A High-Performance Area-Efficient AES Cipher on a Many-Core Platform

Bin Liu and Bevan M. Baas

VLSI Computation Lab
ECE Department
University of California, Davis

November 9th, 2011
Asilomar Conference on Signals, Systems and Computers

Outline

- *Advanced Encryption Standard*
- Targeted Fine-Grained Many-Core Platform
- Implementations of AES Cipher
- Comparison with Related Work

Advanced Encryption Standard

- AES is a symmetric block encryption algorithm
- Plaintext: 128 bits, a 4-by-4 byte array
- Four basic operations in the main loop
  - SubBytes
  - ShiftRows
  - MixColumns
  - AddRoundKey

Length of *round key* (bits)	Number of Rounds (N_r)
128	10
192	12
256	14

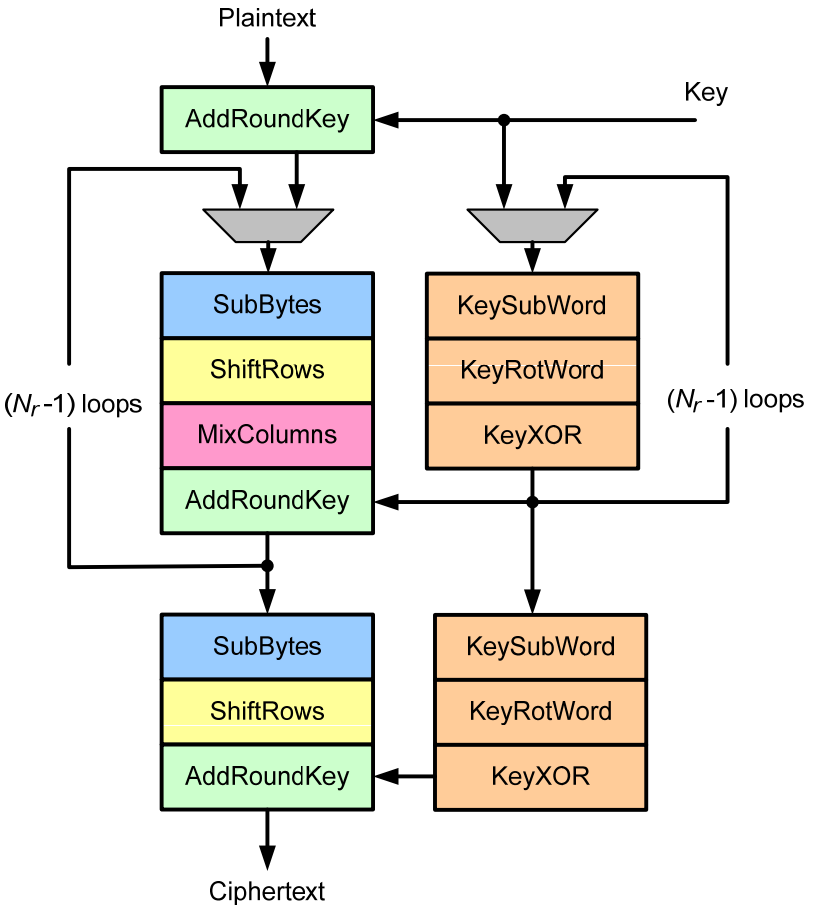

AES Basic Operations

SubBytes: byte substitution from a look up table

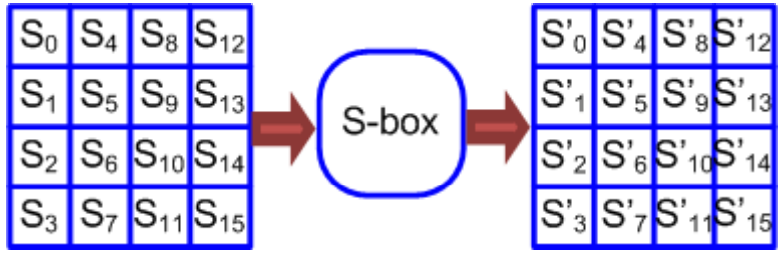

MixColumns: each column multiplies a fixed polynominal over GF(2^8)

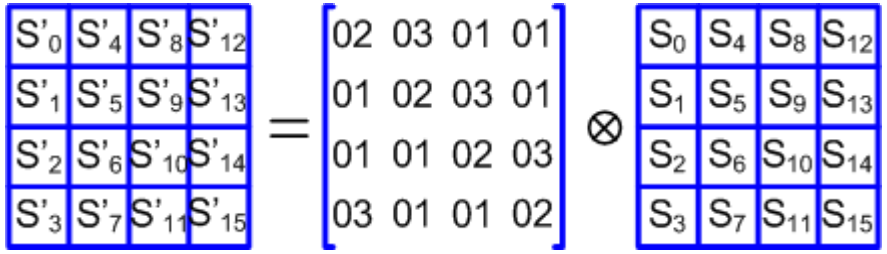

ShiftRows: cyclically shift by one, two and three bytes in the 2nd, 3rd and 4th row

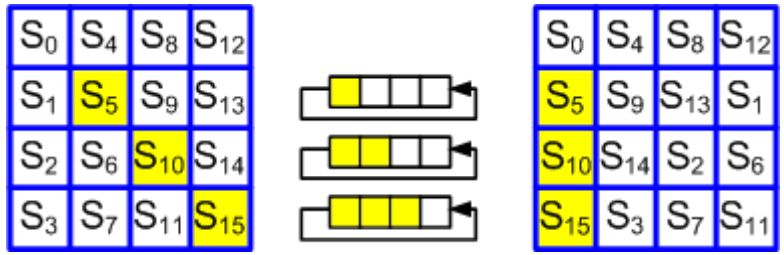

AddRoundKey: round key is added to input using a bitwise XOR operation

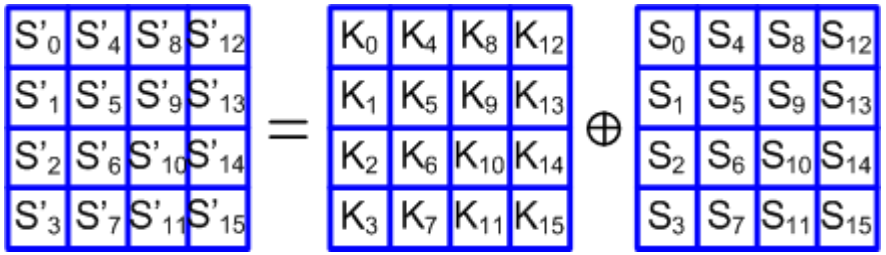

AES Key Expansion

KeySubWord: byte substitution from a look up table for a four-byte word

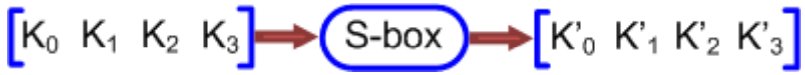

KeyRotWord: left cyclic shift one byte

KeyXOR: every word *w*[*i*] is equal to the bitwise XOR of the previous word, *w*[*i*-1], and the word *Nk* position earlier, *w*[*i*-*Nk*].

Note: *Nk* equals 4, 6 or 8 for the key length of 128, 192 or 256 bits

Outline

- Advanced Encryption Standard
- *Targeted Fine-Grained Many-Core Platform*
- Implementations of AES Cipher
- Comparison with Related Work

Targeted Fine-Grained Many-Core Platform

- **164 homogeneous fine-grained cores**
  - In-order 6-stage pipeline
  - no specialized instructions
  - 128 x 32-bit instruction memory
  - 128 x16-bit data memory
  - Max. frequency 1.2GHz @ 1.3V
  - 0.17 mm^2 in 65nm CMOS
- **On-chip reconfigurable 2D-mesh network**
  - Nearby & long-distance communication

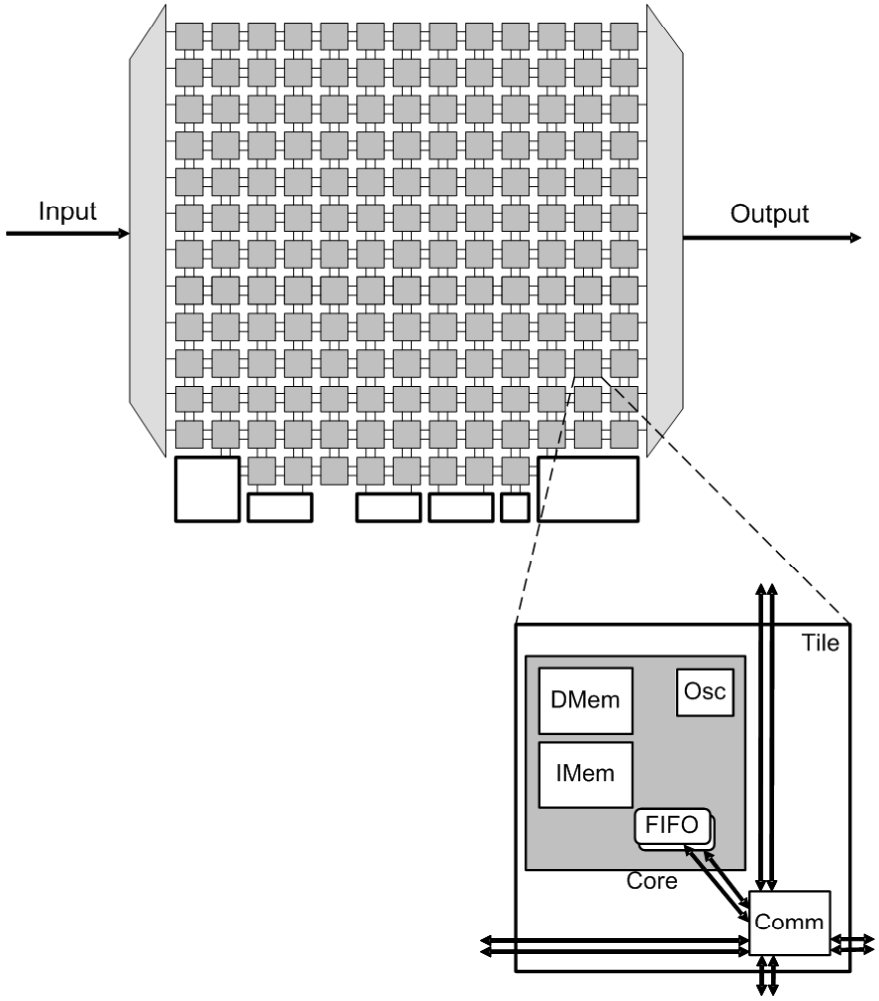

D. Truong et.al, *JSSC*, 2009

Outline

- Advanced Encryption Standard
- Targeted Fine-Grained Many-Core Platform
- *Implementations of AES Cipher*
- Comparison with Related Work

Preliminary Design of AES Cipher

- **(N_r-1) times loop-unrolling is applied to both the main AES algorithm and the key expansion process**
  - Key length = 128 bits, N_r = 10
- **Throughput is 266 clock cycles per block, equaling 16.625 clock cycles per byte**
  - Determined by the *MixColumns* cores.
- **70 cores are used for this implementation**

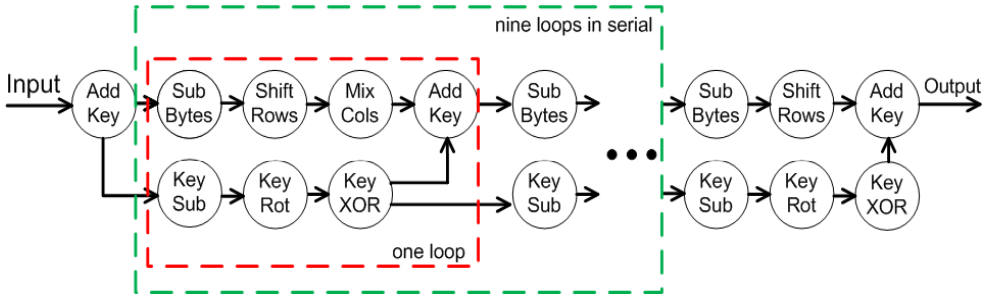

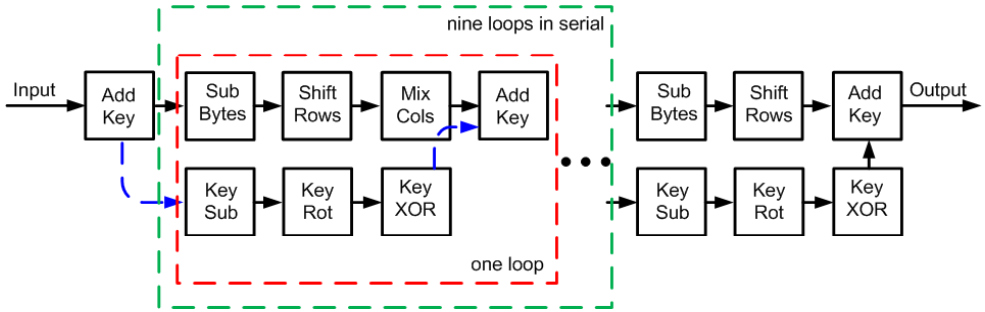

Optimization I: Increasing Throughput

- **Cores running *MixColumns* workloads are 2x slower than other cores, which are the bottlenecks of the design.**
- **Parallelize each *MixColumns* core into two *MixCol-8* cores**
  - Each *MixCol-8* processes two columns (8 bytes) instead of four columns
- **Throughput is increased by 43% (152 cycles per block)**
  - 10 more cores are required

Processor Name	Execution Time for Processing One 128-bit Data Block (Clock Cycles)
SubBytes	**132**
ShiftRows	**38**
MixColumns	**266**
AddRoundKey	**22**
KeySubWord	**56**
KeyRotWord	**26**
KeyXOR	**56**

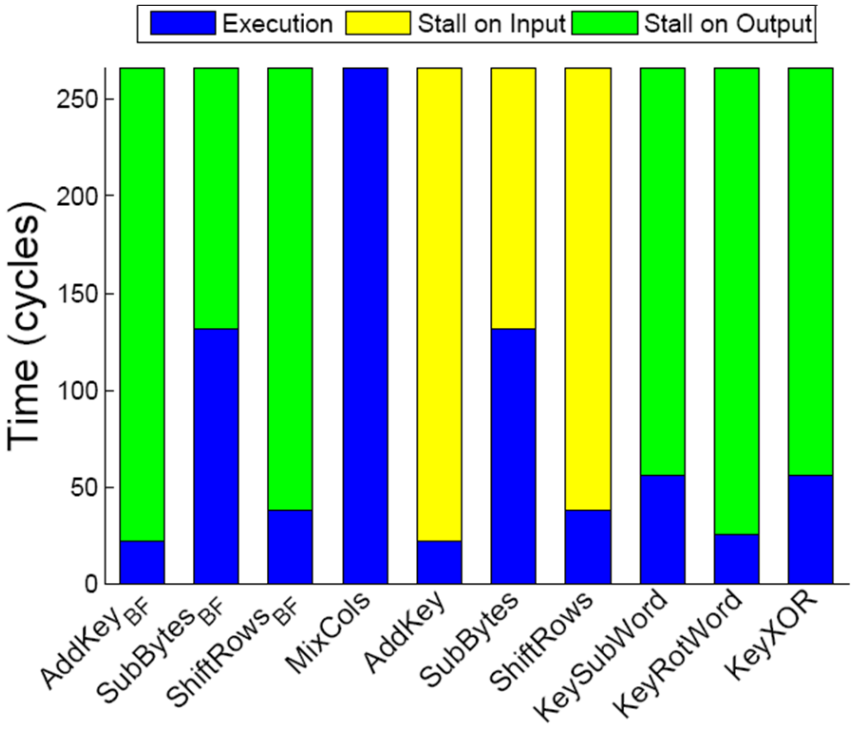

Optimization II: Reducing Cores

- **Before optimization:**
  - ~22% average IMem usage
  - ~43% average DMem usage
- **Combine the neighboring *SubBytes* and *ShiftRows* core into one *SubShift* core**
  - T_{EXE} =148 cycles per data block
  - 80% IMem usage and 100% DMem usage

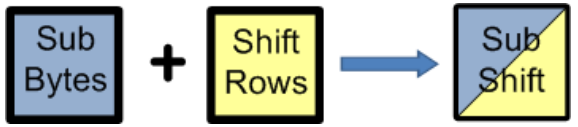

- **Combine the neighboring *KeyRotWord* and *KeyXOR* cores into one *KeyScheduling* core**
  - T_{EXE} =60 cycles per data block
  - 24% IMem usage and 28% DMem usage

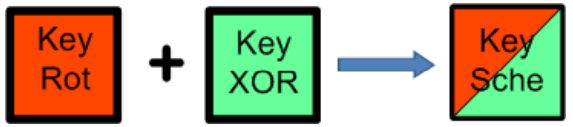

- **Further core merging would reduce the throughput of the design or exceed the memory limitations**

Optimized Design of AES Cipher

- **The optimized cipher achieves a 43% higher throughput**

  (9.5 cycles per data block)
- **The optimized design requires 16% fewer cores (59 cores)**
- **The execution activity of processors for the optimized cipher is more balanced compared with the preliminary design.**

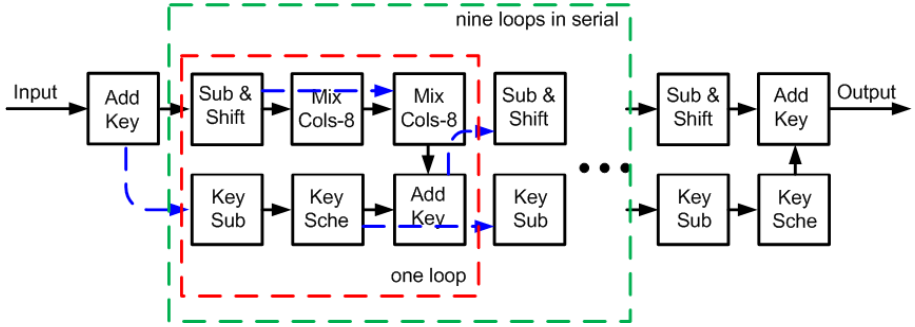

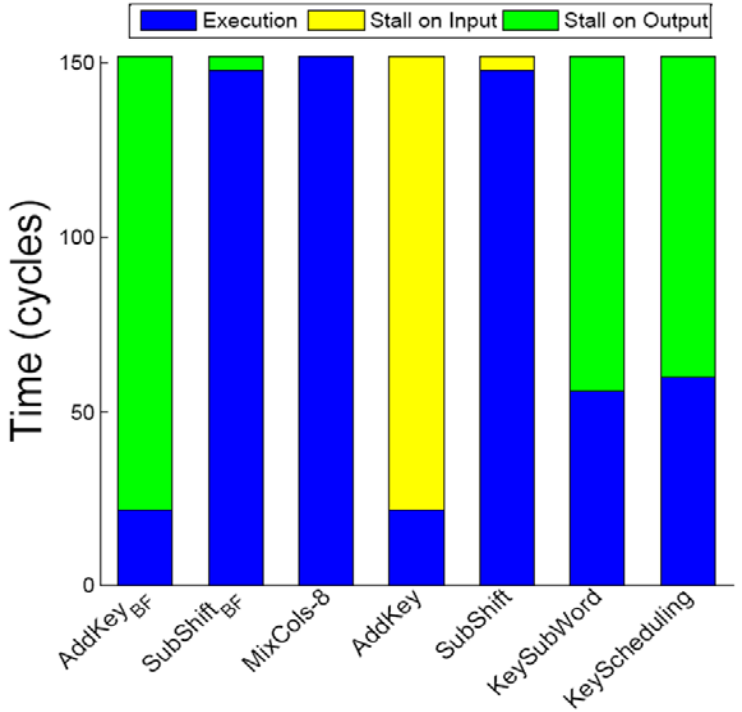

Outline

- Advanced Encryption Standard
- Targeted Fine-Grained Many-Core Platform
- Implementations of AES Cipher
- *Comparison with Related Work*

Comparison with Related Work

Platform	Method	Tech. (nm)	Area (mm^2)	Max Freq. (MHz)	Throughput (cycles/byte)	Scaled Throughput (Mbps)	Scaled Area (mm^2)	Scaled Throughput/Area (Mbps/mm^2)
Pentium 4 561	Bitslice	90	112	3600	16	2492	58.42	42.66
Athlon 64 3500	Bitslice	90	193	2200	10.6	2299	101	22.76
Core 2 Duo E6400	Bitslice	65	111	2130	9.19	1854	111	16.70
Core 2 Quad Q6600 (one core)	Bitslice + SSSE3	65	286/2 = 143	2400	9.32	2060	143	14.41
Core 2 Quad Q9550 (one core)	Bitslice + SSSE3	45	214/4 = 53.5	2830	7.59	2065	112	18.44
Core i7 920 (one core)	Bitslice + SSSE3	45	263/4 = 65.75	2668	6.92	2135	133	16.05
TI C6201		180	NA	200	14.25	311	NA	NA
GeForce 8800 GTX	T-Box	90	484	575	NA	11500	252	45.63
This Work AsAP		65	6.63	1210	9.5	1019	6.63	153.70

- Compared to CPUs, our design achieves 3.6–10.7x higher throughput per chip area
- Compared to DSP, our design achieves 1.5x higher throughput
- Compared to GPU, our design achieves 3.4x higher throughput per chip area

Acknowledgments

- NSF Grant 0430090, 0903549; and CAREER Award 0546907
- SRC GRC Grant 1598, 1971; and CSR Grant 1659
- UC Micro
- ST Microelectronics
- Intel
- Intellasys
- C2S2 Focus Center, one of six reserch centers funded under the Focus Center Research Program (FCRP), a Semiconductor Research Corporation entity.