

REDUCING CONTROL RESPONSE LATENCY WITH DEFINED CROSS-CONTROL BEHAVIOR

[0001] This application is a non-provisional of and claims priority to U.S. Provisional Patent Application No. 61/845,879 filed July 12, 2013. This application is a continuation-in-part of U.S. patent application no. 14/046,819 entitled “Hybrid Systems And Methods For Low-Latency User Input Processing And Feedback” filed October 4, 2013. The entire disclosures of these applications are incorporated herein by reference.

[0002] This application includes material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent disclosure, as it appears in the Patent and Trademark Office files or records, but otherwise reserves all copyright rights whatsoever.

FIELD

[0003] The present invention relates in general to the field of user input, and in particular to user input systems which deliver a low-latency user experience using defined cross-control behavior.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The foregoing and other objects, features, and advantages of the disclosure will be apparent from the following more particular description of embodiments as illustrated in the accompanying drawings, in which reference characters refer to the same parts throughout the various views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating principles of the disclosed embodiments.

[0005] FIG. 1 illustrates a demonstration of the effect of drag latency at 100 ms, 50 ms, 10 ms, and 1 ms in a touch user interface.

[0006] FIG. 2 shows an example of a user interface element for an inbox, where the element has a low latency, low fidelity response to a touch user interaction, as well as a high-latency, high-fidelity response a touch user interaction.

[0007] FIG. 3 shows an example of a user interface of a sliding toggle element. A cursor 310 (represented by the box containing a “cross” character) can be dragged to the target 320 (second empty box, on the right) to activate the UI Element. This element is enabled using both the low latency and high-latency system to provide a touch interaction where moving elements are accelerated 310, thus providing a low-latency experience.

[0008] FIG. 4 shows an illustrative embodiment of a basic architecture of a prototype high-performance touch system used in latency perception studies.

[0009] FIG. 5 shows results of latency perception studies using the prototype device of FIG. 4.

[0010] FIG. 6 shows an example of a user interface element for a button, where the element has a low latency, low fidelity response to a touch user interaction, as well as a high-latency, high-fidelity response a touch user interaction.

[0011] FIG. 7 shows an example of a user interface element for resizable box, where the element has a low latency, low fidelity response to a touch user interaction, as well as a high-latency, high-fidelity response a touch user interaction.

[0012] FIG. 8 shows an example of a user interface element for a scrollable list, where the element has a low latency, low fidelity response to a touch user interaction, as well as a high-latency, high-fidelity response to a touch user interaction.

[0013] FIG. 9 shows an illustrative embodiment of a basic architecture and information flow for a low-latency input device.

[0014] FIG. 10 shows the UI for a volume control. When dragging the slider, a tooltip appears showing a numeric representation of the current setting. This element is enabled using both the low-latency and high-latency system to provide a touch interaction where moving elements are accelerated, thus providing a low-latency experience.

[0015] FIG. 11 shows the system's response for pen input in prior art systems compared to an embodiment of the UI for pen input in the present hybrid feedback user interface system. In the hybrid system, the ink stroke has a low-latency response to pen input, as well as a high-latency response to a pen user input.

[0016] FIG. 12 shows an embodiment of the system where data flows through two overlapping paths through the components of the system to support both high- and low-latency feedback.

[0017] FIG. 13 shows a programming paradigm well known in the art called Model View Controller.

[0018] FIG. 14 shows an embodiment of the system's architecture that supports developing and running applications with blended high and low-latency responses to user input.

[0019] FIG. 15 shows a screenshot illustrating a touch user interface with a dialog box.

[0020] FIG. 16 shows a screenshot illustrating a touch user interface with buttons launching new windows and applications.

DETAILED DESCRIPTION

[0021] The following description and drawings are illustrative and are not to be construed as limiting. Numerous specific details are described to provide a thorough understanding. However, in certain instances, well-known or conventional details are not described in order to avoid obscuring the description. References to one or an embodiment in the present disclosure are not necessarily references to the same embodiment; and, such references mean at least one.

[0022] Reference in this specification to "one embodiment" or "an embodiment" means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the disclosure. The appearances of the phrase "in one embodiment" in various places in the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Moreover, various features are described which may be exhibited by some embodiments and not by others. Similarly, various requirements are described which may be requirements for some embodiments but not for other embodiments.

[0023] This application relates to direct manipulation user interfaces with low-latency such as those disclosed in U.S. Patent Application No. 14/046,819 entitled "Hybrid Systems And Methods For Low-Latency User Input Processing And Feedback" filed October 4, 2013 and U.S. Patent Application No. 61/710,256 entitled "Hybrid Systems And Methods For Low-Latency User Input Processing And Feedback" filed October 5, 2012. This application also relates to the disclosures of U.S. Patent Application No. 13/841,436 filed March 15, 2013 entitled "Low-Latency Touch Sensitive Device," U.S. Patent Application No. 61/798,948 filed March 15, 2013 entitled "Fast Multi-Touch Stylus," U.S. Patent Application No. 61/799,035 filed March 15, 2013 entitled "Fast Multi-Touch Sensor With User-Identification Techniques," U.S. Patent Application No. 61/798,828 filed March 15, 2013 entitled "Fast Multi-Touch Noise Reduction," and U.S.

Patent Application No. 61/798,708 filed March 15, 2013 entitled “Active Optical Stylus.” The entire disclosures of each of these applications are incorporated herein by reference.

[0024] Latency and hybrid systems for low-latency user input processing will first be discussed below, followed by a description of the presently disclosed systems and methods for reducing control response latency with defined cross-control behavior.

Latency and Hybrid Systems

[0025] The present disclosure is directed to systems and methods that provide direct manipulation user interfaces with low latency. Direct physical manipulation of pseudo “real world” objects is a common user interface metaphor employed for many types of input devices, such as those enabling direct-touch input, stylus input, in-air gesture input, as well as indirect devices, including mice, trackpads, pen tablets, etc. For the purposes of the present disclosure, latency in a user interface refers to the time it takes for the user to be presented with a response to a physical input action. Tests have shown that users prefer low latencies and that users can reliably perceive latency as low as 5-10 ms, as will be discussed in greater detail below.

[0026] FIG. 1 illustrates a demonstration of the effect of latency in an exemplary touch user interface at 100 ms (ref. no. 110), 50 ms (ref. no. 120), 10 ms (ref. no. 130), and 1 ms (ref. no. 140) respectively. When dragging an object, increasing latency is reflected as an increasing distance between the user’s finger and the object being dragged (in this case a square user interface element). As can be seen, the effects of latency are pronounced at 100 ms (ref. no. 110) and 50 ms (ref. no. 120), but become progressively less significant at 10 ms (ref. no. 130), and virtually vanish at 1 ms (ref. no. 140). FIG 11 illustrates the effects of latency in an exemplary stylus or pen user interface (1110, 1120). In this example, lag 1120 is visible as an increasing distance between the stylus 1100 tip and the computed stroke 1110. With the introduction of low-latency systems, the distance between the stylus 1100 tip and the computed stroke 1130 would be significantly reduced.

[0027] The systems and methods disclosed in U.S. Patent Application No. 14/046,819 entitled “Hybrid Systems And Methods For Low-Latency User Input Processing And Feedback” provide hybrid touch user interface that provides immediate visual feedback with a latency of less than 10 ms, inter-woven or overlaid with additional visual responses at higher levels of latency. In some embodiments, the designs of these two sets of responses may be designed to be

visually unified, so that the user is unable to distinguish them. In some embodiments, the “low latency” response may exceed 10 ms in latency.

[0028] Latency in a user input device and the system processing its input can have many sources, including the physical sensor that captures touch events; the software that processes touch events and generates output for the display; the display itself; data transmission between components, including bus; data internal storage in either memory stores or short buffers; interrupts and competition for system resources; other sources of circuitry can introduce latency; physical restrictions, such as the speed of light, and its repercussions in circuitry architecture; and mechanical restrictions, such as the time required for a resistive touch sensor to bend back to its ‘neutral’ state.

[0029] Reducing system latency can be addressed through improving latency in one or more of these components. In an embodiment, the presently disclosed systems and methods provide an input device that may achieve 1 ms of latency or less by combining a low-latency input sensor and display with a dedicated processing system. In an embodiment, the presently disclosed systems and methods provide an input device that may achieve 5 ms of latency or less by combining such low-latency input sensor and display with a dedicated processing system. In a further embodiment, the presently disclosed systems and methods provide an input device that may achieve 0.1 ms of latency or less by combining such low-latency input sensor and display with a dedicated processing system. In a further embodiment, the presently disclosed systems and methods provide an input device that may achieve 10 ms of latency or less by combining such low-latency input sensor and display with a dedicated processing system. In an embodiment, in order to achieve such extremely low latencies, the presently disclosed systems and methods may replace conventional operating system (OS) software and computing hardware with a dedicated, custom-programmed field programmable gate array (FPGA) or application-specific integrated circuit (ASIC). In an embodiment, the FPGA or ASIC replaces the conventional OS and computing hardware to provide a low latency response, while leaving a traditional OS and computing hardware in place to provide a higher latency response (which is used in addition in addition to the low latency response). In another embodiment, some or all of the function of the FPGA or ASIC described may be replaced by integrating additional logic into existing components such as but not limited to the graphics processing unit (GPU), input device controller, central processing unit (CPU), or system on a chip (SoC). The low-latency logic can be encoded in hardware, or in software stored-in and/or executed by those or other components.

In embodiments where multiple components are required, communication and/or synchronization may be facilitated by the use of shared memory. In any of these embodiments, responses provided at high or low latency may be blended together, or only one or the other might be provided in response to any given input event.

[0030] In various embodiments, the disclosed systems and methods provide what is referred to herein as “hybrid feedback.” In a hybrid feedback system, some of the basic system responses to input are logically separated from the broader application logic. The result provides a system with a nimble input processor, capable of providing nearly immediate system feedback to user input events, with more feedback based on application logic provided at traditional levels of latency. In some embodiments, these system responses are provided visually. In various embodiments, the low-latency component of a hybrid feedback system may be provided through audio or vibro-tactile feedback. In some embodiments, the nearly immediate feedback might be provided in the same modality as the application-logic feedback. In some embodiments, low-latency feedback may be provided in different modalities, or multiple modalities. An example of an all-visual embodiment is shown in FIG. 2, in this case showing the use of a touch input device. In particular, FIG. 2 shows the result after a user has touched and then dragged an icon 210 representing an inbox. When the user touches the icon 210, a border 220 or other suitable primitive may be displayed. In an embodiment, in an all-visual low-latency feedback, a suitable low-fidelity representation may be selected due to its ease of rendering. In an embodiment, a low-latency feedback may be provided using one or more primitives that can provide a suitable low-fidelity representation. In an embodiment, if the user drags the icon to another place on the touch display 200, a low fidelity border 230 is displayed and may be manipulated (e.g., moved) with a low latency of, for example, 1 ms. Simultaneously, the movement of the icon 210 may be shown with higher latency. In an embodiment, the difference in response between the nearly immediate low-latency response and the likely slower application-logic feedback can be perceived by a user. In another embodiment, this difference in response between the low-latency response and a traditional response is blended and less noticeable or not noticeable to a user. In an embodiment, the nearly immediate feedback may be provided at a lower fidelity than the traditional-path application-logic feedback. In an embodiment, in at least some cases, the low latency response may be provided at similar or even higher fidelity than the application-logic feedback. In an embodiment, the form of low-latency nearly immediate feedback is dictated by application logic, or logic present in the system software (such as the user interface toolkit). For

example, in an embodiment, application logic may pre-render a variety of graphical primitives that can then be used by a low-latency subsystem. Similarly, in an embodiment, a software toolkit may provide the means to develop graphical primitives that can be rendered in advance of being needed by the low latency system. In an embodiment, low-latency responses may be predetermined, or otherwise determined without regard to application and/or system software logic. In an embodiment, individual pre-rendered or partially rendered low-latency responses, or packages of pre-rendered or partially rendered low-latency responses can be pre-loaded into a memory so as to be accessible to the low-latency subsystem in advance of being needed for use in response to a user input event.

[0031] In an embodiment, the modality of low-latency output might be auditory. In an embodiment, the low-latency system may be used, for example, to send microphone input quickly to the audio output system, which may provide users with an “echo” of their own voice being spoken into the system. Such a low-latency output may provide the impression of having the same type of echo characteristics as traditional analog telephones, which allow a user to hear their own voice. In an embodiment, low-latency auditory feedback might be provided in response to user input events (e.g., touch, gesture, pen input, or oral inputs), with a higher latency response provided visually.

[0032] Another illustrative embodiment of a system that employs the present method and system is shown in FIG. 3. In the illustrative system, a cursor 310 (represented by the box containing a “cross” character) can be dragged anywhere on a device’s screen 300. When cursor 310 is dragged to target box 320, the UI action is accepted. If the cursor 310 is dragged elsewhere on the screen 300, the action is rejected. In an embodiment, when dragged, the cursor 310 is drawn with low latency, and thus tracks the user’s finger without perceptible latency. In an embodiment, the target 320 can be drawn with higher latency without impacting user perception. Similarly, in an embodiment, the response 330 of “REJECT” or “ACCEPT” may occur perceptibly later, and thus it can be drawn at a higher latency, e.g., not using the low latency subsystem, without impacting user perception.

[0033] It should be understood that the illustrated embodiment is exemplary. The principles illustrated in FIG. 3 may be applied to any kind of UI element, including all UI elements that are now known, or later developed in the art. Similarly, the principals illustrated in FIG. 3 can be used with substantially any kind of input event on various types of input devices

and/or output devices. For example, in an embodiment, in addition to a “touch” event as illustrated above, input events can include, without limitation, in-air or on-surface gestures, speech, voluntary (or involuntary eye movement, and pen. In an embodiment, once a gesture takes place, the response of any UI element may be bifurcated, where a low-latency response (e.g., a low-fidelity representation of a UI element is presented and responds quickly, for example, in 0.01 ms.), and a non-low-latency response (e.g., a further refined representation of the UI element) is provided with latency commonly exhibited by a system that does not provide accelerated input. In an embodiment, responses may not be split in a hybrid system, and may instead be entirely low latency, with application logic not responsible for the low-latency response otherwise executing with higher latency.

[0034] In an embodiment, touch and/or gesture input events can be achieved using a variety of technologies, including, without limitation, resistive, direct illumination, frustrated total-internal reflection, diffuse illumination, projective capacitive, capacitive coupling, acoustic wave, and sensor-in-pixel. In an embodiment, pen input can be enabled using resistive, visual, capacitive, magnetic, infrared, optical imaging, dispersive signal, acoustic pulse, or other techniques. In an embodiment, gestural input may also be enabled using visual sensors or handheld objects (including those containing sensors, and those used simply for tracking), or without handheld objects, such as with 2D and 3D sensors. Combinations of the sensors or techniques for identifying input events are also contemplated, as are combinations of event types (i.e., touch, pen, gesture, retina movement, etc.) One property technologies to identify or capture input events share is that they contribute to the latency between user action and the system’s response to that action. The scale of this contribution varies across technologies and implementations.

[0035] In a typical multitouch system, there is a path of information flow between the input device and the display that may involve communications, the operating system, UI toolkits, the application layer, and/or ultimately, the audio or graphics controller. Each of these can add latency. Moreover, latency introduced by an operating system, especially a non-real time operating system, is variable. Windows, iOS, OSX, Android, etc. are not real time operating systems, and thus, using these operating systems, there is no guarantee that a response will happen within a certain time period. If the processor is heavily loaded, for example, latency may increase dramatically. Further, some operations are handled at a very low level in the software stack and have high priority. For example, the mouse pointer is typically highly optimized so that

even when the processor is under heavy load, the perceived latency is relatively low. In contrast, an operation such as resizing a photo with two fingers on a touch or gestural system is generally much more computationally intensive as it may require constant rescaling of the image at the application and/or UI toolkit levels. As a result, such operations are rarely able to have a low perceived latency when the processor is under heavy load.

[0036] In a typical multitouch system, the display system (including the graphics system as well as the display itself) may also contribute to latency. Systems with high frame rates may obscure the actual latency through the system. For example, a 60 Hz monitor may include one or more frames of buffer in order to allow for sophisticated image processing effects. Similarly some display devices, such as projectors, include double-buffering in the electronics, effectively doubling the display latency. The desire for 3D televisions and reduced motion artifacts is driving the development of faster LCDs, however, the physics of the liquid crystals themselves make performance of traditional LCD's beyond 480 Hz unlikely. In an embodiment, the low latency system described herein may use an LCD display. In contrast to the performance of an LCD display, OLED or AMOLED displays are capable of response times well below 1ms. Accordingly, in an embodiment, the high performance touch (or gesture) system described herein may be implemented on displays having fast response times, including, without limitation displays based on one or more of the following technologies: OLED, AMOLED, plasma, electrowetting, color-field-sequential LCD, optically compensated bend-mode (OCB or Pi-Cell) LCD, electronic ink, etc.

Latency Perception Studies

[0037] Studies were undertaken to determine what latencies in a direct touch interface users perceive as essentially instantaneous. A prototype device represented in a block diagram in FIG. 4 shows an illustrative embodiment of a basic architecture of a prototype high-performance touch system 400. In an embodiment, the high-speed input device 420 is a multi-touch resistive touch sensor having an active area of 24 cm x 16 cm, and electronics that allow for very high-speed operation. The delay through this sensor is slightly less than 1 ms. In an embodiment, touch data may be transmitted serially over an optical link.

[0038] In the illustrative testing system, the display 460 is a DLP Discovery 4100 kit based on Texas Instruments' Digital Light Processing technology. The illustrative testing system utilizes front-projection onto the touch sensor thus eliminating parallax error that might disturb a

user's perception of finger and image alignment. The DLP projector employed uses a Digital Micromirror Device (DMD), a matrix of mirrors which effectively turns pixels on or off at very high speed. The high speed of the mirrors may be used to change the percentage time on vs. off to create the appearance of continuous colored images. In an embodiment, where only simple binary images are used, these can be produced at an even higher rate. In the illustrative testing system, the projector development system displays 32,000 binary frames/second at 1024x768 resolution with latency under 40 μ s. In the illustrative testing system to achieve this speed, the video data is streamed to the DMD at 25.6 Gbps.

[0039] In the illustrative testing system, to achieve minimal latency, all touch processing is performed on a dedicated FPGA 440 –no PC or operating system is employed between the touch input and the display of low latency output. The DLP kit's onboard XC5VLX50 application FPGA may be used for processing the touch data and rendering the video output. A USB serial connection to the FPGA allows parameters to be changed dynamically. In the illustrative testing system, latency can be adjusted from 1 ms to several hundred ms with 1 ms resolution. Different testing modes can be activated, and a port allows touch data to be collected for analysis.

[0040] In the illustrative testing system, to receive touch data from the sensor 420, the system communicates through a custom high-speed UART. To minimize latency, a baud rate of 2 Mbps can be used, which represents a high baud rate that can be used without losing signal integrity due to high frequency noise across the communication channel. In the illustrative testing system, the individual bytes of compressed touch data are then processed by a touch detection finite state machine implemented on the FPGA 440. The finite-state machine (FSM) simultaneously decodes the data and performs a center-of-mass blob-detection algorithm to identify the coordinates of the touches. In the illustrative testing system, the system is pipelined such that each iteration of the FSM operates on the last received byte such that no buffering of the touch data occurs.

[0041] In the illustrative testing system, the touch coordinates are then sent to a 10-stage variable delay block. Each delay stage is a simple FSM with a counter and takes a control signal that indicates the number of clock cycles to delay the touch coordinate, allowing various levels of latency. The delay block latches the touch sample at the start of the iteration and waits for the appropriate number of cycles before sending the sample and latching the next. The delay block

therefore lowers the sample rate by a factor of the delay count. In an embodiment, to keep the sample rate at a reasonable level, 10 delay stages can be used, so that, for example, to achieve 100 ms of latency, the block waits 10 ms between samples for a sample rate of 100 Hz. In the illustrative testing system, to run basic applications, a MicroBlaze soft processor is used to render the display.

[0042] In an embodiment, the testing system may use a hard coded control FSM in place of the MicroBlaze for improved performance. In an embodiment another soft processor may be used. In the illustrative testing system, the MicroBlaze is a 32-bit Harvard architecture RISC processor optimized to be synthesized in Xilinx FPGAs. The MicroBlaze soft processor instantiation allows the selection of only the cores, peripherals, and memory structures required. In the illustrative testing system, in addition to the base MicroBlaze configuration, an interrupt controller can be used, for example, GPIOs for the touch data, a GPIO to set the variable latency, a BRAM memory controller for the image buffer, and a UART unit to communicate with a PC. In the illustrative testing system, the MicroBlaze is clocked at 100 MHz. The MicroBlaze uses an interrupt system to detect valid touch coordinates. A touch ready interrupt event is generated when valid touch data arrives on the GPIOs from the delay block, and the corresponding image is written to the image buffer. Because of the non-uniform nature of an interrupt-based system, the exact latency cannot be computed, but, by design, it is insignificant in comparison to the 1 ms latency due to the input device.

[0043] In the illustrative testing system, the image buffer is synthesized in on-chip BRAM blocks. These blocks can provide a dual-port high-speed configurable memory buffer with enough bandwidth to support high frame-rate display. In the illustrative testing system, the image buffer is clocked at 200 MHz with a bus width of 128 bits for a total bandwidth of 25.6 Gbps, as needed by the DLP. Finally, the DMD controller continuously reads out frames from the image buffer and generates the signals with appropriate timing to control the DMD.

[0044] In the illustrative testing system, user input is sent simultaneously to a traditional PC, and is processed to produce a traditional, higher latency, response. This higher latency response is output by a traditional data projector, aligned to overlap with the projected lower latency response.

[0045] Studies were conducted to determine the precise level of performance that users are able to perceive when performing common tasks on a touch screen interface. To that end, studies were conducted to determine the just-noticeable difference (JND) of various performance levels. JND is the measure of the difference between two levels of a stimulus which can be detected by an observer. In this case, the JND is defined as the threshold level at which a participant is able to discriminate between two unequal stimuli – one consistently presented at the same level, termed the reference, and one whose value is changed dynamically throughout the experiment, termed the probe. A commonly accepted value for the JND at some arbitrary reference value is a probe at which a participant can correctly identify the reference 75% of the time. A probe value that cannot be distinguished from the reference with this level of accuracy is considered to be “not noticeably different” from the reference.

[0046] Studies were conducted to determine the JND level of the probe latency when compared to a maximum performance of 1 ms of latency, which served as the reference. While such a determination does not provide an absolute value for the maximum perceptible performance, it can serve as our “best case” floor condition against which other levels of latency can be measured, given that it was the fastest speed our prototype could achieve. It was found participants are able to discern latency values that are significantly lower (< 20 ms) which typical current generation hardware (e.g., current tablet and touch computer) provides (~50-200 ms).

[0047] Ten right-handed participants (3 female) were recruited from the local community. Ages ranged between 24 and 40 (mean 27.80, standard deviation 4.73). All participants had prior experience with touch screen devices, and all participants owned one or more touch devices (such as an iOS- or Android- based phone or tablet). Participants were repeatedly presented with pairs of latency conditions: the reference value (1 ms) and the probe (between 1 and 65 ms of latency). Participants dragged their finger from left to right, then right to left on the touch screen display. While any dragging task would have been suitable, left/right movements reduce occlusion in high-latency cases. Participants were asked to move in both directions to ensure they did not “race through” the study. Beneath the user’s contact point, the system rendered a solid white $2\text{ cm} \times 2\text{ cm}$ square as seen in FIG. 1. The speed of movement was left to be decided by the participants. The order of the conditions was randomized for each pair. The study was designed as a two-alternative forced-choice experiment; participants were instructed to choose, within each trial, which case was the reference (1 ms) value and were not

permitted to make a “don’t know” or “unsure” selection. After each pair, participants informed the experimenter which of the two was “faster”.

[0048] In order for each trial to converge at a desired JND level of 75%, the amount of added latency was controlled according to an adaptive staircase algorithm. Each correct identification of the reference value caused a decrease in the amount of latency in the probe, while each incorrect response caused the probe’s latency to increase. In order to reach the 75 % confidence level, increases and decreases followed the simple weighted up-down method described by Kaernbach (Kaernbach, C. 1991. *Perception & Psychophysics* 49, 227-229), wherein increases had a three-fold multiplier applied to the base step size, and decreases were the base step size (initially 8 ms).

[0049] When a participant responded incorrectly after a correct response, or correctly after an incorrect response, this was termed a reversal as it caused the direction of the staircase (increasing or decreasing) to reverse. The step size, initially 8 ms, was halved at each reversal, to a minimum step size of 1 ms. This continued until a total of 10 reversals occurred, resulting in a convergence at 75% correctness. Each participant completed eight staircase “runs.” Four of these started at the minimum probe latency (1 ms) and four at the maximum (65 ms). The higher starting value of the staircase was chosen because it roughly coincides with commercial offerings, and because pilot testing made it clear that this value would be differentiated from the 1 ms reference with near 100% accuracy, avoiding ceiling effects. Staircases were run two at a time in interleaved pairs to prevent response biases that would otherwise be caused by the participants’ ability to track their progress between successive stimuli. Staircase conditions for each of these pairs were selected at random without replacement from possibilities (2 starting levels $\times$ 4 repetitions). The entire experiment, including breaks between staircases, was completed by each participant within a single 1-hour session.

[0050] The study was designed to find the just-noticeable difference (JND) level for latency values greater than 1 ms. This JND level is commonly agreed to be the level where the participant is able to correctly identify the reference 75% of the time. Participant JND levels ranged from 2.38 ms to 11.36 ms, with a mean JND across all participants of 6.04 ms (standard deviation 4.33 ms). JND levels did not vary significantly across the 8 runs of the staircase for each participant. Results for each participant appear in FIG. 5.

[0051] The results show participants were able to discern differences in latency far below the typical threshold of consumer devices (50-200 ms). It is noted that participants were likely often determining latency by estimating the distance between the onscreen object and their finger as it was moved around the touch screen; this is an artifact of input primitives used in UIs (specifically, dragging). Testing a different input primitive (tapping, for example) would exhibit different perceptions of latency. Results confirm that an order-of magnitude improvement in latency would be noticed and appreciated by users of touch devices.

An Architecture for a Low-Latency Direct Touch Input Device

[0052] In an embodiment, a software interface may be designed that enables application developers to continue to use toolkit-based application design processes, but enable those toolkits to provide feedback at extremely low latencies, given the presence of a low-latency system. In an embodiment, the systems and methods outlined in the present disclosure may be implemented on the model-view-controller (“MVC”) model of UI development, upon which many UI toolkits are based. An MVC permits application logic to be separated from the visual representation of the application. In an embodiment, an MVC may include, a second, overlaid *de facto* view for the application. In particular, in an embodiment, touch input receives an immediate response from the UI controls, which is based in part on the state of the application at the time the touch is made. The goal is to provide nearly immediate responses that are contextually linked to the underlying application.

[0053] Previous work on application independent visual responses to touch are completely separate from even the visual elements of the UI, adding visual complexity. In an embodiment, according to the systems and methods outlined herein, a set of visual responses are more fully integrated into the UI elements themselves so as to reduce visual complexity. Thus, in an embodiment, where the particular visuals shown provide a *de facto* “mouse pointer” for touch, the goal is to integrate high performance responses into the controls themselves (i.e., the onscreen content being ‘touched’, ‘manipulated’, etc.), providing a more unified visualization. None the less, in an embodiment, the systems and methods allow the rendering of context-free responses by the low-latency subsystem, which are later merged with responses from the high-latency subsystem. In an embodiment, visuals need not be presented in the same rendering pipeline as the rest of the system’s response. Instead, a system or method which utilizes hybrid

feedback as discussed herein may present lower latency responses to user input in addition to the higher latency responses generated by the traditional system.

[0054] Thus, in an embodiment, accelerated input interactions are designed such that the traditional direct-touch software runs as it would normally, with a high-latency responses, while an additional set of feedback, customized for the UI element, is provided at a lower latency; with a target of user-imperceptible latency. In an embodiment, these two layers are combined by superimposing two or more images. In an embodiment, two combined images may include one projected image from the low-latency touch device, and a second from a traditional projector connected to a desktop computer running custom touch software, receiving input from the low-latency subsystem.

[0055] The two projector solution described above is meant only to serve as one particular embodiment of the more general idea of combining a low latency response and a traditional response. In an embodiment, the visual output from the low and high-latency subsystems are logically combined in the display buffer or elsewhere in the system before being sent to the display, and thus, displayed. In an embodiment, transparent, overlapping displays present the low and high-latency output to the user. In an embodiment, the pixels of a display are interlaced so that some are controlled by the low latency subsystem, and some are controlled by the high-latency sub-system; through interlacing, these displays may appear to a user to overlap. In an embodiment, frames presented on a display are interlaced such that some frames are controlled by the low latency subsystem and some frames are controlled by the high-latency subsystem; through frame interlacing, the display may appear to a user to contain a combined image. In an embodiment, the low-latency response may be generated predominantly or entirely in hardware. In an embodiment, the low-latency response may be generated from input sensor data received directly from the input sensor. In an embodiment, the low-latency response is displayed by having a high bandwidth link to the display hardware.

[0056] In designing a user interface for a low-latency subsystem, one or more of the following constraints may be considered:

- Information: any information or processing needed from the high-latency subsystem in order to form the system's response to input will, necessarily, have high latency, unless such information or processing is e.g., pre-rendered or pre-served.

- Performance: the time allowed for formation of responses in low latency is necessarily limited. Even with hardware acceleration, the design of responses must be carefully performance-driven to guarantee responses meet the desired low latency.
- Fidelity: the fidelity of the rendered low-latency image may be indistinguishable from the higher-latency rendering (indeed, it may be pre-rendered by the high latency system); additional constraints may be placed on fidelity to improve performance, such as, e.g., that visuals are only monochromatic, and/or limited to visual primitives, and/or that the duration or characteristics of audio or haptic responses are limited. Constraints of this type may be introduced by various elements of the system, including acceleration hardware or by the output hardware (such as the display, haptic output device, or speakers).
- Non-Interference: in embodiments where responses are hybridized combinations, some of the application's response may be generated in the low-latency layer, and some in the high-latency layer, a consideration may be how the two are blended, e.g., to provide a seamless response to the user's input. In an embodiment, low-latency responses do not interfere with any possible application response, which will necessarily occur later. In an embodiment, interference may occur between a low-latency response and the traditional response, but the interference may be handled through design, or through blending of the responses.

[0057] In an embodiment, a design process was conducted to create a set of visual UI controls with differentiated low and high latency visual responses to touch. A metaphor was sought which would enable a seamless transition between the two layers of response. These visualizations included such information as object position and state. The designs were culled based on feasibility using the above-described constraints. The final design of such embodiment was based on a heads-up display (HUD) metaphor, similar to the visualizations used in military aircraft. The HUD was suitable, since traditional HUDs are geometrically simple, and it is relatively easy to implement a geometrically simple display at an authentic fidelity. The HUD represents just one example of two visual layers being combined, though in many HUDs, a computerized display is superimposed on video or the "real world" itself. Accordingly, a HUD is generally designed to be non-interfering.

[0058] Based on the HUD metaphor, an exemplary set of touch event and UI element-specific low-latency layer visualizations were developed for a set of UI elements found in many

direct-touch systems. These exemplary elements are both common and representative; their interactions (taps, drags, two-finger pinching) cover the majority of the interaction space used in current direct-touch devices. The low-latency responses developed in such an embodiment are described in Table 1, and they are shown in FIG. 6-8.

Element	Touch Down	Touch Move	Touch Up
Button (FIG. 6)	Bounds outlined 610	(none)	If within bounds, 2 nd outline 620, else none
Draggable/Resizable (FIG. 7)	Bounds outlined 710	Outline changes and moves with input position 720 and/or scales with input gesture 730	Outline 710 fades when high-latency layer catches up
Scrollable List (FIG. 8)	List item outlined 810	If scroll gesture, list edges highlight 830 to scroll distance. If during scroll gesture, edge highlights 840) fade as high- latency layer catches up	If list item selection, outline 820 scales down and fades

Table 1: Accelerated visuals for each element and touch event, which compliment standard high latency responses to touch input.

[0059] These three elements represent broad coverage of standard UI toolkits for touch input. Most higher-order UI elements are composed of these simpler elements (e.g. radio buttons and checkboxes are both “buttons,” a scrollbar is a “draggable/resizable” with constrained translation and rotation). The accelerated input system and method described herein depends on the marriage of visuals operating at two notably different latency levels; this latency difference has been incorporated into the design of low-latency visualizations. In an embodiment, users may be informed of the state of both systems, with a coherent synchronization as the visual layers come into alignment. In an embodiment, a user may be able to distinguish between the high and low latency portions of system feedback. In an embodiment, the visual elements are blended in a manner that provides no apparent distinction between the low-latency response and the traditional response.

[0060] In an embodiment, an application developer utilizes a toolkit to build their application through the normal process of assembling GUI controls. Upon execution, the UI elements bifurcate their visualizations, with high- and low- latency visualizations rendered and overlaid on a single display. An embodiment of information flow through such a system is as shown in FIG. 9. Information flows into the system from an input device 910 and is initially processed by an input processing unit (IPU) 920, programmed via an IPU software toolkit 930. UI events are then processed in parallel by two subsystems, a low-latency, low fidelity subsystem 940, and a high-latency subsystem 950 such as, for example, conventional software running in a conventional software stack. In an embodiment, the low-latency, low fidelity subsystem 940 may be implemented in hardware, such as the FPGA 440 of FIG. 4.

[0061] The bifurcation described in this embodiment creates a fundamental communication problem where any parameterization of the initial responses provided by the low-latency subsystem 940 required by application logic must be defined before the user begins to give input. Any response which requires processing at the time of presentation by the application will introduce a dependency of the low-latency system 940 upon the high-latency system 950, and may therefore introduce lag back into the system. In an embodiment, later stages of the low-latency system's 940 response to input may depend on the high latency subsystem 950. In an embodiment, dependency of the later stages of a low-latency subsystem's 940 response to input on the high latency subsystem 950 is managed such that the dependency does not introduce additional latency. In an embodiment the dependency would be avoided entirely.

[0062] In an embodiment, UI element logic may be built into the low-latency subsystem. Between user inputs, the application executing in the high-latency subsystem 950, has the opportunity to provide parameters for the low-latency subsystem's 940 model of the UI elements. Thus, in an embodiment, the MVC model of UI software design may be extended by providing a separate controller responsible for low-latency feedback. In an embodiment, in the software design, one or more of the following can be specified for each control:

- Element type (e.g., button, draggable object, scrollable list, etc.).
- Bounding dimensions (e.g., x position, y position, width, height, etc.).
- Conditional: additional primitive information (e.g., size of list items in the case of a scrollable list, etc.).

[0063] In an embodiment, logic for a given element-type's response to touch input is stored in the low-latency subsystem 940. Further parameterization of the low-latency subsystem's responses to user input could be communicated in the same manner, allowing a greater degree of customization. In an embodiment, sensor data is processed to generate events (or other processed forms of the input stream), which are then separately distributed to the low-latency subsystem 940 and to the high-latency subsystem 950. Events may be generated at different rates for the low-latency subsystem 940 and high-latency subsystem 950, because the low-latency subsystem is capable of processing events faster than the high-latency subsystem, and sending events to the high-latency sub-system at a high rate may overwhelm that subsystem. The low- and high- latency subsystems' response to user input is therefore independent but coordinated. In an embodiment, one subsystem acts as the "master," setting state of the other subsystem between user inputs. In an embodiment, the relationship between the low- and high- latency subsystems includes synchronization between the two subsystems. In an embodiment, the relationship between the low- and high- latency subsystems includes the ability of the high-latency subsystem to offload processing to the low-latency subsystem 940. In an embodiment, the relationship between the low- and high- latency subsystems includes the ability of the low-latency subsystem 940 to reduce its processing Load and/or utilize the high-latency subsystem 950 for pre-processing or pre-rendering. In an embodiment, a second graphical processing and output system's response is dependent upon a first graphical processing and output system, and state information is passed from the first graphical processing and output system to the second graphical processing and output system. In such embodiments, information passed from the first graphical processing and output system to the second graphical processing and output system is comprised of one or more pieces of data describing one or more of the graphical elements in the user interface. This data may be, e.g., the size, the location, the appearance, alternative appearances, response to user input, and the type of graphical elements in the user interface. The data passed from the first graphical processing and output system to the second graphical processing and output system may be stored in high-speed memory available to the second graphical processing and output system. The passed data may describe the appearance and/or behavior of a button, a slider, a draggable and/or resizable GUI element, a scrollable list, a spinner, a drop-down list, a menu, a toolbar, a combo box, a movable icon, a fixed icon, a tree view, a grid view, a scroll bar, a scrollable window, or a user interface element.

[0064] In an embodiment, an input processing system performs decimation on the user input signals before they are received by one or both of the first or second graphical processing and output systems. The decimated input signals or non-decimated signals are chosen from the set of all input signals based on information about the user interface sent from the first graphical processing and output system. The decimation of input signals may be performed by logically combining the set of input signals into a smaller set of input signals. Logical combination of input signals may be performed through windowed averaging. The decimation considers the time of the user input signals when reducing the size of the set of input signals. The logical combination of input signals can be performed through weighted averaging. In an embodiment, the user input signals received by the first and second graphical processing and output systems have been differentially processed.

[0065] In an embodiment, communication between the high-latency and low-latency layers may be important. Some points which are considered in determining how the high- and low- latency subsystems remain synchronized are described below:

- Latency differences: Low-latency responses may use information about the latency difference between the high- and low- latency layers in order to synchronize responses. In an embodiment, these latency values are static, and thus preprogrammed into the FPGA. In an embodiment where latency levels may vary in either subsystem, it may be advantageous to fix the latency level at an always-achievable constant rather than having a dynamic value that may become unsynchronized, or provide an explicit synchronization mechanism. In an embodiment where latency levels may vary in either subsystem, a dynamic value may be used, however, care should be taken to avoid becoming unsynchronized. In an embodiment where latency levels may vary in either subsystem, an explicit synchronization mechanism may be provided between the subsystems 940, 950.
- Hit testing: Hit testing decisions are often conditional on data regarding the visual hierarchy and properties of visible UI elements. In an embodiment, this consideration can be resolved by disallowing overlapping bounding rectangles, requiring a flat, 'hit test friendly' map of the UI. In an embodiment separate hit testing may provide the necessary information (object state, z-order, and listeners) to the low-latency subsystem. In an embodiment both the low- and high-latency subsystems may conduct hit testing in

parallel. In an embodiment the low-latency subsystem conducts hit testing, and provides the results to the high-latency subsystem.

- Conditional responses: Many interface visualizations are conditional not only on immediate user input, but on further decision-making logic defined in the application logic.

[0066] Two illustrative examples of conditional response logic are as follows: Consider a credit-card purchase submission button, which is programmed to be disabled (to prevent double billing) when pressed, but only upon validation of the data entered into the form. In such a case, the behavior of the button is dependent not only on an immediate user interaction, but is further conditional on additional information and processing. Consider also linked visualizations, such as the one shown in FIG. 10. In this case, feedback is provided to the user not only by the UI element they are manipulating 1010, but also by a second UI element 1020. These examples could be programmed directly into a low-latency subsystem.

[0067] In an embodiment, the division between the high- and low- latency subsystems may be independent of any user interface elements. Indeed, the division of responsibility between the subsystems can be customized based on any number of factors, and would still be possible in systems that lack a user interface toolkit, or indeed in a system which included mechanisms to develop applications both within and without the use of a UI toolkit which might be available. In an embodiment, the division of responsibility between the two subsystems can be dynamically altered while the subsystems are running. In an embodiment, the UI toolkit itself may be included within the low-latency subsystem. The ability to customize responses can be provided to application developers in a number of ways without departing from the systems and methods herein described. In an embodiment, responses may be customized as parameters to be adjusted in UI controls. In an embodiment, responses may be customized by allowing for the ability to provide instructions directly to the low-latency subsystem, in code which itself executes in the low-latency subsystem, or in another high- or low-latency component. In an embodiment, the state of the low-latency subsystem could be set using data generated by application code, e.g., at runtime.

[0068] While many of the examples described above are provided in the context of a touch input, other embodiments are contemplated, including, without limitation, pen input,

mouse input, indirect touch input (e.g., a trackpad), in-air gesture input, oral input and/or other input modalities. The architecture described would be equally applicable to any sort of user input event, including, without limitation, *mixed* input events (i.e., supporting input from more than one modality). In an embodiment, mixed input devices may result in the same number of events being generated for processing by each of the low- and high- latency subsystems. In an embodiment, mixed input devices would be differentiated in the number of events generated, thus, for example, touch input might have fewer events than pen input. In an embodiment, each input modality comprises its own low-latency subsystem. In an embodiment, in systems comprising multiple low-latency subsystems for multiple input modalities, the subsystems might communicate to coordinate their responses. In an embodiment, in systems comprising multiple low-latency subsystems for multiple input modalities, the multiple subsystems may share a common memory area to enable coordination.

Input Processing

[0069] In an embodiment of the invention, low-latency input data from the input hardware is minimally processed into a rapid stream of input events. This stream of events is sent directly to the low-latency sub-system for further processing. Events from this same stream may then be deleted, or the stream may be otherwise reduced or filtered, before being sent to the high-latency subsystem. Events may be generated at different rates for the low-latency subsystem 940 and high-latency subsystem 950 because the low-latency subsystem is capable of processing events faster than the high-latency subsystem, and sending events to the high-latency sub-system at a high rate may overwhelm that subsystem. The low- and high-latency subsystems' response to user input may therefore be independent but coordinated.

[0070] The reduction of events can be optimized. In an embodiment, representative events may be selected among candidate events based on criteria associated with one or more of the application, the UI element, the input device, etc. An example of this for pen input when the user is drawing digital ink strokes might include selecting events which fit best to the user's drawn stroke. Another example for speech input is to favor events where subsequent events in the output stream would have similar volume, thereby "evening out" the sound coming from the microphone. Another example for touch input is to favor events which would result in the output event stream having a consistent speed, providing more "smooth" output. This form of intelligent reduction acts as an intelligent filter, without reducing performance of the high-latency

subsystem. In an embodiment, new events (e.g., consolidated events or pseudo-events) could be generated which represent an aggregate of other events in the input stream. In an embodiment, new events (e.g., corrected events, consolidated events or pseudo-events) may be generated that represent a more desirable input stream, e.g., a correction or smoothing. For example, for in-air gesture input, for every 10 events from the high-speed input device, the high-latency subsystem may be sent the same number or fewer events which provide an “average” of actual input events, thus smoothing the input and removing jitter. New events could also be generated which are an amalgam of multiple “desired” levels of various parameters of an input device. For example, if the intelligent reductions of the tilt and pressure properties of a stylus would result in the selection of different events, a single, new, event object could be created (or one or more existing event objects modified) to include the desired values for each of these properties.

[0071] In an embodiment, an IPU or low-latency subsystem system might be used to provide the high-latency system with processed input information. One or more of methods could be used to coordinate the activities of the two subsystems. These include:

- a. In an embodiment, the low-latency subsystem can respond to all user input immediately, but wait for the user to stop the input (e.g. lifting a finger or pen, terminating a gesture) before providing the input to the high-latency system. This has an advantage of avoiding clogging the system during user interaction while still processing the totality of the data.
- b. In an embodiment, the low-latency system can provide a reduced estimate of input in near real-time; and may optionally store a complete input queue that can be available to the high-latency system upon request.
- c. In an embodiment, user feedback may be divided into two steps. The first, a low-latency feedback, would provide a rough, immediate representation of user input 1130 in FIG. 11. The second, a high-latency system response 1140, could replace the first 1130, whenever the high-latency system is able to compute a refined response, for example after lift-off of the pen 1150 tip. Alternatively, the high latency feedback could be continuously “catching up” to (and possibly subsuming) the low latency feedback.
- d. In an embodiment, the low-latency system can infer simple gesture actions from the input stream, and thus generate gesture events which are included in the input queue in addition to, or replacing, the raw events.

- e. In an embodiment, an IPU or low-latency subsystem can use multiple input positions to predict future input positions. This prediction can be passed along to the high-latency subsystem to reduce its effective latency.
- f. In an embodiment, algorithms which may benefit from additional samples, or earlier detection, are executed in the IPU or low-latency subsystem. In an embodiment, the execution of these events can be limited in time. For example, the initial 50 events can be used to classify an input as a particular finger, or to differentiate between finger and pen inputs. In an embodiment, these algorithms can run continuously.
- g. In an embodiment, the process of the low-latency subsystem passing a stream of events to the high-latency subsystem might be delayed in order to receive and process additional sequential or simultaneous related inputs which might otherwise be incorrectly regarded as unrelated inputs. For example, the letter "t" is often drawn as two separate, but related, strokes. In the normal course, the portion of the input stream passed from the low-latency to the high-latency system would include a "pen up" signal at the end of drawing the first line. In an embodiment, the reduction process waits for the very last frame of input within the sample window to pass along an "up" event, in case the pen is again detected on the display within the window, thus obviating the need for the event.

Hardware Architecture

[0072] In an embodiment, data flows through two overlapping paths through the components of the system to support both high- and low- latency feedback. FIG. 12 shows one such system, which includes an Input Device 1210, an IPU 1220, a System Bus 1230, a CPU 1240 and a GPU 1280 connected to a Display 1290. A User 1200 performs input using the Input Device 1210. This input is sensed by the IPU 1220 which in various embodiments can be either an FPGA, ASIC, or additional software and hardware logic integrated into a GPU 1280, MPU or SoC. At this point, the control flow bifurcates and follows two separate paths through the system. For low-latency responses to input, the IPU 1220 sends input events through the System Bus 1230 to the GPU 1280, bypassing the CPU 1240. The GPU 1280 then rapidly displays feedback to the User 1200. For high-latency response to input, the IPU 1220 sends input events through the System Bus 1230 to the CPU 1240, which is running the graphical application and which may interact with other system components. The CPU 1240 then sends commands via the

System Bus 1230 to the GPU 1280 in order to provide graphical feedback to the User 1200. The low-latency path from Input Device 1210 to IPU 1220 to System Bus 1230 to GPU 1280 is primarily hardware, and operates with low-latency. The high-latency path from Input Device 1210 to IPU 1220 to System Bus 1230 to CPU 1240 back to System Bus 1230 to GPU 1280 is high-latency due to the factors described earlier in this description. In a related embodiment, the Input Device 1210 communicates directly with the GPU 1280 and bypasses the System Bus 1230.

[0073] FIG. 13 shows a familiar programming paradigm called Model View Controller. In this paradigm, the User 1300 performs input on the Controller 1310, which in turn manipulates the Model 1320 based on this input. Changes in the Model 1320 result in changes to the View 1330, which is observed by the User 1300. Some of the latency addressed by the present invention is due to latency in the input, communication among these components, and display of the graphics generated by the View 1330 component.

[0074] FIG. 14 shows an embodiment of an architecture that supports developing and running applications on a system with blended high- and low- latency responses to user input. The User 1400 performs input with the input device 1410. This input is received by the IPU 1420. The IPU 1420 sends input events simultaneously to the Controller 1430 running in the high-latency subsystem via traditional mechanisms and to the ViewModel(L) 1490 running in the low-latency subsystem. Input is handled by the Controller 1430, which manipulates the Model 1440 running in the high-latency subsystem, which may interact with data in volatile memory 1450, fixed storage 1470, network resources 1460, etc. (all interactions that introduce lag). Input events received by the ViewModel(L) 1490 result in changes to the ViewModel(L) which are reflected in changes to the View(L) 1491, which is seen by the User 1400. Changes to the Model 1440 result in changes to the high-latency subsystem's View(H) 1480, which is also seen by the User 1400. In an embodiment, these two types of changes seen by the user are shown on the same display. In an embodiment, these two types of changes are reflected to the user via other output modalities (such as, e.g., sound or vibration). In an embodiment, between inputs, the Model 1440 updates the state of the ViewModel(L) 1490 and View(L) 1491 so that the ViewModel(L) 1490 contains the needed data to present the GUI's components in the correct location on the system's display and so that the ViewModel(L) 1490 can correctly interpret input from the IPU 1420 in the context of the current state of the Model 1440; and so that the View(L)

1491 can correctly generate graphics for display in the context of the current state of the Model 1440.

[0075] By way of example, consider a touch-sensitive application with a button that among its functions responds to a user's touch by changing its appearance indicating that it has been activated. When the application is run, the application reads the location, size, and details of the appearance of the button from memory and compiled application code. The View(H) 1480 code generates the necessary graphics which are presented to the user to display this button. The Model 1440 updates the state of the ViewModel(L) 1490 to record that this graphical element is a button, and that it should change appearances from a "normal" appearance to a "pressed" appearance when touched. The Model 1440 also updates the state of the View(L) 1491 to record the correct appearance for the "normal" and "pressed" states in the ViewModel(L) 1490. This appearance may be a description of low-fidelity graphical elements, or a complete raster to display. In this example, the "pressed" state is represented by a displaying a white box around the button's position.

[0076] A User touches the touch-screen display, and input data describing that touch is received less than 1ms later by the IPU 1420. The IPU 1420 creates an input event representing a touch-down event from the input data and sends this input event to the application Controller 1430. The Controller 1430 manipulates the Model 1440. In this case, the Controller 1430 is indicating to the Model 1440 that the button has been touched and that the application should perform whatever commands are associated with this button. At the same time that the IPU 1420 sends an event to the Controller 1430, it sends an event to the ViewModel(L) 1490 indicating that the button has been touched. The ViewModel(L) 1490 was previously instructed by the Model 1440 as to what to do in the case of a touch, and in this case it responds to the touch event by changing its state to "pressed". The View(L) 1491 responds to this change by displaying a white box around the button, feedback that corresponds to its "pressed" appearance. The change to the Model 1440 that the button is touched causes an update of View(H) 1480, so that it too reflects that button is now touched. The User, who see the output of both View(H) 1480 and View(L) 1491, sees the immediate feedback of their touch by View(L) 1491 followed a fraction of a second later by the feedback from View(H) 1480.

Reducing Control Response Latency

[0077] In typical interactive computer systems which employ a graphical user interface (GUI), the primitives of interaction may be broken down into controls, which may be collected into libraries. The task of application design is often one of assembling collections of known controls to facilitate the desired user interaction. For example, if the developer desires for the user to input their address, he or she might utilize a text field control to enable free-form input of street number and name, a drop-down list of states to force the user to choose from among valid options, and another freeform text field to enable zip-code entry. A button might also be used to allow the user to indicate that they have finished the task of inputting the data. Many systems include libraries which include text fields, drop-down lists, and buttons, among many other types of controls.

[0078] Once the application developer has selected the controls to compose the desired interaction flow, they will typically subsequently design behaviors for those controls. Such behaviors might be specified through the selection of parameters, such as specifying, at design-time, the contents of a drop-down list, or setting a flag to disable a control. Such behaviors might also be specified through computer code. In some embodiments, such code might be added to ‘event handlers’ or ‘callbacks’, which are triggered by the system’s event distribution system when the desired circumstances have been met. For example, an “onClick” callback might be called whenever a control is ‘clicked’. At that time, the code specified by the application developer to control the behavior of the control will execute.

[0079] In some embodiments, the application software might be running on a system where efficiencies are achieved by processing control parameters before the execution of callbacks. For example, setting the ‘enabled’ flag of a control to ‘false’ at design time would typically result in the control being disabled when it is first displayed on screen, without first setting the flag to its default value (such as ‘true’), only to later be reversed in code.

[0080] One method for accelerating input is to implement common controls in hardware, as described above and in U.S. Patent Application No. 14/046,819. In such systems, the parameterization of behaviors, as has been described herein, may result in substantially improved performance over behaviors specified in software code, since the parameterization may be heeded through actions in hardware which are typically faster than those implemented in software.

[0081] While such an approach may yield significantly better performance, it may, in some embodiments, omit interactions which span across controls. For example, consider the window shown in Figure 15. In this figure, we see a dialog message, and a button labeled “OK”. When the user presses the OK button, if accelerated in the method described above and in U.S. Patent Application No. 14/046,819, the button itself may provide an immediate visual indication of response to the user’s touch. The developer of the application shown in Figure 15 may intend this “OK” button to dismiss the window that contains it. In such a case, in some object-oriented implementations, the developer may write source code, in an event handler that will respond to the user’s input by dismissing the window. In some embodiments of the invention described above and in U.S. Patent Application No. 14/046,819, this source code would be executed by the CPU, and thus operate more slowly than the visual response provided by the IPU for the button press itself.

[0082] The presently disclosed system and method for reducing control response latency provide a means by which such cross-control behaviors may also be accelerated. In an embodiment, the logic for a design pattern such as that described above (closing the parent window when the button is clicked) is encoded in hardware, e.g., hardware associated with the input processing unit (IPU) discussed above. In some embodiments, in order to specify that the design pattern should be utilized, the system is configured such that the developer can set a flag on the button control (for example, one named “closeParentWindowWhenClicked”), that would instruct the IPU to use hardware to immediately close the parent window when the user clicks the button, instead of running developer-supplied code which provides the instruction to close the window. In other embodiments, the IPU is implemented as software which runs on a secondary processor, such as the GPU, SOC, micro controller, or a core of the CPU separate from that executing the application software. In other embodiments, the IPU is implemented as software on the same processor/core as other elements, and given high priority by the OS or other resource manager. See FIG. 15. A particular cross-control behavior to be processed by the IPU can be identified at layout time or in response to user input.

[0083] In an embodiment, the presently disclosed system and method for reducing control response latency provides an input subsystem including an input device, the subsystem adapted to output signals in response to user inputs to the input device, the signals comprising information concerning the input. A low-latency subsystem is provided to receive the signals and to generate a low-latency response in response to one or more of the signals, and to send at

least one of the signals to a high-latency subsystem. The signals received by the low-latency subsystem comprise an input which triggers a cross-control behavior. The low-latency subsystem makes an identification of that cross-control behavior, and the low-latency subsystem generates a low-latency response reflecting the cross-control behavior. The high-latency subsystem is adapted to receive at least one of the signals, and to generate a high-latency response in response to one or more of the at least one of the signals. The low-latency subsystem generates the low-latency response with low latency relative to the high-latency subsystem and the high-latency subsystem generates the high-latency response with high latency relative to the low-latency subsystem.

[0084] In an embodiment, the presently disclosed system and method for reducing control response latency provides an input subsystem including an input device, the subsystem adapted to output signals in response to user inputs to the input device, the signals comprising information concerning the input. A low-latency subsystem is provided to receive the signals and to generate a low-latency response in response to one or more of the signals, and to send at least one of the signals to a high-latency subsystem. The signals received by the low-latency subsystem comprise an identification of a cross-control behavior and the generated low-latency response comprises data reflecting the cross-control behavior. The high-latency subsystem is adapted to receive at least one of the signals, and to generate a high-latency response in response to one or more of the signals. The low-latency subsystem generates the low-latency response with low latency relative to the high-latency subsystem and the high-latency subsystem generates the high-latency response with high latency relative to the low-latency subsystem.

Details of Embodiments

Tab Layout

[0085] One commonly used method for organizing large volumes of information and/or controls is to arrange them into a small number of categories. Such categories may then be arranged behind a series of labeled tabs in the interface. In some environments, the user selects a category by touching its corresponding “tab” at the top of the control. The elements associated with that tab label are then displayed.

[0086] A visual indicator such as a colored horizontal bar may be used to highlight the currently selected tab. In some cases, the interface may be designed to appear as if the categories

of content themselves, and not merely the labels, are arrayed next to each other, with each category sized to the width of the display; thus only one full category can be displayed at a time. When changing tabs, the content may perform an animated transition such as a horizontal slide wherein the previously displayed content slides out of view, to be replaced by its immediate neighbor, and if that is not the content of the newly selected tab, further horizontal movement occurs until the desired content is fully in view. In some embodiments, the user may be able to swipe the content to move between tabs, with the content following a drag sideways or animating the transition to the next category upon the performance of a gesture.

[0087] In some embodiments of the presently disclosed system and method for reducing control response latency, custom-designed hardware is employed to render the tab content as a horizontally-arrayed series of bitmaps, while responding to input to change which tab is displayed. At the first CPU rendering of a tab-based layout, the CPU could pass rendered bitmaps of components of the interface to the hardware controller, alongside information concerning their positions within the interface, as well as an indication as to which bitmaps are selectable tab labels and which are the tab's contents.

Buttons Closing Windows

[0088] With reference to FIG. 15, the developer of the application shown in the figure may intend this "OK" button to dismiss the window that contains it. The presently disclosed system and method for reducing control response latency provides a means by which such cross-control behaviors may be accelerated. In an embodiment, the logic for closing the parent window when the button is clicked might be encoded in the IPU. In some embodiments of the invention, in order to specify that the design pattern should be utilized, the developer would set a flag on the button control (perhaps named "closeParentWindowWhenClicked"), that instructs the IPU to execute code that immediately closes the parent window when the user clicks the button, instead of running developer-supplied code which provides the instruction to close the window.

Disabling Buttons

[0089] Another common design pattern is to immediately disable a button when it is clicked, to prevent a user from sending multiple requests. Again, in some object-oriented systems, the developer provides code in the event handler for the "click" event of the button that

instructs the system to disable it upon clicking. In some systems, disabling the button changes its visual appearance, such as “greying out” to make it clear to the user that this button will not respond to subsequent inputs. In some embodiments of the presently disclosed system and method for reducing control response latency, a method is provided by which the developer specifies to the system that, when the button is clicked, it should be disabled, other than by providing source code to be executed in the CPU. In an embodiment, a flag for the button object is provided, perhaps named “disableOnClick”, that, when set to true, instructs the system to immediately disable the button when the user clicks it, rather than waiting to execute source code instructions provided by the developer. In some embodiments, after disabling the button, the system would then process the “click” event handler, executing instructions provided by the developer in source code.

Buttons Disabling Parts of the Interface

[0090] In application design, a common technique is to design an interface comprised of form controls such that one of the controls - typically a checkbox control - is used to enable or disable a set of other controls. When an unchecked checkbox is clicked, it becomes selected and the set of other controls become enabled, allowing the user to click on them or change their values. When the checkbox is clicked a second time, the same set of controls becomes disabled, preventing any changes, and, in some cases, the disabling is reflected through a visual change (typically a greying of the controls’ text).

[0091] One embodiment of the presently disclosed system and method for reducing control response latency would offer a technique for tying the content to be disabled with the control whose value will trigger the disabling. This would allow the IPU to easily identify and adjust the affected controls on processing of a click of the triggering control. By using this technique, the IPU is able to handle the processing and thereby improve the system’s responsiveness to input.

Button Changing Appearance

[0092] When a button performs an action on click, it is sometimes advantageous to provide visual feedback that that action is underway. This is especially true if the action takes several seconds to complete, such as actions that require a network connection. Examples include “refresh” buttons, which sometimes comprise an icon with two circulating arrows. On click,

these buttons sometimes change into circular loading animation icons, an indicator that the refresh routine is underway.

[0093] In some embodiments of the presently disclosed system and method for reducing control response latency, a method is provided by which the developer can specify to the IPU that, when these types of buttons are clicked, they should change to another type of icon, which would also be open for definition. In some embodiments, the CPU would then be responsible for restoring the original icon when the action completed.

Menu or Window Opening Buttons

[0094] Menus are a common technique for organizing buttons with a variety of actions under one heading. See, for example, the application menu shown in Figure 16. In many instances, the menu is made visible or hidden by clicking the menu's label button. In some variations, an animation is initiated to transition the menu from visible to hidden. The menu items displayed in the menu itself are typically arrayed vertically below the menu button.

[0095] In an embodiment of the presently disclosed system and method for reducing control response latency, a mechanism is made available to application developers to allow the menu and its contents, appearance and behavior to be interpreted by the IPU for display and handling of click events on the menu button. This allows tasks typically performed by the CPU to be handled by the IPU, improving performance and interface responsiveness. FIG. 16 shows buttons launching new windows and applications.

[0096] Applications are regularly divided into multiple views. When a user desires to transition from one view to another, he or she usually achieves this by clicking on a button with the functionality to accomplish that task. In some instances, in a similar manner, these buttons may launch other applications. Other buttons launch modal windows that only occupy a portion of the screen. An animation sometimes accompanies the transition to these new views, creating a visual and metaphorical link between the current and next views.

[0097] The presently disclosed system and method for reducing control response latency provides, in an embodiment, a means by which an IPU can recognize these buttons and their functionality, with the goal of transitioning views to a cached bitmap render of the desired new view, if one exists.

Long-Press Contextual Interface

[0100] Some interfaces utilize a long press to trigger the display of contextual interface elements, such as a menu or special menu bar. In some embodiments of the presently disclosed system and method, a parameter is used to specify that a given button or other control should initiate the display of such a menu or menu bar. Another parameter might specify which UI object comprises the menu, and should thus be displayed.

Tab-Stop Advancement

[0101] In some embodiments, a parameter is provided to specify that, when the user presses or makes another gesture on a ‘tab’ or other indicator, the in-focus control should be advanced.

Text Entry

[0102] In some embodiments, a parameter might be included to specify that, when the keyboard is typed-on, or other text entry mechanism utilized, text should be inserted at the text insertion (cursor) point. Conversely, when a text field is selected, the IPU should cause an onscreen keyboard to appear.

Buttons Associated with Opening and Closing Navigation “Drawers”

[0103] A common interface design groups and subsequently hides many navigation items within a sliding “drawer” menu. The drawer has two states, expanded and collapsed. A fixed button is typically located elsewhere in the interface, toggling the drawer’s state upon press.

[0104] In one embodiment of the presently disclosed system and method, a means of specifying that a button has this behavior on a drawer is provided, as well as a way of indicating which control corresponds to the slide out drawer. This would allow the IPU to open and close the drawer in response to button presses (in some embodiments animating the transition between different drawer states).

Text Associated with Slider, Progress and Seek Bars

[0105] Many applications use graphical representations in the form of horizontal bars to indicate the fraction of a task that has been completed, or as an interactive component for inputting a continuously varying parameter. This includes “seek” bars, for example, which allow

a user to move advance or rewind video and audio playback. Many of these kinds of controls have an associated text label that provides a user with feedback on the position of the bar's current value. This information must be synced with the control itself in order to stay up-to-date.

[0106] In one embodiment, a method is provided for indicating the range of values that a bar represents, allowing the IPU to render the text label with a highly responsive, properly updated value.

Synchronized Responses

[0107] Many systems include elements whose movement or behavior it might be desirable to synchronize. For example, one element might be 'draggable', and, when dragged, may show additional elements dragged alongside. For example, in video interfaces, a slider may be shown which allows the user to choose a position in the video by dragging the thumb to a different position. While dragging, a small window may be displayed above the user's finger, showing content from the video. Two aspects of synchronization would be useful: first, ensuring that the window's movement alongside the dragged thumb is instantaneous and linked to ensure consistent spacing relative to the thumb. Second, the video content shown in the window should correspond to the physical location of the thumb within the timeline: thus, the 'frame' property should be synchronized to the 'position' property of the thumb. In some embodiments of the present invention, methods would be provided (such as flags, or fields where the developer can select or specify synchronization targets/sources) to create ad hoc arbitrary synchronizations. These synchronizations would be used by the IPU to ensure a response is synchronized across controls.

Buttons Revealing Search Interface

[0108] Many applications include search functionality. This is often triggered by a button labeled search or indicated with a magnifying glass. When the button is pressed, a search textbox appears in the interface, in many cases with a keyboard accompanying it for text entry.

[0109] In an embodiment of the presently disclosed system and method, a flag would be available for developers to indicate which buttons triggered the display of these search boxes, allowing the IPU to display an accelerated presentation of the search interface on touch of a flagged button.

[0110] Throughout the text of this application, the word “event” is used to describe information describing attributes of user input. This term is used generally, and thus includes embodiments in which event driven architectures are employed (with actual event objects being passed between software elements), as well as more basic input streams in which the “event” being described is simply present in the stream of information. Such events may be, e.g., non-object-orient types of events or object-oriented types events.

[0111] The present system and methods are described above with reference to block diagrams and operational illustrations of methods and devices comprising a computer system capable of receiving and responding to user input. It is understood that each block of the block diagrams or operational illustrations, and combinations of blocks in the block diagrams or operational illustrations, may be implemented by means of analog or digital hardware and computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, ASIC, or other programmable data processing apparatus, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, implements the functions/acts specified in the block diagrams or operational block or blocks. In some alternate implementations, the functions/acts noted in the blocks may occur out of the order noted in the operational illustrations. For example, two blocks shown in succession may in fact be executed substantially concurrently or the blocks may sometimes be executed in the reverse order, depending upon the functionality/acts involved.

[0112] While the invention has been particularly shown and described with reference to a preferred embodiment thereof, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the spirit and scope of the invention.

What is claimed is:

1. A system for processing user input, comprising:

display comprising a visualization of a user interface;

user input device comprising at least one selected from the group consisting of touch input and pen input, the user input device adapted to output signals corresponding to user input;

input processing system operatively connected to the user input device, the input processing system adapted to receive the signals corresponding to user input and to emit signals corresponding to the received signals;

first graphical processing and output system comprising software running in a software stack, the first graphical processing and output system adapted to receive at least some of the emitted signals and to generate high-latency data that can affect the visualization of the user interface on the display and to output the high-latency data for use by the display in response to the at least some of the emitted signals;

second graphical processing and output system implemented at least partially in hardware, the second graphical processing and output system adapted to substantially simultaneously receive at least some of the emitted signals and to generate low-latency data that can affect the visualization of the user interface on the display and to output the generated low-latency data for use by the display in response to the at least some of the emitted signals;

wherein the at least some of the emitted signals received by the second graphical processing and output system comprise an identification of a cross-control behavior, the second graphical processing and output system comprises hardware configured to implement the cross-control behavior, and the generated low-latency data comprises data reflecting the cross-control behavior;

wherein, with respect to the response to the at least some of the emitted signals, the second graphical processing and output system is adapted to output the generated low-latency

data with low latency relative to the output of the high-latency data by the first graphical processing and output system; and

the display being configured to display at least some of the low-latency data and at least some of the high-latency data on the visualization of the user interface.

2. The system for processing user input according to claim 1, wherein the hardware configured to implement the cross-control behavior comprises hardware having logic for implementing the cross-control behavior encoded therein.
3. The system for processing user input according to claim 1, wherein the cross-control behavior comprises interaction between a first control and a second control.
4. The system for processing user input according to claim 3, wherein the first control comprises a button and the second control comprises a window.
5. The system for processing user input according to claim 1, wherein the hardware configured to implement the cross-control behavior comprises hardware configured to render tab content.
6. The system for processing user input according to claim 5, wherein the hardware configured to render tab content is configured to render the tab content as a horizontally-arrayed series of bitmaps and respond to input to change which tab is displayed.
7. The system for processing user input according to claim 1, configured such that the first graphical processing and output system passes rendered bitmaps of components of the user interface to the second graphical processing and output system.
8. The system for processing user input according to claim 1, wherein the input processing system is configured to respond to a flag associated with a control by executing code that implements the cross-control behavior upon receipt of user input.
9. The system for processing user input according to claim 1, wherein the cross-control behavior comprises disabling of a control.
10. The system for processing user input according to claim 1, wherein the cross-control behavior comprises disabling of a part of the user interface upon activation of a control.

11. The system for processing user input according to claim 10, wherein the control is a checkbox.
12. The system for processing user input according to claim 1, wherein the cross-control behavior comprises a change in appearance of a button.
13. The system for processing user input according to claim 1, wherein the cross-control behavior comprises a change in content, appearance or behavior of a menu.
14. The system for processing user input according to claim 1, wherein the cross-control behavior comprises display of at least one contextual interface element.
15. The system for processing user input according to claim 14, wherein the at least one contextual interface element comprises a menu or menu bar.
16. The system for processing user input according to claim 14, wherein the display of at least one contextual interface element is in response to a long-press.
17. The system for processing user input according to claim 1, wherein the cross-control behavior comprises advancement of an in-focus control.
18. The system for processing user input according to claim 1, wherein the cross-control behavior comprises insertion of text.
19. The system for processing user input according to claim 1, wherein the cross-control behavior comprises opening or closing a navigation drawer.
20. The system for processing user input according to claim 1, wherein the cross-control behavior comprises presenting a search interface.
21. A system for processing user input, comprising:

display comprising a visualization of a user interface;

user input device adapted to output signals corresponding to user input;

input processing system operatively connected to the user input device, the input processing system adapted to receive the signals corresponding to user input and to emit a stream of signals corresponding to at least some of the received signals;

first graphical processing and output system adapted to receive at least a portion of the stream of signals and to generate high-latency data that can affect the visualization of the user interface on the display and to output the high-latency data for use by the display in response to at least some of the signals in the at least a portion of the stream of signals received by the first graphical processing and output system;

second graphical processing and output system adapted to receive at least a portion of the stream of signals and to generate low-latency data that can affect the visualization of the user interface on the display and to output the generated low-latency data for use by the display in response to the at least some of the signals in the at least a portion of the stream of signals received by the second graphical processing and output system, the second graphical processing and output system being configured to output the generated low-latency data with low latency relative to the output of the high-latency data by the first output graphical processing and output system;

wherein the at least a portion of the stream of signals received by the second graphical processing and output system comprise an identification of a cross-control behavior and the generated low-latency data comprises data reflecting the cross-control behavior;

and

the display being configured to merge at least some of the low-latency data and at least some of the high-latency data onto the visualization of the user interface.

22. The system for processing user input according to claim 21, wherein the cross-control behavior comprises interaction between a first control and a second control.
23. The system for processing user input according to claim 22, wherein the first control comprises a button and the second control comprises a window.
24. The system for processing user input according to claim 21, wherein the cross-control behavior comprises display of tab content.

25. The system for processing user input according to claim 24, wherein the second graphical processing and output system is configured to render the tab content as a horizontally-arrayed series of bitmaps and respond to input to change which tab is displayed.
26. The system for processing user input according to claim 21, configured such that the first graphical processing and output system passes rendered bitmaps of components of the user interface to the second graphical processing and output system.
27. The system for processing user input according to claim 21, wherein the input processing system is configured to respond to a flag associated with a control by executing code that implements the cross-control behavior upon receipt of user input.
28. The system for processing user input according to claim 21, wherein the cross-control behavior comprises disabling of a control.
29. The system for processing user input according to claim 21, wherein the cross-control behavior comprises disabling of a part of the user interface upon activation of a control.
30. The system for processing user input according to claim 29, wherein the control is a checkbox.
31. The system for processing user input according to claim 21, wherein the cross-control behavior comprises a change in appearance of a button.
32. The system for processing user input according to claim 21, wherein the cross-control behavior comprises a change in content, appearance or behavior of a menu.
33. The system for processing user input according to claim 21, wherein the cross-control behavior comprises display of at least one contextual interface element.
34. The system for processing user input according to claim 33, wherein the at least one contextual interface element comprises a menu or menu bar.
35. The system for processing user input according to claim 33, wherein the display of at least one contextual interface element is in response to a long-press.

36. The system for processing user input according to claim 21, wherein the cross-control behavior comprises advancement of an in-focus control.
37. The system for processing user input according to claim 21, wherein the cross-control behavior comprises insertion of text.
38. The system for processing user input according to claim 21, wherein the cross-control behavior comprises opening or closing a navigation drawer.
39. The system for processing user input according to claim 21, wherein the cross-control behavior comprises presenting a search interface.
40. A system for processing user input, comprising:

input subsystem including an input device, the subsystem adapted to output signals in response to a plurality of user inputs to the input device, the signals comprising information concerning the input;

low-latency subsystem adapted to receive the signals and to generate a low-latency response in response to one or more of the signals, and to send at least one of the signals to a high-latency subsystem;

wherein the signals received by the low-latency subsystem comprise an identification of a cross-control behavior and the generated low-latency response comprises data reflecting the cross-control behavior;

the high-latency subsystem adapted to receive the at least one of the signals, and to generate a high-latency response in response to one or more of the at least one of the signals; and,

wherein the low-latency subsystem generates the low-latency response with low latency relative to the high-latency subsystem and the high-latency subsystem generates the high-latency response with high latency relative to the low-latency subsystem.

41. A system for processing user input, comprising:

input subsystem including an input device, the subsystem adapted to output signals in response to a plurality of user inputs to the input device, the signals comprising information concerning the input;

low-latency subsystem adapted to receive the signals and to generate a low-latency response in response to one or more of the signals, and to send at least one of the signals to a high-latency subsystem;

wherein the signals received by the low-latency subsystem comprise an input which triggers a cross-control behavior, the low-latency subsystem makes an identification of the cross-control behavior, and the low-latency subsystem generates a low-latency response reflecting the cross-control behavior;

the high-latency subsystem adapted to receive the at least one of the signals, and to generate a high-latency response in response to one or more of the at least one of the signals;

wherein the low-latency subsystem generates the low-latency response with low latency relative to the high-latency subsystem and the high-latency subsystem generates the high-latency response with high latency relative to the low-latency subsystem.

42. The system for processing user input according to claim 41, wherein the cross-control behavior comprises interaction between a first control and a second control.

43. The system for processing user input according to claim 42, wherein the first control comprises a button and the second control comprises a window.

44. The system for processing user input according to claim 41, wherein the cross-control behavior comprises display of tab content.

45. The system for processing user input according to claim 44, wherein the low-latency subsystem is configured to render the tab content as a horizontally-arrayed series of bitmaps and respond to input to change which tab is displayed.

46. The system for processing user input according to claim 41, configured such that the high-latency subsystem passes rendered bitmaps of components of the user interface to the low-latency subsystem.
47. The system for processing user input according to claim 41, wherein the input subsystem is configured to respond to a flag associated with a control by executing code that implements the cross-control behavior upon receipt of user input.
48. The system for processing user input according to claim 41, wherein the cross-control behavior comprises disabling of a control.
49. The system for processing user input according to claim 41, wherein the cross-control behavior comprises disabling of a part of the user interface upon activation of a control.
50. The system for processing user input according to claim 49, wherein the control is a checkbox.
51. The system for processing user input according to claim 41, wherein the cross-control behavior comprises a change in appearance of a button.
52. The system for processing user input according to claim 41, wherein the cross-control behavior comprises a change in content, appearance or behavior of a menu.
53. The system for processing user input according to claim 41, wherein the cross-control behavior comprises display of at least one contextual interface element.
54. The system for processing user input according to claim 53, wherein the at least one contextual interface element comprises a menu or menu bar.
55. The system for processing user input according to claim 53, wherein the display of at least one contextual interface element is in response to a long-press.
56. The system for processing user input according to claim 41, wherein the cross-control behavior comprises advancement of an in-focus control.
57. The system for processing user input according to claim 41, wherein the cross-control behavior comprises insertion of text.

58. The system for processing user input according to claim 41, wherein the cross-control behavior comprises opening or closing a navigation drawer.
59. The system for processing user input according to claim 41, wherein the cross-control behavior comprises presenting a search interface.
60. A system for processing user input with low latency, comprising:
- an input device;
 - an input processing unit;
 - a high-latency subsystem;
 - a low-latency subsystem;
 - input processing unit software adapted to generate signals in response to user inputs; and,
 - output device;
- wherein the low-latency subsystem processes a plurality of the signals with low latency relative to the high-latency subsystem and the high-latency subsystem processes a plurality of the signals with high latency relative to the low-latency subsystem;
- wherein the input device, the input processing unit, the high-latency subsystem, the low-latency processing subsystem, the input processing unit software and the output device being arranged such that:
- at least some of the signals are processed in parallel by said low-latency subsystem and said high-latency subsystem, the low-latency subsystem creating a programmable response for output on the output device with low-latency and the high-latency subsystem creating a response for output on the output device with high-latency;
- wherein the plurality of the signals processed by the low-latency subsystem comprise an input from the user, in response to which the low-latency subsystem makes an identification of a cross-control behavior, and the programmable response comprises data reflecting the cross-control behavior; and

the output device outputting one or both of the response from the low-latency subsystem and the response from the high-latency subsystem when the responses are output thereto.

61. The system for processing user input according to claim 60, wherein the cross-control behavior comprises interaction between a first control and a second control.
62. The system for processing user input according to claim 61, wherein the first control comprises a button and the second control comprises a window.
63. The system for processing user input according to claim 60, wherein the cross-control behavior comprises display of tab content.
64. The system for processing user input according to claim 63, wherein the low-latency subsystem is configured to render the tab content as a horizontally-arrayed series of bitmaps and respond to input to change which tab is displayed.
65. The system for processing user input according to claim 60, configured such that the high-latency subsystem passes rendered bitmaps of components of the user interface to the low-latency subsystem.
66. The system for processing user input according to claim 60, wherein the input processing unit is configured to respond to a flag associated with a control by executing code that implements the cross-control behavior upon receipt of user input.
67. The system for processing user input according to claim 60, wherein the cross-control behavior comprises disabling of a control.
68. The system for processing user input according to claim 60, wherein the cross-control behavior comprises disabling of a part of the user interface upon activation of a control.
69. The system for processing user input according to claim 68, wherein the control is a checkbox.
70. The system for processing user input according to claim 60, wherein the cross-control behavior comprises a change in appearance of a button.
71. The system for processing user input according to claim 60, wherein the cross-control behavior comprises a change in content, appearance or behavior of a menu.

72. The system for processing user input according to claim 60, wherein the cross-control behavior comprises display of at least one contextual interface element.
73. The system for processing user input according to claim 72, wherein the at least one contextual interface element comprises a menu or menu bar.
74. The system for processing user input according to claim 72, wherein the display of at least one contextual interface element is in response to a long-press.
75. The system for processing user input according to claim 60, wherein the cross-control behavior comprises advancement of an in-focus control.
76. The system for processing user input according to claim 60, wherein the cross-control behavior comprises insertion of text.
77. The system for processing user input according to claim 60, wherein the cross-control behavior comprises opening or closing a navigation drawer.
78. The system for processing user input according to claim 60, wherein the cross-control behavior comprises presenting a search interface.
79. A system for processing user input with low latency, comprising:
- an input device;
 - output device;
 - an input processing unit;
 - input processing unit software adapted to generate signals in response to user inputs;
 - a high-latency response subsystem adapted to create a high-latency response to at least some of the signals;
 - a low-latency response subsystem implemented at least partially in hardware, the low-latency response subsystem adapted to create a low-latency response to at least some of the signals, the low-latency response subsystem comprising one or more parameters affecting the low-latency response;

wherein the at least some of the signals comprise an identification of a cross-control behavior, the low-latency response subsystem comprises hardware configured to implement the cross-control behavior, and the low-latency response comprises data reflecting the cross-control behavior;

wherein the low-latency response is not of higher latency than the high-latency response, and the high-latency response is not of lower latency than the low-latency response with respect to common signals;

wherein the input device, output device, the input processing unit, input processing unit software, the high-latency response subsystem and the low-latency response subsystem are arranged such that:

the output device outputs the low-latency response and the high-latency response; and

wherein application software can alter the one or more parameters affecting the low-latency response.

80. The system for processing user input according to claim 79, wherein the cross-control behavior comprises interaction between a first control and a second control.
81. The system for processing user input according to claim 80, wherein the first control comprises a button and the second control comprises a window.
82. The system for processing user input according to claim 79, wherein the cross-control behavior comprises display of tab content.
83. The system for processing user input according to claim 82, wherein the low-latency response subsystem is configured to render the tab content as a horizontally-arrayed series of bitmaps and respond to input to change which tab is displayed.
84. The system for processing user input according to claim 79, configured such that the high-latency response subsystem passes rendered bitmaps of components of the user interface to the second graphical processing and output system.

85. The system for processing user input according to claim 79, wherein the input processing unit is configured to respond to a flag associated with a control by executing code that implements the cross-control behavior upon receipt of user input.
86. The system for processing user input according to claim 79, wherein the cross-control behavior comprises disabling of a control.
87. The system for processing user input according to claim 79, wherein the cross-control behavior comprises disabling of a part of the user interface upon activation of a control.
88. The system for processing user input according to claim 79, wherein the control is a checkbox.
89. The system for processing user input according to claim 79, wherein the cross-control behavior comprises a change in appearance of a button.
90. The system for processing user input according to claim 79, wherein the cross-control behavior comprises a change in content, appearance or behavior of a menu.
91. The system for processing user input according to claim 79, wherein the cross-control behavior comprises display of at least one contextual interface element.
92. The system for processing user input according to claim 79, wherein the at least one contextual interface element comprises a menu or menu bar.
93. The system for processing user input according to claim 91, wherein the display of at least one contextual interface element is in response to a long-press.
94. The system for processing user input according to claim 79, wherein the cross-control behavior comprises advancement of an in-focus control.
95. The system for processing user input according to claim 79, wherein the cross-control behavior comprises insertion of text.
96. The system for processing user input according to claim 79, wherein the cross-control behavior comprises opening or closing a navigation drawer.

97. The system for processing user input according to claim 79, wherein the cross-control behavior comprises presenting a search interface.
98. The system for processing user input according to claim 79, wherein the cross-control behavior comprises a synchronized response between at least two interface elements.