

ABSTRACT

FATHM is a large-scale attrition model of two-sided ground combat, including air strikes by Blue on Red. While FATHM has borrowed extensively from other combat models, it also exhibits some unique features, particularly its use of linear programming calls to solve a sequence of air-to-ground engagements that are interspersed among ground-to-ground engagements. In spite of making hundreds of such calls, FATHM can still fight a realistically-scaled theater war in about three minutes on a personal computer. This paper gives the genesis and essential features.

INTRODUCTION

The FAst THEater Model (FATHM) is a two-sided, aggregated joint theater combat model that fuses Air Force style Air-to-Ground attack sortie optimization with Ground-to-Ground Lanchester fire-exchange battles using attrition rates derived from the Army's high-resolution COmbat SAmples GEnerator (COSAGE) model [Jones, 1995]. Figure 1 indicates how the two parts interact. The inputs to FATHM include the outputs (killer-victim scoreboards) of COSAGE, and also a subset of the inputs to the Air Force's Conventional Forces Assessment Model (CFAM) [Burton, 2005]. COSAGE killer-victim scoreboards are summary tables that primarily show who was killed by whom over a two-day battle.

The modeled FATHM war is conducted in three-day periods, with one Air-to-Ground and one Ground-to-Ground action carried out in each period. Period-by-period, damaged and destroyed targets may regenerate, and there may be scheduled reinforcements of attacking platforms, munitions, and new targets. The war progresses in phases whose completion depends on threshold levels of target kills in class categories and limits on phase duration. Each phase has a separate COSAGE input file, so the phases may differ qualitatively from each other. All FATHM inputs and outputs are ASCII flat-files suitable for immediate integration with a host database and spreadsheet analysis.

FATHM's original sponsor was the Joint Staff (J8), which has the problem of

partitioning targets among the armed services as a step preparatory to munitions procurement. For this reason, FATHM's constraints include some whose intention is to ensure that all armed services are equally "stressed", in some sense, by a major regional contingency. FATHM was never actually used for that purpose, however, and subsequently became viewed as a general purpose, theater-scale combat model. While asymmetric warfare is much in the news of late, the United States military is still required to maintain forces as a hedge against possible peer competitors [Department of Defense, 2006]. Theater-scale models of attrition warfare such as FATHM are required to plan for such contingencies.

The Fog of War is present in FATHM only to the extent that it influences the COSAGE killer-victim scoreboards that FATHM relies on. FATHM's utility for studying communications issues is therefore minimal. FATHM's natural domain is instead the study of the quantity and quality of committed forces. Since many scenarios must be evaluated in performing such studies, computational speed has been emphasized throughout FATHM's development.

FATHM is innovative among theater-scale combat models in its repetitive, dynamic use of linear programming to fight the Air-to-Ground battles. Each linear program incorporates data determined by the progress of the war so far, including attrition to Blue ground forces, as will be explained in detail below. The solution of significant linear programs within a dynamic combat model is still rare, but increasingly common because of the capability of modern computers and solvers. FATHM is one example. Another is the Integrated C4ISR Analytic Tool Set (ICATS) [Walsh and Roy, 2001], and there are probably others.

GENERAL STRUCTURE OF FATHM

The details of FATHM are reported in Brown and Washburn [2000]. Here we emphasize only the general structure, particularly the issues that arise in integrating two unrelated combat models.

The Fast Theater Model (FATHM)

Gerald G. Brown

Naval Postgraduate School
gbrown@nps.edu

Alan R. Washburn

Naval Postgraduate School
awashburn@nps.edu

APPLICATION AREAS:
Joint Campaign
Analysis and Modeling,
Simulation and
Wargaming

OR METHODOLOGIES
Linear Programming
and Markov Processes

THE FAST THEATER MODEL (FATHM)

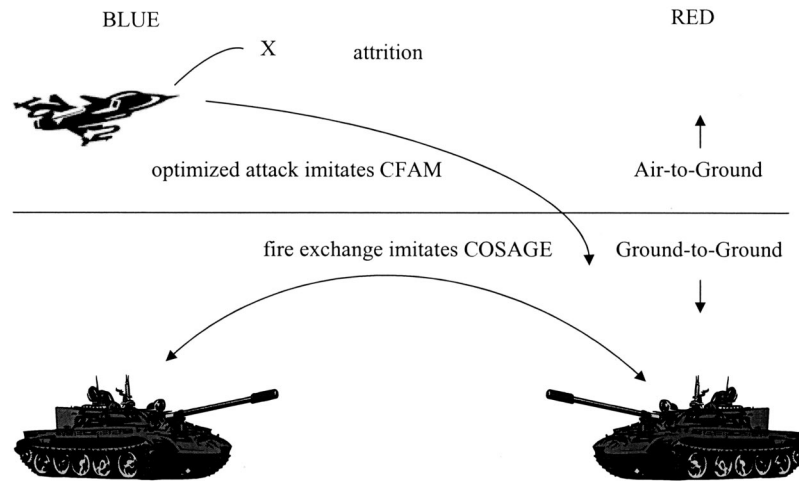


Figure 1. FATHM is a fusion of ground and air warfare.

FATHM sacrifices spatial detail for simplicity and speed, aiming to accommodate a user who must compare many scenarios quickly in the process of making decisions. There is no reference to latitude or longitude in FATHM, for example, so all attempts to represent spatial relationships other than Forward Edge of the Battle Area (FEBA) movement (see below) must be through proliferation of entity types. FATHM nonetheless is capable of dealing with considerable detail in other dimensions.

The Ground-to-Ground combat model in FATHM (hereafter the Ground model) is basically a system of ordinary differential equations (a Lanchester system) solved numerically, except that the state of the system is changed after each Ground battle by a call to a linear program that represents the effects of Blue air strikes. FATHM uses the appropriate COSAGE killer-victim scoreboard to obtain Lanchester coefficients for both aimed and unaimed fire.

The Air-to-Ground part of FATHM (hereafter the Air model) consists of a sequence of sorties by Blue platforms against Red targets. Blue air strikes are conducted similarly to CFAM, using data of the CFAM style even for Navy air assets committed to the battle. The data include dynamic target values that reflect the status of the Ground battle. Most platforms are fixed-wing aircraft, but launchers of expensive munitions such as TOMAHAWK and AT-ACMS are also put in this category. Blue air

platforms are gradually killed in accordance with input attrition data.

Because Blue air supremacy is assumed, only platforms involved in direct attack are modeled. Suppression of enemy air defenses, combat air patrol, and electronic countermeasures, in particular, are not modeled. If sorties for these functions are required in reality, then they must be subtracted from the total made available to the Air model.

The Air model keeps track of munitions expenditure, and will respect any munitions constraints that are imposed. Indeed, one of the functions of FATHM is to measure sensitivity to such constraints. The FATHM objective function flexibly acknowledges the importance of

- ending the current phase quickly;
- assuring an equitable distribution of effort over the services;
- avoiding attrition; and
- killing the Red targets that are killing the most Blue platforms on the ground.

Sea battle is not represented, nor is Air-to-Air battle except by assuming that Blue controls the air throughout. Ground-to-Air battle is represented only in that Blue air strikes in the Air model carry the implication of possible attrition.

Logistic calculations are carried out at the end of each period. Such calculations include Red and Blue reinforcements that have been

scheduled to arrive or move between theaters (FATHM is designed to accommodate more than one theater), Red targets that regenerate after having been killed, and Red targets that change type. An example of the last might be deep tanks not involved in the Ground battle that gradually arrive and become active. A certain fraction of tanks that are deep in one period is converted to tanks that are active in the next. This Markovian method is used in all cases of Red target transformation. It is through this method of entity proliferation (there are “deep tanks” and “active tanks”, instead of just “tanks”) that FATHM is able to achieve some semblance of spatial relationship and movement.

In each three-day time period there are basically three kinds of computations. First the ground battle, then the air battle, and finally the logistic computations that deal with reinforcements and reconstitution. We assert no command-and-control insight in making this particular sequential arrangement. The three functions are imagined to proceed in parallel in reality, so the order in which the computations take place should ideally be of small significance. The length of the three-day time period is a compromise between that goal and the equally important goal of computational speed.

Figure 2 illustrates FATHM’s basic structure. FATHM is a deterministic, time-stepped combat model with no interactive human controls, other than the initial setup.

INTERFACE ISSUES

FATHM is intended to be a fusion of COSAGE and CFAM, but, because the models were not designed to be compatible, there are some significant interfacing problems. Chief among these is nomenclature. The two models have different names for targets, and may also have different levels of detail. A T72 tank to CFAM may correspond to multiple kinds of T72 tanks in COSAGE, or vice versa. FATHM retains both systems of nomenclature, but also introduces a third global name where each entity has a unique index. One Red target type might have global index 14, CFAM index 14, and COSAGE name RT72SM, while a second

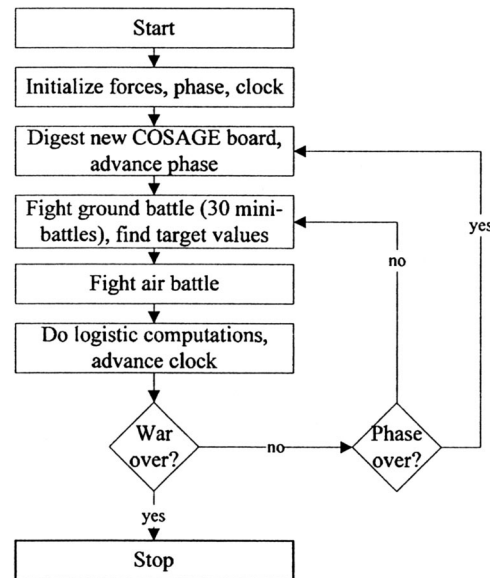


Figure 2. Basic Flow Diagram for FATHM.

target type might have global index 15, CFAM index 14, and COSAGE name RT72S. These two target types would be assumed indistinguishable in the Air model because they have the same CFAM index, and consequently the two types of T72 would be killed in proportion to their numbers on the battlefield. There are also targets acknowledged by one model, but not the other. Command posts, for example, play no role in the Ground model, but nonetheless may serve as targets in the Air model. Another example might be deep tanks that are not yet involved in the Ground battle, but that are within range of air strikes.

Although dealing with multiple systems of nomenclature required a significant part of the development time for FATHM, it will be ignored in the sequel because dealing with it involves a large and distracting complication of notation. See Brown and Washburn (2000) for the full, detailed version. Only the global index will be referred to below.

FATHM is careful to avoid redundant inclusion of the effects of platforms modeled in both COSAGE and CFAM. For example, COSAGE runs already include certain types of air support. FATHM retains some platform types (rotary wing) as part of the Ground battle, while others (fixed wing) are subtracted out of

THE FAST THEATER MODEL (FATHM)

the COSAGE killer-victim scoreboards and included instead in the Air model.

The Ground model does not keep track of munitions expenditure, so certain COSAGE munitions such as ATACMS that are clearly limited in quantity are treated by FATHM as “aircraft” with a very high attrition rate (1.0) per sortie.

CFAM includes the idea that certain killed platforms can come back to life after a time interval long enough to permit repair, whereas COSAGE does not. FATHM retains the distinction, eventually regenerating an input fraction of targets killed in the Air model. Targets killed in the Ground model do not regenerate. There are thus two sources of new targets in FATHM: those scheduled to arrive at a certain time, and those that regenerate after being killed from the air.

Details of the Air and Ground models are outlined below.

THE GROUND MODEL

COSAGE runs are typically for short periods of time such as two days, whereas FATHM battles may last for multiple three-day periods. Although it is unlikely that any platform will be entirely wiped out during the short time of a COSAGE run, that eventuality could very well happen in multiple time periods as a consequence of direct fire. An additional contributor to this possibility is that the Air model may concentrate on only a few target types in any one period. FATHM must therefore be prepared for the possibility that all targets of a particular type will be exhausted.

FATHM imitates COSAGE by including both direct or “aimed” fire (Lanchester’s square-law) and indirect fire (Lanchester’s linear-law). Indirect fire is not influenced by exhaustion of a target type, since the effects of indirect fire vanish naturally when the number of targets becomes 0. Direct fire, however, must be reprogrammed when no targets remain—it does not make sense to be “aiming” at a class of targets that has vanished from the battlefield. The data structures inferred from the killer-victim scoreboards must therefore differ by type of fire. We outline below FATHM’s

method of initially obtaining these, and for modifying them when required by target exhaustion. Only the Blue versus Red case is described, although FATHM carries out parallel computations for Red versus Blue. Variables subscripted for time (t) represent dynamic quantities in the Ground model, or else quantities obtained from a COSAGE killer-victim scoreboard if there is no t subscript. Operations that would result in division by 0 are not necessary and not performed.

Index Use

u	Blue platform type
v	Blue munition type
c or j	Red platform type
d	Red munition type

Direct Fire

B_u = initial number of Blue platforms type u in the COSAGE battle.

R_c = initial number of Red platforms type c in the COSAGE battle.

K_{uvc} = number of kills per day by Blue u shooting munition v at Red platform type c .

S_{uvc} = number of shots per day by Blue u shooting munition v at Red platform type c .

$PKdir_{uvc} = K_{uvc}/S_{uvc}$ = kill probability of each shot.

$drate_{uvc} = S_{uvc}/B_u$ = rate at which each u,v shoots at type c , assuming c is present.

$f_{uvc} = drate_{uvc}/\sum_j drate_{uj}$ = fraction of Blue (u,v)’s fire directed at c .

All of the above are either taken directly from a killer-victim scoreboard or are ratios of such data. At time t , when there are B_{ut} Blue and R_{ct} Red platforms remaining, the Ground model takes the rate at which Blue u,v kills Red type c directly to be:

$$direct_{uvct} = PKdir_{uvc}B_{ut}drate_{uvc} / \sum_{j:R_{jt}>0} f_{uvj}.$$

This amounts to assuming that, if Red platform c gets wiped out, then Blue fire of type (u,v) that was directed against c in the COSAGE run will be proportionally reprogrammed to other targets. If no targets have been wiped out, the sum in the denominator will be 1. Direct fire will not be “wasted” unless the sum in the denominator is 0, in which case $direct_{uvct}$ is taken to be 0 for all c .

Indirect Fire

Because there is no reprogramming issue, it is not necessary to define an indirect kill probability. It suffices to define:

$$irate_{uvc} = K_{uvc} / (B_u R_c) = \text{rate at which each Blue } u, v \text{ kills Red type } c.$$

At time t , the Ground model takes the rate at which Blue type u,v kills Red type c indirectly to be:

$$indirect_{uvct} = irate_{uvc} B_{ut} R_{ct}.$$

Polishing Attrition Rates

If the attrition rates were calculated as described above, and if those rates were then substituted into Lanchester’s equations, the resulting casualties would not be as originally read from the COSAGE killer-victim scoreboard. This is because the COSAGE battle occurs over a significant time interval (call it T , usually two days), whereas the Lanchester time increment is in theory infinitesimal (in practice 0.1 day in FATHM). The COSAGE board reveals attrition over T , but dividing that attrition by T produces only an average or “rough” attrition rate that will lead to disagreement between FATHM and COSAGE if simply inserted into a Lanchester model.

Because the FATHM Ground model is supposed to be a Lanchester model that produces the same results as COSAGE, this potential lack of agreement is unsatisfactory. However, the rough coefficients can be “polished” to make the two models agree. The polishing method can be most easily described using a model that is notationally simpler than the actual Ground model. Suppose it is known that the differential equation $dx/dt = \alpha x(t)$ holds for $t > 0$, with the initial value $x(0)$ and the final value $x(T)$ known, but α unknown. A rough value of α is $\Delta / (x(0)T)$, where $\Delta \equiv x(T) - x(0)$ is the change in x over the time interval $[0, T]$. This is the sort of estimate described above, with Δ being the value read from the COSAGE board. The exact value of α is $\Delta / (\int_0^T x(t)dt)$, but the time record of x is not available. If the time record were available, we could calculate α , or if α were available, we could calculate the time record. The polishing procedure iterates based on that observation. In the algorithm below, n is an iteration index, the attrition rate α_n is associated by numerical integration with the time record $x_n(t)$, and all time records share the same initial value $x(0)$.

1. Initially let α_1 be $\Delta / (x(0)T)$, and set $n = 1$.
2. Use α_n to solve for $x_n(t)$ over the interval $[0, T]$, and let $\delta_n = x_n(T) - x(0)$.
3. If δ_n and Δ are sufficiently close, stop. The polished attrition rate is α_n .
4. Otherwise, let $\alpha_{n+1} = \delta_n / (\int_0^T x_n(t)dt)$, increment n by 1, and go back to step 2.

Upon exit from the procedure, α_n will be whatever parameter makes the change in x over the interval $[0, T]$ be approximately the known value Δ . This recursion works just as well when x is multidimensional, and is the procedure used to polish the initial attrition rate estimates for direct and indirect fire. In practice, step 2 is always executed five times, because that seems to be sufficient to ensure convergence.

Adjusting Frontal Width

Obviously FATHM’s results will be most accurate when FATHM’s Ground battles are

THE FAST THEATER MODEL (FATHM)

most similar to the COSAGE battle from which the Lanchester coefficients are extracted. Nonetheless, FATHM is prepared for certain deviations. FATHM's initial platform numbers need not necessarily equal those of the COSAGE battle. This flexibility is vital, but it also poses a danger.

Suppose, for example, that initial platform numbers in the FATHM battle are all double those of the COSAGE battle. Attrition due to direct fire will double as expected, but attrition due to indirect fire (mainly artillery in most COSAGE battles) will quadruple because indirect attrition is proportional to *both* the number of shooters *and* the number of targets. This would be appropriate if the increased numbers of participants were still enclosed in the same COSAGE battle space, because the density of targets per unit area would be doubled and indirect fire is basically an attack on area, rather than individuals. However, the larger FATHM battle will typically be intended to occur within a larger battle space that will diffuse the targets and thereby reduce the effectiveness of indirect fire. For this reason, the FATHM planner must also provide an additional parameter FEBAWID that represents the frontal width of the battle space in kilometers. If FEBAWID is larger than COSAGEWID, the frontal width of the COSAGE battle, then the indirect fire coefficients are reduced accordingly. The COSAGEWID parameter is included in the preamble of the COSAGE killer-victim scoreboard, and is usually 32.5 kilometers.

To be precise, all indirect fire coefficients for both sides are multiplied by the ratio (COSAGEWID/FEBAWID) before being employed in FATHM's Lanchester battles. Direct fire coefficients are not adjusted. Thus, if the FATHM battle is twice as large as the COSAGE battle in all respects (frontal width as well as platform counts), then direct and indirect attrition will also be doubled.

Only the polished and adjusted coefficients $irate_{uvc}$ and $drate_{uvc}$ are subsequently employed by FATHM. As a result, if FATHM is employed with

- no Air war,
- initial platform numbers that agree with the COSAGE numbers,

- FEBAWID = COSAGEWID, and
- the COSAGE time interval,

then the attrition will agree with the COSAGE killer-victim scoreboard.

Total Attrition

The total rate at which each Blue platform of type u firing munitions of type v kills Red targets of type c at time t is just the sum of direct and indirect terms. There is nothing to prohibit direct and indirect attrition from the same source if the phenomenon occurs in COSAGE. The total rate at which Red targets of type c disappear at time t is thus

$$\begin{aligned} -\frac{d}{dt} R_{ct} &= killrate_{ct} \\ &= \sum_{uv} (direct_{uvct} + indirect_{uvct}), \end{aligned}$$

except that R_{ct} is taken to be 0, rather than negative. There is, of course, a similar expression for the rate at which Blue platforms disappear.

The most complicated part of these evaluations is determining $direct_{uvct}$, which involves a sum over a complicated set. Even so, the arithmetic described above can be done very quickly for hundreds of platforms and munitions. This approach resembles the ATCAL [Concepts Analysis Agency, 1983] model employed by the Army, and in fact FATHM reads the killer-victim scoreboards using FORTRAN code recycled from ATCAL.

With one exception, Red platforms c shooting munition d kill Blue platforms u in an exactly symmetric way. The exception is related to direct Red anti-aircraft fire, and it is needed on account of the assumption of Blue air supremacy in FATHM.

The platforms involved in the Ground model may differ from those in the COSAGE battle, even though it is the COSAGE numbers that are involved in calculating the crucial coefficients $irate_{uvc}$ and $drate_{uvc}$. If the COSAGE board involves platforms that are not in the Ground battle, the effect is as if they were present in 0 number. The implied assumption

about direct fire is that any fire directed at (say) platform RMIG25 by platform UFLAK using munition MUNI in the COSAGE battle will be redirected against other targets in a Ground battle that lacks RMIG25. For example, suppose UFLAK devotes half of its MUNI direct fire to RHELO and half to RMIG25 in the COSAGE battle, but that only RHELO is actually present in the Ground battle. Then, according to the formula for $drate_{UFLAK,MUNI,RHELO}$, the (UFLAK,MUNI) firing rate against RHELO will be effectively doubled (divided by .5) in the Ground battle compared to what it was initially in COSAGE. This is appropriate for Red air targets because it corresponds to the assumption that Blue control of the air makes it possible for platforms like UFLAK to reprogram fire that would otherwise be directed at Red aircraft. However, a symmetric treatment would not be appropriate for Blue air targets.

Red fire against Blue aircraft should not be reprogrammed in the Ground model because this ground fire is the source of the assumed attrition to Blue air in the Air model. In spite of the absence of Blue aircraft in the Ground model, the fire that COSAGE directs against them is retained in FATHM. FATHM accomplishes this with a phantom target for certain Red anti-aircraft systems. The phantom target can cause no damage to Red in the Ground model, but neither can Red shoot it down. For each Red platform of type c and munition of type d , at all times the phantom target attracts whatever fraction of (c,d) 's direct fire is allocated to Blue aircraft in the COSAGE battle. The planner must provide FATHM a "nonentity list" of Blue platform types destined to become phantoms in the Ground battle.

Platform Values

FATHM needs target values mostly because the Air model requires values for the Red targets in order to allocate sorties optimally. A subsidiary use is to calculate FEBA movement, which requires that the forces of both sides be aggregated into a single number, one for Blue and one for Red, in order to determine an overall force ratio.

FATHM's target valuation is partly through inputs (the static part) and partly

through computation (the dynamic part). Static values can be whatever the planner desires, but we imagine them to be simple economic values; i.e., costs. Only the static values are used to determine FEBA movement, with an aggregate value for each side being obtained by summation.

The total values used for Red platforms in the Air model also include a dynamic component. Let the static value of Red platform c be $SVal_c$. The Ground model determines the rate at which each Red platform kills Blue platforms of each type, and therefore the total rate of killing Blue static value, call it $DVal_c(t)$, for Red platform c at time t . The total value of a Red platform is the sum of the static and dynamic parts, except that $DVal_c(t)$ is first multiplied by an input time constant τ . Thus the total value of Red platform c at time t is $RVal_{ct} = SVal_c + \tau DVal_c(t)$, and it is this total value that is used for deciding which targets to engage in the Air model. This target valuation is probably best viewed as a crude attempt to solve what by rights ought to be viewed as an optimal control problem where a tradeoff must be made between defense (shooting at the targets that are hurting you) and offense (shooting at the targets with high static value). The parameter τ can be thought of as bridging an attitude that ranges from offensive when $\tau = 0$ to defensive when τ is large.

Some combat models determine platform values endogenously by using the argument that the value of a platform is just the rate at which it kills value on the other side, almost a circular definition of value. The eigenvalue method and the ATCAL method are examples of this [Caldwell, et al., 2000; Concepts Analysis Agency, 1983]. The beauty of these methods is that they obviate the need for input static target values, but there are also drawbacks. For example, the value of a truck (a common platform type) would be 0 because trucks do not actually shoot at anything, and likewise the value of a communications center would be 0. For these reasons FATHM relies on static values, as well as a dynamic component of value that depends on current combat effectiveness.

During each three-day time period, the Ground model is applied iteratively to a sequence of 30 mini-battles of length 0.1 day,

THE FAST THEATER MODEL (FATHM)

applying Euler's method to approximate the solution of Lanchester's differential equations [e.g., Hamming 1973, p. 382ff]. The last of the thirty mini-battles determines the dynamic values of Red platforms. For each such platform, $DVal_c(t)$ is the total rate at which the average Red platform kills Blue static target value during that period. The total target values are then computed and input to the Air model. In this way Blue Air is encouraged to direct attention to whatever Red platforms are currently proving most troublesome in the Ground battle.

FEBA movement depends entirely on surviving static values. In each period, the static values per platform are combined with the number of surviving platforms of each type to calculate the total value of each side, and the Blue-over-Red value ratio R then determines the rate at which the FEBA moves. The function $f(R)$ that converts the force ratio to a movement rate is a piecewise-linear function whose coefficients are in the input database. It is a symmetric function in that $f(R) = -f(1/R)$, so it suffices to determine $f(R)$ for $R \geq 1.0$. FATHM uses the FEBA movement rate to update the location of the FEBA. The FEBA location does not currently influence the Ground battle in any way, although it could, in principle.

THE AIR MODEL

The Air model is a large linear program that assigns limited aircraft sorties to targets in a variety of conditions in an attempt to simultaneously kill targets, avoid attrition and equalize stress among the services. This section gives the formulation, suppressing references to phase and time period for economy of notation (there are no objectives or constraints that involve multiple time periods). Some of the data are direct input, and other data (target values, for example) are computed between periods based on results of the most recent Ground battle. Because data have lowercase symbols and no time subscript in this formulation, the total target value of Red platform type k is renamed $kvalue_k$.

Subscripts and Sets

$s \in S$ set of services

$p \in P$ set of aircraft platforms
 P_s subset of platforms belonging to service s , a partition of P , $s \in S$
 $m \in M$ set of weapon types
 $k \in K$ set of target types
 $a \in A$ set of attack profiles
 $l \in L$ set of loadouts
 $w \in W$ set of weather states
 $j \in J$ set of target classes
 $k \in K_j$ subset of target types k referenced in target class j

Data

$mxtargk_k$ upper bound on target type k kills (the number available)
 $kvalue_k$ target value for each target of type k
 $mxhours_{pw}$ upper bound on aircraft type p hours used in weather state w
 $used_{kpw}$ hours required for an attack on target k by aircraft p in weather w
 $mxwepns_m$ upper bound on weapon type m use
 cap_p capacity of aircraft p (used only in service equity computations)
 λ multiplier for attrition in objective function
 e_{apmktw} expected kills per sortie of attack profile a , aircraft p , weapon m on target k with loadout l in weather w
 att_{apmktw} expected attrition per sortie as above
 c_{apmktw} weapons used per sortie as above
 $jgoal_j, \overline{jgoal}_j$ lower and upper goals for kills of target class j
 $jpen_j, \overline{jpen}_j$ lower and upper penalties for violating kill goals for target class j
 $sgoal_s, \overline{sgoal}_s$ lower and upper goals for capacity used by service s
 $spen_s, \overline{spen}_s$ lower and upper penalties for violating capacity goals for service s

Variables (all nonnegative)

X_{apmklw}	attacks
$TGTKILLS_k$	targets k killed
$HRSUSED_{pw}$	aircraft platform p hours used in weather state w
$PLTSLOST_p$	aircraft p lost
$WEPUSED_m$	air_weapons m used
$SVCCAP_s$	capacity used by service s
$UNDERKILLS_j$	
$MIDKILLS_j$	
$OVERKILLS_j$	under, slack, and over-kills of target class j
$UNDERCAP_{s'}$	
$MIDCAP_{s'}$	
$OVERCAP_s$	under, slack, and over-achievement of service s goals

Formulation

The number of targets killed is a linear function of the number of sorties assigned, subject to not killing more targets than exist, and similarly for the number of platforms lost, hours used, and weapons used.

$$\begin{aligned}
 KILLS_k: \quad & TGTKILLS_k \\
 & = \sum_{apmklw} e_{apmklw} X_{apmklw} \quad \forall k \\
 & TGTKILLS_k \leq mxtarget_k \quad \forall k \\
 \\
 PLATS_p: \quad & PLTSLOST_p \\
 & = \sum_{amklw} att_{amklw} X_{amklw} \quad \forall p \\
 & PLTSLOST_p \leq mxplats_p \quad \forall p \\
 \\
 WXHOURS_{pw}: \quad & HRSUSED_{pw} \\
 & = \sum_{amkl} used_{kpw} X_{amklw} \quad \forall pw \\
 & HRSUSED_{pw} \leq mxhours_{pw} \quad \forall pw \\
 \\
 WEPNS_m: \quad & WEPUSED_m \\
 & = \sum_{apklw} c_{apklw} X_{apmklw} \quad \forall m \\
 & WEPUSED_m \leq mxwepns_m \quad \forall m
 \end{aligned}$$

The amount of a service's capacity used is a linear function of the number of hours used by the service's platforms.

$$\begin{aligned}
 SERVICE_s: \quad & SVCCAP_s \\
 & = \sum_{p \in P_{s,w}} cap_p HRSUSED_{pw} \quad \forall s
 \end{aligned}$$

For each class of target j , there are soft goals for the number of targets killed and for each service's capacity.

$$\begin{aligned}
 JGOAL_j: \quad & \sum_{k \in K_j} TGTKILLS_k + UNDERKILLS_j \\
 & + MIDKILLS_j - OVERKILLS_j \\
 & = \overline{jgoal}_j \quad \forall j \\
 & MIDKILLS_j \leq \overline{jgoal}_j \\
 & \quad - \underline{jgoal}_j \quad \forall j \\
 \\
 SGOAL_s: \quad & SVCCAP_s + UNDERCAP_s \\
 & + MIDCAP_s \\
 & - OVERCAP_s = \overline{sgoal}_s \quad \forall s \\
 & MIDCAP_s \leq \overline{sgoal}_s - \underline{sgoal}_s \quad \forall s
 \end{aligned}$$

The objective offers terms to balance effort between killing target value, avoiding Blue platform attrition, achieving kills of target classes critical to the war phase, and equitably stressing the services. Attrition is emphasized when input parameter λ is large.

$$\begin{aligned}
 \text{Minimize} \quad & - \sum_k kvalue_k TGTKILLS_k \\
 & + \lambda \sum_p cap_p PLTSLOST_p \\
 & + \sum_j (jpen_j UNDERKILLS_j \\
 & \quad + \overline{jpen}_j OVERKILLS_j) \\
 & + \sum_s (spen_s UNDERCAP_s \\
 & \quad + \overline{spen}_s OVERCAP_s)
 \end{aligned}$$

The Air model is myopic in the sense that the objective function for each period deals only with results in that period. The decision to make the optimization myopic was made only after considerable debate, the alternative being to construct a single monolithic optimization that would include all time periods at once. The myopic optimization is a much smaller linear

program, of course, which contributes to FATHM's speed. The myopic point of view also permits the inclusion of nonlinear operations such as phase advance in the interperiod accounting, without requiring a nonlinear optimization. An additional argument favoring myopia is that monolithic models can exhibit behavior such as holding back munitions to attack targets in the future after their values have increased, something no military planner would contemplate. This last argument could also be turned around, of course—a monolithic model will not make the mistake of using all of its ATACMS missiles in the first period, whereas a myopic model, if unconstrained, might do exactly that. The constraints involving $mxwepns_m$ are included in FATHM to prevent that possibility. Yost [1996] gives a clear accounting of these issues.

PERIOD-TO-PERIOD ACCOUNTING

Interperiod accounting includes inventory constraints to the effect that what enters period $t + 1$ is whatever enters period t , minus losses in period t , plus reinforcements and transformations at the end of period t . Transformations include targets that change type (the aforementioned deep tanks that become active tanks, for example), as well as targets that come back to life after having been killed earlier in the Air model. Because the accounting is conventional, we omit most of the equations for economy. We refer the interested reader to Brown and Washburn (2000).

The data required for transformations are transition rates, as in continuous-time Markov chains (Ross, 2000). For example, the equation for updating the number of targets of type k includes the term $(1 - e^{-len_t/tcon_k})mxregen_k$, where $mxregen_k$ is the number of targets of type k in the regenerating pool at the end of period t . The transition rate from pool to active for type k is the reciprocal of the time constant $tcon_k$ (an input), and len_t is the length of the current time period (three days). The number of targets in the regenerating pool is decremented by the same amount, and augmented by an input fraction of the number of targets killed in the last Air battle. Note that there is only one regener-

ating pool for each target type, rather than one for each target type and each past period. This is in accord with the memoryless property that characterizes transition times in Markov chains.

A time constant also governs the gradual changes in kills per sortie and attrition per sortie in the air war, generally with the former increasing and the latter decreasing with time. Both quantities are initialized at values contained in the CFAM database, and then gradually approach limiting values that are also contained in the database. For example, attrition per sortie for platforms of type p is governed by

$$att_{apmklw} \Leftarrow e^{-len_t/relax_p} att_{apmklw} + (1 - e^{-len_t/relax_p}) attrlast_{apmklw}.$$

Here the time constant is $relax_p$, and the \Leftarrow symbol indicates that the left-hand side is the new value of the attrition coefficient in period $t + 1$ for sorties by aircraft of type p attacking targets of type k with loadout l of weapon m in weather w using profile a . The limiting attrition rate is $attrlast_{apmklw}$. Kills per sortie are governed by a similar equation with the same input time constant.

Each target class includes a phase goal of a certain fraction that must be killed in order to advance the phase. If all of these goals are met, or if the phase has lasted long enough, then the phase advances. It is quite possible that the phase will not advance because of the exhaustion of some munition that is required to kill some target in some particular class (while creating FATHM, the authors have repeatedly learned the old lesson that "amateurs discuss tactics, while professionals discuss logistics"). If all phases are not completed within a specified time interval, then the war is reported to be lost. Circumstances such as this exhibit very high sensitivity to the amount of the crucial munition provided, since the overall result can switch from "lose" to "win" as the amount increases.

IMPLEMENTATION AND EXPERIENCE

With one exception, input files are manipulated within an Excel™ [e.g., Microsoft, 2005a]

workbook *FathmInputs.xls*, which has one sheet for each input file [Brown and Washburn, 2000]. The exception is a large, rarely-edited file in CFAM format that includes effectiveness and attrition data about all of the possible air strike sorties. A macro exports the rest of the data needed by FATHM in the form of comma-delimited text files, including the COSAGE boards that apply to each phase, after which the FATHM executable program (compiled FORTRAN) is invoked.

For each COSAGE board, FATHM extracts direct and indirect firing rates, and then computes and stores the Lanchester coefficients for later use. Next, FATHM reads the rest of the input data files and scrupulously edits them for inconsistencies. These data come from very different sources, so this is the stage where reconciliation and documentation of discrepancies are key. If the input data include a Red target that cannot be killed by any Blue platform, for example, then that fact will be recorded in

FATHM's log file for later perusal by an analyst who wonders why the war never got out of the first phase.

FATHM then commences the first war phase, fighting all the Ground mini-battles that constitute one period, setting up the linear program representing the Air battle, calling the linear program solver, recording actions and outcomes, accomplishing logistics, advancing the time period, and determining phase transitions until the war is over. A complete (detailed, large) history of every action throughout the war is captured in a master *Attacks.csv* database.

A second Excel workbook *Attacks.xls* is used for viewing and summarizing the outputs in *Attacks.csv*. This workbook capitalizes on Excel™ features such as graphics, auto-filtering, and pivot tables. Figure 3 is a typical graph showing kills of various classes of Red targets.

The output file *Attacks.csv* can also be input to a database application that enables queries to




Figure 3: Kills of seven selected Red target types in phase three of a four-phase war, in three-day time periods 4 through 35.

recover such things as overall killer-victim scoreboards.

FATHM has been applied to several large scenarios, including hundreds of platforms, hundreds of thousands of linear programming variables, and about a hundred time periods arranged in four war phases. In other words, the FATHM design works at full scale.

FATHM first checks to see whether the XA callable library [Sunset Software Corporation, 2005] is available, and uses it if so. Using XA, FATHM requires about three minutes to fight a full theater engagement on a 3GHz laptop computer operating under Windows XP™ [Microsoft, 2005b]. If XA is not available, FATHM can also use the intermediate mathematical modeling language GAMS [Brooke, et al., 1998], but doing so may double or triple the running time.

FATHM's use of dynamic target values to link Ground and Air combat turns out to be a significant feature. The quick responsiveness of the Air model to the current situation unfolding on the Ground is striking enough that we wonder how one can achieve such unifying insights by running independent Ground and Air models of combat.

Readers interested in acquiring an executable version of FATHM should contact author Brown.

A Two-Sided, Defender-Attacker Generalization of the Air Model

Seichter [2005] enhances the Air model by endowing Red with the ability to actively prepare for Blue air attacks by deploying dummy targets and anti-aircraft artillery. Red plans these defenses while anticipating that Blue will attack optimally in response to any defensive measure, and Blue can observe defensive preparations and plan Air attacks accordingly. The Air model becomes a mixed integer program — a two-sided Defender-Attacker model [Brown, et al 2005]. Computation times for this generalization are longer, but not much longer.

Possible Future Developments

While FATHM's structure has been dictated by the motivating application, the basic idea of embedding linear programming calls inside of a combat model has considerable flexibility and appeal. The optimization feature provides a direction to combat that makes results both realistic and robust to changes in initial inventories. As computers become faster, it will become practical to incorporate additional details in FATHM or similar models, perhaps details that are geographically motivated, while still providing fast response times.

There is no fundamental reason why the Ground model has to be deterministic. In fact, even the current FATHM Ground model has a clear probabilistic interpretation, and could be run in Monte Carlo mode. Run times would increase considerably, of course, but there would also be the benefit of exposing the stochastic variability of results, or at least the variability within the Ground model. There is no reason why optimization and simulation cannot both be included in the same combat model.

As mentioned earlier, FATHM has no representation of air support operations such as electronic countermeasures. A practical way to include them would be to require that attack sorties be proportionally accompanied by support sorties, and to include the platforms required for support sorties in the database. Since unavailability of supporting assets can be a real limitation on the ability to mount attack sorties, this would increase FATHM's verity without slowing it down significantly.

Acknowledgments

This work has been supported by the Air Force Office of Scientific Research (AFOSR); the Office of Naval Research (ONR); and the Office of the Secretary of Defense, J8 (Joint Chiefs of Staff, Warfare Analysis Division). Our thanks to MAJ Jim McMullin and Mr. John Warren of the Center for Army Analysis for their help with ATCAL, and for contributing the crucial FORTRAN code that reads COSAGE scoreboards. Kirk Yost initiated FATHM, and has

been an advisor throughout its development. Pete Byrne of the joint staff (J8) prepared scenarios of interest, and offered valuable insight as an experienced optimizer, weapons planner, and somebody wearing a Combat Infantryman Badge.

REFERENCES

- Brooke, A., Kendrick, D., Meeraus, A. and Raman, R., 1998, *GAMS: A User's Guide*, GAMS Development Corporation, Washington, DC, September.
- Brown, G.G., Carlyle, W.M., Salmeron, J., and Wood, K., 2005, "Analyzing the Vulnerability of Critical Infrastructure to Attack and Planning Defenses," *Tutorials in Operations Research*, Institute for Operations Research and Management Sciences, September, pp. 102–123.
- Brown, G.G. and Washburn, A.R., 2000, "The Fast Theater Model (FATHM)", NPS-OR-01-002-PR, Naval Postgraduate School, Monterey, CA, October (revised January 2005), 32 pages.
- Brown, G.G., Coulter, D.M., and Washburn, A.R., 1994, "Sortie Optimization and Munitions Planning", *Military Operations Research*, Vol. 1, No. 1, pp. 13–18.
- Burton, S., 2005, personal communication, 31 March.
- Caldwell, W., Hartman, J., Parry, S., Washburn, A.R., and Youngren, M., 2000, *Aggregated Combat Models*, lecture notes, Operations Research Department, Naval Postgraduate School, Monterey, CA, Ch. 6.
- Concepts Analysis Agency, 1983, "ATCAL: An attrition Model using Calibrated Parameters", CAA-TP-83-3.
- Department of Defense, 2006, "Quadrennial Defense Review Report", February 6, 113 pages, <http://www.comw.org/qdr/qdr2006.pdf>, accessed May 9, 2006.
- Hamming, R.W., 1973, *Numerical Methods for Scientists and Engineers*, McGraw-Hill, New York.
- Jones, H.W., 1995, "COSAGE User's Manual, Volume 1 - Main Report. Revision 4", Center for Army Analyses, CAA-D-93-1-REV-4 (August), 112 pages.
- Microsoft Corporation, 2005a, "Microsoft Office", www.office.microsoft.com, accessed March 13, 2005.
- Microsoft Corporation, 2005b, "Windows XP Home Page", www.microsoft.com/windowsxp/default.asp, accessed March 12, 2005.
- Ross, S., 2000, *Introduction to Probability Models* (seventh edition), Academic Press, Ch. 6.
- Seichter, S., 2005, "The FAst THeater Model (FATHM): Optimization of Air-to-Ground Engagements as a Defender-Attacker Model," MS Thesis in Operations Research, December.
- Sunset Software Corporation, 2005, "XA Callable Library", www.sunsetsoft.com, accessed March 13, 2005.
- Walsh, P. and Roy, R., 2001, "Integrated C4ISR Analytical Tool Set", *Research Summaries*, Vol. 8, No. 1, Institute for Defense Analysis.
- Yost, K.A., 1996, "The Time Strike Munitions Optimization Model," NPS-OR-96-001, Naval Postgraduate School, Monterey, CA, (January), 67 pages.