

Formal Logical Environments
for
Building Reliable, High-Performance Software

Christoph Kreitz

Department of Computer Science, Cornell University
Ithaca, NY 14853



WHY FORMAL LOGICAL ENVIRONMENTS?

- **Too many errors in informal arguments**
 - 40–50% of the published results turn out to be wrong
- **We can't afford errors in software development**
 - Errors are **annoying** (Reboot, loss of data, ...)
 - Errors are **expensive** (Pentium bug, Ariane 5 rocket, Mars Climate Orbiter)
 - Errors **cost lives** (Airbus crashes in the early 1990's)
 - ... affects air traffic, banking, government, utilities, schools, e-commerce, ...
- **Current software development methods are unreliable**
 - Tested programs still contain errors
 - Correctness proofs are **tedious** and **error-prone** (if done by hand)



Need formal tools for creating reliable software

FORMAL METHODS TOOLS ARE MOST SUCCESSFUL WHEN ENGAGED AT EARLY STAGES OF SYSTEM DESIGN

● Great potential

- clarifying critical design concepts
- linking abstract and concrete specifications
- detecting subtle errors in design and prototype code
- generating code of components from specifications
- improving system performance


● Requires

- expressive formal language
- knowledge base of formalized facts about systems concepts
- proof environment integrating different reasoning techniques
- collaboration with systems experts in real applications


The NUPRL LPE meets these requirements

- Computational formal logics
- Proof & program development systems
 - The **NUPRL** Logical Programming Environment
 - Fast inference engines + proof search techniques
 - Natural language generation from formal mathematics
 - Program extraction + automated complexity analysis

TYPE THEORY



- Application to reliable, high-performance networks
 - Assigning precise semantics to system software
 - Performance Optimizations
 - Assurance for reliability (verification)
 - Verified System Design



THE NUPRL TYPE THEORY

AN INSTANCE OF MARTIN-LÖF TYPE THEORY

- **Constructive higher-order logic**
 - Reasoning about types, members of types, propositions, functions ...
- **Functional programming language**
 - Similar to core **ML**: polymorphic, with partial recursive functions
- **Open-ended, expressive data type system**
 - Function, product, disjoint union, Π - & Σ -types, atoms ~> programming
 - Integers, lists, inductive types ~> inductive definition
 - Propositions as types, equality type, void, top, universes ~> logic
 - Subsets, subtyping, quotient types ~> mathematics
 - (Dependent) intersection, union, records ~> modules, program composition
 - New types can/will be added as needed*
- **Expressions separate from their types** ~> full λ -calculus
 - ... but must be typeable *in proofs* ~> "total" functions
- **User-defined extensions possible**

SYNTAX ISSUES

- **Uniform term syntax** for all expressions

$opid\{p_i : F_i\}(x_{11}, \dots, x_{m_1 1} \cdot t_1; \dots; x_{1n}, \dots, x_{m_n n} \cdot t_n)$

- *Operator identifier* listed in operator tables
- *Parameters* for base terms (variables, numbers, tokens...)
- *Sub-terms* may contain *bound variables*

No syntactical distinction between types, members, propositions ...

- **Separate term display form**

- Describe visual appearance of terms in “free syntax”
- ↳ Conventional notation, information hiding, auto-parenthesizing, aliases, ...

<i>Internal Term Structure</i>	<i>Display Form</i>
function { } (<i>S</i> ; <i>.T</i>)	<i>S</i> → <i>T</i>
variable { <i>x</i> : <i>v</i> } ()	<i>x</i>
lambda { } (<i>x</i> . <i>t</i>)	$\lambda x . t$
apply { } (<i>f</i> ; <i>t</i>)	<i>f t</i>
⋮	⋮

SEMANTICS MODELS PROOF, NOT DENOTATION

- **(Lazy) evaluation of expressions**

- Identify **canonical expressions** (values)
- Define reducible non-canonical expressions (**redex**)
- Define reduction steps in **redex–contracta table**

<i>canonical</i>	<i>non-canonical</i>	<i>Redex</i>	<i>Contractum</i>
$S \rightarrow T$			
$\lambda x . t$	$\boxed{f} t$	$\boxed{\lambda x . u} t$	$\xrightarrow{\beta} u[t/x]$

- **Judgments: semantical truths about expressions**

- 4 categories: Typehood (**T Type**), Type Equality (**S=T**),
Membership (**t ∈ T**), Member equality (**s=t ∈ T**)
- Defined for **values** of expressions in **semantics tables**

$$S_1 \rightarrow T_1 = S_2 \rightarrow T_2 \quad \text{iff } S_1 = S_2 \text{ and } T_1 = T_2$$

$$\lambda x_1 . t_1 = \lambda x_2 . t_2 \in S \rightarrow T \quad \text{iff } S \rightarrow T \text{ Type and } t_1[s_1/x_1] = t_2[s_2/x_2] \in T$$

for all s_1, s_2 with $s_1 = s_2 \in S$

⋮

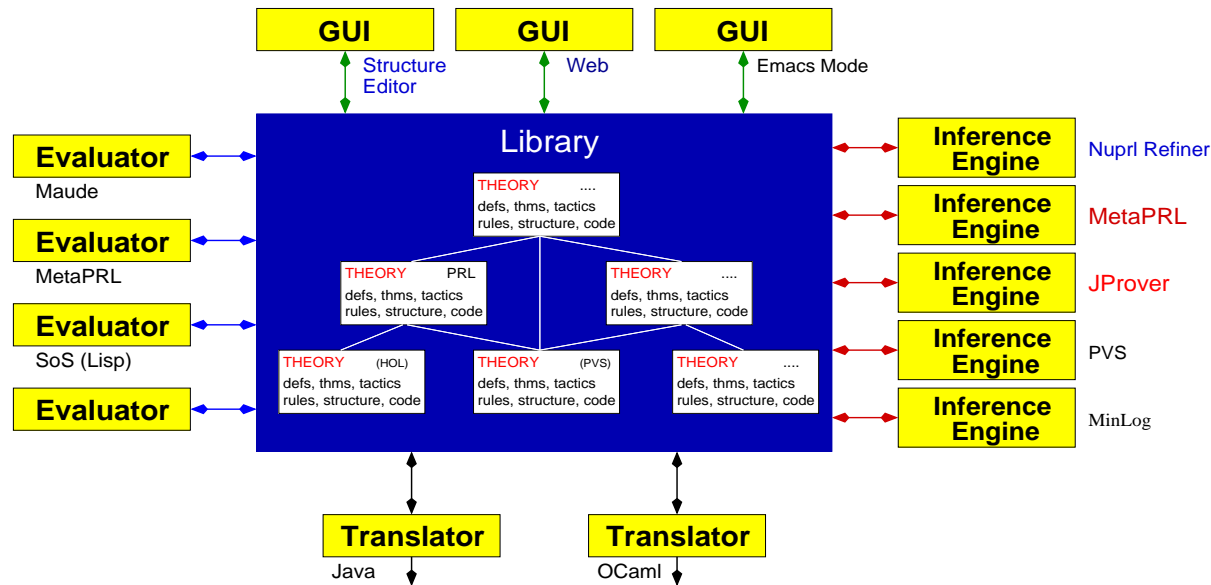
⋮

NUPRL'S PROOF THEORY

- **Sequent** $x_1:T_1, \dots, x_n:T_n \vdash C$ $[\text{ext } t]$
 - “If x_i are variables of type T_i then C has a (yet unknown) member t ”
 - A judgment $t \in T$ is represented as T $[\text{ext } t]$ \rightsquigarrow proof term construction
 - Equality is represented as type $s=t \in T$ $[\text{ext } Ax]$ \rightsquigarrow propositions as types
 - Typehood represented by (cumulative) universes U_i $[\text{ext } T]$
- **Refinement calculus**
 - Top-down decomposition of proof goal \rightsquigarrow interactive proof development
 - Bottom-up construction of proof terms \rightsquigarrow program extraction
$$\Gamma \vdash S \rightarrow T \quad [\text{ext } \lambda x. e] \quad \text{by lambda-formation } x$$
$$\Gamma, x:S \vdash T \quad [\text{ext } e]$$
$$\Gamma \vdash S=S \in U_i \quad [\text{ext } Ax]$$
 - Computation rules \rightsquigarrow program evaluation

About 8–10 inference rules for each data type in NUPRL

NUPRL'S AUTOMATED REASONING ENVIRONMENT




- **Interactive proof development**

- Supports program **extraction** and **evaluation**
- Proof automation through **tactics** & **decision procedures**
- Highly customizable: conservative **language extensions**, **term display**, ...

- **Cooperating processes centered around knowledge base**

- **Library** of formal algorithmic knowledge
- Multiple user interfaces
- External proof engines
- **Asynchronous** & **collaborative** theorem proving



Secure software infrastructure

- **Ensemble Group Communication Toolkit**
 - System optimization and verification, formal component design
- **MediaNet Stream Computation Network** (ongoing)
 - Validation of real-time schedules wrt. resource limitations

THE ENSEMBLE GROUP COMMUNICATION TOOLKIT

Modular group communication system


- Developed by Cornell's [System Group](#) (Ken Birman)
- Used commercially ([BBN](#), [JPL](#), [Segasoft](#), [Alier](#), [Nortel Networks](#))

Architecture: stack of **micro-protocols**

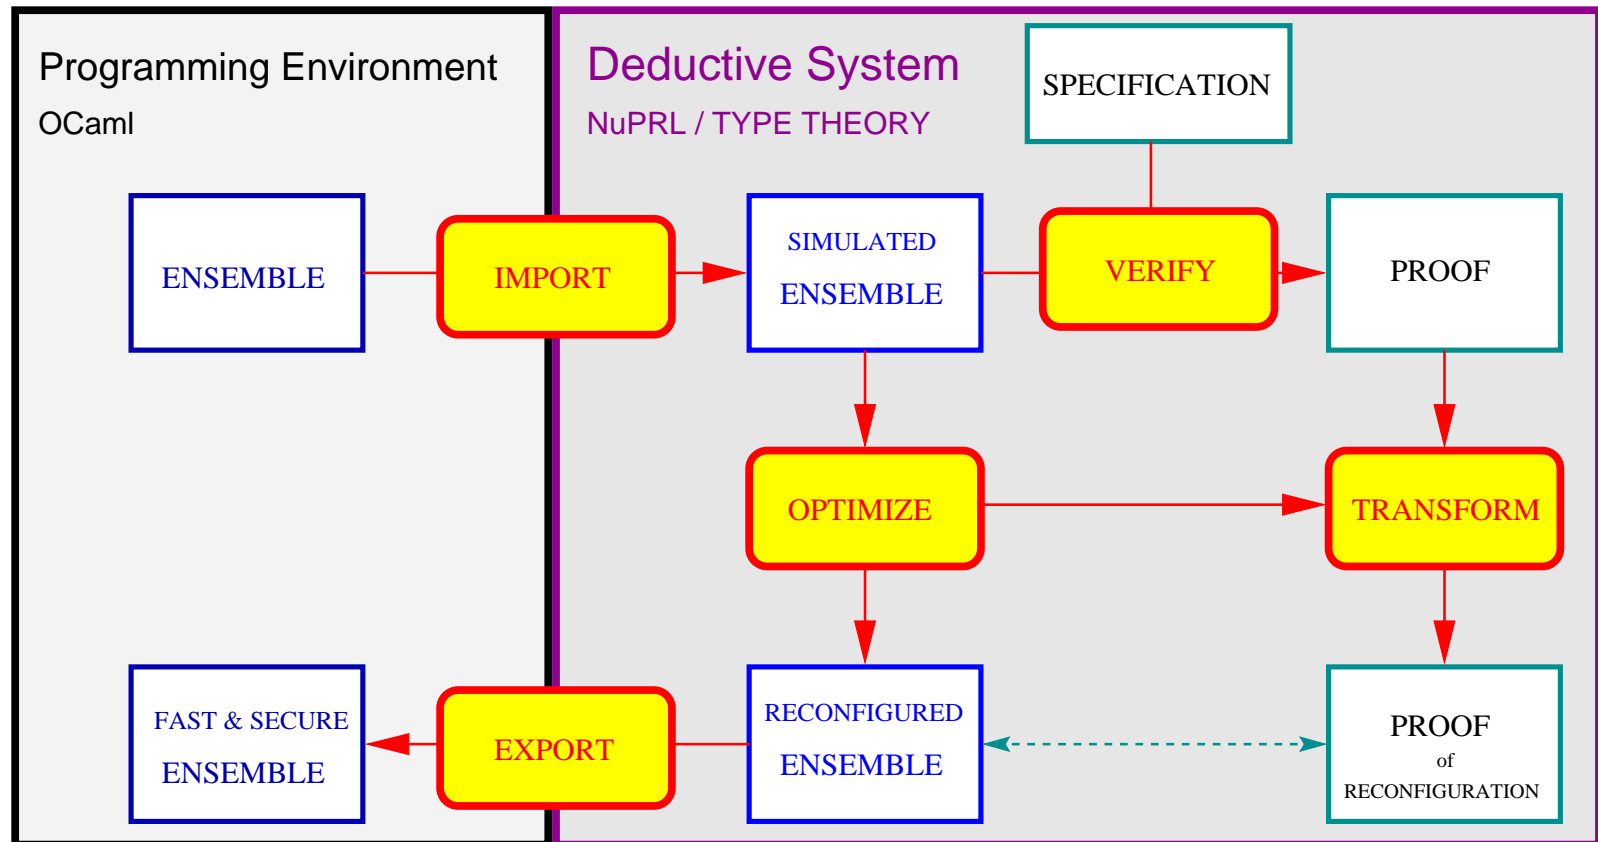
- Select from 60 micro-protocols for specific tasks
- Modules can be [stacked arbitrarily](#)
- Modeled as state/event machines

Implementation in **Objective Caml** ([INRIA](#))

- Easy maintenance (small code, good data structures)
- [Mathematical semantics](#), strict data type concepts
- Efficient compilers and type checkers




FORMAL REASONING ABOUT ENSEMBLE IN NUPRL



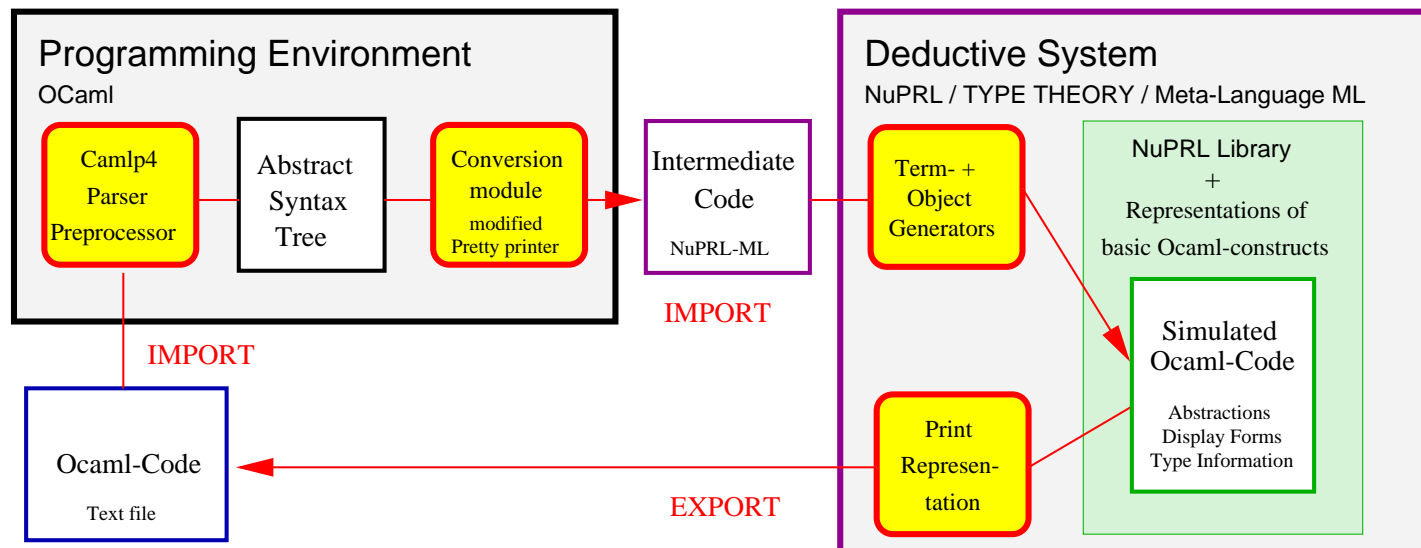
- Formalize semantics of OCAML (CADE 1998, ...)
- Verify protocol components and system configurations (TACAS 1999)
- Formally design and verify new protocols (DISCEX 2001, TPHOLS 2001)
- Optimize performance of configured systems (TACAS 1999, SOSP 1999)


EMBEDDING ENSEMBLE'S CODE INTO NUPRL

ENABLE FORMAL REASONING ON OCAML LEVEL



- **Type-theoretical semantics** of OCAML fragment
- NUPRL **implementation** captures syntax & semantics
- **Programming logic** for OCaml
- **Import and export** mechanisms





VERIFYING SYSTEM PROPERTIES

LINK FOUR LEVELS OF ABSTRACTION

Formalize system specification and code

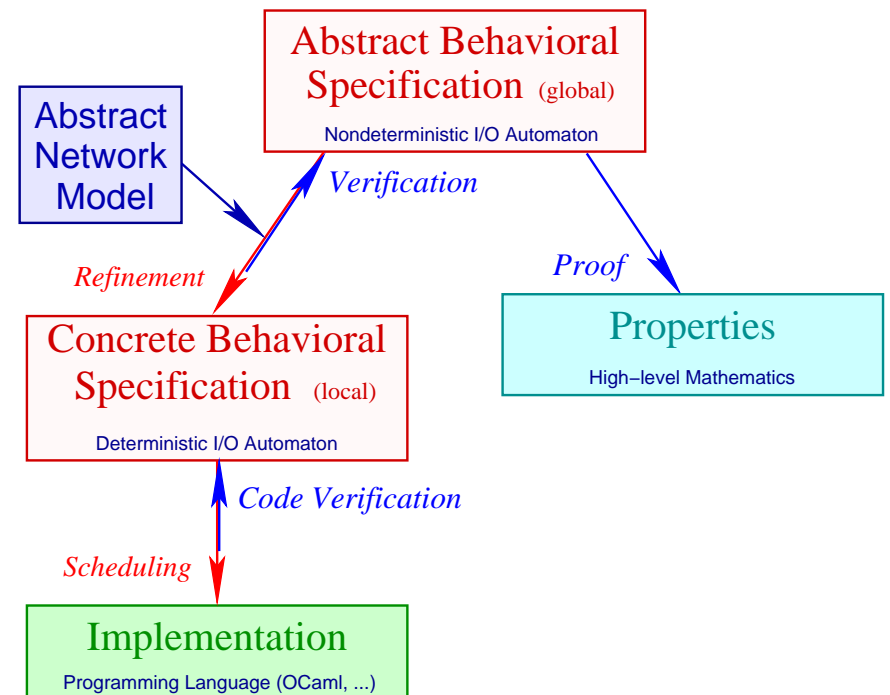
e.g. *“Messages are received in the same order in which they were sent”*

- *“Messages may be appended to global event queue and removed from its beginning”*
- *“Messages whose sequence number is too big will be buffered”*
- ENSEMBLE module Pt2pt.ml: 250 lines of OCAML code

All levels represented in type theory

Verification methodology

- Verify component specifications
(benign assumptions — subtle bug detected)
- Verify systems by composition
(IOA-composition preserves safety properties)
- Weave aspects
- Verify code



FORMAL DESIGN OF ADAPTIVE SYSTEMS

- **Make systems adapt safely to run-time dynamics**


- On-line upgrading, security, performance
- Difficult to design correctly *(distributed migration?)*

- **Generic switching protocol**

- Construct hybrid protocols from simpler ones
- **Normal mode**: interact with one protocol
- **Switching mode**: deliver old messages, buffer new ones

- **Correctness issues**

- *What kind of protocols are switchable at all?*
 - Reliability? Integrity? Confidentiality? Total Order? ...
- *What code invariant guarantees that switchable properties are preserved?*



LPE verification answers both questions

A FORMAL MODEL OF COMMUNICATION

● Communication **property** P

– Predicate on **traces**, i.e. lists of $\text{Send}(p,m)$ and $\text{Deliver}(p,m)$ events

e.g. **Reliable**(tr) $\equiv \forall p,q:\text{PID}.\forall m:\text{Msg}.\text{Send}(p,m) \in tr \Rightarrow \text{Deliver}(q,m) \in tr$

● Characterize **switchable** properties by **meta-properties**

– Predicates on communication properties

– Expressed by relation R between traces tr_u, tr_l above/below a protocol

P is R **property** $\equiv \forall tr_u, tr_l:\text{Trace}.\ (P(tr_l) \wedge tr_u R tr_l) \Rightarrow P(tr_u)$

Examples of meta-properties:

tr_u **safety** $tr_l \quad \equiv \quad tr_u \sqsubseteq tr_l$

tr_u **asynchrony** $tr_l \quad \equiv \quad tr_u$ swap-adjacent_[loc(e)≠loc(e')] tr_l

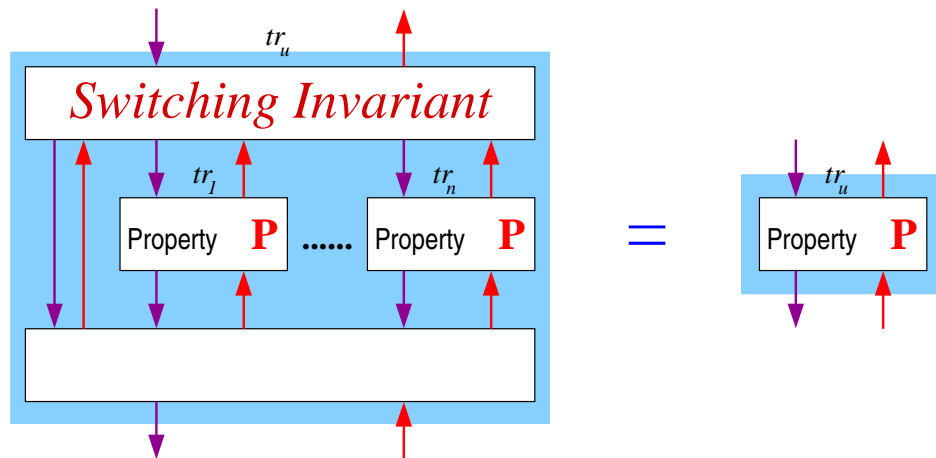
tr_u **delayable** $tr_l \quad \equiv \quad tr_u$ swap-adjacent_[msg(e)≠msg(e') ∧ is-send(e)≠is-send(e')] tr_l

tr_u **send-enabled** $tr_l \quad \equiv \quad \exists e:\text{Events}.\ \text{is-send}(e) \wedge tr_u = tr_l @ [e]$

tr_u **memoryless** $tr_l \quad \equiv \quad \exists e:\text{Events}.\ tr_u = [e_1 \in tr_l \mid \text{msg}(e) \neq \text{msg}(e_1)]$

composable(tr_u, tr_1, tr_2) $\equiv \quad tr_u = tr_1 @ tr_2 \wedge \forall e_1 \in tr_1.\ \forall e_2 \in tr_2.\ \text{msg}(e_1) \neq \text{msg}(e_2)$

VERIFYING THE CORRECTNESS OF SWITCHING



$\text{switchable}(P)$
 $\equiv P$ refines Causality
 $\wedge P$ refines No-replay
 $\wedge P$ is safety property
 $\wedge P$ is asynchrony property
 $\wedge P$ is delayable property
 $\wedge P$ is send-enabled property
 $\wedge P$ is memoryless property
 $\wedge P$ is composable property₃



- Characterize **switch invariant** between tr_u and tr_1, \dots, tr_n
 - tr_u results from joint trace by swapping events with different origin
 - Messages sent by different protocols must be delivered in same order
- Prove that **switchable properties will be preserved**

$\vdash \forall P:\text{TraceProperty}. \forall tr_u, tr_1, \dots, tr_n:\text{Trace}.$
 $\text{switchable}(P) \wedge \text{switch_invariant}(tr_u; tr_1, \dots, tr_n)$
 $\Rightarrow (\forall i \leq n. P(tr_i) \Rightarrow P(tr_u))$



Correct implementation and use of switch

OPTIMIZATION OF PROTOCOL STACKS

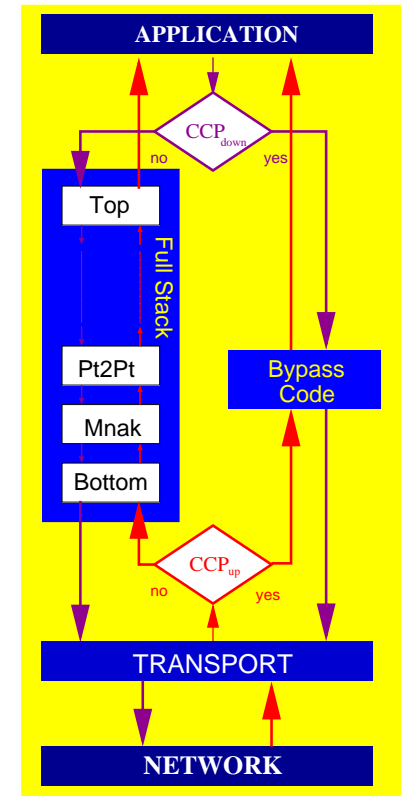


Protocol stacking creates performance loss
 – redundancy, internal communication, large message headers

Possible optimizations

- **Fast-path** for common execution sequences
 - Identify Common Case as Predicate
 - Analyze path of events through stack
 - Isolate code for fast-path and generate *bypass*
 - Insert CCP as runtime switch
- **Header compression** for common messages

Need formal reasoning tools to do this correctly



EXAMPLE PROTOCOL STACK Bottom::Mnak::Pt2pt

Trace downgoing Send events and upgoing Cast events

Bottom (200 lines)

```
let name = Trace.source_file "BOTTOM"
type header = NoHdr | ... | ...
type state = {mutable all_alive : bool ; ... }
let init _ (ls,vs) = {.....}
let hdlrs s (ls,vs)
  {up_out=up;upnm_out=upnm;
   dn_out=dn;dnlm_out=dnlm;dnnm_out=dnnm}
  = ...
  let up_hdlr ev abv hdr =
    match getType ev, hdr with
    | (ECast|ESend), NoHdr ->
      if s.all_alive or not (s.bottom.failed.(getPeer ev))
      then up ev abv
      else free name ev
  | :
  and uplm_hdlr ev hdr = ...
  and upnm_hdlr ev = ...
  and dn_hdlr ev abv =
    if s.enabled then
      match getType ev with
      | ECast -> dn ev abv NoHdr
      | ESend -> dn ev abv NoHdr
      | ECastUnrel -> dn (set name ev[Type ECast]) abv Unrel
      | ESendUnrel -> dn (set name ev[Type ESend]) abv Unrel
      | EMergeRequest -> dn ev abv MergeRequest
      | EMergeGranted -> dn ev abv MergeGranted
      | EMergeDenied -> dn ev abv MergeDenied
      | _ -> failwith "bad down event[1]"
    else (free name ev)
  and dnnm_hdlr ev = ...
  in {up_in=up_hdlr;uplm_in=uplm_hdlr;upnm_in=upnm_hdlr;
      dn_in=dn_hdlr;dnnm_in=dnnm_hdlr}
let l args vs = Layer.hdr init hdlrs args vs
Layer.install name (Layer.init l)
```

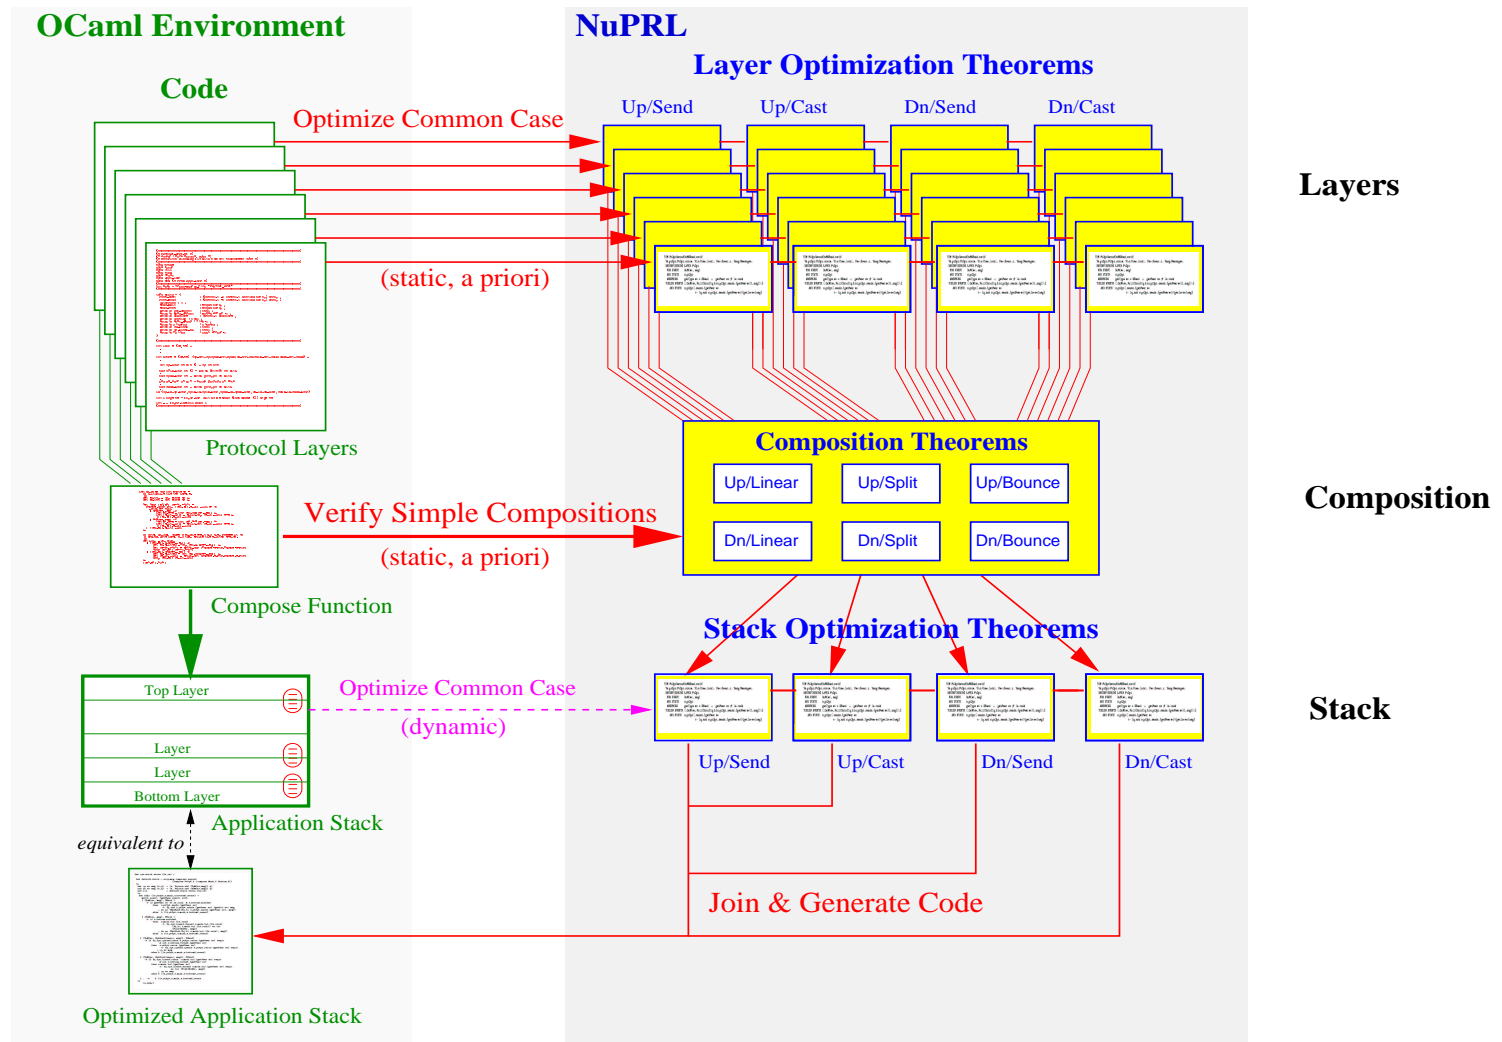
Mnak (350 lines)

```
let init ack_rate (ls,vs) = {.....}
let hdlrs s (ls,vs) { ..... }
  = ...
  let ...
  and dn_hdlr ev abv =
    match getType ev with
    | ECast ->
      let iov = getIov ev in
      let buf = Arraye.get s.buf ls.rank in
      let seqno = Iq.hi buf in
      assert (Iq.opt_insert_check buf seqno) ;
      Arraye.set s.buf ls.rank
        (Iq.opt_insert_doread buf seqno iov abv) ;
      s.acct_size <- s.acct_size + getIovLen ev ;
      dn ev abv (Data seqno)
    | _ -> dn ev abv NoHdr
  :
  :
```

Pt2pt (250 lines)

```
let init _ (ls,vs) = {.....}
let hdlrs s (ls,vs) { ..... }
  = ...
  let ...
  and dn_hdlr ev abv =
    match getType ev with
    | ESend ->
      let dest = getPeer ev in
      if dest = ls.rank then (
        eprintf "PT2PT:%s\nPT2PT:%s\n"
          (Event.to_string ev) (View.string_of_full (ls,vs));
        failwith "send to myself" ;
      ) ;
      let sends = Arraye.get s.sends dest in
      let seqno = Iq.hi sends in
      let iov = getIov ev in
      Arraye.set s.sends dest (Iq.add sends iov abv) ;
      dn ev abv (Data seqno)
    | _ -> dn ev abv NoHdr
  :
  :
```

METHODOLOGY: COMPOSE OPTIMIZATION THEOREMS



1. Use known optimizations of micro-protocols
2. Compose into optimizations of protocol stacks
3. Integrate message header compression
4. Generate code from optimization theorems and reconfigure system

A priori: ENSEMBLE + NUPRL experts

automatic: application designer

automatic: ⋮

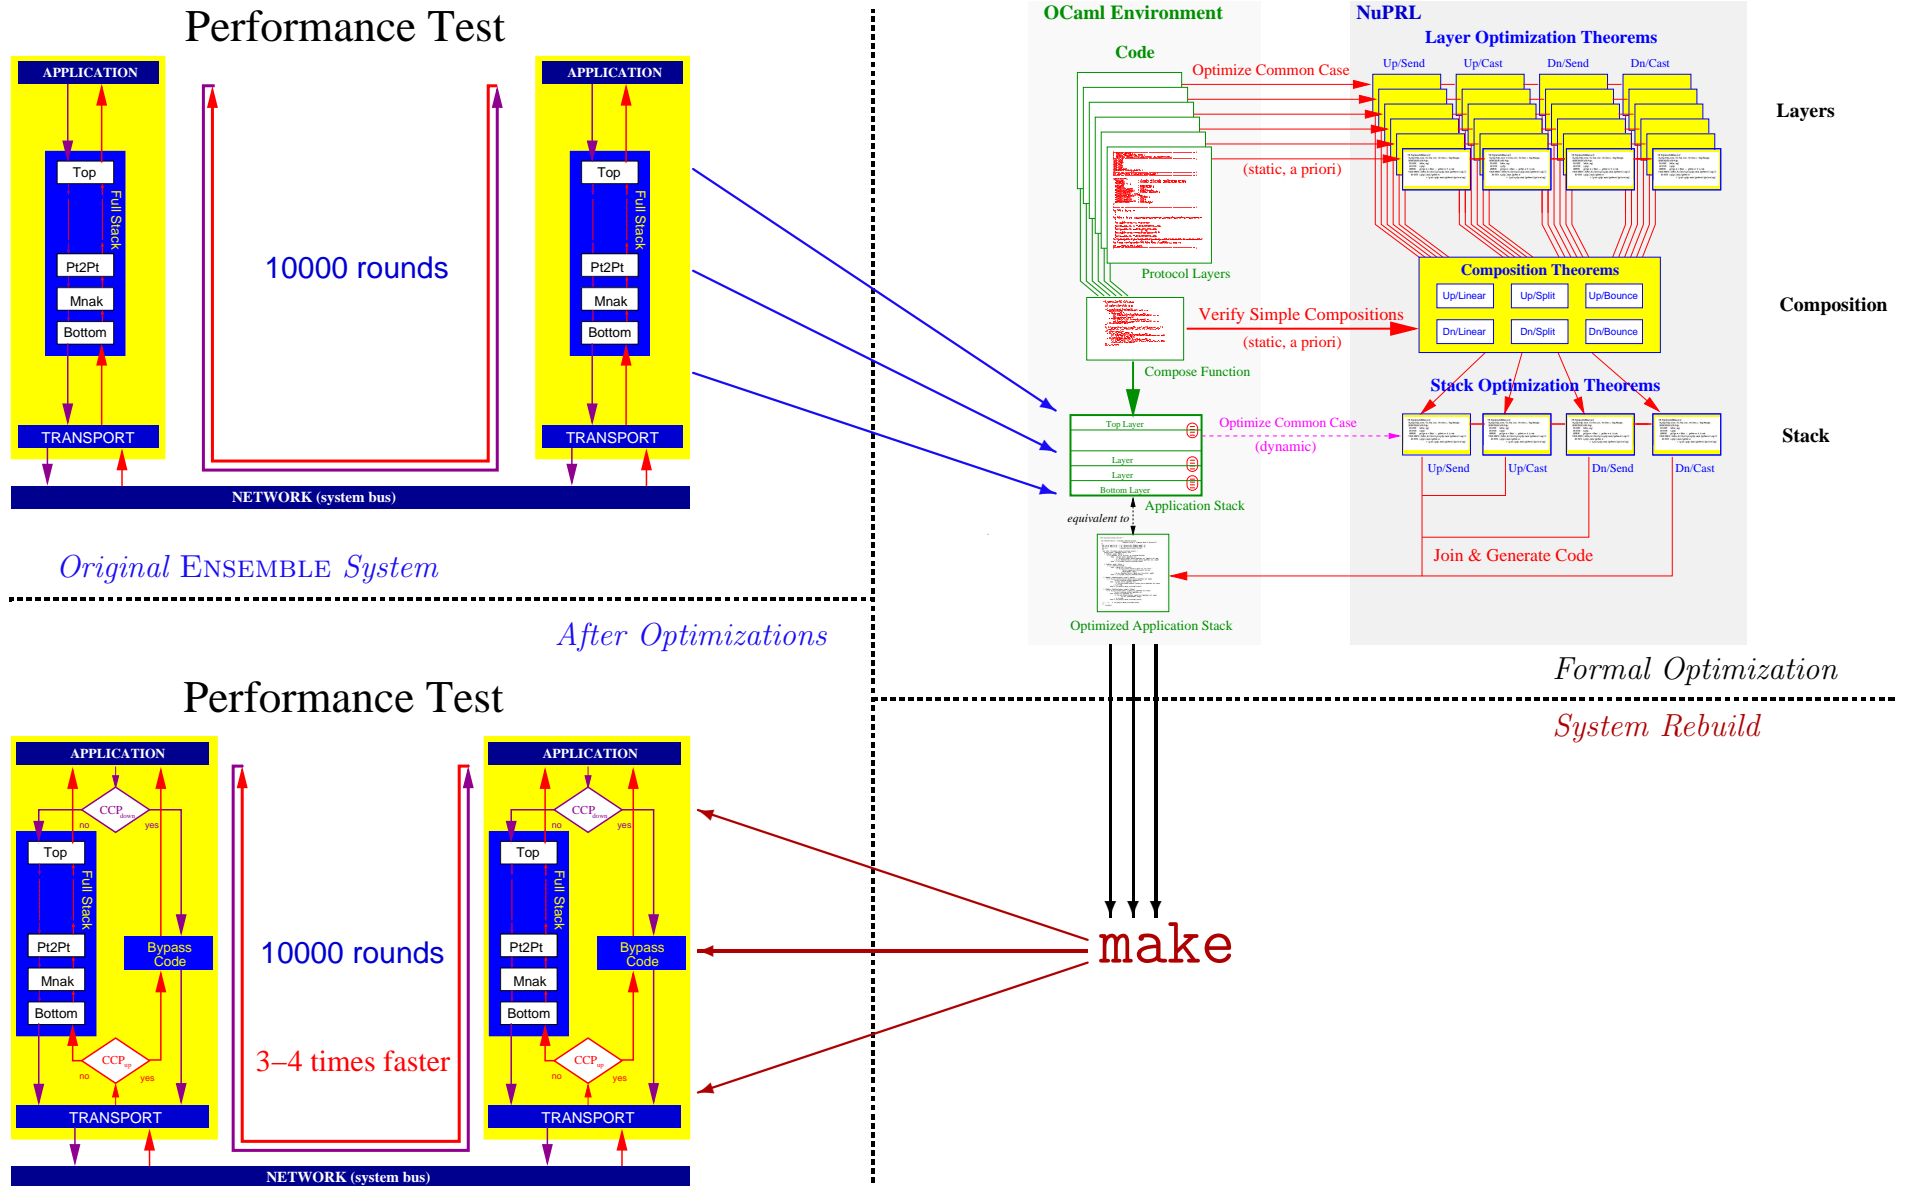
automatic: ⋮

Fast, error-free, independent of programming language

speedup factor 3-10

DEMO: OPTIMIZING A 24-LAYER PROTOCOL STACK

Top::Heal::Switch::Migrate::Leave::Inter::Intra::Elect::Merge::Slander::Sync::Suspect::Stable::Vsync::
 Partial_Appl::Total::Collect::Local::Frag::Pt2ptw::Mflow::Pt2pt::Mnak::Bottom



LESSONS LEARNED

● Results

- Type theory **expressive enough** to formalize today's software systems
- Formal optimization can significantly improve **practical performance**
- Formal verification **reveals errors** even in well-explored designs
- Formal design **reveals** hidden **assumptions** and **limitations** for use of protocols

● Ingredients for success . . .

- Implementation language with **precise semantics**
- **Collaboration** between systems and formal reasoning groups
- **Formal models** of: communication, I/O-automata, programming language
- **Knowledge-based** formal reasoning tools
- **Great colleagues!** **Stuart Allen, Mark Bickford, Ken Birman, Robert Constable, Richard Eaton, Xiaming Liu, Lori Lorigo, Robbert van Renesse**

FUTURE CHALLENGES

● **Advanced reasoning environment**

- Interactive **library** of formal algorithmic knowledge
 - **Archival** capacities (documentation & certification, version control)
 - A variety of **justifications** (levels of trust)
 - Creation of formal and textual **documents**
 - **Meta-reasoning** and **reflection**
- Embed **external library contents**
- Connect additional **proof engines**

Improve cooperation between research groups

● **Learn more from applications**

- Reasoning about **real-time, embedded** systems and **stream computations**
 - verified self-adaptation to changing **resource constraints** (MEDIANET)
- Support **programming languages** with less clean semantics
- Invert reasoning direction: from verification towards **network synthesis**

AREAS FOR STUDY & RESEARCH

- **Courses**

- Applied Logic, Automated Reasoning, PRL Seminar, ...

- **Formal Logics & Type Theory**

- Classes & inheritance, recursive & partial objects, concurrency, real-time
- Meta-reasoning, reflection, relating different logics, ...

- **Theorem Proving Environments**

- Logical accounting, theory modules, interfaces, proof presentation, ...

- **Automated Proof Search Procedures**

- Matrix methods, inductive theorem proving, rewriting, proof planning
- Decision procedures, extended type inference, cooperating provers
- Proof reuse, analogy, distributed proof procedures, ...

- **Applications**

- Formal CS knowledge: graph theory, automata, trees, arrays, ...
- Strategies for program synthesis, verification, and optimization
- Modelling programming languages (OCAML, JAVA, ..)

... Participation in applied research projects