



CHAM

Convergence



Introduction

CHAM

- The conservation equations governing fluid flow are highly non-linear.
- They have to be solved by an iterative procedure, in which the predicted flow fields are continually improved from sweep to sweep.
- Hopefully this process will converge to a solution in which the discretised equations are satisfied with minimal error. This is known as “convergence”.
- But the iterative procedure may turn out to be a vicious circle of ever-increasing and eventually infinite values - divergence!
- This lecture discusses how to control the solution, and how to determine when you have achieved convergence.



Introduction

CHAM

- There are three stages to consider:
 - Do we have convergence? Three possibilities:
 - Obviously has converged 😊
 - Blatantly diverged 😞
 - Not sure 😐
 - If we don't have convergence, why not?
 - What can we do to help get convergence?



Do we have Convergence?

CHAM

- There are a number of ways by which convergence can be assessed:
 1. Graphical convergence plots (spot values and residuals),
 2. Additional graphical plots (maximum absolute corrections)
 3. Quantitative convergence data (RESULT file),
 4. Comparing plots from intermediate dump files.
- It is important to observe, while the solution is in progress, the **graphical convergence plots** which are displayed on the screen.
- These can give a very good indication of whether or not the run is converging.



Graphical convergence plots

CHAM

The default plots indicate

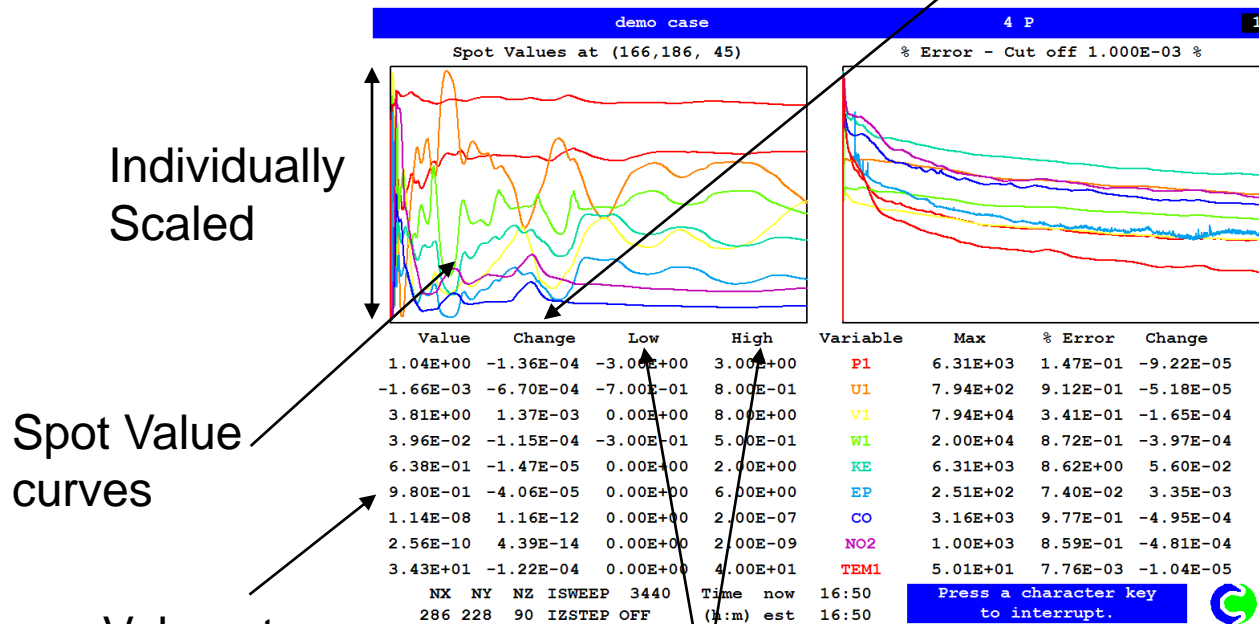
1. the spot values of the variables at the probe position; and
2. the sums of the absolute residuals for each variable - these values are normalised, as we shall see shortly.



Graphical convergence plots

Spot value curves

Change since last sweep



Value at probe location

Values representing top and bottom scale of each value


CHAM



Graphical convergence plots

- The spot value plot typically has three stages:

1. The curves bounce about a lot while the solution gradually establishes the basic flow pattern;
2. They stabilise and move asymptotically toward constant values;
3. They become perfectly horizontal, indicating convergence.




This plot is from a ventilation project.

CHAM



Graphical convergence plots

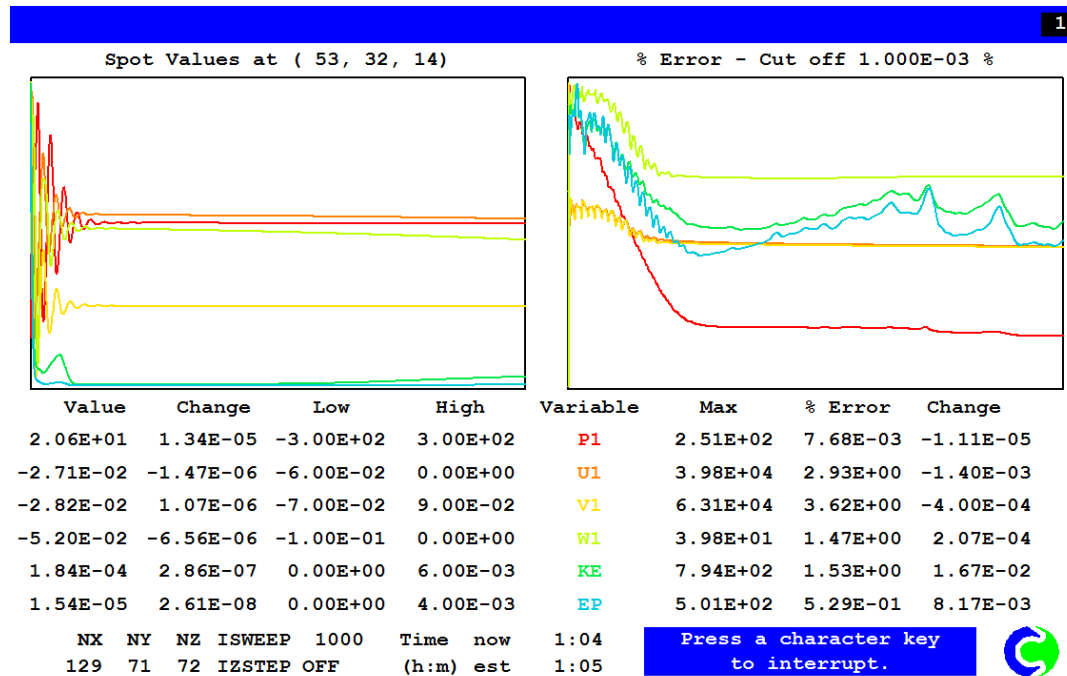
- Here is another example, from an HVAC project.





Graphical convergence plots

- If the spot value curves look suspiciously flat, it is possible that the solution is relaxed too hard (i.e. excessively damped).
- In such a case the solution may not be converged, despite the spot-value curves being almost flat. The residuals may be stuck and not reducing.
- It may then be helpful to consider applying lighter relaxation.




CHAM



Graphical convergence plots

- In a transient simulation, the changes at each time step are likely to be small, and so the spot-value curves will often just comprise stages (2) and (3) above.



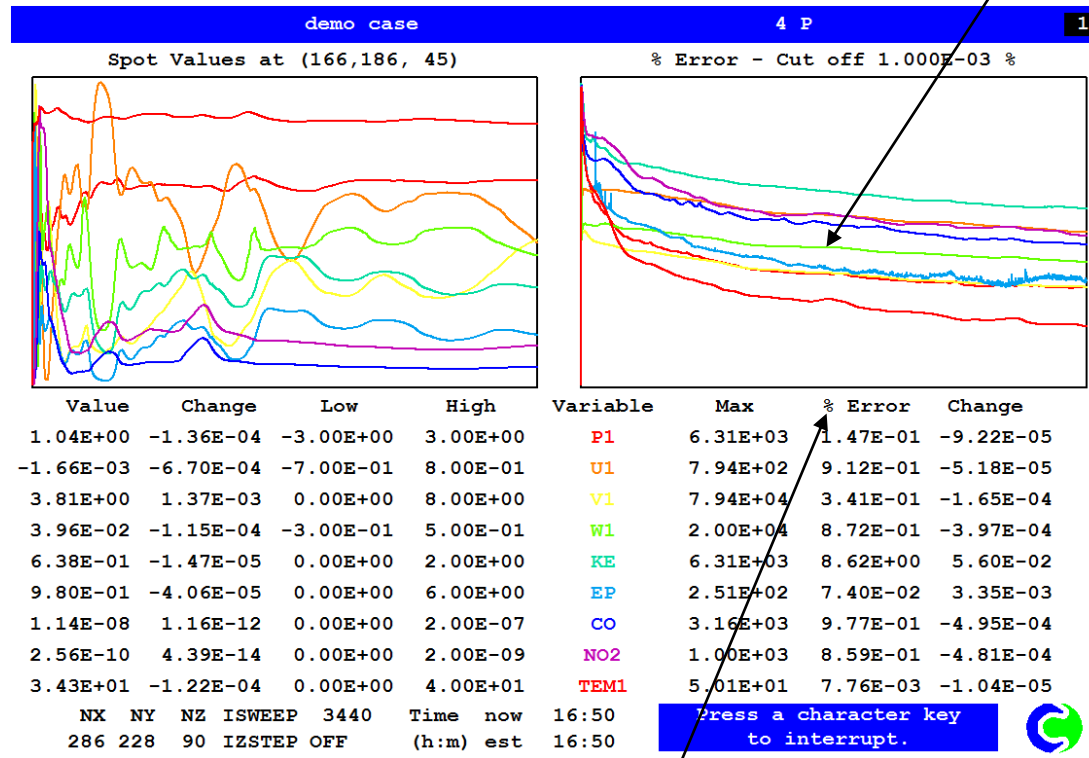
CHAM



Graphical convergence plots

Residual sum curves

Residual sum curves, with a logarithmic scale



Sums of absolute imbalances (“residuals”), normalised, expressed as %

CHAM



Graphical convergence plots

CHAM

- The normalisation reference RESREF is calculated in Earth, based on the flow rates in the domain.
- The residuals have the dimension of flow rate of the variable concerned.

mass continuity (P1 or R1)	kg/s
temperature (TEM1)	watts
velocity components (U1, V1, W1)	newtons
concentrations (e.g. C1, SMOK)	kg/s



Graphical convergence plots


CHAM

- As the run converges, **the residuals curves should move downwards**.
- Perhaps steadily decreasing; possibly "bouncing like a ball down a slope"; or perhaps with occasional spikes; any such pattern is acceptable, provided that the residuals reduce.
- **The "% Error" figures provide a quantitative indication of convergence.**
- The "% Errors" are also shown in the Result file, together with the unnormalised residuals and the RESREF values.



Graphical convergence plots

Residual sum curves



Residuals should move downwards with sweeps.

They can move straight down or “bounce”, provided that they reduce.


The “% Error” figures provide a quantitative indication of convergence. These can also be found in the “Result” file.

CHAM



Graphical convergence plots

RESFAC



This axis corresponds to the value RESFAC, which the user may set.

When all the residuals meet the bottom axis the run terminates.

RESFAC therefore represents a cut-off condition, and should be set appropriately.


Default value is 1.E-4, however a smaller value to prevent premature cut-off may be helpful, e.g. 1.E-5.




Graphical convergence plots

- For simple cases, the residuals curves may suddenly stabilise at near-constant values, with a slight "wiggle".
- This usually indicates excellent convergence!
- The “wiggle” indicates that rounding error has become significant, and therefore convergence cannot be improved – unless you use Double Precision.
- This may only be apparent if RESFAC is set sufficiently low.

Spot Values at (1, 13, 84)



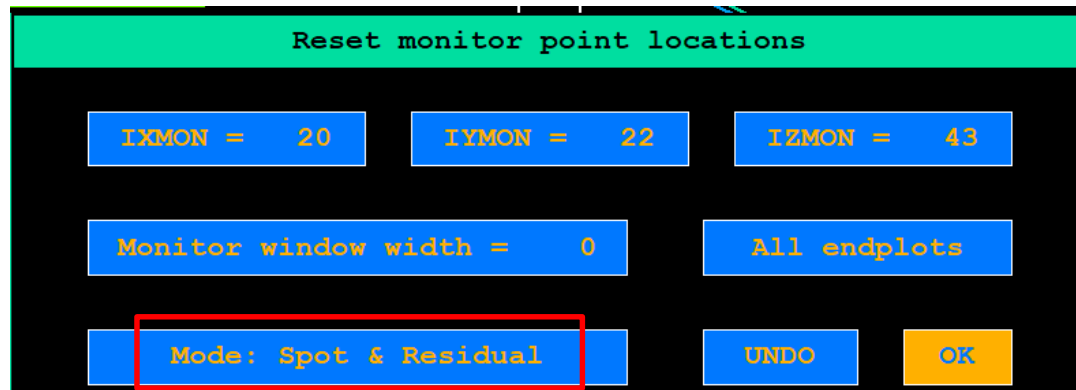
% Error - Cut off 1.000E-05 %





Graphical convergence plots

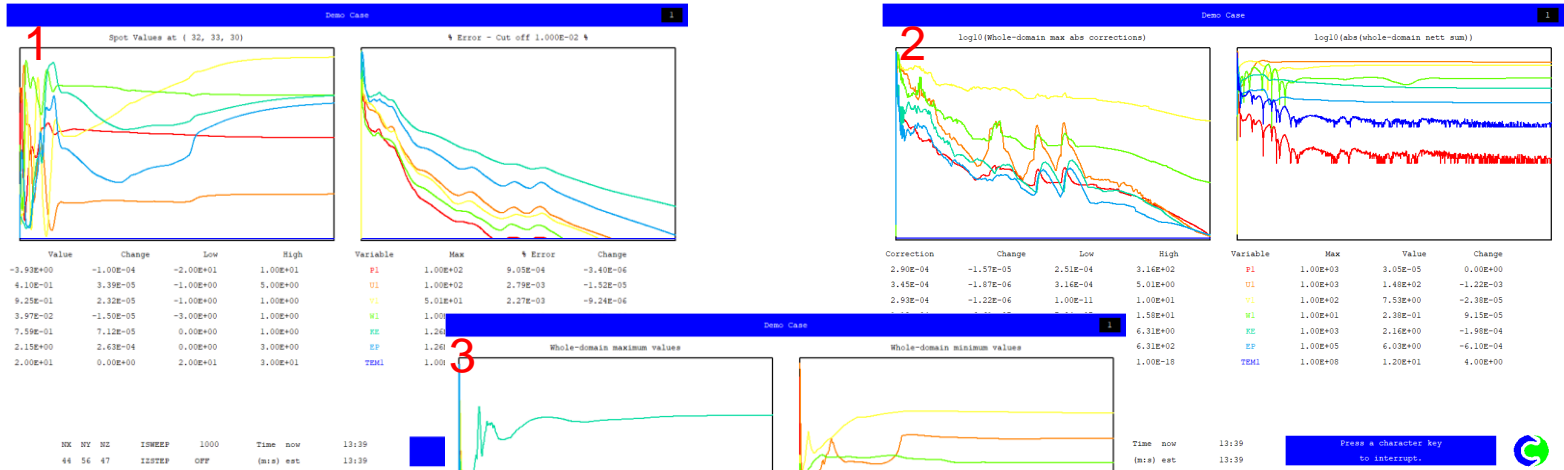
- We can click “Press a character key to interrupt” and then “Monitor” to change to a different convergence screen.





Graphical convergence plots

- There are three convergence screens:
 1. Spot values and residuals,
 2. Max abs corrections and sums of nett sources,
 3. Whole-domain max and min values.

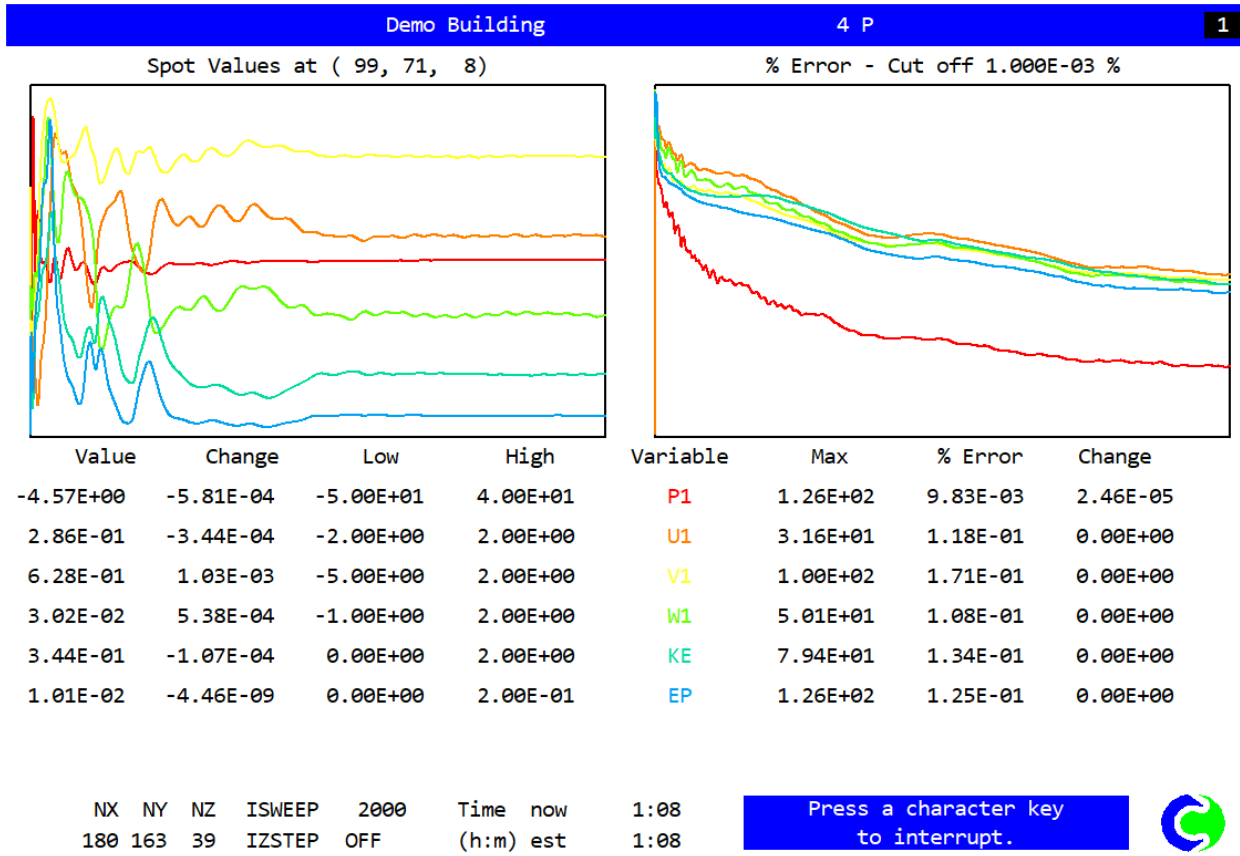


CHAM



Graphical convergence plots

- There are four convergence screens:
- 1. Spot values and residuals.

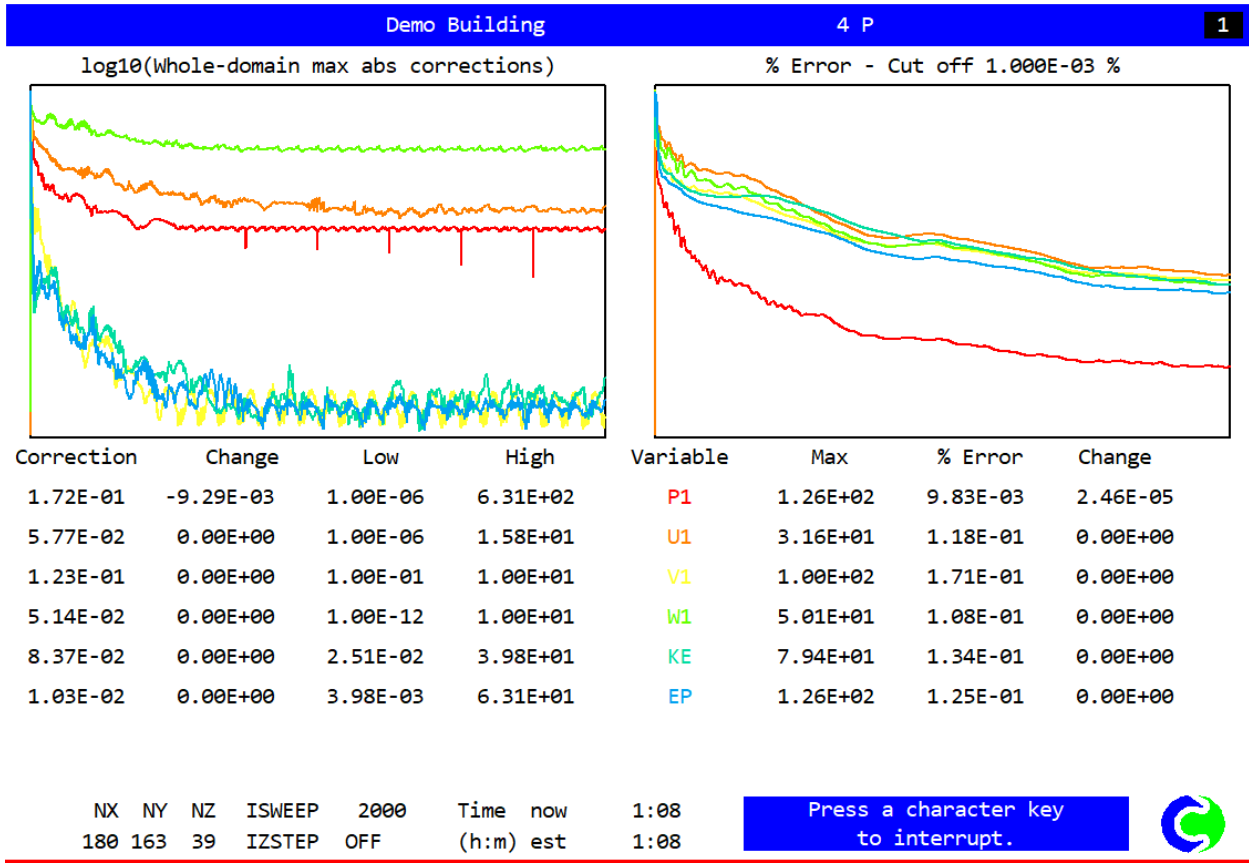


CHAM



Graphical convergence plots

- There are four convergence screens:
- 2. Maximum absolute corrections and residuals.




CHAM



Graphical convergence plots

- There are four convergence screens:
- 3. Domain maximum and minimum values.



NX NY NZ ISWEEP 2000 Time now 1:08
 180 163 39 IZSTEP OFF (h:m) est 1:08

Press a character key
to interrupt.

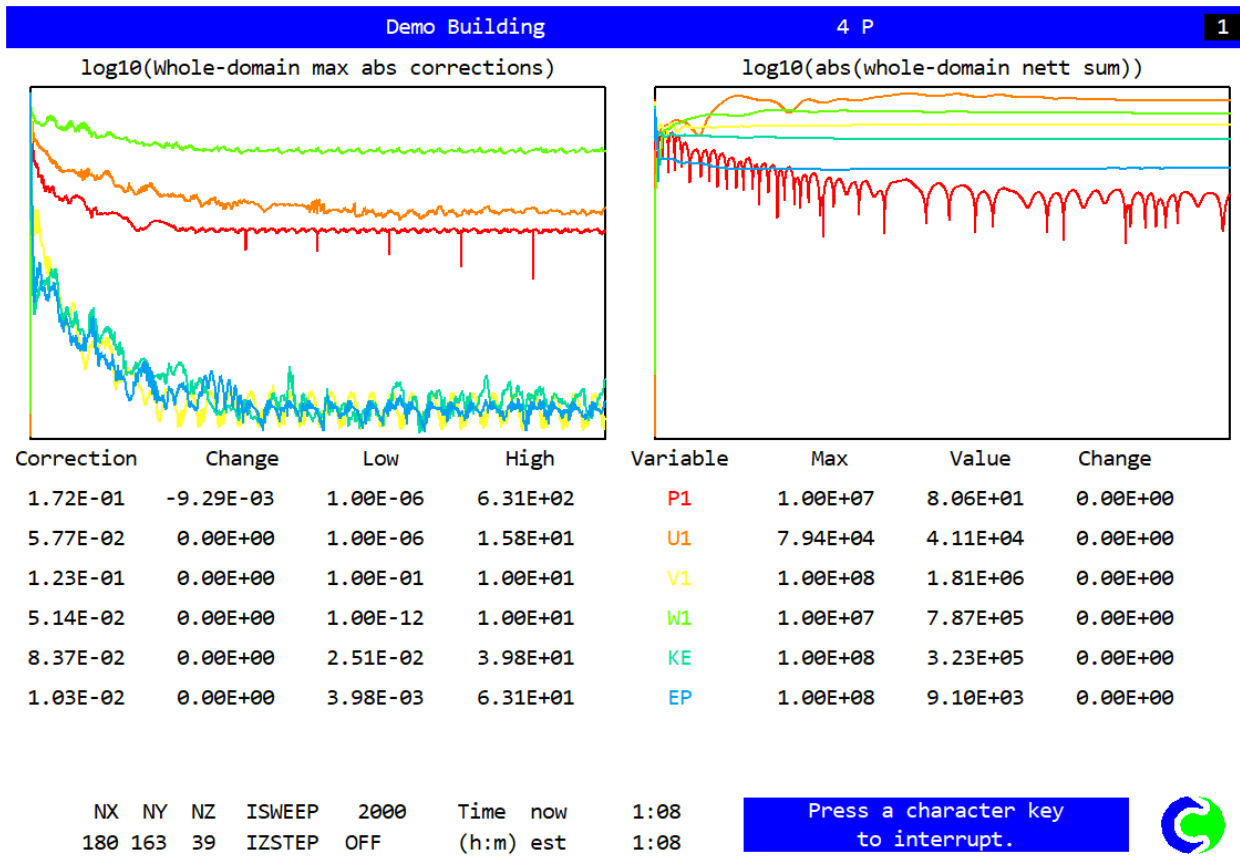


CHAM



Graphical convergence plots

- There are four convergence screens:
- 4. Max abs corrections + whole-domain net sums.



CHAM



Maximum absolute corrections

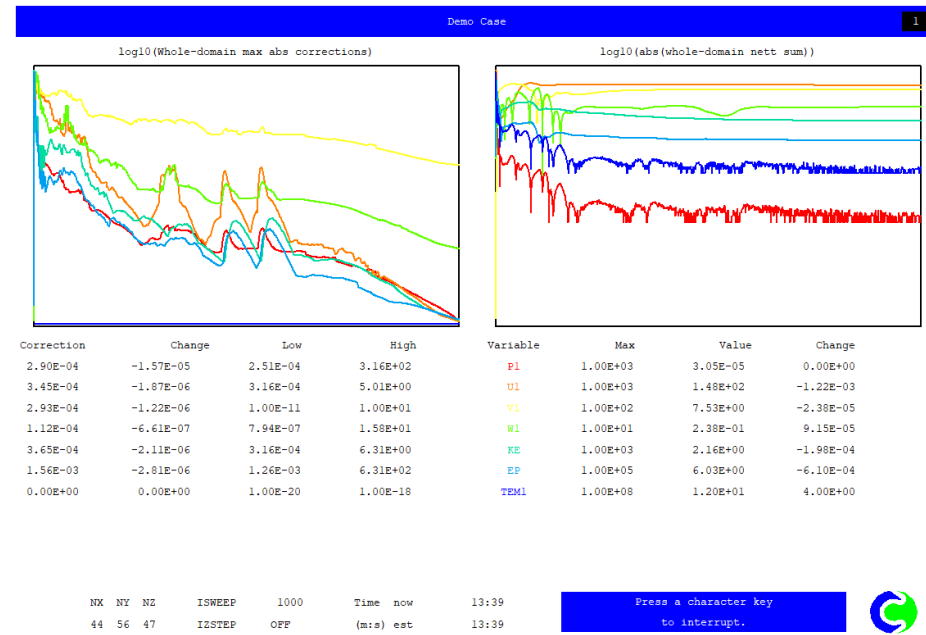
- Probably the most important are the **Spot values**, the **Residuals** and the **Max abs corrections**.
- The “Max abs corrections” are the largest absolute change in the variable at any cell in the domain.
- Note that all four convergence screens are saved when you “save as a case”.



Maximum absolute corrections

CHAM

- Shown on a log scale.
- Applies to a single cell.
- If steadily decreasing, this is a good indicator of convergence.
- If not steadily decreasing, this may not matter - e.g. this might result from a wall function problem at a single cell.
- The maximum abs corrections should be regarded as one of several different indicators of convergence.





The MINIRES file

- Appears at end of run and summarises the convergence information.
- Strongly advise looking at the curves and numbers, not just the colours!



Convergence information for run completed at 13:40:14 on Tuesday, 29 March 2022

on sweep 2000

Title: Demo Building 4 P

Working directory: D:\DRG\demos NEW\2022-01 city demo\WORKING


Variable	Source Sum	Residual Error (%)	Max Correction	Location (IX,IY,IZ)
P1 (Continuity)	80.6 (kg/s)	9.828E-03	0.172 (Pa)	(109, 114, 15)
U1 (X Velocity)	N/A	0.118	5.771E-02 (m/s)	(141, 70, 15)
V1 (Y Velocity)	N/A	0.171	0.123 (m/s)	(109, 113, 15)
W1 (Z Velocity)	N/A	0.108	5.138E-02 (m/s)	(141, 85, 21)
KE (Turbulent Kinetic Energy)	N/A	0.134	8.367E-02 (J/kg)	(142, 70, 14)
EP (Rate of Dissipation of KE)	N/A	0.125	1.030E-02 (J/kg/s)	(142, 71, 14)

An explanation of the above table is given [here](#).

To see the full source and residual printout, click [here](#) then search for "Sources and sinks".

Convergence monitor plots

Spot value & Residual, Minimum & Maximum value, Maximum correction & Residual, Maximum correction & Nett sum



CHAM



Convergence Printout The RESULT file

CHAM

- The **RESULT** file contains important convergence information:
 1. Nett source details for every object and each variable, and nett source imbalances;
 2. Sums of absolute residuals for each variable;
 3. Maximum absolute correction for each variable.
- These data are printed at the end of the solution.
- You can (and should!) arrange to print them during the run.
- Click “Print to RESULT file” in “Output” menu.
- They will now be discussed in turn.



Nett sources and imbalances

CHAM

- The "**Nett Source**" printout shows the inflow or outflow rate, for every object and for each variable.
- The total inflow and outflow rates are also given.
- These are then summed to give the nett source imbalance, i.e. the difference between what flows into and what flows out of the domain.
- This is an important indicator of convergence.
- You should check these for **R1** (mass flow rate), **TEM1** (energy), and any scalar concentrations (e.g. **C1**, **SMOK**, **AGE**).
- Also check that all specified sources are present !
- Do not check **U1**, **V1**, **W1**, **KE**, **EP** imbalances.



Nett sources and imbalances

CHAM

- The example below shows part of the printout of mass nett sources for an actual simulation (kg/s).
- The “IN” objects are mass sources - fixed flow rate.
- Final five are wind flows through the boundaries.
- At each of these, mass inflows and outflows are shown separately.
- Note the small “nett sum” showing good mass balance.

```
Nett source of R1 at patch named: VP49 (IN_187 ) = 1.224930E-03
Nett source of R1 at patch named: VP50 (IN_190 ) = 1.224917E-03
Nett source of R1 at patch named: VP51 (IN_193 ) = 1.224875E-03
Nett source of R1 at patch named: VP52 (IN_196 ) = 5.174413E-04
Nett source of R1 at patch named: VP53 (IN_199 ) = 1.224952E-03
Nett source of R1 at patch named: VP54 (IN_202 ) = 1.224886E-03
Nett source of R1 at patch named: OBCD (DOM_XMIN_O ) = 8.995358E+01 (Mass Out -2.248862E+01 In 1.124422E+02)
Nett source of R1 at patch named: OBCE (DOM_XMAX_O ) = 1.719158E+02 (Mass Out -1.750881E+00 In 1.736667E+02)
Nett source of R1 at patch named: OBCF (DOM_YMIN_O ) = 1.902424E+02 (Mass Out -6.526038E-01 In 1.908951E+02)
Nett source of R1 at patch named: OBD0 (DOM_YMAX_O ) = 8.814846E+01 (Mass Out -1.352015E+02 In 2.233500E+02)
Nett source of R1 at patch named: OBD1 (DOM_ZMAX_O ) = -5.400410E+02 (Mass Out -7.565121E+02 In 2.164713E+02)
pos. sum=540.320557 neg. sum=-540.040955
nett sum=0.279602
```



Nett sources and imbalances

CHAM

- The example below shows heat sources for the same case.
- The “HEAT” objects are heat sources (watts).
- The final five report heat fluxes through the boundaries (mass flux * Cp * temp in deg K).
- The average temperatures of inflowing and outflowing air are shown also (deg C).
- “nett sum” is the overall source imbalance for the energy equation.

```
Nett source of TEM1 at patch named: OCBA (HEAT_183 ) = 1.389819E+04
Nett source of TEM1 at patch named: OCBA (HEAT_186 ) = 1.389814E+04
Nett source of TEM1 at patch named: OCBD (HEAT_189 ) = 1.389810E+04
Nett source of TEM1 at patch named: OCC0 (HEAT_192 ) = 1.389789E+04
Nett source of TEM1 at patch named: OCC3 (HEAT_195 ) = 1.389777E+04
Nett source of TEM1 at patch named: OCC6 (HEAT_198 ) = 1.389800E+04
Nett source of TEM1 at patch named: OCC9 (HEAT_201 ) = 1.389814E+04
Nett source of TEM1 at patch named: OBCE (DOM_XMIN_O ) = 2.753249E+07 (Ave Out 3.257976E+01 In 3.200000E+01)
Nett source of TEM1 at patch named: OBCE (DOM_XMAX_O ) = 5.264343E+07 (Ave Out 3.234735E+01 In 3.200000E+01)
Nett source of TEM1 at patch named: OBCE (DOM_YMIN_O ) = 5.825582E+07 (Ave Out 3.234326E+01 In 3.200000E+01)
Nett source of TEM1 at patch named: OBCE (DOM_YMAX_O ) = 2.682487E+07 (Ave Out 3.323717E+01 In 3.200000E+01)
Nett source of TEM1 at patch named: OBCE (DOM_ZMAX_O ) = -1.659585E+08 (Ave Out 3.277274E+01 In 3.200000E+01)
pos. sum=1.660337E+08 neg. sum=-1.659585E+08
nett sum=7.5152E+04
```



Source imbalances

CHAM

- While a low overall source imbalance is a necessary condition for convergence, it is not sufficient.
- This is because even with a perfect balance of inflow and outflow, there may be "checkerboard" positive and negative errors in the individual cells which cancel out when you add them all up, but represent a significant error in the solution.
- We need a measure which looks at the errors in all the cells – the “sum of the absolute residuals”.



Sums of absolute residuals

CHAM

- The **sums of absolute residuals** are extremely important quantities when it comes to assessing convergence.
- They are displayed on the right of the convergence-monitor screen.
- They are also shown in the RESULT file, just above the “nett source” printout.
- They represent the sum, for each variable, of the individual cell imbalances (“residuals”) over the entire domain.
- When summing we take the absolute residuals, i.e. strip off any minus signs, in order to prevent any cancellation of errors.



Sums of absolute residuals

CHAM

- In the RESULT file, just above the “nett source” printout, you will see a printout looking something like this:

```
whole-field residuals before solution at sweep 4393
with resref values determined by EARTH
& resfac=1.0E-05
variable  resref  (res sum)/resref  (res sum)
P1        3.138E+03  6.159E-03           1.933E+01
U1        5.039E+04  5.531E-03           2.787E+02
V1        5.039E+04  6.104E-03           3.076E+02
W1        5.039E+04  5.124E-03           2.582E+02
KE        2.487E+02   1.209E-01           3.007E+01
EP        3.242E+04  9.339E-04           3.028E+01
CO        1.844E-05   1.798E-02           3.316E-07
TEM1     6.515E+08  1.535E-04           1.000E+05
```

- “res sum” - sum of abs residuals for the variable.
- units: P1 (mass), kg/s; TEM1, watts; U1, V1, W1, Newtons.
- “resref” – reference flow rate determined by Phoenics.
- **Middle column – normalised residual sums** - the same as the values shown on the convergence monitor screen.
- These are what you need to check.



Sums of absolute residuals

CHAM

- If you ask how small the normalised residual sums should be, there is no specific value!
- Of course, the ideal is to achieve whole-field normalised residual-sum errors which are below a defined percentage, for all the solved-for variables.
- However, such a strict criterion is often impractical, and a more pragmatic approach may need to be taken because of budget and time constraints.



Sums of absolute residuals

CHAM

- Based on our general experience, we would often find it acceptable if we can achieve the following:
 - residual errors of the order of 1% or less of the inflow rate, for the mass continuity, energy and other scalar equations, and
 - residual errors of the order of 10% or less of the relevant inflow rate for the momentum equations.
- This is because in general one finds that further reduction of the momentum error tends not to lead to any significant changes in the solution fields.
- If one feels uncomfortable with this level of error, then more sweeps are needed to grind the error down further.



Maximum absolute corrections

CHAM

- The “maximum absolute corrections” have already been mentioned.


```
max abs corrections for solved-for variables
for comparison with maximum & minimum values
variable  max-cor  high  low  ( IX  IY  IZ)
P1        4.064E-02  1.845E+04  -1.495E+02 ( 94, 95, 44)
U1        4.276E-02  9.073E+00  -1.623E+01 ( 126, 120, 66)
V1        5.851E-02  1.738E+01  -2.727E+00 ( 161, 181, 82)
W1        6.327E-02  8.521E+00  -5.809E+00 ( 161, 180, 82)
KE        2.162E-01  1.103E+01  1.000E-10 ( 174, 63, 41)
EP        1.912E+00  3.342E+02  1.000E-10 ( 134, 133, 41)
CO        2.602E-08  5.694E-06  1.000E-10 ( 154, 151, 15)
TEM1     1.727E-01  7.973E+01  3.200E+01 ( 175, 156, 68)
```

- The above table from the Result file shows:
 - the maximum absolute correction,
 - highest value in the domain,
 - lowest value in the domain,
 - (ix,iy,iz) location of max abs correction (i.e. counting cells in each coordinate).



Comparing plots from intermediate dump files

CHAM

- This is a practical and very useful way of checking convergence.
- In “Output” / ”Field dumping” click “Intermediate field dumps” and set a sweep frequency.
- The frequency might be 10 or 20 times through the run, e.g. every 250 sweeps.
- When the case has ended, run the Viewer.
- Click “Use intermediate sweep files” and select a file; plot contours/vectors on a suitable plane.
- Click F8 to produce the same plot at the next dump (or F7 the previous one).
- If the run is converged these successive plots should show no differences.
- Alternatively you can animate the succession of plots using the  button.



Why has it not converged?


CHAM

- If the above tests indicate that convergence has not been reached - or if it is obvious that the run is diverging - we need to ask 'Why?'
- There can be many reasons for poor convergence, such as:
 - Inappropriate relaxation settings
 - Perhaps the solution is not steady!
 - Poorly conditioned grid
 - Insufficient grid to resolve flow details
 - Missing boundary conditions
 - Faulty geometry files
- Here we will discuss the first two of these.



Perhaps the flow is not steady

- Sometimes failure to converge a steady solution can indicate that the flow is in fact oscillatory or unstable.
- Even if you try relaxing harder, the convergence curves might continue to look something like this:



CHAM



Perhaps the flow is not steady

CHAM


- There may just be a small region of oscillation in a largely steady flow, e.g. behind just one building in a wind-flow-over-buildings case.
- Typically we see convergence starting well, but later the spot values continue to “wobble”, and the residuals cease to diminish.
- Can prove the flow unsteady by running transient - if you have the time and budget! (See next slide.)
- But often we might adopt the pragmatic view that this is as good a “steady” solution as we are going to get.



Perhaps the flow is not steady

CHAM

- Simulating transient rather than steady may reveal that the flow is in fact oscillatory, like this:
- (oscillatory 2D flow through a heated screen)





How to get convergence - Relaxation

CHAM

- The most widely-known technique for achieving convergence is the use of **relaxation**.
- This slows down (relaxes) the changes made to the variables from sweep to sweep.
- Relaxation does NOT alter the final solution, only the way in which it is achieved.
- Two types of relaxation are available: **linear** or **false time-step**.
- The whole purpose of relaxation is to slow down changes.
- By default an “automatic” relaxation “**CONWIZ**” is provided.



Automatic relaxation: CONWIZ

- New models are automatically set up with “Automatic Convergence Control” active.
- An alternative name for this is “CONWIZ”.
- Usually the settings activated by CONWIZ will procure convergence.
- Check the reference velocity and reference length settings – e.g. the average inlet velocity and 1/10 of the domain size.



Automatic relaxation: CONWIZ

CHAM

- **However** - the CONWIZ settings are often over-conservative, so convergence may be rather slow.
- In many cases, faster convergence may be obtained if CONWIZ is switched off and the variables for the relaxations are set manually.
- How to do this will now be described.



Linear relaxation

- At each iterative sweep new values ϕ^{new} of variable ϕ are generated. Linear relaxation replaces these with

$$\phi^{\text{new}} = \phi^{\text{old}} + \alpha(\phi^{\text{new}} - \phi^{\text{old}})$$

where ϕ^{old} is the value from the previous sweep;
 α (between 0 and 1) is the relaxation coefficient.

- $\alpha = 0$ implies zero change.
- $\alpha = 1$ implies no relaxation.
- Reducing α gives harder relaxation (i.e. more damping.)



Linear relaxation

- Linear relaxation can be applied to any variable.
- Typical values for α are in the range 0.1 – 0.5.
- $\alpha = 0.5$ would imply rather light relaxation.
- $\alpha = 0.1$ would imply hard relaxation (and would require considerably more sweeps for convergence.)
- Note – we do not usually apply relaxation to the pressure (P1), unless using BFCs.



False time-step relaxation

CHAM

- “False time-step” relaxation is less intuitive, but is generally very effective in procuring convergence.
- We observe that in a transient situation, very little can change in a short enough time.
- This is controlled by the transient term
 $(\text{mass in cell}) \times (\phi^{\text{old}} - \phi^{\text{new}}) / dt$
where ϕ^{old} is the value of ϕ at the previous time step.
- The idea is to add a similar term to the steady equation,
 $(\text{mass in cell}) \times (\phi^{\text{old}} - \phi^{\text{new}}) / dt_f$
where ϕ^{old} is the value of ϕ at the previous sweep.
- dt_f is a “false” time step, with dimension seconds.
(Note - there is no “real” time step in a steady case.)
- dt_f is the user-set parameter, like α for linear relaxation.



False time-step relaxation

CHAM

- False time-step relaxation can only be applied to solved variables.
- dt_f is sometimes written “FALSDT”.
- Values for FALSDT can be in the range $0 - \infty$ seconds (rather a large range!).
- Smaller FALSDTs give harder relaxation (i.e. greater damping).
- FALSDT = 0 would imply no change (very hard relaxation).
- FALSDT = ∞ would imply no relaxation.
- $0 - \infty$ is a very big range! How do we select suitable FALSDT values?



False time-step relaxation

CHAM

- The clue is that FALSDT has the dimension of seconds; so we choose a **characteristic time scale** for the problem. For example...
- Time to flow through/past:
 - single cell < heavy relaxation
 - few cells / typical object < medium “
 - whole domain < light “
- Better to err on the low side (i.e. too heavy), then can experiment with increasing FALSDT in mid-run.
- Suggest increasing by half an order of magnitude at a time, e.g. x3, x10, x30, x100 etc.



Choice of relaxation

- LINEAR can be applied to any variable
- FALSDT only to solved ones (but not P1).

Basic recommendation:

Variable	FALSDT	LINEAR	Comments
U1, V1, W1 (velocities)	✓		Same value for all three.
P1		✓	<u>Do not relax</u> unless using BFC
TEM1, T3 (temperatures)	✓	(✓)	Use LINEAR if you have heat conduction in solids.
C1, SMOK, AGE (concentrations)	✓		
KE, EP		✓	Do not change default

- Trial-and-error, and experience, are the only real guides.



Choice of relaxation

- If buoyancy is insignificant, FALSDT for TEM1 can probably be 10x or 100x the value for the velocities.
- But if buoyancy is significant, set FALSDT for TEM1 the same as the value for the velocities.
- FALSDT for concentrations, smoke or age can probably be 10x or 100x the value for the velocities.
- Relaxation for T3 should be the same as for TEM1.



Choice of relaxation

- Slowly varying spot values may indicate excessively tight relaxation; persistently over-oscillatory spot values may mean insufficiently heavy relaxation.
- BUT don't forget that the spot values need to change rapidly in the initial stage of a run.
- Relaxation can often be lightened after the early stages of a calculation.
- There is no single “right answer” for choice of relaxation.



Choice of relaxation

- Finally, remember that:
- CONWIZ can sometimes converge cases where all attempts at manual relaxation settings have failed.

CHAM



Solver controls

(for experienced users)

CHAM

- LITER and ENDIT specify, for each variable, the number of iterations (within the slab- or whole-field-solver) for the variable, and the solution error at which the iteration should terminate automatically.
- These are set in the “Iteration control” panel.
- Typical values for LITER are 50 for P1, TEM1 and scalars; 10 for U1, V1, W1; 20 for KE, EP.
- 10^{-3} is normal for ENDIT.
- Occasionally, increasing LITER for P1 (e.g. to 200) can be beneficial when the pressure field is slow to reach its correct level (e.g. if two parts of the domain are connected by a link in which there is a large pressure drop).



Variable value limits

(for experienced users)

- VARMIN and VARMAX are used to specify the extreme permissible values of each variable.
- They are set in “Numerics” / “Limits on Variables”
- Unless you are an experienced user it is best not to change these.



Some Summary Advice

- If the relaxation settings are not right(ish), there may be trouble converging.
- If relaxation is too heavy, convergence may be slow - the monitor plots will show very slow changes in spot value and residual.
- If it is too light, the run may diverge!
- If CONWIZ is active, check the settings for reference velocity and reference length – e.g. the average inlet velocity and 1/10 of the domain size.

CHAM



Some Summary Advice

CHAM

- If CONWIZ is not active, the false time step can be estimated, e.g. a fraction of domain size divided by a typical velocity.
- False time steps much smaller than this will tend to 'lock' the solution; much bigger values may have little effect (but may give divergence!).
- Transient cases often converge more easily than steady ones, as the transient terms act like relaxation. You may not need relaxation at all.
- Reducing the time step can help with convergence.
- It is vital to get reasonable convergence on each step of a transient, otherwise subsequent steps may not be reliable.