

Functioning Hardware from Functional Specifications

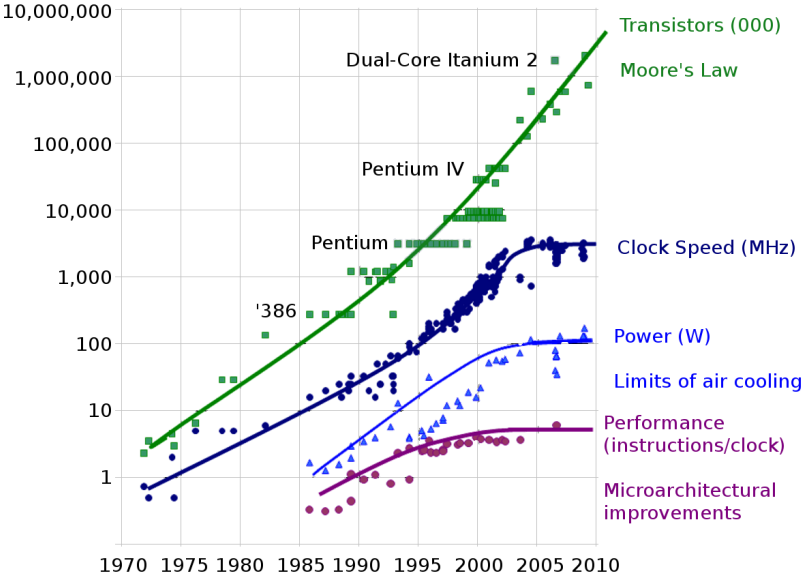
Stephen A. Edwards
Martha A. Kim
Richard Townsend
Kuangya Zhai
Lianne Lairmore

Columbia University

IBM PL Day, November 18, 2014

Where's my 10 GHz processor?


Intel CPU Trends



Sutter, *The Free Lunch is Over*, DDJ 2005.

Data: Intel, Wikipedia, K. Olukotun


The Cray-2: Immersed in Fluorinert



1985 ECL 150 kW

Where's all that power going?
What can we do about it?

Dally: Calculation Cheap; Communication Costly



“Chips are power limited and most power is spent moving data

Performance =
Parallelism

Efficiency =
Locality

Bill Dally's 2009 DAC Keynote, *The End of Denial Architecture*



Parallelism for Performance; Locality for Efficiency



Dally: "Single-thread processors are in denial about these two facts"

We need
different programming paradigms
and
different architectures
on which to run them.

The Future is Wires and Memory




How best to use all those
transistors?

Dark Silicon



Taylor and Swanson's Conservation Cores



Custom datapaths, controllers for loop kernels; uses existing memory hierarchy

Swanson, Taylor, et al. *Conservation Cores*. ASPLOS 2010.

Bacon et al.'s Liquid Metal

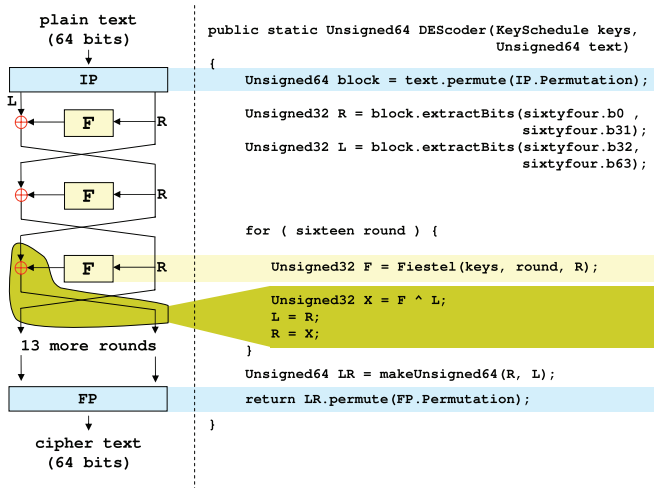



Fig. 2. Block level diagram of DES and Lime code snippet

JITting Lime (Java-like, side-effect-free, streaming) to FPGAs

Huang, Hormati, Bacon, and Rabbah, *Liquid Metal*, ECOOP 2008.

What are we doing about it?

Functional Programs to FPGAs



Functional Programs to FPGAs




Functional Programs to FPGAs




Functional Programs to FPGAs



Functional Programs to FPGAs



Functional Programs to FPGAs




Functional Programs to FPGAs



Why Functional Specifications?

- ▶ Referential transparency/side-effect freedom make formal reasoning about programs vastly easier
- ▶ Inherently concurrent and race-free (Thank Church and Rosser). If you want races and deadlocks, you need to add constructs.
- ▶ Immutable data structures makes it vastly easier to reason about memory in the presence of concurrency



Why FPGAs?

- ▶ We do not know the structure of future memory systems
Homogeneous/Heterogeneous?
Levels of Hierarchy?
Communication Mechanisms?
- ▶ We do not know the architecture of future multi-cores
Programmable in Assembly/C?
Single- or multi-threaded?



Use FPGAs as a surrogate. Ultimately too flexible, but representative of the long-term solution.

The Practical Question

*How do we synthesize hardware
from pure functional languages
for FPGAs?*

Control and datapath are easy; the memory system is interesting.

The Type System: Algebraic Data Types

Types are primitive (Boolean, Integer, etc.) or other ADTs:

```
type ::= Type           Primitive type  
      | Constr Type* | ... | Constr Type* Tagged union
```

Examples:

```
data Intlist = Nil           -- Linked list of integers  
            | Cons Int Intlist
```

```
data Bintree = Leaf Int     -- Binary tree of integers  
            | Branch Bintree Bintree
```

```
data Expr = Literal Int    -- Arithmetic expression  
          | Var String  
          | Binop Expr Op Expr
```

```
data Op = Add | Sub | Mult | Div
```


Example: Huffman Decoder in Haskell

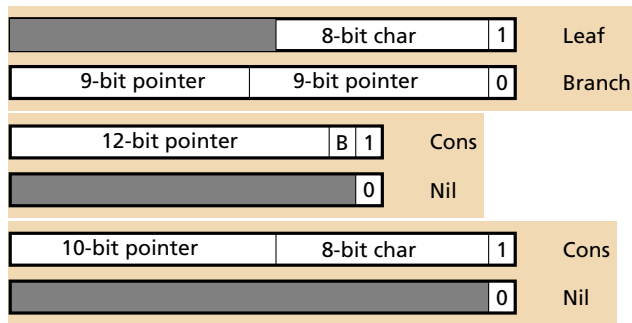
```
data HTree = Branch HTree HTree | Leaf Char
```

```
decode :: HTree → [Bool] → [Char]
```

```
decode table str = bit str table
```

```
where
```


```
bit (False:xs) (Branch l _) = bit xs l           -- 0: left  
bit (True:xs)  (Branch _ r) = bit xs r          -- 1: right  
bit x          (Leaf c)     = c : bit x table   -- leaf  
bit []         _            = []                -- done
```






Memory



Optimizations



Optimizations




Split
Memories
➔




Use Streams


Optimizations




Split
Memories
➔



Use Streams




Unroll for locality




Optimizations


Optimizations




Split
Memories
➔




Use Streams



Unroll for locality



Speculate



Hardware Synthesis: Semantics-preserving steps to a low-level dialect

Removing Recursion: The Fib Example

```
fib n      = case n of
    1      → 1
    2      → 1
    n      → fib (n-1) + fib (n-2)
```


Transform to Continuation-Passing Style

```
fibk n k      = case n of
    1      → k 1
    2      → k 1
    n      → fibk (n-1) (λn1 →
                                fibk (n-2) (λn2 →
                                                k (n1 + n2)))

fib  n      = fibk n (λx → x)
```

Name Lambda Expressions (Lambda Lifting)

```
fibk n k = case n of
  1     → k 1
  2     → k 1
  n     → fibk (n-1) (k1 n k)
```

```
k1 n k n1 = fibk (n-2) (k2 n1 k)
```

```
k2 n1 k n2 = k (n1 + n2)
```

```
k0 x = x
```

```
fib n = fibk n k0
```

Represent Continuations with a Type

```
data Cont = K0 | K1 Int Cont | K2 Int Cont
```

```
fibk n k      = case (n,k) of  
    (1, k) → kk k 1  
    (2, k) → kk k 1  
    (n, k) → fibk (n-1) (K1 n k)
```

```
kk k a      = case (k, a) of  
    ((K1 n k), n1) → fibk (n-2) (K2 n1 k)  
    ((K2 n1 k), n2) → kk k (n1 + n2)  
    (K0,          x ) → x
```

```
fib n      =          fibk n K0
```

Merge Functions

```
data Cont = K0 | K1 Int Cont | K2 Int Cont
```

```
data Call = Fibk Int Cont | KK Cont Int
```

```
fibk z      = case z of
```

```
  (Fibk      1 k) → fibk (KK k 1)
```

```
  (Fibk      2 k) → fibk (KK k 1)
```

```
  (Fibk      n k) → fibk (Fibk (n-1) (K1 n k))
```

```
  (KK (K1 n k) n1) → fibk (Fibk (n-2) (K2 n1 k))
```

```
  (KK (K2 n1 k) n2) → fibk (KK k (n1 + n2))
```

```
  (KK K0      x ) → x
```

```
fib n      =      fibk (Fibk n K0)
```

Add Explicit Memory Operations

`load` :: `CRef` \rightarrow `Cont`

`store` :: `Cont` \rightarrow `CRef`

data `Cont` = `K0` | `K1 Int CRef` | `K2 Int CRef`

data `Call` = `Fibk Int CRef` | `KK Cont Int`

`fibk z` = **case z of**

`(Fibk 1 k)` \rightarrow `fibk (KK (load k) 1)`

`(Fibk 2 k)` \rightarrow `fibk (KK (load k) 1)`

`(Fibk n k)` \rightarrow `fibk (Fibk (n-1) (store (K1 n k)))`

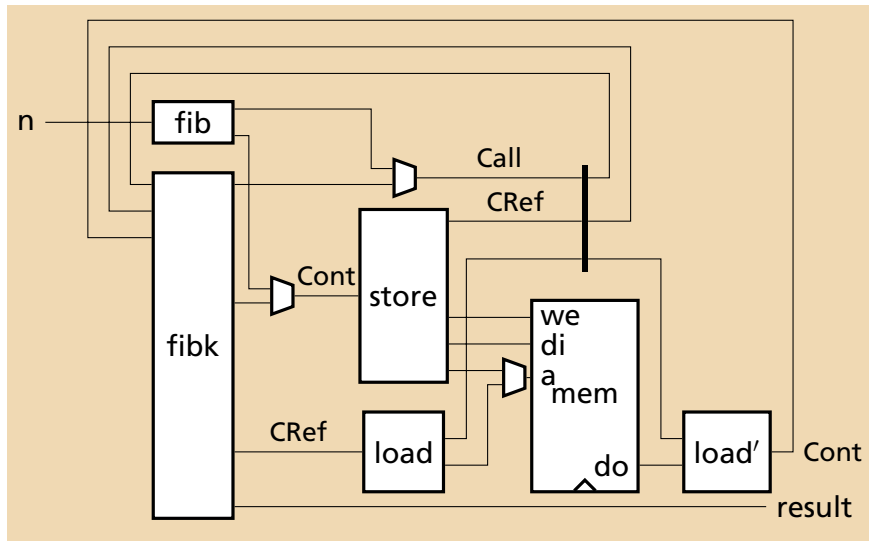
`(KK (K1 n k) n1)` \rightarrow `fibk (Fibk (n-2) (store (K2 n1 k)))`

`(KK (K2 n1 k) n2)` \rightarrow `fibk (KK (load k) (n1 + n2))`

`(KK K0 x)` \rightarrow `x`

`fib n` = `fibk (Fibk n (store K0))`

Syntax-Directed Translation to Hardware



Duplication Can Increase Parallelism

fib 0 = 0

fib 1 = 1

fib n = fib (n-1) + fib (n-2)

Duplication Can Increase Parallelism

```
fib 0 = 0
fib 1 = 1
fib n = fib (n-1) + fib (n-2)
```

After duplicating functions:

```
fib 0 = 0
fib 1 = 1
fib n = fib' (n-1) + fib'' (n-2)
```

```
fib' 0 = 0
fib' 1 = 1
fib' n = fib' (n-1) + fib' (n-2)
```

```
fib'' 0 = 0
fib'' 1 = 1
fib'' n = fib'' (n-1) + fib'' (n-2)
```

Here, *fib'* and *fib''* may run in parallel.

Unrolling Recursive Data Structures

Original Huffman tree type:

```
data Htree = Branch Htree HTree | Leaf Char
```

Unrolled Huffman tree type:

```
data Htree = Branch Htree' HTree' | Leaf Char
```

```
data Htree' = Branch' Htree'' HTree'' | Leaf' Char
```

```
data Htree'' = Branch'' Htree HTree | Leaf'' Char
```




Increases locality: larger data blocks.

A type-aware cache line

- ▶ Dark Silicon is the future: faster transistors; most must remain off

- ▶ Custom accelerators are the future; many approaches

- ▶ Our project: A Pure Functional Language to FPGAs




▶ Algebraic Data Types in Hardware

▶ Optimizations

▶ Removing recursion

Encoding the Types



Huffman tree nodes: (19 bits)



Boolean input stream: (14 bits)



Character output stream: (19 bits)



Syntax-Directed Translation to Hardware

