

Digital Output Data Links from Superconductor Integrated Circuits

Deepnarayan Gupta¹, *Fellow, IEEE*, Saad Sarwana¹, Dmitri Kirichenko¹, Vladimir Dotsenko¹, A. Erik Lehmann¹, Timur V. Filippov¹, Wei-TingWong², Su-Wei Chang², Prasanna Ravindran², and Joseph Bardin²

Abstract— Improving the efficiency in digital data transport, originating from 4 K superconductor integrated circuits (ICs) to standard room-temperature (300 K) electronics, both in terms of speed and power consumption, is vitally important for nearly all applications of superconductor digital electronics. The required data link must include all necessary amplification and conditioning to convert digital data from single flux quantum (SFQ) logic to standard CMOS logic levels. Therefore, one must minimize the overall power consumption for the required total data throughput by taking into account all parts of the link: (a) data transmitters or output drivers with SFQ input, (b) additional amplifiers, (c) transmission lines, and (d) data receivers at room-temperature, such as the ones available on commercial field-programmable gate arrays (FPGAs). Here, we describe two variants of electrical data links, with and without cryogenic semiconductor amplifiers.

Index Terms—Driver circuits, Mixed analog digital integrated circuits, Cryogenic electronics

I. INTRODUCTION

SUPERCONDUCTOR digital and mixed signal electronics have a number of current and future areas of application. It is hard to imagine many applications that will not benefit from better digital data links from superconductor electronics at 4 K to room-temperature (300 K) electronics. Improvements in the energy per bit of information transfer have an immediate positive impact on the systems comprising superconductor digital electronics. At one end of the application space are small-scale digital systems, such as digital-RF receivers [1]. These systems use only one or two integrated circuits (ICs) and are constrained by the cooling capacity of the chosen cryocooler. Even incremental improvements of data links are beneficial as they enable increasing the functionality of the total system while staying within that constraint. The impact of superior data links becomes stronger as one moves to the other end of the spectrum: large-scale supercomputers [2], where energy efficiency is the primary advantage. Therefore, efficiency of digital data links is an enabling factor for large-scale super-

Manuscript receipt and acceptance dates will be inserted here. This work was supported in part by ONR basic research grant and STTR program. (*Corresponding author: Deepnarayan Gupta*)

D. Gupta, S. Sarwana, D. Kirichenko, V. Dotsenko, A. E. Lehmann, and T. Filippov are with HYPRES' Digital-RF Circuits and Systems division, Elmsford, NY 10523 USA (e-mail: dgupta@hypres.com).

W. Ting, S-W. Chang, and J. Bardin are with the Electrical and Computer Engineering Department, University of Massachusetts Amherst, Amherst, MA 01003 USA (e-mail: jbardin@ecs.umass.edu)

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

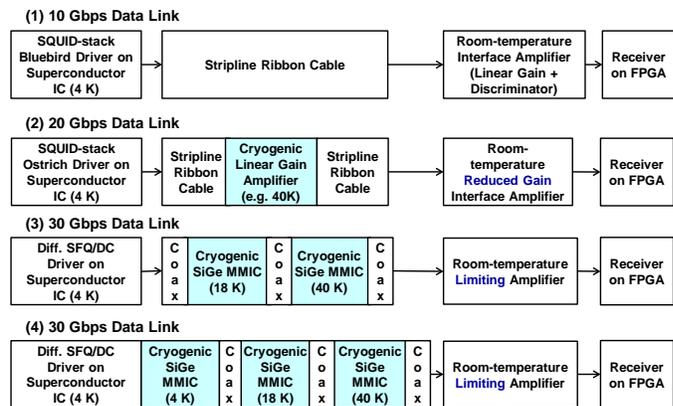


Fig. 1. Different examples of data links are shown: (1) A data link with no cryogenic amplification, (2) one stage of cryogenic amplification with stripline connections to 4 K and room temperature, (3) two stages and (4) three stages of cryogenic amplification at different temperatures starting with an SFQ/DC converter as the output driver. In all cases, the data are received by standard CML logic in the GT* receivers on a Xilinx FPGA (e.g. UltraScale+).

conductor systems. The same applies for streaming digital data processors and network switches.

Another application area is noise-immune digital transport of cryogenic sensor data. Since cryogenics are necessary to obtain the desired sensor functionality, one would expect the use of superconductor digitizers to have the least resistance. In reality, however, the excessive power cost for digital data links has been the biggest factor why that has not been the case during the last two decades.

In this paper, we will primarily address electrical data links between single flux quantum (SFQ) logic, operating near 4 K, and CMOS logic at room-temperature, recognizing that extensions to hybrid electrical-optical data links and other logic families will follow the same paradigm. The link comprises four parts: (a) data transmitters or output drivers with SFQ input in the superconductor electronics domain, (b) semiconductor amplifiers, (c) transmission lines, and (d) data receivers at room-temperature, such as the ones available on commercial field-programmable gate arrays (FPGAs). Here we illustrate this energy accounting analysis for heat loads at different temperatures for different data links that we have developed. The first, called the Bluebird data link after the eponymous SQUID-stack output driver, has been used extensively for its robustness and ease of integration in cryocooled systems. The second is a hybrid-temperature heterogeneous technology

(HTHT) data link [12] that distributes the required additional amplification in one, two, or three stages at different temperatures for faster data links while optimizing overall power consumption (Fig. 1).

II. DIGITAL DATA LINK COMPONENTS

A. Superconductor Output Drivers

Transmission of digital data out of the single flux quantum logic domain requires interface circuitry in the form of output drivers. The design of such drivers has been an active research area and many drivers have been developed [3-8]. We focus on two of these, both dc biased. The first is an SFQ/DC converter, which has a single SQUID attached to a flip-flop. When this flip-flop is in state '1', the SQUID is driven to its voltage state and produces a bunch of SFQ pulses; no output is produced when the flip-flop is in '0' state. Designs use either a reset-set flip-flop (RSFF) or a toggle flip-flop (TFF). The second amplifies the output current of the SFQ/DC converter and modulates a stack of SQUIDs, the voltages of which are summed at the output.

A single SFQ/DC produces a very small signal swing, about 650 μV for MIT-LL SFQ5ee [10] process ($J_c = 10 \text{ kA/cm}^2$) when coupled to a 50 Ω transmission line. Therefore, additional low-noise cryogenic amplification is necessary. The SQUID-stack driver [6] is such a low-noise cryogenic amplifier but it must be placed on the chip at 4 K where the power consumption is costlier than at higher temperatures. Still, the use of SQUID-stack drivers has been ubiquitous in our designs because it provides a compelling advantage in simplifying cryogenic packaging: no additional cryogenic amplifier is necessary.

We compare below a few representative drivers in terms of their speed and power consumption. Starting with original designs that worked only up to 2 Gbps [1], we have been able to dramatically improve the SQUID-stack drivers to approach 30 Gbps, which is close to the speed of a single SQUID. Unlike previous H-tree SQUID-stack drivers [6], these accomplish high speed in a compact geometry by employing a number of design techniques, the most critical being velocity matching of the excitation signal and the output signal in the sequence of SQUIDs in a stack. The excitation signal propagates through a slow-wave structure through the SQUID-stack and the propagation of the outputs down the stack is slowed down to match the excitation signal propagation speed.

The fastest SQUID-stack driver to date, named *Ostrich*, employs complementary drive of two SQUID stacks with ac-

TABLE I
 SUPERCONDUCTOR OUTPUT DRIVERS

Output Drivers	Data rate (Gbps)	Swing @50 Ω (mVpp)	Power (μW)	FOM (fj/bit)	Area (μm^2)
SQUID-stack "Dodo"	6	5	99.0	16.5	22,800
SQUID-stack "BlueBird"	14	5	134.0	9.6	29,400
SQUID-stack "Ostrich"	30	5	157.6	5.3	63,800
Differential SFQ/DC	40	1.3	13.4	0.3	12,000

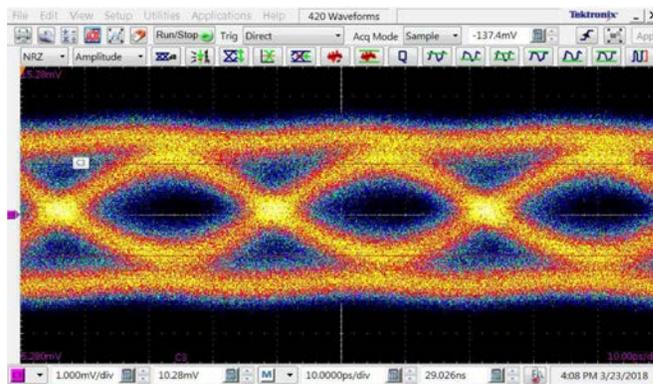


Fig. 2. Measurement of active-reset SQUID-stack (Ostrich) driver, operating at 30 Gbps, using a fast sampling module (80E09B) on Tektronix sampling oscilloscope. The chip was fabricated using the MIT-LL SFQ5ee ($J_c = 10 \text{ kA/cm}^2$) process and tested in ICE-T [9] at 3.7 K with 0.5 m Ag-coated BeCu coaxial cables (Table II).

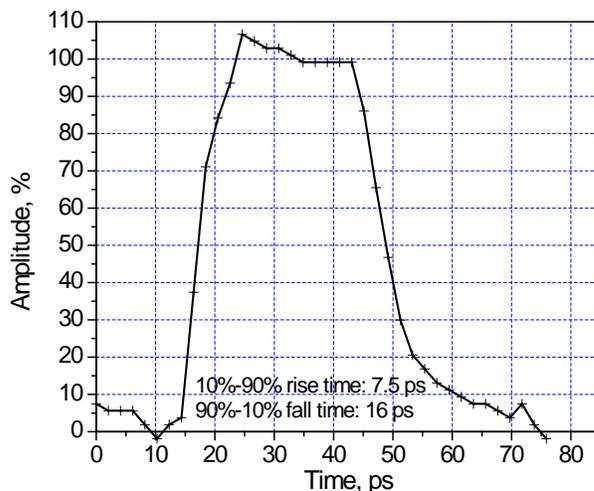


Fig. 3. Direct measurement of the Ostrich driver's output waveform, operating at 30.48 Gbps, was performed using an on-chip sampler. This measurement was done with a liquid He immersion probe at 4.2 K.

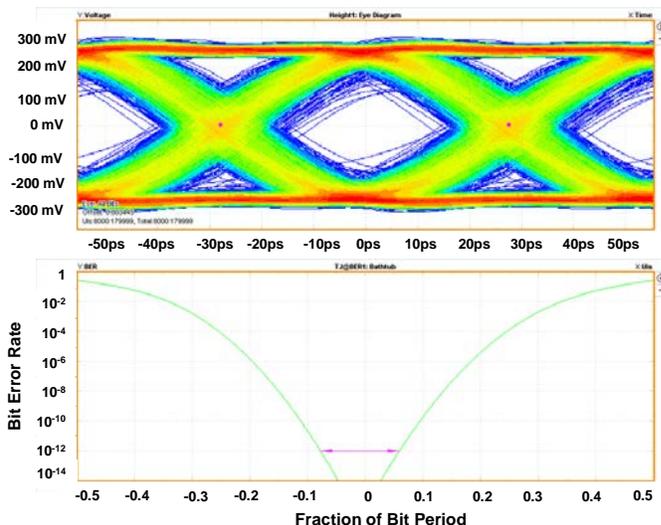


Fig. 4. Amplified output of an Ostrich driver, operating at 18 Gbps, can support bit-error rates of 10^{-12} and can be interfaced to the fast GTY data receivers on a Xilinx UltraScale FPGA.

tive reset. From measurement of rise and fall times with a sampling oscilloscope, we infer that the Ostrich driver can support up to 30 Gbps (Fig. 2). A superconductor sampling circuit was designed on the same chip as the output driver to measure its waveform directly (Fig. 3). Preliminary on-chip sampling measurement at 30.48 Gbps, indicate that rise and fall times are 7.5 ps and 16 ps respectively, confirming that operation in the 20-30 Gbps range should be possible. Currently, the room-temperature interface amplifier limits the data link to 18-22 Gbps (Fig. 4).

The next variant, called *Bluebird*, has been tested up to 16 Gbps as part of a complete data link from superconductor IC at 4 K to an FPGA at room temperature. This represents the best trade-off between speed and power consumption and has been used extensively during the last couple of years. Examples include a set of 8 Bluebird drivers carrying 8 Gbps each to transfer data at an aggregate rate of 64 Gbps from superconductor IC to FPGA. Another example is a set of 24 Bluebird drivers transferring 120 Gbps from a superconductor chip to an FPGA. In the latter example, the speed per channel is only 5 Gbps, which under-utilizes the Bluebird driver. Therefore, another slower driver with lower power consumption, named *Dodo*, was developed. The maximum data rate supported of Dodo is measured to be 6 Gbps. It is also 23% smaller in area than Bluebird and consumes 26% less power.

Another driver that is potentially the fastest is a differential SFQ/DC converter [17], which doubles the effective peak-to-peak swing and offers all the advantages of differential signaling that has caused it to be the standard for high-speed data links. This driver employs a pair of SFQ/DC converters with their RSFF inputs connected in complementary manner (the set input of one connected to the reset input of the other). Such a driver has been proven above 20 Gbps as part of a HTHT link with multiple stages of cryogenic amplification (see section II.C.).

We believe that the speed of those links was not limited by the SFQ/DC driver and data rates up to 50 Gbps can be accomplished, especially using the MIT-LL SFQ5hs (20 kA/cm²) fabrication process [10]. Packaging of the superconductor IC, currently done with pressure contacts on a printed circuit board [9], needs to be improved to support data rates above 30 Gbps. Also, commercially available FPGA data receivers do not support data rates above 33 Gbps. Further deserialization is possible at higher temperature following additional amplification [1].

An RSFQ differential SFQ/DC driver consumes 5.6 μ W for the pair of SFQ/DC converters and another 7.8 μ W (3 mA, 2.6 mV) for the network of JTLs and buffers for a total power consumption of 13.4 μ W. Part of this can be converted to ERSFQ logic for further power savings. Such a driver running at 50 GHz clock frequency (bias voltage = 0.1 mV) will con-

TABLE II
 TRANSMISSION LINES FROM 4 K TO ROOM TEMPERATURE

Transmission Line	3-dB BW (GHz)	Data Rate (Gbps)	4 K Heat load (μ W)	FOM (fJ/bit)
Ag-coated BeCu Coax	30	42.9	2473	57.7
Stripline Version 1	1.5	2.1	440	205.3
Stripline Version 2	4	5.7	740	129.5
Stripline Version 3	7	10.0	1118	111.8

sume 1 μ W for the driver itself and another 0.3 μ W (3 mA, 0.1 mV) for the JTL network.

It is well-known that differential signaling is superior for high-speed data links as they provide greater immunity to rejection of common-mode noise and interference. Differential, both unipolar and bipolar, drivers can be realized with SQUID stacks as well. These could produce 10 mV differential swing from a superconductor chip at the cost of additional power consumption and chip area. However, differential signal transmission doubles the heat load due to transmission lines.

B. Transmission Line

The transmission line of choice for conducting digital data from the superconductor IC to room-temperature has been the flexible stripline ribbon cable [1]. This choice is driven by thermal considerations. Commercially available semi-rigid coaxial cables require much greater metal cross-section for structural integrity than is necessary for electrical transmission. Consequently, they conduct too much heat down from room temperature unless the data rate per link is high. Furthermore, integrating multiple semi-rigid coaxial cables, with good thermal intercept at intermediate temperature stage(s), becomes difficult as the number increases.

As with the drivers, we compare several variants of striplines that offer different trade-offs between heat load and data rate. Using 3-dB bandwidth as a metric, we compare three BeCu-Cu-BeCu stripline ribbons, each with 10 lines. All the ribbon cables have G3PO coaxial connectors at either end for convenient connection with a cryogenic chip module at 4K and a vacuum feedthrough at 300K. The first has a bandwidth of 1.5 GHz and 4.4 mW heat load, or 440 μ W/line. The second and the third 10-stripline cables have higher bandwidth, 4 GHz and 7 GHz respectively, and commensurate 4K heat loads of 7.4 and 11.2 mW respectively. A single coaxial cable with silver-clad BeCu center conductor and BeCu outer conductor has 4K heat load comparable to a 7-GHz 10-stripline ribbon, but has <1 dB loss up to 10 GHz. In all cases, the lines are 11.5 inches long and are thermalized at an intermediate temperature of 40 K that is typical in cryocooled systems with two cold stages. The heat loads at 40K for the single coaxial cable mentioned above is 48.5 mW, whereas the 10-stripline ribbon cables produce 42.8, 61.7, and 68.7 mW for the three versions. If a cryogenic cooler has an additional temperature stage closer to 4 K, the heat load on the 4K stage will drop further.



Fig. 5. Compact 20-channel interface amplifier unit was developed to connect directly to a 20-G3PO coaxial vacuum feedthrough. A pair of twinax cables carry the differential outputs from the two 10-channel circuit boards to the FPGA data receivers.

C. Semiconductor Amplifiers

Data interface to commercial CMOS circuitry at room temperature requires additional amplification to hundreds of mV peak-to-peak swings. These interface digital amplifiers comprise a linear gain stage followed by a discriminator to produce the desired logic levels. In the simplest case, all the additional amplification is performed at room temperature.

A new compact 20-channel datalink amplifier unit was developed and built (Fig. 5). It consists of two 10 channel circuit boards. Each channel contains a low noise linear gain stage (32 dB), and a discriminator that outputs standard CML logic signal. RF connectors allow for direct plugin of this 20-channel amplifier into the 20-coax RF vacuum feedthrough of the cryogenic system. This way of connecting allows eliminating input cables that were typically 24" long and were potentially susceptible to EMI interference. Each of the 10-channel amplifier boards include a commercial interface for a ribbon of differential twin-coaxial cables for direct interface to an FPGA board.

Besides offering greater immunity to interference, this amplifier is 4 times faster, 36 times smaller, 9 times lighter, and consumes less power than its predecessor [1]. The whole unit, including the metal encasing, is $2.18 \times 1.19 \times 2.3$ cu. in. in volume and weighs 0.84 lbs. Its power consumption is 4 W, operating from a single 3.3 V power supply. Together with the Bluebird output driver and the stripline ribbon cable version 3, this interface amplifier has proven data transport from SCE-IC to FPGA at 8 Gbps per channel. In contrast, our previous 17-channel interface amplifier unit [1] had a volume of $14.5 \times 8.5 \times 1.75$ cu. in. and weighed 7.5 lbs.

The data rate per channel is limited by cross-talk from adjacent channels. Single-channel versions, packaged in individual metal enclosures support much higher rates. Using an on-chip 1-bit delta analog-to-digital converter (ADC) as a data source, full signal reconstruction was performed up to 14 Gbps using a single boxed amplifier. By changing the circuit board material to a lower loss dielectric (Fig. 6), the ADC clock frequency could be increased to 16 GHz. Since the delta ADC output needs integration for signal reconstruction, it also integrates all the errors resulting in rise in the noise floor when spectrum is computed. This makes such an ADC an appropriate on-chip data source for data link characterization.

The performance of the aforementioned amplifier may be improved further by cooling part of the linear gain stage (Fig. 7). By separately packaging the first stage of the amplifier for

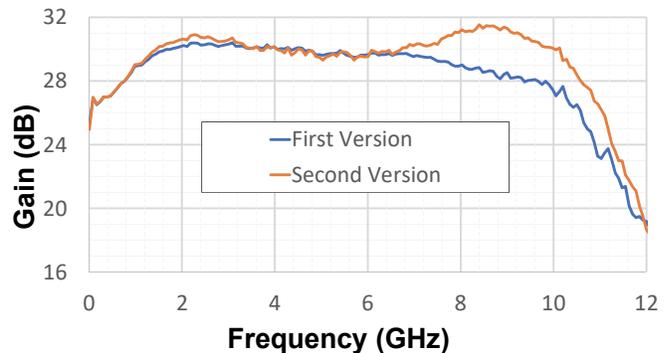


Fig. 6. Frequency response of a two-stage amplifier constructed with commercially available SiGe transistors at room temperature for two different circuit board materials. The design on the lower loss dielectric employs additional inductive peaking to compensate for the attenuation of preceding transmission lines.

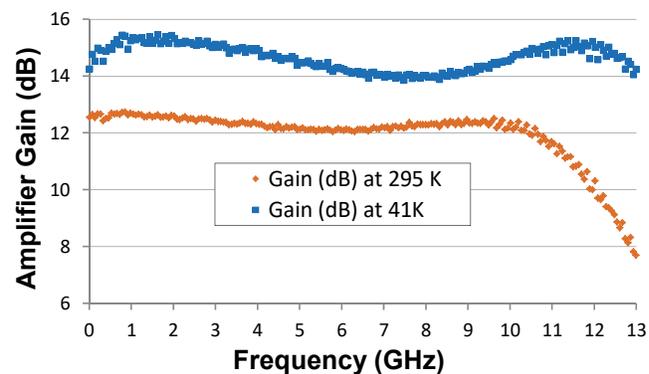


Fig. 7. Frequency response of an amplifier constructed with commercially available SiGe transistors at room temperature (295 K) and at 41 K.

the first stage of the cryocooled system (typically, 35-50 K), the maximum data link speed can be increased.

Distribution of amplification, and its associated power dissipation, along the temperature span between 4 K and room temperature can lead to further optimization of the overall link speed and power consumption. Two designs with 2 and 3 stages of cryogenic semiconductor amplifiers and one stage of room-temperature amplifier were implemented. Both start with the fastest superconductor output driver, a differential SFQ/DC driver. The two variants differ in their approaches towards broadband impedance matching between a low-impedance SFQ/DC driver and the high-impedance semiconductor amplifier chain. The first variant employs a passive match through a 50Ω series resistor integrated on chip at the SFQ/DC converter output and overcomes the 4-K attenuation with additional gain at the warmer temperatures. The second variant employs a low-power differential amplifier stage at 4 K, placed on top of the superconductor chip and connected by wirebonding.

A common superconductor IC for testing these two links was designed with pads for wirebonding and 50Ω series resistors attached to a differential SFQ/DC driver, fed by either an on-chip data source or an off-chip one through a DC/SFQ converter. When testing with the second variant, the lines to the two 50Ω resistors were cut before placing the SiGe chip.

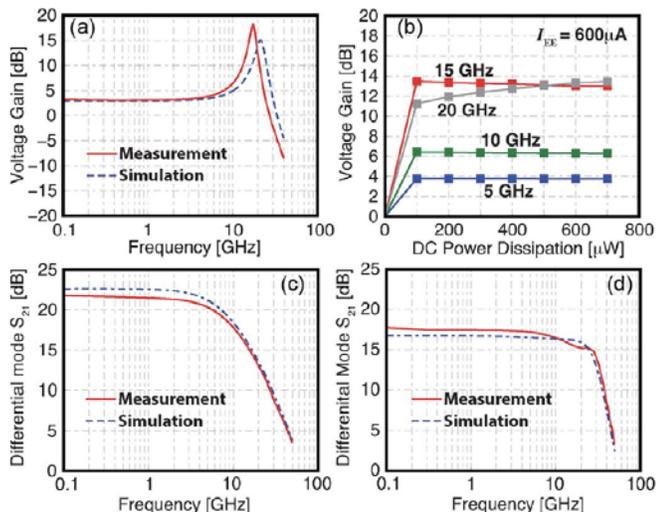


Fig. 8. Active link chipset performance. (a) Measured and simulated differential-mode voltage gain of the 4-K stage at the nominal bias point of $140 \mu\text{W}$. (b) Measured differential-mode voltage gain of the 4-K amplification stage as a function of power dissipation at the nominal bias current. (c) Differential mode S_{21} of the 18 K stage. (d) Differential-mode S_{21} of the 50 K stage.

The first variant comprises two identical gain stages (26 dB each) nominally at 18 K and 40 K. The two-amplifier chain was characterized with an external input attenuated to $400 \mu\text{V}_{\text{pp}}$ and open eye was observed up to 30 Gbps [11]. When connected to a superconductor IC with a differential SFQ/DC driver and a pair of 50Ω coaxial cables at its input and a room-temperature limiting amplifier (Hittite HMC865LC3), an open eye was observed up to 13 Gbps when tested with the on-chip ADC as the data source (Fig. 9). Full signal reconstruction was performed at 10 Gbps with data acquired through a Xilinx UltraScale FPGA (VCU 108). Using an external pseudo-random bit sequence (PRBS), converted to SFQ format using a standard DC/SFQ converter, stable operation was observed up to 22 Gbps.

The second version uses three different amplifier designs at the three stages [12-13]. The first stage has a small (3 dB) nominal gain but takes advantage of the bondwire inductance to obtain strong peaking above 10 GHz (Fig. 8(a)). This helps compensate for the lower bandwidth of the second stage, which consists of a broadband amplifier followed by an open collector amplifier. It is terminated in a differential 100Ω load to provide a broadband impedance match to the third stage, which produces broadband amplification to levels necessary to drive a commercial room-temperature limiting amplifier. With the on-chip ADC as the source, the HTHT link with the active match had an open eye up to 18 Gbps (Fig. 10); full signal reconstruction was done up to 16 Gbps with the FPGA. Stable pattern with external PRBS was observed up to 22 Gbps.

The power consumption of two stages of the HTHT link with passive match is 6.3 mW each. Since power consumption at 4 K is at a high premium, careful measurements were done on the first (4 K) stage of the HTHT link with active match. Remarkably, the gain does not appreciably drop even when the supply voltage is decreased down to 150 mV, resulting in

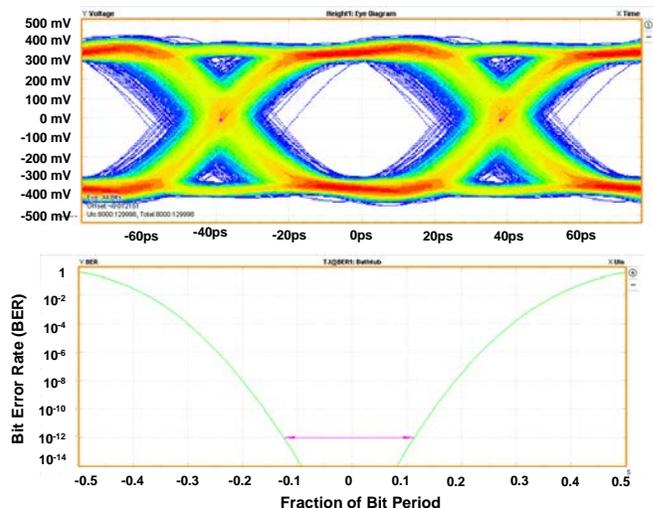


Fig. 9. An eye diagram and a BER “bathtub” curve of a 13 Gbps data link, comprising a differential SFQ/DC driver, a chain of two cryogenic SiGe amplifiers placed at two intermediate temperature stages, a room temperature commercial-off-the-shelf limiting amplifier and GTY receiver on Xilinx UltraScale FPGA, acquired on an Tektronix DSA72004B oscilloscope are shown.

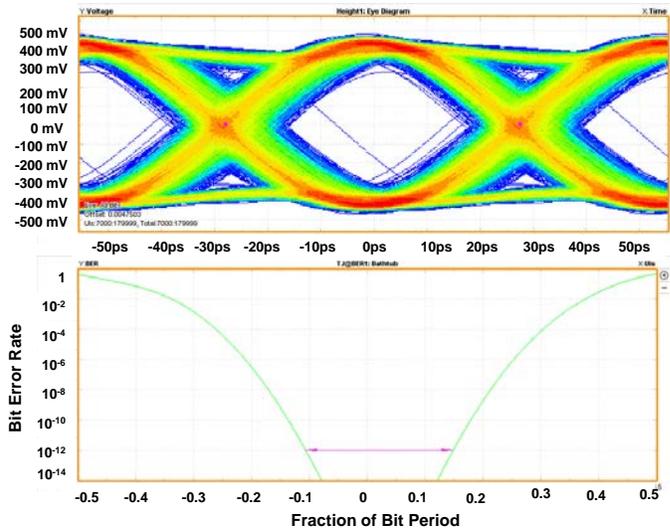


Fig. 10. An eye diagram and a BER “bathtub” curve of a 18 Gbps data link, comprising a differential SFQ/DC driver, a chain of three cryogenic SiGe amplifiers placed at three temperature stages, a room temperature commercial-off-the-shelf limiting amplifier and GTY receiver on Xilinx UltraScale FPGA, acquired on an Tektronix DSA72004B oscilloscope are shown.

power consumption of only $100 \mu\text{W}$. The link test data mentioned above was taken for $140 \mu\text{W}$. The second and the third stages consume 3.3 mW and 29.5 mW respectively.

A comparison of the four different semiconductor amplifier

TABLE III
 COMPARISON OF SEMICONDUCTOR AMPLIFIERS

Example	Data Rate (Gb/s)	4 K Power (mW)	18 K Power (mW)	40 K Power (mW)	300 K Power (mW)	
					Amplifier	Discriminator
RT	18				200	
Cryo-40 + RT	21			7.6	26.4	300
Cryo-18-40 + RT	30		6.3	6.3		300
Cryo-4-18-40 +RT	30	0.14	3.3	29.5		300

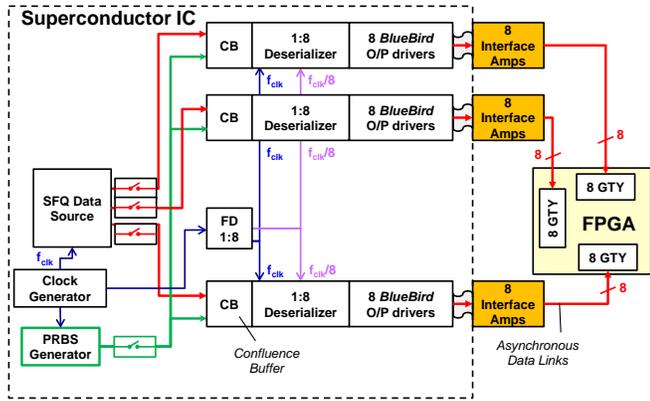


Fig. 11. Data from superconductor IC are transported to a room-temperature FPGA through a set of parallel Bluebird data links. In this example, the deserialization factor is 8. The GTY data receivers on the Xilinx FPGA are asynchronous. Therefore, an on-chip 7-bit pseudo-random bit sequence (PRBS-7) generator is used as a training sequence to synchronize the banks of 8 parallel channels. This configuration with three deserializer banks was demonstrated up to 120 Gbps aggregate data rate, limited by the SFQ data source. With a single deserializer bank, the maximum data transferred were 64 Gbps.

chains is listed in Table III in terms of their data rate and power consumption at different stages.

D. Data Link Receiver

State-of-the-art FPGAs offer several high-speed receivers (and transmitters). The Xilinx Ultrascale and Ultrascale+ FPGAs, which we have used, offer the GT family of asynchronous receivers; they do not need a synchronous clock aligned with the received data stream to operate. Reception of multi-channel asynchronous data using a bank of such receivers requires *channel bonding* [14], for which an on-chip pseudo-random bit sequence (PRBS) generator is used to produce the training sequence. PRBS sequences remain unaltered through the deserialization process and each deserialized PRBS sequence has a pre-determined delay of $2^B/N$, where 2^B-1 is the length of the PRBS and N is the deserialization ratio. As shown in Fig. 11, the output can be selected between the real data source and the PRBS generator through a pair of switches followed by a confluence buffer. Channel bonding is first performed by turning on the switch corresponding to the on-chip PRBS generator and running an alignment engine state machine on the FPGA. Once all the channels are aligned, the switches are toggled and the data flow from the on-chip source to the FPGA.

Once the multi-channel link is set up, it appears to be stable for several hours of testing during a day. Prolonged testing, spanning several days, is planned in the future. Resynchronization takes a fraction of a millisecond and could be automated.

Each GTY receiver on a Xilinx Virtex UltraScale+ VCU118 FPGA board, as used in our experiments, has the following rate-dependent power consumption values when configured for unencoded data: 242 mW for 10 Gbps, 337 mW for 20 Gbps, and 452 mW for 30 Gbps. These latest generation receivers are substantially lower power than the GTY receivers

on the previous generation (Xilinx Virtex UltraScale). For example, the corresponding values are: 399 mW for 10 Gbps, 549 mW for 20 Gbps, and 650 mW for 30 Gbps. There are other receiver alternatives that feature lower maximum data rate but consume less power. For example, an UltraScale GTH receiver consumes only 203 mW for 10 Gbps but operates only up to 16 Gbps with 289 mW.

III. COMPARISON OF DATA LINKS

Depending on the requirements of an application and system constraints, one can design a data link with different combinations of the components. In this section, we describe four examples of data links that have been implemented, compare and contrast their features in terms of data rates and power consumption, and discuss future directions for further development.

The four links, depicted in Fig. 1 are: (1) Bluebird driver with no cryogenic amplification connected with the stripline version 3 (Table II) to the interface amplifier unit (Fig. 5), (2) Ostrich driver with one stage of cryogenic amplification at 40 K with stripline connections to 4 K and room temperature, (3) Differential SFQ/DC driver with passive match and two gain stages at 18 K and 40 K with Ag-clad BeCu coaxial cables for inter-stage connections, and (4) Differential SFQ/DC driver with active match (Fig. 8) at 4 K and two more gain stages at 18 K and 40 K, also with Ag-clad BeCu coaxial cables. The data rates and power consumption at 4, 18, and 40 K are listed in Table III. The 10 Gbps link is the easiest to install in a cryocooled system and is the most mature. We operated this link with on-chip PRBS-7 generator for 30 minutes without any error, implying an error rate $<2 \times 10^{-13}$ with 97% confidence level.

Since the dominant contribution to the 4 K heat load is from the transmission line (TL), it is clear that one can optimize by inserting a cryogenic gain stage to compensate for higher attenuation while operating at a higher data rate. This is evident from the second example where the data rate is pushed to 20 Gbps by using a faster driver (Ostrich) on chip and a 40 K amplifier (Fig. 7). Whereas the Ostrich driver consumes more power than the Bluebird driver, it is a small fraction of the conduction heat load through the stripline ribbon cable. This HTHT approach of distributing the necessary amplification among multiple cryogenic stages can be advantageous, if

TABLE IV
 COMPARISON OF DATA LINK EXAMPLES

Example	Data Rate (Gb/s)	4K Heat Load (mW)			18K Heat Load (mW)			40K Heat Load (mW)		
		TL	Amp	Total	TL	Amp	Total	TL	Amp	Total
1	10	1.12	0.134	1.25				13.0		13.0
2	20	1.12	0.158	1.28				13.0	7.6	20.6
3	30	4.74	0.013	4.75	12.8	6.3	19.1	149.0	6.3	155.3
4	30	4.74	0.153	4.89	12.8	3.3	16.1	149.0	29.5	178.5

TABLE V
 COMPARISON OF SFQ-TO-FPGA DATA LINKS

Example	Data Rate (Gb/s)	4K Heat Load \times 500 (mW)	18K Heat Load \times 100 (mW)	40K Heat Load \times 40 (mW)	300K Heat Load (mW)	Total Equivalent Heat Load at 300 K (W)	FOM (pJ/bit) at 300 K
1	10	626	0	520	442	1.588	159
2	20	638	0	824	642	2.104	105
3	30	2377	1910	6212	752	11.251	375
4	30	2447	1610	7140	752	11.949	398

properly optimized.

From the perspective of minimizing the number of links, which is an increasingly important consideration for systems requiring greater data throughput, it is advantageous to increase data rate. Since FPGA data receivers up to and slightly above 30 Gbps are commercially available, it is highly desirable to establish such a data link. Since the bandwidth of such links is higher, so is their vulnerability to crosstalk and noise. The next two examples (#3 and #4 in Table III) represent fully differential data links to combat this vulnerability at the cost of doubling the transmission line heat load. In case of active match, we have added 13.4 μ W for the differential SFQ/DC driver and 140 μ W for the SiGe amplifier.

As currently implemented, both 30 Gbps links are completely dominated by the use of a pair of coaxial cables between the stages. Although convenient for proof-of-concept experimental work, use of semi-rigid coaxial cables is impractical for extension to multiple parallel data links. Properly designed and manufactured, stripline ribbons for inter-stage connections will lower heat conduction and make integration more practical, especially with a multi-channel chipset.

While the heat load due to metallic coaxial and stripline transmission lines dominate the currently implemented use cases, these are particularly useful for applications with high clock speed and throughput. As an example consider a data stream at 64 Gbps that needs to be transported from a superconductor chip to room-temperature. Until, the development of the Bluebird data link (Example 1 in Table III), this required deserialization by a factor of 32. With the Bluebird link, the deserialization ratio drops to 8. With the higher-speed data links, one can drop the further decrease the deserialization ratio to greatly simplifying many aspects of the system design. The advantages include: (a) reduced complexity and power consumption on chip, (b) simpler cryogenic wiring, and (c) easier interface to data acquisition system (e.g. a state-of-the-art FPGA with built-in fast data receivers). Reducing the number of output lines is particularly advantageous to designs that are limited by the number of input/output pads, rather than circuit area, on a chip.

Proper accounting for heat loads at various temperature stages requires many assumptions regarding the cryogenic system design, heat loads due to other factors such as thermal radiation. However, it is possible to get an idea about the contribution of various components at different temperatures by making some reasonable estimates. In the following analysis,

we assume a large-scale cryocooled system that is much more efficient than the small-scale one. Heat loads at 4, 18, and 40 K are multiplied by 500 W/W, 100 W/W, and 40 W/W respectively [16]. A figure-of-merit (FOM) is derived by dividing the total equivalent heat load at room temperature by the data rate. In this calculation, all the individual component contributions, including the power consumption of the FPGA receiver (GTy on UltraScale+) are counted for completeness.

As system size scales up, it will be necessary to address the conductive heat load. There are two well-known alternatives to metallic transmission lines. The first is the use of high-temperature superconductor (HTS) transmission lines [15] and the second is the use of photonic links [16]. Manufacturing difficulties with HTS transmission lines have so far prevented their use as low-loss wideband transmission lines with low thermal conduction. For an optical link, while the heat load for the fiber-optic cable is negligible, one must account for the power consumed by the electrical-to-optical converter as well as the optical-to-electrical converter.

For the 30 Gbps electrical data links (Table V), the largest contribution is due to the coaxial cables conducting heat down from room temperature to the 40 K stage. Replacing with a pair of striplines will improve the FOM by nearly a factor of 2. One can also convert the signal to the optical domain at the 40 K stage, where the swing is several hundreds of millivolts following amplification. Conceivably, electrical-to-optical conversion may be possible at lower temperatures, which is a subject of future research [16]. An advantage of conversion to the optical domain is transmission on optical fibers which are immune to electromagnetic interference. All of the link examples can be improved by conversion of the 4-40 K transmission line to HTS.

IV. CONCLUSION

We have established the framework for end-to-end data link design by selecting components from four basic categories, starting with the superconductor on-chip driver that takes in SFQ data at 4 K and ending with the data receiver on an FPGA at room temperature. In between, one has to choose appropriate transmission media and any necessary amplification to complete the link. For any given application, one can optimize the link by spanning the temperature range with a mix of transmission and amplification technologies. It is necessary to calculate the end-to-end power consumption for proper comparison between different options. Four proven data link configurations, with data rates ranging from 10-30 Gbps have been analyzed. With the simplest and the slowest of these, we have demonstrated 120 Gbps data transfer from an SFQ circuit to room temperature. No longer limited by data transport, more complex digital circuits can be realized. This is particularly valuable for the smaller-scale cryogenic systems where one can trade power consumption and functionality. Scaling up to larger scale would benefit from further power optimization of the link.

ACKNOWLEDGMENT

We are grateful for the invaluable technical direction provided by Dr. Deborah Van Vechten on all aspects of digital data links.

REFERENCES

- [1] D. Gupta, D. E. Kirichenko, V. V. Dotsenko, R. Miller, S. Sarwana, A. Talalaevskii, J. Delmas, R. J. Webber, S. Govorkov, A. F. Kirichenko, I. V. Vernik, and J. Tang, "Modular, multi-function digital-RF receiver systems," *IEEE Trans. Appl. Supercond.*, vol.21, pp.883-890, June 2011.
- [2] D. S. Holmes, A. L. Ripple, and M. A. Manheimer, "Energy-efficient superconducting computing—power budgets and requirements," *IEEE Trans. Appl. Supercond.*, vol.23, no. 3, June 2013, Art. No. 1701610.
- [3] M. Suzuki et al., "An interface circuit for a Josephson-CMOS hybrid digital system," *IEEE Trans. Appl. Supercond.*, vol. 9, no. 2, pp. 3314–3317, Jun. 1999.
- [4] D. F. Schneider, J. C. Lin, S. V. Polonsky, V. K. Semenov, and C. A. Hamilton, "Broadband interfacing of superconducting digital systems to room temperature electronics," *IEEE Trans. Appl. Supercond.*, vol. 5, no. 2, pp. 3152–3155, Jun. 1995.
- [5] Q. P. Herr, D. L. Miller, A. A. Pesetski, and J. X. Przybysz, "Inductive isolation in stacked squid amplifiers," *IEEE Trans. Appl. Supercond.*, vol. 17, no. 2, pp. 565–568, Jun. 2007.
- [6] A. Inamdar, S. Rylov, S. Sarwana, and D. Gupta, "Superconducting Switching Amplifiers for High Speed Digital Data Links," *IEEE Trans. Appl. Supercond.*, vol. 19, no. 2, pp. 1020 - 1027, June 2009.
- [7] T. Ortlepp, S. Wuensch, M. Schubert, P. Febvre, B. Ebert, J. Kunert, E. Crocoll, H.-G. Meyer, M. Siegel, and F. H. Uhlmann, "Superconductor-to-Semiconductor Interface Circuit for High Data Rates," *IEEE Trans. Appl. Supercond.*, vol. 19, pp. 28-34, February 2009.
- [8] T. Ortlepp, S. R. Whiteley, L. Zheng, X. Meng, and T. Van Duzer, "High-speed hybrid superconductor-to-semiconductor interface circuit with ultra-low power consumption," *IEEE Trans. Appl. Supercond.*, vol. 23, no. 3, p. 1400104, Jun. 2013.
- [9] V.V. Dotsenko, et al., "Integrated Electronics Testbed (ICE-T) for Evaluation of Superconductor and Cryo-Semiconductor Integrated Circuits", in Proc. Int. Cryogenics Eng. Conf. Int. Cryogenic Mater. Conf., 2017.
- [10] S. K. Tolpygo et al., "Advanced fabrication processes for superconducting very large scale integrated circuits", *IEEE Trans. Appl. Supercond.*, vol. 26, no. 3, Apr. 2016.
- [11] Prasana Ravindran, Su-Wei Chang, Deepnarayan Gupta, Amol Inamdar, Vladimir Dotsenko, Saad M. Sarwana, and Joseph C. Bardin, "Power-optimized Temperature-distributed Digital Data Link," *IEEE Trans. Applied Superconductivity*, vol. 25, no. 3, June 2015
- [12] D. Gupta, J. C. Bardin, A. Inamdar, A. Dayalu, S. Sarwana, P. Ravindran, Su-Wei Chang, A. H. Coskun, and M. Ghadiri Sadrabadi, "Low-Power High-Speed Hybrid Temperature Heterogeneous Technology Digital Data Link," *IEEE Trans. Appl. Supercond.*, vol. 23, no. 3, pp. 1701806, June 2013.
- [13] P. Ravindran, Su-Wei Chang, Wei-Ting Wong, S. M. Sarwana, V. Dotsenko, J. Tang, S. Ruotolo, D. Gupta, and Joseph C. Bardin, "Energy Efficient Digital Data Link", *IEEE Trans. Applied Superconductivity*, December 2016.
- [14] A. E. Lehmann et al., "Embedded RSFQ Pseudorandom Binary Sequence Generator for Multichannel High-Speed Digital Data Link Testing and Synchronization", in *IEEE Trans. Appl. Supercond.*, vol. 27, no. 4, pp. 1-6, June 2017, Art no. 1301806.
- [15] A. H. Yonamine et al., "Properties of Individual YBCO Layers in a Two-Layered Design for Energy-Efficient Digital Data Cables," *2013 IEEE 14th International Superconductive Electronics Conference (ISEC)*, Cambridge, MA, 2013, pp. 1-3.
- [16] IARPA SuperCables Broad Agency Announcement, Available: <https://www.iarpa.gov/index.php/research-programs/supercables/supercables-baa>, Accessed on Oct. 25, 2018.
- [17] D. Gupta, T. V. Filippov, A. F. Kirichenko, D. E. Kirichenko, I. V. Vernik, A. Sahu, S. Sarwana, P. Shevchenko, A. Talalaevskii, and O. A. Mukhanov, "Digital Channelizing Radio Frequency Receiver," *IEEE Trans. Appl. Supercond.*, vol. 17, no. 2, pp. 430-437, June 2007