

Egon Börger (Pisa) & Alexander Raschke (Ulm)

Modeling Context-Aware Behavior by Ambient ASMs

Thread Behavior and Thread Management Case Study

Università di Pisa, Dipartimento di Informatica, boerger@di.unipi.it
Universität Ulm, Abteilung Informatik, alexander.raschke@uni-ulm.de

See Ch. 4.2 of Modeling Companion
<http://modelingbook.informatik.uni-ulm.de>

Theme: Modeling context-aware system behavior

Question: How to gently model contexts and programs whose run behavior depends on the context in which the program is executed?

- Often contexts are called environments, not to be confused with *env* understood as interpretation of free variables

In this lecture we use the following 2 examples:

- *Encapsulation of process runs* to separate concurrent process executions, e.g.
 - runs of threads in Java
 - computations of instances of processes which execute similar or even the same program but with different data
- *Encapsulation of state* by scoping disciplines

Idea: *use parameterization* to model context dependency.

The role of parameterization for ASMs

Parameterizing ASM functions and rules permits to directly model:

- partitioning/isolating of states and distributed computations
 - of agents which concurrently execute in heterogeneous contexts
- various forms of information hiding (encapsulation of memory)
 - statically: scope, module, package, library, etc.
 - dynamically: executing agents, threads, process instances, etc.
- patterns of programming, of communication and of control flow
- mobility (environments where agents can move)
- modularity (of specs and property verifications)

Exploit *simplicity of semantical foundation* of parameterization:

$$f(x) = f_{params}(x)$$

in particular when used with implicit (hidden) parameters, supporting conventional implicit parameterization *this*. $f(x) = f(x)$

Ambient ASMs to explicitly support env-sensitive actions

Idea: enrich ASMs by an abstract ambient parameter

- with respect to which the terms involved in a step are evaluated
- for ambients (environments) at whatever level of abstraction
- which can be created, modified, deleted, also at run time

Syntactical construct: **amb** *exp* **in** *P*

where to achieve generality in the widest terms

- *exp* is any expression (term)
- *P* is any (already defined) ASM program

Intended behavior (see definition below):

- PUSH the $eval(exp, S, env, amb)$ of the newly declared ambient expression to the current ambient
- execute *P* in the new ambient

Function classification is extended by ambient (in)dependent fcts.

Ambient dependent functions and their evaluation

AmbDependent(f) (wrt Amb) **iff**

forsome $a, a' \in Amb$ **with** $a \neq a'$ **forsome** x $f(a, x) \neq f(a', x)$

Otherwise f is called *AmbIndependent*. We also say *environment (in)dependent*, hoping that *Ambients* will not be confused with *Environments* (interpretation of free variables in a state).

■ Case *AmbDependent*(f):

$$\text{eval}(f(t_1, \dots, t_n), S, \text{env}, \text{amb}) = f_S(\text{amb}, \text{eval}(t_1, S, \text{env}, \text{amb}), \dots, \text{eval}(t_n, S, \text{env}, \text{amb}))$$

Here f_S is turned into a family of possibly different functions $f_{S, \text{amb}}$.

■ Case *AmbIndependent*(f) (unchanged interpretation of f):

$$\text{eval}(f(t_1, \dots, t_n), S, \text{env}, \text{amb}) = f_S(\text{eval}(t_1, S, \text{env}, \text{amb}), \dots, \text{eval}(t_n, S, \text{env}, \text{amb}))$$

Semantics of ambient ASMs

To avoid a signature blow up by dynamic ambient nesting, we adopt Simone Zenzaro's idea (PhD Thesis, Pisa 2016) to

treat amb as a stack

where new ambient expressions are pushed (passed by value)

- so that for each $f^{(n)}$ one extension $f^{(n+1)}$ suffices which offers an additional argument position for the interpretation of $f^{(n)}$ in a given ambient.

The body of **amb** exp **in** P is then executed with the new stack value.

Yields(**amb** exp **in** P , S , env , amb , U) **if**

$Yields(P, S, env, \text{PUSH}(eval(exp, S, env, amb), amb), U)$

- NB. Often the execution of P (read: the interpretation $f_{S, amb}$) depends only on the top of the stack, i.e. on $eval(exp, S, env, amb)$ (*flat ambient ASMs*, see JCSS 2012).

Special case: flat ambient ASMs

Only the last declared (most recent) ambient is kept

- instead of an ambient stack (of nested ambients) to which each newly declared environment is PUSHED
 - This was the original definition of ambient ASMs in JCSS 2012, which is generalized by the stack interpretation of ambients.

For flat ambients it suffices to bind the value of the *expression* in the ambient declaration

- which is computed in the current state (in the current ambient)

to a logical variable, say *curamb*:

$$(\mathbf{amb}_{\text{flat}} \text{ exp in } P)^* = (\mathbf{let} \text{ curamb} = \text{exp}^* \mathbf{in} P^*)$$

- where exp^* is obtained from *exp* by replacing the names *f* of ambient dependent functions in *exp* by f_{curamb}

Inductive transformation of flat ambient ASMs

■ *transformation of terms*:

– for $AmbIndependent(f)$ define:

$$(f(t_1, \dots, t_n))^* = f(t_1^*, \dots, t_n^*)$$

- logical variables and names of parameterized rules are classified as ambient independent

– for $AmbDependent(f)$ define:

$$(f(t_1, \dots, t_n))^* = f(\text{curamb}, t_1^*, \dots, t_n^*)$$

■ *transformation of ambient ASM rules* (to standard ASM programs):

– for update rules define:

$$(f(t_1, \dots, t_n) = t)^* = (f(t_1, \dots, t_n)^* := t^*)$$

– for ambient rules define (eliminating **amb_{flat}**):

$$(\mathbf{amb}_{\mathbf{flat}} \text{ exp in } P)^* = (\mathbf{let} \text{ curamb} = \text{exp}^* \mathbf{in} P^*)$$

– for the other rules use induction on ASM programs, e.g.

$$(\mathbf{let} x = t \mathbf{in} P)^* = (\mathbf{let} x = t^* \mathbf{in} P^*)$$

Encapsulating computations in distributed runs

Goal: *isolate executions of various tasks by different agents*

- Technical issue: separate concurrent (e.g. multi-core) executions and scheduling from task execution by single agents
- In this lecture we illustrate how to use ambient ASMs for the purpose by two concrete examples (though the principles can be applied also to other distributed systems and concurrent request management schemes):
 - using a component which encapsulates single-thread executions of Java code to model concurrent runs of multiple Java threads
 - modeling Java thread pool management (for J2SE 5.0), using a component which encapsulates to `RUN` a *thread* to `EXECUTE` a *task*

From single-thread to multi-thread Java interpreter

Goal: Define a concurrent `MULTITHREADJAVAINTERPRETER` model which separates the thread instances of the underlying runs of a `SINGLETHREADJAVAINTERPRETER`.

A `SINGLETHREADJAVAINTERPRETER` ASM has been defined in:

- R. Stärk and J. Schmid and E. Börger, *Java and the JVM*. Definition, Verification, Validation. Springer 2001. (Called JBook)

It has been used there (Ch. 7) to define a multi-thread Java interpreter model to execute multiple tasks with shared main and local working memory. It adopts an abstract scheduling mechanism which

- selects each time one *Runnable* thread, out of the current *Thread* class instances in the *heap*, to `RUN` it
- makes a *Synchronizing* or *Notified* thread *Active* before `RUN`ning it

Using ambient ASMs saves the explicit restoring/saving of the current state of a thread when the thread is scheduled to `RUN` or `descheduled`.

Encapsulating single-thread Java code executions

MULTITHREADJAVAINTERPRETER =

if $Runnable(t)$ **then**

ACQUIRELOCKS(t) **seq** **RUN**(t)

where

RUN(t) = **amb_{flat}** t **in** **SINGLETHREADJAVAINTERPRETER**

Restoring/saving current thread state can be skipped by declaring:

- $Amb = Thread$ (i.e. implicit parameterization by threads)
- as $AmbDependent$ the thread state functions used in the ASM **SINGLETHREADJAVAINTERPRETER**, namely:
 - $meth, restbody, pos, locals$ constituting the current frame
 - $frames$ denoting the frame stack
 - $thread$ denoting the executing thread

Mono-core vs multi-core model

- The JBook *execJavaThread* ASM chooses one $t \in \text{Runnable}$ and to $\text{RUN}(t)$ assigns it as the currently executing *thread* $:= t$
 - In JBook *thread* is the agent which executes the `SINGLETHREADJAVAINTERPRETER` (single-core view)
- The `MULTITHREADJAVAINTERPRETER` ASM as defined here triggers a *concurrent ASM run of all threads which are RUNNING*.
 - NB. $\text{RUN}(t)$ can be refined by assigning this execution to a specific computer or core.
 - NB. *Thread* and *Runnable* are dynamic sets.

NB. *Ambient separation supports modular verifications*:

- see ASM-based analysis of C# thread model (LNCS 3052, TCS 343)
- see proofs for conservative theory extensions corresponding to incremental model extensions in D. Batory/E. Börger: Modularizing Theorems for Software Product Lines: The Jbook Case Study. J.UCS 2008

Submachines of MULTITHREADJAVAINTERPRETER

Pro memoria definitions from the JBook:

Runnable(*t*) **iff**

mode(*t*) = *active*

or (*mode*(*t*) = *synchronizing* **and** *locks*(*syncObj*(*t*)) = 0)

or (*mode*(*t*) = *waiting* **and** *locks*(*waitObj*(*t*)) = 0)

ACQUIRELOCKS(*t*) =

if *mode*(*t*) = *synchronizing*(*t*) **then** SYNCHRONIZE(*t*)

if *mode*(*t*) = *notified*(*t*) **then** WAKEUP(*t*)

Active(*q*) := *true*

SYNCHRONIZE(*t*) = Refresh *sync*(*t*) by *syncObj*(*t*)

WAKEUP(*t*) = Reacquire all sync claims on *waitObj*(*t*)

Isolating application-logic task exec from task management

Goal: separate thread management

- creation, deletion and scheduling of threads t to concurrently run tasks
 - assign t to task, decouple t from task, suspend t

and its specification from the execution and application-logic-level specification of tasks

Ex1: J2SE 5.0 (see S. Oaks and H. Wong: *Java Threads*, O'Reilly 2004)

NB. Principles can be applied also to other concurrent request management systems, e.g. web servers

Thread management actions

- assignment of threads to tasks upon TaskEntry
- decoupling of threads from tasks upon TaskCompletion
- creation of threads
- suspension of threads
 - making them idle to possibly RunTaskFromQueue
- deletion of threads
 - if one cannot any more RunTaskFromQueue so that the thread has to Exit

J2SE 5.0 thread pool requirements (1)

corePoolSizeReq. The thread pool should be kept as much as possible within *corePoolSize*. When a new task is submitted and fewer than *corePoolSize* threads are running, a new thread is created to handle the request, even if there are idle threads. If when the task is submitted the pool has reached or exceeds the *corePoolSize* but not the *maxPoolSize* and there are idle threads in the pool, one of them is assigned to run the task.

maxPoolSizeReq. If a task is submitted for execution when the pool is full and all threads in the pool are running, the task is inserted into a *queue*. If the *queue* is full the task is rejected.

QueuePriorityReq. If a task is submitted for execution when *corePoolSize* or more threads are running but the pool is not yet full and no idle thread is available, then a new thread is created and assigned to run the task only if the task cannot be placed to the *queue* without blocking it.

J2SE 5.0 thread pool requirements (2)

RunCompletionReq. If when a thread has completed its current run there are tasks in the queue, the thread is assigned to run one of them. Otherwise it becomes idle.

IdleThreadReuseReq. When there is a task in the queue, an idle thread, if there is one, is assigned to run the task.

IdleThreadExitReq If there are more than *corePoolSize* threads in the pool, excess threads will be terminated if they have been idle for more than the *keepAliveTime*.

THREADPOOLMNGR

- first three requirements ask to `HANDLENEWTASKS`
- last three concern idle threads: how to `HANDLEQUEUEDTASKS`

In both an idle thread may be assigned to run the task in question, but *only one task per thread*:

- abstract from scheduling by letting the `THREADPOOLMNGR` choose in each step one of the two components for execution

`THREADPOOLMNGR` =

one of (`{HANDLENEWTASK, HANDLEQUEUEDTASK}`)

where

one of (*Rules*) =

choose $R \in Rules$ **do** R

HANDLENEWTASK(*task*) component

```
HANDLENEWTASK(task) = if Submitted(task) then  
  if | Pool | < corePoolSize -- first create corePoolSize threads  
    then let t = new (Pool) in RUN(t, task)  
  else if | Pool | < maxPoolSize then -- first try Idle threads  
    choose t ∈ Pool with Idle(t) RUN(t, task)  
    if none if BlockingFreePlaceable(task, queue)  
      then INSERT(task, queue) -- first fill queue  
      else let t = new (Pool) in RUN(t, task)  
    else if forall t ∈ Pool Running(t) then  
      if | queue | < maxQueuesize then INSERT(task, queue)  
      else REJECT(task)  
where RUN(thread, task) = (mode(thread) := running par  
  ambflat task in EXECUTE(pgm(task)))
```

HANDLEQUEUEDTASK component

The *RunCompletionRequirement* seems to establish a reuse priority of *Terminated* threads over idle ones. Thus we try:

- first to FINDTASKFORTERMINATED threads satisfying the RunCompletionReq
- then to FINDTASKFORIDLE threads satisfying the two idle thread requirements

HANDLEQUEUEDTASK =

choose $thread \in Pool$ **with** $Terminated(thread)$

FINDTASKFORTERMINATED($thread$)

if none

choose $thread \in Pool$ **with** $Idle(thread)$

FINDTASKFORIDLE($thread$)

$Terminated(t)$ **iff** $mode(t) = idle$ expresses that a *thread* has completed the run for the assigned task.

FINDTASKFORTERMINATED(*thread*) **submachines**

NB. If there is no task in the queue, a terminated *thread* becomes idle.

FINDTASKFORTERMINATED(*thread*) =

choose $task \in queue$

RUN(*thread*, *task*)

DELETE(*task*, *queue*)

if none

MAKEIDLE(*thread*)

where

MAKEIDLE(*thread*) =

$mode(thread) := idle$

$terminationTime(thread) := now$

-- set timer for $keepAliveTime(thread)$

FINDTASKFORIDLE(*thread*) **submachine**

NB. An idle *thread* may be killed if the pool grew over the *corePoolSize* and the *keepAliveTime(thread)* has expired

FINDTASKFORIDLE(*thread*) =

choose $task \in queue$

RUN(*thread*, *task*)

DELETE(*task*, *queue*)

if none **TRYTOKILL**(*thread*)

where

TRYTOKILL(*t*) =

if $| Pool | > corePoolSize$ **then**

if *Expired*(*aliveTime*(*t*)) **then** **DELETE**(*t*, *Pool*)

Expired(*aliveTime*(*t*)) **iff**

$now - terminationTime(t) > keepAliveTime(t)$

Encapsulating state: scoping disciplined evaluation scheme

Value of id in a $state$ depends on

- the $position$ in the program where id occurs
- an env among those which have a binding for id (i.e. a defined associated value) and whose $scope$ (program part where its bindings are valid) includes pos

$ENVSENSITIVEVAL(id, pos, state) =$

if $Occurs(id, pos, state)$ **then**

let $E = Env \cap \{e \mid inScope(pos, e, state) \wedge HasBinding(e, id)\}$

let $env = select(E)$ -- stands for **choose** $env \in E$

amb env **in** $GETVALUE(id)$

$GETVALUE$ retrieves the value of an id given an env .

Variations: refining $select$ function and $inScope$ predicate

- lexical scoping (Pascal, C), dynamic scoping (Logo), combination of lexical and dynamic scoping (Java)

Encapsulating state: scoping discipline instances

Using *select* to shadow bindings in case *id* is bound in multiple envs all of which are valid at *pos*

Example: stack-based last-in/first-out (LIFO) scope selection

$\text{ENVSENSITIVELIFOEVAL}(id, pos, state) =$
 $\text{ENVSENSITIVEEVAL}(id, pos, state)$

where $select(E)$ is constrained to yield an *env* satisfying

forall $e \in E$ $env \sqsubseteq e$ -- language defined \sqsubseteq

Encapsulating state: TCL scoping discipline

TCL: allows to select any binding of names established in any enclosing scope (possibly hidden by a nearer scope)

TCL*EVAL*(*id*, *pos*, *state*) = **if** *Occurs*(*id*, *pos*, *state*) **then**
let *e* = *selectenv*(*pos*, *state*) **amb** *e* **in** *GETVALUE*(*id*)

where let *n* = *length*(*Env*)

$$\textit{selectenv}(\textit{pos}, \textit{s}) = \begin{cases} \textit{Env}[0] & \text{if } \textit{pos} \text{ is in global } \textit{id} \\ \textit{Env}[\textit{n} - \textit{k}] & \text{if } \textit{pos} \text{ is in upvar } \textit{k} \textit{id} \textit{v} \\ \textit{Env}[\textit{k}] & \text{if } \textit{pos} \text{ is in upvar } \#\textit{k} \textit{id} \textit{v} \\ \textit{Env}[\textit{n}] & \text{otherwise // current scope} \end{cases}$$

upvar *k id v* binds *v* to *id* as bound in the *k*-th scope "up" from the current scope

dto upvar #*k id v* with *k*-th scope "down" from the global one

References

On ambient ASMs:

- E. Börger and A. Raschke: Modeling Companion for Software Practitioners. Springer 2018

<http://modelingbook.informatik.uni-ulm.de>

– Ch.4 contains other applications of ambient ASMs for contex-aware system models and further references

- E. Börger and and A. Cisternino and V. Gervasi: Ambient Abstract State Machines with Applications.
J. Computer and System Sciences 78.3 (2012) 939-959.

On multi-threaded Java and Java thread management:

- R. Stärk and J. Schmid and E. Börger, Java and the JVM. Definition, Verification, Validation. Springer 2000
- S. Oaks and H. Wong: Java Threads. O'Reilly 2004

References on modular property verifications

- R. F. Stärk and E. Börger: An ASM Specification of C# Threads and the .NET Memory Model. LNCS 3052 (2004) 38-60
- R. F. Stärk: Formal specification and verification of the C# thread model. Theoretical Computer Science 343 (2005) 482–508

- D. Batory and E. Börger: Modularizing Theorems for Software Product Lines: The Jbook Case Study. J.UCS 2008

Copyright Notice

It is permitted to (re-) use these slides under the CC-BY-NC-SA licence

<https://creativecommons.org/licenses/by-nc-sa/4.0/>

i.e. in particular under the condition that

- the original authors are mentioned
- modified slides are made available under the same licence
- the (re-) use is not commercial