

DARWIN® 5.1 Release Notes
December 2004
Southwest Research Institute

Summary of New Capabilities

The DARWIN 5.1 release includes the following new features:

- New Stress Intensity Factor Solutions
- 3D Surface Damage Analysis Enhancements
- General Risk Assessment Code Enhancements
- Initial Capability For Parallel Processing
- GUI Batch Processing
- Capability To Execute ANS2NEU Directly From GUI
- Finite Element Filtering

New Stress Intensity Factor Solutions

K solutions were updated for surface cracks in plates (SC17), surface cracks at holes (SC18), and corner cracks at holes (CC08). In the previous release (5.0), a side surface correction factor was used to adjust K for off center cracks. In the current release (5.1), this factor has been replaced by a weight function solution that uses reference solutions based on additional boundary element results.

Predicted life values for CC08 are similar for the 5.0 and 5.1 releases. Predicted life values for SC17 and SC18 for cracks initiating at the center of the plate are also very similar for the 5.0 and 5.1 releases. However, as the initial crack location approaches the edge (i.e., closer to corner), the life values are significantly different for SC18 (5.0 predicts between 20-80% of the life predicted in 5.1, depending on initial crack location). We believe the 5.1 solutions are more accurate, but these differences are currently under investigation.

A new computation scheme was implemented to reduce the amount of memory required during execution of `Flight_Life`. In the previous release (5.0), the reference solution data associated with the weight function approach were either stored in static memory or computed upon the initial call to `Flight_Life`. In the current release (5.1), the reference solutions associated with a specific crack type are stored in binary files that are accessed

by Flight_Life as necessary. This reduces both the overall memory requirement and execution time associated with crack growth life computations.

Note: stress values used for the weight function solutions are based on a polynomial fit of the finite element stress results. For future releases, stress values will be extracted directly from finite element stress results to improve the accuracy of the weight function-based life predictions.

3D Surface Damage Analysis Enhancements

In contrast with inherent materials in which the initial defect distribution (and associated probability of fracture) is based on the volume of material, the occurrence frequency of surface damage-based initial defects is modeled in terms of features per unit area (i.e., based on the number of defects per square in or square mm). The probability of fracture of a feature is based on the conditional probability of fracture (i.e., probability of fracture given that a defect is present), the area of the surface feature, and the area-based surface damage defect distribution.

The DARWIN 5.0 release provided an initial framework for surface damage assessment using 3D finite element results for 8 node brick elements. In the 5.1 release, support is provided for all ANSYS 3D finite element types identified by the TRMD Steering Committee as applicable to gas turbine rotors and disks.

A summary of the 3D element types supported in surface damage analysis mode in DARWIN 5.1 is as follows:

Ansys Element	Shape	Siesta Type
Solid45	Brick	BRI8
	Prism	BRI8
Solid92	Tetrahedron	TETS N10
Solid95	Brick	VANS N20
	Prism	VANS N20
	Tetrahedron	VANS N20
	Pyramid	VANS N20

An improved method for extracting stresses is now performed in the analysis code using 3D shape functions (in the previous 5.0 release, stresses were based on approximate values obtained from 2D slice planes generated in the GUI).

Notes:

1. Since mid-side node stresses are not provided in ANSYS stress results files, 3D stresses are interpolated using values from corner nodes only.
2. The 2D slice results associated with 3D surface damage models created in version 5.0 will need to be regenerated to be consistent with the improved stress extraction method introduced in version 5.1.
3. Surface damage assessment using 3D finite element models is currently limited to a single zone group.

General Risk Assessment Code Enhancements

Unlimited Number of Load Blocks

In previous versions, the maximum number of applied load blocks (steps) was a static value that could not be changed by the user. In the current release, the maximum number of load blocks is a dynamic value based on the number of load blocks in a given input (*.dat) file. Therefore, the number of load steps is limited only by computational resources (it was recently verified for an example with over 5,000 load steps).

Material File (*.mat) Enhancements

The following capabilities were added to the material file

1. Material specific life scatter factor. In previous versions, the life scatter factor was specified in the GUI for each zone. In the current release, the user can specify the life scatter factor directly in the material data file. If present in the material file, the life scatter factor is automatically applied to all zones that use it.
2. Material density. Material density is used in risk computations to convert the number of defects per mass to number of defects per volume of material. In previous versions, it was set to a fixed value for titanium materials. In DARWIN 5.1, the user can specify the material density for application to risk assessment of non-titanium materials.

Initial Capability For Parallel Processing

In previous versions of DARWIN, zone risk computations were performed in series (i.e., in sequential order). An initial parallel processing capability has been developed to perform zone risk computations in parallel, in which the DARWIN input file (*.dat) is

restructured for execution on multiple processors/computers (distributed processing). The capability consists of the following steps:

1. Create DARWIN input file (*.dat) using existing GUI.
2. Restructure DARWIN input file (*.dat) as master and slave files. A new GUI option “create parallel file set” (in the Tools menu) is used to perform this function. The master input file (filename-master.dat) is very similar to a standard (serial) DARWIN input file. It contains the information for all zones and the slave input file names (listed under the *DATABASE keyword). The user specifies a number of slave files (e.g., filename-SK.dat, where “K” is the slave file number and “S” indicates slave file) and the zones contained in each slave file (minimum one zone per slave file).
3. Distribute slave files to multiple processors/remote computers, execute, and retrieve results. The slave files can be executed on any available computers (each slave file contains all of the information needed for input to DARWIN).
4. Execute master file to combine slave file results into a single master results file. Upon execution, the master file searches the slave database files, retrieves the results for each zone, and combines the zone results to obtain the disk risk.

An effective use of the parallel processing capabilities of DARWIN requires a software infrastructure for disseminating and retrieving files between the master and slave computers. This software is available from a variety of vendors.

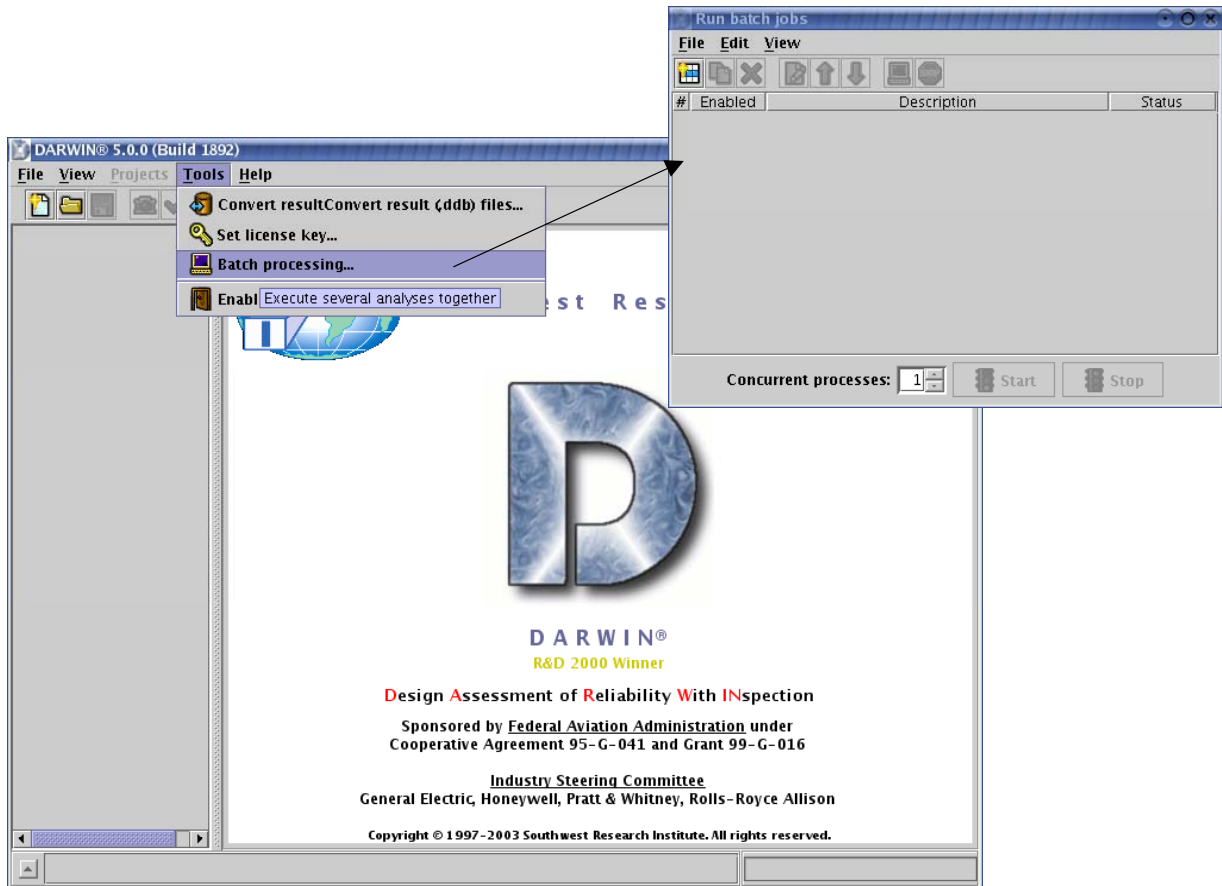
Notes:

1. The results from a parallel processing run will not exactly equal the results from a serial run due differences in the random number sequences applied to the zones. However, the differences will be insignificant if the number of samples is sufficiently large.
 2. The stress multiplying factor and the inspection time random variables are correlated across zones. In order to enforce this correlation for parallel processing, DARWIN now accepts an initial random number seed for both of these random variables in the input file. These seeds are only needed when using the parallel processing capability of DARWIN.
 3. The distributed processing software to automatically distribute jobs and retrieve the results is not included in this feature and must be implemented locally by the user
-

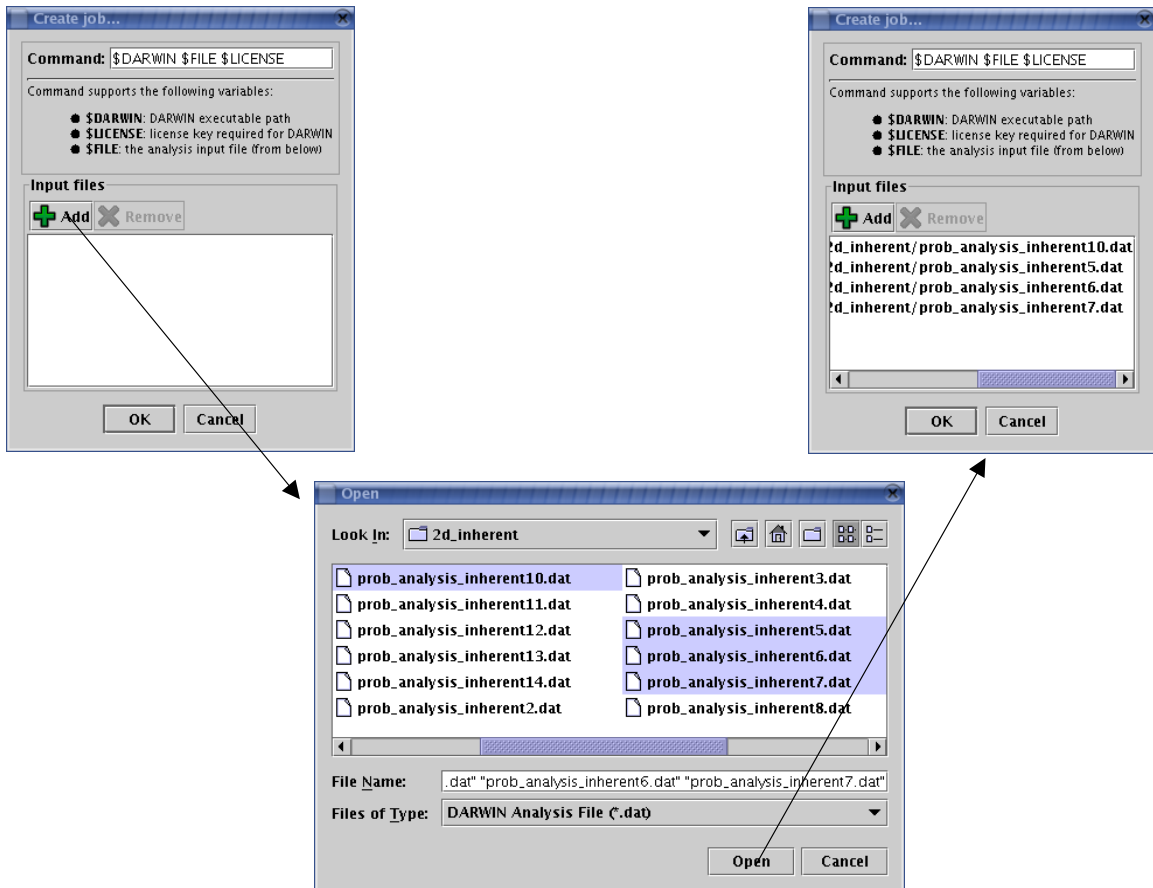
GUI Batch Processing

A new batch processing capability is provided in the 5.1 release that allows the user to execute multiple files in the GUI. This new capability consists of the following steps:

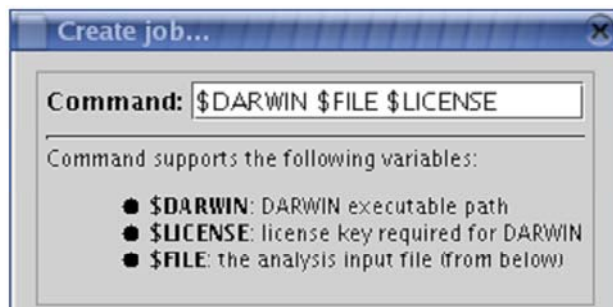
1. Select “Batch processing” feature in tools menu



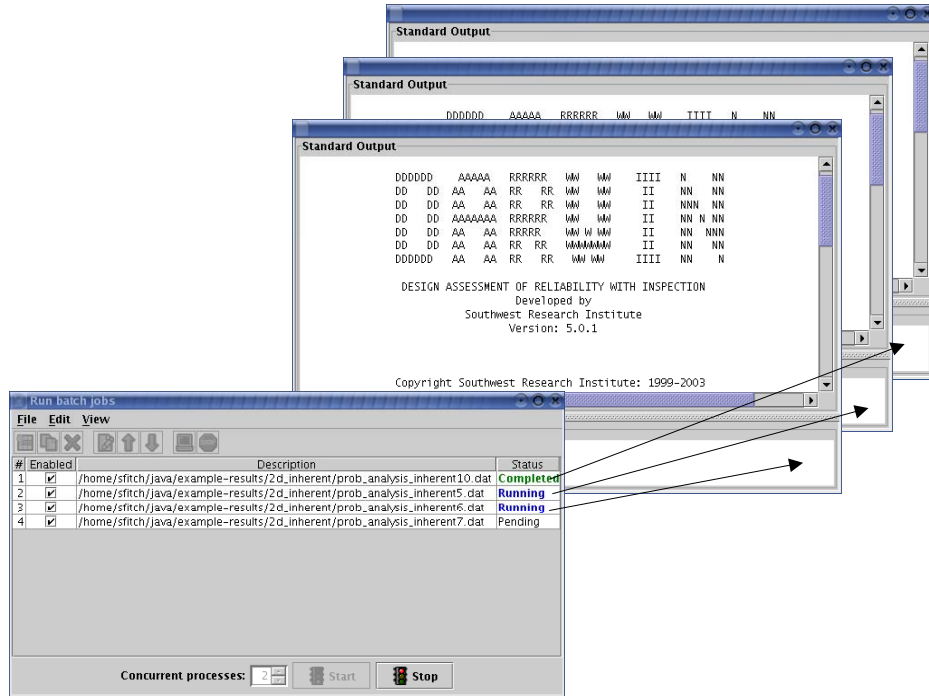
2. A job editor menu will appear - select files to execute via this menu.



3. Adjust the default execution commands as necessary to specify the path to the DARWIN executable and input files.



4. Monitor execution status of all files via the “run batch jobs” editor.



Note: Both the PC and UNIX versions support multiple concurrent processes (i.e., multiple jobs can be executed simultaneously on one or more CPUs).

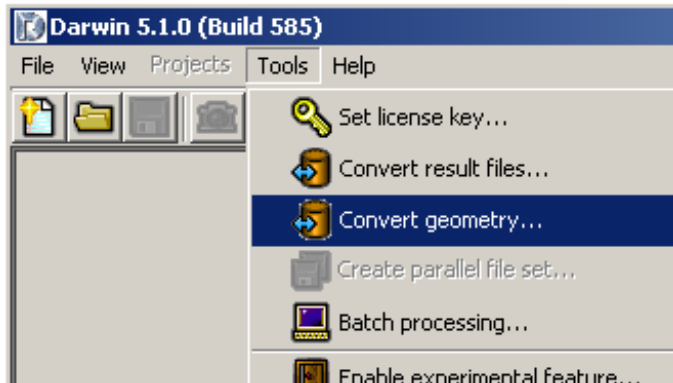
Capability To Execute ANS2NEU Directly From GUI

DARWIN uses a neutral file format to store and retrieve finite element mesh definition and results, stored in the following two files: (1) finite element mesh definition (node locations and element connectivity, *.uif) and (2) finite element results (stress and temperature results, *.uof).

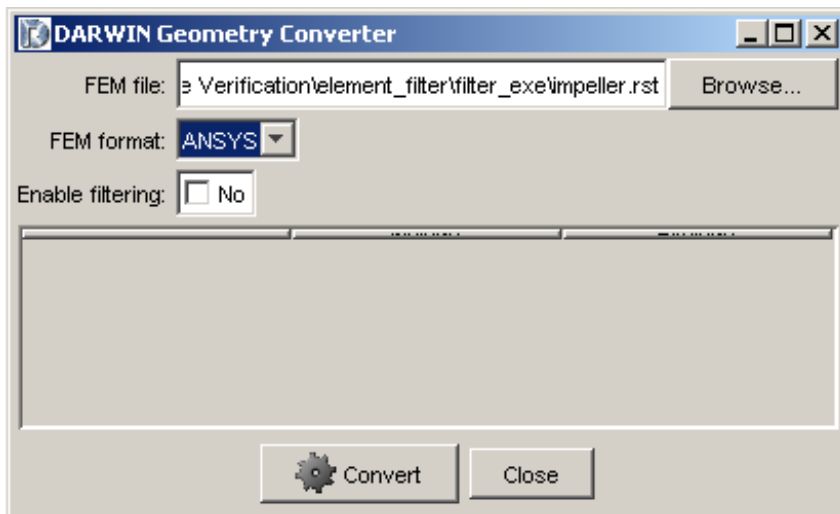
A capability is provided in the 5.1 release to convert finite element results to the neutral file format used in DARWIN. It currently supports conversion of ANSYS finite element results (*.rst), and is available on UNIX and windows platforms (not yet available for LINUX).

To convert an ANSYS results file (*.rst) to the neutral file format (*.uif, *.uof):

1. Click on the “convert geometry” option in the GUI tools menu:



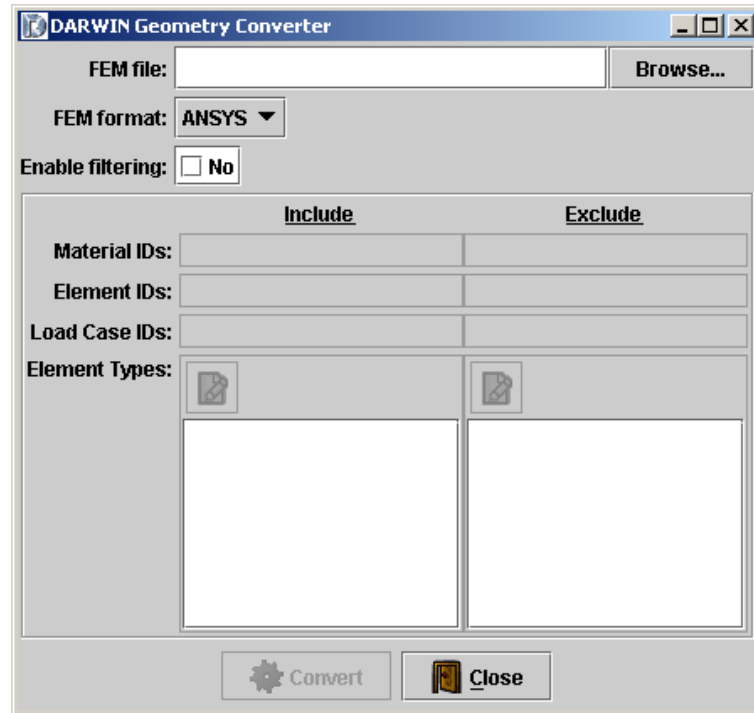
2. A “geometry converter” editor window will appear. Specify the finite element results file to be converted via the “browse” button in the editor window. Note: the geometry converter currently supports ANSYS finite element (*.rst) results format only.



3. To convert the files to neutral format, click on the “convert” button in the geometry converter editor window. The converted files (*.uif, *.uof) will appear in the directory of the original finite element results file.

Finite Element filtering

A capability is provided to allow the user to convert a portion of the original finite element model to the neutral file format. To access this feature, click “yes” in the box near “enable filtering” in the geometry converter editor window. A table will appear as follows:



When element filtering is enabled, the user can specify the material, element, and load case IDs and element types of the results that will be included in the converted neutral files. The syntax for specifying filtering ranges for these sections is identical. An *INCLUDE* directive causes *only* elements with field values that fall *within* the specified ranges to appear in the neutral files; all values outside the given ranges are ignored. An *EXCLUDE* directive causes only elements with field values that fall *outside* the specified ranges to appear in the neutral files; all values within the given ranges are ignored. Both *INCLUDE* and *EXCLUDE* directives may be specified for a given field, but there may only be a single directive of each type within a given field (e.g. cannot specify two *INCLUDE* directives for an `ELEMENT_ID` field).

Additional details regarding this feature are provided in the DARWIN User's Guide.