Silicon Graphics TFP Micro-Supercomputer Chipset

Peter Yan-Tek Hsu

Silicon Graphics Computer Systems
Mountain View, California

TFP is a collaborative development of MIPS Technologies Inc., Toshiba Corporation and Weitek Corporation.

What's New?

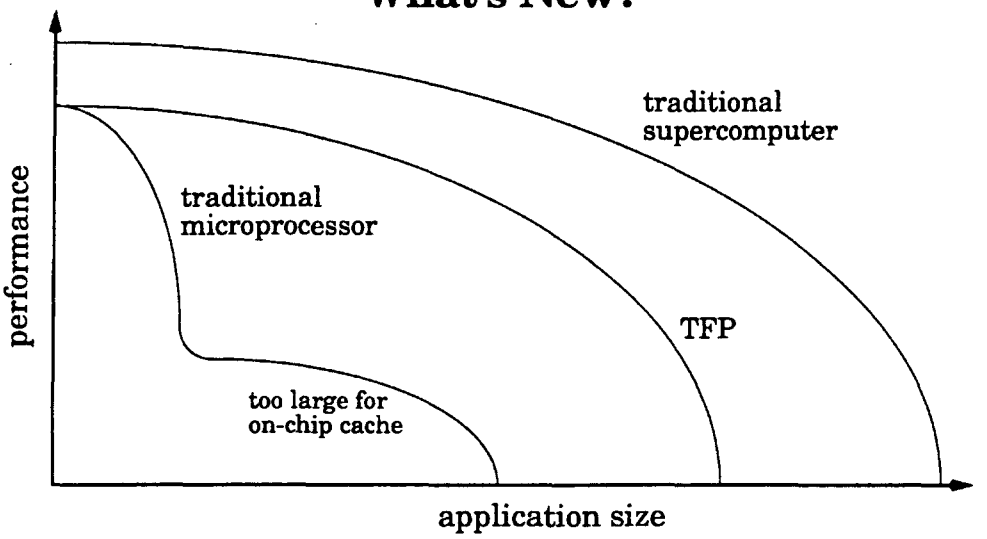

Supercomputer performance at microprocessor prices

- 300 MIPS, 300 MFLOPS peak
- 1.2 GB/s load/store bandwidth
- 8"×7" footprint, < 100 W per processor chipset

Outline

Overview

Integer Pipeline

Superscalar Instruction Dispatching

Branch Prediction

Floating-Point Unit

Split-Level Cache

Interleaved Access Patterns

Chip Facts

Summary

Product Goals

Performance of low-end vector supercomputers

- Two double-precision multiply-adds per cycle (4 FLOPS)
- Two 64-bit load/store operations per cycle
- Multiprocessor capable

Microprocessor-like cost structure

- Few high-integration CMOS chips
- Extensive caching using commodity SRAMs
- TTL signals, PGA packages

Compatibility with SGI product line

- Superscalar implementation of MIPS instruction set
- Scheduled code runs well on R4000 microprocessor

Innovation: Split Level Cache

Architect's dilemma:

- Integer/address computations need fast caches; small ok
- Floating-point needs large, high bandwidth caches; slower ok
- Both must perform well to satisfy real workloads

Solution:

- Single-cycle on-chip cache for integer/address data
- Multi-cycle pipelined off-chip cache for FP data
- Decoupled FPU pipeline hides off-chip cache latency
- Hardware maintains coherence between on/off chip cache

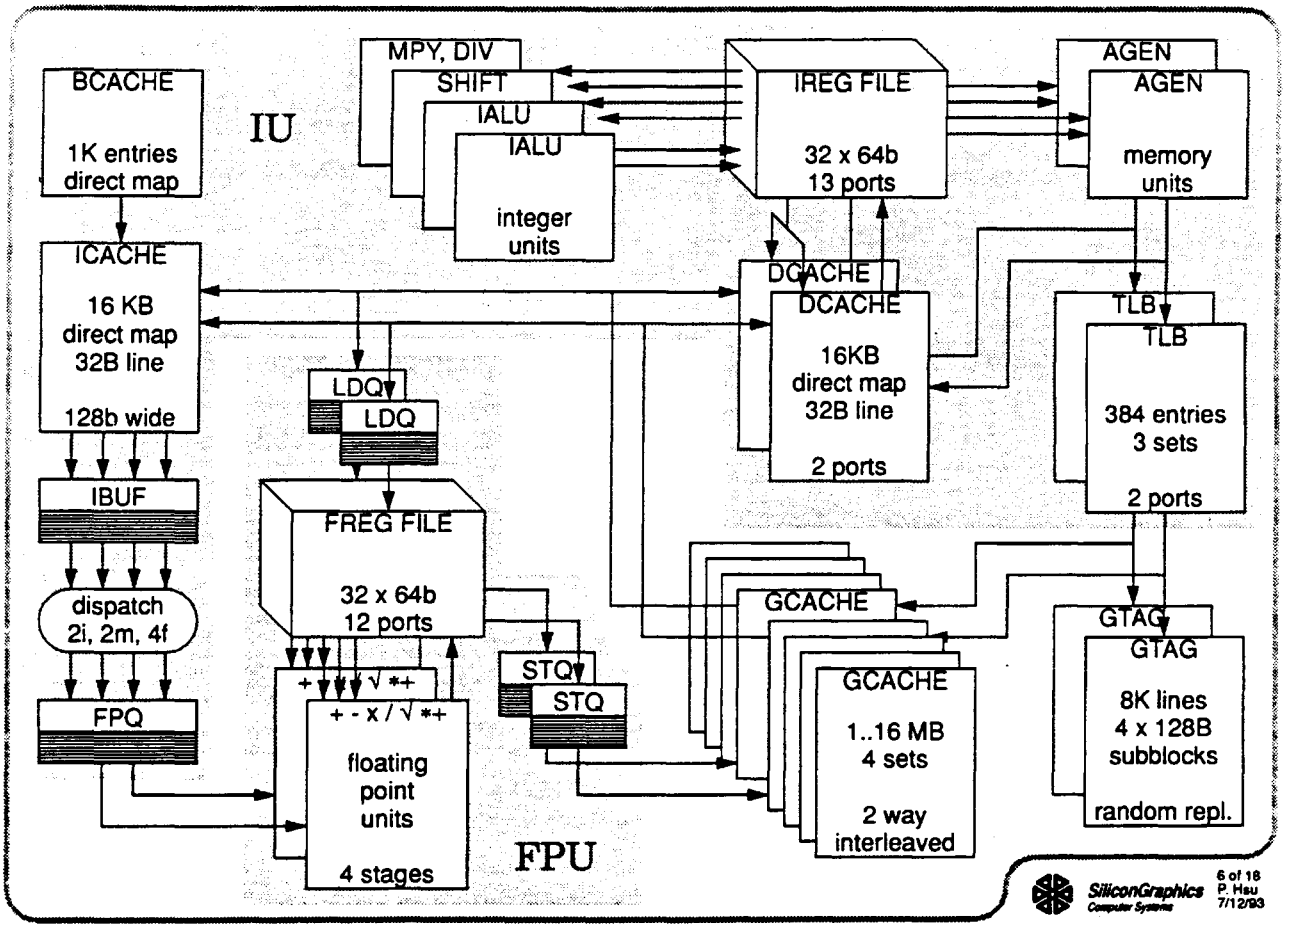

Integer Pipeline

Fetch	Decode	Address	Execute	Writeback
instruction fetch alignment predecode branch prediction	decode scoreboard register file read	generate load/store address	ALU operation data cache access branch resolution TLB lookup exception detection	register file write

- Low-latency RISC pipeline
- No load delay slot
- Data forwarding for dependent load after store in same cycle
- Branch and delay slot can execute together

Instruction Fetch

Insensitive to code alignment

- Executes 0 to 4 instructions each cycle
- Always presents next 4 instructions to decoder

Decoupled instruction prefetch pipeline

- Dispatch decision takes full cycle
- Critical path controls few muxes, not RAM address

Avoid draining pipeline by branch prediction

- "One-bit" prediction scheme
- Implemented as extension of instruction cache RAM
- Gain accuracy by large size

Instruction Alignment Buffer

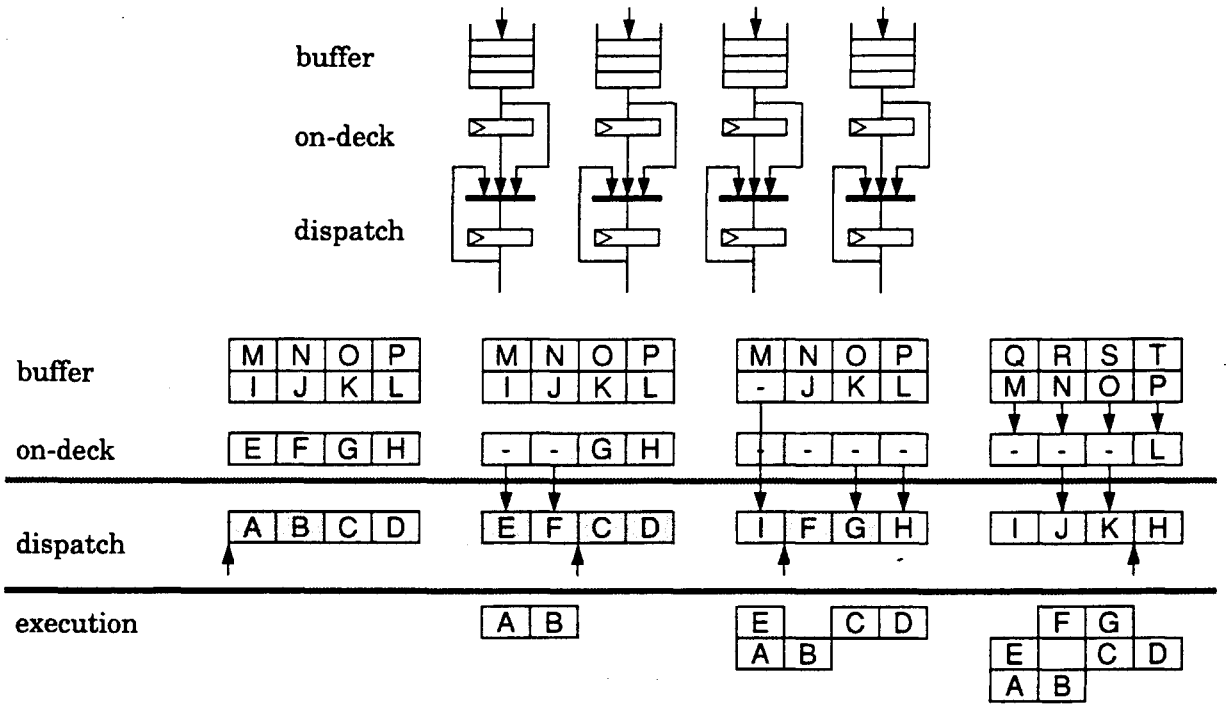

Branch Prediction

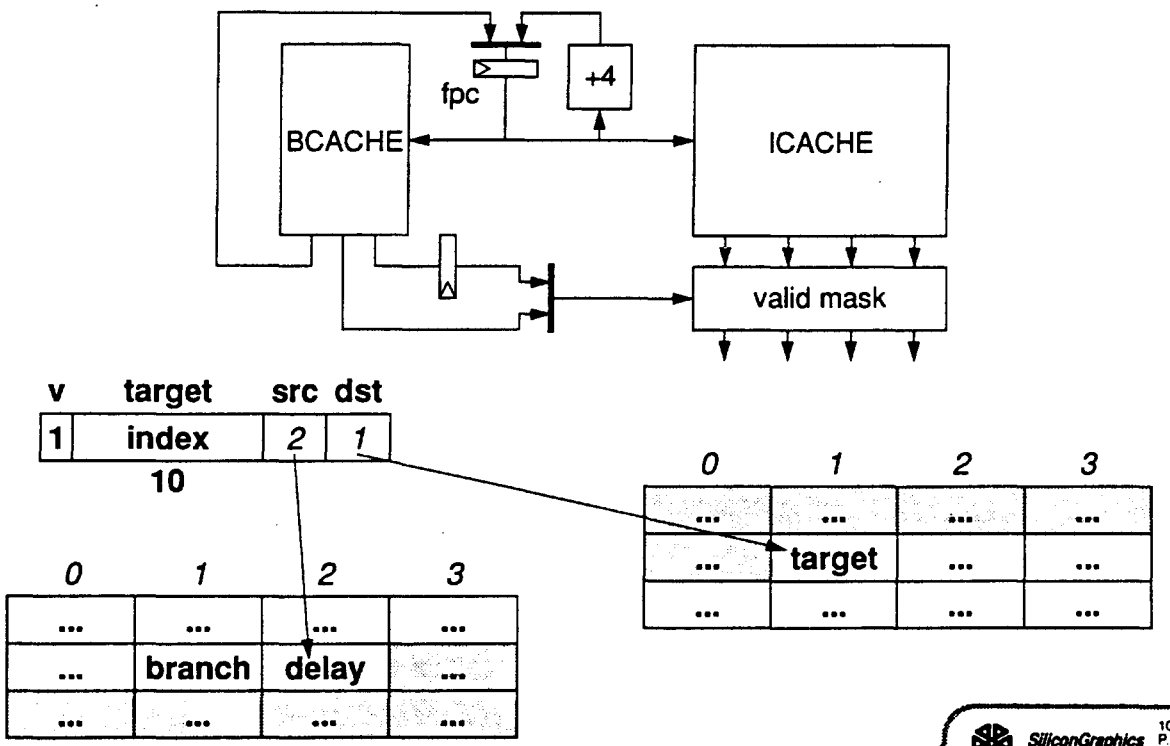

Misprediction Penalty

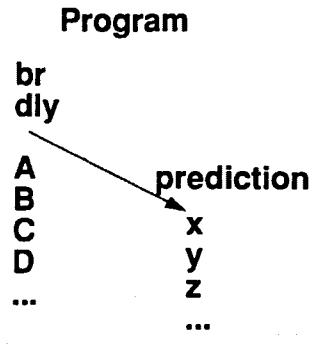

F	D	A	E	
br	...	...	...	
dly	br	...	...	
x	dly	br	...	begin fetching target
y	x	dly	br	
z	y	x	dly	misprediction detected
A	z	y	x	begin fetching good target
B	A	z	y	
C	B	A	z	3 cycle penalty
D	C	B	A	resume execution

- Branch prediction associated with delay slot
- Unified E-stage logic for misprediction, exception, cache miss

Decoupled Floating-Point Unit

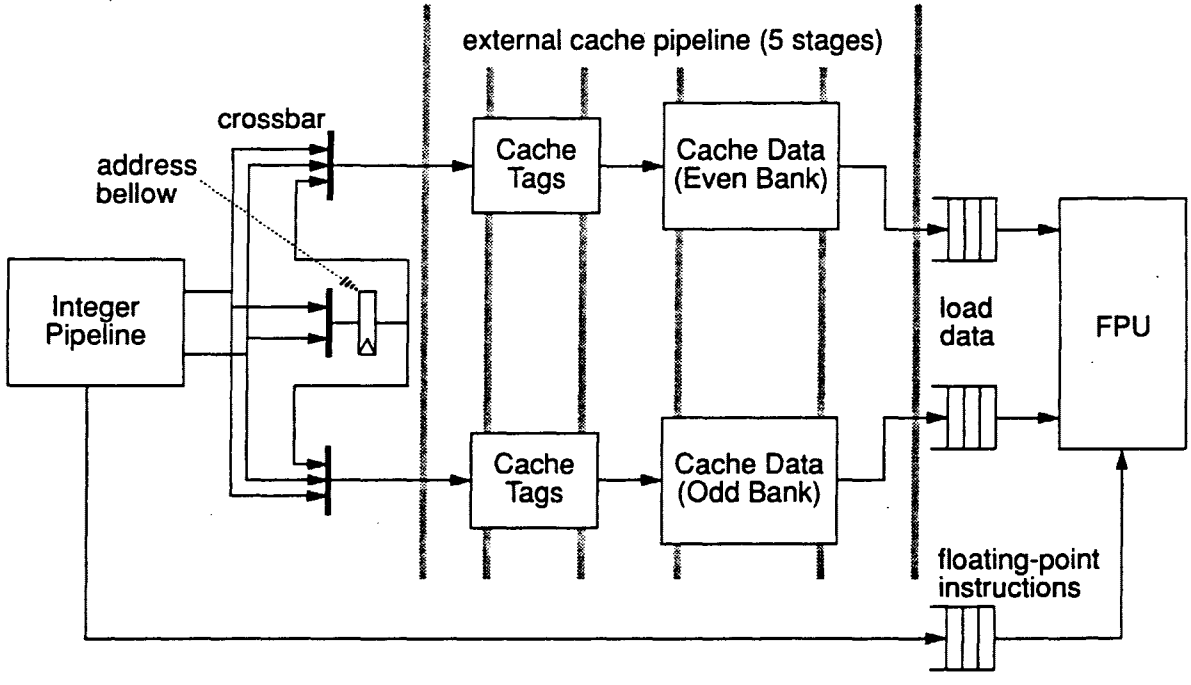

Unbuffered Interleaved Cache

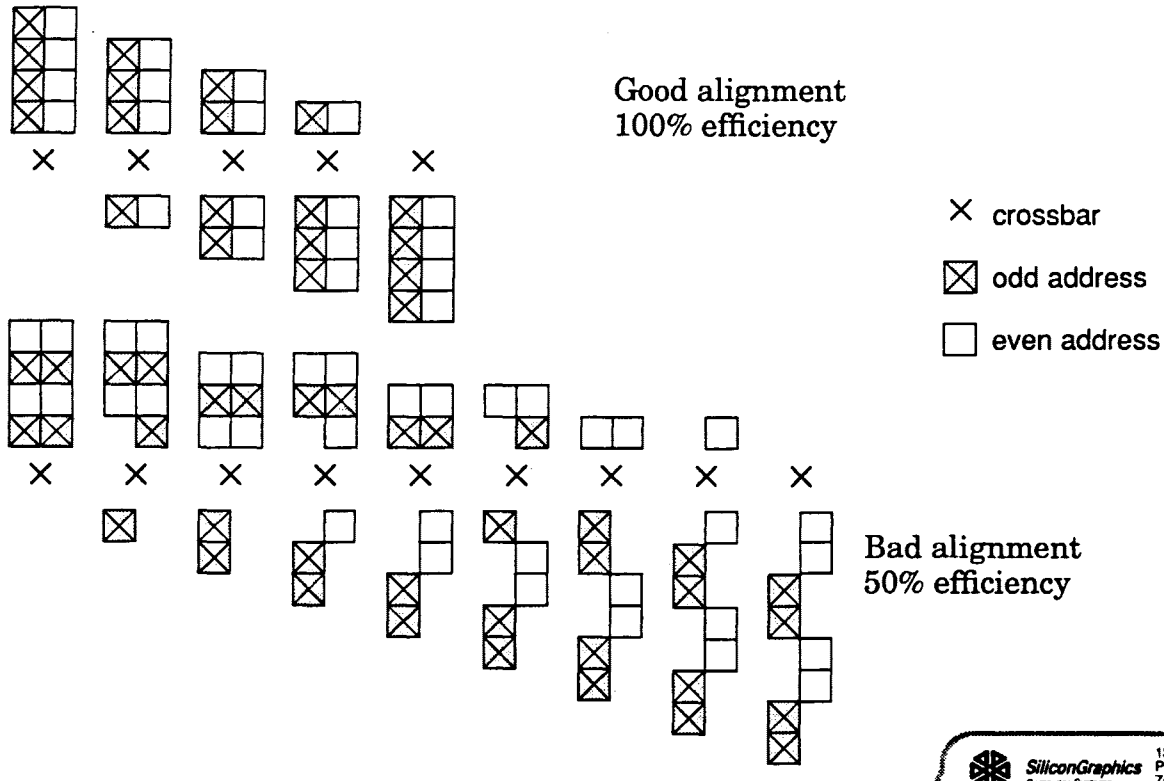

Effect of Address Bellow

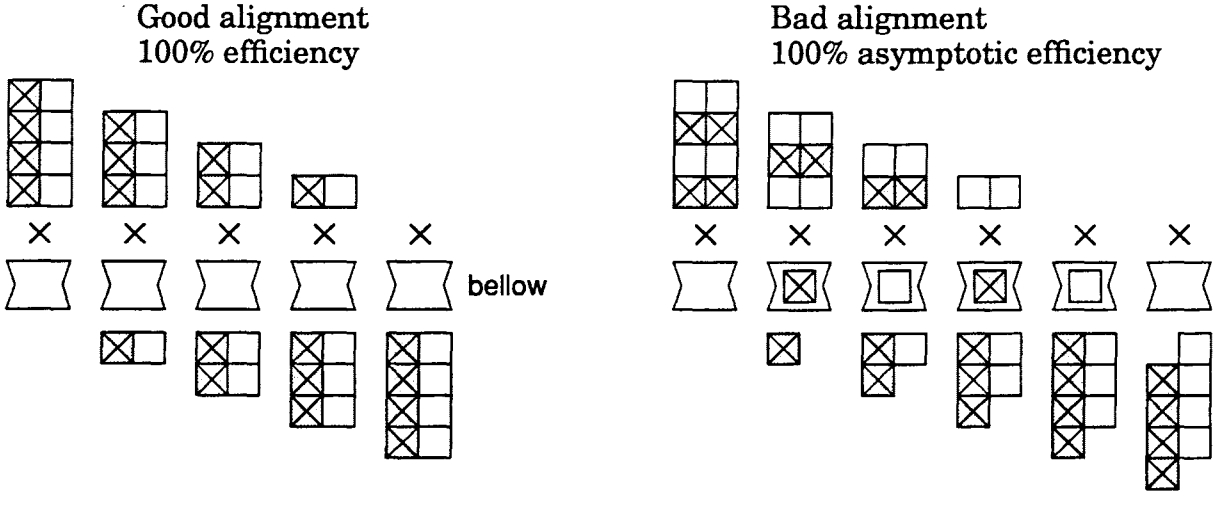

Address bellow resolves bank conflicts

- Compiler worry: Uniform distribution of even/odd references
- Hardware takes care of alignment problems
- Effective for subroutine array parameters

Split Level Cache Coherence

```
struct {
        char Ch[4]
        float Flt
        int Ival
} s
```

```
if ( s.Ch[2] == ... ) {
        s.Flt = ...
}
... = s.Ival
```

Dcache Action

	Hit	Miss
Int St	merge data set valid	fetch & retry (set invalid)
FP St	set invalid	nop

valid states	Cache Line
	...
1 1	... | ... | Ch | ...
1 0	Flt
1 1	Ival
	...

- Valid bit per 32-bit word in cache line
- Efficiently supports mixed integer/FP structures

IU and FPU Facts

Process Technology	0.7µ CMOS (0.5µ L-eff), 3-Metal
Cycle Time	75 MHz
Die Size	17.3 mm × 17.2 mm
Transistor Count	IU 2.6M, FPU 830K
Package	591 pin CPGA, 382 signals
Voltage	3.3 V
Power Dissipation	< 15 W each @ 75 MHz
Interface	LVTTL

IU and FPU Facts (cont)

Issue Rate	Four instructions per cycle
Address Space	2^48 virtual, 2^40 physical
I-Cache, D-cache	16 KB each, direct mapped Single cycle latency
External Cache	1..16 MB, 12 ns Sync. SRAM Five cycle latency, pipelined
TLB	384 entries, 3-way set assoc.
Floating-Point Unit	Two 4-stage mpy-add pipelines Divide latency: 20 cycles

Summary

Supercomputer performance at microprocessor prices
- 300 MIPS, 300 MFLOPS peak
- Handful of CMOS chips

Advanced yet low-cost memory hierarchy
- Split level cache makes optimal use of on-, off-chip RAM
- Address bellow greatly improves interleaved cache efficiency

4-way superscalar
- Delivers bandwidth of vector architecture
- Solves binary compatibility problem