

SASyLF: An Educational Proof Assistant for Language Theory

Jonathan Aldrich Robert J. Simmons Key Shin

School of Computer Science
Carnegie Mellon University

Microsoft
Corporation

Workshop on Mechanizing Metatheory
September 20, 2008





Teaching PL Theory and Proofs is Hard

- Challenges
 - Too easy to get details wrong – never learn concepts
 - Slow feedback loop – wait a week for a corrected homework
 - High time and effort for the payoff
- Tools can potentially help
 - Check the details are correct
 - Provide immediate feedback
 - Raise productivity
- Just like good compilers make programming more accessible



State of the Art

- Educational Tools – great for their domain
 - General Mathematics: EPGY theorem proving environment
 - geometry, linear algebra, number theory, etc.
 - Logic: Tarski's world, Tutch
- Proof assistants and checkers – great for experts
 - Isabelle, HOL, PVS, Coq, ...
 - Benefit: very expressive, good automation, ...
 - Drawback: steep learning curve
 - both for tool and for variable encoding techniques
 - Twelf
 - Higher-order type theory aids in reasoning about programs
 - Avoids variable encodings – potentially more accessible
 - Explicit, machine-checkable proof
 - Small steps students can follow
 - **Still a steep learning curve: based on higher-order type theory**
 - Challenging to use for undergraduate or intro graduate courses



SASyLF Design Goals

- **Easy to learn**
 - Familiar syntax – just like we write on paper
 - Little mathematical context – just inference rules
- **Matches PL teaching approaches**
 - Support for variable binding and α -equivalence
- **Explicit notation**
 - Every step specified completely
- **Rapid feedback**
 - Automated checker
 - Support for partial proofs
- **Good error messages**
 - Local checking and local errors



Outline

- Motivation and Design Goals
- **SASyLF intro: addition commutes**
- Semantics: the λ -calculus and LF
- POPLmark 2A in SASyLF
- Preliminary teaching experience
- Conclusion



Sum Commutes in SASyLF (Definitions)

package edu.cmu.cs.sum;

terminals z s

Definitions

syntax

Natural Numbers $n ::= z \mid s(n)$

$n ::= z$
 $\mid s\ n$

Inference Rules

$$\frac{}{z + n = n} \text{sum-z}$$

$$\frac{n_1 + n_2 = n_3}{s(n_1) + n_2 = s(n_3)} \text{sum-s}$$

judgment sum: $n_1 + n_2 = n_3$

----- sum-z
 $z + n = n$

$n_1 + n_2 = n_3$
----- sum-s
 $s\ n_1 + n_2 = s\ n_3$



Sum Commutes in SASyLF (Lemma)

Theorem [sum-z-right]: $\forall n \quad n + z = n$

theorem sum-z-right: forall n

exists $n + z = n$.

Proof. By induction on n :

Case z : $d1: n + z = n$ **by induction on n :**

Case z is

case z is

$z + z = z$ *by rule sum-z*

$d2: z + z = z$

by rule sum-z

end case

Case $s(n)$:

case $s \ n'$ is

$n + z = n$

by induction hypothesis

$d3: n' + z = n'$

by induction hypothesis on n'

$s(n) + z = s(n)$

by rule sum-s

$d4: s \ n' + z = s \ n'$ **by rule sum-s on $d3$**

end case

q.e.d.

**end induction
end theorem**

Demonstration: Sum Commutes



See `sum.slf` in the `SASyLF/examples` directory



Outline

- Motivation and Design Goals
- SASyLF intro: addition commutes
- **Semantics: the λ -calculus and LF**
- POPLmark 2A in SASyLF
- Preliminary teaching experience
- Conclusion

Semantics of SASyLF



- Based on LF, like Twelf!
 - The interesting bit is how SASyLF syntax encodes LF
- Twelf
 - Judgments as types
 - Proofs as logic programs
- SASyLF
 - Judgments as types
 - Proofs as let-normal functional programs
 - Each intermediate step is named and typed, aiding error messages
- Higher-order abstract syntax
 - SASy - Second-order Abstract Syntax
 - Neither syntax nor judgments can include other judgments as parts
 - Same expressiveness as standard paper notation

From SASyLF to LF/Twelf



terminals fn unit

e : type. One LF type per
tau : type. SASyLF nonterminal

e ::= x
| "(" ")"
| e e
| fn x : tau => e[x]

tau ::= unit
| tau -> tau

From SASyLF to LF/Twelf



terminals fn unit

e : type.
tau : type.

e ::= x

| "(" ")"

| e e

| fn x : tau => e[x]

tau ::= unit

| tau -> tau

RHS of SASyLF grammar
becomes an LF constructor.
Result is LHS nonterminal

From SASyLF to LF/Twelf



terminals fn unit

e : type.
 τ : type.

$e ::= x$
| "(" "("
| $e e$
| fn x : $\tau \Rightarrow e[x]$

Nonterminals in RHS
become arguments of
the nonterminal's type

unit : e .
app : $e \rightarrow e \rightarrow e$.

$\tau ::= \text{unit}$
| $\tau \rightarrow \tau$

From SASyLF to LF/Twelf



terminals fn unit

$e ::= x$ \leftarrow x is a case of e ,
so variable x has type e .
No constructor created.

| "(" "("
| e e \leftarrow $e[x]$ means
 x is a variable
that is bound in e

| fn $x : \text{tau} \Rightarrow e[x]$ \leftarrow Variables appearing in
clause ignored when
producing constructors

$\text{tau} ::= \text{unit}$
| $\text{tau} \rightarrow \text{tau}$

$e : \text{type}$.

$\text{tau} : \text{type}$.

unit : e .

app: $e \rightarrow e \rightarrow e$.

abs: $\text{tau} \rightarrow (e \rightarrow e) \rightarrow e$.

\leftarrow second argument of LF constructor
is a function from type of x to e



From SASyLF to LF/Twelf

terminals fn unit

$e ::= x$
| "(" "("
| e
| fn $x : \tau \Rightarrow e[x]$

$e : \text{type.}$
 $\tau : \text{type.}$

unit : $e.$
app: $e \rightarrow e \rightarrow e.$
abs: $\tau \rightarrow (e \rightarrow e) \rightarrow e.$

$\tau ::= \text{unit}$
| $\tau \rightarrow \tau$

tunit : $\tau.$
arrow : $\tau \rightarrow \tau \rightarrow \tau.$

(fn $x2 : \text{unit} \Rightarrow e1[x2]$)

($e1$ $x2$)

$e1[x2]$ becomes ($e1$ $x2$), where $e1$ has type $e \rightarrow e$

GLR parser produces a parse tree



From SASyLF to LF/Twelf

terminals fn unit

$e ::= x$
| "(" "("
| $e e$
| $\text{fn } x : \text{tau} \Rightarrow e[x]$

$\text{tau} ::= \text{unit}$
| $\text{tau} \rightarrow \text{tau}$

$e : \text{type.}$

$\text{tau} : \text{type.}$

$\text{unit} : e.$

$\text{app} : e \rightarrow e \rightarrow e.$

$\text{abs} : \text{tau} \rightarrow (e \rightarrow e) \rightarrow e.$

$\text{tunit} : \text{tau.}$

$\text{arrow} : \text{tau} \rightarrow \text{tau} \rightarrow \text{tau.}$

The variable declared here is bound here
so we identify the variable here and bind it over here

$(\text{fn } x2 : \text{unit} \Rightarrow e1[x2])$

$[x2] (e1 x2)$

GLR parser produces a parse tree



From SASyLF to LF/Twelf

terminals fn unit

$e ::= x$
| "(" "("
| $e e$
| fn $x : \tau \Rightarrow e[x]$

$e : \text{type.}$
 $\tau : \text{type.}$

unit : $e.$
app: $e \rightarrow e \rightarrow e.$
abs: $\tau \rightarrow (e \rightarrow e) \rightarrow e.$

$\tau ::= \text{unit}$
| $\tau \rightarrow \tau$

tunit : $\tau.$
arrow : $\tau \rightarrow \tau \rightarrow \tau.$

The name x_2 is an extension of x (i.e. a number or ' is added)
and x has type e , so x_2 has type e

(fn $x_2 : \text{unit} \Rightarrow e_1[x_2]$)

$[x_2:e] (e_1 x_2)$

GLR parser produces a parse tree

From SASyLF to LF/Twelf



terminals fn unit

$e ::= x$
| "(" "("
| e
| fn $x : \tau \Rightarrow e[x]$

$e : \text{type.}$
 $\tau : \text{type.}$

$\text{unit} : e.$
 $\text{app} : e \rightarrow e \rightarrow e.$
 $\text{abs} : \tau \rightarrow (e \rightarrow e) \rightarrow e.$

$\tau ::= \text{unit}$
| $\tau \rightarrow \tau$

$\text{tunit} : \tau.$
 $\text{arrow} : \tau \rightarrow \tau \rightarrow \tau.$

$(\text{fn } x2 : \text{unit} \Rightarrow e1[x2])$

$\text{abs tunit } ([x2:e] (e1 x2))$

GLR parser produces a parse tree



From SASyLF to LF/Twelf

Gamma ::= *
| Gamma, x:tau

Judgments as types:
A judgment form in SASyLF
turns into an LF type family

judgment has-type: Gamma |- e : tau \longrightarrow has-type: e \rightarrow tau \rightarrow type
assumes Gamma \longleftarrow

Gamma is not used in the type family;
the assumes clause tells SASyLF it is a
stand-in variable for the LF context



From SASyLF to LF/Twelf

Gamma ::= *
| Gamma, x:tau

judgment has-type: Gamma |- e : tau has-type: e -> tau -> type
assumes Gamma

----- t-var Variable rules have no LF equivalent; they
correspond to using assumptions from the LF
context.

Gamma, x:tau |- x : tau

From SASyLF to LF/Twelf



$\Gamma ::= *$
| $\Gamma, x:\tau$

judgment has-type: $\Gamma \vdash e : \tau$ has-type: $e \rightarrow \tau \rightarrow \text{type}$
assumes Γ

----- t-var
 $\Gamma, x:\tau \vdash x : \tau$

$\Gamma, x_1:\tau \vdash e_1[x_1] : \tau$ $\xrightarrow{\text{t-fn}}$ has-type (e1 x1) tau'
----- t-fn \rightarrow

$\Gamma \vdash \text{fn } x_1 : \tau \Rightarrow e_1[x_1] : \tau \rightarrow \tau$ $\xrightarrow{\text{t-fn}}$ has-type (abs tau ([x1:e] (e1 x1)))
(arrow tau tau')



From SASyLF to LF/Twelf

$\Gamma ::= *$

| $\Gamma, x:\tau$

judgment has-type: $\Gamma \mid - e : \tau$ has-type: $e \rightarrow \tau \rightarrow \text{type}$
 assumes Γ

----- t-var

$\Gamma, x:\tau \mid - x : \tau$

$\Gamma, x1:\tau \mid - e1[x1] : \tau$ '

----- t-fn

$\Gamma \mid - \text{fn } x1 : \tau \Rightarrow e1[x1] : \tau \rightarrow \tau$ '

$\text{t-fn} : \{ x1 : e \}$

has-type $(e1 \ x1) \ \tau$ '

\rightarrow

has-type $(\text{abs } \tau \ ([x1:e] \ (e1 \ x1)))$
 (arrow $\tau \ \tau$ ')



From SASyLF to LF/Twelf

Gamma ::= *

| Gamma, x:tau

judgment has-type: Gamma |- e : tau has-type: e -> tau -> type
assumes Gamma

----- t-var

Gamma, x:tau |- x : tau

t-fn : { x1 : e } { dx1 : has-type x1 tau }
 has-type (e1 x1) tau'

Gamma, x1:tau |- e1[x1] : tau'

----- t-fn

Gamma |- fn x1 : tau => e1[x1] : tau -> tau'

has-type (abs tau ([x1:e] (e1 x1)))
(arrow tau tau')



Case Analysis on Variables

lemma narrow-subtype: forall dsub: Gamma, X <: T |- T1[X] <: T2[X]
forall dsub': Gamma |- T' <: T
exists Gamma, X <: T' |- T1[X] <: T2[X].

dres: Gamma, X <: T' |- T1[X] <: T2[X] **by induction** on dsub:

case rule

----- SA-Var Case: T1[X] is X

d1: Gamma, X <: T |- X <: T

is

d2: Gamma, X <: T' |- X <: T' **by rule** SA-Var

d3: Gamma, X <: T' |- T' <: T **by weakening** on dsub'

d4: Gamma, X <: T' |- X <: T **by rule** SA-Trans on d2, d3

end case



Case Analysis on Variables

lemma narrow-subtype: forall dsub: Gamma, X <: T |- T1[X] <: T2[X]
forall dsub': Gamma |- T' <: T
exists Gamma, X <: T' |- T1[X] <: T2[X].

dres: Gamma, X <: T' |- T1[X] <: T2[X] **by induction** on dsub:

case rule ----- SA-Var Case: T1[X] is X
d1: Gamma, X <: T |- X <: T
is ... end case

case rule ----- SA-Var Case: T1[X] is an arbitrary
other variable X' from Gamma
d1: Gamma2, X' <: T", X <: T |- X' <: T"
is d2: Gamma2, X' <: T" |- X' <: T" **by rule** SA-Var
d3: Gamma2, X' <: T", X <: T' |- X' <: T" **by weakening** on d2
end case



Some Checks in SASyLF

- Case analysis
 - SASyLF tries unifying each inference rule's conclusion with the judgment being analyzed
 - Can't assume anything stronger than case computed by SASyLF (Twelf: input coverage)
 - Can't miss a case (Twelf: input coverage)
- Rule/Lemma application
 - SASyLF unifies rule with sources and conclusion specified by user
 - Can't assume "extra" structure in or equalities between input and output variables (Twelf: output freeness & output coverage)
- Induction Hypothesis
 - Same as rule application, but also the argument we're inducting over must be a subderivation (Twelf: termination)
- Substitution
 - The judgment to be substituted fulfills the other judgment's hypothesis (Twelf: dependent typechecking)
- Exchange
 - The first variable isn't free in the second hypothesis



Outline

- Motivation and Design Goals
- SASyLF intro: addition commutes
- Semantics: the λ -calculus and LF
- **POPLmark 2A in SASyLF**
- Preliminary teaching experience
- Conclusion

POPLmark 2A in SASyLF



syntax

$t ::= x$
 $| \text{lambda } x:T \Rightarrow t[x]$
 $| t t$
 $| \text{lambda } X <: T \Rightarrow t[X]$
 $| t \text{ "[\" } T \text{ "]\"}$

$T ::= X$
 $| \text{Top}$
 $| T \rightarrow T$
 $| \text{all } X <: T \Rightarrow T[X]$

$\text{Gamma} ::= *$

$| \text{Gamma}, x : T$
 $| \text{Gamma}, X <: T$

judgment value: t value

----- V-Abs

$\text{lambda } x:T \Rightarrow t[x]$ value

----- V-Tabs

$\text{lambda } X <: T \Rightarrow t[X]$ value

Limitation: value
 is a judgment, not
 a subset of syntax
 (like Twelf
 solution)

judgment reduce: $t \rightarrow t$

$t1 \rightarrow t1'$

----- E-CtxApp1

$t1 t2 \rightarrow t1' t2$

$t2$ value

----- E-AppAbs

$(\text{lambda } x:T1 \Rightarrow t12[x]) t2 \rightarrow t12[t2]$

Limitation:
 reduction is done
 with context rules,
 not explicit
 evaluation
 contexts (like
 Twelf solution)



POPLmark 2A in SASyLF

judgment has-type: $\Gamma \vdash t : T$
assumes Γ

----- T-Var

$\Gamma, x:T \vdash x : T$

$\Gamma, X <: T1 \vdash t2[X] : T2[X]$

----- T-Tabs

$\Gamma \vdash \lambda X <: T1 \Rightarrow t2[X] : \text{all } X <: T1 \Rightarrow T2[X]$

$\Gamma \vdash t1 : \text{all } X <: T11 \Rightarrow T12[X]$

$\Gamma \vdash T2 <: T11$

----- T-Tapp

$\Gamma \vdash t1 \text{ "["} T2 \text{ "]" } : T12[T2]$

$\Gamma \vdash t : T'$

$\Gamma \vdash T' <: T$

----- T-Sub

$\Gamma \vdash t : T$

Right out of the
POPLmark
specification

POPLmark 2A in SASyLF



judgment subtyping: $\Gamma \vdash T <: T'$
assumes Γ

----- SA-Top

$\Gamma \vdash T <: \text{Top}$

----- SA-Var

$\Gamma, X <: T \vdash X <: T$

$\Gamma \vdash T1 <: T1'$

$\Gamma \vdash T2' <: T2$

----- SA-Arrow

$\Gamma \vdash T1' \rightarrow T2' <: T1 \rightarrow T2$

$\Gamma \vdash T1 <: T1'$

$\Gamma, X <: T1 \vdash T2[X] <: T2'[X]$

----- SA-All

$\Gamma \vdash \text{all } X' <: T1' \Rightarrow T2'[X'] <: \text{all } X <: T1 \Rightarrow T2[X]$

As in the Twelf 2A solution, we use declarative rather than algorithmic subtyping. The algorithmic rules would have a more complicated encoding since the variable use rule is nonstandard.

----- SA-Ref1

$\Gamma \vdash T <: T$

$\Gamma \vdash T1 <: T2$

$\Gamma \vdash T2 <: T3$

----- SA-Trans

$\Gamma \vdash T1 <: T3$

POPLmark 2A in SASyLF



judgment steporvalue: t steporvalue Helper judgments

t value
----- steporvalue-value
t steporvalue

t -> t'
----- steporvalue-steps
t steporvalue

- steporvalue compensates for not having logical “or” built-in (same as Twelf)
- equality isn’t built in either, so we need a judgment
- other similar judgments required

judgment equality: t == t

----- equality
t == t

POPLmark 2A in SASyLF



lemma arrow-sub-arrow:
(similar for Tarrow)

forall dsub: * $\vdash T <: T1 \rightarrow T2$
exists $T == T1' \rightarrow T2'' <: T1 \rightarrow T2$.

lemma canonical-form-lambda :
(similar for type
lambda)

forall dtv : t value
forall dtt : * $\vdash t : T1 \rightarrow T2$
exists t == lambda x : $T1' \Rightarrow t'[x]$.

lemma narrow-subtype:

forall dsub: Gamma, X <: T $\vdash T1[X] <: T2[X]$
forall dsub': Gamma $\vdash T' <: T$
exists Gamma, X <: T' $\vdash T1[X] <: T2[X]$.

theorem progress :

* $\vdash t : T$
exists t steporvalue.

POPLmark 2A in SASyLF



lemma invert-lambda: **forall** dt: * |- lambda x:T11 => t12[x] : T1 -> T
forall dt2: * |- t2 : T2
forall dsub: * |- T2 <: T1
exists * |- t12[t2] : T.

dt12: * |- t12[t2] : T **by induction on dt:**

case rule

d1: * , x:T1 |- t12[x] : T

----- T-Abs

d2: * |- lambda x:T1 => t12[x] : T1 -> T

is

dt2' : * |- t2 : T1 **by rule** T-Sub **on** dt2, dsub

dt12: * |- t12[t2] : T **by substitution on** dt2', d1

end case

POPLmark 2A in SASyLF



```
lemma invert-lambda:      forall dt:
                           forall dt2:
                             forall dsub:
                               exists
dt12: * |- t12[t2] : T by induction on dt:
case rule
  d1: * |- lambda x:T11 => t12[x] : T"
  d2: * |- T" <: T1 -> T
----- T-Sub
  d3: * |- lambda x:T11 => t12[x] : T1 -> T
is
  d4: T" == T1' -> T' <: T1 -> T by lemma arrow-sub-arrow on d2
dt12.* |- t12[t2] : T
case rule
  d5: * |- T1 <: T1'
  d6: * |- T' <: T
----- arrow-sub
  d7: T1' -> T' == T1' -> T' <: T1 -> T
is
  d8: * |- T2 <: T1'
  d9: * |- t12[t2] : T'
dt12.* |- t12[t2] : T
end case
end case analysis
end case
```

* |- lambda x:T11 => t12[x] : T1 -> T
* |- t2 : T2
* |- T2 <: T1
* |- t12[t2] : T.

by rule SA-Trans on dsub, d5
by induction hypothesis on d1, dt2, d8
by rule T-Sub on d9, d6

POPLmark 2A in SASyLF



theorem preservation: **forall** dt: * |- t : T
 forall ds: t -> t'
 exists * |- t' : T.



POPLmark 2A in SASyLF

```
lemma narrow-subtype: forall dsub: Gamma, X <: T |- T1[X] <: T2[X]
  forall dsub': Gamma |- T' <: T
  exists
    Gamma, X <: T' |- T1[X] <: T2[X].
...
d2: Gamma, X <: T, X' <: T2'[X] |- T1"[X][X]" <: T2"[X][X]"
...
d8: Gamma, X <: T', X' <: T2'[X] |- T1"[X][X]" <: T2"[X][X]"
  by induction hypothesis on d2, dsub'
  // does not currently typecheck
  // because X <: T is not rightmost assumption in d2
  // conceptually an easy fix, but haven't done it yet
```

We use **unproved** instead (representing an unchecked claim).

POPLmark 2A in SASyLF



dt2: * |- t2 : T1
d1: * , x:T1 |- t12[x] : T
dt12: * |- t12[t2] : T **by substitution on dt2, d1 // not yet checked (2 others)**

dsub': Gamma |- T' <: T
d3: Gamma, X <: T' |- T' <: T **by weakening on dsub'**
 // not yet checked (1 other place)



POPLmark 2A Summary

- Readable to non-experts
 - Many definitions just an ascii-ization
- A few simplifications
 - All of which are also in the Twelf solution
- More verbose: 1105 lines long
 - compare: Twelf solution is 745 lines
- SASyLF implementation limitations
 - 5 unchecked substitution or weakening
 - 1 “unproved” due to a spurious warning
- Other challenges?
 - Challenge 1: need mutual induction
 - Challenge 2B: should be doable (with condition on distinct fields)
 - Challenge 3: requires proof search (known how to do in LF)



Outline

- Motivation and Design Goals
- SASyLF intro: addition commutes
- Semantics: the λ -calculus and LF
- POPLmark 2A in SASyLF
- **Preliminary teaching experience**
- **Conclusion**

Teaching experience: highlights



- SASyLF used in Spring 2008 graduate Analysis course at CMU
 - One assignment, covering small-step semantics
 - No in-class time was spent teaching the tool
- Results of controlled experiment:
 - 11 of 13 tool users: helped find errors in their proofs
 - 12 of 13 tool users: increased confidence that proofs were correct
 - 14 of 16 no-tool users: wished for earlier feedback on mistakes
- Several students dropped out of tool group
 - Usability issues; many have been fixed
 - 7 of 11 tool users would use the tool again (separate survey)
 - 10 of 11 if usability issues addressed
- Quote: “I actually did the entire assignment on paper first and then moved over to using the tool. I found the paper approach really easy. But once I started using the tool I started understanding the concepts better.”



SASyLF Limitations

- All the same limitations as Twelf
 - e.g. all theorems in $\forall\exists$ form
- Limitations of student-focused design
 - Second-order abstract syntax
 - no intrinsic encodings
 - More readable, but also more verbose
- Some features not yet implemented
 - some checks (substitution, weakening, exchange)
 - mutual induction
- Open questions
 - World subordination and subsumption
 - Advanced LF idioms supported in Twelf

Our goal is to support most or all of what students need to do in an intro course.
SASyLF is (or will be) good enough for some research applications, but don't expect the full power of tools designed for researchers.



SASyLF: The Gateway Drug for Mechanized Metatheory

- Educational proof assistant for PL theory
 - Familiar syntax, simple semantics
 - Support for variable binding
 - Capable of sophisticated proofs
 - Early feedback with local error messages
- Usable now in teaching
 - Can use without a big learning curve
 - Preliminary evidence of student benefit
 - Adopted in John Boyland's current type theory course at U Wisconsin-Milwaukee
- IDE and more features on the way

<http://www.sasylf.org/>