

Distributed Real-Time Processing of Multimedia Data with the P2G Framework

Paul B. Beskow, Håvard Espeland, Håkon K. Stensland, Preben N. Olsen, Ståle Kristoffersen
Espen A. Kristiansen, Carsten Griwodz, Pål Halvorsen
Simula Research Laboratory, Norway and IFI, University of Oslo, Norway

P2G is a framework designed to integrate concepts from modern batch processing frameworks into the world of real-time multimedia processing, where we seek to scale transparently with the available resources. P2G consists of a compiler and run-time that analyzes dependencies dynamically and merges or splits kernel instances based on resource availability and performance monitoring.

KERNEL LANGUAGE

fetch <field F> (age A) [index X] ... [index Z]
Extract a slice of data from a field

store <field F> (age A) [index X] ... [index Z]
Store a slice of data to a field

P2G Kernel language Example

```
Field definitions:
0 int32[] m_data age;
1 int32[] p_data age;

Kernel definitions:
0 init:
1 local int32[] values;
2 int i = 0;
3 for( i < 5; ++i )
4 {
5   put( values, i+10, i );
6 }
7 store m_data(0) = values;
8
9 plus5:
0 age a;
1 index x;
2 local int32 value;
3 fetch value = m_data(a)[x];
4 value += 5;
5 store m_data(a+1)[x] = value;
6
7 print:
0 age a;
1 local int32[] p, m;
2 fetch p = p_data(a);
3 fetch m = m_data(a);
4 for(int i=0; i < extent(p, 0); ++i)
5 {
6   cout << "p: " << get(p, i);
7   cout << "m: " << get(m, i);
8 }
9 %
```

C++ Equivalent

```
void print( int * data, int num )
{
  for( int i = 0; i < num; ++i )
    std::cout << data[i] << " ";
  std::cout << std::endl;
}

int main()
{
  int data[] = { 10, 11, 12, 13, 14 };
  int num = (sizeof data/sizeof *data);
  print( data, num );
  while(true)
  {
    for( int i = 0; i < num; ++i )
      data[i] *= 2;
    print( data, num );
    for( int i = 0; i < num; ++i )
      data[i] += 5;
    print( data, num );
  }
  return 0;
}
```

P2G ARCHITECTURE

The diagram illustrates the P2G architecture. It shows a **Master node** containing a High level scheduler, Instrumentation Manager, and Communication Manager. The Master node receives a workload and partitions it into a **Partition graph**. This graph is refined into a **Refined implicit static dependency graph (RIS-DG)**, which is then used to create a **Dynamically created directed acyclic graph (DC-DAG)**. The DC-DAG is distributed to **Execution nodes**, which contain a Low level scheduler, Instrumentation Daemon, and Communication Manager. The Execution nodes execute the workload and provide instrumentation data back to the Master node.

K-MEANS CLUSTERING

Workload description
K-means clustering is an iterative algorithm for dataset analysis, which aims to partition n datapoints into k clusters in which each datapoint belongs to the cluster with the nearest mean.

Evaluation
The k -means workload was run on a generated dataset of $n=2000$ with $k=100$. The algorithm was not run until convergence, but stopped after 10 iterations, as the point of convergence is not deterministic.

The diagram shows the workflow: **init** (generates n datapoints), **assign** (calculates distances to k centroids), and **refine** (updates centroids). The process repeats until convergence.

Kernel	Instances	Dispatch Time	Kernel Time
init	1	58.00 μ s	9829.00 μ s
assign	2024251	4.07 μ s	6.95 μ s
refine	1000	3.21 μ s	92.91 μ s
print	11	1.09 μ s	379.36 μ s

MOTION JPEG

Workload description
Motion JPEG (MJPEG) is a video coding format used using a sequence of separately compressed JPEG images.

Evaluation
The MJPEG workload was run on the standard test sequence Foreman encoded in CIF resolution. We limited the workload to process 50 frames of video. A single-thread native implementation completed in ~30 sec on the Opteron and ~19 sec on the Core i7.

The diagram shows the workflow: **read/splitYUV** (splits image into Y, U, V), **yDCT**, **uDCT**, **vDCT** (calculate DCT), and **VLC/write** (perform VLC and write to disk).

Kernel	Instances	Dispatch Time	Kernel Time
init	1	69.00 μ s	18.00 μ s
read/splitYUV	51	35.50 μ s	1641.57 μ s
yDCT	80784	3.07 μ s	170.30 μ s
uDCT	20196	3.14 μ s	170.24 μ s
vDCT	20196	3.15 μ s	170.58 μ s
VLC/write	51	3.09 μ s	2160.71 μ s