The Value of Variance

Jesun Firoz, **Martina Barnas,** Marcin Zalewski, and
Andrew Lumsdaine

CREST, Indiana University School of Informatics and Computing,
Bloomington, IN, USA

Center for Research in Extreme Scale Technologies (CREST) runtime and applications for exascale

Outline

- "Theory"
  - National Institute of Standards and Technology (NIST) guidelines
    Law of propagation of uncertainty
  - ADJUSTED SPEEDUP EQUATION

- Experimental Setup
  hardware
  runtimes
  algorithms and input

- Two comments about the algorithms

- Experimental results

- Summary and Conclusions

“THEORY”

National Institute of Standards and Technology (NIST) guidelines

- Consider a quantity being measured, Y, in terms of other quantities, X_i :

$$Y = f(X_1, X_2, \ldots, X_N)$$

- The estimate: $y = f(x_1, x_2, \ldots, x_N)$
- Sample mean of inputs :

$$x_i = \overline{X}_i = \frac{1}{n}\sum_{k=1}^{n} X_{i,k}$$

- Standard deviation as uncertainty measurement

$$u(x_i) = u_i = \sigma_i = \left(\frac{1}{n(n-1)}\sum_{k=1}^{n}(X_{i,k} - \overline{X}_i)^2\right)^{\frac{1}{2}}$$

The law of propagation of uncertainty

- If multiple quantities $X_1, X_2, .., X_N$ are involved in the calculation of estimate y, the combined standard uncertainty is the positive square root of the estimated variance $\sigma^2(y)$ obtained from

$$\sigma^2(y) = \sum_{i=1}^{N} \left(\frac{\partial f}{\partial x_i} \right)^2 \sigma^2(x_i) + 2 \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \frac{\partial f}{\partial x_i} \frac{\partial f}{\partial x_j} \sigma(x_i, x_j).$$

Adjusted Speedup Equation

- Consider average execution time of two algorithms A and B as $\bar{t}_A$ and $\bar{t}_B$, (with uncertainty)
- Speedup equation:

$$S = \frac{\bar{t}_A}{\bar{t}_B}.$$

- Proposed adjusted speedup equation:

$$S_{adj} = S \pm \sigma,$$

where

$$\sigma^2 = \frac{1}{\bar{t}_B^2}{\sigma_A}^2 + \frac{\bar{t}_A^2}{\bar{t}_B^4}{\sigma_B}^2.$$

EXPERIMENTAL SETUP

Experimental Setup - hardware

- Experiments on Indiana University's BigRed 2 Cray XE6/XK7 supercomputer
  - two AMD Opteron 16-core x86 64 CPUs per node
  - 64 GB of RAM per node
  - Gemini interconnect
- 1 to 16 compute nodes for weak scaling (starting from scale 14)
  - 16 threads per node

Experimental Setup –runtime

- (hardware –BigRed2 CRAY)

- two different runtimes:
  - HPX-5 – under development
  - AM++ - less feature rich, but well optimized to balance quick work delivery vs. communication overhead

Experimental Setup – algorithm and data

- (hardware – BigRed2 CRAY)
- (runtime – HPX-5, AM++)

- two different distributed algorithms for SSSP problem
  - KLA with K=2
  - Δ-stepping with Δ=1
  - **NOT OPTIMIZED FOR MAXIMUM PERFORMANCE**!

- Input graph
  - With Graph500 RMAT generator
  - Maximum edge weight 255 and 100
  - Scale 14 (1 node) to scale 18 (16 nodes)

Experimental setup

- (hardware – BigRed2 CRAY)
- (runtime – HPX-5, AM++)
- (algorithms – KLA SSSP, Δ-stepping SSSP)
- (input – Graph500)

- Experiments:
  - 5 runs for each scale
  - 8 problem instances per run (different starting point – source)

TWO COMMENTS ABOUT THE ALGORITHMS

Distributed Algorithms for SSSP

K-Level Asynchronous Algorithm

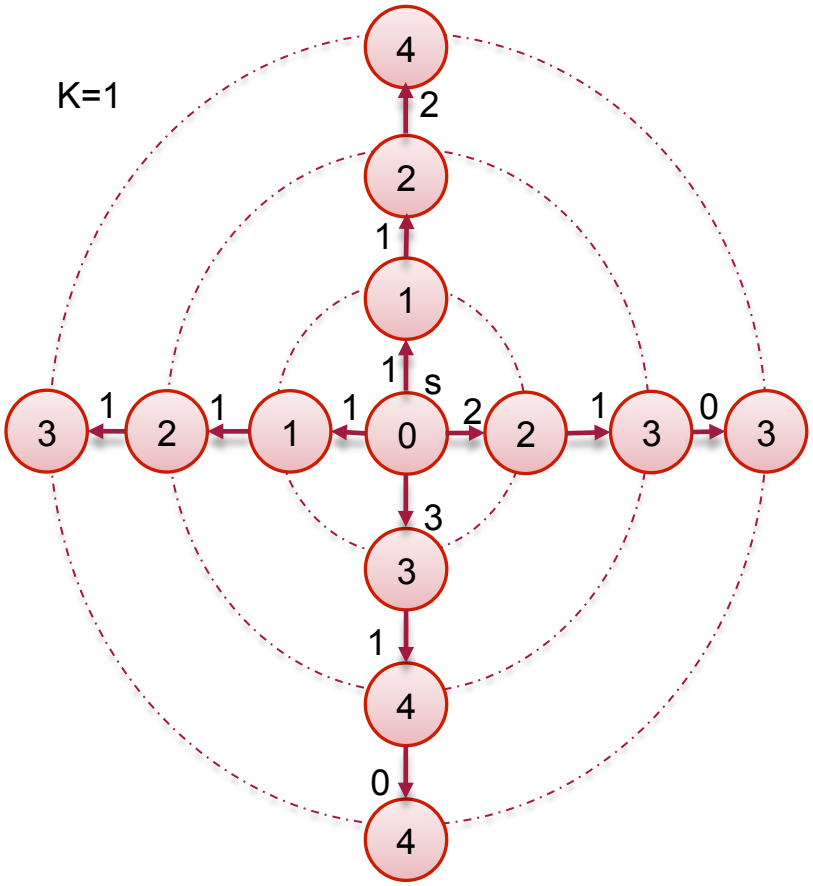

Δ-stepping Algorithm

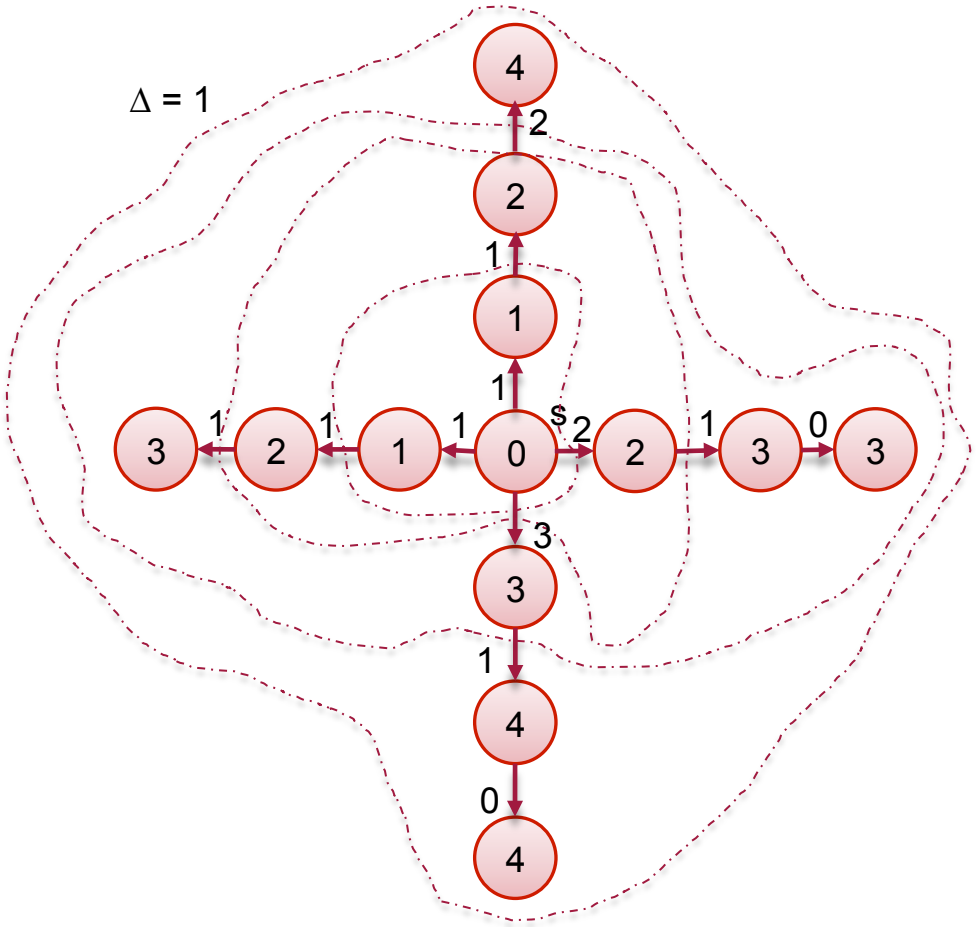

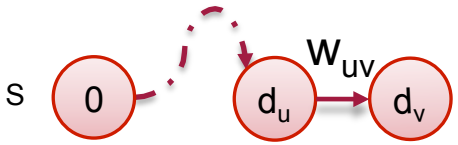

S : Source
d_u :Distance from source to u
w_{uv} : weight between u and v

Distributing work over nodes – Strangler Effect

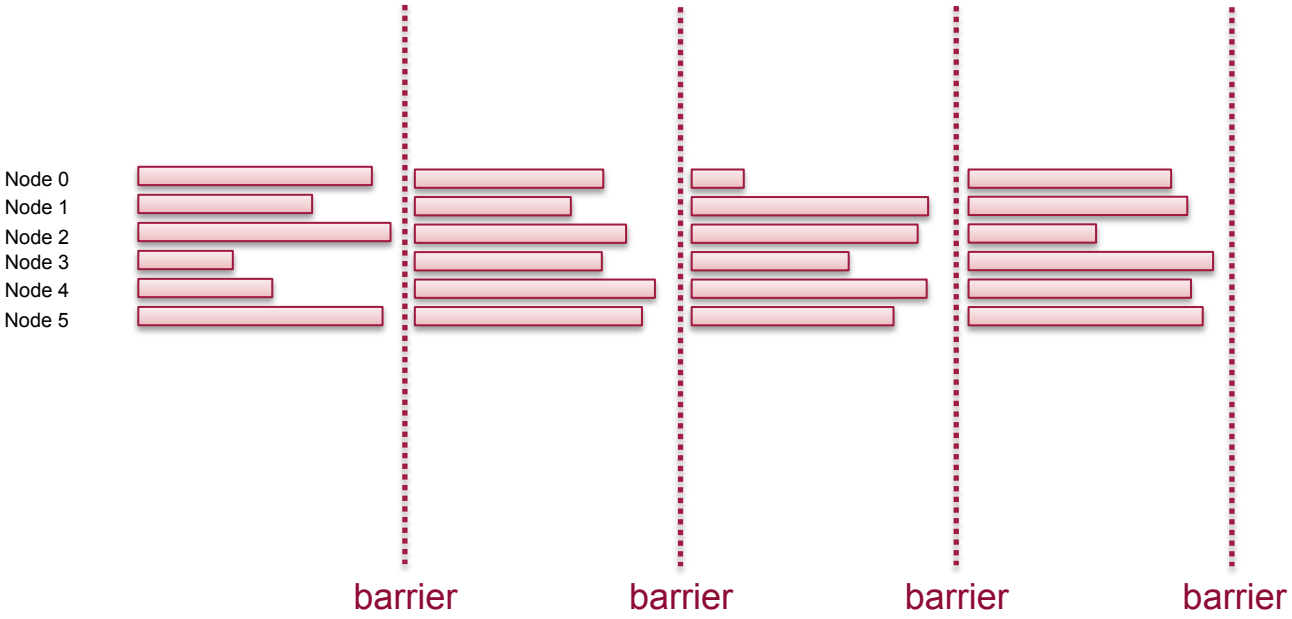

EXPERIMENTAL RESULTS

Reporting Speedup Uncertainty: On HPX-5 Runtime

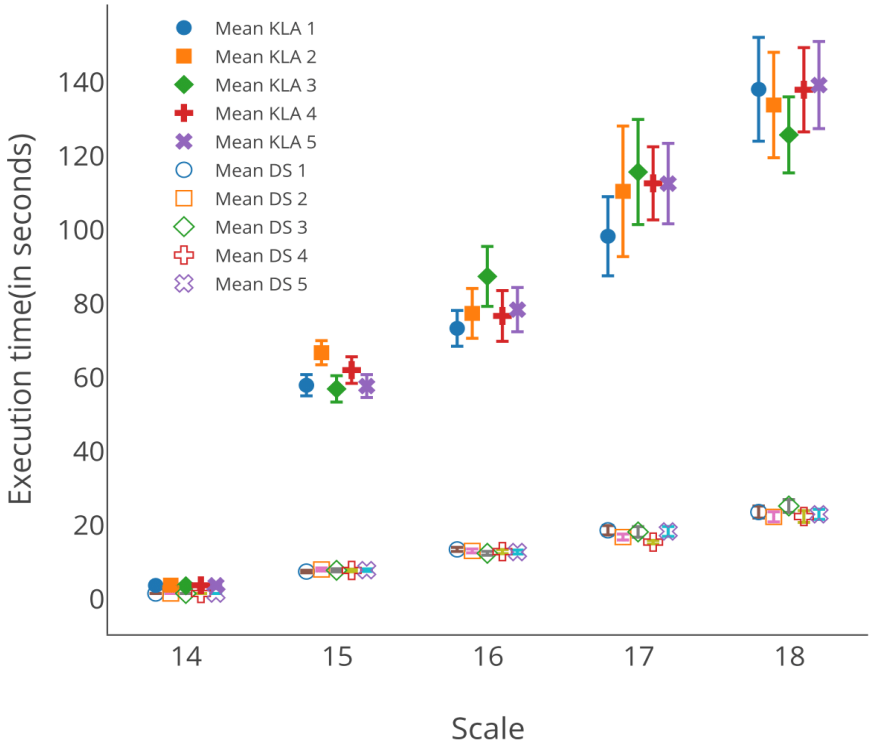

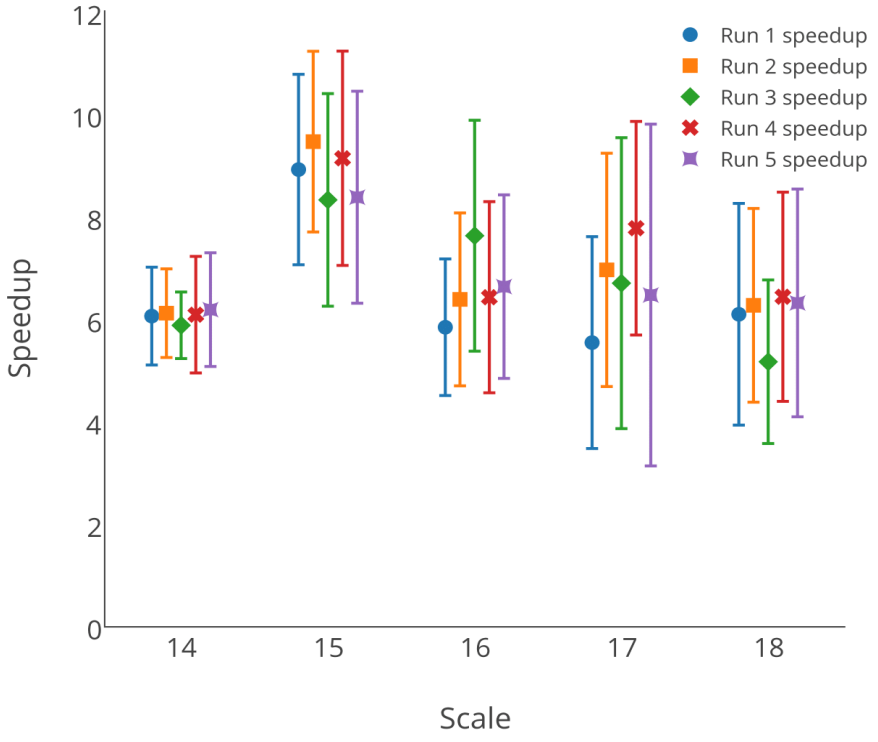

With max edge weight 255

Reporting Speedup Uncertainty: On HPX-5 Runtime (cont.)

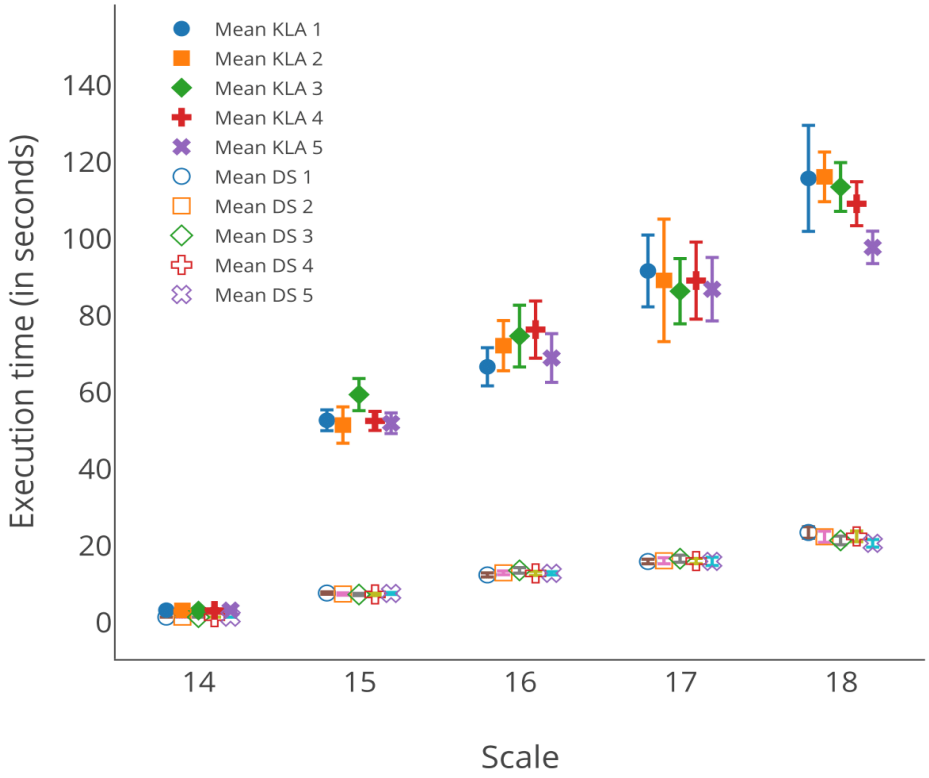

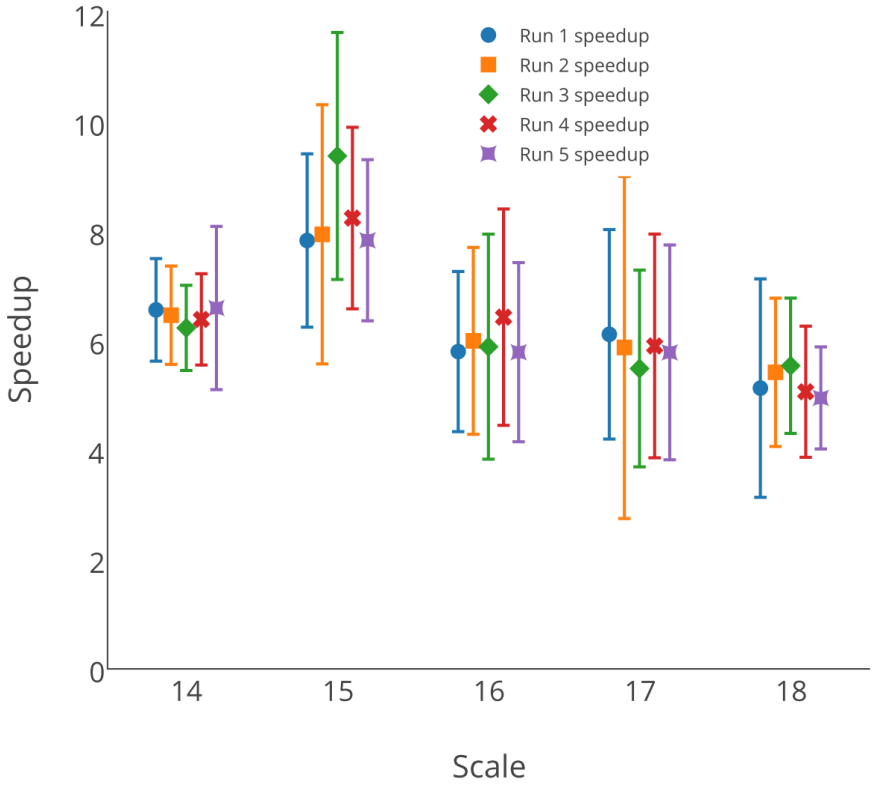

With max edge weight 100

Reporting Speedup Uncertainty: On AM++ Runtime

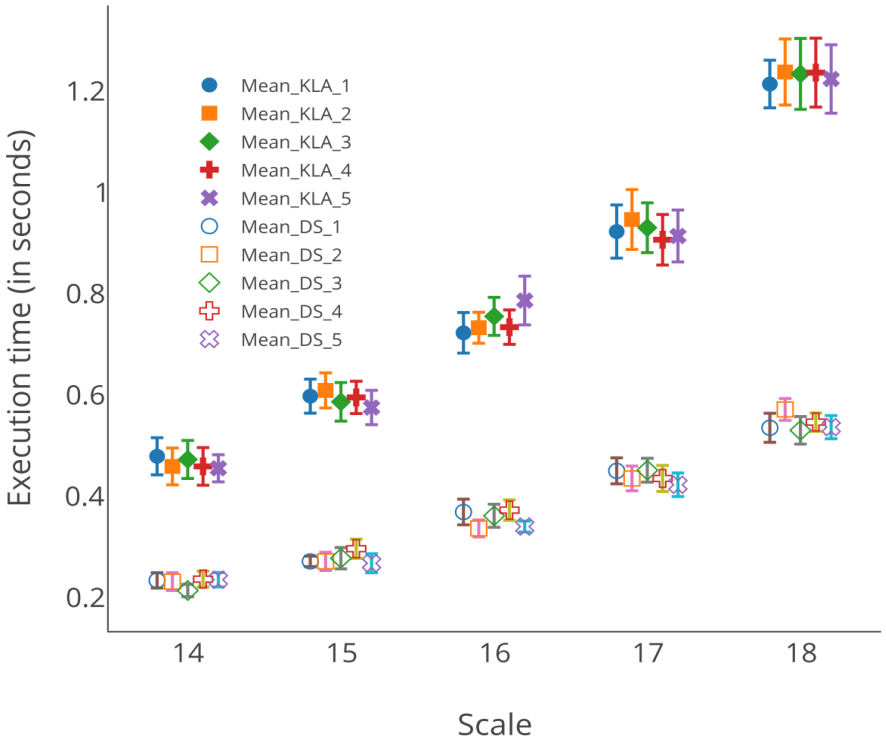

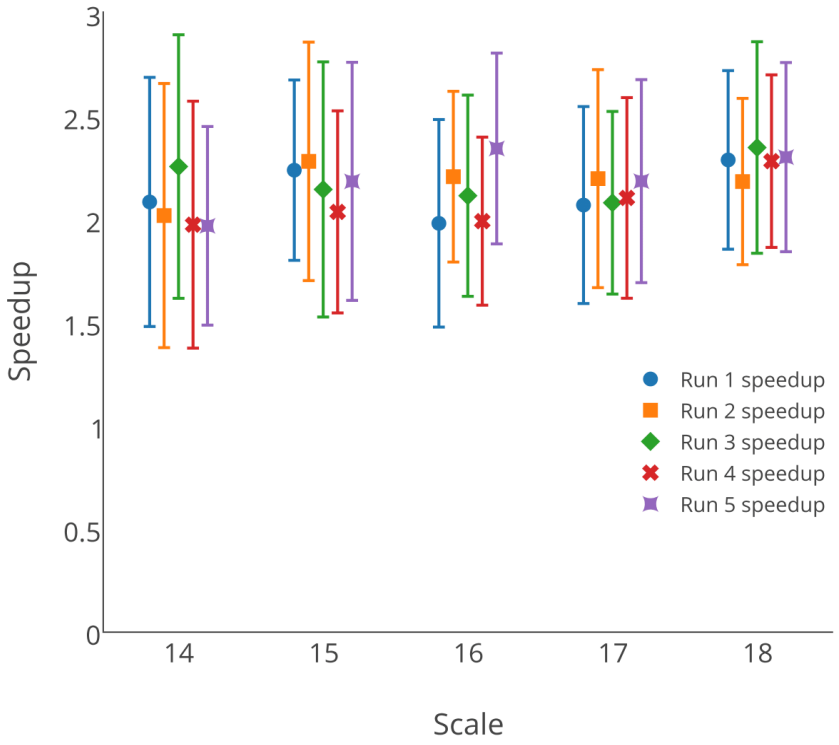

With max edge weight 100

Relative Standard Uncertainty: On AM++ Runtime (execution times)

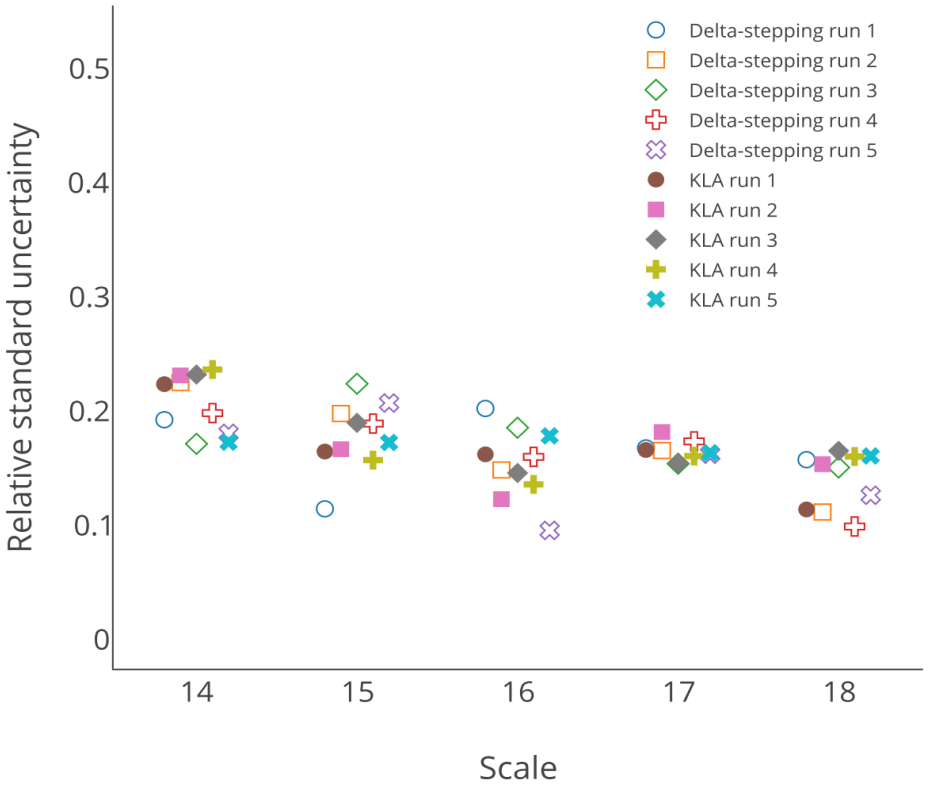

With max edge weight 100

Relative Standard Uncertainty: On HPX-5 Runtime (execution times)

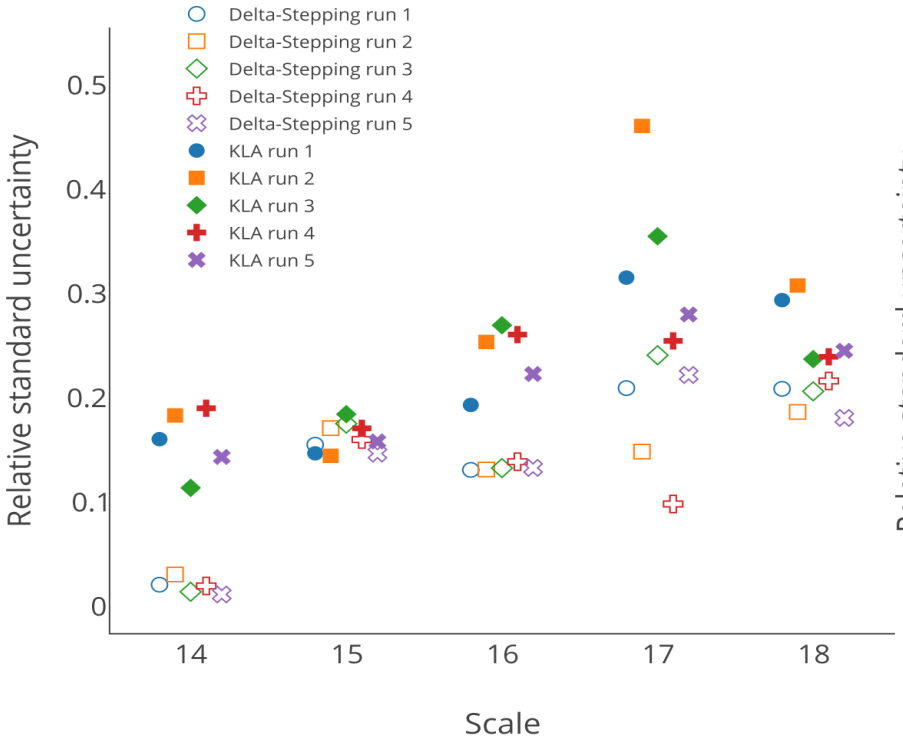

With max edge weight 255

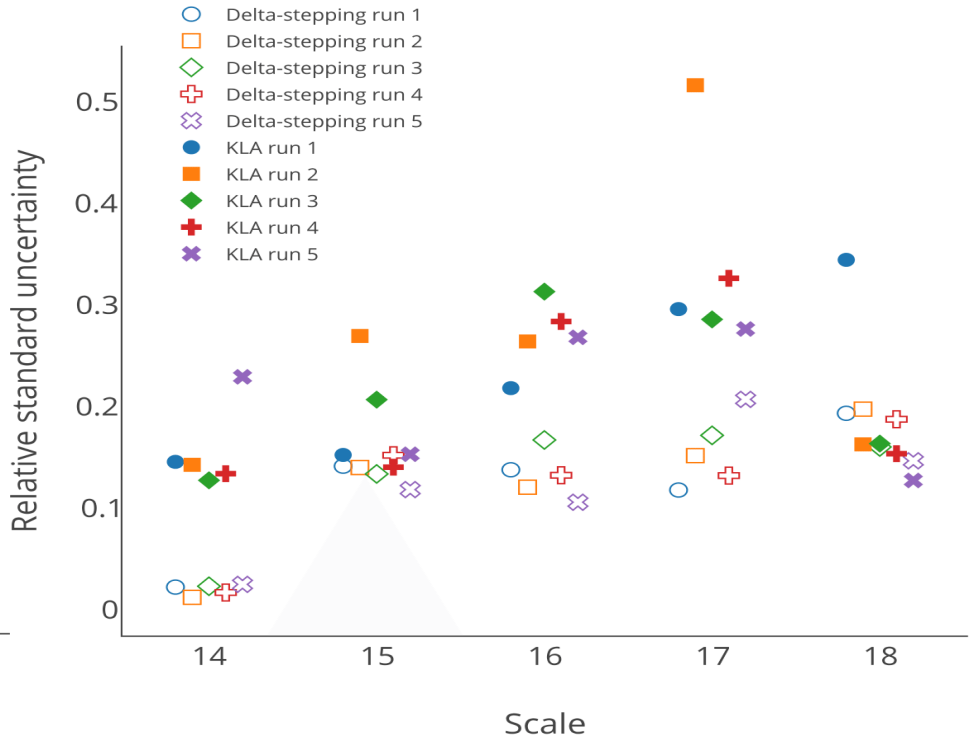

With max edge weight 100

SUMMARY AND CONCLUSIONS

Summary and Conclusions

- We looked at how standard deviation improves our insight from data
  - Simple to include
  - Additional insight that is lacking from aggregate measures (means)
  - Allows to drill down to what is interesting
  - Allows more meaningful comparison across different experiments

- Experiments using two distributed SSSP algorithms and two runtimes on BigRed2 CRAY

- **DISCLAIMER** – NOT about the algorithms and/or runtimes
  - During development of HPX-5; AM++ mature
  - Algorithm parameters not tuned

NOT INDICATIVE OF WHAT ALGORITHMS OR RUNTIMES CAN DO

- ALL ABOUT DATA and METHODOLOGY