

# SAT Solving with GPU-Accelerated Inprocessing

TACAS'21

Muhammad Osama<sup>1</sup> – Anton Wijs<sup>1</sup> – Armin Biere<sup>2</sup>

<sup>1</sup> Eindhoven University of Technology, The Netherlands

<sup>2</sup> Johannes Kepler University, Austria



# Satisfiability (SAT) Solving – Warm up

- A problem of finding a satisfying assignment for a Boolean formula
  - If such assignment exists, the problem is SATISFIABLE (SAT)
  - Otherwise, it's UNSATISFIABLE (UNSAT)
  - Usually presented in Conjunctive Normal Form (CNF)

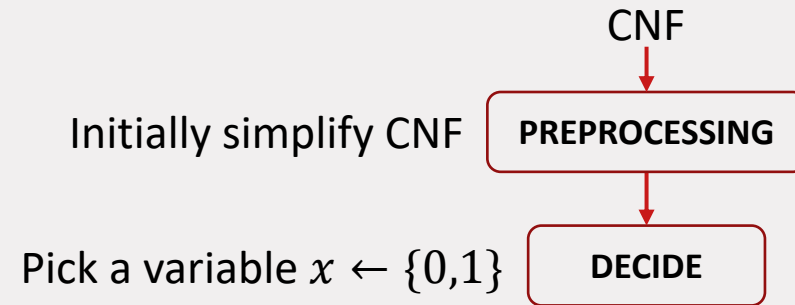
$$F = (a \vee b \vee c) \wedge (\neg a \vee b) \wedge (\neg b \vee c)$$

- Modern SAT solvers use **Conflict-Driven Clause-Learning (CDCL)** search algorithm + simplifications

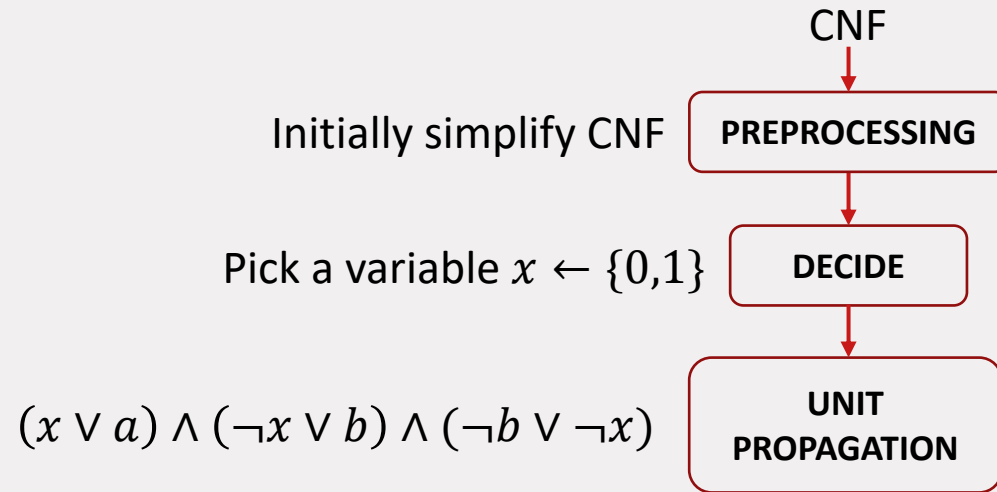
# SAT Solving – Warm up



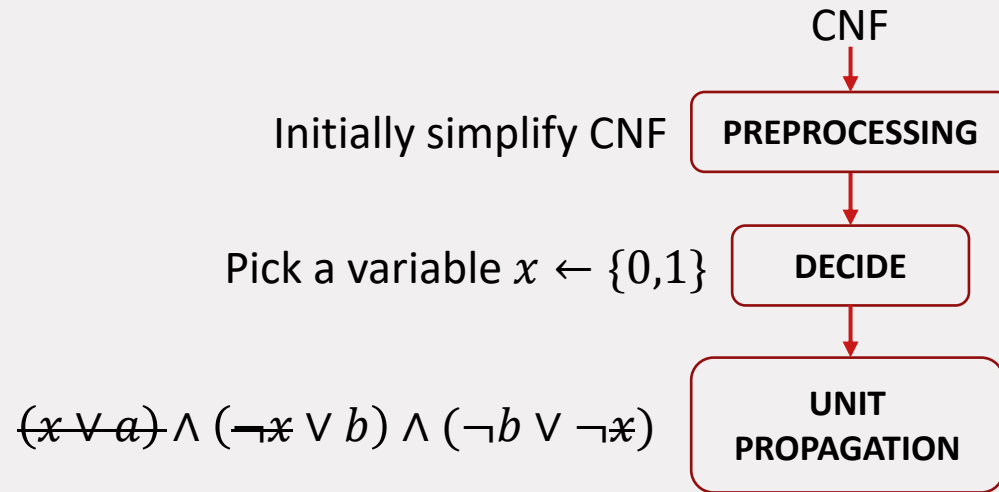
# SAT Solving – Warm up



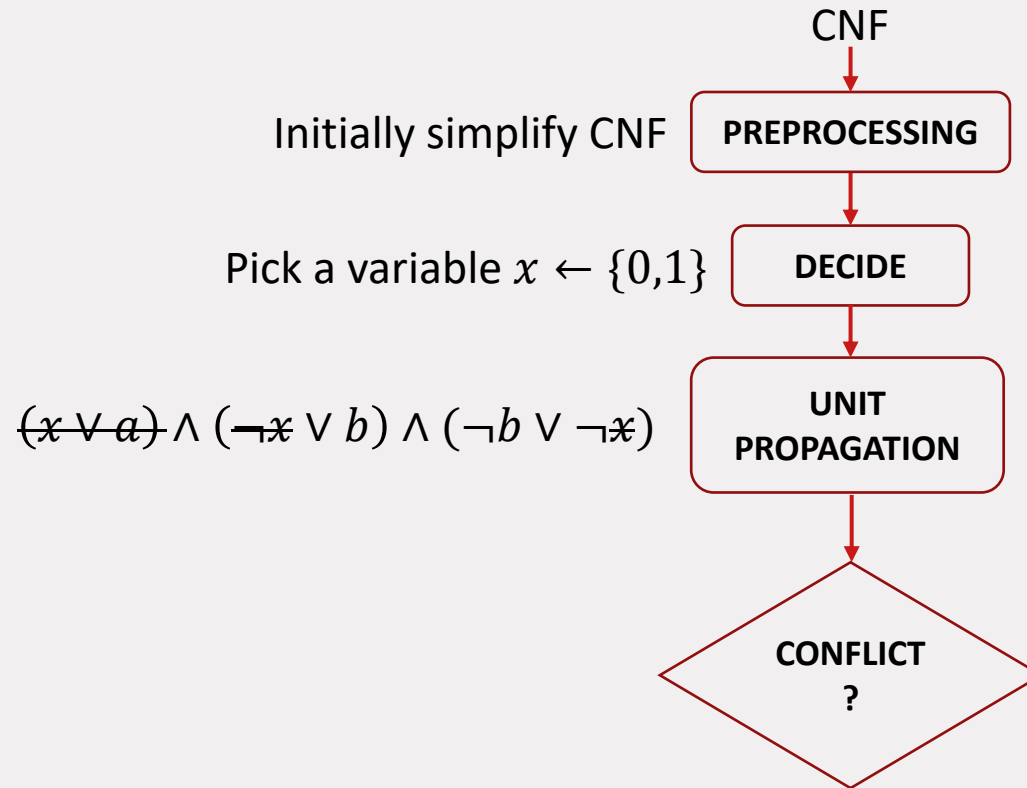
# SAT Solving – Warm up



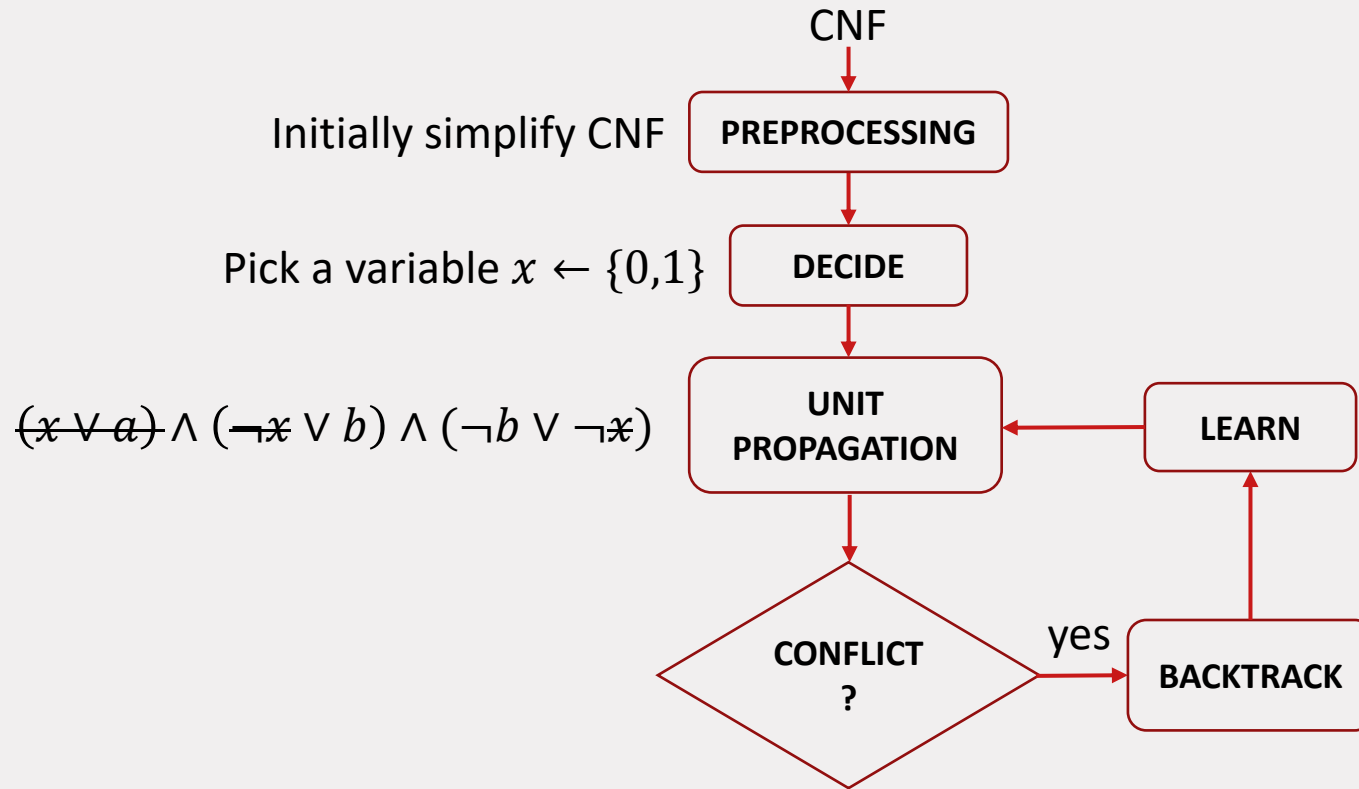
# SAT Solving – Warm up



# SAT Solving – Warm up

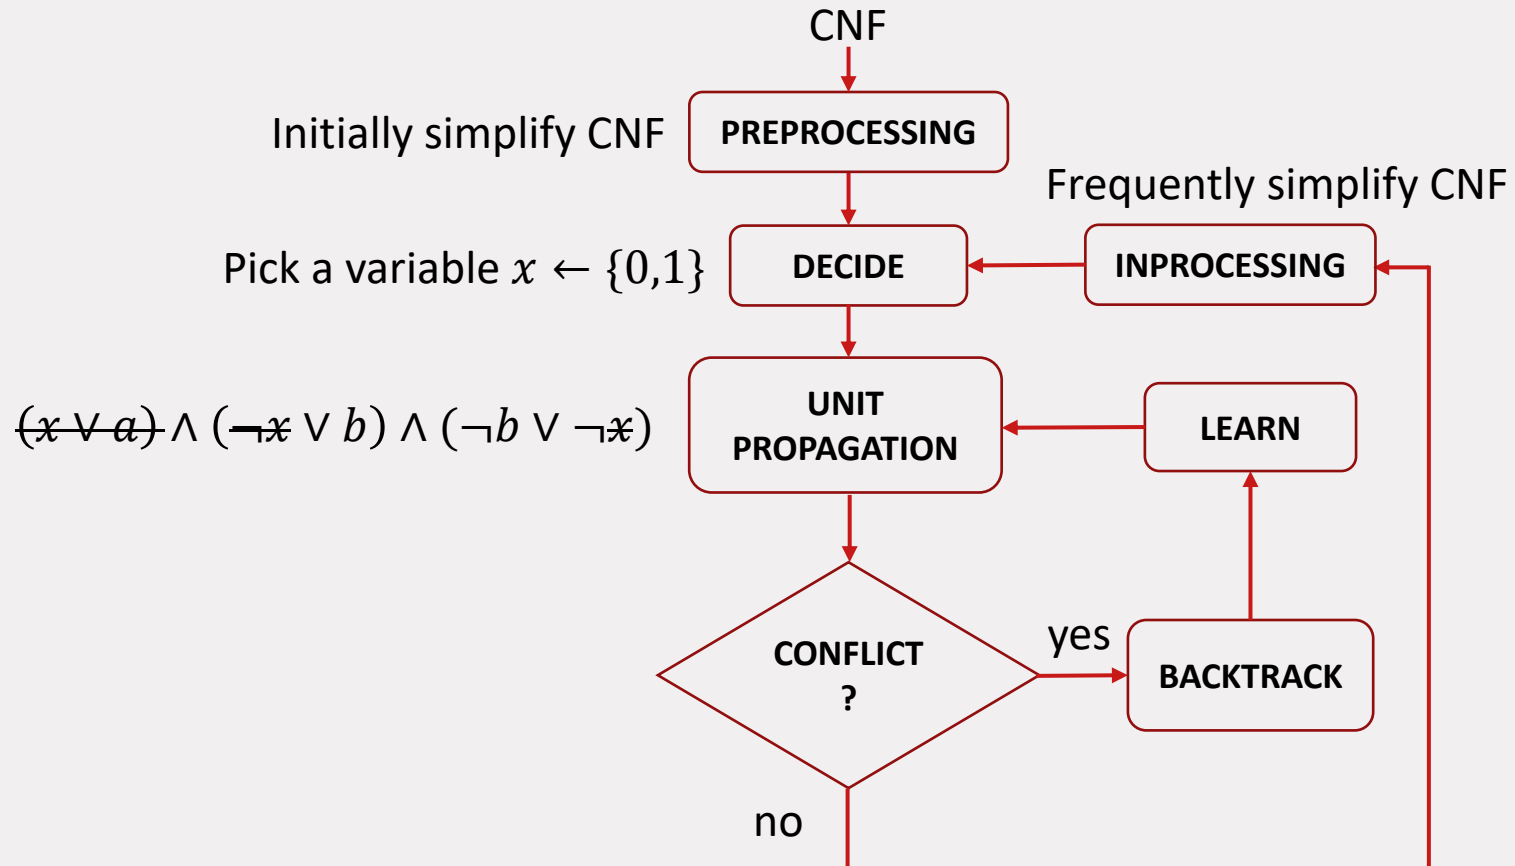


# SAT Solving – Warm up





# SAT Solving – Warm up



# Contributions

## GPU-accelerated Inprocessing

Challenge	Proposed solution
GPU has a very limited memory (24 GBs in powerful GPUs)	✓ New <b>space-efficient data structure</b> ✓ <b>Parallel garbage collection</b>
Our previous work in parallel variable elimination couldn't reproduce results and heavily relies on atomics	✓ Garbage collector designed to preserve the order of clauses ✓ <b>3-phase parallel variable elimination</b>
Many redundant clauses but not all simplifications are suitable for GPU	✓ New simplification method called <b>eager redundancy elimination</b> suitable for data parallelism

# Data Structure

```
class SCLAUSE {  
    char state, used;  
    char flag, added;  
    uint32 lbd, sig;  
    int size;  
    uint32 literals[1];  
};
```

Clause container

```
class CNF {  
    struct {  
        uint32* memory;  
        uint64 size, cap;  
    } clauses;  
    struct {  
        uint64* memory;  
        uint64 size, cap;  
    } references;  
}
```

CNF container

# Data Structure

```
class SCLAUSE {  
    char state, used;  
    char flag, added;  
    uint32 lbd, sig;  
    int size;  
    uint32 literals[1];  
};
```

Clause container

```
class CNF {  
    struct {  
        uint32* memory;  
        uint64 size, cap;  
    } clauses;  
    struct {  
        uint64* memory;  
        uint64 size, cap;  
    } references;  
}
```

CNF container

Example:  $C = \{ a \vee \neg b \vee c \}$

# Data Structure

```
class SCLAUSE {  
    char state, used;  
    char flag, added;  
    uint32 lbd, sig;  
    int size;  
    uint32 literals[1];  
};
```

Clause container

```
class CNF {  
    struct {  
        uint32* memory;  
        uint64 size, cap;  
    } clauses;  
    struct {  
        uint64* memory;  
        uint64 size, cap;  
    } references;  
};
```

CNF container

Example:  $C = \{ a \vee \neg b \vee c \}$

bytes =  $20 + 4(\text{size} - 1) = 28$

buckets = bytes / 4  
= size + 4 = 7

# Data Structure

```
class SCLAUSE {
    char state, used;
    char flag, added;
    uint32 lbd, sig;
    int size;
    uint32 literals[1];
};
```

Clause container

```
class CNF {
    struct {
        uint32* memory;
        uint64 size, cap;
    } clauses;
    struct {
        uint64* memory;
        uint64 size, cap;
    } references;
}
```

CNF container

Example:  $C = \{ a \vee \neg b \vee c \}$

bytes =  $20 + 4(\text{size} - 1) = 28$

buckets = bytes / 4

= size + 4 = 7

# Data Structure

```
class SCLAUSE {
    char state, used;
    char flag, added;
    uint32 lbd, sig;
    int size;
    uint32 literals[1];
};
```

Clause container

```
class CNF {
    struct {
        uint32* memory;
        uint64 size, cap;
    } clauses;
    struct {
        uint64* memory;
        uint64 size, cap;
    } references;
}
```

CNF container

Example:  $C = \{ a \vee \neg b \vee c \}$

bytes =  $20 + 4(\text{size} - 1) = 28$

buckets = bytes / 4

= size + 4 = 7

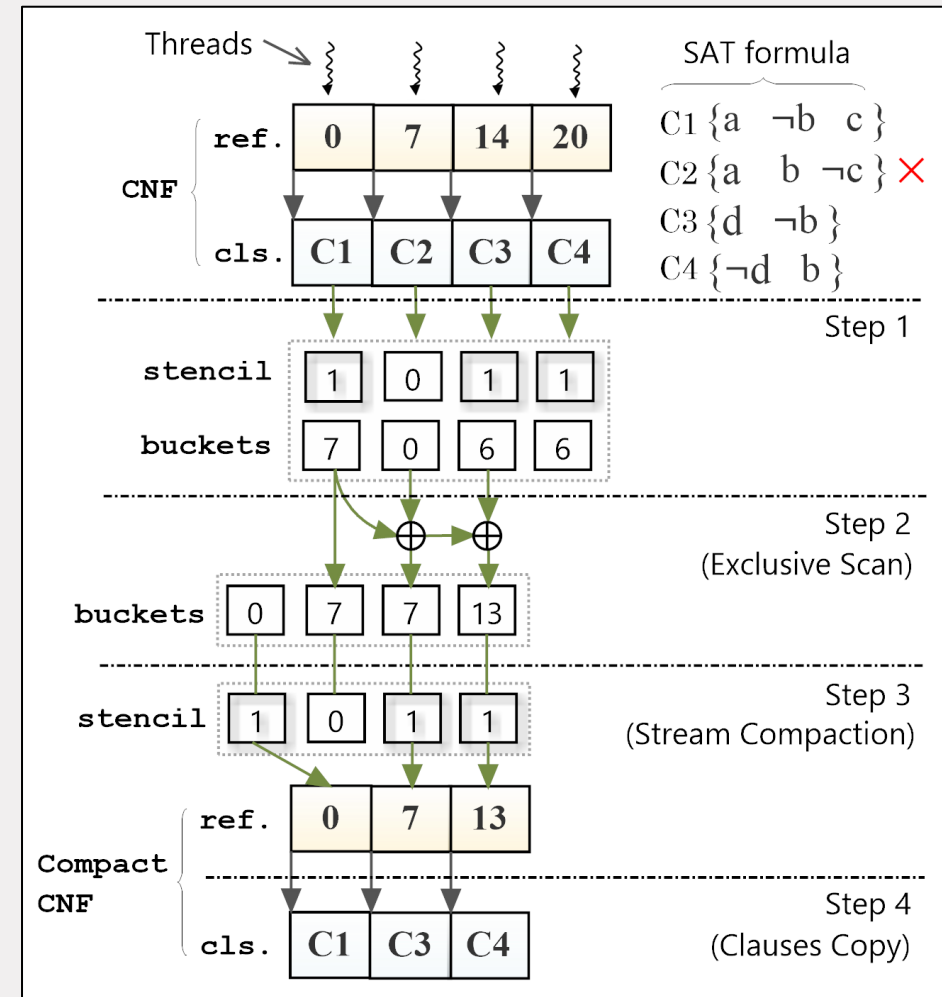
$\rightarrow$  cptr =  
(SCLAUSE\*) (clauses.memory + 0)

references.memory[0] = 0

# Parallel Garbage Collection

- Stencil: a binary array to mark the existence of an element
- Exclusive scan<sup>1</sup>: a parallel prefix sum offset by one position to the right
- Stream compaction<sup>1</sup>: compacts sparse data in parallel to consecutive elements

<sup>1</sup><https://nvlabs.github.io/cub/>





# Bounded Variable Elimination (BVE) - Recap

- Resolution rule: 
$$\frac{(x \vee \ell \vee \dots) \quad (\neg x \vee \ell' \vee \dots)}{(\ell \vee \dots \vee \ell' \vee \dots)}$$

# Bounded Variable Elimination (BVE) - Recap

- Resolution rule: 
$$\frac{(x \vee \ell \vee \dots) \quad (\neg x \vee \ell' \vee \dots)}{(\ell \vee \dots \vee \ell' \vee \dots)}$$
- Example:  $S = (x \vee a) \wedge (x \vee b) \wedge (\neg x \vee c)$

# Bounded Variable Elimination (BVE) - Recap

- Resolution rule: 
$$\frac{(x \vee \ell \vee \dots) \quad (\neg x \vee \ell' \vee \dots)}{(\ell \vee \dots \vee \ell' \vee \dots)}$$
- Example:  $S = (x \vee a) \wedge (x \vee b) \wedge (\neg x \vee c)$

$$S = (a \vee c) \wedge (b \vee c)$$



resolvents

# Bounded Variable Elimination (BVE) - Recap

- Resolution rule: 
$$\frac{(x \vee \ell \vee \dots) \quad (\neg x \vee \ell' \vee \dots)}{(\ell \vee \dots \vee \ell' \vee \dots)}$$
- Gate equivalence reasoning can be also used to find patterns of logical gates (e.g. AND, OR, XOR, ...) and substitute them with their definitions
- Gate substitution reduces the number of resolvents

# Parallel BVE

- In previous work (TACAS19), we proposed our parallel version of BVE on the GPU where threads can eliminate *mutually-independent* variables in parallel

# Parallel BVE

- In previous work (TACAS19), we proposed our parallel version of BVE on the GPU where threads can eliminate *mutually-independent* variables in parallel
- Achieved a speedup of 32x compared to sequential version
- Had two shortcomings:
  - Heavily relies on atomics to add new resolvents
  - Resolvents are added to random positions (side effect of parallel atomics)

# 3-Phase Parallel BVE

# 3-Phase Parallel BVE

0: Initially

- occurrence table ( $T$ ) for all literals in formula ( $S$ ) is created on the GPU
- A list ( $\phi$ ) of independent variables is generated by the variable scheduler on the CPU side
- **For all  $x \in \phi$  in parallel**,  $\text{sort}(T, \text{key})$



# 3-Phase Parallel BVE

0: Initially

- occurrence table ( $T$ ) for all literals in formula ( $S$ ) is created on the GPU
- A list ( $\phi$ ) of independent variables is generated by the variable scheduler on the CPU side
- **For all  $x \in \phi$  in parallel**,  $\text{sort}(T, \text{key})$

1: **For all  $x \in \phi$  in parallel**, Identify the type of elimination and calculate the resolvents info.

$\text{type}, \text{added}, \text{buckets} \leftarrow \text{variableSweep}(\phi, S, T)$

# 3-Phase Parallel BVE

0: Initially

- occurrence table ( $T$ ) for all literals in formula ( $S$ ) is created on the GPU
- A list ( $\phi$ ) of independent variables is generated by the variable scheduler on the CPU side
- **For all  $x \in \phi$  in parallel**,  $\text{sort}(T, \text{key})$

1: **For all  $x \in \phi$  in parallel**, Identify the type of elimination and calculate the resolvents info.

$\text{type}, \text{added}, \text{buckets} \leftarrow \text{variableSweep}(\phi, S, T)$

2: Compute the new positions of resolvents using parallel prefix sum

$\text{buckets} \leftarrow \text{exclusiveSum}(\text{buckets}, \text{size}(\text{clauses}))$  ← old size of all clauses in buckets

$\text{added} \leftarrow \text{exclusiveSum}(\text{added}, \text{size}(\text{references}))$  ← old number of clauses

# 3-Phase Parallel BVE

0: Initially

- occurrence table ( $T$ ) for all literals in formula ( $S$ ) is created on the GPU
- A list ( $\phi$ ) of independent variables is generated by the variable scheduler on the CPU side
- **For all  $x \in \phi$  in parallel**,  $\text{sort}(T, \text{key})$

1: **For all  $x \in \phi$  in parallel**, Identify the type of elimination and calculate the resolvents info.

$\text{type}, \text{added}, \text{buckets} \leftarrow \text{variableSweep}(\phi, S, T)$

2: Compute the new positions of resolvents using parallel prefix sum

$\text{buckets} \leftarrow \text{exclusiveSum}(\text{buckets}, \text{size}(\text{clauses}))$  ← old size of all clauses in buckets

$\text{added} \leftarrow \text{exclusiveSum}(\text{added}, \text{size}(\text{references}))$  ← old number of clauses

3: **For all  $x \in \phi$  in parallel**, resolve or substitute  $x$

$S, \text{units} \leftarrow \text{variableResolvent}(\phi, S, T, \text{type}, \text{added}, \text{buckets})$

# Eager Redundancy Elimination (ERE)

- We observed that unit propagation results in lots of redundancies
- ERE is a combination of resolution + equivalence:

For a given formula  $S$  and  $C_1 \in S, C_2 \in S$  with  $x \in C_1, \neg x \in C_2$  for some variable  $x$ , if there exists a clause  $C \in S$  such that  $C \equiv C_1 \otimes_x C_2$  then  $S = S \setminus C$

- Since we learn clauses during SAT solving, another constraint is checked before the removal of the redundancy


$(C_1 \text{ is LEARNT} \vee C_2 \text{ is LEARNT}) \Rightarrow C \text{ is LEARNT}$

# Eager Redundancy Elimination (ERE)

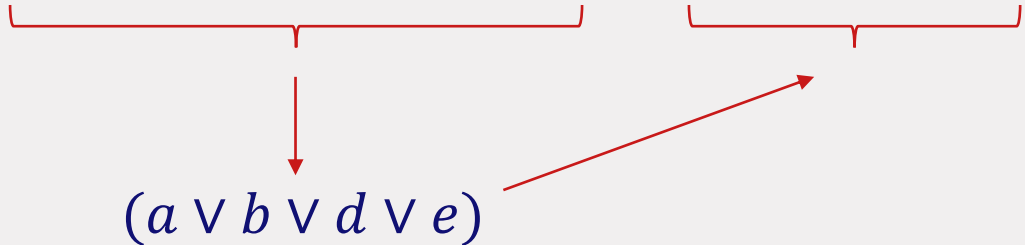
- Example:  $(a \vee b \vee \neg c) \wedge (d \vee e \vee c) \wedge (a \vee b \vee d \vee e)$

# Eager Redundancy Elimination (ERE)

- Example:  $(a \vee b \vee \neg c) \wedge (d \vee e \vee c) \wedge (a \vee b \vee d \vee e)$




# Eager Redundancy Elimination (ERE)

- Example:  $(a \vee b \vee \neg c) \wedge (d \vee e \vee c) \wedge \cancel{(a \vee b \vee d \vee e)}$   


The diagram illustrates the simplification process. Red brackets group the first two terms of the original expression:  $(a \vee b \vee \neg c)$  and  $(d \vee e \vee c)$ . A red arrow points from the resulting expression  $(a \vee b \vee d \vee e)$  back to the third term  $(a \vee b \vee d \vee e)$  in the original expression, indicating that this term is redundant and can be eliminated.

# Eager Redundancy Elimination (ERE)

- Example:  $(a \vee b \vee \neg c) \wedge (d \vee e \vee c) \wedge \cancel{(a \vee b \vee d \vee e)}$




- **For all  $x \in \phi$  in parallel**,  $\text{resolve}(x)$  to obtain the resolvents



# Eager Redundancy Elimination (ERE)

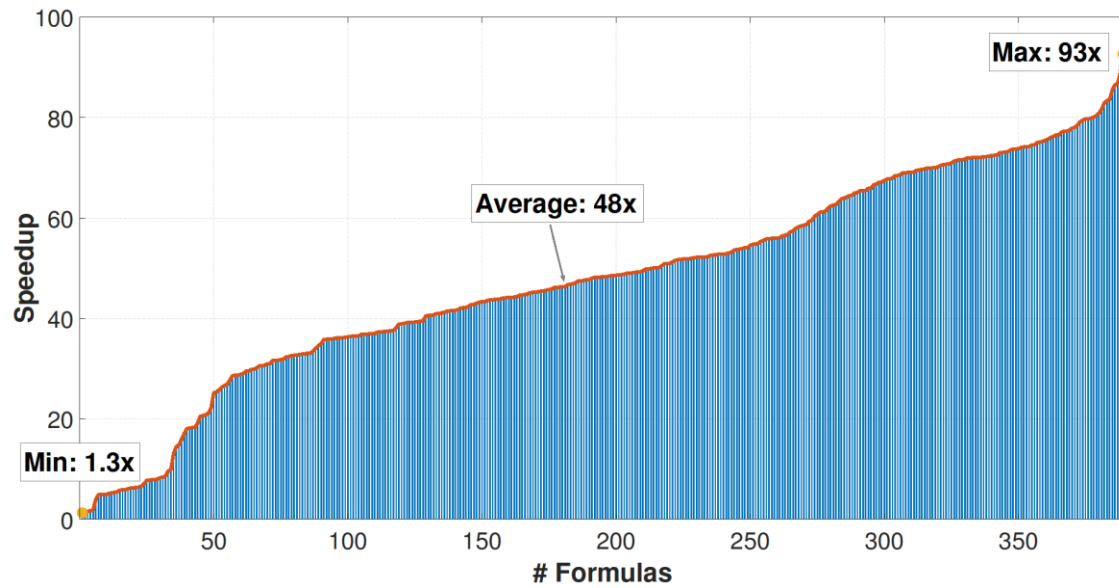
- Example:  $(a \vee b \vee \neg c) \wedge (d \vee e \vee c) \wedge \cancel{(a \vee b \vee d \vee e)}$




- **For all  $x \in \phi$  in parallel**,  $\text{resolve}(x)$  to obtain the resolvents
- For each resolvent, search for an equivalent clause in parallel on the fly

# Benchmarks

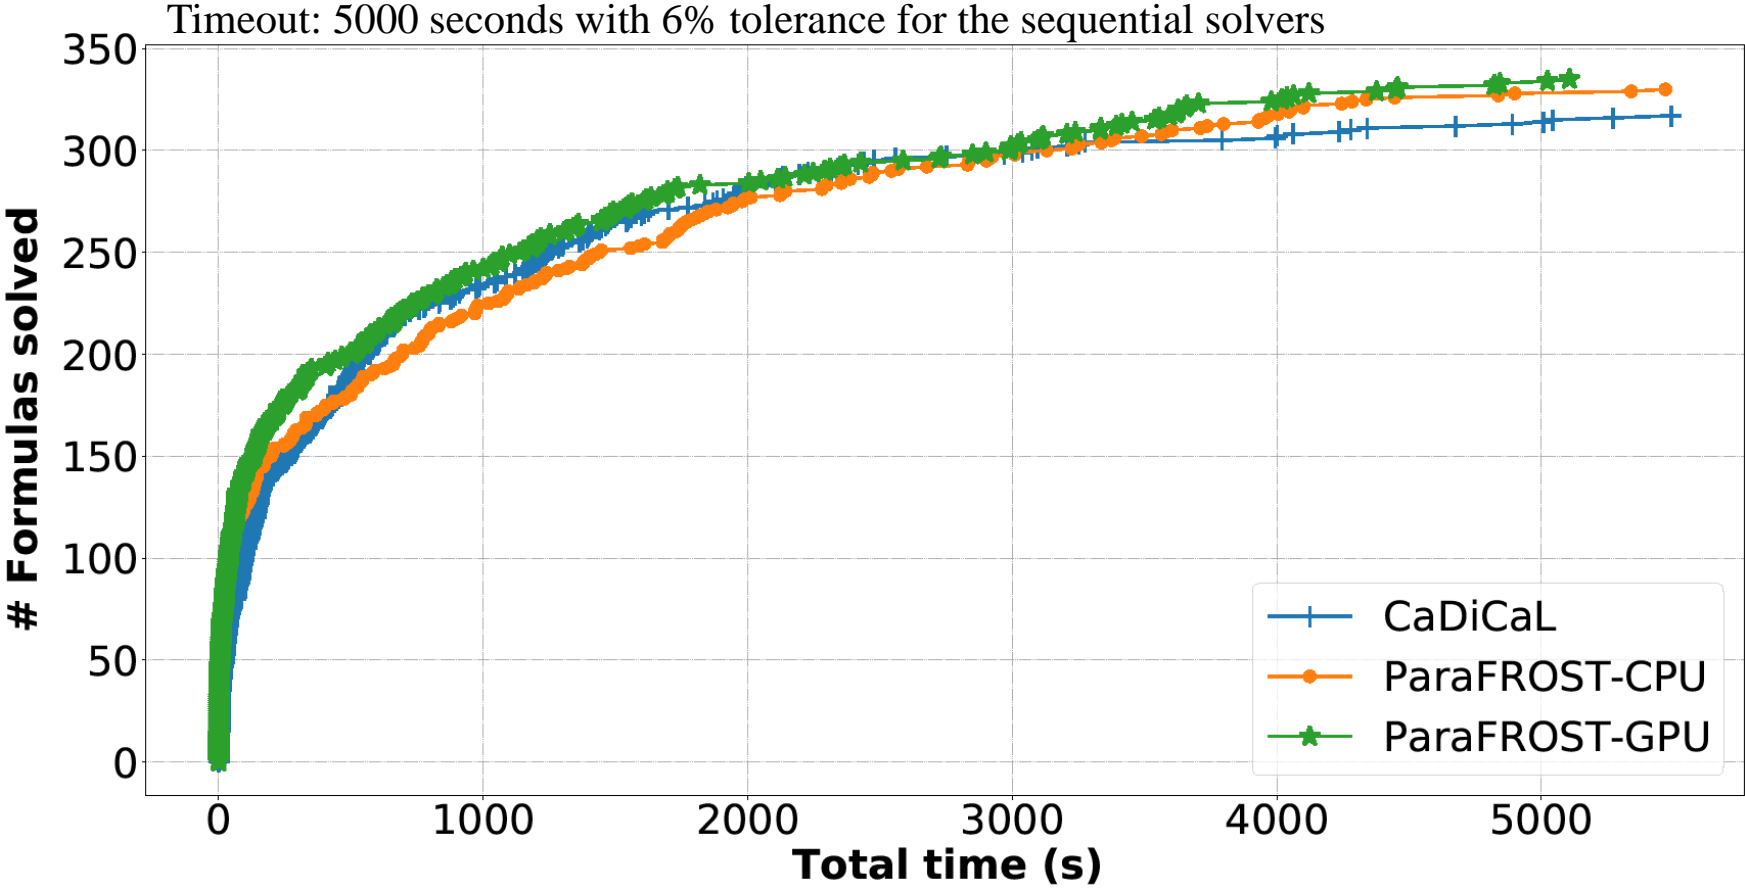
Parallel GC vs a sequential version



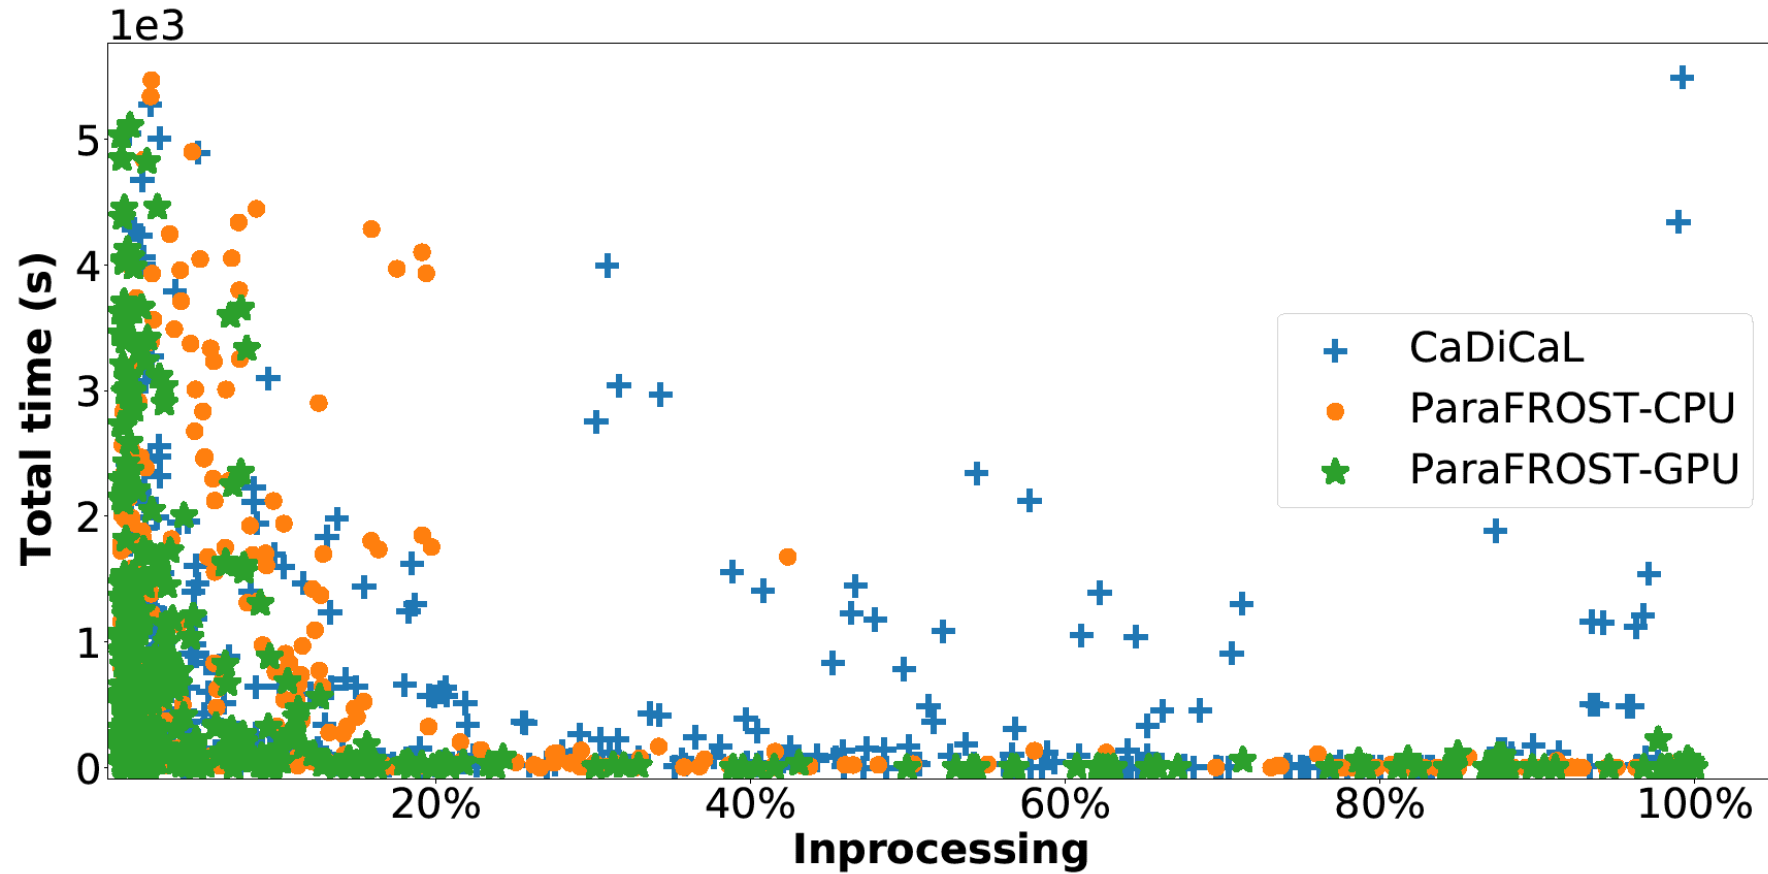
3-Phase BVE vs atomic version



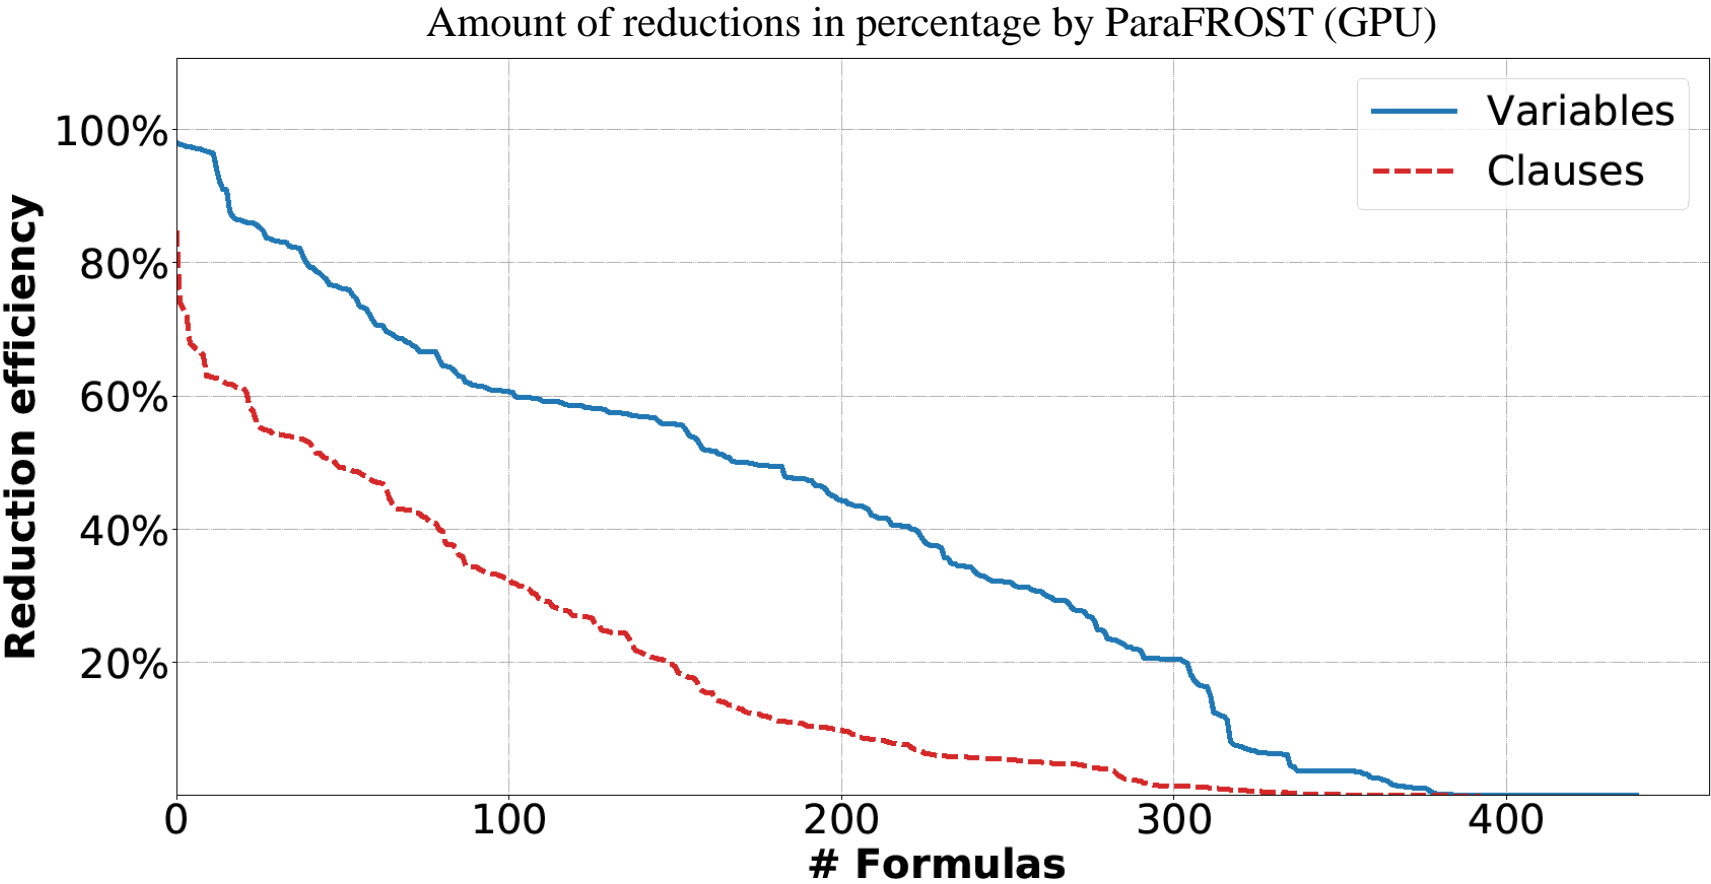
# Benchmarks



# Benchmarks



# Benchmarks



# Summary

- We implemented a new GPU solver with space-efficient data structure
- A parallel garbage collector was introduced for the first time for SAT applications
- The new parallel BVE was twice as fast and results can be reproduced
- Future work may include generating proofs for the GPU work (idea by Armin) and implementing more simplifications on the GPU

# Thank you

Artifact: <https://gears.win.tue.nl/software/parafrost>

ParaFROST: <https://github.com/muhos/ParaFROST>