# Processor Allocation for Optimistic Parallelization of Irregular Programs* 

Francesco Versaci ${ }^{\dagger}$<br>TU Wien \& University of Padova

Keshav Pingali ${ }^{\ddagger}$<br>University of Texas at Austin


#### Abstract

Optimistic parallelization is a promising approach for the parallelization of irregular algorithms: potentially interfering tasks are launched dynamically, and the runtime system detects conflicts between concurrent activities, aborting and rolling back conflicting tasks. However, parallelism in irregular algorithms is very complex. In a regular algorithm like dense matrix multiplication, the amount of parallelism can usually be expressed as a function of the problem size, so it is reasonably straightforward to determine how many processors should be allocated to execute a regular algorithm of a certain size (this is called the processor allocation problem). In contrast, parallelism in irregular algorithms can be a function of input parameters, and the amount of parallelism can vary dramatically during the execution of the irregular algorithm. Therefore, the processor allocation problem for irregular algorithms is very difficult.

In this paper, we describe the first systematic strategy for addressing this problem. Our approach is based on a construct called the conflict graph, which (i) provides insight into the amount of parallelism that can be extracted from an irregular algorithm, and (ii) can be used to address the processor allocation problem for irregular algorithms. We show that this problem is related to a generalization of the unfriendly seating problem and, by extending Turán's theorem, we obtain a worst-case class of problems for optimistic parallelization, which we use to derive a lower bound on the exploitable parallelism. Finally, using some theoretically derived properties and some experimental facts, we design a quick and stable control strategy for solving the processor allocation problem heuristically.


Keywords: Irregular algorithms, Optimistic parallelization, Automatic parallelization, Amorphous data-parallelism, Processor allocation, Unfriendly seating, Turán's theorem.

[^0]
## 1 Introduction

The advent of on-chip multiprocessors has made parallel programming a mainstream concern. Unfortunately writing correct and efficient parallel programs is a challenging task for the average programmer. Hence, in recent years, many projects [14, 10, 3, 20] have tried to automate parallel programming for some classes of algorithms. Most of them focus on regular algorithms such as Fourier transforms [9, 19] and dense linear algebra routines [4]. Automation is more difficult when the algorithms are irregular and use pointer-based data structures such as graphs and sets. One promising approach is based on the concept of amorphous data parallelism [17]. Algorithms are formulated as iterative computations on work-sets, and each iteration is identified as a quantum of work (task) that can potentially be executed in parallel with other iterations. The Galois project [18] has shown that algorithms formulated in this way can be parallelized automatically using optimistic parallelization): iterations are executed speculatively in parallel and, when an iteration conflicts with concurrently executing iterations, it is rolled-back. Algorithms that have been successfully parallelized in this manner include Survey propagation [5], Boruvka's algorithm [6], Delauney triangulation and refinement [12], and Agglomerative clustering [21].

In a regular algorithm like dense matrix multiplication, the amount of parallelism can usually be expressed as a function of the problem size, so it is reasonably straightforward to determine how many processors should be allocated to execute a regular algorithm of a certain size (this is called the processor allocation problem). In contrast, parallelism in irregular algorithms can be a function of input parameters, and the amount of parallelism can vary dramatically during the execution of the irregular algorithm [16]. Therefore, the processor allocation problem for irregular algorithms is very difficult. Optimistic parallelization complicates this problem even more: if there are too many processors and too little parallel work, not only might some processors be idle but speculative conflicts may actually retard the progress of even those processors that have useful work to do, increasing both program execution time and power consumption. This paper presents the first systematic approach to addressing the processor allocation problem for irregular algorithms under optimistic parallelization, and it makes the following contributions.

- We develop a simple graph-theoretic model for optimistic parallelization and use it to formulate processor allocation as an optimization problem that balances parallelism exploitation with minimizing speculative conflicts (Section 2).
- We identify a worst-case class of problems for optimistic parallelization; to this purpose, we develop an extension of Turán's theorem [2] (Section 3).
- Using these ideas, we develop an adaptive controller that dynamically solves the processor allocation problem for amorphous data-parallel programs, providing rapid response to changes in the amount of amorphous data-parallelism (Section 4).


## 2 Modeling Optimistic Parallelization

A typical example of an algorithm that exhibits amorphous data-parallelism is Dalauney mesh refinement, summarized as follows. A triangulation of some planar region is given, containing some "bad" triangles (according to some quality criterion). To remove them, each bad triangle is selected (in any arbitrary order), and this triangle, together with triangles that lie in its cavity, are replaced with new triangles. The retriangulation can produce new bad triangles, but this process can be proved to halt after a finite number of steps. Two bad triangles can be processed in parallel, given that their cavities do not overlap.

[^1]
(i)

(ii)

(iii)

Figure 1: Optimistic parallelization. (i) Nodes represent possible computations, edges conflicts between them. (ii) $m$ nodes are chosen at random and run concurrently. (iii) At runtime the conflicts are detected, some nodes abort and their execution is rolled back, leaving a maximal independent set in the subgraph induced by the initial nodes choice.

There are also algorithms, which exhibit amorphous data-parallelism, for which the order of execution of the parallel tasks cannot be arbitrary, but must satisfy some constraints (e.g., in discrete event simulations the events must commit chronologically). We will not treat this class of problems in this work, but we will focus only on unordered algorithms [16]. A different context in which there is no roll-back and tasks do not conflict, but obey some precedence relations, is treated in [1].

Optimistic parallelization deals with amorphous data-parallelism by maintaining a workset of the tasks to be executed. At each temporal step some tasks are selected and speculatively launched in parallel. If, at runtime, two processes modify the same data a conflict is detected and one of the two has to abort and roll-back its execution. Neglecting the details of the various amorphous data-parallel algorithms, we can model their common behavior at a higher level with a simple graph-theoretic model: we can think a scheduler as working on a dynamic graph $G_{t}=\left(V_{t}, E_{t}\right)$, where the nodes represent computations we want to do, but we have no initial knowledge of the edges, which represent conflicts between computations (see Fig. 11. At time step $t$ the system picks uniformly at random $m_{t}$ nodes (the active nodes) and tries to process them concurrently. When it processes a node it figures out if it has some connections with other executed nodes and, if a neighbor node happens to have been processed before it, aborts, otherwise the node is considered processed, is removed from the graph and some operations may be performed in the neighborhood, such as adding new nodes with edges or altering the neighbors. The time taken to process conflicting and non-conflicting nodes is assumed to be the same, as it happens, e.g., for Dalauney mesh refinement.

### 2.1 Control Optimization Goal

When we run an optimistic parallelization we have two contrasting goals: we both want to maximize the work done, achieving high parallelism, but at the same time we want to minimize the conflicts, hence obtaining a good use of the processors time. (Furthermore, for some algorithms the roll-back work can be quite resource-consuming.) These two goals are not compatible, in fact if we naïvely try to minimize the total execution time the system is forced to use always all the available processors, whereas if we try to minimize the time wasted from aborted processes the system uses only one processor. Therefore in the following we choose a trade-off goal and cast it in our graph-theoretic model.

Let $G=(V, E)$ be a computations/conflicts (CC) graph with $n=|V|$ nodes. When a scheduler chooses, uniformly at random, $m$ nodes to be run, the ordered set $\pi_{m}(\cdot)$ by which they commit can be modeled as a random permutation: if $i<j$ then $\pi_{m}(i)$ commits before $\pi_{m}(j)$ (if there is a conflict between $\pi_{m}(i)$ and $\pi_{m}(j)$ then $\pi_{m}(i)$ commits and $\pi_{m}(j)$
aborts, if $\pi_{m}(i)$ aborted due to conflicts with previous processes $\pi_{m}(j)$ can commit, if not conflicting with other committed processes). Let $k_{t}\left(\pi_{m}\right)$ be the number of aborted processes due to conflicts and $r_{t}\left(\pi_{m}\right) \in[0,1)$ the ratio of conflicting processors observed at time $t$ (i.e. $\left.r_{t}\left(\pi_{m}\right) \triangleq k_{t}\left(\pi_{m}\right) / m\right)$. We define the conflict ratio $\bar{r}_{t}(m)$ to be the expected $r$ that we obtain when the system is run with $m$ processors:

$$
\begin{equation*}
\bar{r}_{t}(m) \triangleq \mathbb{E}_{\pi_{m}}\left[r_{t}\left(\pi_{m}\right)\right] \tag{1}
\end{equation*}
$$

where the expectation is computed uniformly over the possible prefixes of length $m$ of the $n$ nodes permutations. The control problem we want to solve is the following: given $r(\tau)$ and $m_{\tau}$ for $\tau<t$, choose $m_{t}=\mu_{t}$ such that $\bar{r}_{t}\left(\mu_{t}\right) \simeq \rho$, where $\rho$ is a suitable parameter.

Remark 1. If we want to dynamically control the number of processors, $\rho$ must be chosen different from zero, otherwise the system converges to use only one processor, thus not being able to identify available parallelism. A value of $\rho \in[20 \%, 30 \%]$ is often reasonable, together with the constraint $m_{t} \geq 2$.

## 3 Exploiting Parallelism

In this section we study how much parallelism can be extracted from a given CC graph and how its sparsity can affect the conflict ratio. To this purpose we obtain a worst case class of graphs and use it to analytically derive a lower bound for the exploitable parallelism (i.e., an upper bound for the conflict ratio). We make extensive use of finite differences (i.e., discrete derivatives), which are defined recursively as follows. Let $f: \mathbb{Z} \rightarrow \mathbb{R}$ be a real function defined on the integers, then the $i$-th (forward) finite difference of $f$ is

$$
\begin{equation*}
\Delta_{f}^{i}(k)=\Delta_{f}^{i-1}(k+1)-\Delta_{f}^{i-1}(k), \quad \text { with } \Delta_{f}^{0}(k)=f(k) \tag{2}
\end{equation*}
$$

(In the following we will omit $\Delta^{\prime}$ s superscript when equal to one, i.e., $\Delta \triangleq \Delta^{1}$.)
First, we obtain two basic properties of $\bar{r}$, which are given by the following propositions.
Proposition 1. The conflict ratio function $\bar{r}(m)$ is non-decreasing in $m$.
To prove Prop. 1 we first need a lemma:
Lemma 1. Let $\bar{k}(m) \triangleq \mathbb{E}_{\pi_{m}}\left[k\left(\pi_{m}\right)\right]$. Then $\bar{k}$ is a non-decreasing convex function, i.e. $\Delta_{\bar{k}}(m) \geq 0$ and $\Delta_{\bar{k}}^{2}(m) \geq 0$.

Proof. Let $\tilde{k}\left(\pi_{m}, i\right)$ be the expected number of conflicting nodes running $r=m+i$ nodes concurrently, the first $m$ of which are $\pi_{m}$ and the last $i$ are chosen uniformly at random among the remaining ones. By definition, we have

$$
\begin{equation*}
\mathbb{E}_{\pi_{m}}\left[\tilde{k}\left(\pi_{m}, i\right)\right]=\bar{k}(m+i) \tag{3}
\end{equation*}
$$

In particular,

$$
\begin{equation*}
\tilde{k}\left(\pi_{m}, 1\right)=k\left(\pi_{m}\right)+\operatorname{Pr}[(m+1) \text {-th conflicts }] \tag{4}
\end{equation*}
$$

which brings

$$
\begin{equation*}
\bar{k}(m+1)=\mathbb{E}_{\pi_{m}}\left[\tilde{k}\left(\pi_{m}, 1\right)\right]=\bar{k}(m)+\eta \tag{5}
\end{equation*}
$$

with $\eta=\bar{k}(m+1)-\bar{k}(m)=\Delta_{\bar{k}}(m) \geq 0$, hence proving the monotonicity of $\bar{k}$. Consider now

$$
\begin{equation*}
\tilde{k}\left(\pi_{m}, 2\right)=k\left(\pi_{m}\right)+\operatorname{Pr}[(m+1) \text {-th conflicts }]+\operatorname{Pr}[(m+2) \text {-th conflicts }] . \tag{6}
\end{equation*}
$$

If the $(m+1)$-th node does not add any edge, then we have

$$
\begin{equation*}
\operatorname{Pr}[(m+1) \text {-th conflicts }]=\operatorname{Pr}[(m+2) \text {-th conflicts }], \tag{7}
\end{equation*}
$$

but since it may add some edges the probability of conflicting the second time is in general larger and thus $\Delta_{\bar{k}}^{2}(m) \geq 0$.

Prop. 1 Since $\bar{r}(m)=\bar{k}(m) / m$, its finite difference can be written as

$$
\begin{equation*}
\Delta_{\bar{r}}(m)=\frac{m \Delta_{\bar{k}}(m)-\bar{k}(m)}{m(m+1)} \tag{8}
\end{equation*}
$$

Because of Lemma 1 and being $\bar{k}(1)=0$ we have

$$
\begin{equation*}
\bar{k}(m+1) \leq m \Delta_{\bar{k}}(m) \tag{9}
\end{equation*}
$$

which finally brings

$$
\begin{equation*}
\Delta_{\bar{r}}(m)=\frac{m \Delta_{\bar{k}}(m)-\bar{k}(m)}{m(m+1)} \geq \frac{\bar{k}(m+1)-\bar{k}(m)}{m(m+1)}=\frac{\Delta_{\bar{k}}(m)}{m(m+1)} \geq 0 \tag{10}
\end{equation*}
$$

Proposition 2. Let G be a CC graph, with n nodes and average degree $d$, then the initial derivative of $\bar{r}$ depends only on $n$ and $d$ as

$$
\begin{equation*}
\Delta_{\bar{r}}(1)=\frac{d}{2(n-1)} \tag{11}
\end{equation*}
$$

Proof. Since

$$
\begin{equation*}
\Delta_{\bar{r}}(1)=\frac{\Delta_{\bar{k}}(1)-\bar{k}(1)}{2}=\frac{\bar{k}(2)}{2} \tag{12}
\end{equation*}
$$

we just need to obtain $\bar{k}(2)$. Let $\tilde{k}$ be defined as in the proof on Lemma 11 and $\pi_{1}=v$ a node chosen uniformly at random. Then

$$
\begin{equation*}
\bar{k}(2)=\mathbb{E}_{v}[\tilde{k}(v, 1)]=\mathbb{E}_{v}\left[\frac{d_{v}}{n-1}\right]=\frac{\mathbb{E}_{v}\left[d_{v}\right]}{n-1}=\frac{d}{n-1} \tag{13}
\end{equation*}
$$

A measure of the available parallelism for a given CC graph has been identified in [15] considering, at each temporal step, a maximal independent set of the CC graph. The expected size of a maximal independent set gives a reasonable and computable estimate of the available parallelism. However, this is not enough to predict the actual amount of parallelism that a scheduler can exploit while keeping a low conflict ratio, as shown in the following example.
Example 1. Let $G=K_{n^{2}} \cup D_{n}$ where $K_{n^{2}}$ is the complete graph of size $n^{2}$ and $D_{n}$ a disconnected graph of size $n$ (i.e. $G$ is made up of a clique of size $n^{2}$ and $n$ disconnected nodes). For this graph every maximal independent set is maximum too and has size $n+1$, but if we choose $n+1$ nodes uniformly at random and then compute the conflicts we obtain that, on average, there are only 2 independent nodes.

A more realistic estimate of the performance of a scheduler can be obtained by analyzing the CC graph sparsity. The average degree of the CC graph is linked to the expected size of a maximal independent set of the graph by the following well known theorem (in the variant shown in [2] or [22]):

Theorem 1. (Turán, strong formulation). Let $G=(V, E)$ be a graph, $n=|V|$ and let $d$ be the average degree of $G$. Then the expected size of a maximal independent set, obtained choosing greedily the nodes from a random permutation, is at least $s=n /(d+1)$.
Remark 2. The previous bound is existentially tight: let $K_{d}^{n}$ be the graph made up of $s=$ $n /(d+1)$ cliques of size $d+1$, then the average degree is $d$ and the size of every maximal (and maximum) independent set is exactly $s$. Furthermore, every other graph with the same number of nodes and edges has a bigger average maximal independent set.

The study of the expected size of a maximal independent set in a given graph is also known as the unfriendly seating problem [7, 8] and is particularly relevant in statistical physics, where it is usually studied on mesh-like graphs [11]. The properties of the graph $K_{d}^{n}$ has suggested us the formulation of an extension of the Turán's theorem. We prove that the graphs $K_{d}^{n}$ provide a worst case (for a given degree $d$ ) for the generalization of this problem obtained by focusing on maximal independent set of induced subgraphs. This allows, when given a target conflict ratio $\rho$, the computation of a lower bound for the parallelism a scheduler can exploit.

Theorem 2. Let $G$ be a graph with same nodes number and degree of $K_{d}^{n}$ and let $\mathrm{EM}_{m}(G)$ be the expected size of a maximal independent set of the subgraph induced by a uniformly random choice of $m$ nodes in $G$, then

$$
\begin{equation*}
\mathrm{EM}_{m}(G) \geq \mathrm{EM}_{m}\left(K_{d}^{n}\right) \tag{14}
\end{equation*}
$$

To prove it we first need the following lemma.
Lemma 2. The function $\eta_{j}(x) \triangleq \prod_{i=1}^{j}(n-i-x)$ is convex for $x \in[0, n-j]$.
Proof. We prove by induction on $j$ that, for $x \in[0, n-j]$,

$$
\begin{equation*}
\eta_{j}(x) \geq 0, \quad \eta_{j}^{\prime}(x) \leq 0, \quad \eta_{j}^{\prime \prime}(x) \geq 0 \tag{15}
\end{equation*}
$$

Base case Let $\eta_{0}(x)=1$. The properties above are easily verified.
Induction Since $\eta_{j}(x)=\eta_{j-1}(x)(n-j-x)$, we obtain

$$
\begin{equation*}
\eta_{j}^{\prime}(x)=-\eta_{j-1}(x)+(n-j-x) \eta_{j-1}^{\prime}(x), \tag{16}
\end{equation*}
$$

which is non-positive by inductive hypotheses. Similarly,

$$
\begin{equation*}
\eta_{j}^{\prime \prime}(x)=-2 \eta_{j-1}^{\prime}(x)+(n-j-x) \eta_{j-1}^{\prime \prime}(x) \tag{17}
\end{equation*}
$$

is non-negative.
Thm. 2 Consider a random permutation $\pi$ of the nodes of a generic graph $G$ that has the same number of nodes and edges of $K_{d}^{n}$. We assume the prefix of length $m$ of $\pi$ (i.e. $\pi(1), \ldots, \pi(m)$ ) forms the active nodes and focus on the following independent set $\mathrm{IS}_{m}$ in the subgraph induced: a node $v$ is in $\mathrm{IS}_{m}(G, \pi)$ if and only if it is in the first $m$ positions of $\pi$ and it has no neighbors preceding it. Let $b_{m}(G)$ be the expected size of $\operatorname{IS}_{m}(G, \pi)$ averaged over all possible $\pi$ 's (chosen uniformly):

$$
\begin{equation*}
b_{m}(G) \triangleq \mathbb{E}_{\pi}\left[\# \mathrm{IS}_{m}(G, \pi)\right] . \tag{18}
\end{equation*}
$$

Since for construction $b_{m}(G) \leq \operatorname{EM}_{m}(G)$ whereas $b_{m}\left(K_{d}^{n}\right)=\operatorname{EM}_{m}\left(K_{d}^{n}\right)$, we just need to prove that $b_{m}\left(K_{d}^{n}\right) \leq b_{m}(G)$. Given a generic node $v$ of degree $d_{v}$ and a random permutation $\pi$, its probability to be in $\mathrm{IS}_{m}(G, \pi)$ is

$$
\begin{equation*}
\operatorname{Pr}\left[v \in \operatorname{IS}_{m}(G, \pi)\right]=\frac{1}{n} \sum_{j=1}^{m} \prod_{i=1}^{j-1} \frac{n-i-d_{v}}{n-i} . \tag{19}
\end{equation*}
$$

By the linearity of the expectation we can write $b$ as

$$
\begin{align*}
b_{m}(G) & =\frac{1}{n} \sum_{v=v_{1}}^{v_{n}} \sum_{j=1}^{m} \prod_{i=1}^{j-1} \frac{n-i-d_{v}}{n-i}=\mathbb{E}_{v}\left[\sum_{j=1}^{m} \prod_{i=1}^{j-1} \frac{n-i-d_{v}}{n-i}\right],  \tag{20}\\
b_{m}\left(K_{d}^{n}\right) & =\sum_{j=1}^{m} \prod_{i=1}^{j-1} \frac{n-i-d}{n-i}=\sum_{j=1}^{m} \prod_{i=1}^{j-1} \frac{n-i-\mathbb{E}_{v}\left[d_{v}\right]}{n-i} . \tag{21}
\end{align*}
$$

To prove that $\mathrm{EM}_{m}(G) \geq \mathrm{EM}_{m}\left(K_{d}^{n}\right)$ is thus enough showing that

$$
\begin{equation*}
\forall j \quad \mathbb{E}_{v}\left[\prod_{i=1}^{j}\left(n-i-d_{v}\right)\right] \geq \prod_{i=1}^{j}\left(n-i-\mathbb{E}_{v}\left[d_{v}\right]\right) \tag{22}
\end{equation*}
$$

which can be done applying Jensen's inequality [13], since in Lemma 2 we have proved the convexity of $\eta_{j}(x) \triangleq \prod_{i=1}^{j}(n-i-x)$.
Corollary 1. The worst case for a scheduler among the graphs with the same number of nodes and edges is obtained for the graph $K_{d}^{n}$ (for which we can analytically approximate the performance, as shown in \$3.1.

Proof. Since

$$
\begin{equation*}
\bar{r}(m)=\frac{m-\mathrm{EM}_{m}(G)}{m}=1-\frac{1}{m} \operatorname{EM}_{m}(G) \tag{23}
\end{equation*}
$$

the thesis follows.

### 3.1 Analysis of the Worst-Case Performance

Theorem 3. Let $d$ be the average degree of $G=(V, E)$ with $n=|V|$ (for simplicity we assume $n /(d+1) \in \mathbb{N})$. The conflict ratio is bounded from above as

$$
\begin{equation*}
\bar{r}(m) \leq 1-\frac{n}{m(d+1)}\left(1-\prod_{i=1}^{m} \frac{n-d-i}{n+1-i}\right) \tag{24}
\end{equation*}
$$

Proof. Let $s=n /(d+1)$ be the number of connected components in $K_{d}^{n}$. Because of Thm. 2 and Cor. 1 it suffices to show that

$$
\begin{equation*}
\mathrm{EM}_{m}\left(K_{d}^{n}\right)=s\left(1-\prod_{i=1}^{m} \frac{n-d-i}{n+1-i}\right) \tag{25}
\end{equation*}
$$

The probability for a connected component $k$ of $K_{d}^{n}$ not to be accessed when $m$ nodes are chosen is given by the following hypergeometric

$$
\begin{equation*}
\operatorname{Pr}[k \text { not hit }]=\frac{\binom{n-d-1}{m}\binom{d+1}{0}}{\binom{n}{m}}=\prod_{i=1}^{m} \frac{n-d-i}{n+1-i} \tag{26}
\end{equation*}
$$

Let $X_{k}$ be a random variable that is 1 when component $k$ is hit and 0 otherwise. We have that $\mathbb{E}\left[X_{k}\right]=1-\prod_{i=1}^{m} \frac{n-d-i}{n+1-i}$ and, by the linearity of the expectation, the average number of components accessed is

$$
\begin{equation*}
\mathbb{E}\left[\sum_{k=1}^{s} X_{k}\right]=\sum_{k=1}^{s} \mathbb{E}\left[X_{k}\right]=s\left(1-\prod_{i=1}^{m} \frac{n-d-i}{n+1-i}\right) \tag{27}
\end{equation*}
$$

Corollary 2. When $n$ and $m$ increase the bound is well approximated by

$$
\begin{equation*}
\bar{r}(m) \leq 1-\frac{n}{m(d+1)}\left[1-\left(1-\frac{m}{n}\right)^{d+1}\right] \tag{28}
\end{equation*}
$$

Proof. Stirling approximation for the binomial, followed by low order terms deletion in the resulting formula.

Corollary 3. If we set $m=\alpha s=\frac{\alpha n}{d+1}$ we obtain

$$
\begin{equation*}
\bar{r}(m) \leq 1-\frac{1}{\alpha}\left[1-\left(1-\frac{\alpha}{d+1}\right)^{d+1}\right] \leq 1-\frac{1}{\alpha}\left[1-e^{-\alpha}\right] . \tag{29}
\end{equation*}
$$

## 4 Controlling Processors Allocation

In this section we will design an efficient control heuristic that dynamically chooses the number of processes to be run by a scheduler, in order to obtain high parallelism while keeping the conflict ratio low. In the following we suppose that the properties of $G_{t}$ are varying slowly compared to the convergence of $m_{t}$ toward $\mu_{t}$ under the algorithm we will develop (see s4.1), so we can consider $G_{t}=G$ and $\mu_{t}=\mu$ and thus our goal is making $m_{t}$ converge to $\mu$.

Since the conflict ratio is a non-decreasing function of the number of launched tasks $m$ (Prop. 1) we could find $m \simeq \mu$ by bisection simply noticing that

$$
\begin{equation*}
\bar{r}\left(m^{\prime}\right) \leq \rho \leq \bar{r}\left(m^{\prime \prime}\right) \quad \Rightarrow \quad m^{\prime} \leq \mu \leq m^{\prime \prime} . \tag{30}
\end{equation*}
$$

The control we propose is slightly more complex and is based on recurrence relations, i.e., we compute $m_{t+1}$ as a function $F$ of the target conflict ratio $\rho$ and of the parameters which characterize the system at the previous timestep:

$$
\begin{equation*}
m_{t+1}^{F}=F\left(\rho, r_{t}, m_{t}\right) . \tag{31}
\end{equation*}
$$

The initial value $m_{0}$ for a recurrence can be chosen to be 2 but, if we have an estimate of the CC graph average degree $d$, we can choose a smarter value: in fact applying Cor. 3 we are sure that using, e.g., $m=\frac{n}{2(d+1)}$ processors we will have at most a conflict ratio of $21.3 \%$.

Our control heuristic (Algorithm 1) is a hybridization of two simple recurrences. The first recurrence is quite natural and increases $m$ based on the distance between $r$ and $\rho$ :

$$
\begin{equation*}
\text { Recurrence A: } \quad m_{t+1}^{A}=\left(1-r_{t}+\rho\right) m_{t} . \tag{32}
\end{equation*}
$$

The second recurrence exploits some experimental facts. In Fig. 2 we have plotted the conflict ratio functions for three CC graphs with the same size and average degree (note that initial derivative is the same for all the graphs, in accordance with Prop. 2). We see that conflict ratios which reach a high value $\left(\bar{r}(n)>\frac{1}{2}\right)$ are initially well approximated by a straight line (for $m$ such that $\bar{r}(m) \leq \rho=20 \div 30 \%$ ), whereas functions that deviates from this behavior do not raise too much. This suggests us to assume an initial linearity in controlling $m_{t}$, as done by the following recurrence:

$$
\begin{equation*}
\text { Recurrence B: } \quad m_{t+1}^{B}=\frac{\rho}{r_{t}} m_{t} . \tag{33}
\end{equation*}
$$

The two recurrences can be roughly compared as follows (see Fig. 3): Recurrence A has a slower convergence than Recurrence B, but it is less susceptible to noise (the variance that makes $r_{t}$ realizations different from $\bar{r}_{t}$ ). This is the reason for which we chose to merge them in an hybrid algorithm: initially, when the difference between $r$ and $\rho$ is big, we use Recurrence B to exploit its quick convergence and then Recurrence A is adopted, for a finer tuning of the control.

```
Algorithm 1: Pseudo-code of the proposed hybrid control algorithm
    // Tunable parameters
    \(m_{0}=2 ; \quad m_{\max }=1024 ; \quad m_{\min }=2 ;\)
    \(T=4 ; \quad r_{\min }=3 \% ; \quad \alpha_{0}=25 \% ; \quad \alpha_{1}=6 \% ;\)
    // Variables
    \(m \leftarrow m_{0} ; \quad r \leftarrow 0 ; \quad t \leftarrow 0 ;\)
    // Main loop
    while nodes to elaborate \(\neq 0\) do
        \(t \leftarrow t+1 ;\)
        if \(m>m_{\text {max }}\) then \(m \leftarrow m_{\text {max }}\);
        else if \(m<m_{\text {min }}\) then \(m \leftarrow m_{\text {min }}\);
        Launch the scheduler with \(m\) nodes;
        \(r \leftarrow r+\) new conflict ratio;
        if \((t \bmod T)=T-1\) then
            \(r \leftarrow r / T ;\)
            \(\alpha \leftarrow\left|1-\frac{r}{\rho}\right|\);
            if \(\alpha>\alpha_{0}\) then
                if \(r<r_{\text {min }}\) then \(r \leftarrow r_{\text {min }}\);
                \(m \leftarrow\left\lceil\frac{\rho}{r} m\right\rceil ;\)
            else if \(\alpha>\alpha_{1}\) then
                \(m \leftarrow\lceil(1-r+\rho) m\rceil ;\)
            \(r \leftarrow 0 ;\)
```



Figure 2: A plot of $\bar{r}(m)$ for some graphs with $n=2000$ and $d=16$ : (i) the worst case upper bound of Cor. 2 (ii) a random graph (edges chosen uniformly at random until desired degree is reached; data obtained by computer simulation) (iii) a graph unions of cliques and disconnected nodes.


Figure 3: Comparison between two realizations of the hybrid algorithm and one that only uses Recurrence A, for two different random graphs ( $n=2000$ in both cases). The hybrid version has different parameters for $m$ greater or smaller than 20. $\rho$ was chosen to be $20 \%$. The proposed algorithm proves to be both quick in convergence and stable.

### 4.1 Experimental Evaluation

In the practical implementation of the control algorithm we have made the following optimizations:

- Since $r_{t}$ can have a big variance, especially when $m$ is small, we decided to apply the changes to $m$ every $T$ steps, using the averaged values obtained in these intervals, to smooth the oscillations.
- To further reduce the oscillations we apply a change only if the observed $r_{t}$ is sufficiently different from $\rho$ (e.g. more than $6 \%$ ), thus avoiding small variations in the steady state, which interfere with locality exploitation because of the data moving from one processor to another.
- Another problem that must be considered is that for small values of $m$ the variance is much bigger, so it is better to tune separately this case using different parameters (this optimization is not shown in the pseudo-code).
To validate our controller we have run the following simulation: a random CC graph of fixed average degree $d$ is taken and the controller runs on it, starting with $m_{0}=2$. We are interested in seeing how many temporal steps it takes to converge to $m_{t} \simeq \mu$. As can be seen in [15] the parallelism profile of many practical applications can vary quite abruptly, e.g., Delauney mesh refinement can go from no parallelism to one thousand possible parallel tasks in just 30 temporal steps. Therefore, an algorithm that wants to efficiently control the processors allocations for these problems must adapt very quickly to changes in the available parallelism. Our controller, that uses the very fast Recurrence B in the initial phase, proves to do a fast enough job: as shown in Fig. 3 in about 15 steps the controller converges close to the desired $\mu$ value.


## 5 Conclusions and Future Work

Automatic parallelization of irregular algorithms is a rich and complex subject and will offer many difficult challenges to researchers in the next future. In this paper we have focused on the processor allocation problem for unordered data-amorphous algorithms; it would be extremely valuable to obtain similar results for the more general and difficult case of ordered algorithms (e.g., discrete event simulation), in particular it is very hard to obtain good estimates of the available parallelism for such algorithms, given the complex dependencies arising between the concurrent tasks. Another aspect which needs investigation, especially in
the ordered context, is whether some statical properties of the behavior of irregular algorithms can be modeled, extracted and exploited to build better controllers, able to dynamically adapt to the different execution phases.

As for a real-world implementation, the proposed control heuristic is now being integrated in the Galois system and it will be evaluated on more realistic workloads.

## Acknowledgments

We express our gratitude to Gianfranco Bilardi for the valuable feedback on recurrence-based controllers and to all the Galois project members for the useful discussions on optimistic parallelization modeling.

## References

[1] Agrawal, K., Leiserson, C.E., He, Y., Hsu, W.J.: Adaptive work-stealing with parallelism feedback. ACM Trans. Comput. Syst. 26(3), 7:1-7:32 (Sep 2008), http://doi.acm.org/ 10.1145/1394441.1394443
[2] Alon, N., Spencer, J.: The probabilistic method. Wiley-Interscience (2000)
[3] An, P., Jula, A., Rus, S., Saunders, S., Smith, T.G., Tanase, G., Thomas, N., Amato, N.M., Rauchwerger, L.: Stapl: An adaptive, generic parallel C++ library. In: Dietz, H.G. (ed.) LCPC. Lecture Notes in Computer Science, vol. 2624, pp. 193-208. Springer (2001)
[4] Blackford, L.S., Choi, J., Cleary, A., D'Azevedo, E., Demmel, J., Dhillon, I., Dongarra, J., Hammarling, S., Henry, G., Petitet, A., Stanley, K., Walker, D., Whaley, R.C.: ScaLAPACK Users' Guide. Society for Industrial and Applied Mathematics, Philadelphia, PA (1997)
[5] Braunstein, A., Mézard, M., Zecchina, R.: Survey propagation: An algorithm for satisfiability. Random Struct. Algorithms 27(2), 201-226 (2005)
[6] Eppstein, D.: Spanning trees and spanners. In: Sack, J., Urrutia, J. (eds.) Handbook of Computational Geometry, pp. 425-461. Elsevier (2000)
[7] Freedman, D., Shepp, L.: Problem 62-3, an unfriendly seating arrangement. SIAM Review 4(2), p. 150 (1962), http://www.jstor.org/stable/2028372
[8] Friedman, H.D., Rothman, D., MacKenzie, J.K.: Problem 62-3. SIAM Review 6(2), pp. 180-182 (1964), http://www.jstor.org/stable/2028090
[9] Frigo, M., Johnson, S.G.: The design and implementation of FFTW3. Proceedings of the IEEE 93(2), 216-231 (2005), special issue on "Program Generation, Optimization, and Platform Adaptation"
[10] Frigo, M., Leiserson, C.E., Randall, K.H.: The implementation of the Cilk-5 multithreaded language. In: PLDI. pp. 212-223 (1998)
[11] Georgiou, K., Kranakis, E., Krizanc, D.: Random maximal independent sets and the unfriendly theater seating arrangement problem. Discrete Mathematics 309(16), 5120 - 5129 (2009), http://www.sciencedirect.com/science/article/B6V00-4W55T4X-2/2/ 72d38a668c737e68edf497512e606e12
[12] Guibas, L.J., Knuth, D.E., Sharir, M.: Randomized incremental construction of delaunay and voronoi diagrams. Algorithmica 7(4), 381-413 (1992)
[13] Jensen, J.: Sur les fonctions convexes et les inégalités entre les valeurs moyennes. Acta Mathematica 30(1), 175-193 (1906)
[14] Kalé, L.V., Krishnan, S.: Charm++: A portable concurrent object oriented system based on C++. In: OOPSLA. pp. 91-108 (1993)
[15] Kulkarni, M., Burtscher, M., Cascaval, C., Pingali, K.: Lonestar: A suite of parallel irregular programs. In: ISPASS. pp. 65-76. IEEE (2009)
[16] Kulkarni, M., Burtscher, M., Inkulu, R., Pingali, K., Cascaval, C.: How much parallelism is there in irregular applications? In: Reed, D.A., Sarkar, V. (eds.) PPOPP. pp. 3-14. ACM (2009)
[17] Méndez-Lojo, M., Nguyen, D., Prountzos, D., Sui, X., Hassaan, M.A., Kulkarni, M., Burtscher, M., Pingali, K.: Structure-driven optimizations for amorphous data-parallel programs. In: Govindarajan, R., Padua, D.A., Hall, M.W. (eds.) PPOPP. pp. 3-14. ACM (2010)
[18] Pingali, K., Nguyen, D., Kulkarni, M., Burtscher, M., Hassaan, M.A., Kaleem, R., Lee, T.H., Lenharth, A., Manevich, R., Méndez-Lojo, M., Prountzos, D., Sui, X.: The tao of parallelism in algorithms. In: Proceedings of the 32nd ACM SIGPLAN conference on Programming language design and implementation. pp. 12-25. PLDI '11, ACM, New York, NY, USA (2011), http://doi.acm.org/10.1145/1993498.1993501
[19] Püschel, M., Moura, J., Johnson, J., Padua, D., Veloso, M., Singer, B., Xiong, J., Franchetti, F., Gacic, A., Voronenko, Y., Chen, K., Johnson, R., Rizzolo, N.: Spiral: Code generation for dsp transforms. Proceedings of the IEEE 93(2), 232-275 (Feb 2005)
[20] Reinders, J.: Intel threading building blocks. O'Reilly \& Associates, Inc., Sebastopol, CA, USA (2007)
[21] Tan, P.N., Steinbach, M., Kumar, V.: Introduction to Data Mining. Addison-Wesley (2005)
[22] Tao, T.: Additive combinatorics. Cambridge University Press (2006)
[23] Versaci, F., Pingali, K.: Brief announcement: processor allocation for optimistic parallelization of irregular programs. In: Proceedings of the 23 rd ACM symposium on Parallelism in algorithms and architectures. pp. 261-262. SPAA '11, ACM, New York, NY, USA (2011), http://doi.acm.org/10.1145/1989493.1989533
[24] Versaci, F., Pingali, K.: Processor allocation for optimistic parallelization of irregular programs. In: Murgante, B., Gervasi, O., Misra, S., Nedjah, N., Rocha, A., Taniar, D., Apduhan, B. (eds.) Computational Science and Its Applications - ICCSA 2012, Lecture Notes in Computer Science, vol. 7333, pp. 1-14. Springer Berlin / Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-31125-3_1


[^0]:    *The original publication is available at www.springerlink.com [24]
    ${ }^{\dagger}$ Contact author. Email: versaci@par.tuwien.ac.at. This works was partially supported by PAT-INFN Project AuroraScience, by MIUR-PRIN Project AlgoDEEP, and by the University of Padova Projects STPD08JA32 and CPDA099949
    ${ }^{\ddagger}$ pingali@cs.utexas.edu

[^1]:    ${ }^{1}$ A brief announcement of this work has been presented at SPAA'11 [23]

