

KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

METHODS AND DEVICES FOR SELF-INTERFERENCE CANCELATION**Technical Field**

[0001] Various embodiments relate generally to methods and devices for self-interference cancelation.

Background

[0002] Interference cancelation is an important component of radio communication systems, in particular for self-interference cancelation for multi-radio coexistence, frequency division duplexing, and full-duplex radio communication systems. As such radio communication systems may transmit and receive simultaneously, interference may leak from the transmit chain to the receive chain and consequently produce corruptive self-interference signal in signals received at the received chain. While certain largely 'passive' solutions such as special duplexing circuitry may be effective in sufficiently isolating the receive chain from the transmit chain, these solutions may be expensive and thus undesirable for many manufacturers.

[0003] Digital self-interference cancelation may thus offer a lower cost alternative solution. In such self-interference cancelation solutions, one or more adaptive filters may be utilized to model the leakage path from the transmit chain to the receive chain. Accordingly, assuming an accurate model the adaptive filters may be able to produce estimated interference signals from original transmit signals. The receive chain may then subtract these estimated interference signals from received signals, thus canceling the self-interference from the received signals and producing a clean signal that may be largely free of residual self-interference.

Brief Description of the Drawings

[0004] In the drawings, like reference characters generally refer to the same parts throughout the different views. The drawings are not necessarily to scale, emphasis instead generally being placed upon illustrating the principles of the invention. In the following description, various embodiments of the invention are described with reference to the following drawings, in which:

FIG. 1 shows a diagram of communication circuitry;

FIG. 2 shows a flow chart illustrating a self-cancellation interference procedure;

FIG. 3 shows a flow chart illustrating update of a kernel dimension filter;

FIG. 4 shows a flow chart illustrating update of a delay tap dimension filter;

FIG. 5 shows a graphical depiction of a matrix storage scheme;

FIG. 6 shows a diagram of processing circuitry including a signal path and an adaptation engine;

FIG. 7A-7B show a decoupled kernel dimension input signal calculation and decoupled delay tap dimension input signal calculation;

FIG. 8 shows a resource table and hardware pipeline clock schedule;

FIG. 9 shows a diagram of a decoupled kernel dimension input signal calculation circuit;

FIG. 10 shows a diagram of a decoupled delay tap dimension input signal calculation circuit;

FIG. 11 shows diagrams of a correlation and cross-correlation update circuit;

FIG. 12 shows a matrix memory indexing scheme;

FIG. 13 shows a diagram of a DCD circuit;

FIG. 14 shows a diagram of a maximum cross-correlation identification circuit;

FIG. 15 shows a diagram of a memory-based power amplifier estimation;

FIG. 16 shows a first method of performing interference cancelation;

FIG. 17 shows a second method of performing interference cancelation;

FIG. 18 shows a diagram of a bilinear interference cancelation system for a MIMO architecture;

FIG. 19 shows a diagram illustrating adaptation resource sharing for a MIMO architecture;

FIG. 20 shows a diagram of a bilinear interference cancelation architecture for MIMO with joint FIR filter estimation;

FIG. 21 shows a third method of performing interference cancelation; and

FIG. 22 shows a fourth method of performing interference cancelation.

Description

[0005] The following detailed description refers to the accompanying drawings that show, by way of illustration, specific details and embodiments in which the invention may be practiced.

[0006] The word “exemplary” is used herein to mean “serving as an example, instance, or illustration”. Any embodiment or design described herein as “exemplary” is not necessarily to be construed as preferred or advantageous over other embodiments or designs.

[0007] The words “plurality” and “multiple” in the description and the claims expressly refer to a quantity greater than one. The terms “group (of)”, “set [of]”, “collection (of)”, “series (of)”, “sequence (of)”, “grouping (of)”, etc., and the like in the description and in the claims, if any, refer to a quantity equal to or greater than one, i.e. one or more. Any term expressed in plural form that does not expressly state “plurality” or “multiple” refers to a quantity equal to or greater than one. The terms “proper subset”, “reduced subset”, and “lesser

subset” refer to a subset of a set that is not equal to the set, i.e. a subset of a set that contains less elements than the set.

[0008] It is appreciated that any vector and/or matrix notation utilized herein is exemplary in nature and is employed solely for purposes of explanation. Accordingly, it is understood that the approaches detailed in this disclosure are not limited to being implemented solely using vectors and/or matrices, and that the associated processes and computations may be equivalently performed with respect to sets, sequences, groups, etc., of data, observations, information, signals, samples, symbols, elements, etc. Furthermore, it is appreciated that references to a “vector” may refer to a vector of any size or orientation, e.g. including a 1×1 vector (e.g. a scalar), a $1 \times M$ vector (e.g. a row vector), and an $M \times 1$ vector (e.g. a column vector). Similarly, it is appreciated that references to a “matrix” may refer to matrix of any size or orientation, e.g. including a 1×1 matrix (e.g. a scalar), a $1 \times M$ matrix (e.g. a row vector), and an $M \times 1$ matrix (e.g. a column vector).

[0009] A “circuit” as used herein is understood as any kind of logic-implementing entity, which may include special-purpose hardware or a processor executing software. A circuit may thus be an analog circuit, digital circuit, mixed-signal circuit, logic circuit, processor, microprocessor, Central Processing Unit (CPU), Graphics Processing Unit (GPU), Digital Signal Processor (DSP), Field Programmable Gate Array (FPGA), integrated circuit, Application Specific Integrated Circuit (ASIC), etc., or any combination thereof. Any other kind of implementation of the respective functions which will be described below in further detail may also be understood as a “circuit”. It is understood that any two (or more) of the circuits detailed herein may be realized as a single circuit with substantially equivalent functionality, and conversely that any single circuit detailed herein may be realized as two (or more) separate circuits with substantially equivalent functionality. Additionally, references to a “circuit” may refer to two or more circuits that collectively form a single circuit. The term

“circuit arrangement” may refer to a single circuit, a collection of circuits, and/or an electronic device composed of one or more circuits.

[0010] As used herein, “memory” may be understood as a non-transitory computer-readable medium in which data or information can be stored for retrieval. References to “memory” included herein may thus be understood as referring to volatile or non-volatile memory, including random access memory (RAM), read-only memory (ROM), flash memory, solid-state storage, magnetic tape, hard disk drive, optical drive, etc., or any combination thereof. Furthermore, it is appreciated that registers, shift registers, processor registers, data buffers, etc., are also embraced herein by the term memory. It is appreciated that a single component referred to as “memory” or “a memory” may be composed of more than one different type of memory, and thus may refer to a collective component comprising one or more types of memory. It is readily understood that any single memory component may be separated into multiple collectively equivalent memory components, and vice versa. Furthermore, while memory may be depicted as separate from one or more other components (such as in the drawings), it is understood that memory may be integrated within another component, such as on a common integrated chip.

[0011] The term “base station” used in reference to an access point of a mobile communication network may be understood as a macro base station, micro base station, Node B, evolved NodeB (eNB), Home eNodeB, Remote Radio Head (RRH), relay point, etc. As used herein, a “cell” in the context of telecommunications may be understood as a sector served by a base station. Accordingly, a cell may be a set of geographically co-located antennas that correspond to a particular sectorization of a base station. A base station may thus serve one or more cells (or sectors), where each cell is characterized by a distinct communication channel. Furthermore, the term “cell” may be utilized to refer to any of a macrocell, microcell, femtocell, picocell, etc.

[0012] For purposes of this disclosure, radio communication technologies may be classified as one of a Short Range radio communication technology, Metropolitan Area System radio communication technology, or Cellular Wide Area radio communication technology. Short Range radio communication technologies include Bluetooth, WLAN (e.g. according to any IEEE 802.11 standard), and other similar radio communication technologies. Metropolitan Area System radio communication technologies include Worldwide Interoperability for Microwave Access (WiMax) (e.g. according to an IEEE 802.16 radio communication standard, e.g. WiMax fixed or WiMax mobile) and other similar radio communication technologies. Cellular Wide Area radio communication technologies include Global System for Mobile Communications (GSM), Code Division Multiple Access 2000 (CDMA2000), Universal Mobile Telecommunications System (UMTS), Long Term Evolution (LTE), General Packet Radio Service (GPRS), Evolution-Data Optimized (EV-DO), Enhanced Data Rates for GSM Evolution (EDGE), High Speed Packet Access (HSPA), etc., and other similar radio communication technologies. Cellular Wide Area radio communication technologies also include “small cells” of such technologies, such as microcells, femtocells, and picocells. Cellular Wide Area radio communication technologies may be generally referred to herein as “cellular” communication technologies. It is understood that exemplary scenarios detailed herein are demonstrative in nature, and accordingly may be similarly applied to various other mobile communication technologies, both existing and not yet formulated, particularly in cases where such mobile communication technologies share similar features as disclosed regarding the following examples.

[0013] The term “network” as utilized herein, e.g. in reference to a communication network such as a mobile communication network, encompasses both an access section of a network (e.g. a radio access network (RAN) section) and a core section of a network (e.g. a core network section). The term “radio idle mode” or “radio idle state” used herein in reference to

a mobile terminal refers to a radio control state in which the mobile terminal is not allocated at least one dedicated communication channel of a mobile communication network. The term “radio connected mode” or “radio connected state” used in reference to a mobile terminal refers to a radio control state in which the mobile terminal is allocated at least one dedicated uplink communication channel of a mobile communication network.

[0014] Unless explicitly specified, the term “transmit” encompasses both direct (point-to-point) and indirect transmission (via one or more intermediary points). Similarly, the term “receive” encompasses both direct and indirect reception. The term “communicate” encompasses one or both of transmitting and receiving, i.e. unidirectional or bidirectional communication in one or both of the incoming and outgoing directions. The term “calculate” encompass both ‘direct’ calculations via a mathematical expression/formula/relationship and ‘indirect’ calculations via lookup tables and other array indexing or searching operations.

[0015] Radio communication systems that involve simultaneous transmission and reception, including frequency-division duplexing, full duplexing, and multi-radio coexistence, commonly suffer from self-interference that is caused by a device’s own transmitter. As both frequency-division and full duplexing devices transmit and receive simultaneously, such self-interference arises when the transmit signal leaks onto the receive path in the transceiver. Due to the fact that transmit power is considerably higher than receive power, the transmit signal may dominate the received signals and thus lead to a high degree of corruption at the receiver.

[0016] The transmit signal involved in self-interference may conventionally leak from the transmit path to the receive path either through a duplexer and/or over the air. While duplexer leakage may be prevented in through specialized duplexing circuitry such components generally need to be very expensive circuitry in order to provide isolation that is sufficient to avoid crippling self-interference.

[0017] Self-interference cancelation has thus been recognized as a potentially more cost-effective solution. In self-interference cancelation schemes, adaptive filters may be used to ‘model’ the transmit path leakage in order to generate estimated interference signals. These estimated interference signals may then be provided to the receive path, which may subtract the estimated interference signal from the received signal. Depending on the accuracy of the employed adaptive filters, the receiver may be able to counter the leakage signal and thus produce a clean signal that is largely free of interference. Self-interference cancelation may thus relax the requirements for expensive duplexers and allow for most cost-efficient transceiver designs.

[0018] The adaptive filters used in self-interference cancelation techniques may aim to model the radio power amplifier employed in the transmit path. However, as such radio power amplifiers are generally non-linear, the approximation techniques required to generate accurate filter weights may be relatively complex. Accordingly, many self-interference cancelation approaches may model the power amplifier as a set of ‘kernels’, where each kernel is tailored to model a separate nonlinear component of the response of power amplifier response. The self-interference cancelation architecture may then derive a separate adaptive filter tailored to each kernel, where the each input sample for the power amplifier may be decomposed into a kernel sample for each kernel and applied to the corresponding adaptive filter. As such designs may conventionally utilize memory-based filters, each adaptive filter may be applied to each kernel sample in addition to the previous kernel samples, or ‘taps’, in order to generate a filter output. The sum of the adaptive filter output for each kernel may then give the estimated interference signal (where each adaptive filter models the interference contributed by the corresponding kernel), which may be subtracted from the received signal in order to provide a ‘clean’ signal that is conceivably free of self-interference.

[0019] The adaptation engine responsible for dynamically updating the adaptive filter weights may play a critical role in effective self-interference cancelation. These adaptation engines may predominantly examine the ‘clean’ signal in order to evaluate the effectiveness of the interference cancelation, i.e. whether there is any interference residue remaining in the clean signal after the estimated interference signal has been removed. Based on the efficacy of the interference cancelation, the adaptation engines may provide updates to the adaptive filter weights in order to continuously improve the cancelation performance.

[0020] Many adaptation schemes have been developed that offer varying tradeoffs between convergence and area/power. Accordingly, while approaches such as Recursive Least Squares (RLS) designs may offer extremely fast convergence at the expense of large area and power requirements for hardware implementations. Conversely, Least Mean Squares (LMS) may offer comparatively small area and power costs while suffering from poor convergence speed. Various other adaptation schemes and accompanying hardware designs have been proposed that similarly provide different convergence to area and power tradeoffs.

[0021] Many existing adaptation solutions may jointly update filter weights for a two-dimensional filter over both kernels and taps, i.e. by considering the PA response as varying over both taps and kernels. As previously indicated, power amplifier input signals may be decomposed into kernel samples according each of a set of predetermined kernels that collectively model the power amplifier. Each of the adaptive filters may then be tailored to estimate the interference associated with a respective kernel based on the current kernel sample and a number of past kernel samples. The adaptive filters may thus function as Finite Impulse Response (FIR) filters that hold the current and multiple previous input samples, apply a respective filter weight to each input sample, and sum the weighted outputs to obtain the filter output. Accordingly, for a given memory length M , i.e. M total taps, each adaptive filter may apply a respective filter weight (out of M total filter weights) to a respective one of

the M taps. Expressed mathematically, given an length- M input sample vector $\mathbf{X}(t) = [X_t, X_{t-1}, X_{t-2}, \dots, X_{t-M+1}]$ at a given time t , an FIR filter $\mathbf{W}$ with $M \times 1$ weight vector $\mathbf{W} = [w_0, w_1, \dots, w_{M-1}]^T$ may calculate a filter output $e(t)$ as

$$\begin{aligned} e(t) &= \mathbf{X}(t)\mathbf{W} \\ &= w_0X_t + w_1X_{t-1} + w_2X_{t-2} + \dots + w_{M-1}X_{t-M+1} \end{aligned} \quad (1)$$

where X_t is the current input sample and X_{t-1} to X_{t-M+1} are the $M - 1$ previous input samples and each of w_0 to w_{M-1} is a complex weight.

[0022] Accordingly, the adaptive filter may need to store the $M - 1$ previous input samples of $\mathbf{X}(t)$ in addition to the current sample and, for each new input sample, apply the M weights of $\mathbf{W}$ to $\mathbf{X}(t)$ in order to obtain filter output $e(t)$, which may represent the estimated self-interference, or ‘echo’. By selecting proper weights, the FIR filter $\mathbf{W}$ may be able to ‘filter’ $\mathbf{X}(t)$ in order to produce filter output $y(t)$ with certain desired characteristics, such as e.g. to approximate the kernel output in the context of self-interference cancelation.

[0023] As previously detailed, many self-interference cancelation schemes may employ kernelization to model the power amplifier as a set of kernels that each represent a non-linear component of the power amplifier (where the case described above regarding Equation (1) may thus be a non-kernelized or single kernel approach). Accordingly, such self-interference cancelation designs may employ a dedicated adaptive filter for each kernel, where each adaptive filter estimates the self-interference resulting from each respectively assigned kernel. Accordingly, the overall filter may be composed of multiple kernel filters that each model a respective kernel. In other words, for a model using K kernels and M taps, the filter $\mathbf{W}$ may be a $K \times M$ matrix where each row of $\mathbf{W}$ may contain the M weights for a respective kernel filter. Accordingly, $\mathbf{W}$ may be of the form

$$\mathbf{W} = \begin{bmatrix} w_t^{(1)} & w_{t-1}^{(1)} & \cdots & w_{t-M+1}^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ w_t^{(K)} & w_{t-1}^{(K)} & \cdots & w_{t-M+1}^{(K)} \end{bmatrix} \quad (2)$$

where each k -th row for $k = 1, 2, \dots, K$ contains the M filter weights for samples

$X_t^{(k)}, X_{t-1}^{(k)}, \dots, X_{t-M+1}^{(k)}$ of the k -th kernel.

[0024] The filter system may thus need to retain the previous M samples for each kernel (including the current sample); accordingly, $\mathbf{X}(t)$ may thus be defined as a $K \times M$ matrix of the past M samples for each of the K kernels, i.e.

$$\mathbf{X}(t) = \begin{bmatrix} X_t^{(1)} & X_{t-1}^{(1)} & \cdots & X_{t-M+1}^{(1)} \\ \vdots & \vdots & \ddots & \vdots \\ X_t^{(K)} & X_{t-1}^{(K)} & \cdots & X_{t-M+1}^{(K)} \end{bmatrix} \quad (3)$$

[0025] Denoting the kernel samples for the k -th kernel as $\mathbf{X}^{(k)}(t) = [X_t^{(k)} \quad \dots \quad X_{t-M+1}^{(k)}]$

and the kernel weights for the k -th kernel as $\mathbf{W}^{(k)} = [w_t^{(k)} \quad \dots \quad w_{t-M+1}^{(k)}]$, the overall filter output $e(t)$ may be given as

$$e(t) = \sum_{i=1}^K \mathbf{X}^{(i)}(t) \mathbf{W}^{(i)} \quad (4)$$

[0026] The filter output may thus produce a single symbol $e(t)$ at each time t that estimates the self-interference as the sum of the application of each of the K kernel filters to the previous M kernel samples of each kernel. Accordingly, such self-interference cancelation designs (which may be commonly known as “Hammerstein” models) may model the leakage

as the sum of the contributions of each kernel. The estimated interference signal $e(t)$ may then be subtracted from the received signal $y(t)$ (observed at the receive chain). Assuming the filter $\mathbf{W}$ is a suitable approximation of the actual self-interference signal, the resulting ‘clean’ signal $z(t)$ (where $z(t) = y(t) - e(t)$) may only contain a small or even negligible amount of residual self-interference and thus may be properly received and decoded by the receiver.

[0027] Self-interference cancelation designs may generally compute filter weight updates based on analysis of the clean signal $z(t)$ to determine both which filter weights of $\mathbf{W}$ are contributing to the interference residue in $z(t)$ and how to optimally adapt these filter weights to minimize the interference residue for future input samples. As the overall filter $\mathbf{W}$ in such conventional solutions may be a $K \times M$ matrix of filter weights, an optimal self-interference cancelation scheme may adapt all of the KM total filter weights at each input sample.

However, the required computational complexity of updating all KM filter weights may be prohibitively high and, consequently, many self-interference cancelation schemes have explored adaptation schemes that both reduce the number of weights updated at each stage and reduce the computational complexity of the update calculation.

[0028] As detailed herein, an advantageous self-interference cancelation solution of this disclosure may ‘decouple’ the FIR dimension from the kernel dimension in order to allow for computationally manageable adaptation stages. As will be detailed, this decoupling may allow for the FIR dimension to be updated separately from the kernel dimension and may additionally enable the usage of numerous additional computational techniques that may reduce the adaptation complexity. In particular for hardware implementations, the adaptation engine of this disclosure may provide substantial reductions in area and power and may thus prove to be an attractive design feasible with current radio technology. Additionally, the self-

interference cancelation scheme may offer configurability, and accordingly may be able to support variable numbers of taps and kernels for the filter weights.

[0029] FIG. 1 shows a block diagram of communication circuitry **100**, which may implement the self-interference cancelation system of this disclosure and may be a radio communication device (e.g. a mobile terminal, base station, network access point, etc.). As shown in FIG. 1, signal source **102** may generate a baseband transmit signal $s(t)$ for power amplifier **104**, which power amplifier **104** may amplify (following radio modulation, analog-to-digital conversion, and other pre-amplification stages not explicitly shown in FIG. 1) to produce an amplified RF transmit signal. Power amplifier **104** may provide the amplified RF transmit signal to duplexer **106**, which may be placed between power amplifier **104** and receiver **110** in order to allow for power amplifier **104** and receiver **110** to share antenna **108** (which may be a single antenna or an antenna array). Accordingly, in a full-duplex or frequency-division duplexing scheme duplexer **106** may allow power amplifier **104** and receiver **110** to simultaneously transmit and receive, respectively, with antenna **108**. Without loss of generality, communication circuitry **100** may be employed in short range, metropolitan, or cellular radio communication technologies such as 3GPP technologies (LTE, UMTS, GSM, etc.), WiFi, Bluetooth, etc. Communication circuitry **100** may be included in either uplink or downlink devices, including mobile terminals, base station, and network access points.

[0030] In addition to duplexer leakage, transmissions may leak onto the receive chain through the air, e.g. even if no antennas are shared (not explicitly shown in FIG. 1). Accordingly, one or both of the shared use of antenna **108** and wireless leakage may create self-interference that leaks from power amplifier **104** to receiver **110** over leakage path **130** as shown in FIG. 1 (where leakage path **130** may be composed of a wired and/or wireless path). Accordingly, the amplified radio transmit signal produced by power amplifier **104** via

amplification of baseband transmit signal $s(t)$ may appear in the received baseband signal $y(t)$ produced by receiver **110**. Accordingly, the presence of this leakage signal in $y(t)$ may corrupt a desired component of $y(t)$, i.e. the signal actually targeted by receiver **110** (such as a downlink signal received from a base station or an uplink signal received from a mobile terminal).

[0031] While incorporation of specialized duplexing circuitry in duplexer **106** may minimize the self-interference, such may be considerably expensive and thus undesirable in many designs. In order to enable sufficient reception performance by receiver **110**, communication circuitry **100** may instead implement the self-interference cancelation of this disclosure with cancelation circuitry **112**. As will be detailed, cancelation circuitry **112** may be configured to generate an estimated interference signal $e(t)$ from $s(t)$ and subtract $e(t)$ from $y(t)$ to produce clean signal $z(t)$. Assuming generation of $e(t)$ that accurately models leakage path **130**, cancelation circuitry **112** may be able to effectively remove the self-interference signal from $y(t)$ and thus produce clean signal $z(t)$ that is largely free of self-interference. Cancelation circuitry **112** may employ an adaptive filter system in order to model leakage path **130** which, as detailed below, decouples the PA dimension of the filter from the FIR dimension, thus allowing for reduced computational demands during adaptation.

[0032] As will be detailed, in an aspect of this disclosure communication circuitry **100** may be characterized as a communication circuit arrangement including a signal path circuit (e.g. the signal path of cancelation circuitry **112**) configured to separately apply a kernel dimension filter and a delay tap dimension filter to an input signal for an amplifier to obtain an estimated interference signal, a cancelation circuit (e.g. cancelation circuit **122**) configured to subtract the estimated interference signal from a received signal to obtain a clean signal, and a filter update circuit (e.g. the adaptation engine of cancelation circuitry **112**) configured to alternate between updating the kernel dimension filter and the delay tap dimension filter. In another

aspect of this disclosure, communication circuitry **100** may be characterized as a communication circuit arrangement including a subsignal generation circuit (e.g. kernel generation circuit **114**) configured to obtain one or more subsignals from an input signal for an amplifier, each of the one or more subsignals representing a non-linear component of an amplifier response and composed of a plurality of delay taps, a signal path circuit (e.g. the signal path of cancelation circuitry **112**) configured to separately apply a first filter and a second filter to the one or more subsignals to obtain an estimated interference signal, wherein the first filter approximates the amplifier response over the one or more subsignals and the second filter approximates the amplifier response over the plurality of delay taps, a cancelation circuit (e.g. cancelation circuit **122**) configured to subtract the estimated interference signal from the received signal to obtain a clean signal, and a filter update circuit (e.g. the adaptation engine of cancelation circuitry **112**) configured to alternate between updating the first filter and updating the second filter using the clean signal.

[0033] Accordingly, as opposed to utilizing a two-dimensional filter $\mathbf{W}$ as expressed above in Equation (2), cancelation circuitry **112** may utilize a single-dimensional $K \times 1$ power amplifier filter $\mathbf{W}_{PA}$ (PA filter circuit **118**) and a single-dimensional $M \times 1$ FIR filter $\mathbf{W}_F$ (FIR filter circuit **120**). Instead of employing M different weights for each of the K PA kernels, cancelation circuitry **112** may thus decouple the PA dimension from the FIR dimension and utilize a single set of $K \times 1$ PA filter weights in $\mathbf{W}_{PA}$ to collectively model the K PA kernels (over all taps) and similarly utilize a single set of $M \times 1$ FIR filter weights in $\mathbf{W}_F$ to collectively model the M taps (over all kernels). Given this drastic reduction in the amount of filter weights (from KM to $K + M$), cancelation circuitry **112** may considerably reduce the amount of computation needed to adapt the filter weights. Accordingly, cancelation circuitry **112** may present an attractive self-interference cancelation system that offers strong convergence speeds while avoiding excessively complex adaptation calculations.

Hardware area may consequently also be substantially reduced due to the lower computational complexity.

[0034] As shown in FIG. 1, cancelation circuitry **112** may include kernel generation circuit **114**, matrix memory **116**, PA filter circuit **118**, FIR filter circuit **120**, cancelation circuit **122**, PA filter adaptation circuit **124**, FIR filter adaptation circuit **126**, and adaptation switch circuit **128**. The corresponding functionality of the components of cancelation circuitry **112** may be structurally realized/embodied as hardware logic, e.g. as one or more integrated circuits or FPGAs, as software logic, e.g. as one or more processors executing program code that defining arithmetic, control, and I/O instructions stored in a non-transitory computer-readable storage medium, or as a combination of hardware and software logic. Accordingly, while the individual components of cancelation circuitry **112** are depicted separately in FIG. 1, this depiction serves to highlight the operation of cancelation circuitry **112** on a functional level; consequently, one or more of the components of cancelation circuitry **112** may be integrated into a common hardware and/or software element. Additionally, the functionality of cancelation circuitry **112** detailed herein (in particular e.g. the formulas/equations, flow charts, and prose descriptions) may be readily incorporated by skilled persons into program code for retrieval from a non-transitory computer readable medium and execution by a processor.

[0035] Cancelation circuitry **112** may be logically divided into a signal path and an adaptation engine, where the signal path may be tasked with generating the estimated interference signal $e(t)$ and the adaptation engine may be tasked with adapting the filter coefficients of PA filter circuit **118** and FIR filter circuit **120**. Accordingly, the signal path may be composed of kernel generation circuit **114**, matrix memory **116**, PA filter circuit **118**, FIR filter circuit **120**, and cancelation circuit **122** while the adaptation engine may be composed of PA filter adaptation circuit **124**, FIR filter adaptation circuit **126**, and adaptation

switch **128**. Accordingly, the signal path may derive the kernel signals $X_{t:t-M+1}^{(k)}$ for each kernel k for each input sample $s(t)$ and apply PA filter $\mathbf{W}_{PA}$ to the kernel signals $X_{t:t-M+1}^{(k)}$ for each kernel k , thus producing a $1 \times M$ output vector $(\mathbf{W}_{PA}^T \mathbf{X}(t))$, before applying FIR filter $\mathbf{W}_F$ to $\mathbf{W}_{PA}^T \mathbf{X}(t)$ (yielding $\mathbf{W}_{PA}^T \mathbf{X}(t) \mathbf{W}_F$) to obtain the estimated interference samples $e(t)$. The signal path may thus produce a sample of $e(t)$ for each input sample of $s(t)$ using the most recent M samples for each of the K kernels.

[0036] More specifically, kernel generation circuit **114** may obtain baseband transmit signal $s(t)$ from the receive path at the input to power amplifier **104**. Depending on the deployment of cancelation circuitry **112**, signal source **102** may be a baseband modem or similar baseband processing circuit (e.g. in a mobile terminal implementation) or a baseband unit (BBU) or similar baseband processing circuit (e.g. in a base station or access point implementation). Accordingly, kernel generation circuit **114** may tap the input to PA **104** to obtain $s(t)$ and may subsequently apply processing in order to derive the kernel samples $X_t^{(1:K)}$, i.e. the most recent kernel sample for each of the K kernels. As previously noted, kernel generation circuit **114** circuit may rely on a ‘kernelized’ model of PA **104** that estimates the non-linear response of PA **104** by dividing the response into multiple components, i.e. kernels, that collectively sum to model the overall response of PA **104**. Such may include using e.g. Cholesky decomposition or similar established processing solution in order to derive $X_t^{(1:K)}$.

Accordingly, kernel generation circuit **114** may be responsible for generating kernel samples $X_t^{(1:K)}$ from (t) . Various such kernelization techniques are established and are compatible for implementation in kernel generation circuit **114**.

[0037] Kernel generation circuit **114** may thus generate kernel samples $X_t^{(1:K)}$ for each sample $s(t)$ and provide the resulting samples $X_t^{(1:K)}$ to matrix memory **116**. As previously detailed, cancelation circuitry **112** may estimate the self-interference signal $e(t)$ based on the

M most recent samples (including the current sample at time t) of each of the K kernels, i.e. $X_{t:t-M+1}^{(1:K)}$. Accordingly, memory matrix **116** may store the $K \times M$ matrix containing the M most recent samples for each of the K kernels, which may be expressed as $\mathbf{X}(t)$ as defined in Equation (3) at each time t .

[0038] As processing circuitry **118** may generate an estimated self-interference sample $e(t)$ at each time t , matrix memory **116** may update $\mathbf{X}(t)$ with the K new kernel samples $X_t^{(1:K)}$ for each time t and provide $\mathbf{X}(t)$ to PA filter circuit **118**. PA filter circuit **118** may store the K filter weights of $K \times 1$ PA filter weight vector $\mathbf{W}_{PA}$, which may be expressed as

$$\mathbf{W}_{PA} = \begin{bmatrix} w_{PA,1} \\ w_{PA,2} \\ \vdots \\ w_{PA,K} \end{bmatrix} \quad (5)$$

where each $w_{PA,k}$, $k = 1, \dots, K$ is the complex PA weight corresponding to the k -th PA kernel.

[0039] PA filter circuit **118** may then apply $\mathbf{W}_{PA}$ to $\mathbf{X}(t)$ as $\mathbf{W}_{PA}^T \mathbf{X}(t)$ to obtain an $M \times 1$ vector, which PA filter circuit **118** may provide to FIR filter circuit **120**. FIR filter circuit **120** may likewise store the M filter weights of $M \times 1$ filter weight vector $\mathbf{W}_F$, which may be expressed as

$$\mathbf{W}_F = \begin{bmatrix} w_{F,1} \\ w_{F,2} \\ \vdots \\ w_{F,M} \end{bmatrix} \quad (6)$$

where each $w_{F,m}$, $m = M$ is the complex FIR weight corresponding to the m -th tap, i.e. the m -th most recent sample. As leakage path **130** may include (in addition to the non-linear

response of PA **104**) multiple transmit and receive FIR filters, the weights of FIR filter $\mathbf{W}_F$ may therefore model the various taps of each of these FIR filters.

[0040] FIR filter circuit **120** may then apply $\mathbf{W}_F$ to the output vector $\mathbf{W}_{PA}^T \mathbf{X}(t)$ of PA filter circuit **118** as $\mathbf{W}_{PA}^T \mathbf{X}(t) \mathbf{W}_F$ to obtain a single estimated self-interference sample $e(t)$ for each time t . Cancellation circuit **122** may then subtract $e(t)$ from $y(t)$ provided by receiver **110** to obtain clean signal $z(t) = y(t) - e(t)$ for each time t , which may be largely free of self-interference if $e(t)$ closely matches the actual self-interference signal leaking from power amplifier **104** through duplexer **106** on leakage path **130**.

[0041] The efficacy of the self-interference cancellation scheme applied by cancellation circuitry **112** may depend on how accurately $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ model the self-interference over each kernel and tap. Accordingly, cancellation circuitry **112** may employ the adaptation engine composed of PA filter adaptation circuit **124**, FIR filter adaptation circuit **126**, and adaptation switch circuit **128** in order to dynamically update the filter weights $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ of PA filter circuit **118** and FIR filter circuit **120**.

[0042] As the PA and FIR filter weights $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ have been decoupled (compared to the ‘coupled’ weights of $\mathbf{W}$ in Equation (2)), the adaptation engine of cancellation circuitry **112** may be able to update $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ separately, which may substantially reduce the computational complexity of the update process and consequently reduce hardware area and power requirements. Whereas update of the full $K \times M$ matrix $\mathbf{W}$ involved considering all KM filter weights for update, the adaptation engine may only need to consider either K or M total filter weights of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$, respectively, at a time. Additionally, the adaptation engine of cancellation circuitry **112** may employ an adaptation scheme in which $\mathbf{W}_{PA}$ or $\mathbf{W}_F$ (depending on which is currently being updated and which is fixed) is updated for every input sample $s(t)$. Accordingly, $\mathbf{W}_{PA}$ or $\mathbf{W}_F$ may continuously ‘track’ the system in real-time in order to adaptively track changes. Furthermore, such may provide the possibility to re-use

shared processing circuitry to update $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ at disparate times, thus allowing for further hardware reductions.

[0043] FIG. 2 shows method 200 illustrating the general procedure of cancelation circuitry 112, which may iterate over each time t starting at 202. Kernel generation circuit 114 and matrix memory 116 may first generate the KM entries of $\mathbf{X}(t)$ from $s(t)$ based on the most recent kernel samples $X_t^{(1:K)}$ and the past kernel samples $X_{t-1:t-M+1}^{(1:K)}$ in 204. Each set of kernel samples $X_{t:t-M+1}^{(k)}$ may thus be considered a ‘subsignal’ derived from baseband input signal $s(t)$, e.g. according to kernel generation scheme such as based on Cholesky decomposition. PA filter circuit 118 and FIR filter circuit 120 may then apply $\mathbf{W}_{PA}$ and $\mathbf{W}_F$, respectively, to $\mathbf{X}(t)$ as $\mathbf{W}_{PA}^T \mathbf{X}(t) \mathbf{W}_F$ to obtain $e(t)$ in 206. Cancelation circuit 122 may then subtract $e(t)$ from received signal $y(t)$ to obtain clean signal $z(t)$ in 208. If switching circuit 128 has currently selected $\mathbf{W}_{PA}$ for adaptation, PA filter adaptation circuit 124 may update $\mathbf{W}_{PA}$ based on $z(t)$ and $\mathbf{X}(t)$ in 212, i.e. by evaluating the uncanceled leakage in $z(t)$ resulting from $\mathbf{X}(t)$; alternatively, if switching circuit 128 has currently selected $\mathbf{W}_F$ for adaptation, FIR filter adaptation circuit 126 may update $\mathbf{W}_F$ based on $z(t)$ and $\mathbf{X}(t)$ in 214. Method 200 may then iteratively repeat for $t = t + 1$, where switching circuit 128 may periodically change the selection between PA adaptation circuit 124 and FIR adaptation circuit 126 in order to ensure that both $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ are updated over time.

[0044] The decoupling of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ updates may provide considerable reductions in computational complexity due to the reduced number of filter weights, which may also open up the possibility for re-using the same shared processing circuitry to update $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ at alternating times (thus further saving hardware area). Additionally, the decoupling of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ may allow for cancelation circuitry 112 to also employ numerous additional computational techniques that further reduce the computational complexity. As will be further

detailed below, such techniques may include a Dichotomous Coordinate Descent (DCD) update algorithm, simplified maximum magnitude selection, reduced matrix memory storage, preprocessing calculation stages, and clock pipelining schedules. The self-interference cancelation solution of this disclosure may additionally offer configurability in terms of the supported number of taps M and kernels K . Such enhancements are considered optional and may or may not be included in various aspects of this disclosure.

[0045] The general update procedure of the adaptation engine of cancelation circuitry **112** will first be detailed. As $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ are decoupled into two separate vectors, the adaptation engine may update each of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ separately (as in **212** and **214** of method **200** depending on the selection of switching circuit **128**). By updating only one of $\mathbf{W}_{PA}$ or $\mathbf{W}_F$ at a time, the adaptation engine may ‘fix’ one dimension (i.e. either the PA dimension or the FIR dimension) and ignore any dependencies on the fixed dimension during update of the other dimension, thus simplifying the adaptation process. Accordingly, switching circuit **128** may control PA filter adaptation circuit **124** and FIR filter adaptation circuit **126** to alternate between updating the PA weights of $\mathbf{W}_{PA}$ and the FIR weights of $\mathbf{W}_F$. Accordingly, only K or M filter weights will be considered for update at a given time (resulting directly from the decoupling), which may present a drastic reduction over the KM filter weights of $\mathbf{W}$.

[0046] As previously indicated, the adaptation engine may utilize clean signal $z(t)$ in order to decide which filter weights to update. More specifically, as shown in FIG. **1** both PA filter adaptation circuit **124** and FIR adaptation circuit **126** may receive both $z(t)$ and $\mathbf{X}(t)$ (the clean signal sample $z(t)$ and kernel sample matrix $\mathbf{X}(t)$ at time t) and compare $z(t)$ with $\mathbf{X}(t)$ to determine if $z(t)$ is correlated with $\mathbf{X}(t)$. Accordingly, if $z(t)$ is correlated with $\mathbf{X}(t)$, this may indicate that the self-interference cancelation is not effective and that the filter weights need to be adjusted to obtain more accurate estimates in $e(t)$. Depending on the specifics of the adaptation scheme used by the adaptation engine, PA filter adaptation circuit

124 and FIR adaptation circuit **126** may then analyze the relationship between $z(t)$ with $\mathbf{X}(t)$ in order to identify one or more filter weights of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ to update. In optimal adaptation schemes, PA filter adaptation circuit **124** and FIR adaptation circuit **126** may determine an update for each weight of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$, such as an optimal update for each weight to minimize the square error of an error signal (in the case of LMS) or to minimize a cost function (in the case of RLS). Accordingly, in certain aspects of this disclosure, the adaptation engine of cancelation circuitry **112** may update the filter weights of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ with an adaptation scheme, such as e.g. LMS or RLS, which may include switching circuit **128** alternating between update of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ by activating one of either PA filter adaptation circuit **124** or FIR filter adaptation circuit **126** at a time. Accordingly, one of either $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ will be updated during each update iteration at time t and the weight adjustments will be reflected in application of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ to $\mathbf{X}(t + 1)$ by PA filter circuit **118** and FIR filter circuit **120**.

[0047] While application of such optimal update schemes (i.e. to update each filter weight) to the decoupled filters $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ is within the scope of the solution presented in this disclosure, such schemes may be overly complex due to the need to calculate an update for each filter weight. Accordingly, the adaptation engine may instead select a reduced number of filter weights (i.e. only some of the filter weights) to update at a time. For example, in update iterations where PA filter adaptation circuit **124** is activated by switching circuit **128** (and thus FIR filter adaptation circuit **126** is deactivated), PA filter adaptation circuit **124** may compare $z(t)$ and $\mathbf{X}(t)$ and identify a single weight of $\mathbf{W}_{PA}$ that makes the greatest contribution to the self-interference residue remaining in $z(t)$, i.e. that is ‘most accountable’ for un-canceled self-interference in $z(t)$. PA filter adaptation circuit **124** may then only calculate an update for the identified weight of $\mathbf{W}_{PA}$ during the current update iteration. If switching circuit **128** selects to update $\mathbf{W}_{PA}$ again during the next update iteration, PA filter adaptation circuit **124**

may compare $z(t + 1)$ and $\mathbf{X}(t + 1)$ (i.e. the clean signal and kernel matrix for the next time $t + 1$) and again select a single weight of $\mathbf{W}_{PA}$ to update. As the update occurs at the next time $t + 1$, the clean signal $z(t + 1)$ may reflect the change in $\mathbf{W}_{PA}$ as $e(t + 1)$ will directly depend on the updated weight. FIR filter adaptation circuit **126** may similarly select single weights of $\mathbf{W}_F$ to update during each update iteration for which FIR filter adaptation circuit **126** is activated by switching circuit **128**.

[0048] Accordingly, instead of updating all or some of the filter weights of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ during each update iteration, the adaptation engine may select only one weight of $\mathbf{W}_{PA}$ or $\mathbf{W}_F$ (depending on the selection by switching circuit **128**) to update during each update iteration. Such adaptation schemes are commonly known as Coordinate Descent (CD), where updates are made in only a single direction at a time. As previously detailed regarding LMS and RLS, adaptive algorithms may aim to reduce a squared error term or a cost function; accordingly, LMS-CD may update a single weight in order to reduce the squared error term while RLS-CD may update a single weight in order to reduce a cost function. While the adaptation engine of cancelation circuitry **112** may employ any of a number of different adaptation schemes to update $\mathbf{W}_F$ and $\mathbf{W}_{PA}$, an advantageous solution of this disclosure may utilize an RLS-DCD scheme in which a single weight of $\mathbf{W}_F$ or $\mathbf{W}_{PA}$ (depending on the current selection of switching circuit **128**) is updated using a single bit-inversion in order to minimize a cost function derived from $\mathbf{X}(t)$ and $z(t)$. Accordingly, in each update iteration, the appropriate adaptation engine (PA filter adaptation circuit **124** or FIR filter adaptation circuit **126**) may evaluate $\mathbf{X}(t)$ and $z(t)$ to identify which filter weight (of $\mathbf{W}_{PA}$ or $\mathbf{W}_F$) is the worst-offender in terms of uncanceled leakage in $z(t)$ and to identify which bit of a binary representation of the selected weight should be inverted, or ‘flipped’, in order to reduce the cost function. Accordingly, this single-bit inversion update of a single weight may yield the ‘Dichotomous’ Coordinate Descent nomenclature while the minimization of a cost function (namely the

uncanceled residue in $z(t)$ represented by the cross-correlation between $\mathbf{X}(t)$ and $z(t)$) may yield the RLS designation. It is nevertheless noted that other adaptation algorithms may be integrated *mutatis mutandis* in place of the RLS-DCD algorithm to update the decoupled weights of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$.

[0049] FIGs. 3 and 4 further illustrate the update iterations 212 and 214 of method 200, where PA filter adaptation circuit 124 and FIR filter adaptation circuit 126 may respectively update $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ depending on the current selection by switching circuit 128. PA update iteration 212 may involve largely identical calculations to FIR update iteration 214, where the only substantial difference involves the calculation of the decoupled input signals $\mathbf{X}_{PA}(t)$ and $\mathbf{X}_F(t)$. As will be later detailed, the adaptation engine may thus be able to re-use the same shared circuitry for update iterations 212 and 214. Without loss of generality, the context depicted in FIGs. 3 and 4 and detailed below corresponds to a RLS-DCD adaptation algorithm in which where PA filter adaptation circuit 124 and FIR filter adaptation circuit 126 aim to minimize the uncanceled leakage residue in $z(t)$ by selecting a single bit to invert of a selected filter weight of $\mathbf{W}_{PA}$ or $\mathbf{W}_F$. However, other adaptation algorithms, such as based on LMS, CD, gradient descent, etc., may also be employed to update $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ in a decoupled manner while remaining within the scope of this disclosure.

[0050] As shown in FIG. 3, PA filter adaptation circuit 124 may first calculate decoupled PA input signal $\mathbf{X}_{PA}(t)$ in 212a as $\mathbf{X}_{PA}(t) = \mathbf{X}(t)\mathbf{W}_F$, where $\mathbf{X}_{PA}(t)$ is of dimension $1 \times K$ and each k -th element of $\mathbf{X}_{PA}(t)$ is representative of the M taps of the k -th kernel of $\mathbf{X}(t)$. Accordingly, PA filter adaptation circuit 124 may ‘fix’ the weights of $\mathbf{W}_F$ and only pursue update of $\mathbf{W}_{PA}$. As previously indicated, PA filter adaptation circuit 124 may compare $\mathbf{X}(t)$ to $z(t)$ in order to identify which weights of $\mathbf{W}_{PA}$ are contributing to uncanceled leakage in $z(t)$. More specifically, PA filter adaptation circuit 124 may utilize the cross-correlation $\beta_{PA}(t)$ between $\mathbf{X}_{PA}(t)$ and $z(t)$ in order to identify which elements of $\mathbf{X}_{PA}(t)$ are

substantially correlated with $z(t)$, thus allowing PA filter adaptation circuit **124** to identify the weights of $\mathbf{W}_{PA}$ corresponding to the correlated elements of $\mathbf{X}_{PA}(t)$ as requiring update. As PA filter adaptation circuit **124** may need the correlations $\mathbf{R}_{PA}(t)$ between each of the samples of $\mathbf{X}_{PA}(t)$ in order to calculate the proper update, PA filter adaptation circuit **124** may also calculate correlation matrix $\mathbf{R}_{PA}(t)$.

[0051] Accordingly, in **212b** PA filter adaptation circuit **124** may calculate cross-correlation vector $\boldsymbol{\beta}_{PA}(t)$ (of dimension $1 \times M$) in **212b** as

$$\boldsymbol{\beta}_{PA}(t) = \lambda(\boldsymbol{\beta}_{PA}(t-1) - \text{sign}(\beta_{PA,max})\alpha\mathbf{R}_{PA}(t-1)^{(n)}) + \mathbf{X}_{PA}^H(t)z(t) \quad (7)$$

and correlation matrix $\mathbf{R}_{PA}(t)$ as

$$\mathbf{R}_{PA}(t) = \lambda\mathbf{R}_{PA}(t-1) + \mathbf{X}_{PA}^H(t)\mathbf{X}_{PA}(t) \text{ (diag.)} \quad (8)$$

where λ is a forgetting factor, $\text{sign}(\beta_{PA,max})\alpha\mathbf{R}_{PA}(t-1)^{(n)}$ is an update vector that updates $\boldsymbol{\beta}_{PA}(t)$ according to the filter weight update of the n -th element of $\mathbf{W}_{PA}$ from the previous iteration at $t-1$, and (diag.) indicates that only the diagonal elements of $\mathbf{R}_{PA}(t)$ are updated. As will be later detailed, $\mathbf{R}_{PA}(t-1)^{(n)}$ may be the n -th column vector of $\mathbf{R}_{PA}(t-1)$, where n corresponds to the index of $\mathbf{W}_{PA}$ that was updated at the iteration for time $t-1$. As will also be later detailed, the diagonal and non-diagonal elements of $\mathbf{R}_{PA}(t)$ may be performed separately in order to support for an effective hardware pipeline schedule, e.g. where the diagonal elements of $\mathbf{R}_{PA}(t)$ are updated first in **212b** and are later followed by update of the non-diagonal elements of $\mathbf{R}_{PA}(t)$ in **212e** (e.g. where the diagonal elements of $\mathbf{R}_{PA}(t)$ are part of the critical path and thus should be processed before the non-diagonal elements to optimize performance). However, in certain aspects of this disclosure all elements

of $\mathbf{R}_{PA}(t)$ may alternatively be updated in a single procedure. As the update vector $\text{sign}(\beta_{PA,max})\alpha\mathbf{R}_{PA}(t-1)^{(n)}$ depends on the previous update iteration, PA filter adaptation circuit **124** may calculate $\beta_{PA}(t)$ in the first update iteration as $\lambda\beta_{PA}(t-1) + \mathbf{X}_{PA}^H(t)z(t)$, where $\beta_{PA}(t-1)$ is the initialization value of β_{PA} (which may be set to e.g. a zero matrix). $\mathbf{R}_{PA}(t-1)$ may likewise be the initialization value of $\mathbf{R}_{PA}$.

[0052] Accordingly, $\beta_{PA}(t)$ may indicate the cross-correlation between decoupled PA input signal $\mathbf{X}_{PA}(t)$ and clean signal $z(t)$, which thus indicates the residual uncanceled leakage attributed to each weight of $\mathbf{W}_{PA}$. Accordingly, inaccurate weights of $\mathbf{W}_{PA}$ (i.e. weights that do not accurately characterize the corresponding kernel) will produce corresponding elements of $\beta_{PA}(t)$ that have high magnitude. As the goal of PA filter adaptation circuit **124** is to minimize $\beta_{PA}(t)$ (i.e. $\beta_{PA}(t)$ is the RLS cost function), PA filter adaptation circuit **124** may thus aim to update the weights of $\mathbf{W}_{PA}$ in order to reduce $\beta_{PA}(t)$, thus reducing the uncanceled leakage residue in $z(t)$.

[0053] As previously indicated, in an RLS-DCD context PA filter adaptation circuit **124** may select a single weight of $\mathbf{W}_{PA}$ based on $\beta_{PA}(t)$ and invert a single bit of $\mathbf{W}_{PA}$ in order to minimize $\beta_{PA}(t)$. This may be computationally simpler than e.g. updating each weight of $\mathbf{W}_{PA}$ and/or calculating an optimum update for each updated weight of $\mathbf{W}_{PA}$. As the weights of $\mathbf{W}_{PA}$ that make the greatest contribution to uncanceled residue will correspond to the elements of $\beta_{PA}(t)$ with the highest magnitude, PA filter adaptation circuit **124** may in **212c** identify the element $\beta_{PA,max}$ of $\beta_{PA}(t)$ having the highest magnitude and the corresponding index n ($n \in \{1, \dots, K\}$) of $\beta_{PA,max}$ within $\beta_{PA}(t)$, which PA filter adaptation circuit **124** may then apply to update the n -th weight of $\mathbf{W}_{PA}$. As the elements of $\beta_{PA}(t)$ may be complex, in an optimal scenario PA filter adaptation circuit **124** may in **212c** calculate the Euclidean norm of each element of $\beta_{PA}(t)$ and identify the element with the highest magnitude, i.e.

$\max(|\beta_{PA}(t)|)$. However, as such may involve the computation of a squares and square roots ($|a + jb| = \sqrt{a^2 + b^2}$), PA filter adaptation circuit **124** may simplify the computation in **212c** by identifying the element of $\beta_{PA}(t)$ with the highest real or imaginary part as $\beta_{PA,max}$ (located at index n in $\beta_{PA}(t)$), i.e.

$$\beta_{PA,max} = \max\{|Im\{\beta_{PA}(t)\}|, |Re\{\beta_{PA}(t)\}|\} \quad (9)$$

[0054] Accordingly, while Equation (9) is an approximation such may still allow PA filter adaptation circuit **124** to identify an element of $\beta_{PA}(t)$ with the highest real or imaginary part, which may likely be one of the highest if not the highest-magnitude elements of $\beta_{PA}(t)$. PA filter adaptation circuit **124** may thus identify $\beta_{PA,max}$ and the corresponding index n of $\beta_{PA,max}$ in $\beta_{PA}(t)$, which may correspond to the n -th weight of $\mathbf{W}_{PA}$ thus resulting in PA filter adaptation circuit **124** identifying the weight of $\mathbf{W}_{PA}$ that is selected for update.

[0055] PA filter adaptation circuit **124** may then update the n -th weight of $\mathbf{W}_{PA}$, $w_{PA,n}$, based on $\beta_{PA,max}$ and correlation matrix $\mathbf{R}_{PA}(t)$. In a conventional RLS-CD update solution (i.e. non-Dichotomous), PA filter adaptation circuit **124** may update $w_{PA,n}$ in **212d** as $w_{PA,n} = w_{PA,n} \mp \beta_{PA,max}/R_{PA(n,n)}$, i.e. by performing a division of $\beta_{PA,max}$ by the correlation value $R_{PA(n,n)}$ of the n -th element of $\mathbf{X}_{PA}(t)$ with itself (located at the n -th diagonal term of $\mathbf{R}_{PA}$), which may provide the optimal minimization of $\beta_{PA}(t)$ as achieved by update of a single weight. Without loss of generality, PA filter adaptation circuit **124** may alternatively avoid the computational complexity of the divisional operation by instead selecting a single bit of w (i.e. to flip from a 0 to 1 or vice versa). As $R_{PA(n,n)}$ is located on the diagonal of $\mathbf{R}_{PA}(t)$, PA filter adaptation circuit **124** may thus need to ensure that the diagonal entries of $\mathbf{R}_{PA}(t)$ are up-to date at **214d**. As the non-diagonally entries of $\mathbf{R}_{PA}(t)$ will then be used in

214e to update $\mathbf{R}_F(t)$, PA filter adaptation circuit **124** may separate the update of $\mathbf{R}_{PA}(t)$ into diagonal and non-diagonal portions in order to enable a pipelined clock schedule that shares computational elements.

[0056] As expressed in the algorithmic logic of **212d** in FIG. 3, PA filter adaptation circuit **124** may check each of the M_b bits of $w_{PA,n}$ in order to identify which of the bits would optimally minimize $\beta_{PA}(t)$ if inverted. Accordingly, for e.g. $M_b = 8$, this may amount to a subtraction of 1, 2, 4, 8, ..., or 128 from $w_{PA,n}$ (i.e. $2^0, 2^1, \dots, 2^7$). As the optimal update to $w_{PA,n}$ is $-\beta_{PA,max}/R_{PA(n,n)}$, PA filter adaptation circuit **124** may aim to determine in **212d** which of the possible subtraction values from 2^{M_b-1} to 2^0 that $\beta_{PA,max}/R_{PA(n,n)}$ is closest to. Accordingly, starting with e.g. $\alpha = 2^{M_b}$ (although the initiating value of α may be a design parameter and other values other than the most significant bit could be chosen, such as e.g. the second-most significant bit, third-most significant bit, etc.), PA filter adaptation circuit **124** may iterate from $l = 1$ to $l = M_b$ (e.g. $M_b = 8$) in order to determine which $\frac{\alpha}{2}$ satisfies the condition $|\beta_{PA,max}| > \frac{\alpha}{2} R_{PA(n,n)}$ and, upon identifying the appropriate $\frac{\alpha}{2}$, take α as the update value for $w_{PA,n}$. By identifying the first α for which $|\beta_{PA,max}| > \frac{\alpha}{2} R_{PA(n,n)}$, or $\frac{|\beta_{PA,max}|}{R_{PA(n,n)}} > \frac{\alpha}{2}$, PA filter adaptation circuit **124** may identify the α closest in value to $\beta_{PA,max}/R_{PA(n,n)}$, i.e. the optimum update of $w_{PA,n}$ to minimize $\beta_{PA}(t)$. In other words, starting from a possible set of update values 2^{M_b-1} to 2^0 , PA filter adaptation circuit **124** may identify which update value is closest to the optimum update value $\beta_{PA,max}/R_{PA(n,n)}$ and select the identified update value as α . Flipping the l -th bit (where l gives the iteration for which $|\beta_{PA,max}| > \frac{\alpha}{2} R_{PA(n,n)}$ is satisfied) may thus produce an update of α (positive or negative depending on the sign of $\beta_{PA,max}$).

[0057] Upon identifying update factor α , PA filter adaptation circuit **124** may flip the l -th LSB bit of $w_{PA,n}$ to update $w_{PA,n}$ as

$$w_{PA,n} = w_{PA,n} + \text{sign}(\beta_{F,\max})\alpha \quad (10)$$

where $\text{sign}(\beta_{F,\max})$ preserves the sign of $\beta_{F,\max}$ to ensure that $w_{PA,n}$ is updated in the proper direction.

[0058] Accordingly, upon completion of the update of $w_{PA,n}$ in **212d**, PA filter adaptation circuit **124** may have selected the ‘worst offender’ $w_{PA,n}$ of $\mathbf{W}_{PA}$, i.e. the weight of $\mathbf{W}_{PA}$ with the highest contribution (based on having the maximum real or imaginary part) to the uncanceled residue in $z(t)$ and invert a single bit of the selected weight $w_{PA,n}$ in order to minimize $\beta_{PA}(t)$, which in connection also minimizes the leakage residue in $z(t)$. PA filter adaptation circuit **124** may then complete the update iteration for time t by updating the non-diagonal elements of $\mathbf{R}_{PA}(t)$ in **212e** as

$$\mathbf{R}_{PA}(t) = \lambda \mathbf{R}_{PA}(t-1) + \mathbf{X}_{PA}^H(t) \mathbf{X}_{PA}(t) \quad (\text{non} - \text{diag}) \quad (11)$$

[0059] As previously indicated, as the update of $\beta_{PA}(t)$ in **212b** requires the off-diagonal elements of $\mathbf{R}_{PA}(t-1)$ for the n -th column $\mathbf{R}_{PA}(t-1)^{(n)}$, PA filter adaptation circuit **124** may perform the off-diagonal update in **212e** separately in order to re-use computational hardware for both **212b** and **212e**.

[0060] Assuming switching circuit **128** proceeds to $t+1$ with PA filter adaptation circuit **124** still selected, PA filter adaptation circuit **124** may repeat **212** for $t = t+1$. Switching circuit **128** may be configured to maintain the update selection of PA filter adaptation circuit **124** or FIR filter adaptation circuit **126** for multiple iterations before switching, which may

include maintaining the selection of either PA filter adaptation circuit **124** or FIR filter adaptation circuit **126** until the selected filter $\mathbf{W}_{PA}$ or $\mathbf{W}_F$ converges or switching between PA filter adaptation circuit **124** or FIR filter adaptation circuit **126** according to a predetermined cyclic number of update iterations (e.g. every other iteration, every 3 iterations, every 10 iterations, etc.). Accordingly, PA filter adaptation circuit **124** may continue to update $\mathbf{W}_{PA}$ during each iteration by calculating a new $\mathbf{X}_{PA}(t)$ based on the new $\mathbf{X}(t)$ matrix in **212a**, updating $\boldsymbol{\beta}_{PA}(t)$ and $\mathbf{R}_{PA}(t)$ in **212b**, identifying the maximum element $\beta_{PA,max}$ of $\boldsymbol{\beta}_{PA}(t)$ and corresponding index n in **212c**, selecting a bit to flip of $w_{PA,n}$ by determining α in **212d**, and completing update of $\mathbf{R}_{PA}(t)$ in **212e**.

[0061] Conversely, when switching circuit **128** selects FIR filter adaptation circuit **126** for update, FIR filter adaptation circuit **126** may execute the FIR update iteration **214** as shown in FIG. 4. FIR filter adaptation circuit **126** may thus calculate decoupled FIR input signal $\mathbf{X}_F(t)$ in **214a** as $\mathbf{X}_F(t) = \mathbf{W}_{PA}^T \mathbf{X}(t)$, where $\mathbf{X}_F(t)$ is of dimension $1 \times M$ and each m -th element represents the joint contribution of all K kernels to the m -th taps $\mathbf{X}(t)$. Likewise to PA update iteration **212**, FIR filter adaptation circuit **126** may be configured to evaluate the cross-correlation $\boldsymbol{\beta}_F(t)$ between $\mathbf{X}_F(t)$ and $z(t)$ in order to identify which taps of $\mathbf{X}_F(t)$ are most correlated with $z(t)$, thus identifying which taps of $\mathbf{W}_F$ are the ‘worst-offenders’ in terms of uncanceled residual leakage in $z(t)$. As these taps are accordingly the least accurate, in accordance with RLS-DCD FIR filter adaptation circuit **126** may proceed to select the least accurate tap $w_{F,n}$ and flip a single bit of $w_{F,n}$ to reduce $\boldsymbol{\beta}_F(t)$, thus reducing the uncanceled leakage in $z(t)$.

[0062] Accordingly, in **214b** FIR filter adaptation circuit **126** may calculate cross-correlation vector $\boldsymbol{\beta}_F(t)$ (of dimension $1 \times M$) as

$$\boldsymbol{\beta}_F(t) = \lambda(\boldsymbol{\beta}_F(t-1) - \text{sign}(\beta_{F,max})\alpha\mathbf{R}_F(t-1)^{(n)}) + \mathbf{X}_F^H(t)z(t) \quad (12)$$

and correlation matrix $\mathbf{R}_F(t)$ as

$$\mathbf{R}_F(t) = \lambda\mathbf{R}_F(t-1) + \mathbf{X}_F^H(t)\mathbf{X}_F(t) \text{ (diag.)} \quad (13)$$

where, likewise to PA update iteration **212**, λ is a forgetting factor, $\text{sign}(\beta_{F,max})\alpha\mathbf{R}_F(t-1)^{(n)}$ is an update vector that updates $\boldsymbol{\beta}_F(t)$ according to the filter weight update of the n -th element $w_{F,n}$ of $\mathbf{W}_F$ from the previous iteration at $t-1$ (with the n -th column of $\mathbf{R}_F(t)$), and (diag.) indicates that only the diagonal elements of $\mathbf{R}_F(t)$ are updated. Likewise to above, FIR filter adaptation circuit **126** may update the diagonal and non-diagonal elements of $\mathbf{R}_F(t)$ separately in order to support a hardware pipeline clock schedule that re-uses computational hardware for **214b** and **214e**; however, in various aspects of this disclosure FIR filter adaptation circuit **126** may alternatively update all elements of $\mathbf{R}_F(t)$ in a single procedure. As the update vector $\text{sign}(\beta_{F,max})\alpha\mathbf{R}_F(t-1)^{(n)}$ depends on the previous update iteration, FIR filter adaptation circuit **124** may calculate $\boldsymbol{\beta}_F(t)$ in the first update iteration as $\lambda\boldsymbol{\beta}_F(t-1) + \mathbf{X}_F^H(t)z(t)$, where $\boldsymbol{\beta}_F(t-1)$ is the initialization value of $\boldsymbol{\beta}_F$ (which may be set to e.g. a zero matrix). $\mathbf{R}_F(t-1)$ may likewise be the initialization value of $\mathbf{R}_F$.

[0063] As in PA update iteration **212**, $\boldsymbol{\beta}_F(t)$ may represent the cross-correlation between each tap of decoupled FIR input signal $\mathbf{X}_F(t)$ and clean signal $z(t)$. Accordingly, FIR filter adaptation circuit **126** may aim to identify the element of $\boldsymbol{\beta}_F(t)$ with the highest magnitude, which may point to the element of $\mathbf{W}_F$ that is least accurate. FIR filter adaptation circuit **126** may thus identify index n of the maximum-valued element of $\boldsymbol{\beta}_F(t)$ in **214c** as

$$\beta_{F,max} = \max\{|Im\{\beta_F(t)\}|, |Re\{\beta_F(t)\}|\} \quad (14)$$

[0064] As detailed above regarding **212c**, identifying the element $\beta_{F,max}$ of $\beta_F(t)$ with the highest real or imaginary part may be less computationally complex than calculating the magnitude (e.g. L2 norm) of each element of $\beta_F(t)$. The approximation of Equation (14) may nevertheless generally identify $\beta_{F,max}$ as an element of $\beta_F(t)$ that has one of the highest magnitudes.

[0065] FIR filter adaptation circuit **126** may thus identify $\beta_{F,max}$ and the index n of $\beta_{F,max}$ within $\beta_F(t)$, where index n also corresponds to the index of the weight $w_{F,n}$ of $W_F(t)$ that is the ‘worst offender’ and thus needs to be updated. Likewise as to **212d**, FIR filter adaptation circuit **126** may in **214d** select a single bit of $w_{F,n}$ to flip in order to minimize $\beta_F(t)$. As the optimum update value of $w_{F,n}$ is given as $\beta_{F,max}/R_{F(n,n)}$, FIR filter adaptation circuit **126** may identify the scale factor $\alpha \in (2^0, 2^1, \dots, 2^{M_b-1})$ (each corresponding to flipping the l -th bit of $w_{F,n}$ for $l = 1, 2, \dots, M_b$) to adjust $w_{F,n}$ by.

[0066] Accordingly, as shown in FIG. 4, in **214d** FIR filter adaptation circuit **126** may check the condition $|\beta_{F,max}| > \frac{\alpha}{2} R_{F(n,n)}$ for $\alpha = 2^{M_b-1}$ to 2^0 and, if the condition is satisfied, take the current value of α as being closest to the optimum update value $\beta_{F,max}/R_{F(n,n)}$ and update $w_{F,n}$ as

$$w_{F,n} = w_{F,n} + \text{sign}(\beta_{F,max})\alpha \quad (15)$$

thus updating W_F to minimize the uncanceled leakage residue indicated by $\beta_F(t)$.

[0067] FIR filter adaptation circuit **126** may then update the off-diagonal elements of $R_F(t)$ as

$$\mathbf{R}_F(t) = \lambda \mathbf{R}_F(t-1) + \mathbf{X}_F^H(t) \mathbf{X}_F(t) \quad (\text{non-diag.}) \quad (16)$$

[0068] As in the case of PA filter adaptation circuit **124** noted above, FIR filter adaptation circuit **126** may then proceed to update $\mathbf{W}_F(t)$ in the next iteration for $t = t + 1$ with $\mathbf{X}(t + 1)$ and $z(t + 1)$ if switching circuit **128** maintains the selection of FIR filter adaptation circuit **126**.

[0069] As update of PA filter $\mathbf{W}_{PA}$ and FIR filter $\mathbf{W}_F$ are decoupled, each of PA filter adaptation circuit **124** and FIR filter adaptation circuit **126** may be able to compute updates to $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ in an isolated environment, i.e. without having to account for the other dimension. The computational requirements required for each update may be substantially reduced in comparison with the two-dimensional case of the $K \times M$ filter matrix $\mathbf{W}$ noted above regarding the existing solution. In particular, the computational complexity of the correlation matrices $\mathbf{R}_{PA}$ and $\mathbf{R}_F$ may be markedly diminished. More specifically, in the existing solution the adaptation engine may be required to compute a $KM \times KM$ correlation matrix $\mathbf{R}$ that represents the correlations between each M taps of each K kernels. In contrast, the adaptation engine of cancellation circuitry **112** may only need to calculate $K \times K$ PA correlation matrix $\mathbf{R}_{PA}$ and $M \times M$ FIR correlation matrix $\mathbf{R}_F$. Furthermore, instead of applying the $K \times M$ filter $\mathbf{W}$ to $\mathbf{X}(t)$ as in the existing solution, the signal path of cancellation circuitry **112** may only apply a $K \times 1$ PA filter vector $\mathbf{W}_{PA}$ and $M \times 1$ FIR filter vector $\mathbf{W}_F$.

[0070] Various advantageous modifications have been noted above regarding cancellation circuitry **112**, including the use of RLS-DCD (as opposed to gradient descent, traditional coordinate descent, etc.) and identifying the maximum magnitude on account of the highest real or imaginary component. These simplifications are optional, and may or may not be utilized along with the decoupled filter application and adaptation while still remaining within

the scope of this disclosure. The adaptation engine may optionally also employ a variety of further techniques in order to reduce the required computational complexity, including reduced matrix storage, preprocessing calculation stages, and specific pipelining schedules that allow for re-use of hardware during different clock cycles.

[0071] FIG. 5 shows a matrix memory storage scheme that the adaptation engine of cancelation circuitry **112** may employ in order to both reduce the amount of required memory and required calculations during each PA and FIR update iteration in **212** and **214**, respectively. As previously indicated, matrix memory **116** may store the M most recent samples for each of the K kernels, which as shown in FIG. 1 may also be provided to PA filter adaptation circuit **124** and FIR filter adaptation circuit **126** for calculation of $\mathbf{X}_{PA}(t)$, $\mathbf{X}_F(t)$, $\mathbf{R}_{PA}(t)$, $\mathbf{R}_F(t)$, $\boldsymbol{\beta}_{PA}(t)$, and $\boldsymbol{\beta}_F(t)$.

[0072] FIG. 6 shows a more detailed architecture of cancelation circuitry **112**, which may correspond to the matrix storage scheme of FIG. 5. As shown in FIG. 6, kernel generation circuit **114** may provide the kernel outputs $X_t^{(1:K)}$ for time t to the signal path and adaptation engine of cancelation circuitry **112**. In the implementation of FIG. 6, matrix memory **116** may be located in the adaptation engine and accordingly may store the KM samples $X_{t:t-M+1}^{(1:K)}$ for use in updating $\mathbf{W}_{PA}$ and $\mathbf{W}_F$. This $K \times M$ storage is thus reflected in the matrix storage scheme of FIG. 5. Likewise, FIR filter memory **624** and PA filter memory **626** may each store the respective M and K weights of $\mathbf{W}_F$ and $\mathbf{W}_{PA}$, which may similarly be stored at filter memory **610** of the signal path.

[0073] In a general description of the operation of cancelation circuitry **112** as shown in FIG. 6, processing element **606** may apply $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ to $\mathbf{X}(t)$ as previously detailed. In order to reduce hardware complexity and area, processing element **606** may utilize a pipelined approach in order to apply $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ to $\mathbf{X}(t)$. In the example of FIG. 6, processing element **606** may be composed of e.g. 8 complex multipliers and 7 complex adders and may

apply $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ to $\mathbf{X}(t)$ over four cycles (although it is noted that numerous other implementations including a full, single-cycle (i.e. non-pipelined) implementation are within the scope of this disclosure). As noted above, processing element **606** may calculate $e(t)$ as $\mathbf{W}_{PA}^T \mathbf{X}(t) \mathbf{W}_F$, which processing element **606** may perform over e.g. four clock cycles for $M = 24$ and $K = 8$ as

- a. Clock cycle 1: Compute $X_{F,t}(t) = \sum_{i=1}^8 X_t^{(i)}(t) w_{PA,i}$
- b. Clock cycle 2: Compute $EEC_{acc(m=1:8)} = \sum_{i=1}^8 X_{F,t-i+1}(t) w_{F,i}$
- c. Clock cycle 3: Compute $EEC_{acc(m=1:16)} = \sum_{i=9}^{16} X_{F,t-i+1}(t) w_{F,i} + EEC_{acc(m=1:8)}$
- d. Clock cycle 4: Compute $EEC_{acc(m=1:24)} = \sum_{i=17}^{24} X_{F,t-i+1}(t) w_{F,i} + EEC_{acc(m=1:16)}$

where $e(t) = EEC_{acc(m=1:24)}$, $X_{F,t}(t)$ is the sample of $\mathbf{X}_F(t)$ for time t , and $EEC_{acc(m=1:24)}$ denotes the accumulated estimated echo.

[0074] Accordingly, processing element **606** may calculate the most recent sample $X_{F,t}(t)$ of $\mathbf{X}_F(t)$ in clock cycle 1 by applying $\mathbf{W}_{PA}$ to the K samples of $\mathbf{X}(t)$ for time t , i.e. the kernel samples for each of the K kernels. As denoted in clock cycle 1 above, processing element **606** may apply 8 complex multipliers to calculate $X_t^{(i)}(t) w_{PA,i}$ for $i = 1, \dots, 8$ and provide each intermediate product to accumulation circuit **614**, which may sum each intermediate product to obtain $X_{F,t}(t)$. Accumulation circuit **614** may then provide each $X_{F,t}(t)$ to register **602**, which may hold the M samples of $\mathbf{X}_F(t)$, e.g. as a first-in-first-out buffer where the oldest sample from time $t - M$ is pushed out and the newest sample for time t is entered.

Accordingly, the signal path of processing circuit **112** may avoid performing the entire multiplication of $\mathbf{X}_F(t) = \mathbf{W}_{PA}^T \mathbf{X}(t)$ and instead calculate a single sample $X_{F,t}(t)$ at each time t and store the results in register **602**.

[0075] Processing element **606** may then calculate $e(t)$ in clock cycles 2-4 by applying $W_F(t)$ to $X_F(t)$. As detailed above for clock cycles 2-4, processing element **606** may employ the 8 complex multipliers to calculate 8 samples $X_{F,t-i+1}W_{F,i}$ for $i = 1:8$, $i = 9:16$, and $i = 17:24$. Processing element **606** may provide each of the intermediate products $X_{F,t-i+1}W_{F,i}$ to accumulation circuit **616**, which may sum each of the 8 intermediate products for each clock cycle to obtain $e(t)$ as the accumulated estimated echo $EEC_{acc(m=1:24)}$. Accumulation circuit **616** may then provide $e(t)$ to cancelation circuit **122**, which may then subtract $e(t)$ from received signal $y(t)$ to obtain clean signal $z(t)$. Depending on the accuracy of W_{PA} and W_F , clean signal $z(t)$ may be largely free of self-interference.

[0076] As shown in FIG. 6, multiplexers **604** and **612** may provide the appropriate selection of 8 samples from $X_F(t)$, $X(t)$, W_F , and W_{PA} to processing element **606** in accordance with the equations provided above for clock cycles 1-4. The multiplexer selection signals may be provided by FSM **608**, which may utilize a finite state machine logic to cycle through the appropriate multiplexer selection signals.

[0077] As previously detailed above regarding FIGs. 1-4, the adaptation engine of cancelation circuitry **112** may be responsible for adaptively updating W_{PA} and W_F based on $X(t)$ and $z(t)$ in order to enhance the estimation accuracy of $e(t)$ and thus improve the self-interference cancelation offered by cancelation circuitry **112**. As shown in FIG. 6, matrix memory **116** may receive the kernel samples $X_t^{(1:K)}$ at each time t from kernel generation circuit **114** and may subsequently store each of the kernel outputs from the M previous samples, i.e. $X_{t:t-M+1}^{(1:K)}$, as $X(t)$. As illustrated in the matrix storage scheme of FIG. 5, matrix memory **114** may store the entire $K \times M$ entries of $X(t)$.

[0078] Matrix memory **116** may provide the current $X(t)$ to PA update circuit **622** and FIR update circuit **628**, which may be respectively responsible for calculating $X_{PA}(t)$ and $X_F(t)$

in **212a** and **214a** based on $\mathbf{X}(t)$ and $\mathbf{W}_F$ and $\mathbf{W}_{PA}$ provided by FIR filter memory **624** and PA filter memory **626**. As shown in FIG. 6, PA update circuit **622** and FIR update circuit **628** may provide $\mathbf{X}_{PA}(t)$ and $\mathbf{X}_F(t)$ to PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **630**, which may each respectively store the K and M samples of $\mathbf{X}_{PA}(t)$ and $\mathbf{X}_F(t)$.

[0079] PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **630** may then calculate $\mathbf{R}_{PA}(t)$, $\boldsymbol{\beta}_{PA}(t)$, $\mathbf{R}_F(t)$, add $\boldsymbol{\beta}_F(t)$ in **212b** and **214b**. However, as shown in the matrix storage scheme of FIG. 5, PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **630** may avoid calculating and storing the entire $K \times K$ and $M \times M$ matrices for $\mathbf{R}_F(t)$ and $\mathbf{R}_{PA}(t)$ and may instead store only the upper row of $\mathbf{R}_F(t)$ and the upper triangle of $\mathbf{R}_{PA}(t)$, thus conserving considerable storage and reducing the computational demands.

[0080] More specifically, as $\mathbf{R}_{PA}(t)$ is the correlation between each of the kernels of $\mathbf{X}_{PA}(t)$ with one another, $\mathbf{R}_{PA}(t)$ will be a Hermitian matrix, where the lower triangle of off-diagonal elements is equal to the complex conjugate of the upper triangle of off-diagonal elements. Accordingly, instead of calculating the entire $M \times M$ matrix for $\mathbf{R}_{PA}(t)$, PA kernel adaptation circuit **620** may instead in **212b** calculate the upper triangle of $\mathbf{R}_{PA}(t)$ (including the diagonal elements) and consequently only store the upper triangle elements of $\mathbf{R}_{PA}(t)$. For any subsequent calculations involving lower triangle elements of $\mathbf{R}_{PA}(t)$, PA kernel adaptation circuit **620** may take the complex conjugate of the corresponding upper-triangle element in order to obtain the lower triangle element.

[0081] Regarding $\mathbf{R}_F(t)$, as noted above $\mathbf{R}_F(t)$ may be the correlation between each of the taps of $\mathbf{X}_F(t)$ with one another. While $\mathbf{R}_F(t)$ may also be Hermitian, each row of $\mathbf{R}_F(t)$ may share considerable statistical similarities given the relationship between each tap and the fact that the kernel dimension is fixed during update of $\mathbf{W}_F$. In other words, the first row of $\mathbf{R}_F(t)$,

which gives the correlations between the first tap of $\mathbf{X}_F(t)$ and each of the remaining $M - 1$ taps of $\mathbf{X}_F(t)$, may be increasingly similar to the other rows of $\mathbf{R}_F(t)$, which may likewise give the correlations between a given tap of $\mathbf{X}_F(t)$ and the other taps of $\mathbf{X}_F(t)$. As each of the taps are proximate in time, it follows that each row may be approximately similar to the other rows; accordingly, without loss of generality, as opposed to calculating the entire $M \times M$ matrix or even the upper triangle of $\mathbf{R}_F(t)$ FIR delay taps adaptation circuit **630** may instead calculate and store only a single row of $\mathbf{R}_F(t)$, e.g. the first row of $\mathbf{R}_F(t)$. FIR delay taps adaptation circuit **630** may then utilize the single row of $\mathbf{R}_F(t)$ as an approximation when using $R_{F(n,n)}$ and $R_F(t - 1)^{(n)}$. While this may act as an approximation, FIR delay taps adaptation circuit **630** may reduce storage and computational requirements through only storing a single row of $\mathbf{R}_F(t)$.

[0082] Accordingly, in **212d** and **214d** when $\mathbf{R}_{PA}(t)$ and $\mathbf{R}_F(t)$ are needed in determining update factor α , PA kernel adaptation circuit **620** may simply retrieve $R_{PA(n,n)}$ (which is a diagonal element and thus will be explicitly stored as part of the upper triangle) while FIR delay taps adaptation circuit **630** may retrieve $R_{F(1,1)}$ regardless of n . As the actual value of $R_{F(n,n)}$ gives the correlation between the n -th tap and itself, $R_{F(1,1)}$ may serve as a suitable approximation as the correlation of the $n = 1$ tap with itself. As can be seen from the matrix storage scheme illustrated in FIG. 5, this implementation may lead to considerable reductions in both storage and computational demands on the adaptation engine.

[0083] As shown in FIG. 6, PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **630** may each provide index n and weight $w_{PA,n}/w_{F,n}$ to controller/FSM **618**, which may proceed to instruct filter memory **610** to update the corresponding n -th element of $\mathbf{W}_{PA}/\mathbf{W}_F$ with the specified weight $w_{PA,n}/w_{F,n}$. As such may involve a DCD update scheme of only a single bit, controller/FSM **618** may simply identify the weight index n and bit index

m that needs to be flipped, thus prompting filter memory **610** to flip the m -th bit of the n -th weight of $\mathbf{W}_{PA}/\mathbf{W}_F$. Controller/FSM **618** may additionally include switching circuit **128** and thus may be responsible for selecting which of PA filter adaptation circuit **124** or FIR filter adaptation circuit **126** to activate.

[0084] In addition to the matrix storage scheme, the adaptation engine of cancelation circuitry **112** may additionally reduce the computational requirements involved in calculating $\mathbf{X}_{PA}(t)$ and $\mathbf{X}_F(t)$ at PA update circuit **622** and FIR update circuit **628**. FIG. 7A shows an illustration of the matrix multiplication $\mathbf{X}_{PA}(t) = \mathbf{X}(t)\mathbf{W}_F$, where each k -th element of $\mathbf{X}_{PA}(t)$ for $k = 1, \dots, K$ may be the dot product of $\mathbf{W}_F$ with the k -th row $\mathbf{X}(t)$. As denoted in FIG. 7A, the k -th element $X_{PA,k}(t)$ of $\mathbf{X}_{PA}(t)$ may then be given as

$$X_{PA,k}(t) = X_t^{(k)}w_{F,1} + X_{t-1}^{(k)}w_{F,2} + \dots + X_{t-M+1}^{(k)}w_{F,M} \quad (17)$$

where $w_{F,m}$ gives the m -th tap of $\mathbf{W}_F$.

[0085] Accordingly, only the first sum term $X_t^{(k)}w_{F,1}$ will depend on the current input sample at time t while all of the remaining sum terms $X_{t-1}^{(k)}w_{F,2}, X_{t-2}^{(k)}w_{F,3}, \dots, X_{t-M+1}^{(k)}w_{F,M}$ depend on past samples from times $t-1$ and previous. Accordingly, as opposed to requiring the calculation of KM total sum terms required to calculate all of the K elements of $\mathbf{X}_{PA}(t)$ after receiving the most recent $X_t^{(1:K)}$, PA update circuit **622** may pre-calculate and sum the past sum terms $X_{t-1}^{(k)}w_{F,2}, X_{t-2}^{(k)}w_{F,3}, \dots, X_{t-M+1}^{(k)}w_{F,M}$ for $k = 1, \dots, K$ at an earlier time and, after calculating $w_{F,1}X_t^{(1:K)}$, simply add the preprocessing sum terms to $w_{F,1}X_t^{(1:K)}$ to obtain the full $\mathbf{X}_{PA}(t)$. PA update circuit **622** may therefore only need to have sufficient hardware to calculate K sum terms for $X_t^{(1:K)}w_{F,1}$, and thus may re-use this hardware to calculate the other sum terms at an earlier time (which may additionally require memory for K extra terms as

shown for $X_{PA_pre}(t)$ in FIG. 5). While the required computations may be the same, this may substantially reduce the hardware area requirements of PA update circuit 622, thus directly reducing manufacturing costs.

[0086] FIR update circuit 628 may also drastically simplify the calculation of $\mathbf{X}_F(t)$ in 214a. As shown in FIG. 7B, since $\mathbf{W}_{PA}$ is fixed during update of $\mathbf{W}_F$, each of the M elements of $\mathbf{X}_F(t)$ may be time-delayed weighted inputs of one another. Accordingly, FIR update circuit 628 may not need to calculate all M samples of $\mathbf{X}_F(t)$ for each time t ; instead, FIR update circuit 628 may utilize a first-in-first-out buffer approach (similar to as detailed above regarding buffer 602) and calculate $X_{F,t}(t)$, i.e. the sample of $\mathbf{X}_F(t)$ for time t , to be stored with the previous $M - 1$ samples of $\mathbf{X}_F(t)$. The oldest sample of $\mathbf{X}_F(t)$, i.e. from time $t - M$, may then be discarded every time a new sample $X_{F,t}(t)$ is calculated. As such may reduce the required computations from MK multiplications and $M(K - 1)$ additions to K multiplications and $K - 1$ additions, FIR update circuit 628 may save considerably reduce complexity and the required hardware area.

[0087] Accordingly, PA update circuit 622 may employ pre-processing in order to reduce the hardware requirements for calculating $\mathbf{X}_{PA}(t)$ while FIR update circuit 628 may utilize the time-delayed weighted relationship to reduce the computational requirements for calculating $\mathbf{X}_F(t)$. As noted above, these modifications are optional and may or may not be implemented in various aspects of this disclosure.

[0088] Due to the complexity involved in self-interference estimation, a key constraint in many self-interference cancelation designs is hardware area. As noted above, the preprocessing modification of PA update circuit 622 may reduce the number of multiplication and addition hardware elements needed as PA update circuit 622 may only need to perform K multiplications with the current kernel samples $X_t^{(1:K)}$ followed by K additions with each of the K preprocessing sum terms to obtain $\mathbf{X}_{PA}(t)$. PA update circuit 622 may then perform the

$K(M - 1)$ complex multiplications and $K(M - 2)$ complex additions by re-using the same hardware during other clock cycles. The adaptation engine of cancelation circuitry **112** may utilize further hardware computation schedules in order to re-use hardware and thus reduce the hardware requirements of cancelation circuitry **112**. As noted above, such hardware scheduling is optional and thus may or may not be utilized in various aspects of this disclosure.

[0089] FIG. **8** depicts an exemplary hardware computation schedule for PA adaptation circuit **124** (PA update circuit **622**, PA kernel adaptation circuit **620**, and controller/FSM **618**) and FIR adaptation circuit **126** (FIR update circuit **628**, FIR delay taps adaptation circuit **630**, and controller/FSM **618**). Tables **810** and **830** show the total numbers of multipliers and adders (as will be further detailed) that the adaptation engine may utilize for each operation while hardware schedules **820** and **840** depict the computational operations executed by PA adaptation circuit **124** and FIR adaptation circuit **126** during each clock cycle. As shown in hardware schedules **820** and **840**, both PA adaptation circuit **124** and FIR adaptation circuit **126** may utilize 5 clock cycles in order to complete a single update of $\mathbf{W}_{PA}$ and $\mathbf{W}_F$ for a given sample t .

[0090] As shown in PA branch schedule **820**, PA update circuit **622** may calculate the preprocessing sum terms $X_{t-1}^{(k)} w_{F,2}, X_{t-2}^{(k)} w_{F,3}, \dots, X_{t-M+1}^{(k)} w_{F,M}$ for $k = 1, \dots, K$ of $\mathbf{X}_{PA}(t)$ as noted above regarding Equation (17) in ‘hidden’ clock cycles (denoted as shaded cycles in hardware schedule **820**), which may be distributed at any time prior to the start of calculations for sample time t . As denoted in the exemplary context of FIG. **8**, M and K may be set to $M = 24$ and $K = 8$; accordingly, PA update circuit **622** may reuse the same computation hardware over multiple clock cycles. As shown in hardware schedule **820**, PA update circuit **622** may utilize e.g. $2(M - 1)$ complex multipliers and $2(M - 2)$ complex adders in order to support calculation of the preprocessing sum terms $X_{t-1}^{(k)} w_{F,2}, X_{t-2}^{(k)} w_{F,3}, \dots, X_{t-M+1}^{(k)} w_{F,M}$ for

two k indices at a time; e.g. for $k = 1,2$ in clock cycle 0, $k = 3,4$ in clock cycle 1, $k = 5,6$ in clock cycle 2, and $k = 7,8$ in clock cycle 3. The amount of computation hardware may be adjusted to various different with varying tradeoffs between the number of clock cycles and hardware area.

[0091] PA update circuit **622** may then apply a subset of the same complex multipliers and adders to calculate the sum terms for the kernels of time t , $X_t^{(1:K)}$, as $w_{F,1}X_t^{(1:K)}$, during clock cycle 0 and add the sum terms to the preprocessing sum terms to obtain $\mathbf{X}_{PA}(t)$ (**212a**). As each element $\mathbf{X}_{PA}(t)$ is needed to calculate $\beta_{PA}(t)$, PA kernel adaptation circuit **620** may begin the update of $\beta_{PA}(t)$ (**212b**) in clock cycle 1. In order to reduce the hardware requirements, PA kernel adaptation circuit **620** may update $\beta_{PA}(t)$ over both clock cycles 1 and 2, thus only requiring 4 complex multipliers and 4 complex adders as denoted in table **810** while still avoiding any data dependencies. PA kernel adaptation circuit **620** may also begin identifying $\beta_{PA,max}$ and n (**212c**) in clock cycle 2 with the elements of $\beta_{PA}(t)$ calculated in clock cycle 1 and completing the identification of $\beta_{PA,max}$ and n in clock cycle 3 with the remaining elements of $\beta_{PA}(t)$ calculated in clock cycle 2. PA kernel adaptation circuit **620** may simultaneously begin updating $\mathbf{R}_{PA}(t)$ (**212b**) in clock cycle 3, which may be completed in clock cycle 4 (**212e**). As previously indicated, PA kernel adaptation circuit **620** may update the diagonal elements of $\mathbf{R}_{PA}(t)$ in the clock cycle 3 and the non-diagonal elements of $\mathbf{R}_{PA}(t)$ in clock cycle 4 in order to avoid data dependency issues related to the use of $R_{PA(n,n)}$ in identifying $\beta_{PA,max}$.

[0092] While completing update of $\mathbf{R}_{PA}(t)$ in clock cycle 4, PA kernel adaptation circuit **620** may concurrently perform the DCD update scheme in order to identify α , i.e. the bit of $w_{PA,n}$ that will be flipped. PA kernel adaptation circuit **620** may provide n and α to controller/FSM **618**, which may complete the update of $\mathbf{W}_{PA}$ by updating $w_{PA,n}$ according to

α during the next clock cycle, which PA update circuit **622** may concurrently utilize as clock cycle 0 to calculate $w_{F,1}X_{t+1}^{(1:K)}$ for the next time $t + 1$.

[0093] FIR adaptation circuit **126** may similarly utilize a specific 5-clock cycle as shown in hardware schedule **840**. As FIR update circuit **628** may only need to perform M multiplications and $M - 1$ additions to calculate $\mathbf{X}_F(t)$ (due to the time-delayed weighted relationship noted above), FIR adaptation circuit **126** may not need to perform any preprocessing stages in hidden clock cycles. Accordingly, FIR update circuit **628** may apply M complex multipliers and $M - 1$ complex adders in clock cycle 0 to calculate $\mathbf{X}_F(t)$ (**214a**). FIR delay taps adaptation circuit **630** may similarly use a staggered clock cycle scheme in clock cycles 1-4 in order to update $\beta_F(t)$, $\mathbf{R}_F(t)$, and identify n and $\beta_{F,max}$ (**214b** and **214c**). FIR delay taps adaptation circuit **630** may then apply the DCD scheme to determine α (**214d**) in clock cycle 4 and provide n and α to controller/FSM **618** for update of $w_{F,n}$ in the following clock cycle 0 of the next update. FIR delay taps adaptation circuit **630** may likewise calculate the diagonal elements of $\mathbf{R}_F(t)$ in clock cycle 3 and the non-diagonal elements of $\mathbf{R}_F(t)$ in clock cycle 4 in order to avoid data dependency issues.

[0094] The schedules and hardware resource tables depicted in FIG. **8** may thus provide an efficient implementation that offers an advantageous tradeoff between latency and hardware area. Numerous other schedules and hardware architectures are also within the scope of this disclosure.

[0095] The individual components of cancelation circuitry **112** detailed above may be structurally realized/embodyed as hardware logic, e.g. as one or more integrated circuits or FPGAs, as software logic, e.g. as one or more processors executing program code that defining arithmetic, control, and I/O instructions stored in a non-transitory computer-readable storage medium, or as a combination of hardware and software logic. Accordingly, while the individual components of cancelation circuitry **112** are depicted separately in FIG. **1** and FIG.

6, this depiction merely serves to highlight the operation of cancelation circuitry 112 on a functional level; consequently, one or more of the components of cancelation circuitry 112 may be integrated into a common hardware and/or software element.

[0096] Without loss of generality, FIGs. 9-11 illustrate exemplary internal configurations of PA update circuit 622, FIR update circuit 628, PA kernel adaptation circuit 622, and FIR delay taps adaptation circuit 630.

[0097] FIG. 9 shows the internal computational and memory elements of PA update circuit 622. Corresponding to table 810, PA update circuit 622 may include $2(M - 1)$ complex multipliers and $2(M - 2)$ complex adders for the preprocessing $\mathbf{X}_{PA}(t)$ calculations and K complex multipliers and complex adders for the current $\mathbf{X}_{PA}(t)$ calculation. PA update circuit 622 may additionally include the K element memories for $\mathbf{X}_{PA_pre}(t)$ and $\mathbf{X}_{PA}(t)$ as shown in FIG. 9. FIR update circuit 628 shown in FIG. 10 may include the M element $\mathbf{X}_F(t)$ memory and K complex multipliers and $K - 1$ complex adders for $\mathbf{X}_F(t)$.

[0098] FIG. 11 shows correlation update circuit 1110 and cross-correlation update circuit 1120. As the correlation and cross-correlations of PA update iteration 212 and FIR update iteration 214 may be identical, both PA kernel adaptation circuit 620 and FIR delay taps adaptation circuit 630 may include multiple correlation update circuits 1110 and cross-correlation update circuits 1120 in order to perform the calculations needed to obtain $\mathbf{R}_{PA}(t)$, $\boldsymbol{\beta}_{PA}(t)$, $\mathbf{R}_F(t)$, and $\boldsymbol{\beta}_F(t)$. The number of correlation update circuits 1110 and cross-correlation update circuits 1120 included in both PA kernel adaptation circuit 620 and FIR Delay taps adaptation circuit 630 may depend on the M , K , and various factors related to clock cycle/pipeline schedules and hardware area. As shown in FIG. 11, both correlation update circuit 1110 and cross-correlation update circuit 1120 may include a local FSM to direct the clock cycle scheduling, e.g. in accordance with a hardware schedule such as shown in FIG. 8), a set of processing elements (PE) composed of either correlation PEs 1130 or

cross-correlation PEs **1140**, a matrix memory for holding $\mathbf{R}(t)/\boldsymbol{\beta}(t)$, and a scale control circuit. The capacity and contents of the matrix memories may correspond to whether correlation update circuit **1110** and cross-correlation update circuit **1120** is assigned to PA kernel adaptation circuit **620** or FIR delay taps adaptation circuit **830**.

[0099] The number of PEs included in each of correlation update circuit **1110** and cross-correlation update circuit **1120** may be dictated by the number of matrix elements that need to be calculated in parallel. As each correlation PE **1130** and cross-correlation PE **1140** may be configured to update a single element of $\mathbf{R}(t)/\boldsymbol{\beta}(t)$, PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **630** may each include multiple correlation update circuits **1110** and multiple cross-correlation update circuits **1120** in order to support calculation of multipole updates to $\mathbf{R}(t)/\boldsymbol{\beta}(t)$ in parallel. Such may consequently depend on the desired clock cycle scheduling, such as shown in FIG. **8**. As detailed in table **810**, PA kernel adaptation circuit **620** may include e.g. 18 correlation update circuits **1110** and e.g. 4 cross-correlation update circuits **1120** for $K = 8$ and $M = 24$ in order to enable a 5 clock cycle update iteration. Such may allow PA kernel adaptation circuit **620** to calculate the 36 total required complex multiplications and additions to update $\mathbf{R}_{PA}(t)$ spread over two clock cycles (18 complex multiplications and additions per clock cycle) and to calculate the 8 total required complex multiplications and additions to update $\boldsymbol{\beta}_{PA}(t)$ spread over two clock cycles (4 per complex multiplications and additions per clock cycle).

[0100] As noted in table **830**, FIR delay taps adaptation circuit **630** may need to include e.g. 12 correlation update circuits **1110** and e.g. 12 cross-correlation update circuits **1120**. Such may thus allow FIR delay taps adaptation circuit **630** to calculate the 24 total required complex multiplications and additions to update $\mathbf{R}_F(t)$ spread over two clock cycles (12 complex multiplications and additions per clock cycle) and to calculate the 12 total required

complex multiplications and additions to update $\beta_{FIR}(t)$ spread over two clock cycles (12 per complex multiplications and additions per clock cycle).

[0101] Due to the equivalence in the calculations for PA and FIR updates and the decoupled operation mode, PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **630** may in certain aspects of this disclosure be configured to share the same correlation update circuit **1110** and cross-correlation update circuit **1120**, i.e. to re-use the same hardware at different times depending on which of PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **630** is currently updating W_{PA} or W_F , respectively. As PA kernel adaptation circuit **124** and FIR adaptation circuit **126** operate separately, such may only require that the shared correlation update circuit **1110** and shared cross-correlation update circuit **1120** include the proper number of correlation PEs **1130** and cross-correlation PEs **1140** to support both PA and FIR updates, e.g. 18 correlations PEs **1130** and 12 cross-correlation PEs **1140**.

[0102] As shown in FIG. 11, each correlation PE **1130** may receive two samples of $X_{PA}(t)/X_F(t)$ to multiply. In order to simplify the input data switching overhead, correlation update circuit **1110** may utilize a smart memory indexing and cycle allocation scheme. FIG. 12 illustrates the smart memory indexing and cycle allocation scheme for an exemplary implementation where correlation update circuit **1110** includes 18 correlation PEs **1130** and utilizes two cycles to calculate $R_{PA}(t)$ for $K = 8$ and $M = 24$. As denoted by the shading of table **1200**, correlation update circuit **1110** may optimize the calculation of each element of $R_{PA}(t)$ by ensuring that each correlation PE **1130** calculates an element of $R_{PA}(t)$ in the first cycle that is in the same column or row as the element of $R_{PA}(t)$ calculated in the second cycle. As either the column or row remains the same, one input of each correlation PE **1130** may remain the same for both the first and second cycles (as this input will be same element of $X_{PA}(t)$). Such is shown by the numerical indices of table **1200**, which indicate which of the

18 total correlation PEs **1130** will calculate the corresponding element of $R_{PA}(t)$.

Additionally, correlation update circuit **1110** may calculate the diagonal elements of $R_{PA}(t)$ in the first cycle, thus eliminating the data dependency that can occur in the pipeline schedule shown in FIG. 8. Correlation update circuit **1110** may thus utilize the memory index shown in table **1210** to index the memory for $R_{PA}(t)$.

[0103] The optional enable signal en at each of correlation PEs **1130** and cross-correlation PEs **1140** may allow for configurable operation, e.g. for cancelation circuitry **112** to be utilized for any $K' \leq K$ kernels and $M' < M$ taps. Accordingly, the enable signal en may be utilized to disable all unused correlation PEs **1130** and cross-correlation PEs **1140**, e.g. $K - K'$ or $M - M'$ unused correlation PEs **1130** and cross-correlation PEs **1140**.

[0104] While optional, the scale control circuits may be included in order to maintain the stability between W_{PA} and W_F . As both dimensions are decoupled in cancelation circuitry **112**, it may be possible for the weights to become biased towards either PA or FIR, i.e. where one of W_{PA} and W_F becomes much stronger than the other. Accordingly, the scale control circuitry may be configured to detect if the weights become biased towards one of W_{PA} and W_F and, if so, apply a scale factor in order to counter the bias.

[0105] Each of PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **830** may additionally include DCD circuit **1310** as shown in FIG. 13, which may be configured to execute the DCD update process to identify α (**212d/214d**). As for the other modifications introduced above, this internal configuration of DCD circuit **1310** is exemplary and other realizations are within the scope of this disclosure. While mathematically expressed as a for loop in **212d** and **214d**, DCD circuit **1310** may be implemented with a parallel architecture, where DCD circuit **1310** may evaluate each value of $\alpha = 2^0, 2^1, \dots, 2^{M_b-1}$ in parallel with e.g. one comparator per hypothesis to evaluate the comparison $|\beta_{max}| > \frac{\alpha}{2} R_{(n,n)}$ of **212d/214d**.

As only one comparator may be needed, this may offer a significant tradeoff between

hardware and throughput, thus considerably favoring a parallel architecture. DCD circuit **1310** may also be configurable to support up to M_b bits as the individual α comparators may be left unused.

[0106] PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **830** may additionally include maximum cross-correlation detection circuit **1410** shown in FIG. **14**, which may be configured to identify β_{max} for $\beta_{PA}(t)$ and $\beta_F(t)$ in **212c/214c** by identifying the element of $\beta_{PA}(t)/\beta_F(t)$ with the largest real or imaginary part using the comparator and multiplexer logic shown in FIG. **14**. Maximum cross-correlation detection circuit **1410** may additionally preserve the sign of β_{max} in order to update w_n in the proper direction. Generally speaking, PA kernel adaptation circuit **620** and FIR delay taps adaptation circuit **830** may each thus be composed of correlation update circuit **1110**, cross-correlation update circuit **1120**, DCD circuit **1310**, and maximum cross-correlation detection circuit **1410**.

[0107] Accordingly, any of the various modifications noted above may be implemented as part of cancelation circuitry **112**, which may offer a variety of advantages including reduced computational demands and latency, reduced hardware area, reduced memory requirements, etc. Cancelation circuitry **112** may thus provide a highly effective self-interference cancelation system that may be realizable with current technology. Although M and K may be scaled to any values, $M = 24$ and $K = 8$ may present an attractive balance between estimation accuracy and complexity. When combined with various of the optional modifications noted above, such may provide a system that has as much as an 80% hardware reduction compared to the existing ‘non-decoupled’ solution (i.e. with updates of $\mathbf{W}$ of dimension $K \times M$). As noted above, cancelation circuitry **112** may additionally be implemented as a ‘configurable’ design where cancelation circuitry **112** may be realized to support up to M taps and K kernels but can be dynamically configured to perform self-interference cancelation with any $M' \leq M$ taps and $K' \leq K$ kernels.

[0108] The self-interference cancelation scheme detailed thus far may utilize a memoryless model of power amplifier **104** (where the FIR weights $\mathbf{W}_F$ are designed to model the various transmit and receive FIR filters in leakage path **130**). However, although power amplifiers may generally be designed to be memoryless, power amplifiers may in practical scenarios exhibit unintended memory effects. These memory effects may be caused by imperfections such as charging and discharging junctions, thermal variations, and other physical phenomena and may cause time-dependencies on the power amplifier output. Accordingly, cancelation circuitry **112** may be slightly modified in order to model such PA memory effects by adding $L > 1$ taps for each kernel. FIG. **15** shows an exemplary configuration in which $\mathbf{W}_{PA}$ is implemented as $K \times L$ (with $L > 1$) matrix where each kernel is thus modeled over L taps (as opposed to the memoryless design previously detailed with $L = 1$). Using $\mathbf{W}_{PA}$ defined as a long vector

$$\mathbf{W}_{PA} = \begin{bmatrix} w_{PA,1,1} \\ w_{PA,1,2} \\ \dots \\ w_{PA,1,L} \\ w_{PA,2,1} \\ \dots \\ w_{PA,2,L} \\ \dots \\ w_{PA,K,1} \\ \dots \\ w_{PA,K,L} \end{bmatrix} \quad (18)$$

PA filter circuit **118** may apply $\mathbf{W}_{PA}$ to $\mathbf{X}(t)$ to obtain $\mathbf{X}_F$ defined as

$$\mathbf{X}_F = [X_{F,t}, X_{F,t-1}, \dots, X_{F,t-M+1}] \quad (19)$$

where each entry $X_{F,j}$ is calculated as

$$X_{F,j} = W_{PA}^H \begin{bmatrix} X_t^{(1)} \\ X_{t-1}^{(1)} \\ \dots \\ X_{t-L+1}^{(1)} \\ X_t^{(2)} \\ X_{t-1}^{(2)} \\ \dots \\ X_{t-L+1}^{(2)} \\ \dots \\ X_t^{(K)} \\ X_{t-1}^{(K)} \\ \dots \\ X_{t-L+1}^{(K)} \end{bmatrix} \quad (20)$$

[0109] PA filter adaptation circuit **124** may thus be configured to adapt each of the KL weights of W_{PA} in order to fully approximate the L taps of the PA response.

[0110] FIG. **16** shows method **1600** of method of performing interference cancelation. As shown in FIG. **16**, method **1600** includes separately applying a kernel dimension filter and a delay tap dimension filter to an input signal for an amplifier to obtain an estimated interference signal (**1610**), subtracting the estimated interference signal from a received signal to obtain a clean signal (**1620**), and alternating between updating the kernel dimension filter and updating the delay tap dimension filter using the clean signal (**1630**).

[0111] FIG. **17** shows method **1700** of performing interference cancelation. As shown in FIG. **17**, method **1700** includes obtaining one or more subsignals derived from an input signal for an amplifier (**1710**), each of the one or more subsignals representing a non-linear component of an amplifier response and composed of a plurality of delay taps, separately applying a first filter and a second filter to the one or more subsignals to obtain an estimated

interference signal (1720), wherein the first filter approximates the amplifier response over the one or more subsignals and the second filter approximates the amplifier response over the plurality of delay taps (1730), subtracting the estimated interference signal from a received signal to obtain a clean signal, and alternating between updating the first filter and updating the second filter using the clean signal (1740).

[0112] In one or more further exemplary aspects of the disclosure, one or more of the features described above in reference to FIGS. 1-15 may be further incorporated into method 1600 and/or 1700. In particular, method 1600 and/or 1700 may be configured to perform further and/or alternate processes as detailed regarding cancelation circuitry 112 and/or communication circuitry 100.

[0113] Furthermore, the implementations of cancelation circuitry 112 shown in FIGs. 1 and 4 include only a single signal path, which may be appropriate for a Single Input Single Output (SISO) configuration with one receive chain and one transmit chain. FIGs. 18-20 detail several different configurations that may be employed in MIMO cases having multiple transmit and receive chains. Each receive chain may therefore need to include self-interference cancelation circuitry in order to remove leakage from received signals; however, as each of the transmit chains may produce leakage onto each receive chain, the self-interference cancelation circuitry for each receive chain may need to estimate and cancel leakage from each transmit chain. Accordingly, each of the receive chains may need to have a different $\mathbf{W}_F$ and $\mathbf{W}_{PA}$ (assuming use of a bilinear cancelation architecture) for each transmit chain. For example, for an $N \times N$ MIMO scheme each of the N receive chain may need to have N FIR filters $\mathbf{W}_{F,1:N}$ and N PA filters $\mathbf{W}_{PA,1:N}$, i.e. one per transmit chain, in order to produce estimated interference signals $e_{1:N}(t)$ for summation to obtain the combined interference signal $e(t)$ (where each receive chain produces a separate $e(t)$ from separate $e_{1:N}(t)$). As each transmit chain utilizes a different power amplifier and has a different data5

stream, each receive chain may additionally generate a separate kernel matrix $\mathbf{X}_{1:N}(t)$ for each transmit chain to utilize along with the clean signal $z(t)$ to calculate updates to $\mathbf{W}_{F,1:N}$ and $\mathbf{W}_{PA,1:N}$. Such may be extended to MIMO schemes of any dimensions with any number of transmit receive chains.

[0114] Such MIMO applications may unavoidably increase the signal path circuitry and memory requirements for an employed SIC scheme due to the need to store and apply a separate $\mathbf{W}_{F,1:N}$, $\mathbf{W}_{PA,1:N}$, and $\mathbf{X}_{1:N}(t)$ per each of the N receive chains; accordingly each receive chain may have a separate kernel generation circuit **114**, matrix memory **116**, PA filter circuit **118**, and FIR filter circuit **120** for each of the N transmit chains. FIG. **18** illustrates such a configuration for cancelation circuitry **1800** that may be employed in the case of two receive chains, such as e.g. for a 2x2 MIMO scheme. As shown in FIG. **18**, cancelation circuitry **1800** may therefore include multiple ‘duplicates’ of certain components in order to perform the estimation of the leakage signals from each transmit chain, such as kernel generation circuits **114a** and **114b**, matrix memories **116a** and **116b**, PA filter circuits **118a** and **118b**, FIR filter circuits **120a** and **120b**.

[0115] As the power amplifiers have separate input transmit signals in $s_1(t)$ and $s_2(t)$ and additionally utilize different antennas, cancelation circuitry **1800** may utilize a separate kernel generation circuit for each power amplifier in kernel generation circuits **114a** and **114b**.

Accordingly, kernel generation circuits **114a** and **114b** may each both receive the input transmit signals $s_1(t)$ and $s_2(t)$ at each time t for the power amplifiers of both transmit chains and generate kernel samples $X_t^{(1,1:K)}$ and $X_t^{(2,1:K)}$ to provide to matrix memories **116a** and **116b**. Matrix memories **116a** and **116b** may each retain the past M kernel samples for each transmit chain as $K \times M$ matrices $\mathbf{X}_1(t)$ and $\mathbf{X}_2(t)$, which each may provide respectively to PA filter circuits **118a** and **118b** for preprocessing. PA filter circuits **118a** and **118b** may then apply the respective PA filters $\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$ as $\mathbf{W}_{PA,1}^T \mathbf{X}_1(t)$ and

$\mathbf{W}_{PA,2}^T \mathbf{X}_2(t)$, which may each be a $K \times 1$ weight vector that models the K kernels of the respective power amplifier (and each may contain respective PA weights $w_{PA,k}^{(b)}$ for $k = 1, \dots, K$, where $b = 1$ or 2 depending on which PA filter). PA filter circuits **118a** and **118b** may then provide the resulting $\mathbf{X}_{PA,1}(t)$ and $\mathbf{X}_{PA,2}(t)$ to FIR filter circuits **120a** and **120b**. FIR filter circuits **120a** and **120b** may then apply respective FIR filters $\mathbf{W}_{F,1}$ and $\mathbf{W}_{F,2}$ (and each may contain respective FIR weights $w_{F,m}^{(b)}$ for $m = 1, \dots, M$) where $b = 1$ or 2 depending on which FIR filter) as $\mathbf{W}_{PA,1}^T \mathbf{X}_1(t) \mathbf{W}_{F,1}$ and $\mathbf{W}_{PA,2}^T \mathbf{X}_2(t) \mathbf{W}_{F,2}$ to obtain the estimated interference signals $e_1(t)$ and $e_2(t)$ resulting from the first and second transmit chains, respectively. Summation circuit **1802** may then add both estimated interference signals to obtain the overall estimated interference signal $e(t)$ (for the first receive chain) as $e(t) = e_1(t) + e_2(t)$. Cancellation circuit **122** may then subtract $e(t)$ from the receive signal $y(t)$ for the first receive chain to obtain clean signal $z(t)$, which may largely contain exclusively desired data assuming accurate filter models.

[0116] FIG. 18 may also include adaptation circuits **124a/126a** and **124b/126b** (respectively composed of PA filter adaptation circuit **124a**/FIR filter adaptation circuit **126a** and PA filter adaptation circuit **124b**/FIR filter adaptation circuit **126b**), which may each perform the adaptation techniques detailed above (e.g. a bilinear technique such as RLS-DCD, further detailed regarding FIGs. 2 and 3 for PA update iteration **212** and FIR update iteration **214**) for the respective transmit chains to continuously update the FIR and PA filters to enable accurate interference estimation and cancelation. As the adaptations are separate, adaptation circuits **124a/126a** and **124b/126b** may run independently from one another to update the FIR and PA filters (where FIR and PA update iterations for each are alternated by fixing one dimension and adapting the other).

[0117] In order to reduce the amount of hardware and requisite area for cancelation circuitry **1800**, cancelation circuitry **1800** may share adaptation circuitry instead of employing separate adaptation circuits each dedicated to a respective transmit chain. Accordingly, instead of utilizing separate adaptation circuits **124a/126a** and **124b/126b**, cancelation circuitry **1800** may instead utilize only one adaptation circuit, e.g. adaptation circuit **124a/126a**, that may alternate between updating $W_{PA,1}$ and $W_{F,1}$ for the first transmit chain and $W_{PA,2}$ and $W_{F,2}$ for the second transmit chain and thus handle the adaptation responsibilities for the filters of both transmit chains. For example, adaptation circuit **124a/126a** may update $W_{PA,1}$ and $W_{F,1}$ (based on $X_1(t)$ and $z(t)$, which PA adaptation circuit **124a** and FIR adaptation circuit **126a** may additionally update by alternating in a decoupled switching fashion) for a predefined cycle of samples or until a convergence criteria is reached (e.g. for $\beta_{F,1}(t)$ and/or $\beta_{PA,1}(t)$ for the first transmit chain) before switching to updating $W_{PA,2}$ and $W_{F,2}$ (based on $X_2(t)$ and $z(t)$, which PA adaptation circuit **124a** and FIR adaptation circuit **126a** may additionally update by alternating in a decoupled switching fashion) for a predefined cycle of samples or until a convergence criteria is reached (e.g. for $\beta_{F,2}(t)$ and/or $\beta_{PA,2}(t)$ for the first transmit chain). As complete convergence may occur after several hundred samples, such may enable cancelation circuitry **1800** to employ a single adaptation engine in adaptation circuit **124a/126a** to handle filter adaptation for interference cancelation of multiple transmit chains, thus considerably reducing hardware.

[0118] As each of the N receive chains may require an instance of cancelation circuitry **1800**, this adaptation engine sharing may still require N adaptation circuits **124a/126a**, where each adaptation circuit **124a/126a** handles the leakage estimation for N transmit chains for the respective receive chain. Accordingly, in order to further reduce hardware requirements, the self-interference cancelation architecture may expand the adaptation engine sharing between receive chains. FIG. **19** shows an example of such in which a single adaptation engine (e.g.

adaptation circuit **124a/126a**, which contains at least one PA adaptation circuit **124a** and one FIR adaptation circuit **126a**) may handle the adaptation for multiple instances of cancelation circuitry **1800** where each instance of cancelation circuitry **1800** is responsible for leakage estimation and cancelation for one receive chain. Accordingly, the adaptation engine may alternate between adapting $\mathbf{W}_{PA,n}$ and $\mathbf{W}_{F,n}$, $n = 1, \dots, N$, for each of the N receive chains, e.g. with switching based on a predefined number of samples or a convergence criteria, and thus may enable a significant reduction in hardware area. While such adaptation resource sharing may necessarily produce some performance degradation compared to dedicated adaptation resources, this may be offset by the resulting hardware area reductions.

[0119] As shown in FIG. **18**, cancelation circuitry **1800** may employ separate PA and FIR filters to model the leakage from each transmit chain at the receive chain associated with cancelation circuitry **1800**. In the case of PA filters $\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$, such may be necessary in order to effectively model the separate power amplifiers in each transmit chain; for example, even if the same power amplifier model or type is used for both power amplifiers, different filters may be needed to model each power amplifier due to the differing input streams in $s_1(t)$ and $s_2(t)$ and differing antennas. However, in the case of FIR filters $\mathbf{W}_{F,1}$ and $\mathbf{W}_{F,2}$, cancelation circuitry **1800** may be modified to utilize a single FIR filter $\mathbf{W}_{F,1}$ to jointly model the FIR dimension of both transmit chains. More specifically, if the transmit chain FIR filters that $\mathbf{W}_{F,1}$ and $\mathbf{W}_{F,2}$ are intended to model are similar (e.g. have a similar response), cancelation circuitry **1800** may collectively model the FIR filters of both transmit chains with only a single FIR filter $\mathbf{W}_F$. In other words, if the signal paths for each transmit chain utilize separate FIR filters or FIR filters with different weights, the taps of each transmit chain may need to be estimated separately from one another. However, if each transmit chain uses similar FIR filters, the interference cancelation circuitry may be simplified considerably

as the similar FIR filters may be estimated jointly. Such may further reduce hardware area requirements and/or reduce the degree of resource sharing.

[0120] FIG. 20 shows cancelation circuitry 2000, which may utilize a single joint FIR filter 120c and thus may be suited for use in scenarios where the N transmit chains (e.g. two in the context of FIG. 2) contain similar FIR filters and are thus suitable for joint estimation.

Cancelation circuitry 2000 may thus be deployed in each receive chain in order to estimate and cancel leakage resulting from the transmit chains.

[0121] As opposed to separately estimating the tap filter weights for each transmit chain, cancelation circuitry 2000 may instead jointly estimate the tap filter weights for FIR filter circuit 120c, which may function equivalently to FIR filter circuit 120 of FIG. 1 in the SISO case. Cancelation circuitry 2000 may apply $\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$ to $\mathbf{X}_1(t)$ and $\mathbf{X}_2(t)$ to obtain $\mathbf{X}_{F,1}(t)$ and $\mathbf{X}_{F,2}(t)$, which cancelation circuitry 2000 may sum at summation circuit 2002 to obtain $\mathbf{X}_{FIL}(t)$. Cancelation circuitry 2000 may then apply joint FIR filter $\mathbf{W}_F$ to $\mathbf{X}_{FIL}(t)$ to obtain $e(t)$ for subsequent cancelation from $y(t)$ to obtain clean signal $z(t)$.

[0122] FIR filter adaptation circuit 126c of adaptation circuit 124a/124b/126c may function equivalently to FIR filter adaptation circuit 126 as detailed in FIG. 4 regarding FIR update iteration 214 using $\mathbf{X}_{FIL}(t)$ and $z(t)$ and accordingly calculate the correlation $\mathbf{R}_{FIL}(t)$ of $\mathbf{X}_{FIL}(t)$ and the cross-correlation $\boldsymbol{\beta}_{FIL}(t)$ of $\mathbf{X}_{FIL}(t)$ and $z(t)$ to identify updates to weights of $\mathbf{W}_F$. As the FIR filters of each transmit chain are assumed to be similar, the updates calculated by FIR filter adaptation circuit 126c may be able to accurately characterize the leakage path for both transmit chains and accordingly may continue to offer effective interference cancelation.

[0123] As the power amplifiers for each transmit chain are different (and may have different input signals and antennas), the non-linearities of each PA may need to be estimated separately. Adaptation circuit 124a/124b/126c may therefore need to adapt both FIR filters

$\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$ (one per transmit chain) and thus may contain two PA adaptation circuits **124a** and **124b**, which may each adapt PA filters $\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$ equivalent to PA filter adaptation circuit **126** as detailed in FIG. 3 regarding PA update iteration **212**, i.e. by applying joint FIR filter $\mathbf{W}_F$ to filter matrices $\mathbf{X}_1(t)$ and $\mathbf{X}_2(t)$ to respectively obtain $\mathbf{X}_{PA,1}(t)$ and $\mathbf{X}_{PA,2}(t)$ and updating $\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$ accordingly. Similar as detailed above regarding alternating between updating the PA and FIR dimensions, adaptation circuit **124a/124b/126c** may alternate between updating joint FIR filter $\mathbf{W}_F$ and PA filters $\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$, where adaptation circuit **124a/124b/126c** may fix $\mathbf{W}_F$ while updating $\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$ for multiple input samples before switching to fix $\mathbf{W}_{PA,1}$ and $\mathbf{W}_{PA,2}$ while updating $\mathbf{W}_F$.

[0124] As cancelation circuitry **2000** may utilize a single joint FIR filter $\mathbf{W}_F$ and accompanying FIR filter circuit **120c**, cancelation circuitry **2000** may significantly reduce the amount of updates compared to both cancelation circuitry **1800** and other adaptation schemes, i.e. $M + 2K$ updates per receive chain (with two transmit chains) compared to $2(M + K)$ updates and $2MK$ updates. Cancelation circuitry **2000** may thus be suitable for deployment in architectures where the FIR filters of each transmit chain are similar. Returning to the context of FIG. 19, adaptation circuit **124a/124b/126c** may additionally be utilized for adaptation resource sharing between multiple receive chains, where adaptation circuit **124a/124b/126c** may alternate between updating $\mathbf{W}_{PA,n}$ and $\mathbf{W}_{F,n}$, $n = 1, \dots, N$, for each of the N receive chains, e.g. with switching based on a predefined number of samples or a convergence criteria, and thus may enable a significant reduction in hardware area.

[0125] In particular for MIMO applications, stability of the cancelation may become an issue. Such may be exacerbated in the case of cancelation circuitry **1800** due to the parallel operation of two adaptation blocks, which may result in the adaptation blocks struggling between one another and becoming unbalanced. Additionally, the bilinear architecture may allow for a greater dynamic range of $\mathbf{W}_F$ and $\mathbf{W}_{PA}$ compared to the conventional

‘Hammerstein’ architecture. While the bilinear algorithms detailed herein may perform effectively when $\mathbf{W}_F$ and $\mathbf{W}_{PA}$ are balanced, ‘peaky’ or ‘noisy’ samples may drive the adaptation out of reasonable range and thus impair performance. Accordingly, either the FIR weights of $\mathbf{W}_F$ or the PA weights of $\mathbf{W}_{PA}$ may become much larger than the other dimension; while the resulting mathematical calculations may generally produce the same $e(t)$, noisy samples may produce instability (as one of $\mathbf{W}_F$ or $\mathbf{W}_{PA}$ will be very large). Additionally, fixed-point numerical representations (e.g. in the nature of the bit-level inputs shown in FIG. 6) may eventually run into overflow problems when one of $\mathbf{W}_F$ or $\mathbf{W}_{PA}$ grows to be too large. Accordingly, the update of correlation matrices $\mathbf{R}_{PA}(t)$ and $\mathbf{R}_F(t)$ in **212b/212e** and **214b/214e** may be modified in order to introduce a regularization factor, which may counter such imbalances between $\mathbf{W}_F$ and $\mathbf{W}_{PA}$ and ensure that cancelation performance remains stable. Specifically, instead of the correlation updates detailed in Equations (8) and (13), PA filter adaptation circuit **124** and FIR filter adaptation circuit **126** may update $\mathbf{R}_{PA}(t)$ and $\mathbf{R}_F(t)$ as

$$\mathbf{R}_{PA}(t) = \lambda \mathbf{R}_{PA}(t-1) + \mathbf{X}_{PA}^H(t) \mathbf{X}_{PA}(t) + \epsilon_{PA} \mathbf{I} \quad (21)$$

$$\mathbf{R}_F(t) = \lambda \mathbf{R}_F(t-1) + \mathbf{X}_F^H(t) \mathbf{X}_F(t) + \epsilon_F \mathbf{I} \quad (22)$$

where ϵ_{PA} and ϵ_F are small positive numbers that serve as a regularization factor and $\mathbf{I}$ is the identity matrix.

[0126] By including regularization factors ϵ_{PA} and ϵ_F in the correlation update, self-interference cancelation schemes (cancelation circuitry **112**, cancelation circuitry **1800**, and/or cancelation circuitry **2000**) may counter imbalances between the PA and FIR dimension filters. Regularization factors ϵ_{PA} and ϵ_F may need to carefully selected since the strength

thereof will determine the strength of the regularization effect; for example, if ϵ_{PA} or ϵ_F is too large, convergence will suffer. Although both ϵ_{PA} and ϵ_F are design parameters that may vary on a case-by-case basis, ϵ_{PA} and ϵ_F may be in the order of e.g. 2^{-4} or 2^{-8} (although such is exemplary and any value could be used) and may need to be selected based on information about the received signal and/or the expected strength of the interference. Although SISO architectures may additionally utilize Equations (21) and (22), regularization may be more important in MIMO architectures. Different regularization factors ϵ_{PA} and ϵ_F may be chosen for the cancelation circuitry of each receive chain (including $\epsilon \mathbf{I} = \mathbf{0}$, i.e. a zero matrix that provides no regularization) and may be adjusted in order to ensure update stability.

[0127] FIG. 21 shows method **2100** of performing interference cancelation. As shown in FIG. 21, method **2100** includes applying a first kernel dimension filter to a first input signal to estimate a first kernel dimension interference signal derived from a first amplifier (**2110**), applying a second kernel dimension filter to a second input signal to estimate a second kernel dimension interference signal derived from a second amplifier (**2120**), applying a joint delay tap dimension filter to a combination of the first kernel dimension interference signal and the second kernel dimension interference signal to obtain an estimated joint interference signal (**2130**), and removing the estimated joint interference signal from a received signal to obtain a clean signal (**2140**).

[0128] FIG. 22 shows method **2200** of performing interference cancelation. As shown in FIG. 22, method **2200** includes estimating a first interference signal produced by a first amplifier with a first kernel dimension filter and a first delay tap dimension filter (**2210**), estimating a second interference signal produced by a second amplifier with a second kernel dimension filter and a second delay tap dimension filter (**2220**), subtracting a combination of the first interference signal and the second interference signal from a received signal to obtain a clean signal (**2230**), and alternating between a kernel update phase including updating the

first kernel dimension filter and the second kernel dimension filter and a delay tap update phase including updating the first delay tap dimension filter and the second delay tap dimension filter (2240).

[0129] The terms “user equipment”, “UE”, “mobile terminal”, “user terminal”, etc., may apply to any wireless communication device, including cellular phones, tablets, laptops, personal computers, wearables, multimedia playback and other handheld electronic devices, consumer/home/office/commercial appliances, vehicles, and any number of additional electronic devices capable of wireless communications.

[0130] While the above descriptions and connected figures may depict electronic device components as separate elements, skilled persons will appreciate the various possibilities to combine or integrate discrete elements into a single element. Such may include combining two or more circuits for form a single circuit, mounting two or more circuits onto a common chip or chassis to form an integrated element, executing discrete software components on a common processor core, etc. Conversely, skilled persons will recognize the possibility to separate a single element into two or more discrete elements, such as splitting a single circuit into two or more separate circuits, separating a chip or chassis into discrete elements originally provided thereon, separating a software component into two or more sections and executing each on a separate processor core, etc.

[0131] It is appreciated that implementations of methods detailed herein are demonstrative in nature, and are thus understood as capable of being implemented in a corresponding device. Likewise, it is appreciated that implementations of devices detailed herein are understood as capable of being implemented as a corresponding method. It is thus understood that a device corresponding to a method detailed herein may include one or more components configured to perform each aspect of the related method.

[0132] The following examples pertain to further aspects of this disclosure:

[0133] Example 1 is a communication circuit arrangement including a first kernel dimension filter circuit configured to apply a first kernel dimension filter to a first input signal to estimate a first kernel dimension interference signal from a first amplifier, a second kernel dimension filter circuit configured to apply a second kernel dimension filter to a second input signal to estimate a second kernel dimension interference signal from a second amplifier, a joint delay tap dimension filter configured to apply a joint delay tap dimension filter to a combination of the first kernel dimension interference signal and the second kernel dimension interference signal to obtain an estimated joint interference signal, and a cancelation circuit configured to remove the estimated joint interference signal from a received signal to obtain a clean signal.

[0134] In Example 2, the subject matter of Example 1 can optionally further include one or more filter adaptation circuits configured to alternate between a kernel update phase during which the one or more filter adaptation circuits update the first kernel dimension filter and second kernel dimension filter and a delay tap update phase during which the one or more filter adaptation circuits update the joint delay tap dimension filter.

[0135] In Example 3, the subject matter of Example 1 or 2 can optionally be configured as a radio communication device and further including one or more receive chains, a first transmit chain including the first amplifier, a second transmit chain including the second amplifier, and one or more antennas coupled to the one or more receive chains and the first transmit chain.

[0136] In Example 4, the subject matter of Example 1 or 2 can optionally include wherein the first amplifier is configured to amplify the first input signal to form a first amplified signal and the second amplifier is configured to amplify the second input signal to form a second amplified signal and wherein the one or more antennas are configured to transmit the first amplified signal and the second amplified signal.

[0137] In Example 5, the subject matter of Example 3 or 4 can optionally include wherein a first receive chain of the one or more receive chains is configured to receive the received signal via the one or more antennas, wherein the estimated joint interference signal includes approximation of leakage from the first transmit chain and the second transmit chain to the first receive chain.

[0138] In Example 6, the subject matter of Example 5 can optionally include wherein the one or more filter adaptation circuits are further configured to update one or more kernel dimension filters or one or more delay tap dimension filters of at least a second receive chain of the one or more receive chains.

[0139] In Example 7, the subject matter of Example 5 or 6 can optionally include wherein the joint delay tap dimension filter is configured to estimate a delay tap response of one or more filters that are common to a leakage path from the first transmit chain and the second transmit chain to the first receive chain.

[0140] In Example 8, the subject matter of any one of Examples 3 to 7 can optionally include wherein the first input signal and the second input signal are transmit layers of a Multiple Input Multiple Output (MIMO) communication scheme.

[0141] In Example 9, the subject matter of any one of Examples 1 to 8 can optionally include wherein the one or more filter adaptation circuits are configured to update the first kernel dimension filter, the second kernel dimension filter, and the joint delay tap dimension filter in real-time.

[0142] In Example 10, the subject matter of any one of Examples 1 to 9 can optionally include wherein the one or more filter adaptation circuits are configured to update the first kernel dimension filter, the second kernel dimension filter, and the joint delay tap dimension filter by iteratively updating the first kernel dimension filter, the second kernel dimension

filter, and the joint delay tap dimension filter over each of a plurality of samples of the first input signal and the second input signal.

[0143] In Example 11, the subject matter of any one of Examples 1 to 10 can optionally further include a first kernel generation circuit configured to process the first input signal to derive a first plurality of kernel signals from the first input signal, each of the first plurality of kernel signals approximating a non-linear component of a leakage path of the first amplifier, wherein the first kernel dimension filter circuit is configured to apply the first kernel dimension filter to the first input signal by applying the first kernel dimension filter to the first plurality of kernel signals.

[0144] In Example 12, the subject matter of Example 11 can optionally include wherein each of the first plurality of kernel signals is composed of a first plurality of delay taps, wherein each weight of the first kernel dimension filter corresponds to a respective one of the first plurality of kernel signals and each weight of the joint delay tap dimension filter corresponds to a respective one of the first plurality of delay taps.

[0145] In Example 13, the subject matter of Example 12 can optionally include wherein the first kernel dimension filter is configured to approximate the leakage path of the first amplifier over the first plurality of kernel signals, and wherein the joint delay tap dimension filter is configured to approximate the leakage path of the first amplifier over the first plurality of delay taps.

[0146] In Example 14, the subject matter of any one of Examples 1 to 13 can optionally further include a second kernel generation circuit configured to process the second input signal to derive a second plurality of kernel signals from the second input signal, each of the second plurality of kernel signals approximating a non-linear component of a leakage path of the second amplifier, wherein the second kernel dimension filter circuit is configured to apply

the second kernel dimension filter to the second input signal by applying the second kernel dimension filter to the second plurality of kernel signals.

[0147] In Example 15, the subject matter of Example 14 can optionally include wherein each of the second plurality of kernel signals is composed of a second plurality of delay taps, wherein each weight of the second kernel dimension filter corresponds to a respective one of the second plurality of kernel signals and each weight of the joint delay tap dimension filter corresponds to a respective one of the second plurality of delay taps.

[0148] In Example 16, the subject matter of Example 15 can optionally include wherein the second kernel dimension filter is configured to approximate the leakage path of the second amplifier over the second plurality of kernel signals and wherein the joint delay tap dimension filter is configured to approximate the leakage path of the second amplifier over the second plurality of delay taps.

[0149] In Example 17, the subject matter of any one of Examples 1 to 16 can optionally include wherein the first kernel dimension filter, the second kernel dimension filter, and the joint delay tap dimension filter are vectors.

[0150] In Example 18, the subject matter of any one of Examples 1 to 17 can optionally include wherein the one or more filter adaptation circuits are configured to update the first kernel dimension filter, the second kernel dimension filter, and the joint delay tap dimension filter based on the first input signal, the second input signal, and the clean signal.

[0151] In Example 19, the subject matter of any one of Examples 2 to 18 can optionally include wherein the one or more filter adaptation circuits are configured to alternate between the kernel update phase and the delay tap update phase by comparing the combination of the first kernel dimension interference signal and the second kernel dimension interference signal to the clean signal to identify at least one weight of the joint delay tap dimension filter to update during the delay tap update phase.

[0152] In Example 20, the subject matter of Example 1 can optionally include wherein the one or more filter adaptation circuits are configured to compare the combination of the first kernel dimension interference signal and the second kernel dimension interference signal to the clean signal to identify at least one weight of the joint delay tap dimension filter to update by determining a cross-correlation vector between the combination of the first kernel dimension interference signal and the second kernel dimension interference signal and the clean signal, and identifying a first weight of the joint delay tap dimension filter to update based on the cross-correlation vector.

[0153] In Example 21, the subject matter of Example 20 can optionally include wherein the one or more filter adaptation circuits are configured to update exclusively the first weight of the joint delay tap dimension filter.

[0154] In Example 22, the subject matter of Example 20 or 21 can optionally include wherein the one or more filter adaptation circuits are configured to update the first weight of the joint delay tap dimension filter according to a recursive least squares (RLS) optimization scheme.

[0155] In Example 23, the subject matter of Example 22 can optionally include wherein the one or more filter adaptation circuits are configured to update exclusively a single bit of the first weight of the joint delay tap dimension filter according to recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0156] In Example 24, the subject matter of any one of Examples 2 to 23 can optionally include wherein the one or more filter adaptation circuits are configured to alternate between the kernel update phase and the delay tap update phase by first applying the joint delay tap dimension filter to the first input signal to obtain a first filtered input signal, first comparing the first filtered input signal to the clean signal to identify at least one weight of the first kernel dimension filter to update, second applying the joint delay tap dimension filter to the

second input signal to obtain a second filtered input signal, and second comparing the second filtered input signal to the clean signal to identify at least one weight of the second kernel dimension filter to update, wherein the first applying, first comparing, second applying, and second comparing are performed during the kernel update phase.

[0157] In Example 25, the subject matter of Example 1 can optionally include wherein the one or more filter adaptation circuits are configured to compare the first filtered input signal to the clean signal to identify at least one weight of the first kernel dimension filter to update by determining a cross-correlation vector between the first filtered input signal and the clean signal, identifying a first weight of the first kernel dimension filter to update based on the cross-correlation vector.

[0158] In Example 26, the subject matter of Example 25 can optionally include wherein the one or more filter adaptation circuits are configured to update exclusively the first weight of the first kernel dimension filter.

[0159] In Example 27, the subject matter of Example 25 or 26 can optionally include wherein the one or more filter adaptation circuits are configured to update the first weight of the first kernel dimension filter according to a recursive least squares (RLS) optimization scheme.

[0160] In Example 28, the subject matter of Example 27 can optionally include wherein the one or more filter adaptation circuits are configured to update exclusively a single bit of the first weight of the first kernel dimension filter according to recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0161] In Example 29, the subject matter of any one of Examples 1 to 28 can optionally include wherein the one or more filter adaptation circuits include a first kernel dimension filter adaptation circuit configured to perform the updates to the first kernel dimension filter

and further include a different second kernel dimension adaptation circuit configured to perform the updates to the second kernel dimension filter.

[0162] In Example 30, the subject matter of any one of Examples 1 to 28 can optionally include wherein the one or more filter adaptation circuits include a shared kernel dimension update filter compared to alternate between updating the first kernel dimension filter and the second kernel dimension filter.

[0163] Example 31 is a communication circuit arrangement including a signal path circuit configured to estimate a first interference signal produced by a first amplifier with a first kernel dimension filter and a first delay tap dimension filter and further configured to estimate a second interference signal produced by a second amplifier with a second kernel dimension filter and a second delay tap dimension filter, a cancelation circuit configured to subtract a combination of the first interference signal and the second interference signal from a received signal to obtain a clean signal, and one or more filter adaptation circuits configured to alternate between a kernel update phase during which the one or more filter adaptation circuits update the first kernel dimension filter and the second kernel dimension filter and a delay update phase during which the one or more filter adaptation circuits update the first delay tap dimension filter and the second delay tap dimension filter.

[0164] In Example 32, the subject matter of Example 31 can optionally be configured as a radio communication device and further including one or more receive chains, a first transmit chain including the first amplifier, a second transmit chain including the second amplifier, and one or more antennas coupled to the one or more receive chains and the first transmit chain.

[0165] In Example 33, the subject matter of Example 32 can optionally include wherein a first receive chain of the one or more receive chains is configured to receive the received signal via the one or more antennas, wherein the combination of the first interference signal

and the second interference signal includes approximation of leakage from the first transmit chain and the second transmit chain to the first receive chain.

[0166] In Example 34, the subject matter of Example 33 can optionally include wherein the one or more filter adaptation circuits are further configured to update one or more kernel dimension filters or one or more delay tap dimension filters of at least a second receive chain of the one or more receive chains.

[0167] In Example 35, the subject matter of any one of Examples 32 to 34 can optionally include wherein the first transmit chain and the second transmit chain are layers of a Multiple Input Multiple Output (MIMO) communication scheme.

[0168] In Example 36, the subject matter of any one of Examples 31 to 35 can optionally include wherein the first kernel dimension filter and the first delay tap dimension filter are configured to approximate a leakage path from the first amplifier and the second kernel dimension filter and the second delay tap dimension filter are configured to approximate a leakage path from the second amplifier.

[0169] In Example 37, the subject matter of any one of Examples 31 to 36 can optionally include wherein the one or more filter adaptation circuits are configured to update the first kernel dimension filter, the first delay tap dimension filter, the second kernel dimension filter, and the second delay tap dimension filter in real-time.

[0170] In Example 38, the subject matter of any one of Examples 31 to 37 can optionally include wherein the first kernel dimension filter, the second kernel dimension filter, the first delay tap dimension filter, and the second delay tap dimension filter are vectors.

[0171] In Example 39, the subject matter of any one of Examples 31 to 38 can optionally include wherein the signal path circuit is configured to estimate the first interference signal by applying the first kernel dimension filter and the first delay tap dimension filter to a first input signal for the first amplifier to obtain the first interference signal and configured to estimate

the second interference signal by applying the second kernel dimension filter and the second delay tap dimension filter to a second input signal for the second amplifier to obtain the second interference signal.

[0172] In Example 40, the subject matter of Example 39 can optionally include wherein the one or more filter adaptation circuits are configured to update the first kernel dimension filter, the first delay tap dimension filter, the second kernel dimension filter, and the second delay tap dimension filter by iteratively updating the first kernel dimension filter, the first delay tap dimension filter, the second kernel dimension filter, and the second delay tap dimension filter over each of a plurality of samples of the first input signal and the second input signal.

[0173] In Example 41, the subject matter of Example 39 or 40 can optionally further include a first kernel generation circuit configured to process the first input signal to derive a first plurality of kernel signals from the first input signal, each of the first plurality of kernel signals approximating a non-linear component of a leakage path of the first amplifier, wherein the signal path circuit is configured to apply the first kernel dimension filter to the first input signal by applying the first kernel dimension filter to the first plurality of kernel signals.

[0174] In Example 42, the subject matter of Example 41 can optionally include wherein each of the first plurality of kernel signals is composed of a plurality of delay taps, wherein each weight of the first kernel dimension filter corresponds to a respective one of the first plurality of kernel signals and each weight of the first delay tap dimension filter corresponds to a respective one of the plurality of delay taps.

[0175] In Example 43, the subject matter of Example 42 can optionally include wherein the first kernel dimension filter is configured to approximate the leakage path of the first amplifier over the first plurality of kernel signals and the first delay tap dimension filter is configured to approximate the leakage path of the first amplifier over the first plurality of delay taps.

[0176] In Example 44, the subject matter of any one of Examples 39 to 43 can optionally further include a second kernel generation circuit configured to process the second input signal to derive a second plurality of kernel signals from the second input signal, each of the second plurality of kernel signals approximating a non-linear component of a leakage path of the second amplifier, wherein the signal path circuit is configured to apply the second kernel dimension filter to the second input signal by applying the second kernel dimension filter to the second plurality of kernel signals.

[0177] In Example 45, the subject matter of Example 44 can optionally include wherein each of the second plurality of kernel signals is composed of a second plurality of delay taps, wherein each weight of the second kernel dimension filter corresponds to a respective one of the second plurality of kernel signals and each weight of the second delay tap dimension filter corresponds to a respective one of the second plurality of delay taps.

[0178] In Example 46, the subject matter of Example 45 can optionally include wherein the second kernel dimension filter is configured to approximate the leakage path of the second amplifier over the second plurality of kernel signals and the second delay tap dimension filter is configured to approximate the leakage path of the second amplifier over the second plurality of delay taps.

[0179] In Example 47, the subject matter of any one of Examples 39 to 46 can optionally include wherein the one or more filter adaptation circuits are configured to update the first kernel dimension filter and the first delay tap dimension filter based on the first input signal and the clean signal and configured to update the second kernel dimension filter and the second delay tap dimension filter based on the second input signal and the clean signal.

[0180] In Example 48, the subject matter of any one of Examples 39 to 47 can optionally include wherein the one or more filter adaptation circuits are configured to alternate between the kernel update phase and the delay tap update phase by first applying the first kernel

dimension filter to the first input signal to obtain a first filtered input signal, first comparing the first filtered input signal to the clean signal to identify at least one weight of the first delay tap dimension filter to update, second applying the second kernel dimension filter to the second input signal to obtain a second filtered input signal, and second comparing the second filtered input signal to the clean signal to identify at least one weight of the second delay tap dimension filter to update, wherein the first applying, first comparing, second applying, and second comparing are performed during the delay tap update phase.

[0181] In Example 49, the subject matter of Example 48 can optionally include wherein the one or more filter adaptation circuits are configured to compare the first filtered input signal to the clean signal to identify the at least one weight of the first delay tap dimension filter to update by determining a cross-correlation vector between the first filtered input signal and the clean signal, identifying a first weight of the first delay tap dimension filter to update based on the cross-correlation vector.

[0182] In Example 50, the subject matter of Example 49 can optionally include wherein the one or more filter adaptation circuits are configured to update exclusively the first weight of the first delay tap dimension filter.

[0183] In Example 51, the subject matter of Example 49 or 50 can optionally include wherein the one or more filter adaptation circuits are configured to update the first weight of the first delay tap dimension filter according to a recursive least squares (RLS) optimization scheme.

[0184] In Example 52, the subject matter of Example 51 can optionally include wherein the one or more filter adaptation circuits are configured to update the first weight of the first delay tap dimension filter according to a recursive least square (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0185] In Example 53, the subject matter of any one of Examples 48 to 52 can optionally include wherein the one or more filter adaptation circuits include a first delay tap dimension filter adaptation circuit configured to perform the updates to the first delay tap dimension filter and further include a different second delay tap dimension adaptation circuit configured to perform the updates to the second delay tap dimension filter.

[0186] In Example 54, the subject matter of any one of Examples 48 to 52 can optionally include wherein the one or more filter adaptation circuits include a shared delay tap dimension update filter compared to alternate between updating the first delay tap dimension filter and the second delay tap dimension filter

[0187] In Example 55, the subject matter of any one of Examples 39 to 52 can optionally include wherein the one or more filter adaptation circuits are configured to alternate between the kernel update phase and the delay tap update phase by first applying the first delay tap dimension filter to the first input signal to obtain a first filtered input signal, first comparing the first filtered input signal to the clean signal to identify at least one weight of the first kernel dimension filter to update, second applying the second delay dimension filter to the second input signal to obtain a second filtered input signal, and second comparing the second filtered input signal to the clean signal to identify at least one weight of the second kernel dimension filter to update, wherein the first applying, first comparing, second applying, and second comparing are performed during the kernel update phase.

[0188] In Example 56, the subject matter of Example 55 can optionally include wherein the one or more filter adaptation circuits are configured to compare the first filtered input signal to the clean signal to identify at least one weight of the first kernel dimension filter to update by determining a cross-correlation vector between the first filtered input signal and the clean signal, identifying a first weight of the first kernel dimension filter to update based on the cross-correlation vector.

[0189] In Example 57, the subject matter of Example 56 can optionally include wherein the one or more filter adaptation circuits are configured to update exclusively the first weight of the first kernel dimension filter.

[0190] In Example 58, the subject matter of Example 56 or 57 can optionally include wherein the one or more filter adaptation circuits are configured to update the first weight of the first kernel dimension filter according to a recursive least squares (RLS) optimization scheme.

[0191] In Example 59, the subject matter of Example 58 can optionally include wherein the one or more filter adaptation circuits are configured to update exclusively a single bit of the first weight of the first kernel dimension filter according to recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0192] In Example 60, the subject matter of any one of Examples 55 to 59 can optionally include wherein the one or more filter adaptation circuits include a first kernel dimension filter adaptation circuit configured to perform the updates to the first kernel dimension filter and further include a different second kernel dimension adaptation circuit configured to perform the updates to the second kernel dimension filter.

[0193] In Example 61, the subject matter of any one of Examples 55 to 59 can optionally include wherein the one or more filter adaptation circuits include a shared kernel dimension update filter compared to alternate between updating the first kernel dimension filter and the second kernel dimension filter.

[0194] Example 62 is a method of performing interference cancelation including applying a first kernel dimension filter to a first input signal to estimate a first kernel dimension interference signal derived from a first amplifier, applying a second kernel dimension filter to a second input signal to estimate a second kernel dimension interference signal derived from a second amplifier, applying a joint delay tap dimension filter to a combination of the first

kernel dimension interference signal and the second kernel dimension interference signal to obtain an estimated joint interference signal, and removing the estimated joint interference signal from a received signal to obtain a clean signal.

[0195] In Example 63, the subject matter of Example 62 can optionally further include alternating between a kernel update phase including updating the first kernel dimension filter and the second kernel dimension filter and a delay tap update phase including updating the joint delay tap dimension filter.

[0196] In Example 64, the subject matter of Example 62 or 63 can optionally further include amplifying the first input signal with the first amplifier to form a first amplified signal, amplifying the second input signal with the second amplifier to form a second amplified signal, and wirelessly transmitting the first amplified signal and the second amplified signal.

[0197] In Example 65, the subject matter of Example 64 can optionally further include receiving the received signal at a receive chain, where the estimated joint interference signal includes approximation of leakage from the first amplifier and the second amplifier to the receive chain.

[0198] In Example 66, the subject matter of Example 65 can optionally include wherein the joint delay tap estimation filter estimates a delay tap response of one or more filters that are common to a leakage path from both the first amplifier and the second amplifier to the receive chain.

[0199] In Example 67, the subject matter of any one of Examples 64 to 66 can optionally include wherein the first input signal and the second input signal are transmit layers of a Multiple Input Multiple Output (MIMO) communication scheme.

[0200] In Example 68, the subject matter of any one of Examples 63 to 67 can optionally include wherein the kernel update phase and the delay tap update phase include updating the

first kernel dimension filter, the second kernel dimension filter, and the joint delay tap dimension filter in real-time.

[0201] In Example 69, the subject matter of any one of Examples 63 to 68 can optionally include wherein the kernel update phase and the delay tap update phase include iteratively updating the first kernel dimension filter, the second kernel dimension filter, and the joint delay tap dimension filter over each of a plurality of samples of the first input signal and the second input signal.

[0202] In Example 70, the subject matter of any one of Examples 62 to 69 can optionally further include processing the first input signal to derive a first plurality of kernel signals from the first input signal, each of the first plurality of kernel signals approximating a non-linear component of a leakage path of the first amplifier, and wherein applying the first kernel dimension filter to the first input signal to estimate the first kernel dimension interference signal includes applying the first kernel dimension filter to the first plurality of kernel signals.

[0203] In Example 71, the subject matter of Example 70 can optionally include wherein each of the first plurality of kernel signals is composed of a first plurality of delay taps, wherein each weight of the first kernel dimension filter corresponds to a respective one of the first plurality of kernel signals and each weight of the joint delay tap dimension filter corresponds to a respective one of the first plurality of delay taps.

[0204] In Example 72, the subject matter of Example 71 can optionally include wherein the first kernel dimension filter is configured to approximate the leakage path of the first amplifier over the first plurality of kernel signals and the joint delay tap dimension filter is configured to approximate the leakage path of the first amplifier over the first plurality of delay taps.

[0205] In Example 73, the subject matter of any one of Examples 62 to 72 can optionally further include processing the second input signal to derive a second plurality of kernel

signals from the second input signal, each of the second plurality of kernel signals approximating a non-linear component of a leakage path of the second amplifier, and wherein applying the second kernel dimension filter to the second input signal to estimate the second kernel dimension interference signal includes applying the second kernel dimension filter to the second plurality of kernel signals.

[0206] In Example 74, the subject matter of Example 73 can optionally include wherein each of the second plurality of kernel signals is composed of a second plurality of delay taps, wherein each weight of the second kernel dimension filter corresponds to a respective one of the second plurality of kernel signals and each weight of the joint delay tap dimension filter corresponds to a respective one of the second plurality of delay taps.

[0207] In Example 75, the subject matter of Example 74 can optionally include wherein the second kernel dimension filter is configured to approximate the leakage path of the second amplifier over the second plurality of kernel signals and the joint delay tap dimension filter is configured to approximate the leakage path of the second amplifier over the second plurality of delay taps.

[0208] The method of any one of Examples 62 to 75, wherein the first kernel dimension filter, the second kernel dimension filter, and the joint delay tap dimension filter are vectors.

[0209] In Example 77, the subject matter of any one of Examples 63 to 76 can optionally include wherein the kernel update phase and the delay tap update phase include updating the first kernel dimension filter, the second kernel dimension filter, and the joint delay tap dimension filter based on the first input signal, the second input signal, and the clean signal.

[0210] In Example 78, the subject matter of any one of Examples 63 to 77 can optionally include wherein the delay tap update phase includes updating the joint delay tap dimension filter by comparing the combination of the first kernel dimension interference signal and the

second kernel dimension interference signal to the clean signal to identify at least one weight of the joint delay tap dimension filter to update.

[0211] In Example 79, the subject matter of Example 78 can optionally include wherein comparing the combination of the first kernel dimension interference signal and the second kernel dimension interference signal to the clean signal to identify the at least one weight of the joint delay tap dimension filter to update includes determining a cross-correlation vector between the combination of the first kernel dimension interference signal and the second kernel dimension interference signal and the clean signal, and identifying a first weight of the joint delay tap dimension filter to update based on the cross-correlation vector.

[0212] In Example 80, the subject matter of Example 79 can optionally include wherein identifying the first weight of the joint delay tap dimension filter to update based on the cross-correlation vector includes updating exclusively the first weight of the joint delay tap dimension filter.

[0213] In Example 81, the subject matter of Example 79 can optionally include wherein identifying the first weight of the joint delay tap dimension filter to update based on the cross-correlation vector includes updating the first weight of the joint delay tap dimension filter according to a recursive least squares (RLS) optimization scheme.

[0214] In Example 82, the subject matter of Example 81 can optionally include wherein identifying the first weight of the joint delay tap dimension filter to update based on the cross-correlation vector includes updating exclusively a single bit of the first weight of the joint delay tap dimension filter according to recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0215] In Example 83, the subject matter of any one of Examples 63 to 82 can optionally include wherein the kernel update phase includes applying the joint delay tap dimension filter to the first input signal to obtain a first filtered input signal, comparing the first filtered input

signal to the clean signal to identify at least one weight of the first kernel dimension filter to update, applying the joint delay tap dimension filter to the second input signal to obtain a second filtered input signal, and comparing the second filtered input signal to the clean signal to identify at least one weight of the second kernel dimension filter to update.

[0216] In Example 84, the subject matter of Example 83 can optionally include wherein comparing the first filtered input signal to the clean signal to identify the at least one weight of the first kernel dimension filter to update includes determining a cross-correlation vector between the first filtered input signal and the clean signal, identifying a first weight of the first kernel dimension filter to update based on the cross-correlation vector.

[0217] In Example 85, the subject matter of Example 84 can optionally include wherein identifying the first weight of the first kernel dimension filter to update based on the cross-correlation vector includes updating exclusively the first weight of the first kernel dimension filter.

[0218] In Example 86, the subject matter of Example 84 can optionally include wherein identifying the first weight of the first kernel dimension filter to update based on the cross-correlation vector includes updating the first weight of the first kernel dimension filter according to a recursive least squares (RLS) optimization scheme.

[0219] In Example 87, the subject matter of Example 84 can optionally include wherein identifying the first weight of the first kernel dimension filter to update based on the cross-correlation vector includes updating exclusively a single bit of the first weight of the first kernel dimension filter according to a recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0220] Example 88 is a method of performing interference cancelation including estimating a first interference signal produced by a first amplifier with a first kernel dimension filter and a first delay tap dimension filter, estimating a second interference signal produced by a second

amplifier with a second kernel dimension filter and a second delay tap dimension filter, subtracting a combination of the first interference signal and the second interference signal from a received signal to obtain a clean signal, and alternating between a kernel update phase including updating the first kernel dimension filter and the second kernel dimension filter and a delay tap update phase including updating the first delay tap dimension filter and the second delay tap dimension filter.

[0221] In Example 89, the subject matter of Example 88 can optionally further include amplifying a first input signal with the first amplifier, amplifying a second input signal with the second amplifier, and wirelessly transmitting the amplified first input signal and the amplified second input signal.

[0222] In Example 90, the subject matter of Example 89 can optionally further include receiving the received signal at a receive chain, wherein the combination of the first interference signal and the second interference signal includes approximation of leakage from the first power amplifier and the second power amplifier to the receive chain.

[0223] In Example 91, the subject matter of Example 89 can optionally include wherein the first input signal and the second input signal are transmit layers of a Multiple Input Multiple Output (MIMO) communication scheme.

[0224] In Example 92, the subject matter of any one of Examples 88 to 91 can optionally include wherein the first kernel dimension filter and the first delay tap dimension filter are configured to approximate a leakage path from the first amplifier and the second kernel dimension filter and the second delay tap dimension filter are configured to approximate a leakage path from the second amplifier.

[0225] In Example 93, the subject matter of any one of Examples 88 to 92 can optionally include wherein the kernel update phase and the delay tap update phase include updating the

first kernel dimension filter, the second kernel dimension filter, the first delay tap dimension filter, and the second delay tap dimension filter in real-time.

[0226] In Example 94, the subject matter of any one of Examples 88 to 92 can optionally include wherein the first kernel dimension filter, the second kernel dimension filter, the first delay tap dimension filter, and the second delay tap dimension filter are vectors.

[0227] In Example 95, the subject matter of any one of Examples 88 to 94 can optionally include wherein estimating the first interference signal produced by the first amplifier with the first kernel dimension filter and the first delay tap dimension filter includes applying the first kernel dimension filter and the first delay tap dimension filter to a first input signal for the first amplifier to obtain the first interference signal, and wherein estimating the second interference signal produced by the second amplifier with the second kernel dimension filter and the second kernel dimension filter includes applying the first kernel dimension filter and the second kernel dimension filter to a second input signal for the second amplifier to obtain the second interference signal.

[0228] In Example 96, the subject matter of Example 95 can optionally include wherein the kernel update phase and the delay tap update phase include updating the first kernel dimension filter, the first delay tap dimension filter, the second kernel dimension filter, and the second delay tap dimension filter by iteratively updating the first kernel dimension filter, the first delay tap dimension filter, the second kernel dimension filter, and the second delay tap dimension filter over each of a plurality of samples of the first input signal and the second input signal.

[0229] In Example 97, the subject matter of any one of Examples , further including can optionally include the first input signal to derive a first plurality of kernel signals from the first input signal, each of the first plurality of kernel signals approximating a non-linear component of a leakage path of the first amplifier, and wherein applying the first kernel

dimension filter and the first delay tap dimension filter to the first input signal for the first amplifier to obtain the first interference signal includes applying the first kernel dimension filter to the first plurality of kernel signals.

[0230] In Example 98, the subject matter of Example 97 can optionally include wherein each of the first plurality of kernel signals is composed of a first plurality of delay taps, wherein each weight of the first kernel dimension filter corresponds to a respective one of the first plurality of kernel signals and each weight of the first delay tap dimension filter corresponds to a respective one of the first plurality of delay taps.

[0231] In Example 99, the subject matter of Example 98 can optionally include wherein the first kernel dimension filter is configured to approximate the leakage path of the first amplifier over the first plurality of kernel signals and the first delay tap dimension filter is configured to approximate the leakage path of the first amplifier over the first plurality of delay taps.

[0232] In Example 100, the subject matter of any one of Examples 95 to 99 can optionally further include processing the second input signal to derive a second plurality of kernel signals from the second input signal, each of the second plurality of kernel signals approximating a non-linear component of a leakage path of the second amplifier, and wherein applying the second kernel dimension filter and the second delay tap dimension filter to the second input signal for the second amplifier to obtain the second interference signal includes applying the second kernel dimension filter to the second plurality of kernel signals.

[0233] In Example 101, the subject matter of Example 100 can optionally include wherein each of the second plurality of kernel signals is composed of a second plurality of delay taps, wherein each weight of the second kernel dimension filter corresponds to a respective one of the second plurality of kernel signals and each weight of the second delay tap dimension filter corresponds to a respective one of the second plurality of delay taps.

[0234] In Example 102, the subject matter of Example 98 can optionally include wherein the second kernel dimension filter is configured to approximate the leakage path of the second amplifier over the second plurality of kernel signals and the second delay tap dimension filter is configured to approximate the leakage path of the first amplifier over the second plurality of delay taps.

[0235] In Example 103, the subject matter of any one of Examples 95 to 102 can optionally include wherein the kernel update phase and the delay tap update phase include updating the first kernel dimension filter and the first delay tap dimension filter based on the first input signal and the clean signal and updating the second kernel dimension filter and the second kernel dimension filter based on the second input signal and the clean signal.

[0236] In Example 104, the subject matter of any one of Examples 95 to 103 can optionally include wherein the delay tap update phase includes applying the first kernel dimension filter to the first input signal to obtain a first filtered input signal, comparing the first filtered input signal to the clean signal to identify at least one weight of the first delay tap dimension filter to update, applying the second kernel dimension filter to the second input signal to obtain a second filtered input signal, and comparing the second filtered input signal to the clean signal to identify at least one weight of the second delay tap dimension filter to update.

[0237] In Example 105, the subject matter of Example 104 can optionally include wherein comparing the first filtered input signal to the clean signal to identify the at least one weight of the first delay tap dimension filter to update includes determining a cross-correlation vector between the first filtered input signal and the clean signal, identifying a first weight of the first delay tap dimension filter to update based on the cross-correlation vector.

[0238] In Example 106, the subject matter of Example 105 can optionally include wherein identifying the first weight of the first delay tap dimension filter to update based on the cross-

correlation vector includes updating exclusively the first weight of the first delay tap dimension filter.

[0239] In Example 107, the subject matter of Example 105 or 106 can optionally include wherein identifying the first weight of the first delay tap dimension filter to update based on the cross-correlation vector includes updating the first weight of the first delay tap dimension filter according to a recursive least squares (RLS) optimization scheme.

[0240] In Example 108, the subject matter of Example 105 or 106 can optionally include wherein identifying the first weight of the first delay tap dimension filter to update based on the cross-correlation vector includes updating the first weight of the first delay tap dimension filter according to a recursive least square (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0241] In Example 109, the subject matter of any one of Examples 95 to 108 can optionally include wherein the kernel update phase includes applying the first delay tap dimension filter to the first input signal to obtain a first filtered input signal, comparing the first filtered input signal to the clean signal to identify at least one weight of the first kernel dimension filter to update, applying the second delay dimension filter to the second input signal to obtain a second filtered input signal, and comparing the second filtered input signal to the clean signal to identify at least one weight of the second kernel dimension filter to update.

[0242] In Example 110, the subject matter of Example 109 can optionally include wherein comparing the first filtered input signal to the clean signal to identify the at least one weight of the first kernel dimension filter to update includes determining a cross-correlation vector between the first filtered input signal and the clean signal, identifying a first weight of the first kernel dimension filter to update based on the cross-correlation vector.

[0243] In Example 111, the subject matter of Example 110 can optionally include wherein identifying the first weight of the first kernel dimension filter to update based on the cross-

correlation vector includes updating exclusively the first weight of the first kernel dimension filter.

[0244] In Example 112, the subject matter of Example 110 or 111 can optionally include wherein identifying the first weight of the first kernel dimension filter to update based on the cross-correlation vector includes updating the first weight of the first kernel dimension filter according to a recursive least squares (RLS) optimization scheme.

[0245] In Example 113, the subject matter of Example 110 or 111 can optionally include wherein identifying the first weight of the first kernel dimension filter to update based on the cross-correlation vector includes updating the first weight of the first kernel dimension filter according to a recursive least square (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0246] Example 114 is a method of performing interference cancelation including separately applying a kernel dimension filter and a delay tap dimension filter to an input signal for an amplifier to obtain an estimated interference signal, subtracting the estimated interference signal from a received signal to obtain a clean signal, and alternating between updating the kernel dimension filter and updating the delay tap dimension filter using the clean signal.

[0247] In Example 115, the subject matter of Example 114 can optionally include wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter using the clean signal includes updating the kernel dimension filter in real-time over a first plurality of digital time samples of the input signal and updating the delay tap dimension filter in real-time over a second plurality of digital time samples of the input signal.

[0248] In Example 116, the subject matter of Example 114 can optionally include wherein the input signal includes a plurality of digital time samples, and wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter using the

clean signal includes updating the kernel dimension filter or the delay tap dimension filter over each of the plurality of digital time samples.

[0249] In Example 117, the subject matter of Example 114 can optionally include wherein separately applying the kernel dimension filter and the delay tap dimension filter to the input signal for the amplifier to obtain the estimated interference signal includes processing the input signal to derive a plurality of kernel signals from the input signal, wherein each of the plurality of kernel signals approximates a non-linear component of a response of the amplifier, and separately applying the kernel dimension filter and the delay tap dimension filter to the plurality of kernel signals to obtain the estimated interference signal.

[0250] In Example 118, the subject matter of Example 117 can optionally include wherein each of the plurality of kernel signals is composed of a plurality of delay taps, wherein each weight of the kernel dimension filter corresponds to a respective one of the plurality of kernel signals and each weight of the delay tap dimension filter corresponds to a respective one of the plurality of delay taps.

[0251] In Example 119, the subject matter of Example 118 can optionally include wherein the number of plurality of kernel signals is configurable or the number of plurality of delay taps is configurable.

[0252] In Example 120, the subject matter of Example 118 can optionally include wherein the kernel dimension filter approximates the response of the amplifier over the plurality of kernel signals and the delay tap dimension filter approximates the response of the amplifier over the plurality of delay taps.

[0253] In Example 121, the subject matter of any one of Examples 114 to 120 can optionally include wherein the kernel dimension filter and the delay tap dimension filter are vectors.

[0254] In Example 122, the subject matter of any one of Examples 114 to 120 can optionally further include amplifying the input signal with the amplifier and transmitting the amplified input signal with a radio antenna.

[0255] In Example 123, the subject matter of Example 120 can optionally further include receiving the received signal with the radio antenna, wherein the estimated interference signal approximates leakage from the amplified input signal that is contained in the received signal.

[0256] In Example 124, the subject matter of any one of Examples 114 to 123 can optionally include wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter includes updating the kernel dimension filter and the delay tap dimension filter based on the plurality of kernel signals and the clean signal.

[0257] In Example 125, the subject matter of any one of Examples 117 to 123 can optionally include wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter includes selecting between the kernel dimension filter and the delay tap dimension filter to identify a current filter to update and a fixed filter to hold constant, applying the fixed filter to the plurality of kernel signals to obtain a decoupled input signal, and comparing the decoupled input signal to the clean signal to identify at least one weight of the current filter to update.

[0258] In Example 126, the subject matter of Example 125 can optionally include wherein comparing the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update includes determining a cross-correlation vector between the decoupled input signal and the clean signal, identifying a first weight of the current filter to update based on the cross-correlation vector, and exclusively updating the first weight of the current filter.

[0259] In Example 127, the subject matter of Example 126 can optionally include wherein identifying the first weight of the current filter to update based on the cross-correlation vector

includes identifying a maximum-valued element of the cross-correlation vector, and identifying the weight of the current filter with a corresponding element index to the maximum-valued element of the cross-correlation vector as the first weight.

[0260] In Example 128, the subject matter of Example 127 can optionally include wherein identifying the maximum-valued element of the cross-correlation vector includes identifying the element of the cross-correlation vector with the largest real component or the largest imaginary component as the maximum-valued element.

[0261] In Example 129, the subject matter of Example 125 can optionally include wherein comparing the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update includes determining a cross-correlation vector between the decoupled input signal and the clean signal, and updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector.

[0262] In Example 130, the subject matter of Example 129 can optionally include wherein updating the at least one weight of the current filter to reduce the magnitude of the cross-correlation vector includes updating the at least one weight of the current filter to reduce the magnitude of the cross-correlation vector according to a coordinate descent optimization scheme.

[0263] In Example 131, the subject matter of Example 129 can optionally include wherein updating the at least one weight of the current filter to reduce the magnitude of the cross-correlation vector includes updating the at least one weight of the current filter to reduce the magnitude of the cross-correlation vector according to a recursive least squares optimization scheme.

[0264] In Example 132, the subject matter of Example 129 can optionally include wherein updating the at least one weight of the current filter to reduce the magnitude of the cross-correlation vector includes updating the at least one weight of the current filter to reduce the

magnitude of the cross-correlation vector according to a recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0265] In Example 133, the subject matter of Example 129 can optionally include wherein updating the at least one weight of the current filter to reduce the magnitude of the cross-correlation vector includes exclusively inverting a single bit of the first weight to reduce the magnitude of the cross-correlation vector.

[0266] In Example 134, the subject matter of Example 133 can optionally include wherein updating the at least one weight of the current filter to reduce the magnitude of the cross-correlation vector further includes evaluating one or more candidate bit inversions of the first weight to identify which of the one or more candidate bit inversions is closest to a predefined numerical difference, and identifying the single bit according to which of the one or more candidate bit inversions is closest to the predefined numerical difference.

[0267] In Example 135, the subject matter of Example 133 can optionally include wherein exclusively inverting the single bit of the first weight to reduce the magnitude of the cross-correlation vector includes inverting the single bit of the first weight according to a dichotomous coordinate descent optimization scheme.

[0268] In Example 136, the subject matter of Example 125 can optionally include wherein selecting between the kernel dimension filter and the delay tap dimension filter to identify the current filter to update and the fixed filter to hold constant includes selecting the kernel dimension filter as the current filter and the delay tap dimension filter as the fixed filter, and wherein applying the fixed filter to the plurality of kernel signals to obtain the decoupled input signal includes applying previous delay taps of the delay tap dimension filter to previous delay taps of the plurality of kernel signals during a preprocessing stage that occurs before receiving a most recent delay tap of the plurality of kernel signals, and after receiving the

most recent delay tap of the plurality of kernel signals, applying a most recent delay tap of the delay tap dimension filter to the most recent delay tap of the plurality of kernel signals.

[0269] In Example 137, the subject matter of Example 136 can optionally further include performing the preprocessing stage during hidden clock cycles that occur before the most recent delay tap of the plurality of kernel signals is obtained.

[0270] In Example 138, the subject matter of Example 125 can optionally include wherein selecting between the kernel dimension filter and the delay tap dimension filter to identify the current filter to update and the fixed filter to hold constant includes selecting the delay tap dimension filter as the current filter and the kernel dimension filter as the fixed filter, and wherein applying the fixed filter to the plurality of kernel signals to obtain the decoupled input signal includes exclusively applying the kernel dimension filter to samples of the plurality of kernel signals corresponding to a single delay tap to obtain a first element of the decoupled input signal.

[0271] In Example 139, the subject matter of Example 138 can optionally include wherein the remaining elements of the decoupled input signal are time-delayed weighted versions of the first element of the decoupled input signal.

[0272] In Example 140, the subject matter of any one of Examples 117 to 135 can optionally include wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter includes selecting between the kernel dimension filter and the delay tap dimension filter to identify a current filter to update and a fixed filter to hold constant, applying the fixed filter to the plurality of kernel signals to obtain a decoupled input signal, determining a cross-correlation vector between the decoupled input signal and the clean signal and determining a correlation matrix of the decoupled input signal, and updating at least one weight of the current filter based on the cross-correlation vector and the correlation matrix.

[0273] In Example 141, the subject matter of Example 140 can optionally include wherein determining the cross-correlation vector between the decoupled input signal and the clean signal includes determining some elements of the cross-correlation vector during a first clock cycle with first calculation circuitry and determining other elements of the cross-correlation vector during a second clock cycle with the same first calculation circuitry, or determining some elements of the correlation matrix during a third clock cycle with second calculation circuitry and determining other elements of the correlation matrix during a second clock cycle with the same second calculation circuitry.

[0274] In Example 142, the subject matter of Example 140 can optionally include wherein selecting between the kernel dimension filter and the delay tap dimension filter to identify the current filter to update and the fixed filter to hold constant includes selecting the kernel dimension filter as the current filter and the delay tap dimension filter as the fixed filter, and wherein determining the cross-correlation vector between the decoupled input signal and the clean signal and determining the correlation matrix of the decoupled input signal includes exclusively determining the upper-triangle elements of the correlation matrix.

[0275] In Example 43, the subject matter of Example 142 can optionally include wherein the correlation matrix is a Hermitian matrix.

[0276] In Example 144, the subject matter of Example 140 can optionally include wherein selecting between the kernel dimension filter and the delay tap dimension filter to identify the current filter to update and the fixed filter to hold constant includes selecting the delay tap dimension filter as the current filter and the kernel dimension filter as the fixed filter, and wherein determining the cross-correlation vector between the decoupled input signal and the clean signal and determining the correlation matrix of the decoupled input signal includes exclusively determining a single row of the correlation matrix.

[0277] In Example 145, the subject matter of Example 144 can optionally include wherein a plurality of rows including the single row of the correlation matrix are statistically similar.

[0278] In Example 146, the subject matter of any one of Examples 114 to 139 can optionally include wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter includes using shared circuitry to update the kernel dimension filter and re-using the same shared circuitry to update the delay tap dimension filter.

[0279] In Example 147, the subject matter of any one of Examples 114 to 139 can optionally include wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter includes using shared circuitry to update the kernel dimension filter and re-using the same shared circuitry to update the delay tap dimension filter according to a pipeline clock schedule.

[0280] In Example 148, the subject matter of Example 147 can optionally include wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter includes updating the kernel dimension filter at one or more first times and updating the delay tap dimension at one or more different second times.

[0281] Example 149 is a communication circuitry arrangement configured to perform the method of any one of Examples 114 to 148.

[0282] In Example 150, the subject matter of Example 149 can optionally be configured as a radio communication device.

[0283] Example 151 is a non-transitory computer readable medium storing instructions that when executed by a controller of a radio communication device direct the radio communication device to perform the method of any one of Examples 114 to 148.

[0284] Example 152 is a method of performing interference cancelation including obtaining one or more subsignals derived from an input signal for an amplifier, each of the one or more subsignals representing a non-linear component of an amplifier response and composed of a plurality of delay taps, separately applying a first filter and a second filter to the one or more subsignals to obtain an estimated interference signal, wherein the first filter approximates the amplifier response over the one or more subsignals and the second filter approximates the amplifier response over the plurality of delay taps, subtracting the estimated interference signal from a received signal to obtain a clean signal, and alternating between updating the first filter and updating the second filter using the clean signal.

[0285] In Example 153, the subject matter of Example 152 can optionally include wherein alternating between updating the first filter and updating the second filter using the clean signal includes updating the first filter in real-time over a first plurality of digital time samples of the input signal and updating second first filter in real-time over a second plurality of digital time samples of the input signal.

[0286] In Example 154, the subject matter of Example 152 can optionally include wherein the input signal includes a plurality of digital time samples, and wherein alternating between updating the first filter and updating the second filter using the clean signal includes updating the first filter or the second filter over each of the plurality of digital time samples.

[0287] In Example 155, the subject matter of Example 152 can optionally include wherein each of the one or more subsignals correspond to a kernel of the amplifier.

[0288] In Example 156, the subject matter of Example 152 or 155 can optionally include wherein the first filter and the second filter are vectors.

[0289] In Example 157, the subject matter of any one of Examples 152 to 156 can optionally include wherein each weight of the first filter corresponds to a respective one of the

one or more subsignals and each weight of the second filter corresponds to a respective delay tap of the plurality of delay taps.

[0290] In Example 158, the subject matter of any one of Examples 152 to 157 can optionally include wherein the number of one or more subsignals is configurable the number of the plurality of delay taps is configurable.

[0291] In Example 159, the subject matter of any one of Examples 152 to 158 can optionally further include amplifying the input signal with the amplifier and transmitting the amplified input signal with a radio antenna.

[0292] In Example 160, the subject matter of Example 159 can optionally further include receiving the received signal with the radio antenna, wherein the estimated interference signal approximates leakage from the amplified input signal that is contained in the received signal.

[0293] In Example 161, the subject matter of any one of Examples 152 to 160 can optionally include wherein alternating between updating the first filter and updating the second filter includes updating the first filter and the second filter based on the one or more subsignals and the clean signal.

[0294] In Example 162, the subject matter of any one of Examples 152 to 160 can optionally include wherein alternating between updating the first filter and updating the second filter includes selecting between the first filter and the second filter to identify a current filter to update and a fixed filter to hold constant, applying the fixed filter to the one or more subsignals to obtain a decoupled input signal, and comparing the decoupled input signal to the clean signal to identify at least one weight of the current filter to update.

[0295] In Example 163, the subject matter of Example 162 can optionally include wherein comparing the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update includes determining a cross-correlation vector between the decoupled input signal and the clean signal, identifying a first weight of the current filter to

update based on the cross-correlation vector, and exclusively updating the first weight of the current filter.

[0296] In Example 164, the subject matter of Example 163 can optionally include wherein identifying the first weight of the current filter to update based on the cross-correlation vector includes identifying a maximum-valued element of the cross-correlation vector, and identifying the weight of the current filter with a corresponding element index to the maximum-valued element of the cross-correlation vector as the first weight.

[0297] In Example 165, the subject matter of Example 164 can optionally include wherein identifying the maximum-valued element of the cross-correlation vector includes identifying the element of the cross-correlation vector with the largest real component or the largest imaginary component as the maximum-valued element.

[0298] In Example 166, the subject matter of Example 162 can optionally include wherein comparing the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update includes determining a cross-correlation vector between the decoupled input signal and the clean signal, and updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector.

[0299] In Example 167, the subject matter of Example 166 can optionally include wherein updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector includes updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a coordinate descent optimization scheme.

[0300] In Example 168, the subject matter of Example 166 can optionally include wherein updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector includes updating the at least one weight of the current filter to reduce a

magnitude of the cross-correlation vector according to a recursive least squares optimization scheme.

[0301] In Example 169, the subject matter of Example 166 can optionally include wherein updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector includes updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0302] In Example 170, the subject matter of Example 166 or 169 can optionally include wherein updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector includes exclusively inverting a single bit of the first weight to reduce a magnitude of the cross-correlation vector.

[0303] In Example 171, the subject matter of Example 170 can optionally include wherein updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector further includes evaluating one or more candidate bit inversions of the first weight to identify which of the one or more candidate bit inversions is closest to a predefined numerical difference, and identifying the single bit according to which of the one or more candidate bit inversions is closest to the predefined numerical difference.

[0304] In Example 172, the subject matter of Example 170 can optionally include wherein exclusively inverting the single bit of the first weight to reduce a magnitude of the cross-correlation vector includes inverting the single bit of the first weight according to a dichotomous coordinate descent optimization scheme.

[0305] In Example 173, the subject matter of Example 162 can optionally include wherein selecting between the first filter and the second filter to identify the current filter to update and the fixed filter to hold constant includes selecting the first filter as the current filter and the second filter as the fixed filter, and wherein applying the fixed filter to the one or more

subsignals to obtain the decoupled input signal includes applying previous delay taps of the second filter to previous delay taps of the one or more subsignals during a preprocessing stage that occurs before receiving a most recent delay tap of the one or more subsignals, and after receiving the most recent delay tap of the one or more subsignals, applying a most recent delay tap of the second filter to the most recent delay tap of the one or more subsignals.

[0306] In Example 174, the subject matter of Example 173 can optionally further include performing the preprocessing stage during hidden clock cycles that occur before the most recent delay tap of the one or more subsignals is obtained.

[0307] In Example 175, the subject matter of Example 162 can optionally include wherein selecting between the first filter and the second filter to identify the current filter to update and the fixed filter to hold constant includes selecting the second filter as the current filter and the first filter as the fixed filter, and wherein applying the fixed filter to the one or more subsignals to obtain the decoupled input signal includes exclusively applying the first filter to samples of the one or more subsignals corresponding to a single delay tap of the one or more subsignals to obtain a first element of the decoupled input signal.

[0308] In Example 176, the subject matter of Example 175 can optionally include wherein the remaining elements of the decoupled input signal are time-delayed weighted versions of the first element of the decoupled input signal.

[0309] In Example 177, the subject matter of any one of Examples 152 to 172 can optionally include wherein alternating between updating the first filter and updating the second filter includes selecting between the first filter and the second filter to identify a current filter to update and a fixed filter to hold constant, applying the fixed filter to the one or more subsignals to obtain a decoupled input signal, determining a cross-correlation vector between the decoupled input signal and the clean signal and determining a correlation matrix

of the decoupled input signal, and updating at least one weight of the current filter based on the cross-correlation vector and the correlation matrix.

[0310] In Example 178, the subject matter of Example 177 can optionally include wherein determining the cross-correlation vector between the decoupled input signal and the clean signal includes determining some elements of the cross-correlation vector during a first clock cycle with first calculation circuitry and determining other elements of the cross-correlation vector during a second clock cycle with the same first calculation circuitry, or determining some elements of the correlation matrix during a third clock cycle with second calculation circuitry and determining other elements of the correlation matrix during a second clock cycle with the same second calculation circuitry.

[0311] In Example 179, the subject matter of Example 177 can optionally include wherein selecting between the first filter and the second filter to identify a current filter to update and a fixed filter to hold constant includes selecting the first filter as the current filter and the second filter as the fixed filter, and wherein determining the cross-correlation vector between the decoupled input signal and the clean signal and determining the correlation matrix of the decoupled input signal includes exclusively determining the upper-triangle elements of the correlation matrix.

[0312] In Example 180, the subject matter of Example 179 can optionally include wherein the correlation matrix is Hermitian matrix.

[0313] In Example 181, the subject matter of Example 177 can optionally include wherein selecting between the first filter and the second filter to identify the current filter to update and the fixed filter to hold constant includes selecting the second filter as the current filter and the first filter as the fixed filter, and wherein determining the cross-correlation vector between the decoupled input signal and the clean signal and determining the correlation matrix of the

decoupled input signal includes exclusively determining the a single row of the correlation matrix.

[0314] In Example 182, the subject matter of Example 181 can optionally include wherein a plurality of rows including the single row of the correlation matrix are statistically similar.

[0315] In Example 183, the subject matter of any one of Examples 152 to 182 can optionally include wherein alternating between updating the first filter and updating the second filter includes using shared circuitry to update the first filter and re-using the same shared circuitry to update the second filter.

[0316] In Example 184, the subject matter of any one of Examples 152 to 182 can optionally include wherein alternating between updating the first filter and updating the second filter includes using shared circuitry to update the first filter and re-using the same shared circuitry to update the second filter according to a pipeline clock schedule.

[0317] In Example 185, the subject matter of any one of Examples 152 to 182 can optionally include wherein alternating between updating the first filter and updating the second filter includes updating the first filter at one or more first times and updating the delay tap dimension filter at one or more different second times.

[0318] Example 186 is a communication circuitry arrangement configured to perform the method of any one of Examples 152 to 185.

[0319] In Example 187, the subject matter of Example 186 can optionally be configured as a radio communication device.

[0320] Example 188 is a non-transitory computer readable medium storing instructions that when executed by a controller of a radio communication device direct the radio communication device to perform the method of any one of Examples 152 to 185.

[0321] Example 189 is a communication circuit arrangement including a signal path circuit configured to separately apply a kernel dimension filter and a delay tap dimension filter to an

input signal for an amplifier to obtain an estimated interference signal, a cancelation circuit configured to subtract the estimated interference signal from a received signal to obtain a clean signal, and a filter update circuit configured to alternate between updating the kernel dimension filter and the delay tap dimension filter using the clean signal.

[0322] In Example 190, the subject matter of Example 189 can optionally be configured as a radio communication device and further including a receive chain, a radio antenna, and a transmit chain including the amplifier.

[0323] In Example 191, the subject matter of Example 190 can optionally include wherein the amplifier is configured to amplify the input signal and the radio antenna is configured to transmit the amplified input signal.

[0324] In Example 192, the subject matter of Example 190 or 191 can optionally include wherein the receive chain is configured to receive the received signal via the radio antenna, wherein the estimated interference signal approximates leakage from the transmit chain to the receive chain.

[0325] In Example 193, the subject matter of Example 189 can optionally include wherein the filter update circuit is configured to alternate between updating the kernel dimension filter and updating the delay tap dimension filter using the clean signal by updating the kernel dimension filter in real-time over a first plurality of digital time samples of the input signal and updating the delay tap dimension filter in real-time over a second plurality of digital time samples of the input signal.

[0326] In Example 194, the subject matter of Example 189 can optionally include wherein the input signal includes a plurality of digital time samples, and wherein alternating between updating the kernel dimension filter and updating the delay tap dimension filter using the clean signal includes updating the kernel dimension filter or the delay tap dimension filter over each of the plurality of digital time samples.

[0327] In Example 195, the subject matter of any one of Examples 189 to 194 can optionally further include a kernel generation circuit configured to process the input signal to derive a plurality of kernel signals from the input signal, each of the plurality of kernel signals approximating a non-linear component of a response of the amplifier wherein the signal path circuit is configured to separately apply the kernel dimension filter and the delay dimension filter to the input signal by separately applying the kernel dimension filter and the delay tap dimension filter to the plurality of kernel signals to obtain the estimated interference signal.

[0328] In Example 196, the subject matter of Example 195 can optionally include wherein each of the plurality of kernel signals is composed of a plurality of delay taps, wherein each weight of the kernel dimension filter corresponds to a respective one of the plurality of kernel signals and each weight of the delay tap dimension filter corresponds to a respective one of the plurality of delay taps.

[0329] In Example 197, the subject matter of Example 196 can optionally include wherein the signal path circuit and the filter update circuit are configured to support an adjustable number of kernel signals or an adjustable number of delay taps.

[0330] In Example 198, the subject matter of Example 196 can optionally include wherein the kernel dimension filter approximates the response of the amplifier over the plurality of kernel signals and the delay tap dimension filter approximates the response of the amplifier over the plurality of delay taps.

[0331] In Example 199, the subject matter of any one of Examples 189 to 198 can optionally include wherein the kernel dimension filter and the delay tap dimension filter are vectors.

[0332] In Example 200, the subject matter of any one of Examples 189 to 199 can optionally include wherein the filter update circuit is configured to alternate between updating the kernel dimension filter and updating the delay tap dimension filter by updating the kernel

dimension filter and the delay tap dimension filter based on the input signal and the clean signal.

[0333] In Example 201, the subject matter of any one of Examples 195 to 200 can optionally include wherein the filter update circuit is configured to alternate between updating the kernel dimension filter and the delay tap dimension filter by selecting between the kernel dimension filter and the delay tap dimension filter to identify a current filter to update and a fixed filter to hold constant, applying the fixed filter to the plurality of kernel signals to obtain a decoupled input signal, and comparing the decoupled input signal to the clean signal to identify at least one weight of the current filter to update.

[0334] In Example 202, the subject matter of Example 201 can optionally include wherein the filter update circuit is configured to compare the decoupled input signal to identify the at least one weight of the current filter to update by determining a cross-correlation vector between the decoupled input signal and the clean signal, identifying a first weight of the current filter to update based on the cross-correlation vector, and exclusively updating the first weight of the current filter.

[0335] In Example 203, the subject matter of Example 202 can optionally include wherein the filter update circuit is configured to identify the first weight of the current filter to update based on the cross-correlation vector by identifying a maximum-valued element of the cross-correlation vector, and identifying the weight of the current filter with a corresponding element index to the maximum-valued element of the cross-correlation vector as the first weight.

[0336] In Example 204, the subject matter of Example 203 can optionally include wherein the filter update circuit is configured to identify the maximum-valued element of the cross-correlation vector by identifying an element of the cross-correlation vector with the largest real component or the largest imaginary component as the maximum-valued element.

[0337] In Example 205, the subject matter of Example 201 can optionally include wherein the filter update circuit is configured to compare the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update by determining a cross-correlation vector between the decoupled input signal and the clean signal, and updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector.

[0338] In Example 206, the subject matter of Example 205 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a recursive least squares optimization scheme.

[0339] In Example 207, the subject matter of Example 205 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0340] In Example 208, the subject matter of Example 205 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a coordinate descent optimization scheme.

[0341] In Example 209, the subject matter of Example 205 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by exclusively inverting a single bit of the first weight to reduce a magnitude of the cross-correlation vector.

[0342] In Example 210, the subject matter of Example 209 can optionally include wherein the filter update circuit is further configured to update the at least one weight of the current filter to reduce the cross-correlation vector by evaluating one or more candidate bit inversions of the first weight to identify which of the one or more candidate bit inversions is closest to a predefined numerical difference, and identifying the single bit according to which of the one or more candidate bit inversions is closest to the predefined numerical difference.

[0343] In Example 211, the subject matter of Example 209 can optionally include wherein the filter update circuit is configured to exclusively invert the single bit of the first weight to reduce the magnitude of the cross-correlation vector by inverting the single bit of the first weight according to a dichotomous coordinate descent optimization scheme.

[0344] In Example 212, the subject matter of Example 201 can optionally include wherein the filter update circuit is configured to select between the kernel dimension filter and the delay tap dimension filter to identify the current filter to update and the fixed filter to hold constant by selecting the kernel dimension filter as the current filter and the delay tap dimension filter as the fixed filter, and wherein the filter update circuit is configured to apply the fixed filter to the plurality of kernel signals to obtain the decoupled input signal by applying previous delay taps of the delay tap dimension filter to previous delay taps of the plurality of kernel signals during a preprocessing stage that occurs before receiving a most recent delay tap of the plurality of kernel signals, and after receiving the most recent delay tap of the plurality of kernel signals, applying a most recent delay tap of the delay tap dimension filter to the most recent delay tap of the plurality of kernel signals.

[0345] In Example 213, the subject matter of Example 212 can optionally include wherein the filter update circuit is configured to perform the preprocessing stage during hidden clock cycles that occur before the most recent delay tap of the plurality of kernel signals is obtained.

[0346] In Example 214, the subject matter of Example 201 can optionally include wherein the filter update circuit is configured to select between the kernel dimension filter and the delay tap dimension filter to identify the current filter to update and the fixed filter to hold constant by selecting the delay tap dimension filter as the current filter and the kernel dimension filter as the fixed filter, and wherein the filter update circuit is configured to apply the fixed filter to the plurality of kernel signals to obtain the decoupled input signal by exclusively applying the kernel dimension filter to samples of the plurality of kernel signals corresponding to a single delay tap to obtain a first element of the decoupled input signal.

[0347] In Example 215, the subject matter of Example 214 can optionally include wherein the remaining elements of the decoupled input signal are time-delayed weighted versions of the first element of the decoupled input signal.

[0348] In Example 216, the subject matter of any one of Examples 192 to 215 can optionally include wherein the filter update circuit is configured to alternate between updating the kernel dimension filter and updating the delay tap dimension filter by selecting between the kernel dimension filter and the delay tap dimension filter to identify a current filter to update and a fixed filter to hold constant, applying the fixed filter to the plurality of kernel signals to obtain a decoupled input signal, determining a cross-correlation vector between the decoupled input signal and the clean signal and determining a correlation matrix of the decoupled input signal, and updating at least one weight of the current filter based on the cross-correlation vector and the correlation matrix.

[0349] In Example 217, the subject matter of Example 216 can optionally include wherein the filter update circuit is configured to determine the cross-correlation vector between the decoupled input signal and the clean signal by determining some elements of the cross-correlation vector during a first clock cycle with first calculation circuitry and determining other elements of the cross-correlation vector during a second clock cycle with the same first

calculation circuitry, or determining some elements of the correlation matrix during a third clock cycle with second calculation circuitry and determining other elements of the correlation matrix during a second clock cycle with the same second calculation circuitry.

[0350] In Example 218, the subject matter of Example 216 can optionally include wherein the filter update circuit is configured to select between the kernel dimension filter and the delay tap dimension filter to identify the current filter to update and the fixed filter to hold constant by selecting the kernel dimension filter as the current filter and the delay tap dimension filter as the fixed filter, and wherein the filter update circuit is configured to determine the cross-correlation vector between the decoupled input signal and the clean signal and determining the correlation matrix of the decoupled input signal by exclusively determining the upper-triangle elements of the correlation matrix.

[0351] In Example 219, the subject matter of Example 218 can optionally include wherein the correlation matrix is a Hermitian matrix.

[0352] In Example 220, the subject matter of Example 216 can optionally include wherein the filter update circuit is configured to select between the kernel dimension filter and the delay tap dimension filter to identify the current filter to update and the fixed filter to hold constant by selecting the delay tap dimension filter as the current filter and the kernel dimension filter as the fixed filter, and wherein the filter update circuit is configured to determine the cross-correlation vector between the decoupled input signal and the clean signal and determining the correlation matrix of the decoupled input signal by exclusively determining a single row of the correlation matrix.

[0353] In Example 221, the subject matter of Example 220 can optionally include wherein a plurality of rows including the single row of the correlation matrix are statistically similar.

[0354] In Example 222, the subject matter of any one of Examples 189 to 221 can optionally include wherein the filter update circuit is configured to alternate between updating

the kernel dimension filter and updating the delay tap dimension filter by using shared circuitry to update the kernel dimension filter and re-using the same shared circuitry to update the delay tap dimension filter.

[0355] In Example 223, the subject matter of any one of Examples 189 to 221 can optionally include wherein the filter update circuit is configured to alternate between updating the kernel dimension filter and updating the delay tap dimension filter by using shared circuitry to update the kernel dimension filter and re-using the same shared circuitry to update the delay tap dimension filter according to a pipeline clock schedule.

[0356] In Example 224, the subject matter of Example 223 can optionally include wherein the filter update circuit is configured to alternate between updating the kernel dimension filter and updating the delay tap dimension filter by updating the kernel dimension filter at one or more first times and updating the delay tap dimension at one or more different second times.

[0357] Example 225 is a communication circuit arrangement including a subsignal generation circuit configured to obtain one or more subsignals from an input signal for an amplifier, each of the one or more subsignals representing a non-linear component of an amplifier response and composed of a plurality of delay taps, a signal path circuit configured to separately apply a first filter and a second filter to the one or more subsignals to obtain an estimated interference signal, wherein the first filter approximates the amplifier response over the one or more subsignals and the second filter approximates the amplifier response over the plurality of delay taps, a cancelation circuit configured to subtract the estimated interference signal from the received signal to obtain a clean signal, and a filter update circuit configured to alternate between updating the first filter and alternating the second filter using the clean signal.

[0358] In Example 226, the subject matter of Example 225 can optionally be configured as a radio communication device and further including a receive chain, a radio antenna, and a transmit chain including the amplifier.

[0359] In Example 227, the subject matter of Example 226 can optionally include wherein the amplifier is configured to amplify the input signal and the radio antenna is configured to transmit the amplified input signal.

[0360] In Example 228, the subject matter of Example 226 or 227 can optionally include wherein the receive chain is configured to receive the received signal via the radio antenna, wherein the estimated interference signal approximates leakage from the transmit chain to the receive chain.

[0361] In Example 229, the subject matter of any one of Examples 225 to 228 can optionally include wherein the filter update circuit is configured to alternate between updating the first filter and the second filter by updating the first filter in real-time over a first plurality of digital time samples of the input signal and updating the second filter in real-time over a second plurality of digital time samples of the input signal

[0362] In Example 230, the subject matter of any one of Examples 225 to 228 can optionally include wherein the input signal includes a plurality of digital time samples, and wherein alternating between updating the first filter and updating the second filter using the clean signal includes updating the first filter or the second filter over each of the plurality of digital time samples.

[0363] In Example 231, the subject matter of any one of Examples 225 to 228 can optionally include wherein each of the one or more subsignals correspond to a kernel of the amplifier.

[0364] In Example 232, the subject matter of any one of Examples 225 to 231 can optionally include wherein the first filter and the second filter are vectors.

[0365] In Example 233, the subject matter of any one of Examples 225 to 232 can optionally include wherein each of weight of the first filter corresponds to a respective one of the one or more subsignals and each weight of the second filter corresponds to a respective delay tap of the plurality of delay taps.

[0366] In Example 234, the subject matter of any one of Examples 225 to 233 can optionally include wherein the signal path circuit and the filter update circuit are configured to support an adjustable number of subsignals or an adjustable number of delay taps.

[0367] In Example 235, the subject matter of any one of Examples 225 to 234 can optionally include wherein the filter update circuit is configured to alternate between updating the first filter and updating the second filter by updating the first filter and the second filter based on the one or more subsignals and the clean signal.

[0368] In Example 236, the subject matter of any one of Examples 225 to 234 can optionally include wherein the filter update circuit is configured to alternate between updating the first filter and updating the second filter by selecting between the first filter and the second filter to identify a current filter to update and a fixed filter to hold constant, applying the fixed filter to the one or more subsignals to obtain a decoupled input signal, and comparing the decoupled input signal to the clean signal to identify at least one weight of the current filter to update.

[0369] In Example 237, the subject matter of Example 236 can optionally include wherein the filter update circuit is configured to compare the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update by determining a cross-correlation vector between the decoupled input signal and the clean signal, identifying a first weight of the current filter to update based on the cross-correlation vector, and exclusively updating the first weight of the current filter.

[0370] In Example 238, the subject matter of Example 237 can optionally include wherein the filter update circuit is configured to identify the first weight of the current filter to update based on the cross-correlation vector by identifying a maximum-valued element of the cross-correlation vector, and identifying the weight of the current filter with a corresponding element index to the maximum-valued element of the cross-correlation vector as the first weight.

[0371] In Example 239, the subject matter of Example 238 can optionally include wherein the filter update circuit is configured to identify the maximum-valued element of the cross-correlation vector by identifying the element of the cross-correlation vector with the largest real component or the largest imaginary component as the maximum-valued element.

[0372] In Example 240, the subject matter of Example 236 can optionally include wherein the filter update circuit is configured to compare the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update by determining a cross-correlation vector between the decoupled input signal and the clean signal, and updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector.

[0373] In Example 241, the subject matter of Example 240 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a coordinate descent optimization scheme.

[0374] In Example 242, the subject matter of Example 240 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a recursive least squares optimization scheme.

[0375] In Example 243, the subject matter of Example 240 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

[0376] In Example 244, the subject matter of Example 240 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by exclusively inverting a single bit of the first weight to reduce a magnitude of the cross-correlation vector.

[0377] In Example 245, the subject matter of Example 244 can optionally include wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector further by evaluating one or more candidate bit inversions of the first weight to identify which of the one or more candidate bit inversions is closest to a predefined numerical difference, and identifying the single bit according to which of the one or more candidate bit inversions is closest to the predefined numerical difference.

[0378] In Example 246, the subject matter of Example 244 can optionally include wherein the filter update circuit is configured to exclusively invert the single bit of the first weight to reduce a magnitude of the cross-correlation vector by inverting the single bit of the first weight according to a dichotomous coordinate descent optimization scheme.

[0379] In Example 247, the subject matter of Example 236 can optionally include wherein the filter update circuit is configured to select between the first filter and the second filter to identify the current filter to update and the fixed filter to hold constant by selecting the first filter as the current filter and the second filter as the fixed filter, and wherein the filter update circuit is configured to apply the fixed filter to the one or more subsignals to obtain the

decoupled input signal by applying previous delay taps of the second filter to previous delay taps of the one or more subsignals during a preprocessing stage that occurs before receiving a most recent delay tap of the one or more subsignals, and after receiving the most recent delay tap of the one or more subsignals, applying a most recent delay tap of the second filter to the most recent delay tap of the one or more subsignals.

[0380] In Example 248, the subject matter of Example 247 can optionally include wherein the filter update circuit is configured to perform the preprocessing stage during hidden clock cycles that occur before the most recent delay tap of the one or more subsignals is obtained.

[0381] In Example 249, the subject matter of Example 236 can optionally include wherein the filter update circuit is configured to select between the first filter and the second filter to identify the current filter to update and the fixed filter to hold constant by selecting the second filter as the current filter and the first filter as the fixed filter, and wherein applying the fixed filter to the one or more subsignals to obtain the decoupled input signal by exclusively applying the first filter to samples of the one or more subsignals corresponding to a single delay tap of the one or more subsignals to obtain a first element of the decoupled input signal.

[0382] In Example 250, the subject matter of Example 249 can optionally include wherein the remaining elements of the decoupled input signal are time-delayed weighted versions of the first element of the decoupled input signal.

[0383] In Example 251, the subject matter of any one of Examples 225 to 246 can optionally include wherein the filter update circuit is configured to alternate between updating the first filter and updating the second filter by selecting between the first filter and the second filter to identify a current filter to update and a fixed filter to hold constant, applying the fixed filter to the one or more subsignals to obtain a decoupled input signal, determining a cross-correlation vector between the decoupled input signal and the clean signal and determining a

correlation matrix of the decoupled input signal, and updating at least one weight of the current filter based on the cross-correlation vector and the correlation matrix.

[0384] In Example 252, the subject matter of Example 251 can optionally include wherein the filter update circuit is configured to determine the cross-correlation vector between the decoupled input signal and the clean signal by determining some elements of the cross-correlation vector during a first clock cycle with first calculation circuitry and determining other elements of the cross-correlation vector during a second clock cycle with the same first calculation circuitry, or determining some elements of the correlation matrix during a third clock cycle with second calculation circuitry and determining other elements of the correlation matrix during a second clock cycle with the same second calculation circuitry.

[0385] In Example 253, the subject matter of Example 251 can optionally include wherein the filter update circuit is configured to select between the first filter and the second filter to identify a current filter to update and a fixed filter to hold constant by selecting the first filter as the current filter and the second filter as the fixed filter, and wherein the filter update circuit is configured to determine the cross-correlation vector between the decoupled input signal and the clean signal and determining the correlation matrix of the decoupled input signal by exclusively determining the upper-triangle elements of the correlation matrix.

[0386] In Example 254, the subject matter of Example 253 can optionally include wherein the correlation matrix is Hermitian matrix.

[0387] In Example 255, the subject matter of Example 251 can optionally include wherein the filter update circuit is configured to select between the first filter and the second filter to identify the current filter to update and the fixed filter to hold constant by selecting the second filter as the current filter and the first filter as the fixed filter, and wherein the filter update circuit is configured to determine the cross-correlation vector between the decoupled input

signal and the clean signal and determining the correlation matrix of the decoupled input signal by exclusively determining the a single row of the correlation matrix.

[0388] In Example 256, the subject matter of Example 255 can optionally include wherein a plurality of rows including the single row of the correlation matrix are statistically similar.

[0389] In Example 257, the subject matter of any one of Examples 225 to 256 can optionally include wherein the filter update circuit is configured to alternate between updating the first filter and updating the second filter by using shared circuitry to update the first filter and re-using the same shared circuitry to update the second filter.

[0390] In Example 258, the subject matter of any one of Examples 225 to 256 can optionally include wherein the filter update circuit is configured to alternate between updating the first filter and updating the second filter by using shared circuitry to update the first filter and re-using the same shared circuitry to update the second filter according to a pipeline clock schedule.

[0391] In Example 259, the subject matter of any one of Examples 225 to 256 can optionally include wherein the filter update circuit is configured to alternate between updating the first filter and updating the second filter by updating the first filter at one or more first times and updating the delay tap dimension filter at one or more different second times.

[0392] All acronyms defined in the above description additionally hold in all claims included herein.

[0393] While the invention has been particularly shown and described with reference to specific embodiments, it should be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention as defined by the appended claims. The scope of the invention is thus indicated by the appended claims and all changes which come within the meaning and range of equivalency of the claims are therefore intended to be embraced.

Claims

What is claimed is

1. A communication circuit arrangement comprising:
 - a first kernel dimension filter circuit configured to apply a first kernel dimension filter to a first input signal to estimate a first kernel dimension interference signal from a first amplifier;
 - a second kernel dimension filter circuit configured to apply a second kernel dimension filter to a second input signal to estimate a second kernel dimension interference signal from a second amplifier;
 - a joint delay tap dimension filter configured to apply a joint delay tap dimension filter to a combination of the first kernel dimension interference signal and the second kernel dimension interference signal to obtain an estimated joint interference signal; and
 - a cancelation circuit configured to remove the estimated joint interference signal from a received signal to obtain a clean signal.
2. The communication circuit arrangement of claim 1, further comprising:
 - one or more filter adaptation circuits configured to alternate between a kernel update phase during which the one or more filter adaptation circuits update the first kernel dimension filter and second kernel dimension filter and a delay tap update phase during which the one or more filter adaptation circuits update the joint delay tap dimension filter.
3. The communication circuit arrangement of claim 1, configured as a radio communication device and further comprising one or more receive chains, a first transmit

chain comprising the first amplifier, a second transmit chain comprising the second amplifier, and one or more antennas coupled to the one or more receive chains and the first transmit chain.

4. The communication circuit arrangement of claim 3, wherein a first receive chain of the one or more receive chains is configured to receive the received signal via the one or more antennas, wherein the estimated joint interference signal comprises approximation of leakage from the first transmit chain and the second transmit chain to the first receive chain.

5. The communication circuit arrangement of any one of claims 1 to 4, further comprising a first kernel generation circuit configured to process the first input signal to derive a first plurality of kernel signals from the first input signal, each of the first plurality of kernel signals approximating a non-linear component of a leakage path of the first amplifier, wherein the first kernel dimension filter circuit is configured to apply the first kernel dimension filter to the first input signal by applying the first kernel dimension filter to the first plurality of kernel signals.

6. The communication circuit arrangement of any one of claims 2 to 4, wherein the one or more filter adaptation circuits are configured to alternate between the kernel update phase and the delay tap update phase by:

comparing the combination of the first kernel dimension interference signal and the second kernel dimension interference signal to the clean signal to identify at least one weight of the joint delay tap dimension filter to update during the delay tap update phase.

7. The communication circuit arrangement of any one of claims 2 to 4, wherein the one or more filter adaptation circuits are configured to alternate between the kernel update phase and the delay tap update phase by:

first applying the joint delay tap dimension filter to the first input signal to obtain a first filtered input signal;

first comparing the first filtered input signal to the clean signal to identify at least one weight of the first kernel dimension filter to update;

second applying the joint delay tap dimension filter to the second input signal to obtain a second filtered input signal; and

second comparing the second filtered input signal to the clean signal to identify at least one weight of the second kernel dimension filter to update, wherein the first applying, first comparing, second applying, and second comparing are performed during the kernel update phase.

8. A communication circuit arrangement comprising:

a signal path circuit configured to estimate a first interference signal produced by a first amplifier with a first kernel dimension filter and a first delay tap dimension filter and further configured to estimate a second interference signal produced by a second amplifier with a second kernel dimension filter and a second delay tap dimension filter;

a cancelation circuit configured to subtract a combination of the first interference signal and the second interference signal from a received signal to obtain a clean signal; and

one or more filter adaptation circuits configured to alternate between a kernel update phase during which the one or more filter adaptation circuits update the first kernel dimension filter and the second kernel dimension filter and a delay update phase during which the one or

more filter adaptation circuits update the first delay tap dimension filter and the second delay tap dimension filter.

9. The communication circuit arrangement of claim 8, configured as a radio communication device and further comprising one or more receive chains, a first transmit chain comprising the first amplifier, a second transmit chain comprising the second amplifier, and one or more antennas coupled to the one or more receive chains and the first transmit chain.

10. The communication circuit arrangement of claim 8 or 9, wherein the signal path circuit is configured to estimate the first interference signal by applying the first kernel dimension filter and the first delay tap dimension filter to a first input signal for the first amplifier to obtain the first interference signal and configured to estimate the second interference signal by applying the second kernel dimension filter and the second delay tap dimension filter to a second input signal for the second amplifier to obtain the second interference signal.

11. The communication circuit arrangement of claim 10, wherein the one or more filter adaptation circuits are configured to alternate between the kernel update phase and the delay tap update phase by:

first applying the first kernel dimension filter to the first input signal to obtain a first filtered input signal;

first comparing the first filtered input signal to the clean signal to identify at least one weight of the first delay tap dimension filter to update;

second applying the second kernel dimension filter to the second input signal to obtain a second filtered input signal; and

second comparing the second filtered input signal to the clean signal to identify at least one weight of the second delay tap dimension filter to update, wherein the first applying, first comparing, second applying, and second comparing are performed during the delay tap update phase.

12. The communication circuit arrangement of claim 10, wherein the one or more filter adaptation circuits are configured to alternate between the kernel update phase and the delay tap update phase by:

first applying the first delay tap dimension filter to the first input signal to obtain a first filtered input signal;

first comparing the first filtered input signal to the clean signal to identify at least one weight of the first kernel dimension filter to update;

second applying the second delay dimension filter to the second input signal to obtain a second filtered input signal; and

second comparing the second filtered input signal to the clean signal to identify at least one weight of the second kernel dimension filter to update, wherein the first applying, first comparing, second applying, and second comparing are performed during the kernel update phase.

13. A communication circuit arrangement comprising:

a signal path circuit configured to separately apply a kernel dimension filter and a delay tap dimension filter to an input signal for an amplifier to obtain an estimated interference signal;

a cancelation circuit configured to subtract the estimated interference signal from a received signal to obtain a clean signal; and

a filter update circuit configured to alternate between updating the kernel dimension filter and the delay tap dimension filter using the clean signal.

14. The communication circuit arrangement of claim 13, configured as a radio communication device and further comprising a receive chain, a radio antenna, and a transmit chain including the amplifier.

15. The communication circuit arrangement of claim 13, wherein the filter update circuit is configured to alternate between updating the kernel dimension filter and updating the delay tap dimension filter using the clean signal by:

updating the kernel dimension filter in real-time over a first plurality of digital time samples of the input signal and updating the delay tap dimension filter in real-time over a second plurality of digital time samples of the input signal.

16. The communication circuit arrangement of any one of claims 13 to 15, further comprising a kernel generation circuit configured to process the input signal to derive a plurality of kernel signals from the input signal, each of the plurality of kernel signals approximating a non-linear component of a response of the amplifier

wherein the signal path circuit is configured to separately apply the kernel dimension filter and the delay dimension filter to the input signal by separately applying the kernel dimension filter and the delay tap dimension filter to the plurality of kernel signals to obtain the estimated interference signal.

17. The communication circuit arrangement of claim 16, wherein the filter update circuit is configured to alternate between updating the kernel dimension filter and the delay tap dimension filter by:

selecting between the kernel dimension filter and the delay tap dimension filter to identify a current filter to update and a fixed filter to hold constant;

applying the fixed filter to the plurality of kernel signals to obtain a decoupled input signal; and

comparing the decoupled input signal to the clean signal to identify at least one weight of the current filter to update.

18. The communication circuit arrangement of claim 17, wherein the filter update circuit is configured to compare the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update by:

determining a cross-correlation vector between the decoupled input signal and the clean signal; and

updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector.

19. The communication circuit arrangement of claim 18, wherein the filter update circuit is configured to update the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector by:

updating the at least one weight of the current filter to reduce a magnitude of the cross-correlation vector according to a recursive least squares (RLS) dichotomous coordinate descent (DCD) optimization scheme.

20. A communication circuit arrangement comprising:

a subsignal generation circuit configured to obtain one or more subsignals from an input signal for an amplifier, each of the one or more subsignals representing a non-linear component of an amplifier response and composed of a plurality of delay taps;

a signal path circuit configured to separately apply a first filter and a second filter to the one or more subsignals to obtain an estimated interference signal, wherein the first filter approximates the amplifier response over the one or more subsignals and the second filter approximates the amplifier response over the plurality of delay taps;

a cancelation circuit configured to subtract the estimated interference signal from the received signal to obtain a clean signal; and

a filter update circuit configured to alternate between updating the first filter and alternating the second filter using the clean signal.

21. The communication circuit arrangement of claim 20, configured as a radio communication device and further comprising a receive chain, a radio antenna, and a transmit chain including the amplifier.

22. The communication circuit arrangement of claim 20 or 21, wherein the filter update circuit is configured to alternate between updating the first filter and the second filter by:

updating the first filter in real-time over a first plurality of digital time samples of the input signal and updating the second filter in real-time over a second plurality of digital time samples of the input signal.

23. The communication circuit arrangement of claim 20 or 21, wherein the filter update circuit is configured to alternate between updating the first filter and updating the second filter by:

selecting between the first filter and the second filter to identify a current filter to update and a fixed filter to hold constant;

applying the fixed filter to the one or more subsignals to obtain a decoupled input signal; and

comparing the decoupled input signal to the clean signal to identify at least one weight of the current filter to update.

24. The communication circuit arrangement of claim 23, wherein the filter update circuit is configured to compare the decoupled input signal to the clean signal to identify the at least one weight of the current filter to update by:

determining a cross-correlation vector between the decoupled input signal and the clean signal;

identifying a first weight of the current filter to update based on the cross-correlation vector; and

exclusively updating the first weight of the current filter.

25. The communication circuit arrangement of claim 24, wherein the filter update circuit is configured to identify the first weight of the current filter to update based on the cross-correlation vector by:

identifying a maximum-valued element of the cross-correlation vector; and

identifying the weight of the current filter with a corresponding element index to the maximum-valued element of the cross-correlation vector as the first weight.




FIG 2




FIG 3




FIG 4





FIG 5



X (Shift Reg)	$KM + 2K + M$
R_F	M
B_F	M
R_{PA}	$K(K+1)/2$
B_{PA}	K
Total (I+Q)	$2(K^2 + 2.5K + KM + 1.5M)$




FIG 7A







FIG 7B



$$X_F(t) = W_{PA}^T \cdot X(t)$$

FIG 8






FIG 10




FIG 11




FIG 12




FIG 13




FIG 14




FIG 15




FIG 16




FIG 17




FIG 18





FIG 19






FIG 21




FIG 22

