Notes on Bunched Focusing

Kevin Donnelly

4 / 15 / 2004

1 Bunched Focusing

We propose a focusing method for **BI** logic that generalizes the focusing method for intuitionistic and linear logic to a recursive tree-structured focusing. Normal intuitionistic and linear focusing can been seen as the depth 1 restriction of this kind of focusing. Where flat-context focusing had two hypothetical fields, representing the top level context and depth 1 focus respectively. Our bunch focusing will have a variable number of hypothetical fields, depending on the current depth of focus, with the maximum being the depth of the whole top-level bunch.

The proof rules for **BI** necessitate this because for the $-\!\!*L$ and $\supset L$ rules we need to know the contents of the context the formula being focused on is in. Rather than keeping track of where in the bunch some formula came from in an ad-hoc way, we will only allow one to focus one level at a time, by choosing one bunch in the current deepest focused (or top-level if nothing is focused) context. When we come to a formula we apply a chain of synchronous rules to it as usual, then fold it back into the next higher level when it is no longer synchronous.

An example sequent in this calculus looks like:

$$\Delta;\Delta' \mid_+ \Gamma,\Gamma',\psi \mid_* \Delta'';\Delta''' \mid_+ \ldots \mid_* \Delta^{2n'};\varphi \mid_+ \varphi \overset{s}{\Longrightarrow} \cdot \mid R$$

Where $\Delta = \Gamma,\Gamma'$ and $\Gamma' = \Delta'';\Delta'''$ etc. It should also be possible to formulate a calculus that has sequents of the form:

$$\Delta;\Delta'(\Gamma_1,\ldots,\Gamma_n,\varphi) \mid_{\Delta;\Delta'(\Gamma_1,\ldots,\Gamma_n,-)} \varphi \overset{s}{\Longrightarrow} \cdot \mid R$$

Where $\Gamma_1, ..., \Gamma_n, \varphi$ is a maximal multiplicative subbunch of $\Delta'(\Gamma_1, ..., \Gamma_n, \varphi)$. This formulation is very similar to normal focusing except that we have to keep track of what context the focus formula came from. The main advantage to the first formulation is that it gives an inherent restriction on the order that the formulas can be focused on so that it is clear that we never have to interleave focusing on formulas from far away parts of the tree (if it turns out to be complete, which I suspect it is, given the local nature of the rules).

We group the **BI** propositions as left or right synchronous if their left or right rules are uninvertible, and left or right asynchronous otherwise.

$$\begin{array}{rrcl} \text{Left synchronous propositions} & L & ::= & P \mid A \supset B \mid A \mathrel{-\!\!*} B \\ \text{Right synchronous propositions} & R & ::= & P \mid A \vee B \mid \bot \mid A * B \end{array}$$

$$\begin{array}{rrcl} \text{Passive additive antecedents} & \Sigma & ::= & \cdot \mid \Sigma; L \mid \Sigma; \Sigma' \mid (\Pi) \\ \text{Passive multiplicative antecedents} & \Pi & ::= & \cdot \mid \Pi, L \mid \Pi, \Pi' \mid (\Sigma) \end{array}$$

We also ought to be able to treat $A \wedge B$ and $\top$ as left or right synchronous and I as right synchronous whenever it is convenient to do so. A bunch is synchronous if and only if all of its its leaf formulas are left synchronous.

Sequents now consist of four judgements. In our sequents there are still two active zones which are the formulas around the arrow, and two passive zones, one of which grows with the depth of focus. Sequents are written as one of:

$$\begin{array}{ll} \Sigma; \Sigma' \mid_+ \Pi, \Pi' \mid_* ... \mid_+ \Omega \Longrightarrow A \mid \cdot & \Sigma; \Sigma' \mid_+ \Pi, \Pi' \mid_* ... \mid_* \Omega \Longrightarrow A \mid \cdot \\ \Sigma; \Sigma' \mid_+ \Pi, \Pi' \mid_* ... \mid_+ \Omega \Longrightarrow \cdot \mid R & \Sigma; \Sigma' \mid_+ \Pi, \Pi' \mid_* ... \mid_* \Omega \Longrightarrow \cdot \mid R \end{array}$$

The most significant rules for understanding this focusing method are LL and focus$L$ which are:

$$\frac{\ldots\ (\Sigma; L) \mid_+ \Omega \overset{a}{\Longrightarrow} \cdot \mid R \quad (L \text{ synchronous bunch})}{\ldots\ \Sigma \mid_+ \Omega; L \overset{a}{\Longrightarrow} \cdot \mid R}\ LL+$$

$$\frac{\ldots\ (\Pi, L) \mid_* \Omega \overset{a}{\Longrightarrow} \cdot \mid R \quad (L \text{ synchronous bunch})}{\ldots\ \Pi \mid_* \Omega, L \overset{a}{\Longrightarrow} \cdot \mid R}\ LL*$$

For $LL+$ and $LL*$, the L can be either a left synchronous singleton formula or a more complex bunch whose leaves are all left synchronous.

$$\frac{\ldots\ (\Sigma; L)\ |_+\ L \overset{s}{\Longrightarrow} \cdot \mid R}{\ldots\ (\Sigma; L)\ |_+\ \cdot \overset{a}{\Longrightarrow} \cdot \mid R}\ \text{focus}L+ \qquad \frac{\ldots\ (\Pi, L)\ |_*\ L \overset{s}{\Longrightarrow} \cdot \mid R}{\ldots\ (\Pi, L)\ |_*\ \cdot \overset{a}{\Longrightarrow} \cdot \mid R}\ \text{focus}L*$$

Blur$L$ is essentially unchanged.

$$\frac{\ldots\ \Sigma\ |_+\ A \overset{a}{\Longrightarrow} \cdot \mid R \quad (A \text{ not synchronous})}{\ldots\ \Sigma\ |_+\ A \overset{s}{\Longrightarrow} \cdot \mid R}\ \text{blur}L+$$

$$\frac{\ldots \Pi\ |_*\ A \overset{a}{\Longrightarrow} \cdot \mid R \quad (A \text{ not synchronous})}{\ldots \Pi\ |_*\ A \overset{s}{\Longrightarrow} \cdot \mid R}\ \text{blur}L*$$

The corresponding right rules RR, focus$R$ and blur$R$ are basically the same as for intuitionistic or linear focusing. And most of the other rules are obvious from this basis. We basically need to use the intuitionistic focusing rules when the final focusing step is additive (i.e. $\ldots \Sigma\ |_+\ \varphi \Longrightarrow \cdot \mid R$) and multiplicative linear focusing rules when the final focusing step is multiplicative.

Here are a few of the harder rules for example:

$$\frac{\ldots\ \Pi\ |_*\ \psi \overset{s}{\Longrightarrow} \cdot \mid R \quad \cdot\ |_\#\ \Gamma \overset{a}{\Longrightarrow} \varphi \mid \cdot}{\ldots\ (\Pi, \Gamma)\ |_*\ \varphi \mathrel{-\!\!*} \psi \overset{s}{\Longrightarrow} \cdot \mid R}\ \mathrel{-\!\!*} L*$$

$$\frac{\ldots\ \Sigma\ |_+\ \psi \overset{s}{\Longrightarrow} \cdot \mid R \quad \cdot\ |_\#\ \Gamma \overset{a}{\Longrightarrow} \varphi \mid \cdot}{\ldots\ (\Sigma; \Gamma)\ |_+\ \varphi \supset \psi \overset{s}{\Longrightarrow} \cdot \mid R}\ \supset L+$$

$$\frac{\emptyset_a; \Gamma|_+ \cdot \overset{s}{\Longrightarrow} A \mid \cdot \quad \Delta|_\# \cdot \overset{s}{\Longrightarrow} B \mid \cdot}{\Gamma, \Delta|_* \cdot \overset{s}{\Longrightarrow} A * B \mid \cdot}\ *R*$$

For obvious reasons (look at the sequent calculus rules) there are no $\mathrel{-\!\!*} L+$ or $\supset L*$ rules. Also, I have used $|_\#$ in some places to mean that there are really two rules, one for when the top-level connective of the relevent bunch is additive, in which case $|_\# = |_+$, and one for when it is multiplicative, so $|_\# = |_*$.

In this system the weakening rule allows one to drop sequences of $|_+$'s in the backward direction (but, we also have to change the rest of the focus fields to drop the appropriate subbunches):

$$\frac{[\ldots\ |_\#]'\ \Gamma' \ldots\ \overset{s}{\Longrightarrow} \cdot \mid R}{[\ldots\ |_\#]\ \Gamma; \Gamma'\ |_+\ \Gamma' \ldots\ \overset{s}{\Longrightarrow} \cdot \mid R}\ W$$

Where $[\ldots\ |_\#]'$ is the same as $[\ldots\ |_\#]$ but without the new bunch (Γ).

This reflects the fact that we can add subbunches connected by ; while building a proof in the forward direction. It is important to note that in rule W, the $\Gamma;\Gamma'$ can be the leftmost part of the sequent, so we can drop initial additive focus ($|_+$) sequences. Contraction looks basically the same as in the sequent system.

$$\frac{[\ldots\ |_{\#}]'\,\Gamma;\Gamma\ |_+\ \Gamma\ |_{\#}\ \ldots\ \stackrel{a}{\Longrightarrow}\cdot\mid R}{[\ldots\ |_{\#}]\,\Gamma\ |_{\#}\ \ldots\ \stackrel{a}{\Longrightarrow}\cdot\mid R}\ C$$

Where $[\ldots\ |_{\#}]'$ is the same as $[\ldots\ |_{\#}]$ but with an extra occurrence of Γ. Again, the initial "$\ldots\ |_{\#}$" can be empty. Clearly we want a weakening-free and contraction free formulation for purposes of implementation so we need the rules to take these into account. This could be messy in this notation. One possibility would be to allow the rules ignore sequences of additive focus ($|_+$) steps in the sequents when comparing to the rule premises. Contraction will come for free from our representation. The one tricky case is when we need to contract one branch to make it equal to another. I think it may be enough to check that whenever a branch is added to an additive context (or sequents are combined additively in a rule) it is not a subbunch or superbunch of any other in the same context. The subbunch ($\leq_{\mathbf{BI}}$) relation is defined recursively as follows:

$$\frac{}{A \leq_{\mathbf{BI}} A}$$

$$\frac{A \leq_{\mathbf{BI}} A'}{A \leq_{\mathbf{BI}} A';B} \qquad \frac{A \leq_{\mathbf{BI}} A' \quad B \leq_{\mathbf{BI}} B'}{A,B \leq_{\mathbf{BI}} A',B'}$$

Which means, in our n-ary mutually recusive tree representation, one bunch is a subbunch of another if it is the same in multiplicative nodes and a subset in additive nodes. I.e. (a,b) is a subbunch of ((a,(b;c));d) but not of (a,b,c).

However, weakening might not be a bad thing. Sungwoo's method seems to essentially use weakening, which is what justifies the "bubbling up" of rules. We may be able to show some correspondece between the two.