

The VerCors Tool Set

Verification of Parallel and Concurrent Software

Stefan Blom Saeed Darabi
Marieke Huisman Wytse Oortwijn

Formal Methods and Tools, University of Twente

September 20, 2017


Outline

- 1 Introduction
- 2 Deterministic Parallelism
- 3 GPU Kernels
- 4 Model Abstractions
- 5 Conclusion

Outline

- 1 Introduction
- 2 Deterministic Parallelism
- 3 GPU Kernels
- 4 Model Abstractions
- 5 Conclusion

Concurrent programming is *error-prone*



Mars climate orbiter

Concurrent programming is *error-prone*




Mars climate orbiter




THERAC-25

Automated program verification



The VerCors Tool Set



Verifying Java programs: *example*

```
class Counter {  
    int count;  
  
    public void incr(int y) {  
        this.count := this.count + y;  
    }  
}
```


Verifying Java programs: *example*

```
class Counter {  
    int count;  
  
    requires Perm(this.count, write);  
    public void incr(int y) {  
        this.count := this.count + y;  
    }  
}
```

Writing permission for *this.count* is required

Verifying Java programs: *example*

```
class Counter {  
    int count;  
  
    requires Perm(this.count, write);  
    ensures Perm(this.count, write);  
    public void incr(int y) {  
        this.count := this.count + y;  
    }  
}
```


Writing permission for *this.count* is ensured

Verifying Java programs: *example*

```
class Counter {  
    int count;  
  
    requires Perm(this.count, write);  
    ensures Perm(this.count, write);  
    ensures this.count = old(this.count) + y;  
    public void incr(int y) {  
        this.count := this.count + y;  
    }  
}
```

this.count is correctly incremented

The VerCors Tool Set - *Overview*



Outline

- 1 Introduction
- 2 Deterministic Parallelism**
- 3 GPU Kernels
- 4 Model Abstractions
- 5 Conclusion

Deterministic parallel programming

Concurrency as optimisation

Sequential program
annotated by
compiler directives

compiled into

Multi-threaded programs
or
GPU kernels

Deterministic parallel programming

Concurrency as optimisation

Sequential program
annotated by
compiler directives

compiled into

Multi-threaded programs
or
GPU kernels

Compiler directives

Give hints to the compiler to know *where* and *how* to parallelise sequential code.

Deterministic parallel programming

Concurrency as optimisation

Sequential program
annotated by
compiler directives

compiled into

Multi-threaded programs
or
GPU kernels

Compiler directives

Give hints to the compiler to know *where* and *how* to parallelise sequential code. Some implementations are:

- OpenACC
- OpenMP
- PENCIL, etc.

Deterministic parallel programming

Concurrency as optimisation

Sequential program
annotated by
compiler directives

compiled into

Multi-threaded programs
or
GPU kernels

Compiler directives

Give hints to the compiler to know *where* and *how* to parallelise sequential code. Some implementations are:

- OpenACC
- OpenMP (*for C*)
- PENCIL, etc.

Verifying OpenMP programs - *NFM 2017*

```
for (int i = 0; i < N; i++) {  
    C[i] = A[i];  
}
```

```
for (int i = 0; i < N; i++) {  
    D[i] = C[i + 1] + B[i];  
}
```

Verifying OpenMP programs - *NFM 2017*

```
for (int i = 0; i < N; i++) {  
    C[i] = A[i];  
}
```

```
for (int i = 0; i < N; i++) {  
    D[i] = C[i + 1] + B[i];  
}
```

```
for (int i = 0; i < N; i++) {  
    C[i] = A[i];  
}
```

```
for (int i = 0; i < N; i++) {  
    D[i] = C[i] + B[i];  
}
```

Verifying OpenMP programs - *NFM 2017*

```
#pragma omp parallel {  
  #pragma omp for  
  for (int i = 0; i < N; i++) {  
    C[i] = A[i];  
  }  
  
  #pragma omp for  
  for (int i = 0; i < N; i++) {  
    D[i] = C[i + 1] + B[i];  
  }  
}
```

```
for (int i = 0; i < N; i++) {  
  C[i] = A[i];  
}  
  
for (int i = 0; i < N; i++) {  
  D[i] = C[i] + B[i];  
}
```

Verifying OpenMP programs - *NFM 2017*

```
#pragma omp parallel {  
  #pragma omp for  
  for (int i = 0; i < N; i++) {  
    C[i] = A[i];  
  }  
  
  #pragma omp for  
  for (int i = 0; i < N; i++) {  
    D[i] = C[i + 1] + B[i];  
  }  
}
```

```
#pragma omp parallel {  
  #pragma omp for  
  schedule(static) nowait  
  for (int i = 0; i < N; i++) {  
    C[i] = A[i];  
  }  
  
  #pragma omp for  
  schedule(static) nowait  
  for (int i = 0; i < N; i++) {  
    D[i] = C[i] + B[i];  
  }  
}
```

Verifying OpenMP programs - *NFM 2017*

```
#pragma omp parallel {  
  #pragma omp for  
  for (int i = 0; i < N; i++) {  
    C[i] = A[i];  
  }  
  
  #pragma omp for  
  for (int i = 0; i < N; i++) {  
    D[i] = C[i + 1] + B[i];  
  }  
}
```

```
#pragma omp parallel {  
  #pragma omp for  
  schedule(static) nowait  
  for (int i = 0; i < N; i++) {  
    C[i] = A[i];  
  }  
  
  #pragma omp for  
  schedule(static) nowait  
  for (int i = 0; i < N; i++) {  
    D[i] = C[i] + B[i];  
  }  
}
```

Verifying loop parallelisations - *FASE 2015*

Specifying a contract for every loop iteration, capturing its resources

Verifying loop parallelisations - *FASE 2015*

Specifying a contract for every loop iteration, capturing its resources

$$\overline{\vdash \text{for}(\text{int } i = 0; i < N; i++)\{ S_i \}}$$

Verifying loop parallelisations - *FASE 2015*

Specifying a contract for every loop iteration, capturing its resources

$$\frac{\vdash \{P_0\} S_0 \{Q_0\} \quad \dots \quad \vdash \{P_{N-1}\} S_{N-1} \{Q_{N-1}\}}{\vdash \text{for}(\text{int } i = 0; i < N; i++)\{S_i\}}$$

Verifying loop parallelisations - *FASE 2015*

Specifying a contract for every loop iteration, capturing its resources

$$\frac{\vdash \{P_0\} S_0 \{Q_0\} \quad \dots \quad \vdash \{P_{N-1}\} S_{N-1} \{Q_{N-1}\}}{\vdash \{P_0 * \dots * P_{N-1}\} \text{ for } (\text{int } i = 0; i < N; i++) \{ S_i \} \{Q_0 * \dots * Q_{N-1}\}}$$

Deterministic parallelism: *iteration contracts*

```
void clear(int A[], int len) {  
    for (int i = 0; i < len; i++) {  
        A[i] = 0;  
    }  
}
```

Deterministic parallelism: *iteration contracts*

```
requires  $A \neq \text{null} \wedge \text{length}(A) = \text{len};$   
void clear(int A[], int len) {  
    for (int i = 0; i < len; i++) {  
        A[i] = 0;  
    }  
}
```

The input array A may not be null

Deterministic parallelism: *iteration contracts*

```
requires  $A \neq \text{null} \wedge \text{length}(A) = \text{len};$   
requires (forall* int  $k$ ;  $0 \leq k < \text{len}$ ;  $\text{Perm}(A[k], \text{write})$ );  
void clear(int  $A[]$ , int  $\text{len}$ ) {  
    for (int  $i = 0$ ;  $i < \text{len}$ ;  $i++$ ) {  
         $A[i] = 0$ ;  
    }  
}
```

Permission is required to write to every element of A

Deterministic parallelism: *iteration contracts*

```
requires  $A \neq \text{null} \wedge \text{length}(A) = \text{len};$   
requires (forall* int  $k; 0 \leq k < \text{len}; \text{Perm}(A[k], \text{write})$ );  
ensures (forall* int  $k; 0 \leq k < \text{len}; \text{Perm}(A[k], \text{write})$ );  
ensures (forall int  $k; 0 \leq k < \text{len}; A[k] = 0$ );  
void clear(int  $A[], \text{int } \text{len}$ ) {  
    for (int  $i = 0; i < \text{len}; i++$ ) {  
         $A[i] = 0;$   
    }  
}
```

A “cleared” array is ensured

Deterministic parallelism: *iteration contracts*

```
requires  $A \neq \text{null} \wedge \text{length}(A) = \text{len}$ ;  
requires (forall* int  $k$ ;  $0 \leq k < \text{len}$ ;  $\text{Perm}(A[k], \text{write})$ );  
ensures (forall* int  $k$ ;  $0 \leq k < \text{len}$ ;  $\text{Perm}(A[k], \text{write})$ );  
ensures (forall int  $k$ ;  $0 \leq k < \text{len}$ ;  $A[k] = 0$ );  
void clear(int  $A[]$ , int  $\text{len}$ ) {  
    for (int  $i = 0$ ;  $i < \text{len}$ ;  $i++$ )  
        requires  $\text{Perm}(A[i], \text{write})$  {  
             $A[i] = 0$ ;  
        }  
}
```

Iteration i requires permission to write to $A[i]$

Deterministic parallelism: *iteration contracts*

```
requires  $A \neq \text{null} \wedge \text{length}(A) = \text{len}$ ;  
requires (forall* int  $k$ ;  $0 \leq k < \text{len}$ ; Perm( $A[k]$ , write));  
ensures (forall* int  $k$ ;  $0 \leq k < \text{len}$ ; Perm( $A[k]$ , write));  
ensures (forall int  $k$ ;  $0 \leq k < \text{len}$ ;  $A[k] = 0$ );  
void clear(int  $A[]$ , int  $\text{len}$ ) {  
    for (int  $i = 0$ ;  $i < \text{len}$ ;  $i++$ )  
        requires Perm( $A[i]$ , write);  
        ensures Perm( $A[i]$ , write) {  
             $A[i] = 0$ ;  
        }  
    }  
}
```

Iteration i ensures permission to write to $A[i]$

Deterministic parallelism: *iteration contracts*

```

requires  $A \neq \text{null} \wedge \text{length}(A) = \text{len}$ ;
requires (forall* int  $k$ ;  $0 \leq k < \text{len}$ ; Perm( $A[k]$ , write));
ensures (forall* int  $k$ ;  $0 \leq k < \text{len}$ ; Perm( $A[k]$ , write));
ensures (forall int  $k$ ;  $0 \leq k < \text{len}$ ;  $A[k] = 0$ );
void clear(int  $A$ [], int  $\text{len}$ ) {
    for (int  $i = 0$ ;  $i < \text{len}$ ;  $i++$ )
        requires Perm( $A[i]$ , write);
        ensures Perm( $A[i]$ , write);
        ensures  $A[i] = 0$  {
             $A[i] = 0$ ;
        }
    }
}

```

Iteration i “cleared” $A[i]$

Deterministic parallelism: *verified examples*

- 1 **Histogram:** calculating the *histogram* of a matrix.
- 2 **Loop dependencies:** several examples with forward/backward *loop dependencies*.
- 3 **Reductions:** several examples with *reduction patterns* (e.g. summations).
- 4 **Loop vectorisation:** several examples with *loop vectorisations*.

Deterministic parallelism: *verified examples*

- 1 **Histogram:** calculating the *histogram* of a matrix.
- 2 **Loop dependencies:** several examples with forward/backward *loop dependencies*.
- 3 **Reductions:** several examples with *reduction patterns* (e.g. summations).
- 4 **Loop vectorisation:** several examples with *loop vectorisations*.

Deterministic parallelism: *verified examples*

- 1 **Histogram:** calculating the *histogram* of a matrix.
- 2 **Loop dependencies:** several examples with forward/backward *loop dependencies*.
- 3 **Reductions:** several examples with *reduction patterns* (e.g. summations).
- 4 **Loop vectorisation:** several examples with *loop vectorisations*.


See our *FASE2015* and *NFM2017* papers for more details

Outline


- 1 Introduction
- 2 Deterministic Parallelism
- 3 GPU Kernels**
- 4 Model Abstractions
- 5 Conclusion

GPU computing model

CPU (Multiple Cores)



GPU (Hundreds of Cores)



source <https://www.mathworks.com>

Single-Instruction-Multiple-Data (*SIMD*)

```
--kernel void add1(__global int A[]) {  
    // get thread identifier  
    int tid = get_global_id(0);  
  
    // increase entry by one  
    A[tid] = A[tid] + 1;  
}
```

Single-Instruction-Multiple-Data (SIMD)

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```

```

..kernel void add1(..global int A[]) {
  // get thread identifier
  int tid = get_global_id(0);

  // increase entry by one
  A[tid] = A[tid] + 1;
}

```


Single-Instruction-Multiple-Data (*SIMD*)

<i>tid</i> = 0	<i>tid</i> = 1	<i>tid</i> = 2	<i>tid</i> = 3	<i>tid</i> = 4
<i>tid</i> = 5	<i>tid</i> = 6	<i>tid</i> = 7	<i>tid</i> = 8	<i>tid</i> = 9
<i>tid</i> = 10	<i>tid</i> = 11	<i>tid</i> = 12	<i>tid</i> = 13	<i>tid</i> = 14
<i>tid</i> = 15	<i>tid</i> = 16	<i>tid</i> = 17	<i>tid</i> = 18	<i>tid</i> = 19

Single-Instruction-Multiple-Data (*SIMD*)

$A[0] = A[0] + 1$	$A[1] = A[1] + 1$	$A[2] = A[2] + 1$	$A[3] = A[3] + 1$	$A[4] = A[4] + 1$
$A[5] = A[5] + 1$	$A[6] = A[6] + 1$	$A[7] = A[7] + 1$	$A[8] = A[8] + 1$	$A[9] = A[9] + 1$
$A[10] = A[10] + 1$	$A[11] = A[11] + 1$	$A[12] = A[12] + 1$	$A[13] = A[13] + 1$	$A[14] = A[14] + 1$
$A[15] = A[15] + 1$	$A[16] = A[16] + 1$	$A[17] = A[17] + 1$	$A[18] = A[18] + 1$	$A[19] = A[19] + 1$

Thread specifications

```
__kernel void add1(__global int A[]) {  
    // get thread identifier  
    int tid = get_global_id(0);  
  
    // increase entry by one  
    A[tid] = A[tid] + 1;  
}
```

Thread specifications

```
requires A ≠ null;
__kernel void add1(__global int A[]) {
    // get thread identifier
    int tid = get_global_id(0);

    // increase entry by one
    A[tid] = A[tid] + 1;
}
```

The input array may not be null

Thread specifications

```
requires A ≠ null;
requires Perm(A[get_global_id(0)], write);
__kernel void add1(__global int A[]) {
    // get thread identifier
    int tid = get_global_id(0);

    // increase entry by one
    A[tid] = A[tid] + 1;
}
```

Thread with ID i requires *write permission* for $A[i]$

Thread specifications

```
requires A ≠ null;
requires Perm(A[get_global_id(0)], write);
ensures Perm(A[get_global_id(0)], write);
__kernel void add1(__global int A[]) {
    // get thread identifier
    int tid = get_global_id(0);

    // increase entry by one
    A[tid] = A[tid] + 1;
}
```

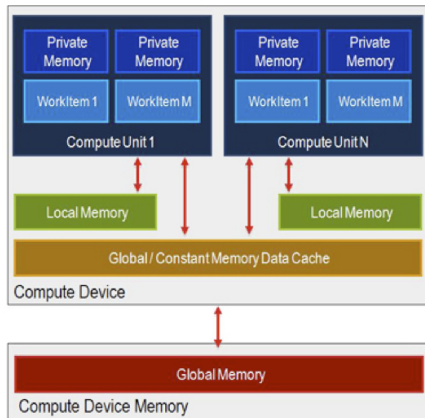
Thread with ID i ensures *write permission* for $A[i]$

Thread specifications

```
requires  $A \neq \text{null}$ ;  
requires Perm( $A[\text{get\_global\_id}(0)]$ , write);  
ensures Perm( $A[\text{get\_global\_id}(0)]$ , write);  
ensures  $A[\text{get\_global\_id}(0)] = \text{old}(A[\text{get\_global\_id}(0)]) + 1$ ;  
__kernel void add1(__global int A[]) {  
    // get thread identifier  
    int tid = get_global_id(0);  
  
    // increase entry by one  
     $A[\text{tid}] = A[\text{tid}] + 1$ ;  
}
```

Thread with ID i has incremented $A[i]$

GPU architecture: *threads & workgroups*



Barrier specifications

```
--kernel void addcomplex(--global int A[]) {  
    // get thread identifier  
    int tid = get_global_id(0);  
  
    // perform the updates  
    A[tid] = A[tid] + 1;  
    if (tid > 0) {  
        A[tid - 1] = 2 * A[tid - 1];  
    }  
}
```

Barrier specifications

```
__kernel void addcomplex(__global int A[]) {  
    // get thread identifier  
    int tid = get_global_id(0);  
  
    // perform the updates  
    A[tid] = A[tid] + 1;  
  
    barrier(global);  
  
    if (tid > 0) {  
        A[tid - 1] = 2 * A[tid - 1];  
    }  
}
```

A barrier is required here for correctness!

Barrier specifications

```
requires Perm(A[get_global_id(0)], write);
__kernel void addcomplex(__global int A[]) {
    // get thread identifier
    int tid = get_global_id(0);

    // perform the updates
    A[tid] = A[tid] + 1;

    barrier(global);

    if (tid > 0) {
        A[tid - 1] = 2 * A[tid - 1];
    }
}
```

Thread with ID i can *write* to $A[i]$ and *read* from $B[i]$

Barrier specifications

```

requires Perm(A[get_global_id(0)], write);
__kernel void addcomplex(__global int A[]) {
    // get thread identifier
    int tid = get_global_id(0);

    // perform the updates
    A[tid] = A[tid] + 1;

    barrier(global) {
        requires Perm(A[tid], write);
        ensures tid > 0  $\Rightarrow$  Perm(A[tid - 1], write);
    };

    if (tid > 0) {
        A[tid - 1] = 2 * A[tid - 1];
    }
}

```

The barrier *consumes* all read permissions

Barrier specifications

```
requires Perm(A[get_global_id(0)], write);
__kernel void addcomplex(__global int A[]) {
    // get thread identifier
    int tid = get_global_id(0);

    // perform the updates
    A[tid] = A[tid] + 1;

    barrier(global) {
        requires Perm(A[tid], write);
        ensures tid > 0  $\Rightarrow$  Perm(A[tid - 1], write);
    };

    if (tid > 0) {
        A[tid - 1] = 2 * A[tid - 1];
    }
}
```

The barrier *redistributes* all consumed permissions

GPU kernels: *verified examples*

- 1 **Parallel prefix sum:** data-race freedom of a *parallell prefix sum* algorithm.
- 2 **Summations:** kernels with *atomic*, *barriers*, and *reductions*.
- 3 **Floats:** support for *floating point* arithmetics.

GPU kernels: *verified examples*

- 1 **Parallel prefix sum:** data-race freedom of a *parallell prefix sum* algorithm.
- 2 **Summations:** kernels with *atomic*, *barriers*, and *reductions*.
- 3 **Floats:** support for *floating point* arithmetics.

Outline

- 1 Introduction
- 2 Deterministic Parallelism
- 3 GPU Kernels
- 4 Model Abstractions**
- 5 Conclusion

Static verification: *deductive verification*

```

requires  $P_1 * P_2$ ;
ensures  $Q_1 * Q_2$ ;
void main(args) {
     $S_1 \parallel S_2$ ;
}
  
```

$$\frac{\frac{\dots}{\vdash \{P_1\} S_1 \{Q_1\}} \quad \frac{\dots}{\vdash \{P_2\} S_2 \{Q_2\}}}{\vdash \{P_1 * P_2\} S_1 \parallel S_2 \{Q_1 * Q_2\}}$$

$$\frac{}{\vdash \{P_1 * P_2\} \text{main}(\textit{args}) \{Q_1 * Q_2\}}$$

Static verification: *deductive verification*

```

requires  $P_1 * P_2$ ;
ensures  $Q_1 * Q_2$ ;
void main(args) {
     $S_1 \parallel S_2$ ;
}

```

$$\frac{\frac{\dots}{\vdash \{P_1\} S_1 \{Q_1\}} \quad \frac{\dots}{\vdash \{P_2\} S_2 \{Q_2\}}}{\vdash \{P_1 * P_2\} S_1 \parallel S_2 \{Q_1 * Q_2\}}$$

$$\frac{}{\vdash \{P_1 * P_2\} \text{main}(\textit{args}) \{Q_1 * Q_2\}}$$

Strong correctness guarantees

Static verification: *deductive verification*

```

requires  $P_1 * P_2$ ;
ensures  $Q_1 * Q_2$ ;
void main(args) {
     $S_1 \parallel S_2$ ;
}
  
```

$$\frac{\frac{\dots}{\vdash \{P_1\} S_1 \{Q_1\}} \quad \frac{\dots}{\vdash \{P_2\} S_2 \{Q_2\}}}{\vdash \{P_1 * P_2\} S_1 \parallel S_2 \{Q_1 * Q_2\}}$$

$$\frac{}{\vdash \{P_1 * P_2\} \text{main}(\textit{args}) \{Q_1 * Q_2\}}$$

Hard to construct a proof

Static verification: *model checking*


```
void main(args) {  
     $S_1 \parallel S_2$ ;  
}
```

Static verification: *model checking*


```
void main(args) {  
     $S_1 \parallel S_2$ ;  
}
```

Abstract model
capturing program
behaviour


Static verification: *model checking*




Static verification: *model checking*




Static verification: *model checking*




Static verification: *model checking*



Static verification: *model checking*




Static verification: *model checking*




Good automation

Static verification: *model checking*




State-space explosion problems

Static verification: *model checking*




Is the program abstraction correct?

Our approach: *deductive + algorithmic verification*

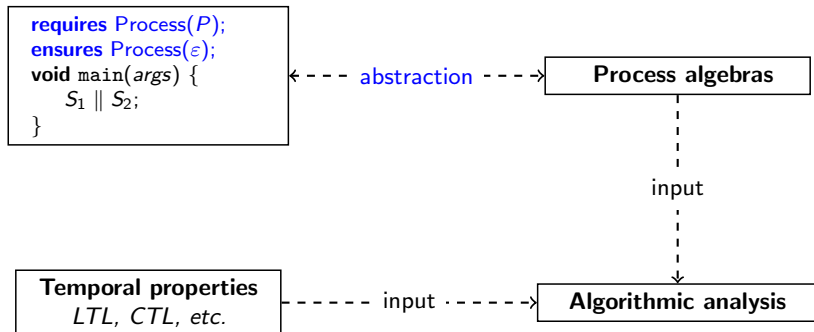


Our approach: *deductive + algorithmic verification*




Abstract models are *process algebra terms*

Our approach: *deductive + algorithmic verification*




Deductively verifying *correctness* of the abstraction

Our approach: *deductive + algorithmic verification*



The properties are encoded as *process/action contracts*

Our approach: *deductive + algorithmic verification*



Applying model checking result *in the program logic*

Example program: *concurrent counting*

Concrete program

```
class Counter {  
  public int val;  
  
  static void add2(Counter c) {  
    parallel {  
      atomic-add(c.val, 1);  
    } and {  
      atomic-add(c.val, 1);  
    }  
  }  
}
```

Program abstraction

Example program: *concurrent counting*

Concrete program

```
class Counter {
  public int val;

  ensures c.val = old(c.val) + 2;
  static void add2(Counter c) {
    parallel {
      atomic-add(c.val, 1);
    } and {
      atomic-add(c.val, 1);
    }
  }
}
```

Program abstraction

How to verify this property?

Example program: *concurrent counting*

Concrete program

```
class Counter {  
  public int val;  
  
  static void add2(Counter c) {  
    parallel {  
      atomic-add(c.val, 1);  
    } and {  
      atomic-add(c.val, 1);  
    }  
  }  
}
```

Program abstraction

```
action incr;
```

The `incr` action *abstracts* the `atomic-add`'s.

Example program: *concurrent counting*

Concrete program

```
class Counter {  
  public int val;  
  
  static void add2(Counter c) {  
    parallel {  
      atomic-add(c.val, 1);  
    } and {  
      atomic-add(c.val, 1);  
    }  
  }  
}
```

Program abstraction

```
action incr;  
  
process parincr() := incr || incr;
```

The `parincr` process *abstracts* the `add2` program.

Example program: *concurrent counting*

Concrete program

```
class Counter {  
  public int val;  
  
  static void add2(Counter c) {  
    parallel {  
      atomic-add(c.val, 1);  
    } and {  
      atomic-add(c.val, 1);  
    }  
  }  
}
```

Program abstraction

```
guard true;  
effect x = old(x) + 1;  
action incr;  
  
process parincr() := incr || incr;
```

Contracts are used to encode the *properties of interest*.

Example program: *concurrent counting*

Concrete program

```

class Counter {
  public int val;

  static void add2(Counter c) {
    parallel {
      atomic-add(c.val, 1);
    } and {
      atomic-add(c.val, 1);
    }
  }
}

```

Program abstraction

```

guard true;
effect x = old(x) + 1;
action incr;

requires true;
ensures x = old(x) + 2;
process parincr() := incr || incr;

```

Contracts are used to encode the *properties of interest*.

Example program: *concurrent counting*

Concrete program

```

class Counter {
  public int val;

  requires Process(parincr);
  requires c.val ~ x;
  static void add2(Counter c) {
    parallel {
      atomic-add(c.val, 1);
    } and {
      atomic-add(c.val, 1);
    }
  }
}

```

Program abstraction

```

guard true;
effect x = old(x) + 1;
action incr;

requires true;
ensures x = old(x) + 2;
process parincr() := incr || incr;

```

Linking concrete program to abstract model.

Example program: *concurrent counting*

Concrete program

```

class Counter {
  public int val;

  requires Process(parincr);
  requires c.val ~ x;
  static void add2(Counter c) {
    parallel {
      act incr do atomic-add(c.val, 1);
    } and {
      act incr do atomic-add(c.val, 1);
    }
  }
}

```

Program abstraction

```

guard true;
effect x = old(x) + 1;
action incr;

requires true;
ensures x = old(x) + 2;
process parincr() := incr || incr;

```

Linking abstract actions to concrete implementation.

Example program: *concurrent counting*

Concrete program

```

class Counter {
  public int val;

  requires Process(parincr);
  requires c.val ~ x;
  ensures Process( $\epsilon$ );
  static void add2(Counter c) {
    parallel {
      act incr do atomic-add(c.val, 1);
    } and {
      act incr do atomic-add(c.val, 1);
    }
  }
}

```

Program abstraction

```

guard true;
effect x = old(x) + 1;
action incr;

requires true;
ensures x = old(x) + 2;
process parincr() := incr || incr;

```

Linking concrete program to abstract model.

Example program: *concurrent counting*

Concrete program

```

class Counter {
  public int val;

  requires Process(parincr);
  requires c.val ~ x;
  ensures Process( $\varepsilon$ );
  ensures c.val = old(c.val) + 2;
  static void add2(Counter c) {
    parallel {
      act incr do atomic-add(c.val, 1);
    } and {
      act incr do atomic-add(c.val, 1);
    }
  }
}

```

Program abstraction

```

guard true;
effect x = old(x) + 1;
action incr;

requires true;
ensures x = old(x) + 2;
process parincr() := incr || incr;

```

Using the result from analyzing all traces of `parincr`.

Model abstractions: *underlying theory*



Formalised using the Coq proof assistant

Model abstractions: *verified examples*

- 1 **Concurrent counting**: verifying that a concurrent counter computes the correct result.
- 2 **No-send-after-read**: verifying that confidential data can not be send after being received.
- 3 **Parallel GCD**: verifying that a parallel GCD algorithm computes the correct result.
- 4 **Locking protocol**: verifying that a lock implementation adheres to the intended locking protocol.

Model abstractions: *verified examples*

- 1 **Concurrent counting**: verifying that a concurrent counter computes the correct result.
- 2 **No-send-after-read**: verifying that confidential data can not be send after being received.
- 3 **Parallel GCD**: verifying that a parallel GCD algorithm computes the correct result.
- 4 **Locking protocol**: verifying that a lock implementation adheres to the intended locking protocol.

Model abstractions: *verified examples*

- 1 Concurrent counting:** verifying that a concurrent counter computes the correct result.
- 2 No-send-after-read:** verifying that confidential data can not be send after being received.
- 3 Parallel GCD:** verifying that a parallel GCD algorithm computes the correct result.
- 4 Locking protocol:** verifying that a lock implementation adheres to the intended locking protocol.

See our *VSTTE2017* paper for more details

Outline

- 1 Introduction
- 2 Deterministic Parallelism
- 3 GPU Kernels
- 4 Model Abstractions
- 5 Conclusion**

Conclusion

Future directions

- 1 Verifying *distributed software*.
- 2 Automatic generation of *annotations*.
- 3 Generating *meaningful* error messages.

Conclusion

Future directions

- 1 Verifying *distributed software*.
- 2 Automatic generation of *annotations*.
- 3 Generating *meaningful* error messages.

VerCors Verification Tool Set

- Automated verification of parallel and concurrent software.
- More information: <http://utwente.nl/vercors>.
- Download & try online: <https://github.com/utwente-fmt/vercors>.