December 5th-6th, 2018 | San Jose, CA

Applying SIMD Optimizations to DPCLS

Harry van Haaren
Intel

OVERVIEW

- Introduction
  - Vectorization
  - Miniflow
  - Datapath Classifier
- AVX-512 DPCLS
  - Graphical Walkthrough
- Future Work

VECTORIZATION

- **More work per CPU cycle**
- SIMD : **S**ingle **I**nstruction **M**ultiple **D**ata
- AVX-512
  - 512-bit wide registers
  - “Lanes” concept
    - 512 bits split
    - 8 lanes x 64 bits
    - A lane is a packet

MINIFLOW

- **What is a miniflow?**
  - Structure to represent metadata
  - Packets and Subtables have miniflows

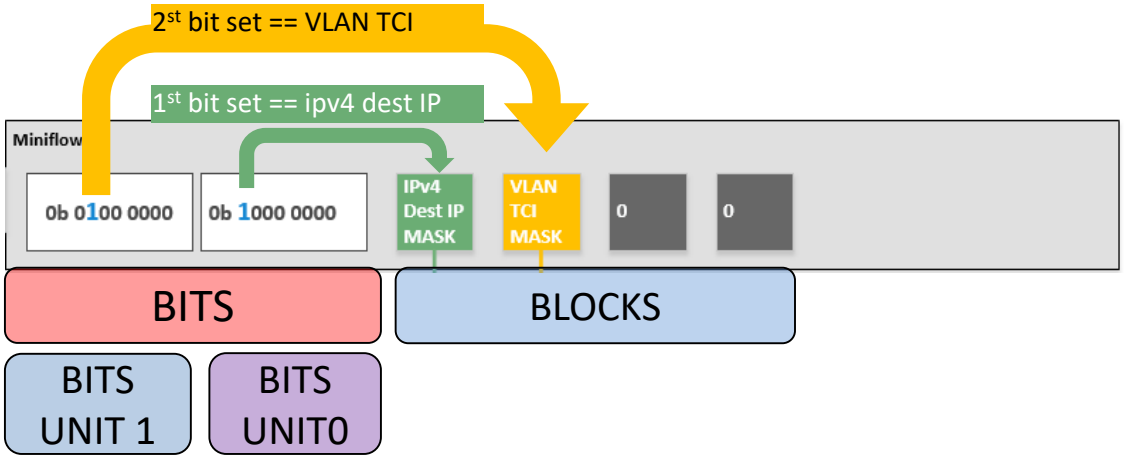

DATA PATH CLASSIFER

- Iterate Subtable BITS
  - Find Corresponding Packet Block, Mask with Subtable

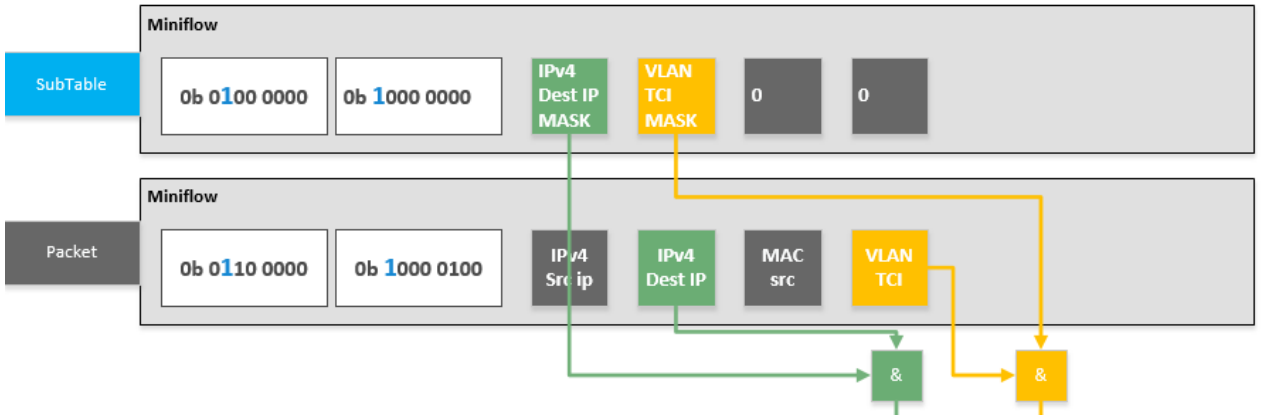

CURRENT CODE

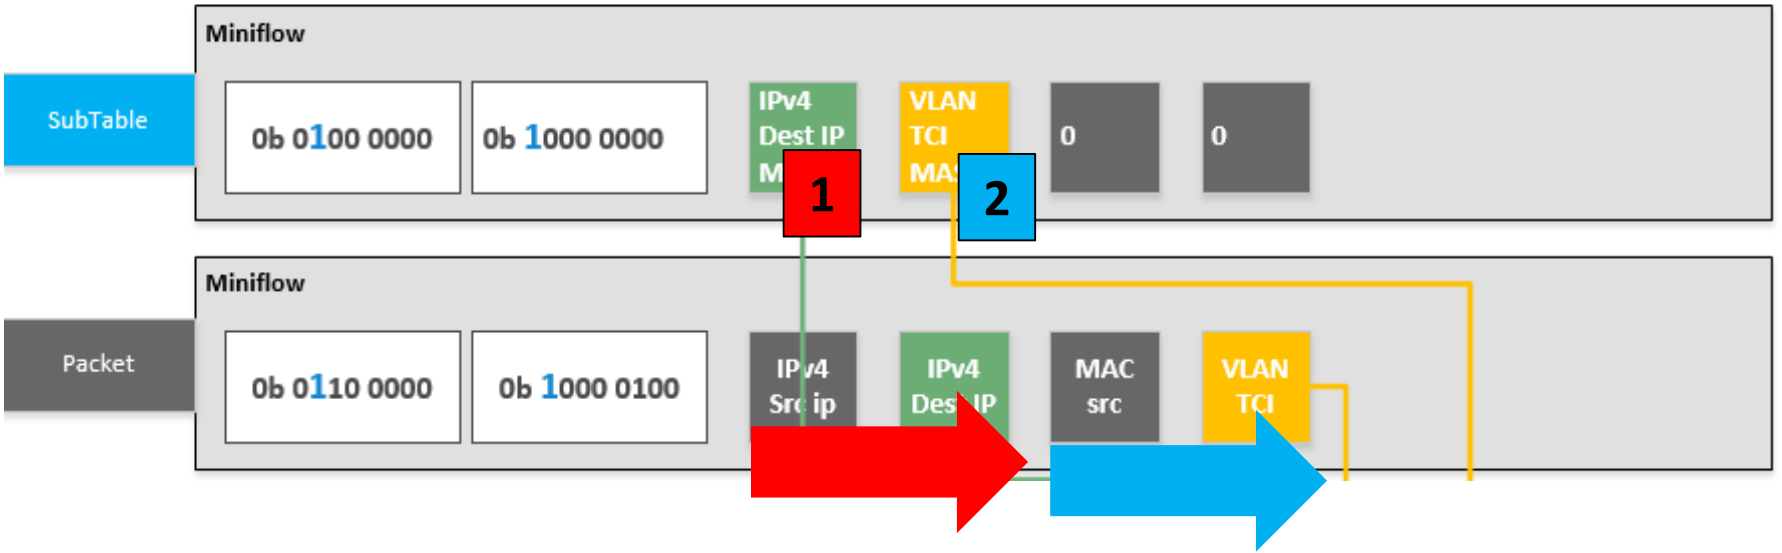

BITS-TO-BLOCKS

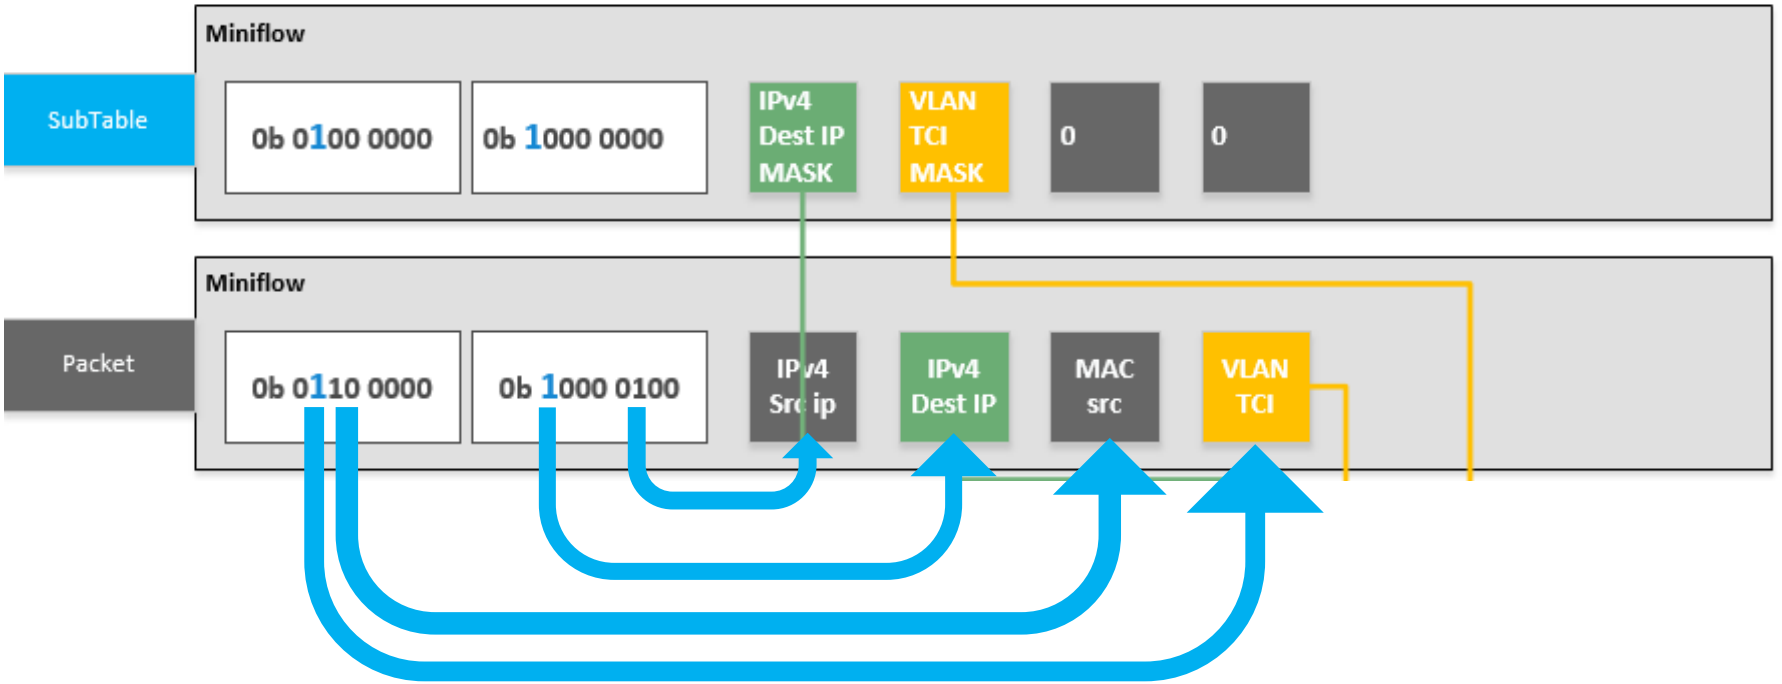

POPCOUNT()

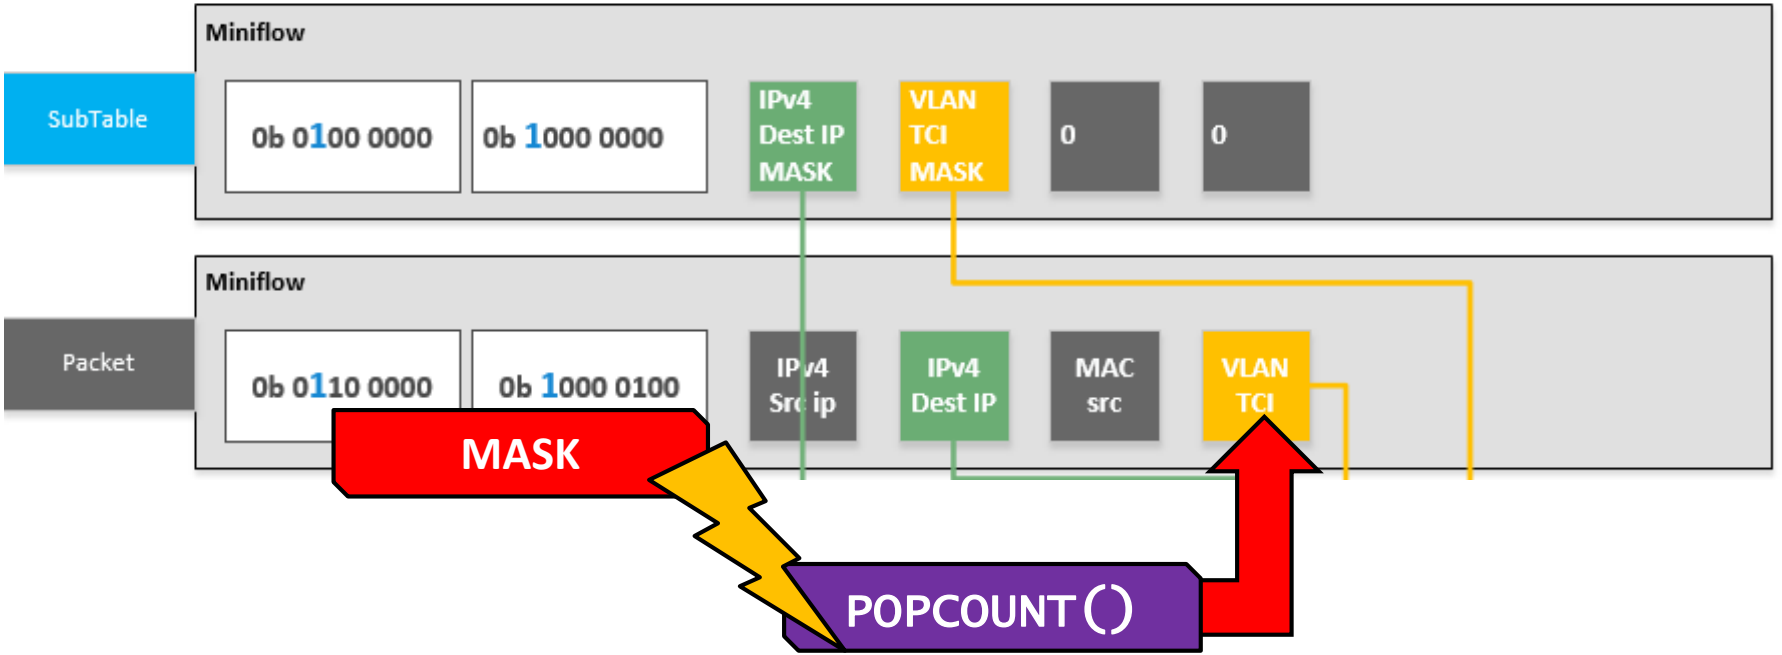
Miniflow
SubTable
0b 0100 0000
0b 1000 0000
IPv4 Dest IP MASK
VLAN TCI MASK
0
0
Miniflow
Packet
0b 0110 0000
0b 1000 0100
IPv4 Src ip
IPv4 Dest IP
MAC src
VLAN TCI
MASK
POPCOUNT ()

POPCOUNT() + AVX-512

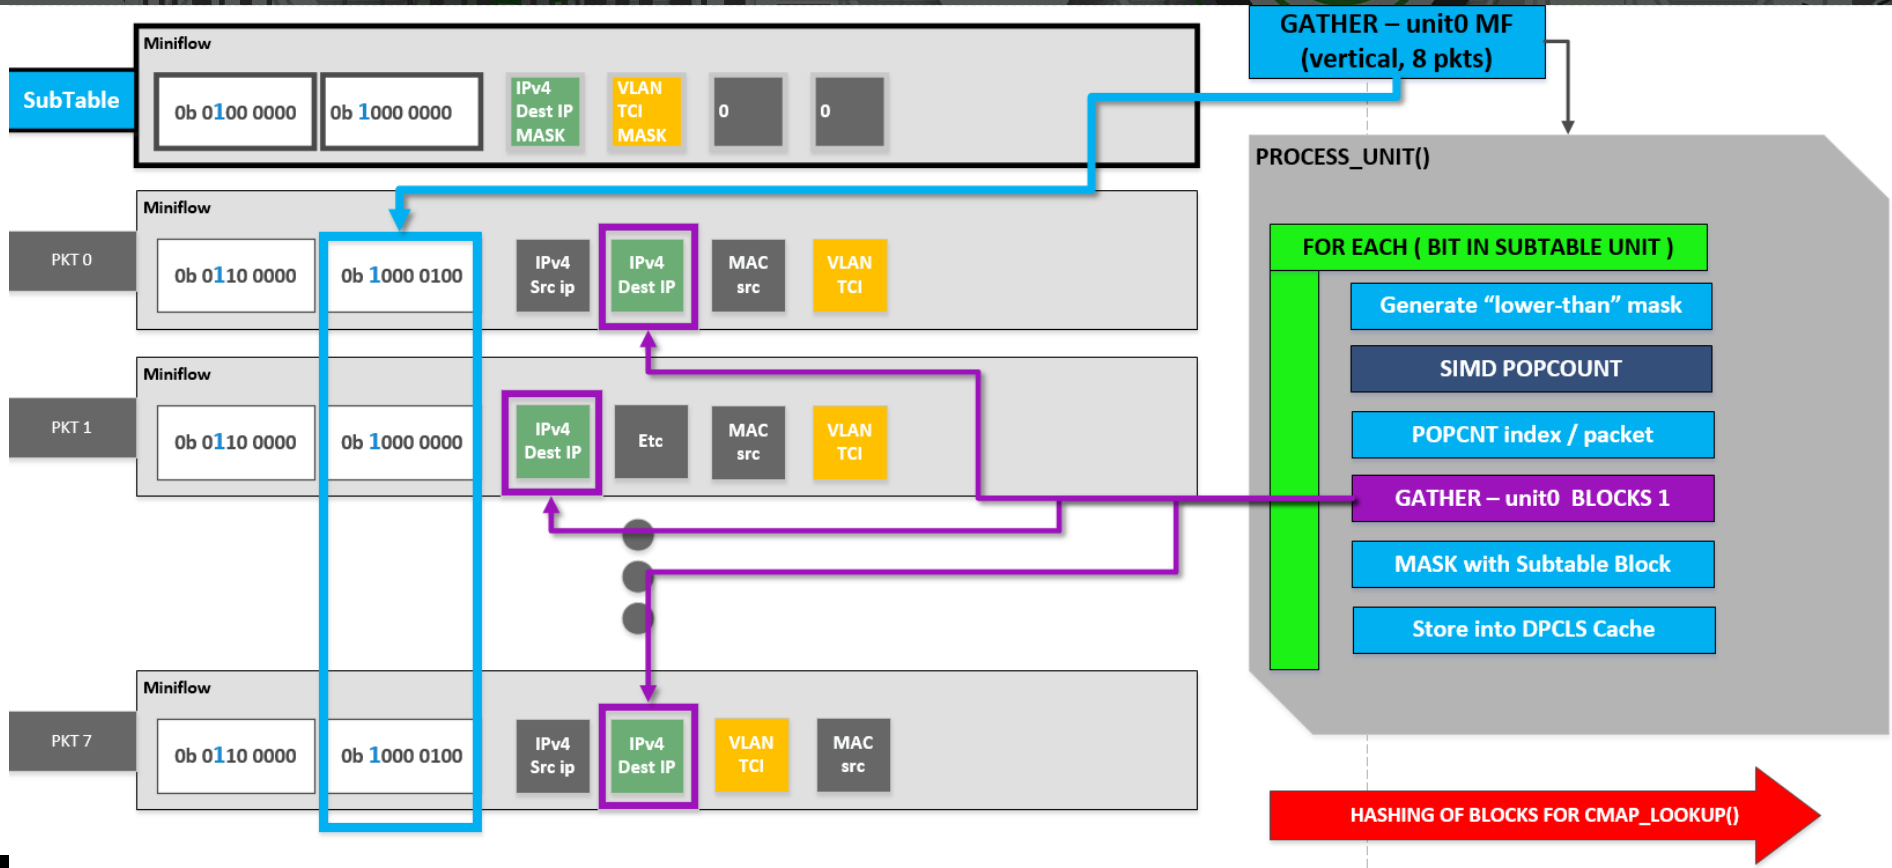

FUTURE WORK

- **Per-Subtable Function pointers**
  - Flexible selection of Optimized Functions
  - Using Context of Subtable Miniflow
    - Benefits scalar code too!

- **Optimize Cache Utilization**
  - EMC – Fast / Cache trashes > X flows
  - SMC – Fast / More Efficient cache usage
  - DPCLS – Slow(er) / Consistent w flow-count

CALL TO ACTION

- **Community to Drive SW Optimizations**
  - **Require well defined use-cases to optimize**
  - **Request to community to contribute use-cases**

- **Tuning Knobs**
  - **OVS community : “Less is More”**
  - **Suggestion : “Use CPU to go faster” button?**
    - **Enables AVX in DPCLS**
    - **Other CPU specific opts?**

? Questions ?

Harry van Haaren
harry.van.haaren@intel.com

Patchwork series 77134 : “dpcls subtable miniflow optimizations”

MINIFLOW

MINIFLOW LEDGEND:

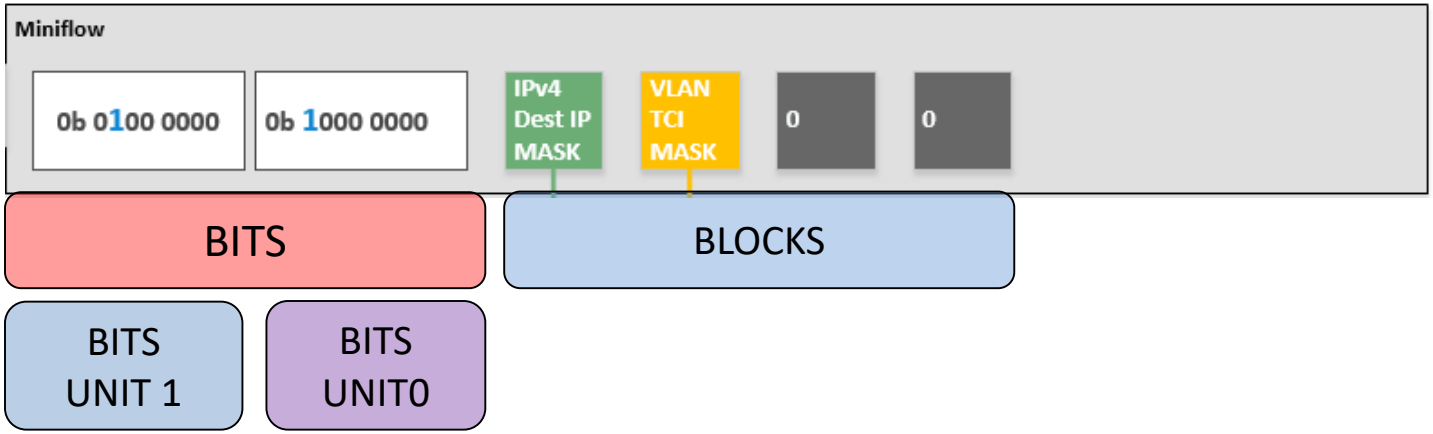

Notices & Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.

No computer system can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/benchmarks .

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit http://www.intel.com/benchmarks .

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.