ESSENCE

Simple. Parallel.

Jason Worth Martin

Department of Mathematics and Statistics
James Madison University

First SHA-3 Candidate Conference

ESSENCE: Classification Categories

- Narrow Pipe (two sizes, just like SHA-2)
- Davies-Meyer
- Internal Block Cipher uses Non-Linear Feedback Shift Register
- Merkle-Damgård at lowest level
- Merkle Tree at mid-level
- Merkle-Damgård at top level to combine tree roots

ESSENCE Structure

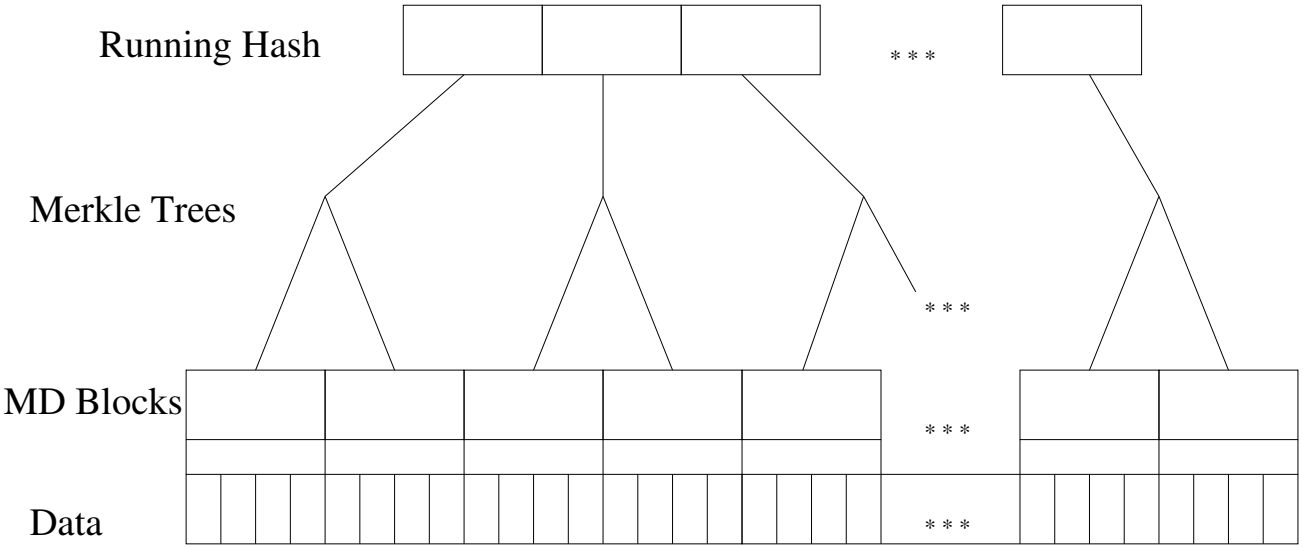

- Hash MD Blocks (each 1MByte) separately with varying IV
- Merkle Trees combine MD Block hashes
- Running Hash combines roots (uses separate IV and Final Block)

ESSENCE Block Cipher

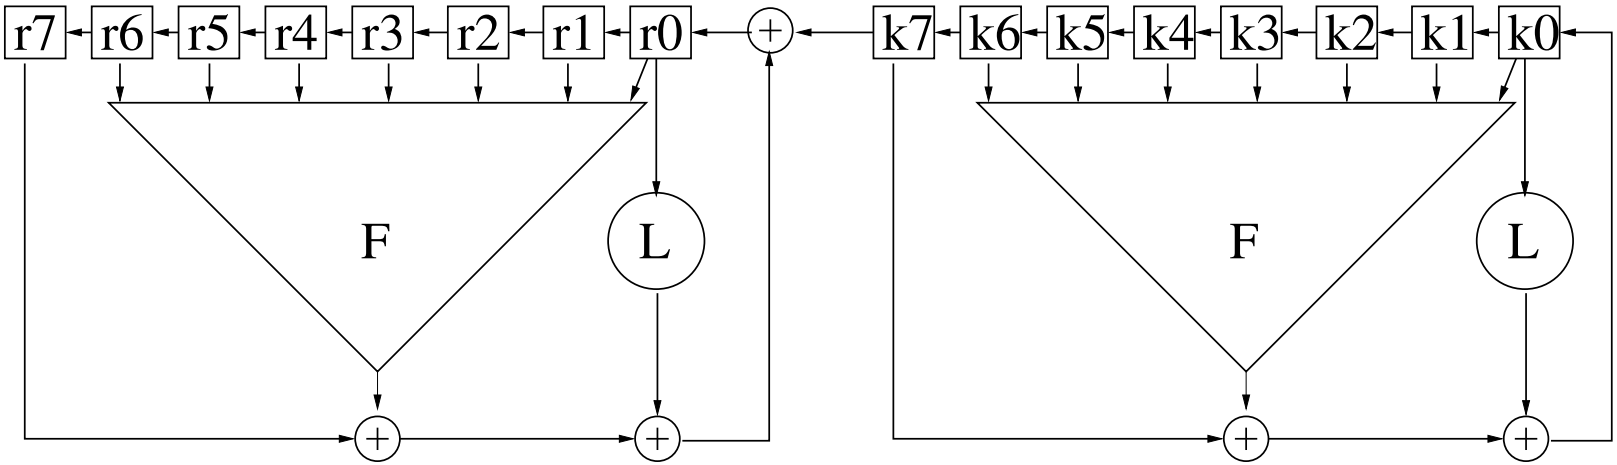

- Compression function is Davies-Meyer
- Boxes all represent 32-bit or 64-bit words
- *F* is non-linear, applied bit-wise
- *L* is linear, applied to a 32-bit or 64-bit word
- Can run entirely within register file on Core 2

F Function

$$F(a,b,c,d,e,f,g) =$$

$$abcdefg + abcdef + abcefg + acdefg + abceg + abdef + abdeg + abefg + acdef + acdfg + acefg + adefg + bcdfg + bdefg + cdefg + abcf + abcg + abdg + acdf + adef + adeg + adfg + bcde + bceg + bdeg + cdef + abc + abe + abf + abg + acg + adf + adg + aef + aeg + bcf + bcg + bde + bdf + beg + bfg + cde + cdf + def + deg + dfg + ad + ae + bc + bd + cd + ce + df + dg + ef + fg + a + b + c + f + 1$$

- Makes single lane register sequence De Bruijn
- Has maximal linear complexity
- Optimized against differential and linear attacks

L Function

L functions implemented via LFSR in Galois Configuration

- The 32-bit *L* function is stepped 32 times, the 64-bit *L* function is stepped 64 times
- The generating polynomials are primitive
- Every bit of input used in approx. 50% of output
- Every bit of output depends on approx 50% of input

Differential and Linear Analysis

- Extremely simple design facilitates analysis
- Differential and linear analysis available in specification
- Block cipher becomes resistant to standard differential and linear attacks in 24 steps
- Default number of steps in block cipher is 32

Algebraic Analysis

Observations of non-randomness in the ESSENCE compression function, by Nicky Mouha, Søren S. Thomsen, Meltem Sönmez Turan

- Found fixed point (1-cycle) for block cipher shift register leading to free-start pre-image for the all zero output in compression function
- Proved that there are no 2, 3, or 4-cycles other than the repeated 1-cycle.
- Described slid pairs for input/output
- Currently no clear way to extend these results to the full hash function

Serial Performance

Serial performance on small-word architectures is slower than SHA-2. Why?

- ESSENCE does not use arithmetic mod 2^{32} or 2^{64}
- ESSENCE does not use S-boxes
- Deliberate serial performance sacrifice for simplicity
- ESSENCE has overhead to support vector registers and parallel implementation

But...

Parallel Performance on dual core x86_64 CPUs

Platform	Implementation	Hash Size	cycles/byte
Vista 64 Core 2 Dual Core	Serial C-only	224	63.7
		256	63.6
		384	64.2
		512	64.2
	OpenMP and Assembly	224	19.7
		256	19.5
		384	23.5
		512	23.5

Parallel Performance on quad core x86_64 CPUs

Platform	Implementation	Hash Size	cycles/byte
Linux 64 Xeon Quad Core	OpenMP and Assembly	224	10.3
		256	9.9
		384	12.1
		512	12.1

Parallel Performance on eight core x86_64 CPUs

Platform	Implementation	Hash Size	cycles/byte
Mac OS X Xeon 8 Core	OpenMP and Assembly	224	4.99
		256	4.99
		384	5.85
		512	5.85

Timing Methods

- These are sustained through-put measurements using large messages ($>$16MB) which do not fit in cache
- Timings include all setup time, even time spent by OS to create and manage threads
- Timings include all memory access
- For consistency, all timing results are in terms of Time Stamp Counter.

Conclusion: Performance Bound by I/O

- Performance measurements on a large SMP machine (24 cores, 128 GBytes of RAM) indicate that the real limit to performance is bus speed
- Parallel hardware existing today demonstrates that the performance limit is I/O not compute power

ESSENCE is Simple and Parallel

- Simplicity
  - Uses simple design to facility analysis
  - Uses old, well-studied, constructions
  - Published external analysis already available
- Parallelism
  - Use bit-wise non-linear function to take advantage of large registers in modern CPUs
  - Use tree structure to take advantage of multiple cores
  - Current parallel implementations are limited by memory and bus speed, not computation speed

ESSENCE in Second Round

If ESSENCE is still unbroken by the next SHA-3 conference, it should be considered for round 2 of the contest because it is simple enough to facilitate analysis and parallelizable enough to get high performance.