Universal, Scalable Factorized Power Architecture for New Processors and Memory Loads

Uptime Institute, New York, NY, May 2010

Abstract:
HPC (High Performance Computing) processors and memory loads continue to increase power demands at lower voltages while performance density and total cost of ownership challenges increase.

Efficient HV (high voltage) AC or DC distribution and power conversion to the loads – whether POWER processors, GPUs (Graphic Processing Unit), x86 systems, memory or solid-state storage – are essential to maximize datacenter metrics such as MegaFLOPS / Watt or PUE (Power Usage Effectiveness).

This presentation highlights Factorized Power Architecture solutions for the latest HPC systems, converting from HV to sub-1V loads.

Stephen Oliver, Vicor
soliver@vicr.com, +1 978 749 3256

The Cost of Power Inefficiency

- Going from ‘A’ to ‘B’
  - From AC to the load – the complete chain
- Power is ‘lost’
  - #1: Along the way (distribution losses, connector losses)
  - #2: Stepping down (conversion losses)
- Leading to more heat and more cost

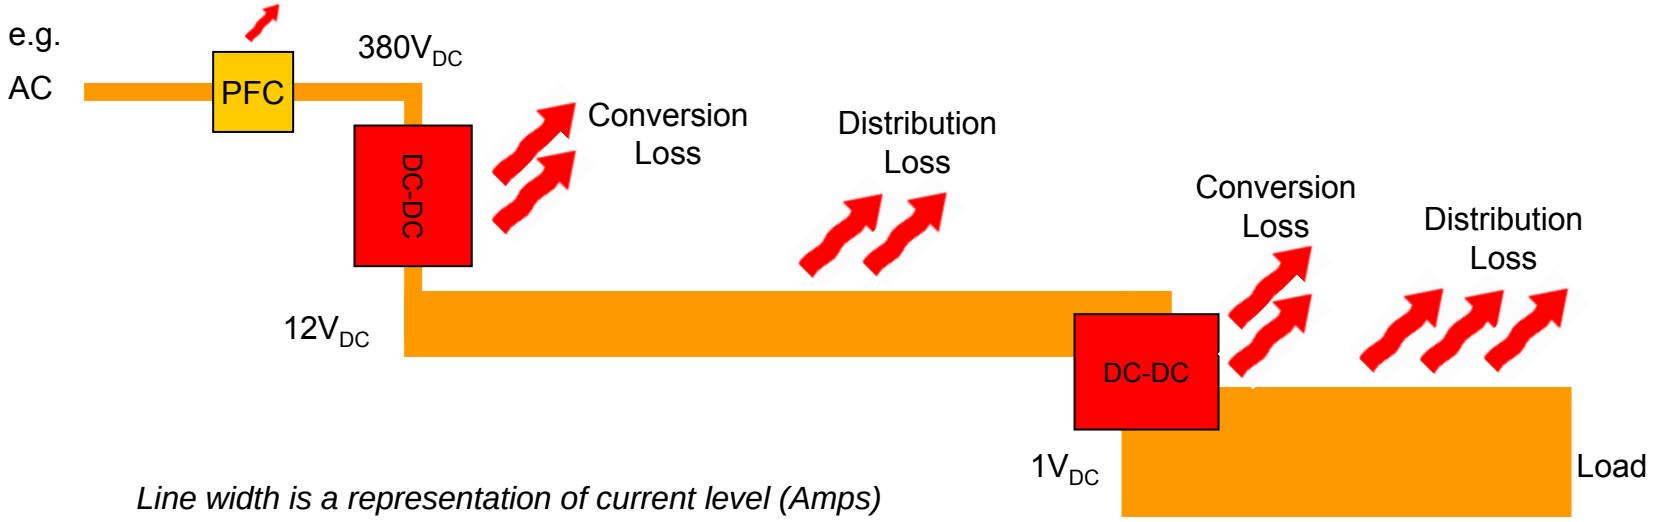

Line width is a representation of current level (Amps)

More Heat, More Cost

- More Heat
  - Larger, heavier heatsinks
  - More air-conditioning
  - Lower reliability
  - Larger, heavier systems, etc.

- Higher Cost
  - Cost of the cooling system
  - Cost of electricity year after year ($/MWhr)
  - Cost of fuel (airborne systems)
  - ‘Total Cost of Ownership’ (TCO)

- Worse Metrics
  - MegaFLOPS / Watt
  - PUE (Power Usage Effectiveness).

High Efficiency Power Distribution

- In large systems, transmission/distribution loss is a problem
  - Power loss is related to current^2 ($P_{LOSS} = I^2R$)
  - So, keep the voltage high (current low), to as close to the load as possible
    - *12V is very high loss (connector loss, pcb distribution loss, etc.)*
    - *48V is 16x better (i.e. 16x lower losses) than 12V*
    - *380V is 1,000x better than 12V*

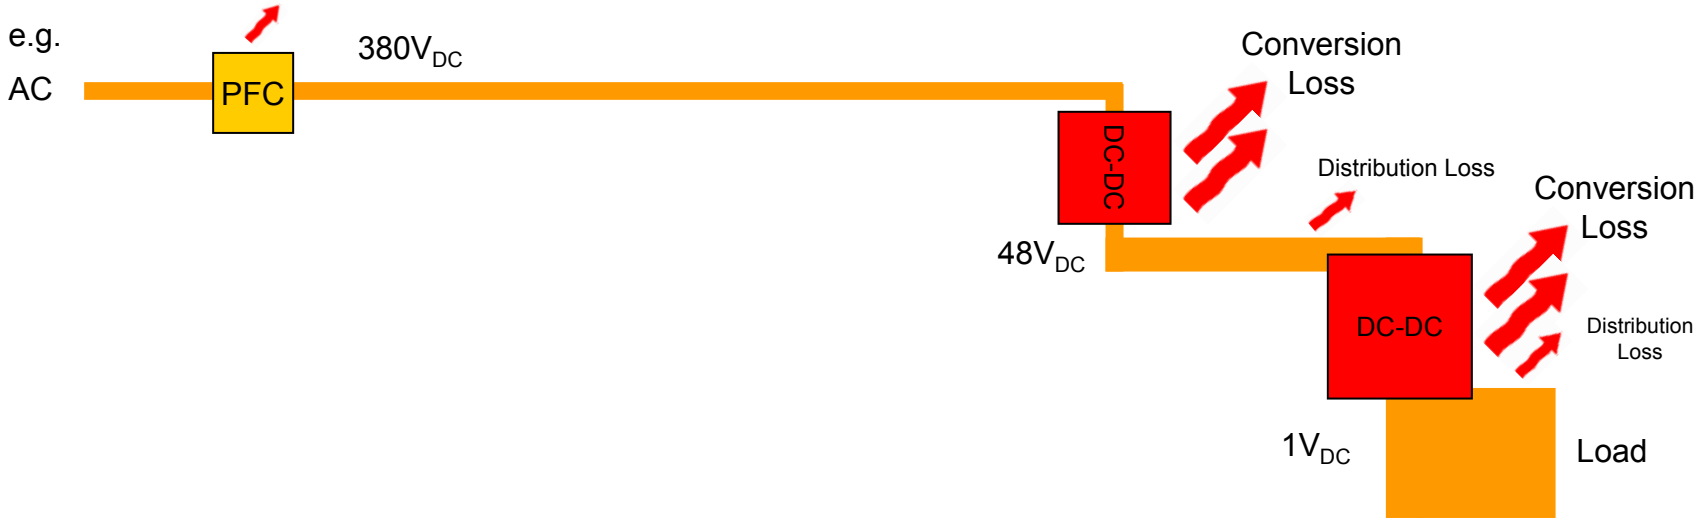

HVDC Distribution Today:

- 350-400V in *datacenters*
  - Telecom (NTT, Hitachi, NEC, etc.)
  - Computing (Hitachi, LBNL, Intel, etc.)
- 350-400V in *rack*
  - Computing (IBM Power6, Power 7, Cray CX1, Dell LC, Verari, etc.)
- 270V in *airframe*
  - Aerospace (Raytheon, Lockheed Martin, Harris, etc.)
- 48V in *rack* or *shelf*
  - Telecom (Huawei, Alcatel, etc.)
  - Computing (IBM Blue Gene/P, etc.)
  - ATE (Advantest, Teradyne, KLA, etc.)
  - Solid-state (LED) Lighting / Display

High Efficiency Power Conversion

- Efficiency is driven by:
  - Topology
    - High efficiency building blocks – smaller, less heatsinking, *even closer to the load*
  - Minimum conversion stages (eliminate duplicate functions)
- Choice of Topology is driven by:
  - System architect's understanding of Total Cost of Ownership
  - Time to design / time to market / availability of conversion building blocks
  - System-wide benefits, e.g.
    - Light-load efficiency
    - Soft-switching / high frequency operation reduces filter size / loss / complexity
    - Low impedance converters enabling bulk capacitance elimination

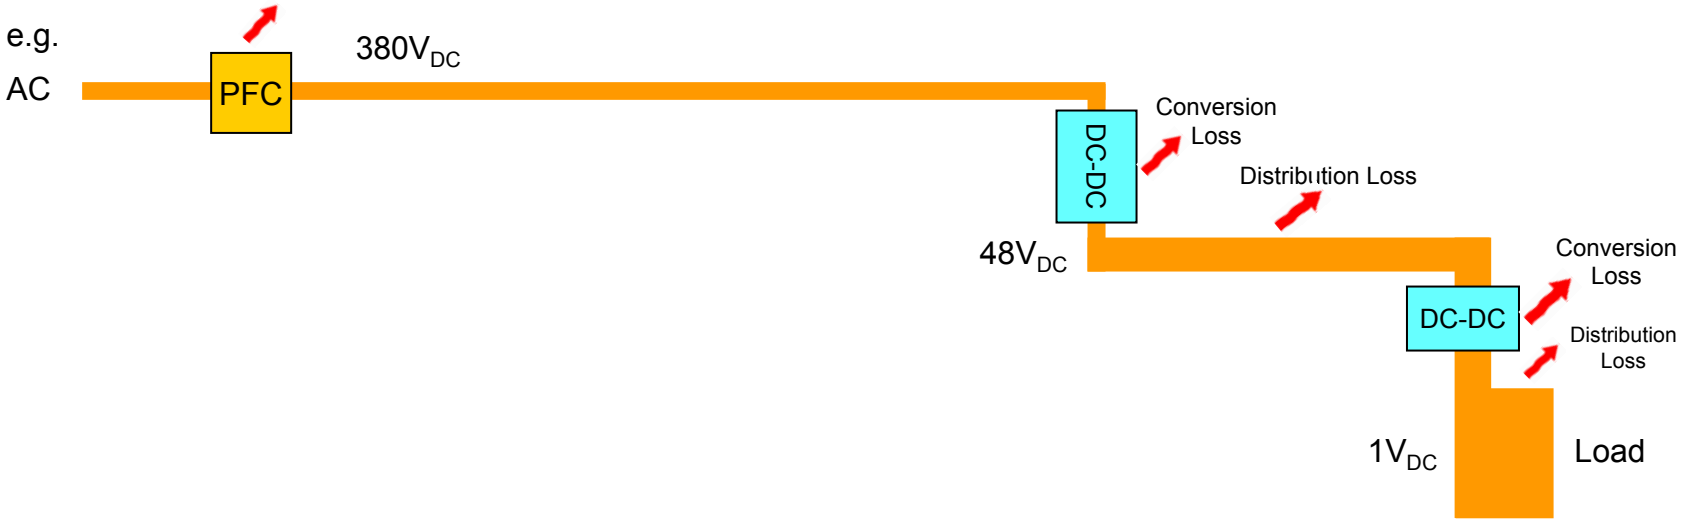

Factorized Power Architecture (FPA) and V•I Chips

- A power supply has 2 functions: ***Regulation*** & ***Voltage Transformation***
- In an old, traditional system (e.g. synch buck, forward, half-bridge)
  - Regulation & Voltage Transformation are ***connected***
  - Difficult to improve performance, consolidation lead to poor *system* performance
- In Factorized Power Architecture (FPA)
  - Regulation & Voltage Transformation are ***separated*** (***'Factorized'***)
  - Easy to improve performance
  - Good system performance

½-chip
Up to 200W or 50A
0.64 x 0.87 x 0.265 in (0.55 in²), 0.25 oz
1.63 x 2.2 x 0.67 cm (7.2 cm²), 8 g

- FPA uses V•I Chips
  - Small, powerful components
    - Highest power density
    - High efficiency
    - Fast, flexible design
    - Fast transient response

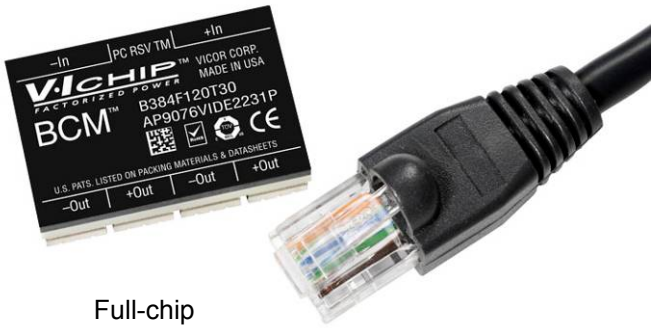

Full-chip
Up to 400W or 130A
1.28 x 0.87 x 0.265 in (1.1 in²), 0.5 oz
3.25 x 2.2 x 0.67 cm (7.2 cm²), 15 g

Why V•I Chip?

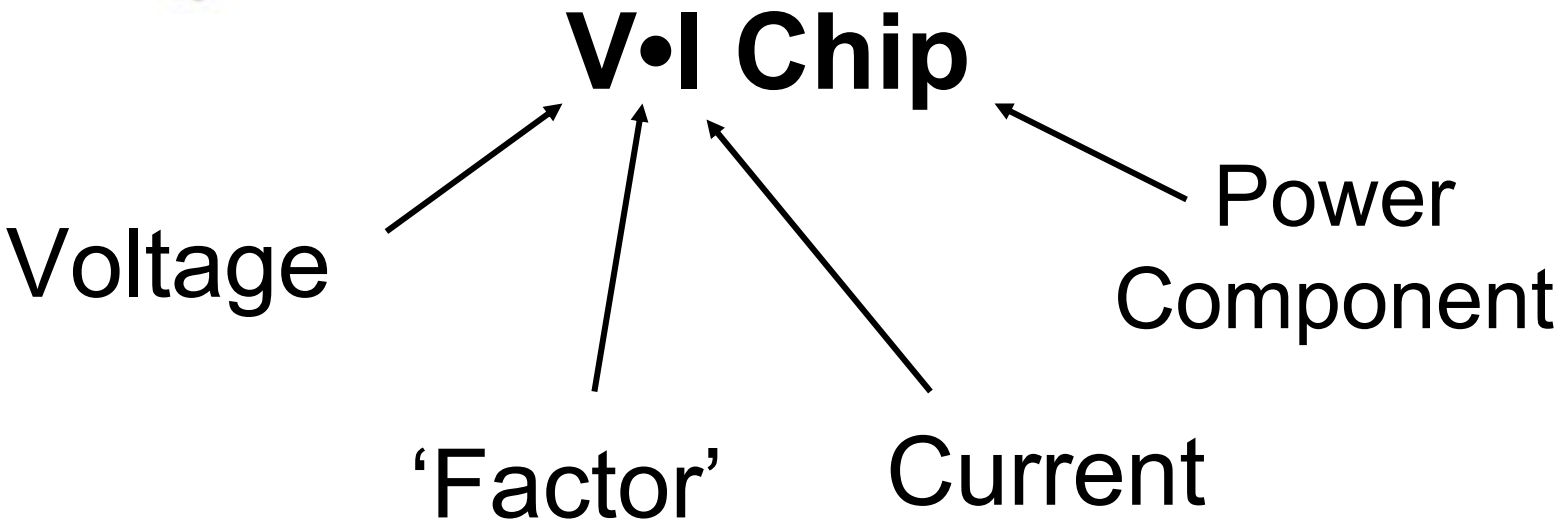

High Efficiency Bus Conversion

- Bus Converter Module (BCM)
  - Isolated, unregulated
  - Voltage transformer / current multiplier

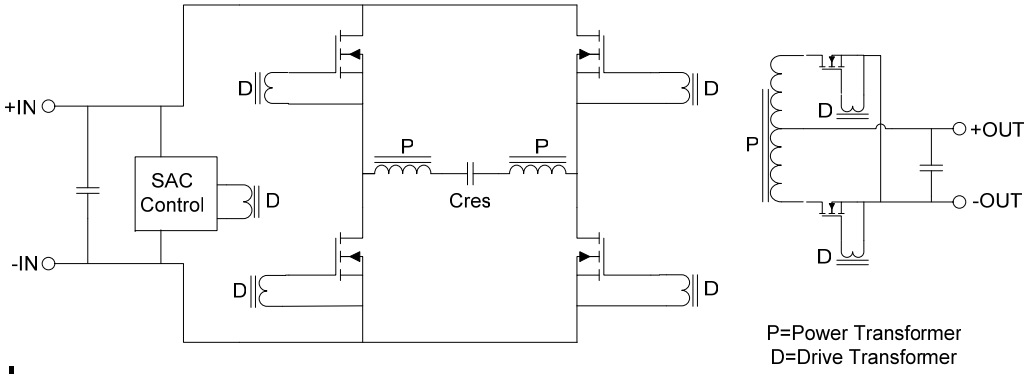

- Sine Amplitude Converter Topology
  - ZVS, ZCS, >1Mhz switching frequency

- 380V to 48V, 330W at 96.5% in only 1.1in^2

System Benefits: Light-Load Efficiency

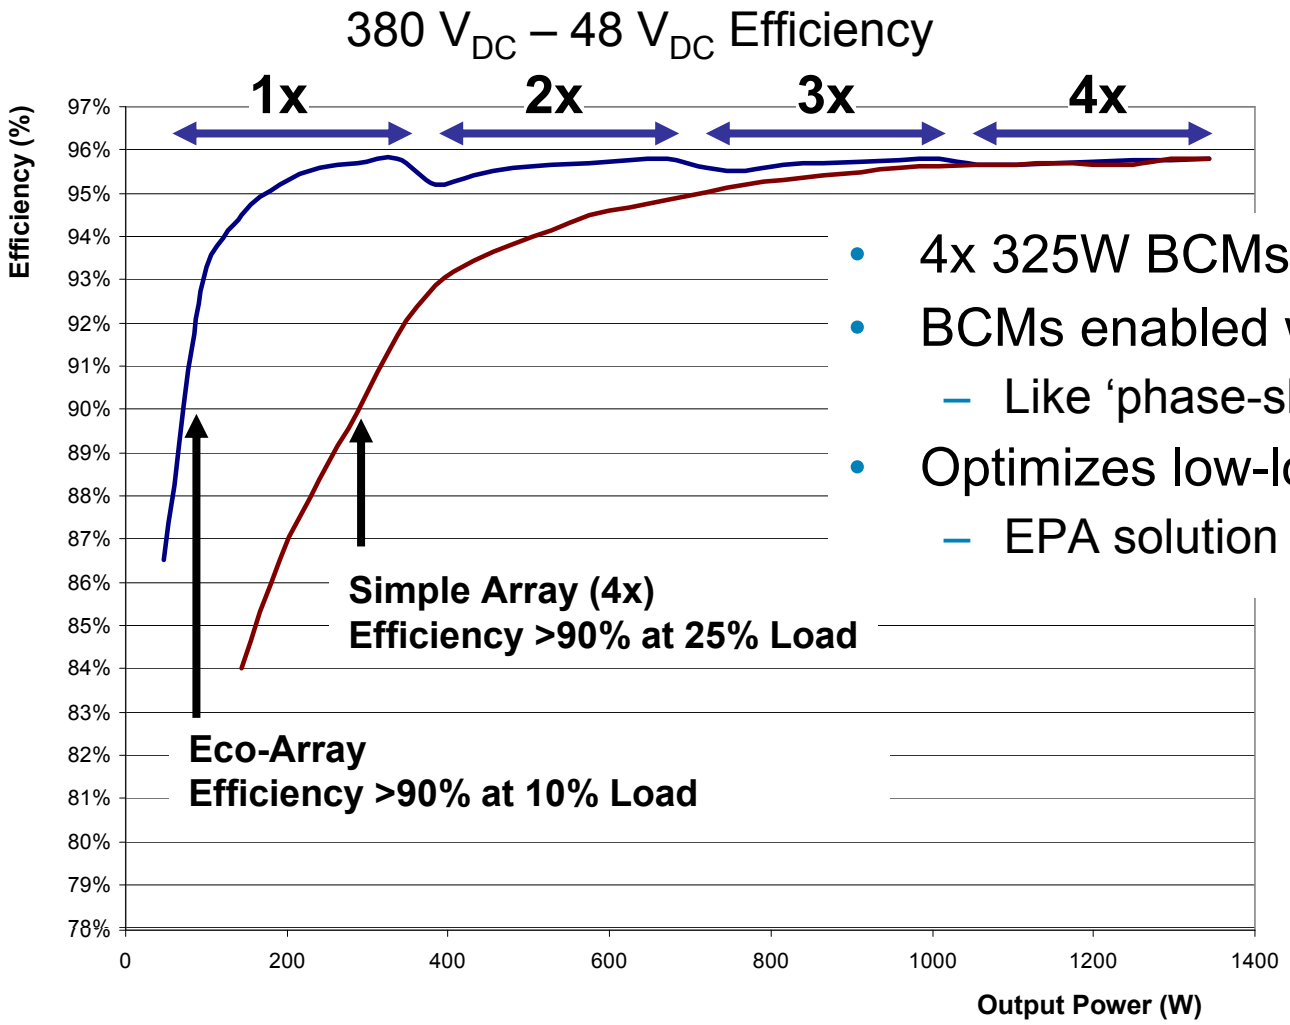

- 4x 325W BCMs in parallel (1,300W)
- BCMs enabled when needed
  - Like ‘phase-shedding’
- Optimizes low-load efficiency
  - EPA solution

VICOR

48V to Regulated, <1V Options

- Traditionally, 48V to 'low voltage load' conversion has been a 2-step approach
  - e.g. 48 - 12V with a 'brick' then 12 - xV with a synchronous buck converter
- As load voltages decrease…
  - Processors now <1V (VR12 spec down to 0.5V), memory <1.2V

… the 'duty-cycle' limitations of the synch buck mean that the topology has reached an 'asymptotic' efficiency level

  - 12 – 3V is a 4:1 'synchronous FET' to 'control FET' ratio
  - 12 – 0.9V is more than 13:1

- Features like phase shedding, multi-phase switching, higher levels of controller integration, etc. achieve smaller and smaller advances

- The resonant 'Sine Amplitude Converter' (as used in the BCM), operating on a 'classic' transformer offers the highest efficiency method to convert from 48V directly to a sub-1V level

Direct 48V-to-Load Conversion

Pre-Regulator Module (PRM)

- Non–isolated, regulator
- ZVS Buck – Boost Topology
  - ZVS, >1MHz switching
- Performance
  - 400W in 1.1 in^2
  - 200W in 0.55 in^2
  - Power Density >1,300 W/in^3
  - Efficiency >97% at full power
- Input : Unregulated 48V (e.g.)
- Output : Regulated into VTM

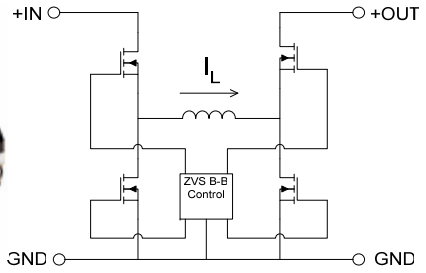

Voltage Transformation Module (VTM)

- Isolated, voltage transformer
- Sine Amplitude Converter Topology
  - ZVS, ZCS, >1Mhz switching frequency
- Performance
  - 130A in 1.2 in^2
  - 50A in 0.55 in^2
  - Current Density >100 A/in^2
  - Efficiency >93% at 1.2V, 115A
- Input : Regulated from PRM
- Outputs : 0.15 - 55V, up to 130A

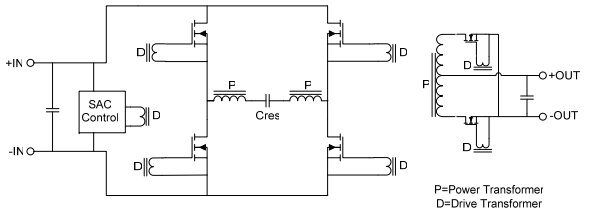

- *PRM + VTM = Isolated, regulated, voltage transformation direct to load*

Direct 48V-to-Intel Processor (VR12)

- ‘Voltage Regulator’ (VR) spec sets voltage for peak processor performance
- FPA is a ‘pure power’ solution
  - Digital or analog loop
  - Any industry standard controller (with access to VID)
  - Re-use existing control / housekeeping solutions

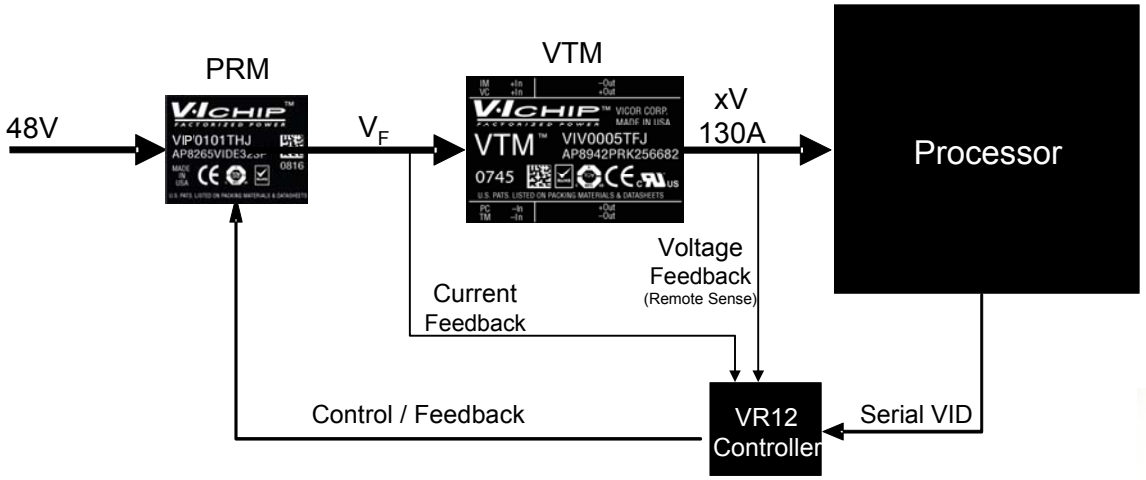

- VR12 Test Board
  - Single 48V feed to board, no bulk caps on V•I Chips
  - ‘Factorized’ core powertrain, low profile VRD solution

The FPA Advantage

- Applying FPA to higher power systems highlights the size, efficiency and value of the V•I Chips

- **Efficiency** ↑ 9%
- **Power Loss** ↓ 45%
- **Size** ↓ 45%

- Save **Money**:
  - –$30 per year, per processor
  - –$600,000 per year, per datacenter

- Save **Energy** (Oil)
  - –2,600 barrels per year, per datacenter

- Save the **Environment**
  - –2,800 tons CO_2 per year, per datacenter

Typical 20,000 processors per data center

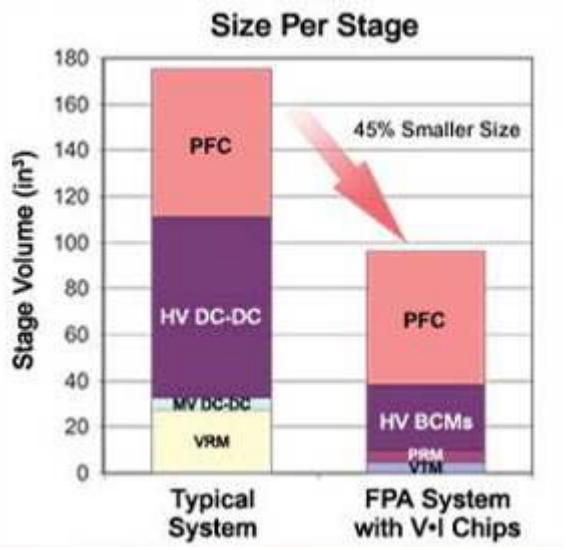

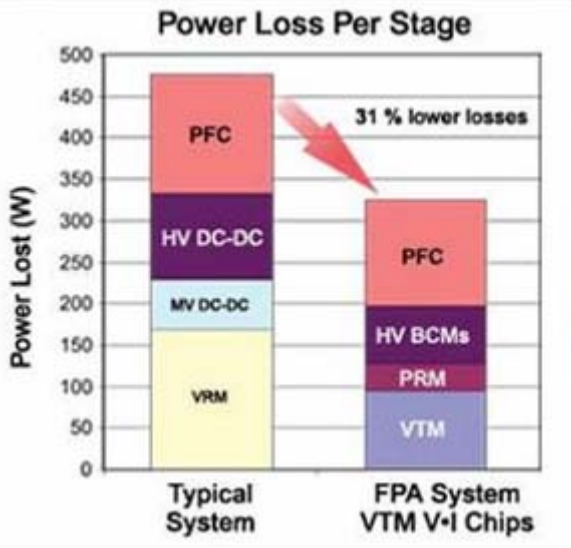

Cray ‘CX1’ / Dell ‘LC’ with V•I Chips

- 380V to 12V with 2x 300W HV BCM bus converters

380V →12V Conversion

Hot-swap, In-rush, etc.

2x 300W HV BCMs

IBM Blue Gene/P with V•I Chips

- Blue Gene/P with Vicor holds 12 of the top 20 in the Green500 (Mflops/Watt)
- Direct 48V-to-load Conversion
  - 1x PRM +2x VTMs per card
    - 48V input, 1.2V output, 200A for processors and 1.8V, 160A for memory

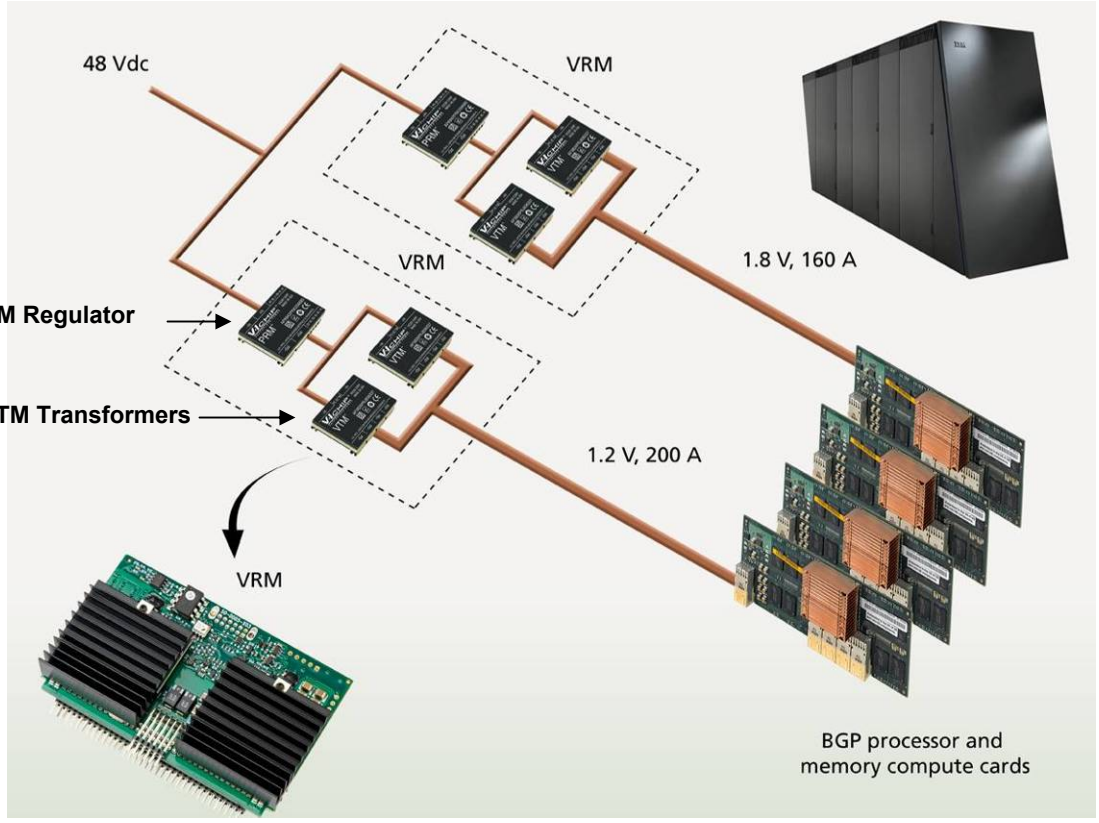

IBM Power 7 with V•I Chips

- PRM regulators ‘factorized’ away from VTM transformers
  - High efficiency distribution, high efficiency power conversion at load

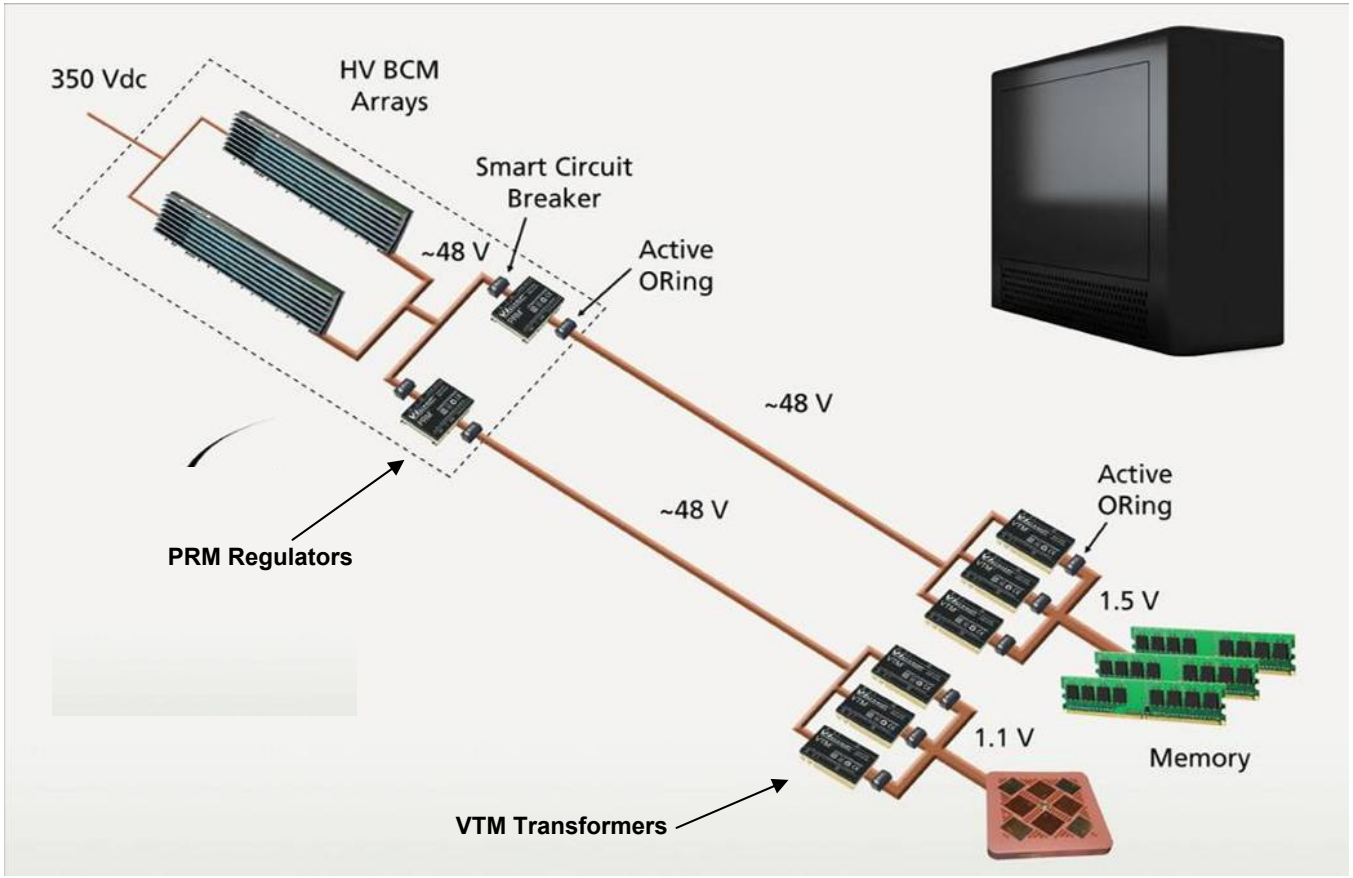

IBM p7-IH (DARPA/PERCS) with V•I Chips

DC converter and Control Assembly
“DCCA”

350VDC to 44V Regulated Output
Using BCM, PRM V•I Chips

20 kW per board

Node (top side)
8x POWER7 MCMs, DRAM, Optical hubs

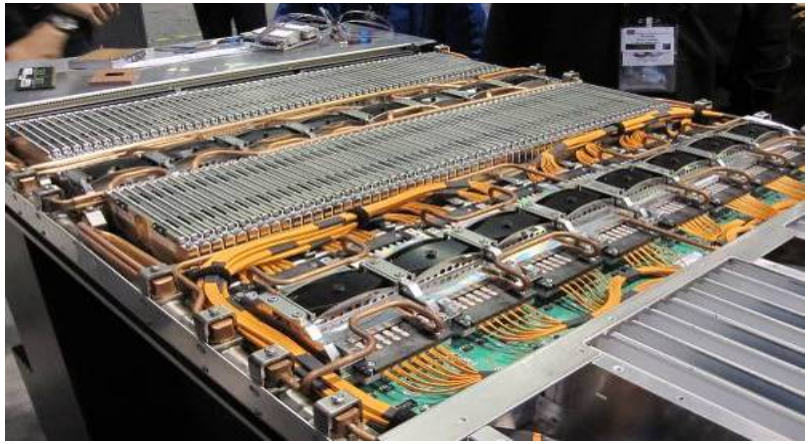

Node (bottom side)
44V to 1.xV using multiple 130A VTM V•I Chips

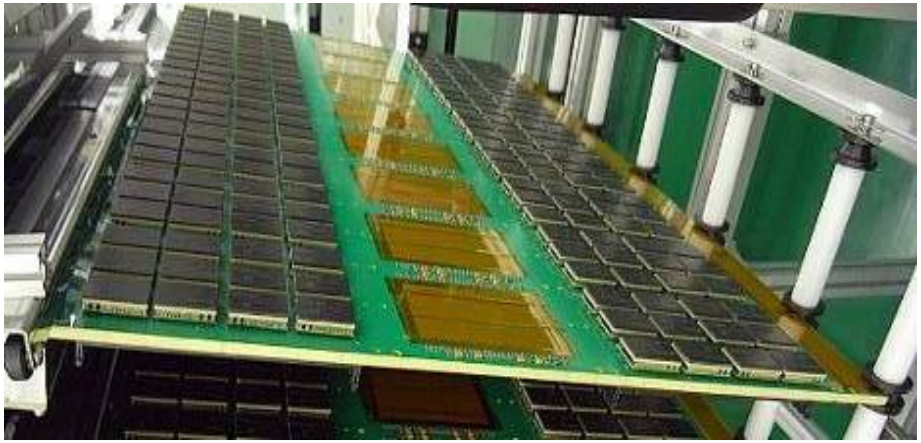

Universal, Scalable Factorized Power Architecture for New Processors and Memory Loads

In Summary:

- The keys:
  - High(er) voltage distribution
  - High efficiency power converters
  - Intelligent exploitation of key features to achieve system benefits

- FPA and V•I Chips
  - Highest system efficiency
  - Smallest power system size

Thank You

- What are your questions / comments?

Stephen Oliver, Vicor
soliver@vicr.com, +1 978 749 3256