Secrets of the MIT Mystery Hunt: An Exploration of the Theory Underlying the Construction of a Multi-Puzzle Contest

by Mark Louis Gottlieb

Submitted to the Program in Writing and Humanistic Studies
in partial fulfillment of the requirements for the degree of

Bachelor of Science in Humanities and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 1998
[June, 1998]

Author.....
Program in Writing and Humanistic Studies
May 8, 1998

Certified by.....
Edward Barrett
Senior Lecturer
Thesis Supervisor

Accepted by.....
James Paradis
Head, Program in Writing and Humanistic Studies

Secrets of the MIT Mystery Hunt: An Exploration of the Theory Underlying the Construction of a Multi-Puzzle Contest

by Mark Louis Gottlieb

Submitted to the Program in Writing and Humanistic Studies
on May 8, 1998 in partial fulfillment of the
requirements for the degree of
Bachelor of Science in Humanities and Engineering

Abstract

This is an exploration of the rules and guidelines that underlie the structure of a multi-puzzle contest (a competition consisting of one large puzzle made up of a number of smaller constituent puzzles). The MIT Mystery Hunt, a multi-puzzle contest held on campus each January, is the second-largest annual event of this nature in the United States.

The theories put forth in this exploration were culled from personal experience. Having played in four MIT Mystery Hunts and constructed two, as well as participating in other multi-puzzle contests such as the *Miami Herald* Tropic Hunt, the *Random House $10,000 Trivia Challenge*, and the National Puzzlers' League convention extravaganza, I am familiar with the specific format in question. Furthermore, I have a firsthand understanding, from both sides of the contest, of what is necessary and what is optional and, more importantly, what works and what does not.

It was found that the structural framework of a multi-puzzle contest is rather loose; most of the guidelines and elements are optional. The only necessary aspects are the most basic structural components: an endgame and a number of individual puzzles. However, while a multi-puzzle contest can technically work without a majority of the elements discussed, many of these must be included for such a contest to be a successful form of entertainment. The most vital aspects are theme and variety; others that should be included are puzzles that use the available geography and experimental puzzles. Ultimately, the puzzle maker should design the multi-puzzle contest he would most like to participate in himself.

Thesis Supervisor: Edward Barrett

Title: Senior Lecturer

Table of Contents

1 Introduction

I love puzzles. As a puzzle lover, I have found nothing in the puzzle world as thrilling, exhausting, complicated, or cathartic as the MIT Mystery Hunt, a weekend-long marathon puzzle competition held each January at the Massachusetts Institute of Technology. Many years ago, I innocently read about this event. That led me to participate as a competitor; that, in turn, led me to create one. Finally, after competing in four and constructing two (the 1995 and 1997 Hunts), I was compelled to explore this labyrinthine form of entertainment in depth.

The MIT Mystery Hunt is an amateur-written contest that prides itself on its "no rules" format; anything is allowed. But is that really true? Every Hunt has a complex underlying structure, and I reasoned that there must be some consistencies from Hunt to Hunt, in the form of rules that must be applied, principles that should be applied, and suggestions that are often applied. This document is an exploration of the theory involved in the construction of an MIT Mystery Hunt and, in general, any multi-puzzle contest.

The introductory section looks at the general topics of what a multi-puzzle contest is and what a puzzle maker is, then moves on to look at puzzle construction. This is a launching point into discussions of the specific issues involved in the creation of a multi-puzzle contest. These discussions are divided into a section dealing with the contest as a whole and two more sections regarding the contest's two constituent pieces, the endgame and the individual puzzles. The theory presented focuses on underlying structure rather than actual puzzle construction, because such construction choices are limitless: all crossword puzzles have the same structure, but every one of the millions in existence is unique and more are made every day.

1.1 What is a Multi-Puzzle Contest?

A multi-puzzle contest is a form of public entertainment in which the audience gets to be full participants. Its structure is that of a large puzzle consisting of a number of smaller puzzles. Depending on its construction parameters, such a contest may last hours or months, and its geographic scope may limit it to a single room, spread it out across a city, or open it up to the broad expanses of cyberspace.

The multi-puzzle contest is a rather unique, and largely undocumented, entertainment form. Due to the difficulty involved in creating them and the fact that any large-scale event of this nature can never be run a second time, there don't exist too many of them. There are, however, a few highly popular competitions that have become annual traditions. The largest multi-puzzle contest is The Tropic Hunt. A one-day event backed by the *Miami Herald* newspaper, it has been held each November in Miami since 1984 and it draws roughly 10,000 participants each year. The MIT Mystery Hunt is even older; it was begun in 1980. An amateur-written, weekend-long contest, it draws a few hundred participants every January. Since 1981, the National Puzzlers' League's annual convention has always been capped off with a multi-puzzle contest held on its final evening. Sometimes books will come out with this theme. *The Random House $10,000 Trivia Challenge*, written by Henry Hook, was published in 1995. In addition, there are always a number of informal multi-puzzle contests held at gaming conventions and the like.

What type of event is a multi-puzzle contest? In some ways it is like a marathon. It is a race; the person (or team) that finishes first receives a prize. It is open to the general

populace, but unlike other forms of entertainment, such as a movie, audience members must take an active role in order to receive the full entertainment value possible.

In some ways it is like a play. There is a basic narrative structure set up by the puzzle contest. Suspense, intrigue, and an unknown ending exist both on paper, in the framework of the puzzle, and in real life, in the larger milieu of the contest. There is, of course, a cast of characters in real life consisting of a participant himself (the protagonist), his competitors, and the Puzzlemasters in charge of running the contest. Some contests include an additional cast of fictional characters in the story that often accompanies, and ties together, the puzzles.

In some ways it is like a mystery story. In a mystery, the simple order of life is shattered by some sort of horrific act, usually a murder. The detective, through a series of logical deductions, discovers and removes the aberration, thus reinstating the societal order. This genre has always been popular because the reader takes a good deal of vicarious pleasure through the acts of the sleuth. A multi-puzzle contest allows the participant to play, firsthand, the role of detective as he decodes secret messages, fills in empty grids, uncovers hidden clues, and, ultimately, deduces how all the pieces fit into place.

Finally, in some ways it is simply like a puzzle. As in all puzzles, the goal is to grasp the unknown and make sense of it. This is an intrinsic goal of human nature, and that is why solving a puzzle is a pleasurable, satisfying experience. This explains why millions of people solve *The New York Times* crossword puzzle every day, and this explains why thousands of people, many from hundreds, if not thousands, of miles away, flock each year to the Tropic Hunt.

1.2 Relationship Between Puzzle Maker and Puzzle Solver

The role of the puzzle maker is a tricky one. Like a playwright or composer, a puzzle maker doesn't just create material that will entertain an audience; rather, he crafts material through which he will indirectly interact with each individual solver. But unlike other writers of creative material, this isn't a passive interaction: the puzzle maker is directly challenging the solver's mental agility. He dares the solver to defeat him and his creation.

It's not quite a fair challenge, though, because the puzzle solver must defeat the puzzle maker. If a puzzle proves to be too difficult for a solver, then the creator, though triumphant in the duel of wits, has failed because the solver didn't have any fun and wound up frustrated. This is a difficult tenet for a fledgling puzzle creator to grasp, but once he has resigned his ego to the fact that the solvers who attempt his lovingly crafted puzzles ought to prevail over them, he has mastered the fundamental step of puzzle making. Skill, craft, construction technique: these are all secondary to the mindset necessary for creating challenging yet beatable puzzles.

Simply by constructing a puzzle, the puzzle maker enters into a standard unwritten contract with the puzzle solver. This is a familiar relationship with specific boundaries, and it is not unlike other creator/patron relationships such as the one between filmmaker and filmgoer. One facet of the relationship is that the solver blindly puts his trust in the puzzle maker. Among other things, the solver trusts that the puzzle maker is aware of the elements of the relationship, including the very trust currently at work. This sets up a loop of circular logic (the solver trusts the puzzle maker to know that the solver trusts the puzzle maker to know...) that, however paradoxically, binds the puzzle maker to the contract.

As stated above, the solver rightly expects to be mentally challenged and entertained. Other aspects of the trust the solver has in the puzzle maker: he trusts the puzzle will be error-free, that it is labeled with the correct difficulty level, and that the puzzle has a guaranteed solution. The notion of a guaranteed solution is rather basic, but it is the most psychologically powerful element of a puzzle; it is extremely important for a solver to be assured that he can win before delving into the arcane construction of a puzzle. The solver needs to know there's a way out before stepping inside. Part of the responsibility in this relationship falls onto the solver, of course. He shouldn't expect the same standards from *The New York Times* crossword puzzle and the *TV Guide* crossword.

The role of the constructor of a multi-puzzle contest is, for the most part, no different from the role of any other puzzle constructor; he is just working on a different scale. On one hand, the scale is larger because the overall puzzle is so much bigger and more complex than a typical puzzle (like a crossword). On the other hand, the scale is smaller and more intimate: the event is often directly run by the puzzle creator, so there is frequent interaction between him and the solvers. Thus the relationship between puzzle maker and puzzle solver becomes more personal, which is beneficial to both parties. For the puzzle maker, it is gratifying to witness people enjoying his work; in any other situation, the puzzle maker is distantly removed from the solvers. For the solver, he may take an active role in the unwritten contract, rather than the passive role he must usually accept. If the puzzle maker isn't living up to his side of the agreement (for example, if there is a mistake or if the solver isn't having fun), the solver may confront him on the problem and they may work it out together. Due to the ability to interact so closely, the terms of the contract are relaxed in this situation; a mistake is more readily tolerated in a multi-puzzle contest (where it can be corrected on the spot) than in other

puzzle milieus, like a magazine (where a solver must wait for the next issue, which is almost certainly too late, to find corrigenda).

1.3 Puzzle Construction

The issues facing the constructor of a multi-puzzle contest may be broken up into three categories:

- Global concerns
- The endgame
- Individual puzzles

Global concerns involve considerations that affect the entire contest. These are general issues that must be grappled with in the conceptual stage of the contest before any actual puzzle construction may even begin. As these issues are pervasive and will, ultimately, have tangible effects on the shape of the puzzles, they must continually be addressed throughout the construction stage.

The actual multi-puzzle contest is made up of two distinct parts. Most of the contest falls under the heading of individual puzzles. The smaller, but more important, part is the endgame. This is the puzzle that takes the individual puzzles and combines them into a single entity. The structure of the endgame dictates the structure of the entire contest, and, when puzzle construction begins, it is the endgame that must be made first. Without the endgame, the contest is nothing but disconnected individual puzzles.

This isn't to say that the individual puzzles are unimportant. They make up the bulk of the contest. These puzzles are the arms, legs, nose, spleen, and every other body part to the endgame's brain. Sure, the body would survive without an arm or a few toes, but it is the great mass of separate elements that gives the brain its function. The solvers of a multi-puzzle contest will spend the vast majority of their time working on the individual puzzles, and many, if not most, will never make it as far as the endgame. Although the endgame is structurally

more important, the individual puzzles are most of what a multi-puzzle contest is.

2 Global Concerns

There are a number of concerns necessary to the creation of a multi-puzzle contest that have an effect on the entire contest. These considerations must be taken into account during the conceptualization process, for they will affect the overall contest, and they must be kept in mind throughout puzzle construction, for they may affect each individual puzzle. These global concerns are:

- Logistical parameters
- Contest format
- Theme
- Use of specific knowledge
- Use of the available geography
- Puzzlemaster involvement
- Experimental puzzles

The first three concerns have more of an effect on the make-up of the contest as a whole. The logistical parameters that exist for any particular contest impose constraints on the event before the puzzle maker can; he must be aware of them and either abide by their limits or figure out ways around them.

The format of the contest may be one of two options: the individual puzzles may be cascaded or compiled. This decision must be made immediately because the two formats are very different and every other aspect of the contest is affected by this choice.

The theme is an extremely important aspect of a multi-puzzle contest. Technically optional, it is a purely cosmetic element of the contest. But although it serves no deep structural purpose, it has a surface-level one: it ties the individual puzzles together with a motif or a story line. Since the endgame doesn't come into play until the end, the theme is the

only thing that binds the individual puzzles into a cohesive whole during the course of the contest.

The rest of the global concerns can be more specifically applied. They may still have an overall effect on the contest, but these may also affect certain individual puzzles. The puzzle maker must be opportunistic, and one way to do so is in the exploitation of specific knowledge that the contest solvers possess. Whether in the entire contest or just a few of the individual puzzles, targeting an audience makes a multi-puzzle contest more personal and relevant to the participants.

Similarly, the puzzle maker should be opportunistic in the use of the geography available to him. Each multi-puzzle contest is held in a specific location, and each location has its own quirks and features. Taking advantage of these interesting geographic aspects makes for a unique and well-rounded contest.

Determining the right level of Puzzlemaster (PM) involvement is a tricky issue. Getting the PM's involved in the contest makes it more personable and allows the PM's to have some fun as well. Too much PM involvement can backfire, though; the contest needs to be able to stand on its own.

Finally, multi-puzzle contests are ripe for the inclusion of experimental puzzles. They add uniqueness, innovation, and daring to the event. Unfortunately, as experiments, they are prone to failure. The puzzle maker needs to be aware of the risks involved and the potential rewards of success in this area.

2.1 Logistical Parameters

A constructor of a multi-puzzle contest will find that there are a number of logistical parameters that he must grapple with. These constraints must be dealt with first; they affect design considerations even before the constructor has a chance to start designing. The effect that some parameters, like team size, will have on a contest is relatively simple to grasp. Obviously, a contest intended to be solved by teams will be larger and more intricate than one designed for individuals; exactly how large and how intricate depends on the number of people on each team. Some parameters, though, have much more complex effects on contest design.

The strongest outside force that can have an effect on the design of a multi-puzzle contest is its own history. Normally this isn't a large concern; most such events are one-shot deals that have no pervasive history or are annual occurrences that are written by the same people each year. Sometimes, however, an annual multi-puzzle contest is written by a different group of people each year. Contests that operate in this manner include the MIT Mystery Hunt and the National Puzzlers' League convention extravaganza. In these circumstances, a history lesson is crucial. While the constructors may change each time, a large number of participants play in these contests year after year, and they have certain expectations that must be met. Some traditional elements, like an overall structure, a standard puzzle form, or even an overarching principle, must be retained. These are usually general aspects, and they have come to represent the contests they appear in. Everything else – all specific details – must be different from what was in the previous year's event (and, preferably, the events from the few years before that as well). For example, four of the past five MIT Mystery Hunts have had a

scavenger hunt as one of the individual puzzles; regular participants have enjoyed them and have come to expect one in each new event. More importantly, the prize for winning the Hunt is, and always has been, a coin, and the Hunt possesses a "no rules" attitude. To change either of these things would be to destructively tamper with the very identity of the Mystery Hunt. On the other hand, if any specific Hunt contained a backwards-written cryptogram, the following year's Hunt better not have one. Puzzle solvers hate solving the same puzzle twice.

The geographic scope and the time range of a puzzle contest are two other basic, and important, logistical constraints. Puzzle contests run the gamut of geographic boundaries: they may be localized to within a single room or they may sprawl over an entire state. I have personally played in contests that have occurred in the space of a hotel (an informal contest by Henry Rathvon and Emily Cox run during an American Crossword Puzzle Tournament), a college campus (the MIT Mystery Hunt), and the downtown area of a major city (the *Miami Herald* Tropic Hunt). There are also puzzle contests that take place in no geographical space; that is, they occur completely on paper and have no reliance on a specific location, such as *The Random House $10,000 Trivia Challenge*. A relatively recent phenomenon is that of puzzle contests taking place in virtual space over the Internet.

For puzzle contests, there is a direct and rigid correlation between space and time. In short, the more space the contest takes up, the more time it should take. Most of the reasons for this are purely logistical. If a contest takes place entirely within one room, it should last at most five hours; people will get restless and start feeling cooped up. If a contest is spread out over a number of city blocks, it should last at least five hours; it will take people a good bit of time just to physically travel the distances involved. The opposite is true as well: a short

contest can't be spread over a wide area or it will be more of a foot race than a puzzle race; a long contest confined to a small room will incite cabin fever.

A contest with no space constraints also has no time constraints. By its design, such a contest is intended to be solved at home at a leisurely pace. However, unlike the other puzzle contests, this type of event cannot be a race. A fair race must have a uniform starting time, and the only way to accomplish that is to have all of the entrants gathered together at the start of the event. By definition, that is not the case here. *The Random House $10,000 Trivia Challenge* is a good example of this scenario. This multi-puzzle contest was sold in book form at retail bookstores across the country. If it was run as a race, an unfair advantage would be given to people who could acquire the book faster (if, for example, urban stores had the book in stock before rural stores). Book acquisition is obviously not one of the puzzles involved in the contest; no reward should be given for doing it skillfully. A race format would also benefit vacationing or unemployed people; there should be no penalty for having to go to work. Instead, any participant who correctly solved the entire book by a predetermined date advanced to the next round; a series of tie-breaker puzzles narrowed the field of entrants until there was only one left. It is worth noting that a random draw wasn't used to pick a winner; this was a contest based solely on merit.

2.2 Contest Format

The first decision that must be made regarding the construction of a multi-puzzle contest is its basic format. How will the individual puzzles that make up the contest interact with one another to form a single large puzzle? There are two choices: the puzzles can be compiled or cascaded.

In a compilation, each individual puzzle stands alone; they may be solved independently of one another. Generally, all of these individual puzzles are handed out at the beginning of the contest and are solved in a bunch. When enough of the puzzles have been solved, their answers are collected together and used as inputs to a final puzzle called the endgame. The answer to the endgame is the answer to the contest.

In a cascade, each puzzle after the first is dependent on the previous one. The solvers must complete each puzzle before moving on to the next. A contest may be set up as a cascade by either content or format. In a content-based cascade, each puzzle after the first lacks some amount of necessary information that is provided by successfully solving the previous puzzle. With this perpetual transfer of information, each puzzle flows into the next one. This type of contest rarely has an endgame; rather, the answer to the last puzzle in line is the answer to the contest. In a format-based cascade, each puzzle is independent, but the Puzzlemasters enforce the cascade style by handing out one puzzle at a time. The solvers only receive the next puzzle when they have completed the one they were working on. This kind of contest usually has an endgame that ties together the answers to the individual puzzles, but sometimes the last puzzle's answer is the contest's final solution.

The compilation is a great deal more common than the cascade. A series of cascaded puzzles is quite elegant from a puzzle aesthetics point of view: it truly is one large, cohesive puzzle. As such, it is much more difficult to construct. Both the compilation and the cascade are large puzzles made up of a number of smaller constituent puzzles, but, to a constructor, once a compilation's endgame has been made, the contest is nothing more than a number of free-standing puzzles that have nothing, with the possible exceptions of theme and/or a motif, to do with one another. These puzzles may be constructed simultaneously by different people, or they may simply be thought of as a number of small tasks. The cascade can never be broken into small pieces because each puzzle is heavily intertwined with two others (unless it is the first one or the last one). Each puzzle must require enough information from the previous one that it cannot be solved without it, and each puzzle must need to be solved before the information required for the next puzzle can be ascertained.

Another reason it is infrequently seen is that a cascade (either format-based or content-based) is more restrictive regarding playing conditions. It is ill-suited to be a team competition because only one puzzle is being worked on at any time; it is difficult and uncomfortable for ten people to try to solve the same puzzle together. A cascade can really only serve as a contest for single players or pairs. By its very nature, it cannot be very long; a cascaded contest intended to last two days would be practically impossible to construct. (I am loath to state that it is fully impossible, for telling a puzzle maker that he cannot do something is the most enticing invitation possible.) A cascade is best suited to be a short, in-person contest lasting only a few hours, or else a long contest in book form.

Another reason a cascade is rarely seen as a contest is that the danger of getting stuck is too great. In a compilation, there are many different puzzles being worked on at once, so if

a competitor gets stuck on one, he can skip it and work on the rest. Depending on the contest, he may be able to complete the competition without it. In all likelihood, someone else will win, but the competitor in question will have been able to continually work on puzzles until the contest ends. In a cascade, there is only one puzzle on the table at any time, and the solver cannot move on to the next puzzle without completing the current one. If he gets stuck on any particular puzzle, the contest grinds to a halt for him. He can't advance and he can't work on something different, so he'll just wind up staring at the puzzle that is unsolvable to him and getting very frustrated. He won't have any fun, and this is the worst result possible; it is a failure for the puzzle creator. Sometimes this problem is solved by setting up deadlines: if a contestant hasn't solved a puzzle by its specified time, he may receive the answer and the next puzzle in the sequence.

By far, the standard format for multi-puzzle contests is the compilation. Throughout the rest of this paper, it is assumed that all contests discussed are compilations; the cascade is simply too rare. It is not unheard of, however, and sometimes, albeit even more rarely, a hybrid puzzle contest is created that exhibits traits of both formats. The 1998 MIT Mystery Hunt was made up of four sections in a format-based cascade: the sections were independent, but each team had to prove it had solved one section before receiving the next section (or wait until a deadline had passed.) Each of these sections was a compilation-style puzzle made up of a number of individual puzzles and an endgame.

2.3 Theme

Theme, it can be argued, is both the most important and least essential element of a multi-puzzle contest. The theme is the overarching motif of the contest that serves as its narrative backbone. Theme provides a contest with structure, consistency, and familiarity to the solvers. Of course, this is all a side effect – a contest's theme is generally window dressing ancillary to the puzzles themselves. The mechanics of the competition would work equally as well if there was no theme at all.

Many multi-puzzle contests do operate without a theme. *The Random House $10,000 Trivia Challenge* book was a straightforward collection of puzzles. To the extent that it had one, its theme was itself: the concept that tied all of the puzzles together was that they were bound into the same special puzzle book. The same applied to *The Random House $10,000 Crossword Challenge*.

The *Miami Herald* Tropic Hunt is also themeless. Each contest consists of five larger-than-life puzzles, but no motif connects them. The designers of the contest try to make each puzzle a huge concept unto itself, rather than using a single huge concept to connect the separate elements. The 1996 Tropic Hunt contained such outsized puzzles as a rebus cut into pieces and plastered onto the sides of the monorail cars that loop the city, a conceptual puzzle consisting of a few guys in St. Louis Cardinals uniforms having a barbecue on a wooden deck (they were a deck of Cards), and a Eureka!-style puzzle made up of a page in the *Herald*, a bumper sticker, and a live mock-presidential debate. Each puzzle is its own separate world, and no theme is large enough in concept to contain them all. To create such a theme would constrict the limitless enormity to which the individual puzzles are currently allowed to grow.

In a sense, however, there is an unwritten theme to the Tropic Hunt: Miami. Miami is the one thing large enough (barely) to hold all of the puzzles. In addition, a large percentage of the thousands of participants travel into the city from out of town, out of state, and even out of the country. For these nonresidents, their discovery of Miami occurs while they are playing in the Tropic Hunt, so to them, the spirit of the city gives the contest a lot of its flavor. Miami cannot constitute a true theme, though, because it has a greatly diminished effect on those participants who live in Miami's environs. It is only the travelers for whom the experience of the contest and the experience of Miami are fully intertwined. Furthermore, the concept of Miami doesn't fit all of the criteria I ascribe to theme; it isn't a narrative construction.

As demonstrated, multi-puzzle contests can work quite well without a theme structure. The two examples above are atypical of this genre, however. Most such contests consist of a series of pencil puzzles and are run in person, in one place. Although this type of contest can also be constructed without a theme, it shouldn't be. The 1994 MIT Mystery Hunt was themeless; it was just a bunch of puzzles to solve. Although it worked quite well, it was bland. The puzzles were disconnected and might as well have been a random assortment from a puzzle magazine. The 1995 MIT Mystery Hunt had the theme of the board game Clue. The puzzles were set in different rooms of the Boddy mansion (Dining Room, Conservatory, etc.) and were peopled with suspicious characters (Colonel Mustard, Mrs. Peacock, etc.). The puzzles had been constructed by Mr. Boddy, who knew he would be killed, and the ultimate objective was to figure out who killed him, where, and how. (It turned out that Professor Plum did it in the Dining Room with the Revolver.) The difference between the two contests was palpable: the former was stale while the latter created a rich fantasy world the solvers could delve into.

2.4 Using Specific Knowledge

A strategy of use in constructing multi-puzzle contests, or, really, in constructing any puzzle, is to specifically target the skills and interests of the intended solvers. This principle is usually not applicable, though, as most puzzles are intended to be solved by a general public. Puzzles published in puzzle magazines must not require any specific knowledge of the solver; they must be accessible by the entire puzzle-solving community. Puzzles found in niche publications may be more narrowly construed. Mike Shenk's *Wall Street Journal* crossword focuses on business-related themes. The *Journal* is a global paper, so these puzzles, while targeting a specific audience, must still be solvable by anyone who regularly does crosswords. The smaller the audience gets, the more constrained a puzzle may be. Eric Albert, a freelance puzzle constructor, has made cryptic crosswords for *The Scientist*, a journal for life science professionals, in which the grids are completely filled with technical scientific terms. Such a puzzle would sustain heightened interest from the subscribers of the journal, but it would be practically unsolvable by anyone else.

Multi-puzzle contests behave the same way. Events that draw people out of the general populace, such as the *Miami Herald* Tropic Hunt, can't require solvers to know information that isn't widely accepted as general knowledge, hasn't been provided, or isn't readily accessible. Competitions that are intended for a narrower audience base may call for the use of skills and information that the average contest entrant, but not the average person, can be reasonably expected to have. In fact, these competitions are encouraged to include such things. The MIT Mystery Hunt ought to require specific knowledge about MIT, such as its campus and its history. Without this, the contest is generic; why have it at MIT at all? The

narrower an audience gets, the more specific knowledge should be used. If a multi-puzzle contest is held at a gaming convention, there better be some questions about Dungeons & Dragons and Magic: The Gathering.

There are limits to the amount of specialized knowledge required in any one event. For example, a puzzle included in the MIT Mystery Hunt may require highly technical knowledge that would only be known by a major in a specific field. This is acceptable because such majors may be found on campus. But very few of these types of puzzles should be part of a Hunt; most of them should be puzzles that the participants can solve by themselves (or with the help of some reference books).

As to reference books, it is perfectly acceptable for a puzzle to assume that the solver is in possession of certain books. Any puzzle contest may employ information from the Bible or the local phone book; an MIT Mystery Hunt puzzle may reference a specific page in the MIT course bulletin. If it cannot be guaranteed, with utmost certainty, that each participant (or participating team) is already in possession of the exact book in question, then that book should not be an integral part of a puzzle. The 1998 MIT Mystery Hunt contained a puzzle that referenced section numbers in Roget's Thesaurus, a book it can be assumed each team had a copy of. Unfortunately, different editions of the thesaurus were numbered in different ways, so every team that had the wrong version of the book couldn't solve the puzzle.

There is also a limit to the amount of generally known specific knowledge that can be included in any one contest. The 1996 MIT Mystery Hunt was based on the book *Goedel, Escher, Bach: An Eternal Golden Braid.* Obviously, this is not a book everyone owns, but the creators of the Hunt publicized this information well beforehand, so that wasn't a problem. The problem was that, like the title of the book, the Hunt was divided into three categories:

one-third of the puzzles were about math, one-third about art, and one-third about music. Including a few music-based puzzles in a multi-puzzle contest is fine; it's even encouraged for the sake of diversity. But I, as a participant in this contest, found the number of music puzzles to be far too great. Music may be classed as general knowledge, but it is also a specific skill that some people have and others don't. As I have no musical knowledge or abilities whatsoever, I found that one-third of the puzzles were instantly inaccessible to me.

2.5 Using the Geography

Just as a multi-puzzle contest should take advantage of any special skills or knowledge possessed by the audience, it should take advantage of any unique opportunities presented by the geography covered by the contest. In this respect, a puzzle constructor ought to be opportunistic: he should try to use anything afforded by the specific circumstances of the competition he is constructing. A puzzle maker naturally strives to be creative and unique; he wants his work to stand out. A puzzle that requires a team to follow a hieroglyphics trail around MIT's Green Building and McDermott Court, as was included in the 1993 Mystery Hunt, cannot possibly be played anywhere else.

The two biggest annual multi-puzzle contests make great use of the geography they have been blessed with. The *Miami Herald* Tropic Hunt utilizes locations and features all over the city and traditionally includes a map that is key to finding the final solution. In the past, the Tropic Hunt has involved local parks, a soccer stadium, the Goodyear Blimp, and the Metrorail, Miami's downtown monorail system. These features add to the larger-than-life feeling of the contest; it's an experience you just can't get if all the puzzles simply existed on paper.

The best thing going for the MIT Mystery Hunt is MIT itself. According to Eric Albert, in his article "The Great Annual MIT Mystery Hunt" (*Games*, July 1991),

> The Massachusetts Institute of Technology seems to have been designed for treasure hunts. Its main buildings, constructed in the early part of [the twentieth] century, are interconnected in a bewildering maze of passages, skywalks, and tunnels. Halls suddenly slope, or change direction, or stop, hinting at some architectural compromise now long-forgotten. Harshly-lit basements lurk beneath, with dingy subbasements below them. Even the doors are peculiar.

> Some are half-size. Some lead nowhere. Some bear inscriptions such as "Department of Alchemy" or "Shelob's Lair."

This is a fertile breeding ground for puzzles, riddle trails, and other Mystery Hunt components. The 1993 Hunt included a puzzle that led participants to wander around the campus corridors looking for unique MIT artifacts such as a machine called the "Braillemboss" and a piece of artwork named "Elmo-Mit". The 1994 Hunt had a puzzle consisting of twenty-five photos of unique views around campus; each had a letter airbrushed out. One of the puzzles in the 1997 Hunt consisted of a bunch of doctored photos depicting an MIT legend: a game of Tetris played in Building Fifty-four's grid-like array of room lights. The 1996 Hunt went farther afield and sent participants out to one of Boston's public parks. And the endgame to the 1993 Mystery Hunt was, quite literally, the campus itself. (See Chapter 3.1.2 for more details.) The campus presents so many fascinating opportunities that it's a shame when they aren't taken advantage of. This was one of the biggest knocks against the 1997 Hunt: since none of the principal constructors (myself included) lived within two hundred miles of MIT during the year we were creating the contest, very little MIT-oriented material was included in the competition.

"Because it's there" isn't the only reason to make good use of the available geography in a multi-puzzle contest, be it located in Miami, on a college campus, or just in someone's house. There's an extremely important second reason: it allows the participants to move around. The value of this aspect is directly related to the time length of the contest in question. If it's a three-hour event, then the competitors can be reasonably expected to be entertained just sitting at their tables. If it's a weekend-long event, like the MIT Mystery Hunt, than any excuse for leaving the room where a participant has been holed up for two days with

fifteen other people is a blessed relief. Getting out and wandering around the campus is both interesting (think "playing detective") and physically invigorating. Of course, this principle also applies to three-hour contests; the situation just isn't as desperate.

2.6 Puzzlemaster Interaction

The extent to which a Puzzlemaster (PM) should be directly involved in the contest is a tricky judgment to make. Puzzle makers, by their nature, are very competitive people. They thrive off of the power derived from creating a new world and immersing other people in it. In the realm of the puzzle, the constructor is a god. He has created the world and all the rules that govern it. He is the only one that fully understands the rules, the only one that can change them, and the only one who knows all the answers. Rarely can a person wield such complete control over other people who have voluntarily entered this one-sided pact, and it can be intoxicating.

After the contest is started, the PM often has very little to do. All of his work came prior to the start of the contest; once it starts to work on its own, the PM can't keep constructing it. Of course, there are still extremely important tasks for the PM to take care of: there may be errors in the puzzles that must be corrected as they are found and hints ought to be given when a puzzle proves to be too difficult. But both of these tasks require waiting around until one of the competitors needs help; the PM is now in a passive role. He is also stuck in a base of operations (somewhere where he can be easily reached when a problem arises) while everyone else is running around and solving puzzles. Now that the world he created is up and running; now that he is, in this little realm, the god – everyone gets to play except him. He doesn't even get to watch the competitors play, which is extremely frustrating because that's the reason he has created this contest. (This dilemma is only applicable to larger-scale contests. In a three-hour contest, the geography is limited by time (no one can go very far), and the puzzle maker does get the joy of watching people play his game.)

Due to this frustration, the PM often creates ways to get himself directly involved in the contest. This is acceptable, even encouraged, but only up to a point. The person who has worked so hard in the creation of this event certainly deserves to have some fun himself. The competitors often enjoy having some direct contact with the person responsible for putting them through such torturous fun. And the involvement of the PM can open up the possibility of some interactive or experimental puzzles: things that can't be done with just pencil and paper.

A terrific way for the PM's to get involved in a puzzle contest is by confirming correct answers. This is traditionally done in the MIT Mystery Hunt, and it is beneficial to everyone involved. For the PM's, it allows them to periodically hear from each team and be aware of their progress. This helps the PM's assess the individual puzzles (Which ones aren't getting solved?) and the entire Hunt (Will it be completed on schedule?) This interaction lets the PM's know which teams could use help and which puzzles need hints or have errors. The solvers also gain from this policy. They certainly benefit from a better overall contest, as well as from the ability to gauge their progress against other teams, from the knowledge of which puzzles are proving to be overly difficult, and from simply keeping in contact with the people running the show. But the greatest asset to the solvers is certainty: the peace of mind that a finished puzzle has been solved correctly and the awareness that a puzzle they think is complete actually isn't. Without this knowledge, one wrong puzzle answer could destroy a team during the endgame as they search through the entire multi-puzzle contest looking for their mistake.

Too much involvement by the PM can have dire consequences, however. Increased reliance on people heightens the possibility of human error. It means the contest is less free-standing. If a team reaches a point where they need the PM at 6:00 AM, and he's asleep, the

competition grinds to a halt. If the PM is up by 9:00, the team in question has unfairly lost three precious hours. The 1998 MIT Mystery Hunt was guilty of PM over-reliance. The entire contest was divided into four parts; a team would receive a section when they had solved the previous section. This ran the danger of the scenario previously mentioned. In addition, there were a number of puzzles that required direct PM involvement:

- A scavenger hunt that required the PM's to check the items.
- A type of race where a competitor had to run from phone to phone on campus. A PM would tell the racer where to go and how much time he had to get there. When the time elapsed, the PM would call the designated phone. If the racer was there, the PM would tell him where to go next.
- A puzzle reliant on a certain book. The catch was that there were only two of these books, they were kept at the PM's base, and a competitor could only look at one for ten minutes at a time.

These activities, combined with the separate packet handouts and the usual answer confirmations, error checks, and necessary hints spread the PM's out too thin. Even though there were a number of them helping to run the contest, when they were trying to keep track of eight teams running the phone relay simultaneously, it was practically impossible to get through to them regarding any other topic.

This same corps of PM's committed the most egregious sin possible by a Puzzlemaster: they directly affected who won and who lost the contest. During a contest, the role of the Puzzlemaster is that of referee. His job is to keep the contest running smoothly and make sure all of the participants are having fun. He should correct flaws when he finds them, but otherwise he should let the puzzles take care of themselves and keep the competition on an even playing field, to the best of his judgment. In this contest, the PM's were afraid of a blowout win by the first-place team (Setec Astronomy, of which I was a member), so they deliberately slowed them down whenever possible. For example, a certain set of puzzles was

scheduled to be handed out at midnight unless a team could solve enough puzzles to get it early. Setec Astronomy got it at 8:30, giving them a 3 1/2-hour lead. The PM's then changed their schedule from 12:00 to 10:00, cutting the lead by two hours. In another case, they withheld a vital piece of information, without which a team could not advance to the final puzzle, for a number of hours. Setec Astronomy got to the point where they couldn't continue without this clue, got stuck, and waited. The second-place team (Iliaphay) closed the gap entirely, then they too got stuck. Only then did the PM's give out the clue. These lead-cutting tactics not only prevented a blow-out, they cost Setec Astronomy the victory altogether. Iliaphay was able to complete the final puzzle faster and won the event. Just as a referee should not determine who wins a football game, a Puzzlemaster should not determine who wins his contest.

2.7 Experimental Puzzles

Most multi-puzzle contests are informal affairs. They are amateur events created by people who love puzzles and want an outlet for their rampant creativity, and they are competed in by people who love puzzles and desire an interesting challenge. It is a fertile environment for the creation of unique or experimental puzzles. The promise of these innovations is frequently an enticing draw into participation.

Not all multi-puzzle contests should include experimental puzzles. Some are particularly well-suited for them: the *Miami Herald* Tropic Hunt is founded on a sense of wackiness, and the MIT Mystery Hunt, an event run by amateurs, prides itself on its no-holds-barred diabolical nature. More "serious" puzzle contests should shy away from off-the-wall risks. The puzzle extravaganza held during the National Puzzlers' League annual convention is a rather succinct competition. Furthermore, its participating teams are randomly generated from among the attendees, so it is quite common for team members to be on unfamiliar terms with one another. Since the extravaganza creators are certain to be puzzle geniuses, innovative structures are par for the course, but truly experimental concepts are best left out. *The Random House $10,000 Challenge* books, since they were published works intended for mass consumption, rightly avoided anything bizarre.

For a puzzle constructor, the multi-puzzle contest is really the only public format where he can try out the crazy ideas he couldn't get away with anyplace else. By his very nature, a puzzle maker is an exhibitionist of creativity; he dreams up tangled enigmas and gives them to total strangers to play with. The platform of the multi-puzzle contest gives him the opportunity to shed the constricting limitations he must endure when creating puzzles for

publication. (It is worth mentioning that the grand majority of all puzzles that are published are crosswords and word searches, the two most staid types of puzzles there are.) The puzzle maker also relishes the challenge of doing something that either no one else has accomplished or that no one else has even thought of before.

There are two distinct kinds of experimental puzzles: bonuses and innovations. Innovations are the true experiments – these are the contest segments (often puzzles, but not always) that have never been successfully done as part of a contest. They are very risky and have a high degree of failure. Bonuses, on the other hand, aren't risky. They are contest segments that will work, but are bizarre enough that no participant could have expected them. A good example of a bonus was Puzzle #0 from the 1994 MIT Mystery Hunt: it wasn't included in the packet with the rest of the puzzles – it was that day's *New York Times* crossword puzzle. The crossword had been written by the Hunt's author, Eric Albert, and his friend Will Shortz, the *New York Times* crossword puzzle editor, made sure it was printed on the correct day. An example of an innovation was Puzzle #1 from the 1995 MIT Mystery Hunt. This puzzle (which I created) consisted of twelve sandwich bags, each containing three pieces of a certain breakfast cereal; the challenge was to identify the cereals. To make things more difficult, I used limited edition cereals (Christmas Crunch, a variety of Cap'n Crunch) and deceptive selection (i.e., the bag with Alpha-Bits samples only contained the letters "O" and "D" so they looked like misshapen Cheerios). This puzzle was true to form for an innovation: it was a crowd-pleaser (no one had ever solved a puzzle at least partially based on the sense of taste before), but it performed poorly (no one could solve it without help).

Innovations aren't limited to individual puzzles; they often occur on a much larger scale. Past MIT Mystery Hunts have employed experimental puzzles as major structural

elements. In the 1995 Hunt (the theme of which was the game Clue), roughly half of the individual puzzles had secondary answers that were completely separate from their normal answers. Hidden somewhere along the course of solving these puzzles was a Clue item, such as "Mrs. Peacock" or "Billiard Room." One person, one weapon, and one room were left out of the entire Hunt; these formed the final solution to the game of Clue (Who killed Mr. Boddy, how, and where?) and fed directly into the final stage of the endgame. The 1996 Hunt contained 36 individual puzzles that were evenly divided into three endgames – but each endgame required 13 inputs. A series of instructions revealed the locations of the missing inputs: the last input to the first endgame could only be found after solving the second endgame, the last input to the second endgame could only be found after solving the third endgame, and the last input to the third endgame could only be found after solving the first endgame. Thus the puzzle makers had created a paradox and forced the solvers to crack one of the endgames by hacking into it without having all of its inputs.

Using innovations like this is highly risky; if the experiment fails, then a large portion of the multi-puzzle contest's structure collapses . In the 1996 Hunt, if none of the endgames could be reasonably solved with only 12 of 13 inputs, there would have been a big problem. The hidden Clue items in the 1995 contest did cause a problem when teams had trouble finding them; the problem was rectified when I (as the Puzzlemaster) revealed which puzzles contained Clue items and which didn't. Whenever a large part of the structure is experimental, it is wise to expect difficulty and deal with it beforehand rather than hoping it goes well and troubleshooting on the fly if it doesn't. The 1997 MIT Mystery Hunt contained an experimental section in which every puzzle was impossible. The key to this section was to realize that although the individual puzzles were unsolvable, the entire collection was solvable

as a whole: each page contained a letter or a number hidden in the same way and the entire packet spelled out a message. This was a dangerous construct form the start since unsolvable puzzles are virtually guaranteed to frustrate solvers. To alleviate any problems before they had a chance to occur, this section was designated as optional and handed out two hours before the real Mystery Hunt was scheduled to start. Any teams that successfully solved it within the two hours were rewarded with a head start on the real Hunt; any teams that didn't solve it got to advance anyway.

3 The Endgame

When creating a multi-puzzle contest, or even a multi-puzzle puzzle, the constructor must start at the end. This means starting with the final answer and the structure of the endgame.

The endgame is the linchpin in this type of structure. It is the puzzle that takes all of the answers to the individual puzzles as inputs and gives the solver the final solution as its output. It is the unifying puzzle that ties everything together. As such, its importance cannot be underestimated. This puzzle has a palpable emotional weight: it is the catharsis of the contest, the sprint at the end of a distance race. All of the effort a solver has invested of himself into the contest is concentrated into this one last puzzle.

Of course, the role of the endgame in the structure of a contest differs from one contest to the next. In a long, in-person race like the MIT Mystery Hunt, it has an added weight because it is the only puzzle that the participants must solve in order to win. A solver can detour around any individual puzzle by solving enough of the rest, but all roads lead to the endgame. On the other hand, the endgame loses importance due to the nature of this style of contest: since this particular contest is such a marathon, the teams are usually rather spaced out by the end of the event and only one or two teams ever get to work on the endgame. The fastest team solves the endgame (and wins the Mystery Hunt) before most teams even get to look at it.

3.1 Structural Tradeoffs

There are a number of tradeoffs to make when constructing the endgame. Some of the considerations that need to be evaluated regarding the endgame's structure include whether it is:

- Handed out with the rest of the puzzles or only given when needed
- Hidden or clearly labeled
- In one part or in many parts
- Simple or complex

Each choice has its merits and can work well, depending on the overall structure of the contest. A key point is that many of the choices do not work in combination. For example, a complex, in-depth endgame that is only handed out to the solvers when they get up to the point where they need it is a recipe for disaster. A very similar situation to this was built into the structure of the 1998 MIT Mystery Hunt. Each of the four sections of the contest were released in stages, and two-thirds of the way through the event, the Puzzlemasters (PM's) became panicked that only the front-running team would get to see the last batch of puzzles. This would shortchange both the majority of the solvers (who wouldn't be given the opportunity to play with the entire Mystery Hunt) and the PM's themselves (who would have spent days constructing puzzles only one team would ever solve). Given the difficulty of these final puzzles and the time required to solve them, the PM's fears were unfounded. They didn't know that, though, so they unfairly held the first-place team back to give everyone a chance to play with these puzzles. This particular batch of puzzles was not the endgame of this contest, but in this structure it could have been and it serves as a good example of what not to do.

3.1.1 Handed Out vs. Withheld

The example given in the previous section explores the possible dangers of withholding the endgame until the solvers need it. This is not always a dangerous option, however. The actual endgame of the 1998 Mystery Hunt was withheld from each team until they reached the point where it was the only puzzle left, and this structure was implemented quite well. The Puzzlemasters were aware that a maximum of two teams would ever get the chance to solve the endgame, so they constructed a very simple one. This was an effective approach: the teams that never received it didn't miss anything exciting and the puzzle constructors were able to devote their time and energy towards the puzzles that every participant would have a chance to enjoy.

As shown, keeping the endgame from solvers until they need it may work out nicely or it may backfire, depending on the other parameters of the endgame. There are two other options regarding the timing of endgame distribution: withholding it until a specific time, as the *Miami Herald* Tropic Hunt does, and simply including it with all the other puzzles at the beginning.

The more popular choice in this matter in recent MIT Mystery Hunts has been to include the endgame (or at least the first part of it) with the rest of the individual puzzles. The advantage to doing so is related to the rule of thumb that dictates that a solver should be able to complete the Hunt after solving about 80% of the individual puzzles. Since the solver has the endgame in front of him, he can work on it while he is completing the general part of the Hunt. Each new answer produced by the individual puzzles completes the endgame puzzle a little more, until eventually the solver can work around the gaps and solve the whole thing.

Since he is continually working on it, the solver, by definition, completes the endgame with the fewest number of individual puzzle answers that he could. By not needing all of them, the solver avoids the headache of staring at the puzzles he is stuck on. In this structure, in direct contrast to a Hunt in which the endgame is only given out when needed, a simple or transparent puzzle makes a poor endgame. Since the solver is working on this puzzle for days while the individual puzzles are being solved, the format lends itself to a complex endgame. In addition, the endgame should have a high enough difficulty level such that the fewest number of individual puzzle answers needed to crack it is around 75 to 80 percent of them.

The *Miami Herald* Tropic Hunt's endgame is always withheld until a predetermined time, at which point it is released over the radio. (It is repeated a few minutes later to make sure that everyone gets it.) The reasons for this are quite practical. In the very first Tropic Hunt in 1984, puzzle sites were spread out from Palm Beach to South Miami, 80 miles away. In this contest, there was no waiting for the endgame: all the pieces of the puzzle were readily available and the first team to reach, and solve, each part would win. This had the detrimental effect of encouraging speeds upwards of 100 miles per hour as the players tried to reach each site as fast as possible. The two-tiered approach was implemented in the second year to make reckless speeding unnecessary. It didn't matter whether a team finished all of the available puzzles by 1:00 or by 4:00; since they'd still have to wait for the radio message, they might as well not risk their lives on the road. As time passed, the geographic scope of the contest shrank: in the 1996 Tropic Hunt, which I participated in, each puzzle site could be reached by foot or by monorail from any of the other puzzles. Even so, the withheld endgame was retained as standard Tropic Hunt structure.

3.1.2 Hidden vs. Clearly Labeled

Obviously, this distinction applies only to those endgames distributed with the rest of the puzzles. In some contests, such as the 1994 and 1996 MIT Mystery Hunts, the endgame is clearly labeled. In others, such as the 1995 and 1997 MIT Mystery Hunts, the endgame is hidden. In rare cases, as in the 1993 MIT Mystery Hunt, the endgame (one-half of it, in this case) isn't included with the other puzzles at all because it is implicit.

The choice here is mainly a matter of taste for the constructor. A labeled endgame, complete with instructions, is gratifying to the solver because it serves as an anchor for the entire contest. In the 1994 MIT Mystery Hunt, the endgame was a series of instructions that took the solver around the MIT campus and led to the location of the coin. A number of important words in these instructions were excerpted and replaced by different words, thus making the trail impossible to follow. Each individual puzzle answer restored one word in these instructions. There was some ambiguity involved, so a team needed to have a bunch of puzzle answers before the final trail started to make sense, but after that point, each additional puzzle that was solved made the trail one step clearer. After roughly three-quarters of the puzzle answers were attained, the trail was clear enough for teams to start sending parties out to try to follow it. (It took a greater number of answers to follow it successfully, though.)

By contrast, the 1995 Mystery Hunt endgame (or, more accurately, the first piece of it) was disguised as one of the individual puzzles. This was a breakthrough point: if a team made the mental leap and identified the endgame for what it was, they could continue. If not, they couldn't. This puzzle wasn't designed to grind teams to a halt; it was just a different method of reaching the recommended number of 80% of individual puzzles solved. As teams

had fewer and fewer unsolved puzzles remaining, the endgame would undoubtedly be among them. The increased attention spent on it (since there weren't many puzzles left to focus on), the extended time spent on it already (since it had been there from the beginning), and the pile of puzzle answers that team's didn't know what to do with were intended to combine into a scenario conducive to figuring out the true nature of the endgame puzzle. This construct proved to be a bit too tricky in execution: only a handful of teams ever found the endgame, meaning the rest were frustrated by their inability to do anything with the individual puzzle answers they had accumulated. It is worth mentioning that the endgame was designed to be solvable with about 80% of the other puzzle answers in hand, so if a team discovered the endgame immediately, the rule of thumb would still hold.

The 1993 Mystery Hunt had a rather unique endgame structure. It had two parts: one part was a cryptogram that decoded to be a path leading from an unidentified starting location to the hiding place of the coin, and the other part – the unique part – revealed the path's starting point. This second part was hidden, both in the sense that it wasn't explicitly identified and that it had a Eureka! quality to it: it required a mental leap to figure it out. It was different from all other such endgames in that it had no external structure to disguise; the endgame consisted of the individual puzzle answers and nothing else. The answer to each individual puzzle or puzzle cluster was a path encompassing part of the MIT campus or, in one case, an entire MIT building. The endgame was to realize that these paths (and the building) spelled something when they were all drawn out on a campus map: they formed the room number "2-105", the correct place to begin the final trail. In this startlingly elegant construction, the endgame could not be puzzled out by working on the framework the individual puzzle answers fit into; here, the individual puzzle answers *were* the framework.

3.1.3 One Part vs. Many Parts

The choice involving whether an endgame will consist of one part or of many parts is closely related to the choice of whether it will be simple or complex. It's not the same issue; a one-part endgame may be quite involved, while an endgame with five different steps may consist of five very straightforward steps. Usually, though, a simple endgame has only one part and a complex endgame has a number of them. Even so, the two issues will be discussed separately, as a different set of considerations affects each of these two decisions.

A one-step endgame is a clear concept to grasp. An endgame with more than one step can take on many different structures, though. The 1995 MIT Mystery Hunt's endgame was a series of five cascaded puzzles. The 1996 MIT Mystery Hunt had three concurrent endgame puzzles whose answers tied together (and led to some more puzzles). The 1998 MIT Mystery Hunt was made up of four sections, each one a multi-puzzle contest (complete with endgame) all its own; the overall Hunt had an endgame that tied together the answers from the four sectional endgames.

The reasons for designing an endgame with numerous steps are different for each structure. The effectiveness of each structure is equally varied. In the 1995 Hunt, the decision to create such a long, cascaded group of puzzles was driven by a desire to have a complex endgame (which is explored in the following section) and to be a reaction against the previous year's endgame, a one-step, clearly labeled puzzle. It was an elegant structure, but it fell at the end of the longest Mystery Hunt ever. By the time the leading teams started in on the endgame puzzles, they had been working for days; they were exhausted. In addition, people started dropping out late Sunday evening because they had work the next morning, which left

these tired teams depleted of vital personnel. It was very frustrating for each team to crack what they hoped was the last puzzle – the culmination of days' worth of labor – only to find, instead of the final answer, yet another puzzle.

The structure of the 1996 Hunt's endgame was dictated by the theme: *Goedel, Escher, Bach: An Eternal Golden Braid.* The Hunt was divided into thirds: twelve of the thirty-six individual puzzles were math-related, twelve had to do with art, and twelve were musical in nature. Each subset of twelve puzzles had its own endgame. This was a well-designed structure, both because it was consistent with the theme in a heavily theme-driven Hunt and because it allowed a team to simultaneously work on a solvable endgame section and on the remaining individual puzzles (targeted for a different endgame section) at the end of the contest. Beyond this point, the Hunt suffered the same ills the 1995 Hunt did: solving the endgame puzzles only led to more endgame puzzles, and this took its toll on the exhausted teams.

The overall structure employed in the 1998 MIT Mystery Hunt had its problems. For instance, one of the main reasons to participate in this contest is because, from a content perspective, it is the largest single event made up of fully-interlocking puzzles found anywhere. When it is divided into four separate pieces, it loses some of this cachet. The structure also had its benefits, though, and the greatest one was the opportunity for every participating team to experience the joy of solving an endgame. Normally, unless the endgame in a competition is withheld until a predetermined time or distributed with the rest of the puzzles and clearly labeled, most teams will never get to work on solving an endgame. Furthermore, unless a team happens to be the winner of a multi-puzzle contest, its members will never get to solve an endgame. By placing endgames at the one-quarter, one-half, and

three-quarters points of the contest, as well as two at the end (one for the last quarter and one for the entire contest), the Puzzlemasters guaranteed that each participating team would solve at least two endgames in the course of the Mystery Hunt.

The alternative to multiple endgames is one endgame, and the main benefit of this is straightforwardness. There is a pristine elegance in employing an uncomplicated ending to cap off a grueling mess of puzzles. After all, one endgame puzzle is all that is truly needed. The main drawback here is that such a structure may be too straightforward; people don't enter multi-puzzle contests because they want to avoid mental labyrinths and obfuscation.

3.1.4 Simple vs. Complex

The decision whether to make an endgame simple or complex is the most basic choice regarding its structure; everything else follows from this outcome. It is also not a black and white decision; there is no border line where simple ends and complex begins. Like the rest of the tradeoffs, the "right" choice is a subjective matter. For instance, a thirty-minute endgame to cap off a forty-eight hour puzzle marathon may be a terrible anticlimax or a heaven-sent relief. But "simple" doesn't necessarily imply easy; rather, it implies that the endgame is straightforward.

The 1995 MIT Mystery Hunt endgame, which I constructed, was anything but straightforward. The first step was a double crostic disguised as one of the individual puzzles. Half of the individual puzzle answers were used here, and the resultant message gave solvers some instructions and the location of another puzzle. This new puzzle was a code that, when solved, gave detailed instructions of how to use the rest of the individual puzzle answers to pinpoint specific hallways on the MIT campus. The room numbers in these hallways were then circled in a number search puzzle. When that was completed, the grid was to be played as a maze (with the circled numbers acting as walls); tracing a path from start to finish delineated a numeric string. This string was decoded using a separate semi-endgame puzzle (it used information – but not actual answers – from about half of the individual puzzles). Once decoded, it spelled out a phone number. This was the number of a centrally-located campus pay phone that the coin was inside of. One of the drawbacks of creating such a complicated structure was that no one understood it except me, so I had no choice but to act as the sole Puzzlemaster. This was a great hardship not only on myself but on all of the teams who,

during the running of the contest, needed to contact me but couldn't get through the communications bottleneck. Another drawback was that the teams that entered this stage were very tired and this structure drained their already-depleted energy as it just kept going and going. Finally, although five teams started the endgame (a relatively good number), only two finished it; the man-hours spent constructing the latter endgame puzzles could have been used to make puzzles every team would get to play. What it had going for it was that, from a puzzle aesthetics standpoint, it was a gorgeous piece of layered intricacy. Is that justification enough? At the ungodly hour of 4:30 A.M. Monday morning, 62 hours since the start of the contest, five different teams were still playing, so something must have been working well.

The 1996 Mystery Hunt had a complex endgame structure, as described in the last section. But each of the three main endgame puzzles inside the structure was complex in its own right. The nature of two of these puzzles was somewhat misguided in the first place: one was made up of thirteen paragraphs encoded in a complex cryptographic algorithm, and the other consisted of thirteen paragraphs encoded in a complex anagrammatic algorithm. The only way to decode them was with adequate computer programming skills and the ability to type in pages of gibberish without a single typo. Granted, computer programmers aren't rare at MIT, but if a team didn't already have one, they had the additional problem of trying to recruit one. The third puzzle was musically oriented and awfully subjective; it was too complicated to be test-solved and, as a result, wound up being impossible to solve. The PM's had to scrap this endgame – the culmination of one-third of the entire Hunt – and give out the answer. (See Chapter 4.5.3 for more details.)

This brings to bear the main problem with complex endgames (and a general problem with any endgame, really): they're very difficult to test solve. The problem isn't so much that

the endgame puzzles are that tough; it's that the situation is so hard to recreate. It's impossible to predict which individual puzzle answers teams will have and which they won't. Obviously, no one can test solve an entire multi-puzzle contest unless it's designed to only last a few hours. Therefore the endgame, if it's tested at all, must be tested in conjunction with assumptions about the solvers' actions, which can lead to a world of trouble. For various reasons, every MIT Mystery Hunt from 1994-1998 has had some sort of error in the endgame. Only once (1995) was the problem corrected before any teams reached it.

This is not to say that simple endgames are problem-free, but they are certainly easier to check. This is also not to say that they are preferable; it depends on the situation. A three-hour multi-puzzle contest must have a simple endgame; there isn't enough time allotted for anything else. The 1993 MIT Mystery Hunt had a simple, elegant endgame. The individual puzzles combined to yield a room number and a (very tricky) cryptogram gave the path to follow from that room. This straightforward structure proved to be wonderfully effective. By contrast, I, as a participant, was disappointed by the 1994 Hunt's endgame. Like the rest of the Hunt, it was effective, but I felt it was too simple. I had come to expect more complexity, more deviousness; that was the reputation of the MIT Mystery Hunt. Judged on its own standards, it was a good puzzle; it just wasn't what it should have been.

Ultimately, the question of simple endgame or complex endgame comes down to consistency. The 1993 and 1994 MIT Mystery Hunts were straightforward and they had straightforward endgames. The 1995 and 1996 Hunts were twisting, labyrinthine constructs and their endgames followed suit.

3.2 Construction Guidelines

The endgame is the heart and soul of the multi-puzzle contest. Without it, the contest is merely an assortment of puzzles. While the theme is the external structure that ties the elements of the multi-puzzle contest together, the endgame is the internal structure that melds the disparate pieces into a cohesive whole. As such, the key to constructing a multi-puzzle contest is the construction of the endgame – after that is complete, all that's left are the individual puzzles, and these may be created individually. They may be made by different people; they may be copied from puzzle magazines; they may have nothing to do with one another.

What traits should the endgame have? As postulated in the preceding section, the appropriate endgame is one that fits into the context of the rest of the multi-puzzle contest and, with consistency of form and style, brings the event to its conclusion. However, although this is true from the puzzle solver's perspective, it is completely backwards from the point of view of the puzzle maker. The contest designer must first construct the endgame, then, rather than ensuring that the endgame is consistent with everything else, he must construct everything else to be consistent with the endgame.

Therefore, structurally, the endgame should be a microcosm of what the puzzle maker expects the contest to be as a whole. The puzzle constructor is aware of the external constraints facing him, which may include the geographic scope of the contest, its length in time, the number of people who will be participating, and the size of each team. Any of these parameters that have not been already defined by outside sources must now be determined by the puzzle maker. The resultant set of parameters, combined with the puzzle maker's personal

preferences and abilities, will loosely define the structure of the entire contest; this in turn will loosely define the structure of the endgame.

Once the puzzle maker has a basic conceptual notion of the overall contest and of the endgame, his next step is to start casting the endgame's actual form. There are two things that are necessary before actual puzzle construction may begin: a more rigorous understanding of the structure involved and the puzzle's answer. For the former consideration, the puzzle maker must make decisions regarding the structural tradeoffs of the endgame, although this is usually more a matter of the puzzle maker being aware of the type of puzzle he wants to create than a process of consciously working with a checklist of possibilities. For the latter, the puzzle maker needs to determine the final answer to the entire contest.

Determining the final answer is a simple but non-trivial task. Its main purpose is proof of puzzle completion, so the final answer can be practically anything. Depending on the contest, this is usually either a simple answer (such as a word, phrase, or number) that is presented to the Puzzlemasters, an instruction or set of instructions the solver must complete, or a geographic location the solver must go to. But just because the answer can be anything doesn't mean it should; the choice shouldn't be arbitrary. A random final answer ensures that it cannot be guessed, so only a team that deserves to win will happen upon it – but if it is completely disjointed from the rest of the competition, it is an anticlimactic and unsatisfactory ending. The final answer should be something directly related to the rest of the contest (such as a phrase or, better, a pun with thematic relevance) or something wacky enough to be fun in its own right. An example of this is the answer from a small contest designed by Emily Cox and Henry Rathvon and run at the 1996 American Crossword Puzzle Tournament. Not exactly a multi-puzzle contest, it was more like a single puzzle broken up into little pieces and

strewn about the hotel where the tournament was held. After teams ran around, pieced the puzzle together, and solved it, they were instructed to play a march for the judges... on kazoos. Thus a couple hundred of the country's most prestigious members of the puzzle community wound up parading around, kazooing the assorted works of Sousa, and having a tremendous amount of fun. (A bonus to this construct was that the judges were able to award prizes to both the team that solved the puzzle the fastest and the team that kazooed the best; usually, if your team isn't made up of four crackerjack solvers, you're assured of not winning a prize.) The MIT Mystery Hunt always has a random final answer: the location of the coin is purely geographic and nearly impossible to tie into the theme. But the answer is often bizarre and, therefore, worthwhile. Past Mystery Hunt coins have wound up taped to a fire door that slid out of a wall, inside a floppy disk inside a computer, and taped to the underside of a drawer of fossilized worms. Such answers can go too far or be poorly designed, though: one coin was in a wallet checked at the desk of MIT's athletic center, but it could only be recovered during the hours the desk was open; a coin taped to the wall of an elevator shaft resulted in one team occupying the elevator for an hour while the second team to figure out the location could only stand outside the shaft and wait; and a coin placed inside a fish inside a plastic bag labeled "Biohazard" was a bad idea from the start: opening a bag so labeled should never be asked of anyone.

Once the endgame's answer and structural format have been determined, the endgame puzzle itself may be created. The final consideration, and one that must be kept in mind during the puzzle's construction, is what kind of answers the constructor wants for the individual puzzles. By this point, the number of individual puzzles has been determined, so the number of inputs into the endgame has been set. What the inputs look like, though, is up in the air.

Sometimes the answers' form is dictated by the specific multi-puzzle contest: the answers to the Tropic Hunt's individual puzzles are always numbers. Usually, though, the puzzle constructor has free reign. It should be kept in mind that some answer forms are safer than others: if each puzzle yields a common English word but a solver winds up with "BAJKET", he knows he's done something wrong; if each puzzle yields a zip code and a solver gets "02138" instead of the intended answer of "02318", he may not realize he's made an error until it's too late. There are ways around this problem; for example, the MIT Mystery Hunt Puzzlemasters allow participants to call them and confirm whether or not their answers are correct. In the Mystery Hunt, a puzzle may have a nonsensical answer because a solver can immediately determine, by a phone call, if he is right or if he has made a mistake. Even so, answers that are meaningless alphanumeric strings are frowned upon for their inelegance; they'll show up every now and again, but the majority of individual answers are common words or names. It is much more satisfying for a solver to be certain he's solved a puzzle correctly simply upon finishing it without having to ask the puzzle's creator.

On a related note, the individual puzzle answers don't have to all belong to the same set (it's not the case that they all have to be common words or else they all have to be cities, or numbers, or pathways, etc.). The answers usually do belong to the same set, though. One reason for this is that such consistency is elegant. A more compelling reason is that a puzzle's inputs will naturally be quite similar to one another. Through the course of constructing the endgame, the individual puzzle answers will be formed, and as they are all members of the same group (endgame puzzle inputs), they will likely have the same properties and the same form. Dissimilar individual puzzle answers are usually inputs to two different endgame puzzles.

What, exactly, should an endgame puzzle look like? This is left to the individual creativity, construction skills, and whims of the puzzle's creator. It can be any type of puzzle that requires inputs, but no particular puzzle (double crostic, riddle trail, word search, etc.) is necessarily any more effective than any other. Of course, it need not be a standard puzzle form; it can be something completely new. In the end, it doesn't matter what kind of puzzle it is – the constraints and parameters that define the puzzle are more important than the particular genre it happens to be. The puzzle maker must pick or design a form that fits his criteria (keeping in mind that many will); the rest of the multi-puzzle contest will follow.

4 Individual Puzzles

When making a multi-puzzle contest, the creation of the individual puzzles is where the real creative freedom comes in. There are very few hard and fast rules to follow. There are guidelines, of course, but they foster, rather than impede, imagination and innovation. A key to creating a successful contest lies in striking a healthy balance between consistency and variety, and a good understanding of these guidelines will bring that about. The relevant issues are as follows:

- Group characteristics
- Individual characteristics
- The importance of variety
- The necessity of equally weighted puzzles
- Quality control

The first two categories explore the few general characteristics that are necessary to the individual puzzles: first, as a whole, and second, as separate entities. Variety in the individual puzzles, like theme for the overall contest, is an optional but strongly recommended (and very powerful) aspect of a multi-puzzle contest. The need for the individual puzzles to all be worth the same amount in the overall structure is an issue centered at the point where these puzzles meet the endgame; it is the individual puzzle answers that must be equally weighted. Finally, enabling a high standard of quality can only make a contest better.

4.1 Group Characteristics

The entire body of a multi-puzzle contest, with the exception of the endgame, is a collection of individual puzzles. Since this collection can be treated as a whole, it is worth exploring. What characteristics are necessary for this group? What traits does it need to possess? Mainly, the group needs to take the solvers the right amount of time. The intended length of the contest and the number of people on each team will have been predetermined, so the puzzle group must be designed to take the fastest team an acceptable amount of time. Other than that, there really aren't any traits the group must possess. The individual puzzles can be a random amalgam of completely unconnected puzzles and the overall contest would work just fine.

The reason it would work is because of the presence of a theme. The theme's main purpose is to provide a solid background that each of the individual puzzles is nestled into. As stated in Chapter 2.3, a multi-puzzle contest can still function without a theme. In that case, a group puzzle consistency goes a long way towards creating a cohesive contest. The 1994 MIT Mystery Hunt was themeless, but its puzzles were all straightforward, no-frills affairs that truly felt as if they belonged together. The same principle applies to a *Miami Herald* Tropic Hunt: it's themeless, but its puzzles all share a wacky, larger-than-life sensibility. If a multi-puzzle contest has no theme and no individual puzzle consistency, it will work structurally, but it will lack a layer of interest vital to a form of entertainment.

Of course, even in a themed multi-puzzle contest, the individual puzzles shouldn't necessarily have nothing to do with one another. The overall contest can only be bolstered by a pervasive consistency of form throughout the individual puzzles. The 1997 Mystery Hunt

had the theme of Elvis's exploits as detailed in a supermarket tabloid, but it also had a motif of exploiting the things solvers take for granted about puzzles. The Hunt started with a bonus section of puzzles that were impossible to solve. In the main section, individual puzzles included a crossword that broke a cardinal rule of crossword design (the answers were contained in the clues, i.e., "Rarely seen" was the clue for "rare"), a number sequence that was a large palindrome (not a sequence at all), and a cryptogram made up not of words but of spelled-out letters ("aitch" and "double u", e.g.) which then had to be re-decoded. Only a few of the individual puzzles behaved this way, but they made the entire Hunt tighter.

If a collection of individual puzzles does possess any group characteristics, they should be derived directly from its endgame. From a solver's point of view, any consistencies shared by the bulk of the contest shouldn't be abandoned at the climax of the event. If a motif is conceptualized before puzzle construction begins, it should be included in the endgame. Otherwise, since a puzzle maker must design the endgame first, its traits are the ones that become included in the individual puzzles. The puzzles of the 1994 Hunt, as mentioned, were straightforward; its endgame was an uncomplicated one-page puzzle. The puzzles of the 1997 Hunt, as mentioned, played with solver's expectations; its endgame was made up of two puzzles masquerading as normal puzzles (one in the regular Hunt and one in the impossible section) and hints hidden, among other places, in the jokey answer page to the impossible section.

4.2 Individual Characteristics

Like the individual puzzles as a group, there really are no stringent characteristics that these puzzles must adhere to on their own. They are, after all, *individual* puzzles. This is where a puzzle maker can let his imagination run wild and create any puzzle his heart desires. As discussed in Chapter 2.7, bizarre or experimental puzzles are encouraged in this platform.

Creating the individual puzzles isn't a total free-for-all, however. By the time the puzzle constructor reaches this point, he must already know the general logistical constraints that form the boundaries of the contest. The construction of the endgame has determined the quantity of the individual puzzles. The constructor should also know the approximate length of time the entire contest is expected to take and the number of people on each team. These factors determine the average difficulty level of the individual puzzles.

The average difficulty level is important to keep in mind during puzzle construction, but not every puzzle should try to meet it. It is only intended to be an average: lots of puzzles should be more difficult; an equal number should be easier. It is important to note that the difficulty level of a puzzle often, but not always, translates directly into the amount of time needed to solve that puzzle. Ideally, this is the case. Sometimes, a puzzle will have a relatively simple concept but take a disparate amount of busy-work time to solve. A good analogy to this is a third-grade mathematical word problem (and how much fun it is to solve). This type of puzzle ought to be avoided.

While there aren't really any rules dictating what the individual puzzles should be, there is one regarding what they should not be: no two individual puzzles should be the same kind – unless, of course, that is a feature of the contest. An all-crossword multi-puzzle contest

should, of course, contain nothing but crosswords. In a more normal contest, each of the individual puzzles should be different. Granted, this is not always possible. Some MIT Mystery Hunts have contained over 35 individual puzzles, at which point coming up with new genres is a real strain. In such a case, the inclusion of a couple of the same puzzle type is acceptable, as long as each one is a separate variation on the form. For example, the same contest could include a cryptogram in which the message is written backwards and a crossword puzzle whose clues are cryptographic, but further uses of the same puzzle form start to get redundant.

Other than that, there are no necessary characteristics for the individual puzzles. There are a number of suggestions and guidelines regarding related topics, but these generally affect just a couple of the individual puzzles in any given contest. For example, some of the individual puzzles should take advantage of any available interesting areas of geography, but not all of them should, and none of them have to. These considerations are discussed in Chapters 2.4-2.7. There are also a couple of principles that apply to individual puzzle construction; the importance of variety is the topic of Chapter 4.3 and the need for all individual puzzles to be worth the same amount is addressed in Chapter 4.4.

4.3 Variety

In a multi-puzzle contest, it is true that no single one of the individual puzzles is tremendously important, especially when compared to the endgame. Still, the individual puzzles make up the bulk of the contest. The endgame is the rich, tantalizing dessert; the individual puzzles are the actual dinner, and that's what the solvers fill up on. Which is more enticing: a twelve-course meal with a wide range of cuisines, or a meal consisting of potato salad, baked potatoes, mashed potatoes, scalloped potatoes, and french fries?

The key to the individual puzzles is variety. This is a totally optional key, however, and the principle underlying the importance of variety is quite similar to the one regarding theme (Chapter 2.3). A multi-puzzle contest can work just fine without variety, just as an all-potato dinner could satiate a diner. But the contest is intended as a form of entertainment, and entertainment shouldn't be repetitive or boring.

A basic application of variety in a multi-puzzle contest is in the types of puzzles that it contains. A high-quality contest will have as many different genres of puzzles as can be included. The 1995 MIT Mystery Hunt, for example, contained word, math, logic, trivia, pop culture, visual, jigsaw, music, and geography puzzles, as well as a maze, a three-dimensional dodecahedron puzzle, a couple of puzzles on videotape, a riddle trail around the MIT campus, a scavenger hunt, and a puzzle in which teams were given pieces of breakfast cereals to identify. There was something for everybody; any solver was free to concentrate on his favorite type of puzzle while ignoring the kinds he disliked. Even people who weren't puzzle solvers got to play as they helped their teams in things like the riddle trail and scavenger hunt. In an open event like this, the point is to entertain as many people as possible.

Of course, there are always constraints on the individual puzzles section: the geographic scope, time limit, intended audience, and other such parameters set limits that the contest must abide by. But within these bounds, the puzzles in the contest ought to cover every square inch of territory possible. For instance, the puzzles should range over the entire skill level. The constructor needs to focus some puzzles at the level of difficulty the overall contest is trying to attain, but there should always be a range of easier puzzles (as warm-ups or for people who aren't expert solvers but want to play) and an equal number of brain stumpers (since this type of contest attracts puzzle aficionados, and they relish a challenge). Puzzles that take advantage of the geography and the skills of the audience as well as experimental puzzles are encouraged, but not all puzzles should do this: some should, some shouldn't.

Variety should be instilled in even the most tightly constrained contests. If a multi-puzzle contest is designed to include only double crostics, there can still be a wide range of difficulty. In addition, the puzzles can still all be different even if they're all the same: perhaps one is written backwards while another has no spaces; maybe one puzzle's clues can be encoded in a cryptogram while another's are cryptic clues; maybe some of the puzzles have themes, like movies or geography. In a mix like that, even a straightforward puzzle is different. Every solver has unique interests and abilities; variety ensures the maximum number of fully satisfied, amply entertained participants.

4.4 Equal Weight

After the number of puzzles and the expected difficulty level of each, the most basic constraint involved in individual puzzle design is that all puzzles must be equally weighted in importance. No single one of the individual puzzles can function as a linchpin; that's not only poor design, it's a sure way to court disaster.

The individual puzzle portion of a contest must work in one of two ways: either the solvers need the answers from every single puzzle to succeed, or they need a general percentage of them. The difference lies in the scope of the overall contest. If the contest is short (three hours), small (five puzzles), or incredibly long (*The Random House $10,000 Challenge* books each gave solvers a few months to complete them), then the victor should, necessarily, be someone who has solved all of the challenges laid before him. If one of these puzzles is more important or disproportionately difficult, then the contest is no longer about who can solve all of the puzzles; it is about who can solve that specific puzzle. The rest are secondary, and that subverts the whole purpose of a multi-puzzle contest.

The MIT Mystery Hunt, and contests like it, behave differently than the type described above: they need to be set up so a team that completes approximately eighty percent of the individual puzzles can hack out the rest and arrive at the final answer. This is necessary because it is a long, difficult amateur event. The creator must assume that some of the puzzles will be fatally flawed and need to be skipped. A more likely circumstance is that a few puzzles will be somewhat flawed. Certain teams, based on the strengths and areas of expertise of their members, will be able to overcome a puzzle's poor construction and solve it anyway. But while these teams will be strong in one area, they will be weak in another; they will be

stumped by a different puzzle – one that other teams will have been able to solve. It is both unfair and unwise to assume that all teams can solve all puzzles, or even that all teams can solve the same puzzles.

This is where the issue of equally weighted puzzles comes into play. If any specific one of the individual puzzles is necessary to winning the contest, then the creator runs the risk that this puzzle will stump some teams and be solved by others. If this necessary puzzle is very hard, requires specific expertise in a skill some people have and some find completely foreign (such as music), or is flawed, then it is entirely possible that a team could solve every other puzzle in the contest and be unable to win. At the same time, a different team could solve 75% of the puzzles, and this one, and be able to solve the overall contest. If this is the case, then victory is predicated on chance, not skill; in order to win, a team would have had to solve the "right" puzzle. An unfair construction such as this is sure to draw the ire of competitors who rightly feel that they've been gypped.

Sometimes, however, the construction of the endgame dictates that there must be a general puzzle that is necessary to solve. Having a disparately weighted puzzle answer is acceptable as long as the puzzle maker is aware of the situation and handles it acceptably. When constructing the 1995 MIT Mystery Hunt, I had to deal with this. The endgame puzzle in question was a double crostic. The answers to the general puzzles were the entries that were rearranged to form a message that instructed the solver what to do next. A double crostic is an effective way to combine the individual puzzle answers: each entry is worth about the same as the others and the answer message is readable after 75-80 percent of the entries have been filled in. This message had a snag, though: there were numbers in it. The numbers were crucial to applying the instructions; they were MIT building numbers. The message read,

in part, "To advance you must overlay US and MIT maps. Match east vertex of 34 with Pulaski, TN; SE corner of 2 with Rock Rapids, IA; and point 66 to Bad Axe, MI." A solver that had the letters to the entire message, but not the numbers, would be completely stuck.

Exacerbating the problem further, all five digits in the message were in the same double crostic entry; that is, they were all in the answer to one puzzle. This was a necessary consequence of the double crostic construct. It was important for each double crostic entry (individual puzzle answer) to be a common English word, phrase, or name. (The fact that not all of them were was a separate problem.) Were the five digits to be split up and incorporated into separate entries, those entries, each consisting of roughly seven letters and one digit, would be awfully strained as recognizable phrases. The puzzle was much more elegant if the digits were kept together as a single entry.

There was another facet of elegance to keeping the digits together as one entry: the five-digit numeric string would look like a zip code. Since many other puzzle answers in this Mystery Hunt were city names or zip codes, this puzzle answer blended in nicely.

Still, this left the problem of a necessary puzzle mixed in with the individual puzzles, and that was a dangerous ploy. It was crucial that no team get stumped on this puzzle; therefore, this puzzle had to be the easiest one in the entire contest. Even that was a risky solution; often, a puzzle that is intended to be (and expected to be) easy winds up being much trickier than expected. A puzzle constructor is oftentimes blind to the difficulty of a puzzle. It isn't because he knows the answer to a puzzle – it's because he knows the solving technique. In this situation, it can be difficult for the constructor to grasp that the mental concept that seems so intuitive to him may be quite elusive for someone else. When a contest creator needs an easy puzzle, he must take great pains to ensure that it is, in fact, easy.

My solution to this problem was to appropriate a previously published puzzle out of a magazine. That puzzle was made by a professional, modified by editors, and deemed worthy of nationwide publication – it was guaranteed to be a high-quality, solvable puzzle. (The issue of puzzle plagiarism is explored in Chapter 4.5.2.) I adapted it to give me the answer I needed and, to make it feel new, changed some wording to include up-to-the-minute pop culture references. I left the actual puzzle mechanism unchanged. This strategy worked – the puzzle was one of the first three that each group solved.

4.5 Quality

It isn't enough that a multi-puzzle contest has a bunch of individual puzzles. It should have good individual puzzles. Quality control is important to any creative endeavor, and puzzle making is no different. One way to ensure high quality during the running of a contest is for the Puzzlemasters to confirm correct answers (as discussed in Chapter 2.6) Three other ways are as follows:

- Test solving
- Plagiarism
- Fixing errors

Test solving and appropriation of previously published puzzles are both methods of quality control that may be enacted during a contest's construction stage; fixing errors is a crucial skill for maintaining high standards after a multi-puzzle contest has already started. While both test solving and error fixing are standard practices that are universally applicable when making puzzles, the plagiarism is not. Plagiarism is a moral issue that is usually frowned upon; only in certain circumstances is it practical to consider.

4.5.1 Test Solving

Of paramount importance to any form of entertainment is quality, and puzzles are no exception. As manuscripts have proofreaders and plays have rehearsals, puzzles have test solvers. It is crucial for a puzzle to be successfully test solved before it is released to the public.

Fortunately, test solvers are generally easy to come by. A multi-puzzle contest is a large project, so it is rarely written by just one person. A puzzle constructor is, by nature, a puzzle enthusiast, so any member of the project would probably be willing to test solve puzzles written by his fellow contest designers. Test solving is not only a necessary part of the job, it's often a welcome break from puzzle construction. When outside help is needed, any friend who likes puzzles and will be unable to participate in the contest will do. If the project is a professional one, like the *Random House $10,000 Challenge* books, then the duty of test solving and editing the puzzles falls to the publisher, who will have someone capable of this on staff or will hire a qualified freelancer. In cases like these, the author will frequently have passed the materials through his own test solvers before giving them to the publisher anyway.

What does a test solver test for? Primarily, he is making sure the puzzle is simply solvable and error-free. If the test solver can solve the puzzle, then the contestants can solve the puzzle, and that's the most important thing. The test solver must be thorough to the point of being picayune. In a logic puzzle, it is not enough that he finds the intended solution – he must prove that that is the only possible solution. In a crossword, it is not enough that he can fill in the grid correctly by using the crossing letters to complete obscure words – he must look up those words and make sure the definitions are correct. If a puzzle fails, the test solver

must then, depending on the organizational structure of the group making the contest, either fix the puzzle himself or give the puzzle back to its constructor with an explanation of what went wrong. In either case, once the puzzle is amended, it must be test solved again by a different person.

There are two distinct steps to puzzle quality: one is that the puzzle works, the second is that the puzzle works well. The secondary aspects of a puzzle a test solver checks for are that is it fun, robust, and fulfills the role it is intended for. Is a prospective puzzle dull or tedious to solve? Is it too arcane? Does it have an acceptable difficulty level? Does it attain its projected difficulty level? A rigorous test solver ought to answer these questions and ones like them. As a surrogate contest participant, he should be encouraged to express any opinions he may have because real participants are likely to have the same ones. The puzzle's constructor or the contest organizer shouldn't just request a thumbs-up or thumbs-down from a test solver; rather, he should ask for a full critique of the puzzle. Corrections, adjustments, and full rewrites are up to the discretion of the puzzle makers, but they aren't possible without the efforts of a good test solver.

4.5.2 Plagiarism

It is my belief that in an informal setting plagiarism of puzzles is acceptable. The right circumstances are limited, of course, but appropriating other people's puzzles can have a great upside.

It should go without saying that taking a puzzle created by someone else, putting your name on it, and publishing it as part of a multi-puzzle contest is dishonest and illegal. Even leaving the original creator's name on it but neglecting to get permission to reprint it is against the law. If a puzzle creator receives compensation for his work, he better legitimately possess the title to all puzzles involved.

In semi-formal circumstances, plagiarism is still frowned upon. Take, for example, the puzzle extravaganza that takes place on the last night of the National Puzzlers' League annual convention. This is an amateur competition – the constructors volunteer to create it and it is treated as a friendly contest. But given the setting, the participants rightly expect to see unique puzzles custom-made specifically for this event.

Borrowing puzzles is acceptable when someone who is not a puzzle constructor wants to create a quality muli-puzzle contest. Generally, this contest would be played by a group of friends; it would be a casual party activity. Previously published puzzles culled from books or magazines are guaranteed to be high-quality and virtually error-free. They have already been test-solved and edited, and done so by professionals. Constructing a series of puzzles to form a contest is a major undertaking for an experienced puzzle maker. Someone who's never made a puzzle before would otherwise be unable to make such an event. With the luxury of using already created puzzles, all the contest constructor would need to do is devise an endgame. If

he is enterprising enough, he may also create a theme to tie the puzzles together and/or construct a puzzle or two with some local relevance (i.e., a scavenger hunt for items from the neighborhood).

The most formal a multi-puzzle contest can get where some plagiarism is acceptable is the level of the MIT Mystery Hunt. This is a prestigious, highly-attended annual event, but it is still an amateur competition. The puzzle constructors don't volunteer for the task; they are thrust into creating the Hunt by virtue of winning the previous one. This often means that none of the constructors have ever made a puzzle before. They always learn; the winners are, by definition, a team of brilliant puzzle lovers. This is the type of project that's usually still being written a few hours before it's scheduled to start, however. When the majority of the puzzles involved are unique and the endgame and overall structure are arcanely sublime, the existence of a few borrowed puzzles is generally overlooked by the participants, especially when they've been modified from their original form. Such puzzles are rarely used in the Mystery Hunt; in most years the entire Hunt is brand new. In the first year I was in charge of creating the Hunt, I actually plagiarized a few puzzles myself. Five out of the forty-one 1995 MIT Mystery Hunt puzzles had been previously published; nearly every one of them was modified in a significant way. In four of the five cases the reason I borrowed them was because I needed sophisticated pieces of puzzle construction (such as a filled crossword grid and a cryptodivision problem) that, as a puzzle making novice, I was incapable of creating. Each of the puzzle parts I took was merely the foundation for a new, different puzzle. The fifth puzzle, the one that was unmodified, was the puzzle mentioned in Chapter 4.4. I needed a puzzle that was guaranteed, beyond a shadow of a doubt, to be solvable, enticing, and easy. The best way for me to accomplish this was by using one that had previously been published

in *Games* magazine. This strategy was successful; each team solved the puzzle in question shortly after the start of the Hunt.

4.5.3 Fixing Errors

Ideally, through conscientious construction and rigorous test-solving, a puzzle maker will run an error-free multi-puzzle competition. Unfortunately, this goal is often missed, although usually only in amateur competitions. Published works cannot have errors in them. In high-profile competitions like the Tropic Hunt, an error would cause a near-riot among the thousands who participate each year. Errors are likely to occur in contests on the scale of the MIT Mystery Hunt: something written by inexperienced puzzle makers that involves not more than a couple hundred participants.

The lack of experience on the part of the puzzle makers certainly contributes to the number of mistakes present in any particular Hunt. Another factor is the time constraint the puzzle makers are under: the Hunt is usually finished being written just a few hours before it is scheduled to start. There often isn't enough time to test solve all of the puzzles. In addition, due to its nature, it is often unfeasible to test solve the endgame. Finally, errors aren't the end of the world here: the type of structure present in the Mystery Hunt makes it easy for the Puzzlemasters to deal with errors once they arise. All of the participating teams are in the same geographic area, are working on the contest at the same time, and are in constant contact with the PM's through multiple channels of communication. When a team detects an error and reports it to the PM's, it is relatively simple for them to alert all other teams to its presence.

Of course, this doesn't mean errors are acceptable; one error may ruin an entire puzzle, thus wasting hours of work put in by the puzzle constructor and depriving the players of the

joy of solving it. Neither does it mean that fixing an error is always a simple task. Correcting a typo may be, but far worse problems may arise.

The first problem with an error is faced by the team that discovers it first (and any others who have been trying to solve the puzzle in question). The puzzle makers have put forth an implicit guarantee that all of the puzzles are solvable. When one of them isn't, it causes a lot of frustration for the player who, through no fault of his own, has gotten hopelessly stuck on it.

The second problem is faced if the team that discovers the error first chooses not to report it to the PM's. They may feel that since they wasted a lot of time agonizing over the puzzle before discovering the problem, why should they help their competitors catch up? If this kind of thinking is going on, then there is a problem in the basic philosophy pervading the contest. The primary purpose of the competition is entertainment; no participant should maliciously wish a flawed puzzle upon any other participant. Should this occur, the PM's must address the players' attitude and sense of fair play.

Once an error is reported, the issue becomes how to deal with it. The first thing the PM's must do is verify that the puzzle in question is flawed. If so, a corrective measure, if there is one, must be created. Often, the team that finds the error will have, in the course of realizing that it is an error, figured out how to fix it. If not, then the PM's must do so immediately. All teams need to be notified of the problem and its correction as soon as possible.

Sometimes, however, there is no simple solution. When a puzzle is so innately flawed that it cannot be salvaged or when a contest is nearing its end and there just isn't time for the PM's to figure out a corrective measure for a broken puzzle, the puzzle in question must be

sacrificed for the greater good. The PM's must be able to swallow their pride and throw away one puzzle so the entire contest can be successful. This is painfully difficult to do; the puzzle was someone's creation that, more likely than not, took many hours of hard work to construct. It isn't easy to abandon, but the PM's cannot afford to be stubborn. A single individual puzzle could be a major sticking point for the participants, who would then be forced to try to solve the endgame with one less input than was intended. Even if the contest was designed not to require all of the individual puzzle answers, 80% of 20 puzzles (16) jumps to 84.2% of solvable puzzles (16 out of 19) when one is eliminated. On top of all of this, the PM's must make their decision quickly; this is, after all, a race, and every minute counts.

What happens when an error strikes a large structural element rather than a mere individual puzzle? This is, unfortunately, quite common. As stated in Chapter 3.1.4, every MIT Mystery Hunt from 1994 through 1998 has had some sort of endgame problem. In the worst case, there is no correction made and the wrong team wins. In 1994, a last-minute run-through before the Hunt began revealed a ghost error: there was actually nothing wrong with the puzzle, but a faulty test solve made it appear as if there were. The PM's "corrected" this puzzle, but their adjustment introduced an error that hadn't been there before. This error manifested itself at the very end of the endgame by instructing the first-place team to go to a nonexistent room number. They went to the hallway where the room number would have appeared and were stuck there until the second-place team caught up. By this time the PM's had realized their mistake, but felt it was too late to correct it; fixing the error at that point would result in a foot race to the location of the coin. The PM's declared the Hunt a tie and gave the correction to the participant who called them first. This person was the captain of the

second-place team (me); with the new information I found the coin first and kept it. And although members of both teams were granted the rights to creating the 1995 Hunt, only my team acted on it. The outcome of the 1998 Hunt (see Chapter 2.6 for details) was similar: the team that should have finished second wound up with the coin. Although this situation was caused by PM errors rather than puzzle errors, the question of how to deal with the problem remained. In this case, the PM's awarded the rightful first-place team with an honorary Mystery Hunt coin; the team with the real coin retained both it and the right to construct the 1999 Hunt. There is no good way to resolve this sort of mess; the front-running team shouldn't lose because of someone else's mistake, but the PM's certainly can't take the coin away from the team that recovers it. In the end, possession is nine-tenths of the law, so the team that has the coin in hand is the winner; the team that gets gypped out of a victory should be placated as much as possible. The situation would actually be easier if the stakes were cash or merchandise, since then both teams could win a grand prize – unless, like the Mystery Hunt's prize of prestige and the right to create the next one, it could not be duplicated (for example, an expensive trip). Without any satisfactory resolution, it is crucial for the PM's to take every step possible to prevent this type of ending from occurring.

Two of the other Mystery Hunt endgame gaffes turned out to be minor. The error in the 1995 Hunt's endgame was noticed by the PM's and corrected long before any teams reached it. The 1997 Hunt's endgame errors could have caused major problems, but the first-place team had such an insurmountable lead that the hold-up didn't cost them. The error in the 1996 Hunt was major and it was handled extremely well by the Puzzlemasters.

This Hunt had three separate endgame puzzles; only one was broken. The team in the lead (Chaos; again, this was my team) had made great headway into the flawed puzzle; no

other team was able to begin solving it. Chaos held approximately a five-hour lead over the second-place team based solely on the progress made on this endgame. This puzzle hadn't been test solved, so Chaos was actually the first group of people to attempt it. The puzzle was riddled with errors and subjective musical interpretations, so Chaos had worked closely with the PM's on this puzzle for a number of hours. (This collaboration allowed the PM's to correct their puzzle for other solvers and it allowed Chaos to make progress on it.) After a while, it became clear to both parties that this puzzle could not be successfully solved; meanwhile, the rest of the teams were complaining that they couldn't even begin to try. This endgame, the culmination of one-third of the contest, would have to be scrapped: but how to do it fairly? First, the PM's decided that they would throw out the puzzle and release the answer – a very tough choice. Next, since Chaos would be adversely affected by this action, the PM's informed them of their decision and asked them for their permission to go ahead with it. Chaos stalled for as long as possible (it was in their best interests to do so), but had no choice but to eventually give in to the firm stance of the PM's. Finally, the PM's released the answer to Chaos, waited thirty minutes to respect Chaos's lead, then gave everyone else the answer. Chaos went on to win the Mystery Hunt by that thirty-minute lead.

5 Conclusions

There is a solid set of rules defining the structure underlying an MIT Mystery Hunt or any other multi-puzzle contest, but these rules have so many choices and optional elements associated with them that the structure turns out to be very loose. This is where the Mystery Hunt gets its "no rules" attitude: there are, in fact, rules, but they can be manifested in so many different ways that it looks like there aren't.

There are barely any concrete elements in a multi-puzzle contest's structure. There must be an endgame and a number of individual puzzles. Construction of this type of contest must start from the end, with the final answer and the application of any external constraints, then work forwards. The endgame has to be designed first. But even the most rigid rules have exceptions to them. One of the individual puzzles may be designed before the endgame and retrofitted to produce an answer that works in the endgame. The notion that all individual puzzles must be weighted equally is about as solid as a rule gets in this format, but a puzzle maker can make an exception to even that, given enough advance planning.

But while the rigid laws governing the construction of a multi-puzzle contest are scant, the applicable theory is not. It just happens that most guidelines are prefaced with "should," not "must," and that is as it should be for a creative endeavor. The puzzle maker needs enough creative leeway to make a unique, special contest that suits his particular style.

So the operative question isn't "What makes a multi-puzzle contest?" but "What makes an ideal multi-puzzle contest?" The simple answer is that a perfect multi-puzzle contest is one in which every participant has fun. How to do that is subjective and every puzzle maker, as well as every participant, has his own ideas on the right way to do this.

I believe the way to achieve this is through richness of form. I have found that the two most important elements of a multi-puzzle contest are theme and variety. These are technically optional elements: a bunch of puzzles don't need to have a story in order for their answers to combine in an effective endgame; neither do they need to all be different. But it is the story that elevates the multi-puzzle contest from being just a puzzle to being an entire realm all its own and it is the variety that ensures that as many participants as possible will get to solve a puzzle ideally suited to his personal taste and skill level.

The ideal multi-puzzle contest will enchant and delight every participant in a different way. Theme and variety go a long way towards meeting this goal; so does opportunism. Specifically targeting an audience's skills and knowledge will naturally captivate that audience. Using the geography makes a contest more relevant and more interesting. Innovations, bonuses, and other experimental puzzles enhance the aura of originality that a multi-puzzle contest should already have. It thrills a puzzler to no end to not just witness, but participate in, a puzzle that's truly unique – something they've never seen before and will probably never see again.

Finally, the ideal multi-puzzle contest will enchant and delight the Puzzlemasters themselves. Unlike the participants, they're not in it just for fun; they have a lot of work to do. From contest conceptualization to puzzle construction, test solving to error correcting on the fly, the PM's primary job, right up to the time the contest is solved, is to make sure the solvers are enjoying themselves. But they should also make sure they are enjoying themselves. One way to do this is with some excuses for interaction with the participants built into the contest. A more important way is to have fun during the creation of the contest. This isn't always easy; puzzle making is a mentally taxing undertaking. But if the contest construction and test

solving process is imbued with fun, then the contest itself will be as well. The puzzle maker has ample opportunity to express himself while making the contest; the rules are, after all, full of options. The choices made during construction, from types of puzzles included to what the theme is, even to issues as structural as the endgame tradeoffs, are personal decisions that should reflect the puzzle maker's tastes and sensibilities. The puzzle maker should design the contest he would most like to play in.