

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
5 July 2007 (05.07.2007)

PCT

(10) International Publication Number
WO 2007/076313 A2

(51) International Patent Classification:
G06F 9/44 (2006.01)

(21) International Application Number:
PCT/US2006/062216

(22) International Filing Date:
18 December 2006 (18.12.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/743,055 20 December 2005 (20.12.2005) US
11/611,847 15 December 2006 (15.12.2006) US

(71) Applicant (for all designated States except US): **NEC LABORATORIES AMERICA, INC.** [US/US]; 4 Independence Way, Suite 200, Princetown, NJ 08540 (US).

(72) Inventors: **KAHLON, Vineet**; 2606 Fox Run Drive, Plainsboro, NJ 08536 (US). **GUPTA, Aarti**; 157 Broadmead Street, Princeton, NJ 08540 (US). **SINHA, Nishant**; 5700 Bunkerhill Street, Apt. 1507, Pittsburgh, PA 15206 (US).

(74) Agent: **BROSEMER, Jeffery**; Brosemer & Kolefas, One Bethany Road, Hazlet, NJ 07730 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: SYMBOLIC MODEL CHECKING OF CONCURRENT PROGRAMS USING PARTIAL ORDERS AND ON-THE-FLY TRANSACTIONS



(57) Abstract: A set of techniques for analyzing concurrent programs that combines the power of symbolic model checking to explore large state spaces, and partial order and transaction-based reduction techniques to manage the size of explored state space.

SYMBOLIC MODEL CHECKING OF CONCURRENT PROGRAMS USING PARTIAL ORDERS AND ON-THE-FLY TRANSACTIONS

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of United States Provisional Application No. 60/743,055 filed 20 DEC 2005 the entire contents of which are incorporated by reference as if set forth at length herein..

FIELD OF THE INVENTION

[0002] This invention relates generally to the field of computer software and in particular it pertains to a software verification methodology for concurrent programs.

BACKGROUND OF THE INVENTION

[0003] The widespread use of concurrent software in modern computing systems necessitates the development of effective verification methodologies for multi-threaded programs. As can be appreciated however, subtle interactions between threads makes multi-threaded software behaviorally complex and particularly hard to analyze and - as a result - formal methodologies are employed for their debugging. Not surprisingly, model checking – both symbolic and explicit state – for the verification of concurrent software has been an active area of research.

[0004] Explicit state model checkers, such as Verisoft (See e.g., P.Godefroid, “*Model Checking For Programming Languages Using Verisoft*”, POPL '97, pp.174-186, 1997) explore an enumeration of the states and transitions of the concurrent program under study. Additional techniques such as state hashing for compaction of state representations, and partial order methods are typically used to avoid exploring all of the interleavings and transitions of constituent threads. And while these techniques have proven to be effective at state space reduction, they do not address scalability problems that arise due to state explosion when model checking large-scale concurrent programs.

NL-05-046

2

[0005] Symbolic model checkers – on the other hand – avoid an explicit enumeration of the state space by using symbolic representations of sets and states and transitions. One successful approach in this regard was the use of Binary Decision Diagrams (BDDs) to succinctly represent large state spaces for the purpose of model checking (See, e.g.,

5 K.L.McMillan, “*Symbolic Model Checking: An Approach To The State Explosion Problem*, Kluwer Academic Publishers, 1993). Subsequently, Boolean Satisfiability (SAT)-based techniques have become popular both for finding software bugs using SAT-based bounded model checking (BCC) and generating proofs via SAT-based unbounded model checking (UMC).

10

[0006] Given their importance, techniques that improved upon or extended the applicability of model checking would represent a significant advance in the art.

SUMMARY OF THE INVENTION

15 [0007] We have developed, in accordance with the principles of the invention, methodology which advantageously leverages the synergy which results from combining partial order techniques to reduce the state space of a system to be explored with the power of symbolic model checking techniques to explore large state spaces. In sharp contrast to existing methods that employ BDDs which encode the entire state of a given

20 concurrent program thereby producing a state space explosion – the method of the present invention provides the freedom to use any technique of choice – either SAT or BDD-based. As those skilled in the art will readily appreciate, such an approach is much more scalable than the prior art approach(es) which required the use of BDDs.

25 [0008] According to an aspect of the present invention, a given concurrent program is translated into a circuit-based (finite-state) model. Accordingly, a finite model for each individual thread is obtained wherein each variable of the thread is represented in terms of a vector of binary-valued *latches* and a Boolean next-state function (or relation) for each latch. Next – using a scheduler – the circuits for the individual threads are

30 composed into one single circuit for the entire concurrent program. Verification is then performed on this circuit and partial order techniques are incorporated into the framework

NL-05-046

3

by statically augmenting the circuit-based Boolean encoding of the concurrent program with additional constraints. According to an aspect of the present invention – these constraints restrict the transitions explored from each global state to a minimal *conditional stubborn set* of that state.

5

[0009] Viewed from yet another aspect, the present invention provides an improved method for identifying transactions *on-the-fly* that is based upon analyzing patterns of lock acquisitions. In sharp contrast, prior art methods employ lockset based analysis. As those skilled in the art will appreciate, lockset based methods for state space reduction exploit the ability of locks to enforce mutually exclusive access to regions of code encapsulated between locking and unlocking operations. Such prior art lockset methods rely on the assumption that a concurrent program follows a lock discipline in accessing shared variables, i.e., that all accesses to a shared variable sh are protected by the same lock I_{sh} .

15

[0010] According to an aspect of the present invention however, patterns of lock acquisitions are analyzed – rather than locksets – thereby producing a demonstrably more comprehensive result. In addition, a method according to the present invention does not require nor rely on a concurrent program exhibiting lock discipline. Consequently, the present invention permits the use of lock-based reductions for a broader class of concurrent programs.

20

[0011] Viewed from yet another aspect, the present invention permits the transparent incorporation of lock-pattern based transactions into partial order reductions by improved conditional dependency detection via the addition of extra constraints – which are not incorporated into the transition relation a-priori but dynamically while unrolling the executions of the threads. As a result, increased granularity of transitions due to transactions can be captured as a reduction in the sizes of conditional stubborn sets of states.

25

30

Printed: 01/19/2007

SPEC

PCT/US06622

NL-05-046

4

[0012] Finally, the present invention provides a new approach for model checking concurrent programs that combines the power of symbolic techniques with partial order reduction and on-the-fly transactions while – at the same time – retaining the flexibility to employ a broad arsenal of model checking techniques – both SAT and BDD-based – to
5 check not just reachability but richer classes of linear temporal problems as well.

BRIEF DESCRIPTION OF THE DRAWING

[0013] A more complete understanding of the present invention may be realized by reference to the accompanying drawing in which:

10

[0014] FIG. 1 is a program segment showing threads T_1 FIG 1(a) and T_2 FIG 1(b) with unprotected access to x ;

[0015] FIG. 2 is a program segment showing threads T_1 FIG 1(a) and T_2 FIG 1(b) with
15 unprotected access to x illustrating the identification of transactions in the absence of lock discipline; and

[0016] FIG. 3 is a block diagram depicting an overview of the present invention.

20 DETAILED DESCRIPTION

[0017] The following merely illustrates the principles of the invention. It will thus be appreciated that those skilled in the art will be able to devise various arrangements which, although not explicitly described or shown herein, embody the principles of the invention
25 and are included within its spirit and scope.

[0018] Furthermore, all examples and conditional language recited herein are principally intended expressly to be only for pedagogical purposes to aid the reader in understanding the principles of the invention and the concepts contributed by the inventor(s) to

Printed 01/19/2007

SPEC

PCTUS06622

NL-05-046

5

furthering the art, and are to be construed as being without limitation to such specifically recited examples and conditions.

5 [0019] Moreover, all statements herein reciting principles, aspects, and embodiments of the invention, as well as specific examples thereof, are intended to encompass both structural and functional equivalents thereof. Additionally, it is intended that such equivalents include both currently known equivalents as well as equivalents developed in the future, i.e., any elements developed that perform the same function, regardless of structure.

10

[0020] Thus, for example, it will be appreciated by those skilled in the art that the diagrams herein represent conceptual views of illustrative structures embodying the principles of the invention.

15 [0021] By way of additional theoretical background, we consider concurrent systems having a finite number of processes or threads where each thread is a deterministic sequential program written in a language such as C. As is known, threads may interact with each other using communication/synchronization objects like shared variables, locks and semaphores.

20

[0022] Formally, we define a concurrent program CP as a tuple (T, V, R, s_0) where $T = \{T_1, \dots, T_n\}$ denotes a finite set of threads, $V = \{v_1, \dots, v_m\}$ a finite set of shared variables and synchronization objects with v_i taking on values from the set V_i , R the transition relation and s_0 , the initial state of CP . Each thread T_i is represented by a control flow graph of the sequential program it executes, and is denoted by the pair (C_i, R_i) , where C_i denotes the set of control locations of T_i and R_i its transition relation.

25

[0023] A global state s of CP is a tuple:

$$(s[1], \dots, s[n], v[1], \dots, v[m]) \in S = C_1 \times \dots \times C_n \times V_1 \times \dots \times V_m;$$

NL-05-046

6

where $s[i]$ represents the current control location of thread T_i and $v[j]$ the current value of variable v_j . The global state transition diagram of CP is defined to be the standard interleaved composition of the transition diagrams of the individual threads. Thus each global transition of CP results by firing a local transition of the form (a_i, g, u, b_i) where

5 a_i and b_i are control location in some thread $T_i = (C_i, R_i)$ with $(a_i, b_i) \in R_i$; g_i is a guard which is a Boolean-valued expression on the values of local variables of T_i and global variables in V ; and u is a function that encodes how the value of each global variable and each local variable of T_i is updated.

- 10 [0024] A transition $t = (a_i, g, u, b_i)$ of thread T_i is enabled in state s iff $s[i] = a_i$ but g need not be true in s , then we simply say that t is *scheduled* in s . We write $s \xrightarrow{t} s'$ to mean that the execution of t leads from states s to s' . Given a transition $t \in T$, we use $proc(t)$ to denote the process executing t . Finally, we note that each concurrent program CP with a global state space S defines the global transition system $A_g = (S, \Delta, s_0)$, where
- 15 $\Delta \subseteq S \times S$ is the *transition relation* defined by $(s, s') \in \Delta$ iff $\exists t \in T : s \xrightarrow{t} s'$; and s_0 is the initial state of CP .

Lock Synchronization Based Reductions

- [0025] We begin our discussion through the use of motivating examples. Consider a
- 20 concurrent program CP shown as a program segment in **FIG 1**. With reference to that **FIG 1**, we see two threads T_1 and T_2 shown in **FIG 1(a)** and **FIG 1(b)**, respectively. We note that x , which is the only variable shared among the two threads is unprotected at control location 5b and protected by lock lk at all other locations. Since x is unprotected at all locations where it is accessed, it does not satisfy the lock discipline
- 25 mentioned earlier and (See, e.g., “*Model Checking Multi-Threaded Distributed JAVA Programs*”, authored by Scott D. Stoller and which appeared in International Journal On

NL-05-046

7

Software Tools For Technology Transfer, 4(1), pp.71-91, Oct. 2002) which will therefore force a context switch before locations 3a and 3b.

5 [0026] Consider a global state s of CP with threads T_1 and T_2 at control locations 3a and 3b respectively. A key observation is that starting at global state s of CP , 3a does not interfere with 3b and 5b even though 5b is unprotected. This is due to the fact that for T_2 to execute 3b it has to acquire lk currently held by T_1 . But in order for T_1 to release lk , it must first execute 3a.

10 [0027] Thus starting at s , CP is forced to execute 3a before 3b. As a result no context switch is required before 3a. However, in the global state s' with T_1 and T_2 at control locations 3a and 5b respectively, the transitions 3a and 5b do interfere with each other thereby forcing a context switch before 3a. As can be appreciated by those skilled in the art, even though shared variables need not follow a locking discipline globally, 15 there are still identifiable portion of the state space where locking discipline is followed. Thus a context driven analysis allows us to define transactions locally – *on-the-fly* – where prior art methods – because of their reliance on global analysis – fail to do so.

20 [0028] Taking this further, we can now show that transactions may be identified even in the absence of lock discipline – local or global. With reference now to **FIG 2**, there is shown program segment threads T_1 and T_2 in **FIG 2(a)** and **FIG 2(b)** respectively each having unprotected access to x . We let CP be a concurrent program comprising these two threads T_1 and T_2 both sharing variable x as shown.

25 [0029] Consider a global state s of CP with threads T_1 and T_2 in control locations 6a and 1b, respectively. Observe that starting at s , the transitions at control locations 6a and 6b cannot interfere with each other even though they access the same shared variable x . This is because in order for thread T_2 to reach location 6b from location 1b it has to traverse the local path 1b, 2b, 3b, 4b, 5b, along which it has to acquire (and release)

NL-05-046

8

lock lk_1 currently held by T_1 . In order for that to happen, T_1 must release lk_1 for which it must execute transition 6a. As a result, transition 6a is forced to be executed before transition 6b. Thus no context switch is required before location 6a.

- 5 [0030] One key observation to be made here is that even though disjoint sets of locks were held at locations 6a and 6b, it was the set of locks that needed to be acquired by T_2 in order to transit from 1b to 6b (even though some of these locks were released before reaching 6b) that prevented 6a and 6b from interfering with each other. A traditional, prior-art, lockset-based analysis such as presented in [Sto02,FQ03] would treat 6a and
10 6b as conflicting transitions (as x does not follow locking discipline) and force a context switch before these locations.

- [0031] Consequently, those skilled in the art will recognize that a conflict analysis based on lock acquisition *patterns* according to the present invention is more refined than one
15 based on locksets.

Transactions VIA Persistent Sets

- [0032] We may now show how to integrate lock-pattern based on-the-fly transactions with partial order reduction in a transparent fashion by capturing the increased granularity
20 of transitions due to transactions as a reduction in the sizes of the conditional stubborn sets of states. This is accomplished by ensuring that if in a global state s , a thread T_i is executing a transaction then, in the persistent set of s , we include only one transition, viz., the transition of T_i that fires next along the transaction being executed. This ensures that once the first transition of a transaction is executed, by a thread T_i then no other process
25 can be scheduled unless all transitions of the transaction finish firing.

- [0033] State space reduction using partial order techniques is obtained by exploring from each individual state only those transitions that belong to a persistent set of that individual state instead of all the enabled transitions. Although there are many ways to

NL-05-046

9

compute persistent sets, a method of computing conditional stubborn sets usually generates those with small cardinality. For our purposes herein, we use standard terminology from the theory of partial order reductions and the algorithm for computing conditional stubborn sets (See, e.g., P.Godenfroid, “*Partial Order Methods For The*
 5 *Verification of Concurrent Programs: An Approach To The State Explosion Problem*”, LNCS 1032, Springer-Verlag, 1996) which we denote by $Algo_1$.

[0034] . We begin by recalling the following definition:

10 [0035] **Might-be-first-to-interfere:** Let op and op' be two operations on the same object O and s be a reachable state. The relation $op \triangleright_s op'$ holds if there exists a sequence

$s = s_1 \xrightarrow{t_1} s_2 \xrightarrow{t_2} \dots \xrightarrow{t_n} s_{n+1}$ of transitions in A_G such that $\forall 1 \leq i < n : \forall op''$ on O used by t_i :
 op and op' are dependent in s_n .

15 [0036] For each local transition $a \xrightarrow{g} b$ of a thread, we let $used(t)$ denote the set of operations on variables and synchronization objects executed during the execution of t . A conditional stubborn set of state s of A_G can then be calculated as follows:

1. Initialize $T_s = \{t\}$, where t is some enabled transition in s .
2. For each $t = a \xrightarrow{g} b \in T_s$
 - (a) If t is disabled in s ,
 - i. If $T_j = \text{proc}(t)$ and $s[j] \neq a$ then add to T_s all transitions t' of T_j of the form $c \xrightarrow{g'} a$, or
 - ii. Choose a condition c_j in the guard g of t that evaluates to false in s ; then, for all operations op used by t to evaluate c_j , add to T_s all transitions t' such that $\exists op' \in \text{used}(t') : op \triangleright_s op'$
 - (b) If t is enabled in s add to T_s all transitions t' such that $\text{proc}(t) \neq \text{proc}(t')$ and $\exists op \in \text{used}(t), \exists op' \in \text{used}(t') : op \triangleright_s op'$
3. Repeat step 2 until no more transitions can be added in T_s . Then return all transitions in T_s that are enabled in s .

Algo₁ for Computing Conditional Stubborn Sets

- [0037] In *Algo₁* dependencies between transitions, arising out of operations on shared communication objects are captured using the $\triangleright_s$ relation which captures for each operation op used by a transition in a state s which other operations *might be first to interfere with op from the current state s* . In practice, to avoid exploration of the state space of the program at hand, static analysis is employed in order to compute a relation, $\triangleright_s^{st}$, which is an over-approximation of $\triangleright_s$. Towards that end, we say that two operations op and op' are *statically dependent* if they access a common shared variable such that at least one of the accesses is a write operation. Then $\triangleright_s^{st}$, is defined as follows.

- [0038] **Definition:** Let op and op' be two operations on a common shared variable and s is a reachable state of A_G . The relation $\triangleright_s^{st} op'$ holds iff there exist distinct threads T_i and T_j such that there exists (1) a transition of T_i scheduled – but not necessarily enabled – at s using op , and (2) a local path $x: p_0 \xrightarrow{i_1} \dots \xrightarrow{i_n} p_n$ of T_j such that p_0 is the

NL-05-046

11

local state of T_j in s , $\forall 1 \leq k < n : \forall op''$ is used by $t_k : op$ and op'' are not statically dependent, t_n uses op' , and op and op' are statically dependent.

[0039] To incorporate on-the-fly transactions, we modify the above definition of $\triangleright_s^{st}$ to
 5 obtain a new relation $\triangleright_s^{lp} \subseteq \triangleright_s^{st}$ by adding (in accordance with our discussion above),
 the extra constraint that none of the locks held by T_i in s is acquired (and possibly
 released) by T_j along x . Note that since $\triangleright_s^{lp}$ is more constrained it enforces fewer
 dependencies between operations than $\triangleright_s^{st}$ thus resulting in smaller conditional stubborn
 sets. As a result, certain interleavings are “weeded out” to produce the effect of
 10 executing transactions.

[0040] Indeed - in the example given in FIG 2 - in global state s , if op and op' are the
 operations $x=0$ and $x=1$ at locations 6a and 6b, respectively, then $op \triangleright_s^{st} op'$ but
 $\neg(op \triangleright_s^{lp} op')$. Thus, using $\triangleright_s^{lp}$ instead of $\triangleright_s^{st}$ to compute conditional stubborn sets
 15 removes transition 1b from the conditional stubborn set s of thereby preventing a
 context switch before 6a.

[0041] Formally, $\triangleright_s^{lp}$ is defined as follows.

20 [0042] **Definition (might-be-the-first-to-interfere-modulo-lock-acquisition)** Let op
 and op' be two operations on a common shared variable and s a reachable state of A_G . The relation
 $op \triangleright_s^{lp} op'$ holds iff there exist distinct threads T_i and T_j such that there exists: (1) a transition of T_i
 scheduled (although not necessarily enabled) at s using op and (2) a local path $x : p_0 \xrightarrow{t_1} \dots \xrightarrow{t_n} p_n$ of T_j
 such that $\forall 1 \leq k < n : \forall op''$ used by $t_k : op$ and op' are not statically dependent, t_n uses op' , and op
 25 and op' are statically dependent and no lock held by T_i in s is acquired by T_j along x .

NL-05-046

12

[0043] Now, if we let $Algo_2$ be the result of replacing $\triangleright_s$ in $Algo_1$ by $\triangleright_s^{st}$ and $Algo_3$ the result of replacing $\triangleright_s^{st}$ in line 2 . (b) . i of $Algo_2$ by $\triangleright_s^b$. Then the following two results state that $Algo_3$ does advantageously compute a conditional stubborn set that is smaller than one computed by $Algo_2$. Note however, that although we used a specific relation $\triangleright_s^{st}$ for computing dependencies statically, one can of course incorporate on-the-fly transactions with any other *implementation* of $\triangleright_s$ by merely adding the extra condition regarding lock acquisition patterns, as above.

[0044] **Theorem 1.** *All sets T_s that are computed by $Algo_3$ are conditional stubborn sets of s .*

[0045] **Proof Sketch:** Let $t = a \xrightarrow{g} b$ executed by thread T_i belong to T_s . Let $w = s_1 \xrightarrow{t_1} s_2 \xrightarrow{t_2} \dots \xrightarrow{t_n} s_{n+1}$ be a sequence of transitions of A_G such that t is dependent with t_n in s_n . We need to show that at least one of $t_1, \dots, t_n$ is in T_s . Without loss of generality, we may assume that for $1 \leq i < n$, t is independent with t_i in s_i and t_n is dependent with t in s_n , else we can pick an appropriate prefix of w .

[0046] First, we assume that t is disabled in s . Since t is disabled in s , and s_n is the first state along w in which t is dependent (with t_n), we have that t is enabled in s_{n+1} . Since t is disabled in s , either $s[i] \neq a$, or a condition c in guard g evaluates to false in s . In the first case, since t is enabled in s_{n+1} , there exists a transition t_j fired along w , of the form $d \rightarrow a$ labeled with some guard g' . But then executing step 2 . (a) . i of $Algo_3$ would cause t_j to be included in T_s .

[0047] In the second case, there exists a transition t_j that changes the value of c from false to true by changing the output of an operation op used to evaluate c , i.e., by

NL-05-046

13

performing an operation op' dependent with op in s_j . Let t_j be the first such transition occurring along w . Clearly, op' is statically dependent with op . By definition of $\triangleright_s^{st}$, we have $op \triangleright_s^{st} op'$, and so $t_j \in T_s$ by step 2.a.(ii).

- 5 [0048] We may now consider the case where t is enabled in s . From the facts that: (i) for $1 \leq j \leq n-1$, t is independent with t_j in s_j , and (ii) t is enabled in s , we have that for $1 \leq j \leq n-1$, t is enabled in s_j . This implies that the thread T_i does not execute any transition along w , otherwise – since T_i is deterministic – we conclude that t is the first transition that T_i executes along w .

10

[0049] As can be appreciated, this would force T_i out of its current local state thereby disabling t and thereby contradicting the above observation. Note that here we assumed that executing a transition takes a process out of its current local state, i.e., there are no self loops in a program thread – which is a reasonable assumption for software programs.

15

[0050] Now, since t and t_n are dependent in s_n , it implies that $\exists op \in used(t), \exists op' \in used(t_n) : op$ and op' are dependent in s_n and therefore are also statically dependent. If we let t_j be the first transition along w that uses an operation op'' dependent op . Note also that there does not exist a lock l held by T_i at s such that l

- 20 has to be acquired before t_j is executed along w . Otherwise, l must first be released by T_i thus forcing T_i to execute a transition contradicting our observation made above that T_i does not execute any transition along w . Thus we have $op \triangleright_s^{lp} op''$ and hence $t_j \in T_s$ by step 2.b.(i).

- 25 [0051] **Theorem 2.** *For all transitions t that are enabled in s , for all persistent sets $Algo_2$ that can be returned by $Algo_2$, there exists a run of $Algo_3$ that returns a persistent set $Algo_3(t) \subseteq Algo_2$.*

NL-05-046

14

[0052] **Proof Sketch:** From the definition of relation $\triangleright_s^{lp}$, it follows that $\triangleright_s^{lp}$ is included in $\triangleright_s^{st}$. Thus the set T_s returned by $Algo_3$ is always a subset of the one returned by $Algo_2$ provided the same choices are made in case of nondetermination.

5 Software Modeling For Concurrent C Programs

[0053] Translating Individual Threads Into Circuits

[0054] We may now describe how - using F-Soft - we first obtain a circuit-based model of each thread, under the assumption of bounded data and bounded control (recursion)
 10 (See, e.g., F.Ivancic et al. "Model Checking C Programs Using F-Soft", In ICCD, 2005). Briefly, we begin with a C program and apply a series of source-to-source transformations to simplify complex C expressions into smaller but equivalent subsets of C. Next, all arrays and structs are "flattened" by replacing them with collections of simple scalar variables, and then build an internal memory representation of the program
 15 by assigning to each scalar variable a unique number representing its memory address.

[0055] Variables that are adjacent in C program memory are given consecutive memory addresses in our model; which advantageously facilitates modeling of pointer arithmetic. The heap is modeled as a finite array, by adding a simple implementation of `malloc()`
 20 that returns pointers into this array.

[0056] For handling pointer accesses, we first perform a "points-to" analysis to determine the set of variables that a pointer variable can point to. Then, we convert each indirect memory access, through a pointer or an array reference, to a direct memory access. For
 25 example, if we determine that pointer `p` can point to variables `a, b, ..., z` at a given program location, we rewrite a pointer read `*(p+i)` as a conditional expression of the form:

$$((p+i) == \&a ? a : ((p+i) == \&b ? b : \dots)),$$

where `&a, &b, ...` are the numeric memory addresses we assigned to the variables

30 `a, b, ...,` respectively.

NL-05-046

15

[0057] Nonrecursive function calls are handled by inlining exactly once, and replacing that particular function's return by a set of goto-s conditioned upon the unique call site id stored on that function's entry. Bounded recursive functions are modeled by introducing a bounded call stack. While we aim for accurate modeling of all C, practical modeling
5 requires making approximations.

[0058] Accordingly, large arrays are truncated. Writes to elements above a certain index are ignored, and reads from these elements yield non-deterministic values. Floating-point values are approximated by modeling their integral parts only. The simplified program
10 includes scalar variables of simple types (Boolean, enumerated, integer). This is compiled using standard techniques into its control flow graph (CFG). T

[0059] Those skilled in the art will recognize that the CFG representation can be viewed as a finite state machine with state vector (pc, V) , where pc denotes an encoding of the
15 basic blocks, and V is a vector of integer-valued program variables. We then construct symbolic transition relations for pc , and for each data variable appearing in the program. For pc , the transition relation reflects the guarded transitions between basic blocks in the CFG counter. For a data variable, the transition relation is built from expressions assigned to the variable in various blocks. Finally, we construct a symbolic representation
20 of these transition relations resembling a hardware circuit. For the pc variable, we allocate $\lceil \log N \rceil$ latches, where N is the total number of basic blocks. For each C program variable, we allocate a vector of n latches, where n is the bit width of the variable. At the end, we obtain a circuit-based model of each thread of the given concurrent program, where each variable of the thread is represented in terms of a vector
25 of binary-valued *latches* and a Boolean next-state function (or relation) for each latch.

[0060] Building The Circuit For The Concurrent Program

[0061] Given the circuit C_i for each individual thread T_i , we may now show how to get the circuit C for the concurrent program CP comprised of these threads. In the case
30 where local variables with the same name occur in multiple threads, to ensure consistency

NL-05-046

16

we prefix the name of each local variable of thread T_i with `thread i`. Next, for each thread we introduce a gate `execute_i` indicating whether P_i has been scheduled to execute in the next step of CP or not.

- [0062] For each latch l , we let $next-state_i(l)$ denote the next state function of l in circuit C_i . Then in circuit C , the next state value of latch `thread_i_l` corresponding to a local variable of thread T_i , is defined to be $next-state_i(\text{thread\_i\_l})$ if `execute_i` is true, and the current value of `thread_i_l`, otherwise. If, on the other hand, latch l corresponds to a shared variable, then $next-state(l)$ is defined to be $next-state_i(l)$, where `execute_i` is true. Note that we need to ensure that `execute_i` is true for exactly one thread T_i . Towards that end, we implement a scheduler which determines in each global state of CP which one of the signals `execute_i` is set to *true* and thus determines the semantics of thread composition.

[0063] Conditional Stubborn Sets based Persistent Sets

- [0064] To incorporate partial order reduction, we need to ensure that from each global state s only transitions belonging to a conditional stubborn set of s are explored. We let R and R_i denote the transitions relations of CP and T_i , respectively. If CP has n threads, we introduce the n -bit vector *cstub* which identifies a conditional stubborn set for each global state s , i.e., in $s, cstub_i$ is *true* for exactly those threads T_i such that the (unique) transition of T_i - enabled at s - belongs to the same minimal conditional stubborn set of s . Then:

$$R(s, s') = \bigvee_{1 \leq i \leq n} ((\text{execute\_i}) \wedge cstub_i(s) \wedge R_i(s, s')).$$

[0065] The *cstub* vector may be computed as follows:

25

1. For each shared variable x and thread T_i , we introduce a latch $touch\_now(T_i, x)$ which is true at control location pc_i of T_i iff T_i accesses x at control location pc_i . This can be done via static analysis of the CFG of T_i by determining at which

NL-05-046

17

control locations x was accessed and taking a disjunction for those values of pc_i .

2. For each shared variable x , and thread T_i , introduce the latch $touch\text{-}now\text{-}later(T_i, x)$, which is true at control location pc_j . Thus, computing $touch\text{-}now\text{-}later(T_i, x)$ involves deciding the reachability of pc'_j , and since it cannot be computed exactly without exploring the entire state space A_G of CP , we over-approximate it by performing a context-sensitive analysis of the control-flow graph of T_j . We set $touch\text{-}now\text{-}later\text{-}pair(T_j, x)$ to *true* in control pc_j if for some control pc'_j in the control flow graph of T_j , x is accessed at pc'_j .
3. For distinct threads T_i and T_j , the relation $conflict_i(j)$ is then defined as $\bigvee_{x \in V_{sh}} (touch\text{-}now(T_i, x)(pc_i) \wedge touch\text{-}now\text{-}later(T_j, x)(pc_j))$, where pc_i and pc_j are the control locations of T_i and T_j , respectively, in the current global state and V_{sh} is the set of shared variables of CP .
4. Using a circuit to compute transitive closures, for each i , starting with $J_i = \{i\}$ we compute the closure of J_i under the *conflict* relation defined above.
5. We build a circuit to compute the index *min* such that the cardinality of J_{min} is the least among the sets $J_1, \dots, J_n$. Finally $\forall 1 \leq i \leq n$, set $cstub_i = 1$ iff $i \in J_{min}$. Note that in the implementation we need to pick only one set with the least cardinality.

[0066] Cycle Detection: We first identify *sticky* transitions for all potential global cycles. We then force a conflict for the process containing the sticky locations with all other processes via the encoding below.

5

[0067] More particularly, we let $sticky(pc)$ be a predicate evaluating to true iff location pc has been marked sticky. Then, for global state s , we define

$$conflict_i(j) = sticky(pc_i) \vee (touch\text{-}now(T_i, x)(pc_i) \wedge touch\text{-}now\text{-}later(T_j, x)(pc_j))$$

NL-05-046

18

where pc_m is the current control location of T_m in s . In other words, if pc_i is sticky then thread T_i is said to conflict with all other threads. This implies that either a thread T_k - with smaller conflict set J_k - would be chosen for the persistent set computation or a full expansion would be forced.

5

[0068] Those skilled in the art will now recognize that this reduction is sound, since any cycle in the global state space can be projected on to one or more local cycles in the control flow graph of the individual threads. By forcing a full expansion inside each (potential) local cycle with the help of sticky transitions, we advantageously ensure that there is no global cycle such that a thread transition is postponed at each state of the cycle. Therefore this encoding allows the model checker to explore a conservative over-approximation of the representative (minimal) set of interleavings of the given threads. Although the reduced model remains sound, the number of interleavings considered may decrease dramatically with the number of annotated sticky transitions.

15

[0069] Encoding Lock Pattern Based Reduction

[0070] In order to incorporate transactions *on-the-fly*, we advantageously have augmented the predicate *touch-now-later*, to generate the new predicate *touch-know-later-LS* that also includes lock acquisition pattern information. For control locations pc_i and pc'_i of thread T_i , we let $paths(pc_i, pc'_i)$ denote the set of paths in the CFG of T_i starting from pc_i that may reach pc'_i . For each $\pi \in paths(pc_i, pc'_i)$ of T_i , let $lockPred(\pi)$ be a formula denoting the set of locks acquired (and possibly released) among π , e.g., $lk_1 = T_1 \wedge lk_2 = T_1$.

25 [0071] Let *touch-now-later-pair* $(T_j, x)(pc') \wedge AP_x(pc_j, cp'_j)$, where

$AP_x(pc_i, pc'_i) = \bigvee_{\pi \in paths(pc_i, pc'_i)} lockPred(\pi)$. Let $CLP(T_i, s)$ denote a formula encoding the ownership of locks T_i in global state s . Then the relation *touch-now-LS* (T_i, x) is

obtained from *touch-now-later-pair*(T_i, x) by quantifying out pc'_i in conjunction with the $CLP(T_i, s)$, i.e.,

$$touch\text{-}now\text{-}LS(T_i, x)(pc_i) = \left(\exists pc'_i \text{touch}\text{-}now\text{-}later\text{-}pair(T_i, x)(pc_i, pc'_i) \right) \wedge CLP(T_i, s)$$

5 [0072] Therefore, $touch\text{-}now\text{-}LS(T_i, x)(pc_i)$ is true if there is a location pc'_i accessing a shared variable x that is reachable from pc_i via a local path π in T_i such that no lock held in s is acquired along π . We evaluate $lockPred(\pi)$ using a context sensitive static analysis of the CFG of T_i .

10 [0073] With the theoretical basis in place we may now summarize our inventive method which is shown in a block diagram in **FIG 3**. In particular, and with reference to that figure, a number of individual threads **310[1]...310[n]** which comprise a concurrent multi-threaded program are reduced into a like number of reduced threads **320[1]...320[n]** through a static analysis including a number of a variety of known
15 methods including slicing, range analysis and constant folding. These reduced threads are further translated into a circuit-based (finite state) model **330[1]...330[n]** for each individual thread respectively where each variable of the thread is represented in terms of a vector of binary-valued latches and a Boolean next-state function (or relation) for each latch.

20 [0074] The individual circuits **330[1]...330[n]** are combined by a scheduler into a single circuit for the entire concurrent program to which constraints are added **350** for partial order reduction, on-the-fly lockset reduction, acquisition history reduction and/or synchronous execution and constraints are added. Finally, the circuit is verified using
25 symbolic model checking **360**.

[0075] The Daisy Case Study

[0076] We have employed our method of the present invention to find bugs in the Daisy file system which those skilled in the art will recognize as a benchmark for analyzing the

NL-05-046

20

efficacy of different concurrent program verification methodologies for verifying concurrent programs. Daisy is a Java implementation of a toy file system where each file is allocated a unique inode that stores the file parameters and a unique block which stores data. One interesting feature of Daisy is that it has fine grained locking in that access to each file, inode or block is guarded by a dedicated lock. Moreover, the acquire and release of each of these locks is guarded by a 'token' lock. Consequently control locations in the program might possibly have multiple open locks and furthermore the acquire and release of a given lock can occur in different procedures.

10 [0077] Currently F-Soft only accepts programs written in C so we first manually translate the Daisy code which is written in Java into C. Furthermore, to reduce the model sizes, we truncated the sizes of the data structures modeling the disk, inodes, blocks, file names, etc., which were not relevant to the race conditions we checked, resulting in a sound and complete *small-domain* reduction. We have shown the existence of the race conditions
15 described below and noted in the art.

[0078] The efficacy of our techniques can be evaluated from the fact that our model checking methodology according to the present invention is able to detect these race conditions in Daisy in a fully automatic fashion directly on the source code without any
20 code structuring/abstractions beyond redefining the constants as discussed above.

[0079] Daisy maintains an allocation area where for each block in the file system a bit is assigned 0 or 1 accordingly as the block has been allocated to a file or not. But each disk operation reads/writes an entire byte. Two threads accessing two different files might
25 access two different blocks. However since bytes are not guarded by locks in order to set their allocation bits these two different threads may access the same byte in the allocation block containing the allocation bit for each of these locks thus setting up a race condition.

[0080] The verification statistics we observed are as follows: We ran our experiments on
30 a machine with an Intel Pentium4 3.20GHz processor and 2GB RAM. Each run was given a timeout of 2 days and had a memout of 2GB. Witnesses for the above race

NL-05-046

21

condition were found in two cases, those corresponding to blocks 0 and 1, and those due to blocks 1 and 2. In sharp contrast, when using purely interleaved scheduling, we failed to find either witness because of a "memout" at depth 15.

5 [0081] When only partial order reduction was employed, was found using SAT-based BMC at unroll depth 122 in 36707 sec and 999MB while incorporating on-the-fly transactions drastically reduced the time and memory usage to 1283sec and 122MB, respectively. The second witness was found at depth 151. Using partial order reduction techniques alone took 145176 sec and 1870 MB, while adding transactions reduced it to
10 5925 sec and 902 MB.

[0082] In Daisy reading/writing a particular byte on the disk is broken down into two operations: a seek operation that mimics the positioning of the head and a read/write operation that transfers the actual data. Due to this separation between seeking and data
15 transfer a race condition may occur. For example, reading two disk locations, say n and m , we must make sure that $seek(n)$ is followed by $read(n)$ without $seek(m)$ or $read(m)$ scheduled in between. In this case a witness was found at depth 48. Using partial order reduction alone took 2.99 sec and 5.7 MB while adding transactions reduced it to 2.89 sec and 5.5 MB. For this example also BMC on the completely interleaved model failed to
20 find a witness because of a memout at depth 20

[0083] Advantageously, and as can be readily appreciated by those skilled in the art - for deep bugs techniques that leverage the use of on-the-fly transactions combined with partial order reduction greatly outperform those which use only partial order reduction --
25 both in terms of time taken and memory used.

[0084] At this point, while we have discussed and described our invention using some specific examples, those skilled in the art will recognize that my teachings are not so limited. Accordingly, our invention should be only limited by the scope of the claims
30 attached hereto.

What is claimed is:

- 1 1. A computer implemented method for analyzing a concurrent program comprising the
2 steps of:
 - 3 generating a model of the concurrent program; and
 - 4 verifying the concurrent program through the use of a symbolic model checker;
- 5 THE METHOD CHARACTERIZED IN THAT
- 6 the model is reduced through the application of a lock acquisition history analysis.
- 1 2. The method claim 1 further CHARACTERIZED IN THAT:
2 the acquisition history analysis reduces the number of stubborn sets .
- 1 3. The method of claim 2, further CHARACTERIZED IN THAT:
2 the concurrent program need not exhibit any substantial lock discipline.
- 1 4. The method of claim 3 further CHARACTERIZED IN THAT:
2 a set of transactions are determined based upon the lock acquisition history
3 analysis and information about the determined transactions are used to further
4 reduce the number of stubborn sets.
- 1 5. The method of claim 4 wherein any constraints of the stubborn sets are represented
2 symbolically.
- 1 6. The method of claim 5 wherein the model of the concurrent program is represented
2 symbolically in circuit-form.
- 1 7. A computer implemented method for analyzing a concurrent program comprising a
2 number of individual threads, said method comprising the steps of:
 - 3 generating a model of the concurrent program; and
 - 4 verifying the concurrent program through the use of a symbolic model checker;

5 THE METHOD CHARACTERIZED IN THAT
6 the model is reduced through the application of a lock acquisition history analysis
7 wherein said lock acquisition history analysis is performed on a per-thread
8 basis.

1 8. The method claim 7 further CHARACTERIZED IN THAT:
2 the acquisition history analysis reduces the number of stubborn sets .

1 9. The method of claim 8 further CHARACTERIZED IN THAT:
2 the concurrent program need not exhibit any substantial lock discipline.

1 10. The method of claim 9 further CHARACTERIZED IN THAT:
2 a set of transactions are determined based upon the lock acquisition history
3 analysis and information about the determined transactions are used to further
4 reduce the number of stubborn sets.

1 11. The method of claim 10 wherein any constraints of the stubborn sets are represented
2 symbolically.

1 12. The method of claim 11 wherein the model of the concurrent program is represented
2 symbolically in circuit-form.




FIG. 1

2/3

1a:	a = 1;	1b:	b = 0;
2a:	lock(lk1);	2b:	lock(lk1);
3a:	lock(lk2);	3b:	lock(lk2);
4a:	y = 1;	4b:	z = 2;
5a:	unlock(lk2);	5b:	unlock(lk1);
6a:	x = 0;	6a:	x = 1;
7a:	unlock(lk1);	7a:	unlock(lk2);

2(a) 2(b)

FIG. 2




FIG. 3