NSF CRI CNS-0551504, ATLAS Support and Development

R. Clint Whaley

whaley@cs.utsa.edu

www.cs.utsa.edu/~whaley

University of Texas at San Antonio
Department of Computer Science

August 23, 2007

NSF CRI CNS-0551504 Details

Community Resource Development

- $100,000 total budget, all for student support
- Just finished first of two years
- Presently supporting one PhD student, Tony Castaldo
- Not hardware acquisition: Community Resource Development:
  - ATLAS support and development

  → What is ATLAS? ⇒ next slide

What is ATLAS?

- Provides high performance linear algebra routines:
  - BLAS, some LAPACK
- Automatically adapts itself to differing architectures using empirical techniques

Why is ATLAS needed?

- Well-tuned linear algebra routine runs orders of magnitude faster than generic implementation
- Hand-tuning is architecture specific
- No such thing as enough compute speed for many scientific codes

ATLAS Usage and Progress

Usage/Tech Transfer

- Scientific simulation:
  - physics, chemistry, biology, astronomy, engineering, math
- Almost all supercomputers
- Many OSes include ATLAS:
  - OS X, most Linux & BSDs
- Most PSEs:
  - Maple, Mathematica, Matlab, Octave
- Multitude of software:
  - GSL, HPL, SciPy, R, some compilers, etc.

Award Progress

- 20 developer releases
- 1^{st} stable release in almost 4 years this summer
- More than 58,000 direct downloads during award
  - Most users get ATLAS via repackegers
- Error analysis research

$\longrightarrow$ next slide

New Research
Reducing Error in FP Computations

- Needed to guarantee stability in face of transformation
- Increasingly important as memory and performance expand
- Submitted to SISC, available as working note:
  → `http://www.cs.utsa.edu/research/tr/2007/CS-TR-2007-002.pdf`
- Initial work covers dot product, which accounts for most of the error of many algorithms:
  - Research undertaken by Tony Castaldo (PhD candidate),
    - Faculty: R. Clint Whaley, Anthony Chronopoulos
  - Produce tighter bound for forward error, improved notation,
  - Survey error prop of several of most important algorithms,
  - Present a new class of dp (**superblock**) which subsumes these:
    - Also produces new algorithms, allowing tuning of error and mem/perf tradeoffs
  - Statistical studies show large difference in practice.

Error Studies on Mixed Sign Data

Averaged over 10,000 random vectors

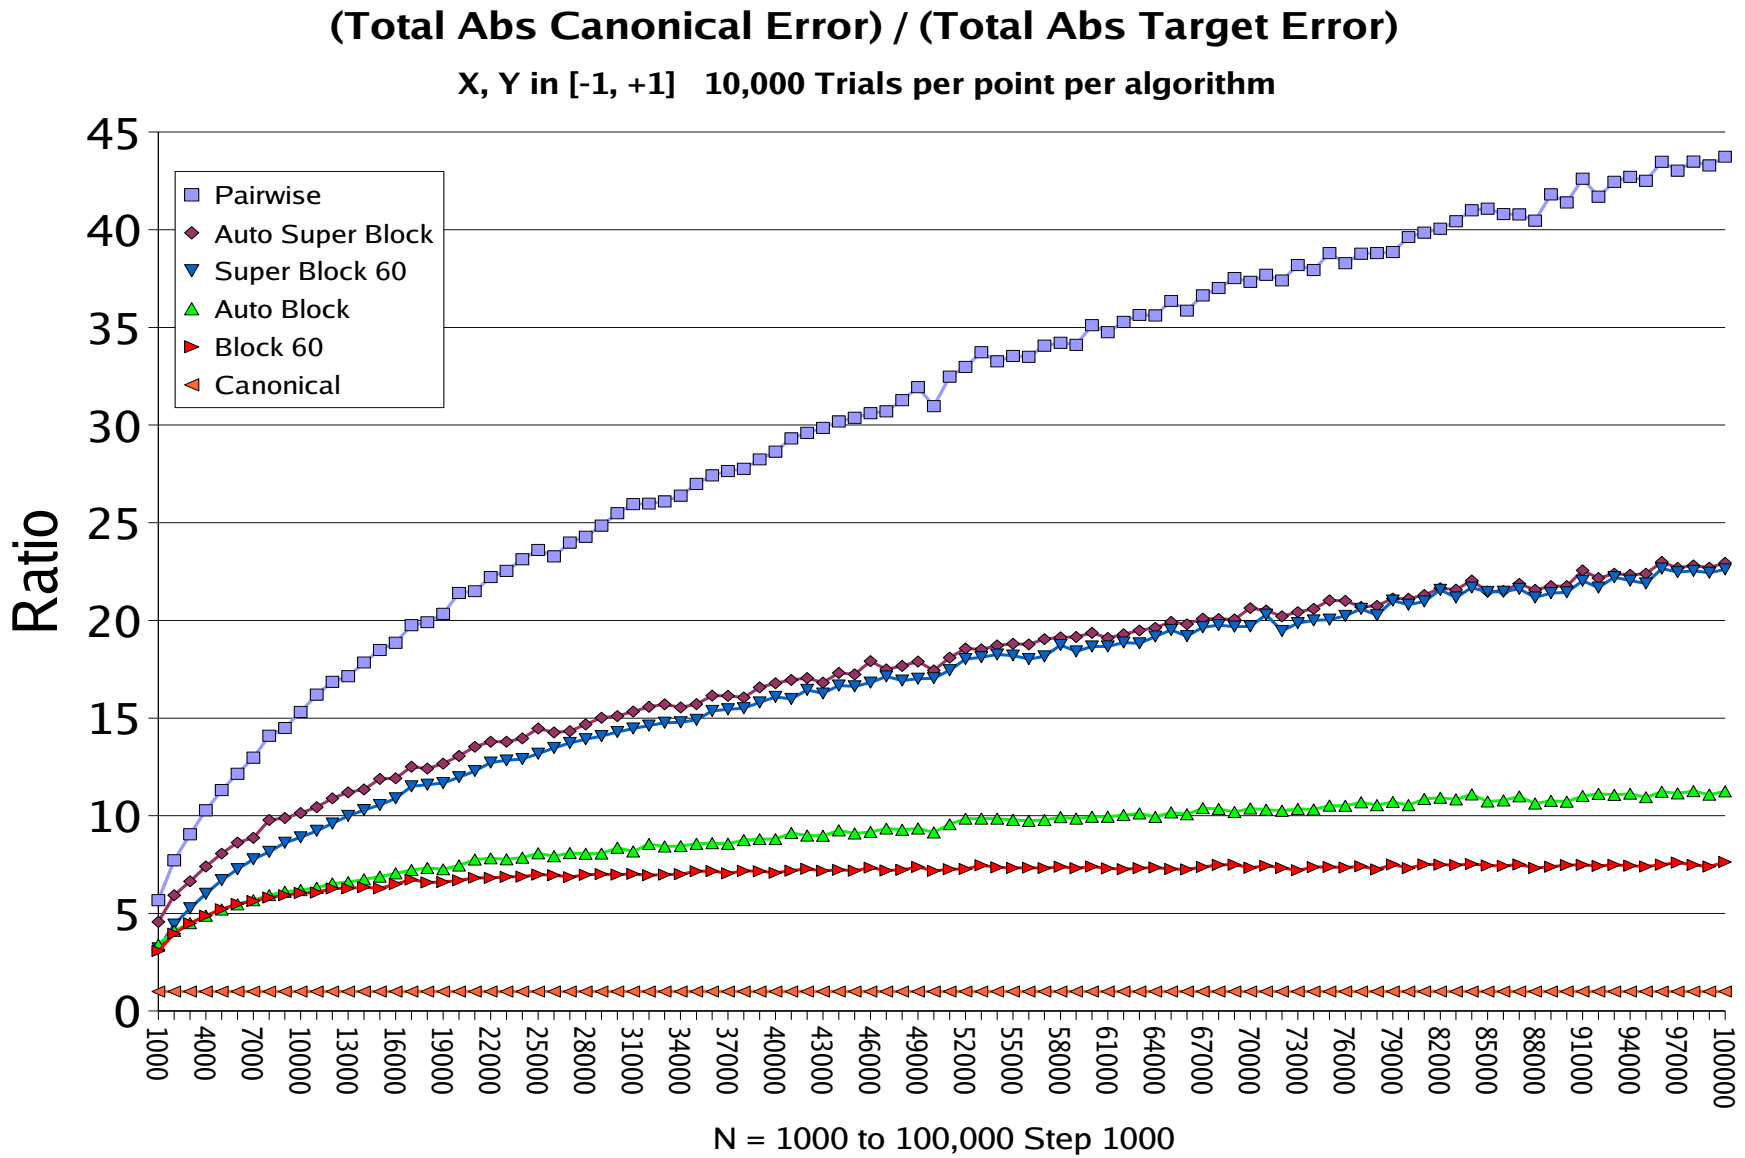

Error Studies on Same Sign Data

Averaged over 10,000 random vectors

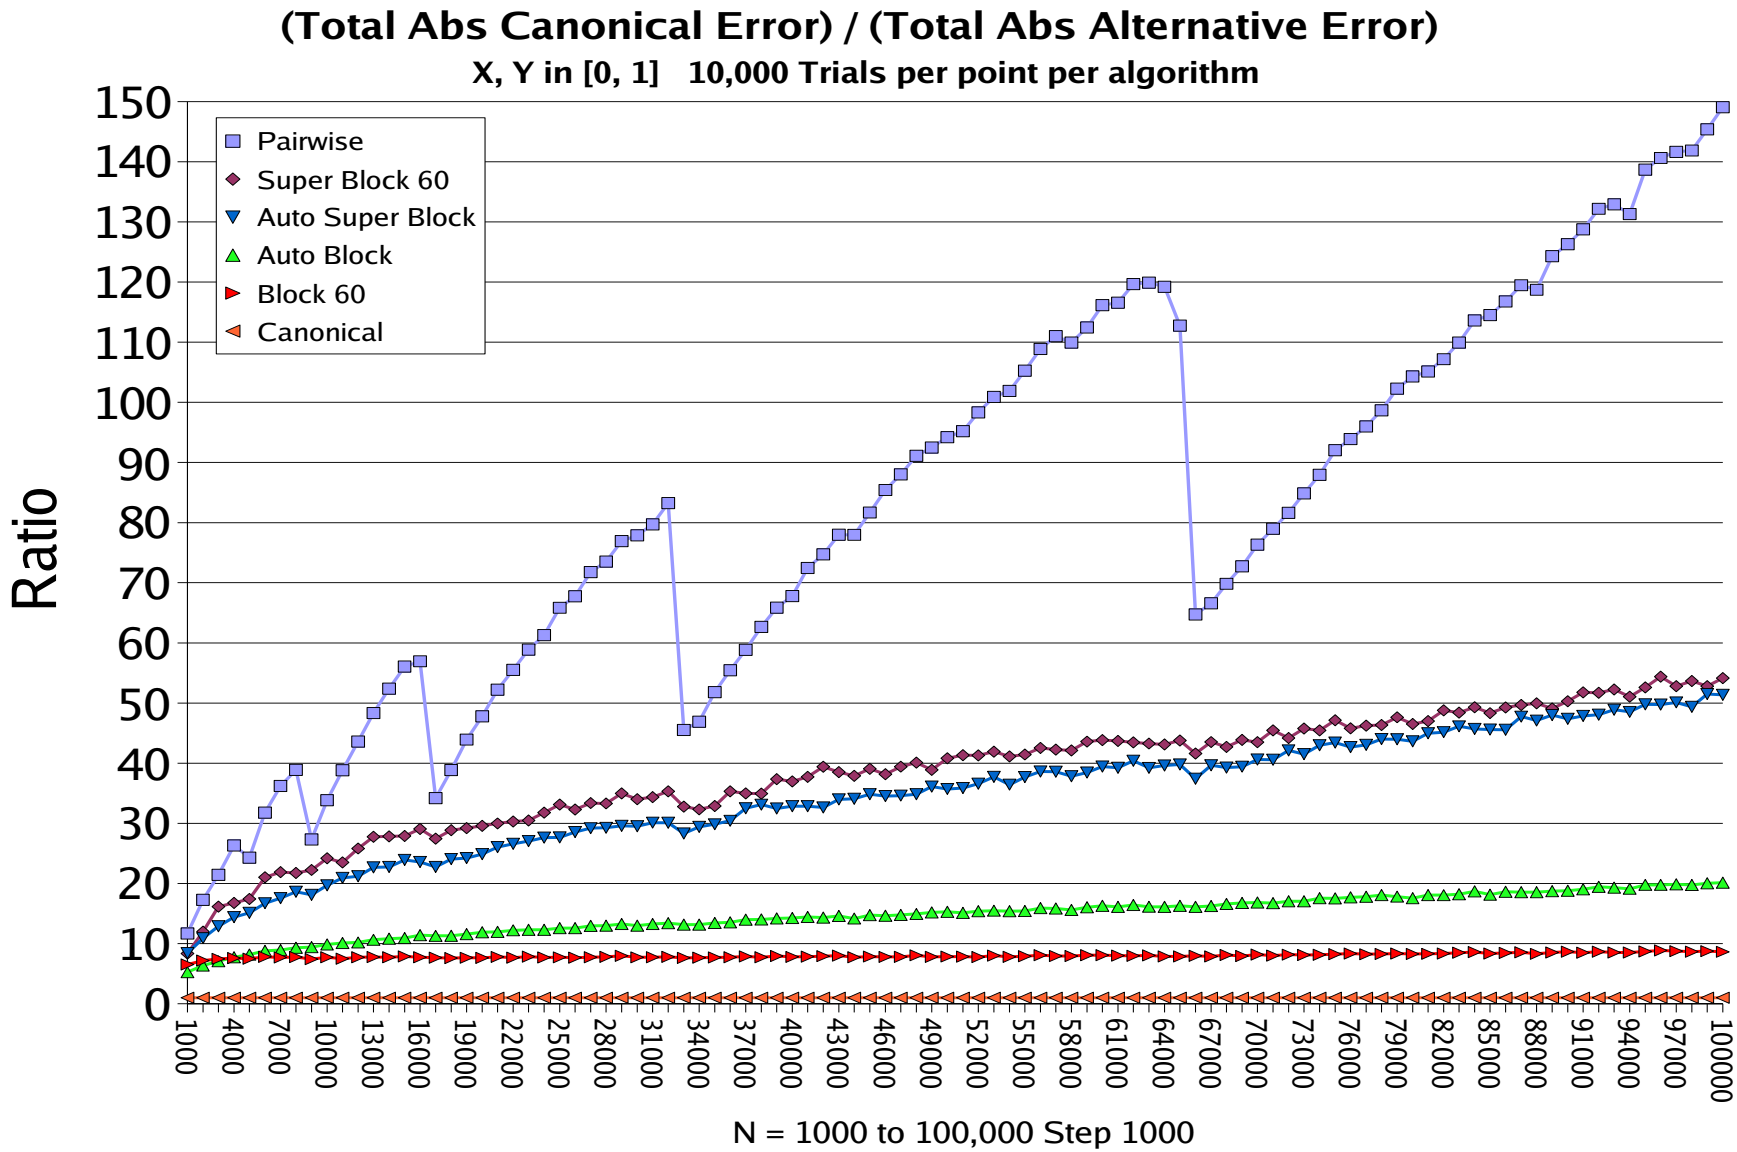

Further details

- Working note:

  `http://www.cs.utsa.edu/research/tr/2007/CS-TR-2007-002.pdf`
- ATLAS homepage:

  `http://math-atlas.sourceforge.net/`
- Author homepages:

  `http://www.cs.utsa.edu/~whaley/`

  `http://www.cs.utsa.edu/~atc/`

  `http://www.cs.utsa.edu/~castaldo/`

Absolute Error by Algorithm - Mixed Sign

Averaged over 10,000 random vectors

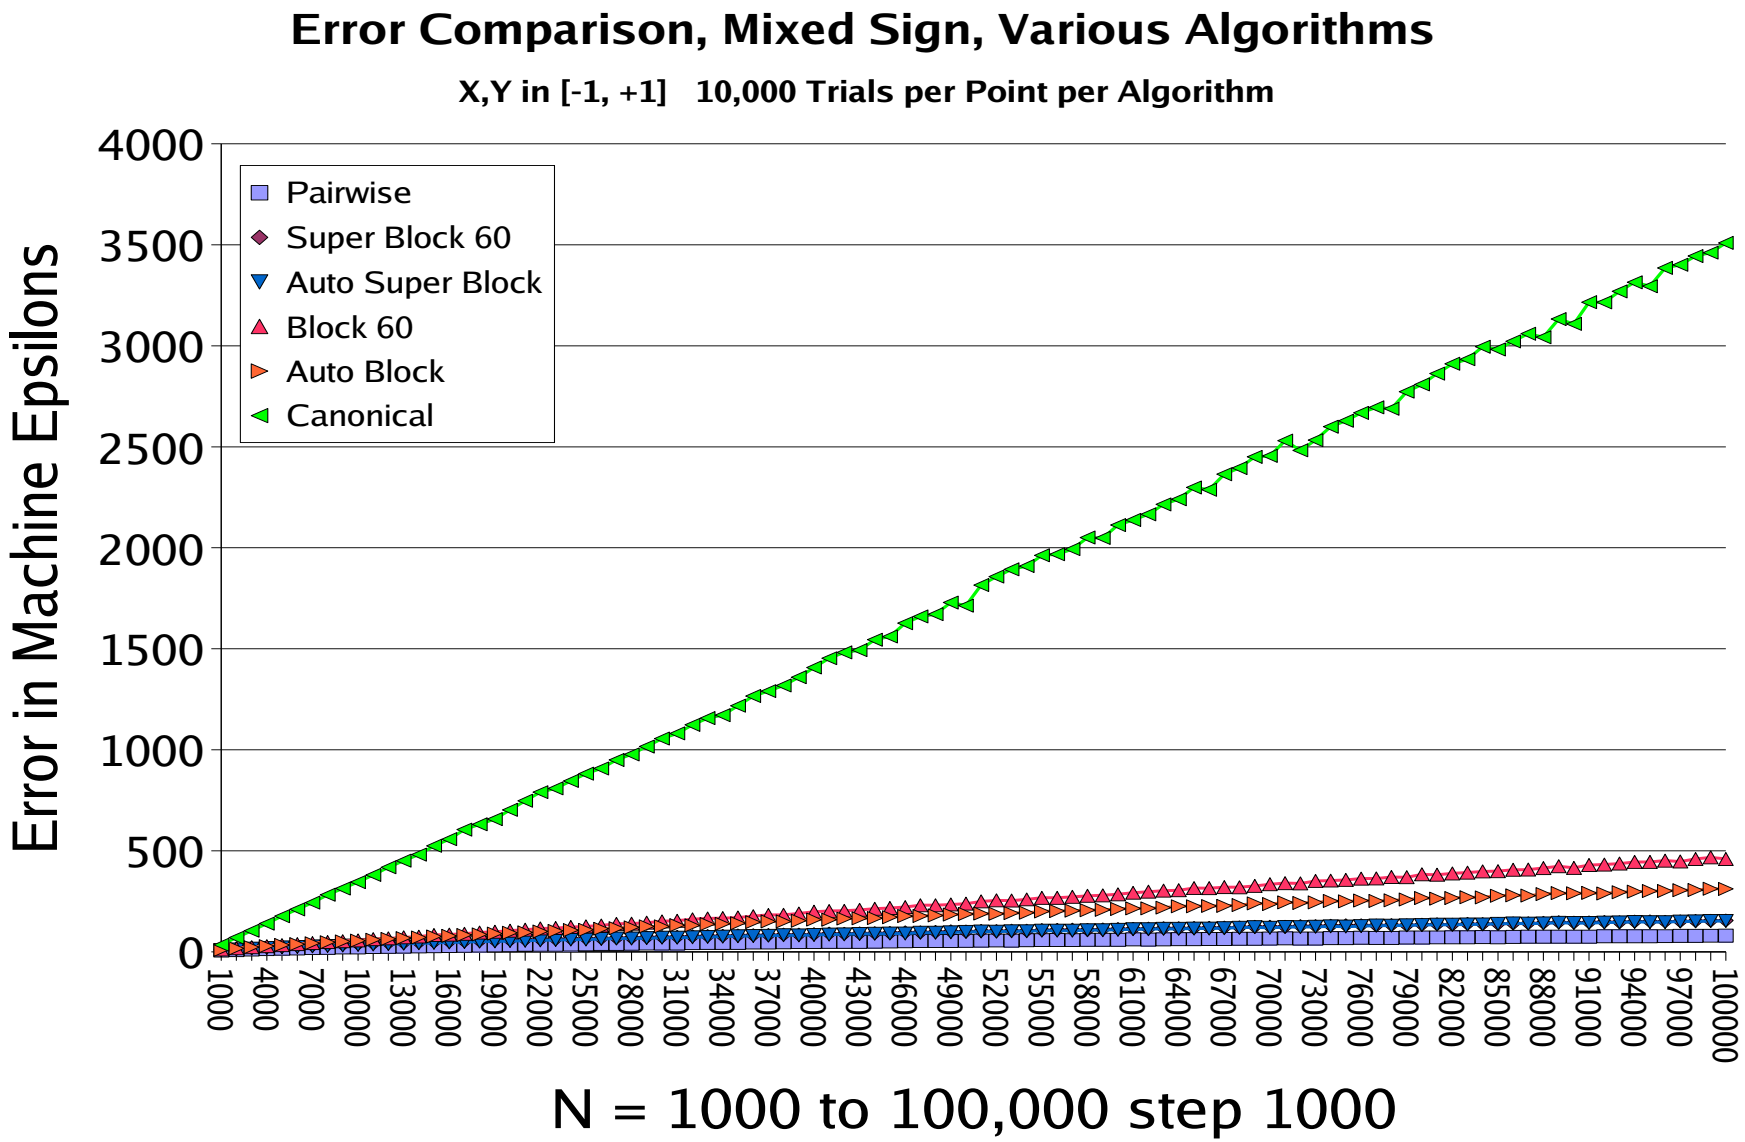

Absolute Error by Algorithm - All Positive

Averaged over 10,000 random vectors