



Industrial Evaluation of Integrated Performance Analysis and Equation Model Debugging for Equation-Based Models

Å. Kinnander¹ M. Sjölund² A. Pop²

¹*Siemens Turbo Machinery AB, Finspång, Sweden. E-mail: ake.kinnander@outlook.com*

²*Department of Computer and Information Science, Linköping University, Linköping, Sweden. E-mail: {martin.sjolund,adrian.pop}@liu.se*

Abstract

The ease of use and the high abstraction level of equation-based object-oriented (EOO) languages such as Modelica has the drawback that performance problems and modeling errors are often hard to find. To address this problem, we have earlier developed advanced performance analysis and equation model debugging support in the OpenModelica tool. The aim of the work reported in this paper is to perform an independent investigation and evaluation of this equation model performance analysis and debugging methods and tool support on industrial models.

The results turned out to be mainly positive. The integrated debugger and performance analyzer locates several kinds of errors such as division by zero, chattering, etc., and greatly facilitates finding the equations that take most of the execution time during simulation.

It remains to further evaluate the performance profiler and debugger on even larger industrial models.

Keywords: profiler, debugging, Modelica, industrial, OpenModelica

1 Introduction

The development of today's complex products requires integrated environments and equation-based object-oriented declarative (EOO) languages such as Modelica (Modelica Association, 2014; Fritzson, 2015) for modeling and simulation.

The increased ease of use, the high abstraction, and the expressivity of such languages are very attractive properties. However, the drawback of this high-level approach is that understanding the causes of unexpected behavior, slow performance, and numerical errors of certain simulation models is very difficult, in particular for users who are not experts in simulation methods.

Therefore Pop et al. (2014) have recently developed an advanced equation model debugger for the Modelica

language, as part of the OpenModelica (Open Source Modelica Consortium, 2016) tool. This is quite different from debuggers of conventional algorithmic programming language debuggers (Stallman et al., 2014; Nethercote and Seward, 2007; Zeller, 2009). Pop and Fritzson (2005) developed a debugger for the algorithmic subset of Modelica and Bunus (2004) developed a debugger that analyzes the causes of over-constrained or under-constrained systems of equations. The new debugger is also based on the recent development of the advanced bootstrapped OpenModelica compiler (Sjölund et al., 2014).

The applications used for evaluation perform simulation of combined cycle power plants. This involves the dynamics of water cycling from water to steam and back while streaming in different flow regimes through

pipes, valves and volumes, affecting the heat transfer from the flue gases. To handle these rather complicated phenomena including boiling and condensation in and on tubes, accurate dynamic models often require high computation power, efficient programming as well as a good balance between accuracy and computational speed in the aspect of simulation purposes. The performance analyzer, also called profiler, which is a tool that informs where in user equations CPU power is spent and gives thereby possibility to evaluate different mathematical methods and make deliberated trading between accuracy and computational speed. As described in Sjölund (2015), the techniques used when profiling Modelica equation-based models are quite different from profiling of general programs (Graham et al., 1983). Some earlier more limited approaches to profiling Modelica models are presented by Huhn et al. (2011) and Schulze et al. (2010).

The integrated equation model debugger has been evaluated by the designers and performs well on both small and big models. However, an independent evaluation of the integrated performance analyzer and debugger by industrial users on industrial problems was still missing. Such an evaluation is the main topic of this paper. We have earlier made a preliminary industrial evaluation only of the debugging functionality (Kinnander et al., 2016). This paper presents an evaluation of the integrated performance analysis and debugging methods and tool, including a slightly updated version of the debugging evaluation results presented by Kinnander et al. (2016).

The rest of the paper is structured as follows: first the errors to be investigated and models to be evaluated are briefly presented. Section 2 introduces the debugger tests in more detail. Section 3 presents debugging of errors in the logarithmic temperature calculation whereas Section 4 presents debugging of errors due to bad initial values. Section 5 presents the performance analyzer and its use. Finally, Section 6 presents conclusions.

1.1 Errors to be Investigated

In order to investigate different types of errors that could be expected to occur, a small and simple evaporator model is used. This has been fetched from a larger model used for transient analysis of combined power plants by Siemens Industrial Turbomachinery AB in Finspång, Sweden. The following errors are to be investigated:

1. Division by zero
2. Errors in the average logarithmic temperature difference used for heat transfer calculation:

- a) Inlet temperature difference = 0
- b) Inlet temperature difference = outlet temperature difference.

3. Boiling in the evaporator that causes halt of simulation progress by much too small time steps (stiffness)

4. Various test of bad initial values, with variation of pressure, temperatures, flows and masses in the different parts of the process.

The model selected is a simplified model of an error free model, hence the above test will be deliberately inserted and the debugging tool will only be examined by its outputs, while a sharper application for a real model development where errors are unknown and the debugger support for identifying them will be more apparent, will be carried out later. The reason for this is the limited time available for testing, and that a sharp application will only provide stochastic errors and could thereby not be planned in time.

1.2 Models for the Debugger and Performance Profiler Evaluation

The evaporator test model shown in Figure 1 will be used for the investigation, containing an instance (Evap) of the Evaporator model shown in the middle of the connection diagram.

It consists of an evaporator model that has flue gases as heating source and water as coolant, producing dry steam to the steam sink. The steam production is decided by the heat from the flue gases, the enthalpy (temperature and pressure) of the water source (FWpump), and the steam extraction to the steam sink (SteamSink) that in turn is tracking the evaporators drum pressure with a negative bias of 0.1 to 1 bar. The model has 1110 equations.

The evaporator model is designed according to Figure 2.

The evaporator model has a drum model (Drum), a heat exchanger (Hex) and a level controller (DLC) that controls level by a control valve (FW_CV). The level controller is no actual water level controller in length units (m), instead it controls the amount of water (kg) in the drum. The thermal capacity of the drum metals is represented of a heat capacity model (DrumWallHeatMass) but the insulation towards surroundings are assumed to be perfect, i.e. no heat losses to the outside of the drum. This model has 809 equations.

The drum model is the volume model from the Fluid library manipulated to only let steam exit, i.e., always perfect separation. The heat exchanger is according to Figure 3.

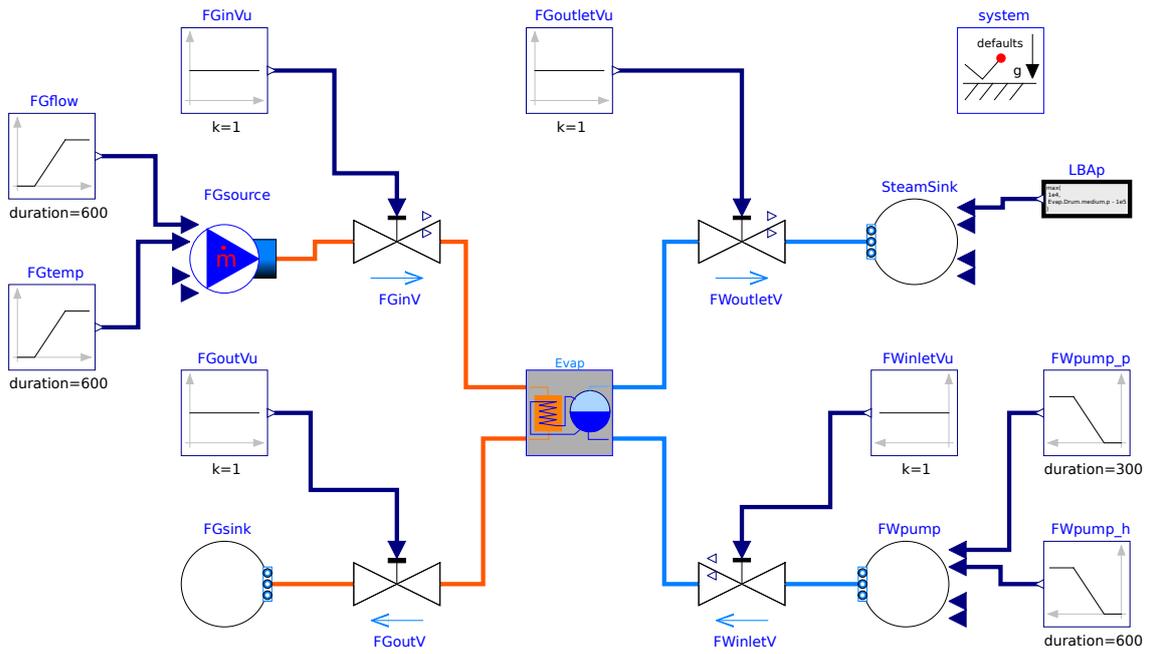


Figure 1: Evaporator test model

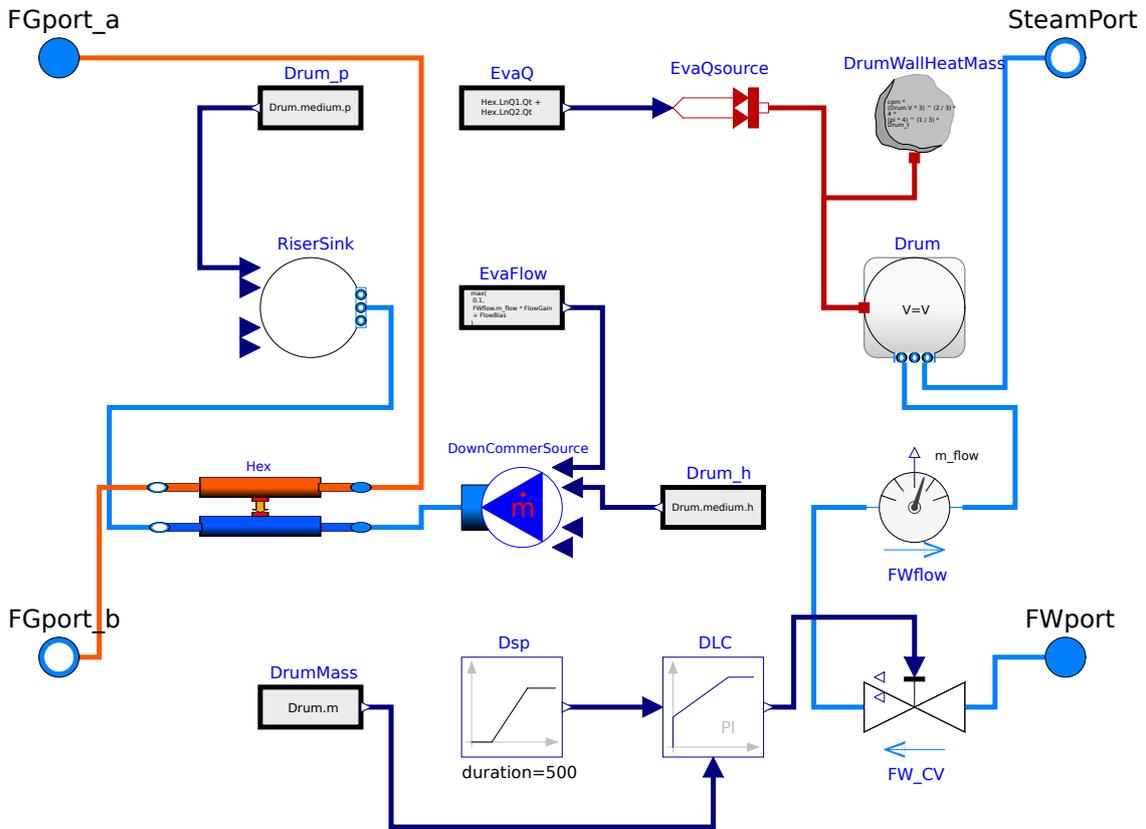


Figure 2: Evaporator model

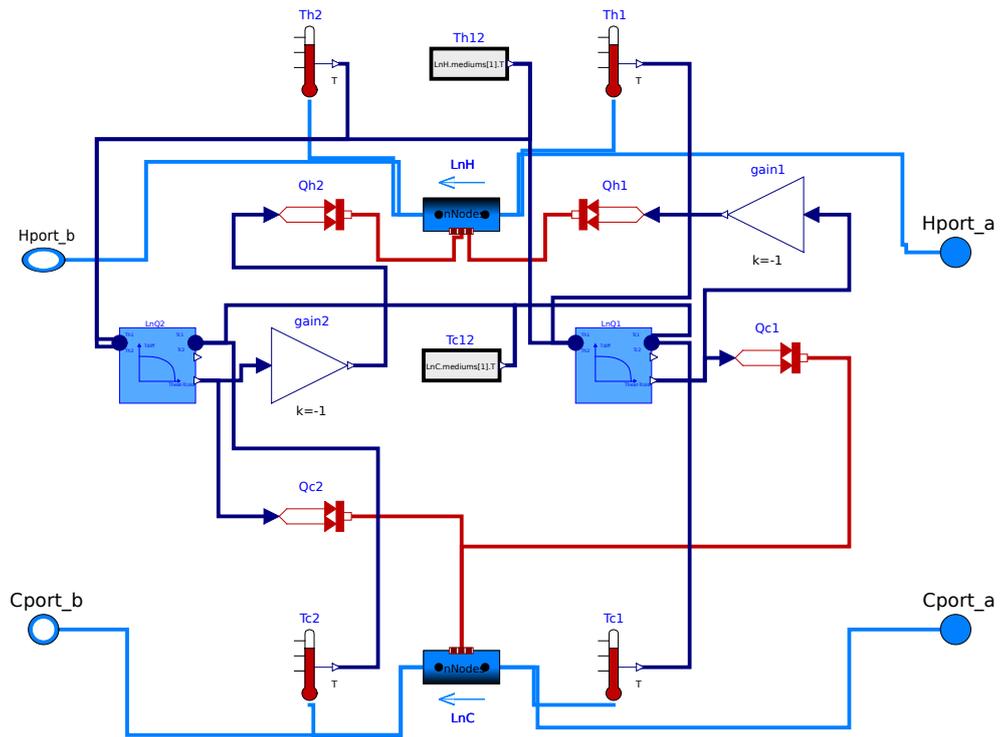


Figure 3: Heat exchanger

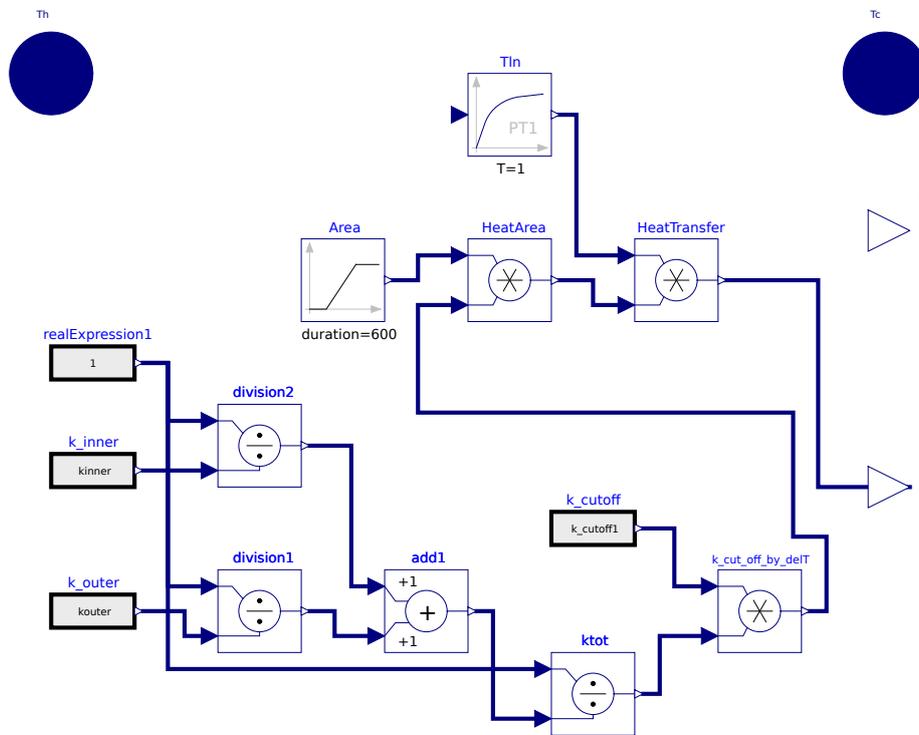


Figure 4: Heat transfer calculation model

The heat exchanger has a pipe model (LnH) that corresponds to the flue gas channel and a pipe model (LnC) that corresponds to the pipe bundle carrying the water to be heated. The heat exchanger has parallel flows. In this simplified model the overall heat transfer coefficient ($\text{J}/\text{m}^2/\text{K}$) is a constant, i.e., it takes not configuration or medium properties into account. The driving temperature difference is calculated for respectively inlet and outlet sections of the pipe bundles (i.e. they are configured with 2 nodes each) by the models LnQ1 and LnQ2. The temperatures are measured besides inlets and outlets for both pipe models also in the middle (Tc12 and Th12). This model has 580 equations.

Finally, the heat calculation models LnQ1 and LnQ2 are according to Figure 4. All other models are from the standard Modelica library.

The model has a low pass filter with an input connected in the text layer where the logarithmic temperature difference from the connectors Th and Tc which both are connecting both inlet and outlet temperatures respectively medium side, is calculated. The output of the model is the calculated heat transfer (W). The overall heat transfer coefficient k_{tot} is calculated from parameters representing the heat transfer coefficient from flue gas to metal, k_{outer} , and from metal to water, k_{inner} . In the present full size boiler model those parameters are replaced with connectors that provide more accurate values, based on medium and flow properties calculated in separate models. The heat transfer area is a ramp with selectable duration and height values, to be adopted to what is needed from initializing aspect (duration of suitable size respectively to the real heat transfer area).

This model contributes with 39 equations of the total of 1110 equations.

2 Debugger Tests

2.1 Activation of Debugger

The debugger is activated by setting the flag « Launch transformational debugger ». After a successful simulation the output windows are containing the following information (Figure 5).

The simulation output window contains assertion violation messages that are false, because the enthalpy flow H (W) has too narrow range in the Standard Modelica library. It ought to be at least 10 times as big.

This violation has no influence on the simulation result (might there be an unnecessary delay?). The window shows with a green bar that 100 % of simulation is done and the blue text that it has been successful. The transformational debugger window shows all variables in the variable browsers window and all equations in

the equation browser window, as found in the simulation code. All other frames in the debugger are empty.

2.2 Division by Zero by Parameter Setting

The test is done by setting parameter k_{inner} to zero. The simulation output window displays the following messages (Figure 6).

The simulation output window gives the required information that simulation crashed at initialization due to an assertion that avoids division by zero and this is caused by $k_{\text{inner}}=0$. The debugger window looks as before but after clicking debug more in the simulation output window it looks as in (Figure 6).

The equation browser marks initial equation with index 102: `Evap.Hex.LnQ1.add1.u1 := DIVISION(1.0, Evap.Hex.LnQ1.kinner)`, which is the same information as in simulation output window. The frame denoted `Defines` gives the variable that becomes undefined by the zero division. The frame denoted `Depends` give the variable name `Evap.Hex.LnQ1.kinner`. For this error the debugger gives all information necessary.

2.3 Division by Zero by Time Function

The k_{inner} variable is replaced by a time function that ramps it down to zero in 100 seconds. This results in a never ending simulation.

The solver manages to pass the 100 s time point where k_{inner} is zero and a division by zero occurs. No plots are available but the ramp proceeds to negative values for k_{inner} . The solver has skipped the exact 100 s time point, but then continued into other problems, due to the negative k_{inner} value. On the passage it has however produced two messages about zero division at time 100 when they occurred. In the case of the user being unaware of the division problem, the large amount of output in the simulation output window hides those messages.

For a ramped denominator passing zero, the debugger is not optimal in case the solver manages to pass the critical point and that consecutive errors then hide the information from the user. A solution would be the option to let the user decide if division by zero should be accepted or not, i.e., the solver should then interrupt and save when any denominator having a passage of zero.

2.4 Division by Zero due to Mismatch in Parameter Settings

By deselection of the heat transfer used in the LnC pipe in the Hex model (Figure 3) the test model still checks OK, but now with only 1106 equations instead of 1110.

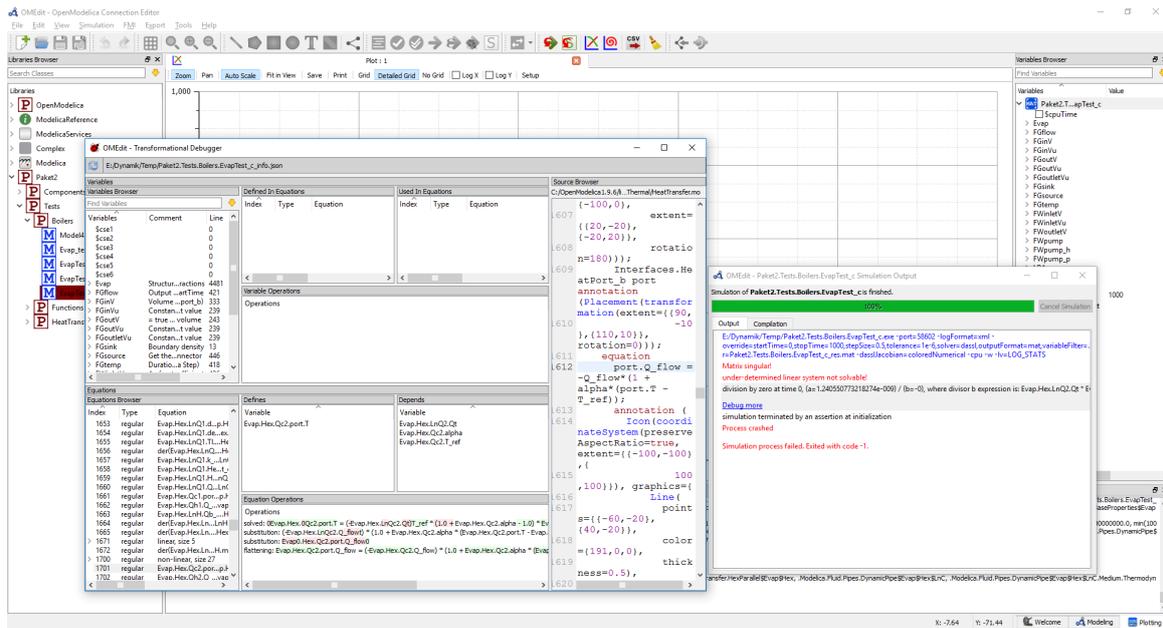


Figure 7: Outputs after no use of heat transfer specified but corresponding heat transfer connectors anyway connected

Simulation gives the following simulation output window and transformational debugger window (Figure 7).

The simulation crashes at initialization due to division by zero where the denominator is involving the variables Q_t and α . A check of the model reveals that α is the heat transfer coefficient to surroundings and is deliberately zero. That is, this model is impossible to run although it checks OK. The debugger points to the equation numbered 1258. Marking that equation points to a source code line 1612 where the Modelica standard library (MSL) component `PrescribedHeatFlow` (PHF) equals the heat ports heat flow (Q_Flow) to the connector input Q_Flow (the prescribed flow). This equation contains also a dependence on α , but with $\alpha=0$ the equation should be OK: $port.Q_Flow=-Q_Flow$. The PHF model calculates its internal temperature in order to be able to have the prescribed heat flow by this equation (index 1258). This temperature should be the same as the connected pipes temperature as there is no thermal resistance in the connection itself so why is it calculated in this way? One could not be sure that any user would understand that this equation is unexpected and caused by a wrong parameter setting in the pipe model.

It would have been better if there would be a consistency check between connectors and their use in the model. One way would be not to show the connector unless it is activated, or give an error message if it is anyway connected. Could the variables browser shed some light over the problem? The variables browser is

presented in (Figure 8) after the port temperature has been clicked.

It gives the same information augmented with the initial equations. A question is why the debugger is not pointing to the initial equations indices 276 and 771, as the simulation output window tells that the error appeared at initialization.

3 Errors in the Logarithmic Temperature Difference Calculation

Errors in the logarithmic temperature difference calculation should be treated the same way as the division by zero. Interesting is however, if the solver also for this type of errors manage to pass the critical point as their time duration could be expected to be very short. Basically this investigation is more an investigation of the solver and not the debugger but the debugger will be activated and therefore also this investigation is a part of the paper.

3.1 Temperature Differences Passing Zero

This test is achieved by removing the numerical fences that prevent zero crossing. Unfortunately, it turns out that there are no crossings that passes $\Delta T=0$ for the case simulated, and the test needs further work to be

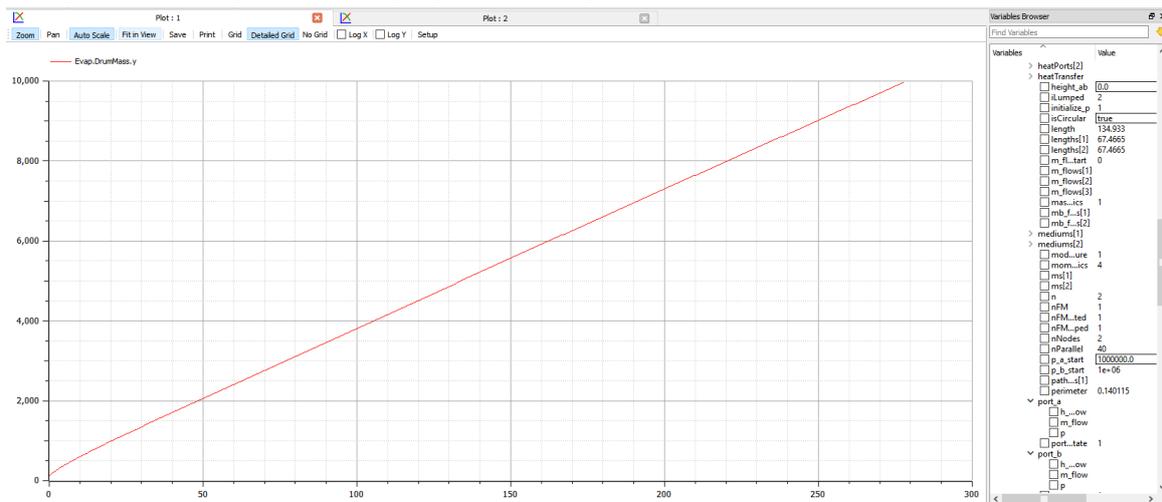


Figure 9: Plot of drum mass at blocked drum outlet

There is one call to the clock before executing the equation block or function call and one call to the clock after execution of the block. Associated with each call is a counter that keeps track on how many times this function was triggered for the given time step. Similarly, each call is associated with clock data one variable for the total time spent in the block for all time steps, and one variable for the total time spent in the block for the current time step. These calls to the clock are in the code generation as a macro set that is generated if profiling is desired this means zero overhead unless profiling is explicitly enabled.

Profiling can be enabled for all equations and function calls. With profiling enabled only for equation blocks (SCCs) and functions, the overhead cost is low compared to the cost of solving most nonlinear systems of equations, which is more suitable for real-time simulation.

The instrumentation is performed by compiling a model with additional code generated to query real-time clocks at appropriate places.

The integrated performance profiler functionality is accessible from the GUI of the transformational debugger (Figure 10).

In the foreground is the window Transformational Debugger that invoked by the selection of « Profiler All » in the simulation setup options for simulation flags. Behind that the « Simulation Output » window, giving the important information that 100% of specified simulation time is successfully reached, but also various warnings in red with an option to debug more. (The output window is better always placed in forefront, to see the progress of the simulation, or alternatively the plot window if it plots the progress continuously or at least frequently). In the background the

OpenModelica OMEdit main window is visible, now displaying the plotting interface.

A closer look at the Transformational Debugger and Profiling window (Figure 11) gives the following:

To the left the integrated transformational debugger and performance profiler displays two browser windows, one for variables and one for equations. Clicking on a variable in the variable browser window (middle left) will show the following:

- where it is defined
- where it is used
- operations made
- the source code line that defines the variable value will be highlighted (source browser to the right)

The contents of displayed lines are truncated to allow all windows above to fit into the same screen. Expanding/resizing windows and scrolling can display all hidden text. Thus, knowledge about where a variable is used becomes instantly available which substantially helps to analyze the consequences of a model change.

The index presented provides the link to the equations which are used in the model and presented in the equation browser. Clicking on a line in the equation browser window (lower left) gives:

- which variable the equation defines
- what variables it depends upon
- what operations that are made for the equation
- where in the source code the equation is placed

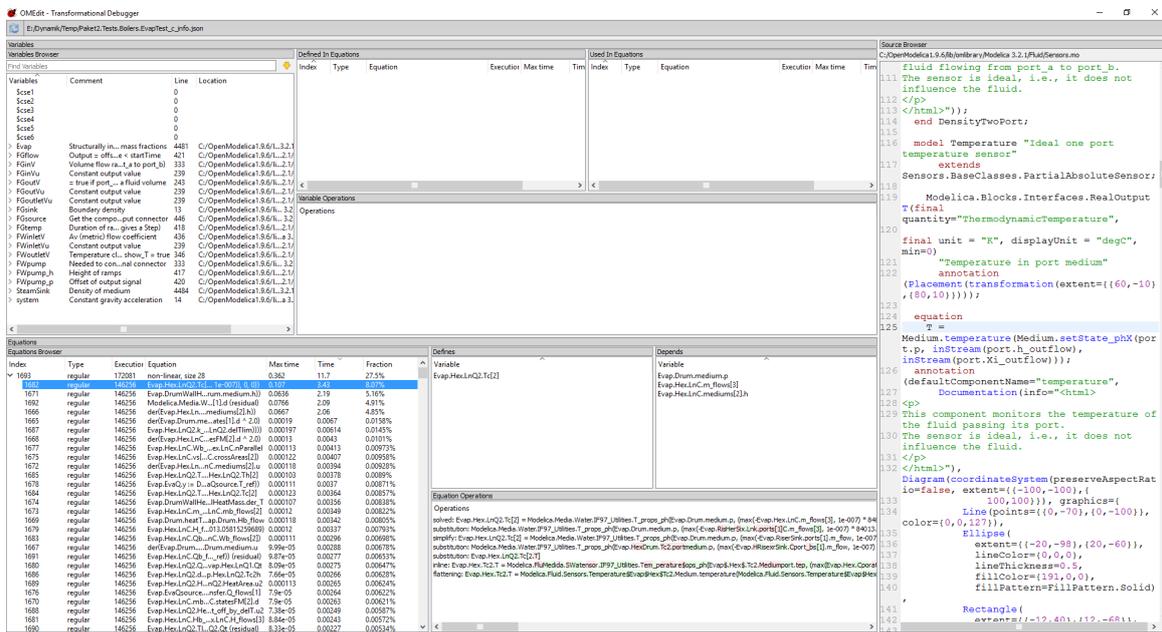


Figure 12: Using the performance profiler to find the computationally heaviest equation after sorting according to used time (table lower left), see close-up in Figure 13.

Index	Type	Execution	Equation	Max time	Time	Fraction
1693	regular	172081	non-linear, size 28	0.362	11.7	27.5%
1682	regular	146256	Evap.Hex.LnQ2.Tc[... 1e-007)], 0, 0)	0.107	3.43	8.07%
1671	regular	146256	Evap.DrumWallH...um.medium.h)	0.0636	2.19	5.16%
1692	regular	146256	Modelica.Media.W...[1].d (residual)	0.0766	2.09	4.91%
1666	regular	146256	der(Evap.Hex.Ln....mediums[2].h)	0.0667	2.06	4.85%
1665	regular	146256	der(Evap.Drum.me...ates[1].d ^ 2.0)	0.00019	0.0067	0.0158%
1687	regular	146256	Evap.Hex.LnQ2.k...LnQ2.delTlim))))	0.000197	0.00614	0.0145%
1668	regular	146256	der(Evap.Hex.LnC...esFM[2].d ^ 2.0)	0.00013	0.0043	0.0101%
1677	regular	146256	Evap.Hex.LnC.Wb...ex.LnC.nParallel	0.000113	0.00413	0.00973%
1675	regular	146256	Evap.Hex.LnC.vs[...C.crossAreas[2].u	0.000122	0.00407	0.00958%
1672	regular	146256	der(Evap.Hex.Ln...nC.mediums[2].u	0.000118	0.00394	0.00928%
1685	regular	146256	Evap.Hex.LnQ2.T...Hex.LnQ2.Th[2]	0.000103	0.00378	0.0089%
1678	regular	146256	Evap.EvaQ.y := D...aQsource.T.ref)	0.000111	0.0037	0.00871%
1684	regular	146256	Evap.Hex.LnQ2.T...Hex.LnQ2.Tc[2]	0.000123	0.00364	0.00857%
1674	regular	146256	Evap.DrumWallHe...HeatMass.der_T	0.000107	0.00356	0.00838%
1673	regular	146256	Evap.Hex.LnC.m...LnC.mb_flows[2]	0.00012	0.00349	0.00822%
1669	regular	146256	Evap.Drum.heatT...ap.Drum.Hb_flow	0.000118	0.00342	0.00805%
1679	regular	146256	Evap.Hex.LnC.H_f...013.05815259689)	0.00012	0.00337	0.00793%

Figure 13: Transformational debugger and profiling window

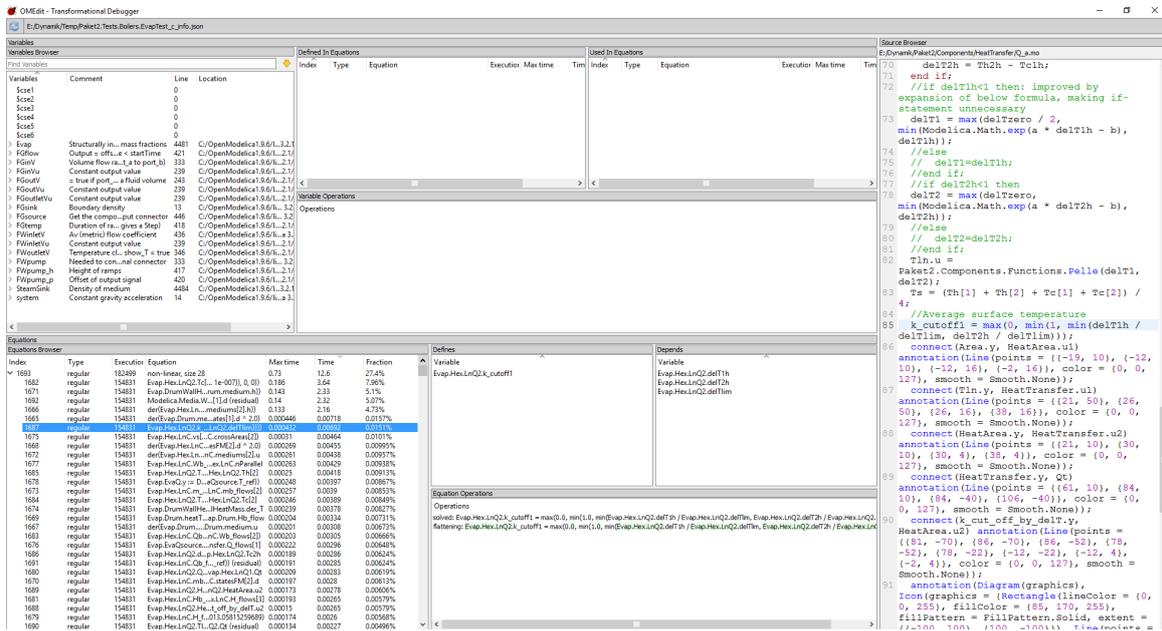


Figure 14: Performance profiler output after a minor model adjustment. See also close-up in the Figure 15.

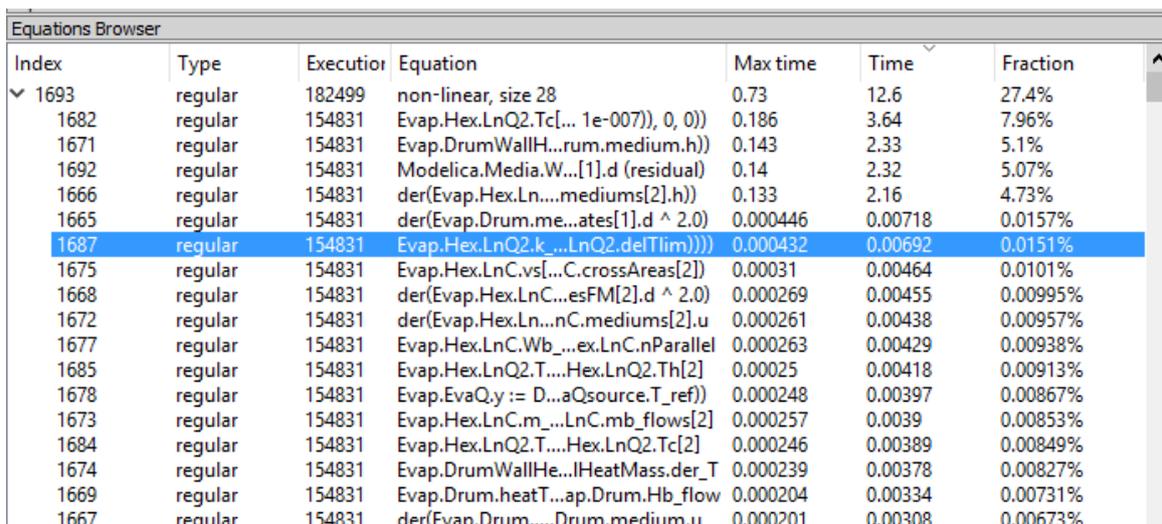


Figure 15: Close-up of performance profiler output after a minor model adjustment. The 28 equations at Index 1693 use 27.4% of the time.

The equation browser gives the link between the C-code equation solved and the source code in the Modelica source files. This browser also displays columns for execution times, number of executions and, total execution time, both as fractions of total simulation time and in seconds. Sorting by the fraction of execution time used with the largest value at the top of the table presents the equations in order according to the time consumed.

As shown in Figure 12 and Figure 13 the profiler displays that Index 1693, a strongly connected equation subsystem that defines the derivative of the enthalpy and pressure of the drum together with the derivative of the enthalpy of the outlet volume of the evaporator pipes (`LnC[2]`), is by far the equation that takes most of the CPU power, 27.5 % while next equation takes 5 % (not shown in figure above). As there are three variables defined by this equation no equation is pointed out in the Modelica file (shown in Source Browser is a previous clicked index—a clearing of that window should be considered when selected variables and equations not corresponds to a line in the source code) nor are any dependencies shown.

Index 1693 constitutes 28 other variables (size 28) and clicking on 1693 reveals indices from 1665 to 1692 (28 equations). Those are the used equations by the solver and clicking on any of them shows the corresponding source code, and what variables the equation is dependent upon and the operations made.

The conclusion from this is that the Drum is the computationally heaviest part of the model, and improvements of its equations should have the best chance of improving performance and reducing the simulation time.

Exploring the 28 equations (Figure 13) reveals only one that is provided by the user (amongst the equations using 85 % of the calculation time). The rest are equations from the Modelica Fluid library. To see the impact by that equation on the performance, one of its parameters was changed. The equation has a parameter `delTlim` that is limiting the heat transfer calculation preventing that the temperature differences become equal, thereby causing a simulation crash. Increasing `delTlim` value from 0.1 to 0.5 °C, which deteriorates the accuracy of the heat transfer calculation, gives the result that equations with index 1693 makes a minor reduction from 27.5 % to 27.4 % of the total time (Figure 14 and Figure 15), but the simulation time (major part of the total time) is anyway reduced from 30.0 to 28.5 s (not shown).

From that change the user could conclude that changing `delTlim`, which rather drastically deteriorates the accuracy of the heat transfer, the resulting improvement on the performance for index 1693 is neg-

ligible, but the change any way influences the total time a bit more substantially.

In general our experience is that the performance analyzer/profiler is a very important tool for model performance optimization since it is very easy to see the link between a slow execution and the equations used. This information is very helpful when changing the model in order to speed up its simulation.

6 Conclusions

A basic property of the debugger is to assist in case of numerical problems and violation of assertions like numerical ranges that a certain variable is expected never to exceed. Then the debugger points out the equation causing the problem. In the work reported in this paper only small to medium-sized (1000-equation size) industrial models have been tested, which demonstrates that the debugger and performance analyzer work well and give significant assistance to the user. The debugger has also been briefly evaluated by Pop et al. (2014) on 11116-equation size models in the Modelica Standard Library such as V6Engine. To really evaluate the benefits of the debugger, but also its functionality, it should also be applied to larger industrial models.

The following conclusions were made from the tests with the transformational debugger and performance analyzer for equation models:

- The debugger works well to find zero denominators that are parameters.
- The debugger does not come into play automatically if the zero denominator is only a momentarily value, as the solver managed work around such time points in so far tested simulations. However, it catches the problem in the simulation output window, and gives a message that by clicking opens the transformational debugger window which displays the concerned equation. However, there is a risk that this is unnoticed as the solver continues and could generate a lot of consequential or other messages that could hide the zero denominator messages. It would be preferable if the simulation output window could aggregate messages of the same type into one, expandable, line, thereby giving a better overview of all the types of messages the simulation has generated.
- A zero denominator caused by structural model errors, like connection to not used connectors (this should not pass the model checking) the debugger points to the causing equation. One could not ask more of the transformational debugger, but

the OpenModelica model check or model building could be made to prevent such mistakes.

- In case of numerical problems causing long execution times the debugger points to the equations that have problems, but to understand the exact problem, plots of variables could be necessary—hence the result file should always be generated, regardless if the simulation is interrupted by solver or manually. This is not the case in the tested version for all the tests.
- The performance analyzer/profiler is a very important tool for model optimization as it is possible to easily see the link between a slow execution and the equations used. Compared to the alternative method where only the CPU curve together with plots of all other variables are available for guidance in the process of finding a good spot in the model to improve, the profiler makes the process of performance optimization of models radically shorter.

Acknowledgments

This work was partially supported by the ITEA2 MODRIO project and the ITEA3 OPENCPS project via the Swedish Government (Vinnova), by the RTISIM project funded by Vinnova, and by Siemens Turbo Machinery AB.

References

- Bunus, P. *Debugging techniques for Equation-Based languages*. Doctoral thesis No 873, Linköping University, Department of Computer and Information Science, 2004. URL <http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-35555>.
- Fritzson, P. *Principles of Object-Oriented Modeling and Simulation with Modelica 3.3: A Cyber-Physical Approach*. Wiley-IEEE Press, 2 edition, 2015.
- Graham, S., Kessler, P., and McKusick, M. An execution profiler for modular programs. *Software: Practice and Experience*, 1983. 13(8):671–685. doi:10.1002/spe.4380130803.
- Huhn, M., Sjölund, M., Chen, W., Schulze, C., and Fritzson, P. Tool support for Modelica real-time models. In C. Clauß, editor, *Proceedings of the 8th International Modelica Conference*. Linköping University Electronic Press, 2011. doi:10.3384/ecp11063537.
- Kinnander, Å., Sjölund, M., and Pop, A. Industrial evaluation of an efficient equation model debugger in OpenModelica. In *Proceedings of 9th EUROSIM Congress on Modelling and Simulation*. 2016.
- Modelica Association. Modelica: A unified object-oriented language for physical systems modeling, language specification version 3.3 revision 1. 2014. URL <http://www.modelica.org/>.
- Nethercote, N. and Seward, J. Valgrind: a framework for heavyweight dynamic binary instrumentation. In *Proceedings of the 2007 ACM SIGPLAN conference on Programming language design and implementation*, PLDI '07. pages 89–100, 2007. doi:10.1145/1250734.1250746.
- Open Source Modelica Consortium. Openmodelica. 2016. URL <https://openmodelica.org/>.
- Pop, A. and Fritzson, P. A portable debugger for algorithmic Modelica code. In G. Schmitz, editor, *Proceedings of the 4th International Modelica Conference*. 2005.
- Pop, A., Sjölund, M., Ashgar, A., Fritzson, P., and Casella, F. Integrated Debugging of Modelica Models. *Modeling, Identification and Control*, 2014. 35(2):93–107. doi:10.4173/mic.2014.2.3.
- Schulze, C., Huhn, M., and Schüler, M. Profiling of Modelica real-time models. In P. Fritzson, E. Lee, F. Cellier, and D. Broman, editors, *Proceedings of the 3rd International Workshop on Equation-Based Object-Oriented Modeling Languages and Tools*. Linköping University Electronic Press, pages 23–32, 2010. URL <http://www.ep.liu.se/ecp/047/>.
- Sjölund, M. *Tools and Methods for Analysis, Debugging, and Performance Improvement of Equation-Based Models*. Doctoral thesis No 1664, Linköping University, Department of Computer and Information Science, 2015. doi:10.3384/diss.diva-116346.
- Sjölund, M., Fritzson, P., and Pop, A. Bootstrapping a Compiler for an Equation-Based Object-Oriented Language. *Modeling, Identification and Control*, 2014. 35(1):1–19. doi:10.4173/mic.2014.1.1.
- Stallman, R., Pesch, R., Shebs, S., et al. *Debugging with GDB*. Free Software Foundation, 2014. URL <http://www.gnu.org/software/gdb/documentation/>.
- Zeller, A. *Why Programs Fail: A Guide to Systematic Debugging*. Morgan Kaufmann Publishers Inc., 2nd edition, 2009.