omega]} \supset E}{\Gamma_0, A[\omega] \vdash \Diamond C[\alpha]} \Diamond I$$

This completes the case analysis, and hence the proof. □

Having proved this theorem, we can complete the proof that $\Gamma_0 \vdash \Diamond C[\alpha]$:

$$\dfrac{\overline{\Gamma_0 \vdash \Diamond A[\alpha]}\ hyp \quad \begin{array}{c}\textbf{Theorem 1} \\ \forall w'.\, \alpha \prec w' \Rightarrow \Gamma_0, A[w'] \vdash \Diamond C[\alpha]\end{array}}{\Gamma_0 \vdash \Diamond C[\alpha]} \Diamond E$$

The most significant thing to notice here is that proofs aren't simple tree-like structures anymore; the second premise of $\Diamond E$ in the proof tree above is satisfied not by another proof tree but by a *theorem*. This particular way of understanding higher-order formulations of judgments is not particularly new; our use of it follows Noam Zeilberger's. To slightly misquote Zeilberger's "Focusing and Higher-Order Abstract Syntax,"

> "We hope to make the case that this higher-order formulation should be taken at face value — interpreted constructively, it demands a mapping from proofs of $w \prec w'$ to proofs of $\Gamma, A[w'] \vdash C[w]$" (see [19, p. 361] for the original quote).

This is exactly the point that the example above tries to draw out: the way we prove that there is a mapping from proofs of the judgment $w \prec w'$ to proofs of the judgment $\Gamma, A[w'] \vdash C[w]$ is to, well, prove the statement "for all w', if $w \prec w'$ then $\Gamma, A[w'] \vdash C[w]$." However, because we get to prove this statement using all the familiar machinery of whatever logic we use to prove theorems, our notion of "proof" has gone from a fairly innocent set of trees to something much more complex. This is what we were foreshadowing in the introduction when we said "the metalogic is interwoven with the proof theory." Luckily, the foundation for this kind of system can be found in Martin Löf's theory of *iterated inductive definitions* [8], and these sorts of logical systems can be represented and reasoned about straightforwardly in logical frameworks like Agda. In fact, we would claim that some of the proofs in this paper can be expressed in Agda more naturally (and certainly more concisely) than they can be expressed on paper.

2.3 Reflection over provability

Došen: "not A holds at w if and only if A doesn't hold at any world accessible from w"[3]

$$\frac{\forall w'.\, w \prec w' \Rightarrow \neg(\Gamma \vdash A[w'])}{\Gamma \vdash \not\Box A[w]}\ \not\Box I$$

3 Natural deduction

$$\frac{}{\Gamma, A[w] \vdash A[w]}\ hyp \qquad \frac{\Gamma, A[w] \vdash B[w]}{\Gamma \vdash A \supset B[w]}\ \supset I \qquad \frac{\Gamma \vdash A \supset B[w] \quad \Gamma \vdash A[w]}{\Gamma \vdash B[w]}\ \supset E$$

$$\frac{w \prec w' \quad \Gamma \vdash A[w']}{\Gamma \vdash \Diamond A[w]}\ \Diamond I \qquad \frac{\Gamma \vdash \Diamond A[w] \quad \forall w'.\, w \prec w' \Rightarrow \Gamma \vdash A[w'] \Rightarrow \Gamma \vdash C[w]}{\Gamma \vdash C[w]}\ \Diamond E$$

$$\frac{\forall w'.\, w \prec w' \Rightarrow \Gamma \vdash A[w']}{\Gamma \vdash \Box A[w]}\ \Box I \qquad \frac{\Gamma \vdash \Box A[w] \quad (\forall w'.\, w \prec w' \Rightarrow \Gamma \vdash A[w']) \Rightarrow \Gamma \vdash C[w]}{\Gamma \vdash C[w]}\ \Box E$$

$$\frac{w \prec w' \quad \neg(\Gamma \vdash A[w'])}{\Gamma \vdash \Diamond A[w]}\ \not\Box I \qquad \frac{\Gamma \vdash \not\Box A[w] \quad \forall w'.\, w \prec w' \Rightarrow \neg(\Gamma \vdash A[w']) \Rightarrow \Gamma \vdash C[w]}{\Gamma \vdash C[w]}\ \not\Box E$$

$$\frac{\forall w'.\, w \prec w' \Rightarrow \neg(\Gamma \vdash A[w'])}{\Gamma \vdash \not\Diamond A[w]}\ \not\Diamond I \qquad \frac{\Gamma \vdash \neg A[w] \quad (\forall w'.\, w \prec w' \Rightarrow \neg(\Gamma \vdash A[w'])) \Rightarrow \Gamma \vdash C[w]}{\Gamma \vdash C[w]}\ \neg E$$

Theorem 2 (Substitution)**.**

4 Sequent calculus

Theorem 3 (Cut admissibility)**.**

Theorem 4 (Identity)**.**

Theorem 5 (Equivalence of sequent calculus and natural deduction)**.**

5 Equivalence of sequent calculus and natural deduction

6 Axioms

7 Conclusion

References

[1] Andreoli, J.M.: Logic programming with focusing proofs in linear logic. J. Logic Computation 2(3), 297–347 (June 1992)

[2] Appel, A.W., Melliès, P.A., Richards, C.D., Vouillon, J.: A very modal model of a modern, major, general type system. In: Proceedings of the 34th annual symposium on Principles of Programming Languages (POPL'07). pp. 109–122. ACM (2007)

[3] Došen, K.: Negation in the light of modal logic. In: Gabbay, D.M. (ed.) What is negation?, pp. 77–86. Kluwer Academic Publishers (1999)

[4] Dummett, M.: The Logical Basis of Metaphysics. Harvard University Press, Cambridge, Massachusetts (1991)

[5] Gabbay, D.M.: Modal provability foundations for negation by failure. In: Extensions of Logic Programming. pp. 179–222. Springer LNCS 475 (1991)

[6] Harper, R.: Practical foundations for programming languages (2010), working draft, available online: http://www.cs.cmu.edu/~rwh/plbook/book.pdf

[7] Licata, D.R., Zeilberger, N., Harper, R.: Focusing on binding and computation. In: IEEE Symposium on Logic in Computer Science (2008)

[8] Martin-Löf, P.: *Hauptsatz* for the intuitionistic theory of iterated inductive definitions. In: Fenstad, J.E. (ed.) Proceedings of the Second Scandinavian Logic Symposium. pp. 179–216. North Holland, Amsterdam (1971)

[9] Martin-Löf, P.: On the meanings of the logical constants and the justifications of the logical laws. Nordic Journal of Philosophical Logic 1(1), 11–60 (1996)

[10] Norell, U.: Towards a practical programming language based on dependent type theory. Ph.D. thesis, Chalmers University of Technology (2007)

[11] Norell, U.: Towards a practical programming language based on dependent type theory. Ph.D. thesis, Chalmers University of Technology (2007)

[12] Pfenning, F.: Lecture notes on tethered semantics (Mar 2010), lecture notes from 15-816: Modal Logic

[13] Pfenning, F., Davies, R.: A judgmental reconstruction of modal logic. Mathematical Structures in Computer Science 11, 511–540 (2001), notes to an invited talk at the *Workshop on Intuitionistic Modal Logics and Applications (IMLA'99)*, Trento, Italy, July 1999

[14] Pfenning, F., Simmons, R.J.: Substructural operational semantics as ordered logic programming. In: Proceedings of the 24th Annual Symposium on Logic in Computer Science (LICS'09). pp. 101–110. IEEE Computer Society, Los Angeles, California (Aug 2009)

[15] Reed, J.: A judgmental deconstruction of modal logic (May 2009), submitted for publication

[16] Richards, C.D.: The Approximation Modality in Models of Higher-Order Types. Ph.D. thesis, Princeton University (2010)

[17] Simpson, A.K.: The Proof Theory and Semantics of Intuitionistic Modal Logic. Ph.D. thesis, University of Edinburgh (1994)

[18] Whaley, J., Avots, D., Carbin, M., Lam, M.S.: Using datalog with binary decision diagrams for program analysis. In: Yi, K. (ed.) Proceedings of the 3nd Asian Symposium on Programming Languages and Systems (APLAS'05). pp. 97–118. Springer-Verlag LNCS 3780 (2005)

[19] Zeilberger, N.: Focusing and higher-order abstract syntax. In: Proceedings of the 35th Annual Symposium on Principles of Programming Languages (POPL'08). pp. 359–369. ACM, New York, NY, USA (2008)