

Computer-Aided Reasoning for Software

CSSE507

Practical Applications of SAT

Today

Past 2 lectures

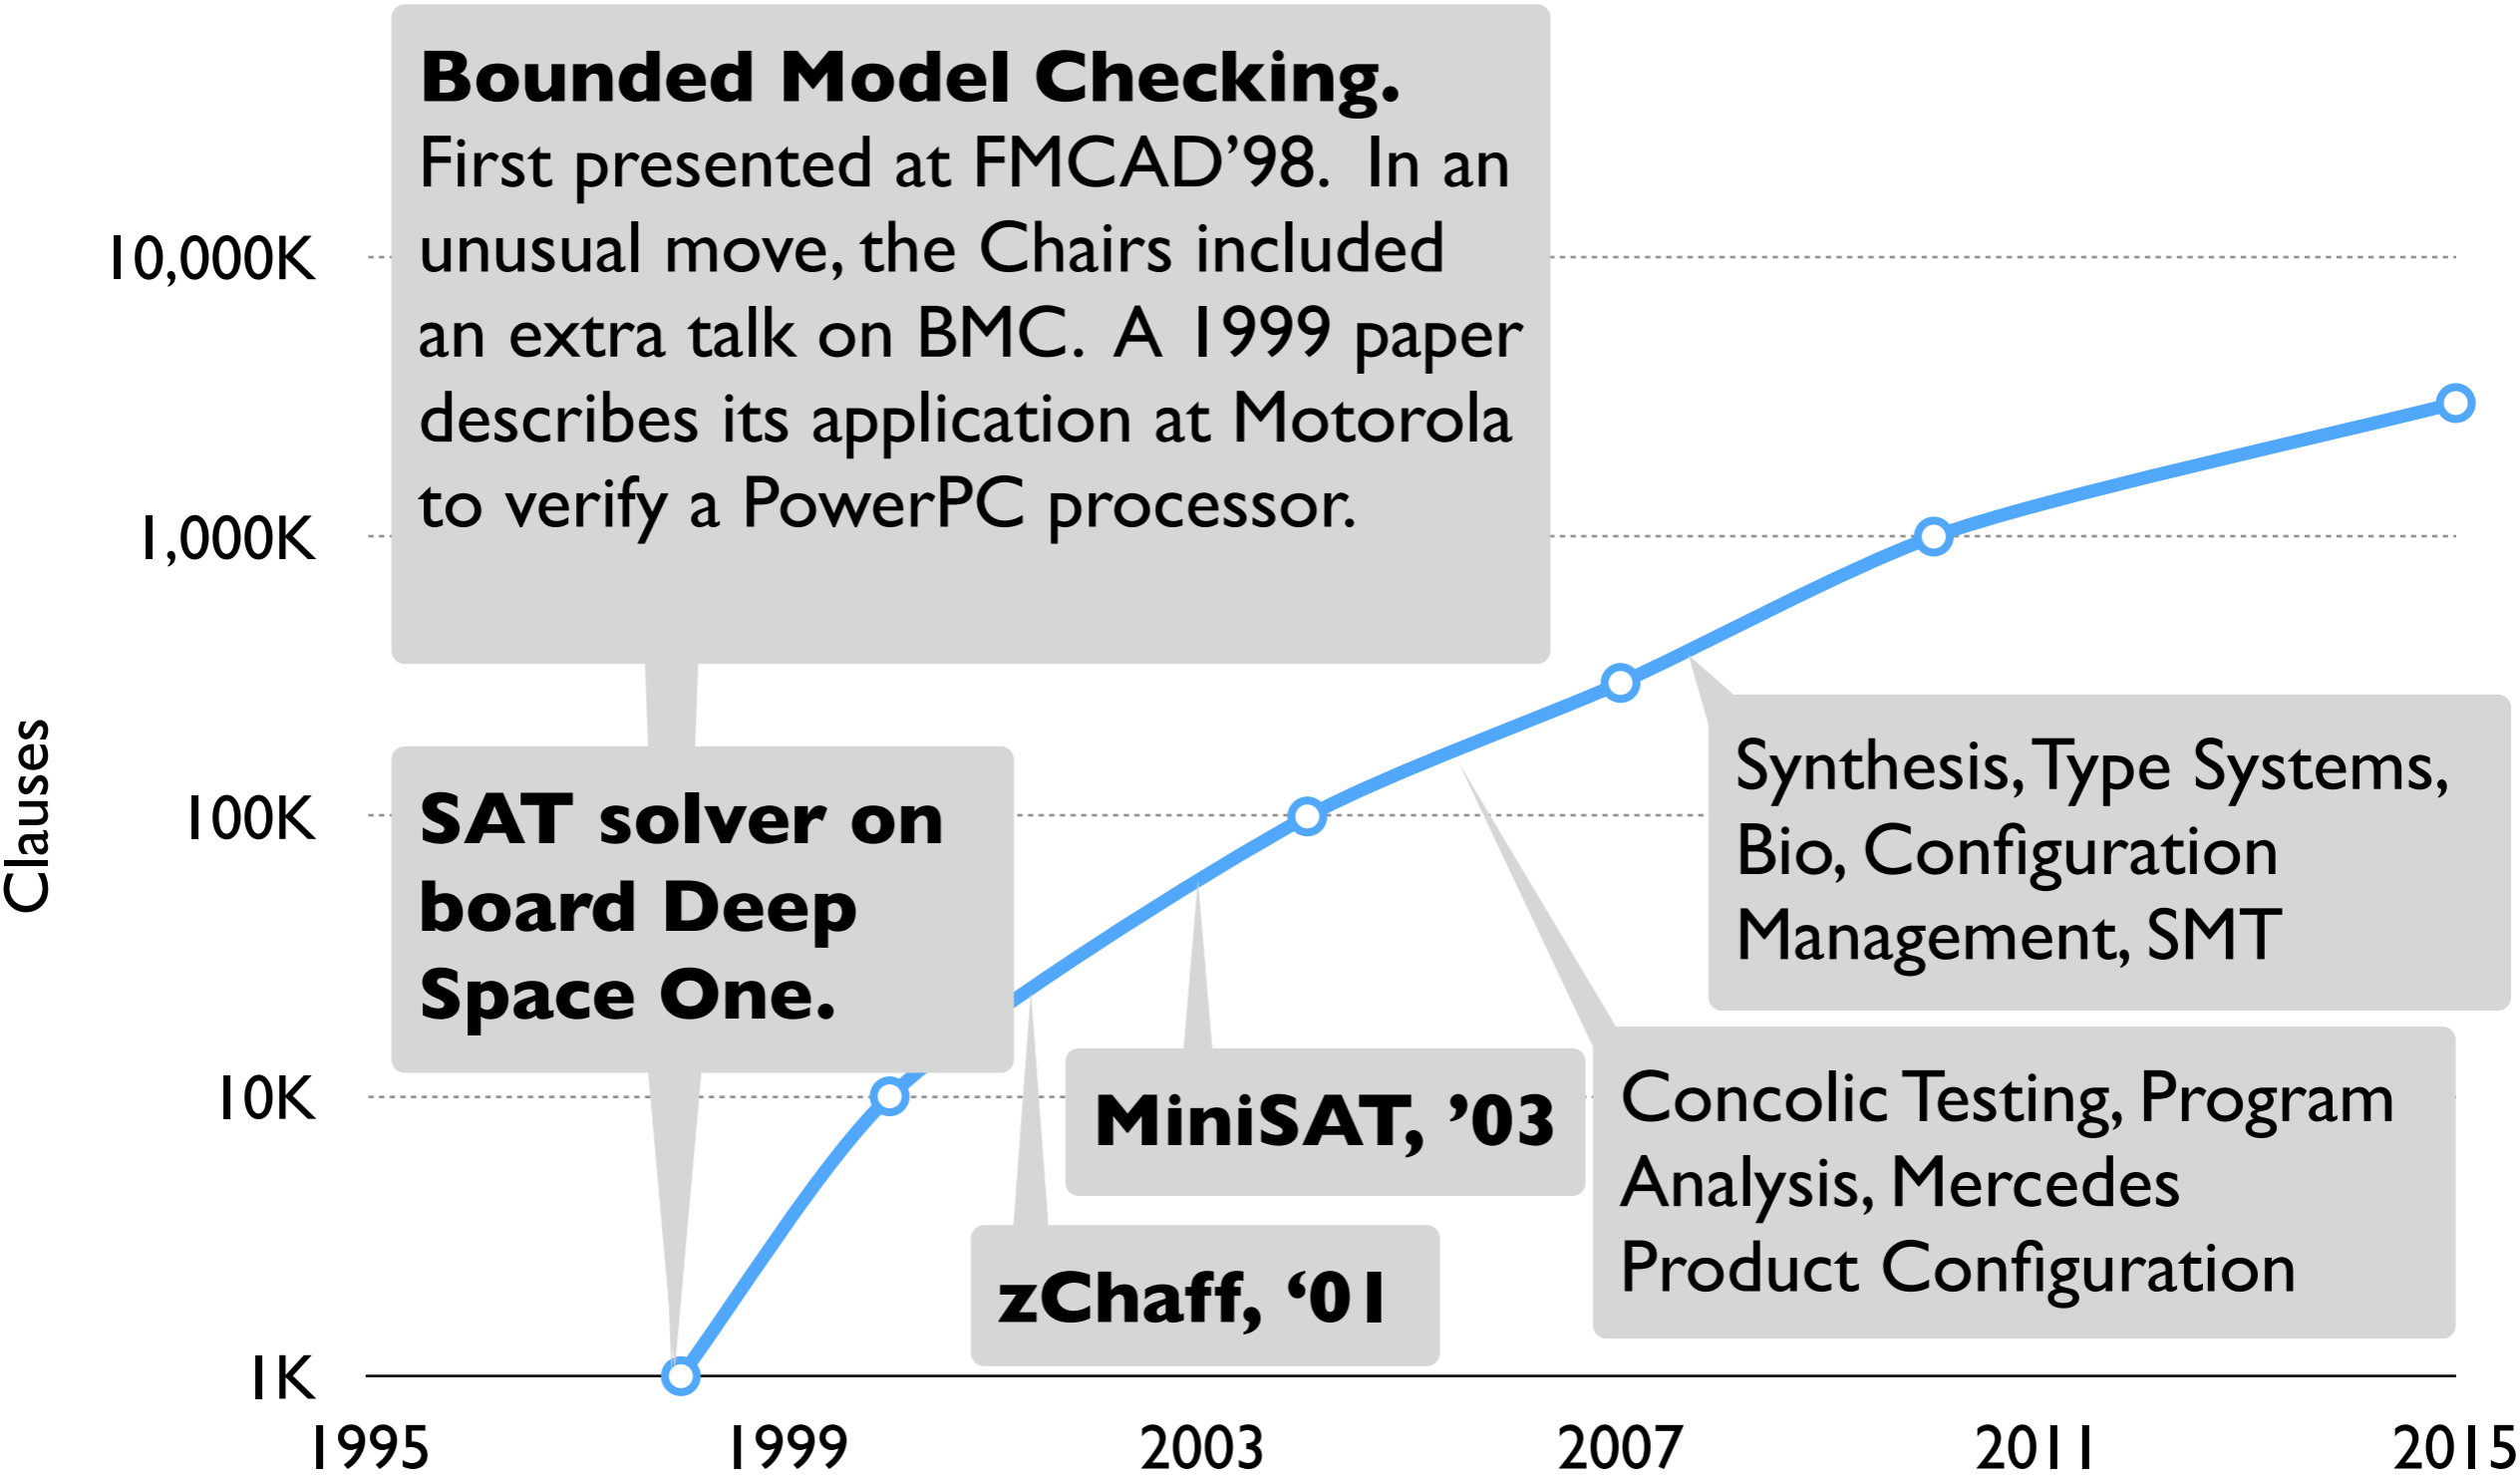
- The theory and mechanics of SAT solving

Today

- Practical applications of SAT
- Variants of the SAT problem
- Motivating the next lecture on SMT



A brief history of SAT solving and applications



Based on a slide from Vijay Ganesh

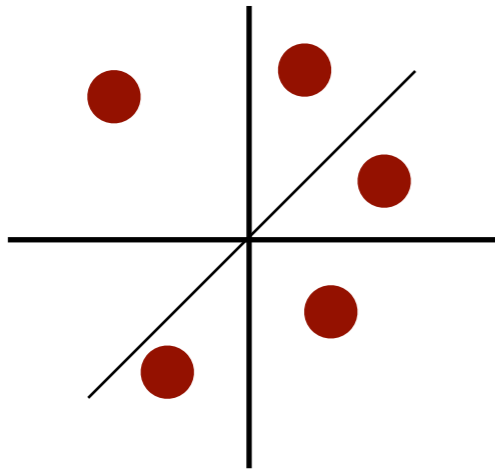
Bounded Model Checking (BMC) & Configuration Management

Bounded Model Checking (in general)

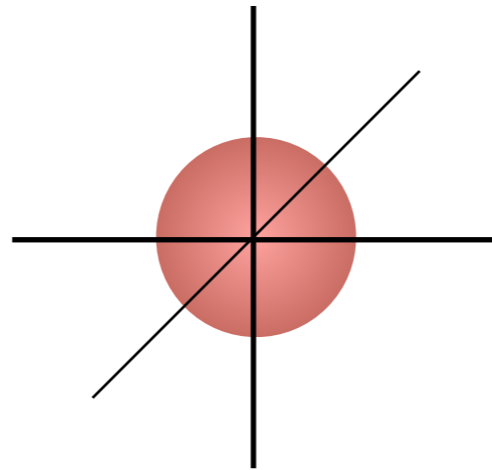
Given a system and a property, BMC checks if the property is satisfied by all executions of the system with $\leq k$ steps, on all inputs of size $\leq n$.

We will focus on **safety properties** (i.e., making sure a bad state, such as an assertion violation, is not reached).

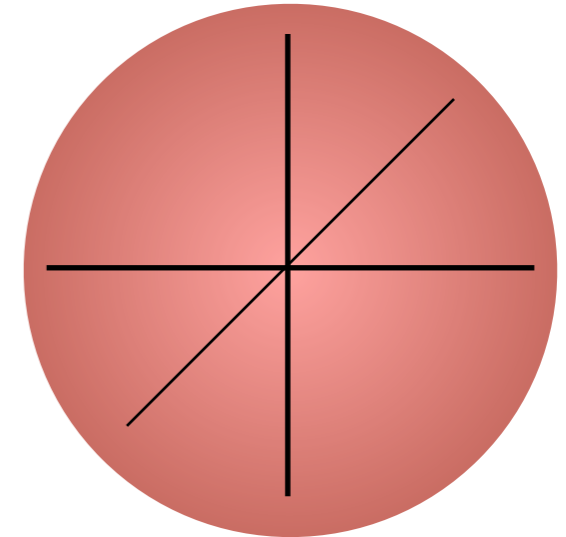
Bounded Model Checking (in general)



Testing: checks a few executions of arbitrary size



BMC: checks all executions of size $\leq k$



Verification: checks all executions of every size

low confidence
low human labor

The **small scope hypothesis**: most bugs can be triggered with small inputs and executions.

high confidence
high human labor

BMC by example



BMC by example

```
int daysToYear(int days) {
    int year = 1980;
    while (days > 365) {
        int oldDays = days;
        if (isLeapYear(year)) {
            if (days > 366) {
                days -= 366;
                year += 1;
            }
        } else {
            days -= 365;
            year += 1;
        }
        assert days < oldDays;
    }
    return year;
}
```

The Zune Bug: on December 31, 2008, all first generation Zune players from Microsoft became unresponsive because of this code. What's wrong?

Infinite loop triggered on the last day of every leap year.

A desired safety property: the value of the days variable decreases in every loop iteration.

BMC step 1 of 4: finitize loops

```
int daysToYear(int days) {
    int year = 1980;
    if (days > 365) {
        int oldDays = days;
        if (isLeapYear(year)) {
            if (days > 366) {
                days -= 366;
                year += 1;
            }
        } else {
            days -= 365;
            year += 1;
        }
        assert days < oldDays;
        assert days <= 365;
    }
    return year;
}
```

- Unwind all loops k times (e.g., $k=1$), and add an **unwinding assertion** at the end.
- If a CEX violates a program assertion, we have found a buggy behavior of length $\leq k$.
- If a CEX violates an unwinding assertion, the program has no buggy behavior of length $\leq k$, but it may have a longer one.
- If there is no CEX, the program is correct for all k !

BMC step 1 of 4: finitize loops & inline calls

```
int daysToYear(int days) {
    int year = 1980;
    if (days > 365) {
        int oldDays = days;
        if (isLeapYear(year)) {
            if (days > 366) {
                days -= 366;
                year += 1;
            }
        } else {
            days -= 365;
            year += 1;
        }
        assert days < oldDays;
        assert days <= 365;
    }
    return year;
}
```

Assume call to isLeapYear is inlined (replaced with the procedure body). We'll keep it for readability.

BMC step 2 of 4: eliminate side effects

```
int daysToYear(int days) {  
    int year = 1980;  
    if (days > 365) {  
        int oldDays = days;  
        if (isLeapYear(year)) {  
            if (days > 366) {  
                days -= 366;  
                year += 1;  
            }  
        } else {  
            days -= 365;  
            year += 1;  
        }  
        assert days < oldDays;  
        assert days <= 365;  
    }  
    return year;  
}
```

BMC step 2 of 4: eliminate side effects

```
int days;  
int year = 1980;  
if (days > 365) {  
    int oldDays = days;  
    if (isLeapYear(year)) {  
        if (days > 366) {  
            days = days - 366;  
            year = year + 1;  
        }  
    } else {  
        days = days - 365;  
        year = year + 1;  
    }  
    assert days < oldDays;  
    assert days <= 365;  
}  
return year;
```

Convert to **Static Single Assignment** (SSA) form:

BMC step 2 of 4: eliminate side effects

```
int days0;
int year0 = 1980;
if (days0 > 365) {
    int oldDays0 = days0;
    if (isLeapYear(year0)) {
        if (days0 > 366) {
            days1 = days0 - 366;
            year1 = year0 + 1;
        }
    } else {
        days3 = days0 - 365;
        year3 = year0 + 1;
    }
    assert days4 < oldDays0;
    assert days4 <= 365;
}
return year5;
```

Convert to **Static Single Assignment** (SSA) form:

- Replace each assignment to a variable v with a definition of a fresh variable v_i .
- Change uses of variables so that they refer to the correct definition (version).

BMC step 2 of 4: eliminate side effects

```
int days0;  
int year0 = 1980;  
boolean g0 = (days0 > 365);  
int oldDays0 = days0;  
boolean g1 = isLeapYear(year0);  
boolean g2 = days0 > 366;  
days1 = days0 - 366;  
year1 = year0 + 1;  
days2 =  $\phi$ (g1 && g2, days1, days0);  
year2 =  $\phi$ (g1 && g2, year1, year0);  
days3 = days0 - 365;  
year3 = year0 + 1;  
days4 =  $\phi$ (g1, days2, days3);  
year4 =  $\phi$ (g1, year2, year3);  
assert days4 < oldDays0;  
assert days4 <= 365;  
year5 =  $\phi$ (g0, year4, year0);  
return year5;
```

Convert to **Static Single Assignment** (SSA) form:

- Replace each assignment to a variable v with a definition of a fresh variable v_i .
- Change uses of variables so that they refer to the correct definition (version).
- Make conditional dependences explicit with gated ϕ nodes.

BMC step 2 of 4: eliminate side effects

```
int days0;  
int year0 = 1980;  
boolean g0 = (days0 > 365);  
int oldDays0 = days0;  
boolean g1 = isLeapYear(year0);  
boolean g2 = days0 > 366;  
days1 = days0 - 366;  
year1 = year0 + 1;  
days2 = φ(g1 && g2, days1, days0);  
year2 = φ(g1 && g2, year1, year0);  
days3 = days0 - 365;  
year3 = year0 + 1;  
days4 = φ(g1, days2, days3);  
year4 = φ(g1, year2, year3);  
assert days4 < oldDays0;  
assert days4 <= 365;  
year5 = φ(g0, year4, year0);  
return year5;
```

```
int days0;  
int year0 = 1980;  
if (days0 > 365) {  
    int oldDays0 = days0;  
    if (isLeapYear(year0)) {  
        if (days0 > 366) {  
            days1 = days0 - 366;  
            year1 = year0 + 1;  
        }  
    } else {  
        days3 = days0 - 365;  
        year3 = year0 + 1;  
    }  
    assert days4 < oldDays0;  
    assert days4 <= 365;  
}  
return year5;
```

BMC step 3 of 4: convert into equations

```
int days0;
int year0 = 1980;
boolean g0 = (days0 > 365);
int oldDays0 = days0;
boolean g1 = isLeapYear(year0);
boolean g2 = days0 > 366;
days1 = days0 - 366;
year1 = year0 + 1;
days2 = φ(g1 && g2, days1, days0);
year2 = φ(g1 && g2, year1, year0);
days3 = days0 - 365;
year3 = year0 + 1;
days4 = φ(g1, days2, days3);
year4 = φ(g1, year2, year3);
assert days4 < oldDays0;
assert days4 <= 365;
year5 = φ(g0, year4, year0);
return year5;
```

We can now read off the equations that encode the program semantics, and the assertions to be checked.

BMC step 3 of 4: convert into equations

```
int year0 = 1980;  
boolean g0 = (days0 > 365);  
int oldDays0 = days0;  
boolean g1 = isLeapYear(year0);  
boolean g2 = days0 > 366;  
days1 = days0 - 366;  
year1 = year0 + 1;  
days2 = φ(g1 && g2, days1, days0);  
year2 = φ(g1 && g2, year1, year0);  
days3 = days0 - 365;  
year3 = year0 + 1;  
days4 = φ(g1, days2, days3);  
year4 = φ(g1, year2, year3);  
assert days4 < oldDays0;  
assert days4 <= 365;
```

We can now read off the equations that encode the program semantics ...

BMC step 3 of 4: convert into equations

```
year0 = 1980;  
g0 = (days0 > 365);  
oldDays0 = days0;  
g1 = isLeapYear(year0);  
g2 = days0 > 366;  
days1 = days0 - 366;  
year1 = year0 + 1;  
days2 =  $\varphi(g_1 \ \&\& \ g_2, \text{days}_1, \text{days}_0)$ ;  
year2 =  $\varphi(g_1 \ \&\& \ g_2, \text{year}_1, \text{year}_0)$ ;  
days3 = days0 - 365;  
year3 = year0 + 1;  
days4 =  $\varphi(g_1, \text{days}_2, \text{days}_3)$ ;  
year4 =  $\varphi(g_1, \text{year}_2, \text{year}_3)$ ;  
assert days4 < oldDays0;  
assert days4 <= 365;
```

We can now read off the equations that encode the program semantics ...

BMC step 3 of 4: convert into equations

```
year0 = 1980 ∧  
g0 = (days0 > 365) ∧  
oldDays0 = days0 ∧  
g1 = isLeapYear(year0) ∧  
g2 = days0 > 366 ∧  
days1 = days0 - 366 ∧  
year1 = year0 + 1 ∧  
days2 = φ(g1 ∧ g2, days1, days0) ∧  
year2 = φ(g1 ∧ g2, year1, year0) ∧  
days3 = days0 - 365 ∧  
year3 = year0 + 1 ∧  
days4 = φ(g1, days2, days3) ∧  
year4 = φ(g1, year2, year3) ∧  
assert days4 < oldDays0;  
assert days4 <= 365;
```

We can now read off the equations that encode the program semantics ...

BMC step 3 of 4: convert into equations

```
year0 = 1980 ∧  
g0 = (days0 > 365) ∧  
oldDays0 = days0 ∧  
g1 = isLeapYear(year0) ∧  
g2 = days0 > 366 ∧  
days1 = days0 - 366 ∧  
year1 = year0 + 1 ∧  
days2 = ite(g1 ∧ g2, days1, days0) ∧  
year2 = ite(g1 ∧ g2, year1, year0) ∧  
days3 = days0 - 365 ∧  
year3 = year0 + 1 ∧  
days4 = ite(g1, days2, days3) ∧  
year4 = ite(g1, year2, year3) ∧  
assert days4 < oldDays0;  
assert days4 <= 365;
```

We can now read off the equations that encode the program semantics ...

BMC step 3 of 4: convert into equations

$$\begin{aligned} & \text{year}_0 = 1980 \wedge \\ & g_0 = (\text{days}_0 > 365) \wedge \\ & \text{oldDays}_0 = \text{days}_0 \wedge \\ & g_1 = \text{isLeapYear}(\text{year}_0) \wedge \\ & g_2 = \text{days}_0 > 366 \wedge \\ & \text{days}_1 = \text{days}_0 - 366 \wedge \\ & \text{year}_1 = \text{year}_0 + 1 \wedge \\ & \text{days}_2 = \text{ite}(g_1 \wedge g_2, \text{days}_1, \text{days}_0) \wedge \\ & \text{year}_2 = \text{ite}(g_1 \wedge g_2, \text{year}_1, \text{year}_0) \wedge \\ & \text{days}_3 = \text{days}_0 - 365 \wedge \\ & \text{year}_3 = \text{year}_0 + 1 \wedge \\ & \text{days}_4 = \text{ite}(g_1, \text{days}_2, \text{days}_3) \wedge \\ & \text{year}_4 = \text{ite}(g_1, \text{year}_2, \text{year}_3) \wedge \\ & (\neg(\text{days}_4 < \text{oldDays}_0) \vee \\ & \quad \neg(\text{days}_4 \leq 365)) \end{aligned}$$

We can now read off the equations that encode the program semantics, and the assertions to be checked.

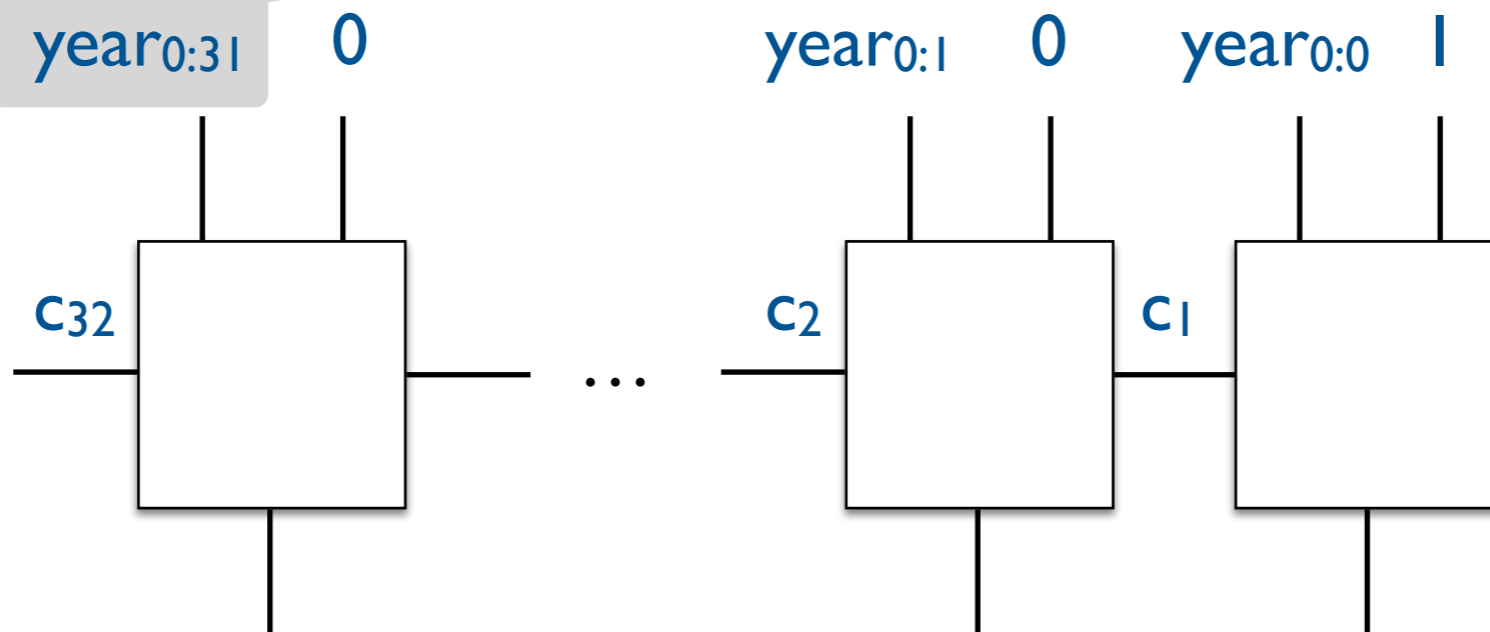
A solution to this formula is a sound *counterexample*: an interpretation for all logical variables that satisfies the program semantics (for up to k unwindings) but violates at least one of the assertions.

BMC step 4 of 4: convert into CNF

$$\text{year}_1 = \text{year}_0 + 1$$

$$\text{year}_0 = \underset{31 \ 30 \ 29}{000} \dots \underset{2 \ 1 \ 0}{000}$$

Represent numbers as arrays of bits.
Use one boolean variable per bit for each number.



Construct an adder circuit for $\text{year}_0 + 1$.

$$\text{year}_{1:31} \leftrightarrow s_{31} \wedge \dots \wedge \text{year}_{1:1} \leftrightarrow s_1 \wedge s_0 \leftrightarrow \text{year}_{1:0}$$

Introduce new clauses to constrain bits in year_1 to match bits in the sum.

BMC counterexample for k=1

```
int daysToYear(int days) {  
    int year = 1980;  
    while (days > 365) {  
        if (isLeapYear(year)) {  
            if (days > 366) {  
                days -= 366;  
                year += 1;  
            }  
        } else {  
            days -= 365;  
            year += 1;  
        }  
    }  
    return year;  
}
```

days = 366

Bounded Model Checking (BMC) & Configuration Management

Configuration Management

Given a configuration, consisting of a set of components, their dependencies, and conflicts:

- Decide if a new component can be added to the configuration.
- Add the component while optimizing some linear function.
- If the component cannot be added, find a way to add it by removing as few conflicting components from the current configuration as possible.

maven

eclipse

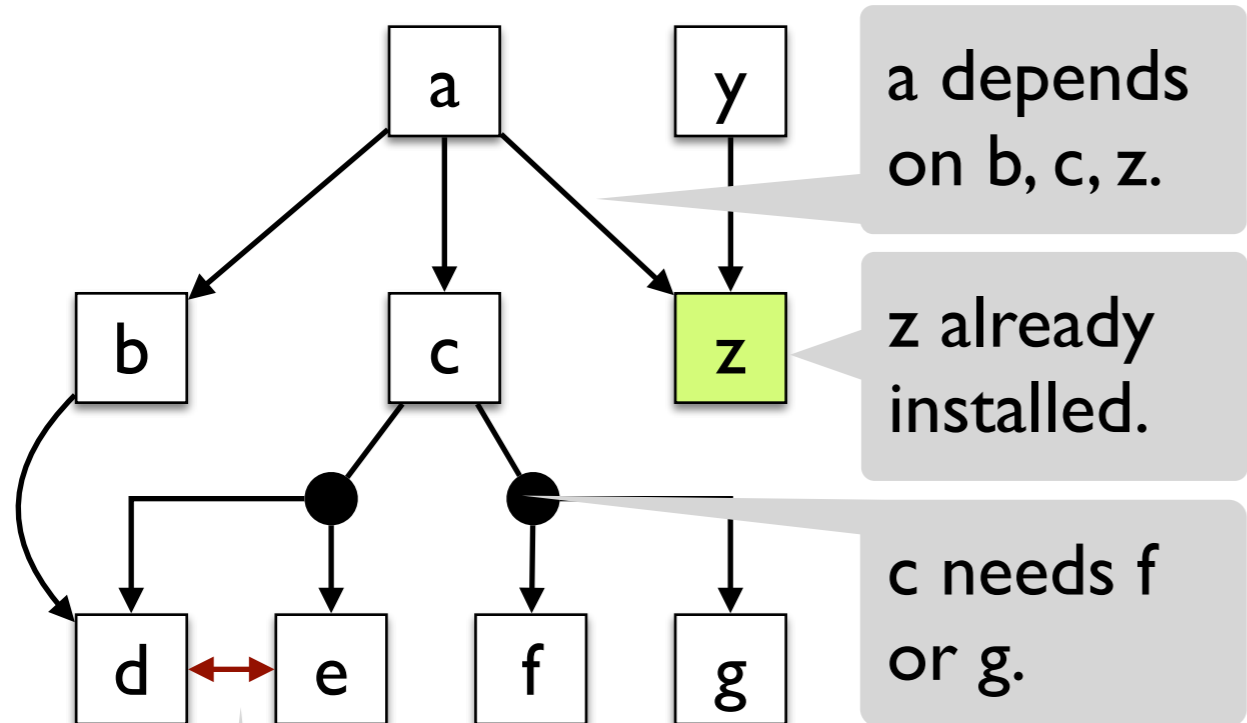


SAT

Pseudo-Boolean Constraints

Partial (Weighted) MaxSAT

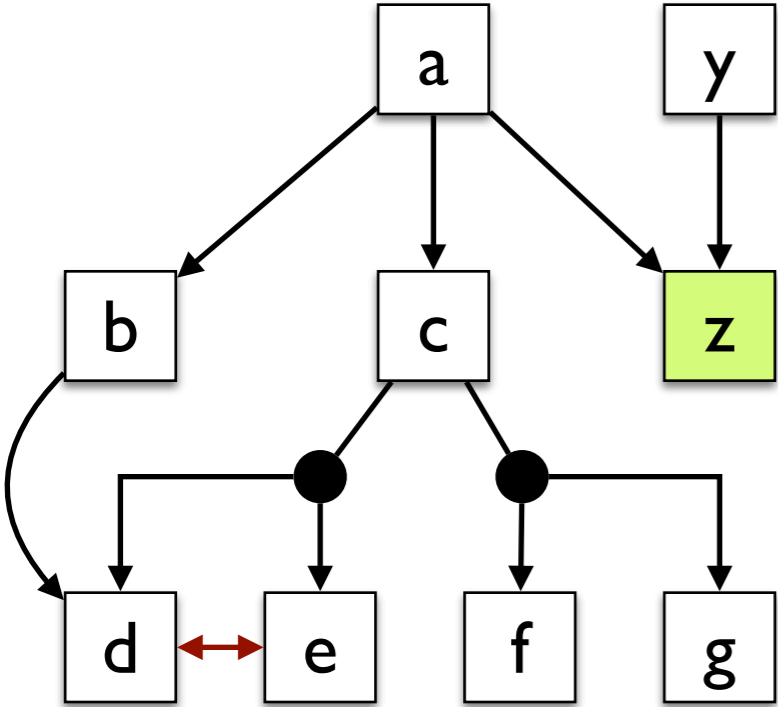
Deciding if a component can be installed



To install a, CNF constraints are:

$$\begin{aligned} &(\neg a \vee b) \wedge (\neg a \vee c) \wedge (\neg a \vee z) \wedge \\ &(\neg b \vee d) \wedge \\ &(\neg c \vee d \vee e) \wedge (\neg c \vee f \vee g) \wedge \\ &(\neg d \vee \neg e) \wedge \\ &(\neg y \vee z) \wedge \\ &a \wedge z \end{aligned}$$

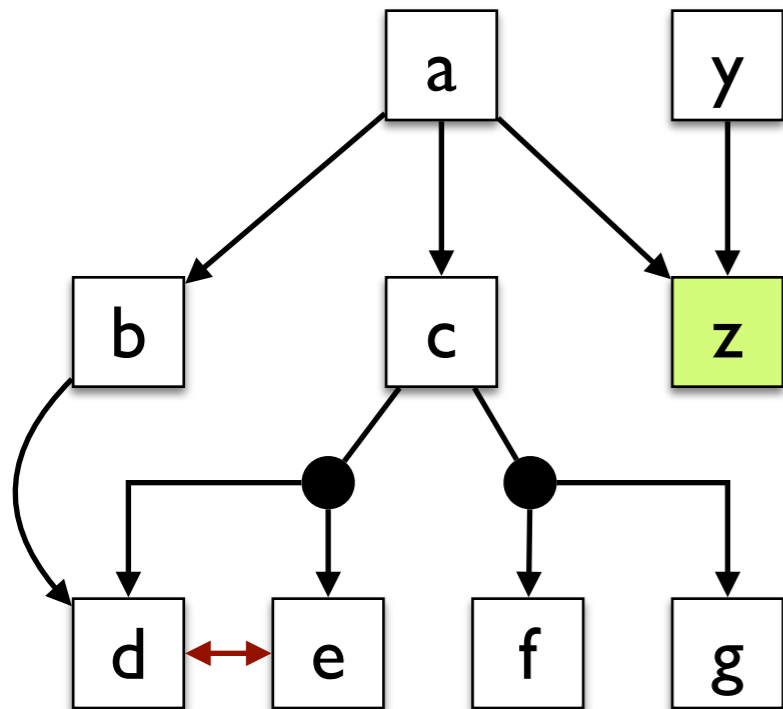
Optimal installation



Assume f and g are 5MB and 2MB each, and all other components are 1MB. How to install a, while minimizing total size?

$$\begin{aligned}
 &(\neg a \vee b) \wedge (\neg a \vee c) \wedge (\neg a \vee z) \wedge \\
 &(\neg b \vee d) \wedge \\
 &(\neg c \vee d \vee e) \wedge (\neg c \vee f \vee g) \wedge \\
 &(\neg d \vee \neg e) \wedge \\
 &(\neg y \vee z) \wedge \\
 &a \wedge z
 \end{aligned}$$

Optimal installation



Assume f and g are 5MB and 2MB each, and all other components are 1MB. How to install a, while minimizing total size?

Pseudo-boolean solvers accept a linear function to minimize, in addition to a (weighted) CNF.

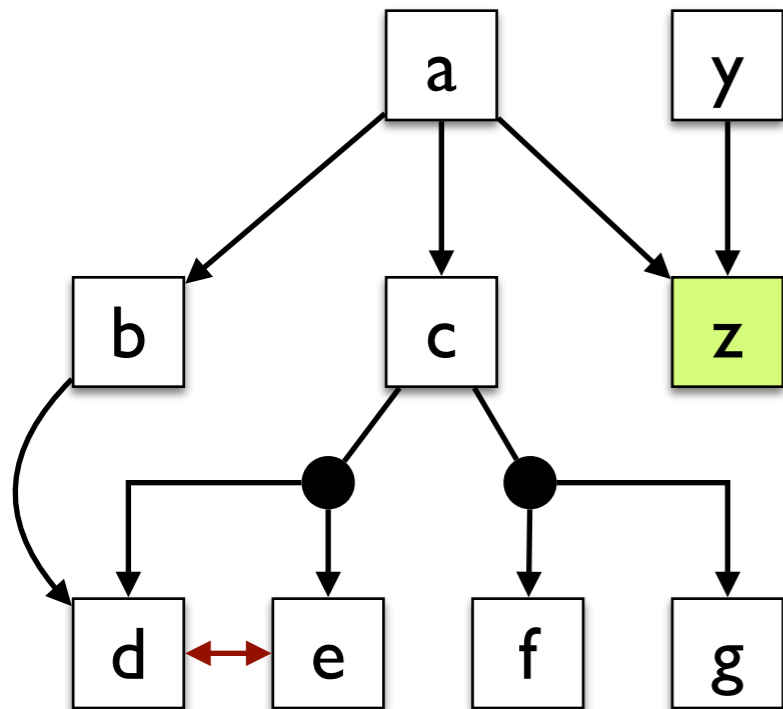
$$\mathbf{min} \ c_1x_1 + \dots + c_nx_n$$

$$a_{11}x_1 + \dots + a_{1n}x_n \geq b_1 \wedge \dots \wedge$$

$$a_{k1}x_1 + \dots + a_{kn}x_n \geq b_k$$

$$\begin{aligned} &(\neg a \vee b) \wedge (\neg a \vee c) \wedge (\neg a \vee z) \wedge \\ &(\neg b \vee d) \wedge \\ &(\neg c \vee d \vee e) \wedge (\neg c \vee f \vee g) \wedge \\ &(\neg d \vee \neg e) \wedge \\ &(\neg y \vee z) \wedge \\ &a \wedge z \end{aligned}$$

Optimal installation

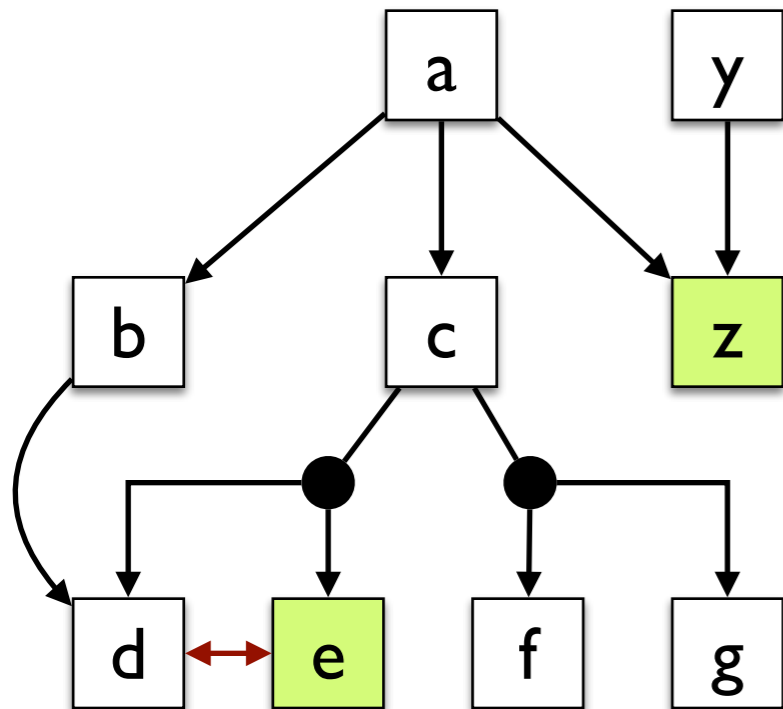


Assume f and g are 5MB and 2MB each, and all other components are 1MB. How to install a, while minimizing total size?

Pseudo-boolean solvers accept a linear function to minimize, in addition to a (weighted) CNF.

$$\begin{aligned} \mathbf{min} \quad & a + b + c + d + e + 5f + 2g + y + 0z \\ & (-a + b \geq 0) \wedge (-a + c \geq 0) \wedge (-a + z \geq 0) \wedge \\ & (-b + d \geq 0) \wedge \\ & (-c + d + e \geq 0) \wedge (-c + f + g \geq 0) \wedge \\ & (-d + -e \geq -1) \wedge \\ & (-y + z \geq 0) \wedge \\ & (a \geq 1) \wedge (z \geq 1) \end{aligned}$$

Installation in the presence of conflicts



a cannot be installed because it requires b, which requires d, which conflicts with e.

Partial MaxSAT solver takes as input a set of **hard** clauses and a set of **soft** clauses, and it produces an assignment that satisfies all hard clauses and the greatest number of soft clauses.

To install a, while minimizing the number of removed components, Partial MaxSAT constraints are:

hard: $(\neg a \vee b) \wedge (\neg a \vee c) \wedge (\neg a \vee z) \wedge$
 $(\neg b \vee d) \wedge$
 $(\neg c \vee d \vee e) \wedge (\neg c \vee f \vee g) \wedge$
 $(\neg d \vee \neg e) \wedge (\neg y \vee z) \wedge a$

soft: $e \wedge z$

Summary

Today

- SAT solvers have been used successfully in many applications & domains
- But reducing problems to SAT is a lot like programming in assembly ...
- We need higher-level logics!

Next lecture

- On to richer logics: introduction to Satisfiability Modulo Theories (SMT)