

Isabelle Formalization of Hedge-Constrained pre^* and DPNs with Locks

Peter Lammich

January 30, 2009

Abstract

Dynamic Pushdown Networks (DPNs) are a model for concurrent programs with recursive procedures and thread creation. We formalize a true-concurrency semantics for DPNs. Executions of this semantics have a tree structure. We show the relation of our semantics to the original interleavings semantics. We then show how to compute predecessor sets of regular sets of configurations w.r.t. tree-regular constraints on the execution.

Acquisition histories have been introduced by Kahlon et al. to model-check parallel pushdown systems with well-nested locks, but without thread creation. We generalize acquisition histories to be used with DPNs. For this purpose, our tree-based semantics can be naturally applied. Moreover, the generalized acquisition histories enable us to characterize the (tree-based) executions that have a schedule that is valid w.r.t. locks, thus obtaining an algorithm to compute lock-sensitive predecessor sets.

Contents

1	Introduction	3
2	Labeled transition systems	5
2.1	Definitions	6
2.2	Basic properties of transitive reflexive closure	6
2.2.1	Appending of elements to paths	7
2.2.2	Transitivity reasoning setup	7
2.2.3	Monotonicity	7
2.2.4	Special lemmas for reasoning about states that are pairs	8
2.2.5	Invariants	8
3	Dynamic Pushdown Networks	8
3.1	Model Definition	8

4	Semantics	9
4.1	Interleaving Semantics	9
4.2	Tree Semantics	10
4.2.1	Scheduler	12
5	Predecessor Sets	15
5.1	Hedge-Constrained Predecessor Sets	15
6	DPN Semantics on Lists	17
6.1	Definitions	17
6.2	Theorems	18
6.2.1	Representation of Single Processes	18
6.2.2	Representation of Configurations	18
6.2.3	Step Relation on List-Configurations	21
6.3	Predecessor Sets on List-Semantics	22
7	Automata for Execution Hedges	22
8	Computation of Hedge-Constrained Predecessor Sets	24
8.1	Correctness of Reduction	25
8.2	Effectiveness of Reduction	27
8.2.1	Definitions	27
8.2.2	Theorems	28
8.3	What Does This Proof Tell You ?	30
9	DPNs With Locks	30
9.1	Model	30
9.2	Interleaving Semantics	31
9.3	Tree Semantics	32
9.4	Equivalence of Interleaving and Tree Semantics	32
10	Well-Nestedness of Locks	32
10.1	Well-Nestedness Condition on Paths	33
10.2	Well-Nestedness of Configurations	35
10.2.1	Auxilliary Lemmas about <i>wn-c</i>	35
10.3	Well-Nestedness Condition on Trees	37
10.4	Well-Nestedness of Hedges	38
10.4.1	Auxilliary Lemmas about <i>wn-h</i>	38
10.4.2	Relation to Path Condition	40
10.5	Well-Nestedness and Tree Scheduling	41
11	Acquisition Structures	43
11.1	Utilities	43
11.1.1	Combinators for <i>option-datatype</i>	43
11.2	Acquisition Structures	44

11.2.1	Parallel Composition	44
11.2.2	Acquisition Structures of Scheduling Trees and Hedges	45
11.3	Consistency of Acquisition Structures	47
11.3.1	Minimal Elements	50
11.3.2	Well-Nestedness and Acquisition Structures	51
11.4	Soundness of the Consistency Condition	52
11.5	Precision of the Consistency Condition	52
11.5.1	Custom Size Function	52
12	DPNs with Initial Configuration	58
12.1	DPNs with Initial Configuration	59
12.1.1	Reachable Configurations	59
13	Property Specifications	59
13.1	Specification Formulas	60
13.2	Semantics	60
13.3	Examples	60
13.3.1	Conflict analysis	61
13.3.2	Bitvector analysis	61
14	Hedge Constraints for Acquisition Histories	62
14.1	Locks Encoded in Control State	62
14.2	Characterizing Schedulable Execution Hedges	64
14.3	Checking Specifications Using <i>prehc</i> Δ <i>Hls</i>	66
15	Monitors (aka Block-Structured Locks)	66
15.1	Non-Trivial Instance of a Well-Nested DPN	69
16	Conclusion	71
16.1	Trusted Code Base	71

1 Introduction

Writing parallel programs has become popular in the last decade. However, writing correct parallel programs is notoriously difficult, as there are many possibilities for concurrency related bugs. These are hard to find and hard to reproduce due to the nondeterministic behaviour of the scheduler. Hence there is a strong need for formal methods to verify parallel programs and help find concurrency related bugs. A formal model for parallel programs, that has been studied in the last few years, are dynamic pushdown networks (DPNs) [2], a generalization of pushdown systems, where a rule may have the additional side effect of creating a new process, that is then executed in parallel. Analysis of DPNs is usually done w.r.t. to an interleaving semantics, where an execution is a sequence of rule applications. The interleaving

semantics models the execution on a single processor, that performs one step at a time and may switch the currently executing process after every step. However, these interleaved executions do not have nice language theoretic properties, what makes them difficult to reason about. For example, it is undecidable whether there exists an execution with a given regular property. Moreover, executions of the interleaving semantics are not suited to track properties of specific processes, e.g. acquired locks.

In the first part of this formalization, we define a semantics that models an execution as a partially ordered set of steps, rather than a (totally ordered) sequence of steps. This partial ordering only reflects the ordering between steps of the same process and the causality due to process creation, i.e. steps of a created process must be executed after the step that created the process. However, it does not enforce any ordering between steps of processes running in parallel. The interleaved executions can be interpreted as topological sorts of the partial ordering. For executions of DPNs the partial ordering has a tree shape, where thread creation steps have at most two successors and pushdown steps have at most one successor. We formally define these executions as list of trees (called execution hedges).

The key concept of model-checking DPNs is to compute the set of predecessor configurations of a set of configurations. Configurations of DPNs are represented as words over control- and stack- symbols, and for a regular set of configurations, the set of predecessor configurations is regular as well and can be computed efficiently [2]. Predecessor computations can be used for various interesting analysis, like kill/gen analysis on bitvectors [2] and context-bounded model checking [1]. Our approach extends the predecessor computation by additionally allowing tree-regular constraints on the executions. The counterpart for the interleaving semantics, i.e. predecessor computations with (word-)regular constraints on the interleaved executions, is not effective.

In the second part of this formalization, we extend DPNs by adding mutual exclusion via well-nested locks. Locks are a commonly used synchronization primitive to manage shared resources between processes. A process may acquire and release a lock, and, at any time, each lock may be owned by at most one process. If a process wants to acquire a lock already owned by another process, it has to wait until the lock is released. We assume that locks are used in a well-nested fashion, i.e. a process has to release locks in the reversed order of acquisition. Note that in practice locks are commonly used in a well-nested fashion, e.g. the synchronized-blocks of Java guarantee well-nested lock usage. Also note that for non-well-nested locks, even simple reachability problems are undecidable [4]. Parallel pushdown processes with well-nested locks have been analyzed using acquisition histories [4, 3]. We generalize this technique to DPNs. Our generalization is non-trivial, as the original technique is defined for a model where only two parallel processes that both exist at the beginning of the execution need to

be considered, while we have a model with unboundedly many processes that may be created at any point of the execution. The generalized acquisition histories allow us to characterize the executions, that are consistent w.r.t. lock usage, by a tree-regular set. Applying the results from the first part of this paper yields an algorithm for computing lock-sensitive predecessor sets with tree-regular constraints.

This formalization accompanies a paper that is currently in preparation. Thus the proofs in this work partially depend on unpublished results that are currently in the process of submission. The following are the most notable results proven in this formalization:

- We present a tree-based view on DPN executions, and an efficient predecessor computation with tree-regular constraints.
- We generalize the concept of acquisition histories to programs with process creation.
- We characterize lock-sensitive executions by tree-regular constraints, thus obtaining an algorithm for computing lock-sensitive predecessor sets.

However, this formalization also has its limits. In particular, it does not include:

- A formalization of operations on automata or tree automata, that would allow to generate executable code.
- A formalization of the saturation algorithm for computing predecessor sets of DPNs [2] — another prerequisite for generating executable code. We have an unpublished formalization of this saturation algorithm, that we will adapt to the latest version of Isabelle and publish in near future.
- Due to the first two limitations, we cannot give a formal proof that shows that our methods are, indeed, executable. However, we prove some lemmas that give strong evidence that our methods are effective and could be implemented in principle.

2 Labeled transition systems

```
theory LTS  
imports Main  
begin
```

Labeled transition systems (LTS) provide a model of a state transition system with named transitions.

2.1 Definitions

An LTS is modeled as a ternary relation between start configuration, transition label and end configuration

types (c, a) LTS = $(c \times a \times c)$ set

Transitive reflexive closure

inductive-set

$trcl :: (c, a)$ LTS \Rightarrow (c, a) list LTS

for t

where

$empty[simp]: (c, [], c) \in trcl\ t$

| $cons[simp]: \llbracket (c, a, c') \in t; (c', w, c'') \in trcl\ t \rrbracket \Longrightarrow (c, a\#w, c'') \in trcl\ t$

2.2 Basic properties of transitive reflexive closure

lemma $trcl\ empty\ cons: (c, [], c') \in trcl\ t \Longrightarrow (c = c')$

$\langle proof \rangle$

lemma $trcl\ empty\ simp[simp]: (c, [], c') \in trcl\ t = (c = c')$

$\langle proof \rangle$

lemma $trcl\ single[simp]: ((c, [a], c') \in trcl\ t) = ((c, a, c') \in t)$

$\langle proof \rangle$

lemma $trcl\ uncons: (c, a\#w, c') \in trcl\ t \Longrightarrow \exists ch. (c, a, ch) \in t \wedge (ch, w, c') \in trcl\ t$

$\langle proof \rangle$

lemma $trcl\ uncons\ cases: \llbracket$

$(c, e\#w, c') \in trcl\ S;$

$!!ch. \llbracket (c, e, ch) \in S; (ch, w, c') \in trcl\ S \rrbracket \Longrightarrow P$

$\rrbracket \Longrightarrow P$

$\langle proof \rangle$

lemma $trcl\ one\ elem: (c, e, c') \in t \Longrightarrow (c, [e], c') \in trcl\ t$

$\langle proof \rangle$

lemma $trcl\ unconsE[cases\ set, case\ names\ split]: \llbracket$

$(c, e\#w, c') \in trcl\ S;$

$!!ch. \llbracket (c, e, ch) \in S; (ch, w, c') \in trcl\ S \rrbracket \Longrightarrow P$

$\rrbracket \Longrightarrow P$

$\langle proof \rangle$

lemma $trcl\ pair\ unconsE[cases\ set, case\ names\ split]: \llbracket$

$((s, c), e\#w, (s', c')) \in trcl\ S;$

$!!sh\ ch. \llbracket ((s, c), e, (sh, ch)) \in S; ((sh, ch), w, (s', c')) \in trcl\ S \rrbracket \Longrightarrow P$

$\rrbracket \Longrightarrow P$

$\langle proof \rangle$

lemma $trcl\ concat: !! c. \llbracket (c, w1, c') \in trcl\ t; (c', w2, c'') \in trcl\ t \rrbracket$

$\Longrightarrow (c, w1@w2, c'') \in trcl\ t$

$\langle proof \rangle$

lemma $trcl\ unconcat: !! c. (c, w1@w2, c') \in trcl\ t$

$\implies \exists ch. (c, w1, ch) \in trcl\ t \wedge (ch, w2, c') \in trcl\ t$
 <proof>

2.2.1 Appending of elements to paths

lemma *trcl-rev-cons*: $\llbracket (c, w, ch) \in trcl\ T; (ch, e, c') \in T \rrbracket \implies (c, w @ [e], c') \in trcl\ T$
 <proof>

lemma *trcl-rev-uncons*: $(c, w @ [e], c') \in trcl\ T$
 $\implies \exists ch. (c, w, ch) \in trcl\ T \wedge (ch, e, c') \in T$
 <proof>

lemma *trcl-rev-uncons-cases*: \llbracket
 $(c, w @ [e], c') \in trcl\ T;$
 $!!ch. \llbracket (c, w, ch) \in trcl\ T; (ch, e, c') \in T \rrbracket \implies P$
 $\rrbracket \implies P$
 <proof>

lemma *trcl-rev-induct*[*induct set, consumes 1, case-names empty snoc*]: $!!c'. \llbracket$
 $(c, w, c') \in trcl\ S;$
 $!!c. P\ c \llbracket c;$
 $!!c\ w\ c'\ e\ c''. \llbracket (c, w, c') \in trcl\ S; (c', e, c'') \in S; P\ c\ w\ c' \rrbracket \implies P\ c\ (w @ [e])\ c''$
 $\rrbracket \implies P\ c\ w\ c'$
 <proof>

lemma *trcl-rev-cases*: $!!c\ c'. \llbracket$
 $(c, w, c') \in trcl\ S;$
 $\llbracket w = []; c = c' \rrbracket \implies P;$
 $!!ch\ e\ wh. \llbracket w = wh @ [e]; (c, wh, ch) \in trcl\ S; (ch, e, c') \in S \rrbracket \implies P$
 $\rrbracket \implies P$
 <proof>

lemma *trcl-cons2*: $\llbracket (c, e, ch) \in T; (ch, f, c') \in T \rrbracket \implies (c, [e, f], c') \in trcl\ T$
 <proof>

2.2.2 Transitivity reasoning setup

declare *trcl-cons2*[*trans*] — It's important that this is declared before *trcl-concat*,
 because we want *trcl-concat* to be tried first by the transitivity reasoner

declare *cons*[*trans*]

declare *trcl-concat*[*trans*]

declare *trcl-rev-cons*[*trans*]

2.2.3 Monotonicity

lemma *trcl-mono*: $!!A\ B. A \subseteq B \implies trcl\ A \subseteq trcl\ B$
 <proof>

lemma *trcl-inter-mono*: $x \in trcl\ (S \cap R) \implies x \in trcl\ S \quad x \in trcl\ (S \cap R) \implies x \in trcl\ R$
 <proof>

2.2.4 Special lemmas for reasoning about states that are pairs

lemmas *trcl-pair-induct* = *trcl.induct*[*of* (*xc1,xc2*) *xb* (*xa1,xa2*), *consumes* 1, *split-format* (*complete*), *case-names empty cons*]

lemmas *trcl-rev-pair-induct* = *trcl-rev-induct*[*of* (*xc1,xc2*) *xb* (*xa1,xa2*), *consumes* 1, *split-format* (*complete*), *case-names empty snoc*]

2.2.5 Invariants

lemma *trcl-prop-trans*[*cases set*, *consumes 1*, *case-names empty steps*]: \llbracket
 $(c,w,c') \in \text{trcl } S;$
 $\llbracket c=c'; w=[] \rrbracket \implies P;$
 $\llbracket c \in \text{Domain } S; c' \in \text{Range } (\text{Range } S) \rrbracket \implies P$
 $\rrbracket \implies P$
<proof>

end

3 Dynamic Pushdown Networks

theory *DPN*
imports *Main common/LTS*
begin declare *predicate2I*[*HOL.rule del*, *Pure.rule del*]

3.1 Model Definition

A *Dynamic Pushdown Network* (DPN) [2] is a system of pushdown rules over states from $'Q$ and stack symbols from \mathbb{T} , where each pushdown rule may spawn additional processes. Rules are labeled by elements of type $'L$

datatype ($'P, \mathbb{T}, 'L$) *pushdown-rule* =
NOSPAWN $'P \ \mathbb{T} \ 'L \ 'P \ \mathbb{T} \ \text{list} \ (\ _ \ _ \ _ \ _ \ _ \ 51) \ |$
SPAWN $'P \ \mathbb{T} \ 'L \ 'P \ \mathbb{T} \ \text{list} \ 'P \ \ \mathbb{T} \ \text{list} \ (\ _ \ _ \ _ \ _ \ _ \ \# \ _ \ _ \ 51)$

notation *NOSPAWN* ($_ \ _ \ _ \ _ \ _ \ 51$)

notation *SPAWN* ($_ \ _ \ _ \ _ \ _ \ \# \ _ \ _ \ 51$)

types ($'Q, \mathbb{T}, 'L$) *dpn* = ($'Q, \mathbb{T}, 'L$) *pushdown-rule set*

We fix the finiteness assumption of the set of rules in a locale. Note that we do not assume the base types of states, stack symbols, or labels to be finite. However, the finiteness assumption of the set of rules implies that the sets of *used* control states, stack symbols, and labels are finite.

locale *DPN* =
fixes $\Delta :: ('Q, \mathbb{T}, 'L) \ \text{dpn}$
assumes *ruleset-finite*[*simp*, *intro!*]: *finite* Δ

end

4 Semantics

theory *Semantics*
imports *DPN RegSet-add*
begin

In this theory, we define an interleaving and a tree-based semantics of DPNs. We show the equivalence of the two semantics.

4.1 Interleaving Semantics

The interleaving semantics models the execution of a DPN on a single processor, that makes one step at a time, and may switch the currently executed process after each step. This is the original semantics of DPNs [2].

The interleaving semantics is formalized by means of a labeled transition system. A single process is modeled as a pair of its control state and its stack.

A configuration of the DPN is modeled as a list of processes. Note that we use lists of processes here, rather than multisets, to enable representation of configurations as regular sets, as required by the algorithms of [2].

types

(Q, T) *pconf* = $'Q \times T$ *list*
 (Q, T) *conf* = (Q, T) *pconf list*

The (single-) step relation *dpntr* of the interleavings semantics is defined as the least solution of the following constraints:

inductive-set *dpntr* :: (Q, T, L) *dpn* $\Rightarrow ((Q, T)$ *conf* $\times L \times (Q, T)$ *conf) *set*
for Δ **where***

— A non-spawning step modifies a single pushdown process according to a non-spawning rule in the DPN:

dpntr-no-spawn:

$(p, \gamma \hookrightarrow_1 p', w) \in \Delta \implies$
 $(c1@(p, \gamma \# r) \# c2, l, c1@(p', w @ r) \# c2) \in \text{dpntr } \Delta \mid$

— A spawning step modifies a pushdown process according to a spawning rule in the DPN and adds the spawned process immediately before the spawning process:

dpntr-spawn:

$(p, \gamma \hookrightarrow_1 ps, ws \# p', w) \in \Delta \implies$
 $(c1@(p, \gamma \# r) \# c2, l, c1@(ps, ws) \# (p', w @ r) \# c2) \in \text{dpntr } \Delta$

We denote the reflexive, transitive closure of the single-step relation by *dpntrc*:

abbreviation *dpntrc* $M == \text{trcl } (\text{dpntr } M)$

4.2 Tree Semantics

Now we regard a true concurrency semantics, where an execution does not contain the interleaving between independent steps. When starting at a single process, we model such an execution as a tree, where each node corresponds to an applied step. A node corresponding to a non-spawning step has one successor, a node corresponding to a spawning step has two successors. We annotate the leafs of the tree by the configuration of the reached process.

When starting at a configuration consisting of (a list of) multiple processes, we model the execution as a list of multiple execution trees, one for each process.

datatype (Q, Γ, L) *ex-tree* =
 $NLEAF (Q, \Gamma) pconf \mid$
 $NNOSPAWN L (Q, \Gamma, L) ex-tree \mid$
 $NSPAWN L (Q, \Gamma, L) ex-tree (Q, \Gamma, L) ex-tree$

types (Q, Γ, L) *ex-hedge* = (Q, Γ, L) *ex-tree list*

inductive *tsem*

$:: (Q, \Gamma, L) dpn \Rightarrow (Q, \Gamma) pconf \Rightarrow (Q, \Gamma, L) ex-tree \Rightarrow (Q, \Gamma) conf \Rightarrow bool$

for Δ where

tsem-leaf[*simp*, *intro!*]:

$tsem \Delta pw (NLEAF pw) [pw] \mid$

tsem-nospawn:

$\llbracket (p, \gamma \xrightarrow{l} p', w) \in \Delta; tsem \Delta (p', w @ r) t c' \rrbracket \Longrightarrow$
 $tsem \Delta (p, \gamma \# r) (NNOSPAWN l t) c' \mid$

tsem-spawn:

$\llbracket (p, \gamma \xrightarrow{l} ps, ws \# p', w) \in \Delta; tsem \Delta (ps, ws) ts cs; tsem \Delta (p', w @ r) t c' \rrbracket \Longrightarrow$
 $tsem \Delta (p, \gamma \# r) (NSPAWN l ts t) (cs @ c')$

inductive *hsem*

$:: (Q, \Gamma, L) dpn \Rightarrow (Q, \Gamma) conf \Rightarrow (Q, \Gamma, L) ex-hedge \Rightarrow (Q, \Gamma) conf \Rightarrow bool$

for Δ where

hsem-empty[*simp*, *intro!*]: $hsem \Delta [] [] [] \mid$

hsem-cons: $\llbracket tsem \Delta \pi t cf'; hsem \Delta c h c' \rrbracket \Longrightarrow hsem \Delta (\pi \# c) (t \# h) (cf' @ c')$

In the following we show some basic facts about the *tsem*- and *hsem*-relations.

lemma *hsem-empty-h*[*simp*]:

$hsem \Delta c [] c' \longleftrightarrow c = [] \wedge c' = []$

<proof>

lemma *hsem-length*: $hsem \Delta c h c' \Longrightarrow length\ c = length\ h$

<proof>

The hedges and configurations of the hedge semantics can be concatenated.

lemmas $hsem-cons-single = hsem-cons$ [**where** $cf'=[\pi]$, *simplified, standard*]

lemma $hsem-conc$: $\llbracket hsem \Delta c1 h1 c1'; hsem \Delta c2 h2 c2' \rrbracket \implies$
 $hsem \Delta (c1@c2) (h1@h2) (c1'@c2')$
<proof>

lemmas $hsem-conc-lel = hsem-conc$ [*OF - hsem-cons*]

lemmas $hsem-conc-leel = hsem-conc$ [*OF - hsem-cons*[*OF - hsem-cons*]]

lemma $tsem-not-empty$ [*simp*]: $\neg tsem \Delta \pi t \llbracket$
<proof>

lemma $hsem-empty-simps1$ [*simp*]:
 $hsem \Delta \llbracket h c' \longleftrightarrow (h=[] \wedge c'=[])$
 $hsem \Delta c h \llbracket \longleftrightarrow (c=[] \wedge h=[])$
<proof>

lemma $hsem-id$ [*simp, intro!*]: $hsem \Delta c (map\ NLEAF\ c) c$
<proof>

lemmas $hsem-id'$ [*simp, intro!*] = $hsem-id$ [*of - \pi#c, simplified, standard*]

Given a partition of the starting configuration, we can construct a corresponding partition of the hedge and the final configuration.

lemma $hsem-split'$:
 $\llbracket hsem \Delta (c1@c2) h c' \rrbracket \implies \exists h1 h2 c1' c2'.$
 $h=h1@h2 \wedge c'=c1'@c2' \wedge$
 $hsem \Delta c1 h1 c1' \wedge hsem \Delta c2 h2 c2'$
<proof>

lemma $hsem-split$ [*consumes 1*]: $\llbracket hsem \Delta (c1@c2) h c';$
 $!!h1 h2 c1' c2'.$
 $\llbracket h=h1@h2; c'=c1'@c2'; hsem \Delta c1 h1 c1'; hsem \Delta c2 h2 c2' \rrbracket \implies P$
 $\rrbracket \implies P$
<proof>

lemma $hsem-single$:
 $\llbracket hsem \Delta [\pi] h c'; !!t. \llbracket h=[t]; tsem \Delta \pi t c' \rrbracket \implies P \rrbracket \implies P$
<proof>

lemma $hsem-split-single$ [*consumes 1*]: $\llbracket hsem \Delta (\pi#c2) h c';$
 $!!t1 h2 c1' c2'.$
 $\llbracket h=t1\#h2; c'=c1'@c2'; tsem \Delta \pi t1 c1'; hsem \Delta c2 h2 c2' \rrbracket \implies P$
 $\rrbracket \implies P$
<proof>

lemma $hsem-lel$: $\llbracket hsem \Delta (c1@\pi#c2) h c';$
 $!!h1 t h2 c1' ct' c2'. \llbracket$
 $h=h1@t\#h2; c'=c1'@ct'@c2';$
 $hsem \Delta c1 h1 c1'; tsem \Delta \pi t ct'; hsem \Delta c2 h2 c2' \rrbracket$

$\] \Longrightarrow P$
 $\] \Longrightarrow P$
 <proof>

Given a partition of the hedge, we can construct a corresponding partition of the initial and final configuration.

lemma *hsem-split-h'*: $\llbracket hsem \Delta c (h1 @ h2) c' \rrbracket \Longrightarrow$
 $\exists c1 c2 c1' c2'. c = c1 @ c2 \wedge c' = c1' @ c2' \wedge$
 $hsem \Delta c1 h1 c1' \wedge hsem \Delta c2 h2 c2'$

<proof>

lemma *hsem-split-h*:

$\llbracket hsem \Delta c (h1 @ h2) c';$
 $!!c1 c2 c1' c2'.$
 $\llbracket c = c1 @ c2; c' = c1' @ c2'; hsem \Delta c1 h1 c1'; hsem \Delta c2 h2 c2' \rrbracket \Longrightarrow P$
 $\] \Longrightarrow P$
 <proof>

lemma *hsem-single-h*:

$\llbracket hsem \Delta c [t] c'; !!p. \llbracket c = [p]; tsem \Delta p t c' \rrbracket \Longrightarrow P \rrbracket \Longrightarrow P$
 <proof>

lemmas *hsem-split-h-single* = *hsem-split-h*[**where** ?h1 . 0 = [t], *simplified, standard*]

lemma *hsem-lel-h*: $\llbracket hsem \Delta c (h1 @ t \# h2) c';$

$!!c1 p c2 c1' ct' c2'. \llbracket$
 $c = c1 @ p \# c2; c' = c1' @ ct' @ c2';$
 $hsem \Delta c1 h1 c1'; tsem \Delta p t ct'; hsem \Delta c2 h2 c2' \rrbracket$
 $\] \Longrightarrow P$
 $\] \Longrightarrow P$
 <proof>

4.2.1 Scheduler

The scheduler maps execution hedges to compatible label sequences. This is done by eating up the given hedge from the roots to the leafs, until all non-leaf nodes have been consumed. From an ordering point of view, the hedge represents a partial ordering on the steps, and the scheduler maps this ordering to the set of all its topological sorts.

An execution hedge is called *final* if it solely consists of leaf nodes.

inductive *final-t where*

[*simp, intro!*]: *final-t* (*NLEAF pw*)

lemma [*simp, intro!*]:

$\neg final-t$ (*NNOSPAWN l t*)
 $\neg final-t$ (*NSPAWN l ts t*)
 <proof>

abbreviation $final == list\text{-}all\ final\text{-}t$

Final execution hedges contain no steps, hence they do not change the configuration.

lemma $final\text{-}tsem\text{-}nostep$: $\llbracket final\text{-}t\ t; tsem\ \Delta\ pw\ t\ c' \rrbracket \implies c' = [pw]$
 $\langle proof \rangle$

lemma $final\text{-}hsem\text{-}nostep$: $\llbracket final\ h; hsem\ \Delta\ c\ h\ c' \rrbracket \implies c' = c$
 $\langle proof \rangle$

As described above, the scheduler eats up the execution hedge from the roots to the leafs, until there are no inner nodes remaining, i.e. the hedge is final.

inductive $sched :: ('Q, T, 'L)\ ex\text{-}hedge \Rightarrow 'L\ list \Rightarrow bool$ **where**

$sched\text{-}final$: $final\ h \implies sched\ h\ []$ |

$sched\text{-}nospawn$:

$sched\ (h1 @ t \# h2)\ w \implies sched\ (h1 @ (NNOSPAWN\ l\ t) \# h2)\ (l \# w)$ |

$sched\text{-}spawn$:

$sched\ (h1 @ ts \# t \# h2)\ w \implies sched\ (h1 @ (NSPAWN\ l\ ts\ t) \# h2)\ (l \# w)$

inductive-set $sched\text{-}rel :: (('Q, T, 'L)\ ex\text{-}hedge, 'L)\ LTS$ **where**

$sched\text{-}rel\text{-}nospawn$: $((h1 @ (NNOSPAWN\ l\ t) \# h2), l, h1 @ t \# h2) \in sched\text{-}rel$ |

$sched\text{-}rel\text{-}spawn$: $((h1 @ (NSPAWN\ l\ ts\ t) \# h2), l, (h1 @ ts \# t \# h2)) \in sched\text{-}rel$

definition $sched'\ h\ ll == (\exists h'. (h, ll, h') \in trcl\ sched\text{-}rel \wedge final\ h')$

lemma $sched\text{-}alt1$: $sched\ h\ ll \implies sched'\ h\ ll$
 $\langle proof \rangle$

lemma $sched\text{-}rel\text{-}alt2$: $\llbracket (h, ll, h') \in trcl\ sched\text{-}rel; final\ h' \rrbracket \implies sched\ h\ ll$
 $\langle proof \rangle$

lemma $sched\text{-}alt$: $sched'\ h\ ll \longleftrightarrow sched\ h\ ll$
 $\langle proof \rangle$

We now show some basic facts about the scheduler.

lemma $sched\text{-}empty\text{-}seq[simp]$: $sched\ h\ [] \longleftrightarrow final\ h$
 $\langle proof \rangle$

lemma $sched\text{-}empty\text{-}hedge[simp]$: $sched\ []\ ll \longleftrightarrow ll = []$
 $\langle proof \rangle$

lemma $sched\text{-}empty\text{-}empty[simp, intro!]$: $sched\ []\ []$ $\langle proof \rangle$

lemma $sched\text{-}final\text{-}simp[simp]$: $final\ h \implies sched\ h\ c \longleftrightarrow c = []$
 $\langle proof \rangle$

In the following few lemmas we derive an induction scheme that reasons about hedges in the way they are consumed by the scheduler

fun *sched-ind-size* **where**
sched-ind-size (*NLEAF* π) = 0 |
sched-ind-size (*NNOSPAWN* $l\ t$) = *Suc* (*sched-ind-size* t) |
sched-ind-size (*NSPAWN* $l\ ts\ t$) = *Suc* (*sched-ind-size* ts + *sched-ind-size* t)

abbreviation *sched-ind-sizeh* h == *listsum* (*map* *sched-ind-size* h)

lemma *sched-ind-h-cases*[*consumes* 1, *case-names* *NOSPAWN SPAWN*]:

[[*sched-ind-sizeh* h > 0;
 !! $h1\ l\ t\ h2$. $h = h1 @ (NNOSPAWN\ l\ t) \# h2 \implies P$;
 !! $h1\ ts\ t\ h2\ l$. $h = h1 @ (NSPAWN\ l\ ts\ t) \# h2 \implies P$
]] $\implies P$
 <proof>

lemma *sched-ind-helper*:

[[!! h . *final* $h \implies P\ h$;
 !! $h1\ t\ h2\ l$. $P\ (h1 @ t \# h2) \implies P\ (h1 @ (NNOSPAWN\ l\ t) \# h2)$;
 !! $h1\ ts\ t\ h2\ l$. $P\ (h1 @ ts \# t \# h2) \implies P\ (h1 @ (NSPAWN\ l\ ts\ t) \# h2)$;
sched-ind-sizeh $h = k$
]] $\implies P\ h$
 <proof>

lemma *sched-ind*[*case-names* *FINAL NOSPAWN SPAWN*]:

[[!! h . *final* $h \implies P\ h$;
 !! $h1\ t\ h2\ l$. $P\ (h1 @ t \# h2) \implies P\ (h1 @ (NNOSPAWN\ l\ t) \# h2)$;
 !! $h1\ ts\ t\ h2\ l$. $P\ (h1 @ ts \# t \# h2) \implies P\ (h1 @ (NSPAWN\ l\ ts\ t) \# h2)$
]] $\implies P\ h$
 <proof>

Every tree/hedge has at least one schedule. From an ordering point of view, this is because hedge-structures are acyclic, and thus have always at least one topological sort. However, using the inductive definition of the scheduler, the proof of this lemma is by straightforward induction.

lemma *exists-schedule*: [[!! ll . *sched* $h\ ll \implies P$]] $\implies P$
 <proof>

Next, we want to show that the true concurrency semantics corresponds to the interleaving semantics. For this purpose, we show that we have an execution with labeling sequence ll in the interleaving semantics if and only if there is an execution h in the true concurrency semantics that has ll in its set of schedules.

The next two lemmas show the two directions of this claim.

lemma *sched-correct1*: $(c, ll, c') \in dpntrc\ \Delta \implies \exists h. hsem\ \Delta\ c\ h\ c' \wedge sched\ h\ ll$
 <proof>

lemma *sched-correct2*: [[*sched* $h\ ll$; *hsem* $\Delta\ c\ h\ c'$]] $\implies (c, ll, c') \in dpntrc\ \Delta$
 <proof>

Finally, we formulate the correspondance between the interleaving and the true concurrency semantics as a single equivalence:

theorem *sched-correct*: $(c, ll, c') \in dpntrc \Delta \longleftrightarrow (\exists h. hsem \Delta c h c' \wedge sched h ll)$
 ⟨proof⟩

As any hedge has at least one schedule, we always get an interleaving execution from a hedge execution:

lemma *obtain-schedule*:

[[$hsem \Delta c h c'$;
 !! $ll. [(c, ll, c') \in dpntrc \Delta; sched h ll]$]] $\implies P$
] $\implies P$
 ⟨proof⟩

5 Predecessor Sets

Following [2], we define the set of immediate predecessors $pre \Delta C$ and predecessors $pre^* \Delta C$ of a set of configurations C . The set of immediate predecessors contains those configurations from that we can reach (a configuration in) C with exactly one step. The set of predecessors contains those configurations from that we can reach C with an arbitrary number of steps, including no steps at all (i.e. pre^* is reflexive).

Computing predecessor sets is the key to model checking and analysis of DPNs, see [2] for details.

definition $pre \Delta C' == \{ c . \exists l c'. c' \in C' \wedge (c, l, c') \in dpntr \Delta \}$

definition *pre-star* (pre^*) **where**

$pre^* \Delta C' == \{ c . \exists ll c'. c' \in C' \wedge (c, ll, c') \in dpntrc \Delta \}$

5.1 Hedge-Constrained Predecessor Sets

For a set of configurations C' and a set of execution hedges H , we define the *hedge-constrained predecessor set* of C' w.r.t. H as the set of those configurations from that we can reach C' with an execution hedge in H .

definition $prehc \Delta H C' == \{ c . \exists h c'. h \in H \wedge c' \in C' \wedge hsem \Delta c h c' \}$

lemma *prehcI*: $[[h \in H; c' \in C'; hsem \Delta c h c']] \implies c \in prehc \Delta H C'$
 ⟨proof⟩

lemma *prehcE*:

[[$c \in prehc \Delta H C'$; !! $h c'. [[h \in H; c' \in C'; hsem \Delta c h c']] \implies P$]] $\implies P$
 ⟨proof⟩

The hedge-constrained predecessor set is monotonic in the constraint

lemma *prehc-mono*: $H \subseteq H' \implies prehc \Delta H C' \subseteq prehc \Delta H' C'$
 ⟨proof⟩

The hedge-constrained predecessor set without constraints is the same as the original predecessor set.

lemma *prehc-triv-is-pre-star*: $prehc \Delta UNIV C' = pre^* \Delta C'$
 ⟨proof⟩

The hedge-constrained predecessor set is always a subset of the unconstrained predecessor set.

lemma *prehc-subset-pre-star*: $prehc \Delta H C' \subseteq pre^* \Delta C'$
 ⟨proof⟩

We can use a hedge-constraint to express immediate predecessor sets.

definition *Hpre* :: ('P, T, 'L) *ex-hedge set* **where**

$$Hpre == \{ hl1 @ t \# hl2 \mid hl1 \ t \ hl2 \ lab \ ts \ t'. \\ final \ hl1 \ \wedge \ final \ hl2 \ \wedge \ final-t \ ts \ \wedge \ final-t \ t' \ \wedge \\ (t = NNOSPAWN \ lab \ t' \vee t = NSPAWN \ lab \ ts \ t') \}$$

lemma *HpreI-nospawn*:

$$\llbracket final \ h1; \ final \ h2; \ final-t \ t \rrbracket \implies h1 @ NNOSPAWN \ lab \ t' \# h2 \in Hpre$$

⟨proof⟩

lemma *HpreI-spawn*:

$$\llbracket final \ h1; \ final \ h2; \ final-t \ ts; \ final-t \ t \rrbracket \implies h1 @ NSPAWN \ lab \ ts \ t' \# h2 \in Hpre$$

⟨proof⟩

lemmas *HpreI = HpreI-nospawn HpreI-spawn*

lemma *HpreE*[*cases set, consumes 1, case-names nospawn spawn*]:

$$\llbracket h \in Hpre; \\ !!h1 \ lab \ t' \ h2. \llbracket \\ h = h1 @ NNOSPAWN \ lab \ t' \# h2; \ final \ h1; \ final \ h2; \ final-t \ t' \\ \rrbracket \implies P; \\ !!h1 \ lab \ ts \ t' \ h2. \llbracket \\ h = h1 @ NSPAWN \ lab \ ts \ t' \# h2; \\ \ final \ h1; \ final \ h2; \ final-t \ ts; \ final-t \ t' \\ \rrbracket \implies P \\ \rrbracket \implies P$$

⟨proof⟩

In order to show that *Hpre* is correct, we first show that it exactly admits the schedules of length one.

lemma *Hpre-length1*: $\llbracket h \in Hpre; \ sched \ h \ ll \rrbracket \implies length \ ll = 1$
 ⟨proof⟩

lemma *Hpre-length2*: $\llbracket \ sched \ h \ ll; \ length \ ll = 1 \rrbracket \implies h \in Hpre$
 ⟨proof⟩

theorem *Hpre-length*: $\llbracket \ sched \ h \ ll \rrbracket \implies h \in Hpre \longleftrightarrow length \ ll = 1$
 ⟨proof⟩

It is then straightforward to show that $prehc \Delta Hpre = pre \Delta$

lemma *Hpre-correct1*: $c \in prehc \Delta Hpre \ C' \implies c \in pre \Delta C'$

<proof>

lemma *Hpre-correct2*: $c \in \text{pre} \Delta C' \implies c \in \text{prehc} \Delta \text{Hpre} C'$
<proof>

theorem *Hpre-correct*: $\text{prehc} \Delta \text{Hpre} = \text{pre} \Delta$
<proof>

end

6 DPN Semantics on Lists

theory *ListSemantics*
imports *Semantics*
begin

The interleaving semantics works on configurations that are lists of process configurations.

However, in [2] a DPN configuration is represented as a sequence of control and stack symbols. Each process starts with a control symbol, followed by its stack symbols. The configuration is simply a concatenation of processes. This representation allows the notion of a regular set of configurations as a set of configurations accepted by a FSM.

In this theory, we adopt this representation of configurations, define a semantics directly over this representation, and show that this representation is isomorphic to ours for sequences starting with a control symbol. Note that sequences starting with a stack symbol have no meaningful interpretation, as each process's configuration has to start with a control symbol.

6.1 Definitions

We separate stack and control symbols using a datatype with two constructors:

datatype $('Q, T)$ *cl-item* = *CTRL* 'Q | *STACK* T
types $('Q, T)$ *cl* = $('Q, T)$ *cl-item list*

The mapping from configurations to list-based configurations is straightforward:

fun *pc2cl* :: $('Q, T)$ *pconf* \Rightarrow $('Q, T)$ *cl* **where**
pc2cl (p,w) = *CTRL* p # *map STACK w*

definition *c2cl* :: $('Q, T)$ *conf* \Rightarrow $('Q, T)$ *cl* **where**
c2cl c == *concat (map pc2cl c)*

abbreviation *c2cl-abbrev* :: $('Q, T)$ *conf* \Rightarrow $('Q, T)$ *cl*
— This abbreviation is just for convenience

where

$c2cl\text{-}abbrev\ c == concat\ (map\ pc2cl\ c)$

Valid single-process configurations are those that start with a control symbol followed by a list of stack symbols:

definition $pcvalid == \{CTRL\ p\#\#map\ STACK\ w\ | \ p\ w.\ True\}$

Valid configurations are those that start with a control symbol:

definition $clvalid == \{\}\ \cup\ \{CTRL\ p\#\#c\ | \ p\ c.\ True\}$

We also define the step relation directly on list representation of configurations:

inductive-set $cltr :: ('Q,\Gamma,'L)\ dpn \Rightarrow (('Q,\Gamma)\ cl \times 'L \times ('Q,\Gamma)\ cl)\ set$

for Δ **where**

cltr-no-spawn:

$$\begin{aligned} & \llbracket (p,\gamma \hookrightarrow_l p',w) \in \Delta \rrbracket \implies \\ & \quad (c1@[CTRL\ p,\ STACK\ \gamma]@c2, \\ & \quad \quad l, \\ & \quad \quad c1@CTRL\ p'\#\#(map\ STACK\ w)@c2 \\ & \quad) \in cltr\ \Delta \mid \end{aligned}$$

cltr-spawn:

$$\begin{aligned} & \llbracket (p,\gamma \hookrightarrow_l ps,ws \# p',w) \in \Delta \rrbracket \implies \\ & \quad (c1@[CTRL\ p,\ STACK\ \gamma]@c2, \\ & \quad \quad l, \\ & \quad \quad c1@CTRL\ ps\#\#(map\ STACK\ ws)@CTRL\ p'\#\#(map\ STACK\ w)@c2 \\ & \quad) \in cltr\ \Delta \end{aligned}$$

6.2 Theorems

lemma $inj\text{-}STACK[simp,\ intro!]:\ inj\ STACK\ \langle proof \rangle$

6.2.1 Representation of Single Processes

lemma $pc2cl\text{-}not\text{-}empty[simp]:\ pc2cl\ \pi \neq []\ \langle proof \rangle$

lemma $pc2cl\text{-}inj[simp,\ intro!]:\ inj\ pc2cl\ \langle proof \rangle$

lemmas $pc2cl\text{-}inj\text{-}simp[simp] = inj\text{-}eq[OF\ pc2cl\text{-}inj]$

lemma $pc2cl\text{-}valid[intro!,simp]:\ pc2cl\ \pi \in pcvalid\ \langle proof \rangle$

lemma $pc2cl\text{-}surj: \llbracket \pi l \in pcvalid; !!\pi.\ \pi l = pc2cl\ \pi \implies P \rrbracket \implies P\ \langle proof \rangle$

6.2.2 Representation of Configurations

We start with a bunch of simplification rules and other auxilliary lemmas:

lemma *stack-no-ctrl1*[simp]:
 $\text{map STACK } w \neq c1 @ \text{CTRL } p \# c2$
 ⟨proof⟩

lemmas *stack-no-ctrl2*[simp] = *stack-no-ctrl1*[symmetric]

lemma *map-stack-ne-cCc1* [simp]:
 $\text{map STACK } w \neq c @ \text{CTRL } s \# c'$
 ⟨proof⟩

lemmas *map-stack-ne-cCc2*[simp] = *map-stack-ne-cCc1*[symmetric]

lemmas *map-stack-ne-add-simps*[simp] =
 map-stack-ne-cCc1 [where $c=[]$, simplified]
 map-stack-ne-cCc1 [where $c=[a]$, simplified, standard]

lemma *map-STACK-eq-map-STACK-simp*[simp]:
 $\text{map STACK } w @ \text{CTRL } p \# cl = \text{map STACK } w' @ \text{CTRL } p' \# cl' \longleftrightarrow$
 $w'=w \wedge p'=p \wedge cl'=cl$
 ⟨proof⟩

lemma *map-stack-ne-pc2cl*[simp]:
 $\text{map STACK } w \neq c @ \text{pc2cl } \pi @ c'$
 $c @ \text{pc2cl } \pi @ c' \neq \text{map STACK } w$
 ⟨proof⟩

lemmas *map-stack-ne-pc2cl-add-simps*[simp] =
 $\text{map-stack-ne-pc2cl}$ [where $c=[]$, simplified]

lemma *map-STACK-eq-map-STACK-add-simps*[simp]:
 $\text{map STACK } w @ \text{CTRL } p \# cl = \text{map STACK } w' @ \text{pc2cl } \pi' @ cl' \longleftrightarrow$
 $w=w' \wedge p=\text{fst } \pi' \wedge cl=\text{map STACK } (\text{snd } \pi') @ cl'$
 $\text{map STACK } w' @ \text{pc2cl } \pi' @ cl' = \text{map STACK } w @ \text{CTRL } p \# cl \longleftrightarrow$
 $w=w' \wedge p=\text{fst } \pi' \wedge cl=\text{map STACK } (\text{snd } \pi') @ cl'$
 ⟨proof⟩

lemma *c2cl-simps*[simp]:
 $c2cl [] = []$
 $c2cl (\pi \# c) = \text{pc2cl } \pi @ c2cl c$
 $c2cl (c1 @ c2) = c2cl c1 @ c2cl c2$
 ⟨proof⟩

lemma *c2cl-empty*[simp]:
 $c2cl c = [] \longleftrightarrow c=[]$
 $[] = c2cl c \longleftrightarrow c=[]$
 ⟨proof⟩

lemma *c2cl-start-with-ctrl*[simp]:

$c2cl\ c \neq STACK\ \gamma\ \#cl$
 $STACK\ \gamma\ \#cl \neq c2cl\ c$
 ⟨proof⟩

lemma *c2cl-start-with-ctrl-map*:

$w \neq [] \implies c2cl\ c \neq map\ STACK\ w$
 $w \neq [] \implies map\ STACK\ w \neq c2cl\ c$
 ⟨proof⟩

lemma *map-stack-c2cl-eq-simps*[simp]:

$map\ STACK\ w\ @\ c2cl\ c = map\ STACK\ w'\ @\ c2cl\ c' \longleftrightarrow w=w' \wedge c2cl\ c=c2cl\ c'$
 ⟨proof⟩

lemma *c2cl-s-cl-eqE*:

$\llbracket STACK\ \gamma\ \#cl = map\ STACK\ w\ @\ c2cl\ c;$
 $\quad !!wr.\ \llbracket w=\gamma\ \#wr; cl = map\ STACK\ wr\ @\ c2cl\ c \rrbracket \implies P$
 $\rrbracket \implies P$
 ⟨proof⟩

lemma *c2cl-first-processE*:

$\llbracket c2cl\ c = CTRL\ p\ \#cl2;$
 $\quad !!w\ c2\ cl2'.\ \llbracket c=(p,w)\ \#c2; cl2=(map\ STACK\ w)\ @\ cl2'; c2cl\ c2=c2cl\ c' \rrbracket \implies P$
 $\rrbracket \implies P$
 ⟨proof⟩

lemma *c2cl-find-process1*:

$\llbracket c2cl\ c = cl1\ @\ CTRL\ p\ \#cl2;$
 $\quad !!c1\ w\ c2.\ \llbracket c=c1\ @\ (p,w)\ \#c2; cl2=(map\ STACK\ w)\ @\ c2cl\ c2;$
 $\quad \quad \quad cl1=c2cl\ c1$
 $\rrbracket \implies P$
 $\rrbracket \implies P$
 ⟨proof⟩

Then we show that our representation mapping is injective and surjective on valid configurations.

lemma *c2cl-inj*[simp, intro!]: *inj c2cl*

⟨proof⟩

lemmas *c2cl-inj-simps*[simp] = *inj-eq*[OF *c2cl-inj*]

lemmas *c2cl-img-Int*[simp] = *image-Int*[OF *c2cl-inj*]

lemma *c2cl-valid*[simp, intro!]: $c2cl\ c \in cvalid$

⟨proof⟩

lemma *c2cl-surj*: $\llbracket cl \in cvalid; !!c.\ cl=c2cl\ c \implies P \rrbracket \implies P$

<proof>

6.2.3 Step Relation on List-Configurations

lemma *cltr-pres-valid*: $(cl, l, cl') \in cltr \Delta \implies cl \in clvalid \iff cl' \in clvalid$
<proof>

lemma *dpntr-is-cltr*: $\llbracket (c, l, c') \in dpntr \Delta \rrbracket \implies (c2cl \ c, l, c2cl \ c') \in cltr \Delta$
<proof>

lemma *cltr-is-dpntr*: $\llbracket (c2cl \ c, l, c2cl \ c') \in cltr \Delta \rrbracket \implies (c, l, c') \in dpntr \Delta$
<proof>

The following theorem formulates the equivalence of the original semantics and the list-based semantics.

theorem *cltr-eq-dpntr*: $(c2cl \ c, l, c2cl \ c') \in cltr \Delta \iff (c, l, c') \in dpntr \Delta$
<proof>

The next two lemmas ease the derivation of executions of the original semantics from executions of the list-based semantics.

lemma *cltr2dpntr-fwd*:
 $\llbracket (c2cl \ c, l, cl') \in cltr \Delta; \\ !!c'. \llbracket cl' = c2cl \ c'; (c, l, c') \in dpntr \Delta \rrbracket \implies P \rrbracket \implies P$
<proof>

lemma *cltr2dpntr-bwd*:
 $\llbracket (cl, l, c2cl \ c') \in cltr \Delta; \\ !!c. \llbracket cl = c2cl \ c; (c, l, c') \in dpntr \Delta \rrbracket \implies P \rrbracket \implies P$
<proof>

Finally, we give some lemmas to directly reason about the transitive closure of the step relation:

lemma *cltr-is-dpntrc*:
 $(c2cl \ c, l, c2cl \ c') \in trcl (cltr \Delta) \implies (c, l, c') \in dpntrc \Delta$
<proof>

lemma *dpntrc-is-cltr*:
 $(c, l, c') \in dpntrc \Delta \implies (c2cl \ c, l, c2cl \ c') \in trcl (cltr \Delta)$
<proof>

theorem *cltr-eq-dpntrc*:
 $(c2cl \ c, l, c2cl \ c') \in trcl (cltr \Delta) \iff (c, l, c') \in dpntrc \Delta$
<proof>

lemma *cltrc-pres-valid*:
 $(cl, w, cl') \in trcl (cltr \Delta) \implies cl \in clvalid \iff cl' \in clvalid$
<proof>

lemma *cltr2dpntrc-fwd*:

$$\llbracket (c2cl\ c, l, cl') \in trcl\ (cltr\ \Delta);$$

$$\quad !!c'. \llbracket cl' = c2cl\ c'; (c, l, c') \in dpntrc\ \Delta \rrbracket \implies P$$

$$\rrbracket \implies P$$
<proof>

lemma *cltr2dpntrc-bwd*:

$$\llbracket (cl, l, c2cl\ c') \in trcl\ (cltr\ \Delta);$$

$$\quad !!c. \llbracket cl = c2cl\ c; (c, l, c') \in dpntrc\ \Delta \rrbracket \implies P$$

$$\rrbracket \implies P$$
<proof>

6.3 Predecessor Sets on List-Semantics

We also define predecessor sets for the list-semantics:

definition *precl* (*pre_{cl}*) **where**

$$pre_{cl}\ \Delta\ C' == \{ c . \exists l\ c'. c' \in C' \wedge (c, l, c') \in cltr\ \Delta \}$$

definition *precl-star* (*pre*_{cl}*) **where**

$$pre^*_{cl}\ \Delta\ C' == \{ c . \exists ll\ c'. c' \in C' \wedge (c, ll, c') \in trcl\ (cltr\ \Delta) \}$$

And show that they are equivalent to their counterparts defined over the original semantics:

lemma *precl-is-pre*: $pre_{cl}\ \Delta\ (c2cl' C) = c2cl'(pre\ \Delta\ C)$
<proof>

lemma *precl-star-is-pre-star*: $pre^*_{cl}\ \Delta\ (c2cl' C) = c2cl'(pre^*\ \Delta\ C)$
<proof>

end

7 Automata for Execution Hedges

theory *HedgeAutomata*
imports *Main Semantics*
begin

In this section we define hedge automata that accept execution hedges.

A hedge automaton consists of a set of states, an regular *initial language* of state sequences and a set of transitions. Transitions are either leaf transitions that label a leaf node with a state if the configuration at the leaf node is contained in some (regular) language, or non-spawning or spawning transitions, that label a spawning or non-spawning node respectively with a state depending on the states of the successor nodes.

In this formalization, we model the initial language and the regular languages at the leafs just at sets. However, if we want an executable representation, we need to model real automata there. This is planned to be done in the future.

datatype ('S,'P,'T,'L) *ha-rule* =
 HAR-LEAF 'S 'P 'T list set |
 HAR-NOSPAWN 'S 'L 'S |
 HAR-SPAWN 'S 'L 'S 'S

types ('S,'P,'T,'L) *ha* = 'S list set × ('S,'P,'T,'L) *ha-rule* set

In order to model acceptance of a hedge, we define a relation between trees and states with which we can label those trees. We then extend this relation to hedges.

inductive *lab*

:: ('S,'P,'T,'L) *ha-rule* set ⇒ ('P,'T,'L) *ex-tree* ⇒ 'S ⇒ bool

for *H* **where**

lab-leaf:

[[HAR-LEAF *s p W* ∈ *H*; *w* ∈ *W*]] ⇒ *lab H (NLEAF (p,w)) s* |

lab-nospawn:

[[HAR-NOSPAWN *s l s'* ∈ *H*; *lab H t s'*]] ⇒ *lab H (NNOSPAWN l t) s* |

lab-spawn:

[[HAR-SPAWN *s l ss s'* ∈ *H*; *lab H ts ss*; *lab H t s'*]] ⇒
lab H (NSPAWN l ts t) s

inductive *labh* :: ('S,'P,'T,'L) *ha-rule* set ⇒ ('P,'T,'L) *ex-hedge* ⇒ 'S list ⇒ bool

for *H* **where**

labh-empty[*simp*, *intro!*]: *labh H [] []* |

labh-cons: [[*lab H t s*; *labh H h σ*]] ⇒ *labh H (t#h) (s#σ)*

lemma *labh-empty*[*simp*]:

labh H [] σ ↔ *σ* = []

labh H h [] ↔ *h* = []

⟨*proof*⟩

lemma *labh-length*: *labh H h σ* ⇒ *length h* = *length σ*

⟨*proof*⟩

The language of a hedge automaton consists of those hedges whose roots can be labeled with a state sequence in the initial language.

definition *languh* :: ('S,'P,'T,'L) *ha* ⇒ ('P,'T,'L) *ex-hedge* set **where**

languh HA == { *h* . ∃ *σ* ∈ *fst HA*. *labh (snd HA) h σ* }

end

8 Computation of Hedge-Constrained Predecessor Sets

```

theory CrossProd
imports ListSemantics HedgeAutomata
begin

```

In this section we show how to compute predecessor sets with regular hedge constraints. The computation is done by reduction to the computation of the unconstrained predecessor set. The reduction uses a cross-product like approach, computing a product-DPN of the original DPN and the hedge automaton, and a product regular set of the original regular set and the hedge-automaton's leaf rules.

This theory uses a list-based representation of DPN-configurations, where the type of a configuration is a list of control- and stack-symbols. This type is less structured than the original type of configurations, that is lists of pairs of control symbol and stack. However, it admits handling configurations as words, and sets of configurations as (regular) languages.

This theory does not use a formalization of regular languages, nor does it generate executable code. Instead, regular sets are modeled as sets. The effectiveness proofs show representations that only contain operations well-known to preserve regularity. However, an implementation of those operations is not formalized.

The cross-product DPN simulates the rules of the hedge-automaton via its transitions, the current state of the hedge automaton is stored in the DPN's state:

inductive-set

$$xdpn :: ('P, \tau, 'L) \text{ dpn} \Rightarrow ('S, 'P, \tau, 'L) \text{ ha-rule set} \Rightarrow ('P \times 'S, \tau, 'L) \text{ dpn}$$

for Δ H **where**

xdpn-nospawn:

$$\begin{aligned} & \llbracket (p, \gamma \hookrightarrow_l p', w) \in \Delta; \text{HAR-NOSPAWN } s \ l \ s' \in H \rrbracket \Longrightarrow \\ & ((p, s), \gamma \hookrightarrow_l (p', s'), w) \in xdpn \ \Delta \ H \ | \end{aligned}$$

xdpn-spawn:

$$\begin{aligned} & \llbracket (p, \gamma \hookrightarrow_l ps, ws \ \# \ p', w) \in \Delta; \text{HAR-SPAWN } s \ l \ ss \ s' \in H \rrbracket \Longrightarrow \\ & ((p, s), \gamma \hookrightarrow_l (ps, ss), ws \ \# \ (p', s'), w) \in xdpn \ \Delta \ H \end{aligned}$$

The *xdpn-nospawn*-rule adds a transition rule to the cross-product DPN for each original non-spawning transition rule and hedge automaton rule that could be used to label the node generated by this transition rule. Analogously, the *xdpn-spawn*-rule adds a transition rule to the cross-product DPN for spawning rules.

We now define operators to map configurations of the cross-product DPN to configurations of the original DPN and sequences of states of the hedge automaton.

abbreviation

$proj-c1 :: ('P \times 'S, \mathbb{T}) \text{ conf} \Rightarrow ('P, \mathbb{T}) \text{ conf}$ **where**
 $proj-c1 \text{ cx} == \text{map } (\lambda((p,s),w). (p,w)) \text{ cx}$

abbreviation

$proj-c2 :: ('P \times 'S, \mathbb{T}) \text{ conf} \Rightarrow 'S \text{ list}$ **where**
 $proj-c2 \text{ cx} == \text{map } (\lambda((p,s),w). s) \text{ cx}$

We also have to define a mapping for execution hedges, because the labeling of the leafs is different:

fun $proj-t1 :: ('P \times 'S, \mathbb{T}, 'L) \text{ ex-tree} \Rightarrow ('P, \mathbb{T}, 'L) \text{ ex-tree}$ **where**
 $proj-t1 (NLEAF ((p,s),w)) = NLEAF (p,w) \mid$
 $proj-t1 (NNOSPAWN l t) = NNOSPAWN l (proj-t1 t) \mid$
 $proj-t1 (NSPAWN l ts t) = NSPAWN l (proj-t1 ts) (proj-t1 t)$

Next we define how to transform the target set, that contains the configurations of that we want to compute the predecessors.

The new target set contains the configurations of the original target set with all labelings that may be done by leaf-rules of the hedge automaton:

— Process labeled by a leaf-rule:

abbreviation

$xdpnCLP H == \{ ((p,s),w). \exists W. HAR-LEAF s p W \in H \wedge w \in W \}$

— Configuration labeled by leaf-rules:

abbreviation

$xdpnCL H == \{ \text{cx} . (\forall ((p,s),w) \in \text{set } \text{cx}. ((p,s),w) \in \text{xdpnCLP } H) \}$

— New target set:

definition

$xdpnC C H == \{ \text{cx} . \text{proj-c1 } \text{cx} \in C \} \cap \text{xdpnCL } H$

Finally we define how to transform the computed predecessor set in order to get a set of configurations of the original DPN. This phase consists of two operations: First, we have to restrict the configurations to those that are accepted by the hedge automaton's initial language, and then we have to project away the hedge-automaton's states to get a configuration of the original DPN. In the following definition, these two steps are combined:

definition

$projH :: 'S \text{ list set} \Rightarrow ('P \times 'S, \mathbb{T}) \text{ conf set} \Rightarrow ('P, \mathbb{T}) \text{ conf set}$ **where**
 $projH H0 Cx == \{ \text{proj-c1 } \text{cx} \mid \text{cx}. \text{cx} \in Cx \wedge \text{proj-c2 } \text{cx} \in H0 \}$

8.1 Correctness of Reduction

In this section we show that our reduction is correct, i.e. that we really get the hedge-constrained predecessor set by computing the predecessor set of the cross-product DPN and a transformed target set, and then applying the $projH$ -operator to the result.

We first need to introduce a combination operator that combines an original DPN's configuration and a list of hedge automaton states to a cross-product DPN's configuration.

abbreviation $cx\ s\ c\ \sigma == \text{zipf } (\lambda(p,w)\ s.\ ((p,s),w))\ c\ \sigma$

lemma $\text{proj-cxs1}[simp]$: $\text{length } c = \text{length } \sigma \implies \text{proj-c1 } (cx\ s\ c\ \sigma) = c$
 ⟨proof⟩

lemma $\text{proj-cxs2}[simp]$: $\text{length } c = \text{length } \sigma \implies \text{proj-c2 } (cx\ s\ c\ \sigma) = \sigma$
 ⟨proof⟩

lemma $\text{cx\ s\ proj}[simp]$: $cx\ s\ (\text{proj-c1 } cx)\ (\text{proj-c2 } cx) = cx$
 ⟨proof⟩

lemma xdpn\ c\ proj : $cx \in \text{xdpnC } C\ H \implies \text{proj-c1 } cx \in C$
 ⟨proof⟩

We now prove the two directions of our main goal. Each direction requires 2 lemmas, the first one for a single tree and the second one for a hedge.

lemmas $\text{tsem-induct-x} =$
 $\text{tsem.induct}[\text{ where } ?x1.0 = ((p,s),w), \text{ split-format } (\text{complete}),$
 $\text{ consumes } 1, \text{ case-names } \text{tsem-leaf } \text{tsem-nospawn } \text{tsem-spawn}$
 $]$

lemmas $\text{tsem-induct-p} =$
 $\text{tsem.induct}[\text{ where } ?x1.0 = (p,w), \text{ split-format } (\text{complete}),$
 $\text{ consumes } 1, \text{ case-names } \text{tsem-leaf } \text{tsem-nospawn } \text{tsem-spawn}$
 $]$

lemma xdpn-correct1-t :
 $\llbracket \text{tsem } (xdpn\ \Delta\ H)\ ((p,s),w)\ t\ c';\ c' \in \text{xdpnCL } H \rrbracket \implies$
 $\text{tsem } \Delta\ (p,w)\ (\text{proj-t1 } t)\ (\text{proj-c1 } c') \wedge \text{lab } H\ (\text{proj-t1 } t)\ s$
 ⟨proof⟩

lemma xdpn-correct1 :
 $\llbracket \text{hsem } (xdpn\ \Delta\ H)\ c\ h\ c';\ c' \in \text{xdpnCL } H \rrbracket \implies$
 $\text{hsem } \Delta\ (\text{proj-c1 } c)\ (\text{map } \text{proj-t1 } h)\ (\text{proj-c1 } c') \wedge$
 $\text{labh } H\ (\text{map } \text{proj-t1 } h)\ (\text{proj-c2 } c)$
 ⟨proof⟩

lemma xdpn-correct2-t :
 $\llbracket \text{tsem } \Delta\ (p,w)\ t\ c';\ \text{lab } H\ t\ s \rrbracket \implies$
 $\exists tx\ cx'. \text{tsem } (xdpn\ \Delta\ H)\ ((p,s),w)\ tx\ cx' \wedge$
 $cx' \in \text{xdpnCL } H \wedge \text{proj-t1 } tx = t \wedge$
 $\text{proj-c1 } cx' = c'$
 ⟨proof⟩

lemma xdpn-correct2 :
 $\llbracket \text{hsem } \Delta\ c\ h\ c';\ \text{labh } H\ h\ \sigma \rrbracket \implies$
 $\exists hx\ cx'. \text{hsem } (xdpn\ \Delta\ H)\ (cx\ s\ c\ \sigma)\ hx\ cx' \wedge$

$$\begin{aligned}
& cx' \in xdpnCL H \wedge \\
& (map\ proj-t1\ hx) = h \wedge \\
& proj-c1\ cx' = c'
\end{aligned}$$

<proof>

Finally we use the lemmas proven above to show our main goal, i.e. a representation of the hedge-constrained predecessor set w.r.t. the language of a hedge automaton by means of the sequential pre^* -operator and the cross-product construction.

theorem *xdpn-correct*:

$$prehc\ \Delta\ (langu\ (H0,H))\ C' = projH\ H0\ (\ pre^*\ (xdpn\ \Delta\ H)\ (xdpnC\ C'\ H))$$

<proof>

8.2 Effectiveness of Reduction

In this section we give indication that the cross-product construction is computable for regular target sets.

The new set of rules $xdpn$ can be computed if the set of dpn rules and the set of hedge automaton transitions are finite, as the definition of $xdpn$ is not recursive and each LHS depends on only one element of each set. However, as said above, we do not provide executable code here.

In [2], a configuration is represented as a sequence of control and stack symbols, each process starting with a control symbol followed by its stack. For sequences that start with a control symbol, this representation is isomorphic to our representation (cf. Section 6.2.3). As regular sets of configurations are best defined on this list-based semantics, we also show the effectiveness of our construction on the list-based semantics.

This section, especially the proofs of the Theorems, are rather technical. The theorems itself show how to compute the new target configuration and the projection from the computed predecessor set using only operations well-known to preserve regularity (in this case intersection, union, concatenation, star, and substitution) as well as some sets that are obviously regular. However, no formal proof of regularity or effectiveness is given.

8.2.1 Definitions

This function defines the projection operator from the extended to the original configuration:

fun *fp-cl1* **where**

$$\begin{aligned}
fp-cl1\ (CTRL\ (p,s)) &= CTRL\ p\ | \\
fp-cl1\ (STACK\ \gamma) &= STACK\ \gamma
\end{aligned}$$

This function maps a hedge-automaton state to the regular set of all process configurations labeled with that state. Note that the sets $\{[CTRL\ (p, s)]\ | p. True\}$ and $\{[STACK\ \gamma]\ | \gamma. True\}$ are obviously regular.

definition *fp-inv-subst2* **where**

$fp\text{-inv}\text{-subst2 } s = conc \{ [CTRL (p,s)] \mid p. True \} (star \{ [STACK \gamma] \mid \gamma. True \})$

The projection operator can be written using substitution, projection (a special form of substitution), and intersection.

The intuitive idea is, that $subst \text{ fp}\text{-inv}\text{-subst2 } H0$ is the set of all configurations with a hedge-automaton labeling sequence that is accepted by $H0$.

definition $projH\text{-cl} :: 'S \text{ list set} \Rightarrow ('Q \times 'S, \Gamma) \text{ cl set} \Rightarrow ('Q, \Gamma) \text{ cl set}$ **where**
 $projH\text{-cl } H0 \text{ Clx} = lang\text{-proj } fp\text{-cl1} (subst \text{ fp}\text{-inv}\text{-subst2 } H0 \cap (Clx))$

The derivation of the new target set is done by first characterizing all sets of cross-product configurations whose leafs are labeled correctly according to the leaf rules of the hedge automaton. Note that there are only finitely many leaf-rules, hence the union below is over a finite set. Moreover, the language W at a leaf rule is regular by default, the operation $map \text{ STACK } \text{'}$ - is a projection and the operation $op \# (CTRL (p,s)) \text{'}$ - is a concatenation. Hence all the operations below are effective.

definition $xdpnCL\text{-cl} :: ('S, 'P, \Gamma, 'L) \text{ ha-rule set} \Rightarrow ('P \times 'S, \Gamma) \text{ cl set}$ **where**
 $xdpnCL\text{-cl } H = star (\bigcup \{ op \# (CTRL (p,s)) \text{' } (map \text{ STACK } \text{' } W) \mid$
 $s \text{ p } W. \text{ HAR-LEAF } s \text{ p } W \in H \}$
 $)$

Having characterized all configurations that are correctly labeled, one gets the new target set by intersecting them with all configurations that correspond to the old target set:

definition $xdpnC\text{-cl}$
 $:: ('P, \Gamma) \text{ cl set} \Rightarrow ('S, 'P, \Gamma, 'L) \text{ ha-rule set} \Rightarrow ('P \times 'S, \Gamma) \text{ cl set}$
where
 $xdpnC\text{-cl } Cl \text{ H} = lang\text{-inv}\text{-proj } fp\text{-cl1 } Cl \cap xdpnCL\text{-cl } H$

In order to compute $prehc \Delta (langh (H0, H)) \text{ C}'$, we map C' to its corresponding regular set of list-based configurations $c2cl \text{ C}'$ and apply the list-based operations for cross-product, predecessor set and projection on it:

definition $prehc\text{-cl}$
 $:: ('Q, \Gamma, 'L) \text{ dpn} \Rightarrow ('S, 'Q, \Gamma, 'L) \text{ ha} \Rightarrow ('Q, \Gamma) \text{ cl set} \Rightarrow ('Q, \Gamma) \text{ cl set}$
where
 $prehc\text{-cl } \Delta \text{ HA } Cl' =$
 $projH\text{-cl} (fst \text{ HA}) (pre^*_{cl} (xdpn \Delta (snd \text{ HA})) (xdpnC\text{-cl } Cl' (snd \text{ HA})))$

8.2.2 Theorems

lemma $fp\text{-cl1}\text{-map}\text{-stack}\text{-id}[simp]$: $map \text{ fp}\text{-cl1} (map \text{ STACK } w) = map \text{ STACK } w$
 $\langle proof \rangle$

lemma $fp\text{-cl1}\text{-stack}\text{-id}[simp]$: $fp\text{-cl1 } s = \text{STACK } \gamma \iff s = \text{STACK } \gamma$
 $\langle proof \rangle$

lemma *fp-cl1-eq-map-stack[simp]*:
 $map\ fp-cl1\ la = map\ STACK\ w \longleftrightarrow la = map\ STACK\ w$
 $\langle proof \rangle$

lemma *star-STACK[simplified,simp]*:
 $star\ \{[STACK\ \gamma] \mid \gamma.\ True\} = \{map\ STACK\ w \mid w.\ True\}$
 $\langle proof \rangle$

lemma *proj-c1-effective*: $c2cl\ (proj-c1\ c) = map\ fp-cl1\ (c2cl\ c)$
 $\langle proof \rangle$

lemma *fp-inv-subst2I[intro!, simp]*:
 $CTRL\ (p,s)\#map\ STACK\ w \in fp-inv-subst2\ s$
 $\langle proof \rangle$

lemma *fp-inv-subst2E*:
 $\llbracket cl \in fp-inv-subst2\ s; !!p\ w.\ cl = CTRL\ (p,s)\#map\ STACK\ w \implies P \rrbracket \implies P$
 $\langle proof \rangle$

Idea of the operation on the original representations of configurations:

lemma *projH-effective'*:
 $projH\ H0\ Cx = lang-proj\ (\lambda((p,s),w).\ (p,w))$
 $\quad\quad\quad (lang-inv-proj\ (\lambda((p,s),w).\ s)\ H0 \cap Cx)$
 $\langle proof \rangle$

Correctness of the list-level operation:

theorem *projH-effective*: $c2cl\ 'projH\ H0\ Cx = projH-cl\ H0\ (c2cl\ 'Cx)$
 $\langle proof \rangle$

lemma *c2cl-empty-rev*: $[] = c2cl\ []$ $\langle proof \rangle$

theorem *xdpnCL-effective*: $c2cl\ '(xdpnCL\ H) = xdpnCL-cl\ H$
 $\langle proof \rangle$

lemma *inv-proj-c1-effective*:
 $c2cl\ '\{cx.\ proj-c1\ cx \in C\} = lang-inv-proj\ fp-cl1\ (c2cl\ 'C)$
 $\langle proof \rangle$

theorem *xdpnC-effective*: $c2cl\ '(xdpnC\ C\ H) = xdpnC-cl\ (c2cl\ 'C)\ H$
 $\langle proof \rangle$

theorem *prehc-effective*:

$c2cl \text{ ' } prehc \Delta (lanch (H0,H)) C' = prehc-cl \Delta (H0,H) (c2cl \text{ ' } C')$

<proof>

8.3 What Does This Proof Tell You ?

In order to believe that our construction is effective, you have to believe that the RHS of Theorem *prehc-effective* is really effective.

The effectiveness of the *pre** - computation is shown in [2], and we have also an unpublished formal proof of the algorithm presented there. We are planning to adapt this proof to our model definition and the latest Isabelle version in near future, and then publish it.

The effectiveness of the involved automata computations is well-known. In a future version of this formalization, we plan to formalize or adopt an automata library and use it to generate executable code.

end

9 DPNs With Locks

theory *LockSem*

imports *DPN Semantics*

begin

In this theory, we define an extension of DPNs, where synchronization of the processes via a finite set of locks is allowed.

For this purpose, we assume that the rules are labeled with lock operations.

9.1 Model

- If a label has either no effect on locks, we allow it to be labeled by some other generic type *'L*. Otherwise, the label indicates either the acquisition or the release of a lock:

datatype $(\text{'L}, \text{'X}) \textit{lockstep} = LNone \text{'L} \mid LAcq \text{'X} \mid LRel \text{'X}$

- Abbreviation for the datatype of a DPN with locks:

types $(\text{'P}, \text{T}, \text{'L}, \text{'X}) \textit{ldpn} = (\text{'P}, \text{T}, (\text{'L}, \text{'X}) \textit{lockstep}) \textit{dpn}$

We encode DPNs with locks in a locale.

To save some case distinctions in proofs, we assume that only non-spawning rules are labeled with lock operations.

locale $LDPN = DPN +$

constrains

$\Delta :: (\text{'P}, \text{T}, \text{'L}, \text{'X} :: \textit{finite}) \textit{ldpn}$

assumes

spawn-no-locks: $\llbracket (p, \gamma \xrightarrow{a} ps, ws \# p', w) \in \Delta; \forall l. a = LNone \ l \implies P \rrbracket \implies P$

begin

lemma *snl-simps*[*simp*, *intro!*]:

$(p, \gamma \hookrightarrow_{LAcq} x \ ps, ws \ \# \ p', w) \notin \Delta$
 $(p, \gamma \hookrightarrow_{LRel} x \ ps, ws \ \# \ p', w) \notin \Delta$
 $\langle proof \rangle$

lemma *X-finite*: *finite* (*UNIV*::'*X set*) $\langle proof \rangle$

end

9.2 Interleaving Semantics

The following predicate models the step-relation on the set of allocated locks:

inductive *lock-valid* :: '*X set* \Rightarrow ('*L*, '*X*) *lockstep* \Rightarrow '*X set* \Rightarrow *bool* **where**

— A *LNone*-step does not change the set of allocated locks:

lw-none: *lock-valid* *X* (*LNone* *l*) *X* |

— A *LAcq*-step adds the acquired lock to the set of locks. It is only executable if the lock was not allocated before:

lw-acquire: *lock-valid* (*X* - {*x*}) (*LAcq* *x*) (*insert* *x* *X*) |

— A *LRel*-step removes the released lock from the set of locks. It is only executable if the lock was allocated before:

lw-release: *lock-valid* (*insert* *x* *X*) (*LRel* *x*) (*X* - {*x*})

lemma *lock-valid-simps*[*simp*]:

lock-valid *X* (*LNone* *l*) *X'* \longleftrightarrow *X* = *X'*

lock-valid *X* (*LAcq* *x*) *X'* \longleftrightarrow *X'* = *insert* *x* *X* \wedge *x* \notin *X*

lock-valid *X* (*LRel* *x*) *X'* \longleftrightarrow *X* = *insert* *x* *X'* \wedge *x* \notin *X'*

$\langle proof \rangle$

Configurations of the lock-sensitive step-relation consists of the list of processes and the set of currently acquired locks. Note that, at this point in the formalization, we do not make any assumptions on which process may release a lock, or on well-nestedness of locks.

That is, we allow a process releasing a lock that it has not acquired before, or locks being used in non-well-nestedness fashion.

However, in Section 10, we formalize such assumptions.

The lock-sensitive step-relation is the intersection of the original step-relation and the step-relation on allocated locks.

definition *ldpntr*

:: ('*P*, '*T*, '*L*, '*X*) *ldpn* \Rightarrow (('*P*, '*T*) *conf* \times '*X set*, ('*L*, '*X*) *lockstep*) *LTS*

where

ldpntr Δ = { ((*c*, *X*), *l*, (*c'*, *X'*)) . (*c*, *l*, *c'*) \in *dpntr* Δ \wedge *lock-valid* *X* *l* *X'* }

abbreviation *ldpntrc* Δ == *trcl* (*ldpntr* Δ)

lemma *ldpntr-subset*: ((*c*, *X*), *w*, (*c'*, *X'*)) \in *ldpntr* Δ \implies (*c*, *w*, *c'*) \in *dpntr* Δ

$\langle proof \rangle$

lemma *ldpntrc-subset*: ((*c*, *X*), *w*, (*c'*, *X'*)) \in *ldpntrc* Δ \implies (*c*, *w*, *c'*) \in *dpntrc* Δ

$\langle proof \rangle$

9.3 Tree Semantics

For the tree semantics, we only need to redefine the scheduler, such that it keeps track of the allocated locks.

— Abbreviation for type of execution trees and hedges with locks:

types $(\prime Q, \Gamma, \prime L, \prime X)$ *lex-tree* = $(\prime Q, \Gamma, (\prime L, \prime X)$ *lockstep*) *ex-tree*
types $(\prime Q, \Gamma, \prime L, \prime X)$ *lex-hedge* = $(\prime Q, \Gamma, (\prime L, \prime X)$ *lockstep*) *ex-hedge*

— The definition of the lock-sensitive scheduler is straightforward:

inductive *lsched*
 $:: (\prime Q, \Gamma, \prime L, \prime X)$ *lex-hedge* $\Rightarrow \prime X$ *set* $\Rightarrow (\prime L, \prime X)$ *lockstep list* \Rightarrow *bool*
where
lsched-final: *final* *h* \Longrightarrow *lsched* *h* *X* \square |
lsched-nospawn:
 \llbracket *lsched* (*h1*@*t*#*h2*) *X* *w*; *lock-valid* *X* *l* *X* $\rrbracket \Longrightarrow$
 \textit{lsched} (*h1*@(*NNOSPAWN* *l* *t*)#*h2*) *X* (*l*#*w*) |
lsched-spawn:
 \llbracket *lsched* (*h1*@*ts*#*t*#*h2*) *X* *w*; *lock-valid* *X* *l* *X* $\rrbracket \Longrightarrow$
 \textit{lsched} (*h1*@(*NSPAWN* *l* *ts* *t*)#*h2*) *X* (*l*#*w*)

— Obviously, a lock-sensitive schedule is also a schedule of the original scheduler:

lemma *lsched-is-sched*: *lsched* *h* *X* *ll* \Longrightarrow *sched* *h* *ll*
 \langle *proof* \rangle

9.4 Equivalence of Interleaving and Tree Semantics

— Straightforward adoption of proof of *sched-correct1*

lemma *lsched-correct1*:
 $((c, X), ll, (c', X')) \in \textit{ldpntrc}$ $\Delta \Longrightarrow \exists h. \textit{hsem}$ Δ *c* *h* *c'* \wedge *lsched* *h* *X* *ll*
 \langle *proof* \rangle

lemma *lsched-correct2*:
 \llbracket *lsched* *h* *X* *ll*; *hsem* Δ *c* *h* *c'* $\rrbracket \Longrightarrow \exists X'. ((c, X), ll, (c', X')) \in \textit{ldpntrc}$ Δ
 \langle *proof* \rangle

theorem *lsched-correct*:
 $(\exists X'. ((c, X), ll, (c', X')) \in \textit{ldpntrc}$ $\Delta) \longleftrightarrow (\exists h. \textit{hsem}$ Δ *c* *h* *c'* \wedge *lsched* *h* *X* *ll*)
 \langle *proof* \rangle

end

10 Well-Nestedness of Locks

theory *WellNested*
imports *DPN Semantics LockSem*
begin

Well-nestedness of locks is the property that no locks are re-acquired by the same process and a released locks is always the last one that was acquired and not yet released by the releasing process. Usually, these two properties are called non-reentrance and well-nestedness.

In this theory, we formulate a sufficient condition for well-nestedness, that regards every possible lock-insensitive run of the DPN from some initial configuration. We then define an equivalent condition on execution hedges.

Note that our condition may rule out DPNs where some non-well-nested runs are blocked by deadlocks or other lock-induced effects. However, important classes of programs, in particular programs that use locks in a block-structured way (like synchronized-blocks in Java), always satisfy our condition.

Further work required at this point is to formalize a program analysis or some sufficient conditions (like block-structured lock-acquisition [monitors]) for well-nestedness. We would then be able to prove some non-trivial DPNs to have well-nested configurations, thus giving a stronger indication that the well-nestedness assumption is correct. In the current state, we have no formal proof that the well-nestedness assumption is correct, i.e. an uncorrect well-nestedness assumption, e.g. a too strict one, would affect the scope of all our proofs that use this assumption. In the worst case, there would be no well-nested DPNs at all (or only trivial ones).

10.1 Well-Nestedness Condition on Paths

We first define the set of all paths that may occur from a process. We collect local paths and environment paths.

ppairs (q,w) *False* l means that there is a local path l from process (q,w) .

ppairs (q,w) *True* l means that we can reach a spawn step from process (q,w) that spawns a process having path " l ".

inductive *ppairs*

$:: ('P, \Gamma, 'L, 'X) \text{ldpn} \Rightarrow ('P, \Gamma) \text{pconf} \Rightarrow \text{bool} \Rightarrow ('L, 'X) \text{lockstep list} \Rightarrow \text{bool}$

for Δ **where**

ppairs-empty: *ppairs* Δ (q,w) *False* $[]$ |

ppairs-prepend1:

$\llbracket (q, \gamma \hookrightarrow_a q', w) \in \Delta; \text{ppairs } \Delta (q', w @ r) \text{ False } l \rrbracket \Longrightarrow$
 $\text{ppairs } \Delta (q, \gamma \# r) \text{ False } (a \# l)$ |

ppairs-mvenv1:

$\llbracket (q, \gamma \hookrightarrow_a q', w) \in \Delta; \text{ppairs } \Delta (q', w @ r) \text{ True } l \rrbracket \Longrightarrow$
 $\text{ppairs } \Delta (q, \gamma \# r) \text{ True } l$ |

ppairs-prepend2:

$\llbracket (q, \gamma \hookrightarrow_a qs, ws \# q', w) \in \Delta; \text{ppairs } \Delta (q', w @ r) \text{ False } l \rrbracket \Longrightarrow$
 $\text{ppairs } \Delta (q, \gamma \# r) \text{ False } (a \# l)$ |

ppairs-mvenv2: $\llbracket (q, \gamma \hookrightarrow_a qs, ws \# q', w) \in \Delta; \text{ppairs } \Delta (q', w @ r) \text{ True } l \rrbracket \Longrightarrow$

$$\begin{aligned} & \text{ppairs } \Delta (q, \gamma \# r) \text{ True } l \mid \\ \text{ppairs-genenv: } & \llbracket (q, \gamma \hookrightarrow_a qs, ws \# q', w) \in \Delta; \text{ppairs } \Delta (qs, ws) x l \rrbracket \implies \\ & \text{ppairs } \Delta (q, \gamma \# r) \text{ True } l \end{aligned}$$

This function checks whether a path is well-nested by using a lock stack.

fun $wn-p :: ('L, 'X) \text{lockstep list} \Rightarrow 'X \text{list} \Rightarrow \text{bool}$ **where**
 $wn-p \ [] \ \mu = \text{distinct } \mu \mid$
 $wn-p (LAcq \ x \# l) \ \mu \longleftrightarrow wn-p \ l (x \# \mu) \mid$
 $wn-p (LRel \ x \# l) \ \mu \longleftrightarrow (\exists \mu'. \ \mu = x \# \mu' \wedge x \notin \text{set } \mu' \wedge wn-p \ l \ \mu') \mid$
 $wn-p (-\# l) \ \mu \longleftrightarrow wn-p \ l \ \mu$

A process π is defined to be well-nested w.r.t. some initial lock stack μ if all reachable path – local paths and environment paths – are well-nested.

definition $wn-\pi \ \Delta \ \pi \ \mu ==$
 $\text{case } \pi \text{ of } (p, w) \Rightarrow$
 $\forall l. (\text{ppairs } \Delta (p, w) \text{ False } l \longrightarrow wn-p \ l \ \mu) \wedge$
 $(\text{ppairs } \Delta (p, w) \text{ True } l \longrightarrow wn-p \ l \ [])$

Introduction and elimination rules for $wn-\pi$

lemma $wn-\pi I$:
 \llbracket
 $\quad \text{!!}l. \text{ppairs } \Delta (q, w) \text{ False } l \implies wn-p \ l \ \mu;$
 $\quad \text{!!}l. \text{ppairs } \Delta (q, w) \text{ True } l \implies wn-p \ l \ []$
 $\rrbracket \implies wn-\pi \ \Delta (q, w) \ \mu$
 $\langle \text{proof} \rangle$

lemma $wn-\pi E$:
 $\llbracket wn-\pi \ \Delta (q, w) \ \mu;$
 $\quad \llbracket$
 $\quad \quad \text{!!}l. \text{ppairs } \Delta (q, w) \text{ False } l \implies wn-p \ l \ \mu;$
 $\quad \quad \text{!!}l. \text{ppairs } \Delta (q, w) \text{ True } l \implies wn-p \ l \ []$
 $\quad \rrbracket \implies P$
 $\rrbracket \implies P$
 $\langle \text{proof} \rangle$

We have set up the definitions such that well-nestedness w.r.t a lock stack implies distinctness of this lock stack.

lemma $wn-p\text{-distinct}$: $wn-p \ l \ \mu \implies \text{distinct } \mu$
 $\langle \text{proof} \rangle$

lemma $wn-\pi\text{-distinct}$: $wn-\pi \ \Delta \ \pi \ \mu \implies \text{distinct } \mu$
 $\langle \text{proof} \rangle$

Well-nestedness is preserved by steps:

lemma $wn-\pi\text{-none}$:
 $\llbracket (q, \gamma \hookrightarrow_{(LNone \ l)} q', w) \in \Delta; wn-\pi \ \Delta (q, \gamma \# r) \ \mu \rrbracket \implies wn-\pi \ \Delta (q', w @ r) \ \mu$
 $\langle \text{proof} \rangle$

lemma (in $LDPN$) $wn-\pi\text{-spawn1}$:
 $\llbracket (q, \gamma \hookrightarrow_a qs, ws \# q', w) \in \Delta; wn-\pi \ \Delta (q, \gamma \# r) \ \mu \rrbracket \implies wn-\pi \ \Delta (q', w @ r) \ \mu$

$\langle proof \rangle$
lemma *wn- π -spawn2*:
 $\llbracket (q, \gamma \hookrightarrow_a qs, ws \# q', w) \in \Delta; wn-\pi \Delta (q, \gamma \# r) \mu \rrbracket \implies wn-\pi \Delta (qs, ws) \llbracket \rrbracket$
 $\langle proof \rangle$
lemma *wn- π -acq*:
 $\llbracket (q, \gamma \hookrightarrow_{LAcq} x q', w) \in \Delta; wn-\pi \Delta (q, \gamma \# r) \mu \rrbracket \implies wn-\pi \Delta (q', w @ r) (x \# \mu)$
 $\langle proof \rangle$
lemma *wn- π -rel*:
assumes $A: (q, \gamma \hookrightarrow_{LRel} x q', w) \in \Delta \quad wn-\pi \Delta (q, \gamma \# r) \mu$ **and**
 $C: !!\mu'. \llbracket \mu = x \# \mu'; x \notin set \mu'; wn-\pi \Delta (q', w @ r) \mu' \rrbracket \implies P$
shows P
 $\langle proof \rangle$
lemma (in LDPN) *wn- π -preserve*:
 $\llbracket (q, \gamma \hookrightarrow_l q', w) \in \Delta; wn-\pi \Delta (q, \gamma \# r) xs; !!xs'. wn-\pi \Delta (q', w @ r) xs' \implies P \rrbracket \implies P$
 $\llbracket (q, \gamma \hookrightarrow_l qs, ws \# q', w) \in \Delta; wn-\pi \Delta (q, \gamma \# r) xs; !!xs'. \llbracket wn-\pi \Delta (q', w @ r) xs'; wn-\pi \Delta (qs, ws) \llbracket \rrbracket \rrbracket \implies P$
 $\llbracket \rrbracket \implies P$
 $\langle proof \rangle$

10.2 Well-Nestedness of Configurations

The locks of a list of lock stacks

abbreviation *locks- μ* :: $'X list list \Rightarrow 'X set$ **where**
 $locks-\mu \mu == list-collect-set set \mu$

A configuration $c = \pi_1 \dots \pi_n$ is well-nested w.r.t. a list $\mu = s_1 \dots s_n$ of lock stacks ($wn-h h \mu$), iff all π_i are well-nested w.r.t. stack s_i and μ is consistent, i.e. contains no duplicate locks.

fun *wn-c where*

$wn-c \Delta \llbracket \rrbracket \longleftrightarrow True$ |
 $wn-c \Delta (\pi \# c) (xs \# \mu) \longleftrightarrow$
 $wn-c \Delta c \mu \wedge set xs \cap locks-\mu \mu = \{\}$ $\wedge wn-\pi \Delta \pi xs$ |
 $wn-c \Delta - - \longleftrightarrow False$

10.2.1 Auxilliary Lemmas about *wn-c*

lemma *wn-c-simps[simp]*:

$wn-c \Delta c \llbracket \rrbracket \longleftrightarrow c = \llbracket \rrbracket$
 $wn-c \Delta \llbracket \rrbracket \mu \longleftrightarrow \mu = \llbracket \rrbracket$
 $\langle proof \rangle$

lemma *wn-c-length*: $wn-c \Delta c \mu \implies length c = length \mu$
 $\langle proof \rangle$

lemma *wn-c-prepend-c*:

$$\begin{aligned} & \llbracket wn-c \Delta (\pi\#c) \mu; \\ & \quad !!xs \ \mu'. \llbracket \mu=xs\#\mu'; wn-c \Delta c \ \mu'; \\ & \quad \quad \quad set \ xs \cap \ locks-\mu \ \mu' = \{\}; wn-\pi \Delta \pi \ xs \\ & \quad \quad \quad \rrbracket \implies P \\ & \rrbracket \implies P \\ & \langle proof \rangle \end{aligned}$$

lemma *wn-c-prepend- μ* :

$$\begin{aligned} & \llbracket wn-c \Delta c \ (xs\#\mu); \\ & \quad !!\pi \ c'. \llbracket c=\pi\#c'; wn-c \Delta c' \ \mu; \\ & \quad \quad \quad set \ xs \cap \ locks-\mu \ \mu = \{\}; wn-\pi \Delta \pi \ xs \\ & \quad \quad \quad \rrbracket \implies P \\ & \rrbracket \implies P \\ & \langle proof \rangle \end{aligned}$$

lemma *wn-c-append-c-helper*:

assumes

$$\begin{aligned} & A: wn-c \Delta c \ \mu \quad c1@c2=c \ \mathbf{and} \\ & C: !!\mu1 \ \mu2. \llbracket \mu=\mu1@\mu2 \wedge wn-c \Delta c1 \ \mu1 \wedge wn-c \Delta c2 \ \mu2 \wedge \\ & \quad \quad \quad \locks-\mu \ \mu1 \cap \ locks-\mu \ \mu2 = \{\} \\ & \quad \quad \quad \rrbracket \implies P \end{aligned}$$

shows P

$\langle proof \rangle$

lemma *wn-c-append-c*:

$$\begin{aligned} & \llbracket wn-c \Delta (c1@c2) \ \mu; \\ & \quad !!\mu1 \ \mu2. \llbracket \mu=\mu1@\mu2 \wedge wn-c \Delta c1 \ \mu1 \wedge wn-c \Delta c2 \ \mu2 \wedge \\ & \quad \quad \quad \locks-\mu \ \mu1 \cap \ locks-\mu \ \mu2 = \{\} \rrbracket \implies P \\ & \rrbracket \implies P \\ & \langle proof \rangle \end{aligned}$$

lemma *wn-c-append- μ -helper*:

assumes

$$\begin{aligned} & A: wn-c \Delta c \ \mu \quad \mu1@\mu2=\mu \ \mathbf{and} \\ & C: !!c1 \ c2. \llbracket c=c1@c2 \wedge wn-c \Delta c1 \ \mu1 \wedge wn-c \Delta c2 \ \mu2 \wedge \\ & \quad \quad \quad \locks-\mu \ \mu1 \cap \ locks-\mu \ \mu2 = \{\} \rrbracket \implies P \end{aligned}$$

shows P

$\langle proof \rangle$

lemma *wn-c-append- μ* :

$$\begin{aligned} & \llbracket wn-c \Delta c \ (\mu1@\mu2); \\ & \quad !!c1 \ c2. \llbracket c=c1@c2 \wedge wn-c \Delta c1 \ \mu1 \wedge wn-c \Delta c2 \ \mu2 \wedge \\ & \quad \quad \quad \locks-\mu \ \mu1 \cap \ locks-\mu \ \mu2 = \{\} \rrbracket \implies P \\ & \rrbracket \implies P \\ & \langle proof \rangle \end{aligned}$$

lemma *wn-c-appendI*:

$$\begin{aligned} & \llbracket wn-c \Delta c1 \ \mu1; wn-c \Delta c2 \ \mu2; \locks-\mu \ \mu1 \cap \ locks-\mu \ \mu2 = \{\} \rrbracket \implies \\ & \quad wn-c \Delta (c1@c2) \ (\mu1@\mu2) \end{aligned}$$

$\langle proof \rangle$

lemma *wn-c-prependI*:

$\llbracket wn-\pi \Delta \pi \ xs; wn-c \Delta c \ \mu; set \ xs \cap locks-\mu \ \mu = \{\} \rrbracket \implies wn-c \Delta (\pi \# c) (xs \# \mu)$
 $\langle proof \rangle$

lemma *wn-c-singlecE*: $\llbracket wn-c \Delta [\pi] \ \mu; !!xs. \llbracket \mu=[xs]; wn-\pi \Delta \pi \ xs \rrbracket \implies P \rrbracket \implies P$

$\langle proof \rangle$

lemma *wn-c-split-aux*:

assumes

WN: $wn-c \Delta c \ \mu$ **and**

HFM $T[simp]$: $c=c1 @ \pi \# c2$ **and**

$C: !!\mu1 \ xs \ \mu2. \llbracket \mu=\mu1 @ xs \# \mu2; wn-\pi \Delta \pi \ xs; wn-c \Delta c1 \ \mu1; wn-c \Delta c2 \ \mu2;$
 $locks-\mu \ \mu1 \cap set \ xs = \{\}; locks-\mu \ \mu1 \cap locks-\mu \ \mu2 = \{\};$
 $set \ xs \cap locks-\mu \ \mu2 = \{\}$
 $\rrbracket \implies P$

shows P

$\langle proof \rangle$

Well-nestedness of configurations is preserved by lock-sensitive steps.

lemma (in *LDPN*) *wnc-preserve-singlestep*:

assumes

$A: ((c, locks-\mu \ \mu), l, (c', X')) \in ldpntr \Delta \quad wn-c \Delta c \ \mu$ **and**

$C: !!\mu'. \llbracket X'=locks-\mu \ \mu'; wn-c \Delta c' \ \mu' \rrbracket \implies P$

shows P

$\langle proof \rangle$

lemma (in *LDPN*) *wnc-preserve*:

assumes $A: ((c, locks-\mu \ \mu), ll, (c', X')) \in ldpntrc \Delta \quad wn-c \Delta c \ \mu$ **and**

$C: !!\mu'. \llbracket X'=locks-\mu \ \mu'; wn-c \Delta c' \ \mu' \rrbracket \implies P$

shows P

$\langle proof \rangle$

10.3 Well-Nestedness Condition on Trees

Now we define well-nestedness on scheduling trees. Note that scheduling trees that contain spawn steps with locks interaction are not well-nested.

We define two equivalent formulations of well-nestedness of a tree:

fun *wn-t'* :: (P, T, L, X) *lex-tree* $\implies X$ *list* $\implies bool$ **where**

$wn-t' (NLEAF \ \pi) \ \mu \longleftrightarrow distinct \ \mu \mid$

$wn-t' (NNOSPAWN (LNone \ l) \ t) \ \mu \longleftrightarrow wn-t' \ t \ \mu \mid$

$wn-t' (NSPAWN (LNone \ l) \ ts \ t) \ \mu \longleftrightarrow wn-t' \ t \ \mu \wedge wn-t' \ ts \ \square \mid$

$wn-t' (NNOSPAWN (LAcq \ x) \ t) \ \mu \longleftrightarrow wn-t' \ t \ (x \# \mu) \wedge x \notin set \ \mu \mid$

$wn-t' (NNOSPAWN (LRel \ x) \ t) \ \mu \longleftrightarrow$

$(\exists \mu'. \mu = x \# \mu' \wedge wn-t' \ t \ \mu' \wedge x \notin set \ \mu') \mid$

$wn-t' \ - \ - \longleftrightarrow False$

inductive *wn-t* :: (P, T, L, X) *lex-tree* $\implies X$ *list* $\implies bool$ **where**

$$\begin{aligned}
& \text{distinct } \mu \implies \text{wn-t } (NLEAF \ \pi) \ \mu \mid \\
& \text{wn-t } t \ \mu \implies \text{wn-t } (NNOSPAWN \ (LNone \ l) \ t) \ \mu \mid \\
& \llbracket \text{wn-t } t \ \mu; \text{wn-t } ts \ \square \rrbracket \implies \text{wn-t } (NSPAWN \ (LNone \ l) \ ts \ t) \ \mu \mid \\
& \llbracket \text{wn-t } t \ (x\#\mu); x \notin \text{set } \mu \rrbracket \implies \text{wn-t } (NNOSPAWN \ (LAcq \ x) \ t) \ \mu \mid \\
& \llbracket \text{wn-t } t \ \mu; x \notin \text{set } \mu \rrbracket \implies \text{wn-t } (NNOSPAWN \ (LRel \ x) \ t) \ (x\#\mu)
\end{aligned}$$

inductive lock-valid-xs where

$$\begin{aligned}
& \text{distinct } xs \implies \text{lock-valid-xs } (LNone \ l) \ xs \ xs \mid \\
& \llbracket \text{distinct } xs; x \notin \text{set } xs \rrbracket \implies \text{lock-valid-xs } (LRel \ x) \ (x\#xs) \ xs \mid \\
& \llbracket \text{distinct } xs; x \notin \text{set } xs \rrbracket \implies \text{lock-valid-xs } (LAcq \ x) \ xs \ (x\#xs)
\end{aligned}$$

The two formulations of well-nestedness of trees are, indeed, equivalent:

lemma wnt-eq-wnt': $\text{wn-t } t \ \mu = \text{wn-t}' \ t \ \mu$
 $\langle \text{proof} \rangle$

Well-nestedness of trees also implies distinctness of the lock stacks

lemma wnt-distinct: $\text{wn-t } t \ \mu \implies \text{distinct } \mu$
 $\langle \text{proof} \rangle$

lemma wnt-distinct': $\text{wn-t}' \ t \ ms \implies \text{distinct } ms$
 $\langle \text{proof} \rangle$

lemma all-t-wnt-distinct: $\forall t \ c'. \text{tsem } \Delta \ (q,w) \ t \ c' \longrightarrow \text{wn-t } t \ \mu \implies \text{distinct } \mu$
 $\langle \text{proof} \rangle$

10.4 Well-Nestedness of Hedges

The well-nestedness property of a hedge expresses that each tree is well-nested, and the allocated locks of the trees are consistent.

Consistency of a list of lock stacks. $\mu = s_1 \dots s_n$ is consistent, iff all s_i are distinct and $\forall i \ j. i \neq j \longrightarrow \text{set } s_i \cap \text{set } s_j = \{\}$.

fun cons- μ :: $'X \text{ list list} \Rightarrow \text{bool}$ **where**

$$\begin{aligned}
& \text{cons-}\mu \ \square \longleftrightarrow \text{True} \mid \\
& \text{cons-}\mu \ (xs\#\mu) \longleftrightarrow \text{cons-}\mu \ \mu \wedge \text{distinct } xs \wedge \text{set } xs \cap \text{locks-}\mu \ \mu = \{\}
\end{aligned}$$

A hedge $h = t_1 \dots t_n$ is well-nested w.r.t. a list $\mu = s_1 \dots s_n$ of lock stacks ($\text{wn-h } h \ \mu$), iff all t_i are well-nested w.r.t. stack s_i and μ is consistent.

fun wn-h where

$$\begin{aligned}
& \text{wn-h } \square \ \square \longleftrightarrow \text{True} \mid \\
& \text{wn-h } (t\#h) \ (xs\#\mu) \longleftrightarrow \text{wn-h } h \ \mu \wedge \text{set } xs \cap \text{locks-}\mu \ \mu = \{\} \wedge \text{wn-t}' \ t \ xs \mid \\
& \text{wn-h } - \ - \longleftrightarrow \text{False}
\end{aligned}$$

lemma cons- μ -append[simp]:

$$\text{cons-}\mu \ (\mu1 @ \mu2) \longleftrightarrow \text{cons-}\mu \ \mu1 \wedge \text{cons-}\mu \ \mu2 \wedge \text{locks-}\mu \ \mu1 \cap \text{locks-}\mu \ \mu2 = \{\}$$

$\langle \text{proof} \rangle$

10.4.1 Auxilliary Lemmas about wn-h

lemma wn-h-simps[simp]:

lemma *wn-h-appendI*:

$$\llbracket \text{wn-h } h1 \ \mu1; \text{wn-h } h2 \ \mu2; \text{locks-}\mu \ \mu1 \cap \text{locks-}\mu \ \mu2 = \{\} \rrbracket \Longrightarrow \\ \text{wn-h } (h1@h2) \ (\mu1@mu2) \\ \langle \text{proof} \rangle$$

lemma *wn-h-prependI*:

$$\llbracket \text{wn-t}' \ t \ xs; \text{wn-h } h \ \mu; \text{set } xs \cap \text{locks-}\mu \ \mu = \{\} \rrbracket \Longrightarrow \text{wn-h } (t\#h) \ (xs\#\mu) \\ \langle \text{proof} \rangle$$

lemma *wn-h-singlehE*: $\llbracket \text{wn-h } [t] \ \mu; !!xs. \llbracket \mu=[xs]; \text{wn-t}' \ t \ xs \rrbracket \Longrightarrow P \rrbracket \Longrightarrow P$
 $\langle \text{proof} \rangle$

Auxilliary lemma to split the list of lock-stacks w.r.t. to that a hedge is well-nested by some tree in that hedge.

lemma *wn-h-split-aux*:

assumes

WN: *wn-h* *h* μ **and**

HFMT[*simp*]: $h=h1@t\#h2$ **and**

C: $!!\mu1 \ xs \ \mu2. \llbracket$

$$\mu=\mu1@xs\#\mu2;$$

$$\text{wn-t}' \ t \ xs; \text{wn-h } h1 \ \mu1; \text{wn-h } h2 \ \mu2;$$

$$\text{locks-}\mu \ \mu1 \cap \text{set } xs = \{\}; \text{locks-}\mu \ \mu1 \cap \text{locks-}\mu \ \mu2 = \{\};$$

$$\text{set } xs \cap \text{locks-}\mu \ \mu2 = \{\}$$

$$\rrbracket \Longrightarrow P$$

shows *P*

$\langle \text{proof} \rangle$

10.4.2 Relation to Path Condition

We show that the notion of well-nestedness on paths and trees are equivalent, i.e. a configuration is well-nested w.r.t. a lock stack μ if and only if all trees from that configuration are well-nested w.r.t. μ .

A process π is well-nested w.r.t. some stack of locks μ , if all its execution trees are well-nested w.r.t. μ :

definition *wn- π -t* $\Delta \ \pi \ xs == (\forall t \ c'. \text{tsem } \Delta \ \pi \ t \ c' \longrightarrow \text{wn-t } t \ xs)$

definition *wn-c-h* $\Delta \ c \ \mu == (\forall h \ c'. \text{hsem } \Delta \ c \ h \ c' \longrightarrow \text{wn-h } h \ \mu)$

lemma *wn- π -tI*[*intro?*]: $\llbracket !!t \ c'. \text{tsem } \Delta \ \pi \ t \ c' \Longrightarrow \text{wn-t } t \ xs \rrbracket \Longrightarrow \text{wn-}\pi\text{-t } \Delta \ \pi \ xs$
 $\langle \text{proof} \rangle$

lemma *wn-c-hI*[*intro?*]: $\llbracket !!h \ c'. \text{hsem } \Delta \ c \ h \ c' \Longrightarrow \text{wn-h } h \ \mu \rrbracket \Longrightarrow \text{wn-c-h } \Delta \ c \ \mu$
 $\langle \text{proof} \rangle$

lemma *wn- π -t-distinct*: $\text{wn-}\pi\text{-t } \Delta \ \pi \ \mu \Longrightarrow \text{distinct } \mu$
 $\langle \text{proof} \rangle$

lemma *wn-c-h-prepend1*: **assumes** A : $wn-c-h \Delta (\pi \# c) (xs \# \mu)$
shows $wn-\pi-t \Delta \pi xs \quad wn-c-h \Delta c \mu \quad set\ xs \cap locks-\mu \mu = \{\}$
 $\langle proof \rangle$

lemma *wn-c-h-prepend2*:
 $\llbracket wn-\pi-t \Delta \pi xs; wn-c-h \Delta c \mu; set\ xs \cap locks-\mu \mu = \{\} \rrbracket \implies$
 $wn-c-h \Delta (\pi \# c) (xs \# \mu)$
 $\langle proof \rangle$

lemma *wn-c-h-prepend[simp]*:
 $wn-c-h \Delta (\pi \# c) (xs \# \mu) \longleftrightarrow$
 $wn-\pi-t \Delta \pi xs \wedge wn-c-h \Delta c \mu \wedge set\ xs \cap locks-\mu \mu = \{\}$
 $\langle proof \rangle$

lemma *wn-c-h-empty[simp]*: $wn-c-h \Delta c [] \longleftrightarrow (c=[])$ $\langle proof \rangle$

lemma *wn-c-h-prepend-c*:
 $\llbracket wn-c-h \Delta (\pi \# c) \mu;$
 $!!xs \ \mu'. \llbracket \mu = xs \# \mu'; wn-\pi-t \Delta \pi xs; wn-c-h \Delta c \mu';$
 $set\ xs \cap locks-\mu \mu' = \{\} \rrbracket \implies P$
 $\rrbracket \implies P$
 $\langle proof \rangle$

lemma *wn-c-h-simps[simp]*: $wn-c-h \Delta [] \mu \longleftrightarrow (\mu=[])$
 $\langle proof \rangle$

lemma (in *LDPN*) *wn π 2wnt*: $\llbracket tsem \Delta (q,w) t c'; wn-\pi \Delta (q,w) \mu \rrbracket \implies wn-t\ t\ \mu$
 $\langle proof \rangle$

lemma (in *LDPN*) *wnt2wnp*:
 $\llbracket ppairs \Delta (q,w) en\ l; \forall t\ c'. tsem \Delta (q,w) t\ c' \longrightarrow wn-t\ t\ \mu \rrbracket \implies$
 $(\neg en \longrightarrow wn-p\ l\ \mu) \wedge (en \longrightarrow wn-p\ l\ [])$
 $\langle proof \rangle$

theorem (in *LDPN*) *wn π -eq-wn π t*: $wn-\pi \Delta \pi \mu \longleftrightarrow wn-\pi-t \Delta \pi \mu$ $\langle proof \rangle$

theorem (in *LDPN*) *wnc-eq-wnc-h*: $wn-c \Delta c \mu \longleftrightarrow wn-c-h \Delta c \mu$
 $\langle proof \rangle$

10.5 Well-Nestedness and Tree Scheduling

In this section we show that well-nestedness is invariant under the tree scheduling relation. This is important, as it shows that we cannot reach non-well-nested trees from well-nested ones.

lemma *wnt-preserve-nospawn*:
 $\llbracket lock-valid\ (set\ xs)\ l\ X'; wn-t'\ (NNOSPAWN\ l\ t)\ xs \rrbracket \implies$
 $\exists xs'. X' = set\ xs' \wedge lock-valid-x\ l\ xs\ xs' \wedge wn-t'\ t\ xs'$
 $\langle proof \rangle$

lemma *wn-h-preserve-nospawn*:

$\llbracket \text{lock-valid } (\text{locks-}\mu \ \mu) \ l \ X'; \text{wn-h } (h1@(\text{NNOSPAWN } l \ t)\#h2) \ \mu \rrbracket \implies$
 $\exists \mu'. X'=\text{locks-}\mu \ \mu' \wedge \text{wn-h } (h1@t\#h2) \ \mu'$
<proof>

All-in-one lemma for reasoning about a non-spawning step on a well-nested hedge. In words: If we make a non-spawning step on a well-nested hedge:

- We can split the list of lock stacks according to the tree that made the step,
- The lock stack of the tree that made the step changes according to the label (cf. *lock-valid-xs*),
- And the resulting hedge is well-nested w.r.t. the new locks, too.

lemma *wn-h-split-nospawn*:

assumes

A: *lock-valid* (*locks-μ μ*) *l Xh* *wn-h* (*h1@(\text{NNOSPAWN } l \ t)\#h2*) *μ* **and**

C: $\exists \mu1 \ xs \ \mu2 \ xsh. \llbracket$

$\mu=\mu1@xs\#\mu2;$

$Xh=\text{locks-}\mu \ \mu1 \cup \text{set } xsh \cup \text{locks-}\mu \ \mu2;$

lock-valid-xs *l xs xsh*;

wn-t' (*NNOSPAWN l t*) *xs*;

wn-t' *t xsh*;

wn-h *h1 μ1*;

wn-h *h2 μ2*;

wn-h (*h1@t\#h2*) (*μ1@xsh\#\mu2*);

locks-μ μ1 \cap *set xs* = {};

locks-μ μ1 \cap *set xsh* = {};

locks-μ μ1 \cap *locks-μ μ2* = {};

locks-μ μ2 \cap *set xs* = {};

locks-μ μ2 \cap *set xsh* = {}

$\rrbracket \implies P$

shows *P*

<proof>

lemma *wn-h-preserve-spawn*:

$\llbracket \text{lock-valid } (\text{locks-}\mu \ \mu) \ l \ X'; \text{wn-h } (h1@(\text{NSPAWN } l \ ts \ t)\#h2) \ \mu \rrbracket \implies$
 $\exists \mu'. X'=\text{locks-}\mu \ \mu' \wedge \text{wn-h } (h1@ts\#t\#h2) \ \mu'$

<proof>

lemma *wn-h-preserve-spawn'*:

$\llbracket \text{lock-valid } (\text{locks-}\mu \ \mu) \ l \ X'; \text{wn-h } (h1@(\text{NSPAWN } l \ ts \ t)\#h2) \ \mu \rrbracket \implies$
 $\exists \mu1 \ xs \ \mu2. \mu=\mu1@xs\#\mu2 \wedge X'=\text{locks-}\mu \ \mu1 \cup \text{set } xs \cup \text{locks-}\mu \ \mu2 \wedge$
 $\text{wn-h } (h1@ts\#t\#h2) \ (\mu1@[]\#xs\#\mu2)$

<proof>

lemma *wn-h-preserve-rel*:
 $\llbracket (h, l, h') \in \text{sched-rel}; \text{lock-valid } (\text{locks-}\mu \ \mu) \ l \ X'; \text{wn-h } h \ \mu;$
 $\quad \llbracket \mu'. \llbracket X' = \text{locks-}\mu \ \mu'; \text{wn-h } h' \ \mu \rrbracket \implies P$
 $\rrbracket \implies P$
 $\langle \text{proof} \rangle$

lemma *wn-h-spawn-simps[simp]*:
 $\neg \text{wn-h } (h \ @ \ (\text{NSPAWN } (\text{LAcq } x) \ ts \ t) \ # \ h') \ \mu$
 $\neg \text{wn-h } (h \ @ \ (\text{NSPAWN } (\text{LRel } x) \ ts \ t) \ # \ h') \ \mu$
 $\langle \text{proof} \rangle$

lemmas *wn-h-spawn-simps-add[simp]* =
 $\text{wn-h-spawn-simps}[\text{where } h=[], \text{simplified}]$
 $\text{wn-h-spawn-simps}[\text{where } h=[tx], \text{simplified}, \text{standard}]$

lemma *wn-h-spawn-imp-LNoneE*:
 $\llbracket \text{wn-h } (h \ @ \ (\text{NSPAWN } l \ ts \ t) \ # \ h') \ \mu; \llbracket \! \! \! ll. \ l = \text{LNone } ll \implies P \rrbracket \implies P$
 $\langle \text{proof} \rangle$

end

11 Acquisition Structures

theory *Acqh*
imports *Main Semantics WellNested SpecialLemmas*
begin

11.1 Utilities

11.1.1 Combinators for *option-datatype*

Extending a function to option datatype, where *None* indicates failure

fun *opt-ext1* :: ('a \Rightarrow 'b option) \Rightarrow 'a option \Rightarrow 'b option **where**
 $\text{opt-ext1 } f \ \text{None} = \text{None} \mid$
 $\text{opt-ext1 } f \ (\text{Some } x) = f \ x$

fun *opt-ext2* :: ('a \Rightarrow 'b \Rightarrow 'c option) \Rightarrow 'a option \Rightarrow 'b option \Rightarrow 'c option
where
 $\text{opt-ext2 } f \ \text{None} \ - = \text{None} \mid$
 $\text{opt-ext2 } f \ - \ \text{None} = \text{None} \mid$
 $\text{opt-ext2 } f \ (\text{Some } x) \ (\text{Some } y) = f \ x \ y$

lemma *opt-ext2-simps[simp]*:
 $\text{opt-ext2 } f \ x \ \text{None} = \text{None} \langle \text{proof} \rangle$

lemma *opt-ext2-alt*:

$$\text{opt-ext2 } f \ x \ y = (\text{case } x \text{ of } \begin{array}{l} \text{None} \Rightarrow \text{None} \mid \\ \text{Some } xx \Rightarrow (\text{case } y \text{ of } \begin{array}{l} \text{None} \Rightarrow \text{None} \mid \\ \text{Some } yy \Rightarrow f \ xx \ yy \end{array} \end{array})$$

$$\langle \text{proof} \rangle$$

11.2 Acquisition Structures

Acquisition structures are an abstraction of scheduling trees, that are sufficient to decide whether a tree is schedulable. The basic concept of acquisition structures was invented by Kahlon et al. [4, 3] as abstraction of a linear execution of a single pushdown system. We extend this concept here to scheduling trees of DPNs.

An acquisition or release history is a partial map from locks to set of locks. This is the same representation as in [3]. Another, equivalent representation is as a set of locks and a graph on locks.

An acquisition structure is a triple of a release history, a set of locks and an acquisition history.

types

$$'X \text{ ah} = 'X \Rightarrow 'X \text{ set option}$$

$$'X \text{ as} = 'X \text{ ah} \times 'X \text{ set} \times 'X \text{ ah}$$

This is a collection of the common split-lemmas required when reasoning about acquisition histories

lemmas *eahl-splits* = *option.split-asm list.split-asm prod.split-asm split-if-asm*

11.2.1 Parallel Composition

fun *as-comp* :: $'X \text{ as} \Rightarrow 'X \text{ as} \Rightarrow 'X \text{ as option}$ **where**

$$\text{as-comp } (l, u, e) \ (l', u', e') = (\text{if } \text{dom } l \cap \text{dom } l' = \{\} \wedge \text{dom } e \cap \text{dom } e' = \{\} \text{ then } \text{Some } (l++l', u \cup u', e++e') \text{ else } \text{None})$$

definition *as-comp-op*

$$:: 'X \text{ as option} \Rightarrow 'X \text{ as option} \Rightarrow 'X \text{ as option (infixr } \parallel \text{ 56) where}$$

$$\text{op } \parallel == \text{opt-ext2 as-comp}$$

lemma *as-comp-op-simps[simp]*:

$None \parallel x = None$
 $x \parallel None = None$
 $Some\ a \parallel Some\ b = as-comp\ a\ b$
 $\langle proof \rangle$

lemma *as-comp-assoc-helper*:

$(Some\ x \parallel Some\ y) \parallel Some\ z = Some\ x \parallel Some\ y \parallel Some\ z$
 $\langle proof \rangle$

lemma *as-comp-assoc*: $(x \parallel y) \parallel z = x \parallel y \parallel z$

$\langle proof \rangle$

interpretation *as-comp-acz*: $ACIZ[op \parallel Some\ (empty, \{\}, empty)\ None]$

$\langle proof \rangle$

lemma *as-comp-SomeE*:

$\llbracket h1 \parallel h2 = Some\ (l, u, e);$
 $!!!l1\ u1\ e1\ l2\ u2\ e2. \llbracket h1=Some\ (l1, u1, e1); h2=Some\ (l2, u2, e2);$
 $dom\ l1 \cap dom\ l2 = \{\}; dom\ e1 \cap dom\ e2 = \{\};$
 $l=l1++l2; u=u1 \cup u2; e=e1++e2$
 $\rrbracket \implies P$

$\rrbracket \implies P$

$\langle proof \rangle$

11.2.2 Acquisition Structures of Scheduling Trees and Hedges

This function adds a set of locks to every entry in a release history. On graph interpretation, this corresponds to adding edges from any initially released lock to any lock in X .

definition *l-add-use* :: $'X\ ah \Rightarrow 'X\ set \Rightarrow 'X\ ah$ **where**

l-add-use $l\ X == \lambda x. case\ l\ x\ of\ None \Rightarrow None \mid Some\ Y \Rightarrow Some\ (Y \cup X)$

This function removes an initially released lock x from the release history. On graph interpretation, this corresponds to removing the node x from the graph.

definition *l-remove* :: $'X\ ah \Rightarrow 'X \Rightarrow 'X\ ah$ **where**

l-remove $l\ x == \lambda y. if\ y=x\ then\ None\ else\ l\ y$

The acquisition history of a tree is defined inductively over the tree structure. Note that we assume that spawn steps have no lock operation. For spawn steps with an operation on locks, the acquisition structure is defined to be *None*. We further assume that a tree contains no two initial releases of the same lock. In this case, its acquisition structure has no meaning any more. However, if an execution tree contains two final acquisitions of the same lock, its acquisition structure is defined to be *None*.

Intuitively, the release history maps all locks that are initially released to the set of locks that have to be used before the initial release. The set of

used locks contains the locks that are used by the execution tree (But not the locks that are only initially released or finally acquired). The acquisition history maps all locks that are finally acquired to the set of locks that have to be used after the final acquisition.

fun *as* :: ('P, T, 'L, 'X) *lex-tree* \Rightarrow 'X *as option* **where**

```

as (NLEAF  $\pi$ ) = Some (empty, {}, empty) |
as (NNO SPAWN (LNone l) t) = as t |
as (NSPAWN (LNone l) ts t) = as ts || as t |
as (NNO SPAWN (LAcq x) t) = (
  case as t of
    None  $\Rightarrow$  None |
    Some (l, u, e)  $\Rightarrow$ 
      if  $x \in \text{dom } l$  then
        Some (l-add-use (l-remove l x) {x}, insert x u, e)
      else if  $x \notin \text{dom } e$  then
        Some (l, u, e(x $\mapsto$ u))
      else
        None
  ) |
as (NNO SPAWN (LRel x) t) = (
  case as t of
    None  $\Rightarrow$  None |
    Some (l, u, e)  $\Rightarrow$  Some (l(x $\mapsto$ {}), u, e)
  ) |
as - = None

```

The acquisition structure of a hedge is the parallel composition of the acquisition structures of its trees. The acquisition structure of the empty hedge is the identity acquisition structure *Some* (*empty*, {}, *empty*).

fun *ash* :: ('P, T, 'L, 'X) *lex-hedge* \Rightarrow 'X *as option* **where**

```

ash [] = Some (empty, {}, empty) |
ash (t#h) = as t || ash h

```

lemma *l-add-use-dom[simp]*: *dom* (*l-add-use* *l* *X*) = *dom* *l*
 <proof>

lemma *l-add-use-empty[simp]*: *l-add-use* *empty* *X* = *empty*
 <proof>

lemma *l-add-use-eq-empty[simp]*: *l-add-use* *f* *X* = *empty* \iff *f* = *empty*
 <proof>

lemma *l-add-use-add[simp]*:
l-add-use (*l*++*l'*) *X* = *l-add-use* *l* *X* ++ *l-add-use* *l'* *X*
 <proof>

lemma *l-add-use-le*: *l* \leq *l-add-use* *l* *X*
 <proof>

lemma *l-remove-add[simp]*: $l\text{-remove } (l1 ++ l2) m = l\text{-remove } l1 m ++ l\text{-remove } l2 m$
 ⟨proof⟩

lemma *l-remove-no-eff[simp]*: $x \notin \text{dom } l \implies l\text{-remove } l x = l$
 ⟨proof⟩

lemma *l-remove-dom[simp]*: $\text{dom } (l\text{-remove } l x) = \text{dom } l - \{x\}$
 ⟨proof⟩

lemma *l-remove-app[simp]*:
 $l\text{-remove } l x x = \text{None}$
 $x \neq x' \implies l\text{-remove } l x x' = l x'$
 ⟨proof⟩

lemma *l-remove-eq-empty*: $l\text{-remove } l x = \text{empty} \implies \text{dom } l \subseteq \{x\}$
 ⟨proof⟩

lemma *l-remove-le-l [simp]*: $l\text{-remove } l x \leq l$
 ⟨proof⟩

lemma *as-ran-e-le-u*: $as t = \text{Some } (l, u, e) \implies \bigcup \text{ran } e \subseteq u$
 ⟨proof⟩

lemma *ash-le-u*: $ash h = \text{Some } (l, u, e) \implies \bigcup \text{ran } e \subseteq u$
 ⟨proof⟩

lemma *ash-final[simp]*: $\text{final } h \implies ash h = \text{Some } (\text{empty}, \{\}, \text{empty})$
 ⟨proof⟩

lemma *ash-append[simp]*: $ash (h1 @ h2) = ash h1 \parallel ash h2$
 ⟨proof⟩

lemma *ash-LNone-simps[simp]*:
 $ash (h1 @ \text{NSPAWN } (LNone l) ts t \# h2) = ash (h1 @ ts \# t \# h2)$
 $ash (h1 @ \text{NNOSPAWN } (LNone l) t \# h2) = ash (h1 @ t \# h2)$
 ⟨proof⟩

11.3 Consistency of Acquisition Structures

The consistency criterium of an acquisition structure decides whether the corresponding hedge can be scheduled. Note that we currently do not check this criterium during construction of the acquisition structure, but only at the end, for the completely constructed acquisition structure.

The consistency criterium has two parts. The first part is a generalization of the $\neg \exists m_1, m_2. m_1 \in h_1(m_2) \wedge m_2 \in h_2(m_1)$ -condition of [4]. There, the condition was checked for two separate acquisition histories h_1 and h_2 that

resulted from executions of two independent pushdown systems. Here, we have one execution described as a tree. This criterium can be interpreted as checking acyclicity of a graph defined by the acquisition histories. In [4], every possible cycle has length two, hence their condition is sufficient. In our setting, a cycle may have arbitrary length (bounded only by the number of locks), hence we use a general cyclicity check.

The acquisition and release histories encode a graph between locks. For an acquisition history e , the graph contains an edge (x, x') if x has to be finally acquired before x' is used, that is if $x \in \text{dom } e \wedge x' \in \text{the } (e x)$

For a release history l , the graph contains an edge (x, x') if x has to be used before x' is initially released, that is if $x' \in \text{dom } l \wedge x \in \text{the } (l x')$

definition $\text{agraph} :: 'X \text{ ah} \Rightarrow ('X \times 'X)$ set **where**
 $\text{agraph } e == \{ (x, x') . x \in \text{dom } e \wedge x' \in \text{the } (e x) \}$

definition $\text{rgraph} :: 'X \text{ ah} \Rightarrow ('X \times 'X)$ set **where**
 $\text{rgraph } l == \{ (x, x') . x' \in \text{dom } l \wedge x \in \text{the } (l x') \}$

lemma agraph-alt : $\text{agraph } e = \{ (x, x') . \exists X'. e x = \text{Some } X' \wedge x' \in X' \}$
 $\langle \text{proof} \rangle$

lemma rgraph-alt : $\text{rgraph } l = \{ (x, x') . \exists X. l x' = \text{Some } X \wedge x \in X \}$
 $\langle \text{proof} \rangle$

For the same map, the acquisition graph is the converse of the release graph. This lemma makes reasoning simpler at some points, as acquisition and release histories have the same type, and cyclicity is equivalent for a graph and its converse.

lemma $\text{agraph-rgraph-converse}$: $\text{agraph } h = (\text{rgraph } h)^{-1}$
 $\langle \text{proof} \rangle$

lemma agraph-add-union :
 $\llbracket \text{dom } e \cap \text{dom } e' = \{\} \rrbracket \implies \text{agraph } (e ++ e') = \text{agraph } e \cup \text{agraph } e'$
 $\langle \text{proof} \rangle$

lemma rgraph-add-union :
 $\llbracket \text{dom } l \cap \text{dom } l' = \{\} \rrbracket \implies \text{rgraph } (l ++ l') = \text{rgraph } l \cup \text{rgraph } l'$
 $\langle \text{proof} \rangle$

lemma $\text{agraph-domain-simp}[\text{simp}]$:
 $\text{Domain } (\text{agraph } h) = \text{dom } h - \{ x . h x = \text{Some } \{\} \}$
 $\langle \text{proof} \rangle$

lemma $\text{agraph-range-simp}[\text{simp}]$: $\text{Range } (\text{agraph } h) = \bigcup \text{ran } h$
 $\langle \text{proof} \rangle$

lemma $\text{rgraph-domain-simp}[\text{simp}]$: $\text{Domain } (\text{rgraph } h) = \bigcup \text{ran } h$
 $\langle \text{proof} \rangle$

lemma $\text{rgraph-range-simp}[\text{simp}]$:

$Range (rgraph\ h) = dom\ h - \{ x . h\ x = Some\ \{\} \}$
 ⟨proof⟩

lemma *graph-empty[simp]*:

$agraph\ empty = \{\}$
 $rgraph\ empty = \{\}$
 ⟨proof⟩

lemma *rgraph-add-use*: $rgraph\ (l\text{-add-use}\ l\ X) = rgraph\ l \cup X \times dom\ l$
 ⟨proof⟩

lemma *rgraph-remove*: $rgraph\ (l\text{-remove}\ l\ x) = rgraph\ l - UNIV \times \{x\}$
 ⟨proof⟩

lemma *rgraph-upd*: $x \notin dom\ l \implies rgraph\ (l(x \mapsto X)) = rgraph\ l \cup X \times \{x\}$
 ⟨proof⟩

lemmas *rgraph-ops* = *rgraph-add-use* *rgraph-remove* *rgraph-upd*

lemma *agraph-upd*: $x \notin dom\ e \implies agraph\ (e(x \mapsto X)) = agraph\ e \cup \{x\} \times X$
 ⟨proof⟩

lemmas *agraph-ops* = *agraph-upd*

lemma *rgraph-mono*: $l \leq l' \implies rgraph\ l \subseteq rgraph\ l'$
 ⟨proof⟩

lemma *agraph-mono*: $e \leq e' \implies agraph\ e \subseteq agraph\ e'$
 ⟨proof⟩

An acquisition or release history is consistent, iff its graph is acyclic.

abbreviation *cons-rh* :: $'X\ ah \Rightarrow bool$ **where** *cons-rh* $h == acyclic\ (rgraph\ h)$

abbreviation *cons-ah* :: $'X\ ah \Rightarrow bool$ **where** *cons-ah* $h == acyclic\ (agraph\ h)$

abbreviation *cons-h* == *cons-rh*

As noted above, the cyclicity criterion is equivalent for a graph and its converse, such that we can use *cons-h* for both, acquisition and release histories.

lemma *cons-ah-rh-eq*:

$cons\text{-ah}\ e = cons\text{-h}\ e$
 $cons\text{-rh}\ r = cons\text{-h}\ r$
 ⟨proof⟩

lemma *cons-h-empty[simp]*: *cons-h* *empty*

⟨proof⟩

lemma *cons-h-add*:

$\llbracket dom\ h \cap dom\ h' = \{\};\ cons\text{-h}\ (h++h') \rrbracket \implies cons\text{-h}\ h$
 $\llbracket dom\ h \cap dom\ h' = \{\};\ cons\text{-h}\ (h++h') \rrbracket \implies cons\text{-h}\ h'$

<proof>

lemma *cons-h-antimono*: $\llbracket l \leq l'; \text{cons-h } l' \rrbracket \implies \text{cons-h } l$
<proof>

lemma *cons-h-update*:
assumes *A*: $\text{cons-h } h \quad X \cap \text{insert } x (\text{dom } h) = \{\}$
shows $\text{cons-h } (h(x \mapsto X))$
<proof>

lemma *cons-h-update2*:
assumes *A*: $\text{cons-h } h \quad x \notin \text{dom } h \quad x \notin X \quad x \notin \bigcup \text{ran } h$
shows $\text{cons-h } (h(x \mapsto X))$
<proof>

lemma *cons-h-remove*: $\text{cons-h } l \implies \text{cons-h } (l\text{-remove } l \ m)$
<proof>

lemma *cons-h-add-use*: $\llbracket m \notin \text{dom } l; \text{cons-h } l \rrbracket \implies \text{cons-h } (l\text{-add-use } l \ \{m\})$
<proof>

lemma *cons-h-add-remove*: $\text{cons-h } l \implies \text{cons-h } (l\text{-add-use } (l\text{-remove } l \ m) \ \{m\})$
<proof>

lemma *cons-h-add-remove-partial*:
 $\llbracket m \notin \text{dom } l1; \text{cons-h } (l1 ++ l2) \rrbracket \implies$
 $\text{cons-h } (l1 ++ l\text{-add-use } (l\text{-remove } l2 \ m) \ \{m\})$
<proof>

The consistency condition for acquisition structures checks available locks in addition to consistency of the acquisition and release histories.

fun *cons-as* :: $'X \text{ as} \Rightarrow 'X \text{ set} \Rightarrow \text{bool}$ **where**
cons-as $(l, u, e) \ \xi \longleftrightarrow$
 $u \cap (\xi - \text{dom } l) = \{\} \wedge \text{dom } e \cap (\xi - \text{dom } l) = \{\} \wedge \text{cons-h } l \wedge \text{cons-h } e$

lemma *cons-as-antimono*: $\llbracket \text{cons-as } h \ \xi; \xi' \subseteq \xi \rrbracket \implies \text{cons-as } h \ \xi'$
<proof>

fun *cons* **where**
cons $\text{None } X = \text{False} \mid$
cons $(\text{Some } (l, u, e)) \ X = \text{cons-as } (l, u, e) \ X$

11.3.1 Minimal Elements

lemma *finite-acyclic-wf*: $\llbracket \text{finite } r; \text{acyclic } r \rrbracket \implies \text{wf } r$
<proof>

The minimal elements of acquisition and release histories corresponds to those final acquisitions or initial releases that can safely be scheduled as

next step — for an acquisition history without blocking any further locks usage and for a release history without requiring usage of already acquired locks.

abbreviation $rh\text{-}min\ l\ m == m \in dom\ l \wedge dom\ l \cap the\ (l\ m) = \{\}$

abbreviation $ah\text{-}min\ e\ m == m \in dom\ e \wedge m \notin \bigcup ran\ e$

lemma $rh\text{-}min\text{-}alt$:

$rh\text{-}min\ l\ m = (case\ l\ m\ of\ None \Rightarrow False \mid Some\ M \Rightarrow dom\ l \cap M = \{\})$
 $\langle proof \rangle$

There exists a minimal element in a consistent release history. Note that this lemma depends on the set of locks being finite, as assumed by the *LDPN* locale.

theorem (in *LDPN*) $cons\text{-}h\text{-}ex\text{-}rh\text{-}min$:

fixes $l :: 'X\ ah$

assumes $A: l \neq empty \quad cons\text{-}h\ l$

shows $\exists m. rh\text{-}min\ l\ m$

$\langle proof \rangle$

There exists a minimal element in a consistent acquisition history.

Note that this lemma depends on the set of locks being finite, as constrained by the *LDPN* locale.

theorem (in *LDPN*) $cons\text{-}h\text{-}ex\text{-}ah\text{-}min$:

fixes $e :: 'X\ ah$

assumes $A: e \neq empty \quad cons\text{-}h\ e$

shows $\exists m. ah\text{-}min\ e\ m$

$\langle proof \rangle$

11.3.2 Well-Nestedness and Acquisition Structures

Only locks that are on the lock-stack can be initially released:

lemma $wn\text{-}t\text{-}dom\text{-}l\text{-}lower\text{-}\mu$:

$\llbracket wn\text{-}t'\ t\ \mu; as\ t = Some\ (l, u, e) \rrbracket \Longrightarrow dom\ l \subseteq set\ \mu$

$\langle proof \rangle$

lemmas $wn\text{-}dom\text{-}l\text{-}empty = wn\text{-}t\text{-}dom\text{-}l\text{-}lower\text{-}\mu[of\ _, simplified]$

lemma $wn\text{-}h\text{-}dom\text{-}l\text{-}lower\text{-}\mu$:

$\llbracket wn\text{-}h\ h\ \mu; ash\ h = Some\ (l, u, e) \rrbracket \Longrightarrow dom\ l \subseteq locks\text{-}\mu\ \mu$

$\langle proof \rangle$

Due to well-nestedness, if a lock x is left, all locks that are above this lock on the stack are left, too. This lemma expresses leaving a lock by means of the domain of the release-history. Moreover, the release histories of the locks released before are smaller or equal than the release history of x , and do not contain x .

lemma $wn\text{-}t\text{-}dom\text{-}l\text{-}stack$: $\llbracket wn\text{-}t'\ t\ \mu; as\ t = Some\ (l, u, e); x \in dom\ l \rrbracket \Longrightarrow$

$\exists \mu 1 \ \mu 2. \ \mu = \mu 1 @ x \# \mu 2 \wedge \text{set } \mu 1 \subseteq \text{dom } l \wedge$
 $(\forall x' \in \text{set } \mu 1. \ l \ x' \leq l \ x \wedge$
 $(\text{case } l \ x' \text{ of None} \Rightarrow \text{True} \mid \text{Some } lx' \Rightarrow x \notin lx' \wedge x' \notin lx'))$
 $)$
 <proof>

lemma *wn-t-dom-l-stack'*: $\llbracket \text{wn-t}' \ t \ \mu; \text{ as } t = \text{Some } (l, u, e); x \in \text{dom } l \rrbracket \Longrightarrow$
 $\exists \mu 1 \ \mu 2. \ \mu = \mu 1 @ x \# \mu 2 \wedge \text{set } \mu 1 \subseteq \text{dom } l \wedge$
 $(\forall x' \in \text{set } \mu 1. \ l \ x' \leq l \ x \wedge x \notin \text{the } (l \ x') \wedge x' \notin \text{the } (l \ x'))$
 <proof>

11.4 Soundness of the Consistency Condition

context *LDPN*

begin

The consistency condition for acquisition structures is sound, i.e. if a hedge h is schedulable with initial locks X , and is well-nested w.r.t. a lock stack list μ containing the locks from X , then the acquisition structure of h is consistent w.r.t. X .

theorem *acqh-sound*:

$\llbracket \text{lsched } h \ X \ w; \text{ wn-h } h \ \mu; X = \text{locks-}\mu \ \mu \rrbracket \Longrightarrow$
 $\exists l \ u \ e. \ \text{ash } h = \text{Some } (l, u, e) \wedge \text{cons-as } (l, u, e) \ (\text{locks-}\mu \ \mu)$

— The proof works by induction over the schedule, in each induction step prepending a step to the schedule.

For steps that have performed operation on locks, the proof is straightforward.

If the first step of the execution is a release of a lock, the acquisition history of the new hedge (with prepended release step at one tree) remains consistent. Acyclicity is preserved, as the release-step is the first step of the execution. Consistency w.r.t. used locks is also preserved.

If the first step of the execution is an acquisition step, we further have to distinguish whether it is a usage or a final acquisition.

<proof>

end

11.5 Precision of the Consistency Condition

11.5.1 Custom Size Function

In the following we construct a custom size function for hedges that is suited to do induction over hedges. This size function decreases on any step done on the hedge.

fun *list-size'* **where**

$\text{list-size}' \ f \ [] = (0 :: \text{nat}) \mid$
 $\text{list-size}' \ f \ (a \# l) = f \ a + \text{list-size}' \ f \ l$

fun *size-t* **where**

$\text{size-t} \ (\text{NLEAF } \pi) = \text{Suc } 0 \mid$

$size-t (NNOSPAWN lab t) = Suc (size-t t) \mid$
 $size-t (NSPAWN lab ts t) = Suc (size-t ts + size-t t)$

lemma *list-size'-conc[simp]*: $list-size' f (a@b) = list-size' f a + list-size' f b$
 ⟨proof⟩

abbreviation *hedge-size* :: $('P, \Gamma, 'L, 'X) lex-hedge \Rightarrow nat$ **where**
 $hedge-size h == list-size' size-t h$

lemma *hedge-size-zero[simp]*: $hedge-size h = 0 \longleftrightarrow h=[]$
 ⟨proof⟩

This function checks whether a lock is released in the current execution tree, and returns the set of locks that are acquired before this lock is released. Note that this function ignores the lock-effect of labels of spawn-nodes, as we assume that spawn-nodes have no lock-operation.

fun *closing* :: $'X \Rightarrow ('P, \Gamma, 'L, 'X) lex-tree \Rightarrow 'X set option$ **where**
 $closing x (NLEAF \pi) = None \mid$
 $closing x (NSPAWN lab ts t) = closing x t \mid$
 $closing x (NNOSPAWN (LNone nlab) t) = closing x t \mid$
 $closing x (NNOSPAWN (LAcq x') t) = ($
 $case closing x t of None \Rightarrow None \mid$
 $Some X \Rightarrow Some (insert x' X)$
 $) \mid$
 $closing x (NNOSPAWN (LRel x') t) = (if x=x' then Some \{\} else closing x t)$

Function that checks whether a tree starts with the acquisition of a lock that is used (i.e. not finally acquired) and returns all the locks that are used from the acquisition to to the release of that lock:

fun *closing'* **where**
 $closing' (NNOSPAWN (LAcq x) t) = closing x t \mid$
 $closing' - = None$

The following functions define the set of locks that are acquired at the roots of a tree/hedge. This function is used in the case of the precision proof, where all the roots of the hedge are either leaves or final acquisitions.

fun *rootlocks-t* **where**
 $rootlocks-t (NNOSPAWN (LAcq x) t) = \{x\} \mid$
 $rootlocks-t - = \{\}$
fun *rootlocks* **where**
 $rootlocks [] = \{\} \mid$
 $rootlocks (t \# h) = rootlocks-t t \cup rootlocks h$

lemma *rootlocks-conc[simp]*: $rootlocks (h1@h2) = rootlocks h1 \cup rootlocks h2$
 ⟨proof⟩

lemma *rootlocks-split*:
 $\llbracket x \in rootlocks h; !!h1 t h2. h=h1@NNOSPAWN (LAcq x) t\#h2 \implies P \rrbracket \implies P$

<proof>

If a lock x is closed (before it is acquired), the value of the release history for x is precisely the set of used locks before x is closed. Closing x before it is acquired is expressed by well-nestedness w.r.t. a lock-stack that contains x .

lemma *closing-dom-l*:

$\llbracket wn-t' t (xs1 @ x \# xs2); closing\ x\ t = Some\ Xu; as\ t = Some\ (l, u, e) \rrbracket \implies$
 $l\ x = Some\ Xu$

<proof>

A lock must not be used before it is closed.

lemma *wn-closing-ni*: $\llbracket wn-t' t (\mu1 @ x \# \mu2); closing\ x\ t = Some\ Xu \rrbracket \implies x \notin Xu$

<proof>

This lemma gives properties of the acquisition structure after an acquisition step of a lock usage. It is used in the case when there is a tree starting with a usage, to reason about the acquisition structure after the root node of this tree has been scheduled.

lemma *wn-closing-as-fmt*:

assumes A : $wn-t' (NNOSPAWN (LAcq\ x)\ t)\ \mu$
 $as\ (NNOSPAWN (LAcq\ x)\ t) = Some\ (l, u, e)$
 $closing\ x\ t = Some\ Xu$

assumes C : $!!l'\ u'$. $\llbracket as\ t = Some\ (l', u', e); l' \leq l(x \mapsto Xu);$
 $u = insert\ x\ u'; dom\ l' = insert\ x\ (dom\ l)$
 $\rrbracket \implies P$

shows P

<proof>

A lock that occurs in the release history is closed in the execution tree, using the locks as described in the RH.

lemma *dom-l-closing*:

$\llbracket as\ t = Some\ (l, u, e); wn-t' t\ \mu; l\ x = Some\ Xu \rrbracket \implies closing\ x\ t = Some\ Xu$
<proof>

If a tree starts with a final acquisition of x , its release history is empty and the acquisition history of x contains all the used locks.

With Lemma *as-ran-e-le-u* we then also have that the ranges of the acquisition histories contain precisely the used locks.

lemma *ncl-as-fmt-single*:

assumes A : $wn-t' (NNOSPAWN (LAcq\ x)\ t)\ \mu$
 $closing' (NNOSPAWN (LAcq\ x)\ t) = None$
 $as\ (NNOSPAWN (LAcq\ x)\ t) = Some\ (l, u, e)$

shows $u = \bigcup ran\ e$ $l = empty$ $e\ x = Some\ u$

<proof>

This lemma describes properties of the acquisition structure of a tree after a final acquisition has been scheduled.

lemma *ncl-as-fmt-single'*:

assumes A : $wn-t' (NNOSPAWN (LAcq\ x)\ t)\ \mu$
 $closing' (NNOSPAWN (LAcq\ x)\ t) = None$
 $as (NNOSPAWN (LAcq\ x)\ t) = Some\ (l,u,e)$
assumes C : $!!e'. \llbracket as\ t = Some\ (empty,\ u,\ e')$;
 $u = \bigcup\ ran\ e; l = empty$;
 $e = e'(x \mapsto u); x \notin dom\ e'$
 $\rrbracket \implies P$

shows P

<proof>

The acquisition structure of a hedge whose trees start with final acquisitions or are leafs has a special structure:

- The release history is empty.
- The ranges of the acquisition histories contain precisely the used locks.
- The acquisition histories for the locks at the roots of the hedge contain precisely the used locks.
- The acquisition histories are defined for the locks at the roots of the hedge.

The first proposition follows because an initial release cannot come after a final acquisition due to well-nestedness. The second and third propositions follow as the roots of the hedge precede every other node in the hedge. The fourth proposition follows directly from the assumption that every root node that acquired a lock is a final acquisition.

lemma *ncl-as-fmt*:

\llbracket
 $wn-h\ h\ \mu; ash\ h = Some\ (l,u,e);$
 $!!Q\ t. \llbracket t \in set\ h; !!x\ t'. t = NNOSPAWN\ (LAcq\ x)\ t' \implies Q;$
 $!!p\ w. t = NLEAF\ (p,w) \implies Q$
 $\rrbracket \implies Q;$
 $\forall t \in set\ h. closing'\ t = None$
 $\rrbracket \implies l = empty \wedge u = \bigcup\ ran\ e \wedge$
 $\bigcup\ ran\ (e \mid 'rootlocks\ h) = \bigcup\ ran\ e \wedge$
 $rootlocks\ h \subseteq dom\ e$

<proof>

This lemma makes explicit the case-distinction along which the precision proof is done. The cases are:

final All trees are leaf nodes.

spawn There is a tree starting with a *NSPAWN* x - node.

none There is a tree starting with a *NNOSPAWN* $LNone$ - node.

release There is a tree starting with a *NNOSPAWN* (*LRel x*)-node.

acquire All trees start with a *NNOSPAWN* (*LAcq x*)-node or are leafs. At least one tree is no leaf.

lemma *h-cases*[*case-names final spawn none release acquire*]:

assumes *C*:

final h $\implies P$

$!!h1\ lab\ ts\ t\ h2. h=h1@NSPAWN\ lab\ ts\ t\#h2 \implies P$

$!!h1\ t\ nlab\ h2. h=h1@NNOSPAWN\ (LNone\ nlab)\ t\#h2 \implies P$

$!!h1\ x\ t\ h2. h=h1@NNOSPAWN\ (LRel\ x)\ t\#h2 \implies P$

$\llbracket !!Q\ t. \llbracket t \in set\ h; !!x\ t'. t=NNOSPAWN\ (LAcq\ x)\ t' \implies Q;$

$!!p\ w. t=NLEAF\ (p,w) \implies Q$

$\rrbracket \implies Q;$

$!!Q. \llbracket !!t' x. NNOSPAWN\ (LAcq\ x)\ t' \in set\ h \implies Q \rrbracket \implies Q$

$\rrbracket \implies P$

shows *P*

<proof>

This lemma determines the tree within a hedge whose release history contains a specific lock.

lemma *ash-find-l-t*[*consumes 2*]:

$\llbracket ash\ h = Some\ (l,u,e); x \in dom\ l;$

$!!h1\ t\ h2\ l1\ u1\ e1\ l2\ u2\ e2. \llbracket$

$h=h1@t\#h2; l=l1++l2; u=u1 \cup u2; e=e1++e2;$

$as\ t = Some\ (l1,u1,e1); ash\ h1 \parallel ash\ h2 = Some\ (l2,u2,e2);$

$x \in dom\ l1; dom\ l1 \cap dom\ l2 = \{\}; dom\ e1 \cap dom\ e2 = \{\}$

$\rrbracket \implies P$

$\rrbracket \implies P$

<proof>

This lemma determines the tree within a hedge whose acquisition history contains a specific lock.

lemma *ash-find-e-t*[*consumes 2*]:

$\llbracket ash\ h = Some\ (l,u,e); x \in dom\ e;$

$!!h1\ t\ h2\ l1\ u1\ e1\ l2\ u2\ e2. \llbracket$

$h=h1@t\#h2; l=l1++l2; u=u1 \cup u2; e=e1++e2;$

$as\ t = Some\ (l1,u1,e1); ash\ h1 \parallel ash\ h2 = Some\ (l2,u2,e2);$

$x \in dom\ e1; dom\ l1 \cap dom\ l2 = \{\}; dom\ e1 \cap dom\ e2 = \{\}$

$\rrbracket \implies P$

$\rrbracket \implies P$

<proof>

Auxilliary lemma to split the acquisition history of a hedge by some tree in that hedge.

lemma *ash-split-aux*:

assumes *AS*: *ash h = Some (l,u,e)* **and**

HFMT[*simp*]: *h=h1@t#h2* **and**

C: $!!l1\ u1\ e1\ l2\ u2\ e2. \llbracket$

$$\begin{aligned}
& l=l1++l2; u=u1\cup u2; e=e1++e2; \text{ as } t = \text{Some } (l1,u1,e1); \\
& \text{ash } h1 \parallel \text{ash } h2 = \text{Some } (l2,u2,e2); \\
& \text{dom } l1 \cap \text{dom } l2 = \{\}; \text{ dom } e1 \cap \text{dom } e2 = \{\}
\end{aligned}$$

]] $\Rightarrow P$

shows P

<proof>

Auxilliary lemma that combines *ash-split-aux* and *wn-h-split-aux*.

lemma *wn-ash-split-aux*:

assumes

WN: *wn-h* h μ **and**

AS: *ash* $h = \text{Some } (l,u,e)$ **and**

HFMT[*simp*]: $h=h1@t\#h2$ **and**

C: $!!\mu1\ xs\ \mu2\ l1\ u1\ e1\ l2\ u2\ e2.$ [[

$\mu=\mu1@xs\#\mu2; l=l1++l2; u=u1\cup u2; e=e1++e2;$

wn-t' $t\ xs; \text{wn-h } h1\ \mu1; \text{wn-h } h2\ \mu2;$

as $t = \text{Some } (l1,u1,e1); \text{ash } h1 \parallel \text{ash } h2 = \text{Some } (l2,u2,e2);$

locks- μ $\mu1 \cap \text{set } xs = \{\}; \text{locks-}\mu\ \mu1 \cap \text{locks-}\mu\ \mu2 = \{\};$

set $xs \cap \text{locks-}\mu\ \mu2 = \{\}; \text{dom } l1 \cap \text{dom } l2 = \{\}; \text{dom } e1 \cap \text{dom } e2 = \{\}$

]] $\Rightarrow P$

shows P

<proof>

context *LDPN*

begin

Precision of the acquisition structure construction, i.e. for a well-nested hedge, a consistent acquisition history implies a schedule.

theorem *acqh-precise*:

fixes $h::('P, 'T, 'L, 'X)$ *lex-hedge*

assumes *A*: *ash* $h = \text{Some } (l,u,e)$ *cons-as* (l,u,e) $(\text{locks-}\mu\ \mu)$ *wn-h* $h\ \mu$

shows $\exists w. \text{lsched } h (\text{locks-}\mu\ \mu) w$

— The proof is done by induction on the size of the hedge.

Given a non-empty hedge, it constructs the first step of the schedule and shows that the acquisition structure remains consistent.

It considers the following cases:

- If the hedge contains a root that has no effect on locks, this root is scheduled. Those steps can always be scheduled, as the acquisition structure and the set of acquired locks do not change.
- If the hedge contains a root that initially releases a lock x , it is scheduled. A release can always be scheduled, as it cannot block. The new acquisition structure remains consistent: The acquisition history is unchanged, the release history decreases (the lock x is removed). Consistency is preserved, as the lock x does not occur in the set of acquired locks any more.
- If the hedge contains only roots that are lock acquisitions or leaves, we further distinguish whether some of the roots are usages, or there are only final acquisitions.

- If some of the roots are usages, we can find a usage where the used locks are disjoint from the domain of the release history (Due to acyclicity of the RH). Intuitively, this is a usage where the required locks are already released. This usage could be scheduled as a whole, without changing the RH, AH or set of acquired locks, and only decreasing the set of used locks. However, we chose another way here and show that scheduling only the first acquisition step of the usage also preserves consistency of the AS. We chose this approach in order to not having to formalize the scheduling of a usage. We assume that this simplifies formalization overhead (Perhaps at the cost of increased proof complexity).
- If all of the roots are leafs or final acquisitions, due to acyclicity of the AH, we can select a final acquisition that acquires a lock that is not used in the rest of the hedge. Scheduling this acquisition preserves consistency of the AS.

<proof>

The following is the main theorem of this section. It states the correctness of the acquisition structure construction. For all non-empty hedges that are well-nested w.r.t. a list of lock-stacks with locks X , the existence of a schedule starting with locks X is equivalent to the consistency of the hedge's acquisition history w.r.t. X .

lemma *acqh-correct'*:

fixes $h::('P, 'T, 'L, 'X)$ *lex-hedge*

shows $\llbracket wn-h\ h\ \mu \rrbracket \implies$

$(\exists w. \text{lsched } h\ (\text{locks-}\mu\ \mu)\ w) \longleftrightarrow$

$(\exists l\ u\ e. \text{ash } h = \text{Some } (l, u, e) \wedge \text{cons-as } (l, u, e)\ (\text{locks-}\mu\ \mu)$

)

<proof>

theorem *acqh-correct*:

fixes $h::('P, 'T, 'L, 'X)$ *lex-hedge*

assumes $WN: wn-h\ h\ \mu$

shows $(\exists w. \text{lsched } h\ (\text{locks-}\mu\ \mu)\ w) \longleftrightarrow \text{cons } (\text{ash } h)\ (\text{locks-}\mu\ \mu)$

<proof>

end

end

12 DPNs with Initial Configuration

theory *DPN-c0*

imports *WellNested*

begin

12.1 DPNs with Initial Configuration

In the following locale, we fix a DPN with an initial configuration, and a list of lock-stacks. We assume that the initial configuration is well-nested w.r.t. the list of lock-stacks.

This is the model we are able to analyze with our acquisition history based techniques, that assume well-nestedness.

Note that we – up to now – do not show that there exists a non-trivial instance of this locale. Such a proof would support the trust in that the model we formalize here is really the intended model.

```

locale LDPN-c0 = LDPN +
  constrains  $\Delta :: ('P, T, 'L, 'X :: finite) \text{ldpn}$ 
  fixes c0 :: ('P, T) conf           – Initial configuration
  fixes  $\mu 0 :: 'X \text{list list}$        – Locks held at the start configuration
  assumes wellnested: wn-c  $\Delta$  c0  $\mu 0$  – Start configuration must be well-nested
begin

```

12.1.1 Reachable Configurations

definition *reachable* == { *c* . $\exists w. (c0, w, c) \in \text{dpntrc } \Delta$ }

definition *reachablels* == { (*c, X*) . $\exists w. ((c0, \text{locks-}\mu \mu 0), w, (c, X)) \in \text{ldpntrc } \Delta$ }

lemma *reachablels-subset*: (*c, X*) \in *reachablels* $\implies c \in$ *reachable*
 <proof>

lemma *reachable-wn*:

$\llbracket (c, X) \in \text{reachablels}; !!\mu. \llbracket \text{wn-c } \Delta \text{ } c \mu; X = \text{locks-}\mu \mu \rrbracket \implies P \rrbracket \implies P$
 <proof>

lemma *reachablels-triv[simp]*: (*c0*, *locks-}\mu \mu 0*) \in *reachablels*
 <proof>

end

end

13 Property Specifications

theory *Specification*

imports *DPN-c0 Semantics LockSem common/SublistOrder*

begin

We develop a formalism that allows a concise and readable notation for a class of properties that are checkable via cascaded predecessor computations.

A specification consists of a list of atoms, where each atom either restricts the current configuration or describes some step.

13.1 Specification Formulas

The base element of a property is an atom, that describes a step or restricts the current configuration

```
datatype ('Q,Γ,'L,'X) spec-atom =
  — Restrict current configuration to be in a specified set
  SPEC-RESTRICT ('Q,Γ) conf set |
  — Go forward one step, using a rule with labels from a specified set
  SPEC-STEP ('L,'X) lockstep set |
  — Go forward any number of steps, using rules with labels from a specified set
  SPEC-STEPS ('L,'X) lockstep set
```

A property is a list of atoms

```
types ('Q,Γ,'L,'X) spec = ('Q,Γ,'L,'X) spec-atom list
```

13.2 Semantics

The semantics of a property specification Φ w.r.t. the current DPN is modelled by a transition relation $spec-tr \ \Phi$, that contains all pairs (c, c') of configurations, such that there is a path between c and c' satisfying the property.

```
context LDPN
begin
  fun spec-tr where
    spec-tr [] = Id |
    spec-tr (SPEC-RESTRICT C #  $\Phi$ ) = {(c, c') . (c, c') ∈ spec-tr  $\Phi$  ∧ fst c ∈ C} |
    spec-tr (SPEC-STEP L #  $\Phi$ ) =
      {(c, c') . ∃ l ∈ L. ∃ ch. (c, l, ch) ∈ ldpntr  $\Delta$  ∧ (ch, c') ∈ spec-tr  $\Phi$ } |
    spec-tr (SPEC-STEPS L #  $\Phi$ ) =
      {(c, c') . ∃ ll ∈ lists L. ∃ ch. (c, ll, ch) ∈ ldpntrc  $\Delta$  ∧ (ch, c') ∈ spec-tr  $\Phi$ }
end
```

```
context LDPN-c0
begin
```

In most cases, it suffices to check whether there is a path matching the specification from the initial configuration.

```
definition model-check-ref  $\Phi$  == (c0, locks-μ μ0) ∈ Domain (spec-tr  $\Phi$ )
end
```

13.3 Examples

In this section, we present two short examples to justify the usefulness of our property specifications.

13.3.1 Conflict analysis

Given two stack symbols $u, v \in \Gamma$, conflict analysis asks whether a configuration c is reachable that has a conflict between u and v .

A configuration has a conflict between u and v , iff it contains a process with top stack symbol u and another (different) process with top stack symbol v .

context *LDPN-c0*
begin

$atUV\ u\ v$ is the set of configurations that have a conflict between u and v .

definition $atUV\text{-ordered}\ u\ v == \{ c . \exists q\ r\ q'\ r'. [(q, u \# r), (q', v \# r')] \leq c \}$

definition $atUV\ u\ v == (atUV\text{-ordered}\ u\ v) \cup (atUV\text{-ordered}\ v\ u)$

The following property specification describes all executions reaching a conflict:

definition $conflict\text{-spec}\ u\ v ==$
 $[SPEC\text{-STEPS}\ UNIV, SPEC\text{-RESTRICT}\ (atUV\ u\ v)]$

The following definition is a direct definition of a conflict between u and v being reachable from an initial configuration $[(qmain, [\gamma main])]$:

definition $has\text{-conflict}\text{-ref}\ u\ v == \exists (c, X) \in reachables. c \in atUV\ u\ v$

The next lemma shows that the direct definition of a conflict matches the property specification:

lemma $has\text{-conflict}\text{-ref}\ u\ v \longleftrightarrow model\text{-check}\text{-ref}\ (conflict\text{-spec}\ u\ v)$
 $\langle proof \rangle$

end

13.3.2 Bitvector analysis

Given a set of generator labels $G :: 'L\ set$, a set of killer labels $K :: 'L\ set$ and a stack symbol $u :: \Gamma$, bitvector analysis asks whether there is a path to a configuration that has process being at u , such that the path executes a generator rule, and after that no killer rule is executed.

context *LDPN-c0*
begin

For a stack symbol, $u \in \Gamma$, the set $atU\ u$ is the set of all configurations that have a process with u at the top of the stack.

definition $atU\ u == \{ c . \exists q\ r. (q, u \# r) \in set\ c \}$

The following property specification describes all paths that lead to u and have the bit set:

definition $bitvector\text{-fwd}\text{-spec}\ G\ K\ u ==$

```

[ SPEC-STEPS UNIV,
  SPEC-STEP G,
  SPEC-STEPS (UNIV-K),
  SPEC-RESTRICT (atU u)
]

```

The following is the direct definition of bitvector analysis:

definition *bitvector-fwd-ref* $G K u ==$
 $\exists c1 X1 lg c2 X2 ll c3 X3 q r.$
 $(c1, X1) \in \text{reachableIs} \wedge$
 $((c1, X1), lg, (c2, X2)) \in \text{ldpntr} \Delta \wedge$
 $lg \in G \wedge$
 $((c2, X2), ll, (c3, X3)) \in \text{ldpntrc} \Delta \wedge$
 $ll \in \text{lists} (UNIV-K) \wedge$
 $(q, u \# r) \in \text{set } c3$

This lemma shows that the direct definition matches the property specification:

lemma *bitvector-fwd-ref* $G K u \longleftrightarrow$
 $\text{model-check-ref} (\text{bitvector-fwd-spec } G K u)$
 $\langle \text{proof} \rangle$

end
end

14 Hedge Constraints for Acquisition Histories

theory *As-hc*
imports *Acqh WellNested DPN-c0 Specification*
begin

This theory formulates the set of execution hedges that have a lock-sensitive schedule, and shows how to use hedge-constrained predecessor set computations to compute property specifications based on cascaded predecessor sets.

14.1 Locks Encoded in Control State

For this section, we make the assumption that the set of locks is encoded in the control state of the DPN. We formalize this by means of a locale.

locale *EncodedLDPN* = *LDPN* +
— The states of the DPN are tuples of some states $'P$ and sets of locks:
constrains $\Delta :: ('P \times 'X \text{ set}, 'T, 'L, 'X :: \text{finite}) \text{ldpn}$
constrains $c0 :: ('P \times 'X \text{ set}, 'T) \text{conf}$
constrains $\mu0 :: 'X \text{ list list}$
— A step of the DPN transforms the locks as expected:

assumes *encoding-correct-nospawn*:

$((p, X), \gamma \hookrightarrow_l (p', X'), w) \in \Delta \implies \text{lock-valid } X \text{ l } X'$

assumes *encoding-correct-spawn1*:

$((p, X), \gamma \hookrightarrow_l (ps, Xs), ws \# (p', X'), w) \in \Delta \implies \text{lock-valid } X \text{ l } X'$

— A freshly spawned process initially owns no locks:

assumes *encoding-correct-spawn2*:

$((p, X), \gamma \hookrightarrow_l (ps, Xs), ws \# (p', X'), w) \in \Delta \implies Xs = \{\}$

begin

lemmas *encoding-correct-spawn* = *encoding-correct-spawn1* *encoding-correct-spawn2*

lemmas *encoding-correct* = *encoding-correct-nospawn* *encoding-correct-spawn*

lemma *encoding-correct-nospawn'*:

$(p, \gamma \hookrightarrow_l p', w) \in \Delta \implies \text{lock-valid } (\text{snd } p) \text{ l } (\text{snd } p')$
 $\langle \text{proof} \rangle$

lemma *encoding-correct-spawn'*:

assumes *A*: $(p, \gamma \hookrightarrow_l ps, ws \# p', w) \in \Delta$
shows *lock-valid* $(\text{snd } p) \text{ l } (\text{snd } p')$ $\text{snd } ps = \{\}$
 $\langle \text{proof} \rangle$

lemma *encoding-correct-spawn2'*:

$(p, \gamma \hookrightarrow_l ps, ws \# p', w) \in \Delta \implies \text{snd } ps = \{\}$
 $\langle \text{proof} \rangle$

lemma *ec-preserve-singlestep*:

assumes

A: $((c, \text{locks-}\mu \ \mu), l, (c', X')) \in \text{ldpntnr } \Delta \quad \text{wn-c } \Delta \ c \ \mu$
 $\text{map } (\text{snd} \circ \text{fst}) \ c = \text{map set } \mu$ **and**

C: $!!\mu'. \llbracket \text{wn-c } \Delta \ c' \ \mu'; X' = \text{locks-}\mu \ \mu';$
 $\text{map } (\text{snd} \circ \text{fst}) \ c' = \text{map set } \mu'$
 $\rrbracket \implies P$

shows *P*

$\langle \text{proof} \rangle$

lemma *ec-preserve*:

assumes

A: $((c, \text{locks-}\mu \ \mu), ll, (c', X')) \in \text{ldpntrc } \Delta \quad \text{wn-c } \Delta \ c \ \mu$
 $\text{map } (\text{snd} \circ \text{fst}) \ c = \text{map set } \mu$ **and**

C: $!!\mu'. \llbracket X' = \text{locks-}\mu \ \mu'; \text{wn-c } \Delta \ c' \ \mu'; \text{map } (\text{snd} \circ \text{fst}) \ c' = \text{map set } \mu' \rrbracket \implies P$

shows *P*

$\langle \text{proof} \rangle$

The following abbreviates the locks owned by a configuration:

abbreviation *locks-c* $c == \text{list-collect-set } (\text{snd} \circ \text{fst}) \ c$

lemma *locks-μ-mapset*: $\text{locks-}\mu \ \mu = \bigcup \text{set } (\text{map set } \mu)$

<proof>

lemma *locks-c-mapset*: $locks-c\ c = \bigcup set\ (map\ (snd \circ fst)\ c)$
<proof>

end

locale *EncodedLDPN-c0* = *EncodedLDPN* + *LDPN-c0* +

— The states of the DPN are tuples of some states $'P$ and sets of locks:

constrains $\Delta :: ('P \times 'X\ set, \Gamma, 'L, 'X :: finite)\ ldpn$

constrains $c0 :: ('P \times 'X\ set, \Gamma)\ conf$

constrains $\mu0 :: 'X\ list\ list$

— The locks encoded in the initial configuration correspond to the locks in the initial list of lock-stacks:

assumes *encoding-correct-start*:

$map\ (snd \circ fst)\ c0 = map\ set\ \mu0$

begin

Reachable configurations are well-nested w.r.t. a lock-stack corresponding to the locks encoded in the control states of the processes

lemma *reachable-ec*:

$\llbracket (c, X) \in reachablels;$

$!!\mu. \llbracket wm-c\ \Delta\ c\ \mu; X = locks-\mu\ \mu; map\ (snd \circ fst)\ c = map\ set\ \mu \rrbracket \implies P$

$\rrbracket \implies P$

<proof>

Due to our assumptions, a reachable configuration always encodes the locks that are also used by the lock-sensitive semantics.

theorem *reachable-locks*: $(c, X) \in reachablels \implies locks-c\ c = X$

<proof>

14.2 Characterizing Schedulable Execution Hedges

In order to characterize schedulable execution hedges, we have to first characterize the locks allocated at the roots of an execution hedge. This can be done by deriving the locks at the roots from the control states annotated at the leafs.

fun *lock-eff* :: $('L, 'X)\ lockstep \Rightarrow 'X\ set \Rightarrow 'X\ set$ **where**

lock-eff (LNone nlab) X = X |

lock-eff (LAcq x) X = insert x X |

lock-eff (LRel x) X = X - {x}

fun *lock-eff-inv* :: $('L, 'X)\ lockstep \Rightarrow 'X\ set \Rightarrow 'X\ set$ **where**

lock-eff-inv (LNone nlab) X = X |

$lock\text{-}eff\text{-}inv (LAcq\ x)\ X = X - \{x\}$ |
 $lock\text{-}eff\text{-}inv (LRel\ x)\ X = insert\ x\ X$

fun $rlocks\text{-}t :: ('P \times 'X\ set, \Gamma, 'L, 'X)\ lex\text{-}tree \Rightarrow 'X\ set$ **where**
 $rlocks\text{-}t (NLEAF\ \pi) = (case\ \pi\ of\ ((p, X), w) \Rightarrow X)$ |
 $rlocks\text{-}t (NNOSPAWN\ l\ t) = lock\text{-}eff\text{-}inv\ l\ (rlocks\text{-}t\ t)$ |
 $rlocks\text{-}t (NSPAWN\ l\ ts\ t) = lock\text{-}eff\text{-}inv\ l\ (rlocks\text{-}t\ t)$

abbreviation $rlocks\text{-}h :: ('P \times 'X\ set, \Gamma, 'L, 'X)\ lex\text{-}hedge \Rightarrow 'X\ set\ list$ **where**
 $rlocks\text{-}h\ h == map\ rlocks\text{-}t\ h$

lemma $tsem\text{-}locks: tsem\ \Delta\ \pi\ t\ c' \Longrightarrow snd\ (fst\ \pi) = rlocks\text{-}t\ t$
 $\langle proof \rangle$

lemma $hsem\text{-}locks: hsem\ \Delta\ c\ h\ c' \Longrightarrow map\ (snd\ o\ fst)\ c = rlocks\text{-}h\ h$
 $\langle proof \rangle$

Next, we have to characterize the execution hedges with consistent acquisition histories w.r.t. the set of allocated locks.

definition $Hls\ h == cons\ (ash\ h)\ (\bigcup\ set\ (rlocks\text{-}h\ h))$

theorem $reachable\text{-}hls\text{-}char:$

assumes $A: (c, X) \in reachablels$ $hsem\ \Delta\ c\ h\ c'$
shows $(\exists w. lsched\ h\ X\ w) \longleftrightarrow Hls\ h$
 $\langle proof \rangle$

Now we can put it all together and show correctness of lock-sensitive predecessor computation

lemma $lsprestar1:$

assumes
 $REACH: (c, X) \in reachablels$ **and**
 $PRE: c \in prehc\ \Delta\ Hls\ C'$
shows $\exists c' \in C'. \exists ll\ X'. ((c, X), ll, (c', X')) \in ldpntrc\ \Delta$
 $\langle proof \rangle$

lemma $lsprestar2:$

assumes
 $REACH: (c, X) \in reachablels$ **and**
 $MEM: c' \in C'$ **and**
 $PATH: ((c, X), ll, (c', X')) \in ldpntrc\ \Delta$
shows $c \in prehc\ \Delta\ Hls\ C'$
 $\langle proof \rangle$

theorem $lsprestar:$

assumes $REACH: (c, X) \in reachablels$
shows $c \in prehc\ \Delta\ Hls\ C' \longleftrightarrow (\exists c' \in C'. \exists ll\ X'. ((c, X), ll, (c', X')) \in ldpntrc\ \Delta)$
 $\langle proof \rangle$

14.3 Checking Specifications Using $prehc \Delta Hls$

We now show that we can use our construction to check for property specifications (cf. `Specification.thy`).

We first have to construct a hedge-constraint for execution hedges that contain a restricted set of labels.

```
fun isLab :: ('L,'X) lockstep set  $\Rightarrow$  ('Q,'T,'L,'X) lex-tree  $\Rightarrow$  bool where
  isLab L (NLEAF  $\pi$ )  $\longleftrightarrow$  True |
  isLab L (NNOSPAWN l t)  $\longleftrightarrow$  l  $\in$  L  $\wedge$  isLab L t |
  isLab L (NSPAWN l ts t)  $\longleftrightarrow$  l  $\in$  L  $\wedge$  isLab L ts  $\wedge$  isLab L t
```

```
abbreviation HLab L == { h . list-all (isLab L) h }
```

```
lemma final-h-is-lab[simp]: final h  $\Longrightarrow$  list-all (isLab L) h
  <proof>
```

```
lemma HLab-correct: sched h ll  $\Longrightarrow$  h  $\in$  HLab L  $\longleftrightarrow$  ll  $\in$  lists L
  <proof>
```

```
lemmas HLab-correct' = HLab-correct[OF lsched-is-sched]
```

Then we can show how to check property specifications using $prehc$.

```
fun mc-pre :: ('P  $\times$  'X set, 'T, 'L, 'X) spec  $\Rightarrow$  ('P  $\times$  'X set, 'T) conf set where
  mc-pre [] = UNIV |
  mc-pre (SPEC-RESTRICT C #  $\Phi$ ) = C  $\cap$  mc-pre  $\Phi$  |
  mc-pre (SPEC-STEP L #  $\Phi$ ) = prehc  $\Delta$  (Hls  $\cap$  Hpre  $\cap$  HLab L) (mc-pre  $\Phi$ ) |
  mc-pre (SPEC-STEPS L #  $\Phi$ ) = prehc  $\Delta$  (Hls  $\cap$  HLab L) (mc-pre  $\Phi$ )
```

```
lemma mc-pre-correct-aux:
  (c,X)  $\in$  reachablels  $\Longrightarrow$  c  $\in$  mc-pre  $\Phi$   $\longleftrightarrow$  (c,X)  $\in$  Domain (spec-tr  $\Phi$ )
  <proof>
```

```
theorem mc-pre-correct: c0  $\in$  mc-pre  $\Phi$   $\longleftrightarrow$  model-check-ref  $\Phi$ 
  <proof>
```

end

end

15 Monitors (aka Block-Structured Locks)

```
theory Monitors
imports LockSem WellNested As-hc
begin
```

We model monitors by binding locks to stack symbols, and making some restrictions on rules:

- A rule labeled by *LNone* must not change the allocated locks, nor must it push or pop stack symbols associated with locks.
- An acquisition rule must be a rule that pushes a stack-symbol with the acquired lock, and does not change the locks of the stack-symbol at the bottom.
- A release rule must be a rule that pops a stack-symbol with the released lock.

One purpose of this theory is, that it gives strong evidence that our model is not too restrictive. This is done by defining an introduction rule for encoded DPNs with initial configurations that only depends on local properties of the rules and the initial configuration.

— Lock-stack encoded into stack

definition *lstackm-s* :: ($\mathbb{T} \rightarrow 'X$) \Rightarrow $\mathbb{T} \Rightarrow 'X$ list **where**
lstackm-s mon γ = (case mon γ of *None* \Rightarrow [] | *Some* $x \Rightarrow [x]$)

lemma *lstackm-s-simps*[*simp*]:

mon γ = *None* \Longrightarrow *lstackm-s mon* γ = []
mon γ = *Some* $x \Longrightarrow$ *lstackm-s mon* γ = [x]
 ⟨*proof*⟩

fun *lstackm* :: ($\mathbb{T} \rightarrow 'X$) \Rightarrow \mathbb{T} list \Rightarrow $'X$ list **where**
lstackm mon [] = [] |
lstackm mon ($\gamma \# s$) = *lstackm-s mon* γ @ *lstackm mon* s

lemma *lstackm-conc*[*simp*]:

lstackm mon ($s @ s'$) = *lstackm mon* s @ *lstackm mon* s'
 ⟨*proof*⟩

lemma *lstack-spawn-empty*[*simp*]:

[[$(\forall \gamma s \in \text{set } w. \text{mon } \gamma s = \text{None})$]] \Longrightarrow *lstackm mon* w = []
 ⟨*proof*⟩

locale *MDPN* = *EncodedLDPN* +

constrains

Δ :: ($'P \times 'X$ set, $\mathbb{T}, 'L, 'X$:: finite) *ldpn*

fixes *mon* :: $\mathbb{T} \Rightarrow 'X$ option — Maps stack symbols to associated monitors

assumes

locks-lnone-pop-nospawn:

$(p, \gamma \hookrightarrow_{LNone} a \ p', []) \in \Delta \Longrightarrow$ *mon* γ = *None* **and**

locks-lnone-pop-spawn:

$(p, \gamma \hookrightarrow_l ps, ws \ \# \ p', []) \in \Delta \Longrightarrow$ *mon* γ = *None* **and**

locks-lnone-nospawn:

$(p, \gamma \hookrightarrow_{LNone} a \ p', w@[\gamma']) \in \Delta \implies \text{mon } \gamma' = \text{mon } \gamma \wedge$
 $(\forall \gamma s \in \text{set } w. \text{mon } \gamma s = None) \text{ and}$

locks-lnone-spawn:

$(p, \gamma \hookrightarrow_l ps, ws \ \# \ p', w@[\gamma']) \in \Delta \implies \text{mon } \gamma' = \text{mon } \gamma \wedge$
 $(\forall \gamma s \in \text{set } w. \text{mon } \gamma s = None) \text{ and}$

locks-spawn:

$(p, \gamma \hookrightarrow_l ps, ws \ \# \ p', w) \in \Delta \implies (\forall \gamma s \in \text{set } ws. \text{mon } \gamma s = None) \text{ and}$

locks-acquire:

$\llbracket (p, \gamma \hookrightarrow_{LAcq} x \ p', w) \in \Delta;$
 $!!w' \ \gamma 2 \ \gamma 1. \llbracket w = w'@[\gamma 1, \gamma 2]; \text{mon } \gamma 2 = \text{mon } \gamma; \text{mon } \gamma 1 = \text{Some } x;$
 $(\forall \gamma s \in \text{set } w'. \text{mon } \gamma s = None)$
 $\rrbracket \implies P$
 $\rrbracket \implies P \text{ and}$

locks-release:

$(p, \gamma \hookrightarrow_{LRel} x \ p', w) \in \Delta \implies w = [] \wedge \text{mon } \gamma = \text{Some } x$

begin

abbreviation *lstack-s* == *lstackm-s mon*

abbreviation *lstack* == *lstackm mon*

lemma *lstack-lnone-nospawn:*

$\llbracket (p, \gamma \hookrightarrow_{LNone} a \ p', w) \in \Delta \rrbracket \implies \text{lstack } (\gamma \# r) = \text{lstack } (w @ r)$
 $\langle \text{proof} \rangle$

lemma *lstack-lnone-spawn:*

$\llbracket (p, \gamma \hookrightarrow_a ps, ws \ \# \ p', w) \in \Delta \rrbracket \implies \text{lstack } (\gamma \# r) = \text{lstack } (w @ r)$
 $\langle \text{proof} \rangle$

lemma *well-nested-t:*

assumes *CONS*: *distinct (lstack (snd π))*
assumes *H*: *tsem $\Delta \ \pi \ t \ c'$*
assumes *COINC*: *snd (fst π) = set (lstack (snd π))*
shows *wn-t' t (lstack (snd π))*
 $\langle \text{proof} \rangle$

lemma *well-nested-h:*

assumes *CONS*: *cons- μ (map (lstack \circ snd) c)*
assumes *H*: *hsem $\Delta \ c \ h \ c'$*
assumes *COINC*: *map (snd \circ fst) c = map (set \circ lstack \circ snd) c*
shows *wn-h h (map (lstack \circ snd) c)*
 $\langle \text{proof} \rangle$

theorem *well-nested:*

assumes *CONS*: *cons- μ (map (lstack \circ snd) c)*
assumes *COINC*: *map (snd \circ fst) c = map (set \circ lstack \circ snd) c*
shows *wn-c $\Delta \ c$ (map (lstack \circ snd) c)*

<proof>

This theorem can be used to show that an MDPN along with a consistent start configuration is a DPN with well-nested lock usage, as described by the locale *EncodedLDPN-c0*.

theorem *EncodedLDPN-c0-intro*[*intro?*]:
 assumes *start-config-cons*: *cons-μ μ0*
 assumes *start-config-coinc*: *map (snd ∘ fst) c0 = map set μ0*
 assumes *start-config-match*: *map (lstack ∘ snd) c0 = μ0*
 shows *EncodedLDPN-c0 Δ c0 μ0*
<proof>

end

theorem *EncodedLDPN-c0-intro-external*:
 assumes *MDPN*: *MDPN Δ mon*
 assumes *start-config-cons*: *cons-μ μ0*
 assumes *start-config-coinc*: *map (snd ∘ fst) c0 = map set μ0*
 assumes *start-config-match*: *map (lstackm mon ∘ snd) c0 = μ0*
 shows *EncodedLDPN-c0 Δ c0 μ0*
<proof>

15.1 Non-Trivial Instance of a Well-Nested DPN

In this section, we define a non-trivial Well-nested DPN by hand. This gives strong evidence that our model assumptions are not too restrictive.

We start by introducing some finite set of locks that we can use in our programs:

typedef *t-my-locks* = {*1..6::nat*} *<proof>*

instance *t-my-locks::finite*
<proof>

definition *l1* :: *t-my-locks* **where** *l1* = *Abs-t-my-locks (1::nat)*

definition *l2* :: *t-my-locks* **where** *l2* = *Abs-t-my-locks (2::nat)*

lemma [*simp, intro!*]: *l1 ≠ l2* *l2 ≠ l1*
<proof>

The following rules correspond to a by-hand translation of the (nonsense) program:

```
procedure p1:  
  sync l1 {  
    sync l2 {  
      spawn p1  
      spawn p2  
    }  
  }
```

```

}

procedure p2:
  if ? then
    spawn p2
    call p2
  else
    sync l2 {
      sync l1 {
        spawn p1
      }
    }
}

```

definition $my\Delta :: (nat \times t\text{-my-locks set}, nat, unit, t\text{-my-locks}) \text{ ldpn where}$

```

myΔ = {
  ((0, {}), 1) ↦LAcq l1 (0, {l1}), [2, 3],
  ((0, {l1}), 2) ↦LAcq l2 (0, {l1, l2}), [4, 5],
  ((0, {l1, l2}), 4) ↦LNone () (0, {}), [1]#(0, {l1, l2}), [6],
  ((0, {l1, l2}), 6) ↦LNone () (0, {}), [11]#(0, {l1, l2}), [7],
  ((0, {l1, l2}), 7) ↦LRel l2 (0, {l1}), [],
  ((0, {l1}), 5) ↦LRel l1 (0, {}), [],
  ((0, {}), 3) ↦LNone () (0, {}), [],

  ((0, {}), 11) ↦LNone () (0, {}), [11]#(0, {}), [12],
  ((0, {}), 12) ↦LNone () (0, {}), [11, 13],
  ((0, {}), 11) ↦LAcq l2 (0, {l2}), [14, 13],
  ((0, {l2}), 14) ↦LAcq l1 (0, {l1, l2}), [16, 17],
  ((0, {l1, l2}), 16) ↦LNone () (0, {}), [1]#(0, {l1, l2}), [18],
  ((0, {l1, l2}), 18) ↦LRel l1 (0, {l2}), [],
  ((0, {l2}), 17) ↦LRel l2 (0, {}), [],
  ((0, {}), 13) ↦LNone () (0, {}), []
}

```

definition $my\text{-mon} :: nat \Rightarrow t\text{-my-locks option where}$

```

my-mon s = (
  if s=1 then None
  else if s=2 then Some l1
  else if s=3 then None
  else if s=4 then Some l2
  else if s=5 then Some l1
  else if s=6 then Some l2
  else if s=7 then Some l2
  else if s=11 then None
  else if s=12 then None
)

```

```

else if s=13 then None
else if s=14 then Some l2
else if s=15 then None
else if s=16 then Some l1
else if s=17 then Some l2
else if s=18 then Some l1
else None
)

```

It is straightforward to show that this is an MDPN

interpretation $MDPN[my\Delta\ my-mon]$
<proof>

And with the stuff proven above, we also get that this program is a well-nested LDPN w.r.t. the start configuration $[(\emptyset::'a, \{\}), [1::'c]]$, which corresponds to starting with procedure **p1**.

interpretation $EncodedLDPN-c0[my\Delta\ [(\emptyset, \{\}), [1]]]$ $[\square]$
<proof>

end

16 Conclusion

We formalized a tree-based semantics for DPNs, where executions are modeled as hedges, that reflect the ordering of steps of each process and the causality due to process creation, but enforce no ordering between steps of processes running in parallel. We have shown how to efficiently compute predecessor sets of regular sets of configurations with tree-regular constraints on the execution hedges, by encoding a hedge-automaton into the DPN, thus reducing the problem to unconstrained predecessor set computation.

We have then formalized a generalization of acquisition histories to DPNs, and have shown its correctness. We have demonstrated how to use the generalized acquisition histories to describe the set of execution hedges, that have a lock-sensitive schedule, as a regular set. Thus we could use the techniques for hedge-constrained predecessor set computation to also compute lock-sensitive, hedge-constrained predecessor sets. Finally, we have defined a class of properties that can be computed using cascaded predecessor computations, and have applied our techniques to decide those properties for DPNs.

16.1 Trusted Code Base

In this section we shortly characterize on what our formal proof depends, i.e. how to interpret the information contained in this formal proof and the fact that it is accepted by Isabelle.

First of all, you have to trust the theorem prover and its axiomatization of HOL, the ML-platform, the operating system software and the hardware it runs on. All this components are able to cause false theorems to be proven.

Next, most of the theorems proven here have some implicit and explicit assumptions. The most critical assumptions are the assumptions of the locales, namely *DPN*, *LDPN*, *LDPN_c0*, and *encodedLDPN*. It is not formally proven that these assumptions make sense, and the locales really admit useful models. In Section 15 we give an example for a non-trivial DPN and formally prove that it satisfies our assumptions. This gives some evidence that our assumptions are not too restrictive.

The next crucial point – already discussed in the introduction – is, that we at some points claim that our methods are executable. However, we do not derive any executable code, and even if we did, the Isabelle code-generator can only guarantee *partial* correctness, i.e. correctness under the assumption of termination. At this point, the belief in the existence of executable methods depends on the belief in that the model-checking functions, i.e. the function *mc-pre* in *As-hc.thy* is effective for regular sets, and the result is a regular set again, such that we can check $c_0 \in \text{mc} - \text{pre}\Phi$ as required by Theorem *mc-pre-correct*, using the saturation algorithm of [2].

However, we prove some theorems that support this belief by showing how the required operations can be decomposed to operations that are well-known to be effective and to preserve regularity.

References

- [1] A. Bouajjani, J. Esparza, S. Schwoon, and J. Strejcek. Reachability analysis of multithreaded software with asynchronous communication. In *Proc. of FSTTCS'05*, pages 348–359. Springer, 2005.
- [2] A. Bouajjani, M. Müller-Olm, and T. Touili. Regular symbolic analysis of dynamic networks of pushdown systems. In *Proc. of CONCUR'05*. Springer, 2005.
- [3] V. Kahlon and A. Gupta. An automata-theoretic approach for model checking threads for LTL properties. In *Proc. of LICS 2006*, pages 101–110. IEEE Computer Society, 2006.
- [4] V. Kahlon, F. Ivancic, and A. Gupta. Reasoning about threads communicating via locks. In *Proc. of CAV 2005*, pages 505–518. Springer, 2005.