

Designing Reliable, High-Performance Networks

...with the NUPRL Logical Programming Environment

Christoph Kreitz

Department of Computer Science, Cornell University

Ithaca, NY 14853



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 capable of integrating different reasoning techniques
- Collaboration between systems and formal methods experts in real applications

TYPE THEORY



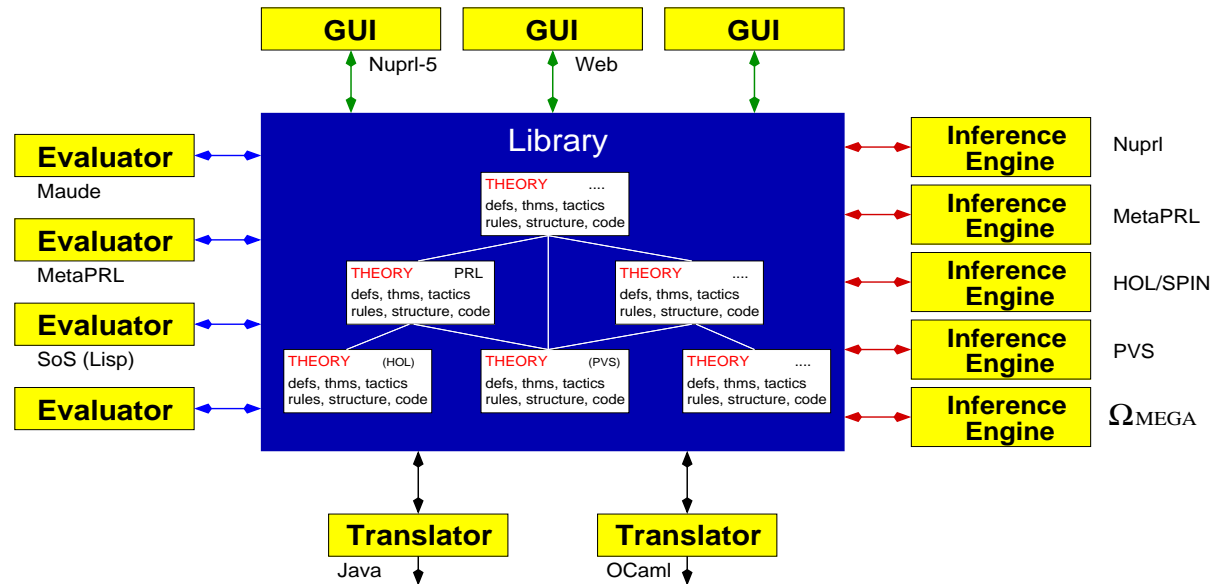
The NUPRL LPE meets these requirements

NUPRL'S FORMAL LOGIC: COMPUTATIONAL TYPE THEORY

- **Logic for constructive reasoning**
- **Open-ended, expressive 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
- **Top-down refinement calculus** ~> interactive proof development
 - Sequent calculus + computation rules + extract terms ~> program development
- **Expressions separate from their types** ~> full λ -calculus
 - ... but must be typeable *in proofs* ~> "total" functions
- **Uniform internal notation + display forms** ~> "free syntax"
- **User-defined extensions possible**

NUPRL'S AUTOMATED REASONING ENVIRONMENT

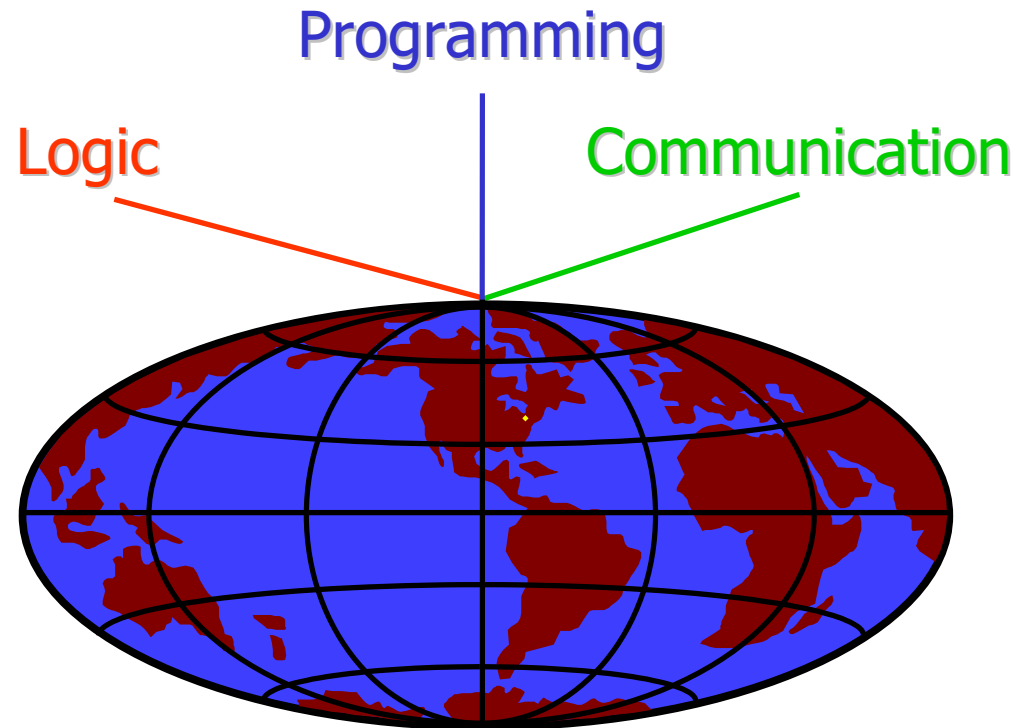


- **Interactive proof development**

- Supports program **extraction** (**synthesis**) and **evaluation**
- Proof automation through **tactics** & **decision procedures**
- Highly customizable: **language extensions**, **term display**, **system structure**,...

- **Cooperating processes centered around knowledge base** (CADE 2000)

- Large **library** of formal algorithmic knowledge
- **Asynchronous**, **distributed** & **collaborative** theorem proving
- Multiple user interfaces: **proof editor**, **structured term editor**, **web browser**
- **External proof engines**: **MetaPRL**, **JProver** (TPHOLs 2000, IJCAR 2001)



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

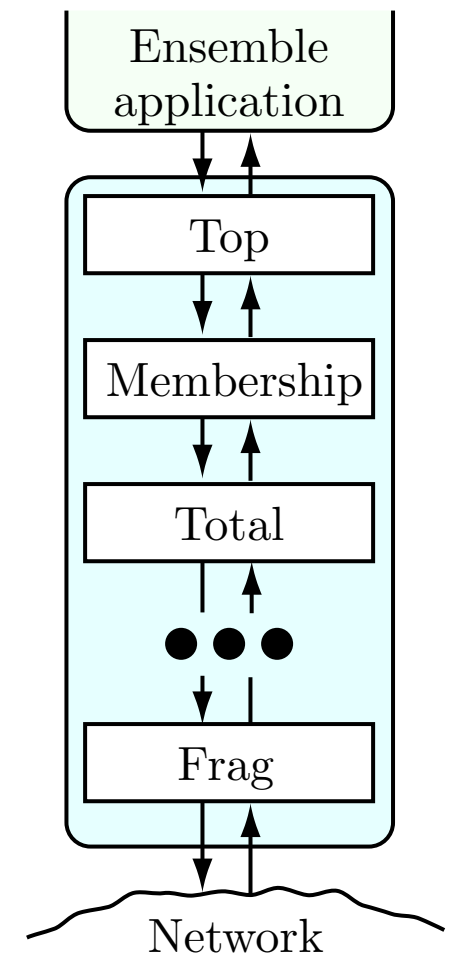
- Developed by Cornell's [System Group](#) (Ken Birman)
- Used commercially ([BBN](#), [JPL](#), [Segasoft](#), [Alier](#), [Nortel Networks](#))

Architecture: stack of **micro-protocols**

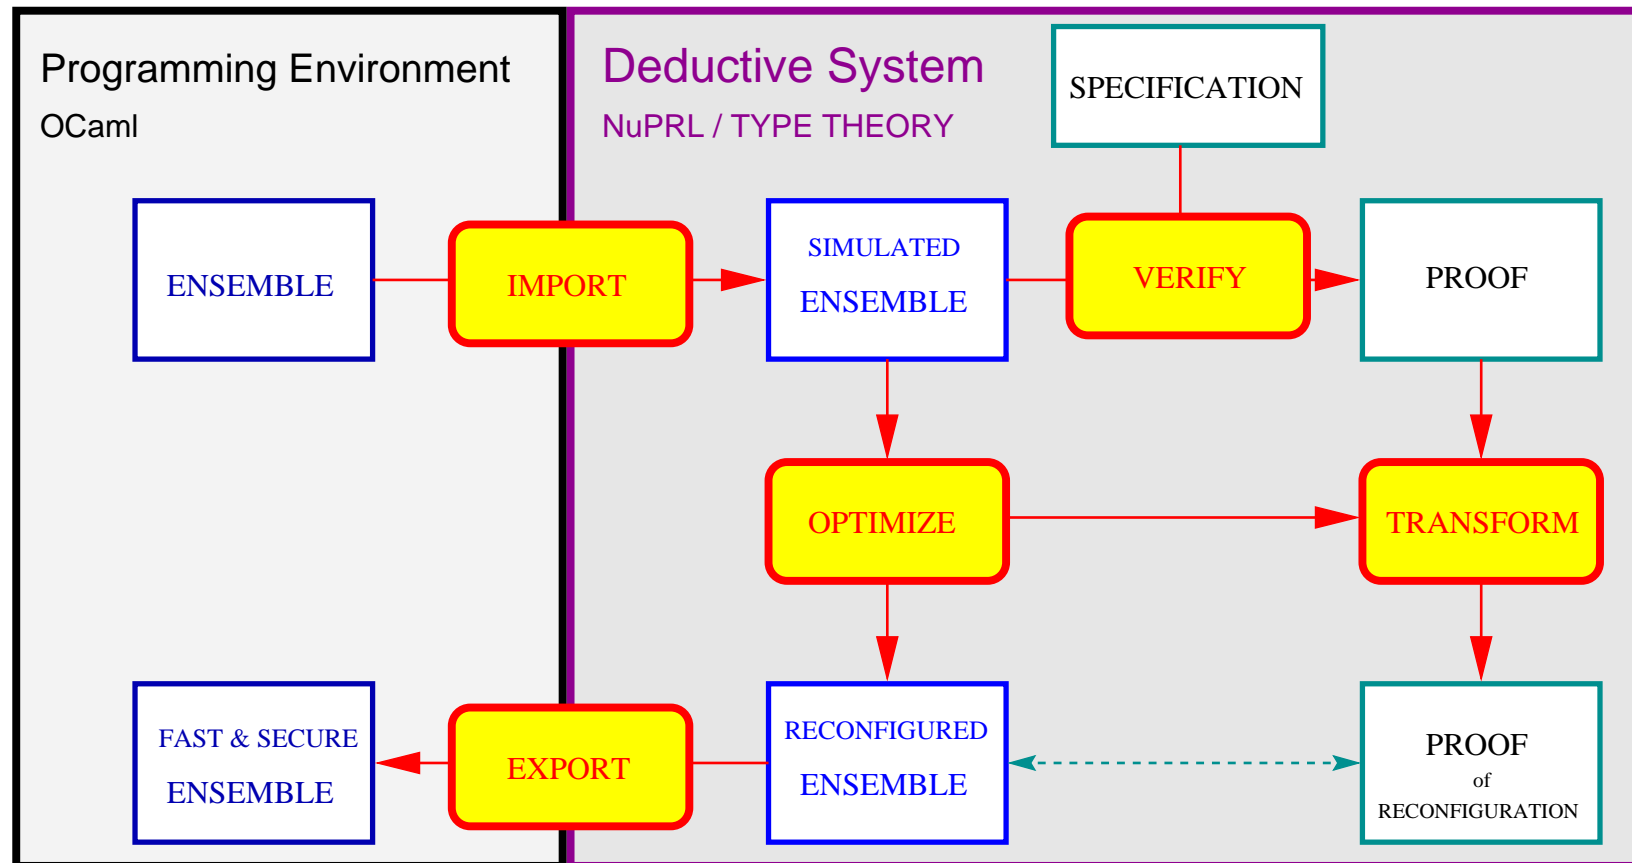
- Select from more than 60 micro-protocols for specific tasks
- Modules can be [stacked arbitrarily](#)
- Modeled as state/event machines

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



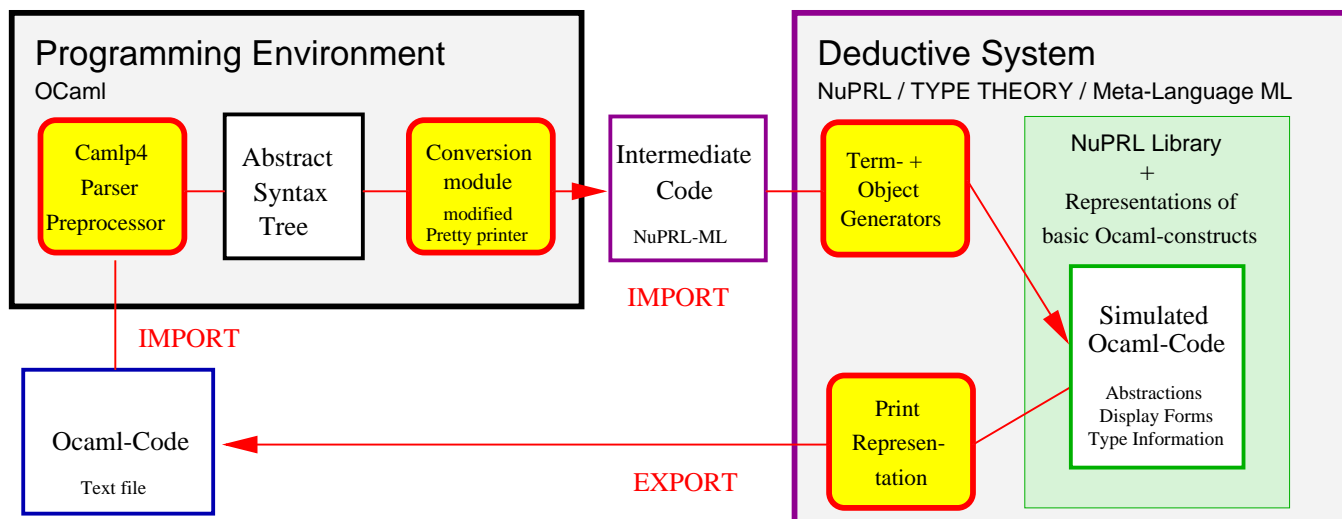
- Formalize semantics of OCAML (CADE 1998, ...)
- Optimize performance of configured systems (TACAS 1999, SOSP 1999)
- Verify protocol components and system configurations (TACAS 1999)
- Formally design and verify new protocols (DISCEX 2001, TPHOLS 2001)

EMBEDDING ENSEMBLE'S CODE INTO NUPRL

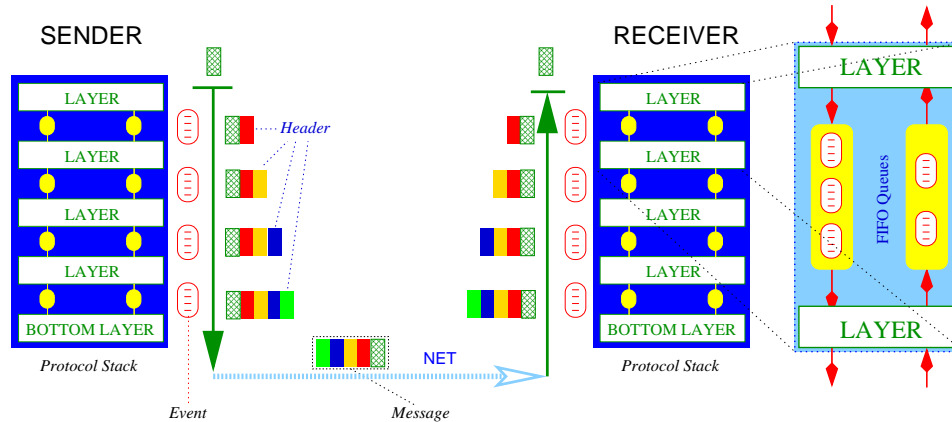

ENABLE FORMAL REASONING ON OCAML LEVEL



- **Type-theoretical semantics** of OCAML
 - Pattern matching, exceptions, references, modules, ... \mapsto type theory
- **Implementation** in NUPRL
 - OCAML semantics \mapsto **abstractions** OCAML syntax \mapsto **display forms**
- **Programming logic** for OCaml
 - Derived inference **rules** for reasoning about OCAML code
- **Import and Export** mechanisms
 - Actual system code available for formal reasoning in NUPRL



OPTIMIZATION OF PROTOCOL STACKS

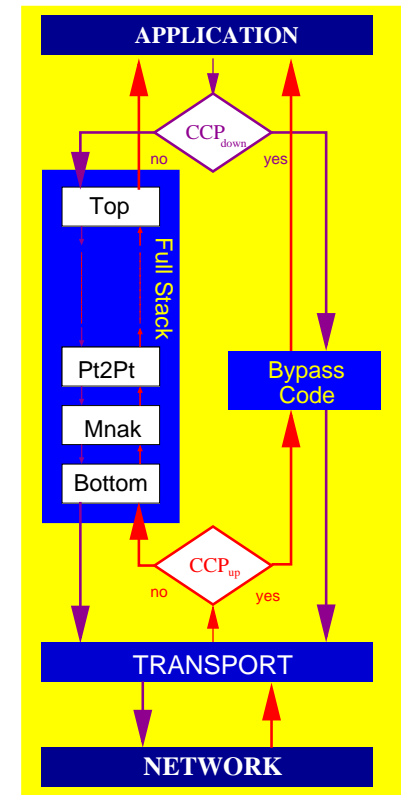


Protocol stacking creates performance loss
 – redundancy, internal communication, large message headers

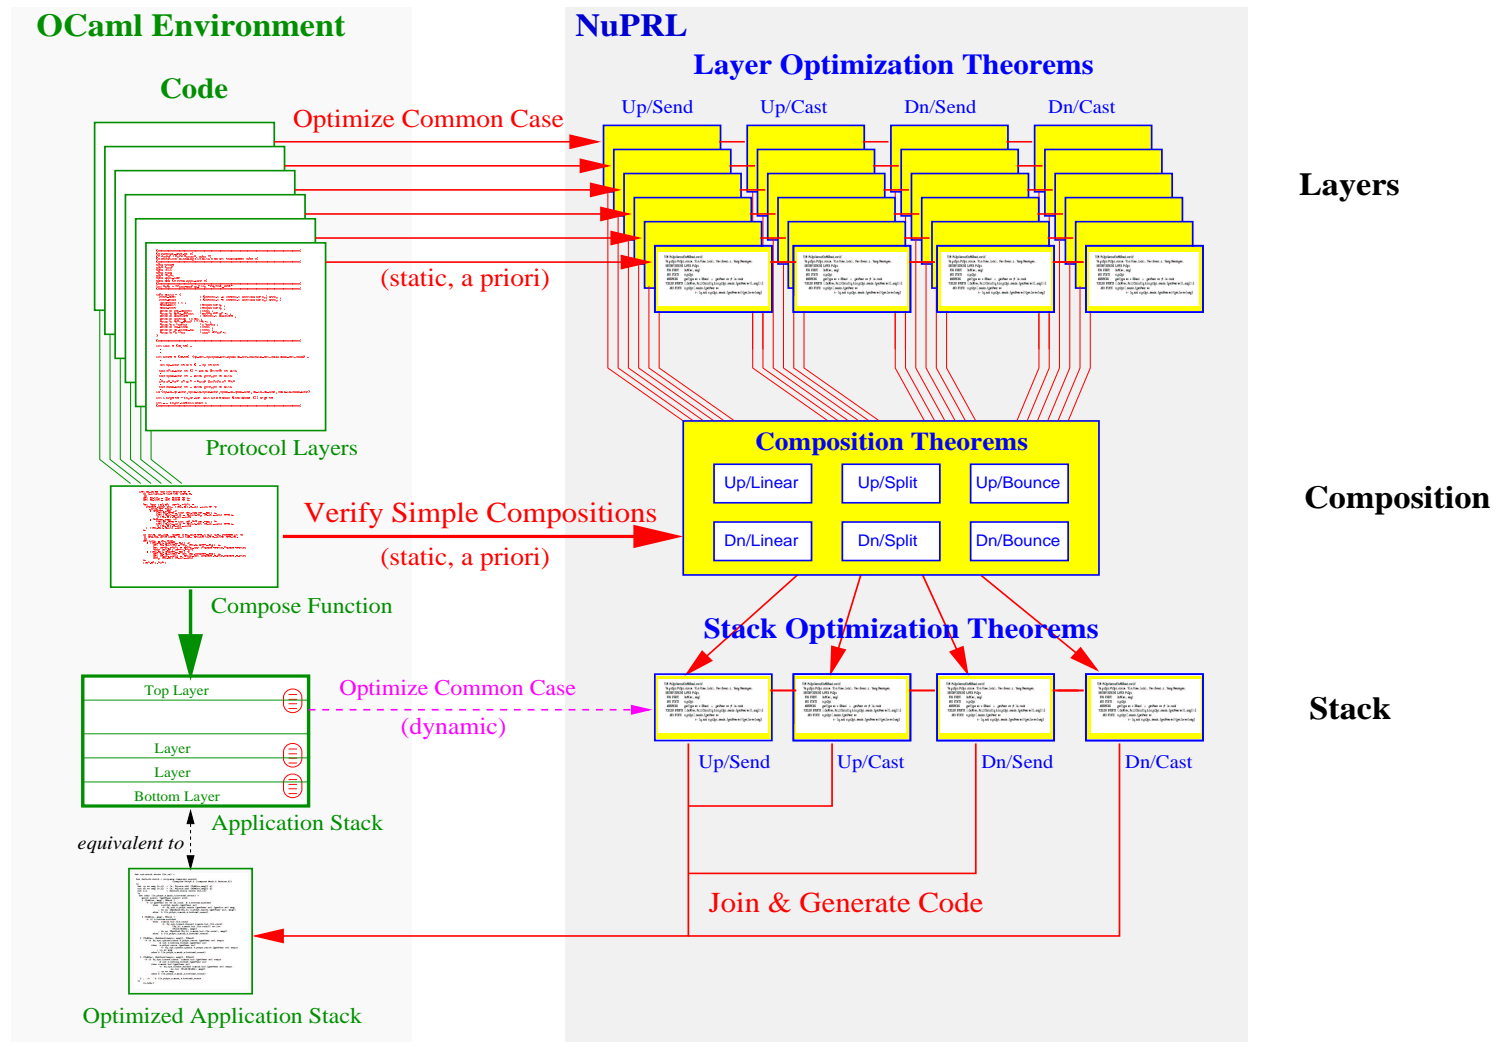
Possible optimizations

- **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



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: ⋮

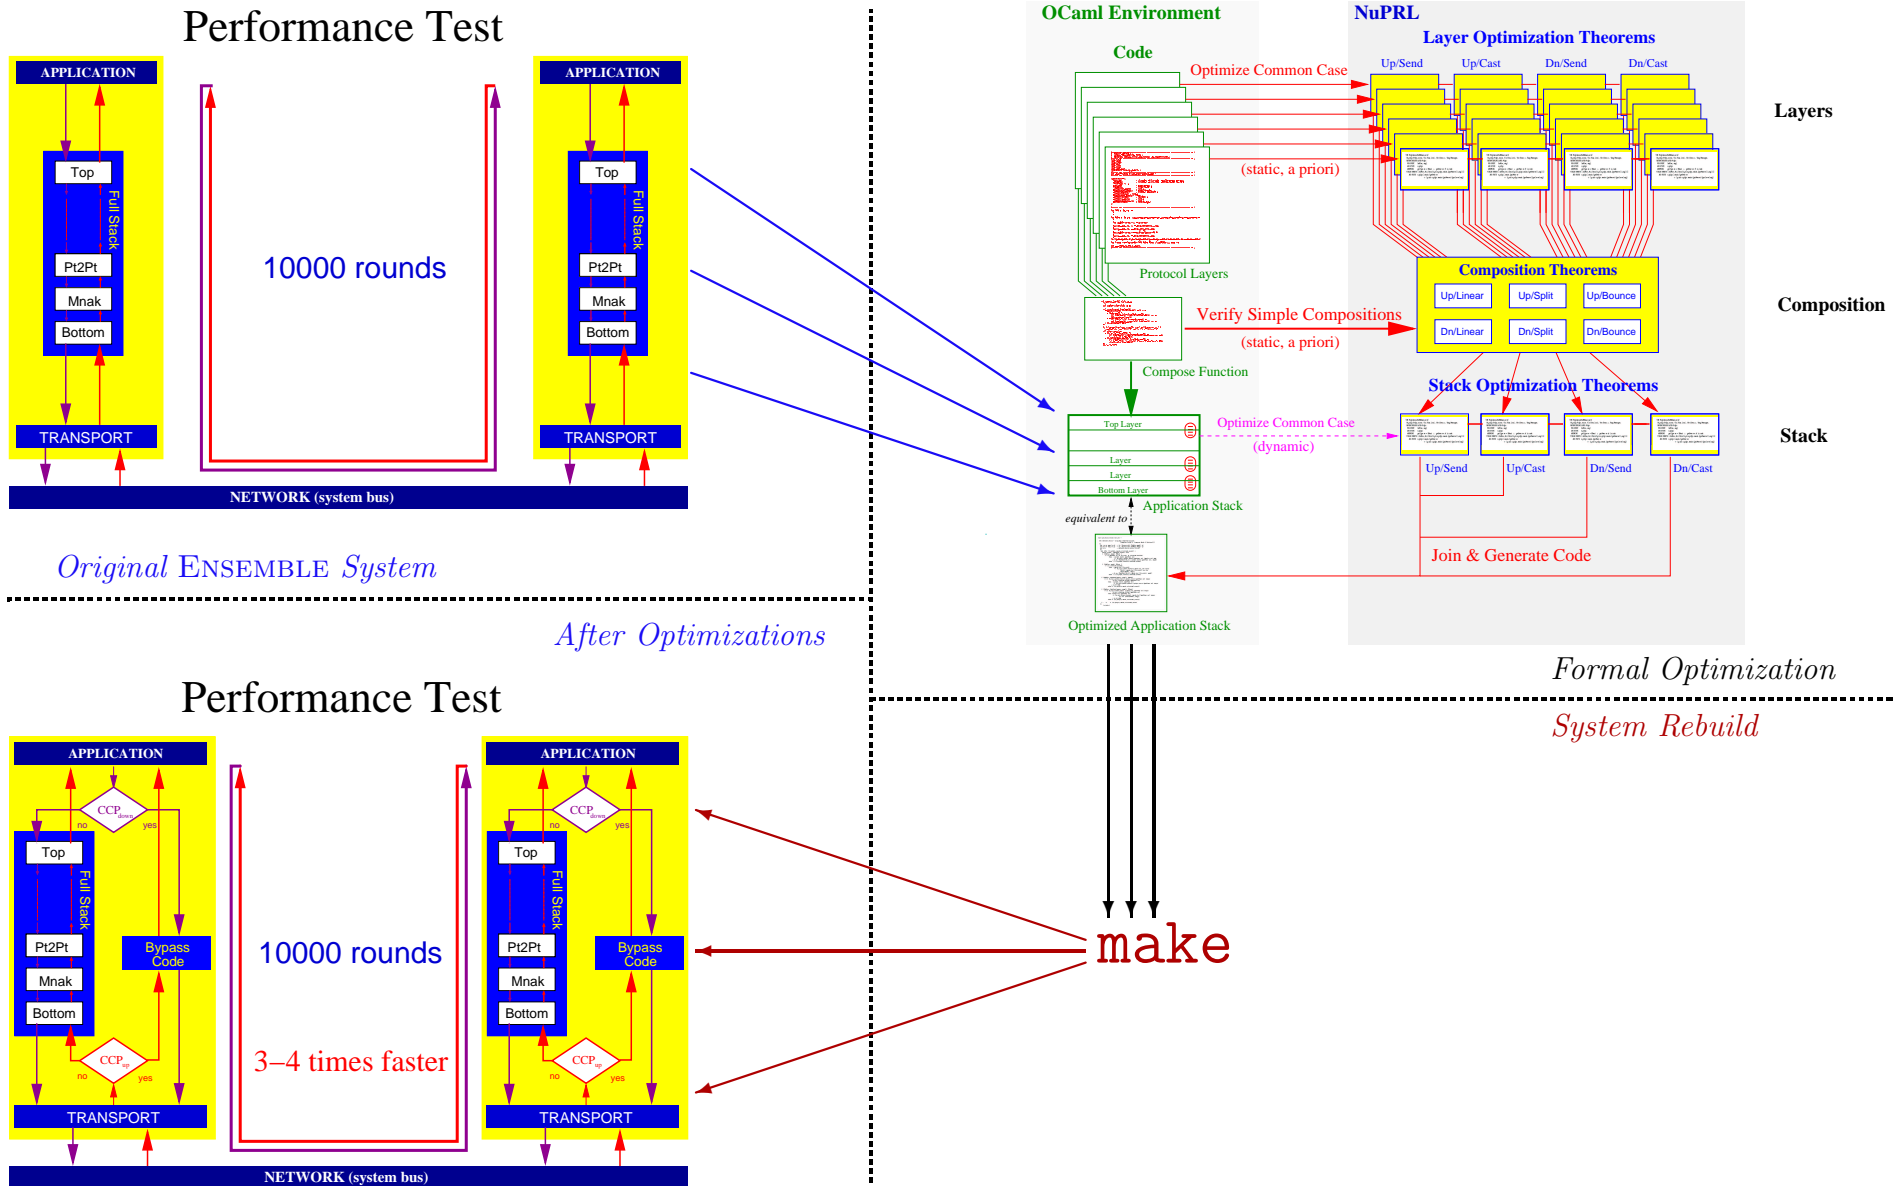
automatic: ⋮


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





VERIFICATION AND SYNTHESIS

LINK FOUR LEVELS OF ABSTRACTION

Formalize system specification and code

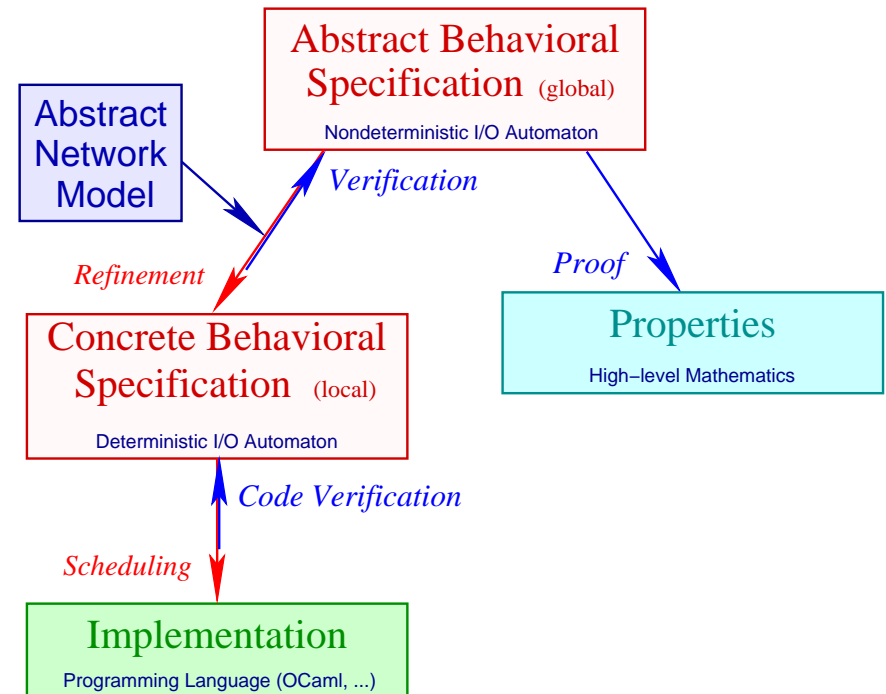
e.g. *“Messages are received in the same order in which they were sent”*

- *“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)
- Verify systems by composition (IOA-composition preserves safety properties)
- Weave aspects
- Verify code



Reasoning direction can be reversed into network synthesis

FORMAL DESIGN OF ADAPTIVE SYSTEMS

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

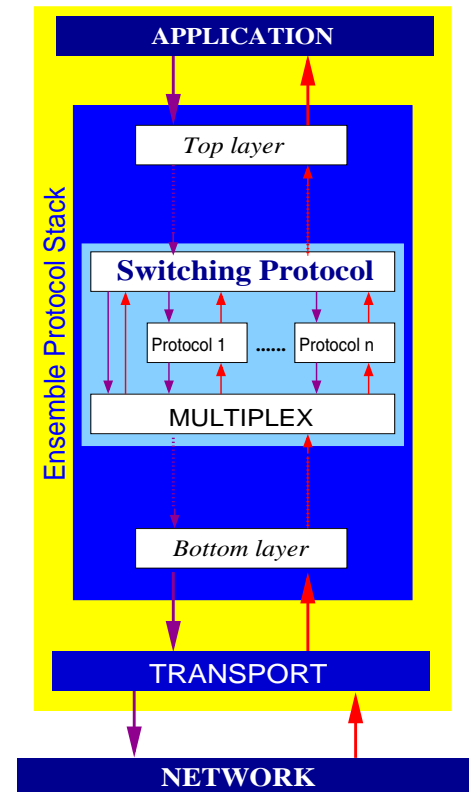
- On-line upgrading, security, performance
- Difficult to design correctly *(distributed migration?)*

- **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?*



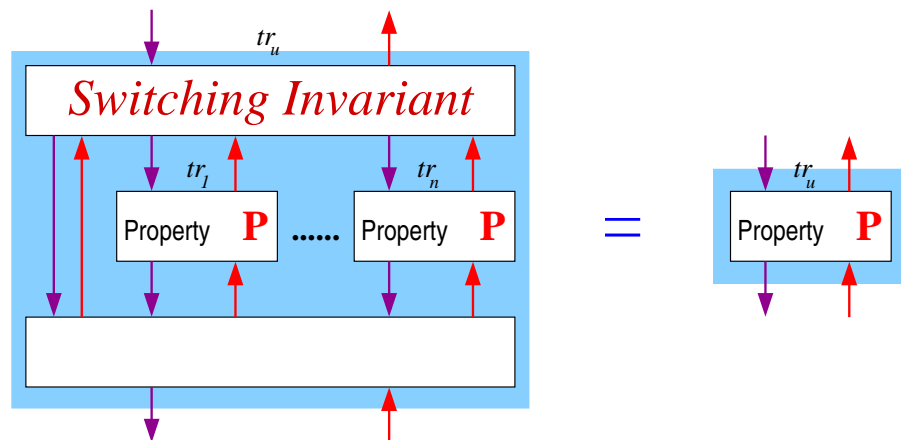
LPE verification answers both questions

VERIFYING THE CORRECTNESS OF SWITCHING

- Develop Formal Model of Communication

- Communication **property** P : predicate on **traces** (lists of Send/Deliver events)
- **Meta-property**: relation between traces that preserves properties

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



switchable(P)

- $\equiv P$ refines Causality
- $\wedge P$ refines No-replay
- $\wedge P$ safety property
- $\wedge P$ asynchrony property
- $\wedge P$ delayable property
- $\wedge P$ send-enabled property
- $\wedge P$ memoryless property
- $\wedge P$ 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 all switchable properties will be preserved



Correct implementation and use of switch

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
- ⇒ **NUPRL is capable of supporting real design at reasonable pace**

● Ingredients for success . . .

- Implementation language with **precise semantics**
- Employing formal methods at **every design stage**
- **Collaboration** between systems and formal reasoning groups
- **Formal models** of: communication, I/O-automata, programming language
- **Knowledge-based** approach: large library of algorithmic knowledge
- **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**: **PVS**, **HOL**, **MinLog**, ...

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**
- Support **programming languages** with less clean semantics
- Invert reasoning direction: from verification towards **network synthesis**