

Controlled Owicki-Gries Concurrency:

Reasoning about the Preemptible eChronos Embedded Operating System

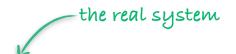
MARS 2015

<u>June Andronick</u>, Corey Lewis, Carroll Morgan November 2015



Aim





Formal model of the **eChronos** embedded RT OS

which requires **Concurrency** reasoning

the challenge

for which we use **Owicki-Gries** method



eChronos







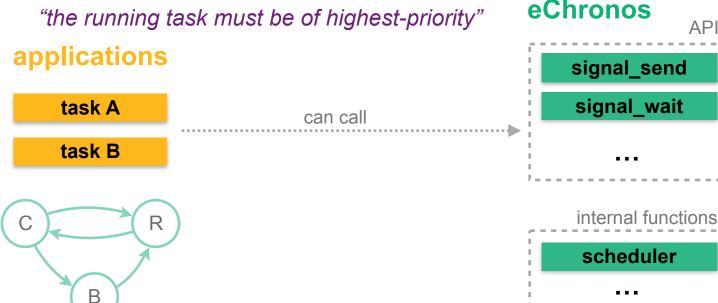
Small real-time OS library (~500 SLOC) What:

(Joint development with Breakaway Consulting)

Where: Embedded devices, with limited resources, no memory-protection

Job: provides synchronisation primitives

> schedules tasks according to priorities "the running task must be of highest-priority"



C=current

R=runnable B=blocked

eChronos







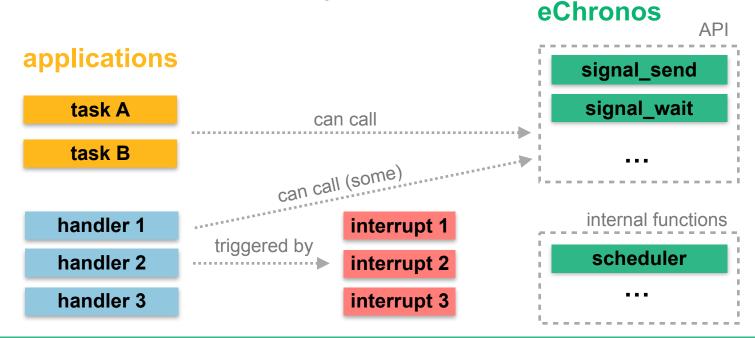
What: Small real-time OS library (~500 SLOC)

(Joint development with Breakaway Consulting)

Where: Embedded devices, with limited resources, no memory-protection

Job: • provides synchronisation primitives

schedules tasks according to priorities



eChronos



Characteristics: → small, fast

Aim: ▶ verified

Challenges:
• tasks can be preempted by another task

OS can be interrupted by external event

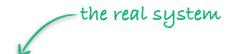


Our approach applies to OS systems with these characteristics:

- tasks can be preempted by higher priority tasks
- OS code can be interrupted by external event

Aim





Formal model of the **eChronos** embedded RT OS

which requires **Concurrency** reasoning

the challenge

for which we use **Owicki-Gries** method





What: Extension of Hoare logic to shared-variable parallel programs

(Suzanne Owicki and David Gries, 1976)

cobegin
$$S1 / |S2 / | ... | |Sn$$
 coend await B then S

Why: Small, fast



How:

```
Program P
{is_odd x}
x:=x+1;
{is_even x}
x:=x+1;
{is_odd x}
```

P is (sequentially) correct

Local correctness

P' does not interfere with P Interference freedom

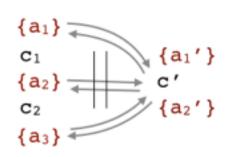
P is globally correct



From a parallel composition of fully annotated programs, generates correctness verification conditions

▶ local correctness prove each {a_i} c_i {a_{i+1}}

• interference freedom for each assertion a in P, and each command c' in P', prove that {a ∧ a'} c' {a}



- ! quadratic explosion of proof obligations, not compositional
 - For our system, interleaving is more controlled
 use automation power of modern theorem provers



Leonor Prensa Nieto

2002



Aim:

use OG to model interleaving between
 eChronos code, interrupt code, and tasks

Challenges:

- tasks are not 1st class citizens
- concurrency is uncontrolled

Contributions



Formal model of eChronos

using Owicki-Gries method

Challenges:

- tasks can be preempted
- OS can be interrupted

Challenges:

- ▶ tasks are not 1st class citizens
- concurrency is uncontrolled

Contributions



Tasks as 1st class citizens in O-G "AWAIT-painting"



Controlled interleaving model

Hardware API model

3

Formal model of eChronos

to prove that the running task is always the highest-priority runnable task

Formalised in Isabelle/HOL theorem prover





Tasks as 1st class citizens in O-G "AWAIT-painting"

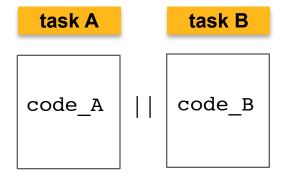
Controlled interleaving model

Hardware API model

Formal model of eChronos
to prove that the running task is always
the highest-priority runnable task

Context switching between tasks





No arbitrary concurrency between all these programs Can only switch from A to B if B become the active task

Context switching between tasks



task A task B a₁; a₂; a₃; b₁; b₂; b₃;

- → We introduce: Variable AT (Active Task)
- → We "AWAIT-paint" all statements:

```
AWAIT AT=A THEN a<sub>1</sub>;

AWAIT AT=A THEN a<sub>2</sub>;

AWAIT AT=B THEN b<sub>1</sub>;

AWAIT AT=B THEN b<sub>2</sub>;

AWAIT AT=B THEN b<sub>3</sub>;
```

➡ We automate this with an "await_paint task-id code" command:

```
await_paint A code_A; || await_paint B code_B;
```

→ We model context switch: context_switch task_id = AT:=task_id;

Tasks as 1st class citizens



Owicki-Gries

1975

Leonor Prensa Nieto 2002



+

AT

await_paint

context_switch



Tasks as 1st class citizens in O-G "AWAIT-painting"

Controlled interleaving model

Hardware API model

Formal model of eChronos
to prove that the running task is always
the highest-priority runnable task

Recall eChronos



applications

task A

task B

handler 1

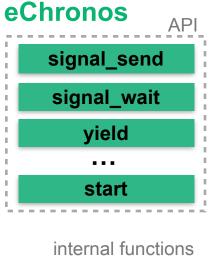
handler 2

handler 3

interrupt 1

interrupt 2

interrupt 3

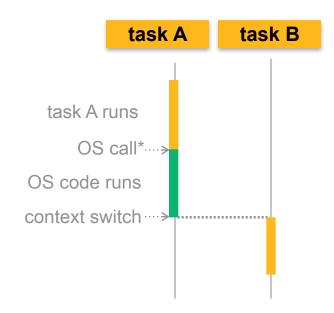


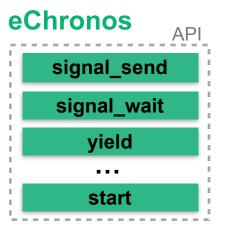
scheduler



A task runs until:

1 calling a OS
API function
that changes
runnable tasks



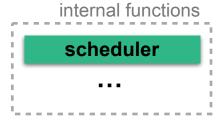


* Note: no mode switch between OS and task in such constrained hardware.

handler 2

handler 1

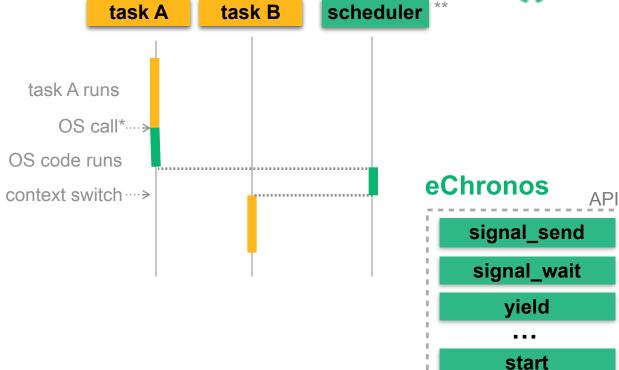
handler 3





A task runs until:

1 calling a OS
API function
that changes
runnable tasks



* Note: no mode switch between OS and task in such constrained hardware.

** Note: using ARM "supervisor call" (svc) mechanism handler 1 handler 2

handler 3

internal functions
scheduler
...

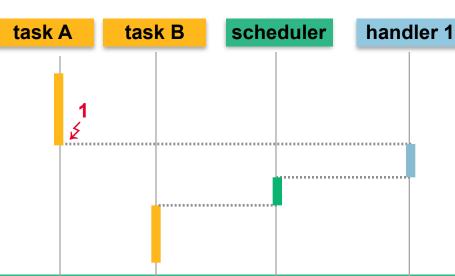


A task runs until:

1 calling a OS
API function
that changes
runnable tasks

task A runs
OS call ····>
OS code runs
context switch ····>

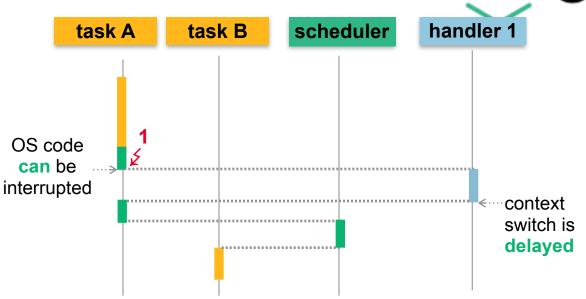
being interrupted by a handler that changes runnable tasks





A task runs until:

3 calling a OS
API function
that gets
interrupted by
a handler that
changes
runnable tasks

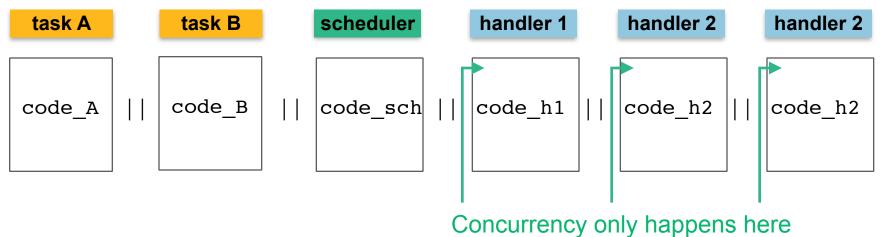


Our approach applies to OS systems with these characteristics:

- tasks can be preempted by higher priority tasks
- OS code can be interrupted by external event

Controlled Interleaving in OG





We do not have arbitrary concurrency between all these programs

Other interleaving is controlled by hardware instructions

context-switch return-from-interrupt interrupt masking

Modelling application tasks



task A

task B

code A

code B





```
WHILE True DO
 await paint A code A | await paint B code B
END
```

WHILE True DO AWAIT AT=A THEN a_1 ; AWAIT AT=A THEN a_2 ;

```
END
```

```
WHILE True DO
END
```

```
WHILE True DO
  AWAIT AT=B THEN b_1;
  AWAIT AT=B THEN b_2;
```

Interleaving can only happen if one instruction is a call to an OS function calling the scheduler calling context switch

or if an interrupt happens

OD

Modelling application tasks

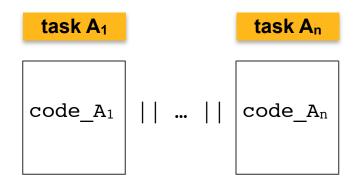


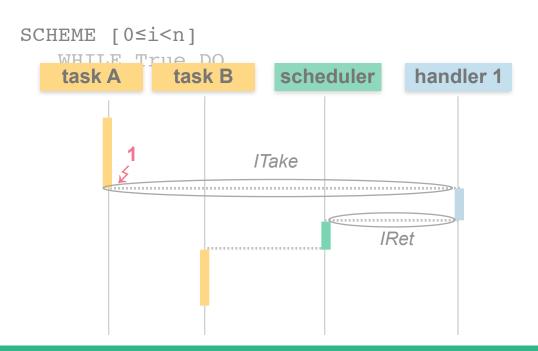
$\begin{array}{|c|c|c|c|c|c|}\hline \textbf{task } \textbf{A}_1 & & \textbf{task } \textbf{A}_n \\ \hline \\ \textbf{code}_\textbf{A}_1 & || & ... & || & \textbf{code}_\textbf{A}_n \\ \hline \end{array}$

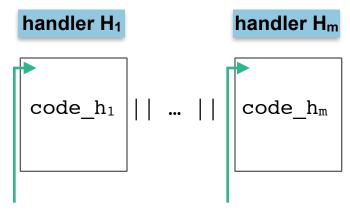
→ We can generalise to n tasks

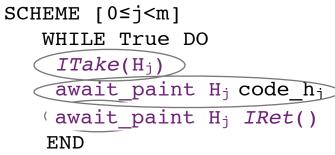
```
SCHEME [0≤i<n]
WHILE True DO
await_paint A<sub>i</sub> code_A<sub>i</sub>
END
```







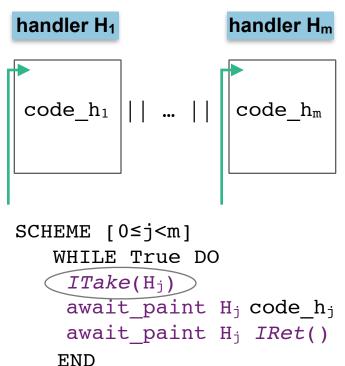






task A₁ code_A₁ || ... || code_A_n

$$ITake(X) \equiv AT := X$$





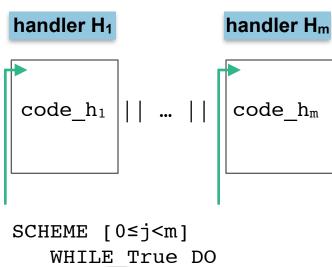
$$ITake(X) \equiv \underset{\overset{\bullet}{\text{AWAIT}}}{\text{AMAIT}} X \in EI^{****}$$

$$THEN$$

$$AT := X$$

- What if X is masked?
- New variable El (Enabled Interrupts)
- New hardware functions

```
IntDisable(X) \equiv EI := EI - X
IntEnable(X) \equiv EI := EI \cup X
```



```
WHILE True DO

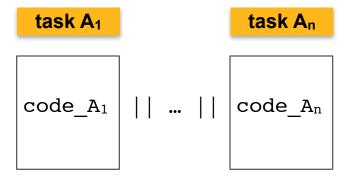
ITake(H<sub>j</sub>)

await_paint H<sub>j</sub> code_h<sub>j</sub>

await_paint H<sub>j</sub> IRet()

END
```





```
ITake(X) ≡ AWAIT X ∈ EI
THEN

**.push AT ATstack;
AT:= X
```

```
handler H<sub>m</sub>

code_h<sub>1</sub> || ... || code_h<sub>m</sub>
```

```
SCHEME [0≤j<m]

WHILE True DO

ITake(H<sub>j</sub>)

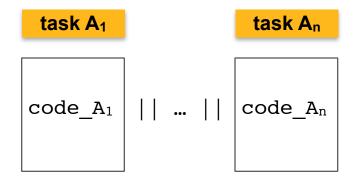
await_paint H<sub>j</sub> code_h<sub>j</sub>

await_paint H<sub>j</sub> IRet()

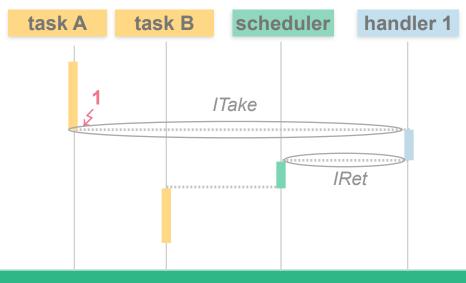
END
```

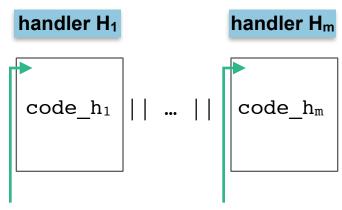
- → How to return to previsouly running task eventually?
- New variable ATstack





$$IRet() \equiv AT := sched;$$

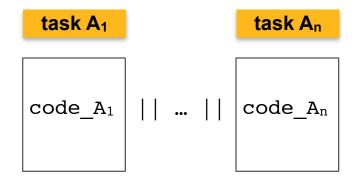


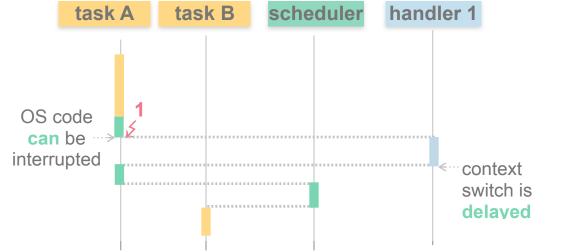


```
SCHEME [0≤j<m]
WHILE True DO

ITake(H<sub>j</sub>)
await_paint H<sub>j</sub> code_h<sub>j</sub>
await_paint H<sub>j</sub> IRet()
END
```







```
handler H<sub>m</sub>

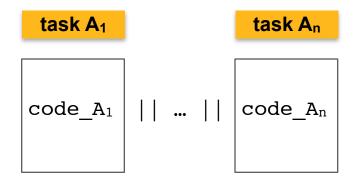
code_h<sub>1</sub> || ... || code_h<sub>m</sub>
```

```
SCHEME [0≤j<m]
WHILE True DO

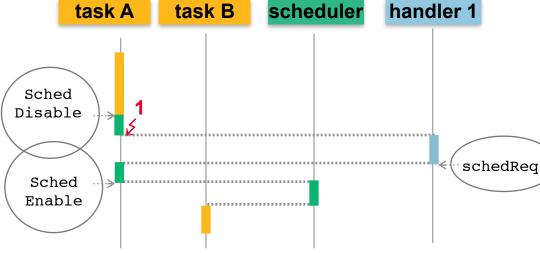
ITake(H<sub>j</sub>)
await_paint H<sub>j</sub> code_h<sub>j</sub>
await_paint H<sub>j</sub> IRet()
END
```

- New hardware functions
 SchedDisable() = EI:=EI-sched
 SchedEnable() = EI:=EI∪sched
- → New flag schedReq



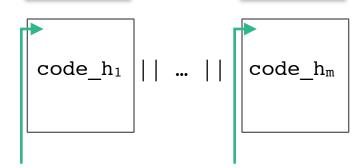


$$IRet() \equiv \widehat{AT} := sched;$$



handler H₁

handler H_m

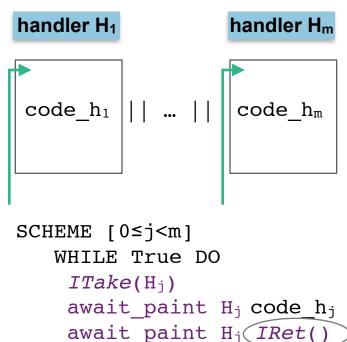


```
SCHEME [0≤j<m]
WHILE True DO

ITake(H<sub>j</sub>)
await_paint H<sub>j</sub> code_h<sub>j</sub>
await_paint H<sub>j</sub> IRet()
END
```

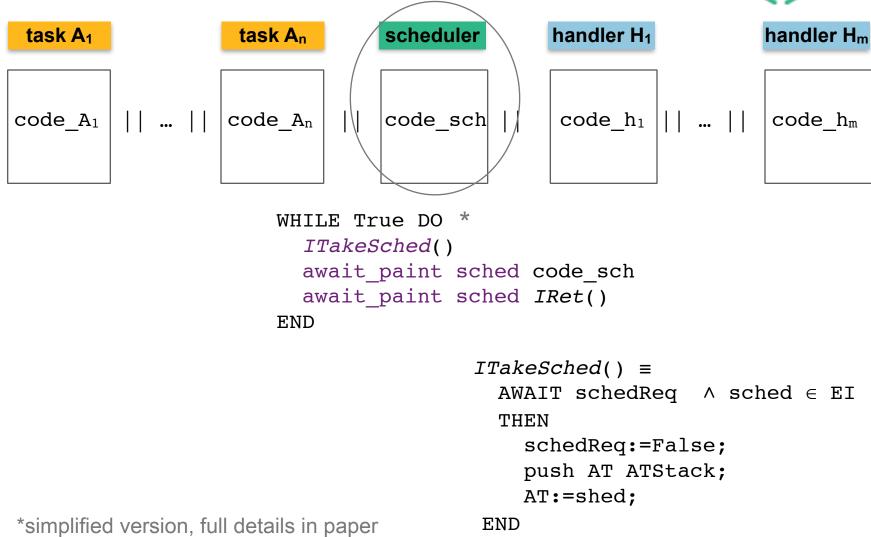
- New hardware functions
 SchedDisable() ≡ EI:=EI-sched
- SchedEnable() = EI:=EIUsched
- → New flag schedReq





END





Our model of interleaving and HW API



```
task A<sub>1</sub>
```

SCHEME [0≤i<n]

END

WHILE True DO

await-paint Ai code-Ai

task An

scheduler

```
WHILE True DO
  ITakeSched()
 await paint sched code sch
  await paint sched IRet()
END
```

handler H₁ ··· handler H_m

```
SCHEME [0≤j<m]
    WHILE True DO
     ITake(H<sub>1</sub>)
     await paint Hi code hi
     await paint H<sub>j</sub> IRet()
    END
```

```
Variables AT, ATStack, EI, schedReq
               IntEnable(X)
               IntDisable(X)
              SchedEnable()
              SchedDisable()
                 ITake(X)
                  IRet()
               ITakeSched()
```



Tasks as 1st class citizens in O-G "AWAIT-painting"

Controlled interleaving model

Hardware API model

Formal model of eChronos to prove that the running task is always the highest-priority runnable task

Instantiation to eChronos



```
handler H<sub>1</sub> ... handler H<sub>m</sub>
    task A<sub>1</sub>
                                        scheduler
                   task An
                               WHILE True DO
                                                                SCHEME [0≤j<m]
SCHEME [0≤i<n]
                                  ITakeSched()
                                                                   WHILE True DO
  WHILE True DO
                                 await-paint sched code-shed
                                                                     ITake(H<sub>i</sub>)
     await-paint Ai code-Ai
                                 await-paint sched IRet()
                                                                    await-paint Hi code hi
   END
                               END
                                                                     await-paint H; IRet()
                                                                    END
                                             code-sch ≡
code A<sub>i</sub> ≡
                                                  nextT:=None;
  SchedDisable();
                                                  WHILE nextT=None
  R := changeRunnable(R);
                                                     DO
   SchedEnable();
                                                        Etmp := E;
                                                        R := handleEvents Etmp R;
                                                        E := E - Etmp;
                                                        nextT:= schedPolicy(R);
  code h_i \equiv
                                                      OD;
    E := changeEvents();
                                                    context switch(nextT);
     schedReq:=True;
```

Summary



1

Tasks as 1st class citizens in O-G "AWAIT-painting"

await_paint task_id code

2

Controlled interleaving model

Hardware API model

AT, ATStack, EI, schedReq

IntEnable(X)

IntDisable(X)

SchedEnable()

SchedDisable()

ITake(X)

IRet()

ITakeSched()



3

Formal model of eChronos

to prove that the running task is always the highest-priority runnable task

 $code_A_i \equiv ...$

 $code_h_i \equiv ...$

code_sch ≡ ...



Future work



Tasks as 1st class citizens in O-G "AWAIT-painting"

Controlled interleaving model

Hardware API model

Formal model of eChronos

to prove that the running task is always the highest-priority runnable task



Prove the model correctly abstracts eChronos code

Prove this property holds on our eChronos model

Informal arguments of validity



AWAIT statements do not exist in the implementation!

→ only use them to represent atomicity enforced by hardware

Introduced variables do no exist in the implementation!

→ only modified by hardware API functions

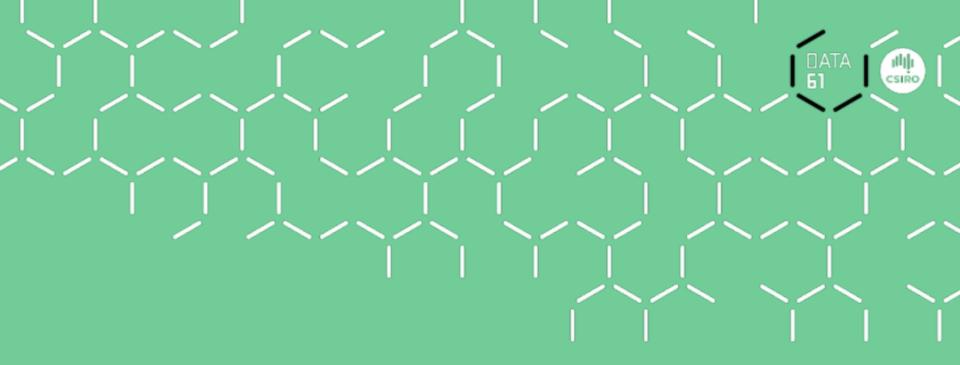
Preliminary arguments of practicality



We have started proving the correctness of the scheduling behaviour

- → proof done* for an initial version of the model
- → ~10,000 verification conditions generated
- → down to ~500 by removing redundant conditions automatically
- → down to ~10 after automatic discharge by Isabelle/HOL

*almost ;-)



Thank you