Secondary Side-Channel Wireline Communication Using Transmitter Clock Frequency Modulation

Yi Fan Zhang, Joshua Liang, Shayan Shahramian, Behzad Dehlaghi, Ryan Bespalko, Michael O'Farrel, Dustin Dunwell, Davide Tonietto, and Anthony Chan Carusone

***Abstract*—A secondary communication link, or side-channel, is proposed, which uses binary frequency shift keying to modulate the frequency of a high-speed wireline transmitter clock. This side-channel data is then received using the corresponding receiver's conventional clock and data recovery (CDR) circuit. An analysis of the CDR loop parameters demonstrates that, within certain limits, the side-channel does not impact the signal integrity of the primary high-speed data. Measurements were made on a side-channel prototype using a 56-Gb/s PAM-4 (half-rate 14-GHz clock) transceiver in 7-nm FinFET technology through a 15-dB loss channel. Over 10^4 side-channel bits were transmitted at 50 kb/s, and the side-channel remained error-free with no visible impact on the high-speed link.**

***Index Terms*—Back-channel, clock and data recovery (CDR), frequency modulation, phase locked loop (PLL), wireline communication.**

I. Introduction

In high-speed wireline communication links, limited channel bandwidth requires equalization in the transmitter (TX) and receiver (RX) to compensate for frequency-dependent loss [1]. Typically, the RX would employ adaptation algorithms to optimize their highly programmable equalizers to minimize bit error rate (BER) and power during operation [2], [3]. Highly programmable TXs are also available; however, without channel information, they can only be initialized to nominal settings at start-up, and cannot be dynamically adapted during operation [4].

TX adaptation can enable higher data rates while keeping power low. For example, accurately setting the TX equalizer in [4] allowed the use of a simplified and very low-power RX architecture in [5]. To facilitate this, a secondary communication link, i.e., a back-channel, would be required to permit the RX to transmit control bits embedding information, such as ISI, channel response, BER, process variation, and temperature back to the TX. While this can be implemented with another physical channel [6], the cost of increasing pin count is very high. Thus, the ability to communicate two sets of independent data over the same channel is needed.

Previously presented back-channel techniques suffer from various limitations. In [7], the common-mode voltage of a differential high-speed signal is modulated to transmit information for bidirectional communication. However, this technique limits the swing of the differential signal, which, at sub-1-V supply voltages, can impair signal integrity. Another method is to use frequency division multiplexing to introduce a secondary channel at baseband by AC coupling the primary channel [8], [9]. However, this method requires additional passive components and/or power-hungry radio-frequency topologies, which admits cross-talk between the primary and secondary data. Finally, in [10], a secondary channel is created by modulating the TX clock phase to transmit additional data. While this method alleviates the drawbacks in [7] and [8], it requires injecting a large phase shift outside the RX clock and data recovery (CDR) circuit's tracking bandwidth for the secondary data to be visible. This process would introduce phase errors that are untrackable by the CDR, which degrades jitter tolerance.

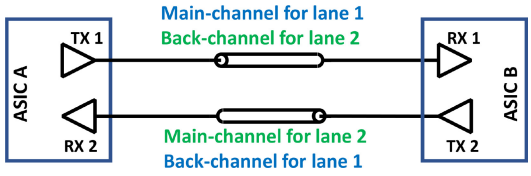

Fig. 1. Application of the proposed side-channel for bidirectional communications.

In this letter, a secondary side-channel architecture is proposed, which transmits data by modulating the TX clock frequency using binary frequency shift keying (BFSK) and recovers the data using the RX's CDR to track the frequency offsets. Unlike [7] and [8], this architecture does not impact signal swing, as it modulates the data edges rather than its voltage level, and has minimal cross-talk between the high-speed channel and the secondary channel. Whereas in [10], clock phase modulation introduced additional jitter to the high-speed link, we show in Section III that by respecting specific limits of the CDR, the BFSK side-channel will not impact on the high-speed data's signal integrity.

While the proposed side-channel is not bidirectional on its own, it can be implemented as a back-channel using a transceiver pair, as illustrated in Fig. 1, where one serial link's side-channel carries the back-channel data for the other link. Although inserting control bits directly into the high-speed data would be an option in this setup, it is generally impractical to do so because it would require new communication protocols. Without modifying the primary data, the proposed side-channel architecture creates a low-speed and nondisruptive secondary channel for communication between the high-speed physical-layer TX and RX.

II. Side-Channel Architecture

The system architecture of the side-channel is shown in Fig. 2. The TX clock frequency can be modulated in various ways. Using the TX frequency synthesizer (PLL), either the clock divider ratio or the PLL's reference clock frequency can be modulated; the latter method is prototyped for measurement in Section IV. Frequency modulation on the PLL can be implemented in firmware without additional

Manuscript received October 21, 2019; revised November 26, 2019; accepted December 5, 2019. Date of publication December 16, 2019; date of current version January 8, 2020. This article was approved by Associate Editor Samuel Palermo. This work was supported by the Natural Sciences and Engineering Research Council of Canada under Grant CRPDJ 505827-16. *(Corresponding author: Yi Fan Zhang.)*

Yi Fan Zhang and Anthony Chan Carusone are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada (e-mail: danny.zhang@isl.utoronto.ca).

Joshua Liang, Shayan Shahramian, Behzad Dehlaghi, Ryan Bespalko, Michael O'Farrel, and Dustin Dunwell are with Huawei Technologies, Markham, ON L3R5A4, Canada.

Davide Tonietto is with Huawei Technologies, Kanata, ON K2K 3J1, Canada.

Digital Object Identifier 10.1109/LSSC.2019.2959779

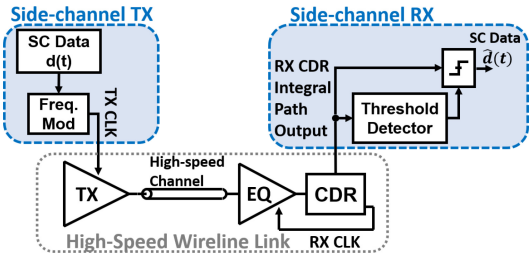

Fig. 2. Frequency-modulated side-channel architecture.

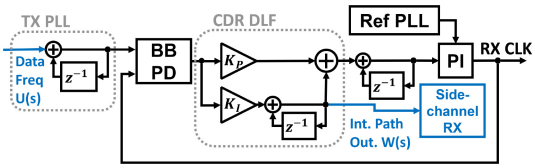

Fig. 3. Phase model of the BFSK side-channel showing the details of the digital Bang–Bang CDR.

high-speed circuitry on virtually any high-speed link. However, this implementation modulates all links sharing the same PLL, which, in multichannel applications, would require communication protocols to identify individual lanes using serial numbers. Alternatively, frequency modulation can also be implemented using a phase interpolator (PI), if present, within the TX [11] by rotating the phase code of the PI at a constant rate. This method allows for independent frequency modulation of each link.

The side-channel bitstream is demodulated using the RX's CDR. Fig. 3 shows the phase model of a digital Bang–Bang CDR widely used in high-speed links. When the CDR locks, the integral path output of its digital loop filter (DLF) will settle to a digital value corresponding to the TX frequency. This value can be sampled by the side-channel RX to recover the side-channel data. It will be shown in Section IV that a digital slicer can be adequate to demodulate the BFSK data, without requiring additional high-speed circuitry.

The raw values of the CDR integral path output corresponding to side-channel's 1s and 0s will vary based on the high-speed data rate, the nominal frequency offset between the TX and RX clocks, the side-channel BFSK amplitude, and the CDR architecture. A "threshold detector" can be incorporated to determine the threshold to demodulate the side-channel data by computing the mean of the integral path output over the previous n-bits of the side-channel data.

III. Side-Channel Design Parameters

Two key design parameters of the side-channel are its data rate and BFSK amplitude. To ensure the side-channel is functional and does not impair the high-speed link, we must ensure all frequency changes are trackable by the high-speed RX's CDR.

Fig. 3 shows the side-channel data path through a digital Bang–Bang CDR. The side-channel system transfer function is shown in (1), where K_{PD} is the linearized phase detector (BB PD) gain, K_{PI} is the PI gain, K_P is the proportional path gain, K_I is the integral path gain, and T_S is the simulation time-step of the CDR [12], [13][1]

$$\frac{W(s)}{U(s)} \approx \frac{(K_{\mathrm{PD}}K_I)}{s^2 + \left(\frac{K_{\mathrm{PD}}K_{\mathrm{PI}}K_P}{T_S}\right)s + \left(\frac{K_{\mathrm{PD}}K_{\mathrm{PI}}K_I}{T_S^2}\right)}. \tag{1}$$

[1]The transfer function is approximated using the inverse-bilinear transform, $z^{-1} = [(1-(sT_S)/2)/(1+(sT_S)/2)]$, assuming $sT_S \ll 1$ [13].

TABLE I
Side-Channel Simulation Parameters, Approximated Based on the PAM-4 56 Gb/s Test Chip in Section IV

f_{ref}	T_S	K_{PD}	K_{PI}
14 GHz	2 ps	$\pi/4$	$2\pi/128$
K_P	K_I	Δf_{max}	$\lvert\delta f/\delta t\rvert_{max}$
2.5×10^{-4}	7.5×10^{-9}	300 ppm	70 MHz/μs
side-channel f_{-3dB}		side-channel data rate	
2 MHz		1 Mb/s	

This results in a 3-dB bandwidth

$$f_{-3\mathrm{dB}} \approx \frac{K_I}{2\pi K_P T_S}. \tag{2}$$

In order for the side-channel data to be recoverable, its data rate must be within the bandwidth in (2), which is normally lower than the bandwidth of the CDR.

The BFSK amplitude, Δf, which can be expressed in parts per million (ppm), must be greater than the self-generated jitter of the high-speed link, but within the tracking range of the CDR. The clock phase changes introduced by the BFSK modulation can be tracked by the proportional path of the CDR DLF up to a maximum frequency offset of

$$\Delta f_{\mathrm{max}} = \frac{K_P K_{\mathrm{PD}} K_{\mathrm{PI}}}{f_{\mathrm{ref}} T_S} \tag{3}$$

based off of the first order loop dynamics of the CDR, where f_{ref} is the reference clock frequency of the CDR [14].

If higher BFSK amplitudes are desired, the integral path can track additional frequency offsets that are greater than Δf_{max} as long as the rate of change in frequency, or frequency modulation slope, is low. Specifically, the slope of the rising and falling edges of the side-channel BFSK data must be below

$$\left|\frac{\delta f}{\delta t}\right|_{\mathrm{max}} = \frac{K_I K_{\mathrm{PD}} K_{\mathrm{PI}}}{2T_S} \tag{4}$$

which the maximum rate of change of the integral path in a second-order Bang–Bang CDR loop [14].

These limits are demonstrated in Fig. 4, which shows the behavioural simulation results of a side-channel link using the parameters listed in Table I. For BFSK amplitudes below $\Delta f_{\mathrm{max}} = 300$ ppm, the peak phase error of the high-speed link in Fig. 4(a) is close to zero regardless of the frequency modulation slope, which, in turn, does not cause any visible jitter tolerance degradation as shown in Fig. 4(b). Conversely, for BFSK amplitudes above Δf_{max}, the side-channel becomes limited by its frequency modulation slope from (4). As shown in Figs. 4(a) and (c), as long as the frequency modulation slope stays below $|\delta f/\delta t|_{\mathrm{max}} = 70$ MHz/μs, it will not cause additional phase error or jitter tolerance degradation.

Therefore, there are two ways to ensure that the side-channel will not disrupt the high-speed link:

1) ensure the side-channel BFSK amplitude is within the tracking range of the proportional path, Δf_{max};
2) ensure that the BFSK modulation slope, which also affects the data rate of the side-channel, is within the range of the integral path, $|\delta f/\delta t|_{\mathrm{max}}$.

Admittedly, since CDR bandwidths are typically in the range of a few MHz, the side-channel is restricted to low data-rates to ensure no impact on the high-speed data signal integrity. However, this low data-rate serves an essential purpose: enabling communication between the physical-layer TX and RX.

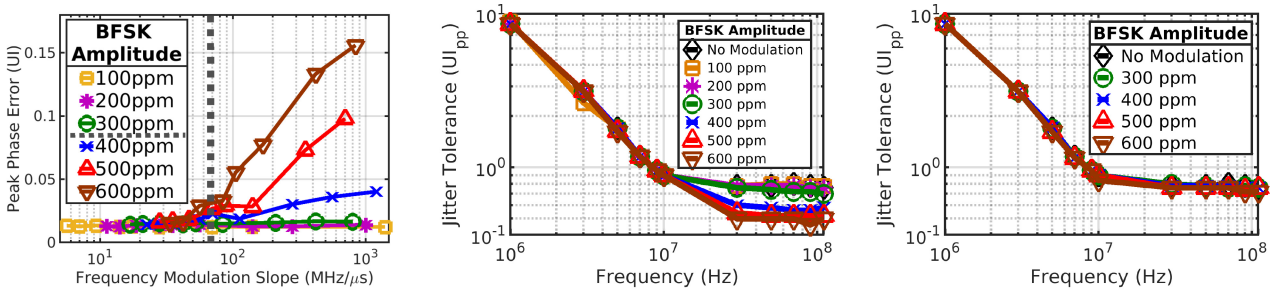

Fig. 4. Simulation results of various side-channel BFSK amplitudes using parameters presented in Table I. (a) Peak phase error versus frequency modulation slope. (b) Jitter tolerance of the CDR with frequency modulation slopes above 10^3 MHz/μs. (c) Jitter tolerance of the CDR with 40-MHz/μs side-channel frequency modulation slope.

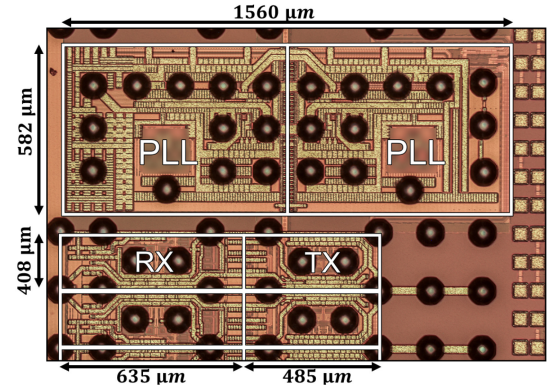

Fig. 5. 7-nm FinFET transceiver micrograph.

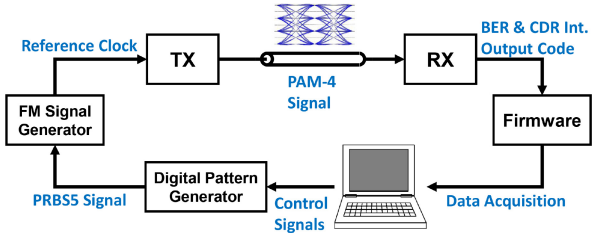

Fig. 6. Side-channel measurement test setup.

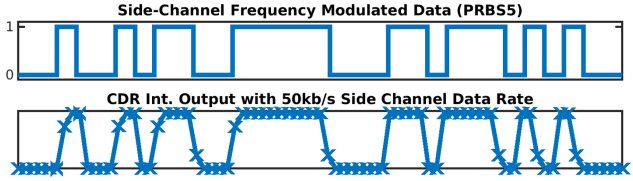

Fig. 7. Transmitted and received 50-ppm, 50 kb/s side-channel data using PRBS5.

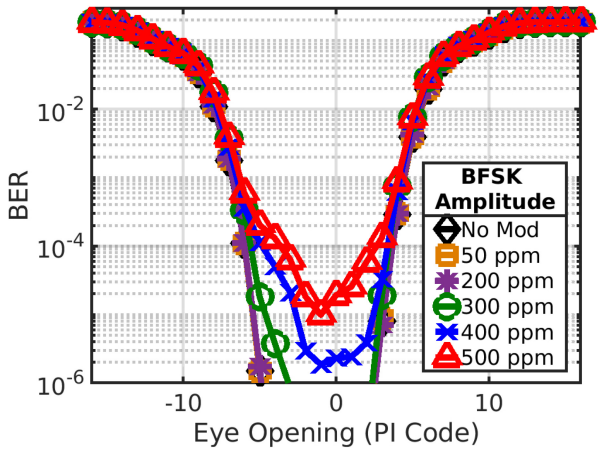

Fig. 8. Bathtub curves for the high-speed data for different BFSK amplitudes at 50 kb/s, over 10^4 side-channel bits.

IV. Measurement Results

A prototype of the proposed side-channel link is tested on a 7-nm FinFET 56 Gb/s PAM-4 transceiver from [2]. Fig. 5 shows the chip's die photo, and Fig. 6 shows the test setup. A digital pattern generator is used to generate a pseudo-random bit sequence (PRBS) of length $(2^5 - 1)$, which modulates an external reference clock signal generated by a Keysight N5172B FM signal generator. This frequency-modulated reference clock signal is connected to the PLL of the high-speed TX to drive a PAM-4 TX at 14 GHz, transmitting a PRBS31 sequence at 56 Gb/s. The high-speed channel has approximately 15-dB insertion loss at 14 GHz including cables, connectors, an ISI board, and an evaluation board. The high-speed data's BER and the RX CDR integral path output code are recorded in software.

Fig. 7 shows the measured integral path output in comparison with the transmitted 50-kb/s PRBS5 data pattern at a BFSK amplitude of 50 ppm. It can be observed that, even with a small BFSK amplitude of 50 ppm, the bit pattern is clearly distinguishable from the raw CDR integral path output code. The side-channel bitstream is demodulated by subsampling the integral path output (at a fixed subsampling ratio of the recovered high-speed clock) and comparing it to a fixed threshold in software. Over 10^4 consecutive bits were communicated through the side-channel without error. Unfortunately, in our prototype, the side-channel data-rate was limited by the digital sampling rate of the data acquisition firmware to fewer than 100 kb/s, but higher data-rates can be enabled with specific side-channel logic embedded in the chip.

Fig. 8 shows the measured bathtub curve of the high-speed signal for various side-channel BFSK amplitudes. The frequency modulation slope of the transmitted side-channel data was not controlled, so very fast (i.e., faster than $|\delta f/\delta t|_{\max}$) data transitions are assumed. The measured bathtub curve shows that below 300 ppm BFSK amplitude, the horizontal eye opening remains undisturbed, showing no detriment to the high-speed channel's signal integrity.

Fig. 9 compares measurement against simulations and theory by sweeping the CDR's proportional path gain constant, K_P, and computing $\Delta f_{\max}$. The theoretical $\Delta f_{\max}$ is calculated using (3), and

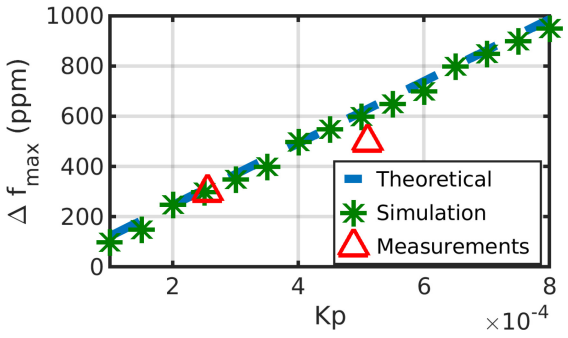

Fig. 9. $\Delta f_{\max}$ versus CDR proportional path constant K_P.

the simulations are based on the parameters in Table I. The simulated $\Delta f_{\max}$ is identified as the maximum BFSK amplitude for which the peak phase error is below 0.04 UI (Fig. 4), and the measured $\Delta f_{\max}$ is identified as the maximum BFSK amplitude for which the BER is below 10^{-6} (Fig. 8). The data points match well, with a small discrepancy of around 50 ppm in the measurements at $K_P = 5 \times 10^{-4}$ due to factors, such as the channel insertion loss and the nonlinearity of the Bang–Bang phase detector unaccounted for in simulation. However, this can still allow us to infer the accuracy of the models in Section III in demonstrating the limitations of the side-channel.

V. Conclusion

A side-channel architecture is presented which permits the simultaneous transmission of a secondary data stream over an existing high-speed serial link for low-speed communication between the physical-layer TX and RX. This architecture uses BFSK to modulate the TX clock frequency and recovers the data using an RX's CDR without any additional high-speed hardware. Analysis and simulations place limits on the side-channel's parameters to ensure no impact on the high-speed data's signal integrity. The technique was demonstrated on a 7-nm FinFET 56-Gb/s quarter-rate PAM-4 transceiver. After transmitting over 10^4 side-channel bits at 50 kb/s using 50-ppm BFSK amplitude, no error or disturbance on the high-speed data was observed.

References

[1] C. Loi *et al.*, "6.5 A 400Gb/s transceiver for PAM-4 optical direct-detect application in 16nm FinFET," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2019, pp. 120–121.

[2] S. Shahramian *et al.*, "30.5 a 1.41pJ/b 56Gb/s PAM-4 wireline receiver employing enhanced pattern utilization CDR and genetic adaptation algorithms in 7nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2019, pp. 482–483.

[3] M.-A. LaCroix *et al.*, "6.2 A 60Gb/s PAM-4 ADC-DSP transceiver in 7nm CMOS with SNR-based adaptive power scaling achieving 6.9pJ/b at 32dB loss," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2019, pp. 114–115.

[4] C. Menolfi *et al.*, "A 112Gb/S 2.6pJ/b 8-tap FFE PAM-4 SST TX in 14nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2018, pp. 104–105.

[5] A. Cevrero *et al.*, "6.1 A 100Gb/s 1.1pJ/b PAM-4 RX with dual-mode 1-tap PAM-4 / 3-tap NRZ speculative DFE in 14nm CMOS FinFET," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2019, pp. 112–113.

[6] J. T. Stonick, G.-Y. Wei, J. L. Sonntag, and D. K. Weinlader, "An adaptive PAM-4 5-Gb/s backplane transceiver in 0.25-μm CMOS," in *Proc. IEEE Custom Integr. Circuits Conf.*, May 2002, pp. 363–366.

[7] V. Stojanović *et al.*, "Autonomous dual-mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 1012–1026, Apr. 2005.

[8] M. Bichan, M. Hossain, and A. C. Carusone, "Frequency-division bidirectional communication over chip-to-chip channels," *IEEE Trans. Adv. Packag.*, vol. 32, no. 2, pp. 298–305, May 2009.

[9] J. Ko, J. Kim, Z. Xu, Q. Gu, C. Chien, and M. F. Chang, "An RF/baseband FDMA-interconnect transceiver for reconfigurable multiple access chip-to-chip communication," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2005, pp. 338–339.

[10] X. Wang, T. Liu, S. Guo, M. A. Thornton, and P. Gui, "A 2.56-Gb/s serial wireline transceiver that supports an auxiliary channel in 65-nm CMOS," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, to be published.

[11] P.-C. Chiang, H.-W. Hung, H.-Y. Chu, G.-S. Chen, and J. Lee, "2.3 60Gb/s NRZ and PAM4 transmitters for 400GbE in 65nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2014, pp. 42–43.

[12] J. Lee, K. S. Kundert, and B. Razavi, "Analysis and modeling of bang-bang clock and data recovery circuits," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1571–1580, Sep. 2004.

[13] F. Gardner, *Phaselock Techniques*, 3rd ed. Hoboken, NJ, USA: Wiley, 2005.

[14] R. C. Walker, "Designing Bang–Bang PLL's for clock and data recovery in serial data transmission systems," in *Phase-Locking in High-Performance Systems*, B. Razavi, Ed., Hoboken, NJ, USA: Wiley, 2003, pp. 34–45.