

Multi-core programming with automatic parallelisation

Tim Harris and Satnam Singh Microsoft Research Cambridge



Moore's law; the free lunch



Source: table from http://www.intel.com/technology/mooreslaw/index.htm



The free lunch is over





Using multi-cores: space scheduling



- Partition resources between concurrent apps or VMs
- Machine consolidation
- Preserve responsiveness
- VMs for stronger isolation



Using multi-cores: control parallelism



- Find controlparallelism within an app
- Multiple threads in a shared address space
- Handle different clients
- GC, indexing, etc concurrent with application work



Using multi-cores: data parallelism



- Exploit data parallelism in large computations
- OpenMP
- Nested dataparallelism

What about existing client code?

Microsoft

- Centred around interactions with a single user...
 - No easy parallelism through multiple clients
- Not written with parallelism in mind...
 No explicit threading for performance
 No calls onto parallel libraries



Implicit parallelism







The key problems

• Getting the granularity right

• Keeping it transparent

```
for (int i = 0; i < 16; i++) {
    total += x[i];
}
```



- Haskell
- Limit study: is there anything out there?
- Making it more real: what does the granularity look like?
- Actual implementation: can we get a free lunch after all?

Introduction - Haskeliwar, Windows Interior Function



Microsoft^{*}

Research

Edit this page | Discuss this page | Page history | What links here | Related changes

Go Search

Introduction

Categories: Tutorials | Language

Haskell is a computer programming language. In particular, it is a *polymorphically typed, tazy,* purely functional language, quite different from most other programming languages. The language is named for Haskell Brooks Curry, whose work in mathematical logic serves as a foundation for thonal languages. Haskell is based on *lambda calculus*, hence the lambda we use as a logo.

Contents Initel

λ http://www.haskell.org/haskelliv#i/Introduction

Random page | Special pages

Haskell | Wiki community | Recent changes

A Introduction - HaskelWiki

"purely functional" : computing a value doesn't have side effects

> 3.7 Strong glue 3.8 Powerful abstractions 3.9 Built-in memory management 4 When C is better 4.1 Functional vs imperative 5 What is Haskell? 5 Ones among use functional programming?

A toy example: Fibonacci

100

× 90 WD

WIND a m

Microsoft'



Running the toy example

× 90 40 100

WIND a m

Microsoft^{*}





- Haskell
- Limit study: is there anything out there?
- Making it more real: what does the granularity look like?
- Actual implementation: can we get a free lunch after all?



Thunk lifecycle





Thunk dependencies





Potential parallelism







Limit study





- Haskell
- Limit study: is there anything out there?
- Making it more real: what does the granularity look like?
- Actual implementation: can we get a free lunch after all?

Which thunks to select?

Microsoft

- Select based on thunk allocation site
- Binary spark / no-spark decision
- Threshold for needed/not-needed (3/4)
- Consider total work for allocation site

Simulation, select by alloc site

100

w

× 90

Microsoft^{*}





- Haskell
- Limit study: is there anything out there?
- Making it more real: what does the granularity look like?
- Actual implementation: can we get a free lunch after all?

Avoiding duplicate work

100

Microsoft

Research

 By default GHC *does not* prevent concurrent evaluation of the same thunk

× 90

- But here the penalty is bad
 - We're selecting big thunks
 - Independence of thunk evaluation is compromised
- We add a level of locking to avoid duplication



Avoiding speculation of unsafe I/O

- Dynamically detect attempts to perform unsafe I/O while speculating
- Upon detection suspend speculation of current thunk
- Defined using new I/O action which has no side effect in application thread but suspends under speculation

Measured performance

100

Microsoft^{*}





Microsoft^{*}



Batcher's Parallel Sorter





Parallelism in the sorter



two :: ([a] -> [b]) -> [a] -> [b] two r = halve >-> **toBoth** r r >-> unhalve



Related work

- Parallel function programming [Hammond 94].
- Speculative evaluation in Multilisp [Osborne 90].
- par and seq [Roe 91].
- Parallel strategies [Trinder et. al. 98].
- Combining non-strictness with eager evaluation: Id90 and pH [Nikhil and Arvind 01].
- Fine grain concurrency: Monsson [Traub 91].
- Optimistic evaluation [Ennals and Peyton-Jones 03].
- GranSim [Loidl 98].



- Can we use a single program execution that avoids the need for a profile-collection phase?
- Impact of adjusting optimization level
- Use of feedback information to guide manual parallelization



Conclusions

- For selected programs an automatic speedup of 10-80% can be achieved
- Simulator can identify programs that are not amenable to automatic parallelization
- Not a silver bullet but does provide an interesting approach for parallelizing single application across 2, 3 or 4 cores