

Contents

- Editorial
- Chairman's Chat
- Events
- UK Conference
24th September,
London
- Annual General
Meeting
- SIR Development
- SIR on the Web
- SIR Tip
- Silly Section
- Committee
Members

Editorial

It's been six months since the last issue of SIR Reporter, but we've not been idle. We've lined up a one-day conference in London on 24th September at a subsidised price of £40 per head! There will be a demonstration of the soon-to-arrive new version, SIR XS, as well as a mix of user papers, tips and tricks and a discussion forum, and of course a chance to catch up with other users. More details inside.

Also in this issue is an update from SIR Oz on progress with SIR XS and an outline of some exciting proposals they have for replacing SIR/Master with a new 'on-disk' lock strategy. SIR Oz would welcome your views on this – especially if you currently use SIR/Master.

If you've not been a regular visitor to the SIR Oz website, then the final feature in this issue should convince you to try it: there is a vast wealth of info available on bugs, workarounds, fixes and downloads. A regular browse is useful: if you have a problem with SIR you may find that someone else has already notified it and a fix may be sitting waiting for you.

Finally, thanks go to the contributors for this issue: Tony Reardon and Dave Doulton; and to my daughter Ruth for emailing me the extract for the Silly Section (author unknown).

Kathy Brooks
Kathy@forvus.co.uk

SIR is at <http://www.sir.com.au>
UK SIR Users Group is at <http://www.soton.ac.uk/~sug>

Chairman's Chat

Welcome to the first 2004 issue of SIR Reporter. This year is the 20th that I have been using SIR. What other product has users that have been using it regularly for so long? I know that I am not unique in using the product for this length of time as both my trainers on my SIR training course are still users and they obviously started well before me. I believe this proves what a good product it is. It is a shame we cannot convince others of its benefits.

This year there will be a one-day conference in London. This is mainly because the expected SIR/XS will not be fully available this year due to development time and work to fix some problems with SIR2002 – all of which are now fixed. See the article later in this issue about the update and bug fix web site. However there are a couple of other reasons: one being that there was to be an International Conference in New York this year but this has been postponed to May next year. The other is that we are rather short of committee members with time to spare to do the organising. If you would like to return to 3day conferences please consider volunteering or find someone else to volunteer to assist us on the committee. It does not take much when there are many of us but is hard work for a few.

Dave Doulton

Events

24 th September 2004	UK Users Group and AGM in London
May 2005	International SIR Users Conference in New York

UK Conference 24th September, London

We have been fortunate to secure space for our next one-day UK conference and AGM at the Royal Statistical Society in London. This means that we can offer a remarkably low fee of £40 per delegate and we hope this will bring a good turn-out.

Provisional programme

- 10.00 Coffee and registration
- 10.30 Welcome and AGM
- 11.00 User Papers and Tips&Tricks
- 12.15 Lunch
- 10.30 Demonstration of SIR/XS
- 14.15 Further user papers Q&A and open forum

Book online now at:

<http://www.soton.ac.uk/~sug>

If you are interested in presenting a user paper please contact me:

Kathy Brooks
Kathy@forvus.co.uk

Annual General Meeting

The Annual General Meeting of the SIR UK Users Group will take place on Friday 24th September 2004 at 10.30am at the Royal Statistical Society, 12 Errol Street, London, EC1Y 8LX.

The Agenda will be as follows:

1. Apologies for Absence
2. Minutes of the last Annual General Meeting
3. Matters Arising
4. Chair's report (Dave Doulton)
5. Treasurer's report (Randy Banks)
6. Motion proposed by the committee

*"That Article 4a of the Byelaws of the SIR UK Users Group be amended to read: A general meeting of the group shall be held at least once in **two** years. Written notification of the time and place shall be sent to all full members at least thirty days in advance."*

7. Elections

The committee currently consists of:

- *Kathy Brooks and Fran Williams who were elected at the 2003 AGM and continue in office until 2005 (Bylaws article 6(b))*
- *Randy Banks (Treasurer), Patrick Brown, Dave Doulton (Chair), Adrian Hodgson and John Lemon who were elected at the 2002 AGM and are now due for re-election (Bylaws article 6(b))*
- *Mo Reardon who is the ex-officio member appointed by the SIR distributor (Bylaws article 5(b)).*

There are thus up to four vacant positions to fill by elections (Bylaws article 5(a))

Randy Banks, Patrick Brown, Dave Doulton, Adrian Hodgson and John Lemon are eligible to stand again for committee membership (Bylaws articles 5(c) and 6(d)). Other nominations for committee membership will be received by the current chair (Dave Doulton) up to 10.30am on 24th September 2004 (Bylaws article 6(e)).

8. A.O.B.

SIR Development

Some of you may remember the days when bugs were not fixed until the next major version of SIR was released. Problems would be reported and acknowledged as an 'SPR' or Software Performance Report and workarounds would be suggested but a proper fix would have to wait. Some time ago we changed that approach and now our highest priority is to diagnose and fix reported problems in the current release.

SIR is a large and complex system with a wide range of features and sometimes bugs remain undetected by our testing procedures or by customer use for quite some time. Some bugs are relatively easy to reproduce once they are identified; others are difficult to pin down. Some bugs are simple programming errors that are fixed with changes to a few lines of code while others can be deeper problems requiring considerable work.

A key task before any problem can be solved is to reproduce it in as simple a form as possible and we hope that anyone finding what they think might be a bug can help. If we are sent simple working examples of problems, it means we can quickly pin down what is happening. If all we get is a description, then we must create an example and we might fail to reproduce the bug because we are lacking some key piece of information about the exact circumstance that gave rise to the problem.

A recent example might serve as an illustration of key information being overlooked. We had a report that secondary indexes were sometimes being dropped when a database was reloaded. This did not seem to happen on our tests. It was suggested that the problem only happened after a schema change that required a restructure but this didn't seem to be anything we could reproduce either. It was only when we were told that it only happened if the reload did not reload the procedure file (because it already existed) that we could find what was going wrong.

We have just released SIR2002 revision 35 and it is strongly recommended that you upgrade to this version. A following article '*SIR on the Web*' discusses information on the SIR website and the complete set of fixes in revision 35 is listed there. A major issue addressed by revision 35 is the updating of secondary indexes under Master.

The interaction of multiple users creating records with secondary indexes through master gave rise to problems under certain circumstances. Even though some users had experienced problems in this area, it took considerable effort to come to a testing set-up where we could consistently reproduce the problem. Once we were able to re-create the circumstances, we could then fix the problem.

This particular problem was deep and complex and required extensive work to overcome. In essence, when a new record is created, it is first added to the database with all of its keys but with no data. You might think of this as what happens on a REC IS command in PQL. Then real data values are set in the record in memory by commands in the record block. Then the record is re-written. You might think of this as what happens at the END REC IS command. If the record participated in a secondary index, then it was added to the index with dummy values when initially added to the database. If multiple users added records at the same time, this could give rise to conflicts on unique indexes and this was at the heart of the problems.

A number of other issues were dealt with at the same time. For example, you may notice that when you upgrade to the latest revision, the Verify

process now checks and reports on all secondary indexes.

The work on revision 35 has been extensive and has taken priority over SIR/XS development. In addition, since the two versions are now quite different, the changes for revision 35 have to be applied individually to SIR/XS.

One development that has come out of the work on Master is a proposal to replace Master in SIR/XS completely and to adopt a different strategy to concurrent updating.

Multiple Updaters in SIR/XS

It is proposed to replace the current MASTER technology with a new 'on-disk' lock strategy to control multiple simultaneous update processes.

The standard behavior would be as today - standalone. If a database is opened in stand-alone mode no locks are maintained. Any number of stand-alone readers can access a database. If a stand-alone writer opens a database, no other readers or writers are allowed - if any are already accessing the database when the stand-alone writer tries to open the database, the open will fail.

If a database is opened in shared mode then locks are maintained for readers and writers. Schema changes and certain utilities require stand-alone write access. Although batch data input and spreadsheet updates currently do not run under master, it is proposed that these will be able to operate in a shared environment.

We will allow a system setting that everything was shared (equivalent to switching to master mode) or allow this to be specified as clauses on commands. If a database is open in one mode (stand-alone or shared) then opens in the other mode will fail (if we find it is already open in shared mode, we could convert open for normal processes automatically to shared mode unless some new STANDALONE clause specified). A shared environment automatically encompasses transaction processing. A transaction is a set of related updates that are delayed as far as any other process is concerned until the end of the transaction.

While this will require considerable development effort it will have a number of benefits:

- No separate master to start/stop.
- The process is simpler and involves no inter-process communication. Swapping disk I/O for inter-process communication

should be much faster. There is no single threaded choke point.

- Logic does not have to be duplicated into a separate product (master) with the attendant problems of master specific bugs.
- The strategy should scale more easily allowing more simultaneous processes to access a database.
- Databases should be more robust. Physical updates only happen as a consistent group.
- Implements transaction processing in a simple way – there is no need to roll back incomplete transactions.
- A sharing reader will see updates that have happened since the reader began running but will only see complete transactions.
- Some extra utilities will run in shared mode

Currently a standalone read-only retrieval running at the same time as master sees a consistent 'old' version of the database, that is one without any updates done through master. Under the proposed new system, the retrieval would see the updated records as soon as the transaction completed.

Tony Reardon
tony@sir.com.au

SIR on the Web

Have you ever wondered what happens when you click on the SIR on the Web or press Ctrl+Shift+W in SIR2002. Well if you have why have you not done it?

If you do you will find your self at the web page <http://www.sir.com.au/users/>

Which looks like this

[SIR Home](#)

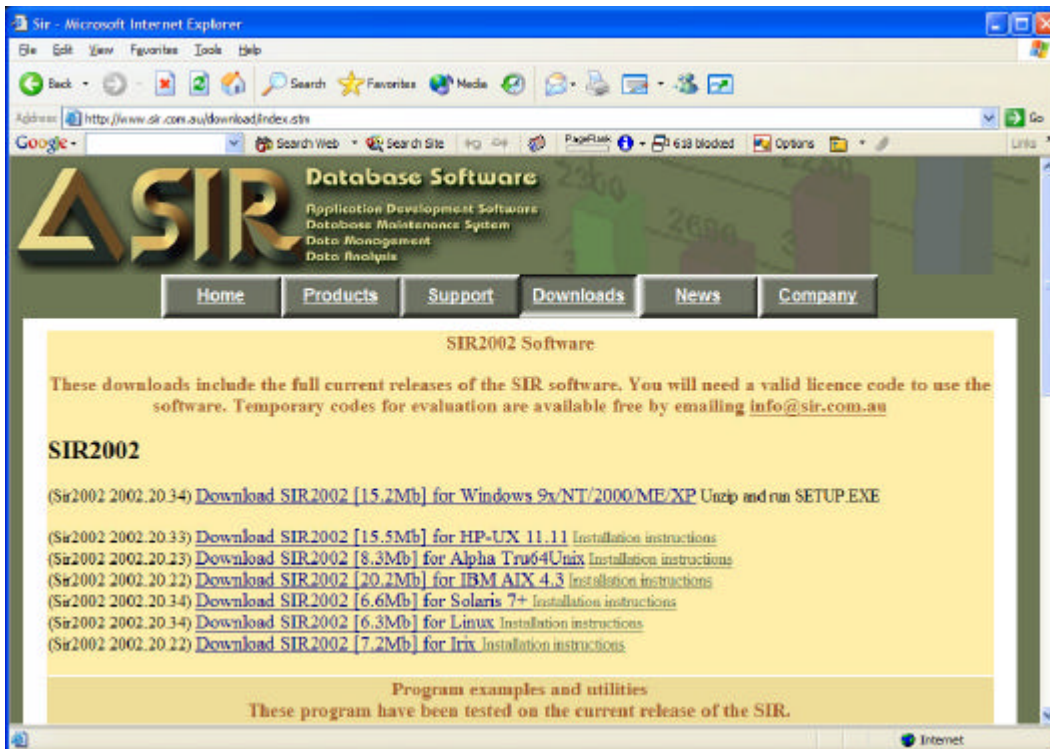
[Check latest versions and updates.](#) Compare the versions listed here with that reported in your Help, about SIR... Dialog.

[SIR2000 Problem Reports, Workarounds, Q&A](#)

[SIR2002 Problem Reports, Workarounds, Q&A](#)

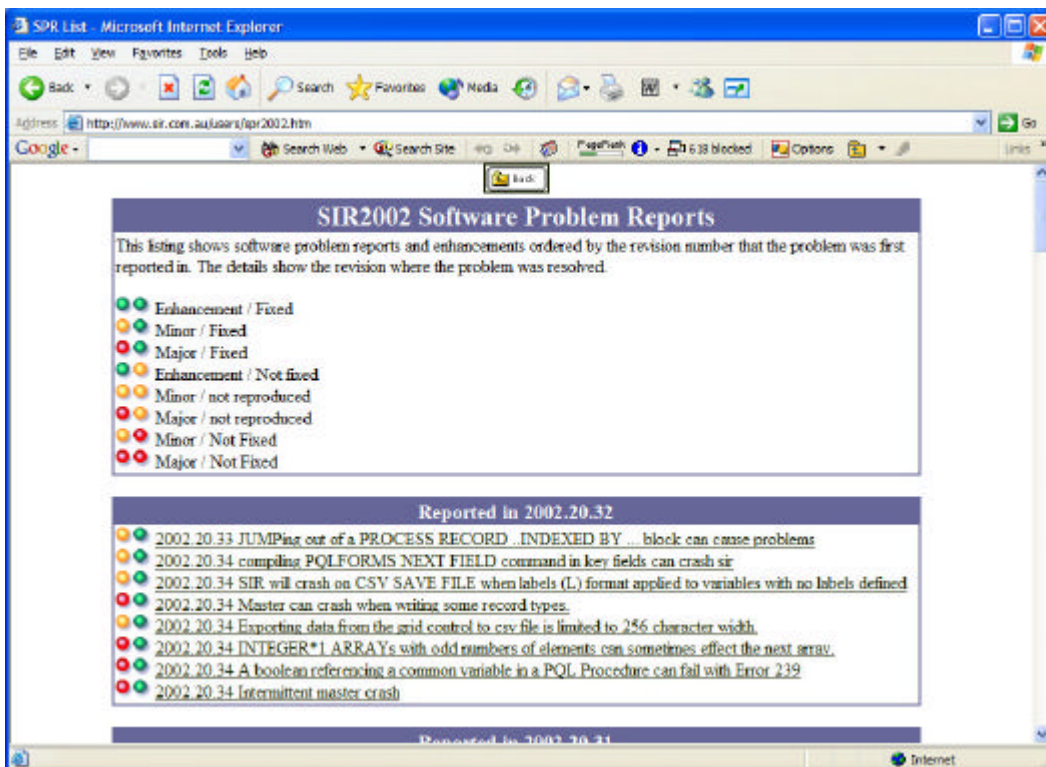
The features I am going to mention here are the [Check latest versions and updates](#) link and the [SIR2002 Problem Reports, Workarounds, Q&A](#) link. These should be regular places to check if you want to make sure that you are up to date with bug fixes and features.

The [Check latest versions and updates](#) link takes you to a page that looks like this



This shows you the latest release and allow you to download it along with instructions on how to install it.

The [SIR2002 Problem Reports, Workarounds, Q&A](#) link takes you to a page that will look like this



which shows the status of all bugs/problems/enhancements. The coloured lights show first whether it is major, minor or an enhancement and secondly whether it is not reproduced, not fixed or fixed. As of the time of writing all were fixed.

The entries are arranged in blocks showing which version the problem was reported in and within each block the level at which it was fixed. Clicking on an entry gives a screen with more detail often giving a description followed by an example program demonstrating the problem and a suggested workaround where available.

Some examples follow

A bug in indexed process recs classified as moderate the page shows the date and version in which it was found and the date and version it was fixed in. Along with an example program and workaround.

The screenshot shows a web browser window with a title bar and navigation buttons (Back, Home, Next). The main content area displays a bug report for the issue "JUMPing out of a PROCESS RECORD ..INDEXED BY ... block can cause problems". The report is categorized as "Moderate" and "Closed - Fixed". It provides the following details:

- Found:** 2002.20.32 09 FEB 2004
- Fixed:** 2002.20.33 13 FEB 2004

DETAILS

JUMPing out of a PROCESS RECORD ... INDEXED BY ... block can give you an execution error
message:
Database return code 8035 Record operation performed on case block.

After this if the program is run again, the thing hangs using tons of CPU and can't be interrupted with the 'Break' key. Ctl-Alt-Del to end the task must be used. If the retrieval is run only once then try to exit SIR you get a Windows ABEND notice.

EXAMPLE PROGRAM

```
retrieval
process cases
process rec employes
. get vaxs name
. write name
. execute subprocedure test
end rec
end case
=
subprocedure test
process rec employes indexed by name via (name)
. if(qgender=2)jump mylabel
end rec
mylabel:
end subprocedure
0
end retrieval
```

WORKAROUND

Use an EXIT RECORD command rather than a JUMP.

The next example shows a severe problem that was fixed within 2 days.

INTEGER*1 ARRAYS with odd numbers of elements can sometimes effect the next array.

Severe - Data loss Closed - Fixed Found:2002.20.32 05 APR 2004 Fixed:2002.20.34 07 APR 2004

DETAILS

Integer*1 arrays with odd numbers of elements can effect the first element of an array defined immediately after.

The example program is OK with INTEGER*1 ARRAY A (3) but produces bad results when the A array is size 5, 9, 13 etc.

EXAMPLE PROGRAM

```

PROGRAM
INTEGER*1 ARRAY A (5)
INTEGER*4 ARRAY B (3)
SET B * (2)
SET A * (1)
WRITE "A :" A(1) A(2) A(3)
WRITE "B :" B(1) B(2) B(3)
END PROGRAM

```

Correct output:
Start program translation
Start program execution
A : 1 1 1
B : 2 2 2
End program execution

Bad output:
Start program translation
Start program execution
A : 1 1 1
B : 1 2 2
End program execution

WORKAROUND

Set the second array after the first, or change the integer*1 to *2 or add an element to ensure the array has an even number of elements.

The next shows an enhancement that has been added to the system so you are missing out in functionality if you have not installed this version or later.

New function MEMINFO gets status of a member

New Feature Closed - Fixed Found:2002.20.28 ***** Fixed:2002.20.30 11 DEC 2003

DETAILS

New Feature. MEMINFO function returns information about a named member.

EXAMPLE PROGRAM

```

n = meminfo(fully_qualified_member_name,info_type)
Specify a fully qualified member name (i.e. with family/password if necessary) and use a member type (e.g. :T) if referring to specific member when possibly others of same name. This is either a quoted string or a string variable.
info_type can take values 1 - 8
1 = member count
Returns: 0 - does not exist,
1 :I,
2 :K
4 :O
8 :V
Additive values are returned if multiple versions present e.g. 3 = :T + :E
(If you specify a qualified member name with a member type, will return either 0 or 1,2,4,8 if member is present/absent regardless as to whether other versions exist)

2 = member password (public)
Returns: 0 = No password, 1 = Password + Public, 2 = Password + Not public

3 = Length in bytes of member
4 = Creation date of member
5 = Creation time of member
6 = Modification date of member
7 = Modification time of member
8 = Member out of date or not (Only really applies to :E or :O. If not fully qualified name, applies to first of any matching :E or :O or :V if present. Returns 0 for :T. Returns for :E
0 Not out of date
-1 Wrong version of SIR (This version is specific to the stored format - does not change with minor versions and may not change with major versions.)
-2 Default database not connected

```

I hope this has encouraged you to visit the web pages to keep yourself informed and free of known bugs. There is nothing worse than falling foul of a bug you could have easily avoided.

Dave Doulton

SIR Tip

A tip on being able easily to switch off some output or other code.

Here is a program that creates a csv file from record 1

```
retrieval
process cases
get vars all
process rec 1
get vars all
<1>write all
perform procs
end process rec
end process cases
csv save file filename='demo.txt'/variables=all
```

This program as it stands works fine and writes out the values before performing the procedure. However if it is called with the parameter \$c \$ the write is commented out.

Eg include file demo.pql (\$c \$) or call demo(\$c \$)

The \$'s are needed to make sure the space after c is passed through.

Dave Doulton

Silly Section

News for maths lovers

At New York's Kennedy airport today, an individual later discovered to be a public school teacher was arrested trying to board a flight while in possession of a ruler, a protractor, a setsquare, a slide rule, and a calculator. At a morning press conference, Attorney General John Ashcroft said he believes the man is a member of the notorious Al-Gebra movement.

He is being charged by the FBI with carrying weapons of math instruction. "Al-Gebra is a fearsome cult" Ashcroft said. "They desire average solutions by means and extremes, and sometimes go off on tangents in a search of absolute value. They use secret code names like "x" and "y" and refer to themselves as "unknowns", but we have determined they belong to a common denominator of the axis of medieval with coordinates in every country. " As the Greek philanderer Isosceles used to say, there are 3 sides to every triangle," Ashcroft declared.

When asked to comment on the arrest, President Bush said, "If G-d had wanted us to have better weapons of math instruction, He would have given us more fingers and toes. "I am gratified that our government has given us a sine that it is intent on protracting us from these math-dogs who are willing to disintegrate us with calculus disregard. Murky statisticians love to inflict plane on every sphere of influence," the President said, adding, "Under the circumferences, we must differentiate their root, make our point, and draw the line." President Bush warned, "These weapons of math instruction have the potential to decimal everything in their math on a scalene never before seen unless we become exponents of a Higher Power and begin to factor-in random facts of vertex."

Attorney General Ashcroft said, "As our Great Leader would say, read my ellipse. Here is one principle he is uncertainty of though they continue to multiply, their days are numbered as the hypotenuse tightens around their necks."

Anon

SIR UK User Group Committee Members 2003/2004

Randy Banks (Treasurer)

ISER
University of Essex
Wivenhoe Park
Colchester CO4 3SQ
Tel: 01206 873067
Fax: 01206 873151
email: randy@essex.ac.uk

Kathy Brooks (Reporter)

Forvus
53 Clapham Common
South Side
London SW4 9BX
Tel: 020 7819 1012
Fax: 020 7819 1010
email: kathy@forvus.co.uk

Patrick Brown

Institute of Occupational Medicine
8 Roxburgh Place
Edinburgh EH8 9SU
Tel: 0131 667 5131
Fax: 0131 667 0136
email: Patrick.Brown@iomhq.org.uk

Dave Doulton (Chair)

University of Southampton
Computing Services
Highfield
Southampton SO17 1BJ
Tel: 023 8059 3541
Fax: 023 8059 3131
email: D.C.Doulton@soton.ac.uk

Adrian Hodgson

ORC International
5th Floor City Point
701 Chester Road
Stretford
Manchester M32 0RW
Tel: 0161 877 6781 (switchboard)
Fax: 0161 872 3997
Email: adrian.hodgson@orc.co.uk

John S. Lemon

Aberdeen University Computing Centre
Edward Wright Building
Dunbar Street
Aberdeen AB24 3QY
Tel: 01224 273350
Fax: 01224 273372
Email: j.s.lemon@abdn.ac.uk

Mo Reardon/Tony Reardon

SIR PTY LTD
312 Mona Vale Road
Terrey Hills
NSW 2084, Australia
Tel: 00612 9450 2354
Fax: 00612 9475 1430
email: mo@sir.com.au
email: tony@sir.com.au

Frances Williams (Secretary)

ISER
University of Essex
Wivenhoe Park
Colchester CO4 3SQ
Tel: 01206 873568
Fax: 01206 873151
email: fwill@essex.ac.uk

You can email the whole committee by addressing your message to **sug.comm@essex.ac.uk**

