Managing Conflicts among Decision-Makers in Multiobjective Design and Operations*

Victor M. Zavala
Mathematics and Computer Science Division
Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439, USA

Abstract

We present a systematic framework to manage conflicts among multiple decision makers (stakeholders) arising in the multiobjective design and operations of process systems. Addressing such situations is particularly relevant in sustainability studies because many conflicting social, environmental, and economic objectives need to considered. The proposed framework factors in the opinion of the stakeholders and computes a compromise solution that seeks to minimize a measure of their dissatisfactions. We propose to use conditional-value-at-risk (CVaR) as a measure of dissatisfaction as this provides a generalization of average and worst-case metrics considered previously in the literature. In addition, the use of CVaR enables us to shape the distribution of dissatisfactions and to avoid extreme conservativeness of worst-case solutions. A key advantage of the proposed framework is that it does not require the computation of a Pareto front and can thus be used to address problems with many stakeholders and objectives. Examples are presented to illustrate the concepts.

Keywords: multiobjective, stakeholders, disagreement, decision making.

1 Introduction

Almost any decision-making activity must resolve conflicts among multiple stakeholders. Conflicts arise because stakeholders have different opinions and perceptions on the economic, environmental, and safety metrics (objectives) that should be used and/or on how they should be prioritized [1]. As an example, in designing an infrastructure that supports an urban area; the community, local

*Preprint Number ANL/MCS-P5269-0115

government, and environmental groups would like to find a solution that minimizes the competing objectives of investment, project duration, and environmental impact. Some stakeholders prefer to minimize environmental impact, some other prefer to minimize investment, some others value project duration and investment equally, and some other are indifferent. In other words, stakeholders disagree on priorities. In addition, when expressing their preferences, stakeholders are often not aware of how much a certain metric (e.g., environmental impact) should be sacrificed to improve another one (e.g., investment). Moreover, metrics are often ambiguous, in the sense that they mean different things to different stakeholders. This, in fact, is a key issue in the design of sustainability metrics [2]. Consequently, when stakeholders opinions are not systematically managed, they can leave a subset of stakeholders strongly dissatisfied. This situation can ultimately delay consensus reaching and lead to arbitrary decisions.

The most popular approach for dealing with conflicting objectives consists on computing the set of Pareto solutions (often called the Pareto front) and let an *expert* make a final decision by choosing a "suitable" Pareto solution in the set (compromise solution) [3, 4, 5]. This approach has two important disadvantages: (i) it is ambiguous by assuming that the judgement of a single expert is used to obtain the compromise and (ii) the complexity of computing the Pareto set is exponential in the number of objectives. Consequently, choosing a compromise solution can be cumbersome if not impossible, particularly when many conflict metrics must be considered. For instance, once a Pareto set is computed and the trade-offs are obtained, an expert can try to factor in the opinion of the stakeholders when picking a solution or a group of stakeholders will negotiate and try to reach consensus based on the observed trade-offs. Such negotiations, however, are often performed in non-systematic ways and they are particularly complicated when many metrics and stakeholders are involved. Consequently, it is necessary to develop decision-making frameworks that factor in the opinion of multiple experts in more systematic ways and that are capable of computing compromise solutions without explicit enumeration of trade-offs (i.e., without computing the Pareto set).

Another approach commonly used in multiobjective decision-making is to give equal priority to all objectives (i.e., weighting all objectives equally). This approach is equivalent to picking a specific point in the Pareto set and therefore is ambiguous. Moreover, the approach is also unreliable because, depending on the strength of the trade-offs, a slight modification of the weights can yield drastically different solutions [6]. In other words, this approach does not capture the shape of the Pareto set and thus might neglect solutions that yield high returns for one objective with few sacrifices for others. Another popular approach in multiobjective decision-making is to prioritize objectives, as proposed in [7]. This approach, however, also assumes that a single expert is involved in creating the priority hierarchy and consequently it is ambiguous.

Ambiguity can be mitigated by considering the opinion (judgement) of multiple stakeholders when obtaining a compromise solution. An interesting multistakeholder approach was recently presented by [8]. Here, the authors assume that stakeholders are polled to provide *priority rules* to be followed. From these rules, a unique set of weights that satisfy such rules is computed and these weights are used to obtain a compromise solution. A disadvantage of this approach is that it can yield situations in which no unique feasible weights can be obtained that satisfy all the stakeholders' rules. In addition, this approach does not provide insights into the level of dissatisfaction of the stakeholders with a given compromise decision.

In this work, we present an optimization framework that systematically quantifies and mitigates dissatisfactions among stakeholders. The idea consists of factoring the opinion of the multiple stakeholders in the form of weights (instead of rules). Consequently, compared with the approach presented in [8], the proposed framework provides more flexibility. The framework is an extension of the robust optimization approach proposed in [9] in which a compromise decision is obtained by minimizing the maximum dissatisfaction among the stakeholders. A key advantage of the robust approach is that it provides a metric to quantify stakeholder dissatisfaction. In addition, it does not require the computation of a Pareto front and can thus be used to address problems with many objectives and stakeholders. We generalize this approach by considering average and conditional-value-at-risk (CVaR) metrics. This enable us to shape the distributions of the stakeholder dissatisfactions and capture the statistics of the stakeholder population more effectively. We argue that this feature is advantageous in certain applications. In addition, generalizing the robust approach using CVaR and average metrics enable us to provide utopia-tracking interpretations of the different metrics in a common setting.

The proposed approach provides a systematic procedure to inform decision-makers about the influence of their opinions on the final decision and can help decision-makers *reassess* their priorities and thus resolve and quantify the *cost of conflict*. Examples are presented to illustrate the concepts.

2 Approach

Consider a set of objectives functions $\mathcal{O} := \{1...O\}$ and the corresponding objective function vector $\mathbf{f}(x)^T = [f_1(x), f_2(x)..., f_O(x)]^T$. Consider also a set of stakeholders $\mathcal{S} := \{1..S\}$ and that each stakeholder $s \in \mathcal{S}$ prioritizes the objectives according to the weight vector $\mathbf{w}_s \in \Re^O$. We define the elements of weight vector $\mathbf{w}_s$ as $w_{s,i},\ i \in \mathcal{O}$ and we assume that the weight vectors satisfy

$\sum_{i\in\mathcal{O}} w_{s,i} = 1,\ i \in \mathcal{O}$. Note that this definition of weight vectors assumes that a proper scaling of the objectives has taken place so that the range of all objectives $f_i(\cdot),\ i \in \mathcal{S}$ is $[0,1]$. This can be done by scaling the objectives using the coordinates of the so-called utopia point (the point at which each objective is minimized independently). For more details, we refer the reader to [10, 11].

In a sustainability context, objectives can be of social nature (human health hazard, safety hazard, jobs created), economic nature (net present value, return of investment, initial investment, budget allocations), and environmental nature (ecotoxicity, global warming potential, energy intensity, resource use) [1, 3]. Stakeholders can involve government (federal, state, and local agencies); society (communities, advocacy groups); industry (investors, managers, technology provides); and so on [9].

Our framework implicitly covers situations in which a stakeholder s wishes to consider a single objective function. In this case, the stakeholder will set one of the weights $w_{s,i}$ to one and the condition $\sum_{i\in\mathcal{O}} w_{s,i} = 1,\ i \in \mathcal{O}$ guarantees that the rest of the weights should be set to zero.

A key observation that we make is that, if the stakeholder population is finite, we can interpret the weight vectors $\mathbf{w}_s$ as samples from a probability distribution with finite support. In other words, the weight vectors can be interpreted as weight samples from the population of stakeholders. It is natural that each stakeholder $s \in \mathcal{S}$ seeks to solve its individual weighted optimization problem (based on her/his individual priorities):

$$\min_x \ \mathbf{w}_s^T \mathbf{f}(x) = \sum_{i\in\mathcal{O}} w_{s,i} f_i(x) \tag{2.1a}$$

$$\text{s.t. } g(x) \leq 0. \tag{2.1b}$$

Here, the constraint vector $g(x)$ includes operational constraints and/or system models. The solution of problem (2.1) will yield an optimal solution x_s^* and a weighted cost for stakeholder s that we denote as $\mathbf{w}_s^T \mathbf{f}_s^* := \mathbf{w}_s^T \mathbf{f}(x_s^*)$. This weighted cost is *ideal* or *utopian* in the sense that it assumes that stakeholder s does not have to compromise with the rest of the stakeholders.

When compromise is needed, as is often the case, we define the *dissatisfaction of stakeholder* s at an arbitrary compromise decision x as $d_s(x) := \mathbf{w}_s^T(\mathbf{f}(x) - \mathbf{f}_s^*)$. From optimality of x_s^* and of the associated weighted cost $\mathbf{w}_s^T \mathbf{f}_s^*$ we have that $d_s(x) \geq 0$ for all x and for all $s \in \mathcal{S}$. Consider now that two arbitrary decisions $\bar{x}, x$ yield $d_s(\bar{x}) < d_s(x)$ for a given stakeholder s. Thus, stakeholder s will be more satisfied under decision $\bar{x}$ than under decision x. Because of disagreement, however, another stakeholder s' might be less satisfied under decision $\bar{x}$ than under decision x (i.e., $d_{s'}(\bar{x}) > d_{s'}(x)$). We thus have that, given a compromise decision x, we can measure the *disagreement* among stakeholders by using a measure of the dissatisfactions $d_s(x),\ s \in \mathcal{S}$. Note that the case in which no disagreement at decision x can only occur when $d_s(x) = 0$ for all $s \in \mathcal{S}$. In the presence of disagreements among stakeholders, however, this situation cannot occur.

Our objective is to find a compromise decision x that minimizes a measure of the dissatisfactions $d_s(x), s \in \mathcal{S}$. We can think of this problem as one of shaping the distribution of the dissatisfactions. For convenience, we define the vector of dissatisfactions $\mathbf{d}(x)^T := [d_1(x), d_2(x), ..., d_S(x)]^T$.

The most straightforward alternative to managing disagreements consists of minimizing the average dissatisfaction among the stakeholders. This is done by solving the problem,

$$\min_x \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \mathbf{w}_s^T (\mathbf{f}(x) - \mathbf{f}_s^*) \tag{2.2a}$$

$$\text{s.t. } g(x) \leq 0. \tag{2.2b}$$

Note that, because $d_s(x) \geq 0$ for all $s \in \mathcal{S}$ and x, we have that problem (2.2) is also equivalent to,

$$\min_x \frac{1}{|\mathcal{S}|} \|\mathbf{d}(x)\|_1 = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} d_s(x) \tag{2.3a}$$

$$\text{s.t. } g(x) \leq 0. \tag{2.3b}$$

In other words, the solution of problem (2.2) can be interpreted as a compromise solution relative to an utopia point given by the collection of the ideal stakeholder weighted costs $\mathbf{w}_s^T \mathbf{f}_s^*$. This definition of utopia point is not to be confused with the traditional definition used in multiobjective optimization in which the utopia point is given by the minimization of individual objectives [11].

Another way to address disagreement consists of minimizing the worst (largest) dissatisfaction among the stakeholders. In other words, we find a solution under which the dissatisfaction of the most dissatisfied stakeholder is minimized. This is done by solving the robust optimization problem,

$$\min_x \max_{s \in \mathcal{S}} \left\{ \mathbf{w}_s^T (\mathbf{f}(x) - \mathbf{f}_s^*) \right\} \tag{2.4a}$$

$$\text{s.t. } g(x) \leq 0. \tag{2.4b}$$

This formulation was proposed in [9]. It is well-known that the minimax problem (2.4) can be reformulated as,

$$\min_{x} \quad \eta \tag{2.5a}$$

$$\text{s.t. } \mathbf{w}_s^T(\mathbf{f}(x) - \mathbf{f}_s^*) \leq \eta, \; s \in \mathcal{S} \tag{2.5b}$$

$$g(x) \leq 0. \tag{2.5c}$$

The optimal value of η is the worst dissatisfaction. Because $d_s(x) \geq 0$, a solution x of problem (2.4) also solves the problem,

$$\min_{x} \frac{1}{|\mathcal{S}|}\|\mathbf{d}(x)\|_\infty = \frac{1}{|\mathcal{S}|}\max_{s\in\mathcal{S}}\{d_s(x)\} \tag{2.6a}$$

$$\text{s.t. } g(x) \leq 0. \tag{2.6b}$$

Because we can assume that the stakeholders polls are obtained from a finite population, we can measure the disagreement by using a risk metric such as the conditional value at risk (CVaR) [12]. To this end we solve the following problem:

$$\min_{x} \mathrm{CVaR}_\alpha \left[\mathbf{w}_s^T(\mathbf{f}(x) - \mathbf{f}_s^*)\right] \tag{2.7a}$$

$$\text{s.t. } g(x) \leq 0. \tag{2.7b}$$

Here $\alpha \in [0, 1]$ is the probability level. This problem can be reformulated as [12],

$$\min_{x,\nu,\phi_s} \quad \frac{1}{|\mathcal{S}|}\sum_{s\in\mathcal{S}}\left(\frac{1}{1-\alpha}\phi_s + \nu\right) \tag{2.8a}$$

$$\text{s.t. } \mathbf{w}_s^T(\mathbf{f}(x) - \mathbf{f}_s^*) - \nu \leq \phi_s, \; s \in \mathcal{S} \tag{2.8b}$$

$$\phi_s \geq 0, \; s \in \mathcal{S} \tag{2.8c}$$

$$g(x) \leq 0. \tag{2.8d}$$

This approach penalizes the large dissatisfactions in the $(1-\alpha)$ tail of the distribution. In other words, for a given decision x, computing the CVaR of vector $\mathbf{d}(x)$ is equivalent to arrange the dissatisfactions $d_s(x)$ in increasing order, take the $(1-\alpha)$ largest elements (the tail), and we average them. The CVaR minimization problem thus finds the decision x under which the average of the $(1-\alpha)$ largest elements are minimized. Consequently, one can show that the CVaR solution converges to the robust solution as $\alpha \to 1$ (we take only the largest element corresponding to the most dissatisfied stakeholder) and to the average solution as $\alpha \to 0$ (we average all the elements) [13]. Consequently, the CVaR solution has the important property that it covers the spectrum of solutions between the average and robust solutions and can help us shape the distribution of dissatisfactions. This is important, as CVaR allows us to prevent the extreme conservatism of the worst-case solution and to shape the distribution of stakeholders. For instance, in some circumstances we would like to explore if a decision changes when we minimize the worst-case dissatisfaction and when we discard the $(1-\alpha)$ tail of largest dissatisfactions. If the decision does not change, it would imply that the opinion of some stakeholders does not influence the decision.

3 Illustrative Examples

In this section we present a couple of examples to demonstrate the applicability of the presented concepts.

3.1 Generation Expansion

Consider a decision-making setting in which a community (stakeholders) needs to decide among three technologies (denoted as I, II, and III) for power generation. In doing so, the community must satisfy a given demand while trading off three objectives: minimize electricity cost (denoted as C), minimize carbon emissions (denoted as E), and minimize land use (denoted as L). Table 1 lists the coefficients for cost, emissions, and land use for the three technologies.

Table 1: Emissions, cost, and land use for each technology.

Technology	E	C	L
I	100	10	100
II	50	50	50
III	75	50	25

The coefficients are dimensionless and are used only to represent relative magnitudes of different technologies. Technology I has high emissions, low cost, and high land use (relative to the others). Technology II has low emissions, high cost, and medium land use. Technology III has medium emissions, medium cost, and low land use.

The weighted multi-objective optimization problem can be formulated as form:

$$\min w_C C + w_E E + w_L L \tag{3.9a}$$

$$\text{s.t. } C = y_I C_I + y_{II} C_{II} + y_{III} C_{III} \tag{3.9b}$$

$$E = y_I E_I + y_{II} E_{II} + y_{III} E_{III} \tag{3.9c}$$

$$L = y_I L_I + y_{II} L_{II} + y_{III} L_{III} \tag{3.9d}$$

$$D = y_I P_I + y_{II} P_{II} + y_{III} P_{III} \tag{3.9e}$$

$$y_I, y_{II}, y_{III} \in \{0, 1\}. \tag{3.9f}$$

Here, y_I, y_{II}, and y_{III} denote the decisions to install technology I,II, or III, respectively. Symbol D denotes the electricity demand and P_I, P_{II}, and P_{III} denote the power supplied by each technology. For simplicity, we assume that $P_I = P_{II} = P_{III} = 10$ and we set $D = 10$. Note that the demand constraint (3.9e) implies that only one technology must be installed. The objectives (C, E, L) are all normalized by their best and worst possible values (these can be obtained from Table 1) so that their value lie in the range $[0, 1]$. For instance, we rescale objective E as,

$$E \leftarrow \frac{100 - E}{E - 50}. \tag{3.10}$$

In Table 2 we present the average and worst-case solutions under four different polls from 100 stakeholders. We assume that the polls are designed in such a way that the stakeholders express four different opinions: 1) their only priority is emissions, 2) their only priority is cost, 3) their only priority is land use, and 4) all three objectives are equally important. In a first poll we have {50%,50%,0%,0%}; in a second poll we have {49%,51%,0%,0%}, in a third poll we have {25%,25%,25%,25%}, and in a fourth poll we have {0%,0%,0%,100%}. The first poll indicates that 50% of stakeholders give full priority to minimize emissions and 50% give full priority to minimize cost. In the second poll, the number of stakeholders giving full priority to minimize cost dominates by 1% the number of stakeholders giving full priority to minimize emissions. In the third poll 25% of the stakeholders give full priority to emissions, 25% give full priority to cost, 25% give full priority to land use, and 25% give equal priority to all objectives. The fourth poll correspond to the special case in which all stakeholders give equal priority to minimize all objectives. In other words, in the fourth poll we have *perfect agreement among stakeholders.*

From the first three polls we can see that the robust strategy achieves the same worst-case dissatisfaction for all technologies. In other words, the three technologies are optimal regardless of the polls. Under perfect agreement (fourth poll), on the other hand, technology II is optimal and the worst-case dissatisfaction is zero. From the first poll, we can see that the average strategy predicts that technologies I and II are equally optimal. This is expected because we have the same number of stakeholders giving priority to emissions and cost; consequently, the solutions are indistinguishable. For the second poll, technology II is optimal because the number of stakeholders giving priority to cost is larger (by one vote) than those giving priority to emissions. For the third poll we have the less obvious result that technology II is optimal.

By comparing the results for the robust and average strategies we can obtain some important insights. First note that in the fourth poll with perfect agreement the robust and the average solutions are equal, as expected. From the rest of the polls we can see that the robust solution is insensitive to the statistics of the polls. While this feature might seem desirable at a first sight, it is not likely to be accepted by stakeholders because it implies that their opinions do not influence the solution (even if there is a majority of stakeholders). In fact, in this example, even a poll with a distribution of $\{1\%, 99\%, 0\%, 0\%\}$ will give the same robust solution. The average strategy, on the other hand, does account for the statistics of the stakeholder polls but it cannot guarantee minimization of the worst dissatisfaction, as the robust strategy does.

Table 2: Compromise solutions under different polls.

Poll	Strategy	Compromise Solution
{50%,50%,0%,0%}	Average	y_I, y_{II}
	Robust	y_I, y_{II}, y_{III}
{49%,51%,0%,0%}	Average	y_I
	Robust	y_I, y_{II}, y_{III}
{25%,25%,25%,25%}	Average	y_{II}
	Robust	y_I, y_{II}, y_{III}
{0%,0%,0%,100%}	Average	y_{II}
	Robust	y_{II}

3.2 Energy-Comfort Management in Buildings

One of the objectives of an energy management system is to minimize energy subject to thermal comfort constraints of a population of occupants (stakeholders) [6, 14]. Thermal comfort is difficult to enforce because perceptions vary significantly from individual to individual as a result of variations in factors such as metabolic rate (e.g., activity, gender, race, age), building location (e.g., next to air damper, next to window), and clothing level. To address this disagreement, we can poll the opinion of occupants about their temperature preferences.

Consider thus the following stakeholder problem in which we seek to minimize energy demand while satisfying the stakeholder j temperature constraint:

$$\min E(T) \tag{3.11a}$$

$$\text{s.t. } T \leq T_s,\ (\lambda_s). \tag{3.11b}$$

Here, $E(\cdot)$ is the building energy that is a function of occupant's s temperature requirement T_s and λ_s is the Lagrnage multiplier of the comfort constraint (3.11b). For simplicity, we assume that energy is a quadratic function of the difference between the building temperature and the ambient temperature,

$$E(T) := (T - T_{amb})^2. \tag{3.12}$$

We set the ambient temperature to 35 ^oC and we assume that $T_s \leq T_{amb}$. Clearly, $\lambda_s = \frac{\partial E}{\partial T_s}$ and $\lambda_s < 0$ if the objective and the comfort constraint are in conflict (i.e., energy increases as we decrease the temperature requirement T_s). In other words, an occupant with a lower temperature requirement requires more cooling energy. Consequently, λ_s can be interpreted as a comfort price. We can thus formulate the weighted multi-objective problem,

$$\min w_{s,1}E(T) + w_{s,2}T, \tag{3.13}$$

with $w_{s,1} := \frac{1}{1-\lambda_s}$ and $w_{s,2} := \frac{-\lambda_s}{1-\lambda_s}$. This problem is equivalent to (3.11).

We consider a poll of temperature preferences for $S = 1,000$ occupants in a building. The average preference is 22 ^oC, the minimum temperature preference is 15 ^oC, and the maximum preference is 29 ^oC. All temperatures are below ambient temperature and we thus simulate a situation in which energy is used for cooling. Because of this, if comfort is not a concern, the energy required will be zero and the building will be set to ambient temperature. This also implies that, as T_s is increased, the comfort price λ_s will decrease and will be zero at T_{amb}. We thus have that the weight $w_{s,2}$ will tend to zero and $w_{s,1}$ will tend to one, reflecting the fact that a larger T_s implies a lower priority on comfort.

Consequently, polling temperature preferences can be interpreted as polling priorities of energy and comfort among occupants.

In Figure 1 we present the corresponding comfort weights w_T ($w_{s,2}$) for the occupants while in Figure 2 we present the Pareto curve of temperature against energy demand. Each point along the front is obtained by solving problem (3.11) for each stakeholder s. Note that these points represent the ideal (non-achievable) situation in which each stakeholder can reach their desired temperature preference without having to compromise with the rest of the stakeholders. In Figure 2 we also present the solution of different compromise decisions. The vertical line indicate the solution in which the stakeholders compromise naively by averaging their temperature preferences (average is 22 ^oC). This solution corresponds to solving the energy minimization problem,

$$\min E(T) \tag{3.14a}$$

$$\text{s.t. } T \geq \frac{1}{|\mathcal{S}|}\sum_{s\in\mathcal{S}} T_s. \tag{3.14b}$$

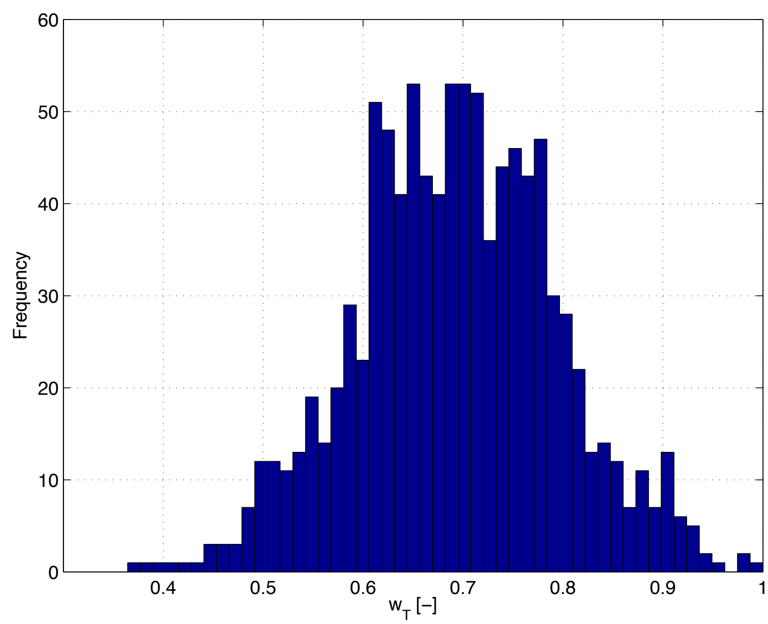

Figure 1: Occupants weights for temperature-energy trade-off.

Note that this *naive* approach does not capture energy in the stakeholders opinions. The black dot next to the naive solution represents the compromise solution obtained by minimizing the average dissatisfaction given by (2.2). The solutions obtained by averaging preferences and minimizing the

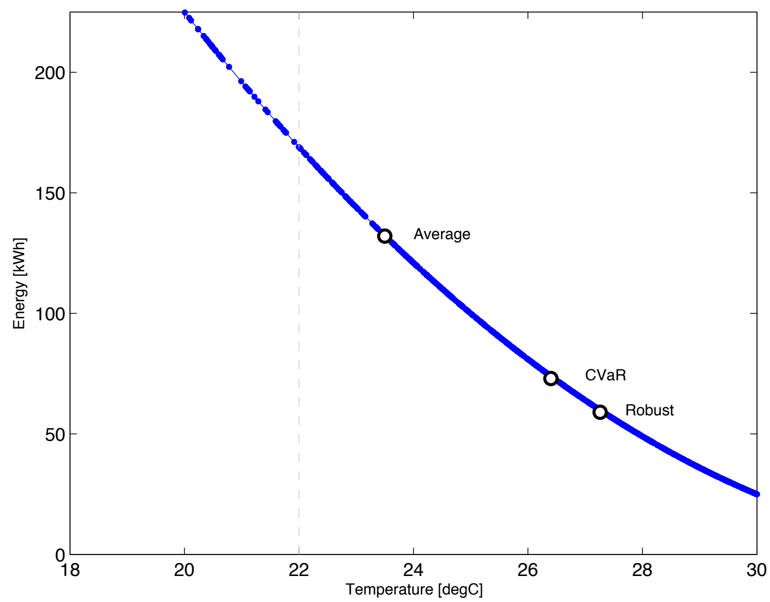

Figure 2: Pareto front and compromise decisions. Vertical line is naive approach.

average dissatisfaction do not coincide. The reason is that averaging temperatures is not equivalent to averaging dissatisfactions (dissatisfactions factor in energy and not only temperature).

The compromise solutions located on the right-most end of the Pareto front represent are those obtained by minimizing CVaR for $\alpha = 95\%$ given by (2.7); and the decision obtained by minimizing the worst dissatisfaction among the stakeholders given by (2.4). Note that the compromise decisions move toward the maximum temperature preference (warmer building) as robustness is increased. Consequently, the CVaR and robust approaches yield much lower energy demands than do the naive and average approaches. From the naive approach perspective (without factoring in energy in the opinions) this result is counterintuitive because one would expect that at a higher temperature more people would be dissatisfied. From a dissatisfaction perspective as defined, however, a colder building yields much larger dissatisfactions because some occupants actually care about energy. This illustrates how a more systematic management of stakeholder opinions can yield more efficient (and nonintuitive) solutions.

In Figure 3 we present histograms for the dissatisfactions of all the stakeholders. In the top graph we present the dissatisfactions when the stakeholders compromise by minimizing the average dissatisfaction. Note the pronounced tail of large dissatisfactions. In the middle graph we present the dissatisfactions when the stakeholders compromise by minimizing CVaR and in the lower graph we present the dissatisfactions when the stakeholders compromise by minimizing the worst dissatisfaction. Note that the tail of the distribution of dissatisfactions is reduced by CVaR and the robust approach reduces the tail further by penalizing the largest dissatisfaction. We can thus see that, in this application, the robust approach provides important benefits compared to the average approach. Moreover, the CVaR approach provides a mechanism to relax the worst-case solution.

Arguably, the need to choose among average, CVaR, and worst case metrics introduces some ambiguity in the decision-making process in the sense that the stakeholders must also agree that such a metric is the appropriate one. The level of ambiguity, however, is significantly reduced compared with standard multi-objective approaches that assume a single decision-maker picking an arbitrary point on the Pareto front. Moreover, the proposed approach has the additional advantage that it can manage many opinions and objective functions in a systematic manner.

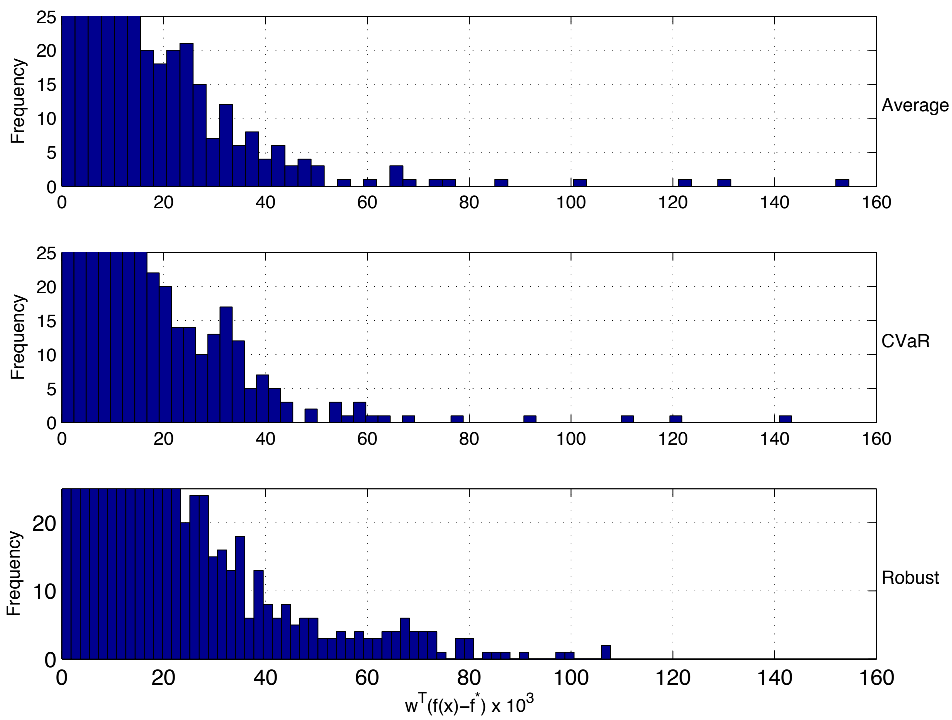

Figure 3: Dissatisfaction of stakeholders under different compromise decisions.

4 Conclusions and Future Work

We have presented a framework to manage conflicts among multiple decision-makers. The framework enables the computation of compromise solutions in the presence of many objectives and stakeholders preferences without having to compute the Pareto set. The framework also provides a systematic procedure to manage conflicts by using quantifiable metrics of disagreement among stakeholders.

We highlight that the framework proposed can manage objective functions in either deterministic or stochastic settings. For instance, one can trade-off mean profit and profit variance. Our framework, however, does not account for situations in which stakeholders can change their preferences based on possible scenarios, as discussed in [9]. We will extend our framework to consider this possibility in future work. It is also necessary to consider formulations under which the stakeholders not only provide their preferences in terms of weights but also in terms of goals. This will give rise to interesting goal-oriented multi-objective formulations. We are also interested in understanding under what conditions the CVaR compromise solution gives a Pareto solution. This is motivated from the fact that the average and robust metrics have utopia-tracking interpretations (under different norms). We will look for a similar definition in the CVaR case.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357. The author thanks Sanjay Mehrotra for suggesting the building energy management example.

References

[1] Ruiz-Mercado GJ, Smith RL, Gonzalez MA. Sustainability indicators for chemical processes: I. Taxonomy. Industrial & Engineering Chemistry Research. 2012;51(5):2309–2328.

[2] Sikdar SK. Sustainable development and sustainability metrics. AIChE journal. 2003;49(8):1928–1932.

[3] You F, Tao L, Graziano DJ, Snyder SW. Optimal design of sustainable cellulosic biofuel supply chains: multiobjective optimization coupled with life cycle assessment and input–output analysis. AIChE Journal. 2012;58(4):1157–1180.

[4] El-Halwagi AM, Rosas C, Ponce-Ortega JM, Jiménez-Gutiérrez A, Mannan MS, El-Halwagi MM. Multiobjective optimization of biorefineries with economic and safety objectives. AIChE Journal. 2013;59(7):2427–2434.

[5] Fu Y, Diwekar UM, Young D, Cabezas H. Process design for the environment: A multi-objective framework under uncertainty. Clean Products and Processes. 2000;2(2):92–107.

[6] Zavala VM. Real-time resolution of conflicting objectives in building energy management: an utopia-tracking approach. In: In proceedings of the 5th national conference of IBPSA-USA; 2012. p. 1–6.

[7] Mendoza GA, Prabhu R. Multiple criteria decision making approaches to assessing forest sustainability using criteria and indicators: a case study. Forest Ecology and Management. 2000;131(1):107–126.

[8] Smith RL, Ruiz-Mercado GJ. A method for decision making using sustainability indicators. Clean Technologies and Environmental Policy. 2014;16(4):749–755.

[9] Hu J, Mehrotra S. Robust and stochastically weighted multiobjective optimization models and reformulations. Operations research. 2012;60(4):936–953.

[10] Miettinen K. Nonlinear multiobjective optimization. vol. 12. Springer; 1999.

[11] Zavala VM, Flores-Tlacuahuac A. Stability of multiobjective predictive control: A utopia-tracking approach. Automatica. 2012;48(10):2627–2632.

[12] Rockafellar RT, Uryasev S. Optimization of conditional value-at-risk. Journal of risk. 2000;2:21–42.

[13] Rockafellar RT, Uryasev S. Conditional value-at-risk for general loss distributions. Journal of Banking & Finance. 2002;26(7):1443–1471.

[14] Morales-Valdez P, Flores-Tlacuahuac A, Zavala VM. Analyzing the effects of comfort relaxation on energy demand flexibility of buildings: A multiobjective optimization approach. Energy and Buildings. 2014;85(0):416 – 426.