

Breaking The I/O Bottleneck

Karl C. Hansen, CEO
Brilliant Points, Inc.
Karl.Hansen@BrilliantPoints.com

Executive Summary

A system is presented which enables breaking the existing I/O bottleneck in digital circuitry. Designers can choose between similar bandwidth with significantly reduced I/O connections, wiring, and board layout issues, or significant bandwidth increase across the roughly the same I/O and wiring. The system achieves the same noise immunity as differential signaling while encoding significantly more bits per state-change than current methods, using combinations of techniques from wireless communications, three-phase power transmission, and QAM encoding. The system can provide significant savings in board or circuit real estate, I/O power consumption, or both while maintaining or even increasing bandwidth of a design.

Decades of increases in the transistor density have far outstripped our ability to get information on and off of chips. The need to have a secure and reliable connection which remained so through a wide variety of operating conditions, while retaining the ability to assemble chips into larger modules, placed hard minimums on the feature sizes of the I/O connections. Unfortunately the number of transistors we are able to pack onto a chip continues to increase at or near the Moore's Law curve, while we are approaching physical limits to both the count and size of reliable off-chip connections.

Many approaches have helped counter this problem.

Multiplexed data/address busses served for a period of time, but with processors commonly using 64-bit data and address busses, and some using 128, 256 or more bits, even multiplexing is unable to keep up with the growth in processing power. Furthermore, the increase in processor speed has made even the short time required to multiplex data and addresses onto the same pins a "long" operation compared to the processing ability of the chip.

The original ubiquitous DIP package was replaced with miniDIP, and in turn surface-mount technologies like SOP, TSOP, QFP, Pin-Grid Array (PGA), Ball-Grid Array (BGA) and others, all in the attempt to further miniaturize I/O connections and to further increase the information density across the on/off chip interfaces.

As clocking speeds increased, coupling between signals increased, leading to adoption of differential-pair interconnections for transferring high-speed data, but this was kind of a two-steps-forward-one-step-backward approach, because now two I/O connections were required for every I/O path.

XAUI and other "grouped differential busses" were created to enable very high speed clocking of the I/O path, but even those are reaching limits as the data transfer speeds approach 40 gigabits-per-second (GBPS) to 100 GBPS. A separate clock (typically 1.25GHz or a multiple) synchronizes transmitters and receivers at each end of the XAUI bus. 8b/10b or 64b/66b encoding is used to aid in synchronization.

The ever-present wired Ethernet connections use combinations of pulse-amplitude modulation and trellis-coded-modulation to accomplish encoding of 2 bits per transition in 1000BASE-T, but Ethernet does so at the cost of reduced state-to-state noise immunity from rail-to-rail differential-pair signaling.

As more and more systems begin to incorporate on-the-fly audio and video encoding, it is likely that even combinations of all of the best current approaches will be taxed to keep pace with the ever-increasing bandwidth demands of people, the internet, and the processors which deliver the data.

In the RF world there has existed a fairly powerful method of encoding data called Quadrature-Amplitude Modulation. In this system, a sine wave has both its phase and amplitude changed simultaneously to encode information. A QAM diagram has phase and amplitude (or Q and I) axis, with marked "allowable" locations for phase/amplitude combinations which form what is called a QAM "Constellation". An integer number is always associated with a QAM system which indicates the number of stations in the constellation. We will generically refer to an arbitrary QAM system with "N" stations as "nQAM".

For example, a typical nQAM with 16 stations in its constellation (16QAM) would appear something like the following diagram, Figure 1. The vertical axis represents the "Q" (Quadrature or Phase) dimension

while the horizontal axis represents the “I” (Inphase or Amplitude) dimension. Each station within the constellation is assigned a value which represents the bits transmitted when that station is visited. In the example below, each station represents four bits from 0000 to 1111, as labeled. The station rows and columns are typically numbered in a GRAY-code so that only single-bit-changes occur between adjacent row or column stations.

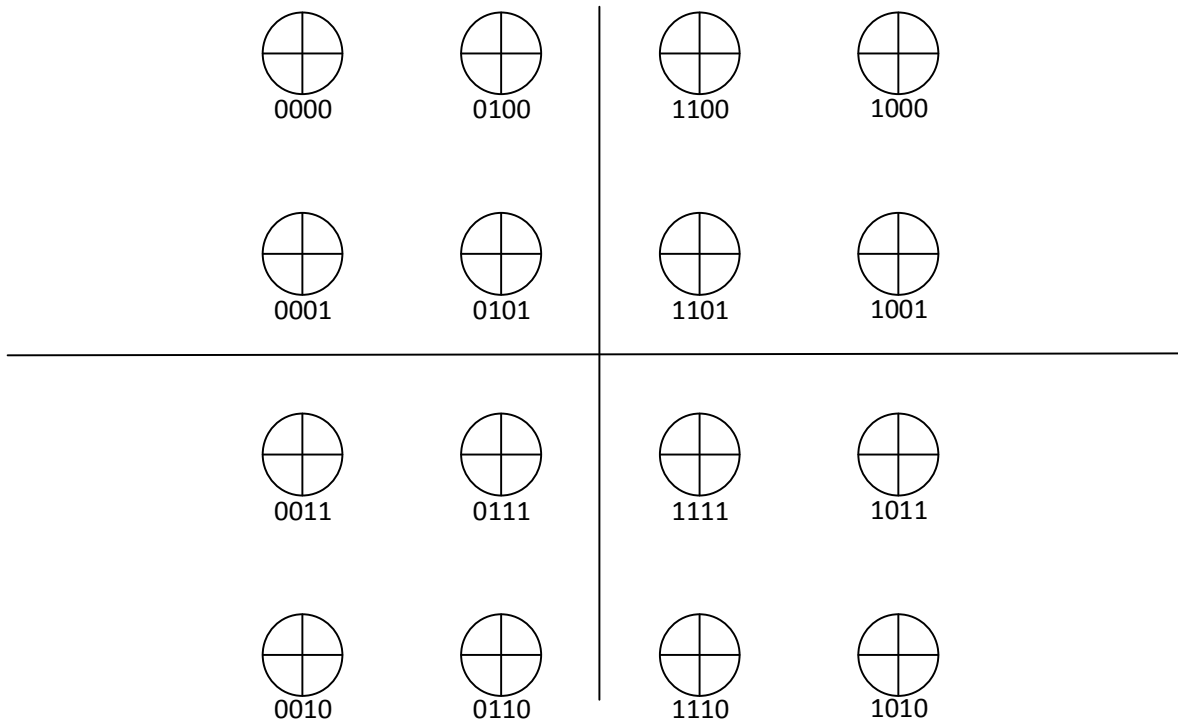


Figure 1 - Example 16-QAM Constellation

QAM transfer of data requires a base sine wave, then equipment for near-instant modification of the phase and/or amplitude of the sine wave in order to transition from one constellation station to another. The typical method involves generating phase-locked sine and cosine waves, using the Q value to modulate the amplitude of the cosine wave, and the I value to modulate the amplitude of the sine wave, then combining the two using a mixer to create the transmitted RF output. The bandwidth required for transmitting the data is determined by the symbol rate and Nyquist’s Theorem.

Multiple simultaneous streams of information can be transferred by using wideband transmitters/receivers, and using multiple center-carrier sine waves of different frequencies, separated by at least the bandwidth of each individual stream.

The advantages of QAM applied to a digital world have begun to be sensed by many in the field. Recent increases in demand for internet bandwidth-to-the-house have taxed the limits of earlier modulation techniques over twisted-pair copper, which is ubiquitous due to its use for POTS (Plain Old Telephone Service) delivery. The International Telecommunications Union (ITU) has responded to this rapid ramp-up in bandwidth demand with first DSL, then VDSL (G993.1) and in 2007, VDSL2 (G993.2), for using RF over twisted-pair and coding techniques such as QAM to create multiple subscriber bands over a single twisted pair.

These methods, while intriguing, are substantial overkill for a typical digital system (e.g. single card, or backplane-interconnected cards) where most of the information is transferred point-to-point with fan-out handled by nodes dedicated to such rather than using multi-drop lines. At very high speeds, the stubs associated with multiple listeners become difficult to manage, and line reflections, in-coupled noise, and other issues quickly drive high-speed data systems to full-mesh point-to-point designs.

Furthermore, attempting to size-reduce the required encoders, transmitters, receivers, and decoders for VDSL and/or VDSL2 systems so that they could be used as generic chip-to-chip interfaces is beyond problematic.

The need exists for a communication method which can substantially increase the capacity of existing I/O without consuming significant on-chip resources or power budgets. The method described herein will permit such, providing an order-of-magnitude or more increase in I/O data rates over chip-to-chip busses or wired networks without an increase in pin-count or increases in power budget.

Wired-nQAM

Transmitting simple nQAM over a differential pair in the digital world would require transmitting a differential sine wave signal of sufficiently high frequency that the nQAM modulation-induced bandwidth remains well above DC. For a 10 gigabit-per-second data rate, using 256QAM to encode eight bits per symbol, a symbol rate of 1.25GHz is required, requiring a multi-gigahertz carrier (and inverse) whose phase is precisely controlled. The existing XAUI interfaces achieve this bit rate by using a 64-bit parallel bus clocked at 156.25MHz, but transmitting a modulated carrier which **contains** this information density would require a much higher line frequency.

Since we are already approaching physical limits for wire-based frequencies, increasing the base frequency significantly above the current bit rates is not a practical solution.

If, instead of transmitting a modulated carrier, we treat the Q axis as representing the voltage or current on one of the pair of wires, and the I axis as representing the voltage or current on the other of the pair, we can encode a QAM-like signal as a pair of voltages or currents on a set of paired wires. In a 16QAM constellation, each group of four bits would encode a pair of voltages or currents, one on the Q-line, the other on the I-line. These voltages or currents can be coupled onto the respective wires using a good voltage-follower design or switchable current-source design. As in the XAUI approach, a separate clock which identifies symbol intervals is connected to the devices at each end of the pair to synchronize transmitters and receivers. This approach of representing the location on an axis from a QAM constellation by a voltage (or current)

The system can be extended to a 3-D nQAM constellation by adding a third wire with its own voltage or current states (3-wire-nQAM), or 4-D nQAM with a fourth wire (4-wire-nQAM), as desired. Note that a 4-wire-nQAM system, with only four voltage or current states per wire, can transmit one eight-bit symbol per transition, using only four wires instead of the 16 required by an eight-bit XAUI interface.

This is intriguing, but raises a few issues which could impact whether this can be done.

- Generating precision voltages or currents is required to represent the different stations in the constellation.
- With differential pairs, capacitive coupling is used. Continuous transmission of a single bit value could saturate the coupling capacitors at either end of the pair, so encoding methods such as 8b/10b with embedded commands have been created to allow bit randomization/inversion to prevent this.
- Differential pairs are used because the EM fields from the pair self-cancel *if impedance is correctly controlled and wire lengths match closely.*

Precision Voltages

As long as the stations in the nQAM constellation represent voltages which are offset by easily-generated gaps (e.g. on the order of a diode-drop), generating precision voltage/current references on both transmitter and receiver is fairly straight-forward. While thermal gradients between transmit and receive devices may impact the matching from device to device, protocols analogous to 8b/10b, 64b/66b, or Transition Minimized Differential Signaling (TMDS) can be used to send packets of data which exercises the full range on the transmitter so the receiver can perform local calibration and adjust to drift across time/temperature.

If stations represent currents, precision current sources can be created which can be summed through a precision Norton-style amplifier to create the pin current. Analogous encoding methods can be used for data transport and similar calibration and drift compensation can be done in a current-based design.

Pin driver design has already been briefly discussed. In one implementation, a bank of precision voltages is generated, and a low-impedance-output, high-impedance-input voltage-follower is switched via pass-gates from one voltage to another by a decoder fed by the bits which control that particular pin from the set of bits associated with the currently transmitting symbol. The follower is designed to generate an edge compliant with design specifications.

Capacitor Saturation

In XAUIs and similar capacitive-coupled data transfer systems, encoding systems such as 8b/10b or 64b/66b are used for the data to control DC balance, emissions, transition density, etc.

In this system, for DC balance a given wire must transition from one sign to opposite sign regularly, similar to a differential pair design, but the capacitor is no longer switching between charge and discharge. Now it is switching between RATES of charge and discharge. A similar encoding system can readily be devised which monitors the transmitted bit stream and shuffles bit order, bit encoding, etc., in order to maintain DC balance, emissions, effective transition density, etc., analogous to the 8b/10b, 64b/66b, and other encoding systems mentioned or existing in the industry.

In a wired-QAM system using more than two wires, e.g. a 4-wire-nQAM system the “dimension” controlled by a given wire can be changed in order to help achieve DC balance across all wires.

Emissions

The potential increase in EM emissions is perhaps the most problematic, but also perhaps the easiest to address. Digital systems are prone to significant spectral emissions, especially as processor and data rates have shot upwards. In order to significantly reduce emissions, many modern designs run the differential pairs for XAUIs and other high speed connections on interior layers of a card, with ground and/or voltage planes above and below the pairs to guarantee low radiated emissions from the high speed lines. The same technique can be used with this approach to keep any increased emissions contained between the planes within the card.

Emissions are also related to the edge speeds on the transmitted signals. Since this technique permits transmission of multiple bits per symbol (AKA voltage change on one or more wires) the symbol rate may be able to be reduced, permitting a reduction in edge rates, which in turn reduces radiated emissions. Pin drivers are already designed to drive specified loads with guaranteed edge rates. The voltage-follower approach is a straight-forward modification of existing designs.

Emissions can also be reduced by only permitting one wire to change state at a time. This has the advantage on the receiver side of allowing assumption that any voltage change in the “stable” line maps to a common-mode change in the moving line, as most transients common-mode couple onto adjacent lines. In a 256QAM system this means that a symbol-change really only transmits four bits of information because from any given station the next symbol can only be a side-to-side or up-and-down movement depending on which wire is allowed to move during the current symbol window.

Finally, in the layout of the signals, a ground wire can be laid between each symbol-carrying wire-cluster. This will prevent coupling of signals from the outside wires of a given cluster to the outside wires of an adjacent cluster.

Of these three areas of concern, emissions is probably the most difficult to closely control in the method discussed thus far because there is no longer a true differential pair whose EM emissions cancel in the far field. For lower data rates, however, the approach discussed thus far may suffice and offers **enormous potential for bandwidth gain** across existing I/O boundaries with low circuit-area impact on the transmit/receive sides of the boundary. One could also add differential signaling without loss of generality where each differential pair transmits a positive and a negative signal for the given nQAM dimension. While this would double the number of wires, the bandwidth increase by using the non-binary encoding still reduces the total from existing designs.

It should be noted that the standard oscilloscope-based “eye diagram” for determining signal quality is no longer available in this system, but it is replaced by a plot of the constellation points. Over a long dwell with every constellation point visited, such an X-Y graph will show tight points for systems with high fidelity and broad/blurry/smeared points for systems with low fidelity. The blurrier (or broader) the points in the constellation, the lower the fidelity.

We have examined wired-nQAM over more than two wires. ***A variation on the encoding using three wires reveals a system which not only offers significant bandwidth increase, but also returns to the self-cancelling EM fields of a true differential system.***

Triphase Wired-nQAM

First a few well-understood but unrelated observations: **First**, one of the major sources of radiated emissions from digital equipment is switching transients. As a bit on a CMOS integrated circuit changes state from 0 to 1 or from 1 to 0, the nature of the design has a brief moment when there is a fairly high current path from power to ground, resulting in a current (and emissions) spike. This is true even on differential pairs, where the two wires are switched to opposite states. During the transition there are high $\frac{dV}{dT}$ and $\frac{dI}{dT}$ transients in both wires. With perfect switching these transients would be exactly in synch and of opposite sign and would still cancel in the far field. Unfortunately switching is rarely perfect so switching transients radiate even from well-matched differential pairs. **Second**, power generation companies have long used three-phase wires for long-distance power transmission, with each successive wire carrying a voltage or current which is 120° phase advanced or retarded from the previous wire. For a three-wire pair, this gives relative phasing of 0°, 120°, and 240°, or alternately, 0°, 120°, and -120°. As long as demand on each wire is identical, analysis shows that power per unit time is constant. **Finally**, the bulk of the power consumed in CMOS integrated circuits is consumed during switching. This is related to the first observation, but applies to power consumption not radiated emissions. In this case, the voltage across the CMOS device remains fairly constant (with good capacitive decoupling) but very large $\frac{dI}{dT}$ transients occur at every gate which changes binary state.

These observations, coupled with the recognized ability to use wired-nQAM with more than two wires allows us to create a system which has potential **to transmit data from one device to another device using constant-power and near zero radiated emissions**. With careful design we can approach the desired state of no switching transients associated with transition from one transmitted symbol to another, and no more power fluctuations on either transmitter or receiver associated with data transfer. Furthermore, as die-shrink continues, it may be possible to use this system for data transfers **both external and internal** to complex devices. And finally, **with appropriate choice of wire-drivers, it may be possible to achieve all of this with lower total power consumption than with existing I/O methods**.

1. We use three wires to transmit the signal.
2. We continue to use one nQAM axis (typically the I-axis) to represent the amplitude of the nQAM signal **as if the data were being transmitted via a modulated RF sine wave**. We will refer to this non-existent sine wave as a “virtual-sine-wave” (VSW).
3. We use the other nQAM axis (typically the Q-axis) to represent the phase the VSW would have if it were actually transmitted, exactly like RF nQAM.

Now, arbitrarily designate one of the three wires as the “Q” wire, representing a 0° phase shifted virtual sine wave. One of the remaining wires is designated as the “R” wire, representing the same virtual sine wave, but at a location 120° phase shifted along the wave. The remaining wire is designated as the “S” wire, representing the same virtual sine wave (VSW), but at a location 240° phase shifted along the wave. For PCB layouts, the three wires can be oriented as planar, vertical, or triangle cross-sections, with typical arrangements shown below. **Just as with differential pairs, controlled impedance is critical**.

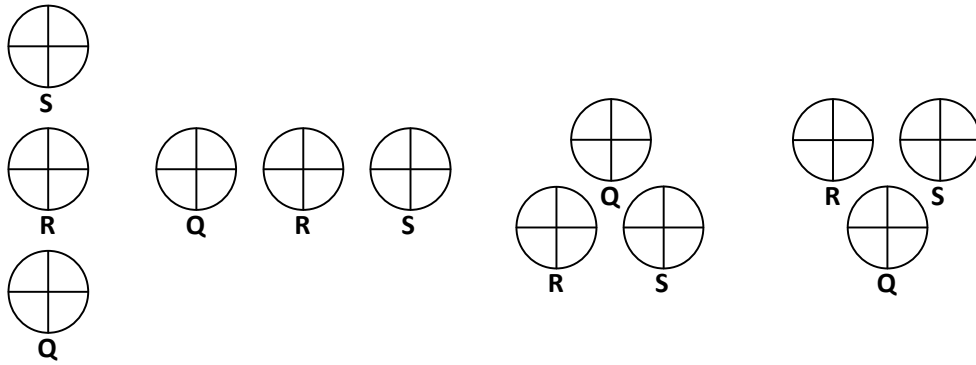


Figure 2 Possible QRS trace configurations

In translating the nQAM constellation to voltage/current values for QRS, a table like the one below is used. This example shows the QRS voltage coding for a voltage triphase-16QAM system. The 16 station codings are split with two bits for amplitude and two bits for phase. For every station, the QRS wires receive a voltage or current representing the voltage or current value of the VSW at appropriate phase offsets from the VSW represented by the station amplitude/phase values.

The Q wire directly encodes the Virtual Sine Wave amplitude or phase as a current or voltage value appropriate for the current station in the QAM constellation.

The R wire advances (or retards) by 120° along the Virtual Sine Wave and encodes the amplitude or phase at that location.

The S wire advances by 240° or retards by -120° along the Virtual Sine Wave and encodes the amplitude or phase at that location.

Shown below is a lookup table for QRS values for a triphase wired-36QAM constellation where amplitude for the stations are spaced at four equal intervals +0.25, +0.50, +0.75, and +1.0. (Note this facilitates scaling to other voltage/current ranges.) The phase for the stations are equally spaced around the 360° circle, at 45° intervals. This gives eight “radials” and four positions on each radial for a total of 32 “stations” in the QAM constellation, for a five-bit coding system over three wires. Note that if negative amplitudes are used, then the phasing choices must be such that a negative amplitude does not map onto another phase radial that is 180 degrees from the current one. We will assume a voltage-oriented design in the following discussion without loss of generality.

Amp	Phase	Q	R (+120)	S (-120)	QR	QS	RS
0.25	0	0	0.216506	-0.21651	-0.21651	0.216506	0.433013
0.25	45	0.176777	0.064705	-0.24148	0.112072	0.418258	0.306186
0.25	90	0.25	-0.125	-0.125	0.375	0.375	0
0.25	135	0.176777	-0.24148	0.064705	0.418258	0.112072	-0.30619
0.25	180	3.06E-17	-0.21651	0.216506	0.216506	-0.21651	-0.43301
0.25	225	-0.17678	-0.0647	0.241481	-0.11207	-0.41826	-0.30619
0.25	270	-0.25	0.125	0.125	-0.375	-0.375	0
0.25	315	-0.17678	0.241481	-0.0647	-0.41826	-0.11207	0.306186
0.5	0	0	0.433013	-0.43301	-0.43301	0.433013	0.866025
0.5	45	0.353553	0.12941	-0.48296	0.224144	0.836516	0.612372

0.5	90	0.5	-0.25	-0.25	0.75	0.75	0
0.5	135	0.353553	-0.48296	0.12941	0.836516	0.224144	-0.61237
0.5	180	6.13E-17	-0.43301	0.433013	0.433013	-0.43301	-0.86603
0.5	225	-0.35355	-0.12941	0.482963	-0.22414	-0.83652	-0.61237
0.5	270	-0.5	0.25	0.25	-0.75	-0.75	0
0.5	315	-0.35355	0.482963	-0.12941	-0.83652	-0.22414	0.612372
0.75	0	0	0.649519	-0.64952	-0.64952	0.649519	1.299038
0.75	45	0.53033	0.194114	-0.72444	0.336216	1.254774	0.918559
0.75	90	0.75	-0.375	-0.375	1.125	1.125	0
0.75	135	0.53033	-0.72444	0.194114	1.254774	0.336216	-0.91856
0.75	180	9.19E-17	-0.64952	0.649519	0.649519	-0.64952	-1.29904
0.75	225	-0.53033	-0.19411	0.724444	-0.33622	-1.25477	-0.91856
0.75	270	-0.75	0.375	0.375	-1.125	-1.125	0
0.75	315	-0.53033	0.724444	-0.19411	-1.25477	-0.33622	0.918559
1	0	0	0.866025	-0.86603	-0.86603	0.866025	1.732051
1	45	0.707107	0.258819	-0.96593	0.448288	1.673033	1.224745
1	90	1	-0.5	-0.5	1.5	1.5	0
1	135	0.707107	-0.96593	0.258819	1.673033	0.448288	-1.22474
1	180	1.23E-16	-0.86603	0.866025	0.866025	-0.86603	-1.73205
1	225	-0.70711	-0.25882	0.965926	-0.44829	-1.67303	-1.22474
1	270	-1	0.5	0.5	-1.5	-1.5	0
1	315	-0.70711	0.965926	-0.25882	-1.67303	-0.44829	1.224745

The phase arrangement for any given constellation can be arranged to give best separation and decidability for a given design or system. In this second encoding if a transmitter were sending to a receiver where only the ground is shared, the receiver could auto-align by determining the absolute min/max voltages on the Q wire and using those to scale the R & S wires to recover the constellation.

A couple of observations:


- Because this system is modeled after the ubiquitous tri-phase power transmission system, and the layout has controlled impedance, we know a priori that the steady-state power at every station is *identical* or very nearly so, for a given amplitude.
- While a transition from one station to another station may produce little or no voltage/current change in one wire of the QRS trio, one or more of the other wires experiences a larger-range digital-like voltage swing.
- External transients will tend to couple equally onto all three wires, which still permits decoding a station based on the relative values between the various wires rather than their absolute individual values.
- Because the positions of the stations in the constellation are arbitrary, they may be relocated to achieve voltage/current ratios which are easier to generate and decode, rather than using forced equal spacing.
- Using the difference values between QR, QS, and SR gives common-mode noise rejection and still gives a good set of detection values.

The identical power-per-station for a given Amplitude, gives us strong implication that we *may* be able to achieve constant power *during* transitions as well, achieving a digital transmission system with zero switching emissions.

As noted previously, the traditional “eye diagram” associated with differential pairs is no longer directly available. Further note that in many possible encodings, NO STATION HAS THE FULL AMPLITUDE VALUE. As long as the QRS stream visits each station with near-same probability, however, recovery of the constellation and the associated data is possible even with severe signal drop down a transmission line, similar to how it is done with RF nQAM systems.

Transmission-side Implementation For Triphase Wired-nQAM

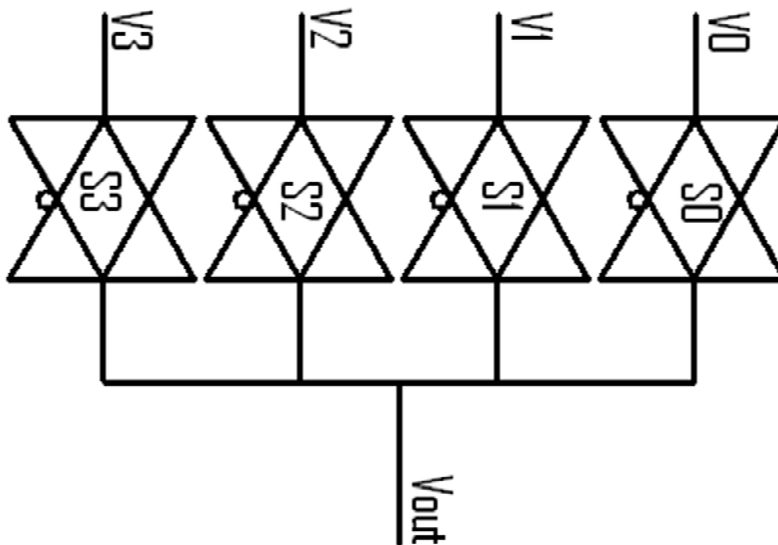
In one implementation, a standard R-2R constant-current switch ladder is used to create a voltage which is fed to a precision voltage follower/inverter:



V_r is connected to the internal reference voltage, and the follower/Inverter directly drives the wire. A look-up table drives the switching bits to select the voltage appropriate for the QRS wire being driven. One advantage of the R-2R ladder is that the various switches controlled by bits can be designed as make-before-break (e.g. MAX4625 circuit) which allows constant current flow without switching transients at the driving side.

If using the table above, V_r would be connected to +1, and the number of “driving” bits would be selected to give reasonable approximations of the different Q, R, and S values from the table.

Another implementation of the transmission side uses fixed voltage sources and CMOS transmission gates to select the appropriate voltage source from which to drive the line:




Here, multiple voltages, represented by $V_0...V_3$ (readily extended to an arbitrary number of voltages) are connected to a set of CMOS transmission gates $S_0...S_3$. The desired voltage is selected by using a decoder to select one specific transmission gate to enable. A design such as this one would internally generate a voltage reference tree which contains all of the required voltages in the Q, R, and S columns in the table above.

Either of these designs permits reasonable accuracy in V_{out} for driving each of the Q, R, and S wires.

Receive-side Implementation For Triphase Wired-nQAM

There are many possible methods of decoding specific voltages or currents on the wires in this design. What follows is a discussion of a couple of possible approaches. On the receive side, high-impedance comparators can be used to quickly determine a given wire's approximate voltage/current. Because this is a digital system we do not need a high precision determination of the voltage/current. Another approach would use low-bit flash ADC convertors to quickly find the values. Strictly comparing each wire independently against a series of references, however, would have problems with common-mode

noise spikes coupled onto the wires in the WQAM system. Instead of comparing each wire against fixed references, instead pairs of wires may be compared.



When $R_1 = R_2 = R_3$ the circuit acts as a differential follower:

$$V_{out} = V_2 - V_1$$

In the Triphase Wired-nQAM system, there would be three such differential amplifiers, computing the differences for QR, RS, and QS pairs.


In the first amplifier, Q connects to V2 and R to V1, and Vout produces the values in the QR column of the above table.

In the second amplifier, Q connects to V2 and S to V1, and Vout produces the values in the QS column of the above table.

In the third amplifier, R connects to V2 and S connects to V1, and Vout produces the values in the RS column of the above table.

Analysis of the columns QR, QS, and RS from the above table allows one to determine a set of comparator values which can uniquely identify a bit encoding based on the values from the three differential amplifiers. If we select the following comparator voltages: -1.7, -1.55, -1.4, -1.25, -1.1, -0.95, -0.8, -0.65, -0.5, -0.35, -0.2, -0.05, 0.05, 0.2, 0.35, 0.5, 0.65, 0.8, 0.95, 1.1, 1.25, 1.4, 1.55, and 1.7, we can use a bank of 20 voltage comparators on the QR, QS, and RS differential outputs to create a digital

estimate of the value on the given differential output. Note that it may be possible to obtain a unique set of detection values with a smaller optimized set of detection voltages. This set is shown as an easy to follow example.



In the diagram below, IN is connected to one of QR, QS, or RS. One comparator is used for each of the selected comparator voltages. In this example, there would be 20 comparators in the bank, and 20 “Match” or “M” values from the bank’s outputs.

The values from the bank of comparators attached to the QR output would form the column shown as “QR Match Set.”

The values for the comparators attached to the QS output form the column “QS Match Set.”

The values for the comparators attached to the RS output form the column “RS Match Set.”

Max	Min	Del+	Del-	Bits	Tag	QR Match Set	QS Match Set	RS Match Set
15	10	0.75	0.6	00000	RS	11111111110000000000000000	11111111111111000000000000	11111111111111100000000000
15	13	0.3	0.15	00001	-	11111111111111000000000000	11111111111111100000000000	11111111111111100000000000
15	12	0.45	0.3	00010	-	11111111111111100000000000	11111111111111110000000000	1111111111111111000000000000
15	10	0.75	0.6	00011	QR	11111111111111100000000000	11111111111111100000000000	11111111110000000000000000
14	9	0.75	0.6	00100	QR	11111111111111100000000000	11111111111000000000000000	11111111110000000000000000
11	9	0.3	0.15	00101	-	11111111111000000000000000	11111111110000000000000000	11111111110000000000000000
12	9	0.45	0.3	00110	-	11111111100000000000000000	11111111100000000000000000	11111111111000000000000000
14	9	0.75	0.6	00111	RS	11111111100000000000000000	11111111110000000000000000	11111111111100000000000000
18	9	1.35	1.2	01000	-	11111111100000000000000000	11111111111111000000000000	11111111111111111100000000
18	14	0.6	0.45	01001	-	11111111111111000000000000	11111111111111111000000000	11111111111111110000000000
17	12	0.75	0.6	01010	-	11111111111111110000000000	11111111111111111000000000	11111111111100000000000000
18	8	1.5	1.35	01011	-	11111111111111111000000000	11111111111111000000000000	11111111000000000000000000
15	6	1.35	1.2	01100	-	11111111111111100000000000	11111111100000000000000000	11111100000000000000000000
10	6	0.6	0.45	01101	-	11111111110000000000000000	11111100000000000000000000	11111110000000000000000000
12	7	0.75	0.6	01110	-	11111110000000000000000000	11111110000000000000000000	11111111111100000000000000
16	6	1.5	1.35	01111	-	11111100000000000000000000	11111111110000000000000000	11111111111111110000000000
21	8	1.95	1.8	10000	-	11111111000000000000000000	11111111111111100000000000	11111111111111111111110000
21	14	1.05	0.9	10001	-	11111111111111000000000000	11111111111111111111100000	11111111111111111100000000
20	12	1.2	1.05	10010	-	11111111111111111111100000	11111111111111111111100000	11111111111100000000000000
21	6	2.25	2.1	10011	-	11111111111111111111110000	11111111111111000000000000	11111100000000000000000000
16	3	1.95	1.8	10100	-	11111111111111100000000000	11111111000000000000000000	11100000000000000000000000
10	3	1.05	0.9	10101	-	11111111110000000000000000	11100000000000000000000000	11111100000000000000000000
12	4	1.2	1.05	10110	-	11110000000000000000000000	11110000000000000000000000	11111111111100000000000000
18	3	2.25	2.1	10111	-	11100000000000000000000000	11111111110000000000000000	11111111111111111100000000
24	6	2.7	2.55	11000	-	11111100000000000000000000	11111111111111111000000000	11111111111111111111111111
23	15	1.2	1.05	11001	-	11111111111111100000000000	11111111111111111111111110	11111111111111111111100000
22	12	1.5	1.35	11010	-	11111111111111111111111000	11111111111111111111111000	11111111111100000000000000
23	4	2.85	2.7	11011	-	11111111111111111111111110	11111111111111000000000000	11110000000000000000000000
18	0	2.7	2.55	11100	-	11111111111111111100000000	11111100000000000000000000	00000000000000000000000000
9	1	1.2	1.05	11101	-	11111111100000000000000000	10000000000000000000000000	11110000000000000000000000
12	2	1.5	1.35	11110	-	11000000000000000000000000	11000000000000000000000000	11111111111100000000000000
20	1	2.85	2.7	11111	-	10000000000000000000000000	11111111100000000000000000	11111111111111111111100000

In each of the “Match Set” columns there is a string of zero-or-more 1’s followed by zero-or-more 0’s.

The “Max” column gives the maximal 1’s length span across all three Match Set columns.

The “Min” column gives the minimal 1’s length span across all three Match Set columns.

The “Bits” column gives a binary 5-bit encoding for the given row.

Note that value pairs in the Max and Min columns are very nearly unique across all 32 rows, with the exception of two pairs of rows: (15,10), appearing in rows with bit encodings of 00000 and 00011, and (14,9), appearing in rows with bit encodings of 00100 and 00111.

The Del+ column gives the typical “eye” spacing for detection of the given row using the difference between the maximal 1’s span and the minimal 1’s span for that row. The spacing is based on the average spacing of the comparator voltages, in this case approximately 0.15 per comparator.

The Del- column shows the “eye” spacing if the maximal span is just OVER the last ‘1’ comparator detection threshold, while the minimal span is just UNDER the first ‘0’ comparator threshold, which reduces the “eye” by about 0.15 in this example.

It can be seen that for all rows there is at least 0.15 separation in the “eye” worst case and typically 0.3 separation, which is readily detectable in a digital system using common-mode noise rejection as in this approach.

The two pairs mentioned above, however, both have dual-hits for the (Max,Min) values.

Looking at the QR, QS, and RS patterns for the (15,10) pair, we see that the maximal 1 span for bit row 00000 occurs in the RS column, while in the bit row 00011 it occurs in the QR column, making it trivial to distinguish between the two. Similarly, for the (14,9) pair, we see the maximal 1 span for bit row 00100 occurs in the QR column, while for bit row 00111 it occurs in the RS column.

A Mild Optimization Exercise

The following two tables show the same design, updated by modifying only the phasing of the constellation so that it starts at 30° and uses the same 45° spacing, and using a reduced set of only 12 comparator voltages, a design results with similar robustness but with about ½ of the required receive-side circuitry.

The phases are 30°, 75°, 120°, 165°, 210°, 255°, 300°, and 345°.

The comparator voltages are -1.4, -1.18, -1, -0.55, -0.28, -0.05, 0.05, 0.28, 0.55, 1, 1.18, and 1.4.

As with the previous design there are Min/Max conflicts which are readily resolved: (8,4), (9,1), (9,3), and (11,3) all of which are trivially resolved by large differences between the QR and RS columns, so again a simple lookup table decodes the five bits.

Breaking The I/O Bottleneck

Amp	Phase	Q	R (+120)	S (-120)	QR	QS	RS
0.25	30	0.125	0.125	-0.25	0	0.375	0.375
0.25	75	0.241481	-0.0647	-0.17678	0.306186	0.418258	0.112072
0.25	120	0.216506	-0.21651	0	0.433013	0.216506	-0.21651
0.25	165	0.064705	-0.24148	0.176777	0.306186	-0.11207	-0.41826
0.25	210	-0.125	-0.125	0.25	0	-0.375	-0.375
0.25	255	-0.24148	0.064705	0.176777	-0.30619	-0.41826	-0.11207
0.25	300	-0.21651	0.216506	3.06E-17	-0.43301	-0.21651	0.216506
0.25	345	-0.0647	0.241481	-0.17678	-0.30619	0.112072	0.418258
0.5	30	0.25	0.25	-0.5	0	0.75	0.75
0.5	75	0.482963	-0.12941	-0.35355	0.612372	0.836516	0.224144
0.5	120	0.433013	-0.43301	0	0.866025	0.433013	-0.43301
0.5	165	0.12941	-0.48296	0.353553	0.612372	-0.22414	-0.83652
0.5	210	-0.25	-0.25	0.5	0	-0.75	-0.75
0.5	255	-0.48296	0.12941	0.353553	-0.61237	-0.83652	-0.22414
0.5	300	-0.43301	0.433013	6.13E-17	-0.86603	-0.43301	0.433013
0.5	345	-0.12941	0.482963	-0.35355	-0.61237	0.224144	0.836516
0.75	30	0.375	0.375	-0.75	0	1.125	1.125
0.75	75	0.724444	-0.19411	-0.53033	0.918559	1.254774	0.336216
0.75	120	0.649519	-0.64952	0	1.299038	0.649519	-0.64952
0.75	165	0.194114	-0.72444	0.53033	0.918559	-0.33622	-1.25477
0.75	210	-0.375	-0.375	0.75	0	-1.125	-1.125
0.75	255	-0.72444	0.194114	0.53033	-0.91856	-1.25477	-0.33622
0.75	300	-0.64952	0.649519	9.19E-17	-1.29904	-0.64952	0.649519
0.75	345	-0.19411	0.724444	-0.53033	-0.91856	0.336216	1.254774
1	30	0.5	0.5	-1	0	1.5	1.5
1	75	0.965926	-0.25882	-0.70711	1.224745	1.673033	0.448288
1	120	0.866025	-0.86603	0	1.732051	0.866025	-0.86603
1	165	0.258819	-0.96593	0.707107	1.224745	-0.44829	-1.67303
1	210	-0.5	-0.5	1	0	-1.5	-1.5
1	255	-0.96593	0.258819	0.707107	-1.22474	-1.67303	-0.44829
1	300	-0.86603	0.866025	1.23E-16	-1.73205	-0.86603	0.866025
1	345	-0.25882	0.965926	-0.70711	-1.22474	0.448288	1.673033

			Max	Min	Del+	Del-	Bits	Tag	QR Match Set	QS Match Set	RS Match Set
111111000000	111111110000	111111110000	8	6	0.3	0.15	00000	-	111111000000	111111100000	111111100000
111111110000	111111110000	111111100000	8	7	0.15	0	00001	-	111111100000	111111100000	111111100000
111111110000	111111100000	111110000000	8	5	0.45	0.3	00010	-	111111100000	111111100000	111110000000
111111110000	111110000000	111100000000	8	4	0.6	0.45	00011	QR	111111100000	111100000000	111100000000
111111000000	111100000000	111100000000	6	4	0.3	0.15	00100	-	111111000000	111100000000	111100000000
111100000000	111100000000	111110000000	5	4	0.15	0	00101	-	111100000000	111100000000	111110000000
111100000000	111110000000	111111100000	7	4	0.45	0.3	00110	-	111100000000	111100000000	111111000000
111100000000	111111100000	111111110000	8	4	0.6	0.45	00111	RS	111100000000	111111100000	111111100000
111111000000	111111111000	111111111000	9	6	0.45	0.3	01000	-	111111000000	111111110000	111111110000
111111111000	111111111000	111111100000	9	7	0.3	0.15	01001	-	111111110000	111111110000	111111100000
111111111000	111111110000	111100000000	9	4	0.75	0.6	01010	-	111111110000	111111100000	111100000000
111111111000	111110000000	111000000000	9	3	0.9	0.75	01011	QR	111111110000	111100000000	111000000000
111111000000	111000000000	111000000000	6	3	0.45	0.3	01100	-	111111000000	111000000000	111000000000
111000000000	111000000000	111110000000	5	3	0.3	0.15	01101	-	111000000000	111000000000	111110000000
111000000000	111100000000	111111100000	8	3	0.75	0.6	01110	-	111000000000	111000000000	111111100000
111000000000	111111100000	111111111000	9	3	0.9	0.75	01111	RS	111000000000	111111100000	111111110000
111111000000	111111111100	111111111100	10	6	0.6	0.45	10000	-	111111000000	111111111000	111111111000
111111111000	111111111110	111111110000	11	8	0.45	0.3	10001	-	111111110000	111111111000	111111100000
111111111110	111111111000	111000000000	11	3	1.2	1.05	10010	QR	111111111000	111111110000	111000000000
111111111000	111100000000	100000000000	9	1	1.2	1.05	10011	QR	111111110000	111100000000	100000000000
111111000000	110000000000	110000000000	6	2	0.6	0.45	10100	-	111111000000	110000000000	110000000000
111000000000	100000000000	111100000000	4	1	0.45	0.3	10101	-	111000000000	100000000000	111100000000
100000000000	111000000000	111111111000	9	1	1.2	1.05	10110	RS	100000000000	111000000000	111111110000
111000000000	111111110000	111111111110	11	3	1.2	1.05	10111	RS	111000000000	111111100000	111111111000
111111000000	111111111111	111111111111	12	6	0.9	0.75	11000	-	111111000000	111111111111	111111111111
111111111110	111111111111	111111110000	12	8	0.6	0.45	11001	-	111111111100	111111111111	111111100000
111111111111	111111111000	111000000000	12	3	1.35	1.2	11010	-	111111111111	111111110000	111000000000
111111111110	111100000000	000000000000	11	0	1.65	1.5	11011	-	111111111100	111100000000	000000000000
111111000000	000000000000	000000000000	6	0	0.9	0.75	11100	-	111111000000	000000000000	000000000000
100000000000	000000000000	111100000000	4	0	0.6	0.45	11101	-	100000000000	000000000000	111100000000
000000000000	111000000000	111111111000	9	0	1.35	1.2	11110	-	000000000000	111000000000	111111110000
100000000000	111111110000	111111111111	12	1	1.65	1.5	11111	-	100000000000	111111100000	111111111111

Summary

The given design, without any optimization, provides a robust mechanism for decoding five bits per transition over a trio of wires. When compared with a single bit per transition of a differential pair, we find that in 6 wires using differential pairs we can transmit three bits per transition, while in this system we can transfer TEN bits per transition, for a bit density increase of over 3 times. This comes at the expense of additional circuitry.

Every differential pair already has effectively a driver per wire. The resistors needed to form the R-2R switch or to create the desired voltage values for the pass-gate implementation require very little area. Pass-gates are similarly small, low-current devices, so there is minimal impact on the transmit side.

On the receive side, the example implementation uses one differential amplifier per pair, which is analogous the differential receiver for a standard differential pair, so there is a 50% increase in amplifiers (9 amplifiers for six wires instead of the current 6 amplifiers for six wires) and in chip area for the amplifiers.

The bank of 12 comparators per amplifier is new, and the output match sets requires a simple 32-entry 36-to-5 bit lookup table to translate them into the appropriate five-bit value. All of this circuitry is fast and small. Optimization of the Match voltages can produce a design which uses significantly fewer comparators.

Note that if more noise is present in the system a direct-match can be replaced by a statistical similarity match similar to PRML (Partial Response Maximum Likelihood) techniques to extract the most probable 5-bit result.

The resulting design provides a system which is capable of matching existing differential-pair transmission speeds, but provides over 3x the number of bits per IO pad on an integrated circuit.

With a small amount of rework, a 10 gigabit-per-second design becomes a 50 gigabit-per-second design at the same clock speeds.

A 15-bit differential bus, with 30 wires, capable of transferring 15 bits per transition, becomes a set of 10 triphase nQAM sets, each transferring 5 bits in this example, for a capacity of 50 bits per transition. Conversely, the same 15-bit bus with 30 wires could be replaced with three triphase nQAM sets requiring only 9 wires instead of the 30 for differential signaling, freeing up an enormous 21 wires for use to expand the design.

With additional precision on the voltages and the decoding, the constellation size could readily be increased beyond the demonstrated 32 stations.

With triphase Wired-nQAM, the existing differential advantages of constant power (for a given amplitude) are retained, and with modification of the driving amplifiers, a design which preserves constant power across amplitude changes can be created, giving a significant reduction in switching transients associated with data transfer from one integrated circuit to another. Backplanes and card connectors, also a bottleneck source, could be redesigned to use wire triplets, with either significant

reduction in connector pin count, significant increase in bandwidth over the same size connector, or a combination of both.

With Wired-nQAM, at slower switching speeds, with each wire representing one dimension in an N-dimensional constellation, and each dimension represented by K values, an N-wire system can carry K^N values instead of 2^N in existing digital systems. If K is eight (2^3), for example the N-wire capacity becomes 2^{3N} on the same N-wire bus. An 8-wire bus can now carry 2^{24} values instead of only 2^8 . Existing differential-pair style connectors and backplanes can still be used, with slower switching speeds due to the increased bit carrying capacity over the same conductors. A 10 gigabit design requiring a 64-bit bus clocked at 156.25MHz could be redesigned with $K = 8$ and a voltage range of +2 to -2V (2.0, 1.43, 0.857, 0.286, -0.286, -0.857, -1.43, -2.0) so that each pair carries three bits per transition instead of a single bit. The same 10 gigabit rate could be achieved with either 1/3 of the wires at the same clock rate, or the same number of wires could be used and the clock rate be reduced by 1/3 to about 53MHz, significantly reducing emissions.

Note that existing techniques such as Trellis-Modulation-Coding to improve noise immunity can readily be applied to this system. In 1000BASE-T Ethernet, for example, a 2/3 TCM approach allows encoding 2 bits per transition using PAM-5. With this system, one could readily design an increased bandwidth system which used bi-directional capacitive-coupled signaling over CAT5/CAT6 cabling that only required three of the four twisted pairs in such cabling if a QRS design were used, and using all four twisted pairs if a QRST design were used. If the current CAT5/CAT6 wiring were replaced with twisted-trio wiring (using techniques analogous to those used in CAT5/CAT6 wiring), each trio could be used for bi-directional signaling with five or more bits per symbol instead of the current two. A six-wire cable could carry ten or more bits per symbol instead of the eight-wire-and-8-bits-per-symbol of current 1000BASE-T systems. Conversely, a nine-wire cable could easily carry 15 or more bits, nearly doubling the bandwidth with only one more wire. If the constellation is increased to 64, 128, or 256 stations, with TCM coding to recover noise immunity, one could reach capacities of 8 bits per symbol per trio, and the nine-wire cable could carry three times the bandwidth at the same symbol rate as 1000BASE-T.

Using this system, it now becomes possible to significantly increase the bandwidth available chip-to-chip and over wired networks, bringing I/O pathways into line with the growth of processing power.

The I/O bottleneck encountered by parallel high-speed systems has been broken and a path forward shown which will prevent IO pad size from being the integrated circuit bottleneck for the foreseeable future.