Hybrid Analysis and its Application to Thread Level Parallelization

Lawrence Rauchwerger



Thread (Loop) Level Parallelization

- Thread level Parallelization
 - Extracting parallel threads from a sequential program
- Manual
 - Very expensive for sequential legacy code: MPI OpenMP
- Automatic
 - So far not very successful due to:
 - weak static analysis
 - Decisions depend on dynamic values (input dependent)

What is a parallel loop ?



Parallelizable loop







flow, $i \rightarrow i+1$

Static Data Dependence Analysis: Linear Reference Patterns



- Geometric view: Polytope model
 - Some convex body contains no integral points
- Simplified solutions: GCD Test, Banerjee Test etc
 - Potentially overly conservative
- General solution: Presburger formula decidability
 - Omega Test: Precise, potentially slow

Restricted to linear addressing and control: mostly small kernels ⁴

Static Data Dependence Analysis: Nonlinear Reference Patterns



Paraso

No known solution in general Common cases:

Indirect access (subscripted subscripts)

Symbolic Unknowns

Uninterpreted Function

Collard, Feautrier '95)

Symbols (Pugh, Wonnacott '95)

5

• Fuzzy Array Dataflow (Barthou,

- Control dependence on data values
- Recurrence without closed form

- Linear Approximation
 - Maslov '94; Creusillet, Irigoin '96
 - Range Test (Blume, Eigenmann '94)
 - Monotonicity (Wu, Cohen, Padua '01)

- User assistance

- Uninterpreted Function Symbols (Pugh, Wonnacott '95)
- SUIF Explorer (Liao '00)

Most nonlinear cases not solvable statically: need run time analysis

Alternative: Run-time Analysis



```
READ *, N
DO j=1,N
a(j)=a(j+40)
ENDDO
```

Linear, very simple, but not decidable statically!

Solution: LRPD Run-time Test (Rauchwerger and Padua '95)

Instrument all relevant memory references

Analyze the resulting trace at run time



Compile-time vs. Run-time

Compile Time

PROs

No run-time overhead

CONs: too conservative when

- Input/computed values
 - Indirection
 - Control
- Weak symbolic analysis
 - Complex recurrences
- Impractical symbolic analysis
 - Combinatorial explosion

Run-time, reference by reference

• PROs

- Always finds answers

CONs

- Run-time overhead proportional to dynamic memory reference count
- Unnecessary work: Ignores partial compile-time results

Why Did We Fail ?

- Static Analysis cannot be sufficient (weak and/or input sensitive)
 - Dynamic Analysis (misunderstood as only speculation) is not a substitute for Static Analysis (too costly)
 - No program level representation suitable for representing statically unanalyzable code



Framework: Hybrid memory reference analysis Application: Automatic parallelization





Aggregation of Linear References Across an Iteration Space



LMAD = Linear Memory Access Descriptor (Hoeflinger '98) Multidimensional, strided intervals

Gate Operator Postpone Analysis Failure

The truth value of (x=0) is not known



Recurrence Operator

Paraso

Postpone Static Analysis Failure due to Nonlinear Reference Pattern





Uniform Set of References (USR)

Paraso

Closed under composition at program level



LMAD = Start + [Stride₁:Span₁, Stride₂:Span₂, ...]



Data Dependences



- Given:
 - Loop expression: j = 1, N
 - Per-iteration aggregated descriptors RO₁, WF₁, RW₁



- At compile-time: •
 - RO \cap WF evaluates to $\emptyset \Rightarrow$ independent
 - RO \cap WF evaluates to a set that is not empty \Rightarrow dependent
 - All other cases: run-time dependence test

Similar equations for privatization and reduction recognition





Novel Static/Dynamic Interface: Predicate DAG



Represents any dynamic condition for loop parallelization

- T = { Logical Expression, \land , \lor , \bigotimes_{\land} , \bigotimes_{\lor} , Θ , Recurrence, Call Site, Library Routine} N = { PDAG } S = PDAG
- $P = \{ PDAG \rightarrow Logical Expression \\ PDAG \rightarrow PDAG \land PDAG \\ PDAG \rightarrow PDAG \lor PDAG \\ PDAG \rightarrow PDAG \otimes Recurrence \\ PDAG \rightarrow PDAG \otimes Recurrence \\ PDAG \rightarrow PDAG \Theta Call Site \\ PDAG \rightarrow Library Routine \}$



- Expressive: arbitrary dependence questions
- Inexpensive: evaluates quickly at run-time

PDAG Extraction



Output: PDAG P

Paraso

Such that:



Optimistic: Sufficient predicate: $P \Rightarrow D = \emptyset$ Pessimistic: Necessary predicate: $D = \emptyset \Rightarrow P$, or $\overrightarrow{P} \Rightarrow D \neq \emptyset$

DO j=1,100
 A(ind₁(j)) = A(ind₂(j))
 ENDDO

Fallback

Parasol



Fallback Solutions

Trace based Speculation

Possibly Aggregated

Hybrid Analysis: Run Time

Parasol



Evaluation Methodology

- Parasol
- Automatic parallelization using Hybrid Analysis in Polaris
 - Analyses/Transformations: Dependence analysis, privatization, reduction, pushback
 - Candidate loop selection: Profiling
 - No scheduling or memory locality optimization
 - Simple dynamic mechanism to rule out very small loops



- Experiment Setup
 - SPEC 2000 2006, PERFECT, Other SPEC
 - -Dual Core CoreDuo Intel, Lenovo X60s
 - Dual Core AMD, quad socket Sun

Polaris-HA vs. Ifort Coverage: PERFECT/SPEC89/92





Polaris-HA vs. Ifort Coverage: PERFECT Benchmarks





Polaris-HA vs. Ifort Coverage: SPEC 89/92





Polaris-HA vs. Ifort Coverage: SPEC2000/06





Speedups SPEC2000/06:

Paraso

AMD dual core/quad socket



Speedups Perfect/Spec89/92: Intel Core Duo





Speedups PERFECT: Intel CoreDuo





Speedups SPEC 89/92: Intel Core Duo





Conclusions: HA + Autopar

- Hybrid memory reference analysis is a general framework for optimization
- Representation is crucial
 - -- USR
 - Closed-form representation that tolerates analysis failure
 - PDAG:
 - Input sensitivity of optimization decisions
 - Continuum of compile-time to run-time solutions
- Efficient automatic parallelization
 - Good Speedups on FP benchmark applications

http://parasol.tamu.edu/compilers/ha