

Third SHA-3 Candidate Conference

Performance Discussion

Bill Burr

burr44@gmail.com

23 March, 2012

How Did We Get Here?

- 2004-2005 New cryptanalysis
 - Wang, Biham, Joux, Kelsey....
 - Cast doubt on existing hash standards
- 2005-2006 NIST Hash Fun. Workshops
- 2007 NIST organized SHA-3 competition
 - Wanted a very secure alternative to SHA-2
 - 64 candidates submitted 31 Oct. 2008
- Mar. 2012 third SHA-3 Candidate Conf.
 - Now down to five “finalist” hash functions


What have we learned?

- A lot about hash functions
 - Cryptanalysis (applies to block ciphers too)
 - New, stronger constructions
- No free lunches
 - Collision resistance takes a lot of computation
 - Improving MD security increases state & computation
- SHA-2 is not bad
 - No apparent threat to collision resistance
 - Overall performance is fairly competitive

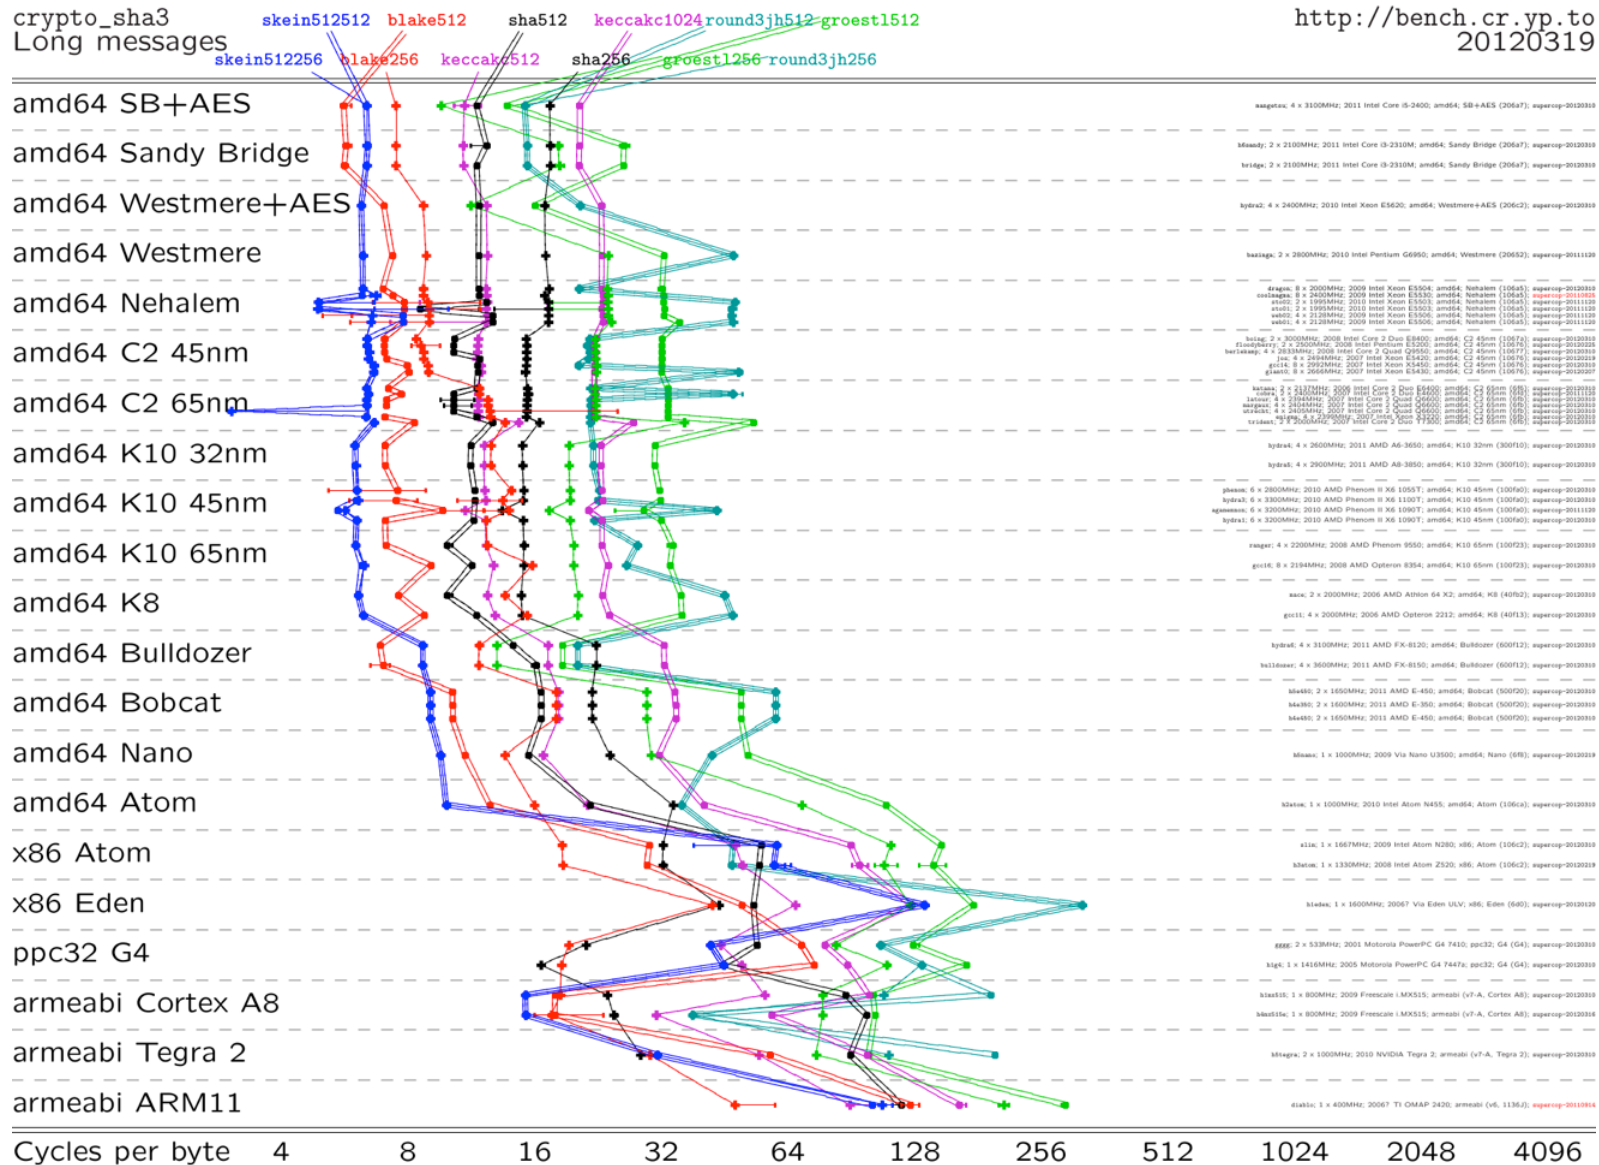
The Big Question

- Which Candidate best complements SHA-2?
 - All candidates have higher 2nd-preimage resistance than SHA-2 & fix the generic limitations of MD, but
 - SHA-2 is not apparently broken
 - SHA-2 collision resistance seems fine
 - SHA-2 performance overall is respectable
 - Some candidates have significantly better performance on some common platforms
 - More readily exploitable parallelism, but
 - A standard tree hashing mode may diminish this advantage
 - SHA-2 performance isn't the worst in any category
 - Some candidates offer extras
 - Wide block cipher, authenticated encryption

eBASH 18 Nov. 2011: SHA-3 Finalists + SHA2, long message



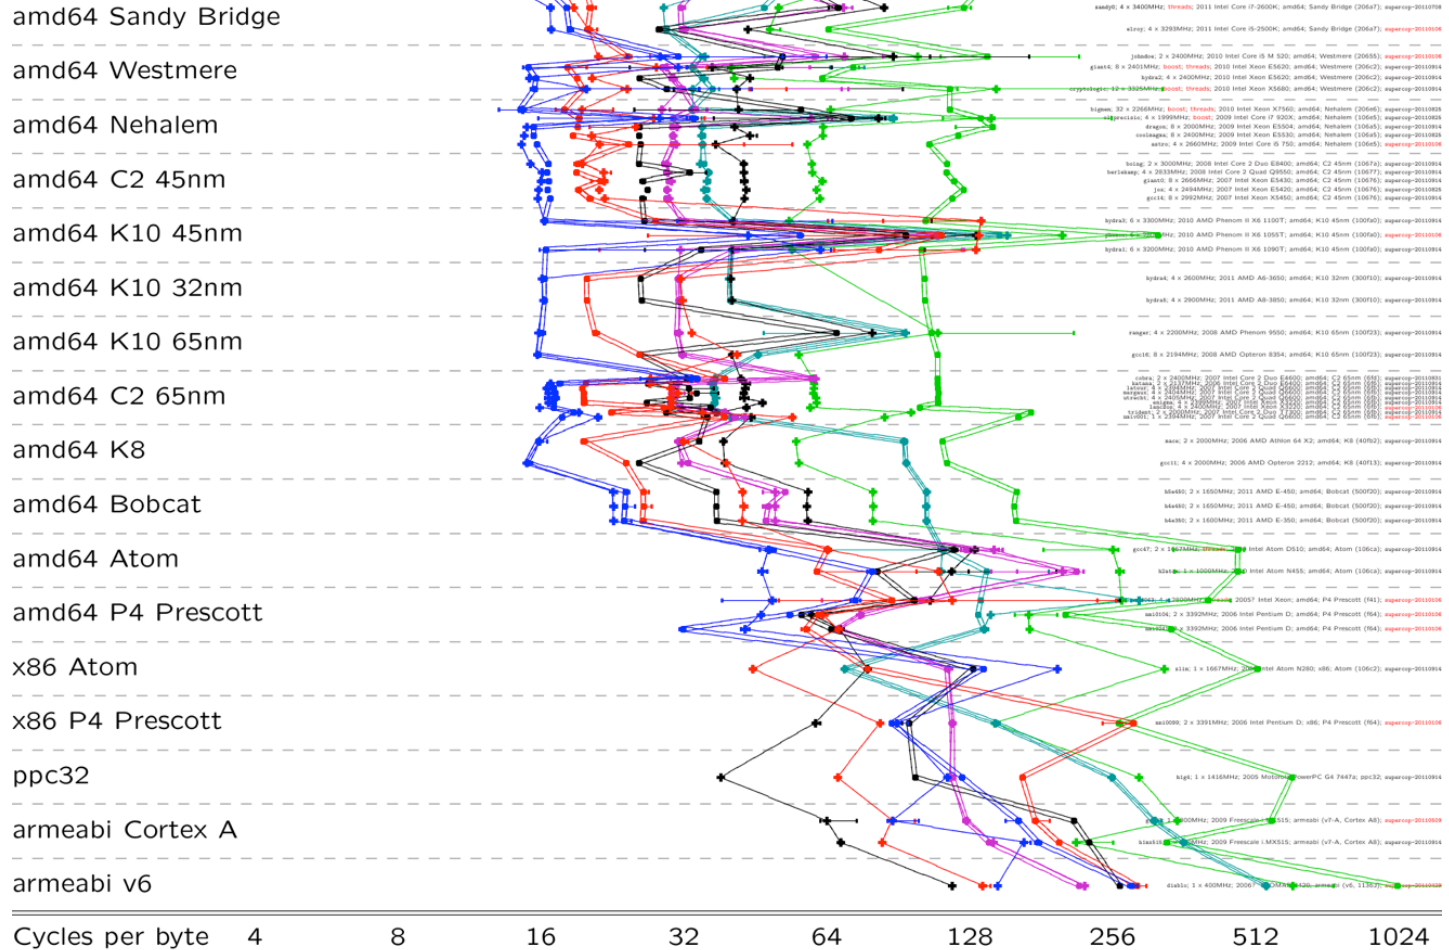
eBASH 19 Mar. 2012: SHA-3 Finalists + SHA2, long message



eBASH 18 Nov. 2011: SHA-3 Finalists + SHA2, 64-byte msg

crypto_sha3
64 bytes

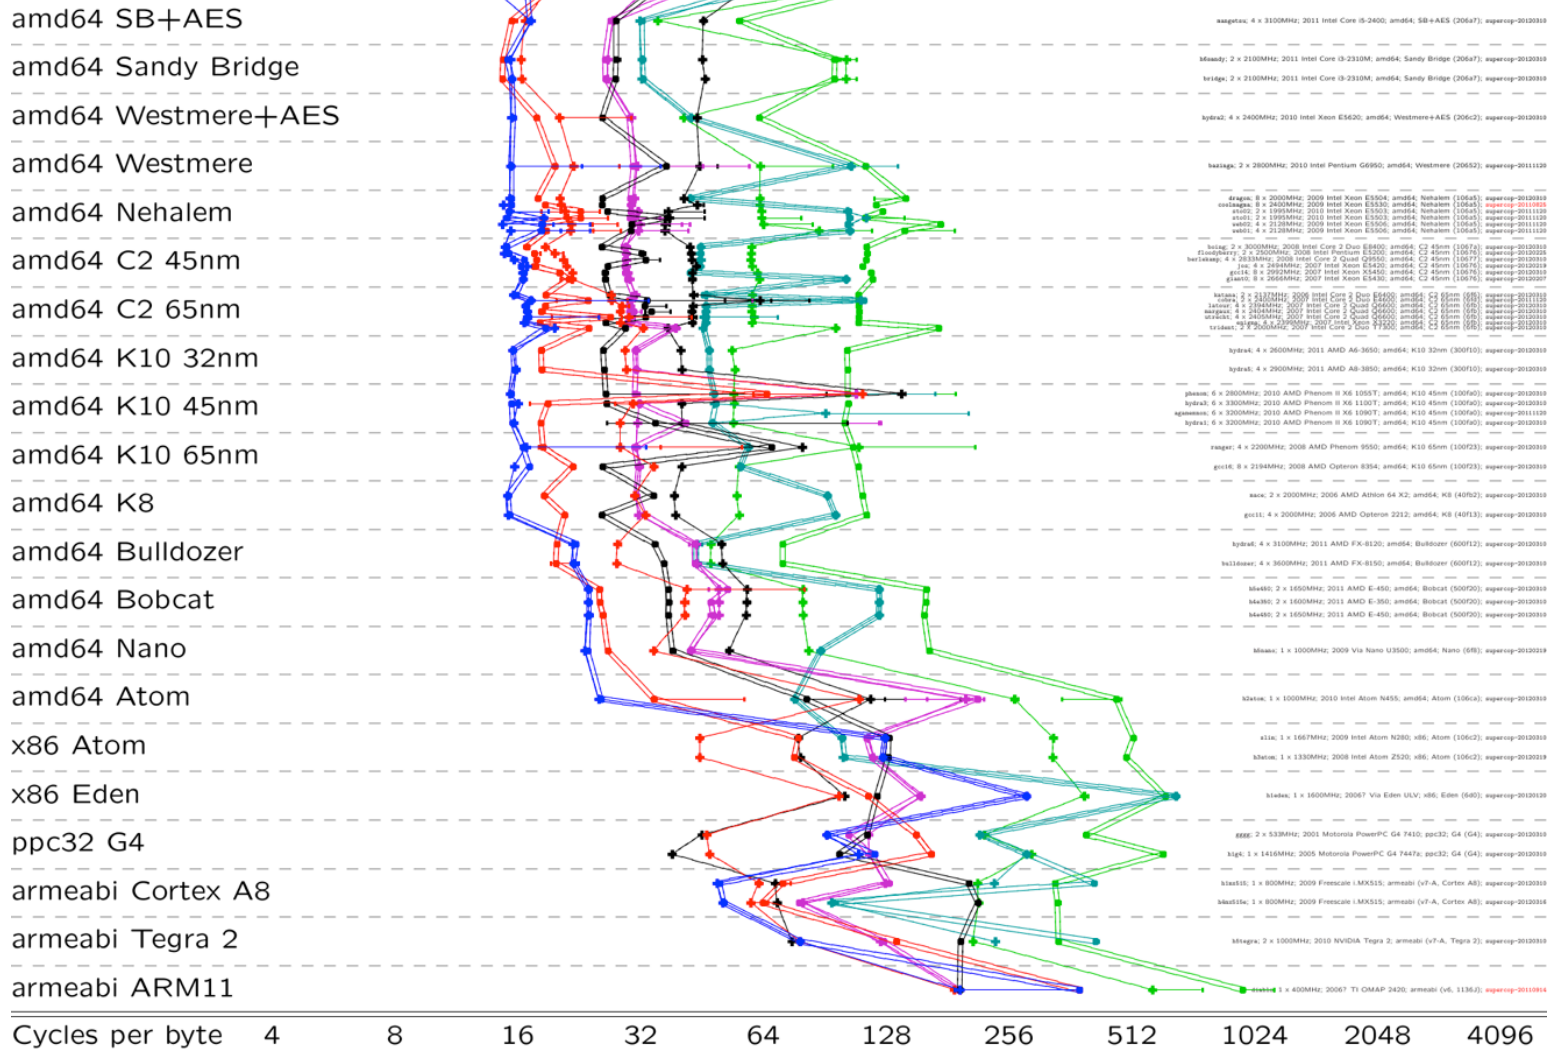
<http://bench.cr.yj.to>
20111118



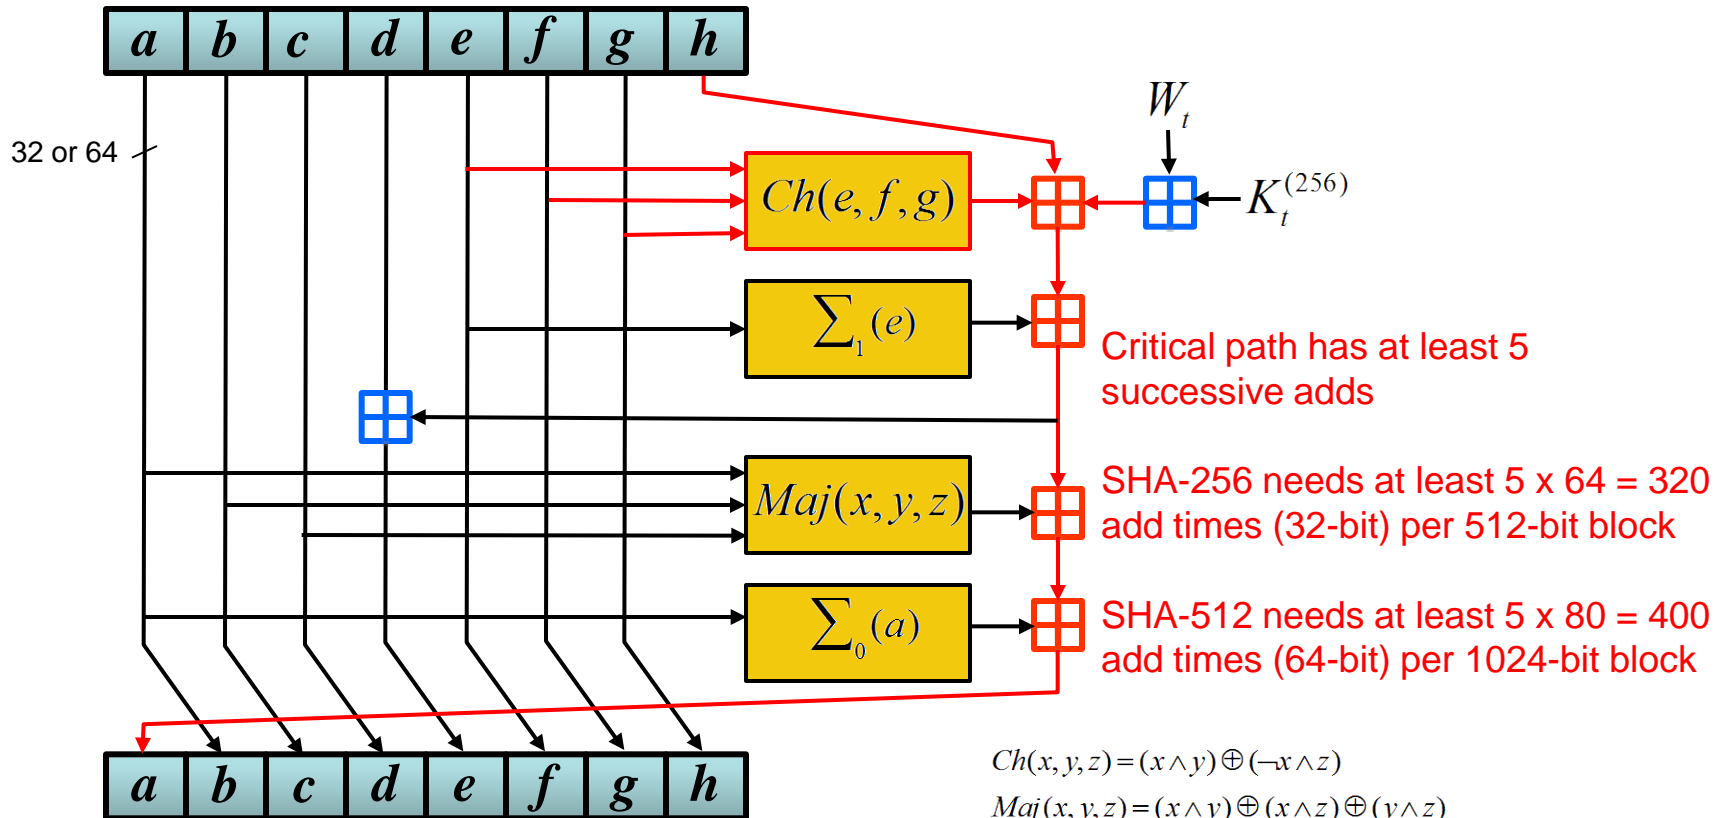
eBASH 19 Mar. 2012: SHA-3 Finalists + SHA2, 64-byte msg

crypto_sha3
64 bytes

http://bench.cr.yp.to
20120319



SHA-2 Round Function



$$Ch(x, y, z) = (x \wedge y) \oplus (\neg x \wedge z)$$

$$Maj(x, y, z) = (x \wedge y) \oplus (x \wedge z) \oplus (y \wedge z)$$

$$\Sigma_0^{256}(x) = ROTR^2(x) \oplus ROTR^{13}(x) \oplus ROTR^{22}(x)$$

$$\Sigma_1^{256}(x) = ROTR^6(x) \oplus ROTR^{11}(x) \oplus ROTR^{25}(x)$$


SHA-256 shown, SHA-512 has different ROTR constants

This figure is based on one in the Wikipedia article on SHA-2

SKEIN Round

72 Rounds

- 4 Parallel MIX ops/rnd
- 4 rounds shown
- Key added in every 4 rounds
- Rotations cycle every 8 rounds
- Easy to visualize how this vectorizes.



Performance Questions

- Which candidate's performance best complements SHA-256 & SHA-512?
- What performance weakness would really hurt? What applications are most performance sensitive and what SHA-3 candidate has a weakness that would affect the adoption of SHA-3 for current or future applications?

Performance Questions

- If we have a tree hashing mode, does speed of a single thread matter a lot?
- What performance issues haven't we yet considered?
- Should we give the same weight to 512 and 256-bit performance?
 - If the 512-bit variant is faster should we chop it as NIST did with SHA-512/256?

Performance Questions

- Divide the world into unconstrained & constrained implementations and into hardware & software. Then:
 - Which quadrants are most and least critical to the success of SHA-3?
 - Which constraints are most critical?
 - How “constrained” is an ARM with NEON?
 - Will NEON help all candidates similarly?
 - Which candidates would be helped by vector 64-bit rotates?
 - Do we have any good way to get or infer energy per bit hashed?
- Are there coming applications that could jump right into SHA-3 without a transition from SHA-1 or SHA-2?