

"JAGUAR"

AMD's Next Generation Low Power x86 Core

Jeff Rupley, AMD Fellow
Chief Architect / Jaguar Core
August 28, 2012



TWO X86 CORES TUNED FOR TARGET MARKETS

“Bulldozer Family”

Performance
& Scalability


Mainstream Client and Server Markets



“Cat Family”

Flexible,
Low Power
& Small

Small Die Area



Low Power Markets




Cloud Clients Optimized



"JAGUAR" DESIGN GOALS

- Improve on "Bobcat": performance in a given power envelope
 - More IPC
 - Better Frequency at given Voltage
 - Improve power efficiency thru clock gating and unit redesign
- Update the ISA/Feature Set
- Increase Process Portability



ISA/FEATURE SET

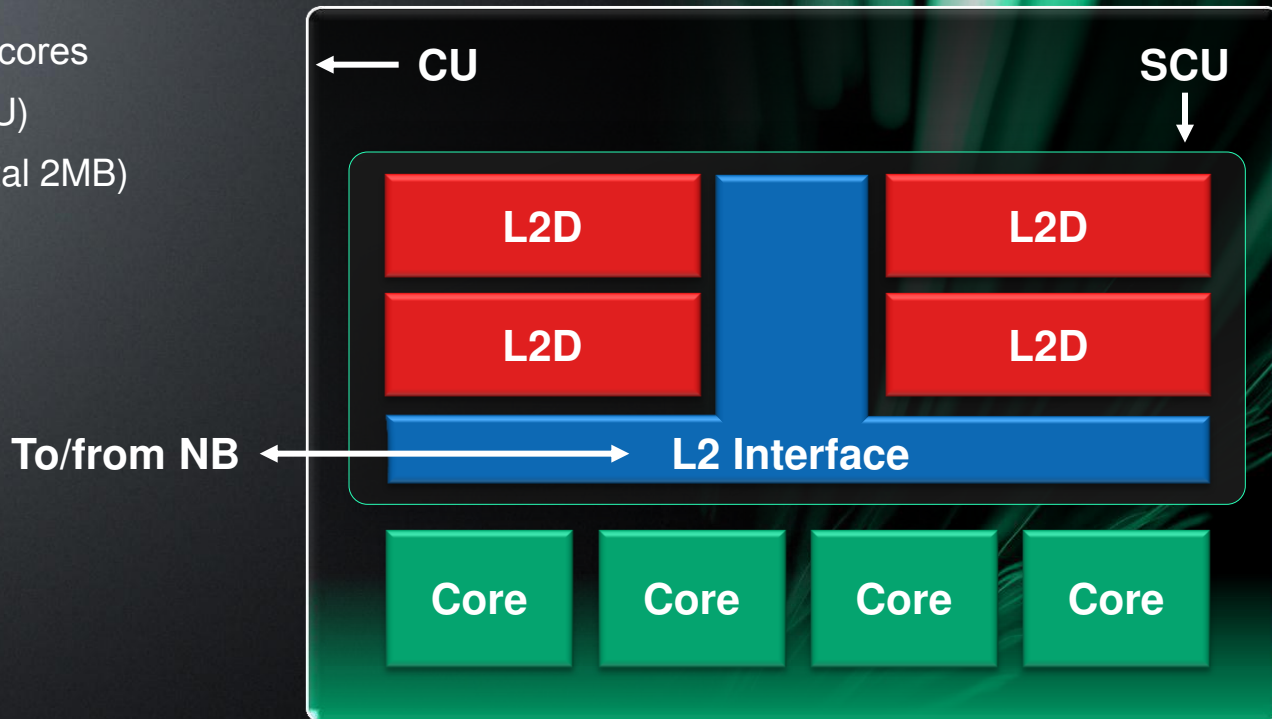
- ISA: “Bobcat” baseline of AMD64 x86 ISA w/ SSE1-SSSE3, SSE4A

“Jaguar” added:

- SSE4.1, SSE4.2
 - AES, CLMUL
 - MOVBE
 - AVX, XSAVE/XSAVEOPT
 - F16C, BMI1
- 40 bit physical address capable
 - Improved virtualization

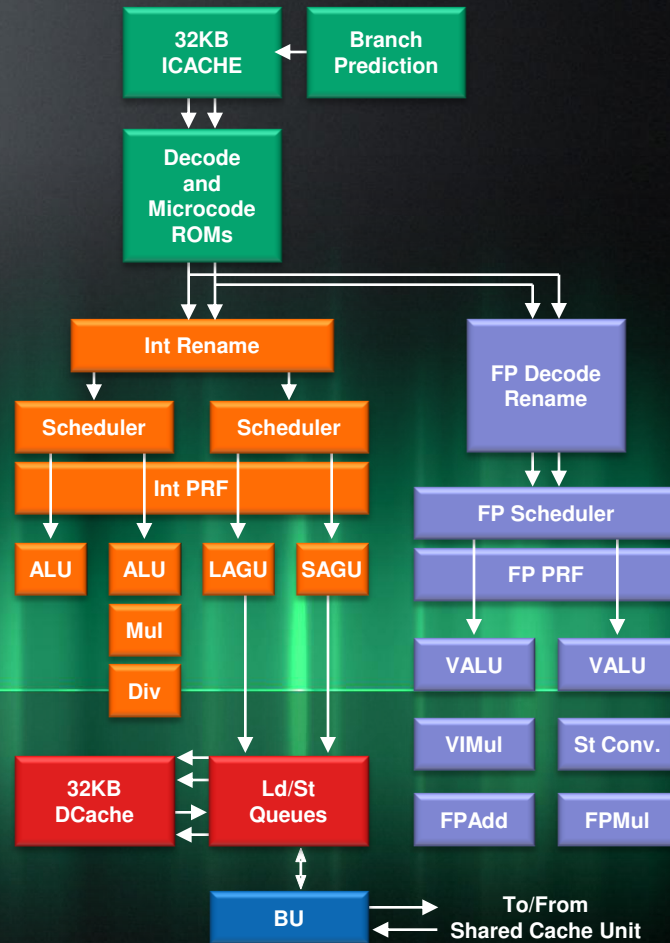
"JAGUAR" COMPUTE UNIT (CU)

- 4 Independent "Jaguar" cores
- Shared Cache Unit (SCU)
 - 4 L2 Data Banks (total 2MB)
 - L2 Interface Tile



"JAGUAR" CORE

Micro-Architecture

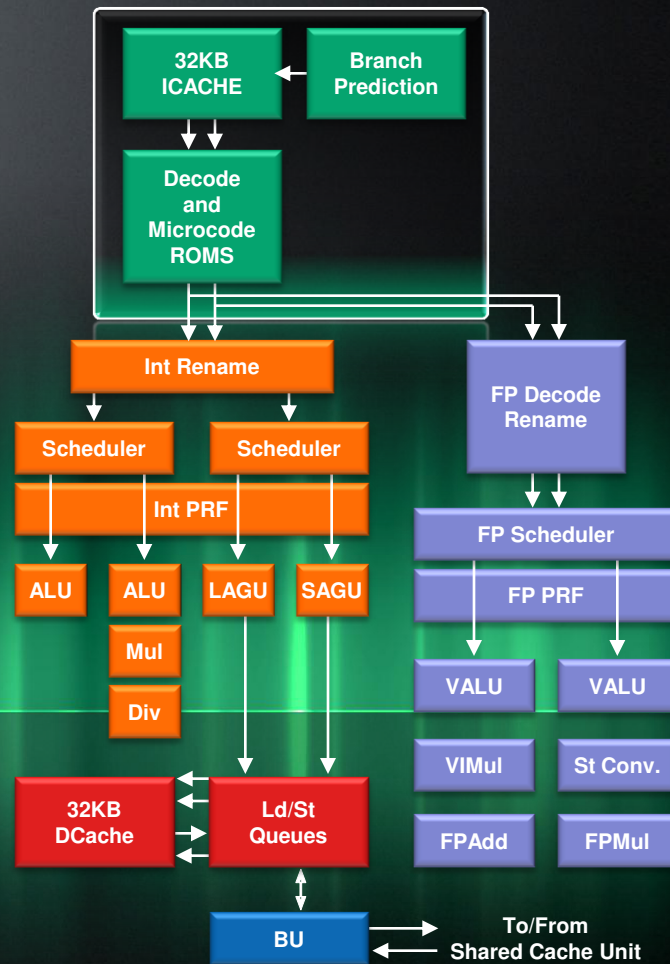


"JAGUAR" CORE

Frontend

Like "Bobcat":

- IC: 32KB, 2way
- Itlb: 512 4KB pages
- Layered branch predictor w/ state of the art conditional predictor
- 32B fetch
- 2-instruction decode



"Jaguar" Enhancements:

- 4x32B IC loop buffer for power
- Improved IC prefetcher for IPC
- Grew IB for improved fetch/decode decoupling
- Added decode stage for frequency

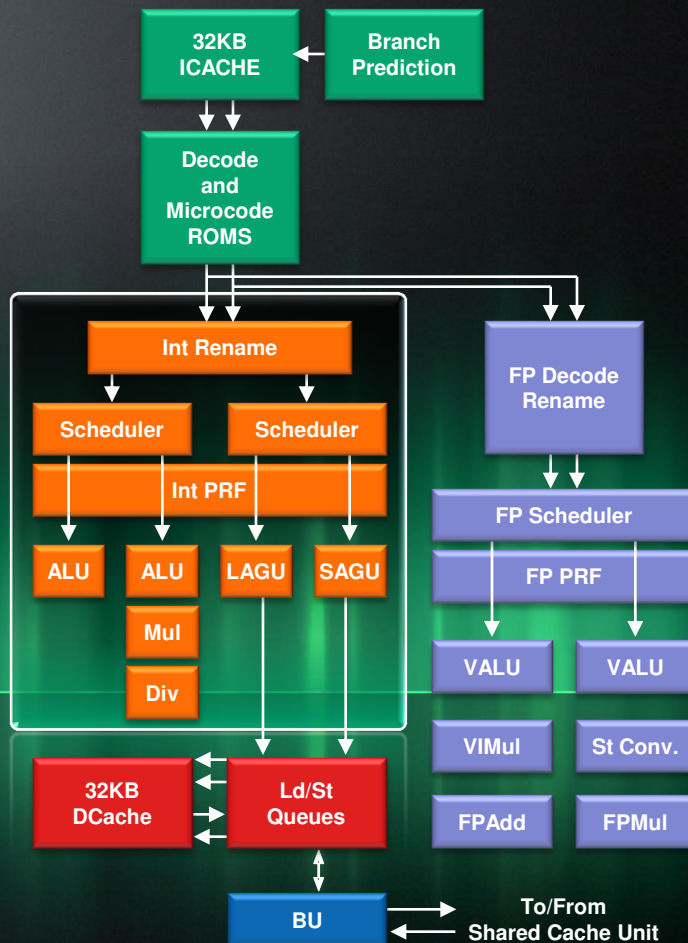
"JAGUAR" CORE

Integer Execution

Like "Bobcat":

Schedulers can issue

- 2 ALU
- 1 LD AGU
- 1 ST AGU



"Jaguar" Enhancements:


- New hardware divider (leveraged from Llano)
- New/improved cops: CRC32/SSE4.2, BMI1, POPCNT, LZCNT
- More OOO resources
Larger schedulers, ROB

"JAGUAR" CORE

Floating Point Unit

Like "Bobcat":

- 2 wide FP decode
- OOO scheduler
- 2 execution pipes



"Jaguar" Enhancements:

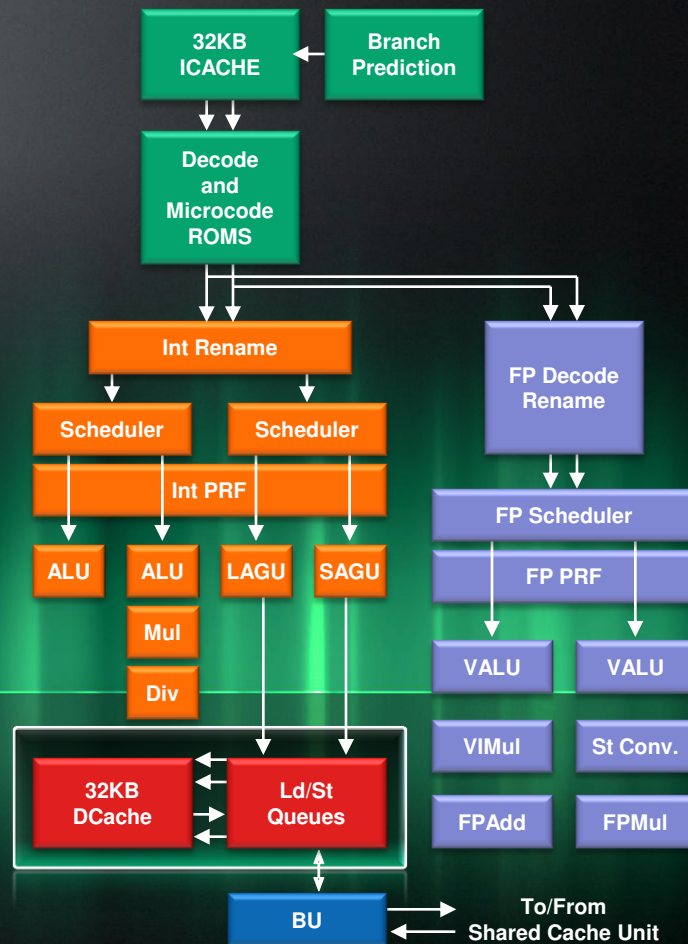
- 128b native hardware
 - 4 SP muls + 4 SP adds
 - 1 DP mul + 2 DP adds
- ISA: many new COPs
 - 256b AVX supported by double pumping 128b hardware
- New Zero Optimizations
- Second FPRF stage for frequency

"JAGUAR" CORE

Load/Store, Data Cache

Like "Bobcat":

- DC: 32KB, 8way
- L2DTLB: 512 4KB pages
- 8-stream DC prefetcher
- OOO LS



"Jaguar" Enhancements:

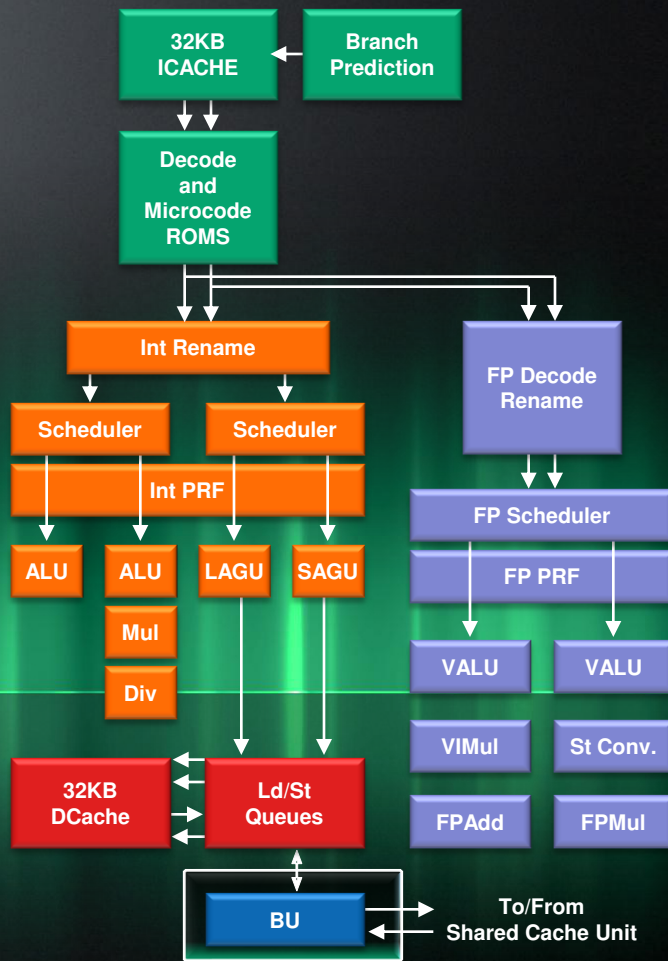
- Ld/St Queues redesign:
 - Improved OOO picker
 - Improved STLF
 - Less store data shuffling
 - More OOO resources
- Enhanced Tablewalks
- 128b data path to FPU

"JAGUAR" CORE

Bus Unit

Like "Bobcat":

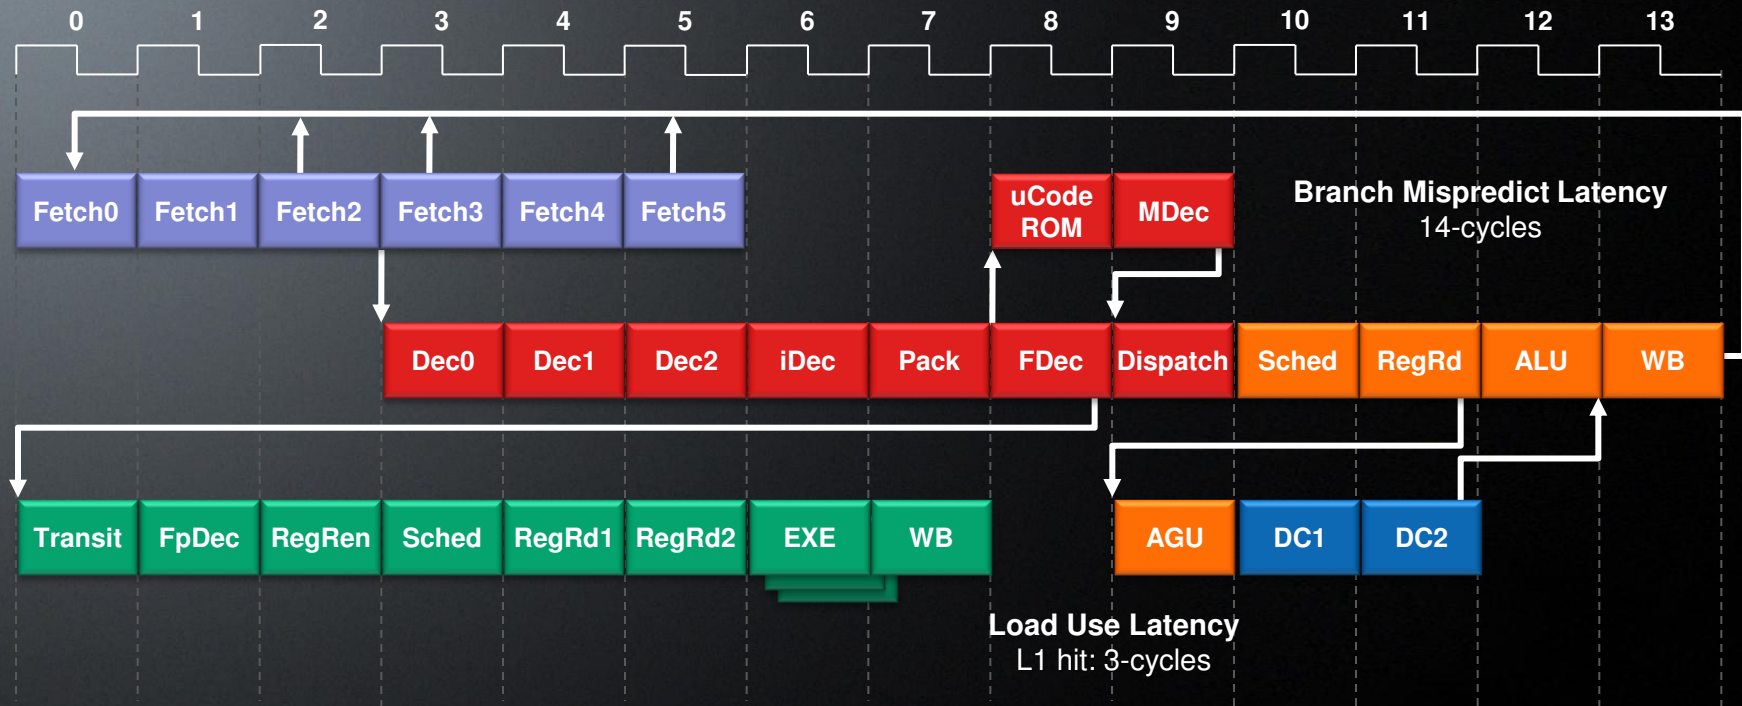
- BU interfaces between Core (I\$,D\$) and L2\$/NB



"Jaguar" Supports:

- 8 DC miss/prefetch
- 3 IC miss/prefetch
- Improved Write Combining with 4 WCB data buffers


"JAGUAR" CORE PIPELINE




Microarchitectural Frequency improvement over "Bobcat": >10%
One additional cycle branch mispredict latency vs. "Bobcat"




"JAGUAR" CORE FLOOR PLAN



CORE FLOOR PLAN COMPARISON



“Bobcat” core in 40nm = 4.9 mm²
7 core macros, 2 L2 macros,
3 clock macros



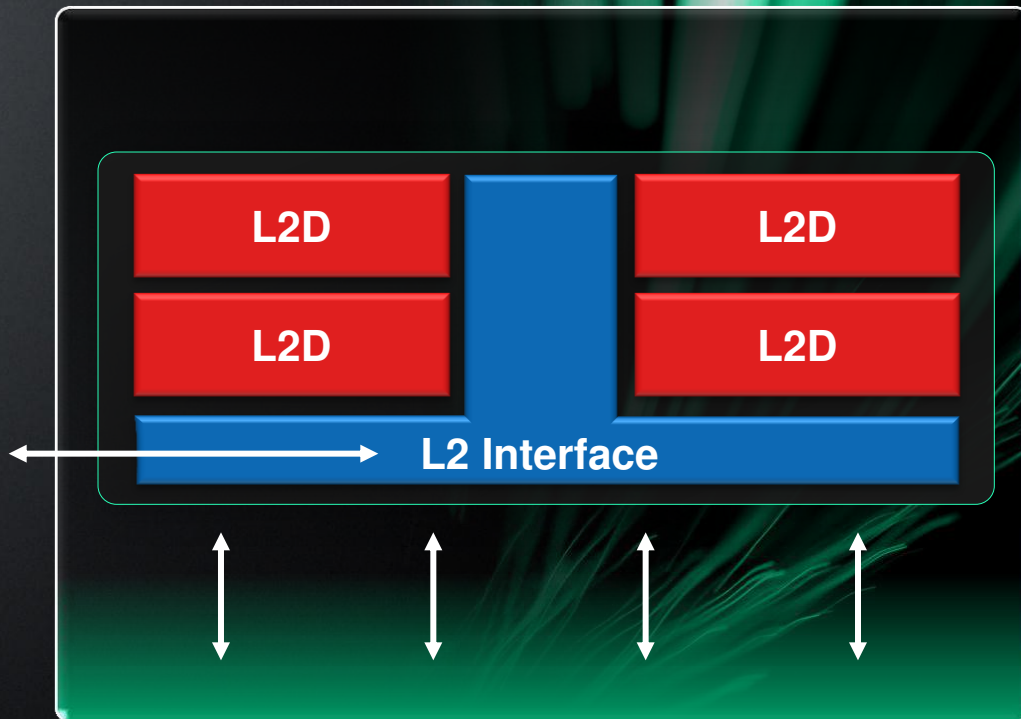
“Jaguar” core in 28nm = 3.1 mm²
3 core macros, 1 L2 macro,
1 clock macro



"JAGUAR" SHARED CACHE UNIT


Shared Cache is major design addition in "Jaguar"

- Supports 4 cores
- Total shared 2MB, 16-way
 - Supported by 4 L2D banks of 512KB each
- L2 cache is inclusive – allows using L2 tags as probe filter
 - Any line in a Core L1 instruction or data cache must be in the L2



"JAGUAR" L2 INTERFACE


- All connections routed thru L2 interface
- L2 tags reside in interface block
 - Divided into 4 banks
 - L2D bank lookup only after L2 tag hit
- L2 Interface block runs at core clock
 - L2D's run at half clock for power, only clocked when required
- New L2 stream prefetcher per core
 - Allows improved bandwidths & IPC
- Up to a total of 24 paired read + write transactions in flight
- 16 additional L2 snoop queue entries
 - Allows for handling coherent probes at high bandwidth



"JAGUAR" C6

- Any Core can independently go into CC6 power gating
 - Optimized microcode routines and hardware allow for fast CC6 entry/exit
 - Shared L2 leaves more cache for the remaining active cores (IPC)
- Last core in the compute unit to be power gated flushes shared L2 in preparation for full C6. Hardware engines added to improve L2 flush times.

Relative C6 Latencies Under Normalized Conditions



"JAGUAR" POWER

- Many blocks redesigned for improved power efficiency
 - IC Loop Buffer, Store Queue, L2 clocks, etc.
- Clock power usage scrubbed, including improved dynamic clock gating:


	"Bobcat" IPC	"Jaguar" IPC	"Bobcat" Core Gater Efficiency	"Jaguar" Core Gater Efficiency
Halt	0.00	0.00	91.8	98.8
Apps	0.95	1.10	89.7	92.3
"Bobcat" Virus	1.74	1.78	84.6	87.1
"Jaguar" Virus	0.81	1.86	85.7	85.0

- Increased frequency capability allows choices:
 - Higher frequency -> higher performance
 - Same frequency at lower voltage -> lower power



SUMMARY

- ISA enhancements
- Increased process portability
- Estimated typical IPC improvement: >15%*
- Frequency improvement: >10%*
- Dynamic power efficiency improvements



*Based on internal AMD modeling using benchmark simulations

DISCLAIMER & ATTRIBUTION

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. There is no obligation to update or otherwise correct or revise this information. However, we reserve the right to revise this information and to make changes from time to time to the content hereof without obligation to notify any person of such revisions or changes.

NO REPRESENTATIONS OR WARRANTIES ARE MADE WITH RESPECT TO THE CONTENTS HEREOF AND NO RESPONSIBILITY IS ASSUMED FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

ALL IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE ARE EXPRESSLY DISCLAIMED. IN NO EVENT WILL ANY LIABILITY TO ANY PERSON BE INCURRED FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Trademark Attribution

©2012 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Phenom, Radeon, and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Microsoft and DirectX are registered trademarks, of Microsoft Corporation in the United States and/or other jurisdictions. Other names used in this presentation are for identification purposes only and may be trademarks of their respective owners.

