

John Krause

1310 E. University Ave, Georgetown, TX 78626

(512) 868-4462

johnckrause@gmail.com

SUMMARY

John Krause's most recent position was that of Telco Platforms Systems Architect for Hewlett-Packard. He has extensive experience in the computer and digital networks architecture and design. He has worked at all levels of design and architecture: system, module, board, and chip. He has over 24 years experience, and has lead teams that successfully shipped over 15 products in that time. John has been the architect and a developer on over 8 major computer chips. He was a member of the architectural team for ServerNet, and is named on many of the patents for ServerNet and Infiniband. He is currently named on 21 patents overall.

Positions held by Mr Krause include but are not limited to:

- Systems Architect for Hewlett-Packard Telco Platforms.
- Lead Architect for Hewlett-Packard Neoview Business Intelligence Platforms.
- Lead Architect for Hewlett-Packard NonStop ServerNet fabric architecture.
- Lead Architect for Hewlett-Packard NonStop Telco platforms.
- Architect for Tandem S-Series fault tolerant, scalable system.
- Fabric Architect for Tandem Modular Architecture
- ServerNet architect.
- Chip level architect and development lead for HP NonStop S-Series CPU chipsets.
- Architect and development lead for Tandem ServerNet switching chips.
- Architect and development lead for Tandem Quad Modular Reduncy Equalizer.
- Project lead for Carroll Touch standard and custom infrared scanning products.
- Architect and project lead for Carroll Touch Smart-Y.
- Project lead for an instruction compression tool, Accelerated Solutions.

PROFESSIONAL HISTORY

- **Hewlett-Packard** **2003-2009**

**Business Critical Systems (BCS):
Telco Platforms Systems Architect**

Maintained ownership for the BCS Telco platform definition and roadmaps. Setting of requirements for future product releases. Definition of Telco-specific products, and selection of products being adopted from other divisions. Technical response to customer queries and requests for quotes, and technical support for customer issues. Definition of the HP Telco alarm panel and architectural definition of the HP Telco clock distribution system. Setting of requirements for the Telco manageability system. Investigational work for scoping future technologies and longer term directions.

Owner for technical scoping of HP healthcare and military markets for potential investment opportunities.

NeoView business intelligence platform architect and owner for the platform roadmap. Architect for system fabric, racking strategies, and system cabling. Definition of NeoView expansion architecture. Definition of patch panels and cabling system.

**NonStop Division (NSD):
Senior Member Technical Staff**

NonStop Telco platform architecture and roadmap, architecture, definition, and specification of the NonStop NS1600CG. Architecture and definition of the NonStop NS16200CG Telco storage subsystem, Telco maintenance LAN, alarm panel.

NeoView business intelligence platform architect. NeoView fabric topology definition and implementation. Neoview platform architecture and roadmap. Definition and topology of Ethernet based front end LAN. Introduction of new methods and chip specifications to reduce ServerNet network congestion. Creation of fault tolerant manageability LAN, racking strategy definition. Definition and specification of architectural changes made for I/O performance increases.

Low end product architecture and definition for NonStop “P-Switchless” system. Specification and topology definition for a limited scalability low end system.

- **Compaq Computers (merged with HP in 2003)** **1999-2003**
Senior Member Technical Staff

Architected, specified, and lead implementation for NeoView fabric topology. This topology significantly increased performance over heritage NonStop systems, increasing overall NeoView platform performance.

Lead architectural and implementation efforts for the HP NonStop Telco System (NS16000 CG). Efforts included the definition of the system architecture, setting overall system requirements, definition of the I/O architecture and network mapping of I/O devices, architecture and definition of a

new telco power system, modification of a seismic rack, setting requirements for the Telco maintenance LAN, oversight of the telco system through NEBs testing.

Architected and specified the Voter Module for the Compaq NonStop NS16000. Work included definition of the general mechanical approach, power system, cooling system, serviceability strategy and features, panel interfaces, and means of handling high speed signaling.

Architected and lead implementation teams for a three-stage topology upgrade roadmap; first step was an immediate scalability increase, the second stage a significant scalability increase and performance increase based on the same chipsets, and a third long term performance increase available with the next generation of ServerNet chipsets. The first stage used existing switching modules. The second stage involved the architecture and lead of the implementation team for new switching modules, architected to be upgradeable with the next generation of chipsets. Set requirements for the following generation of chipsets. Set general mechanical approach, power system requirements, cooling approach and requirements, service strategy, panel interfaces, and manageability requirements for the switching modules.

Served on the NonStop Modular Architecture committee, authored the System Area Network (fabric) chapter of the Modular Architecture specifications. NonStop fabric architect.

Support of Tiger Team created and assigned by Division Vice President to find and fix 4 system problems creating significant numbers of customer issues. Of the division wide effort, John personally found 2 of the 4 design flaws and oversaw design changes to fix them.

Architected and proposed System Area Network topology changes for improved NonStop node performance.

- **Qualis Design** **1998-1999**
Principal Engineer

Developed and high level simulation model for ATM switch chipset. Model served as means architectural verification and as a simulation model to support test suite development.

Recoded communications interface ASIC for a customer with gate-level compile issues, large numbers of instantiated gates that needed to be moved to a new library. Recoding was done with custom tools in perl used to improve coding style and reduce the number of instantiated gates.

Taught Qualis Verilog class.

- **Compaq Computers** **1997-1998**
Senior Member Technical Staff

Micro-architecture and implementation lead for "Curator" FPGA in the NonStop S88000 processor, a chip designed to keep two MIPS CPUs in lockstep despite MIPS design flaws causing the CPUs to leave lockstep. Hardware specification, specification of the software

interfaces, setting verification requirements and priorities, support of laboratory efforts, and physical design requirements. Design blocks included a block read engine that synchronized reads between processors with minimum latency, multiple internal block internace protocols.

Micro-architecture of a series of two PCI-X based ASICs serving as ServerNet Host Bus Adaptors (HBAs). Defined external interfaces, internal bus structures and protocols, supermodule interfaces, onchip and side memory requirements, deadlock avoidance and recovery algorithms, clocking requirements. Authored architectural specifications.

Micro-architecture and lead developer of the “CCC” ASIC, a MIPS cluster bus agent, providing a performance improvement over the previous agent. Sole designer.

- **Tandem Computers (Merged with Compaq) 1990-1997
Project Engineer**

Micro-architecture and lead developer for a MIPS Cluster Coherency Agent (“CCA”), and chip allowing a 4-way SMP or 2-way fully self-checked SMP. The CCA provided the primary memory path for the system, and the ServeNet interface for clustering and I/O. Authored chip specifications, architected internal bus structures, set performance requirements, supermodule design and testing approach. Design blocks included the main processor interface, processor coherency logic, I/O coherency controls.

Developed a fault tolerant clocking FPGA for QMR self-checked processors. Provided clock source for all four processors in the QMR, and was capable of surviving the failure of one of the (identical) FPGAs or clock sources. Sole designer.

Author, ASIC coding standards for Tandem/Compaq ASIC development group. Coding standards allowed a large group of developers to work differing parts of a chip and

Architected and developed board and FPGA for a rate matching device to allow ServerNet connection between standard speed devices and slower emulation devices, allowing lab development with mixed systems. Sole designer.

Micro-architecture and lead developer for the first ServerNet switching ASIC, a six port network switch. Authored chip architecture and specification document, set internal bus structures, number of ports. Set novel verification approach for randomized high stress environements within Tandem, and was lead verification engineer. Wrote compiler scripts (Synopsys compiler). Chip was used as a test vehicle for physical layout, backend tools, and foundry interfaces. Design blocks included the output path, crossbar biased arbitration (refer to US patent 5694121), duplex control state machines, self-checking coalescing logic.

Initiated verification approach for randomized high stress simulation environments, simultaneous hardware/software development. Authored Verilog coding standards. Process flow and tools evaluation for Tandem/Compaq ASIC group.

- **ROLM/IBM** **1987-1990**
Project Engineer

Project leadership for new phone architecture that allowed pluggable phone options, creating an option market for ROLM telephones.

Owner of ROLMLink telephony interface hardware, a core digital telephony technology for ROLM/IBM.

Lead architect and designer in a 2500 gate, high volume interface ASIC between ROLMLink, a microprocessor, and other telephony functions.

- **Carroll Touch** **1985-1987**
Project Engineer

Project leadership for scanning infrared touch systems. Definition of standard product designs, direct work with customers for custom designs.

Shipped multiple scanning infrared touch systems, providing hardware and firmware designs, oversight of the mechanical design, and project scheduling. Architected, designed, and wrote embedded code for intelligent RS-232 “Y”. Innovated design approaches to adopt weaker optics, significantly reducing product cost.

- **Accelerated Solutions** **1984**
Project Engineer

Project leadership machine instruction code optimizer. Accelerated Solutions was a startup company producing a simulation accelerator. The processors were capable of executing 4 instructions in parallel. The code optimizer was capable of re-ordering machine instructions from a single threaded form into a parallel execution form.

PATENTS

- 5547849 *Synchronization of data between two processing elements*
Provides a means of checking two processing units (CPUs) for lockstep synchronization by checking the I/O streams of those units. Protocol errors on one side automatically leads to the use of the processing unit on the other side until a repair is made.
- 5631955 *Option bus for use in telephony systems*
Provides a means of connection option modules to a digital telephone system, allowing those modules access to voice, data, and control streams.
- 5694121 *Latency reduction in routing arbitration*
Provides a means of arbitration within a networking system that allows system software to configure bandwidth ratios among different ports of the network switches.

- 5751932 *Fail fast, fail functional, fault tolerant multiprocessor system*
A duplicated, self-checked and fault tolerance system architecture. This system is scalable via a System Area Network (SAN) in both processing units and I/O devices. This patent received a Compaq Computers Key Patent Award in 2002.
- 5751955 *Method to synchronize a pair of CPUs*
Provides a method of synchronizing two processor (reference 5547849) so that they may be placed in lockstep and run synchronously. Data transfer for the synchronization of memory is done over a System Area Network.
- 5790776 *Apparatus for detecting divergence between a pair of duplexed, synchronized processor elements*
Provides a method of detecting differences between I/O streams of two synchronized processors, and provides related error recovery mechanisms to allow the system to survive on a remaining good processor.
- 5867501 *Encoding for communicating data and commands*
Provides a link layer data encoding scheme used by ServerNet I links. The encoding scheme is a 9 bit scheme allowing symbols to represent either 8 bit data or control/command characters. Both data and control/command symbols have error checking.
- 5914953 *Network message routing for deadlock prevention*
Describes key elements of a network layer definition, including a means of routing data packets with reliable delivery among processing and I/O nodes on a System Area Network in a deadlock free manner.
- 6151689 *Detecting and isolating errors in data communications*
Patent describes a set of packet error checking mechanisms used to isolate a faulty component within a large system of processing and I/O components connected by a System Area Network.
- 6157967 *Method of data communications flow control using BUSY/READY*
Patent describes a link layer flow control mechanism that allows packet buffering on a symbol basis, further allowing wormhole routing in the network. The flow control mechanism is independent of packet size, and is easily configured to work with different buffer sizes and link length.
- 6175882 *Network system for a first module port auto configuring same mode as a second module port*

Provides a means of detecting the rate of an incoming data stream that may be sent at one of several rates, and automatically adjusting to that rate.

- 6249756 Hybrid flow control

Provides a flow control mechanism compatible with the one described in 6157967, yet makes more optimal use of the available buffer space needed to accommodate cable lengths and other contributors to data latency.

- 6690757 High speed interconnect adapter with lane de-skew

Provides a means of initializing and synchronizing ports with groups of serial lanes, each with independent received clocks, bit timing, and multiple transmit rates such that the separate lanes may act together as a single port.

- 6728909 Data communication with speculative reception of data in a data processing system

Describes a means by which a network end node receiving a packet may initiate processing and permissions checking on that packet before the packet has been fully received and before it is fully checked for being complete and correct.

- 6765922 Speculative transmit for system area network latency reduction

Describes a means by which a network host adaptor may start transmitting a packet with data before all of the data for the packet is received from memory. Means of marking the packet as being known as erroneous are provided for those cases in which the remainder of data from memory is not received in a timely fashion.

- 6865231 High speed interconnect adapter having automated crossed differential pair correction

Provides a means by which an inversion in high speed board level signals may be automatically compensated within the receiving logic of a network node. This allows flexibility in the routing of high speed signals, allowing them to meet high speed routing constraints more easily.

- 6870814 Link extenders with error propagation and reporting

Provides a means by which an inversion in high speed board level signals may be automatically compensated within the receiving logic of a network node. This allows flexibility in the routing of high speed signals, allowing them to meet high speed routing constraints more easily.

- 6882656 Speculative transmit for system area network latency reduction

Patent adds further claims to 6765922, which describes a means by which a network host adaptor may start transmitting a packet with data before all of the data for the packet is received from memory. Once a packet is terminated erroneously, existing data may be retransmitted. If the remainder of data from memory never arrives, there must be an abandonment of the first part of the data.

- 6961347 High-speed interconnection link having automated lane reordering

Provides a means by which port capable of receiving multiple serial streams of data (“lanes”) simultaneously will automatically order the lanes appropriately during a configuration period based on known data values on the lanes during configuration time. This allows flexibility in the cabling and in the high speed signal routing, reducing the need to cross wiring, and making it easier to meet high speed signal constraints.

- 7016971 Congestion management in a distributed computer system... (injection rate control)

Provides a means by which network routing nodes may give feedback to network end nodes during periods of high congestion. The endnodes may respond by reducing the rate at which they transmit packets into the network.

- 7197100 High speed interconnection adapter having automated lane de-skew

Provides a means by which port capable of receiving multiple serial streams of data (“lanes”) simultaneously will automatically accommodate for latency differences in the transmission paths of the lanes. This allows flexibility in the cabling, eliminating any need to carefully match cable lengths among the lanes, and allowing design flexibility in other parts of the transmission path.

PUBLICATIONS

- Internal (proprietary): 19 white papers, 14 product specifications, 6 product proposals
- External: Contributor to Qualis Re-use Guide (regarding intellectual property reuse)
- Courses: Contributor to the Qualis design re-use course.

EDUCATION

- BSEE with Honors, The University of Texas at Austin
- MSEE, The University of Texas at Austin

ORGANIZATIONS

- IEEE, Senior Member.