11. Process analysis approach to industrial ecology

Urmila Diwekar and Mitchell Small

Industrial ecology is the study of the flows of materials and energy in industrial and consumer activities, of the effects of these flows on the environment, and of the influence of economic, political and social factors on the use, transformation and disposition of resources (White 1994). Industrial ecology applies the principles of material and energy balance, traditionally used by scientists and engineers to analyze well-defined ecological systems or industrial unit operations, to more complex systems of natural and human interaction. These systems can involve activities and resource utilization over scales ranging from single industrial plants to entire sectors, regions or economies. In so doing, the laws of conservation must incorporate a number of interacting economic, social and environmental processes and parameters. Furthermore, new methods and data are required to identify the appropriate principles and laws of thermodynamics at these higher levels of aggregation (Ayres 1995a, 1995b).

Figure 11.1 presents a conceptual framework for industrial ecology applied at different scales of spatial and economic organization, evaluating alternative management options using different types of information, tools for analysis and criteria for performance evaluation. As one moves from the small scale of a single unit operation or industrial production plant to the larger scales of an integrated industrial park, community, firm or sector, the available management options expand from simple changes in process operation and inputs to more complex resource management strategies, including integrated waste recycling and re-use options. Special focus has been placed on implementing the latter via industrial symbiosis, for example, through the pioneering work of integrating several industrial and municipal facilities in Kalundborg, Denmark (Ehrenfeld and Gertler 1997). At higher levels of spatial and economic organization, for example, at national and, in recent years, global scales of management, policy may be implemented through the tools of regulation, economic incentives, taxation, trade policy and international agreements.

To evaluate the full range of options illustrated in Figure 11.1, highly quantitative information on chemical properties, thermodynamic constants and constraints are needed, as are data relating to firm, sector, national and global resource utilization and conversion. However, these data are often unavailable or difficult to obtain, and more qualitative, order-or-magnitude information must be developed and used. These different types of information are used in developing mass and energy balances, formulating process simulation tools and optimizing process designs. For the latter, multiple objectives and performance criteria must be considered. At the local scale, performance measures include conversion efficiency, throughput, cost and safety. While these factors remain applicable

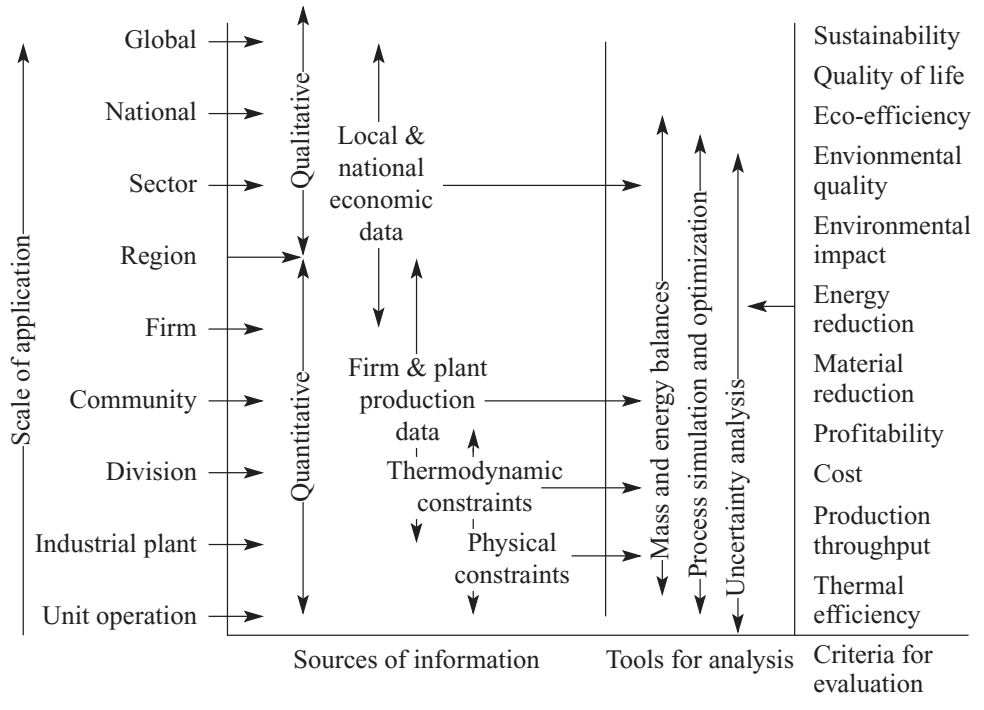

Figure 11.1 A conceptual framework for a process analysis approach to industrial ecology

at larger scales, broader metrics of overall resource use, the quality of the environment and the sustainability of economic activities are also considered. Industrial ecology provides a framework for integrating across these multiple scales of problem aggregation, assessment and analysis.

Process simulation provides a potentially useful basis upon which to begin to build assessments of this type. It is the purpose of this chapter to present an overview of the state of the art of methods for process simulation and optimization needed to develop a process-modeling approach to industrial ecology. To do this, we describe the development and state of the art of process simulators, their approach to mass and energy balance calculation, alternatives methods for linking multiple unit operations and processes, and methods for estimating parameters, searching and optimizing the design space, and incorporating multiple, possibly conflicting, objectives and performance uncertainty.

Commercial process simulators were first made available in the late 1970s and have since been used extensively by the chemical process industry to track unit operation performance and component flows. Current simulators are equipped with detailed process and cost models, and include elaborate physical property databanks. However, they lack several capabilities needed to provide a complete economic and environmental assessment. Problem formulation and system representation for industrial ecology applications require the characterization of material, energy and information flows and reservoirs, often at a combination of local, regional and global scales. Even for a narrowly defined

production process the necessary information for the full system may be highly dispersed among various organizations and organizational units (for example, see the analysis of a printed circuit board assembly process by Sheng and Worhach 1998). Such problems are only multiplied when dealing with multiple firms, industrial sectors or whole economies, and multiplied again when environmental impacts are added to the equation. A multi-objective approach to design under uncertainty is proposed to begin to address these assessment challenges.

PROCESS SIMULATION: AN ECOLOGICAL PERSPECTIVE

Process simulation involves the utilization of computer software resources to develop an accurate, representative model of a chemical process in order to understand its behavior in response to different inputs and control. In the past, process simulation was mainly concerned with the development of sophisticated unit operation blocks to predict mass flows of principal components through a process. In recent years, environmental consciousness has led to demands for tracking trace components (for example, resulting from fugitive emissions) that have an impact on environmental health and compliance, as well as major product and process components. Coupled with this demand for higher resolution models is the need for sophisticated computer-aided process design tools to identify low-cost, environmentally friendly solutions in the presence of considerable uncertainty. This calls for an integrated hierarchy of models, including modules with a high degree of detail for individual unit operations and process engineering activities, to simpler modules for analyzing system interactions at higher scales, with material flows and symbiotic interactions often controlled by exogenous factors, market forces or government regulation.

Many industries, both private and public, are involved in the transformation of raw material to useful products and by-products (some of which may be environmentally unacceptable). Several use process simulation tools to model their core production processes. These include chemical industries involved in the processing of organic and inorganic materials, the electric power industry involved in the transformation of fossil fuel to energy, and municipal treatment plants involved in the transformation of dirty to clean water. Effective facility operation is dependent upon accurate process simulation for assessing the material and energy flows through the process, so that the required thermal, environmental and economic performance can be assessed. These same process simulation tools have the potential to address programs and strategies to improve material and energy flows at higher scales of economic aggregation, providing guidance for industry, governments and citizens wishing to improve efficiency, sustainability and environmental quality through pollution prevention, material re-use, waste recycling, and material and energy conservation.

Process Simulation Tools

To understand how process simulation is used to model and design complex systems, the key components of a process simulation software package are identified and reviewed. The essential building blocks of a process simulator or 'flowsheeting' package include the following.

- Thermodynamic models: these are models developed to predict the different physical properties of the components under process conditions.
- Unit module models: these are routines that simulate the different unit operations (distillation, mixing, splitting, heat exchange and so on) and processes (reactions, mass and energy transfer, head loss).
- Data bank: the data on component physical properties, reaction rates and cost coefficients.

In addition to these, there are mathematical routines for numerical computations and cost routines for performing an economic analysis of the process.

Process simulation software can be classified as 'sequential modular', 'equation oriented' or 'simultaneous modular'. Traditionally, most simulators have adopted a sequential modular approach. With this approach, individual modules are developed for each unit operation and process. Output stream values are computed for each, given the input stream values and the equipment parameters. Each unit module in a flowsheet is solved sequentially. The overall flowsheet calculations in a sequential modular simulator follow a hierarchy. Thermodynamic models and routines are at the bottom of this hierarchy, followed by the unit operation modules that perform the necessary material and energy balances, based on the thermodynamic property routines. At the next level design specifications dictate iterative calculations around the units, superseded by the recycle iterations for stream convergence. Program utilities, such as parameter estimation and optimization, occupy the highest level in the calculation hierarchy in the sequential modular framework.

Equation oriented simulators define and solve a set of simultaneous non-linear equations that represent the process modules, mass and energy balances in the process. Although these simulators are more flexible in terms of information flow, they are more difficult to construct, and it is often difficult to diagnose errors when they occur. The simultaneous modular approach utilizes individual modules for each unit operation and process, as in the sequential modular approach, but attempts to establish a more immediate link among the inputs, outputs and operations of these individual modules. This is accomplished by defining a set of linear equations that approximately relate the outputs for each module to a linear combination of its input values. These equations are solved simultaneously in the simultaneous modular approach.

While efforts are under way to develop and advance equation oriented and simultaneous modular software systems for education and research applications, most of the currently available, widely applied commercial simulators are sequential modular in nature. However, as indicated in Table 11.1, significant effort has been made in recent years to develop and disseminate equation-oriented packages. There are no commercial simulators that use the simultaneous modular approach as yet.

Process simulators are also classified on the basis of their level of temporal aggregation; that is, whether the processes being considered are steady-state or dynamic in nature. Accordingly, steady-state and dynamic simulators are both available for modeling continuous processes. The sequential modular simulators shown in Table 11.1 are steady-state simulators. The equation-oriented simulators in the table can be used for both dynamic and steady-state analysis, but are mostly used for dynamic simulations.

The following example illustrates the use of ASPEN, a sequential modular simulator, to model the steady-state behavior of a benzene production process.

Table 11.1 Process simulation tools

Simulation Package	Type
FLOWTRAN	Sequential modular
FLOWPACK II	Sequential modular
PRO II (previously PROII)	Sequential modular
ASPEN Plus	Sequential modular
SPEEDUP	Equation-oriented
ASCEND	Equation-oriented
MODELLA	Equation-oriented
gPROMS	Equation-oriented

Modeling Benzene Production

The hydrodealkylation (HDA) of toluene to produce benzene is often used as a benchmark for demonstrating chemical process synthesis methods. The HDA process has been extensively studied by Douglas (1988) using a hierarchical design/synthesis approach. The problem presented and solved here is based on the flowsheet structure analyzed by Diwekar *et al*. (1992), which involved the selection of the flowsheet configuration and some of the operating conditions that maximize profit. The flowsheet for this case study is described below.

The primary reaction of the HDA process is

$$C_6H_5CH_3 + H_2 \rightarrow C_6H_6 + CH_4.$$

In addition to this desired reaction, an undesired reaction

$$2C_6H_6 \leftrightarrow C_6H_5 + H_2$$

also occurs. These homogeneous gas phase reactions occur in the range of 894 °K and 974 °K. A molar ratio of at least 5:1 hydrogen to aromatics is maintained to prevent coking. The reactor effluents must be quenched to 894 °K to prevent coking in the heat exchanger following the reactor.

The HDA flowsheet is shown in Figure 11.2. In this process, benzene is formed by the reaction of toluene with hydrogen. The hydrogen feed stream has a purity of 95 per cent (the rest is methane) and is mixed with a fresh inlet stream of toluene, a recycled toluene stream and a recycled hydrogen stream. The feed mixture is heated in a furnace before being fed to an adiabatic reactor. The reactor effluent contains unreacted hydrogen and toluene, benzene (the desired product), diphenyl and methane; it is quenched and subsequently cooled in a flash separator to condense the aromatics from the non-condensable hydrogen and methane. The vapor stream from the flash unit contains hydrogen that is recycled. The liquid stream contains traces of hydrogen and methane that are separated from the aromatics in a secondary flash unit. The liquid stream from the secondary flash unit consists of benzene, diphenyl and toluene. It is separated in two distillation columns. The first column separates the product,

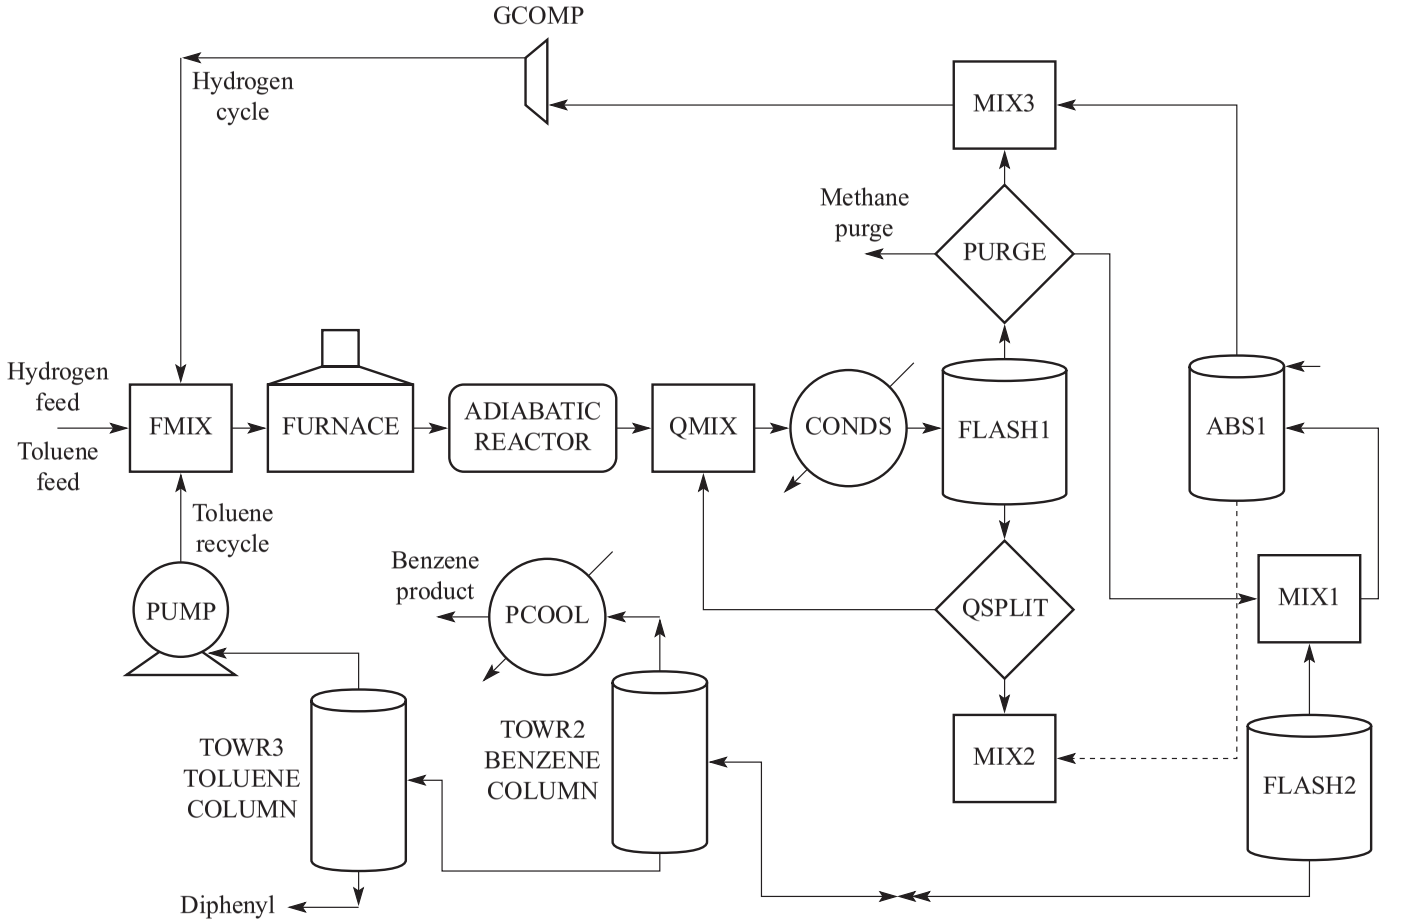

Figure 11.2 The process flowsheet for the production of benzene through the hydrodealkylation of toluene

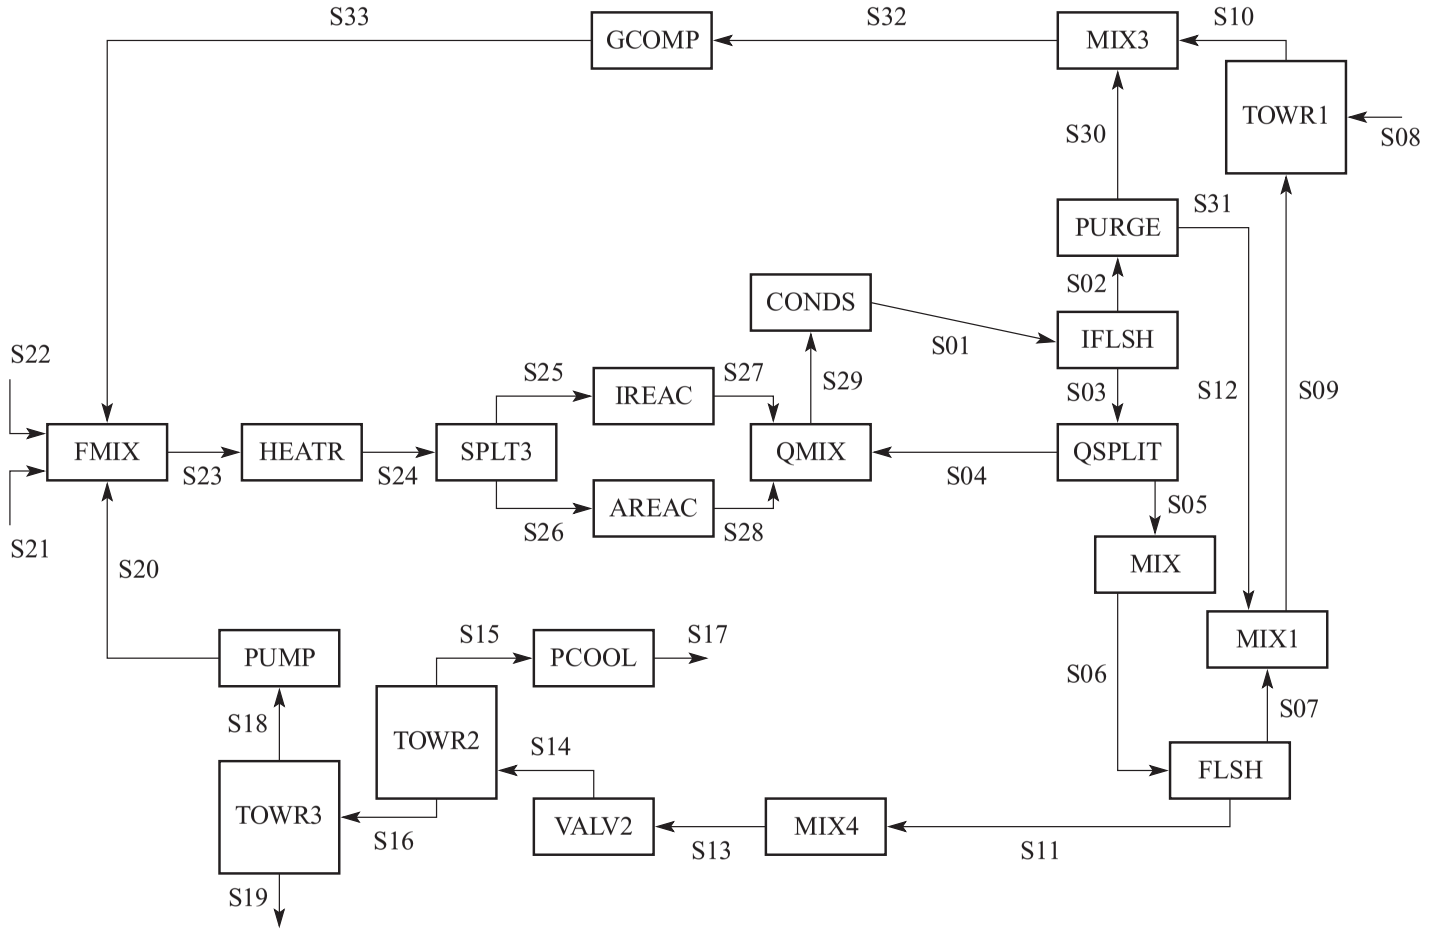

Figure 11.3 ASPEN representation of the HDA process

benzene, from diphenyl and toluene, while the second separates the diphenyl from toluene. The toluene is recycled back into the reactor.

Figure 11.3 presents the ASPEN representation of this flowsheet where unit operation blocks, including splitters, separators and reactors, are used as building blocks to track the material and energy streams through the complete process. Material and energy balances are computed around each unit and the system state variables are calculated, including component flows and system thermodynamic properties like enthalpy, entropy and so on, as shown in Table 11.2.

Table 11.2 Sample results for the HDA flowsheet simulation

Unit Operation Block Results

FLASH:2-OUTL (FLASH2): FLASH
INPUT STREAM(S): S01
OUTPUT STREAM(S): S02 S03
PROPERTY OPTION SET SYSOP0

Mass and Energy Balance

Conventional components		In	Out	Relative diff.
H_2	LBMOL/HR	2047.48	2047.41	3.66387e-05
CH_4	LBMOL/HR	2414.51	2414.54	−1.39574e-05
C_6H_6	LBMOL/HR	374.131	374.139	−2.16750e-05
C_7H_8	LBMOL/HR	227.842	227.837	2.21647e-05
$C_{12}H_{10}$	LBMOL/HR	16.8401	16.8394	4.22304e-05
Total balance				
MOLE	LBMOL/HR	5080.8	5080.76	7.66967e-06
MASS	LB/HR	95679.1	95679.5	−4.68125e-06
ENTHALPY	BTU/HR	−6.60813e+07	−6.60813e+07	1.48284e-05

Stream results

Stream ID		S01	S02	S03
From:		CONDS	FLASH	FLASH
To:		FLASH	PURGE	QSPLIT
H_2	LBMOLE/HR	2047.4828	2046.7644	0.6433
CH_4	LBMOLE/HR	2414.5081	2390.8028	23.7389
C_6H_6	LBMOLE/HR	374.1306	17.8004	356.3382
C_7H_8	LBMOLE/HR	227.8422	3.5885	224.2486
$C_{12}H_{10}$	LBMOLE/HR	16.84	0.096816	16.8336
TOTAL	LBMOLE/HR	5080.8038	4458.9569	621.8079
TEMP	DEGREES F	100	100	100
PRES	PSIA	465	465	465
ENTHALPY	BTU/LBMOLE	−13006	−16927	15105
V	FRACTION	0.8776	1.0	0.0
L	FRACTION	0.1223	0.0	1.0
ENTROPY	BTU/LBMOLE-R	−21.6285	−15.6219	−64.704
DENSITY	LBMOLE/CUFT	0.0463	0.0774	0.6473
AVG MW		18.8314	9.9133	82.785

Consider the extension of this simulation technology to track flows in a multi-plant or multi-sector analysis to achieve industrial symbiosis. Now, instead of unit level balances, each plant or sector is represented as a building block for the complete analysis. This can be achieved either by simplifying each process as a simple reactor unit operation block in a process network, as done by Chang and Allen (1997) in their analysis of 428 chemical processes that produce or consume 224 chemicals, or by using detailed recipes from chemical engineering textbooks and carrying out detailed elemental balances (Ayres and Ayres 1997). While the first approach offers a simplified solution to a large, complex problem, the second approach can address the problem of data inconsistencies and dispersion by constraining the system. This approach is described next.

THERMODYNAMIC AND OTHER CONSTRAINTS

One of the major problems in including industrial ecological concepts in design is the problem of data inconsistencies and dispersion. Even for a narrowly defined production process, the necessary information is highly dispersed and in various forms. These inconsistencies can be attributed to one or more of the following factors: (a) non-comparable units of measurements, (b) uncertainties in the assumptions, (c) confidential and non-verifiable data and data from unreliable sources, (d) measurement uncertainties, and (e) data violating laws of physics.

The first law of thermodynamics for conservation of mass and energy is applicable to every process network. It is, therefore, applicable to every firm and every industry that is in a steady state. This means that, for every process or process chain, the mass inputs must equal the mass of outputs, including wastes. Moreover, in many processes, non-reactive chemical components, such as process water and atmospheric nitrogen, can also be independently balanced. Thus various independent material balance constraints may have to be satisfied for each process. In short, systematic use of material balance conditions can increase the accuracy of empirical data by reducing error bounds (Ayres 1995a, 1995b). Alternatively, the material balance conditions can be used to 'fill in' missing data. Furthermore, material balance conditions are not the only basis for data augmentation. Energy conservation, constitutive relationships or statistical methods can also be used.

Process simulators are based on mass and energy balance principles. They utilize thermodynamic models and data, and hence are ideally suited for imposing these constraints on the available data. However, the constraints and data involved are not restricted to mass and energy balance principles, and are available in various forms. For example, it is common practice to report undetectable quantities of emissions in terms of the detection limit (or least count) of the measuring instrument (specifying that the data may be less than or equal to the detection limit). Sometimes the data are reported in order-of-magnitude terms (for example, refer to Case 3 in Ayres 1995a, where the Benzo(a)pyrene content is reported to be much smaller than 0.0001). Furthermore, discrete, categorical information about the occurrence or non-occurrence of particular reactions, or the presence or absence of reaction by-products, may be available.

Given that knowledge is available in various forms (for example, quantitative models for material and energy balances, order-of-magnitude information, qualitative information and logical information), a unified framework that incorporates information of each

type in its inference is desirable. Optimization methods combined with artificial intelligence techniques, as proposed in Kalagnanam and Diwekar (1994), provide such a framework, in which information can be represented as inequality constraints. Unlike numerical methods for solving equations (equality constraints), optimization methods can handle both equality and inequality conditions and hence can be used to make inferences from data in various forms.

Defining Objectives and Goals

As stated earlier, methods for assessing economic impacts and profitability have been available for a number of years. However, methods and measures for characterizing environmental impacts and sustainability are as yet in their infancy. Recent attempts at defining ecological impacts for use in life cycle assessment and similar industrial ecology applications include the environmental burden system by ICI (Wright *et al.* 1998), sustainability indicators by Tyteca (1999), ecological risk indicators described by Koenig and Cantlon (2000), exergy as a unifying indicator for material and energy transformation (Ayres 1995b), environmental damage indices (DeCicco and Thomas 1999) and the generalized waste reduction (WAR) algorithm (Cabezas *et al*. 1997; Cabezas *et al*. 1999; Young and Cabezas 1999). The WAR algorithm uses a series of indices characterizing different environmental, social and economic impacts. With WAR the potential environmental impact is defined in terms of the pollution index, calculated by multiplying the mass of each pollutant emitted by a measure of its potential impact, then summing over all pollutants. This index is a carefully constructed function encompassing a comprehensive list of human health and environmental impacts for each chemical (see Table 11.3). However, like the other methods described above, the WAR index provides a highly simplified representation of environmental impacts. For example, effects of pollutants emitted to different media are not differentiated in the WAR algorithm. Chemical exergy content likewise provides only a partial insight into environmental impact, since it cannot be directly linked to toxicity to humans or other organisms. Nonetheless, these impact assessment methods provide a first-order qualitative indication of the environmental damage and hence a useful staring point for analysis.

Table 11.3 The potential environmental impact categories used within the WAR algorithm

Local Toxicological		Global Atmospheric	Regional Atmospheric
Human	Ecological		
Human toxicity potential by ingestion (HTPI)	Aquatic toxicity potential (ATP)	Global warming potential (GWP)	Acidification, or acid rain potential (ARP)
Human toxicity potential by exposure, dermal and inhalation (HTPE)	Terrestrial toxicity potential (TTP)	Ozone depletion potential (ODP)	Photochemical oxidation potential or smog formation potential (PCOP)

Recently the WAR algorithm was added to the ASPEN simulator to allow consideration of the eight environmental impacts shown in Table 11.3. This was easily done, since chemical simulators keep track of mass balance and emissions information required for calculation of these indices. Similarly, the unified indicator based on exergy proposed by Ayres (1995b) is readily computed using process simulation technology, since most commercial simulators have a unit operation block based on the 'concept of Gibbs free energy minimization'.

Once different environmental impacts are calculated, they must be weighted and balanced against each other, as well as other concerns, such as cost and long-term sustainability. These multiple, often conflicting, goals pose significant challenges to process optimization and design. How can designs be identified that best satisfy multiple objectives? Multi-objective optimization algorithms provide a particularly useful approach, aimed at determining the *set* of non-dominant/non-dominated ('Pareto') designs where a further improvement for one objective can only be made at the expense of another. This determines the set of potentially 'best' designs and explicitly identifies the trade-offs between them. This is in contrast to cost–benefit analysis, which deals with multiple objectives by identifying a single fundamental objective and then converting all the other objectives into this single currency. The multi-objective approach is particularly valuable in situations where there are a large number of desirable and important production, safety and environmental objectives which are not easily translated into dollars. Formulation of a process simulation and optimization model with multiple objectives is illustrated in the following section, with particular application to the HDA benzene synthesis problem.

ECOLOGICAL AND ECONOMIC CONSIDERATIONS: A MULTI-OBJECTIVE OPTIMIZATION PROBLEM

As stated earlier, algorithms such as WAR provide a first approximation of environmental objectives. However, the various environmental impact indices and economic objectives are so different in terms of evaluation, quantification and magnitude that the choice of relative weights for environmental and economic impacts depends upon the decision makers' perspectives. Thus it is necessary to provide decision makers with the complete economic and environmental surface, so that they can understand the full set of alternatives and the trade-offs among them in terms of the desired objectives.

A Multi-objective Optimization Framework

As is well known, mathematics cannot isolate a unique optimum when there are multiple competing objectives. Mathematics can at most aid designers to eliminate alternatives dominated by others, leaving a number of alternatives in what is called the 'Pareto set' (Hwang *et al*. 1980). For each of these alternatives, it is impossible to improve one objective without sacrificing the value of another, relative to some other alternatives in the set. From among the dominating solutions, it is then a value judgment by the customer to select which alternative is the most appropriate. At issue is an effective means of finding the members of the Pareto set for a problem, especially when there are more than two or

three objectives, the analysis per design requires significant computations to complete and there are a very large number of feasible alternatives.

For example, consider the generalized WAR algorithm which expresses environmental objectives in terms of potential impact indices combined together using weighting factors, a_i. This formulation used in the WAR algorithm can be easily expressed in terms of the weighting method for a multi-objective optimization problem where different weights are assigned to obtain the Pareto surface. The generalized formulation is shown below:

$$\text{Min } \dot{I}_{out}^{(NP)} = \sum_{i=1}^{EnvCat} \alpha_i \dot{I}_i^{(NP)}$$

subject to:

mass and energy balance constraints
decision variables bounds

where $\dot{I}_i^{(NP)}$ is the rate of potential environmental impact generation for all the NP products and $\dot{I}_{tot}^{(NP)}$ is the weighted sum across these.

The basic strategy of the weighting method is to transform the multi-objective problem into a series of single objective problems with weighting factors assigned to each objective. The Pareto set can be derived by solving the large number of optimization problems created by modifying the weighting factors of the objectives. However, the major disadvantage of using the weighting method is its inefficiency that arises because of the large number of optimization problems that must be solved for the different linear combinations of objectives. It is also difficult to direct and limit the search to the region of the non-dominated surface which the decision maker prefers.

The constraint method is another technique for generating an approximation of the Pareto set. The basic strategy also is to transform the multi-objective problem into a series of single objective problems. A purely algorithmic solution approach (Cohon 1978) is to select one of the objectives to maximize (for example, profit) while each of the other objectives (for example, potential environmental impacts) is turned into an inequality constraint with a parametric right-hand side ($\varepsilon_1, \ldots, \varepsilon_p$). It is important to note that each optimal design in the Pareto set derived from a combination of α_i ($i=1, \ldots, EnvCat$) by the weighting method, can be alternatively generated from a corresponding combination of ε_i ($i=1, \ldots, EnvCat$) by the constraint method. One is a mapping of the other. For example, the upper bound of ε_i used in the constraint method correspond to $\alpha_i=0$ in the weighting method, and the lower bound of ε_i used in the constraint method corresponds to $\alpha_i=+\infty$ in the weighting method. In other words, the upper and lower bounds of ε_i cover the entire range between 0 and $+\infty$ for α_i. The constraint method offers the advantages of better control over exploration of the non-dominated set and of locating points anywhere along the non-dominated surface.

The constraint method based on profit

$$\textbf{max} \qquad Profit$$

subject to:

$\dot{I}_i^{(NP)} \leq \epsilon_i, i=1, \ldots, EnvCat$
mass and energy balance constraints
decision variables bounds

By selectively decreasing each ε_i and rerunning the optimization (resulting in a lower maximum profit), the analyst explicitly identifies the trade-off between the profit that must be forgone to achieve improved environmental performance in environmental category i (that is, lower $\dot{I}_i^{(NP)}$). Of course, such trade-offs only occur on this outer envelope of the Pareto surface; the method generates designs where *mutual* improvement in both environmental quality and profit have been achieved to the maximum extent possible.

Determination of the Pareto set for various potential environmental impacts and economic objectives, using either the weighting or the constraint method, requires solution of a large number of optimization problems. For example, if there are six objectives and five of them are evaluated over 10 levels, we must solve 100000 optimization problems. To circumvent this problem, a new and efficient multi-objective optimization algorithm based on the constraint method has been developed (Fu and Diwekar, 2001). Figure 11.4 illustrates the major features of this algorithm. In the outer loop, problem inputs are specified and the optimization problem is defined in terms of the objective function and

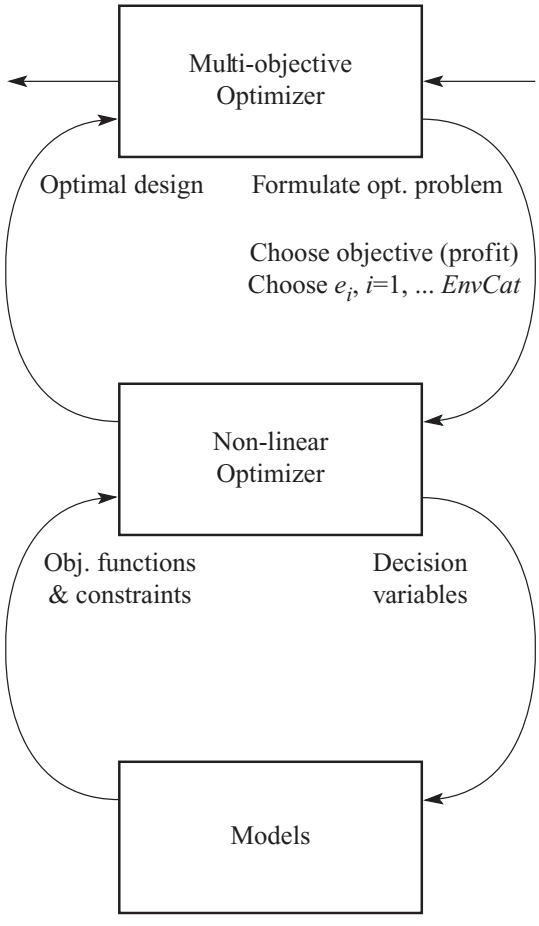

Figure 11.4 A generalized multi-objective optimization framework

constraints that will be systematically varied. In the middle loop a non-linear optimization program is called to search the decision variable (design) space for the optimal design, returning the value of the objective function and noting which constraints are met and the amount of slack for each, indicating which are binding. The non-linear optimizer requires multiple calls to the process simulation model in the inner loop of the algorithm. The overall search and bounding method is designed to search and map out efficiently the Pareto set of design alternatives. In the initial application of the method and the example that follows, only representative results are presented in terms of the bounds of the different objectives considered. In particular, the Pareto set is approximated by obtaining optimal designs with two values of ε_i ($i=1, \ldots$, *EnvCat* and $i=$ profit) for each objective function ($\dot{I}_i^{(NP)}$, $i=1, \ldots$, *EnvCat*, and Profit). As such, as a first step in obtaining the overall Pareto set, each environmental objective is characterized in terms of its relative trade-off with respect to profit.

Benzene Production: a Benchmark Multi-objective Example

As the first step toward a multi-objective analysis with broad consideration of ecological protection, the benchmark problem of hydrodealkylation of toluene to form benzene is again considered here. The simulation model for this flowsheet was described earlier. The objective is to illustrate the benefits of using the multi-objective optimization framework to obtain alternatives with minimum environmental impacts and maximum profit.

The important control parameters – molar flow rate of the hydrogen feed, molar flow rate of the toluene feed, furnace temperature and conversion of the adiabatic reactor – are chosen as the decision variables for this multi-objective analysis. The different objectives include maximizing the annualized profit and minimizing the different environmental impacts of the output as calculated by the WAR algorithm, subject to the following process and product constraints:

- the benzene production rate must be maintained at 120kmol/hr;
- the adiabatic reactor must have a volume less than 500m^3;
- the hydrogen feed must have a purity of 95 per cent.
- the purity of the benzene product is at least 95 per cent.

The potential environmental impact indices for each of the chemicals present in the HDA process are listed in Table 11.4. The ozone depletion potential (ODP) indices and acid rain potential (ARP) indices for all components are zero, hence there is no need to include them as separate objectives. Furthermore, the indices for all components of human toxicity potential by ingestion (HTPI) and terrestrial toxicity potential (TTP) are equivalent, and the optimal solutions for minimizing or maximizing them are the same; hence they are listed together. The reduced problem thus includes six total objectives: HTPI or TTP, HTPE, ATP, GWP, PCOP and an economic objective. The economic objective is represented in terms of the annualized profit. The cost model (Diwekar *et al.* 1992) is represented by linear, fixed-charge costs. For details of this case study, please see Fu *et al.* (2000).

In the HDA process, benzene is the desired product and diphenyl, which is also formed during this process, can be either treated as a pollutant or sold as a by-product. These two cases are considered here.

Table 11.4 Potential environmental impact indexes for the components in the HDA process

	HTPI	HTPE	ATP	TTP	GWP	ODP	PCOP	ARP
Hydrogen	0	0	0	0	0	0	0	0
Methane	0	0	0.057	0	0.0035	0	0.014	0
Benzene (product)	0	0	0	0	0	0	0	0
Toluene	0.078	2.2e-06	0.065	0.078	0	0	1.2	0
Diphenyl (pollutant)	0.12	0.0016	0.88	0.12	0	0	0	0
Diphenyl (by-product)	0	0	0	0	0	0	0	0

Case 1: diphenyl as a pollutant

Figure 11.5 shows the results of 10 different optimal designs for the HDA process by minimizing and maximizing each of the nine objective functions and removing the duplicate designs. These designs represent a first approximation to the complete Pareto surface consisting of many such designs. While this is only an initial exploration of the design space, the results in Figure 11.5 do show that profit does not conflict with the environmental criteria in all cases, and one can find an optimal design that is effective in meeting both economic and environmental objectives. This is due to the non-convex nature of the objective surface for the HDA process. From the figure, it can be seen that designs 1 and 2 are likely to be deemed superior to the others since (a) the profit of design no. 1 is the highest of all 10 designs and its environmental impacts (except for the value of PCOP) are lower than designs 3 to 10; and (b) design no. 2 has the best environmental performance (except PCOP), and its profit is only 7.4 per cent less than that of design no. 1. Further, the results

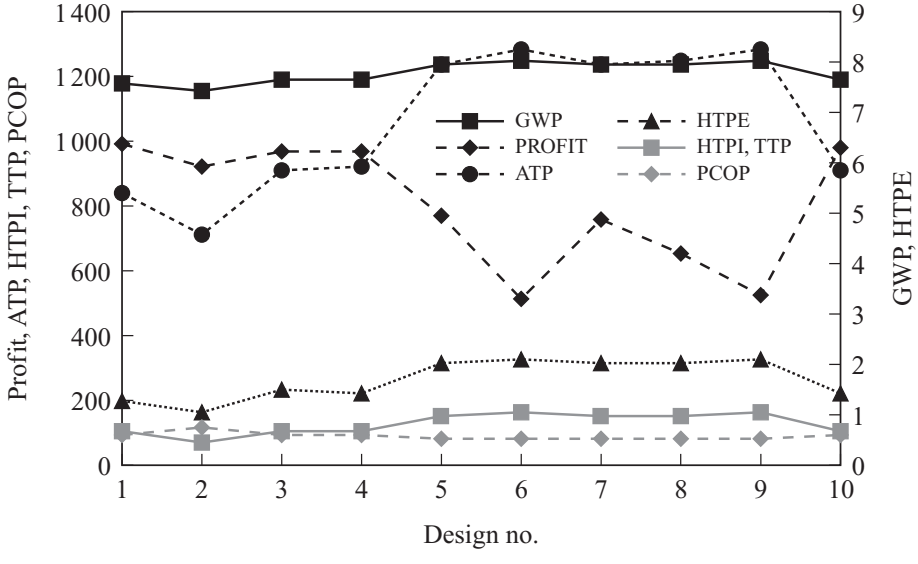

Figure 11.5 Approximate Pareto set for the HDA process multi-objective optimization (case 1: diphenyl as a pollutant)

indicate that designs 6 and 8 are the worst designs as they have the lowest profit and high environmental impacts.

Case 2: diphenyl as a by-product

Figure 11.6 presents the results of 10 different designs for the HDA process considering diphenyl as a by-product (that can be sold in the market). Again these results were obtained by minimizing and maximizing each of the six objective functions using the non-linear optimizer, and then removing the duplicate designs.

As in the previous case, profit follows a trend similar to that of a number of the projected environmental impacts: HTPI, TTP, PCOP and HTPE indicating once again the potential for designs that are both economically and environmentally attractive. However, the desirable and undesirable designs suggested by the analysis do differ from those derived for the previous case. When diphenyl is treated as a marketable product, rather than a pollutant, the undesirable designs can clearly be eliminated; however, the best designs are more difficult to identify. For example, designs 2, 6 and 8 in Figure 11.6 are likely to be eliminated when compared to design no.5, because designs 2, 6 and 8 have higher environmental impacts (along most dimensions, though a number of the environmental impact indices exhibit only a small amount of variation between designs, once diphenyl is removed from consideration as a pollutant) and lower profit than does design no.5.

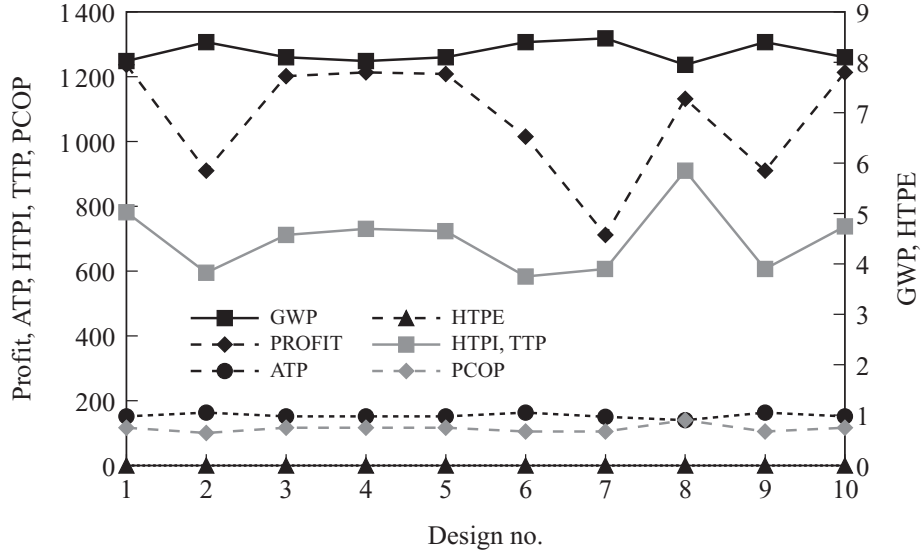

Figure 11.6 Approximate Pareto set for the HDA process multi-objective optimization (case 2: diphenyl as a by-product)

The multi-objective framework presented in this example helps to identify choices for the designer amongst the different economic and environmental objectives considered. In this case, we suggest that especially good designs will be those that have (a) higher profit and lower environmental impacts, (b) lower environmental impacts at the expense of

small profit loss, (c) higher profit at the expense of slightly higher environmental impacts, and (d) lower environmental impacts for some categories at the expense of slightly higher environmental impacts in others. It is then a value judgment by the decision maker(s) to determine which design among these is the most appropriate.

The analysis presented thus far has been *deterministic*: relationships between the system design and economic and environmental performance are assumed to be known and modeled with certainty. In reality major uncertainties are usually present, and these can have a significant effect on the results. Methods for addressing uncertainty in process simulation, design and optimization are considered in the following section, as the next major challenge to implementing efficient, environmentally conscious process design.

A MULTI-OBJECTIVE OPTIMIZATION FRAMEWORK UNDER UNCERTAINTY

A probabilistic or stochastic modeling procedure involves (a) specifying the uncertainties in key input parameters in terms of probability distributions, (b) sampling the distribution of the specified parameter in an iterative fashion, (c) propagating the effects of uncertainties through the process flowsheets, and (d) applying statistical techniques to analyze the results (Diwekar and Rubin 1991).

As regards specifying uncertainty using probability distributions, to accommodate the diverse nature of uncertainty, different probability distribution functions can be used (Morgan and Henrion 1990; Taylor 1993; Cullen and Frey 1999). Some of the representative distributions are shown in Figure 11.7. The type of distribution chosen for an uncertain variable reflects the amount of information that is available. For example, the uniform and log uniform distributions represent an equal likelihood of a value lying anywhere within a specified range, on either a linear or a logarithmic scale, respectively. A normal (Gaussian) distribution reflects a symmetric, but decreasing, probability of a parameter value above or below its mean value. Normal distributions often result from the summation of multiple errors, and are often used to represent small measurement errors. In contrast, log normal distributions are positively skewed (with a heavy upper tail) and often result from multiplicative, order-of-magnitude variation or errors. Triangular distributions indicate a higher probability towards the mid-range of values, but may be specified to be symmetric, positively or negatively skewed. A beta distribution provides a wide range of shapes and is a very flexible means of representing variability over a fixed range. The standard beta distribution, for random variables between zero and one, is often used to represent uncertainty in a chemical mixture fraction or a product or process failure probability. Finally, in some special cases, empirical, user-specified distributions can be used to represent arbitrary characterizations of uncertainty (for example, fixed probabilities of discrete values based on observed samples).

Once probability distributions are assigned to the uncertain parameters, the next step is to sample the uncertain, multi-variable parameter domain. Alternatively, one can use collocation-based methods to derive a response surface of the actual uncertainty surface (Tatang 1994). Although this method can require significantly fewer runs than a sampling method, one needs to have substantial knowledge of the model, and discontinuities or

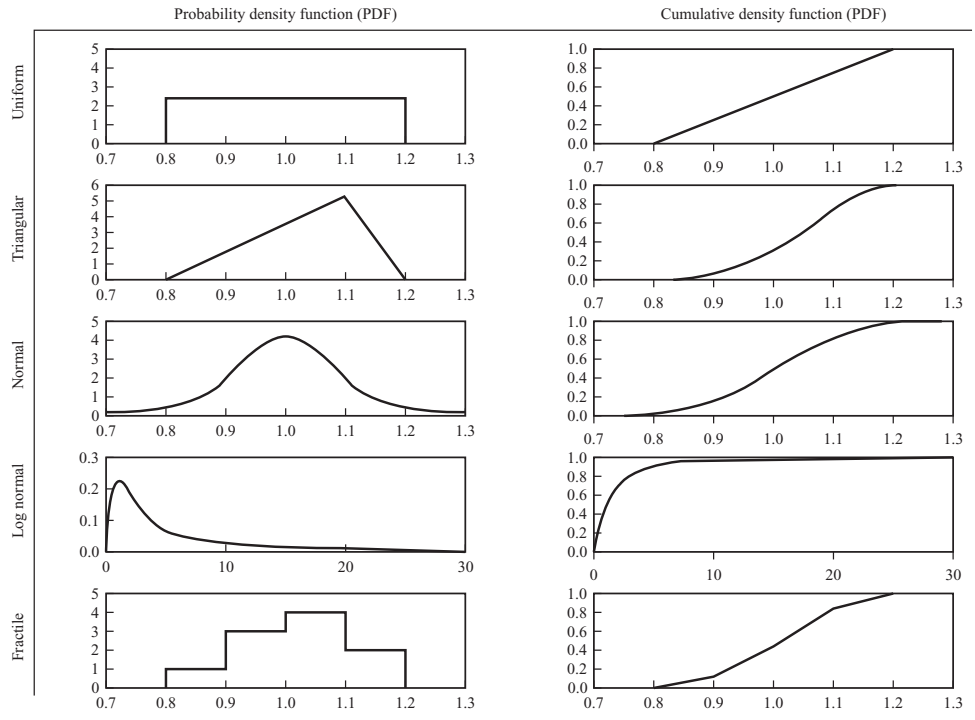

Figure 11.7 Probabilistic distribution functions for stochastic modeling

non-smoothness can result in erroneous results. Thus sampling methods provide the most generally applicable approach, and are discussed in more detail below.

One of the most widely used techniques for sampling from a probability distribution is the Monte Carlo method, which samples the parameter *in a purely random* manner, that is, all samples are independent and identically distributed from the overall target distribution. The main advantage of the Monte Carlo method is that the simulation results can be analyzed using classical methods of statistical estimation and inference. Nevertheless, in most applications, the actual relationship between successive points in a sample has no physical significance; hence the randomness/independence of successive samples used in approximating the target distribution is not critical. In such cases, constrained or stratified sampling techniques can allow better representation of the target distribution with a much small sample size.

Latin hypercube sampling is one form of stratified sampling that can yield more precise estimates of the distribution function. Here the range of each uncertain parameter X_i is sub-divided into non-overlapping intervals of equal probability. One value from each interval is selected at random with respect to the probability distribution within the interval. The n-values thus obtained for X_1 are paired in a random manner (that is, equally likely combinations) with n-values of X_2. These n-values are then combined with n-values of X_3 to form n-triplets, and so on, until n k-tuplets are formed. The main drawback of this stratification scheme is that, while it is uniform in one dimension, it does

not ensure uniformity properties in k dimensions. Recently, an efficient sampling technique based on Hammersley points (Hammersley sequence sampling – HSS) has been developed (Diwekar and Kalagnanam 1997). This method uses an optimal design scheme for placing the n-points on a k-dimensional hypercube. This scheme ensures that the sample set is more representative of the population, showing uniformity properties in multi-dimensions, unlike Monte Carlo, Latin hypercube and its variant, the Median Latin hypercube sampling technique. It has been found that the HSS technique is at least three to 100 times more efficient than LHS and Monte Carlo techniques and hence is a preferred technique for uncertainty analysis, as well as optimization under uncertainty. Limitations in the effectiveness of methods with high uniformity (such as HSS) can occur when the uncertain parameters exhibit highly periodic properties or effects, however, such cases are expected to be unusual in most process design applications.

An Efficient Multi-objective Optimization under Uncertainty

The following is a brief description of an efficient multi-objective optimization framework under uncertainty based on the HSS technique. The details of the algorithms can be found in Fu and Diwekar (2001). Figure 11.8 shows the generalized framework of multi-objective optimization under uncertainty. Once again, as in Figure 11.4, the outer multi-objective optimization framework is used to formulate a number of optimization problems to generate optimal alternative solutions within the Pareto set. The innermost loop incorporates uncertainty by converting the deterministic model to a stochastic one. In the innermost loop, the HSS technique is used to generate distributions of uncertain parameters, which then map into the probability distribution of corresponding objective functions and constraints computed by the model. In the outermost loop the HSS technique is also employed to formulate combinations of right-hand sides for the constraint method so that the minimum number of optimization problems can be identified and solved to attain an accurate representation of the whole Pareto set. By using an efficient sampling technique for uncertainty analysis and for multi-objective optimization, and an efficient iteration between the optimizer and the sampling loop, this approach allows significant computational savings, bringing a number of real-world, large-scale problems that were previously unsolvable within reach for effective design and optimization. The case study for benzene production, used earlier to illustrate the usefulness of the multi-objective optimization framework, is now used to illustrate the effect of considering uncertainty in determining the solution surface.

Benzene Production: Uncertainty Results

As a first step, with the same HDA flowsheet, we ignore the uncertainties in economic and other input parameters. We again consider the case with diphenyl assumed to be a useful byproduct, so that the potential environmental impact values for the diphenyl stream are set to zero. To allow for uncertainty, all non-zero potential environmental impacts are assumed uncertain with the log-normal distributions shown in Table 11.5. The common standard deviation of the logs in each case implies that the uncertainty range (plus or minus three standard deviations) for each of the potential environmental impact values ranges from a factor of 10 below to a factor of 10 above the deterministic values assumed

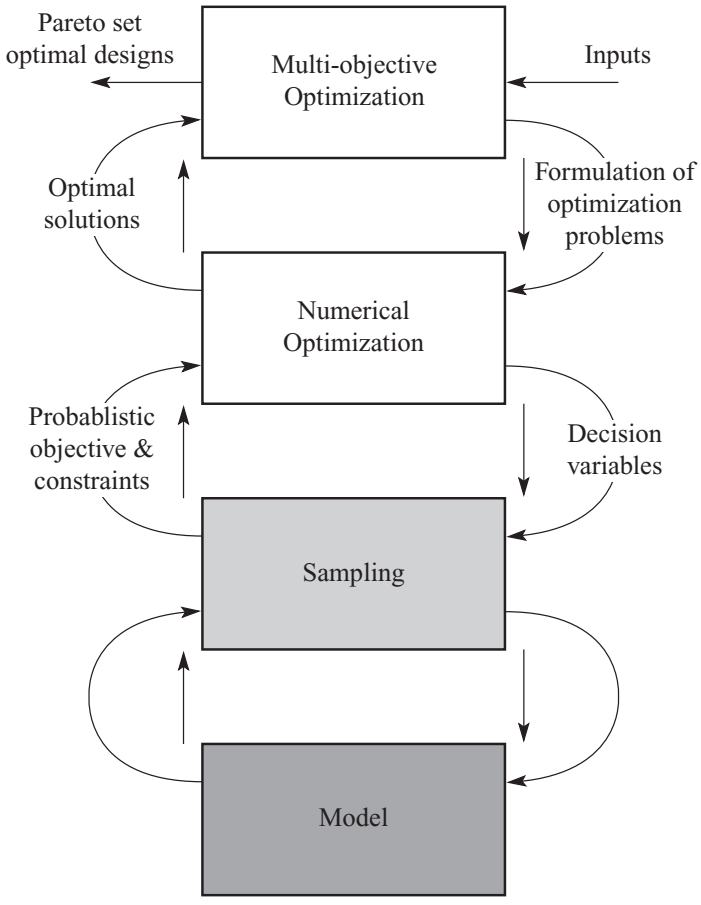

Figure 11.8 The multi-objective optimization under uncertainty framework

in Table 11.4. Two-order of magnitude uncertainty in environmental impacts is not uncommon, given the highly diverse and aggregate nature of the environmental indices considered.

Since there are no uncertainties factored into the economic objective, the same optimal design is obtained when only profit maximization is considered, as was obtained for the deterministic case. However, the results for the environmental objectives are no longer single values; rather they now follow probability distributions. The results for these are described by the cumulative distribution functions shown in Figure 11.9. Also shown in this figure (with a vertical line in each case) are the environmental impact index values determined for the previous, deterministic case. The results indicate that there is a 59 per cent probability that PCOP will be higher than the deterministic estimate shown by the horizontal line, and similarly 54 per cent, 52 per cent, 53 per cent and 51 per cent probabilities for HTPI, HTPE, ATP and GWP, respectively, being higher than the deterministic estimates when uncertainty factors in the potential environmental impacts are considered. While useful knowledge, this type of uncertainty characterization, after the

Table 11.5 Uncertainty quantification in environmental impacts indices for the components in the HDA process

	HTPI	HTPE	ATP	TTP	GWP	ODP	PCOP	ARP
Hydrogen	0	0	0	0	0	0	0	0
Methane	0	0	U3	0	U5	0	U6	0
Benzene (product)	0	0	0	0	0	0	0	0
Toluene	U1	U2	U4	U1	0	0	U7	0
Diphenyl (by-product)	0	0	0	0	0	0	0	0

		Parameters (of natural log)	
Uncertainty factors	Type of distribution	Mean	Standard deviation
U1	Lognormal	2.551	0.7675
U2	Lognormal	−13.0271	0.7675
U3	Lognormal	−2.7336	0.7675
U4	Lognormal	−5.655	0.7675
U5	Lognormal	−6.5713	0.7675
U6	Lognormal	−4.2687	0.7675
U7	Lognormal	0.1823	0.7675

design is specified, does little to indicate how explicit consideration of uncertainty may have redirected the design in the first place. For this, an integrated procedure, such as that show in Figure 11.8, is needed.

One approach for using the information on environmental impact uncertainty explicitly in the design optimization is to define a probabilistic objective function in terms of the mean (environmental impact and/or cost), the probability of exceeding certain values of these, the variance, or the median value of the objectives, depending on the decision maker's choice. For illustration purposes, we choose the mean value of each potential environmental impact to include as part of the objective function.

Figure 11.10 shows the different mean potential environmental impacts and profit for 10 optimal designs generated as an approximation to the Pareto set under uncertainty. The trends for the potential environmental impacts are similar to those determined for the deterministic case. This can be attributed to the fact that we have considered only the uncertainties in the environmental impacts for each component and these quantities are related to each environmental objective via a linear function. However, even from this first stage analysis, it is apparent that the relative effects of uncertainties on each objective function are different. This is illustrated in Figure 11.11, which demonstrates that, while the uncertainties in environmental impact have little impact on the profit, the mean environmental impacts are higher in the case where uncertainties are explicitly considered.

This benzene production case study is carried out entirely in the ASPEN simulator environment and provides a first step toward the process analysis approach to industrial ecology presented in this chapter. The major result is that environmental objectives need not conflict with economic benefits, as is often believed. This approach can be easily extended to industrial symbiosis. For example, Chang and Allen (1997) used multi-objective optimization combined with simplified material and energy balance models to identify chemical

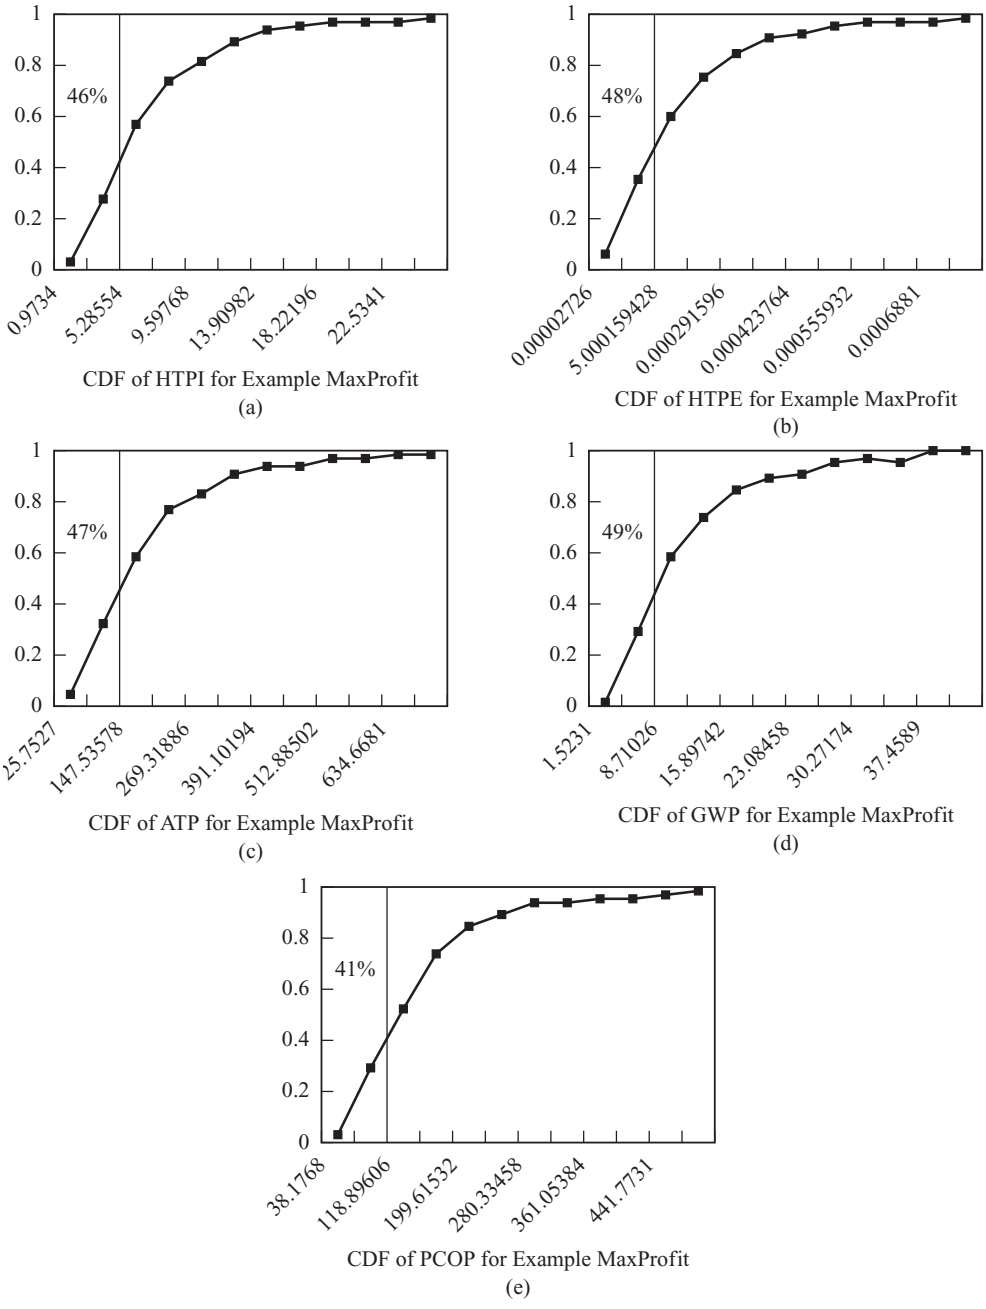

Figure 11.9 Uncertainty quantification in environmental impacts indices for the case study

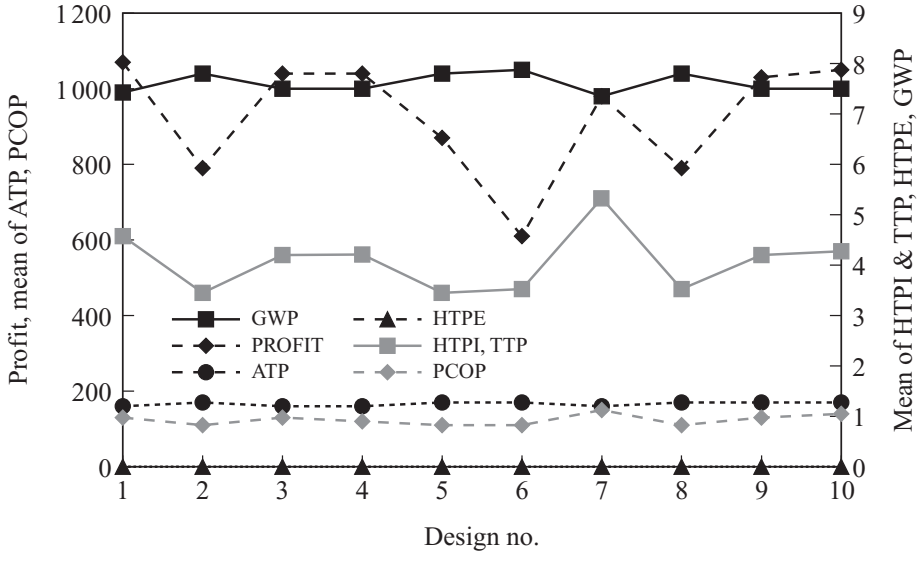

Figure 11.10 Approximation of Pareto set for the uncertainty case

manufacturing technologies for chlorine use for various industrial systems. However, to address the question of inaccuracies in the models, and lack of data, the problem of uncertainty (not considered by Chang and Allen) must be dealt with using methods such as those presented herein.

CONCLUSIONS

This chapter has presented a conceptual framework for a process analysis approach to industrial ecology. Current process simulation technology based on mass and energy balance principles can provide a unified framework for this approach. The capabilities of existing process simulation tools and their deficiencies in performing this task have been elucidated. A multi-objective optimization framework provides a mechanism to include the multiple, often conflicting, goals associated with industrial ecology. However, to address the issues of accuracy and relative weights assigned to these goals one must wrestle with the problem of uncertainty – in this case addressing how to value different environmental impacts, some of which are well characterized and some highly speculative. Uncertainty analysis coupled with the multi-objective framework can be truly beneficial in this context. This framework can also provide a basis for dealing with the problem of dispersed and scarce data, given that there is little or no commercial experience with industrial symbiosis, or with applying industrial ecology at larger scales, in practice. While the case study of benzene production illustrates the usefulness of the process analysis approach to industrial ecology using multi-objective optimization under uncertainty, we expect that applications at higher levels of economic aggregation, at the plant, community, national and even global level, will one day provide comparable insights into broader strategies for improving economic and environmental sustainability.

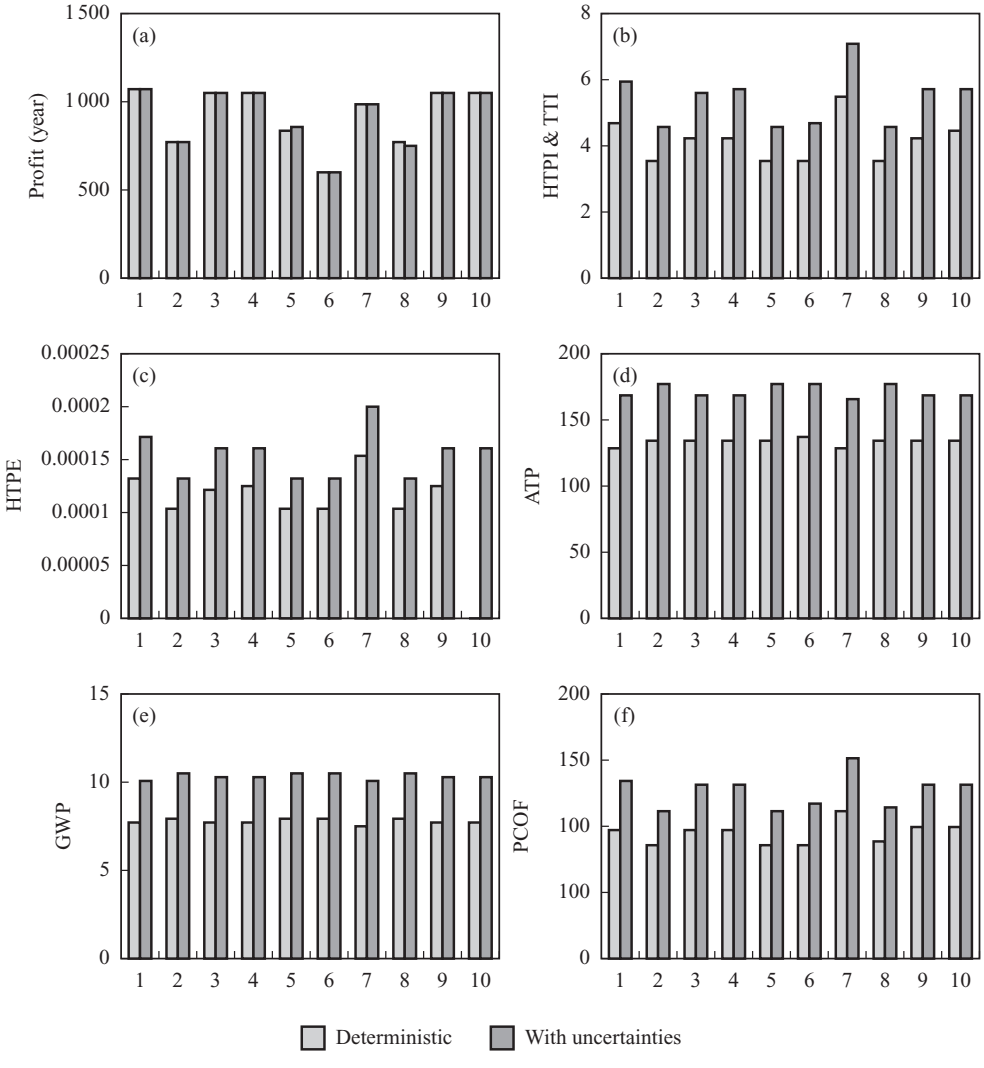

Figure 11.11 Relative effects of uncertainties on different objectives