



Microsoft Research

Faculty Summit



Parallelizing Sequential Algorithms

Madan Musuvathi
Microsoft Research

joint work with
Saeed Maleki, Univ. of Illinois, Urbana-Champaign
Todd Mytkowicz, Microsoft Research

Hardware is Parallel



But, many important algorithms are 'inherently sequential'

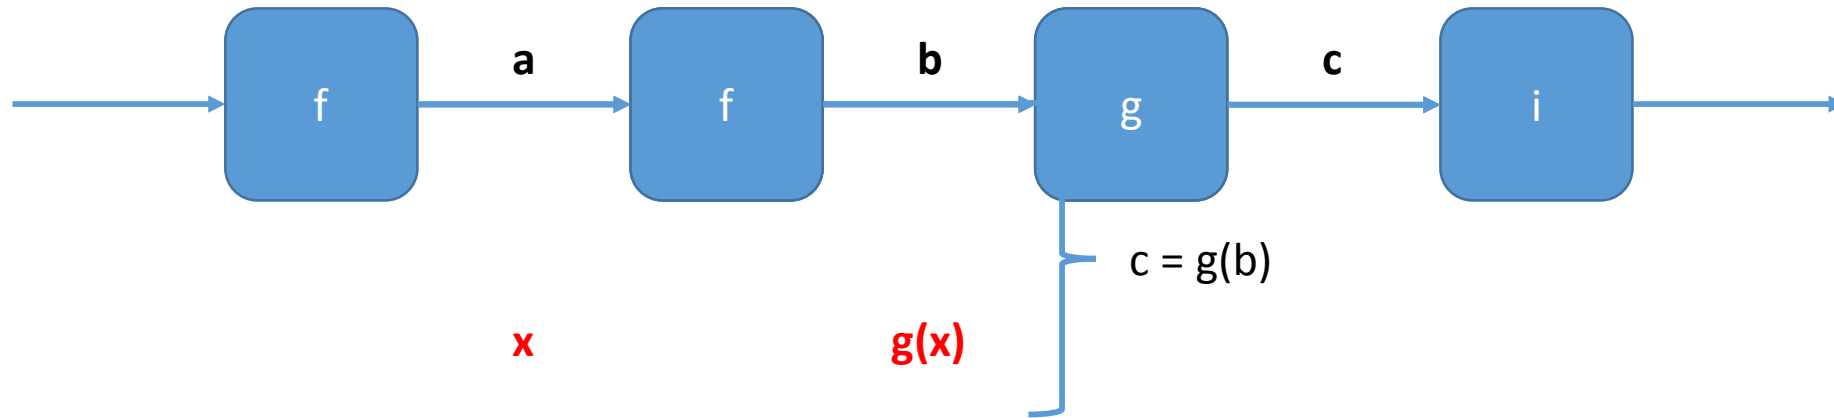
Motivating Problem

Grep terabytes in seconds

Sequentially reading terabytes from disk takes hours

Grep implementations are sequential

Breaking Sequential Dependences



Two Parallel Algorithms

Finite State Machines

30x faster on 12 cores

Dynamic Programming

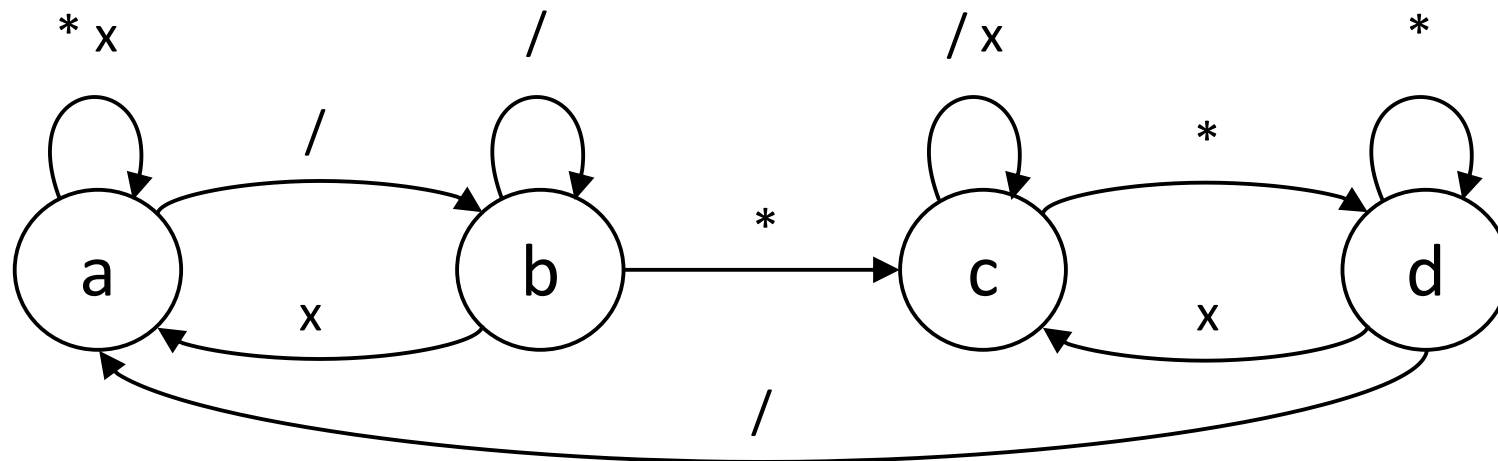
Fastest software Viterbi decoder

Higher throughput on 16 cores than a commercial FPGA implementation

Parallel Finite State Machines

FSM Applications

- grep (regex matching)
- lex (tokenization)
- Dictionary-based decoding (e.g. Huffman decoding)
- Text encoding/decoding (e.g. UTF8)



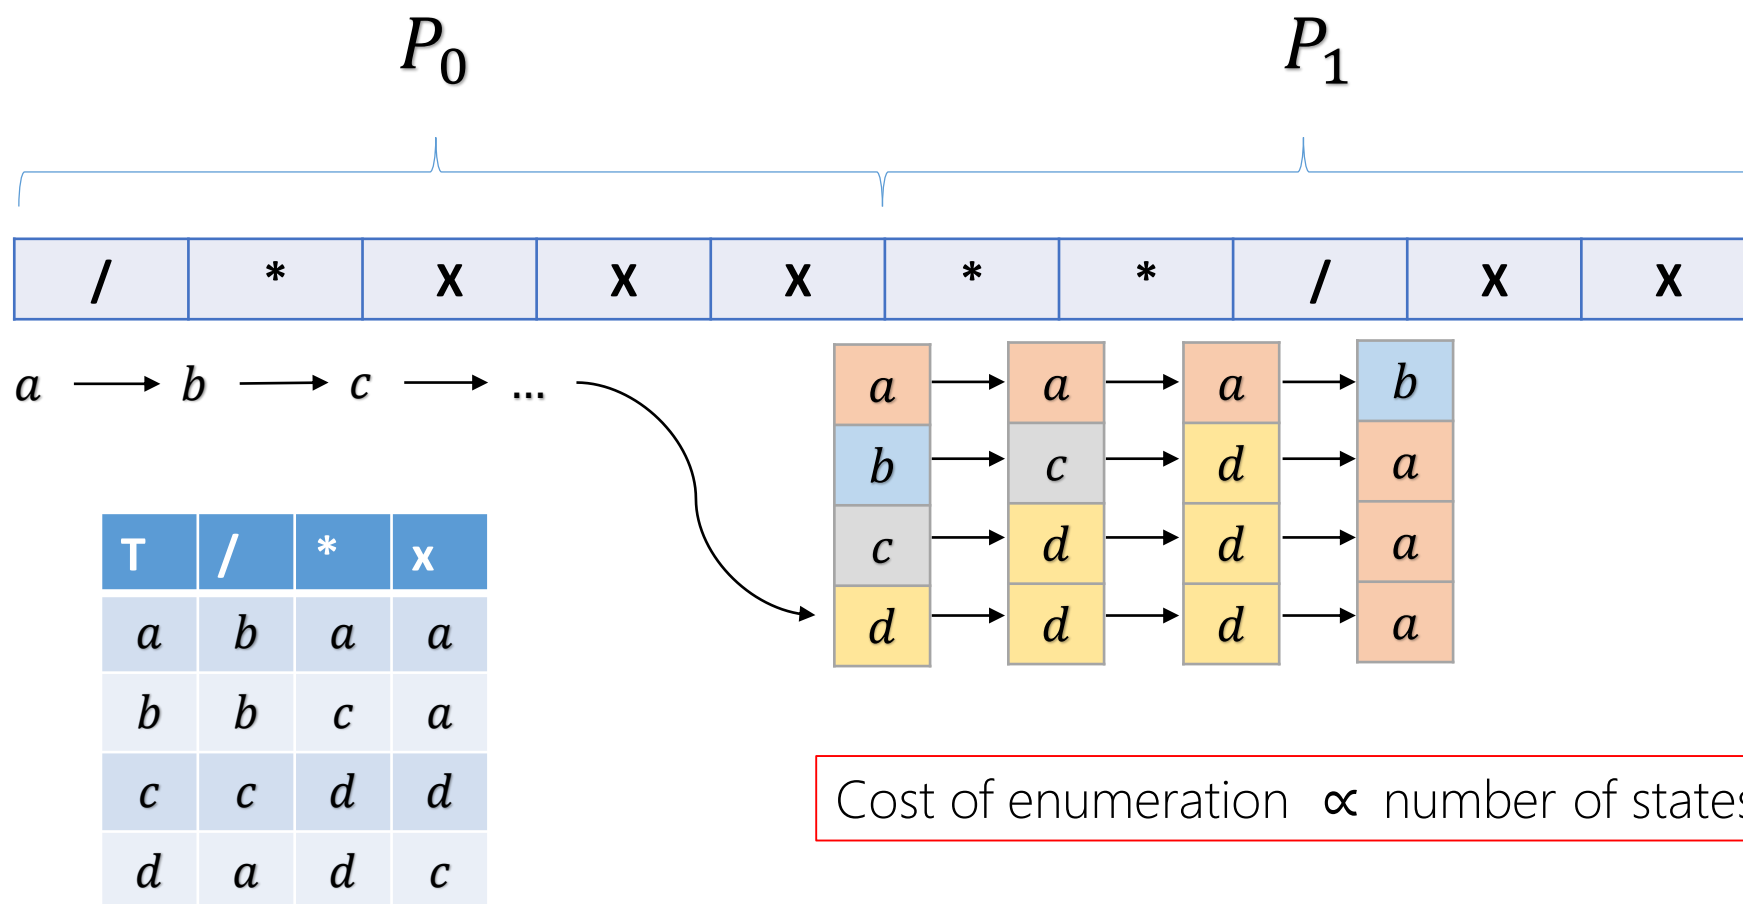
Data Dependence limits ILP, SIMD, and multicore parallelism

T	/	*	x
<i>a</i>	<i>b</i>	<i>a</i>	<i>a</i>
<i>b</i>	<i>b</i>	<i>c</i>	<i>a</i>
<i>c</i>	<i>c</i>	<i>d</i>	<i>d</i>
<i>d</i>	<i>a</i>	<i>d</i>	<i>c</i>

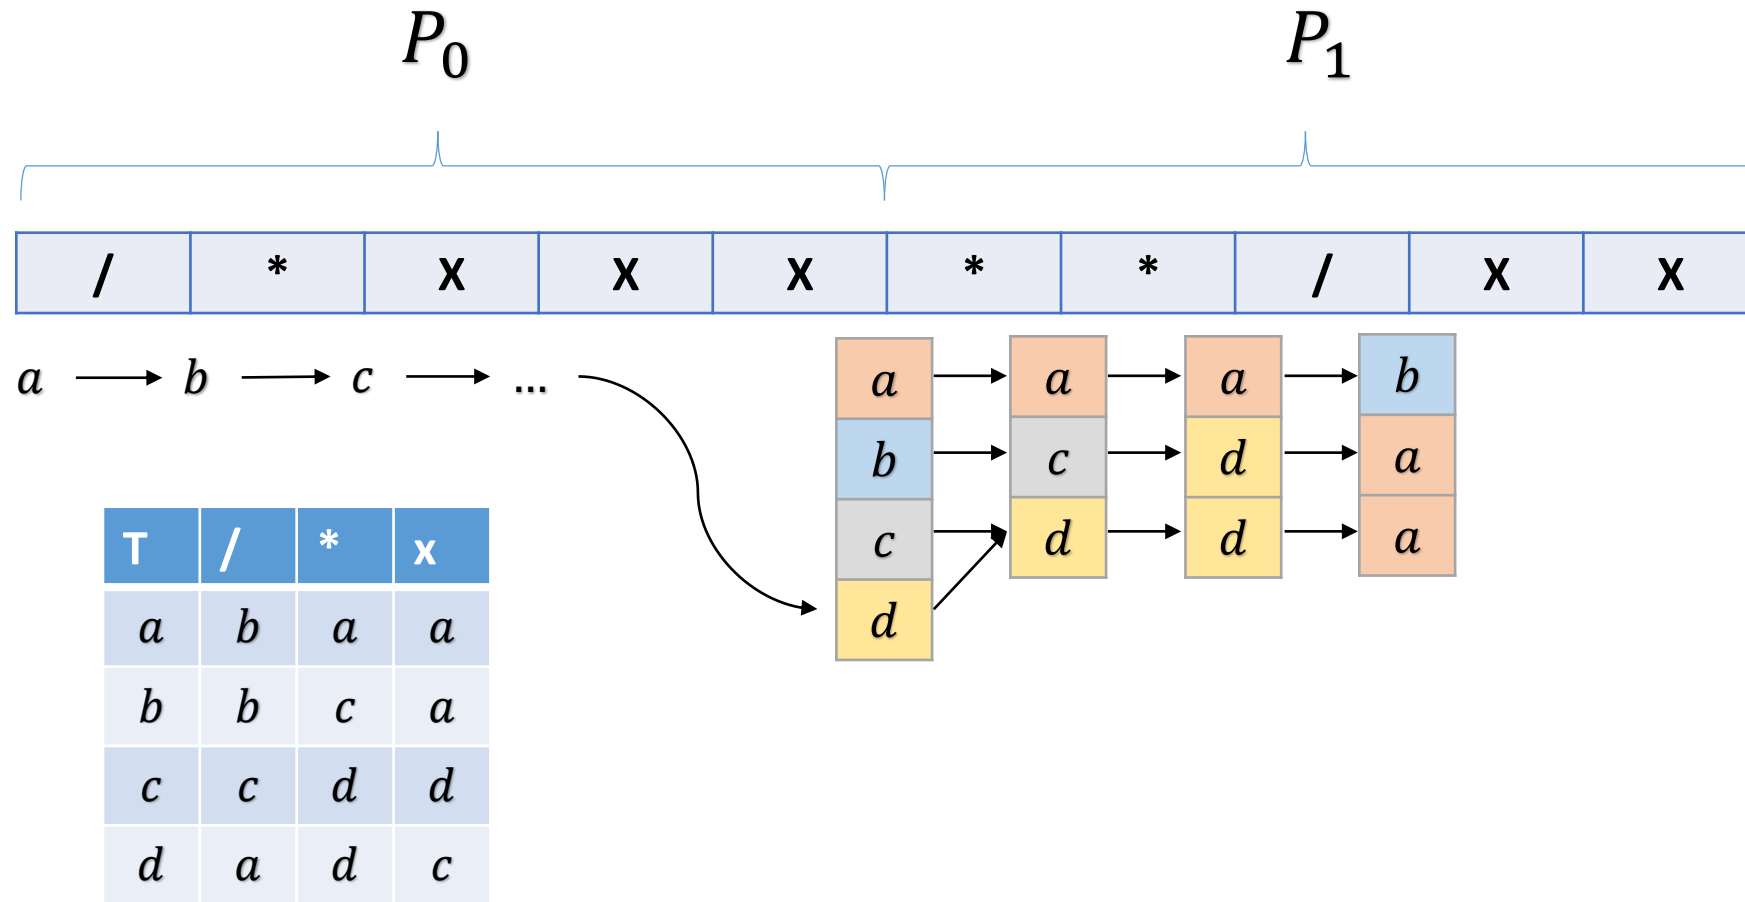
```

state = a;
foreach(input in)
  state = T[in][state];
  
```

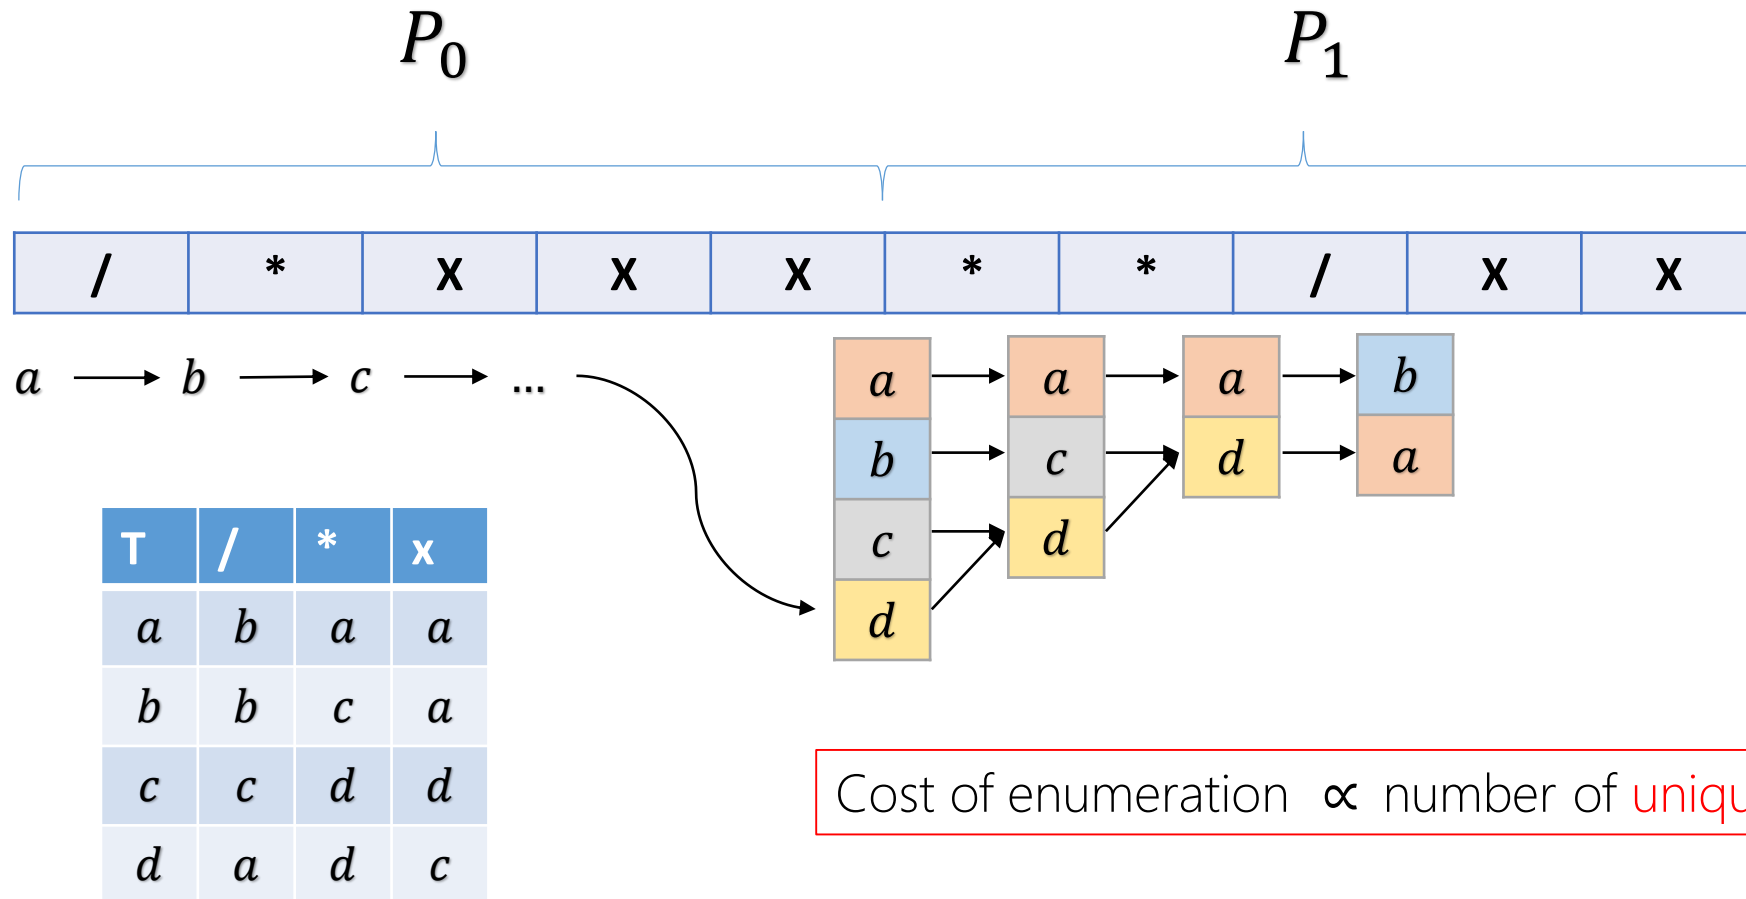
Breaking Dependences with Enumeration



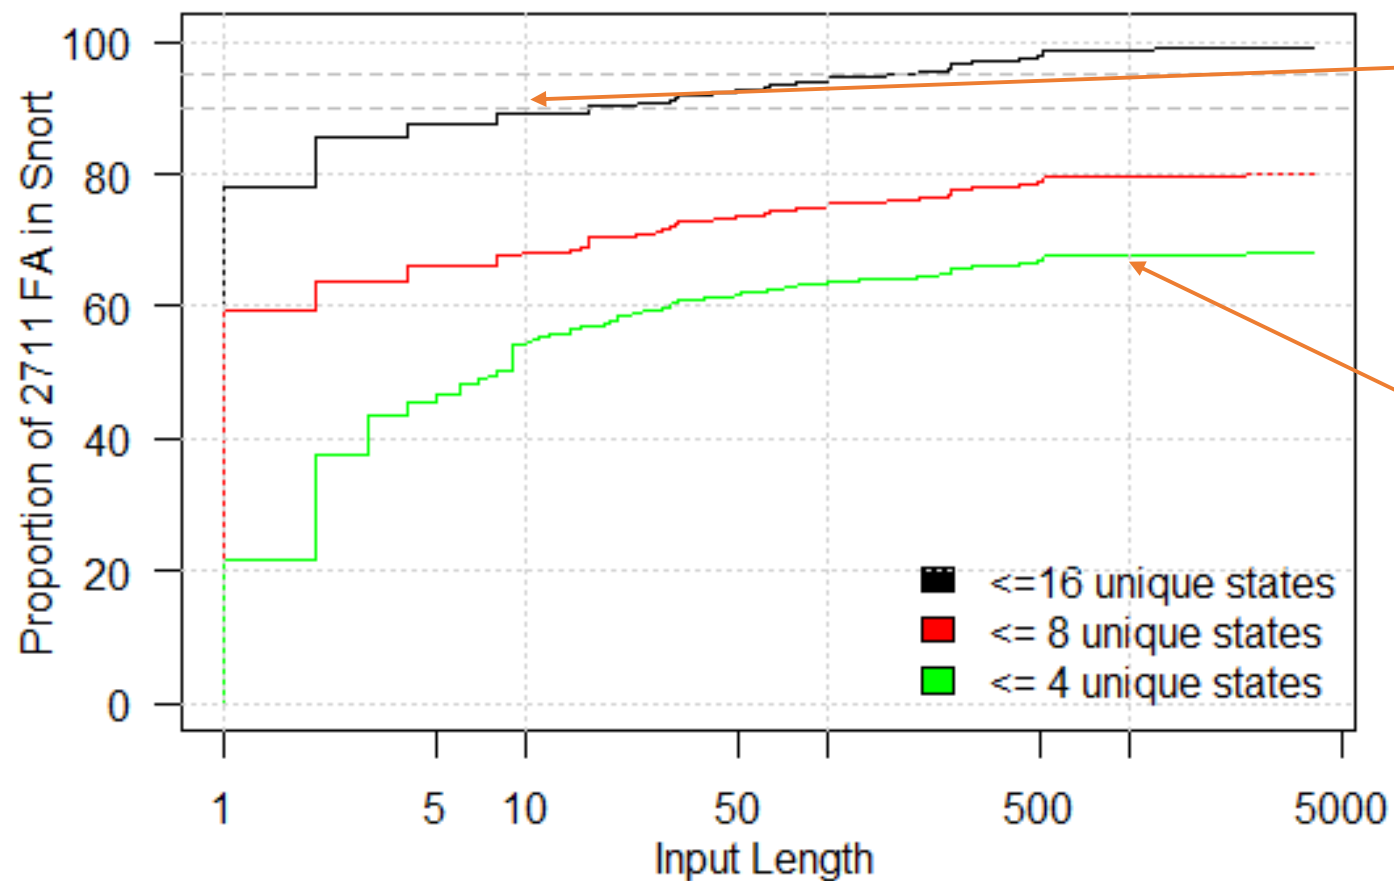
Convergence Reduces the Cost of Enumeration



Convergence Reduces the Cost of Enumeration



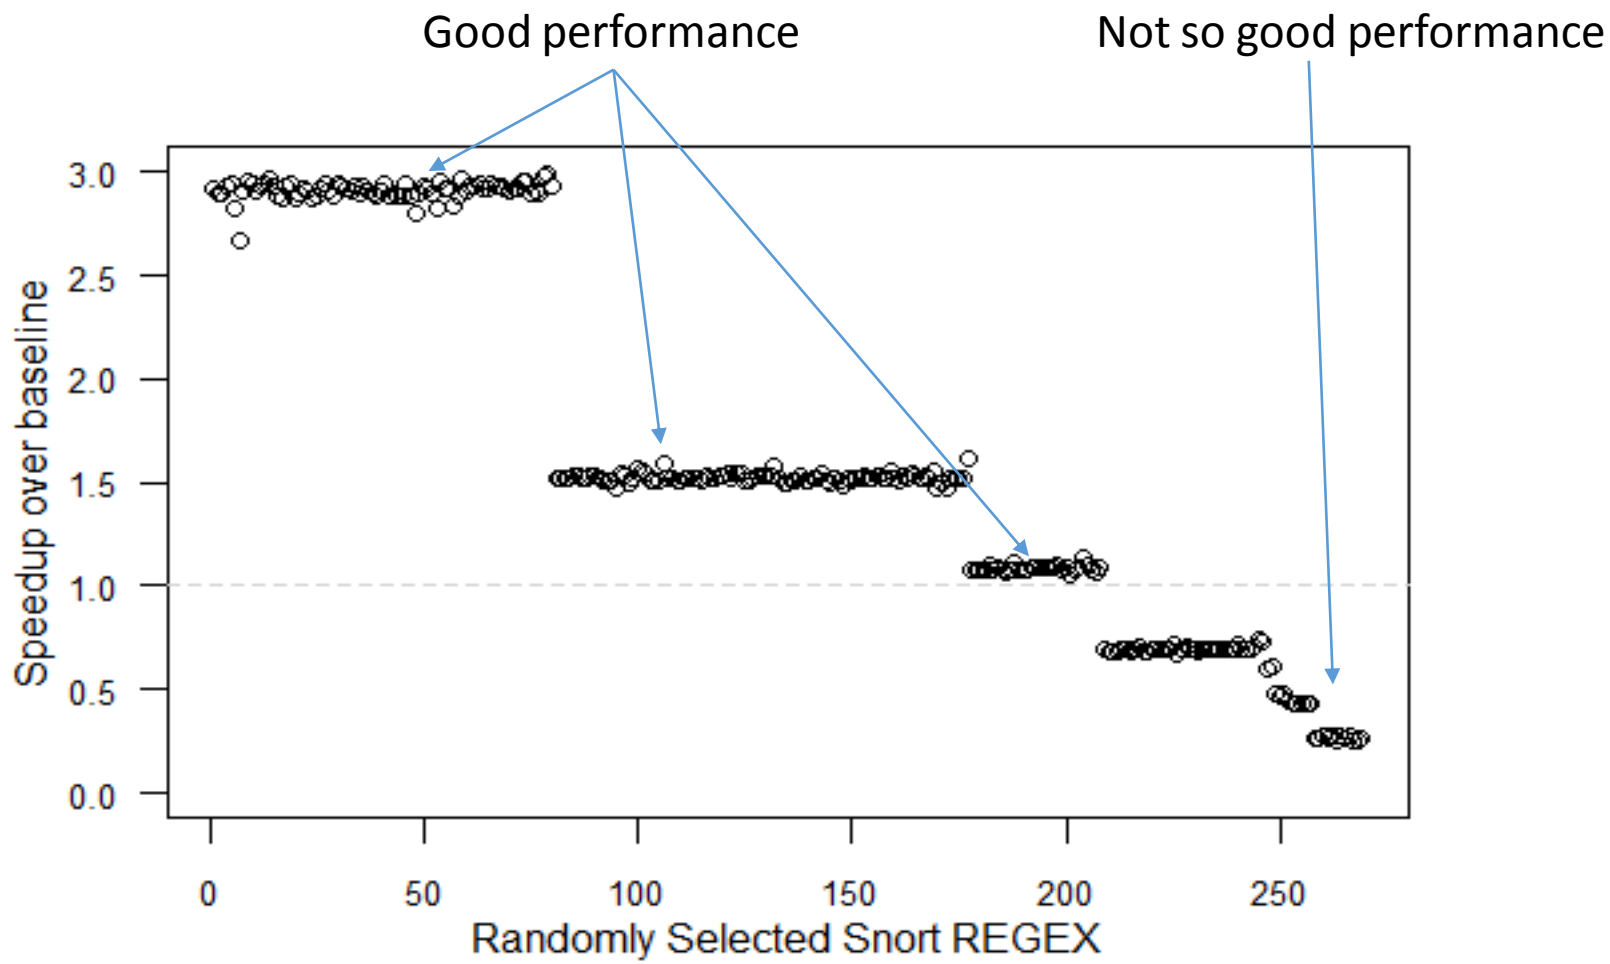
Convergence for Worst-Case Inputs



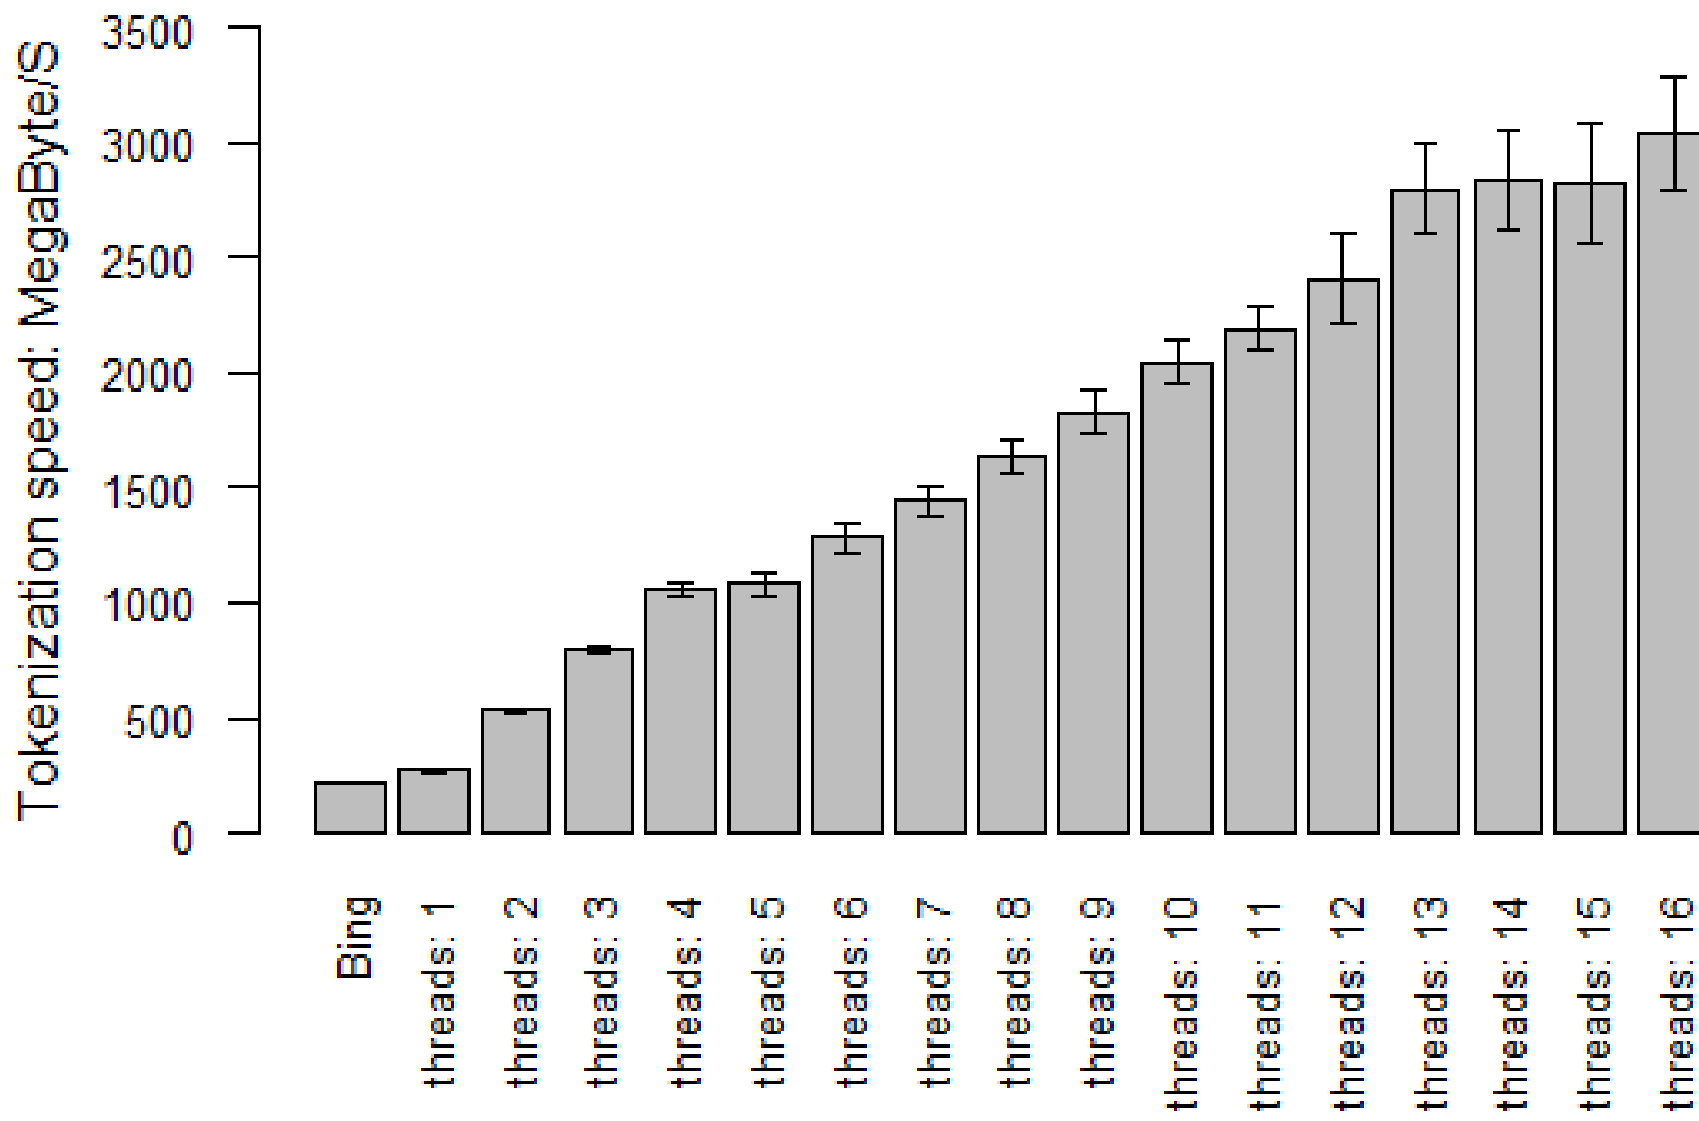
90% of FSMs converge to ≤ 16 states after 10 steps

Only 65% of FSMs converge to ≤ 4 states even after 5000 steps

Single-Core Performance when using SIMD

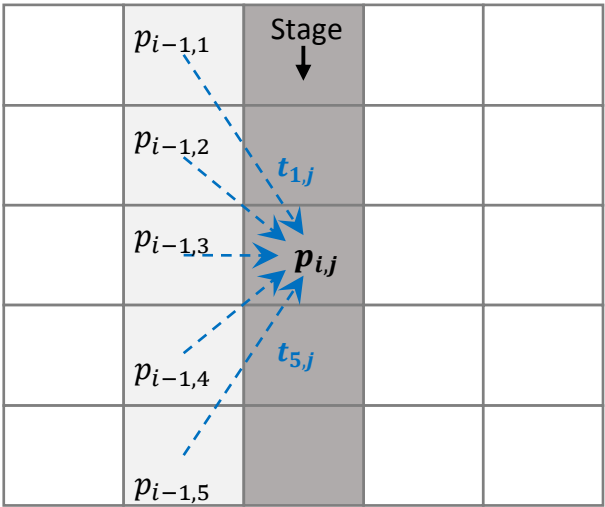


Multicore Performance for Bing Tokenization



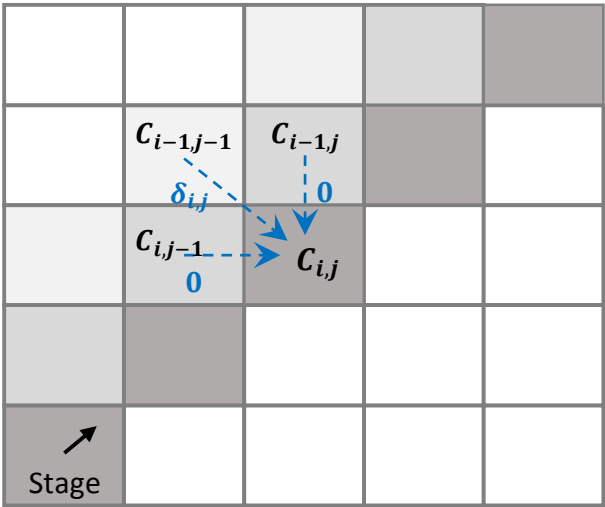
Parallelizing Dynamic Programming

Optimization Problems Solved using Dynamic Programming



$$p_{i,j} = \max_k (p_{i-1,k} * t_{k,j})$$

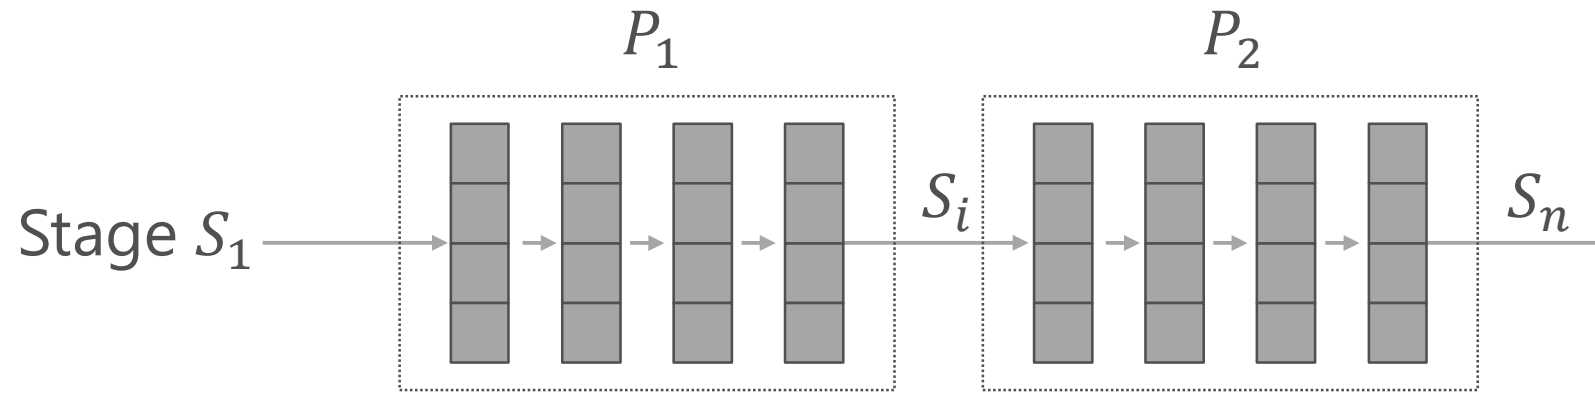
Viterbi



$$c_{i,j} = \max \begin{cases} c_{i-1,j-1} + \delta_{i,j} \\ c_{i,j-1} \\ c_{i-1,j} \end{cases}$$

diff

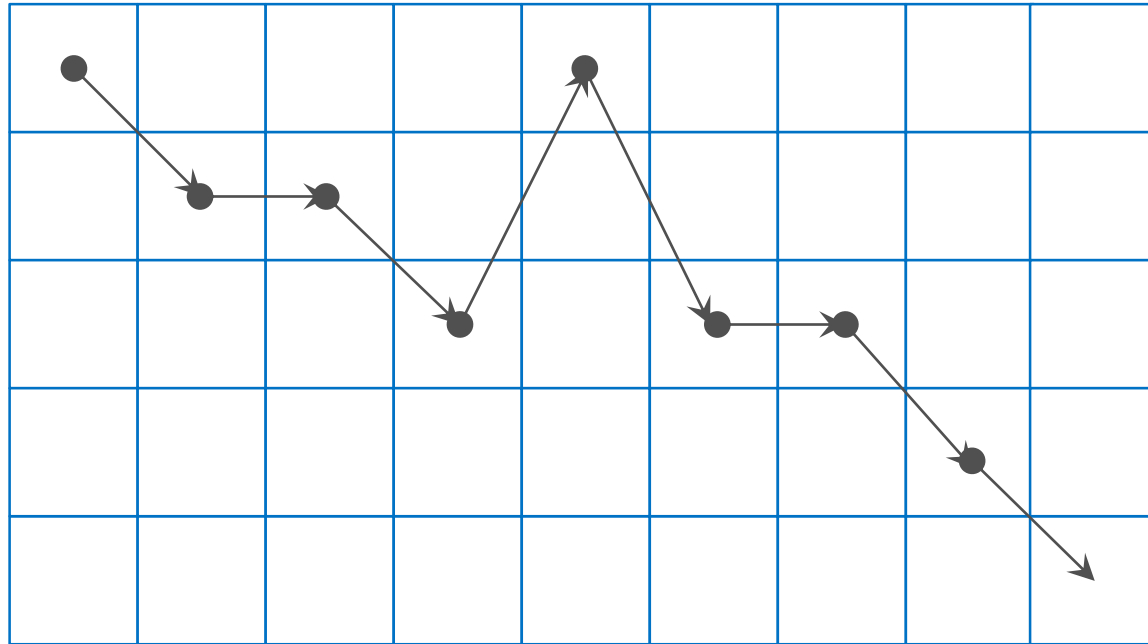
Parallelize Across Stages



Assume recurrence is of the form

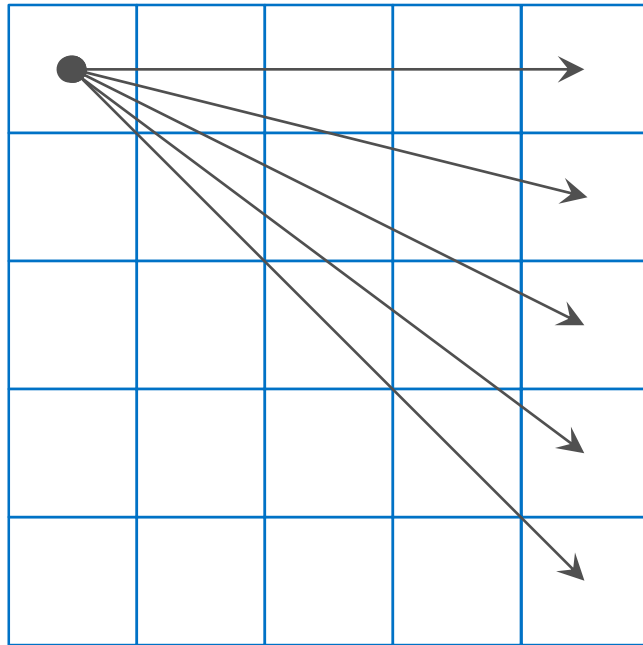
$$s_i[j] = \min_k (C_{ijk} + s_{i-1}[k])$$

Optimization problem =
Finding shortest path in some graph

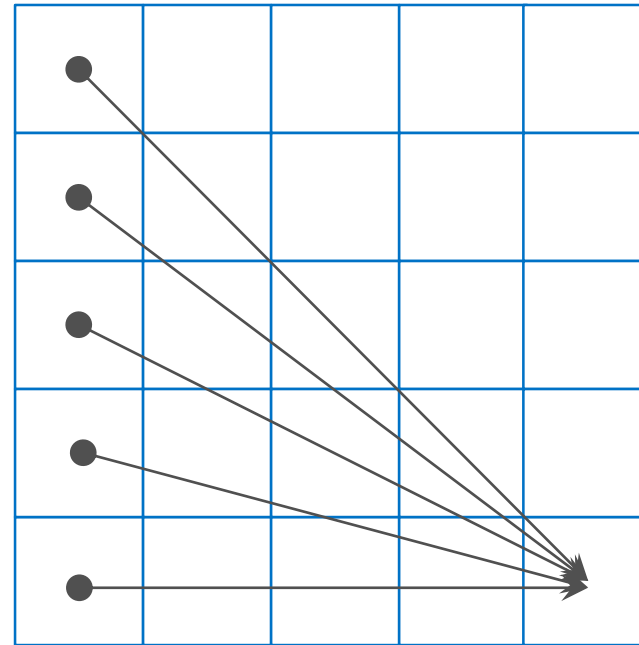


Break Dependences with All-Pair Shortest Paths

source to
all boundary nodes

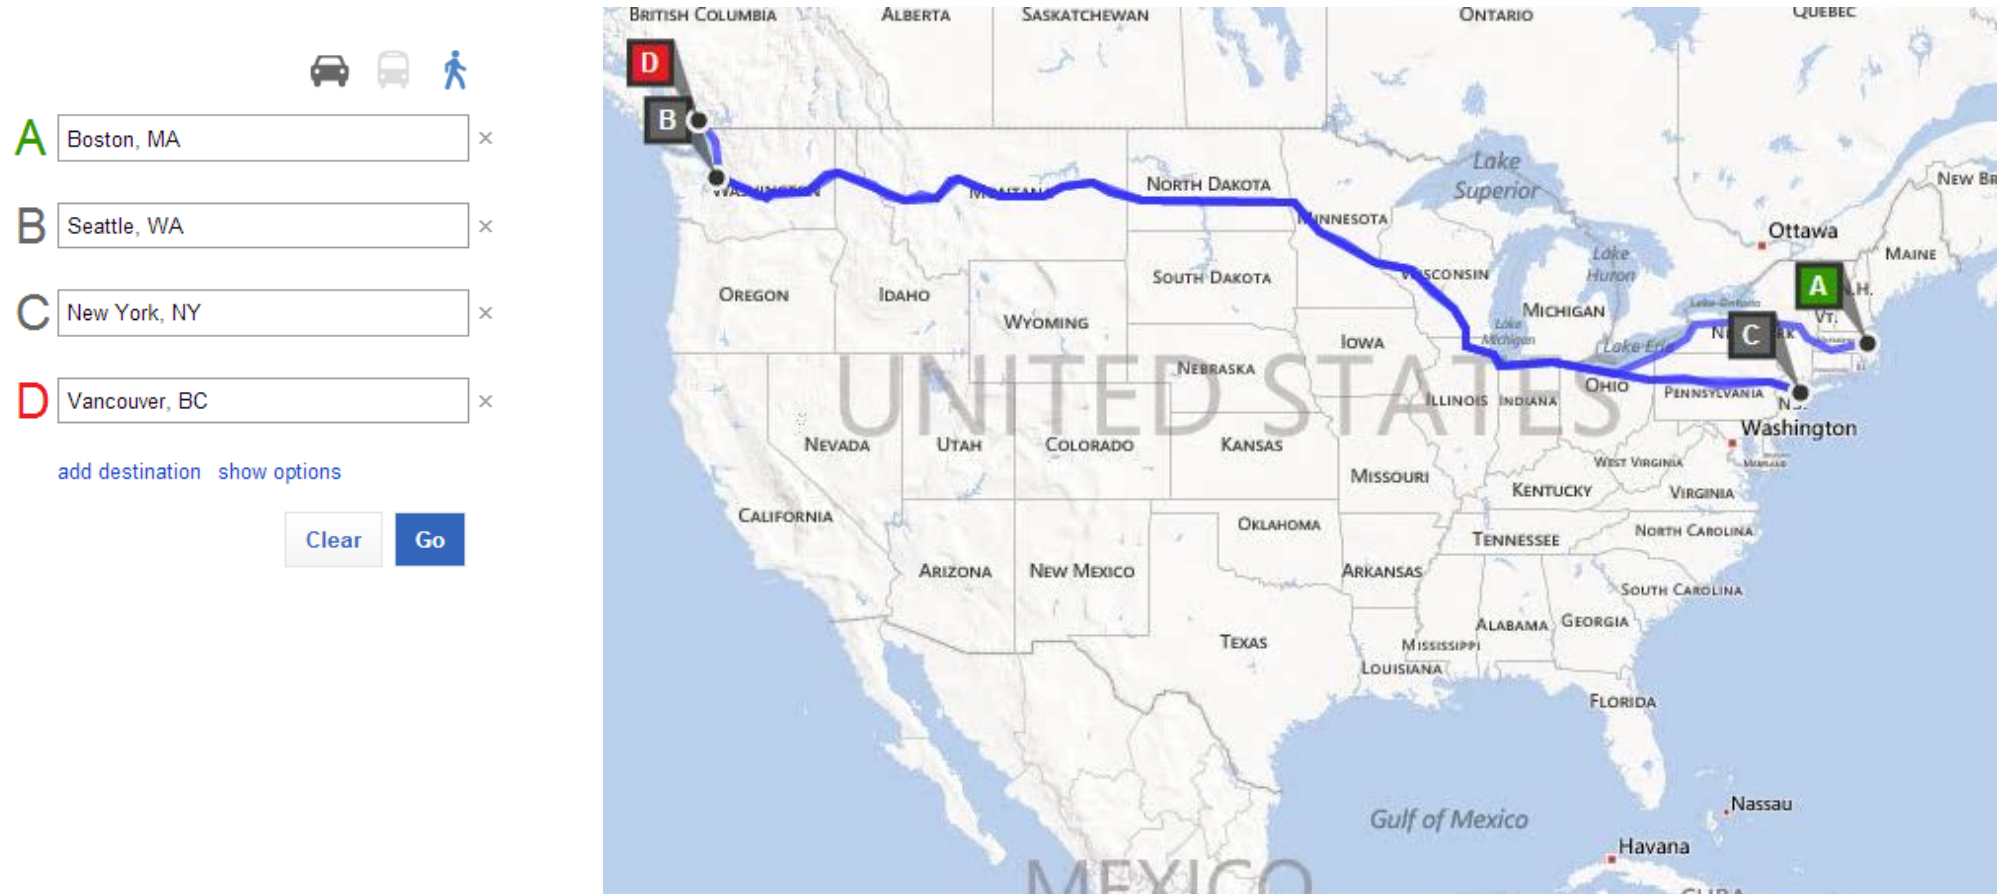


all boundary nodes
to destination



Overhead \propto stage size

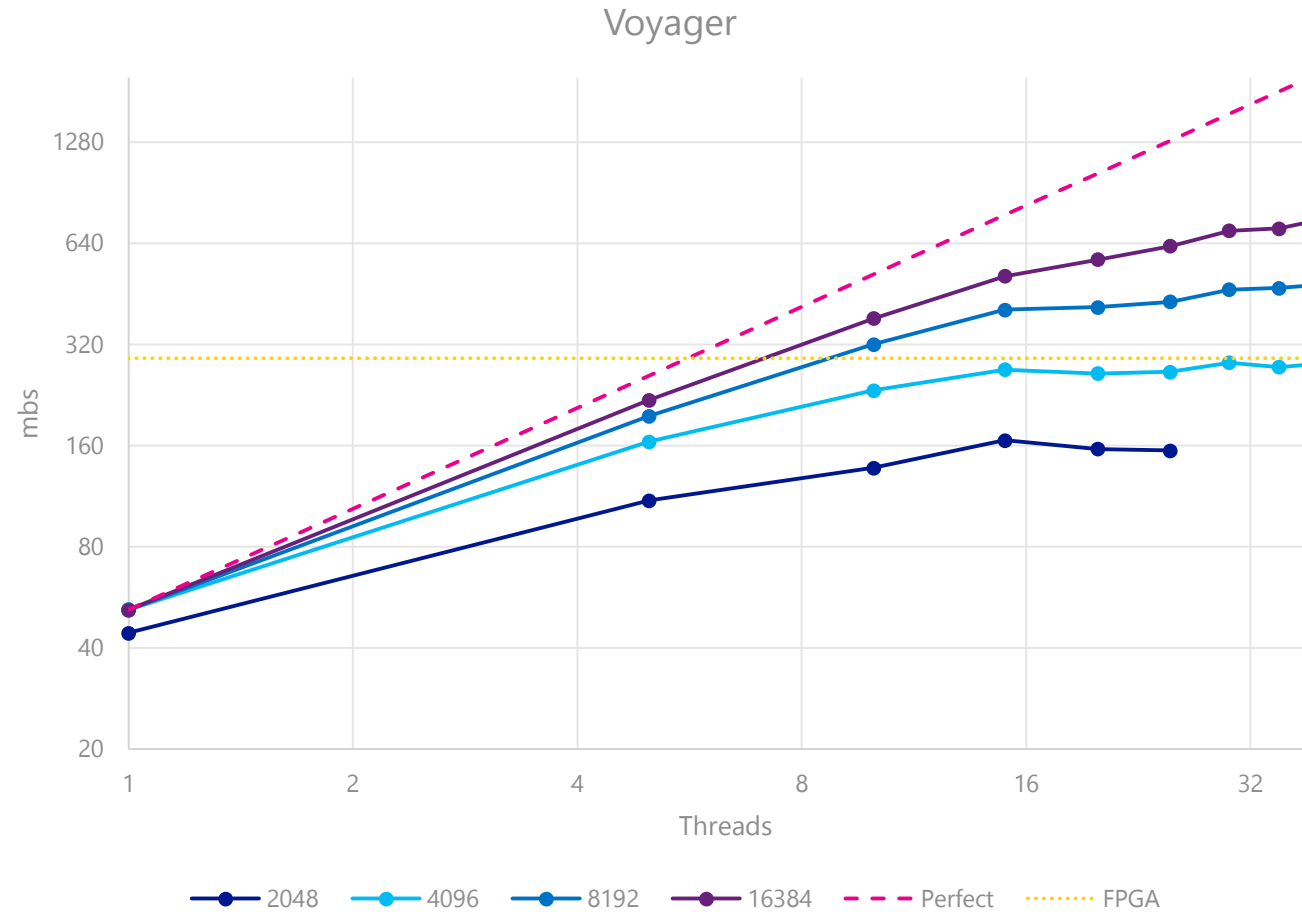
All-Pair Shortest Paths Converge



Convergence in LCS



Results – Viterbi Decoder



Conclusions

Parallel algorithms for FSMs and dynamic programming

Using dynamic convergence properties

Can we break dependences for other algorithms?

Parsing

Iterative machine learning

Graph algorithms

Can we automatically parallelize across dependences?



Save the planet and return
your name badge before you
leave (on Tuesday)

