

# **Power Awareness through Selective Dynamically-Optimized Traces**

**Roni Rosner, Yoav Almog, Naftali Schwartz, Avi Mendelson,  
Micha Moffie, Ari Schmorak and Ronny Ronen**

**Intel Labs, Haifa**

**Compiler & Architecture Seminar at IBM Haifa  
November 2003**



<http://www.intel.com/research/mrl/research/arch.htm>

**Intel Labs**  
IDC/MTL Architecture Research

# Agenda

***PARROT =***  
***Power-aware ARchitecture Running***  
***Optimized Traces***

- Concept and micro-architecture
- Performance and energy results
- Dynamic optimizations



# Trends in a Nut-Shell

- ☺ Ideally, each process technology generation provides
  - Reduction in transistor switching energy by 3X
  - Reduction in transistor delay time by 1.5X
  - Reduction in transistor switching power by 2X.
- ☺ More instructions per cycle, more cycles per second
  - ➔ gains performance but consumes power...
- ☹ We pay more (energy) than we get (performance)
  - More work per instruction (transistor switches)
  - More instructions per task (deep speculations)
  - More leakage power
  - ➔ Higher energy consumption per task!
- Parrot tries to change the balance!

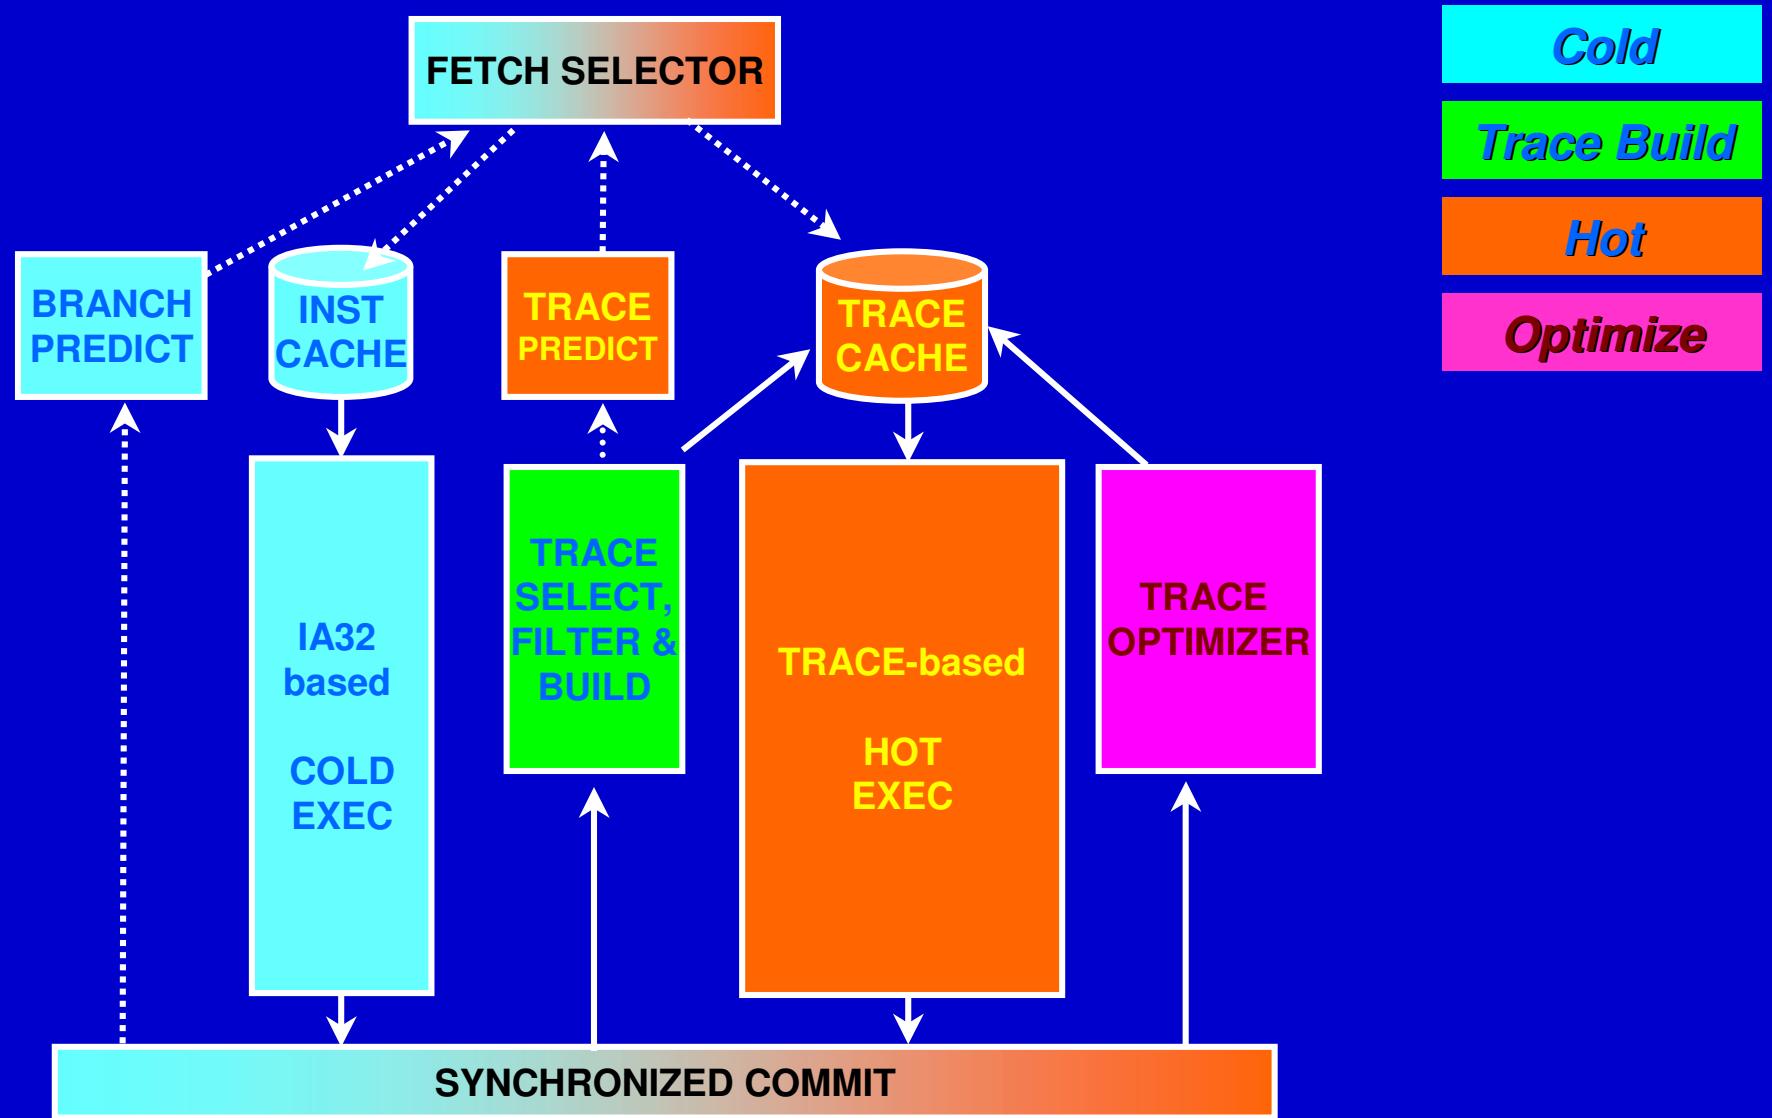
# Parrot Concept and Microarchitecture

Microarchitecture to identify hot traces, optimize them, execute efficiently – awareness of both performance and power

- Extend the well known cold/hot (10 code / 90 execution) paradigm
  - frequently used code behaves differently
    - More regular, predictable
    - Coarser granularity (longer sequences of instructions)
- Parrot Principles
  - **Reuse** → trace-cache centric
  - **Dynamic optimizations** → more performance with less energy, hw-oriented
  - **Focus** → invest where it pays
  - **Asymmetric decoupling** → hybrid front-end, cold and hot execution pipelines
  - **Transparency** → immune to s/w compatibility, overcome legacy



# Parrot 4-Phase Scheme



# Parrot in Context

## ■ Software techniques:

- JIT, profile-based compilation, binary-translators (Dynamo, Daisy, FX32!, IA32 Execution Layer for IA64), code morphing (Transmeta)

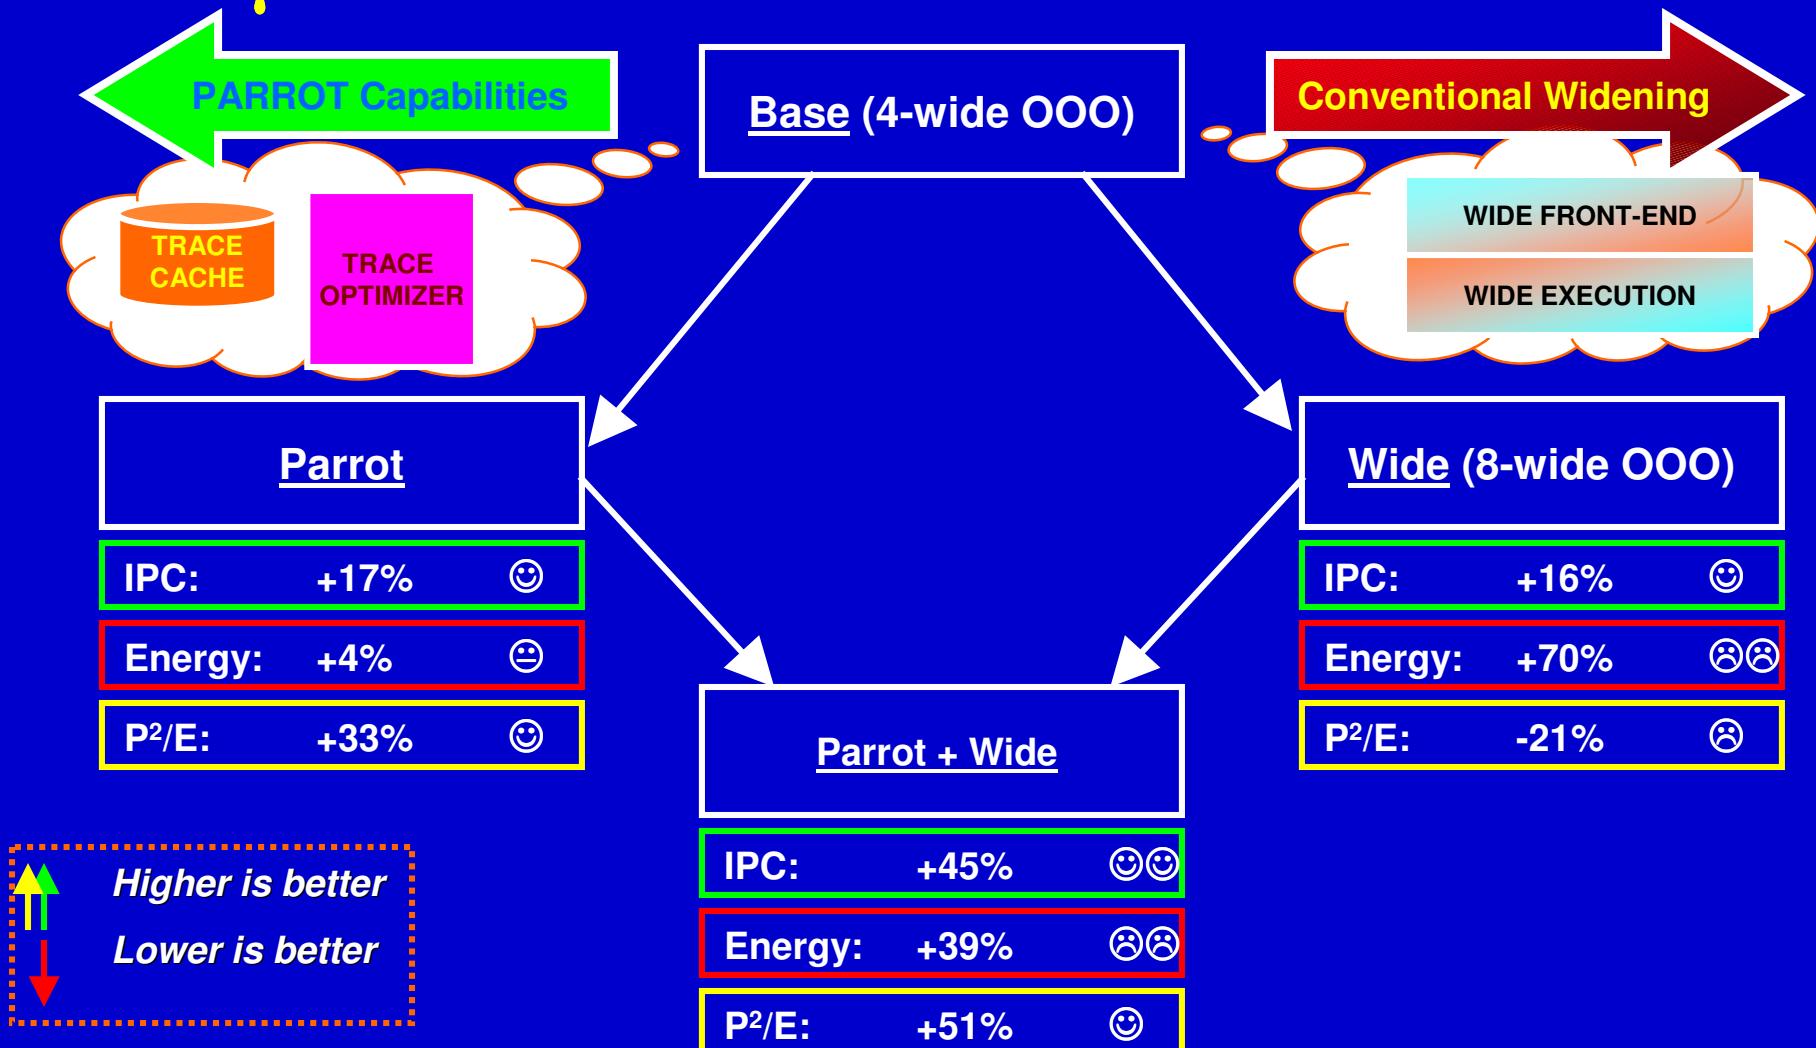
## ■ Micro architectural techniques

- Front end optimizations
  - » Pentium® 4, UOP cache, RePlay (U of Illinois)
- Asymmetric execution
  - » Turbo-Scalar [Shen], DIVA [Austin]

## ■ Our uniqueness

- Filtering [PACT'01]
- Smart trace selection [ICS'03]
- Full IA32 context, including X87 floating-point model
- Optimization for specialized execution [paper submitted]
- Putting it *all* together [paper submitted]

# $\mu$ Architecture Trade-Offs



**PARROT exhibits comparable performance and better power-awareness (P<sup>2</sup>/E)**

# Dynamic Optimizations

## Optimization Classes

- Trace selection
  - Loop unrolling
  - Procedure inlining
- Basic transformation
  - Virtual renaming, SSA
  - FP-stack flattening
  - Handle partial regs
  - Color regs: *LI*, *tmp*, *LO*
  - Replace *CTI* by *ASSERT*
- Generic optimizations
  - Logic, arithmetic simplifications
  - Propagation of values, registers, conditions
  - Dead code elimination
- Core-specific
  - Uop fusion
  - SIMDification
  - Semi-dynamic scheduling

## Benefits of Optimizations

- Less uops
  - Higher performance
  - Reduced energy consumption
- Reduced dependencies
  - Higher ILP → higher performance

## Efficiency of Optimizations

- Utilization
  - High rate of *execute / optimize*
- Sensitivity
  - High benefits at low frequency of optimization



# Optimization Example [SysMark, MS Word]

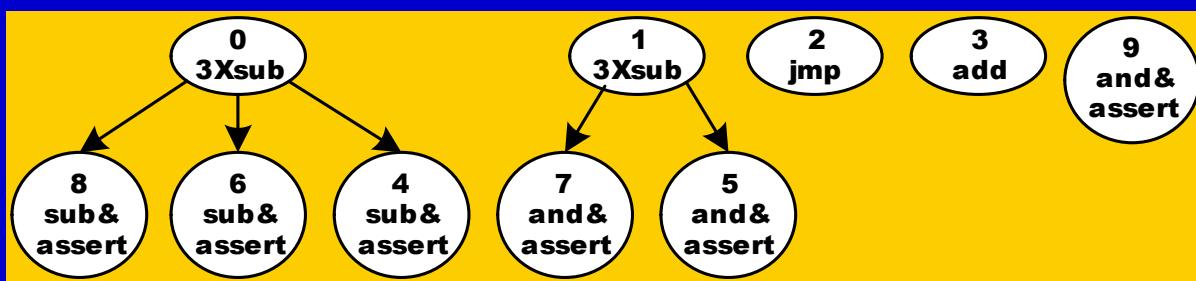
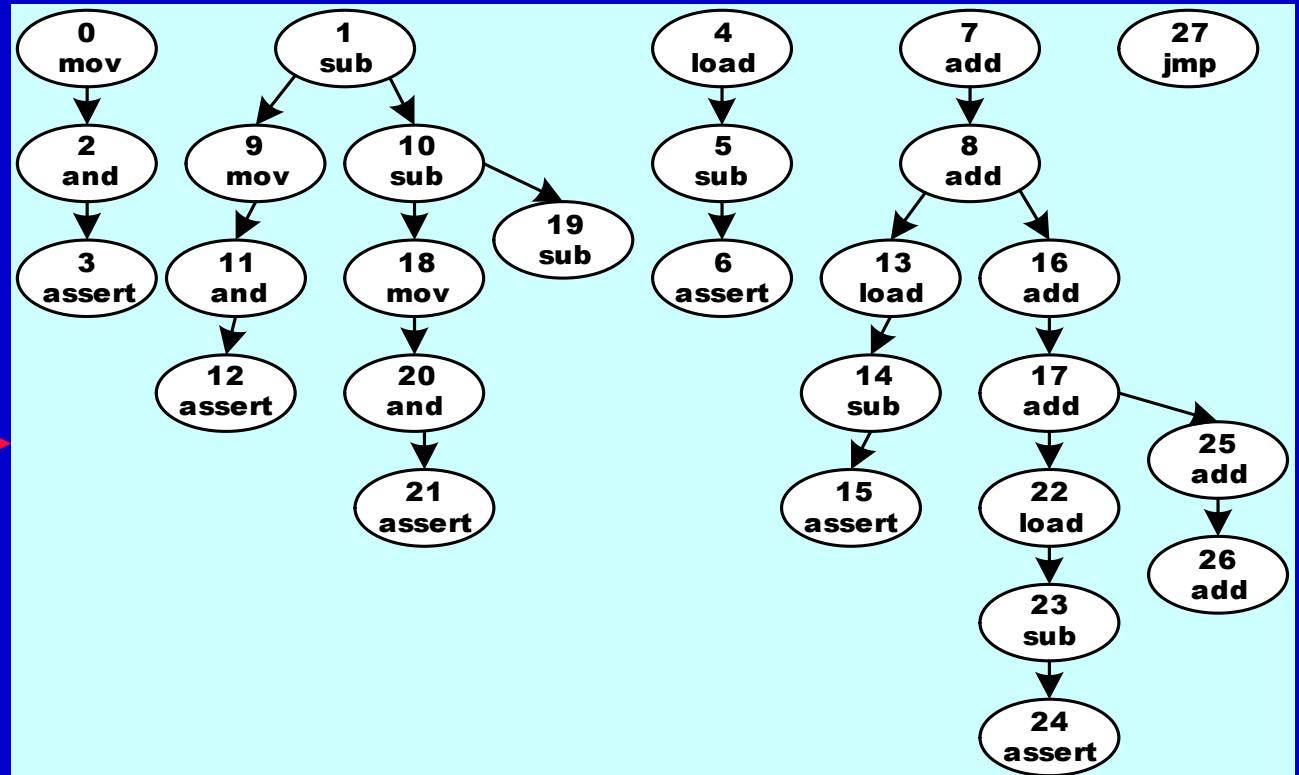
IA32 Instructions	Original PARROT uops	uops after basic-transformations	uops after optimizations
<pre> 0. <b>mov</b> eax ← edi 1. <b>dec</b> edi 2. <b>testl</b> eax, eax 3. <b>je</b> +0x14 4. <b>mov</b> cx ← [esi] 5. <b>cmp</b> cx, 256 6. <b>ja</b> 0x2c645e 7. <b>inc</b> esi 8. <b>inc</b> esi 9. <b>jmp</b> -23 -----</pre>	<pre> 0. eax ← <b>mov</b> (edi) 1. edi, eflags ← <b>sub</b> (edi, 1) 2. eflags ← <b>and</b> (eax, eax) 3. <b>cond_jmp</b>(e,eflags) 4. cx ← <b>load</b> (ds, esi) 5. eflags ← <b>sub</b> (cx, 256) 6. <b>cond_jmp</b> (nbe, eflags) 7. esi, eflags ← <b>add</b> (esi, 1) 8. esi, eflags ← <b>add</b> (esi, 1) 9. <b>jmp</b>(0x3004d8a5)</pre>	<pre> 0. v_a(0) ← <b>move</b> (li_edi) 1. v_b(0), v_flags(0) ← <b>sub</b> (li_edi, 1) 2. v_flags(1) ← <b>and</b> (v_a(0), v_a(0)) 3. <b>assert_cond</b> (v_flags.z(1), e, ntaken) 4. v_c(1) ← <b>load</b> (li_ds, li_esi) 5. v_flags(2) ← <b>sub</b> (v_c(1), 256) 6. <b>assert_cond</b>(v_flags.cz(2),nbe, ntaken) 7. v_b(1), v_flags(3) ← <b>add</b> (li_esi, 1) 8. v_b(2), v_flags(4) ← <b>add</b> (v_b(1), 1) -----</pre>	<pre> 0. v_c(1),v_c(3), lo_cx ← 3X<b>load</b>(li_ds,li_esi) 1. v_b(0),v_a(2),v_b(3),lo_eax, lo_edi ← 3X<b>sub</b>(li_edi,1, li_edi,2, li_edi, 3) 2. <b>jmp</b> (0x3004d8a5) 3. lo_esi ← <b>add</b>(li_esi, 0x6) 4. lo_eflags ← <b>sub&amp;assert</b>(lo_cx,256,nbe,ntaken) 5. <b>and&amp;assert</b>(v_b(3),v_b(3), e, ntaken) 6. <b>sub&amp;assert</b>(v_c(3), 256, nbe, ntaken) 7. <b>and&amp;assert</b>(v_b(0), v_b(0),e, ntaken) 8. <b>sub&amp;assert</b>(v_c(1), 256, nbe, ntaken) 9. <b>and&amp;assert</b>(li_EDI, li_EDI, e, ntaken)</pre>
<pre> 10. <b>mov</b> eax ← edi 11. <b>dec</b> edi 12. <b>testl</b> eax, eax 13. <b>je</b> +0x14 14. <b>mov</b> cx ← [esi] 15. <b>cmp</b> cx, 256 16. <b>ja</b> 0x2c645e 17. <b>inc</b> esi 18. <b>inc</b> esi 19. <b>jmp</b> -23 -----</pre>	<pre> 10. eax ← <b>mov</b> (edi) 11. edi, eflags ← <b>sub</b> (edi, 1) 12. eflags ← <b>and</b> (eax, eax) 13. <b>cond_jmp</b>(e,eflags) 14. cx ← <b>load</b> (ds, esi) 15. eflags ← <b>sub</b> (cx, 256) 16. <b>cond_jmp</b> (nbe, eflags) 17. esi, eflags ← <b>add</b> (esi, 1) 18. esi, eflags ← <b>add</b> (esi, 1) 19. <b>jmp</b>(0x3004d8a5)</pre>	<pre> 9. v_a(2) ← <b>move</b> (v_b(0)) 10. v_b(3), v_flags(5) ← <b>sub</b> (v_b(0), 1) 11. v_flags(6) ← <b>and</b> (v_a(2), v_a(2)) 12. <b>assert_cond</b> (v_flags.z(6),e, ntaken) 13. v_c(3) ← <b>load</b> (li_ds,v_b(2)) 14. v_flags(7) ← <b>sub</b> (v_c(3), 256) 15. <b>assert_cond</b>(v_flags.cz(7),nbe, ntaken) 16. v_b(4), v_flags(8) ← <b>add</b> (v_b(2), 1) 17. v_b(5), v_flags(9) ← <b>add</b> (v_b(4), 1) -----</pre>	
<pre> 20. <b>mov</b> eax ← edi 21. <b>dec</b> edi 22. <b>testl</b> eax, eax 23. <b>je</b> +0x14 24. <b>mov</b> cx ← [esi] 25. <b>cmp</b> cx, 256 26. <b>ja</b> 0x2c645e 27. <b>inc</b> esi 28. <b>inc</b> esi 29. <b>jmp</b> -23 -----</pre>	<pre> 20. eax ← <b>mov</b> (edi) 21. edi, eflags ← <b>sub</b> (edi, 1) 22. eflags ← <b>and</b> (eax, eax) 23. <b>cond_jmp</b>(e,eflags) 24. cx ← <b>load</b> (ds, esi) 25. eflags ← <b>sub</b> (cx, 256) 26. <b>cond_jmp</b> (nbe, eflags) 27. esi, eflags ← <b>add</b> (esi, 1) 28. esi, eflags ← <b>add</b> (esi, 1) 29. <b>jmp</b>(0x3004d8a5)</pre>	<pre> 18. lo_eax ← <b>move</b> (v_b(3)) 19. lo_edi,v_flags(10)←<b>sub</b>(v_b(3), 1) 20. v_flags(11)←<b>and</b>(lo_eax,lo_eax) 21. <b>assert_cond</b>(v_flags.z(11), e, ntaken) 22. lo_cx ←<b>load</b>(li_ds,v_b(5)) 23. lo_eflags ←<b>sub</b>(lo_cx, 256) 24. <b>assert_cond</b>(lo_eflags(65),nbe, ntaken) 25. v_b(6),v_flags(12) ← <b>add</b> (v_b(5), 1) 26. lo_esi,v_flags(13)←<b>add</b>(v_b(6), 1) 27. <b>jmp</b> (0x3004d8a5)</pre>	



# Optimization Example

Dependency graph  
of the trace after  
renaming & SSA,  
*before optimizations*

- 28 micro-ops
- tree-height 7

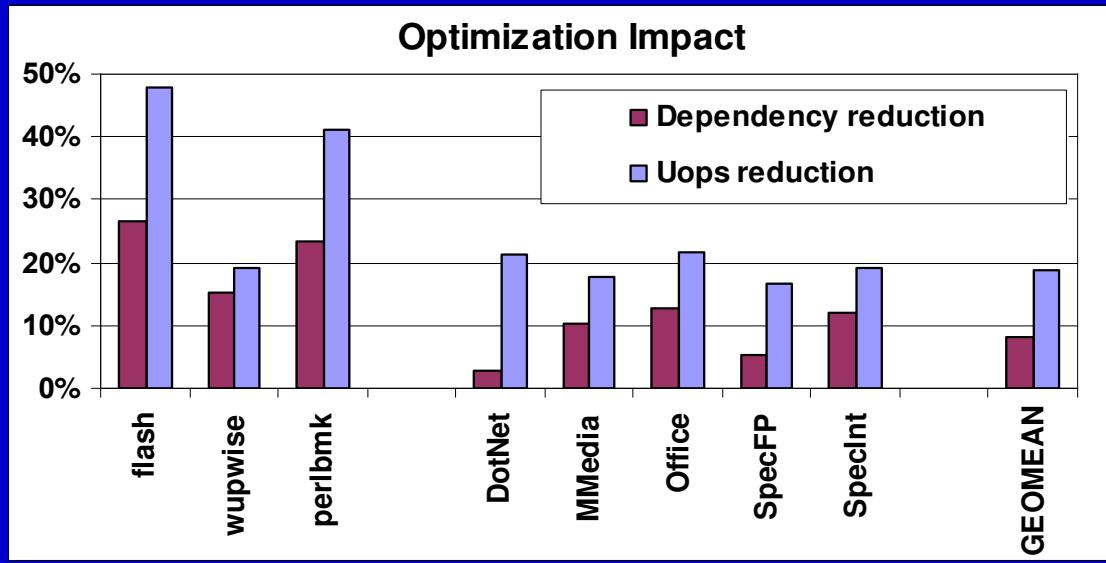


Dependency graph  
of the trace *after*  
*optimizations*

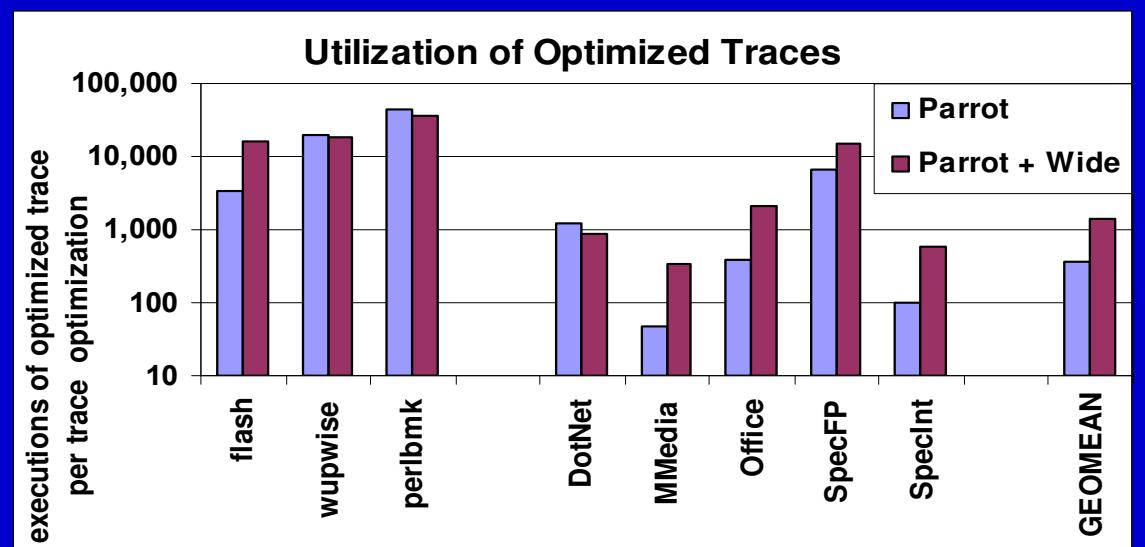
- 10 micro-ops
- Tree-height 2



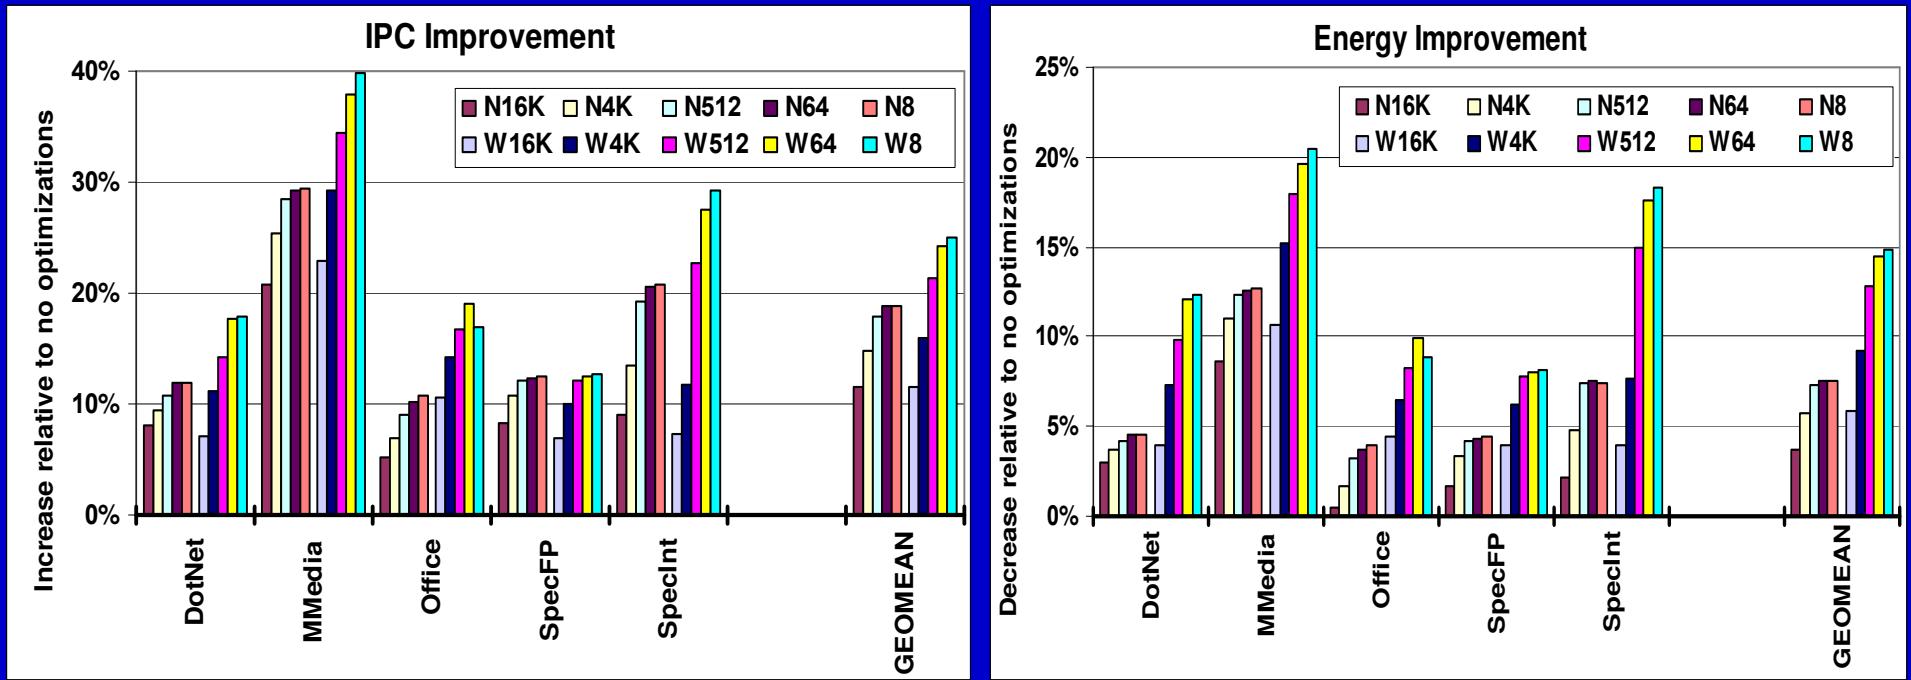
# Optimizer Quality



Higher is better →



# Optimizer Sensitivity: Impact of Filter's Threshold on IPC, Energy



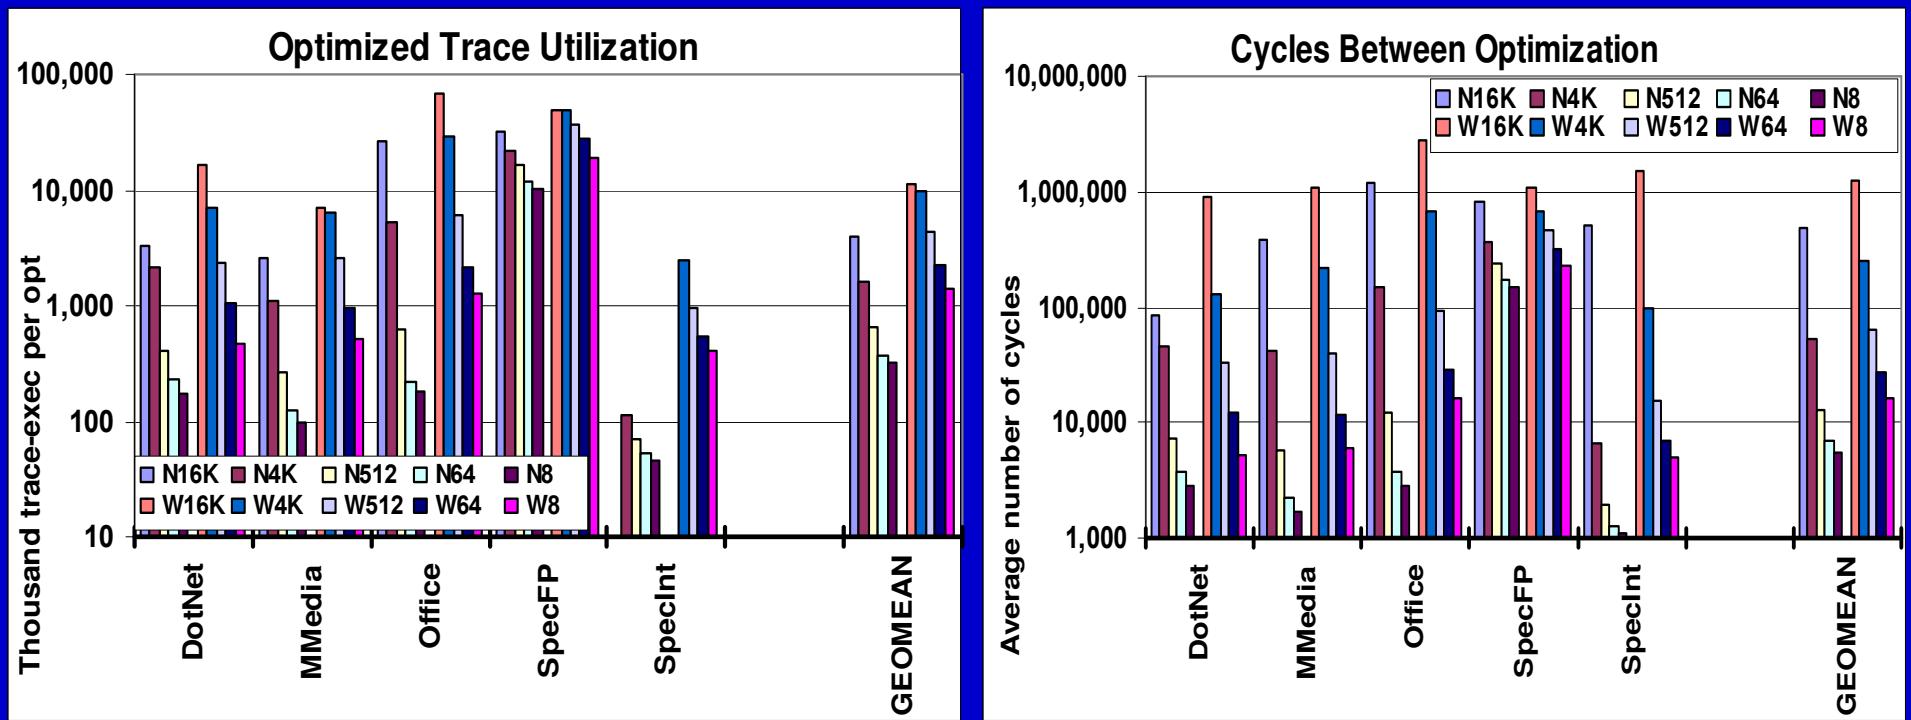
- Higher bars are better for overall performance, energy consumption
- Higher threshold  $j$  represent less sensitivity to optimizer cost, latency
- Tradeoff point: “knee” where smaller threshold has a low marginal contribution to performance ~512

## Legend

$N_j$	- Parrot, optimizer threshold is $j$
$W_j$	- Parrot+Wide, optimizer threshold is $j$

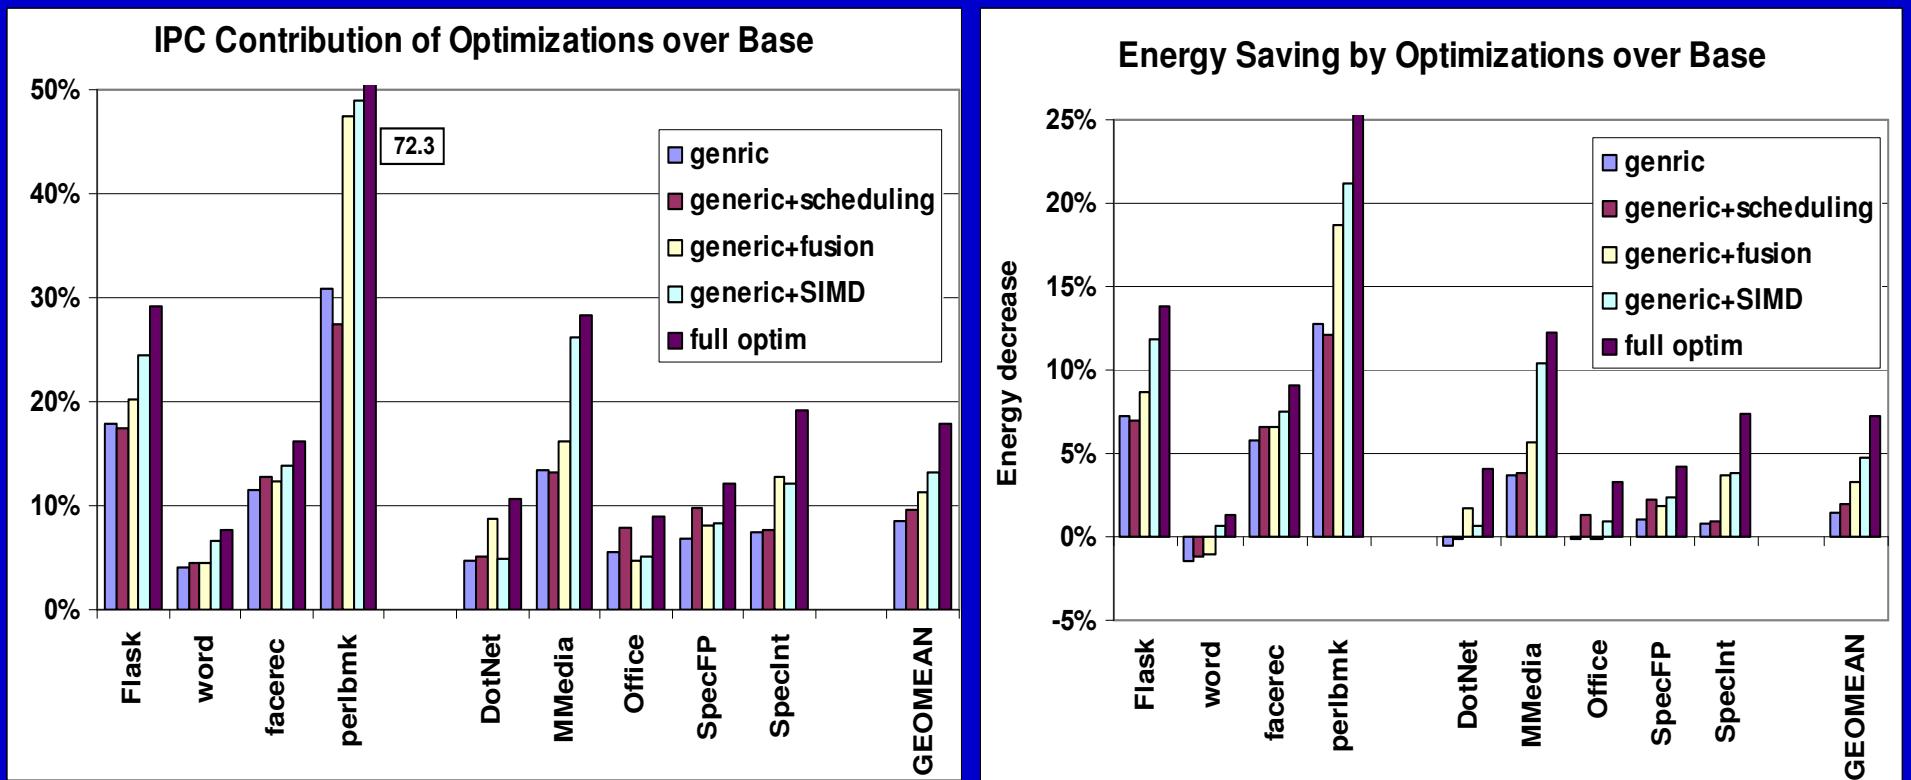


# Optimizer Sensitivity: Impact of Filter's Threshold on IPC, Energy



*Relaxed design of optimizer is feasible!!!*

# Optimizations Breakdown: Generic vs. Core-Specific



Higher is better

# Conclusions and Future

## Major conclusions

- Parrot-style μarchitecture presents a power-aware alternative to conventional design
- Parrot μarch shows low sensitivity to optimizer latency, greediness  
→ feasibility of relaxed optimizer
- Core specific optimizations are central for hardware-based dynamic optimizations

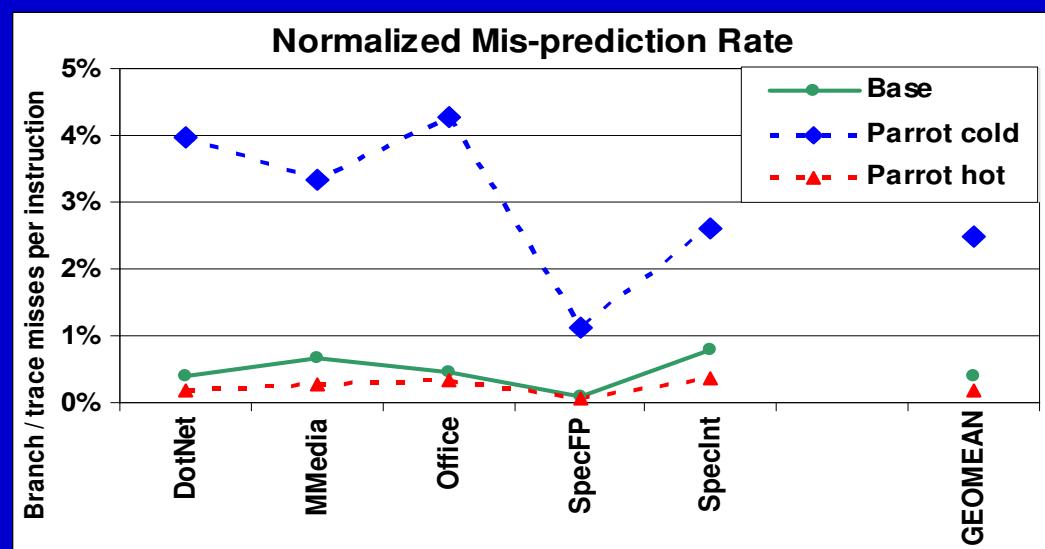
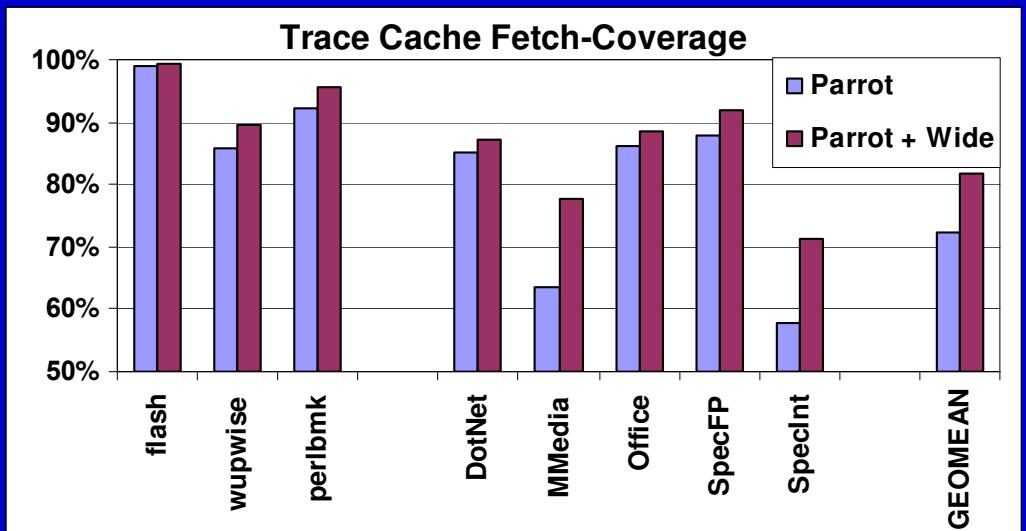
## Interesting Future Directions

- Microarchitecture
  - Improved selection, prediction
  - Specialized traces, value prediction
  - Multi-threading
- Architecture and software
  - Interaction with compilers, profiling, JIT
  - Quality of service
- Improved dynamic optimizations

# Backup Slides



# Front-End Characteristics



Higher is better  
Lower is better