

Profiling HPC Application Resilience using DisCVar

Stephen Penton

Mentors: Dr. Nathan DeBardeleben and Terry Grové

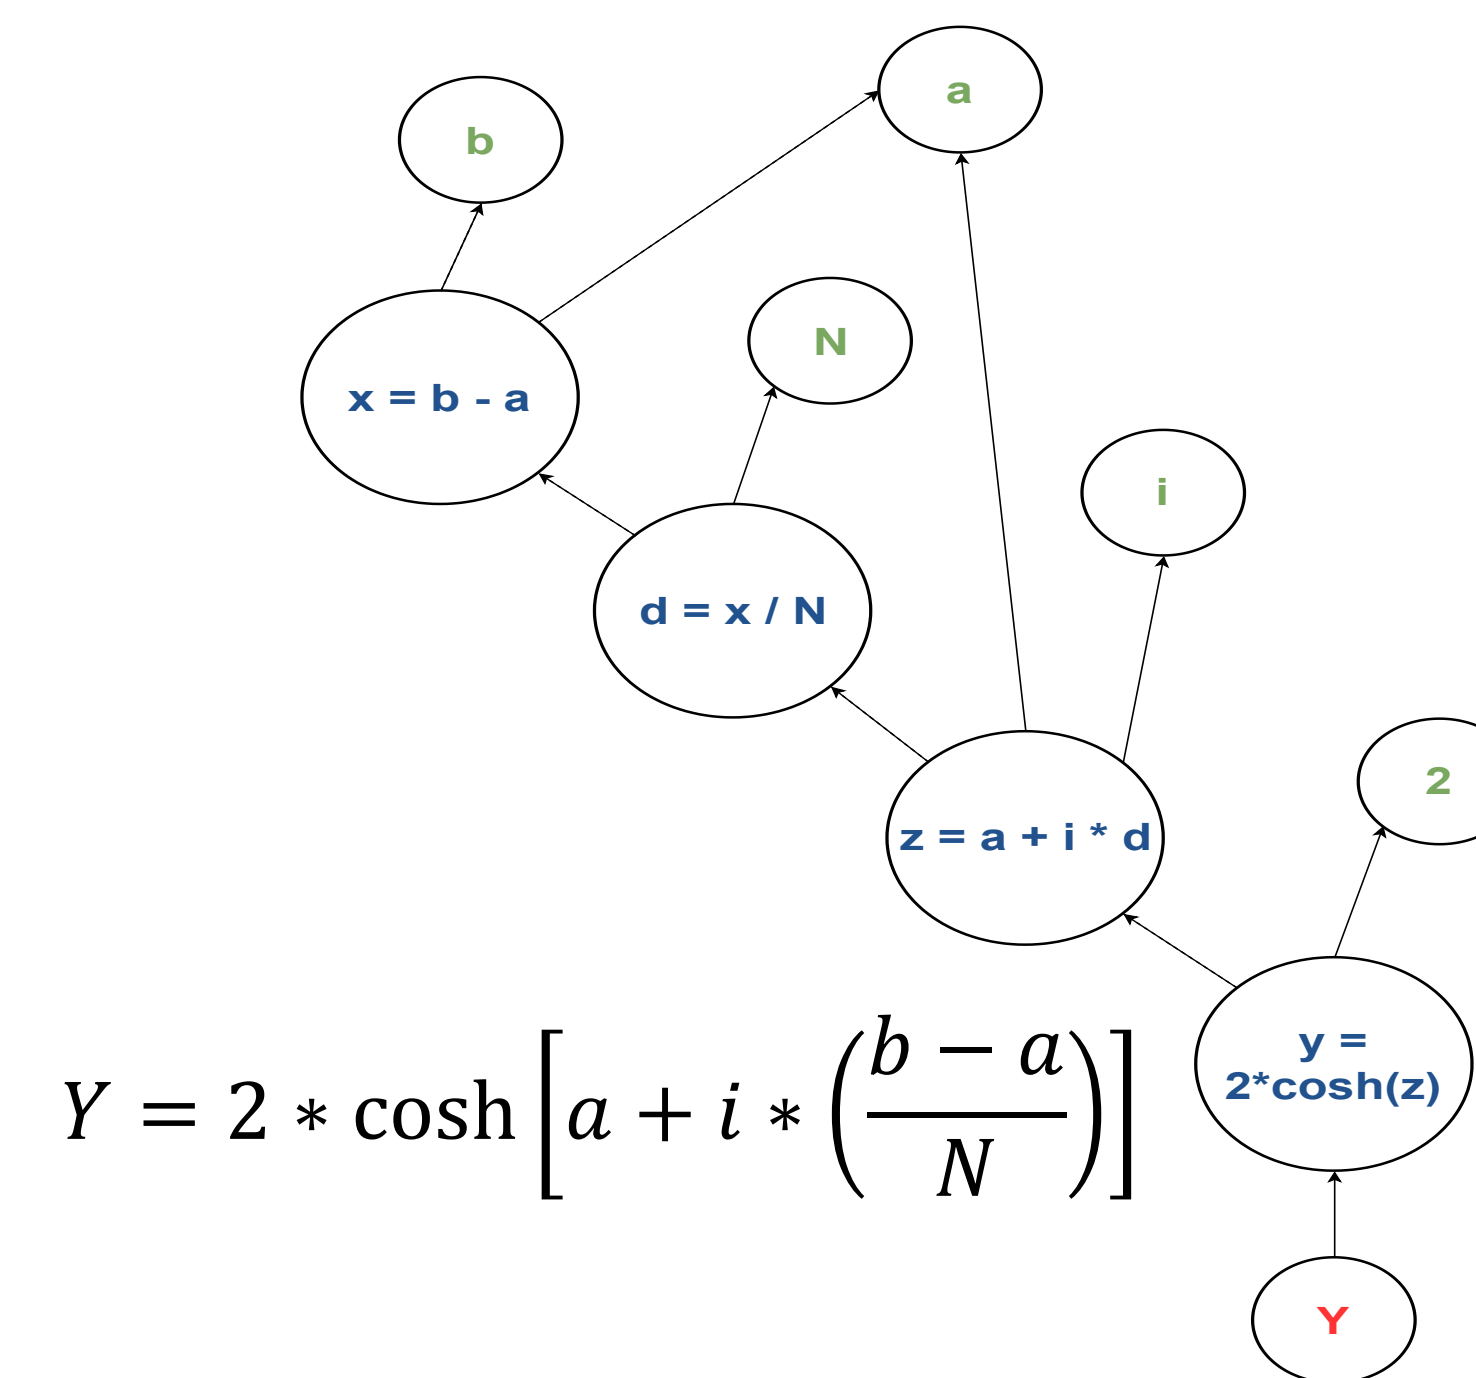
Motivation

Application resilience is an important part of future HPC systems. Doing so through exhaustive fault injection campaigns comes at a cost. Exhaustive campaigns are not practical for certain HPC applications. Thus, using tools such as DisCVar would allow us to:

- Identify application variables susceptible to faults
- Design targeted fault injection campaigns

Background

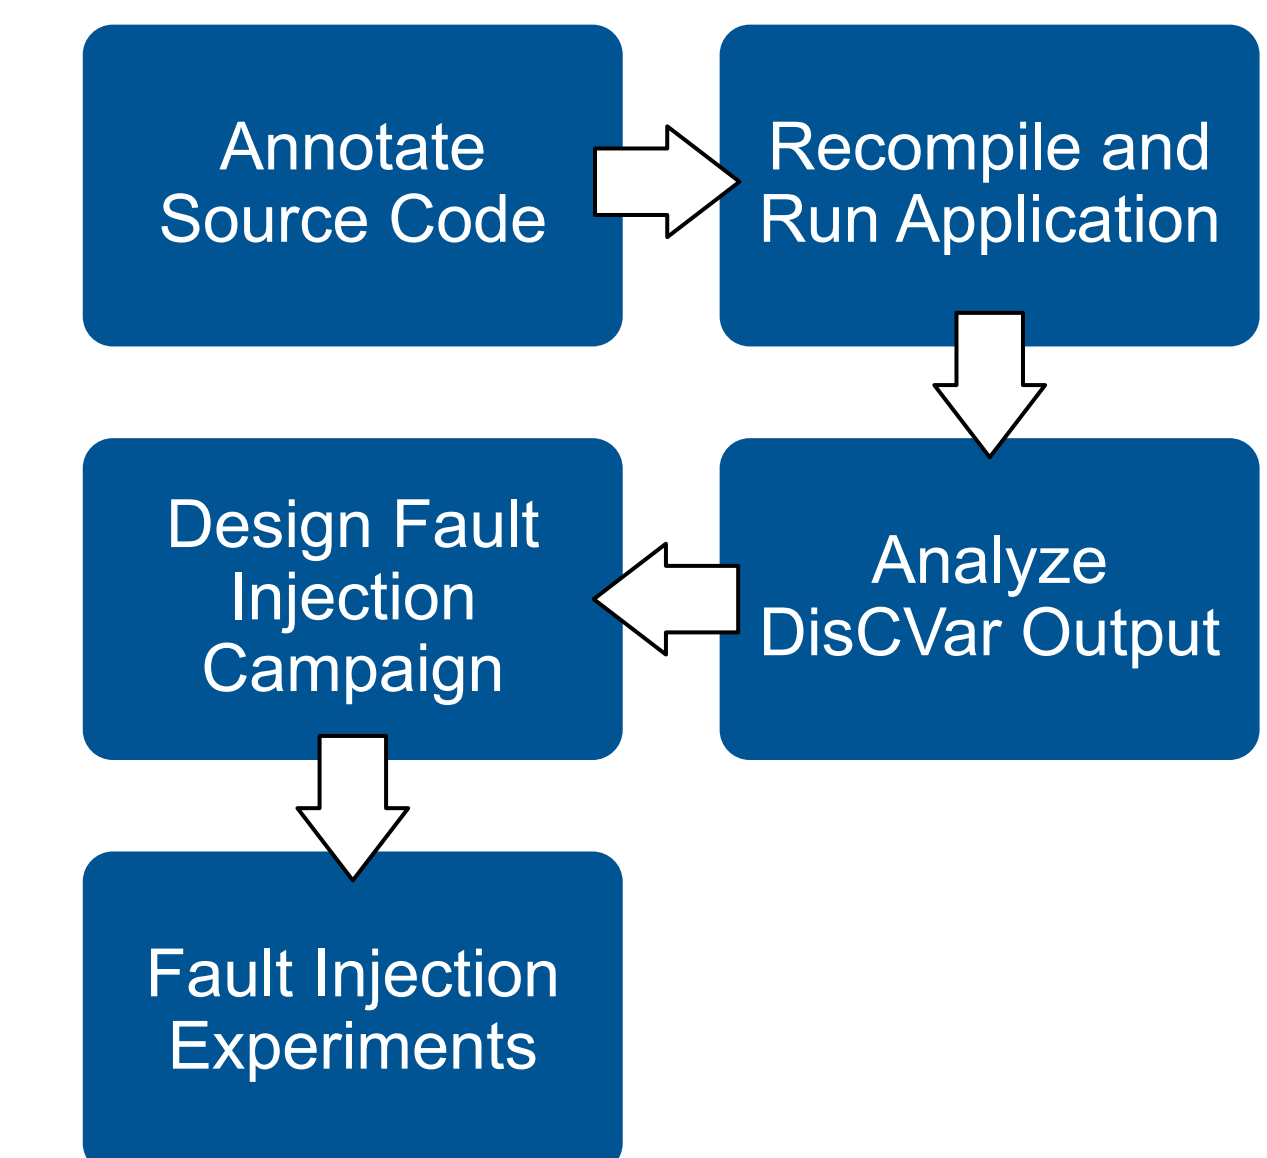
DisCVar is a tool developed at LLNL that uses Algorithmic Differentiation (AD) to identify critical variables within an application's source code. AD is a method of finding the derivative of a program by applying the chain rule. A partial derivative of the output is found with respect to input and intermediate variables. These are then used to rank variables according to how critical they are.



Methodology

- Annotate source code
- Recompile application
- Run to generate DisCVar output

These results can then be used to design fault injection campaigns. In this work, we applied DisCVar to a numerical integration application and VPIC, an HPC particle-in-cell simulation.



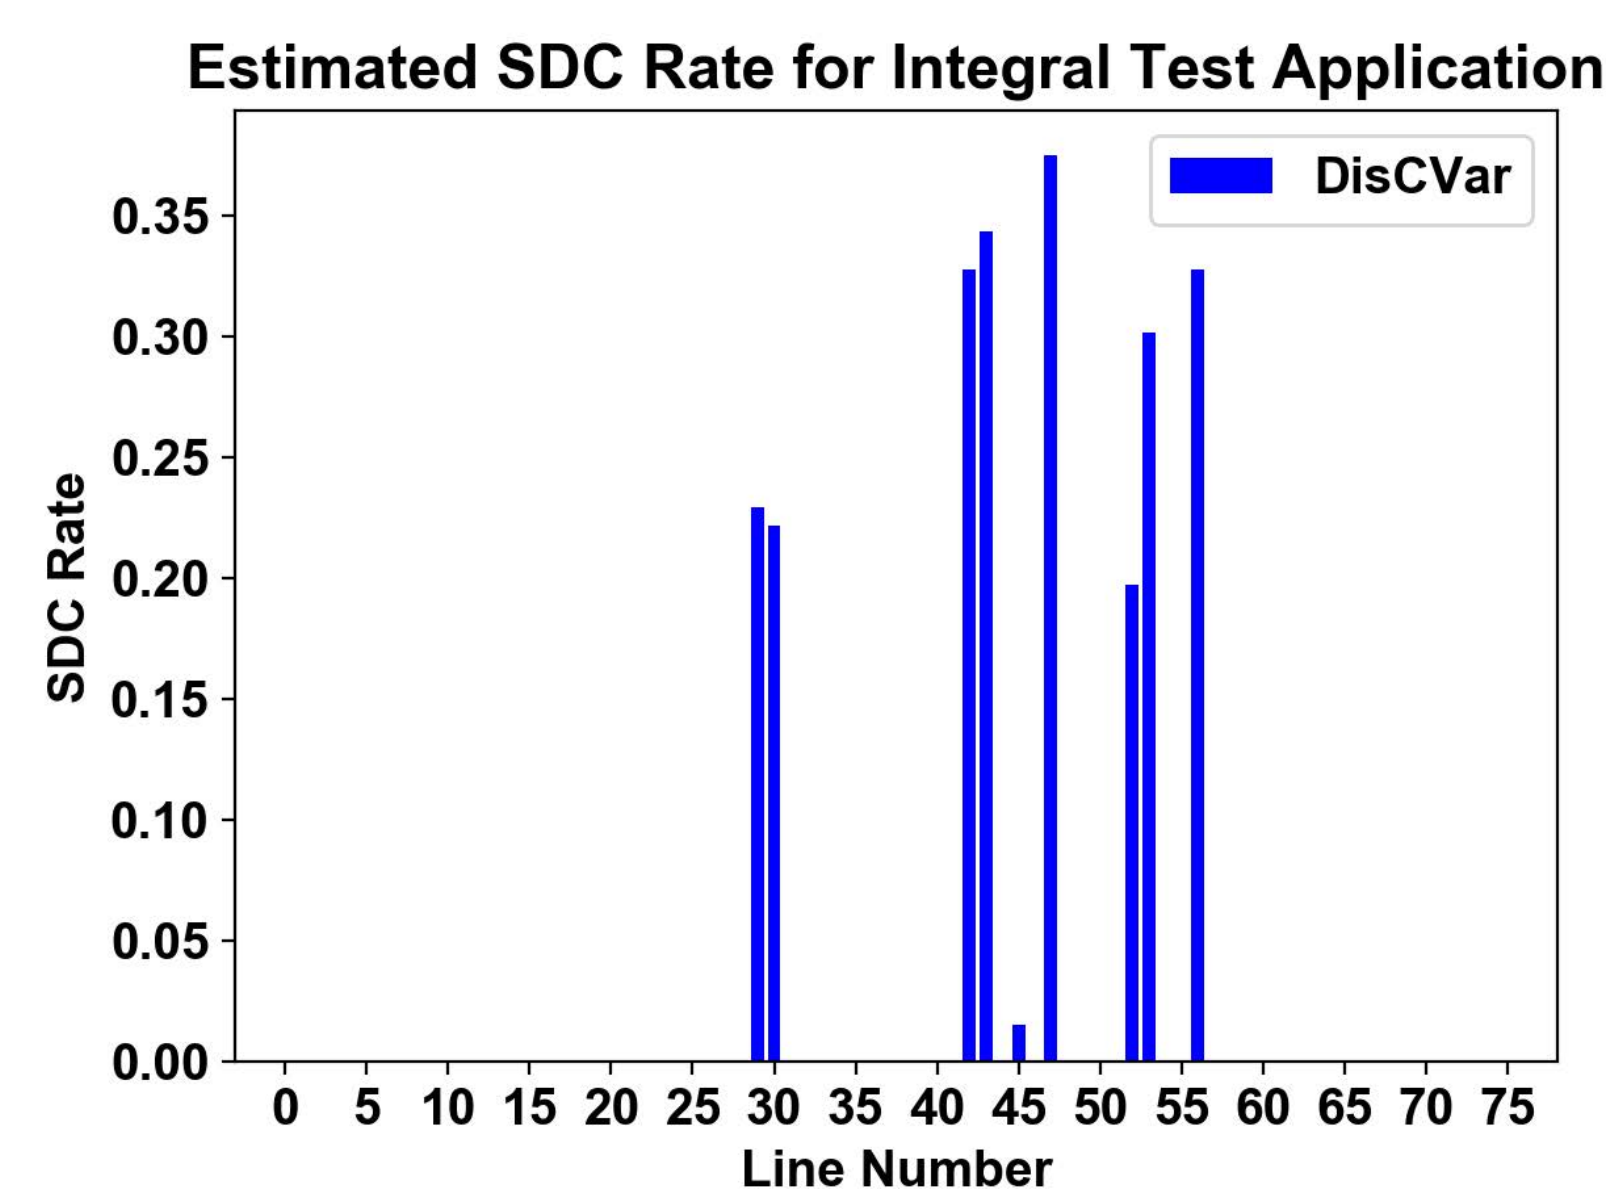
Analysis

Variable criticality is determined by an estimated silent data corruption (SDC) rate calculated via the derivatives from the AD process and a given threshold. Running DisCVar on our integration application allows for the source code to be highlighted based on SDC rates for each line. Lines of note include 42, 43, and 47.

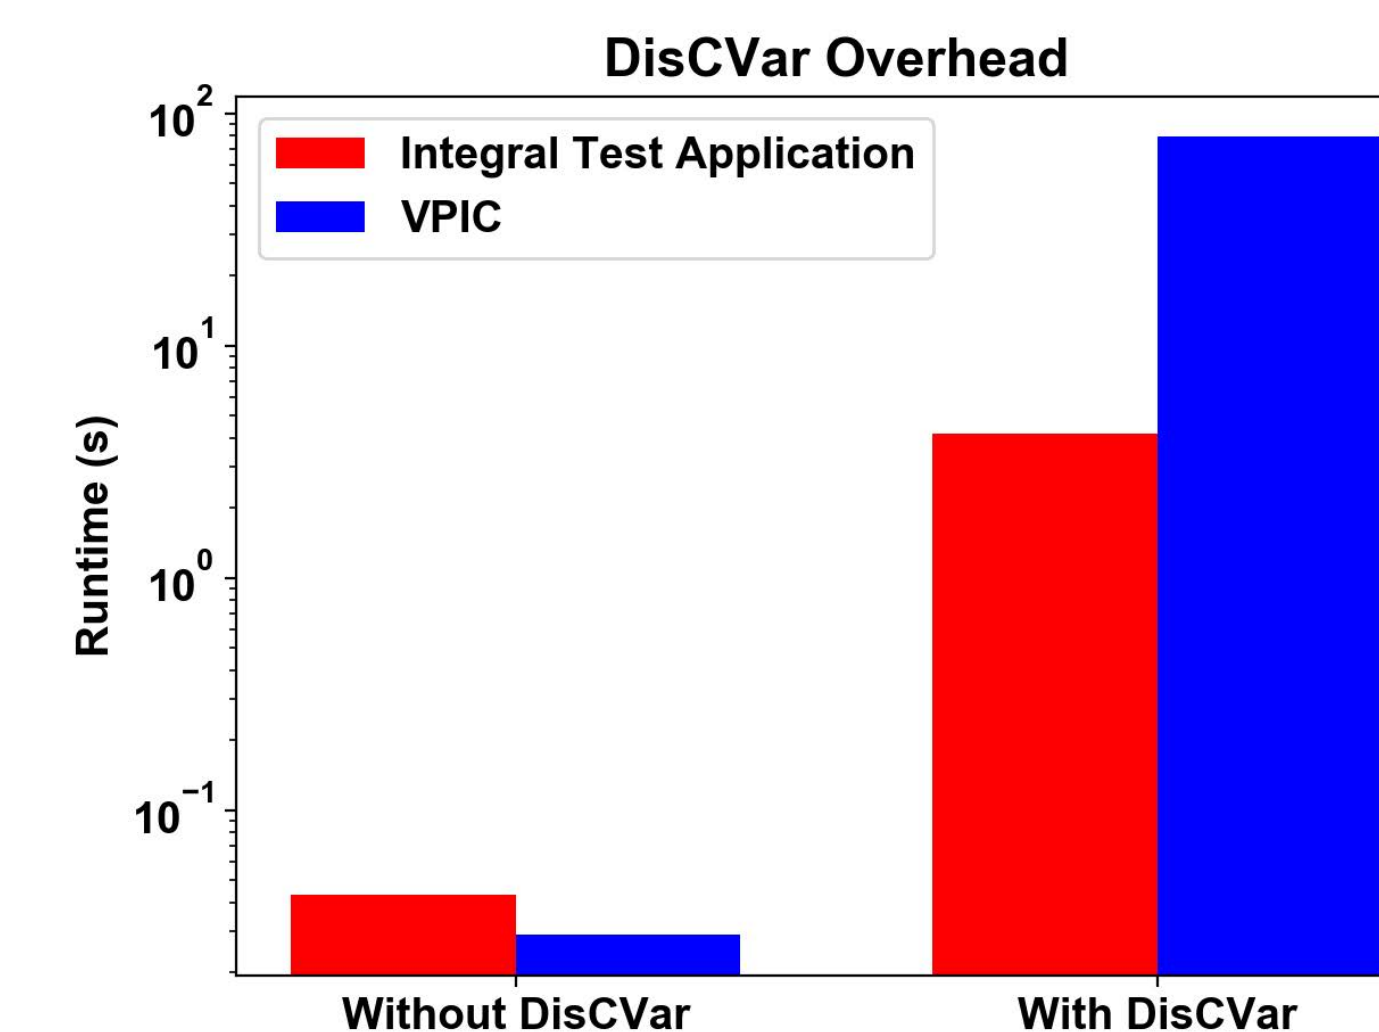
```

42  a = -2;
43  b = 3;
44
45  area = 0;
46
47  double dx = (b - a)/N; // width of each subinterval
    
```

DisCVar does introduce its own overhead, albeit for only a single run of the application. We see a runtime increase of approximately 2 orders of magnitude in our integration application; VPIC experienced an increase greater than 3 orders of magnitude.



The figure on the left describes the DisCVar estimated SDC rates by line number for our integration test application. The non-zero values range from 0.02 to 0.38. Of these 9 SDC rates, 7 are greater than 0.20. DisCVar's output facilitates the identification of the most critical variables within an application.



Future Work

A continuation of this work would include:

- Exhaustive FSEFI fault injection campaign on integral test application
- Further testing of DisCVar on VPIC
- Design and execution of FSEFI campaign on VPIC using DisCVar results.