


d e
The App





**MAKING THE MOST
OF COLOR ON 1-BIT
DEVICES**

**THE TEXTBOX YOU'VE
ALWAYS WANTED**

**MAKING YOUR
MACINTOSH SOUND
LIKE AN ECHO BOX**

**SIMPLE TEXT
WINDOWS VIA THE
TERMINAL MANAGER**

**TRACKS: A NEW
TOOL FOR
DEBUGGING DRIVERS**

**USING THE
PALETTE MANAGER
OFF-SCREEN**

**BACKGROUND-ONLY
APPLICATIONS IN
SYSTEM 7**

MACINTOSH Q & A

APPLE II Q & A

**NEW FEATURE:
KON & BAL'S
PUZZLE PAGE**



EDITORIAL STAFF

Editor-in-Cheek *Caroline Rose*

Technical Buckstopper *Dave J*

Review Board *Pete "Luke" Alex*

Derossi, C. K. Haun, Larry K

Shebanow, Gregg Williams

Managing Editor *Monica Meff*

Assistant Managing Editor *An*

Contributing Editors *Lorraine*

Toni Haskell, Judy Helfand, L

Rilla Reynolds, Leslie Steere, C

Indexer *Ira Kleinberg*

Manager, Developer Support S

Communications *David Kr*

ART & PRODUCTION

Production Manager *Hartley L*

Art Director *Diane Wilcox*

Technical Illustration *Geoff M*

John Ryan

Formatting *Forbes Mill Press*

Printing *Wolfer Printing Comp*

Film Preparation *Aptos Post, In*

Production *PrePress Assembly*

Photography *Sharon Beals, De*

Katbleen Siemont

Circulation Management *Dav*

Robinson
Cander, Chris
Rosenstein, Andy

ert
a Wilczynski
Anderson,
Rebecca Pepper,
Carol Westberg

Systems and
athwohl

ON

esser


Cormack,

ny, Inc.

nc.

mis Hescoc,

id Wilson



To create this cover, Hal Rucker and Cleo Huggins bought the nicest-looking fruit they could find, photographed it and scanned in a slide, manipulated the scan with Adobe Photoshop, and blended in a dithered version of it. Delicious!

develop, *The Apple Technical Journal*, is a quarterly publication of the Developer Support Systems and Communications group.

The *Developer CD Series* disc for February 1992 or later contains this issue and all back issues of *develop* along with the code that the articles describe. The contents of this disc, which includes other handy software and documentation, can also be found on AppleLink.

EDITORIAL We want your two cents! **2**

LETTERS Your praise and your scorn. **4**

ARTICLES **Making the Most of Color**

Daniel Lipton A two-part article on color printing machines, and the theory and practice of color.

The TextBox You've Always Wanted

Here's a replacement for TextEdit that offers international compatibility. **12**

Making Your Macintosh Sound Like a Pro

Learn how to use double buffering to make your Macintosh sounds. **48**

Simple Text Windows via the Terminal Manager

The Terminal Manager (in the Macintosh Toolbox) offers output capabilities in your applications. **56**

Tracks: A New Tool for Managing Disk Space

Managing disk space with device drivers is a pain. This new driver, greatly easing your dilemma. **64**

COLUMNS **Graphics Hints From Forrest**

by Forrest Tanaka Can you get better performance from screen ports? Well, yes, but there are some caveats. **70**

Be Our Guest: Backgrounds

Haun Faceless background graphics. C. K. shows you the basics. **78**

The Veteran Neophyte: A Guide to the Macintosh

Macintoshes are not only great for studying, they're fun. At least Dave thinks so. **82**

KON & BAL's Puzzle Page

and Bruce Leak Are there any more bugs? A debugging puzzle. **90**

Q & A Answers to your product development questions.

Macintosh Q & A 85

Apple II Q & A 100

INDEX 106

© 1992 Apple Computer, Inc. All rights reserved. AppleTalk, EtherTalk, GS/OS, ImageWriter, LaserWriter, and TokenTalk are trademarks of Apple Computer, Inc., registered in the U.S. and other countries. develop, Finder, Macintosh Coprocessor, and the Terminal Manager, Tools Advisor, and TrueType are trademarks of Apple Computer, Inc. licensed to Clarisc Systems Inc. MacWrite is a registered trademark of Apple Computer, Inc. CompuServe, Inc. Internet and VAX are trademarks of Digital Equipment Corporation. International Business Machines Corp. Lotus is a registered trademark of Lotus Development Corporation. Linotype is a registered trademark of Linotype Company. Microsoft is a registered trademark of Microsoft Corporation. Concepts, Inc. Sony is a registered trademark of Sony Electronics Inc. UNIX System is a registered trademark of UNIX System Laboratories.

or on 1-Bit Devices by **Konstantin Othmer** and
 Article: how to create color PICTs on black-and-white
 practice of dithering. **7**

ays Wanted by **Bryan K. ("Beaker") Ressler**
 eBox, with better performance, more flexibility, and
 What more do you want? **31**

Sound Like an Echo Box by **Rich Collyer**
 efering techniques to simultaneously record and play

ia the Terminal Manager by **Craig Hotchkiss**
 e Communications Toolbox) provides handy text
 application with virtually no effort. **60**

Debugging Drivers by **Brad Lowe** Debugging
 tool provides an easy way to log information from your
 bugging woes. **68**

rrrest: Using the Palette Manager Off-Screen
 ou use the Palette Manager to manage colors in off-
 here are some caveats. **29**

und-Only Applications in System 7 by **C. K.**
 tasks provide a handy way out of some sticky situations.
58

Silicon Surprise by **Dave Johnson** Computers
 g complex systems, they *are* complex systems. Or at

ge: It's Just a Computer by **Konstantin Othmer**
 demons in Kon's computer? Or is it just a simple
 to tickle your brain. **103**

velopment questions.

reserved. Apple, the Apple logo, APDA, Apple IIGS, AppleLink, AppleShare,
 ter, LaserWriter, LocalTalk, MacApp, Macintosh, MPW, MultiFinder, and
 utor, Inc., registered in the U.S. and other countries. A/ROSE, Balloon Help,
 Platform, Macintosh Quadra, QuickDraw, QuickTime, SNA•ps, Sound
 re trademarks of Apple Computer, Inc. HyperCard is a registered trademark
 is Corp. Adobe Photoshop and PostScript are registered trademarks of Adobe
 ademark of Claris Corp. CompuServe is a registered trademark of
 rademarks of Digital Equipment Corp. IBM is a registered trademark of
 notronic is a trademark, and Helvetica and Times are registered trademarks,
 uted trademark of Microsoft Corp. Nisus is a trademark of Paragon
 mark of Sony Corporation. NuBus is a trademark of Texas Instruments. UNIX
 n Laboratories, Inc.



CAROLINE ROSE

Dear Readers,

Let's talk about *develop*: what you can do for it. This job is done by reading on and giving

Originally, *develop* was the accompanying CD, it was you, the developer, could be compatible with future systems were written primarily by Apple.

But other types of articles. Most notable was the group which source code was not by our readers, yet the overall favorable. So we've moved *possible*. We still make even in future systems.

Recently we've had some ideas, not code. Our current good Apple products, we to see. There are some Apple would like to know just what Developer Technical Support let us at *develop* know directly

Regarding who writes the Apple engineers, there's no We'd like to encourage a fellow developers—some showcase and release your offer something those *others* the assurance of future code prose shine so brilliantly

2

CAROLINE ROSE (AppleLink writing computer documentation Jobs was barely a teen. When moved in down the block from as a writer and then a programmer no notice—until they asked if she what even then was known as Around the time she completed tome, Steve left Apple to form

help turn your raw material into a polished piece—or tread lightly if you need. We'll give your article that professional look and find humor. So, if you're willing, please send me your ideas or our from there.

Back to the subject of your opinions about *develop*: Many of you Associates and Partners have by now been formally surveyed about support-related materials, of which *develop* is only one shining like to hear from the rest of you, however informally. I can't tell how important your opinions are and how much they'll affect *develop*. Express yourself! Tell us what's good or bad about this journal's delivery, or anything else. We're all ears.

Issue 8 ended with this trivia question: What word was used to describe the action of pressing a button on that first mouse? If you've gotten as of this writing, is "bug." Maybe you'll do better: The original hardcover *Inside Macintosh* Volumes I-III had a pattern of Macintosh computers across its endpapers (those heavy sheets at the end of hardcover books). What broke this pattern, and why?

SUBSCRIPTION INFORMATION

Use the order form on the last page of this issue to subscribe to *develop*. Please address all subscription-related inquiries to *develop*, Apple Computer, Inc., P.O. Box 531, Mt. Morris, IL 61054 (or AppleLink DEV.SUBS). •

BACK ISSUES

For information about how to obtain them, see the order form on the last page. Back issues are also on the *Develop* website.

ightly on it if that's all
eel without killing the
tlines, and we'll take it

ou who are Apple
on how you rate various
g example. We'd also
overemphasize how
elop's future. So please,
l's content, format,

instead of "click" to
The answer, which none
do better on this next
ad a running pattern of
s at the very beginning
hy?



Caroline Rose
Editor

ut back issues of *develop* and
see the reverse of the order
e of this issue. Back issues
Developer CD Series disc. •

LETTERS

CURLING UP WITH D

Regarding your editorial "I don't agree with you on liking 'copy' to be able to curl up and try to understand something the first time. I can always go to a computer and try examples. To lay back and put up my feet under a quilt in bed is more likely to let concepts sink in and come of my own, to spring forth with ideas."

I like the idea of sending the disc separately in its own case. I have never had a mangled disc. I like the magazine itself that has been on a trip. I received Issue 8 with the disc and hope that the disc is safe. The mailing label on the disc was torn half off, not torn, but detached. Flapping in the breeze, so I'm sure.

Keep up the good work on the magazine. I look forward to the next time.

—Robert Redmond

Thanks for your letter. It's always heartening to hear from developers who agree with me on this, but I understand the difference. Your opinions do matter.

The disc is now in a separate case and not mailed under a separate label. It should have arrived wrapped in a separate case. We'll send you the disc right away. Sorry about that.

—Caroline Rose

TEXT FORMATS GALORE

There's a problem with the Technical Notes which I hope is increasing frequency in all electronic publications. The

4

PLEASE WRITE!

We welcome timely letters to the editor, especially from readers reacting to what we publish in *develop*. Letters should be addressed to Caroline Rose (caroline@develop-related questions, to Dan Johnson, Apple Computer, Inc., 20525 Linden Avenue, M/S 75-2B, Cupertino, CA 95014. E-mail: CROSE or JOHNSON.DK). A

DEVELOP

in Issue 8: I
to have a "hard
up with when
something for the
to the
es or ideas. But
y feet or nestle
ore relaxing to
develop on their
clarity later.

the disc
—though I
e problem. It is
as a rougher
thout the disc
not far behind.
back cover was
ached.

o to speak.
on the
to it each

not only
velopers who
it makes a
count.

te case, but it's
e label. They
ed cozily together.
t away. Sorry

MORE

he Macintosh
m finding with
Apple's
The only word

processor I use is Nisus, with which I can read MS Word 3.0 and 4.0 files without buying a Microsoft product. This may be an unreasonable prejudice, but I bet it isn't uncommon.

But Nisus can't decipher fast-saved Word files. This means, I suspect, that the entire set of new Macintosh technical publications is unavailable to me. Worse, I fear that the next Developer CD is going to have lots of files with new, valuable, and (for me) hidden information.

I know Apple is serious about electronic distribution of technical documents. I'm sure fast-saving in Word is a great convenience to the authors, but surely using a format not widely readable defeats the purpose of the exercise. I don't object to standardizing on Word 3.0 or 4.0, so long as that format—and not Microsoft's convenience variant—is actually used.

Could you *please* ask your authors, when providing documents for publication, to use an accessible format?

—Fritz Anderson

Thank you for alerting us to this problem. It was a snafu on our part. None of the files should have been fast-saved in Word.

We know that having text documents in Word and MacWrite® represents a bias toward these products. Unfortunately our alternatives are limited and we'll probably have to continue using these products until the spring.

The good news is that we're working on a new text formatting tool. This tool will be available on the CD and will be able to open, search, and print text documents available on the CD. The dilemma of how

the editors,
ng to articles that
should be
r, if technical
Dave Johnson) at
Mariani Avenue,
5014 (AppleLink:
ll letters should

include your name and company name as well as your address and phone number. Letters may be excerpted or edited for clarity (or to make them say what we wish they did). •

to make every document available to every developer has been a topic of discussion for some time. We're hoping this will solve the problem.

Again, thank you for your input. Developer feedback is the fuel of change around here. Keep it coming.

—Sharon Flowers

NEW AND IMPROVED CD

I just received Issue 8 of *develop*, and was pleased to find that the developer's CD has improved. Is this new?

—Mike Caputo

Yes, starting with Issue 8 the CD is not just Developer Essentials, but the entire Developer CD Series disc (of which Developer Essentials is just a subset).

—Caroline Rose

SUBMITTING TO DEVELOP

First of all, I'd like to say that I'm a big fan of *develop*. The combination of excellent technical articles (with required humor) and a CD-ROM of other developer materials is unmatched. At least the flak surrounding the CD-ROM has finally calmed down in the Letters section. I've always liked the idea from the start even though I purchased a CD-ROM drive only last week.

I'm writing to find out if *develop* accepts articles from non-Apple employees. I haven't looked through the back issues to see if there were any, but none come to mind. If so, do you have a style guide for writing articles?

Keep up the great work!

—Paul-Marcel S.

Thank you for your input. We'll continue to develop; it's always a work in progress. We'll send you a big fat letter that editors love.

We do indeed accept articles from non-Apple employees (see this issue for details). We have a vast array of articles, including prospective author guidelines, introductory documents, a set of detailed instructions, and even a MicroWorld format. We'll stand by you and send the rest as you wish.

On the subject of your question, it may experience some difficulty with Apple's dropping of the Macintosh II series. Notes from the Macintosh Associates and Palettes are what developers tend to use.

—Caroline Rose

ART ILLEGIBLE

Figure 1 from *Macintosh on Futures* (Issue 8) printed under *Software* on a Macintosh II on a CD-ROM. It's just as illegible as the ideas?

The HyperCard format is convenient for finding articles but is tedious to page at one time. If windows are not more, HyperCard is from a CD. And it's fast nor intuitive.

—Steve Tyler

at work!

St-Onge

*ur kind words about
ays a pleasure to receive
n. Yours is the type of
in chief dream of.*

*pt articles from non-Apple
s issue's Editorial). We
of materials ready for
rs, including an
ment, a short submission
iled author's guidelines,
soft Word template for
icle in a develop-like
rt with the intro and then
u need it.*

*the CD-ROM controversy,
a revival as a result of
printed develop and Tech
onthly mailing to
rtners. We'd like to hear
hink about that.*

E ON-LINE

Michael Gough's article
(p. 7) was illegible when
printed on a LaserWriter II
NTX. The text was illegible
on the screen. Any

HyperCard® format on the CD is
clipping to pages and
horrible for seeing all of a
page, since HyperCard's
pages are not resizable. Further-
more, HyperCard is slow—especially
when searching is neither
fast nor accurate.

I looked into the problem and a mistake was made when the version of Issue 7 was created. The EPS format is normally opened in an application that interprets the file and then saved as a PICT. This followed, with the result that the PICT was only an approximation and not very legible. This will be on the CD.

FEELING LOST? SEE THE MACINTOSH

To a Macintosh developer starting a new project, the range of equipment available can seem daunting. To an experienced engineer suffering the constant barrage of catalogs, technical brochures, and advertisements, it can feel safest to hang on to familiar tools, whatever their shortcomings. But what's available? What systems or tools might help you get your project done? Developer University has released the Macintosh Development Tools Advisor to help answer these questions.

The Tools Advisor offers a broad array of information. As a hypertext system, it tailors the data it presents to your particular interests and demands. The Advisor incorporates comprehensive technical data on over 80 programming tools—compilers and languages, debuggers and prototypers, CASE tools, and multimedia packages. It also includes essays on critical topics such as object-oriented programming, Apple events, and System 7. In preparing the Tools Advisor, Developer University collected a considerable body of catalog-style information on products available.

But a catalog is rarely sufficient. It's not enough to read lists of capabilities as recorded by manufacturers. You need to know how the tools get used in actual projects. So the Tools Advisor provides a collection of stories by programmers who use the tools it describes. These stories provide a real feel for the product. They're sometimes critical, warning of potential hazards and

6

You can obtain a copy of the Tools Advisor through APDA. The Advisor can also be found on the Developer University disc. To use the Tools Advisor, you need a Macintosh with System 6.0.5 or later, HyperCard 2.0 or later, and a CD-ROM drive. •

...d found out that
...he electronic
...d: Art that's in
...ned in an
...PostScript® and
...process wasn't
...t the conversion
...cimation, and so
...e fixed in Issue 7

Regarding the HyperCard format, a lot of people agree with you. We're working on an alternate viewing mechanism—but this mechanism may not apply to develop for a while yet. Meanwhile, HyperCard's windows are in fact resizable. If you're not able to resize them, your memory partition for HyperCard is probably not large enough; try increasing it.

—Caroline Rose

THE DEVELOPMENT TOOLS ADVISOR

... shortcomings of particular tools. They're also often
... inspiring in explaining how particular achievements
... were made. To help you find stories most appropriate to
... you, the Advisor lets you match a loose profile of your
... needs and wants to stories by developers with similar
... backgrounds and tasks.

To augment its profiles of programming tools, critical
... essays, and developers' stories, the Tools Advisor
... includes a glossary that describes exactly what technical
... and trade terms mean and what they imply to a
... development effort. Glossary entries and cross-
... references let you navigate the intricate terrain of
... technical information without losing sight of your
... particular interests.

Two versions of the Tools Advisor are available. The disk-
... based edition includes screen shots and comprehensive
... data on programming tools in a range of categories as
... well as technical details on the Macintosh and on
... Macintosh programming in general. The CD-ROM
... edition of the Tools Advisor adds demonstration versions
... of dozens of tools; for instance, you can take a
... multimedia tool for a test drive as you learn about
... animations and about other developers' experiences
... with that product.


We hope that with the Tools Advisor guiding you, you
... won't feel lost any more.

the Tools

... disk-based version
... loper CD Series
... you'll need a
... or later,
... hard disk. To use
... course also need

MAKING THE MOST OF COLOR ON 1-BIT DEVICES

Macintosh developers face the challenge of how to develop software for—many of them—older Macintoshes with Color QuickDraw. This article shows how to draw adequately on the lower-end Macintoshes. The technique, when running on the high-end Macintoshes, is a simple and elegant solution to the problem of color development and testing. The final product's outcome more satisfying, and the time to process them for display is less.



**KONSTANTIN OTHMER
AND DANIEL LIPTON**

Suppose you're writing a program that is supposed to work on all Macintosh computers. You want to use QuickDraw with the original QuickDraw support for bitmaps, thus severing you from the support for color. In our continuing quest for a solution to this problem with original QuickDraw, we have come up with a technique to process them for display, although it does not resolve single pixels, which requires the accompanying sample code (code available with you).

SAVING COLOR IMAGES

The key to saving color images with QuickDraw is a transcript of the PICT created on one Macintosh. The version of system software on the Macintosh Plus you can draw

KONSTANTIN OTHMER has wanted to be a professional athlete since he was a young boy. He has a photograph to appear in *Sports Illustrated* is his ultimate goal, as long as he can remember. Unfortunately, his athletic career at college was in the NCAA's Division II, which is often overlooked by *SI*'s editors, and they've missed his virtuosity on the slopes of Tahoe, Vail, and Red Lodge. So Konstantin has had to scale down his dream, setting his sights on becoming a developer. He's making the pages of *develop* instead

*ed with the dilemma of which platform to
achines with the original QuickDraw or those
—can always choose to write code that runs
end machines and gives additional functionality
gher-end machines. While this sounds like a
on, it generally requires a great deal of
effort. To make this effort easier and the
we offer techniques to save color images and
on 1-bit (black-and-white) devices.*

rogram that controls a 24-bit color scanner and you'd like
computers. The problem you'll run into is that machines
(those based on the 68000 microprocessor) only have
erely crippling the potential of your scanner. But don't
est to add Color QuickDraw functionality to machines
ve worked out techniques to save color images and
it in black and white, on the latter machines. We've also
address the problem of a laser printer's inability to
results in distorted image output. This article and the
on the *Developer CD Series* disc) share these techniques

AGES

es is using pictures. Recall that a picture (or PICT) in
calls to routines that draw something—anything. A
osh can be displayed on any other Macintosh (provided
e on the machine doing the displaying is the same as or
machine that created the picture). For example, on a
a PICT containing an 8-bit image that was created on a

nted his
strated for as
ately, his
III, which is
d somehow
ki slopes at
's had to
ghts on
d. Here he's

gotten to try on various alter egos. To come up
with his latest persona, he spent a few late nights
in a secret Apple lab with skilled pixel surgeon
Jim Batson. •

7

Macintosh II. With System 7, 8-bit `pixMaps` on machines with 24-bit displays are displayed as 1-bit images.

Creating a picture normally involves:

1. Call `OpenPicture`.
2. Perform the drawing.
3. Call `ClosePicture`.

The catch is that the only `OpenPicture` routines available on the Macintosh are those that create 1-bit pictures. This procedure on a machine with 24-bit displays converts `pixMaps` into a picture, since you can't create an 8-bit `PICT`. But that's exactly what would be needed to create a `PICT` containing deep `pixMaps`. With this ability, you could create a Color QuickDraw machine with original QuickDraw machine with original QuickDraw.

To get around the limitation, a routine called `CreatePICT2` to make a `PICT` that an application can display. The routine does whether creating your own `PICT` or one who directly modify private `PICT` data. Please ease your mind, see "But I Don't Want to Create a `PICT`."

The parameters to `CreatePICT2` are: `picRgn`, `stdBits`. The `picRgn` is not used if `stdBits` is not zero. The `picRgn` is not used if `stdBits` is not zero.

The first thing the routine does is to allocate that amount of storage for the `PICT` returning a `NIL PicHandle` if the memory is not available. Rather than writing out the `PICT` to disk, the specified number of bytes is written to the `picRgn`. Essentially, you need an `OpenPicture` routine that can create a `PICT` containing deep `pixMaps`.

At this point the size of the `PICT` is known, well the `pixMap` will complete the `PICT` picture frame. Next is the `PICT` data with a header that has version information. (C) ignore the header data. (C)

DANIEL LIPTON (a.k.a. "The Iguana Guy") is a two-and-a-half-year veteran of the Software Imaging Group, where he's working on the next generation of printing for the Macintosh. When he's not thinking about printing, he enjoys taking in a good flick, sipping coffee, and his iguana, "Iggy" (who's never been to the zoo). Dan for the time she nearly froze in the cargo compartment of a 747),

n 7, you can even display PICTs containing 16-bit and 32-bit with original QuickDraw. (Of course, they will only be there.)

ly requires three steps:

to begin picture recording.

ng commands you want to record.

to end picture recording.

drawing commands that can be recorded into a picture are on a Macintosh on which your application is running. Thus, using the original QuickDraw provides no way to save color information there's no call to draw a `pixMap`. In other words, you can't draw a `pixMap` on a Macintosh Plus and see it in color on a Macintosh II. This would make a developer's life easier—the ability to create a `pixMap` on a machine without Color QuickDraw. It would capture a color image in its full glory for someone with a monitor to see, while still being able to display a 1-bit version on a machine without Color QuickDraw.

In addition to the normal procedure, we came up with a routine to manually create a PICT containing color information. Your application can create a picture using `DrawPicture`. Now, you may be wondering if creating pictures is advisable. After all, Apple frowns on developers who create data structures, and isn't that what's going on here? To quote a friend: "Don't I Need a License to Do This?"

The `PICT2` routines are similar to those for the QuickDraw bottleneck. The difference is that `CreatePICT2` returns a `PicHandle` and does

what it does is calculate a worst-case memory scenario and return a `PicHandle`. If the memory isn't available, the routine aborts, returning `noPic`. You could easily extend this routine to spool the picture to a file if memory is not available, but that's left as an exercise for you. (*Hint:* To save the data inline as is done here, call a function that saves a `pixMap` in the picture. Have that routine write the data to disk. This is equivalent to the `putPicData` bottleneck.)

Since the picture is not known (since there's no way to know how big it will be) so we simply skip the `picSize` field and put out the `picHeader`. `CreatePICT2` creates version `$02FF` pictures, while `OpenPicture` expects version `$FFFF`. This version of the header tells QuickDraw to use `OpenPicture`, available originally in 32-Bit QuickDraw

PostScript Kid") is
of Apple's System
re he's working on
software for the
going backward, he
pending time with
quite forgiven
ze to death in the
and writing zany

new lyrics to classic tunes (his "Working in the
Print Shop Blues" is well known to his coworkers).
Most of all, Dan enjoys building and flying model
airplanes, and he's recently joined the
competition circuit. In fact, when asked what he'd
really like to do with his life, Dan replies:

```
sunny { { { { hours 8 { flying } for } rather_be }  
dayforall } } if *
```

BUT DON'T I NEED A LICENSE TO DO TH

The reason Apple doesn't want developers modifying data structures is that it makes it hard to change them in the future. For example, early Macintosh programs locked handles by manually setting the high bit of the handle rather than calling HLock. This caused numerous compatibility problems when the 32-bit-clean Memory Manager was introduced.

So what gives? What if Apple changes OpenPicture so that it creates a totally different data format—won't the manually created pictures break?

Calm down, because the answer is no. The difference between creating your own pictures and directly modifying other data structures is that Apple can't make the current picture data format obsolete without invalidating users' data that exists on disk. Just as you can still call DrawPicture on version 1 pictures and everything works, you will always be able to call DrawPicture on existing version 2 pictures, regardless of the format of pictures created in the future.

version 1.2 and in Color QuickDraw in System 7, still creates but the header version is now \$FFFE and contains picture res

In addition, the bounds of the clipping region of the current picture. Without this, the default clipping region is wide open. QuickDraw have trouble drawing pictures with wide-open cl

Next we put out an opcode—either \$98 (PackBitsRect) or \$99 depending on whether the pixMap is indexed or direct. Then dstRect, and mode are put in the picture using the (are you re PutOutPixMapSrcRectDstRectAndMode routine. Finally, either PutOutPackedDirectPixData or PutOutPackedIndexedPixData the pixel data.

There's an important difference between indexed and direct pix baseAddr field is skipped when putting out indexed pixMaps for direct pixMaps. This is done because machines without su pixMaps (opcode \$9A) read a word from the picture, skip tha

IS?

One possible pitfall is that you might create a picture with subtle compatibility risks that draws on the existing system software but breaks at some future date. To minimize the chances of such an occurrence, you should compare the pictures you generate with those that QuickDraw generates in identical circumstances. You must be able to account for any and all differences.

Creating your own pixMaps (as our example code does) is definitely in the gray area between risky and outright disastrous behavior, and you shouldn't do it. Then why would an article written by two upstanding citizens do such a thing? The answer is that the pixMaps used by this code are kept private; they're never passed as arguments to a trap. We could just as easily have called them something else, but pixMaps work for what we're doing, so we used them. If you want to pass a pixMap to a trap, you can generate it using the NewPixMap call (not available on machines with original QuickDraw) or let other parts of Color QuickDraw, like OpenCPort, generate it.

version \$02FF pictures,
(resolution information.)

port are put in the
n, and some versions of
ipping regions.

A (DirectBitsRect),
the pixMap, srcRect,
ready for this?)
ther
ata is called to put out

pixMaps here. The
and is set to \$000000FF
upport for direct
t many bytes, and

continue picture parsing. baseAddr, the number of ends the picture playback

An interesting fact buried the value of packType. All to be unpacked. Thus, yo the pixMap should get wh "Compression" (*develop* Iss compression schemes use packing schemes lose no i in the lead article in *devel* sacrifice some image qual

Anyway, these routines su any pixMap with rowByte bit direct pixMaps with ro support 16-bit pixMaps.

Finally, the end-of-pictur actually used.

```
PicHandle CreatePICT2(
    short mode)
{
    PicHandle    myPic;
    short        myRowByte;
    short        *picPtr;
    short        iii;
    long         handleSiz

#define CLIPSIZE 12
#define PIXMAPPRECSIZE
#define HEADERSIZE 40
#define MAXCOLORTABLES
#define OPCODEMISCSIZE
#define ENDOFPICTSIZE
#define PICSIZE PIXMAP
        ENDOFPICTSIZE + OPC

    myRowBytes = srcBits
/* Allocate worst-case
    myPic = (PicHandle)
        ((myRowBytes/127
        - srcBits->bounc
```

When such a machine encounters the \$000000FF bytes skipped is \$0000 and the next opcode is \$00FF, which is a graceful exit from a tough situation.

in the PutOutPixMapSrcRectDstRectAndMode routine is that all in-memory pixMaps (that aren't in a picture) are assumed to be 1-bit. You can set the packType field to specify the type of packing to be used when put in a picture. "The Low-Down on Image Compression" (Issue 6, page 43) gives details of the different pixMap packing methods used by QuickDraw. Note that all of QuickDraw's existing packing methods are of low image quality. QuickTime (the new INIT described in detail in *MacWorld* Issue 7) adds many new packing methods, most of which are designed to achieve much higher compression.

QuickDraw only supports the default packing formats: 1 (or unpacked) for 1-bit maps less than 8 rows, 0 for all other indexed pixMaps, and 4 for 32-bit maps with rowBytes greater than 8. Note that these routines do not

put the opcode out and the handle is resized to the amount

```
PixMap *srcBits, Rect *srcRect, Rect *dstRect,
```

```
};
```

```
};
```

```
50
```

```
SIZE 256*8+8
```

```
2+8+8+2 /* opcode+srcRect+dstRect+mode */
```

```
2
```

```
PRECSIZE + HEADERSIZE + MAXCOLORTABLESIZE + \  
CODEMISCSIZE + CLIPSIZE
```

```
cs->rowBytes & 0x3fff;
```

```
memory scenario using PackBits packing. */
```

```
NewHandle(PICSIZE + (long)
```

```
7)+2+myRowBytes)*(long)(srcBits->bounds.bottom  
ds.top));
```

```

if(!myPic)
    return(0);

/* Skip picSize and put out picFrame (10 bytes). */
picPtr = (short *) (((long)*myPic) + 2);
*picPtr++ = dstRect->top;
*picPtr++ = dstRect->left;
*picPtr++ = dstRect->bottom;
*picPtr++ = dstRect->right;

/* Put out header (30 bytes). This could be done from
taken from an existing picture. */
*picPtr++ = 0x11;      /* Version opcode. */
*picPtr++ = 0x2ff;    /* Version number. */
*picPtr++ = 0xC00;    /* Header opcode. */
*picPtr++ = 0xFFFF;  /* Put out PICT header vers
*picPtr++ = 0xFFFF;

/* The rest of the header is ignored--0 it out. */
for(iii = 10; iii > 0; iii--)
    *picPtr++ = 0;    /* Write out 20 bytes of 0.

/* Put out current port's clipping region. */
*picPtr++ = 0x01;    /* Clipping opcode. */
*picPtr++ = 0x0A;    /* Clipping region only has
*picPtr++ = (**thePort->clipRgn).rgnBBox.top;
*picPtr++ = (**thePort->clipRgn).rgnBBox.left;
*picPtr++ = (**thePort->clipRgn).rgnBBox.bottom;
*picPtr++ = (**thePort->clipRgn).rgnBBox.right;

HLock(myPic);
if(srcBits->pixelType == RGBDirect)
{
    /* Must be 32-bits/pixel */
    /* Put out opcode $9A, DirectBitsRect. */
    *picPtr++ = 0x9A;
    *picPtr++ = 0; /* BaseAddr for direct pixMaps
    *picPtr++ = 0xFF;
    PutOutPixMapSrcRectDstRectAndMode(srcBits, &picPtr,
        dstRect, mode);
    if(PutOutPackedDirectPixData(srcBits, &picPtr)
        goto errorExit; /* Nonzero indicates an
}
else
{
    /* Put out opcode $98, PackBitsRect. */
    *picPtr++ = 0x98;

```

a resource or

ion. */

*/

bounds rectangle. */

is 0x000000FF. */

cPtr, srcRect,

)
error. */

```

        PutOutPixMapSrcF
            dstRect, mode
        if(PutOutPackedI
            /* Nonzero in
            goto errorEx:

    }
    HUnlock(myPic);

/* All done! Put out e
    *picPtr++ = 0x00FF;

/* Size handle down to
    handleSize = (long)
    SetHandleSize(myPic)
    /* Write out pictur
    *((short *) *myPic)
    return(myPic);

errorExit:
    DisposHandle(myPic)
    return(0);
}

```

Just remember that it's not
The reason is that although
it's not a persistent data st
application.

The subroutines the Crea
CreatePICT2 are on the

PROCESSING COI

The remainder of this art
(black-and-white) devices

There are many techniqu
color resources are limite
routines for determining
color space. For example,
Picture Utilities can tell y
CreatePICT2 routine jus
using the Picture Utilities

```
RectDstRectAndMode(srcBits, &picPtr, srcRect,  
e);  
IndexedPixData(srcBits, &picPtr))  
ndicates an error. */  
it;
```

```
nd-of-picture opcode, $00FF. */
```

```
the amount actually used. */  
picPtr - (long) *myPic;  
, handleSize);  
e size. */  
= (short) handleSize;
```

```
;
```

It is not advisable to pass a `pixMap` you create yourself to a trap. Although it's unlikely, the format of a `pixMap` could change (since its structure, as a picture is); this would then break your

code that uses the `atePICT2` routine calls as well as some sample code that uses the *Developer CD Series* disc.

COLOR IMAGES FOR DISPLAY

This article focuses on processing color images for display on 1-bit monitors, both monitors and laser printers.

There are several techniques for representing a full-color image on a monitor when the available color palette is limited. The Picture Utilities Package (new in System 7) offers several techniques for selecting optimal colors to use when displaying a `pixMap` in a limited color palette. For example, if you want to display a 32-bit image on an 8-bit monitor, the package will select the 256 best colors to use to display the image. The package also provides a routine that is described here that creates a picture that you can legally analyze and display on a monitor.

You can also use the techniques of thresholding and of dithering in three varieties: error diffusion, ordered, and random. Ordered dithering, as halftoning, is particularly useful for producing images to be printed on a printer. We'll examine each of these techniques in turn.

USING A 50% THRESHOLD

The first technique that leaps to mind when one is faced with a color picture on a 1-bit screen is to convert each color to a luminance value and then compare that value to a threshold value to determine whether or not to set the corresponding bit. It turns out that green contributes the most to the luminance and blue the least. Red, green, and blue contribute approximately 30%, 59%, and 11% to the luminance. Thus, our formula to convert an RGB value to a luminance value is

$$\text{Luminance} = (30*\text{RED} + 59*\text{GREEN} + 11*\text{BLUE})/100$$

If the resulting luminance is 128 (50% of 256) or greater, the bit is set to white; otherwise it's set to black. This technique produces the result shown in Figure 1: a gray gradient and a lovely picture of one of the authors. Note that the gray occurs at the source pixel resolution. Thus, even though the original picture produce Konenna is 300 dpi, the thresholded picture appears to have a resolution of 150 dpi. In contrast, the techniques of error-diffusion dithering and halftoning discussed in the following pages occur at the destination device resolution.

The results shown in Figure 1 are far from ideal. The gray gradient is a solid black rectangle beside a white rectangle, and the picture of Konenna is completely devoid of detail.

USING ERROR-DIFFUSION DITHERING

The major problem with the threshold algorithm is that a great deal of information is thrown away. The luminance is calculated as a value between 0 and 255, but the information we use is whether it's 128 or greater.

An easy fix is to preserve the overall image lightness by maintaining a running error and then passing the error onto neighboring pixels. Both original and QuickDraw have dithering algorithms built in for precisely this purpose. The original is true—while a dither flag cannot be passed explicitly to any one pixel, a picture containing a color bit image created using dither mode will be dithered. QuickDraw machine will dither when drawn with original QuickDraw. The error is calculated as

$$\text{Error} = \text{Requested Intensity} - \text{Closest Available Intensity}$$

For a black-and-white destination, the closest available intensity is either 0 (black) or 255 (white). The requested intensity is the luminance of the color pixel.

ing, of which there are
dithering, also known
e printed on a laser

displaying a color
ce and then use a
ponding pixel. It turns
e contributes the least.
d 11%, respectively, to
o a luminance becomes

pixel is set to white;
s shown in Figure 1 for
te that thresholding
output device used to
to be 72 dpi. In
oning discussed on the

adations end up as a
onenna, while still cute,

at deal of information is
0 and 255, but the only

aining an error term
iginal and Color
his purpose. (Yes, it's
iginal QuickDraw trap,
ode on a Color
uickDraw.) The error is

ity is either 0 (black) or
current pixel plus some




Figure 1

Gray Gradations and Kone

part of the error term of
among all surrounding p
uses a shortcut. In origin
scan lines and to the left
technique, except it push
to the pixel immediately
QuickDraw at monitor r

This form of dithering is
each pixel is thresholded
across the image in some
produces very pleasing re
accurately rendering a sin
this; laser printers are no
laser printer, a different t

USING ORDERED DITHI

There are two kinds of la
printer (such as some of t
process) starts the image



enna Pictured Using 50% Threshold

surrounding pixels. Ideally, the error term is spread evenly
pixels. But to maintain acceptable performance, QuickDraw
al QuickDraw, the error term is pushed to the right on even
on odd scan lines. Color QuickDraw uses the same
es only half the error to the left or right, and the other half
below. The result of using this technique in Color
resolution for the two test images is shown in Figure 2.

normally referred to as error diffusion. That is to say that
at 50%, but the error incurred in that process is distributed
manner, thus minimizing information loss. Error diffusion
results when the device being drawn onto is capable of
ngle dot at the image resolution. Monitors are quite good at
t. If you want your application's output to look good on a
technique is called for.

ERING (HALFTONING)

user printers: write-white and write-black. A write-white
the high-end Linotronic printers that use a photographic
out black and uses the laser to turn off pixels. A write-black




Figure 2

Gray Gradations and Konenna Dithered at Monitor Resolution

printer (such as Apple's LaserWriter) starts the image out white with the laser. Since the pixels are thought of as being square or round, neither process can accurately turn on or off single pixels.

Generally, the circle generated by the laser beam is slightly bigger than the computer "sees" it, to guarantee that all space is covered (the effect of this with a write-black printer is that the black dots touch the individual pixels, causing any 1-bit image drawn at device resolution to appear dark. The effect with a write-white printer is that the black dots do not touch the individual pixels, causing any 1-bit image drawn at device resolution to appear too light. If the area of the circle is 20% greater than the area of the square, the percentage of unwanted toner, or error, for a single pixel is 20%.




Figure 3

A Laser's Idea of a Square Pixel



te and turns on pixels
and the laser beam is
els.

gger than the pixel as
(see Figure 3). The
end to be bigger than
resolution to appear too
ots tend to be smaller
evice resolution to
he individual pixel, the
%.

Because the error is introduced by two or more pixels are dithered, the error is reduced to the perimeter of the cell. For a single pixel is 20%, two pixels is 40% error, and four pixels in a cell is 80% error.

Ordered dithering, or halftone, is a technique for clumping pixels. Pixels are dithered in a regular order and the luminance is averaged over a larger area. This way that clumps of pixels are dithered. This allows us to minimize the error. The dithering order is determined by the frequency of the things get deep, so put on a regular grid. The following to use the sample.

About the dither matrix
The dither matrix is a 2x2 matrix of gray or primary colors, which effectively lower the device resolution. The shades that we can print. For a single pixel on the page, we've lost the resolution. Now we have 2^4 or 16 different shades. The dither pattern has anywhere from 0 to 100%. In fact, for a 2x2 matrix, 0%, 25%, 50%, 75%, and 100% are the cell. The dither matrix is a 2x2 matrix. The dither matrixes represent the five possible dither matrixes using the algorithm. The *CD Series* disc has a complete list of dither matrixes.

If we construct a matrix with a threshold of 50, we use (2x2 for the described dither matrix). For a threshold of 100, we can use this matrix. The dither matrix in the pattern corresponds to the threshold. For gray, we turn on all the dither matrixes or equal to 50. The position of the dither matrix pattern, as shown in Figure 16.

The dither matrix is used to dither the threshold described earlier. The dither matrix element has a value of 50. The dither matrix where (x, y) is the device resolution. The dither matrix.

It turns out that the spatial frequency of the dither matrix. For resolution, the dither matrix. For resolution, the dither matrix.

duced only at the black/white boundaries, it's reduced when drawn next to each other. Then the percentage of error is of the pixel group. So in the case where the error for a pixels drawn next to each other would have only a 15.5% square would have only a 10.25% error in the area covered.

ftoning, minimizes the dot-to-pixel error just described by re turned on and off in a specific order in relation to each of the source image. The order can be specified in such a next to each other are turned on as the luminance decreases. ze the effects of the laser printer's dot-to-pixel error. The hat's known as a dither matrix. (*Warning:* From here on out, n your waders. You don't really need to understand all the ble code we provide.)

x. With a dither matrix, to render intermediate shades of re sacrifice spatial resolution for shading—that is, we ce's dots-per-inch rating while increasing the number of For example, if we use a 2x2 cell of 300-dpi dots for every owered the spatial resolution of the device to 150 dpi but we ent patterns to choose from for each one of the pixels. Each m 0 to 4 of the 300-dpi dots blackened, or a density between he 16 possible patterns there are only five possible densities: d 100%, corresponding to 0, 1, 2, 3, and 4 dots blackened in ix determines which five of the possible patterns to use to e densities. It's left to you as an exercise to generate these thm we provide below. (The sample code on the *Developer* only useful example.)


with the same dimensions as the dot cell that we're going to d case) so that the matrix contains the values 25, 50, 75, and ix to determine each of the five possible patterns. Each dot ds to a position in the matrix. To generate a pattern for 50% ots in the pattern with corresponding matrix values less than ion of the values in the matrix determines the shape of the re 4.

l to render an image in much the same way as the 50% er. In fact, that process uses a 1x1 dither matrix whose single %. The dither matrix is sampled with $(x \bmod m, y \bmod n)$, pixel location and (m, n) is the width and height of the dither

al resolution of the device isn't really reduced by the size of gions that are all black, for example, the resolution remains

25	50
75	100

Dither matrix



50% gray pattern

Figure 4

A 2x2 Dither Matrix

the device resolution. Each pixel in the device is still sampled source image.

The basic algorithm for doing an ordered dither of an image following:

For all device pixels x, y :

- $s_x, s_y = \text{transform}(x, y)$ where transform maps device pixel coordinates to source pixel coordinates
- If $\text{sourceLuminance}(s_x, s_y) > \text{ditherMatrix}[x \bmod m, y \bmod n]$ then $\text{device-dot}(x, y) = \text{black}$

The code on the *Developer CD Series* disc is an elaboration on

As stated before, the position of the various values in the dither patterns that various luminances generate. A general way to spot use a spot function, as the PostScript interpreter does. If the matrix is thought to be a continuous space whose domain is 0 to 1 in both directions, $\text{spot-function}(x, y)$ will return some value that ultimately into a luminance threshold in the matrix. If the desired pattern from the center as the luminance decreases (known as a cluster), $\text{spot-function}(x, y)$ is simply the distance from (x, y) to the center of the dither matrix would be generated from the spot function as follows:

```

for  $i = 1$  to  $m$ 
   $x = i/m$ 
  for  $j = 1$  to  $n$ 
     $y = j/n$ 
     $\text{matrix}[i, j] = \text{spot-function}(x, y)$ 

```

The result of this process is that the matrix contains the spot values we really want in the matrix are threshold values for the luminance.

back to a pixel in the

onto a page becomes the

ixel

mod n],

this basic algorithm.

er matrix determines the
pecify this order is to
rectangle of the dither
-1 in the x and y
ately can be converted
n is a dot that grows
ered-dot halftone), *spot-*
the cell (0.5, 0.5). The
ollows:

function's results. What
nance. The spot

function result is converted to a one-dimensional array A , generated from 1 to $m*n$. Then, repeated until the desired threshold matrix is reached, the code uses numbers that are percentages assume that the relationship between image and

Ordered dithering is generated by a matrix of size n is the number of cells (or patterns) produced patterns are oriented at an angle. For example, the frequency (if the angle is 0°).

Because of the way our brain perceives color (not at 45° angles), it's better to use the dither matrix itself is rotated. The dither matrix in such a way to achieve the effect of being rotated requires us to "tile" an area enclosing a part of the rotated image. This rectangle must be

Suppose we want to halftone an image at an angle of 45° . At 0° there are 5 possible shades of gray. For a 5×5 cell, we approximate the desired area by the vectors $(0, 5)$ and $(5, 0)$ by 45° and $(-4, 4)$. Since the magnitude of the frequency achieved will be 5 , the angle is due to the need to

Here's the basic algorithm:

1. The halftone cell is defined by the vectors (x_1, y_1) and (x_2, y_2) .
2. A , the area of the required dither matrix is $A = x_1 * y_1 + x_2 * y_2$. The vertical dimension is $Q = GCD(x_2, x_1)$.
3. For every point in the image we want to find its position in the repeated halftone pattern. (See Figure 6.) Calculate the position of the point in the halftone pattern.

The source of step 2 in the algorithm is from "An Optimum Algorithm for Halftone Dithering for Displays and Hard Copies" by Robert F. Holladay, from the *Proceedings of the International Information Display*, Vol. 21, 1981.

ed as follows: Treating the dither matrix as a one-
erate a sort vector V such that $A[V[i]]$ is sorted as i goes
placing all of the values in A with $V[i] * 100/(m*n)$ will yield
trix, with each value being a percentage of luminance. (The
re more computer-friendly than percentages.) These
the device is capable of accurately rendering a single pixel.
ed by a gamma function to more accurately produce a linear
ge luminance and pixel density.

erally done at a specific angle and frequency. The frequency
dither matrixes) per inch and the angle refers to how the
ented with respect to the device grid. In the preceding
f printing on a 300-dpi device) is 150 cells per inch and the

rains work (our eyes tend to pick up patterns at 90° angles
desirable to orient these patterns at arbitrary angles. Since
never rotated with respect to the device, we must generate
a way that it contains enough repetitions of the rotated cell
eing rotated itself. In other words, because a square device
ea with 0° rectangles, we need to find a 0° rectangle
tated pattern that forms a repeatable tile. For some angles of
ay be much larger than the pattern itself.

one to a 300-dpi device at a frequency of 60 cells per inch
, the dither matrix would be 5x5 (300/60), yielding 26
However, as Figure 5 illustrates, we need an 8x8 matrix to
angle. These dimensions are found by rotating the vectors
and pinning them to integers, yielding the vectors (4, 4) and
ide of the vector (4, 4) is $4*\sqrt{2}$, the actual halftone
e $300/(4*\sqrt{2})$, around 53. The error in frequency and
o pin the vectors to integer space.


n for computing the dither matrix:

is specified by the parallelogram composed of
) and (x_2, y_2) and based at $(0, 0)$.


modified halftone cell, is $(x_1*y_2) - (x_2*y_1)$. For the
atrix, the horizontal dimension is A/P and the
n is A/Q , where $P = \text{GCD}(y_2, y_1)$ and

a the matrix, which is in (x, y) orthogonal space,
s relative position in the space of one of the
cells, defined by the vectors (x_1, y_1) and (x_2, y_2) .
all this point (u, v) . The transformation is

above algorithm is
alftone Generation
' by Thomas M.
s of the Society for
No. 2, 1980. •



Desired matrix shape



Actual larger matrix

Figure 5

Approximating the Desired Angle

$u = A*x + B*y, v = C*x + D*y$. Since the point (x_2, y_2) in is the point $(1, 0)$ in halftone cell space and the point $(0, 1)$ in halftone cell space, the coefficients A, B are found by solving the following simultaneous linear

$$\begin{aligned} A*x_1 + B*y_1 &= 0 \\ C*x_1 + D*y_1 &= 1 \\ A*x_2 + B*y_2 &= 1 \\ C*x_2 + D*y_2 &= 0 \end{aligned}$$

We compute the dither matrix in the rotated case as follows:


For each position in the matrix (i, j) :

- Get (x, y) the center of the matrix point (i, j)

$$\begin{aligned} x &= i + 0.5 \\ y &= j + 0.5 \end{aligned}$$

- Transform (x, y) to a point in halftone cell space (u, v)

$$\begin{aligned} u &= A*x + B*y \\ v &= C*x + D*y \end{aligned}$$



x repeated four times

(x, y) space
 (x_1, y_1) is the
 P , C , and D
equations:




Figure 6

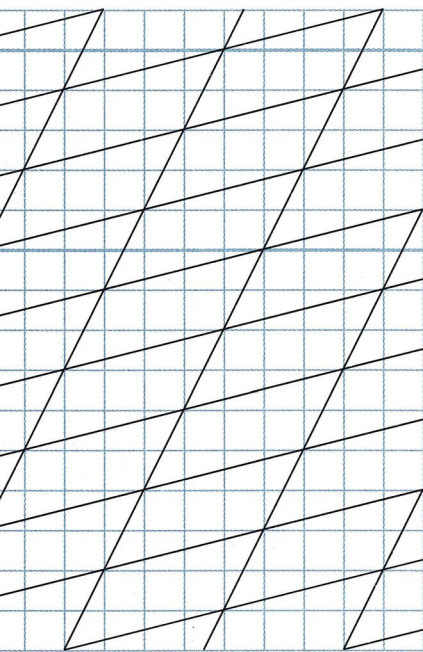
Transforming a Halftone Cell

u and v now express vectors. Therefore position as if the $(0, 0)$ cell.

- $Z = \text{spot-function}(u, v)$
- Find the index of the cell that $u = x, v = y$. If Z is a table. (Note that ϵ is an allowable epsilon of error.)
- $\text{matrix}[i, j] = \text{index}$

Find the order of recurrence (described earlier in chapter 2). Reassign values of matrix

Figure 7 shows our example of luminances, filled in. A lu



(x_2, y_2)

$$A * x_1 + B * y_1 = 0$$

$$C * x_1 + D * y_1 = 1$$

$$A * x_2 + B * y_2 = 1$$

$$C * x_2 + D * y_2 = 0$$

$$A = \frac{y_1}{x_2 * y_1 - x_1 * y_2}$$

$$B = \frac{x_1}{-x_2 * y_1 + x_1 * y_2}$$

$$C = \frac{y_2}{-x_2 * y_1 + x_1 * y_2}$$

$$D = \frac{x_2}{x_2 * y_1 - x_1 * y_2}$$

||

Express the point (x, y) as multiples of the two cell
 e, the fractional parts of u and v represent the
 particular halftone cell at the point (x, y) were the

$$u - \text{floor}(u), v - \text{floor}(v)$$

the record (containing fields $x, y,$ and Z) such
 the record doesn't exist, enter u, v, Z into the
 the equality between $[u, v]$ and $[x, y]$ requires an
 difference to account for fixed-point round-off

ords sorted by values of Z ; store order in sort vector
 connection with converting the spot function result).
 matrix based upon sort vector.

ple matrix with values from 0 through 255, representing
 minance from an image with this range could be sampled

directly against the matrix. The values in this matrix are those used for a 300-dpi, 60-line-per-inch, 45° halftone. As in Figure 7, the matrix is repeated four times for the sake of clarity, with the 45° halftone pattern. The position of any particular number in the matrix relative to the center corresponds exactly to the relative position of that same number in a 45° rotated halftone cell is compared to an unrotated dither matrix.

7	111	183	239	231	191	71	15	7	111	183	239	231	191	71	15
87	47	135	175	199	159	55	127	87	47	135	175	199	159	55	127
215	143	39	79	119	63	151	207	215	143	39	79	119	63	151	207
247	223	103	31	23	95	167	255	247	223	103	31	23	95	167	255
231	191	71	15	7	111	183	239	231	191	71	15	7	111	183	239
199	159	55	127	87	47	135	175	199	159	55	127	87	47	135	175
119	63	151	207	215	143	39	79	119	63	151	207	215	143	39	79
23	95	167	255	247	223	103	31	23	95	167	255	247	223	103	31
7	111	183	239	231	191	71	15	7	111	183	239	231	191	71	15
87	47	135	175	199	159	55	127	87	47	135	175	199	159	55	127
215	143	39	79	119	63	151	207	215	143	39	79	119	63	151	207
247	223	103	31	23	95	167	255	247	223	103	31	23	95	167	255
231	191	71	15	7	111	183	239	231	191	71	15	7	111	183	239
199	159	55	127	87	47	135	175	199	159	55	127	87	47	135	175
119	63	151	207	215	143	39	79	119	63	151	207	215	143	39	79
23	95	167	255	247	223	103	31	23	95	167	255	247	223	103	31

Figure 7

Our Example Matrix With Luminance Values Filled In


This particular example leads us to some other interesting points: that QuickDraw patterns are 8x8 matrixes, just like our example; that we can halftone other QuickDraw primitives besides pixMaps with a pattern on a non-PostScript device (provided that pattern stretching is disabled by setting the bPatScale field in the print record to 0) and achieve a look similar to what a PostScript device would give us.

that would actually be
 e 5, the matrix is
 e cells overlaid. The
 45° cell it falls in
 er in any of the other
 reated with an

71	15
55	127
151	207
167	255
183	239
135	175
39	79
103	31
71	15
55	127
151	207
167	255
183	239
135	175
39	79
103	31

possibilities. It turns out
 ble. This means that we
 en drawing to a 300-dpi
 abled, by setting the
 milar to what a

Here's how. Suppose we w
 from 0 to 255. We simply
 correspond to the cells in
 pattern (shown in Figure 8
 primitive to get the halfto
 aligned to the origin of th
 will not generate undesira
 nature of the clustered do
 extent possible at the reso



8x8 pattern

7	111
87	47
215	143
247	223
231	191
199	159
119	63
23	95
7	111
87	47
215	143
247	223
231	191
199	159
119	63
23	95

Figure 8

Pattern for an Image With a

Figure 9 shows the gray g
 error-diffusion dithering o
 difference in print quality
 "Printing: Ideal Versus Re


want to paint a region with a luminance of 150 on the scale create a QuickDraw pattern in which all of the 1 bits the 8x8 matrix that are greater than or equal to 150. This 8) can then be used to paint any region or other QuickDraw effect. Furthermore, because QuickDraw patterns are drawn on the graffPort, separate objects drawn touching one another do not have visible seams, even when drawn with different shades. The QuickDraw pattern is such that gradations appear continuous to the resolution of the device.

183	239	231	191	71	15	7	111	183	239	231	191	71	15
135	175	199	159	55	127	87	47	135	175	199	159	55	127
39	79	119	63	151	207	215	143	39	79	119	63	151	207
103	31	23	95	167	255	247	223	103	31	23	95	167	255
71	15	7	111	183	239	231	191	71	15	7	111	183	239
55	127	87	47	135	175	199	159	55	127	87	47	135	175
151	207	215	143	39	79	119	63	151	207	215	143	39	79
167	255	247	223	103	31	23	95	167	255	247	223	103	31
183	239	231	191	71	15	7	111	183	239	231	191	71	15
135	175	199	159	55	127	87	47	135	175	199	159	55	127
39	79	119	63	151	207	215	143	39	79	119	63	151	207
103	31	23	95	167	255	247	223	103	31	23	95	167	255
71	15	7	111	183	239	231	191	71	15	7	111	183	239
55	127	87	47	135	175	199	159	55	127	87	47	135	175
151	207	215	143	39	79	119	63	151	207	215	143	39	79
167	255	247	223	103	31	23	95	167	255	247	223	103	31

Halftone dots created by repetition of 8x8 pattern

Luminance of 150

gradations and Konenna printed on a laser printer using compared with halftoning using the 8x8 matrix. The difference is radical. For more commentary on this difference, see the next page.




Gray gradations dither



Konenna dither

Figure 9

Gray Gradations and Konenna Dithered and Halftoned at Laser P



Gray gradations halftone



Konenna halftone

Printer Resolution

PRINTING: IDEAL VERSUS REAL

We've already talked about the error introduced in printing by the fact that the laser beam is round while the pixel is square. Many other factors also can make the transfer of toner to paper deviate from the ideal. Sources of error include differences in inks, papers, printer drums, and even humidity. Additionally, a printer's behavior changes over time as the drum wears. Compensating for all these factors to achieve ideal images would require constant calibration and recalibration of the printer.

An error appears most pronounced in the final print when imaging directly at device resolution, as Figure 9 shows. Halftoning hides much of this error and produces reasonably uniform results among printers with varying degrees of error.

The tonal reproduction curves (known as TRC or gamma curves) shown in Figure 10 indicate the gray levels produced by the Apple LaserWriter when dithering and halftoning. Note that with dithering, the measured

About the code. And no our sample code is pixel-b Thus, the performance is s optimized version of this c scan-line rather than a pix only supports input pixMa accept pixMaps of other d

The first routine we need current pixel. The LUMV 255 using the 30%-59%-

```
long LUMVAL(Ptr pPixel
{
    long    red, g

    if (pMap->pixels
        red = (long)
        green = (long)
```

luminance of an image remains dark much longer than with halftoning as requested luminance increases, due to the error when each pixel is printed. Of particular interest is the point on the dither curve right at 50% luminance. The measured luminance is actually darker than when 44% luminance is requested. The reason is that with a 50% dither, every other pixel is drawn, maximizing the effect of the laser error.

While the TRC curve for the halftone print doesn't match the ideal curve, it's much closer to the ideal than is the dither curve. To get the halftone even closer to ideal, you could adjust the luminance calculation by the amount indicated by the halftone TRC to compensate. Indeed, most image-processing applications perform this TRC adjustment to compensate for the nonlinearities of the output device. See *Designing Cards and Drivers for the Macintosh Family*, Second Edition (Addison-Wesley, 1990) for more information about how gamma correction works on the Macintosh II family for monitors.

...w, about the code. To illustrate the principle of dithering, ...ased—that is, the calculations are done on a pixel basis. ...luggish. A real-world commercial application would use an ...code. One way to do this is to make the routines work on a ...el basis. Also note that the routine that does the halftoning ...aps of 8 or 32 bits. It would be easy to extend the routine to ...epts.

...is one that calculates the luminance given a pointer to the ...VAL routine returns a long luminance in the range of 0 to ...11% formula described previously.

```
, PixMapPtr pMap)
```

```
reen, blue;
```

```
ize == 32) {
```

```
unsigned char)*(++pPixel);          /* Skip alpha,  
                                     get red. */  
y)(unsigned char)*(++pPixel);      /* Get green. */
```




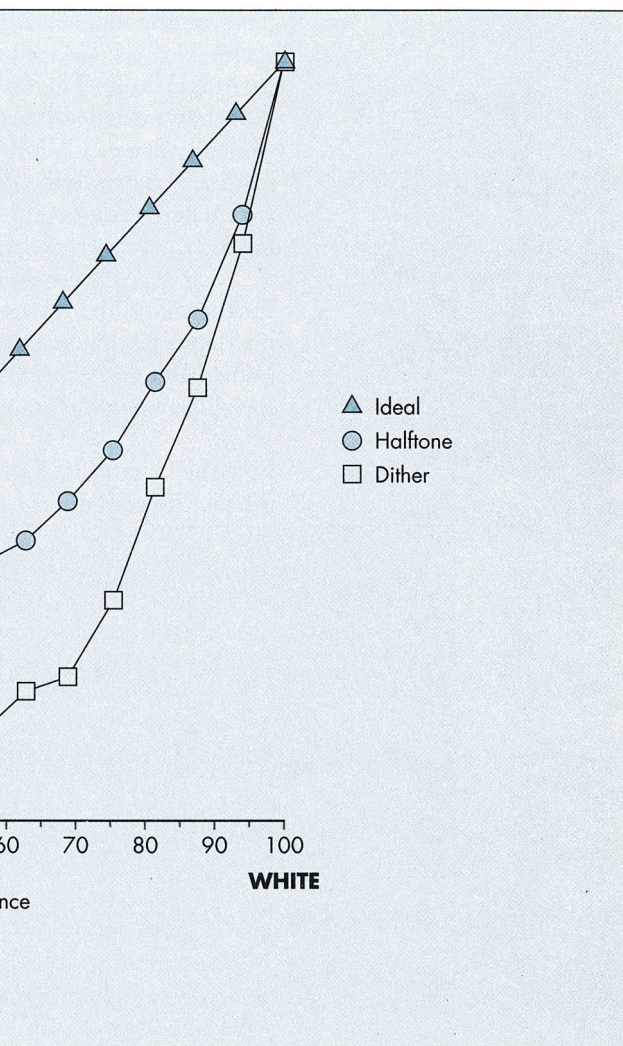
Figure 10

TRC Curves for the LaserWriter

```

    blue = (long)(unsigned char)*(++pPixel);
    return((30 * red + 59 * green + 11 * blue)/100);
} else if (pMap->pixelSize == 8) {
    RGBColor*    theColor;
    theColor = &((*pMap->pmTable)->ctTable[ (
        char)*pPixel ].rgb);
    return( (30 * (theColor->red >> 8) + 59 * (
        theColor->green >> 8) + 11 * (theColor->blue >> 8))/100);
} /* End if */
} /* LUMVAL */

```



```

/* Get blue. */
100);

unsigned
theColor->green >>

```

The routine that actually
than taking a `PixMapPtr` a
`PixMapHandle`. This enabl
did when we called `Create`
example, from a `GWorld`)
routine knows whether it
we created it) or if it must
for the `LockPixels` and `G`

Furthermore, the `Halftone`
`pixMap` is 72 dpi (screen r
(same `hRes` and `vRes`). Yo
`Resolution` parameter, bu

Like the `CreatePICT2` ro
the picture contains a 1-b

The prototype for the `Ha`

```
PicHandle HalftonePixM  
    short Resolution);
```

The source code for the c
disc.

USING RANDOM DITHER

Random dithering is yet a
discussed last, however, b

The method is simple. It's
earlier. The only differenc
values are compared to a r
that the probability of any
proportional to the lumin
point.

This method has three lin
expensive operation that v
except at very high resolu
bad reception on a black-
number generator that's v

Ironically, this least frequ
physical process of photog
composed of pixels. How

does the halftoning is the `HalftonePixMap` routine. Rather than as the `CreatePICT2` routine did, this routine takes a `Boolean` parameter that enables us to pass in either a `pixMap` we create manually (as we did with `CreatePICT2`) or a `PixMapHandle` that `QuickDraw` creates (for `GetPixMapHandle`). We must distinguish which one we pass in so that the routine can access the fields of the `pixMap` directly (which it can if we pass a `pixMap`) or use `QuickDraw` to access the fields. This is relevant only for routines that use `GetPixBaseAddr`.

The `HalftonePixMap` routine assumes the resolution of the source image (the resolution of the `pixMap`) and only supports devices with square pixels. You can pass in the resolution of the destination device in the `device` parameter, but it must be greater than or equal to 72 dpi.

The `HalftonePixMap` routine, `HalftonePixMap` returns a `PicHandle`. In this case, you can display it using `DrawPicture`.

The `HalftonePixMap` routine is

```
HalftonePixMap(PixMapHandle hSource, Boolean qdPixMap,
```

The complete routine can be found on the *Developer CD Series*.

DITHERING

Another kind of dither useful for drawing images. It's useful because of its inherent limitations.

The method is much the same as the 50% threshold method described earlier. The difference is that instead of being compared to 50%, the luminance of the pixel is compared to a random number between 0 and 100%. The effect of this is that the probability of a dot in the device image being turned on is directly proportional to the luminance of the pixel in the source image at the corresponding location.


There are three limitations. First, calculating a random number is an expensive operation that we would not want to do for every device pixel. Second, images dithered in this manner appear very noisy, like a black-and-white TV. And third, this method requires a random number generator that is very good at producing a uniform distribution.

The most commonly used method of dithering most accurately models the halftone process of photography. Photographic film is like laser printing in that it's a continuous tone. However, the pixels are grains of silver rather than toner.

Additionally, there are tens of thousands of grains per inch rather than the 300 per inch we're used to with laser printers. The lower the ASA speed, the higher the grain density.

The place on a film where a photon strikes one of these silver grains on the film is developed (which is why you get negatives). Since photons are really small, the likelihood of a single photon striking one of these grains is very low. However, the brighter the light, the more photons there are, so the probability of striking one of those silver grains increases in proportion to the luminance. Thus, we see how random dithering simulates photographic film.

Figure 11 shows the image of a frog's head produced using a halftone matrix as compared with using a 72-dpi random dither. You can see that the dithered image looks like a really grainy photograph.



Halftoned with 8x8 at 72 dpi

Figure 11

Frog's Head, Halftoned and Randomly Dithered


HASTA LA VISTA, BABY

This article has addressed several issues. First, the problem of rendering on machines with original QuickDraw was overcome by showing how to manually create a PICT, which can then be rendered by calling

her than the 300 dots
rating of the film, the

grains turns black when
photons are really, really,
the grains of silver is
here are; so the
roportion to the
otography.

lftoning with an 8x8
an see that the randomly



Randomly dithered at 72 dpi

saving deep pixMaps
ing you how to
g DrawPicture. Such a

PICT can be exported by QuickDraw machine.

Second, several solutions for black-and-white devices with a 50% threshold or error-diffusion are a way to get around the problem. Random dithering has proven to be producing color images on

Thanks to these techniques, the need not be limited to Color QuickDraw. The necessary code is small (a few lines) and runs very high. Now get to work!

WANT TO READ

If you'd like to delve more into the world of color display, check out the following:

- "An Optimum Algorithm for Color Display" by Thomas M. Holladay, *IEEE Transactions on Computers*, Vol. 21, No. 2, 1980.
- *Digital Halftoning* by Robert F. Schaefer, Ph.D. thesis done at MIT. It's extremely thorough and covers many different ways.
- *Fundamentals of Image Processing* by R. C. Gonzalez (Addison-Wesley, 1987). It's as thorough as Ulichney's.

And then, of course, there are the books:

- *Programming with QuickDraw* by G. Othmer (Addison-Wesley, 1988). It's the Macintosh.
- *Debugging Macintosh* by R. Straus (Addison-Wesley, 1988). It's software, including in

an application so that it can be viewed in color on a Color

to the problem of displaying and printing color images on were discussed. Images can be displayed on screen using a diffusion dithering. Ordered dithering (halftoning) provides a problem of the laser printer's inability to resolve single pixels. Practical limitations but represents yet another alternative for on black-and-white devices.

es, the market for applications that deal with color images color QuickDraw machines and PostScript printers. The (and already written for you) and the gain in functionality is work on those applications!

► MORE?

re deeply into the mysteries of processing color images for following:

thm for Halftone Generation for Displays and Hard Copies" day, in the *Proceedings of the Society for Information Display*, 0.

Robert Ulichney (MIT Press, 1987). This book, based on a MIT, is devoted entirely to discussing halftoning algorithms; it's and includes many example images halftoned in different

Interactive Computer Graphics by J. D. Foley and A. Van Dam (1982). The standard text on computer graphics. Not nearly as y, but has a solid discussion of the basics.


two books all Macintosh programmers should own:

QuickDraw by Dave Surovell, Frederick Hall, and Konstantin esley, 1992). Everything you need to know about graphics on

Macintosh Software with MacsBug by Konstantin Othmer and Jim ley, 1991). Everything you need for debugging Macintosh n-depth discussions of a number of the Macintosh managers.

LOCAL REVIEWERS

Dave Williams •



FORREST TANAKA

GRAPHICS HINTS FROM FORREST

USING THE PALETTE MANAGER OFF-SCREEN

Most people who've done any graphics programming on the Macintosh are aware of the Palette Manager, because it's the documented way to control the on-screen color environment, and perhaps because my cohorts and I in Developer Technical Support keep going on about how right the world would be if everyone used it. In an effort to follow the rules as best they can, some people have taken the Palette Manager so much to heart that they use it not only with windows, but with off-screen cGrafPorts as well—something that isn't heard about very much. Some of these people have concluded that all the features of the Palette Manager apply just as well to off-screen cGrafPorts as they do to windows. Logical enough, right?

Well, that's the kick; whether this is logical or not, the truth is that only a small part of the Palette Manager works with off-screen cGrafPorts. Specifically, the pmCourteous usage mode and the pmWhite and pmBlack usage-mode modifiers work fine when they're used in a palette that's attached to an off-screen cGrafPort, but the pmTolerant, pmAnimated, and pmExplicit usage modes do not. In this column, I'll describe how you can take advantage of the Palette Manager features that work off-screen and how you can simulate the features that don't work.

The pmCourteous usage mode seems pretty useless to a lot of people because it has no effect on the current

FORREST TANAKA has been playing Developer Technical Support as one of the graphics support people for slightly more than two years. "It amazes me still," he says, "that the more you learn about the Macintosh graphics tools, the farther off total understanding seems to be." Outside of DTS, he likes to ride his bike, and uses it to commute the three blocks to his office ("Hey, it's faster than driving the three blocks!"), and he likes to try getting his radio-controlled car to act as if it's actually controlled. •

color environment. But in general, making a palette full of pmCourteous colors is a lot better than hard-coding RGBColors into your code. Instead of hard-coding colors, make a palette of courteous colors—as many entries as you need colors—and save it as a 'pltt' resource. When your application runs, call SetPalette to attach this palette to your off-screen cGrafPort. When you need to use a color while drawing into this GrafPort, pass the desired color's palette index to PmForeColor or PmBackColor, and then draw. This is better than hard-coding colors because you or a software localizer can easily change the colors by changing the 'pltt' resource—no code changes are necessary.

The pmWhite and pmBlack usage-mode modifiers are new with System 7; they let you specify whether you want a particular palette entry to map to white or black in a black-and-white graphics environment. By default, colors whose average color-component value is larger than 32767 are mapped to white and other colors are mapped to black. (If you use RGBForeColor, Color QuickDraw also checks to see whether your specified color is different from your background color but maps to your background color; if so, Color QuickDraw uses the complement of the color you specified so that your drawing is visible over the background.) By specifying that a palette entry is pmCourteous + pmBlack or pmCourteous + pmWhite, you can control which colors map to black and to white when there aren't enough colors available. This applies to palettes attached to off-screen cGrafPorts as well as to palettes attached to windows.

Those are the Palette Manager features that do work off-screen. Now I'll talk about the features that don't and what you can do to get the same effect.

The pmExplicit usage mode is handy when you want to draw using a pixel value without knowing or caring what color that pixel value represents. With this mode you can easily show the colors in a screen's color table, and you can also draw into a pixel image with a specific value even though you specify the color for that value elsewhere.

When you have a palette that's attached to an off-screen `cGrafPort`, `pmExplicit` colors are interpreted as `pmCourteous` colors. Instead of using a palette, you should convert your pixel value to an `RGBColor` and use this as the foreground or background color. Set the current `GDevice` to your off-screen `GDevice` so that the color environment is set; then pass your pixel value to `Index2Color`, which is documented on page 141 of *Inside Macintosh* Volume V. `Index2Color` converts your pixel value to the corresponding `RGBColor`, which you can pass to `RGBForeColor` or `RGBBackColor`, and then you can draw. The result is that your pixel value is drawn into the destination pixel image.

Both the `pmAnimated` and `pmTolerant` usage modes are used to modify the color environment, and both are interpreted as `pmCourteous` when they're in a palette that's attached to an off-screen `cGrafPort`. The most important difference between the two usage modes is in the style of color-table arbitration that they do—`pmTolerant` gives the front window the colors it needs, while `pmAnimated` additionally makes sure that nothing outside the front window is drawn in its colors. Color-table arbitration doesn't apply off screen, so the `pmAnimated` and `pmTolerant` usage modes can be unified into "I want to change my off-screen colors."

Changing the colors in an off-screen color environment means changing its color table; the most straightforward way to do this is to modify the contents of the color table directly. That is, get your off-screen color table's handle and then directly assign new values to the `rgb` fields in its `CSpecArray`. You could also assign a whole new color table to the off-screen environment by assigning the new one to the `pmTable` field of the off-screen `pixMap`. Either way, you have to

30

For more details about changing or replacing off-screen color tables, see the October 1991 version of Macintosh Technical Note #120, "Principia Off-Screen Graphics Environments." •


tell Color QuickDraw what you've done by updating the changed color table's ctSeed field. The next time you draw into your off-screen graphics environment, Color QuickDraw detects your change by comparing the ctSeed of your changed color table against the iTabSeed of the current GDevice's inverse table, and it rebuilds the inverse table according to the changed color table. You can update the ctSeed field by assigning to it the return value of GetCTSeed, which is documented on page 143 of *Inside Macintosh* Volume V. If the 32-Bit QuickDraw extensions are available, you can update a color table's ctSeed simply by passing the color table to CTabChanged, documented on page 17-26 of *Inside Macintosh* Volume VI.

If you have a GWorld and you want to replace its color table, you should call UpdateGWorld, passing it a new color table. UpdateGWorld makes sure that all the cached parts of a GWorld are properly updated, which is tough to do any other way. If you don't pass any flags to UpdateGWorld, it's within its rights to destroy your existing GWorld's image. But if you pass the clipPix or stretchPix flag, UpdateGWorld is obligated to keep your existing image, and it tries to reproduce the existing image in the new colors as best it can.

To wrap up, you can use the Palette Manager with off-screen graphics environments, but you'll only be able to use the pmCourteous usage mode and the pmWhite and pmBlack usage-mode modifiers. But that's not to cast aspersions on these features, because they can be very handy for both on-screen and off-screen drawing. The pmExplicit, pmTolerant, and pmAnimated usage modes don't work for off-screen drawing, but there are easy ways to simulate those features without the Palette Manager and without risking future compatibility.

THE TEXTBOX YOU'VE ALWAYS WANTED

NeoTextBox is an alternative to the standard TextBox. NeoTextBox provides full TrueType features while using only three routines that comprise 100 bytes yet offer a 40% performance improvement in most cases.



**BRYAN K. ("BEAKER")
RESSLER**

In the deepest, darkest corner of the Macintosh toolbox, there's an extremely useful routine called `TextBox`.

```
pascal void TextBox(void *rect, void *font, void *style, void *text);
```

Given a rectangle and some text, `TextBox` will draw the text in the font, style, and size specified.

Anyone who's tried to write a word processor knows how difficult it is to draw text. Perhaps that's why `TextBox` takes a rectangle and a font, creates a new `TERec` with `TENew`, calls `TESetText` to create a temporary `Text` object, and then calls `TEDrawText`. `TextBox` then calls `TEDispose` to dispose of the `Text` object. By using `TEDrawText` to draw the text, `TextBox` avoids doing a lot of overhead.

Despite its pass-the-buck implementation, `TextBox` has several advantages. Perhaps most important, it can draw text in systems like Japanese and Arabic. Another handy side effect is that `TextBox` can draw text in a font that is not supported by `DrawText`, and can therefore be used to draw text in a font designed specifically for drawing. `TextBox` also works well.

BRYAN K. RESSLER, or "Beaker" as he is known at Apple, is one of our twisted engineers who seems to be convinced that anything is possible on a Macintosh. If something has already been done, it can be done better. He received his BSCS from the University of California and wrote commercial MIDI applications before coming to Apple. Beaker wrote many of the programs used for testing TrueType fonts.

tive to the TextEdit utility routine TextBox. -justification capability and the option to use retaining all the advantages of TextBox. The ise NeoTextBox compile to fewer than 900 performance increase over TextBox in common

of the TextEdit chapter in *Inside Macintosh* Volume I, routine called TextBox.

text, long length, Rect *box, short just)

xt, TextBox word wraps the text inside the rectangle, size specified in the current grafPort.

rap text knows that it's not as easy as it first appears. takes the approach it does: to perform its task, TextBox New, sets up the rectangles in the record, and calls ary handle to a copy of the text you provided to EUpdate to wrap and draw the text, and finally TERec. By calling TextEdit to do the text wrapping and g any hard work. Unfortunately, it also incurs quite a bit

ementation, TextBox's use of TextEdit has several ortant, TextBox works correctly with non-Roman script ic without the need for any extra programming. at updates in TextEdit degenerate into calls to e recorded into QuickDraw pictures. TextBox was ng static text items in dialog boxes and performs this

as he's
d software
d that
and if it's
better. He got
ornia, Irvine,
ions before
y of the
onts. When

he's not on a coding frenzy, he writes *noncommercial* MIDI applications, tries to have a life, and keeps a consistent blood-caffeine level so high you need scientific notation to express it. •

So TextBox is great—if you want better performance. You want full justification (instead of using whizzy TrueType callouts in drawing mode). You can't use it when you're drawing. You don't like that 32K text limit. A single call to TextBox. And you know the height it used, and where to draw something below the text. It's a TextBox.

Well, this is your lucky day.

ENTER NEOTEXTBOX

NeoTextBox is the TextBox you've always wanted. NeoTextBox is on the average, considerably more flexible than TextBox.

- NeoTextBox allows for full justification by default (same behavior as TextEdit). It adjusts for character width, ascender, x-height, or descent line; or for line height.
- NeoTextBox provides for full justification.
- NeoTextBox never erases or, if you wish, it can erase.
- NeoTextBox returns the bounding box of the text.
- NeoTextBox can return the position of the last character drawn.
- NeoTextBox can draw the text.

NeoTextBox gives you all the features of TextBox. It is completely compatible with TextBox (just like TextEdit). It's easy to use (just like TextEdit). It's easy to get TextBox-like behavior.

Let's take a look at the package.

```
short NeoTextBox(unsig  
    Rect *wrapBox, sho  
    short *lhUsed)
```

ou're drawing dialog boxes. But you want more. You want
want more flexibility. You want to control line height. You
stead of only left, center, and right alignment). You want to
s when they're available. You want to control the text
stand the way TextBox always erases (and therefore isn't too
ng to printers—it slows printing way down). Yeah, and you
imitation either. You want to word wrap *War and Peace* in a
and you'd like some useful information back, too, like the line
e the last line of text was drawn, so that you can draw
c. And, of course, you want to retain the advantages of

ay.

BOX

ox you've always wanted (and didn't even have to ask for).
verage 33% faster than an equivalent call to TextBox. Plus, it's
le:

ws a line height specification. You can ask for the
avior as TextBox); use variable line height, which
ters that extend beyond the font's standard ascent
r specify a line height in points.

ides left, center, and right alignment and full

r erases the rectangle it's drawing into. It lets you
sh, draw a colored background.

urns the total number of lines in the wrapped text.

return, via VAR parameters, the vertical pen
t line of text and the line height that was used to

l this extra functionality, yet retains the advantages of
language independent and uses the Script Manager heavily
asy to call, and if you don't want all the spiffy new features,
ke behavior with a free performance increase.

parameters for NeoTextBox.

igned char *theText, unsigned long textLen,
rt align, short lhCode, short *endY,

The first two parameters, the `Text` and `textLen`, are analogous length parameters: they specify the text to be wrapped. Note that `Text` is a Pascal string—it's a pointer to the first printable character.

The third and fourth parameters, `wrapBox` (`box` in `TextBox`) and `align`, go back to `NeoTextBox`'s ancestor. Just as in `TextBox`, `wrapBox` specifies the box within which you're wrapping text, and the `align` parameter specifies the alignment. In addition to the standard `TextEdit` alignments `teFlushLeft`, `teFlushRight` (see "Text Alignment Constants for System 7"), and `teFlushDefault`, a new one is defined—`ntbJustFull`. It performs full justification in whatever script is used for the current script.

The fifth parameter, `lhCode`, specifies how the line height is determined. The default line height is derived via a call to `GetFontInfo`. The behavior is the same as `TextBox`. If `lhCode` is less than 0, the line height is determined by the height of the tallest character in the text that's being drawn, extending above and below the baseline (see "SetPreserveGlyph With TrueType"). If `lhCode` is greater than 0, the value of `lhCode` itself specifies the line height. For example, you can draw 12-point text in 16-point lines.

The last two parameters, `endY` and `lhUsed`, are reference parameters. `endY` is used to retrieve the vertical position of the last line of text and the `lhUsed` parameter is used to draw the text, respectively. The `endY` parameter can be used to

TEXT ALIGNMENT CONSTANTS FOR SYSTEM 7

Before System 7, there was a conflict between the names of the text alignment constants and their actual behavior. To help make applications more portable, the new constants for Roman scripts, `teJustLeft` was interpreted as the default text alignment for the current script rather than forcing text to be aligned on the left. For example, on a Hebrew system, a `TextBox` call with a `just` parameter would actually use the default justification for Hebrew, which is `teJustRight`.

To overcome this conflict, new constants were introduced in System 7. See Table 1.

Table 1

Text Alignment Constants

New Constant	Old Constant	Value	Meaning
<code>teFlushLeft</code>	<code>teForceLeft</code>	-2	Align text on the left
<code>teFlushRight</code>	<code>teJustRight</code>	-1	Align text on the right
<code>teFlushDefault</code>	<code>teJustLeft</code>	0	Use conventional justification
<code>teCenter</code>	<code>teJustCenter</code>	1	Center text for the width

to TextBox's text and
that theText isn't a

and align, also hearken
specifies the rectangle
specifies the alignment.
eCenter, and
a new alignment is
r manner is appropriate

erived. If lhCode is 0,
is gives the same
derived by
tend the most above
e Fonts"). Finally, if
e line height. For

meters that allow you
ine height that was
e very useful if you

ITEM 7

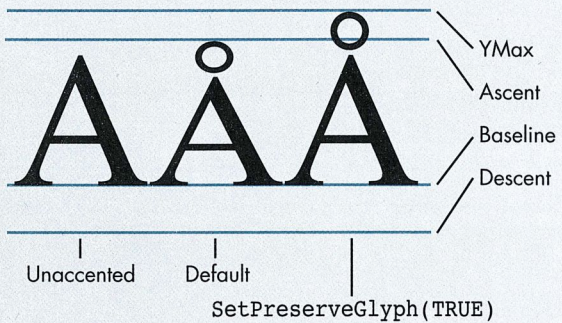
text alignment
compatible with non-
ment appropriate for
ft as specified. For
eter of teJustLeft would
ight.

tem 7, as shown in

he left for all scripts
he right for all scripts
nal alignment for script
all scripts

SETPRESERVEGLYPH WITH TRUETYPE

Before TrueType, all characters in all fonts fit beneath the font's ascent line and above the descent line, like the default characters shown in Figure 1. Bitmapped fonts were drawn so that diacriticals, like the angstrom over the A in Ångström, would fit beneath the ascent line. To do this, the letterform had to be distorted. With the advent of TrueType, this "feature" can be controlled, because TrueType fonts carry outline data that's true to the original design (hence the name TrueType).



Example shown: 72-point Times

Figure 1

How SetPreserveGlyph Affects Line Height


intend to draw anything
text was drawn. To find
negative lhCode, use the
two parameters if you do

NeoTextBox returns the
off because they were be
overflowed wrapBox by
wrapBox.bottom. If you
the height of wrapBox by

FONTS

Since most applications expect characters to fit beneath the ascent line and above the descent line, QuickDraw transforms characters in TrueType fonts to force them within those bounds. To override this transformation and preserve the original glyph shape, use the Font Manager call `SetPreserveGlyph(TRUE)`. After this call, TrueType fonts will be drawn as shown to the right in Figure 1.

Preserving the glyph shape makes it possible to take advantage of `NeoTextBox`'s variable line height feature.



Example shown: 72-point Symbol

below the text, since it tells you exactly where the last line of text is. You can also pass `endY` to find out what the actual derived line height was if you used a `wrapBox` parameter. Pass `nil` for either or both of these last two parameters if you don't want this extra information.

`total` is the total number of lines in the text. That includes lines clipped by the bottom of `wrapBox`. You can tell whether the text is wrapped by checking whether the value returned in `endY` is greater than `wrapBox`. If you want to know how many lines fit in `wrapBox`, simply divide `total` by the value returned in `lhUsed`.

REQUIREMENTS

NeoTextBox uses some advanced Script Manager routines that require System 6 or later. NeoTextBox assumes they're available, so the program checks that it's running on System 6 or later via a Gestalt call.

NeoTextBox requires one global variable, a Boolean named `gHasTrueType` to be set to TRUE if the TrueType trap (\$A854) is available, or FALSE if the development environment provides glue for Gestalt, you can use `Gestalt` to set up `gHasTrueType`:

```
#define kTrueTypeTrap 0x54 /* The TrueType trap number */
#define kUnimplTrap 0x9f /* The "unimplemented" trap number */
long gResponse;

if (Gestalt(gestaltFontMgrAttr, &gResponse) == noErr)
    gHasTrueType = BitTst(&gResponse, 31-gestaltOutlineAttr);
else {
    gHasTrueType = (NGetTrapAddress(kTrueTypeTrap, ToolTrap) &
        NGetTrapAddress(kUnimplTrap, ToolTrap));
}
```

THE BASIC ALGORITHM

NeoTextBox does a lot. But, in order to appease the programmer and avoid work, we allow the Script Manager to do the hard parts. (do full justification in Arabic?) In short, here's how NeoTextBox

1. It saves the current `grafPort`'s clipping region and clips the text we're drawing into.
2. It calculates the appropriate line height with the function `NTBLineHeight`.
3. It calls the Script Manager routine `StyledLineBreak` to find the line-break point in the input text.
4. It draws each line with the function `NTBDraw`.
5. It advances the pen down one line.
6. When there's no more text, it restores the clipping region and returns the appropriate values.

It sounds simple, doesn't it? That's because `StyledLineBreak` does all the work. It knows how to find word breaks in whatever script we're using, and it's smart, too. For instance, in English, it knows that it's OK to break a word if necessary. It uses rules that are provided by the installed script manager. It takes the appropriate actions. Let's take a closer look at the code.

t are available only in
ake sure your main
stalt or SysEnviron

HasTrueType. It should
FALSE if not. If your
use the following lines

```
er */  
ap number */
```

```
Fonts);
```

```
Trap) !=
```

ner's natural desire to
(Do you know how to
ox gets its job done:

to the box

on

find each

on and

oes all the work. It
. StyledLineBreak is
reak a hyphenated word
pt systems, so it always
de.

THE NEOTEXTBOX

The source code for Neo starts in the NeoTextBox we come upon them.

Here's the NeoTextBox c

```
short NeoTextBox(unsigned
    Rect *wrapBox, short
    short *lhUsed)
{
    RgnHandle
    StyledLineBreakCod
    Fixed
    Fixed
    short
    long
    unsigned short
    short
    unsigned short
    long
    unsigned char
    unsigned char
```

Many of these variables a detail later. The most im contains the line break co lineBytes, which are retu the current vertical pen l

GET READY

NeoTextBox, like TextBo it's safest to save the clip width of wrapBox, becau which is used in calls to S appropriate text alignme

```
GetClip((oldClip = Ne
ClipRect(wrapBox);
boxWidth = wrapBox->r
fixedMax = Long2Fix(
if (align == teFlushD
    align = GetSysJust
```

TEXT FUNCTION

TextBox that's shown here is written in MPW C 3.2. We'll use a text function and break out to a couple of utility functions when

declaration and local variables:

```
signed char *theText, unsigned long textLen,  
short align, short lhCode, short *endY,  
  
oldClip;          /* Saved clipping region */  
breakCode;       /* From StyledLineBreak */  
fixedMax;        /* boxWidth in fixed point */  
wrapWid;         /* Width to wrap within */  
boxWidth;        /* Width of box */  
lineBytes;       /* Number of bytes in one line */  
lineHeight;      /* Calculated line height */  
curY;            /* Current vert pen location */  
lineCount;       /* Number of lines we've drawn */  
textRemaining;   /* Number of bytes of text left */  
*lineStart;      /* Pointer to start of a line */  
*textEnd;        /* Pointer to end of input text */
```

are used in the call to `StyledLineBreak`, which is explained in the next section. Important variables to know about here are `breakCode`, which is the code returned by each call to `StyledLineBreak`; `lineStart` and `lineEnd`, which are returned by `StyledLineBreak` to specify a single line; and `curY`, which is the current pen location.

`wrapBox`, clips to `wrapBox`. Since this is a general-purpose routine, we save the clipping region, then restore it at the end. We calculate the clipping region since it's used a lot, and convert it to fixed point as `fixedMax`, and pass `StyledLineBreak` as a VAR parameter. Also, we retrieve the default alignment if the user has requested default alignment.

```
wRgn());  
right - wrapBox->left;  
(long)boxWidth);  
default)  
();
```

DETERMINE THE LINE HEIGHT

Now we need to determine the appropriate line height. NeoTNTBLineHeight to perform this function, passing the text pointer, the wrap rectangle, the caller-specified line height code, and the current vertical pen location. NTBLineHeight calculates and returns the line height and calculates the correct starting pen location. Here's the NTBLineHeight function:

```
unsigned short NTBLineHeight(unsigned char *theText,
    unsigned long textLen, Rect *wrapBox, short lhCode)
{
    short          asc, desc;
    FontInfo       fInfo;
    Point          frac;
    unsigned short lineHeight;

    GetFontInfo(&fInfo);
    if (lhCode < 0) {
        /* lhCode < 0 means "variable line height", so
        /* TrueType font use OutlineMetrics, otherwise
        frac.h = frac.v = 1;
        if (gHasTrueType && IsOutline(frac, frac)) {
            OutlineMetrics((short)textLen, theText, frac,
                &desc, nil, nil, nil);
            lineHeight = MAXOF(fInfo.ascent, asc)
                + MAXOF(fInfo.descent, -desc) + fInfo.leading;
            *startY = wrapBox->top + MAXOF(fInfo.ascent,
                + fInfo.leading);
        } else {
            /* Punt to "default" if we can't use TrueType
            lineHeight = fInfo.ascent + fInfo.descent + fInfo.leading;
            *startY = wrapBox->top + fInfo.ascent + fInfo.leading;
        }
    } else if (lhCode == 0) {
        /* lhCode == 0 means "default line height." */
        lineHeight = fInfo.ascent + fInfo.descent + fInfo.leading;
        *startY = wrapBox->top + fInfo.ascent + fInfo.leading;
    } else {
        /* lhCode > 0 means "use this line height" so w
        lineHeight = lhCode;
        *startY = wrapBox->top + lhCode + fInfo.leading;
    }
    return(lineHeight);
}
```

textBox calls
winter, the text length,
the address of curY, the
returns the line height
TBLineHeight function:

```
, short *startY)
```

```
if it's a */  
use default. */
```

```
, frac, &asc,
```

```
ding;  
asc)
```

```
e. */  
fInfo.leading;  
o.leading;
```

```
fo.leading;  
leading;
```

```
e trust 'em. */
```

```
;
```

Remember, there are three

- Variable line height. This is the first. If the TrueType font is a variable TrueType font, `OTL` height (see “Descent” for a variety of information) and the lowest descent, `asc` and `desc`. Then `asc` or `asc`, and if TrueType isn’t available, we punt to the font, we punt to the
- If `lhCode` is 0, the sum of the ascent, `GetFontInfo` call.
- Finally, if `lhCode` is 0, the line height. In this case, the line height.

Each of the three line heights is used on the line height and width

Back in `NeoTextBox`, we use `lineHeight` and `curY`:

```
lineHeight = NTBLineHeight;
lineCount = 0;
lineStart = theText;
textEnd = theText + textLength;
textRemaining = textLength;
```

DESCENT INTO

Descent is the amount of text that descends below the baseline. When you calculate the number of points below the baseline, in the industry it’s more common to use the term *descent*.

In an attempt to be more precise, we use its descent values as needed. Note the sign of descent.

the possible line height codes:

nt (specified by an lhCode less than 0) is handled
ype trap is available and this particular font is a
OutlineMetrics is called to determine the line
ent Into Hell”). OutlineMetrics can return a
tion, but we really only want the highest ascent
cent, which are returned in the local variables
n we choose whichever is higher, the default
whichever is lower, the default descent or desc. If
available or the particular font isn't a TrueType
ne default line height.

default line height is used. This is defined as the
descent, and line gap (leading) derived by a

is greater than 0, the caller is providing a specific
s case, NTBLineHeight returns lhCode as the

ght calculation methods also figures the correct startY based
apBox->top.

call NTBLineHeight to set up our local variables

```
ight(theText, textLen, wrapBox, lhCode, &curY);
```

```
textLen;
```

```
en;
```

HELL

of space that should be allocated for a font below the text
ll GetFontInfo, the value returned for descent is a positive
y the baseline. Although this is convenient, in the typographic
on to represent descent values as *negative* numbers.

the typographically useful, TrueType's OutlineMetrics call returns
negative numbers. So, to avoid a descent into hell, remember to
t values when mixing calls to GetFontInfo and OutlineMetrics.

Here we also set up some other local variables. The variable `lineCount` is the number of lines we've drawn. The pointer `lineStart` points to the beginning of the current line, which initially is the beginning of the text. The variable `textEnd` is a pointer to just beyond the end of the input text and is used for determining when all is used up. Finally, the variable `textRemaining` keeps track of how much input text remain to be processed.

THE BREAK-DRAW LOOP

Now `NeoTextBox` is ready to break lines and draw the text. The following do-while loop:

```
do {
    lineBytes = 1;
    wrapWid = fixedMax;

    breakCode = StyledLineBreak(lineStart, textRemaining,
        textRemaining, 0, &wrapWid, &lineBytes);

    NTBDraw(breakCode, lineStart, lineBytes, wrapBox,
        boxWidth);

    curY += lineHeight;
    lineStart += lineBytes;
    textRemaining -= lineBytes;
    lineCount++;
} while (lineStart < textEnd);
```

If this looks simple, that's because it is. Anyone who's tried to write a text editor knows that it's a difficult task. Making the algorithm compatible with different systems complicates the matter even more. Fortunately, we have `StyledLineBreak`, which in this case makes our lives a *lot* easier.

The workhorse: `StyledLineBreak`. First we set `lineBytes` to 1. Then we call `StyledLineBreak` that this is the first "script run" on this line. Since this is the first script run, we always reset `lineBytes` at the top of the loop. Also, `wrapWid` is always `fixedMax` (which was previously initialized to the fixed-point width of the rectangle). `WrapWid` tells `StyledLineBreak` the width within which to wrap and returns how much of the line is left (if any) after wrapping. We reset it at the top of the loop each time).

Now we call `StyledLineBreak`. We provide a pointer to the beginning of this line, the number of bytes of text remaining, the wrap width, and a pointer to a variable where `StyledLineBreak` puts the number of bytes remaining. `StyledLineBreak` does the hard work of finding word boundaries, calculating word widths, and handling special cases, all in an internationally com-

lineCount records the
the beginning of the
variable textEnd is a
testing when the text is
now many bytes of

this task is performed by

ng, 0,

align, curY,

write code to wrap text
e with different script
ve the Script Manager,

1, signaling to
Since we have only one
o, we reset wrapWid to
t width of the wrap
which to wrap the text
(that's why we have to

ginning of the text for
h, and the address
in this line.
es, adding up character
npatible way.

After `StyledLineBreak` returns, the text begins at `lineStart`, and the circumstances of the line break are as follows:

```
The head and in fr  
English writer that  
point is therefore c  
letters in a time wh  
the problem to an  
The Shannon Text  
wayultramegasupe  
docious sentence.
```

Figure 2
Line Break Codes

Usually, `StyledLineBreak` returns a word boundary. The break was too long to fit on a single line. `StyledLineBreak` returns the line length for the given width. These line lengths are:

Draw the text with NTBD
line, `NeoTextBox` calls `NTBD` to draw the line of text, the length of the line, the current vertical pen location, and the `NTBD` function:

```
#define kReturnChar '\r'

void NTBDDraw(StyledLineBreak *line,
              long lineBytes, Rect box,
              short boxWidth)
{
    unsigned long blackLen;
    short slen;

    blackLen = VisibleLength(line, boxWidth);
    slen = (short)lineBytes;
    if (blackLen > slen)
        slen = blackLen;
    DrawText(line, box, slen, slen);
}
```

returns, lineBytes tells us the length of the current line and breakCode has a line break code that tells us the type of break, as shown in Figure 2.

horizontal attack on an
 the character of this
 another method for the
 men whom ever told
 unexpected. ←
 is a strange,
 percalafragilisticexpiala

Break Code

```
smBreakWord
smBreakWord
smBreakWord
smBreakWord
smBreakWord
smBreakWord
smBreakWord
smBreakChar
smBreakOverflow
```

Text returns smBreakWord, indicating that it broke the line on a word boundary. The break code smBreakChar says that it encountered a word that was broken on a single line and was forced to break in the middle of a word. The code smBreakOverflow if you run out of text before filling the line. The break codes help determine how to draw the text.

TextDraw. After StyledLineBreak figures the length of the line, TextDraw to draw the line. NeoTextBox passes a pointer to the start of the line in bytes, the wrap rectangle, the alignment, the text position, and the width of the wrap rectangle. Let's take a look at

```
0x0d
TextDraw(int breakCode, unsigned char *lineStart,
         int *wrapBox, short align, short curY,
         int charLen; /* Length of non-white characters */
         int slop; /* Number of pixels of slop for */
         int justify; /* full justification */
         int lineLength=length(lineStart, lineBytes);
```

```

if (align == ntbJustFull) {
    slop = boxWidth - TextWidth(lineStart, 0, blackLen);
    MoveTo(wrapBox->left, curY);
    if (breakCode == smBreakOverflow ||
        *(lineStart + (lineBytes - 1)) == kReturn)
        align = GetSysJust();
    else DrawJust(lineStart, blackLen, slop);
}
switch(align) {
    case teFlushLeft:
    case teFlushDefault:
        MoveTo(wrapBox->left, curY);
        break;
    case teFlushRight:
        MoveTo(wrapBox->right - TextWidth(lineStart,
            blackLen), curY);
        break;
    case teCenter:
        MoveTo(wrapBox->left + (boxWidth - TextWidth(
            lineStart, blackLen)) / 2, curY);
        break;
}
if (align != ntbJustFull)
    DrawText(lineStart, 0, lineBytes);
}

```

NTBDraw's job is to move the pen and draw the text as indicated by the `align` parameter, `align`, and the line break code, `breakCode`. NTBDraw determines the visible length of the line with a call to the Script Manager routine `VisibleLength`. This excludes white-space characters at the end of the line. What about white space characters? Well, that depends on the script. `VisibleLength` returns the number of visible characters and which are not for the current script. `VisibleLength` also returns the appropriate length in bytes, which is stored in the local variable `blackLen`.

When `align` is `ntbJustFull`, we need to determine whether the carriage return character (`$0D`) at the end, because a line with a carriage return (for example, the last line in a paragraph) should always be drawn with left alignment, rather than fully justified.

Looking back at the break codes for different types of lines shows that the line that ends with the carriage return (denoted graphically in the illustration) returns a line break code of `smBreakWord`, whereas a line that ends with a carriage return returns `smBreakOverflow`. As you can see, `StyledLineBreak` returns `smBreakOverflow` when a line is the last line of a paragraph. Therefore, every line that returns `smBreakWord` must be checked for a carriage return.

```
blackLen);
```

```
inChar)
```

```
, 0,
```

```
h(lineStart, 0,
```

ated by the alignment
Draw first calculates the
routine VisibleLength.
What exactly are white-
length knows which
, and returns an
ble blackLen.

e current line has a
n a carriage return (for
with the default system

own in Figure 2, notice
hically in the
e you might expect it to
pects the caller to know
ne whose break code is

NTBDraw looks at the last character. Since the carriage return character is guaranteed never to occur, there's no need for NTBDraw to have to test whether the character is a carriage return. NTBDraw can proceed directly to the last character.

We now know whether the text is justified. We can calculate the amount of width to use for the text. The Manager routine DrawJustifiedText does this for this script. (In Arabic, for example, text is justified differently than for Roman scripts.) NTBDraw overrides the align parameter to justify the text.

For the left, right, and center alignment, NTBDraw calls DrawJustifiedText appropriately, and a DrawJustifiedText routine (blackLen) helps correctly calculate the amount of text and full justification.

Update the variables. After drawing the text, NTBDraw updates the variables and loops around to draw the next line.

```
    curY += lineHeight;
    lineStart += lineBy;
    textRemaining -= lineLen;
    lineCount++;
} while (lineStart < textLen);
```

First, we add lineHeight to curY, the current vertical pointer to the beginning of the current line. TextRemaining is the amount of text in the current line, and lineCount is the number of lines of the text, the whole break.

RETURN SOME VALUES

Now that NeoTextBox has calculated the amount of text, some useful values to the caller.

```
    if (endY)
        *endY = curY - lineHeight;
    if (lhUsed)
        *lhUsed = lineHeight;
```

NeoTextBox returns these values. The caller can get TextBox-like behavior by setting endY and lhUsed. If the caller doesn't want a return value, set endY and lhUsed to NULL.

st byte in the line it's drawing to see if it's a carriage return. character (\$0D) falls into the control-code range, it's as the low byte of a two-byte character. This frees us from the last character in the line is two-byte and allows us to t byte.

the current line has a carriage return or not. If not, we white-space slop remaining in the line, then call the Script t to draw the text fully justified—whatever that means for instance, full justification is performed completely n text.) If the current line *does* end in a carriage return, we ter with the default system alignment and fall through.

ter alignments, the switch statement moves the pen /Text call is made to draw the text. The visible length (in calculate the pen position for right and center alignment

After NTBDraw returns, we need to update a bunch of local again.

```
tes;  
neBytes;  
extEnd);
```

o curY, setting us up for the next line. LineStart, the of a line, gets updated to the character after the end of the ng gets reduced by the number of bytes consumed by the nt gets incremented. If lineStart still hasn't run off the end k-draw process is repeated.

s done such a fine job wrapping the text, it's time to return caller.

```
ineHeight;  
ight;
```

e values only if the caller wants them. This makes it easy to from NeoTextBox without having to do any work: if you just pass nil instead of providing the address of a variable.

CLEAN UP AND WE'RE DONE

The only thing left to do is a little cleanup, and we're outta here.

```
SetClip(oldClip);
DisposeRgn(oldClip);

return(lineCount);
}
```

We restore the clipping region, dispose of our saved region, and

CALLS TO NEOTEXTBOX

One of the best features of NeoTextBox is that you can easily substitute it for the TextBox you're currently making to TextBox. If that's all you want to do, here's an occurrence that looks like this:

```
TextBox(textPtr, textLen, &wrapBox, justify);
```

with this:

```
{
    EraseRect(&wrapBox);
    NeoTextBox(textPtr, textLen, &wrapBox, justify, 0,
}
}
```

To use NeoTextBox in place of TextBox, you pass 0 for lhCode and nil for endY and lhUsed, and ignore the return value. If you use NeoTextBox in your program and just do the substitution above, every NeoTextBox call is the average 33% faster than the old TextBox call. If you use NeoTextBox in your program, it means a real performance increase.

You can use NeoTextBox in more ways than just as a direct substitution for TextBox performance. It does, after all, have whizzy new features that you can't get from TextBox. Take a look at a more sophisticated call to NeoTextBox that uses some of its features:

```
short UseNTB(void)
{
    Rect          wrapBox;
    RGBColor      ltBlue;
    Handle        textHdl;
    long          textLen;
    short         numLines = 0;
    short         endY, lineHt;
```

re.

and return lineCount.

substitute it for calls
o, replace every

```
nil, nil);
```

e (default line height)
ou add NeoTextBox to
extBox call will be on
extBox a lot, that can

stitution to improve
TextBox never had. Let's
es some of its unique

```

/* Set up our RGBColor
SetRect(&wrapBox, 1
ltBlue.red = 39321;
ltBlue.green = 5242
ltBlue.blue = 65535

/* Paint the background
PenNormal();
RGBForeColor(&ltBlue
PaintRect(&wrapBox)
ForeColor(blackColor
TextFont(helvetica)
TextFace(0); TextM

/* Retrieve some text
textHdl = GetResource
if (textHdl) {
    textLen = GetHandleSize(textHdl);
    /* Be sure to lock the handle
    HLock(textHdl);

    /* Wrap text and return number of lines
    numLines = NeoTextWrap(
        18, &endY, &
    HUnlock(textHdl);

    /* Beep if text overflows
    if (endY > wrapBox.bottom)
        SysBeep(1);

    /* Prove we know how to move the cursor
    MoveTo(wrapBox.left,
    Line(20, 0);
}
return(numLines);
}

```

This sample function draws text from a TEXT resource into a window. The text is fully justified. If the text overflows the box, a beep is heard. If subsequent text might be

Here's an example using

```

color and wrapBox. */
(0, 10, 110, 110);

18;
;

ground, then set up the port text parameters. */

e);
;
or);
; TextSize(12);
ode(srcOr);

text for us to draw. */
rce('TEXT', 128);

andleSize(textHdl);
ock the handle. NeoTextBox can move memory! */

d set numLines, endY, and lineHt. */
extBox(*textHdl, textLen, &wrapBox, ntbJustFull,
lineHt);
);

verflows wrapBox. */
Box.bottom)

w where the text ended by drawing a line. */
left, endY + lineHt);

```

draws a 100-by-100-pixel box in light blue, then wraps text
 into the rectangle, ORing the text over the blue background.
 12-point Helvetica®, with 18-point line spacing. If the text
 sounds. A small line is drawn at the baseline where
 drawn.

NeoTextBox with variable line height and TrueType fonts:

```

void UseVariableLineHeight(Rect *wrapBox, short align
{
    Boolean    oldPreferred, oldPreserve;
    Handle     textHdl;
    long       textLen;

    if (gHasTrueType) {
        oldPreferred = GetOutlinePreferred();
        oldPreserve = GetPreserveGlyph();
        SetOutlinePreferred(TRUE);
        SetPreserveGlyph(TRUE);
    }

    textHdl = GetResource('TEXT', 128);
    textLen = GetHandleSize(textHdl);
    HLock(textHdl);
    NeoTextBox(*textHdl, textLen, wrapBox, align, -1,
    HUnlock(textHdl);

    if (gHasTrueType) {
        SetOutlinePreferred(oldPreferred);
        SetPreserveGlyph(oldPreserve);
    }
}

```

Notice that we save the current settings of the Font Manager's PreserveGlyph flags. This allows us to be transparent to the OutlinePreferred to TRUE, we are ensured of using TrueType bitmapped fonts are available. By setting PreserveGlyph to TRUE, accurate glyph shapes and measurements (see "SetPreserveGlyphs" on page 34). Calling NeoTextBox with -1 as its lhCode use variable line height, which results in the difference shown

lhCode = 0

1

The Head And in Frontal Att
Therefore Another Method F
Problem to An Unexpected.

The Head
Therefore
Problem t

Default line height

Va

Figure 3

Using Variable Line Height

LIMITATIONS AND

NeoTextBox is a nice alternative, but it could benefit from improvements. Here are some limitations and adding us

32K TEXT SIZE LIMIT

All you *War and Peace* fans know that the 32K text limitation that TextEdit has is not only 32K of text in one call, but also arises from the OutlineManager. The Manager can only handle 32K of text at a time. NeoTextBox implements variable line wrapping with a NeoTextBox call (knock y

DON'T FORGET TO ERASE

Perhaps this isn't really a limitation, but a NeoTextBox call. You need to call NeoTextBox as shown earlier in the section.

SCREEN-ONLY OPTIMIZATION

If you know you'll be using NeoTextBox, you won't be using it to draw. If you don't care about the performance, the wrap rectangle intersects the screen, return.

If you don't need the return value, you can break-draw loop terminate.

SPECIAL ONE-LINE CASE

In Macintosh computers, NeoTextBox is a worthwhile addition to NeoTextBox. If you use DrawBoxWidth, simply use DrawBoxWidth. TextBox wrapping code. TextBox items, which are often on

DON'T DRAW OFF THE

It might make NeoTextBox wrapBox->bottom + lineCount (total number of lines), but it's clipped.

D POSSIBLE ENHANCEMENTS

Alternative to TextBox, but it has its limitations and areas that need improvement. Following are some suggestions for overcoming the limitations of the useful features.

As it stands out there need to do a little work. NeoTextBox shares the limitations of TextBox, but not for the same reason. TextBox can wrap text all because it uses TextEdit. In NeoTextBox, the limitation is the TextMetrics call, which is used in deriving variable line height and text. Heavy-duty Tolstoy types could remove the code that calculates line height and subsequently word wrap most novels in a single call (you yourself out).

BASE

One limitation, but you can't simply replace a TextBox call with a NeoTextBox call. You need to call EraseRect explicitly if you want TextBox behavior, and you need to call NeoTextBox.

OPTIMIZATIONS

Using NeoTextBox only for screen applications (that is, you don't send it to a printer port), you can make a few optimizations. If you need return values, you can use RectInRgn to check whether the text is within the current port's visRgn; if it doesn't, you can simply return a value giving the number of total lines, you can make the text wrap when curY exceeds wrapBox->bottom + lineHeight.

BASE

With 256K ROMs, TextBox has a feature that might be a NeoTextBox. If the TextWidth of the input text is less than wrapText, draw the text and don't bother with any of the other features. It has this feature because it's used for dialog box statText and the line.

END OF WRAPBOX

Draw faster if NTBDraw isn't called when curY is greater than wrapBox->bottom + lineHeight. You'd still have to wrap all the text (to determine the text you wouldn't be drawing text that you know will be

MAKE SAVING/RESTORING THE CLIPPING REGION OPTI

It might be useful to be able to set up some complex clipping. NeoTextBox wrap as usual but clip its text to whatever the clip invocation. You could add a Boolean swapClip parameter to c

STYLED NEOTEXTBOX

With considerable effort, NeoTextBox could be extended to h multiscrypt text. Since StyledLineBreak, the workhorse of Ne be used with styled text, such an enhancement is possible.

CONCLUSION

Once you start using NeoTextBox, you'll find it ending up in tests on a Macintosh IIfx running System 7, NeoTextBox was faster than TextBox, 33% faster on the average. Performance font, screen depth, and the ratio of wrapping to drawing. For on an 8-bit screen, NeoTextBox is 40% faster than TextBox. T reason to use it. Plus, it has features you can't get out of TextB

Perhaps the moral of this article is if you don't like some featur OS go ahead and write your own. But you'll be doing yourself a lot more compatible in the future—if you can find lower-level OS facilities to aid you in your task, rather than recoding the

So go ahead and whip NeoTextBox into your application. Enjoy performance and new features. And if there's something you c there and change it. Make NeoTextBox the TextBox *you've* alv

THANKS TO OUR TECHNICAL REVIEWERS

Sue Bartalo, John Harvey, Joe Ternasky •

ONAL

region and have
pping region is set to at
ontrol this.


andle styled and
oTextBox, is designed to

all your applications. In
between 25% and 50%
varies depending on
left-aligned Geneva text
That alone is a good
Box at all.

ure of the Toolbox or
f a favor—and you'll be
el system, Toolbox, or
entire feature yourself.

oy the improved
don't like, go right in
ways wanted!

MAKING YOUR MACINTOSH SOUND LIKE AN ECHO BOX



RICH COLLYER

Happy notes for sound on the Developer CD and record sounds at record into one buffer between the buffers. I out elements of this co

We all know that the Mac clever programming you 2BufRecordToBufCmd in application (sans interface time that you're playing purpose is to educate you definitive road to becomi

In addition to the main routines and a completion code out of this article.

CONSTANT COM

Before I get into the same in the application.

GETTING A HANDLE O

The kMilliSecondsOfSo sound the application should number of milliseconds, used to calculate the size sound effect you're after, 400,000 or so. If you set time for the completion high end of the range, or

RICH COLLYER is just your year Developer Technical Support often heard screaming at his soothing accompaniment of B KOME radio, he's honed his fine point dodging (and casting arrows at Apple, and he actual degree from Cal Poly with a s computational fluid dynamics.

d buffs: As you'll see from the sample code provided on the Developer CD Series disc, you can make your Macintosh play two buffers at the same time, simply by using double buffering to record into one buffer while playing a second buffer, and then flipping the buffers. If you want to take things a few steps further, pull the sample code and tailor them to suit your own acoustic needs.

Macintosh is a sound machine, so to speak, but with a little programming you can turn it into an echo box as well. The sample code included on the *Developer CD Series* disc is just a small routine (see Listing 1) that demonstrates one way to record sounds at the same time and play them. There are other ways to achieve the same goal, but my intention here is to point you about the Sound Manager, not to lead you down the path of creating your own recording studio.

The routine, `2BufRecordToBufCmd` includes various setup and cleanup routines. For easy reading, I've left out any unnecessary

COMMENTS

In the sample code itself, here are a few of the constants you'll run into

CONSTANTS

The `kSndBufSize` constant is used to declare how many milliseconds of sound should be recorded before it starts to play back. The smaller the value, the more quickly the sound is played back. This constant is the size of the 'snd' buffer handles (just the data). Depending on the value of `kMilliSecondsOfSound` can range from 50 milliseconds to less than 50, you risk problems: there may not be enough time for the routine to finish executing before it's called again. On the other hand, if the application's available memory limits the size. The

run-of-the-mill three-
port veteran: He's
computer to the
lazy and Bob on
archery skills to a
g) the slings and
ally admits to a
specialty in
We let you in on

his outdoor adventures last time he wrote for us
and he claims most of his indoor adventures
aren't appropriate *develop* material, but we have
it on good authority that he lives with carnivorous
animals, if that's any clue. He's also a confirmed
laserdisc and CD addict; he keeps promising to
start a recovery program for those of us with the
same affliction just as soon as he finishes writing
that next sample . . . •

smaller the value, of course, the faster the buffers fill up and produce an echo effect you'll get. A millisecond value of 1000 provides a delay between record and echo, which I've found is good for general experimentation to find the effect you like. (Beware of feedback, both from the system and from anyone who's in close enough proximity to "enjoy" the sound secondhand.)

YOUR HEAD SIZE, AND OTHER #DEFINES

The next three constants (`kBaseHeaderSize`, `kSynthSize`, and `kCmdSize`) parse the sound header buffers in the routine `FindHeaderSize`. `kBaseHeaderSize` is the number of bytes at the top of all 'snd' headers that aren't recorded by the application itself. While the number of bytes isn't really of interest, it's used to parse the header in order to find the part of the sound header that's in the `bufferCmd`. How much you parse off the top is determined by `kBaseHeaderSize` and the type of file; for the purposes of this code, however, the constants concerned with are the 'snd' resources. The second constant, `kSynthSize`, is of one 'synth'. In the calculations of the header, I find out how many 'synth's and multiply that number by `kSynthSize`. The last constant, `kCmdSize`, is one command, which is used in the same way as `kSynthSize`. (This code is derived from *Inside Macintosh* Volume VI, page 22-20.)

ERROR CHECKING WITH EXITWITHMESSAGE

`2BufRecordToBufCmd` includes error checking, but only as a result of the commercialization of the product. If the present code detects an error, it calls `ExitWithMessage` routine, which displays a dialog box that tells you where the error occurred and what the error was. Closing this dialog box ends the application, at which point you have to start over again. Note that calling `ExitWithMessage` at interrupt time could be fatal, since it uses `Move` to move memory. For errors that could occur at interrupt time, I use `ExitToShell` instead.

USING THE SOUND INPUT DRIVER

Use of the sound input driver is fairly well documented in *Inside Macintosh* Volume VI, Chapter 22 (pages 22-58 through 22-68 and 22-92 through 22-98). This is a little overview of what `2BufRecordToBufCmd` does at this point. The reason why. When you use sound input calls at the low level (not using `SndRecordToFile`), you need to open the sound input driver. The routine `SPBOpenDevice` just opens the driver, which the user selects via the sound device menu.

```
gError = SPBOpenDevice (kDefaultDriver, siWritePermissions);
```

To open the driver, you call `SPBOpenDevice` and pass in a couple of parameters. The first parameter is a driver name. It doesn't really matter what the name of the driver is; it simply needs to be the user-selected driver.

ay back, and the faster
a one-second delay
use. You'll want to
th from your machine
he experimentation

(CmdSize) are used to
kBaseHeaderSize is
needed in the
erest here, you need to
that you'll pass to the
the format of the
ver, all you need to be
kSynthSize, is the size
many 'snth's there are,
CmdSize, is the size of
These equations are

placeholder for future
an error, it calls the
ls you more or less
dialog box quits the
that calling
routines that might
DebugStr is used

de Macintosh Volume
h 22-99), but here's a
nt in the routine, and
g SndRecord or
This section of the code

```
tion, &gSoundRefNum);
```

uple of simple
ally matter what the
river, so the code passes

GESTALT YOUR

You do need to check two things. First of all, your machine must be trying to record sounds. Second, the sound input capabilities of the hardware. Gestalt Manager will give you the following information:

Second, your hardware must be capable of recording for sound input and one of the returned feature variables will be non-zero.

The following code shows how to check the variable. (I didn't use the variable name shown in the code.)

```
err = Gestalt (gestaltFeatureID, &feature);
if (!err) {
    if (feature & (1 << kGestaltSoundInput))
        //This Macintosh has sound input
    if (feature & (1 << kGestaltSoundOutput))
        //This Macintosh has sound output
    if (feature & (1 << kGestaltSoundInputOutput))
        //This Macintosh has both
    if (feature & (1 << kGestaltSoundInputOnly))
        //This Macintosh has only input
    if (feature & (1 << kGestaltSoundOutputOnly))
        //This Macintosh has only output
}
```

in nil (which is what kDeviceSoundRefNum tells the driver you'd like to use). You can enable the application to use the device by calling SPBSetDeviceRefNum(gSoundRefNum). This part of the code asks specific questions about the hardware to ensure that nothing went wrong. If there is an error, ExitWithMessage, and then return.

```
gError = SPBSetDeviceRefNum(gSoundRefNum);
```

Continuous recording is supported on the Quadra 700 and 900 that have the sound input levels to their normal levels.

MACHINE

Two rather critical sound attributes for 2BufRecordToBufCmd. You must have a sound input driver. There's very little point in running the sample if the machine is being run on a machine that doesn't have a sound input driver. Checking bit 5 of the returned feature variable with the Gestalt function gives you this handy bit of information.

The sample needs to support stereo sound, since you need one channel for sound output. Check for this attribute by checking bit 0 of the Gestalt feature variable.

Here's how you can test all of the bits returned in the feature variable (this is code in my sample.)

```
GestaltSoundAttr, &feature);

    (<< gestaltStereoCapability))
    Macintosh Supports Stereo (test bit 0)
    (<< gestaltStereoMixing))
    Macintosh Supports Stereo Mixing (test bit 1)
    (<< gestaltSoundIOMgrPresent))
    Macintosh Has the New Sound Manager (test bit 3)
    (<< gestaltBuiltInSoundInput))
    Macintosh Has Built-in Sound Input (test bit 4)
    (<< gestaltHasSoundInputDevice))
    Macintosh Supports Sound Input (test bit 5)
```

defaultDriver translates into). The constant siWritePermission gives you read/write permission to the sound input driver. This will actually use the recording calls. The last parameter is the driver parameter is needed later in the sample so that you can ask the driver that's open. The error checking is just to make sure everything is working; if something did go wrong, the code goes to the error handling routine and the sample quits.

```
Info (gSoundRefNum, siContinuous, (Ptr) &contOnOff);
```

is activated here to avoid a "feature" of the new Macintosh which gives you a slowly increasing ramp of the sound level each time you call SPBRecord. The result in

2BufRecordToBufCmd is a pause and gradual increase in the s buffers as the buffers are being played. Continuous recording only on the first buffer, where it's almost unnoticeable.

BUILDING 'SND ' BUFFERS

Now that the sound input driver is open, the code can get the build the 'snd ' buffers. As its name implies, 2BufRecordToBu The reason is sound (no pun intended): The code basically use method to record and play the buffers. The code doesn't tell t play the sound until the recording completion routine has bee have to worry about playing a buffer before it has been filled w code also does not restart the recording until the previous buff

INFORMATION, PLEASE

To build the sound headers, you need to get some information driver about how the sound data will be recorded and stored. the GetSoundDeviceInfo routine, which looks for information (the number of samples per second at which the sound is recor (the sample size of the sound being recorded—8 bits per samp CompressionType (see “Putting on the Squeeze”), the Number of sound input channels, normally 1), and the DeviceBufferInf internal buffers).

This code (minus the error checking) extracts these values from driver.

```
gError = SPBGetDeviceInfo (gSoundRefNum, siSampleRate,
    (Ptr) &gSampleRate);

gError = SPBGetDeviceInfo (gSoundRefNum, siSampleSize,
    (Ptr) &gSampleSize);

gError = SPBGetDeviceInfo (gSoundRefNum, siCompression,
    (Ptr) &gCompression);

gError = SPBGetDeviceInfo (gSoundRefNum, siNumberChan,
    (Ptr) &gNumberOfChannels);

gError = SPBGetDeviceInfo (gSoundRefNum, siDeviceBuff,
    (Ptr) &gInternalBuffer);

value = kMillisecondsOfSound;
gError = SPBMillisecondsToBytes (gSoundRefNum, &value);
gSampleAreaSize = (value / gInternalBuffer) * gIntern
```

ound volume between
gives you this ramp

information it needs to
fCmd uses two buffers.
es a double-buffer
ne machine to start to
n called, so you don't
with recorded data. The
fer has started to play.

from the sound input
That's the function of
n about the SampleRate
rded), the SampleSize
le is normal), the
erChannels (the number
fo (the size of the

m the sound input

nType,

nels,

erInfo,

);
alBuffer;

PUTTING ON THE SQUEEZE

If you want to use compression for `2BufRecordToBufCmd`, keep in mind that the Sound Manager basically supports three types of sound compression: none at all, which is what I'm using, and MAC3 and MAC6, which are Mac compression types for 3:1 and 6:1 compression, respectively.

If you set the compression, the sound data is compressed after the interrupt routine is called (if you have one) and

Opening the sound input...
`siSampleRate`, `siSampleS`,
`siDeviceInfo` are co...
the `SPBGetDeviceInfo` c...
pointer to a global variab...
the requested information

The last bit of work the c...
headers is to convert the...
buffer. To do this, the rou...
down the resulting value...
to bypass a bug connecte...
sound input device, which...
buffer is not a multiple of...
size not only avoids this p...
record data into your buf

Now the code has the inf...
space, I've made a short r...
has to do is call this routi...
appropriate data.

IT'S A SETUP

The first line of code in t...
Memory Manager to allo...
data buffer and an estima...
checking (see the code its...
header. Setting up the so...
call, `SetupSndHeader`, to...
`SetupSndHeader` only on

before the Sound Manager internal buffers are moved to the application's sound buffers.

You have a couple of options for playing back a compressed sound. Either the `bufferCmd` or `SndPlay` will decompress the sounds on the fly. If you need to decompress a sound yourself, you'll want to call the Sound Manager routine `Exp1to3` or `Exp1to6` (depending on the compression you were using).

driver gives you the `gSoundRefNum`. The values `siSize`, `siCompressionType`, `siNumberChannels`, and constants defined in the `SoundInput.h` file; these constants tell all what information you want. The last parameter is a pointer to a `SoundInput` structure. The `SPBGetDeviceInfo` call uses this parameter to return information.

code needs to do before it's ready to start building the 'snd' buffers. The constant `kMillisecondsOfSound` is the sample size of the sound. The routine needs to call `SPBMillisecondsToBytes` and then round up to a multiple of the size of the internal sound buffer. This is done to avoid the continuous recording feature of Apple's built-in sound driver. It will collect garbage rather than audio data if the recording buffer is full of the device's internal buffer. Creating a buffer of the right size solves the problem, but also enables the input device to more efficiently transfer data.

information it needs to build the sound buffers. To save code, the routine that builds the buffers and their headers. All the code needed for each of the buffers it needs and pass in the

The `SetupSounds` routine is fairly obvious. It simply calls the `SPBGetDeviceInfo` routine to get the requested handles, based on the known size of the buffers (and the requested maximum size for the header, and does some error checking (if the handle is good, the routine builds the 'snd' buffers. The sound buffer requires building the header by making a simple call to the Sound Manager. There's a small problem with calling `SPBGetDeviceInfo`, however: When you call it, you don't know how big the

sound header is, so you just give the call the buffer, along with size. When the call returns with the header built, one of the values that's the number of bytes in the sample—will be wrong. (It's correct, but the data in the header will not be.) To correct this, when your recording is complete and then put the correct number of bytes in the header, at which time you'll know how much data there is to prevent misinformation in the header won't affect your recording, only

Once the header's built, the code resets the size of the handle, (to avoid fragmentation of the heap), and locks it down. It's important to lock the handles in this way; otherwise the Sound Manager will move them working with out from under itself.

```
*bufferHandle = NewHandle (gSampleAreaSize + kEstimatedHeaderSize);
gError = SetupSndHeader (*bufferHandle, gNumberOfChannels,
    gSampleSize, gCompression, kMiddleC, 0, headerSize);
SetHandleSize (*bufferHandle, (Size) *headerSize + gSampleAreaSize);
MoveHHI (*bufferHandle);
HLock (*bufferHandle);
```

TELLING IT WHERE TO GO

The next part of the program allocates and initializes a sound record structure `gRecordStruct`. This structure tells the sound input call how to do what it wants it to do.

The first instruction is obvious: it simply creates a new pointer to where the structure can be stored.

```
gRecordStruct = (SPBPtr) NewPtr (sizeof (SPB));
```

The recording call will need to know where it can find the open file, so next it needs the reference number to the driver (`gSoundRefNum`). The subsequent three lines of code inform the recording call how much data to record into. Here, you could either give the call a count value or a number of milliseconds are available for recording, or give it the size of the buffer. In this code, it's easiest to just make the `bufferLength` the same as the `milliseconds` value. The code then tells the recording call how much data as it's recorded.

```
gRecordStruct->inRefNum = gSoundRefNum;
gRecordStruct->count = gSampleAreaSize;
gRecordStruct->milliseconds = 0;
```

a 0 value for the buffer values in the header—the header size will be, you simply wait until of bytes directly into the playback. The playback.

moves the handle high important to lock down the sound buffers it's

```
edHeaderSize);
```

```
nels, gSampleRate,  
);
```

```
ampleAreaSize);
```

input parameter block, do what the code

r into which the

en sound input driver, (refNum). The much buffer space it has ue, tell it how many he sound buffer. For s the count and ignore where to put the sound

```
gRecordStruct->bufferL
gRecordStruct->bufferP
gRecordStruct->complet
gRecordStruct->interru
gRecordStruct->userLon
gRecordStruct->error =
gRecordStruct->unused1
```

The recording call also means that the call is done asynchronously (you can call this routine later on.) You can check the driver periodically to see if the routine SPBGetRecord has finished recording, you can check if the buffer has been filled. For this code, the recording is done because you are likely you are to prevent p

The userLong field is a global variable you'll need in order to handle the completion routine. As you don't need an interrupt routine using the unused1 field.

You'd need to use an interrupt routine before compression, or before the "Interruptions").

TIME TO CHANNEL

Just before the code jumps to the channel generally is not a big deal. You can use no interpolation.

ROUTINE INTERRUPTIONS

The interrupt routine gives you a chance to manipulate the sound data before any sound compression is done. For some of the operations that you may want to carry out inside the interrupt routine, you'll need access to the A5 world of the application, which is why I stored 2BufRecordToBufCmd's A5 value in the userLong field of gRecordStruct.

```
length = gSampleAreaSize;
ptr = (Ptr) ((*bufferHandle) + gHeaderLength);
completionRoutine = (ProcPtr) MyRecComp;
interruptRoutine = nil;
g = SetCurrentA5();
0;
= 0;
```

needs to know what to do when it's finished recording. Since obviously, it needs a completion routine. (I'll talk more about *could* leave out the completion routine and just poll the if it's finished recording. To do that, you'd repeatedly call `rdStatus`, and when the status routine informed you that you'd restart the recording and play the buffer that had just however, it's better to know as soon as possible when the the more quickly you can restart the recording, the more pauses between recordings.

ood place to store `2BufRecordToBufCmd`'s A5 value, which ve access to the application's global variables from the ou can see, the rest of the fields are set to 0. The code routine. There's also no point in passing an error back or

errupt routine if you wanted to change the recorded sound before the completion routine was called (see "Routine

s into the main loop, it needs to open a sound channel. This , but for `2BufRecordToBufCmd`, I initialized the channel to

For more information about sound interrupt routines, take a look at *Inside Macintosh* Volume VI, page 22–63.

Warning: Don't try to accomplish too much in an interrupt routine. In general, you'll want interrupts to be minimal, and possibly written in assembly language, to avoid unnecessary compiler-generated code.


```
gError = SndNewChannel (&gChannel, sampledSynth, init
```

Interpolation causes clicks between the sound buffers when the buffers wrap back, which can be a rather annoying addition to your recording (if you're going for that samba beat).

JUST FOR THE RECORD

To start recording, all the code needs to do now is call the low-level routine, pass in `gRecordStruct`, and tell it that it wants the recording to proceed asynchronously.

```
gError = SPBRecord (gRecordStruct, true);
```

LOOP THE LOOP

The main loop of this code is a simple while loop that waits until the user presses a button or an error occurs in the recording, at which time the loop terminates.

```
/* main loop of the app */  
while (!Button() || (gRecordStruct->error < noErr));
```

ROUTINE COMPLETION

You don't want a completion routine to do much, generally, since it takes a long time and keeps your system locked up while it's running. The completion routine, one of which has four parts to itself.

The first part of the completion routine sets its A5 value to be the same as the A5 value of the application. This gives you access to the application's A5 value from the completion routine.

```
storeA5 = SetA5 (inParamPtr->userLong);
```

If the completion routine weren't broken into two parts here, the current optimization scheme would cause a problem at this point: the completion routine would be pointed to in an address register as an offset of A5 from the current value of A5. To set A5 to your application's A5 value, and you'd get garbage. Therefore, it's necessary to restore your A5 value (part 1 of the completion routine) and then call the secondary completion routine to actually do the work.

Before the routine does any work, it needs to make sure that there are no problems with the recording. If there were errors, the code would return from the completion routine without doing anything.

```
if (gRecordStruct->error < 0)  
    return;
```

NoInterp, nil);

they're played back to
ing (unless, of course,

y-level recording
ording to occur

until the mouse button is
application quits.

since it's run at interrupt
there are three parts to this

the same as the A5
on's global variables

the MPW C compiler
ess to global arrays
before you had a chance
e information.
e completion routine)
all the work.

there have not been any
rops out of the

Next the routine prepares correcting the header's length in `gRecordStruct`, which now

```
header = (SoundHeaderP  
         gHeaderSize);  
header->length = gReco
```

Once the header's been fixed, the routine to play the sound.

```
PlayBuffer (gBufferHan
```

The last part of the real code is recording. To do this, the routine rebuilds `gRecordStruct` to include setting the correct `bufferLength` is correct. Call `SPBRecord` to restart the

```
#define NextBuffer(x)  
  
gWhichRecordBuffer = N  
gRecordStruct->bufferP  
    gDataStart);  
gRecordStruct->millise  
gRecordStruct->count =  
gRecordStruct->bufferL  
  
err = SPBRecord (gReco
```

The last piece of the command routine started.

```
storeA5 = SetA5 (store
```

PLAY TIME

The code in the `PlayBuffer` routine set up the command parameters to know what channel to play the sound structure by telling `SndDoCommand` along with `SndDoCommand` then play the call, false, basically tells the

s the header of the buffer, which has just been filled, by length field. This field needs to be set to the count field of w contains the actual number of bytes recorded.

```
tr) (*(gBufferHandle[gWhichRecordBuffer]) +
```

```
rdStruct->count;
```

ked, the code just sends the buffer handle off to the play (See "Play Time" for a full explanation of the play routine.)

```
dle[gWhichRecordBuffer]);
```

ompletion routine prepares gRecordStruct to start the next code needs to select the correct buffer to record to and reflect any changes. The macro NextBuffer performs an hichRecordBuffer to make it either 1 or 0. The changes t buffer to record to and checking to see that the Once the structure is reset, the code makes the next call to recording.

```
(x ^= 1)
```

```
extBuffer (gWhichRecordBuffer);
```

```
tr = (*(gBufferHandle[gWhichRecordBuffer]) +
```

```
conds = 0;
```

```
gSampleAreaSize;
```

```
length = gSampleAreaSize;
```

```
rdStruct, true);
```

pletion routine resets A5 to what its value was when the

```
A5);
```

er routine is very simple Sound Manager code. All it does is meters and call SndDoCommand. The routine needs to y into and what buffer to play, so the code sets up the local it which buffer to play, and sends that local structure to with the necessary channel information (gChannel). plays the sound. The last parameter in the SndDoCommand ne Sound Manager to always insert the command in the

channel's queue: if the queue is full, SndDoCommand will wait and insert the command before returning.

```
localSndCmd.cmd = bufferCmd;
localSndCmd.param1 = 0;
localSndCmd.param2 = (long) ((*bufferHandle) + gHeader);
gError = SndDoCommand (gChannel, &localSndCmd, false);
```

If you wanted to send the sounds to a different machine to be played, you can simply replace the code in the the PlayBuffer routine with IPC Toolbox calls telling a second machine to play the buffers.

CLEANING UP AFTER THE SHOW

Once the code finds the mouse button down or discovers that the recording and exits the main loop, there's only one last thing to do. The first part of cleaning up is to close the sound input driver. Before closing the driver, you need to make sure it's not in use; the routine SPBSStopRecording stops recording.

```
gError = SPBStopRecording (gSoundRefNum);
SPBCloseDevice (gSoundRefNum);
```

Next you need to dispose of the handles and pointers you've been sending them on their way, however, you have to make sure they were allocated, so the code checks to see whether or not the handle

```
for (index = 0; index < kNumberOfBuffers; ++index)
    DisposeHandle (gBufferHandle[index]);
DisposePtr ((Ptr) gRecordStruct);
```

Last but not least, the code disposes of the sound channel for you. The quitNow flag clears the sound queue before the channel is closed.

```
gError = SndDisposeChannel (gChannel, true);
```

COMPOSE YOURSELF

So now you know a little bit more about doing basic sound input. I've fielded many questions about clicks, pauses between buffers, and how to resolve and built into 2BufRecordToBufCmd. The specific code here may not apply to what you're interested in doing right now, but the sound input driver or are interested in continuous recording may be useful to you in some other application. You've heard the saying "you like and leave the rest"? Sound advice (so to speak).

THANKS TO OUR TECHNICAL REVIEWERS

Neil Day, Kip Olson, and Jim Reekes, who burned the midnight oil ripping this code to shreds and putting it back together again. •

it until there's space to

```
erSize);
```


played, you could
C or Communications

an error occurred in
ng to do: clean up. The
ore you can close the
topRecording stops the

een using. Before
at they have been
s and pointer are nil.

you. Setting the
sed.

put at a low level. I've
nd so on, which I've
echniques I've outlined
ow, but if you're using
ng, parts of this sample
the saying "take what



BE OUR GUEST

BACKGROUND-ONLY APPLICATIONS IN SYSTEM 7

C. K. HAUN

One of the least heralded new features of System 7, but nonetheless a very important one, is full support for faceless background applications (FBAs). An FBA is a full-fledged application that's invisible to the user. It has its own event loop, and it receives time and some events like any other application, but it doesn't have a menu bar, windows, dialogs, or other graphic components. An FBA is a normal file of type 'APPL'.

FBAs are, by a stretch of the imagination, similar to UNIX® daemons. The purpose of an FBA is to provide services to other applications or to monitor the system. For instance, an application that periodically checks your hard drive for files that haven't been backed up lately is a perfect candidate for FBA status. Thus, an FBA can be a silent partner to your application, INIT, cdev, desk accessory, or driver.

An FBA is the best way to provide certain services. For example, an FBA paired with a desk accessory can enable the DA to send Apple events, something a DA cannot usually do. (See the AECDEV/AEDAEMON sample in the snippets provided with the DTS Sample Code on the *Developer CD Series* disc.) An FBA can replace an INIT that patches traps to get time and provides services, or it can replace a driver that depended on periodic run messages to operate. Converting to an FBA not only frees you from having to patch to get the time you need, but also gives you a fully supported and documented interface and design.

58

C. K. HAUN has been programming Apple computers since 1979, writing commercial education, utility, and game applications for the Apple II, IIGs, and Macintosh, with some occasional dark forays into the Big Blue world. (It paid the rent.) He currently works in Developer Technical Support and focuses mainly on Apple events and the Edition Manager. Besides working to provide the best possible support to developers, he's been trying to organize the Silicon Valley chapter of Heck's Moofers, a motorcycle club devoted to the precept that computer nerds on