

A BSP algorithm for the state space construction of security protocols

Frédéric Gava, Michael Guedj, Franck Pommereau

Laboratory of **A**lgorithms, **C**omplexity and **L**ogic (LACL)
University of Paris-East

Security protocols

- apparent simplicity but notoriously error-prone
- need for formal proofs
 - multiple sessions
 - in presence of an attacker
 - discover attacks (man-in-the-middle, replay, ...)
- explicit model-checking
 - automated
 - exhaustive analysis
 - on finite scenarios
 - quick state explosion
- need for faster computation, with more memory

Protocol example: Otway-Rees

Goal: distribute a fresh shared symmetric key K_{ab} between agents A and B through trusted server S

- 1 $A \rightarrow B$ $M, A, B, \{N_a, M, A, B\}K_{as}$
- 2 $B \rightarrow S$ $M, A, B, \{N_a, M, A, B\}K_{as}, \{N_b, M, A, B\}K_{bs}$
- 3 $S \rightarrow B$ $M, \{N_a, K_{ab}\}K_{as}, \{N_b, K_{ab}\}K_{bs}$
- 4 $B \rightarrow A$ $M, \{N_a, K_{ab}\}K_{as}$

Security requirement: secrecy of K_{ab}

Dolev-Yao attacker

- 1 agents send messages to the network
- 2 spy captures messages
 - learns by recursive decomposition/decryption
 - forges new messages from learnt information
 - using only allowed operations (perfect cryptography)
- 3 spy delivers messages (including the original one)

Otway-Rees: known replay attack



Exploits a typing error: A accepts $\{M, A, B\}$ as fresh key K_{ab}

- 1 $A \rightarrow I(B)$ $M, A, B, \{N_a, M, A, B\}K_{as}$
- 2 $B \rightarrow S$ $M, A, B, \{N_a, M, A, B\}K_{as}, \{N_b, M, A, B\}K_{bs}$
- 3 $S \rightarrow B$ $M, \{N_a, K_{ab}\}K_{as}, \{N_b, K_{ab}\}K_{bs}$
- 4 $I(B) \rightarrow A$ $M, \{N_a, M, A, B\}K_{as}$


Tools: SNAKES and BSP-Python

- coloured Petri nets for protocol and scenarios models
- SNAKES (LACL)
 - coloured Petri nets library
 - colour domain = Python language
 - ABCD algebra for protocols
 - send/receive sequences in parallel
- BSP-Python (CNRS Physic Orleans)
 - BSP library for Python
 - used for global exchanges
- LACL cluster \Rightarrow 40 CPU BSP machine
- short cycles: algorithm \mapsto implementation \mapsto benchmarks
- not efficient but accurate for algorithms comparison


High Level Petri Nets



$$\begin{aligned}
 S &\stackrel{\text{def}}{=} \{s_1, s_2\} \\
 T &\stackrel{\text{def}}{=} \{t\} \\
 \ell &\stackrel{\text{def}}{=} \{s_1 \mapsto \mathbb{N}, s_2 \mapsto \{\bullet\}, t \mapsto x > 0, (s_1, t) \mapsto \{x\}, \\
 &\quad (s_2, t) \mapsto \emptyset, (t, s_1) \mapsto \{x-1\}, (t, s_2) \mapsto \{\bullet\}\}
 \end{aligned}$$




Algebra of High Level Petri Nets (1)



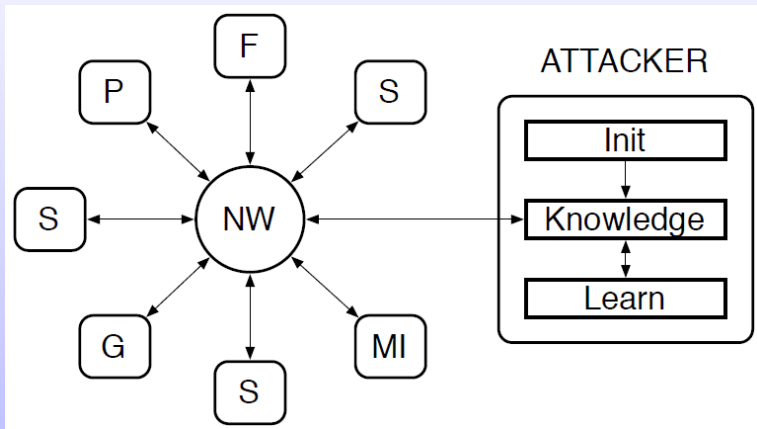
Algebra of High Level Petri Nets (2)

[count-(x), count+(x+1), shift?(j), buf+(j+x) if $x < 10$]



- communication inter-agent using a buffer (the network)
- each agent is a sequential process using its own buffers and the network

Model of Attacker



Learn is a Python code for dolev-Yao inductives rules !

Python ? Are you serious ?

- Advantages of Python:
 - untyped and interpreted
 - many efficient data-structures
 - SNAKE In Python
- Drawbacks of Python:
 - untyped and interpreted
 - non-determinism on list/set iterations
 - hashing a value is partially machine dependant

Labelled transitions systems

- initial state s_0
- successors given by $\text{succ}(s)$
- transition $s \rightarrow s'$ whenever $s' \in \text{succ}(s)$
- inductive computation of the state space
 - as a graph (explicit LTS)
 - as a set of reachable states
- we present sets for simplicity


Sequential algorithm

```
1: todo  $\leftarrow \{s_0\}$ 
2: known  $\leftarrow \emptyset$ 
3: while todo  $\neq \emptyset$  do
4:   pick s from todo
5:   known  $\leftarrow \textit{known} \cup \{s\}$ 
6:   for  $s' \in \textit{succ}(s)$  do
7:     if  $s' \notin \textit{known} \cup \textit{todo}$  then
8:       todo  $\leftarrow \textit{todo} \cup \{s'\}$ 
9:     end if
10:  end for
11: end while
```

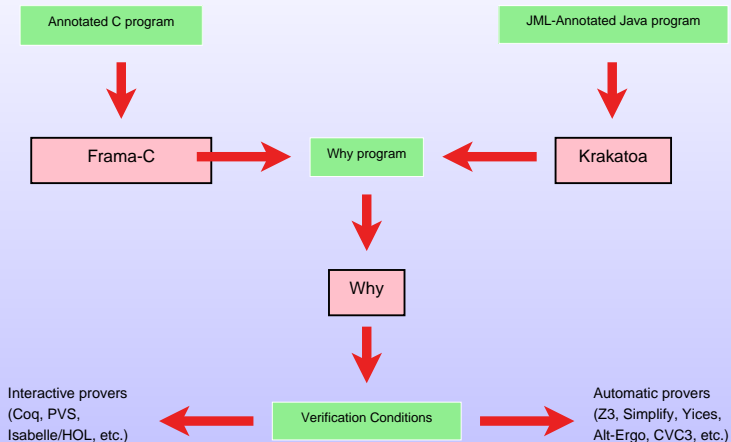
Sequential algorithm (2)

```
1:  $todo \leftarrow \{s_0\}$ 
2:  $known \leftarrow \emptyset$ 
3: while  $todo \neq \emptyset$  do
4:   pick  $s$  from  $todo$ 
5:    $known \leftarrow known \cup \{s\}$ 
6:    $todo \leftarrow (todo \cup succ(s)) \setminus known$ 
7: end while
```

Do you truth me ? Is there a bug ?



The Why tools



- Annotated programs (small algorithmic/logic language)
- generated proof obligations for theorem provers
- need axiomatisation for set/list etc.

BSP model




Figure: A BSP super-step

Building a BSP algorithm

Four steps:

- 1 start from a naive algorithm
- 2 improve the locality of computation
- 3 improve the memory consumption
- 4 balance computation

Naive BSP algorithm

- partition function cpu to place states onto processors
 - hash the state (modulo number of processors)
 - most used approach
- each processor i computes $succ(s)$ iff $cpu(s) = i$
- other states are sent to their owners
- stop when no processor has computed new states

Main loop

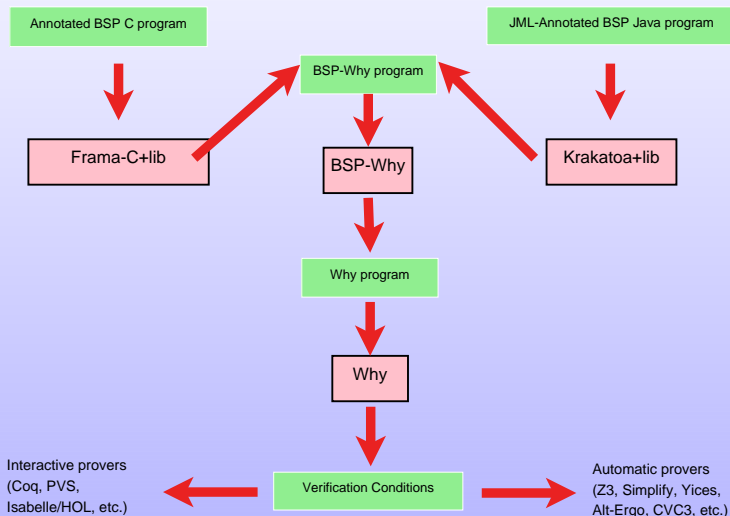
```
1: todo  $\leftarrow \emptyset$ 
2: total  $\leftarrow 1$ 
3: known  $\leftarrow \emptyset$ 
4: if cpu(s0) = myid then
5:   todo  $\leftarrow todo \cup \{s_0\}$ 
6: end if
7: while total > 0 do
8:   tosend  $\leftarrow Successor(known, todo)$ 
9:   todo, total  $\leftarrow BSP\_Exchange(known, tosend)$ 
10: end while
```

Successor(*known*, *todo*)

```
1: tosend  $\leftarrow \emptyset$ 
2: while todo  $\neq \emptyset$  do
3:   pick s from todo
4:   known  $\leftarrow$  known  $\cup$  {s}
5:   for  $s' \in \text{succ}(s)$  do
6:     if  $s' \notin \text{known} \cup \text{todo}$  then
7:       if  $\text{cpu}(s') = \text{mypid}$  then
8:         todo  $\leftarrow$  todo  $\cup$  {s'}
9:       else
10:        tosend  $\leftarrow$  tosend  $\cup$  {( $\text{cpu}(s')$ , s')}
11:      end if
12:    end if
13:  end for
14: end while
15: return tosend
```

Trying to prove its correctness

BSP-WHY, an extension of WHY for BSP algorithms:



Problems with naive algorithm

- a new state is very likely to be sent
 - **no locality** of the computations
 - **too much communication**
- protocols yield **structured** models
- we can exploit their **properties** for more efficient algorithms

Assumptions about protocol models

- 1 finite and fixed number of agents
- 2 states are functions $\mathcal{L} \rightarrow \mathcal{D}$ (locations \rightarrow data)
- 3 $\mathcal{L}_R \subseteq \mathcal{L}$ defines *reception locations*
- 4 $succ = succ_R \uplus succ_L$ such that $succ_L$ is the identity on \mathcal{L}_R
- 5 cpu_R such that $s_1|_{\mathcal{L}_R} = s_2|_{\mathcal{L}_R} \Rightarrow cpu_R(s_1) = cpu_R(s_2)$
 - use a hash on \mathcal{L}_R

\Rightarrow cross-transitions correspond to receptions in the protocol

General assumptions, easy to implement

Improve local computation

Successor(*known*, *todo*) :

- 1: *tosend* $\leftarrow \emptyset$
- 2: **while** *todo* $\neq \emptyset$ **do**
- 3: **pick** *s* **from** *todo*
- 4: *known* \leftarrow *known* \cup {*s*}
- 5: **for** $s' \in \text{succ}_L(s) \setminus \text{known}$ **do**
- 6: *todo* \leftarrow *todo* \cup {*s'*}
- 7: **end for**
- 8: **for** $s' \in \text{succ}_R(s) \setminus \text{known}$ **do**
- 9: *tosend* \leftarrow *tosend* \cup {(*cpu*_R(*s'*), *s'*)}
- 10: **end for**
- 11: **end while**
- 12: **return** *tosend*

```

while todo  $\neq \emptyset$  do
  pick s from todo
  known  $\leftarrow$  known  $\cup$  {s}
  for  $s' \in \text{succ}(s)$  do
    if  $s' \notin \text{known} \cup \text{todo}$  then
      if cpu(s') = myid then
        todo  $\leftarrow$  todo  $\cup$  {s'}
      else
        tosend  $\leftarrow$  tosend  $\cup$  {(cpu(s'), s')}
      end if
    end if
  end for
end while

```

Improve memory consumption

- super-steps match protocols progression (receptions)
- known states cannot be reached in a future super-step
- at each super-step, we can **dump** memory

```
1: todo  $\leftarrow \emptyset$ 
2: total  $\leftarrow 1$ 
3: known  $\leftarrow \emptyset$ 
4: if  $\text{cpu}(s_0) = \text{mypid}$  then
5:   todo  $\leftarrow \text{todo} \cup \{s_0\}$ 
6: end if
7: while  $\text{total} > 0$  do
8:   tosend  $\leftarrow \text{Successor}(\text{known}, \text{todo})$ 
9:   dump(known)
10:  todo, total  $\leftarrow \text{BSP\_Exchange}(\text{known}, \text{tosend})$ 
11: end while
```

Balancing the computation

- 1 efficient parallel computation \iff good workload balance
- 2 number of computed states cannot be anticipated
- 3 but, related to the number of received states
- 4 introduce a rebalancing step

7: **while** $total > 0$ **do**

8: $tosend \leftarrow Successor(known, todo)$

9: **dump**($known$)

10: $todo, total \leftarrow BSP_Exchange(Balance(tosend))$





11: **end while**

- define cpu_B as cpu_R for infinitely many processors
- exchange local histograms of $cpu_B \Rightarrow$ global histogram H
- balance $tosend$ wrt H (approximate bin packing)





Methodology

- run two scenarios (small and large) for:
 - 1 Needham-Schroeder mutual authentication
 - 2 Yahalom key sharing and mutual authentication with TTP
 - 3 Otway-Rees key sharing with TTP
 - 4 Kao-Chow key sharing and mutual authentication
- measure on each processor
 - 1 local computation time (yellow)
 - 2 synchronisation time, *i.e.*, waiting time (red)
 - 3 communication time (blue)

Needham-Schroeder (top) and Yahalom (bottom)



Otway-Rees (top) and Kao-Chow (bottom)



NS and Yahalom

Scenario	Naive	Balance	Nb_states
NS_1-2	0m50.222s	0m42.095s	7807
NS_1-3	115m46.867s	61m49.369s	530713
NS_2-2	112m10.206s	60m30.954s	456135







Scenario	Naive	Balance	Nb_states
Y_1-3-1	12m44.915s	7m30.977s	399758
Y_1-3-1_2	30m56.180s	14m41.756s	628670
Y_1-3-1_3	481m41.811s	25m54.742s	931598
Y_2-2-1	2m34.602s	2m25.777s	99276
Y_3-2-1	COMM	62m56.410s	382695
Y_2-2-2	2m1.774s	1m47.305s	67937

Otway-Rees and Woo and Lam Pi

Scenario	Naive	Balance	Nb_states
OR_1-1-2	38m32.556s	24m46.386s	12785
OR_1-1-2_2	196m31.329s	119m52.000s	17957
OR_1-1-2_3	411m49.876s	264m54.832s	22218
OR_1-2-1	21m43.700s	9m37.641s	1479

Scenario	Naive	Balance	Nb_states
WLP_1-1-1	0m12.422s	0m9.220s	4063
WLP_1-1-1_2	1m15.913s	1m1.850s	84654
WLP_1-1-1_3	COMM	24m7.302s	785446
WLP_1-2-1	2m38.285s	1m48.463s	95287
WLP_1-2-1_2	SWAP	55m1.360s	946983

Memory consumption for NS, WLP and Y



Conclusion

- BSP algorithms for state space generation
- exploit structural properties of security protocols to
 - reduce communications
 - anticipate the number of super-steps
 - decrease local storage
 - balance the workload
- properties easily computable on Petri net models
- our algorithm is
 - efficient
 - **scalable**
 - simple
 - **provable**
- benchmarks on realistic and non-trivial examples

Perspectives

- ongoing work
 - new benchmarks of other protocols and scenarios
 - correctness proofs using BSP Hoare logic (BSP-WHY)
- future work
 - more complex protocol from SPREAD project
 - LTL/CTL* model checking, exploiting state space locality
 - certified implementation ?