

## SATURN 11.6 Release Notes (11.6.03E)

This SATURN 11.6.03E, released in July 2024, replaces 11.6.03A.

SATURN 11.6.03E is the second full 11.6 release, superseding the last full release in June 2023 (11.6.03A) and all intermediate versions made available under the SATURNnext programme.

Simulation is unchanged from 11.6.03A and it produces assignment results identical to 11.6.03A in all the test runs we have undertaken. Some runs that previously hit exception errors and failed in 11.6.03A will now run through.

### Contents

<b>SATURN 11.6 Release Notes (11.6.03E)</b>	<b>1</b>
IMPORTANT change – SATPIJA/SATME2 and SUBFIX=T option.	1
IMPORTANT change – P1X - GRAF.DAT – affects ALL VERSIONS OF SATURN	2
CONTINUED warning – Select Link Analyses (SLA)	4
Summary of other changes from 11.6.03A	4
Reminder of the major changes in 11.6 from 11.5	5
SATURN EXES – Multi-core implementation	5
SATURN manual	6
SatWin	7
SatView / UFS2UFV	7
MSHELL	7
Technical Support	7

### IMPORTANT change – SATPIJA/SATME2 and SUBFIX=T option.

The main reason for this update is related to matrix estimation.

At the SATURN User Group Meetings in March 2024, we warned that we were looking into possible problems in the apportionment of the "fixed" and "passq" flows (flow passed from previous period that remain in the network) to the different vehicle classes, used to adjust the presented counts to give the target count to approach relevant for the matrix we are adjusting. As long as these flows are small, any "errors" in the calculation of the estimated target will also be small, so it is probably only a critical affect for highly congested models.

The final outcome is that we have only had to improve the estimation of the target in the case of using the option SUBFIX=T and with splits by class or level. Our understanding is that most users use the alternative SUBPQ=T option, and matrix estimation runs using this option are unaffected.

At the time of the UGMs, we feared that the problem affected both these options, so it is possibly worth expanding on our investigations.

Initial analysis showed that the apportionment was definitely not being calculated in the best manner for the SUBFIX=T option, and within the adjacent code, it seemed that the SUBPQ=T was similarly not adjusted. So logically both options had an "error".

Part of the confusion for users is that the adjustments of the targets are deduced within the heart of SATPIJA, and transferred through the SATPIJA (.UFP) file, allowing for such things as the difference on queued flow on the count site links in question. Though some data is presented in the SATME2 outputs (.LPM, .ME2), it is difficult to summarise the effects, and there is a presumption that simple factoring of those figures fully explains the difference between target and original count. Unfortunately, it doesn't. Note that this presentational problem has NOT been addressed in this update. So, the SATPIJA output, .UFP file, contains much more than just the P<sub>ija</sub> factors.

Two methods of checking were undertaken. One of the questions was, should we have been able to see the effects of such an error from the outputs. And we now believe the answer is, not obviously, as the whole matrix estimation process is a combination of iterative processes. We've looked at results from practical applications, with and without the adjustments, and find both "plausible" – with the counts broadly being matched by compensating adjustments in passq flows and the current assignment.

So, it was back to looking at the process in detail. The original independent verification of the SATPIJA calculations had been lost over time, and painstaking re-verification of the calculations was performed. Finally, this showed perfect matching of the SUBFIX=T figures with the new correction, but surprisingly the SUBPQ=T version NOT matching. This then went into a wider investigation of the code beyond where we "fixed" it, and low-and-behold we found that the passq flows had been adjusted on reading the data from files (as that was the only use of that data) but the fixed not (as the total fixed was used in its true form elsewhere in the program). This was well documented at the read in stage, but not at the point of application. So, the simple solution was to remove the new SUBPQ=T double application, and add further commentary at the point of applying the SUBFIX=T adjustment.

If only this software maintenance was easy!

## IMPORTANT change – P1X - GRAF.DAT – affects ALL VERSIONS OF SATURN

P1X uses the GRAF.DAT file to define the basic layout of the screen. This file was previously setup for a 4:3 ratio screen, but most users now use 16:9 screens (and have been putting up with an elongated view). We have updated the GRAF.DAT to, by default, better reflect the network extent for 16:9 screen users.

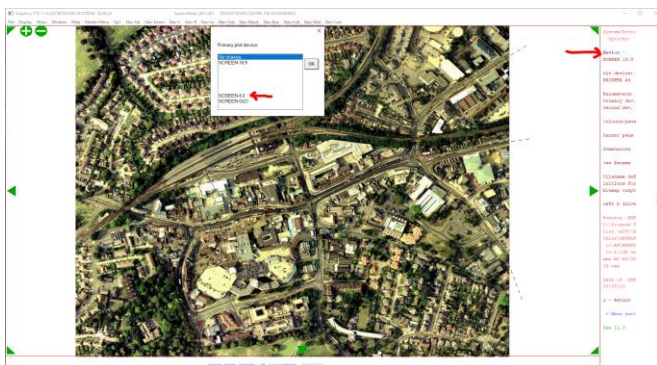
Note that this change in GRAF.DAT does not only affect the current release, but all previous releases.

However, for those with a 4:3 ratio screen or preference for the old setting, a facility in P1X has been re-invoked that allows you to switch between screen settings. The following example is based on the Epsom test network, and on a set up that works best in 16:9 mode, so switching to 4:3 mode will give a distorted picture.


To change the screen setting, from the "Master Menu", first select "System/devic".



Then, click on "Device" to open up the selection menu, select the "SCREEN 4:3" version.




For 11.6.03E, this will directly give you ...




which, in this case is the distorted version, as the better aspect ratio for this machine is 16:9.


For versions earlier than 11.6.03E however, there is an issue with P1X not refreshing the screen correctly, and you will see something like this:



The reshaping has begun, but the menu has moved off the right-hand side. [For other aspect changes, you may find that the menu has moved in from the right, and you will actually see two menu panels]. There is an “odd” way to force the refresh. From the top banner menu, select File, then Print Graphics,,,

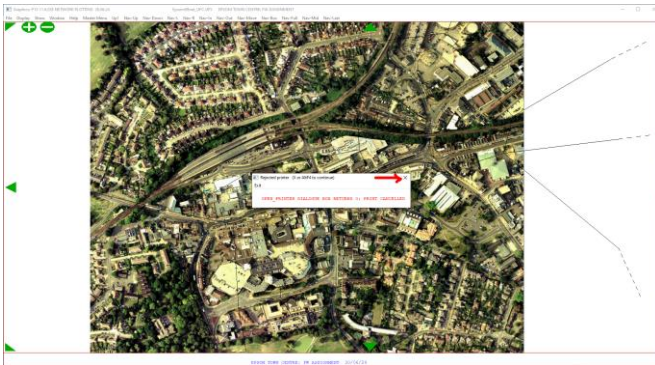


and just hit the cancel button...






which throws up the message...



which you can clear by clicking on the X, to get...



In this case, again, it is the distorted version, as the better aspect ratio for this machine is 16:9.

Note, we have also added a facility in SatWin (see below) to change the default GRAF.DAT file in use, but this may require help from your organisations' IT department.

## CONTINUED warning – Select Link Analyses (SLA)

The SLA routines for generating link and/or matrix-based outputs for a **single** link, node or turn using demand flows are consistent and replicate the main assignment results using either single or multi-core routines. The testing work suggests that the outputs based on **unfactored demand** flows are now generally correct for the more complex situations (screenlines etc) - this is **not** the default selection which is factored. The major problem is that for more complicated SLA, it is difficult to identify an appropriate factor, eg for screenlines, each selected could require a different factor. It is important to note that, if an appropriate factor cannot be calculated, an unfactored matrix is created, with only a warning within the LP file.

**Therefore, users are strongly advised not to rely on any other Select Link Analysis outputs unless additional independent verification checks are undertaken.**

Note that we have recently received a support query on whether 11.6 SLA does work with elastic demand assignments. We are investigating. Note that there have not been any significant changes in the simulation between 11.5 and 11.6, so, while investigation continues, the affected user is continuing to use 11.5.

## Summary of other changes from 11.6.03A

SATLOOK – when skimming "cost" matrices from a forest (using interactive menus or key-strokes), if a stacked matrix for all user classes was created, it had very wrong answers in all but user class 1, in all 11.5 and 11.6 versions to date. The work around was to skim each user-class individually – as these gave the correct results – and merge the matrices later. This long undiscovered problem (only necessary for a few specific analyses) has been corrected for 11.6.03E. The standard SATURN skimming **batch files** were (and, of course, continue to be) correct.

There have been no other major changes. A few fixes have been included, mainly where exception errors have been hit, P1X menus have gone off screen, etc.

A full list of the issues found, and bugs fixed, can be found in [Appendix E-Latest](#), the up-to-date version of which is available on our website.

No changes have been made in regards to UFF file creation and use, so the following commentary still is pertinent:

- Generally - UFF files DO NOT handle COMpliance factors - BUT will work when compliance is handled in separate user-classes. This will not be addressed before 11.7.
- SATLOOK skimming analyses with UFF files
  - They are taking a long time to run - multi-cored UFC++ are running significantly faster - so at the moment it is not recommended to use UFF files in skimming.
- SATPIJA analyses can work with UFF files
  - BUT we recommend using the SATPIJA batch file when doing so (and not the SATPIJA\_MC batch file) as the "IPART/NPART" was not envisaged as necessary with UFFs.
  - Similarly, the ALLIJ option does not work with UFFs.
  - We have not done a full investigation on the time efficiency of having to create UFF files and currently use them in single core, compared with the multi-cored UFC/UFC++ versions.

### Reminder of the major changes in 11.6 from 11.5

The major change from 11.5 is the introduction of new tree handling methods, which may have some speeding up of the convergence within SATALL, but are primarily intended to make considerably faster all later path-based analyses, such as select link, skimming, cordoning etc.

The update of the tree building has basically taken three forms:

One, the further expansion of using Dijkstra tree building algorithms, and a faster enhancement of that method (called **Celeste**) which, for large networks, can speed up assignment in all stages, including the convergence process. There is also a new process, **Lucien**, which as a path modification method can usefully be employed in the Area Charging process, as the diverting traffic can be treated as a diversion from the unaffected (paying or compliant) trees built.

The second, is the introduction of an update of the method of creating the "SAVEIT" assignments, the post-convergence assignment. This reduces the number of iterations that have to be saved in the cost path files (UFC) to reproduce the converged assignment. Though they have small differences from the previous UFC files, and are internally identified as UFC++ files, as the process is intended to replace the existing method, they are still given the .UFC extension.

The third form is the consideration that computing systems can now handle large files effectively (improved input/output etc), and instead of rebuilding the trees every time to do path analysis (using UFCs) we can now explore methods of directly storing the trees. We have implemented a method of storing the trees efficiently (but still large files) and these files are called UFF files. They are now provided as a replacement to the old UFO process (which had their difficulties in use, and the approximations used in generating them, and the difficulty in interpreting them when trying to do more complex path following, meant we were often surprised where the process could fail to give appropriate answers). Note, as with all brand-new methods, we are learning about the use of these new files. For select link analysis, UFF appear to be very fast in use and give similar answers to the UFC-based ones. However, currently, for skimming the analysis is slower than a multi-cored UFC equivalent. We're learning and refining, so please keep an eye on UFF results.

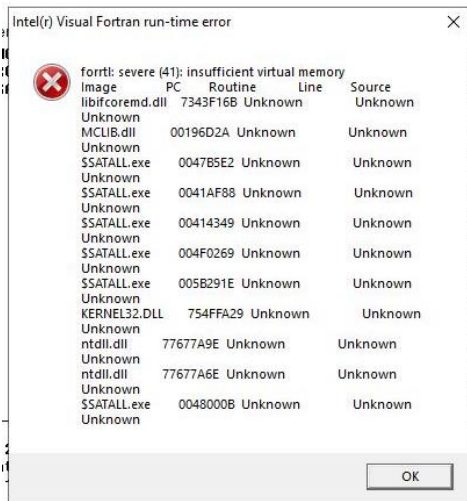
Most analyses requiring routing information can now take advantage of the new UFF files where feasible. Some, such as "X-SLA" which requires rebuilding of paths, cannot use UFFs.

**If you are new to 11.6, we recommend you read the 11.6.03A Release Notes for how to transition from running 11.5 to 11.6.**

### SATURN EXES – Multi-core implementation

If you have a multi-core version of SATURN, multiple threads are used for normal assignment in SATALL and SATUFE, for both TAC and non-TAC networks (ie creating UFC including UFC++ files). P1X and SATDB are both multi-cored when using UFC files, for SLAs (in P1X) and SATRAP/One-Song-to-the-Tune-of-Another (in SATDB and P1X). However, the use of UFFs should reduce the need for multi-coring the analysis programs.

Note though, with increased demands for memory in the multi-core programs and the machines having more available threads, we now are sometimes seeing the following error message.



Looking at the "Intel(r) Visual Fortran run-time error" window above, the detail message is "insufficient virtual memory".

This arises because the program is not acquiring sufficient virtual memory to run the MC version of the software, and doesn't mean that there is necessarily anything fundamentally wrong with the program, just that the machine isn't meeting the requirements to run.

When running in multi-core mode, the program has to replicate memory for each thread it is using. Our programs default to running with as many threads as possible (to a limit of 32). What exactly controls the virtual memory it can find is not easily quantifiable but is dependent on what physical memory the machine has, what also is running, how any paging is set up, etc etc.

Typically, you may get this message when running on a machine capable of running more threads. It may work on a machine handling 20 threads with no problem, but moving to one running 24 threads will require it finding 20% more virtual memory and potentially fail.

However, we can control the number of threads that are used by the environmental variable OMP\_NUM\_THREADS. All other things being equal (which they probably won't be), the solution on the 24-core machine may be to limit it to 20 threads. This can be done by including the following in your batch files / typing the following into the command line:

```

---
set OMP_NUM_THREADS=20
call SATURN_net mat
set OMP_NUM_THREADS=
---

```

or using the OMP NUM THREADS field in the Tools tab in SatWin.

You may need to experiment to determine what the appropriate setting for each machine / program is.

## SATURN manual

There have not been any changes to the 11.6 manual, with the exception of a formatting error in Section 6 being corrected, and updates made to Appendix E (Latest Issues) and Appendix N (Guidance Notes) to reflect the current situation. The "no changes" also means that we have **not** updated the manual to refer to the new entities of AtkinsRéalis and saturnsoftware@atkinsrealis.com. The installer will overwrite existing versions of the manual on your machine with the latest version of the 11.6 manual / help files.

## SatWin

This has been updated, though there has been no major change from the 11.6.03A release.

The changes are:

- MSHELL module
  - Updated so simple "browse" and "get" tools can pick up files selected in the normal SatWin manner
- Tool tab
  - Added SATURNacc environmental variable (to help us when investigating problems)
  - Added GRAF.DAT preference selection table

This means that the default SCREEN setting when opening P1X can be set to 4:3 instead of 16:9, but with many IT departments locking down access to software, you may have to ask them to run this for you with administrator privileges.

The new version will, of course, equally work with 11.6.03E or 11.6.03A (though you should migrate to 11.6.03E), and continues to function with older versions of SATURN.

## SatView / UFS2UFV

11.6.03E continues to be compatible with SatView v1.34.

## MSHELL

An updated version of MSHELL is provided with this release - see Appendix U of the manual for basic information.

It has got many more features for investigatory handling of UF files – but not yet "complete" enough in scope for a full product launch. But, do not be surprised to be asked to run MSHELL to help the SATURN team when investigating support queries.

Note - MSHELL may not run on machines with only 8 Gb of RAM. Unfortunately, it just does not start, with no discernible error to trap and hence message to report. In our tests, it does run on machines with 16 Gb of RAM or greater.

## Technical Support

If you require technical support, please do not hesitate to contact us at [saturnsoftware@atkinsrealis.com](mailto:saturnsoftware@atkinsrealis.com).

If you have any suggestions for what you would like to see in future versions of SATURN, please let us know. Please also let us know quickly of any problems you have, and we will look into them.

In the meantime, we wish you a successful continuation with SATURN 11.6 and thank you for your ongoing support.

**David Swain**

24/07/24      11.6.03E Release Notes.docx