

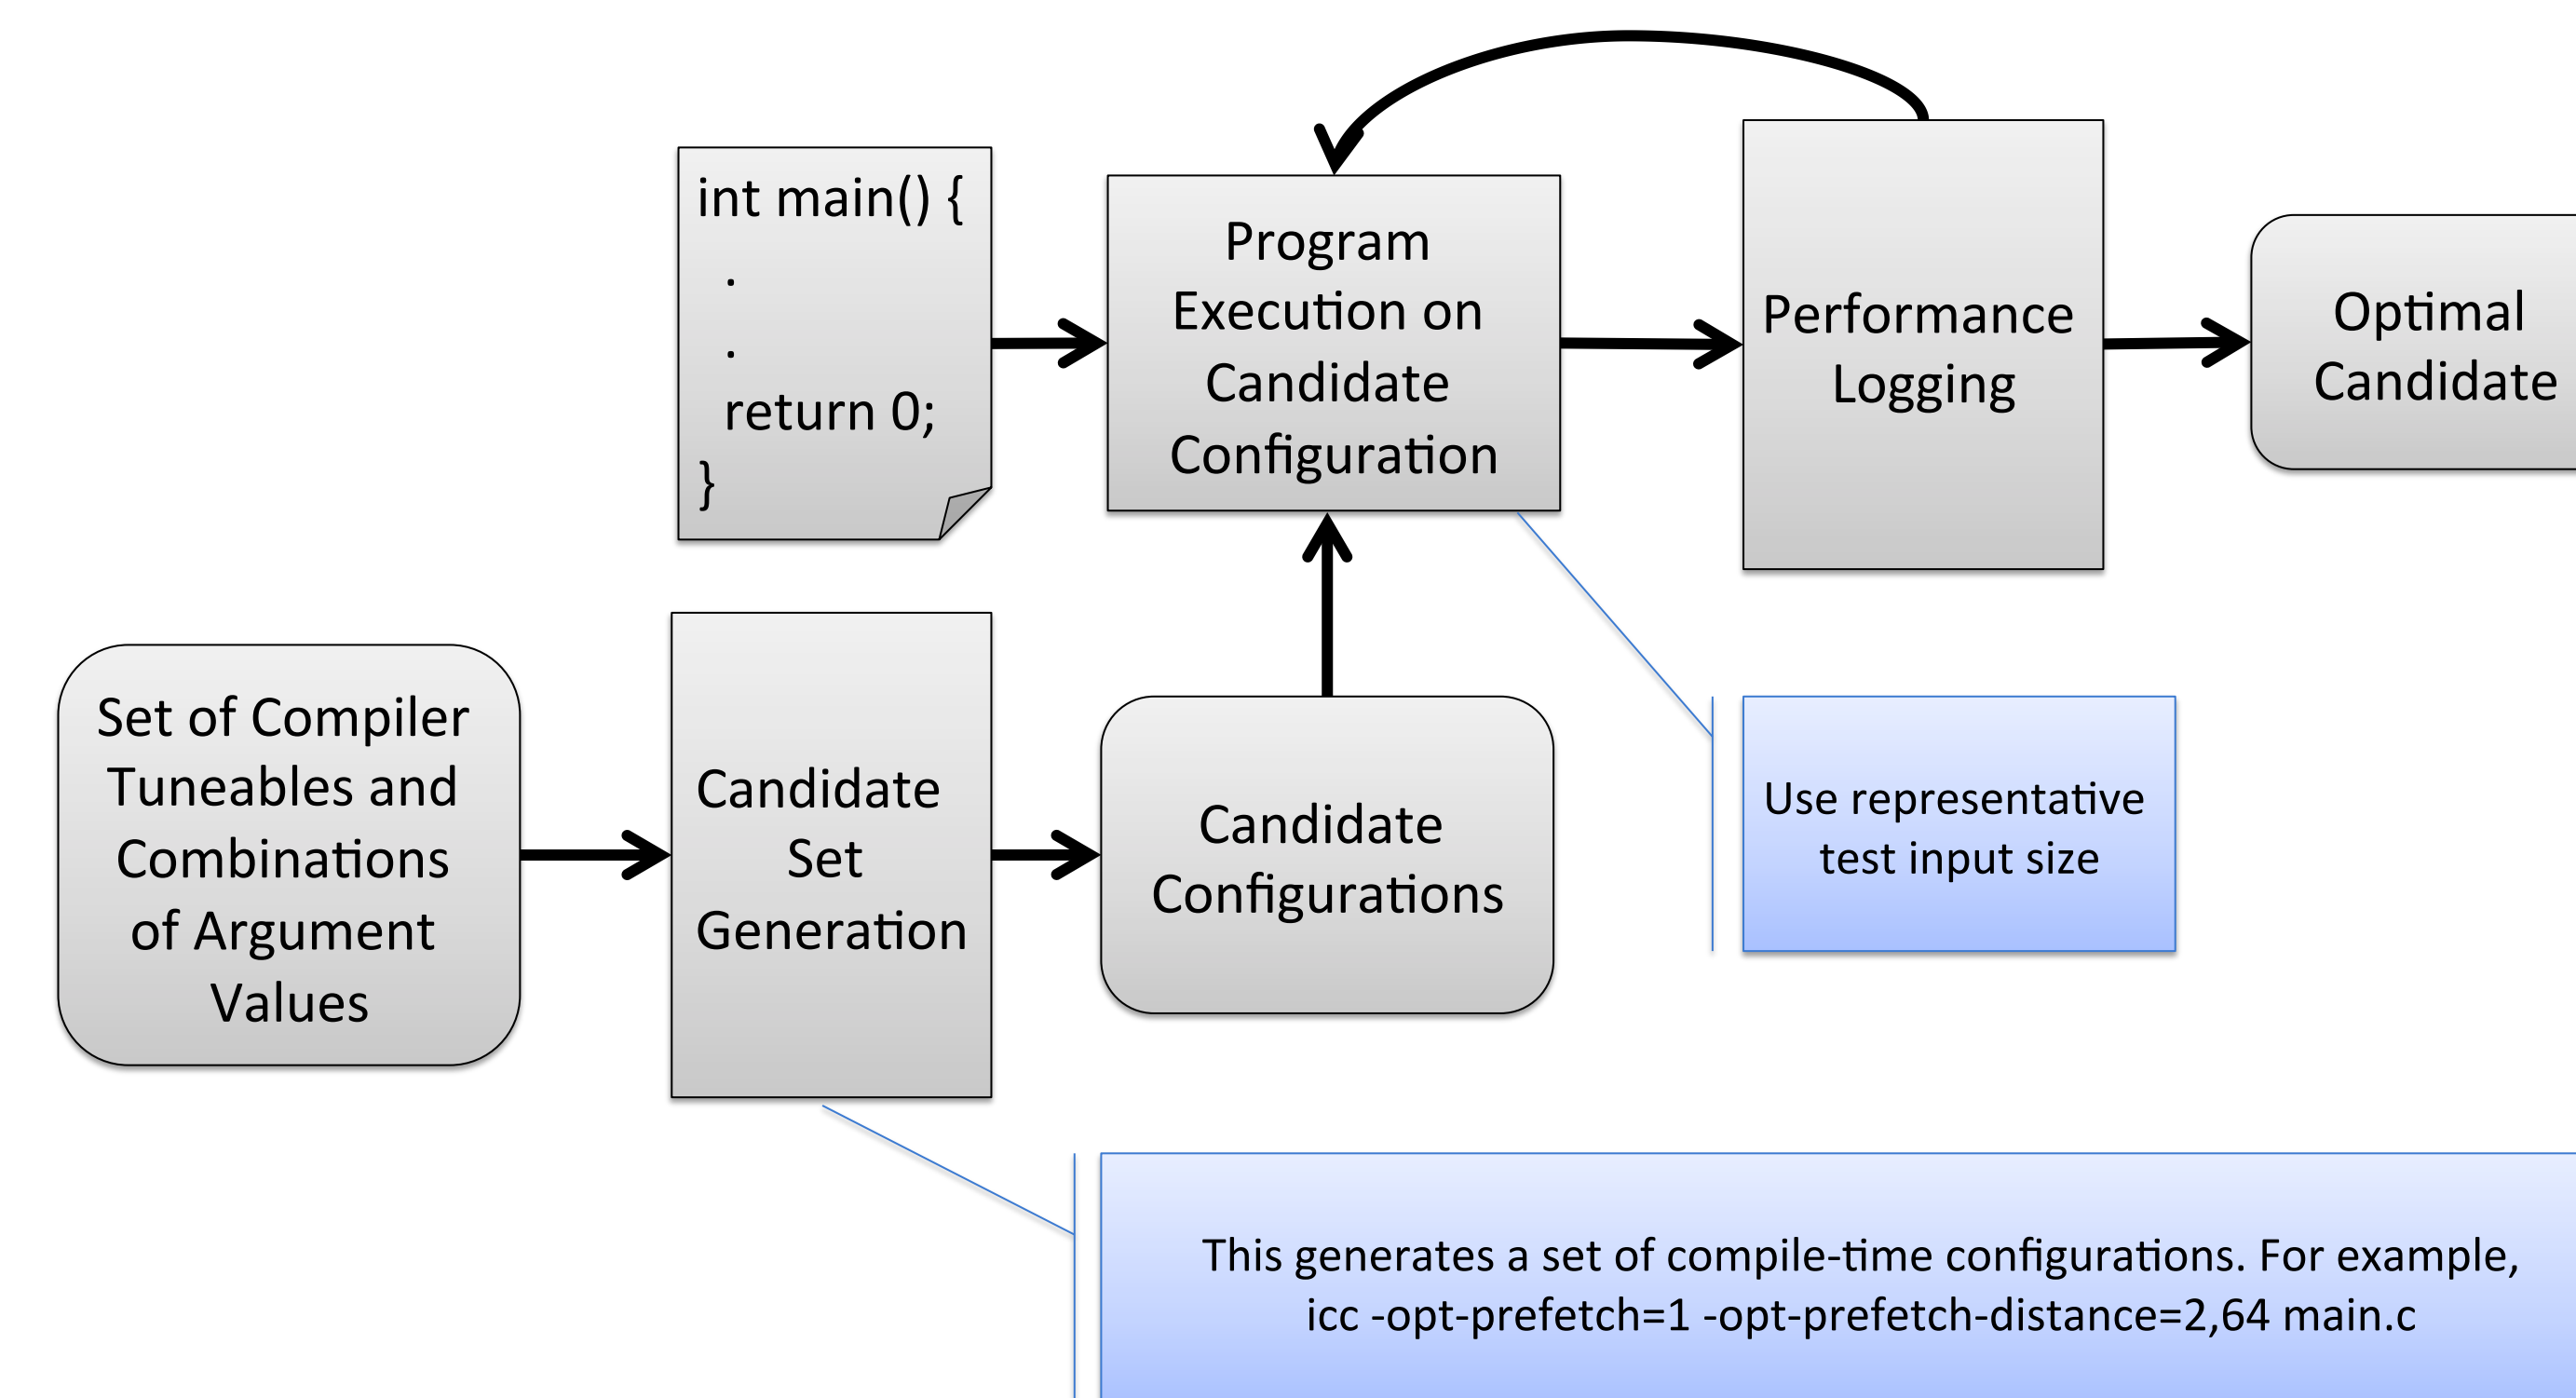
Prefetching is a well-known technique that is used to hide memory latency. Modern compilers analyze the program and insert prefetch instructions in the compiled binary. The Intel C Compiler (ICC) allows the programmer to specify two parameters that can help the compiler insert more accurate and timely prefetch instructions. The two parameters are `-opt-prefetch` and `-opt-prefetch-distance`. When unspecified, ICC uses default heuristics. In this work, we present the results of autotuning the two mentioned parameters and the its effect on performance and energy. The motivation for this stems from the understanding that most programmers choose to leave the default parameters unchanged, but significant improvements can be achieved by tuning these parameters. Choosing these parameters using analysis by hand can be challenging and time consuming as it requires knowledge of memory access patterns as well as significant time investment. We have developed a simple autotuning framework for the Xeon Phi architecture that automatically tunes these two parameters for any given program. We have used the framework on 4 memory intensive programs and gained up to 1.47 speedup and 1.39 greenup.

## Compiler Tuneables

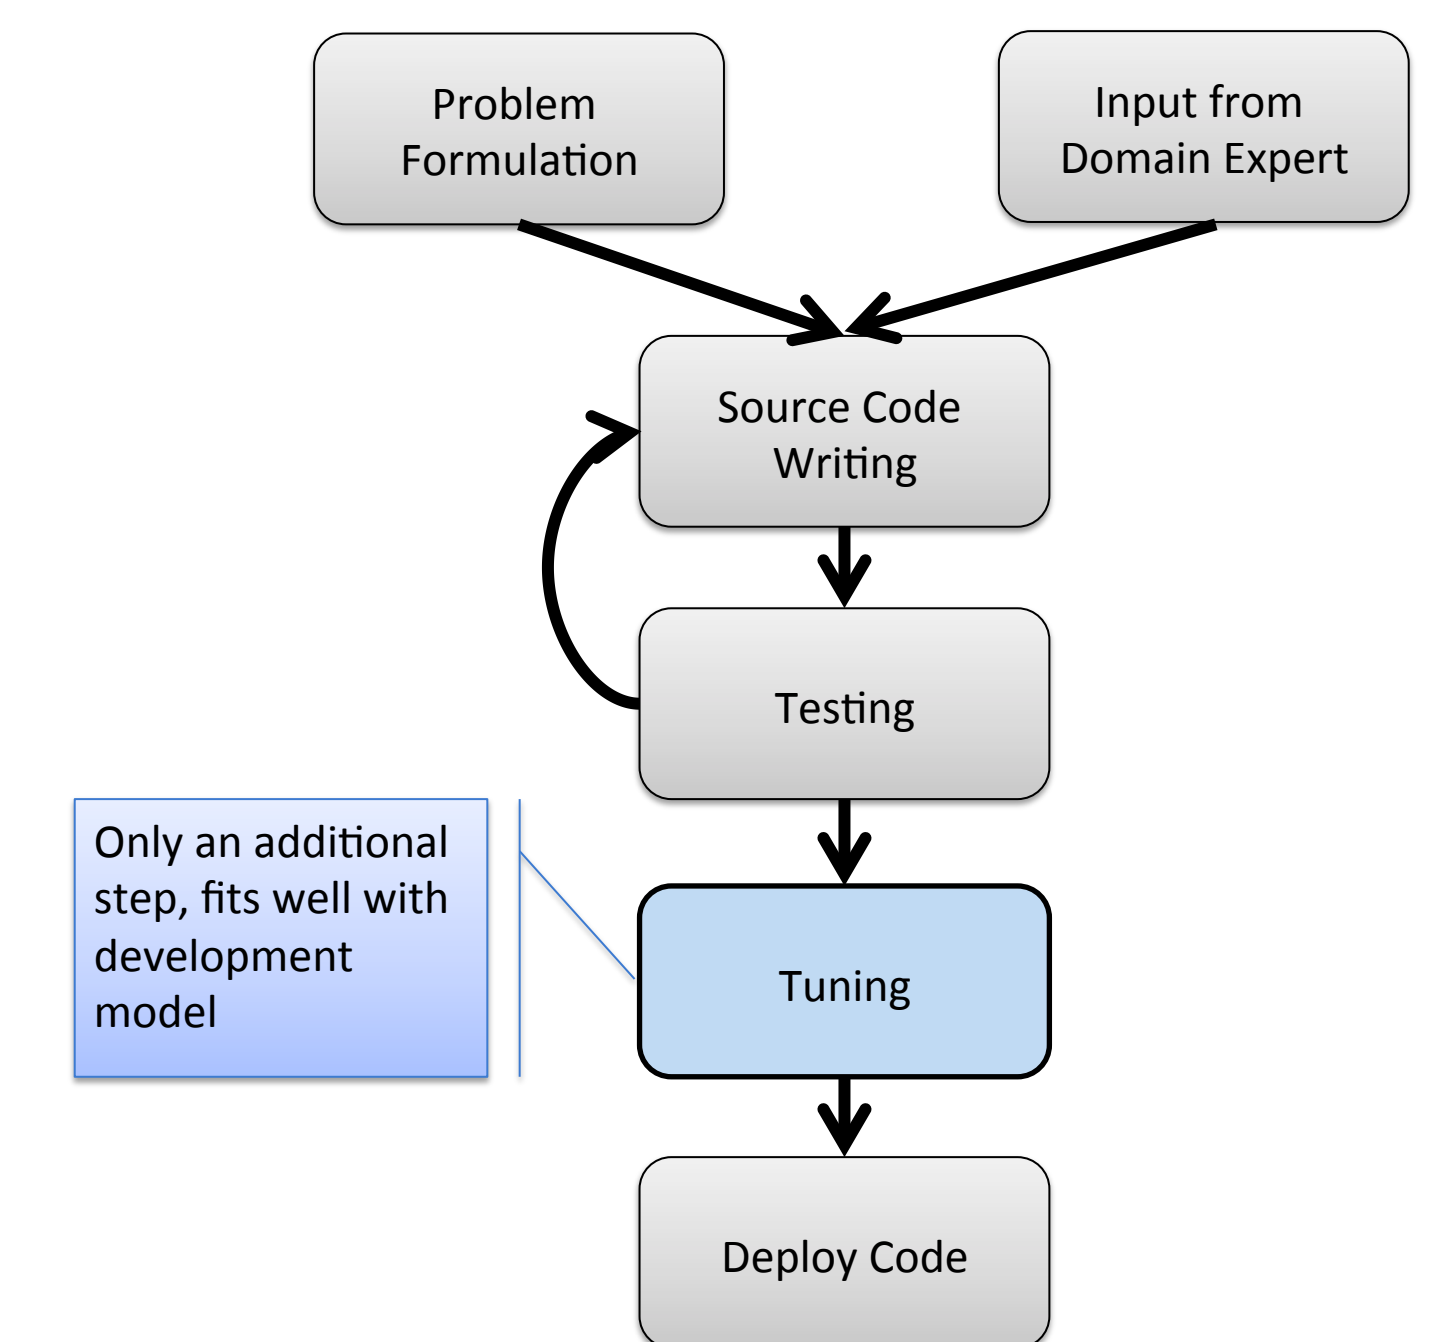
**-opt-prefetch:** This option in ICC can take 5 values - 0 through 4. A value of 0 means no prefetch instructions are added. Non-zero values indicate the amount of prefetching done.

**-opt-prefetch-distance:** This option takes two arguments, n1 and n2, where n2 is optional. The values of n1 and n2 represent the prefetch distance in terms of loop iterations. n1 represents the prefetch distance for prefetches from memory to L2, where n2 represents the same for prefetches from L2 to L1. In our framework we use values 2, 4, 8 ... 1024 for both arguments.

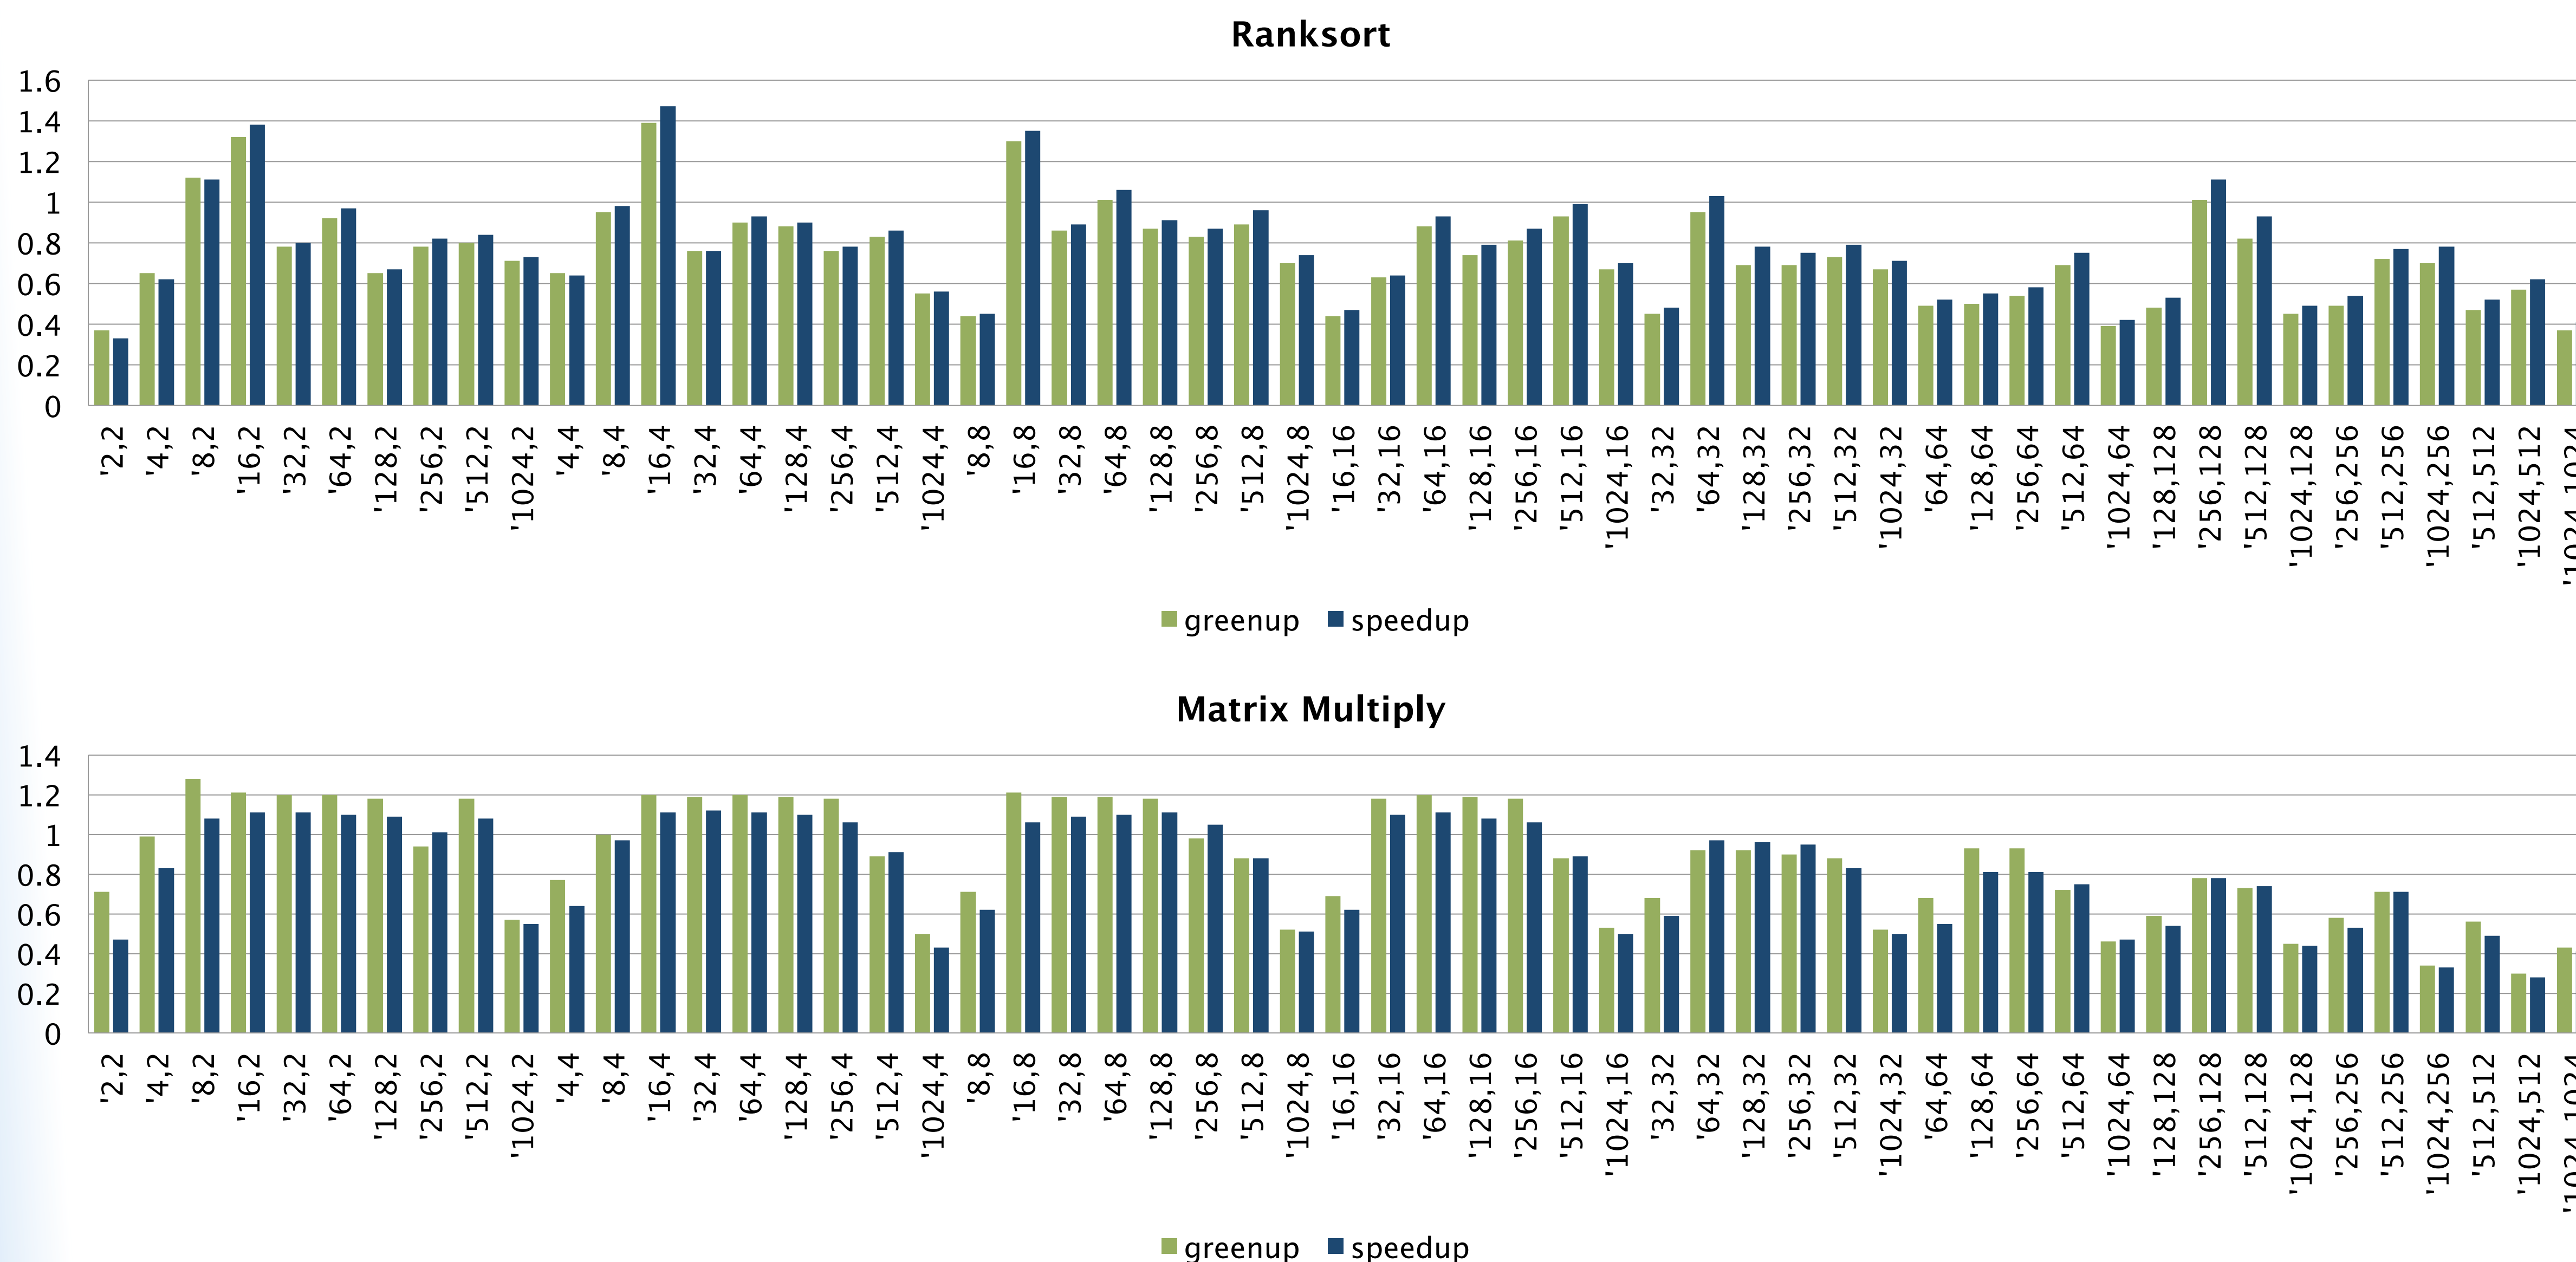
## Prefetch Configuration Selection Framework



## Scientific Program Development Model



## Results



The effect of tuning parameters `-opt-prefetch-distance` on two programs. These two program exhibit more interesting patterns than the other two. Additionally, the `-opt-prefetch-distance` argument has more effect than the `-opt-prefetch` argument.

## Program Description

1. Matrix multiplication: multiplication of two square matrices of size 5K each
2. Mandelbrot Fractal: computation of  $4001 \times 4001$  fractal with pixel depth of 256
3. Rank sort: sorting of 1.2M integers
4. Topological SSSP: shortest path calculation of NYC roadmap with 264,346 nodes and 733,846 edges

## Analysis

- Changing compile-time configurations affect both performance and energy
- When n1 is equal to n2 or is much larger than n2, performance and energy suffer a large penalty
- Smaller values of n2 are better
- Performance and energy appear to be more sensitive to n1