

ANALYSIS AND OPTIMIZATION FOR PROCESSING GRID-SCALE XML DATASETS

Michael R. Head
Ph.D. Candidate

Grid Computing Research Laboratory
Department of Computer Science
Binghamton University
mike@cs.binghamton.edu

Tuesday, May 12, 2009

OUTLINE

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- PixIMAL: Parallel Approach for Processing XML

4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

```
<?xml version="1.0" encoding="UTF-8"?>
<ns1:MoleculeType xsd:type="ns1:MoleculeType"
    xmlns:ns1="http://nbcr.sdsc.edu/chemistry/types"
    xmlns:xsd="http://www.w3.org/2001/XMLSchema">
<moleculeName xsi:type="xsd:string"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
    1kzk
</moleculeName>
<moleculeRadius xsi:type="xsd:double" xsi:nil="true"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"/>
<atom xsi:type="ns1:AtomType"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
    <fieldName xsi:type="ns1:FieldNameType">ATOM</fieldName>
    ...
</atom>
<atom xsi:type="ns1:AtomType"
    ...
</atom>
...
</ns1:MoleculeType>
```

OUTLINE

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- Piximal: Parallel Approach for Processing XML

4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

XML DEFINED

- Text based (usually UTF-8 encoded)
- Tree structured
- Language independent
- Generalized data format

MOTIVATION FROM SOAP

- Generalized RPC mechanism (supports other models, too)
- Broad industrial support
- Web Services on the Grid
 - OGSA: Open Grid Services Architecture
 - WSRF: Web Services Resource Framework
- At bottom, SOAP depends on XML

IMPORTANCE OF HIGH PERFORMANCE XML PROCESSORS

- Becoming standard for many scientific datasets
 - HapMap - mapping genes
 - Protein Sequencing
 - NASA astronomical data
 - Many more instances

EXPLOSION OF DATA

- Enormous increase in data from sensors, satellites, experiments, and simulations*
- Use of XML to store these data is also on the rise
- XML is in use in ways it was never really intended (GB and large size files)

BENCHMARK MOTIVATION

- Scientific applications place a wide range of requirements on the communication substrate and data formats.
- Simple and straightforward implementations can have a severe performance impact.

OUTLINE

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- Piximal: Parallel Approach for Processing XML

4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

PREVALENCE OF PARALLEL MACHINES

- All new high end and mid range CPUs for desktop- and laptop-class computers have at least two cores
- The future of AMD and Intel performance lies in increases in the number of cores
- Despite extant SMP machines, many classes of software applications remain single threaded
 - Multi-threaded programming considered “hard”

XML AND MULTI-CORE

- Most string parsing techniques rely on a serial scanning process
- **Challenge:** Existing (singly-threaded) XML parsers are already very efficient (Zhang et al 2006)

OUTLINE

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- Piximal: Parallel Approach for Processing XML

4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

CONTRIBUTIONS

- We present the design and implementation of a comprehensive benchmark suite for XML and SOAP implementations with standard mechanisms to quantify, compare, and evaluate the performance of each toolkit and study the strengths and weaknesses for a wide range of use case scenarios.
- We present an analysis of pre-fetching and piped implementation techniques that aim to offset disk I/O costs while processing large-scale XML datasets on multi-core CPU architectures.

CONTRIBUTIONS CONTINUED

- We propose techniques to modify the lexical analysis phase for processing large-scale XML datasets to leverage opportunities for parallelism. (PiXIMAL)
- We present an analysis of the scalability that can be achieved with our proposed parallelization approach as the number of processing threads and size of XML-data is increased.
- We present an analysis on the usage of various *states* in the processing automaton to provide insights on why the performance varies for differently shaped input data files.

PUBLICATIONS

- ``A Benchmark Suite for SOAP-based Communication in Grid Web Services,’’ in *The Proceedings of Supercomputing 2005*
- ``Benchmarking XML Processors for Applications in Grid Web Services,’’ in *The Proceedings of Supercomputing 2006*
- ``Approaching a Parallelized XML Parser Optimized for Multi-Core Processors,’’ in *The Proceedings of SOCP 2007*, workshop held in conjunction with HPDC 2007
- ``Parallel Processing of Large-Scale XML-Based Application Documents on Multi-core Architectures with PiXiMaL,’’ in *The Proceedings e-Science 2008*
- ``Performance Enhancement with Speculative Execution Based Parallelism for Processing Large-scale XML-based Application Data,’’ to appear in *The Proceedings of HPDC 2009*

THESIS STATEMENT

In this thesis we present a comprehensive benchmark suite that facilitates the study of the strengths and weaknesses of XML and SOAP toolkits for a wide range of use case scenarios.

We propose a parallel processing model for some application-based large-scale XML datasets that can effectively leverage opportunities for parallelism in emerging multi-core CPU architectures.

OUTLINE

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- PIXIMAL: Parallel Approach for Processing XML

4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

SOAP BENCHMARK SUITE

- Defines a set of operations to implement within a SOAP toolkit
- Tests both serialization and deserialization of a variety of data structures over a range of input sizes
 - Simple types: integers, strings, and floats
 - Base64 encoded data
 - Complex types: event streams, mesh interface objects

OUTLINE

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- PIXIMAL: Parallel Approach for Processing XML

4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

XML BENCHMARK SUITE

- ➊ A chosen set of XML documents
 - Low level probes
 - Application-based benchmarks
- ➋ A driver application for each XML processor
 - Runs the parser on the input, but does not act on the data
 - Eliminates application-level performance differences
 - One for each interface style (SAX/DOM)

OUTLINE

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- PIXIMAL: Parallel Approach for Processing XML

4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

READAHEAD/RUNAHEAD

- Explore OS level caching effects
- Offload disk input to another thread/core
- Improved the performance of an existing high performance parser by using a separate thread to read the input into cache

OUTLINE

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- **PIXIMAL: Parallel Approach for Processing XML**

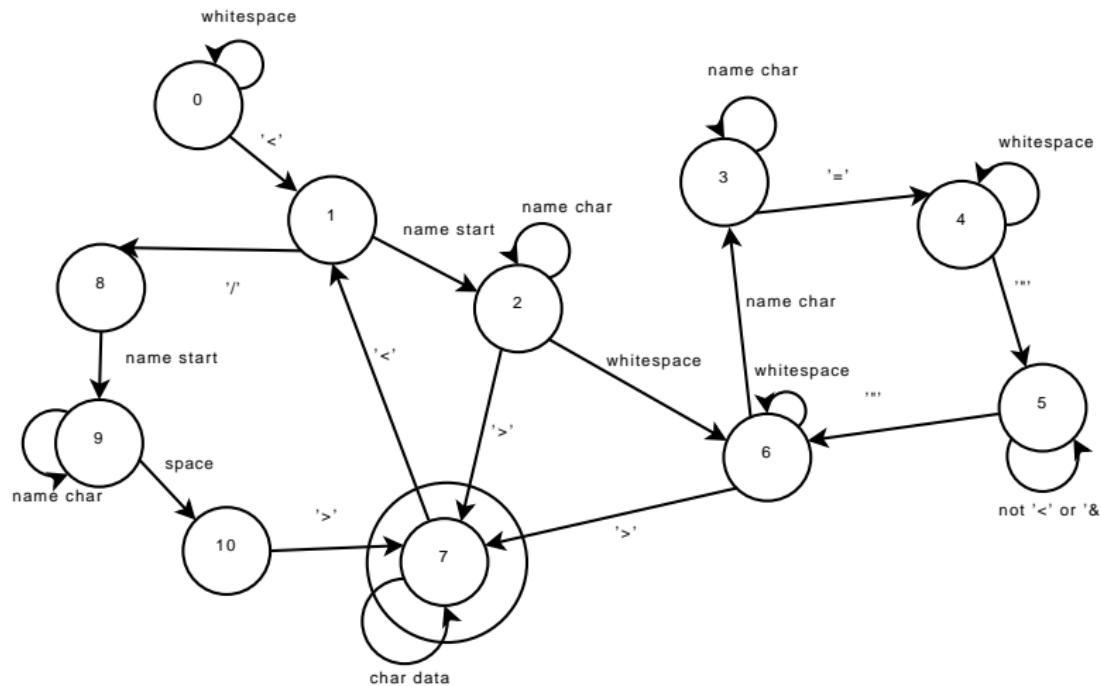
4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

TOKEN-SCANNING WITH A DFA

- DFA-based table-driven scanning is both popular and fast
 - (or at least performance-competitive with other techniques)
- Input is read *sequentially* from start to finish
 - Each character is used to transition over states in a DFA
 - Transition may have associated actions
 - Supports languages that are not “regular”
- Commonly used in high performance XML parsers, such as TDX (C) and Piccolo (Java)
 - Amenable to SAX parsing
 - PIXIMAL-DFA uses this approach

DFA USED IN PIXIMAL-DFA



PARALLEL SCANNING WITH A DFA?

- DFA-based scanning \implies sequential operation
- Desire: run multiple, concurrent DFAs throughout the input
 - Generally not possible because the start state would be unknown

OVERCOMING SEQUENTIALITY WITH AN NFA

- Problem: start state is unknown
- Solution: assume every possible state is a start state
 - Construct an NFA from the DFA used in **PIXIMAL-DFA**
 - Such an NFA can be applied on any substring of the input
- **PIXIMAL-NFA** is the parser that does all of this:
 - Partition input into segments
 - Run **PIXIMAL-DFA** on the initial segment
 - Run NFA-based parsers on subsequent partition elements
 - Fix up transitions at partition boundaries and run queued actions

PIXIMAL-NFA's PARAMETERS

- *split_percent*:
 - The portion of input to be dedicated to the first element of the partition, expressed as a percentage of the total input length
- *number_of_threads*:
 - The number of threads to use on a run

PRELIMINARY RESEARCH QUESTIONS

- Is there enough memory bandwidth to allow multiple automata to concurrently feed each thread its input?
- Processing each character along several paths through the NFA is costly: how does this work scale with the size of the initial DFA?
 - (E-science 2008)
- Does the overhead of queuing the NFA actions cost an acceptable amount compared with the cost of DFA-parsing the first partition element?
 - (HPDC 2009)

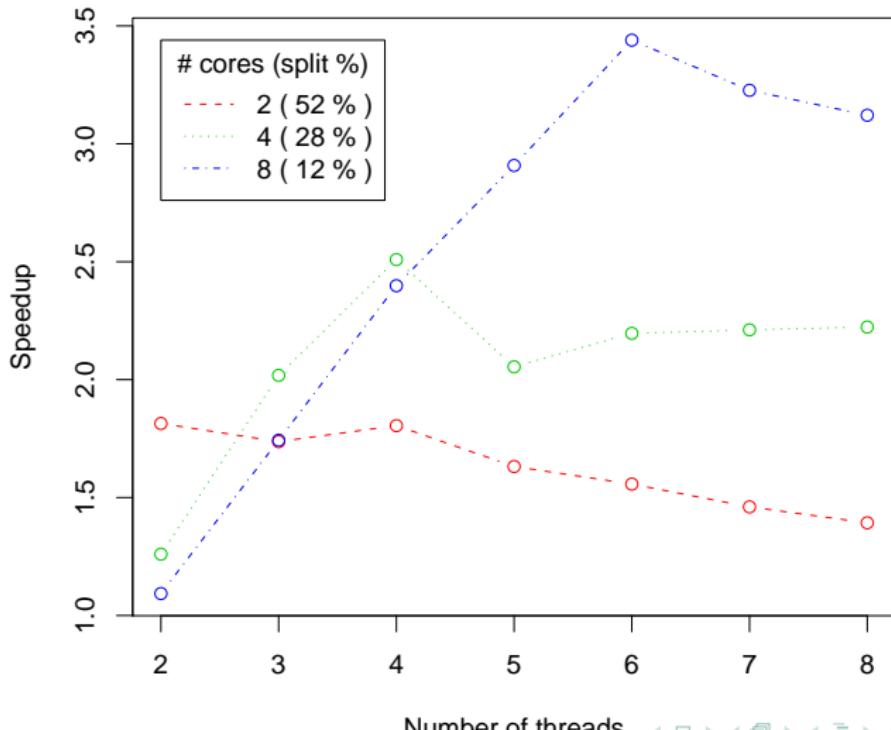
MEMORY BANDWIDTH TEST

- Models the work of partitioning the input the way PIXIMAL-NFA does
 - File I/O is via `mmap` (2)
- A thread is created for each partition element which accumulates each character
- A variety of `split_percents` and `number_of_thread` are chosen
 - Total time to read a large input a fixed number of times is measured
 - Input file is `SwissProt.xml`, which is 109 MB in size

MEMORY BANDWIDTH TEST – EXPERIMENTAL SETUP

- Run several machines, each from a homogeneous class running 64-bit versions of Linux
 - 2x uniprocessor: 3.2 Ghz Intel Xeon (uniprocessor), 4 GB RAM, Linux kernel 2.6.15, GNU Lib C 2.3.6, GCC 4.0.3
 - 2x dual core: 2.66 Ghz Intel Xeon 5150 (dual core) CPUs, 8 GB RAM, Linux kernel 2.6.18, GNU Lib C 2.3.6, GCC 4.1.2
 - 2x quad core: 2.33 Ghz Intel Xeon E5354 (quad-core) CPUs, 8 GB RAM, Linux kernel 2.6.18, GNU Lib C 2.3.6, GCC 4.1.2
- 4 nodes used from the 2x UP cluster, 10 from each of the other two
- Results for each class are averaged across all runs

BANDWIDTH IS NOT A BOTTLENECK UP TO 6 CORES



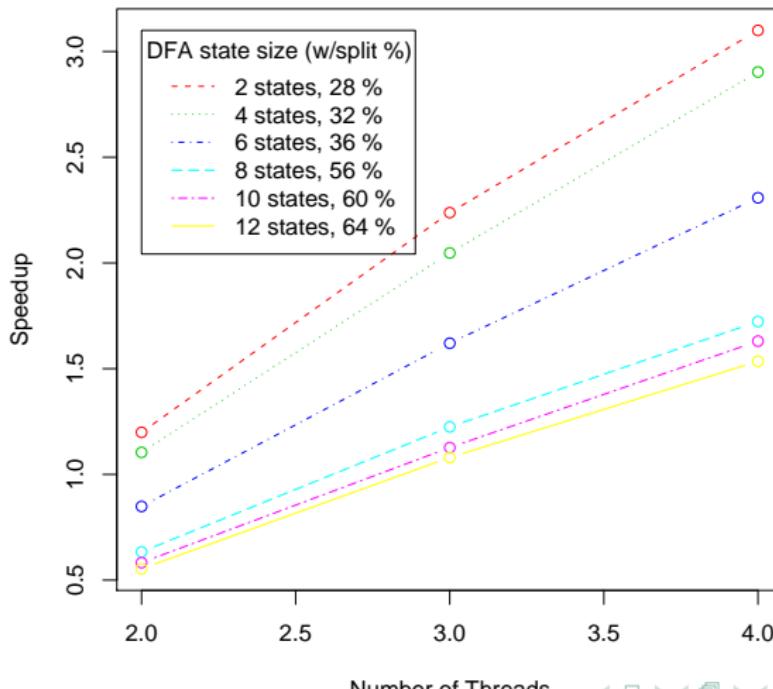
CONCLUSIONS FROM MEMORY BANDWIDTH TESTS

- Even when doing very little per-character processing, performance gains possible by adding threads
- Returns do diminish rapidly
- More cores lead to smoother results

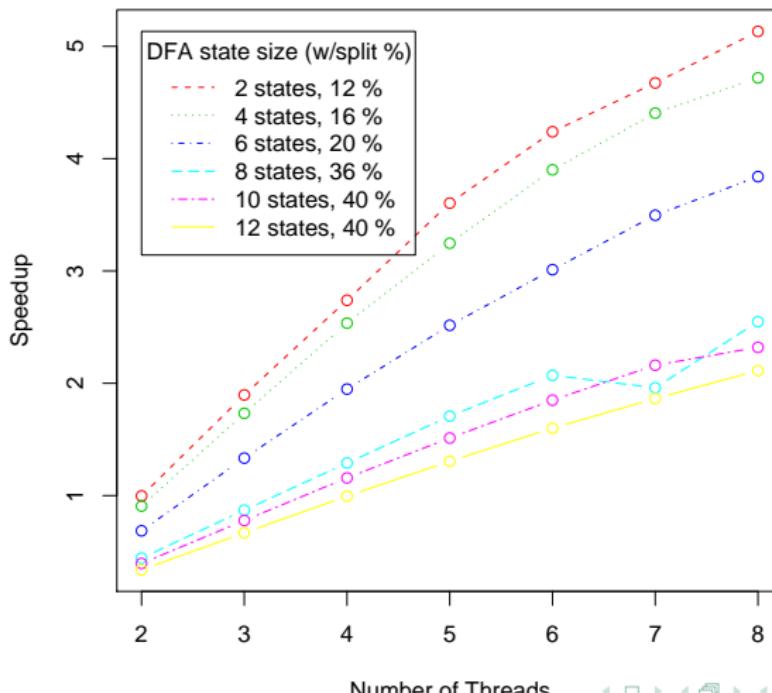
STATE SCALABILITY TEST

- Models the additional work done by the NFA threads by following multiple execution paths through the table
- Each NFA thread now must remember the state and calculate the next state for each character and for each start state
 - The DFA need only remember and calculate one state per input character
- Does not model the memory used, actions stored, or garbage state elimination
- Goal: to find a balance point for DFA size
 - + increased complexity of the recognized language
 - - more work for the NFA to do, more space required for table

2× DC



2× QC - BEST SPEEDUP FOR DFA SIZES



CONCLUSIONS FROM STATE SCALABILITY TEST

- The extra work of pushing characters through the multiple execution paths of the NFA is not in itself a limiting factor
- There is a “sweet spot” for DFA size: around 6-7 states which allows for the greatest language complexity and the best scalability
 - This is a crossover point where the $O(N)$ extra NFA work overcomes the $O(1)$ work of simply reading the input

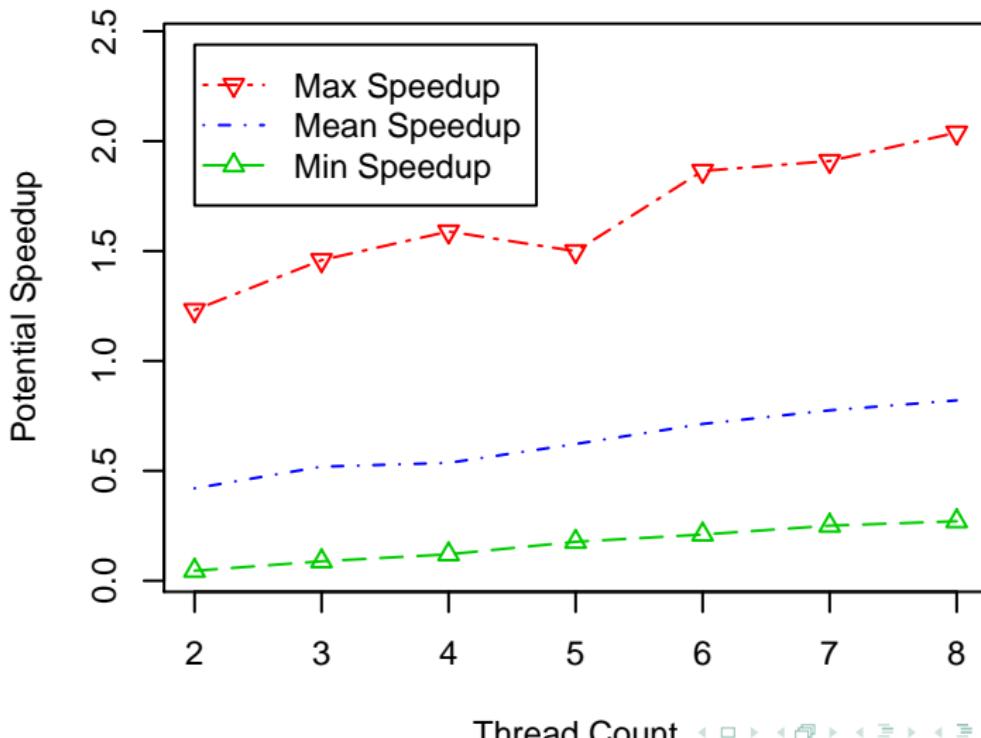
SERIAL NFA TESTS

- Test hypothesis: the extra work required by using an NFA is offset by dividing processing work across multiple threads
- Run each automaton-parser sequentially and independently
- Divide the work as usual, with a range of *split_percents* and *number_of_threads*
- Time each component independently
- Completely parses the input, generating the correct sequence of SAX events
- The maximum time for all components to complete (plus fix up time) represents an upper bound on the time PIXIMAL-NFA would take with components running concurrently

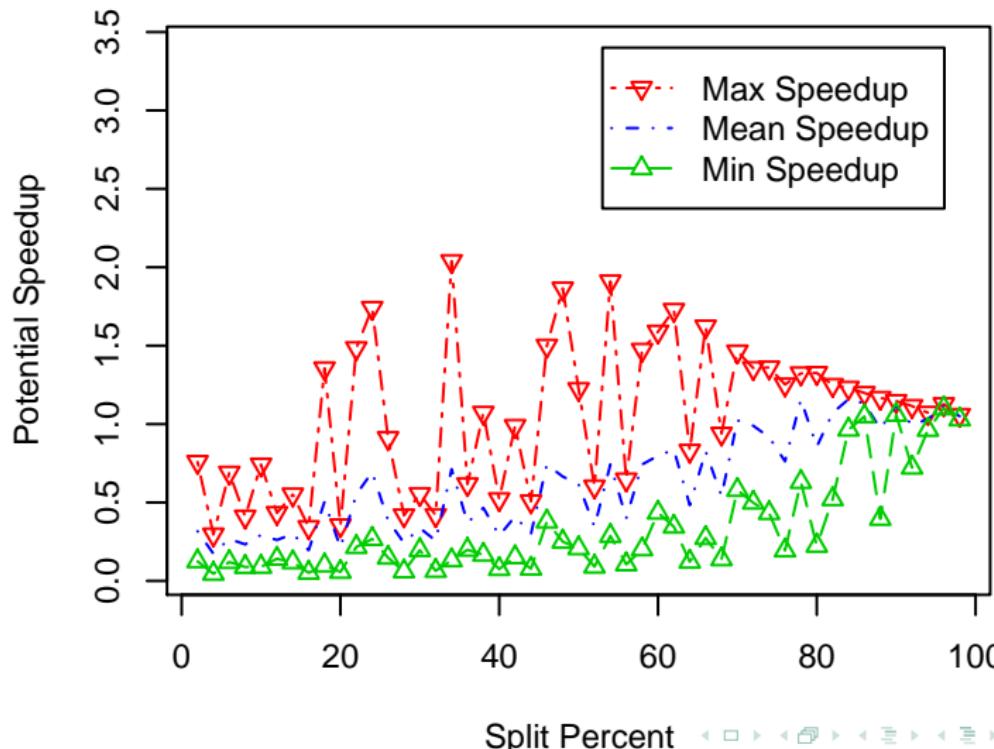
DIFFERENCES FROM PREVIOUS TESTS

- Entirely sequential (no concurrency)
- Full XML parsing takes place
- Input file is different
 - “Interop” test from SOAPBench and XMLBench
 - SOAP-encoded arrays of various data types: integers, strings, and MIOs
 - Array size is scaled between 10 and 50,000 elements for each type

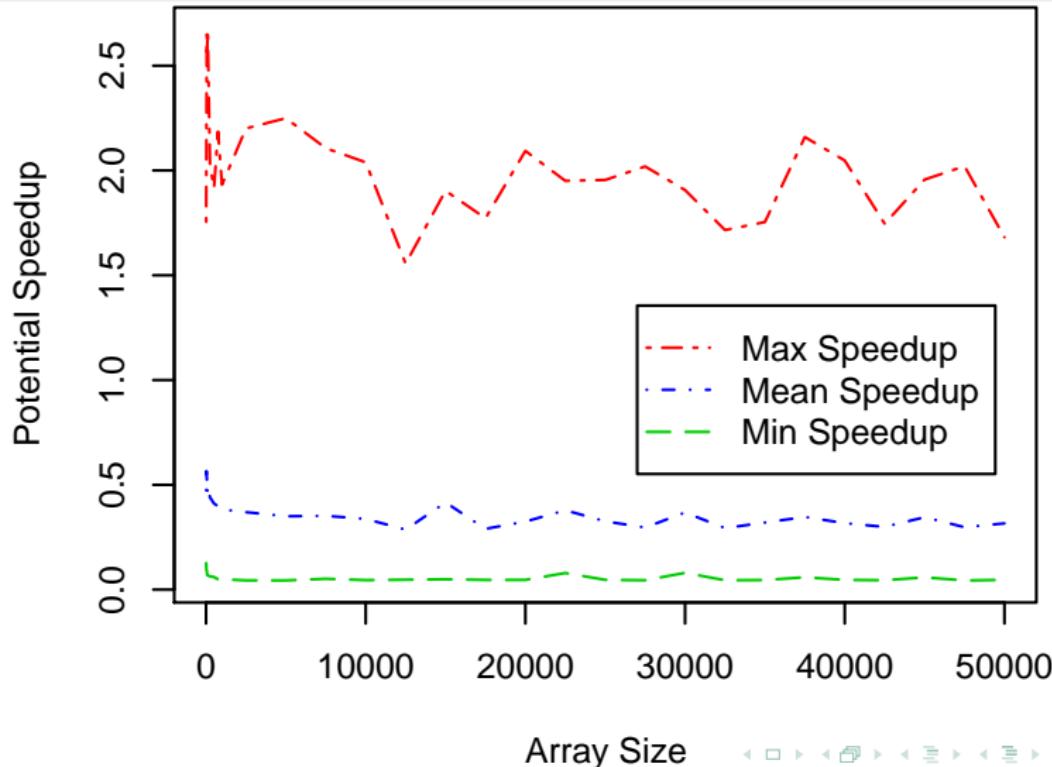
MODEST SPEEDUP SCALABILITY FOR 10,000 INTEGERS



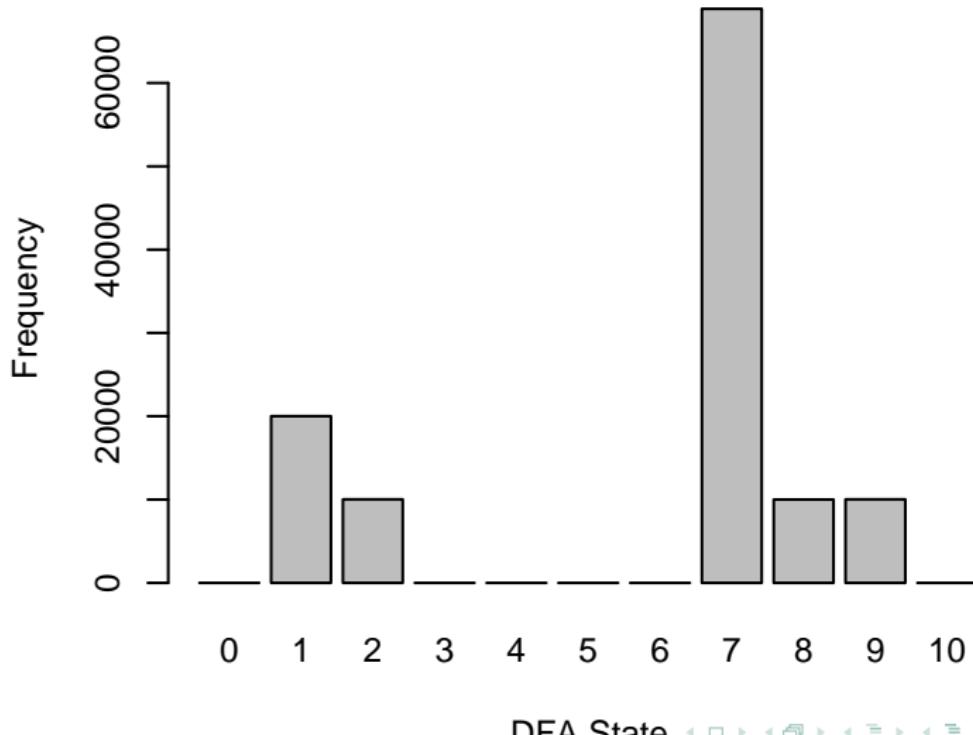
Split_Percent CRITICAL FOR SPEEDUP FOR 10,000 INTEGERS



INCONSISTENT SPEEDUP OVER A RANGE OF ARRAY LENGTHS



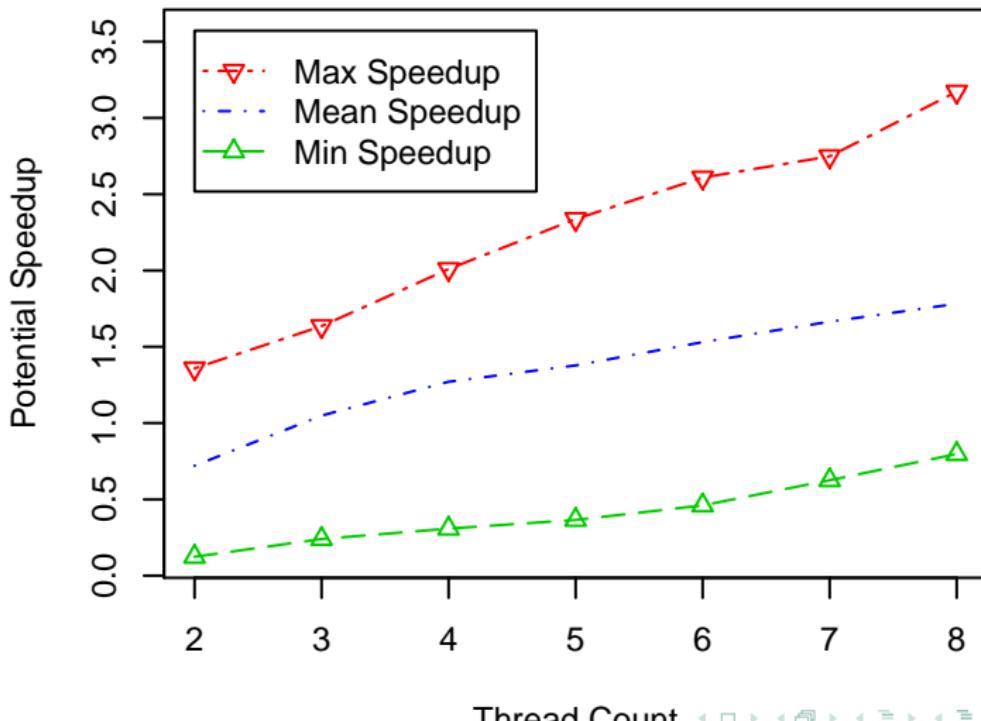
CHARACTERS IN 10,000 INTEGERS IN A RANGE OF STATES



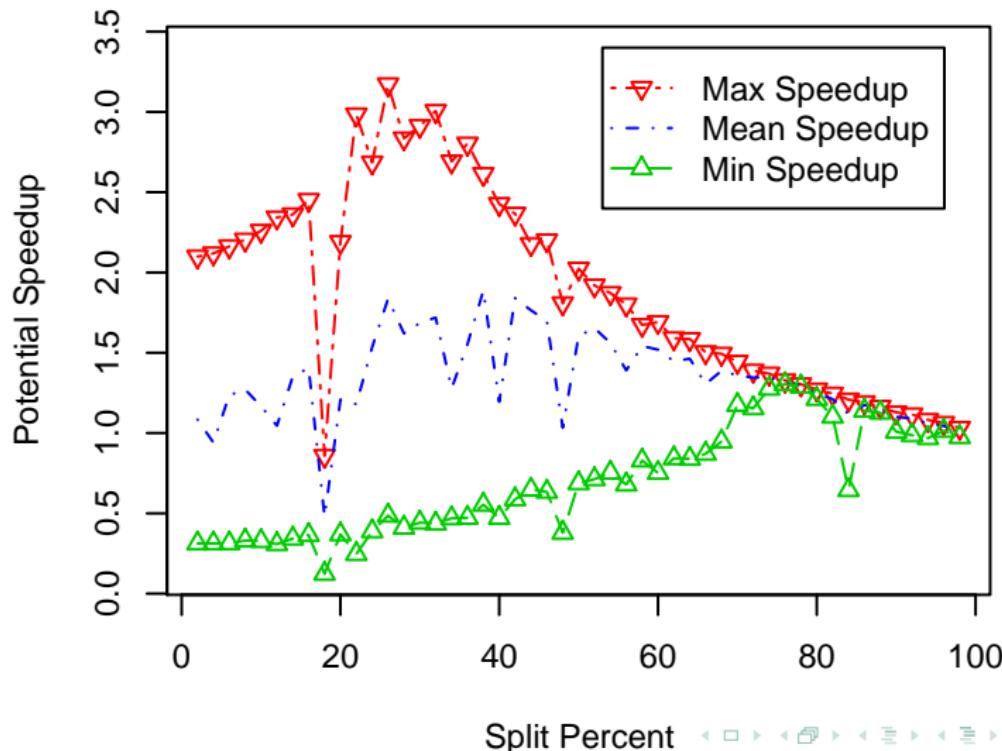
CONCLUSIONS FROM INTEGER RESULTS

- Speedup is possible in this case
- Choice of split point is critical for achieving any speedup at all
- Characters in content sections account for roughly 60% of the input characters
- Input is 117 KB in length
- Consists mainly of
...<i>1234</i><i>1235</i><i>1236</i>...

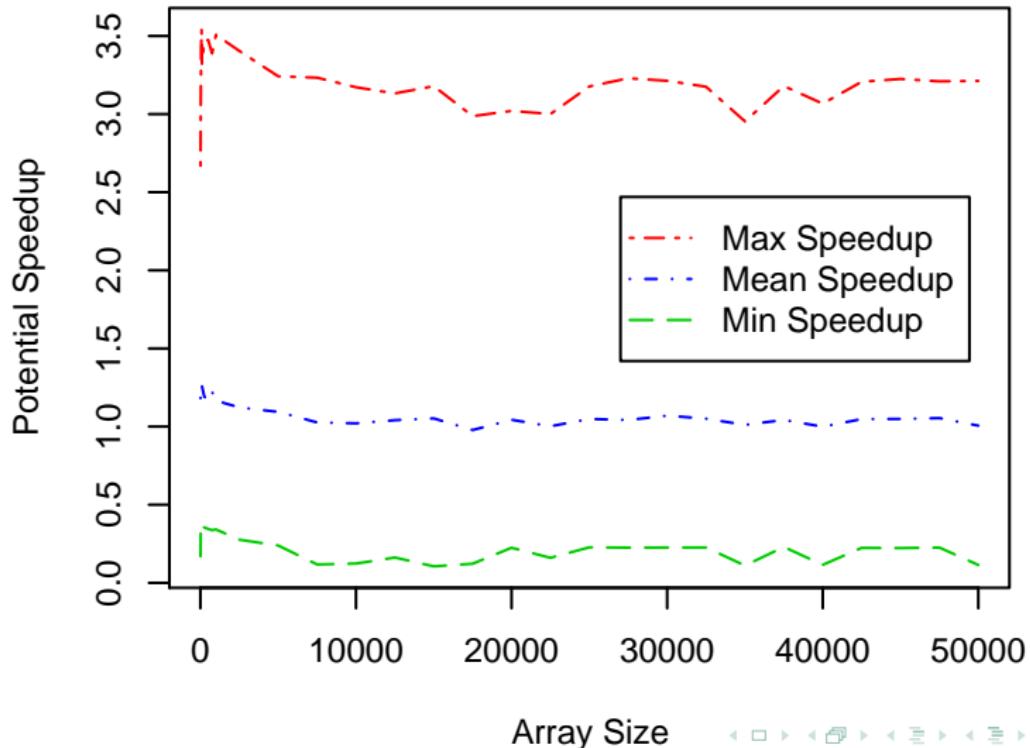
SPEEDUP IMPROVES WITH *Thread_Count* FOR 10,000 STRINGS



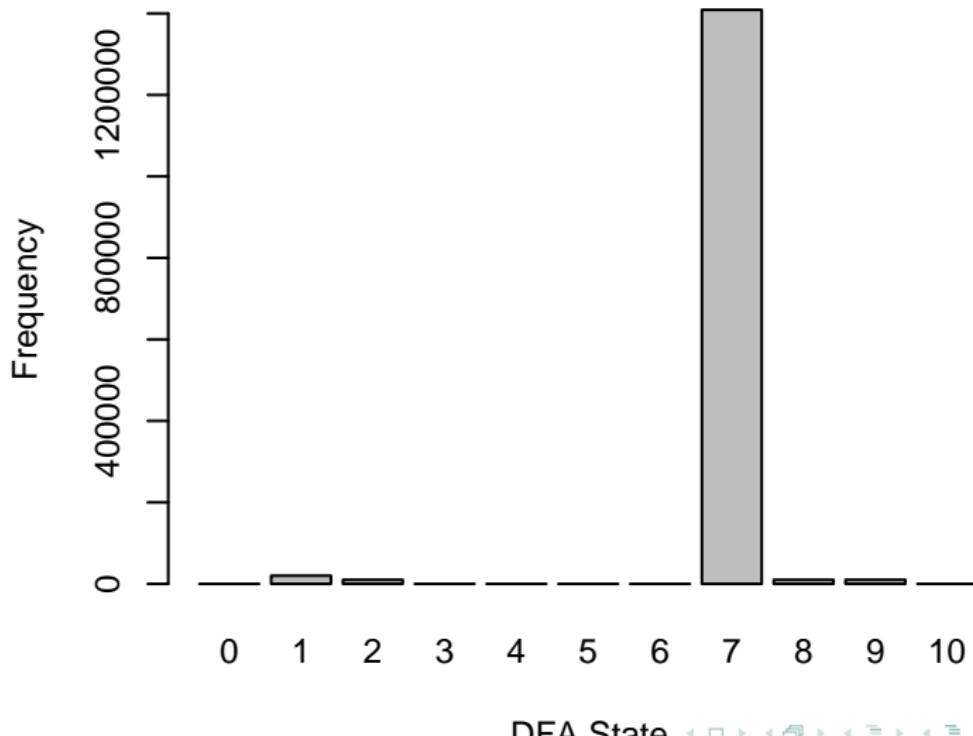
Split_Percent LESS CRITICAL FOR 10,000 STRINGS



CONSISTENT SPEEDUP OVER A RANGE OF INPUT SIZES



CHARACTERS IN 10,000 STRINGS ARE MAINLY IN CONTENT



CONCLUSIONS FROM STRING RESULTS

- This sort of input is much more amenable to this approach
 - In maximum potential speedup achieved
 - In number of cases where speedup is > 1
- Split point is much less important here
- Characters in content sections account for roughly 99% of the input characters
- Input is 1.4 MB in size (though similar results are seen in inputs that are 117 KB)
- Consists mainly of ...<i>String content for the array element number 0. This is long to test the hypothesis that longer content sections are better for the NFA.</i>...

CONCLUSIONS FROM SERIAL NFA TEST

- Shape of the input strongly determines the efficacy of the PIXIMAL approach
 - MIO has similar state usage and mix of content and tags as the integer and PIXIMAL has a similar performance profile there
 - PIXIMAL works well on inputs with longer content sections punctuated by short tags
- Starting in a content section helps because the ‘<’ character eliminates a large number of execution paths through the NFA
 - If ‘>’ could be treated similarly by the parser, starting in a tag would be less harmful

PXML: A BETTER LANGUAGE FOR PIXIMAL

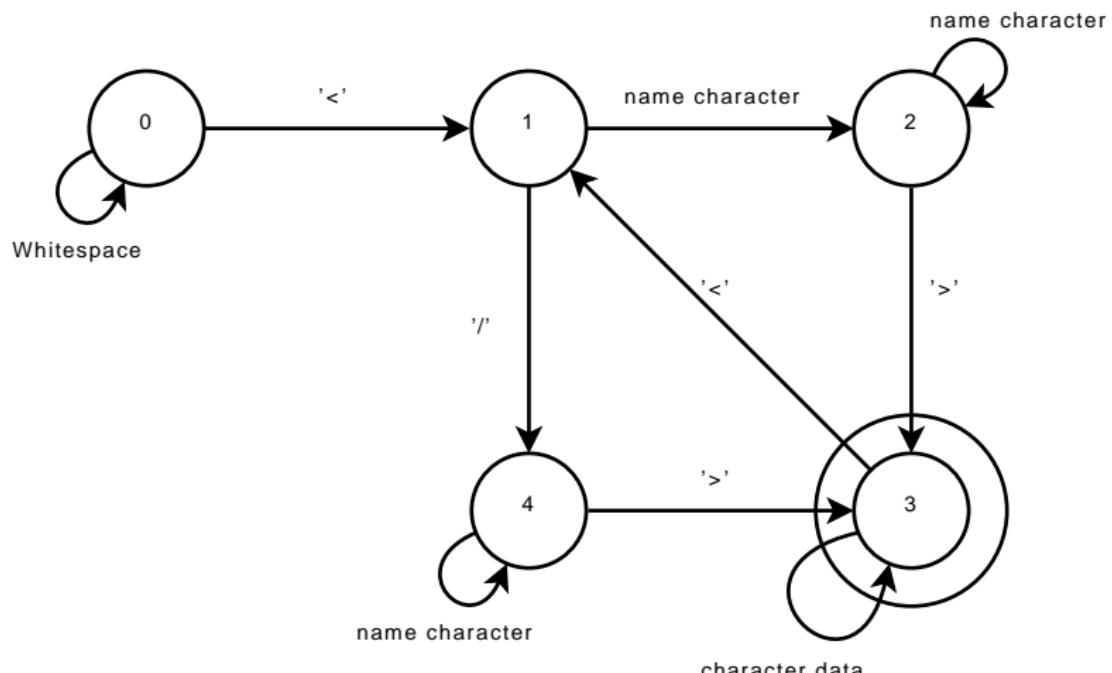
Goal: Improve PIXIMAL performance

- Reduce DFA size
- Increase the number of paths that lead to contradictions

Restrict XML (as supported in PIXIMAL) in the following ways:

- **Disallow attributes:** Transform into nested elements
- **Disallow whitespace in tags:** Without attributes, these are completely unnecessary
- **Disallow ‘>’ in content sections:** Unnecessary in any case
- **Ignore distinction between characters that start a name and the rest**

DFA FOR PIXIMAL-PXML



RELATED WORK IN HIGH PERFORMANCE XML PROCESSING

- Look-aside buffers/String caching (gsoap, XPP)
- Trie data structure with schema-specific parser (Chiu et al 02, Engelen 04)
- One pass table-driven recursive descent parser (Zhang et al 2006)
- Pre-scan and schedule parser (Lu et al 2006)
- Parallelized scanner, scheduled post-parser (Pan et al 2007)

CONCLUSIONS

- Existing XML and SOAP toolkits make limited use of multiple cores
- Scientific applications strain existing XML infrastructure
- Pre-caching mechanisms can improve performance of existing parsers
- A parallel parsing approach is necessary to achieve increased parser performance as document sizes grow
- 5-6 states is a good size for a Piximal DFA
- Restricting XML slightly should provide better performance at a low semantic cost
- Piximal's applicability is dependent on the characteristics of the input file

LIMITATIONS

- PThread overhead during concurrent runs
- Restrictions on XML format
 - Namespaces
 - CDATA
 - Unicode
 - Processing Instructions
 - Validation
- Optimal splitting algorithm unknown

SUMMARY

1 INTRODUCTION AND MOTIVATION

- XML and SOAP
- Ubiquity of Multi-processing Capabilities
- Contributions

2 SOAP AND XML BENCHMARKS

- SOAPBench
- XMLBench

3 PARALLEL XML

- Investigating System Cache Effects
- PixIMAL: Parallel Approach for Processing XML

4 RELATED WORK

5 CONCLUSIONS AND FUTURE WORK

Thank you for your time.

Questions?

EXTRA SLIDES

The following slides are additional and not part of the presentation.

PROPOSED WORK

RE-RUN BENCHMARKS, NORMALIZE ANALYSIS AND PLOTTING

SOAPBench and XMLBench results should be re-run. Plots should be rebuilt to match the rest of the figures.

- XMLBench is available for researchers to download and use
- SOAPBench is available, but cannot support all the tested SOAP toolkits due to their proprietary nature

ANALYZE A BROADER RANGE OF DATA FROM THE SERIAL NFA TEST

The serial NFA tests show a small portion of the data collected in that test. There is a wealth of information to uncover about the efficacy of this approach in the data.

- Data and analysis is available in our repository and will be posted to a web site shortly

PROPOSED WORK CONTINUED

INVESTIGATE MEMORY ALLOCATION ISSUES

Heap contention is a well known problem for applications with concurrent memory allocations. We plan to investigate the effect of a variety of allocators on PixIMAL. During PixIMAL development, we encountered some issues involving the the performance of malloc once a thread (even a thread with an empty *start_routine*) was created. We plan to investigate and report on this in detail.

- Have initial results (HPDC 2009), potential for broader investigation remains

PROPOSED WORK CONTINUED

DEFINE CHARACTERISTICS OF A RESTRICTED SUBSET OF XML DOCUMENTS: “PXML”

Based on the above results, we can design a language which works best with PixIMAL-NFA. Potential targets include eliminating ‘>’ from content sections, removing CDATA sections, disallowing extra whitespace in tags, and perhaps eliminating attributes altogether.

- Briefly described in Chapter 5, Section 4 of the thesis document
- A formal grammar was not considered necessary for the scope of the thesis

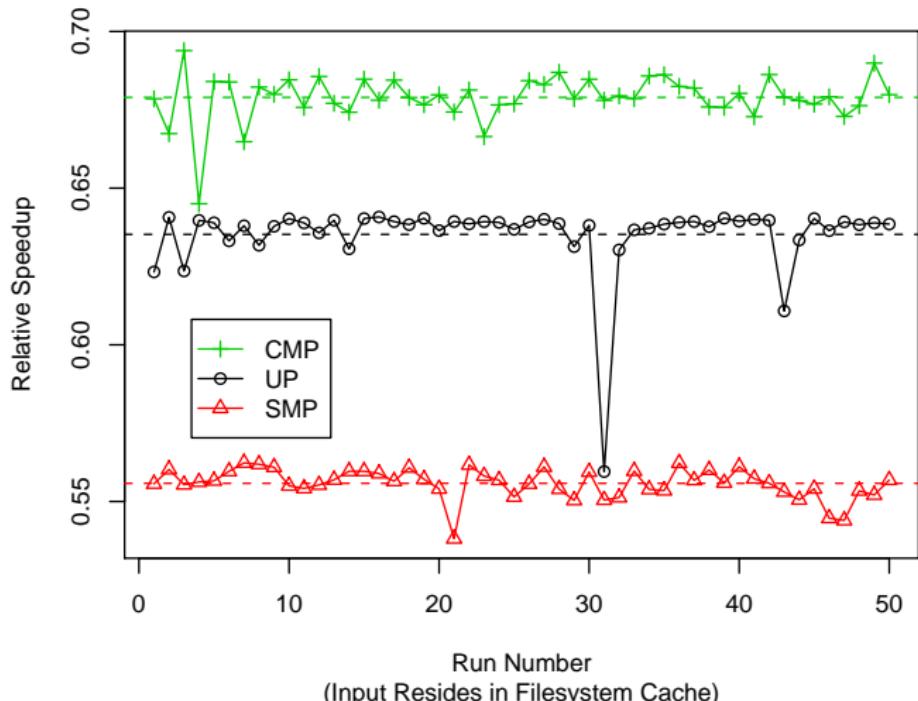
OVERCOMING SEQUENTIALITY WITH AN NFA

- Problem: start state is unknown
- Solution: assume every possible state is a start state
 - Construct an NFA from the DFA used in PixIMAL-DFA
 - 1 Mark every state as a start state
 - 2 Remove all the garbage state and all transitions to it
 - 3 Create a queue for each start state to store actions that should be performed
 - Such an NFA can be applied on any substring of the input
- PixIMAL-NFA is the parser that does all of this:
 - Partition input into segments
 - Run PixIMAL-DFA on the initial segment
 - Run NFA-based parsers on subsequent partition elements
 - Fix up transitions at partition boundaries and run queued actions

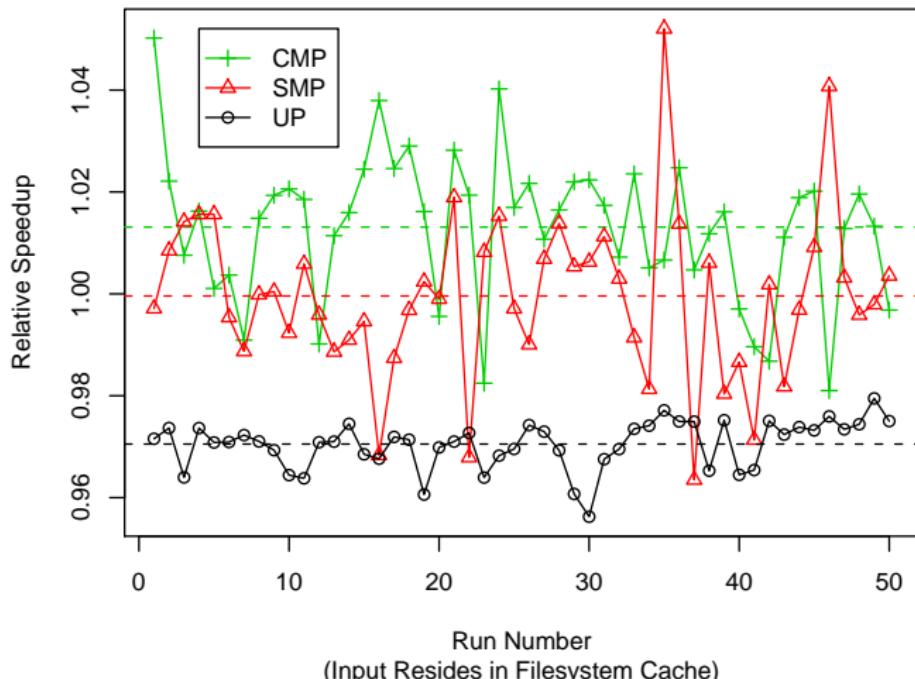
PIXIMAL-DFA IMPLEMENTATION DETAILS

- `mmap(2)` s input file to save memory
- Uses `{length, pointer}` string representation
 - Strings (for tagnames, attribute values) point into the mapped memory
 - All the way through the SAX-style event interface
- DFA is encoded as two tables
 - Table of "next" state numbers indexed by state number and input character
 - Table of boolean "action required" indicators indexed by "current" state and "next" state
 - Action required \implies a function is called to decode and execute the required action
- DFA table is generated at compile time using a separate generator program

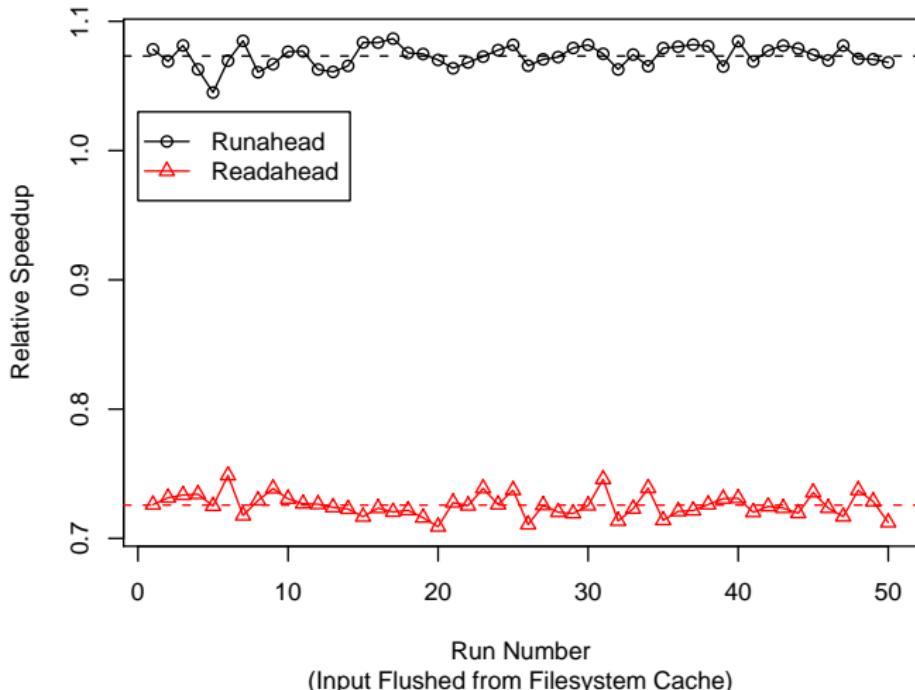
Speedup for the Readahead Parser Relative to Architecture



Speedup for the Runahead Parser Relative to Architecture



Speedup for the CMP Architecture Relative to Parser Type



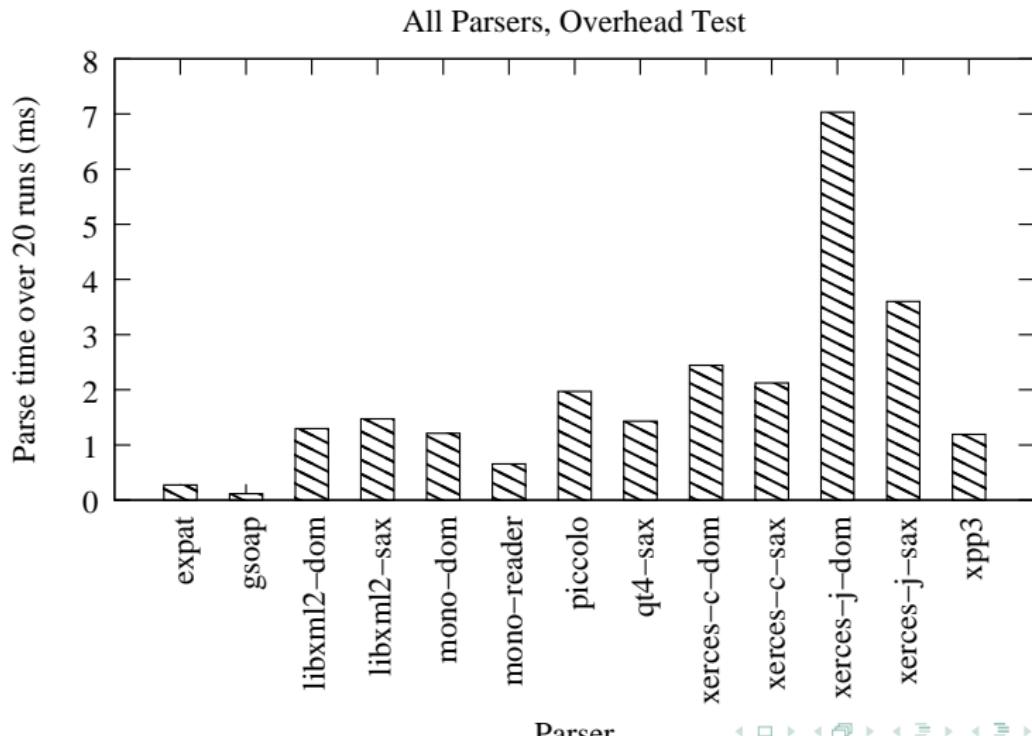
BENCHMARK PROBES

- Overhead test
 - Minimal XML document
 - (header plus one self-closing element)
- Buffering
 - Repeated use of *xsi:type* attributes
- Namespace management
 - Gratuitous use of *xml/ns* attributes
- SOAP payloads
 - “Interop” test: arrays of integer, string, double, MIO, event objects

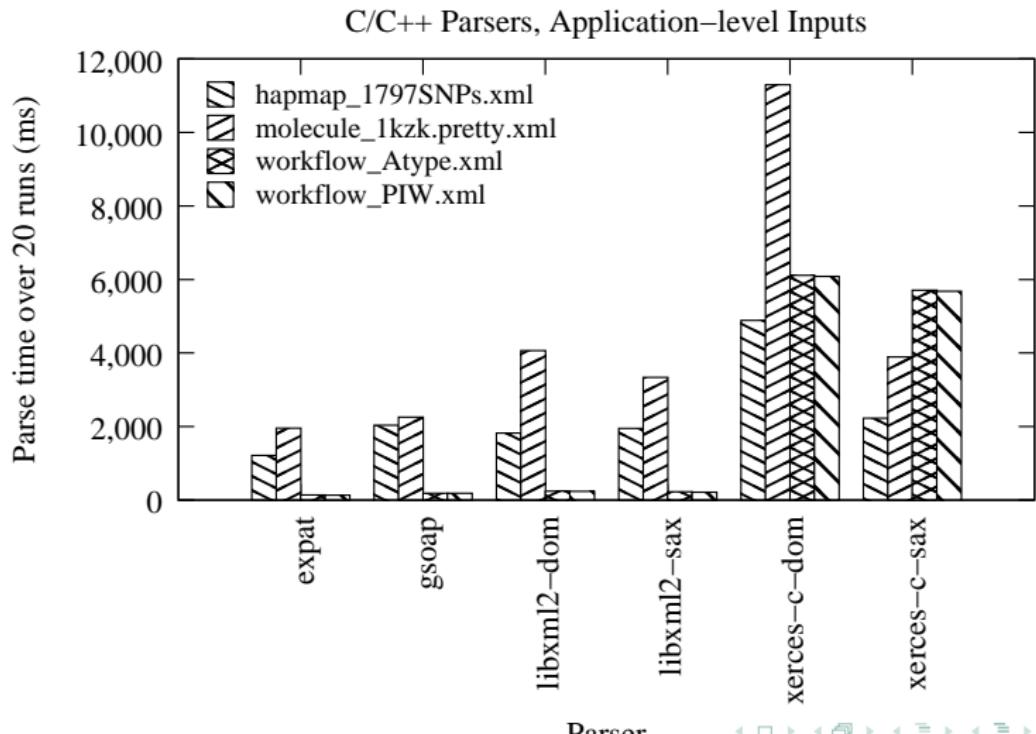
BENCHMARKS FOR SELECTED APPLICATIONS

- Ptolemy Workflow documents (which Kepler uses)
- Genetic data files
 - (Large) files from the International HapMap Project
- Molecular data
- Mesh interface objects, event streams (WSMG)
- WS-Security documents

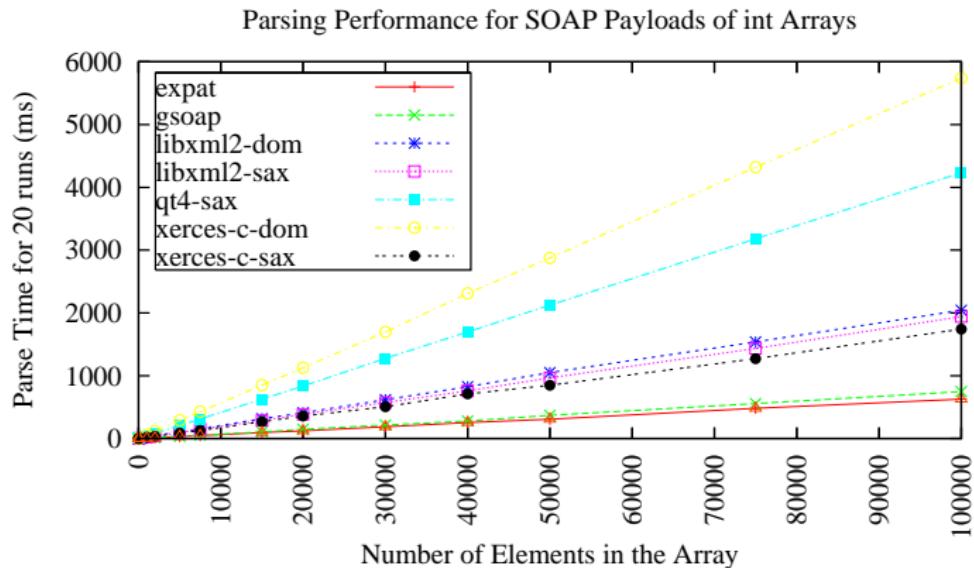
OVERHEAD OF EACH PARSER



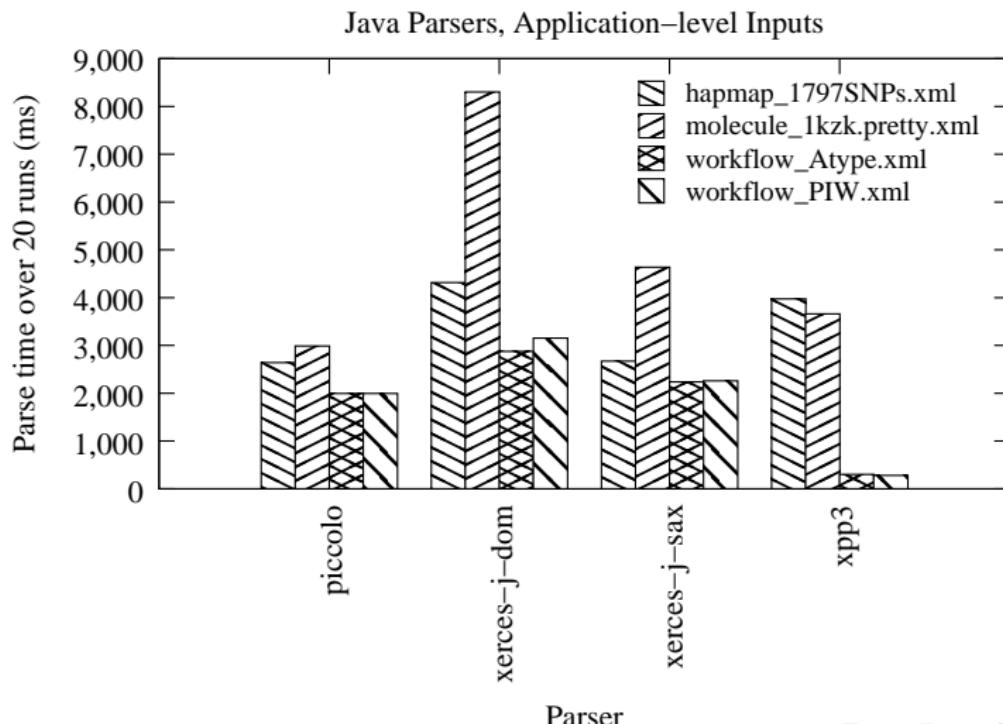
PERFORMANCE OF C AND C++-BASED PARSERS



C PARSER PERFORMANCE OVER SOAP PAYLOADS



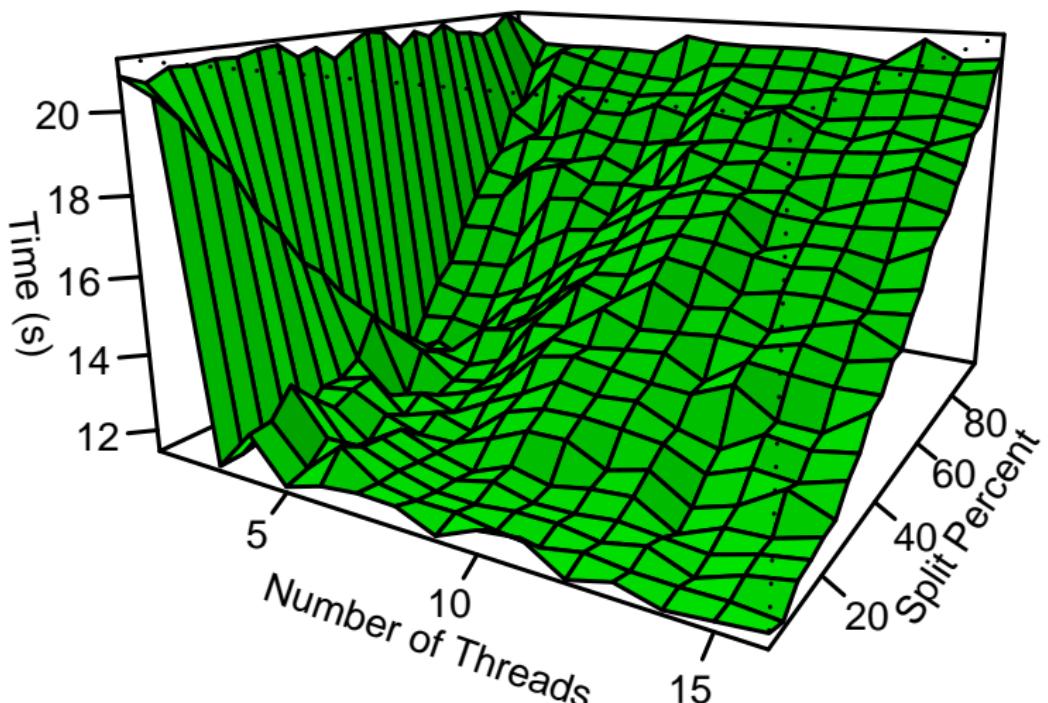
PERFORMANCE OF JAVA-BASED PARSERS



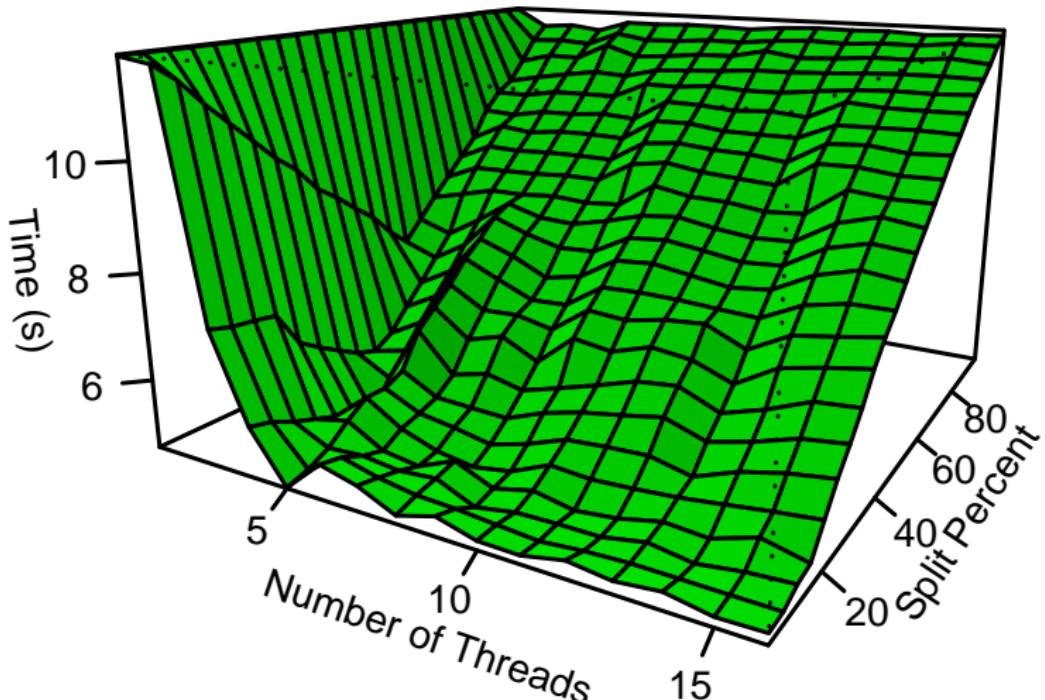
XMLBENCH CONCLUSIONS

- Low overhead \implies gSOAP and Expat, XPP3
- gSOAP performs well with namespaces due to look-aside buffers
- Piccolo and XPP3 have comparable performance in Java

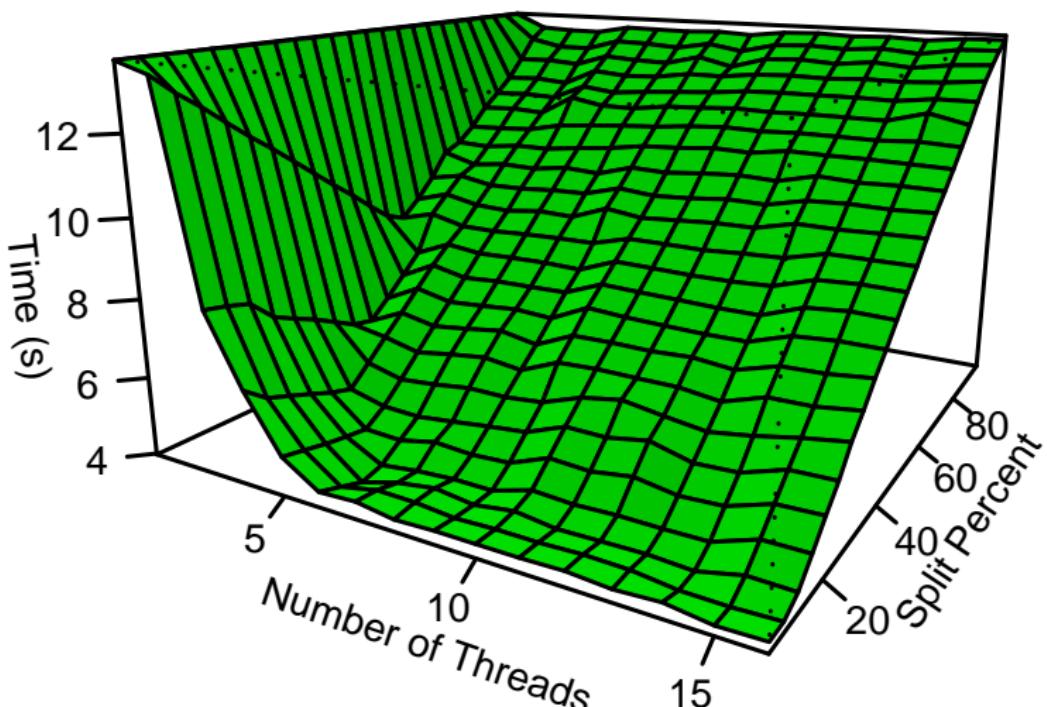
2× UP OVERALL RESULTS



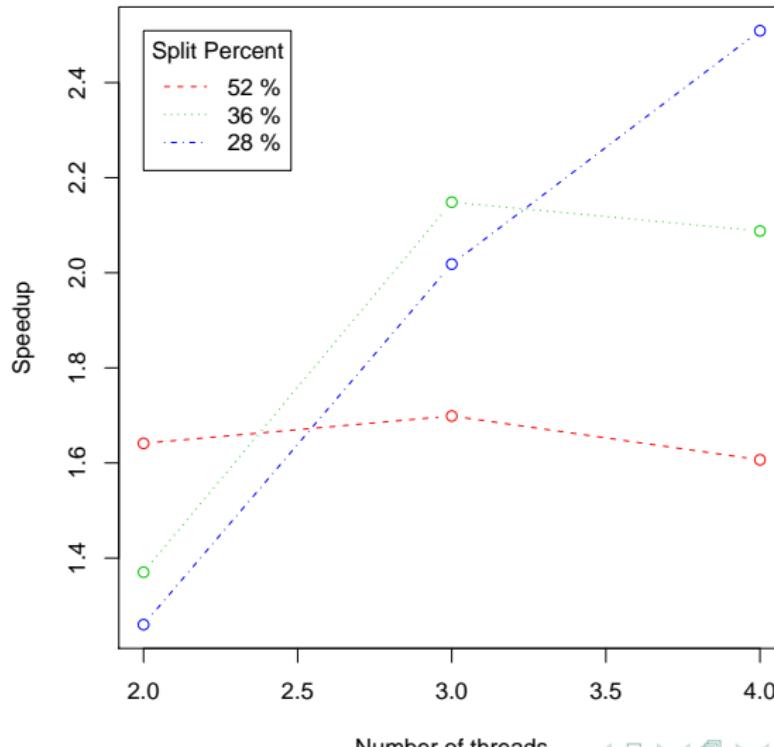
2× DC OVERALL RESULTS

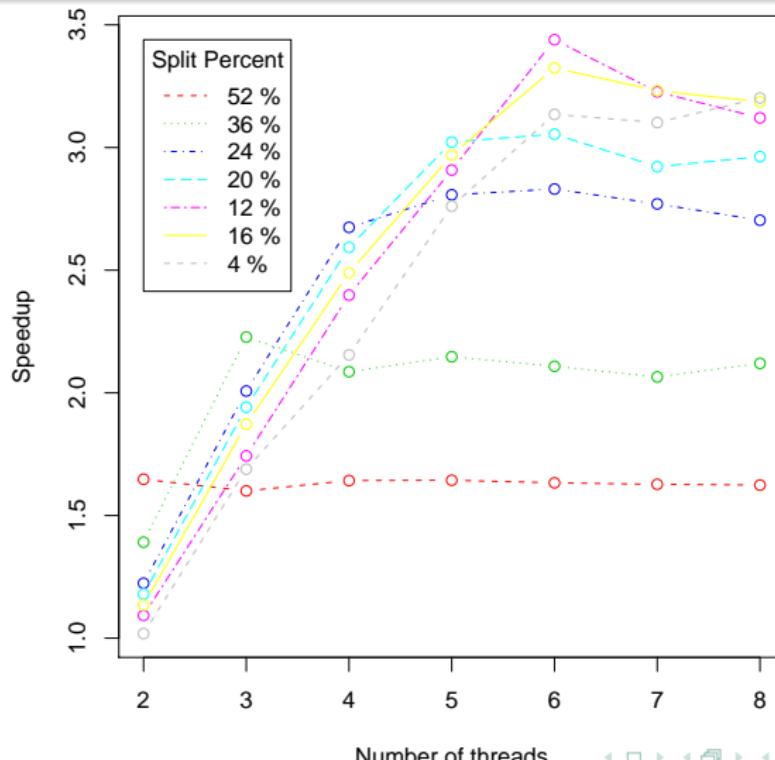


2x QC OVERALL RESULTS



2× DC SPEEDUP FOR BEST *split_percents*

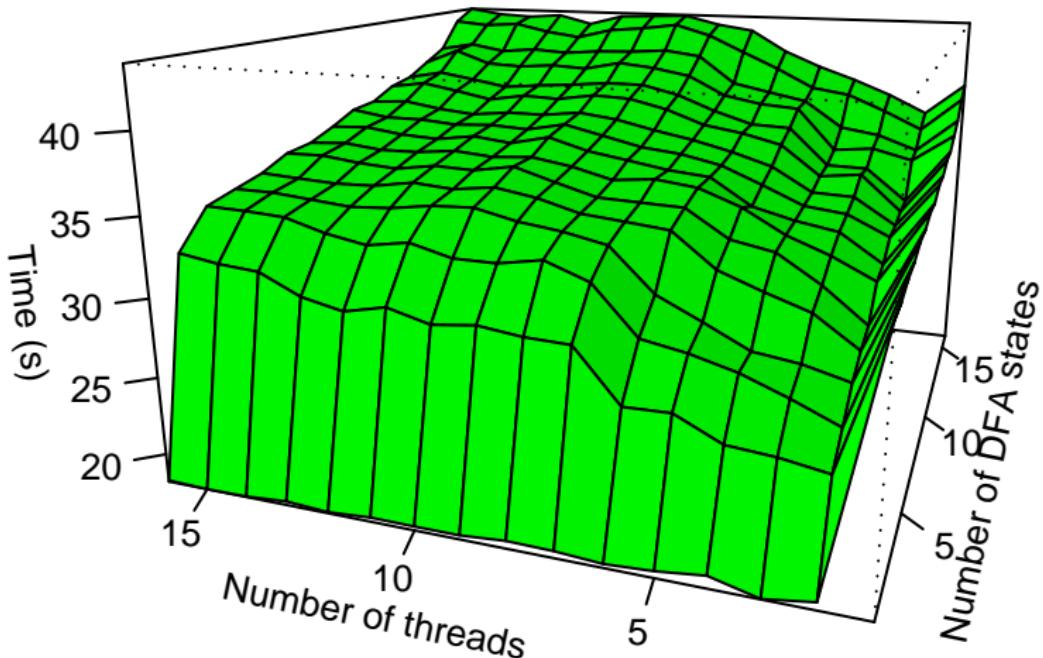


2× QC SPEEDUP FOR BEST *split_percents*

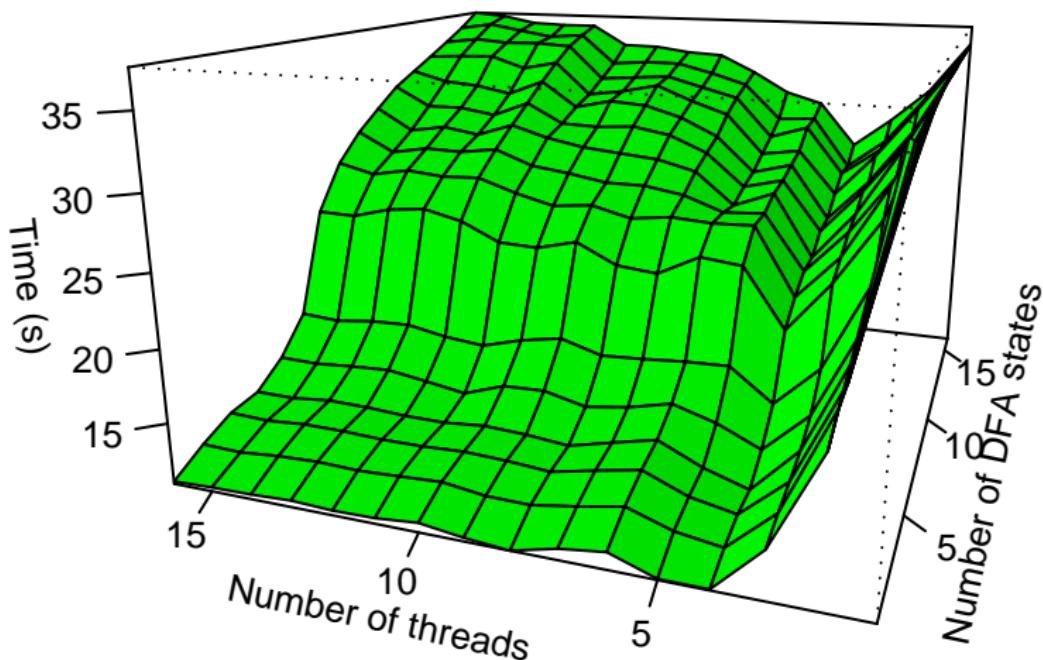
CONCLUSIONS FROM SPEEDUP CROSS SECTIONS

- Reaffirmation that speedup is possible
- Returns diminish for these machines at around 6 threads
- Overall, access to main memory is not an immediate bottleneck
- Putting the results from the best *split_percents* for each architecture...

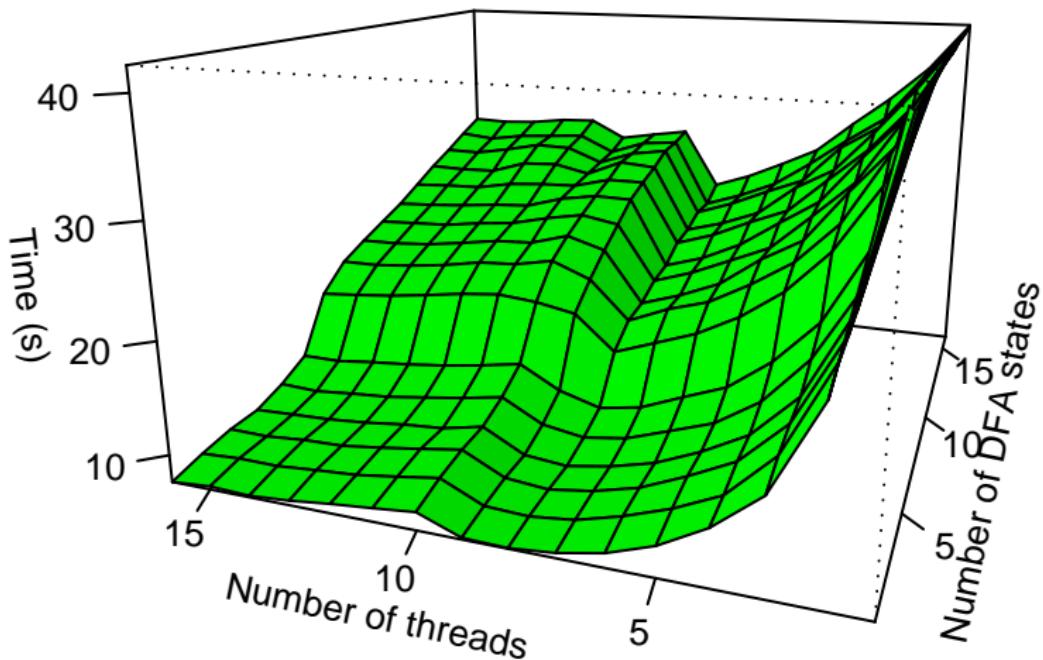
2× UP OVERALL RAW RESULTS



2× DC OVERALL RESULTS – BEST TIMES



2x QC OVERALL RESULTS – BEST TIMES



CONCLUSIONS FROM STATE SCALABILITY OVERALL RESULTS

- Two major conclusions:
 - The speedup on the 2× quad-core machines appears stable as the number of threads increases
 - There is a significant steepening when the DFA has 6-7 states
- Performance reaches its max when the number of threads match the number of processing cores available
 - Each new thread adds substantial extra work compared with the memory bandwidth test
- Plotting speedup for certain *split_percents*

XML PERFORMANCE LIMITATIONS

- Compared to “legacy” formats
 - Text-based
 - Lacks any “header blocks” (ex. TCP headers), so must scan every character to tokenize
 - Numeric types take more space and conversion time
 - Lacks indexing
 - Unable to quickly skip over fixed-length records

LIMITATIONS OF XML

- Poor CPU and space efficiency when processing scientific data with mostly numeric data (Chiu et al 2002)
- Features such as nested namespace shortcuts don't scale well with deep hierarchies
 - May be found in documents aggregating and nesting data from disparate sources
- Character stream oriented (not record oriented): initial parse inherently serial
- Still ultimately useful for sharing data divorced of its application

READING AHEAD

- Introduce two parsers which extend the existing, high performance **Piccolo** parser (Head et al 2006)
 - **Runahead:** opens two file descriptors for the input file
 - Start a thread that repeatedly calls `read()` on one of the file descriptors
 - Pass the other file descriptor to the existing Piccolo parser in the main thread
 - **Readahead:** opens one file descriptor for the input file, and one pipe
 - Start a thread that reads from the file descriptor and writes to the pipe
 - Pass the pipe to the existing Piccolo parser in the main thread

TEST RUN

- Run each parser (**Piccolo**, **Runahead**, and **Readahead**) on a large (GB-scale) XML file
 - Specifically, a protein sequence database file, `psd7003.xml`
- No user code is run for any SAX event -- just the parser itself is tested
- File cache is cleared between each run running a separate process that reads multiple gigabyte files
- Each test is run 50 times for each parser
- Hotspot is warmed by running the parser on another input file with identical content before timing begins

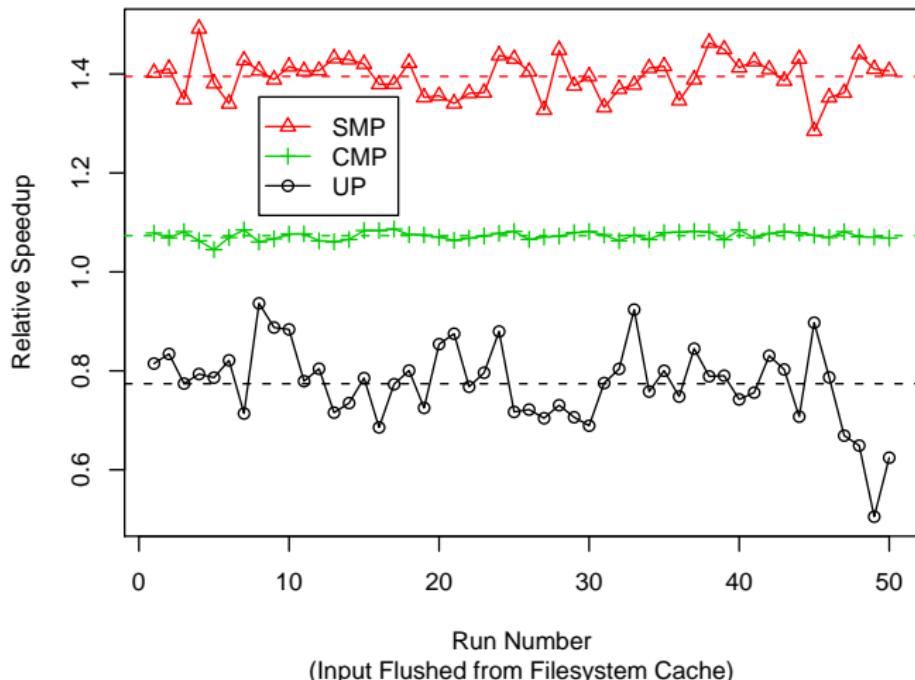
Two Environmental Conditions Tested

- Architectures
 - **UP**: Classic Uniprocessor P4-based machine (Dell workstation)
 - **SMP**: Classic Symmetrical MultiProcessing P4-based machine (has server-class I/O system) (IBM e-server)
 - **CMP**: Modern Chip MultiProcessing Core 2 Duo-based machine (Dell workstation)
- System conditions
 - **Cached**: The input file is read (hence loaded into the system file cache) before timing begins
 - **Uncached**: The input file is not read before timing begins (and flushed between each run)

DATA ANALYSIS

- Speedup for both of the proposed parsers is computed to compare across architectures
- Baseline value is computing by averaging the times for each run of the unmodified **Piccolo** parser
- Speedup for each run is computed by dividing the baseline by the time at each test point

Speedup for the Runahead Parser Relative to Architecture



READAHEAD CONCLUSIONS

- On systems with available memory and an available processing core with fresh inputs, this approach can provide some performance wins.

COMPARISON WITH EXPAT

Input file	Expat	Piximal-dfa	Piximal-nfa
psd-7003	15.51	17.47	14.18

TABLE: Parse time, in seconds per parse, of high performance parsers

COMPARISON BETWEEN GLIBC AND TCMALLOC

