



Clutter and Multitouch

# 1 Contents

2	<b>Clutter and Multitouch</b>	<b>2</b>
3	Introduction . . . . .	2
4	Multi-touch . . . . .	2
5	The multi-touch stack . . . . .	2
6	Requirements . . . . .	5
7	Approach . . . . .	7
8	Risks . . . . .	10
9	Smooth panning . . . . .	11
10	Requirements . . . . .	11
11	Approach . . . . .	11
12	Risks . . . . .	11
13	Documentation . . . . .	11
14	Design notes . . . . .	12

## 15 Clutter and Multitouch

### 16 Introduction

17 This document explains Collabora’s design about several issues related to the  
18 main UI toolkit used in Apertis: Clutter.

### 19 Multi-touch

20 This section describes the support for multi-touch (MT) events and gestures in  
21 the Apertis middle-ware. It will be explained which requirements Collabora will  
22 support and the general architecture of MT on Linux and X.Org.

23 When we mention MT in this document, we refer to the ability of applications  
24 to react to multiple touch event streams at the same time. By gestures, we refer  
25 to the higher-level abstraction that groups individual touch events in a single  
26 meaningful event.

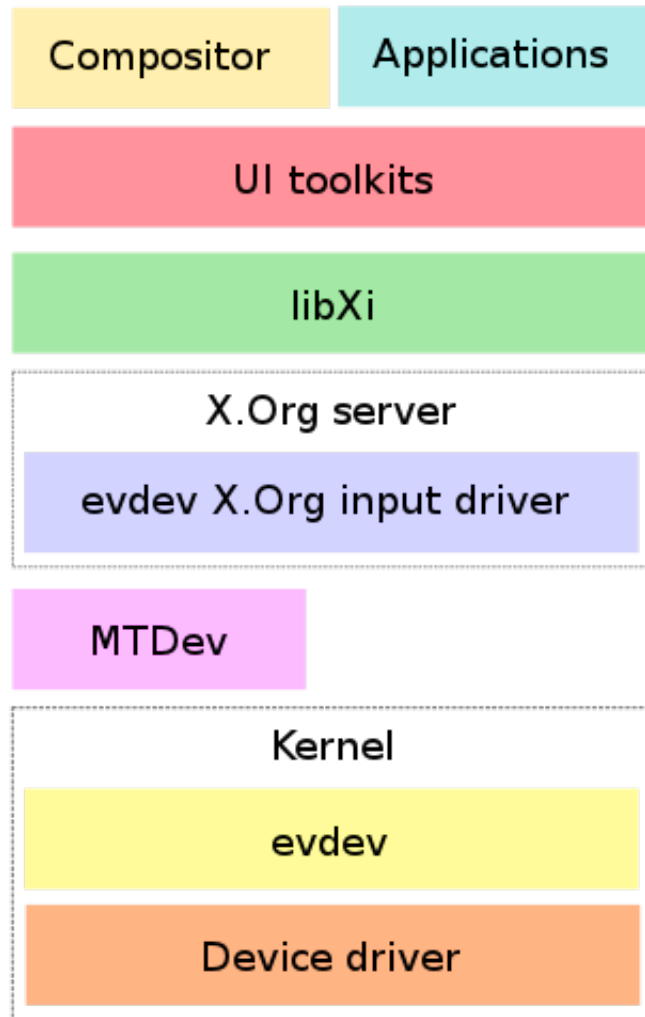
27 At this point of time (Q1 2012), MT and gesture functionality is implemented  
28 in several consumer products but is only starting to be available in frameworks  
29 of general availability such as X.Org and HTML. As will be explained later, this  
30 is reflected in X.Org just being released with [MT functionality](http://lists.x.org/archives/xorg-announce/2012-March/001846.html)<sup>1</sup>, Clutter not  
31 having MT support in a release yet, and the lack of high level gesture support  
32 in X.Org-based toolkits. In the same vein, MT is not yet standardized in the  
33 web. This design will discuss the challenges posed by these facts and ways of  
34 overcoming them.

### 35 The multi-touch stack

36 For the purposes of this document, the MT stack on X.Org is layered as follows:

---

<sup>1</sup><http://lists.x.org/archives/xorg-announce/2012-March/001846.html>



37

38 **Compositor**

39 The compositor has to be able to react to gestures that may happen anywhere  
 40 in the display. The X server usually delivers events to the window where they  
 41 happen, so the compositor overrides this behavior by telling the X server that  
 42 it has interest in all events regardless of where they happen (specifically, it does  
 43 so by registering a passive grab on the root window).

44 The compositor will receive all events and decide for each whether it should  
 45 be handled at this level, or let it pass through to the application to which the  
 46 underlying window belongs.

47 For touch events, this isn't done for individual events but rather for touch  
48 sequences. A touch sequence is the series of touch events that starts when a  
49 finger is placed on the screen, and finished when it is lifted. Thus, a touch  
50 sequence belongs to a single finger, and a gesture can be composed by as many  
51 touch sequences as fingers are involved.

52 There is a period of time during which the compositor inspects the touch events  
53 as they come and decides whether it should handle this particular gesture, or  
54 if it should be ignored and passed to the application. If the compositor decides  
55 the later, the events that had been already delivered to the compositor will be  
56 replayed to the application.

57 The period of time that the compositor needs to decide whether a gesture should  
58 be handled by applications should be as small as possible, in order to not disrupt  
59 the user experience.

60 An application can tell the X server that it doesn't want to have to wait until  
61 the compositor has let the touch sequence pass. In that case, either the gesture  
62 shouldn't cause any visible effects in the UI, or should be reverted in case the  
63 compositor ends up deciding to handle the touch sequence by itself.

## 64 **Applications**

65 Widgets inside applications can react to MT events in a way similar to how  
66 they react to single-touch events. Additionally, some toolkits provide additional  
67 functionality that make it easier to react to gestures. Widgets can either react  
68 to a few predefined gestures (tap, panning, pinch, etc.), or they can implement  
69 their own gesture recognizers by means of the lower level MT events.

## 70 **UI toolkits**

71 As mentioned before, UI toolkits provide API for reacting to MT events and  
72 usually also functionality related to gestures. Because MT is so new to X.Org,  
73 UI toolkits based on X.Org don't implement yet everything that applications  
74 would need regarding MT and gestures, so for now additional work needs to  
75 happen at the application level.

## 76 **libXi**

77 This library allow toolkits to communicate with the XInput extension in the  
78 X.Org server. Communication with the X.Org server is asynchronous and com-  
79 plex, so having a higher level library simplifies this interaction.

## 80 **X.Org server**

81 The X.Org server delivers input events to each application based on the coordi-  
82 nates of the event and the placement of the application windows.

83 **evdev X.Org input driver**

84 This input driver for X.Org uses udev to discover devices and evdev to get input  
85 events from the kernel and posts them to the X.Org server.

86 If it is needed to apply a jitter-reduction filter and it's impossible to do so in  
87 the kernel device driver, then we recommend to patch the evdev X.Org input  
88 driver.

89 **MTDev**

90 For device drivers that use the legacy MT protocol as opposed to the new slots  
91 protocol, the X.Org input driver will use libmtdev to translate events from the  
92 old protocol (type A) to the new one (type B).

93 See [http://www.kernel.org/doc/Documentation/input/multi-  
touch-protocol.txt](http://www.kernel.org/doc/Documentation/input/multi-<br/>94 touch-protocol.txt)

95 **Kernel event driver (evdev)**

96 This kernel module will emit input events in a protocol that the evdev X.Org  
97 input driver can understand.

98 **Kernel device driver**

99 This kernel module is hardware-dependent and will interface with the hardware  
100 and pass input events to the evdev event driver.

101 **Requirements**

102 **Multitouch event handing in Clutter applications**

103 Clutter will have APIs for reacting to multi-touch input events and for recog-  
104 nizing gestures.

105 The Apertis middleware will have the support that Clutter requires to provide  
106 MT and gestures functionality, as described later in this document.

107 Though it is expected that Clutter will eventually provide support for recog-  
108 nizing a few basic gestures such as taps and panning, more advanced gestures  
109 will have to be implemented outside Clutter. In the Apertis case, recognizing  
110 additional gestures will have to be done by the applications themselves or by  
111 the SDK API.

112 New gestures will be developed using the gesture framework in Clutter, regard-  
113 less of whether the gesture recognition is implemented in the SDK, compositor  
114 or applications. In other words, new gestures can be developed making use of  
115 the Clutter API, but the code that implements them can belong to applications,  
116 compositors or libraries. No modifications to Clutter are required to implement  
117 new gestures.

118 **Full-screen event handling in applications**

119 Applications should be able to handle events anywhere in the screen even if  
120 their windows don't cover the whole of it. For example, there may be windows  
121 belonging to the system such as a launcher panel or a status bar in the screen,  
122 but a gesture that starts in one of the auxiliary windows should be handled by  
123 the focused application.

124 **Multi-touch event handling in Mutter**

125 The compositor based on Mutter will be able to react to multi-touch input  
126 events and recognize gestures using the same Clutter API as applications.

127 The compositor will be able to register for MT sequences before applications  
128 get them, so it can claim ownership over them in case a system-wide gesture  
129 is detected. Even then, applications will be able to get all events that happen  
130 in their windows though care needs to be taken in case the compositor ends up  
131 claiming ownership.

132 **Multitouch event handling in web applications**

133 Although there are no approved specifications yet on how browsers should ex-  
134 pose MT events to web content, some browsers have started already to provide  
135 experimental APIs and some websites are using them. Most notable are the  
136 Safari browser on iOS and websites that target specifically iPhone and iPad  
137 devices, though nowadays other WebKit-based browsers implement MT events  
138 and more websites are starting to use them.

139 A [spec](#)<sup>2</sup> is being drafted by the W3C on base on the WebKit implementation,  
140 but attention must be paid to the fact that because it's still a draft, it may  
141 change in ways that are incompatible with WebKit's implementation and the  
142 later may not always be in sync with the spec.

143 **Support two separate touchscreens**

144 The Apertis middleware will be able to drive two screens, each with multi-touch  
145 support.

146 **Support out-of-screen touch events**

147 The reactive area of the touchscreen may extend past the display and the user  
148 will be able to interact with out-of-screen UI elements.

149 **Actors with bigger reactive area**

150 So the UI is more tolerant to erratic touch events (caused for example by a  
151 bumpy road), some actors will be reactive past the boundaries of their represen-  
152 tation in the screen.

---

<sup>2</sup><http://dvcs.w3.org/hg/webevents/raw-file/tip/touchevents.html>

153 **Approach**

154 **Multitouch event handing in Clutter applications**

155 MT support in X.Org is very recent, so Collabora will have to update several  
156 components (mainly belonging to X) in the stack because Precise is not going  
157 to ship with the versions that are needed.

158 Clutter 1.8 had support for single-touch gestures. In 1.10, support for multi-  
159 touch event handling landed, and it is expected that for 1.12 (August 2012),  
160 support for multi-touch gestures will be added. It's also planned to have support  
161 for rotation and pinch gestures in [Clutter 1.12](#)<sup>3</sup>.

162 **Full-screen event handing in applications**

163 In order for applications to be able to handle events anywhere in the screen  
164 even if their windows do not cover the whole of it, applications will have to set  
165 grabs on the other visible windows, even if they don't belong to the application  
166 process.

167 So in the case that the currently-focused application is displayed along with a  
168 launcher panel and a status bar, the application process should set a grab on  
169 those windows for the events that it is interested in. When another application  
170 becomes focused, the first one releases the grab and the second takes it.

171 In order to make sure that the second application takes the grab as soon as  
172 possible, it should try calling `XIGrabTouchBegin` repeatedly until it stops failing  
173 with `BadAccess` (this will happen once the previously focused application has  
174 removed its grab).

175 So the compositor can still handle system-level gestures as explained in 2.3.3,  
176 it will have to set a grab on an invisible window that is the parent of each  
177 additional window.

178 This is so complex because this scenario is a bit removed from the design of  
179 event delivery in X. In Wayland, as the compositor is also the display server  
180 and has total control over event delivery, it could redirect events to application  
181 windows instead of to the panels.

182 **Multi-touch event handling in Mutter**

183 Current releases of Mutter use XLib for event management, which doesn't sup-  
184 port multi-touch. To Collabora's knowledge, there aren't as of yet any Mutter-  
185 based products that make use of system-wide multi-touch gestures.

186 See [https://wiki.gnome.org/GnomeOS/Design/Whiteboards/  
187 Touchscreen#Mutter\\_problems](https://wiki.gnome.org/GnomeOS/Design/Whiteboards/Touchscreen#Mutter_problems)

188 Collabora has [modified](#)<sup>4</sup> Mutter to allow plugins to register for touch events and

<sup>3</sup><http://wiki.clutter-project.org/wiki/ClutterRoadMap#1.12>

<sup>4</sup><http://blog.tomeuvizoso.net/2012/09/multi-touch-gestures-in-gnome-shell.html>

189 to pass them to Clutter so subclasses of ClutterGestureAction can be used.

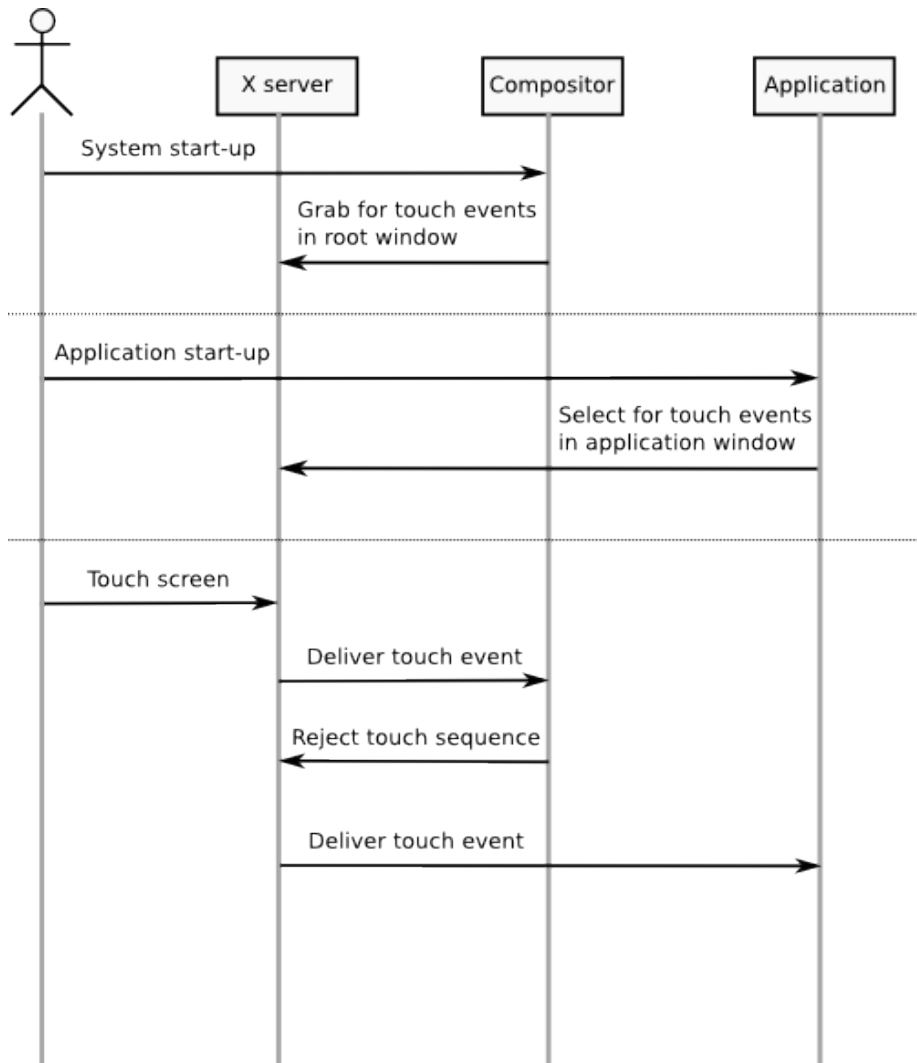
190 Though applications are able to start receiving MT events even before the com-  
191 positor rejects ownership, Collabora recommends that applications don't do that  
192 and instead that all efforts are directed towards having the compositor recognize  
193 gestures as fast as possible.

194 Otherwise, it will be very hard to avoid glitches in the user experience when the  
195 compositor decides to handle a gesture and the application already started to  
196 react to it.

197 By limiting system-wide gestures to those with 4 or 5 fingers as iOS 5 does  
198 (what Apple calls *multitasking gestures*), the compositor should be able to decide  
199 whether to take ownership of a MT sequence with a minimal delay and without  
200 applications having to do anything on their side. Gestures with less than 4  
201 fingers will be considered as intended for applications and the compositor will  
202 decline ownership immediately.

203 This diagram illustrates how the different components interact when handling  
204 touch events:





205

206 If there is a window that gets placed on top of the others (specifically, the  
 207 “busy” indicator animation) and it shouldn’t get any events, it can be marked  
 208 as such with the XShapeCombineRegion call, passing an empty region and the  
 209 ShapeInput destKind.

210 See <http://article.gmane.org/gmane.comp.kde.devel.kwin/19992>

211 This way, any events that the compositor doesn’t intercept will be delivered to  
 212 the currently-focused application.

213 **Multi-touch event handing in web applications**

214 Collabora will implement the port specific bits of MT in WebKit-Clutter to en-

215 sure that it has state-of-the-art WebKit MT support, but won't be changing the  
216 behavior of the generic WebKit implementation nor working on the specification  
217 level.

218 Collabora won't be implementing any high-level gesture support because they  
219 are far from being specified and its support in browsers is very experimental.

### 220 **Support two separate touchscreens**

221 There is support in X.Org for arbitrarily matching input devices to screens,  
222 though the configuration *isn't straightforward*<sup>5</sup>. Collabora will test the simul-  
223 taneous support of 2 touch-screens and produce documentation about how to  
224 configure X.Org in this regard during the development phase of the project.

### 225 **Support out-of-screen touch events**

226 Regarding out-of-screen events support, we propose writing a daemon that lis-  
227 tens for events from the touchscreen kernel driver and that, based on its config-  
228 uration, translates those to key presses for special key codes that correspond to  
229 each button.

230 As the X evdev driver will also get those events, it has to be configured to ignore  
231 touch events outside the screen area.

232 In the SDK, a wrapper around Xephyr will be provided that synthesizes key  
233 events when buttons around the Xephyr screen are pressed, simulating the ones  
234 in the real hardware.

### 235 **Actors with bigger reactive area**

236 Clutter actors tell the Clutter framework in which area of the screen they are  
237 sensitive to pointer events. This area usually matches the area that the actor  
238 uses to display itself, but the actor could choose to mark a bigger area as its  
239 reactive area.

240 Though the Clutter maintainer has recommended this approach, he warns that  
241 the optimization that culls actors based on their paint volumes might get in the  
242 way in this case.

243 Collabora will verify that this works and communicate with the upstream com-  
244 munity in case any problem is discovered.

### 245 **Risks**

246 The *DDX*<sup>6</sup> driver provided by the hardware vendor should support having a  
247 frame-buffer that is bigger than the actual display resolution, for the out-of-  
248 screen touch events.

---

<sup>5</sup><http://www.x.org/wiki/XInputCoordinateTransformationMatrixUsage>

<sup>6</sup><http://dri.freedesktop.org/wiki/DDX>

## 249 **Smooth panning**

250 This section proposes an improvement to the kinetic scrolling functionality in Mx  
251 so that panning is smooth even when the input device’s resolution is very low.  
252 This is going to affect only to Clutter applications that use MxKineticScrollView  
253 as the container of scrollable areas.

254 The problem with input devices with low resolution is that as the finger moves  
255 during panning, the motion events received by the application are similarly low-  
256 resolution (i.e., occurring infrequently). Given that Mx currently updates the  
257 position in the scrolled view to match the position of the finger, the panning  
258 movement appears “jumpy”.

## 259 **Requirements**

260 When panning, the position in the scrolled view will smoothly interpolate to the  
261 last finger position, instead of jumping there straight away. The visual effect  
262 would be that of the scroll position lagging slightly behind the finger. The lower  
263 the resolution of the touch screen, the bigger the lag.

## 264 **Approach**

265 Collabora will rewrite the part of the [MxKineticScrollView](http://docs.clutter-project.org/docs/mx/stable/MxKineticScrollView.html)<sup>7</sup> widget that  
266 tracks the finger position when panning, ideally using the function  
267 [mx\\_adjustment\\_interpolate](http://docs.clutter-project.org/docs/mx/stable/MxAdjustment.html#mx-adjustment-interpolate)<sup>8</sup> to animate the movement along the path  
268 between the current position and the finger position. The time function  
269 ([ClutterAlpha](http://developer.gnome.org/clutter/stable/ClutterAlpha.html)<sup>9</sup>) used in the interpolation animation will be configurable, as  
270 well as the speed at which the scrolling position will follow the finger.

## 271 **Risks**

272 Upstream could decide to reject this feature when it is proposed for inclusion  
273 because of the substantial added complexity to a widget (MxKineticScrollView)  
274 that is already pretty complex. However, preliminary discussions with the Mx  
275 maintainers show that they are interested in gaining this functionality.

276 Another risk is Intel not funding all the work in Clutter that it has committed  
277 to. In that case, Collabora may need to do the work.

## 278 **Documentation**

279 The following items have been identified for future documentation work later in  
280 the project:

---

<sup>7</sup><http://docs.clutter-project.org/docs/mx/stable/MxKineticScrollView.html>

<sup>8</sup><http://docs.clutter-project.org/docs/mx/stable/MxAdjustment.html#mx-adjustment-interpolate>

<sup>9</sup><http://developer.gnome.org/clutter/stable/ClutterAlpha.html>

- 281 • Add a section to the Clutter Cookbook about implementing a ClutterGest-  
282 tureAction subclass.
- 283 • Best practices about MT gestures in user experience, both system-wide  
284 and application gestures. Compare these guidelines with the equivalent  
285 ones in iOS and Android.
- 286 • Best practices about performance with Clutter and Mx (including how  
287 to write containers which are responsive independently of the number of  
288 children and how to drag actors across the screen).
- 289 • Best practices about using and writing reusable UI components (including  
290 Mx widgets), and explicitly these subjects (to be specified further at a later  
291 stage):
  - 292 – Panning actors
  - 293 – Finger moves outside the “active area” of an actor (e.g., moving but-  
294 ton of timeline in video very fast)
  - 295 – Snap forward/backward to final position
  - 296 – Expand/shrink of groups (sliding)
- 297 • Best practices about writing applications in which functionality and UI  
298 are separate, so derivatives can be written by replacing the UI and without  
299 having to modify the rest of the application.

### 300 **Design notes**

301 The following items have been identified for future investigation and design work  
302 later in the project and are thus not addressed in this design:

- 303 • Support two separate touchscreens
- 304 • Support out-of-screen touch events
- 305 • Implement jitter reduction (there already has an algorithm), taking into  
306 account that the input driver may be a binary blob
- 307 • Implement zoom via pitch gestures in Google Earth without having access  
308 to its source code