

Automatic Partitioning of Sequential Applications Driven by Domain-Independent Kernels

José M. Andi3n, Manuel Arenaz, and Juan Touri3o

Computer Architecture Group
Department of Electronics and Systems
University of A Coru3a
Spain



Index



- Introduction: Motivation & Foundations
- Domain-Independent Kernel-based IR
 - Kernel-based Data Dependence Graph (KDDG)
 - Kernel-based Control Flow Graph (KCFG)
- Automatic Partitioning
- Case Studies
 - Sobel Edge Filter
 - EQUAKE from SPEC CPU2000
- Conclusions & Future Work

Index



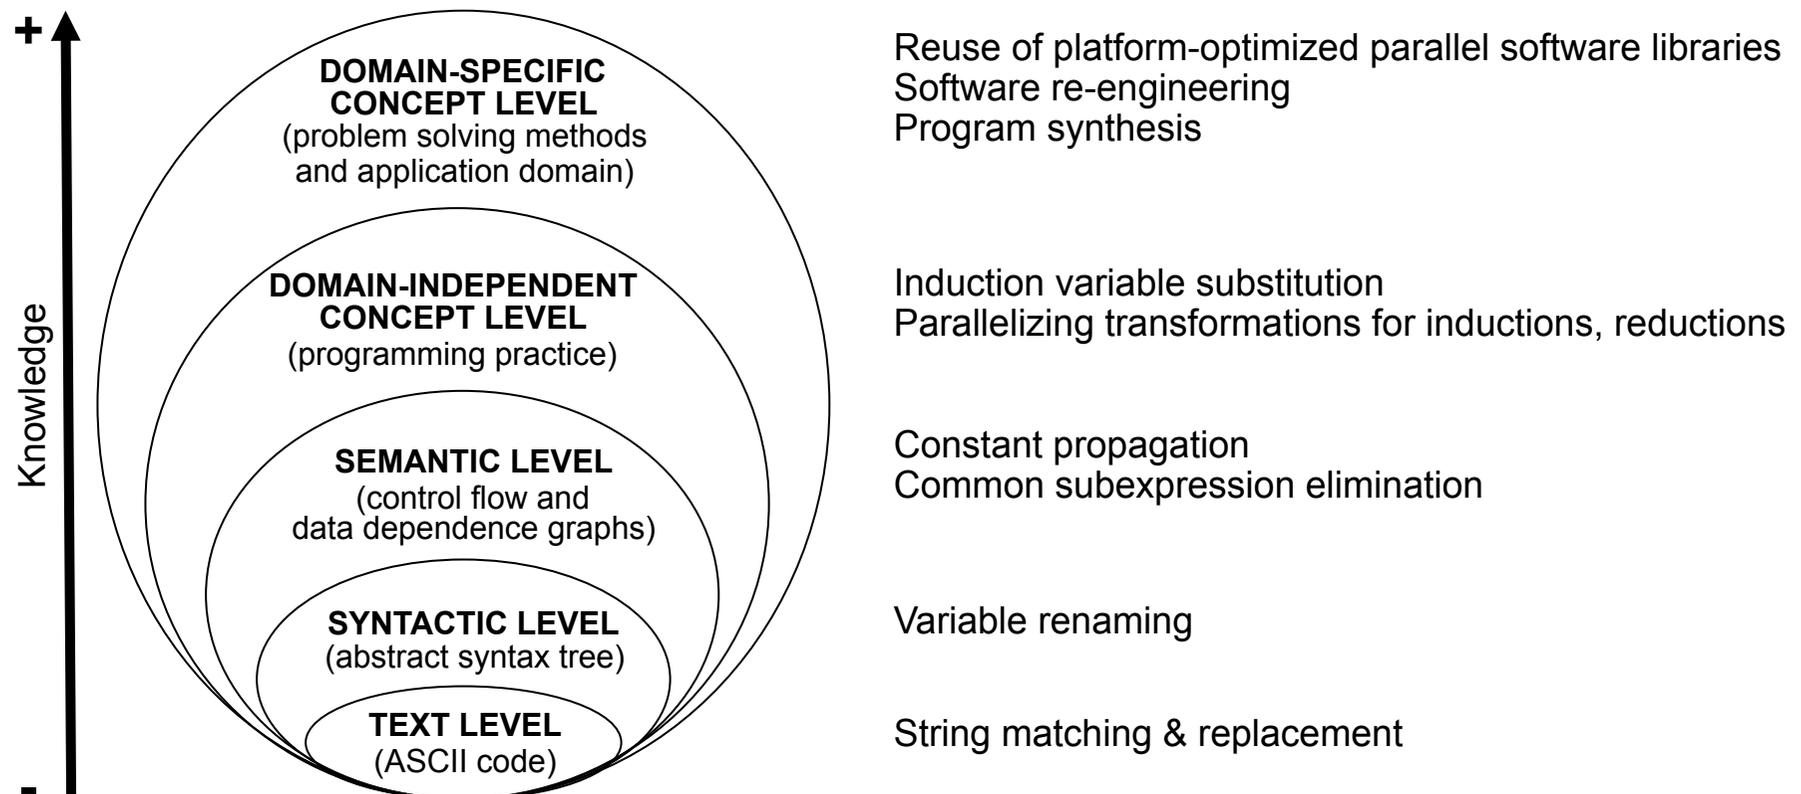
- **Introduction: Motivation & Foundations**
- Domain-Independent Kernel-based IR
 - Kernel-based Data Dependence Graph (KDDG)
 - Kernel-based Control Flow Graph (KCFG)
- Automatic Partitioning
- Case Studies
 - Sobel Edge Filter
 - EQUAKE from SPEC CPU2000
- Conclusions & Future Work

Motivation



- Emergence and widespread use of multicore and manycore systems
- Partitioning of sequential applications
 - By the developer
 - Domain-specific languages
 - Complex and time-consuming
 - In-depth knowledge about application and computer architecture
- We propose
 - New kernel-based IR
 - Automatic approach to application partitioning

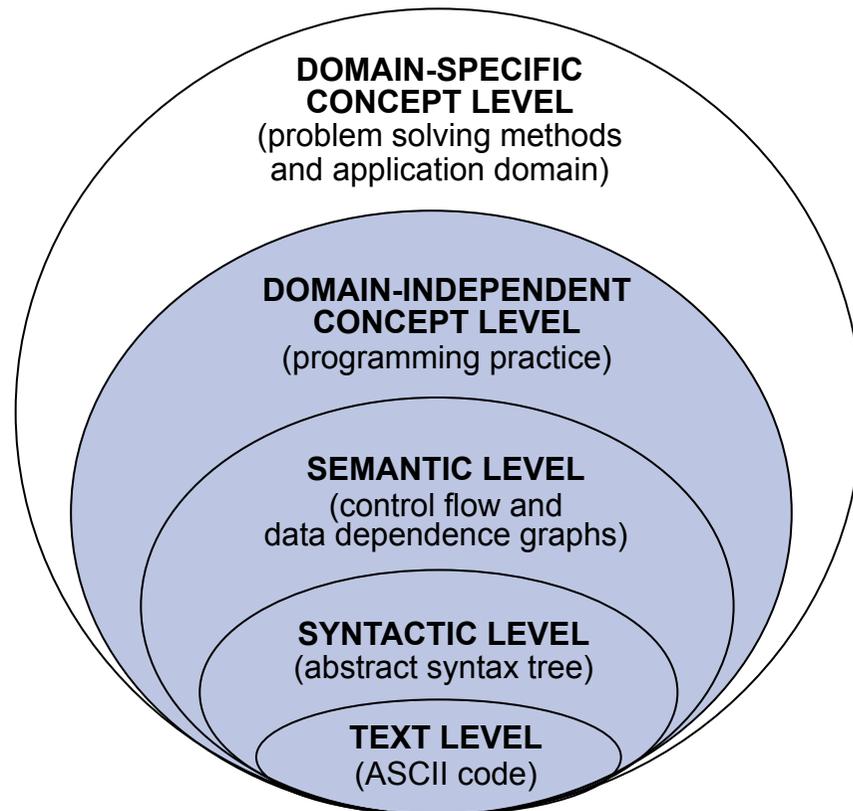
Domain-Independent Computational Kernels



The XARK Compiler Framework



- General and extensible solution for automatic kernel recognition at the domain-independent concept level.
- Properties:
 - Completeness: scalars/arrays/pointers, ifs-endifs
 - Robustness: different versions of a kernel
 - Delocalization: statements spread over the source code
 - Uniqueness: one code, one kernel
 - Extensibility: user-defined kernels



M. Arenaz, J. Touriño and R. Doallo: "XARK: An eXtensible framework for Automatic Recognition of computational Kernels", *ACM Trans. Program. Lang. Syst.*, 30(6):1-56, October 2008

Kernel Recognition



```
for(i) {
  Anext = ...
  sum0 = A[Anext]...
  Anext++;
  while(Anext) {
    col = Acol[Anext];
    sum0 += A[Anext]...
    w[col][0] += ...
    Anext++;
  }
  w[i][0] += sum0;
}
```

Excerpt of smvp() from EQUAKE



Forest of ASTs + DDG + CFG

Index



- Introduction: Motivation & Foundations
- **Domain-Independent Kernel-based IR**
 - Kernel-based Data Dependence Graph (KDDG)
 - Kernel-based Control Flow Graph (KCFG)
- Automatic Partitioning
- Case Studies
 - Sobel Edge Filter
 - EQUAKE from SPEC CPU2000
- Conclusions & Future Work

Kernel-based Data Dependence Graph (KDDG)

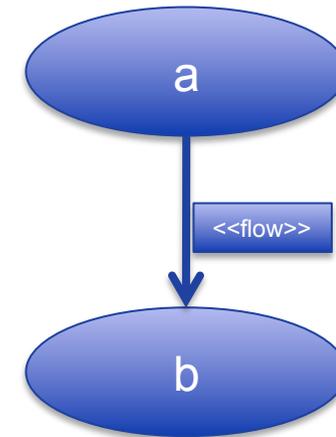


- KDDG = $\langle N, E \rangle$
- N = set of kernels $K(x_1 \dots x_n)$
 - $K(x_1 \dots x_n)$ = sentences $s_1 \dots s_n$ that define $x_1 \dots x_n$ in Gated Single Assignment (GSA) form
 - Header of K: s^h that dominates remaining statements
 - Latch of K: s^l that posdominates remaining statements
- E = set of edges $K(x_1 \dots x_n) \rightarrow K(y_1 \dots y_m)$
 - Connect statements of different kernels

Kernel-based Control Flow Graph (KCFG)



- $K(x_1 \dots x_n) \rightarrow K(y_1 \dots y_m)$ is a flow dependence if a dominance relationship exists
- K_1 dominates K_2 if and only if for all s^2 in K_2 there exists s^1 in K_1 such that
 1. s^1 and s^2 in same BB, then s^1 precedes s^2or
 2. s^1 in BB_1 and s^2 in BB_2 , then BB_1 dominates BB_2 in Dominator Tree



```
a = 5;
b = a + 1;
```

```
a = 5;
if (c == 0) {
  b = a + 3;
} else {
  b = a + 1;
}
```

```
if (c == 0) {
  a = 5;
  b = a + 3;
} else {
  a = 2;
  b = a + 1;
}
```

```
if (c == 0) {
  a = 5;
  b = a + 3;
} else {
  b = 2;
}
```

Construction of the KCFG



1. Group kernels into execution scopes

```
5: procedure COMPUTE_EXECUTION_SCOPES
6:   compute hierarchy of loops
7:   foreach kernel  $K$  in the KDDG do
8:      $bb\_dom$  = basic block of CFG that contains a stmt of  $K$  (excluding  $\mu$ -stmt)
9:     foreach statement  $s^K$  in  $K$  do
10:      if  $s^K$  is not a  $\mu$ -statement then
11:         $bb\_s^K$  = basic block of CFG that contains  $s^K$ 
12:        if  $bb\_s^K$  dominates  $bb\_dom$  then
13:           $bb\_dom = bb\_s^K$ 
14:        end if
15:      end if
16:    end for
17:     $L$  = innermost enclosing loop of  $bb\_dom$ 
18:    if  $L$  includes loop indices that address the output variable of  $K$  then
19:       $K.execution\_scope = L$ 
20:    end if
21:  end for
22:  remove non-attached loops from hierarchy
23: end procedure
```

Construction of the KCFG



2. Search for flow dependences

```
24: procedure DETECT_FLOW_DEPENDENCES
25:   foreach kernel-level dependence  $K_1 \rightarrow K_2$  of the KDDG do
26:      $L_1 = \text{execution\_scope}(K_1); L_2 = \text{execution\_scope}(K_2)$ 
27:     if ( $L_1.\text{parent} == L_2.\text{parent}$ ) && ( $L_1$  precedes  $L_2$  in the hierarchy) then
28:       mark  $K_1 \rightarrow K_2$  as flow dependence
29:     else if  $K_1.\text{latch}$  dominates  $K_2.\text{header}$  then
30:       mark  $K_1 \rightarrow K_2$  as flow dependence
31:     else
32:       mark = true
33:       foreach  $s^{K_2} \in K_2$  (excluding  $\mu$ -stmt) do
34:         dom_stmt_found = false
35:         foreach  $s^{K_1} \in K_1$  (excluding  $\mu$ -stmt) do
36:            $BB_1 = \text{basic\_block}(s^{K_1}); BB_2 = \text{basic\_block}(s^{K_2})$ 
37:           if ( $BB_1 == BB_2$ ) && ( $s^{K_1}$  precedes  $s^{K_2}$ ) then
38:             dom_stmt_found = true; break
39:           else if ( $BB_1 \neq BB_2$ ) && ( $BB_1$  dominates  $BB_2$ ) then
40:             dom_stmt_found = true; break
41:           end if
42:         end for
43:         mark = mark && dom_stmt_found
44:       end for
45:       if mark == true then
46:         mark  $K_1 \rightarrow K_2$  as flow dependence
47:       end if
48:     end if
49:   end for
50: end procedure
```

Index



- Introduction: Motivation & Foundations
- Domain-Independent Kernel-based IR
 - Kernel-based Data Dependence Graph (KDDG)
 - Kernel-based Control Flow Graph (KCFG)
- **Automatic Partitioning**
- Case Studies
 - Sobel Edge Filter
 - EQUAKE from SPEC CPU2000
- Conclusions & Future Work

Automatic Partitioning



- Kernel-based IR exposes multiple levels of parallelism
 - intra- and inter-kernel
- Modern hardware architectures also expose multiple levels of parallelism
 - cluster, multicores, Intel SSE or AMD 3DNow!
- Kernel-based IR complexity is lower than statement-based IR
 - Exhaustive search (vs. heuristics)

Automatic Partitioning



□ Initialization

- Kernels with low computational load → non-splittable → SIMD-like vector instructions
- Merge consecutive execution scopes with one flow dependence

□ Search best partitioning

- Recursive function
- Bottom-up traversal of KCFG looking for splittable kernels to be mapped to the system

Automatic Partitioning



Map kernels2arch

- Create as many tasks as needed to fill-in a given number of processing elements with a given set of splittable kernels
- Estimate cost
 - Computational load of the kernels
 - Computational capacity of processing elements
 - Amount of data that needs to be transferred
 - Synchronization
 - Etc.

Index



- Introduction: Motivation & Foundations
- Domain-Independent Kernel-based IR
 - Kernel-based Data Dependence Graph (KDDG)
 - Kernel-based Control Flow Graph (KCFG)
- Automatic Partitioning
- **Case Studies**
 - **Sobel Edge Filter**
 - EQUAKE from SPEC CPU2000
- Conclusions & Future Work

Case Study 1: Sobel



- The Sobel edge filter detects those pixels whose intensity is very different from the intensity of their neighbors
- Simple, but widely used in image processing and computer vision



Sobel

```

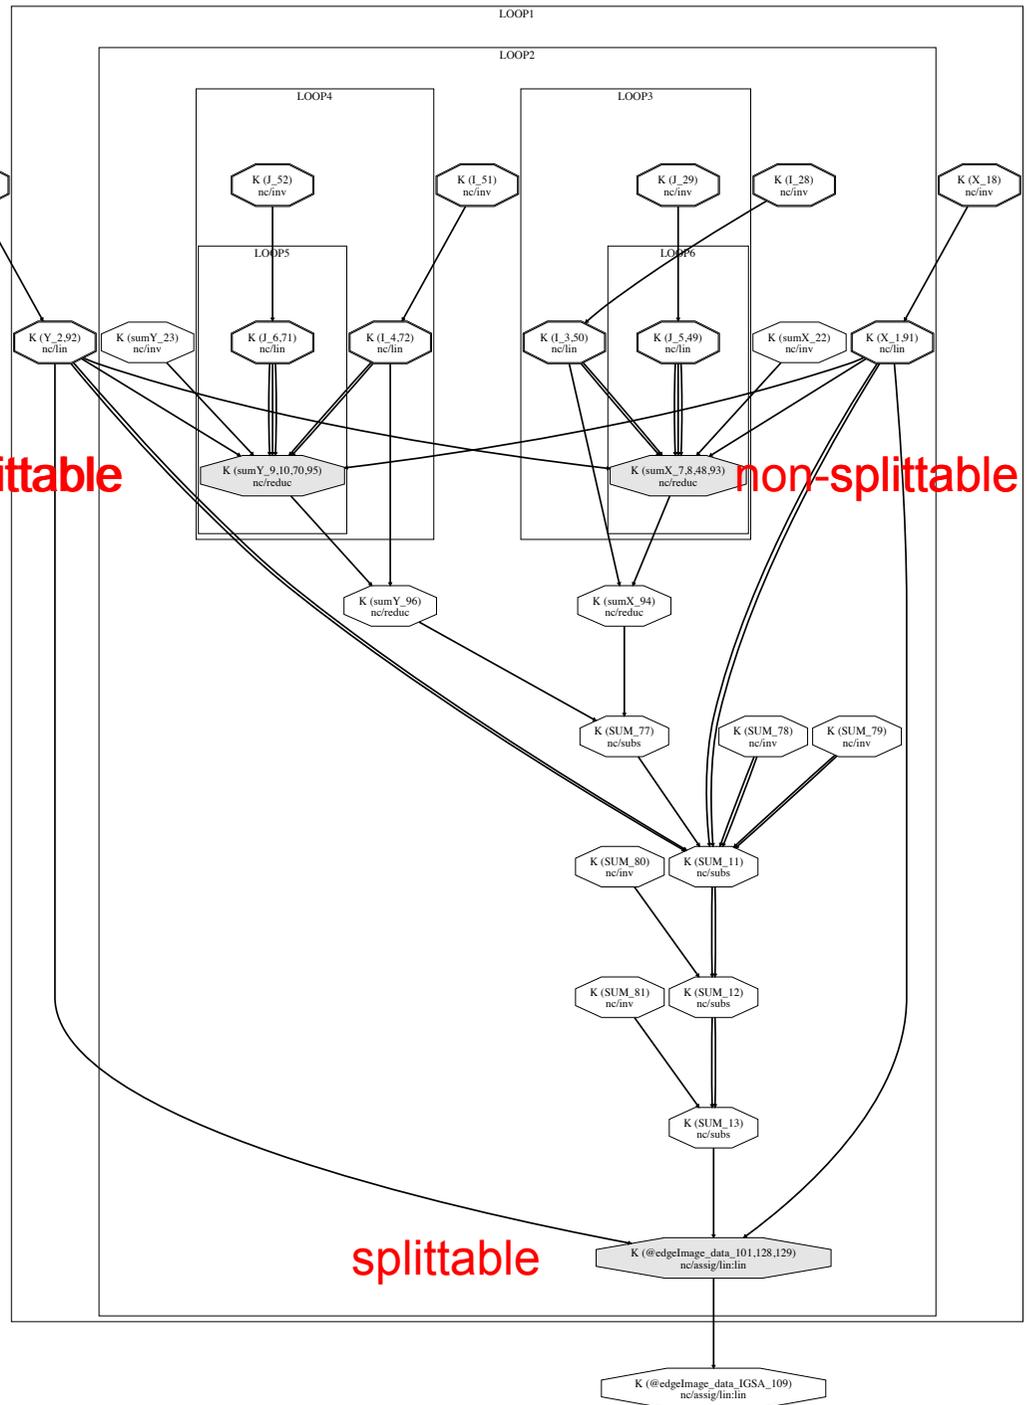
1 void gradient_aprox (long *sum, unsigned char **data,
2                     int cols, int Y, int X, int G[3][3])
3 {
4     int I, J;
5     for (I=-1; I<=1; I++)
6         for (J=-1; J<=1; J++)
7             (*sum) = (*sum) +
8                 (int)((**data) + X + I + (Y + J)*cols) * G[I+1][J+1];
9 }
10
11 int main(void)
12 {
13     int originalImage_rows, originalImage_cols,
14         edgeImage_rows, edgeImage_cols;
15     unsigned char* originalImage_data, edgeImage_data;
16     int X, Y, I, J, GX[3][3], GY[3][3];
17     long sumX, sumY, SUM;
18
19     for (Y=0; Y<=(originalImage_rows-1); Y++) {
20         for (X=0; X<=(originalImage_cols-1); X++) {
21             sumX = 0;
22             sumY = 0;
23
24             if (Y==0 || Y==originalImage_rows-1)
25                 SUM = 0;
26             else if (X==0 || X==originalImage_cols-1)
27                 SUM = 0;
28             else {
29                 gradient_aprox (&sumX, originalImage_data,
30                               originalImage_cols, Y, X, GX);
31                 gradient_aprox (&sumY, originalImage_data,
32                               originalImage_cols, Y, X, GY);
33                 SUM = abs(sumX) + abs(sumY);
34             }
35             if (SUM>255) SUM=255;
36             if (SUM<0) SUM=0;
37
38             *(edgeImage_data + X + Y*originalImage_cols) =
39                 255 - (unsigned char)(SUM);
40         }
41     }
42 }

```

non-splittable

non-splittable

splittable



Index



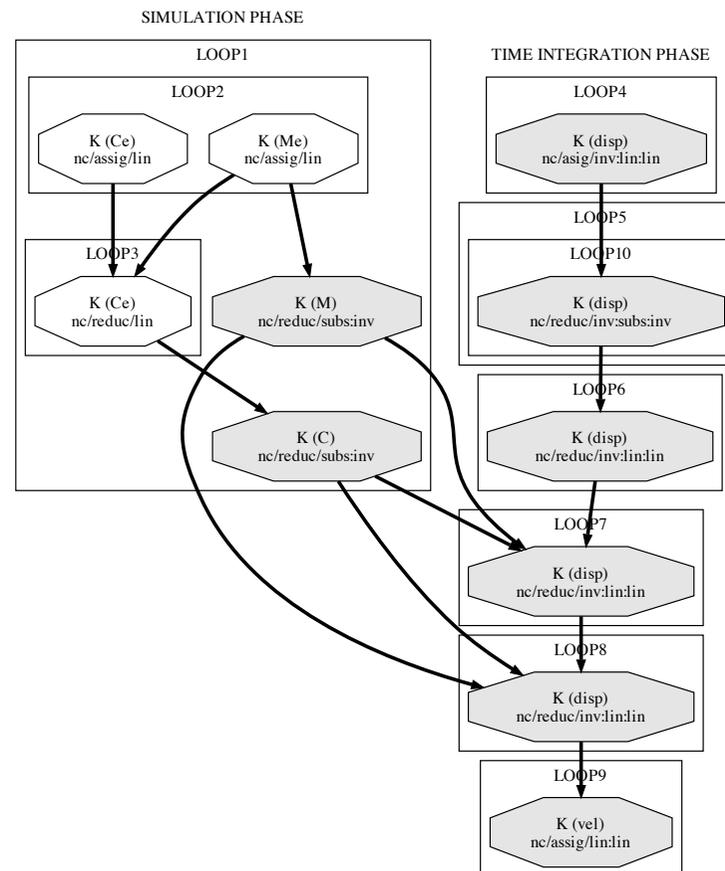
- Introduction: Motivation & Foundations
- Domain-Independent Kernel-based IR
 - Kernel-based Data Dependence Graph (KDDG)
 - Kernel-based Control Flow Graph (KCFG)
- Automatic Partitioning
- **Case Studies**
 - Sobel Edge Filter
 - **EQUAKE from SPEC CPU2000**
- Conclusions & Future Work

Case Study 2: EQuAKE



- An example of full-scale application is EQuAKE, from SPEC CPU2000
- Simulation of seismic waves in large, highly heterogeneous valleys
- Finite element method
 - Simulation phase
 - Time integration phase
- 70 % of execution time is consumed by smvp()

Case Study 2: EQUAKE



Initialization
merges the
execution scopes
LOOP4-9

Index



- Introduction: Motivation & Foundations
- Domain-Independent Kernel-based IR
 - Kernel-based Data Dependence Graph (KDDG)
 - Kernel-based Control Flow Graph (KCFG)
- Automatic Partitioning
- Case Studies
 - Sobel Edge Filter
 - EQUAKE from SPEC CPU2000
- **Conclusions & Future Work**

Conclusions & Future Work



- Definition of a new kernel-based IR
 - Domain and codification-style independent
 - Inspired by standard statement-based IRs
 - Exposes multiple levels of parallelism
- Partitioning algorithm takes advantage of multiple levels in IR and architecture
 - Kernel-based IR allows exhaustive search
- Work in progress
 - Port of XARK from Polaris to GCC
 - From F77 to C, Fortran, and even C++, Java...
 - XARK is built on top of GSA form
 - Interprocedural GSA on top of GIMPLE-SSA
 - Implementation of automatic partitioning algorithm

Conclusions & Future Work



□ Future Work

- Obtention of hardware characteristics
- Estimation of the cost of a partition
- Experimental evaluation with representative interprocedural implementations of well-known benchmarks

Automatic Partitioning of Sequential Applications Driven by Domain-Independent Kernels

José M. Andi3n, Manuel Arenaz, and Juan Touri3o

Computer Architecture Group
Department of Electronics and Systems
University of A Coru3a
Spain

