Steam Utility Systems are not 'Business as Usual' for Chemical Process Simulators

J. Currie * **D.I. Wilson** ** **N. Depree** *** **B.R. Young** ****
S. Zainal Abidin † **L. Karim** †

* *Electrical & Electronic Engineering,*
AUT University, Auckland, New Zealand
(e-mail: jonathan.currie@aut.ac.nz*).*
** *Industrial Information & Control Centre,*
AUT University, Auckland, New Zealand
*** *Light Metals Research Centre,*
University of Auckland, Auckland, New Zealand
**** *Industrial Information & Control Centre,*
University of Auckland, Auckland, New Zealand
† *Group Technology Solutions,*
PETRONAS, Kuala Lumpur, Malaysia

Abstract: The operating costs of a typical refinery utility system are significant such that small operational improvements can result in large annual savings. This paper introduces a new utility modelling package designed for PETRONAS, to allow in house modelling and optimization of their refinery and processing plant utility systems. Using this package, engineers can construct a rigorous utility system model which can calculate the actual cost of their utility system, and how operational changes can reduce the bottom line. This paper also contrasts the fundamental differences in approach between utility systems and a steady-state chemical process simulator and provides methods to resolve these differences.

Keywords: Utility, Steam, Modelling, Simulation

1. INTRODUCTION

Industrial utility systems such as steam or electricity form an integral part of most processing plants, supplying the process demands for heat and power, Varbanov et al. (2004a). Utility energy is the largest managed operating cost for the hydrocarbon process industry according to Fernandez-Polanco and Richard (2004), yet Eastwood (2002) argues from both a control and operation point of view, it typically takes a back seat to production.

Within the utility system, cogeneration is a typical modern feature, where steam turbines and/or gas turbines can be used to provide both mechanical and electrical power, either for the plant itself, or if economically viable, to sell energy back to the national grid. However the cost of fuel to generate steam or run a gas turbine can be significant, with large plants costing hundreds of millions of dollars per year to run, as is purchasing external power from the local grid, Varbanov et al. (2004b).

Searching for the most efficient way to operate a utility system given varying electricity and fuel prices, changes in process demands, and equipment availability is a complex problem, Fernandez-Polanco et al. (2004), Eastwood and Bealing (2004), where the optimal selection of running equipment is not always directly obvious. This provides an excellent opportunity for a rigorous utility system model to be built, which can not only provide run time information, but also forms the basis for optimization.

Utility system modelling and optimization has been explored by multiple authors and methods, with perhaps the most well known being Pinch, Linnhoff and Hindmarsh (1983). Using graphical curves, networks of heat exchangers (HENs) can be optimized based on the'pinching' of two composite curves, created from the hot and cold sides of the network. An alternative method is the R-curve proposed by Kenney (1984), which was further explored by Kimura and Zhu (2000) building on the co-generation targeting methodology developed by Mavromatis and Kokossis (1998). This method focused on the redistribution of steam flows through the system's steam turbines to maximise the steam boiler fuel efficiency.

Introducing constraints on maximum or minimum flows, site-wide power requirements or a refinery fuel balance and the methods described so far will fail to establish the optimum, nor will they provide a simulation framework suitable as a commercial platform, Fien (2008), Varbanov et al. (2004a). In order to identify practical and therefore realistic cost savings, the model must take into account system constraints and variable equipment efficiencies, which requires a dedicated utility modelling package complete with rigorous thermodynamics and a flexible system structure.

PETRONAS, the Malaysian state owned oil and gas company, recognised the economical and operational advantages of building and maintaining utility models of their refineries and processing plants, and commissioned the Industrial Information and Control Centre (I^2C^2) to develop the steam utility package described within this paper. While several other commercial utility modelling packages exist such as ProSteam from KBC, Ariane from ProSim, or Aspen Utilities Operations, PETRONAS decided to use iCON as the simulation platform. iCON is a steady state and dynamic process simulator finely tuned to PETRONAS needs and has been in use since 2003. iCON is based on Virtual Materials Group's VMGSim. The package therefore was to be an 'add on' and thus the core design was the development of utility unit operations and methodology of solving a utility system, rather than steady state solvers, graphical interfaces and so forth which already existed.

This paper details the development of iCON Utility Optimizer (iUO); a utility modelling add on package designed for iCON, and describes some of the issues and difficulties that arose during development.

2. ICON UTILITY OPTIMIZER

As expected from most simulation packages, iUO was designed to perform automatic mass and energy balances, provide a live graphical PFD interface with drag and drop unit operations, as well as a rigorous thermodynamic model for combustion calculations and state properties. What was unique about the iUO package was that by building on an existing chemical process simulator, which was already heavily used by process engineers, installed knowledge of modelling within iCON was retained by these engineers and the learning curve could be substantially shortened. Another advantage of this arrangement was also the ability to model both process and utility systems within the one package, enabling more accurate economic calculations if required.

The PFD of the resulting package is exemplified in Figure 1, which shows a three header closed-loop steam system. The system is built using the custom 'Steam' unit operations developed for the package, and connected using the familiar drag unit operation and drop into flowsheet procedure familiar (and expected by) most process engineers today when using process simulators.

2.1 Mass and Energy Balance

The base iCON distribution is a rigorous process simulation package such that mass and energy balances within a unit operation are automatically closed. If the user supplies specifications which are inconsistent with calculated results, these will be flagged by the software as errors so the user can quickly identify the problem. To further aid in this all flows are represented by colour coded streams, indicating vapour (cyan), liquid (blue), two-phase (pink) and zero-flow (grey). This is a substantial advantage over a spreadsheet based modelling approach which demands considerable user discipline in order to be assured that mass balances are closed and errors are not introduced inadvertently.

2.2 Steam Unit Operations

Steam specific unit operations are built by assembling collections of fundamental process simulation unit operations such as heaters, mixers, balances, and expanders. Using basic process insight, simple models can be built, such as the Simple Boiler in Figure 2, which is two heaters, one splitter, one energy balance and two fraction blocks. By building unit operations in this way we ensure that the entire unit operation's mass and energy balance is complete, as all internal operations must meet the rules of a normal process model.

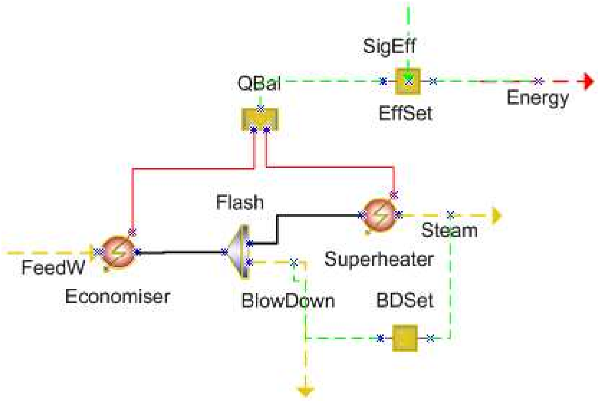

Fig. 2. An exploded view of the Simple Boiler.

A second advantage of this method is the automatic degree of freedom assessment by the iCON solver engine. For example a steam turbine can be specified by flow or by duty, but the same unit operation can fulfill both roles, because the solver will automatically calculate the remaining unknown based on the user's specification.

2.3 Combustion Unit Operations

Combustion unit operations inside the boilers, furnaces and turbines used VMG's Advanced Peng-Robinson (APR) thermodynamic model and a reaction furnace to automatically model the combustion process of the fuel and air mixture specified by the user. These two components automatically complete the combustion stoichiometry and reaction kinetics, enabling accurate estimation of combustion emissions and available heat duty of a given fuel air mix. These were used in the Advanced Boiler, Fired Furnace, Gas Turbine, and other unit operations.

Using the advanced reaction kinetics, the Advanced Gas Turbine predicts NO_x production based on the current operating conditions, and the user can trial varying levels of steam injection to view the $DeNO_x$ effect. Estimation of SO_3 generation for sulfurous fuels is also modelled, such that the acid dew point for a furnace flue gas could be estimated.

Unit operations were built using the same methodology as steam only unit operations, but used split thermodynamic models to model the fuel and steam sections separately, enabling more accurate calculations of state properties.

2.4 Parallel Unit Operations

A common unique feature of utility models is the large number of parallel unit operations, such as steam turbines which operate at the same conditions, but are separated

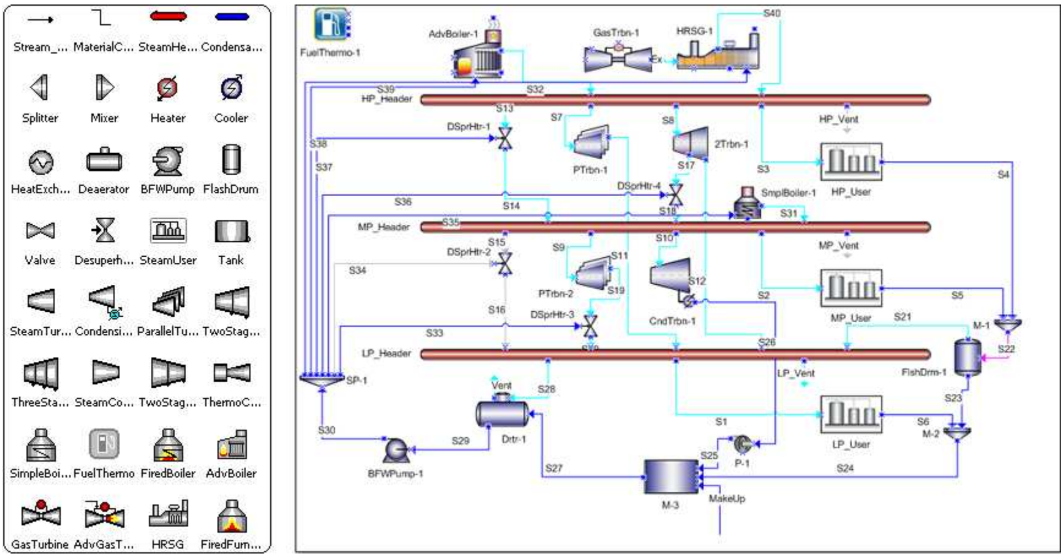

Fig. 1. An example of a steam system comprising of multi-stage turbines, gas turbines, heat recovery steam generators and the advanced boiler. Also shown is the 'Steam' palette for dropping unit operations onto the PFD.

due to sheer size, redundancy, process requirements, or installation requirements. Although a steam turbine may be diagrammatically represented on the PFD as one turbine, there may be in fact ten, twenty or more steam turbines of varying sizes and efficiencies, connected to the same steam temperature and pressure, and all connected into the same header below.

This provided the iUO package with the ability to lump all of these parallel unit operations as one steam turbine, appropriately sized and specified such as to accurately reflect the outlet enthalpy and power of all turbines. The advantage gained here is computational efficiency, because within iCON the steam turbine unit operation required an iterative internal solver to converge the unit. If solved individually this would dominate the computation time when compared to the remainder of the model's operations.

Our solution to the construction of these large groups of steam turbines was to provide an Excel based calculation tool which could connect live data back to iCON, using the iCON ActiveX connection. A user could enter the required number of steam turbines, together with their operating specifications into a preformatted table, collect live temperature and pressure values from iCON, and then calculate an equivalent parallel turbine within Excel. The results could then be automatically exported back to iCON such that in effect, the iCON model only ever solved one turbine for that group, resulting in substantial computational efficiency gains. Figure 3 illustrates a completed steam turbine table.

A second benefit to this arrangement was the ability to control the re-solving of the parallel turbine. A number of case studies run on the models may not affect the incoming enthalpy to this parallel group, such that there was no need to re-solve all the steam turbines at every iteration. When the header temperature or pressure did change, or one of the individual turbines was switched on/off, then the user could resolve the parallel turbine as required.

HP - MP Process Drivers

Equivalent Turbine:	
Power	3600.00 kW
Efficiency	70%
Mass Loss Fraction	0.000

0 Slow Roll = 0.3 Hot = 0.1	Turbine Base Power	Isentropic Efficiency	Steam Flow To Match	Steam Flow In	Steam Flow Out	Turbine Shaftwork	Power Import
1.0	1200.00 kW	70%		16.87 ton/h	16.87 ton/h	1200.00 kW	
1.0	1200.00 kW	70%		16.87 ton/h	16.87 ton/h	1200.00 kW	
1.0	1200.00 kW	70%		16.87 ton/h	16.87 ton/h	1200.00 kW	
			Totals:	50.61 ton/h	50.61 ton/h	3600.00 kW	0.00 kW

Fig. 3. The iUO Excel add in showing the parallel turbine calculation table.

This functionality was built into iUO for steam boilers, steam turbines, steam users (process), furnaces and gas turbines.

3. INDUSTRIAL APPLICATION

iCON Utility Optimizer was benchmarked against two major utility systems: a Methyl Tertiary Butyl Ether (MTBE) refinery and a Liquified Natural Gas (LNG) plant.

Twenty case studies designed to reduce operating costs were run across each model in order to validate the model results. Cases included: reducing the deaerator pressure; optimizing the selection of running steam turbines against electric motors, and installation of Steam Turbo Generators (STGs) to reduce power import requirements. For each case a number of key variables were checked across the model, which included: total running cost; total steam production; header temperatures and let down flows.

As these were hypothetical case studies the values were not compared to actual operating measurements, but were compared with values which were currently being used for a Strategic Energy Review (SER). Typical savings of 1–2% were identified and the results matched easily within predetermined tolerances for the iUO package.

4. UNIQUE CHARACTERISTICS OF UTILITY SYSTEM MODELS

From the outset it would seem that building on a process simulator would be relatively straight forward; the user interface already exists, the thermodynamic models exist or were considered simple, fundamental unit operations already exist, among other benefits. However it quickly became evident that steam utility systems exhibit some unique characteristics that distinguish them from a typical chemical process model.

4.1 Multi-Point Modelling

Modelling for optimization, as was intended with the utility models, requires a new paradigm for both process and utility system models. The model environment must now be able to solve for another operating point, no matter what inputs are changed. For a utilities model, input specifications can be a wide variety of things such as steam temperatures, steam pressures, operating equipment, as well as boiler fuel selection. If any one of these inputs are changed, the resulting solution should be of acceptable accuracy, especially if compared with the actual operating plant. We call this strategy *multi-point* process modelling.

In a multi-point model, all unknowns must be solved by the equations of the model, so that internal specifications are either directly calculated by unit operations or manually entered equations. It is typical that not all internal specifications can be analytically solved, so that a numerical iteration strategy is commonly required. While this adds extra computation time to the model, the model can now represent the plant across a wider range of operating conditions and therefore is now suitable for optimization.

4.2 Single Component Systems

As most process models in the hydrocarbon industry contain multiple components, it was noted that iCON would not automatically infer composition data across a mixer unit operation until mass or molar flows were also entered. This is a default arrangement because calculated compositions (mass or mole fractions) are dependent on flows through the mixer.

However in the case of single component systems such as H_2O in steam systems, once all input stream compositions have been entered, (as 100%) it is logical to assume that the mixer will now calculate that the only option for the output will at that same composition. As shown in Fig. 4, the mixer and input streams have failed to solve even though the unit is fully specified. As the mixer forms an integral part of a utility model (within headers, desuperheaters and deaerators), this problems requires each one these unit operations to have the output composition specified automatically by the package.

4.3 Water and Steam Thermodynamic Model

The thermodynamic model for steam within iCON is an implementation of Steam 1995, and is a modification of the 1995 steam tables available in NIST's Refprop system developed by Virtual Materials Group. This property package is based on a complex expression of the Helmholtz energy and therefore uses temperature and volume as natural variables instead of pressure and temperature and although very accurate it is not ideally suited for process calculations.

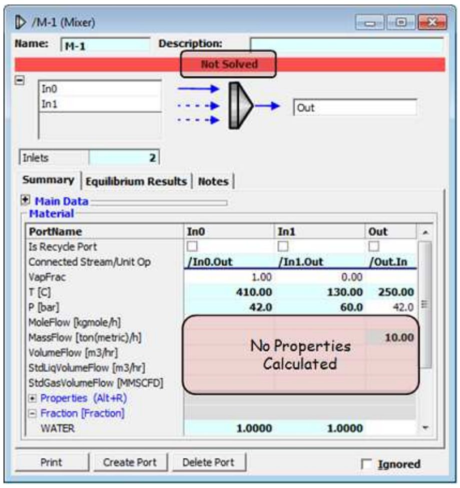

Fig. 4. Example indicating compositions are not automatically inferred on output streams.

Currently under development for the package are the IAPWS 1997 correlations IAPWS (2007), which include region based polynomial regressions with temperature and pressure pairings. Initial testing has shown these to be up to 14 times faster than the original Steam 1995 implementation.

4.4 Zero Flows

A unique characteristic of steam utility systems is the common situation where the flow of a stream drops to zero. Examples of zero-flow situations include by-pass valves connected to large steam turbines, the substitution of an electric motor instead of a steam turbine, the closing of vents or let-downs when the steam demand changes, or when equipment is shut-down for service.

The difficulty with a zero flows is the subsequent calculation of the thermodynamic properties of the stream, and what can be reliably inferred from the stream. As the flow is zero, the energy of the stream is also zero, which now infers that an energy balance is meaningless. Take for example a mixer with two input ports, and one output port, which is similar in form to a utility system desuperheater.

$$H_{\text{out}} = \frac{H_1 M_1 + H_2 M_2}{M_{\text{out}}}$$

Given arbitrary input pressures and temperatures, and specified outlet flow of zero, $M_{\text{out}} = 0$ (such as in the case where the steam header below requires no steam), the outlet temperature can now be entered as any value within the thermodynamic model limits, and the solution is mathematically valid. If the unit operation were to solve for the outlet enthalpy, the result $H_{\text{out}} = \infty$, causes a thermodynamic exception.

The result of this equation within iCON is shown in Fig. 5 where the mixer appears solved, however both input streams have failed to solve.

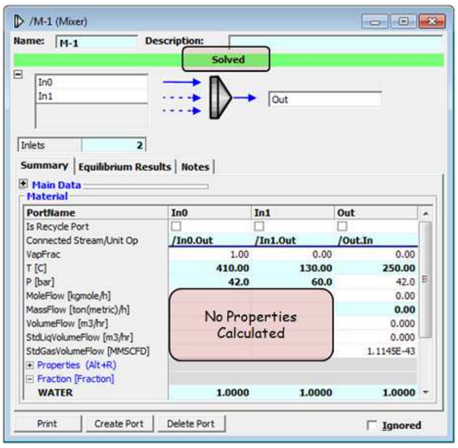

Fig. 5. Example showing zero flow solving issues.

In addition to mixers suffering from zero flow problems, other typical unit operations such as compressors and expanders were also noted to fail to solve with a zero flow. This meant that when a steam turbine was turned off,

$$\text{Power} = M_{\text{out}} \left(H_{\text{in}} - H_{\text{out}} \right), \text{ with } M_{\text{out}} = 0$$

the model would no longer solve as ΔH is underfined. Because the output stream from the expander would not have solved, no information can be propagated subsequently through the flowsheet.

In order to overcome this hurdle our solution was to bias all flows to a small number if a zero flow was encountered entering or exiting a unit operation identified with problems solving with discrete flows. This allowed the thermodynamic model to still calculate the dependent properties of the stream, while being below the threshold identified as a zero flow stream by iCON (identified graphically by a grey stream). This value was also chosen so as to remain within an appropriate dynamic range to avoid numerical issues with the simulator engine. A preferable solution to this would be to remove the equations altogether, and to use a user defined constant or previous solution. This would be possible in an equation based simulator such as Modelica, but not in a modular simulator such as iCON.

5. IMPLEMENTATION IN A CHEMICAL PROCESS SIMULATOR

A critical requirement for the steam utility package is the ability to not only arrange the pre-built unit operations and thereby construct a utility model, but also to implement strategies to ensure efficient and robust solving of these models.

5.1 Closed Loop Steam Systems

The term closed-loop when applied to a steam system context typically refers to the fact that the condensate return is also considered, meaning the complete loop from steam production, through process use, and through condensate collection is modelled. In reality, there are a number of loops present within a utility model, as shown in Fig. 6.

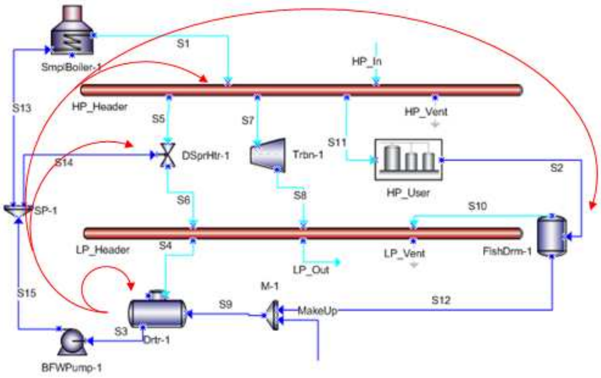

Fig. 6. Two Header steam system with calculation loops shown.

In order for the numerical solver to have some chance of establishing a solution, each of these loops must contain an initial estimate of temperature, mass flow, or enthalpy depending on the structure of the loop. All models built with iUO are assumed to be closed loop, thus an iterative recycle/root solver was required for every model.

5.2 Recycle Solver

As detailed in Section 4.1 utility models must be multipoint models, and thus require a number of internal variables to be established iteratively. As with most modular process simulation packages, iCON contains an iterative recycle solver for converging closed loops, in our case using a Broyden based algorithm.

When applying a recycle estimate to a utility model, it was discovered that the direction of the calculation flow (and thus the direction of information flow) must not be constrained to a single direction, as is the default in a process simulator Cota (2003). Taking the steam header temperature as an example, an estimate of temperature provides two solving paths: the first is the outlet temperatures to all equipment connected downstream of the header; and secondly the temperature completes the mass and energy balance of the upstream equipment, dictating the required mass flow of steam (deficit) from the header/boiler above.

Based on the above, the single direction information propagation of a recycle estimate was not sufficient for converging a utility model. Our solution was to implement what is known as a ‘controller’ in iCON. Simply put this is an iterative non-linear root solver, which adds the estimate to the ‘sheet convergence manager’. Controllers can be given a solve level, thus multiple controllers can be solved simultaneously (i.e. a multi-variable root solver), sequentially with different solve levels, or with some combination of the two.

By adding a controller to each variable we had identified as an estimate, together with heuristics for the solve level for each unit operation and a robust non linear solver, we skipped the need to add any recycles to our utility models, regardless of how many loops were created.

5.3 Solver Variable Identification

Key to solving these models so that the user did not inadvertently over specify specifications which would normally vary within typical optimization studies was the identification of degrees of freedom within a utility model. These were variables that initial estimates were required in order to begin solving, and would then be iterated until they converged, in much the same way a recycle works, but using controllers as detailed.

The heuristic we came up with to identify the estimation points follows from typical process simulation modelling. If the variable met the following rules, then it would be estimated:

(1) Expected to change between different optimization runs.
(2) After completing all specifications would still remain unsolved.
(3) If multiple variables remain, choose the one which is expected to vary least.

An example is the steam header unit operation, which would automatically balance (steam in = steam out). The obvious unknown is the amount of steam required from the above header/boiler. However this would not be known until all steam demands exiting the header were known, which required the header temperature to be known. In turn, the header temperature would not be known until the header feed (steam required) was known, and thus we can identify the need for an estimate with initial condition.

The problem is we can estimate either the steam mass flow required, or the header temperature. This was the need for rule 3, as we know from experience the temperature should not vary by more than a few degrees Celsius, while the mass flow may vary widely with process and operating conditions. With numerical conditioning and computational efficiency in mind we estimate the header temperature and thus the header will now iterate to converge the mass balance across it.

Other unit operations required taking a system-wide approach to looking for estimates. One of these was the deaerator, which closes two large loops: one between the lowest steam header and the Boiler Feed Water (BFW) return back to the steam boilers, and the second between the condensate return and the BFW return to the steam boilers. Each loop requires an estimate, and thus two extra estimates are added for each deaerator. This can be seen in Figure 6.

5.4 Solver Level Identification

A feature of the controller solver used for converging model estimates was the ability to parallelize (multi-variable) or constrain to sequential solving for multiple variables. iCON allows a solve level between 0 and 20, with variables on level 20 being solved first.

In order to determine which level an estimate should be placed on, the following heuristic rules were used:

- Lv 10: Variables which contain model wide loops, e.g. condensate return
- Lv 8: Variables which are specific to a unit operation, but influence other unit operations, e.g. header temperature
- Lv 1-5: Variables which are specific to a unit operation, and have no influence on a other unit operations, e.g. boiler air flow

Appropriate gaps between levels were used so that future rules or exceptions could be interleaved.

6. CONCLUSIONS

This paper has introduced the iCON Utility Optimizer utility modelling package and described the advanced functionality built for PETRONAS. Furthermore, it has detailed the unique characteristics of utility systems, and what problems can occur when building them in a package originally designed only for chemical process simulation. Solutions for these problems have been presented, together with heuristics for modelling utility systems within a process simulation environment.

ACKNOWLEDGMENTS:

Financial support to this project from the Industrial Information and Control Centre, Faculty of Engineering, AUT University and The University of Auckland, New Zealand is gratefully acknowledged.

REFERENCES

Raul Cota. *Development of an Open Source Process Simulator*. PhD thesis, University of Calgary, 2003.

Alan Eastwood. ProSteam – A Structured Approach to Steam System Improvement. In *24th IETC*, Houston, 2002.

Alan Eastwood and Chris Bealing. Optimizing the day-to-day operation of utility systems. *Energy Engineering*, 101(1):7–25, 2004.

Diego Fernandez-Polanco and A. Tainshm Richard. OptiSteam reduces operating costs at Spanish Refinery. *Driving Competitive Advantage*, 1Q04, 2004.

Diego Fernandez-Polanco, Alan Eastwood, and Nicola Knight. Real-time utility system optimization. In *Twenty-Sixth Industrial Energy Technology Conference*, Houston, Texas, 2004.

Gert Jan Fien. Make the most of it. *The Chemical Engineer*, 2008.

IAPWS. Revised Release on the IAWPS Industrial Formulation 1997 for the Thermodynamic Properties of Water and Steam, 2007.

W.F. Kenney. *Energy Conservation in Process Industries.* London, 1984.

Hiroyuki Kimura and X. X. Zhu. R-curve concept and its application for industrial energy management. *Industrial & Engineering Chemistry Research*, 39(7):2315–2335, 2000.

B. Linnhoff and E. Hindmarsh. The pinch design method for heat exchanger networks. *Chemical Engineering Science*, 38(5):745–763, 1983.

S. P. Mavromatis and A. C. Kokossis. Conceptual optimisation of utility networks for operational variations–I: Targets and level optimisation. *Chemical Engineering Science*, 53(8):1585–1608, 1998.

P. Varbanov, S. Perry, Y. Makwana, X. X. Zhu, and R. Smith. Top-level analysis of site utility systems. *Chemical Engineering Research and Design*, 82(6):784–795, 2004a.

P. S. Varbanov, S. Doyle, and R. Smith. Modelling and optimization of utility systems. *Chemical Engineering Research and Design*, 82(5):561–578, 2004b.