

## Programming for Performance – Part 2

#### **XMT Performance-Tuning Tools**





## The C/C++ Compiler: Initial Comments

- The XMT programmer needs to think about *algorithms* and the *compiler*.
- Running XMT programs is, practically speaking, dependent on the XMT C/C++ compiler.
- Programming the XMT for performance is a "negotiation" with the compiler.



## The XMT C/C++ Compiler: Optimizing Loop-Level Parallelism

– with thanks to Mike Ringenburg

## Outline



#### Introduction to loop parallelism

- Conditions necessary for parallelism
- Single processor, multiprocessor, and loop future parallelism

#### Pragmas to assist parallelization

- The noalias pragma and the restrict type qualifier
- The no dependence pragma
- The assert parallel pragma

**Compiler transformations to augment parallelism** 

- Scalar expansion
- Reductions/Recurrences
- Nested parallelism and loop collapse
- A parallelization example

## When Will the Compiler Parallelize a Loop?

The compiler attempts to parallelize your loops if:

- 1. It can figure out how to compute the number of iterations prior to executing the loop
- 2. It can prove that there are no dependences between iterations
- 3. There are no function calls with unknown side effects (e.g., output)
- 4. The loop has a simple structure (e.g., no multiple exits)

Pragmas are promises made by the user that help the compiler establish that these conditions hold.





#### This loop parallelizes:

```
void foo(int n) {
    int i;
    int my_array[n];
    for (i = 0; i < n; i++) {
        my_array[i] = i;
     }
    return;
}</pre>
```





#### This loop does not:

```
void foo(int *a, int *b) {
    int i;
    for (i = 0; i < 10000; i++) {
        a[i] = b[i];
    }
}</pre>
```

#### a and b may point to overlapping memory

foo(x+5000, x);



#### There are three forms of loop parallelism available on the XMT: single processor, multiprocessor, and loop futures.

- You can select a preferred mode with a compile flag
  - -par for multiprocessor (this is the default), -par1 for single processor, and -parfuture for loop futures
- In multiprocessor mode, the compiler will sometimes choose to use single processor if it judges that the amount of work and iterations are too small to justify the overhead of multiprocessor.
- You can override the parallelization mode with a pragma
- Parallelization mode is determined on a per-region basis



```
#pragma mta loop single processor
for (int i = 0; i < small_size; i++)
    a[i] = b[i];</pre>
```

- Use multiple threads on a single processor.
- Very low overheard.
- Good for shorter loops where the time saved by going parallel does not justify the expense of more heavyweight forms of parallelism.



#pragma mta loop multiprocessor
for (int i = 0; i < big\_size; i++)
 a[i] = b[i];</pre>

- Use multiple threads on multiple processors.
- Higher overhead.
- Allows you to take advantage of all the resources of the machine.

### **Loop Future Parallelism**

```
#pragma mta loop future
for (i = firstNode; i < lastNode; i++) {
    int nbr = Neighbors[i];
    int v = int_fetch_add(&Visited[nbr], 1);
    if (v == 0) BFS(nbr, A);
}</pre>
```

- Loop futures are a highly dynamic style of loop parallelism
  - For those familiar with futures, this is *not* just a loop of futures
  - Compiler still manages threads and schedules iterations
- Highest overhead form of loop parallelism
- The only form of parallelism where the number of assigned threads can increase dynamically
- Good for recursive-style loops with highly variable workloads

## Outline



#### Introduction to loop parallelism

- Conditions necessary for parallelism
- Single processor, multiprocessor, and loop future parallelism

#### Pragmas to assist parallelization

- The noalias pragma and the restrict type qualifier
- The no dependence pragma
- The assert parallel pragma

#### **Compiler transformations to augment parallelism**

- Scalar expansion
- Reductions/Recurrences
- Nested parallelism and loop collapse
- A parallelization example

## Using Pragmas to Help Find Parallelism

The XMT compiler supports a number of pragmas that can be used to give the compiler additional information about loops and the variables referenced inside them. The most commonly used are:

- 1. pragma mta assert noalias
- 2. pragma mta assert no dependence
- 3. pragma mta assert parallel

#### The compiler treats these pragmas as promises by the user

- The compiler trusts what you tell it
- If you give incorrect information, and the compiler relies on it, your program may not run correctly.

## The noalias Pragma and restrict

- Promises that the listed variables are not aliased with any other variables.
- Must appear within the scope and after the declarations of the listed variables.
- Only need to use once per variable (not once per loop).

- Promises that the listed variables are not aliased with any other variables.
- Must appear within the scope and after the declarations of the listed variables.
- Only need to use once per variable (not once per loop).
- Can also use restrict pointers to get the same affect.

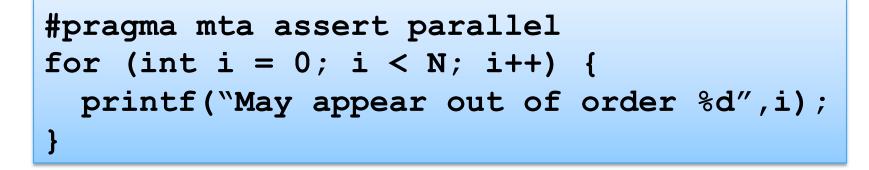
## The no dependence Pragma (or nodep)

- Promises that any memory location accessed in the loop via any variable on the no dependence list is accessed by *exactly* one iteration of the loop
- Appears immediately before a loop
- Variables must be noalias or restrict pointers

```
#pragma mta assert no dependence
for (int i = 0; i < N; i++) {
    IA[i][1] = IA[i][INDEX[i]];
}</pre>
```

- Promises that any memory location accessed in the loop via any variable on the no dependence list is accessed by *exactly* one iteration of the loop
- Appears immediately before a loop
- Variables must be noalias or restrict pointers
- Can also use with no variable list. This makes the pragma apply to all memory references in the loop (and doesn't require noalias pragmas).





- Promises that the iterations of the loop can safely be executed concurrently without any synchronization.
- Does not force the compiler to parallelize the loop, but it is a strong suggestion.
- Should only be used when other techniques to get your loop to parallelize fail. It limits the types of optimizations and transformations the compiler can perform on the loop.
  - You are only asserting that the loop is parallel as written.
  - Compiler worries that loop transformations may invalidate that.

#### Outline



#### Introduction to loop parallelism

- Conditions necessary for parallelism
- Single processor, multiprocessor, and loop future parallelism

#### Pragmas to assist parallelization

- The noalias pragma and the restrict type qualifier
- The no dependence pragma
- The assert parallel pragma

**Compiler transformations to augment parallelism** 

- Scalar expansion
- Reductions/Recurrences
- Nested parallelism and loop collapse
- A parallelization example

## **Compiler Transformations for Parallelism**

The compiler will attempt to restructure code to find or enhance parallelism:

- Scalar expansion
- Reductions
- Loop collapse

You can view the ways the compiler restructured your code in Canal (text-based) or in the Canal report of Apprentice2 (GUI-based).



This loop can not be parallelized as written because of dependences between the reads and writes of t in different iterations (writing t in one iteration may overwrite the value of t from another iteration before it is used):

```
int t;
for (i = 0; i < n; ++i) {
   t = sqrt(b[i]);
   ...
   a[i] = t + 5;
}
```



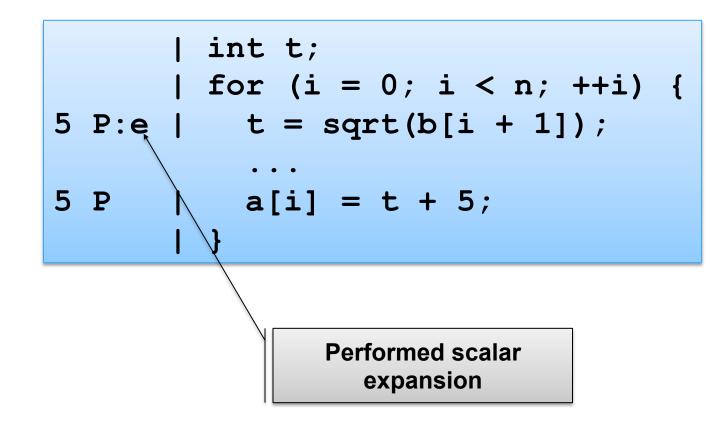
This loop can not be parallelized as written because of dependences between the reads and writes of t in different iterations (writing t in one iteration may overwrite the value of t from another iteration before it is used):

```
int t;
for (i = 0; i < n; ++i) {
  t[i] = sqrt(b[i]);
   ...
  a[i] = t[i] + 5;
}
```

The compiler solves this by converting the scalar integer t into an array of integers



# Viewing this loop in the Canal report of Apprentice2, we see:



### Reductions



The compiler attempts to recognize loops that calculate sums, products, minimums, and maximums over an array. E.g.:

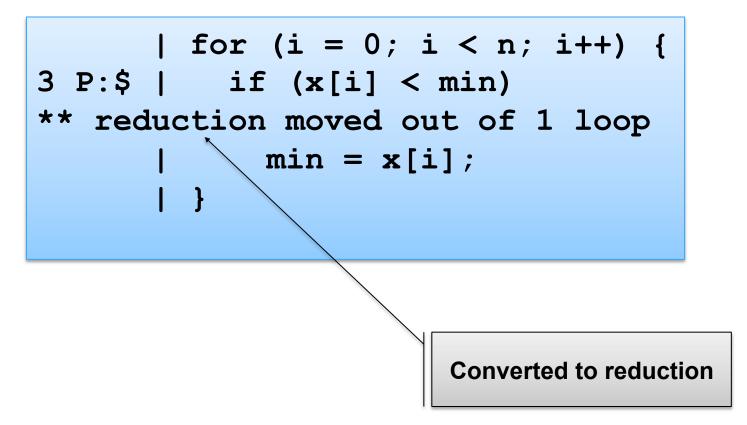
#### The compiler converts these to reductions

- Each thread computes the min/max/sum/product over a sub-section of the array.
- Threads then combine results to determine the final value.

#### Reductions



# Viewing this in the Canal report of Apprentice2, we see:





- How do we handle nested parallel loops?
- Option 1: Go parallel for the outer loop, and then again for the inner loop.
  - Inefficient there is a significant overhead to going parallel. If we nest, then every iteration of the outer loop has to pay that overhead.
  - Limits the effectiveness of the load balancing obtained by some of the scheduling methods.

## Loop Collapse



#### **Option 2: Loop collapse.**

- Convert the nested pair of parallel loops to a single parallel loop that simulates the execution of the nested loops.
  - Create a new parallel loop to calculate the total number of iteration of the inner loop (across all iterations of the outer loop).
  - Convert the pair of loops into a single loop where each iteration corresponds to a distinct outer/inner iteration pair.

#### Often a big performance win.

}

```
// t[i] = total # of inner loop iterations
// in first i iterations of outer loop
t[0] = 0;
for (i = 0; i < size_x; i++)
  t[i + 1] = t[i] + num bars[i];
for (k = 0; k < t[size x]; k++) {
  // Set i to index of largest element of t
  // less than k (use binary search)
  i = max element less than(t, k);
  j = k - t[i];
```

x[i] += bar[i + j]; // original loop body

'm' indicates loop collapse occurred

#### Outline



#### Introduction to loop parallelism

- Conditions necessary for parallelism
- Single processor, multiprocessor, and loop future parallelism

#### Pragmas to assist parallelization

- The noalias pragma and the restrict type qualifier
- The no dependence pragma
- The assert parallel pragma

**Compiler transformations to augment parallelism** 

- Scalar expansion
- Reductions/Recurrences
- Nested parallelism and loop collapse
- A parallelization example

#### **An Example**



```
bool foo(int *a, int *b, int n,
         int sought, int *old val) {
  int i;
  for (i = 0; i < n; i++) {
    if (b[i] == sought)
      break;
    a[i] = b[i];
  }
  return (i < n);
}
```

## An Example (2)



## An Example (3)



```
bool foo(int *a, int *b, int n,
         int sought, int *old_val) {
  int i;
  int found index = n;
  for (i = 0; i < n; i++) {
    if (b[i] == sought)
      if (i < found index)
        found index = i;
  }
  for (int i = 0; i < found index; i++)</pre>
    a[i] = b[i];
  return (found index < n);
}
```



## An Example (5)



```
bool foo(int *a, int *b, int n,
         int sought, int *old val) {
#pragma mta assert noalias *a
  int i;
  int found index = n;
  for (i = 0; i < n; i++) {
    if (b[i] == sought) {
      if (i < found index) {
        found index = i;
  }
  for (int i = 0; i < found index; i++)
    a[i] = b[i];
  return (found index < n);</pre>
}
```

## An Example (6)



```
#pragma mta assert noalias *a
         int i;
         int found index = n;
         for (i = 0; i < n; i++) {
3 P:$
           if (b[i] == sought) {
** reduction moved out of 1 loop
             if (i < found index) {</pre>
               found index = i;
         }
         for (int i = 0; i < found index; i++)
5 P
           a[i] = b[i];
```

# Summary



Loop parallelism is an important technique for obtaining good performance on the XMT.

The compiler will automatically parallelize loop if it can establish that it is safe to do so.

 Safe means that parallelization will preserve the correct program behavior.

Pragmas may be used to assist the compiler in proving safety.

The compiler will also attempt to aggressively transform loops to make them safe to parallelize.



## **Canal parallel annotations**

Code	Description
Р	The compiler parallelized the loop automatically.
р	An assertion caused the compiler to parallelize the loop.
D	The compiler parallelized the loop, even though it looks like there is a dependence between iterations.
L	The loop is a linear recurrence or reduction that the compiler parallelized.

# CRAY

#### **Canal sequential annotations**

Code	Description
-	The compiler could not parallelize the loop.
S	The marked statement prevented the compiler from parallelizing the loop.
S	The compiler could have parallelized the loop, but didn't because the loop had too few iterations.
Х	The compiler didn't parallelize the loop, because the loop isn't inductive.

#### **Additional annotations**

Code	Description
U	The compiler completely unrolled the loop.
е	The compiler expanded the scalar variable in the statement to a vector that it distributed across the loop iterations.
\$	The compiler implemented the memory update in this statement as an atomic operation, using full-empty bits.



### **Exercise 3**

## /mnt/lustre/Workshop/Exercise3



## Lunch



# **Overview**



## Apprentice2

- GUI application for debugging performance problems
- Consists of one or more reports based on how your program was compiled and executed
- Canal Report
  - Feedback from the compiler
  - Insight on how latent parallelism was exploited
  - Information on expected resource utilization and scheduling
- Tview Report
  - Hardware counter plots
  - Actual performance of your application
  - Runtime trap information for detecting hotspots
- Bprof Report
  - Profile tables in terms of instructions issued and memory references

# Outline



## Two sample workflows

- Provide an overview of all the visual elements of Apprentice2
- Highlight the interactions between the reports
- Parallelizing and tuning a radix sort
- Removing a hotspot from a BFS-like queue

## Write code

- Sequentially sorts on each byte of an eight byte word, starting with the least significant byte and preserving the relative ordering from partial sort to partial sort.
- Count the elements with the same value at the current byte position:

```
for (i = 0; i < size; ++i) {
    cnt[MTA_BIT_PACK(~mask, src[i])]++;
}</pre>
```

## Compile

• cc -o radix.v1 sorted.c radix.v1.c

#### Post-process

- pproc radix.v1
- Run Apprentice2
- app2 radix.v1.ap2 [additional .ap2 files]

## Modify code

- Add user trace events to further identify areas of interest
- Move the elements to their new relative positions

```
#pragma mta trace "shift elements"
for (i = 0; i < size; ++i) {
    j = pos[MTA_BIT_PACK(~mask, src[i])]++;
    dst[j] = src[i];
}</pre>
```

#### Recompile with tracing and profiling enabled

- cc -trace -profile -o radix.v1 sorted.c radix.v1.c
- Run the application and post-process
- mtarun -trace -pproc radix.v1

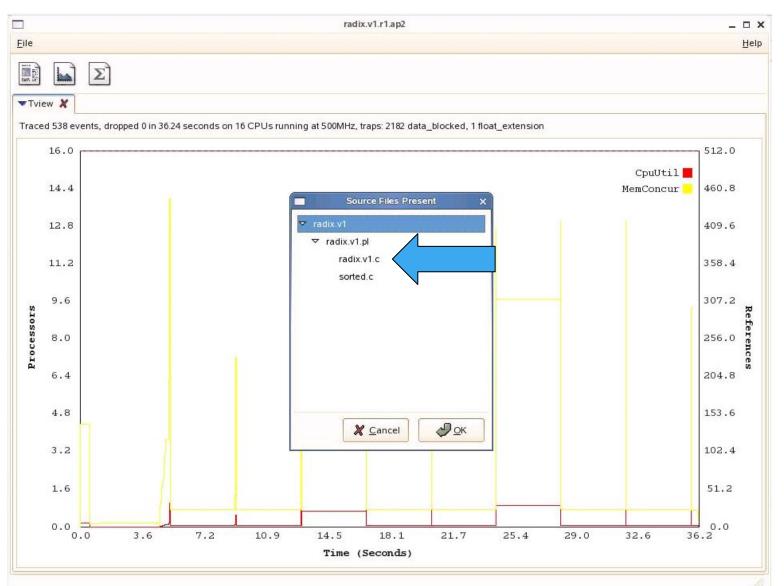
## Run Apprentice2

• app2 sorted.v1.ap2





Source Files Pr rmat2file x	
rmat2file.pl	
./mtgl/dynamic_array.h	(graph.o)
./mtgl/dynamic_array.h	(rmat2file.o)
./mtgl/graph_adapter.hpp	
./mtgl/graph.h	
./mtgl/mtgl_adapter.hpp	
./mtgl/mtgl_boost_graph.hpp	
/mtgl/rand_fill.hpp	
./mtgl/util.h	(graph.o)
./mtgl/util.h	(rmat2file.o)
./mtgl/xmt_hash_table.hpp	
/opt/mta-pe/6.4.0/include/c++/exception	(graph.o)
/opt/mta-pe/6.4.0/include/c++/exception	(rmat2file.o)
/opt/mta-pe/6.4.0/include/c++/iostream	(graph.o)
/opt/mta-pe/6.4.0/include/c++/iostream	(rmat2file.o)
/opt/mta-pe/6.4.0/include/c++/new	(rmat2file.o)
/opt/mta-pe/6.4.0/include/c++/new	(graph.o)
/opt/mta-pe/6.4.0/include/c++/stdexcept	(graph.o)
/opt/mta-pe/6.4.0/include/c++/stdexcept	(rmat2file.o)
/opt/mta-pe/6.4.0/include/c++/typeinfo	(graph.o)
/opt/mta-pe/6.4.0/include/c++/typeinfo	(rmat2file.o)
/opt/mta-pe/6.4.0/include/c++/xiosbase	(graph.o)
/opt/mta-pe/6.4.0/include/c++/xiosbase	(rmat2file.o)
/opt/mta-pe/6.4.0/include/c++/xlocale	(graph.o)
/opt/mta-pe/6.4.0/include/c++/xlocale	(rmat2file.o)
/opt/mta-pe/6.4.0/include/c++/xstring	(graph.o)
/opt/mta-pe/6.4.0/include/c++/xstring	(rmat2file.o)
test/graph.cpp	
test/rmat2file.cpp	
test/static_graph.c	
	X Cancel





	radix.v1.r1.ap2				_ 🗆 🗙
<u>F</u> ile					Help
Tview 🗶 🔻 Ca	anal 🗶				
Line Notes	Code	Issues	MemRefs	Count 1	raps 🔺
	#include "radix.h"				=
	<pre>#include <stdio.h></stdio.h></pre>				
*	unsigned* radix_sort(unsigned* array, unsigned size) {	8650583	8250137	42	
5	unsigned* src = array;	1	1	1	
	<pre>unsigned* dst = (unsigned*)malloc(size*sizeof(unsigned));</pre>	8	5	1	
	unsigned buckets = 1 << 8;			1	
	unsigned* cnt = (unsigned*)malloc(buckets*sizeof(unsigned));	259	256	1	
▶ Loops	a an a an ar a an ar a an	-		4	
Details					



	radix.v1.r1.ap2				_ 🗆 🗙
Eile	1				<u>H</u> elp
Σ 🗋					
Tview 🗶 🔻 Ca	nal 🗶 🎽				
Line Notes	Code	Issues	MemRefs	Count	Traps
	#include "radix.h"				=
	<pre>#include <stdio.h></stdio.h></pre>				
*	unsigned* radix_sort(unsigned* array, unsigned size) {	8650583	8250137	42	
5	unsigned* src = array;	1	1	1	
	<pre>unsigned* dst = (unsigned*)malloc(size*sizeof(unsigned));</pre>	8	5	1	
	unsigned buckets = 1 << 8;			1	
	unsigned* cnt = (unsigned*)malloc(buckets*sizeof(unsigned));	259	256	1	
▶ Loops	the state of the s	1		•	
Details					
Details					
<u></u>					



	radix.v1.r1.ap2				_ 🗆 🗙
Eile					Help
Σ					
Tview 🗶 🔻 Can					1000
Line Notes	Code	Issues	MemRefs	Count 1	raps
	#include "radix.h"				=
	<pre>#include <stdio.h></stdio.h></pre>				
*		8650583	8250137		
5	unsigned* src = array;	1	1	1	
	<pre>unsigned* dst = (unsigned*)malloc(size*sizeof(unsigned));</pre>	8	5	1	
				1	
		259	256	1	-
↓ Loops		1			
Details					
<pre>     Tview X Canal X     Ine Notes Code     Issues MemRefs     #include "radix.h"     #include <stdio.h>     *     unsigned* radix_sort (unsigned* array, unsigned size) {          unsigned* src = array;          unsigned* src = array;          unsigned* dst = (unsigned*)malloc(size*sizeof(unsigned));</stdio.h></pre>					
15					1



radix.v1.r1.ap2				-	
Eile				1	Help
Line Notes	Issues	MemRefs		Traps	
Columns	10	3	1		
1 - Change Font trace "count elements"	16	16	8		
Panel Actions = ); i < size; ++i) (	152434	12943			H
30 6 SI: Panel Help A BIT_PACK(-mask, src[1])]++;	24796094	8108779	8000000		H
2 - #pragma mta trace "find positions"	40243	944	1 8		
Loops	10	16	0		•
Details					
Loop 6 in radix_sort at line 29 in loop 5 Loop summary: 4 instructions, 0 floating point operations 1 loads, 1 stores, 0 reloads, 0 spills, 1 branches, 0 calls Loop 5 in radix_sort in loop 4 in parallel phase 1 dynamically scheduled, variable chunks, min size = 7					
Parallel Region 4 in radix_sort multiple processor implementation requesting at least 30 streams					
Loop 2 in radix_sort at line 22 expecting 8 iterations					



	rad <mark>ix v1 r</mark> 1 ap2				- 0
ile					He
Σ					
Tview 🗶 🔻 Car					
Line Notes	Code	Issues	MemRefs	Count	Traps
	3	10	3	1	
1 - 2 -	#pragma mta trace "count elements"	16	16	8	
	for (i = 0; i < size; ++i) (	152434	12943	3841	-
30 6 SP:\$	<pre>cnt[MTA_BIT_PACK(-mask, src[i])]++;</pre>	24796094	8108779	8000000	
	)			1	
2 -	#pragma mta trace "find positions"	40243	944	8	
1 - ▶ Loops	non [0] = 0.	40	16	•	
Details	noningential.				
Loop 6 in radix_sort Loop summary 1 loads, 1 s Loop 5 in radix_sort in parallel phas	: 4 instructions, 0 floating point operations stores, 0 reloads, 0 spills, 1 branches, 0 calls t in loop 4				
	a radix_sort sor implementation east 30 streams				
Loop 2 in radix_sort expecting 8 iter					



Tview Canal					H
Tview					
Tview					
2000 CONTRACT 10					
ine Notes C	Code	lssues	A REAL PROPERTY AND A REAL PROPERTY.	and the second second	Traps
	3	10	3	1	
1 - 2 -	#pragma mta trace "count elements"	16	16	8	
	for (i = 0; i < size; ++i) (	152434	12943	3841	
30 6 SP:\$	<pre>cnt[MTA_BIT_PACK(-mask, src[i])]++;</pre>	24796094	8108779	8000000	
	1			1	
2 -	#pragma mta trace "find positions"	40243	944	8	
1 -	nor[0] = 0.	40	16	0	3
> Loops					
	line 29 in loop 5 instructions, 0 floating point operations ores, 0 reloads, 0 spills, 1 branches, 0 calls				
oop 5 in radix_sort in in parallel phase 1 dynamically sche					
Parallel Region 4 in ra multiple processo requesting at leas	or implementation				
.oop 2 in radix_sort at expecting 8 iterati					



	radix.v1.r1.ap2				_ □ >
<u>F</u> ile					Help
Σ					
▼Tview 🗶 ▼Car	nal 🗶				
Line Notes	Code	Issues	MemRefs	Count	Traps
	3	10	3	1	
1 - 2 -	#pragma mta trace "count elements"	16	16	8	
	for (i = 0; i < size; ++i) {	152434	12943	3841	=
30 6 SP:\$	<pre>cnt[MTA_BIT_PACK(-mask, src[i])]++;</pre>	6094	8108779	8000000	
	1			1	
2 -	#pragma mta trace "find positions"	40243	944	8	
1 - ▶ Loops	nos [0] = 0.	40	16	0	•
Details	HURDON HURD				
Loop 6 in radix_sort Loop summary: 1 loads, 1 s Loop 5 in radix_sort in parallel phas	: 4 instructions, 0 floating point operations stores, 0 reloads, 0 spills, 1 branches, 0 calls : in loop 4				
Parallel Region 4 in multiple proces requesting at le	sor implementation				
Loop 2 in radix_sort expecting 8 iter					



	radix.v1.r1.ap2				_ 🗆
<u>F</u> ile					Help
Tview 🗶 🔻 Ca	nal 🗶				
Line Notes	Code	Issues	MemRefs	Count	Traps
	)	10	3	1	
1 - 2 -	#pragma mta trace "count elements"	16			
	for (i = 0; i < size; ++i) {	152434			-
30 6 SP:\$	<pre>cnt[MTA_BIT_PACK(-mask, src[i])]++;</pre>	24796094	8108779		
	)			1	
2 -	#pragma mta trace "find positions"	40243	944	8	
1 - ▶ Loops	Dog [0] = 0.	40	16	0	•
Details	manananan .				
Loop 6 in radix_sort Loop summary 1 loads, 1 Loop 5 in radix_sort in parallel phas dynamically so	: 4 instructions, 0 floating point operations stores, 0 reloads, 0 spills, 1 branches, 0 calls t in loop 4 se 1 .heduled, variable chunks, min size = 7				
	ssor implementation east 30 streams				
Loop 2 in radix_sort expecting 8 iter					

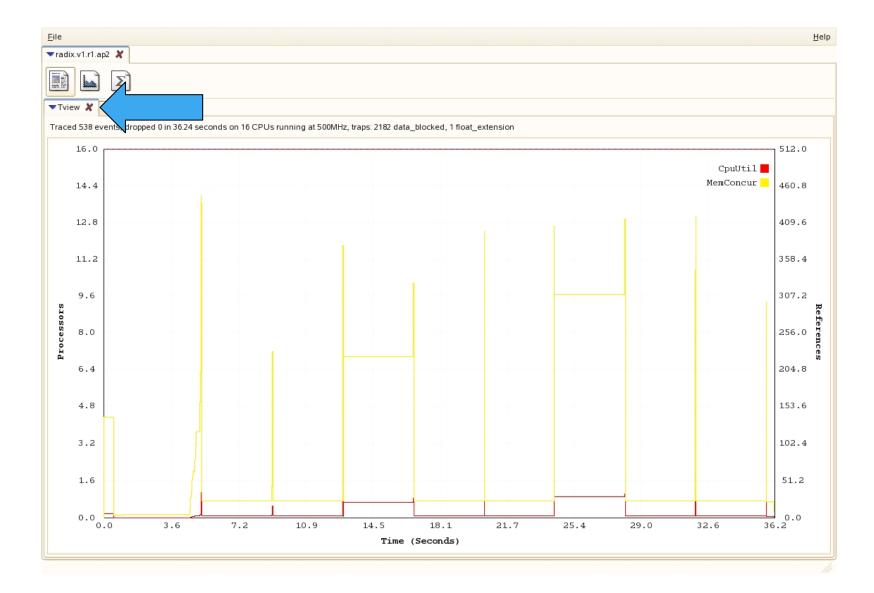


-ile	radix.v1.r1.ap2				-
🕶 Tview 🗶 🔍 Ca	inal 🗶				
Line Notes	Code	lssues	MemRefs	Count	Traps
	)	10	3	1	
1 - 2 -	#pragma mta trace "count elements"	16	16	8	
	for (i = 0; i < size; ++i) (	152434	12943	3841	
30 6 SP:\$	<pre>cnt[MTA_BIT_PACK(-mask, src[i])]++;</pre>	24796094	8108779	8000000	
	)			1	
2 -	#pragma mta trace "find positions"	40243	944	8	
1 - ♦ Loops		40	16	•	
Details	Representation -				
Loop summary 1 loads, 1 Loop 5 in radix_sor in parallel pha					
	n radix_sort ssor implementation east 30 streams				
Loop 2 in radix_sor expecting 8 ite					

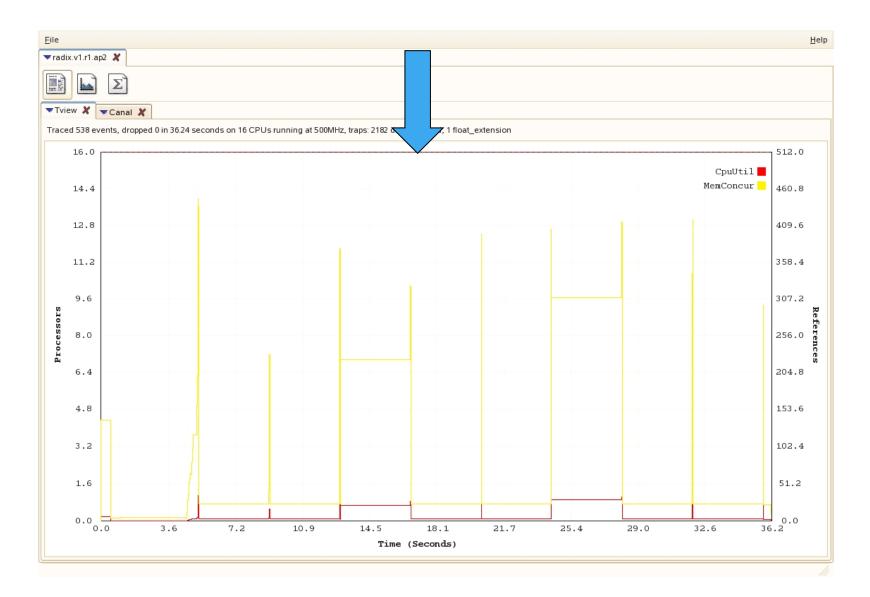
CR	ar
	10.22

				- 0
Eile				He
Tview 🗶 🔽 Canal 🗶				
Line Notes Code	Issues	MemRefs	Count	Traps
)	10	3	1	0
1				
2 - #pragma mta trace "count elements"	16	16	8	
for $(i = 0; i < size; ++i)$ (	152434	12943	3841	5
<pre>30 6 SP:\$ cnt[MTA_BIT_PACK(~mask, src[i])]++;</pre>	24796094	8108779	8000008	-
▼ Loops				
Loop 1				-
▼ Loop 2				
Loop 3				
▼ Parallel Region 4				
▼ Parallel Region 4				
▼ Loop 5				
Loop 5     Loop 6				
Loop 5      Loop 6      Loop 7				
Loop 5     Loop 6				
✓ Loop 5     Loop 7     Loop 7     Loop 8  Details Loop 6 in radix_sort at line 29 in loop 5 Loop summary: 4 instructions, 0 floating point operations 1 loads, 1 stores, 0 reloads, 0 spills, 1 branches, 0 calls Loop 5 in radix_sort in loop 4     in parallel phase 1     dynamically scheduled, variable chunks, min size = 7				
✓ Loop 5     Loop 7     Loop 7     Loop 8  Details Loop 6 in radix_sort at line 29 in loop 5 Loop summary: 4 instructions, 0 floating point operations     1 loads, 1 stores, 0 reloads, 0 spills, 1 branches, 0 calls Loop 5 in radix_sort in loop 4     in parallel phase 1     dynamically scheduled, variable chunks, min size = 7  Parallel Region 4 in radix_sort				
✓ Loop 5     Loop 7     Loop 7     Loop 8  Details Loop 6 in radix_sort at line 29 in loop 5 Loop summary: 4 instructions, 0 floating point operations 1 loads, 1 stores, 0 reloads, 0 spills, 1 branches, 0 calls Loop 5 in radix_sort in loop 4     in parallel phase 1     dynamically scheduled, variable chunks, min size = 7				

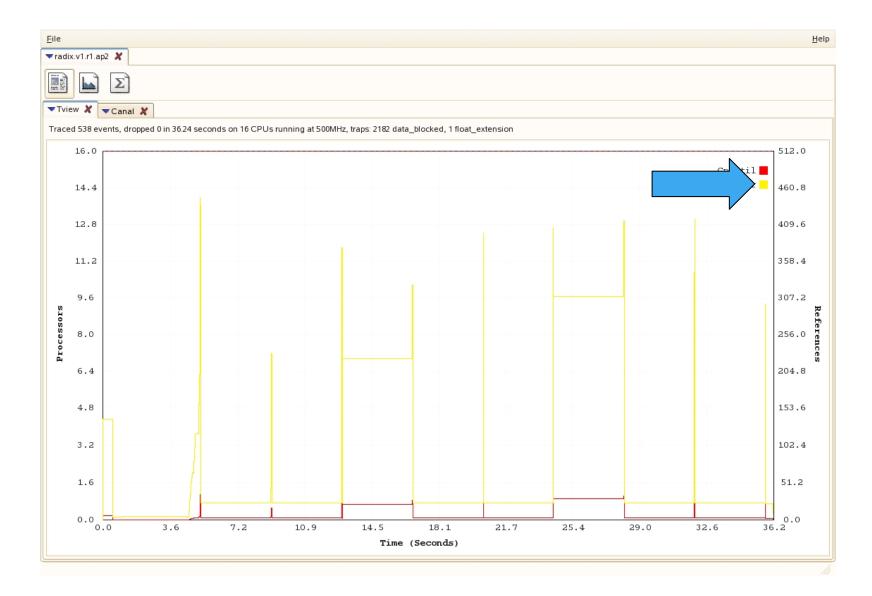


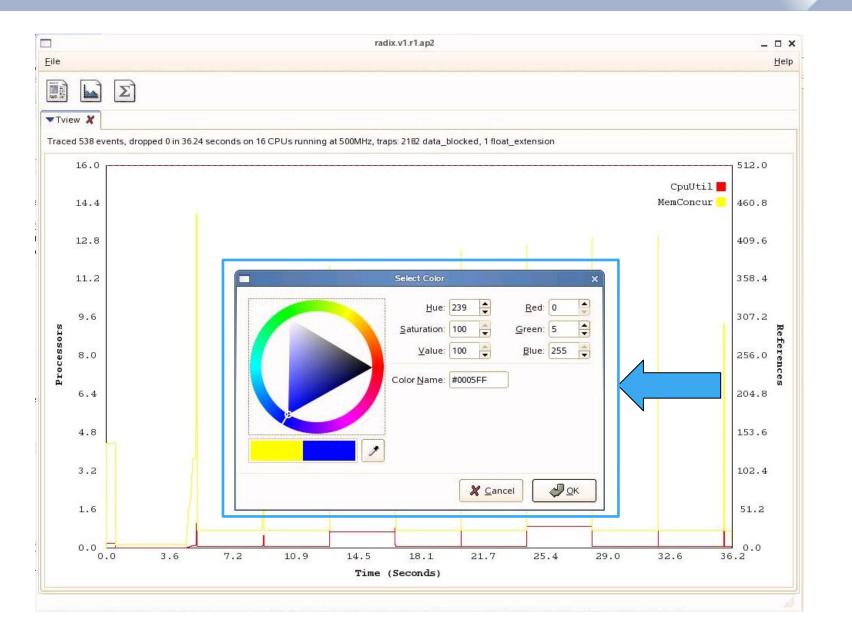


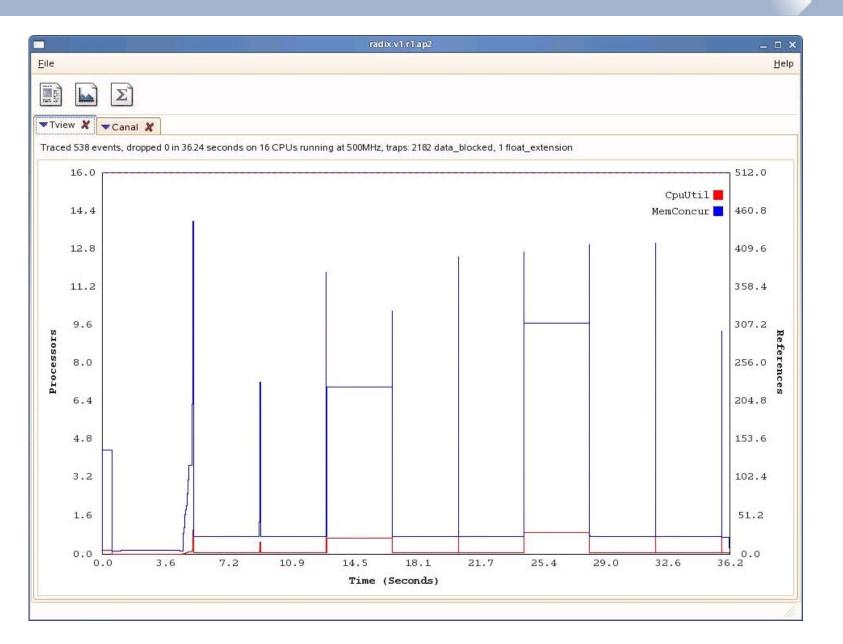


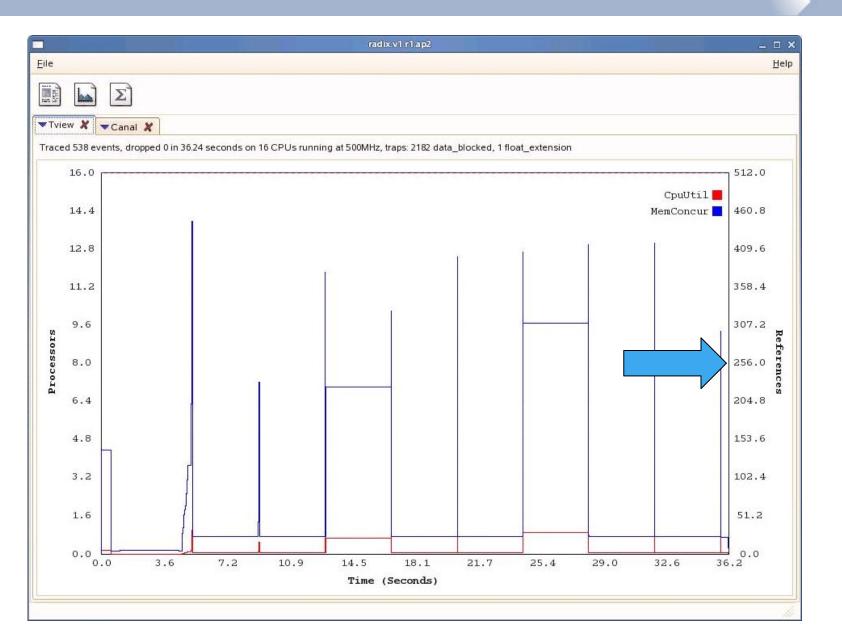


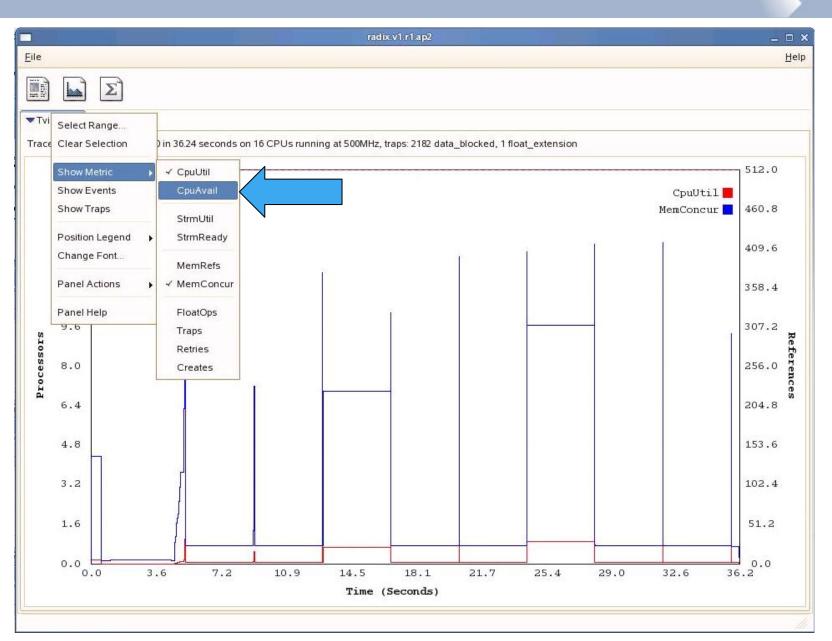


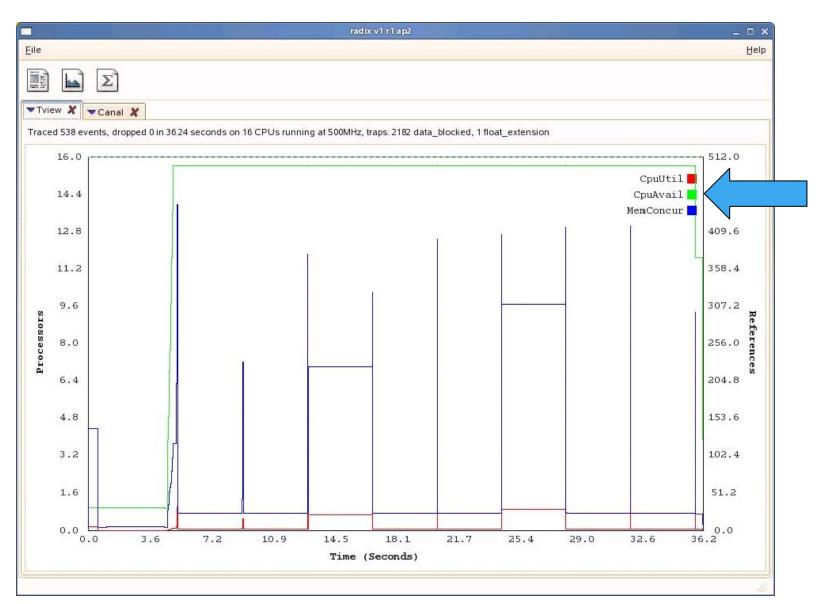




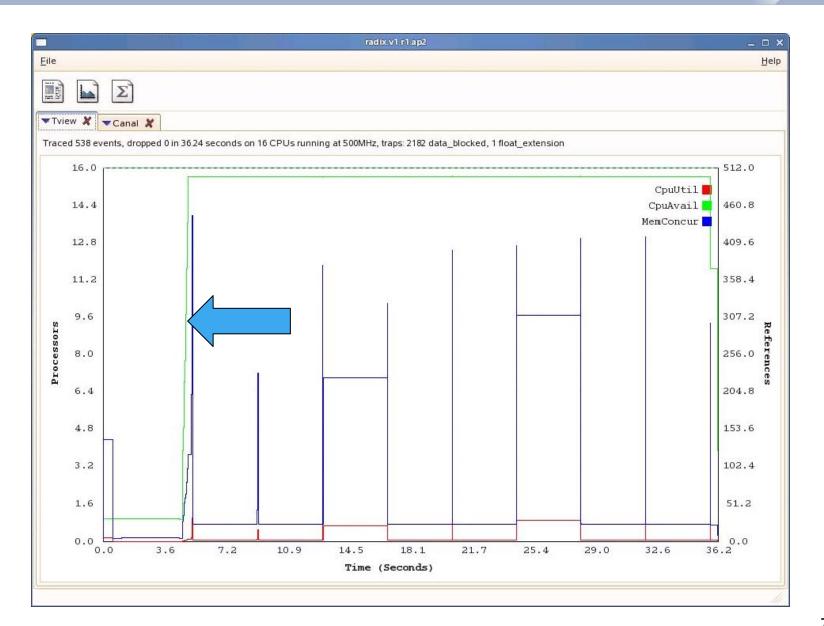


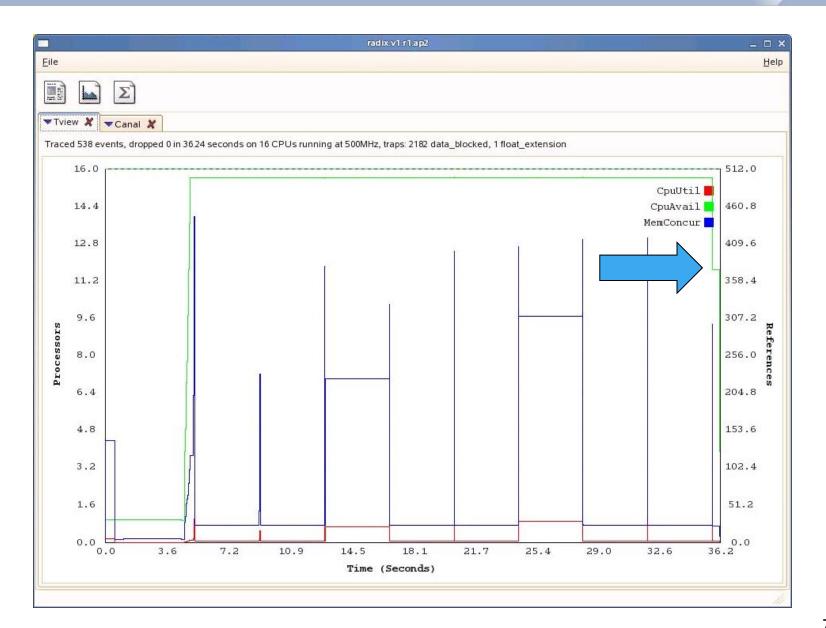




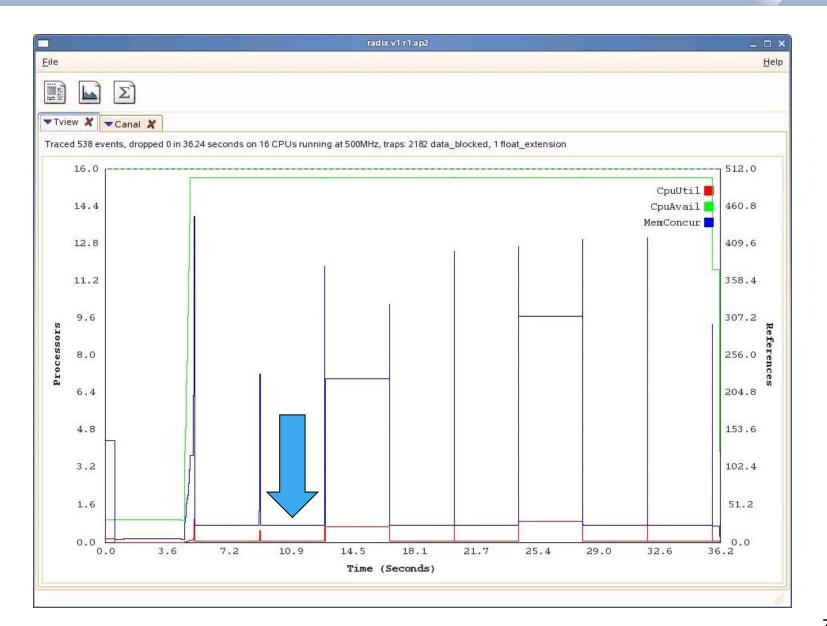










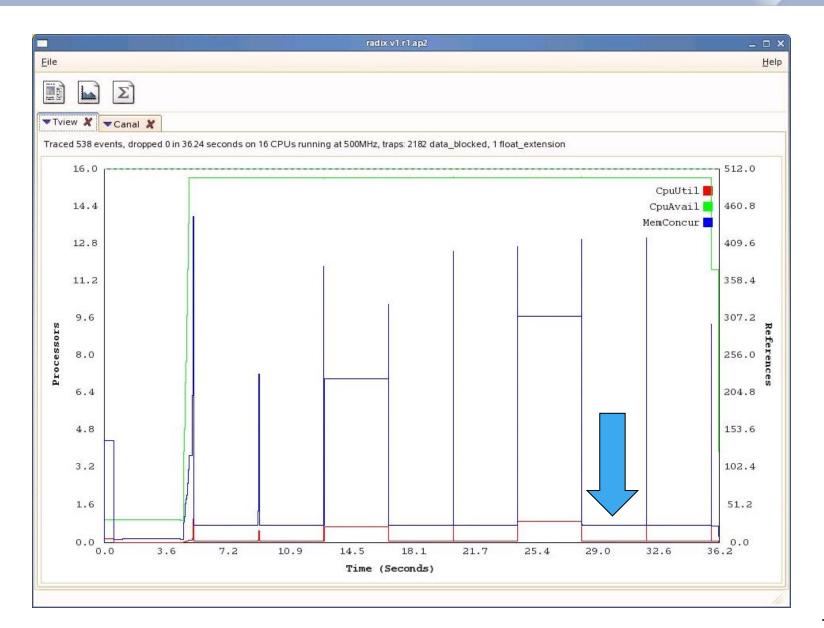




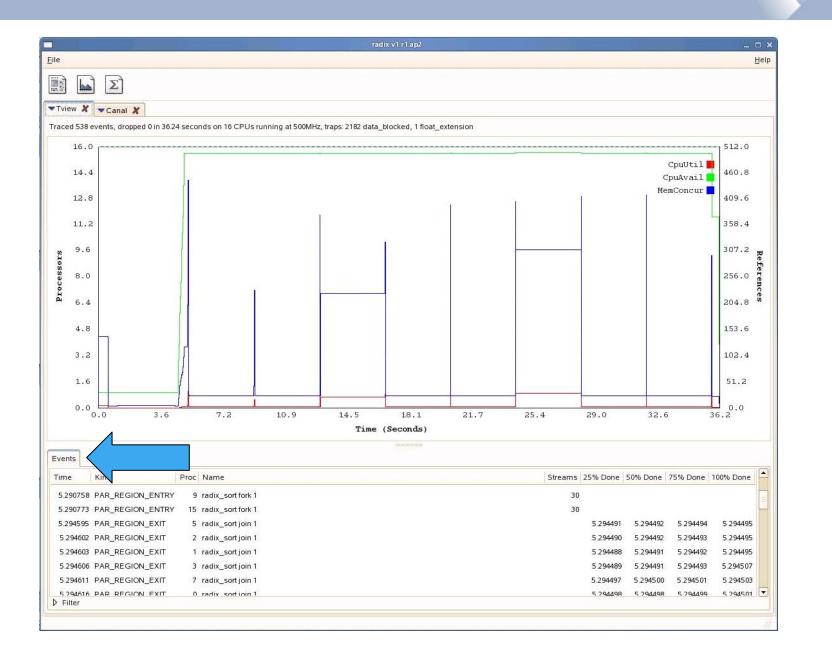


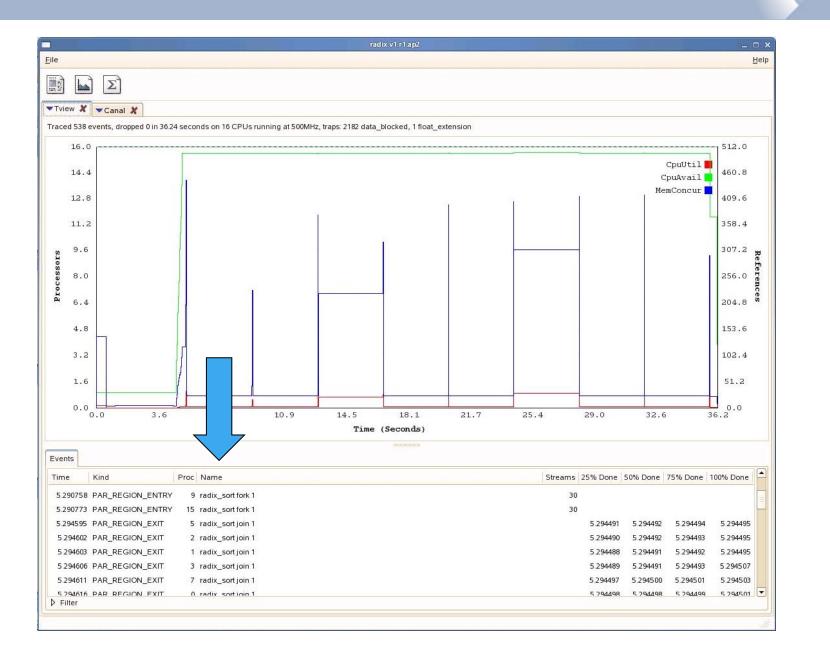


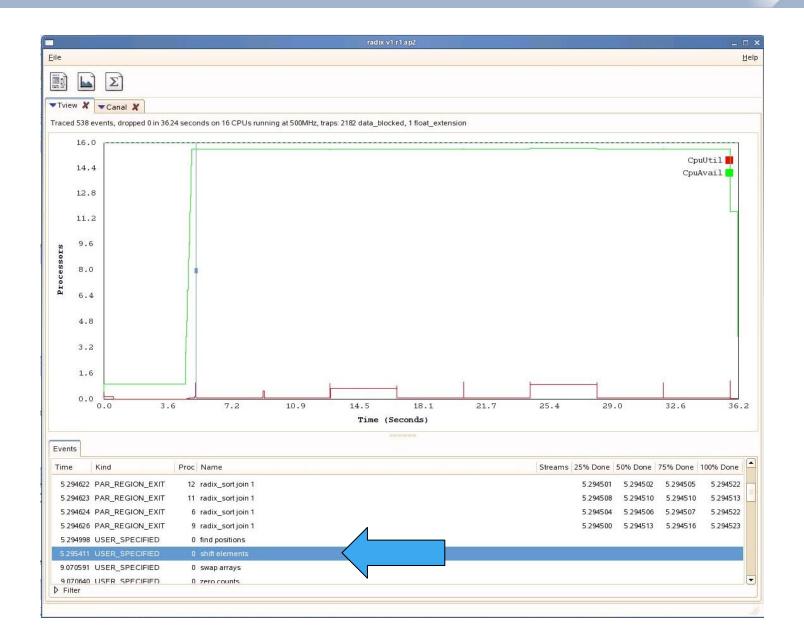


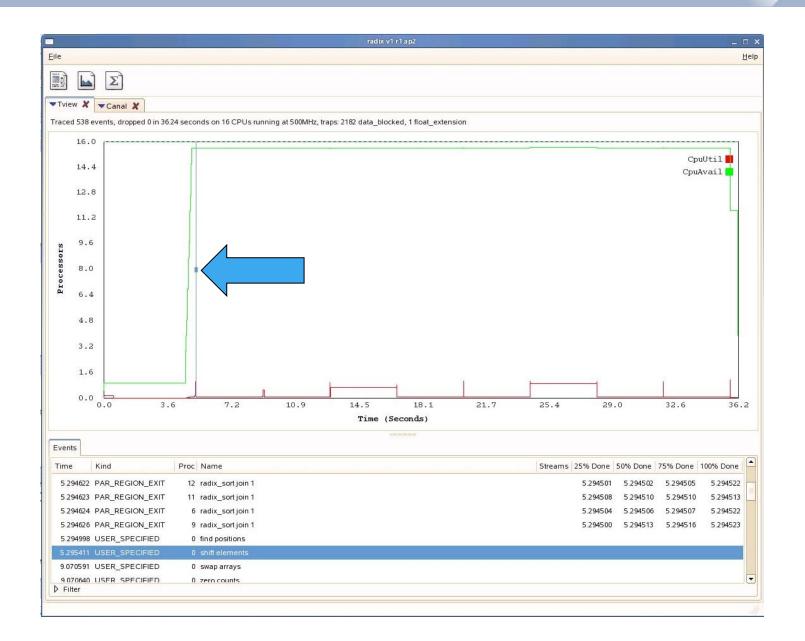


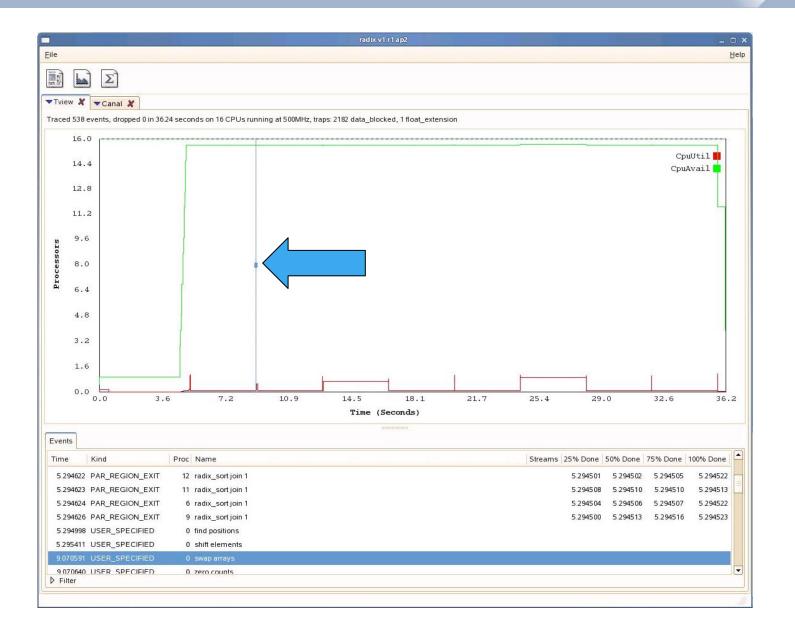


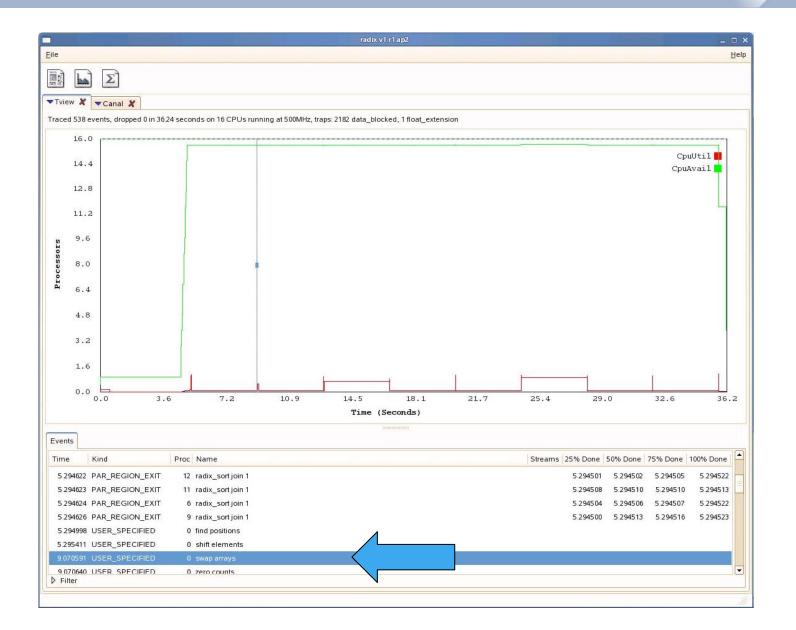












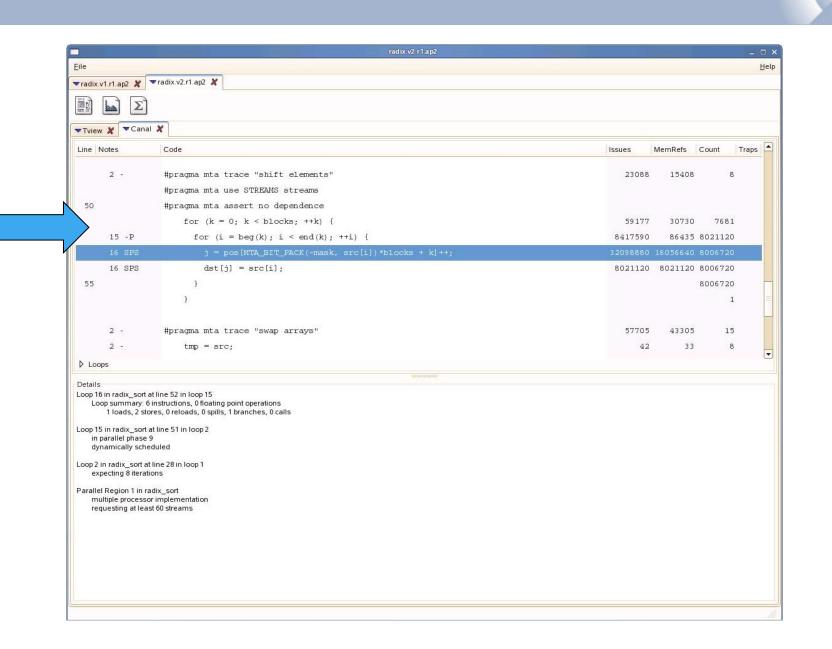
8 SS       j = pos[MTA_BIT_PACK(-mask, src[i])]++;       32000032 16000024 8000000         8 SS       dst[j] = src[i];       24000018 8000009 8000000         15       2 -       #pragma mta trace "swap arrays"       8       8         15       2 -       tmp = src;       24       8       8         2 S*       src = dst;       16       16       8       8         30       2 S*       dst = tmp;       9       8       8         30       2 S*       mask <<= 8;       8       8       8	8 SS       j = pos[MTA_BIT_PACK(-mask, src[i])]++;       32000032 16000024 8000000         8 SS       dst[j] = src[i];       24000018 8000009 8000000         }       1         2 -       #pragma mta trace "swap arrays"       8       8       8         2 -       tmp = src;       24       8       8         2 S*       src = dst;       2       16       16         2 S*       mask <<= 8;       8       8       8         0 ps	1 - 2 -	#pragma mta trace "shift elements"	2048	2048	
8 SS       dst[j] = src[i];       24000018       8000009       8000009         15 2 - #pragma mta trace "swap arrays"       8       8       8       8         2 - tmp = src;       24       8       8       8       8       8         2 S*       src = dst;       16       16       16       8<	8 SS       dst[j] = src[i];       24000018       8000009       8000009         2 -       #pragma mta trace "swap arrays"       8       8       8         2 -       tmp = src;       24       8       8         2 S*       src = dst;       25*       16       16       8         2 S*       mask <<= 8;	40			diene	
}       }         45       2 - #pragma mta trace "swap arrays"       8	}       3       8					
15       2 -       #pragma mta trace "swap arrays"       8	2 -       #pragma mta trace "swap arrays"       8	0.55		24000010	0000003	
2 -     tmp = src;     24     8     1       2 S*     src = dst;     16     16     16       2 S*     dst = tmp;     9     8     1	2 -       tmp = src;       24       8       1         2 S*       src = dst;       16       16       16       1         2 S*       dst = tmp;       9       8       1         2 S*       mask <<= 8;					
2 S*       src = dst;       16       16         2 S*       dst = tmp;       9       8         50 2 S*       mask <<= 8;	2 S*       src = dst;       16       16         2 S*       dst = tmp;       9       8         2 S*       mask <<= 8;	45 2 -	#pragma mta trace "swap arrays"	8	8	
2 S* dst = tmp; 9 8 50 2 S* mask <<= 8; 8 8 Loops	2 S* dst = tmp; 9 8 2 S* mask <<= 8; 8 8 0005	2 -	<pre>tmp = src;</pre>	24	8	
50 2 S* mask <<= 8; 8 8 Loops	2 S* mask <<= 8; 8 8					
Loops	nops	2 5*	dst = tmp;	9	8	
Loops	nops	50 0 0*	mak 2</td <td></td> <td></td> <td></td>			
				0	0	
		▷ Loops				
		Details				

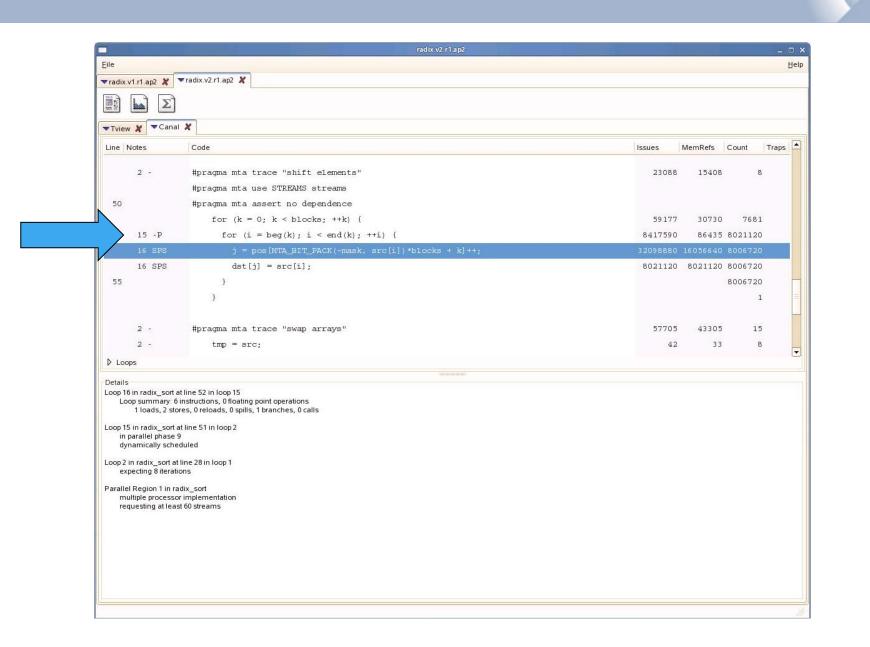


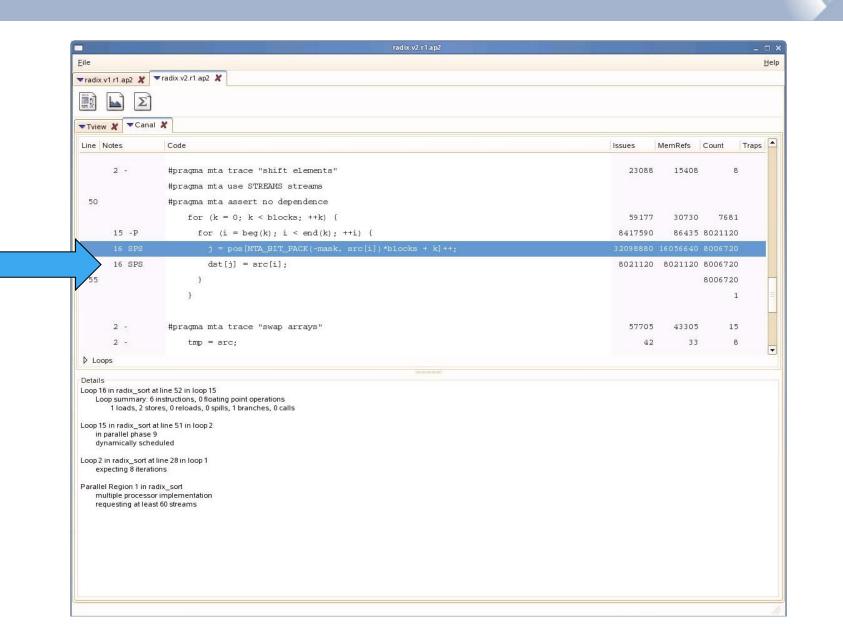
## **Further modify code**

 Split the array up into sequential blocks, allowing us to replicate the buckets into which we determine the element counts and new positions.

```
#pragma mta assert no dependence
      for (k = 0; k < blocks; ++k) {
        for (i = beg(k); i < end(k); ++i) {
          j = pos[MTA BIT PACK(~mask,
              src[i])*blocks + k]++;
          dst[j] = src[i];
Recompile, run, post-process, and run
  Apprentice2
```





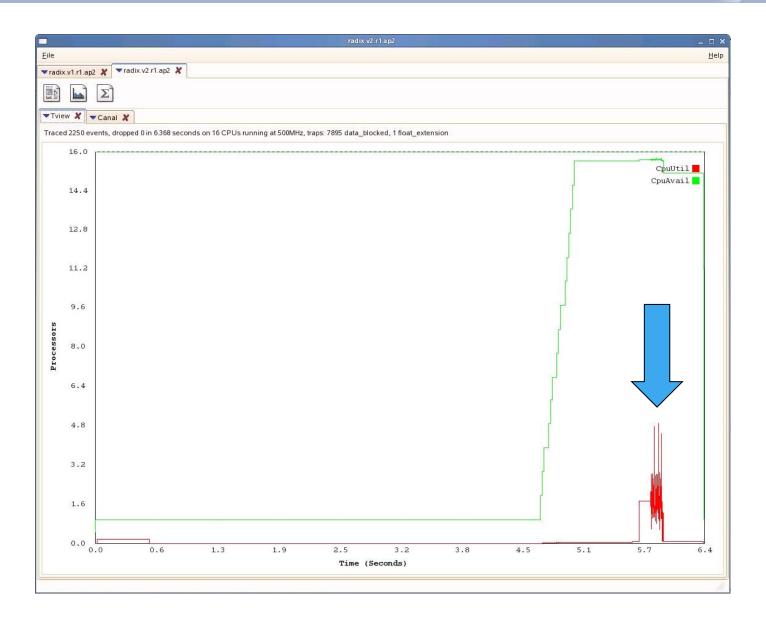


	radix v2.r1.ap2	زوز محيدين أترجعه العا			_ 🗆 🗙
<u>F</u> ile					Help
🔻 radix.v1.r1.ap2 🗶	🛛 radix.v2.r1.ap2 🦹				
Σ					
Tview 🗶 🔻 Canal	*				
Line Notes	Code	Issues	MemRefs	Count	Traps
2 -	#pragma mta trace "shift elements"	23088	15408	8	
	#pragma mta use STREAMS streams				
50	#pragma mta assert no dependence				
	for $(k = 0; k < blocks; ++k)$ {	59177	30730	7681	
15 -P	for $(i = beg(k); i < end(k); ++i) $ {	8417590	86435	8021120	
16 SPS	<pre>j = pos[MTA_BIT_PACK(~mask, src[i])*blocks + k]++;</pre>	32098880	16056640	8006720	
16 SPS	<pre>dst[j] = src[i];</pre>	8021120	8021120	8006720	
55	3			8006720	-
	3			1	-
2 -	#pragma mta trace "swap arrays"	57705	43305	15	
2 -	tmp = src;	42	33	8	•
D Loops					U
Details	1010000000				
	Instructions, 0 floating point operations es, 0 reloads, 0 spills, 1 branches, 0 calls t line 51 in loop 2 duled line 28 in loop 1 ons dix_sort implementation				
<u> </u>					

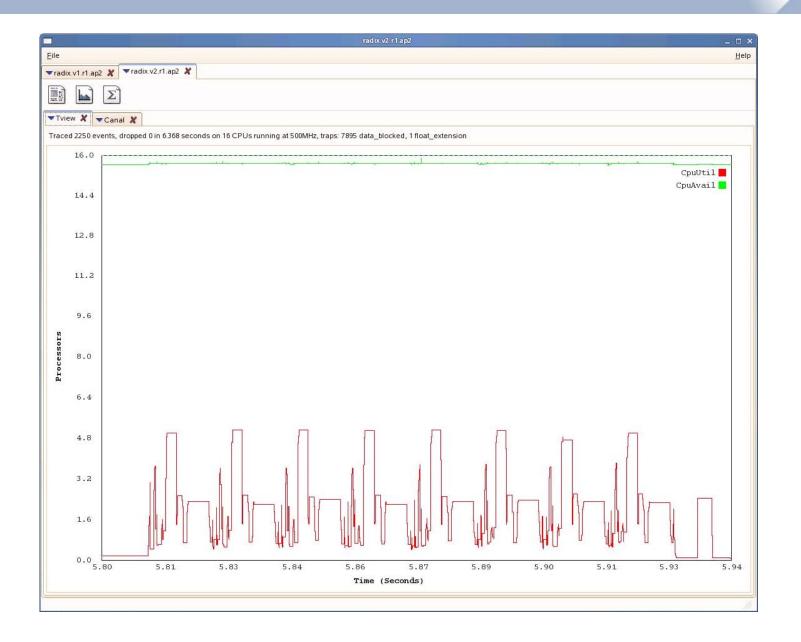
radix.v1.r1.ap2 ★ Tradix.v2.r1.ap2 ★		radix.v2.r1.ap2				- 0
Notes         Code         Issues         MemRets         Count         Traps           2 -         Hpragma mta trace "shift elements"         23088         15408         8           50         Hpragma mta use STREMIS streams         23088         15408         8           50         Hpragma mta use STREMIS streams         59177         30730         7681           15 -P         for (k = 0; k < blocks; +tk) (         59177         30730         7681           15 -P         for (k = 0; k < blocks; +tk) (         8417590         86435         8021120           16 SPS         j = pos[WTA_PETC_FACK(-mask, arc[1])*blocks + k]++;         32058880         16056640         8006720           16 SPS         dst[j] = src[i];         8021120         8021120         802120         8006720           1         1         .         .         .         .         .         .           2 -         tmp = src;         .         .         .         .         .         .           2 -         tmp = src;         .         .         .         .         .         .           2 -         tmp = src;         .         .         .         .         .         .						H
Notes       Code       Haues       MemRets       Count       Traps         2       -       #pragma mta trace "shift elements"       #pragma mta use STREAMS streams       23088       15408       8         50       #pragma mta use STREAMS streams       59177       30730       7681         50       #pragma mta assert no dependence       for (k = 0; k < blocks; ++k) (       8417590       86435       8021120         15       -P       for (i = beg(k); i < end(k); ++i) (       8417590       86435       8006720         16       SPS       j = pos[MTR_BIT_PACK(-mask, src(i))*blocks + k]++;       32098880       1605640       8006720         16       SPS       dst[j] = src[i];       8001120       8006720       8006720         2       .       #pragma mta trace "swap arrays"       57705       43305       15         2       .       tmp = arc;       42       33       8         Octops       .	radix.v1.r1.ap2 🗶	▼radix.v2.r1.ap2 🗶				
Notes       Code       Hasses       MemRets       Count       Traps         2 -       Hpragma mta trace "shift elements" Hpragma mta use STREAMS streams       23088       15408       8         50       Hpragma mta use STREAMS streams       59177       30730       7681         50       Hpragma mta assert no dependence for (k = 0; k < blocks; ++k) (       59177       30730       7681         15 -P       for (i = beg(k); i < end(k); ++i) (       8417590       86435       8006720         16 SPS       det[j] = erc[i];       32098880       16056640       8006720         55       )       .       8001120       8006720       8006720         3       .       .       .       .       8006720       8006720         3       .       .       .       .       .       .       .         2 -       tmp = arc;       .       .       .       .       .       .         2 -       tmp = arc;       .       .       .       .       .       .         2 -       tmp = arc;       .       .       .       .       .       .         2 -       tmp = arc;       .       .       .       .       .	Σ					
2 -       #pragma mta trace "shift elements" #pragma mta use STREAMS streams         50       #pragma mta use STREAMS streams         50       for (x = 0; x < blocks; +tk) {         15 -P       for (i = beg(k); i < end(k); ++i) {         16 SPS       dst[j] = prof[NTA_BIT_PACK(-mask, pro[i])*blocks + k]++;         32098880       16056400         55       )         2 -       #pragma mta trace "swap arrays"         2 -       tmp = prc;         2 -       tmp = src;         2	Tview 🗶 🔻 Cana	al 🗶				
#pragma mta use STREAMS streams       #pragma mta assert no dependence       for (k = 0; k < blocks; ++k) (       59177       30730       7681         15 -P       for (i = 0; k < blocks; ++k) (       8417590       86435       8021120         16 SPS       j = pos (MTA_BIT_PACK(-mask, src[i]) *blocks + k)++;       32098880       16056640       8006720         16 SPS       dst[j] = src[i];       8021120       8006720       1         2 -       #pragma mta trace "swap arrays"       1       1       1         2 -       #pragma mta trace "swap arrays"       57705       43305       15         2 -       #pragma mta trace "swap arrays"       57705       43305       15         2 -       tmp = src;       second       second       second       second         2 -       #pragma mta trace "swap arrays"       1       3       8       second       second       second       second       second       secondors	ine Notes	Code	Issues	MemRefs	Count	Traps
50       #pragma mta assert no dependence for (k = 0; k < blocks; ++k) {	2 -	#pragma mta trace "shift elements"	23088	15408	8	
for (k = 0; k < blocks; ++k) {		#pragma mta use STREAMS streams				
15 -P       for (i = beg(k); i < end(k); ++i) (	50	#pragma mta assert no dependence				
16 SPS       j = pos (MTA_PIT_PACK(-mask, src[i])*blocks + k]++;       3209880 16056640 8006720         16 SPS       dst[j] = src[i];       8021120       8021120       802120       8006720         55       }       }       1       1       1       1         2 -       #pragma mta trace "swap arrays"       57705       43305       15       1         2 -       tmp = src;       1       57705       43305       15         2 -       tmp = src;       2       33       8         2 -       tmp = src;       1       57705       43305       15         2 -       tmp = src;       1       57705       43305       15         2 -       tmp = src;       1       57705       43305       15         2 -       tmp = src;       1       57705       43305       15         2 -       tmp = src;       1       57705       43305       15         2 -       tmp = src;       1       57705       43305       15         2 -       tmp = src;       1       57705       43305       15         2 -       tmp = src;       1       1       1       1         2 -       tmp = s		for $(k = 0; k < blocks; ++k)$ {	59177	30730	7681	
16 SPS       dst[j] = src[i];       8021120       8026720         55       )       8006720         j       1         2 -       #pragma mta trace "swap arrays"       57705       43305       15         2 -       tmp = src;       24       33       8         b Loops       5000000000000000000000000000000000000	15 -P	for $(i = beg(k); i \le end(k); ++i)$ {	8417590	86435	8021120	
55       }         2 -       #pragma mta trace "swap arrays"       57705       43305       15         2 -       tmp = src;       42       33       8         > Loops         55       10       10         Details          55       5       5       10         Loops 15 in radix_sort at line 52 in loop 15       Loop summary. 6 instructions, 0 floating point operations 11 loads, 2 stores, 0 releads, 0 splits, 1 branches, 0 calls   <	16 SPS	<pre>j = pos[MTA_BIT_PACK(~mask, src[i]) *blocks + k]++;</pre>	32098880	16056640	8006720	
} 1   2 - #pragma mta trace "swap arrays" 57705 43305 15   2 - tmp = src; 42 33 8   Loops   Image: Im	16 SPS	<pre>dst[j] = src[i];</pre>	8021120	8021120	8006720	
2 - #pragma mta trace "swap arrays" 57705 43305 15   2 - tmp = src; 42 33 8 <b>&gt; Loops</b> Details   Details   comment   Details   comment   Details   comment   Details   company: 6 instructions, 0 floating point operations   1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls 57705 43305   Loop 15 in radix_sort at line 51 in loop 2 company: 6 instructions, 0 floating point operations 57705   1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls company: 6 instructions, 0 floating point operations   1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls company: 6 instructions, 0 floating point operations   1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls company: 6 instructions, 0 floating point operations   1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls company: 6 instructions, 0 floating point operations   1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls company: 6 instructions, 0 floating point operations   1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls company: 6 instructions, 0 floating point operations   1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls company: 6 instructions, 0 floating point operations	55	3			8006720	
2 - tmp = src; 42 33 8 ▷ Loops Details Loop summary: 6 instructions, 0 floating point operations 1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls .coop 15 in radix_sort at line 51 in loop 2 in parallel phase 9 dynamically scheduled .coop 2 in radix_sort at line 28 in loop 1 expecting 8 iterations Parallel Region 1 in radix_sort multiple processor implementation		3			1	
Loops  Loops  Loop 16 in radix_sort at line 52 in loop 15 Loop summary: 6 instructions, 0 floating point operations 1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls Loop 15 in radix_sort at line 51 in loop 2 in parallel phase 9 dynamically scheduled Loop 2 in radix_sort multiple processor implementation	2 -	#pragma mta trace "swap arrays"	57705	43305	15	
Details Loop 16 in radix_sort at line 52 in loop 15 Loop summary: 6 instructions, 0 floating point operations 1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls Loop 15 in radix_sort at line 51 in loop 2 in parallel phase 9 dynamically scheduled Loop 2 in radix_sort at line 28 in loop 1 expecting 8 iterations Parallel Region 1 in radix_sort multiple processor implementation	2 -	<pre>tmp = src;</pre>	42	33	8	
Loop 16 in radix_sort at line 52 in loop 15 Loop summary: 6 instructions, 0 floating point operations 1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls Loop 15 in radix_sort at line 51 in loop 2 in parallel phase 9 dynamically scheduled Loop 2 in radix_sort at line 28 in loop 1 expecting 8 iterations Parallel Region 1 in radix_sort multiple processor implementation	Loops					
Loop 16 in radix_sort at line 52 in loop 15 Loop summary: 6 instructions, 0 floating point operations 1 loads, 2 stores, 0 reloads, 0 spills, 1 branches, 0 calls Loop 15 in radix_sort at line 51 in loop 2 in parallel phase 9 dynamically scheduled Loop 2 in radix_sort at line 28 in loop 1 expecting 8 iterations Parallel Region 1 in radix_sort multiple processor implementation	Details					
	Loop summary: 1 loads, 2 st 1 oradix_sort in parallel phase dynamically sch .cop 2 in radix_sort a expecting 8 iterat Parallel Region 1 in r multiple process	6 instructions, 0 floating point operations ores, 0 reloads, 0 spills, 1 branches, 0 calls at line 51 in loop 2 9 eduled tt line 28 in loop 1 tions iadix_sort or implementation				

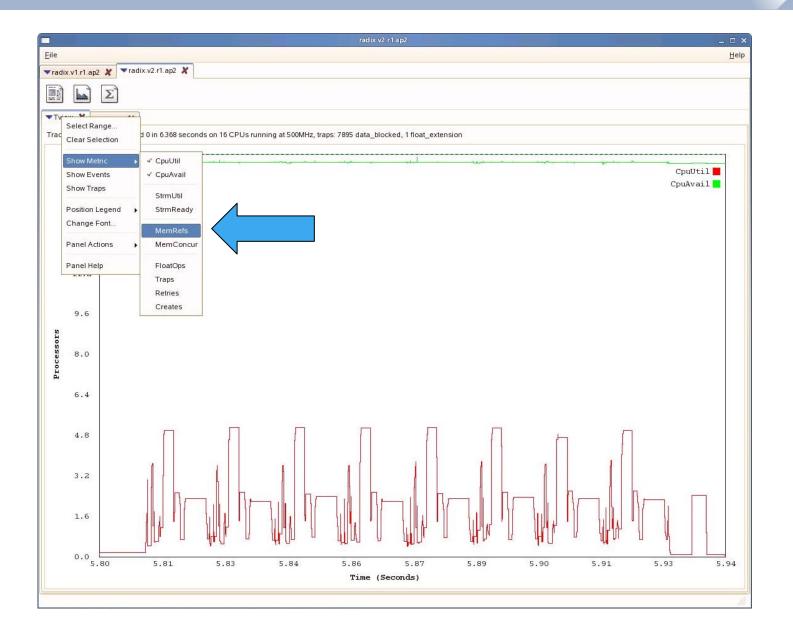
	radix v2.r1 ap2				- 0
ile					Ц
radix.v1.r1.ap2 🗶	🔻 radix.v2.r1.ap2 🦹				
Σ					
Tview 🗶 🔻 Cana	1 X				
Line Notes	Code	Issues	MemRefs	Count	Traps
2 -	#pragma mta trace "shift elements"	23088	15408	8	
	#pragma mta use STREAMS streams				
50	#pragma mta assert no dependence				
	for $(k = 0; k < blocks; ++k)$ {	59177	30730	7681	
15 -P	for $(i = beg(k); i < end(k); ++i)$ {	8417590	86435	8021120	
16 SPS	<pre>j = pos[MTA_BIT_PACK(~mask, src[i])*blocks + k]++;</pre>	32098880	16056640	8006720	
16 SPS	<pre>dst[j] = src[i];</pre>	8021120	8021120	8006720	
55	)			8006720	ŕ
	3			1	1
2 -	#pragma mta trace "swap arrays"	57705	43305	15	
2 -	tmp = src;	42	33	8	
Loops					[
Details					
1 loads, 2 st oop 15 in radix_sort in parallel phase dynamically sch oop 2 in radix_sort a expecting 8 iterat Parallel Region 1 in r	5 instructions, 0 floating point operations ores, 0 reloads, 0 spills, 1 branches, 0 calls at line 51 in loop 2 9 eduled t line 28 in loop 1 ions adix_sort or implementation				



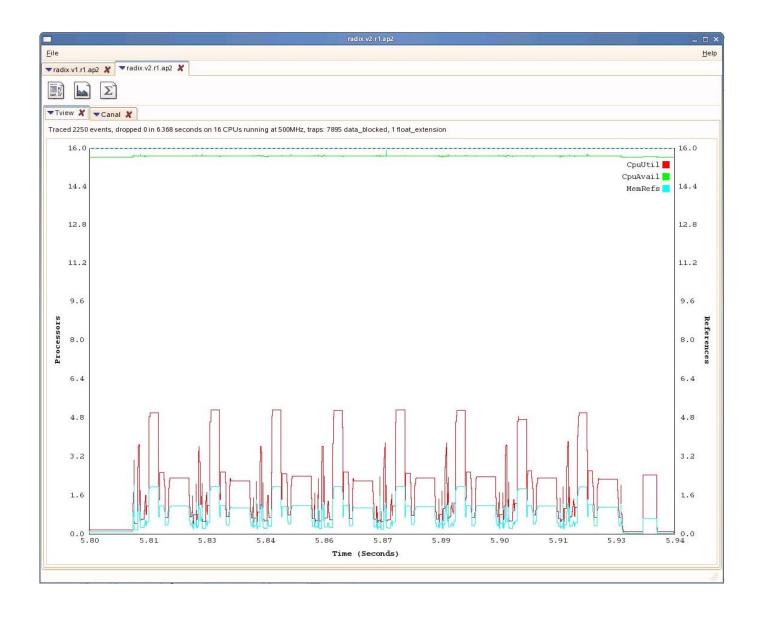




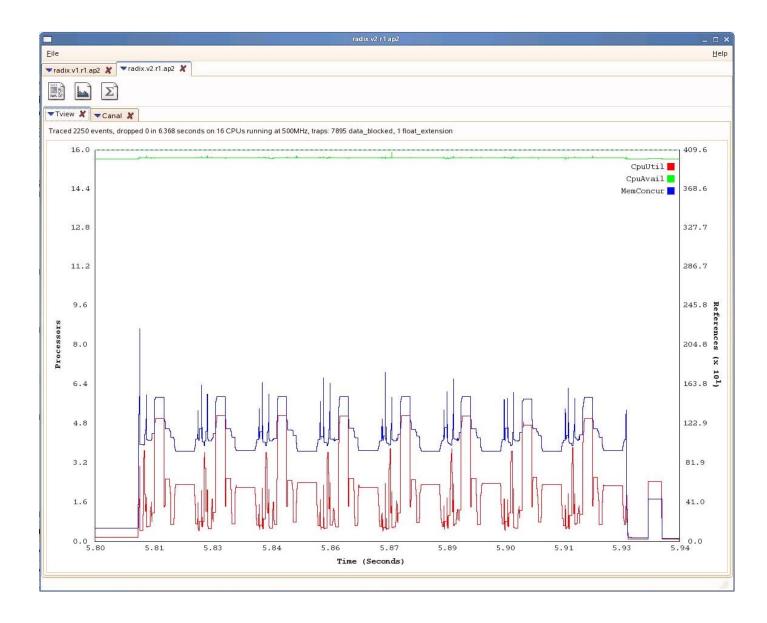














adix.v1.r1.ap2 🗶 🔻 radix.v2.r1.	ap2 🗶									
Tview 🗶 🔻 Canal 🗶										
ced 2250 events, dropped 0 in 6.	368 seconds on 10	6 CPUs running at	500MHz, tra	ps: 7895 data_blocked	, 1 float_extensio	n				
16.0										204 8
1402 1412 1412 1412 1412 1402 1412 1412										Jtil
									MeinCon	Jtil   143.4 122.9 102.4
6.4									Henco	ncur 102.4 102.4 81.9 61.4 20.5 0.0
1.6										20:5
0.0 0.6	1.3	3 1.5	9		3.2	3.8	4.5	5.1	5.7	6.4
				Time (Sec						
				1111						
ents Traps										
pe Data Result Co	de Retry Op Cod	e Count Rate	Time	Destination Register	Data Address	Program Counte	r			
TA_BLOCKED RETRY_LIMIT	LOAD	1 ir	nf 5.816743	r23	0x4203015ad20					0x404c4
TA_BLOCKED RETRY_LIMIT	LOAD	1 ir	nf 5.822485	r14	0x420301774f0					0x404b0
TA_BLOCKED RETRY_LIMIT	LOAD	1 ir	nf 5.823629	r23	0x420301603b0					0x404c4
TA_BLOCKED RETRY_LIMIT	LOAD	1 ir	nf 5.831658	r23	0x4203016c050					0x404c4
TA_BLOCKED RETRY_LIMIT	LOAD	1 ir	nf 5.833836	r5	0x42030162fd0					0x404b1
TA_BLOCKED RETRY_LIMIT	LOAD	2 0.00000	0 5.835164	r5	0x42030162190					0x404b1
TA_BLOCKED RETRY_LIMIT	LOAD	2 0.00000	0 5.835835	r23	0x4203017daf0					0x404c4
TA_BLOCKED RETRY_LIMIT	LOAD		nf 5.838070							0x404b1
TA_BLOCKED RETRY_LIMIT	LOAD	2 0.00000	0 5.838158	r23	0x420301667a0					0x404c4
TA_BLOCKED RETRY_LIMIT	LOAD	1 ir	nf 5.838771	r23	0x4203016bbc0					0x404c4
TA_BLOCKED RETRY_LIMIT	LOAD	1 ir	nf 5.843265	r23	0x4203015b490					0x404c4
TA_BLOCKED RETRY_LIMIT	LOAD		nf 5.845259		0x4203015c770					0x404b0
TA_BLOCKED RETRY_LIMIT	LOAD		0 5.847962		0x42030174010					0x404c5
TA_BLOCKED RETRY_LIMIT	LOAD		0 5.848210							0x404c4
ATA_BLOCKED RETRY_LIMIT	LOAD		nf 5.851762	r5						0x404b1
TA_BLOCKED RETRY_LIMIT	LOAD		0 5.852285		0x42030164570					0x404b0
TA_BLOCKED RETRY_LIMIT	LOAD		nf 5.853986							0x404c4
TA_BLOCKED RETRY_LIMIT	LOAD		nf 5.854300							0x404b0
ATA_BLOCKED RETRY_LIMIT	LOAD		nf 5.855008							0x404c4
ATA BLOCKED RETRY LIMIT	LOAD		0 5.855172	r23	0x42030174130 0x42030161d50					0x404c4 0x404c4
ATA_BLOCKED RETRY_LIMIT	LOAD	1 ir	nf 5.856899							

CRAY

<u>F</u> ile				Help
▼radix.v	v1.r1.ap2 🗶 🔻 r	radix.v2.r1.ap2 🗶		
	Σ			
Tview	v 🗶 🔻 Canal 🔉	K		
Line N	Votes	Code	Issues	MemRefs Count Traps
	8	for $(k = 0; k < blocks; ++k)$ {	3429662	1454708 46081 138
	9 -P 6 -P	for $(i = beg(k); i < end(k); ++i)$ {	9206993	416852 8000000
	12 SPP:m\$	<pre>cnt[MTA_BIT_PACK(~mask, src[i])*blocks + k]++;</pre>	24528394	8528392 8000000
		}		8000000
40		}		1
	2 -	#pragma mta trace "find positions"	134928	30768 8
	2 S	pos[0] = 0;	16	8 8
	13	for (i = 1; i < blocks*buckets; ++i) {	3152310	288061 30721 248
45	14 SL	pos[i] = pos[i - 1] + cnt[i - 1];	9846680	
		}		1
	~		03000	
D Loo	ps			
sta	in radix_sort at lin age 1 of recurrence op summary: 12 in			
in p rec Loop 2 i	in radix_sort in loc parallel phase 5 currence control lo in radix_sort at lin pecting 8 iterations	pop, chunk size = 1024 e 28 in loop 1		
Parallel mu	I Region 1 in radix litiple processor in questing at least 6	<_sort mplementation		



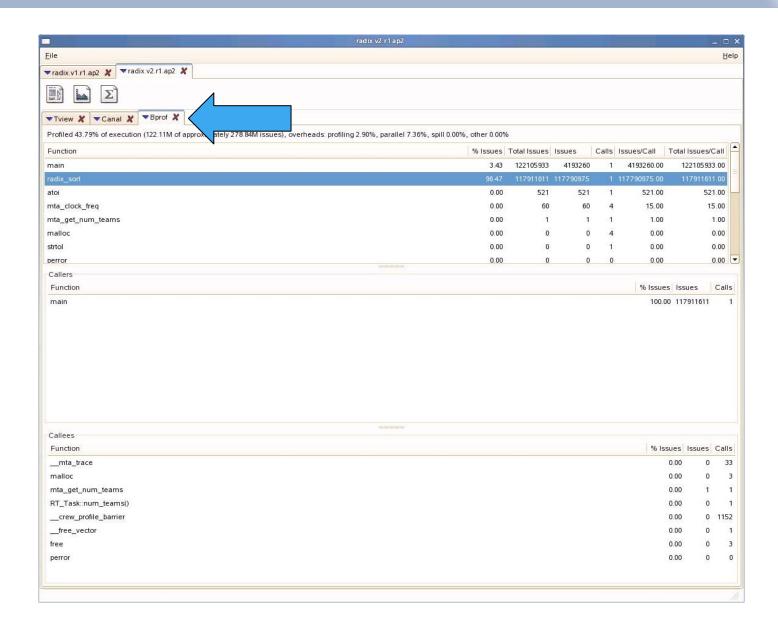
lix.v1.r1.ap2 🎗											
view 🗶 🔻 Car	nal 🗶										
ced 2250 events	, dropped 0 in 6.36	8 seconds	on 16 CPUs runnin	g at 500MHz, tra	aps: 7895 data_block	ed, 1 float_exter	nsion				
0400400400400 6401-1000400400400 111110000400100										Meircon	acur   122.4 %
- 5:8 <u>-</u>	0.6		1.3	1.9	2.5	3.2	3.8	4.5	5.1	5.7	6.4 J
					Time (S	econds)					-
vents Traps											
			pCode Count Pro			Source					Line
ATA_BLOCKED		LOAD LOAD	248 117		radix.v2.pl radix.v2 radix.v2.pl radix.v2						44 36
ATA_BLOCKED		LOAD	21		radix.v2.pl radix.v2						36
ATA_BLOCKED		LOAD	21		radix.v2.pl radix.v2						45
Collate by progr	am counter										



•	_										
dix.v1.r1.ap2 🎗	▼radix.v2.r1.ap2	×									
view 🗶 🔻 Ca	nal ¥										
ced 2250 events		cocondo	on 16 CPU c run	aing at E00MHz tr	nc: 790E da	ta_blocked, 1 float_exte	ncion				
		seconds	on to CPOSitum	ing at 500 winz, the	aps. 7055 uz	ita_blocked, Tilbat_exte	IISIOII				
0482 0482 0482 0482 0482 0482 0482 0482											204:8
04801600480160 64401600480160 64401600480160 14411978648710										SpuU	122.9
8.0 6.4 4.8										MeirCond	
3.2											01.4 41.0 20.5 0.0
0.0	0.6		1.3	1.9	2.5	3.2	3.8	4.5	5.1	5.7	6.4 20.5 0.0
					Ti	.me (Seconds)					
ents Traps											
pe	Data Result Code	Retry Op	Code Count	Program Counter	Library	Module Source					L
TA_BLOCKED	RETRY_LIMIT	LOAD	248	0x404c4c	radix.v2.pl	radix.v2.o radix.v2.c					
TA_BLOCKED	RETRY_LIMIT	LOAD	117	0x404b05	radix.v2.pl	radix.v2.o radix.v2.c					
TA_BLOCKED	RETRY_LIMIT	LOAD	21	0x404b18	radix.v2.pl	radix.v2.o radix.v2.c					
ATA_BLOCKED	RETRY_LIMIT	LOAD	21	0x404c5d	radix.v2.pl	radix.v2.o radix.v2.c					
Collate by prog											

	radix v2.r1.ap2		_ 0
ile			He
radix.v1.r1.ap2 🗶	🔻 radix.v2.r1.ap2 🗶		
Σ			
Tview 🗶 🔻 Cana			$\checkmark$
Line Notes	Code	Issues	MemRefs Count Traps
	Frank.		
2 -	#pragma mta trace "shift elements"	23088	15408 8
	#pragma mta use STREAMS streams		
50	#pragma mta assert no dependence		
	for $(k = 0; k < blocks; ++k)$ {	59177	30730 7681
15 -P	for $(i = beg(k); i < end(k); ++i)$ {	8417590	86435 8021120
16 SPS	<pre>j = pos[MTA_BIT_PACK(~mask, src[i]) *blocks + k]++;</pre>	3.2098880	16056640 8006720
16 SPS	<pre>dst[j] = src[i];</pre>	8021120	8021120 8006720
55	1		8006720
	)		1
2 -	#pragma mta trace "swap arrays"	57705	43305 15
2 -	tmp = src;	42	
▷ Loops			
Details			
1 loads, 2 sto Loop 15 in radix_sort in parallel phase dynamically sche Loop 2 in radix_sort al expecting 8 iterat Parallel Region 1 in ra	instructions, 0 floating point operations ores, 0 reloads, 0 spills, 1 branches, 0 calls at line 51 in loop 2 9 siduled t line 28 in loop 1 ions adix_sort or implementation		

radix v2 r1 ap2				_ = >
				Help
▼radiy v2.r1.ap2 🗶				
Code	Issues	MemRefs	Count	Traps
#pragma mta trace "shift elements"	23088	15408	8	
#pragma mta use STREAMS streams				
#pragma mta assert no dependence				
for $(k = 0; k < blocks; ++k)$ {	59177	30730	7681	
for $(i = beg(k); i < end(k); ++i)$ {	8417590	86435	8021120	
<pre>j = pos[MTA_BIT_PACK(~mask, src[i]) *blocks + k] ++;</pre>	32098880	16056640	8006720	
<pre>dst[j] = src[i];</pre>	8021120	8021120	8006720	
3			8006720	-
)			1	
#pragma mta trace "swap arrays"	57705	43305	15	
<pre>tmp = src;</pre>	42	33	8	-
				Ū
At line 52 in loop 15 instructions, 0 floating point operations ores, 0 reloads, 0 spills, 1 branches, 0 calls at line 51 in loop 2 9 eduled t line 28 in loop 1 tons adix_sort or implementation st 60 streams				
	<pre>radiy 2.r1.ap2 * Code #pragma mta trace "shift elements" #pragma mta use STREAMS streams #pragma mta use STREAMS streams #pragma mta assert no dependence for (k = 0; k &lt; blocks; ++k) ( for (i = beg(k); i &lt; end(k); ++i) (</pre>	<pre>*rait 211.92 *  Code Issues  #pragma mta trace "shift elements" #pragma mta use STREAMS streams #pragma mta assert no dependence for (k = 0; k &lt; blocks; ++k) ( for (i = beg(k); i &lt; end(k); ++i) (</pre>	radiu 2.11 ap2       MemRefs.         Code       Issues       MemRefs.         #pragma mta trace "shift elements"       23088       15408         #pragma mta use STREAMS streams       #pragma mta use STREAMS streams       \$9177       30730         for (k = 0; k < blocks; ++k) (	Code       Issues       MemRefs       Count         #pragma mta trace "shift elements"       23088       15408       8         #pragma mta use STEEAMS streams       9177       30730       7681         #pragma mta assert no dependence       for (k = 0; k < blocks; ++k) (





	radix v2.r1.ap2						_ 0	×
Eile							He	lp
🔽 radix.v1.r1.ap2 🧩 🔍 radix.v2.r1.ap2 🗶								
Tview X Canal X Bprof X								
Profiled 43.79% of execution (122.11M of approximately 278.84M issues), overheads: prof	filing 2.90%, parallel 7.36%, spill 0.00%, other 0.00%	6						
Function	% Issues	Total Iss	Issues	Calls	Issues/Call	Total Issues	Call	-
main	3.43	122105933	4193260	1	4193260.00	12210593	33.00	
radix_sort	96.47	117911611	117790975	1	117790975.00	11791161	11.00	-
atoi	0.00	521	521	1	521.00	52	21.00	
mta_clock_freq	0.00	60	60	4	15.00	1	15.00	
mta_get_num_teams	0.00	1	1	1	1.00		1.00	
malloc	0.00	0	0	4	0.00		0.00	
strtol	0.00	0	0	1	0.00		0.00	
berror	0.00	0	0	0	0.00		0.00	•
Callers								
Function						s Issues 0 117911611	Calls	3
Callees								
Function					% Iss	ues Issues	Calls	5
mta_trace						0.00 0	33	3
malloc						0.00 0	3	3
mta_get_num_teams						0.00 1	1	Ē
RT_Task::num_teams()						0.00 0	1	1
crew_profile_barrier						0.00 0	1152	2
free_vector						0.00 0	c 1	Ł
free						0.00 0	3	3
perror						0.00 0	C	0

	radix v2.r1.ap2						14	_ 0
le								F
radix.v1.r1.ap2 🦹 🔻 radix.v2.r1.ap2 🦹								
Tview 🗶 🔻 Canal 🗶 TBprof 🗶				、				
rofiled 43.79% of execution (122.11M of approximately 278.84M issues), overheads	profiling 2.90%, parallel 7.36%, spill 0.	00%, other 0.00%	C.	$\checkmark$				
unction		% Issues	Total Issues	Issues	Calls	Issues/Call	Total Issues	Call
nain		3.43	122105933	4193260	1	4193260.00	1221059	33.00
adix_sort		96.47	117911611	117790975	1	117790975.00	1179116	11.00
toi		0.00	521	521	1	521.00	5	21.00
nta_clock_freq		0.00	60	60	4	15.00		5.00
nta_get_num_teams		0.00	1	1	1	1.00		1.00
nalloc		0.00	0	0	4	0.00		0.00
trtol		0.00	0	0	1	0.00		0.00
error		0.00	0	0	0	0.00		0.00
Callers								
							s Issues 0 117911611	-
Callers Function								-
Callers Function	AREADORNIUM							-
Callers Function main						100.0		
Callees						100.0	0 117911611	Ca
Callees Function						100.0 % Iss	ues Issues	Ca
Callees Function Callees Functiontrace						100.0 % Iss	ues Issues	Са
Callees Function Callees Functionmta_trace malloc						100.0 % Iss	ues Issues 0.00 C	Са
Callees Function Callees Functiontrace malloc mta_get_num_teams						100.0	ues Issues 0.00 C 0.00 C 0.00 1	Са
Callers Callees Function Callees Functionmta_trace malloc mta_get_num_teams RT_Task::num_teams()						100.0	ues Issues 0.00 C 0.00 C 0.00 1 0.00 C	Ca 11
Callers Function Callees Function Callees Functionmta_trace malloc mta_get_num_teams RT_Task::num_teams()crew_profile_barrier						100.0	ues Issues 0.00 C 0.00 C 0.00 C 0.00 C 0.00 C	<b>Ca</b>

						v2 r1 ap2							- 6
le													Н
radix.v1.r1.ap2 🗶 🔻 radix.v	v2.r1.ap2 🗶												
Tview 🗶 🔻 Canal 🗶 🔻	• Issues												
Profiled 43.79% of execution (1		78.	.84M issues	s), overheads: p	profiling 2.90%	6, parallel 7.36	%, spill 0.00%, other 0.00	%					
unction	mennels						% Issues	Total Issues	Issues	Calls	Issues/Call	Total Issue	s/Call
nain	Hide Callers		/				3.43	122105933	4193260	1	4193260.00	122105	933.00
adix_sort	Hide Callees						96.47	117911611	117790975	1	117790975.00	11791	611 00
loi	Panel Actions						0.00	521	521	1	521.00		521.00
nta_clock_freq							0.00	60	60	4	15.00		15.00
ita_get_num_teams	Panel Help	-					0.00	1	1	1	1.00		1.00
nalloc							0.00	0	0	4	0.00		0.00
trtol							0.00	0	0	1	0.00		0.00
error							0.00	0	0	0	0.00		0.00
Callers													
												25 Issues	
												and the second s	
nain						10000000						and the second s	
aliees -unction											100.0	sues Issue	es Cal
nain Callees Function mta_trace											100.0	sues Issue 0.00	es Cal
allees Function mta_trace malloc											100.0	00 1179116 sues Issue 0.00 0.00	es Cal 0
allees -unction mta_trace nalloc nta_get_num_teams						ANNOUNCE					100.0	sues Issue 0.00 0.00 0.00	rs Cal 0 3 1
allees unction _mta_trace nalloc tta_get_num_teams tT_Task::num_teams()						annonenenenenenenenen erreteren err					100.0	sues Issue 0.00 0.00 0.00 0.00 0.00	es Ca 0 0 1 0
allees Sunction _mta_trace nalloc nta_get_num_teams tT_Task::num_teams() _crew_profile_barrier											100.0	sues Issue 0.00 0.00 0.00 0.00 0.00 0.00 0.00	rs Cal 0 : 1 0 0 11:
main Callees Function mta_trace malloc mta_get_num_teams RT_Task::num_teams() crew_profile_barrier free_vector											100.0	sues Issue 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.	s Cal 0 : 1 0 11! 0 11! 0
main Callees Function mta_trace malloc mta_get_num_teams RT_Task::num_teams() crew_profile_barrier free_vector ree											100.0	sues Issue 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.	s Cal 0 3 0 1 0 115 0 0
Function main Callees Functiontrace malloc mta_trace malloc mta_get_num_teams RT_Task::num_teams()crew_profile_barrierfree_vector free perror											100.0	sues Issue 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.	rs Cal 0 3 0 1 0 0 115 0 0

radix.v2.r1.ap								- 🗆 >
								Help
v1.r1.ap2 🗶 🔻 radix.v2.r1.ap2 🗶								
Canal X Bprof X				7				
143.79% of execution (122.11M of approximately 278.84M issues), overheads: profiling 2.90%, paralle	17.36%, spill 0.00%, other 0	.00%	$\checkmark$					
n	% Memrefs	Total Memrefs	Memrefs	Calls	Memrefs/Call	Total	Memrefs/0	all
	1.94	54017433		1	1049689.00		54017433	
ort	97.93	52966727	52900352	4	52900352.00		52966727	.00
	0.00	519	519	1	519.00		51	9.00
ock_freq	0.00	24	24	4	6.00			5.00
et_num_teams	0.00	1	1	1	1.00		10	00.1
	0.00	0	0	4	0.00		8	0.00
	0.00	0	0	1	0.00			0.00
	0.00	0	0	0	0.00			0.00
	0.00							
on							Memrefs 52966727	Calls 1
ar Alba and S								111
on								111
s and a set of a set						100.00	52966727	1
s and the second						100.00		Calls
s and a set of a set						mrefs	S2966727 Memrefs	1
s and a second s						mrefs 0.00	S2966727 Memrefs	Calls 33
s sono						mrefs 0.00 0.00	52966727 Memrefs 0 0	Calls 33 3
s source						mrefs 0.00 0.00 0.00 0.00	Memrefs 0 0 1	Calls 33 3 1
s on s s on trace trace trace trace trace teams sisk::num_teams()						mrefs 0.00 0.00 0.00 0.00 0.00 0.00	Memrefs 0 0 1 0	Calls 33 3 1 1
s on s s on trace  trace  tet_num_teams .tsk::num_teams() w_profile_barrier						mrefs 0.00 0.00 0.00 0.00 0.00 0.00 0.00	Memrefs 0 0 1 0 0 0	Calls 33 3 1 1 1152

CRA

	radix.v2.r1 ap2	-
ile		He
🔽 radix.v1.r1.ap2 🗶 🔍 radix.v2.r1.ap2 🎗		
Tview X Canal X Bprof X		
Profiled 43.79% of execution (122.11M of approximately 278.84M issues), overhead	s: profiling 2.90%, parallel 7.36%, spill 0.00%, other 0.00%	
Function	% Memrefs Total Memrefs Memrefs Calls Memrefs/Call Total Memr	efs/Call
main	1.94 54017433 1049689 1 1049689.00 5401	7433.00
radix_sort	97.93 52966727 52900352 1 52900352 00 5296	6727.00
atoi	0.00 519 519 1 519.00	519.00
mta_clock_freq	0.00 24 24 4 6.00	6.00
mta_get_num_teams	0.00 1 1 1 1.00	1.00
malloc	0.00 0 0 4 0.00	0.00
strtol	0.00 0 0 1 0.00	0.00
perror	0.00 0 0 0.00	0.00
Callers		
Callers Function main	% Memrefs Memr 100.00 52966	
Function	% Memrefs Memr 100.00 52966	
Function	% Memrefs Memr	
Function main	% Memrefs Memr 100.00 52966	127 1
Function main	% Memrefs         Memr           100.00         52966	127 1
Function main Callees Function	% Memrefs     Memrefs       100.00     52966	efs Calls 0 33
Function main Callees Functionmta_trace	% Memrefs     Memrefs       100.00     529663	efs Calls 0 33 0 3
Function Callees Functionmta_trace malloc	% Memrets         Memrets           100.00         52966	efs Calls 0 33 0 3 1 1
Function main Callees Functionmta_trace mailoc mta_get_num_teams	% Memrets         Memrets           100.00         52966	efs Calls 0 33 0 3 1 1 0 1
Function main Callees Functionmta_trace malloc mta_get_num_teams RT_Task::num_teams()	% Memrets         Memrets           100.00         52966	efs Calls 0 33 0 3 1 1 0 1 0 1152
Function main Callees Functionmta_trace malloc mta_get_num_teams RT_Task::num_teams()crew_profile_barrier	% Memrets         Memrets           100.00         52966           ************************************	efs Calls 0 33 0 3 1 1 0 1 0 1152

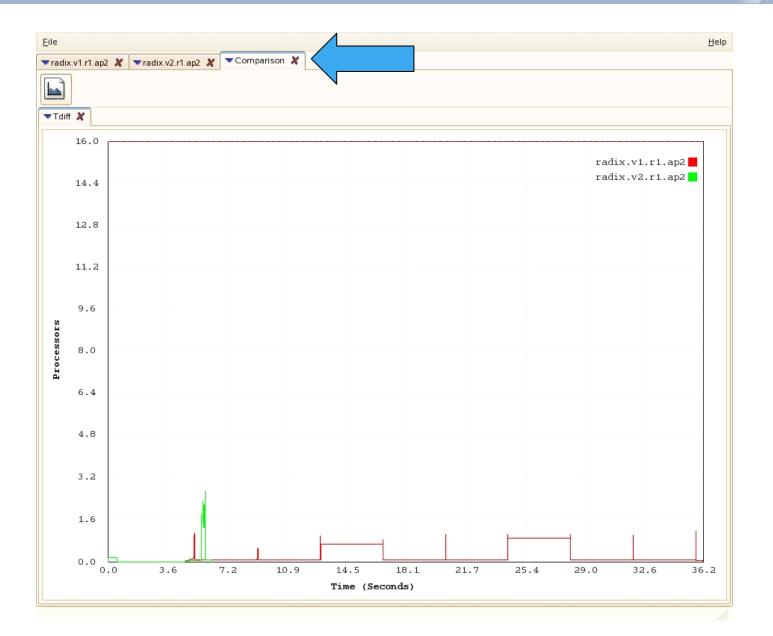
CRA

	radix v2.r1.ap2	- 0
		Щe
adix.v1.r1.ap2 🦹 🔻 radix.v2.r1.ap2 🦹		
Σ		
Tview 🗶 💌 Canal 🗶 🔍 Bprof 🗶		
rofiled 43.79% of execution (122.11M of approximately 278.84M issues), overheads	profiling 2.90%, parallel 7.36%, spill 0.00%, other 0.00%	
unction	% Memrefs Total Memrefs Memrefs Calls Memrefs/Call Total Me	emrefs/Call
nain	1.94 54017433 1049689 1 1049689.00 5	64017433.00
adix_sort	97.93 52966727 52900352 1 52900352 00 53	2966727.00
loi	0.00 519 519 1 519.00	519.00
ta_clock_freq	0.00 24 24 4 6.00	6.00
ita_get_num_teams	0.00 1 1 1 1.00	1.00
nalloc	0.00 0 0 4 0.00	0.00
trtol	0.00 0 0 1 0.00	0.00
error	0.00 0 0 0.00	0.00
unction	% Memrefs   Me 100.00 529	
unction		
unction		
allees	100.00 529	966727
allees	100.00 529 Jacobardon	emrefs Call
allees unction mta_trace	100.00 529 Weighter	emrefs Call
allees unction 	100.00 529 Memmente % Memrefs Me 0.00	emrefs Call 0 3 0
allees unction mta_trace nalloc tta_get_num_teams	100.00 529 Memmets Me 0.00 0.00	emrefs Call 0 3 0 1
allees Function	100.00 529 www	emrefs Call 0 3 0 1 0
Function main Callees Function mta_trace malloc mta_get_num_teams RT_Task::num_teams() crew_profile_barrier	100.00 529 	emrefs Call 0 3 0 1
Callers Function main Callees Functionmta_trace malloc mta_get_num_teams RT_Task::num_teams()crew_profile_barrierfree_vector free	100.00 529 ************************************	emrefs Call 0 3 0 1 0 0 115

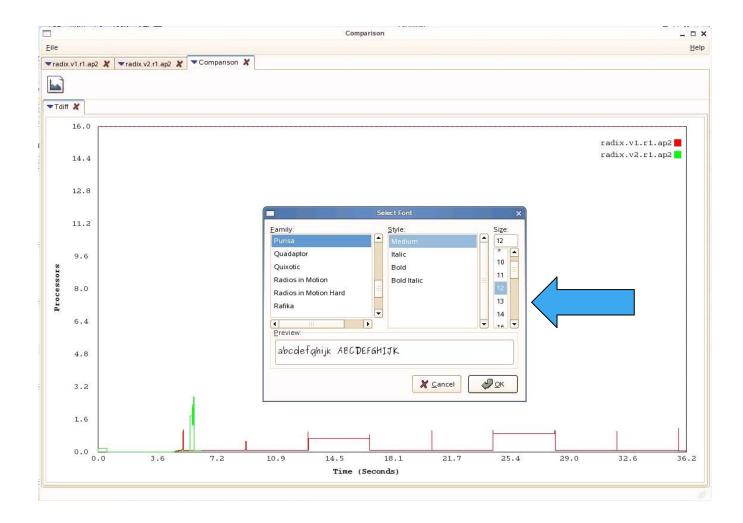
CRA

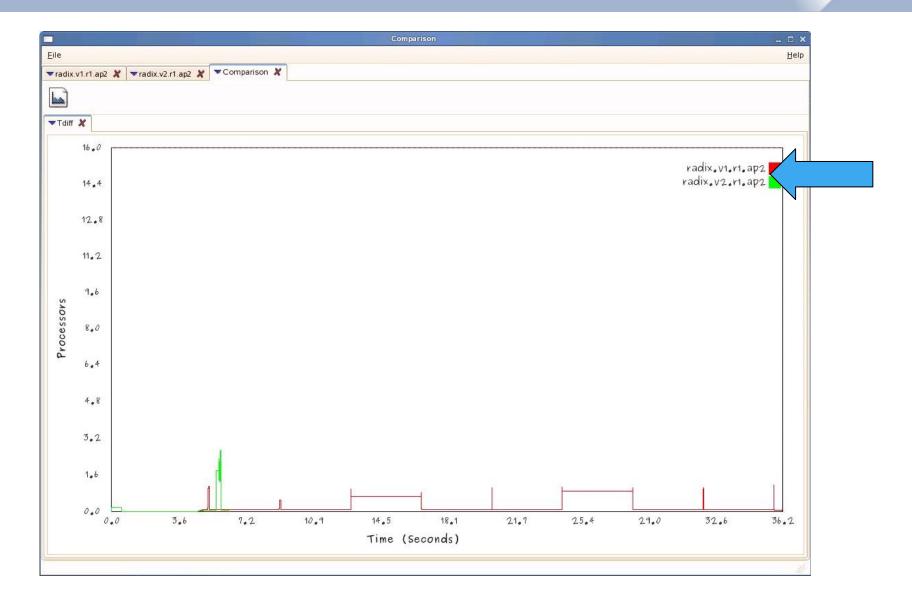
	radix v2.r1.ap2		-
Eile			He
radix.v1.r1.ap2 🗶 🔽 radix.v2.r1.ap2 🗶			
Tview X Canal X Bprof X			
Profiled 43.79% of execution (122.11M of approximately 278.84M issues), overhe	ds: profiling 2.90%, parallel 7.36%, spill 0.00%, other 0.00%		
Function	% Memrefs Total Memrefs Memrefs Calls Memr	efs/Call Total Memre	fs/Call
nain	1.94 54017433 1049689 1 104	9689.00 54017	433.00
radix_sort	97.93 52966727 52900352 1 5290	0352.00 52966	727.00
atoi	0.00 519 519 1	519.00	519.00
nta_clock_freq	0.00 24 24 4	6.00	6.00
mta_get_num_teams	0.00 1 1 1	1.00	1.00
malloc	0.00 0 0 4	0.00	0.00
strtol	0.00 0 1	0.00	0.00
perror	0.00 0 0	0.00	0.00
Callers			
		% Memrefs Memre 100.00 5296672	
		Contraction of the second states	1
	MUNITIPALITY I	Contraction of the second states	
Callees	MANNA ANA ANA ANA ANA ANA ANA ANA ANA AN	100.00 5296672	27
Callees Function	PROMISSION IN CONTRACTOR OF CONT	100.00 5296672 % Memrefs Memre	fs Call
Callees Function mta_trace	AND DESCRIPTION OF	100.00 5296672 % Memrefs Memre 0.00	fs Calls 0 3:
Callees Function mta_trace malloc	PROMINING STATES	100.00 5296672 % Memrefs Memre 0.00 0.00	fs Call: 0 3: 0 :
Callees Function mta_trace malloc mta_get_num_teams		100.00 5296672 % Memrefs Memre 0.00 0.00 0.00	fs Calls 0 3: 0 : 1
Callees Function mta_trace malloc mta_get_num_teams RT_Task::num_teams()		100.00 5296672 % Memrefs Memre 0.00 0.00 0.00 0.00 0.00	fs Calls 0 3: 0 : 1 :
Callees Function mta_trace malloc mta_get_num_teams RT_Task::num_teams() crew_profile_barrier		100.00 5296672 % Memrefs Memre 0.00 0.00 0.00 0.00 0.00 0.00 0.00	fs Calls 0 3: 0 : 1 : 0 115;
Callees Function mta_trace malloc mta_get_num_teams RT_Task::num_teams() crew_profile_barrier free_vector		100.00 5296672 % Memrefs Memre 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.	fs Calls 0 3: 1 0 115; 0
Callees Function mta_trace mailoc mta_get_num_teams RT_Task::num_teams() crew_profile_barrier		100.00 5296672 % Memrefs Memre 0.00 0.00 0.00 0.00 0.00 0.00 0.00	fs Calls 0 3: 0 3: 0 3: 0 3: 0 3: 0 3: 0 3: 0 3:

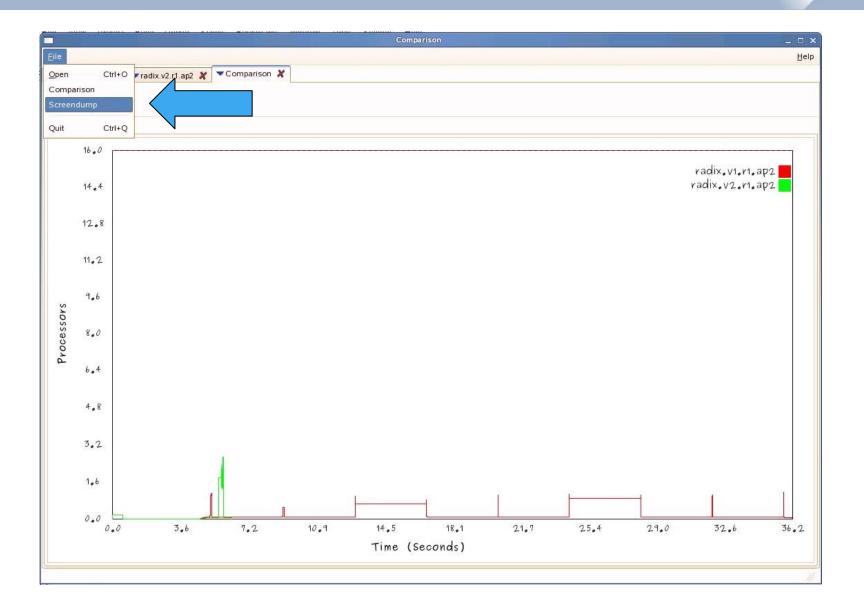








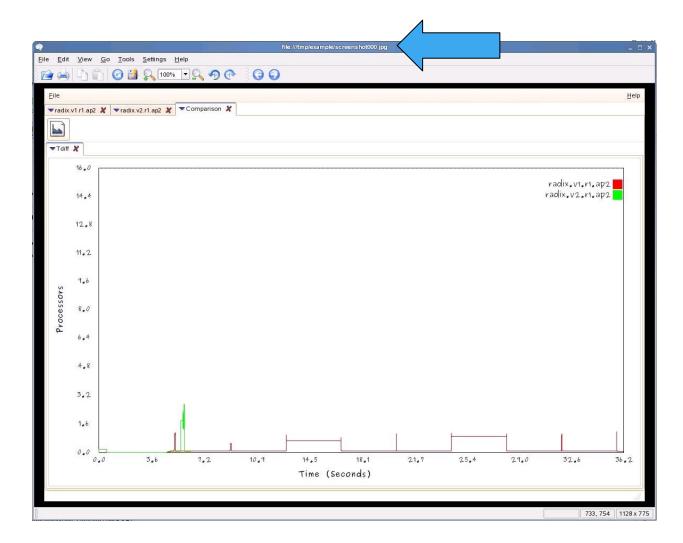




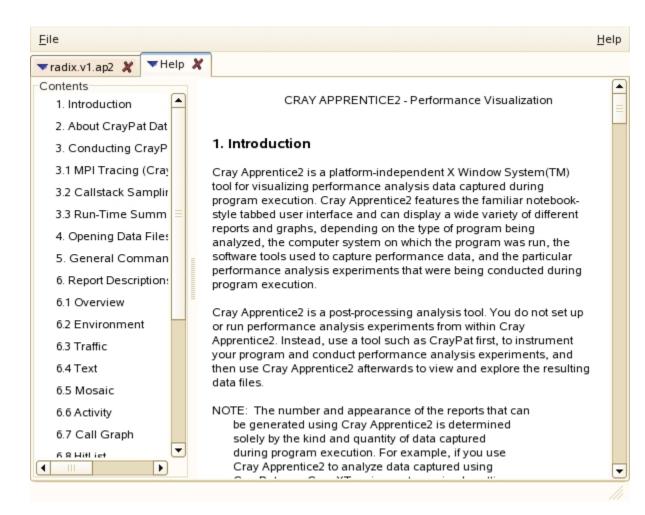












# Workflow: Queue Append I

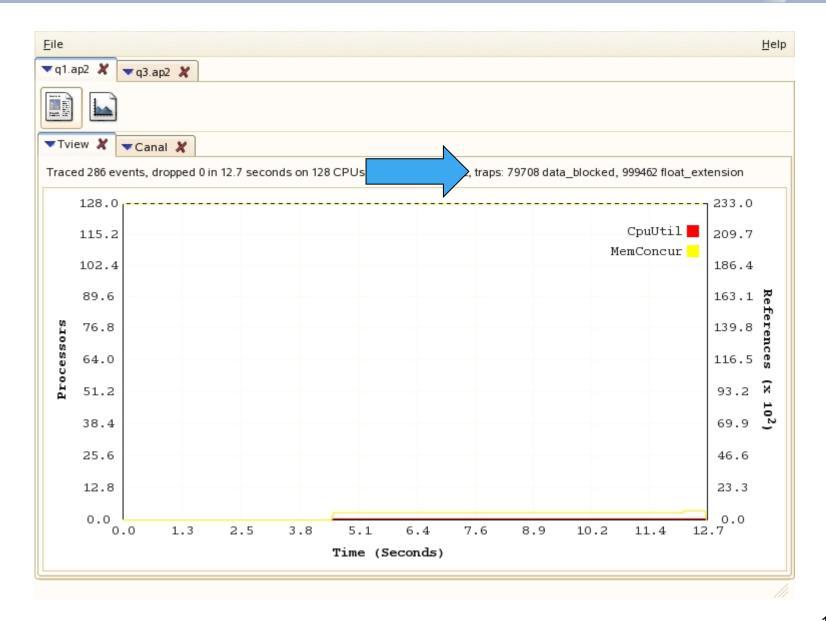
### Write code

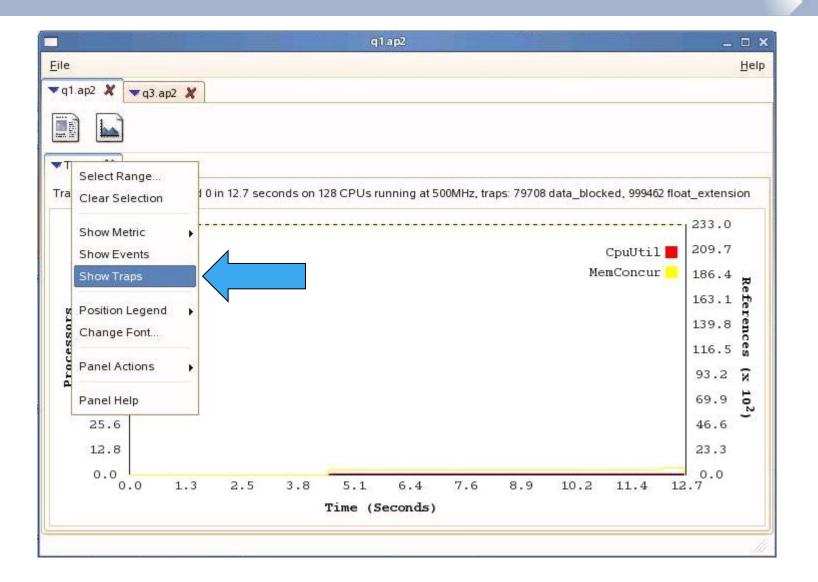
 Append elements onto a queue like what might be done when looking for unvisited nodes in the adjacency lists for set of nodes during a parallelized BFS.

```
unsigned sync k = 0;
#pragma mta assert parallel
#pragma mta use 100 streams
for (unsigned i = 0; i < n; ++i) {
  for (unsigned j = 0; j < cnt[i]; ++j) {
    dst[int_fetch_add(&k, 1)] = src[idx[i]+j];
  }
}
```

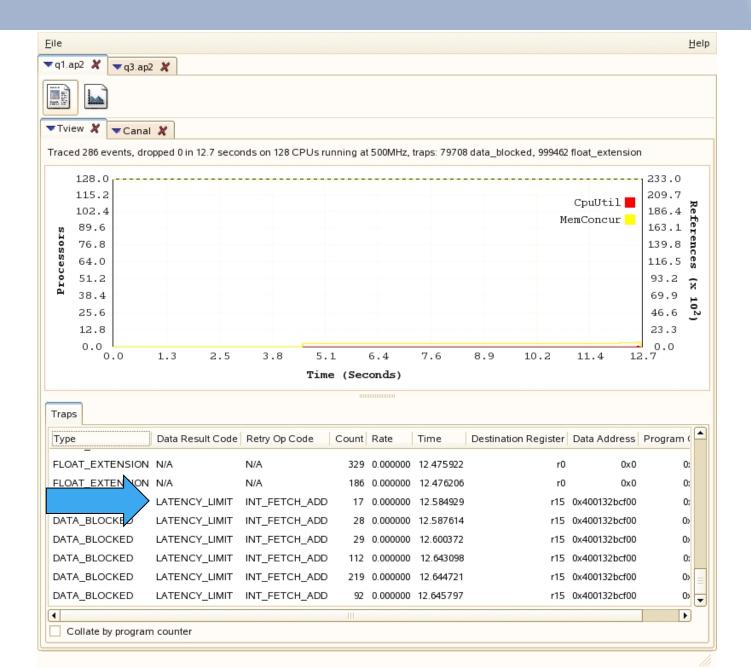
Compile, run, post-process, and run Apprentice2

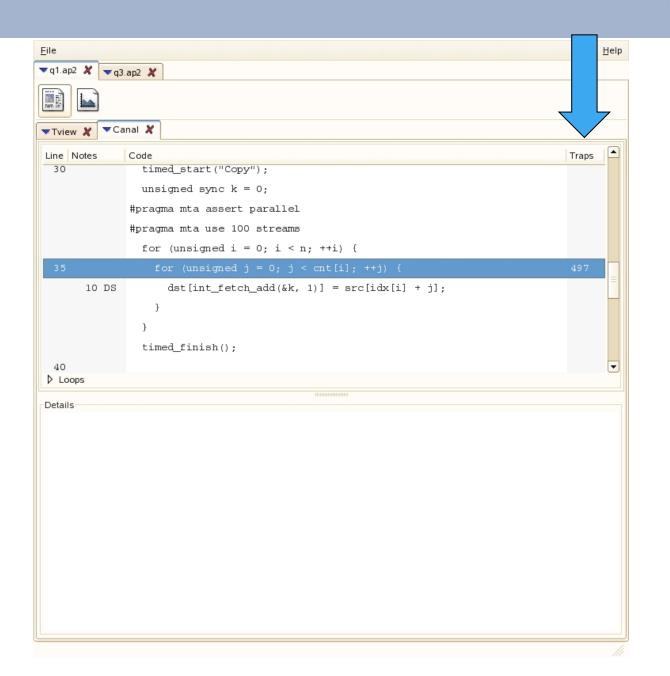












# Workflow: Hotspot Queue II

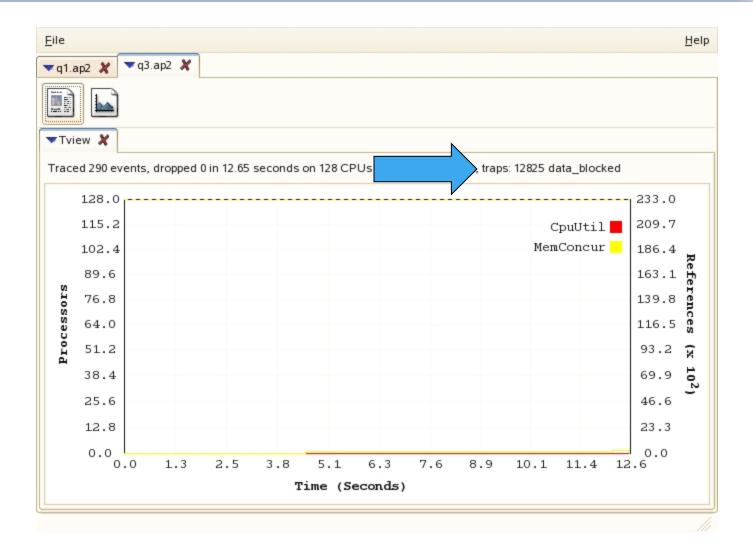
# CRAY

### Modify code

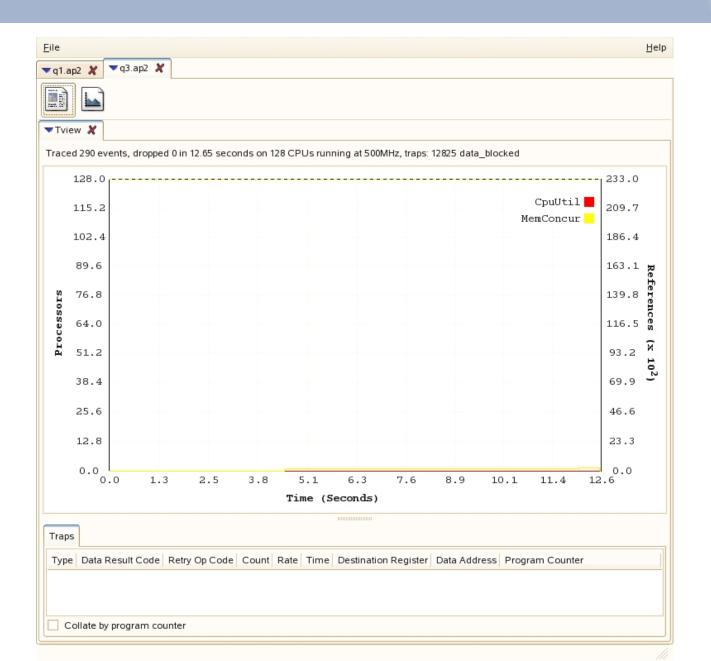
 Grab multiple input and output elements at a time, allowing blanks

```
for (unsigned t = 0; t < blocks; ++t) {
      unsigned 1 = 0;
       for (unsigned i = beg(t); i < end(t); ++i) {
         for (unsigned j = 0; j < cnt[i]; ++j) {
           if (mod(1, BLOCKSIZE) == 0) {
             l = int fetch add(&k, BLOCKSIZE);
           }
          dst[l++] = src[idx[i] + j];
Compile, run, post-process, and run Apprentice2
```











<u>F</u> ile			Help
▼ q1.ap2	2 🗶 ▼q3.a	ap2 💥	
			Ļ
<b>T</b> view	v 🗶 🔻 Can	al 🗶	$\checkmark$
Line N	Notes		Traps 📥
		unsigned sync k = 0;	
		#pragma mta assert parallel	
		#pragma mta use 100 streams	
		for (unsigned t = 0; t < blocks; ++t) {	
45	9 p	unsigned $1 = 0;$	
	10 рХ 9 р	for (unsigned $i = beg(t); i < end(t); ++i)$ {	=
		for $(unsigned j = 0; j < cnt[i]; ++j) $ {	
	11 DXS	if (mod(1, BLOCKSIZE) == 0) {	
	11 DXS 11 DXS		
50	11 DXS	if (mod(1, BLOCKSIZE) == 0) {	•
▶ Loo	11 DXS	<pre>if (mod(1, BLOCKSIZE) == 0) {     l = int_fetch_add(&amp;k, BLOCKSIZE);</pre>	•
Details Loop 11 Loo	11 DXS ops 1 in copy at lin op not pipelin op summary: 2 loads, 2 s	<pre>if (mod(1, BLOCKSIZE) == 0) {     1 = int_fetch_add(&amp;k, BLOCKSIZE); }</pre>	•



#pragma mta trace [on | off | default]

- Doesn't work unless compiled with the -trace flag

#pragma mta trace level n

- Trace only functions that contain at least *n* lines
- Affects the rest of the source code in this file

#pragma mta trace ``string that identifies this
point in the source code"

 Resource usage data is recorded at this point and associated with this character string in the Traceview textual output

# Managing very large trace/ap2 files

- Created mechanism using temporary to help offset memory usage on the login nodes
  - XMT login nodes lack the ability to swap
- Root name for temp files set by environment variable
  - export APP2\_SWAPFILE = /mnt/lustre/users/app2
  - Will generate files with names like
    - /mnt/lustre/users/app2.xxxxx

Temp files cleaned when exiting Apprentice2

# Summary



### Apprentice2

- GUI based system allows rapid assimilation of information
- Interaction between reports assists problem detection
- Canal Report
  - Establishes a dialogue with the compiler that allows the creation of highly parallelized and optimized code

#### Tview Report

- Quickly identify underperforming code sections
- Visualize resource utilization and concurrency
- Pinpoint memory hotspots
- Bprof Report
  - Target specific functions running the least efficiently

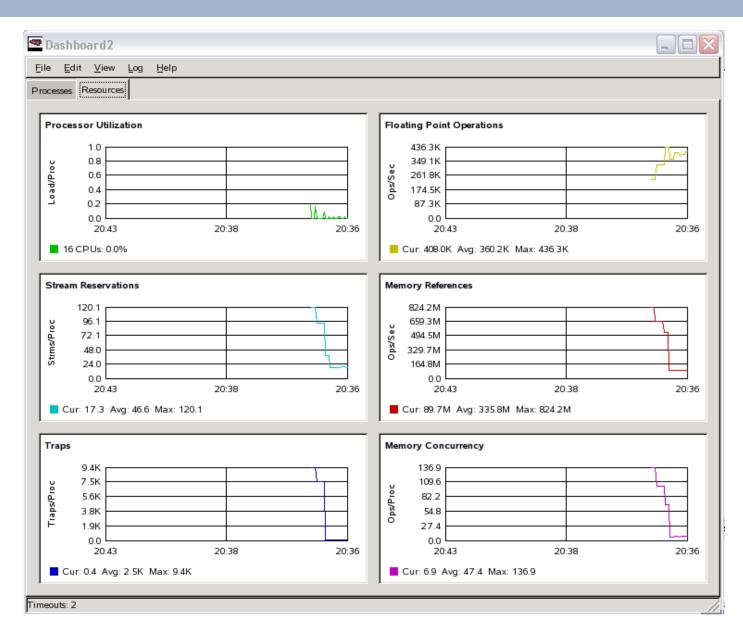
### **Additional Tools and Approaches**





### **Dash: a Real-Time Resource Monitor**





### **Accessing the Hardware Counters**

unsigned issues, memrefs, concur, streams, traps, retries; mta reserve task event counter ( RT ANY COUNTER, RT TRAP ); /\* other counters are available without reserving them \*/ issues = mta get task counter(RT ISSUES); memrefs= mta get task counter(RT MEMREFS); concur= mta get task counter(RT CONCURRENCY); streams= mta get task counter(RT STREAMS); traps = mta get task counter(RT TRAP); retries = mta get task counter(RT MEM RETRY); /\* (continued on next slide) \*/

### **Accessing the Hardware Counters**



```
double start_time = timer();
```

```
int result = bigParallelFunction();
```

```
double elapsed_time = start_time - timer();
```

```
issues = mta_get_task_counter(RT_ISSUES) - issues;
memrefs= mta_get_task_counter(RT_MEMREFS) - memrefs;
concur= mta_get_task_counter(RT_CONCURRENCY) - concur;
streams= mta_get_task_counter(RT_STREAMS) - streams;
traps = mta_get_task_counter(RT_TRAP) - traps;
retries = mta_get_task_counter(RT_MEM_RETRY) - retries;
/* printf() probably goes here */
```

# **Example: Using Hardware Counters**

• Customer was observing significant performance difference between the two loops.

```
- Loop 1:
    for (i = 0; i < no_of_edges; i++) {
        ends[i] = endpoints[i] + 1;
    }
- Loop 2:
    for (i = 0; i < no_of_edges; i++) {
        endpoints[i] = ends[i] % BILLION;
    }
```

- Significant point:
  - Endpoints is an array of unsigned integers, initialized using prand\_int.
  - Elements of endpoints are then randomly distributed 64-bit unsigned integers over the range of all possible 64-bit unsigned integers.

#### Loop1: time copying to ends 6.654852e-02

```
for (i = 0; i < no_of_edges; i++) {
    ends[i] = endpoints[i] + 1;
}</pre>
```

ticks: 33257303, secs: 0.066515, issues: 104509017, memrefs: 101770916, traps: 0, retries: 717757, concurrency: 2542, streams: 755

#### Loop2: time modulo 4.534740e+01

```
for(i = 0; i < no_of_edges; i++) {
    endpoints[i] = ends[i] % BILLION;
}</pre>
```

ticks: 22673682155, secs: 45.347364, issues: 32993662421, memrefs: 15585715642, traps: 49976693, retries: 18913196115, concurrency: 811, streams: 610

# CRAY

## XMT only provides a 53-bit integer divide

# Details:

- If either x or y is more than 53 bits, the float extension will be raised and the trap handler will complete the operation successfully.
- Sign extension is handled correctly, so for example, INV\_DIV\_CHOP/INV\_DIV\_FLOOR support values in the range [-2<sup>53</sup>, 2<sup>53</sup>-1], whereas UNS\_DIV supports values in the range [0, 2<sup>53</sup>-1].

The performance issue was that the modulo operator was being applied to integers generated to be random across all 64-bit unsigned integers.

#### Modified code snippet which avoids 64-bit divide

```
unsigned int mask = (1 << 53) - 1;
#pragma mta assert nodep *endpoints
#pragma mta assert nodep *ends
for (i = 0; i < no_of_edges; i++) {
    endpoints[i] = (ends[i] & mask) % N2;
}
```

#### Loop1: time copying to ends 6.654919e-02

```
for (i = 0; i < no_of_edges; i++) {
    ends[i] = endpoints[i] + 1;
}</pre>
```

ticks: 33258200, secs: 0.066516, issues: 104491617, memrefs: 101756261, traps: 0, retries: 705688, concurrency: 2541, streams: 754

#### Loop2: time modulo 8.084661e-02

```
unsigned int mask = (1 << 53) - 1;
for (i = 0; i < no_of_edges; i++) {
      endpoints[i] = (ends[i] & mask) % N2;
}
```

ticks: 40407611, secs: 0.080815, issues: 155276319, memrefs: 102072724, traps: 0, retries: 673585, concurrency: 1939, streams: 720

### Looking at Assembly Code

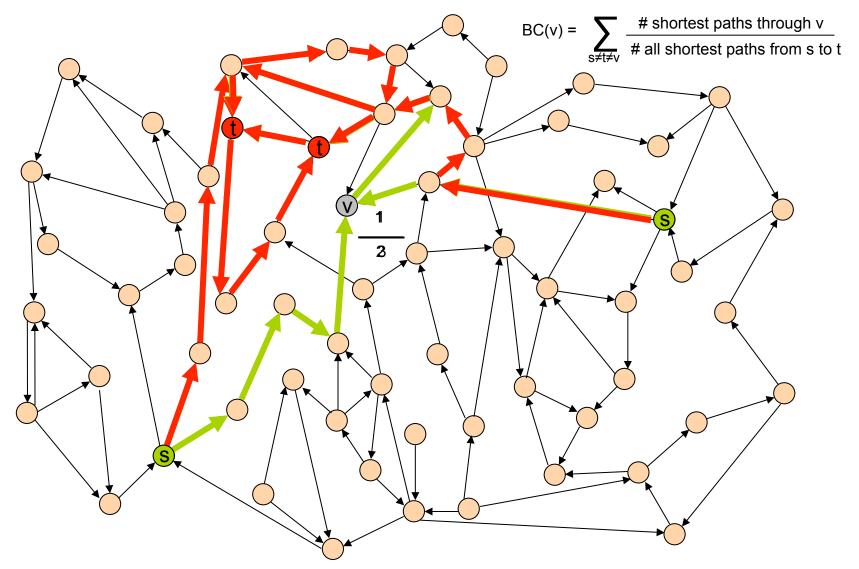
- From the command line dis executableName
- From within mdb
  - mdb> disass

0x004ac9ec	(inst 2 (STORE_DISP r25 r1 (* 8 2)) (NOP))
0x004ac9ed	(inst 1 (STORE_DISP r19 r1 (* 8 9)) (NOP))
0x004ac9ee	(inst 0 (STORE_DISP r21 r1 (* 8 12)) (NOP))
0x004ac9ef	(inst 0 (NOP) (INT_ADD_IMM r30 r0 15) (CLOCK r26 r0))
0x004ac9f0	(inst 0 (REG_LOAD_DISP r25 r31 (* 8 10)) (BIT_AND r30 r26 r30))
0x004ac9f1	(inst 0 (REG_LOAD_DISP r3 r31 (* 8 10)) (INT_ADD_IMM r25 r25 104))
0x004ac9f2	(inst 2 (LOAD_DISP r28 r31 (* 8 19)) (INT_ADD_IMM r3 r3 96))
0x004ac9f3	(inst 0 (LOAD r25 r25) (INT_ADD_IMM r29 r0 216) (REG_MOVE r10 r9))
0x004ac9f4	(inst 3 (LOAD r3 r3) (TARGET_RESTORE t5 r25) (REG_MOVE r5 r27))
0x004ac9f5	(inst 0 (LOAD_AC_INDEX r28 r28 ( (fe_control FE_FUTURE)) r30) (TARGET_DISP t1



### SSCA #2, Kernel 4, *Betweenness Centrality*

### **Betweenness Centrality of a Vertex**



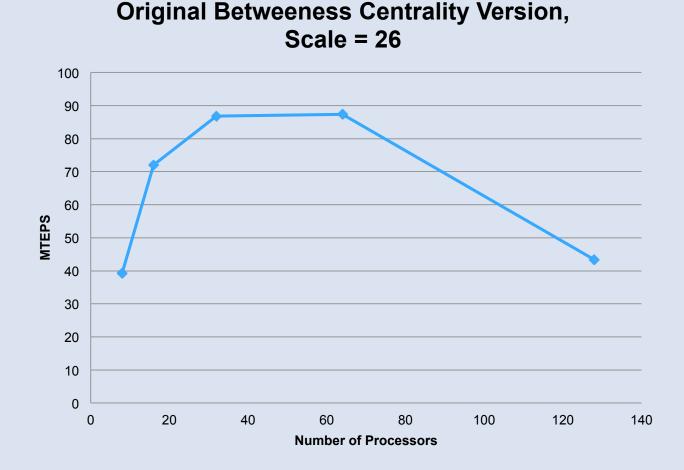
## **Our Original Implementation**

CRAY

Large data structures => sequential outer loop

```
/* Use |Vs| nodes to compute centrality values */
  for (s = 0; (s < NV) \& (Vs > 0); s ++) {
  •••
#pragma mta assert no dependence
      for (j = QHead[nQ - 1]; j < QHead[nQ]; j++) {
          ...
          int myStart = start[v];
          int myEnd = start[v + 1];
#pragma mta assert no dependence
          for (k = myStart; k < myEnd; k++) {
              ...
```

#### **Original Version Didn't Scale Well**



Note: MTEPS=millions of traversed edges per second, estimate defined as #vertices \* 7 / execution time.

## **Trying to Scale Better**

- CRAY
- With larger systems, we thought we could run a few iterations of the outermost loop in parallel
- First attempt:

```
#define BFS_THREADS 16
...
#pragma mta assert parallel
  for(num_threads=0; num_threads < BFS_THREADS; num_threads ++){
  for(;;) {
    start_vertex = int_fetch_add(&Vs_ptr,1);
    if (start_vertex > Vs -1) break;
    ...
```

```
#pragma mta assert no dependence
    for (j = QHead[nQ - 1]; j < QHead[nQ]; j++) {</pre>
```

## **Resulting Canal Listing**



- Performance was horrible.
- Canal:

... 7 pXX | for (j = QHead[nQ - 1]; j < QHead[nQ]; j++) { ...

- What happened:
  - Compiler assumed there was sufficient parallelism in the outer loop.

## Second Try



• Using "loop future" parallelism

```
#pragma mta assert parallel
#pragma mta loop future
for(num_threads=0; num_threads < BFS_THREADS; num_threads ++){
...</pre>
```

- Performance was bad.
- Canal output:

Parallel region 7 in cenTrality in loop 6 Single processor implementation

Parallelism was confined within BFS\_THREADS processors.

## **Third Approach**



#### • Use "future" variables

#define BFS THREADS 16

...

•••

}

future int thread\_id[BFS\_THREADS];

for (num\_threads=0; num\_threads < BFS\_THREADS; num\_threads++) {
 touch (&thread\_id[num\_threads]);</pre>

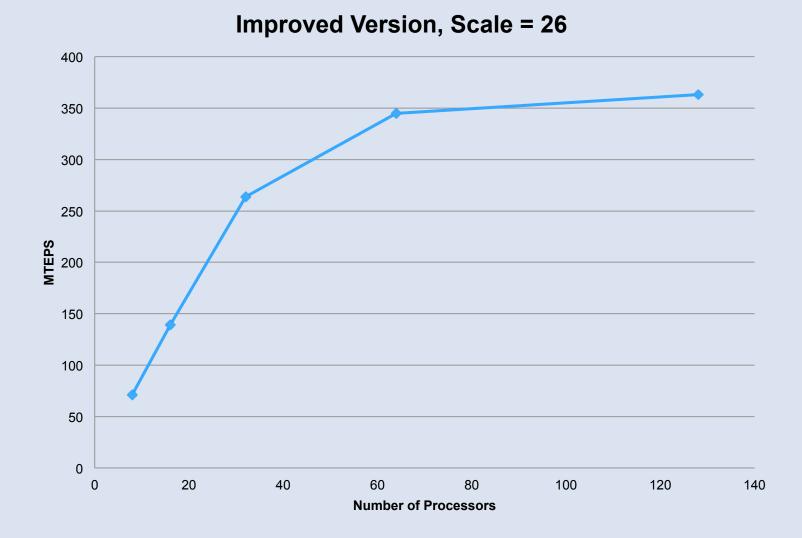


#### **#pragma mta max** *n* **processors**

- Limits the number of processors used on a multiprocessor loop to n
- *n* must be a compile-time integer constant > 0
- for collapsible loop nests, the max processors value for the collapsed loop is the same as that specified for the outer loop

```
/* Use at most 4 processors. */
#pragma mta max 4 processors
for(i = 0; i < size; i++) {
    array[i] += array[i] + (size + i);
}</pre>
```

#### Much Improved Performance and Scaling



CRA

#### **Lessons Learned**

CRAY

- We used future variable parallelism because
  - The amount of parallelism in the outermost loop was small
  - The amount of parallelism in the inner loops was large and dynamic
- We used #pragma mta max n processor to control how wide each of the inner parallel loops could grow
  - Better load balancing across outer loops
  - Avoid hotspotting in the Queues

## This also works



- Use "mta loop future" pragma on outer loop, but have inner loop inside subroutine call
  - Loop parallelism inside Process\_cenTrality subroutine call is now multi-processor
  - OK to do it this way, but not the intended use of the loop future pragma. More targeted for large loop bounds.

```
#define BFS_THREADS 16
...
// Outer loop to run independent BFS processes
#pragma mta loop future
for (num_threads=0; num_threads < BFS_THREADS; num_threads++) {
    Process_cenTrality(G, BC, Vs, &Vs_ptr, ...
...</pre>
```



#### **#pragma mta max** *n* **processors**

- Limits the number of processors used on a multiprocessor loop to n
- *n* must be a compile-time integer constant > 0
- for collapsible loop nests, the max processors value for the collapsed loop is the same as that specified for the outer loop

```
/* Use at most 4 processors. */
#pragma mta max 4 processors
for(i = 0; i < size; i++) {
    array[i] += array[i] + (size + i);
}</pre>
```



#### **#pragma mta max concurrency c**

- Limits the number of processors used by a multiprocessor loop to max(1, c/<num\_streams\_per\_processor>), where <num\_streams\_per\_processor> is the number of streams the compiler requests for each processor used by the parallel loop.
- Limits the number of <u>streams</u> used by a single processor parallel loop to min(c, <max\_streams\_per\_processor>)
- c is a compile-time integer constant > 0

```
/* Use at most 512 streams across all processors. */
#pragma mta max concurrency 512
for(i = 0; i < size; i++) {
    array[i] += array[i] + (size + i);
}</pre>
```

## **Parallelism-Limiting Pragmas**



Using them together...

```
/* Use at most 512 streams across all processors or
    at most 8 processors, whichever is smaller */
#pragma mta max concurrency 512
#pragma mta max 8 processors
for(i = 0; i < size; i++) {
    array[i] += array[i] + (size + i);
  }
```

 Loop future loops can only use the max concurrency pragma:

```
/* Create at most 512 futures. */
#pragma mta loop future
#pragma mta max concurrency 512
for(i = 0; i < size; i++) {
    array[i] += array[i] + (size + i);
}</pre>
```

## New Pragmas in XMT PE Release 1.4 (1)

## **#pragma mta for all streams**

starts a parallel region

}

- executes the statement or block of statements exactly once for each stream allocated to the region
- acts like an "assert parallel" pragma
- can be used in conjunction with "use n streams" (but no guarantee that many will be allocated)

#pragma mta use 100 streams

**#pragma mta for all streams** 

```
{ //do parallel stuff
```

## New Pragmas in XMT PE Release 1.4 (2)

#### #pragma mta for all streams *i* of *n*

- Sets n to the total number of streams executing the region
- Variable *i* is a unique per-stream identifier;  $0 \le i \le n-1$

```
int istr, ntotal;
```

int check\_in\_array[SOME\_BIG\_NUMBER];

```
for( istr=0; istr<SOME_BIG_NUMBER; istr++ )
    check_in_array[istr] = 0;</pre>
```

#pragma mta for all streams istr of ntotal
{ check\_in\_array[istr] = 1;
 printf("Stream %d of %d checked in.\n", istr, ntotal);



# Generating Graphs on the XMT



#### **Generating synthetic graphs**

- Random graphs
- R-MAT graphs

## Random Graph



- Or Erdös-Renyi graph
- Given the number of vertices n, average number of outedges per vertex (out-degree) x,
- Return a randomly-generated graph with *n* vertices with a uniformly-distributed out-degree with average x
- Random graphs are pretty "well-behaved"
  - Partition reasonably well

### What Our Random Graph Generator Does

- SSCA #2 Handout 1
- Given desired # vertices (int), and desired average outdegree (double), generates random graph
- How it works:
  - Generate array, length numEdges, with random vertex IDs in it

randNeighbors(Neighbors, nN, numEdges,(double\*)Marked, nN);

– Histogram the IDs

numNeighbors[Neighbors[i]]++;

- This decides how many out-edges each vertex has.
- Now generate another array of numEdges random vertex IDs

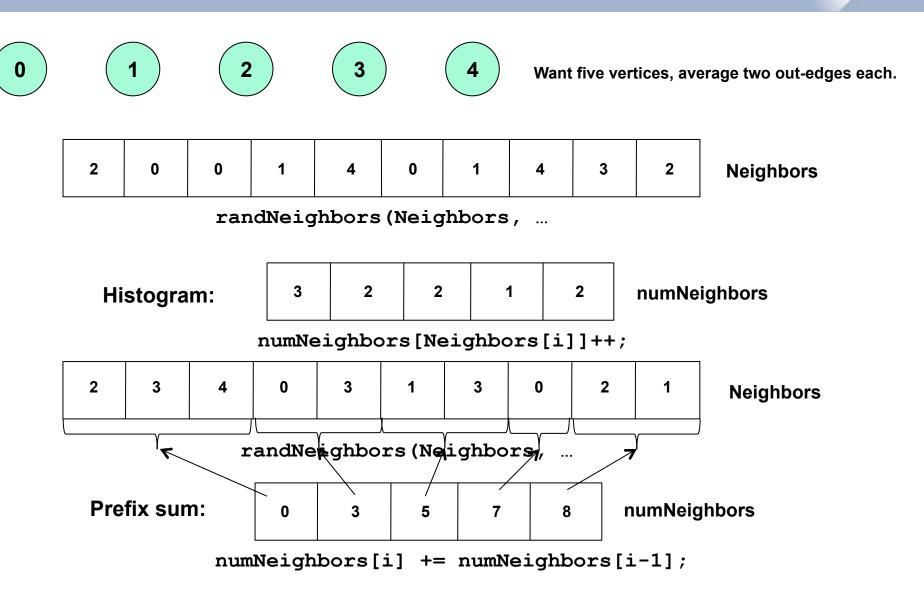
randNeighbors(Neighbors, nN, numEdges,(double\*)Marked, nN);

- Prefix-sum the numNeighbors array

numNeighbors[i] += numNeighbors[i-1];

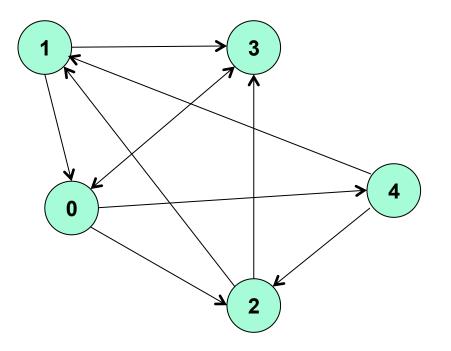
- numNeighbors[i] now holds the starting index in the Neighbors array of the neighbors of vertex i

#### **Random Graph Generator Example**



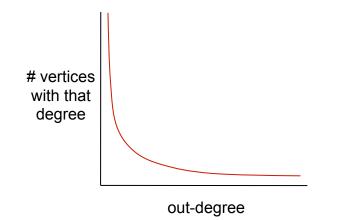
#### **Resulting Graph**



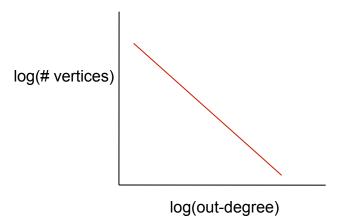


#### **Power Law Graphs**

Number of vertices with out-degree x = cx<sup>-β</sup>



Straight line on a log-log graph, with slope -β



## **About Power Law Graphs**

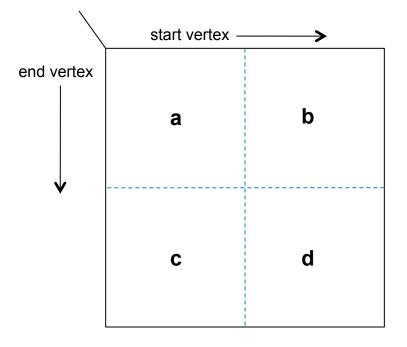
- A few vertices with huge out-degree ("heavy-tailed" distribution)
- Contain clusters of "communities" (connectivity inside > connectivity outside)
- Found to be much closer fit to many "real world" social networks
  - WWW
  - Internet router graphs
  - Citation graphs
  - Facebook
  - ...
- Small diameter but high connectivity ("six degrees of separation")
- AKA "self-similar", "scale-free" graphs

#### **Popular Power Law Graph Generator**

- R-MAT = "recursive matrix"
- Invented by Chakrabarti, Yahoo! Research and Falutsos, CMU
- Four probability parameters, a, b, c and d

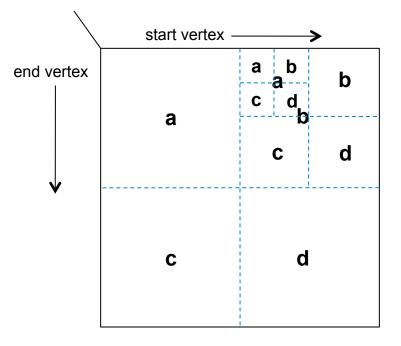
– Three, really: d = 1.0 – (a + b + c)

Divide the graph adjacency matrix into four quadrants



#### **Power Law Graph Generator**

- Apply recursively until you reach a single cell
- Our implementation tweaks a, b, c and d ± at most 10% every iteration
- It also uses a hash table to eliminate duplicate edges
- Typical parameters are a=.55, b=.19, c=.19, d=.07
- SSCA #2 Handout 2 point to file instead???





#### Parallel I/O Using Snapshot/Restore

## **Avoiding Graph Generation**

- Generating synthetic data is often slow e.g. graph generation
- We try to avoid it, when possible, by using Snapshot-Restore.
- STEPS:
  - What you put in the source code: SSCA #2 Handout 3
  - How you initialize for using the parallel file system

In the current configuration, the parallel file system is / mnt/lustre, a Lustre file system When fully configured, matterhorn will attach to a DVS filesystem

Snapshot/restore will work identically on both systems



#### The #includes needed, and some file names

#include <luc/luc\_exported.h>
#include <snapshot/client.h>

const char VE\_FILENAME[] =
 "/scratch/dmizell/mydata/rmat28\_vertex\_edge\_info.data";
const char SV2\_FILENAME[] =

"/scratch/dmizell/mydata/rmat28\_sv2\_snapshot.data"; const char EV2 FILENAME[] =

"/scratch/dmizell/mydata/rmat28\_ev2\_snapshot.data"; const char START FILENAME[] =

"/scratch/dmizell/mydata/rmat28\_start\_snapshot.data"; const char WEIGHT\_FILENAME[] =

"/scratch/dmizell/mydata/rmat28\_weight\_snapshot.data";



#### **Initializing the Snapshot library**

```
// Initialize Snapshot Library
// The SWORKER_EP environment variable is read at this point.

if (err = snap_init() != SNAP_ERR_OK)
{
fprintf(stderr,
"Failed to initialize snapshot library. Error %d. \n", err);
}
```

## CRAY

#### Writing files

```
/* Memory allocated for the vertex/edge info data structure */
int *veptr = (int *)malloc(2*sizeof(int));
veptr[0] = G->numVertices; veptr[1] = G->numEdges;
```

// All file system workers must be able to access the specified path. numbytes = (size\_t)(2\*sizeof(int)); err = snap\_snapshot ((char \*)VE\_FILENAME, veptr, numbytes, &snapError);

//... more file writes...



#### **Reading files**

```
// Restore the Graph from disk
numbytes = (size_t)(2*sizeof(int));
err = snap_restore ((char *)VE_FILENAME, veptr, numbytes, &snapError);
G->numVertices = veptr[0]; G->numEdges = veptr[1];
//...
// All file system workers must be able to access the specified path.
numbytes = (size_t)(G->numEdges*sizeof(int));
err = snap_restore ((char *)SV2_FILENAME, G->startVertex, numbytes, &snapError);
//... more file reads...
```



#### **Exercise 4**

#### /mnt/lustre/Workshop/Exercise4

Dataflow exercise /mnt/lustre/Workshop/Homework/wavefront

Integer Sort /mnt/lustre/Workshop/Homework/sort



#### The End!

#### Thanks!



#### End of Day 2



## Backup Slides



#### Making Reductions and Recurrences Run in Parallel

Thanks to Jon Berry, Sandia

#### **Example 1**



#### Summing the absolute values of integers

• First try (compiler didn't parallelize it)

```
int total=0;
for (int i=0; i<n; i++) {
    if (v[i] < 0) {
        total += -v[i];
    } else {
        total += v[i];
    }
}
```

## Second Try



#### **Parallelized successfully**

```
int total=0;
for (int i=0; i<n; i++) {
    int incr = (v[i] < 0) * -v[i] + (v[i] >= 0) * v[i];
    total += incr;
}
```

## **A Conditional Reduction**

This didn't parallelize:

• This did:

