



European Processor Initiative to deliver a RISC-V based vector accelerator. We provide a HW and SW ecosystem to enable co-design before the chip arrives.



We also develop an evaluation methodology to study potential optimizations. Our optimization techniques are CPU-agnostic and provide benefits in other architectures.



Short reasons for long vectors in HPC CPUs: a study based on RISC-V DOI: 10.1145/3624062.3624231



Gerard Oliva (corresponding author) gerard.oliva@bsc.es



Pablo Vizcaino pablo.vizcaino@bsc.es



Urs Ganse urs.ganse@helsinki.fi



Filippo Mantovani filippo.mantovani@bsc.es



Explicit loops

Marta garcia-Gasulla marta.garcia@bsc.es



Vlasiator is a novel, computationally intensive simulation that models ions differently for more accurate, noiseless space weather predictions. Space weather describes the Sun's variable effects on near-Earth space via solar wind, impacting technology like satellites and power grids.

Code structure



vector load a What the compiler sees vector load b What the programmer codes (overloaded operators) vector add vector store res1 Vec foo(Vec a, Vec b, Vec c){ Vec foo(Vec a, Vec b, Vec c){ Vec res1, res, res3; Vec res1, res, res3; vector load res1 for(int i=0;i<VECL;++i) res1[i]=a[i]+b[i];</pre> res1 = a+b;vector mul.f for(int i=0;i<VECL;++i) res2[i]=res1[i]*0.5;</pre> vector store res2 res2 = res1*0.5; for(int i=0;i<VECL;++i) res3[i]=res2[i]-c[i];</pre> res3 = res2-c; return res3; return res3; / vector load res2 vector load c vector sub vector store res3 8 memory instr 3 arithmetic instr Better approach: Explicit loops vector load a Vec foo(Vec a, Vec b, Vec c){ vector load b Vec res1, res, res3; vector add vector mul.f for(int i=0;i<VECL;++i){ ---</pre> vector load c res1[i] = a[i]+b[i]; vector sub res2[i] = res1[i]*0.5; vector store res3 res3[i] = res2[i]-c; 4 memory instr return res3; 3 arithmetic instr Autovec 📕 Explicit Loops 1.00E+8 1.00E+7

Vector class

Templated C++ class Defines common vector operations Vector length is always known at compile time Vlasiator code heavily depends on Vec

Original: Written for fixed-length vectors template <class T> static inline Vec4Simple<T> operator + (const Vec4Simple<T> &l, const Vec4Simple<T> &r){ return Vec4Simple<T>(l.val[0]+r.val[0], l.val[1]+r.val[1], l.val[2]+r.val[2], l.val[3]+r.val[3]);

Phases





VecX < 16 >

Vector length is no longer limited by the implementation

Extended version: Vector-length agnostic

template <size_t N, class T> static inline VecX<N, T> operator + (const VecX<N, T> &l, const VecX<N, T> &r){ VecX<N, T> vec; PRAGMAS for(unsigned long int i=0; i<l.get_size(); ++i){</pre> vec.val[i] = l.val[i] + r.val[i];

return vec;

Increase Average Vector Length

Avg.

Metric definition



Compile time parameter VECL Affects: Problem size $VECL \le 64$

VECL=8

	vadd	vfadd	vfmul	vle	vse	vfsub	vmflt
Avg. VL	8	8	8	8	8	8	8
#Instr.	1	2	4	21	14	2	3

VECL=16									
	vadd	vfadd	vfmul	vle	vse	vfsub	vmflt		v
Avg. VL	16	16	16	16	16	16	16	Avg. VL	
#Instr.	1	2	4	21	14	2	3	#Instr.	

e	P q	roposal roblem	: Increa size an	se VE d lev	ECL to erage	o incren e long v	nent ectors
	J .				5	3	
			VE	CI.=:	32		
	vadd	vfadd	vfmul	vle	vse	vfsub	vmflt

	vadd	vfadd	vfmul	vle	vse	vfsub	vmflt	
Avg. VL	32	32	32	32	32	32	32	
#Instr.	1	2	4	21	14	2	3	

	VECL=64							
	vadd	vfadd	vfmul	vle	vse	vfsub	vmflt	
Avg. VL	64	64	64	64	64	64	64	
#Instr.	1	2	4	21	14	2	3	



Increase in Vector Instruction Mix

Metric definition



Code analysis

inline Vec [add/mul/sub/...](Vec op1, Vec op2)

Vec R1 = add_vector(A,C); Vec R2 = mul_vector(A,B); Vec R3 = sub_vector(C,R2)

Arguments passed by value Functions are called back-to-back

Proposal: Pass arguments by reference inline Vec [add/mul/sub/...](Vec & op1, Vec & op2)





Conclusions and next steps





a Auto-vectorization alone does not provide significant speedup

Adding explicit loops achieves a speedup of slightly more than 4x

Eliminating scalar copies (in addition to explicit loops) increases the speedup to slightly more than 7x

These speedups were achieved using a vector length of 64 (VECL = 64) Limiting factor: allow VECL greater than 64 Possible benefit: Avg. VL up to 256

Further analysis of Vector Mix:

- Eliminate non-computing instructions like mask creation and data shuffling - Compile and run/emulate with rvv1.0 vector specification

Run the full Vlasiator application to study the effect of optimizations derived from the mini-app

Evaluate portability of code changes to other CPU architectures