

PERFORMANCE ENHANCEMENT WITH SPECULATIVE EXECUTION BASED PARALLELISM FOR PROCESSING LARGE-SCALE XML-BASED APPLICATION DATA

Michael R. Head and Madhusudhan Govindaraju

Grid Computing Research Laboratory
Department of Computer Science
Binghamton University

<http://www.cs.binghamton.edu/~{mike,mgovinda}>

HPDC 2009

Thursday, June 11, 2009

OUTLINE

1 INTRODUCTION

- Large XML Data
- Ubiquity of Multi-processing Capabilities
- SAX-based parsing

2 PARALLEL XML

- PIXIMAL: Parallel Approach for Processing XML
- Serial NFA Tests

3 CONCLUSIONS

- Final Remarks

OUTLINE

1 INTRODUCTION

- Large XML Data
- Ubiquity of Multi-processing Capabilities
- SAX-based parsing

2 PARALLEL XML

- PIXIMAL: Parallel Approach for Processing XML
- Serial NFA Tests

3 CONCLUSIONS

- Final Remarks

XML

- 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

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

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)

OUTLINE

1 INTRODUCTION

- Large XML Data
- Ubiquity of Multi-processing Capabilities
- SAX-based parsing

2 PARALLEL XML

- PIXIMAL: Parallel Approach for Processing XML
- Serial NFA Tests

3 CONCLUSIONS

- Final Remarks

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

XML AND MULTI-CORE

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

OUTLINE

1 INTRODUCTION

- Large XML Data
- Ubiquity of Multi-processing Capabilities
- SAX-based parsing

2 PARALLEL XML

- PIXIMAL: Parallel Approach for Processing XML
- Serial NFA Tests

3 CONCLUSIONS

- Final Remarks

SAX-STYLE XML PARSING

- Sequential processing model
 - Program invokes parser with a set of callback functions
 - Parser scans input from start to finish
 - `<element attributes...>`
 - *content*
 - `</element>`
 - Invokes callbacks in file order
 - `startElement()`
 - `content()`
 - `endElement()`

OUTLINE

1 INTRODUCTION

- Large XML Data
- Ubiquity of Multi-processing Capabilities
- SAX-based parsing

2 PARALLEL XML

- PIXIMAL: Parallel Approach for Processing XML
- Serial NFA Tests

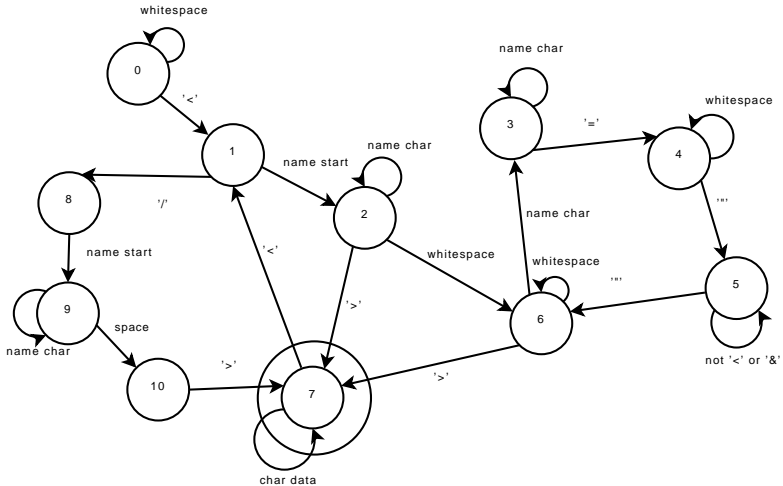
3 CONCLUSIONS

- Final Remarks

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



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

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
 - 1 Mark every state as a start state
 - 2 Remove all the garbage state and all transitions to it
 - 3 Create an 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-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
 - The final $(100 - \textit{split_percent})\%$ of the input is divided evenly across the remaining $(\textit{number_of_threads} - 1)$ partitions
 - The final partition element gets up to $\textit{number_of_threads} - 2$ fewer characters

OUTLINE

1 INTRODUCTION

- Large XML Data
- Ubiquity of Multi-processing Capabilities
- SAX-based parsing

2 PARALLEL XML

- PIXIMAL: Parallel Approach for Processing XML
- Serial NFA Tests

3 CONCLUSIONS

- Final Remarks

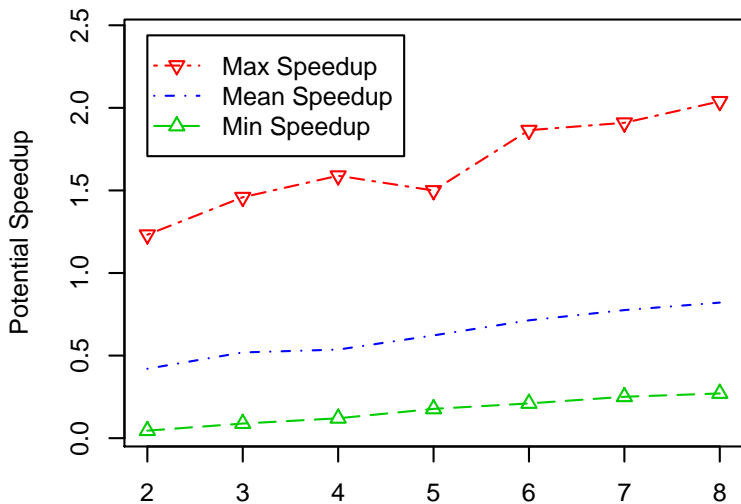
SERIAL NFA TESTS

- 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

TEST CONDITIONS


- Synthetic data
 - Arrays of Integers, Strings, Mesh Interface Objects
 - SOAP encoded
 - Same as previously presented in benchmarks
- Across a cluster (taking mean of results)
- Range of input sizes
- Range of parameters (*split_percent*, *number_of_threads*)

MODEST SPEEDUP SCALABILITY FOR 10,000 INTEGERS

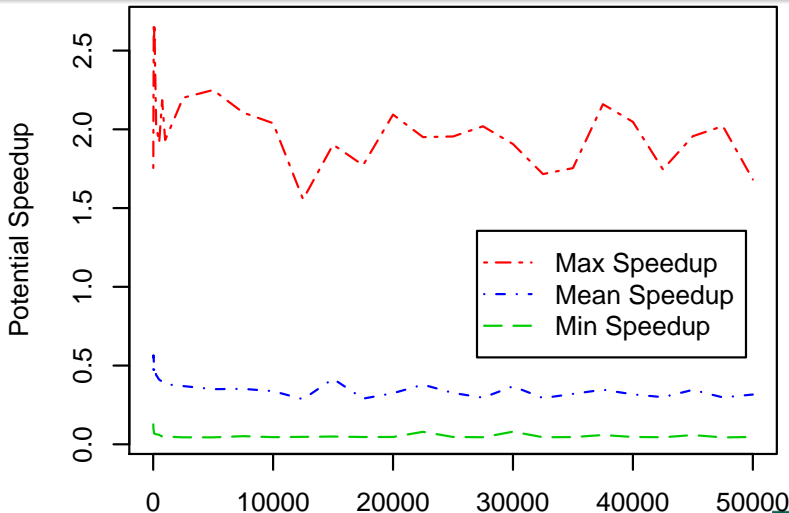


Thread Count


Split_Percent CRITICAL FOR SPEEDUP FOR 10,000 INTEGERS



INCONSISTENT SPEEDUP OVER A RANGE OF ARRAY LENGTHS

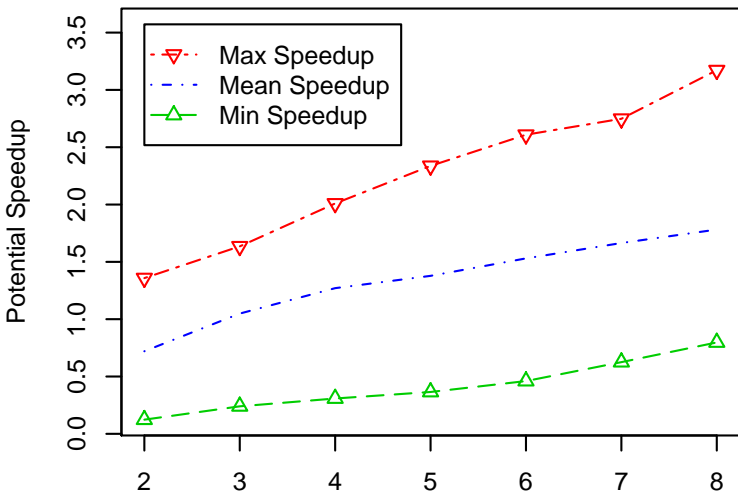


CHARACTERS IN 10,000 INTEGERS IN A RANGE OF STATES

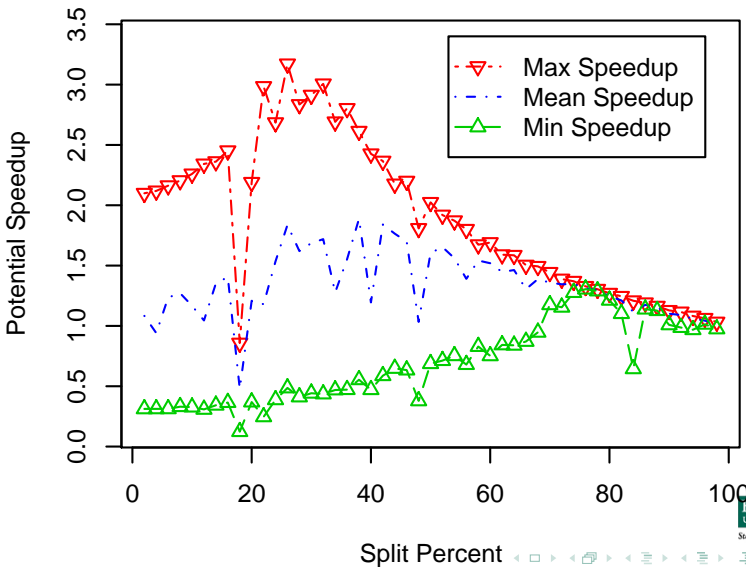


CONCLUSIONS FROM INTEGER RESULTS

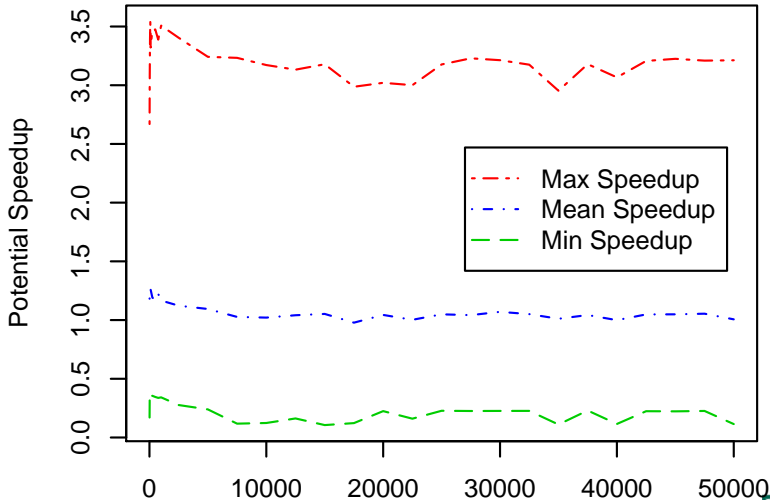
- Speedup is possible in this case
- Choice of split point is critical for achieving any speedup at all
- Characters in content sections account for roughly 60% of the input characters
- Input is 117 KB in length
- Consists mainly of
...<i>1234</i><i>1235</i><i>1236</i>...

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

OUTLINE

1 INTRODUCTION

- Large XML Data
- Ubiquity of Multi-processing Capabilities
- SAX-based parsing

2 PARALLEL XML

- PIXIMAL: Parallel Approach for Processing XML
- Serial NFA Tests

3 CONCLUSIONS

- Final Remarks

CONCLUSIONS

- Scientific applications strain existing XML infrastructure
- A parallel parsing approach is necessary to achieve increased parser performance as document sizes grow
- Restricting XML slightly should provide better performance at a low semantic cost
- PIXIMAL's applicability is dependent on the characteristics of the input file

SUMMARY

1 INTRODUCTION

- Large XML Data
- Ubiquity of Multi-processing Capabilities
- SAX-based parsing

2 PARALLEL XML

- PIXIMAL: Parallel Approach for Processing XML
- Serial NFA Tests

3 CONCLUSIONS

- Final Remarks

Thank you for your time.

Questions?

EXTRA SLIDES

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

LIMITATIONS

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

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)

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



PERSPECTIVE PLOT FOR 10,000 INTEGERS



PERSPECTIVE PLOT FOR 10,000 STRINGS

