

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
29 April 2010 (29.04.2010)

(10) International Publication Number
WO 2010/048051 A2

(51) International Patent Classification:
G06F 9/455 (2006.01) *G06F 3/048* (2006.01)
G06F 3/041 (2006.01)

PRADHAN, Samir S.; c/o Microsoft Cororation, One Microsoft Way, Redmond, Washington 98052-6399 (US). TEED, Jennifer A.; c/o Microsoft Cororation, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(21) International Application Number:
PCT/US2009/060977

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(22) International Filing Date:
16 October 2009 (16.10.2009)

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

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/258,439 26 October 2008 (26.10.2008) US

(71) Applicant (for all designated States except US): MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: TOWNSEND, Reed L.; c/o Microsoft Cororation, One Microsoft Way, Redmond, Washington 98052-6399 (US). TU, Xiao; c/o Microsoft Cororation, One Microsoft Way, Redmond, Washington 98052-6399 (US). SCOTT, Bryan D.; c/o Microsoft Cororation, One Microsoft Way, Redmond, Washington 98052-6399 (US). TORSET, Todd A.; c/o Microsoft Cororation, One Microsoft Way, Redmond, Washington 98052-6399 (US). SYKES, Kenneth W.; c/o Microsoft Cororation, One Microsoft Way, Redmond, Washington 98052-6399 (US).

Declarations under Rule 4.17:

[Continued on next page]

(54) Title: MULTI-TOUCH OBJECT INERTIA SIMULATION




Figure 3

(57) Abstract: The inertia system provides a common platform and application-programming interface (API) for applications to extend the input received from various multi-touch hardware devices to simulate real-world behavior of application objects. To move naturally, application objects should exhibit physical characteristics such as elasticity and deceleration. When a user lifts all contacts from an object, the inertia system provides additional manipulation events to the application so that the application can handle the events as if the user was still moving the object with touch. The inertia system generates the events based on a simulation of the behavior of the objects. If the user moves an object into another object, the inertia system simulates the boundary characteristics of the objects. Thus, the inertia system provides more realistic movement for application objects manipulated using multi-touch hardware and the API provides a consistent feel to manipulations across applications.



- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

MULTI-TOUCH OBJECT INERTIA SIMULATION

BACKGROUND

[0001] A tablet PC, or pen computer, is a notebook or slate-shaped mobile computer, equipped with a touch screen or graphics tablet/screen hybrid

5 technology that allows the user to operate the computer with a stylus, digital pen, or fingertip instead of a keyboard or mouse. Tablet PCs offer a more natural form of input, as sketching and handwriting are a much more familiar form of input than a keyboard and mouse, especially for people who are new to computers. Tablet PCs can also be more accessible because those who are physically unable to
10 type can utilize the additional features of a tablet PC to be able to interact with the electronic world.

[0002] Multi-touch (or multitouch) denotes a set of interaction techniques that allow computer users to control graphical applications using multiple fingers or input devices (e.g., stylus). Multi-touch implementations usually include touch

15 hardware (e.g., a screen, table, wall, and so on) and software that recognizes multiple simultaneous touch points. Multi-touch stands in contrast to traditional touch screens (e.g., computer touchpad, ATM, shopping kiosk) that only recognize one touch point at a time. Multi-touch hardware can sense touches using heat, finger pressure, high capture rate cameras, infrared light, optic
20 capture, tuned electromagnetic induction, ultrasonic receivers, transducer microphones, laser rangefinders, shadow capture, and other mechanisms. Many applications for multi-touch interfaces exist and application designers and users are proposing even more. Some uses are individualistic (e.g., Microsoft Surface, Apple iPhone, HTC Diamond). As a new input method, multi-touch offers the
25 potential for new user experience paradigms.

[0003] An application cannot use multi-touch hardware without an interface for the application software to receive information from the multi-touch hardware.

Unfortunately, each multi-touch hardware device includes its own proprietary interface and application authors must have specific knowledge of a hardware
30 device to write software that works with the device. For example, a multi-touch hardware provider may provide a kernel-mode driver and a user-mode application interface through which user-mode software applications can communicate with the multi-touch hardware to receive touch information. An application author

writes software that communicates with the user-mode application interface, but the application author's software works only with that multi-touch hardware. A computer user with a different multi-touch hardware device cannot use the application author's software unless the application author produces a different

5 version of the software that operates correctly with the computer user's device.

This produces a very limited potential market for application authors, reduces the incentive to write applications supporting multi-touch interactions, and keeps the cost of the most popular devices high for which the greatest number of applications is available.

10 **[0004]** Another problem is the difficulty for applications to determine a user's intentions based on touch input received from multi-touch hardware. Touch input may be received as a list of coordinates where the hardware senses touch input at any given time. Each application has to include software to interpret the coordinates and determine the user's intention. In addition, the user's intention
15 may extend beyond the actual touch input received. The user may expect virtual objects to behave how they do in the physical world. For example, a user may expect to be able to "toss" a file from one side of the desktop to another by flicking his/her finger. This type of movement is not supported by current multi-touch applications, which would expect the user to drag his/her finger from one side of
20 the screen all the way to the other. Even if an application provides support for this type of movement, other applications could not benefit from it and thus application authors would have to repeat the work of the first application author to offer the same functionality in their applications.

SUMMARY

25 **[0005]** The inertia system provides a common platform and application-programming interface (API) for applications to extend the input received from various multi-touch hardware devices to simulate real-world behavior of objects. The manipulations received by the application only describe the movement of an object based on the movement of contacts with the multi-touch hardware.
30 However, to move naturally, objects should also exhibit physical characteristics such as elasticity and deceleration. When a user lifts all contacts from an object, the inertia system provides additional manipulation events to the application so that the application can handle the events as if the user was still moving the object with touch. However, the inertia system actually generates the events based on a

simulation of the behavior of the objects. If the user moves an object into another object, the inertia system sends manipulation events based on the boundary characteristics of the objects. Thus, the inertia system provides more realistic movement for application objects that a user manipulates using multi-touch

5 hardware and the API provides a consistent feel to manipulations across applications.

[0006] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the
10 claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Figure 1 is a block diagram that illustrates components of the inertia system, in one embodiment.

15 **[0008]** Figure 2 is a data flow diagram that illustrates a typical operating environment of the inertia system and the flow of data between components, in one embodiment.

[0009] Figure 3 is a display diagram that illustrates an application object manipulated by user touch, in one embodiment.

20 **[0010]** Figure 4 is a flow diagram that illustrates the input loop processing of a multi-touch application using the inertia system to handle manipulation events, in one embodiment.

[0011] Figure 5 is a flow diagram that illustrates the processing of the inertia system when the system receives touch input, in one embodiment.

25 **[0012]** Figure 6 is a flow diagram that illustrates the processing of a multi-touch application using the inertia system to process inertia events, in one embodiment.

[0013] Figure 7 is a flow diagram that illustrates the processing of the simulation component of the inertia processing system, in one embodiment.

DETAILED DESCRIPTION

30 **[0014]** The inertia system provides a common platform and API for applications to extend the input received from various multi-touch hardware devices to simulate real-world behavior of objects. For example, real world objects do not typically stop moving when a user stops pushing them, but rather exhibit some inertia and keep moving until friction slows them finally to a stop. In some

embodiments, the touch input first goes through a process to interpret the movement of one or more contacts as manipulations. Manipulations map more directly to user intentions than do individual touch inputs and add support for basic transformation of objects using multiple touch contacts. An application can use

5 manipulations to support rotating, resizing, and translating multiple objects (e.g., photos) at the same time. The manipulations may be described as two-dimensional (2D) affine transforms that contain rotation, scale (e.g., zoom), and translation (e.g., pan) information.

[0015] Each touch of the multi-touch hardware is called a contact. For example,

10 when a user sets his/her finger on the multi-touch hardware, moves his/her finger around, and lifts his/her finger, that series of events is a single contact. For example, if the user moves two contacts closer together or further apart, the system may determine that the user is scaling (e.g., zooming into or out from) an object. As another example, if the user moves multiple contacts in a circular

15 motion, then the system may interpret the movement as a rotation of an object.

Each application can define objects that are relevant differently, so it is up to the application to attach an instance of the system (called a manipulation processor) to each object that a user can manipulate using touch input within the application.

For example, a photo browsing application may attach a manipulation processor

20 to each displayed photo, so that the user can move the photos around, scale the photos, rotate the photos, and so forth.

[0016] The manipulations handled by the application only describe the

movement of an object based on the movement of contacts. However, to move naturally, objects should also exhibit physical characteristics such as elasticity and

25 deceleration. When a user lifts all contacts from an object, the inertia system provides additional manipulation events to the application so that the application can handle the events as if the user was still moving the object with touch.

However, the inertia system actually generates the events based on a simulation of the behavior of the objects. For example, if the user lifted the contacts while

30 the object had a velocity in a particular direction, then inertia system continues sending events that indicate that the object is moving in that direction, slowing down over time as the object decelerates. If the user moves an object into another object, such as the edge of the screen, the inertia system sends manipulation events based on the boundary characteristics of the objects. For

example, if an application author defines two objects as being elastic, then the two objects may bounce off each other when a user moves the objects into each other. Thus, the inertia system provides more realistic movement for application objects that a user manipulates using multi-touch hardware and the API provides

5 a consistent feel to manipulations across applications.

[0017] Figure 1 is a block diagram that illustrates components of the inertia system, in one embodiment. The inertia system 100 includes a hardware interface 110, one or more manipulation processors 120, an input transformation component 130, a simulation component 140, and an application interface 150.

10 Each of these components is described in further detail herein.

[0018] The hardware interface 110 communicates with the hardware to receive touch contacts and movements. The hardware interface 110 may include several subcomponents that work together to provide touch input information. For example, the operating system may provide a common driver model for multi-15 touch hardware manufacturers to provide touch information for their particular hardware. The operating system may translate touch information received through this model into window messages (e.g., WM_TOUCH described herein) and pass these messages to the application. Thus, the hardware interface 110 may involve the coordination of the hardware, a hardware driver, and an operating

20 system layer. The result is a series of messages to the inertia system that identify a particular contact (e.g., touch of a finger), and the coordinates of the contact over time. For example, the operating system may provide a message when a new contact is set down on the multi-touch hardware, a message each time the contact moves, and a message when the contact is lifted away from the multi-25 touch hardware.

[0019] One or more manipulation processors 120 use the input transformation component 130 to interpret movement of each contact associated with a particular application object. The manipulation processor 120 may determine that a user is using multiple contacts to perform a single action. For example, a user could

30 touch a photo with all five fingers of one hand and twist his/her hand to indicate an intention to rotate the photo. The manipulation processor 120 receives five separate contacts (one for each finger) and the change in coordinates of each contact as the user rotates his/her hand. The manipulation processor 120 determines that each contact is grabbing the same object and performing the

same rotation. The system will inform the application that the user rotated the object, but the application can ignore whether the user used two, five, or any particular number of fingers or other contacts to perform the rotation. This greatly simplifies the authoring of the application because the application author can

5 handle those types of manipulations that are relevant to the application and leave it to the inertia system to interpret the meaning of each low-level touch input received from the multi-touch hardware.

[0020] The manipulation processor 120 uses the input transformation

component 130 to make determinations about the meaning of received

10 movements of various contacts, both alone and in concert. For example, if a user is manipulating a photo with two fingers, which creates two corresponding input contacts, then the manipulation processor 120 uses the input transformation component 130 to determine the meaning of relative movements between the two contacts. If the two contacts move apart, then the input transformation component

15 130 may determine that the user is scaling the object to change the object's size.

If the two contacts rotate, then the input transformation component 130 may

determine that the user is rotating the object. If the two contacts both slide in a particular direction, then the input transformation component 130 may determine the user is panning the object to a new location. Although each type of movement

20 is discussed separately, note that a user can make all three types of movements at the same time, and the input transformation processor can report the overall transformation to the application. For example, a user can rotate, scale, and pan an object all in one motion.

[0021] The simulation component 140 simulates the continued movement of an

25 application object after the user stops touching the object based on initialization parameters and constraints defined for the object. An application may initialize the simulation component 140 with the final state of the manipulation processor 120 associated with the object. The application may also define various characteristics of the object, such as how the object's boundaries should behave.

30 The simulation component 140 uses techniques based on physics to simulate the behavior of the object for a period after the user releases the object. For example, the simulation component 140 may continue to fire notifications to the application in the same form as the manipulation events received by the application while the user was moving the object. The application can then focus on reacting to the

movement of the object rather than being concerned with what actions (user or physical) caused the object to move. Those of ordinary skill in the art will recognize numerous well-known techniques for simulating the equivalent physical behavior of virtual objects in software.

5 **[0022]** The application interface 150 communicates with the application to receive information and provide manipulation transforms to the application. The application interface 150 receives initialization information from the application. The initialization information may specify which types of transforms the application object supports for a particular object and associated manipulation processor as
10 well as initialization data for the simulation component 140 when the user is no longer moving the object. For example, some application objects may support scaling but not rotation. The initialization information may also specify a pivot point of the object. The inertia system provides manipulation transforms to the application through the application interface. For example, when the inertia
15 system receives low-level touch input that the system interprets as a recognized transform (e.g., a rotation), the system fires an event to notify the application about the manipulation. The application processes the manipulation transform to modify the object based on the transform. For example, if the user rotated the object, then the application may store the new orientation of the object to use the
20 next time the application displays the object. As another example, if the object continued to rotate after the user released it based on calculations of the simulation component 140, then the application may store the new orientation of the object.

[0023] The computing device on which the system is implemented may include a
25 central processing unit, memory, input devices (e.g., keyboard and pointing devices), output devices (e.g., display devices), and storage devices (e.g., disk drives). The memory and storage devices are computer-readable media that may be encoded with computer-executable instructions that implement the system, which means a computer-readable medium that contains the instructions. In
30 addition, the data structures and message structures may be stored or transmitted via a data transmission medium, such as a signal on a communication link. Various communication links may be used, such as the Internet, a local area network, a wide area network, a point-to-point dial-up connection, a cell phone network, and so on.

[0024] Embodiments of the system may be implemented in various operating environments that include personal computers, server computers, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, digital cameras, network PCs,

5 minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and so on. The computer systems may be cell phones, personal digital assistants, smart phones, personal computers, programmable consumer electronics, digital cameras, and so on.

[0025] The system may be described in the general context of computer-

10 executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, and so on that perform particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various
15 embodiments.

[0026] Figure 2 is a data flow diagram that illustrates a typical operating environment of the inertia system and the flow of data between components, in one embodiment. A multi-touch hardware device produces inputs 210 through a hardware interface. For example, the hardware may send the inputs 210 to an

20 operating system through a software driver provided by the hardware manufacturer. The hardware interface provides input events 220 to an application 230. For example, an application may inform the operating system that the application 230 supports multi-touch user input and register to receive messages related to multi-touch user input. The application 230 receives low-level touch

25 input information as input changes 240 and forwards the input changes 240 to a manipulation system 250. For example, the input changes 240 may describe each movement of one or more touch contacts with the hardware using a set of coordinates that indicate each contact's current position and other movement characteristics. The manipulation system 250 interprets the input changes 240

30 and notifies the application 230 of one or more manipulation events 260 that indicate higher-level manipulations that the user is performing on a displayed object. For example, if the movement of the contacts indicates that the user intends to rotate the object, the manipulation events 260 indicate a degree of rotation.

[0027] When the user is done moving the object (e.g., when the application receives a notification that each contact touching an object has been removed from the touch hardware), the application 230 sends initialization information 270 to the inertia system 280. The inertia system 280 determines a next position of

5 the object and provides inertia events 290 similar to the manipulation events 260 that the manipulation system 250 provided when the user was moving the object.

The application 230 also provides a driving timer to periodically call the inertia system 280 to provide the next position of the object through inertia events 290.

10 The application 230 processes the inertia events in a way similar to manipulation events.

[0028] Although the diagram illustrates that the application first receives touch input and passes the touch input to the manipulation system and inertia, in some embodiments, these systems receive touch input directly from the hardware interface, interpret the touch input, and provides interpreted manipulation events

15 to the application. Likewise, the application may not know that a separate inertia system 280 provides inertia events after a user stops moving an object with touch, but rather may receive events from one interface during the time the user is moving the object and afterwards when the object is moving based on inertia.

This represents an alternative architecture that provides similar resultant 20 functionality but gives the application less control over the processing of the input.

For example, the application may not be able to define individual application objects to which the system attaches individual manipulation processors. The RTS plug-in described herein is one example of this alternative architecture for the system.

25 **[0029]** Figure 3 is a display diagram that illustrates an application object manipulated by user touch, in one embodiment. An application may simultaneously display and receive touch input for many such objects. For example, an operating system shell application may display one or more document objects stored on the user's computer desktop. In the display 310, the

30 document object 320 represents a document on the user's desktop that the user wants to drag to the recycle bin 330 using touch. The user performs a flicking motion of the document object 320 that results in system processing a first manipulation location 340, second manipulation location 350, and third manipulation location 360. The application receives the first manipulation location

340 when the user initially touches the document object 320 with one or more fingers (i.e., contacts). The application receives the second manipulation location 350 as the user slides his/her fingers across the screen. The application receives the third manipulation location when the user lifts his/her fingers from the screen.

5 The arrows 365 represent the vectors of the document object's 320 movement.

[0030] Without inertia, the document object 320 would stop at the third manipulation location 360, which is likely not what the user intends. The inertia system provides additional manipulation locations to the application as if the user was still touching and moving the document object 320 based on the document

10 object's 320 velocity when the user releases the document object 320. The application receives the first inertia-based manipulation location 370 when the application initializes the inertia system and calls the inertia system's processing function for the first time. The application receives the second inertia-based manipulation location 380 as the application continues to call the inertia system's 15 processing function. Because the final manipulation location 380 of the document object 320 is over the recycle bin 330, the application processes the contact between the two objects (e.g., by placing the document object 320 in the recycle bin 330). In the example illustrated, even though the inertia system decelerates the movement of the document object 320, the document object 320 is still able to 20 move a fair distance across the display 310 based on the high initial velocity of the user's movement of the document object 320 at the start.

[0031] Figure 4 is a flow diagram that illustrates the input loop processing of a multi-touch application using the inertia system to handle manipulation events, in one embodiment. In block 410, the application receives low-level touch input. For

25 example, an operating system or instance of the inertia system receives touch contact information from multi-touch hardware and forwards the touch contact information to the application. In block 420, the application identifies the object to which the input applies. For example, the application may hit test the coordinates of the received input by comparing the coordinates with the coordinates of each 30 application object displayed by the application. If the touch input is within the boundaries of a displayed application object, then the application determines that the touch input applies to that object. In block 430, the application sends the received touch input and the information about the identified application object to a manipulation API for invoking the inertia system (see Figure 5). For example,

the application may create a numeric identifier for each application object and pass the numeric identifier to the inertia system each time touch input corresponds to that object.

[0032] In block 440, the application receives a manipulation event from the

5 inertia system that describes one or more manipulations of the identified application object. For example, the application may receive an event describing a 2D affine transform of the application object. Note that block 440 is illustrated serially after block 430 for simplicity of illustration. In practice, the application may receive many touch input events before the inertia system notifies the application
10 with a manipulation event. There is not necessarily a one-to-one mapping of touch input events to manipulation events. Because manipulation events represent a higher-level interpretation of low-level touch inputs, multiple touch inputs may make up a single manipulation event. In block 450, the application handles the received manipulation event. For example, if the received
15 manipulation event is a rotation, then the application may rotate the application object on the screen and store the application objects new location for use when the application object is displayed again. The inertia system frees the application from performing steps specific to a particular multi-touch hardware device or even from knowing which hardware device is providing the multi-touch input. In
20 addition, the inertia system frees the application from processing individual contact movement and allows the application to focus on processing transforms at the application object level.

[0033] In block 460, the application waits for the next touch input. For example, the application may call an operating system provided message API, such as

25 GetMessage on Microsoft Windows that waits for the next message to be delivered to the application's message queue. In decision block 470, if the application receives the next touch input, then the application loops to block 410 to process the input, else the application loops to block 460 to continue waiting for further input. When the application closes, the application exits the input loop (not
30 shown).

[0034] Figure 5 is a flow diagram that illustrates the processing of the inertia system when the system receives touch input, in one embodiment. In block 505, the system receives touch input along with information identifying an application object with which the touch input is associated. For example, the touch input may

include coordinates or other location information of one or more touch contacts, and the application object information may include an identifier that the application assigned to a particular displayed object that the touch input is over on the multi-touch hardware. In block 510, the system identifies a manipulation processor

5 associated with the application object. In decision block 520, if the system has not previously associated a manipulation processor with the application object, then the system continues at block 530, else the system continues at block 540. In block 530, the system creates a manipulation processor and associates it with the application object, then continues at block 540.

10 [0035] In decision block 540, if the received touch input indicates that the application received a new contact (e.g., a touch down event), then the system continues at block 550, else the system continues at block 560. For example, a user may make initial contact of a finger with an on-screen object, or set down another finger (i.e., contact) on a previously touched object. In block 550, the

15 system adds the new contact to the list of contacts associated with the manipulation processor, and then continues at block 560. In decision block 560, if the received touch input indicates that the application received notification that a touch contact was removed (e.g., a touch up event), then the system continues at block 570, else the system continues at block 580. For example, the user may lift

20 one or more fingers from a previously touched object. In block 570, the system removes the contact from the list of contacts associated with the manipulation processor, and then continues at block 580. In block 580, the system processes the touch input to determine any manipulations represented by the touch input. For example, touch movement may indicate a rotation or translation manipulation,

25 while touch contact removal may indicate completion of a manipulation. In block 590, the system fires a manipulation event to send transform information describing the manipulation to the application. For example, the system may provide a degree of angular rotation of the object to the application. After block 590, these steps conclude.

30 [0036] Figure 6 is a flow diagram that illustrates the processing of a multi-touch application using the inertia system to process inertia events, in one embodiment. In block 610, the application determines that the user has released an object. For example, following the handling of a manipulation event in block 450 of Figure 4, the application may receive an indication that the manipulation is complete or that

the user has lifted all contacts that were touching an application object. In block 620, the application initializes the inertia system. For example, the application may pass a reference to the manipulation processor that was handling the movement of the object and other initialization information. In block 630, the

5 application sets a timer that will drive the inertia processing period of the inertia system. For example, the application may set a timer that will fire every 100 milliseconds to process the next movement increment of an object. In block 640, the application calls the processing function of the inertia system (see Figure 7). For example, the inertia system may provide a "Process" function that the
10 application calls to inform the inertia system that it is time to perform simulation for the period since the last simulation period.

[0037] In block 650, the application receives one or more inertia events that describe manipulations of the object (e.g., rotation, translation, and/or scaling) based on simulated inertia. For example, if the object was traveling in a particular
15 direction, the application may receive an inertia event that describes a translation manipulation in that direction. As another example, if the object was expanding when the user released it, the application may receive an inertia event that describes a scaling manipulation. Note that block 650 is illustrated serially after block 640 for simplicity of illustration. In practice, the application may call the

20 inertia processing function several times before the inertia system notifies the application with an inertia event. There is not necessarily a one-to-one mapping of calls to the processing function and inertia events. On the other hand, the inertia system may notify the application of multiple inertia events after a single call to the processing function.

25 **[0038]** In block 660, the application handles the received inertia event based on the meaning (e.g., an effect) of the manipulation in the context of the particular application. For example, if the received inertia event is a rotation, then the application may rotate the application object on the screen and store the application objects new location for use when the application displays the
30 application object again. In decision block 670, if the inertia events are complete, then these steps conclude, else the system continues at block 680. The inertia system may inform the application that a particular simulated manipulation is complete as a return value from the process function or through the notifications provided to the application (e.g., through a Component Object Model (COM) event

interface). In block 680, the application waits for the next firing of the timer, then loops to block 640 to call the inertia system processing function.

[0039] Figure 7 is a flow diagram that illustrates the processing of the simulation component of the inertia processing system, in one embodiment. In block 710,

5 the component receives initial simulation parameters. For example, an application or manipulation processor may provide the final state of an application object when a user stopped touching the object. In block 720, the component initializes a simulation engine that performs calculations based on physics to determine the behavior of an object based on the parameters. For example, the simulation engine may provide realistic deceleration or elasticity behavior for application objects set in motion by user touch input. In block 730, the component receives a process call from the application. The application or other component drives the simulation process by repeatedly calling a processing function at regular intervals to move the simulation forward. The simulation component may also internally 10 generate the timer.

[0040] In block 740, the component simulates movement of the object based on the initial parameters, any previous processing, and the time passed since the last process call. The process call may also provide a timestamp that indicates the time that the application wants the simulation to use. This allows the application

20 to simulate application behavior in other than real-time (e.g., for application testing or debugging). In decision block 750, if the movement is complete, then the component continues at block 760, else the component continues at block 770.

The component may determine that the movement is complete based on factors such as whether the object is still moving or whether the object movement has

25 fallen below a certain threshold. In block 760, the component sets a completion flag on the next inertia event. In block 770, the component fires an inertia event to send transform information describing the current movement (e.g., as a manipulation) to the application. For example, the system may provide a degree of angular rotation of the object to the application. After block 770, these steps 30 conclude.

[0041] In some embodiments, the inertia system receives object constraints from the application. For example, the application may define the elasticity of an object, friction coefficient (to determine how an object decelerates), boundary characteristics of the object, and so forth. For example, an application author may

define rigid objects that the user can move and a bouncy application window edge, so that objects moved into the window edge bounce back from the window edge when the user releases them.

[0042] In some embodiments, the inertia system receives initial object state information from a manipulation system that was tracking the movement of the object when the user was manipulating the object with touch input. For example, the manipulation system may track the current position of each object, the historical movement of the object, the linear and angular velocity of the object, and so forth. The application author can provide the output of the manipulation to the inertia system to initialize the inertia system, so that the inertia system can smoothly continue the past movement of the object and slow it down based on appropriate physics and characteristics of the object.

[0043] In some embodiments, the inertia system receives limits on the movement of objects from the application. For example, the application author may define an upper bound on the distance that an object can move once a user releases the object. As another example, the application may define an upper bound on how long the object can move once a user releases the object. These and other limits allow the application author to adjust the inertia system to suit the types of objects manipulated by the application and to enhance the user experience with the application.

[0044] In some embodiments, the inertia system does not provide additional movement for objects with movement below a predefined threshold. The threshold may be configurable by the application. For example, the inertia system may have a particular object linear or angular velocity below which the system will not continue movement of the object after the user releases the object. If the object is not moving very fast when the user releases it, the user may expect that the object will stay put and not continue to move. The threshold allows the application or author of the inertia system to determine the level of movement after manipulation that provides a good user experience.

[0045] In some embodiments, the inertia system receives instructions to simulate movement incrementally from the application. For example, the inertia system may provide a "Process" or "DoWork" function that the application calls to instruct the inertia system to perform a portion of the overall simulation. The inertia system may expect the application to set a timer or otherwise periodically

call the function to cause the inertia system to simulate movement over time according to a natural timeline. The application can affect the characteristics of the manipulation events provided by the inertia system by varying how often the application calls the function. In other embodiments, the inertia system uses an

5 internal timer to provide manipulation events on a regular schedule until each object has stopped moving (e.g., due to deceleration or other simulated forces).

[0046] In some embodiments, the inertia system is part of a message-based operating system, and the system receives messages related to touch input that the operating system receives from the hardware. For example, using a paradigm

10 similar to WM_MOUSEMOVE for mouse messages, future versions of Microsoft Windows may provide a WM_TOUCH message that contains low-level touch movement information received from multi-touch hardware. The operating system

may also provide finer grained messages, such as WM_TOUCHDOWN (when a new contact is made with the multi-touch hardware), WM_TOUCHMOVE (when

15 an existing contact moves), and WM_TOUCHUP (when a contact is lifted from the multi-touch hardware). An application that receives a WM_TOUCH-related message can invoke the inertia system and pass the message to the inertia system for interpretation and processing. The application then receives higher-level events that represent the inertia system's interpretation of the manipulation

20 intended by the user based on the received low-level touch movement information.

[0047] In some embodiments, the inertia system receives low-level touch movement information from specialized hardware, such as a real-time stylus. For example, the Microsoft Tablet PC Software Development Kit (SDK) provides a

25 real-time stylus (RTS) component that application authors can extend with hooks. RTS hooks receive input from the RTS hardware and can perform processing on the received input. The inertia system may provide a hook that an application can insert into the RTS component to automatically process RTS and other input to manipulate application objects as described herein. The RTS hook provides a

30 different way for the inertia system to receive input, but the inertia system interprets input and fires events to the application describing manipulations implied by the input as previously described. A user may use a combination of stylus and touch input. For example, the user may draw an object with the stylus and then rotate the object using his/her fingers.

[0048] In some embodiments, the inertia system is part of a common control that an application can invoke to provide a common user interface. Microsoft Windows provides common controls for displaying lists, trees, buttons, and so forth.

Likewise, the inertia system may provide a multi-touch based control for

5 manipulating application objects in the ways described herein. For example, the system may provide a scatter control that allows the user to display one or more objects and manipulate the objects. The scatter control handles processing of low-level touch input and associating the input with a particular application object, and the application receives events from the control to handle the manipulations
10 of the application objects. For example, if the control indicates that the user resized an object, then the application may store the objects new size.

[0049] In some embodiments, the inertia system performs the processing described herein in three dimensions. Although two-dimensional multi-touch hardware is described herein, those of ordinary skill in the art will recognize that

15 the processing of the system described herein can be applied equally well to three-dimensional (3D) manipulations if hardware is available to provide coordinate movement in three dimensions. For example, hardware that detects pressure or uses cameras to detect 3D movement of a user's fingers could provide the coordinates of movement in the third dimension to the inertia system,
20 and the inertia system could then produce 3D transforms that describe manipulations (e.g., rotation, scaling, and translation) of objects in multiple 3D directions.

[0050] The following table defines one API that the inertia system provides to applications for providing inertia-based movement to application objects following
25 user touch-based movement of the objects.

Properties:	
BoundaryBottom	Limits how far towards the bottom of the screen the target object can move.
BoundaryLeft	Limits how far towards the left of the screen the target object can move.
BoundaryRight	Limits how far towards the right of the screen the target object can move.
BoundaryTop	Limits how far towards the top of the screen the target object can move.
DesiredAngularDeceleration	Specifies the desired rate that the target object will stop spinning in radians per millisecond.

DesiredDeclaration	Specifies the desired rate at which translation operations will decelerate.
DesiredDisplacement	Specifies the desired distance that the object will travel.
DesiredExpansion	Specifies the desired change in the object's average radius.
DesiredExpansionDeceleration	Specifies the rate at which the object will stop expanding.
ElasticMarginBottom	Specifies the bottom region for bouncing the target object.
ElasticMarginLeft	Specifies the leftmost region for bouncing the target object.
ElasticMarginRight	Specifies the rightmost region for bouncing the target object.
InitialAngularVelocity	Specifies the rotation of the target when movement begins.
InitialOriginX	Gets or puts the property designating the horizontal position of a target object. This property specifies the starting horizontal location for a target with inertia.
InitialOriginY	Gets or puts the property designating the vertical location for a target object. This property specifies the starting vertical location for a target with inertia.
InitialRadius	Specifies the distance from the edge of the target to its center before the object was changed.
InitialTimestamp	Specifies the starting timestamp for a target object with inertia.
InitialVelocityX	Specifies the initial movement of the target object on the horizontal axis.
InitialVelocityY	Specifies the initial movement of the target object on the vertical axis.
Methods:	
HRESULT Reset();	Initializes the processor with initial timestamp.
HRESULT Process([out] BOOL* completed);	Performs calculations for the given tick and can raise the Delta or Completed event depending on whether extrapolation is completed or not. If extrapolation finished at the previous tick, the method is no-op.
HRESULT ProcessTime([in] DWORD timestamp, [out] BOOL* completed);	Performs calculations for the given tick and can raise the Delta or Completed event depending on whether extrapolation is completed or not. If extrapolation finished at the previous tick, the method is no-op.
HRESULT Complete();	Raises the Completed event.
HRESULT CompleteTime([in] DWORD timestamp);	Processes the given tick and raises the Completed event.

Events:	
HRESULT ManipulationStarted([in] FLOAT x, [in] FLOAT y)	Handles the event for when a manipulation begins.
HRESULT ManipulationDelta([in] FLOAT x, [in] FLOAT y, [in] FLOAT translationDeltaX, [in] FLOAT translationDeltaY, [in] FLOAT scaleDelta, [in] FLOAT expansionDelta, [in] FLOAT rotationDelta, [in] FLOAT cumulativeTranslationX, [in] FLOAT cumulativeTranslationY, [in] FLOAT cumulativeScale, [in] FLOAT cumulativeExpansion, [in] FLOAT cumulativeRotation)	Handles events that happen when a manipulated object changes
HRESULT ManipulationCompleted([in] FLOAT x, [in] FLOAT y, [in] FLOAT cumulativeTranslationX, [in] FLOAT cumulativeTranslationY, [in] FLOAT cumulativeScale, [in] FLOAT cumulativeExpansion, [in] FLOAT cumulativeRotation)	Handles the event when manipulation finishes.

[0051] In the table above, the inertia system may provide the listed events on the same interface on which an application was previously receiving events based on user movement.

5 **[0052]** From the foregoing, it will be appreciated that specific embodiments of the inertia system have been described herein for purposes of illustration, but that various modifications may be made without deviating from the spirit and scope of the invention. For example, although the system has been described in the context of multi-touch manipulations, the system provides simulation of inertia that could be used in other contexts, such as games and other areas where simulation
10 is commonly used. Accordingly, the invention is not limited except as by the appended claims.

CLAIMS

I/We claim:

1. A computer-implemented method for providing realistic movement of objects manipulated using multi-touch input, the method comprising:
 - 5 determining 610 that a user has released an application object by removing one or more contacts from a multi-touch input device;
 - invoking 640 an inertia API to process a simulation of movement of the application object for a current inertia processing period, wherein the inertia API provides an application-independent platform for simulating realistic movement
 - 10 independent of a type of the application object;
 - receiving 650 an inertia event that describes a manipulation of the object based on simulated inertia;
 - handling 660 the received inertia event based on an effect of the manipulation in a context of the application by modifying the application object.
 - 15 2. The method of claim 1 further comprising, after determining that the user released the object, initializing the inertia system by passing a reference to a manipulation processor that was handling the movement of the object before the user released the object.
 3. The method of claim 1 further comprising, after determining that the user released the object, initializing the inertia system by passing one or more parameters describing a state of the object when the user released the object.
 4. The method of claim 1 wherein receiving the inertia event comprises receiving information describing a 2D affine transform of the application object.
 5. The method of claim 1 wherein the manipulation described by the inertia event includes at least one of a rotation manipulation, a translation manipulation, and a scaling manipulation.
 - 25 6. The method of claim 1 wherein the received inertia event is a rotation and wherein handling the received inertia event comprises rotating the application object on a display.
 - 30 7. The method of claim 1 further comprising setting a timer that determines a length of the inertia processing period, and wherein the application invokes the inertia API at each firing of the timer.

8. The method of claim 7 further comprising receiving an indication from the inertia API that the simulation is complete, and expiring the timer when the indication is received.

9. The method of claim 1 wherein receiving the inertia event comprises

5 receiving a notification through a COM event interface.

10. A computer system for handling touch input from multi-touch hardware, the system comprising:

10 a hardware interface 110 configured to communicate with the multi-touch hardware to receive touch contact information and movements of the touch

10 contacts;

one or more manipulation processors 120 configured to manage interpretation of movement of each contact associated with a particular application object;

15 an input transformation component 130 configured to interpret a meaning of received movements of various contacts to produce manipulations of application objects;

15 a simulation component 140 configured to simulate continued movement of the application object after a user stops touching the object;

20 an application interface 150 configured to communicate with the application to receive contact movement information and provide manipulation transforms to the application.

11. The system of claim 10 wherein the simulation component is further configured to receive initialization parameters from an application through the application interface.

25 12. The system of claim 11 wherein the initialization parameters are provided by a manipulation processor that was managing the application object before the user stopped touching the object.

13. The system of claim 10 wherein the simulation component is further configured to receive boundary constraints of the application object.

30 14. The system of claim 10 wherein the simulation component is further configured to apply laws of physics to determine the continued movement of the object based on a past velocity and direction of movement of the object.

15. The system of claim 10 wherein the input transformation component and simulation component generate events in a similar format so that the application can handle events from both components similarly.

16. A computer-readable storage medium encoded with instructions for
5 controlling a computer system to simulate movement of an application object previously moved by multi-touch input, by a method comprising:

receiving 710 one or more initial simulation parameters that provide a last state of the application object when a user released the application object;

10 initializing 720 a simulation engine that performs calculations based on physics to determine the behavior of the application object object based on the initial simulation parameters;

receiving 730 an indication that a current simulation period is due for moving a simulation forward;

15 simulating 740 movement of the application object based on the initial parameters, any previous processing, and the time passed since any previous simulation period; and

firing 770 an inertia event to send transform information describing a current movement of the application object to the application. For example, the system may provide a degree of angular rotation of the object to the application.

20 17. The computer-readable medium of claim 16 wherein the simulation engine simulates realistic deceleration behavior for the application object after it was set in motion by the user touching the object.

18. The computer-readable medium of claim 16 wherein the simulation engine simulates realistic elasticity behavior for the application object when the
25 application object overlaps with another element on a display.

19. The computer-readable medium of claim 16 wherein receiving the indication that the current simulation period is due comprises receiving a call from the application to a simulation processing function.

20. The computer-readable medium of claim 16 further comprising determining
30 whether the simulation is complete and if the simulation is complete, informing the application that the simulation is complete.




Figure 1




Figure 2





Figure 3

4/7

*Figure 4*






Figure 5

6/7

**Figure 6**

7/7

**Figure 7**