


The Cilkprof Scalability Profiler



Tao B. Schardl Bradley C. Kuszmaul I-Ting Angelina Lee*
William M. Leiserson Charles E. Leiserson

MIT Computer Science and Artificial Intelligence Laboratory

*Washington University in St. Louis

SPAA 2015

Quicksort

C++ quicksort:

```
01 void qsort(int64_t array[], size_t n,  
02           size_t l, size_t h) {  
03     // ... base case ...  
04     size_t part;  
05     part = partition(array, n, l, h);  
06     qsort(array, n, l, part);  
07     qsort(array, n, part, h);  
08  
09 }  
10  
11 int main(int argc, char* argv[]) {  
12     // ... initialization ...  
13     qsort(array, n, 0, n);  
14     // ... use array ...  
15     return 0;  
16 }
```

Parallel quicksort using Cilk

The `cilk_spawn` and `cilk_sync` keywords expose parallel work.

C++ quicksort:

```
01 void qsort(int64_t array[], size_t n,  
02           size_t l, size_t h) {  
03     // ... base case ...  
04     size_t part;  
05     part = partition(array, n, l, h);  
06     qsort(array, n, l, part);  
07     qsort(array, n, part, h);  
08  
09 }  
10  
11 int main(int argc, char* argv[]) {  
12     // ... initialization ...  
13     qsort(array, n, 0, n);  
14     // ... use array ...  
15     return 0;  
16 }
```

Cilk parallel quicksort:

```
01 void qsort(int64_t array[], size_t n,  
02           size_t l, size_t h) {  
03     // ... base case ...  
04     size_t part;  
05     part = partition(array, n, l, h);  
06     cilk_spawn qsort(array, n, l, part);  
07     qsort(array, n, part, h);  
08     cilk_sync;  
09 }  
10  
11 int main(int argc, char* argv[]) {  
12     // ... initialization ...  
13     qsort(array, n, 0, n);  
14     // ... use array ...  
15     return 0;  
16 }
```

Parallel quicksort using Cilk

The `cilk_spawn` and `cilk_sync` keywords expose parallel work.

- The `cilk_spawn` allows the two recursive `qsort` calls to execute in parallel.
- The `cilk_sync` waits for both recursive `qsort` calls return.

Cilk parallel quicksort:

```
01 void qsort(int64_t array[], size_t n,  
02           size_t l, size_t h) {  
03     // ... base case ...  
04     size_t part;  
05     part = partition(array, n, l, h);  
06     cilk_spawn qsort(array, n, l, part);  
07     qsort(array, n, part, h);  
08     cilk_sync;  
09 }  
10  
11 int main(int argc, char* argv[]) {  
12     // ... initialization ...  
13     qsort(array, n, 0, n);  
14     // ... use array ...  
15     return 0;  
16 }
```

Running Cilk parallel quicksort

The parallel quicksort code can be compiled and run similarly to its serial counterpart.

Using ICC:

```
$ icpc -O3 qsort.cpp -o qsort  
$ ./qsort -n 100000000
```

Using GCC:

```
$ g++ -O3 qsort.cpp -o qsort -fcilkplus  
$ ./qsort -n 100000000
```

Using Cilk Plus/LLVM:

```
$ clang++ -O3 qsort.cpp -o qsort -fcilkplus -ldl  
$ ./qsort -n 100000000
```

Running Cilk parallel quicksort

The parallel quicksort code can be compiled and run similarly to its serial counterpart.

Using ICC:

```
$ icpc -O3 qsort.cpp -o qsort
$ ./qsort -n 100000000
```

Using GCC:

```
$ g++ -O3 qsort.cpp -o qsort -fcilkplus
$ ./qsort -n 100000000
```

Using Cilk Plus/LLVM:

```
$ clang++ -O3 qsort.cpp -o qsort -fcilkplus -ldl
$ ./qsort -n 100000000
```

Questions: How fast is this code? Does it speed up on multiple processors? How parallel is it?

Prior art: Cilkview


Cilkview [HLL10] analyzes the scalability of a Cilk program.

```
$ icpc -O3 qsort.cpp -o qsort
$ cilkview --trials=all -- ./qsort -n 100000000
```

- Cilkview instruments the executable binary of the Cilk program (with a little help from ICC).
- Cilkview executes the program serially while keeping track of the logical series-parallel relationships between instructions.
- Cilkview incurs only a constant-factor slowdown over the program's serial running time.

Cilkview's output for qsort(100M)

Speedup of qsort




```
Work : 36,272,478,614 instructions
Span : 1,621,934,437 instructions
...
Parallelism : 22.36
...
```

Note: This output has been simplified for didactic purposes.

Work, span, and parallelism

Cilkview models the Cilk program performance as follows:

- Let T_P denote the execution time on P processors.
- **Work** is serial execution time T_1 .
- **Speedup** on P processors is $T_1/T_P \leq P$.
 - **Linear speedup** occurs when $T_1/T_P = P$.
- **Span** is critical-path length T_∞ .
- **Parallelism** is T_1/T_∞ , the maximum possible speedup.

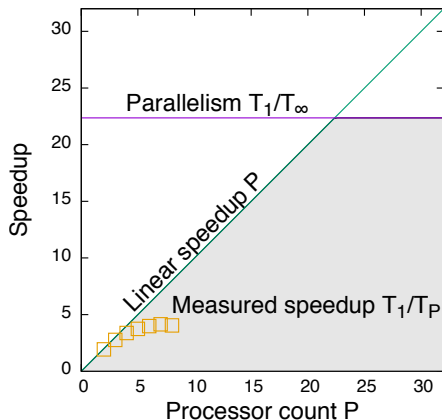


As a practical matter, a Cilk program should exhibit **parallel slackness** of 10 — it should exhibit $\geq 10P$ parallelism.

Work, span, and parallelism of qsort(100M)

Cilkview tells us that `qsort` does not exhibit much parallelism.

Speedup of qsort




Work : 36,272,478,614 instructions
Span : 1,621,934,437 instructions
...
Parallelism : 22.36
...

Where is the bottleneck?

Cilkview does not tell us where the scalability bottleneck is in `qsort`.

```
01 void qsort(int64_t array[], size_t n,
02           size_t l, size_t h) {
03     // ... base case ...
04     size_t part;
05     part = partition(array, n, l, h);
06     cilk_spawn qsort(array, n, l, part);
07     qsort(array, n, part, h);
08     cilk_sync;
09 }
10
11 int main(int argc, char* argv[]) {
12     // ... initialization ...
13     qsort(array, n, 0, n);
14     // ... use array ...
15     return 0;
16 }
```

Speedup of `qsort`



Our contribution: Cilkprof

Cilkprof profiles the parallelism of a Cilk program.

```
$ clang++ -O3 -g qsort.cpp -o qsort -fcilkplus -ldl -fcilktool-instr-c -lcilkprof  
$ ./qsort -n 100000000
```

- We modified the Cilk Plus/LLVM compiler to instrument functions, spawns, and syncs in a Cilk program.
- We implemented Cilkprof as a library to link into an instrumented Cilk program.
- Running the instrumented program linked with Cilkprof produces a spreadsheet attributing portions of the program's work and span to different call sites.
 - *No user interface*

Cilkprof's output for qsort(100M)

Cilkprof gathers work and span measurements for every call site.

<i>File</i>	<i>Line</i>	<i>Top-caller work on work</i>	<i>Local work on work</i>	<i>Top-caller span on span</i>	<i>Local span on span</i>
qsort.cpp	5	0.6	10.4	0.6	3.3
qsort.cpp	6	1.5	3.2	0.0	0.0
qsort.cpp	7	14.3	2.7	0.0	0.0
qsort.cpp	13	16.3	0.0	3.3	0.0

Cilkprof's output for qsort(100M)

Cilkprof gathers work and span measurements for every call site.

	<i>On work (gigacycles)</i>		<i>On span (gigacycles)</i>	
	<i>Top-caller work</i>	<i>Local work</i>	<i>Top-caller span</i>	<i>Local span</i>
01 void qsort(<i>/*...*/</i>) {				
02 <i>/* ... base case ... */</i>				
03 part = partition(<i>/*...*/</i>);	0.6	10.4	0.6	3.3
04 cilk_spawn qsort(<i>/*...*/</i>);	1.5	3.2	0.0	0.0
05 qsort(<i>/*...*/</i>);	14.3	2.7	2.8	0.0
06 cilk_sync;				
07 }				
08				
09 int main(<i>/*...*/</i>) {				
10 <i>// ... initialization ...</i>				
11 qsort(<i>/*...*/</i>);	16.3	0.0	3.3	0.0
12 <i>// ... use array ...</i>				
13 }				

Compiler instrumentation

We used compiler instrumentation for its efficiency.

	<i>Instrumentation</i>		
	<i>Sampling</i>	<i>Binary</i>	<i>Compiler</i>
<i>Example tools</i>	gprof, pprof, perf	valgrind, DynamoRio, Cilkview (via Pin)	gcov, tsan, Cilkprof
<i>Overhead</i>	✓ None	× Lots	~ Some
<i>Properties</i>	✓ No recompilation necessary. × Statistical; don't know how to measure span via sampling.	✓ No recompilation necessary. ~ Creates measurement error, but can still measure parallelism well enough.	~ Requires recompilation. ~ Creates measurement error, but can still measure parallelism well enough.

Cilkprof's algorithm

Cilkprof implements an efficient serial algorithm.

- Cilkprof profiles a Cilk program with work T_1 in $\Theta(T_1)$ time.
 - *Constant overhead*, independent of number of call sites.
- Remarkably, in the same asymptotic time Cilkview takes to measure work and span for the whole program, Cilkprof measures work and span for *every call site*.
 - For video encoding, in the asymptotic time Cilkview takes to measure two values, Cilkprof can measure those values for ≈ 3000 call sites.

Cilkprof's empirical performance

Cilkprof is efficient in practice.

<i>Benchmark</i>	<i>Description</i>	<i>Overhead</i>
mm	Matrix multiplication	0.99
dedup	Compression	1.03
lu	LU decomposition	1.04
strassen	Strassen	1.06
heat	Heat diffusion	1.07
cilksort	Mergesort	1.08
pbfs	Breadth-first search	1.10
fft	Fast Fourier transform	1.15
quicksort	Quicksort	1.20
nqueens	<i>n</i> -Queens	1.27
ferret	Image similarity	2.04
leiserchess	Game-tree search	3.72
collision	Collision detection	4.37
cholesky	Cholesky decomposition	4.54
hevc	H265 video coding	6.25
fib	Fibonacci	7.36

Cilkprof incurs the following overheads:

- A geometric-mean slowdown of $1.9\times$.
- A maximum slowdown of $7.4\times$.

Cilkprof performs favorably to similar debugging tools.

Outline

- 1 Case study: qsort
- 2 Case study: pbfs
- 3 Profiling the work and span
- 4 Conclusion

Outline

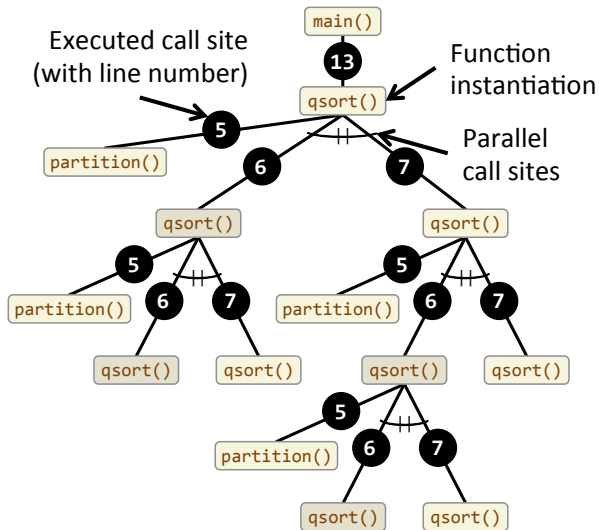
- 1 Case study: qsort
- 2 Case study: pbf s
- 3 Profiling the work and span
- 4 Conclusion

Cilkprof's profile for qsort


	<i>On work (gigacycles)</i>		<i>On span (gigacycles)</i>	
	<i>Top-caller work</i>	<i>Local work</i>	<i>Top-caller span</i>	<i>Local span</i>
01 void qsort(<i>/*...*/</i>) {				
02 <i>/* ... base case ... */</i>				
03 part = partition(<i>/*...*/</i>);	0.6	10.4	0.6	3.3
04 cilk_spawn qsort(<i>/*...*/</i>);	1.5	3.2	0.0	0.0
05 qsort(<i>/*...*/</i>);	14.3	2.7	2.8	0.0
06 cilk_sync;				
07 }				
08				
09 int main(<i>/*...*/</i>) {				
10 <i>// ... initialization ...</i>				
11 qsort(<i>/*...*/</i>);	16.3	0.0	3.3	0.0
12 <i>// ... use array ...</i>				
13 }				

Invocation trees

Cilkprof interprets the program execution in terms of its *invocation tree*.



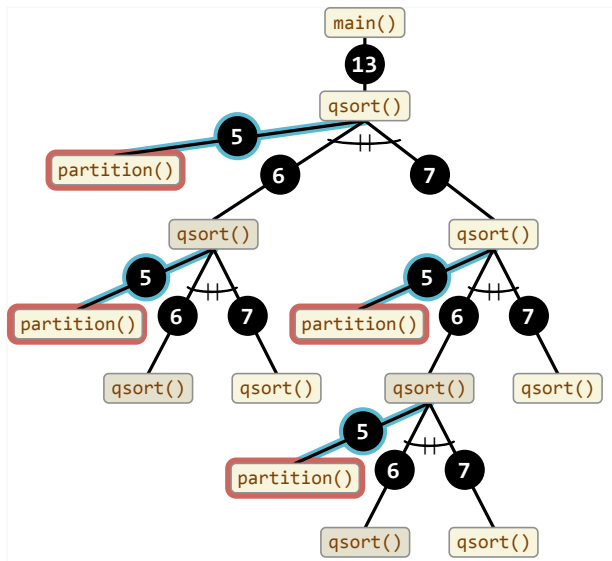
Separately profiling work and span



Cilkprof breaks down both the work and the span of the program.

- Instantiations **on the work** are all instantiations in the computation.
- Instantiations **on the span** are the instantiations on the critical path.

Handling multiple executions of a call site




How does Cilkgprof aggregate measurements from multiple executions of the same call site?

Simple idea: Just add them.

Problem: This strategy overcounts measurements for recursive functions.

Handling multiple executions of a call site: local




Idea: Measure every execution of each call site, but only record the *local* computation of the instantiations.

- Exclude the computation performed by child instantiations.

This is similar to `gprof`'s "self time."

Handling multiple executions of a call site: local



Idea: Measure every execution of each call site, but only record the **local** computation of the instantiations.

- Exclude the computation performed by child instantiations.

This is similar to gprof's "self time."

Finding the scalability bottleneck in qsort

	<i>On work (gigacycles)</i>		<i>On span (gigacycles)</i>	
	<i>Top-caller work</i>	<i>Local work</i>	<i>Top-caller span</i>	<i>Local span</i>
01 void qsort(<i>/*...*/</i>) {				
02 <i>/* ... base case ... */</i>				
03 part = partition(<i>/*...*/</i>);	0.6	10.4	0.6	3.3
04 cilk_spawn qsort(<i>/*...*/</i>);	1.5	3.2	0.0	0.0
05 qsort(<i>/*...*/</i>);	14.3	2.7	2.8	0.0
06 cilk_sync;				
07 }				
08				
09 int main(<i>/*...*/</i>) {				
10 <i>// ... initialization ...</i>				
11 qsort(<i>/*...*/</i>);	16.3	0.0	3.3	0.0
12 <i>// ... use array ...</i>				
13 }				

Finding the scalability bottleneck in qsort

Property: For the topmost qsort instantiation, its **top-caller work** equals its local work plus **the total top-caller work of its children**.

```

01 void qsort(/*...*/) {
02     /* ... base case ... */
03     part = partition(/*...*/);
04     cilk_spawn qsort(/*...*/);
05     qsort(/*...*/);
06     cilk_sync;
07 }
08
09 int main(/*...*/) {
10     // ... initialization ...
11     qsort(/*...*/);
12     // ... use array ...
13 }

```

<i>On work (gigacycles)</i>		<i>On span (gigacycles)</i>	
<i>Top-caller work</i>	<i>Local work</i>	<i>Top-caller span</i>	<i>Local span</i>

0.6	10.4	0.6	3.3
1.5	3.2	0.0	0.0
14.3	2.7	2.8	0.0
16.3	0.0	3.3	0.0

Finding the scalability bottleneck in qsort

Property: For the topmost qsort instantiation, its **top-caller work** equals its local work plus **the total local work of its children**.

```

01 void qsort(/*...*/) {
02     /* ... base case ... */
03     part = partition(/*...*/);
04     cilk_spawn qsort(/*...*/);
05     qsort(/*...*/);
06     cilk_sync;
07 }
08
09 int main(/*...*/) {
10     // ... initialization ...
11     qsort(/*...*/);
12     // ... use array ...
13 }

```

<i>On work (gigacycles)</i>		<i>On span (gigacycles)</i>	
<i>Top-caller work</i>	<i>Local work</i>	<i>Top-caller span</i>	<i>Local span</i>

0.6	10.4	0.6	3.3
1.5	3.2	0.0	0.0
14.3	2.7	2.8	0.0
16.3	0.0	3.3	0.0

Finding the scalability bottleneck in qsort

	<i>On work (gigacycles)</i>		<i>On span (gigacycles)</i>	
	<i>Top-caller work</i>	<i>Local work</i>	<i>Top-caller span</i>	<i>Local span</i>
01 void qsort(<i>/*...*/</i>) {				
02 <i>/* ... base case ... */</i>				
03 part = partition(<i>/*...*/</i>);	0.6	10.4	0.6	3.3
04 cilk_spawn qsort(<i>/*...*/</i>);	1.5	3.2	0.0	0.0
05 qsort(<i>/*...*/</i>);	14.3	2.7	2.8	0.0
06 cilk_sync;				
07 }				
08				
09 int main(<i>/*...*/</i>) {				
10 <i>// ... initialization ...</i>				
11 qsort(<i>/*...*/</i>);	16.3	0.0	3.3	0.0
12 <i>// ... use array ...</i>				
13 }				

Finding the scalability bottleneck in qsort

Property: For the topmost qsort instantiation, its **top-caller span** equals its local span plus **the total top-caller span of its children**.

```

01 void qsort(/*...*/) {
02     /* ... base case ... */
03     part = partition(/*...*/);
04     cilk_spawn qsort(/*...*/);
05     qsort(/*...*/);
06     cilk_sync;
07 }
08
09 int main(/*...*/) {
10     // ... initialization ...
11     qsort(/*...*/);
12     // ... use array ...
13 }

```

<i>On work (gigacycles)</i>		<i>On span (gigacycles)</i>	
<i>Top-caller work</i>	<i>Local work</i>	<i>Top-caller span</i>	<i>Local span</i>
0.6	10.4	0.6	3.3
1.5	3.2	0.0	0.0
14.3	2.7	2.8	0.0
16.3	0.0	3.3	0.0

Finding the scalability bottleneck in qsort

Property: For the topmost qsort instantiation, its **top-caller span** equals its local span plus **the total local span of its children**.

```

01 void qsort(/*...*/) {
02     /* ... base case ... */
03     part = partition(/*...*/);
04     cilk_spawn qsort(/*...*/);
05     qsort(/*...*/);
06     cilk_sync;
07 }
08
09 int main(/*...*/) {
10     // ... initialization ...
11     qsort(/*...*/);
12     // ... use array ...
13 }

```

<i>On work (gigacycles)</i>		<i>On span (gigacycles)</i>	
<i>Top-caller work</i>	<i>Local work</i>	<i>Top-caller span</i>	<i>Local span</i>


0.6	10.4	0.6	3.3
1.5	3.2	0.0	0.0
14.3	2.7	2.8	0.0
16.3	0.0	3.3	0.0

Outline

- 1 Case study: qsort
- 2 Case study: pbfs**
- 3 Profiling the work and span
- 4 Conclusion

Months of struggle with pbf s...

We were stumped for months trying to pinpoint a scalability bottleneck in `pbf s`, a Cilk program to perform parallel breadth-first search.



- A back-of-the-envelope calculation suggests that `pbf s` can achieve parallelism of 200–400.
- Cilkview measured the parallelism of `pbf s` to be only 12.

Solved in two hours with Cilkprof

Using a prototype of Cilkprof, we were able to pinpoint and fix the scalability bottleneck in `pbfs` in 2 hours.

- Sorting Cilkprof's output revealed three main contributors to the span:

`parseBinaryFile()`: Routine to read the input graph.

`Graph()`: Constructor for the graph data structure.

`pbfs_proc_Node()`: Base case of the primary recursive routine.

- We found and fixed a mistuned constant in `pbfs_proc_Node()`, improving the parallelism of `pbfs` by a factor of 5.

Outline

- 1 Case study: qsort
- 2 Case study: pbf s
- 3 Profiling the work and span**
- 4 Conclusion

Computing work and span: variables

Cilkprof augments an algorithm that incrementally computes the work and span of a whole program execution.

For each instantiation F :

Variables

Work $F.w$

$F.p$

Span $F.l$

$F.c$

Invariant: When F returns:

- $F.w$ stores its work.
- $F.p$ stores its span.

Computing work and span: variables

Gilkprof augments an algorithm that incrementally computes the work and span of a whole program execution.

For each instantiation F :

Variables

Work $F.w$

$F.p$

Span $F.l$

$F.c$

Invariant: When F returns:

- $F.w$ stores its work.
- $F.p$ stores its span.

Corollary: When `main` returns:

- `main.w` stores the computation's work.
- `main.p` stores the computation's span.

Computing work and span: sum and max

Cilkprof computes work and span incrementally by taking sums and (something like) maxes of the work and span variables.

<i>F</i> spawns or calls <i>G</i> : 1 let <i>G.w</i> = 0 2 let <i>G.p</i> = 0 3 let <i>G.l</i> = 0 4 let <i>G.c</i> = 0	Called <i>G</i> returns to <i>F</i> : 5 <i>G.p</i> += <i>G.c</i> 6 <i>F.w</i> += <i>G.w</i> 7 <i>F.c</i> += <i>G.p</i>
Spawned <i>G</i> returns to <i>F</i> : 8 <i>G.p</i> += <i>G.c</i> 9 <i>F.w</i> += <i>G.w</i> 10 if <i>F.c</i> + <i>G.p</i> > <i>F.l</i> 11 <i>F.l</i> = <i>G.p</i> 12 <i>F.p</i> += <i>F.c</i> 13 <i>F.c</i> = 0	<i>F</i> syncs: 14 if <i>F.c</i> > <i>F.l</i> 15 <i>F.p</i> += <i>F.c</i> 16 else 17 <i>F.p</i> += <i>F.l</i> 18 <i>F.c</i> = 0 19 <i>F.l</i> = 0

This pseudocode performs the following types of operations:

- Initialization
- Sum
- Max (or something like it)

Computing work and span: sum and max

Cilkprof computes work and span incrementally by taking sums and (something like) maxes of the work and span variables.

<i>F</i> spawns or calls <i>G</i> : 1 let <i>G.w</i> = 0 2 let <i>G.p</i> = 0 3 let <i>G.l</i> = 0 4 let <i>G.c</i> = 0	Called <i>G</i> returns to <i>F</i> : 5 <i>G.p</i> += <i>G.c</i> 6 <i>F.w</i> += <i>G.w</i> 7 <i>F.c</i> += <i>G.p</i>
Spawned <i>G</i> returns to <i>F</i> : 8 <i>G.p</i> += <i>G.c</i> 9 <i>F.w</i> += <i>G.w</i> 10 if <i>F.c</i> + <i>G.p</i> > <i>F.l</i> 11 <i>F.l</i> = <i>G.p</i> 12 <i>F.p</i> += <i>F.c</i> 13 <i>F.c</i> = 0	<i>F</i> syncs: 14 if <i>F.c</i> > <i>F.l</i> 15 <i>F.p</i> += <i>F.c</i> 16 else 17 <i>F.p</i> += <i>F.l</i> 18 <i>F.c</i> = 0 19 <i>F.l</i> = 0

This pseudocode performs the following types of operations:

- Initialization
- Sum
- Max (or something like it)

Computing work and span: sum and max

Cilkprof computes work and span incrementally by taking sums and (something like) maxes of the work and span variables.

<p><i>F</i> spawns or calls <i>G</i>:</p> <pre> 1 let <i>G.w</i> = 0 2 let <i>G.p</i> = 0 3 let <i>G.l</i> = 0 4 let <i>G.c</i> = 0 </pre>	<p>Called <i>G</i> returns to <i>F</i>:</p> <pre> 5 <i>G.p</i> += <i>G.c</i> 6 <i>F.w</i> += <i>G.w</i> 7 <i>F.c</i> += <i>G.p</i> </pre>
<p>Spawned <i>G</i> returns to <i>F</i>:</p> <pre> 8 <i>G.p</i> += <i>G.c</i> 9 <i>F.w</i> += <i>G.w</i> 10 if <i>F.c</i> + <i>G.p</i> > <i>F.l</i> 11 <i>F.l</i> = <i>G.p</i> 12 <i>F.p</i> += <i>F.c</i> 13 <i>F.c</i> = 0 </pre>	<p><i>F</i> syncs:</p> <pre> 14 if <i>F.c</i> > <i>F.l</i> 15 <i>F.p</i> += <i>F.c</i> 16 else 17 <i>F.p</i> += <i>F.l</i> 18 <i>F.c</i> = 0 19 <i>F.l</i> = 0 </pre>

This pseudocode performs the following types of operations:

- Initialization
- Sum
- Max (or something like it)

Computing work and span: sum and max

Cilkprof computes work and span incrementally by taking sums and (something like) maxes of the work and span variables.

<i>F</i> spawns or calls <i>G</i> : 1 let <i>G.w</i> = 0 2 let <i>G.p</i> = 0 3 let <i>G.l</i> = 0 4 let <i>G.c</i> = 0	Called <i>G</i> returns to <i>F</i> : 5 <i>G.p</i> += <i>G.c</i> 6 <i>F.w</i> += <i>G.w</i> 7 <i>F.c</i> += <i>G.p</i>
Spawned <i>G</i> returns to <i>F</i> : 8 <i>G.p</i> += <i>G.c</i> 9 <i>F.w</i> += <i>G.w</i> 10 if <i>F.c</i> + <i>G.p</i> > <i>F.l</i> 11 <i>F.l</i> = <i>G.p</i> 12 <i>F.p</i> += <i>F.c</i> 13 <i>F.c</i> = 0	<i>F</i> syncs: 14 if <i>F.c</i> > <i>F.l</i> 15 <i>F.p</i> += <i>F.c</i> 16 else 17 <i>F.p</i> += <i>F.l</i> 18 <i>F.c</i> = 0 19 <i>F.l</i> = 0

This pseudocode performs the following types of operations:

- Initialization
- Sum
- Max (or something like it)

Computing work and span profiles

Key Idea: Attach a profile to each work and span variable.

For each instantiation F :

Variables

Work $F.w$

Span $F.p$
 $F.l$
 $F.c$

Invariant: When F returns:

- $F.w$ stores its work.
- $F.p$ stores its span.

Corollary: When `main` returns:

- `main.w` stores the computation's work.
- `main.p` stores the computation's span.

Computing work and span profiles

Key Idea: Attach a profile to each work and span variable.

For each instantiation F :

	<i>Variables</i>	<i>Profiles</i>
<i>Work</i>	$F.w$	$F.w.prof$
	$F.p$	$F.p.prof$
<i>Span</i>	$F.l$	$F.p.prof$
	$F.c$	$F.c.prof$

Invariant: When F returns:

- $F.w$ stores its work.
- $F.p$ stores its span.

Corollary: When `main` returns:

- `main.w` stores the computation's work.
- `main.p` stores the computation's span.

Computing work and span profiles

Key Idea: Attach a profile to each work and span variable.

For each instantiation F :

	<i>Variables</i>	<i>Profiles</i>
<i>Work</i>	$F.w$	$F.w.prof$
	$F.p$	$F.p.prof$
<i>Span</i>	$F.l$	$F.p.prof$
	$F.c$	$F.c.prof$

Invariant: When F returns:

- $F.w$ stores its work.
- $F.p$ stores its span.
- $F.w.prof$ stores the *profile* of F 's work.
- $F.p.prof$ stores the *profile* of F 's span.

Corollary: When `main` returns:

- `main.w` stores the computation's work.
- `main.p` stores the computation's span.

Computing work and span profiles

Key Idea: Attach a profile to each work and span variable.

For each instantiation F :

	<i>Variables</i>	<i>Profiles</i>
<i>Work</i>	$F.w$	$F.w.prof$
	$F.p$	$F.p.prof$
<i>Span</i>	$F.l$	$F.p.prof$
	$F.c$	$F.c.prof$

Invariant: When F returns:

- $F.w$ stores its work.
- $F.p$ stores its span.
- $F.w.prof$ stores the *profile* of F 's work.
- $F.p.prof$ stores the *profile* of F 's span.

Corollary: When `main` returns:

- `main.w` stores the computation's work.
- `main.p` stores the computation's span.
- `main.w.prof` stores the *profile* of the computation *on the work*.
- `main.p.prof` stores the *profile* of the computation *on the span*.

Maintaining a profile

The `prof` data structure supports the following operations:

- `INIT()`: Initialize a `prof R` to be empty.
- `ASSIGN(R, R')`: Replace the contents of `prof R` with that of `prof R'`, then discard the contents of R' .
- `UNION(R, R')`: Update the `prof R` element-wise with the contents of the R' , then discard the contents of R' .
- `UPDATE($R, \langle s, v \rangle$)`: If no record v' associated with call site s already exists in R , store $\langle s, v \rangle$ into R . Otherwise, store $\langle s, v' + v \rangle$.

Computing work and span profiles: UNION and ASSIGN

Called G returns to F :

```

1   $G.p += G.c$ 
2
3   $F.w += G.w$ 
4
5   $F.c += G.p$ 
6

```

Spawned G returns to F :

```

7  if  $F.c + G.p > F.l$ 
8       $F.l = G.p$ 
9
10      $F.p += F.c$ 
11
12      $F.c = 0$ 

```

When variables are summed, Cilkprof uses UNION to combine their profiles.

When variables are maxed, Cilkprof uses ASSIGN to discard the profile of the smaller variable.

- UNION is also used when the max-like pseudocode adds variables.

Computing work and span profiles: UNION and ASSIGN

Called G returns to F :

```

1   $G.p += G.c$ 
2   $\text{UNION}(G.p.\text{prof}, G.c.\text{prof})$ 
3   $F.w += G.w$ 
4   $\text{UNION}(F.w.\text{prof}, G.w.\text{prof})$ 
5   $F.c += G.p$ 
6   $\text{UNION}(F.c.\text{prof}, G.p.\text{prof})$ 

```

Spawned G returns to F :

```

7  if  $F.c + G.p > F.l$ 
8      $F.l = G.p$ 
9
10      $F.p += F.c$ 
11
12      $F.c = 0$ 

```

When variables are summed, Cilkprof uses **UNION** to combine their profiles.

When variables are maxed, Cilkprof uses **ASSIGN** to discard the profile of the smaller variable.

- **UNION** is also used when the max-like pseudocode adds variables.

Computing work and span profiles: UNION and ASSIGN

Called G returns to F :

```

1   $G.p += G.c$ 
2  UNION( $G.p.prof, G.c.prof$ )
3   $F.w += G.w$ 
4  UNION( $F.w.prof, G.w.prof$ )
5   $F.c += G.p$ 
6  UNION( $F.c.prof, G.p.prof$ )

```

Spawned G returns to F :

```

7  if  $F.c + G.p > F.l$ 
8       $F.l = G.p$ 
9      ASSIGN( $F.l.prof, G.p.prof$ )
10  $F.p += F.c$ 
11 UNION( $F.p.prof, F.c.prof$ )
12  $F.c = 0$ 

```

When variables are summed ,
Cilkprof uses UNION to combine their profiles.

When variables are maxed ,
Cilkprof uses ASSIGN to discard the profile of the smaller variable.

- UNION is also used when the max-like pseudocode adds variables.

Cilkprof asymptotic running time

Theorem: If each operation on a **prof** data structure takes $\Theta(1)$ time, then Cilkprof executes a given Cilk program with work T_1 in $\Theta(T_1)$ total time.

An intuitive prof data structure

Intuitively, a **prof** is just a hashtable mapping call sites to work and span values.

A	(20,9)
C	(6,6)
E	(2,2)
D	(10,7)
B	(2,2)

Problem: Combining two hashtables must be done element-wise, which takes linear time. This data structure increases Cilkprof's overhead to $\Theta(S)$ in the worst case.

An intuitive prof data structure

Intuitively, a **prof** is just a hashtable mapping call sites to work and span values.

A	(20,9)
C	(6,6)
E	(2,2)
D	(10,7)
B	(2,2)

Problem: Combining two hashtables must be done element-wise, which takes linear time. This data structure increases Cilkprof's overhead to $\Theta(S)$ in the worst case.

- Merging hashtables is wasteful when they don't contain many entries.

An intuitive prof data structure

Intuitively, a **prof** is just a hashtable mapping call sites to work and span values.


A	(20,9)
C	(6,6)
E	(2,2)
D	(10,7)
B	(2,2)

Problem: Combining two hashtables must be done element-wise, which takes linear time. This data structure increases Cilkprof's overhead to $\Theta(S)$ in the worst case.

- Merging hashtables is wasteful when they don't contain many entries.
- If the table contains few entries, then it's more efficient to simply log what gets added to the table in, e.g., a linked-list.

The actual prof data structure

The **prof** data structure transforms between a linked-list and a hashtable.



- When the linked-list gets too large — $\Theta(S)$ entries — convert it to a hashtable.
- An amortization argument justifies that all operations on this **prof** data structure take $\Theta(1)$ amortized time.

Outline

- 1 Case study: qsort
- 2 Case study: pbfs
- 3 Profiling the work and span
- 4 Conclusion**

Summary

Cilkprof is a scalability profiler for Cilk programs.

- We modified the Cilk Plus/LLVM compiler to instrument functions, spawns, and syncs in a Cilk program.
Available from <https://github.com/neboat/{llvm,clang}>.
- We implemented Cilkprof as a library to link into an instrumented Cilk program.
Available from <https://github.com/neboat/cilktools>.
- Running the instrumented program linked with Cilkprof produces a spreadsheet attributing portions of the program's work and span to each call site.
- Cilkprof profiles a Cilk program with work T_1 in $\Theta(T_1)$ time.
- Cilkprof incurs a geometric-mean slowdown of $1.9\times$ and a maximum slowdown of $7.4\times$.

