ECOLE NORMALE SUPERIEURE

Solving Scheduling Problems

by Simulated Annealing

Olivier CATONI

LMENS - 96 - 20

Solving Scheduling Problems
by Simulated Annealing

Olivier CATONI

Laboratoire de Mathématiques de l'Ecole Normale Supérieure
45 rue d'Ulm 75230 PARIS Cedex 05

Tel : (33)(1) 44 32 00 00
Adresse électronique : catoni@dmi.ens.fr

Solving Scheduling Problems by Simulated Annealing

Olivier Catoni

DIAM – Laboratoire de Mathématiques, URA 762 du CNRS
Département de Mathématiques et d'Informatique
Ecole Normale Supérieure
45 rue d'Ulm, 75 230 PARIS CEDEX 5

Summary. We define a general methodology to deal with a large family of scheduling problems. We consider the case when some of the constraints are expressed through the minimization of a loss function. We study with more details a benchmark example consisting in some jigsaw puzzle problem with additional constraints. This is the occasion to discuss some algorithmic issues typical of scheduling problems, such as the apparition of small unused gaps or the representation of proportionality constraints. This is also an opportunity to carry on an experimental comparison between the Metropolis algorithm, simulated annealing and the iterated energy transformation method, and to see whether asymptotical theoretical results are a good guide towards practically efficient algorithms.

Mathematics Subject Classification (1991): 60J10, 90C42

Introduction

The aim of this paper is to describe a general strategy to deal with scheduling problems and to illustrate its use on an example: the resolution of jigsaw puzzles. We will assume we can put our scheduling problem in the form of a task assignment problem, and we will turn it into the minimization of a cost function defined on a suitable search space. This cost function will be minimized by a Monte Carlo algorithm of the Metropolis kind: either simulated annealing or our recently introduced iterated energy transformation method. We have already studied some of the theoretical aspects of these two methods in previous papers (see [5], [6]).

We have chosen to experiment on a jigsaw puzzle problem with rectangular pieces, because this is a typical instance of the kind of difficulties encountered when building time-tables, and because it is in itself a difficult problem (it is N.P.-complete), which deserves a special interest. In the course of this experimentation, we will compare four algorithms: a randomized descent algorithm (the Metropolis dynamic at temperature zero), the Metropolis algorithm, simulated annealing and the iterated energy transformation algorithm.

1. An abstract task assignment framework

Let B be a finite set of tasks. Let E be a set of resources needed to perform these tasks. The set E may be any kind of set, a finite set, a domain in $\mathbb{R}^N$ etc.. In applications it can represent various things, such as a set of people who are to perform the tasks, in which case it is natural to see it as a finite set, or it can represent also space and time needed for the tasks, in which case it is sometimes natural to see it as a domain in $\mathbb{R}^n$. More often it is a product space of both kinds. Anyhow, we will only consider a finite collection of subsets of E, therefore it will be always possible to consider that E is a finite set from the theoretical point of view. This is reasonable since a computer can only handle a finite number of possible ways to allocate resources, and also because in many problems of the time-table type, continuous quantities, such as time, are discretized (for instance when one tries to schedule lectures, they are usually constrained to start at full hours). Anyhow the reader should think of E as of a large set and our methods, inherited from statistical mechanics, are precisely meant to cope with a large state space.

The abstract scheduling problem we will consider is to allocate to each task in B a set of resources in a way which satisfies a set of constraints.

At this level of generality, we will not represent the constraints by equations or logical relations, we will merely view them as a subset $\mathcal{S}$ of $\mathcal{P}(B \times E)$ (where $\mathcal{P}(A)$ is the set of subsets of the set A). We will call $\mathcal{S}$ the "solution space". A solution x in $\mathcal{S}$ is a subset of the product space $B \times E$. We will use the notations π_B and π_E for projections on B and E. The fact that $(b, e) \in x$ means that the task b uses the resource e. The set of resources used by b is $\pi_E(\pi_B^{-1}(b) \cap x)$, for which we will use the functional notation $x(b)$.

We will assume that each solution $x \in \mathcal{S}$ is a complete assignment, in the sense that all the tasks are scheduled:

$$\pi_B(x) = B \quad \text{for any } x \in \mathcal{S}.$$

Our scheduling problem is to construct a solution x belonging to the solution space $\mathcal{S}$.

2. The jigsaw puzzle example

This example is meant to be a benchmark where the main algorithmic issues of scheduling problems are present.

The set of resources E will be a discretized rectangular frame

$$E = \{0, \ldots, M-1\} \times \{0, \ldots, N-1\} \subset \mathbb{Z}^2.$$

The set of tasks B will be the set of pieces of the jigsaw puzzle. Each piece r has a rectangular shape defined by its width w_r and by its height h_r. The constraint is that pieces should not overlap. Thus the solution space is

$$\mathcal{S} = \{x \subset B \times E \ : \ x(r) = [a_r, a_r + w_r[\times[b_r, b_r + h_r[, \quad (a_r, b_r) \in \mathbb{Z}^2, \quad r \in B,$$
$$\text{and } x(r) \cap x(r') = \emptyset, \quad r \neq r' \in B\}.$$

The problem is to build the jigsaw puzzle, that is to construct $x \in \mathcal{S}$. Although the shape of pieces is very simple, this problem can be seen to be very complex. In fact, it is easy to see that it is N.P. complete, because it contains the partition problem among its instances. Indeed, the partition of given integers $\{c_1, \ldots, c_N\}$ into two sets I and J such that

$$\sum_{i \in I} c_i = \sum_{j \in J} c_j$$

can be viewed as a jigsaw puzzle with N pieces of width c_i, and height 1, and a frame of width $\dfrac{1}{2} \sum_{i=1}^{N} c_i$ and height 2.
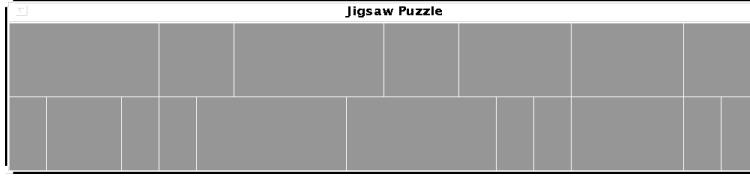
3. A method of resolution based on the Metropolis dynamic

In this section we will sketch a methodology to solve the abstract problem of section 1. The general idea is to perform a random search for a solution in a state space larger than the solution space. This search space should be easy to describe and easy to search by a Markov chain performing a succession of elementary moves. Of course, we will not use a Markov chain which samples uniformly the search space, because usually the search space we will be able to build will be very large when compared to the solution space, and drawing points at random in the search space would lead to discover a solution hopelessly seldom.

Instead we will use a Markov chain with rare transitions, whose invariant measure is concentrated in a neighbourhood of the solution space. This optimization technique is well known, but its improvement is still a subject of active research. The prototype algorithm we will start from is the Metropolis dynamic at low temperature. The Metropolis dynamic has been designed to simulate statistical mechanics systems, and not for optimization purposes. In order to improve its performance as an optimization algorithm, some speed-up techniques have been proposed. The most famous one is simulated annealing. We have also proposed recently another one, which we called "the iterated energy transformation method"[6]. We will describe and use both of these.

3.1. Choice of a search space

The first step of the method is to choose a search space $\tilde{\mathcal{S}}$ containing the solution space $\mathcal{S}$. The most popular way to construct $\tilde{\mathcal{S}}$ is to relax some constraints about the solution and to measure instead how much the constraints have been violated by a score function one has afterwards to minimize. For instance in circuits placement applications (one of the earliest applications of simulated annealing) the constraint that circuits should not overlap is often relaxed, and the overlapping of circuits is instead merely discouraged by some score function of the surface of the overlap. Our strategy will be somehow of the same kind, with the difference that we will not relax a constraint which is specific to the problem. Instead, we will allow partial solutions, where only some proportion of the tasks have been scheduled. Defining partial solutions is usually very easy and very natural. Most of the time, this is how the problem is posed from the beginning. Indeed the constraints come usually from incompatibilities between tasks, such as sharing the same resource or needing to be performed in a given order, and can be expressed without assuming that all the tasks are already scheduled.

From the technical point of view, we will assume that the search space (the space of partial solutions) satisfies the following properties:

- The empty solution is in the search space: $\emptyset \in \tilde{\mathcal{S}}$.
- There is a path from the empty solution leading to any partial solution $x \in \tilde{\mathcal{S}}$ along which tasks are scheduled one after the other. This can be expressed in the following way:

    For any $x \in \tilde{\mathcal{S}}$, $x \neq \emptyset$, there is $b \in \pi_B(x)$ such that $x \setminus \pi_B^{-1}(b) \in \tilde{\mathcal{S}}$.

- All complete solutions in the search space satisfy the constraints. In other words, the solution space is exactly made of the complete solutions of the search space. This is expressed by the following equation:

$$\mathcal{S} = \{x \in \tilde{\mathcal{S}} \; : \; \pi_B(x) = B\}.$$

Let us notice that the "best" choice for $\tilde{S}$ would be $\{x \cap \pi_B^{-1}(C) \; : \; x \in \mathcal{S}, C \subset B\}$, the set of all partial solutions contained in global solutions. Anyhow, this set is in practical situations never defined by simple relations, because when you have scheduled some of the tasks, it is never possible (except for trivial problems) to foretell whether there will remain suitable resources to schedule the remaining ones. Therefore the search space $\tilde{S}$ is most of the time much broader than $\mathcal{S}$ and contains many dead ends.

3.2. Building the dynamic: constructions and destructions

The next idea is to define on $\tilde{S}$ two kinds of dynamics, a *constructive dynamic* and a *destructive dynamic*. These two random dynamics are characterized by two Markov matrices q_C and q_D,

$$
\begin{aligned}
q_C &: \quad \tilde{\mathcal{S}} \times \tilde{\mathcal{S}} \longrightarrow [0,1] \\
q_D &: \quad \tilde{\mathcal{S}} \times \tilde{\mathcal{S}} \longrightarrow [0,1]
\end{aligned}
$$

We will assume that the transitions allowed by q_C consist either in keeping the current partial solution or in scheduling one more task. In a similar way the transitions allowed by q_D consist in unscheduling a given number of tasks. We allow to unschedule more than one task at a time, because it is in some situations more sensible to do so. For instance, if many tasks have to share the same resource, it may sometimes speed-up the allocation process to unschedule all of them at the same time (think of students sharing the same teacher ...).

This conception of constructions and destructions can be expressed by the following equations, where $|A|$ is the number of elements in the finite set A:

$$
(1) \quad
\begin{cases}
- \ \{(x,y) \ : \ q_C(x,y) > 0,\, x \neq y\} \\
\qquad = \{(x,y) \ : \ y \cap \pi_B^{-1}(\pi_B(x)) = x,\, |\pi_B(y)| = |\pi_B(x)| + 1\}, \\
- \ \{(x,y) \ : \ q_C(y,x) > 0,\, x \neq y\} \\
\qquad\qquad \subset \{(x,y) \ : \ q_D(x,y) > 0\} \\
\qquad\qquad\qquad \subset \bigcup_{n=1}^{+\infty} \{(x,y) \ : \ q_C^n(x,y) > 0\}.
\end{cases}
$$

Let us remark that usually constructions will decompose into two steps, one being to choose an unscheduled task $b \in B \setminus \pi_B(x)$ and the second one being to try to allocate to it a set of resources. This second step is sometimes unsuccessful (either because it is impossible or the proper allocation has not been discovered), therefore, as a rule, we have $q_C(x,x) > 0$ for a substantial number of partial solutions. On the contrary, destructions are simple moves, where you have only to choose a scheduled task b in $\pi_B(x)$ and to remove it. Therefore, as a rule, we will have $q_D(x,x) = 0$, except when $x = \emptyset$, for which $q_D(\emptyset, \emptyset) = 1$.

When the two above assumptions are satisfied, the whole search space $\tilde{\mathcal{S}}$ can be constructed by q_C starting from the empty solution $\emptyset$, and reversely, any solution can be shrunk to the empty solution by successive applications of q_D. More precisely, the following proposition holds:

Proposition 1.

$$
\begin{aligned}
\tilde{\mathcal{S}} &= \bigcup_{n=0}^{+\infty} \{y \ : \ q_C^n(\emptyset, y) > 0\} \\
&= \bigcup_{n=0}^{+\infty} \{x \ : \ q_D^n(x, \emptyset) > 0\}
\end{aligned}
$$

3.3. Building the cost function

Now we will build a cost, or energy, function defined on the search space, which penalizes partial solutions: we will call it $U : \tilde{\mathcal{S}} \longrightarrow \mathbb{R}$. Namely we will require the following properties to hold:

(2)
$$
\begin{cases}
- \; \arg\min_{x\in\tilde{\mathcal{S}}} U(x) = \mathcal{S}, \\
- \; \text{There is a positive constant} \\
\quad \gamma \text{ such that } U(y) \geq U(x) + \gamma \\
\quad \text{when } x \neq y \text{ and } q_D(x,y) > 0.
\end{cases}
$$

A typical example for U is

$$
U(x) = \mu(B \setminus \pi_B(x)),
$$

where μ is a positive measure on B. In this case the assumptions on U are satisfied and the largest choice of γ is

$$
\gamma = \min_{b\in B} \mu(b).
$$

3.4. Building a Metropolis dynamic

From Prop. 1, we see that any Markov matrix of the form $\lambda q_C + (1-\lambda)q_D$ with $\lambda \in]0,1[$ is irreducible. Therefore a straightforward way to build a Metropolis dynamic would be to consider the Markov matrix

$$
p_T(x,y) = \begin{cases}
(\lambda\, q_C(x,y) + (1-\lambda)\, q_D(x,y))e^{-(U(y)-U(x))^+/T}, & x \neq y \\
1 - \sum_{z\neq x} p_T(x,z), & x = y.
\end{cases}
$$

In fact we can do better, using the fact that we know in advance that during a construction the energy will decrease and that during a destruction the energy will increase by a quantity at least equal to γ. This avoids applying uselessly the kernel q_D at low temperatures, in situations where we know that it will most of the time generate a move to be rejected.

More precisely, we will use the following Markov matrix:

(3)
$$
\begin{aligned}
p_T(x,y) = \lambda\, &e^{-\gamma/T} q_D(x,y)e^{-(U(y)-U(x)-\gamma)^+/T} \\
&+ q_C(x,y)\left(1 - \lambda \sum_{z\in\tilde{\mathcal{S}}} q_D(x,z)e^{-(U(z)-U(x)-\gamma)^+/T-\gamma/T}\right),
\end{aligned}
$$

where λ is again a positive parameter in the interval $0 \leq \lambda \leq 1$. The presence of λ avoids that destructions should always be chosen at high temperatures. Usually we will take $\lambda = 1/2$ or $\lambda = 1$. The positive part in $(U(y)-U(x)-\gamma)^+$ is needed only to cover the case when $x = y$.

The computer implementation of this Metropolis dynamic is the following: starting from the state x,

– First flip a coin with odds $\lambda e^{-\gamma/T}$ and $1 - \lambda e^{-\gamma/T}$ to decide whether to try a destruction or not.
– In case a destruction is tried,
  – choose a transition (x, y), drawing y according to the probability distribution $q_D(x, y)$.
  – then flip a second coin with odds $\exp -((U(y) - U(x) - \gamma)^+/T)$ and $1 - \exp -((U(y) - U(x) - \gamma)^+/T)$ to decide whether to apply this move or not.
– If the answer to one of the two previous tosses was *no*, then choose a transition (x, y) where y is chosen according to the distribution $q_C(x, y)$ and apply it.

The hypothesis we made about q_C, q_D and U are what is needed to prove the following proposition:

Proposition 2. *For any temperature $T > 0$, the matrix p_T is an irreducible Markov matrix.*

Considering the rate function $V : \tilde{\mathcal{S}} \times \tilde{\mathcal{S}} \rightarrow \mathbb{R}_+ \cup \{+\infty\}$ defined by

$$V(x, y) = \begin{cases} (U(y) - U(x))^+ & \text{if } q_D(x, y) + q_C(x, y) > 0 \text{ and } x \neq y, \\ +\infty & \text{otherwise} \end{cases}$$

we see that there is a positive constant κ such that whenever $x, y \in \tilde{\mathcal{S}}$, $x \neq y$,

$$\kappa\, e^{-V(x,y)/T} \leq p_T(x, y) \leq \frac{1}{\kappa}\, e^{-V(x,y)/T}.$$

Moreover V satisfies the weak reversibility condition of Hajek-Trouvé with respect to U. More precisely if $\Gamma_{x,y}$ is the set of paths from x to y, we put for any $\gamma = (\gamma_1 = x, \ldots, \gamma_r = y) \in \Gamma_{x,y}$

$$H(\gamma) = \max_{i=1,\ldots,r-1} U(\gamma_i) + V(\gamma_i, \gamma_{i+1})$$

and

$$H(x, y) = \min_{\gamma \in \Gamma_{x,y}} H(\gamma).$$

The weak reversibility condition of Hajek-Trouvé states that for any $x, y \in \tilde{\mathcal{S}}$

$$H(x, y) = H(y, x).$$

Due to this reversibility property, U is a quasi-potential for p_T. We mean by this statement that the (unique) invariant probability measure μ_T of p_T satisfies for some positive constant α (independent of T) and for any $x \in \tilde{\mathcal{S}}$

$$\alpha \leq \mu_T(x) e^{(U(x) - \min U)/T} \leq 1/\alpha.$$

Corollary 1. *We can build optimization algorithms based on p_T following the results of Catoni [6] and Trouvé [20]. More precisely, for any fixed value of the temperature T, the homogeneous Markov chain with transition matrix p_T is a generalized Metropolis algorithm with quasi-potential function U. In the same way, for any decreasing sequence of temperatures $(T_n)_{n \in \mathbb{N}}$, the non-homogeneous Markov chain $(X_n)_{n \in \mathbb{N}}$ on $\tilde{\mathcal{S}}$ with transitions*

$$P(X_n = y \; : \; X_{n-1} = x) = p_{T_n}(x, y)$$

is a generalized simulated annealing algorithm. Its behaviour has been studied in [20], [22] and is very similar to the behaviour of classical simulated annealing as studied in [5].

We can also apply the iterated energy transformation method to p_T, which will be described in a further section of this paper and is studied in [6].

Proof:

The only non straightforward point to check is the Hajek-Trouvé weak reversibility condition. Let us consider $x, y \in \tilde{\mathcal{S}}$ and $\gamma \in \Gamma_{x,y}$. We build a path from y to x in the following way. Replace any edge $(z, t) \in \gamma$ by the edge (t, z) if $q_C(z, t) > 0$ or $p_T(z, t) = 0$. If non of the above two conditions is true, this means that $q_D(z, t) > 0$, then there is a path $\varphi \in \Gamma_{t,z}$ such that $q_C(u, v) > 0$ for any edge $(u, v) \in \varphi$, and we replace (z, t) by φ. The path φ is such that $H(\varphi) = U(z) + V(z, t) = U(t)$ because for any $(u, v) \in \varphi$, $U(v) < U(u)$ and $V(u, v) = 0$. Therefore by concatenating all these reversed edges and paths in reverse order we get a path $\psi \in \Gamma_{y,x}$ such that $H(\psi) \leq H(\varphi)$. Therefore $H(y, x) \leq H(x, y)$ and consequently $H(y, x) = H(x, y)$. $\square$

Remarks:

- In the search space we consider, there is a natural starting point for optimization algorithms, which is the empty schedule $\emptyset$.
- In many scheduling problems, it is not known in advance whether a complete solution exists, or whether one can possibly be found within the available computer time. Our method has the advantage to find at least a partial solution, where some proportion of the tasks are scheduled in a coherent way. This is not the case if other constraints are relaxed as is usually done. For instance, if the aim is to schedule the lectures in a University, a solution where some lectures share the same room at the same time has no practical interest, whereas a solution where some proportion of the lectures are scheduled in a coherent way can be applied.
- A slight variant of the present set-up is the case when the search space satisfies condition (1), but one does not know whether it is possible to schedule all the tasks, and wants instead to schedule as many tasks as possible. In this situation the energy can weight (through a positive measure) the relative importance of tasks.

In the three following sections, we are going to recall briefly some theoretical results about the speed of convergence of three optimization algorithms.

3.5. Rate of convergence of the Metropolis algorithm

In this section we consider the canonical process $(X_n)_{n \in \mathbb{N}}$ on the canonical space $(\tilde{\mathcal{S}}^{\mathbb{N}}, \mathcal{B})$ defined by $X_n(x) = x_n$, where $\mathcal{B}$ is the sigma field generated by the events depending on a finite number of coordinates.

For any temperature $T \in \mathbb{R}_+$, P_T will be the probability distribution on $(\tilde{\mathcal{S}}^{\mathbb{N}}, \mathcal{B})$ of a Markov chain with transition matrix p_T (where p_T is as in Prop. 2). Under this distribution, $(X_n)_{n \in \mathbb{N}}$ is a Metropolis algorithm and has the following convergence speed:

Proposition 3. *There exists a positive constant d, depending only on the choice of the search space $\tilde{\mathcal{S}}$, of the constructive and destructive dynamics q_C and q_D and of the parameter λ, $0 < \lambda < 1$, such that for any energy function U satisfying the hypothesis (2) of section 3.3, for any positive constant η,*

$$\max_{x \in \tilde{\mathcal{S}}} P_T(U(X_N) \geq U_{\min} + \eta \mid X_0 = x)$$

$$\leq d \left(\exp - \left(\frac{N}{d} e^{-H_1/T} \right) + e^{-\eta/T} \right),$$

where $H_1(V)$ is the first critical depth of the rate function V defined in Prop. 2. The exponent $H_1(V)/T$ is optimal when η is small, and when T tends to 0 and N tends to $+\infty$. With the notations of proposition 2,

$$H_1(V) = \max_{x \notin \mathcal{S}} \min_{y \in \mathcal{S}} H(x, y) - U(x).$$

As a consequence, considering $1/T = \dfrac{1}{H_1} \log \left(\dfrac{N H_1}{d \eta \log N} \right)$, we see that there is a constant d (independent of U and η), such that

$$\inf_{T \in \mathbb{R}_+} \max_{x \in \tilde{\mathcal{S}}} P_T(U(X_N) \geq U_{\min} + \eta \mid X_0 = x)$$

$$\leq d \left(\frac{d \eta}{H_1(V)} \frac{\log N}{N} \right)^{\eta/H_1(V)}.$$

Moreover the exponent $\eta/H_1(V)$ is optimal for small enough values of $\eta \in (U(E) - U_{\min})$.

For a proof see for instance Cot and Catoni [8].

This proposition shows that the convergence speed of the Metropolis algorithm is slow when there are states with energies close to $U_{\min}$. Indeed if one wants to study the convergence to $\mathcal{S}$, one has to choose

$$\eta = \min\{U(x) - U_{\min} , x \in \tilde{\mathcal{S}} \setminus \mathcal{S}\}.$$

If η is small, then it will reflect on the exponent $\eta/H_1(V)$.

This is a theoretical justification for the introduction of simulated annealing, which will not suffer from this draw-back, when proper robust cooling schedules are used.

3.6. Rate of convergence of simulated annealing

We consider now a non-increasing triangular sequence $T_1^N \geq T_2^N \geq \cdots \geq T_N^N$ of temperatures and the measure $P_{(T_1^N, \ldots, T_N^N)}$ on $\tilde{\mathcal{S}}^N$ of the non-homogeneous Markov chain with transitions

$$P_{(T_1^N, \ldots, T_N^N)}(X_n = y \mid X_{n-1} = x) = p_{T_n^N}(x, y).$$

The rate of convergence of such an algorithm has been studied in [5] and [20] (translated into English in [22]). We give here a simple result, for more precise estimates, we refer to the original papers.

Proposition 4 (Catoni, Trouvé). *There is a positive constant K such that*

$$K^{-1}N^{-D^{-1}} \leq \inf_{T_1^N \geq \cdots \geq T_N^N} \max_{x \in \tilde{\mathcal{S}}} P_{(T_1^N, \ldots, T_N^N)}(U(X_N) > U_{\min} \mid X_0 = x) \leq KN^{-D^{-1}},$$

where the constant $D = D(V)$ is the difficulty of the rate function V. With the notations of proposition 2, the definition of $D(V)$ is

$$D(V) = \max_{x \in \tilde{\mathcal{S}} \backslash \mathcal{S}} \min_{y \in \mathcal{S}} \frac{H(x, y) - U(x)}{U(x) - \min U}.$$

For any $A > 0$, there is a positive constant K such that the triangular exponential schedule

$$T_n^N = \frac{1}{A}\left(\frac{A}{(\log N)^2}\right)^{n/N}$$

gives a convergence speed of

$$\max_{x \in \tilde{\mathcal{S}}} P_{T^N}(U(X_N) > U_{\min} \mid X_0 = x) \leq K\left(\frac{\log N \, \log \log N}{N}\right)^{D(V)^{-1}},$$

for N large enough.

In the case of simulated annealing, we have $\epsilon \asymp_{\log} \left(\dfrac{1}{N}\right)^{1/D}$ (meaning that the logarithms of both sides of this equation are equivalent when N tends to infinity). The important feature of this theoretical result is that the exponent $1/D$ is independent of the precision with which we want to reach $U_{\min}$, but depends on the contrary only on the structure of the local minima of U.

One other interesting point is that the exponential triangular cooling schedule $T_n^N = A^{-1}\left(\dfrac{A}{(\log N)^2}\right)^{n/N}$ is robust: it gives a convergence rate with the optimal exponent $1/D(V)$ for any energy function U.

3.7. Rate of convergence of the energy transformation method

We introduced in [6] the iterated energy transformation method as another mean to discourage uphill moves from low energy states more than from high energy states. In simulated annealing this effect is produced by an exogenous control of the temperature parameter: in "typical" successful runs of simulated annealing, the energy of the current state is moving downwards on the average, and in the same time uphill moves are more and more discouraged. In the iterated energy transformation method, a temporary hypothesis is made about the value of $U_{\min}$, and a concave transformation is applied to U on the basis of this hypothesis. Then the algorithm is run at constant temperature using the transformed energy. This produces the desired effect of discouraging more uphill moves from low energy states. Of course, in the beginning, the hypothesis about $U_{\min}$ is necessarily grossly underestimated, so that the energy transform is not very efficient, but after some iterations, it can be improved (this will work with a probability close to one) depending on the values of the energies of the explored states.

The convergence of the lower bound estimate for $U_{\min}$ towards the true value of $U_{\min}$ is exponentially fast (with a probability close to one), and therefore, the energy transformation is quickly tuned to an efficient value.

The iterated energy transformation method applied to our problem describes as follows. For any strictly concave, strictly increasing energy transformation $F : [U_{\min}, +\infty[\longrightarrow \mathbb{R}$, we consider the Markov matrix

$$p_F(x, y) = \lambda e^{(F(U(x)) - F(U(x) + \gamma))} q_D(x, y) e^{(F(U(x) + \gamma) - F(U(y)))^-}$$

$$+ q_C(x, y) \left(1 - \lambda \sum_z q_D(x, z) e^{F(U(x)) - F(U(x) + \gamma) + (F(U(x) + \gamma) - F(U(z)))^-} \right).$$

Consider for any positive constant α, any real shift τ such that $U_{\min} + \tau > 0$, any positive temperature T, the transformation

$$F_{\alpha, T, \tau}(U) = \alpha\, U + \frac{1}{T} \log(U + \tau).$$

Let us introduce the simplified notation $\overline{p}_{\alpha, T, \tau} = p_{F_{\alpha, T, \tau}}$.

Given parameters $M \in \mathbb{N}$, (number of iterations performed with each energy transform), two real numbers $\rho > 0$ and $\eta_0 \geq 0$, (two parameters for the update of the shift τ), and an initial lower bound $\delta < U_{\min}$, we consider the canonical process $(X_n)_{n \in \mathbb{N}}$ on $\tilde{S}^{\mathbb{N}}$ with probability distribution $P_{\alpha, T, M, \rho, \eta_0}$ defined by the following conditional distributions:

$$P_{\alpha, T, M, \rho, \eta_0}(X_n = y \mid (X_0, \ldots, X_{n-1}) = (x_0, \ldots, x_{n-1})) =$$
$$\overline{p}_{\alpha, T, \tau_n(x_0, \ldots, x_{n-1})}(x_{n-1}, y).$$

with

$$\left\{ \begin{array}{rcl} \tau_r & = & \eta_0 - \delta, \quad 0 < r \leq M \\[2mm] \tau_{kM+r} & = & \tau_{kM} - \dfrac{1}{1+\rho} \left(\min_{n, n \leq kM} U(X_n) + \tau_{kM} \right) + \eta_0, \quad 0 < k, 0 < r \leq M. \end{array} \right.$$

We have proved in [6] the following theorem:

Theorem 1 (Catoni). *For any fixed $\alpha > 0$ the family of processes described above satisfies for some positive constants B and K, for any $r \in \mathbb{N}$,*

$$\max_{x \in \tilde{S}} P_{M, T}(U(X_{rM}) \geq \eta \mid X_0 = x) \leq \epsilon,$$

where

$$\begin{array}{rcl} M & = & B \left(\dfrac{\epsilon}{K\,r} \right)^{-\log(1 + \tilde{D}_{\eta_0})/\log(1+\rho)} \log \dfrac{K\,r}{\epsilon}, \\[3mm] & = & \dfrac{B}{T} \log(1+\rho) \left(1 + \tilde{D}_{\eta_0} \right)^{1/T}, \\[3mm] \eta & = & U_{\min} + \rho \left(\dfrac{\rho}{1+\rho} \right)^{r-1} (U_{\min} - \delta + \eta_0) + \eta_0 \rho (1+\rho), \end{array}$$

and where the constant $\tilde{D}_{\eta_0}(V)$ is given by

$$\tilde{D}_{\eta_0}(V) \;=\; \max_{x \in \tilde{\mathcal{S}} \backslash \mathcal{S}} \min_{y \in \mathcal{S}} \frac{H(x,y) - U(x)}{U(x) - U_{\min} + \eta_0}$$

$$<\; D(V).$$

Corollary 2.

$$\limsup_{N \to +\infty} (\log N)^{-2} \log \inf_{T,M,\eta_0,\rho} P(U(X_N) > U_{\min} \mid X_0 = x) \le -\frac{1}{4 \, \log(1+D)}.$$

The interest of this theorem lies mainly in its corollary, which shows that a proper tuning of the parameters leads to a faster scale of convergence speed than the one achieved by simulated annealing (see [6]). This remark of course deals with the comparison of two long runs of both algorithms. For repeated trials of bounded length, which we will consider in a following paragraph, the question of knowing which algorithm is faster is open.

We will discuss practical means of choosing the parameters in connection with the jigsaw puzzle benchmark.

4. Solving jigsaw puzzles

We will illustrate on jigsaw puzzles the different steps of the general method of resolution.

First of all, we have to choose a search space. This will be the set of partial solutions where only some of the pieces are put in the frame.

$$\tilde{\mathcal{S}} \;=\; \{x \subset B \times E \,:\, x(r) = [a_r, a_r + w_r[\times [b_r, b_r + h_r[,$$

$$r \in \pi_B(x), x(r) \cap x(r') = \emptyset, r \ne r' \in \pi_B(x)\}.$$

Let us define now $q_C(x, .)$, the constructive dynamic starting from state x:

- First choose $r \in B \setminus \pi_B(x)$ according to the uniform distribution on this set.
- Then choose $(z, t) \in E \setminus \pi_E(x)$ according to the uniform distribution.
- Then try to expend this germ to a rectangle $[a_r, a_r + w_r[\times [b_r, b_r + h_r[$ of the desired size by adding alternatively a column to the left (or else to the right) and a line to the top (or else to the bottom). If it is not possible to grow the germ to its final size, just abandon the construction.
- Then draw a number k at random in the interval $[0, \text{max\_drift}[$ and move the location of $[a_r, b_r[$ k steps along the direction $(-1,-1)$ (that is to the upper left corner, according to usual image indexing) if there is enough room to do so, or else move it as far as possible in this direction (until it bumps into other pieces).

The last two actions are better described by the following self-explanatory pseudo-C code, where $[a, c\,[\times [b, d\,[$ is the current germ:

```
int expend() {
  a=z;  c=z+1;  b=t;  d=t+1;
  while((test1=(c-a<w))||(test2=(d-b<h))) {
  if (test1&&grow_left()&&grow_right()) return 1;
  if (test2&&grow_up()&&grow_down()) return 1;
  }
```

```
   for (k=rand(0,max_drift);k;k--) {
   if (move_left()&move_up()) break;
   } return 0;
}
```

where the functions `expend()`, `grow_left()`, `grow_right()`, `grow_up()`, `grow_down()`, `move_left()`, and `move_up()` return 0 on success and 1 on failure.

The destructive dynamic q_D is simpler:

- Draw $r \in \pi_B(x)$ at random,
- Form $y = x \cap \pi_B^{-1}(B \setminus \{r\})$, the partial solution where the piece labeled "r" has been removed from the frame.

The mechanism which was chosen for the constructions is meant to discourage the formation of small gaps between pieces. If nothing were done, when the discretization step of the grid is fine, small gaps would be left between the pieces with a large probability and a complete solution to the puzzle, where pieces necessarily stick together, would never be discovered.

We have now to choose an energy function. Here again we will discourage the formation of gaps between pieces by introducing a term proportional to the contact-length. By contact-length, we mean the sum of the contact-lengths between pieces and between pieces and the edge of the frame.

Let μ be the counting measure on E. We take

$$U(x) = -\mu(\pi_E(x)) - \alpha \times \text{contact-length}.$$

For this choice of U, we can take the constant γ in equation (2) to be equal to the size of the smallest piece:

$$\gamma = \min_{r \in B} w_r \, h_r.$$

5. Minimizing a loss function

5.1. Statement of the problem

We will discuss in the next two sections the case when some loss function $V : \mathcal{S} \longrightarrow \mathbb{R}$ has to be minimized on the state space $\mathcal{S}$ of global solutions of a task assignment problem. We consider the same framework as in the first section, with the difference that the problem is now to find a solution x belonging to $\arg\min_{y \in \mathcal{S}} V(y)$.

5.2. A general method of resolution

We will extend the method of section 3 to deal with a loss function.

The two first steps, building the search space and the constructive and destructive dynamics will be the same as in section 3.

The change comes from the choice of the energy function. First we need to extend the loss function V to the search space $\tilde{\mathcal{S}}$ of partial solutions. Ideally, we would like to use the extension $\overline{V} : \tilde{\mathcal{S}} \longrightarrow \mathbb{R}$ defined by

$$\overline{V}(x) = \min\{V(y) \; : \; y \in \mathcal{S}, x \subset y\}.$$

Usually this is not an easily computable function, but in many situations, there is a natural way to define a loss function for partial solutions. A simple way to do so, if there is nothing else at hand, is to set $V(y) = c$ for $y \in \tilde{\mathcal{S}} \setminus \mathcal{S}$, where c is a constant and $c \geq \max_{x \in \mathcal{S}} V(x)$. Then we build a compound energy function

$$W(x) = \alpha\, U(x) + V(x), \quad x \in \tilde{\mathcal{S}},$$

where the real positive coefficient α is chosen such that for some positive constant γ

$$\begin{cases} \arg\min_{x \in \tilde{\mathcal{S}}} W(x) \subset \mathcal{S} \\ W(y) - W(x) \leq -\gamma < 0, \quad x \subset y, \; x \neq y \in \tilde{\mathcal{S}}. \end{cases}$$

These conditions are always satisfied for α large enough. However, the difficulty D of the energy landscape, related to the performance of simulated annealing, tends to $+\infty$ when α tends to $+\infty$. Therefore, it is better to keep α as small as possible. In the next section, we will give an example for which we can take α arbitrarily small, and even $\alpha = 0$ if we are satisfied with $\gamma = 0$.

Equipped with this new energy function, we can proceed just as in the simpler case of section 3.

5.3. Some example of useful loss function

Often in task assignment problems, we would like some resources to be distributed according to some prescribed distribution. For instance in a time-table problem, we may want to schedule an equal number of hours in each week of the year.

This can be formalized in the following way. We consider first some function $\Phi : B \times E \longrightarrow F$ where F is a finite set (which may be the discretization of a domain in $\mathbb{R}^n$). Typically, Φ will be the projection on the time axis in a time-table problem. Then we consider a target distribution ρ defined on F. Let us consider some reference measure μ on $B \times E$ (such as the counting measure). To each partial solution $x \subset B \times E$, we may associate the restriction μ_x of μ to x, defined by

$$\mu_x(A) = \mu(x \cap A).$$

This induces a measure $\mu_x \circ \Phi^{-1}$ on F. The constraint we would like to represent by a loss function is that we would like $\mu_x \circ \Phi^{-1}$ to be approximately proportional to the reference measure ρ. This can be reflected in a loss function of the type

$$V(x) = \int h\left(\frac{\mu_x \circ \Phi^{-1}}{\rho}\right) d\rho,$$

where $h(x) = (1-x)^2$ or $h(x) = 1 - x + x\log x$. The function h is in both cases strictly convex, satisfies $h(1) = h'(1) = 0$ and h' is strictly increasing, therefore $\mu_x \circ \Phi^{-1} = \rho$ if and only if $V(x) = 0$ and the minimum of $V(x)$ on the set $\mu_x(B \times E) = constant$ is attained when $\mu_x \circ \Phi^{-1}$ is proportional to ρ, when this is feasible.

The following proposition holds:

Proposition 5. *Assume that the total weight $\mu_x(B \times E)$ of any solution is a function of the tasks to be scheduled only. This means that there is a measure $\tilde{\mu}$ on B such that*

$$\mu_x(B \times E) = \mu(x) = \tilde{\mu}(\pi_B(x)), \quad x \in \tilde{\mathcal{S}}.$$

Then for all global solutions $x \in \mathcal{S}$, $\mu_x(B \times E) = \mu(x) = \tilde{\mu}(B)$ is a constant.

Assume moreover that the measure ρ defining the constraint is such that $\rho(F) \geq \tilde{\mu}(B)$ and assume also that

$$\left\{ x \in \mathcal{S} \ : \ \frac{\mu_x \circ \Phi^{-1}}{\rho} \equiv constant \right\} \neq \emptyset.$$

Then

$$\arg \min_{s \in \tilde{\mathcal{S}}} V(x) = \left\{ s \in \mathcal{S} \ : \ \frac{\mu_x \circ \Phi^{-1}}{\rho} \equiv constant \right\},$$

meaning that the partial solutions minimizing V are exactly the global solutions x for which $\mu_x \circ \Phi^{-1}$ is proportional to the constraint ρ.

The assumptions of the proposition will be satisfied when $\mu_x(B \times E)$ measures the amount of assigned resources and the amount of resources to be allocated to a task depends only on the task and not on the way it is scheduled. Typically for instance, the number of hours of a course of teaching will be prescribed in advance and will not depend on the choice of a schedule for the lectures.

Now let us make the supplementary assumption that $\dfrac{\mu_x \circ \Phi^{-1}}{\rho} \leq 1$ for any $x \in \tilde{\mathcal{S}}$. We can always make this assumption true by increasing ρ by a suitable multiplicative factor (at least when ρ is strictly positive on F). In some cases we may on the contrary want to restrict $\tilde{\mathcal{S}}$ by adding the new constraint $\dfrac{\mu_x \circ \Phi^{-1}}{\rho} \leq$ 1. This will be done when the constraint has a practical meaning for the problem. For instance, if $\mu_x \circ \Phi^{-1}$ measures the number of lectures taking place in each hour of time in the week, we may want to fix ρ to a constant equal to the total number of available lecture rooms, add the constraint $\dfrac{\mu_x \circ \Phi^{-1}}{\rho} \leq 1$, to indicate that there is to be enough rooms to schedule all the lectures, and use the loss function $\int h\left(\dfrac{\mu_x \circ \Phi^{-1}}{\rho}\right) d\rho$ to indicate that we would like the rooms to be evenly occupied during the week (in a weekly time-table problem).

If the assumption $\dfrac{\mu_x \circ \Phi^{-1}}{\rho} \leq 1$, $x \in \tilde{\mathcal{S}}$ holds, then only the decreasing part of h is used, and the loss function V is always increasing during a destruction and decreasing during a construction. Therefore if γ is the constant corresponding to U in Eq. (2), we will have

$$W(y) \geq W(x) + \alpha \, \gamma, \quad x, y \in \tilde{\mathcal{S}}, \quad x \neq y, q_D(x, y) > 0.$$

6. The practical issue of the choice of parameters

In practical situations, the critical constants of the energy landscape are usually unknown. Therefore it is not possible to rely on the theoretical results we recalled in preceding sections to choose the parameters of algorithms. In the following paragraphs, we explain how we set the parameters in the experiments about jigsaw puzzles.

### 6.1. Simulated annealing

The cooling schedule can be written as

$$\frac{1}{T_n^N} = \beta_{\min} \left(\frac{\beta_{\max}}{\beta_{\min}}\right)^{n/N}.$$

We choose $\beta_{\min}$ and $\beta_{\max}$ by looking at the repartition function of the energies of the explored states in simulations at constant temperatures. We keep a value of $\beta_{\min}$ for which the slope of the repartition function stays large up to the largest values of the energies, meaning that states with high energies have a significant probability to be explored. For $\beta_{\max}$, we require on the contrary a repartition function concentrated on the lowest energy values.

The theory tells us that we can safely underestimate $\beta_{\min}$ and overestimate $\beta_{\max}$, which makes their choice possible from a qualitative inspection of repartition functions.

Here are two examples of repartition functions, corresponding to values of $\beta_{\min}$ and $\beta_{\max}$ which have been retained during the experiments.

rep h.t.,beta=0.0005

rep l.t.,beta=0.001
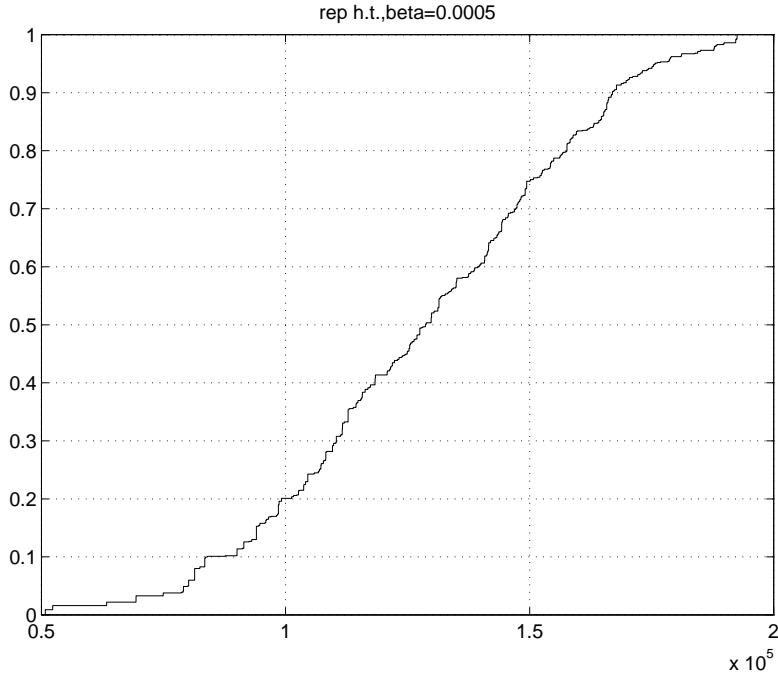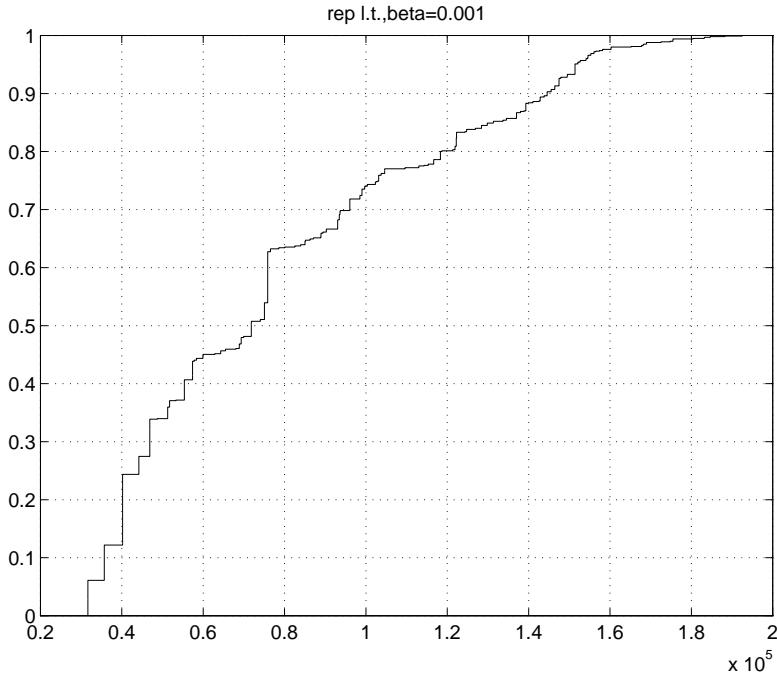
6.2. The iterated energy transformation method

In this case, the choice of parameters is may-be less straightforward. The analogy with simulated annealing can serve as a guideline: the high temperature regime corresponds to the case $\tau = \tau_1$ (i.e. to the first energy transform used). The low temperature regime corresponds to

$$\tau = (1 + \rho)\eta_0 - U_{\min}.$$

In order to test the behaviour of the algorithm in these two configurations, we make a short test using a small value of ρ ($\rho << 1$). The law of evolution of τ_k shows that for a small value of ρ, the algorithm will quickly switch from the high temperature regime during the first step, to a low temperature regime, during the following steps. In fact the value $\rho = 0$ may even sometimes be used. However when this is done, the algorithm sometimes encounters a state with a non defined energy transform too quickly, and there are not enough iterations to compute a reliable repartition function for the low temperature regime. This problem, when encountered, can be circumvented by using a low but non zero value of ρ.

We compute the repartition function of energies during the first step of the test run and during the last. The first function describes the equivalent of the "high temperature regime", and is tuned by the choice of the constant α and of the temperature parameter T, the second function corresponds to the "low temperature regime", and is tuned by a proper choice of $\tilde{\eta}_0 = \eta_0(1 + \rho)$.

Once these two choices are made, there remains a free parameter, namely ρ.

The theory [6] indicates an optimal choice of ρ of order $\sqrt{N}$ and an optimal choice of $r = N/M$ of order $\sqrt{N}\log(N)$. On the other hand, as soon as $\log(1 + \rho) > 1$, the convergence rate will be better than for simulated annealing. This indicates that a large value of ρ may safely be chosen, and that r can then be set to make

$$(U_{\min} - \gamma + \eta_0)\rho \left(\frac{\rho}{1 + \rho}\right)^{(r-1)}$$

small. This will ensure a small dependence of the final value of the shift τ_N with respect to its initial value $\tau_1 = \delta - \eta_0$.

6.3. Repeated optimizations

In this paragraph, we will consider that N iterations are to be divided into N/M trials of length M, and that we will keep the best solution found out of these N/M trials. In this context, the probability of failure in the worst case with respect to the starting point of each trial is $\epsilon_1(M)^{N/M}$, where

$$\epsilon_1(M) = \max_{x \in \tilde{\mathcal{S}}} P\left(U(X_N) > U_{\min} \mid X_0 = x\right) .$$

The first remark to be made (see Azencott [2]) is that for all the algorithms we have considered $\lim_{M \to +\infty} \frac{1}{M} \log \epsilon_1(M) = 0$. Therefore, when N is large enough, the optimal value for M is independent of N.

In this paragraph we will discuss the choice of the length M of each run of the algorithm. For simulated annealing, we can, on the basis of the theoretical bound on the probability of failure, namely $\left(\frac{A}{M^\alpha}\right)^{N/M}$ for N iterations divided into N/M runs of length M, conjecture that an overestimation of M will be relatively harmless, whereas an underestimation would be more penalizing. This can be seen on the derivative

$$\frac{\partial}{\partial M}\left(\frac{A}{M^\alpha}\right)^{N/M} = \left(\frac{A}{M^\alpha}\right)^{N/M} \frac{N\left(\alpha(\log M - 1) - \log A\right)}{M^2},$$

but is may be more vividly illustrated by a small numerical application. If we take for example $A = e^4$, $\alpha = 1$ and $N = 1000$, and if we put $\epsilon(M) = \left(\frac{A}{M^\alpha}\right)^{N/M}$, we see that

$$\min_M \epsilon(M) = \epsilon(148) \simeq 0.0012,$$

and that for this optimal value the probability of failure in each run is $\simeq 0.37$.

Here are some values taken by $\epsilon(M)$

| M	74	100	148	300	500	1000
$\epsilon(M)$	0.016	0.0024	0.0012	0.0034	0.012	0.055

and a graphic of this function:
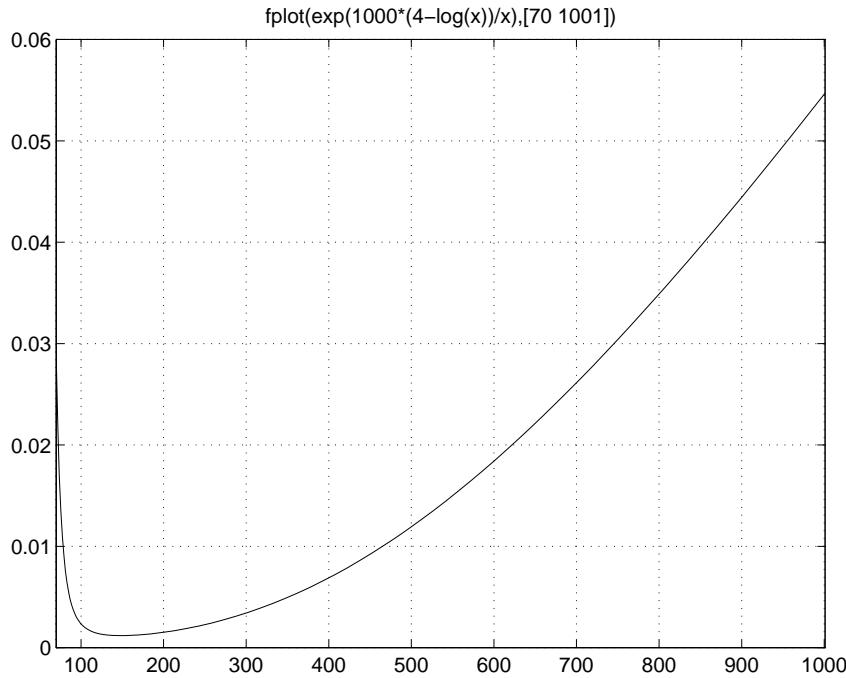
These figures show that, as far as this rough theoretical bound is a good guideline, there is a clear benefit in performing multiple runs instead of one long run, but that an overestimation of a factor two of the length of each run is relatively harmless. We remark also that a quite low confidence level for each run is favourable in this example where the difficulty is 1.

The same kind of reasoning would also hold for the theoretical bound of order $\epsilon(M) = \left(\dfrac{A}{M^{\alpha \log M}}\right)^{N/M}$ obtained for the iterated energy transformation method. In this case the derivative of the confidence level $\epsilon(M)$ is

$$\frac{\partial}{\partial M}\epsilon(M) = \epsilon(M)\frac{N\left(\alpha((\log M)^2 - 2\log M) - \log A\right)}{M^2}.$$

This is a plot of this function for some choice of the parameters α, N and A:

fplot(exp(5000*(2 − log(x)*log(x)/16)/x),[350 5001])
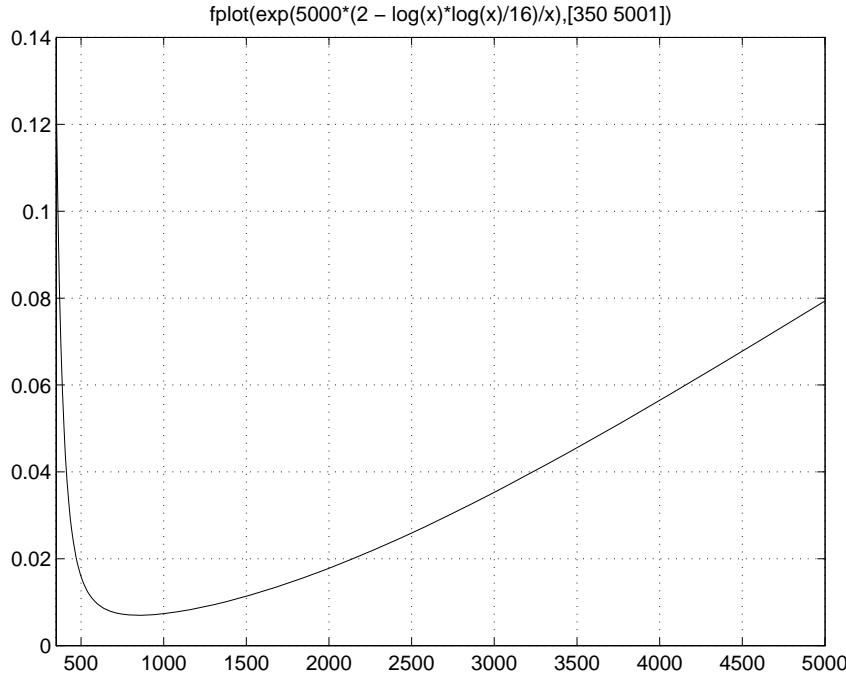
The tolerance with respect to an overestimation of M is even better than for simulated annealing.

For a comparison between repeated searches and interacting parallel searches, we refer to Graffigne [14] and to Azencott and Graffigne [3].

7. Experimental results

We tried to solve two kinds of puzzles: a small "tight" puzzle with 9 pieces and no loss function, and a big "loose" 60 pieces puzzle with a loss function. By "tight" we mean that there is just enough room in the frame to put all the pieces, and by "loose" we mean on the contrary that there is some extra-room left in the frame, the difficulty being then to minimize the loss function.

7.1. Small "tight" jigsaw puzzle

Our small jigsaw puzzle is a 9 pieces problem. The algorithm we used to solve it corresponds to the description given in section 4. The frame is a 40×50 grid. The size of the pieces are $(14, 27)$, $(8, 36)$, $(8, 9)$, $(6, 14)$, $(34, 5)$, $(18, 9)$, $(22, 9)$, $(18, 21)$, $(18, 15)$. The problem has several solutions, due to symmetry properties. This is one solution:

Jigsaw Puzzle
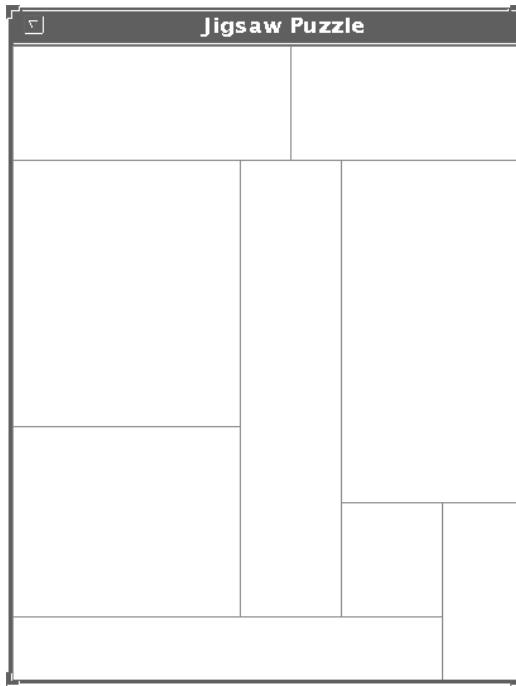
The parameters of the algorithms were set using the heuristics described in section 6.

We performed 40 runs of the simulated annealing algorithm and the same number of runs or the I.E.T. algorithm. For each algorithm we computed the repartition function of the energy of the best solution encountered during each run, and the repartition function of the energy of the final state of the algorithm. Of course, the former repartition function is always above the latter, therefore we can unambiguously plot them on the same diagram. In order to perform a "fair" comparison, we allowed the same number of iterations in both cases, namely $N = 5000$ iterations per run.

Performance of simulated annealing

rep of ener and ener_min
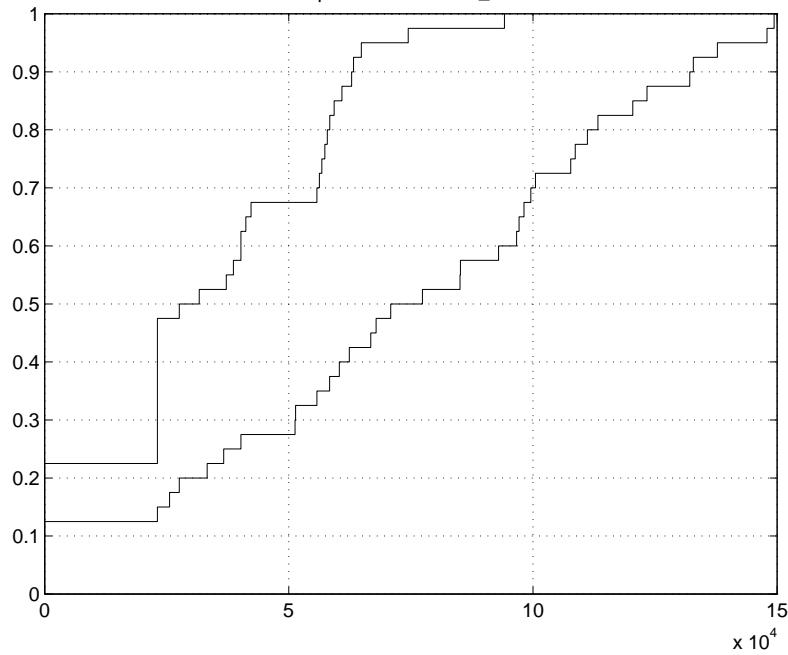
Performance of the iterated energy transformation method

rep of ener and ener_min
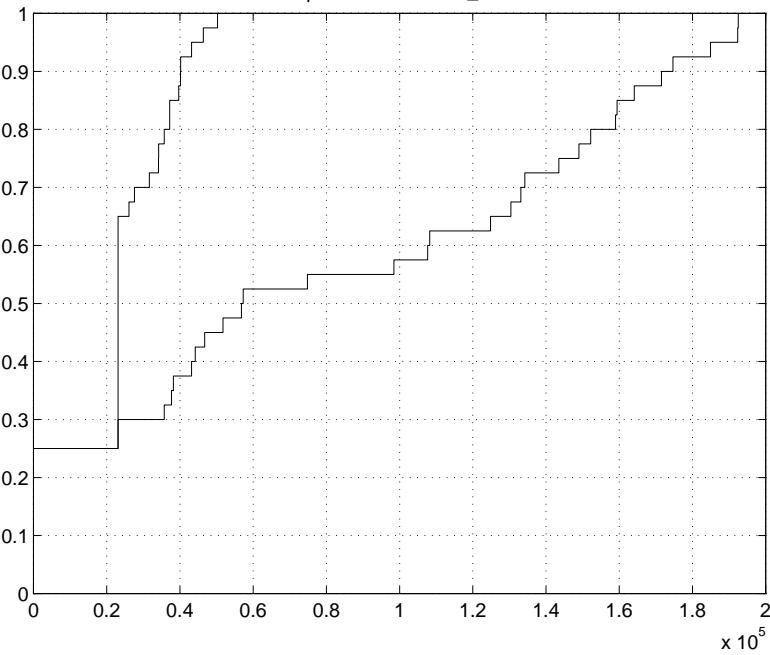
The results are of the same order, with some advantage in favour of the I.E.T. method. This is especially true when the energy of the final state is considered. An interpretation of this fact is that the I.E.T. algorithm is more efficient in preventing the process from leaving the global minimum once it has reached it.

We were also able to check the influence of the drift towards the upper left corner. In the two previous experiments, the maximum number of steps of the drift (the constant `max_drift` in the pseudo code of section 4) is 10. We have also tried a maximum number of steps of 50, for simulated annealing. We got the following improvement in the performance on 40 runs:

rep of ener and ener_min
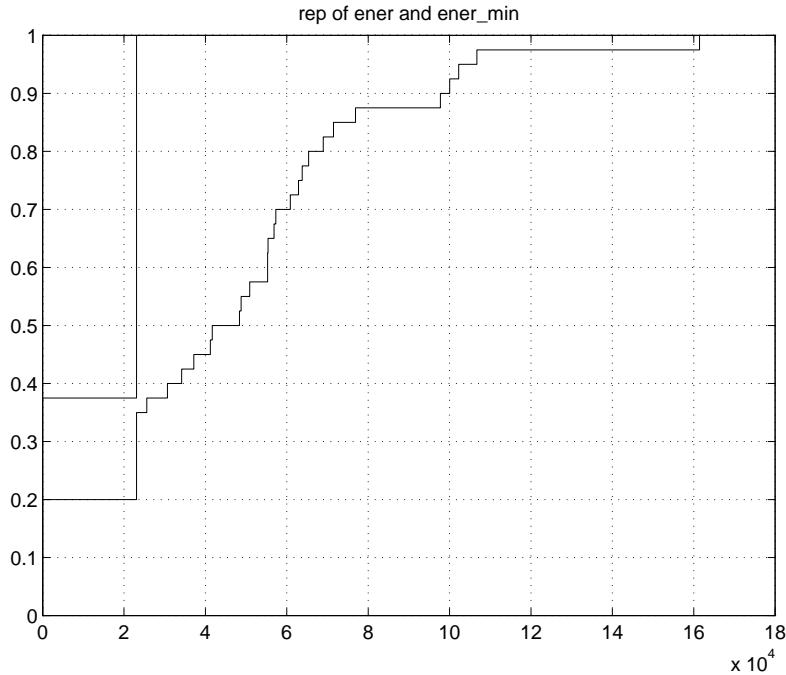
7.2. A big "loose" jigsaw puzzle

Our big jigsaw puzzle has 60 pieces, covering an area of 230 unit squares. The frame is a grid of size 30×10. The sizes of the pieces are the following:

number of pieces	width	height
15	3	2
15	2	1
5	5	2
5	2	4
20	1	1

The loss function is of the type described in the previous section. The function Φ here is the projection on the second axis, $\Phi((r, a, b)) = b$, $(r, a, b) \in B \times E$, so that the constraint indicates how much of each line the pieces should fill. On the following diagram, we have plotted the constraint function ρ.
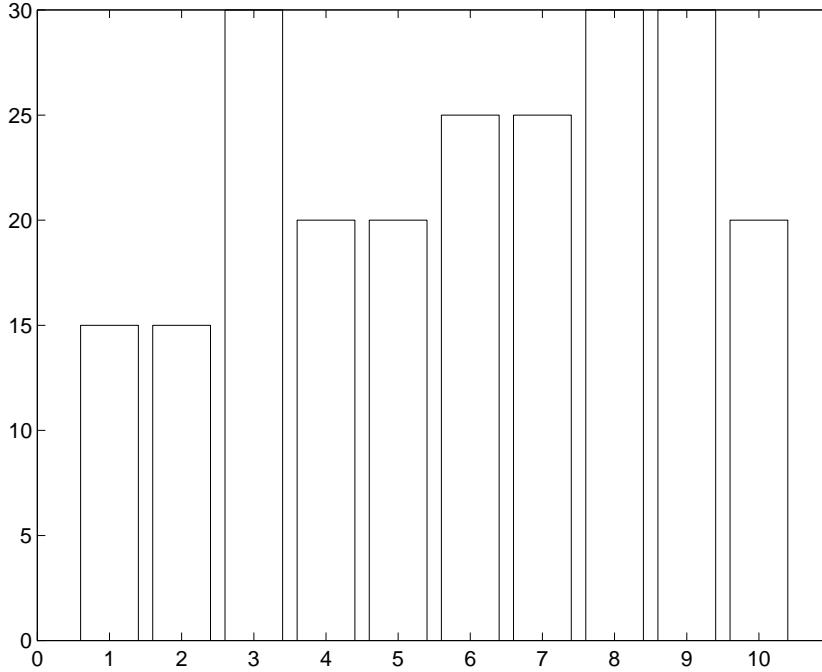
With this choice of ρ, the constraint is tight, meaning that $\rho(F)$ is equal to the area of the pieces. When we use tight constraints, we build problems of the partition type, which are therefore N.P. complete. We chose the size of the pieces such that the set of global solutions is not empty. However, for a 60 pieces problem, it is very difficult to find a (complete) solution.

We have chosen a coarse discretization step to keep the difficulty of the problem to a reasonable level, since we had to switch off the vertical drift. Indeed keeping a vertical drift would have decreased the stability of minimizing configurations in a non favourable way.

We tried two kinds of energies. In order to have a point of comparison, we tried to use the simple energy $U(x) = -\mu(x) + \max_{y \in \tilde{\mathcal{S}}} \mu(y)$, where μ is the counting measure.

Then we tried a compound energy $W(x) = U(x) + \alpha V(x)$ for a large value of α and for $V(x) = \int \left(1 - \dfrac{\mu_x \circ \Phi}{\rho} \right)^2 d\rho$.

Eventually, we tried to relax the constraint, changing ρ to $\tilde{\rho} = 6/5 \times \rho$.

7.2.1. Experiments with a simple energy function

In order to have a point of comparison, we recorded first the performance of repeated relaxations. The relaxation algorithm we used corresponds to a choice of $\lambda = 0$, or equivalently to a choice of $\beta = +\infty$ in the Metropolis algorithm.

Then we considered the Metropolis algorithm for different values of λ and of β. We tried $\lambda = 1$ and $\lambda = 0.5$, two "natural" choices for λ. The former let us

inhibit destructions only according to the energy increment, whereas the latter let constructions and destructions have equal frequencies at infinite temperature.

The first conclusion we reached was that a significant improvement over the relaxation scheme could be obtained, using the Metropolis algorithm with a moderate number of steps. We compared relaxation with 300 steps (for which convergence was always reached) with Metropolis with $\lambda = 1$, $\beta = 1$ and $N = 4000$. In order to compare methods using the same number of iterations, we repeated Metropolis 20 times and the relaxation algorithm $\left\lfloor \dfrac{4000 \times 20}{300} \right\rfloor = 266$ times. On the following diagram we plotted the repartition functions of the best solution found for each of the 20 runs of Metropolis (dashed lines) with the best 20 results out of the 266 runs of the relaxation algorithm (solid lines).

comp. of ener_min
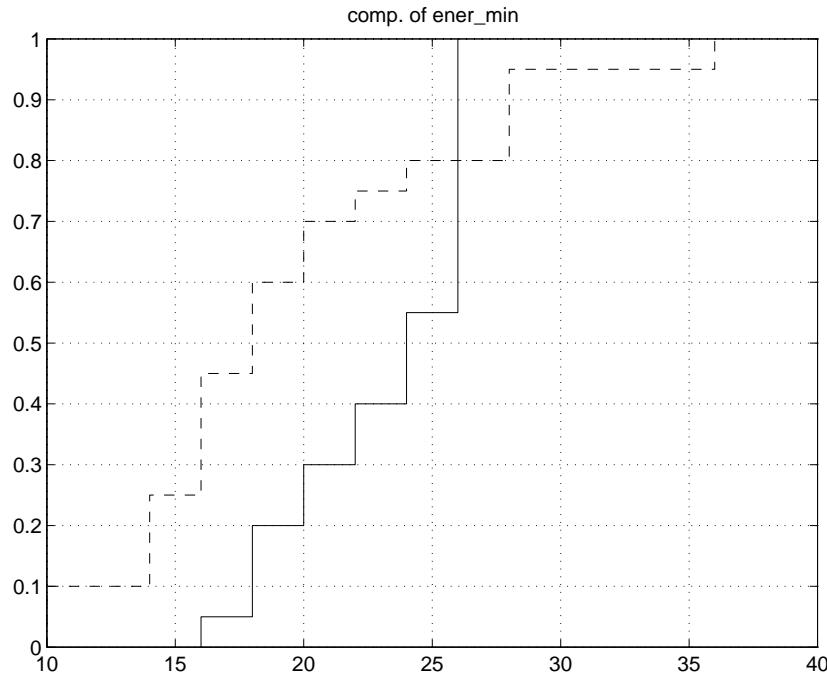
We obtained very suggestive evolutions for the Metropolis algorithm, such as the following:

Energy/time for the Metropolis algorithm
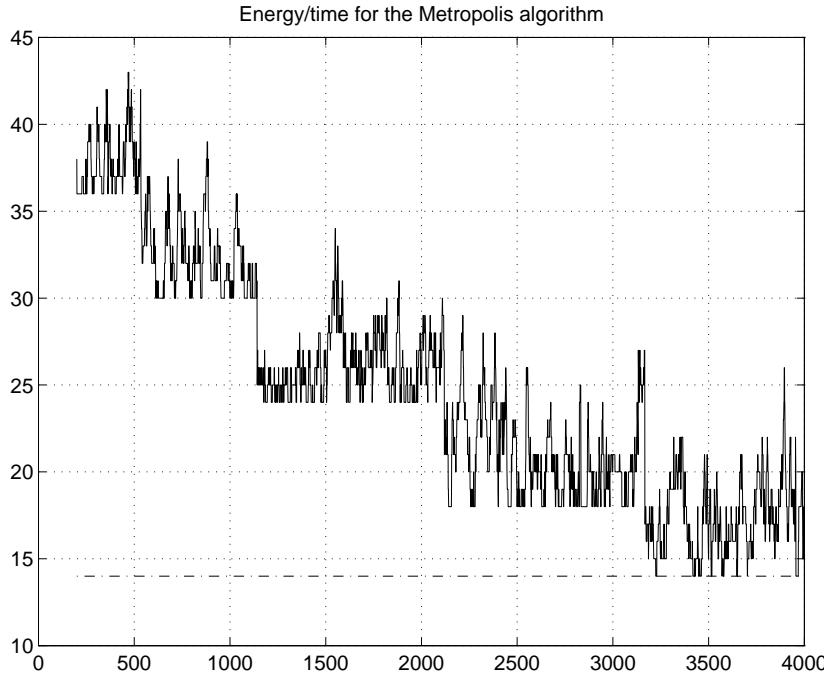
On this plot of $u_n = U(X_n)$ for $n = 300, \ldots, 4000$, we see the "staircase" shape of the trajectories of the Metropolis algorithm. The algorithm "falls" into deeper and deeper maximal cycles of the domain $\tilde{\mathcal{S}} \setminus \mathcal{S}$. We refer to [4] for a theoretical study of the exit path of the Metropolis algorithm from a domain at low temperature. For a study of the trajectories of simulated annealing algorithms, we refer to [5] and [20], which rely on more complex but also more general induction proofs which cover the time inhomogeneous case. For a semi-group approach of the same question in the continuous time case, we refer to [18] and [19].

The energy evolution can be decomposed into a decreasing part $\underline{u}_n = \min_{k \leq n} u_k$ and a "wandering" part $\overline{u}_n = u_n - \underline{u}_n$, as in the following diagram:

Decomposition of energy/time for the Metropolis algorithm
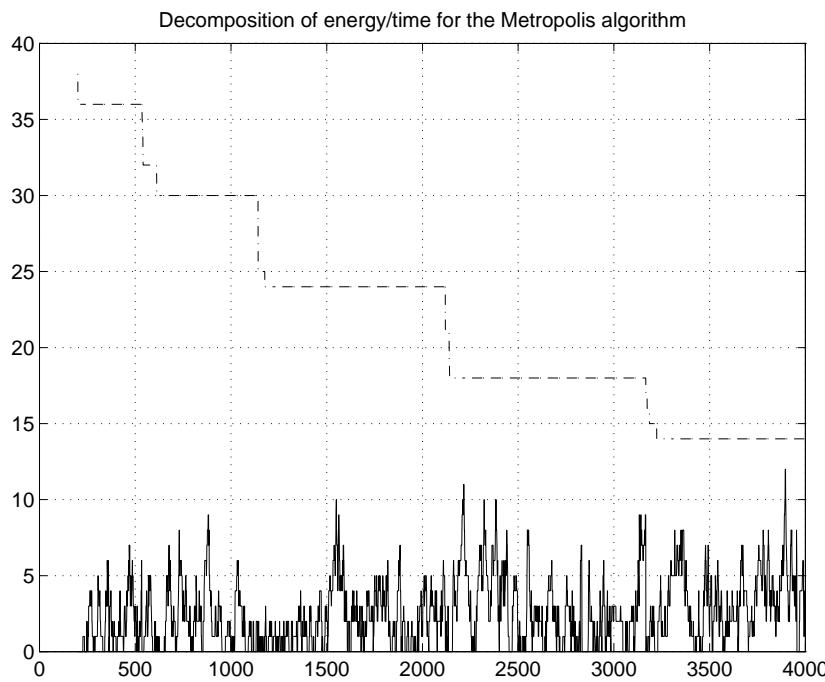
The repartition function of the wandering part gives informations about the depth of secondary attractors from which the algorithm is able to escape within the time of the simulation. It is a useful tool to choose the inverse temperature parameter β. This is the repartition function corresponding to the preceding plot:

Repartition of the stationary part (Met. alg.)
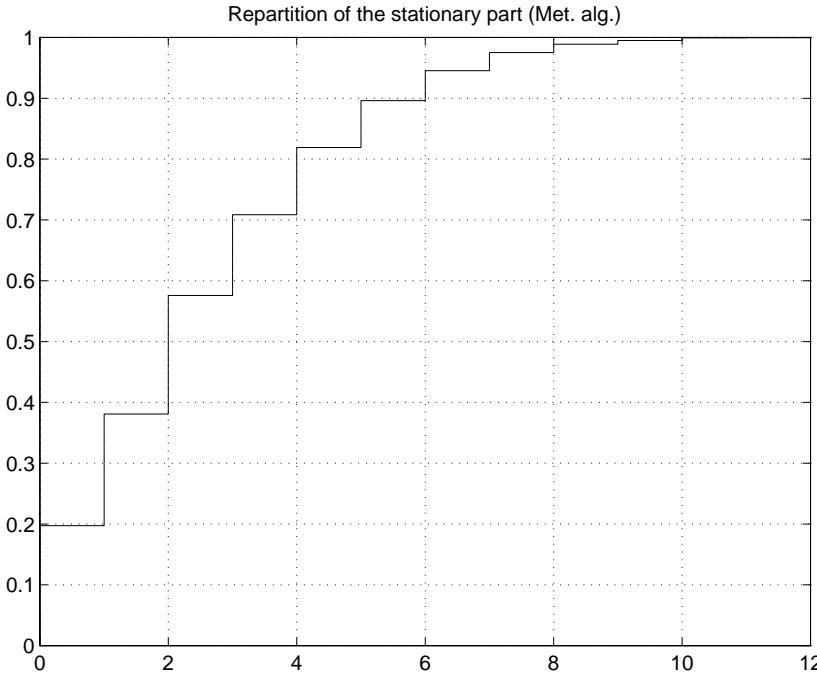
The best results for the Metropolis algorithm of time length $N = 4000$ were obtained for $\beta = 1$ and $\lambda = 1$ or for $\beta = 0.8$ and $\lambda = 0.5$. This shows that in this case, the choice of λ is not crucial. In the following, we will use $\lambda = 1$, because we can hope to take a better advantage of the discrimination made by the energy function between small and big pieces when we use this value of λ.

Then we used the Metropolis algorithm and simulated annealing on long time intervals. Namely we took $N = 20000$, $\beta_{\min} = 0.7$, $\beta_{\max} = 1.1$ for simulated annealing and $\beta = 1$ for the Metropolis algorithm. On ten runs of each algorithm, we could notice a clear gain in performance in favour of simulated annealing.

We tried eventually to get a better improvement using the I.E.T. algorithm. Since the state space is already rather large, we followed the idea introduced in [6] to use transformations $F_{\alpha,T,\tau}$ with a non zero value of α.

We took $\alpha = 0.3$, $\beta = 1/T = 30$, $\dfrac{1}{1+\rho} = 0.5$, $r = 4$ and $\eta_0 = 15$. We obtained the following comparative results for the best energy found in each of 10 runs of each algorithm. The mean values are 16.2 for the Metropolis algorithm (solid lines), 11.6 for simulated annealing (dashed lines) and 10.9 for the I.E.T. algorithm (dash-dot lines). The repartition functions are plotted on the next diagram:

comp. of ener_min
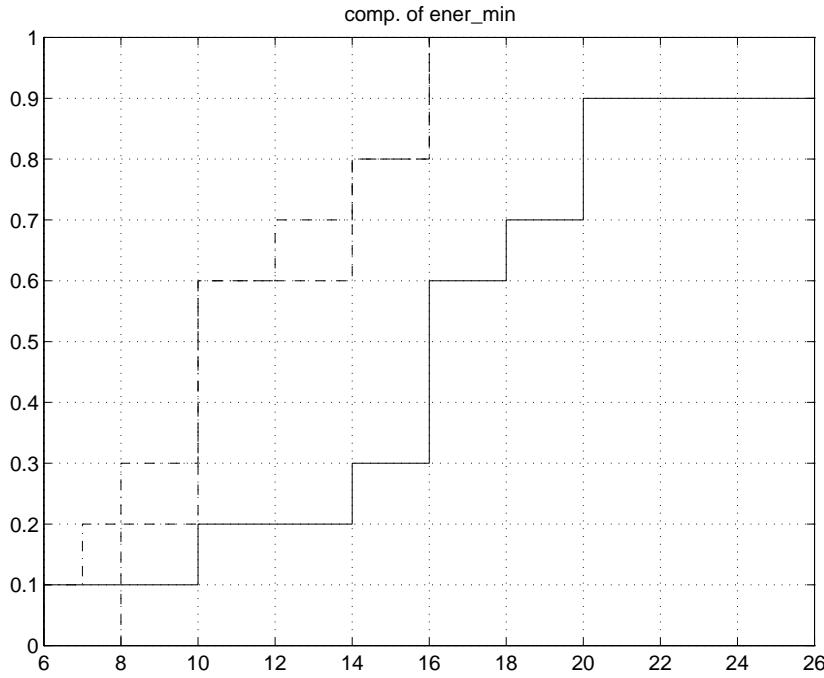
7.2.2. Experiments with a compound energy function

We used the energy

$$W(x) = U(x) + \alpha V(x),$$

with a huge value of $\alpha = 10000$.

The range of this energy is very large, when compared with the previous one, since $W_{\max} = 2300230$, whereas $W_{\min} = 0$ and removing a piece of size 1×1 from a complete solution in a line of weight $\rho(y) = 30$ costs $\Delta W \simeq 334.33$. Therefore we may expect more spectacular improvements from the speed-up techniques.

We tried different temperatures for the Metropolis algorithm with $N = 20000$. The best results were obtained when $\beta = 8 \times 10^{-4}$. On 10 runs, the average best value was 15853.

Using simulated annealing with $\beta_{\min} = 10^{-4}$, $\beta_{\max} = 10^{-3}$, we improved the performance on the average, as shown in the next diagram. On 10 runs, the average best energy value was 8765.

We obtained some more improvement using the I.E.T. algorithm (with $\gamma = 5 \times 10^{-5}$, $\beta = 10$ and $\eta_0 = 2000$). On 10 runs the average best energy value was 6280.

This is a diagram of the repartition functions of the best energy value for ten runs of the Metropolis algorithm (solid lines), simulated annealing (dashed lines), and the I.E.T. algorithm (dash-dot lines).

comp. of ener_min
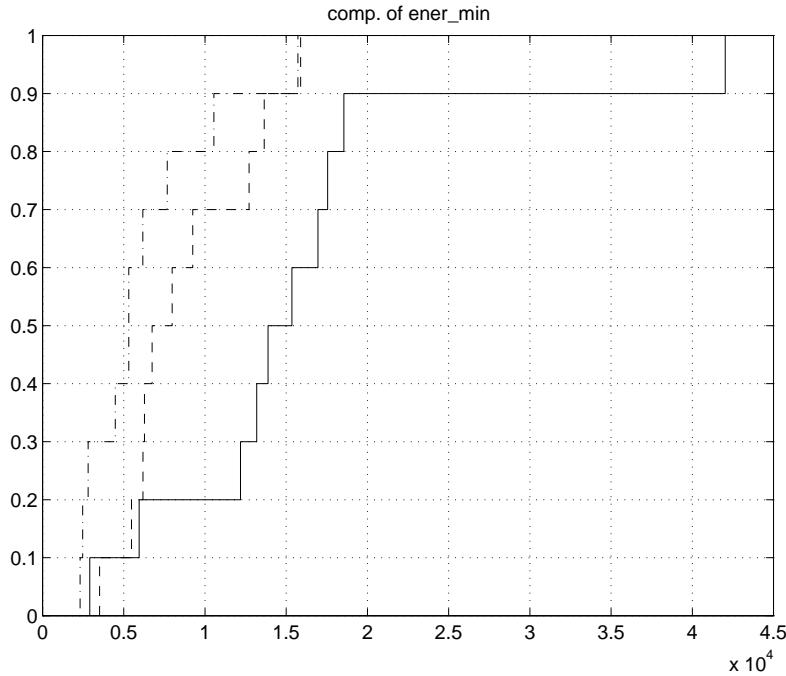
7.2.3. Experiments with a relaxed constraint

We explored also an alternative in the optimization design, which consists in replacing ρ by $\tilde{\rho} = \frac{6}{5}\rho$. We considered accordingly a larger search space $\tilde{\mathcal{S}}$ where the constraint $\frac{\mu_x \circ \Phi^{-1}}{\rho} \leq 1$ is relaxed to $\frac{\mu_x \circ \Phi^{-1}}{\tilde{\rho}} \leq 1$. We took again a compound energy of the type $W(x) = U(x) + \alpha\, V(x)$, with $\alpha = 10000$. The range of W is between $W_{\min} = 76666.66$ and $W_{\max} = 2760230$.

On this example, we can perform the same kind of comparison as in the case of tight constraints. We made 10 runs of length $N = 20000$ of each algorithm. The average of the best energy value found in each run is 8731 for the Metropolis algorithm, 8685 for simulated annealing, and 8567 for the I.E.T. algorithm. This is a diagram of the corresponding repartition functions (solid lines for the Metropolis algorithm, dashed lines for simulated annealing and dash-dot lines for the I.E.T. algorithm).

comp. of ener_min
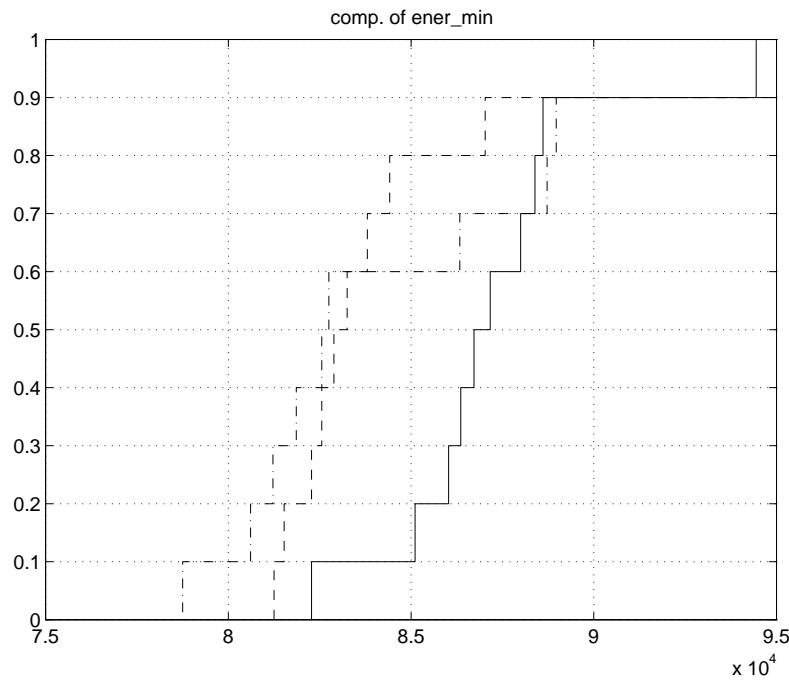
The best solution was found by the I.E.T. algorithm. It has an energy of $W(x) = 78750$ and is shown below.


Jigsaw Puzzle
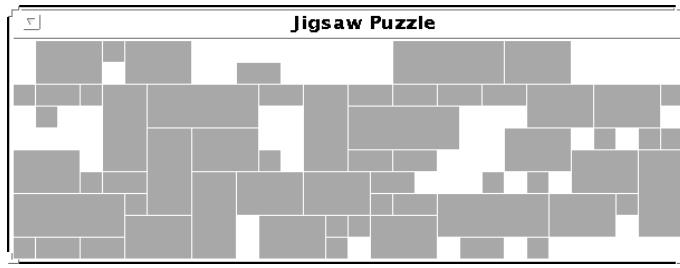
In this solution, all the pieces are set in the frame. We can judge of the quality of the solution with respect to the proportionality constraint on the following diagram, where we have plotted $\tilde{\rho}$ (dashed lines), the measure expressing the constraint, and $\mu_x \circ \Phi^{-1}$ (solid lines), giving the number of unit squares actually filled on each line by the solution. The optimum would be $\mu_x \circ \Phi^{-1} = \rho = 5/6 \times \tilde{\rho}$. We are not too far from that: the two entries $\mu_x \circ \Phi^{-1}(2)$ and $\mu_x \circ \Phi^{-1}(4)$ are one unit too large, and $\mu_x \circ \Phi^{-1}(9)$ is two units short from the optimum. This is the best approximation to an optimal solution we were able to compute on this example. This seems to show that relaxing slightly the constraint and introducing the loss function $V(x)$ in the energy eases the optimization process.
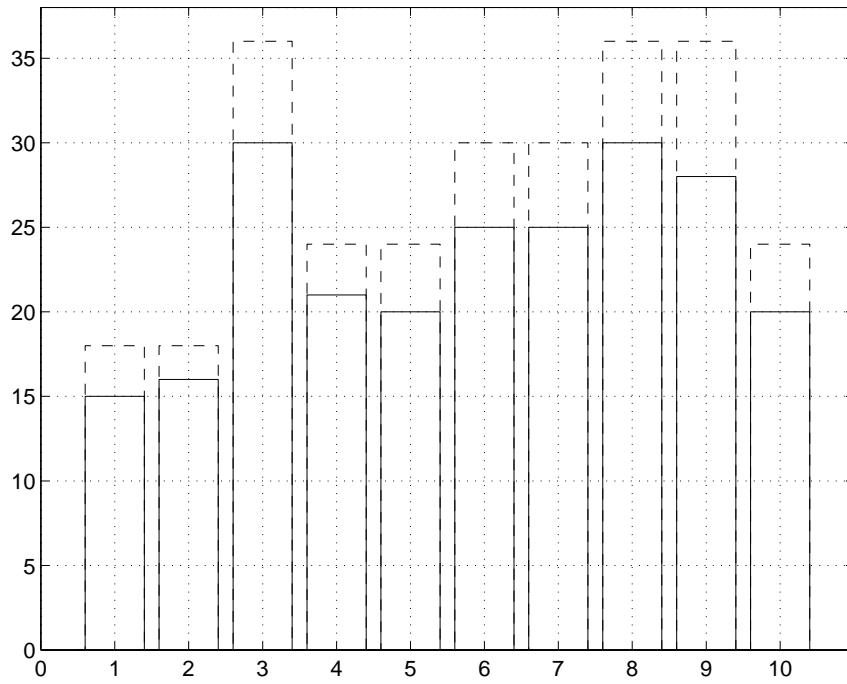
This should be compared with the best solution found without relaxing the constraint:
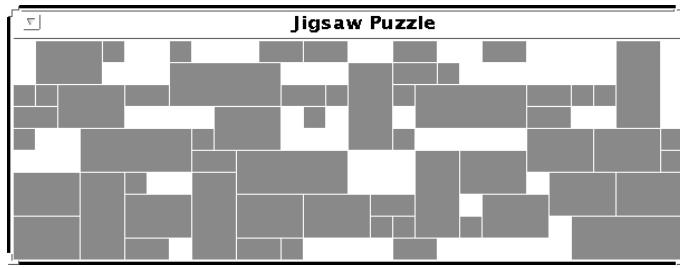
and its constraint diagram:
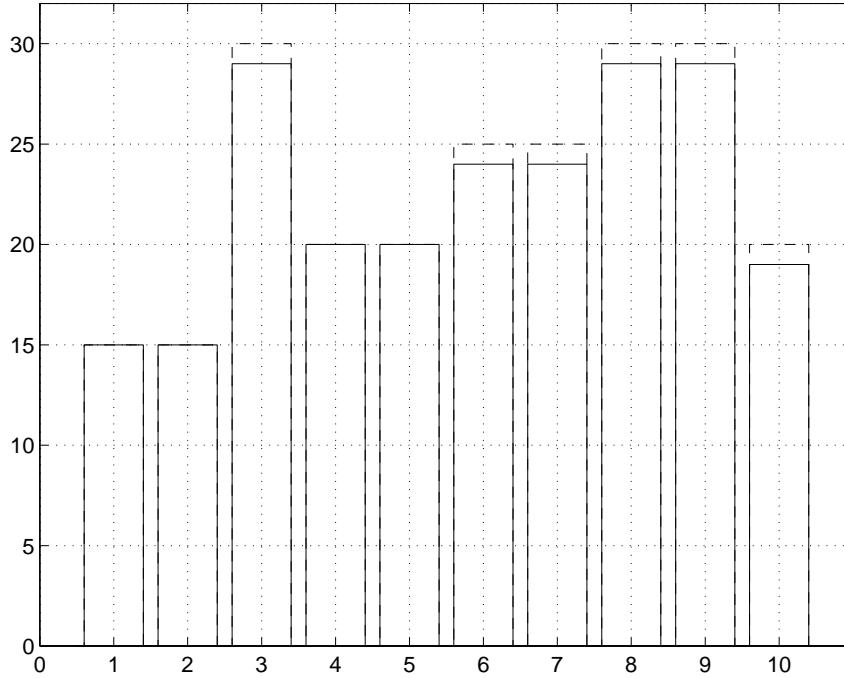
For this solution $U(x) = 6$. Solutions of energy $U(x) = 6$ were also found using the simple energy U to guide the search. Therefore the advantage of introducing the V component in the energy function is not obvious when the constraint is really tight.

It is also interesting to consider typical energy evolutions of those three algorithms. On the following diagrams, we have plotted the sequence

$$u_n = U(X_n)$$

As we have already mentioned, these sequences of energy values can be decomposed into a decreasing component

$$\underline{u}_n = \min\{u_k \ : \ k \le n\},$$

and a wandering component

$$\overline{u}_n = u_n - \underline{u}_n.$$

The repartition functions of $(\overline{u}_n, n = 1, \ldots, N)$ can help to set properly the parameters. It indicates the depth of the attractors from which the algorithm is able to escape.

Energy/time for the Metropolis algorithm

Energy/time for Simulated Annealing
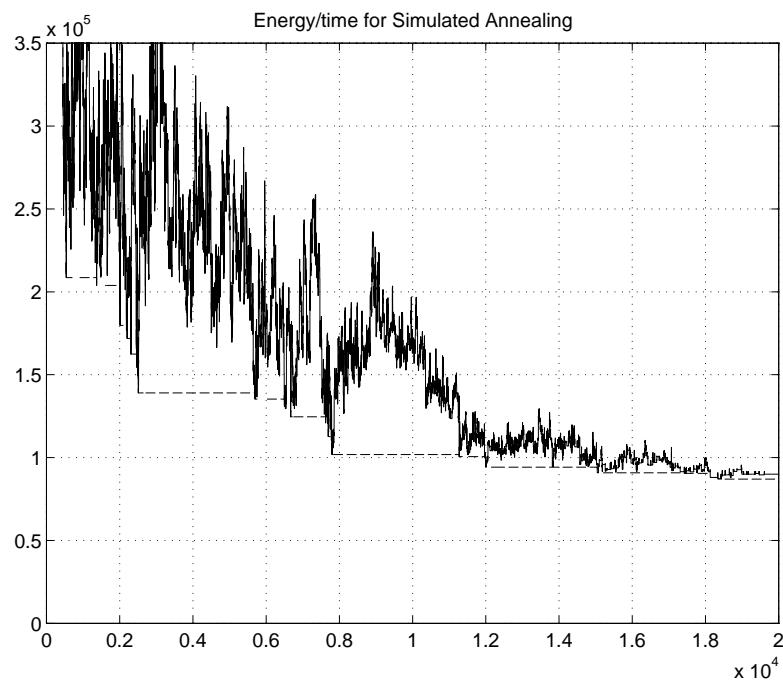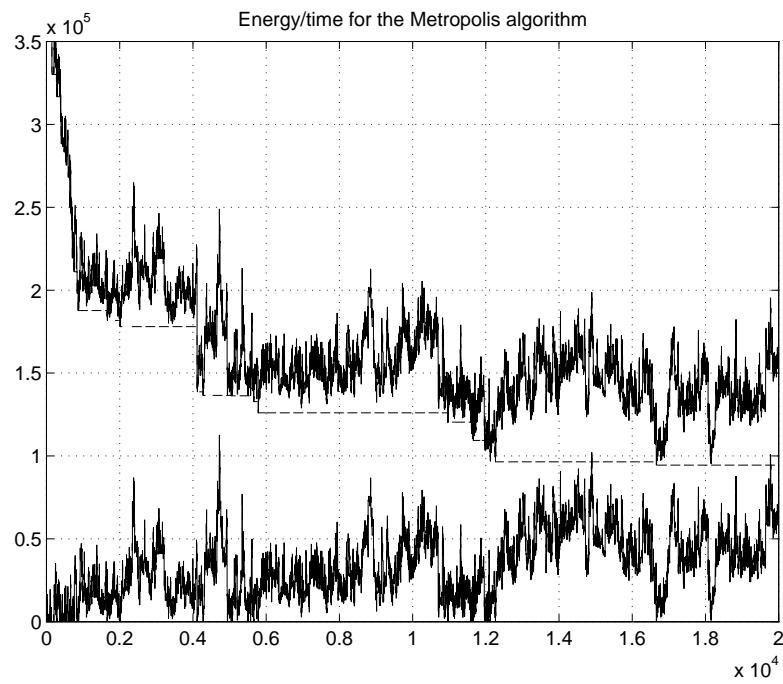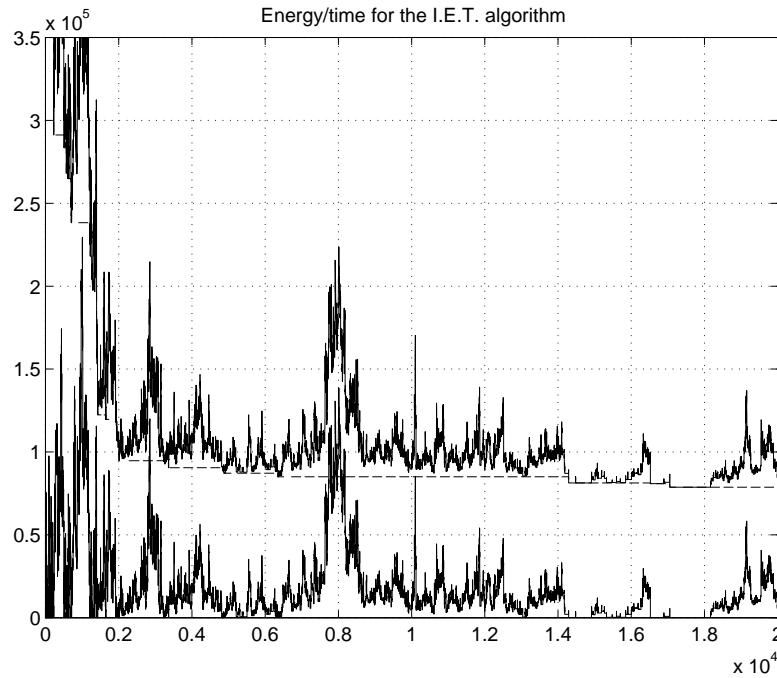
It is interesting to compare the energy evolutions of the three algorithms. The comparison between the Metropolis algorithm and simulated annealing shows clearly that the temperature used in Metropolis is too low during the 4000 first iterations and too high during the 8000 last iterations. As for the I.E.T. algorithm, we can see that the fluctuations of the wandering part are decreasing with time, as in the case of simulated annealing, but that the evolution of the energy is more unstable: it can go up and down faster (in other words its peaks are sharper). This explains why it is able to sample more efficiently a state space containing many local minima.

Conclusion

We touched in this paper three related topics with various degrees of generality. The first one was to bring experimental evidences comforting theoretical results about the behaviour of algorithms. We wanted to show that the theory was not concerned with "never reached" asymptotic and led to the same qualitative ranking of performances as an experimental benchmark would do. Our second aim was to describe a general purpose methodology to deal with scheduling tasks. We insisted on two problems which are likely to be encountered in many situations: the creation of small gaps in the allocation of resources and the way to handle "proportionality constraints".

The third aspect of the paper is to account for experiments on a benchmark of the "jigsaw puzzle" type. This was the occasion to be confronted with the practical problem of the choice of parameters and of optimization design options (such as relaxing some of the constraints). Our conclusion on this third

point is that we have acquired some know-how about the choice of parameters, which we tried to reflect in section 6, but that we have presently no systematic rule to choose them. We worked very much in a trial and error way, looking at the repartition functions we mentioned to guide our intuition. A trial and error procedure is somehow justified by the theoretical result that many trials of moderate length are preferable to a long one. This gives the opportunity to tune the parameters trial after trial.

Anyhow we have to admit that the choice of parameters requires some skill, especially for simulated annealing and the I.E.T. algorithms, where there are more than one parameter to tune. What we did not find too hard to do was, starting from a given Metropolis algorithm at inverse temperature β_{Met}, to find $\beta_{\min} < \beta_{\mathrm{Met}} < \beta_{\max}$ for which simulated annealing performs better than Metropolis. Then we could get some more improvement using the I.E.T. algorithm, where again we chose the parameters in relation with those used for simulated annealing. We are not sure at all that this is the best way to tune simulated annealing or the I.E.T. algorithm, but it shows at least that the theoretical gains of one algorithm upon the previous one could be obtained in practice.

Another positive result of these experiments is that it is possible to get good, if not optimal, solutions even in the case when very non-monotonous evolutions of the energy are needed, as it is the case here, since the only way to move a piece of the puzzle is to remove it and put it somewhere else afterwards, a succession of two moves the first of which implies an energy increase.

Of course we have touched in this paper only a limited number of questions. For instance, we let open the practical question of the best choice of parameters for simulated annealing and for the I.E.T. algorithm, since we used only robust "all purposes" set of parameters, namely exponential temperature sequences in the case of simulated annealing and logarithmic energy transforms for the I.E.T. algorithm. Another question we let purposely in the dark is the choice of elementary moves. Although it is clear that a benefit can be obtained from the use of more complex compound moves, we felt such an investigation would have been too much dependent on the precise examples we chose to study. We rather tried to lay the stress on general ideas and tools, with the hope that they could be useful in a variety of situations.

Remark about the bibliography: [21] and [22] are partial English translations of [20].

References

1. Azencott Robert (1988) Simulated Annealing, *Séminaire Bourbaki 40ième année, 1987-1988* **697**.

2. Azencott Robert (1992) Sequential Simulated Annealing: Speed of Convergence and Acceleration Techniques, in *Simulated Annealing: Parallelization Techniques*, R. Azencott Ed., chap. 1, Wiley Interscience.

3. Azencott Robert and Graffigne Christine (1992) Parallel Annealing by Periodically Interacting Multiple Searches: Acceleration Rates, in *Simulated Annealing: Parallelization Techniques*, R. Azencott Ed., chap. 1, Wiley Interscience.

4. Catoni Olivier and Cerf Raphael (1995) The Exit Path of a Markov Chain with Rare Transitions preprint, avaible at http://www.dmi.ens.fr, to appear in ESAIM: P&S, http://www.emath.fr/Maths/Ps/ps.html.

5. Catoni Olivier (1992) Rough Large Deviation Estimates for Simulated Annealing: Application to Exponential Schedules, *The Annals of Probability*, Vol. 20, nb. 3, pp. 1109 - 1146.

6. Catoni Olivier (1994) The Energy Transformation Method fo the Metropolis Algorithm Compared with Simulated Annealing, *preprint, rapport de recherche du LMENS, to appear in Prob. Th. and Rel. Fields*, Nov. 1994,

7. Catoni Olivier (1991) Exponential Triangular Cooling Schedules for Simulated Annealing Algorithms: a case study, *Applied Stochastic Analysis, Proceedings of a US-French Workshop, Rutgers University, April 29 - May 2, 1991*, Karatzas I. and Ocone D. eds., Lecture Notes in Control and Information Sciences No 177, Springer Verlag, 1992.

8. Cot Cécile and Catoni Olivier (1996) Piecewise Constant Triangular Cooling Schedules for Generalized Simulated Annealing Algorithms *preprint, research report of the LMENS, available at http://www.dmi.ens.fr.*

9. Deuschel J.D. and Mazza C. (1994) L^2 convergence of time nonhomogeneous Markov processes: I. Spectral Estimates, *The annals of Applied Probability*, vol. 4, no. 4, 1012-1056.

10. Diaconis Persi and Stroock Daniel (1991) Geometric Bounds for Eigenvalues of Markov Chains, *The Annals of Applied Probability*, Vol. 1, No 1, 36 - 61.

11. Freidlin, M. I. and Wentzell, A. D. (1984). *Random Perturbations of Dynamical Systems.* Springer, New York.

12. Garey M. R. and Johnson D. S. (1979) *Computers and Intractability: a guide to the theory of NP - completeness*, Freeman, New-York.

13. S. Geman, D. Geman, *Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images*, I.E.E.E. Transactions on Pattern Analysis and Machine Intelligence, 6, 721- 741, 1984.

14. Graffigne Christine (1992) Parallel Annealing by Periodically Interacting Multiple Searches: An Experimental Study, in *Simulated Annealing: Parallelization Techniques*, R. Azencott Ed., chap. 1, Wiley Interscience.

15. Holley R. and Stroock D. (1988) Annealing via Sobolev inequalities, *Comm. Math. Phys.*, 115:553-559.

16. S. Kirkpatrick, C. D. Gelatt and M. P. Vecchi, *Optimization by simulated annealing*, Science, 220, 621-680, 1983.

17. Miclo L. (1991) Evolution de l'énergie libre. Application à l'étude de la convergence des algorithmes du recuit simulé. *Doctoral Dissertation*, Université d'Orsay, February 1991.

18. Miclo Laurent (1995) Sur les problèmes de sortie discrets inhomogènes *The Annals of Applied Probability, to appear.*

19. Miclo Laurent (1995) Sur les temps d'occupations des processus de Markov finis inhomogènes à basse température, *submitted to Stochastics and Stochastics Reports.*

20. Trouvé Alain (1993) Parallélisation massive du recuit simulé, *Doctoral Dissertation*, Université Paris 11, January 5 1993.

21. Trouvé Alain (1994) Cycle Decomposition and Simulated Annealing, *S.I.A.M. J. Control Optim., 34(3)*, 1996.

22. Trouvé Alain (1995) Rough Large Deviation Estimates for the Optimal Convergence Speed Exponent of Generalized Simulated Annealing Algorithms, *to appear in Ann. Inst. H. Poincaré, Probab. Statist.*, 32(2), 1996.