SAND2018-5969C

Sandia National Laboratories

Convolutional Back Projection on the S1 Reduced Precision Processor

Michael Holzrichter and Randy Spaulding – Sandia National Laboratories

Concept

We conjecture that if less precise computations can be tolerated, then more parallelism can be obtained from a given number of transistors. This means increased processing performance without increasing size, weight and power (SWaP). We explored this concept in the context of SAR image formation using the S1 reduced precision processor from Singular Computing. We assessed the S1's potential by forming SAR images using the Convolutional Back Projection algorithm.

The S1 Reduced Precision Processor

- DARPA-funded, first-generation prototype reduced precision processor
- 2112 simple computing elements (called "approximate processing elements" or APEs) per chip
- Each APE has 8 2-byte general registers and 512 2-byte words of memory (approximately 2MB per chip total).
- 200MHz clock, one instruction executed per clock cycle
- 2-byte word size

The S1 Programmers Model

- SIMD parallelism: all APEs execute the same instruction stream in lock-step.
- Datatypes:
    1. Boolean
    2. 2-byte integers
    3. 14-bit "approx" (similar to half-precision floats)
- Precision of "approx" good to 1 part in 64
- Core machine instructions:
    - Integer: add, subtract, bit shift left, bit shift right
        - (no integer multiply or divide)
    - Approx: add, subtract, multiply, divide, square root
- Programmer:
    - Manually manages memory and register usage and
    - Programs the S1 at the machine instruction level (except for common constructs such as loops).

Convolutional Back Projection Results

- Images computed on S1 and Intel (double precision) are indistinguishable to human eye. (See Figures 1 and 2.)
- The difference image (Figure 3) shows
    - The horizontal streaks off the corner reflectors are from multiplicative noise. These streaks are a result of the S1 and its reduced precision computations.
    - The vertical corduroy pattern is due to our computing phase using a linear interpolator; this would go away with a cubic interpolator.
    - Difference rarely exceeds 1 dBsm (Figure 4).
- The memory and I/O bandwidth required by Convolutional Back Projection greatly exceeds that provided by the S1.
- The 678ms it takes to form a 1k x 1k image on a Tegra X1 is less than the 1434ms we used to fetch an already formed image from a 50MHz S1 development system. (Table 1)

Image Size	256 x 256	512 x 512	1024 x 1024
Initialize	867	876	910
Process phase histories:			
• Evaluate control points	3028	11268	43710
• Compute interpolation parameters	2883	11495	45899
• Distribute interpolation parameters	2458	9544	37853
• Perform interpolation	2200	2251	2401
• Distribute phase histories	10528	20680	41228
Form image	748	746	750
Fetch Image	90	358	1434
Total	**22829**	**57246**	**174213**

Table 1: Breakdown of time (in ms) of CBP on S1 (running at 50 MHz)

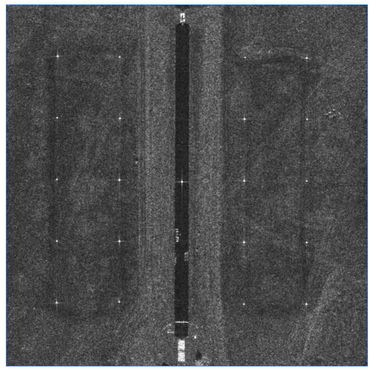

Figure 1: Image formed an S1 system with 16 S1 chips.

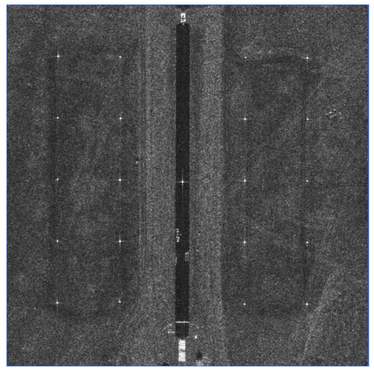

Figure 2: Image formed using double precision on Intel processor

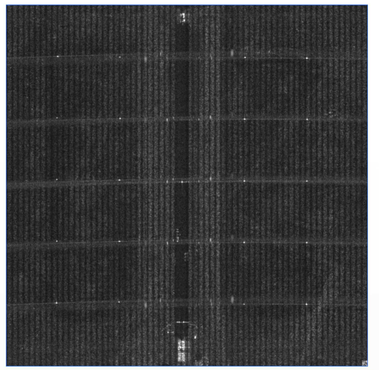

Figure 3: Difference between S1 and Intel images (scaled by a factor of 100)

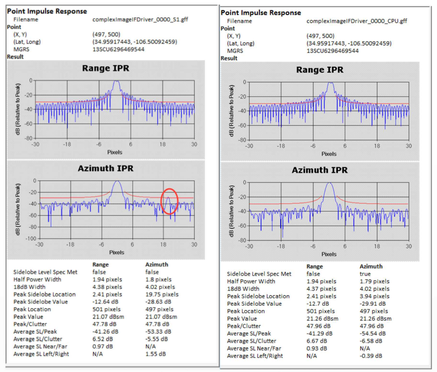

Figure 4: Impulse response of center corner reflector of S1 (left) and Intel (right) images

Conclusions

- It is possible to form reasonable quality SAR images on a reduced-precision processor.
- The Convolutional Back Projection algorithm requires more memory and I/O bandwidth than provided by the S1.
- Since the S1 is a proof-of-concept prototype, we expect its successors will have much more memory and I/O bandwidth.

Acknowledgements

- Joint Munitions Program Technical Coordinating Group XIII: Casey Clark (NSWC-Dahlgren), Tom Lewis (AFRL-Eglin), Martin Heimbeck (AMRDEC), and Brian Smith (AMRDEC)
- Joe Bates of Singular Computing
- Dale Dubbert, Derek West and Kurt Sorensen of Sandia Labs

U.S. DEPARTMENT OF ENERGY
National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology &Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
SAND No. _____

Sandia National Laboratories