

Superscalar Performance without Superscalar Overheads



UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

Ali Mustafa Zaidi & David Greaves

The 'Dark Silicon' Challenge

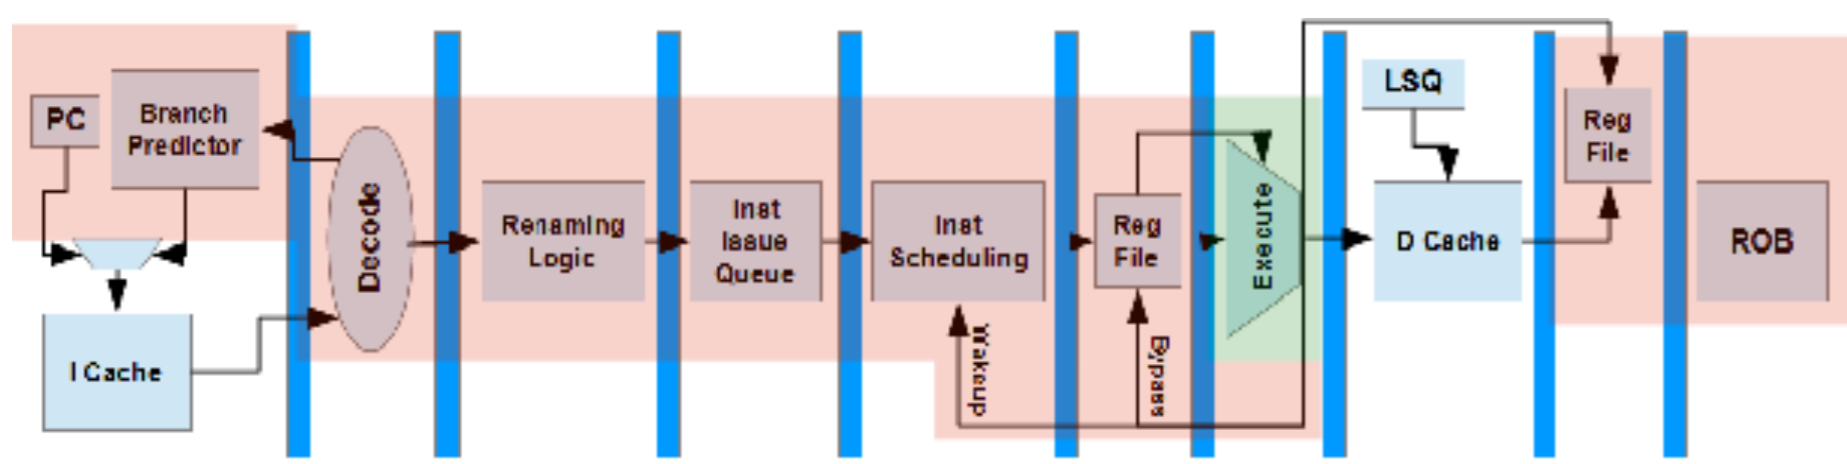
In the 'Dark Silicon Era', achievable speedups will be far below the exponential growth suggested by Moore's Law, primarily due to constrained power budgets. To mitigate the effects of Dark Silicon, we must dramatically improve the energy efficiency of computation.

Computer architects are increasingly relying on custom hardware to achieve high performance with very high efficiency. Unfortunately, there are several problems:

(1) **Amenability**: custom hardware currently exhibits poor performance on irregular, control-flow intensive applications, or those with low data parallelism.

(2) **Programmability**: developing custom hardware is far more difficult & time consuming than writing code for a sequential processor.

Thus to provide a balance between energy, performance and programmability, we must build heterogeneous designs: combining complex, power hungry superscalar processors with energy efficient custom hardware.



$$EPI_{CPU} = \frac{E_{I-CacheWarmup}}{InstructionCount_{Static}} + \frac{(E_{Execute} + E_{CPUOverheads})}{InstructionCount_{Dynamic}}$$

Dataflow Computing for Efficiency and Performance

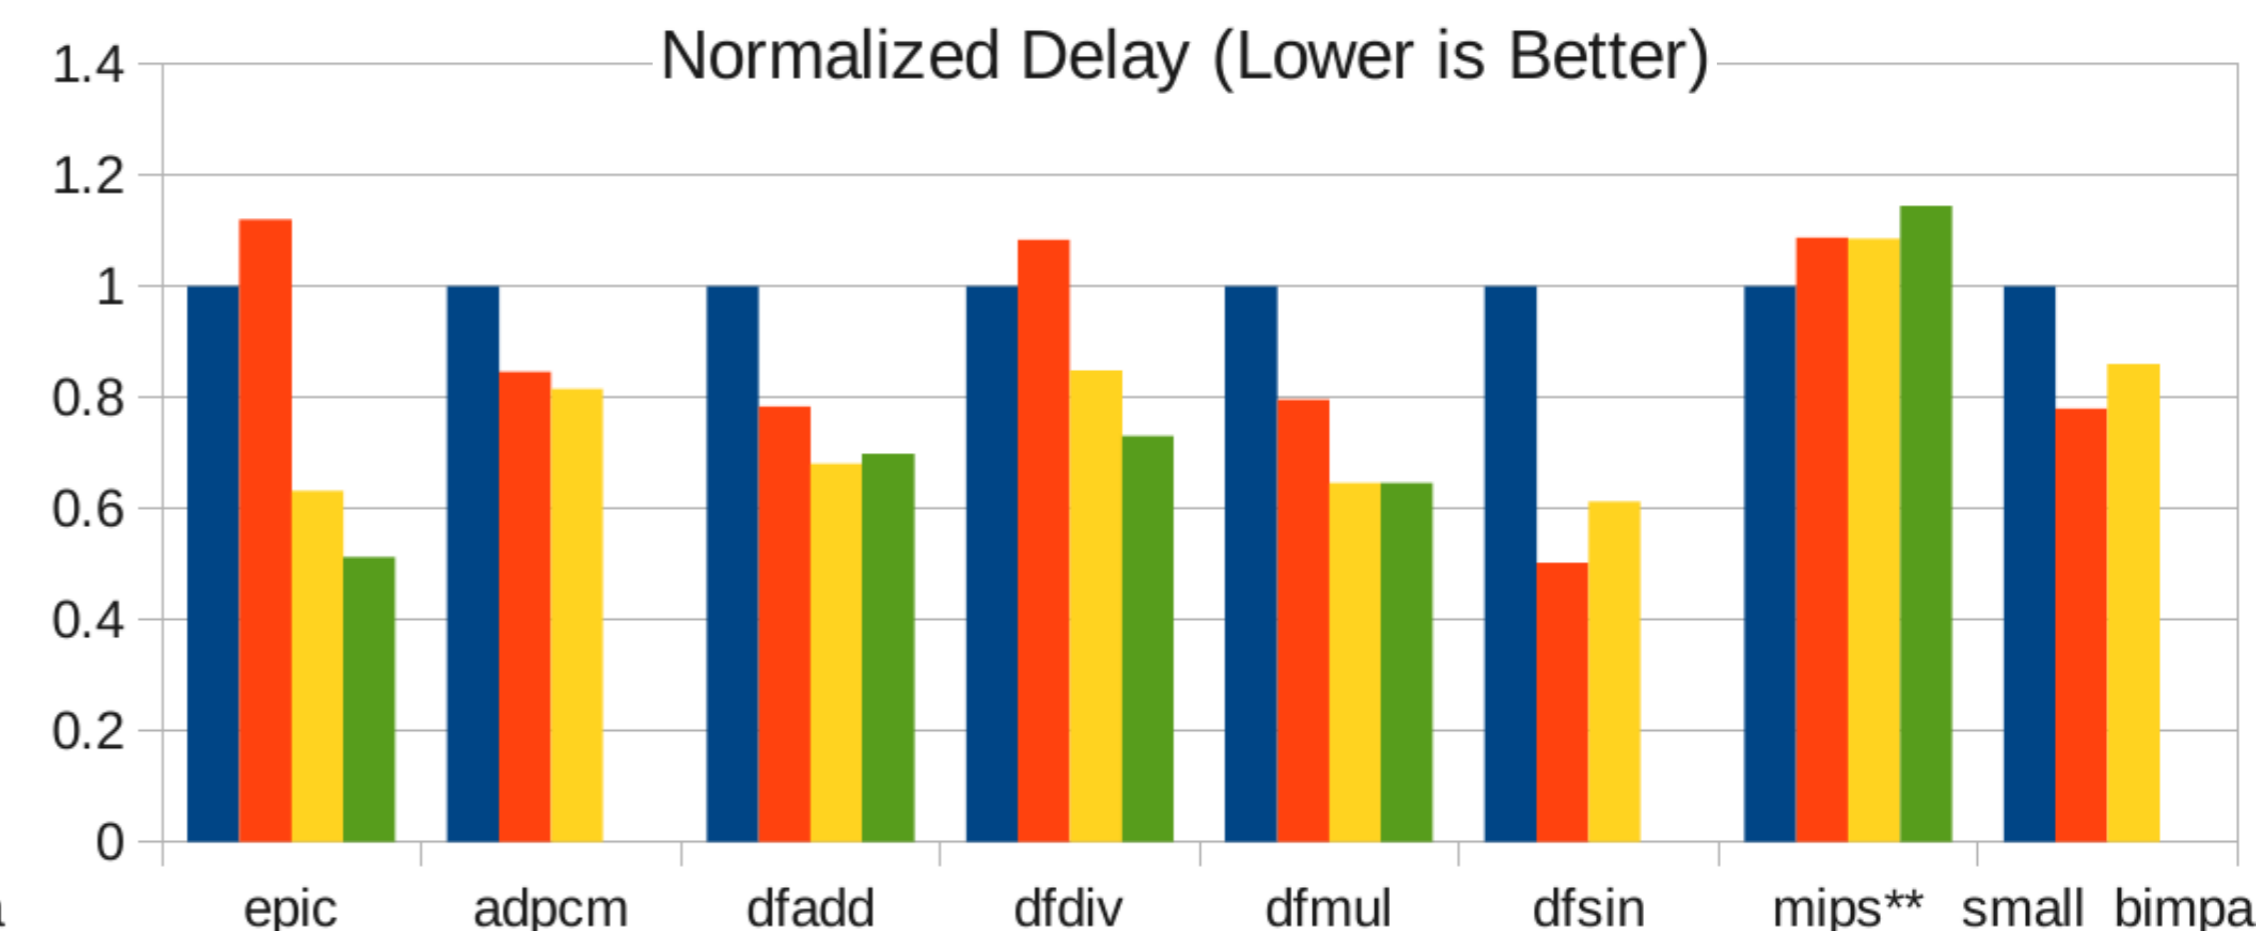
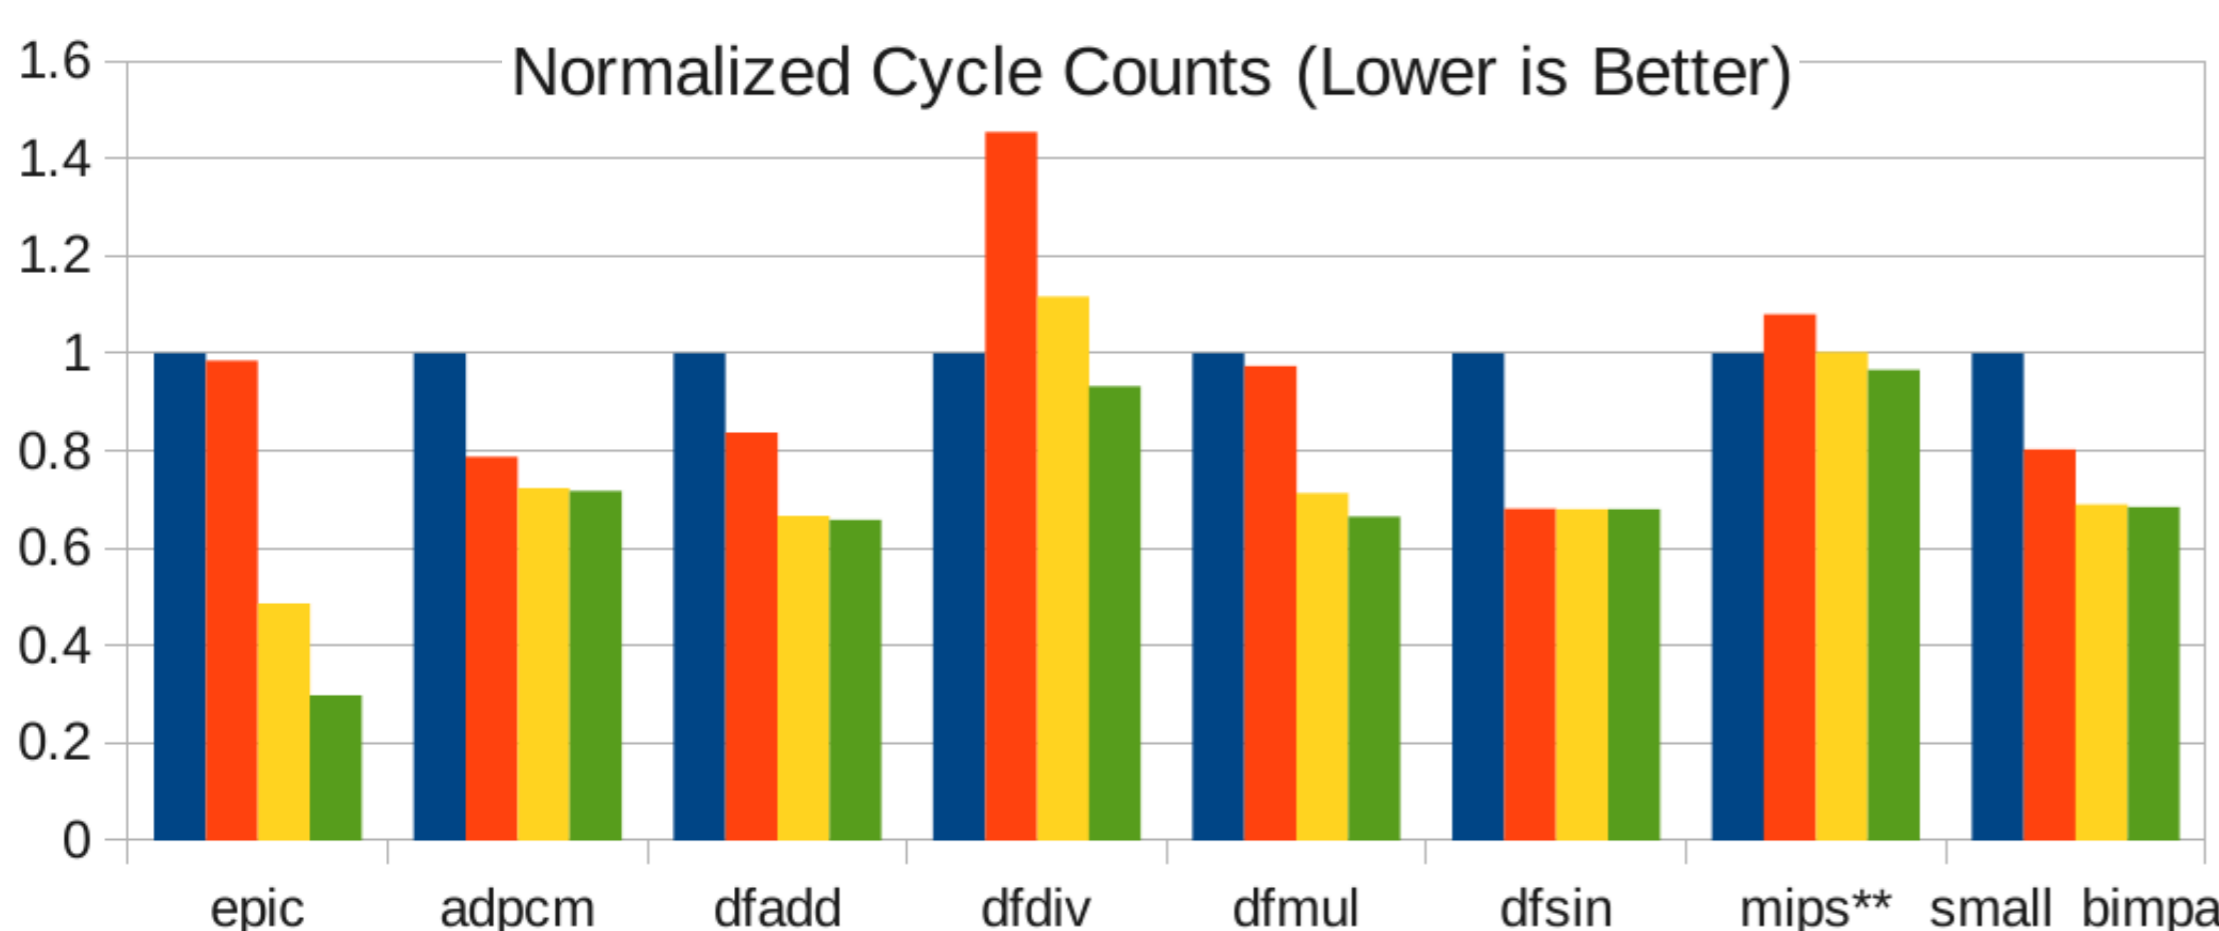
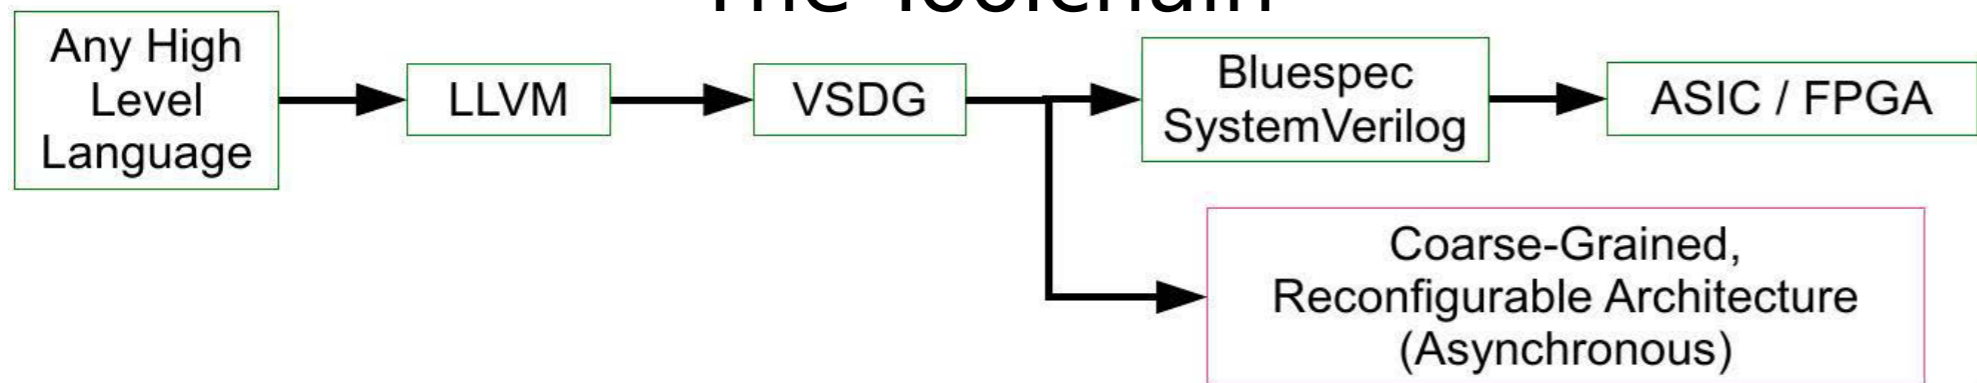
Our research goal is to address the Programmability and Amenability issues of custom hardware, enabling a broader, more energy-efficient design space for heterogeneous systems. We do this by:

(1) Developing a new program representation (IR) - the 'Value-State Flow Graph (VSFG)' that is inherently dataflow oriented, exposing much greater concurrency than existing IRs like the conventional Control Flow Graph (CFG).

(2) Providing a new High-level synthesis (HLS) toolchain that converts input high-level language code into our dataflow IR and then implements it as custom dataflow hardware.


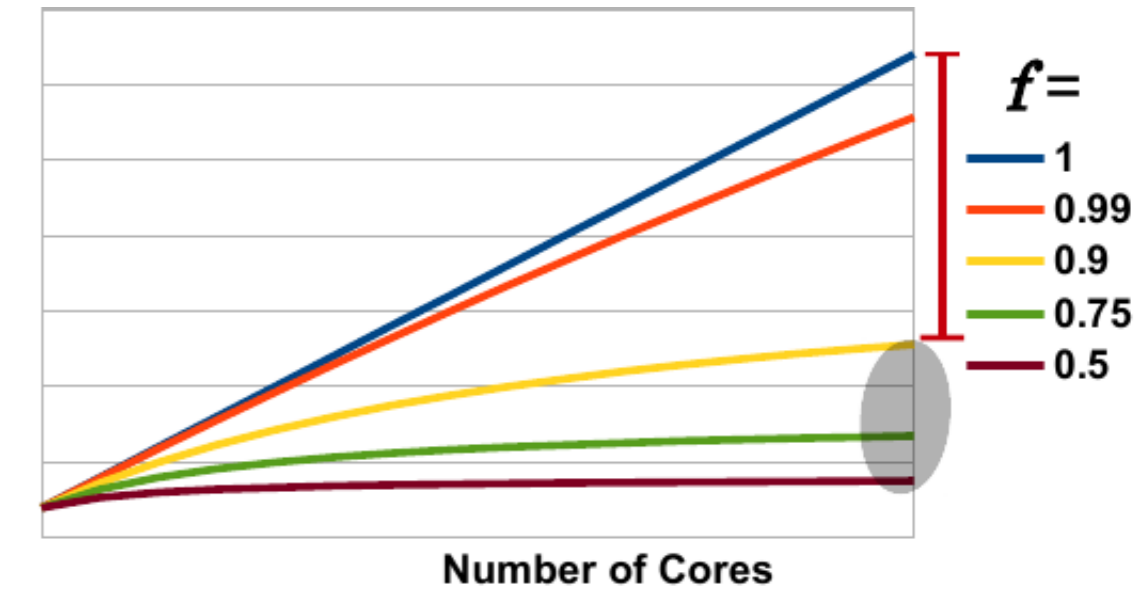
(3) Unlike conventional HLS tools, but like superscalar processors, our generated custom hardware uses *dynamic execution scheduling* to provide further performance improvements on control-flow intensive code.

The Toolchain

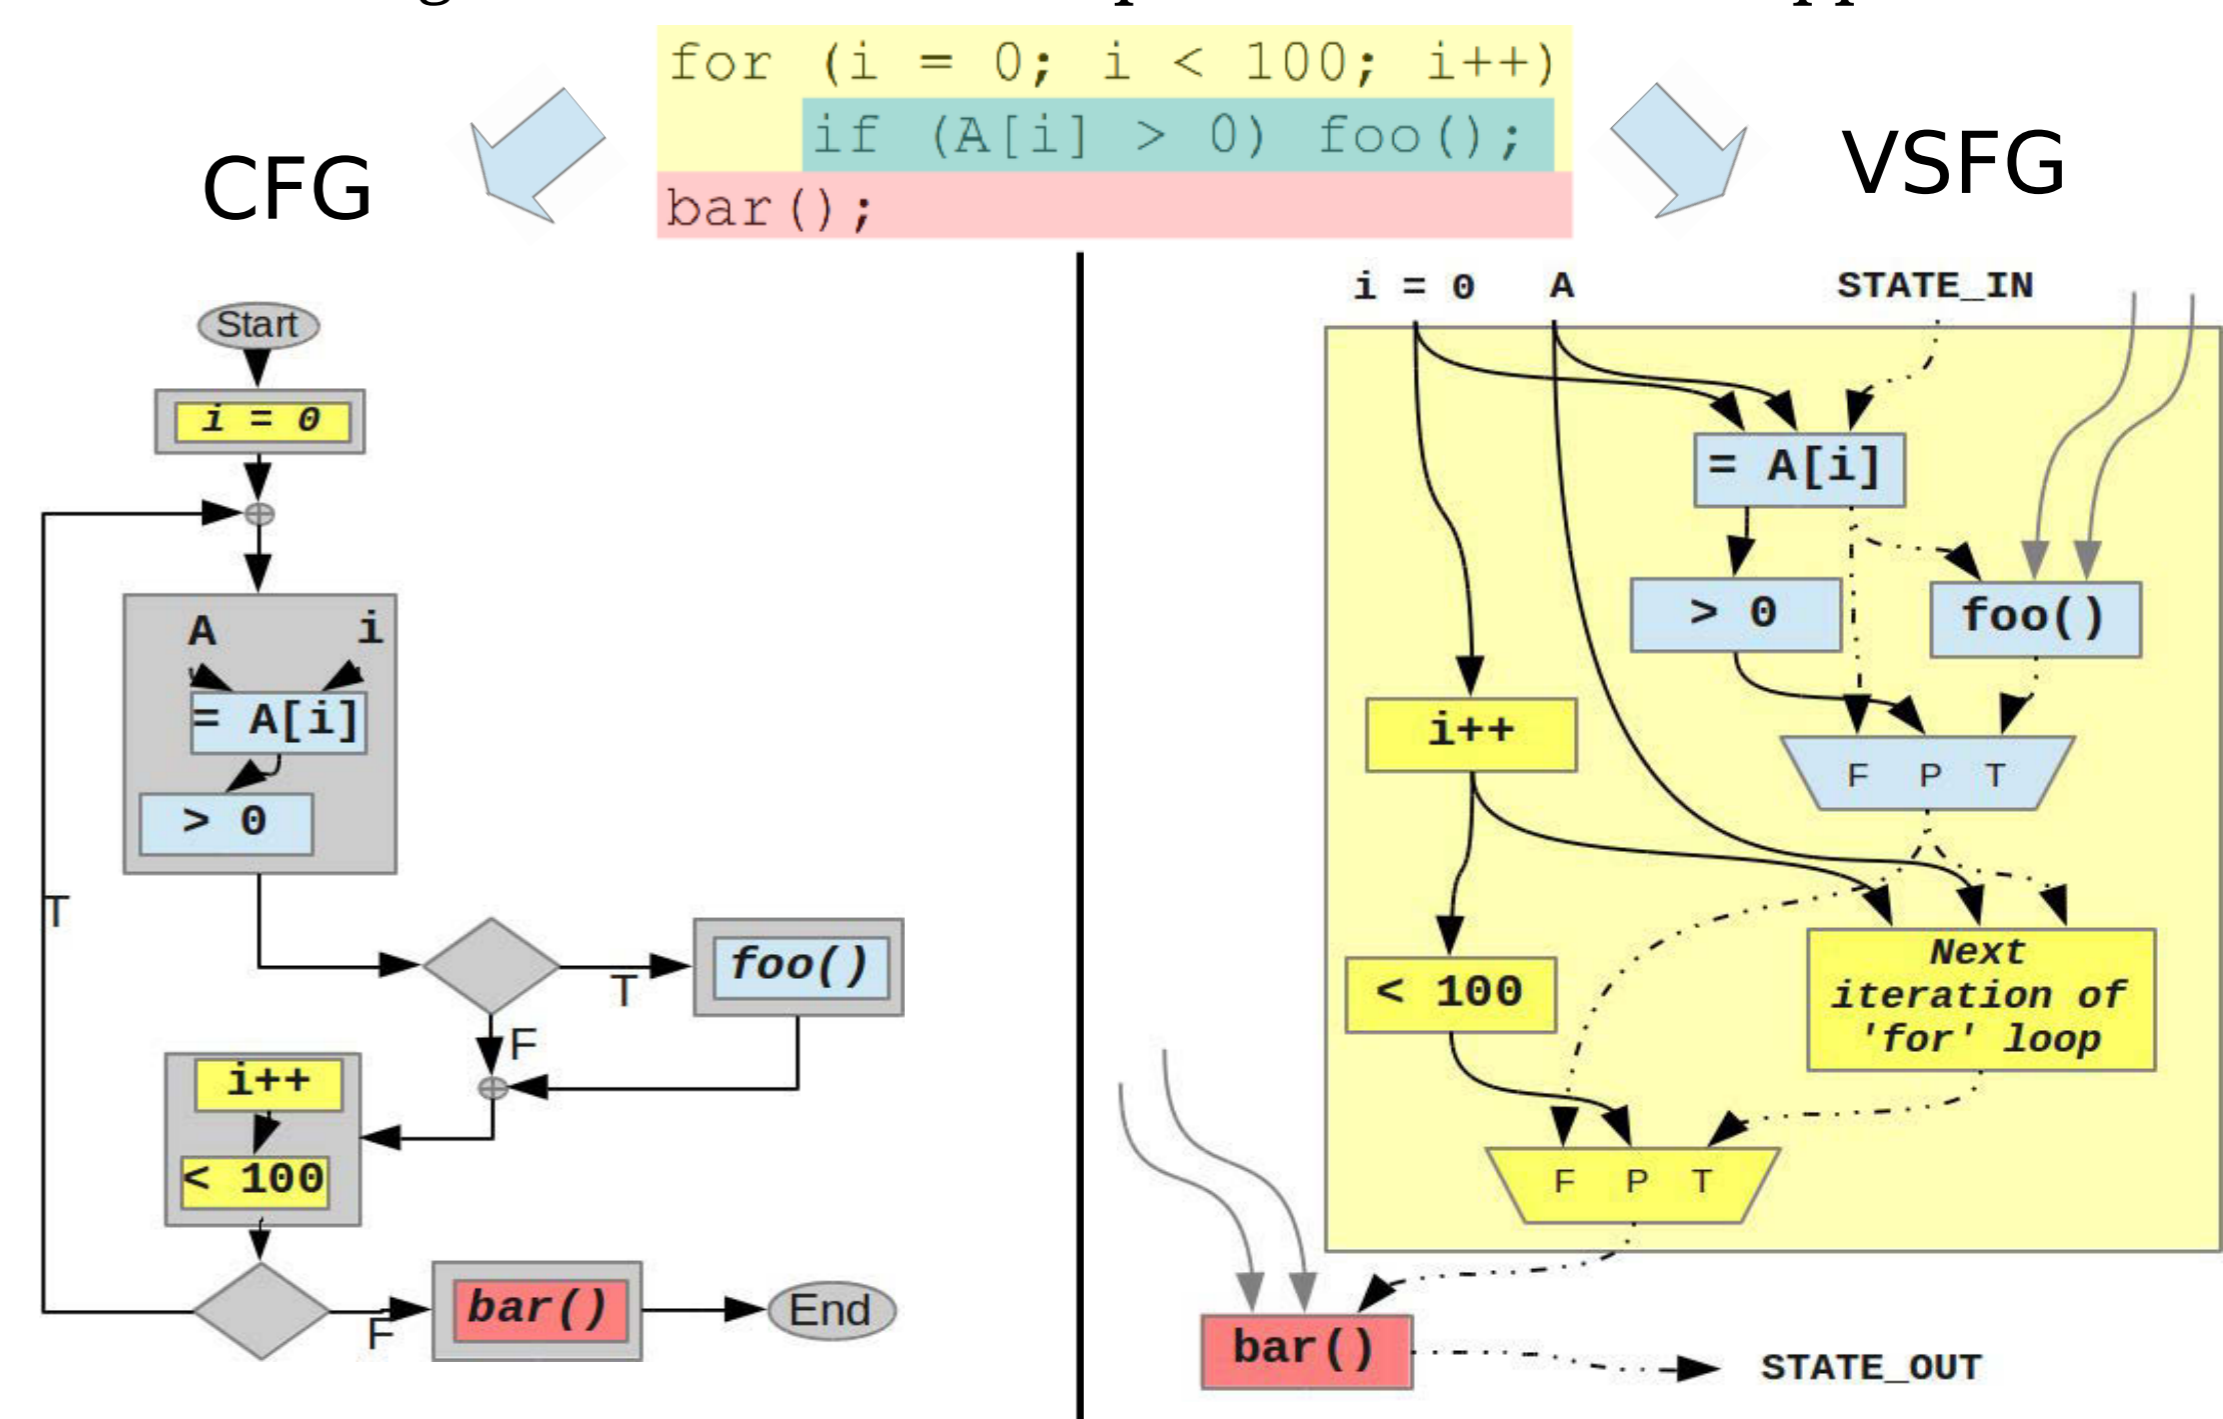


The Dark Silicon Problem

$$Speedup = \frac{1}{\frac{(1-f)}{S} + \frac{f}{N*S}}$$



Superscalar processors achieve high performance by internally converting from the sequential Von Neumann to the highly concurrent Dataflow computational model: (1) Branch prediction is used to overcome control dependences, (2) Register Renaming is used to overcome name dependences (3) Out-of-order Issue, Bypass & forwarding logic is used to accelerate true dependences. However, performing this conversion at runtime in a general-purpose processor incurs over 10-1000x greater energy overheads than the equivalent custom / reconfigurable hardware implementation of an application.



The VSFG exposes more concurrency by statically converting control flow to dataflow

The Results

We compare against LegUp, an existing HLS tool from Toronto. The VSFG allows us to perform aggressive loop unrolling. VSFG results shown for 0, 1, & 3 degrees of unrolling.