Developers' Newsletter

No 26 - May 1992

Acorn

General News

Queen's Award

Acorn Computer Group is a proud recipient of The Queen's Award for Technological Achievement for 1992. Acorn gains the award for the development of the ARM 32 bit Reduced Instruction Set (RISC) Microprocessor.

It was in 1983 that Acorn Computers first saw the need for a low-cost computer architecture to achieve the levels of performance that were only then achievable by 32-bit CISC (complex instruction set computer) designs, which were extremely expensive.

By 1985 a small team had an initial design of a new processor working in silicon. This was the world's first commercial implementation of the RISC approach to microprocessor design.

The first Acorn computer to utilise this new processor was launched in 1987, whilst two years later saw the birth of the low-cost A3000 computer which became the UK's fourth best selling computer within 18 months of launch. By 1991 five times as many A3000s were selling into the UK education market as any other computer.

To date, some 180,000 ARM-powered machines have been sold by Acorn.

Financial Results

An excerpt from the Chairman's statement:
"Acorn returned to profitability for the second half of 1991, and for the full year, following the previously reported disappointments of the first half. Sales regained the levels of the previous year, gross margins were improved and operating expenses well controlled. With interest costs falling as the group's cash position improved, the second half provided a net profit of £694,000 to give a full-year profit of £274,000.

Over the year, the group's opening bank overdraft of £6.5m was improved to a cash balance of £1.4m at 31 December 1991."

It is recognised at Acorn that how important and valuable is the part played by third-party hardware and software suppliers. Increasing stress is laid on the importance of our partnership with you, the Developers.

New fax number

There is now a fax machine located in the Support Group area. The number is:

0223 415439

Please use this number to contact us. The general Fulbourn Road fax number is still available, but contacting Developer Support via the Support Group number should result in quicker delivery.

Dates for your diary

A Games applications developers' conference took place on 21 May 1992. Both technical and marketing issues were discussed and those who attended contributed to what appeared to be a very useful day. Acorn's Sales and Marketing representatives took away some useful feedback to follow-up and the discussion of the draft Games Application Note will result in a released version. A copy of the revised application note will be sent to all Games Conference attendees.

Forthcoming events
Hardware developers' forum - Wednesday 17 June
General Registered Developers' seminars - 15/16 July

We will be writing to invite to the meetings those Registered Developers for whom we believe the content would be relevant. If you think we may not be aware of your activities in a certain area and that, as a result, we are in danger of missing you, please do contact us.

Terminology

Developers frequently enquire about how they should refer to Acorn's range in advertising and in manuals. It is appropriate to use 'Acorn's Archimedes range of computers' or 'Acorn's RISC OS-based computers' to refer to all models - Archimedes 305, 310 and 440, A3000, 400/1 series, A540 and A5000.

New Acorn products for Developers

As Acorn introduces new computer products to the marketplace, a number will be set aside for Developers who make early purchases.

Registered Developers will have the opportunity to place orders, at a special price, within a three month period from launch of a new machine. These special Developer offers will be on a one-per-company basis.

The usual ordering procedure will apply (payment with order, sent to Developer Support quoting Developer Registration number).

RISC OS 3 status

The ROM software and Application Disc software for the general release of RISC OS 3 has now been frozen. Many of you will have been involved in testing your applications with RISC OS 3.00, or later versions in EPROM, obtained through Developer Support. The frozen version of the product will be released as version 3.10 later in the year for

use across Acorn's 32 bit product range. Please remember that although the ROM code may now be frozen, it takes many months of ROM lead-time and kit assembly etc. before large quantities of the product can be made available to end-users.

Developers who have taken advantage of the Developer RISC OS 3 EPROM upgrades (with the promise, from Acorn, that full-release upgrade packs will be available to them at £25 plus VAT) will be given the opportunity to obtain version 3.10 in EPROM form, ahead of end-user ROM availability, by paying for their full packs in June/ July. These 3.10 EPROMs can later be exchanged for the full packs when these are available.

A3000 Technical Reference Manual

This publication is currently undergoing rework from Issue 1 to Issue 2. Please contact Developer Support if you have the manual and would like copies of the updated pages.

Updated Acorn SCSI card EPROM

An upgrade product, AKA33, has been produced for the Acorn SCSI card. The EPROM now includes the release version of CDFS, in addition to an updated version of SCSIFS and SCSI Utilities. Developers who have purchased the A540 Developer system will automatically be sent this upgrade to replace the pre-release EPROM sent to support the Cumana CD-ROM drive.

Other Developers may purchase the upgrade through Developer Support at £15.56 + VAT (£19.95 + VAT Retail). In addition to including version 2.10 of CDFS, the revised EPROM contains SCSIFS version 1.10 which provides more efficient transfer algorithms for some file sizes.

PC Emulator

Version 1.8 of the PC Emulator is now in stock as product code AKA70. The differences between this version and 1.7 include:

-Full VGA support (both hardware and BIOS)
-LIM 3.2 Expanded memory support
-A single task only configuration for small machines
-MDA emulation has been removed
-A few bug fixes

The RRP remains the same as for previous versions of the emulator, at £99.00 (exc VAT).

The product includes the new version of the emulator, MS DOS 3.30, a revised manual and MS DOS CD ROM extensions (MSCDEX 2.20).

Now that the PC Emulator version 1.8 contains full VGA support, CD-ROMs using VGA graphics, including some of those listed under the NCET scheme, can be fully utilised.

Memory requirements
The memory requirements for version 1.8 (when running without VGA support) are the same as for version 1.7. DOS applications can be run in single-tasking mode in 1 MByte and in multitasking or single tasking mode in 2 MBytes or more.

When configured to use VGA support, some applications may need more than 2 MBytes to run in multitasking mode but will still run in single-tasking mode; this includes the CD ROM titles for which you need 4 Mbytes to run in multitasking mode.

Upgrades can be obtained as follows:
•There will be an upgrade to 1.8 (AKA71) for existing owners of the PC Emulator available from June 1st until September 1st, 1992.

•The upgrade contains the 1.8 PC Emulator and manual only, with no CD ROM extensions or DOS disc: users retain their current DOS disc.

•The price for upgrading from versions 1.6 and 1.7 to 1.8 is £9.00 (exc VAT)

•The price for upgrading from versions earlier than 1.6 is £29.00 (exc VAT).

•The upgrade is available for all existing owners of the PC Emulator, whether originally supplied with MS DOS or DR DOS.

•There is a special upgrade offer for Educational institutions which are using version 1.6 or 1.7. This allows institutions to purchase one unit of the upgrade for £9, and copy the disc enough times to replace all their units of 1.6 and 1.7.

Manuals
The PC Emulator 1.8 manual (AKJ35) is available separately at £8.00. (no VAT payable) through Acorn's network of dealers.

How to obtain your upgrade
If you are unsure about whether or not you need to upgrade your existing PC Emulator, please contact Acorn Direct's technical help-line on 0933 279300. To take advantage of the upgrade offer please send your name and address, your original PC Emulator disc and your payment to: PC Emulator 1.8, Acorn Direct, 13 Denington Road, Wellingborough, Northants, NN8 2RL. Most charge cards, postal orders, and cheques made payable to 'Acorn Direct' are accepted. Official orders (accompanied by one disc only) are acceptable from educational institutions.

Important note: do not send your MS DOS or DR DOS discs back; these do not need to be replaced and are not part of the upgrade offer.

New version of !Request v1.06

The version of !Request (v1.05) that was supplied with Newsletter 25 has a small fault in it which inhibits its use on RISC OS 2 machines. A new version has been included on the Developer Disc 26 enclosed with this newsletter. You should upgrade your master copy.

Distribution of System resources

To date, any Developer wishing to distribute !System, its modules and !Sysmerge with his own products has required a Binary Distribution Agreement to do so. Registered Developers may now ship these Acorn copyright items without the formality of an agreement provided that they comply with certain requirements:

- They distribute !System together with !Sysmerge and !Scrap
- They ship only the latest versions released by Acorn (the latest current versions for RISC OS 2.00 are provided on the enclosed disc)
- Acorn's copyright is acknowledged in the documentation
- Only appropriate Acorn modules are shipped in the modules directory.

Licences for Squash and ResourceFS modules

These RISC OS 3 modules are not directly suitable for use with RISC OS 2.00. Versions for use with RISC OS 2.00 machines are available. If you wish to distribute them with your products, you will require a licence to do so.

Squash and ResourceFS can be provided to you, on request, under the terms of a software evaluation agreement.

RAM available to applications under RISC OS 3.10

Developers should note that the general release of RISC OS 3, version 3.10, provides approximately 30K less RAM for applications than was available under RISC OS 2.00. It is, however, greater than that provided by RISC OS 3.00 in the A5000. Developers are asked urgently to review future, current and previous software titles to check whether they can work with the amount of RAM available under RISC OS 3.10. Where there are problem areas, please decide what action to take to handle customer enquiries and needs when using version 3.10. It is hoped that program changes can be made and upgrades provided by developers where problems exist. Acorn is investigating the provision of a small user utility to help recover, in a controlled way, areas of RAM such as FontCache which may enable large applications to continue to work. If you need assistance in understanding memory usage in the machine, making modifications to your application or making controlled recovery of unused RAM in the machine, please contact Developer Support. An Application Note for Games Developers provides guidelines for memory usage and is now available on request.

Screen Mode independence

Some Developers will have received a questionaire recently concerned with application RAM space needs and whether screen modes 27 & 28 (VGA resolution) are supported. With the increasing interest in the use of VGA-only monitors amongst our customers, it is perhaps worth reminding all Developers of the need to provide screen-mode independent applications wherever possible. In this way, customers can choose to use PAL TV monitors (e.g. modes 12 & 15) where budgets are constrained, but can move on later to VGA or multiscan monitors and select modes 27 & 28 to obtain higher resolution and less screen flicker.

See further notes on screen modes in the Technical News section of this newsletter.

Feedback from customers

Acorn has received a number of complaints over the past few months from customers who have contacted suppliers about products but have received no response to their communications.

To quote one person's letter, "...it does make a very poor impression and, with tough competition, offering good service is essential. It is difficult to buy the software in the first place and, what is more worrying, is the feeling that after sales service is not going to be provided.... You make good computers but with such an amateurish service being provided it would be hard to recommend them."

We know that many of you provide excellent service to customers. In the current economic climate, however, no one can afford to be complacent. We would ask you all to take a look at your response to customers and to implement any changes which seem necessary.

Marketing News

Education

General Manager UK Education

Peter Talbot joined Acorn in March as General Manager for Education. Peter previously worked for Commodore, heading up their Education marketing team. In a recent News Release highlighting a report about the success of the BBC A3000 in the UK Schools computer market, Peter commented that he was delighted with the figures presented and went on to emphasise the value provided by software companies and other support organisations.

Acorn Software Roadshow

The May Roadshow visited venues in Aberdeen, Stirling, Middlesbrough, Sandwell and South London. This strenuous, but very successful, event was much appreciated by those who attended. Roadshows give Educational Advisers, teachers and others involved in education the opportunity to see, at first-hand, some of the range of applications and hardware available for use with RISC OS-based computers. Only RISC OS-compliant (or, in very exceptional circumstances, compatible) software is shown on Roadshows.

Further Roadshows are planned.

ACCESS IT

The 1992 ACCESS IT conference takes place in Nottingham on 3rd and 4th July.

Consumer

Acorn User Show
This year's show takes place on October 16, 17 and 18 at the Wembley Exhibition Centre - but in a larger hall than in 1991!

A revised product directory will be published for the show. Please ensure that you have sent any corrections and revisions to Developer Support before the end of June.

Special Needs

Our colleagues in Acorn Australia are compiling a special needs resources database and invite information from Developers who wish to be included. Please send any relevant details to Krystyna Boughey here at Fulbourn Road for forwarding to Australia.

International

Exhibitions
Acorn was represented at Didac in Switzerland and at ECCO in Canada in early May.

Australia and New Zealand
The supply of third party products to Australia and New Zealand, via Ian Sewell here at Acorn in Cambridge, has got off to an excellent start. Some suppliers' products have proved very popular. Please understand that the choice of software is entirely that of the customers, not of Acorn. If customers require your products, they order them. If they don't, perhaps the customers do not consider them relevant for their purposes.

Corporate

Library of software for exhibitions
You are probably aware that Acorn often uses third party software on exhibition stands and at various events. Until now, it has been the practice of the organisers to contact relevant suppliers each time to request copies of software and for permission to use it and then to return the packs.

Colleagues in the Events Unit will, over the coming months, be contacting software suppliers to give them the opportunity to provide software (or demo versions of software) for event and exhibition purposes. The intention is that the Events Unit should hold a library of items (entirely separate from the library of third party software which resides in Developer Support) which they have permission to install on a SCSI disc, as appropriate, for any event which they may be organising without neeessarily contacting the supplier on each occasion.

Brian Salter, Corporate Affairs Manager, has asked the Events Unit to update their library every six months. You may wish to update their copies more frequently to ensure that the software shown on Acorn's stands is up-to-date. Brian also asks me to remind you that the choice of what software is relevant for any particular event resides, in the last instance, with Acorn.

Technical News

RISC OS 3 and Font SWIs

In Newsletter No 19, the following paragraph appeared:
"Conversion of Acorn Font file contents to Draw path
Some applications software makes assumptions on the current RISC OS outline font file format when converting fonts to Draw paths. Some aspects of this may change in the future, and a new SWI call has been introduced to manage this in future Font Managers. Applications writers should start using this SWI call, even though a "SWI not known" error will be returned at present. The application should continue with its present technique when the "SWI not known" error is reported, otherwise it should let the new SWI manage the reading of the format. Please request a copy of the SWI if you make use of font file conversions."

At that time, we were not in a position to be more explicit. You should be aware that, for RISC OS 3, it is essential that the foregoing information is taken into account and the SWI (details follow) implemented.

Font_SwitchOutputToBuffer (SWI &4009E)

```
Font_SwitchOutputToBuffer
In: R0 = flags
  if R1 <= 0, then bits 0..31 are reserved (must be 0)
  if R1 > 0:
     bit 0 set => update R1, but don't store anything
     bit 1 set => apply 'hints' to the outlines
     bit 2 set => include skeleton lines if required ; NYI
     bit 3 set => produce sprites for bitmapped
                                        characters ; NYI
     bit 4 set => give error if bitmapped characters
                  occur (this bit overrides bit 3)
     bits 5..31 reserved (must be
  R1 -> word-aligned buffer,
    or 8 initially to count the space required for a buffer
    or 0 to switch back to normal
    or -1 to leave state unaltered (i.e. enquire about
                                        current status)
    if buffer, [R1,#0]=0, [R1,#4] = size remaining from
                                        R1+8 onwards
Out:  R0 = previous flag settings
      R1 = previous buffer pointer (updated)
```

After this call, any calls to Font_Paint will be redirected into the buffer, as a Draw file structure. Each letter painted will be treated as a separate filled object, with the colours specified in the paint command.

If R0 bit 0 is set, output is not actually sent to the buffer, but the pointer is updated. This allows the size of the required buffer to be computed properly before allocating the space for it. Note that if R0 bit 0 is set, R1 must be greater than 0 initially (a value of 8 is suitable, since the buffer must allow space for the terminator and free space counter on the end). After filling the buffer, the data is between (original R1) and (final R1 minus 1), ie. on exit [R1 -> 0, <size remaining>] and these two words do not form part of the output data itself.

The rubout box(es) and any underlining are also sent to the buffer as a series of filled outlines. These will be in the correct order so as to be behind any characters which overlap them. The output will also take into account matrix transformations, font and colour changes, explicit movements, justification and kerning.

If R0 bit 1 is set, the character outlines have hints applied to them (at the current size) - this means that they are not really suitable for scaling later on.

If R0 bit 2 is set, the character objects consist of a group of two objects - one is the filled outline, and the second is the stroked skeleton.

Any characters which are only available as bitmaps will either be errored (if R0 bit 4 is set), ignored (if R0 bit 3 is clear), or represented as bitmap objects in the output (either 1-bpp or 4-bpp, with a palette to match the output colours).

In this way Draw should be able to turn on buffering, then proceed to draw an arbitrary text column in the appropriate position and size, ending up with a series of Draw objects which represent the same thing.

The set of objects that the Font Manager produces could easily be converted into a group by wrapping them suitably.

CodeCheck

With Newsletter No 24 you received copies of a program (called CodeCheck) from Barry Mallett to test for illegal code sequences in your software. Although these code sequences cause unreliability problems in existing hardware, they are much more significant on new ARM processors. Acorn intends using such a processor in the future. If you haven't already done so, please test your code for these illegal sequences as soon as possible. In case you cannot locate your copy of the program, a copy is provided on Developers' Disc 26 enclosed with this newsletter.

Expansion Card Specification

As a result of feedback from existing users of this specification, Acorn has updated the specification to Issue 3. Additional areas of design guidance have been added, together with detailed corrections or additions to bus timing charts. Copies of the updated specification were mailed to some Registered Developers just prior to this Newsletter. Please let Developer Support know if you need a copy of the revised document and have not received it.

Safety Testing Guidelines

We have produced an A5 booklet covering various aspects of testing with safety in mind. These will be provided to attendees at the Hardware Forum to be held during the summer. Please contact Developer Support if you require a copy prior to the meeting.

SCSI cards interrupts

Whilst investigating CD-ROM system performance, in particular when Acorn Replay is being used, we have found a problem with some third party SCSI cards. In the current competitive scene, with Acorn third parties looking for the ultimate in SCSI card burst data rate, overall system performance has sometimes suffered. Some cards are typically taking control of interrupts for too long, preventing other concurrent applications, e.g. Replay, from having adequate access to computing time. As a result, taking the Acorn SCSI card as a reference, the user of some third party SCSI card sees an overall decrease in performance, rather than the expected increase. We would ask all third parties to review the overall system performance of their cards and not to go just for the maximum file transfer rate. In addition to reviewing the interrupt behaviour of the card firmware, the use of suitable hardware buffering should also be considered if not already used.

Alternative screen modes

Please note that RISC OS 3.00 defines modes up to and including mode 46 which clash with some user-defined screen modes, intended for use with RISC OS 2.00, in the range 24 upwards, the result of which is a scrambled screen. This particular problem also affects the 'ModeEx' example provided with the 'DDE', which defines a mode 29.

Modes 0 to 63 are reserved for Acorn, modes 64 to 95 for third party applications and modes 96 to 127 for user defined modes.

Available mode numbers are so few that the decision was made, in the early days of RISC OS, not to reserve specific modes for Developers who applied for them. The first applicant wanted most of those available!

VGA monitors and beyond: the Low Resolution Modes

Only UK TV standard monitors and Multiscan monitors can manage the RISC OS modes like mode 12. Both of these classes of monitors are going out of fashion and their replacements can't scan at the 15625Hz line rate needed for these modes. Similarly, special MultiSync modes such as Mode 20 and Mode 39/40 which take advantage of the general scanrate capability of Multiscan monitors will not be possible on SuperVGA and XGA monitors.

In the future, monitors will be tuned to EGA/VGA/ SuperVGA/XGA etc. and are unlikely to do anything else reliably. Indeed, the cheapest monitors are already in this situation.

Within a horizontal line rate (such as an EGA/VGA monitor) there is still much choice - the number of bits per pixel and the number of pixels per line can be varied and the number of lines per frame may be variable. For example the EGA/VGA monitor supports 350, 400 and 480 lines per frame.

EGA/VGA adaptors use this capability to provide 320x200, 640x350, 640x400, 360x480 and 640x480 (colour) plus 320x350, 720x350, 720x400 (mono character mapped).

RISC OS provides fewer EGA/VGA modes: RISC OS 2 provides 640x480, RISC OS 3 adds 640x350 (and 640x200).

Please remember that customers may choose any monitor from the variety on offer. Your products need to be able to cope with this.

Acorn File Formats Working Group

Various working parties have been established comprising representatives from Acorn and from Registered Developers with the aim of discussing and producing recommendations on issues of general interest.

The File Formats Working Group met to consider generic text and numeric file formats. The results of their deliberations follow:

Archimedes Paragraph Format

General

The following section describes a simple file format which is suitable for Archimedes applications to use to help with the interchange of free text.

The paragraph file format is of value to users of multiple word processors and desktop publishing programs. It enables users to pass textual files between applications without having to worry about manually adding in or stripping out linefeeds.

The description given is not a formal specification but should give developers enough information to implement Paragraph format successfully. Several programs support it already. Additionally, recommendations are made to developers which go beyond the main standard but which will make life easier for users.

Status

This specification has no formal status as yet. If approved by the File Formats Working Group and Acorn, the specification will be put in the public domain. The working party recommends that the information should also be inserted as an appendix in forthcoming editions of the RISC OS Programmers Reference Manual and, possibly, the RISC OS Style Guide.

The following applications are able to load and save files in Paragraph Format:

EasiWriter
Genesis
Impression
Ovation
PenDown
PipeDream

Feedback on this document should be sent to Robert Macmillan at Colton Software, 2 Signet Court, Swanns Road, Cambridge CB5 8LA.

Description

1) Text is represented by the ASCII characters 32-255.

2) There is a linefeed (LF) at the end of each paragraph. On loading, the importing application is expected to format the text to the appropriate right margin or frame. Importing applications must be able to cope with paragraphs of any length.

3) Tab characters (ASCII 9) are used to denote column information. The columns are left-aligned and at a position determined by the receiving program, hopefully under the control of the user. Each line in a table is considered to be a paragraph and so has a linefeed at its end. As an example the following table

L1F2		L1F4
L3F2	L3F3	
L4F2		L4F4

would be represented by

[09]L1F2[09][09]L1F4[0A][0A][09]L3F2[09]L3F3[0A][09]L4F2[09][09]L4F4

There is an implicit assumption that text in tables does not over-run tab stops. Since there is no information about either the positions of tab-stops or the font sizes of the text, if text is exported which does over-run tab-stops, it may not be possible to format the text correctly in the importing application. However, if text does not over-run tab-stops when exported, the user will be able to align his table correctly in the importing application, simply by widening the tab-stops.

Applications which are both importers and exporters should be able to export a paragraph format file in the same format as the file had when it was loaded.

There is no other page layout information.

4) There is no typeface information.

5) It is wise for applications to accept any of LF or CR (carriage return) or CR-LF or LF-CR as valid paragraph separators so that text from other systems can be imported. It is also wise to provide options for saving CR and CR-LF in order that files may be saved to applications on other computers.

6) Control characters other than LF, CR and tab may be faulted, or accepted as text as considered appropriate.

CSV File Format

General

CSV (Comma Separated Values) is a simple file format which can be used to ease the transfer of numerical data between applications. Many Archimedes applications already support variants of CSV format.

A CSV file consists of the character codes 9 (tab), 10 (line feed), 13 (carriage return), 32-126, and 128-255. Each line in the file consists of a sequence of numeric or string data 'items', which are separated by commas. The entire file is interpreted as a rectangular array of integer or string items.

A syntax for the file can be loosely defined as follows:

```
CSV-file     ::= line [newline line]* [newline]
newline      ::= "\n" | "\x0d" | "\n\x0d" | "\x0d\n"
line         ::= comment-line | CSV-line
comment-line ::= "#" [any except one of "\n\x0d"]*
CSV-line     ::= element ["," element]*
```

```
  whitespace      ::= [" " | "\t"]*
  element         ::= whitespace [(integer | number |
quoted-string | unquoted-string) whitespace]
  integer         ::= [ "+" | "-" ] digit*
  number          ::= integer [ "." digit* ] [("e" | "E")
integer]
  digit           ::= one of "0123456789"
  quoted-string ::= "\""  [quoted-character*] "\""
  quoted-character ::= (any except one of "\"\\\t\n\x0d")
| two-quote-marks | C-escape
  two-quote-marks ::= "\"\""
  C-escape ::= "\\n" | "\\\\" | "\\"" | "\\t" | ("\\x" hex-digit
hex-digit)
  unquoted-string ::= unquoted-character*
  unquoted-character ::= any except one of ",\"\n\x0d"
  hex-digit ::= one of "0123456789abcdefABCDEF"
```

In this grammar ANSI C string constants are used to denote terminals, [] denotes optional elements, * denotes arbitrary repetition, () denotes grouping, | denotes alternatives.

Newline usually appears as just a "\n" character (code 10), as in a RISC OS plain Text file. Alternative forms are permitted, however, in order to enable data import from other platforms.

The aim of this document is to help establish a standard format so that all applications interested in simple numeric data can load files which have been saved by the others.

Related work has been done by a group under the direction of the NCET and is summarised in the technical bulletin entitled 'Resources for data monitoring and control for science education' published in March 1991. The document specifies a file format called Software Independent Data format (SID). A SID file is very similar to a CSV file but includes a header at the start which provides extra information.

Ideally, there would be common 'CSV' format across the range of applications covered by productivity tools and data monitoring and control. In practice, the recommendations given in the information below from the File Formats Working Group differs in some respects from the NCET SID file format. This is partly a reflection of the different backgrounds of the two groups. The motivation behind the SID format was to unite the manufacturers of control hardware and software and to extend an already existing control format for one computer platform to others used in Education. The experience which has gone into this CSV proposal has come primarily from developers of spreadsheets, databases and graphics packages. The differences between the two formats are generally small but currently appear justified for the two environments.

Status

Whilst many existing packages already exchange CSV files successfully there is no absolute standard from which to work. The recommendation in this document is a mixture of the de facto standard which has evolved from a number of packages for Acorn's Archimedes range of computers, and some extra recommendations which help act as a catch-all. Where there are differences with the NCET description these are marked in brackets [].

For the future, we need to work with NCET to produce as much convergence as possible. When a stable state has been reached and approval obtained from both the File Formats Working Group and the Acorn Key Developers Group, the format will be put in the public domain. It is recommended that it would also be inserted as an appendix in forthcoming editions of the RISC OS Style Guide and/or the RISC OS Programmers Reference Manual.

Feedback on this document should be sent to Robert Macmillan at Colton Software, 2 Signet Court, Swanns Road, Cambridge CB5 8LA.

Description

A CSV file consists of numbers and textual labels, with fields separated by commas and lines by a line separator. The following line demonstrates most of the characteristics that applications should support:

```
"text string 1",123.456,"text""string 2", +7.8E-9,"text,string 3"  {LS}
      (1)           (2)          (3)          (4)          (5)          (6)
```

1. Text strings appear in double quotes. However, it is also desirable for importers to be able to read text strings without quotes. [The SID description says little about labels and makes no allowance for quotes. The implication is that labels are field names rather than field data, which is truer of the control world than of the database world.]

2. Numeric fields do not appear in quotes. Numbers that are in quotes should be read as textual labels.

3. Double quotes characters can be included in strings and will be represented as two double quotes characters. The example given should be interpreted as text"string 2. [See 1 above]

4. Scientific notation should be understood, with both capital E and lower case e being acceptable. [Unspecified in the SID documentation.]

The full format of a number is

(i) optional spaces, followed by
(ii) optional + or - sign, followed by
(iii) a sequence of digits, optionally containing a decimal point, followed by
(iv) an optional letter e or E which, if present, is followed by
(v) an integer (the exponent) which is optionally signed.

Implementation note for C programmers: this specification conforms with the ANSI C standard library function strtod().

5. Commas may appear in strings without ambiguity. [but not in SID files]

6. Applications exporting files for other Archimedes packages should export a linefeed (LF) as the line (record) separator. It is desirable for importers to understand any of carriage return (CR), LF or CR-LF or LF-CR as valid line separators. This helps reading CSV files created on other systems. Applications wishing to export files to BBC Model B and Master Series

computers may wish to provide the user with an option to save only CR, and applications intending files for the PC world may wish to provide an option for CR-LF.

[SID files are defined to have CR-LF as the record separator. This is more a reflection of the practice of DOS packages than RISC OS packages. The best situation is for applications to accept any separator on loading and to offer a choice on saving.]

7. White space (spaces, tabs) either side of comma separators should be ignored. Spaces within numbers are erroneous. Applications may regard control codes other than LF and CR as erroneous. [Undefined in SID]

Some applications export tab separated values (TSV) files. Importing applications may provide an option to read tabs rather than commas as the field separator.

8. CSV importers, such as spreadsheet graphics packages, often expect the first line of the file to provide field titles. Exporters may take advantage of this although it would be wise for importers to be able to cope with a purely numeric file as might come from a data logging package, for example. [SID prefers upon field-title information to be encapsulated in the SID header.]

9. No application owns the CSV filetype and so no application should claim a CSV file if the user double-clicks on it. [Similarly SID files]

10. All textual strings which are introduced by a double quotes character should be terminated with a double quotes character. If an application encounters a line separator apparently in the middle of a string it should terminate the string at that point and start a new line. This is an erroneous CSV file and the receiving application may accept or reject it.

11. Lines within CSV files need not have the same number of fields. [Since one of the two obligatory header lines in SID is the data size (number of records, number of fields), this is presumably not allowed in a SID file.]

12. A trailing comma at the end of the line is neither necessary nor erroneous. [See 11.]

13. A line separator should exist at the end of the last line, but receiving applications should accept a file without one. [Unclear from the SID specification.]

14. Two adjacent commas implies a missing item. A line without items (with or without commas) is interpreted as a blank row in the value array (rather than being ignored). A comment line is ignored. A line ending with a comma implies that there is a missing element following the comma: this may affect the reader program's impression of the overall size of the array. (Conversely lines are separated by newlines, but the presence of a newline at the end of the file does not imply a blank line after it, because programs and humans can be lazy about terminating plain text files.)

15. Because of the diversity in CSV's history, your reading code may well meet files that do not conform to this definition. It is better to read as best you can than to produce an error complaining about a malformed file: in practice, this will probably serve the user better. Specifically, if you detect an error the conventional action is to skip to the next comma or newline, and try to continue parsing from there.

The exception to this is that numeric overflow is best reported to the user, as ignoring it might lead to serious miscalculations.

16. It is recommended that when producing program output you should always use quoted strings rather than unquoted ones. They are easier to output because some strings cannot be represented as an unquoted string (e.g. "123").

17. There is no limit to the allowable length for a string item. Reading programs are not expected to deal with integers that do not fit in a 32-bit integer, or numbers that do not fit in an IEEE 64-bit number.

The above is the suggested format for CSV files. The SID format is set and should not be extended. SID and CSV have separate filetypes.

CD-ROM

If you are developing a CD-ROM, it is important to remember that many of Acorn's educational customers have networks.

Developers of CD-ROM discs which make use of quantities of real-time information, such as with Acorn Replay, may need to consider supplying multiple copies of CD-ROMs for use with individual, local CD-ROM drives.

RISC_OSLib source code

Over the last year or so, Developers have often told us that they have developed their own modules of code to supplement the code already supplied in RISC_OSLib. We would like to review these code modules to determine the areas of extension needed in RISC_OSLib. If suitable, and if you are prepared to share them, we would also like to consider adding these code modules to the source Library already provided to Developers on request. A questionaire will be sent to known active users of C during May to see what items of code are suitable and can be made available to other Developers through Acorn Developer Support. Do contact Developer Support if you do not receive a questionnaire during May.

RISC_OSLib source release

Developer Support have just completed a new source release for RISC_OSLib that is available, on request, to Developers who use Desktop C. A Developer requiring the source no longer needs to sign a Non-Disclosure Agreement but must enter into a Software Evaluation Agreement with Acorn.

The source release also comes with a few patched source files that provide fixes to certain areas (some of which are described later in the newsletter).

Items collected during RISC OS 3 testing

Type 11 Objects

Drawfiles produced using the RISC OS 3.00 version of Draw save files with appended Type 11 objects. This can cause problems with other applications which accept Drawfiles but which fail to ignore unrecognised objects.

Applications which accept Draw files but do not allow them to be exported are 'safe' in discarding additional information which they do not handle. Applications which export Draw files should retain all the information so that the file can be read into an application which does use the additional parts.

The problem of Draw holding multiple type 11 objects is a bug and has been fixed for the RISC OS 3 version 3.10 upgrade release.

Please note that the Draw file format is extensible. If any applications in the field do not cope with extended formats, such applications will need to be fixed. For example Acorn's DTP application displayed this fault and a patch can be found on the Support Disc for RISC OS 3.

ADFSBuffers

There are a few known problems with the use of ADFSBuffers under RISC OS 2 and RISC OS 3.00. The problems are as follows:

RISC OS 2

Data not correctly read back from a file when the following sequence is followed:

open
truncate
close
reopen
write
read

The sequence here is the truncation, close, reopen, write and then read sequence.

RISC OS 3.00

Address exceptions and other errors on a disc whose sector size is less than 1K.

The problems that occur in RISC OS 2 (this includes RISC OS 2.01 as supplied with the A540) only occur in the extreme circumstances described above.

If a user reports problems with RISC OS 3.00 (e.g. corrupt files) then you should instruct him to configure ADFSBuffers to be 0, as described on page 10 of the current RISC OS 3 Release Note. This workaround will solve the ADFSBuffers problem.

The problem is not present in RISC OS 3 version 3.10.

Drawfiles and RISC_OSLib

There are a couple of problems with RISC_OSLib when rendering Drawfiles using the supplied Draw routines. These are described below:

Printing

There are some mistakes in the 'draw_render_diag' code that can cause print jobs to hang when the Draw object contains text columns. The reason for this problem is that RISC_OSLib makes calls to ColourTrans_ReturnFontColours and ColourTrans_SetFontColour with handles of -1 for the current font instead of 0. This can cause errors like ' Illegal font handle' during printing from which the user cannot escape.

This fault has been fixed in RISC_OSLib and the source file fix is available as part of the RISC_OSLib source release. A patch module called 'ColourUtils2' has been written which can be distributed to customers if they encounter problems similar to that of the above.

ColourUtils2 is included on Developer Disc 26 which accompanies this newsletter.

Rendering Drawfiles with coloured text

There is another small bug in the Draw file routines from RISC_OSLib (which has been introduced with Desktop C) which causes text column colours to become corrupted if text columns happen to contain coloured text. This fault has been fixed in RISC_OSLib and a source file fix is available as part of the RISC_OSLib source release.

!DrawEx

The example application !DrawEx, as supplied with Desktop C, has a small fault in it with regard to mode independence. The fault is that DrawEx does not call ' wimpt_checkmode()' to cache mode variables on start up nor after a mode change. This can cause sprites and system font text to be displayed incorrectly.

The way around this problem is to call ' wimpt_checkmode()' on start up and after mode changes.

PrintSpr and !PrintSpr

Some quite serious flaws have been discovered in example code previously supplied to some of you by Developer Support. We apologise for the inconvenience this may have caused some of you and thank those who pointed out the bug.

The example applications PrintSpr and !PrintSpr can actually cause the machine to hang when printing starts. The main reason for this fault is that both programs call PDriver_PageSize to read the printable area of the page and then specify this as the rectangle. The problem lies with the fact that PDriver_PageSize returns values in millipoints whereas PDriver_GiveRectangle expects its rectangles to be specified in OS Units. Both applications do the conversion between the different coordinate systems.

This fault has now been fixed in both applications and a new copy of PrintSpr has been included on Developer Disc 26 supplied with this newsletter.

OS_SerialOp bug in RISC OS 3.00

RISC OS 3.00 has a small problem with the OS_SerialOp 5 and 6 SWI calls which are used to read/write the

transmit/receive baud rates. The fault occurs where the previous baud rate is returned in R0 instead of in R1. The following code should get around the problem :-

```
OS_Sig% = INKEY(-256)
IF OS_Sig% <> 163 THEN
  SYS"OS_SerialOp",6,-1 TO ,OldBaudRate%
ELSE
  SYS"OS_SerialOp",6,-1 TO OldBaudRate%
ENDIF
```

This fault has been fixed in RISC OS 3.10.

Revised Released Modules list

The list of latest released modules has been extended to include RISC OS 3.00 modules. The new list is enclosed.

Please note that the list includes the modules present in the RISC OS 3.00 ROM set as fitted to A5000s. These modules are intended to work as a set, and have been tested as such. While some may have the same name as modules present in the RISC OS 2 operating system, they are not for use in such a system, and the reliability of such a system cannot be guaranteed. Some modules, for example irqutils, are merely shells to prevent the overloading of the RISC OS 2 IRQUtils module. Others may rely on enhanced facilities in other modules which are not present in the RISC OS 2 equivalent of that module.

Correction

PDriver_DeclareFont - Developers' Newsletter No 25, November 91

The above states that you should call PDriver_DeclareFont and if you get a "SWI not known" error, you should assume you are using an old printer driver and carry on printing.

This is no longer considered the best approach.

Bit 29 in the printer driver features word should be checked. If it is set (ie. 1), then the driver supports PDriver_DeclareFont and it should be called. If it is not set (0), then you should not call the SWI as the driver does not support it.

!Killer - further clarification

The latest issue of Archive has published some incorrect information about !Killer, which may lead to some confused customers telephoning for clarification. Here is the correct situation:

Acorn have released !Killer in two versions, 1.17 and 1.26. 1.17 was on the 'Extras Extras' disc, and 1.26 is being propagated in the following ways:

a) On the April RISC User magazine disc
b) On the April Archimedes World cover disc
c) It has been sent to all Registered Developers
d) It is available to registered dealers and educationalists from Acorn
e) It is available from Archive, instead of version 1.17

Version 1.26 should be considered to have succeeded 1.17. Anyone with 1.17 should be advised to obtain 1.26 and use that instead. Version 1.26 contains an integral 'timebomb'. No version of 1.26 will work after 30 June 1992.

Pineapple Software (UK 081 599 1476) will be selling future versions, i.e. those *above* 1.26. Version 1.35 is the current version.

The enclosed disc

Developers' disc No 26 includes the following items:

!Reserve v 1.00 - for reserving app names etc.
!Request v 1.06 - revised filetype request application
Codecheck - for checking for 'illegal' sequences in code
ColourUtil - patch for RISC OS 3 drawfile printing
PrintSpr - for printing sprites directly from an application
SystResour - current System, Sysmerge, Scrap and modules

The software on the disc is provided in good faith to Registered Developers. It has not necessarily been extensively testedand is provided "as is"; Acorn Computers Limited ("Acorn") makes no warranty, express or implied, of the merchantability of this software or its fitness for any particular purpose. In no circumstances shall Acorn be liable for any damage, loss of profits, or any indirect or consequential loss arising out of the use of this software or inability to use this software, even if Acorn has been advised of the possibility of such loss.

News from non Acorn sources

RM05 Archimedes i860 Accelerator

RM05 is a joint development between Nexus Electronics and Riverside Machines who manufacture the Accelerators for other platforms. RM05 is basically a very fast maths accelerator podule for Archimedes. To give you some idea of its capability, it is several thousand times faster than an Archimedes system running the FPE.

The RM05 is an Intel i860XR based accelerator podule for Acorn Archimedes computers. It is suitable for any Archimedes with full width expansion bus. Features include:

Archimedes expansion podule: no external hardware required;
64-bit i860XR processor: 25, 33, or 40MHz (50, 66, or 80Mflops);
2 to 48Mbytes fast DRAM: zero-wait-state same page operation;
16-bit fast bidirectional FIFO communication with Archimedes bus;
Onboard APX-compatible kernel in PROM: runs standard compilers;
Inmos link and subsystem control for network connectivity;
Easy expansion via industry-standard TRAM daugherboards;
Fast I/O expansion bus;
160 Mbytes/s processor data bus transfer rate.