



ELSEVIER

Available at
www.ComputerScienceWeb.com
 POWERED BY SCIENCE @ DIRECT®

 **The Journal of
 Systems and
 Software**

The Journal of Systems and Software xxx (2003) xxx-xxx

www.elsevier.com/locate/jss

Roam, a seamless application framework

Hao-hua Chu *, Henry Song, Candy Wong, Shoji Kurakake, Masaji Katagiri

DoCoMo Communications Laboratories USA, Inc., 181 Metro Drive, Suite 300, San Jose, CA 95110 USA

Received 23 December 2002; accepted 28 December 2002

6 Abstract

7 One of the biggest challenges in future application development is device heterogeneity. In the future, we expect to see a rich
 8 variety of computing devices that can run applications. These devices have different capabilities in processors, memory, networking,
 9 screen sizes, input methods, and software libraries. We also expect that future users are likely to own many types of devices. De-
 10 pending on users' changing situations and environments, they may choose to switch from one type of device to another that brings
 11 the best combination of application functionality and device mobility (size, weight, etc.). Based on this scenario, we have designed
 12 and implemented a seamless application framework called the *Roam system* that can both assist developers to build multi-platform
 13 applications that can run on heterogeneous devices and allow a user to move/migrate a running application among heterogeneous
 14 devices in an effortless manner. The Roam system is based on partitioning of an application into components and it automatically
 15 selects the most appropriate adaptation strategy at the component level for a target platform. To evaluate our system, we have
 16 created several multi-platform Roam applications including a Chess game, a Connect4 game, and a shopping aid application. We
 17 also provide measurements on application performance and describe our experience with application development in the Roam
 18 system. Our experience shows that it is relatively easy to port existing applications to the Roam system and runtime application
 19 migration latency is within a few seconds and acceptable to most non-real-time applications.

20 © 2003 Elsevier Science Inc. All rights reserved.

21

22 1. Introduction

23 The era of PC-dominated applications is about to
 24 end. Now-a-days, we see widespread use of mobile de-
 25 vices that have sufficient computing and networking
 26 capabilities to run a variety of feature-rich applications.
 27 They come in a variety of form factors, including smart
 28 pagers, cell phones, PDAs (Pocket PCs), handheld PCs,
 29 car navigation systems, and notebook PCs. An average
 30 person may already own multiple such devices. It is
 31 expected that in the near future, the number of such
 32 devices will far exceed the number of desktop PCs. As a
 33 consequence we expect that the predominant software
 34 platform will shift from PCs to mobile devices. This
 35 creates a challenge for developers to build applications
 36 that can run on different mobile device platforms. To
 37 address this challenge, we believe that there is a need to

provide an application framework for building multi- 38
 platform mobile applications. 39

We also expect a mobile user may choose to switch 40
 from one type of device to another type of device, in the 41
 middle of using an application, in order to access nec- 42
 essary application functionality or to become more 43
 mobile, based on changing situations, environments, or 44
 needs. For example, a user starts planning a vacation 45
 online using a desktop computer in his/her office. In the 46
 middle of planning, he/she receives an urgent call and 47
 must leave the office for a meeting at a remote site. The 48
 user would like to continue planning the trip on the bus 49
 or during break time between meetings. Given the need 50
 for mobility, he/she switches to a lightweight, mobile 51
 device (PDA or cell phone) away from the office. Based 52
 on this scenario, we believe that when mobile users 53
 switch devices, there is a need to allow them to move any 54
 running application effortlessly between devices. 55

Based on these two needs, developers will be required 56
 to write *seamless applications*. We define a seamless 57
 application to be an application that can run on heter- 58
 ogeneous devices and migrate at runtime among heter- 59

* Corresponding author.

E-mail addresses: haochu@docomolabs-usa.com (H.-h. Chu),
cshyus@docomolabs-usa.com (H. Song), wong@docomolabs-usa.com
 (C. Wong), kurakake@docomolabs-usa.com (S. Kurakake), katagiri@docomolabs-usa.com (M. Katagiri).

60 ogeous devices. There has been an abundance of re-
61 search in the area of mobile agents that allow applica-
62 tions to move from one host to another at runtime with
63 little or no loss of execution states (Aramira, 1999;
64 Acharya et al., 1997; Banavar et al., 2000; Fünfrocken,
65 1998; Gray et al., 1997; Lange and Oshima, 1998; Mil-
66 ojicic et al., 1998; Mitsubishi Electric ITA Horizon
67 Systems Laboratory, 1998; ObjectSpace, xxxx; Peine
68 and Stolpmann, 1997; Strasser et al., 1996; White et al.,
69 1997). These systems are mostly built on top of the Java
70 virtual machine (JVM), which provides the advantage of
71 a common runtime environment. However, these sys-
72 tems make an important assumption of *device homoge-*
73 *neity*. For example, the underlying hosts/devices must be
74 PC or PC-like devices with enough hardware/software
75 (HW/SW) capabilities to run the standard JVM. Given
76 this device homogeneity assumption and Java's "Write
77 Once, Run Anywhere" programming model, it is rela-
78 tively straightforward to realize runtime application
79 migration among similar devices. However, device ho-
80 mogeneity is not a realistic assumption in today's mobile
81 computing environment where there exists a wide range
82 of mobile devices and information appliances with dif-
83 ferent HW/SW capabilities. For an example, a cell
84 phone or a PDA in comparison to a desktop computer
85 has a slower processing speed, less memory, little or no
86 permanent storage, slower and unreliable network con-
87 nectivity, smaller screen size, and limited input capa-
88 bilities. To address device heterogeneity, Java introduces
89 mutually incompatible Java 2 Micro Edition (J2ME)
90 profiles and configurations for different classes of mobile
91 devices. As a result, Java's "Run Anywhere" ideal does
92 not apply to the device heterogeneous environment.

93 1.1. Challenges

94 The focus of the Roam system is to address the *device*
95 *heterogeneity* problem in runtime application migration.
96 Fig. 1 shows different aspects of this problem that a Java
97 application developer may face when developing a
98 seamless application.

- 99 • *Incompatible Java virtual machine configurations and*
100 *profiles*: Java provides incompatible virtual machine
101 (VM) configurations and profiles for different classes
102 of mobile device. These VM configurations and pro-
103 files support different subsets of APIs that application
104 developers can use to build applications. A device
105 with more capable HW (e.g., more memory and faster
106 processors) can typically support a VM with a
107 more extensive Java class library. If an application
108 uses APIs provided by a Java profile or configuration
109 of a more capable platform, this application may not
110 be able to run on less capable platforms that do not
111 support these APIs.

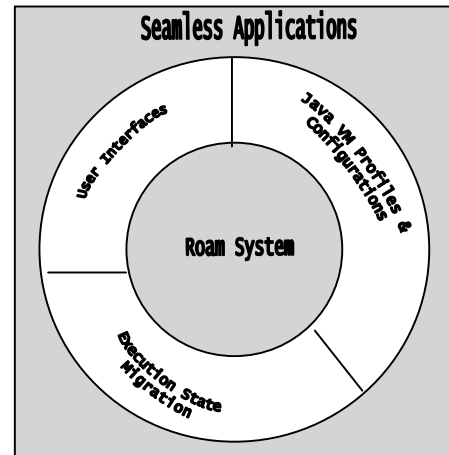


Fig. 1. Challenges for application migration in heterogeneous device environment.

- *Execution state migration*: In a runtime application migration from a source device to a target device, the execution state of an application needs to be captured on the source device, transferred to a target device and restored. The application execution state includes heap, stack, network sockets, file I/O state, and other state information. Device heterogeneity brings an additional challenge in execution state migration. Given the incompatible configurations and profiles on different device platforms, application developers may choose to create multiple device-dependent implementations (M-DD) for an application component (e.g., one implementation for each specific device platform), or they can use an automated tool that can transform an application component into code that can be run on each platform. This means that, if an application has two different implementations (either hand-coded or transformed) on two different platforms, the captured runtime execution state on the source device may not correspond to the different implementation on the target device.
- *User interface*: Different platforms have different display sizes and input methods. A DoCoMo 503i cell phone has an approximately 120 × 130 pixel display and a small numeric keypad; a Compaq iPaq PocketPC has a 320 × 240 pixel display and a stylus; and an average notebook PC has a 1024 × 768 pixel display, a keyboard and a pointing device. Display size can change the presentation and layout of the GUI components. A large display (e.g., desktop PC) can accommodate big GUI components, and many GUI components can be presented at once in the same window or page. On the other hand, a small display can only accommodate highly compact GUI components, and only a few GUI components can be displayed at once. In addition, the constraints of the devices' display sizes and input methods can affect

149 the range of tasks that are appropriate on each plat-
 150 form. For example, a notebook PC or a Blackberry
 151 pager with a keyboard is suitable for performing
 152 tasks that involve text entry. On the other hand, a cell
 153 phone with only a numeric keypad is prohibitively
 154 difficult to use for tasks involving text entry.

155 1.2. General approach

156 To address the challenges described in the previous
 157 section, the Roam system provides a seamless applica-
 158 tion framework for developers to build *resource-aware*
 159 seamless applications. By resource-aware, we mean that
 160 applications are aware of the underlying device capa-
 161 bilities (Java VM configuration and profile, display size,
 162 input method, memory, network, etc.), which may
 163 change when they migrate between devices. In order to
 164 migrate applications between devices with different ca-
 165 pabilities, *adaptation* is needed. The Roam system uti-
 166 lizes the following three adaptation strategies:

- 167 • *Dynamic instantiation*: An application is divided into
 168 M-DD components that perform the same function,
 169 but each component is implemented/re-implemented
 170 for different platforms. At migration time or runtime,
 171 the set of components that best fit the target device
 172 capabilities is selected for instantiation. For example,
 173 the Java VM configuration and profile supported on
 174 the target device is considered as a part of the target
 175 device capabilities in the Roam system. A developer
 176 can provide two device-dependent implementations
 177 of the same component, one for cell phones that sup-
 178 port the J2ME DoCoMo Java (DoJa) profile, and an-
 179 other for notebook PCs that support the Java 2
 180 Standard Edition (J2SE) VM. When an application
 181 is started on or migrated to a cell phone, the Roam
 182 system *dynamically instantiates* the cell phone imple-
 183 mentation.
- 184 • *Offloading computation*: The Roam system allows ap-
 185 plications to make use of distributed computing re-
 186 sources to run software components that are
 187 beyond the target device capabilities. These compo-
 188 nents may require a more capable Java VM, or have
 189 certain CPU, memory, or network requirements, such
 190 that it is necessary to dispatch them to remote servers
 191 that can better support these requirements. Consider
 192 an application that has a component implemented to
 193 run only on the J2SE VM. When this application is

started on or migrated to a cell phone, the Roam sys- 194
 tem finds a remote server capable of running the J2SE 195
 VM and offloads that software component to that re- 196
 mote server. 197

- *Transformation*: The Roam system allows device-in- 198
 dependent components to be transformed at runtime 199
 to fit the target device capabilities. It provides a de- 200
 vice-independent GUI toolkit that a developer can 201
 use to build user interface (UI) components. At mi- 202
 gration time or runtime, these device-independent 203
 UI components are transformed to run on the target 204
 device. 205

The Roam system is a framework to assist developers 206
 in building applications that can adapt to heterogeneous 207
 devices. At design time, a developer must make the de- 208
 cision on how to partition an application into separate 209
 components, and for each component, whether to pro- 210
 vide *multiple device-dependent* implementations (M- 211
 DD), a *single device-dependent* implementation (S-DD), 212
 or a *single device-independent* presentation (S-DI). At 213
 runtime or migration time, the Roam system applies the 214
 most appropriate adaptation strategy, shown in Table 1, 215
 based on the developer's design decision and the target 216
 device capabilities. 217

We believe that some adaptation strategies are better 218
 suited for certain types of application components. For 219
 some UI components, we believe that transformation 220
 adaptation may be the most appropriate, given that 221
 most UI components cannot be offloaded to a remote 222
 server for execution. However, device-independent UI 223
 authoring sometimes does not result in high-quality in- 224
 terfaces. In these cases, M-DD UI components are de- 225
 veloped, and dynamic instantiation adaptation is 226
 preferred. This decision is up to the UI developer. For 227
 application logic components, we believe that S-DD 228
 implementation, using offloading computation adapta- 229
 tion is often most suitable, because most application 230
 logic components can be offloaded to remote servers for 231
 execution. 232

In order to suspend and resume the execution of 233
 mobile applications across heterogeneous devices, the 234
 Roam system must be able to capture execution state on 235
 the source device, serialize and transfer it to the target 236
 device, and then restore it. Because M-DD components 237
 may be dynamically instantiated, and the implementa- 238
 tion of each is different, execution state must be cap- 239
 tured using a *device-independent state representation*, 240

Table 1
 Summary of adaptation strategies

Adaptation strategies	Application components	Intended for
Dynamic instantiation	Multiple device-dependent implementations (M-DD)	Customized, high-quality UI components
Offloading computation	Single device-dependent implementation (S-DD)	Application logic components
Transformation	Single device-independent presentation (S-DI)	UI components

241 and converted into representations specific to each
242 component implementation using *execution state trans-*
243 *formation*. The Roam system requires developers to
244 provide execution state transformation logic for M-DD
245 components. There is no need for a developer to provide
246 execution state transformations for S-DI or S-DD
247 components that have a single implementation. When
248 Roam system applies transformation adaptation to an
249 S-DI component, it automatically generates the corre-
250 sponding execution state transformation.

251 Another important requirement is *Security*. Proper
252 security policy must be enforced to ensure that only
253 authorized applications are allowed to migrate from one
254 device to another.

255 1.3. Organization

256 We organize the remainder of the paper as follows:
257 Section 2 presents the design of the Roam system, Sec-
258 tion 3 describes its implementation, APIs and a sample
259 application, Section 4 presents related work, and Section
260 5 draws conclusions.

261 2. Design

262 The Roam system architecture is shown in Fig. 2. A
263 *Roamlet* is an application that runs in the Roam system.
264 A Roamlet can migrate between any two connected
265 devices that have the Roam system running. The arch-
266 itecture contains three main components: *Roam*
267 *Agent*, *Roamlet*, and *HTTP Server*. Roam agents must
268 be installed and executed on both the source and target
269 devices before a Roamlet can run, and before any
270 Roamlet migration can occur. The flow for a Roamlet
271 migration from a PC device to a PDA device is described
272 as follows:

- 273 1. The Roam agent on the source device first negotiates
274 with the Roam agent on the target device. The nego-
275 tiation involves exchanges of the target device capa-
276 bilities, the device capabilities needed by each
277 application component, and the code base URL
278 where the Roamlet component byte code can be
279 downloaded from. Based on the exchanged informa-
280 tion, the Roam agent decides the appropriate adapta-
281 tion strategy for each component.

- 282 2. The Roam agent on the target device downloads the
283 necessary Roamlet class byte code from the HTTP
284 server for all application components that will be in-
285 stantiated on the target device.
286 3. The Roamlet on the source device serializes its execu-
287 tion state and sends it to the Roam agent on the tar-
288 get device. The Roam agent may perform execution
289 state transformation if an application component is
290 transformed or dynamically instantiated.
291 4. The Roam agent instantiates the Roamlet on the tar-
292 get device.

293 Roamlets are based on a *component-based program-*
294 *ming model*: they are built as components that can be
295 distributed across multiple devices, and can be com-
296 posed together into a working application through well-
297 defined interfaces. Fig. 3 shows an example of a
298 Roamlet partitioned into three components: a GUI
299 component with device-independent representation (S-
300 DI), a second GUI component with two device-depen-
301 dent implementations on PC and PDA (M-DD), and an
302 application logic component with a PC implementation
303 (S-DD). The component-based programming model al-
304 lows the Roam system to apply a different adaptation
305 strategy on different application component, such as
306 offloading an S-DD component to a remote device,
307 dynamically instantiating an M-DD component of a
308 different implementation, and transforming an S-DI
309 component without affecting other components in a
310 Roamlet. This requires each component to invoke pro-
311 cedure calls of any other components regardless of
312 whether they are running on remote or local hosts. In
313 our current Roam system, the component-based pro-
314 gramming model is built on top of Java Remote Method
315 Invocation (RMI).

316 Roamlets and Roam agents use *resource specifica-*
317 *tions* to describe the device requirements for each com-
318 ponent implementation, and the capabilities of a host
319 device, respectively. At migration or load time, the
320 Roam agent can compare the device requirements from
321 the application components with the target device ca-
322 pabilities and decide the best adaptation strategy for
323 each application component. Note that components in a
324 Roamlet may have different device requirements, so a
325 Roamlet is required to provide a separate set of device
326 requirements for each component implementation. The
327 current Roam system supports specification of a limited
328 set of device capabilities, including Java VM configura-
329 tion or profile, display size and input method, as
330 shown in Table 2. For example, device capabilities for a
331 PC are {J2SE VM, 1024 × 780 pixels, keyboard and
332 mouse}. The device requirements for the PDA imple-
333 mentation of the GUI component shown in Fig. 3 can
334 be {Personal Java VM, 320 × 240 pixels, stylus}, and the
335 PC implementation of the GUI component can be
336 {J2SE VM, 1024 × 780 pixels, keyboard and mouse}.




Fig. 2. Roam system architecture.

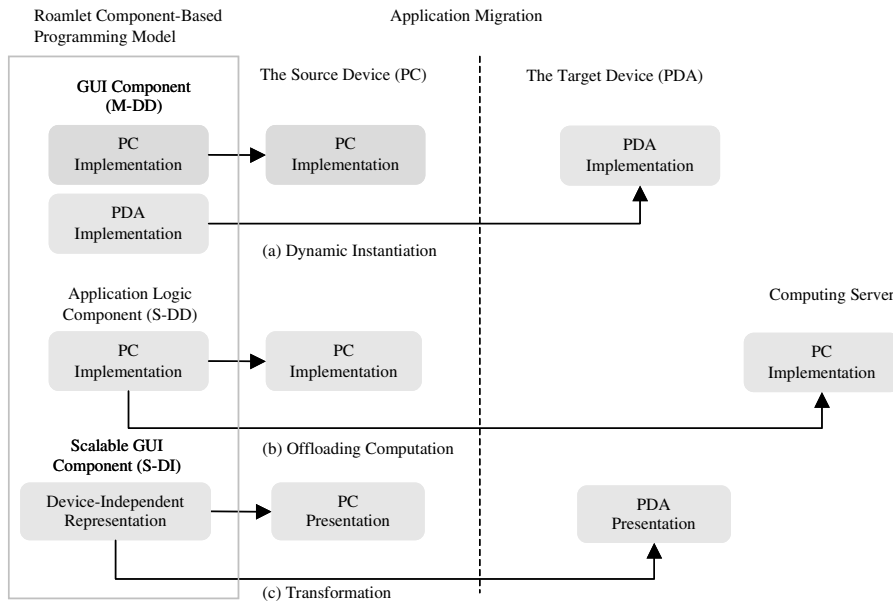


Fig. 3. Roamlet component-based programming model and different adaptation strategies for Roamlet components.

Table 2
Specification for device capabilities

	Types of Java VM	Display sizes	Input methods
PC device capability	J2SE	1024 × 780	Mouse and keyboard
Pocket PC device capability	PersonalJava	320 × 240	Stylus and virtual keyboard
Cell phone device capability	J2ME/DoJa profile	100 × 100	Keypad

337 Since the scalable GUI (SGUI) component has a device-
338 independent representation, its device requirements can
339 contain a range of capabilities.

340 2.1. Dynamic instantiation

341 In dynamic instantiation, a developer provides M-
342 DD implementations for a Roamlet component. The
343 rule for selecting the most suitable device-dependent
344 implementation is that the component implementation
345 that best matches (but not exceeds) the target device
346 capability is selected. An example of a Roamlet mi-
347 grating from a PC to a PDA is shown in Fig. 3(a). The
348 Roam agent finds that the PDA implementation of the
349 second GUI component exactly matches the target de-
350 vice (PDA) capability and instantiates it.

351 Using dynamic instantiation, Roamlet developers can
352 program different Roamlet behaviors in separate code
353 segments. At migration time the Roam system loads
354 only the code segment that exhibits a behavior suited for
355 the target device capability. Note that dynamic instan-
356 tiation is different from programming different behav-
357 iors in the same code segment, which is not always
358 desirable in Java. For this approach to work, the
359 Roamlet has to be programmed to run on the least ca-
360 pable device (the *lowest-common denominator* approach)

because the code segment that runs on different VM
361 configurations and devices can only assume the most
362 primitive API libraries provided by the least capable
363 VM and devices. This is obviously undesirable for ap-
364 plication programmers who would like to take advan-
365 tage of rich class libraries provided by more capable
366 devices.

367 The Connect4 Roamlet is a multi-platform game that
368 we implemented to illustrate dynamic instantiation ad-
369 aptation. Its GUI component provides M-DD presen-
370 tations, one matching the PDA device capability and
371 one matching the PC device capability. When Connect4
372 is instantiated on the PC, its PC presentation is selected
373 for instantiation (Fig. 4(a)). When the user later mi-
374 grates the Connect4 game to a PDA, its PDA presen-
375 tation is selected for instantiation (Fig. 4(b)). The PDA
376 presentation is a scaled down version of the PC pre-
377 sentation with some features and buttons removed to fit
378 the smaller PDA display size (such as “undo” and “new
379 game”).

380 Dynamic instantiation creates a challenging situa-
381 tion, where a Roamlet is instantiated with a different
382 implementation on the source device from that on the
383 target device, as shown in Fig. 3(a). The execution state
384 of the PC GUI component on the source device does not
385 correspond to the PDA GUI component instantiated on
386

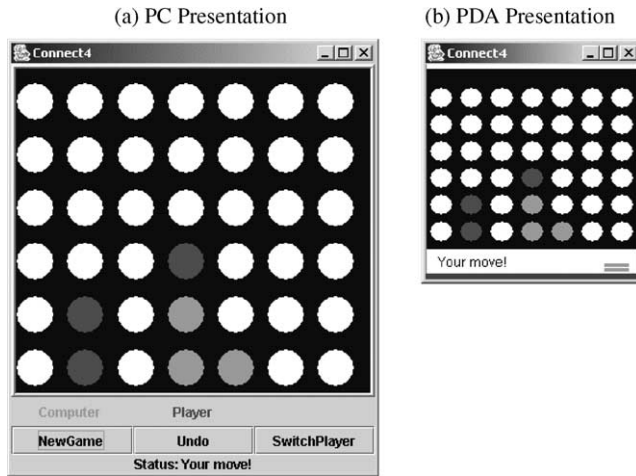


Fig. 4. Screenshots of M-DD presentations of Connect4 Roamlet: one presentation for PC and one presentation for PDA.

387 the target device. For example, the PC GUI component
388 might use radio buttons to present a choice, while the
389 PDA GUI component might use a drop-down list instead.
390 In order to migrate the execution state of the PC
391 GUI component to the PDA GUI component, execu-
392 tion state must be transformed before it can be restored
393 on the target device (e.g., mapping the value of the radio
394 button to the drop-down list). This state transformation
395 logic is provided by developers.

396 2.2. Offloading computation

397 Offloading computation allows a Roamlet to delegate
398 the execution of computation- and memory-intensive
399 components to any remote server that has the device
400 capabilities to run them. Offloading computation is in-
401 tended for an application component that has a S-DD
402 implementation. Fig. 3(b) shows offloading computation
403 for an S-DD application logic component. After the
404 Roam agent finds that the application logic component
405 requires a PC-equivalent device, which exceeds the PDA
406 host's capabilities, it attempts to find a computing server
407 (PC) where it can offload this application logic compo-
408 nent. Both the Roamlet users and the Roam agent can
409 specify a list of servers that can accept offloadable
410 components. When a server is found, the Roam agent
411 directs the application logic component to be migrated
412 to that server. A Roam agent can apply offloading
413 computation to an M-DD component if none of its
414 device-dependent implementations can run on the target
415 device. For example, if an M-DD component has two
416 implementations on PC and PDA and the target device
417 is a cell phone, dynamic instantiation will fail. In this
418 case, the Roam agent will find the most capable server
419 on the list, and offload an appropriate M-DD compo-
420 nent to it. The precedence rule for applying different
421 adaptation strategies is that a Roam agent first applies

dynamic instantiation adaptation to an application 422
component. If the dynamic instantiation adaptation 423
fails, it will apply offloading computation adaptation. 424

425 The Connect4 game in Fig. 4 contains a S-DD ap-
426 plication logic component that is offloadable. The ap-
427 plication logic component keeps track of the state of the
428 game and implements an artificial intelligence (AI) en-
429 gine that computes the computer's next move. Since the
430 AI could be computationally- and memory-intensive, it
431 requires a device with a fast processor and sufficient
432 memory. When the Connect4 game is running on the
433 PC, the application logic component is instantiated lo-
434 cally. When the Connect4 game is migrated to the PDA,
435 the application logic component is offloaded to a com-
436 puting server.

437 Not all components in a Roamlet are offloadable.
438 The Roam system allows a Roamlet to specify its
439 components as either non-offloadable or offloadable
440 components. Non-offloadable components are required
441 to run on the target device. For example, GUI compo-
442 nents and security-sensitive components are rarely off-
443 loaded. If the target device does not have the capability
444 to run the non-offloadable components, the application
445 migration simply fails and an error is presented to the
446 user. A Roamlet can also specify components as *reverse-*
447 *offloadable* or *non-reverse-offloadable*. Consider the same
448 example in Fig. 3(b). At a later time, the user wants to
449 migrate the Roamlet back to the PC from the PDA. If
450 the application logic component is reverse-offloadable, it
451 will be pulled back from the computing server to the
452 target PC device; otherwise, the application logic com-
453 ponent will stay on the computing server. It is a trade-off
454 between paying a one-time transfer cost of migrating the
455 application logic component (in the case of reverse-off-
456 loadable) vs. the continuous cost of communicating with
457 that remote component in a Roamlet (in the case of non-
458 reverse-offloadable).

459 2.3. Transformation

460 Transformation is a runtime UI generation process
461 from a device-independent representation constructed
462 by developers at design time to device-dependent pre-
463 sentations. Fig. 3(c) shows transformation of a device-
464 independent UI component into a PDA presentation
465 when the Roamlet migrates to a PDA. To build a de-
466 vice-independent representation and to generate device-
467 dependent presentations, we provide a device-independ-
468 ent GUI library and a runtime transformation tool,
469 which we called the SGUI toolkit.

470 The design of the SGUI toolkit is shown in Fig. 5. It
471 contains three modules: *SGUI library*, *transformation*
472 *manager*, and *render manager*. The process flow goes as
473 follows. At design time, UI designers use widgets pro-
474 vided by the SGUI library to construct device inde-
475 pendent UI presentations. At runtime, the




Fig. 5. SGUI toolkit.

476 transformation manager transforms a device-indepen-
477 dent representation into a device-dependent presenta-
478 tion that satisfies the target device constraints (display
479 size and the available input methods). The render
480 manager displays the device- presentation on the target
481 device.

482 2.3.1. Shop hunter application

483 “Shop hunter” is a hypothetical multi-platform ap-
484 plication that helps a user to buy items anytime, any-
485 where. We will use the shop hunter application to
486 illustrate transformation adaptation supported by the
487 SGUI toolkit. The shop hunter application allows a user
488 to input new shopping items, to view a list of shopping
489 items, and to search for directions to nearby stores with

490 the cheapest prices. We have prototyped a device-inde-
491 pendent representation of the shop hunter UI using our
492 SGUI toolkit, and applied transformation to generate
493 device-specific presentations for notebook PCs, PDAs,
494 and cell phones which are shown in Fig. 6.

495 The first noticeable difference between different de-
496 vice-specific presentations is how the device’s display-
497 able screen size changes the *layout* and *pagination*. The
498 notebook PC presentation fits on one page, the PDA
499 presentation is laid out on three different pages, and the
500 cell phone presentation is broken up into six pages. The
501 second noticeable difference between the screenshots is
502 that the “sort” UI components are different between the
503 notebook/Pocket PC presentation and the cell phone
504 presentation. The notebook and Pocket PC presenta-
505 tions use a radio button group and the cell phone pre-
506 sentation uses a more compact drop-down list. The third
507 noticeable difference is that the UI components corre-
508 sponding to “add shopping item” is missing on the cell
509 phone presentation, because this task involves text en-
510 tries and is therefore inconvenient on a cell phone.

511 2.3.2. SGUI library

512 The SGUI library contains two modules: the *virtual*
513 *widget library* and the *SGUI events*. The virtual widget

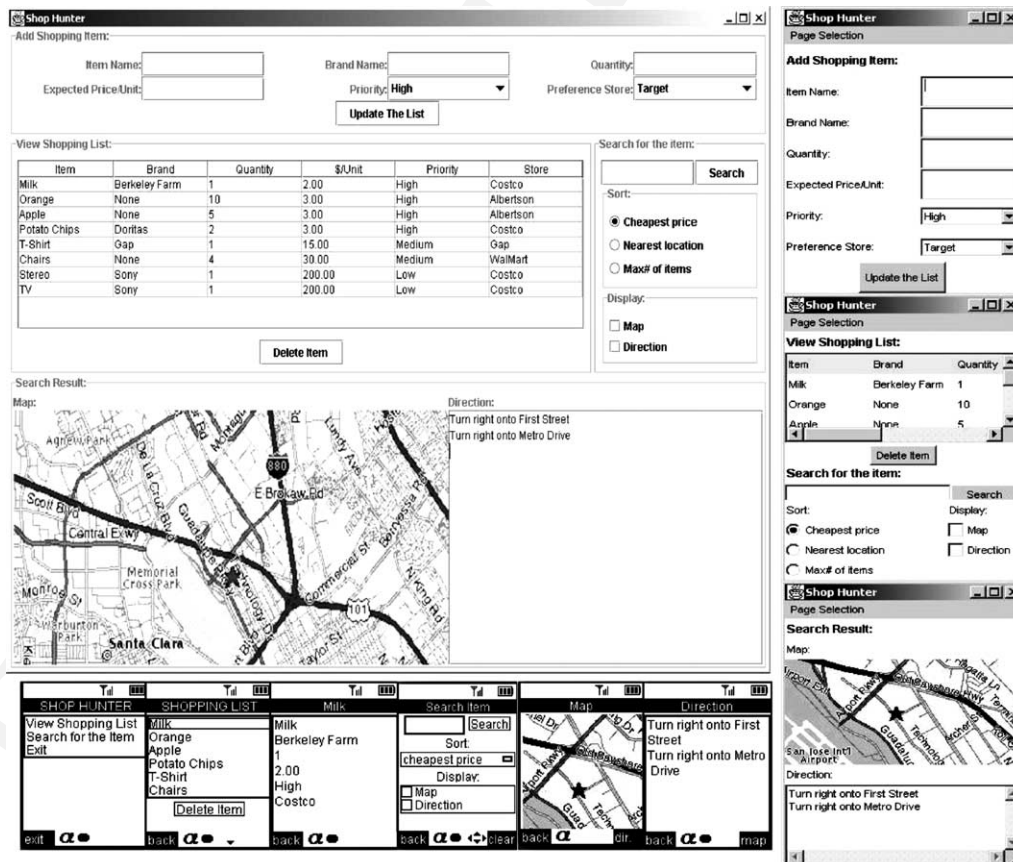


Fig. 6. Shop hunter application.

514 library contains a comprehensive set of UI widgets
515 similar to Java's Swing library. It includes widgets such
516 as labels, buttons, menus, text fields, scroll bars, tables,
517 etc. The virtual widget library is implemented for each
518 Roam-supported device platform. Since the virtual
519 widget library is based on Swing, its APIs look very
520 similar to Swing APIs.

521 The SGUI events are abstractions of user interaction
522 events such as mouse, keypad, or stylus input. Since
523 different device platforms support different input meth-
524 ods, the SGUI library describes mappings between de-
525 vice-dependent events and abstract device-independent
526 events. This abstraction allows UI developers to handle
527 these events regardless of which platform a UI is run-
528 ning on. For example, an abstract *action* event associ-
529 ated with a button press can be generated from a mouse
530 click on a PC, a tap from stylus on a Pocket PC, or a
531 select-key press on a cell phone. The event is realized as
532 a device-dependent event at runtime, and then delivered
533 to the application.

534 For some applications, mapping between device-de-
535 pendent GUI events to device-independent events may
536 not always be possible due to platform constraints. For
537 example, an application may contain an interactive map
538 that can zoom in or out when a user clicks on a specific
539 point in the map image. On a PC, a mouse event can be
540 used to capture the specific location of the mouse click.
541 However, a cell phone lacks a pointing input method; as
542 a result, it is not possible to map an equivalent event on
543 a cell phone. To address the need for device-dependent
544 events, the SGUI library also retains a set of events for
545 device-dependent input methods. This allows a scalable
546 application to enhance its UI on a particular device
547 platform. Since device-dependent events are generated
548 only on some specific device platforms, they should not

549 be used to implement any core features of UI that are
550 expected to be available on any devices.

2.3.3. Device-independent representation

551 UI designers specify two kinds of information in the
552 device-independent representation—*task model* and
553 layout. Fig. 7 depicts the top portion of the device-in-
554 dependent representation for the shop hunter applica-
555 tion. The device-independent representation has a tree-
556 like structure. The child nodes of the root node are task
557 nodes representing different tasks. For the shop hunter
558 application, there are three tasks: add shopping item,
559 view items, and search item. A task node can have sub-
560 task nodes, and so forth. The add shopping item task
561 contains a scalable button called update list, and logical
562 panel A. The button is a scalable UI component instan-
563 tiated from the virtual widget library. A logical
564 panel is a presentation unit that groups related child
565 logical panels and scalable UI widgets. A spatial layout
566 can be specified on the relative positions of widgets and
567 panels contained in logical panels. For the add shopping
568 item task node, UI designers would specify that logical
569 panel A be placed at the top, and update list button at
570 the bottom. Logical panel A contains six smaller logical
571 panels A1 through A6 arranged in a 2 × 3 grid.

572 UI designers can specify *task preferences* on task
573 nodes. The task preference tells which platforms with
574 which input methods are suitable to display this task. In
575 the shop hunter application, a UI designer would specify
576 the Add Shopping Item task to be suitable for displaying
577 on notebook PCs or PDAs with virtual keyboards, but
578 not for cell phones. As a result, UI components corre-
579 sponding to Add Shopping Item task will not be dis-
580 played on cell phones as shown in Fig. 6.

581 In addition to layout, UI designers can specify other
582 layout properties on logical panels, such as *layout pri-*
583

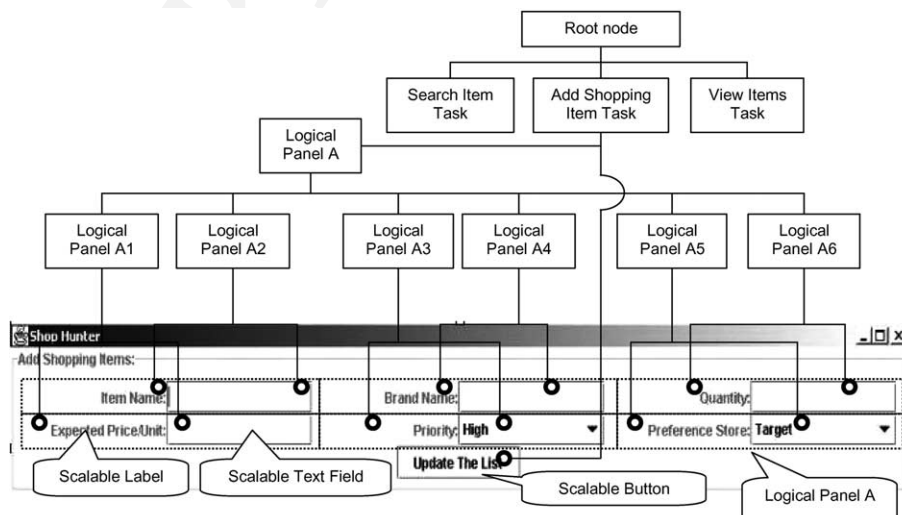


Fig. 7. Device-independent representation for the top portion of shop hunter application.

584 *ority, split-ability* and *title*. Layout priority decides the
585 order in which logical panels will be placed on pages.
586 Split-ability decides if descendant nodes of a logical
587 panel can be placed on separate pages. By default, all
588 logical panels are not splittable. Titles of logical panels
589 are used to create a navigation menu bar for moving
590 between pages. For example, in the Pocket PC presen-
591 tation of the shop hunter application of Fig. 6, the
592 navigation menu bar Page Selection contains the titles of
593 logical panels as menus.

594 Integration of the layout structure and the task model
595 in the device-independent representation tree poses a
596 limitation that the task model hierarchy must not con-
597 flict with the layout structure hierarchy. It is not possible
598 to specify a task that places its widgets in the same
599 logical panel as widgets from another task. For example,
600 the view items task cannot place a widget under Logical
601 Panel A1 because it is also under Add Shopping Item
602 task. To address this limitation, we are looking into
603 separating the task model from the layout structure into
604 two presentation trees similar to Eisenstein et al. (2001).

605 Note that UI designers must construct device-inde-
606 pendent representations for both static and dynamic
607 content. In the following sections, we explain how the
608 transformation manager transforms a device-indepen-
609 dent representation into a device-specific presentation.

610 2.3.4. Transformation manager

611 The design of the transformation manager is shown
612 in Fig. 8. It contains three sub-modules: *task manager*,
613 *layout algorithm* and *widget transformation*. The first
614 step in transformation manager involves the task man-
615 ager. Its job is to choose the appropriate tasks for the
616 target device platform, and to remove widgets belonging
617 to tasks not appropriate for the target device platform.
618 For example, if UI designers specify that the Add
619 Shopping Item task in the shop hunter application is not
620 suitable for cell phones; widgets corresponding to the
621 Add Shopping Item task are removed by the task man-
622 ager and not displayed on the cell phone presentation.

623 The second step in the transformation manager in-
624 volves the layout algorithm working with widget trans-
625 formation. Its job is to paginate the presentation into
626 separate pages according to the layout specified by UI
627 designers and also to meet the constraint that a page
628 cannot exceed the target device's displayable screen size.




Fig. 8. Transformation manager design.

2.3.5. Layout algorithm

629 We define two requirements for our layout algorithm. 630
631 First, the generated presentation should have reasonably 632
633 high quality without requiring significant device-specific 634
635 customization from UI developers. One implication of 636
637 high quality and the display size constraint is that the 638
639 generated presentation must fit the target display with- 640
641 out scrolling, because scrolling generally degrades us- 642
643 ability. Second, the algorithm must be fast, so that it 644
645 does not cause significant presentation generation delay. 646
647 There have been many methods on how to layout wid- 648
649 gets for different display sizes. Some of these methods, 650
651 such as the one proposed by Masui (1994), are highly 652
653 computationally complex. Some methods require too 654
655 much information from developers: for example, Hu- 656
657 manoid (Szekely et al., 1992) asks developers many 658
659 layout-related questions before it can generate a final 660
661 presentation. Among all proposed methods, T_{EX} (Lin- 662
663 ton, 1989) is the most promising method in formatting 664
665 two dimensional box-like GUI widgets (Masui, 1994; 666
667 Olsen et al., 2000). T_{EX} allows each widget to report its 668
669 desired size for positioning. 670

671 Fortunately, Java has default layout managers that 672
673 are similar to T_{EX}. Java also allows UI developers to 674
675 specify the desired location of each widget through a set 676
677 of predefined layout constraints. In order to meet our 678
679 first requirement, we propose to have only one set of 680
681 layout specification from UI developer, and this set of 682
683 layout specification is used to generate layouts for other 684
685 platforms. The specification is the same as the Grid Bag 686
687 Layout Constraint in Java. We recommend UI devel- 688
689 opers to specify the layout according to the presentation 690
691 on the largest device display, such as a presentation on a 692
693 PC, for reasons detailed in Section 3.6. The layout al- 694
695 gorithm will try to follow that specification as closely as 696
697 possible when it is creating layouts for other platforms 698
699 with smaller display sizes. 700

701 The general strategy for the layout algorithm is to 702
703 start from the lowest nodes in the device-independent 704
705 representation tree, and then work its way up to the root 706
707 node. When it encounters a non-splittable logical panel 708
709 that cannot fit the display size using the specified grid 710
711 bag layout, the layout algorithm applies *flow layout*. We 712
713 choose flow layout because it is simple, and it involves 714
715 minimum computation while maintaining a reasonable 716
717 presentation (Vanderdonckt, 1995). Since using flow 718
719 layout may not generate a high quality layout, it is *only* 720
721 applied when the specified grid bag layout fails. If flow 722
723 layout also fails on a non-splittable logical panel, the 724
725 layout algorithm may apply widget transformation to 726
727 find smaller alternative presentations. The detailed steps 728
729 of the layout algorithm are described as follows: 730

- 731 1. The algorithm starts by finding the *lowest-level un-* 732
733 *splittable node* with the highest layout priority in the 734
735 presentation tree as shown in Fig. 9. A lowest-level 736

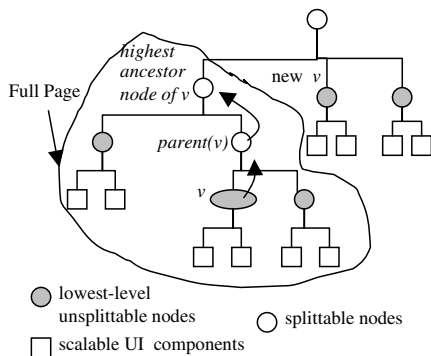


Fig. 9. Layout algorithm.

684 unsplittable node is defined as a node whose child
685 nodes are all widgets (or leaf nodes), and the widgets
686 have to be placed on the same page. The goal is to try
687 to place widgets that are assigned higher layout priori-
688 ties to the first few pages, so that they are easier for
689 users to locate. Denote v as the chosen node.

690 2. A default *style* corresponding to the target device
691 platform is applied to the set of widgets under node
692 v . The style guide sets common properties of widgets
693 such as font size, spacing between components, etc.
694 The purpose of the style is to have a consistent look
695 across different widgets. The set of widgets under
696 node v are laid out on a page according to the grid
697 bag layout constraints specified by the UI developer.
698 The precise size of the page is then calculated.

699 3. If the size of page is bigger than the device display
700 size, the page is *over-filled*. We apply flow layout on
701 child nodes under v , and re-compute the page size.
702 If the page is still *over-filled*, we check if the node is
703 splittable or not. If the node is splittable, we allocate
704 a new page to place extra widgets that cannot fit on
705 the previous page. If the node is non-splittable, wid-
706 get transformation is applied to as many child nodes
707 under v as needed such that the flow layout can fit
708 them on one page. If widget transformation also fails,
709 the layout algorithm will leave some widgets out of
710 the page and notify its users. Details on widget trans-
711 formation are discussed in Section 2.3.7.

712 4. If the size of the page is smaller than the size of the
713 device screen size, the page is *under-filled*. This means
714 that the algorithm can accommodate additional wid-
715 gets from sibling nodes of node v . The algorithm pro-
716 ceeds to find the *highest ancestor node* of node v in the
717 presentation tree as shown in Fig. 9, such that using
718 that ancestor node can pack the maximum number
719 of widgets in a page without violating the display size.
720 Denote the parent node of node v as $PARENT(v)$.
721 The algorithm tries to place the set of widgets under
722 $PARENT(v)$ on a page. This is done by repeating step
723 2, i.e., applying the current style and the grid bag lay-
724 out specification to the widgets under $PARENT(v)$. If

the resulting page is still under-filled, it tries $PAR-$
 $ENT(PARENT(v))$, and so forth until the resulting
page becomes full.

5. If the page is *full*, a new page is allocated for placing remaining widgets. The algorithm first finds another lowest-level unsplittable node with the next highest task priority from the presentation tree that hasn't been placed on any pages. This node becomes the *new node v*, and the algorithm goes to step 2.
6. When all nodes are processed—that is, placed on a page—the algorithm terminates.

For the shop hunter application in Fig. 6, the table showing the shopping list has different device-specific presentations on Pocket PCs and cell phones. For the cell phone presentation, the table is transformed into a selection list using item names as indices. The map image is scaled down on the Pocket PC presentation and on the cell phone presentation. Techniques involving scaling images and graphic objects (W3C, 2000) are outside the scope of this paper.

2.3.6. Page navigation

While the layout algorithm is automatically generating a new layout, it is also constructing a simple navigation menu bar for moving between pages. If the application already has a menu bar, the algorithm will simply add a customizable menu named “Page Selection” by default as the first menu from the left on the menu bar. If the application does not have any menu bar, a new one is created. On the Pocket PC, the menu bar always appears at the top of each page under the Page Selection menu bar as shown in Fig. 6. On the cell phone, the first page is always the top-level task navigation menu. For each generated page, a menu item is created and it will be automatically added as the item of the “Page Selection Menu.” The name for each menu item is from the title property of the occupying logical panel.

2.3.7. Widget transformation

Widget transformation is the transformation from one widget, either primitive or composite, to another that consumes less space in the graphical layout. Transformation is guided by *transformation rules*, similar to rules in a rule-based system. Widget transformation rules are selectively applied to yield a presentation which is optimally usable given display size constraints. For example, Fig. 10 shows the effect of a transformation rule mapping a large, composite widget containing pairs of property labels and value textfields, into a smaller composite widget containing a property list box and a single value textfield. The transformed composite widget is smaller and remains easy to use.




Fig. 10. A multiple-to-multiple transformation.

776 At the core of widget transformation is an algorithm
777 to determine which widgets to transform, and which
778 transformation rules to apply to them.

779 To determine the widgets that are more acceptable to
780 transform, developers specify which widgets are *core*—
781 representing important application functionality, and
782 which are *optional*—less important for use of the ap-
783 plication. These hints are provided for primitive widgets,
784 as well as those that are combined into composite wid-
785 gets. We prefer to transform optional widgets rather
786 than core widgets whenever possible, so that we degrade
787 usability as little as possible (Vanderdonckt, 1995); only
788 when there are no optional widgets, or when the layout
789 size reduction is insufficient to meet display size con-
790 straints, are core widgets transformed.

791 Determining the transformation rules to apply to the
792 widgets is more complex. There are four types of
793 transformation rule:

- 794 • *One-to-one*: transforms from a single primitive widget
795 to another single widget, e.g., a rule transforming a
796 list into a drop-down box.
- 797 • *One-to-multiple*: transforms from a single widget into
798 multiple widgets, e.g., a rule transforming a single ta-
799 ble into multiple lists or drop-down boxes.
- 800 • *Multiple-to-one*: transforms from multiple widgets of
801 the same type into a single widget, e.g., a rule trans-
802 forming a group of radio buttons or checkboxes into
803 a single drop-down list box, and a rule transforming a
804 group of one-line textfields into a single multi-line
805 textarea.
- 806 • *Multiple-to-multiple*: transforms from multiple wid-
807 gets of different types into other multiple widgets,
808 e.g., a rule transforming a group of label and textfield
809 pairs into a drop-down list box and a single textfield,
810 as shown in Fig. 10. Note that a multiple-to-multiple
811 rule is composed of a set of multiple-to-one rules, and
812 a set of classic relationships described in Vander-
813 donckt (1995).

814 Each of these classes of transformation rule has a
815 priority assigned to it, corresponding to the demon-
816 strated usability of widgets transformed by rules in the
817 class. That is, we prefer transformation rules that tend
818 to result in easier-to-use widgets over transformation
819 rules that tend to result in harder-to-use widgets. One-

to-one rules are preferred first, then one-to-multiple, 820
multiple-to-one, and finally multiple-to-multiple. Note 821
that multiple-to-multiple rules can change the UI dras- 822
tically, and degrade usability, compared to one-to-one 823
rules. Conversely, note that one-to-one rules tend to 824
offer less space savings than multiple-to-multiple rules. 825
Our prioritization balances this tradeoff: we want as 826
large—and consequently as usable—a UI as possible 827
given display size constraints, but we accept smaller and 828
less usable UIs as necessary. 829

We limit the set of transformation rules that may be 830
applied by eliminating rules whose original widgets re- 831
quire input methods that are unsupported by the 832
transformed widgets on the target device. For example, 833
if the original UI contains a J2ME DoJa button that has 834
an action triggered on a mouse-in event, any rules 835
transforming that button to a softkey are not applicable, 836
because softkeys do not support any equivalent to 837
mouse-in events. The current SGUI toolkit provides a 838
set of commonly used transformation rules. We also 839
allow developers to add their own transformation rules. 840

2.3.8. SGUI state and event transformation 841

When an application is migrated from a device of one 842
platform to another device of a different platform, the 843
presentation on the target device may use a different set 844
of widgets from the source device, so there is a need to 845
map running state between the source presentation and 846
the target presentation. This is done by state synchroni- 847
zation and restoration between the running state of a 848
device-specific widget and its corresponding device-in- 849
dependent widget. These two mechanisms are supported 850
in the SGUI library to ensure that the device-indepen- 851
dent representation always has the most recent widget 852
state right before and right after any migration. Prior to 853
any application migration, the running state of a device- 854
specific widget is synchronized with its corresponding 855
device-independent widget. After an application migra- 856
tion, the running state of the device-independent widget 857
is restored to the device-specific widget. 858

Event transformation is also needed to support the 859
same UI interaction across different device-dependent 860
presentations. This is done by a bi-directional mapping 861
mechanism between device-dependent events generated 862
from device-dependent widgets and their corresponding 863
device-independent events generated by a device-inde- 864
pendent widget. This mapping mechanism is supported 865
in our SGUI library. 866

2.4. Security 867

We assume that Roam agents have been installed 868
securely and they are trusted entities. We also assume 869
that a user would only migrate applications among his/ 870
her personal devices (not public and not shared) such as 871
cell phones, PDAs, and home PCs. This means that a 872

```
public abstract class Roamlet implements Serializable {  
    public Roamlet(Object[] args);  
    public synchronized final void migrate(String hostname);  
    public synchronized boolean onInitialization();  
    public synchronized boolean onRemoval();  
    public synchronized boolean onArrival();  
    public synchronized boolean exit();  
    static final void instantiate(String className, Object[] initArgs, String codebase);  
};
```

Fig. 11. Roamlet runtime migration API.

873 user should know when to expect an application mi-
874 gration on which device and from which device.

875 Based on these assumptions, we have applied an au-
876 thenticated encryption scheme (Rogaway et al., 2001)
877 that can provide both *privacy* and *authentication*. Pri-
878 vacy means that application execution state is encrypted
879 so that only the intended recipient with the correct
880 password can decrypt it. Authentication means that the
881 intended recipient can distinguish a legitimate applica-
882 tion migration sent by its user from a malicious appli-
883 cation migration sent by other users. Roam's security
884 works as follows. At the start of the application mi-
885 gration, the Roam agent running on the source device
886 prompts the user for a password. This password is a
887 one-time password that is valid only for one application
888 migration. At the same time, it will also generate a one-
889 time number called a *nonce*. The only requirement for
890 the nonce is that it is a new number for each application
891 migration. The nonce can be generated from a random
892 number generator or a simple counter that increments
893 each time it is used. Based on the one-time password and
894 nonce, the application execution state is encrypted as
895 ciphertext. The ciphertext is sent together with the
896 cleartext nonce to the target device. On the target device,
897 the Roam agent asks the user for the same password
898 that was used to encrypt the ciphertext on the source
899 device. Based on the password and the nonce, the Roam
900 agent can decrypt the ciphertext to get the application
901 execution state. If the password used in decryption does
902 not match the password used in encryption, the au-
903 thenticated encryption algorithm returns "INVALID"
904 to the Roam agent and the migration request is denied.

905 3. Implementation

906 We have implemented the Roam system using the
907 Java language and the Java language features: RMI,
908 Serialization, and Reflection. The Roam system cur-
909 rently runs on the PCs with JVM, and on Pocket PC
910 (PDA) devices with Personal Java VM. In the future, we
911 are planning to port the Roam system to other JVMs
912 with RMI and Serialization. The total code size of the
913 Roam system is approximately 15,000 lines, with 8500
914 lines in the SGUI toolkit.

The Roam system provides a set of Roamlet APIs. 915
The core class in the Roamlet APIs is called Roamlet, 916
which all Roamlets must extend. The Roamlet APIs 917
allow a Roamlet to migrate at runtime, to be instanti- 918
ated dynamically, to offload computation, or to be 919
transformed. We describe the essential Roam APIs here. 920

3.1. Roamlet runtime migration API 921

The runtime migration API is encapsulated in a class 922
called Roamlet shown in Fig. 11. All Roamlets must 923
extend the class Roamlet. A Roamlet calls the migrate() 924
method to be dispatched to a target device specified by 925
the hostname parameter. The Roam agent calls the 926
Roamlet's onInitialization() method when it is first in- 927
stantiated on a device. The Roam agent calls the 928
Roamlet's onRemoval() method before it is removed 929
from the source device, and after it migrates successfully 930
to a target device. The Roam agent calls the Roamlet's 931
onArrival() method when it arrives at the target device. 932
A Roamlet calls the exit() method to remove itself from 933
the Roam system. The instantiate() method is called to 934
instantiate a Roamlet on a local host. After a Roamlet is 935
instantiated on a device, the Roam agent creates a sep- 936
arate thread to run it. 937

3.2. Roamlet component description API 938

A Roamlet can describe its components using the 939
API shown in Fig. 12. A Roamlet calls addComponentDesc() 940
to add a component descriptor, RoamletComponentDesc, 941
for each of its components. RoamletComponentDesc 942
describes the type of component (type, which can be M-DD, 943
S-DD, or S-DI), the class name of this component, the class 944
name of the component to be instantiated (classname), and a 945
list of possible implementations for this component. Each 946
implementation of a component is described in a separate 947
ImplDesc, which contains the class name (classname) of the 948
implementation and the device requirement 949
for this implementation (req). 950
951

```

public abstract class Roamlet implements Serializable {
    ... // from Figure 11
    public synchronized boolean addComponentDesc(RoamletComponentDesc desc);
};

public class RoamletComponentDesc {
    public RoamletComponentDesc(String classname, ComponentType type, ImplDesc implDesc[]);
};

public class ImplDesc {
    public ImplDesc(String classname, DeviceRequirement req);
};

```

Fig. 12. Roamlet component description API.

952 3.3. SGUI API

953 The SGUI library provides four packages (sgui,
954 sgui.event, sgui.dd.event, and sgui.ir) for UI developers.
955 The sgui package provides virtual widgets similar to
956 widgets available in the Swing library with prefix S, such
957 as SButton, SCheckbox, STextfield, SLabel, etc. The
958 sgui.event package provides device-independent events
959 with S prefix, such as SActionEvent. The sgui.dd.event
960 package provides device-specific events, such as
961 SMouseEvent. UI developers follow the Java 2 event
962 model to add SGUI event listeners and specify actions
963 upon the reception of SGUI events.

964 The sgui.ir package provides APIs for constructing a
965 device-independent representation. One of the core
966 classes is the LogicalPanelNode shown in Fig. 13. It has
967 an add() method to add virtual widget and logical
968 panels as its child nodes. A LogicalPanelNode can
969 specify a list of properties, such as layout constraint,
970 task preference, layout priority, whether it can be pag-
971 inated across multiple pages, and whether it represents
972 core or optional features.

973 3.4. Chess Roamlet

974 We have taken a Java-based Chess game that is freely
975 available on the Internet and rewritten it as a Roamlet.
976 The Chess game has two modes: multiplayer mode
977 where two users play against each other, and computer
978 mode where a user plays against AI. The mobile envi-
979 ronment is setup as in Fig. 14: a Pocket PC device

running Personal Java VM, a notebook PC running a 980
standard J2SE VM, and they are both connected by an 981
802.11 wireless LAN. The user starts the Chess game in 982
computer mode on the notebook PC and then migrates 983
it to the Pocket PC. There is a third device, a desktop 984
PC, acting as the remote server where the AI component 985
can be offloaded to. 986

The Chess Roamlet contains the following compo- 987
nents: 988

- An application logic component is implemented as an 989
offloadable device-dependent component (S-DD). It 990
runs an AI engine that calculates the computer's next 991
move by building a search tree, which is both compu- 992
tationally and memory-intensive. It has a device re- 993
quirement of a PC. 994
- A GUI component is implemented as a non-offloada- 995
ble device-independent component (S-DI) using the 996
SGUI library. It draws a player board that holds 997
both a chessboard and chat/message boxes where 998
two players can send messages to each other. 999
- A second GUI component is implemented as a non- 1000
offloadable device-dependent component with two 1001
implementations (M-DD). These two implementa- 1002
tions have two different sets of images showing chess 1003
pieces, the set for the PC presentation has larger im- 1004
ages than that of the Pocket PC presentation. 1005

The Chess game in computer mode with PC presen- 1006
tation is shown on the left side of Fig. 15. When the user 1007
migrates the Chess game to the Pocket PC, the first (S- 1008

```

- sgui.SComponent
  |
  +- (device-independent widgets: SButton, SCheckbox, STextfield, SLabel, SPanel, etc)

public class LogicalPanelNode extends Node {
    public LogicalPanelNode(NodeGridBagLayoutConstraint layoutConstraint,
        boolean[] taskPreference,
        int layoutPriority,
        boolean splittability,
        boolean coreFeature);

    public add(Node node);
}

```

Fig. 13. Sample APIs for constructing device-independent representation.




Fig. 14. Experimental setup of Chess Roamlet.

1009 DI) GUI component is transformed into the Pocket PC
1010 presentation shown on the right side of Fig. 15. Since the
1011 player board is too large for the Pocket PC's display,
1012 SGUI layout algorithm paginates it into two pages. The
1013 first page shows the chessboard, and the second page
1014 shows chat and message boxes. A page navigation menu
1015 is added to the right of the menu bar for users to switch
1016 between two pages. The M-DD GUI component is in-
1017 stantiated with the Pocket PC presentation and it loads
1018 the smaller set of images for chess pieces. The applica-
1019 tion logic component is offloaded to the computing
1020 server.

1021 3.5. Experience on development effort and performance

1022 The Chess game was originally implemented to run
1023 on the PC using Java AWT GUI library. We have asked
1024 third party developers, who are new to Roamlet pro-
1025 gramming, to rewrite the Chess game as a Roamlet us-
1026 ing SGUI library. The experience from the third party
1027 developers have told us that most of the efforts in re-
1028 writing the Chess application fall into the following
1029 three categories: (1) modularizing an application into
1030 application logic components and GUI components,

and converting them into Roamlet components, (2)
porting GUI components to use the SGUI library, and
(3) adding multiple device-specific GUI components for
platforms that are not supported by the original Chess
application. Their experience also told us that it is rel-
atively straightforward to separate application logic
components from UI components in a well-written Java
application. The difficult part is to decide the type (S-DI,
S-DD, or M-DD) for each component. Finally, their
experience told us that, since our SGUI library follows
Java Swing library, porting AWT components to use
SGUI components is relatively straightforward.

The original Chess game has approximately 7900
lines of code. After the converting it into a Roamlet, it
has approximately 9000 lines of code, which is about
14% increase in code size. Most of the additional code
comes from code that converts Java objects into
Roamlet components, and the second (M-DD) GUI
component that supports the Pocket PC presentation.
This shows the tradeoff between development cost and
UI quality. That is device-dependent UIs in general have
higher quality, but developers need to create M-DD
presentations for different platforms, resulting in in-
creased application code size. This tradeoff can also be
found in the Connect4 game. When its UI component
has two device-dependent implementations for PC and
PDA, the code size is approximately 1200 lines. The
code size is reduced to 1100 lines when its UI component
is re-implemented as a device-independent component
using SGUI library.

We have measured *migration latency* of application
migration. Migration latency is defined as the elapsed
time from the moment the user accepts the migration
request on the target device, to the moment when the
Roamlet is completely restored on the target device. We
have measured latency for both Connect4 and Chess
games migrating from a notebook PC (Intel Pentium
III-800 CPU running Windows XP) to a Compaq iPAQ
H3870 Pocket PC (StrongARM 206 MHZ CPU running



Fig. 15. Screenshots of Chess game on notebook PC and Pocket PC after migration.

Table 3
Migration latency measurements

Applications	Migration latency			
	Notebook PC to Pocket PC		Pocket PC to notebook PC	
	Context transfer time (s)	Transformation time (s)	Context transfer time (s)	Transformation time (s)
Chess	4.9	10.2	4.2	1.1
Connect4	2.8	6.4	1.2	1.4
Shop hunter	N/A	3.7	N/A	0.9

Linux) and vice versa. We divide the migration latency into two parts: (1) *context transfer time* is the time for saving the execution state on the source device, sending and restoring it on the target device, and (2) *GUI transformation time* is the time for transforming a device-independent GUI component into a target device-specific presentation. The measurements are shown in Table 3.

Measurements show that transformation time dominates the migration latency when migration is from a notebook PC to a Pocket PC, because the layout algorithm is computationally intensive and it runs the target device which is a slower Pocket PC. However, when migration is from a Pocket PC to a notebook PC, the context transfer time dominates the migration latency. This is due to serialization of the Roamlet execution time on a slow Pocket PC. For the shop hunter application, transformation is offloaded to a PC. Measurements show that the transformation manager takes less time to generate a presentation for a bigger display (0.9 s) than for a smaller display (3.7 s). The reason is that generating a presentation for a smaller display is likely to involve testing different paginations and transformation rules.

3.6. End-user usability

Our limited experience on the available Roamlet applications shows that SGUI toolkit can generate consistent presentations across different platforms. There are three aspects of consistencies: task consistency, layout consistency, and transformation consistency. Task consistency means that, unless the developers specify that some tasks are not appropriate on certain platforms, all tasks will be presented across all platforms. Layout consistency means that, if UI components are laid out next to each other on one platform, they are likely to be found adjacent to each other or on adjacent pages on other platforms. Transformation consistency means that users can expect similar transformations to be applied to same type of UI components across applications. We believe that these consistencies contribute to better *learnability* of applications.

Our simple page navigation menu works only for simple UIs but not complex UIs. We believe that we can improve the usability and efficiency of generated page

navigation by incorporating device-specific navigation models provided by SGUI toolkit or explicitly specified by UI developers. To generate better page navigation, the navigation model should consider interactions and relation between tasks and device's input methods.

4. Related work

We divide the related work into two parts: mobile agent systems and multi-platform UIs.

4.1. Mobile agent systems

There has been an abundance of research work in the area of migrateable mobile agent systems. Aglet (Lange and Oshima, 1998) is one of the most well-known Java-based agent systems. It provides a Java development toolkit and libraries for building mobile agents that can move from one computer to another. Like Roamlet, Aglets do not require any modifications to the Java VM. An Aglet is shielded through an Aglet Proxy, which protects the Aglet from unauthorized access. Aglets run in an execution environment called AgletContext, which is like the Roam agent. However, Roamlets differ from Aglets in that Roamlets assume a heterogeneous device environment, whereas the Aglet assumes a homogeneous PC or PC-equivalent environment. Furthermore, the Roam system implements dynamic instantiation and computation apportioning to address the device heterogeneity problem. We have found similar Java-based agent systems in Jumping Beans (Aramira, 1999), MOA (Milojicic et al., 1998), Concordia (Mitsubishi Electric ITA Horizon Systems Laboratory, 1998), Voyager (ObjectSpace, xxxx), Mole (Strasser et al., 1996) and Telescript (White et al., 1997). Voyager is a commercial product by ObjectSpace, and it incorporates features of mobile agents and mobile objects into the ORB and COBRA. MOA performs resource management on migrateable applications as to what system resource they can use on the target device, and it also supports migrateable communication channels (the external running state) so that running communication channels can migrate with the rest of the applications. There are also similar but non-Java-based agent systems, e.g., Agent-Tcl (Gray et al., 1997).

1155 Fünfroeken (Fünfroeken, 1998) and Ara (Peine and
1156 Stolpmann, 1997) describe ways to achieve *strong mi-*
1157 *gration* for Java applications. Strong migration means
1158 that the call stack, which is a part of the thread state that
1159 is not serializable, can also be migrated. Fünfroeken
1160 uses a preprocessor that adds code to capture and to
1161 restore the state of the program stack at various exe-
1162 cution points. Ara proposes modifications to the JVM
1163 interpreters. Since thread state migration is not a specific
1164 challenge for heterogeneous device, the Roam system
1165 supports only weak mobility.

1166 Sumatra (Acharya et al., 1997) is an extension of Java
1167 that supports resource-aware migrateable mobile appli-
1168 cations. Sumatra's resource-awareness is based on a
1169 monitor-feedback-adaptation loop. A resource monitor
1170 watches the level and quality of the system resources,
1171 and provides updates to the applications either contin-
1172 uously or on-demand. Applications can then adapt
1173 based on the resource updates, e.g., migrating some
1174 threads to remote devices to perform the computation.
1175 Sumatra is focused on building adaptive applications in
1176 the PC or PC-equivalent device environment, and ad-
1177 aptation is based on runtime resource level and quality.
1178 This is different from the Roam system, which is focused
1179 on adaptation in a heterogeneous device environment.
1180 Roamlets adapt according to target device capability at
1181 migration time and load time only.

1182 We have found parallel and ongoing research efforts
1183 at IBM (Banavar et al., 2000). The IBM researchers
1184 have proposed a new application model for pervasive
1185 computing. They have described several research chal-
1186 lenges, e.g., device-specific rendering and application
1187 apportioning, which are similar to the problems that the
1188 Roam system is addressing.

1189 4.2. Multi-platform user interfaces

1190 The model-based approach offers an attractive alter-
1191 native to build a high-level tool for multi-platform UIs
1192 (Szekely, 1996). In model-base systems such as Hu-
1193 manoid (Szekely et al., 1993), ITS (Wiecha et al., 1990),
1194 UIDE (Sukaviriya et al., 1993), and Mickey (Olsen,
1195 1989), UI designers would use a declarative language to
1196 specify abstract and high-level *models* that describe what
1197 UI should be. The model-base system would automati-
1198 cally generate low-level UI executable code. The model-
1199 based technique can be exploited for multi-platform UI
1200 generation. For example, UI designers can specify the
1201 high-level models using AIOs (abstract interactive ob-
1202 jects), which are device-independent objects. The AIOs
1203 are mapped to device-dependent CIOs (concrete inter-
1204 active objects) supported by the target platform UI li-
1205 brary. This AIO-CIO selection technique has been
1206 illustrated in model-based systems for purposes other
1207 than multi-platform UI generation. For example, Hu-
1208 manoid uses a "replacement hierarchy" for selecting the

1209 most appropriate presentation template for displaying a
1210 semantic object. UIDE contains constraints on how to
1211 select the most appropriate interface actions to represent
1212 application actions specified by the UI designers in ap-
1213 plication models.

1214 In recent work, RedWhale software (Eisenstein et al.,
1215 2000, 2001) employs the model-based approach by in-
1216 troducing abstract models that specifically target multi-
1217 platform UIs. It is comprised of platform model, pre-
1218 sentation model, and task model. The platform model
1219 allows the UI designers to specify platform constraints
1220 that the UI generation must follow, such as device
1221 screen sizes, input methods, etc. The presentation model
1222 specifies the visual appearance of the generated UI,
1223 which may include the layout of the UI components and
1224 the mapping between AIOs and CIOs. The task model
1225 specifies the decomposition of big tasks into smaller sub-
1226 tasks. The device-dependent UI generation is based on
1227 all three models.

1228 Another recent model-based work is Paternò et al's
1229 ConcurTaskTree (Paternò, 2000). Its task model con-
1230 tains user tasks, abstract tasks, interaction tasks, and
1231 application tasks. Each task is related to each other via a
1232 set of temporal relationships such as enabling, deacti-
1233 vation, and iteration. The ConcurTaskTree is generally
1234 accepted to be a powerful task model to construct UIs.
1235 However, it is mainly targeted to single platform UIs. If
1236 we try to apply ConcurTaskTree to generate multi-
1237 platform UIs, all UIs will have the same interaction
1238 pattern as the interaction tasks are tightly coupled in the
1239 model. However, devices such as cell phones and PDAs
1240 have distinctive interaction patterns. To solve this
1241 problem, our device-independent representation only
1242 has user tasks. The interaction pattern information is
1243 expressed as properties of the task, e.g., task preference,
1244 or as platform-specific interaction models that would
1245 interact with the task model. The latter expression will
1246 be our future work. With our task model, we can easily
1247 customize the interaction pattern for each platform.

1248 Although the model-based approach is an attractive
1249 technique for creating SGUI tools, it has usability and
1250 authoring issues (Myers, 1995) such as loss of fine-
1251 grained control of UI details, the quality of generated
1252 UI, and the time needed for UI designers to learn and
1253 construct models. Our SGUI toolkit combines the
1254 model-based approach with the benefits of traditional
1255 widget-based UI programming, as a solution to some of
1256 these issues. Programming with the SGUI library is very
1257 similar to programming with a traditional widget library
1258 such as Swing, because the SGUI library is consisted of
1259 low-level UI widgets rather than abstract models. The
1260 transformation in SGUI is at the level of widgets and
1261 events, rather than from abstract interaction models to
1262 widgets and events. This has additional benefits that no
1263 code generation is necessary, and that complex models

1264 of *UI Logic* (data flow constraints, sequencing, side ef-
1265 fects, etc.) (Szekely et al., 1992) are also not needed.

1266 Microsoft has a commercial product available called
1267 the .NET Mobile Internet Toolkit (Microsoft Corpo-
1268 ration, xxxx), which targets different mobile devices
1269 ranging from the more capable PocketPC platform to
1270 relatively simple pagers. The Mobile Toolkit does not
1271 use a model-based approach. Rather, it assumes devel-
1272 opers have already implemented desktop PC versions of
1273 the user interface components, and that they want to
1274 port from existing code. The Mobile Toolkit allows
1275 developers to specify the corresponding UI presentation
1276 on mobile platforms, but the layout specification is very
1277 limited. It only offers flow-based layout (rather than the
1278 more sophisticated Grid-Bag layout). In addition, de-
1279 velopers must group widgets together into a *form*, the
1280 function of which is very similar to our LogicalPanel-
1281 Node. However, .NET mobile forms cannot be placed
1282 on the same page even if a page has enough space to
1283 accommodate them.

1284 5. Conclusion

1285 In this paper, we present challenges, design, and im-
1286 plementation of the Roam system. The Roam system is
1287 a seamless application framework for building seamless
1288 applications that can migrate at runtime across hetero-
1289 geneous devices. The Roam system provides adaptation
1290 strategies at the component level, including dynamic
1291 instantiation, offloading computation, and transforma-
1292 tion.

1293 There are many future directions to improve the
1294 Roam system. One problem with the existing SGUI
1295 toolkit is that it is difficult to customize a device-inde-
1296 pendent representation for a particular device. The
1297 reason is that the device-specific customization requires
1298 developers to add transformation rules that are both
1299 device-specific and application-specific. When the de-
1300 velopers change the device-independent model at a later
1301 time, they may also have to update these transformation
1302 rules that are affected by the change.

1303 We also like to support seamless application migra-
1304 tion for real-time applications such as video conferenc-
1305 ing. For example, a use may want to migrate a video
1306 conferencing from a mobile phone to a car navigation
1307 system when he/she is entering a car. This places a real-
1308 time constraint on migration latency. We want to make
1309 sure that the interruption time is minimized during ap-
1310 plication migration. One possible approach is to hide the
1311 migration latency, by delaying the application termina-
1312 tion on the source device, until adaptation has been
1313 applied to the application components on the target
1314 device.

1315 Finally, we would like to support a Roam execution
1316 state repository service where a user can save a Roam

application on one device, and restore it at a later time
on any device. This overcomes the limitation in the
current Roam system where application migration can-
not be suspended. We would like to extend this to col-
laborative applications where each user can save and
restore application execute state individually without
affecting other users.

6. Uncited references

Chu et al. (2001), NTT DoCoMo (2001), Sukaviriya
et al. (1994), Sun Microsystems (1998, 2000a,b) and
Wiebus (xxxx).

Acknowledgements

We would like to thank Dan Rosen for his help in
proof-reading this document. We would also like to
thank reviewers of JSS for their comments.

References

- Acharya, A., Ranganathan, M., Saltz, J., 1997. Sumatra: a language for
resource-aware mobile programs. *Mobile Object Systems, Lecture
Notes in Computer Science* April, 1997.
- Aramira, Inc., 1999. *Jumping Beans™* White Paper, October 1999.
Available from <<http://www.jumpingbeans.com/index.html>>.
- Banavar, G., Beck, J., Gluzberg, E., Munson, J., Susman, J.,
Zukowski, D., 2000. Challenges: an application model for
pervasive computing. In: *Proceedings of the 6th Annual Interna-
tional Conference on Mobile Computing and Networking, Bos-
ton, MA, August 2000*, pp. 266–274.
- Chu, H., Song, H., Wong, C., Kurakake, S., 2001. Seamless Appli-
cations over Roam System, *UbiTools'01* (Part of ACM Ubi-
Comp'01), September 2001. Available from <[http://
choices.cs.uiuc.edu/UbiTools01/](http://choices.cs.uiuc.edu/UbiTools01/)>.
- Eisenstein, J., Vanderdonck, J., Purerta, A., 2000. Adapting to mobile
contexts with user-interface modeling. In: *Proceedings
WMCSA'00, December 2000*.
- Eisentein, J., Vanderdonck, J., Puerta, A., 2001. Applying model-
based techniques to the development of uis for mobile computers.
In: *Proceedings IUI'01, January 2001*.
- Fünfroeken, S., 1998. Transparent migration of Java-based mobile
agents (capturing and reestablishing the state of Java programs).
In: *Proceedings of the 2nd International Workshop on Mobile
Agents, Stuttgart, Germany, September, 1998*.
- Gray, R.S., Kotz, D., Nog, S., Rus, D., Cybenko, G., 1997. Mobile
agents for mobile computing. In: *Proceedings of the Second Aizu
International Symposium on Parallel Algorithms/Architectures
Synthesis, Fukushima, Japan, March, 1997*.
- Lange, D.B., Oshima, M., 1998. Mobile agents with Java: the Aglet
API. *World Wide Web Journal*.
- Linton, A., Vlassides, M., Calder, R., 1989. Composing user interfaces
with interviews. *IEEE Computer* 22 (2).
- Masui, T., 1994. Evolutionary learning of graph layout constraints
from examples. In: *Proc. of ACM UIST'94, November 1994*,
pp.103–108.

- 1368 Microsoft Corporation, Best Practices for the Microsoft Mobile
1369 Internet Toolkit Image Control. Available from <[http://msdn.mi-](http://msdn.microsoft.com/theshow/Episode023/default.asp)
1370 [crosoft.com/theshow/Episode023/default.asp](http://msdn.microsoft.com/theshow/Episode023/default.asp)>.
- 1371 Milojicic, D.S., LaForge, W., Chauhan, D., 1998. Mobile objects and
1372 agents (MOA). In: Proceedings of the 4th USENIX Conference
1373 on Object-Oriented Technologies and Systems (COOTS). Santa
1374 Fe, New Mexico, April, 1998.
- 1375 Mitsubishi Electric ITA Horizon Systems Laboratory, 1998. Mobile
1376 Agent Computing A White Paper, January 1998.
- 1377 Myers, B., 1995. User interface software tools. *ACM Transactions on*
1378 *Computer Human Interaction* 2 (1), 64–103.
- 1379 NTT DoCoMo, Inc., 2001. i-mode Java Content Developer's Guide,
1380 May 2001.
- 1381 ObjectSpace, Inc., Voyager. Available from <[http://www.object-](http://www.object-space.com/products/voyager)
1382 [space.com/products/voyager](http://www.object-space.com/products/voyager)>.
- 1383 Olsen, D., 1989. A programming language basis for user interface
1384 management. In: Proceedings of CHI'89, Austin, Texas, May
1385 1989, pp. 171–176.
- 1386 Olsen, D., Jefferies, S., Nielsen, T., Moyes, W., Fredrickson, P., 2000.
1387 Cross-modal Interaction using XWeb. In: Proc. of ACM UIST'00,
1388 November 2000.
- 1389 Paternò, F., 2000. Model-Based Design and Evaluation of Interactive
1390 Applications. Springer-Verlag, London.
- 1391 Peine, H., Stolpmann, T., 1997. The architecture of the ara platform
1392 for mobile agents. In: Proceedings of the 1st International
1393 Workshop on Mobile Agents, Berlin, Germany, April 1997.
- 1394 Rogaway, P., Bellare, M., Black, J., Krovetz, T., 2001. In: Eighth
1395 ACM Conference on Computer and Communications Security
1396 (CCS-8). ACM Press, pp. 196–205.
- 1397 Strasser, M., Baumann, J., Hohl, F., 1996. Mole—a Java based mobile
1398 agent system. In: 2nd ECOOP Workshop on Mobile Object
1399 Systems, Linz, Austria, July 1996, pp. 28–35.
- 1400 Sukaviriya, P., Foley, J.D., Griffith, T., 1993. A second generation user
1401 interface design environment: the model and the runtime archi-
1402 tecture. In: Proceedings InterCHI'93, April 1993, pp. 375–382.
- Sukaviriya, P., Kovacevic, S., Foley, J.D., Myers, B., Olsen, D., 1403
Schneider-Hufschmidt, M., 1994. Model-based user interfaces,
1404 what are they and why should we care? In: Proceedings UIST'94,
1405 November 1994, pp. 133–135.
- Sun Microsystems, 1998. PersonalJava™ Technology—White Paper,
1406 August. 1407
- Sun Microsystems, 2000. Java™ 2 Platform, Standard Edition White
1408 Paper, June 2000. 1409
- Sun Microsystems, 2000. Java™ 2 Platform Micro Edition(J2ME™)
1410 Technology for Creating Mobile Devices, White Paper, May 2000. 1411
- Szekely, P., 1996. Retrospective and challenges for model-based
1412 interface development. In: Proceedings of 3rd Int. Workshop on
1413 Computer-Aided Design of User Interfaces CADUI'96, June
1414 1996, pp. xxi–xliv. 1415
- Szekely, P., Luo, P., Neches, R., 1992. Facilitating the exploration of
1416 interface design alternatives: the HUMANOID model of interface
1417 design. In: Proceedings CHI'92, May 1992, pp. 507–514. 1418
- Szekely, P., Luo, P., Neches, R., 1993. Beyond interface builders:
1419 model-based interface tools. In: Proceedings InterCHI '93 April
1420 1993, pp. 383–390. 1421
- Vanderdonckt, J., 1995. Knowledge-based systems for automated user
1422 interface generation: the trident experience. Technical Report RP-
1423 95-010, Fac. Univ. de N-D de la Paix, Inst. D'Informatique,
1424 Namur. 1425
- White, J., et al., 1997. System and Method for Distributed Compu-
1426 tation Based upon the Movement, Execution, and Interaction of
1427 Processes in a Network, US patent no. 5603031, February 1997. 1428
- Wiebus, S., Connect4 Applet. Available from <[http://www.physics.a-](http://www.physics.a-delaide.edu.au/~swright/java_apps/Connect4/)
1429 [delaide.edu.au/~swright/java_apps/Connect4/](http://www.physics.a-delaide.edu.au/~swright/java_apps/Connect4/)>. 1430
- Weicha, C., Bennett, W., Boies, S., Gould, J., Green, S., 1990. ITS: a
1431 tool for rapidly developing interactive applications. *ACM Trans-*
1432 *actions on Information Systems* 8 (3), 204–236. 1433
- W3C, 2000. Scalable Vector Graphics (SVG) 1.0 Specification,
1434 November 2000. Available from <<http://www.w3.org/TR/SVG/>>. 1435
1436