Automatic Discovery of Efficient Divide-&-Conquer Algorithms for Dynamic Programming Problems

Pramod Ganapathi (Stony Brook University)

Rezaul Chowdhury, Pramod Ganapathi, Stephen Tschudi, Jesmin Jahan Tithi, Charles Bachmeier, Charles E Leiserson, Armando Solar-Lezama, Bradley C Kuszmaul, and Yuan Tang. ACM Transactions on Parallel Computing (TOPC), 2017





Automatic discovery of efficient algorithms

Automation



ATMs



Vending machines



Ticket machines



e-commerce



Manufacturing



House appliances



Self-driving cars



Robots

Automation in CS



IBM Watson



Mathematica



E-mail



Navigation



Automated testing

d Rooks file latt Vew So Bore w	Index Inde	0 1 9 41 B Serbeau G H
	Robalmanh Channel	-A Q #-
Netwo G		
Chapter X-Tarres Scent	Chapter X	However, strangh are tolde paring as we further, jour the size of a way ner at all assumption to the design which I had its sizes when I tasks the first, I seem of semanting users to the zero data, where is was above large strains band, assumpting, the analyses of try loss mathem to put as sen is the design, and area if design to a most of a field
•	Turnes Goats	had a basis, one most design true to index a rooke transf the index. (in an 1 had been an the other only in one place, resuming, in 1 have attempt described is, one the local, so the disense it must be that their inter-
Courses has been at the C	I carrent see that after this, for the years, are entrandinary thing hap- pened to me, but I liked on its the same traver, in the same postore	see that me very sugar in see other parts of the book, and now I had a beat, I thought of animing bac sating mond the intend.
	The second secon	In the spin of the phi is entering and include and in the spin of
I block my satisfies also is due omy et die wenn, lies a maie, no wand ense my freid, auf kang die base of die san off ke, lies at permage auf dies i dwart now with dies hald k allem wage span die son, bat sever want die son, son las foot, the lidite onch. A		April Manual
· · · · · · · · · · · · · · · · · · ·		
state Spanification		

Text-to-speech converter

Automatic programming



Automatic code generation



Automatic algorithm design



Aim: Automatic algorithm design



Automatic algorithm design is inevitable









Programming is hard

Machine architectures are changing

Problems are increasing

State-of-the-art in automatic programming

- Automated parameter tuning
- Automatic parallelizer and locality optimizer
- Program synthesis
- Machine learning
- Self-improving algorithms

No algorithms that can discover nontrivial (in structure) D&C algorithms

We focus on dynamic programming (DP) problems

Dynamic programs are important

D&C is the best way to implement dynamic programs

I-DP = iterative algorithms Tiled I-DP = blocked iterative algorithms R-DP = recursive divide-and-conquer algorithms

Feature	I-DP	Tiled I-DP	R-DP
Cache-efficiency			
Cache-obliviousness			
Cache-adaptivity			
High parallelism			
Highly optimizable kernels			
Energy efficiency			
Bandwidth efficiency			

Currently, D&C algorithms are designed manually

Sequence alignment Parenthesis problem Gap problem Floyd Warshall's APSP Multi-instance Viterbi Spoken word recog. Function approx.

Computational scientists need efficient dynamic programs

We want to automatically discover efficient and portable DP algorithms

Blackbox implementation of a DP recurrence

Parallelism and cache locality are the key factors to improve performance

Parallelism

Cache locality

Communication is more expensive than computation

Feature	L1	L2	L3	RAM	Disk
Access time	O.5ns	7ns	20ns	100ns	1 ms
Size	32KB	256KB	20MB	32GB	2TB
Cost	-	-	-	\$150	\$15O
Technology	SRAM	SRAM	SRAM	DRAM	Magnetic

Source: Intel and StackExchange forums

I-DP -> Autogen -> R-DP

Example: Parenthesis problem

Parenthesization that minimizes the cost

• Applications:

- Optimal chain matrix multiplication
- RNA secondary structure prediction
- Optimal binary search trees
- Optimal polygon triangulation
- String parsing for context-free grammar CYK algorithm
- Optimal natural join of database tables Selinger algorithm
- Maximum perimeter inscribed polygon
- Offline job scheduling minimizing flow time of jobs

Parenthesis problem recurrence

Input: A string $s_1 s_2 \dots s_n$ $C[i, j] = \text{cost of parenthesization from } s_i \text{ to } s_j$

$$C[i,j] = \begin{cases} \infty & \text{if } 0 \le i = j \le n \\ v_j & \text{if } 0 \le i = j - 1 < n \\ \min_{k \in [i,j]} \{C[i,k] + C[k,j] + w(i,k,j)\} & \text{if } 0 \le i < j - 1 < n \end{cases}$$

Output: *C*[1, *n*]

Dependency structure for the parenthesis problem

Not computed
Computed
Write cell
Read cell

Slow loop DP algorithms (I-DPs)

Loop-Parenthesis(C, n)

- 1. for $i \leftarrow n 1$ downto 1 do
- **2.** for $j \leftarrow i + 2$ to n do
- 3. for $k \leftarrow i$ to j do

4.
$$C[i,j] \leftarrow \min \begin{pmatrix} C[i,k]+C[k,j] \\ +w(i,k,j),C[i,j] \end{pmatrix}$$

- [- 1]]

Parallel-Loop-Parenthesis(C, n)

- 1. for $t \leftarrow 2$ to n 1 do
- **2.** parallel for $i \leftarrow 1$ to n t do
- **3.** $j \leftarrow t + i$

4. for
$$k \leftarrow i$$
 to j do

5.
$$C[i,j] \leftarrow \min \begin{pmatrix} C[i,k] + C[k,j] \\ +w(i,k,j), C[i,j] \end{pmatrix}$$

Serial I-DP

Parallel I-DP

Fast D&C DP algorithm (R-DP)

A(X)if X is a small matrix then A-loop(X)1. 2. else 3. parallel: $A(X_{11})$, $A(X_{22})$ 4. $B(X_{12}, X_{11}, X_{22})$ $\mathbf{B}(X, U, V)$ if X is a small matrix then B-loop(X, U, V)1. else 2. **3. B**(X_{21}, U_{22}, V_{11}) 4. parallel: $C(X_{11}, U_{12}, V_{21}), C(X_{22}, X_{21}, V_{12})$ 5. parallel: $B(X_{11}, U_{11}, V_{11}), B(X_{22}, X_{22}, V_{22})$ 6. $C(X_{12}, U_{12}, X_{22})$ 7. $C(X_{12}, X_{11}, V_{12})$ 8. $B(X_{12}, U_{11}, V_{22})$ C(X, U, V)if X is a small matrix then C-loop(X, U, V)1. else 2. З. parallel: $C(X_{11}, U_{11}, V_{11}), C(X_{12}, U_{11}, V_{12})$ $C(X_{21}, U_{21}, V_{11}), C(X_{22}, U_{21}, V_{12})$ 4. parallel: $C(X_{11}, U_{12}, V_{21}), C(X_{12}, U_{12}, V_{22})$ $C(X_{21}, U_{22}, V_{21}), C(X_{22}, U_{22}, V_{22})$

Fast D&C DP algorithm (R-DP)

I-DP -> Autogen -> R-DP

Loop-Parenthesis(C, n) 1. for $i \leftarrow n - 1$ downto 1 do 2. for $j \leftarrow i + 2$ to n do 3. for $k \leftarrow i$ to j do 4. $C[i, j] \leftarrow \min\begin{pmatrix} C[i, k] + C[k, j] \\ + w(i, k, j), C[i, j] \end{pmatrix}$ I-DP

Autogen

A(X)

- 1. if X is a small matrix then A-loop(X)
- 2. else
- **3.** parallel: $A(X_{11}), A(X_{22})$
- 4. $B(X_{12}, X_{11}, X_{22})$

$\mathbf{B}(X, U, V)$

- 1. if X is a small matrix then B-loop(X, U, V)
- 2. else
- **3.** $B(X_{21}, U_{22}, V_{11})$
- 4. parallel: $C(X_{11}, U_{12}, V_{21}), C(X_{22}, X_{21}, V_{12})$
- 5. parallel: $B(X_{11}, U_{11}, V_{11}), B(X_{22}, X_{22}, V_{22})$
- 6. $C(X_{12}, U_{12}, X_{22})$
- 7. $C(X_{12}, X_{11}, V_{12})$
- 8. $B(X_{12}, U_{11}, V_{22})$

C(X, U, V)

- 1. if X is a small matrix then C-loop(X, U, V)
- 2. else
- **3.** parallel: $C(X_{11}, U_{11}, V_{11}), C(X_{12}, U_{11}, V_{12})$ $C(X_{21}, U_{21}, V_{11}), C(X_{22}, U_{21}, V_{12})$
- 4. parallel: $C(X_{11}, U_{12}, V_{21}), C(X_{12}, U_{12}, V_{22})$ $C(X_{21}, U_{22}, V_{21}), C(X_{22}, U_{22}, V_{22})$

Autogen works on Fractal-DP class of DP problems

One-way sweep property: If a cell x depends on another cell y, then y is fully updated before x reads from y

Fractal property:

Autogen works on Fractal-DP class of DP problems

One-way sweep property: If a cell x depends on another cell y, then y is fully updated before x reads from y

Fractal property:

Fractal property satisfied for:

75% of 16 DP problems in CLRS book60% of 40 DP problems in Lew and Mauch's book

Core idea: Find recursive patterns from DP table updates

Generate DP table accesses for a small DP table. Find recursive patterns in DP table accesses. Build a recursive algorithm around these recursive access patterns.

Step 3. Algorithm-tree labeling

Step 1. Generate a set of DP table updates

Cell-set

Step 1. Generate a set of DP table updates

Small problem, say, n = 64

Generate cell-dependencies in the form below $\langle writecell, readcell_1, \dots, readcell_s \rangle$

Cell-set = set of all cell dependencies

Step 2. Construct an algorithm-tree

	X	

Step 2. Distribute the DP table updates into buckets

Bucketing: Distribute the cell-dependencies into $4 \times 4 \times 4 = 64$ possible buckets

 $\langle writecell, readcell_1, readcell_2 \rangle$

$$(\langle X_{11}, X_{11}, X_{11} \rangle) (\langle X_{11}, X_{11}, X_{11} \rangle) (\langle X_{12}, X_{12}, X_{22} \rangle)$$

X ₁₁	<i>X</i> ₁₂
X ₂₁	X ₂₂

Step 2. Distribute the DP table updates into buckets

Step 2. Combine buckets that writeto and read-from the same regions

Step 2. Combine buckets that writeto and read-from the same regions

Step 2. After combining buckets

Step 3. Give function names to nodes of the algorithm-tree

Steps 2&3. Recursively expand new functions

Steps 2&3. Recursively expand new functions

Steps 2&3. Labeled algorithm-tree

Step 4. Construct algorithm-DAG

Limitations with algorithm-tree: - Order of execution - Parallelism

Solution: Algorithm-DAGs for each function Four rules to construct algorithm-DAGs

Step 4. Four rules to construct algorithm-DAGs

Step 4. Algorithm-DAGs for different recursive functions

Step 4. The DAGs represent an R-DP

A(X)

- 1. if X is a small matrix then A-loop(X)
- 2. else
- **3.** parallel: $A(X_{11}), A(X_{22})$
- 4. $B(X_{12}, X_{11}, X_{22})$

 $\mathbf{B}(X, U, V)$

- 1. if X is a small matrix then B-loop(X, U, V)
- 2. else
- **3. B**(X_{21}, U_{22}, V_{11})
- 4. parallel: $C(X_{11}, U_{12}, V_{21}), C(X_{22}, X_{21}, V_{12})$
- 5. parallel: $B(X_{11}, U_{11}, V_{11})$, $B(X_{22}, X_{22}, V_{22})$ 6. $C(X_{12}, U_{12}, X_{22})$
- 7. $C(X_{12}, V_{12}, X_{22})$
- **8.** $B(X_{12}, U_{11}, V_{22})$

C(X, U, V)

- 1. if X is a small matrix then C-loop(X, U, V)
- 2. else
- 3. parallel: $C(X_{11}, U_{11}, V_{11}), C(X_{12}, U_{11}, V_{12})$ $C(X_{21}, U_{21}, V_{11}), C(X_{22}, U_{21}, V_{12})$ 4. parallel: $C(X_{11}, U_{12}, V_{21}), C(X_{12}, U_{12}, V_{22})$ $C(X_{21}, U_{22}, V_{21}), C(X_{22}, U_{22}, V_{22})$

I-DP -> Autogen -> R-DP

Loop-Parenthesis(C, n) 1. for $i \leftarrow n - 1$ downto 1 do 2. for $j \leftarrow i + 2$ to n do 3. for $k \leftarrow i$ to j do 4. $C[i, j] \leftarrow \min\begin{pmatrix} C[i, k] + C[k, j] \\ + w(i, k, j), C[i, j] \end{pmatrix}$ I-DP

Autogen

A(X)

- 1. if X is a small matrix then A-loop(X)
- 2. else
- **3.** parallel: $A(X_{11}), A(X_{22})$
- 4. $B(X_{12}, X_{11}, X_{22})$

$\mathbf{B}(X, U, V)$

- 1. if X is a small matrix then B-loop(X, U, V)
- 2. else
- **3.** $B(X_{21}, U_{22}, V_{11})$
- 4. parallel: $C(X_{11}, U_{12}, V_{21}), C(X_{22}, X_{21}, V_{12})$
- 5. parallel: $B(X_{11}, U_{11}, V_{11}), B(X_{22}, X_{22}, V_{22})$
- 6. $C(X_{12}, U_{12}, X_{22})$
- 7. $C(X_{12}, X_{11}, V_{12})$
- 8. $B(X_{12}, U_{11}, V_{22})$

C(X, U, V)

- 1. if X is a small matrix then C-loop(X, U, V)
- 2. else
- **3.** parallel: $C(X_{11}, U_{11}, V_{11}), C(X_{12}, U_{11}, V_{12})$ $C(X_{21}, U_{21}, V_{11}), C(X_{22}, U_{21}, V_{12})$
- 4. parallel: $C(X_{11}, U_{12}, V_{21}), C(X_{12}, U_{12}, V_{22})$ $C(X_{21}, U_{22}, V_{21}), C(X_{22}, U_{22}, V_{22})$

Performance analysis

Theory	Parallelism	Work-span model	
Ineory	Cache locality	Cache-oblivious model	
	Processor	Intel Sandy Bridge	
	#Cores	16	
	L1, L2, L3, RAM	32KB, 256KB, 20MB, 32GB	
	Compiler	Intel C++ Compiler	
Practice	Parallelization	Intel Cilk Plus	
	Count cache misses	PAPI	
	I-DP	Parallel + nontrivial opt.	
	Tiled I-DP	Pluto-generated + nontrivial opt.	
	R-DP	Simple impl. + trivial opt.	

R-DPs are efficient in theory

Droblem	I-C	OP R-DP					
Problem	Cache comp.	Parallelism	Cache comp.	Parallelism			
Parenthesis problem	$\Theta(n^3)$	$\Theta(n)$		$\Theta(n^{3}-\log 3)$			
Gap problem	$\Theta(n^2)$	$\Theta(n)$		$O(n \circ)$			
Floyd-Warshall's APSP		$\Theta(n^2/\log n)$	$O\left(\frac{1}{3}\left(\frac{1}{N}\right)\right)$	$\Theta(n^2/\log^2 n)$			
Protein folding	$\Theta(m^3/D)$	O(m)	$\Theta(n^{\circ}/(B\sqrt{M}))$	$\Theta(n^2/\log n)$			
Function approx.	$\Theta(n^{\circ}/B)$	$\Theta(n)$		$\Theta(n^{3-\log 3})$			
Matrix multiplication*							
Multi-instance Viterbi	$\Theta(n^3t/B)$	$\Theta(n^2)$	$\Theta\left(n^{3}t/(B\sqrt{M})\right)$	$\Theta(n^2)$			
LCS / edit distance							
Spoken word recog.		O(m)		$\Omega(m^2 - \log 3)$			
Binomial coefficient		$\Theta(n)$		$O(n \circ)$			
Bitonic TSP	$\Theta(n^2/B)$					$\Theta(n^2/(BM))$	
Bubble sort*				$\Theta(m)$			
Selection sort*		Θ(1)		$\Theta(n)$			
Insertion sort*				$\Theta(n^{2-\log 3})$			

B =block size, M =cache size, * =non-DP problems

R-DPs are efficient in practice

R-DPs are efficient in practice

Problem	Speedup
Parenthesis	18x
Gap	17×
Floyd-Warshall's APSP	6х

$$n = 8192$$
 and $p = 16$ speedup w.r.t. highly parallel and optimized I-DP

Overview of our work

Problem-solving technique

