


Exceptional service in the national interest



The Q Compiler for Verifying High-Consequence Controls

Jon Aytac

6/27/18



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND NO. TBD

Motivation

- Suppose you believe that
 - Designs of High Consequence Systems should come with formal proofs of safety and reliability

Motivation

- Suppose you believe that
 - Designs of High Consequence Systems should come with formal proofs of safety and reliability
 - Or, at least, Designers of High Consequence Systems should be able to check whether their designs satisfy some safety and reliability properties

Motivation

- Suppose you believe that
 - Designs of High Consequence Systems should come with formal proofs of safety and reliability
 - Or, at least, Designers of High Consequence Systems should be able to check whether their designs satisfy some safety and reliability properties
- So you give a presentation to decision makers advocating the adoption of formal methodologies

Motivation

- Suppose you believe that
 - Designs of High Consequence Systems should come with formal proofs of safety and reliability
 - Or, at least, Designers of High Consequence Systems should be able to check whether their designs satisfy some safety and reliability properties
- So you give a presentation to decision makers advocating the adoption of formal methodologies
- They don't have time to read everything, but they like to stay abreast of what's going on in the literature...

and they show you this paper of Vanhoef and Piessens in 2017

Key Reinstallation Attacks: Forcing Nonce Reuse in WPA2

Mathy Vanhoef
imec-DistriNet, KU Leuven
Mathy.Vanhoef@cs.kuleuven.be

Frank Piessens
imec-DistriNet, KU Leuven
Frank.Piessens@cs.kuleuven.be

which describes security vulnerabilities in a protocol *proven secure* in this paper of He et al in 2005

A Modular Correctness Proof of IEEE 802.11i and TLS

Changhua He, Mukund Sundararajan, Anupam Datta, Ante Derek, John C. Mitchell
Electrical Engineering and Computer Science Departments,
Stanford University, Stanford, CA 94305-9045

He et al Proved Security Properties for a Composition of Protocols ...

Our proof consists of separate proofs of specific security properties for 802.11i components - the TLS authentication phase, the 4-Way Handshake protocol and the Group Key Handshake protocol. Using a new form of PCL composition

He et al

... The Attack didn't Violate Those Properties ...

Interestingly, our attacks do not violate the security properties proven in formal analysis of the 4-way and group key handshake. In particular, these proofs state that the negotiated session key remains private, and that the identity of both the client and Access Point (AP) is confirmed [39]. Our attacks do not leak the session

Vanhoef and Piessens

... The Crucial Properties Were Temporal...

key installation. Put differently, their models do not state when a negotiated key should be installed. In practice, this means the same key can be installed multiple times, thereby resetting nonces and replay counters used by the data-confidentiality protocol.

Vanhoef and Piessens

... and the 802.11i amendment didn't specify the protocols in such a way that questions about temporal properties could even be posed ...

The 802.11i amendment does not contain a formal state machine describing how the supplicant must implement the 4-way handshake. Instead, it only provides pseudo-code that describes how, but not when, certain handshake messages should be processed [4, §8.5.6].²

Vanhoef and Piessens

... Worse yet, the Map isn't the Territory

Another somewhat related work is that of Beurdouche et al. [14] and that of de Ruiter and Poll [27]. They discovered that several TLS implementations contained faulty state machines. In particular, certain implementations wrongly allowed handshake messages to be repeated. However, they were unable to come up with example

Vanhoef and Piessens

... Vanhoef and Piessens Proved CounterMeasure Correctness in NuSMV...

Proving the correctness of the above countermeasure is straightforward: we modeled the modified state machine in NuSMV [23], and used this model to prove that two key installations are always separated by the generation of a fresh PTK. This implies the same key is never installed twice. Note that key secrecy and session authentication was already proven in other works [39].

Vanhoef and Piessens

- So then the decision makers say

- So then the decision makers say
- We need a language for the specification of temporal behavior, preferably one designers would actually use

Jedi Mind Trick

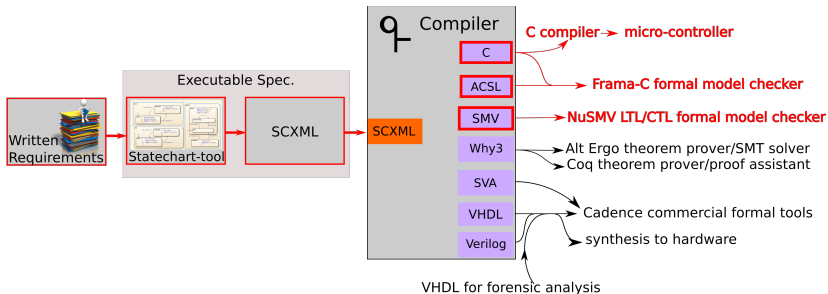
- So then the decision makers say
- We need a language for the specification of temporal behavior, preferably one designers would actually use
- We need to compile these descriptions into languages suitable for proving (e.g. Why3) and checking (e.g. NuSMV) temporal properties about compositions of specifications

- So then the decision makers say
- We need a language for the specification of temporal behavior, preferably one designers would actually use
- We need to compile these descriptions into languages suitable for proving (e.g. Why3) and checking (e.g. NuSMV) temporal properties about compositions of specifications
- For scalability's sake, the compiler should allow compositional reasoning about systems

- So then the decision makers say
- We need a language for the specification of temporal behavior, preferably one designers would actually use
- We need to compile these descriptions into languages suitable for proving (e.g. Why3) and checking (e.g. NuSMV) temporal properties about compositions of specifications
- For scalability's sake, the compiler should allow compositional reasoning about systems
- We need to prove those properties descend to implementations

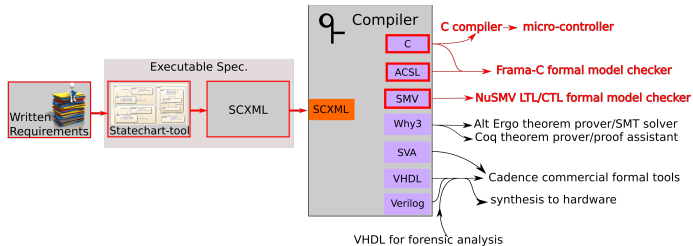
What a Coincidence

At this point, you say, we just so happen to have already written a tool, the Q compiler, to do just that!



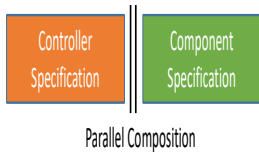
Q Compiler

Designers draw statechart-like diagrams in their specifications, so *Q* captures these in a statechart-like formal language



and compiles them into a multitude of languages. We focus here on how *Q* compiles into ACSL and SMV to enable a workflow with FramaC and NuSMV.

WorkFlow: Boxes and Arrows



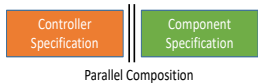
The designer writes abstract specifications for the controller and the component in a Statechart-like language, along with system level properties about their composition

WorkFlow: Boxes and Arrows

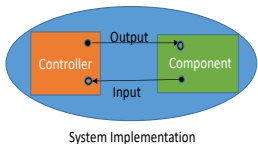


Q then generates an LTL model of the composition of abstract specifications along with the system level properties. These properties may now be checked with model checking tools like NuSMV

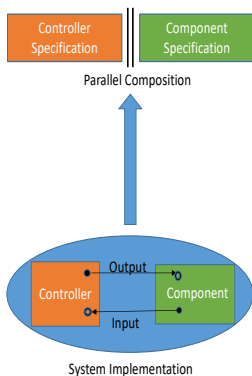
WorkFlow: Boxes and Arrows



The Engineers then implement the controller in C

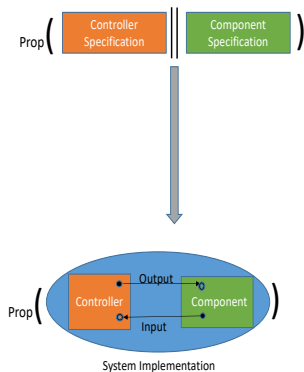


WorkFlow: Boxes and Arrows



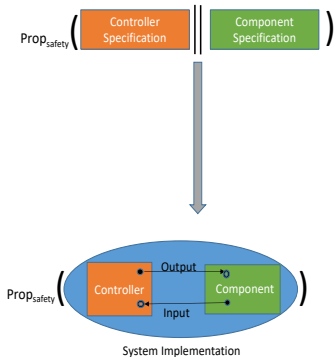
The Engineers need to present some kind of evidence that the specified system abstracts the system implementation

WorkFlow: Boxes and Arrows




In particular, the evidence of abstraction must be sufficient for proofs of properties about the abstract system to imply proofs of those properties about the system implementation

WorkFlow: Boxes and Arrows



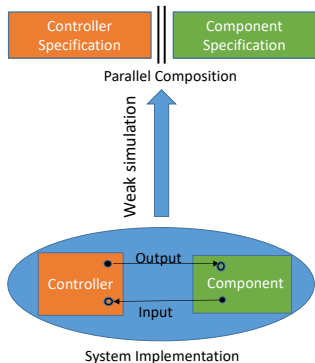
In our context, we are allowed to focus on stuttering invariant safety properties

WorkFlow: Boxes and Arrows



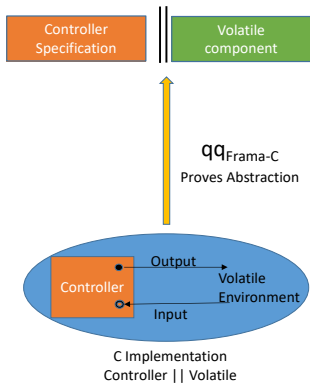
The engineer must therefore present evidence that the composition of specifications *weakly simulates* the composition of implementations

WorkFlow: Boxes and Arrows



Only Frama-C meets our requirements for reasoning about C programs, but Frama-C analyzes *sequential* programs. *Fortunately, our C program interacts with components through memory mapped I/O*

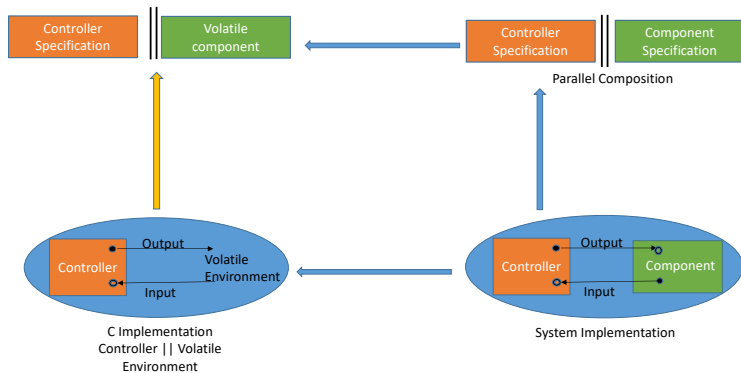
WorkFlow: Boxes and Arrows



So we ask the engineer to present evidence from which Q can construct , in ACSL, a proof obligation that *the composition of the abstract controller specification with a volatile environment weakly simulates the composition of the implementation with a volatile environment*

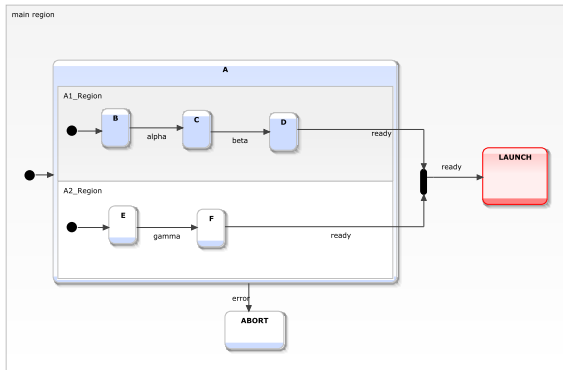
Workflow: Boxes and Arrows

The proof of this yellow, weak simulation arrow is discharged by Frama-C's Weakest Precondition plugin. The trick is to obtain the blue arrows by construction. Then proof of the yellow arrow gives us our goal, the right column, for free



The Boxes are StateChart-like

The designer writes in a Statechart-like language, by which we mean a specification language for reactive programs, inductively defined through hierarchic and parallel composition of transition systems.



The Boxes: LTS

A labelled transition system $P : \text{LTS}$ is a tuple $(S_P, \mathcal{A}_P, \rightarrow_P, P_0)$, with S_P the set of states, \mathcal{A}_P an alphabet, the transition relation $\rightarrow_P: \mathcal{A}_P \rightarrow \mathcal{P}(S_P \times S_P)$ a map from alphabet to relations on states, and P_0 a set of initial states

The Boxes: C Programs as LTS

Example

- a *program point* and an *execution environment* (which maps variables to values) together constitute the *program state* of a C program.

The Boxes: C Programs as LTS

Example

- a *program point* and an *execution environment* (which maps variables to values) together constitute the *program state* of a C program.
- The program is given by the data of a map from program points to commands

The Boxes: C Programs as LTS

Example

- a *program point* and an *execution environment* (which maps variables to values) together constitute the *program state* of a C program.
- The program is given by the data of a map from program points to commands
- Given a program state, the evaluation of the command found at that program state's program point defines a *transition* to new program state.

The Boxes: C Programs as LTS

Example

- a *program point* and an *execution environment* (which maps variables to values) together constitute the *program state* of a C program.
- The program is given by the data of a map from program points to commands
- Given a program state, the evaluation of the command found at that program state's program point defines a *transition* to new program state.
- The *label* of that transition is some predicate on the execution environment.

The Boxes: C Programs as LTS

Example

- a *program point* and an *execution environment* (which maps variables to values) together constitute the *program state* of a C program.
- The program is given by the data of a map from program points to commands
- Given a program state, the evaluation of the command found at that program state's program point defines a *transition* to new program state.
- The *label* of that transition is some predicate on the execution environment.
- In this way, we can think about C programs as labelled transition systems.

The Arrows: Simulation Relations

Definition

- For any $P, Q : \text{LTS}$ with $\mathcal{A}_P = \mathcal{A}_Q$, a relation $R \subseteq S_P \times S_Q$ is a **simulation relation** if and only if $\forall (p, q) \in R, \alpha \in \mathcal{A}_P, p' \in S_P$

$$p \xrightarrow{\alpha}_P p' \Rightarrow \exists q' \in S_Q : (q \xrightarrow{\alpha}_Q q' \wedge (p', q') \in R)$$

Q simulates P , written $P \preceq Q$ or $P \rightarrow Q$, if $P_0 \subset R^{-1}(Q_0)$.

The Arrows: Simulation Relations

Definition

- For any $P, Q : \text{LTS}$ with $\mathcal{A}_P = \mathcal{A}_Q$, a relation $R \subseteq S_P \times S_Q$ is a **simulation relation** if and only if $\forall (p, q) \in R, \alpha \in \mathcal{A}_P, p' \in S_P$

$$p \xrightarrow{\alpha}_P p' \Rightarrow \exists q' \in S_Q : (q \xrightarrow{\alpha}_Q q' \wedge (p', q') \in R)$$

Q simulates P , written $P \preceq Q$ or $P \rightarrow Q$, if $P_0 \subset R^{-1}(Q_0)$.

- The relation R is said to be a **witness** for $P \preceq Q$, and Q is said to be an **abstraction** of P . When the witness is a function, the simulation relation is said to be a **refinement**.

The Arrows: Simulation Relations

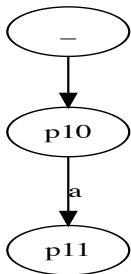
Proposition

This preorder is sound in that $P \preceq Q \Rightarrow \text{trace}(P) \subseteq \text{trace}(Q)$. So, if $P \preceq Q$, then for any $\mathbb{P}_Q : \text{Prop}(Q)$ temporal property on Q and $\mathbb{P}_P : \text{Prop}(P)$ the corresponding temporal property on P , $\mathbb{P}_Q \Rightarrow \mathbb{P}_P$.


An Example

Example

Here, $P_1 \preceq Q_1$



P_1



Q_1

Memory Mapped I/O, Volatile Variables and the Terminal Abstraction

Definition

The terminal abstraction over the alphabet \mathcal{A} is the labeled transition system $1_{\mathcal{A}_0} = (\star, \mathcal{A}, \rightarrow_{1_{\mathcal{A}}}, \star)$ with $\forall \alpha \in \mathcal{A}. \rightarrow_{\mathcal{A}}: \alpha \rightarrow \star \times \star$.

Memory Mapped I/O, Volatile Variables and the Terminal Abstraction

Definition

The terminal abstraction over the alphabet \mathcal{A} is the labeled transition system $1_{\mathcal{A}_0} = (\star, \mathcal{A}, \rightarrow_{1_{\mathcal{A}}}, \star)$ with $\forall \alpha \in \mathcal{A}. \rightarrow_{\mathcal{A}}: \alpha \rightarrow \star \times \star$.

Proposition

The terminal abstraction over a given alphabet abstracts any machine of the same alphabet $\forall P = (S_P, \mathcal{A}_P, \dot{\rightarrow}_P, P_0)$, $P \preceq 1_{\mathcal{A}_P}$

Memory Mapped I/O and the Terminal Abstraction


Example

Let “volatile uint8_t *I;” a declaration of a reference to a volatile variable of type uint8_t. Then,


$\mathcal{A}_I = \{(\star I == 0), \dots, (\star I == 255)\}$, and the presence of a volatile variable in a C program is an **asynchronous** parallel composition with $1_{\mathcal{A}_I}$, which we'll write as, $C \parallel^a 1_{\mathcal{A}_I}$

Proving Spec Simulates Implementation

So we will use the Q compiler and FramaC to prove



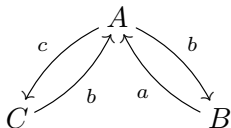
qFrama-C
Proves
Weak Simulation



C Implementation
Controller || Volatile
Environment

Proving Spec Simulates Implementation

So the designer gives a formal specification of their design in our statechart-like language, e.g. the ABC example



The implementer of the C program *should give a witness that their program is simulated by this abstract transition system.*

Proving Spec Simulates Implementation

A witness that Q simulates P ($P \preceq Q$ or $P \rightarrow Q$) decomposes into a relation on labels $R_{\mathcal{A}}$ and a relation on the states $R_{\mathcal{S}}$

$$\begin{array}{ccc}
 \mathcal{A}_P & \xrightarrow{\dot{\rightarrow}_P} & \mathcal{P}(S_P \times S_P) \\
 \downarrow f_{[R_{\mathcal{A}}]} & & \downarrow f_{[R_{\mathcal{S}}]} \\
 \mathcal{A}_Q & \xrightarrow{\dot{\rightarrow}_Q} & \mathcal{P}(S_Q \times S_Q)
 \end{array}$$

such that $f_{[R_{\mathcal{A}}]} \circ \dot{\rightarrow}_P \subseteq \dot{\rightarrow}_Q \circ f_{[R_{\mathcal{A}}]}$.

Proving Spec Simulates Implementation

- For C programs thought of as LTS, the labels are expressions over the execution environment.

```
type simMap = {  
  stateRels: stateRelAtom list;  
  exprRels: exprRelAtom list;  
  inputRels: inputRelAtom list;  
  intVarRels: intVarRelAtom list;  
  valueRels: valueRelAtom list;  
  ...  
}[@@deriving yojson]
```

Proving Spec Simulates Implementation

- For C programs thought of as LTS, the labels are expressions over the execution environment.
- so $R_{\mathcal{A}}$ may be given by a relation on internal variables, a relation on values, and a relation on inputs,

```
type simMap = {  
  stateRels: stateRelAtom list;  
  exprRels: exprRelAtom list;  
  inputRels: inputRelAtom list;  
  intVarRels: intVarRelAtom list;  
  valueRels: valueRelAtom list;  
  ...  
}[@@deriving yojson]
```


Proving Spec Simulates Implementation

- For C programs thought of as LTS, the labels are expressions over the execution environment.
- so $R_{\mathcal{A}}$ may be given by a relation on internal variables, a relation on values, and a relation on inputs,
- So the implementer proposes $R_{\mathcal{A}}$ and $R_{\mathcal{S}}$ via a JSON file described by the OCaml `simMap` type below

```
type simMap = {  
  stateRels: stateRelAtom list;  
  exprRels: exprRelAtom list;  
  inputRels: inputRelAtom list;  
  intVarRels: intVarRelAtom list;  
  valueRels: valueRelAtom list;  
  ...  
}[@@deriving yojson]
```

Proving Spec Simulates Implementation

R_S should be pairs of program states in the abstract and concrete program. For our simple ABC example, the relevant execution context is held by a type

```
struct machine {  
    enum states currState;  
    enum states nextState;  
    uint8_t input;  
};  
typedef struct machine *machine_t;
```

Proving Spec Simulates Implementation

In the ABC example, the particular instance of the type carrying the execution environment part of the program state is declared in the C as

```
struct machine theMac;  
  
machine *theMachine() {  
    return &theMac;  
}
```

so the implementer specifies this in the JSON as

```
"sMInstance": {"typeName": "machine_t", "instanceName": "theMac"},
```

Proving Spec Simulates Implementation

In the ABC example, the initial and final program states in the abstraction correspond in the C program to the entry and exit program points given the execution environment

```
void action_s00(machine_t mac) {
    mac->currState = 0x00;
    printf("State A (0x%02x)\n", mac->currState);
    if(!read_packet(mac)) {
        error();
    }
    if(mac->input == 'b') {
        mac->nextState = SB;
    }
    else if(mac->input == 'c') {
        mac->nextState = SC;
    }
    else {
        printf("Error on input: %c\n", mac->input);
    }
}
```

Proving Spec Simulates Implementation

So the implementer gives the relation on states as

```

"stateRels": [
  {
    "abStates": [
      {"stateName": "SA"}
    ],
    "impStates": [
      {"stateName": "SA", "funcName": "acti\
on_s00", "args": [{"typeName": "machine_t", "instanc\
eName": "mac"}]}
    ]
  },

```

Proving Spec Simulates Implementation

there being, in this example, no internal variables, the map on alphabets factors through a map on input variables

```



```

Proving Spec Simulates Implementation

and values

```
    "valueRels": [  
      {  
        "abValues" : [{"value": "AA"}],  
        "impValues": [{"value": "'a'"}]  
      },  
      {  
        "abValues" : [{"value": "BB"}],  
        "impValues": [{"value": "'b'"}]  
      },  
      {  
        "abValues" : [{"value": "CC"}],  
        "impValues": [{"value": "'c'"}]  
      }  
    ],
```

Proving Spec Simulates Implementation

Q then generates the ACSL annotations positing the pre and post conditions from the specification and the simulation relation. A short clang program searches through the C implementation's code base for the program points specified by the stateRel and annotates:

```

/*@
requires theMac->currState == SA;
behavior behPinner_unit:
  assumes((theMac->currState == SA)&&\true&&\true&&(!(((getInput(theMac)) == (('c')))) &
& (!(((getInput(theMac)) == (('b')))))) && (!(((getInput(theMac)) == (('c')))) && (!(((getInp
ut(theMac)) == (('b')))))));
  ensures((theMac->currState == SA));

behavior behPouter_SA_SBunit:
  assumes((theMac->currState == SA)&&\true&&\true&&((getInput(theMac)) == (('b')));
  ensures((theMac->currState == SB));

behavior behPouter_SA_SCunit:
  assumes((theMac->currState == SA)&&\true&&\true&&((getInput(theMac)) == (('c')));
  ensures((theMac->currState == SC));

behavior behPouter_Pcomplete_SAunit:
  assumes((theMac->currState == SA)&&\true&&\true&&(!(((getInput(theMac)) == (('c')))) && (!\
(((getInput(theMac)) == (('b')))))));
  ensures((theMac->currState == SA));
disjoint behaviors;*/
void action_s00(struct machine *theMac);
  
```


Proving Spec Simulates Implementation

The C program interacts with its environment through memory mapped IO

```
volatile uint8_t *fgetC = (uint8_t *)INPUT_ADDRESS;
int read_packet(machine_t mac) {
    uint8_t c = *fgetC;
    mac->input = (char)c;
    return ('a' <= c && c <= 'z');
}
```

Proving Spec Simulates Implementation

so the implementer must relate the interface to the environment in the abstraction to a memory mapped I/O interface in the implementation of a given type

```
    "interfaceRels": [  
      {  
        "abInt": {"abInterfaceLabel": "input"},  
        "impInt": {"impInterfaceVarPtr": {"typeNa  
e": "uint8_t", "instanceName": "fgetC"}}  
      }  
    ],
```

Proving Spec Simulates Implementation

Q generates the ACSL describing the corresponding terminal abstraction $1_{\mathcal{A}_{fgetc}}$, and a short clang program searches through the codebase for the declaration of a volatile variable of that name and that type

```
volatile uint8_t *fgetC = (uint8_t *)INPUT_ADDRESS;
/*@ ghost //@ requires fgetCArg == fgetC;
@ uint8_t readfgetC(volatile uint8_t *fgetCArg) {
@ static uint8_t injectorfgetCBuffer[256];
@ static uint8_t injectorfgetCBufferCount;
@ for (int i=0; i<256; i++){
@   injectorfgetCBuffer[i]=i;
@ }
@ if (fgetC == fgetCArg)
@   return injectorfgetCBuffer[(injectorfgetCBufferCount++)%256];
};
@ else
@   return 0;
@ }
@ */
/*@ ghost uint8_t injectorfgetCCollector[256];
/*@ ghost uint8_t fgetCCollectorCount = 0;
/*@ ghost //@ requires fgetCArg == fgetC;
@ uint8_t writefgetC(volatile uint8_t *fgetCArg, uint8_t v) {
@   if (fgetCArg == fgetC)
@     return injectorfgetCCollector[(fgetCCollectorCount++)%256];
@   else
@     return 0;
@ }
@ }
```

Proving Spec Simulates Implementation

and defines a collection of ACSL predicates

```
predicate Pinner_unit(integer currState, integer nextState, struct machine * theMac) = (currState == SA) && (nextState == SA) && \true && \true && (!(((getInput(theMac)) == ('c')))) && (!(((getInput(theMac)) == ('b')))) && (!(((getInput(theMac)) == ('c')))) && (!(((getInput(theMac)) == ('b'))));
```

Using the Q Compiler to Prove Spec Simulates Implementation

which form clauses in the transition relation

```
predicate theMacTransitionRelation(integer currState, integer nextState, machi`  
ne_t theMac) = Pinner_unit(currState, nextState, theMac) || Pinner_unit_1(currSt`  
ate, nextState, theMac) || Pinner_unit_2(currState, nextState, theMac) || Pouter`  
_SA_SBUnit(currState, nextState, theMac) || Pouter_SA_SCUnit(currState, nextStat`  
e, theMac) || Pouter_Pcomplete_SAUnit(currState, nextState, theMac) || Pouter_SB`  
_SAUnit(currState, nextState, theMac) || Pouter_Pcomplete_SBUnit(currState, next`  
State, theMac) || Pouter_SC_SBUnit(currState, nextState, theMac) || Pouter_Pcomp`  
lete_SCUnit(currState, nextState, theMac);
```

Using the Q Compiler to Prove Spec Simulates Implementation

in the ABC example, the transition relation in this design is implemented as an actual function on the execution environment

```
void step(machine_t mac) {
  switch(mac->nextState) {
    case SA: action_s00(mac); break;
    case SB: action_s01(mac); break;
    case SC: action_s11(mac); break;
    default: error(); break;
  }
}
```

and the implementer has related the abstract transition function to this program point

```
"transitionFunction": {"funcName": "step", "args": [
{"typeName": "machine_t", "instanceName": "theMac"]}]}
```

Using the Q Compiler to Prove Spec Simulates Implementation

For any invocation of that function found within a while loop, the clang program annotates with the corresponding loop invariant

```
int main() {
    machine_t theMac = theMachine();
    theMac->nextState = SA;
    /*@
     loop assigns *theMac; loop invariant theMacTransitionRelation(\
     \at(theMac->currState, Pre), theMac->currState, theMac);*/
    while(1) {
        step(theMac);
    }
    return 0;
}
```

We then pose this invariant as a proof obligation to FramaC's WP plugin.

Proving LTL properties about the Spec

The Q tool produces from the same abstraction a description of the same transition system as a collection of LTL constraints, e.g. in NuSMV

```

INIT
(st = SA) & (input in {AA, BB, CC, EE});

DEFINE
Pinner_unit := (st = SA) & (next(st) = SA) & (!((!(input) = (CC))) & (!((input) = (BB)))) & (!((input) = (CC))) & (!((input) = (BB)));

Pinner_unit_1 := (st = SB) & (next(st) = SB) & ((input) = (AA)) & (!((input) \
= (AA)));

Pinner_unit_2 := (st = SC) & (next(st) = SC) & ((input) = (BB)) & (!((input) \
= (BB)));

Pouter_SA_SB$unit := (st = SA) & (next(st) = SB) & (input) = (BB);

Pouter_SA_SC$unit := (st = SA) & (next(st) = SC) & (input) = (CC);

Pouter_#complete_SA$unit := (st = SA) & (next(st) = SA) & (!((input) = (CC))) \
& (!((input) = (BB)));

Pouter_SB_SA$unit := (st = SB) & (next(st) = SA) & (input) = (AA);

Pouter_#complete_SB$unit := (st = SB) & (next(st) = SB) & (!((input) = (AA)));

Pouter_SC_SB$unit := (st = SC) & (next(st) = SB) & (input) = (BB);

Pouter_#complete_SC$unit := (st = SC) & (next(st) = SC) & (!((input) = (BB));
  
```


Proving LTL properties about the Spec

Any proof obligations the designer may have posited

```

<!-- Global Properties to Satisfy -->
<!-- System exits SA only if input is b or c -->
<iuymb:pragma key="SMV_LTLN" value="Req001 := G((st=SA & (input!=BB & input!=CC)) -> X(st)

<!-- System exits SB only if input is a -->
<iuymb:pragma key="SMV_LTLN" value="Req002 := G((st=SB & input!=AA) -> X(st)=SB)" />

<!-- System exits SC only if input is b -->
<iuymb:pragma key="SMV_LTLN" value="Req003 := G((st=SC & input!=BB) -> X(st)=SC)" />
  
```

can be checked in NuSMV

```

LTLSPEC NAME
Req001 := G((st=SA & (input!=BB & input!=CC)) -> X(st)=SA);

LTLSPEC NAME
Req002 := G((st=SB & input!=AA) -> X(st)=SB);

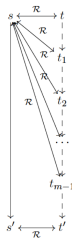
LTLSPEC NAME
Req003 := G((st=SC & input!=BB) -> X(st)=SC);
  
```

What have we proved?

When FramaC's WP plugin discharges these proof obligations, the implementer's `simMap` witnesses *weak simulation*

```
behavior behPouter_SA_SBunit:
  assumes((theMac->currState == SA)&&\true&&\true&&
    ((theMac->input) == (('b')));
```

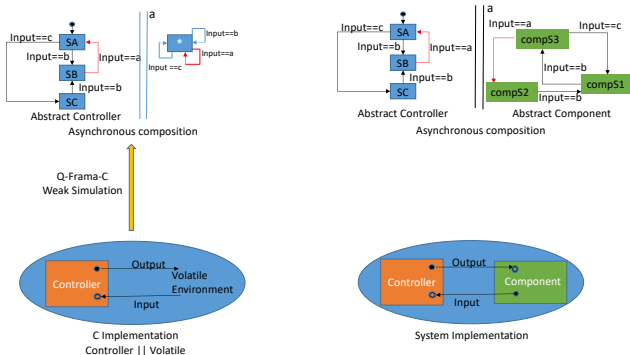
```
ensures((theMac->currState == SB));
```



```
void action_s00(machine_t mac) {
  mac->currState = 0x00;
  printf("State A (0x%02x)\n", mac->currState);
  if(!read_packet(mac)) {
    error();
  }
  if(mac->input == 'b') {
    mac->nextState = SB;
  }
  else if(mac->input == 'c') {
    mac->nextState = SC;
  }
  else {
    printf("Error on input: %c\n", mac->input);
  }
}
```

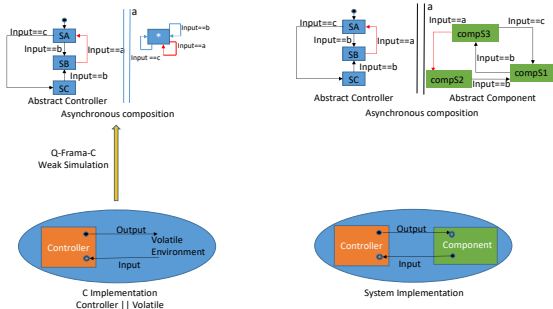
Proving LTL Properties of Compositions with Implementations

- So Q with FramaC proves weak simulation (in yellow) of the implementation asynchronously composed with a volatile environment by the asynchronous composition of the spec with a volatile environment



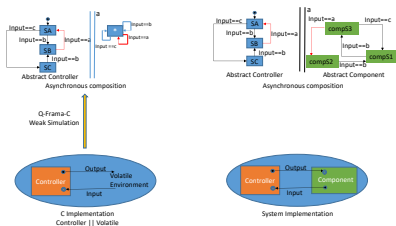
Proving LTL Properties of Compositions with Implementations

- Q with FramaC proves weak simulation (in yellow) of the implementation composed with a volatile environment by the asynchronous composition of the spec with the environment
- Q can generate models and LTL proof obligations for a composition of abstractions (upper right hand corner)



Proving LTL Properties of Compositions with Implementations

- Q with FramaC proves weak simulation (in yellow) of the implementation composed with a volatile environment by the asynchronous composition of the spec with the environment
- Q can generate models and LTL proof obligations for a composition of abstractions (upper right hand corner)
- *What does that say about the implementation (lower right corner)?*



The \parallel between Boxes: Parallel Composition of LTS

To reduce state space, and to model the composition of hardware within clock domains, we actually have in Q a synchronous composition

Definition

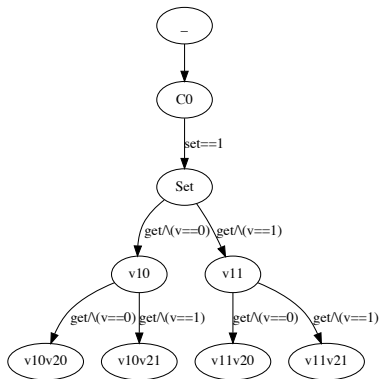
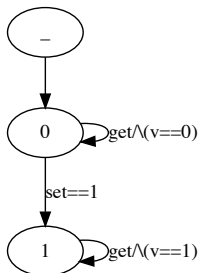
The Synchronous composition of $P, Q : \text{LTS}$, written $P \parallel Q$ is given by the rule

$$\frac{p \xrightarrow{\alpha}_P p' \quad q \xrightarrow{\alpha}_Q q'}{(p, q) \xrightarrow{\alpha}_{P \parallel Q} (p', q')}$$

An example

Example

$\forall q \in S_Q$, the synchronous composition of these two LTS's will never visit $(v10v21, q)$ or $(v11v20, q)$



Moving Arrows past \parallel : Precongruence

Our synchronous \parallel was a convenient choice for reducing statespace and for modeling hardware. More importantly, now, is that it is *compositional*, in that

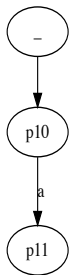
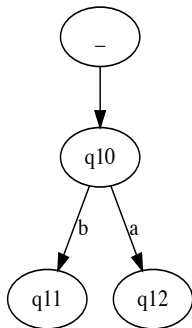
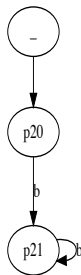
Lemma

Strong simulation \preceq is a pre-congruence over \parallel , that is,
 $\forall O, P, Q : \text{LTS}. P \preceq Q \Rightarrow P \parallel O \preceq Q \parallel O$

An Example

Example

Here, $P_1 \preceq Q_1$, and $P_1 \parallel P_2 \preceq Q_1 \parallel P_2$

 P_1  Q_1  P_2

Synchronous, Asynchronous, Strong, Weak

We can recover asynchronous composition *and* weak simulation through a completion with the right alphabet

Definition

Let \mathcal{A} an alphabet, and let $M_{\mathcal{A}} : \text{LTS} \rightarrow \text{LTS} :: P \mapsto P$ the completion over \mathcal{A} be the transition system $(S_P, \mathcal{A}_P \sqcup \mathcal{A}, \rightarrow, S_0)$ with the new transition relation given by rules

$$\frac{p \xrightarrow{\alpha}_P p'}{p \xrightarrow{\alpha}_{M_{\mathcal{A}}(P)} p'} \qquad \frac{\alpha \in \mathcal{A}}{p \xrightarrow{\alpha}_{M_{\mathcal{A}}(P)} p} \qquad (2.1)$$

Synchronous, Asynchronous, Strong, Weak

Let $P, Q : \text{LTS}$ with alphabets \mathcal{A}_P and \mathcal{A}_Q . We can recover the asynchronous composition from the synchronous composition via the completion once given a decomposition of their meet $\mathcal{A}_P \sqcap \mathcal{A}_Q$ into disjoint \mathcal{A}_{O_P} (labels output by P) and \mathcal{A}_{O_Q} (labels output by Q). Let $\mathcal{A}_{Q/P} := \mathcal{A}_{O_Q} \sqcup (\mathcal{A}_Q \setminus \mathcal{A}_P)$ the asynchronous alphabet of Q over P

Synchronous, Asynchronous, Strong, Weak

Let P, Q : LTS with alphabets \mathcal{A}_P and \mathcal{A}_Q . We can recover the asynchronous composition from the synchronous composition via the completion once given a decomposition of their meet $\mathcal{A}_P \sqcap \mathcal{A}_Q$ into disjoint \mathcal{A}_{O_P} (labels output by P) and \mathcal{A}_{O_Q} (labels output by Q). Let $\mathcal{A}_{Q/P} := \mathcal{A}_{O_Q} \sqcup (\mathcal{A}_Q \setminus \mathcal{A}_P)$ the asynchronous alphabet of Q over P

Proposition

Let Q, P : LTS share alphabet $\mathcal{A}_P = \mathcal{A}_Q$ and let \mathcal{A}_E the alphabet of their environment. Q weakly simulates P (i.e. $P \preceq_W Q$) if and only if $P \preceq M_{\mathcal{A}_{E/Q}}(Q)$, or $P \rightarrow M_{\mathcal{A}_{E/Q}}(Q)$

Synchronous, Asynchronous, Strong, Weak

Lemma

The asynchronous composition is in strong simulation's kernel with the synchronous composition of completions

$$\forall P, Q : LTS.P \parallel^a Q \simeq \mathbf{M}_{\mathcal{A}_{Q/P}}(P) \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(Q)$$

Synchronous, Asynchronous, Strong, Weak

Lemma

The asynchronous composition is in strong simulation's kernel with the synchronous composition of completions

$$\forall P, Q : LTS.P \parallel^a Q \simeq \mathbf{M}_{\mathcal{A}_{Q/P}}(P) \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(Q)$$

Lemma

$$P \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q}) \simeq P$$

Synchronous, Asynchronous, Strong, Weak

Lemma

The asynchronous composition is in strong simulation's kernel with the synchronous composition of completions

$$\forall P, Q : LTS. P \parallel^a Q \simeq \mathbf{M}_{\mathcal{A}_{Q/P}}(P) \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(Q)$$

Lemma

$$P \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q}) \simeq P$$

Lemma

Q asks FramaC-WP to prove

$$\mathbf{M}_{\mathcal{A}_{Q/P}}(P \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q})) \stackrel{\text{qFramaC}}{\leftarrow} \mathbf{M}_{\mathcal{A}_{Q/P}}(C) \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q})$$

Soundness and Compositionality of Q

- By terminality of the terminal abstraction, $1_{\mathcal{A}_Q} \stackrel{!}{\leftarrow} Q$.

Soundness and Compositionality of Q

- By terminality of the terminal abstraction, $1_{\mathcal{A}_Q} \stackrel{!}{\leftarrow} Q$.
- By functoriality of the completion,

$$M_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q}) \stackrel{M_{\mathcal{A}_{P/Q}}}{\leftarrow} M_{\mathcal{A}_{P/Q}}(Q)$$

Soundness and Compositionality of Q

- By terminality of the terminal abstraction, $1_{\mathcal{A}_Q} \stackrel{!}{\leftarrow} Q$.
- By functoriality of the completion,

$$M_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q}) \stackrel{M_{\mathcal{A}_{P/Q}}}{\leftarrow} M_{\mathcal{A}_{P/Q}}(Q)$$
- Therefore, by precongruence of simulation (\succeq or \leftarrow) over \parallel ,

Soundness and Compositionality of Q

- By terminality of the terminal abstraction, $1_{\mathcal{A}_Q} \stackrel{!}{\leftarrow} Q$.

- By functoriality of the completion,

$$M_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q}) \stackrel{M_{\mathcal{A}_{P/Q}}}{\leftarrow} M_{\mathcal{A}_{P/Q}}(Q)$$

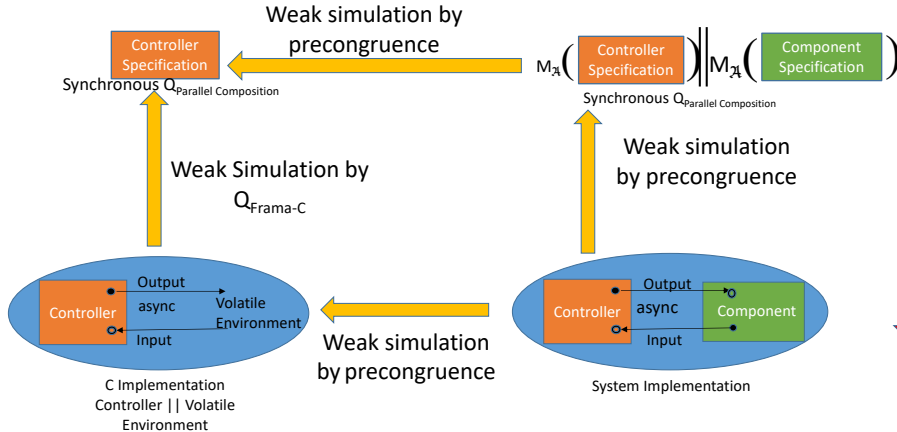
- Therefore, by precongruence of simulation (\succeq or \leftarrow) over \parallel ,

Theorem (Main Theorem)

The following is a commuting diagram of LTS with, as maps of LTS, simulations

$$\begin{array}{ccc}
 M_{\mathcal{A}_{Q/P}}(P \parallel M_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q})) & \stackrel{M_{\mathcal{A}_{Q/P}}(!)}{\leftarrow} & M_{\mathcal{A}_{Q/P}}(P) \parallel M_{\mathcal{A}_{P/Q}}(Q) \\
 \uparrow q_{\text{FramaC}} & & \uparrow \\
 M_{\mathcal{A}_{Q/P}}(C) \parallel M_{\mathcal{A}_{P/Q}}(1_{\mathcal{A}_Q}) & \stackrel{!}{\leftarrow} & M_{\mathcal{A}_{Q/P}}(C) \parallel M_{\mathcal{A}_{P/Q}}(Q)
 \end{array}$$

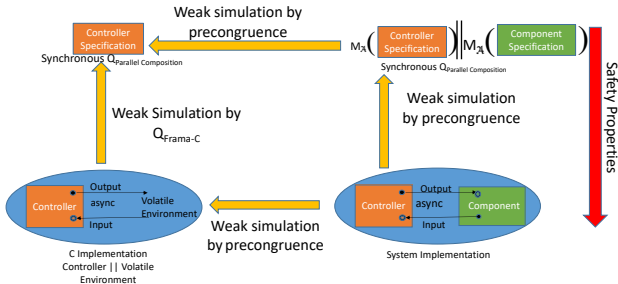
What Was Accomplished



What Was Accomplished

Thus we have accomplished our goal -

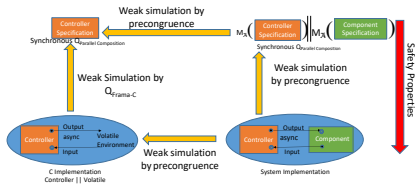
- Given a proposed witness of simulation of a controller implementation by its spec, the Q compiler generates and emplaces the ACSL annotations from which FramaC generates the proof of simulation.



What Was Accomplished

Thus we have accomplished our goal -

- Given a proposed witness of simulation of a controller implementation by its spec, the Q compiler generates and replaces the ACSL annotations from which FramaC generates the proof of simulation.
- The Q compiler computes the *asynchronous* composition of that specification with specifications for other components and constructs LTL models of the compositions along with high level properties in e.g. Why3 if they are to be proved or e.g. NuSMV if they are to be checked.



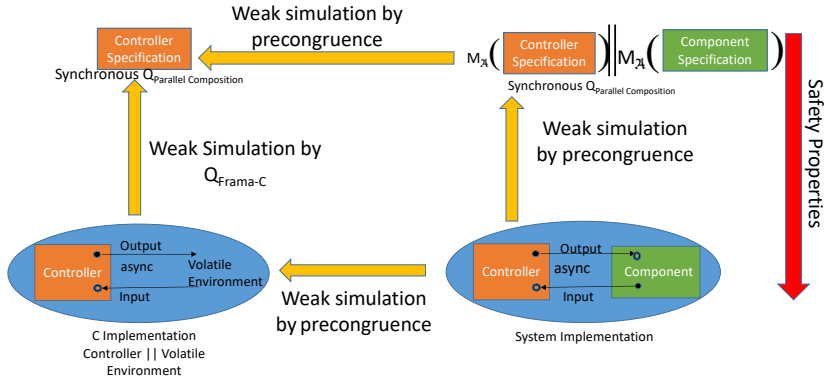
What Was Accomplished

Thus we have accomplished our goal -

- Given a proposed witness of simulation of a controller implementation by its spec, the Q compiler generates and emplaces the ACSL annotations from which FramaC generates the proof of simulation.
- The Q compiler computes the composition of that specification with (the completion of the) specifications for other components, constructs LTL models of the compositions along with high level properties in e.g. Why3 if they are to be proved or e.g. NuSMV if they are to be checked.
- By the Main Theorem, any Safety Properties proved about that composition of specifications implies the same property about the actual System implementation



What Was Accomplished



What is Next for Q?

- Q is currently proprietary. However, there is internal discussion about potentially open sourcing it
- We would like to mechanize the metatheory: ideally Q would provide certificates in, e.g. Coq, that it is doing the right thing. For that it would be nice to have a formal semantics of ACSL.

Thanks!



Robert Armstrong



Jon M. Aytac



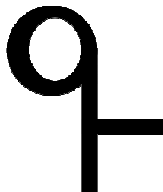
Geoffrey Compton Hulette



Benjamin Curt Nilsen



Philip Alden Johnson-Freyd



Jackson Mayo



Jason Michnovicz



Karla Morris



Ratish J. Punnoose



Andrew Michael Smith

Bonus Slide: What about Aeorai

- There is already a plugin for producing LTL properties about C programs: Aeorai

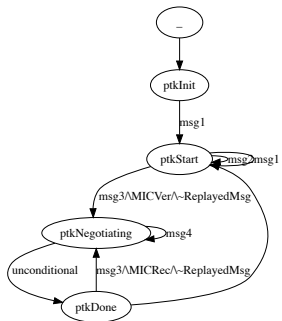
Bonus Slide: What about Aeorai

- There is already a plugin for producing LTL properties about C programs: Aeorai
- We could use this for our purposes only after solving the following problem

Bonus Slide: What about Aeorai

- There is already a plugin for producing LTL properties about C programs: Aeorai
- We could use this for our purposes only after solving the following problem
- Let $\mathbb{P}_{P\parallel^a Q} : \text{Prop}(P\parallel^a Q)$ a temporal property on the asynchronous composition of P and Q . Given the specification of Q , what temporal property $\mathbb{P}_P : \text{Prop}(P)$ suffices to obtain $\mathbb{P}_{P\parallel^a Q}$. This is sometimes called the problem of producing quotient specifications

Bonus Slide: What about WPA?



Bonus Slide: Lattice Structure on Labels

The labels in a C program will have some lattice structure, i.e. a Boolean algebra of expressions over its variables, and this puts some constraints on $\dot{\rightarrow}: \mathcal{A} \rightarrow \mathcal{P}(S \times S)$,

$$\frac{\beta \Rightarrow \alpha \quad p \xrightarrow{\alpha}_C p'}{\quad \beta} \quad \frac{p \xrightarrow{\alpha} p' \quad p \xrightarrow{\beta}_C p'}{\quad \alpha \vee \beta} \quad \frac{}{p \xrightarrow{\perp}_C p'} \quad (2.2)$$

Bonus Slide: Lattice Structure on Labels

The labels in a C program will have some lattice structure, i.e. a Boolean algebra of expressions over its variables, and this puts some constraints on $\dot{\rightarrow}: \mathcal{A} \rightarrow \mathcal{P}(S \times S)$,

$$\frac{\beta \Rightarrow \alpha \quad p \xrightarrow{\alpha}_C p'}{\quad \beta} \quad \frac{p \xrightarrow{\alpha} p' \quad p \xrightarrow{\beta}_C p'}{\quad \alpha \vee \beta} \quad \frac{}{p \xrightarrow{\perp}_C p'} \quad (2.2)$$

So $\dot{\rightarrow}$ is a Galois Connection. $f_{[R_{\mathcal{A}}]}$ must therefore be monotonic for our witness to be viable.

Bonus Slide: Lattice Structure on Labels

- The labels in the abstraction are the expressions over abstract variables Var_Q taking values in abstract value domain \mathbb{V}_Q .

Bonus Slide: Lattice Structure on Labels

- The labels in the abstraction are the expressions over abstract variables Var_Q taking values in abstract value domain \mathbb{V}_Q .
- if $f_{[R, \mathcal{A}]}$ factors through a map of variables $f_{\text{Var}} : \text{Var}_Q \rightarrow \mathcal{P}(\text{Var}_P)$ and a map of their value domains $f_{\mathbb{V}} : \mathbb{V}_Q \rightarrow \mathcal{P}(\mathbb{V}_P)$, then we can know $f_{[R, \mathcal{A}]}$ is monotonic, no further proof obligations required.

Bonus Slide: Lattice Structure on Labels

- The labels in the abstraction are the expressions over abstract variables Var_Q taking values in abstract value domain \mathbb{V}_Q .
- if $f_{[R,\mathcal{A}]}$ factors through a map of variables $f_{\text{Var}} : \text{Var}_Q \rightarrow \mathcal{P}(\text{Var}_P)$ and a map of their value domains $f_{\mathbb{V}} : \mathbb{V}_Q \rightarrow \mathcal{P}(\mathbb{V}_P)$, then we can know $f_{[R,\mathcal{A}]}$ is monotonic, no further proof obligations required.
- Otherwise, checking monotonicity of $f_{[R,\mathcal{A}]}$ requires reasoning about the lattice structure of the algebra of boolean expressions over the execution environment of the C program

So we present the data type provide required to construct the simulation witness R as well as the information to generate $1_{\mathcal{A}_I}$ in a format encouraging the designer to give the witness without incurring additional proof obligations

```
type simMap = {  
  stateRels: stateRelAtom list;  
  exprRels: exprRelAtom list;  
  inputRels: inputRelAtom list;  
  intVarRels: intVarRelAtom list;  
  valueRels: valueRelAtom list;  
  . . .  
}[@@deriving yojson]
```

Completion is not Strong

The completion is not, however, a strong endofunctor on LTS

$$\mathbf{M}_{\mathcal{A}_{Q/P}}(P \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(Q)) \rightarrow \mathbf{M}_{\mathcal{A}_{Q/P}}(P) \parallel \mathbf{M}_{\mathcal{A}_{P/Q}}(Q)$$