

SVM Update

Francois Labonte

11/10/03

Reservoir HLC

- The Reservoir HLC has been release last Wednesday (11/05).
 - To members of DARPA PCA
 - Do we all have access to it?
- Includes a brook2c
- And a simple svm

SVM simulator status

- Being tested with Reservoir produced code
 - Currently not producing valid result for FEM
 - Working on it at the moment
- Problem might be in the way that the compiler does gathers now...
 - Scalar processor copies elements in the gather order, then stream processor loads stream normally.... Not ideal in terms of performance either.

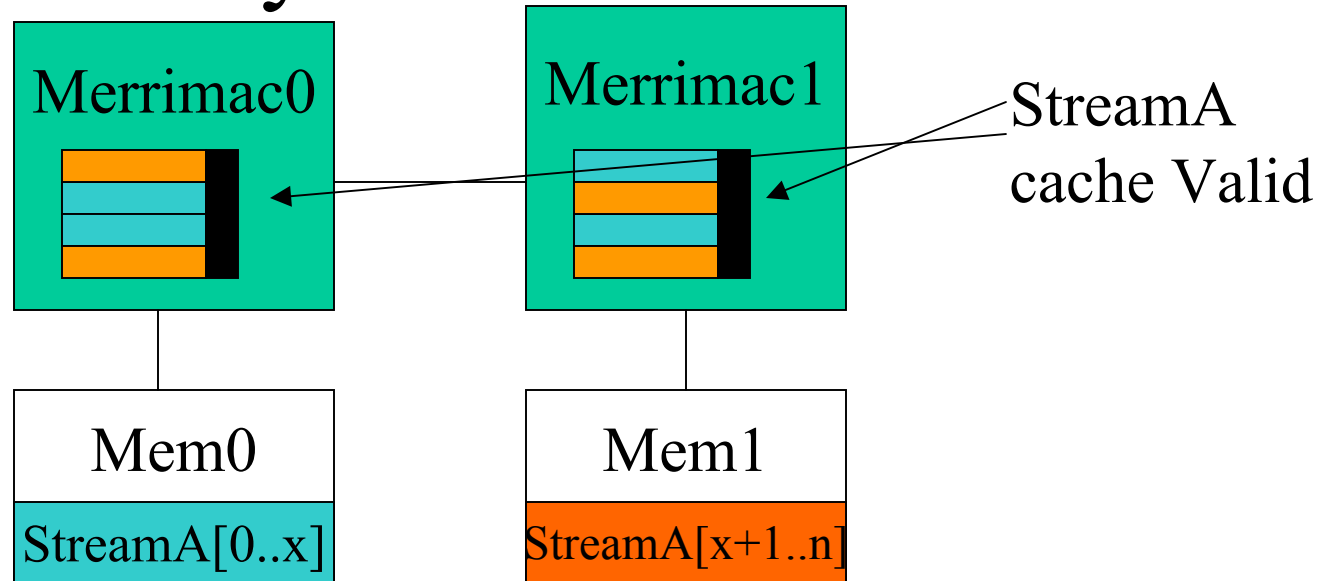
Looking forward

- SVM API
 - ScatterOP, GatherOP need to be added
 - Control of stream Cache behavior
 - Want a stream to be cacheable remotely for a phase of the program, then flushed
- Simulator
 - Work with Reservoir compiler (obviously)
 - Automatic kernel schedule/runt-time eval integration
 - Stream cache integration
 - Scalar processor run-time eval

SVM API

- ScatterOp and GatherOp
 - simple to add to the SVM, they are “black box kernels” that run on the DMA engine
 - Jayanth will have to do most of the work to have the compiler use generate the calls.
- Stream Cache control
 - Need to specify a stream is cacheable,
 - Need to flush a stream from the cache
 - Scatter-Op, GatherOp special case


Stream Cache Case 1: Read-only cache for streams



- Gather on stream with strong affinity (expect re-use within a node, so want local caching)
- When stream gets updated, gang invalidate all lines in cache from stream A
- Could have multiple streams cache like this (one valid bit each) though practically too many streams would cross-pollute each other

Stream Cache Case 2

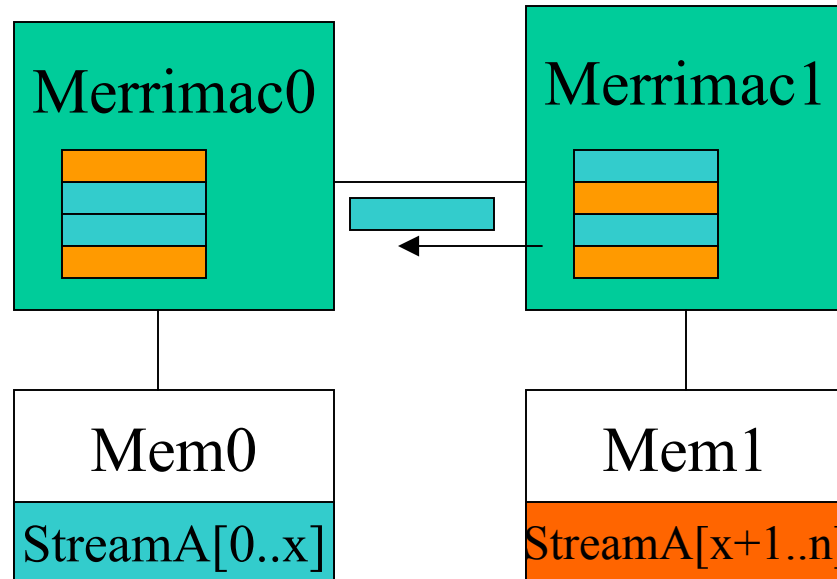
Scatter Add local accumulate



- Scatter add into your own cache, start with 0 if not a local memory address
- When a remote cache line is evicted, scatter add to the local node, or when all scatter add are done... (when is that exactly because you might scatter add for multiple strips) need a scatter add flush command (in the SVM case another black box kernel?)

Stream Cache Case 3

Gather Op local aggregate



- Gather Add would mostly be used to advance a tail pointer
- In the binning case, not a lot of bins, so a lot of addresses in common, and we are adding the same value (sizeof(struct)) each time
- Could aggregate numerous results into stream cache (Send when evicted or at the end of gatherop),
- Need to expand results re-order properly (might make this too complicated)

SVM Simulator: Kernel performance estimates

- Currently kernel run-time estimate is of the form $t = \text{preamble} + \text{loop_size} * \text{stream_size} + \text{postscript}$
- Abhishek told me he was going to work on taking the Reservoir kernel IR and converting it to the kernel scheduler IR (isched) then we will get the kernel schedule to use for the SVM simulator
- Note that this run-time estimate doesn't account for data dependent (conditional stream) type of kernels,

SVM Simulator: Stream Cache integration

- Besides a simple cache simulator, we need :
 - Gang invalidate
 - Local Scatter Add
 - On final scatter add, need to do scatter-add at local node.
- Need to account for cache in performance estimates

Accounting for scalar processor run-time

- Currently, only stream memory operations and kernel run-times are evaluated by the SVM simulator. Missing
 - Scalar processor holding up stream processor
 - Scalar processor memory transactions
 - Scalar caches?
- Probably desirable, the easiest solution would probably be to use Tim Knight's GDB MIPS simulator.

Conclusion

- SVM simulator working with Reservoir...
this week
- Scatter Op/Gather Op ... 2 weeks from now
- Stream Cache ... 4 weeks from now