

# Paral-ITP Front-End Technologies and Isabelle Prover Architecture

Makarius Wenzel  
Univ. Paris-Sud, LRI

October 2013



Project **Paral-ITP** meeting  
ANR-11-INSE-001

# Overview

# Papers

1. M. Wenzel: [READ-EVAL-PRINT in Parallel and Asynchronous Proof-checking](#). User Interfaces for Theorem Provers ([UITP 2012](#)), postproceedings. EPTCS 118, 2013.
2. B. Barras, H. Herbelin, L. Gonzalez Huesca, Y. Regis-Gianas, E. Tassi, M. Wenzel, B. Wolff: [Pervasive Parallelism in Highly-Trustable Interactive Theorem Proving Systems](#). Conference on Intelligent Computer Mathematics ([CICM 2013](#)), 2013. Springer LNCS 7961.
3. C. Lange, M. Caminati, M. Kerber, T. Mossakowski, C. Rowat, M. Wenzel, W. Windsteiger: [A Qualitative Comparison of the Suitability of Four Theorem Provers for Basic Auction Theory](#). Conference on Intelligent Computer Mathematics ([CICM 2013](#)), 2013. Springer LNCS 7961.
4. M. Wenzel: [Shared-Memory Multiprocessing for Interactive Theorem Proving](#). Interactive Theorem Proving ([ITP 2013](#)), 2013. Springer LNCS 7998.

# Releases

## **Isabelle2013 (February 2013)**

- (cf. meeting March 2013)

## **Isabelle2013-1 (November 2013)**

- immutable kernel certificates
- monotonic document updates
- improved management of executions (goal forks, prints)
- asynchronous print functions (e.g. auto sledgehammer)
- query operations (e.g. sledgehammer, find theorems)
- improved multiplatform support (Linux, Windows, Mac OS X)

# Parallel Prover Architecture

# Summary


## **Change of Moore's Law:** (since 2005)

- application developers need to care about explicit parallelism
- **ignoring the challenge:** loose factor 10 on 16 core hardware today
- **mastering the challenge:** catch up with multiplication of cores (doubling every 18 months)

## **Parallel LCF proof processing:**

- problem structure: **practical proof irrelevance**
- programming paradigm: **purely functional programming** in ML (explicit context, immutable data)
- practical problems to achieve actual performance, but feasible

# Overall document structure



```

theory A imports  $B_1 \dots B_n$ 
begin
  :
  inductive P ...
  theorem a: A <proof>
  theorem b: B <proof>
  theorem c: C <proof>
  :
  have  $A \wedge B$ 
  proof
    show A by simp
    show B by blast
  qed
  :
  end
  
```

## Stack of parallel system layers

- Multicore Hardware (Intel, AMD)
- Operating System (Linux, Windows, Mac OS X)
- Poly/ML compiler and runtime system (David Matthews)
- Isabelle/ML futures and parallel skeletons: explicit parallelism
- Isabelle theory and proof processing: implicit parallelism
- Isabelle session build management: implicit parallelism, tree of ML processes
- asynchronous and parallel front-end technology (PIDE)



## Parallel Poly/ML (David Matthews, 2007, 2012)

**Hardware:** regular shared-memory multiprocessor (2–32 cores)

**Operating system:** native POSIX threads (pthreads)

**ML multithreading:** structures *Thread*, *Mutex*, *ConditionVar*

**ML memory management:**

- parallel garbage collection (various stages)
- online sharing of immutable values  
(reduced memory bandwidth requirements)

<http://www.polymml.org>

# Parallel Isabelle/ML

## Future values:

**type**  $\alpha$  *future*

**val** *Future.fork*:  $(\text{unit} \rightarrow \alpha) \rightarrow \alpha$  *future*

**val** *Future.join*:  $\alpha$  *future*  $\rightarrow \alpha$

**val** *Future.cancel*:  $\alpha$  *future*  $\rightarrow \text{unit}$

**strict evaluation:** spontaneous execution via thread-pool

**synchronous exceptions:** propagation within nested task groups

**asynchronous interrupts:** cancellation and signalling of tasks

**nested groups:** implicit block structure of parallel program

**dependencies:** implicit graph of tasks determined statically

# Explicit theory context (Paulson 1989)

## Main judgment:

$$\boxed{\Theta, \Gamma \vdash \varphi}$$

- background theory  $\Theta$   
(polymorphic types, constants, axioms; **global data**)
- proof context  $\Gamma$  (fixed variables, assumptions; **local data**)

## Operations on theories:

- extend and merge:  $\Theta_3 = \Theta_1 \cup \Theta_2 \cup \tau \cup c :: \tau \cup c \equiv t$
- symbolic sub-theory check:  $\Theta_1 \subseteq \Theta_2$
- transfer of results:  $\Theta_1 \subseteq \Theta_2 \implies \Theta_1 \vdash \varphi \implies \Theta_2 \vdash \varphi$

**Key benefit:** timeless and stateless prover kernel

# Proof promises

## Main ideas:

- closed proof constants as **place-holders** for future proofs
- **substitution** by finished proofs – from proper theory context!
- special support for **schematic polymorphism** of proofs

**New kernel inferences:** wrt. proof promise environment  $\Pi$

$$\frac{\text{FV } A = \emptyset \quad \text{TV } A = \{?\bar{\alpha}\}}{\Theta, \{a : A\}, \emptyset \vdash a[?\bar{\alpha}] : A[?\bar{\alpha}]} \text{ (promise)}$$

$$\frac{\Theta, \Pi, \Gamma \vdash p : B \quad \Theta_0, \emptyset, \emptyset \vdash q : A \quad \Theta_0 \subseteq \Theta}{\Theta, \Pi - \{a : A\}, \Gamma \vdash p[a := q] : B} \text{ (fulfill)}$$

# Goal forks

## Main ideas:

- specific infrastructure for **goal-directed proof** (via tactics)
- **global accounting** of forked proofs, avoid flooding by futures
- systematic **tracking of errors**

## ML interfaces:

**val** *Goal.prove*: *Proof.context* → *term* → *tactic* → *thm*

**val** *Goal.prove\_future*: *Proof.context* → *term* → *tactic* → *thm*

- same signature
- same semantics for **successful tactic** (without side-effects)
- same semantics for **failing tactic**, where it indicates global breakdown without local error handling

# Performance and Scalability

## Expected speedup in practice

**Real speedup:**  $\varepsilon(1) / \varepsilon(m)$  for  $m$  cores,  
relating elapsed run-time of sequential vs. parallel application

### Rules of Thumb:


**very naively**  $speedup = m$   
before you make an implementation

**naively**  $speedup \approx 1$   
first attempt with too little parallelism in the application

**asymptotically**  $speedup \approx 0$   
worse than Amdahl's law, excessive overhead for many cores

**realistically**  $speedup \approx 0.66 \times m$ , for reasonable  $m = 4, 8, 16, \dots$

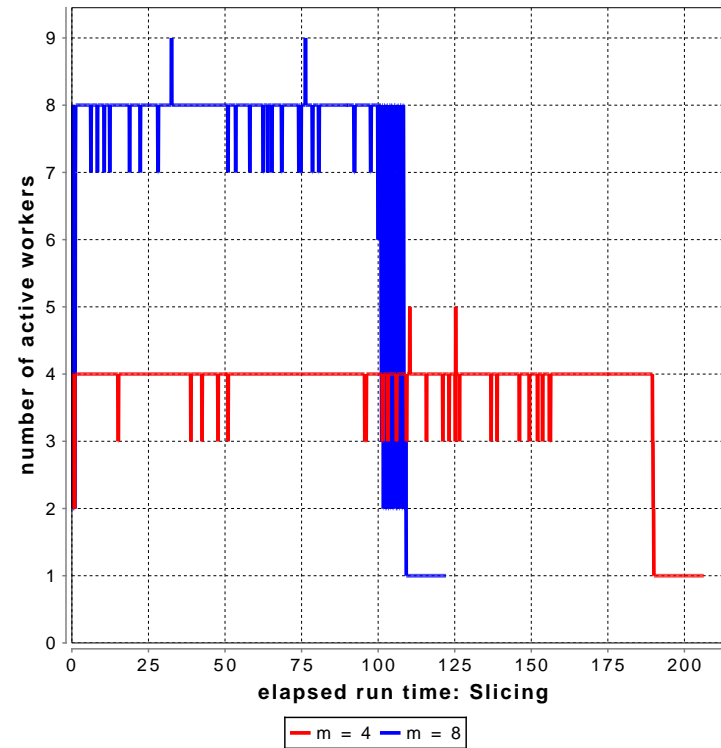
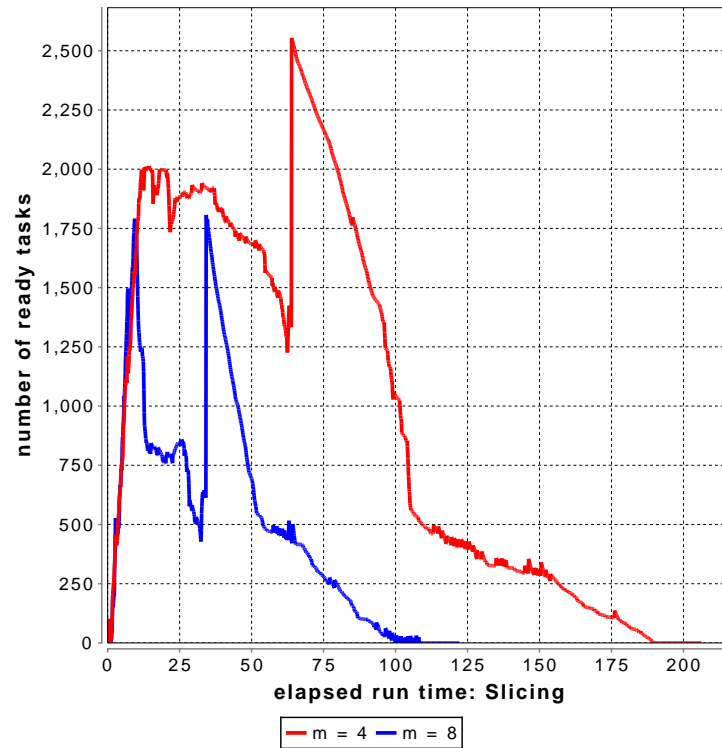
# Isabelle2013 (February 2013) on 8 cores



		$\varepsilon(1)$	$\varepsilon(8)$
A	<i>HOL</i>	210 s	85 s
B	<i>HOL-UNITY</i>	69 s	13 s
C	<i>HOL-Nominal_Examples</i>	526 s	91 s
D	<i>Slicing</i>	720 s	122 s
E	<i>HOL-Decision_Procs</i>	415 s	67 s
F	<i>HOL-Hoare_Parallel</i>	218 s	34 s



# Task queue population and worker thread utilization



# Saturation of threads vs CPUs (March 2013)



## Best speedups so far

### AFP/Slicing:

- factor 6.5: 8 threads on 8 cores  
February 2013 (official Isabelle2013)
- factor 9.5: 16 threads on 8 cores × hyperthreading,  
March 2013 with “prescient scheduling” of tasks (from last run)
- factor 12.5: 20 threads on 32 cores,  
April 2013 (measured by David Rager, Univ. of Texas)

**Asynchronous READ-EVAL-PRINT**

# Command Transactions

## Isolated commands:

- “small” toplevel state  $st$ : *Toplevel.state*
- command transaction  $tr$  as partial function over  $st$   
we write  $st \xrightarrow{tr} st'$  for  $st' = tr\ st$
- general structure:  $tr = read; eval; print$   
(for example  $tr = intern; run; extern$  in LISP)

## Interaction view:

$tr\ st =$

**let**  $eval = read\ src$  **in**      —  $read$  does not use  $st$   
**let**  $(y, st')$  =  $eval\ st$  **in**    — main transaction  
**let**  $() = print\ st'\ x$  **in**  $st'$  —  $print$  does not update  $st'$

**Note:** flexibility in separating  $read; eval; print$

# Document Structure

## Traditional structure:

- **local body:** linear sequence of **command spans**
- **global outline:** directed acyclic graph (DAG) of **theories**

## Notes:

- in **theory:** document consists single linear sequence

$$st \xrightarrow{tr} st' \xrightarrow{tr'} st'' \dots$$

- in **practice:** independent paths in graph important for parallelism

## Approach:

- incremental editing of command sequences
- parallel scheduling of resulting R-E-P phases
- continuous processing while the user is editing

# Document model with immutable versions

- overall *Document.state* with associated *Execution*
- document version contains command structure and assignment of “exec ids” for command transactions
- implicit sharing between versions (content and running commands)
- functional document update

*Document.define\_command: command\_id → src → state → state*

*Document.update: version\_id → version\_id → edit\* → state → state*

*Document.remove\_versions: version\_id\* → state → state*

*edit ≈ insert | remove | dependencies | perspective*

- global execution management

*Execution.start: unit → execution\_id*

*Execution.discontinue: unit → unit*

*Execution.running: execution\_id → exec\_id → bool*

*Execution.fork → exec\_id → (α → unit) → α future*

*Execution.cancel: exec\_id → unit*

# Asynchronous print functions

## Observations:

- cumulative PRINT operations consume more time than EVAL (output of goals is slower than most proof steps)
- PRINT depends on user perspective
- PRINT may diverge or fail
- PRINT augments results without changing proof state
- many different PRINTs may be run independently

## Approach:

- each command transaction is associated with several *exec\_ids*: one *eval* + many *prints*
- document content forms **union of markup**
- print management via **declarative parameters**: startup delay, time-out, task priority, persistence, strictness wrt. eval state



## Application: print proof state

- parameters:  $\{pri = 1, persistent = false, strict = true\}$
- change of perspective invokes or revokes asynchronous / parallel prints spontaneously
- GUI panel follows cursor movement to display content

The screenshot shows the jEdit IDE with a file named Unix.thy. The main editor displays the following Isabelle code:

```
theorem transition_type_safe:
  assumes tr: "root -x→ root'"
  and inv: "∃att dir. root = Env att dir"
  shows "∃att dir. root' = Env att dir"
proof (cases "path_of x")
  case Nil
  with tr inv show ?thesis
  by cases (auto simp add: access_def split: if_splits)
next
  case Cons
  from tr obtain opt where
    "root' = root ∨ root' = update (path_of x) opt root"
  bv cases auto
```

The Output panel on the right shows the current proof state:


```
using this:
  • root -x→ root'
  • ∃att dir.
    root =
      Env att dir
  • path_of x = []

goal (1 subgoal):
  1. ∃att dir.
    root' =
      Env att dir
```

The status bar at the bottom indicates the cursor is at line 468, column 15 (16972/37993) and the file is encoded in UTF-8.

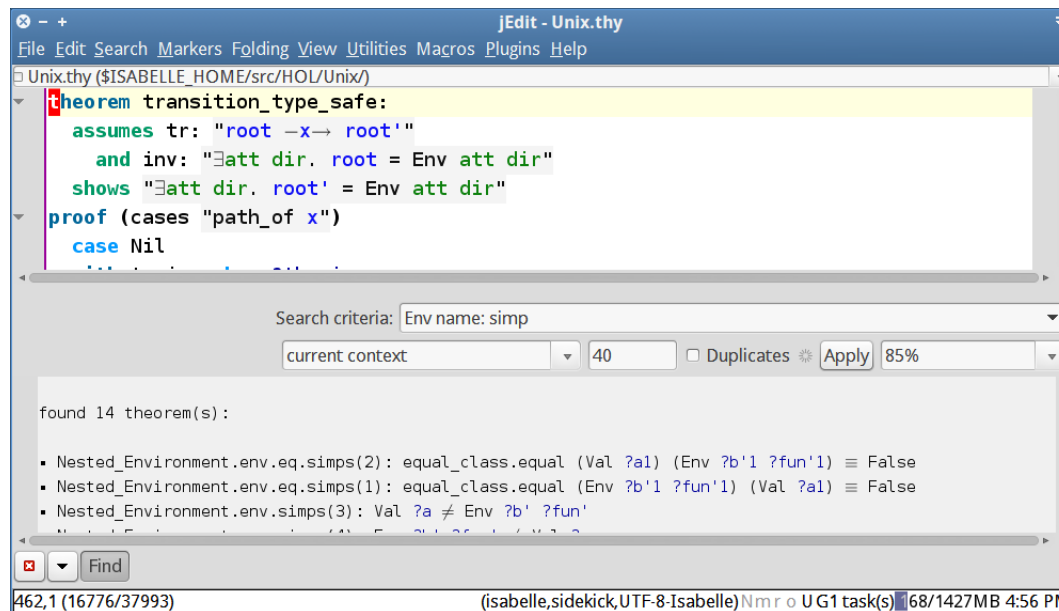
## Application: automatically tried tools

- parameters:  $\{delay = 1s, timeout = 4s, pri = -10, persistent = true, strict = true\}$
- long-running tasks with little output, e.g. automated (dis-)provers
- comment on existing document content via [information message](#)



## Application: query operations with user input

- parameters:  $\{pri = 0, persistent = false, strict = false\}$
- separate infrastructure to manage temporary **document overlays**
- stateful GUI panel with user input, system output, and control of corresponding command transaction (status icon, cancel button)



The screenshot shows the JEdit IDE with a file named Unix.thy. The editor displays a theorem proof:

```
theorem transition_type_safe:
  assumes tr: "root -x→ root'"
  and inv: "∃att dir. root = Env att dir"
  shows "∃att dir. root' = Env att dir"
proof (cases "path_of x")
case Nil
```

Below the editor, a search panel is visible. The search criteria are "Env name: simp". The search results show 14 theorems found, including:

- Nested\_Environment.env.eq.simps(2): equal\_class.equal (Val ?a1) (Env ?b'1 ?fun'1) ≡ False
- Nested\_Environment.env.eq.simps(1): equal\_class.equal (Env ?b'1 ?fun'1) (Val ?a1) ≡ False
- Nested\_Environment.env.simps(3): Val ?a ≠ Env ?b' ?fun'

The status bar at the bottom indicates the file is located at (isabelle,sidekick,UTF-8-Isabelle)Nmr o U G1 task(s) 68/1427MB 4:56 PM.

## Application: Sledgehammer

- heavy-duty query operation, with long-running ATPs and SMTs in the background (local or remote)
- progress indicator (spinning disk)
- clickable output
- implementation: trivial corollary of above concepts

```
File Edit Search Markers Folding View Utilities Macros Plugins Help
Scratch.thy (~/)
Lemma "[a] = [b] => a = b" by (metis the_elem_set)

Provers: e spass remote_vampire z3 remote_e_sine remote_waldmeis
[e] [s] [p] [r] [v] [z] [r] [e] [s] [r] [w] [a] [l] [d] [m] [e] [i] [s]
[ ] Isar proofs [ * ] [Apply] [Cancel] [Locate] [100%]

"e": Try this: by (metis the_elem_set) (9 ms).
"spass": Try this: by (metis list.inject) (15 ms).
"remote_vampire": Try this: by (metis list.inject) (10 ms).
"remote_e_sine": Try this: by (metis list.inject) (10 ms).
"remote_waldmeister": The generated problem lies outside the prover's scope.
"z3": Try this: by (metis list.inject) (8 ms).

Sledgehammer
5,26 (60/42556) (isabelle,sidekick,UTF-8-Isabelle)Nm r o U G 1352/1427M 35:02 PM
```

