Group Decision and Negotiation **14:** 441–460, 2005.
DOI: 10.1007/s10726-005-9002-x

The Graph Model for Conflict Resolution: Past, Present, and Future

D. MARC KILGOUR
Department of Mathematics, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
(E-mail: mkilgour@wlu.ca)

KEITH W. HIPEL
Department of Systems Design Engineering, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada
(E-mail: kwhipel@uwaterloo.ca)

Abstract

The Graph Model for Conflict Resolution is a methodology for the modeling and analysis of strategic conflicts. An historical overview of the graph model is presented, including the basic modeling and analysis components of the methodology, the decision support system GMCR II that is now used to apply it, and the recent initiatives that are currently in various stages of development. The capacity of this simple, flexible system to provide advice to decision-makers facing strategic conflicts is emphasized throughout, and illustrated using a real-life groundwater contamination dispute.

Key words: Graph Model for Conflict Resolution, GMCR II, strategic conflict, stability, equilibrium, coalition analysis, status quo analysis

1. Motivation

There were two major reasons for developing the methodology of the Graph Model for Conflict Resolution, and later its associated decision support system GMCR II. First, there was a perceived demand for a comprehensive methodology to understand conflict decision-making and conflict resolution. The ubiquity of conflicts means that support is needed not only by decision-makers, but also by mediators – who propose resolutions – and policy-makers who determine the structures within which conflicts are played out. The inescapability of conflict implies that the need for the graph model will continue for as long as humans interact.

The second motivation was a view that existing methods failed to provide the kind of analysis and advice that was needed, either because they were too inflexible to be used in most conflict situations, or because they demanded so much information and calibration that they became impractical. The graph model was designed to be simple and flexible, and to have minimal information requirements. At the same time, the suggestions that it makes encourage disputants to 'think outside the box."

This paper summarizes the development of the Graph Model, which began in the early 1980s and continues to the present day. The approach is roughly historical: after a discussion of strategic conflicts (Section 2), the basics of the Graph Model are set out in Section 3 (Past), the characteristics of the Decision Support System GMCR II are described in Section 4 (Present), and various new initiatives and their status with respect to existing software are outlined in Section 5 (Future). Section 6 offers a summary and a few conclusions.

2. What is a Strategic Conflict?

A *strategic conflict is* an interaction of two or more independent decision-makers (DMs), each of whom makes choices that together determine how the state of the conflict evolves, and each of whom has preferences over these possible states (as eventual resolutions). Thus, a strategic conflict is a joint, or interactive, decision problem; there are two or more DMs, each DM has a choice (i.e. two or more alternatives), and every DM is in principle concerned about the others' choices. More specifically, each DM must benefit, or be harmed, according to the choices of at least one other DM, in the sense that that other DM's choices make the eventual resolution more, or less, preferable. It is clear that strategic conflicts are very common in interactions at all levels including personal, family, business, national, and international.

One way to model and analyze a strategic conflict is to use *non-cooperative game theory* (von Neumann and Morgenstern, 1944). A game structure permits the analyst to capitalize on a large and well-developed body of theory, which has established links with economics and Bayesian decision analysis. But to use a non-cooperative game model to analyze a strategic conflict and provide strategic advice imposes constraints which may limit the verisimilitude of the model and the usefulness of the advice. For instance, in a game the order of action of the DMs (called players) must be specified but, in many strategic conflicts such as negotiations, the order of action is not known in advance – deciding when to act is part of the problem. Another requirement is that in a game players' preferences must be represented by real-valued (von Neumann-Morgenstern) utilities, which open up the possibility of mixed strategies (probabilistic mixtures of actions, as opposed to specific actions). But this requirement is a serious drawback for two reasons: utilities are notoriously difficult to measure; and mixed strategies are often hard to interpret as "advice." (Would you really tell your president to toss a coin to decide whether to attack or press for peace?)

The Graph Model for Conflict Resolution provides a methodology for modeling and analyzing strategic conflicts that does not suffer from these problems. It is easy-to-use, flexible, and provides a good understanding of how DMs should choose what to do. Of course, there are alternative systems to model and analyze strategic conflicts that are distinct from non-cooperative game theory; they include metagame analysis (Howard, 1971), conflict analysis (Fraser and Hipel, 1984), drama theory (Howard, 1999), theory of moves (Brams, 1994), and theory of fuzzy moves (Kandel et al., 1998; Li et al., 2001). For a broader view of related approaches and results, see the Encyclopedia section introduced by Hipel (2002). The specific focus of this paper is the Graph Model for Conflict Resolution, which we believe is more flexible, broader in scope, and easier to use than the alternatives.

The original formulation of the Graph Model for Conflict Resolution appeared in Kilgour et al. (1987); the first complete presentation is the text of Fang et al. (1993). It has been applied across a wide range of application areas; examples include environmental management at the local level (Kilgour et al., 2001; Noakes et al., 2003; Hamouda et al., 2004a,b; Li et al., 2006) and the international level (Noakes et al., 2006; Obeidi et al., 2002); labor-management negotiation (Fang et al., 1993, Section 8.5); military and peace-keeping activities (Kilgour et al., 1998); and international negotiations on economic issues (Hipel et al., 2001) and arms control (Obeidi et al., 2006). A complete list of publications is maintained on the website http://www.systems.uwaterloo.ca/Research/CAG/.

3. Past: Basics of the Graph Model for Conflict Resolution

3.1. Graph model definitions

The Graph Model for Conflict Resolution is described in full in Fang et al. (1993) and is summarized here. A Graph Model has four components, as follows:

- $\mathbf{N}$, the set of decision-makers (DMs), where $2 \leq |\mathbf{N}| < \infty$. We write $\mathbf{N} = \{1, 2, \ldots, n\}$.
- $\mathbf{S}$, the set of (distinguishable) states, satisfying $2 \leq |\mathbf{S}| < \infty$. One particular state, s_0, is designated as the *status quo* state.
- For each $i \in \mathbf{N}$, DM i's directed graph $G_i = (\mathbf{S}, A_i)$. The arc set $A_i \subseteq \mathbf{S} \times \mathbf{S}$ has the property that if $(s, t) \in A_i$ then $s \neq t$; in other words, G_i contains no loops. The entries of A_i are the *state transitions* controlled by DM i.
- For each $i \in \mathbf{N}$, a complete binary relation $\succeq_i$ on $\mathbf{S}$ that specifies DM i's *preference over* $\mathbf{S}$. If $s, t \in \mathbf{S}$, then $s \succeq_i t$ means that DM i prefers s to t, or is indifferent between s and t. Following well-established conventions, we say that i *strictly prefers* s to t, written $s \succ_i t$, if and only if $s \succeq_i t$ but $\neg[t \succeq_i s]$ (i.e. it is not the case that $t \succeq_i s$). Also, we say that i is *indifferent* between s and t, written $s \sim_i t$, if and only if $s \succeq_i t$ and $t \succeq_i s$.

The arcs in a DM's graph represent state transitions controlled by the DM; specifically, if $s, t \in \mathbf{S}$ and $s \neq t$, then there is an arc from s to t in DM i's graph, i.e. $(s, t) \in A_i$, if and only if DM i can (unilaterally) force the conflict to change from state s to state t. In this case, we say that t is *reachable* for i from s. Note that all DMs' graphs have the same vertex set, $\mathbf{S}$. A consequence is that relatively small Graph Models can be conveniently described using the *integrated graph* $G = (\mathbf{S}, A_1, A_2, \ldots, A_n))$. Note that the integrated graph is a directed graph (possibly with multiple arcs), in which each arc is labelled with the name of the DM who controls it.

In principle, the Graph Model methodology does not require preference or indifference relations to be transitive. (For example, $\succeq_i$ is *transitive* if, whenever $s_1 \succeq_i s_2$ and $s_2 \succeq_i s_3$, then $s_1 \succeq_i s_3$ also.) Typically when participants begin to think about a dispute, confusion and lack of information may produce intransitive preferences. But intransitive preferences usually disappear over time. If preferences are transitive, then each DM's preference can be used to order the state set $\mathbf{S}$. In other words, for each DM there is a ranking of all states from most preferred to least preferred, possibly including ties as groups of equally preferred states. The assumption of ordinal preferences makes the presentation of a graph model using the integrated graph particularly compact. The decision support system GMCR II assumes that all preferences are transitive.

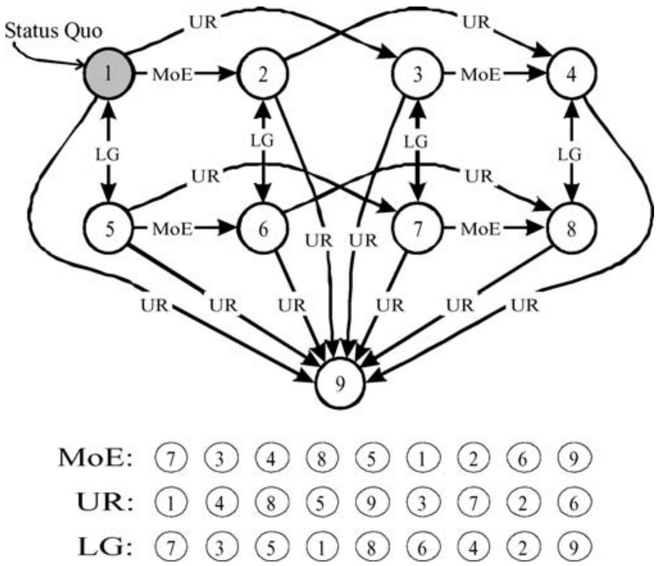

Figure 1. Elmiral graph model (Kilgour et al., 2001, Figure 3).

Figure 1 shows a complete Graph Model; it displays both the integrated graph and the preference orderings. In this particular example, all preferences happen to be strict, but in general this need not be the case.

The Graph Model in Figure 1 is a simple but very useful model of a strategic conflict studied by the authors and their collaborators. Figure 1 is a graph model of a situation that arose in 1991 when a carcinogen was discovered in the underground aquifer from which the town of Elmira, Ontario drew all of its water. The three DMs are the Ontario Ministry of the Environment (MoE), Uniroyal Chemical Limited (UR), and the Local Governments (LG). The strategic conflict centers on responsibility for clean-up of the pollution; at the time point of the model, the Ministry has just issued a control order requiring Uniroyal to clean up the pollution, but Uniroyal has the right to appeal. In the Elmiral model, MoE is considering modifying the control order to make it more acceptable to UR (an option called *Modify*); UR is deciding whether to delay the process by appealing (*Delay*), accept the current version of the control order (*Accept*), or abandon its Elmira facility (*Abandon*); and Local Governments have not yet decided whether to support the Ministry's control order (*Support*). All told, there are nine distinct states in the model.

It is a reflection of the simplicity of the Graph Model that the states form the basis for all the definitions and all of the analysis. States are depicted as circles in Figure 1. The current state of a Graph Model is assumed to be known to all DMs at all times, beginning with the status quo state. At the status quo state (shown as state 1 in the model of Figure 1), MoE is refusing to modify its control order, UR is delaying, and LG has not yet taken a position.

If the current state of a Graph Model is s, then DM i may choose to change the state to any $t \in \mathbf{S}$ that is reachable for i from s (i.e., such that $(s, t) \in A_i$) if any such t exists; but DM i may also choose not to change the state. In Figure 1, for example, all three DMs can move away from the status quo, while no DM can move away from state 9 (which represents the consequences of UR's choice to Abandon). If there are two different DMs who can move the conflict away from *s*, there is no assumption about which one has priority. In most examples, each DM's graph turns out to be transitive (if DM *i* can move from s_1 to s_2 and from s_2 to s_3, then DM *i* can also move – in one step – from s_1 to s_3) but this property is not assumed in the Graph Model methodology.

DM *i*'s preference over **S** represents *i*'s preferences among the states of **S** considered as final outcomes, or resolutions, of the conflict. Thus, in the Elmiral model (Figure 1), Uniroyal most prefers the status quo, state 1, whereas both MoE and LG most prefer state 7, where LG supports MoE's control order and UR accepts it. Note also that state 9, at which UR abandons its Elmira facility, is the least preferred outcome for both MoE and LG.

3.2. Stability analysis of a graph model

Now that we have described how a strategic conflict is modeled using the Graph Model for Conflict Resolution, we turn to the second function of the methodology—analysis. First, we give some definitions. From any state, $s \in \mathbf{S}$, a state that is reachable by DM *i* from s and that DM *i* prefers to *s* is called a (*unilateral*) *improvement* for *i* from *s*, and a state that is reachable by *i* from *s* but is less preferred by *i* than *s* is called a *(unilateral) disimprovement*. For example, in Figure 1, a move by LG from the status quo, state 1, to state 5 is a unilateral improvement, whereas a move by UR from state 1 to state 3 is a unilateral disimprovement.

In the Graph Model for Conflict Resolution, a *stability definition* (or *solution concept*) is a set of rules for calculating whether a decision-maker would prefer to stay at a state or move away from it unilaterally. A stability definition is therefore a model of a DM's strategic approach, or more generally of human behavior in strategic conflict. Of course, different stability definitions may be appropriate for different DMs.

A general principle for stability definitions in a Graph Model with $n = 2$ DMs is that specifying a state, *s*, a DM, *i*, and a particular stability definition is equivalent to specifying a two-person finite extensive-form game of perfect information with a particular structure. In this game, the first move must be a choice by DM *i* to stay at *s* or to move to any of the states reachable for *i* from *s*. If *i* chooses to stay at *s*, the game is over and the outcome is s. If *i* does not stay on the initial move, then there may be additional choices by other DMs (and possibly by *i* again), but at all subsequent decision nodes one alternative is always to stay at the current state, and selecting this alternative always ends the game at that state. Stability definitions differ only in the construction of this auxiliary extensive-form game. For Graph Models with $n > 2$ DMs, stability definitions are generalized in a natural way from the $n = 2$ case.

An *equilibrium* is a state that is stable, according to an appropriate definition, for every DM in a Graph Model. The equilibria are the predicted resolutions of the strategic conflict.

The main stability definitions currently used in Graph Model analysis include Nash Stability (Nash), General Metarationality (GMR), Symmetric Metarationality (SMR), Sequential Stability (SEQ), Limited Move Stability (L_h), and Non-Myopic Stability (NM).

Table 1. Main stability definitions used in the graph model.

Definition	Foresight	Disimprovements	How does focal DM (i) anticipate that other DMs will respond to *i*'s improvement?
Nash	1	Never	None
GMR	2	Sanctions only	Will sanction i's improvement at any cost
SMR	3	Sanctions only	Will sanction i's improvement at any cost
SEQ	2	Never	Will sanction i's improvement, but only using their own improvements
L_h	$h \geq l$	Strategic	Symmetric; others optimize for themselves, just like *i*
NM	∞	Strategic	Symmetric; others optimize for themselves, just like *i*

For complete definitions and original references, see Fang et al. (1993, Ch.3). (In fact, these are the definitions that are implemented in the software GMCR II.) Table 1 describes some features of these definitions that relate them to behavior in conflicts. Foresight, for example, refers to the maximum number of moves foreseen by a DM whose stability calculation follows a particular definition. Nash stability looks one move ahead; the conservative definitions (GMR, SMR, and SEQ) look two or three moves ahead; in L_h-stability, the horizon of foresight is a parameter, *h,* which may equal any positive integer; a state exhibits NM stability if and only if it is L_h-stable for all sufficiently large *h*. Stability definitions also differ with respect to disimprovements: in Nash stability there are none for the focal DM, and for opponents it is not an issue; in GMR and SMR, there are none for the focal DM, but sanctions by other DMs may be disimprovements; in SEQ disimprovements are forbidden for either the focal DM or the opponents; and in L_h $(h > 1)$ and NM, disimprovements are permitted provided they are *strategic,* that is, anticipated to induce other DMs to react in a way that benefits the DM making the move.

The logical relationships among the stability definitions in Table 1 are described in detail in Fang et al. (1993, Ch.5). For instance, a state that is Nash is also GMR, SMR, and SEQ, and a state with any other form of stability must also be GMR. Different stability definitions have features that make them more, or less, appropriate to describe certain DMs. Some of these features are suggested in Table 1: for instance, GMR and SMR describe conservative DMs, who expect to be sanctioned if it is possible for the opponents to do so, no matter how much the opponents may harm themselves in the process. A DM who follows SEQ is almost as conservative; he or she never disimproves, and expects to be sanctioned by the opponents, but only if they can do so without disimproving. By contrast, DMs who follow L_h are calculating and strategic, and see every DM as attempting to optimize – subject, of course, to limitations of foresight. The NM stability definition supposes the ultimate in strategy, but sometimes it is so demanding that no state satisfies it.

Even though a Graph Model is simple and easy to construct, and the Graph Model stability definitions have straightforward characterizations to make them easy to calculate, the computational burden involved in finding states with appropriate forms of stability for every DM – even in a small model – was quickly found to be daunting. For this reason, the software system GMCR (now called GMCR I) was developed. (GMCR I is described by Fang et al., 1993, Appendices A and B). GMCR I is an analysis engine that calculates the stability of every state in a Graph Model, from the point of view of every DM in the model, according to all of the stability definitions listed in Table 1.

The use of GMCR I fostered a philosophical change in the analysis of Graph Models. Instead of first assigning a stability type to each DM and then identifying states that are stable for each DM according to the appropriate definition, it was easy to find all states stable for each DM under a range of definitions, and then to focus on those states with at least some form of stability for every DM. Then only those states that are stable under one or more definitions that might be appropriate to the DM are retained. For example, one looks at Nash, GMR, and SMR for a DM who is cautious and may lack knowledge of other DMs' preferences. If the DM is more confident of other DMs' views, SEQ may be included also. The Limited Move definitions are appropriate for DMs who are more far-sighted and strategic, and who are confident of their knowledge of other DMs' points of view. This approach brings additional information to the analyst, and was quickly discovered to encourage better modeling and deeper analysis.

In practice, it turns out, most Graph Models have a few states that are stable under most definitions (they are now called *strongly stable*), some that are stable under only one or two definitions, and some that are always unstable. Usually there are a few states that are strongly stable for all DMs; almost always, these states are the equilibria that are easiest to interpret as stable resolutions of the underlying strategic conflict.

For example, in the Elmiral model of Figure 1, states 5, 8, and 9 are stable for all DMs under the definitions Nash, GMR, SMR, SEQ, $L(h)$ for $h = 2, 3, \ldots$, and NM. States 1 and 4 are stable for all DMs, but for LG they are stable only under the short-sighted, low-knowledge definitions GMR and SMR. Thus, analysis of the Elmiral model suggests the conflict is likely to end up at either state 5 (similar to the status quo, except that LG supports the control order), state 8 (a compromise in which, despite LG's support MoE modifies the control order and UR accepts it), or state 9 (in which UR abandons the Elmira facility). It should be noted (see Figure 1) that state 9 must be stable in this model, since no DM can move away from it.

4. Present: The Decision Support System GMCR II

4.1. GMCR II's basic structure

The availability of software that analyzed Graph Models quickly, completely, and reliably resulted in an increase in the number and range of applications of the Graph Model methodology, which in turn provided convincing evidence of the utility of the approach. But the need to justify these models and interpret the results of the analysis created the need to analyze even more Graph Models, typically related to the initial models but distinct from them. For example, one might ask of the Elmiral model whether the assumed preferences of Uniroyal – and in particular that state 9 is fifth in UR's preference ordering – affect the conclusions about stability, or whether the DM Local Government makes any difference to the outcome. The natural way to answer such questions is to modify the original model, re-analyze, compare the results, and assess.

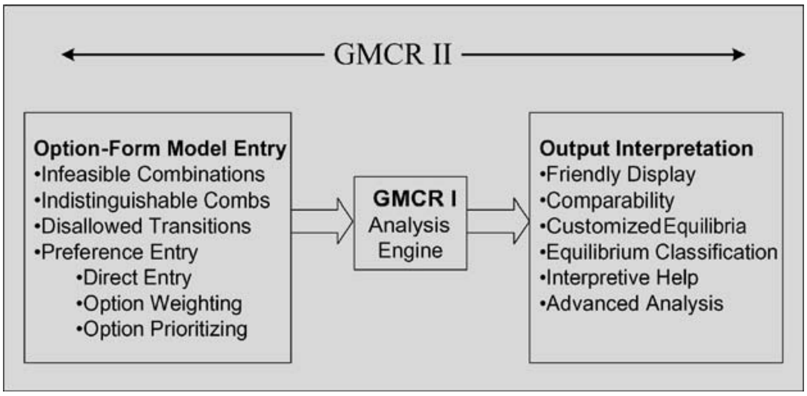

Figure 2. Components of decision support system GMCR II.

Attempts to analyze models derived from an initial model highlighted several bottlenecks in the use of GMCR I. Detailed, time-consuming calculations were usually required to convert a model to the GMCR I data format. Moreover, this format was so opaque that it was difficult to convert the data for one model to another that was conceptually very close. Another problem was evident in output interpretation; the output of GMCR I was not organized to facilitate efficient, in-depth understanding of the analysis results.

The Decision Support System GMCR II was designed to ease these and other problems. (GMCR II is described in detail by Fang et al., 2003 a, b and Hipel et al., 1997.) As indicated in the schematic description in Figure 2, GMCR II is a three-stage system of which the earlier system, GMCR I, forms the second stage. The major advances incorporated into GMCR II are at the initial stage, Model Entry, and the final stage, Output Display. In fact, improvements were also made to the GMCR I Analysis Engine, primarily to enable it to analyze larger models faster.

4.2. Option-form entry of graph models

Option-form entry was an important advance, allowing natural Graph Models for most strategic conflicts to be entered quickly and conveniently, and making it simple and straightforward to adjust existing models. A basic version of option form was developed for metagame analysis (Howard, 1971) and used later in conflict analysis (Fraser and Hipel, 1984) and drama theory (Howard, 1999). However, the additional flexibility of the Graph Model required further developments and adaptations.

Option-form entry avoids direct specification of the states of a Graph Model by listing, for each DM, i, a non-empty finite set O_i representing the *options*, or courses of action, available to i. An option can belong to one and only one DM, so $O = O_1 \cup O_2 \cup \cdots \cup O_n$ represents the set of all options in the model. The default assumption is that a DM can select any subset of its options (including the empty subset); under this assumption, a state is simply a subset of O, which is called an *option combination*, and the set of all states is $\mathbf{S} = 2^O$.

Usually, however, it is a difficult task to specify options in such a way that (1) every option combination is feasible and (2) every option combination represents a distinct state. Moreover, it has been found to be more efficient in practice to give the analyst or modeler free rein to list options without restriction, eliciting the details of any restrictions later on. In entry of the original Elmiral model, for instance, the domain expert specified three options for Uniroyal: Delay, Accept, and Abandon. But obviously these options are not independent of each other (for example, it would be impossible to choose more than one of Delay, Accept, and Abandon), so some option combinations are infeasible.

In GMCR II, two important steps immediately follow specification of the options. First, infeasible option combinations are removed from the model. Second, option combinations that are essentially equivalent are coalesced. For instance, the latter procedure was also used in the Elmiral model; the domain expert felt that all option combinations that included Uniroyal's Abandon option were essentially the same.

GMCR II incorporates flexible procedures that can remove all infeasible option combinations and coalesce all equivalent option combinations in any option form. As well, it includes short-cut procedures that are sufficient for most practical examples. For instance, infeasibilities often occur because options – like Uniroyal's, as discussed above – are mutually exclusive; in many Graph Models, the GMCR II procedure to specify mutually exclusive options is sufficient. Specifying all infeasible option combinations and coalescing all equivalent option combinations determines all states of the model.

Option-form entry assumes by default that the state transitions controlled by a DM correspond to changes in the DM's options, and that any unilateral change of options is allowed. After the states are determined, the GMCR II user is directed to a procedure to specify any disallowed transitions. Again, the most flexible procedures can disallow any specific transition, while short-cut procedures are often sufficient in practice. In the Elmiral model, for instance, a short-cut procedure was used to specify options (like MoE's decision to Modify and UR's decisions to Accept or to Abandon) that, once selected, could not be de-selected.

The remaining Graph Model component is preference – specifically, an ordering of $\mathbf{S}$ for each DM. GMCR II's basic method of preference entry, called Direct Entry, asks the user to rearrange ('drag and drop') states so that they are listed in preference order. Any number of adjacent states can be coded as tied in preference if appropriate. Direct Entry is flexible in that any (transitive) preference can be entered, but it is slow, and for most models it is cumbersome, as screen size is quickly exceeded. On the other hand, if an ordering is approximately correct, then usually it can be quickly adjusted using Direct Entry, which is often called Fine Tuning when used for this purpose.

Because states are entered in option form, they have a structure (each state corresponds to an option combination) that can be used to approximate a DM's preference ordering very quickly. It is often easiest to take advantage of this structure in a brain-storming session. It is recommended that, for all but the smallest models, one of the two more sophisticated GMCR II preference entry procedures be used first to obtain an approximate ranking, which then should be adjusted using Fine Tuning.

Another GMCR II preference-entry procedure is Option Weighting. A numerical weight (positive or negative) is assigned to each option. Then the score of each state is calculated as the sum of the weights of the options selected when the state is represented as an option combination. Finally, states are ordered according to score. Option Weighting is a very simple procedure that can come surprisingly close to many preference orderings encountered in practice.

The most sophisticated of GMCR II's preference-entry procedures is Option Prioritizing. The user enters, in priority order, a sequence of *preference statements*, which are logical statements about options. Each statement must be a true-or-false statement about the options selected at a state, and may contain logical connectives such as 'and,' 'or,' 'not,' 'if,' or 'iff.' Typical statements are "option 3 is selected," "option 4 is not selected," and "both option 3 and option 4 are selected." For state $s \in \mathbf{S}$, each of the statements in the hierarchy is true or false. GMCR II orders the states so that state s precedes state t if and only if the highest priority statement that is true for exactly one of s and t is true for s and false for t. While Option Prioritizing is harder to learn than Option Weighting, most users report that they appreciate the additional flexibility inherent in the procedure. The priority hierarchy of preference statements seems to reflect a natural way of describing the derivation or justification of preferences.

Table 2 shows the use of Option Prioritization to describe the MoE's preferences in the Elmiral model. For easy reference, the left-hand column contains the DMs and options. The middle column lists preference statements in priority order (from most important to least important). This column is exactly what would be entered into GMCR II's Option Prioritization dialog box. The right-hand column of Table 2 interprets each preference statement. Notice that MoE most prefers that UR not abandon its Elmira plant (the negative sign means 'not'). Next most important to MoE is that UR accept the current control order. GMCR II uses these five preference statements to produce the ordering of states (for MoE) shown in Figure 1, where the numbers here refer to the state numbers. In this case, no fine tuning is necessary. Clients find Option Prioritization to be an extremely attractive feature of GMCR II.

Table 2. Option Prioritization for MoE in Elmiral model.

Decision makers and options	Preference statements	Interpretation
MoE		
1. Modify	−4	MoE most prefers that UR not abandon its Elmira facility.
UR		
2. Delay	3	Next, MoE prefers that UR accept the current control order.
3. Accept	−2	Next, MoE prefers that UR not delay the appeal process.
4. Abandon	−1	Next, MoE prefers not to modify the original control order.
LG		
5. Insist	5 iff−1	Finally, MoE prefers that LG support the implementation of the original control order if and only if MoE does not modify it.

This completes the description of model entry in GMCR II. In practice, the system seems to work very well for most strategic conflicts, though we have encountered a few for which it is difficult to frame a Graph Model in option form. But one important success of the option-form entry system is that it is easy to make small changes in a model so that it can be reanalyzed, in order to assess whether the features changed are important to the conclusions.

4.3. GMCR II analysis and output display

The analysis engine of GMCR II is essentially the algorithm of GMCR I, modified to increase speed and capacity. To date, the largest model analyzed using GMCR II is a model of international negotiations over trade in services, originally developed in Hipel et al. (1990), and analzyed using another technology. This model has a six DMs, 21 options, and over 100,000 states; GMCR II determined the stability of every state for every DM according to the SEQ definition in about 8 hours. But for most models of real-world disputes that we have constructed using GMCR II, the analysis results are available in seconds.

For details regarding GMCR II's output displays, see Fang et al. (2003b). Typical of these displays is the GMCR II Equilibria property page; for the Elmiral model, this page is shown as Figure 3. Note that Elmiral is a very small model; states are described using

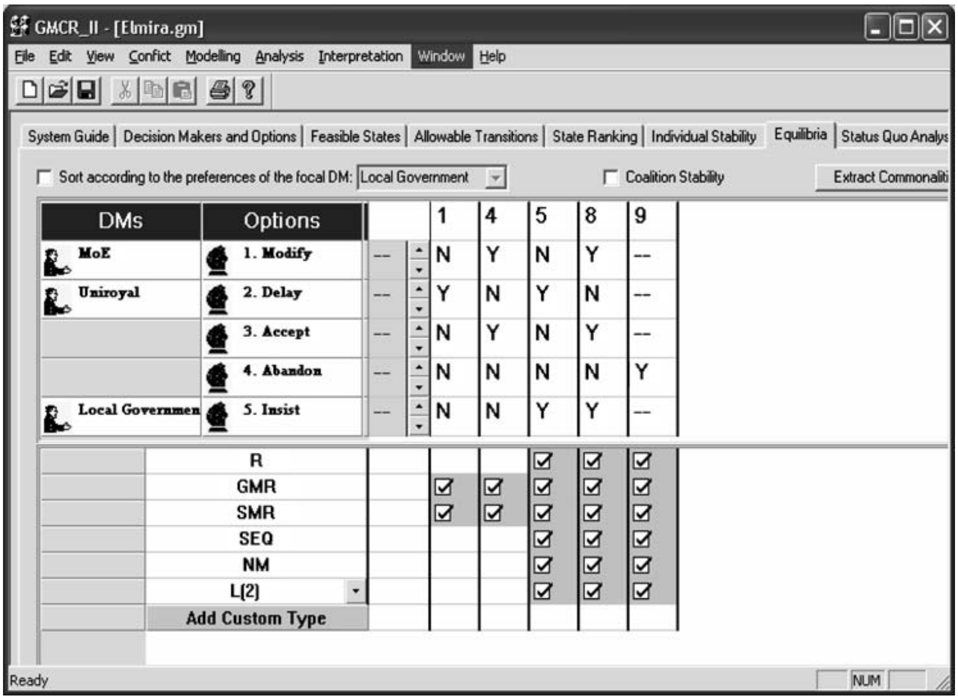

Figure 3. GMCR II equilibria property page.

five options (strictly speaking only four options are necessary), and only five states have any form of stability for all DMs. In Figure 3, 'Y' indicates an option selected by the DM controlling it, while 'N' means the option is not selected, and '-' means either Y or N. Figure 3 shows that states 5, 8, and 9 are stable for all DMs under every definition. (Only these three states can be called strongly stable.) It should be noted that GMCR II also provides an Individual Stability property page, which can be used to find, for each DM, the stability of every state under each of the stability definitions incorporated into the system. Study of this page shows that, while states 1 and 4 are stable for both MoE and UR under all of the GMCR II stability definitions, they are stable for LG only under the GMR and SMR definitions.

The fast analysis and informative displays of GMCR II facilitated a broad range of applications. They also provided a platform for the first form of development of the next generation the graph model, Coalition Analysis, to be described next.

4.4. Coalition analysis

The GMCR II Coalition Analysis algorithm (Kilgour et al., 2001) constitutes a technique of Graph Model analysis beyond the stability analysis introduced in Section 3.2 and Table 1. In fact, the strategic conflict represented by the Elmiral model provides a good example of the importance of coalition analysis; in fact, it was an important motivation. Recall (see Figures 1 and 3) that the model features two strong equilibria – states 5 and 8 – in which Uniroyal does not abandon its Elmira facility. Basic Graph Model stability analysis provides no way to distinguish between these two equilibria – they both have all forms of stability for all DMs. In the actual events in Elmira in 1991, an equilibrium corresponding to state 5 was reached quickly, and remained in place for several months. Then, in a dramatic turn of events, MoE and Uniroyal announced an agreement that effectively shifted the equilibrium to state 8. Local Government was not part of the agreement, and was clearly harmed by it. What happened?

Coalition Analysis offers an answer. If a state is an equilibrium, then no DM is motivated to move away from it. But there is a catch – a DM prefers not to move away in the expectation that the opponents may act together, against the DM's interest, to sanction the move, which thereby stabilizes the original state. But when a coalition of two or more DMs moves away from a state, then not only may the coalition reach better states, but also there are fewer sanctioners. The objective of Coalition Analysis is to detect equilibria (based on what are now called *individual stability calculations*) that are vulnerable to coalition moves.

Several conditions define a coalitionally unstable equilibrium state. First, a *coalition* – a subset of **N** containing two or more DMs – must be able to achieve a state, the *target state*, that all members of the coalition prefer to the initial state. Second, the coalition must be *minimal* with respect to the target state, i.e., no subset of the coalition can reach the target state. (Nonetheless, the move from the initial state to the target state must require at least two moves, by at least two DMs.) Finally, the target state must itself be an equilibrium – otherwise, the DMs in the coalition would have no assurance of obtaining a final state that they all prefer.

Coalition Analysis (activated by a check box on the Equilibria property page, Figure 3) shows that the Elmiral model contains a coalitionally unstable equilibrium, namely state 5. DMs MoE and UR can jointly move the conflict from state 5 to state 8, which they both prefer to state 5 as shown in Figure 1. The primitive Coalition Analysis algorithm in GMCR II carries out this calculation, and shows that state 5 is coalitionally unstable in favor of state 8. However, the algorithm is not very sophisticated; it shows a maximum of one coalitional instability for each equilibrium, and does not indicate the coalition(s) causing the instability, which must be calculated directly from the model. Nonetheless, the Coalition Analysis feature of GMCR II is noteworthy as it was the first success for the Graph Model methodology beyond the traditional individual stability calculations.

4.5. Applying GMCR II

The methodology of the Graph Model for Conflict Resolution and its associated decision support system GMCR II capture the key characteristics of a strategic conflict, so that the strategic issues are better understood and decision makers are better informed. Among the benefits to be gained from this methodology are the following:

- Putting a strategic conflict into perspective,
- Furnishing a systematic structure,
- Facilitating a better understanding of strategic decisions,
- Permitting easy and convenient communication,
- Pointing out where more information is required,
- Understanding strategic implications quickly, before a conflict is resolved or progresses to another phase,
- Identifying stable compromises,
- Providing strategic insights and advice,
- Expeditiously performing extensive sensitivity analyses to determine how model parameters affect conclusions,
- Suggesting optimal decision paths to a specific DM, and
- Identifying opportunities for coalition formation to move to a mutually preferred stable outcome.

Specific advantages of employing Graph Model decision technologies include

- Strategic conflicts with any finite number of decision makers and states can be analyzed (GMCR II can analyze small, medium and relatively large disputes);
- Minimal information is required to perform an analysis (using GMCR II, a model can be calibrated using a list of DMs, a list of options, and relative preference information);
- All possible states (scenarios) are algorithmically generated;
- Moves can be modeled as reversible or irreversible according to the analyst's choice;
- Movements among states are tracked systematically, using each DM's own directed graph of state transitions;

- A rich range of human behavior under conflict conditions can enter into the calculation of stability;
- Equilibria, including compromise resolutions, can be forecast;
- Extensive sensitivity analyses can be conveniently executed.

Further benefits of employing the upgraded version of GMCR II are described in Section 5. Below is a list of situations to which GMCR II can be beneficially applied.

- *Analysis and simulation tool for conflict participants:* A consultant can use GMCR II in simulation and role-playing exercises, for example to help participants to understand the thinking of their competitors.
- *Analysis and communication tool for mediators:* Between sessions, a mediator may use GMCR II to gain insight into how to guide conflicting parties toward a stable win/win resolution. In addition to the avoidance of unstable outcomes, the mediator may see opportunities for side payments which might change preferences so as to stabilize desirable outcomes.
- *Analysis tool for a third party or a regulator:* Others are interested in the outcomes of strategic conflicts, such as representatives of third parties and regulators who frame the rules within which conflicts are played out.

Besides these general situations, GMCR II can be employed in conjunction with many other procedures for negotiation and conflict resolution.

5. Future: New Initiatives

5.1. Status Quo Analysis

The general idea for Status Quo Analysis was conceived early in the development of the Graph Model, and GMCR II provides for it, but the concept was not sufficiently developed to include in the Decision Support System. A consistent and effective set of Status Quo Analysis definitions and algorithms was introduced in Li et al. (2004b, 2005, 2006).

The main idea of Status Quo Analysis is to look forward from the current state, usually the status quo, in order to identify attainable states and to assess how readily they can be attained. Special attention is paid to attaining states known to be stable (for example, states that were found to be equilibria in some prior stability analysis). In a sense, Status Quo Analysis is the reverse of Stability Analysis; Status Quo Analysis is dynamic, following the actual choices of the DMs, whereas Stability Analysis is static, identifying states which, if attained, would be stable.

Several Status Quo Analysis algorithms have been developed. One variation takes account of the DMs' propensity to disimprove, and another optimizes the procedure when the DMs' graphs $G_i = (\mathbf{S}, A_i)$ are transitive. The result is a Status Quo Diagram, which tracks possible moves beginning at the status quo state in order to identify attractors and other states

Table 3. Status Quo table for Elmiral model.

$V_t^{(h)}$	SQ	2	3	9	5	4	6	7	8
$V_T^{(0)}$	√								
$V_T^{(1)}$	√	MoE	UR	UR	LG				
$V_T^{(2)}$	√	MoE	UR	UR	LG	√	√	√	
$V_T^{(3)}$	√			UR	LG	√	√	√	√
$V_T^{(4)}$	√			UR	LG	√	√	√	√

Source. Li et al. (2005).

of interest. Other procedures summarize information about attainability by constructing a Status Quo Table. To illustrate, the basic Status Quo table for the Elmiral model of Figure 1 is given in Table 3.

To interpret Table 3, note that the status quo state, called SQ, is state 1. The states reachable from the status quo (in this case, all states) are listed on the top row of the table. Rows of the table correspond to numbers of moves, $h = 0, 1, 2, 3$, and 4. If there is no entry in the cell corresponding to state s and row h, then state s cannot be reached from the status quo in h moves. If the name of a DM appears in this cell, then state s can be reached from the status quo in h or fewer moves, and the named DM must make the last move (by which state s is actually attained). Finally, if the symbol $\surd$ appears in this cell, then state s can be reached from the status quo in h or fewer moves, and at least two different DMs can make the last move. (It is important to keep track of the last mover because of the Graph Model convention that no DM can move twice in succession. For example, the sequence SQ, 3, 9 would be ruled out in the Elmiral model, since it requires two consecutive moves by UR. But in this model all DMs' graphs are transitive, so the one-move sequence SQ, 9 would be possible.)

Algorithms have been developed and tested for Status Quo Analysis, and several applications have demonstrated how much Status Quo Analysis adds to Stability Analysis. But coding and testing have yet to be done; although GMCR II makes provision for Status Quo Analysis, including it in the software remains a future project.

5.2. Stability analysis with uncertain preferences

In a Graph Model, DM i's preference over $\mathbf{S}$ can be thought of as a pair of binary relations $\{\succ_i, \sim_i\}$ on $\mathbf{S}$, $\succ_i$ indicating i's strict preference, and $\sim_i$, i's indifference. The Graph Model methodology assumes that this pair of relations is *strongly complete* in the sense that, for any states $s, t \in \mathbf{S}$, exactly one of $s \succ_i t$, $s \sim_i t$, or $t \succ_i s$ is true. But in practice, DMs and analysts sometimes lack information about some state comparisons, to the point that they cannot estimate preference. Now, using recently developed definitions (Li et al., 2004a), a partial stability analysis can be carried out even if there is some uncertainty in preferences; moreover, this partial analysis can be updated as additional preference information becomes available.

In this new approach, the preferences of DM i are described by a triple of relations $\{\succ_i, \sim_i, U_i\}$ on **S**, such that $\succ_i$ and $\sim_i$, are interpreted as before, and sU_it means that the relative preference of states s and t is unknown. Under the assumptions that (1) $\succ_i$ is asymmetric; (2) $\sim_i$ is symmetric and reflexive; (3) U_i is symmetric; and (4) the triple $\{\succ_i, \sim_i, U_i\}$ is strongly complete, four different extensions of the Nash, GMR, SMR, and SEQ stability definitions have been developed. Stability under these definitions is effectively standard stability modified to take the unknown preferences into account. Under Nash, GMR, SMR, or SEQ, the strongest of the new definitions never has more stable states than the standard definitions, and the weakest never has fewer.

The new definitions have been shown to be consistent and useful in some applications. Algorithms could be developed now, but a few modeling issues remain to be addressed before computer implementation.

5.3. Stability analysis including strength of preference

One of the great advantages of the Graph Model for Conflict Resolution is that it depends only on ordinal preference information, which is easy to elicit and leads to conclusions that are game-theoretically unimpeachable. But it is natural to wonder whether additional preference information, if available, could be included in a Graph Model so as to give more nuanced conclusions. For example, suppose that a DM is deterred from moving away from a state because of a possible sanction; if that sanction resulted in a state much less preferred than the original, then the stability of the original state would be enhanced, whereas if the sanction resulted in a mildly less preferred state, the stability of the original state would be somewhat weaker.

Hamouda et al. (2004a) introduced a refined preference structure that includes more information about strength of preference, and developed new stability definitions to reflect the refinement. As usual, DM i's preference is described by a pair of binary relations $\{\succ_i, \sim_i\}$ on **S**. But now the strict preference relation $\succ_i$, is split into two relations, $\gg_i$ and $>_i$, in the sense that $s \succ_i t$ if and only if $s \gg_i t$ or $s >_i t$. The interpretation is that $s \gg_i t$ indicates that i strongly prefers s to t, while $s >_i t$ indicates that i mildly prefers s to t. Formally, the relations $\gg_i$, $>_i$, and $\sim_i$, on **S** have the properties that (1) $\gg_i$ is asymmetric; (2) $>_i$ is asymmetric; (3) $\sim_i$ is symmetric and reflexive; and (4) the triple $\{\gg_i, >_i, \sim_i\}$ is strongly complete.

Then changes in the GMR, SMR, and SEQ stability definitions were proposed to reflect this more detailed preference information. Roughly, a state $s \in \mathbf{S}$ is *strongly stable* for DM i if i has improvements from s, but after any improvement the opponents could move to a state t such that $s \gg_i t$. Analogously, a state $s \in \mathbf{S}$ is *weakly stable* for i if i has improvements from s, and after any improvement the opponents could move to a state t such that $s \succ_i, t$, but for at least one improvement from s the sanction t satisfies $s >_i t$. Examples described by Hamouda et al. (2004a) show that information about strength of preference can distinguish levels of stability that are meaningful in practice.

Again, while these new definitions are known to be consistent and believed to be useful, algorithms for computing them are a future project; coding and testing will follow later.

5.4. Perceptual graph models

Coalition Analysis and Status Quo Analysis constitute new ways of analyzing a standard Graph Model. Uncertain Preference and Strength of Preference constitute features that can be added to Graph Models in a way that enhances, rather than compromises, the standard Stability Analysis. Yet another recent initiative is the development of the Perceptual Graph Model, which is a Graph Model in which different DMs perceive different sets of states. A good argument can be made that models in which state perceptions differ offer a descriptive dimension that cannot be achieved by other means (Obeidi et al., 2005). For example, it has been established that the presence of strong negative emotion may prevent a DM from perceiving certain possibilities.

A simple way to produce Perceptual Graph Models that remain amenable to Stability Analysis is now under development. It begins with a complete underlying Graph Model and assumes that each DM perceives only a subset of the states. Then it requires that transitions between perceived states, and preference orderings on perceived states, be inherited from the complete model. Although this approach is far from complete, it provides enough consistency across DMs' models that many established analysis techniques can be applied meaningfully and, we believe, usefully.

5.5. Policy stability

Zeng et al. (2005) develop a novel approach to the analysis of a Graph Model. A *policy* for a DM is a complete plan that specifies the DM's intended move starting at every state in a graph model. Given a profile of policies, a *Policy Stable State (PSS)* is a state $s^* \in S$ such that (1) s^* is an equilibrium in the sense that no DM's policy calls for a move away from s^*; and (2) no DM would prefer to change its policy, given that the policies of the other DMs are fixed. The profile of policies associated with a PSS is a *Policy Equilibrium*. Policy stability is an interesting and useful idea on its own; moreover, examination of its relationship with the standard forms, including Nash, GMR, SMR and SEQ (see Table 1), is providing new perspectives on Stability Analysis in the Graph Model for Conflict Resolution.

6. Summary and Conclusions

Table 4 summarizes the status of the various initiatives, discussed above, that have been undertaken to develop the methodology of the Graph Model for Conflict Resolution.

Section 5 of this review was called "Future", to suggest that the developments and initiatives described there are soon to be implemented in software and used widely in practice. But it is appropriate to end this review with some other perspectives on the future. The authors, and their many colleagues and collaborators, are confident that the Graph Model methodology offers a valuable and flexible tool for the study and understanding of strategic conflict.

Table 4. Current status of graph model initiatives.

Model	Analysis	Current status
Graph model	Individual stability status Quo policy stability	Definitions and algorithms definitions and algorithms definitions
Graph model with option-form entry	Individual stability	In GMCR II
	Coalitional stability	In GMCR II
Graph model with uncertain preference	Individual stability	Definitions and algorithms
Graph model with strength of Preference	Individual stability	Definitions
Perceived graph models	Individual stability	Under development

We believe that strategic conflict is best understood as a process of negotiation—often informal or implicit, and sometimes even ill-structured. The voluminous literature on negotiation includes many calls for systems to analyze the strategic problems of negotiators, but few reports of success. Many have suggested that the natural tool to analyze strategic problems "should be" game theory. But for various reasons, including its insistence on fixed rules of play and its strong assumptions about shared knowledge, game theory was dismissed as "theoretical acrobatics" by those who study negotiation (Raiffa, 1982, p. 6). Later, Raiffa explicated: "[F]or a long time I have found the assumptions made in standard game theory too restrictive for it to have wide applicability [to negotiation] . . . Such limits are hard to swallow in seeking to put this elegant theory into practice" (Raiffa et al., 2002, p.12). In our view, a negotiation is a strategic conflict; as we argued above (Section 2), games are often poor models for strategic conflicts.

The Graph Model for Conflict Resolution models strategic conflicts in a way that avoids the problems of game models. It draws on game theory, but is not game-theoretic in the traditional sense. Its strength is its simplicity and flexibility, both in modeling and analyzing strategic decision problems. The Graph Model incorporates some plausible restrictions on knowledge and rationality, making it appropriate for advising individuals in a multi-decision-maker context. Another advantage is that it has been implemented efficiently in a decision support system, which has led to an extensive list of applications. With this experience has come considerable expertise in Graph Model analysis and application. Yet the Graph Model is a continuing project and, as this article has shown, there is plenty of scope for developments and improvements over the next few decades.

References

Brams, S. J. (1994). *Theory of Moves*, Cambridge, UK: Cambridge University Press.

Fang, L., K. W. Hipel, and D. M. Kilgour. (1993). *Interactive Decision Making: The Graph Model for Conflict Resolution.* New York: Wiley.

Fang, L., K. W. Hipel, D. M. Kilgour, and X. Peng. (2003a). "A Decision Support System for Interactive Decision Making, Part 1: Model Formulation," *IEEE Transactions on Systems, Man and Cybernetics, Part C* 33(1), 42–55.

Fang, L., K. W. Hipel, D. M. Kilgour, and X. Peng. (2003b). "A Decision Support System for Interactive Decision Making, Part 2: Analysis and Output Interpretation," *IEEE Transactions on Systems, Man and Cybernetics, Part C* 33(1), 56–66.

Fraser, N. M. and K. W. Hipel. (1984). *Conflict Analysis: Models and Resolutions*. New York: North Holland.

Hamouda, L., D. M. Kilgour, and K. W. Hipel. (2004a). "Strength of Preference in the Graph Model for Conflict Resolution," *Group Decision and Negotiation* 13, 449–462.

Hamouda, L., K. W. Hipel, and D. M. Kilgour. (2004b). "Shellfish Conflict in Baynes Sound: A Strategic Perspective," *Environmental Management* 34(4), 474–486.

Hipel, K.W. (2002). "Conflict Resolution," Theme Overview Paper, in Conflict Resolution, *Encyclopedia of Life Support Systems (EOLSS),* Oxford, UK: EOLSS Publishers, [http://www.eolss.net].

Hipel, K. W., Fraser, N. M., and Cooper, A. F. (1990). "Conflict Analysis of the Trade in Services Dispute," *Information and Decision Technologies* 16(4), 347–360.

Hipel, K. W., D. M. Kilgour, L. Fang, and X. Peng. (1997). "The Decision Support System GMCR II in Environmental Conflict Management," *Applied Mathematics and Computation* 83(2/3), 117–152.

Hipel, K. W., D. M. Kilgour, L. Fang, and X. Peng. (2001). "Strategic Decision Support for the Services Industry," *IEEE Transactions on Engineering Management* 48(3), 358–369.

Howard, N. (1971). *Paradoxes of Rationality: Theory of Metagames and Political Behaviour.* Cambridge, MA: MIT Press.

Howard, N. (1999). *Confrontation Analysis: How to Win Operations Other Than War.* Pentagon, Washington, DC: CCRP Publications.

Kandel, A., Y. Zhang, P. S. S. Borges. (1998). "Fuzzy Prisoner's Dilemma," *Fuzzy Econ. Rev.* 3(1), 3–20.

Kilgour, D. M., K. W. Hipel, and L. Fang (1987). "The Graph Model for Conflicts," *Automatica* 23(1), 41–55.

Kilgour, D. M., K. W. Hipel, L. Fang, and X. Peng. (1998). "Applying the Decision Support System GMCR II to Peace Operations," in *Analysis for and of the Resolution of Conflict* (A. E. R. Woodcock and D. F. Davis, eds.) Cornwallis Park, NS: Canadian Peacekeeping Press, pp. 29–47.

Kilgour, D. M., K. W. Hipel, L. Fang, and X. Peng. (2001). "Coalition Analysis in Group Decision Support," *Group Decision and Negotiation* 10(2) 159–175.

Li, K. W., Hipel, K. W., Kilgour, D. M., and Fang, L. (2004a). "Preference Uncertainty in the Graph Model for Conflict Resolution," *IEEE Transactions on Systems, Man, and Cybernetics, Part A* 34(4), 507–520.

Li, W., F. Karray, K. W. Hipel, and D. M. Kilgour. (2001). "Fuzzy Approaches to the Game of Chicken," *IEEE Transactions on Fuzzy Systems* 9(4), 608–623.

Li, K. W., D. M. Kilgour, and K. W. Hipel. (2004b). "Status Quo Analysis of the Flathead River Conflict," *Water Resources Research* 40, W05S03.

Li, K. W., D. M. Kilgour, and K. W. Hipel. (2005). "Status Quo Analysis in the Graph Model for Conflict Resolution," *Journal of the Operational Research Society* 56(6), 699–707.

Li, K. W., K. W. Hipel, D. M. Kilgour, and D. J. Noakes. (2006). "Integrating Uncertain Preferences into Status Quo Analysis with Application to an Environmental Conflict," *Group Decision and Negotiation* (this issue).

Noakes, D. J., L. Fang, K. W. Hipel, and D. M. Kilgour. (2003). "An Examination of the Salmon Aquaculture Conflict in British Columbia using the Graph Model for Conflict Resolution," *Fisheries Management and Ecology* 10, 1–15.

Noakes, D. J., L. Fang, K. W. Hipel, and D. M. Kilgour. (2006). "The Pacific Salmon Treaty: A Century of Debate and an Uncertain Future," *Group Decision and Negotiation* (this issue).

Obeidi, A., K. W. Hipel, and D. M. Kilgour. (2002). "Canadian Bulk Water Exports: Analyzing the Sun Belt Conflict using the Graph Model for Conflict Resolution," *Knowledge, Technology, and Policy* 14(4), 145–163.

Obeidi, A., K. W. Hipel, and D. M. Kilgour. (2005). "Perception and Emotion in the Graph Model for Conflict Resolution," *Proceedings of the 2005 IEEE International Conference on Systems, Man and Cybernetics,* Waikoloa, HI, October 2005.

Obeidi, A., K. W. Hipel, and D. M. Kilgour. (2006). "The Role of Emotions in Envisioning Outcomes in Conflict Analysis," *Group Decision and Negotiation* (this issue).

Raiffa, H. (1982). *The Art and Science of Negotiation*. Cambridge, MA: Harvard University Press.
Raiffa, H., J. Richardson, and D. Metcalfe. (2002). *Negotiation Analysis: The Science and Art of Collaborative Decision Making*. Cambridge, MA: Harvard University Press.
von Neumann, J. and O. Morgenstern. (1944). *Theory of Games and Economic Behavior.* Princeton, NJ: Princeton University Press.
Zeng, D.-Z., L. Fang, K. W. Hipel and D. M. Kilgour. (2005). "Policy Stable States in the Graph Model for Conflict Resolution," *Theory and Decision* 57, 345–365.