UNIVERSITY OF TWENTE. formal methods & tools.



Problem Solving with Model Checking Techniques



Jaco van de Pol and Michael Weber

June 12, 2011



ICAPS Tutorial, Freiburg, Germany

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Table of Co	ontents				



- 1 Model Checking in a Nutshell
- 2 Planning Example: Sokoban
- 3 LTSmin Tool Architecture
- 4 Symbolic Algorithms



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Table of C	Contents				



1 Model Checking in a Nutshell

- 2 Planning Example: Sokoban
- **3** LTSmin Tool Architecture
- **4** Symbolic Algorithms



5 Multi-Core Algorithms

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Model Che	cking in	a Nutshell			

- \blacktriangleright Specification of system: logical formula φ
- Implementation of system: Kripke structure
- Question: Does the system meet its specification?



- ► Applications: hardware, software, wetware
- ► Method: (Variations of) Graph Reachability

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
What is N	/lodel Che	ecking?			

Model Checking

- Check if a given model satisfies a given property
- Promise: Automatic answer to combinatorial questions

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
What is N	Model Che	ecking?			

Model Checking

- Check if a given model satisfies a given property
- Promise: Automatic answer to combinatorial questions

Models: discrete dynamics

- software / hardware / embedded systems
- communicating concurrent components
- biological systems, intra/inter cell level

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
What is I	Model Che	ecking?			

Model Checking

- Check if a given model satisfies a given property
- Promise: Automatic answer to combinatorial questions

Models: discrete dynamics

- software / hardware / embedded systems
- communicating concurrent components
- biological systems, intra/inter cell level

Properties: of transition graphs

- invariants, assertions, absence of errors
- absence / presence of event orderings
- complicated fairness restrictions possible

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Application	IS				

Impressive ApplicationsTuring Award 2007

- Numerous case studies have been published:
 - communication and security protocols
 - embedded controllers, e.g. elevators, railways, cars
 - concurrent and distributed algorithms
 - Biology: signaling pathways, gene regulation, differentiation

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Application	S				

Impressive ApplicationsTuring Award 2007

- Numerous case studies have been published:
 - communication and security protocols
 - embedded controllers, e.g. elevators, railways, cars
 - concurrent and distributed algorithms
 - ► Biology: signaling pathways, gene regulation, differentiation
- ► Leading industries rely on model checking for quality control:
 - Intel/IBM's processors go through extensive model checking (they report that it replaced a considerable amount of testing)
 - Microsoft's Static Device Verifier is part of the WDK (3rd party device drivers are checked for interface compliance)

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Application	is				

Impressive ApplicationsTuring Award 2007

- Numerous case studies have been published:
 - communication and security protocols
 - embedded controllers, e.g. elevators, railways, cars
 - concurrent and distributed algorithms
 - ► Biology: signaling pathways, gene regulation, differentiation
- Leading industries rely on model checking for quality control:
 - Intel/IBM's processors go through extensive model checking (they report that it replaced a considerable amount of testing)
 - Microsoft's Static Device Verifier is part of the WDK (3rd party device drivers are checked for interface compliance)

But still, model checking is...

- ▶ Not built into CASE tools, despite its "push-button" nature
- Not available to the average (SME-type) software engineer

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Table of (Contents				



1 Model Checking in a Nutshell

- 2 Planning Example: Sokoban
- **3** LTSmin Tool Architecture
- 4 Symbolic Algorithms



5 Multi-Core Algorithms

Model Checking Sokoban PINS/LTSmin Symbolic Multi-Core Conclusion

Sokoban as you know it



UNIVERSITY OF TWENTE.

Problem Solving with Model Checking Techniques

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Modelling
 Sokoban
 as a Transition
 System
 System
 Symbolic
 Symbolic
 Multi-Core
 Conclusion

States and Transitions

- ► States:
 - ► View each location on the board as a variable x_{i,j}
 - ▶ Possible values: $x_{i,j} \in \{wall, man, block, empty\}$

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Modelling
 Sokoban
 as a Transition
 System
 System
 Sokoban
 Sok

States and Transitions

- ► States:
 - ► View each location on the board as a variable x_{i,i}
 - ▶ Possible values: $x_{i,j} \in \{wall, man, block, empty\}$

Transitions – distinguish moves and pushes in four directions

Move right:

if $x_{i,j} = man$ and $x_{i,j+1} = empty$

then set $x_{i,j} := empty$ and $x_{i,j+1} := man$.

Push down:

if $x_{i,j} = man$ and $x_{i+1,j} = block$ and $x_{i+2,j} = empty$ then set: $x_{i,j} := empty$, $x_{i+1,j} := man$ and $x_{i+2,j} := block$.
 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Modelling
 Sokoban
 as a Transition
 System
 <td

States and Transitions

- ► States:
 - ► View each location on the board as a variable x_{i,i}
 - ▶ Possible values: $x_{i,j} \in \{wall, man, block, empty\}$

Transitions – distinguish moves and pushes in four directions

- Move right:
 - if $x_{i,j} = man$ and $x_{i,j+1} = empty$
 - then set $x_{i,j} := empty$ and $x_{i,j+1} := man$.
- Push down:

if $x_{i,j} = man$ and $x_{i+1,j} = block$ and $x_{i+2,j} = empty$ then set: $x_{i,j} := empty$, $x_{i+1,j} := man$ and $x_{i+2,j} := block$.

- ▶ Initial state: an assignment to all the x_{i,j}
- ▶ Goal state: If for all goal positions (i,j): x_{i,j} = block then do special action finish

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Next: Solve	Sokobar	with Brute	e Force		

Reachability

- This is just the reachability problem
- Property: "Finish action is not reachable"
- Counter-example: trace to a finish-action

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Next:
 Solve
 Sokoban with
 Brute
 Force
 Force

Reachability

- This is just the reachability problem
- Property: "Finish action is not reachable"
- Counter-example: trace to a finish-action

Brute Force Exploration

- On-the-fly:
 - start with the initial state
 - expand newly encountered states
 - stop when the goal is reached
- Breadth-first strategy guarantees the shortest solution
- Limitations:
 - Only feasible for about 10⁹ states
 - General: no utilization of specific structure of Sokoban

Μ	lodel Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
S	Small Pr	oblem: E×	ponential G	Growth		
	screen	states				
	107	10,165				
	1001	127,509				
	207	1 025 014				

387	1,235,214
372	10,992,856
792	117,434,655
747	1,307,942,326
38	12,197,960,188
754	308,479,382,084
2	4,748,854,893,784

Mode	l Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Sm	nall Pr	oblem: Exp	oonential G	rowth		
	screen	states				
	107	10,165				
	1001	127,509				
	387	1,235,214				
	372	10,992,856				
	792	117,434,655				
	747	1,307,942,326				
	38	12,197,960,188	3			
	754	308,479,382,08	34			
	2	4,748,854,893,	784			_
	376	1,066,317,080,	868,510,000			
	45	1,195,159,588,	147,710,000,000			
	71	1,238,482,287,	687,790,000,000	,000		
	551	1,654,592,279,	840,460,000,000	,000,000		
	48	40,215,766,407	,984,300,000,00	0,000,000,000		
	22	3,531,895,015,	833,180,000,000	,000,000,000,0	00	
	10	12,864,741,234	,813,200,000,00	0,000,000,000,	000,000	
	558	218,612,674,52	7,952,000,000,0	00,000,000,000	0,000,000	
	778	9,341,745,200,	574,070,000,000	,000,000,000,0	00,000,000,000	

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 A Glimpse of Hope:
 Transition Locality
 Image: Conclusion
 Conclusion

Locality

- Transitions may depend on a part of the state vector only
- In Sokoban, every transition depends on 2 or 3 variables; independent of the size of the Sokoban screen
- In general, if you learn one transition and (statically) know the dependency matrix, then you can infer many more transitions

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 A Glimpse of Hope:
 Transition Locality
 Image: Conclusion
 Conclusion

Locality

- Transitions may depend on a part of the state vector only
- In Sokoban, every transition depends on 2 or 3 variables; independent of the size of the Sokoban screen
- In general, if you learn one transition and (statically) know the dependency matrix, then you can infer many more transitions

Example matrix for a Sokoban instance (fragment)

+++----

+----+---

- -++----
- -+---+-
- --+---+
- ---+++---

- ► The dependency matrix is sparse (good)
- ► The +'s are often close together
- ► Some +'s are far apart (2D board)
- Heuristic regrouping can help a lot (cf. BDD variable reordering)

Model Checking Sokoban PINS/LTSmin Symbolic Multi-Core Conclusion
Exploiting Locality

By locality, successor states are "much alike":

Locality helps Implementation

- Compression schemes for storing sets of states
- ► Incremental hashing / storage / compression / etc.
- Communicate diffs only, save bandwidth on clusters

Model Checking Sokoban PINS/LTSmin Symbolic Multi-Core Conclusion
Exploiting Locality

By locality, successor states are "much alike":

Locality helps Implementation

- Compression schemes for storing sets of states
- Incremental hashing / storage / compression / etc.
- Communicate diffs only, save bandwidth on clusters

Locality helps Algorithms (orders of magnitude!)

- Cache intermediate results to save computations
- Store sets of states in Binary Decision Diagrams
- Apply transitions to sets of states symbolically

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Model Checking
 Techniques I

Mentioned later

- Multi-core and Grid implementations
- State space reduction (e.g. partial-order reduction)
- Symbolic Model Checking using BDDs
- Adapting the search order of reachability
- Effects of changing the search order:
 - the peak memory of intermediate BDDs is reduced
 - however, a shortest solution is not guaranteed

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Model Che	cking T	echniques II			

- Bounded model checking, based on SAT solving
 - Runs of increasing, fixed length, are encoded into one big satisfiability problem
 - Huge potential, but Sokoban runs are very long

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Model Che	cking T	echniques II			

- Bounded model checking, based on SAT solving
 - Runs of increasing, fixed length, are encoded into one big satisfiability problem
 - Huge potential, but Sokoban runs are very long
- Directed Model Checking
 - Gives priority to transitions towards a "promising" direction
 - Here maybe: number of blocks already in correct position

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Model Che	ecking T	echniques II			

- Bounded model checking, based on SAT solving
 - Runs of increasing, fixed length, are encoded into one big satisfiability problem
 - Huge potential, but Sokoban runs are very long
- Directed Model Checking
 - Gives priority to transitions towards a "promising" direction
 - Here maybe: number of blocks already in correct position
- Ad Hoc techniques for (classes of) puzzles
 - Avoid "deadlock" situations
 - Use high-level planning, e.g. recognize rooms

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Model Che	ecking T	echniques II			

- Bounded model checking, based on SAT solving
 - Runs of increasing, fixed length, are encoded into one big satisfiability problem
 - Huge potential, but Sokoban runs are very long
- Directed Model Checking
 - Gives priority to transitions towards a "promising" direction
 - ► Here maybe: number of blocks already in correct position
- Ad Hoc techniques for (classes of) puzzles
 - Avoid "deadlock" situations
 - Use high-level planning, e.g. recognize rooms

Here is where model checking might learn from planning!

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Beyond R	eachabilit	с у			

More complicated planning problems

- Avoid unwanted situations (still a reachability problem)
 - ▶ in LTL logic: **F** Good becomes (¬Bad) **Until** Good
 - alternatively: restrict the transition relation
- Associate costs with transitions (e.g.: minimal pushes)
 - this requires an adapted search strategy

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Beyond R	eachabilit	y			

More complicated planning problems

- Avoid unwanted situations (still a reachability problem)
 - ► in LTL logic: **F** Good becomes (¬Bad) **Until** Good
 - alternatively: restrict the transition relation
- Associate costs with transitions (e.g.: minimal pushes)
 - this requires an adapted search strategy
- Take into account real time
 - UPPAAL: real-time model checker; used for HRT scheduling
- Compute (optimal) cyclic schedules (mean pay-off games)
 - ► Liveness in logic: **GF** Good, reachability infinitely often

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Beyond Re	eachability	/			

More complicated planning problems

- Avoid unwanted situations (still a reachability problem)
 - ► in LTL logic: **F** Good becomes (¬Bad) **Until** Good
 - alternatively: restrict the transition relation
- Associate costs with transitions (e.g.: minimal pushes)
 - this requires an adapted search strategy
- Take into account real time
 - UPPAAL: real-time model checker; used for HRT scheduling
- Compute (optimal) cyclic schedules (mean pay-off games)
 - ► Liveness in logic: **GF** Good, reachability infinitely often
- Plan in the presence of uncertainty (two-player games)
 - \blacktriangleright and/or graphs with loop restrictions: parity games, $\mu\text{-calculus}$
 - can even take into account stochastic environment (MDP)

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Table of (Contents				



- 1 Model Checking in a Nutshell
- 2 Planning Example: Sokoban
- **3** LTSmin Tool Architecture
- 4 Symbolic Algorithms



5 Multi-Core Algorithms

LTSmin Tool Architecture



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Great Abst	tractions	are Cheap!			

- Automatic model translations are not good enough
- Separate languages and tools via a clean interface (API)
- API should be simple: allow many different languages
- API should be rich: expose structure, enable algorithms



- Automatic model translations are not good enough
- Separate languages and tools via a clean interface (API)
- API should be simple: allow many different languages
- ► API should be rich: expose structure, enable algorithms



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
PINS in a	Nutshell				

Interface based on a Partitioned Next-State function

- ► State: Fixed-size vector of integers
- Partitioned transition relation


Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
PINS in a	Nutshell				

Interface based on a Partitioned Next-State function

- State: Fixed-size vector of integers
- Partitioned transition relation
- Dependency Matrix $[D]_{N \times K}$:
 - ► Language module guarantees: if D_{i,j} = -, then transition group →_i is independent of state slot j.
 - Matrix: statically known (currently)
 - Language module may over-approximate dependencies

 $\langle 3, 5, 5, 4, 1, 3 \rangle_{\mathrm{IV}}$ $\rightarrow = \bigcup_{i}^{K} \rightarrow_{i}$



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Example [Depender	ncy Matrix			
int x=7; process i	o1() {				

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Example D	Dependenc	y Matrix			
<pre>int x=7;</pre>		int y=3;	o() (int z=9;	
process p do	DI() 1	process p do	2() {	do	3() 1
::{x>0 =>	<pre>x;y++}</pre>	::{y>0 =>	y;x++}	::{z>0 =>	z;x++}
od }	~ ,2115	od }	y ,211)	od }	z, , y + + ;

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Example D	Dependenc	y Matrix			
<pre>int x=7;</pre>		int y=3;		int z=9;	
process p	51() {	process p	p2() {	process p	3() {
do		do		do	
::{x>0 =>	• x;y++}	::{y>0 =>	> y;x++}	::{z>0 =>	z;x++}
::{x>0 =>	• x;z++}	::{y>0 =>	> y;z++}	::{z>0 =>	z;y++}
od }		od }		od }	

Default Matrix

$$\begin{array}{cccc}
x & y & z \\
p1 \\
p2 \\
p3 \\
+ & + & + \\
+ & + & +
\end{array}$$

Model Checking S	okoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Example Dep	endency	y Matrix			
<pre>int x=7; process p1() do ::{x>0 => x- ::{x>0 => x- od }</pre>) { ;y++} ;z++}	<pre>int y=3 process do ::{y>0 = ::{y>0 = od }</pre>	; p2() { => y;x++ => y;z++	<pre>int z=9; process p do -} ::{z>0 =: .::{z>0 =: od }</pre>	p3() { > z;x++} > z;y++}
Default Matrix	Be	tter Mat	rix		
$ \begin{array}{ccc} x & y \\ p1 \\ p2 \\ p3 \\ + & + \\ + & + \\ + & + \\ \end{array} $	z p1. + p1. + p2. p2. p3. p3.	$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	z + + + + + + + + +		

Model Checking Sokoban Example Depend	ency	PINS/LTSmi	in S rix	ymbolic	Multi-Core	Conclusion
<pre>int x=7; process p1() { do ::{x>0 => x;y ::{x>0 => x;z od }</pre>	++}	<pre>int y= proces do ::{y>(::{y>(od }</pre>	=3; ss p2() 0 => y- 0 => y-) { ;x++} ;z++}	<pre>int z= proces do ::{z>0 ::{z>0 od }</pre>	9; s p3() { => z;x++} => z;y++}
	Ве <i>p</i> 1. <i>p</i> 1. <i>p</i> 2. <i>p</i> 2. <i>p</i> 3. <i>p</i> 3.	x 1 + 2 + 1 + 2 - 1 2 -	y z + - - + + - + + + + + +	state = $\langle 7, 3, 9 \\ \langle 7, 3, 2 \\ \langle 7, 3, 2 \\ \langle 7, 3, 9 \\ \langle *, 3, 9 \\ \rangle$	$= \langle 7, 3, 9 \rangle$ $= \langle 7, 3, 9 \rangle$ $\Rightarrow \rangle_{\rm IV} \qquad \frac{p_{\rm 1.}}{p_{\rm 1.}}$ $\Rightarrow \rangle_{\rm IV} \qquad \frac{p_{\rm 3.}}{p_{\rm 2.}}$	IV $ \begin{array}{cccc} & & & \\ & & &$

UNIVERSITY OF TWENTE.

Problem Solving with Model Checking Techniques

June 12, 2011 21 / 71

Model Checkir	ng Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
PINS /	/ LTSmin				



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
PINS / LT	Smin				



Basic functions (algorithms call language modules)

- **GetMatrix** : returns the dependency matrix $[D]_{N \times K}$
- ▶ INITSTATE(): returns the initial state vector
- ▶ NEXTSTATE(I,S): successors of state *s* in transition group *i*

	Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
	Table of Co	ontents				
_						



- 1 Model Checking in a Nutshell
- 2 Planning Example: Sokoban
- **3** LTSmin Tool Architecture
- 4 Symbolic Algorithms



5 Multi-Core Algorithms

UNIVERSITY OF TWENTE.

Local Transition Caching



 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Caching Local Transitions (1)
 Conclusion
 <

► Recall local transition in specification:

p3.2: atomic { $z>0 \rightarrow z--; y++$ }

- x y z
- ► Dependency matrix row: *p*3.2 [0 1 1]
- Define projection: $\pi_{p3.2}\langle x, y, z \rangle = \langle y, z \rangle$

Model Checking Sokoban PINS/LTSmin Symbolic Multi-Core Conclusion Caching Local Transitions (1)

X V Z

- Recall local transition in specification:
 - p3.2: atomic { $z>0 \rightarrow z--; y++$ }
- Dependency matrix row: $p3.2 \begin{bmatrix} 0 & 1 & 1 \end{bmatrix}$
- Define projection: $\pi_{p3,2}\langle x, y, z \rangle = \langle y, z \rangle$
- Next, consider two consecutive calls to p3.2:

first call: $\langle x, y, z \rangle$ successor: $\langle x, y', z' \rangle$ project and $\langle y, z \rangle \rightarrow$ store in cache: $\langle v', z' \rangle$

Model Checking Sokoban PINS/LTSmin Symbolic Multi-Core Conclusion Caching Local Transitions (1)

Recall local transition in specification:

p3.2: atomic { $z>0 \rightarrow z--; y++$ }

- Dependency matrix row: $p3.2 \begin{bmatrix} 0 & 1 & 1 \end{bmatrix}$
- Define projection: $\pi_{p3.2}\langle x, y, z \rangle = \langle y, z \rangle$
- Next, consider two consecutive calls to p3.2:

first call: $\langle x, y, z \rangle$ successor: $\langle x, y', z' \rangle$ project and $\langle y, z \rangle \rightarrow$ store in cache: $\langle y', z' \rangle$

second call: $\langle x'', y, z \rangle$ project: cache lookup: $\rightarrow \langle y', z' \rangle$ expand: $\langle x'', v', z' \rangle$

X V Z



Model Checking Sokoban PINS/LTSmin Symbolic Multi-Core Conclusion Caching Local Transitions (1)

Recall local transition in specification:

p3.2: atomic { $z>0 \rightarrow z--; y++$ }

- Dependency matrix row: $p3.2 \begin{bmatrix} 0 & 1 & 1 \end{bmatrix}$
- Define projection: $\pi_{p3,2}\langle x, y, z \rangle = \langle y, z \rangle$
- Next, consider two consecutive calls to p3.2:

first call: $\langle x, y, z \rangle$ successor: $\langle x, y', z' \rangle$ project and $\langle v, z \rangle \rightarrow$ store in cache: $\langle v', z' \rangle$

second call: $\langle x'', y, z \rangle$ project: cache lookup: $\rightarrow \langle y', z' \rangle$ expand: $\langle x'', v', z' \rangle$

x y z



Maintain a memoization table cache[i] for each transition group i



- Caching can save calls to the language module
- Still some work for every concrete state (cache lookup)
- Caching is useful especially:
 - For expressive/inefficient languages
 - When dependency matrices are sparse
- Always uses a bit more memory (tables)

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Static Ma	trix Reor	dering			

PINS Optimization

- Regrouping similar transition groups reduces overhead
- Reordering state variables reduces BDD sizes (a.o.)

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Static Ma	trix Reor	dering			

PINS Optimization

- Regrouping similar transition groups reduces overhead
- Reordering state variables reduces BDD sizes (a.o.)



Multi-Valued Decision Diagrams



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Multi-Valu	ied Decis	sion Diagram	ns		





Every path in the MDD represents a concrete state vector

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Multi_Valu	ed Decis	ion Diagram	nc		





- Every path in the MDD represents a concrete state vector
- ▶ Potential gain in memory saving: exponential (here: $54 \rightarrow 15$)
- Symbolic Reachability: explore sets of states stored as MDDs

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion	
Symbolic I [ICTAC 2008]	Reachab	ility Algorith	ım			

- ► L, V: MDDs for sets of long state vectors (level, visited)
- R_i : MDDs to store transition relation *i* on short vectors
- ► *L_i*, *V_i*: MDDs for sets of short state vectors (level,visited for *i*)

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Symbolic
 Reachability
 Algorithm
 [ICTAC 2008]
 Image: Conclusion
 Image: Conclusic
 Image: Conclusion
 Image: Co

- ► L, V: MDDs for sets of long state vectors (level, visited)
- R_i : MDDs to store transition relation *i* on short vectors
- ► *L_i*, *V_i*: MDDs for sets of short state vectors (level,visited for *i*)

symbolic-reachability()

(1)	$L := \{ \text{INITSTATE}() \}; V := L; \text{all } R_i := \emptyset; \text{all } V_i := \emptyset$
(2)	while $L \neq \emptyset$ do
(3)	for $i \in groups$ do $/*$ enumerate short vectors $*/$
(4)	$L_i := \pi_i([D]_{N imes K}, L) \setminus V_i; V_i := V_i \cup L_i$
(5)	$R_i := R_i \ \cup \ \{(s,s') \mid s \in L_i \land s' \in \operatorname{NextState}(i,s)\}$
(6)	$L := \bigcup_i (R_i(L) \setminus V); V := V \cup L$
(7)	return V

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Symbolic	Reachabi	ility			

Symbolic Reachability with PINS

- Global set of reachable states is computed as fix point
- Stored as a multi-valued decision diagram (MDD)
- ► Learn symbolic sub-groups *R_i* on-the-fly (via NEXTSTATE)

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Symbolic	Reachabi	ility			

Symbolic Reachability with PINS

- Global set of reachable states is computed as fix point
- Stored as a multi-valued decision diagram (MDD)
- ► Learn symbolic sub-groups *R_i* on-the-fly (via NEXTSTATE)

Extensions

- Multiple exploration strategies:
 - Breadth-first: $(T_1 + T_2 + \cdots + T_n)^*$
 - Chaining: $(T_1 \circ T_2 \circ \cdots \circ T_n)^*$
 - Saturation-like: $(T_1^* \circ T_2^* \circ \cdots \circ T_n^*)^*$
 - Full Saturation: $((((T_1^*T_2)^*T_3)^*\cdots)^*T_n)^*$

Static variable reordering boosts performance

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion	
Table of C	Contents					



- 1 Model Checking in a Nutshell
- 2 Planning Example: Sokoban
- 3 LTSmin Tool Architecture
- **4** Symbolic Algorithms



5 Multi-Core Algorithms

Multi-Core Algorithms



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Multi-Core	Reality				

Recent standard hardware (x86)

- Multiple cores per processor, multiple processors
- ► Typical big machine: 48–64 cores, 256 GB shared memory
- ► Communication via L1, L2 caches: cache-coherence protocols

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Multi-Core	"Crisis"				

Intel Processor Clock Speed (MHz)



(Source: Smoothspan)





 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Scalable
 Multi-Core
 Reachability
 (cf. Holzmann, FMCAD 2006)
 Conclusion
 Conclusion

- Exploitable parallelism must double every 2 years (Corollary of Moore's Law)
- (Graph) Reachability is basis of many verification problems
- Multi-Core Model Checking: state-of-the-art not very impressive



 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Scalable
 Multi-Core
 Reachability
 (cf. Holzmann, FMCAD 2006)
 Conclusion
 Conclusion

- Exploitable parallelism must double every 2 years (Corollary of Moore's Law)
- (Graph) Reachability is basis of many verification problems
- Multi-Core Model Checking: state-of-the-art not very impressive



Having trouble with scaling simple (enumerative) reachability?

Then what are the chances to parallelize: Liveness, partial-order reduction, symbolic reachability, ...?
 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Multi-core
 reachability
 Conclusion
 Conclusion</td

Problems for Model Checking: Visited Set

- Parallel access to hash table: correctness and efficiency
- Parallel access requires synchronization lock contention
- ► Graph traversal: Random memory access cache misses
- ► Main problem with cache lines:false sharing

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Experiments:
 SPIN
 5.2.4 (NASA/JPL)
 Conclusion
 Conclusion</td



UNIVERSITY OF TWENTE.





UNIVERSITY OF TWENTE.

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Experiments:
 LTSmin
 Symbolic
 Multi-Core
 Conclusion



UNIVERSITY OF TWENTE.

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Multi-Core	Reacha	bility			

Study

- ► Multi-core reachability: (pseudo) breadth-first, depth-first, ...
- Load balancing

Where is efficiency lost?

- Lock Contention
- Lock Convoying
- Cache-Line Sharing
- Two-Step Dance
- Stampeding

• • • •


Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Measure,	Measure,	Measure			

Analyze & distill thousands of measurements

Guides decisions what to tackle next





Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Reachabilit	v Archit	ectures			





- DiVinE 2.2: Static Partitioning
- BFS, high comm. cost, static load balancing
- SPIN 5.2.4: Shared Storage & Stack Slicing
- DFS, multiple sync. points, integrated load balancing

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Reachability	y Archite	ctures II			

Shared state storage as main sync point

- Flexible reachability algorithms
- Flexible load balancing

Shared Hashtables in Parallel Model Checking (Barnat, Ročkai 2007)

- "our shared hash tables do not scale beyond 8 cores"
- "could not investigate lockless hash table solution"
- "haven't found the cause of the scalability issues with pre-sized tables"



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Lockless H [FMCAD 2010]	ash Tabl	le: Design			

- 1 Investigate Requirements on Shared Storage
- Investigate Hardware Support fine-grained synchronization, caches,
- **3** Exploit LTSmin Infrastructure incremental (Zobrist) hashing



Model Checking Sol	koban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Lockless Hash [FMCAD 2010]	n Table:	Design			

- 1 Investigate Requirements on Shared Storage
- Investigate Hardware Support fine-grained synchronization, caches,
- 3 Exploit LTSmin Infrastructure incremental (Zobrist) hashing
- Hash Table Designed for Model Checking:
 - ► FINDORPUT operation only
 - Don't store pointers, no allocation
 - No resizing!
 - Fixed-size keys (state vectors)
 - Low memory working set improves scalability



 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Lockless Hash Table:
 Layout

 [FMCAD 2010]



See also: Cliff Click (JavaOne 2007 presentation)



Albert L. Zobrist, <u>A New Hashing Method with Application for</u> <u>Game Playing</u>, Tech. Rep. 88, Computer Sciences Department, University of Wisconsin, 1969.

UNIVERSITY OF TWENTE.

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Incremental Hashing For State Vectors
 Vector

- Zobrist matrix Z: filled with random integers
- Dependency Matrix: find modified vector elements
- Difference to Chess: possible values not statically known

Experiments



	Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
	Experimenta	al Setup				
1						

- ▶ 16-core AMD Opteron 8356, 64 GB RAM
- Linux 2.6.32+patch
- ► BEEM model database (250+ models)
 - extensive collection of models for enumerative model checkers
 - ► case studies, protocols, games, planning, synthetic models, ...
 - http://anna.fi.muni.cz/models/
- Statically sized hash tables (no resizing)
- ► Speedups: relative to best sequential(!) tool

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Experiment	s: Sum	mary			



Green LTSmin (UTwente) Red DiVinE 2.2 Blue SPIN 5.2.4

UNIVERSITY OF TWENTE.





UNIVERSITY OF TWENTE.



Static Load Balancing (SLB)

- Feasible for many models
- Threads can run out of work

Synchronous Random Polling (Sanders '97)

- Work stealing/hand-off
- Almost no overhead vs. SLB
- Additional Improvements
 - Shared-Memory Multi-Core
 - Informed Polling
 - Scare off stampeding threads





Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Observatio	ons				

Model Checking Limitations:

- Old Days: Memory
- Availability of large RAM: Time-Outs ("Patience-Out")
- Multi-Core Reachability
 - Analyzing 10 Million states/sec on 16-core AMD
 - Allocation rate: 1 GB/sec

Memory is again Bottleneck



State Vector Compression



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Tree Comp [PDMC 2007, J	pression				

- State vectors highly similar
- Compression via tree of tables, recursive version of SPIN's COLLAPSE
- Multi-Core version: Lockless hash table as building block!
- Increases arithmetic intensity: Super-linear speedups!
- "For Free"! (pays its own way)



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Tree Cor	mpression				
[PDMC 2007	7, JLC 2009]				





- ► Folded vector $\langle 2, 1 \rangle_{\rm FV}$ represents $\langle 0, 0, 1, 0, 1, 0 \rangle_{\rm IV}$ resp. $\langle O, O, X, O, X, O \rangle_{\rm SV}$ (with $O \leftrightarrow 0, X \leftrightarrow 1$)
- Selected tree fringe of grey boxes corresponds to state vector

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Tree Cor	mpression				
[PDMC 2007	7, JLC 2009]				





- $$\label{eq:Folded vector} \begin{split} \blacktriangleright \mbox{ Folded vector } \langle 2,1\rangle_{\rm FV} \mbox{ represents } \langle 0,0,1,0,1,0\rangle_{\rm IV} \mbox{ resp.} \\ \langle O,O,X,O,X,O\rangle_{\rm SV} \mbox{ (with } O\leftrightarrow 0, \ X\leftrightarrow 1) \end{split}$$
- Selected tree fringe of grey boxes corresponds to state vector
- Potential gain (here $54 \rightarrow 42$ entries):
 - ▶ main table of size *N* is only two integers wide
 - small tables of size $\mathcal{O}(\sqrt{N})$ only

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Multi-Core
 Tree
 Compression
 [NASA FM 2011, SPIN 2011]
 Value
 Value

- State vectors highly similar: Exploit combinatorial structure
- Extreme example: Memory/Time usage for firewire_tree.5:

Uncompressed: 14 GB Tree Compression: 96 MB

- Near 100 % efficiency!
- Dependency Matrix: Incremental Tree Compression



Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Observatio	ons				

- Centralized state storage scales at least as well as static state space partitioning
- Shared state storage orthogonal to: search strategy, load balancing, ...
- Simpler architecture
- Arithmetic Intensity (hide memory latency)
- Data layout important (caches), sadly no language support
- Designed from the ground up for Scalability
- Establishing correctness of implementation is a pain



 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Algorithm
 Engineering
 Engine
 Engineering
 Engineering

- ► > 10× improvement over sequential algorithms
- Beats state-of-the-art tools, reopens earlier conclusions
- Research questions guided by experiments
- Engineering approach, repeatable benchmark results



Hash table as building block for multi-core xDDs, liveness checking, ...

Conclusion



 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Model Checking for Planning?

 Conclusion

 <t

Adoption obstacles

- Modeling effort
 - many languages
 - avoid modeling?
- Scalability
 - parallel components
 - data, buffers, . . .
- Complex tools
 - algorithms, heuristics
 - Iow-level details

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Model Checking for Planning?

Adoption obstacles

- Modeling effort
 - many languages
 - avoid modeling?
- Scalability
 - parallel components
 - data, buffers, . . .
- Complex tools
 - algorithms, heuristics
 - Iow-level details

Algorithmic solutions (combinatorics, locality)

- on-the-fly model checking
- symbolic model checking
- bounded model checking
- parallel model checking
- partial-order reduction
- symmetry reduction

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Model Checking for Planning?

Adoption obstacles

- Modeling effort
 - many languages
 - avoid modeling?
- Scalability
 - parallel components
 - data, buffers, . . .
- Complex tools
 - algorithms, heuristics
 - Iow-level details

Algorithmic solutions (combinatorics, locality)

- on-the-fly model checking
- symbolic model checking
- bounded model checking
- parallel model checking
- partial-order reduction
- symmetry reduction

Problem: Algorithms tied to specification languages

- ► No particular technique suits all applications / models
- Users need to rewrite model in different languages

UNIVERSITY OF TWENTE.

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Silver Bul	lets?				

- No clear winner
- BUT: also depends on modelling
- PINS matrix can be used as predictor







Model Checking Sokoban PINS/LTSmin Symbolic Multi-Core Conclusion

Links to External Tools



 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Portfolio of Language Modules
 Final Action and A

μ CRL mCRL2	process algebra process algebra, INESS	(CWI) (TU/e)	
PROMELA	SPIN	(NASA/JPL)	Prototype
PROMELA	SpinJa	(UTwente)	
DVE	DiVinE-cluster	(MU Brno)	
DVE2	DiVinE model checking toolset	(MU Brno)	
ETF	Enumerated Table Format	(LTSmin)	
GNA	Genetic Network Analyzer	(INRIA)	Prototype
ODE	Maple	(EC-MOAN)	Prototype

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Portfolio c	of Tools				

Depth/Breadth-First Enumerative Reachability (spec)2lts-gsea $\langle spec \rangle$ -reach MDD-based symbolic reachability (spec)2lts-mpi Distributed state-space generation (spec)2lts-mc Multi-Core Enumerative Reachability (spec)2torx TorX tester RPC interface Connection to CADP toolset (VASY/INRIA) pins_open ltsmin-mpi Signature-based distributed minimization ce-mpi Orzan's distributed cycle elimination ltsmin-tracepp (Error) trace pretty printer

> Next: Partial-Order Reduction (POR), Linear Temporal Logic (LTL), μ -calculus Saturation,

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Who bene	fits from	LTSmin?			

Useful for end users: larger case studies

- ► Model in a suitable (your favourite) specification language
- Decide later what model checking algorithm to use

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Who bene	fits from	LTSmin?			

Useful for end users: larger case studies

- ► Model in a suitable (your favourite) specification language
- Decide later what model checking algorithm to use

Useful for tool providers: algorithms for free

- Your Domain Specific Specification Language can get HPMC
- Ideally, LTSmin can be viewed as a library

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Who bene	fits from	LTSmin?			

Useful for end users: larger case studies

- ► Model in a suitable (your favourite) specification language
- Decide later what model checking algorithm to use

Useful for tool providers: algorithms for free

- Your Domain Specific Specification Language can get HPMC
- Ideally, LTSmin can be viewed as a library

Useful for researchers: rigorous benchmarking

- LTSmin allows benchmarking algorithms on one model
- LTSmin allows comparing languages on one algorithm

UNIVERSITY OF TWENTE

Model Checking	Sokoban	PINS/LTSmin	Symbolic	Multi-Core	Conclusion
Literature					



Check out: http://fmt.cs.utwente.nl/tools/ltsmin/

- Alfons Laarman, Jaco van de Pol, Michael Weber, Multi-Core LTSmin: Marrying Modularity and Scalability (NFM2011,tool)
- Stefan Blom, Jaco van de Pol, Michael Weber, LTSmin: Distributed and Symbolic Reachability (CAV 2010, tool)
- Alfons Laarman, Jaco van de Pol and Michael Weber, ... (FMCAD 2010) Boosting Multi-Core Reachability Performance with Shared Hash Tables
- Stefan Blom and Jaco van de Pol, Symbolic Reachability for Process Algebras with Recursive Data Types(ICTAC'08)
- Stefan Blom, Bert Lisser, Jaco van de Pol and Michael Weber, A Database Approach to Distributed State-Space Generation (JLC 2009)

 Model Checking
 Sokoban
 PINS/LTSmin
 Symbolic
 Multi-Core
 Conclusion

 Model Checking (and related)
 Venues
 Venues



- CAV Computer Aided Verification
- TACAS Tools and Algorithms for the Construction and Analysis of Systems
- VMCAI Verification, Model Checking, and Abstract Interpretation
 - ATVA Automated Technology for Verification and Analysis
- FMCAD Formal Methods in Computer Aided Design
 - SPIN SPIN Workshop on Model Checking Software
- PDMC Parallel and Distributed Methods in VerifiCation
- MoChArt Model Checking and Artificial Intelligence
 - QEST Quantitative Evaluation of SysTems
- FORMATS Formal Modelling and Analysis of Timed Systems