

Extending ITL with Interleaved Programs for Interactive Verification

Gerhard Schellhorn

joint work with

Bogdan Tofan, Gidon Ernst, Kurt Stenzel,
Wolfgang Reif, Michael Balser, Simon Bäumler

Institute for Software and Systems Engineering
University of Augsburg

TIME, Lübeck, 13.9.2011

Background: Development of Correct Software


General Setting:

- Specification of Software Systems with:
Algebraic Specification, Z,
Abstract State Machines (ASMs)
- Incremental Refinement of Designs:
Algebraic, Data, ASM Refinement
- Verification of refinements:
Tool support with KIV Interactive Verifier

Background: Proving Sequential Programs with KIV

KIV is an interactive theorem prover based on

- Structured algebraic specification of data types with higher-order logic
- Sequent calculus with proof trees
- wp-calculus for ASMs and Java
- Proof principle for sequential programs: symbolic execution (+ induction) [BurSTALL 74] (= incremental computation of strongest postconditions for instructions)



Concurrent systems: What Logic to use?

Define a general logic which

- allows proofs for arbitrary properties: safety, liveness, deadlock, fairness, refinement (trace inclusion)
- can handle systems specifications that use abstract data types
⇒ interactive proving approach
- provides modular support for various forms of concurrency:
Programs with interleaving (“threading”)
Synchronous and asynchronous programs
Harel- and UML-Statecharts
(no encoding to transition systems)

Concurrent systems: What Calculus to use?

Define a calculus where

- proving properties (e.g. contracts) for sequential programs should not be more difficult than using wp-calculus
- compositional reasoning (e.g. rely-guarantee) is supported, as otherwise concurrency generates too many cases

Content of my talk:

- One particular answer to choosing a logic and a calculus, based on ITL [Moszkowski 00].
- Some applications for interleaved programs.


- **The Logic RGITL**
 - Compositional interleaving
 - A semantics with system and environment steps
 - Integration with HOL
- Proof principles in RGITL
 - Symbolic Execution
 - Induction
 - Rely-Guarantee
- Application: Lock-Free Algorithms
 - Motivation
 - Simple Example: Treiber's Stack
 - Linearizability and Lock-Freedom
- Experiences, Future Work

Why base the logic on ITL?

- + ITL directly offers termination/nontermination by using finite & infinite intervals
- + ITL is (easily) compatible with higher-order logic.
- + ITL offers the concept: programs \subseteq formulas.
The semantics of both is a set of intervals.
- Some small extensions are needed:
Is variable M in the program N := t?
Recursive procedures
- ITL does not offer a concept for interleaving.

Interleaving: Informal Semantics

Interleaved program $\{N := N^2; N := N^2\} \parallel N := N + 1$ started with $N = 2$:



Weak Fairness:

$\{\text{while } N \neq 0 \text{ do } N := N + 1\} \parallel N := 0$ terminates

Interleaving and Compositionality

A substitution rule is basic for a calculus to scale:

$$\frac{\alpha \rightarrow A \quad \beta \rightarrow B \quad A \oplus B \rightarrow C}{\alpha \oplus \beta \rightarrow C}$$

- holds in ITL for $\oplus =$ sequential composition and other operators (similar to Hoare calculus)
- ideally, third premise should be trivial
- should hold for \oplus interleaving too!

Example for Noncompositional Interleaving in ITL

In classical ITL:

$$\{\mathbf{while}^* N \neq 0 \mathbf{do} N := 0\} \leftrightarrow \{\mathbf{if}^* N \neq 0 \mathbf{then} N := 0\} \quad (1)$$

(the star indicates, that the test does not take time)

Example for Noncompositional Interleaving in ITL

In classical ITL:

$$\{\mathbf{while}^* N \neq 0 \mathbf{do} N := 0\} \leftrightarrow \{\mathbf{if}^* N \neq 0 \mathbf{then} N := 0\} \quad (1)$$

(the star indicates, that the test does not take time)

Using the substitution rule:

$$\{\mathbf{while}^* N \neq 0 \mathbf{do} N := 0\} \parallel \{\mathbf{while}^* N \neq 1 \mathbf{do} N := 1\} \quad (2)$$

\leftrightarrow

$$\{\mathbf{if}^* N \neq 0 \mathbf{then} N := 0\} \parallel \{\mathbf{while}^* N \neq 1 \mathbf{do} N := 1\} \quad (3)$$

which is **wrong**:

(2) has nonterminating runs, which alternate between the loops

(3) terminates, since at some time $N := 0$ is executed

Example for Noncompositional Interleaving in ITL

In classical ITL:

$$\{\mathbf{while}^* N \neq 0 \mathbf{do} N := 0\} \leftrightarrow \{\mathbf{if}^* N \neq 0 \mathbf{then} N := 0\} \quad (1)$$

(the star indicates, that the test does not take time)

Using the substitution rule:

$$\{\mathbf{while}^* N \neq 0 \mathbf{do} N := 0\} \parallel \{\mathbf{while}^* N \neq 1 \mathbf{do} N := 1\} \quad (2)$$

\leftrightarrow

$$\{\mathbf{if}^* N \neq 0 \mathbf{then} N := 0\} \parallel \{\mathbf{while}^* N \neq 1 \mathbf{do} N := 1\} \quad (3)$$

which is **wrong**:

(2) has nonterminating runs, which alternate between the loops

(3) terminates, since at some time $N := 0$ is executed

The problem is, that equivalence (1) ignores effects of the **environment** of the program

RGITL: Intervals with Environment Steps

Basic idea: environment steps between program steps

- Semantics is based on Intervals $I =$
sequence of states $(I(0), I'(0), I(1), I'(1), \dots)$
- state = valuation of variables
- I has finite (termination!) or infinite length $\# I \in \mathbf{N} \cup \{\infty\}$
- I alternates system steps $(I(0), I'(0)), (I(1), I'(1)), \dots$
with environment steps $(I'(0), I(1)), (I'(1), I(2)), \dots$
(similar to reactive sequences [deRoever 01])
- Programs determine system steps only
- Primed and double primed (flexible) variables are needed:
 X, X', X'' denote the value of X in $I(0), I'(0), I(1)$
($X = X' = X''$ in final states by convention)

Semantics of the Example in RGITL

The semantics of **while* N ≠ 0 do N := 0** now are intervals where N has values ($n_i \neq 0$):

- (0)
- ($n_0, 0, 0$) /* first env step does not change N */
- ($n_0, 0, n_1, 0, 0$) /* env sets N to n_1 */
- ($n_0, 0, n_1, 0, n_2, 0, 0$)
- ...
- Nonterminating run ($n_0, 0, n_1, 0, n_2, 0, \dots$)
- \Rightarrow The two programs are not equivalent
- But: equivalence is provable with environment assumption:
($\square N'' = N'$) \rightarrow
(**{while* N ≠ 0 do N := 0}** \leftrightarrow **{if* N ≠ 0 then N := 0}**)

Extends simply types lambda-expressions with

- static (x) and flexible variables (X,X',X'')
- formulas (= expressions of type bool) with:
 - ◇, □, **until**, **A**, **E** /* all paths/exists path */,
 - , ● /* strong/weak next state */,
 - last** /* termination */, ; /* chop */, * /* star */,
 - ||, ||_{nf} /* weak fair/nonfair interleaving */,
 - p(T;Y) /* procedure call with input an in-out parameters */

TL and HOL operators can be freely mixed


- Expressions are evaluated over algebras (constructed as models of algebraic specs.) and an interval $I = (I(0), I'(0), I(1), \dots)$
- If formula φ evaluates to true, write: $I \models \varphi$
- TL Operators have standard semantics:
 - $(I(0), I'(0), I(1), I'(1), \dots) \models \Box \varphi$
iff for all $n \leq \# I$: $(I(n), I'(n), I(n+1), I'(n+1), \dots) \models \varphi$
 - $I \models \mathbf{A} \varphi$ iff for all J with $J(0) = I(0)$: $J \models \varphi$
 - $I \models \mathbf{last}$ iff $I = (I(0))$
 - $(I(0), I'(0), \dots) \models \exists X. \varphi$
iff ex. (a_0, a'_0, \dots) with $(I(0)[X \leftarrow a_0], I'(0)[X \leftarrow a'_0], \dots) \models \varphi$

Programs in RGITL

- Programs α are formulas too:
 $I \models \alpha \Leftrightarrow$ the **system steps** in I are possible steps of α
- Programs: parallel assignments $\underline{X} := \underline{T}$,
sequential (**let**, **while**, **or**, **choose**, rec. procedures) +
 $\alpha \parallel \beta$ (interleaving), **await** C (block until C holds)
- Programs α are placed in a frame assumption $[\alpha]_{X,Y}$
to indicate which variables are fixed in assignments
(similar to TLA [Lamport 94], but no built-in stuttering)
- $[X := T]_{X,Y} \Leftrightarrow X' = T \wedge Y' = Y \wedge \circ$ **last**
- Typical goal: $\alpha \wedge E \rightarrow P$
“Executing α in environment E satisfies P ”

Semantics of Interleaving

- Interleaving of two programs (or formulas) α and β is defined compositionally, by interleaving individual intervals \Rightarrow substitution rule is valid!
- Assume $I_1 \models \alpha$, $I_2 \models \beta$
- Interleaving gives all intervals I which have
 - Interleaved system steps from I_1 and I_2 (fair)
 - The environment steps of I_1 (I_2) are the relevant alternating sequences of env. steps and system steps of β (α) in I



- Formal def. in paper, including **blocked** steps (tricky):
await $\varphi \equiv \text{while}^* \neg \varphi$ do blocked

- The Logic RGITL
 - Compositional interleaving
 - A semantics with system and environment steps
 - Integration with HOL
- **Proof principles of RGITL**
 - Symbolic Execution
 - Induction
 - Rely-Guarantee
- Application: Lock-Free Algorithms
 - Motivation
 - Simple Example: Treiber's Stack
 - Linearizability and Lock-Freedom
- Experiences, Future Work

Proof principle 1: Symbolic Execution

- Symbolic execution = Step forwards through an interval
- Advantage: **no encoding of programs as transition systems with program counters** (as in Step, TLA or Model checking)
⇒ readable goals
- Symbolic execution is done in two phases:
Unwinding and Stepping to the next state

Symbolic Execution: Unwinding (1)

- Splits formulas φ with $\underline{X} = \text{free}(\varphi)$ into formulas
 - $p(\underline{X}, \underline{X}', \underline{X}'')$ describing the first step
 - $\circ \psi$ describing properties of the rest of the run
- Termination gives formulas of the form $q(\underline{X}) \wedge \mathbf{last}$
- examples:

$$\square \varphi \equiv \varphi \wedge \bullet \square \varphi$$

$$\bullet \varphi \equiv \mathbf{last} \vee \circ \varphi$$

$$[X := T; \alpha]_{\underline{X}, \underline{Y}} \equiv X' = T \wedge \underline{Y}' = \underline{Y} \wedge \circ [\alpha]_{\underline{X}, \underline{Y}}$$

$$[\mathbf{let} X = T \mathbf{in} \alpha]_{\underline{Y}} \equiv \exists X. (X = T \wedge [\alpha]_{\underline{X}, \underline{Y}} \wedge \square X' = X'')$$

$$\begin{aligned} [\mathbf{choose} X \mathbf{with} \psi \quad &\equiv \quad \exists X. (\psi \wedge [\alpha]_{\underline{X}, \underline{Y}} \wedge \square X' = X'') \\ \mathbf{in} \alpha \mathbf{ifnone} \beta]_{\underline{Y}} &\quad \vee \quad (\neg \exists X. \psi) \wedge [\beta]_{\underline{Y}} \end{aligned}$$

Symbolic Execution: Unwinding (2)

To unwind interleaving and compounds unwind subprograms:

- If $\alpha \equiv p(X, X', X'') \wedge \circ \alpha'$ then

$$\{\alpha; \beta\} \equiv p(X, X', X'') \wedge \circ \{\alpha'; \beta\}$$

$$\{\alpha \parallel \beta\} \equiv \{\alpha <\parallel \beta\} \vee \{\alpha \parallel >\beta\}$$

$$\{\alpha <\parallel \beta\} \equiv p(X, X', X'') \wedge \circ \{\alpha' <\parallel \beta\}$$

- If $\alpha \equiv q(X) \wedge \mathbf{last}$ then

$$\{\alpha; \beta\} \equiv q(X) \wedge \beta$$

$$\alpha <\parallel \beta \equiv q(X) \wedge \beta$$

Symbolic Execution: Stepping

- Stepping removes the first step of interval:
Instead of $(l(0), l'(0), l(1), l'(1), \dots)$ consider $(l(1), l'(1), \dots)$
- Use new static variables x_0, x_1 to store $l(0)(X)$ and $l'(0)(X)$ of the old first step in $l(1)(x_0)$ and $l(1)(x_1)$

$$\frac{p(x_0, x_1, X) \wedge \psi}{p(X, X', X'') \wedge \circ \psi} \textit{step} \qquad \frac{q(x_0)}{q(X) \wedge \mathbf{last}} \textit{last}$$

- Effect: computation of the strongest postcondition of the first statement, weakened with environment assumption
 \Rightarrow sequential programs are executed as in wp-calculus
- Temporal properties result in (often non-temporal) additional goals for intermediate states

Proof principle 2: Induction

- Proofs use induction over well-founded orders
- Temporal induction reduced to well-founded induction by:
 - ◊ $\varphi \equiv \exists N. N = N'' + 1$ **until** φ
“There is a number N of steps after which φ holds”
- Note that $N = N'' + 1 \leftrightarrow N'' = N - 1 \wedge N > 0$
- Proof of $\square \varphi$ by contradiction:
Assume a number N of steps after which $\neg \varphi$ holds
Proof is then by well-founded induction over N
- Can be generalized to arbitrary safety properties
(e.g. sequential programs without local variables)

Induction to prove Fairness

- Weak Fairness: In an interleaving $\alpha \parallel \beta$, program α eventually gets a chance to do a step (if not blocked)
- In TLA: separate formula talking about encoded steps with program counters \Rightarrow not an option of RGITL
- Alternative: General transformation of fair to unfair interleaved programs using counters [Apt, Olderog 91]
- In RGITL: Add an “ α is scheduled flag” B :
$$\{B: \alpha \parallel \beta\} \leftrightarrow \{\alpha < \parallel \beta\} \vee (\neg B \wedge \{B: \alpha \parallel > \beta\})$$
- New Axiom: $\{\alpha \parallel \beta\} \leftrightarrow \exists B. \diamond B \wedge \{B: \alpha \parallel \beta\}$
- $\diamond B$ allows induction!
- Unfair interleaving satisfies almost the same axiom:
$$\alpha \parallel_{nf} \beta \equiv (\exists B. \diamond B \wedge \{B: \alpha \parallel_{nf} \beta\}) \vee (\beta \wedge \square (\neg \mathbf{blocked}) \wedge \mathbf{E} \exists \underline{x}. \alpha)$$
- $\mathbf{E} \exists \underline{X}. \alpha$: “there is at least one run of α ” ($\underline{X} = \text{free}(\alpha)$)

Proof principle 3: Compositional Reasoning

- Substitution principle allows to abstract each program in an interleaving to a property
- In particular: Rely/Guarantee rules are expressible
- Guarantee = Predicate for steps of a process $G(X, X')$
- Rely = Predicate on environment steps $R(X', X'')$
- Program α satisfies R/G, iff:



- As a TL formula: $R \xrightarrow{+} G \equiv \neg (R \text{ until } (\neg G))$
(not a special operator as in TLA [Lampert 94]!)

Proof principle 3: Compositional Reasoning

- Basic principle:
 - Prove R_i/G_i for interleaved programs α_i ($i = 1,2$)
 - Prove $G_i \rightarrow R_j$ for $i \neq j$, R_i transitive
 - Then: $\alpha_1 \parallel \alpha_2$ satisfies $\square G_1 \vee G_2$
- Provable by using the substitution principle, with
 $A \equiv R_1 \xrightarrow{+} G_1, B \equiv R_2 \xrightarrow{+} G_2, C \equiv (X' = X'') \xrightarrow{+} G_1 \vee G_2$

$$\frac{\alpha_1 \rightarrow A \quad \alpha_2 \rightarrow B \quad A \parallel B \rightarrow C}{\alpha_1 \parallel \alpha_2 \rightarrow C}$$

- First two premises = Assumptions for the two programs
- Third premise provable by induction, using

$$R \xrightarrow{+} G \leftrightarrow \forall B. \diamond B \rightarrow (R \wedge \neg B) \xrightarrow{+} G$$

Rely-Guarantee Theorem

Theorem

$$(1) \text{ pre} \wedge \text{CO}p_1 \rightarrow R_1 \xrightarrow{+} (G_1 \wedge (\mathbf{last} \rightarrow \text{post}_1))$$

$$(2) \text{ pre} \wedge \text{CO}p_2 \rightarrow R_2 \xrightarrow{+} (G_2 \wedge (\mathbf{last} \rightarrow \text{post}_2))$$

$$(3) G_1 \vee R \rightarrow R_2, G_2 \vee R \rightarrow R_1, G_1 \vee G_2 \rightarrow G$$

$$(4) \text{ reflexive}(G_1, G_2), \text{ transitive}(R_1, R_2)$$

$$(5) \text{ pre} \wedge (R_1 \vee R_2) \rightarrow \text{pre}$$

$$\text{then } \text{pre} \wedge \text{CO}p_1 \parallel \text{CO}p_2 \rightarrow R \xrightarrow{+} (G \wedge (\mathbf{last} \rightarrow \text{post}_1 \wedge \text{post}_2))$$

- similar to [Xu,deRoeper 97] (except cond. (5))
- Their notation for (1): $\text{CO}p_1 \underline{\text{sat}} (\text{pre}, \text{rely}_1, \text{guar}_1, \text{post}_1)$
- Deadlock freedom provable too (using **blocked** \rightarrow *wait*)

- The Logic RGITL
 - Compositional interleaving
 - A semantics with system and environment steps
 - Integration with HOL
- Proof principles of RGITL
 - Symbolic Execution
 - Induction
 - Rely-Guarantee
- **Application: Lock-Free Algorithms**
 - Motivation
 - Simple Example: Treiber's Stack
 - Linearizability and Lock-Freedom
- Experiences, Future Work

Motivation

- Multi-core processors getting more and more common
⇒ Concurrent algorithms more important than ever
- Usually, concurrency is implemented using locks (semaphores, synchronize in Java etc.)
- Lock-free algorithms (also called nonblocking) are an interesting class of algorithms that does not use locks
- Instead they use **CAS instructions** (x86, Sparc, Itanium) or LL/SC (Alpha, PowerPC)


Example: Treiber's Stack

- Defined in [Treiber 86]
- Implementation of a global stack
- Abstract view: Operations APush and APop
- Implementation with algorithms CPush and CPop
- Representation of stack as a linked list.

Treiber's Stack - Push

```
Cpush(v :Data; top : Pointer) {  
  n := new(Node);  
  n.val := v;  
  success := false;  
  while success = false do {  
    tmp := top;  
    /* other process .. */  
    /* .. may change top! */  
    n.next := tmp;  
    CAS(tmp, n, top)}
```

```
CAS(tmp, n, top; success) { /* atomic ! */  
  if* top = tmp  
    then { top := n, success := true; }  
  else success := false; }
```



Treiber's Stack - Push

```
CPush(v :Data; top : Pointer) {
```

```
  n := new(Node);
```

```
  n.val := v;
```

```
  success := false;
```

```
  while success = false do {
```

```
    tmp := top;
```

```
    /* other process .. */
```

```
    /* .. may change top! */
```

```
    n.next := tmp;
```


```
    CAS(tmp, n, top)}
```

```
CAS(tmp, n, top; success) { /* atomic ! */
```

```
  if* top = tmp
```

```
    then { top := n, success := true; }
```


```
  else success := false; }
```



Treiber's Stack - Push

```
Cpush(v :Data; top : Pointer) {  
  n := new(Node);  
  n.val := v;  
  success := false;  
  while success = false do {  
    tmp := top;  
    /* other process .. */  
    /* .. may change top! */  
    n.next := tmp;  
    CAS(tmp, n, top)}
```


```
CAS(tmp, n, top; success) { /* atomic ! */  
  if* top = tmp  
    then { top := n, success := true; }  
  else success := false; }
```



Treiber's Stack - Push

```
CPush(v :Data; top : Pointer) {  
  n := new(Node);  
  n.val := v;  
  success := false;  
  while success = false do {  
    tmp := top;  
    /* other process .. */  
    /* .. may change top! */  
    n.next := tmp;  
    CAS(tmp, n, top)}  
}
```


```
CAS(tmp, n, top; success) { /* atomic ! */  
  if* top = tmp  
    then { top := n, success := true; }  
  else success := false; }
```



Treiber's Stack - Push

```
Cpush(v :Data; top : Pointer) {  
  n := new(Node);  
  n.val := v;  
  success := false;  
  while success = false do {  
    tmp := top;  
    /* other process .. */  
    /* .. may change top! */  
    n.next := tmp;  
    CAS(tmp, n, top)}  
}
```


```
CAS(tmp, n, top; success) { /* atomic ! */  
  if* top = tmp  
    then { top := n, success := true; }  
  else success := false; }
```



Treiber's Stack - Push

```
CPush(v :Data; top : Pointer) {  
  n := new(Node);  
  n.val := v;  
  success := false;  
  while success = false do {  
    tmp := top;  
    /* other process .. */  
    /* .. may change top! */  
    n.next := tmp;  
    CAS(tmp, n, top)}  
}
```


```
CAS(tmp, n, top; success) { /* atomic ! */  
  if* top = tmp  
    then { top := n, success := true; }  
  else success := false; }
```



Treiber's Stack - Push

```
CPush(v :Data; top : Pointer) {  
  n := new(Node);  
  n.val := v;  
  success := false;  
  while success = false do {  
    tmp := top;  
    /* other process .. */  
    /* .. may change top! */  
    n.next := tmp;  
    CAS(tmp, n, top)}  
}
```


```
CAS(tmp, n, top; success) { /* atomic ! */  
  if* top = tmp  
  then { top := n, success := true; }  
  else success := false; }
```



Treiber's Stack - Push

```
CPush(v :Data; top : Pointer) {  
  n := new(Node);  
  n.val := v;  
  success := false;  
  while success = false do {  
    tmp := top;  
    /* other process .. */  
    /* .. may change top! */  
    n.next := tmp;  
    CAS(tmp, n, top)}  
}
```

```
CAS(tmp, n, top; success) { /* atomic ! */  
  if* top = tmp  
  then { ; }  
  else success := false; }
```



Lock-Free Algorithms and their Use

- Principle of lock-free algorithms:
 - read old data structure
 - prepare modified version
 - update with CAS. Retry on failure
- Treiber's Stack is one of the simplest algorithms (inefficient for high loads; better: [Hendler et. al 04])
- Lock-Free Algorithms exist for many data structures: Queues [Michael, Scott 96], Hashtables [Michael 02], [Gao et al 05], Linked Lists [Harris 01], [Heller 05]
- Used for: process queues, indexes of data bases and Web Servers, real-time 3D games, garbage collection
- Java library supports CAS; f implements lock-free data structures

Locks or no locks?

Advantages of using Locks:


- Well understood, uniform principle
⇒ easier to verify than lock-free algorithms
(essentially: verify sequential algorithm)
- Automatic checks for correct use of locks available
- Simple lock-free algorithms are inefficient at high loads:
they waste processor time trying over and over

Disadvantages of using Locks:

- Lock is a bottleneck (pessimistic strategy)
- Deadlocks and priority inversion possible
- What happens when the locking process crashes?

Safety: Linearizability

- Defined in [Herlihy & Wing 90]
- Scenario: Several processes (p, q, r), all running algorithm COp in parallel (e.g. CPush \vee CPop)
- Informal definition: Parallel run must be equivalent to a sequential run of AOp (APush \vee APop)



Decomposition of Linearizability

Theorem (Bäumler et al. 09)

If for all $1 \leq p, q \leq n, p \neq q$:

$$(1) \quad COp_p \rightarrow R_p \xrightarrow{+} G_p$$

$$(2) \quad G_p \rightarrow R_q, \text{ reflexive}(G_p), \text{ transitive}(R_p), R \rightarrow R_p$$

$$(3) \quad COp_p(CS) \wedge \square (R_p \wedge \text{Abs}(CS) = AS \wedge \text{Abs}(CS') = AS') \\ \rightarrow \text{skip}^*; AOp_p(AS); \text{skip}^*$$

then $COp_1^* \parallel \dots \parallel COp_n^* \wedge \square R \rightarrow AOp_1^* \parallel \dots \parallel AOp_n^* \parallel \text{skip}^*$

- COp_p is a concrete algorithm (procedure) that implements an atomic operation AOp_p
- R is the global environment assumption
- Linearizability expressed as special case of refinement
- Most linearizable algorithms allow reduction to two representative processes \Rightarrow reduction proved

Liveness: Lock-Freedom

For Treiber's Stack:

- CPush may have to retry over and over
⇒ one single process might be starved
- Every time a retry is necessary, another CPush/CPop must have succeeded and terminated
- This is true, even if the scheduling is unfair, or when a process crashes

Treiber's stack satisfies property of Lock-Freedom:

*As long as **some** operations are running, **one of them** will terminate*

Decomposition of Lock-Freedom

Theorem (Tofan et al. 10)

If for all $0 \leq p, q, p \neq q$:

$$(1) \quad COP_p \rightarrow R_p \xrightarrow{+} G_p$$

$$(2) \quad G_p \rightarrow R_q, \text{ reflexive}(G_p), \text{ transitive}(R_p), R \rightarrow R_p$$

$$(3) \quad \text{reflexive}(U), \text{ transitive}(U), R \rightarrow R_p \wedge U$$

$$(4) \quad COP_p(CS) \wedge \square R_p$$

$$\rightarrow \square (\neg U(CS, CS') \vee (\square U(CS', CS''))) \rightarrow \diamond \mathbf{last}$$

then $COP_0^* \parallel \dots \parallel COP_n^* \wedge \square R \rightarrow \square \text{progress}$

where *progress* = “some operation active \rightarrow some operation terminates”

- Predicate U (“unchanged”) describes conditions under which $COP_p(CS)$ terminates in environment R_p .
- At any time, COP_p eventually terminates ($\diamond \mathbf{last}$), if:
 - It updates the shared state itself $\neg U(CS, CS')$, or
 - It encounters no interference $\square U(CS', CS'')$
- Theorem holds for weak fair and nonfair interleaving

General Experience with the Calculus

- Symbolic execution is natural to verify even concurrent programs:
 - rest of the program directly visible
 - feels much like debugging
- Main new difficulty for proofs is to determine the correct Relys and Guarantees (similar to invariants) in advance
⇒ Add techniques to automatically infer them
- We've done some significant case studies already:
 - Hazard pointers for lock-free algorithms [⇒ tomorrow]
 - Medical protocols with synchronous parallel hierarchical plans [Procure 06]
- Calculus is not yet as easy to use or automated as the wp-calculus for sequential programs (takes time and experience)

Some Open Issues

- Guarantees often hold in a certain section of the code:
currently boolean variables must be added manually
 \Rightarrow labels would be helpful, but are incompatible with chop:
 $\alpha; \{L : \beta\}$: final state of α and first of β disagree on L
- express general refinement modulo stuttering
- Prove general commuting diagrams for forward and backward simulation (bounded nondeterminism!)
- Completeness in general is open
(complete fragments of ITL and RG)

Proving Lock-Free Algorithms

- Calculus is adequate to show correctness of proof obligations (POs) as well as proving instances of the POs for case studies
- Automation is not as high as in related work:
Automatic checking of linearizability for short operations sequences with model checking [Alur10]
Automatic proofs for some algorithms using RGSep [Vafeiadis01]
- Nevertheless, the algorithms we check are already more difficult than those that have been proved automatically

Current Work on Lock-Free Algorithms

- Support for Heap modularity is often beneficial
⇒ develop library with a lightweight embedding of separation
- Open issue: good frame rule for temporal logic?
- Generalize proof obligations (POs) for linearizability
(POs shown require lin. points within executing thread
⇒ complete POs for arbitrary lin. points

Major Challenge:

Interleaving assumes sequentially consistent memory, but:
Processors use weak memory models
(and Javas much debated memory model is even weaker)