



University
of
St Andrews



Skel: A Streaming Process-based Skeleton Library for Erlang

Archibald Elliott¹ Christopher Brown¹ Marco Danelutto²
Kevin Hammond¹

Email: ashe@st-andrews.ac.uk

¹School of Computer Science, University of St Andrews, Scotland, UK.

²Dept. Computer Science, University of Pisa, Pisa, Italy.

IFL 2012 - Oxford

PARAPHRASE



- ▶ Why we need Parallelism
- ▶ Skeletons are good Abstractions
- ▶ Skeletons in Erlang
- ▶ `skel`'s good Speedups

PARAPHRASE



1. The single-core processor is almost completely obsolete
2. Hardware systems are rapidly moving towards many- and mega-core

By 2019 there will be millions of cores in home desktop machines

– Joe Armstrong

3. Software systems are still not ready:
 - ▶ Programming languages have not caught up
 - ▶ Software practices have not caught up
 - ▶ Programmers have not caught up
4. We need to make programming parallel systems easy

PARAPHRASE

Race Conditions



What happens when you use Pthreads

PARAPHRASE





- ▶ Most Programmers are taught to program sequentially
- ▶ Modifying sequential code will not scale
- ▶ Typical concurrency techniques will not scale
- ▶ Fundamentally, current approaches are too low-level:
 - ▶ You can't program effectively while thinking about deadlocks, race conditions, synchronisation, non-determinism etc.
 - ▶ You can't program effectively directly with threads, message passing, mutexes or shared memory.
 - ▶ You can only program effectively with a different mindset

A graphic at the bottom of the slide with a green, monochromatic aesthetic. It features a central perspective view of a road or tunnel leading into the distance. The word "PARAPHRASE" is written in large, bold, white capital letters across the middle of the image. On either side of the central image are several columns of small, white text that appear to be code snippets or data logs, similar to the terminal output seen in the background image.

PARAPHRASE



- ▶ We need to provide a set of abstractions for the programmer;
- ▶ They need to become second-nature;
- ▶ They need to make it easy to introduce parallelism;
- ▶ They need to make it easy to tune parallelism to gain maximum speedup;

PARAPHRASE



Functional programming can provide the correct abstractions

However, languages take fundamentally different approaches

- ▶ Haskell (GpH) is too implicit:
 - ▶ `par :: a -> b -> b`
 - ▶ `pseq :: a -> b -> b`
- ▶ Erlang is too explicit:
 - ▶ `spawn`
 - ▶ `!` and `receive`


PARAPHRASE



```
skel:run(Skeleton, InputItems).  
% -> OutputItems
```


- ▶ Skeleton – a skeleton
- ▶ InputItems – items to be processed
- ▶ OutputItems – items that have been processed

PARAPHRASE



```
skel:run({seq, fun (X) -> X+1 end},
         [1,2,3,4,5,6]).
% -> [2,3,4,5,6,7]
```


PARAPHRASE



```


Inc      = {seq, fun (X) -> X+1 end},
Double  = {seq, fun (X) -> X*2 end},
skel:run({pipe, [Inc, Double]},
         [1,2,3,4,5,6]).
% -> [4,6,8,10,12,14]

```



```
Inc = {seq, fun(X) -> X+1 end},
skel:run({farm, Inc, 3},
        [1,2,3,4,5,6]).
% -> [2,5,3,6,4,7]
```


PARAPHRASE



```


Inc = {seq, fun(X)-> X+1 end},
skel:run({map, Inc,
         fun erlang:tuple_to_list/1,
         fun erlang:list_to_tuple/1},
        [{1,2},{3,4}]).
% -> [{2,3},{4,5}]

```



```
skel:run({reduce, fun(X,Y) -> X + Y end,
          fun erlang:tuple_to_list/1},
         [{1,2,3,4,5,6},{7,8,9,10,11,12}]).
% -> [21,57]
```

PARAPHRASE



```
Inc = {seq, fun(X) -> X+1 end},
skel:run({feedback, Inc,
         fun(X) -> X < 5 end},
        [1,2,3,4,5,6,7,8,9,10]).
% -> [5,6,7,8,9,10,11,5,5,5]
```

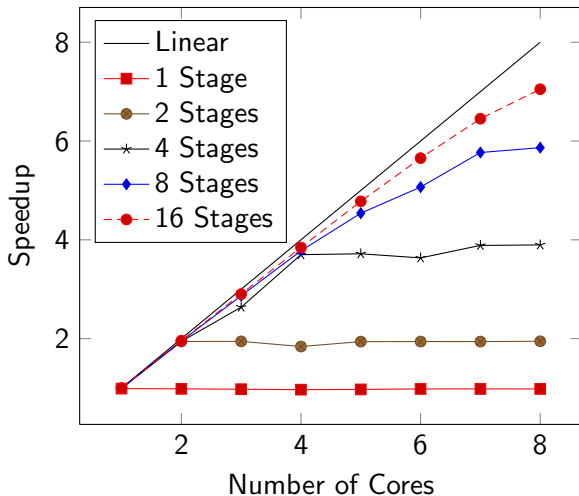
PARAPHRASE



- ▶ Regular task cost: 1ms
- ▶ Task cost $>$ communication cost
- ▶ Constant input number: 100,000
- ▶ 8 Cores: 2x Intel Xeon 4-Core X5355 2.66GHz
- ▶ Erlang R15B01


PARAPHRASE

Results: Pipe



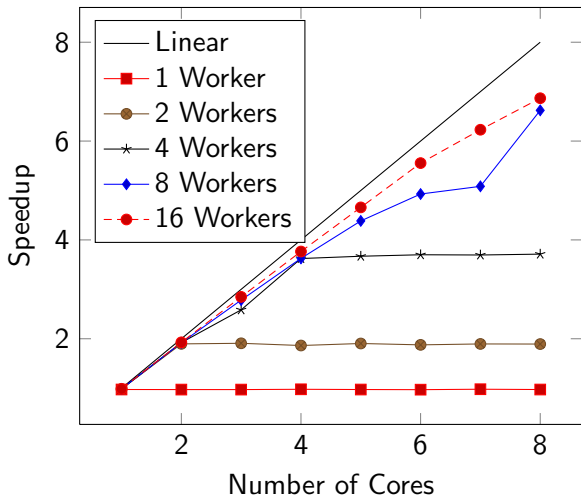
PARAPHRASE

Results: Farm



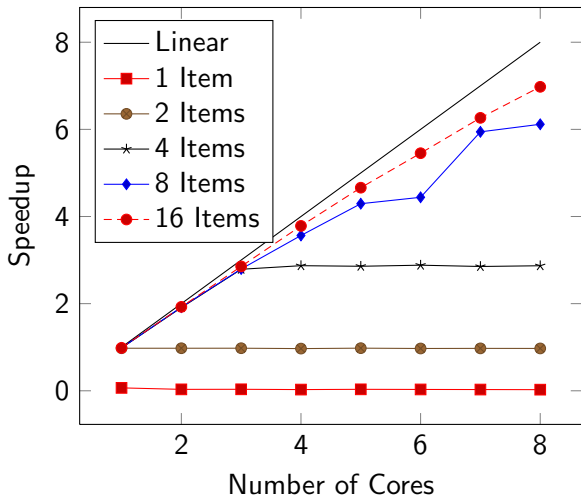
PARAPHRASE

Results: Map



PARAPHRASE

Results: Reduce



PARAPHRASE



- ▶ Adopting a Parallel Mindset
 - ▶ Functional Programming
 - ▶ Algorithmic Skeletons
- ▶ Erlang is a great fit.
 - ▶ Low-level enough for control
 - ▶ High-level enough to allow abstraction
- ▶ skel's speedups prove our implementation is good

PARAPHRASE



- ▶ More Benchmarks
- ▶ Better Speedups
- ▶ Higher-order Skeletons
 - ▶ Divide and Conquer
 - ▶ MapReduce
 - ▶ Genetic Algorithms
 - ▶ ...
 - ▶ Domain-specific Algorithms

PARAPHRASE



THANK YOU

<http://www.paraphrase-ict.eu>

@paraphrase_fp7

PARAPHRASE