

Determining the People Capacity of a Structure

Team 243

May 7, 1999

1 Introduction and Restatement of the Problem

Many public facilities have signs posted in large public rooms, such as lobbies and conference halls, that specify the maximum number of people that the space may be occupied by. Presumably, this number is based on the speed with which people could be evacuated from the structure in an emergency. Elevators and other facilities, such as gymnasiums and swimming pools, also have “maximum capacities” posted, which are based on similar criteria.

In deciding what number to put on such a sign, one must consider two important factors. The most obvious is that of safety: what must the maximum capacity of a structure be in order to minimize the time that it takes every occupant of the building to exit without sustaining injury? Another important issue is that of comfort: how many people can be fit in a room, during a given interval, before the room becomes overheated or the carbon dioxide level in the room rises significantly above normal?

Our analysis considers each of the above two issues, which we will call the “emergency problem” and the “comfort problem,” for different structures and spaces.

We will present two models as possible solutions to the emergency problem, both of which give a method for determining the minimum time T it takes N people to exit a specified structure. Conversely, we will use these methods to determine the maximum number of people N who can exit a structure in a given period of time T . For the comfort problem, we will give estimates of the maximum number N people that can comfortably occupy a given space for a period of time T .

To avoid ambiguity, we will use the following definitions:

- A “structure” is an assortment of interconnected spaces, each of which leads to at least one other space or an exit.
- An “emergency” is a situation that poses sufficient potential or actual harm to the well being of the group within a structure to require its complete evacuation.

- The assumption of “orderly movement” states that no personal injuries or other accidents occur that affect the minimum time T taken to evacuate N people from a given structure.
- A “panic” is a situation in which the assumption of orderly movement does not hold.
- A room is “comfortable” if the quality of its air is acceptable and its temperature falls within a specified range.

2 Further Considerations

One difficulty in developing a model for the emergency problem is deciding how different types of emergencies affect the rate at which people can exit a given structure. A bomb threat and a fire are both pressing reasons to evacuate a building, for instance, but the imminent danger that smoke inhalation poses to the occupants of a structure is greater than the knowledge that, five hours later, a bomb may or may not explode in their vicinity. Likewise, a bomb threat called in five minutes before detonation could cause a panic that, in an overcrowded room, might leave many people injured in the rush to exit whether or not the threat is real. The dynamics of the exiting processes for each of these situations, naturally, will present distinctly different modeling situations.

In addressing the emergency problem, we first consider the case where the assumption of orderly movement holds, and then extend our analysis to what might happen in a panic, where accidents and personal injuries occur that slow the rate of movement within the structure and increase the minimum evacuation time for a group of N people.

3 Assumptions and Hypotheses

The following assumptions were made in tackling the emergency problem:

- The people in our models are a uniform, average adult weight of between 100 and 300 lbs, and are of approximately the same size.
- There are no “security guards” or similar individuals responsible for regulating the evacuation of the structures in our models. That is, every individual has the desire to exit the structure as quickly as possible, and employs the same process for deciding the best route by which to do so.
- The ceilings in our rooms are assumed to be of normal height and the uppermost floors of our structures are not located extremely distant from ground level (i.e. they are not crawl spaces and they are not located at the top of skyscrapers).
- The time it takes for a person to move from one room to another is negligible compared to the time it takes to evacuate all people out of a room.

- The room is situated in a modern building in a town or city. In other words, we do not expect our results to be applicable to submarines, space stations, or other unusual structures.

4 Personal Space Constraints

The simplest constraint on the capacity of any room is space. Each person requires about one square meter (9 square feet) to stand and be able to move around comfortably. So if a room is designed for standing or sitting in an upright chair, an upper bound on the room's capacity is given by dividing its area less any area occupied by furniture by one square meter.

Special cases, such as a rock concert or elevator, in which people are willing to stand closer together can be accommodated by dividing by a smaller amount of personal space, say, 0.75 or 0.5 of a square meter.

5 Evacuation Models

Evacuation models answer these questions:

- Given a room full of people, how long will it take for them to exit?
- What is the risk that someone will be injured during the evacuation? (By being trampled, left in the building, etc.)
- In an emergency, how long do people have to get out of the room?

To answer these questions, we developed several models of the evacuation of a room based on different assumptions about how people move through doors, and what kinds of emergencies are likely to force an evacuation.

5.1 The Constant Rate Model

The constant rate model for room evacuation is based on the following assumptions:

- A door will let people through at a constant flow rate, for example, one person per second.
- The time it takes for a person to get in line at a door is assumed to be negligible compared to the time it takes to evacuate the room.
- To simplify the model, we assume that doors do not become blocked during the evacuation.
- People are crowded around each door. Until the room is almost empty, there are enough people standing close to the door to use it to its full capacity. When someone exits, the crowd pushes forward to fill up the gap.

- People select a door based on what they can see in the room around them and attempt to minimize the time it will take them to exit. More specifically, they tend to go either to the nearest door, or to the door which will allow them to exit the fastest.

First, we analyze a room containing only people, and add furniture later. Similarly, we initially ignore the possibility of a panic to simplify the problem.

5.1.1 Single Room with One Door

Suppose we have a single room with one door. In this case, everyone will try to exit through that door. There will always be enough people at the door to use it to its capacity. Thus, if the door allows people through at a rate of r and there are n people in the room, it will take

$$t = \frac{n}{r} \tag{1}$$

time for the room to empty.

5.1.2 Single Room with Multiple Doors

If the room has multiple doors, each person will initially adopt a strategy of going toward the nearest door. If it becomes clear that one crowd is moving faster than the others, people at the end of slow lines will move to the end of the fast line. In this way, all the doors will be crowded until the room is empty. Suppose there are k doors with flow rates r_1, \dots, r_k and n_1, \dots, n_k people exit through the doors, respectively. If the line at one door looks like it will finish before the others, people will move from the end of their line to the end of the faster line. Thus, all the lines finish at the same time, yielding

$$t = \frac{n_1}{r_1} = \frac{n_2}{r_2} = \dots = \frac{n_k}{r_k}. \tag{2}$$

If we let n be the sum of the n_i , we have

$$n = tr_1 + tr_2 + \dots + tr_k.$$

Defining r to be the total number of people divided by the total time of evacuation and substituting yields

$$r = \frac{n}{t} = r_1 + r_2 + \dots + r_k. \tag{3}$$

This equation states that a room with many doors leading out is equivalent to a room with a single larger door whose flow rate is the sum of the rates of all the smaller doors.

5.1.3 Subroom and Corridor Decomposition

Until now, we have been dealing with an empty room. Now we consider furniture and other obstacles.

First, imagine a dining room with a large number of tables and chairs. See Figure 1. In this case, the furniture restricts people to certain paths as they try to exit, but the assumptions of the open room model still hold. People can generally move in whatever direction they want, there is always a crowd at each door, and each door is able to flow at maximum capacity. It is the combined flow rate of all the doors that determines the evacuation time, as in Equation 3.




Figure 1: A dining room, view from above.

Alternatively, obstacles can divide a room into smaller rooms and corridors, which requires a significantly different model. For example, consider a small lecture hall with rows of seats, a table, and several doors. See Figure 2. If people had to leave, they would most likely walk between the chairs rather than leap over them. So, the single room is broken up by the furniture into smaller “subrooms” and “corridors,” as shown in Figure 3. This situation is different from the dining hall because the furniture more severely restricts the directions people can move in. A person in the hall must first exit a row of seats, then go down one of the outside aisles. If one end of one of the aisles were blocked, it would take longer for the last person on that aisle to exit the room. In the dining hall, a blocked passageway is less critical because there are so many other passages.

Once a room has been broken up into subrooms and corridors, it is useful to think of them each as being separate rooms with doors connecting them, and the evacuation problem becomes one of evacuating a whole complex, not just one room. The lecture hall becomes the diagram in Figure 4. The diagram can be simplified somewhat by combining doors that lead to the same place as in Equation 3. In this case, simple inspection shows that once again the exit




Figure 2: A lecture hall, view from above.




Figure 3: Same hall, gray areas denote corridors of movement.

doors can operate at maximum capacity the whole time, so the time it takes to evacuate the room is determined entirely by their combined flow rate.




Figure 4: Same hall, schematic diagram of subrooms. Circles represent subrooms, lines represent passage from one subroom to the next, and ground symbols represent doors leading to the outside. Each subroom is marked with how many people are in it, and each connection is marked with how many people per second can flow through it. Ground symbols indicate exits that lead completely out of the complex.

For a more complex example, consider the cafeteria floor plan shown in Figure 5.¹ Most of the rooms are connected by open arches that function as doors with large flow rates. It reduces to the schematic diagram shown in Figure 6. Here it is not so clear that the flow rate of the four exit doors determines the evacuation time although some simulations described in the appendix and a method of analysis described in Section 5.1.4 show that this is in fact the case. If we had a large room connected to a lobby by a single, small door, and a large door connecting the lobby to the outside, the evacuation time would be more dependent on the flow of people into the lobby. In other words, sometimes a small interior door is a bottleneck, and sometimes it is not. For a complicated network like the cafeteria, whether or not there is an interior bottleneck is not immediately apparent.

5.1.4 Maximum Flow Model

The evacuation problem for a complex of rooms can be solved if we think about it differently by ignoring the actual number of people in the rooms. Suppose people are constantly flowing out of the complex, and are coming into existence

¹This is based on an actual building on campus.




Figure 5: A large cafeteria, view from above.




Figure 6: Same cafeteria, schematic.

on the inside at the same rate. Think of them falling out of the ceiling as fast as people exit. The rooms will all have a constant number of people in them since they are replaced at the same rate they leave. This problem is one of finding the flow rate of people through a complex.

There is an algorithm for finding the maximum flow through a graph, called the Ford-Fulkerson algorithm. Although tricky to implement, it is simple in concept. Suppose we have a directed graph and each connection has a known maximum capacity, in this case, people per second that can pass through a crowded door. One of the nodes is designated the “source” (people falling from the ceiling) and another is designated the “sink” (the outside). We assign to each connection an amount indicating the actual flow through it. Such an assignment can be improved if there is a path from source to sink in which the flow through every connection can be increased. An assignment is maximal if there is no such path. The Ford-Fulkerson algorithm looks at all possible paths until there are no more improvements to be made.

The time for n people to leave the building can be estimated by dividing n by the maximum flow. To use the Ford-Fulkerson algorithm on a room graph to determine this flow, we must add two nodes. First, a source is connected to all rooms with lines of infinite capacity. Second, a sink node representing the outside is connected to all exits from the complex with connection capacities equal to those of the exit doors.

For a continuation of the cafeteria example, see Figure 7. This graph is marked with a maximum flow. It cannot be improved because all the connections leading to the sink are at their maximum. It confirms that the rate of evacuation is determined by the flow rate of the exit doors, in other words, there are no internal bottlenecks. The same technique can be applied to any room graph.




Figure 7: A large dining hall, graph for Ford-Fulkerson algorithm. The + node represents the source and the bull’s eye represents the sink.

6 Quadratic Rate Model

6.1 Motivation for the Quadratic Model

While the constant rate model is useful, the assumption that the rate at which people move through doors in an emergency is constant is subject to scrutiny. One idea that has been proposed to improve on this model is the linear rate model, which states that the rate at which people can exit a room, $f(t)$, is bound by a linear function of the number of people in the room. The evacuation problem can be stated, in terms of the linear rate model, as:

$$\text{maximize } \int_0^T f(s) ds \quad (4)$$

which is the number of people that can evacuate in time T

$$\text{subject to } 0 < f(t) < a \int_t^T f(s) ds + b, \text{ for } 0 < t < T \quad (5)$$

In the above constraint, the integral $\int_t^T f(s) ds = \int_0^T f(s) ds - \int_0^t f(s) ds$ represents the total number of people evacuated after time T minus the number of people who have been evacuated up to the point of time t ; in other words, it denotes the number of people in the room at time t .

The linear model represents the situation where the number of people in the room has a “forcing” effect on the flow rate through the exits. This is modeled by the constant a ; the greater a is, the greater the forcing effect is. The constant b represents the normal rate of flow provided that the forcing effect is negligible. Provided people are exiting the room in an efficient and orderly manner, the linear model hypothesizes that the maximum flow rate out of the room increases as the number of people in the room increases.

As the author of the paper suggesting this model acknowledges, the linear capacity function model does not represent the fact that for sufficiently large flows, the capacity function representing the upper bound of the flow rate should decrease to zero. (See [3].) He stipulates further that, since this decrease occurs only at the upper end of the range of flows, the linear bound model should result in “no lack of realism.”

We reasoned, however, that the evacuation dynamics of an emergency call for an upper bound model that takes large flow values into consideration—for when, other than an emergency or a panic, would such large flow values occur, and more importantly, when would the question of evacuation time be more crucial?

6.2 Developing the Quadratic Model

To this end, we pose a model that assumes the upper bound of the flow rate to be a quadratic function of the number of people in the room at time t . The

evacuation problem, using the quadratic model, can be stated as follows:

$$\text{maximize } \int_0^T f(s) ds \quad (6)$$

$$\text{subject to } 0 < f(t) < q - r \left(\int_t^T f(s) ds - p \right)^2, \text{ for } 0 < t < T \quad (7)$$

In the above equation, the maximum flow rate q occurs when the room is occupied by some optimal capacity p people. The motivation for the quadratic rests on two assumptions: 1) The upper bound decreases when the number of people in the room is substantially less than p because the time it takes people to walk to and through the exit becomes non-negligible compared to the total time required to evacuate all people from the room. 2) Conversely, when the number of people in the room noticeably exceeds p , the jostling, discomfort, and limitation of movement that occurs reduces the flow rate through the exits.

The value of the constant p for a given room depends on its floor space A in square feet and a critical density d , which is the value in people per square foot beyond which impediment to motion increases and flow efficiency decreases. The value of p can be computed by the equation $p = Ad$. In our discussion of the quadratic model, we assume that $d = 0.75$. What the exact value of d should be, of course, is open to interpretation. We consider this problem briefly in our suggestions for further study at the end of the paper.

The other constant of importance in the quadratic model is r , which is somewhat, but not entirely, analogous to the constant a in the linear model. At this point, it makes sense to rewrite the constraint of Equation 6 as the following:

$$0 < f(t) < q - rp^2 \left(1 - \frac{1}{p} \int_t^T f(s) ds \right)^2, \text{ for } 0 < t < T \quad (8)$$

With this formulation, it follows that the constant a is roughly the same order of magnitude as the quantity rp^2 . We will see shortly that this makes sense in the context of an example.

To solve the evacuation problem using the quadratic model, we will assume that maximum flow occurs. The constraint in Equation 6 thus becomes

$$f(t) = q - r \left(\int_0^T f(s) ds - \int_0^t f(s) ds - p \right)^2, 0 < t < T \quad (9)$$

Differentiating both sides twice with respect to t leads to the following differential equation:

$$f''(t)f(t) - f'(t)^2 + 2rf(t)^3 = 0 \quad (10)$$

With a little help from MAPLE, and using the initial values $f(T) = q - rp^2$ and $f'(T) = 0$, we get the following solution for the flow rate out of the room at time t :

$$f(t) = \left(\frac{q - rp^2}{\cos((t - T)(\sqrt{-qr + r^2p^2}))} \right) \quad (11)$$

From this result, we can compute the maximum number of people N who can exit the room in a time interval T :

$$N(T) = \int_0^T f(t)dt = \left(\frac{-\tan\left(T\sqrt{r(-q + rp^2)}\right)(-q + rp^2)}{\sqrt{r(-q + rp^2)}} \right) \quad (12)$$

Using the above, we can solve directly for the the minimum time T it would take N people to exit a structure:

$$T(N) = \frac{\left(\frac{\arctan\left(-N\sqrt{r(-q + rp^2)}\right)}{-q + rp^2} \right)}{\sqrt{r(-q + rp^2)}} \quad (13)$$

Equations 12 and 13 are the main results of our analysis using the quadratic model. They allow us to address two of the main questions raised in 5—namely, how long it takes N people to evacuate a building and how many people can be evacuated from a building in time T .

6.3 The Relevance of the Quadratic Model to Emergengy Dynamics and “Panic”

The above results bring us to an important point. At the beginning of the paper, we made the distinction between a “panic” and the “assumption of orderly flow.” In a panic, it is assumed that some people will sustain injury, fall down, or disrupt the flow of the crowd in some related way. The justification we have presented for the quadratic model, in the case where the number of people in the room exceeds the optimal value p , assumes something similar: people packed together at a density greater than the critical density will slow each other down in their attempt to evacuate a room. Essentially, therefore, we are assuming that the difference between the impediments to flow caused by crowding and the impediments caused by panic is one of degree. This supposition is reflected in the fact that the quadratic bound for the maximum flow rate as a function of the number of people in the room is concave down, which represents the fact that every additional person added to a crowd in a given room poses a slightly greater obstacle to efficient evacuation.

To illustrate the predictions of the quadratic model for a normal example, suppose the optimum flow rate $q = 6$ people per minute, the number of people at which optimum flow occurs $p = Ad = (1000)(0.75) = 750$ people in a room of size $A = 1000$ square feet, and $T = 6$ minutes in which to evacuate as many people as possible. For purposes of comparison, we will also compute the number of people who can exit in this time predicted by the linear rate model with $b = q = 6$ and $a = .01$. We will take the value of r for the quadratic rate model to be $a/p^2 = .01/(750^2) = 1.8 \times 10^{-8}$.

Doing so yields $N(6) = 540$ for the quadratic model and $N(6) = 557$ for the linear model. It makes sense that these numbers are not too far apart, since we are not dealing with an extreme case where the number of people evacuated from the room greatly exceeds or undercuts the critical value p . More to the point, it seems that when p does not deviate significantly from Ad , this will be usually be the case. However, if we set p , for example, to 10^{-5} and compute the above example with all other things held constant, we get $N(6) = 501$ for the quadratic model. By setting $p = 10^{-4}$, we get $N(6) = 195$. The difference in these predictions demonstrates the effect that increasing p has on the maximum number of people who can be evacuated in a given time predicted by the quadratic model. Rather than forcing people out at a higher rate, as the linear model predicts, the quadratic model suggests that a “forcing effect” caused by the efforts of a packed crowd to evacuate a building in a hurry may actually decrease the total number of people evacuated by causing injuries and inefficient flow.

6.4 Limitations of the Quadratic Model and Further Considerations

Like the linear model, the quadratic model has some important limitations. One is that it is designed to model the evacuation of a space, rather than an entire structure. In our simulation, however, we have managed to apply the quadratic model to a cafeteria on our campus with results that agree with those obtained by the constant rate model. Unfortunately, we did not have the time to compare the predictions of the quadratic rate model and the constant rate or linear model by simulating a “panic” situation in the cafeteria. One further extension of our project would be to simulate a variety of panic situations using the quadratic rate model, the linear model, and the constant rate model and compare the results. Ideally, the quadratic rate model should yield evacuation times for panic situations that are noticeably less than those predicted by the other two models.

Another problem that became apparent in applying the quadratic model to our cafeteria is that when the number of people in the room is significantly less than the value of p , the simulated people would not go through the exits. This was remedied by using a constant rate model when the number of people in the room dropped below 10 and switching to the quadratic rate model when the number of people in the room was above 10.

In our model, we estimated the values of p and d . Since the results given by

our model depend heavily on these two constants, it is important to find ways of estimating them more accurately. This problem presents a number of different possibilities for further study that we will mention briefly in the conclusion.

6.5 How Long is Long Enough?

We have models for how long it takes a certain number of people to evacuate a room or complex, and if it were given that all the people had to be out by time t , it is a matter of back-solving and perhaps some trial and error to determine the largest number n of people that can escape. The difficult part is determining t .

In most cases, t is completely unpredictable. For example, someone might call in a bomb threat and give insufficient warning (maybe fifteen seconds) for more than a fraction of a room to be evacuated. Presumably a bomber intending for people to escape would give reasonable warning. Otherwise there would be no warning at all, in which case there is nothing we can do.

In the case where the air in a room is being polluted with toxins, a fairly good estimate can be made of t . Given a room of volume V , the amount N of air molecules is given by the gas law $PV = NRT$, where P is pressure (1 atmosphere), T is the room temperature in Kelvins, and R is the gas constant. Most toxins of note are lethal in small amounts, so we may assume the pressure is constant. Denote by r the rate in moles per second of toxin being created, and by q the fraction of the air which is toxic. Then:

$$qN = rt \quad \text{or:} \quad t = \left(\frac{qV}{r}\right) \left(\frac{P}{RT}\right) \quad (14)$$

Under room temperature and one atmosphere, $\frac{P}{RT} = 41.4 \frac{\text{mol}}{\text{m}^3}$. Substituting for q the lethal concentration of the toxin yields t .

For a fire, which is far more likely for most rooms, we have a “back-of-the-napkin” calculation for how long the room will fill up with the toxin CO_2 . Consider a wood fire. A rough estimation based on camp fires gives that 1 kg of wood combusts in 15 min or so into 0.25 kg of CO_2 , 0.25 kg of water vapor, and 0.5 kg of ashes. Dividing by the molecular mass of CO_2 , we get that combustion of wood creates about $r = \frac{1}{3}$ mol of CO_2 per minute per kilogram of wood, and t comes out to be 600 s per cubic meter of volume divided by the number of kilograms of fuel. For example, a room 6 m by 12 m by 3 m polluted by the smoke from 100 kg of burning wood reaches 8% of CO_2 in 21 minutes. This estimate is very rough and should be improved to take into consideration the heat and change of pressure over time. Furthermore, the room will become very uncomfortable far before it becomes deadly, so 21 minutes is an overestimate t .

For reference, a room that size can hold perhaps 100 people, and if there is one door operating at one person per second, the room can be evacuated in less than two minutes using the constant rate model.

The example cafeteria has a volume of around 15000 cubic meters, and by the above estimate, it would take 9000 s for it to reach 8% CO_2 if a ton of wood were to burn inside, which is 2.5 hours. By our estimate it can be evacuated in 2.5 minutes.

7 Ventillation Models

The comfort level of people in a room is another consideration for the maximum capacity. According to our research, the following atmospheric conditions worth considering in determining a legal capacity for a crowded room.

- The temperature of the room should be between 65 and 90 degrees Fahrenheit. In particular, the ventillation system should be able to dissipate the heat produced by the bodies of the people inside.
- The amount of various toxins in the atmosphere should be kept to harmless levels. The only one likely to apply to all situations is carbon dioxide (CO_2) which is produced naturally by human respiration. It is recommended (in [2]) that the CO_2 level should be below 0.1%. It must be kept below 8%, a level which can be fatal.
- If smoking is allowed in a room, additional circulation must be allowed for.

According to an air conditioning manual ([2]), human bodies produce heat at a rate of 60 Watts when asleep to 600 Watts when undergoing strenuous activity. Moderate activity yields about 100 Watts. The rate at which a room dissipates heat depends upon its insulation, what sort of windows it has, and the power of any air conditioner that flows through it, and must be determined on a case by case basis. Rooms such as auditoriums which are used for several hours at a time should be able to dissipate 100 Watts per person so that the temperature remains roughly constant.

A certain amount of fresh air is recommended in [2], at least 0.2 liters per second per person. Fresh air dilutes the CO_2 concentration and unpleasant odors. If smoking is allowed in a room, 25 liters per second per person are recommended.

In a tight, enclosed space, the CO_2 produced naturally by human respiration becomes important. According to [1], a normal human breath is about 500 ccs, 4.1% of which is CO_2 , and the breath takes 4 seconds or so. Thus humans produce CO_2 at a rate of 5×10^{-3} mol per second. Equation 14 in Section 6.5 can be used to estimate the time the room can safely be inhabited with a given occupancy.

Consider for example an elevator 3 m by 3 m by 3 m carrying 12 people that becomes stuck and is somehow completely air-tight. The people take up about half its volume. Using Equation 14, it takes about 150 minutes, two and one-half hours, for the CO_2 level to reach 8%. Here, the capacity of the elevator might be limited by the time it takes to get a rescue crew in to open it up. In a city, such a rescue should not take more than a few minutes. Elevators are usually well-vented anyway, so CO_2 build-up will normally not be a significant constraint.

The manual [2] also states that the fraction of oxygen in the air can be allowed to decrease fairly significantly (down to 13%) before it becomes dangerous, so the presence of toxins is the limiting factor.

8 Swimming Pools

In the case of an indoor swimming pool, evacuation is basically the same as for an open room. People can exit the pool itself on all sides, except for weaker swimmers who may have to use a ladder. They should be able to flow through the exit doors as described in Section 5.

In the case of an outdoor swimming pool, evacuation is not much of a consideration. Outdoor pools are evacuated mostly to avoid lightning, and there is usually sufficient warning.

In both cases, personal space is the most important safety issue. In the water, people must move their arms and legs over a greater range of motion to maneuver than is needed for walking on land. Many swimming strokes limit a swimmer's vision and make collisions more likely. Some swimmers have to wear floats, which take up additional space.

For swimming pools, we recommend that the area of the pool be divided by 3 square meters per person (giving each swimmer one meter in all directions to move) to determine the capacity. A large space should left open around diving boards and slides, perhaps a circle of 4 meters,

9 Capacities for Elevators

Elevators usually have very wide doors and only carry a few people at once. Thus, evacuation time is negligible in case of an emergency. (The real time constraint will be getting the people down the stairs and out of the building, which is similar to the room problem.) As mentioned in Section 7, the amount of fresh air available to a small, enclosed place can place some limitation on the number of people who can stay inside for long periods of time. For elevators, this is only a factor if it gets stuck and rescue teams take more than two hours to cut a hole in the wall.

The most important factors would seem to be weight and space. Elevators usually have a weight limit which is supplied by the manufacturer, and a simple elbow-room constraint which can be calculated by dividing the floor area by 0.5 square meters per person.²

10 Article for the Newspaper

There are many important considerations when determining the maximum occupancy of a room. We all realize the desire to fit as many persons into a room as possible, but we also must consider the safety of those occupants. A building filled past its maximum occupancy becomes a modern day Titanic, trapping victims inside its walls.

²The case of an elevator seems fairly straightforward, so we spent most of our time working on rooms and complexes.

Our model has been rigorously tested. We consider both the comfort of the occupants and their ability to leave quickly in an emergency. Comfort is a fairly straightforward measure, allowing for personal space and sufficient ventilation. Building evacuation is far more complex.

With building evacuation, we had to consider the possibility of panic along with the time it takes to evacuate under more structured conditions. Our model takes this into account, and provides quite excellent predictions of the time in which various numbers of occupants could leave a building. Our model is derived analytically from common sense and first principles, but buttressed by actual data. The values of predicted constants have been determined by careful comparison with actual data gathered, ensuring that our model matches the real world. The model gives a good match to other previously used methods of calculating maximum capacity as well.

There is always room for disagreement of course, as some may feel that our calculations allow for too little capacity. We stand behind our calculations. Since our model predicts capacities similar to those posted currently for existing buildings, it seems reasonable. It would be always possible to increase the capacity by increasing the time to evacuate, but that is not always acceptable. We predict based on evacuation times less than ten minutes, which we feel is quick enough for most emergencies, and agrees with our calculations and experience. It strikes a good balance between getting occupants out quickly and allowing as many as possible to enjoy the building.

11 Conclusion

11.1 Strengths and Weaknesses

On the whole, our models are fairly robust, and the quadratic model is a more realistic modeling tool than the linear bound model, since it more accurately simulates panic. Unfortunately, we did not have time to test them on as many actual structures as we would have liked. The quadratic model, additionally, yields questionable results for large values of room occupancy. Hopefully, the model could be modified with more simulations on existing structures.

11.2 Recommendations

There are a number of different possibilities for further study presented by our models. For the quadratic model, we could extend an analysis of how to determine the value of p to a more comprehensive understanding of how the “forcing effect” operates to slow the evacuation of a panicked crowd. Also, we could develop techniques for measuring the value of the critical density d , such as observing how many people can evacuate a building in different time intervals T , and using that data to estimate the critical value at which maximum flow occurs.

For improving our analysis of the comfort problem, we could develop ways to better estimate the time a room takes to become overheated our stuffy.

A The Computer Simulation of Evacuation

To test the evacuation times of complexes of rooms, we wrote a simulation engine in Python. It represents rooms connected by doors. Object oriented programming techniques allow us to use different kinds of doors (always open, sometimes blocked, variable flow rate, etc.) and different strategies of selecting a path out of the building with the same structural models. Each door has a queue of people standing around, waiting to get through. Each time step, all the doors “warp” some number of people into the next room. Then, everyone standing in line is given the opportunity to move to a different line based on their perception of the room. A special room object is designated the “outside” and throws an exception to halt the simulation when a specified number of people have arrived outside. A class diagram for the simulation is given in Figure 8.




Figure 8: Class diagram of the simulation in abbreviated UML. Triangles indicate inheritance. Hairline arrows indicate “creates.”

References

- [1] Hughes, G. M. *Comparative Physiology of Vertebrate Respiration*. Cambridge, MA: Harvard University Press, 1963.

- [2] Jones, W. P. *Air Conditioning Engineering*. London: Edward Arnold Publishers, 1973.
- [3] A report prepared for the Center for Fire Research of the National Engineering Laboratory.