Cycle-Accurate, Distributed Chip Multiprocessor Simulation

Motivation

It is anticipated that by approximately 2010 the number of transistors available on a single chip will exceed one billion.

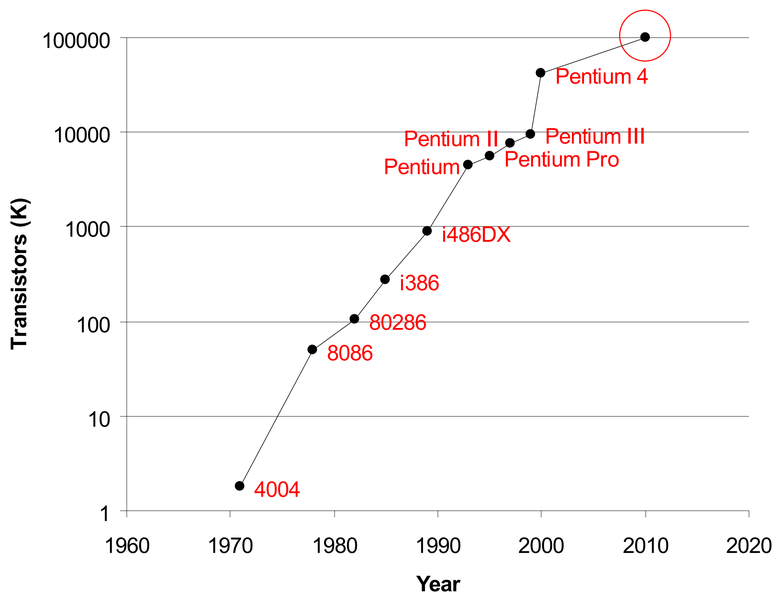

What is the optimal use of a billion transistor chip?

The decrease in feature size notably of wires and transistors will lead to proportionately large wire delays across the chip, which will hamper the progress of future deeply pipelined and more superscalar processor designs, are CMPs a viable alternative?

Chip Multiprocessors (CMP)

CMPs typically consist of a number of simple processor cores on a single chip. The repetition of components duplicated in each core can reduce the overall processor design time.

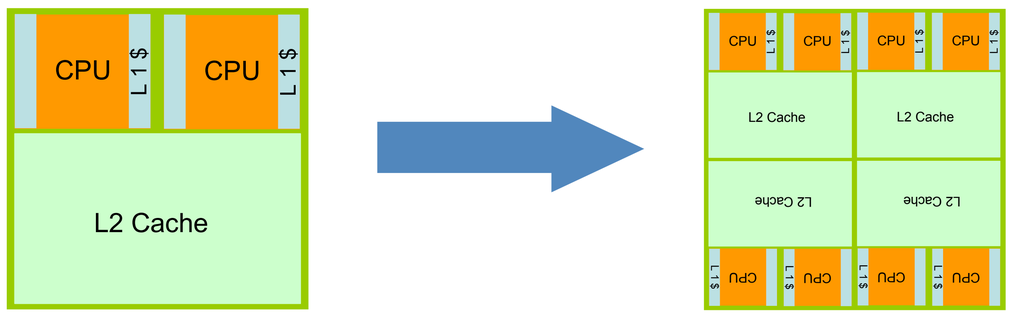

Locality of tightly integrated components within each processor core reduces wire delay effects, as the longest clock signal path is significantly reduced in CMP designs.

The design space for CMPs is vast, and simulation is the only viable method of evaluating it.

Simulation of CMP designs

Simulation is an essential tool for ensuring both the functionality and performance of a new processor design.

Trade-offs can be made between the speed and accuracy of the simulation. However in order to produce reliable results when evaluating CMP architectures, cycle-accurate simulations are required.

When simulating a CMP architecture the simulation time linearly increases with the number of processing cores on the chip.

SIMPA - simulation of multiprocessor architectures

In order to evaluate the design space of a CMP architecture aimed at a billion transistor chip, **SIMPA** a distributed and cycle accurate CMP simulator has been designed. By distributing the simulation we aim to reduce the simulation time.

Distribution in the SIMPA Simulation

The **SIMPA** simulator has been designed to be distributed both globally on separate resources and locally using threading.

Global consistency needs to be maintained in a distributed simulator, and three techniques were considered.

Look-ahead and Null Messages [Chandy & Misra 1978]

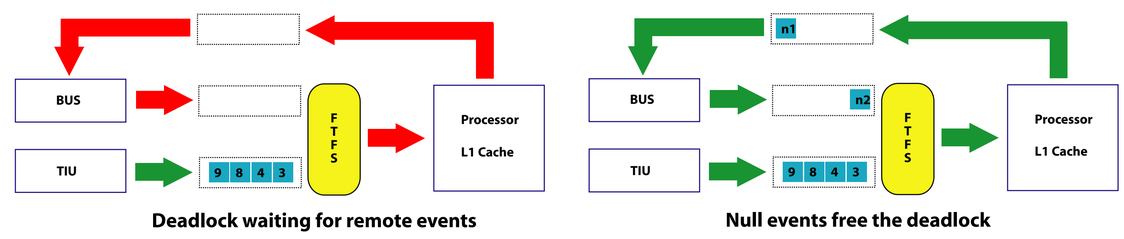

Barrier Synchronization [Fujimoto 2000]

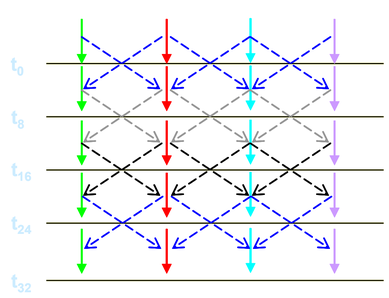

Virtual Time and Rollback [Jefferson 1984, 1990]

The communication latency of the first two conservative schemes is restrictive on the progress of the simulator. The optimistic rollback scheme appears to allow the greatest potential gains.

Extensible simulation platform

The **SIMPA** simulator can be easily reconfigured to test different architectures. An XML document describing the architecture to evaluate is loaded into the simulator at run-time.

```
<!DOCTYPE JamaicaArchitecture SYSTEM "./JamaicaArchitecture.dtd">
<JamaicaArchitecture>
        <Bus />
        <Memory Size="4096" Type="RDR"/>
      <MultiProcessor>
              <Bus />
              <Cache>
                  <TagsAndData Size="64" LineSize="8" Associativity="4"/>
              </Cache>
              <ProcessorCluster>
                  <Bus />
                  <Cache>
                      <TagsAndData />
                  </Cache>
                  <Processor>
                      <TIU />
                      <InstructionCache>
                          <TagsAndData Size="16" LineSize="4" Associativity="4"/>
                      </InstructionCache>
                      <DataCache>
                          <TagsAndData Size="16" LineSize="4" Associativity="4"/>
                      </DataCache>
                  </Processor>
        </ProcessorCluster>
      </MultiProcessor>
</JamaicaArchitecture>
```

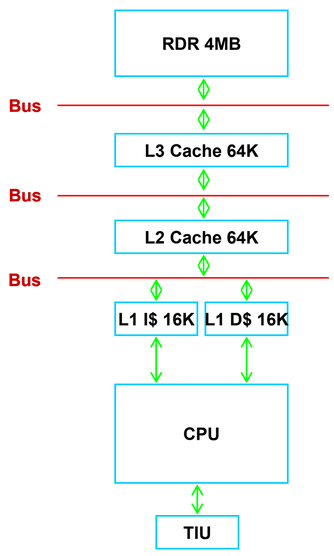

Using pre-built components the architecture can be reconfigured to evaluate scalability, topology or hierarchies for alternative CMP designs.

Current Status

The **SIMPA** simulator currently simulates the JAMAICA chip multiprocessor core, and additional components such as bus, cache and memory are nearing completion.

Future Work

Using the simulator to assess trade-offs between more caches and more cores, adding hardware directed thread level speculation and evaluating work distribution schemes.

The ultimate goal of this work is to find the optimal CMP configuration for a billion transistor budget.

Matthew Horsnell , Advanced Processor Technologies Group , University of Manchester