Disjoint Conditions

CS 536: Science of Programming, Spring 2023

A. Why?

- Combining arbitrary threads are in parallel can yield programs that have surprisingly different results from how each thread works when run sequentially.

- When threads don't interfere with each other's work, we can combine them into a parallel program without worrying about these strange and unexpected kinds of behavior.

- Simply having disjoint parallel programs for threads isn't sufficient to avoid interference problems.

- What's needed are disjoint conditions, which ensure that no thread can interfere with the conditions of another thread.

B. Objectives

After this class, you should know

- What disjoint conditions are, why we need them, and how to recognize them.

- What the disjoint parallelism rule for disjoint parallel programs with disjoint conditions allows.

- That maybe not all apparent interference actually causes problems.

C. Parallelism Rule for Parallel Programs?

- The ***sequentialization proof rule*** for DPPs lets us reason about DPPs, which is nice, but it has two major flaws.

- First, It requires disjoint parallel programs, and we know disjoint parallelism is a strong constraint on what programs we can write.

- Second, to use it sequentialization, we need many intermediate conditions. To prove $\{p\}[S_1 || … || S_n]\{q\}$, we need to prove $\{p\}S_1; … ; S_n\{q\}$, which (if we use wp) means finding a sequence of preconditions $q_1, …, q_n$ and proving $\{p\}\{q_n\}S_1; \{q_{n-1}\}S_2; \{q_{n-2}\} … \{q_1\}S_n\{q\}$. The proofs of $q_1, q_2, … q_n$ can get increasingly complicated because each q_i can depend on all the threads and conditions to its right.

- ***Parallelism Rule***: Ideally, we'd like to take advantage not only of parallel execution of programs but also parallel writing of programs. Can we have n programmers write n sequential threads and then combine them into a parallel program? This behavior describes a ***parallelism rule***. For two threads, the form would be as follows (it generalizes to n threads).

1. $\{p_1\}S_1\{q_1\}$
2. $\{p_2\}S_2\{q_2\}$
3. $\{p_1 \wedge p_2\}[S_1 \mid\mid S_2]\{q_1 \wedge q_2\}$ by (some kind of) parallelism 1, 2

In proof outline format, we write this as $\{p_1 \wedge p_2\}[\{p_1\}S_1\{q_1\} \mid\mid \{p_2\}S_2\{q_2\}]\{q_1 \wedge q_2\}$.

- ***Example 1***: In proof format, we write
    1. $\{x \geq 0\}z := x\{z \geq 0\}$
    2. $\{y \leq 0\}w := -y\{w \geq 0\}$
    3. $\{x \geq 0 \wedge y \leq 0\}[z := x \mid\mid w := -y]\{z \geq 0 \wedge w \geq 0\}$ by ??? parallelism 1, 2

    The proof outline form is

    $\{x \geq 0 \wedge y \leq 0\}$
    $[\{x \geq 0\}z := x\{z \geq 0\}$
    $\mid\mid \{y \leq 0\}w := -y\{w \geq 0\}$
    $]\{z \geq 0 \wedge w \geq 0\}$

- It's nice when this pattern works, but we'll see examples where the pattern produces incorrect programs. To get correctness, we have to impose side conditions that limit how we can combine threads. Altogether, we'll look at a couple of different kinds of side conditions, which work in different situations.

- ***Example 2***: Here are some outlines where using a parallelism rule works. Outlines *(b)* and *(c)* use postcondition weakening to get a final postcondition.
    a. $\{x \geq 0 \wedge y \leq 0\}[\{x \geq 0\}z := x\{z \geq 0\} \mid\mid \{y \leq 0\}w := -y\{w \geq 0\}]\{z \geq 0 \wedge w \geq 0\}$
    b. $\{z = 0\}[\{z = 0\}x := z - 1 \ \{x \leq z = 0\} \mid\mid \{z = 0\}y := z\{y = z = 0\}]\{x \leq z = 0 \wedge y = z = 0\}$
        $\{x \leq y = z = 0\}$
    c. $\{T\}[\{T\}a := x + 1\{a = x + 1\} \mid\mid \{T\}b := x + 2\{b = x + 2\}]\{a = x + 1 \wedge b = x + 2\}\{a + 1 = b\}$
    d. $\{x = y = z = c\}$
        $[\{x = c\}x := x^2\{x = c^2\}$
        $\mid\mid \{y = c\}y := y^2\{y = c^2\}$
        $\mid\mid \{z = c\}z := (z - d) * (z + d)\{z = c^2 - d^2\}$
        $]$
        $\{x = c^2 \wedge y = c^2 \wedge z = c^2 - d^2\}$
        $\{x = y = z + d^2\}$
    e.

- ***Example 3:*** Here are some outlines where using a parallelism rule fails. Outline *(a)* is incorrect because it wants to end with $x = 1 \wedge x = y = 0$ Outline *(b)* doesn't always leave $x > y$ (e.g., if we start with $x = y = 2$). Outline *(c)* never works.
    a. $\{x = 0\}[\{x = 0\}x := 1\{x = 1\} \mid\mid \{x = 0\}y := 0\{x = y = 0\}]\{x = 1 \wedge x = y = 0\}$

b. $\{x \geq y \land y = z\}[\{x \geq y\}x := x + 1\{x > y\} \;||\; \{y = z\}y := y * 2; z := z * 2\{y = z\}]\{x > y \land y = z\}$

c. $\{T\}[\{T\}[x := 2 + 2\{x = 2 + 2\} \;||\; \{T\}x := 5\{x = 5\}]\{x = 2 + 2 \land x = 5\}\{2 + 2 = 5\}$

D. Disjoint Conditions

- Why do the proof outlines in **Example 3** fail? Since outline *(c)* **should** fail because the semantics don't support it. But outlines *(a)* and *(b)* fail even though they have disjoint parallel programs.

- Looking more closely, we see outline *(a)* fails because thread 1 sets x to *1*, which invalidates the precondition $x = 0$ needed by thread 2. Similarly, outline *(b)* fails because thread 2 modifies y, which might invalidate the $x \geq y$ and $x > y$ pre- and postconditions in thread1.

- It's not enough for outlines *(a)* and *(b)* to have disjoint parallel programs. I.e., it's not enough to guarantee that threads don't change each others' states. For correctness, we need each thread to **not modify** the variables that appear in the conditions of other threads.

- **Definition**: $Free(p, ..., q)$ is the set of variables that appear free in any of the predicates $p, ..., q$. (Recall that a variable can have both free and bound occurrences, "bound" meaning "in the scope of a quantifier" for it. If a variable has at least one free occurrence in a predicate, then it is free in that predicate, regardless of whether or not there are bound occurrences.)

- **Definition**: With $\{p_1\}S_1\{q_1\}$ and $\{p_2\}S_2\{q_2\}$, the first triple **interferes with the conditions** of the second triple if $Change(S_1) \cap Free(p_2, q_2) \neq \varnothing$. We may use **disjoint from the conditions** as a synonym for "not interfering with." The two triples have **disjoint conditions** if neither interferes with the conditions of the other. (I.e., $Change(S_1) \cap Free(p_2, q_2) = \varnothing$ and $Change(S_2) \cap Free(p_1, q_1) = \varnothing$.)

- **Definition**: A parallel program outline $\{p\}[\{p_1\}S_1\{q_1\} \;||\; ... \;||\; \{p_n\}S_n\{q_n\}]\{q\}$ has **disjoint conditions** if its threads have pairwise disjoint conditions. I.e., threads i and j have disjoint conditions, for all indexes i and j (with $i \neq j$).

- **Example 4:** All of the outlines in **Example 2** have pairwise disjoint programs and conditions. To show this, we can add a Disjoint Conditions column to the table we used for checking for disjoint programs. (Note that not all the sequential thread outlines are full outlines.) [Note: In the table below, I don't think it makes any difference if you write (e.g.) $w\,y$ or w, y or $\{w, y\}$.]

a. $\{x \geq 0 \land y \leq 0\}[\{x \geq 0\}z := x\{z \geq 0\} \;||\; \{y \leq 0\}w := -y\{w \geq 0\}]\{z \geq 0 \land w \geq 0\}$

i	j	*Change i*	*Vars j*	*Free j*	*Disj Pgm*	*Disj Cond*
1	2	z	$w\,y$	$w\,y$	Y	Y
2	1	w	$x\,z$	$x\,z$	Y	Y

b. *{ z = 0 } [{ z = 0 } x := z + 1 { x ≤ z = 0 } || { z = 0 } y := z { y = z = 0 }] { x ≤ z = 0 ∧ y = z = 0 }*
    { x ≤ y = z = 0 }

i	j	Change i	Vars j	Free j	Disj Pgm	Disj Cond
1	2	x	y z	y z	Y	Y
2	1	y	x z	x z	Y	Y

c. *{ T } [{ T } a := x + 1 { a = x + 1 } || { T } b := x + 2 { b = x + 2 }] { a = x + 1 ∧ b = x + 2 } { a + 1 = b }*

i	j	Change i	Vars j	Free j	Disj Pgm	Disj Cond
1	2	a	b x	b x	Y	Y
2	1	b	a x	a x	Y	Y

d. *{ x = y = z = c }*
    [{ x = c } x := x² { x = c² }
    || { y = c } y := y² { y = c² }
    *|| { z = c } z := (z – d) * (z + d) { z = c² – d² }*
    ] { x = c² ∧ y = c² ∧ z = c² – d² }
    { x = y = z + d² }

i	j	Change i	Vars j	Free j	Disj Pgm	Disj Cond
1	2	x	y	c y	Y	Y
1	3	x	d z	c d z	Y	Y
2	1	y	x	c x	Y	Y
2	3	y	d z	c d z	Y	Y
3	1	z	x	c x	Y	Y
3	2	z	y	c y	Y	Y

- **Example 5:** All of the outlines in **Example 4** lack pairwise disjoint conditions. Outlines *(a)* and *(b)* have disjoint programs but outline *(c)* does not.

a. *{ x = 0 } [{ x = 0 } x := 1 { x = 1 } || { x = 0 } y := 0 { x = y }] { x = 1 ∧ x = y }*

i	j	Change i	Vars j	Free j	Disj Pgm	Disj Cond
1	2	x	y	x y	Y	N
2	1	y	x	x	Y	Y

b. *{ x ≤ y ∧ y = z } [{ x ≤ y } x := x − 1 { x < y } || { y = z } y := y * 2 ; z := z * 2 { y = z }] { x < y ∧ y = z }*

i	j	Change i	Vars j	Free j	Disj Pgm	Disj Cond
1	2	x	y z	y z	Y	Y
2	1	y z	x	x y	Y	N

c. *{ T } [{ T } x := 2 + 2 { x = 2 + 2 } || { T } x := 5 { x = 5 }] { x = 2 + 2 ∧ x = 5 } { 2 + 2 = 5 }*

i	j	Change i	Vars j	Free j	Disj Pgm	Disj Cond
1	2	x	x	x	N	N
2	1	x	x	x	N	N

E. Disjoint Conditions; Disjoint Parallelism Rule

- If two outlines have disjoint programs and conditions, then their evaluations can be arbitrarily interleaved without fear. Having disjoint programs guarantees that no thread modifies the parts of the state used by the calculations of other threads. Having disjoint conditions guarantees that no thread modifies the parts of the state used to determine satisfaction of other threads' conditions.

- Because of this lack of runtime interference, we can use a a parallelism rule for disjoint parallel programs with disjoint conditions. The rule has the form of a parallelism rule (the parallel program uses the conjunctions of its threads' preconditions and postconditions as its precondition and postcondition). Having disjoint programs and conditions is the side condition that guarantees correctness of the rule.

DisjointParallelism Rule

$1.$ $\{p_1\} S_1 \{q_1\}$

$2.$ $\{p_2\} S_2 \{q_2\}$

...

n $\{p_n\} S_n \{q_n\}$

$n+1$ $\{p_1 \wedge p_2 \wedge ... \wedge p_n\} [S_1 \;||\; ... \;||\; S_n] \{q_1 \wedge q_2 \wedge ... \wedge q_n\}$ 　Disjoint Parallelism, $1, 2, ...,$ n

where threads $1, 2, ..., n$ are pairwise disjoint programs with pairwise disjoint conditions

- A reminder: The pairwise tests check thread i for interference with thread j, for all $i, j \in \{1, 2, ..., n\}$ where $i \neq j$.

Example 6: The outlines from Example 2 all have disjoint programs and conditions, so all of them can be proved using the disjoint parallelism rule. The presentations in Example 2 illustrate the uses of the rule. Just to be different, we can write them in nonlinear form. The comments marked * indicate parts of outline that are partial but not full. E.g., $\{y \leq 0\} w := -y \{w \geq 0\}$ is not a full outline; $\{y \leq 0\} w := -y \{y \leq 0 \wedge w = -y\} \{w \geq 0\}$ is; so is $\{y \leq 0\} \{-y \geq 0\} w := -y \{w \geq 0\}$.

a. $\{x \geq 0 \wedge y \leq 0\}$

$[\{x \geq 0\} z := x \{z \geq 0\}$

$|| \{y \leq 0\} w := -y \{w \geq 0\}]$ 　　(*) (Not a full outline for the sequential thread)

$\{z \geq 0 \wedge w \geq 0\}$

b. $\{z = 0\}$ 　　　　　　　　　//I.e., $z = 0 \wedge z = 0$

$[\{z = 0\} x := z - 1 \; \{x \leq z = 0\}$ 　(*)

$|| \{z = 0\} y := z \{y = z = 0\}]$

$\{x \leq z = 0 \wedge y = z = 0\}$

$\{x \leq y = z = 0\}$

c. $\{T\}$

$[\{T\} a := x + 1 \{a = x + 1\}$

$|| \{T\} b := x + 2 \{b = x + 2\}]$

$\{a = x + 1 \wedge b = x + 2\}$

$\{a + 1 = b\}$

d. $\{x = y = z = c\}$
$\{x = c \land y = c \land z = c\}$
$[\{x = c\} x := x^2 \{x = c^2\}$ (*)
$\|\ \{y = c\} y := y^2 \{y = c^2\}$ (*)
$\|\ \{z = c\} z := (z - d) * (z + d) \{z = c^2 - d^2\}$ (*)
$]$
$\{x = c^2 \land y = c^2 \land z = c^2 - d^2\}$
$\{x = y = z + d^2\}$

F. Removing Interference of Conditions

- Sometimes, interference with conditions can be removed by adding logical variables.

***Example* 7**: All of the outlines from Example 3 have interfering conditions, but we can remove the interference for outlines *(a)* and *(b)*. Since *(c)* doesn't have disjoint programs, modifying its conditions won't help.

a. Original: $\{x = 0\} [\{x = 0\} x := 1 \{x = 1\} \| \{x = 0\} y := 0 \{x = y\}] \{x = 1 \land x = y\}$.
The interference is that thread 1 setting x to *1* interferes with the x in $x = y$ in thread2.
Said the other way, the x in $x = y$ refers to the value of x before $x := 1$. Introducing a logical variable for the initial value of x takes care of that. We also have to change the final post-condition to get one that's true.

Without interference: (note the outline for thread 2 is not full).
$\{x = x_0 \land x = 0\}$
$[\{x = x_0 \land x = 0\} x := 1 \{x = 1 \land x_0 = 0\}$
$\|\ \{x = x_0 \land x = 0\} y := 0 \{x_0 = y\}]$
$\{x = 1 \land x - 1 = y\}$

b. Original:
$\{x \leq y \land y = z\} [\{x \leq y\} x := x - 1 \{x < y\} \| \{y = z\} y := y * 2 ; z := z * 2 \{y = z\}] \{x < y \land y = z\}$.
The interference is $y := y * 2$ in thread 2 versus $x \leq y$ and $x < y$ in thread 1. Again, introducing logical variables helps.

Without interference:

$$\{ x = x_0 \wedge y = y_0 \wedge z = z_0 \wedge x \le y \wedge y = z \}$$
$$[\{ x = x_0 \wedge x \le y_0 \}$$
$$x := x - 1$$
$$\{ x_0 \le y_0 \wedge x = x_0 - 1 \}$$
$$|| \{ y = y_0 \wedge z = z_0 \wedge y = z \}$$
$$y := y * 2 ; z := z * 2$$
$$\{ y_0 = z_0 \wedge y = y_0 * 2 \wedge z = z_0 * 2 \}$$
$$]$$
$$\{ (x_0 \le y_0 \wedge x = x_0 - 1) \wedge (y_0 = z_0 \wedge y = y_0 * 2 \wedge z = z_0 * 2) \}$$
$$\{ x < y \div 2 \wedge y = z \}$$

- For another example of removing interference, let's look at a modified version of Example 6.d where we have the initial precondition $x = y = z = c$ but no annotation for the sequential threads. In that case, it would be natural to use $x = y = z = c$ as the precondition for all three threads:

$$\{ x = y = z = c \} \quad \text{(partial outline)}$$
$$[\{ x = y = z = c \} x := x^2 \{ x = c^2 \wedge y = z = c \}$$
$$|| \{ x = y = z = c \} y := y^2 \{ y = c^2 \wedge x = z = c \}$$
$$|| \{ x = y = z = c \} z := (z - d) * (z + d) \{ z = c^2 - d^2 \wedge x = y = c \}$$
$$]$$
$$\{ x = c^2 \wedge y = z = c \wedge y = c^2 \wedge x = z = c \wedge z = c^2 - d^2 \wedge x = y = c \}$$
$$\{ x = y = z + d^2 \}$$

- But this has interference all over the place. E.g., $x := x^2$ interferes with all the conditions $x = \ldots$ in the other threads, and similarly for the assignments to y and z.

- A possible fix that works very well for this example is to look at the sequential annotations and drop everything that isn't directly related to the sequential thread's programs. E.g., with $\{ x = y = z = c \} x := x^2 \{ x = c^2 \wedge y = z = c \}$, we need x and c but not y or z. Removing y and z gives us $\{ x = c \} x := x^2 \{ x = c^2 \}$. Similarly, we can remove x and z from the thread for $y := y^2$ and then x and y from the thread for $z := (z - d) * (z + d)$. The result is exactly Example 6.d:

$$\{ x = y = z = c \} \qquad\qquad \text{(partial outline)}$$
$$[\{ x = c \} x := x^2 \{ x = c^2 \}$$
$$|| \{ y = c \} y := y^2 \{ y = c^2 \}$$
$$|| \{ z = c \} z := (z - d) * (z + d) \{ z = c^2 - d^2 \}$$
$$]$$
$$\{ x = c^2 \wedge y = c^2 \wedge z = c^2 - d^2 \}$$
$$\{ x = y = z + d^2 \}$$

- Another analysis that adds x_0, y_0, and z_0 turns $\{ x = y = z = c \} x := x^2 \{ x = c^2 \wedge y = z = c \}$ into $\{ x = x_0 \wedge x = c \} x := x^2 \{ x_0 = c \wedge x = x_0^2 \}$ and similarly for y and z:

$\{ x = y = z = c \}$ (partial outline)

$[\{ x = x_0 \wedge x = c \}\, x := x^2\, \{ x_0 = c \wedge x = x_0{}^2 \}$

$|| \{ y = y_0 \wedge y = c \}\, y := y^2\, \{ y_0 = c \wedge y = x_0{}^2 \}$

$|| \{ z = z_0 \wedge z = c \}\, z := (z - d) * (z + d)\, \{ z_0 = c \wedge z = z_0{}^2 - d^2 \}$

$]$

$\{ x_0 = c \wedge x = x_0{}^2 \wedge y_0 = c \wedge y = x_0{}^2 \wedge z_0 = c \wedge z = z_0{}^2 - d^2 \}$

$\{ x = c^2 \wedge y = c^2 \wedge z = c^2 - d^2 \}$

$\{ x = y = z + d^2 \}$

G. Not All Changes are Interference

- We'll look into this in more detail next class, but can we loosen the requirement of disjoint conditions? (After all, requiring disjoint threads already disallows too many programs we want to write.)

- The basic observation is that often we can change the values of variables without changing whether a condition is satisfied or not. E.g., the condition $x > 0$ isn't invalidated by setting $x := 1$ or $x := x + 1$ (or $x := x$, for that matter).

- Say a thread includes a condition p (either as a precondition or postcondition; it doesn't matter), and a different thread is about to execute $\{ q_1 \}\, x := e\, \{ q_2 \}$. Surprisingly, q_2 isn't important; what is important is that executing $x := e$ only causes harm if it makes p false. More precisely, if p and q_1 hold and we execute $x := e$, we don't have interference if p is still satisfied after the assignment.

- As long as $\{ p \wedge q_1 \}\, x := e\, \{ p \}$ is valid, we know that whatever change $x := e$ causes, it won't disturb execution of the other thread, at least as far as p holding.

- More generally, with p and $\{ q_1 \}\, S_1\, \{ ... \}$ (we don't care about the postcondition of S_1), we're free of interference if $\{ p \wedge q_1 \}\, S_1\, \{ p \}$ is valid.

- We can build on this observation to get a notion of interference-freedom that works as a side condition for a nicer parallelism rule. It's not trivial because a thread has multiple statements that we have to check against multiple conditions in the other thread. But we can reduce the burden by looking more carefully at exactly what statements we need to check for interference.

- E.g., $x_1 := e_1 ; x_2 := e_2 ; x_3 := e_3$ comprises 5 statements (three assignments and two sequences). We'll have the check the assignments, of course, but do we need to check the two sequences? As it turns out, no, we don't.