Focused proof systems for Intuitionistic and Classical Logics:

Handouts for some lectures on during 11-15 April 2011

ISCL 2011, Bertinoro, Italy
Dale Miller, INRIA Saclay

Abstract: This handout contains a few, specific focused proof systems for intuitionistic logic as well as for classical logic with and without equality and fixed points. These notes are not meant to be self-contained: rather, they will provide some details and references that will be useful during the lectures.

Contents

1 Focusing in intuitionistic logic LJF

[The material in this section is taken from the paper [4].]

A second aspect of focusing proofs is that the synchronous/asynchronous classification of non-atomic formulas must be extended to atomic formulas. The arbitrary assignment of positive (synchronous) and negative (asynchronous) *bias* to atomic formulas can have a major impact on, not the existence of focused proofs, but the shape of focused proofs. For example, consider the Horn clause specification of the Fibonacci series:

$$\mathit{fib}(0,0) \wedge \mathit{fib}(1,1) \wedge \forall n \forall f \forall f' [\mathit{fib}(n,f) \supset \mathit{fib}(n+1,f') \supset \mathit{fib}(n+2,f+f')].$$

If all atomic formulas are given a negative bias, then there exists only one focused proof of $\mathit{fib}(n, f_n)$: this one can be classified as a "backward chaining" proof and its size is exponential in n. On the other hand, if all atomic formulas are given a positive bias, then there is an infinite number of focused proofs all of which are classified as "forward chaining" proofs: the smallest such proof is of size linear in n.

Polarity in intuitionistic logic is defined as follows.

Definition 1 Atoms in LJF are arbitrarily divided between those that are positive and those that are negative. *Positive formulas* are of the following forms: positive atoms, *true*, *false*, $A \wedge^+ B$, $A \vee B$ and $\exists x A$. *Negative formulas* are among negative atoms, $A \wedge^- B$, $A \supset B$ and $\forall x A$. ∎

Notice that if we consider only the "negative connectives" $\wedge^-$, $\supset$, $\forall$ and assign all atoms a negative bias, then LFJ proofs are uniform proofs (with backchaining). Here, goal reduction is the "negative" phase while backchaining is the "positive" phase.

If we allow $\vee$ and $\exists$ in the restricted setting that we described before (in *fohh*), then LJF can derive the completeness of uniform proofs as well.

Describe term representations choices here. Consider simple types only (only implication).

- If the atoms have only negative basis, then simply typed λ-terms are proofs with the familiar "head-normal form": $\lambda x_1 \ldots \lambda x_n.\ (h\ t_1\ \ldots\ t_m)$ $(n, m \geq 0)$.
- If the atoms have only positive basis, then simply type λ-terms are essentially in "administrative normal form": $let x_1 = f_1 \bar{y}_1\ and \ldots x_n = f_n \bar{y}_n int$. Here, $x_1, \ldots, x_n$ are distinct variables and $f_1, \ldots, f_n$ are constants.... (** what about higher-order variables for fs?).

Decision and Reaction Rules

$$\frac{[N,\Gamma] \xrightarrow{N} [R]}{[N,\Gamma] \longrightarrow [R]}\ Lf \qquad \frac{[\Gamma] -_P\!\rightarrow}{[\Gamma] \longrightarrow [P]}\ Rf \qquad \frac{[\Gamma], P \longrightarrow [R]}{[\Gamma] \xrightarrow{P} [R]}\ R_l \qquad \frac{[\Gamma] \longrightarrow N}{[\Gamma] -_N\!\rightarrow}\ R_r$$

$$\frac{[C,\Gamma],\Theta \longrightarrow \mathcal{R}}{[\Gamma],\Theta,C \longrightarrow \mathcal{R}}\ []_l \qquad \frac{[\Gamma],\Theta \longrightarrow [D]}{[\Gamma],\Theta \longrightarrow D}\ []_r$$

Initial Rules

$$\frac{}{[P,\Gamma] -_P\!\rightarrow}\ I_r,\ \text{atomic } P \qquad \frac{}{[\Gamma] \xrightarrow{N} [N]}\ I_l,\ \text{atomic } N$$

Introduction Rules

$$\frac{}{[\Gamma],\Theta,false \longrightarrow \mathcal{R}}\ falseL \qquad \frac{[\Gamma],\Theta \longrightarrow \mathcal{R}}{[\Gamma],\Theta,true \longrightarrow \mathcal{R}}\ trueL \qquad \frac{}{[\Gamma] -_{true}\!\rightarrow}\ trueR$$

$$\frac{[\Gamma] \xrightarrow{A_i} [R]}{[\Gamma] \xrightarrow{A_1 \wedge^- A_2} [R]}\ \wedge^- L \qquad \frac{[\Gamma],\Theta \longrightarrow A \quad [\Gamma],\Theta \longrightarrow B}{[\Gamma],\Theta \longrightarrow A\ \wedge^-\ B}\ \wedge^- R$$

$$\frac{[\Gamma],\Theta,A,B \longrightarrow \mathcal{R}}{[\Gamma],\Theta,A\ \wedge^+\ B \longrightarrow \mathcal{R}}\ \wedge^+ L \qquad \frac{[\Gamma] -_A\!\rightarrow \quad [\Gamma] -_B\!\rightarrow}{[\Gamma] -_{A\ \wedge^+ B}\!\rightarrow}\ \wedge^+ R$$

$$\frac{[\Gamma],\Theta,A \longrightarrow \mathcal{R} \quad [\Gamma],\Theta,B \longrightarrow \mathcal{R}}{[\Gamma],\Theta,A \vee B \longrightarrow \mathcal{R}}\ \vee L \qquad \frac{[\Gamma] -_{A_i}\!\rightarrow}{[\Gamma] -_{A_1 \vee A_2}\!\rightarrow}\ \vee R$$

$$\frac{[\Gamma] -_A\!\rightarrow \qquad [\Gamma] \xrightarrow{B} [R]}{[\Gamma] \xrightarrow{A \supset B} [R]}\ \supset L \qquad \frac{[\Gamma],\Theta,A \longrightarrow B}{[\Gamma],\Theta \longrightarrow A \supset B}\ \supset R$$

$$\frac{[\Gamma],\Theta,A \longrightarrow \mathcal{R}}{[\Gamma],\Theta,\exists yA \longrightarrow \mathcal{R}}\ \exists L \qquad \frac{[\Gamma] -_{A[t/x]}\!\rightarrow}{[\Gamma] -_{\exists xA}\!\rightarrow}\ \exists R \qquad \frac{[\Gamma] \xrightarrow{A[t/x]} [R]}{[\Gamma] \xrightarrow{\forall xA} [R]}\ \forall L \qquad \frac{[\Gamma],\Theta \longrightarrow A}{[\Gamma],\Theta \longrightarrow \forall yA}\ \forall R$$

Figure 1: The Intuitionistic Sequent Calculus LJF. Here, P is positive, N is negative, C is a negative formula or positive atom, and D a positive formula or negative atom. Other formulas are arbitrary. Also, y is not free in Γ, Θ, or R. ∎

2 Focusing in classical logic LKF

[The material in this section is taken from the paper [4] and from [5].]

The inference rules for the LKF focused proof system [4] for classical logic is given in Figure 2.

Structural Rules

$$\frac{\vdash \Theta, C \Uparrow \Gamma}{\vdash \Theta \Uparrow \Gamma, C}\ Store \qquad \frac{\vdash \Theta \Uparrow N}{\vdash \Theta \Downarrow N}\ Release$$

$$\frac{\vdash P, \Theta \Downarrow P}{\vdash P, \Theta \Uparrow \cdot}\ Focus \qquad \frac{}{\vdash \neg P, \Theta \Downarrow P}\ Id \text{ (literal } P\text{)}$$

Introduction of negative connectives

$$\frac{}{\vdash \Theta \Uparrow \Gamma, t^-} \qquad \frac{\vdash \Theta \Uparrow \Gamma, A \quad \vdash \Theta \Uparrow \Gamma, B}{\vdash \Theta \Uparrow \Gamma, A \wedge^- B}$$

$$\frac{\vdash \Theta \Uparrow \Gamma}{\vdash \Theta \Uparrow \Gamma, f^-} \qquad \frac{\vdash \Theta \Uparrow \Gamma, A, B}{\vdash \Theta \Uparrow \Gamma, A \vee^- B} \qquad \frac{\vdash \Theta \Uparrow \Gamma, A}{\vdash \Theta \Uparrow \Gamma, \forall x A}$$

Introduction of positive connectives

$$\frac{}{\vdash \Theta \Downarrow t^+} \qquad \frac{\vdash \Theta \Downarrow A \quad \vdash \Theta \Downarrow B}{\vdash \Theta \Downarrow A \wedge^+ B}$$

$$\frac{\vdash \Theta \Downarrow A_i}{\vdash \Theta \Downarrow A_1 \vee^+ A_2} \qquad \frac{\vdash \Theta \Downarrow A[t/x]}{\vdash \Theta \Downarrow \exists x A}$$

Figure 2: The focused proof system LKF for classical logic. Here, P is positive, N is negative, C is a positive formula or a negative literal, Θ consists of positive formulas and negative literals, and x is not free in Θ, Γ. Endsequents have the form $\vdash \cdot \Uparrow \Gamma$. ∎

Sequents for LKF are divided into *negative sequents* $\vdash \Theta \Uparrow \Gamma$ and *positive sequents* $\vdash \Theta \Downarrow B$, where Θ and Γ are multisets of formulas and B is a formula. (These sequents are formally one-sided sequents: formulas on the left of $\Uparrow$ and $\Downarrow$ are *not* negated as they are in two-sided sequents.) Notice that in this focused proof system, we have reused the term "structural rule" for a different set of rules which formally contains instances of weakening (*Id*) and contraction (*Focus*). Notice

also that in any proof that has a conclusion of the form $\vdash \cdot \Uparrow B$, the only formulas that are to the left of an $\Uparrow$ or $\Downarrow$ occurring in that proof are either positive formulas or negative literals: it is only these formulas that are weakened (in the *Id* rule). The only formulas contracted (in the *Focus* rule) are positive formulas. Thus, although linear logic is not used here directly, non-atomic negative formulas are treated linearly in the sense that they are never duplicated nor weakened in an LKF proof.

Let B be a formula of first-order logic. By a *polarization* of B we mean a formula, say B', where all the propositional connectives are replaced by *polarized versions* of the same connective and where all atomic formulas are assigned either a positive or negative polarity. Thus, an occurrence of the disjunction $\vee$ is replaced by an occurrence of either $\vee^+$ or $\vee^-$; similarly with $\wedge$ and with the logical constants for true t and false f. For simplicity, we shall assume that polarization for atomic formulas is a global assignment to all atomic formulas. Properly speaking, focused proof systems contain *polarized* formulas and not simply formulas.

Theorem LKF is sound and complete for classical logic. More precisely, let B be a first order formula and let B' be a polarization of B. Then B is provable in classical logic if and only if there is an LKF proof of $\vdash \cdot \Uparrow B'$ [4].

Notice that polarization does not affect provability but it does affect the shape of possible LKF proofs. To illustrate an application of the correctness of LKF, we show how it provides a direct proof the following theorem.

Herbrand's Theorem Let B is quantifier-free formula and let $\bar{x}$ be a (non-empty) list of variables containing the free variables of B. The formula $\exists \bar{x} B$ is classically provable if and only if there is a list of substitutions $\theta_1, \ldots, \theta_m$ ($m \geq 1$), all with domain $\bar{x}$, such that the (quantifier-free) disjunction $B\theta_1 \vee \cdots \vee B\theta_m$ is provable (*i.e.*, tautologous).

Proof. The converse direction is straightforward. Thus, assume that $\exists \bar{x} B$ is provable. Let B' be the result of polarizing all occurrences of propositional connectives negatively. By the completeness of LKF, there is an LKF proof Ξ of $\vdash \exists \bar{x} B \Uparrow \cdot$. The only sequents of the form $\vdash \Theta \Uparrow \cdot$ in Ξ are such that Θ is equal to $\{\exists \bar{x} B'\} \cup \mathcal{L}$ for $\mathcal{L}$ a multiset of literals. Such a sequent can only be proved by a *Decide* rule by focusing on either a positive literal in $\mathcal{L}$ or the original formula $\exists \bar{x} B'$: in the latter case, the synchronous phase above it provides a substitution for all the variables in $\bar{x}$. One only needs to collect all of these substitutions into a list $\theta_1, \ldots, \theta_m$ and then show that the proof Ξ is essentially also a proof of $\vdash B'\theta_1 \vee^+ \cdots \vee^+ B'\theta_m \Uparrow \cdot$. ∎

Macro inference rules Focused proof systems such as LKF allow us to change the size of inference rules with which we work. Let us call individual introduction rules "micro-rules". An entire phase within a focused proof can be seen as a

"macro-rule". In particular, consider the following derivation.

$$\dfrac{\dfrac{\vdash \Theta, D \Uparrow N_1 \quad \cdots \quad \vdash \Theta, D \Uparrow N_n}{\vdash \Theta, D \Downarrow D}}{\vdash \Theta, D \Uparrow \cdot}$$

Here, the selection of the formula D for the focus can be taken as selecting among several macro-rules: this derivation illustrates one such macro-rule: the inference rule with conclusion $\vdash \Theta, D \Uparrow \cdot$ and with $n \geq 0$ premises $\vdash \Theta, D \Uparrow N_1, \ldots, \vdash \Theta, D \Uparrow N_n$ (where $N_1, \ldots, N_n$ are negative formulas). We shall say that this macro-rule is positive.

Example 2 Two extremes in proof sizes within LKF. Negative (small size but automatic and exponential) and postive (larger, interactive, smaller). ∎

Similarly, there is a corresponding negative macro-rule with conclusion, say, $\vdash \Theta, D \Uparrow N_i$, and with $m \geq 0$ premises of the form $\vdash \Theta, D, C \Uparrow \cdot$, where C is a multiset of positive formulas or negative literals.

In this way, focused proofs allow us to view the construction of proofs from conclusions of the form $\vdash \Theta \Uparrow \cdot$ as first attaching a positive macro rule (by focusing on some formula in Θ) and then attaching negative inference rules to the resulting premises until one is again to sequents of the form $\vdash \Theta' \Uparrow \cdot$. Such a combination of a positive macro rule below negative macro rules is often called a *bipole* [1].

Focusing can be broken at any point via *delays*. Within LKF, we can define the delaying operators

$$\partial^+(B) = B \wedge^+ t^+ \quad \text{and} \quad \partial^-(B) = B \wedge^- t^-.$$

Clearly, B, $\partial^-(B)$, and $\partial^+(B)$ are all logically equivalent but $\partial^-(B)$ is always negative and $\partial^+(B)$ is always positive. If one wishes to break a positive macro rule resulting from focusing on a given positive formula into smaller pieces, then one can insert $\partial^-(\cdot)$ into that formula. Similarly, inserting $\partial^+(\cdot)$ can limit the size of a negative macro rule. By inserting many delay operators, a focused proof can be made to emulate an unfocused proof.

3 Equality as a logical connective

Consider the two set of introduction rules for equality given in Figure 3: the first set presents (unfocused) left and right introduction rules for = while the second set of rules is formulated as one-sided and focused.

$$\frac{\Gamma\sigma \vdash \Delta\sigma}{\Gamma, s = t \vdash \Delta}\ \dagger \qquad \frac{}{\Gamma, s = t \vdash \Delta}\ \ddagger \qquad \frac{}{\Gamma \vdash \Delta, t = t}$$

$$\frac{\vdash \Theta\sigma \Uparrow \Gamma\sigma}{\vdash \Theta \Uparrow \Gamma, s \neq t}\ \dagger \qquad \frac{}{\vdash \Theta \Uparrow \Gamma, s \neq t}\ \ddagger \qquad \frac{}{\vdash \Theta \Downarrow t = t}$$

Figure 3: Introduction rules for =: the first set of inference rules uses two-sided, focused sequents and the second set of rules uses one-sided, unfocused sequents. Here, $\bar{t}$ is a list of n terms. The † proviso requires the terms s and t to be unifiable and σ to be their most general unifier. The ‡ proviso requires that the terms s and t are not unifiable. ∎

Examples of what is provable: equivalence and congruence. Show that $x = y$ is logically equivalent to Leibniz's equality relationship $[\forall P.\ Px \supset Py]$ (to prove the latter, the higher-order predicate will need to be instantiate with a predicate using equality.

Encoding finite sets: Encode $B = \{a_1, \ldots, a_n\}$ as the predicate expression

$$\hat{B} = [\lambda w.w = a_1 \vee \cdots \vee x = a_n].$$

Let B and C be two finite sets of objects of the same type. Show that in $B \subseteq C$ if and only if $\vdash_C \forall w.\hat{B}w \supset \hat{C}w$.

Mention the problems with unification: usually unification is used to *implement* proof systems: it has not been a rule in the proof system. An implementation of this logic containing equality must now do unification for existential variables (logic variables in the Prolog literature) and universal variables (eigenvariables in the proof theory literature).

4 Fixed points

In order for capture more interesting computational problems, we introduce the fixed point operator μ as a logical connective: in this way, we can define sets and relations recursively. Consider the left and right introduction rules for μ given in Figure 4. Notice that since the left and right introduction rules for μ are the same, μ is *self-dual*: that is, the De Morgan dual of μ is μ. It is possible to have a more expressive proof theory for fixed points that provides also for least and greatest fixed points (see, for example, [3, 2]): in that case, the De Morgan dual of the least fixed point is the greatest fixed point.

$$\frac{\Gamma, B(\mu B)\bar{t} \vdash \Delta}{\Gamma, \mu B\bar{t} \vdash \Delta} \qquad \frac{\Gamma \vdash \Delta, B(\mu B)\bar{t}}{\Gamma \vdash \Delta, \mu B\bar{t}}$$

$$\frac{\vdash \Theta \Uparrow \Gamma, B(\mu B)\bar{t}}{\vdash \Theta \Uparrow \Gamma, \mu B\bar{t}} \qquad \frac{\vdash \Theta \Downarrow B(\mu B)\bar{t}}{\vdash \Theta \Downarrow \mu B\bar{t}}$$

Figure 4: Two sets of introduction rules for μ. The first set is in a two-sided, unfocused sequent calculus and the second set is in a one-sided, focused sequent calculus. Here, B is a formula with $n \geq 0$ variables abstracted and $\bar{t}$ is a list of n terms. ∎

Example 3 Identify the natural numbers as terms involving 0 for zero and s for successor. The following simple logic program defines two predicates on natural numbers.

$$\begin{aligned} \mathit{nat}\ 0 &\subset \mathit{true}. \\ \mathit{nat}\ (s\ X) &\subset \mathit{nat}\ X. \\ \mathit{leq}\ 0\ Y &\subset \mathit{true}. \\ \mathit{leq}\ (s\ X)\ (s\ Y) &\subset \mathit{leq}\ X\ Y. \end{aligned}$$

The predicate *nat* can be written as the fixed point

$$\mu(\lambda p \lambda x.(x = 0) \vee \exists y.(s\ y) = x \wedge p\ y)$$

and binary predicate *leq* (less-than-or-equal) can be written as the fixed point

$$\mu(\lambda q \lambda x \lambda y.(x = 0) \vee \exists u \exists v.(s\ u) = x \wedge (s\ v) = y \wedge q\ u\ v).$$

In a similar fashion, any Horn clause specification can be made into fixed point specifications (mutual recursions requires standard encoding techniques). ∎

These two logical connectives can be added to LKF as follows. First, we classify both = and μ as positive connectives (this choice is forced for equality while μ can be polarized either way). The (one-sided) focused versions of the introduction rules above are given in Figures 3 and 4.

Example Consider proving the positive focused sequent

$$\vdash \Theta \Downarrow (\mathit{leq}\ m\ n \wedge^{+} N_1) \vee^{+} (\mathit{leq}\ n\ m \wedge^{+} N_2),$$

where m and n are natural numbers and *leq* is the fixed point expression displayed above but this time with all occurrences of $\wedge$ and $\vee$ polarized with their positive variants. If both N_1 and N_2 are negative formulas, then there are exactly two possible macro rules: one with premise $\vdash \Theta \Uparrow N_1$ when $m \leq n$ and one with premise $\vdash \Theta \Uparrow N_2$ when $n \leq m$ (thus, if $m = n$, either premise is possible). In this sense, a macro inference rule can contain an entire Prolog-style computation.

Example Macro rules can be built to match many computational situations. Consider, for example, defining simulation as the (greatest) fixed point of the equivalence

$$\mathit{sim}\ P\ Q \equiv \forall P' \forall A[P \xrightarrow{A} P' \supset \exists Q'[Q \xrightarrow{A} Q' \wedge \mathit{sim}\ P'\ Q']].$$

Although the right-hand-side of this definition looks complex, we show how it is possible to see proof search with this formula as being *exactly two* macro inference rules. First, the expression $P \xrightarrow{A} P'$ is, presumably, given via some SOS (structured operational semantic) specifications. Such specifications are simple, syntax-directed inference rules that can be captured as a least fixed point expression. As above, we will view such fixed point expressions as purely positive formulas. Thus, the expression $\forall P' \forall A[P \xrightarrow{A} P' \supset \cdot]$ is a negative macro rule: since all possible actions A and continuations P' must be computed, there are no choices to be made in building a proof for this expression. (Here, we are assuming that the implication $B \supset C$ is rendered as $\neg B \vee^{-} C$ in the polarized setting.) On the other hand, focusing on the expression $\exists Q'[Q \xrightarrow{A} Q' \wedge^{+} \cdot]$ yields a non-invertible, positive macro rule. In this way, the focused proof system is aligned directly with the structure of the actual (model-checking) problem. Notice that if one wishes to communicate a proof of a simulation to a proof checker, no information regarding the use of the negative macro rule needs to be communicated since the proof checker can also perform the computation behind that inference rule (*i.e.*, enumerating all possible transitions of a given process P).

Discuss model checking problems more. Bedwyr's architecture.

5 Induction and co-induction

Proposition. The following inference rules are derivable:

$$\frac{}{\vdash P, P^{\perp}}\ init \qquad \frac{\vdash \Gamma, B(\nu B)\vec{t}}{\vdash \Gamma, \nu B\vec{t}}\ \nu R$$

These results are standard, cf. [6]. The proof of the second one relies on monotonicity and is obtained by applying the ν rule with $B(\nu B)$ as the co-invariant.

Definition We classify as *asynchronous* (resp. *synchronous*) the connectives $\parr$, $\perp$, $\&$, $\top$, $\forall$, $\neq$, ν (resp. $\otimes$, $\mathbf{1}$, $\oplus$, $\mathbf{0}$, $\exists$, $=$, μ). A formula is said to be asynchronous (resp. synchronous) when its top-level connective is asynchronous (resp. synchronous). A formula is said to be *fully asynchronous* (resp. *fully synchronous*) when all of its connectives are asynchronous (resp. synchronous). Finally, a body $\lambda p \lambda \vec{x}.Bp\vec{x}$ is said to be fully asynchronous (resp. fully synchronous) when the formula $Bp\vec{x}$ is fully asynchronous (resp. fully synchronous).

Notice, for example, that $\lambda p \lambda \vec{x}.p\vec{x}$ is fully asynchronous and fully synchronous.

Proposition The following structural rules are admissible provided that B is fully asynchronous:

$$\frac{\vdash \Gamma, \nu B\vec{t}, \nu B\vec{t}}{\vdash \Gamma, \nu B\vec{t}}\ \nu C \qquad \frac{\vdash \Gamma}{\vdash \Gamma, \nu B\vec{t}}\ \nu W$$

Hence, the following structural rules hold for any fully asynchronous formula P:

$$\frac{\vdash \Gamma, P, P}{\vdash \Gamma, P}\ C \qquad \frac{\vdash \Gamma}{\vdash \Gamma, P}\ W$$

Proposition. The following structural rules are admissible provided that B is fully asynchronous:

$$\frac{\vdash \Gamma, \nu B\vec{t}, \nu B\vec{t}}{\vdash \Gamma, \nu B\vec{t}}\ \nu C \qquad \frac{\vdash \Gamma}{\vdash \Gamma, \nu B\vec{t}}\ \nu W$$

Fixed points (where S is closed, $\vec{x}$ is new)

$$\frac{\vdash \Gamma, B(\mu B)\vec{t}}{\vdash \Gamma, \mu B\vec{t}}\ \mu \qquad \frac{\vdash \Gamma, S\,\vec{t} \quad \vdash BS\,\vec{x}, (S\,\vec{x})^{\perp}}{\vdash \Gamma, \nu B\vec{t}}\ \nu \qquad \frac{}{\vdash \mu B\vec{t}, \nu \overline{B}\vec{t}}\ \mu\nu$$

Figure 5: Inference rules for least μ and greatest ν fixed points. ∎

Hence, the following structural rules hold for any fully asynchronous formula P:

$$\frac{\vdash \Gamma, P, P}{\vdash \Gamma, P}\ C \qquad \frac{\vdash \Gamma}{\vdash \Gamma, P}\ W$$

The rules for equality are not surprising. The main novelty here is the treatment of fixed points. Depending on the body, both μ and ν rules can be applied any number of times — but not with any co-invariant concerning ν. Notice for example that an instance of $\mu\nu$ can be η-expanded into a larger derivation, unfolding both fixed points to apply $\mu\nu$ on the recursive occurrences. As a result, each of the fixed point connectives has two rules in the focused system: one treats it as "an atom" and the other one as an expression with "internal structure."

Here, μ is treated during the synchronous phase and ν during the asynchronous phase. (Other choices are possible.) Roughly, what the focused system implies is that if a proof involving a ν-expression proceeds by co-induction on it, then this co-induction can be done at the beginning; otherwise that formula can be ignored in the whole derivation, except for the $\mu\nu$ rule. Focusing on a μ-expression yields two choices: unfolding or applying the initial rule for fixed points. If the body is fully synchronous, the focusing will never be lost. For example, if *nat* is the (fully synchronous) expression $\mu(\lambda nat.\lambda x.\ x = 0 \oplus \exists y.x = s\ y \otimes nat\ y)$, then focusing puts a lot of structure on a proof of $\Gamma \Downarrow nat\ t$: either t is a ground term representing a natural number and Γ is empty, or $t = s^n x$ for some $n \geq 0$ and Γ is $\{(nat\ x)^\perp\}$.

Theorem. The focused system is sound and complete with respect to $\mu\text{MALL}^=$.

Examples We shall now give a few theorems in $\mu\text{MALL}^=$. Although we do not give their derivations here, we stress that all of these examples are proved naturally in the focused proof system. The reader will also note that although $\mu\text{MALL}^=$ is linear, these derivations are intuitive and their structure resemble that of proofs in intuitionistic logic.

We first define a few least fixed points expressing basic properties of natural numbers. We assume two constants z and s of respective types n and $n \to n$. Note that all these definitions are fully synchronous.

$$\begin{aligned}
nat &\stackrel{def}{=} \mu(\lambda nat\lambda x.\ x = z \oplus \exists y.\ x = s\ y \otimes nat\ y)\\
even &\stackrel{def}{=} \mu(\lambda even\lambda x.\ x = z \oplus \exists y.\ x = s\ (s\ y) \otimes even\ y)\\
plus &\stackrel{def}{=} \mu(\lambda plus\lambda a\lambda b\lambda c.\ a = z \otimes b = c\\
&\qquad\qquad \oplus \exists a'\exists c'.a = s\ a' \otimes c = s\ c' \otimes plus\ a'\ b\ c')\\
leq &\stackrel{def}{=} \mu(\lambda leq\lambda x\lambda y.\ x = y \oplus \exists y'.\ y = s\ y' \otimes leq\ x\ y')
\end{aligned}$$

$$half \stackrel{def}{=} \mu(\lambda half \lambda x \lambda h.\ (x = z \oplus x = s\ z) \otimes h = z \\ \oplus\ \exists x' \exists h'.\ x = s\ (s\ x') \otimes h = s\ h' \otimes half\ x'\ h')$$

The following statements are theorems, all of which can be proved by induction. The main insights required for proving these theorems involve deciding which fixed point expression should be introduced by induction: the proper invariant is not the difficult choice here since the context itself is adequate in these cases.

$$\vdash \forall x.\ nat\ x \multimap even\ x \oplus even\ (s\ x)$$
$$\vdash \forall x.\ nat\ x \multimap \forall y \exists z.\ plus\ x\ y\ z$$
$$\vdash \forall x.\ nat\ x \multimap plus\ x\ z\ x$$
$$\vdash \forall x.\ nat\ x \multimap \forall y.\ nat\ y \multimap \forall z.\ plus\ x\ y\ z \multimap nat\ z$$

In the last theorem, the assumption $(nat\ x)^{\perp}$ is not needed and can be weakened (see earlier Proposition). In order to prove $(\forall x.\ nat\ x \multimap \exists h.\ half\ x\ h)$ one has to use a complete induction, *i.e.*, use the strengthened invariant $(\lambda x.\ nat\ x \otimes \forall y.\ leq\ y\ x \multimap \exists h.\ half\ y\ h)$.

A typical example of co-induction involves the simulation relation. Assume that $step : state \to label \to state \to o$ is an inductively defined relation encoding a labeled transition system. Simulation can be defined using the definition

$$sim \stackrel{def}{=} \nu(\lambda sim \lambda p \lambda q.\ \forall a \forall p'.\ step\ p\ a\ p' \multimap \exists q'.\ step\ q\ a\ q' \otimes sim\ p'\ q').$$

Reflexivity of simulation $(\forall p.\ sim\ p\ p)$ is proved easily by co-induction with the co-invariant $(\lambda p \lambda q.\ p = q)$. Instances of *step* are not subject to induction but are treated "as atoms". Proving transitivity, that is,

$$\forall p \forall q \forall r.\ sim\ p\ q \multimap sim\ q\ r \multimap sim\ p\ r$$

is done by co-induction on $(sim\ p\ r)$ with the co-invariant $(\lambda p \lambda r.\ \exists q.\ sim\ p\ q \otimes sim\ q\ r)$. The focus is first put on $(sim\ p\ q)^{\perp}$, then on $(sim\ q\ r)^{\perp}$. The fixed points $(sim\ p'\ q')$ and $(sim\ q'\ r')$ appearing later in the proof are treated "as atoms", as are all negative instances of *step*.

Except for the totality of *half*, all these theorems seem simple to prove using a limited number of heuristics. For example, one could first try to treat fixed points "as atoms", an approach that would likely fail quickly if inappropriate. Second, depending on the "rigid" structure of the arguments to a fixed point expression, one might choose to either unfold the fixed point or attempt to use the surrounding context to generate an invariant.

References

[1] J.-M. Andreoli. Focussing and proof construction. *Annals of Pure and Applied Logic*, 107(1):131–163, 2001.

[2] D. Baelde. *A linear approach to the proof-theory of least and greatest fixed points*. PhD thesis, Ecole Polytechnique, Dec. 2008.

[3] D. Baelde and D. Miller. Least and greatest fixed points in linear logic. In N. Dershowitz and A. Voronkov, editors, *International Conference on Logic for Programming and Automated Reasoning (LPAR)*, volume 4790 of *LNCS*, pages 92–106, 2007.

[4] C. Liang and D. Miller. Focusing and polarization in linear, intuitionistic, and classical logics. *Theoretical Computer Science*, 410(46):4747–4768, 2009.

[5] D. Miller. Finding unity in computational logic. In *ACM-BCS-Visions Conference*, Apr. 2010.

[6] A. Tiu. *A Logical Framework for Reasoning about Logical Specifications*. PhD thesis, Pennsylvania State University, May 2004.