# Design Space Exploration for Programmable Fabrics

Daniel Bates, Alex Bradbury and Robert Mullins

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

## Motivation

Challenges:
- Power limits require a smaller percentage of transistors to be active
- Need to accept device failures
- Poor interconnect scaling
- Amdahl's law
- Design and verification cost
- Complexity
- Greater levels of control flow in new embedded applications

Aims:
- Explore the design space between FPGAs and many-core processors
- Design a simple, robust, any-purpose fabric
- Achieve power-efficiency by specialising:
  - execution datapaths
  - communication
  - memory subsystem
- Purely software-programmable

## Overview

Our architecture, Loki, is an all-purpose, robust, homogeneous computing fabric. It consists of hundreds or thousands of individual cores and memories interconnected by a hierarchical, chip-wide network. It is targeted at embedded computing applications.

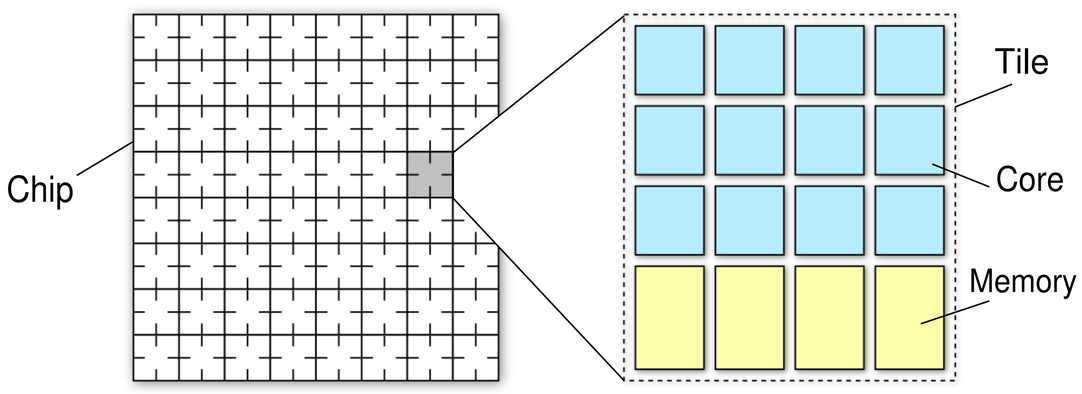


Characteristics:

Loki aims to combine the flexibility of an FPGA approach with the high performance and low power of an ASIC, while remaining a good target for high-level languages. The network-centric design blurs the boundaries between cores, allowing them to share resources.

## Architecture

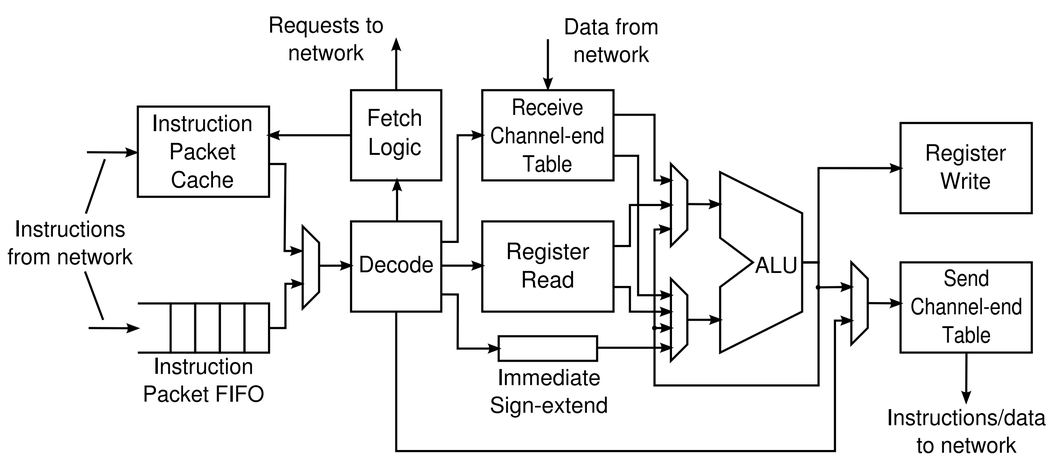


Each core features:
- Simple, 4-stage pipeline
- Register-mapped channel-ends
- Instructions grouped into packets
  - No program counter
- Indirect register access support

Other details:
- Tiled, homogenous architecture
- Composable memory blocks

Homogeneity simplifies:
- Design validation
- Fault-tolerance
- Place and route
- Application mapping
- Scaling

Current status:
- SystemC simulator implemented
- Developing benchmarks

## Target Application Areas

A range of embedded applications will be targeted, with a particular emphasis on those which may become common in the future, or which are traditionally difficult to parallelise.

Traditional:
- Video/image/audio processing
- Wireless network protocols
- Encryption/decryption
- Common operations
  - Searching, sorting, FFT

Emerging:
- Face detection and recognition
  - Emotion recognition
- Voice recognition
- Data mining
- Neural networks

## Execution Patterns

### Local

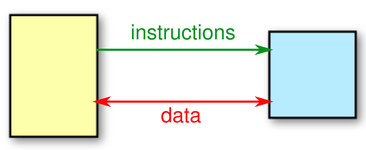


Behaviour is like that of an ordinary RISC core.

### Data-driven

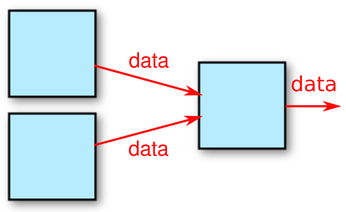


e.g. Streaming datapaths, or program a core to act as a router or memory arbiter.

### Remote

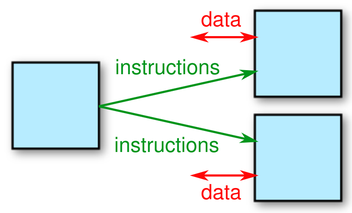


e.g. One core takes control of another to set up SIMD or VLIW virtual processors.

## Software Specialisation

Benefits of heterogeneity can be achieved by specialising the fabric in software:
- Give each core a very small section of code
- Identify unneeded functionality in each core
- Tailor the memory system to the application
- Group many specialised cores and memories together to form a virtual processor

The fabric can be reconfigured at runtime by simply executing different code.

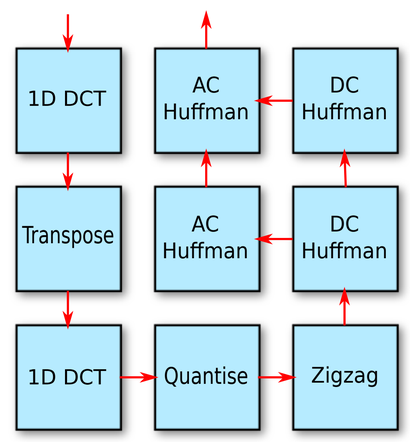


## Design Space Exploration

Determine the best possible design for any combination of constraints:
- Parametrise the design as much as possible
  - Eventually, explore using architectural transformations
- Perform a wide sweep over the design space
  Using the results of exploration, determine the best targets for optimisation and/or innovation

The design space is too large to explore thoroughly:
- It must be carefully constrained, sampled and pruned
- Interesting approaches to exploration include machine learning, genetic algorithms, tabu search, and simulated annealing

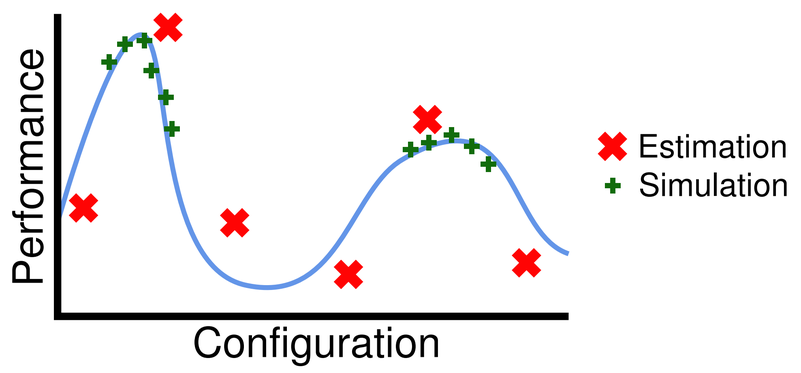


Daniel.Bates@cl.cam.ac.uk
Alexander.Bradbury@cl.cam.ac.uk
Robert.Mullins@cl.cam.ac.uk

Computer Architecture Group
http://www.cl.cam.ac.uk/research/comparch

Supported by EPSRC grant EP/G033110/1

2010-07-14