ISSN 1012-277X S.Afr.J.Ind.Eng.

S A Journal of Industrial Engineering, Vol 5, No 2, December 1992, pp.50-65

LP IN TOWNSHIP PLANNING

Adolf Diegel, Professor of Decision Science
Faculty of Commerce, University of Natal, Durban 4001

Ezio Gori, Horne Glasson Partners, Durban

Initial version presented at EURO XI, Aachen, Germany, 16-19 July 1991

Abstract

A township planner wants to house many people within a given area, without exceeding a tight budget, while yet respecting certain standards. These *specify a variety of house types*. They also imply *indirect effects* such as open space, schools or hospitals: more people use more of these. But as we build more schools, less space remains for houses, so we are left with fewer people, requiring fewer schools.

We show how to capture such simultaneous interactions. We also discuss some unexpected results. These arise notably in the transition from fractional to integer solutions.

Opsomming

'n Dorpsbeplanner probeer soveel moontlik mense behuis in 'n betrokke area, sonder om 'n beperkte begroting te oorskrei, terwyl gegewe standaarde behou moet word. Laasgenoemde maak voorsiening vir 'n *verskeidenheid van behuising*. Dit behels ook *indirekte navolge*, soos oop ruimte, skole of hospitale: groot getalle gehuisveste mense maak meer gebruik hiervan. Maar soos daar meer skole gebou word, word minder ruimte gelaat vir behuising, dus word minder mense behuis, wat minder skole benodig.

Ons toon aan hoe om sulke gelyktydige interaksie te verwerk. Ons bespreek ook sommige onverwagte resultate. Hierdie ontstaan veral in die oorgang van fraksionele na heeltal oplossings.

Introduction: the township planning problem

The *direct* effects of building a house are that it requires money and land, and these are scarce. The *indirect* effects are that, for a group of (private) houses, we must provide some (public) buildings such as civic centers, schools, and similar amenities. These may be financed privately, but they must still fit the land at hand. So the fundamental problem is: given a budget limited in both land and money, how many (private) houses can we build, while yet allowing for all the (public) amenities they will ultimately require?

For example, consider key data of a project handled by Horne Glasson Partners in the Durban area:

Table 1. Direct effects, by house type

House type	Flat	Row	Duplex	Town	
Inhabitants	48	30	12	4	Capacity
Land surface	2600	1800	700	450	650000
Cost of house	312000	186000	66000	34000	55000000

A flat (more exactly, a block of 12 flats housing 4 people each) accommodates 48 in all. It covers 2600 square meters and costs R312000 to build. Similar data apply to the other house types: they house fewer people, but also consume fewer resources. Our task is to maximize the number of people who can live on the land available, without exceeding the budget shown under *capacity*. Suitable formulae appear in Table 2.

Table 2. LP-formulation of people, land and cost per house

	Flat		Row		Duplex		Town			
maximize	48	+	30	+	12	+	4	such that		
	2600	+	1800	+	700	+	450	<	650000	Land
	312000	+	186000	+	66000	+	34000	<	55000000	Cost

The variables are known by their column heading, so we need not repeat them in each inequality. The plus sign can also be omitted to facilitate presentation of larger data sets. Alternatively, we could use the hand-written style

maximize 48*x1 + 30*x2 + 12*x3 + 4*x4, etc.,

In any case, the essential part of translating a data table into inequalities is to insert "less than or equal" before each capacity (using simply < or > to include "or equal").

In short, formulating direct effects is straightforward. It is otherwise for the *indirect* ones: while we can easily see that private houses imply public buildings, it is *not* easy to see *how many* of the latter. This is so because indirect effects arise only as a result of the direct ones. We do not know how many schools to build unless we first know how many people we have housed. Yet as we build more schools, less space remains for houses, so we are left with fewer people, requiring fewer schools....

Indirect effects: composite formulation

How can we formulate indirect effects? Indeed, how can we do so with a view to optimizing/minimizing their impact on the over-all solution, especially when it contains integers?

We can do so by a *composite* method, that is, *compose* or combine interrelated data into compound blocks, then optimize the block rather than its components. Conversely, we deal with each component *individually*, while formulating their interdependence in separate cause-effect statements.

To see the merits of either approach in the context at hand, consider the demand that we build *"at least twice as many row houses as flats"*. Such a demand is *indirect* if we would not build row houses on their own account (because they are cheaper), but only because we have already opted for flats: the row houses are to break up the monotony.

The demand is easy enough if the planners want *exactly* twice as many of one as of the other. We could then plan, not for individual units of *flats* and *row houses*, but for *composite blocks*, "one flat plus two row houses". Such a block covers 2600 + 2*1800 or 6200 square meters and houses 108 people.

Yet such a composite approach could be optimal only by coincidence, if the land available is an exact multiple of 6200 (and if all other requirements solve in corresponding integers). If not, a mix of say 10 flats and 22 row houses improves on the two-to-one ratio (it *is* an improvement if the planners want row houses to begin with). Also, the two extra houses *use land otherwise idle*. This is so because an area admitting just 10 & 22 could not hold 11 & 22; but if we make it 11 & 20 or 11 & 21, we are not getting twice as many of the latter; finally, 10 & 20 or 10 & 21 are clearly not optimal when we can have 10 & 22.

The nature of the composite method becomes even clearer if we maintain the previous two-to-one ratio and also ask for a suitable number of schools, for example. We are to provide one (primary) school for every 720 children. If a two-to-one block houses 108 people, we expect one third to be children, or 36 per block. 720/36 makes 20, or one school per twenty blocks. Finally, a school occupies 40000 square meters, or 2000 per block, so we need 8200 square meters for one flat, two row houses and a classroom for 36 children.
With 650000 square meters available for the project at hand, we can build 650000/8200 or 79.2683 composite units. These contain one flat per unit, so we know that we can build as many flats, then twice as many row houses and 1/20 as many schools. We also house 108 people per block, specifically:

Table 3. Solution based on composite block data

Item	Flat	Row house	School	People	Children
Number	79.2683	158.5366	3.9634	8560.98	2853.6585

This solution is hardly obvious from the raw data, so a composite approach does have its merits. We can add to these that it requires only basic arithmetic (and no knowledge of linear programming). And it helped to clarify our problem.

On the other hand, we cannot use the composite approach for any but *exact* proportions ("twice as many" rather than "*at least* twice as many"), so it is not optimal, in principle. Nor will it obtain integer solutions except by coincidence.

Indirect effects: individual if/then-formulation

The composite approach may be obvious, but it is cumbersome, even with two house types and one type of public building. In fact, we have a dozen house types, secondary schools as well as primary ones, then trade schools and universities. We must include hospitals, places of worship, open spaces, even shopping centers, roads and bus stations. We also face constraints other than land, notably cost, possibly time (if the project must be completed by a certain date).

Such a host of data would become hopelessly entangled if we were to see each as some multiple of another. But the lines remain clean if we translate cause-and-effect into if/then-statements. For example, consider the requirement that we build "*at least twice as many row houses as flats*":

Rows > 2*Flats, *the number of row houses must be twice as large as the number of flats.*

This statement expresses the idea as clearly as we speak the words (yet keeps the arithmetic as simple as the composite method). Its shortcoming is that the variables appear on both sides of the equation. Yet this is quickly rectified:

-2*Flats + 1*Rows > 0 or better, 2*Flats - 1*Rows < 0

We prefer the formulation at right because it aligns neatly with the inequalities defining the original constraints:

Table 4. LP-formulation of simple "twice as many"

maximize 48*Flat + 30*Row

(+)	2600	+	1800	<	650000	Land
(+)	312000	+	186000	<	55000000	Cost
(+)	2	-	1	<	0	Quota, "if/then"

With the same "less than or equal"-sign in all inequalities, we can read the other signs consistently. Specifically, if (+) 2600 under Flat means that a flat *uses* so much land, the (+) 2 means the same: a flat *uses up* so many units from the quota on flats. And -1 under Row means the opposite to +2: each row house *contributes* to the quota. In short, *if* we build flats, *then* we must have 2 row houses for each flat.

The new formulation's essential advantage is that we need no longer combine variables into new, compound blocks. Rather, individual data remain as they stand in Table 2. We merely add a new constraint. If we can add a constraint for flats, we can do so for any variable. And the constraint standing on its own, it can go integer on its own, or remain fractional, as the planners see fit.

Most importantly, an if/then-formulation admits "larger than or equal" rather than "strictly equal", so it guarantees a truly optimal solution. At least, it does so as long as we also consider the aspect of *dominance* discussed below.

Dominant variables

The nature of dominance becomes clear as soon as we examine the unit costs of the main constraints (from Table 1):

Table 5. Unit costs for each constraint

House type / Inhabitants	Flat 48	Row 30	Duplex 12	Town 4	Capacity
Land surface	2600	1800	700	450	650000
Surface/person	54.17*	60.00	58.33	112.50	
Maximum people	12000	10833	11142	5777	
Cost of house	312000	186000	66000	34000	55000000
Cost/person	6500	6200*	5500*	8500	
Maximum people	8461	8870	10000	6470	

Which type of housing is most effective? Clearly a flat is best when it comes to land usage (with 54.16 square meters per person). But since it is built in two or three stories, it is somewhat more expensive. Construction cost per person is lowest for a duplex. Town houses would never be built for the sake of their efficiency (and we will omit them for the time being). On the other hand, row houses *do* use more land than flats, but they are cheaper in their building cost.

Given these facts as well as our budget limits (but ignoring for the moment indirect demands), we compute how many people we can possibly house under either constraint. Clearly *cost dominates land* and, being cheaper, *duplexes dominate other house types*. So, to optimize the data as they stand, we make 55000/66 or 833.3333 duplexes. These house 10000 people, but over 66666 square meters of land remain unused.

This solution may be all for the better, to be sure, but it hardly helps to clarify other relevant aspects of our data. So, both to advance our task and to explain an idea that has merit in its own right, we propose to *neutralize* the cost constraint. This means not simply to omit it, but --

instead of 312000*F + 186000*R + 66000*D < 55000000 Cost

we write 312000*F + 186000*R + 66000*D > 0, accumulate!

In other words, more money can eventually be obtained from other sources, but land is limited absolutely. So we aim to utilize land as effectively as possible, regardless of cost. Yet, while we ignore the cost limit, we still know exactly how much we will require for any given planning proposal.

Mixing house types

With cost being neutralized, we know from Table 5 that flats make the best use of the available land. So again, if we run our data as they stand now, we would build *only* flats. Such a solution is unacceptable. A modern town must mix several house types, not repeat one type endlessly.

Which mix *is* acceptable? An obvious answer does not exist, nor can we derive one by analytic means. What we *can* do is to define various mixes, then explore their effect on the solution as a whole. We do this via any LP-program [1].

For example, here are our simplified raw data and the results of a first run:

Table 6. LP-formulation and solution for simplified data

max 48 Flat + 30 Row + 12 Duplex

2600	1800	700	<	650000	Square meters
312	186	66	>	0	Cost in thousands

People	Flats	Row house	Duplex	Total cost
12000.00	250.0000	None	None	78000.0000

We know already that flats use the smallest land surface per person, so we must build flats to use land most effectively. Our program proposes indeed 650000/2600 or 250 flats. They house 250*48 or 12 thousand people. They also cost a total of 78000 thousand Rand. The answer is trivial, but it is still the basis for evaluating other solutions.

A sea of flats is unacceptable. As a first improvement, we insist on having some of the other house types. How many of each? At least as many as we have flats, for example:

Table 7. As many row houses or duplexes as flats

Constraint: 1*F - 1*R - 1*D < 0, "as many"

People	Flats	Row houses	Duplexes	Total cost
12000.00	250.0000	None	None	78000.0000
11818.18	196.9697	None	196.9697	74454.5455
11814.00	196.0000	1.0000	198.0000	74406.0000

We first show the constraint to be added to those in Table 6 (effectively, *for each flat, we must have a row house or a duplex*). Then we show the previous, unconstrained solution so we can compare one with the other.

If we accept one or the other of the alternatives to flats, we get *only* duplexes. Why so? The latter use less land per person than row houses. Indeed, we get no row houses at all if we build 196.9697 flats. We may be tempted to round *up* to 197, but more flats force the duplexes down to 196, so we no longer get "at least as many duplexes as flats"!

But we can *reduce* the number of flats. This makes room for *more* duplexes and indeed for one row house. Also, it hardly changes the number of people: we lost 186 in 12000, less than 2%. This is so because we adjust *optimally*: what we lose on flats, we gain on other types of housing.

If mixed housing holds roughly as many people as flats only, what about total cost? Surprisingly, it *decreased* (from 78 to 74+, by 4.61%). Or is this surprising? Duplexes *are* cheaper to build, but they use more space per person.

More mixing

Mixed houses accommodate roughly as many people as a single type and, unexpectedly, they cost less. This invites exploration. How else can we mix house types?

We just had as many of one or the other, so the next step is "one and the other": if we plan for 200 flats, then we want also 200 row houses and 200 duplexes, not just the latter on their own. However, now that we propose so many of the other types, we can no longer build 200 flats. No room! But now we know (from Table 4) how to enforce certain proportions, so it remains to examine the quantitative results:

Table 8. All three house types in various proportions

People	Flats	Row houses	Duplexes	Total cost
12000.00	250.0000	None	None	78000.0000
Flats and one or the other				
11818.18	196.9697	None	196.9697	74454.5455
11814.00	196.0000	1.0000	198.0000	74406.0000
One time as many of each as of flats				
11470.59	127.4510	127.4510	127.4510	71882.3529
11466.00	127.0000	127.0000	130.0000	71826.0000
Half a time as many of each as of flats				
11649.35	168.8312	84.4156	84.4156	73948.0519
11646.00	168.0000	85.0000	86.0000	73902.0000
Two times as many of each as of flats				
11289.47	85.5263	171.0526	171.0526	69789.4737
11286.00	85.0000	171.0000	173.0000	69744.0000

We had "as many of one as of the other", indeed "one time as many". How about "two times" or "half a time"?

"Half a time" gains room for 180 people, 11466 to 11646. But this seems negligible, certainly so as it means extra cost as well as again an undesirable dominance of flats.

How about the alternative, "twice as many"? We save a lot of money and we lose room for only 180 people from those we had in a true balance of houses (11466 - 11286). And with so many of the other types, the flats are hardly visible.

Many planners consider two-to-one acceptable, if not ideal, and this we need not argue. We do argue this:

1) If/then-logic can handle *any* ratio, not just two-to-one but say five-to-four, for any pair of house types.

2) Changing such a ratio means surprisingly little change in the main objective (number of people, total cost), *provided all simultaneous changes are done optimally.*

Indirect effects: children and schools

People want not just shelter, but a living environment. So if we build houses, we must leave room for hospitals, shops, schools, parks etc. These must fit into the given space and budget, so we cannot just add them on at will. Rather, as we have more people we also need more schools; but as we build more schools, we have less room for people. We need another set of simultaneous equations to capture these interactions.

For example, continue for the moment with "two to one" as an acceptable house mix. We already know that this leaves room for 11286 people within the given area. How many of these are children? 33% or 1/3, according to the planners.

33% of 11286 makes 3762 children. We could stop here if we just wanted that number. But we also want to build it into our equations so we can subsequently compute the demand for schools and the resulting demand for extra land and money:

Table 9. LP-formulation of house mix and child count

max 48 Flat + 30 Row + 12 Dup + 0 Students

2600	1800	700	0	<	650000	Square meters
312	186	66	0	>	0	Cost thousands
2	-1	0	0	<	0	Two Rows/Flat
2	0	-1	0	<	0	Two Dups/Flat
16	10	4	-1	<	0	Count children

The logic of "if flats, then two row houses" (from Table 4) applies as well to counting the number of children:

2 Flats - 1 Rows < 0, 2 Flats < 1 Rows, or Rows > 2 Flats,

the number of row houses must be twice that of flats, or

the number of children is 16 times that of flats....

It was 48/3 or 16 children per flat, then 10 and 4 per row house and duplex. These numbers reflect 33% of each type's inhabitants, but they can be set more discerningly: flat dwellers may have more or fewer children than other people.

Anyway, if we run the data as they stand, we merely *count* children and we obtain indeed the total of 3762 computed previously (Table 10). How many schools do they require?

The planners tell us that we need one school for 720 pupils. So a naive answer to our question is, 3762/720, at least 5 schools, even 6. This simple division omits that a school takes space, indeed 40000 square meters of land, and money, 792 thousand Rand. If these resources go to schools, they no longer go to housing. This means fewer people in the same area. In turn, fewer people mean fewer schools.

We define these interactions as we count children: instead of placing -1 under *Students*, we place -720 under *Schools*, each school absorbing so many children. Then also we include land and cost per school in the corresponding rows.

Table 10. Accounting for Schools

People	Flats	Row houses	Duplexes	Children	Schools	Total Cost
11289.47	85.5263	171.0526	171.0526	3763.1579	None	69789.4737
11286.00	85.0000	171.0000	173.0000	3762.0000	None	69744.0000

People	Flats	Row houses	Duplexes	Students	Schools	Total Cost
8542.04	64.7124	129.4248	129.4248	2847.3451	3.9546	55937.3894
8508.00	63.0000	126.0000	142.0000	2836.0000	4.0000	55632.0000

We first test the logic of counting children. Indeed, 3762 added up from each house gives the same total as 33% of all people. We then include schools for them. They take so much space that we get fewer houses. Room remains for only 8508 people. These no longer have 3762 children, but only 2836, and *these* finally need just 4 schools!

It is intuitively clear that more schools mean fewer houses. But how can *more* schools mean *less* total cost (it fell from 70 million to just under 56)? The answer is that schools use up so much land, at such a low cost per square meter, that little land remains for the (expensive) dwellings.

Table 11. Cost per surface of various buildings

Building	Flat	Row	Duplex	Town	School
Land surface	2600	1800	700	450	40000
Cost per unit	312000	186000	66000	34000	792000
Cost/surface	*120.00*	103.33	94.29	75.56	*19.80*

A fractional number of schools?

How many students can go to 4 schools? At 720 per school, 2880. This is 44 more than the 2836 we need (Table 11). Or else, we can build 3*720 schools and one of 676, or 4*709. However, if we build small schools, we have more room for houses. Indeed, the small school can become bigger, too!

Run LP again *without* insisting on an integer solution for schools, but *only for houses* (variables 1 to 3):

Table 12. Fractional number of schools (large & small)

People	Flats	Row houses	Duplexes	Students	Schools	Total Cost
8542.04	64.7124	129.4248	129.4248	2847.3451	3.9546	55937.3894
8508.00	63.0000	126.0000	142.0000	2836.0000	4.0000	55632.0000
8538.00	64.0000	129.0000	133.0000	2846.0000	3.9528	55870.6000

By avoiding empty desks we gain room for 30 extra people and for 10 *more* pupils. Their total comes to 2846. They go to 3 big schools, then to a small one of 686. This small school has 10 more pupils than the last of the four big ones!

Critical fractions must not be forced *down*

We removed fractions first by going *all-integer*, with four large schools. We then forced into integers only the units that *must* be integers (houses), while keeping schools fractional. This enabled us to house more people. Indeed, we had *more* students in 3.9528 schools than in four large ones, the latter being partially empty.

Can we have 3.9528 schools? Of course not, but we can build large and small ones, and this idea is intuitively clear. It is less clear that we must not force a fraction *down* to "at most 3 schools", for example. What happens if we do?

Table 13. Fractional number of schools forced *down*

People	Flats	Row houses	Duplexes	Students	Schools	Total Cost
8542.04	64.7124	129.4248	129.4248	2847.3451	3.9546	55937.3894
8508.00	63.0000	126.0000	142.0000	2836.0000	4.0000	55632.0000
8538.00	64.0000	129.0000	133.0000	2846.0000	3.9528	55870.6000
6480.00	None	None	540.0000	2160.0000	3.0000	38016.0000

3 schools of 720 accept at most 2160 pupils. These being 1/3 of all people, the latter cannot exceed 3*2160 or 6480: in limiting the schools, we limited the population! Much land remains unused and it hardly matters which houses we build. Conversely, *forcing schools up to 4 is harmless*. It wastes some space, but hardly affects the solution.

So we must never arbitrarily force a fraction *down*. We go all-integer if that is necessary, and decide, optimally, whether to cut up or down. Or else, we accept a fraction as it stands, then translate it into *large* and *small* units.

Limiting cost

As we *accumulate* cost, so we can *limit* it. For example, we now use just over 55 million Rand. If these are available, we need say no more. If we have *more* funds, we need only say that the excess remains unused. Extra money becomes useful only if we can acquire more land. But LP comes into its own if we have *less* money than we need, say R55 million:

Table 14. Various solutions without/with cost ceiling

People	Flats	Row houses	Duplexes	Students	Schools	Total Cost
8542.04	64.7124	129.4248	129.4248	2847.3451	3.9546	55937.3894
8508.00	63.0000	126.0000	142.0000	2836.0000	4.0000	55632.0000

People	Flats	Row houses	Duplexes	Students	Schools	Cost
8529.54	55.1113	110.2226	214.7934	2843.1808	3.9489	.0133
8526.00	55.0000	111.0000	213.0000	2842.0000	3.9525	(5.6200)
8496.00	56.0000	112.0000	204.0000	2832.0000	4.0000	(64.0000)

Less money means fewer people. How many fewer? That depends on whether we study the fractional or the integer solution:

Fractions: we reduce 55937.3894 by 937.3894 Rand and we have 8542.04 - 8529.54 = 12.50 fewer people, so

12.50/937.3894 = .0133, *Shadow price*, people/Rand, or

937.3894/12.50 = R75 thousand, *marginal cost of one bed.*

Changing the budget from 55000 to 55001 adds .0133 people to the current total -- if we *can* accept fractions. If not, the integer solution makes quite a different story:

Once we insist on R55000 or less, *we cannot spend it all*, but 64 (000) are left over. We spend 54936 thousand Rand.

Upper ceiling on ratio

Going integer with limited funds has a second unexpected result: we build over 200 duplexes, almost four times as many as flats. This may be all for the better, but if we did not like too many flats to begin with, we may feel equally about the other types. None should dominate.

Duplexes tend to be over-produced, so we need not insist on "*at least* twice as many" (the constraint becomes redundant). Rather, make it "*at most* three times as many as flats":

Table 15. Integer data with cost and duplex ceiling

People	Flats	Row houses	Duplexes	Students	Schools	Cost
8496	56	112	204	2832	4	(64)
8454	56	125	168	2818	4	(22)

We still have "at least twice as many row houses" (they rose from 112 to 125, to compensate for fewer duplexes). The main objective is hardly changed (8454 people vs 8496). It would change even less (at 8466) with large and small schools.

Critical or dominant variables

Why did we have so many duplexes after adding a cost ceiling to our data? Because they are cheaper than flats, cheaper in their construction cost per person (whereas flats use land more effectively, see Table 5). In other words, as long as money is unlimited but land *is* limited, we make *only* flats; in the opposite case, we build *only* duplexes:

Table 16. Limits on either land or money, but not both

People	Flats	Row houses	Duplexes	Students	Schools	Total Cost
8942.68	186.3057	None	None	2980.8917	4.1401	61406.3694
8940.00	186.0000	None	1.0000	2980.0000	4.1425	61378.8600
8304.00	173.0000	None	None	2768.0000	5.0000	57936.0000

People	Flats	Row houses	Duplexes	Students	Schools	Total Land
9375.00	None	None	781.2500	3125.0000	4.3403	7204.8611
9372.00	None	None	781.0000	3124.0000	4.3611	7211.4444
9276.00	None	None	773.0000	3092.0000	5.0000	7411.0000

Ignoring the need to mix house types (but including indirect demand for schools and eventually hospitals etc.), we ask:

Given 65 hectares, how much money *can* we spend? Or again,

given 55 million Rand, how much land could we occupy?

The answer depends on whether schools are standard-size or a mix of large and small sizes. But in either case, it clearly favors one type of housing, the *critical* or *dominant* one. We introduce that term to be able to formulate this rule:

an if/then-relation must depend on a dominant variable.

Thus, in the first part of our example (*limited land*), *flats dominate*: we have only flats and no duplexes. Here it makes sense to say "Duplex > 2*Flats", *the number of duplexes must be twice as large as the number of flats*. This constraint is *binding*. It forces duplexes into the solution, in spite of their using land less efficiently than flats.

The same demand, "Duplex > 2*Flats", is useless in the other case (*money limited, land free*). With land freely available, we build only duplexes and no flats. This means *inherently* more duplexes than flats. "Duplex > Flats" cannot create a relation that exists already: the constraint has become ineffective, redundant, or *non-binding*.

So if we had too many duplexes after limiting our budget, it was because we neglected dominance: duplexes being cheaper, we should not build any flats, but *only* duplexes. If we had any flats at all, it is because eventually *both* land *and* money are limited, as we shall see below.

We summarize first if/then-statements respecting dominance:

Table 17. Summary of if/then respecting dominance

Flats dominate and we want twice as many duplexes:

*Duplex > 2*Flat* expresses the idea, but we rewrite

-2*Flat + 1*Duplex > 0 to combine all variables. Then

2*Flat - 1*Duplex < 0, change signs to get "<".

Duplexes dominate and we want half as many flats:

*Flat > 1/2*Duplex* expresses the idea, but we rewrite

1*Flat - 1/2*Duplex > 0 to combine all variables. Then

2*Flat - 1*Duplex > 0, multiply by 2 to kill fraction,

-2*Flat + 1*Duplex < 0, change signs to get "<".

Use one of the statements in the context of other data:

max 48 Flat	+ 12 Dup	+ ...	*other variables, schools etc.*	
2600	700	...	< 650000	Square meters
312	66	...	< 55000	Cost thousands
2	-1	...	< 0	Quota on dominant Flats
-2	1	...	< 0	Quota on dominant Duplex

We transform all relations to "<", *less than or equal to*, because that is how they finally appear in the (numerical) LP-table. It also creates a symmetry easy to remember: For a given 2 to 1 relation, these *same* constants appear in the *same* positions, under Flat and Duplex, but --

a dominant type is positive, the other one is negative.

In other words, it is as easy to insist on "twice as many duplexes" as on "half as many flats": it is a mere matter of making the dominant variable positive -- *once we know the one that dominates.* Unfortunately, this is not clear *before* the fact, from inspecting the data. Rather, we must first run the data *without* insisting on any specific proportions. Then, *after the fact*, when we see unacceptable results, we add a suitable constraint to bring them into line.

The only alternative to the step-by-step procedure, that is, to add constraints of proportion to the initial table, is to insist on *exact* ratios ("exactly twice as many" rather than "at least ..."). However, these do not generally permit as good a solution as the open constraint (Table 18, below).

Several variables dominant

As a rule, each resource implies a dominant variable, the one using that scarce resource best. In our initial example, we had a land and a cost constraint. The former makes flats dominant, the latter, duplexes.

Table 18 summarizes their interaction without special constraints, then with suitable two-to-one ratios:

Table 18. Limits on land *and* money, without/with ratios

People	Flats	Row houses	Duplexes	Students	Schools	Cost
8679.16	85.0472	None	383.0743	2893.0530	4.0181	.0411
8670.00	84.0000	1.0000	380.0000	2890.0000	4.0300	(22.2400)
8640.00	70.0000	None	440.0000	2880.0000	4.0000	(952.0000)

"At least twice as many row houses as flats, half as many flats as duplexes", *dominance*!

People	Flats	Row houses	Duplexes	Students	Schools	Cost
8398.89	63.6280	127.2559	127.2559	2799.6298	3.8884	.1527
8394.00	62.0000	131.0000	124.0000	2798.0000	3.8861	(28.2000)
8382.00	61.0000	133.0000	122.0000	2794.0000	4.0000	(10.0000)

"Exactly twice as many row houses and duplexes as flats"

People	Flats	Row houses	Duplexes	Students	Schools	Cost
8398.89	63.6280	127.2559	127.2559	2799.6298	3.8884	.1527
8316.00	63.0000	126.0000	126.0000	2772.0000	3.8500	(542.8000)
8316.00	63.0000	126.0000	126.0000	2772.0000	4.0000	(424.0000)

We first respect all constraints other than the need to mix house types. As expected, flats *and* duplexes are critical, but row houses are not recommended. They use both land and money less efficiently than the other types.

Shall we leave the solution as it stands, with 70 flats and 440 duplexes? Shall we insist that one type appear *at least* or indeed *exactly* twice as often as another?

In the latter case, we can add equality-constraints to the initial table and solve in one run, but we lose room for 66 people. In the other case, we must make two runs to discover the dominant variables, then add open constraints....

Whatever the planner's decision, we can explore its implications. In particular, we can study various ratios and their indirect effects. We can also provide for large and small schools, or for those of standard size.

The more stringent such demands, the more funds remain idle. Exact proportions allow fewer houses than open proportions. So do schools of equal size as compared to a variable size. And so many houses can accommodate so many pupils. A school too big for them just remains partially empty.

Summary of data

Here are the data as they were run to produce Table 18:

48 Flt + 30 Row + 12 Dup + 0 Std + 0 Sch

26	18	7	0	400	<	6500	Sqm/100
312	186	66	0	792	<	55000	Cost/1000
16	10	4	-1	0	<	0	Children
16	10	4	0	-720	<	0	Pupils for school
/							
/							
2	-1	0	0	0	<	0	RiF
2	0	-1	0	0	<	0	DiF*
-2	0	1	0	0	<	0	FiD

We used three-letter labels such as *Flt* for "Block of flats" or *RiF*, "Row house if Flat". Indeed, we had "Row to be *twice* the number of flats". This way of speaking associates *twice* or 2 correctly with the Flat-column.

We separated constraint groups by two slashes: these exclude subsequent data from a run, but keep them at hand. Here we split constraints of *substance* (land, money, schools) from those of *proportion*. In running the substantive part first, we get to know the *dominant* variables. Once we know these, after the fact, we add the desired if/then-proportions. The advantage of proceeding step-by-step is this:

we obtain generally a better solution than if we add the corresponding (strict) equalities to the initial data;

we add only those few if/then-statements that are *binding*, not the multitude of all possible relationships.

In turn, if a statement does not seem to be effective, it is not always due to poor logic: more likely, *another variable became dominant*. For example, X dominates if we may have it on its own; but if we force Y to go with X, then we get only Z on its own. Z dominates X and Y *bound together* when either X or Y on their own dominate Z! In such a case, simply add an if/then as a function of Z to the previous statements.

Here we discover that constraint DiF *on its own* is binding before we impose a cost-limit. It is redundant otherwise -- unless we keep *both* DiF and FiD: then we get the proportion "*exactly* twice as much of one as of the other".

Strictly speaking, the *Children*-row and the *Student*-column are redundant: they cannot change the solution, but they merely *count* the number of children in any given house-mix. Still, that is useful. It is equally useful that the small numbers in this constraint may speed up the transition from fractional to integer solutions.

Conclusions

We applied well-established LP-techniques in the context of township planning. We endeavored to present these techniques so as to make them accessible to the practitioner.

Specifically, we formulated cause-effect relations such as "if flats, then row houses" or "if houses, then people, then children, then schools". More generally, people want schools and similar public amenities. Yet as we build more schools, less space remains for houses. This means fewer people and less demand for schools. LP *optimizes* such interactions so as to house as many people as possible while yet providing all the schools and amenities they require.

In doing so, we paid particular attention to the transition from *fractional* to *partial* and *all-integer* solutions. *Houses* naturally must appear as integer numbers, but *schools* need not, not if we admit schools of *variable size.* For example, we cannot have 3.50 schools of size 720, but we *can* build three of 720 and one of 360, or four of 605. Such seemingly minor concessions affect optimality.

So does the need to insist on *exact* or *open* proportions ("*at least* twice" or "*at most* twice" instead of "*exactly* twice"). Exact proportions have the advantage that they can be added to the *initial* data set: being exact, naturally they will be solved exactly, regardless of cost. *Open* proportions, on the other hand, are more cost effective, but they must be formulated in terms of *dominant* variables. These are not usually known beforehand, but only from an exploratory run (*without* the new constraint). Thus they require more effort from the planner. Yet given the amounts at stake in a township, that effort is likely to be justified.

We have endeavored to provide the tools so that, if anybody cares to make the effort, he can do so effectively.

References

[1] Any linear programming routine providing for integer solutions will handle the problems discussed here. As a matter of convenience, the authors provide, at nominal cost, an MS-DOS disk with LP and LP7 (for 8087 co-processor) as well as printed documentation, *From simultaneous equations to linear and integer programming*, 214 pp.