

# Supporting Vector Programming on a Bi-Endian Architecture

Bill Schmidt, Ph.D.

Linux on Power Toolchain Development

IBM Linux Technology Center

Michael Gschwind, Ph.D.

System Architecture

IBM Systems and Technology Group


# Outline

- History
- Little Endian on Power
- Endianness (bytes and elements)
- Programming models
- Example vector interfaces and implementations
- Optimization
- Implementation status (gcc, llvm)

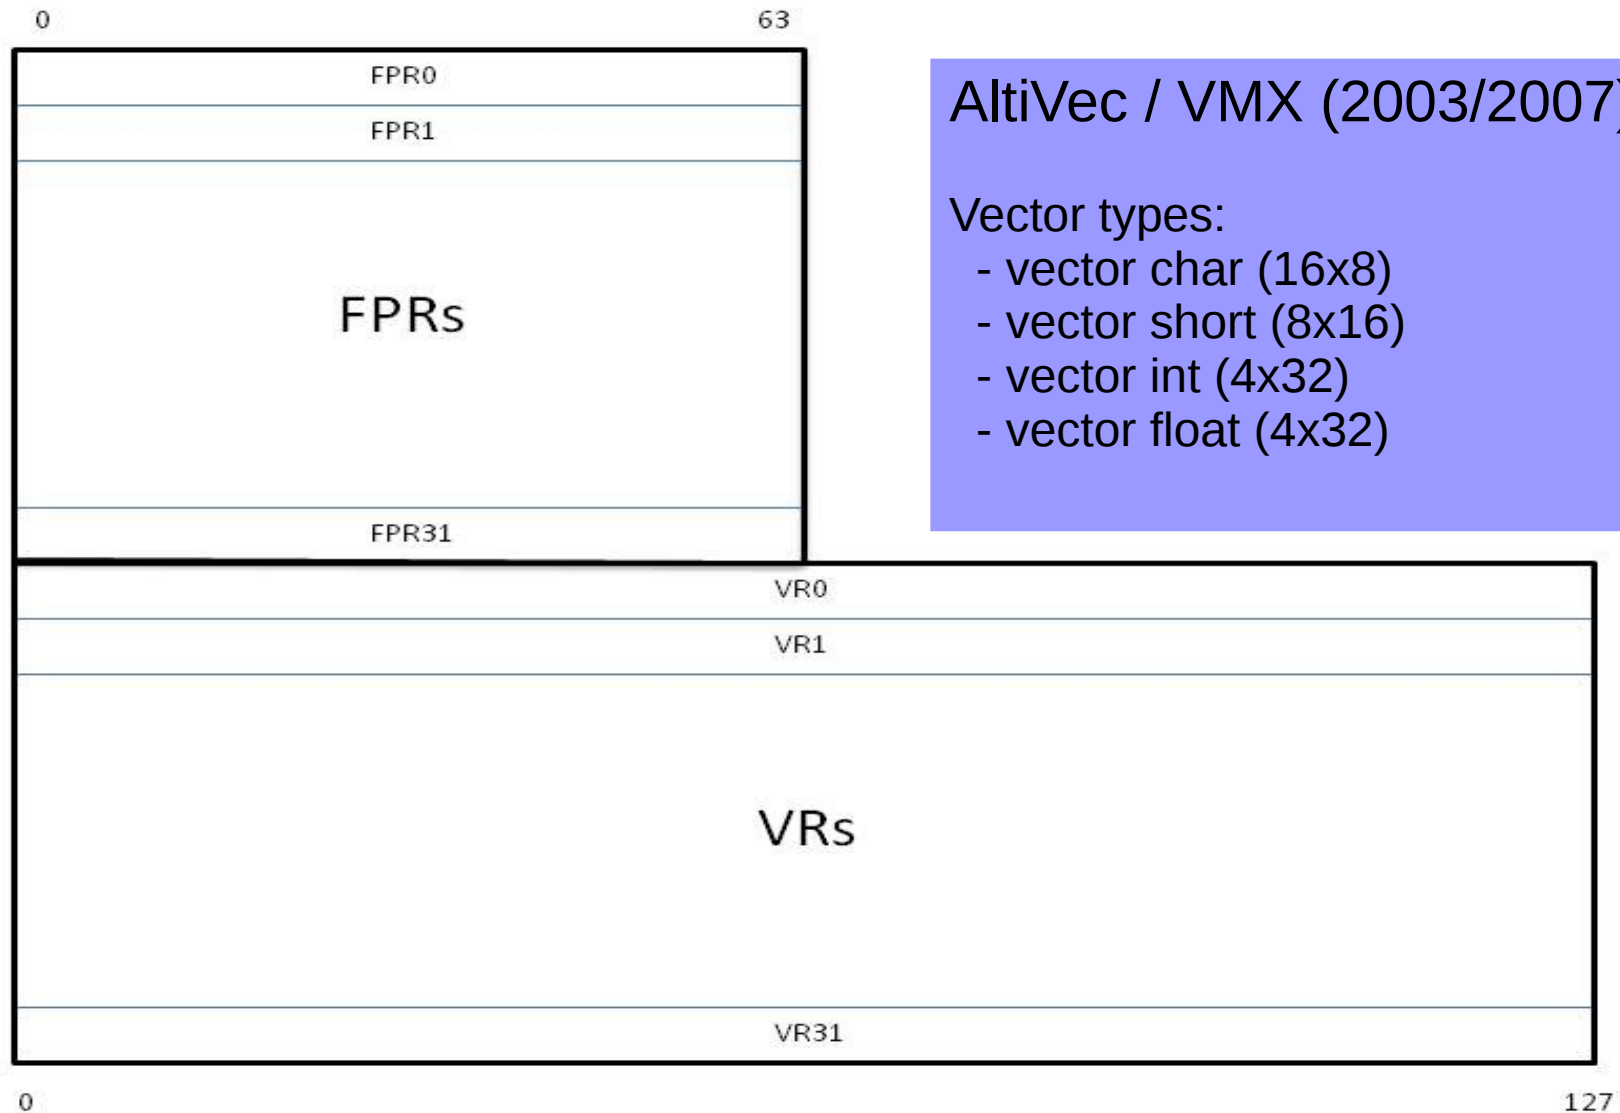
# Outline

- **History**
- Little Endian on Power
- Endianness (bytes and elements)
- Programming models
- Example vector interfaces and implementations
- Optimization
- Implementation status (gcc, llvm)

# History: Floating-point registers



# History: Vector registers

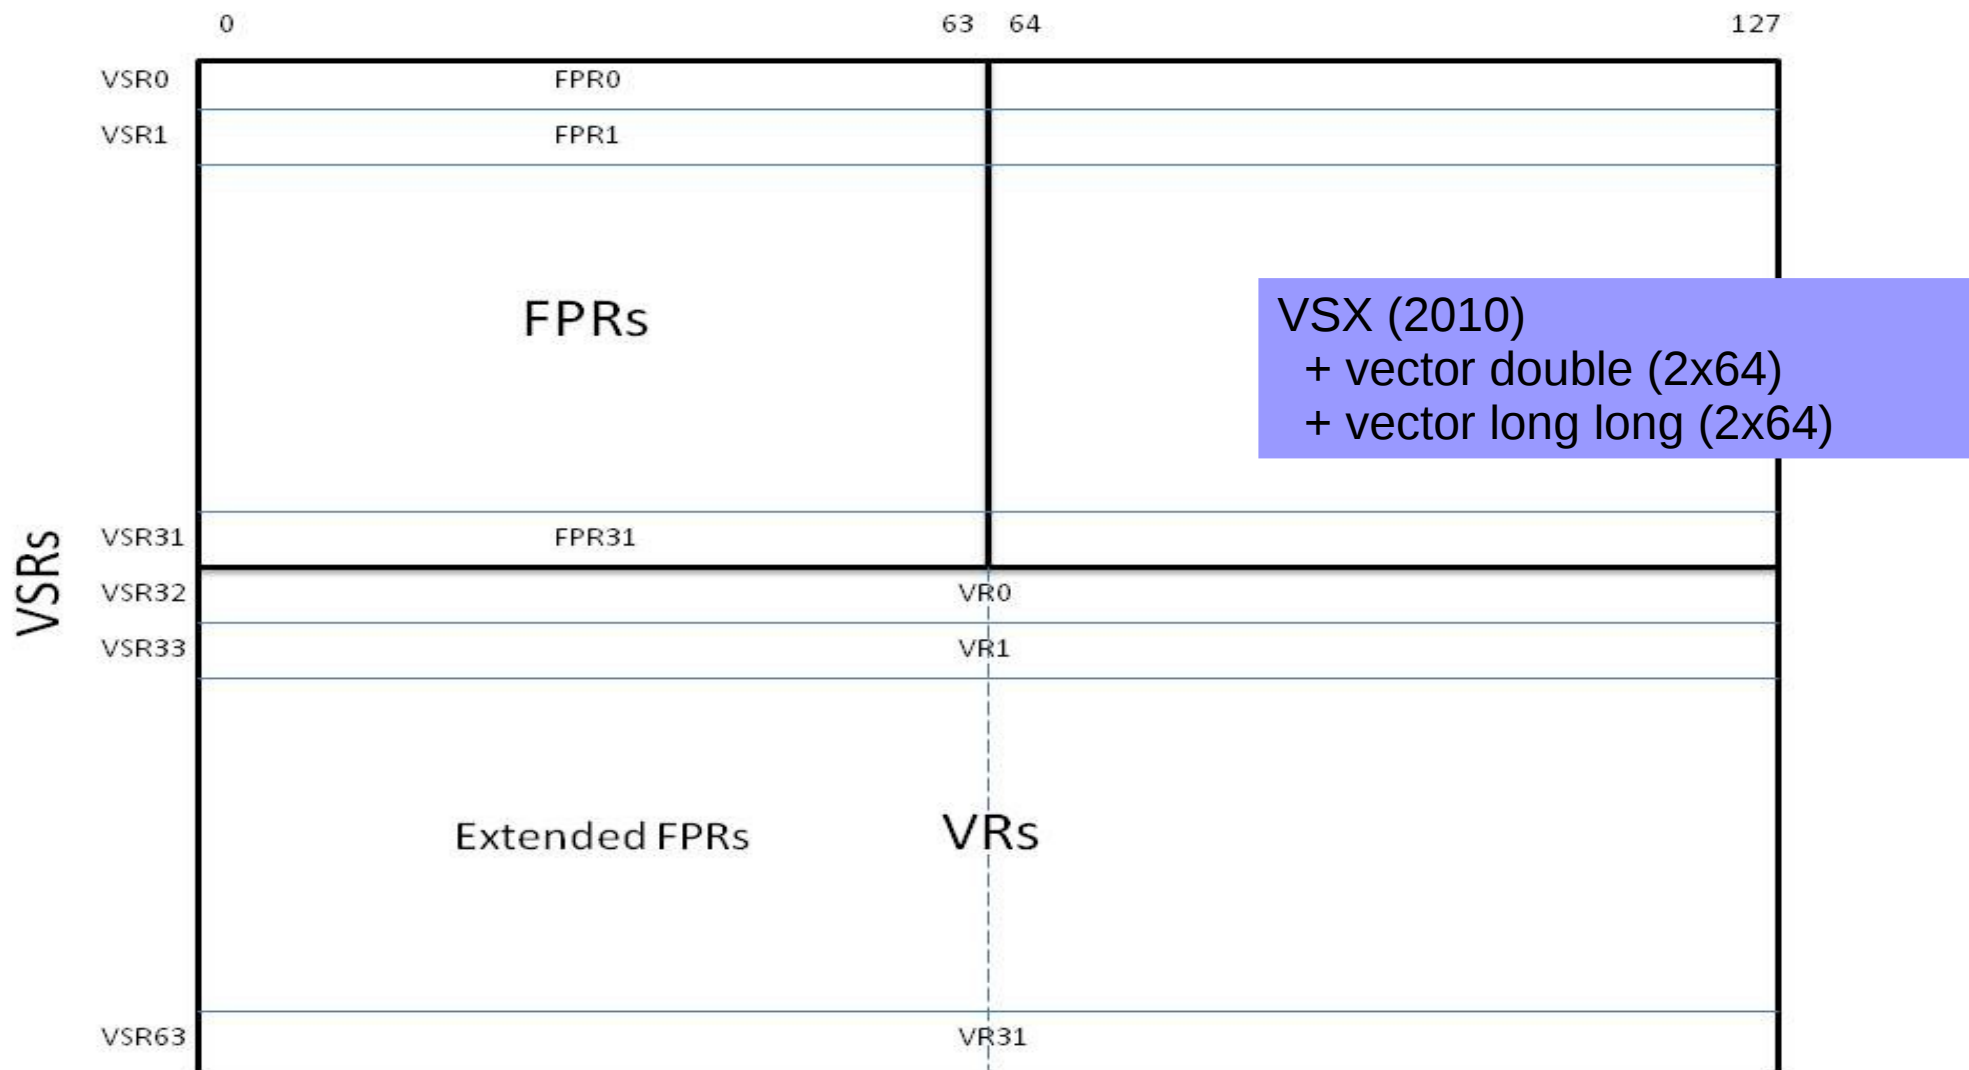


## Altivec / VMX (2003/2007)

### Vector types:

- vector char (16x8)
- vector short (8x16)
- vector int (4x32)
- vector float (4x32)

# History: Unified vector-scalar registers



## History: Power ISA Endianness

- The Power ISA has been bi-endian since its inception (mode selectable by supervisor/hypervisor).
- IBM's server operating systems have always been big endian.
- Server program models for vector programming have also been exclusively big-endian.
  - VMX contains big-endian “biases” (inherent byte and element numbering), except for memory accesses.
  - VSX adds some vector memory access instructions with big-endian biases.
  - Silicon and opcode space limits bi-endian implementation.

# Outline

- History
- **Little Endian on Power**
- Endianness (bytes and elements)
- Programming models
- Example vector interfaces and implementations
- Optimization
- Implementation status (gcc, llvm)



## Little Endian for Linux on Power – Why and How?

- Changes in technology and market forces make a Little Endian offering attractive.
  - Slower growth in single-thread CPU performance
    - Accelerator hardware becomes more attractive.
    - Existing GPGPUs use LE data layout.
  - Growth in Linux share of server market
    - Most Linux apps developed for LE systems.
    - Though porting from LE to BE is usually not difficult, the perception is otherwise – obstacle to “mindshare” – and there are exceptions.
- LE Linux on Power distributions
  - Ubuntu Server (14.04, 14.10, 15.04, ...)
  - SLES 12
  - RHEL 7.2 (not yet released)
- All offerings are 64-bit only.

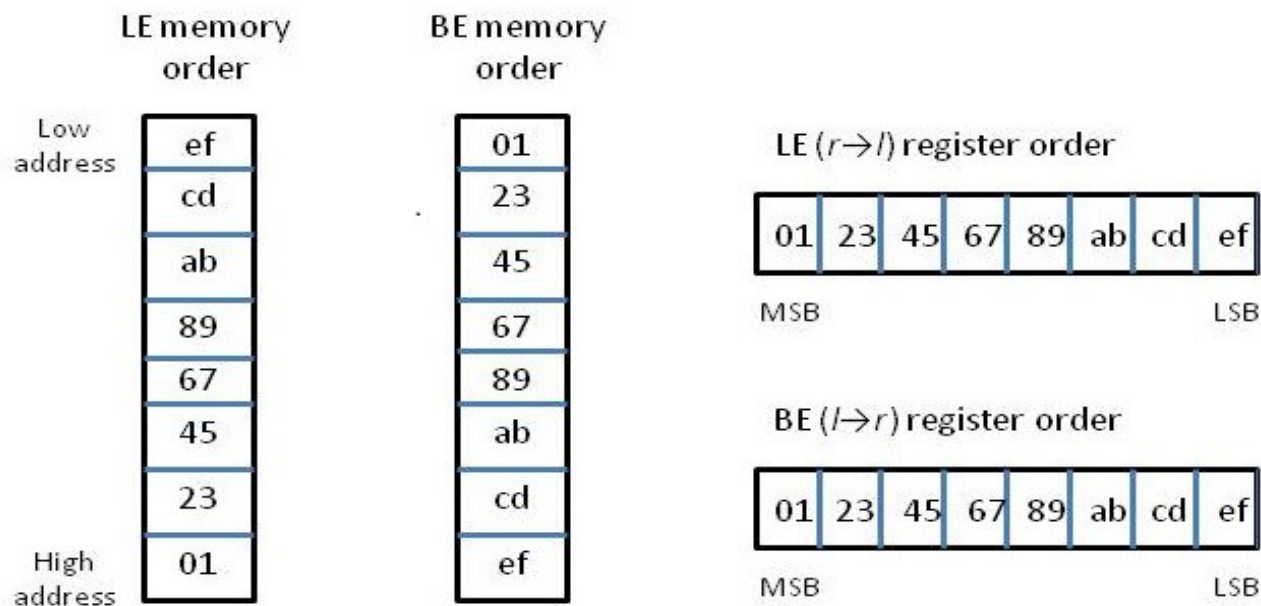
# Outline

- History
- Little Endian on Power
- **Endianness (bytes and elements)**
- Programming models
- Example vector interfaces and implementations
- Optimization
- Implementation status (gcc, llvm)


## Byte order and element order: Scalars

### ■ Scalar values


- In registers, BE and LE scalars are identical.
- BE: Most significant byte (MSB) stored at lowest memory address
- LE: Least significant byte (LSB) stored at lowest memory address



# Byte order and element order: Byte arrays



LE ( $r \rightarrow l$ ) register order



MSB

LSB


BE ( $l \rightarrow r$ ) register order




MSB

LSB


# Byte order and element order: Word arrays



LE ( $r \rightarrow l$ ) register order



BE ( $l \rightarrow r$ ) register order



# Outline

- History
- Little Endian on Power
- Endianness (bytes and elements)
- **Programming models**
- Example vector interfaces and implementations
- Optimization
- Implementation status (gcc, llvm)

## Vector programming models: Existing

- Power supports an extensive vector interface, accessed via `<altivec.h>`.
  - Many hundreds of vector intrinsics, overloaded by type:

```
vector int va, vb, vc;  
va = vec_add (vb, vc);  
vector float vfa, vfb, vfc;  
vfa = vec_add (vfb, vfc);
```
  - Each intrinsic maps to a single VMX or VSX instruction.
  - Supported by GCC, LLVM, and XL compilers on Linux, AIX, and other platforms
  - Uses  $l \rightarrow r$  element order where ordering matters
- How should this model be modified for different consumers on little endian?

## Vector programming models

- Conflicting goals
  - LE programmers expect  $r \rightarrow l$  element ordering.
  - GCC and LLVM do also.
  - But: Experienced POWER programmers are used to  $l \rightarrow r$  ordering.
    - BE vector math libraries created with this assumption
  - Many VMX/VSX instructions encode  $l \rightarrow r$  element numbers.
    - Explicitly: `vspltw v2,v1,0` (duplicate zeroth word from the left)
    - Implicitly: `merge-high/low`, `multiply even/odd`
- So: Two LE programming models provided
  - True-LE: LE memory order,  $r \rightarrow l$  register element order
    - Chosen as default to ease porting of Linux applications
  - Big-on-Little (BoL): LE memory order,  $l \rightarrow r$  register element order
    - Not discussed further today due to time constraints



## Vector programming models: True-LE

- The True-LE programming model asserts that vectors in registers will use r->l element order.
- Many vector intrinsics are “pure SIMD.”
  - Computations operate within “lanes”
  - No changes necessary
- But
  - Implicit and explicit encoding of element numbers are wrong (l → r)
    - Intrinsics must fix these up.
    - No longer a 1-1 correspondence between intrinsics and instructions
  - Vector memory accesses become interesting

## Vector loads (and stores)

- VMX loads are endian-correct (l->r for BE, r->l for LE).
  - But: They enforce alignment on a 16-byte boundary!

```
lvx    va, quadword1
lvx    vb, quadword2
lvsl   vc, real_address
vperm  vd, va, vb, vc
```

- VSX adds lxvd2x (2x64) and lxv4wx (4x32) loads that always load l->r even for LE.
  - BE can just use lxvd2x and lxv4wx for unaligned loads
  - LE must use:

```
lxvd2x  va, real_address
xxswabd vb, va
```

Ouch!



## Vector programming models: True-LE

- The True-LE model is appropriate and useful to most consumers.
  - Apps ported from other LE platforms
    - Straightforward mapping from “old” interface to “new” one using the same element ordering assumptions
  - New apps
    - Programmers coming from an LE environment will find this natural
    - Programmers familiar with the <altivec.h> interface can use it without (much) change.
  - Apps ported from BE Power
    - Most apps can be ported without (much) change.
    - Libraries using lxvd2x and lxvw4x intrinsics can use Big-on-Little model to ease porting.

# Outline

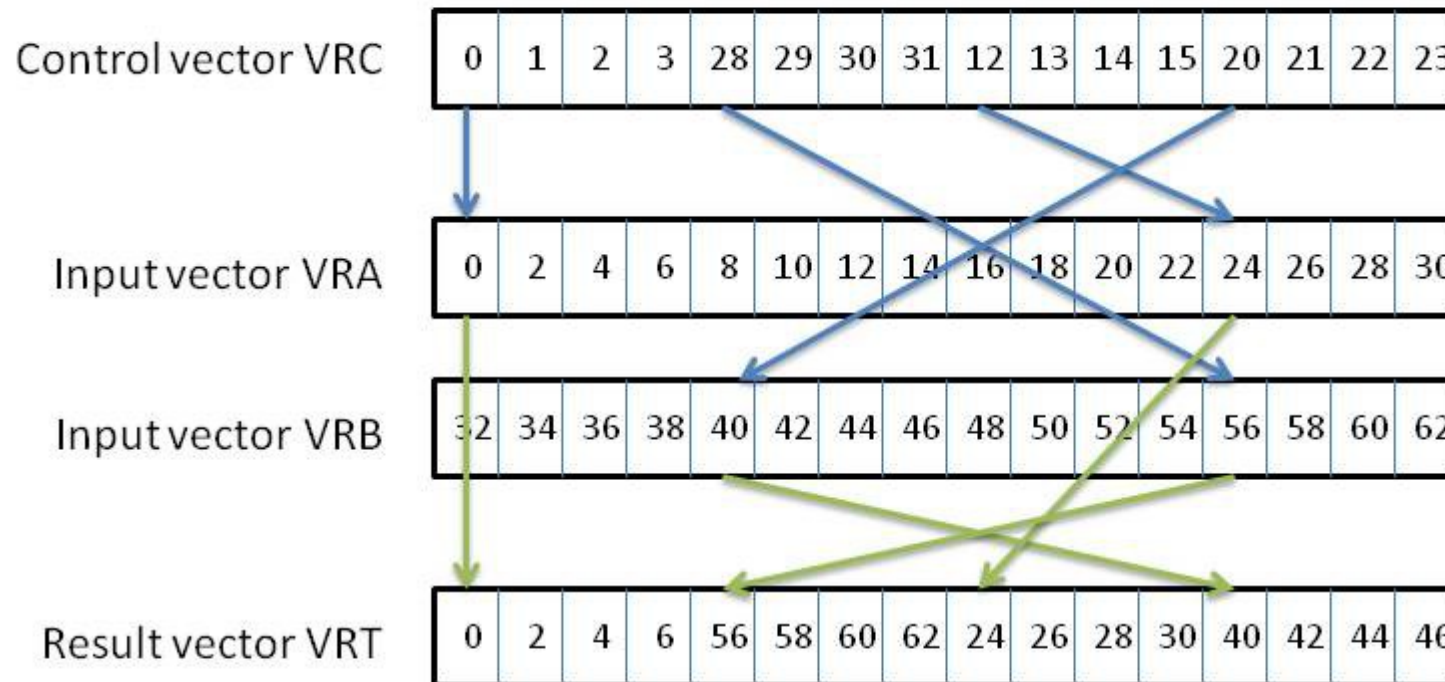
- History
- Little Endian on Power
- Endianness (bytes and elements)
- Programming models
- **Example vector interfaces and implementations**
- Optimization
- Implementation status (gcc, llvm)

## Example: vec\_mergeh



# Example: vec\_perm


## Big Endian



`vperm VRT, VRA, VRB, VRC`

Example: `vec_perm`, cont.

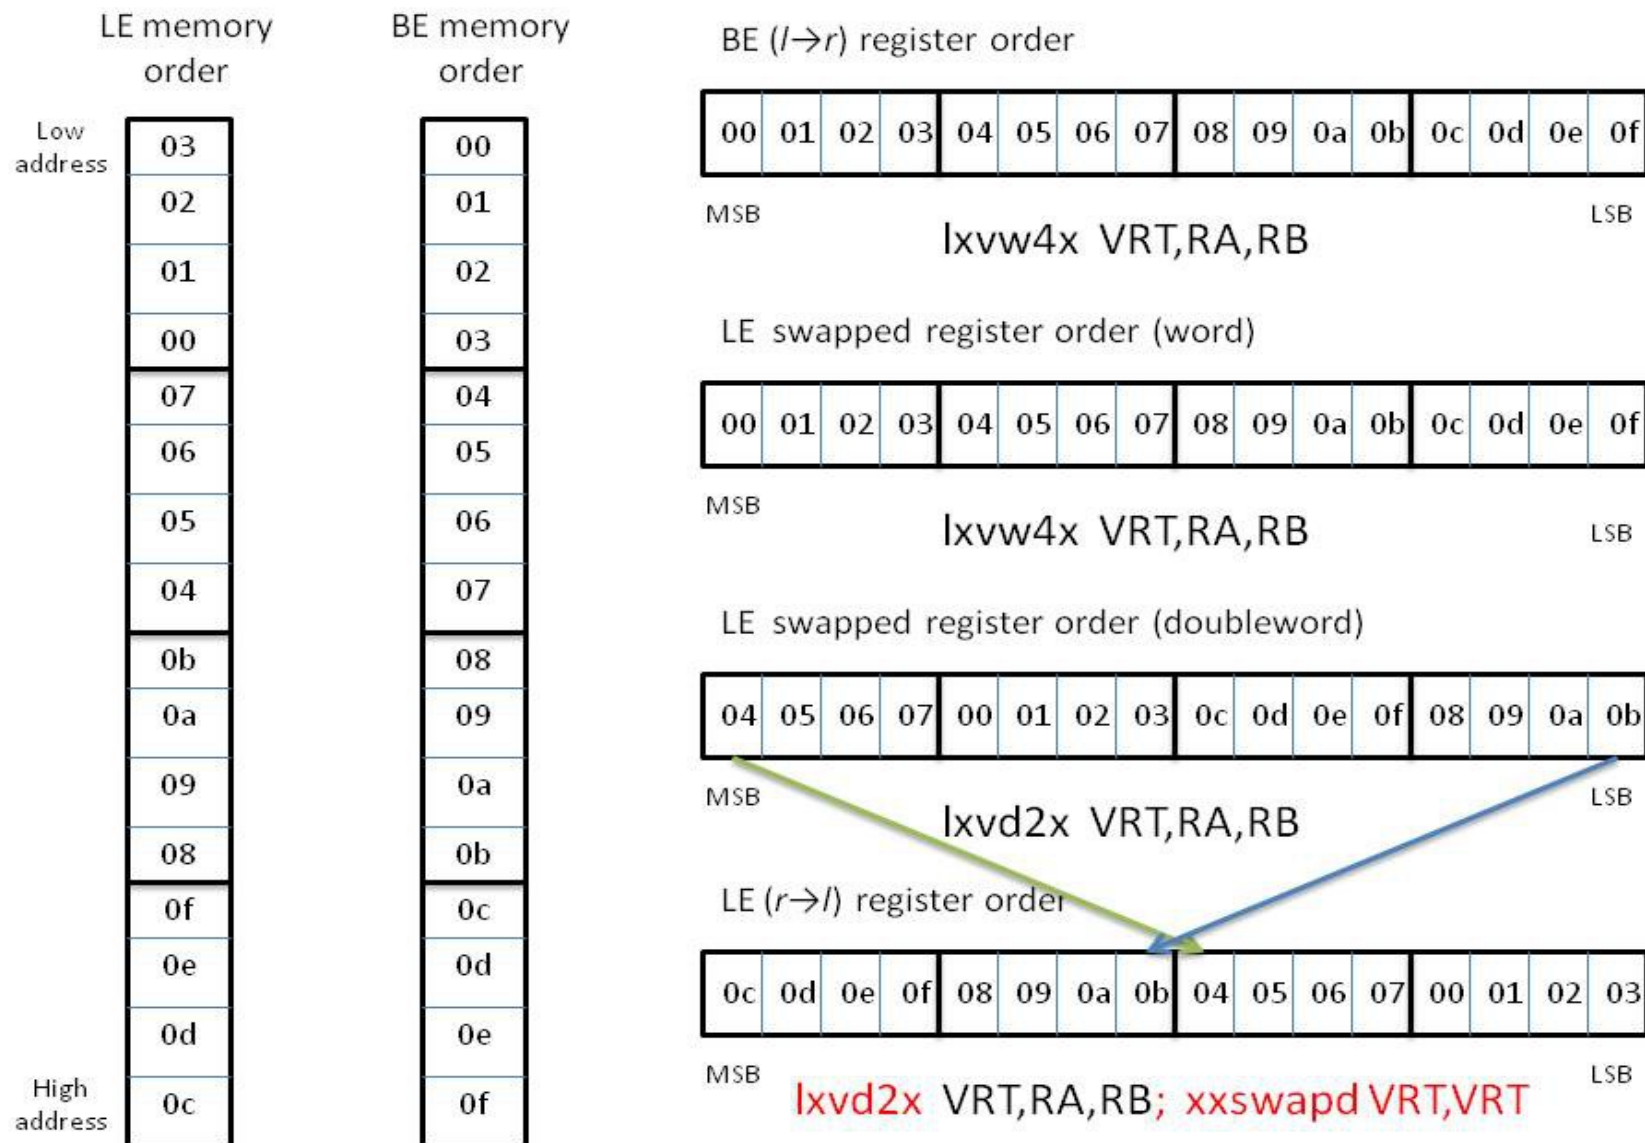
## Little Endian



`vnor` VRC', VRC, VRC

`vperm` VRT, VRB, VRA, VRC'

# Example: `vec_xl` for 4x32 vectors





# Outline

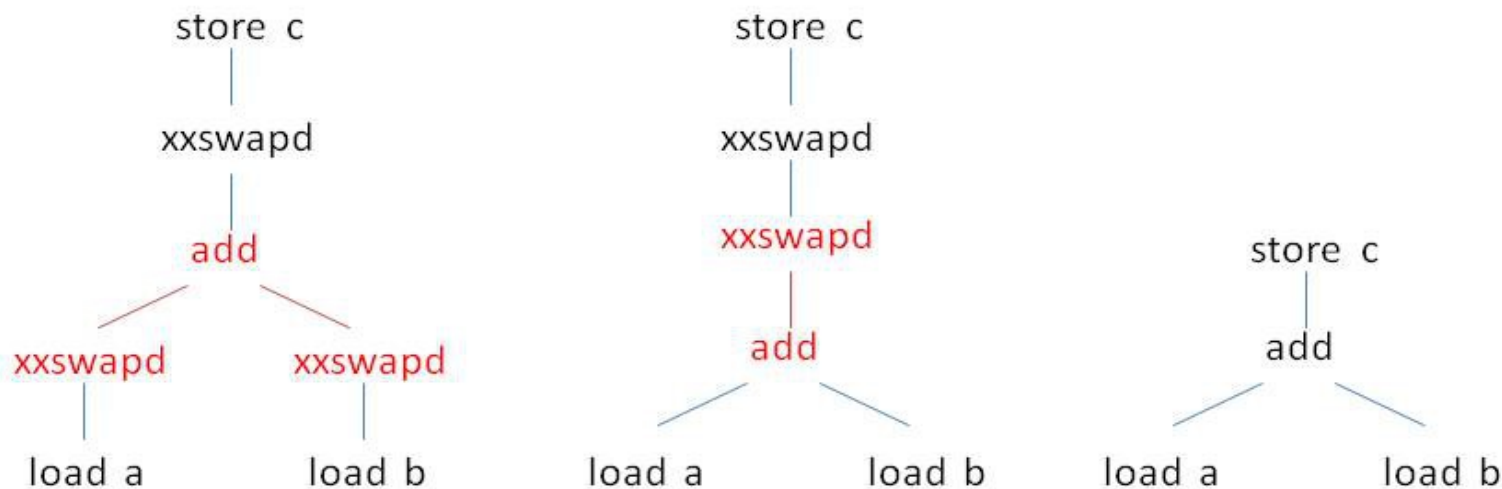
- History
- Little Endian on Power
- Endianness (bytes and elements)
- Programming models
- Example vector interfaces and implementations
- **Optimization**
- Implementation status (gcc, llvm)

## Optimization – swap removal

- VMX load/store instructions may be preferred by compiler when 128-bit alignment guaranteed
  - Operate as expected for little endian
- But poor choice otherwise
  - Longer (but pipelinable) sequence
  - Limited to upper 32 vector registers
- VSX lxvd2x, lxvw4x, stxvd2x, stxvw4x provide good unaligned performance on POWER8, and 64 vector registers.
  - But now we have extra swap costs.
- If we can rid ourselves of the swaps, BE and LE performance are comparable.

## Optimization – swap removal

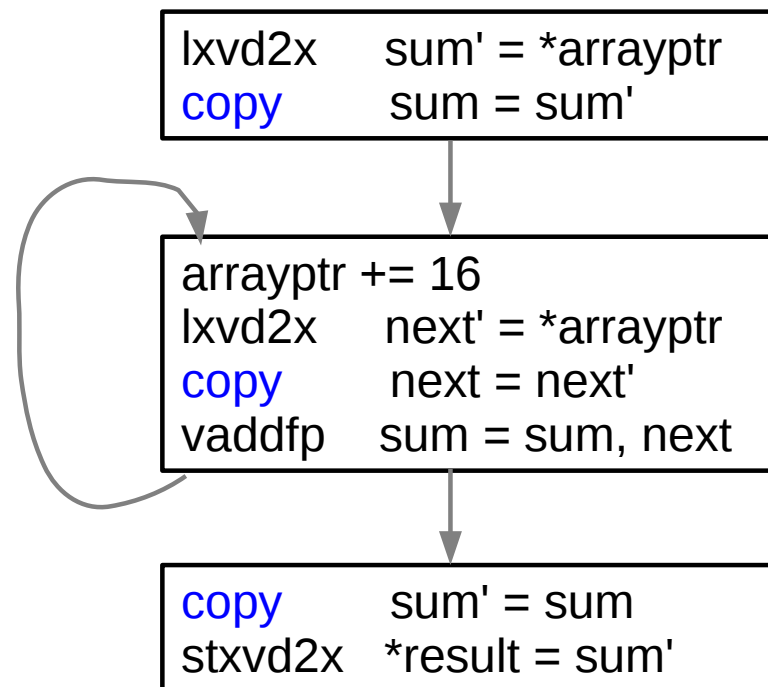
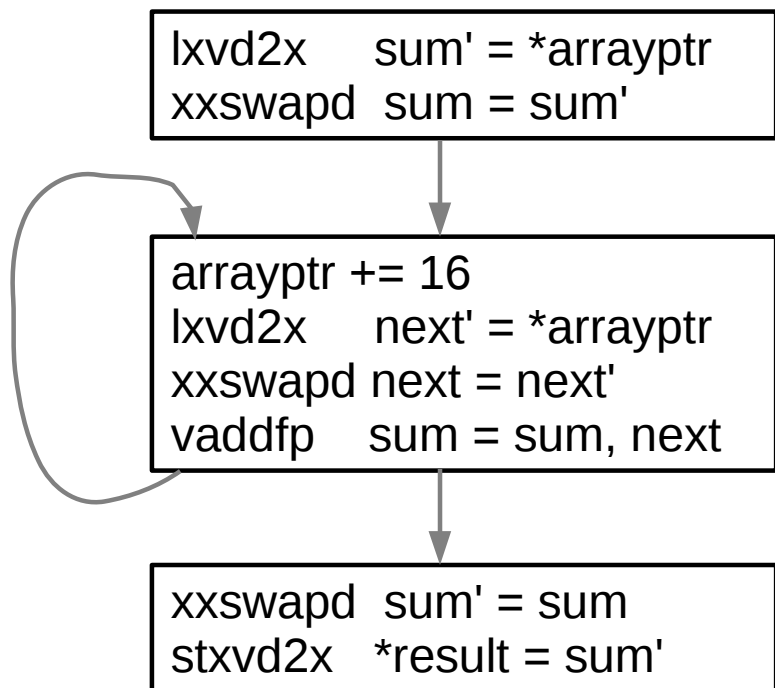
- Observation: Most operations are pure-SIMD and don't care which lanes are used for computation, provided correct memory order is maintained.
  - Thus we look for computations where no swaps are needed.
- “Local” optimization (tree rewrites)



## Optimization – swap removal

- Global optimization
  - Local optimizations work well on extended basic blocks.
  - But loop-carried dependencies, join points interfere
  - Auto-vectorized code is a common case that needs more.
- With du- and ud-chains available, simple to identify maximal vector computations with swaps at the boundaries
  - If all vector computations are lane-insensitive, remove swaps
  - Special handling for ops that are lane-sensitive
    - E.g., splat-from-lane changed to use different lane
  - Cost-modeling if special handling isn't “free”
  - Linear time (union-find)

## Example: sum by lanes



# Outline

- History
- Little Endian on Power
- Endianness (bytes and elements)
- Programming models
- Example vector interfaces and implementations
- Optimization
- **Implementation status (gcc, llvm)**

## Implementation status

- GCC
  - True-LE and BoL models are implemented (4.8+).
  - Global optimization of vector computations is implemented.
    - Most common special-handling cases are covered.
    - No cost modeling
    - Effective in optimizing most code with little effort
  - No local optimization (global is effective)
- LLVM
  - True-LE implemented for VMX (3.5) and VSX (3.7)
  - No plans for BoL model at this time
  - Basic global optimization is implemented, but limited
    - SSA form simplifies
    - LLVM back end is less optimal for vector on BE as well

## Implementation status

- Ported packages/libraries (community)
  - GROMACS
  - Eigen
  - Atlas
  - ...?



# Acknowledgments and Links

- Acknowledgments

- Kit Barton, David Edelsohn, Jinsong Ji, Ke Wen Lin, Joan McComb, Ian McIntosh, Steve Munroe, Hong Bo Peng, Julian Wang, Ulrich Weigand, Rafik Zurob

- Links

- Power ISA 2.07:  
<http://ibm.biz/powerisa>
- Power Architecture 64-Bit ELF V2 ABI Specification:  
<http://ibm.biz/power-abi> (free registration required)
- AltiVec Technology Programming Interfaces Manual (1999):  
[http://www.freescale.com/files/32bit/doc/ref\\_manual/ALTIVECPIM.pdf](http://www.freescale.com/files/32bit/doc/ref_manual/ALTIVECPIM.pdf)

We're hiring!

# Backup

## History: Power ISA Vector Support

- **Altivec (VMX) – Apple/IBM/Motorola alliance**
  - 32 x 128-bit vector registers, aligned memory access only
  - Support for 16xi8, 8xi16, 4xi32, 4xf32, “bool” and “pixel” types
  - Power Mac G5, JS20 blade server (ISA 2.03, PPC970) – **2003**
  - POWER6 (ISA 2.05) - 2007
- **VSX – Vector-scalar extensions**
  - Extended vector register file to 64 x 128-bit, scalar float to 64 x 64-bit
  - Added support for 2xf64 and (limited) 2xi64
  - Unified floating-point and vector register files
  - Unaligned memory access supported but somewhat slow
  - POWER7 (ISA 2.06) - **2010**
- **Expanded VSX and VMX instruction sets**
  - Expanded 2xi64 ops, i128 ops, crypto, logical, decimal, i64/f64 VSX load/store, merge even/odd, ...)
  - Improved unaligned vector access performance
  - POWER8 (ISA 2.07) - **2014**